



# UltraSPARC Architecture 2005

---

*Draft D0.8.7, 27 Mar 2006*

*Privilege Levels: Hyperprivileged,  
Privileged,  
and Nonprivileged*

*Distribution: Public*

Some portions of this specification are undergoing final review;  
please check monthly to see if an updated revision is available  
for download.

Sun Microsystems, Inc.  
4150 Network Circle  
Santa Clara, CA 95054  
U.S.A. 650-960-1300

Part No: 950-4895-07  
Revision: Draft D0.8.7, 27 Mar 2006



Copyright 2002-2005 Sun Microsystems, Inc., 4150 Network Circle • Santa Clara, CA 950540 USA. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd. For Netscape Communicator™, the following notice applies: Copyright 1995 Netscape Communications Corporation. All rights reserved.

Sun, Sun Microsystems, the Sun logo, Solaris, UltraSPARC, and VIS are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

**RESTRICTED RIGHTS:** Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID

---

Copyright 2002–2005 Sun Microsystems, Inc., 4150 Network Circle • Santa Clara, CA 950540 Etats-Unis. Tous droits réservés.

Des parties de ce document est protégé par un copyright© 1994 SPARC International, Inc.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Sun, Sun Microsystems, le logo Sun, Solaris, UltraSPARC et VIS sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.

---

Comments and "bug reports" regarding this document are welcome; they should be submitted to email address: [UST1-editor@sun.com](mailto:UST1-editor@sun.com)



# Contents

---

<b>1</b>	<b>Preface</b> .....	<b>i</b>
1.1	What's New? .....	i
1.2	Acknowledgements .....	ii
<b>2</b>	<b>Document Overview</b> .....	<b>1</b>
2.1	Navigating <i>UltraSPARC Architecture 2005</i> .....	1
2.2	Fonts and Notational Conventions .....	2
2.2.1	Implementation Dependencies .....	4
2.2.2	Notation for Numbers .....	4
2.2.3	Informational Notes .....	4
2.3	Reporting Errors in this Specification .....	5
<b>3</b>	<b>Definitions</b> .....	<b>7</b>
<b>4</b>	<b>Architecture Overview</b> .....	<b>21</b>
4.1	The UltraSPARC Architecture 2005 .....	22
4.1.1	Features .....	22
4.1.2	Attributes .....	23
4.1.2.1	Design Goals .....	24
4.1.2.2	Register Windows .....	24
4.1.3	System Components .....	24
4.1.3.1	Binary Compatibility .....	24
4.1.3.2	UltraSPARC Architecture 2005 MMU .....	25
4.1.3.3	Privileged Software .....	25
4.1.4	Architectural Definition .....	25
4.1.5	UltraSPARC Architecture 2005 Compliance with SPARC V9 Architecture 25	
4.1.6	Implementation Compliance with UltraSPARC Architecture 2005 25	
4.2	Processor Architecture .....	26
4.2.1	Integer Unit (IU) .....	26

4.2.2	Floating-Point Unit (FPU) .....	27
4.3	Instructions .....	27
4.3.1	Memory Access .....	27
4.3.1.1	Memory Alignment Restrictions .....	28
4.3.1.2	Addressing Conventions .....	28
4.3.1.3	Addressing Range .....	28
4.3.1.4	Load/Store Alternate .....	29
4.3.1.5	Separate I and D Memories .....	29
4.3.1.6	Input/Output (I/O) .....	29
4.3.1.7	Memory Synchronization .....	30
4.3.2	Arithmetic / Logical / Shift Instructions .....	30
4.3.3	Control Transfer .....	30
4.3.4	State Register Access .....	31
4.3.5	Floating-Point Operate .....	32
4.3.6	Conditional Move .....	32
4.3.7	Register Window Management .....	32
4.4	Traps .....	32
4.5	Chip-Level Multithreading (CMT) .....	33
<b>5</b>	<b>Data Formats .....</b>	<b>35</b>
5.1	Integer Data Formats .....	36
5.1.1	Signed Integer Data Types .....	37
5.1.1.1	Signed Integer Byte, Halfword, and Word .....	37
5.1.1.2	Signed Integer Doubleword (64 bits) .....	37
5.1.1.3	Signed Integer Extended-Word (64 bits) .....	38
5.1.2	Unsigned Integer Data Types .....	38
5.1.2.1	Unsigned Integer Byte, Halfword, and Word .....	38
5.1.2.2	Unsigned Integer Doubleword (64 bits) .....	39
5.1.2.3	Unsigned Extended Integer (64 bits) .....	39
5.1.3	Tagged Word (32 bits) .....	39
5.2	Floating-Point Data Formats .....	40
5.2.1	Floating Point, Single Precision (32 bits) .....	40
5.2.2	Floating Point, Double Precision (64 bits) .....	41
5.2.3	Floating Point, Quad Precision (128 bits) .....	42
5.2.4	Floating-Point Data Alignment in Memory and Registers .....	43
5.3	SIMD Data Formats .....	43
5.3.1	Uint8 SIMD Data Format .....	44
5.3.2	Int16 SIMD Data Formats .....	44
5.3.3	Int32 SIMD Data Format .....	44
<b>6</b>	<b>Registers .....</b>	<b>47</b>
6.1	Reserved Register Fields .....	48
6.2	General-Purpose R Registers .....	49
6.2.1	Global R Registers .....	49
6.2.2	Windowed R Registers .....	49

6.2.3	Special R Registers	55
6.3	Floating-Point Registers	55
6.3.1	Floating-Point Register Number Encoding	58
6.3.2	Double and Quad Floating-Point Operands	59
6.4	Floating-Point State Register (FSR)	61
6.4.1	Floating-Point Condition Codes (fcc0, fcc1, fcc2, fcc3)	61
6.4.2	Rounding Direction (rd)	62
6.4.3	Trap Enable Mask (tem)	62
6.4.4	Nonstandard Floating-Point (ns)	63
6.4.5	FPU Version (ver)	63
6.4.6	Floating-Point Trap Type (ftt)	63
6.4.7	FQ Not Empty (qne)	66
6.4.8	Accrued Exceptions (aexc)	66
6.4.9	Current Exception (cexc)	67
6.4.10	Floating-Point Exception Fields	68
6.4.11	FSR Conformance	70
6.5	Ancillary State Registers	70
6.5.1	32-bit Multiply/Divide Register (Y) (ASR 0)	72
6.5.2	Integer Condition Codes Register (CCR) (ASR 2)	72
6.5.2.1	Condition Codes (CCR.xcc and CCR.icc)	73
6.5.3	Address Space Identifier (ASI) Register (ASR 3)	74
6.5.4	Tick (TICK) Register (ASR 4)	74
6.5.5	Program Counters (PC, NPC) (ASR 5)	76
6.5.6	Floating-Point Registers State (FPRS) Register (ASR 6)	76
6.5.7	Performance Control Register (PCR <sup>P</sup> ) (ASR 16)	78
6.5.8	Performance Instrumentation Counter (PIC) Register (ASR 17)	79
6.5.9	General Status Register (GSR) (ASR 19)	80
6.5.10	SOFTINT <sup>P</sup> Register (ASRs 20, 21, 22)	80
6.5.10.1	SOFTINT_SET <sup>P</sup> Pseudo-Register (ASR 20)	82
6.5.10.2	SOFTINT_CLR <sup>P</sup> Pseudo-Register (ASR 21)	82
6.5.11	Tick Compare (TICK_CMPR <sup>P</sup> ) Register (ASR 23)	83
6.5.12	System Tick (STICK) Register (ASR 24)	83
6.5.13	System Tick Compare (STICK_CMPR <sup>P</sup> ) Register (ASR 25)	84
6.6	Register-Window PR State Registers	85
6.6.1	Current Window Pointer (CWP <sup>P</sup> ) Register (PR 9)	86
6.6.2	Savable Windows (CANSAVE <sup>P</sup> ) Register (PR 10)	86
6.6.3	Restorable Windows (CANRESTORE <sup>P</sup> ) Register (PR 11)	87
6.6.4	Clean Windows (CLEANWIN <sup>P</sup> ) Register (PR 12)	87
6.6.5	Other Windows (OTHERWIN <sup>P</sup> ) Register (PR 13)	87
6.6.6	Window State (WSTATE <sup>P</sup> ) Register (PR 14)	88
6.6.7	Register Window Management	88
6.6.7.1	Register Window State Definition	88
6.6.7.2	Register Window Traps	89
6.7	Non-Register-Window PR State Registers	90
6.7.1	Trap Program Counter (TPC <sup>P</sup> ) Register (PR 0)	90
6.7.2	Trap Next PC (TNPC <sup>P</sup> ) Register (PR 1)	91

6.7.3	Trap State (TSTATE <sup>P</sup> ) Register (PR 2) .....	92
6.7.4	Trap Type (TT <sup>P</sup> ) Register (PR 3) .....	93
6.7.5	Trap Base Address (TBA <sup>P</sup> ) Register (PR 5) .....	94
6.7.6	Processor State (PSTATE <sup>P</sup> ) Register (PR 6) .....	94
6.7.7	Trap Level Register (TL <sup>P</sup> ) (PR 7) .....	99
6.7.8	Processor Interrupt Level (PIL <sup>P</sup> ) Register (PR 8) .....	101
6.7.9	Global Level Register (GL <sup>P</sup> ) (PR 16) .....	101
6.8	HPR State Registers .....	103
5.7.1	Hyperprivileged State (HPSTATE <sup>H</sup> ) Register (HPR 0) .....	104
5.7.2	Hyperprivileged Trap State (HTSTATE <sup>H</sup> ) Register (HPR 1) .	105
5.7.3	Hyperprivileged Interrupt Pending (HINTP <sup>H</sup> ) Register (HPR 3) .	106
5.7.4	Hyperprivileged Trap Base Address (HTBA <sup>H</sup> ) Register (HPR 5) .	107
5.7.5	Hyperprivileged Implementation Version (HVER <sup>H</sup> ) Register (HPR 6)	107
5.7.6	Hyperprivileged System Tick Compare (HSTICK_CMPR <sup>H</sup> ) Register (HPR 31)	109
<b>6</b>	<b>Instruction Set Overview .....</b>	<b>111</b>
6.1	Instruction Execution .....	111
6.2	Instruction Formats .....	112
6.3	Instruction Categories .....	113
6.3.1	Memory Access Instructions .....	113
6.3.1.1	Memory Alignment Restrictions .....	114
6.3.1.2	Addressing Conventions .....	115
6.3.1.3	Address Space Identifiers (ASIs) .....	120
6.3.1.4	Separate Instruction Memory .....	121
6.3.2	Memory Synchronization Instructions .....	122
6.3.3	Integer Arithmetic and Logical Instructions .....	122
6.3.3.1	Setting Condition Codes .....	122
6.3.3.2	Shift Instructions .....	122
6.3.3.3	Set High 22 Bits of Low Word .....	123
6.3.3.4	Integer Multiply/Divide .....	123
6.3.3.5	Tagged Add/Subtract .....	123
6.3.4	Control-Transfer Instructions (CTIs) .....	123
6.3.4.1	Conditional Branches .....	125
6.3.4.2	Unconditional Branches .....	126
6.3.4.3	CALL and JMWL Instructions .....	126
6.3.4.4	RETURN Instruction .....	126
6.3.4.5	DONE and RETRY Instructions .....	126
6.3.4.6	Trap Instruction (Tcc) .....	127
6.3.4.7	DCTI Couples .....	127
6.3.5	Conditional Move Instructions .....	127
6.3.6	Register Window Management Instructions .....	129
6.3.6.1	SAVE Instruction .....	129

6.3.6.2	RESTORE Instruction . . . . .	129
6.3.6.3	SAVED Instruction . . . . .	130
6.3.6.4	RESTORED Instruction . . . . .	130
6.3.6.5	Flush Windows Instruction . . . . .	131
6.3.7	Ancillary State Register (ASR) Access . . . . .	131
6.3.8	Privileged Register Access. . . . .	131
6.3.9	Floating-Point Operate (FPop) Instructions . . . . .	131
6.3.10	Implementation-Dependent Instructions . . . . .	132
6.3.11	Reserved Opcodes and Instruction Fields . . . . .	132
<b>7</b>	<b>Instructions . . . . .</b>	<b>135</b>
7.30.1	FMUL8x16 Instruction . . . . .	202
7.30.2	FMUL8x16AU Instruction . . . . .	203
7.30.3	FMUL8x16AL Instruction . . . . .	203
7.30.4	FMUL8SUx16 Instruction . . . . .	204
7.30.5	FMUL8ULx16 Instruction . . . . .	204
7.30.6	FMULD8SUx16 Instruction . . . . .	205
7.30.7	FMULD8ULx16 Instruction. . . . .	206
7.33.1	FPAK16 . . . . .	211
7.33.2	FPAK32 . . . . .	212
7.33.3	FPAKFIX . . . . .	214
7.46.1	IMPDEP1 Opcodes . . . . .	236
7.46.1.1	Opcode Formats . . . . .	237
7.46.2	IMDEP2B Opcodes. . . . .	237
7.61.1	Memory Synchronization . . . . .	273
7.61.2	Synchronization of the Virtual Processor . . . . .	274
7.61.3	TSO Ordering Rules affecting Use of MEMBAR. . . . .	274
7.72.1	Exceptions . . . . .	293
7.72.2	Weak versus Strong Prefetches . . . . .	294
7.72.3	Prefetch Variants. . . . .	295
7.72.3.1	Prefetch for Several Reads (f <sub>cn</sub> = 0, 20(14 <sub>16</sub> )) . . . . .	295
7.72.3.2	Prefetch for One Read (f <sub>cn</sub> = 1, 21(15 <sub>16</sub> )) . . . . .	296
7.72.3.3	Prefetch for Several Writes (and Possibly Reads) (f <sub>cn</sub> = 2, 22(16 <sub>16</sub> )) . . . . .	296
7.72.3.4	Prefetch for One Write (f <sub>cn</sub> = 3, 23(17 <sub>16</sub> )). . . . .	297
7.72.3.5	Prefetch Page (f <sub>cn</sub> = 4) . . . . .	297
7.72.4	Implementation-Dependent Prefetch Variants (f <sub>cn</sub> = 16, 18, 19, and 24–31) . . . . .	297
7.72.5	Additional Notes . . . . .	297
<b>8</b>	<b>IEEE Std 754-1985 Requirements for UltraSPARC Architecture 2005. . . . .</b>	<b>379</b>
8.1	Traps Inhibiting Results. . . . .	379
8.2	NaN Operand and Result Definitions . . . . .	380
8.2.1	Untrapped Result in Different Format from Operands . . . . .	380
8.2.2	Untrapped Result in Same Format as Operands . . . . .	381
8.3	Trapped Underflow Definition (uf <sub>m</sub> = 1). . . . .	382

8.4	Untrapped Underflow Definition (ufm = 0)	383
8.5	Integer Overflow Definition	383
8.6	Floating-Point Nonstandard Mode	384
<b>9</b>	<b>Memory</b>	<b>385</b>
9.1	Memory Location Identification	386
9.2	Memory Accesses and Cacheability	386
9.2.1	Coherence Domains	386
9.2.1.1	Cacheable Accesses	387
9.2.1.2	Noncacheable Accesses	387
9.2.1.3	Noncacheable Accesses with Side-Effect	387
9.3	Memory Addressing and Alternate Address Spaces	389
9.3.1	Memory Addressing Types	389
9.3.2	Memory Address Spaces	390
9.3.3	Address Space Identifiers	391
9.4	SPARC V9 Memory Model	393
9.4.1	SPARC V9 Program Execution Model	393
9.4.2	Virtual Processor/Memory Interface Model	395
9.5	The UltraSPARC Architecture Memory Model — TSO	396
9.5.1	Memory Model Selection	397
9.5.2	Programmer-Visible Properties of the UltraSPARC Architecture TSO Model	397
9.5.3	TSO Ordering Rules	398
9.5.4	Hardware Primitives for Mutual Exclusion	400
9.5.4.1	Compare-and-Swap (CASA, CASXA)	400
9.5.4.2	Swap (SWAP)	401
9.5.4.3	Load Store Unsigned Byte (LDSTUB)	401
9.5.5	Memory Ordering and Synchronization	401
9.5.5.1	Ordering MEMBAR Instructions	401
9.5.5.2	Sequencing MEMBAR Instructions	402
9.5.5.3	Synchronizing Instruction and Data Memory	403
9.6	Nonfaulting Load	405
9.7	Store Coalescing	405
<b>10</b>	<b>Address Space Identifiers (ASIs)</b>	<b>407</b>
10.1	Address Space Identifiers and Address Spaces	407
10.2	ASI Values	407
10.3	ASI Assignments	408
10.3.1	Supported ASIs	409
10.4	Special Memory Access ASIs	422
10.4.1	ASIs 10 <sub>16</sub> , 11 <sub>16</sub> , 16 <sub>16</sub> , 17 <sub>16</sub> and 18 <sub>16</sub> (ASI_*AS_IF_USER_*)	422
10.4.2	ASIs 18 <sub>16</sub> , 19 <sub>16</sub> , 1E <sub>16</sub> , and 1F <sub>16</sub> (ASI_*AS_IF_USER_*_LITTLE)	423
10.4.3	ASI 14 <sub>16</sub> (ASI_REAL)	424
10.4.4	ASI 15 <sub>16</sub> (ASI_REAL_IO)	424

10.4.5	ASI 1C <sub>16</sub> (ASI_REAL_LITTLE) . . . . .	425
10.4.6	ASI 1D <sub>16</sub> (ASI_REAL_IO_LITTLE) . . . . .	425
10.4.7	ASIs 22 <sub>16</sub> , 23 <sub>16</sub> , 27 <sub>16</sub> , 2A <sub>16</sub> , 2B <sub>16</sub> , 2F <sub>16</sub> (Privileged Load Integer Twin Extended Word) 425	
10.4.8	ASIs 26 <sub>16</sub> and 2E <sub>16</sub> (Privileged Load Integer Twin Extended Word, Real Addressing) 426	
10.4.9	ASIs 30 <sub>16</sub> , 31 <sub>16</sub> , 36 <sub>16</sub> , 38 <sub>16</sub> , 39 <sub>16</sub> , 3E <sub>16</sub> (ASI_AS_IF_PRIV_*) .	427
10.4.10	ASIs E2 <sub>16</sub> , E3 <sub>16</sub> , EA <sub>16</sub> , EB <sub>16</sub> (Nonprivileged Load Integer Twin Extended Word) 428	
10.4.11	Block Load and Store ASIs. . . . .	429
10.4.12	Partial Store ASIs . . . . .	430
10.4.13	Short Floating-Point Load and Store ASIs. . . . .	430
10.5	ASI-Accessible Registers . . . . .	430
10.5.1	Privileged Scratchpad Registers (ASI_SCRATCHPAD). . . . .	431
10.5.2	Hyperprivileged Scratchpad Registers (ASI_HYP_SCRATCHPAD) 432	
10.5.3	CMT Registers Accessed Through ASIs. . . . .	432
10.5.4	ASI Changes in the UltraSPARC Architecture . . . . .	432
<b>11</b>	<b>Performance Instrumentation . . . . .</b>	<b>435</b>
<b>12</b>	<b>Traps. . . . .</b>	<b>437</b>
12.1	Virtual Processor Privilege Modes . . . . .	438
12.2	Virtual Processor States, Normal Traps, and RED_state Traps . . . . .	440
12.2.1	RED_state . . . . .	441
12.2.1.1	RED_state Execution Environment . . . . .	442
12.2.1.2	RED_state Entry Traps . . . . .	443
12.2.1.3	RED_state Software Considerations . . . . .	444
12.2.1.4	Usage of Trap Levels . . . . .	444
12.2.2	error_state . . . . .	444
12.3	Trap Categories. . . . .	445
12.3.1	Precise Traps . . . . .	445
12.3.2	Deferred Traps . . . . .	445
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	Reset Traps. . . . .	449
12.3.5	Uses of the Trap Categories. . . . .	450
12.4	Trap Control . . . . .	451
12.4.1	PIL Control . . . . .	452
12.4.2	FSR.tem Control. . . . .	452
12.5	Trap-Table Entry Addresses . . . . .	452
12.5.1	Trap-Table Entry Address to Privileged Mode . . . . .	453
12.5.2	Privileged Trap Table Organization . . . . .	454

12.5.3	Trap-Table Entry Address to Hyperprivileged Mode . . . . .	454
12.5.4	Hyperprivileged Trap Table Organization. . . . .	455
12.5.5	Trap Table Entry Address to RED_state . . . . .	455
12.5.6	RED_state Trap Table Organization. . . . .	456
12.5.7	Trap Type (TT) . . . . .	456
12.5.7.1	Trap Type for Spill/Fill Traps. . . . .	469
12.5.8	Trap Priorities . . . . .	469
12.6	Trap Processing. . . . .	470
12.6.1	Normal Trap Processing. . . . .	472
12.6.2	RED_state Trap Processing. . . . .	474
12.6.2.1	Nonreset Traps with TL = MAXTL - 1. . . . .	474
12.6.2.2	Power-On Reset (POR) Traps . . . . .	476
12.6.2.3	Watchdog Reset (WDR) Traps. . . . .	477
12.6.2.4	Externally Initiated Reset (XIR) Traps . . . . .	478
12.6.2.5	Software-Initiated Reset (SIR) Traps . . . . .	479
12.6.2.6	Nonreset Traps When the Virtual Processor Is in RED_state480	
12.7	Exception and Interrupt Descriptions . . . . .	481
12.7.1	SPARC V9 Traps Not Used in UltraSPARC Architecture 2005	490
12.8	Register Window Traps . . . . .	491
12.8.1	Window Spill and Fill Traps . . . . .	491
12.8.2	<i>clean_window</i> Trap. . . . .	492
12.8.3	Vectoring of Fill/Spill Traps . . . . .	492
12.8.4	CWP on Window Traps . . . . .	493
12.8.5	Window Trap Handlers . . . . .	493
<b>13</b>	<b>Interrupt Handling . . . . .</b>	<b>495</b>
13.1	Interrupt Packets. . . . .	496
13.2	Software Interrupt Register (SOFTINT). . . . .	496
13.2.1	Setting the Software Interrupt Register . . . . .	497
13.2.2	Clearing the Software Interrupt Register. . . . .	497
13.3	Interrupt Queues . . . . .	497
13.3.1	Interrupt Queue Registers . . . . .	498
13.4	Interrupt Traps . . . . .	500
13.5	Strand Interrupt ID Register (STRAND_INTR_ID) . . . . .	501
<b>14</b>	<b>Memory Management . . . . .</b>	<b>503</b>
14.1	Virtual Address Translation . . . . .	503
14.2	TSB Translation Table Entry (TTE) . . . . .	505
14.3	Translation Storage Buffer (TSB). . . . .	508
14.3.1	TSB Indexing Support . . . . .	508
14.3.2	TSB Cacheability . . . . .	509
14.3.3	TSB Organization . . . . .	509
14.4	Faults and Traps . . . . .	509

<b>15</b>	<b>Chip-Level Multithreading (CMT)</b>	<b>511</b>
15.1	Overview of CMT	511
15.1.1	CMT Definition	512
15.1.1.1	Background Terminology	512
15.1.1.2	CMT Definition	514
15.1.2	General CMT Behavior	515
15.2	Accessing CMT Registers	515
15.2.1	Classes of CMT Registers	516
15.2.2	Accessing CMT Registers Through ASIs	517
15.3	CMT Registers	518
15.3.1	Strand ID Register (STRAND_ID)	519
15.3.1.1	Exposing Stranding	520
15.3.2	Strand Interrupt ID Register (STRAND_INTR_ID)	520
15.3.2.1	Assigning an Interrupt ID	521
15.3.2.2	Dispatching and Receiving Interrupts	521
15.3.2.3	Updating the Strand Interrupt ID Register	522
15.4	Disabling and Parking Virtual Processors	522
15.4.1	Strand Available Register (STRAND_AVAILABLE)	523
15.4.2	Enabling and Disabling Virtual Processors	523
15.4.2.1	Strand Enable Status Register (STRAND_ENABLE_STATUS)	524
15.4.2.2	Strand Enable Register (STRAND_ENABLE)	524
15.4.2.3	Dynamically Enabling/Disabling Virtual Processors	526
15.4.3	Parking and Unparking Virtual Processors	526
15.4.3.1	Strand Running Register (STRAND_RUNNING)	527
15.4.3.2	Strand Running Status Register (STRAND_RUNNING_STATUS)	530
15.4.4	Virtual Processor Standby (or Wait) State	531
15.5	Reset and Trap Handling	532
15.5.1	Per-Strand Resets (SIR and WDR Resets)	532
15.5.2	Full-Processor Resets (POR and WRM Resets)	532
15.5.2.1	Boot Sequence	533
15.5.3	Partial Processor Resets (XIR Reset)	533
15.5.3.1	XIR Steering Register (XIR_STEERING)	534
15.6	Error Handling in CMT Processors	535
15.6.1	Virtual-Processor-Specific Error Reporting	535
15.6.2	Reporting Errors on Shared Structures	536
15.6.2.1	Error Steering	536
15.6.2.2	Reporting Non-Virtual-Processor-Specific Errors	539
15.7	Additional CMT Software Interfaces	540
15.7.1	Diagnostic/RAS Registers	540
15.7.2	Configuration Registers	540
15.7.3	Performance Registers	540
15.7.4	Booting Support	540
15.8	Performance Issues for CMT Processors	541



# 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. Now, the UltraSPARC Architecture specification provides the first significant update in over 10 years to Sun's SPARC processor architecture.

---

## 1.1 What's New?

For the first time, UltraSPARC Architecture 2005 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:

- the VIS™ 1 and VIS 2 instruction sets and GSR register
- multiple levels of global registers, controlled by the GL register
- MMU architecture
- the new Hyperprivileged mode
- Chip-level Multithreading (CMT) architecture

Plus, now architectural features are tagged with Software Classes and Implementation Classes<sup>1</sup>. Software Classes provide a new, high-level view of the expected architectural longevity and portability of software that references those

<sup>1</sup>. 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

features. Implementation Classes give an indication of how efficiently each feature is likely to be implemented across current and future UltraSPARC 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.

---

## 1.2 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 *UltraSPARC Architecture 2005* specification are the result of **years** of deliberation, review, and feedback from readers of earlier Sun-internal revisions. I would particularly like to acknowledge the following people for their key contributions:

- The UltraSPARC 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 expansion of exception descriptions in every page of the Instructions chapter, major re-writes of 7 chapters and appendices (*Memory*, *Memory Management*, *Performance Instrumentation*, *Resets*, 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 this 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

I hope you find the *UltraSPARC Architecture 2005* specification more complete, accurate, and readable than its predecessors.

— *David Weaver*

UltraSPARC Architecture coordinator and specification editor

Corrections and other comments regarding this specification can be emailed to:  
UA-editor@sun.com



## Document Overview

---

This chapter discusses:

- **Navigating UltraSPARC Architecture 2005** on page 1.
- **Fonts and Notational Conventions** on page 2.
- **Reporting Errors in this Specification** on page 5.

---

### 2.1 Navigating *UltraSPARC Architecture 2005*

If you are new to the SPARC architecture, read Chapter 4, *Architecture Overview*, study the definitions in Chapter 3, *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 UltraSPARC Architecture 2005, note that UltraSPARC Architecture 2005 conforms to the SPARC V9 Level 1 architecture (and most of Level 2), with numerous extensions — particularly with respect to CMT features, VIS instructions, and support for hyperprivileged-mode operation. For additional details, see the following:

- Chapter 3, *Definitions*
- Chapter 5, *Data Formats*, for a description of the supported data formats
- Chapter 6, *Registers*, for a description of the register set
- Chapter 6, *Instruction Set Overview*, for a description of the new instructions
- Chapter 7, *Instructions*, for descriptions of instruction set extensions
- Chapter 8, *IEEE Std 754-1985 Requirements for UltraSPARC Architecture 2005*, for a description of the trap model
- Chapter 9, *Memory*
- Chapter 10, *Address Space Identifiers (ASIs)*, for a complete list of supported ASIs

- Chapter 11, *Performance Instrumentation*
- Chapter 12, *Traps*, for a description of the trap model
- Chapter 13, *Interrupt Handling*, for information on how interrupts are handled
- Chapter 14, *Memory Management*
- Chapter 15, *Chip-Level Multithreading (CMT)*, for a description of new CMT features
- Chapter 16, *Resets*, for a detailed description of resets, `RED_state`, and `error_state`.
- Appendix A, *Opcode Maps*, to see the overall pictures of how the instruction opcodes are mapped
- Appendix B, *Implementation Dependencies*, for descriptions of resolutions of all implementation dependencies
- Appendix C, *Assembly Language Syntax*, to see extensions to the assembly language syntax; in particular, synthetic instructions are documented in this appendix
- Appendix D, *Formal Specification of the Memory Models*

---

## 2.2 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, “`fccn`”, “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 3, *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 (  $\leftarrow$  ) is the assignment operator. For example, “PC  $\leftarrow$  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 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”.
- The 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.

- The 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.

## 2.2.1 Implementation Dependencies

Implementors of UltraSPARC Architecture 2005 processors are allowed to resolve some aspects of the architecture in machine-dependent ways. Each possible implementation dependency is indicated by the notation “**IMPL. DEP. #*nn***: Some descriptive text.” In this specification, the number  $\overline{nn}$  enumerates the dependencies in . References to implementation dependencies are indicated by the notation “(impl. dep. #*nn*)”.

## 2.2.2 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,  $1001_2$ ,  $FFFF\ 0000_{16}$ ). 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).

## 2.2.3 Informational Notes

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

<b>Note</b>	General notes contain incidental information relevant to the paragraph preceding the note.
<b>Programming Note</b>	Programming notes contain incidental information about how software can use an architectural feature.
<b>Implementation Note</b>	An Implementation Note contains incidental information, describing how an UltraSPARC Architecture 2005 processor might implement an architectural feature.
<b>V9 Compatibility Note</b>	Note containing information about possible differences between UltraSPARC Architecture 2005 and SPARC V9 implementations. Such information is relevant to UltraSPARC Architecture 2005 implementations and might not apply to other SPARC V9 implementations.

<b>Forward Compatibility Note</b>	Note containing information about how the UltraSPARC 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.
---	---

---

## 2.3 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: [UA-editor@sun.com](mailto:UA-editor@sun.com)



**Note: This chapter is undergoing final review; please check back later for a copy of UltraSPARC Architecture 2005 containing the final version of this chapter.**

## Definitions

---

This chapter defines concepts and terminology common to all implementations of UltraSPARC Architecture 2005.

- aliased** Said of each of two virtual addresses that refer to the same underlying memory location.
- address space identifier (ASI)** An 8-bit value that identifies an address space. For each instruction or data access, an ASI is associated with the address. *See also* **implicit ASI**.
- 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.
- available (virtual processor)** A virtual processor that is physically present and functional, that can be enabled and used.
- 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.
- BLD** (Obsolete) abbreviation for Block Load instruction; replaced by LDBLOCKF.
- BST** (Obsolete) abbreviation for Block Store instruction; replaced by STBLOCKF.

<b>bypass ASI</b>	An ASI that refers to memory space and for which the MMU does not perform address translation (that is, memory is accessed using a direct physical address).
<b>byte</b>	Eight consecutive bits of data, aligned on an 8-bit boundary.
<b>clean window</b>	A register window in which all of the registers contain 0, a valid address from the current address space, or valid data from the current address space.
<b>CMT</b>	Chip-level MultiThreading (or, as an adjective, Chip-level MultiThreaded). Refers to a processor containing more than one virtual processor.
<b>coherence</b>	A set of protocols guaranteeing that all memory accesses are globally visible to all caches on a shared-memory bus.
<b>completed (memory operation)</b>	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.
<b>consistency</b>	<i>See coherence.</i>
<b>context</b>	A set of translations that defines a particular address space. <i>See also Memory Management Unit (MMU).</i>
<b>context ID</b>	A numeric value that uniquely identifies a particular context.
<b>copyback</b>	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.
<b>CPI</b>	Cycles per instruction. The number of clock cycles it takes to execute an instruction.
<b>core</b>	In an UltraSPARC Architecture processor, may refer to either a virtual processor or a physical processor core.
<b>cross-call</b>	An interprocessor call in a system containing multiple virtual processors.
<b>CTI</b>	Abbreviation for <b>control-transfer instruction</b> .
<b>current window</b>	The block of 24 R registers that is presently in use. The Current Window Pointer (CWP) register points to the current window.
<b>data access (instruction)</b>	A load, store, load-store, or FLUSH instruction.
<b>DCTI</b>	Delayed control transfer instruction.
<b>demap</b>	To invalidate a mapping in the MMU.

<b>denormalized number</b>	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.
<b>deprecated</b>	The term applied to an architectural feature (such as an instruction or register) for which an UltraSPARC 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 UltraSPARC Architecture software and may not be supported in future versions of the architecture.
<b>disable (core)</b>	The process of changing the state of a virtual processor to <i>Disabled</i> , during which all other processor state (including cache coherency) may be lost and all interrupts to that virtual processor will be discarded. <i>See also</i> <b>park</b> and <b>enable</b> .
<b>disabled (core)</b>	A virtual processor that is out of operation (not executing instructions, not participating in cache coherency, and discarding interrupts). <i>See also</i> <b>parked</b> and <b>enabled</b> .
<b>dispatch</b>	To send a previously fetched instruction to one or more functional units for execution. Typically, the instruction is dispatched from a reservation station or other buffer of instructions waiting to be executed. (Other conventions for this term exist, but the this specification attempts to use <i>dispatch</i> consistently as defined here). <i>See also</i> <b>issued</b> .
<b>doublet</b>	Two bytes (16 bits) of data.
<b>doubleword</b>	An 8-byte datum. <b>Note:</b> The definition of this term is architecture dependent and may differ from that used in other processor architectures.
<b>enable (core)</b>	The process of moving a virtual processor from <i>Disabled</i> to <i>Enabled</i> state and preparing it for operation. <i>See also</i> <b>disable</b> and <b>parked</b> .
<b>enabled (core)</b>	A virtual processor that is in operation (participating in cache coherency, but not executing instructions unless it is also <i>Running</i> ). <i>See also</i> <b>disabled</b> and <b>running</b> .
<b>even parity</b>	The mode of parity checking in which each combination of data bits plus a parity bit contains an even number of '1' bits.
<b>exception</b>	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 <i>FSR.tem</i> ) so that the decision on whether or not to apply special processing can be deferred and made by software at a later time. <i>See also</i> <b>trap</b> .
<b>explicit ASI</b>	An ASI that that is provided by a load, store, or load-store alternate instruction (either from its <i>imm_asi</i> field or from the ASI register).

<b>extended word</b>	An 8-byte datum, nominally containing integer data. <b>Note:</b> The definition of this term is architecture dependent and may differ from that used in other processor architectures.
<b>fccn</b>	One of the floating-point condition code fields fcc0, fcc1, fcc2, or fcc3.
<b>floating-point exception</b>	An exception that occurs during the execution of a floating-point operate (FPop) instruction. The exceptions are <i>unfinished_FPop</i> , <i>unimplemented_FPop</i> , <i>sequence_error</i> , <i>hardware_error</i> , <i>invalid_fp_register</i> , or <i>IEEE_754_exception</i> .
<b>F register</b>	A floating-point register. The SPARC V9 architecture includes single-, double-, and quad-precision F registers.
<b>floating-point operate (FPop) instructions</b>	Instructions that perform floating-point calculations, as defined in <i>Floating-Point Operate (FPop) Instructions</i> on page 131. 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.
<b>floating-point trap type</b>	The specific type of a floating-point exception, encoded in the FSR.ftt field.
<b>floating-point unit</b>	A processing unit that contains the floating-point registers and performs floating-point operations, as defined by this specification.
<b>FPop</b>	See floating-point operate (FPop) instructions.
<b>FPRS</b>	Floating-Point Register State register.
<b>FGU</b>	Floating-point and Graphics Unit (which, in most implementations, is a synonym for <b>FPU</b> ).
<b>FPU</b>	Floating-Point Unit.
<b>FSR</b>	Floating-Point Status register.
<b>GL</b>	Global Level register.
<b>GSR</b>	General Status register.
<b>halfword</b>	A 2-byte datum. <b>Note:</b> The definition of this term is architecture dependent and may differ from that used in other processor architectures.
<b>hyperprivileged (software)</b>	Software executing while the processor is in hyperprivileged state.
<b>hyperprivileged (state)</b>	The highest processor privilege state (defined by HPSTATE.hpriv = 1), in which all processor features are accessible.

<b>hypervisor (software)</b>	A layer of software that executes in hyperprivileged processor state. One purpose of hypervisor software (also referred to as “the hypervisor”) is to provide greater isolation between operating system (“supervisor”) software and the underlying processor implementation.
<b>IEEE 754</b>	IEEE Standard 754-1985, the IEEE Standard for Binary Floating-Point Arithmetic.
<b>IEEE-754 exception</b>	A floating-point exception, as specified by IEEE Std 754-1985. Listed within this specification as IEEE_754_exception.
<b>implementation</b>	Hardware or software that conforms to all of the specifications of an instruction set architecture (ISA).
<b>implementation dependent</b>	An aspect of the UltraSPARC 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.
<b>implicit ASI</b>	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).
<b>initiated</b>	<i>Synonym for issued.</i>
<b>instruction field</b>	A bit field within an instruction word.
<b>instruction group</b>	One or more independent instructions that can be dispatched for simultaneous execution.
<b>instruction set architecture</b>	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 UltraSPARC Architecture 2005 instruction set architecture.
<b>integer unit</b>	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.
<b>interrupt request</b>	A request for service presented to a virtual processor by an external device.
<b>inter-strand</b>	Describes an operation that crosses virtual processor (strand) boundaries.
<b>intra-strand</b>	Describes an operation that occurs entirely within one virtual processor (strand).
<b>invalid (ASI or address)</b>	Undefined, reserved, or illegal.
<b>ISA</b>	Instruction set architecture.

**issued** (1) 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.*

(2) An instruction (or sequence of instructions) is said to be *issued* when released from the virtual processor’s instruction fetch unit. Typically, instructions are issued to a reservation station or other buffer of instructions waiting to be executed. (Other conventions for this term exist, but this specification attempts to use “issued” consistently as defined here.)  
*See also dispatched.*

**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** The address translation hardware in an UltraSPARC Architecture implementation that translates 64-bit virtual address into physical addresses. The MMU is composed of the TLBs, ASRs, and ASI registers used to manage address translation. *See also context, physical address, and virtual address.*

**MMU** 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 UltraSPARC Architecture-compliant products. *Synonym: shall.*

<b>next program counter (NPC)</b>	Conceptually, a register that contains the address of the instruction to be executed next if a trap does not occur.
<b>NFO</b>	Nonfault access only.
<b>nonfaulting load</b>	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 UltraSPARC Architecture processor, hardware treats regular and nonfaulting loads identically; the distinction is made in trap handler software). <i>Contrast with speculative load.</i>
<b>nonprivileged</b>	An adjective that describes <ul style="list-style-type: none"> <li>(1) the state of the virtual processor when <code>PSTATE.priv = 0</code>, that is, nonprivileged mode;</li> <li>(2) virtual processor state information that is accessible to software while the virtual processor is in either privileged mode or nonprivileged mode; for example, nonprivileged registers, nonprivileged ASRs, or, in general, nonprivileged state;</li> <li>(3) an instruction that can be executed when the virtual processor is in either privileged mode or nonprivileged mode.</li> </ul>
<b>nonprivileged mode</b>	The mode in which a virtual processor is operating when executing application software (at the lowest privilege level). Nonprivileged mode is defined by <code>PSTATE.priv = 0</code> and <code>HSTATE.hpriv = 0</code> . <i>See also privileged.</i>
<b>nontranslating ASI</b>	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.
<b>normal trap</b>	A trap processed in <code>execute_state</code> (or equivalently, a non- <code>RED_state</code> trap). <i>Contrast with RED_state trap.</i>
<b>NPC</b>	Next program counter.
<b>npt</b>	Nonprivileged trap.
<b>nucleus software</b>	Privileged software running at a trap level greater than 0 ( <code>TL &gt; 0</code> ).
<b>NUMA</b>	Nonuniform memory access.
<b><i>N_REG_WINDOWS</i></b>	The number of register windows present in a particular implementation.
<b>octlet</b>	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.
<b>odd parity</b>	The mode of parity checking in which each combination of data bits plus a parity bit together contain an odd number of ‘1’ bits.
<b>opcode</b>	A bit pattern that identifies a particular instruction.

- optional** A feature not required for UltraSPARC Architecture 2005 compliance.
- PA** Physical address.
- park** The process of suspending a virtual processor from operation. There may be a delay until the virtual processor is parked, but no heavyweight operation (such as a reset) is required to complete the parking process. *See also* **disable** and **unpark**.
- parked** A virtual processor suspended from operation. When parked, a virtual processor does not issue instructions for execution but (if enabled) still maintains cache coherency. *See also* **disabled**, **enabled**, and **running**.
- PC** Program counter.
- PCR** Performance Control register.
- physical address** An address that maps real physical memory or I/O device space. See also **virtual address**.
- physical core** The term *physical processor core*, or just *physical core*, is similar to the term *pipeline* but represents a broader collection of hardware. A physical core includes an execution pipeline plus associated structures, such as caches, that are required for performing the execution of instructions from one or more software threads. A physical core contains one or more virtual processors (strands). The physical core provides the necessary resources for the threads on each strand to make forward progress at a reasonable rate. A multistranded physical core can execute multiple software threads either by time multiplexing or partitioning resources (or any combination thereof). *See also* **pipeline**, **processor**, **strand**, **thread**, and **virtual processor**.
- PIC** Performance Instrumentation Counter.
- PIL** Processor Interrupt Level register.
- pipeline** Refers to an execution pipeline. It is a loose term for the basic collection of hardware needed to execute instructions. A pipeline may be used by one or more strands to execute instructions from one or more threads. *Synonym for* **microcore**. *See also* **physical core**, **processor**, **strand**, **thread**, and **virtual processor**.
- PIPT** Physically indexed, physically tagged (cache).
- POR** Power-on reset.

<b>prefetchable</b>	(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. <i>See side effect.</i>
<b>privileged</b>	An adjective that describes: (1) the state of the processor when PSTATE.priv = 1 and HPSTATE.hpriv = 0, that is, <i>privileged mode</i> ; (2) processor state that is only accessible to software while the processor is in <i>privileged mode</i> ; for example, privileged registers, privileged ASRs, or, in general, privileged state; (3) an instruction that can be executed only when the processor is in <i>privileged mode</i> .
<b>privileged mode</b>	The mode in which a processor is operating when PSTATE.priv = 1 and HPSTATE.hpriv = 0. <i>See also nonprivileged.</i>
<b>processor</b>	The unit on which a shared interface is provided to control the configuration and execution of a collection of strands. A processor contains one or more physical cores, each of which contains one or more strands. On a more physical side, a <i>processor</i> is a physical module that plugs into a system. A processor is expected to appear logically as a single agent on the system interconnect fabric. <i>Synonym for processor module. See also pipeline, physical core, strand, thread, and virtual processor.</i>
<b>processor core</b>	<i>See virtual processor.</i>
<b>processor module</b>	<i>Synonym for processor.</i>
<b>program counter (PC)</b>	A register that contains the address of the instruction currently being executed.
<b>quadword</b>	A 16-byte datum. <b>Note:</b> The definition of this term is architecture dependent and may be different from that used in other processor architectures.
<b>R register</b>	An integer register. Also called a general-purpose register or working register.
<b>RA</b>	Real address.
<b>RAS</b>	(1) Return Address Stack (2) Reliability, Availability, and Serviceability
<b>RAW</b>	Read After Write (hazard)
<b>rd</b>	Rounding direction.
<b>RDPR</b>	Read Privileged Register instruction.

<b>RED_state</b>	Reset, Error, and Debug state. The virtual processor state when HPSTATE.red = 1. A restricted execution environment used to process resets and traps that occur when $TL = MAXTL - 1$ .
<b>RED_state trap</b>	A trap processed in RED_state. <i>Contrast with normal trap.</i>
<b>reserved</b>	<p>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.</p> <p><i>A reserved instruction field must read as 0, unless the implementation supports extended instructions within the field. The behavior of an UltraSPARC Architecture 2005 virtual processor when it encounters a nonzero value in a reserved instruction field is as defined in <i>Reserved Opcodes and Instruction Fields</i> on page 132.</i></p> <p><i>A reserved bit combination within an instruction field is defined in Chapter 7, <i>Instructions</i>. In all cases, an UltraSPARC Architecture 2005 processor must decode and trap on such reserved bit combinations.</i></p> <p><i>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 <i>Reserved Register Fields</i> on page 48).</i></p> <p>Throughout this specification, figures and tables illustrating registers and instruction encodings indicate reserved fields and combinations with an em dash (—).</p>
<b>reset trap</b>	A vectored transfer of control to privileged software through a fixed-address reset trap table. Reset traps cause entry into RED_state.
<b>restricted</b>	Describes an address space identifier (ASI) that may be accessed only while the virtual processor is operating in a privileged mode or hyperprivileged mode.
<b>retired</b>	<p>An instruction is said to be “retired” when one of (instruction) the following two events has occurred:</p> <ol style="list-style-type: none"> <li>(1) A precise trap has been taken, with TPC containing the instruction's address (the instruction has not changed architectural state in this case).</li> <li>(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: <ul style="list-style-type: none"> <li>■ The PC has advanced beyond the instruction.</li> <li>■ 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.</li> <li>■ Stores are visible to all loads in the same instruction stream, including stores to noncacheable locations.</li> </ul> </li> </ol>
<b>RMO</b>	Relaxed memory order.

<b>rs1, rs2, rd</b>	The integer or floating-point register operands of an instruction. <i>rs1</i> and <i>rs2</i> are source registers; <i>rd</i> is the destination register.
<b>RTO</b>	Read to Own (a type of transaction, used to request ownership of a cache line).
<b>RTS</b>	Read to Share (a type of transaction, used to request read-only access to a cache line).
<b>running</b>	A state of a virtual processor in which it is in operation (maintaining cache coherency and issuing instructions for execution) and not <i>Parked</i> .
<b>service processor</b>	A device external to the processor that can examine and alter internal processor state. A service processor may be used to control/coordinate a multiprocessor system and aid in error recovery.
<b>SFAR</b>	Synchronous Fault Address register.
<b>SFSR</b>	Synchronous Fault Status register.
<b>shall</b>	<i>Synonym for must.</i>
<b>should</b>	A keyword indicating flexibility of choice with a strongly preferred implementation. <i>Synonym for it is recommended.</i>
<b>side effect</b>	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. <i>See also prefetchable.</i>
<b>SIMD</b>	Single Instruction/Multiple Data; a class of instructions that perform identical operations on multiple data contained (or “packed”) in each source operand.
<b>SIR</b>	Software-initiated reset.
<b>snooping</b>	The process of maintaining coherency between caches in a shared-memory bus architecture. All cache controllers monitor (snoop) the bus to determine whether they have a copy of the shared cache block.
<b>speculative load</b>	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. <i>Contrast with nonfaulting load.</i>
<b>store</b>	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 <i>Store</i> 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*.

<b>strand</b>	Identifies the hardware state used to hold a software thread in order to execute it. <i>Strand</i> is specifically the software-visible architected state (program counter (PC), next program counter (NPC), general-purpose registers, floating-point registers, condition codes, status registers, ASRs, etc.) of a thread and any microarchitecture state required by hardware for its execution. <i>See also</i> <b>pipeline</b> , <b>physical core</b> , <b>processor</b> , <b>thread</b> , and <b>virtual processor</b> .
<b>subnormal number</b>	<i>Synonym for</i> <b>denormalized number</b> .
<b>superscalar</b>	An implementation that allows several instructions to be issued, executed, and committed in one clock cycle.
<b>supervisor software</b>	Software that executes when the virtual processor is in privileged mode.
<b>suspend</b>	<i>See</i> <b>park</b> .
<b>synchronization</b>	An operation that causes the processor to wait until the effects of all previous instructions are completely visible before any subsequent instructions are executed.
<b>system</b>	A set of virtual processors that share a physical address space.
<b>taken</b>	A control-transfer instruction (CTI) is <i>taken</i> when the CTI alters the control flow by writing a value into NPC other than the default value NPC = 4.  A trap is <i>taken</i> 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.
<b>TBA</b>	Trap base address.
<b>TEE</b>	Thread Execution Engine. <i>Synonym for</i> <b>virtual processor</b> and <b>strand</b> .
<b>thread</b>	A software entity that can be run on hardware. A thread is scheduled, may or may not be actively running on hardware at any given time, and may migrate around the hardware of a system. <i>See also</i> <b>pipeline</b> , <b>physical core</b> , <b>processor</b> , <b>strand</b> , and <b>virtual processor</b> .
<b>TLB</b>	<i>See</i> <b>Translation Lookaside Buffer (TLB)</b> .
<b>TLB hit</b>	The desired translation is present in the TLB.
<b>TLB miss</b>	The desired translation is not present in the TLB.
<b>TPC</b>	Trap-saved program counter.
<b>Translation Lookaside Buffer (TLB)</b>	A cache within an MMU that contains recent partial translations. TLBs speed up closely following translations by often eliminating the need to reread Translation Table Entries from memory.

<b>trap</b>	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 supervisor software through a table, the address of which is specified by the privileged Trap Base Address (TBA) register. <i>See also</i> <b>exception</b> .
<b>TSB</b>	Translation storage buffer. A table of the address translations that is maintained by software in system memory and that serves as a cache of the address translations.
<b>TSO</b>	Total store order.
<b>TTE</b>	Translation Table Entry. Describes the virtual-to-physical translation and page attributes for a specific page in the page table. In some cases, the term is explicitly used for the entries in the TSB.
<b>UA-2005</b>	UltraSPARC Architecture 2005
<b>unassigned</b>	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.
<b>undefined</b>	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, may or may not cause a trap, can vary among implementations, and can vary with 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 privileged mode, or put the virtual processor into an unrecoverable state.
<b>unimplemented</b>	An architectural feature that is not directly executed in hardware because it is optional or is emulated in software.
<b>unpark</b>	The process of bringing a virtual processor out of suspension. There may be a delay until the virtual processor is unparked, but no heavyweight operation (such as a reset) is required to complete the unparking process. <i>See also</i> <b>disable</b> and <b>park</b> .
<b>unparked</b>	<i>Synonym for</i> <b>running</b> .
<b>unpredictable</b>	<i>Synonym for</i> <b>undefined</b> .
<b>uniprocessor system</b>	A system containing a single virtual processor.
<b>unrestricted</b>	Describes an address space identifier (ASI) that can be used in all privileged modes; that is, regardless of the value of PSTATE.priv and HPSTATE.hpriv.
<b>user application program</b>	<i>Synonym for</i> <b>application program</b> .

**VA** Virtual address.

**virtual address** An address produced by a virtual processor that maps all systemwide, program-visible memory. Virtual addresses usually are translated by a combination of hardware and software to physical addresses, which can be used to access physical memory.

**virtual core,  
virtual processor,  
virtual processor core** *Synonyms: virtual processor.*

**virtual processor** The term *virtual processor*, or *virtual processor core*, is used to identify each strand in a processor. A CMT processor contains one or more physical cores, each of which contains one or more virtual processors (strands). Each virtual processor (strand) has its own interrupt ID. At any given time, an operating system can have a different thread scheduled on each virtual processor. *See also pipeline, physical core, processor, strand, and thread.*

**VIS** VIS™ Instruction Set.

**Strand** *Abbreviation for Virtual Processor.*

**WDR** Watchdog reset.

**word** A 4-byte datum. **Note:** The definition of this term is architecture dependent and may differ from that used in other processor architectures.

**WRPR** Write Privileged Register instruction.

**XIR** Externally initiated reset.

## Architecture Overview

---

The UltraSPARC Architecture supports 32- and 64-bit integer and 32-, 64-, and 128-bit floating-point as its principal data types. The 32- 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 *UltraSPARC Architecture 2005* 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 UltraSPARC Architecture 2005 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 UltraSPARC Architecture 2005 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).

---

## 4.1 The UltraSPARC Architecture 2005

This section briefly describes features, attributes, and components of the UltraSPARC Architecture 2005 and, further, describes correct implementation of the architecture specification and SPARC V9-compliance levels.

### 4.1.1 Features

The UltraSPARC Architecture 2005, 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), 16 quad-precision (128-bit) registers, or a mixture thereof.
- **Fast trap handlers** — Traps are vectored through a table.
- **Multiprocessor synchronization instructions** — One instruction performs an atomic read-then-set-memory operation; another performs an atomic exchange-register-with-memory operation; another compares the contents of a register with a value in memory and exchanges memory with the contents of another register if the comparison was equal (compare and swap); two others synchronize the order of shared memory operations as observed by virtual processors.
- **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, UltraSPARC Architecture 2005 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, supports the ability to run multiple simultaneous guest operating systems, and provides more robust handling of error conditions.
- **Multiple levels of global registers** — Instead of the two 8-register sets of global registers specified in the SPARC V9 architecture, UltraSPARC Architecture 2005 provides multiple sets; typically, one set is used at each trap level.
- **Extended instruction set** — UltraSPARC Architecture 2005 provides many instruction set extensions, including the VIS instruction set for "vector" (SIMD) data operations.
- **More detailed, specific instruction descriptions** — UltraSPARC Architecture 2005 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** — Although some details of the UltraSPARC MMU architecture are necessarily implementation-specific, UltraSPARC Architecture 2005 provides a blueprint for the UltraSPARC MMU, including software view (TTEs and TSBs) and MMU hardware control registers.
- **Chip-Level Multithreading (CMT)** — UltraSPARC Architecture 2005 provides a control architecture for highly-threaded processor implementations.

## 4.1.2 Attributes

UltraSPARC Architecture 2005 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, UltraSPARC Architecture 2005 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.

### 4.1.2.1 Design Goals

The UltraSPARC Architecture 2005 architecture is designed to be a target for optimizing compilers and high-performance hardware implementations. This specification documents the UltraSPARC Architecture 2005 and provides a design spec against which an implementation can be verified, using appropriate verification software.

### 4.1.2.2 Register Windows

The UltraSPARC Architecture 2005 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 supervisor software, not user programs, manages the register windows. The supervisor can save a minimum number of registers (approximately 24) during a context switch, thereby optimizing context-switch latency.

## 4.1.3 System Components

The UltraSPARC Architecture 2005 allows for a spectrum of subarchitectures, such as cache system, I/O, and memory management unit (MMU).

### 4.1.3.1 Binary Compatibility

The most important SPARC V9 architectural mandate is binary compatibility of nonprivileged programs across implementations. Binaries executed in nonprivileged mode should behave identically on all SPARC V9 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 SPARC V9 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, the SPARC V9 architecture is binary upward-compatible from SPARC V8 for applications running in nonprivileged mode that conform to the SPARC V8 ABI.

### 4.1.3.2 UltraSPARC Architecture 2005 MMU

Although the SPARC V9 architecture allows its implementations freedom in their MMU designs, UltraSPARC Architecture 2005 defines a common MMU architecture (see Chapter 14, *Memory Management*) with some specifics left to implementations (see processor implementation documents).

### 4.1.3.3 Privileged Software

UltraSPARC Architecture 2005 does not assume that all implementations must execute identical privileged software (operating systems) or hyperprivileged software (hypervisors). Thus, certain traits that are visible to privileged software may be tailored to the requirements of the system.

## 4.1.4 Architectural Definition

The UltraSPARC Architecture 2005 is defined by the chapters and normative 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 normative appendixes.

UltraSPARC Architecture 2005 defines a set of implementations that conform to the SPARC V9 architecture, Level 1.

## 4.1.5 UltraSPARC Architecture 2005 Compliance with SPARC V9 Architecture

UltraSPARC Architecture 2005 fully complies with SPARC V9 Level 1 (nonprivileged). It partially complies with SPARC V9 Level 2 (privileged).

## 4.1.6 Implementation Compliance with UltraSPARC Architecture 2005

Compliant implementations must not add to or deviate from this standard except in aspects described as implementation dependent. Appendix B, *Implementation Dependencies* lists all UltraSPARC Architecture 2005, SPARC V8, and SPARC V9 implementation dependencies. Documents for specific UltraSPARC Architecture 2005 processor implementations describe the manner in which implementation dependencies have been resolved in those implementations.

**IMPL. DEP. #1-V8:** Whether an instruction complies with UltraSPARC Architecture 2005 by being implemented directly by hardware, simulated by software, or emulated by firmware is implementation dependent.

---

## 4.2 Processor Architecture

An UltraSPARC 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 UltraSPARC Architecture virtual processor can run in *nonprivileged* mode, *privileged* mode, or *hyperprivileged* mode. In hyperprivileged mode, the processor can execute any instruction, including privileged instructions. 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 to hyperprivileged software.

### 4.2.1 Integer Unit (IU)

The 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 UltraSPARC Architecture implementation may contain from 72 to 640 general-purpose 64-bit R registers. This corresponds to a grouping of the registers into  $MAXPGL\ MAXGL + 1$  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).

## 4.2.2 Floating-Point Unit (FPU)

The 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. Double-precision values occupy an even-odd pair of single-precision registers, and quad-precision values occupy a quad-aligned group of four single-precision registers.

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.

---

## 4.3 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

These classes are discussed in the following subsections.

### 4.3.1 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 doubleword (64-bit) accesses. Some versions of integer load instructions perform sign extension on 8-, 16-, and 32-bit values as they are loaded into a 64-bit destination register. Floating-point load and store instructions support word, doubleword, and quadword<sup>1</sup> memory accesses.

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

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

### 4.3.1.1 Memory Alignment Restrictions

A memory access on an UltraSPARC 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 114.

### 4.3.1.2 Addressing Conventions

The SPARC V9 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 SPARC V9 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. See *Processor State (PSTATE<sup>P</sup>) Register (PR 6)* on page 94 for information about changing the implicit byte order to little-endian.

Addressing conventions are illustrated in FIGURE 6-2 on page 117 and FIGURE 6-3 on page 119.

### 4.3.1.3 Addressing Range

**IMPL. DEP. #405-S10:** An UltraSPARC Architecture implementation may support a full 64-bit virtual address space or a more limited range of virtual addresses. In an implementation that does 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  $2^{64} - 1$ .

<sup>1</sup> No UltraSPARC Architecture processor currently implements the LDQF instruction in hardware; it generates an exception and is emulated in supervisor software.

### 4.3.1.4 Load/Store Alternate

Versions 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  $00_{16}$ – $7F_{16}$  is restricted to privileged code, and access to alternate spaces  $80_{16}$ – $FF_{16}$  is unrestricted. Some of the ASIs are available for implementation-dependent uses. Supervisor software can use the implementation-dependent ASIs to access special protected registers, such as MMU, cache control, and virtual processor state registers, and other processor- or system-dependent values. See *Address Space Identifiers (ASIs)* on page 120 for more information.

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

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

### 4.3.1.5 Separate I and D 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 translation mechanism and cache and data references use another, although the same main memory is shared.

In such split-memory systems, the coherency mechanism may be split, so a write into data memory is not immediately reflected in instruction memory. For this reason, programs that modify their own code (self-modifying code) and that wish to be portable across all 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 UltraSPARC Architecture virtual processor may or may not have coherent instruction and data caches. Even if it *does*, 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.

### 4.3.1.6 Input/Output (I/O)

The UltraSPARC 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.

### 4.3.1.7 Memory Synchronization

Two instructions are used for synchronization of memory operations: FLUSH and MEMBAR. Their operation is explained in *Flush Instruction Memory* on page 186 and *Memory Barrier* on page 271, respectively.

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

## 4.3.2 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 another arithmetic or logical instruction to create a 32-bit constant in an R register.

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

The integer multiply instruction performs a  $64 \times 64 \rightarrow 64$ -bit operation. The integer division instructions perform  $64 \div 64 \rightarrow 64$ -bit operations. Division by zero causes a trap. Some versions of the 32-bit multiply and divide instructions set the condition codes.

The tagged arithmetic instructions assume that the least-significant two bits of each operand are a data-type tag. These instructions set the integer condition code (icc) and extended integer condition code (xcc) overflow bits on 32-bit (icc) or 64-bit (xcc) arithmetic overflow. In addition, if any of the operands' tag bits are nonzero, icc is set. The xcc overflow bit is not affected by the tag bits.

## 4.3.3 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. A bit in a delayed control-transfer instruction (the *annul bit*) can cause the delay instruction to be annulled (that is, to have no effect) if the branch is not taken (or in the “branch always” case if the branch is taken).

**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 (JMPL) 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  $\pm 8$  Mbytes; the “branch on condition codes with prediction” instruction provides a displacement of  $\pm 1$  Mbyte; the “branch on register contents” instruction provides a displacement of  $\pm 128$  Kbytes; and the CALL instruction’s 30-bit word displacement allows a control transfer to any address within  $\pm 2$  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.

## 4.3.4 State Register Access

The read and write state register instructions read and write the contents of state registers visible to nonprivileged software (Y, CCR, ASI, PC, TICK, and FPRS). The read and write privileged register instructions 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, WSTATE, and VER).

**IMPL. DEP. #8-V8-Cs20:** Ancillary state registers (ASRs) in the range 0–27 that are not defined in UltraSPARC Architecture 2005 are reserved for future architectural use. ASRs in the range 28–31 are available to be used for implementation-dependent purposes.

**IMPL. DEP. #9-V8-Cs20:** Whether each of the implementation-dependent read/write ancillary state register instructions (for ASRs 28–31) is privileged is implementation dependent.

## 4.3.5 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 or two source operands. The groups of instructions that are considered FPops are listed in *Floating-Point Operate (FPop) Instructions* on page 131.

## 4.3.6 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 upon the contents of an integer register. These instructions increase performance by reducing the number of branches.

## 4.3.7 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.

---

# 4.4 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, or the Hyperprivileged Trap Base Register, HTBA). 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 causes hyperprivileged state to be saved in the HTSTATE trap stack. 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.

---

## 4.5 Chip-Level Multithreading (CMT)

An UltraSPARC Architecture implementation may include multiple virtual processor cores on the same processor module to provide a dense, high-throughput system. This may be achieved by having a combination of multiple physical processor cores and/or multiple strands (threads) per physical processor core. Chapter 15, *Chip-Level Multithreading (CMT)*, specifies a common interface between hardware and software for such products, referred to here as chip-level multiprocessors (CMT processors). Chapter 10 addresses issues common to CMT processors, regardless of the microarchitecture of the individual physical processor cores.

The CMT Programming Model describes a set of privileged registers that are used for identification and configuration of CMT processors. Equally important, the CMT Programming Model describes certain behavior that must be common across all UltraSPARC Architecture CMT processors. The set of registers and the common behavior are covered in Chapter 15, *Chip-Level Multithreading (CMT)*.



## Data Formats

---

The UltraSPARC 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 36.
- **Floating-Point Data Formats** on page 40.
- **SIMD Data Formats** on page 43.

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

- **Signed Integer Doubleword (64 bits)** on page 37.
- **Unsigned Integer Doubleword (64 bits)** on page 39.
- **Floating Point, Double Precision (64 bits)** on page 41.
- **Floating Point, Quad Precision (128 bits)** on page 42.

# 5.1 Integer Data Formats

TABLE 5-1 describes the width and ranges of the signed, unsigned, and tagged integer data formats.

**TABLE 5-1** 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 5-2 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 5-2** Integer Doubleword/Extended-word Alignment

Subformat Name	Subformat Field	Memory Address		Register Number	
		Required Alignment	Address (big-endian) <sup>1</sup>	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.

## 5.1.1 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

### 5.1.1.1 Signed Integer Byte, Halfword, and Word

FIGURE 5-1 illustrates the signed integer byte, halfword, and word data formats.




FIGURE 5-1 Signed Integer Byte, Halfword, and Word Data Formats

### 5.1.1.2 Signed Integer Doubleword (64 bits)

FIGURE 5-2 illustrates both components (SD-0 and SD-1) of the signed integer double data format.




FIGURE 5-2 Signed Integer Double Data Format

### 5.1.1.3 Signed Integer Extended-Word (64 bits)

FIGURE 5-3 illustrates the signed integer extended-word (SX) data format.



FIGURE 5-3 Signed Integer Extended-Word Data Format

## 5.1.2 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

### 5.1.2.1 Unsigned Integer Byte, Halfword, and Word

FIGURE 5-4 illustrates the unsigned integer byte data format.





FIGURE 5-4 Unsigned Integer Byte, Halfword, and Word Data Formats

### 5.1.2.2 Unsigned Integer Doubleword (64 bits)

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



**FIGURE 5-5** Unsigned Integer Double Data Format

### 5.1.2.3 Unsigned Extended Integer (64 bits)


FIGURE 5-6 illustrates the unsigned extended integer (UX) data format.



**FIGURE 5-6** Unsigned Extended Integer Data Format

## 5.1.3 Tagged Word (32 bits)

FIGURE 5-7 illustrates the tagged word data format.



**FIGURE 5-7** Tagged Word Data Format

## 5.2 Floating-Point Data Formats

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

### 5.2.1 Floating Point, Single Precision (32 bits)

FIGURE 5-8 illustrates the floating-point single-precision data format, and TABLE 5-3 describes the formats.




FIGURE 5-8 Floating-Point Single-Precision Data Format

TABLE 5-3 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); $f = .000--00$

## 5.2.2 Floating Point, Double Precision (64 bits)

FIGURE 5-9 illustrates both components (FD-0 and FD-1) of the floating-point double-precision data format, and TABLE 5-4 describes the formats.




FIGURE 5-9 Floating-Point Double-Precision Data Format

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

$s$	= sign (1 bit)
$e$	= biased exponent (11 bits)
$f$	= fraction (52 bits)
$u$	= undefined
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); $f = .000--00$

## 5.2.3 Floating Point, Quad Precision (128 bits)

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




FIGURE 5-10 Floating-Point Quad-Precision Data Format

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

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

## 5.2.4 Floating-Point Data Alignment in Memory and Registers

TABLE 5-6 describes the address and memory alignment for floating-point data.

**TABLE 5-6** 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	$f + 1^\diamond$
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	$f + 1^\diamond$
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	$f + 3^\diamond$

\* 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  $\leq 31$ ).

## 5.3 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.,  $\alpha$ , 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.

### 5.3.1 Uint8 SIMD Data Format

The Uint8 SIMD data format consists of four unsigned 8-bit integers contained in a 32-bit word (see FIGURE 5-11).




FIGURE 5-11 Uint8 SIMD Data Format

### 5.3.2 Int16 SIMD Data Formats

The Int16 SIMD data format consists of four signed 16-bit integers contained in a 64-bit word (see FIGURE 5-12).




FIGURE 5-12 Int16 SIMD Data Format

### 5.3.3 Int32 SIMD Data Format

The Int32 SIMD data format consists of two signed 32-bit integers contained in a 64-bit word (see FIGURE 5-13).




FIGURE 5-13 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 49.
- **Floating-Point Registers** on page 55.
- **Floating-Point State Register (FSR)** on page 61.
- **Ancillary State Registers** on page 70. The following registers are included in this category:
  - **32-bit Multiply/Divide Register (Y) (ASR 0)** on page 72.
  - **Integer Condition Codes Register (CCR) (ASR 2)** on page 72.
  - **Address Space Identifier (ASI) Register (ASR 3)** on page 74.
  - **Tick (TICK) Register (ASR 4)** on page 74.
  - **Program Counters (PC, NPC) (ASR 5)** on page 76.
  - **Floating-Point Registers State (FPRS) Register (ASR 6)** on page 76.
  - **Performance Control Register (PCR<sup>P</sup>) (ASR 16)** on page 78.
  - **Performance Instrumentation Counter (PIC) Register (ASR 17)** on page 79.
  - **General Status Register (GSR) (ASR 19)** on page 80.
  - **SOFTINT<sup>P</sup> Register (ASRs 20, 21, 22)** on page 80.
  - **SOFTINT\_SET<sup>P</sup> Pseudo-Register (ASR 20)** on page 82.
  - **SOFTINT\_CLR<sup>P</sup> Pseudo-Register (ASR 21)** on page 82.
  - **Tick Compare (TICK\_CMPR<sup>P</sup>) Register (ASR 23)** on page 83.
  - **System Tick (STICK) Register (ASR 24)** on page 83.
  - **System Tick Compare (STICK\_CMPR<sup>P</sup>) Register (ASR 25)** on page 84.
- **Register-Window PR State Registers** on page 85. The following registers are included in this subcategory:
  - **Current Window Pointer (CWP<sup>P</sup>) Register (PR 9)** on page 86.
  - **Savable Windows (CANSAVE<sup>P</sup>) Register (PR 10)** on page 86.
  - **Restorable Windows (CANRESTORE<sup>P</sup>) Register (PR 11)** on page 87.
  - **Clean Windows (CLEANWIN<sup>P</sup>) Register (PR 12)** on page 87.
  - **Other Windows (OTHERWIN<sup>P</sup>) Register (PR 13)** on page 87.
  - **Window State (WSTATE<sup>P</sup>) Register (PR 14)** on page 88.
- **Non-Register-Window PR State Registers** on page 90. The following registers are included in this subcategory:
  - **Trap Program Counter (TPC<sup>P</sup>) Register (PR 0)** on page 90.
  - **Trap Next PC (TNPC<sup>P</sup>) Register (PR 1)** on page 91.

- Trap State (**TSTATE<sup>P</sup>**) Register (PR 2) on page 92.
- Trap Type (**TT<sup>P</sup>**) Register (PR 3) on page 93.
- Trap Base Address (**TBA<sup>P</sup>**) Register (PR 5) on page 94.
- Processor State (**PSTATE<sup>P</sup>**) Register (PR 6) on page 94.
- Trap Level Register (**TL<sup>P</sup>**) (PR 7) on page 99.
- Processor Interrupt Level (**PIL<sup>P</sup>**) Register (PR 8) on page 101.
- Global Level Register (**GL<sup>P</sup>**) (PR 16) on page 101.
- HPR State Registers on page 103. The following registers are included in this category.
  - Hyperprivileged State (**HPSTATE<sup>H</sup>**) Register (HPR 0) on page 104.
  - Hyperprivileged Trap State (**HTSTATE<sup>H</sup>**) Register (HPR 1) on page 105.
  - Hyperprivileged Interrupt Pending (**HINTP<sup>H</sup>**) Register (HPR 3) on page 106.
  - Hyperprivileged Trap Base Address (**HTBA<sup>H</sup>**) Register (HPR 5) on page 107.
  - Hyperprivileged Implementation Version (**HVER<sup>H</sup>**) Register (HPR 6) on page 107.
  - Hyperprivileged System Tick Compare (**HSTICK\_CMPR<sup>H</sup>**) Register (HPR 31) on page 109.

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

---

## 6.1 Reserved Register Fields

For convenience, some registers in this chapter are illustrated as fewer than 64 bits wide. Any bits not shown (or explicitly marked as reserved) are reserved for future extensions to the architecture.

Such a reserved field within a register reads as zero in current implementations and, when written by software, should only be written with the value of that field previously read from that register or with the value zero.

<b>Programming Note</b>	Software intended to run on future versions of the UltraSPARC Architecture should not assume that reserved register fields will read as 0 or any other particular value.
-------------------------	--

---

## 6.2 General-Purpose R Registers

An UltraSPARC Architecture virtual processor contains an array of general-purpose 64-bit R registers. The array is partitioned into  $MAXGL + 1$  sets of eight *global* registers, plus  $N\_REG\_WINDOWS$  groups of 16 registers each. The value of  $N\_REG\_WINDOWS$  in an UltraSPARC 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 6-1.


### 6.2.1 Global R Registers (A1)

Registers R[0]–R[7] refer to a set of eight registers called the *global* registers (labeled g0 through g7). At any time, one of  $MAXGL + 1$  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 (GL<sup>P</sup>) (PR 16)* on page 101.

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

### 6.2.2 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 6-1 and FIGURE 6-2.



**FIGURE 6-1** General-Purpose Registers (as Visible at Any Given Time)

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 6-1** Window Addressing

Windowed Register Address	R Register Address
<i>in</i> [0] – <i>in</i> [7]	R[24] – R[31]

**TABLE 6-1** Window Addressing

Windowed Register Address	R Register Address
<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 Note**

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 UltraSPARC Architecture, *MAXGL*, ranges from 2 to 15. The number of implemented global register sets is *MAXGL* + 1. The total number of R registers in a given UltraSPARC Architecture implementation is:

$$(N\_REG\_WINDOWS \times 16) + ((MAXGL + 1) \times 8)$$

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

The current window in the windowed portion of R registers is indicated by the current window pointer (CWP) register. The CWP is decremented by the RESTORE instruction and incremented by the SAVE instruction.




FIGURE 6-2 Three Overlapping Windows and Eight Global Registers

**Overlapping Windows.** Each window shares its *ins* with one adjacent window and its *outs* with another. The *outs* of the CWP - 1 (**modulo**  $N\_REG\_WINDOWS$ ) window are addressable as the *ins* of the current window, and the *outs* in the current window are the *ins* of the CWP + 1 (**modulo**  $N\_REG\_WINDOWS$ ) window. The *locals* are unique to each window.

Register address  $o$ , where  $8 \leq o \leq 15$ , refers to exactly the same *out* register before the register window is advanced by a SAVE instruction (CWP is incremented by 1 (**modulo**  $N\_REG\_WINDOWS$ )) as does register address  $o+16$  after the register window is advanced. Likewise, register address  $i$ , where  $24 \leq i \leq 31$ , refers to exactly the same

*in* register before the register window is restored by a RESTORE instruction (CWP is decremented by 1 (**modulo**  $N\_REG\_WINDOWS$ )) as does register address  $i-16$  after the window is restored. See FIGURE 6-2 on page 52 and FIGURE 6-3 on page 54.

To application software, the virtual processor appears to provide an infinitely-deep stack of register windows.

<b>Programming Note</b>	Since the procedure call instructions (CALL and JMPL) do not change the CWP, a procedure can be called without changing the window. See the section “Leaf-Procedure Optimization” in <i>Software Considerations</i> , contained in the separate volume <i>UltraSPARC Architecture Application Notes</i>
-------------------------	---

Since CWP arithmetic is performed modulo  $N\_REG\_WINDOWS$ , the highest-numbered implemented window overlaps with window 0. The *outs* of window  $N\_REG\_WINDOWS - 1$  are the *ins* of window 0. Implemented windows are numbered contiguously from 0 through  $N\_REG\_WINDOWS - 1$ .

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 SAVE<sup>1</sup>. 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 (CLEANWIN<sup>P</sup>) Register (PR 12)* on page 87, *Savable Windows (CANSAVE<sup>P</sup>) Register (PR 10)* on page 86, and *Restorable Windows (CANRESTORE<sup>P</sup>) Register (PR 11)* on page 87.

<b>Implementation Note</b>	An UltraSPARC Architecture virtual processor supports the 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.
----------------------------	--

<sup>1</sup>. 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.

Register Window Management Instructions on page 129 describes how the windowed integer registers are managed.



$$\text{CANSAVE} + \text{CANRESTORE} + \text{OTHERWIN} = N\_REG\_WINDOWS - 2$$

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 6-3** Windowed R Registers for  $N\_REG\_WINDOWS = 8$

In FIGURE 6-3,  $N\_REG\_WINDOWS = 8$ . The eight *global* registers are not illustrated.  $CWP = 0$ ,  $CANSAVE = 4$ ,  $OTHERWIN = 1$ , 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.

## 6.2.3 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.

---

## 6.3 Floating-Point Registers A2

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 6-2. 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 6-2** 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	Bits	Register	Assembly Language
F <sub>S</sub> [0]	%f0	63:32	F <sub>D</sub> [0]	%d0	127:64	F <sub>Q</sub> [0]	%q0
F <sub>S</sub> [1]	%f1	31:0					
F <sub>S</sub> [2]	%f2	63:32	F <sub>D</sub> [2]	%d2	63:0		
F <sub>S</sub> [3]	%f3	31:0					
F <sub>S</sub> [4]	%f4	63:32	F <sub>D</sub> [4]	%d4	127:64	F <sub>Q</sub> [4]	%q4
F <sub>S</sub> [5]	%f5	31:0					
F <sub>S</sub> [6]	%f6	63:32	F <sub>D</sub> [6]	%d6	63:0		
F <sub>S</sub> [7]	%f7	31:0					
F <sub>S</sub> [8]]	%f8	63:32	F <sub>D</sub> [8]	%d8	127:64	F <sub>Q</sub> [8]	%q8
F <sub>S</sub> [9]	%f9	31:0					
F <sub>S</sub> [10]	%f10	63:32	F <sub>D</sub> [10]	%d10	63:0		
F <sub>S</sub> [11]	%f11	31:0					
F <sub>S</sub> [12]	%f12	63:32	F <sub>D</sub> [12]	%d12	127:64	F <sub>Q</sub> [12]	%q12
F <sub>S</sub> [13]	%f13	31:0					
F <sub>S</sub> [14]	%f14	63:32	F <sub>D</sub> [14]	%d14	63:0		
F <sub>S</sub> [15]	%f15	31:0					
F <sub>S</sub> [16]	%f16	63:32	F <sub>D</sub> [16]	%d16	127:64	F <sub>Q</sub> [16]	%q16
F <sub>S</sub> [17]	%f17	31:0					
F <sub>S</sub> [18]	%f18	63:32	F <sub>D</sub> [18]	%d18	63:0		
F <sub>S</sub> [19]	%f19	31:0					
F <sub>S</sub> [20]	%f20	63:32	F <sub>D</sub> [20]	%d20	127:64	F <sub>Q</sub> [20]	%q20
F <sub>S</sub> [21]	%f21	31:0					
F <sub>S</sub> [22]	%f22	63:32	F <sub>D</sub> [22]	%d22	63:0		
F <sub>S</sub> [23]	%f23	31:0					

**TABLE 6-2** 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	Register	Assembly Language
F <sub>S</sub> [24]	%f24	63:32	F <sub>D</sub> [24]	%d24	127:64	F <sub>Q</sub> [24]	%q24
F <sub>S</sub> [25]	%f25	31:0					
F <sub>S</sub> [26]	%f26	63:32	F <sub>D</sub> [26]	%d26	63:0		
F <sub>S</sub> [27]	%f27	31:0					
F <sub>S</sub> [28]	%f28	63:32	F <sub>D</sub> [28]	%d28	127:64	F <sub>Q</sub> [28]	%q28
F <sub>S</sub> [29]	%f29	31:0					
F <sub>S</sub> [30]	%f30	63:32	F <sub>D</sub> [30]	%d30	63:0		
F <sub>S</sub> [31]	%f31	31:0					
		63:32	F <sub>D</sub> [32]	%d32	127:64	F <sub>Q</sub> [32]	%q32
		31:0					
		63:32	F <sub>D</sub> [34]	%d34	63:0		
		31:0					
		63:32	F <sub>D</sub> [36]	%d36	127:64	F <sub>Q</sub> [36]	%q36
		31:0					
		63:32	F <sub>D</sub> [38]	%d38	63:0		
		31:0					
		63:32	F <sub>D</sub> [40]	%d40	127:64	F <sub>Q</sub> [40]	%q40
		31:0					
		63:32	F <sub>D</sub> [42]	%d42	63:0		
		31:0					
		63:32	F <sub>D</sub> [44]	%d44	127:64	F <sub>Q</sub> [44]	%q44
		31:0					
		63:32	F <sub>D</sub> [46]	%d46	63:0		
		31:0					
		63:32	F <sub>D</sub> [48]	%d48	127:64	F <sub>Q</sub> [48]	%q48
		31:0					
		63:32	F <sub>D</sub> [50]	%d50	63:0		
		31:0					

**TABLE 6-2** 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 <sub>D</sub> [52] %d52	127:64	F <sub>Q</sub> [52] %q52
		31:0		63:0	
		63:32	F <sub>D</sub> [54] %d54	63:0	
		31:0			
		63:32	F <sub>D</sub> [56] %d56	127:64	F <sub>Q</sub> [56] %q56
		31:0		63:0	
		63:32	F <sub>D</sub> [58] %d58	63:0	
		31:0			
		63:32	F <sub>D</sub> [60] %d60	127:64	F <sub>Q</sub> [60] %q60
		31:0		63:0	
		63:32	F <sub>D</sub> [62] %d62	63:0	
		31:0			

## 6.3.1 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 labeled 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 6-3.

**TABLE 6-3** 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}

<b>SPARC V8 Compatibility Note</b>	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 UltraSPARC Architecture virtual processor, using the encoding in TABLE 5-3.
--	--

## 6.3.2 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.

<b>Programming Note</b>	<p>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:</p> <ol style="list-style-type: none"> <li>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 <i>must</i> be properly aligned.</li> <li>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.</li> <li>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: <ul style="list-style-type: none"> <li>■ 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</li> <li>■ 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 (<i>LDDF_mem_not_aligned</i>) on some implementations, possibly slowing down program execution significantly.</li> </ul> </li> </ol>
-------------------------	--

<b>Programming Note</b>	<p>If an UltraSPARC Architecture 2005 implementation does not implement a particular quad floating-point arithmetic operation in hardware and an invalid quad register operand is specified, per FSR.ftt priorities in TABLE 6-7, the <i>fp_exception_other</i> exception occurs with FSR.ftt = 3 (unimplemented_FPop) instead of with FSR.ftt = 6 (invalid_fp_register).</p>
-------------------------	---

<b>Implementation Note</b>	<p>UltraSPARC Architecture 2005 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 <i>fp_exception_other</i> exception with FSR.ftt = 3 (unimplemented_FPop).</p>
----------------------------	--

## 6.4 Floating-Point State Register (FSR)

The Floating-Point State register (FSR) fields, illustrated in FIGURE 6-4, contain FPU mode and status information. The lower 32 bits of the FSR are read and written by the STFSR and LDFSR instructions; all 64 bits of the FSR are read and written by the STXFSR and LDXFSR instructions, respectively. FSR.ver, FSR.ftt, and the reserved (“—”) fields of FSR are not modified by LDFSR or LDXFSR.




FIGURE 6-4 FSR Fields

Bits 63–38, 29–28, 21–20, and 12 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 61 through 70 describe the remaining fields in the FSR.

### 6.4.1 Floating-Point Condition Codes (fcc0, fcc1, fcc2, fcc3) (A1)

The four sets of floating-point condition code fields are labeled fcc0, fcc1, fcc2, and fcc3 (fcc*n* 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 fcc*n* fields in the FSR, as selected by the compare instruction. The fcc*n* fields are read and written by STXFSR and LDXFSR instructions, respectively. 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 6-5,  $f_{rs1}$  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 6-4** Floating-Point Condition Codes (`fccn`) Fields of FSR

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

**TABLE 6-5** Floating-Point Condition Codes (`fccn`) Fields of FSR

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

## 6.4.2 Rounding Direction (`rd`) (A1)

Bits 31 and 30 select the rounding direction for floating-point results according to IEEE Std 754-1985. TABLE 6-6 shows the encodings.

**TABLE 6-6** Rounding Direction (`rd`) Field of FSR

<code>rd</code>	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 80 for further details.

## 6.4.3 Trap Enable Mask (`tem`) (A1)

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 (`cxex`). See FIGURE 6-5 on page 69. 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.

## 6.4.4 Nonstandard Floating-Point (ns)

On an UltraSPARC Architecture 2005 processor, *FSR.ns* is a reserved bit; it always reads as 0 and writes to it are ignored. (impl. dep. #18-V8)

## 6.4.5 FPU Version (ver) (A1)

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

For each SPARC V9 IU implementation (as identified by its *VER.impl* field), there may be one or more FPU implementations, or none. This field 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 is read-only; it cannot be modified by the *LDFSR* and *LDXFSR* instructions.

## 6.4.6 Floating-Point Trap Type (ftt) (A1)

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*, *STXFSR*, 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 *STXFSR* instruction. The *LDFSR* and *LDXFSR* instructions do not affect *FSR.ftt*.

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

**Programming Note** Neither LDFSr nor LDXFSr 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.

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 6-7. 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 6-7** 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
unimplemented_FPop	10	3	<code>fp_exception_other</code>
invalid_fp_register	20	6	<code>fp_exception_other</code>
unfinished_FPop	30	2	<code>fp_exception_other</code>
IEEE_754_exception	40	1	<code>fp_exception_ieee_754</code>
<i>Reserved</i>	—	4, 5, 7	—
(none detected)	—	0	—

IEEE\_754\_exception, unimplemented\_FPop, and unfinished\_FPop 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* or *unimplemented\_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* or *unimplemented\_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.

**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 (unimplemented\_FPop).** The *unimplemented\_FPop* floating-point trap type indicates that the virtual processor decoded an *FPop* that it does not implement in hardware. In this case, the *cexc* field is unchanged.

For example, all quad-precision *FPop* variations in an UltraSPARC Architecture 2005 virtual processor cause an *fp\_exception\_other* exception, setting *FSR.ftt = unimplemented\_FPop*.

<b>Forward Compatibility Note</b>	The next revision of the UltraSPARC Architecture is expected to eliminate “ <i>unimplemented_FPop</i> ”, to simplify handling of unimplemented instructions. At that point, all conditions which currently cause <i>fp_exception_other</i> with <i>FSR.ftt = 3</i> will cause an <i>illegal_instruction</i> exception, instead. <i>FSR.ftt = 3</i> and the trap type associated with <i>fp_exception_other</i> will become reserved for other possible future uses.
---	---

#### **ftt = 4 (Reserved).**

<b>SPARC V9 Compatibility Note</b>	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 UltraSPARC 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).**

<b>SPARC V9 Compatibility Note</b>	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 UltraSPARC 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.

<b>Implementation Note</b>	Per FSR.ftt priorities in TABLE 6-7, if an UltraSPARC Architecture 2005 processor does not implement a particular quad FPop in hardware, that FPop generates an <i>fp_exception_other</i> exception with FSR.ftt = 3 (unimplemented_FPop) instead of <i>fp_exception_other</i> with FSR.ftt = 6 (invalid_fp_register), regardless of the specified F registers.
--------------------------------	---

### 6.4.7 FQ Not Empty (qne) (Y2)

Since UltraSPARC Architecture virtual processors do not implement a floating-point queue, FSR.qne always reads as zero and writes to FSR.qne are ignored.

### 6.4.8 Accrued Exceptions (aexc) (A1)

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 6-6 on page 69.

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` is written with the appropriate value when an `LDFSR` or `LDXFSR` instruction is executed.

## 6.4.9 Current Exception (`cexc`) **(A1)**

`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 6-5 on page 69.

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

The `cexc` bits are set as described in *Floating-Point Exception Fields* on page 68, 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 6-8 (where “✓” indicates “exception was detected” and “x” indicates “don’t care”):

TABLE 6-8 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
-	✓ <sup>1</sup>	✓ <sup>1</sup>	x	0	0	no	0	1	1
✓ <sup>2</sup>	-	✓ <sup>2</sup>	0	x	0	no	1	0	1
-	-	✓	x	x	1	yes	0	0	1
-	✓ <sup>1</sup>	✓ <sup>1</sup>	x	0	1	yes	0	0	1
-	✓	-	x	1	x	yes	0	1	0
-	✓	✓	x	1	x	yes	0	1	0
✓ <sup>2</sup>	-	✓ <sup>2</sup>	1	x	x	yes	1	0	0
✓ <sup>2</sup>	-	✓ <sup>2</sup>	0	x	1	yes	0	0	1

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

<sup>2</sup> 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.

## 6.4.10 Floating-Point Exception Fields (A1)

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




FIGURE 6-6 Accrued Exception Bits (aexc) Fields of FSR




FIGURE 6-7 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 UltraSPARC Architecture strand detects tininess before rounding (impl. dep. #55-V8-Cs10). See *Trapped Underflow Definition (ufm = 1)* on page 382 and *Untrapped Underflow Definition (ufm = 0)* on page 383 for additional details.

**Division by zero (dzc, dza).**  $X \div 0.0$ , 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.

## 6.4.11 FSR Conformance

An UltraSPARC 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 `fp_exception_other` (with `FSR.ftt = unfinished_FPop` or `unimplemented_FPop`) and `IEEE_754_exception` floating-point trap types properly. Thus, a user application program always sees an FSR that is fully compliant with IEEE Std 754-1985.

## 6.5 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. UltraSPARC Architecture virtual processors implement the ASRs summarized in TABLE 6-9 and defined in the following subsections.

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

**TABLE 6-9** ASR Register Summary

ASR number	ASR name	Register	Read by Instruction(s)	Written by Instruction(s)
0	Y <sup>D</sup>	Y register (deprecated)	RDY <sup>D</sup>	WRY <sup>D</sup>
1	—	<i>Reserved</i>	—	—
2	CCR	Condition Codes register	RDCCR	WRCCR
3	ASI	ASI register	RDASI	WRASI

**TABLE 6-9** ASR Register Summary (Continued)

ASR number	ASR name	Register	Read by Instruction(s)	Written by Instruction(s)
4	TICK <sup>Pnpt</sup>	TICK register	RDTICK <sup>Pnpt</sup> , RDPR <sup>P</sup> (TICK)	WRPR <sup>P</sup> (TICK)
5	PC	Program Counter (PC)	RDPC	(all instructions)
6	FPRS	Floating-Point Registers Status register	RDFPRS	WRFPS
7–14	—	<i>Reserved</i>	—	—
15	—	<i>Reserved</i>	—	—
16–31	—	non-SPARC V9 ASRs	—	—
16	PCR <sup>P</sup>	Performance Control registers (PCR)	RDPCT <sup>P</sup>	WRPCT <sup>P</sup>
17	PIC <sup>P</sup>	Performance Instrumentation Counters (PIC)	RDPICT <sup>Ppic</sup>	WRPICT <sup>Ppic</sup>
18	—	Implementation dependent (impl. dep. #8-V8-Cs20, 9-V8-Cs20)	—	—
19	GSR	General Status register (GSR)	RDGSR, FALIGNDATA, many VIS and floating-point instructions	WRGSR, BMASK, SIAM
20	SOFTINT_CLR <sup>P</sup>	(pseudo-register, for "Write 1s Clear" to SOFTINT register, ASR 22)	—	WRSOFTINT_CLR <sup>P</sup>
21	SOFTINT_SET <sup>P</sup>	(pseudo-register, for "Write 1s Set" to SOFTINT register, ASR 22)	—	WRSOFTINT_SET <sup>P</sup>
22	SOFTINT <sup>P</sup>	per-virtual processor Soft Interrupt register	RDSOFTINT <sup>P</sup>	WRSOFTINT <sup>P</sup>
23	TICK_CMPR <sup>P</sup>	Tick Compare register	RDTICK_CMPR <sup>P</sup>	WRTICK_CMPR <sup>P</sup>
24	STICK <sup>Pnpt</sup>	System Tick register	RDSTICK <sup>Pnpt</sup>	WRSTICK <sup>H</sup>
25	STICK_CMPR <sup>P</sup>	System Tick Compare register	RDSTICK_CMPR <sup>P</sup>	WRSTICK_CMPR <sup>P</sup>
26–31	—	Implementation dependent (impl. dep. #8-V8-Cs20, 9-V8-Cs20)	—	—

## 6.5.1 32-bit Multiply/Divide Register (Y) (ASR 0) E3

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, MULSc, SDIV, SDIVcc, UDIV, UDIVcc, RDY, and WRY) be avoided. For suitable substitute instructions, see the following pages: for the multiply instructions, see page 367; for the multiply step instruction, see page 281; for division instructions, see page 364; for the read instruction, see page 300; and for the write instruction, see page 370.

The low-order 32 bits of the Y register, illustrated in FIGURE 6-8, contain the more significant word of the 64-bit product of an integer multiplication, as a result of either a 32-bit integer multiply (SMUL, SMULcc, UMUL, UMULcc) instruction or an integer multiply step (MULSc) 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 6-8 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.

## 6.5.2 Integer Condition Codes Register (CCR) (ASR 2) A1

The Condition Codes Register (CCR), shown in FIGURE 6-9, contains the integer condition codes. The CCR register may be explicitly read and written by the RDCCR and WRCCR instructions, respectively.





FIGURE 6-9 Condition Codes Register

### 6.5.2.1 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  $0000\ 0000\ \text{FFFF}\ \text{FFFF}_{16}$ , 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 6-10.



**FIGURE 6-10** 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`) two’s complement notation for the last instruction that modified the integer condition codes; 1 = overflow, 0 = no 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 6-10).

**TABLE 6-10** Setting of Carry (Borrow) bits for Subtraction That Sets CCs

Unsigned Comparison of Operand Values	Setting of Carry bits in CCR
$R[rs1]\{31:0\} \geq R[rs2]\{31:0\}$	<code>CCR.icc.c</code> $\leftarrow$ 0
$R[rs1]\{31:0\} < R[rs2]\{31:0\}$	<code>CCR.icc.c</code> $\leftarrow$ 1
$R[rs1]\{63:0\} \geq R[rs2]\{63:0\}$	<code>CCR.xcc.c</code> $\leftarrow$ 0
$R[rs1]\{63:0\} < R[rs2]\{63:0\}$	<code>CCR.xcc.c</code> $\leftarrow$ 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 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 (x $cc$ ).** 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.

### 6.5.3 Address Space Identifier (ASI) Register (ASR 3) A1

The Address Space Identifier register (FIGURE 6-11) 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  $00_{16}$  to  $7F_{16}$  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  $00_{16}$  to  $7F_{16}$  or an instruction executing in privileged mode attempts an access using an ASI the range  $30_{16}$  to  $7F_{16}$ , a *privileged\_action* exception is generated. See Chapter 10, *Address Space Identifiers (ASIs)* for details.




FIGURE 6-11 Address Space Identifier Register

### 6.5.4 Tick (TICK) Register (ASR 4) A1

FIGURE 6-12 illustrates the TICK register.

TICK<sup>npt</sup>




FIGURE 6-12 TICK Register

The counter field of the TICK register is a 63-bit counter that counts strand clock cycles. Bit 63 of the TICK register is the nonprivileged trap (npt) bit, which controls access to the TICK register by nonprivileged software.

Privileged and hyperprivileged software can always read the TICK register with either the RDPR or RTICK instruction.

Privileged software cannot write to the TICK register; an attempt to do so (with the WRPR instruction) results in an *illegal\_instruction* exception. Hyperprivileged software can always write to the TICK register with the WRPR instruction (there is no distinct WRTICK instruction).

Nonprivileged software can read the TICK register by using the RTICK instruction, but only when nonprivileged access to TICK is enabled (TICK.npt = 0) by hyperprivileged software. If nonprivileged access is disabled (TICK.npt = 1), an attempt by nonprivileged software to read the TICK register causes a *privileged\_action* exception. Nonprivileged software cannot write the TICK register. An attempt by nonprivileged software to read the TICK register using the privileged RDPR instruction causes a *privileged\_opcode* exception.

TICK.npt is set to 1 by a power-on reset trap. The value of TICK.counter is reset to 0 after a power-on reset trap.

After the TICK register is written, reading the TICK register returns a value incremented (by 1 or more) from the last value written, rather than from some previous value of counter. The number of counts between a write and a subsequent read does not accurately reflect the number of strand cycles between the write and the read. Software may rely only on read-to-read counts of the TICK register for accurate timing, not on write-to-read counts.

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.

**Programming Note** TICK.npt may be used by a secure operating system to control access by user software to high-accuracy timing information. The operation of the timer might be emulated by the trap handler, which could read TICK.counter and “fuzz” the value to lower accuracy.

## 6.5.5 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.

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

The Floating-Point Registers State (FPRS) register, shown in FIGURE 6-13, contains control information for the floating-point register file; this information is readable and writable by nonprivileged software.




FIGURE 6-13 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.

<b>Programming Note</b>	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.

**IMPL. DEP. #403-S10(a):** An UltraSPARC Architecture 2005 virtual processor may set *FPRS.du* pessimistically; that is, it may be set whenever an FPop is issued, even though no destination F register is modified. The specific conditions under which a dirty bit is set pessimistically are implementation dependent.

**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.

**IMPL. DEP. #403-S10(b):** An UltraSPARC Architecture 2005 virtual processor may set *FPRS.dl* pessimistically; that is, it may be set whenever an FPop is issued, even though no destination F register is modified. The specific conditions under which a dirty bit is set pessimistically are implementation dependent.

<b>Implementation Note</b>	If an instruction that normally writes to the F registers is executed and causes an <i>fp_disabled</i> exception, an UltraSPARC Architecture 2005 implementation still sets the “dirty” bit ( <i>FPRS.du</i> or <i>FPRS.dl</i> ) corresponding to the destination register to ‘1’.
----------------------------	--

<b>Forward Compatibility Note</b>	It is expected that in future revisions to the UltraSPARC Architecture, if an instruction that normally writes to the F registers is executed and causes an <i>fp_disabled</i> exception the “dirty” bit ( <i>FPRS.du</i> or <i>FPRS.dl</i> ) corresponding to the destination register will be left <i>unchanged</i> .
-----------------------------------	---

## 6.5.7 Performance Control Register (PCR<sup>P</sup>) (ASR 16) D2

The PCR is used to control performance monitoring events collected in counter pairs, which are accessed via the Performance Instrumentation Counter (PIC) register (ASR 17) (see page 79). Unused PCR bits read as zero; they should be written only with zeroes or with values previously read from them.

When the virtual processor is operating in privileged mode (PSTATE.priv = 1 and HPSTATE.hpriv = 0) or hyperprivileged mode (HPSTATE.hpriv = 1), PCR may be freely read and written by software.

When the virtual processor is operating in nonprivileged mode (PSTATE.priv = 0), an attempt to access PCR (using a RDPCR or WRPCR instruction) results in a *privileged\_opcode* exception (impl. dep. #250-U3-Cs10).

The PCR is illustrated in FIGURE 6-14 and described in TABLE 6-11.



FIGURE 6-14 Performance Control Register (PCR) (ASR 16)

**IMPL. DEP. #207-U3:** The values and semantics of bits 47:32, 26:17, and bit 3 of the PCR are implementation dependent.

TABLE 6-11 PCR Bit Description

Bit	Field	Description
47:32	—	These bits are implementation dependent (impl. dep #207-U3).
26:17	—	These bits are implementation dependent (impl. dep. #207-U3).
16:11	su	Six-bit field selecting 1 of 64 event counts in the upper half (bits {63:32}) of the PIC.
9:4	sl	Six-bit field selecting 1 of 64 event counts in the lower half (bits {31:0}) of the PIC.
3	—	This bit is implementation dependent (impl. dep. #207-U3).
2	ut	User Trace Enable. If set to 1, events in nonprivileged (user) mode are counted.
1	st	System Trace Enable. If set to 1, events in privileged (system) mode are counted.
		<b>Notes:</b> If both PCR.ut and PCR.st are set to 1, all selected events are counted. If both PCR.ut and PCR.st are zero, counting is disabled. PCR.ut and PCR.st are global fields which apply to all PIC pairs.
0	priv	Privileged. Controls access to the PIC register (via RDPIC or WRPIC instructions). If PCR.priv = 0, an attempt to access PIC will succeed regardless of the privilege state (PSTATE.priv). If PCR.priv = 1, access to PIC is restricted to privileged software; that is, an attempt to access PIC while PSTATE.priv = 1 will succeed, but an attempt to access PIC while PSTATE.priv = 0 will result in a <i>privileged_action</i> exception.

## 6.5.8 Performance Instrumentation Counter (PIC) Register (ASR 17) A2

PIC contains two 32-bit counters that count performance-related events (such as instruction counts, cache misses, TLB misses, and pipeline stalls). Which events are actively counted at any given time is selected by the PCR register.

The difference between the values read from the PIC register at two different times reflects the number of events that occurred between register reads. Software can only rely on the difference in counts between two PIC reads to get an accurate count, not on the difference in counts between a PIC write and a PIC read.

PIC is normally a nonprivileged-access, read/write register. However, if the priv bit of the PCR (ASR 16) is set, attempted access by nonprivileged (user) code causes a *privileged\_action* exception.

Multiple PICs may be implemented. Each is accessed through ASR 17, using an implementation-dependent PIC pair selection field in PCR (ASR 16) (impl. dep. #207-U3). Read/write access to the PIC will access the picu/picl counter pair selected by PCR.

The PIC is described below and illustrated in FIGURE 6-15.

Bit	Field	Description
63:32	picu	32-bit counter representing the count of an event selected by the su field of the Performance Control Register (PCR) (ASR 16).
31:0	picl	32-bit counter representing the count of an event selected by the sl field of the Performance Control Register (PCR) (ASR 16).



FIGURE 6-15 Performance Instrumentation Counter (PIC) (ASR 17)

**Counter Overflow.** On overflow, the effective counter wraps to 0, SOFTINT register bit 15 is set to 1, and an interrupt level 15 trap is generated if not masked by PSTATE.ie and PIL. The counter overflow trap is triggered on the transition from value FFFF FFFF<sub>16</sub> to value 0.

## 6.5.9 General Status Register (GSR) (ASR 19) A1

The General Status Register<sup>1</sup> (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 299) and written by the WRGSR instruction (see *Write Ancillary State Register* on page 369).

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 6-16 and described in TABLE 6-12.




FIGURE 6-16 General Status Register (GSR) (ASR 19)

TABLE 6-12 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: <table style="margin-left: auto; margin-right: auto; border-collapse: collapse;"> <thead> <tr> <th style="padding: 2px;">irnd</th> <th style="padding: 2px;">Round toward ...</th> </tr> </thead> <tbody> <tr> <td style="padding: 2px;">0</td> <td style="padding: 2px;">Nearest (even, if tie)</td> </tr> <tr> <td style="padding: 2px;">1</td> <td style="padding: 2px;">0</td> </tr> <tr> <td style="padding: 2px;">2</td> <td style="padding: 2px;">+ ∞</td> </tr> <tr> <td style="padding: 2px;">3</td> <td style="padding: 2px;">− ∞</td> </tr> </tbody> </table>	irnd	Round toward ...	0	Nearest (even, if tie)	1	0	2	+ ∞	3	− ∞
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.										

## 6.5.10 SOFTINT<sup>P</sup> Register (ASRs 20 D2, 21 D2, 22 D1)

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

<sup>1</sup> This register was (inaccurately) referred to as the "Graphics Status Register" in early UltraSPARC implementations

SOFTINT can be read with a RDSOFTINT instruction (see *Read Ancillary State Register* on page 299) and written with a WRSOFTINT, WRSOFTINT\_SET, or WRSOFTINT\_CLR instruction (see *Write Ancillary State Register* on page 369). 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 6-17 and described in TABLE 6-13.




FIGURE 6-17 SOFTINT Register (ASR 22)

TABLE 6-13 SOFTINT Bit Description

Bit	Field	Description
16	sm	When the STICK_CMPR (ASR 25) register's int_dis (interrupt disable) field is 0 (that is, System Tick Compare is enabled) and its stick_cmpr 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_CMPR<sup>P</sup>) Register (ASR 25)</i> on page 84 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><b>Notes:</b> A level-14 interrupt (<i>interrupt_level_14</i>) can be triggered by SOFTINT.sm, SOFTINT.tm, 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 &lt; n) and (PSTATE.ie = 1 and (HPSTATE.hpriv = 0)).</p>		
0	tm	When the TICK_CMPR (ASR 23) register's int_dis (interrupt disable) field is 0 (that is, Tick Compare is enabled) and its tick_cmpr field matches the value in the TICK register, then the tm ("TICK match") field in SOFTINT is set to 1 and a level-14 interrupt ( <i>interrupt_level_14</i> ) is generated. See <i>Tick Compare (TICK_CMPR<sup>P</sup>) Register (ASR 23)</i> on page 83 for details. SOFTINT.tm can also be directly written to 1 by software.

Setting any of SOFTINT.sm, SOFTINT.int\_level{13} (SOFTINT{14}), or SOFTINT.tm to 1 causes a level-14 interrupt (*interrupt\_level\_14*). However, those three bits are independent; setting any one of them does not affect the other two.

See *Software Interrupt Register (SOFTINT)* on page 496 for additional information regarding the SOFTINT register.

### 6.5.10.1 SOFTINT\_SET<sup>P</sup> Pseudo-Register (ASR 20) (D2)

A Write State register instruction to ASR 20 (WRSOFTINT\_SET) atomically sets selected bits in the privileged SOFTINT Register (ASR 22) (see page 80). 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 6-18 illustrates the SOFTINT\_SET pseudo-register.




FIGURE 6-18 SOFTINT\_SET Pseudo-Register (ASR 20)

### 6.5.10.2 SOFTINT\_CLR<sup>P</sup> Pseudo-Register (ASR 21) (D2)

A Write State register instruction to ASR 21 (WRSOFTINT\_CLR) atomically clears selected bits in the privileged SOFTINT register (ASR 22) (see page 80). 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 6-19 illustrates the SOFTINT\_CLR pseudo-register.




FIGURE 6-19 SOFTINT\_CLR Pseudo-Register (ASR 21)

## 6.5.11 Tick Compare (TICK\_CMPR<sup>P</sup>) Register (ASR 23) D1

The privileged TICK\_CMPR register allows system software to cause a trap when the TICK register reaches a specified value. Nonprivileged accesses to this register cause a *privileged\_opcode* exception (see *Exception and Interrupt Descriptions* on page 481).

After a power-on reset trap, the `int_dis` bit is set to 1 (disabling Tick Compare interrupts) and the value of the `tick_cmpr` field is undefined.

The TICK\_CMPR register is illustrated in FIGURE 6-20 and described in TABLE 6-14.




FIGURE 6-20 TICK\_CMPR Register

TABLE 6-14 TICK\_CMPR Register Description

Bit	Field	Description
63	<code>int_dis</code>	Interrupt Disable. If <code>int_dis</code> = 0, TICK compare interrupts are enabled and if <code>int_dis</code> = 1, TICK compare interrupts are disabled.
62:0	<code>tick_cmpr</code>	Tick Compare Field. When this field exactly matches the value in <code>TICK.counter</code> and <code>TICK_CMPR.int_dis</code> = 0, <code>SOFTINT.tm</code> is set to 1. This has the effect of posting a level-14 interrupt to the virtual processor, which causes an <i>interrupt_level_14</i> trap when ( <code>PIL</code> < 14) and ( <code>PSTATE.ie</code> = 1 and <code>HPSTATE.hpriv</code> = 0). The level-14 interrupt handler must check <code>SOFTINT{14}</code> , <code>SOFTINT{0}</code> ( <code>tm</code> ), and <code>SOFTINT{16}</code> ( <code>sm</code> ) to determine the source of the level-14 interrupt.

## 6.5.12 System Tick (STICK) Register (ASR 24) D1

The System Tick (STICK) register provides a counter that is synchronized across a system, useful for timestamping. The counter field of the STICK register is a 63-bit counter that increments at a rate determined by a clock signal external to the processor.

Bit 63 of the STICK register is the nonprivileged trap (`npt`) bit, which controls access to the STICK register by nonprivileged software.

The STICK register is illustrated in FIGURE 6-21 and described below.




FIGURE 6-21 STICK Register

Hyperprivileged software can always read the STICK register with the RDSTICK instruction and write it with the WRSTICK instruction.

Privileged software can always read the STICK register with the RDSTICK instruction. 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 by using the RDSTICK instruction, but only when nonprivileged access to STICK is enabled (STICK.npt = 0) by hyperprivileged software. If nonprivileged access is disabled (STICK.npt = 1), 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.

After the STICK register is written, reading the STICK register returns a value incremented (by 1 or more) from the last value written, rather than from some previous value of counter.

After a power-on reset trap, STICK.npt is set to 1 and the value of STICK.counter is undefined.

**Note** | The STICK register is unaffected by any reset other than a power-on reset.

### 6.5.13 System Tick Compare (STICK\_CMPR<sup>P</sup>) Register (ASR 25) D2

The privileged STICK\_CMPR register allows system software to cause a trap when the STICK register reaches a specified value. Nonprivileged accesses to this register cause a *privileged\_opcode* exception (see *Exception and Interrupt Descriptions* on page 481).

After a power-on reset trap, the int\_dis bit is set to 1 (disabling System Tick Compare interrupts), and the stick\_cmpr field is undefined.

The System Tick Compare Register is illustrated in FIGURE 6-22 and described in TABLE 6-15.




FIGURE 6-22 STICK\_CMPRP Register

TABLE 6-15 STICK\_CMPRP Register Description

Bit	Field	Description
63	int_dis	Interrupt Disable. If set to 1, STICK_CMPRP interrupts are disabled.
62:0	stick_cmpr	System Tick Compare Field. When this field exactly matches STICK.counter and STICK_CMPRP.int_dis = 0, SOFTINT.sm is set to 1. This has the effect of posting a level-14 interrupt to the virtual processor, which causes an <i>interrupt_level_14</i> trap when (PIL < 14) and (PSTATE.ie = 1). The level-14 interrupt handler must check SOFTINT{14}, SOFTINT{0} (tm), and SOFTINT{16} (sm) to determine the source of the level-14 interrupt.

## 6.6 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.

**IMPL. DEP. #126-V9-Ms10:** Privileged registers CWP, CANSAVE, CANRESTORE, OTHERWIN, and CLEANWIN contain values in the range 0 to  $N\_REG\_WINDOWS - 1$ . An attempt to write a value greater than  $N\_REG\_WINDOWS - 1$  to any of these registers causes an implementation-dependent value between 0 and  $N\_REG\_WINDOWS - 1$  (inclusive) to be written to the register. Furthermore, an attempt to write a value greater than  $N\_REG\_WINDOWS - 2$  violates the register window state definition in *Register Window State Definition* on page 88.

Although the width of each of these five registers is architecturally 5 bits, the 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. All five registers should have the same width.

For UltraSPARC Architecture 2005 processors,  $N\_REG\_WINDOWS = 8$ . Therefore, each

register window state register is implemented with 3 bits, the maximum value for CWP and CLEANWIN is 7, and the maximum value for CANSAVE, CANRESTORE, and OTHERWIN is 6. When these registers are written by the WRPR instruction, bits 63:3 of the data written are ignored.

For details of how the window-management registers are used, see *Register Window Management Instructions* on page 129.

**Programming Note** CANSAVE, CANRESTORE, OTHERWIN, and CLEANWIN must never be set to a value greater than  $N\_REG\_WINDOWS - 2$  on an UltraSPARC Architecture virtual processor. Setting any of these to a value greater than  $N\_REG\_WINDOWS - 2$  violates the register window state definition in *Register Window State Definition* on page 88. Hardware is not required to enforce this restriction; it is up to system software to keep the window state consistent.

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

## 6.6.1 Current Window Pointer (CWP<sup>P</sup>) Register (PR 9)

(D1)

The privileged CWP register, shown in FIGURE 6-23, is a counter that identifies the current window into the array of integer registers. See *Register Window Management Instructions* on page 129 and Chapter 12, *Traps*, for information on how hardware manipulates the CWP register.




FIGURE 6-23 Current Window Pointer Register

## 6.6.2 Savable Windows (CANSAVE<sup>P</sup>) Register (PR 10)

(D1)

The privileged CANSAVE register, shown in FIGURE 6-24, 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 6-24 CANSAVE Register, Figure 5-24, page 88

### 6.6.3 Restorable Windows (CANRESTORE<sup>P</sup>) Register (PR 11) (D1)

The privileged CANRESTORE register, shown in FIGURE 6-25, 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 6-25 CANRESTORE Register

### 6.6.4 Clean Windows (CLEANWIN<sup>P</sup>) Register (PR 12) (D1)

The privileged CLEANWIN register, shown in FIGURE 6-26, contains the number of windows that can be used by the SAVE instruction without causing a *clean\_window* exception.



FIGURE 6-26 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.

### 6.6.5 Other Windows (OTHERWIN<sup>P</sup>) Register (PR 13) (D1)

The privileged OTHERWIN register, shown in FIGURE 6-27, 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 6-27 OTHERWIN Register

## 6.6.6 Window State (WSTATE<sup>P</sup>) Register (PR 14) D1

The privileged WSTATE register, shown in FIGURE 6-28, 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 6-28 WSTATE Register

## 6.6.7 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 85. Those registers are affected by the instructions described in *Register Window Management Instructions* on page 129. Privileged software can read/write these state registers directly by using RDPR/WRPR instructions.

### 6.6.7.1 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 6-3 on page 54 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 6-3, 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 6-3) 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 6-3) 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 6-3) 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.

<b>Programming Note</b>	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 $N\_REG\_WINDOWS - 1$ .
-------------------------	---

### 6.6.7.2 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 491 for a detailed description of how fill, spill, and *clean\_window* traps support register windowing.

## 6.7 Non-Register-Window PR State Registers

The registers described in this section are visible only to software running in privileged or hyperprivileged mode (that is, when `PSTATE.priv = 1` or `HPSTATE.hpriv = 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.

### 6.7.1 Trap Program Counter (TPC<sup>P</sup>) Register (PR 0) D1

The privileged Trap Program Counter register (TPC; FIGURE 6-29) contains the program counter (PC) from the previous trap level. There are `MAXTL` 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
<b>TPC<sub>1</sub><sup>P</sup></b>	pc_high62 (PC{63:2} from trap while TL = 0)	00
<b>TPC<sub>2</sub><sup>P</sup></b>	pc_high62 (PC{63:2} from trap while TL = 1)	00
<b>TPC<sub>3</sub><sup>P</sup></b>	pc_high62 (PC{63:2} from trap while TL = 2)	00
⋮	⋮	⋮
<b>TPC<sub>MAXTL</sub><sup>P</sup></b>	pc_high62 (PC{63:2} from trap while TL = MAXTL - 1)	00

63
2 1 0

FIGURE 6-29 Trap Program Counter Register Stack

After a power-on reset, the contents of `TPC[1]` through `TPC[MAXTL]` are undefined. During normal operation, the value of `TPC[n]`, where `n` is greater than the current trap level (`n > TL`), is undefined.

TABLE 6-16 lists the events that cause TPC to be read or written.

**TABLE 6-16** Events that involve TPC, when executing with TL =  $n$ .

Event	Effect
Trap	TPC[ $n + 1$ ] $\leftarrow$ PC
RETRY instruction	PC $\leftarrow$ TPC[ $n$ ]
RDPR (TPC)	R[rd] $\leftarrow$ TPC[ $n$ ]
WRPR (TPC)	TPC[ $n$ ] $\leftarrow$ <i>value</i>
Power-on reset (POR)	All TPC values are left undefined

## 6.7.2 Trap Next PC (TNPC<sup>P</sup>) Register (PR 1) D1

The privileged Trap Next Program Counter register (TNPC; FIGURE 6-29) is the next program counter (NPC) from the previous trap level. There are  $MAXTL$  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 <sub>1</sub> <sup>P</sup>	npc_high62 (NPC{63:2} from trap while TL = 0)	00
TNPC <sub>2</sub> <sup>P</sup>	npc_high62 (NPC{63:2} from trap while TL = 1)	00
TNPC <sub>3</sub> <sup>P</sup>	npc_high62 (NPC{63:2} from trap while TL = 2)	00
⋮	⋮	⋮
TNPC <sub>MAXTL</sub> <sup>P</sup>	npc_high62 (NPC{63:2} from trap while TL = $MAXTL - 1$ )	00

63 2 1 0

**FIGURE 6-30** Trap Next Program Counter Register Stack

After a power-on reset, the contents of TNPC[1] through TNPC[ $MAXTL$ ] are undefined. During normal operation, the value of TNPC[ $n$ ], where  $n$  is greater than the current trap level ( $n > TL$ ), is undefined.

TABLE 6-17 lists the events that cause TNPC to be read or written.

**TABLE 6-17** 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[ $n$ ] $\leftarrow$ <i>value</i>
Power-on reset (POR)	All TNPC values are left undefined

## 6.7.3 Trap State (TSTATE<sup>P</sup>) Register (PR 2) D1

The privileged Trap State register (TSTATE; FIGURE 6-31) contains the state from the previous trap level, comprising the contents of the GL, CCR, ASI, CWP, and PSTATE registers from the previous trap level. There are *MAXTL* 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.

	RW	RW	RW	R	RW	R	RW							
TSTATE <sub>1</sub> <sup>P</sup>	gl (GL from TL = 0)	ccl (CCR from TL = 0)	asi (ASI from TL = 0)	—	pstate (PSTATE from TL = 0)	—	cwp (CWP from TL = 0)							
TSTATE <sub>2</sub> <sup>P</sup>	gl (GL from TL = 1)	ccl (CCR from TL = 1)	asi (ASI from TL = 1)	—	pstate (PSTATE from TL = 1)	—	cwp (CWP from TL = 1)							
TSTATE <sub>3</sub> <sup>P</sup>	gl (GL from TL = 2)	ccr (CCR from TL = 2)	asi (ASI from TL = 2)	—	pstate (PSTATE from TL = 2)	—	cwp (CWP from TL = 2)							
: <sup>P</sup>	:	:	:	:	:	:	:							
TSTATE <sub>MAXPTL</sub> <sup>P</sup>	gl (GL from TL = MAXPTL - 1)	ccr (CCR from TL = MAXPTL - 1)	asi (ASI from TL = MAXPTL - 1)	—	pstate (PSTATE from TL = MAXPTL - 1)	—	cwp (CWP from TL = MAXPTL - 1)							
TSTATE <sub>MAXPTL+1</sub> <sup>H</sup>	gl (GL from TL = MAXPTL)	ccr (CCR from TL = MAXPTL)	asi (ASI from TL = MAXPTL)	—	pstate (PSTATE from TL = MAXPTL)	—	cwp (CWP from TL = MAXPTL)							
: <sup>H</sup>	:	:	:	:	:	:	:							
TSTATE <sub>MAXTL</sub> <sup>H</sup>	gl (GL from TL = MAXTL - 1)	ccr (CCR from TL = MAXTL - 1)	asi (ASI from TL = MAXTL - 1)	—	pstate (PSTATE from TL = MAXTL - 1)	—	cwp (CWP from TL = MAXTL - 1)							
	42	40	39	32	31	24	23	21	20	8	7	5	4	0

TABLE 6-18

FIGURE 6-31 Trap State (TSTATE) Register Stack

After a power-on reset the contents of TSTATE[1] through TSTATE[MAXTL] are undefined. 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 of the addition of additional bits in the PSTATE register in the UltraSPARC Architecture, a 13-bit PSTATE value is stored in TSTATE instead of the 10-bit value specified in the SPARC V9 architecture.


TABLE 6-19 lists the events that cause TSTATE to be read or written.

**TABLE 6-19** 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[ $n$ ] $\leftarrow$ <i>value</i>
Power-on reset (POR)	All TSTATE values are left undefined

## 6.7.4 Trap Type (TT<sup>P</sup>) Register (PR 3) D1

The privileged Trap Type register (TT; see FIGURE 6-32) contains the trap type of the trap that caused entry to the current trap level. On a reset trap, the TT register contains the trap type of the reset (see TABLE 12-2 on page 444). There are  $MAXTL$  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 6-32** Trap Type Register Stack

After a power-on reset the contents of TT[1] through TT[ $MAXTL - 1$ ] are undefined and TT[ $MAXTL$ ] = 001<sub>16</sub>. During normal operation, the value of TT[ $n$ ], where  $n$  is greater than the current trap level ( $n > TL$ ), is undefined.


TABLE 6-20 lists the events that cause TT to be read or written.

**TABLE 6-20** Events that involve TT, when executing with TL = n.

Event	Effect
Trap	$TT[n + 1] \leftarrow (\text{trap type})$
RDPR (TT)	$R[\text{rd}] \leftarrow TT[n]$
WRPR (TT)	$TT[n] \leftarrow \text{value}$
Power-on reset (POR)	TT values $TT[1]$ through $TT[\text{MAXTL} - 1]$ are left undefined; $TT[\text{MAXTL}] \leftarrow 001_{16}$ .

## 6.7.5 Trap Base Address (TBA<sup>P</sup>) Register (PR 5) D1

The privileged Trap Base Address register (TBA), shown in FIGURE 6-33, 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 6-33** 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 453.

## 6.7.6 Processor State (PSTATE<sup>P</sup>) Register (PR 6) D1

The privileged Processor State register (PSTATE), shown in FIGURE 6-34, contains control fields for the current state of the virtual processor. There is only one instance of the PSTATE register per virtual processor.



**FIGURE 6-34** PSTATE Field

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.

**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 120. 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. During a virtual processor trap to hyperprivileged mode, the `PSTATE.tle` bit is *not* copied into `PSTATE.cle` of the new `PSTATE` register and is unused.

**Memory Model (mm).** This 2-bit field determines the memory model in use by the virtual processor. The defined values for an UltraSPARC Architecture virtual processor are listed in TABLE 6-21.

**TABLE 6-21** `PSTATE.mm` Encodings

<code>mm</code> 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 UltraSPARC Architecture processor is implementation dependent. If the  $10_2$  model is supported, then when

PSTATE.mm = 10<sub>2</sub> the implementation must correctly execute software that adheres to the RMO model described in *The SPARC Architecture Manual-Version 9*. If the 11<sub>2</sub> 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.

<b>SPARC V9 Compatibility Notes</b>	<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 UltraSPARC 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 UltraSPARC Architecture 2005 implementations. All software written to run correctly under RMO will run correctly under TSO on an UltraSPARC Architecture 2005 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 V9 processor, by masking out (zeroing) bits 63:32 of virtual addresses at appropriate times.

When PSTATE.am = 0, the full 64 bits of all instruction and data addresses are *preserved* at all times.

When PSTATE.am = 1, bits 63:32 of instruction and data virtual addresses are masked out (treated as 0).

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

Instances in which the more-significant 32 bits of a virtual address **are masked** include:

- Before any (virtual, real, or physical) data address is sent out of the virtual processor (notably, to the memory system, which includes MMU, internal caches, and external caches); this includes ASI accesses using `ASI_AS_IF_USER*` in privileged or hyperprivileged mode.
- 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 virtual address is sent to a watchpoint comparator

<b>Programming Note</b>	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 a bypassing ASI (`ASI_*REAL_*` or `ASI_*PHYS_*`) is used in a load or store instruction (see `ASI 1416`, `ASI_REAL`, for an example).
- When an exception occurs and any address (virtual, real, or physical) is written to the Data Synchronous Fault Address register (`DSFAR`) (impl.dep. #241-U3)

<b>Programming Note</b>	If a memory access is initiated when <code>PSTATE.am = 1</code> , the memory system will only see a 32-bit memory address. Therefore, if such a memory access causes an exception or error, the memory system will (is only able to) report a 32-bit address in the <code>DSFAR</code> register (64-bit address with the more-significant 32 bits set to 0).
-------------------------	--

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

<b>Forward Compatibility Note</b>	This behavior is expected to change in the next revision of the architecture, such that implementations will explicitly mask out (not preserve) the more-significant 32 bits, in this case.
-----------------------------------	---

- When NPC is incremented to `NPC + 4` during execution of an instruction that is not a taken control transfer

<b>Forward Compatibility Note</b>	This behavior is expected to change in the next revision of the architecture, such that implementations will explicitly mask out (not preserve) the more-significant 32 bits, in this case.
-----------------------------------	---

- When a WRPR instruction writes to TPC[TL] or TNPC[TL]

<b>Programming Note</b>	Since writes to PSTATE are nondelayed (see page 94), a change to PSTATE.am can affect the address of the next instruction executed. Specifically, if a WRPR to the PSTATE register changes the value of PSTATE.am from '0' to '1', and the more-significant 32 bits of NPC when the WRPR began execution were nonzero, then the next instruction that executes after the WRPR will not be from the address in NPC when the WRPR began execution but rather from that address truncated to a 32-bit address (NPC with its more-significant 32 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 executed<sup>1</sup>, 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).

<b>Programming Note</b>	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].
-------------------------	---

**Privileged Mode (priv).** When PSTATE.priv = 1 and HPSTATE.hpriv = 0, the virtual processor is operating in privileged mode.

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

When HPSTATE.hpriv = 1, the virtual processor is operating in hyperprivileged mode, independent of the state of PSTATE.priv. Hyperprivileged mode provides a superset of the capabilities of privileged 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.

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

Outstanding disrupting exceptions that are destined for hyperprivileged mode can only cause a trap when the virtual processor is not in hyperprivileged mode, or when it is in hyperprivileged mode and `PSTATE.ie = 1`. At all other times, they are held pending. For more details, see *Conditioning of Disrupting Traps* on page 448

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

## 6.7.7 Trap Level Register (TL<sup>P</sup>) (PR 7) D1

The privileged Trap Level register (TL; FIGURE 6-35) 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 6-35 Trap Level Register

The maximum valid value that the TL register may contain is *MAXTL*, 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 visible to privileged software). In a typical implementation *MAXPTL* = *MAXPGL* (see impl. dep. #401-S10). The architectural parameter *MAXTL* is a constant for each implementation; its legal values are from 4 to 7 (supporting from 4 to 7 levels of saved trap state). Architecturally, *MAXPTL* must be  $\geq 2$ , *MAXTL* must be  $\geq 4$ , and *MAXTL* must be  $> MAXPTL$ .

In an UltraSPARC Architecture 2005 implementation, *MAXPTL* = 2 and *MAXTL* = 6. See Chapter 12, *Traps*, for more details regarding the TL register.

After a power-on rest (POR), TL is set to *MAXTL*.

The effect of writing to TL with a WRPR instruction is summarized in TABLE 6-22.

**TABLE 6-22** Effect of WRPR of Value  $x$  to Register TL

Value $x$ Written with WRPR	Privilege Level when Executing WRPR		
	Nonprivileged	Privileged	Hyperprivileged
$x \leq MAXPTL$	<i>privileged_opcode</i> exception	TL $\leftarrow x$	TL $\leftarrow x$
$MAXPTL < x \leq MAXTL$		TL $\leftarrow MAXPTL$ (no exception generated)	
$x > MAXTL$		TL $\leftarrow MAXTL$ (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 UltraSPARC Architecture implementation only needs to implement sufficient bits in the TL register to encode the maximum trap level value. In an implementation where  $MAXTL \leq 7$ , bits 63:3 of data written to the TL register using the WRPR instruction are ignored; only the least-significant three bits (bits 2:0) of TL are actually written. For example, if  $MAXTL = 6$ , writing a value of  $09_{16}$  to the TL register causes a value of  $1_{16}$  to actually be stored in TL.

**Implementation Note**  $MAXPTL = 2$  for all UltraSPARC Architecture 2005 processors. Writing a value between 3 and 7 to the TL register in privileged mode causes a 2 to be stored in TL.

**Implementation Note**  $MAXTL = 6$  for all UltraSPARC Architecture 2005 processors. Writing a value of 7 to the TL register in hyperprivileged mode causes a 6 to be stored in TL.

**Programming Note** Although it is possible for hyperprivileged software to set  $TL > MAXPTL$  for privileged or nonprivileged software<sup>†</sup>, an UltraSPARC Architecture virtual processor's behavior when executing with  $TL > MAXPTL$  outside of hyperprivileged mode is undefined.

Although it is possible for privileged or hyperprivileged software to set  $TL > 0$  for nonprivileged software<sup>†</sup>, an UltraSPARC Architecture virtual processor's behavior when executing with  $TL > 0$  in nonprivileged mode is undefined.

<sup>†</sup> by executing a WRPR to TSTATE followed by DONE instruction or RETRY instruction or a JMPL/WRHPR instruction pair.

## 6.7.8 Processor Interrupt Level (PIL<sup>P</sup>) Register (PR 8)

(D1)

The privileged Processor Interrupt Level register (PIL; see FIGURE 6-36) 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 459 for a list of exception and interrupt priorities.




FIGURE 6-36 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.

## 6.7.9 Global Level Register (GL<sup>P</sup>) (PR 16) (D1)

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

FIGURE 6-37 illustrates the Global Level register.




FIGURE 6-37 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 *MAXGL*, where *MAXGL* 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 visible to privileged software). In a typical implementation *MAXPGL* = *MAXPTL* (see impl. dep. #101-V9-CS10). The architectural parameter *MAXGL* is a constant for each implementation; its legal values are from 3 to 7 (supporting from 4 to 8 sets of global registers). Architecturally, *MAXPGL* must be  $\geq 2$  and *MAXGL* must be  $> \text{MAXPGL}$ .

In all UltraSPARC Architecture 2005 implementations,  $MAXPGL = 2$  and  $MAXGL = 3$ .  $MAXGL$  (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  $MAXGL$  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,  $MAXTL$  and  $MAXGL$  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  $GL$ .

In hyperprivileged mode, attempting to write a value greater than  $MAXGL$  to the  $GL$  register causes  $MAXGL$  to be written to  $GL$ .

When a  $DONE$  or  $RETRY$  instruction is executed in *privileged* mode and  $HTSTATE[TL].hpstate.hpriv = 0$  (which will cause the  $DONE$  or  $RETRY$  to return the virtual processor to nonprivileged or privileged mode), the value of  $GL$  restored from  $TSTATE[TL]$  saturates at  $MAXPGL$ . That is, if the value in  $TSTATE[TL].gl$  is greater than  $MAXPGL$ , then  $MAXPGL$  is substituted and written to  $GL$ . This protects against non-hyperprivileged software executing with  $GL > MAXPGL$ .

<b>Programming Note</b>	Although it is possible for hyperprivileged software to set $GL > MAXPGL$ for privileged or nonprivileged software <sup>†</sup> , executing with $GL > MAXPGL$ outside of hyperprivileged mode is an illegal state and the behavior of a virtual processor in that state is undefined. <sup>†</sup> by executing a $WRPR$ that modifies $GL$ , followed by a $JMPL/WRHPR$ instruction pair (it is not possible to set $GL > MAXPGL$ using $DONE/RETRY$ )
-------------------------	---

The effect of writing to GL with a WRPR instruction is summarized in TABLE 6-23.

**TABLE 6-23** Effect of WRPR to Register GL

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

If  $\text{MAXGL} < \text{MAXTL}$ , then there are fewer sets of global registers than trap levels. In this case, if a trap occurs while  $\text{GL} = \text{MAXGL}$ , GL will have the same value before the trap occurs and in the software that handles the trap. Trap handler software must detect this case and safely save any global register before the trap handler writes to it. The Hyperprivileged Scratchpad registers (see *Privileged Scratchpad Registers (ASI\_SCRATCHPAD)* on page 431) may be useful in such cases.

**Programming Note**

An UltraSPARC Architecture implementation only needs to implement sufficient bits in the GL register to encode the maximum global level value. In an implementation where  $\text{MAXGL} \leq 7$ , bits 63:3 of data written to the GL register using the WRPR instruction are ignored; only the least-significant three bits (bits 2:0) are actually written to GL. For example, if  $\text{MAXGL} = 7$ , writing a value of  $9_{16}$  to the TL register causes a value of  $1_{16}$  to actually be stored in GL.

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.

During power-on reset (POR), the value of GL is set to MAXGL. During all other resets, GL is incremented (the same behavior as TL).



## 6.8 HPR State Registers

The registers described in this section can be directly accessed with the hyperprivileged WRHPR and RDHPR instructions.



2. Write a 0 to HPSTATE.red with a WRHPR instruction.

**Programming Note** Software should not write 0 to HPSTATE.red in the delay slot of a DCTI (e.g. JMPL instruction). Exiting RED\_state using a DONE or RETRY instruction avoids this problem entirely.

**Programming Note** HPSTATE.hpriv = 0 and HPSTATE.red = 1 is an undefined operational state. Therefore, care should be taken never to write that combination of values to HPSTATE.

**Hyperprivileged mode (hpriv).** When HPSTATE.hpriv = 1, the virtual processor is operating in hyperprivileged mode and ignores PSTATE.priv.

When HPSTATE.hpriv = 0, the processor is operating in privileged or nonprivileged mode, as determined by PSTATE.priv.

See the Programming Note on page 373, recommending that a WRHPR instruction that changes HPSTATE.priv never be executed in the delay slot of a DCTI instruction.

**Trap Level Zero trap enable (tlz).** When HPSTATE.tlz = 0, generation of *trap\_level\_zero* exceptions is disabled. When all three of the following conditions exist, a *trap\_level\_zero* exception is generated:

- HPSTATE.tlz = 1 (generation of *trap\_level\_zero* is enabled)
- the virtual processor is in nonprivileged or privileged mode (HPSTATE.hpriv = 0)
- the trap level (TL) register's value is zero (TL = 0)

**Programming Note** The purpose of *trap\_level\_zero* is to improve efficiency when descheduling a virtual processor. When a descheduling event occurs and the virtual processor is executing in privileged mode at TL > 0, hyperprivileged software can choose to enable the *trap\_level\_zero* exception (set HPSTATE.tlz ← 1) and return to privileged mode, enabling privileged software to complete its TL > 0 processing. When privileged code returns to TL = 0, this exception enables the hyperprivileged code to regain control and deschedule the virtual processor with low overhead.

## 5.7.2 Hyperprivileged Trap State (HTSTATE<sup>H</sup>) Register (HPR 1)

The Hyperprivileged Trap State register (HTSTATE; FIGURE 5-38) contains the hyperprivileged state from the previous trap level, comprising the contents of the HPSTATE register from the previous trap level. There are *MAXTL* instances of the HTSTATE register, but only one is accessible at a time. The current value in the TL register determines which instance of HTSTATE is accessible.

$HTSTATE_1^H$	—	HPSTATE from TL = 0
$HTSTATE_2^H$	—	HPSTATE from TL = 1
$HTSTATE_3^H$	—	HPSTATE from TL = 2
	:	:
$HTSTATE_{MAXTL}^H$	—	HPSTATE from TL = $MAXTL - 1$
63		11 10 0

**FIGURE 5-38** Hyperprivileged Trap State Register

An attempt to read or write the HTSTATE register when TL = 0 causes an *illegal\_instruction* exception.

After a power-on reset the contents of HTSTATE[1] through HTSTATE[MAXTL] are undefined. During normal operation the value of HTSTATE[n], when n is greater than the current trap level ( $n > TL$ ), is undefined.

TABLE 5-24 lists the events that cause HTSTATE to be read or written.

**TABLE 5-24** Events that involve HTSTATE, when executing with TL = n.

Event	Effect
Trap	$HTSTATE[n + 1]\{10:0\} \leftarrow HPSTATE$
DONE instruction	$HPSTATE \leftarrow HTSTATE[n]\{10:0\}$
RETRY instruction	$HPSTATE \leftarrow HTSTATE[n]\{10:0\}$
RDHPR (HTSTATE)	$R[rd] \leftarrow HTSTATE[n]$
WRHPR (HTSTATE)	$HTSTATE[n] \leftarrow value$
Power-on reset (POR)	All HTSTATE values are left undefined

### 5.7.3 Hyperprivileged Interrupt Pending (HINTP<sup>H</sup>) Register (HPR 3)

The hyperprivileged HINTP register provides a mechanism for hyperprivileged software to determine that an *hstick\_match* interrupt is pending while PSTATE.ie = 0 and to clear the interrupt without having to first set PSTATE.ie = 1 and take a disrupting trap.

When HINTP.hsp = 1, a match between STICK and HSTICK\_CMPR has occurred with the match enabled (see *System Tick Compare (STICK\_CMPR<sup>P</sup>) Register (ASR 25)* on page 84), causing an *hstick\_match* exception to be generated.

When HINTP.hsp = 0, no interrupt is pending due to a match between STICK and HSTICK\_CMPR. An interrupt might have existed earlier but was cleared by software.

The format of the HINTP register is illustrated in FIGURE 5-39.




FIGURE 5-39 Hyperprivileged Interrupt Pending (HINTP) Register Format

## 5.7.4 Hyperprivileged Trap Base Address (HTBA<sup>H</sup>) Register (HPR 5) (D1)

The Hyperprivileged Trap Base Address register (HTBA), shown in FIGURE 5-40, provides the most significant 50 bits (bits 63:14) of the physical address used to select the trap vector for a trap that is to be serviced in hyperprivileged mode. The least significant 14 bits of HTBA always read as zero, and writes to them are ignored.




FIGURE 5-40 Hyperprivileged Trap Base Address Register

Details on how the full address for a trap vector is generated, using HTBA and other state, are provided in *Trap-Table Entry Address to Hyperprivileged Mode* on page 454.

**IMPL. DEP. #406-S10:** It is implementation dependent whether all 50 bits of HTBA{63:14} are implemented or if only bits  $n-1:14$  are implemented. If the latter, writes to bits 63: $n$  are ignored and when HTBA is read, bits 63: $n$  read as sign-extended copies of the most significant implemented bit, HTBA{ $n - 1$ }.

See Chapter 12, *Traps*, for more details on trap vectors.

## 5.7.5 Hyperprivileged Implementation Version (HVER<sup>H</sup>) Register (HPR 6) (D2)

The Hyperprivileged Implementation Version register, shown in FIGURE 5-41, specifies the fixed parameters pertaining to a particular processor implementation and mask set. The HVER register is read-only, readable by the RDHPR instruction in hyperprivileged mode.




FIGURE 5-41 Hyperprivileged Implementation Version Register

**IMPL. DEP. #104-V9:** HVER.manuf contains a 16-bit manufacturer code. This field is optional and if not present shall read as 0. HVER.manuf may indicate the original supplier of a second-sourced processor. It is intended that the contents of HVER.manuf track the JEDEC semiconductor manufacturer code as closely as possible. If the manufacturer does not have a JEDEC semiconductor manufacturer code, SPARC International will assign a value for HVER.manuf.

**IMPL. DEP. #13-V8:** HVER.impl uniquely identifies an implementation or class of software-compatible implementations of the architecture. Values FFF0<sub>16</sub>–FFFF<sub>16</sub> are reserved and are not available for assignment.

HVER.mask specifies the current mask set revision and is chosen by the implementor. It generally increases numerically with successive releases of the processor but does not necessarily increase by 1 for consecutive releases.

**Implementation Note** | Conventionally, this field is die-specific, with bits 31:28 indicating the major mask revision number and bits 27:24 indicating the minor mask revision number.

HVER.maxgl contains the maximum number of levels of global register sets supported by an implementation (impl. dep. #401-S10), that is, *MAXGL*, the maximum value that the GL register may contain.

HVER.maxtl contains the maximum number of trap levels supported by an implementation (impl. dep. #101-V9-CS10), that is, *MAXTL*, the maximum value of the contents of the TL register.

HVER.maxwin contains the maximum index number available for use as a valid CWP value in an implementation; that is, HVER.maxwin contains the value *N\_REG\_WINDOWS* – 1 (impl. dep. #2-V8).

**SPARC V9 Compatibility Note** | The SPARC V9 VER register was replaced in the UltraSPARC Architecture by the hyperprivileged HVER register.

## 5.7.6 Hyperprivileged System Tick Compare (HSTICK\_CMPR<sup>H</sup>) Register (HPR 31)

The Hyperprivileged System Tick Compare (HSTICK\_CMPR) register allows hyperprivileged software to cause an *hstick\_match* interrupt when the STICK register reaches a specified value while HSTICK\_CMPR.int\_dis = 0. While executing in hyperprivileged mode and PSTATE.ie = 0, the interrupt is masked.

The Hyperprivileged System Tick Compare Register is illustrated in FIGURE 5-42.




FIGURE 5-42 HSTICK\_CMPR Register

The fields of HSTICK\_CMPR are described in TABLE 5-25.

TABLE 5-25 Bit Description of HSTICK\_CMPR Register

Bit(s)	Field Name	Description
63	int_dis	Interrupt Disable. If int_dis = 1, <i>hstick_match</i> interrupts are disabled; if 0, they are enabled (may occur).
62:0	hstick_cmpr	Hyperprivileged System Tick Compare Field. When HSTICK_CMPR.int_dis = 0 and the value in HSTICK_CMPR.hstick_cmpr exactly matches the value in STICK.counter, HINTP.hsp is set to 1. After that, if HINTP.hsp remains set to 1, the next time that hyperprivileged interrupts are unmasked (HPSTATE.hpriv = 0 or PSTATE.ie = 1), an <i>hstick_match</i> exception will occur.

After a power-on reset trap, the int\_dis bit is set to 1 (disabling Hyperprivileged System Tick Compare interrupts), and the hstick\_cmpr field is undefined.



## 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 111.
  - **Instruction Formats** on page 112.
  - **Instruction Categories** on page 113.
- 

### 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 to 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 120) 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

Instructions are encoded in four major 32-bit formats and several minor formats, as shown in FIGURE 6-1. For detailed formats for specific instructions, see individual instruction descriptions in the *Instructions* chapter.

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



$op = 01_2$ : *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

UltraSPARC 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. The LDFSR/LDXFSR and the STFSR/STXFSR are special load and store instructions that load or store the floating-point status 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.

<b>Programming Note</b>	For some instructions, by using <code>simm13</code> , any location in the lowest or highest 4 Kbytes of an address space can be accessed without using 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) or Store (STBLOCKF) 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 which 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 which 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 Note** | Although the architecture provides for the *LDQF\_mem\_address\_not\_aligned* exception, UltraSPARC Architecture 2005 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 Note** | Although the architecture provides for the *STQF\_mem\_address\_not\_aligned* exception, UltraSPARC Architecture 2005 implementations do not currently generate it.

### 6.3.1.2 Addressing Conventions

An UltraSPARC 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 using selected ASIs. It is also possible to change the default byte order for implicit data accesses. See *Processor State (PSTATE<sup>P</sup>) Register (PR 6)* on page 94 for more information.<sup>1</sup>

**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
<b>byte</b>	A load/store byte instruction accesses the addressed byte in both big- and little-endian modes.
<b>halfword</b>	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.

<sup>1</sup> 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.

**TABLE 6-1** Big-endian Addressing Conventions

<b>Term</b>	<b>Definition</b>
<b>word</b>	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.
<b>doubleword or extended word</b>	<p>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.</p> <p>For the deprecated integer load/store twin word instructions (LDTW, LDTWA<sup>†</sup>, 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.</p> <p><sup>†</sup>Note that the LDTXA instruction, which is not an LDTWA operation but does share LDTWA's opcode, is <i>not</i> deprecated.</p>
<b>quadword</b>	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.




**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
<b>byte</b>	A load/store byte instruction accesses the addressed byte in both big- and little-endian modes.
<b>halfword</b>	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.
<b>word</b>	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.
<b>doubleword or extended word</b>	<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<sup>†</sup>, 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><sup>†</sup>Note that the LDTXA instruction, which is not an LDTWA operation but does share LDTWA's opcode, is <i>not</i> deprecated.</p>
<b>quadword</b>	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.


**Byte**





**Halfword**







**Word**



**Doubleword /  
Extended word**



**Quadword**



**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. Note that in hyperprivileged mode, all accesses are performed using physical addresses, so there is no implicit ASI in hyperprivileged mode.

**TABLE 6-3** ASIs Used for Data Accesses and Instruction Fetches in Nonprivileged and Privileged Modes

Access Type	TL	PSTATE.cle	ASI Used
Instruction Fetch	= 0	any	ASI_PRIMARY
	> 0	any	ASI_NUCLEUS*
Non-alternate-space Load, Store, or Load-Store	= 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.

See also *Memory Addressing and Alternate Address Spaces* on page 389.

ASIs  $00_{16}$  through  $7F_{16}$  are restricted; only software with sufficient privilege is allowed to access them. ASIs  $00_{16}$ - $2F_{16}$  are accessible by both privileged and hyperprivileged software, while ASIs  $30_{16}$ - $7F_{16}$  are accessible only by hyperprivileged software. An attempt to access a restricted ASI by insufficiently-privileged software results in a *privileged\_action* exception (impl. dep #103-V9-Ms10(6)). ASIs  $80_{16}$  through  $FF_{16}$  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 (HPSTATE.hpriv, PSTATE.priv)	Result of ASI Access
$00_{16}$ - $2F_{16}$	Restricted (Privileged)	Nonprivileged (0,0)	<i>privileged_action</i> exception
		Privileged (0,1)	Valid access
		Hyperprivileged (1,x)	Valid access
$30_{16}$ - $7F_{16}$	Restricted (Hyperprivileged)	Nonprivileged (0,0)	<i>privileged_action</i> exception
		Privileged (0,1)	<i>data_access_exception</i> exception
		Hyperprivileged (1,x)	Valid access
$80_{16}$ - $FF_{16}$	Unrestricted	Nonprivileged (0,0)	Valid access
		Privileged (0,1)	Valid access
		Hyperprivileged (1,x)	Valid access

**IMPL. DEP. #29-V8:** Some UltraSPARC Architecture 2005 ASIs are implementation dependent. See TABLE 10-1 on page 409 for details.

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

An UltraSPARC 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 598 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 355 and *Tagged Subtract* on page 361 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)
- Unconditional branch
- Call and link (CALL)
- Jump and link (JMPL, 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 the next program counter, 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, depending upon whether the transfer is taken or not taken. Annulled instructions have no effect upon the program-visible state, nor can they cause a trap.

**Programming Note** 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.

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 UltraSPARC 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.

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

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
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

**TABLE 6-5** Control-Transfer Characteristics (Continued)

Instruction Group	Address Form	Delayed	Taken	Annul Bit	New PC	New NPC
CALL	PC-relative	Yes	—	—	NPC	EA
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 two bits to create a byte displacement, and adding the result to the contents of the PC.
- **Register-indirect effective address** — A register-indirect effective address computes its target address as either  $R[rs1] + R[rs2]$  if  $i = 0$ , or  $R[rs1] + \text{sign\_ext}(simm13)$  if  $i = 1$ .
- **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 is used depends on the privilege level — 7 bits are used in nonprivileged mode and 8 bits are used in privileged and hyperprivileged modes. 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 470 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 register 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 or hyperprivileged 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.

<b>Programming Note</b>	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 UltraSPARC Architecture; no new software should place a DCTI in the delay slot of another DCTI, as on future UltraSPARC Architecture implementations that construct may execute either slowly or differently than the programmer assumes it will.

<b>SPARC V8 and SPARC V9 Compatibility Note</b>	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 is deprecated</i> .
---	--

## 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 may 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:

```
fmovdg           %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 following 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
```

which eliminates 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 UltraSPARC Architecture. The privileged registers affected by these instructions are described in *Register-Window PR State Registers* on page 85.

### 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  $\neq$  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).

<b>Programming Note</b>	<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 executing 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 315 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 307 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 70.

## 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 or two 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 update one of the *fcc<sub>n</sub>* fields of the FSR instead.

The term “FPop” refers to instructions in the FPop1, 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 127.

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 128.

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<s|d|q>`, `FMOV<s|d|q>`, `FMOVcc<s|d|q>`, `FMOVr<s|d|q>`, and `FNEG<s|d|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` or `FSR.ftt = unimplemented_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 65 to see which instructions can produce an *fp\_exception\_other* exception (with `FSR.ftt = unfinished_FPop`). See *ftt = 3 (unimplemented\_FPop)* on page 65 to see which instructions can produce an *fp\_exception\_other* exception (with `FSR.ftt = unimplemented_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 UltraSPARC Architecture, the IMPDEP1 opcode space is used by VIS instructions.

In the UltraSPARC Architecture, IMPDEP2 is subdivided into IMPDEP2A and IMPDEP2B. IMPDEP2A remains implementation dependent. The IMPDEP2B opcode space is reserved for implementation of floating-point multiply-add/multiply-subtract instructions.

## 6.3.11 Reserved Opcodes and Instruction Fields

If a conforming UltraSPARC Architecture 2005 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.

- *{r=1}* If the instruction bit pattern does not encode an extension to the instruction set, but would decode as a valid instruction if nonzero bits in reserved instruction field(s) were ignored (read as 0):
  - The recommended behavior is to generate an *illegal\_instruction* exception (or, for FPop, an *fp\_exception\_other* exception with FSR.ftt = 3 (unimplemented\_FPop)).
  - Alternatively, the implementation can ignore the nonzero reserved field bits and execute the instruction as if those bits had been zero.
- *{r=1}* If the instruction bit pattern does not encode an extension to the instruction set and would still not decode as a valid instruction if nonzero bits in reserved instruction field(s) were ignored, then the instruction bit pattern is invalid and causes an exception. Specifically, attempting to execute an FPop instruction (see *Floating-Point Operate* on page 32) causes an *fp\_exception\_other* exception (with FSR.ftt = unimplemented\_FPop); attempting to execute any other invalid instruction bit pattern causes an *illegal\_instruction* exception.

<b>Forward Compatibility Note</b>	To further enhance backward (and forward) binary compatibility, the next revision of the UltraSPARC Architecture is expected to require an <i>illegal_instruction</i> exception to be generated by any instruction bit pattern that encodes neither a known UltraSPARC Architecture instruction nor an implementation-specific extension instruction (including those with nonzero bits in reserved instruction fields).
-----------------------------------	--

*{r>1}* See Appendix A, *Opcode Maps*, for an enumeration of the reserved instruction bit patterns (opcodes).

<b>Implementation Note</b>	As described above, implementations are strongly encouraged, but not strictly required, to trap on nonzero values in reserved instruction fields.
----------------------------	---

<b>Programming Note</b>	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” (“—”).
-------------------------	---



**Note: This chapter is undergoing final review; please check back later for a copy of UltraSPARC Architecture 2005 containing the final version of this chapter.**

## Instructions

---

*UltraSPARC Architecture 2005* extends the standard SPARC V9 instruction set with additional classes of instructions:

- Enhanced functionality:
  - Instructions for alignment (*Align Address* on page 147)
  - Array handling (*Three-Dimensional Array Addressing* on page 150)
  - Byte-permutation instructions ( )
  - Edge handling (*Edge Handling Instructions* on pages 168 and 170)
  - Logical operations on floating-point registers (*F Register Logical Operate (1 operand)* on page 224)
  - Partitioned arithmetic (*Fixed-point Partitioned Add* on page 216 and *Fixed-point Partitioned Subtract* on page 221)
  - Pixel manipulation (*FEXPAND* on page 184, *FPACK* on page 210, and *FPMERGE* on page 219)
- Efficient memory access
  - Partial store (*Store Partial Floating-Point* on page 341)
  - Short floating-point loads and stores (*Store Short Floating-Point* on page 344)
  - Block load and store (*Block Load* on page 245 and *Block Store* on page 328)
- Efficient interval arithmetic: SIAM (*Set Interval Arithmetic Mode* on page 319) and all instructions that reference GSR.im

TABLE 7-2 provides a quick index of instructions, alphabetically by architectural instruction name.

TABLE 7-3 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

<b>Superscript</b>	<b>Meaning</b>
D	Deprecated instruction
H	Hyperprivileged instruction
N	Nonportable instruction
P	Privileged instruction
P <sub>ASI</sub>	Privileged action if bit 7 of the referenced ASI is 0
P <sub>ASR</sub>	Privileged instruction if the referenced ASR register is privileged
P <sub>npt</sub>	Privileged action if PSTATE.priv = 0 and (S)TICK.npt = 1
P <sub>PIC</sub>	Privileged action if PCR.priv = 1

**TABLE 7-2** UltraSPARC Architecture 2005 Instruction Set - Alphabetical (1 of 2)

Page	Instruction				
146	ADD (ADDcc)	193	FMOV(s,d,q)cc	249	LDQF
146	ADDC (ADDCcc)	198	FMOV(s,d,q)R	252	LDQFA <sup>PASI</sup>
147	ALIGNADDRESS[_LITTLE]	207	FMUL(s,d,q)	240	LDSB
148	ALLCLEAN	201	FMUL8(SU,UL)x16	242	LDSBA <sup>PASI</sup>
149	AND (ANDcc)	201	FMUL8x16	240	LDSH
150	ARRAY(8,16,32)	201	FMUL8x16(AU,AL)	242	LDSHA <sup>PASI</sup>
154	Bicc	201	FMULD8(SU,UL)x16	258	LDSHORTF
156	BMASK	227	FNAND{s}	260	LDSTUB
157	BPcc	209	FNEG(s,d,q)	261	LDSTUBA <sup>PASI</sup>
160	BPr	227	FNOR{s}	240	LDSW
156	BSHUFFLE	225	FNOT(1,2){s}	242	LDSWA <sup>PASI</sup>
162	CALL	224	FONE{s}	263	LDTXA <sup>N</sup>
163	CASA <sup>PASI</sup>	227	FORNOT(1,2){s}	266	LDTW <sup>D</sup>
163	CASXA <sup>PASI</sup>	227	FOR{s}	268	LDTWA <sup>D, PASI</sup>
166	DONE <sup>P</sup>	210	FPACK(16,32, FIX)	260	LDUB
168	EDGE(8,16,32){L}cc	216	FPADD<16,32>[S]	242	LDUBA <sup>PASI</sup>
170	EDGE(8,16,32){L}N	219	FPMERGE	240	LDUH
231	F(s,d,q)TO(s,d,q)	221	FPSUB<16,32>[S]	242	LDUHA <sup>PASI</sup>
229	F(s,d,q)TOi	207	FsMULd	240	LDUW
229	F(s,d,q)TOx	228	FSQRT(s,d,q)	242	LDUWA <sup>PASI</sup>
171	FABS(s,d,q)	225	FSRC(1,2){s}	240	LDX
172	FADD(s,d,q)	233	FSUB(s,d,q)	242	LDXA <sup>PASI</sup>
173	FALIGNDATA	227	FXNOR{s}	249	LDXFSR
227	FANDNOT(1,2){s}	227	FXOR{s}	271	MEMBAR
227	FAND{s}	234	FxTO(s,d,q)	275	MOVcc
174	FBfcc <sup>D</sup>	224	FZERO{s}	279	MOVr
176	FBPfcc	235	ILLTRAP	281	MULSc <sup>D</sup>
181	FCMP(s,d,q)	236	IMPDEP2A	283	MULX
178	FCMP*<16,32>	236	IMPDEP2B	284	NOP
181	FCMPE(s,d,q)	238	INVALW	285	NORMALW
183	FDIV(s,d,q)	239	JMPL	286	OR (ORcc)
207	FdMULq	245	LDBLOCKF	286	ORN (ORNcc)
184	FEXPAND	249	LDDF	287	OTHERW
185	FiTO(s,d,q)	252	LDDFA <sup>PASI</sup>	288	PDIST
186	FLUSH	249	LDF	289	POPC
190	FLUSHW	252	L DFA <sup>PASI</sup>	291	PREFETCH
191	FMOV(s,d,q)	256	LDFSR <sup>D</sup>	291	PREFETCHA <sup>PASI</sup>

**TABLE 7-2** UltraSPARC Architecture 2005 Instruction Set - Alphabetical (2 of 2)

Page	Instruction				
299	RDASI	332	STDF	372	WRPR <sup>P</sup>
299	RDAsI <sup>PASR</sup>	335	STDFA <sup>PASI</sup>	369	WRSOFTINT_CLR <sup>P</sup>
299	RDCCR	332	STF	369	WRSOFTINT_SET <sup>P</sup>
299	RDFPRS	335	STFA <sup>PASI</sup>	369	WRSOFTINT <sup>P</sup>
299	RDGSR	339	STFSR <sup>D</sup>	369	WRSTICK_CMPR <sup>P</sup>
302	RDHPR <sup>H</sup>	323	STH	369	WRSTICK <sup>P</sup>
299	RDPC	324	STHA <sup>PASI</sup>	369	WRTICK_CMPR <sup>P</sup>
299	RDPCR <sup>P</sup>	341	STPARTIALF	369	WRY <sup>D</sup>
299	RDPIC <sup>PIC</sup>	332	STQF	377	XNOR (XNORcc)
303	RDPR <sup>P</sup>	335	STQFA <sup>PASI</sup>	377	XOR (XORcc)
299	RDSOFTINT <sup>P</sup>	344	STSHORTF		
299	RDSTICK_CMPR <sup>P</sup>	346	STTW <sup>D</sup>		
299	RDSTICK <sup>Pnpt</sup>	348	STTWA <sup>D, PASI</sup>		
299	RTICK_CMPR <sup>P</sup>	323	STW		
299	RTICK <sup>Pnpt</sup>	324	STWA <sup>PASI</sup>		
307	RESTORED <sup>P</sup>	323	STX		
305	RESTORE <sup>P</sup>	324	STXA <sup>PASI</sup>		
309	RETRY <sup>P</sup>	332	STXFSR		
311	RETURN	351	SUB (SUBcc)		
315	SAVED <sup>P</sup>	351	SUBC (SUBCcc)		
313	SAVE <sup>P</sup>	353	SWAPA <sup>D, PASI</sup>		
364	SDIV <sup>D</sup> (SDIVcc <sup>D</sup> )	352	SWAP <sup>D</sup>		
283	SDIVX	355	TADDcc		
317	SETHI	356	TADDccTV <sup>D</sup>		
318	SHUTDOWN <sup>D,P</sup>	358	Tcc		
319	SIAM	361	TSUBcc		
320	SIR <sup>H</sup>	362	TSUBccTV <sup>D</sup>		
321	SLL	364	UDIV <sup>D</sup> (UDIVcc <sup>D</sup> )		
321	SLLX	283	UDIVX		
367	SMUL <sup>D</sup> (SMULcc <sup>D</sup> )	367	UMUL <sup>D</sup> (UMULcc <sup>D</sup> )		
321	SRA	369	WRASI		
321	SRAX	369	WRAsI <sup>PASR</sup>		
321	SRL	369	WRCCR		
321	SRLX	369	WRFPRS		
323	STB	369	WRGSR		
324	STBA <sup>PASI</sup>	372	WRHPR <sup>H</sup>		
327	STBAR <sup>D</sup>	369	WRPCR <sup>P</sup>		
328	STBLOCKF	369	WRPIC <sup>PIC</sup>		

**TABLE 7-3** Instruction Set - by Functional Category (1 of 6)

Instruction	Category and Function	Page	Ext. to V9?
<b>Data Movement Operations, Between R Registers</b>			
MOVcc	Move integer register if condition is satisfied	275	
MOVr	Move integer register on contents of integer register	279	
<b>Data Movement Operations, Between F Registers</b>			
FMOV(s,d,q)	Floating-point move	191	
FMOV(s,d,q)cc	Move floating-point register if condition is satisfied	193	
FMOV(s,d,q)R	Move f-p reg. if integer reg. contents satisfy condition	198	
FSRC(1,2){s}	Copy source	225	VIS 1
<b>Data Conversion Instructions</b>			
FiTO(s,d,q)	Convert 32-bit integer to floating-point	185	
F(s,d,q)TOi	Convert floating point to integer	229	
F(s,d,q)TOx	Convert floating point to 64-bit integer	229	
F(s,d,q)TO(s,d,q)	Convert between floating-point formats	231	
FxTO(s,d,q)	Convert 64-bit integer to floating-point	234	
<b>Logical Operations on R Registers</b>			
AND (ANDcc)	Logical <b>and</b> (and modify condition codes)	149	
OR (ORcc)	Inclusive- <b>or</b> (and modify condition codes)	286	
ORN (ORNcc)	Inclusive- <b>or not</b> (and modify condition codes)	286	
XNOR (XNORcc)	Exclusive- <b>nor</b> (and modify condition codes)	377	
XOR (XORcc)	Exclusive- <b>or</b> (and modify condition codes)	377	
<b>Logical Operations on F Registers</b>			
FAND{s}	Logical <b>and</b> operation	227	VIS 1
FANDNOT(1,2){s}	Logical <b>and</b> operation with one inverted source	227	VIS 1
FNAND{s}	Logical <b>nand</b> operation	227	VIS 1
FNOR{s}	Logical <b>nor</b> operation	227	VIS 1
FNOT(1,2){s}	Copy negated source	225	VIS 1
FONE{s}	One fill	224	VIS 1
FOR{s}	Logical <b>or</b> operation	227	VIS 1
FORNOR(1,2){s}	Logical <b>or</b> operation with one inverted source	227	VIS 1
FXNOR{s}	Logical <b>xnor</b> operation	227	VIS 1
FXOR{s}	Logical <b>xor</b> operation	227	VIS 1
FZERO{s}	Zero fill	224	VIS 1
<b>Shift Operations on R Registers</b>			
SLL	Shift left logical	321	
SLLX	Shift left logical, extended	321	
SRA	Shift right arithmetic	321	
SRAX	Shift right arithmetic, extended	321	

**TABLE 7-3** Instruction Set - by Functional Category (2 of 6)

<b>Instruction</b>	<b>Category and Function</b>	<b>Page</b>	<b>Ext. to V9?</b>
SRL	Shift right logical	321	
SRLX	Shift right logical, extended	321	
<b><i>Special Addressing Operations</i></b>			
ALIGNADDRESS[_LITTLE]	Calculate address for misaligned data	147	<b>VIS 1</b>
ARRAY(8,16,32)	3-D array addressing instructions	150	<b>VIS 1</b>
FALIGNDATA	Perform data alignment for misaligned data	173	<b>VIS 1</b>
<b><i>Control Transfers</i></b>			
Bicc	Branch on integer condition codes	154	
BPcc	Branch on integer condition codes with prediction	157	
BPr	Branch on contents of integer register with prediction	160	
CALL	Call and link	162	
DONE <sup>P</sup>	Return from trap	166	
FBfcc <sup>D</sup>	Branch on floating-point condition codes	174	
FBPfcc	Branch on floating-point condition codes with prediction	176	
ILLTRAP	Illegal instruction	235	
JMPL	Jump and link	239	
RETRY <sup>P</sup>	Return from trap and retry	309	
RETURN	Return	311	
SIR <sup>H</sup>	Software-initiated reset	320	
Tcc	Trap on integer condition codes	358	
<b><i>Byte Permutation</i></b>			
BMASK	Set the GSR.mask field	156	<b>VIS 2</b>
BSHUFFLE	Permute bytes as specified by GSR.mask	156	<b>VIS 2</b>
<b><i>Data Formatting Operations on F Registers</i></b>			
FEXPAND	Pixel expansion	184	<b>VIS 1</b>
FPACK(16,32, FIX)	Pixel packing	210	<b>VIS 1</b>
FPMERGE	Pixel merge	219	<b>VIS 1</b>
<b><i>Memory Operations to/from F Registers</i></b>			
LDBLOCKF	Block loads	245	<b>VIS 1</b>
STBLOCKF	Block stores	328	<b>VIS 1</b>
LDDF	Load double floating-point	249	
LDDFA <sup>PASI</sup>	Load double floating-point from alternate space	252	
LDF	Load floating-point	249	
LDFA <sup>PASI</sup>	Load floating-point from alternate space	252	
LDQF	Load quad floating-point	249	
LDQFA <sup>PASI</sup>	Load quad floating-point from alternate space	252	
LDSHORTF	Short floating-point loads	258	<b>VIS 1</b>

**TABLE 7-3** Instruction Set - by Functional Category (3 of 6)

<b>Instruction</b>	<b>Category and Function</b>	<b>Page</b>	<b>Ext. to V9?</b>
STDF	Store double floating-point	332	
STDFA <sup>PASI</sup>	Store double floating-point into alternate space	335	
STF	Store floating-point	332	
STFA <sup>PASI</sup>	Store floating-point into alternate space	335	
STPARTIALF	Partial Store instructions	341	VIS 1
STQF	Store quad floating point	332	
STQFA <sup>PASI</sup>	Store quad floating-point into alternate space	335	
STSHORTF	Short floating-point stores	344	VIS 1
<b>Memory Operations — Miscellaneous</b>			
LDFSR <sup>D</sup>	Load floating-point state register lower	256	
LDFSR	Load floating-point state register	249	
MEMBAR	Memory barrier	271	
PREFETCH	Prefetch data	291	
PREFETCHA <sup>PASI</sup>	Prefetch data from alternate space	291	
STFSR <sup>D</sup>	Store floating-point state register	339	
STXFSR	Store extended floating-point state register	332	
<b>Atomic (Load-Store) Memory Operations to/from R Registers</b>			
CASA <sup>PASI</sup>	Compare and swap word in alternate space	163	
CASXA <sup>PASI</sup>	Compare and swap doubleword in alternate space	163	
LDSTUB	Load-store unsigned byte	260	
LDSTUBA <sup>PASI</sup>	Load-store unsigned byte in alternate space	261	
SWAP <sup>D</sup>	Swap integer register with memory	352	
SWAPA <sup>D, PASI</sup>	Swap integer register with memory in alternate space	353	
<b>Memory Operations to/from R Registers</b>			
LDSB	Load signed byte	240	
LDSBA <sup>PASI</sup>	Load signed byte from alternate space	242	
LDSH	Load signed halfword	240	
LDSHA <sup>PASI</sup>	Load signed halfword from alternate space	242	
LDSW	Load signed word	240	
LDSWA <sup>PASI</sup>	Load signed word from alternate space	242	
LDTXA <sup>N</sup>	Load integer twin extended word from alternate space	263	VIS 2+
LDTW <sup>D, PASI</sup>	Load integer twin word	266	
LDTWA <sup>D, PASI</sup>	Load integer twin word from alternate space	268	
LDUB	Load unsigned byte	260	
LDUBA <sup>PASI</sup>	Load unsigned byte from alternate space	242	
LDUH	Load unsigned halfword	240	
LDUHA <sup>PASI</sup>	Load unsigned halfword from alternate space	242	

**TABLE 7-3** Instruction Set - by Functional Category (4 of 6)

<b>Instruction</b>	<b>Category and Function</b>	<b>Page</b>	<b>Ext. to V9?</b>
LDUW	Load unsigned word	240	
LDUWA <sup>PASI</sup>	Load unsigned word from alternate space	242	
LDX	Load extended	240	
LDXA <sup>PASI</sup>	Load extended from alternate space	242	
STB	Store byte	323	
STBA <sup>PASI</sup>	Store byte into alternate space	324	
STBAR <sup>D</sup>	Store barrier	327	
STTW <sup>D</sup>	Store twin word	346	
STTWA <sup>D, PASI</sup>	Store twin word into alternate space	348	
STH	Store halfword	323	
STHA <sup>PASI</sup>	Store halfword into alternate space	324	
STW	Store word	323	
STWA <sup>PASI</sup>	Store word into alternate space	324	
STX	Store extended	323	
STXA <sup>PASI</sup>	Store extended into alternate space	324	
<b><i>Floating-Point Arithmetic Operations</i></b>			
FABS(s,d,q)	Floating-point absolute value	171	
FADD(s,d,q)	Floating-point add	172	
FDIV(s,d,q)	Floating-point divide	183	
FdMULq	Floating-point multiply double to quad	207	
FMUL(s,d,q)	Floating-point multiply	207	
FNEG(s,d,q)	Floating-point negate	209	
FsMULd	Floating-point multiply single to double	207	
FSQRT(s,d,q)	Floating-point square root	228	
FSUB(s,d,q)	Floating-point subtract	233	
<b><i>Floating-Point Comparison Operations</i></b>			
FCMP* <sup>&lt;16,32&gt;</sup>	Compare four 16-bit signed values or two 32-bit signed values	178	<b>VIS 1</b>
FCMP(s,d,q)	Floating-point compare	181	
FCMPE(s,d,q)	Floating-point compare (exception if unordered)	181	
<b><i>Register-Window Control Operations</i></b>			
ALLCLEAN	Mark all register window sets as "clean"	148	
INVALW	Mark all register window sets as "invalid"	238	
FLUSHW	Flush register windows	190	
NORMALW	"Other" register windows become "normal" register windows	285	
OTHERW	"Normal" register windows become "other" register windows	287	
RESTORE <sup>P</sup>	Restore caller's window	305	
RESTORED <sup>P</sup>	Window has been restored	307	

**TABLE 7-3** Instruction Set - by Functional Category (5 of 6)

<b>Instruction</b>	<b>Category and Function</b>	<b>Page</b>	<b>Ext. to V9?</b>
SAVE <sup>P</sup>	Save caller's window	313	
SAVED <sup>P</sup>	Window has been saved	315	
<i>Miscellaneous Operations</i>			
FLUSH	Flush instruction memory	186	
IMPDEP2A	Implementation-dependent instructions	236	
IMPDEP2B	Implementation-dependent instructions (reserved)	236	
NOP	No operation	284	
SHUTDOWN <sup>D,P</sup>	Shut down the virtual processor	318	<b>VIS 1</b>
<i>Integer SIMD Operations on F Registers</i>			
FPADD<16,32>[S]	Fixed-point partitioned add	216	<b>VIS 1</b>
FPSUB<16,32>[S]	Fixed-point partitioned subtract	221	<b>VIS 1</b>
<i>Integer Arithmetic Operations on R Registers</i>			
ADD (ADDcc)	Add (and modify condition codes)	146	
ADDC (ADDCcc)	Add with carry (and modify condition codes)	146	
MULSc <sup>D</sup>	Multiply step (and modify condition codes)	281	
MULX	Multiply 64-bit integers	283	
SDIV <sup>D</sup> (SDIVcc <sup>D</sup> )	32-bit signed integer divide (and modify condition codes)	364	
SDIVX	64-bit signed integer divide	283	
SMUL <sup>D</sup> (SMULcc <sup>D</sup> )	Signed integer multiply (and modify condition codes)	367	
SUB (SUBcc)	Subtract (and modify condition codes)	351	
SUBC (SUBCcc)	Subtract with carry (and modify condition codes)	351	
TADDcc	Tagged add and modify condition codes (trap on overflow)	355	
TADDccTV <sup>D</sup>	Tagged add and modify condition codes (trap on overflow)	356	
TSUBcc	Tagged subtract and modify condition codes (trap on overflow)	361	
TSUBccTV <sup>D</sup>	Tagged subtract and modify condition codes (trap on overflow)	362	
UDIV <sup>D</sup> (UDIVcc <sup>D</sup> )	Unsigned integer divide (and modify condition codes)	364	
UDIVX	64-bit unsigned integer divide	283	
UMUL <sup>D</sup> (UMULcc <sup>D</sup> )	Unsigned integer multiply (and modify condition codes)	367	
<i>Integer Arithmetic Operations on F Registers</i>			
FMUL8x16	8x16 partitioned product	201	<b>VIS 1</b>
FMUL8x16(AU,AL)	8x16 upper/lower $\alpha$ partitioned product	201	<b>VIS 1</b>
FMUL8(SU,UL)x16	8x16 upper/lower partitioned product	201	<b>VIS 1</b>
FMULD8(SU,UL)x16	8x16 upper/lower partitioned product	201	<b>VIS 1</b>
<i>Miscellaneous Operations on R Registers</i>			
POPC	Population count	289	
SETHI	Set high 22 bits of low word of integer register	317	
<i>Miscellaneous Operations on F Registers</i>			

**TABLE 7-3** Instruction Set - by Functional Category (6 of 6)

Instruction	Category and Function	Page	Ext. to V9?
EDGE(8,16,32){L}cc	Edge handling instructions (and modify condition codes)	168	VIS 1
EDGE(8,16,32){L}N	Edge handling instructions	170	VIS 2
PDIST	Pixel component distance	288	VIS 1
<b>Control and Status Register Access</b>			
RDASI	Read ASI register	299	
RDAsr <sup>PASR</sup>	Read ancillary state register	299	
RDCCR	Read Condition Codes register (CCR)	299	
RDFPRS	Read Floating-Point Registers State register (FPRS)	299	
RDGSR	Read General Status register (GSR)	299	
RDPC	Read Program Counter register (PC)	299	
RDP <sup>PCR</sup>	Read Performance Control register (PCR)	299	
RDPIC <sup>PIC</sup>	Read Performance Instrumentation Counters register (PIC)	299	
RDHPR <sup>H</sup>	Read hyperprivileged register	302	
RDP <sup>P</sup>	Read privileged register	303	
RDSOFTINT <sup>P</sup>	Read per-virtual processor Soft Interrupt register (SOFTINT)	299	
RDSTICK <sup>Pnpt</sup>	Read System Tick register (STICK)	299	
RDSTICK_CM <sup>P</sup> PR <sup>P</sup>	Read System Tick Compare register (STICK_CM <sup>P</sup> PR <sup>P</sup> )	299	
RDTICK <sup>Pnpt</sup>	Read Tick register (TICK)	299	
RDTICK_CM <sup>P</sup> PR <sup>P</sup>	Read Tick Compare register (TICK_CM <sup>P</sup> PR <sup>P</sup> )	299	
SIAM	Set interval arithmetic mode	319	VIS 2
WRASI	Write ASI register	369	
WRAsr <sup>PASR</sup>	Write ancillary state register	369	
WRCCR	Write Condition Codes register (CCR)	369	
WRFPRS	Write Floating-Point Registers State register (FPRS)	369	
WRGSR	Write General Status register (GSR)	369	
WRP <sup>PCR</sup>	Write Performance Control register (PCR)	369	
WRPIC <sup>PIC</sup>	Write Performance Instrumentation Counters register (PIC)	369	
WRHPR <sup>H</sup>	Write hyperprivileged register	372	
WRP <sup>P</sup>	Write privileged register	372	
WRSOFTINT <sup>P</sup>	Write per-virtual processor Soft Interrupt register (SOFTINT)	369	
WRSOFTINT_CLR <sup>P</sup>	Clear bits of per-virtual processor Soft Interrupt register (SOFTINT)	369	
WRSOFTINT_SET <sup>P</sup>	Set bits of per-virtual processor Soft Interrupt register (SOFTINT)	369	
WRTICK_CM <sup>P</sup> PR <sup>P</sup>	Write Tick Compare register (TICK_CM <sup>P</sup> PR <sup>P</sup> )	369	
WRSTICK <sup>P</sup>	Write System Tick register (STICK)	369	
WRSTICK_CM <sup>P</sup> PR <sup>P</sup>	Write System Tick Compare register (STICK_CM <sup>P</sup> PR <sup>P</sup> )	369	
WRY <sup>D</sup>	Write Y register	369	

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

**(1) Instruction Table.** This lists the instructions that are defined in the subsection, including the values of the field(s) that uniquely identify the instruction(s), assembly language syntax, and software and implementation classifications for the instructions. (*description of the Software Classes [letters] and Implementation Classes [digits] will be provided in a later update to this specification*)

**(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 UltraSPARC Architecture implementation encounters nonzero values in these fields, its behavior is as defined in *Reserved Opcodes and Instruction Fields* on page 132.

**Note** | Instruction classes are subject to change, and are not yet defined in this document. The classes will be defined in a later draft of this document and in the meantime are subject to change.

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

**(4) Exceptions.** The exception that can occur as a consequence of attempting to execute the instruction(s). Exceptions due to an *instruction\_access\_exception*, *fast\_instruction\_access\_MMU\_miss*, *WDR*, and interrupts are not listed because they can occur on any instruction. An FPop that is not implemented in hardware generates an *fp\_exception\_other* exception with *FSR.ftt* = *unimplemented\_FPop* when executed. A non-FPop 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 469), 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<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i>	<b>A1</b>
ADDcc	01 0000	Add and modify cc's	addcc <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i>	<b>A1</b>
ADDC	00 1000	Add with 32-bit Carry	addc <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i>	<b>A1</b>
ADDCcc	01 1000	Add with 32-bit Carry and modify cc's	addccc <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i>	<b>A1</b>



*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*

# ALIGNADDRESS

## 7.2 Align Address VIS 1

Instruction	opf	Operation	Assembly Language Syntax	Class
ALIGNADDRESS	0 0001 1000	Calculate address for misaligned data access	<code>alignaddr <i>reg<sub>rs1</sub></i>, <i>reg<sub>rs2</sub></i>, <i>reg<sub>rd</sub></i></code>	A1
ALIGNADDRESS_LITTLE	0 0001 1010	Calculate address for misaligned data access little-endian	<code>alignaddr1 <i>reg<sub>rs1</sub></i>, <i>reg<sub>rs2</sub></i>, <i>reg<sub>rd</sub></i></code>	A1



**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 FALIGNDATA 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 on page 173

# ALLCLEAN

## 7.3 Mark All Register Window Sets “Clean”

Instruction	Operation	Assembly Language Syntax	Class
ALLCLEAN <sup>P</sup>	Mark all register window sets as “clean”	allclean	C1



*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.

This instruction allows window manipulations to be atomic, without the value of *N\_REG\_WINDOWS* being visible to privileged software and without an assumption that *N\_REG\_WINDOWS* is constant (since hyperprivileged software can migrate a thread among virtual processors, across which *N\_REG\_WINDOWS* may vary).


*Exceptions* *illegal\_instruction* (not implemented in hardware in UltraSPARC Architecture 2005)  
*privileged\_opcode*

*See Also* INVALIDW on page 238  
NORMALW on page 285  
OTHERW on page 287  
RESTORED on page 307  
SAVED on page 315

# AND, ANDN

## 7.4 AND Logical Operation

Instruction	op3	Operation	Assembly Language Syntax	Class
AND	00 0001	<b>and</b>	and <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i>	<b>A1</b>
ANDcc	01 0001	<b>and</b> and modify cc's	andcc <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i>	<b>A1</b>
ANDN	00 0101	<b>and not</b>	andn <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i>	<b>A1</b>
ANDNcc	01 0101	<b>and not</b> and modify cc's	andncc <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i>	<b>A1</b>



*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].

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*

# ARRAY<8|16|32>

## 7.5 Three-Dimensional Array Addressing

VIS 1

Instruction	opf	Operation	Assembly Language Syntax	Class
ARRAY8	0 0001 0000	Convert 8-bit 3D address to blocked byte address	<code>array8 <i>reg<sub>rs1</sub></i>, <i>reg<sub>rs2</sub></i>, <i>reg<sub>rd</sub></i></code>	<b>C3</b>
ARRAY16	0 0001 0010	Convert 16-bit 3D address to blocked byte address	<code>array16 <i>reg<sub>rs1</sub></i>, <i>reg<sub>rs2</sub></i>, <i>reg<sub>rd</sub></i></code>	<b>C3</b>
ARRAY32	0 0001 0100	Convert 32-bit 3D address to blocked byte address	<code>array32 <i>reg<sub>rs1</sub></i>, <i>reg<sub>rs2</sub></i>, <i>reg<sub>rd</sub></i></code>	<b>C3</b>



*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 7-4. Illegal values produce undefined results in the destination register, R[rd].

**TABLE 7-4** 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.

# ARRAY<8|16|32>

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  $(0,0,0)$  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:

# ARRAY<8|16|32>

The (x,y,z) coordinates are input via a single 64-bit integer organized in R[rs1] as shown in FIGURE 7-2.

Z integer	Z fraction	Y integer	Y fraction	X integer	X fraction
63 55 54	44 43	33 32	22 21	11 10	0

**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  $512 \times 512 \times 32$  or a  $512 \times 512 \times 256$  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.


UPPER			MIDDLE			LOWER		
Z	Y	X	Z	Y	X	Z	Y	X
20 + 2n	17 + 2n	17 + n	17	9	5	4	2	0

**FIGURE 7-3** Three-Dimensional Array Blocked-Address Format (ARRAY8)

UPPER			MIDDLE			LOWER			0
Z	Y	X	Z	Y	X	Z	Y	X	
21 + 2n	18 + 2n	18 + n	18	10	6	5	3	1	0

**FIGURE 7-4** Three-Dimensional Array Blocked-Address Format (ARRAY16)

# ARRAY<8|16|32>



**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 7-5** 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

# Bicc Instructions

## 7.6 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 <sup>†</sup> { , a } label	A1
BE	0001	Branch on Equal	Z	be <sup>‡</sup> { , 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 <sup>◊</sup> { , a } label	A1
BCS	0101	Branch on Carry Set (Less Than, Unsigned)	C	bcs <sup>∇</sup> { , 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

<sup>†</sup> synonym: bnz

<sup>‡</sup> synonym: bz

<sup>◊</sup> synonym: bgeu

<sup>∇</sup> 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.

# Bicc Instructions

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.

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*.

*Exceptions*      None

# BMASK / BSHUFFLE

## 7.7 Byte Mask and Shuffle VIS 2

Instruction	opf	Operation	Assembly Language Syntax	Class
BMASK	0 0001 1001	Set the GSR.mask field in preparation for a subsequent BSHUFFLE instruction	bmask <i>reg<sub>rs1</sub>, reg<sub>rs2</sub>, reg<sub>rd</sub></i>	<b>C3</b>
BSHUFFLE	0 0100 1100	Permute 16 bytes as specified by GSR.mask	bshuffle <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i>	<b>C3</b>



*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*

## 7.8 Branch on Integer Condition Codes with Prediction (BPcc)

Instruction	cond	Operation	cc Test	Assembly Language Syntax	Class
BPA	1000	Branch Always	1	$ba\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>
BPN	0000	Branch Never	0	$bn\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>
BPNE	1001	Branch on Not Equal	<b>not</b> Z	$bne\ddagger\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>
BPE	0001	Branch on Equal	Z	$be\ddagger\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>
BPG	1010	Branch on Greater	<b>not</b> (Z <b>or</b> (N <b>xor</b> V))	$bg\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>
BPLE	0010	Branch on Less or Equal	Z <b>or</b> (N <b>xor</b> V)	$ble\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>
BPGE	1011	Branch on Greater or Equal	<b>not</b> (N <b>xor</b> V)	$bge\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>
BPL	0011	Branch on Less	N <b>xor</b> V	$bl\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>
BPGU	1100	Branch on Greater Unsigned	<b>not</b> (C <b>or</b> Z)	$bgu\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>
BPLEU	0100	Branch on Less or Equal Unsigned	C <b>or</b> Z	$bleu\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>
BPCC	1101	Branch on Carry Clear (Greater than or Equal, Unsigned)	<b>not</b> C	$bcc\hat{\diamond}\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>
BPCS	0101	Branch on Carry Set (Less than, Unsigned)	C	$bcs\forall\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>
BPPOS	1110	Branch on Positive	<b>not</b> N	$bpos\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>
BPNEG	0110	Branch on Negative	N	$bneg\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>
BPVC	1111	Branch on Overflow Clear	<b>not</b> V	$bvc\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>
BPVS	0111	Branch on Overflow Set	V	$bvs\{,a\}\{,ptl,pn\} \ i\_or\_x\_cc, \ label$	<b>A1</b>

† synonym:  $bnz$     ‡ synonym:  $bz$     ◊ synonym:  $bgeu$     ∇ synonym:  $blu$



cc1	cc0	Condition Code
0	0	icc
0	1	—
1	0	xcc
1	1	—

# BPcc

**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.

- **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. In no case does a Branch Never 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*.

An attempt to execute a BPcc instruction with `cc0 = 1` (a reserved value) causes an *illegal\_instruction* exception.

*Exceptions* *illegal\_instruction*


## **BPcc**

*See Also*      Branch on Integer Register with Prediction (BPr) on page 160

# BPr

## 7.9 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} reg<sub>rs1</sub>, label</code>	<b>A1</b>
BRLEZ	010	Branch on Register Less Than or Equal to Zero	$R[rs1] \leq 0$	<code>brlez {,a}{,pt ,pn} reg<sub>rs1</sub>, label</code>	<b>A1</b>
BRLZ	011	Branch on Register Less Than Zero	$R[rs1] < 0$	<code>brlz {,a}{,pt ,pn} reg<sub>rs1</sub>, label</code>	<b>A1</b>
—	100	<i>Reserved</i>	—	—	—
BRNZ	101	Branch on Register Not Zero	$R[rs1] \neq 0$	<code>brnz {,a}{,pt ,pn} reg<sub>rs1</sub>, label</code>	<b>A1</b>
BRGZ	110	Branch on Register Greater Than Zero	$R[rs1] > 0$	<code>brgz {,a}{,pt ,pn} reg<sub>rs1</sub>, label</code>	<b>A1</b>
BRGEZ	111	Branch on Register Greater Than or Equal to Zero	$R[rs1] \geq 0$	<code>brgez {,a}{,pt ,pn} reg<sub>rs1</sub>, label</code>	<b>A1</b>



\* Although SPARC V9 implementations should cause an *illegal\_instruction* exception when bit 28 = 1, many early implementations ignored the value of this bit and executed the opcode as a BPr instruction even if bit 28 = 1.

**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).

# BPr

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 ( $000_2$  or  $100_2$ ) causes an *illegal\_instruction* exception.

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*

*See Also*        Branch on Integer Condition Codes with Prediction (BPcc) on page 157

# CALL

## 7.10 Call and Link

Instruction	OP	Operation	Assembly Language Syntax	Class
CALL	01	Call and Link	<code>call label</code>	<b>A1</b>



*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)

*Exceptions* None

*See Also* JMPL on page 239

## 7.11 Compare and Swap

Instruction	op3	Operation	Assembly Language Syntax	Class
CASA <sup>P<sub>ASI</sub></sup>	11 1100	Compare and Swap Word from Alternate Space	casa [reg <sub>rs1</sub> ] imm_asi, reg <sub>rs2</sub> , reg <sub>rd</sub> casa [reg <sub>rs1</sub> ] %asi, reg <sub>rs2</sub> , reg <sub>rd</sub>	<b>A1</b>
CASXA <sup>P<sub>ASI</sub></sup>	11 1110	Compare and Swap Extended from Alternate Space	casxa [reg <sub>rs1</sub> ] imm_asi, reg <sub>rs2</sub> , reg <sub>rd</sub> casxa [reg <sub>rs1</sub> ] %asi, reg <sub>rs2</sub> , reg <sub>rd</sub>	<b>A1</b>



### Description

Concurrent processes use these 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.

# CASA / CASXA

A compare-and-swap operation does *not* imply any memory barrier semantics. When compare-and-swap is used for synchronization, the same consideration should be given to memory barriers as if a load, store, or swap instruction were used.

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.

An attempt to execute a CASXA or CASA instruction when  $i = 1$  and instruction bits 12:5 are nonzero causes an *illegal\_instruction* exception.

A *mem\_address\_not\_aligned* exception is generated if the address in R[rs1] is not properly aligned.

In nonprivileged mode (PSTATE.priv = 0 and HPSTATE.hpriv = 0), if bit 7 of the ASI is 0, CASXA and CASA cause a *privileged\_action* exception. In privileged mode (PSTATE.priv = 1 and HPSTATE.hpriv = 0), if the ASI is in the range  $30_{16}$  to  $7F_{16}$ , CASXA and CASA cause a *privileged\_action* exception.

<b>Compatibility Note</b>	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.
---------------------------	--

<b>Programming Note</b>	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 <i>Synthetic Instructions</i> 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 *data\_access\_exception*.

---

#### ASIs valid for CASA and CASXA instructions

ASI_NUCLEUS	ASI_NUCLEUS_LITTLE
ASI_AS_IF_USER_PRIMARY	ASI_AS_IF_USER_PRIMARY_LITTLE
ASI_AS_IF_USER_SECONDARY	ASI_AS_IF_USER_SECONDARY_LITTLE
ASI_REAL	ASI_REAL_LITTLE
ASI_PRIMARY	ASI_PRIMARY_LITTLE
ASI_SECONDARY	ASI_SECONDARY_LITTLE

# CASA / CASXA


## *Exceptions*

*illegal\_instruction*  
*mem\_address\_not\_aligned*  
*privileged\_action*  
*VA\_watchpoint*  
*data\_access\_exception*  
*fast\_data\_access\_MMU\_miss*  
*data\_access\_MMU\_miss*  
*data\_access\_MMU\_error*  
*fast\_data\_access\_protection*

# DONE

## 7.12 DONE

Instruction	op3	Operation	Assembly Language Syntax	Class
DONE <sup>P</sup>	11 1110	Return from Trap (skip trapped instruction)	done	C1



*Description* The DONE instruction restores the saved state from TSTATE[TL] (GL, CCR, ASI, PSTATE, and CWP), HTSTATE[TL] (HPSTATE), 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.

When a DONE instruction is executed in *privileged* mode and HTSTATE[TL].hpstate.hpriv = 0 (which will cause the DONE to return the virtual processor to nonprivileged or privileged mode), the value of GL restored from TSTATE[TL] saturates at MAXPGL. That is, if the value in TSTATE[TL].gl is greater than MAXPGL, then MAXPGL is substituted and written to GL. This protects against non-hyperprivileged software executing with  $GL > MAXPGL$ .

If software writes invalid or inconsistent state to TSTATE or HTSTATE before executing DONE, virtual processor behavior during and after execution of the DONE instruction is undefined.

The DONE instruction does not provide an error barrier, as MEMBAR #Sync does (impl. dep. #215-U3).

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.

# DONE

**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.

**Exceptions.** In privileged mode (PSTATE.priv = 1 and HPSTATE.hpriv = 0) or hyperprivileged mode (HPSTATE.hpriv = 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 and HPSTATE.hpriv = 0), an attempt to execute DONE causes a *privileged\_opcode* exception.

<b>Implementation</b>	In nonprivileged mode, <i>illegal_instruction</i> exception due to TL = 0
<b>Note</b>	does not occur. The <i>privileged_opcode</i> exception occurs instead, regardless of the current trap level (TL).

A *trap\_level\_zero* disrupting trap can occur upon the *completion* of a DONE instruction, if the following three conditions are true after DONE has executed:

- *trap\_level\_zero* exceptions are enabled (HPSTATE.tlz = 1),
- the virtual processor is in nonprivileged or privileged mode (HPSTATE.hpriv = 0), and
- the trap level (TL) register's value is zero (TL = 0)

*Exceptions*      *illegal\_instruction*  
                      *privileged\_opcode*

*trap\_level\_zero*

*See Also*        RETRY on page 309

# EDGE<8|16|32>{L}cc

## 7.13 Edge Handling Instructions VIS 1

Instruction	opf	Operation	Assembly Language Syntax †		Class
EDGE8cc	0 0000 0000	Eight 8-bit edge boundary processing	edge8cc	<i>reg<sub>rs1</sub></i> , <i>reg<sub>rs2</sub></i> , <i>reg<sub>rd</sub></i>	<b>C3</b>
EDGE8Lcc	0 0000 0010	Eight 8-bit edge boundary processing, little-endian	edge8lcc	<i>reg<sub>rs1</sub></i> , <i>reg<sub>rs2</sub></i> , <i>reg<sub>rd</sub></i>	<b>C3</b>
EDGE16cc	0 0000 0100	Four 16-bit edge boundary processing	edge16cc	<i>reg<sub>rs1</sub></i> , <i>reg<sub>rs2</sub></i> , <i>reg<sub>rd</sub></i>	<b>C3</b>
EDGE16Lcc	0 0000 0110	Four 16-bit edge boundary processing, little-endian	edge16lcc	<i>reg<sub>rs1</sub></i> , <i>reg<sub>rs2</sub></i> , <i>reg<sub>rd</sub></i>	<b>C3</b>
EDGE32cc	0 0000 1000	Two 32-bit edge boundary processing	edge32cc	<i>reg<sub>rs1</sub></i> , <i>reg<sub>rs2</sub></i> , <i>reg<sub>rd</sub></i>	<b>C3</b>
EDGE32Lcc	0 0000 1010	Two 32-bit edge boundary processing, little-endian	edge32lcc	<i>reg<sub>rs1</sub></i> , <i>reg<sub>rs2</sub></i> , <i>reg<sub>rd</sub></i>	<b>C3</b>

† 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* 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.

EDGE8Lcc, EDGE16Lcc, and EDGE32Lcc are little-endian versions of EDGE8cc, EDGE16cc, and EDGE32cc. 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 Partial Store instruction (see *Partial Store* on page 298) on 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] and the right edge mask is computed from the 3 least significant bits of R[rs2], according to TABLE 7-6.
2. If a 32-bit address masking is disabled (PSTATE.am = 0, 64-bit addressing) and the upper 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.

## EDGE<8|16|32>{L}cc

3. If 32-bit address masking is enabled (PSTATE.am = 1, 32-bit addressing) 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.

The integer condition codes are set per the rules of the SUBcc instruction with the same operands (see *Subtract* on page 303).

TABLE 7-6 lists edge mask specifications.

**TABLE 7-6** Edge Mask Specification

Edge Size	R[rsn] {2:0}	Big Endian		Little Endian	
		Left Edge	Right Edge	Left Edge	Right Edge
8	000	1111 1111	1000 0000	1111 1111	0000 0001
8	001	0111 1111	1100 0000	1111 1110	0000 0011
8	010	0011 1111	1110 0000	1111 1100	0000 0111
8	011	0001 1111	1111 0000	1111 1000	0000 1111
8	100	0000 1111	1111 1000	1111 0000	0001 1111
8	101	0000 0111	1111 1100	1110 0000	0011 1111
8	110	0000 0011	1111 1110	1100 0000	0111 1111
8	111	0000 0001	1111 1111	1000 0000	1111 1111
16	00x	1111	1000	1111	0001
16	01x	0111	1100	1110	0011
16	10x	0011	1110	1100	0111
16	11x	0001	1111	1000	1111
32	0xx	11	10	11	01
32	1xx	01	11	10	11

*Exceptions*     *illegal\_instruction*

*See Also*        EDGE(8,16,32){L}N on page 170

# EDGE<8|16|32>{L}N

## 7.14 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 $reg_{rs1}, reg_{rs2}, reg_{rd}$	<b>C3</b>
EDGE8LN	0 0000 0011	Eight 8-bit edge boundary processing, little-endian, no CC	edge8ln $reg_{rs1}, reg_{rs2}, reg_{rd}$	<b>C3</b>
EDGE16N	0 0000 0101	Four 16-bit edge boundary processing, no CC	edge16n $reg_{rs1}, reg_{rs2}, reg_{rd}$	<b>C3</b>
EDGE16LN	0 0000 0111	Four 16-bit edge boundary processing, little-endian, no CC	edge16ln $reg_{rs1}, reg_{rs2}, reg_{rd}$	<b>C3</b>
EDGE32N	0 0000 1001	Two 32-bit edge boundary processing, no CC	edge32n $reg_{rs1}, reg_{rs2}, reg_{rd}$	<b>C3</b>
EDGE32LN	0 0000 1011	Two 32-bit edge boundary processing, little-endian, no CC	edge32ln $reg_{rs1}, reg_{rs2}, reg_{rd}$	<b>C3</b>



*Description* EDGE8[L]N, EDGE16[L]N, and EDGE32[L]N operate identically to EDGE8[L]cc, EDGE16[L]cc, and EDGE32[L]cc, respectively, but do not set the integer condition codes.

See *Edge Handling Instructions* on page 168 for details.

*Exceptions* *illegal\_instruction*

*See Also* EDGE<8,16,32>[L]cc on page 168

# FABS

## 7.15 Floating-Point Absolute Value

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FABSs	11 0100	0 0000 1001	Absolute Value Single	<code>fabss</code> <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></i>	<b>A1</b>
FABSd	11 0100	0 0000 1010	Absolute Value Double	<code>fabsd</code> <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></i>	<b>A1</b>
FABSq	11 0100	0 0000 1011	Absolute Value Quad	<code>fabsq</code> <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></i>	<b>C3</b>



**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** UltraSPARC Architecture 2005 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.

**Exceptions** *illegal\_instruction*  
*fp\_disabled*  
*fp\_exception\_other* (`FSR.ftt = unimplemented_FPop` (FABSq))

# FADD

## 7.16 Floating-Point Add

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FADDs	11 0100	0 0100 0001	Add Single	fadds <i>freq<sub>rs1</sub></i> , <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></i>	A1
FADDd	11 0100	0 0100 0010	Add Double	faddd <i>freq<sub>rs1</sub></i> , <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></i>	A1
FADDq	11 0100	0 0100 0011	Add Quad	faddq <i>freq<sub>rs1</sub></i> , <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></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** UltraSPARC Architecture 2005 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.

If the FPU is enabled, FADDq causes an *fp\_exception\_other* (with FSR.ftt = unimplemented\_FPop), since that instruction is not implemented in hardware in UltraSPARC Architecture 2005 implementations.


**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 UltraSPARC Architecture 2005*.

**Exceptions** *illegal\_instruction*  
*fp\_disabled*  
*fp\_exception\_other* (FSR.ftt = unimplemented\_FPop (FADDq))  
*fp\_exception\_other* (FSR.ftt = unfinished\_FPop)  
*fp\_exception\_ieee\_754* (OF, UF, NX, NV)

# FALIGNDATA

## 7.17 Align Data VIS 1




Instruction	opf	Operation	Assembly Language Syntax	Class
FALIGNDATA	0 0100 1000	Perform data alignment for misaligned data	<code>faligndata freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>

*Description* FALIGNDATA concatenates the two 64-bit floating-point registers specified by `rs1` and `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 specified by `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`).

`GSR.align` is normally set by a previous `ALIGNADDRESS` instruction.

`GSR.align` 101



**FIGURE 7-6** FALIGNDATA

A byte-aligned 64-bit load can be performed as shown below.

<code>alignaddr</code>	<code>Address, Offset, Address !set GSR.align</code>
<code>ldd</code>	<code>[Address], %d0</code>
<code>ldd</code>	<code>[Address + 8], %d2</code>
<code>faligndata</code>	<code>%d0, %d2, %d4 !use GSR.align to select bytes</code>

If the FPU is not enabled (`FPRS.fef = 0` or `PSTATE.pef = 0`) or if no FPU is present, an attempt to execute an `FALIGNDATA` instruction causes an *fp\_disabled* exception.

*Exceptions* `fp_disabled`

*See Also* Align Address on page 147

## FBfcc (Deprecated)

### 7.18 Branch on Floating-Point Condition Codes (FBfcc)

The FBfcc instructions are deprecated and should not be used in new software. The FBPFcc instructions should be used instead.

Opcode	cond	Operation	fcc Test	Assembly Language Syntax	Class
FBA <sup>D</sup>	1000	Branch Always	1	fba{, a} label	A1
FBN <sup>D</sup>	0000	Branch Never	0	fbn{, a} label	A1
FBU <sup>D</sup>	0111	Branch on Unordered	U	fbu{, a} label	A1
FBG <sup>D</sup>	0110	Branch on Greater	G	fbg{, a} label	A1
FBUG <sup>D</sup>	0101	Branch on Unordered or Greater	G or U	fbug{, a} label	A1
FBL <sup>D</sup>	0100	Branch on Less	L	fbl{, a} label	A1
FBUL <sup>D</sup>	0011	Branch on Unordered or Less	L or U	fbul{, a} label	A1
FBLG <sup>D</sup>	0010	Branch on Less or Greater	L or G	fblg{, a} label	A1
FBNE <sup>D</sup>	0001	Branch on Not Equal	L or G or U	fbne <sup>†</sup> {, a} label	A1
FBE <sup>D</sup>	1001	Branch on Equal	E	fbe <sup>‡</sup> {, a} label	A1
FBUE <sup>D</sup>	1010	Branch on Unordered or Equal	E or U	fbue{, a} label	A1
FBGE <sup>D</sup>	1011	Branch on Greater or Equal	E or G	fbge{, a} label	A1
FBUGE <sup>D</sup>	1100	Branch on Unordered or Greater or Equal	E or G or U	fbuge{, a} label	A1
FBLE <sup>D</sup>	1101	Branch on Less or Equal	E or L	fble{, a} label	A1
FBULE <sup>D</sup>	1110	Branch on Unordered or Less or Equal	E or L or U	fbule{, a} label	A1
FBO <sup>D</sup>	1111	Branch on Ordered	E or L or G	fbo{, a} label	A1

<sup>†</sup> synonym: fbnz

<sup>‡</sup> synonym: fbz



**Programming Note** To set the annul (a) bit for FBfcc instructions, append “, a” to the opcode mnemonic. For example, use “fbl, a label”. In the preceding table, braces around “, a” signify that “, a” is optional.

*Description* Unconditional and Fcc branches are described below:

## FBfcc (Deprecated)

- **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**(disp22))” 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**(disp22))”. 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.

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.

*Exceptions*     *fp\_disabled*

# FBPfcc

## 7.19 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}{,pt ,pn} %fccn, label</code>	<b>A1</b>
FBPN	0000	Branch Never	0	<code>fbn{,a}{,pt ,pn} %fccn, label</code>	<b>A1</b>
FBPU	0111	Branch on Unordered	U	<code>fbu{,a}{,pt ,pn} %fccn, label</code>	<b>A1</b>
FBPG	0110	Branch on Greater	G	<code>fbg{,a}{,pt ,pn} %fccn, label</code>	<b>A1</b>
FBPUG	0101	Branch on Unordered or Greater	G or U	<code>fbug{,a}{,pt ,pn} %fccn, label</code>	<b>A1</b>
FBPL	0100	Branch on Less	L	<code>fb{,a}{,pt ,pn} %fccn, label</code>	<b>A1</b>
FBPUL	0011	Branch on Unordered or Less	L or U	<code>fbul{,a}{,pt ,pn} %fccn, label</code>	<b>A1</b>
FBPLG	0010	Branch on Less or Greater	L or G	<code>fb{,a}{,pt ,pn} %fccn, label</code>	<b>A1</b>
FBPNE	0001	Branch on Not Equal	L or G or U	<code>fbne<sup>†</sup>{,a}{,pt ,pn} %fccn, label</code>	<b>A1</b>
FBPE	1001	Branch on Equal	E	<code>fbe<sup>‡</sup>{,a}{,pt ,pn} %fccn, label</code>	<b>A1</b>
FBPUE	1010	Branch on Unordered or Equal	E or U	<code>fbue{,a}{,pt ,pn} %fccn, label</code>	<b>A1</b>
FBPGE	1011	Branch on Greater or Equal	E or G	<code>fbge{,a}{,pt ,pn} %fccn, label</code>	<b>A1</b>
FBPUGE	1100	Branch on Unordered or Greater or Equal	E or G or U	<code>fbuge{,a}{,pt ,pn} %fccn, label</code>	<b>A1</b>
FBPLE	1101	Branch on Less or Equal	E or L	<code>fb{,a}{,pt ,pn} %fccn, label</code>	<b>A1</b>
FBPULE	1110	Branch on Unordered or Less or Equal	E or L or U	<code>fbule{,a}{,pt ,pn} %fccn, label</code>	<b>A1</b>
FBPO	1111	Branch on Ordered	E or L or G	<code>fb{o}{,a}{,pt ,pn} %fccn, label</code>	<b>A1</b>

<sup>†</sup> synonym: `fbnz`      <sup>‡</sup> 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>

# FBPfcc

**Programming Note** | To set the annul (a) bit for FBPfcc instructions, append “,a” to the opcode mnemonic. For example, use “fbl,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.

- **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.

## Exceptions

*fp\_disabled*

# FCMP\* <16|32> (SIMD)

## 7.20 SIMD Signed Compare VIS 1

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class
FCMPLE16	0 0010 0000	Four 16-bit compare; set R[rd] if $src1 \leq src2$	f64	f64	i64	<code>fcmp1e16 <math>freq_{rs1}, freq_{rs2}, reg_{rd}</math></code>	<b>C3</b>
FCMPNE16	0 0010 0010	Four 16-bit compare; set R[rd] if $src1 \neq src2$	f64	f64	i64	<code>fcmpne16 <math>freq_{rs1}, freq_{rs2}, reg_{rd}</math></code>	<b>C3</b>
FCMPLE32	0 0010 0100	Two 32-bit compare; set R[rd] if $src1 \leq src2$	f64	f64	i64	<code>fcmp1e32 <math>freq_{rs1}, freq_{rs2}, reg_{rd}</math></code>	<b>C3</b>
FCMPNE32	0 0010 0110	Two 32-bit compare; set R[rd] if $src1 \neq src2$	f64	f64	i64	<code>fcmpne32 <math>freq_{rs1}, freq_{rs2}, reg_{rd}</math></code>	<b>C3</b>
FCMPGT16	0 0010 1000	Four 16-bit compare; set R[rd] if $src1 > src2$	f64	f64	i64	<code>fcmpgt16 <math>freq_{rs1}, freq_{rs2}, reg_{rd}</math></code>	<b>C3</b>
FCMPEQ16	0 0010 1010	Four 16-bit compare; set R[rd] if $src1 = src2$	f64	f64	i64	<code>fcmp1eq16 <math>freq_{rs1}, freq_{rs2}, reg_{rd}</math></code>	<b>C3</b>
FCMPGT32	0 0010 1100	Two 32-bit compare; set R[rd] if $src1 > src2$	f64	f64	i64	<code>fcmpgt32 <math>freq_{rs1}, freq_{rs2}, reg_{rd}</math></code>	<b>C3</b>
FCMPEQ32	0 0010 1110	Two 32-bit compare; set R[rd] if $src1 = src2$	f64	f64	i64	<code>fcmp1eq32 <math>freq_{rs1}, freq_{rs2}, reg_{rd}</math></code>	<b>C3</b>



**Description** Either four 16-bit signed values or two 32-bit signed values in  $F_D[rs1]$  and  $F_D[rs2]$  are compared. The 4-bit or 2-bit condition-code results are stored in the least significant bits of the integer register R[rd]. The least significant 16-bit or 32-bit compare result corresponds to bit zero of R[rd].

**Note** Bits 63:4 of the destination register R[rd] are set to zero for 16-bit compares. Bits 63:2 of the destination register R[rd] are set to zero for 32-bit compares.

For FCMPGT{16,32}, each bit in the result is set to 1 if the corresponding signed value in  $F_D[rs1]$  is greater than the signed value in  $F_D[rs2]$ . Less-than comparisons are made by swapping the operands.

For FCMPLE{16,32}, each bit in the result is set to 1 if the corresponding signed value in  $F_D[rs1]$  is less than or equal to the signed value in  $F_D[rs2]$ . Greater-than-or-equal comparisons are made by swapping the operands.

For FCMPEQ{16,32}, each bit in the result is set to 1 if the corresponding signed value in  $F_D[rs1]$  is equal to the signed value in  $F_D[rs2]$ .

# FCMP\* <16|32> (SIMD)

For FCMPNE{16,32}, each bit in the result is set to 1 if the corresponding signed value in  $F_D[rs1]$  is not equal to the signed value in  $F_D[rs2]$ .

FIGURE 7-7 and FIGURE 7-8 illustrate 16-bit and 32-bit pixel comparison operations, respectively.




FIGURE 7-7 Four 16-bit Signed Fixed-point SIMD Comparison Operations




FIGURE 7-8 Two 32-bit Signed Fixed-point SIMD 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 SIMD signed compare operation can be used directly by both integer operations (for example, partial stores) and partitioned conditional moves.

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.

## FCMP\* <16|32> (SIMD)

*Exception*      *fp\_disabled*

*See Also*      STPARTIALF on page 341

# FCMP<s|d|q> / FCMPE<s|d|q>

## 7.21 Floating-Point Compare

Instruction	opf	Operation	Assembly Language Syntax	Class
FCMPs	0 0101 0001	Compare Single	<code>fcmps %fccn, freg<sub>rs1</sub>, freg<sub>rs2</sub></code>	<b>A1</b>
FCMPd	0 0101 0010	Compare Double	<code>fcmpd %fccn, freg<sub>rs1</sub>, freg<sub>rs2</sub></code>	<b>A1</b>
FCMPq	0 0101 0011	Compare Quad	<code>fcmpq %fccn, freg<sub>rs1</sub>, freg<sub>rs2</sub></code>	<b>C3</b>
FCMPEs	0 0101 0101	Compare Single and Exception if Unordered	<code>fcmpes %fccn, freg<sub>rs1</sub>, freg<sub>rs2</sub></code>	<b>A1</b>
FCMPEd	0 0101 0110	Compare Double and Exception if Unordered	<code>fcmped %fccn, freg<sub>rs1</sub>, freg<sub>rs2</sub></code>	<b>A1</b>
FCMPEq	0 0101 0111	Compare Quad and Exception if Unordered	<code>fcmpeq %fccn, freg<sub>rs1</sub>, freg<sub>rs2</sub></code>	<b>C3</b>



cc1	cc0	Condition Code
0	0	<code>fcc0</code>
0	1	<code>fcc1</code>
1	0	<code>fcc2</code>
1	1	<code>fcc3</code>

*Description* 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 (`fccn`) as shown below.

fcc value	Relation
0	$freg_{rs1} = freg_{rs2}$
1	$freg_{rs1} < freg_{rs2}$
2	$freg_{rs1} > freg_{rs2}$
3	$freg_{rs1} ? freg_{rs2}$ (unordered)

The “?” in the preceding table means that the comparison is unordered. The unordered condition occurs when one or both of the operands to the compare is a signalling or quiet NaN.

## FCMP<s|d|q> / FCMPE<s|d|q>

The “compare and cause exception if unordered” (FCMPEs, FCMPEd, and FCMPEq) instructions cause an invalid (NV) exception if either operand is a NaN.

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 UltraSPARC 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 UltraSPARC 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 UltraSPARC 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** UltraSPARC Architecture 2005 processors do not implement in hardware the instructions that refer to quad-precision floating-point registers. An attempt to execute FCMPq or FCMPEq generates *fp\_exception\_other* (with `FSR.ftt = unimplemented_FPop`), 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.

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for UltraSPARC Architecture 2005*.

*Exceptions*

- illegal\_instruction*
- fp\_disabled*
- fp\_exception\_ieee\_754* (NV)
- fp\_exception\_other* (`FSR.ftt = unimplemented_FPop` (FCMPq, FCMPEq only))

# FDIV<s|d|q>

## 7.22 Floating-Point Divide

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FDIVs	11 0100	0 0100 1101	Divide Single	<code>fdivs <i>freq<sub>rs1</sub></i>, <i>freq<sub>rs2</sub></i>, <i>freq<sub>rd</sub></i></code>	<b>A1</b>
FDIVd	11 0100	0 0100 1110	Divide Double	<code>fdivd <i>freq<sub>rs1</sub></i>, <i>freq<sub>rs2</sub></i>, <i>freq<sub>rd</sub></i></code>	<b>A1</b>
FDIVq	11 0100	0 0100 1111	Divide Quad	<code>fdivq <i>freq<sub>rs1</sub></i>, <i>freq<sub>rs2</sub></i>, <i>freq<sub>rd</sub></i></code>	<b>C3</b>



*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`.

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.

If the FPU is enabled, `FDIVq` causes an *fp\_exception\_other* (with `FSR.ftt = unimplemented_FPop`), since that instruction is not implemented in hardware in UltraSPARC Architecture 2005 implementations.

**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 UltraSPARC Architecture 2005*.

*Exceptions*

- illegal\_instruction*
- fp\_disabled*
- fp\_exception\_other* (`FSR.ftt = unimplemented_FPop` (`FDIVq` only))
- fp\_exception\_other* (`FSR.ftt = unfinished_FPop` (`FDIVs`, `FDIVd`))
- fp\_exception\_ieee\_754* (`OF`, `UF`, `DZ`, `NV`, `NX`)


# FEXPAND

## 7.23 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 freg_rs2, freg_rd</code>	<b>C3</b>



*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 each result to a 16-bit fixed value.
2. Store the result in the destination register,  $F_D[rd]$ .

**Programming Note** FEXPAND performs the inverse of the FPACK16 operation.

In an UltraSPARC Architecture 2005 implementation, this instruction is not implemented in hardware, causes an *illegal\_instruction* exception, and is emulated in software.

*Exceptions* *illegal\_instruction*

*See Also* FPMERGE on page 219  
FPACK on page 210

# FiTO<s|d|q>

## 7.24 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 <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i>	A1
FiTOd	11 0100	0 1100 1000	Convert 32-bit Integer to Double	—	f32	f64	fitod <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i>	A1
FiTOq	11 0100	0 1100 1100	Convert 32-bit Integer to Quad	—	f32	f128	fitoq <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i>	C3



**Description** FiTOs, FiTOd, and FiTOq convert the 32-bit signed integer operand in floating-point register F<sub>S</sub>[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** UltraSPARC Architecture 2005 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.

If the FPU is enabled, FiTOq causes an *fp\_exception\_other* (with FSR.ftt = unimplemented\_FPop), since that instruction is not implemented in hardware in UltraSPARC Architecture 2005 implementations.

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for UltraSPARC Architecture 2005*.


**Exceptions** *illegal\_instruction*  
*fp\_disabled*  
*fp\_exception\_other* (FSR.ftt = unimplemented\_FPop (FiTOq))  
*fp\_exception\_ieee\_754* (NX (FiTOs only))

# FLUSH

## 7.25 Flush Instruction Memory

Instruction	op3	Operation	Assembly Language Syntax†	Class
FLUSH	11 1011	Flush Instruction Memory	flush [address]	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 writes (stores) to a memory location containing an instruction (self-modifying code<sup>1</sup>), a potential memory consistency problem arises, which is addressed by the FLUSH instruction. Use of FLUSH ensures that instruction and data memory are synchronized after instruction memory has been modified.

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  $P_{\text{FLUSH}}$  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  $P_{\text{FLUSH}}$  appear to execute after any loads, stores, and atomic load-stores to that address issued by  $P_{\text{FLUSH}}$  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 was a store operation (see *Memory Barrier* on page 271).

If  $i = 0$ , the effective address operand for the FLUSH instruction is “R[rs1] + R[rs2]”; if  $i = 1$ , it is “R[rs1] + sign\_ext (simm13)”. The three least-significant three bits of the effective address are ignored; that is, the effective address always refers to an aligned doubleword.

<sup>1</sup> 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 UltraSPARC 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-Cs20:** 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.

<b>Programming Note</b>	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

<b>Programming Notes</b>	<ol style="list-style-type: none"><li>1. Typically, FLUSH is used in self-modifying code. The use of self-modifying code is discouraged.</li><li>2. If a program includes self-modifying code, to be portable it <i>must</i> 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.</li></ol>
--------------------------	--

# FLUSH

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.
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 or hyperprivileged 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.

## Exceptions

*illegal\_instruction*

*fast\_data\_access\_MMU\_miss* (TLB miss with hardware tablewalk disabled)  
(impl. dep. #409-S10-Cs20)

# FLUSH

*data\_access\_MMU\_miss* (TLB miss with hardware tablewalk enabled)  
(impl. dep. #409-S10-Cs20)

*fast\_data\_access\_protection*

# FLUSHW

## 7.26 Flush Register Windows

Instruction	op3	Operation	Assembly Language Syntax	Class
FLUSHW	10 1011	Flush Register Windows	<code>flushw</code>	<b>A1</b>



*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 129.

**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, 18:14, or 12:0 are nonzero causes an *illegal\_instruction* exception.

*Exceptions* *illegal\_instruction*  
*spill\_n\_normal*  
*spill\_n\_other*

# FMOV

## 7.27 Floating-Point Move

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FMOV <sub>s</sub>	11 0100	0 0000 0001	Move (copy) Single	<code>fmovs</code> <i>reg<sub>rs2</sub>, reg<sub>rd</sub></i>	<b>A1</b>
FMOV <sub>d</sub>	11 0100	0 0000 0010	Move (copy) Double	<code>fmovd</code> <i>reg<sub>rs2</sub>, reg<sub>rd</sub></i>	<b>A1</b>
FMOV <sub>q</sub>	11 0100	0 0000 0011	Move (copy) Quad	<code>fmovq</code> <i>reg<sub>rs2</sub>, reg<sub>rd</sub></i>	<b>C3</b>



**Description** FMOV copies the source floating-point register(s) to the destination floating-point register(s), unaltered.

FMOV<sub>s</sub>, FMOV<sub>d</sub>, and FMOV<sub>q</sub> 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.

**Note** UltraSPARC Architecture 2005 processors do not implement in hardware instructions that refer to quad-precision floating-point registers. An attempt to execute an FMOV<sub>q</sub> 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.

If the FPU is enabled, an attempt to execute an FMOV<sub>q</sub> instruction causes an *fp\_exception\_other* (with FSR.ftt = unimplemented\_FPop), since that instruction is not implemented in hardware in UltraSPARC Architecture 2005 implementations.

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for UltraSPARC Architecture 2005*.

**Exceptions** *illegal\_instruction*  
*fp\_disabled*  
*fp\_exception\_other* (FSR.ftt = unimplemented\_FPop (FMOV<sub>q</sub> only))


# FMOV

*See Also*      *F Register Logical Operate (2 operand)* on page 225

# FMOVcc

## 7.28 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	<code>fmovsicc %icc, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FMOVDicc	00 0010	Move Floating-Point Double, based on 32-bit integer condition codes	<code>fmovdicc %icc, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FMOVQicc	00 0011	Move Floating-Point Quad, based on 32-bit integer condition codes	<code>fmovqicc %icc, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>C3</b>
FMOVsxcc	00 0001	Move Floating-Point Single, based on 64-bit integer condition codes	<code>fmovsxcc %xcc, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FMOVDxcc	00 0010	Move Floating-Point Double, based on 64-bit integer condition codes	<code>fmovdxcc %xcc, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FMOVQxcc	00 0011	Move Floating-Point Quad, based on 64-bit integer condition codes	<code>fmovqxcc %xcc, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>C3</b>
FMOVsfcc	00 0001	Move Floating-Point Single, based on floating-point condition codes	<code>fmovsfcc %fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FMOVdfcc	00 0010	Move Floating-Point Double, based on floating-point condition codes	<code>fmovdfcc %fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FMOVQfcc	00 0011	Move Floating-Point Quad, based on floating-point condition codes	<code>fmovqfcc %fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>C3</b>



# FMOVcc

Encoding of the **cond** Field for F.P. Moves Based on Integer Condition Codes (**icc** or **xcc**)

<b>cond</b>	<b>Operation</b>	<b>icc / xcc Test</b>	<i>icc/xcc</i> name(s) in <b>Assembly Language Mnemonics</b>
1000	Move Always	1	a
0000	Move Never	0	n
1001	Move if Not Equal	<b>not Z</b>	ne (or nz)
0001	Move if Equal	Z	e (or z)
1010	Move if Greater	<b>not (Z or (N xor V))</b>	g
0010	Move if Less or Equal	Z or (N xor V)	le
1011	Move if Greater or Equal	<b>not (N xor V)</b>	ge
0011	Move if Less	N xor V	l
1100	Move if Greater Unsigned	<b>not (C or Z)</b>	gu
0100	Move if Less or Equal Unsigned	(C or Z)	leu
1101	Move if Carry Clear (Greater or Equal, Unsigned)	<b>not C</b>	cc (or geu)
0101	Move if Carry Set (Less than, Unsigned)	C	cs (or lu)
1110	Move if Positive	<b>not N</b>	pos
0110	Move if Negative	N	neg
1111	Move if Overflow Clear	<b>not V</b>	vc
0111	Move if Overflow Set	V	vs

# FMOVcc

Encoding of the **cond** Field for F.P. Moves Based on Floating-Point Condition Codes (**fccn**)

<b>cond</b>	<b>Operation</b>	<b>fcc<sub>n</sub> Test</b>	<b>fcc name(s) in Assembly Language Mnemonics</b>
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

Encoding of **opf\_cc** Field (also see TABLE E-10 on page 484)

<b>opf_cc</b>	<b>Instruction</b>	<b>Condition Code to be Tested</b>
100 <sub>2</sub>	FMOV(S,D,Q)icc	icc
110 <sub>2</sub>	FMOV(S,D,Q)xcc	xcc
000 <sub>2</sub>	FMOV(S,D,Q)fcc	fcc0
001 <sub>2</sub>		fcc1
010 <sub>2</sub>		fcc2
011 <sub>2</sub>		fcc3
101 <sub>2</sub>	<i>(illegal_instruction</i> exception)	
111 <sub>2</sub>		

# FMOVcc

*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** | UltraSPARC Architecture 2005 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* = 101<sub>2</sub> or 111<sub>2</sub> 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.

If the FPU is enabled, an attempt to execute an FMOVQicc, FMOVQxcc, or FMOVQfcc instruction causes an *fp\_exception\_other* (with FSR.ftt = unimplemented\_FPop), since that instruction is not implemented in hardware in UltraSPARC Architecture 2005 implementations.

# 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  
    fmovd1e %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 = unimplemented\_FPop (opf\_cc = 101<sub>2</sub> or 111<sub>2</sub>))  
                  *fp\_exception\_other* (FSR.ftt = unimplemented\_FPop (FMOVQ instructions only))

# FMOVR

## 7.29 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	<b>A1</b>
FMOVRsLEZ	010	0 0101	Move Single if Register ≤ 0	R[rs1] ≤ 0	<b>A1</b>
FMOVRsLZ	011	0 0101	Move Single if Register < 0	R[rs1] < 0	<b>A1</b>
—	100	0 0101	<i>Reserved</i>	—	—
FMOVRsNZ	101	0 0101	Move Single if Register ≠ 0	R[rs1] ≠ 0	<b>A1</b>
FMOVRsGZ	110	0 0101	Move Single if Register > 0	R[rs1] > 0	<b>A1</b>
FMOVRsGEZ	111	0 0101	Move Single if Register ≥ 0	R[rs1] ≥ 0	<b>A1</b>
—	000	0 0110	<i>Reserved</i>	—	—
FMOVRdZ	001	0 0110	Move Double if Register = 0	R[rs1] = 0	<b>A1</b>
FMOVRdLEZ	010	0 0110	Move Double if Register ≤ 0	R[rs1] ≤ 0	<b>A1</b>
FMOVRdLZ	011	0 0110	Move Double if Register < 0	R[rs1] < 0	<b>A1</b>
—	100	0 0110	<i>Reserved</i>	—	—
FMOVRdNZ	101	0 0110	Move Double if Register ≠ 0	R[rs1] ≠ 0	<b>A1</b>
FMOVRdGZ	110	0 0110	Move Double if Register > 0	R[rs1] > 0	<b>A1</b>
FMOVRdGEZ	111	0 0110	Move Double if Register ≥ 0	R[rs1] ≥ 0	<b>A1</b>
—	000	0 0111	<i>Reserved</i>	—	—
FMOVRqZ	001	0 0111	Move Quad if Register = 0	R[rs1] = 0	<b>C3</b>
FMOVRqLEZ	010	0 0111	Move Quad if Register ≤ 0	R[rs1] ≤ 0	<b>C3</b>
FMOVRqLZ	011	0 0111	Move Quad if Register < 0	R[rs1] < 0	<b>C3</b>
—	100	0 0111	<i>Reserved</i>	—	—
FMOVRqNZ	101	0 0111	Move Quad if Register ≠ 0	R[rs1] ≠ 0	<b>C3</b>
FMOVRqGZ	110	0 0111	Move Quad if Register > 0	R[rs1] > 0	<b>C3</b>
FMOVRqGEZ	111	0 0111	Move Quad if Register ≥ 0	R[rs1] ≥ 0	<b>C3</b>



# FMOVR

---

**Assembly Language Syntax**

---

fmovr{s,d,q}z *reg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub>* (synonym: fmovr{s,d,q}e)  
fmovr{s,d,q}lez *reg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub>*  
fmovr{s,d,q}lz *reg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub>*  
fmovr{s,d,q}nz *reg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub>* (synonym: fmovr{s,d,q}ne)  
fmovr{s,d,q}gz *reg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub>*  
fmovr{s,d,q}gez *reg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub>*

---

## 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** | UltraSPARC Architecture 2005 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 = 000<sub>2</sub> or 100<sub>2</sub> 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.

If the FPU is enabled, an attempt to execute an FMOVRq instruction causes an *fp\_exception\_other* (with FSR.ftt = unimplemented\_FPop), since that instruction is not implemented in hardware in UltraSPARC Architecture 2005 implementations.

# FMOVR

**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 :

<u>Branch</u>	<u>Test</u>
FMOVRNZ	<b>not</b> Z
FMOVRZ	Z
FMOVRGEZ	<b>not</b> N
FMOVRLZ	N
FMOVRLEZ	N <b>or</b> Z
FMOVRGZ	N <b>nor</b> Z

*Exceptions*    *fp\_disabled*  
*fp\_exception\_other* (FSR.ftt = unimplemented\_FPop (rcond = 000<sub>2</sub> or 100<sub>2</sub>))  
*fp\_exception\_other* (FSR.ftt = unimplemented\_FPop (FMOVRq))

# FMUL (partitioned)

## 7.30 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 <i>reg<sub>rs1</sub></i>, <i>reg<sub>rs2</sub></i>, <i>reg<sub>rd</sub></i></code>	<b>C3</b>
FMUL8x16AU	0 0011 0011	Unsigned 8-bit by signed 16-bit upper $\alpha$ partitioned product	f32	f32	f64	<code>fmul8x16au <i>reg<sub>rs1</sub></i>, <i>reg<sub>rs2</sub></i>, <i>reg<sub>rd</sub></i></code>	<b>C3</b>
FMUL8x16AL	0 0011 0101	Unsigned 8-bit by signed 16-bit lower $\alpha$ partitioned product	f32	f32	f64	<code>fmul8x16al <i>reg<sub>rs1</sub></i>, <i>reg<sub>rs2</sub></i>, <i>reg<sub>rd</sub></i></code>	<b>C3</b>
FMUL8SUx16	0 0011 0110	Signed upper 8-bit by signed 16-bit partitioned product	f32	f64	f64	<code>fmul8sux16 <i>reg<sub>rs1</sub></i>, <i>reg<sub>rs2</sub></i>, <i>reg<sub>rd</sub></i></code>	<b>C3</b>
FMUL8ULx16	0 0011 0111	Unsigned lower 8-bit by signed 16-bit partitioned product	f32	f64	f64	<code>fmul8ulx16 <i>reg<sub>rs1</sub></i>, <i>reg<sub>rs2</sub></i>, <i>reg<sub>rd</sub></i></code>	<b>C3</b>
FMULD8SUx16	0 0011 1000	Signed upper 8-bit by signed 16-bit partitioned product	f32	f32	f64	<code>fmuld8sux16 <i>reg<sub>rs1</sub></i>, <i>reg<sub>rs2</sub></i>, <i>reg<sub>rd</sub></i></code>	<b>C3</b>
FMULD8ULx16	0 0011 1001	Unsigned lower 8-bit by signed 16-bit partitioned product	f32	f32	f64	<code>fmuld8ulx16 <i>reg<sub>rs1</sub></i>, <i>reg<sub>rs2</sub></i>, <i>reg<sub>rd</sub></i></code>	<b>C3</b>



**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.

In an UltraSPARC Architecture 2005 implementation, these instructions are not implemented in hardware, cause an *illegal\_instruction* exception, and are emulated in software.

*Exceptions* *illegal\_instruction*

# FMUL (partitioned)

## 7.30.1 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 7-10 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.




FIGURE 7-10 FMUL8x16 Operation

# FMUL (partitioned)

## 7.30.2 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  $\alpha$  pixel component value). FIGURE 7-11 illustrates the operation.




FIGURE 7-11 FMUL8x16AU Operation

## 7.30.3 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 7-12 illustrates the operation.




FIGURE 7-12 FMUL8x16AL Operation

# FMUL (partitioned)

## 7.30.4 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 7-13 illustrates the operation.




FIGURE 7-13 FMUL8SUx16 Operation

## 7.30.5 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 7-14 illustrates the operation; CODE EXAMPLE 7-1 exemplifies the operation.

## FMUL (partitioned)




FIGURE 7-14 FMUL8ULx16 Operation

### CODE EXAMPLE 7-1 16-bit × 16-bit 16-bit Multiply

<code>fmul8sux16</code>	<code>%f0, %f1, %f2</code>
<code>fmul8ulx16</code>	<code>%f0, %f1, %f3</code>
<code>fpadd16</code>	<code>%f2, %f3, %f4</code>

## 7.30.6 FMULD8SUx16 Instruction

FMULD8SUx16 multiplies the most significant (“upper”) 8 bits of each 16-bit signed value in  $F_S[rs1]$  by the corresponding signed 16-bit fixed-point value in  $F_S[rs2]$ . 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 7-15 illustrates the operation.




FIGURE 7-15 FMULD88SUx16 Operation

# FMUL (partitioned)

## 7.30.7 FMULD8ULx16 Instruction

FMULD8ULx16 multiplies the unsigned least significant (“lower”) 8 bits of each 16-bit value in  $F_S[rs1]$  by the corresponding 16-bit fixed-point signed integer in  $F_S[rs2]$ . 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 7-16 illustrates the operation; CODE EXAMPLE 7-2 exemplifies the operation.




FIGURE 7-16 FMULD8ULx16 Operation

CODE EXAMPLE 7-2 16-bit x 16-bit 32-bit Multiply

```
fmuld8sux16    %f0, %f1, %f2
fmuld8ulx16    %f0, %f1, %f3
fpadd32        %f2, %f3, %f4
```

## 7.31 Floating-Point Multiply

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FMULs	11 0100	0 0100 1001	Multiply Single	<code>fmuls <i>freq<sub>rs1</sub></i>, <i>freq<sub>rs2</sub></i>, <i>freq<sub>rd</sub></i></code>	<b>A1</b>
FMULd	11 0100	0 0100 1010	Multiply Double	<code>fmuld <i>freq<sub>rs1</sub></i>, <i>freq<sub>rs2</sub></i>, <i>freq<sub>rd</sub></i></code>	<b>A1</b>
FMULq	11 0100	0 0100 1011	Multiply Quad	<code>fmulq <i>freq<sub>rs1</sub></i>, <i>freq<sub>rs2</sub></i>, <i>freq<sub>rd</sub></i></code>	<b>C3</b>
FsMULd	11 0100	0 0110 1001	Multiply Single to Double	<code>fsmuld <i>freq<sub>rs1</sub></i>, <i>freq<sub>rs2</sub></i>, <i>freq<sub>rd</sub></i></code>	<b>A1</b>
FdMULq	11 0100	0 0110 1110	Multiply Double to Quad	<code>fdmulq <i>freq<sub>rs1</sub></i>, <i>freq<sub>rs2</sub></i>, <i>freq<sub>rd</sub></i></code>	<b>C3</b>



*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** UltraSPARC Architecture 2005 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.

If the FPU is enabled, an attempt to execute an `FMULq` or `FdMULq` instruction causes an *fp\_exception\_other* (with `FSR.ftt = unimplemented_FPop`), since that instruction is not implemented in hardware in UltraSPARC Architecture 2005 implementations.

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for UltraSPARC Architecture 2005*.


## FMUL<s|d|q>

*Exceptions*

- illegal\_instruction*
- fp\_disabled*
- fp\_exception\_other* (FSR.ftt = unimplemented\_FPop (FMULq, FdMULq only))
- fp\_exception\_other* (FSR.ftt = unfinished\_FPop)
- fp\_exception\_ieee\_754* (any: NV; FMUL(s,d,q) only: OF, UF, NX)

## 7.32 Floating-Point Negate

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FNEGs	11 0100	0 0000 0101	Negate Single	<code>fnegs <i>freg<sub>rs2</sub></i>, <i>freg<sub>rd</sub></i></code>	<b>A1</b>
FNEGd	11 0100	0 0000 0110	Negate Double	<code>fnegd <i>freg<sub>rs2</sub></i>, <i>freg<sub>rd</sub></i></code>	<b>A1</b>
FNEGq	11 0100	0 0000 0111	Negate Quad	<code>fnegq <i>freg<sub>rs2</sub></i>, <i>freg<sub>rd</sub></i></code>	<b>C3</b>



*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** UltraSPARC Architecture 2005 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.

*Exceptions* *illegal\_instruction*  
*fp\_disabled*  
*fp\_exception\_other* (`FSR.ftt = unimplemented_FPop` (FNEGq only))

# FPAK

## 7.33 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	fpack16 <i>reg<sub>rs2</sub></i> , <i>reg<sub>rd</sub></i>	<b>C3</b>
FPAK32	0 0011 1010	Two 32-bit packs into 8 unsigned bits	f64	f64	f64	fpack32 <i>reg<sub>rs1</sub></i> , <i>reg<sub>rs2</sub></i> , <i>reg<sub>rd</sub></i>	<b>C3</b>
FPAKFIX	0 0011 1101	Four 16-bit packs into 16 signed bits	—	f64	f32	fpackfix <i>reg<sub>rs2</sub></i> , <i>reg<sub>rd</sub></i>	<b>C3</b>



*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.

In an UltraSPARC Architecture 2005 implementation, these instructions are not implemented in hardware, cause an *illegal\_instruction* exception, and are emulated in software.

*Exceptions* *illegal\_instruction*

*See Also* FEXPAND on page 184  
FPMERGE on page 219

# FPAK

## 7.33.1 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 7-17 illustrates the FPAK16 operation.




FIGURE 7-17 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
```

# FPAK

```
// 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};
```

## 7.33.2 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 7-18 illustrates the FPAK32 operation.




FIGURE 7-18 FPAK32 Operation

This operation, illustrated in FIGURE 7-18, 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.

# FPACK

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;
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});
```

# FPAK

## 7.33.3 FPAKFIX

FPAKFIX 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 7-19 illustrates the FPAKFIX operation.




FIGURE 7-19 FPAKFIX 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
```

# FPACK

```
trunc_signed_value ← tmp{(31+GSR.scale):16};
if (trunc_signed_value < -32768)
    signed_16bit_result ← -32768;
else if (trunc_signed_value > 32767)
    signed_16bit_result ← 32767;
else
    signed_16bit_result ← trunc_signed_value{31:16};
```

# FPADD

## 7.34 Fixed-point Partitioned Add vis 1

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class
FPADD16	0 0101 0000	Four 16-bit adds	f64	f64	f64	f_padd16 <i>freg<sub>rs1</sub></i> , <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i>	<b>A1</b>
FPADD16S	0 0101 0001	Two 16-bit adds	f32	f32	f32	f_padd16s <i>freg<sub>rs1</sub></i> , <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i>	<b>A1</b>
FPADD32	0 0101 0010	Two 32-bit adds	f64	f64	f64	f_padd32 <i>freg<sub>rs1</sub></i> , <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i>	<b>A1</b>
FPADD32S	0 0101 0011	One 32-bit add	f32	f32	f32	f_padd32s <i>freg<sub>rs1</sub></i> , <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i>	<b>A1</b>



*Description* FPADD16 (FPADD32) performs four 16-bit (two 32-bit) partitioned additions between the corresponding fixed-point values contained in the source operands ( $F_D[rs1]$ ,  $F_D[rs2]$ ). The result is placed in the destination register,  $F_D[rd]$ .

The 32-bit versions of these instructions (FPADD16S and FPADD32S) perform two 16-bit or one 32-bit partitioned additions.

Any carry out from each addition is discarded and a 2's-complement arithmetic result is produced.




FIGURE 7-20 FPADD16 Operation

# FPADD




FIGURE 7-21 FPADD32 Operation




FIGURE 7-22 FPADD16S Operation




FIGURE 7-23 FPADD32S Operation

# FPADD


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*

# FPMERGE

## 7.35 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<sub>rs1</sub></i>, <i>freg<sub>rs2</sub></i>, <i>freg<sub>rd</sub></i></code>	<b>C3</b>



*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 7-24 illustrates the operation.




FIGURE 7-24 FPMERGE Operation

	%d0	R1	G1	B1	A1	R2	G2	B2	A2	} packed representation
	%d2	R3	G3	B3	A3	R4	G4	B4	A4	
<code>fpmerge %f0, %f2, %d4</code>	<code>!r1</code>	R3	G1	G3	B1	B3	A1	A3	} intermediate	
<code>fpmerge %f1, %f3, %d6</code>	<code>!r2</code>	R4	G2	G4	B2	B4	A2	A4		
<code>fpmerge %f4, %f6, %d0</code>	<code>!r1</code>	R2	R3	R4	G1	G2	G3	G4	} planar representation	
<code>fpmerge %f5, %f7, %d2</code>	<code>!B1</code>	B2	B3	B4	A1	A2	A3	A4		
<code>fpmerge %f0, %f2, %d4</code>	<code>!r1</code>	B1	R2	B2	R3	B3	R4	B4	} intermediate	
<code>fpmerge %f1, %f3, %d6</code>	<code>!G1</code>	A1	G2	A2	G3	A3	G4	A4		
<code>fpmerge %f4, %f6, %d0</code>	<code>!R1</code>	G1	B1	A1	R2	G2	B2	A2	} packed representation	
<code>fpmerge %f5, %f7, %d2</code>	<code>!R3</code>	G3	B3	A3	R4	G4	B4	A4		

# FPMERGE

## CODE EXAMPLE 7-3 FPMERGE

In an UltraSPARC Architecture 2005 implementation, these instructions are not implemented in hardware, cause an *illegal\_instruction* exception, and are emulated in software.


*Exceptions*      *illegal\_instruction*

*See Also*        FPACK on page 210  
                  FEXPAND on page 184

# FPSUB

## 7.36 Fixed-point Partitioned Subtract VIS 1

Instruction	opf	Operation	s1	s2	d	Assembly Language Syntax	Class
FPSUB16	0 0101 0100	Four 16-bit subtracts	f64	f64	f64	fpsub16 <i>freq<sub>rs1</sub></i> , <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></i>	<b>A1</b>
FPSUB16S	0 0101 0101	Two 16-bit subtracts	f32	f32	f32	fpsub16s <i>freq<sub>rs1</sub></i> , <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></i>	<b>A1</b>
FPSUB32	0 0101 0110	Two 32-bit subtracts	f64	f64	f64	fpsub32 <i>freq<sub>rs1</sub></i> , <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></i>	<b>A1</b>
FPSUB32S	0 0101 0111	One 32-bit subtract	f32	f32	f32	fpsub32s <i>freq<sub>rs1</sub></i> , <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></i>	<b>A1</b>



*Description* FPSUB16 (FPSUB32) performs four 16-bit (two 32-bit) partitioned subtractions between the corresponding fixed-point values contained in the source operands ( $F_D[rs1]$ ,  $F_D[rs2]$ ). The values in  $F_D[rs2]$  are subtracted from those in  $F_D[rs1]$ , and the result is placed in the destination register,  $F_D[rd]$ .

The 32-bit versions of these instructions (FPSUB16S and FPSUB32S) perform two 16-bit or one 32-bit partitioned subtractions.

Any carry out from each subtraction is discarded and a 2's-complement arithmetic result is produced.




FIGURE 7-25 FPSUB16 Operation

# FPSUB




FIGURE 7-26 FPSUB32 Operation




FIGURE 7-27 FPSUB16S Operation




FIGURE 7-28 FPSUB32S Operation

## FPSUB


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*

## F Register 1-operand Logical Ops

### 7.37 F Register Logical Operate (1 operand) VIS 1

Instruction	opf	Operation	Assembly Language Syntax	Class
FZERO	0 0110 0000	Zero fill	<code>fzero</code> <i>freg<sub>rd</sub></i>	<b>A1</b>
FZEROs	0 0110 0001	Zero fill, 32-bit	<code>fzeros</code> <i>freg<sub>rd</sub></i>	<b>A1</b>
FONE	0 0111 1110	One fill	<code>fone</code> <i>freg<sub>rd</sub></i>	<b>A1</b>
FONEs	0 0111 1111	One fill, 32-bit	<code>foness</code> <i>freg<sub>rd</sub></i>	<b>A1</b>



*Description* FZERO and FONE 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[s] or FONE[s] instruction causes an *fp\_disabled* exception.


*Exceptions* *illegal\_instruction*  
*fp\_disabled*

*See Also* F Register 2-operand Logical Operations on page 225  
F Register 3-operand Logical Operations on page 227

## F Register 2-operand Logical Ops

### 7.38 F Register Logical Operate (2 operand) VIS 1

Instruction	opf	Operation	Assembly Language Syntax	Class
FSRC1	0 0111 0100	Copy $F_D[rs1]$ to $F_D[rd]$	<code>fsrc1</code> <i>freq<sub>rs1</sub>, freq<sub>rd</sub></i>	<b>A1</b>
FSRC1s	0 0111 0101	Copy $F_S[rs1]$ to $F_S[rd]$ , 32-bit	<code>fsrc1s</code> <i>freq<sub>rs1</sub>, freq<sub>rd</sub></i>	<b>A1</b>
FSRC2	0 0111 1000	Copy $F_D[rs2]$ to $F_D[rd]$	<code>fsrc2</code> <i>freq<sub>rs2</sub>, freq<sub>rd</sub></i>	<b>A1</b>
FSRC2s	0 0111 1001	Copy $F_S[rs2]$ to $F_S[rd]$ , 32-bit	<code>fsrc2s</code> <i>freq<sub>rs2</sub>, freq<sub>rd</sub></i>	<b>A1</b>
FNOT1	0 0110 1010	Negate (1's complement) $F_D[rs1]$	<code>fnot1</code> <i>freq<sub>rs1</sub>, freq<sub>rd</sub></i>	<b>A1</b>
FNOT1s	0 0110 1011	Negate (1's complement) $F_S[rs1]$ , 32-bit	<code>fnot1s</code> <i>freq<sub>rs1</sub>, freq<sub>rd</sub></i>	<b>A1</b>
FNOT2	0 0110 0110	Negate (1's complement) $F_D[rs2]$	<code>fnot2</code> <i>freq<sub>rs2</sub>, freq<sub>rd</sub></i>	<b>A1</b>
FNOT2s	0 0110 0111	Negate (1's complement) $F_S[rs2]$ , 32-bit	<code>fnot2s</code> <i>freq<sub>rs2</sub>, freq<sub>rd</sub></i>	<b>A1</b>



**Description** The standard 64-bit versions of these instructions perform one of four 64-bit logical operations on the 64-bit floating-point 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(s) or FNOT1(s) 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], FNOT1[s], FSRC1[s], or FNOT1[s] instruction causes an *fp\_disabled* exception.

**Programming Note** FSRC1s (FSRC1) functions similarly to FMOVs (FMOVd), except that FSRC1s (FSRC1) does not modify the FSR register while FMOVs (FMOVd) update some fields of FSR (see *Floating-Point Move* on page 191). Programmers are encouraged to use FMOVs (FMOVd) instead of FSRC1s (FSRC1) whenever practical.

**Exceptions** *illegal\_instruction*  
*fp\_disabled*

## F Register 2-operand Logical Ops

*See Also*

*Floating-Point Move* on page 191


F Register 1-operand Logical Operations on page 224

F Register 3-operand Logical Operations on page 227

## F Register 3-operand Logical Ops

### 7.39 F Register Logical Operate (3 operand) VIS 1

Instruction	opf	Operation	Assembly Language Syntax		Class
FOR	0 0111 1100	Logical <b>or</b>	<code>for</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FORs	0 0111 1101	Logical <b>or</b> , 32-bit	<code>fors</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FNOR	0 0110 0010	Logical <b>nor</b>	<code>fnor</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FNORs	0 0110 0011	Logical <b>nor</b> , 32-bit	<code>fnors</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FAND	0 0111 0000	Logical <b>and</b>	<code>fand</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FANDs	0 0111 0001	Logical <b>and</b> , 32-bit	<code>fands</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FNAND	0 0110 1110	Logical <b>nand</b>	<code>fnand</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FNANDs	0 0110 1111	Logical <b>nand</b> , 32-bit	<code>fnands</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FXOR	0 0110 1100	Logical <b>xor</b>	<code>fxor</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FXORs	0 0110 1101	Logical <b>xor</b> , 32-bit	<code>fxors</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FXNOR	0 0111 0010	Logical <b>xnor</b>	<code>fxnor</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FXNORs	0 0111 0011	Logical <b>xnor</b> , 32-bit	<code>fxnors</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FORNOT1	0 0111 1010	( <b>not</b> F[rs1]) <b>or</b> F[rs2]	<code>fornot1</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FORNOT1s	0 0111 1011	( <b>not</b> F[rs1]) <b>or</b> F[rs2], 32-bit	<code>fornot1s</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FORNOT2	0 0111 0110	F[rs1] <b>or</b> ( <b>not</b> F[rs2])	<code>fornot2</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FORNOT2s	0 0111 0111	F[rs1] <b>or</b> ( <b>not</b> F[rs2]), 32-bit	<code>fornot2s</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FANDNOT1	0 0110 1000	( <b>not</b> F[rs1]) <b>and</b> F[rs2]	<code>fandnot1</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FANDNOT1s	0 0110 1001	( <b>not</b> F[rs1]) <b>and</b> F[rs2], 32-bit	<code>fandnot1s</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FANDNOT2	0 0110 0100	F[rs1] <b>and</b> ( <b>not</b> F[rs2])	<code>fandnot2</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>
FANDNOT2s	0 0110 0101	F[rs1] <b>and</b> ( <b>not</b> F[rs2]), 32-bit	<code>fandnot2s</code>	<code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>A1</b>



**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 224  
F Register 2-operand Logical Operations on page 225

# FSQRT<s|d|q> Instructions

## 7.40 Floating-Point Square Root

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FSQRTs	11 0100	0 0010 1001	Square Root Single	<code>fsqrts <i>freq<sub>rs2</sub>, freq<sub>rd</sub></i></code>	<b>A1</b>
FSQRTd	11 0100	0 0010 1010	Square Root Double	<code>fsqrtd <i>freq<sub>rs2</sub>, freq<sub>rd</sub></i></code>	<b>A1</b>
FSQRTq	11 0100	0 0010 1011	Square Root Quad	<code>fsqrtq <i>freq<sub>rs2</sub>, freq<sub>rd</sub></i></code>	<b>C3</b>



**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** UltraSPARC Architecture 2005 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.

If the FPU is enabled, an *fp\_exception\_other* (with FSR.ftt = unimplemented\_FPop) exception occurs, since the FSQRT instructions are not implemented in hardware in UltraSPARC Architecture 2005 implementations.

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for UltraSPARC Architecture 2005*.

**Exceptions** *illegal\_instruction*  
*fp\_disabled*  
*fp\_exception\_other* (FSR.ftt = unimplemented\_FPop (FSQRT is not implemented in hardware))

# F<s|d|q>TOi

## 7.41 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	<i>fstox</i> <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></i>	<b>A1</b>
FdTOx	0 1000 0010	Convert Double to 64-bit Integer	—	f64	f64	<i>fdtox</i> <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></i>	<b>A1</b>
FqTOx	0 1000 0011	Convert Quad to 64-bit Integer	—	f128	f64	<i>fqtox</i> <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></i>	<b>C3</b>
FsTOi	0 1101 0001	Convert Single to 32-bit Integer	—	f32	f32	<i>fstoi</i> <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></i>	<b>A1</b>
FdTOi	0 1101 0010	Convert Double to 32-bit Integer	—	f64	f32	<i>fdtoi</i> <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></i>	<b>A1</b>
FqTOi	0 1101 0011	Convert Quad to 32-bit Integer	—	f128	f32	<i>fqtoi</i> <i>freq<sub>rs2</sub></i> , <i>freq<sub>rd</sub></i>	<b>C3</b>



*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 F<sub>D</sub>[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 F<sub>S</sub>[rd].

The result is always rounded toward zero; that is, the rounding direction (rd) field of the FSR register is ignored.

**Note** UltraSPARC Architecture 2005 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.

If the FPU is enabled, FqTOi and FqTOx cause *fp\_exception\_other* (with FSR.ftt = unimplemented\_FPop), since those instructions are not implemented in hardware in UltraSPARC Architecture 2005 implementations.

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 383.

## F<s|d|q>TOi

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for UltraSPARC Architecture 2005*.

### *Exceptions*

*illegal\_instruction*

*fp\_disabled*


*fp\_exception\_other* (FSR.ftt = unimplemented\_FPop (FqTOx, FqTOi only))

*fp\_exception\_ieee\_754* (NV, NX)

# F<s|d|q>TO<s|d|q>

## 7.42 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<sub>rs2</sub>, freq<sub>rd</sub></i>	<b>A1</b>
FsTOq	11 0100	0 1100 1101	Convert Single to Quad	—	f32	f128	<code>fstoq</code> <i>freq<sub>rs2</sub>, freq<sub>rd</sub></i>	<b>C3</b>
FdTOs	11 0100	0 1100 0110	Convert Double to Single	—	f64	f32	<code>fdtos</code> <i>freq<sub>rs2</sub>, freq<sub>rd</sub></i>	<b>A1</b>
FdTOq	11 0100	0 1100 1110	Convert Double to Quad	—	f64	f128	<code>fdtoq</code> <i>freq<sub>rs2</sub>, freq<sub>rd</sub></i>	<b>C3</b>
FqTOs	11 0100	0 1100 0111	Convert Quad to Single	—	f128	f32	<code>fqtos</code> <i>freq<sub>rs2</sub>, freq<sub>rd</sub></i>	<b>C3</b>
FqTOd	11 0100	0 1100 1011	Convert Quad to Double	—	f128	f64	<code>fqtod</code> <i>freq<sub>rs2</sub>, freq<sub>rd</sub></i>	<b>C3</b>



*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** UltraSPARC Architecture 2005 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.

If the FPU is enabled, `FsTOq`, `FdTOq`, `FqTOs`, and `FqTOd` cause *fp\_exception\_other* (with `FSR.ftt = unimplemented_FPop`), since those instructions are not implemented in hardware in UltraSPARC Architecture 2005 implementations.

`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.

## F<s|d|q>TO<s|d|q>

*Untrapped Result in Different Format from Operands* on page 380 defines the rules for converting NaNs from one floating-point format to another.

**Note** For FdTOs and FsTOd, 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 UltraSPARC Architecture 2005*.

### Exceptions

*illegal\_instruction*

*fp\_disabled*

*fp\_exception\_other* (FSR.ftt = unimplemented\_FPop (FsTOq, FqTOs, FdTOq, and FqTOd only))

*fp\_exception\_other* (FSR.ftt = unfinished\_FPop)

*fp\_exception\_ieee\_754* (NV)

*fp\_exception\_ieee\_754* (OF, UF, NX (FqTOd, FqTOs, and FdTOs))

# FSUB

## 7.43 Floating-Point Subtract

Instruction	op3	opf	Operation	Assembly Language Syntax	Class
FSUBs	11 0100	0 0100 0101	Subtract Single	<code>fsubs</code> $freq_{rs1}, freq_{rs2}, freq_{rd}$	<b>A1</b>
FSUBd	11 0100	0 0100 0110	Subtract Double	<code>fsubd</code> $freq_{rs1}, freq_{rs2}, freq_{rd}$	<b>A1</b>
FSUBq	11 0100	0 0100 0111	Subtract Quad	<code>fsubq</code> $freq_{rs1}, freq_{rs2}, freq_{rd}$	<b>C3</b>



*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** UltraSPARC Architecture 2005 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.

If the FPU is enabled, `FSUBq` causes an *fp\_exception\_other* (with `FSR.ftt = unimplemented_FPop`), since that instruction is not implemented in hardware in UltraSPARC Architecture 2005 implementations.

**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 UltraSPARC Architecture 2005*.

*Exceptions* *illegal\_instruction*  
*fp\_disabled*  
*fp\_exception\_other* (`FSR.ftt = unimplemented_FPop` (`FSUBq`))  
*fp\_exception\_other* (`FSR.ftt = unfinished_FPop`)  
*fp\_exception\_ieee\_754* (`OF, UF, NX, NV`)

# FxTO(<s|d|q>

## 7.44 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<sub>rs2</sub>, freq<sub>rd</sub></i></code>	<b>A1</b>
FxTOd	11 0100	0 1000 1000	Convert 64-bit Integer to Double	—	i64	f64	<code>fxtod <i>freq<sub>rs2</sub>, freq<sub>rd</sub></i></code>	<b>A1</b>
FxTOq	11 0100	0 1000 1100	Convert 64-bit Integer to Quad	—	i64	f128	<code>fxtoq <i>freq<sub>rs2</sub>, freq<sub>rd</sub></i></code>	<b>C3</b>



**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** UltraSPARC Architecture 2005 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.

If the FPU is enabled, FxTOq causes an *fp\_exception\_other* (with FSR.ftt = unimplemented\_FPop), since that instruction is not implemented in hardware in UltraSPARC Architecture 2005 implementations.

For more details regarding floating-point exceptions, see Chapter 8, *IEEE Std 754-1985 Requirements for UltraSPARC Architecture 2005*.

**Exceptions** *illegal\_instruction*  
*fp\_disabled*  
*fp\_exception\_other* (FSR.ftt = unimplemented\_FPop (FxTOq only))  
*fp\_exception\_ieee\_754* (NX (FxTOs and FxTOd only))

# ILLTRAP

## 7.45 Illegal Instruction Trap

Instruction	op	op2	Operation	Assembly Language Syntax	Class
ILLTRAP	00	000	<i>illegal_instruction</i> trap	<i>illtrap</i> <i>const22</i>	<b>A1</b>



*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 Note** | Except for its name, this instruction is identical to the SPARC V8 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*

# IMPDEP

## 7.46 Implementation-Dependent Instructions

Instruction	op3	op4	Operation	Class
IMPDEP1	11 0110	(any)	Implementation-Dependent Instruction 1	<b>N3</b>
IMPDEP2A	11 0111	0	Implementation-Dependent Instruction 2A	<b>N3</b>
IMPDEP2B	11 0111	1, 2, 3	Implementation-Dependent Instruction 2B	<b>N3</b>



*Description* **IMPL. DEP. #106-V9:** The IMPDEP2A opcode space is completely implementation dependent. Implementation-dependent aspects of IMPDEP2A instructions include their operation, the interpretation of bits 29–25, 18–7, and 4–0 in their encodings, and which (if any) exceptions they may cause.

IMPDEP2B opcodes are reserved; see *IMDEP2B Opcodes* on page 237.

See “Implementation-Dependent and Reserved Opcodes” in the “Extending the UltraSPARC Architecture” section of the separate document *UltraSPARC Architecture Application Notes*, for information about extending the instruction set by means of implementation-dependent instructions.

**Compatibility Note** IMPDEP2A and IMPDEP2B are subsets of the SPARC V9 IMPDEP2 opcode space. The IMPDEP1 opcode space from SPARC V9 is occupied by various VIS instructions in the UltraSPARC Architecture, so it should not be used for implementation-dependent instructions.

*Exceptions* implementation-dependent (IMPDEP2A, IMPDEP2B)

### 7.46.1 IMPDEP1 Opcodes VIS 1, 2

All operands of instructions using IMPDEP1 opcodes are in floating-point registers, unless otherwise specified. Pixel values are stored in single-precision floating point registers and fixed values are stored in double-precision floating point registers, unless otherwise specified.

**Note** All instructions, regardless of whether they use floating-point registers or integer registers, leave FSR.cexc and FSR.aexc unchanged.

# IMPDEP

## 7.46.1.1 Opcode Formats

Most of the VIS instruction set maps to the opcode space reserved for the Implementation-Dependent Instruction 1 (op3 = IMPDEP1 =  $36_{16}$ ) instructions.

## 7.46.2 IMDEP2B Opcodes

No instructions are currently encoded in the IMPDEP2B opcode space; it is a reserved opcode space.

# INVALW

## 7.47 Mark Register Window Sets as “Invalid”

Instruction	Operation	Assembly Language Syntax	Class
INVALW <sup>P</sup>	Mark all register window sets as “invalid”	invalw	C1



**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. This instruction allows window manipulations to be atomic, without the value of  $N\_REG\_WINDOWS$  being visible to privileged software and without an assumption that  $N\_REG\_WINDOWS$  is constant (since hyperprivileged software can migrate a thread among virtual processors, across which  $N\_REG\_WINDOWS$  may vary).

In an UltraSPARC Architecture 2005 implementation, these instructions are not implemented in hardware, cause an *illegal\_instruction* exception, and are emulated in software.


**Exceptions** *illegal\_instruction* (not implemented in hardware in UltraSPARC Architecture 2005)

**See Also** ALLCLEAN on page 148  
NORMALW on page 285  
OTHERW on page 287  
RESTORED on page 307  
SAVED on page 315

# JMPL

## 7.48 Jump and Link

Instruction	op3	Operation	Assembly Language Syntax	Class
JMPL	11 1000	Jump and Link	<code>jmp1 address, reg<sub>rd</sub></code>	<b>A1</b>



**Description** The JMPL instruction causes a register-indirect delayed control transfer to the address given by “R[rs1] + R[rs2]” if *i* field = 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.

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]. (closed impl. dep. #125-V9-Cs10)

**Exceptions** *illegal\_instruction*  
*mem\_address\_not\_aligned*

**See Also** CALL on page 162  
Bicc on page 154  
BPCC on page 160

# LD

## 7.49 Load Integer

Instruction	op3	Operation	Assembly Language Syntax	Class
LDSB	00 1001	Load Signed Byte	ldsb [address], reg <sub>rd</sub>	A1
LDSH	00 1010	Load Signed Halfword	ldsh [address], reg <sub>rd</sub>	A1
LDSW	00 1000	Load Signed Word	ldsw [address], reg <sub>rd</sub>	A1
LDUB	00 0001	Load Unsigned Byte	ldub [address], reg <sub>rd</sub>	A1
LDUH	00 0010	Load Unsigned Halfword	lduh [address], reg <sub>rd</sub>	A1
LDUW	00 0000	Load Unsigned Word	lduw <sup>†</sup> [address], reg <sub>rd</sub>	A1
LDX	00 1011	Load Extended Word	ldx [address], reg <sub>rd</sub>	A1

<sup>†</sup> 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 104). 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.

If the effective address is not halfword-aligned, an attempt to execute an LDUH or LDSH causes a *mem\_address\_not\_aligned* exception. If the effective address is not word-aligned, an attempt to execute an LDUW or LDSW instruction causes a *mem\_address\_not\_aligned* exception. If the effective address is not doubleword-aligned, an attempt to execute an LDX instruction causes a *mem\_address\_not\_aligned* exception.

**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 266 for details.

# LD

*Exceptions*

- illegal\_instruction*
- mem\_address\_not\_aligned* (all except LDSB, LDUB)
- VA\_watchpoint*
- data\_access\_exception*
- fast\_data\_access\_MMU\_miss*
- data\_access\_MMU\_miss*
- data\_access\_MMU\_error*

# LDA

## 7.50 Load Integer from Alternate Space

Instruction	op3	Operation	Assembly Language Syntax		Class
LDSBA <sup>PASI</sup>	01 1001	Load Signed Byte from Alternate Space	ldsba	[regaddr] imm_asi, reg_rd ldsba [reg_plus_imm] %asi, reg_rd	A1
LDSHA <sup>PASI</sup>	01 1010	Load Signed Halfword from Alternate Space	ldsha	[regaddr] imm_asi, reg_rd ldsha [reg_plus_imm] %asi, reg_rd	A1
LDSWA <sup>PASI</sup>	01 1000	Load Signed Word from Alternate Space	ldswa	[regaddr] imm_asi, reg_rd ldswa [reg_plus_imm] %asi, reg_rd	A1
LDUBA <sup>PASI</sup>	01 0001	Load Unsigned Byte from Alternate Space	lduba	[regaddr] imm_asi, reg_rd lduba [reg_plus_imm] %asi, reg_rd	A1
LDUHA <sup>PASI</sup>	01 0010	Load Unsigned Halfword from Alternate Space	lduha	[regaddr] imm_asi, reg_rd lduha [reg_plus_imm] %asi, reg_rd	A1
LDUWA <sup>PASI</sup>	01 0000	Load Unsigned Word from Alternate Space	lduwa†	[regaddr] imm_asi, reg_rd lduwa [reg_plus_imm] %asi, reg_rd	A1
LDXA <sup>PASI</sup>	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 access is privileged if bit 7 of the ASI is 0; otherwise, it is not privileged. 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 268 for details.

An attempt to execute a load integer from alternate space instruction when i = 0 and instruction bits 12:5 are nonzero causes an *illegal\_instruction* exception.

# LDA

If the effective address is not halfword-aligned, an attempt to execute an LDUHA or LDSHA instruction causes a *mem\_address\_not\_aligned* exception. If the effective address is not word-aligned, an attempt to execute an LDUWA or LDSWA instruction causes a *mem\_address\_not\_aligned* exception. If the effective address is not doubleword-aligned, an attempt to execute an LDXA instruction causes a *mem\_address\_not\_aligned* exception.

In nonprivileged mode (PSTATE.priv = 0 and HPSTATE.hpriv = 0), if bit 7 of the ASI is 0, these instructions cause a *privileged\_action* exception. In privileged mode (PSTATE.priv = 1 and HPSTATE.hpriv = 0), if the ASI is in the range 30<sub>16</sub> to 7F<sub>16</sub>, these instructions cause a *privileged\_action* exception.

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 *data\_access\_exception* exception.

---

ASIs valid for LDSBA, LDSHA, LDSWA, LDUBA, LDUHA, and LDUWA	
ASI_NUCLEUS	ASI_NUCLEUS_LITTLE
ASI_AS_IF_USER_PRIMARY	ASI_AS_IF_USER_PRIMARY_LITTLE
ASI_AS_IF_USER_SECONDARY	ASI_AS_IF_USER_SECONDARY_LITTLE
ASI_REAL	ASI_REAL_LITTLE
ASI_REAL_IO	ASI_REAL_IO_LITTLE
ASI_PRIMARY	ASI_PRIMARY_LITTLE
ASI_SECONDARY	ASI_SECONDARY_LITTLE
ASI_PRIMARY_NO_FAULT	ASI_PRIMARY_NO_FAULT_LITTLE
ASI_SECONDARY_NO_FAULT	ASI_SECONDARY_NO_FAULT_LITTLE

LDXA can be used with any ASI (including, but not limited to, the above list), 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 *data\_access\_exception* exception.

---

ASIs invalid for LDXA (cause <i>data_access_exception</i> exception)	
24 <sub>16</sub> (aliased to 27 <sub>16</sub> , ASI_LDTX_N)	2C <sub>16</sub> (aliased to 2F <sub>16</sub> , ASI_LDTX_NL)
22 <sub>16</sub> (ASI_LDTX_AIUP)	2A <sub>16</sub> (ASI_LDTD_AIUP_L)
23 <sub>16</sub> (ASI_LDTX_AIUS)	2B <sub>16</sub> (ASI_LDTX_AIUS_L)
26 <sub>16</sub> (ASI_LDTX_REAL)	2E <sub>16</sub> (ASI_LDTX_REAL_L)
27 <sub>16</sub> (ASI_LDTX_N)	2F <sub>16</sub> (ASI_LDTX_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
ASI_PST8_PRIMARY	ASI_PST8_PRIMARY_LITTLE
ASI_PST8_SECONDARY	ASI_PST8_SECONDARY_LITTLE
ASI_PST16_PRIMARY	ASI_PST16_PRIMARY_LITTLE
ASI_PST16_SECONDARY	ASI_PST16_SECONDARY_LITTLE
ASI_PST32_PRIMARY	ASI_PST32_PRIMARY_LITTLE
ASI_PST32_SECONDARY	ASI_PST32_SECONDARY_LITTLE

# LDA

---

## ASIs invalid for LDXA (cause *data\_access\_exception* exception)

ASI_FL8_PRIMARY	ASI_FL8_PRIMARY_LITTLE
ASI_FL8_SECONDARY	ASI_FL8_SECONDARY_LITTLE
ASI_FL16_PRIMARY	ASI_FL16_PRIMARY_LITTLE
ASI_FL16_SECONDARY	ASI_FL16_SECONDARY_LITTLE
ASI_BLOCK_COMMIT_PRIMARY	ASI_BLOCK_COMMIT_SECONDARY
E2 <sub>16</sub> (ASI_LDTX_P)	EA <sub>16</sub> (ASI_LDTX_PL)
E3 <sub>16</sub> (ASI_LDTX_S)	EB <sub>16</sub> (ASI_LDTX_SL)
ASI_BLOCK_PRIMARY	ASI_BLOCK_PRIMARY_LITTLE
ASI_BLOCK_SECONDARY	ASI_BLOCK_SECONDARY_LITTLE

*Exceptions*     *mem\_address\_not\_aligned* (all except LDSBA and LDUBA)  
*privileged\_action*  
*VA\_watchpoint*  
*data\_access\_exception*  
*fast\_data\_access\_MMU\_miss*  
*data\_access\_MMU\_miss*  
*data\_access\_MMU\_error*


*See Also*       LD on page 240  
                 STA on page 324

# LDBLOCKF

## 7.51 Block Load VIS 1

The LDBLOCKF instruction is intended to be a processor-specific instruction, which may or may not be implemented in future UltraSPARC Architecture implementations. Therefore, it should only be used in platform-specific dynamically-linked libraries, in hyperprivileged software, or in software created by a runtime code generator that is aware of the specific virtual processor implementation on which it is executing.

Instruc-tion	ASI Value	Operation	Assembly Language Syntax	Class
LDBLOCKF	16 <sub>16</sub>	64-byte block load from primary address space, user privilege	ldda [regaddr] #ASI_BLK_AIUP, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	<b>B2</b>
LDBLOCKF	17 <sub>16</sub>	64-byte block load from secondary address space, user privilege	ldda [regaddr] #ASI_BLK_AIUS, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	<b>B2</b>
LDBLOCKF	1E <sub>16</sub>	64-byte block load from primary address space, little-endian, user privilege	ldda [regaddr] #ASI_BLK_AIUPL, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	<b>B2</b>
LDBLOCKF	1F <sub>16</sub>	64-byte block load from secondary address space, little-endian, user privilege	ldda [regaddr] #ASI_BLK_AIUSL, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	<b>B2</b>
LDBLOCKF	F0 <sub>16</sub>	64-byte block load from primary address space	ldda [regaddr] #ASI_BLK_P, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	<b>B2</b>
LDBLOCKF	F1 <sub>16</sub>	64-byte block load from secondary address space	ldda [regaddr] #ASI_BLK_S, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	<b>B2</b>
LDBLOCKF	F8 <sub>16</sub>	64-byte block load from primary address space, little-endian	ldda [regaddr] #ASI_BLK_PL, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	<b>B2</b>
LDBLOCKF	F9 <sub>16</sub>	64-byte block load from secondary address space, little-endian	ldda [regaddr] #ASI_BLK_SL, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	<b>B2</b>



*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.

# LDBLOCKF

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.

The block load instruction is intended to support fast block-copy operations.

<b>Programming Note</b>	LDBLOCKF is intended to be a processor-specific instruction (see the warning at the top of page 245). If LDBLOCKF <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 #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 (LDBLOCKF is not required to follow TSO memory ordering)
- Whether LDBLOCKF 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 appears to execute out of order, or follow LoadLoad ordering (with respect to older loads, younger loads, and other LDBLOCKFs)
- Whether LDBLOCKF follows register-dependency interlocks, as do ordinary load instructions
- Whether LDBLOCKFs to non-cacheable locations are (a) strictly ordered, (b) not strictly ordered and cause an *illegal\_instruction* exception, or (c) not strictly ordered and silently execute without causing an exception (option (c) is strongly discouraged)
- Whether *VA\_watchpoint* exceptions are recognized on accesses to all 64 bytes of a LDBLOCKF (the recommended behavior), or only on the first eight bytes
- Whether the MMU ignores the side-effect bit (TTE.e) for LDBLOCKF accesses

# LDBLOCKF

<b>Programming Note</b>	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.

<b>Implementation Note</b>	In all UltraSPARC Architecture implementations, the MMU ignores the side-effect bit (TTE.e) for LDBLOCKF 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 instruction causes an *fp\_disabled* exception.

If the least significant 6 bits of the effective memory address in an LDBLOCKF instruction are nonzero, a *mem\_address\_not\_aligned* exception occurs.

In nonprivileged mode (PSTATE.priv = 0 and HPSTATE.hpriv = 0), if bit 7 of the ASI is 0 (ASIs 16<sub>16</sub>, 17<sub>16</sub>, 1E<sub>16</sub>, and 1F<sub>16</sub>), LDBLOCKF causes a *privileged\_action* exception.

An access caused by LDBLOCKF may trigger a *VA\_watchpoint* exception (impl. dep. #410-S10).

<b>Implementation Note</b>	LDBLOCKF shares an opcode with LDDEFA and LDSHORTF; it is distinguished by the ASI used.
----------------------------	--

*Exceptions*

- illegal\_instruction*
- fp\_disabled*
- mem\_address\_not\_aligned*
- privileged\_action*
- VA\_watchpoint* (impl. dep. #410-S10)
- data\_access\_exception*
- fast\_data\_access\_MMU\_miss*

# LDBLOCKF

*data\_access\_MMU\_miss*  
*data\_access\_MMU\_error*

*See Also*      STBLOCKF on page 328

# LDF / LDDF / LDQF / LDXFSR

## 7.52 Load Floating-Point

Instruction	op3	rd	Operation	Assembly Language Syntax	Class
LDF	10 0000	0–31	Load Floating-Point Register	ld [address], freg <sub>rd</sub>	A1
LDDF	10 0011	‡	Load Double Floating-Point Register	ldd [address], freg <sub>rd</sub>	A1
LDQF	10 0010	‡	Load Quad Floating-Point Register	ldq [address], freg <sub>rd</sub>	C3
LDXFSR	10 0001	1	Load Floating-Point State Register	ldx [address], %fsr	A1
—	10 0001	2–31	Reserved		

‡ 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).

The load floating-point state register instruction (LDXFSR) 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*, *flt*, *qne*, or reserved fields of FSR (see page 61).

**Programming Note** For future compatibility, software should only issue an LDXFSR instruction with a zero value (or a value previously read from the same field) written into any reserved field of FSR.

These load floating-point instructions access memory using the implicit ASI (see page 104).

If  $i = 0$ , the effective address for these instructions is “ $R[rs1] + R[rs2]$ ” and if  $i = 1$ , the effective address is “ $R[rs1] + \text{sign\_ext}(\text{simm13})$ ”.

# LDF / LDDF / LDQF / LDXFSR

**Exceptions.** An attempt to execute an LDF, LDDF, LDQF, or LDXFSR instruction when  $i = 0$  and instruction bits 12:5 are nonzero causes an *illegal\_instruction* exception. An attempt to execute an instruction encoded as  $op = 2$ ,  $op3 = 21_{16}$ , and  $rd > 1$  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, LDQF, or LDXFSR instruction causes an *fp\_disabled* exception.

If the effective address is not word-aligned, an attempt to execute an LDF instruction causes a *mem\_address\_not\_aligned* exception. If the effective address is not doubleword-aligned, an attempt to execute an LDXFSR instruction causes a *mem\_address\_not\_aligned* exception.

LDDF requires only word alignment. However, if the effective address is word-aligned but not doubleword-aligned, an attempt to execute an LDDF instruction 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)).

LDQF requires only word alignment. However, if the effective address is word-aligned but not quadword-aligned, an attempt to execute an LDQF instruction 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)).

<b>Programming Note</b>	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 <i>not</i> properly aligned.
-------------------------	---

An attempt to execute an LDQF instruction when  $rd\{1\} \neq 0$  causes an *fp\_exception\_other* ( $FSR.ftt = \text{invalid\_fp\_register}$ ) exception.

<b>Implementation Note</b>	Since UltraSPARC Architecture 2005 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 $FSR.ftt = \text{invalid\_fp\_register}$ ) 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.

# LDF / LDDF / LDQF / LDXFSR

**IMPL. DEP. #44-V8-Cs10(a):** 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.

<b>Implementation</b>	LDXFSR shares an opcode with the LDFSR instruction (and possibly with other implementation-dependent instructions); they are differentiated by the instruction rd field. An attempt to execute the op = 11 <sub>2</sub> , op3 = 10 0001 <sub>2</sub> opcode with an invalid rd value (rd > 1) causes an <i>illegal_instruction</i> exception.
<b>Note</b>	

*Exceptions*

- illegal\_instruction*
- fp\_disabled*
- LDDF\_mem\_address\_not\_aligned*
- mem\_address\_not\_aligned*
- fp\_exception\_other* (FSR.ftt = invalid\_fp\_register (LDQF only))
- VA\_watchpoint*
- data\_access\_exception*
- fast\_data\_access\_MMU\_miss*
- data\_access\_MMU\_miss*
- data\_access\_MMU\_error*

*See Also*

- Load Floating-Point from Alternate Space* on page 252
- Load Floating-Point State Register* on page 256
- Store Floating-Point* on page 332

# LDFA / LDDFA / LDQFA

## 7.53 Load Floating-Point from Alternate Space

Instruction	op3	rd	Operation	Assembly Language Syntax	Class
LDFA <sup>PASI</sup>	11 0000	0–31	Load Floating-Point Register from Alternate Space	ldfa [regaddr] imm_asi, freg <sub>rd</sub> ldfa [reg_plus_imm] %asi, freg <sub>rd</sub>	<b>A1</b>
LDDFA <sup>PASI</sup>	11 0011	‡	Load Double Floating-Point Register from Alternate Space	lddda [regaddr] imm_asi, freg <sub>rd</sub> lddda [reg_plus_imm] %asi, freg <sub>rd</sub>	<b>A1</b>
LDQFA <sup>PASI</sup>	11 0010	‡	Load Quad Floating-Point Register from Alternate Space	ldqda [regaddr] imm_asi, freg <sub>rd</sub> ldqda [reg_plus_imm] %asi, freg <sub>rd</sub>	<b>C3</b>

‡ 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 (LDFA) 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 LDFA, LDDFA, or LDQFA instruction causes an *fp\_disabled* exception.

LDFA causes a *mem\_address\_not\_aligned* exception if the effective memory address is not word-aligned.

**V9 Compatibility** | LDFA, LDDFA, and LDQFA cause a *privileged\_action* exception if  
**Note** | `PSTATE.priv = 0` and bit 7 of the ASI is 0.

# LDFA / LDDFA / LDQFA

LDDFA requires only word alignment. However, if the effective address is word-aligned but not doubleword-aligned, LDDFA 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(b)).

LDQFA requires only word alignment. However, if the effective address is word-aligned but not quadword-aligned, LDQFA 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(b)).

An attempt to execute an LDQFA instruction when  $rd\{1\} \neq 0$  causes an *fp\_exception\_other* (with  $FSR.ftt = \text{invalid\_fp\_register}$ ) exception.

<b>Implementation Note</b>	Since UltraSPARC Architecture 2005 processors do not implement in hardware instructions (including LDQFA) that refer to quad-precision floating-point registers, the <i>LDQF_mem_address_not_aligned</i> and <i>fp_exception_other</i> (with $FSR.ftt = \text{invalid\_fp\_register}$ ) 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.
----------------------------	---

<b>Programming Note</b>	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 <i>not</i> properly aligned.
-------------------------	---

In nonprivileged mode ( $PSTATE.priv = 0$  and  $HPSTATE.hpriv = 0$ ), if bit 7 of the ASI is 0, this instruction causes a *privileged\_action* exception. In privileged mode ( $PSTATE.priv = 1$  and  $HPSTATE.hpriv = 0$ ), 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 *data\_access\_exception* exception.

---

## ASIs valid for LDFA and LDQFA

ASI_NUCLEUS	ASI_NUCLEUS_LITTLE
ASI_AS_IF_USER_PRIMARY	ASI_AS_IF_USER_PRIMARY_LITTLE
ASI_AS_IF_USER_SECONDARY	ASI_AS_IF_USER_SECONDARY_LITTLE
ASI_REAL	ASI_REAL_LITTLE
ASI_REAL_IO	ASI_REAL_IO_LITTLE
ASI_PRIMARY	ASI_PRIMARY_LITTLE

# LDFA / LDDFA / LDQFA

---

## ASIs valid for LDFA and LDQFA

ASI_SECONDARY	ASI_SECONDARY_LITTLE
ASI_PRIMARY_NO_FAULT	ASI_PRIMARY_NO_FAULT_LITTLE
ASI_SECONDARY_NO_FAULT	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 *data\_access\_exception* exception.

---

## ASIs valid for LDDFA

ASI_NUCLEUS	ASI_NUCLEUS_LITTLE
ASI_AS_IF_USER_PRIMARY	ASI_AS_IF_USER_PRIMARY_LITTLE
ASI_AS_IF_USER_SECONDARY	ASI_AS_IF_USER_SECONDARY_LITTLE
ASI_REAL	ASI_REAL_LITTLE
ASI_REAL_IO	ASI_REAL_IO_LITTLE
ASI_PRIMARY	ASI_PRIMARY_LITTLE
ASI_SECONDARY	ASI_SECONDARY_LITTLE
ASI_PRIMARY_NO_FAULT	ASI_PRIMARY_NO_FAULT_LITTLE
ASI_SECONDARY_NO_FAULT	ASI_SECONDARY_NO_FAULT_LITTLE

**Behavior with Partial Store ASIs.** ASIs C0<sub>16</sub>–C5<sub>16</sub> and C8<sub>16</sub>–CD<sub>16</sub> are only defined for use in Partial Store operations (see page 341). 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<sub>16</sub>–C5<sub>16</sub> or C8<sub>16</sub>–CD<sub>16</sub> (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 *data\_access\_exception*, *mem\_address\_not\_aligned*, or *LDDF\_mem\_address\_not\_aligned* exception.
2. If the memory address is correctly aligned, the virtual processor generates a *data\_access\_exception*.

**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 and LDSHORTF instructions; it is distinguished by the ASI used.

## LDFA / LDDFA / LDQFA

*Exceptions*

- illegal\_instruction*
- fp\_disabled*
- LDDF\_mem\_address\_not\_aligned*
- mem\_address\_not\_aligned*
- fp\_exception\_other* (FSR.ftt = invalid\_fp\_register (LDQFA only))
- privileged\_action*
- VA\_watchpoint*
- fast\_data\_access\_MMU\_miss*
- data\_access\_MMU\_miss*
- data\_access\_MMU\_error*

*See Also*


- Load Floating-Point* on page 249
- Block Load* on page 245
- Store Short Floating-Point* on page 344
- Store Floating-Point into Alternate Space* on page 335

# LDFSR - Deprecated

## 7.54 Load Floating-Point State Register

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 <sup>D</sup>	10 0001	0	Load Floating-Point State Register Lower	ld [address], %fsr	<b>C2</b>



**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 upper 32 bits of FSR are unaffected by LDFSR. LDFSR does not alter the *ver*, *ftt*, *qne*, or reserved fields of FSR (see page 61).

**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 120).

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 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.w

# LDFSR - Deprecated

<b>Implementation</b>	LDFSR shares an opcode with the LDXFSR 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 ( $rd > 1$ ) causes an <i>illegal_instruction</i> exception.
<b>Note</b>	


*Exceptions*

- illegal\_instruction*
- fp\_disabled*
- mem\_address\_not\_aligned*
- VA\_watchpoint*
- fast\_data\_access\_MMU\_miss*
- data\_access\_MMU\_miss*
- data\_access\_MMU\_error*

# LDSHORTF

## 7.55 Short Floating-Point Load VIS 1

Instruction	ASI Value	Operation	Assembly Language Syntax		Class
LDSHORTF	D0 <sub>16</sub>	8-bit load from primary address space	ldda	[regaddr] #ASI_FL8_P, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	C3
LDSHORTF	D1 <sub>16</sub>	8-bit load from secondary address space	ldda	[regaddr] #ASI_FL8_S, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	C3
LDSHORTF	D8 <sub>16</sub>	8-bit load from primary address space, little-endian	ldda	[regaddr] #ASI_FL8_PL, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	C3
LDSHORTF	D9 <sub>16</sub>	8-bit load from secondary address space, little-endian	ldda	[regaddr] #ASI_FL8_SL, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	C3
LDSHORTF	D2 <sub>16</sub>	16-bit load from primary address space	ldda	[regaddr] #ASI_FL16_P, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	C3
LDSHORTF	D3 <sub>16</sub>	16-bit load from secondary address space	ldda	[regaddr] #ASI_FL16_S, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	C3
LDSHORTF	DA <sub>16</sub>	16-bit load from primary address space, little-endian	ldda	[regaddr] #ASI_FL16_PL, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	C3
LDSHORTF	DB <sub>16</sub>	16-bit load from secondary address space, little-endian	ldda	[regaddr] #ASI_FL16_SL, freg <sub>rd</sub> ldda [reg_plus_imm] %asi, freg <sub>rd</sub>	C3



*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.

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 8-bit load places the loaded value in the least significant byte of F<sub>D</sub>[rd] and zeroes in the most-significant three bytes of F<sub>D</sub>[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<sub>D</sub>[rd] and zeroes in the more-significant halfword of F<sub>D</sub>[rd]. A 16-bit LDSHORTF from an address that is not halfword-aligned (an odd address) causes a *mem\_address\_not\_aligned* exception.

# LDSHORTF

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 instruction (see *Align Address* on page 147) to assemble or store 64 bits from noncontiguous components.

**Implementation Note** LDSHORTF shares an opcode with the LDBLOCKF and LDDFA instructions; it is distinguished by the ASI used.

In an UltraSPARC Architecture 2005 implementation, these instructions are not implemented in hardware, cause a *data\_access\_exception* exception, and are emulated in software.


*Exceptions*

- VA\_watchpoint*
- data\_access\_exception*
- fast\_data\_access\_MMU\_miss*
- data\_access\_MMU\_miss*

# LDSTUB

## 7.56 Load-Store Unsigned Byte

Instruction	op3	Operation	Assembly Language Syntax	Class
LDSTUB	00 1101	Load-Store Unsigned Byte	ldstub [address], reg <sub>rd</sub>	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 104). 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*
- data\_access\_exception*
- fast\_data\_access\_MMU\_miss*
- data\_access\_MMU\_miss*
- data\_access\_MMU\_error*
- fast\_data\_access\_protection*

# LDSTUBA

## 7.57 Load-Store Unsigned Byte to Alternate Space

Instruction	op3	Operation	Assembly Language Syntax	Class
LDSTUBA <sup>PASI</sup>	01 1101	Load-Store Unsigned Byte into Alternate Space	ldstuba [regaddr] imm_asi, reg_rd ldstuba [reg_plus_imm] %asi, reg_rd	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 and HPSTATE.hpriv = 0), if bit 7 of the ASI is 0, this instruction causes a *privileged\_action* exception. In privileged mode (PSTATE.priv = 1 and HPSTATE.hpriv = 0), 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 *data\_access\_exception* exception.

### ASIs valid for LDSTUBA

ASI_NUCLEUS	ASI_NUCLEUS_LITTLE
ASI_AS_IF_USER_PRIMARY	ASI_AS_IF_USER_PRIMARY_LITTLE
ASI_AS_IF_USER_SECONDARY	ASI_AS_IF_USER_SECONDARY_LITTLE
ASI_REAL	ASI_REAL_LITTLE
ASI_PRIMARY	ASI_PRIMARY_LITTLE
ASI_SECONDARY	ASI_SECONDARY_LITTLE

# LDSTUBA

*Exceptions*

- privileged\_action*
- VA\_watchpoint*
- data\_access\_exception*
- fast\_data\_access\_MMU\_miss*
- data\_access\_MMU\_miss*
- data\_access\_MMU\_error*
- fast\_data\_access\_protection*

# LDTXA

## 7.58 Load Integer Twin Extended Word from Alternate Space VIS 2+

The LDTXA instructions are not guaranteed to be implemented on all UltraSPARC Architecture implementations. Therefore, they should only be used in platform-specific dynamically-linked libraries, in hyperprivileged software, or in software created by a runtime code generator that is aware of the specific virtual processor implementation on which it is executing.

Instruction	ASI Value	Operation	Assembly Language Syntax †	Class
LDTXA <sup>N</sup>	22 <sub>16</sub>	Load Integer Twin Extended Word, as if user (nonprivileged), Primary address space	<code>ldtxa [regaddr] #ASI_LDTX_AIUP, regrd</code>	<b>N1</b>
	23 <sub>16</sub>	Load Integer Twin Extended Word, as if user (nonprivileged), Secondary address space	<code>ldtxa [regaddr] #ASI_LDTX_AIUS, regrd</code>	<b>N1</b>
	26 <sub>16</sub>	Load Integer Twin Extended Word, real address	<code>ldtxa [regaddr] #ASI_LDTX_REAL, regrd</code>	<b>N1</b>
	27 <sub>16</sub>	Load Integer Twin Extended Word, nucleus context	<code>ldtxa [regaddr] #ASI_LDTX_N, regrd</code>	<b>N1</b>
	2A <sub>16</sub>	Load Integer Twin Extended Word, as if user (nonprivileged), Primary address space, little endian	<code>ldtxa [regaddr] #ASI_LDTX_AIUP_L, regrd</code>	<b>N1</b>
	2B <sub>16</sub>	Load Integer Twin Extended Word, as if user (nonprivileged), Secondary address space, little endian	<code>ldtxa [regaddr] #ASI_LDTX_AIUS_L, regrd</code>	<b>N1</b>
	2E <sub>16</sub>	Load Integer Twin Extended Word, real address, little endian	<code>ldtxa [regaddr] #ASI_LDTX_REAL_L, regrd</code>	<b>N1</b>
	2F <sub>16</sub>	Load Integer Twin Extended Word, nucleus context, little-endian	<code>ldtxa [regaddr] #ASI_LDTX_NL, regrd</code>	<b>N1</b>
LDTXA <sup>N</sup>	E2 <sub>16</sub>	Load Integer Twin Extended Word, Primary address space	<code>ldtxa [regaddr] #ASI_LDTX_P, regrd</code>	<b>N1</b>
	E3 <sub>16</sub>	Load Integer Twin Extended Word, Secondary address space	<code>ldtxa [regaddr] #ASI_LDTX_S, regrd</code>	<b>N1</b>
	EA <sub>16</sub>	Load Integer Twin Extended Word, Primary address space, little endian	<code>ldtxa [regaddr] #ASI_LDTX_PL, regrd</code>	<b>N1</b>
	EB <sub>16</sub>	Load Integer Twin Extended Word, Secondary address space, little-endian	<code>ldtxa [regaddr] #ASI_LDTX_SL, regrd</code>	<b>N1</b>

† 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.

# LDTXA



*Description* ASIs  $26_{16}$ ,  $2E_{16}$ ,  $E2_{16}$ ,  $E3_{16}$ ,  $F0_{16}$ , and  $F1_{16}$  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].

ASIs  $E2_{16}$ ,  $E3_{16}$ ,  $F0_{16}$ , and  $F1_{16}$  perform an access using a virtual address, while ASIs  $26_{16}$  and  $2E_{16}$  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* and *PA\_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 a *data\_access\_exception* 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.

**Programming Note** In hyperprivileged mode, an access to ASI  $E2_{16}$ ,  $E3_{16}$ ,  $F0_{16}$ , or  $F1_{16}$  is performed using physical (not virtual) addressing.

The virtual processor MMU does not provide virtual-to-real translation for ASIs  $26_{16}$  and  $2E_{16}$ ; 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.

# LDTXA

A *mem\_address\_not\_aligned* trap is taken if the access is not aligned on a 128-byte boundary.

<b>Implementation</b>	LDTXA shares an opcode with the “i = 0” variant of the (deprecated) LDTWA instruction. See <i>Load Integer Twin Word from Alternate Space</i> on page 268.
<b>Note</b>	

*Exceptions*

- illegal\_instruction*
- mem\_address\_not\_aligned*
- privileged\_action*
- VA\_watchpoint* (impl. dep. #413-S10)
- data\_access\_exception*
- fast\_data\_access\_MMU\_miss*
- data\_access\_MMU\_miss*
- data\_access\_MMU\_error*
- PA\_watchpoint* (impl. dep. #413-S10)
- data\_access\_error*


# LDTW (Deprecated)

## 7.59 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 <sup>D</sup>	00 0011	Load Integer Twin Word	ldtw [address], reg <sub>rd</sub>	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 104). 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.

## LDTW (Deprecated)

<b>SPARC V9 Compatibility Note</b>	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.

A successful LDTW instruction operates atomically.

*Exceptions*

- unimplemented\_LDTW*
- illegal\_instruction*
- mem\_address\_not\_aligned*
- VA\_watchpoint*
- data\_access\_exception*
- fast\_data\_access\_MMU\_miss*
- data\_access\_MMU\_miss*
- data\_access\_MMU\_error*

*See Also*

- LDW/LDX on page 240
- STTW on page 346

# LDTWA (Deprecated)


## 7.60 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 <sup>D, P<sub>ASI</sub></sup>	01 0011	Load Integer Twin Word from Alternate Space	ldtwa [ <i>regaddr</i> ] <i>imm_asi</i> , <i>reg<sub>rd</sub></i> ldtwa [ <i>reg_plus_imm</i> ] % <i>asi</i> , <i>reg<sub>rd</sub></i>	‡

† 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 (00<sub>16</sub>-7F<sub>16</sub>); **D2** for unrestricted ASIs (80<sub>16</sub>-FF<sub>16</sub>)



**Description** The load integer twin word from alternate space instruction (LDTWA) 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 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.

# LDTWA (Deprecated)

**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.

<b>Programming Note</b>	LDTWA 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.  If LDTWA is emulated in software, an LDXA instruction should be used for the memory access in the emulation code in order to preserve atomicity.
-------------------------	--

<b>SPARC V9 Compatibility Note</b>	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 and *HPSTATE.hpriv* = 0), if bit 7 of the ASI is 0, these instructions cause a *privileged\_action* exception. In privileged mode (*PSTATE.priv* = 1 and *HPSTATE.hpriv* = 0), if the ASI is in the range  $30_{16}$  to  $7F_{16}$ , 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 *data\_access\_exception* exception (impl. dep. #300-U4-Cs10).

---

#### ASIs valid for LDTWA

ASI_NUCLEUS	ASI_NUCLEUS_LITTLE
ASI_AS_IF_USER_PRIMARY	ASI_AS_IF_USER_PRIMARY_LITTLE
ASI_AS_IF_USER_SECONDARY	ASI_AS_IF_USER_SECONDARY_LITTLE
ASI_REAL	ASI_REAL_LITTLE
ASI_REAL_IO	ASI_REAL_IO_LITTLE
$22_{16}\ddagger$ (ASI_LDTX_AIUP)	$2A_{16}\ddagger$ (ASI_LDTD_AIUP_L)
$23_{16}\ddagger$ (ASI_LDTX_AIUS)	$2B_{16}\ddagger$ (ASI_LDTX_AIUS_L)
$24_{16}\ddagger$ (aliased to $27_{16}$ , ASI_LDTX_N)	$2C_{16}\ddagger$ (aliased to $2F_{16}$ , ASI_LDTX_NL)
$26_{16}\ddagger$ (ASI_LDTX_REAL)	$2E_{16}\ddagger$ (ASI_LDTX_REAL_L)
$27_{16}\ddagger$ (ASI_LDTX_N)	$2F_{16}\ddagger$ (ASI_LDTX_NL)

# LDTWA (Deprecated)

---

## ASIs valid for LDTWA

ASI_PRIMARY	ASI_PRIMARY_LITTLE
ASI_SECONDARY	ASI_SECONDARY_LITTLE
ASI_PRIMARY_NO_FAULT	ASI_PRIMARY_NO_FAULT_LITTLE
ASI_SECONDARY_NO_FAULT	ASI_SECONDARY_NO_FAULT_LITTLE

E2 <sub>16</sub> ‡ (ASI_LDTX_P)	EA <sub>16</sub> ‡ (ASI_LDTX_PL)
E3 <sub>16</sub> ‡ (ASI_LDTX_S)	EB <sub>16</sub> ‡ (ASI_LDTX_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 263.

If this ASI is used with the opcode for LDTWA and  $i = 1$ , behavior is undefined.

**Programming Note** Nontranslating ASIs (see page 407) 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 *data\_access\_exception* exception (impl. dep. #300-U4-Cs10).

**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 263.

*Exceptions*

- unimplemented\_LDTW illegal\_instruction*
- mem\_address\_not\_aligned*
- privileged\_action*
- VA\_watchpoint*
- data\_access\_exception*
- fast\_data\_access\_MMU\_miss*
- data\_access\_MMU\_miss*
- data\_access\_MMU\_error*


*See Also*

- LDWA/LDXA on page 242
- STTWA on page 348

# MEMBAR

## 7.61 Memory Barrier

Instruction	op3	Operation	Assembly Language Syntax	Class
MEMBAR	10 1000	Memory Barrier	membar <i>membar_mask</i>	<b>A1</b>



*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 load instructions LDSTUB[A], SWAP[A], CASA, and CASX[A]) and stores (including store 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 7-7 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.

# MEMBAR

TABLE 7-7 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. Equivalent to the deprecated STBAR instruction.
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.

The `cmask` field is encoded in bits 6 through 4 of the instruction. Bits in the `cmask` field, described in TABLE 7-8, 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 7-8 MEMBAR cmask Encodings

Mask Bit	Function	Assembly Language Name	Description
cmask{2}	Synchronization barrier	#Sync	All operations (including nonmemory reference operations) appearing prior to the MEMBAR must have been performed and the effects of any exceptions be visible before any instruction after the MEMBAR may be initiated.
cmask{1}	Memory issue barrier	#MemIssue	All memory reference operations appearing prior to the MEMBAR must have been performed before any memory operation after the MEMBAR may be initiated.
cmask{0}	Lookaside barrier	#Lookaside	A store appearing prior to the MEMBAR must complete before any load following the MEMBAR referencing the same address can be initiated.

A MEMBAR instruction with both `mmask` = 0 and `cmask` = 0 is functionally a NOP.

For information on the use of MEMBAR, see *Memory Ordering and Synchronization* on page 401 and *Programming with the Memory Models* contained in the separate volume *UltraSPARC Architecture Application Notes*. For additional information about the memory models themselves, see Chapter 9, *Memory*.

# MEMBAR

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  $m\text{mask} = 8_{16}$  and  $c\text{mask} = 0_{16}$  (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 and STBAR<sup>D</sup>; it is distinguished by  $rs1 = 15$ ,  $rd = 0$ ,  $i = 1$ , and bit 12 = 0.

## 7.61.1 Memory Synchronization

The UltraSPARC 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 UltraSPARC Architecture implementation may define the operation of each MEMBAR variant in any manner that provides the required semantics.

**Implementation Note** For an UltraSPARC Architecture virtual processor that only provides TSO memory ordering semantics, three of the ordering MEMBARs would normally be implemented as NOPs. TABLE 7-9 shows an acceptable implementation of MEMBAR for a TSO-only UltraSPARC Architecture implementation.

**TABLE 7-9** MEMBAR Semantics for TSO-only implementation

MEMBAR variant	Preferred Implementation
#StoreStore, STBAR	NOP
#LoadStore	NOP
#StoreLoad	#Sync
#LoadLoad	NOP
#Sync	#Sync
#MemIssue	#Sync
#Lookaside	#Sync

If an UltraSPARC 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.

# MEMBAR

## 7.61.2 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.

During synchronization, if a disrupting trap condition due to a hardware error is detected and external interrupts are enabled, the disrupting trap will occur before the instruction after the synchronizing instruction is executed. In this case, the PC value saved in TPC during trap entry will be the address of the instruction after the synchronizing instruction.

<b>Programming Note</b>	Completion of a MEMBAR #Sync instruction does <i>not</i> guarantee that data previously stored has been written all the way out to external memory (that is, that cache writebacks to external memory have completed). Software cannot rely on that behavior. There is no mechanism in the UltraSPARC Architecture that allows software to wait for all previous stores to be written to external memory (that is, for cache writebacks to completely drain).
-------------------------	---

## 7.61.3 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 398.

*Exceptions*      *illegal\_instruction*

# MOVcc

## 7.62 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	mova <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>
MOVN	10 1100	0000	Move Never	0	movn <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>
MOVNE	10 1100	1001	Move if Not Equal	<b>not</b> Z	movne <sup>†</sup> <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>
MOVE	10 1100	0001	Move if Equal	Z	move <sup>‡</sup> <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>
MOVG	10 1100	1010	Move if Greater	<b>not</b> (Z <b>or</b> N <b>xor</b> V)	movg <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>
MOVLE	10 1100	0010	Move if Less or Equal	Z <b>or</b> (N <b>xor</b> V)	movle <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>
MOVGE	10 1100	1011	Move if Greater or Equal	<b>not</b> (N <b>xor</b> V)	movge <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>
MOVL	10 1100	0011	Move if Less	N <b>xor</b> V	movl <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>
MOVGU	10 1100	1100	Move if Greater, Unsigned	<b>not</b> (C <b>or</b> Z)	movgu <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>
MOVLEU	10 1100	0100	Move if Less or Equal, Unsigned	(C <b>or</b> Z)	movleu <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>
MOVCC	10 1100	1101	Move if Carry Clear (Greater or Equal, Unsigned)	<b>not</b> C	movcc <sup>◇</sup> <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>
MOVCS	10 1100	0101	Move if Carry Set (Less than, Unsigned)	C	movcs <sup>∇</sup> <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>
MOVPOS	10 1100	1110	Move if Positive	<b>not</b> N	movpos <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>
MOVNEG	10 1100	0110	Move if Negative	N	movneg <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>
MOVVC	10 1100	1111	Move if Overflow Clear	<b>not</b> V	movvc <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>
MOVVS	10 1100	0111	Move if Overflow Set	V	movvs <i>i_or_x_cc, reg_or_imm11, reg_rd</i>	<b>A1</b>

<sup>†</sup> *synonym: movnz*

<sup>‡</sup> *synonym: movz*

<sup>◇</sup> *synonym: movgeu*

<sup>∇</sup> *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	movfa %fccn, reg_or_imm11, regrd	A1
MOVFN	10 1100	0000	Move Never	0	movfn %fccn, reg_or_imm11, regrd	A1
MOVFU	10 1100	0111	Move if Unordered	U	movfu %fccn, reg_or_imm11, regrd	A1
MOVFG	10 1100	0110	Move if Greater	G	movfg %fccn, reg_or_imm11, regrd	A1
MOVFUG	10 1100	0101	Move if Unordered or Greater	G or U	movfug %fccn, reg_or_imm11, regrd	A1
MOVFL	10 1100	0100	Move if Less	L	movfl %fccn, reg_or_imm11, regrd	A1
MOVFUL	10 1100	0011	Move if Unordered or Less	L or U	movful %fccn, reg_or_imm11, regrd	A1
MOVFLG	10 1100	0010	Move if Less or Greater	L or G	movflg %fccn, reg_or_imm11, regrd	A1
MOVFNE	10 1100	0001	Move if Not Equal	L or G or U	movfne <sup>†</sup> %fccn, reg_or_imm11, regrd	A1
MOVFE	10 1100	1001	Move if Equal	E	movfe <sup>‡</sup> %fccn, reg_or_imm11, regrd	A1
MOVFUE	10 1100	1010	Move if Unordered or Equal	E or U	movfue %fccn, reg_or_imm11, regrd	A1
MOVFGE	10 1100	1011	Move if Greater or Equal	E or G	movfge %fccn, reg_or_imm11, regrd	A1
MOVFUGE	10 1100	1100	Move if Unordered or Greater or Equal	E or G or U	movfuge %fccn, reg_or_imm11, regrd	A1
MOVFLE	10 1100	1101	Move if Less or Equal	E or L	movfle %fccn, reg_or_imm11, regrd	A1
MOVFULE	10 1100	1110	Move if Unordered or Less or Equal	E or L or U	movfule %fccn, reg_or_imm11, regrd	A1
MOVFO	10 1100	1111	Move if Ordered	E or L or G	movfo %fccn, reg_or_imm11, regrd	A1

<sup>†</sup> synonym: movnz      <sup>‡</sup> 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.



# MOVcc

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)

## 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.

# MOVcc

If 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

## 7.63 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<sup>†</sup> reg<sub>rs1</sub>, reg_or_imm10, reg<sub>rd</sub></code>	<b>A1</b>
MOVrLEZ	10 1111	010	Move if Register Less Than or Equal to Zero	$R[rs1] \leq 0$	<code>movrLEZ reg<sub>rs1</sub>, reg_or_imm10, reg<sub>rd</sub></code>	<b>A1</b>
MOVrLZ	10 1111	011	Move if Register Less Than Zero	$R[rs1] < 0$	<code>movrLZ reg<sub>rs1</sub>, reg_or_imm10, reg<sub>rd</sub></code>	<b>A1</b>
—	10 1111	100	<i>Reserved (illegal_instruction)</i>			—
MOVrNZ	10 1111	101	Move if Register Not Zero	$R[rs1] \neq 0$	<code>movrNZ<sup>‡</sup> reg<sub>rs1</sub>, reg_or_imm10, reg<sub>rd</sub></code>	<b>A1</b>
MOVrGZ	10 1111	110	Move if Register Greater Than Zero	$R[rs1] > 0$	<code>movrGZ reg<sub>rs1</sub>, reg_or_imm10, reg<sub>rd</sub></code>	<b>A1</b>
MOVrGEZ	10 1111	111	Move if Register Greater Than or Equal to Zero	$R[rs1] \geq 0$	<code>movrGEZ reg<sub>rs1</sub>, reg_or_imm10, reg<sub>rd</sub></code>	<b>A1</b>

<sup>†</sup> *synonym: movre*      <sup>‡</sup> *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.

# MOVr

**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.

<u>Move</u>	<u>Test</u>
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<sub>2</sub> or 100<sub>2</sub> causes an *illegal\_instruction* exception.


*Exceptions*     *illegal\_instruction*

# MULScc - Deprecated

## 7.64 Multiply Step

The MULScc instruction is deprecated and should not be used in new software. The MULX instruction should be used instead.

Opcode	op3	Operation	Assembly Language Syntax	Class
MULScc <sup>D</sup>	10 0100	Multiply Step and modify cc's	<code>mul<sub>scc</sub> reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></code>	<b>Y3</b>



*Description* MULScc treats the less-significant 32 bits of R[rs1] and the less-significant 32 bits of the Y register as a single 64-bit, right-shiftable doubleword register. The least significant bit of R[rs1] is treated as if it were adjacent to bit 31 of the Y register. The MULScc instruction performs an addition operation, based on the least significant bit of Y.

Multiplication assumes that the Y register initially contains the multiplier, R[rs1] contains the most significant bits of the product, and R[rs2] contains the multiplicand. Upon completion of the multiplication, the Y register contains the least significant bits of the product.

**Note** | In a standard MULScc instruction,  $rs1 = rd$ .

MULScc operates as follows:

1. If  $i = 0$ , the multiplicand is R[rs2]; if  $i = 1$ , the multiplicand is **sign\_ext**(simm13).
2. A 32-bit value is computed by shifting the value from R[rs1] right by one bit with “CCR.icc.n xor CCR.icc.v” replacing bit 31 of R[rs1]. (This is the proper sign for the previous partial product.)
3. If the least significant bit of  $Y = 1$ , the shifted value from step (2) and the multiplicand are added. If the least significant bit of the  $Y = 0$ , then 0 is added to the shifted value from step (2).

# MULScc - Deprecated

4. MULScc writes the following result values:

Register field	Value written by MULScc
CCR.icc	updated according to the result of the addition in step (3) above
R[rd]{63:32}	<i>undefined</i>
R[rd]{31:0}	the least-significant 32 bits of the sum from step (3) above
Y	the previous value of the Y register, shifted right by one bit, with Y{31} replaced by the value of R[rs1]{0} prior to shifting in step (2)
CCR.xcc	<i>undefined</i>

5. The Y register is shifted right by one bit, with the least significant bit of the unshifted R[rs1] replacing bit 31 of Y.


An attempt to execute a MULScc instruction when  $i = 0$  and instruction bits 12:5 are nonzero causes an *illegal\_instruction* exception.

*Exceptions*      *illegal\_instruction*

# MULX / SDIVX / UDIVX

## 7.65 Multiply and Divide (64-bit)

Instruction	op3	Operation	Assembly Language		Class
MULX	00 1001	Multiply (signed or unsigned)	mulx	<i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i>	A1
SDIVX	10 1101	Signed Divide	sdivx	<i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i>	A1
UDIVX	00 1101	Unsigned Divide	udivx	<i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i>	A1



*Description* MULX computes “R[rs1] × R[rs2]” if *i* = 0 or “R[rs1] × **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] ÷ R[rs2]” if *i* = 0 or “R[rs1] ÷ **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:

$$8000\ 0000\ 0000\ 0000_{16} \div \text{FFFF FFFF FFFF FFFF}_{16} = 8000\ 0000\ 0000\ 0000_{16}.$$

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*

# NOP

## 7.66 No Operation

Instruction	op2	Operation	Assembly Language Syntax	Class
NOP	100	No Operation	<code>nop</code>	<b>A1</b>



*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 only one guaranteed to be implemented efficiently across all implementations.

*Exceptions* None

# NORMALW

## 7.67 NORMALW

Instruction	Operation	Assembly Language Syntax	Class
NORMALW <sup>P</sup>	“Other” register windows become “normal” register windows	normalw	C1



**Description** NORMALW<sup>P</sup> 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.

This instruction allows window manipulations to be atomic, without the value of *N\_REG\_WINDOWS* being visible to privileged software and without an assumption that *N\_REG\_WINDOWS* is constant (since hyperprivileged software can migrate a thread among virtual processors, across which *N\_REG\_WINDOWS* may vary).

In an UltraSPARC Architecture 2005 implementation, this instruction is not implemented in hardware, causes an *illegal\_instruction* exception, and is emulated in software.


**Exceptions** *illegal\_instruction* (not implemented in hardware in UltraSPARC Architecture 2005)

**See Also** ALLCLEAN on page 148  
INVALW on page 238  
OTHERW on page 287  
RESTORED on page 307  
SAVED on page 315

# OR

## 7.68 OR Logical Operation

Instruction	op3	Operation	Assembly Language Syntax	Class
OR	00 0010	Inclusive <b>or</b>	<code>or</code> $reg_{rs1}, reg_{or\_imm}, reg_{rd}$	<b>A1</b>
ORcc	01 0010	Inclusive <b>or</b> and modify cc's	<code>orcc</code> $reg_{rs1}, reg_{or\_imm}, reg_{rd}$	<b>A1</b>
ORN	00 0110	Inclusive <b>or not</b>	<code>orn</code> $reg_{rs1}, reg_{or\_imm}, reg_{rd}$	<b>A1</b>
ORNcc	01 0110	Inclusive <b>or not</b> and modify cc's	<code>orncc</code> $reg_{rs1}, reg_{or\_imm}, reg_{rd}$	<b>A1</b>



*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

## 7.69 OTHERW

Instruction	Operation	Assembly Language Syntax	Class
OTHERW <sup>P</sup>	“Normal” register windows become “other” register windows	otherw	C1



*Description* OTHERW<sup>P</sup> 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.

This instruction allows window manipulations to be atomic, without the value of *N\_REG\_WINDOWS* being visible to privileged software and without an assumption that *N\_REG\_WINDOWS* is constant (since hyperprivileged software can migrate a thread among virtual processors, across which *N\_REG\_WINDOWS* may vary).

In an UltraSPARC Architecture 2005 implementation, this instruction is not implemented in hardware, causes an *illegal\_instruction* exception, and is emulated in software.

*Exceptions* *illegal\_instruction* (not implemented in hardware in UltraSPARC Architecture 2005)

*See Also* ALLCLEAN on page 148  
INVALW on page 238  
NORMALW on page 285  
RESTORED on page 307  
SAVED on page 315

# PDIST

## 7.70 Pixel Component Distance VIS 1

Instruction	opf	Operation	Assembly Language Syntax	Class
PDIST	0 0011 1110	Distance between eight 8-bit components, with accumulation	<code>pdist freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code>	<b>C3</b>



*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 uses  $F_D[rd]$  as both a source and a destination register. Typically, PDIST is used for motion estimation in video compression algorithms.


In an UltraSPARC Architecture 2005 implementation, this instruction is not implemented in hardware, causes an *illegal\_instruction* exception, and is emulated in software.

*Exceptions* *illegal\_instruction*

# POPC

## 7.71 Population Count

Instruction	op3	Operation	Assembly Language Syntax	Class
POPC	10 1110	Population Count	popc <i>reg_or_imm</i> , <i>reg_rd</i>	<b>D3</b>



*Description* POPC counts the number of one bits in R[rs2] if *i* = 0, or the number of one 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” in a register. A ‘C’-language program illustrating how POPC can be used for this purpose follows:

```
int ffs(zz) /* finds first 1 bit, counting from the LSB */
unsigned zz;
{
    return popc ( zz ^ (~ (-zz)) ); /* for nonzero zz */
}
```

Inline assembly language code for `ffs()` is:

```
neg    %IN, %M_IN      ! -zz (2's complement)
xnor   %IN, %M_IN, %TEMP ! ^ ~ -zz (exclusive nor)
popc   %TEMP, %RESULT  ! result = popc(zz ^ ~ -zz)
movrz  %IN, %g0, %RESULT ! %RESULT should be 0 for %IN=0
```

where *IN*, *M\_IN*, *TEMP*, and *RESULT* are integer registers.

Example computation:

```
IN = ...00101000 !1st '1' bit from right is
-IN = ...11011000 ! bit 3 (4th bit)
~ -IN = ...00100111
IN ^ ~ -IN = ...00001111
popc(IN ^ ~ -IN) = 4
```

# POPC

**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    %IN, %DEST
cmp     %IN, -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 *IN*, *TEMP*, and *DEST* are integer registers.

In an UltraSPARC Architecture 2005 implementation, this instruction is not implemented in hardware, causes an *illegal\_instruction* exception, and is emulated in software.

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*

# PREFETCH


## 7.72 Prefetch

Instruction	op3	Operation	Assembly Language Syntax	Class
PREFETCH	10 1101	Prefetch Data	<code>prefetch [address], prefetch_fcn</code>	<b>A1</b>
PREFETCHA <sup>PASI</sup>	11 1101	Prefetch Data from Alternate Space	<code>prefetcha [regaddr] imm_asi, prefetch_fcn</code> <code>prefetcha [reg_plus_imm] %asi, prefetch_fcn</code>	<b>A1</b>

### PREFETCH



### PREFETCHA



**TABLE 7-10** Prefetch Variants, by Function Code

fcn	Prefetch Variant
0	(Weak) Prefetch for several reads
1	(Weak) Prefetch for one read
2	(Weak) Prefetch for several writes and possibly reads
3	(Weak) Prefetch for one write
4	Prefetch page
5–15 (05 <sub>16</sub> –0F <sub>16</sub> )	Reserved ( <i>illegal_instruction</i> )
16 (10 <sub>16</sub> )	Implementation dependent (NOP if not implemented)
17 (11 <sub>16</sub> )	Prefetch to nearest unified cache
18–19 (12 <sub>16</sub> –13 <sub>16</sub> )	Implementation dependent (NOP if not implemented)
20 (14 <sub>16</sub> )	Strong Prefetch for several reads
21 (15 <sub>16</sub> )	Strong Prefetch for one read
22 (16 <sub>16</sub> )	Strong Prefetch for several writes and possibly reads
23 (17 <sub>16</sub> )	Strong Prefetch for one write
24–31 (18 <sub>16</sub> –1F <sub>16</sub> )	Implementation dependent (NOP if not implemented)

# PREFETCH

*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 movement of a block of data containing the addressed byte from memory toward the virtual processor or creates an address mapping.

**Implementation Note** A PREFETCH[A] instruction may be used by software to:

- prefetch a cache line into a cache
- prefetch a valid address translation into a TLB
- invalidate a cache line that may have caused a correctable error during a load instruction.

If  $i = 0$ , the effective address operand for the PREFETCH instruction is “R[rs1] + R[rs2]”; if  $i = 1$ , it is “R[rs1] + **sign\_ext** (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 (including a possible hardware tablewalk to load a TLB entry), 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.

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 Note** Software may prefetch 64 bytes beginning at an arbitrary address 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

## 7.72.1 Exceptions

Prefetch instructions PREFETCH and PREFETCHA generate exceptions under the conditions detailed in TABLE 7-11. Only the implementation-dependent prefetch variants (see TABLE 7-10) may generate an exception under conditions not listed in this table; the predefined variants only generate the exceptions listed here.

**TABLE 7-11** Behavior of PREFETCH[A] Instructions Under Exceptional Conditions (1 of 2)

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 <sub>16</sub> ..7F <sub>16</sub> , while in nonprivileged mode ( <i>privileged_action</i> condition)	executes as NOP
any	PREFETCHA	reference to an ASI in range 30 <sub>16</sub> ..7F <sub>16</sub> , while in privileged mode ( <i>privileged_action</i> condition)	executes as NOP
0-3 (weak)	PREFETCH[A]	condition detected for MMU miss ( <i>data_access_MMU_miss</i> or <i>fast_data_access_MMU_miss</i> )	executes as NOP
0-3 (weak)	PREFETCH[A]	condition detected for <i>data_access_MMU_error</i>	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 ((TTE.cp = 0) or ((fcn = 0) and TTE.cv = 0)), or (TTE.e = 1)	executes as NOP
4, 20-23 (strong)	PREFETCH[A]	prefetching the requested data would be a very time-consuming operation (condition detected for <i>data_access_MMU_miss</i> )	executes as NOP
4, 20-23 (strong)	PREFETCH[A]	prefetching the requested data would be a time-consuming operation (condition detected for <i>fast_data_access_MMU_miss</i> )	executes as NOP
4, 20-23 (strong)	PREFETCH[A]	condition detected for <i>data_access_MMU_error</i> , <i>hw_corrected_error</i> , or <i>sw_recoverable_error</i> (impl. dep. #_?_)	<i>data_access_MMU_error</i> , <i>hw_corrected_error</i> , or <i>sw_recoverable_error</i>

# PREFETCH

**TABLE 7-11** Behavior of PREFETCH[A] Instructions Under Exceptional Conditions (2 of 2)

fcn	Instruction	Condition	Result
5-15 (05 <sub>16</sub> -0F <sub>16</sub> )	PREFETCH[A]	(always)	<i>illegal_instruction</i>
16-31 (18 <sub>16</sub> -1F <sub>16</sub> )	PREFETCH[A]	variant unimplemented	executes as NOP

---

**ASIs valid for PREFETCHA (all others are invalid)**

ASI_NUCLEUS	ASI_NUCLEUS_LITTLE
ASI_AS_IF_USER_PRIMARY	ASI_AS_IF_USER_PRIMARY_LITTLE
ASI_AS_IF_USER_SECONDARY	ASI_AS_IF_USER_SECONDARY_LITTLE
ASI_PRIMARY	ASI_PRIMARY_LITTLE
ASI_SECONDARY	ASI_SECONDARY_LITTLE
ASI_PRIMARY_NO_FAULT	ASI_PRIMARY_NO_FAULT_LITTLE
ASI_SECONDARY_NO_FAULT	ASI_SECONDARY_NO_FAULT_LITTLE
ASI_REAL	ASI_REAL_LITTLE

## 7.72.2 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) it’s willing for the underlying hardware to 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 requests 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<sup>1</sup> 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

<sup>1</sup>. such as a *fast\_data\_access\_MMU\_miss* trap, plus subsequently filling the cache line at the requested address

# PREFETCH

may pay the cost of performing the time-consuming operation so it can finish initiating the requested data movement. Behavioral differences among loads and prefetches are compared in TABLE 7-12.

**TABLE 7-12** Comparative Behavior of Load and Weak Prefetch Operations

Condition	Behavior	
	Load	Prefetch
On a $\mu$ TLB miss, is an MMU access performed?	Yes	Yes
Upon detection of <i>fast_data_access_MMU_miss</i> exception...	Traps	NOP‡
Upon detection of <i>privileged_action</i> , <i>data_access_exception</i> , <i>data_access_protection</i> , <i>PA_watchpoint</i> , or <i>VA_watchpoint</i> exception...	Traps	NOP‡
If page table entry has $cp = 0$ , $e = 1$ , and $cv = 0$ for Prefetch for Several Reads	Traps	NOP‡
If page table entry has $nfo = 1$ for a non-NoFault access...	Traps	NOP‡
If page table entry has $w = 0$ for any prefetch for write access ( $fcn = 2, 3, 22$ , or $23$ )...	Traps	NOP‡
Upon detection of fatal error or disrupting error conditions...	Traps	Traps
Instruction blocks until cache line filled?	Yes	No

## 7.72.3 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 UltraSPARC Architecture 2005.

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 UltraSPARC 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).

### 7.72.3.1 Prefetch for Several Reads ( $fcn = 0, 20(14_{16})$ )

The intent of these variants is to cause movement of data into the cache nearest the virtual processor.

# PREFETCH

There are Weak and Strong versions of this prefetch variant;  $\text{fcn} = 0$  is Weak and  $\text{fcn} = 20$  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.

## 7.72.3.2 Prefetch for One Read ( $\text{fcn} = 1, 21(15_{16})$ )

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;  $\text{fcn} = 1$  is Weak and  $\text{fcn} = 21$  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.

## 7.72.3.3 Prefetch for Several Writes (and Possibly Reads) ( $\text{fcn} = 2, 22(16_{16})$ )

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;  $\text{fcn} = 2$  is Weak and  $\text{fcn} = 22$  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.

# PREFETCH

## 7.72.3.4 Prefetch for One Write (fcn = 3, 23(17<sub>16</sub>))

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; fcn = 3 is Weak and fcn = 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.

## 7.72.3.5 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.

## 7.72.4 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.

## 7.72.5 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.

# PREFETCH

**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 or hyperprivileged code should be reasonable (for example, in handling ECC errors, no page prefetching is allowed within code that handles page faults). The benefits of prefetching should be available to most privileged code.

**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*  
                      *data\_access\_MMU\_error*

# RDasr

## 7.73 Read Ancillary State Register

Instruction	rs1	Operation	Assembly Language Syntax	Class
RDY <sup>D</sup>	0	Read Y register ( <i>deprecated</i> )	rd %y, reg <sub>rd</sub>	<b>C2</b>
—	1	<i>Reserved</i>		
RDCCR	2	Read Condition Codes register (CCR)	rd %ccr, reg <sub>rd</sub>	<b>A1</b>
RDASI	3	Read ASI register	rd %asi, reg <sub>rd</sub>	<b>A1</b>
RD <sup>P</sup> TICK <sup>npt</sup>	4	Read TICK register	rd %tick, reg <sub>rd</sub>	<b>A1</b>
RDPC	5	Read Program Counter (PC)	rd %pc, reg <sub>rd</sub>	<b>B2</b>
RDFPRS	6	Read Floating-Point Registers Status (FPRS) register	rd %fprs, reg <sub>rd</sub>	<b>A1</b>
—	7–14	<i>Reserved</i>		
See text	15	STBAR, MEMBAR or <i>Reserved</i> ; see text		
RDPCR <sup>P</sup>	16	Read Performance Control registers (PCR)	rd %pcr, reg <sub>rd</sub>	<b>A1</b>
RD <sup>P</sup> PIC <sup>pic</sup>	17	Read Performance Instrumentation Counters register (PIC)	rd %pic, reg <sub>rd</sub>	<b>A1</b>
—	18	<i>Reserved</i> (impl. dep. #8-V8-Cs20, 9-V8-Cs20)		
RDGSR	19	Read General Status register (GSR)	rd %gsr, reg <sub>rd</sub>	<b>A1</b>
—	20–21	<i>Reserved</i> (impl. dep. #8-V8-Cs20, 9-V8-Cs20)		
RDSOFTINT <sup>P</sup>	22	Read per-virtual processor Soft Interrupt register (SOFTINT)	rd %softint, reg <sub>rd</sub>	<b>N2</b>
RD <sup>P</sup> TICK_CM <sup>P</sup> R <sup>P</sup>	23	Read Tick Compare register (TICK_CM <sup>P</sup> R <sup>P</sup> )	rd %tick_cmpr, reg <sub>rd</sub>	<b>N2</b>
RD <sup>P</sup> STICK <sup>npt</sup>	24	Read System Tick Register (STICK)	rd %sys_tick, reg <sub>rd</sub>	<b>N2</b>
RD <sup>P</sup> STICK_CM <sup>P</sup> R <sup>P</sup>	25	Read System Tick Compare register (STICK_CM <sup>P</sup> R <sup>P</sup> )	rd %sys_tick_cmpr, reg <sub>rd</sub>	<b>N2</b>
—	26–27	<i>Reserved</i> (impl. dep. #8-V8-Cs20, 9-V8-Cs20)		
—	28–31	Implementation dependent (impl. dep. #8-V8-Cs20, 9-V8-Cs20)		



# 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 FPods 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, 7–14, 18, 20–21, and 26–27.

**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 or hyperprivileged (impl. dep. #9-V8-Cs20), and
- whether an attempt to execute the instruction causes an *illegal\_instruction* exception.

**Implementation Note** | See the section “Read/Write Ancillary State Registers (ASRs)” in *Extending the UltraSPARC Architecture*, contained in the separate volume *UltraSPARC Architecture Application Notes*, for a discussion of extending the SPARC V9 instruction set using read/write ASR instructions.

**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 UltraSPARC Architecture, since the PSR, WIM, and TBR registers do not exist.

See *Ancillary State Registers* on page 70 for more detailed information regarding ASR registers.

# RDasr

**Exceptions.** An attempt to execute a RDasr instruction when any of the following conditions are true causes an *illegal\_instruction* exception:

- $rs1 = 15$  and  $rd \neq 0$  (reserved for future versions of the architecture)
- $rs1 = 1, 7-14, 18, 20-21, \text{ or } 26-27$  (reserved for future versions of the architecture)
- instruction bits 13:0 are nonzero

An attempt to execute a RDPCR (impl. dep. #250-U3-Cs10), RDSOFTINT, RDTICK\_CMPR, RDSTICK, or RDSTICK\_CMPR instruction in nonprivileged mode (PSTATE.priv = 0 and HPSTATE.hpriv = 0) causes a *privileged\_opcode* exception (impl. dep. #250-U3-Cs10).

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.

In nonprivileged mode (PSTATE.priv = 0 and HPSTATE.hpriv = 0), the following cause a *privileged\_action* exception:

- execution of RDTICK when TICK.npt = 1
- execution of RDSTICK when STICK.npt = 1
- execution of RDPIC when nonprivileged access to PIC is disabled (PCR.priv = 1)

**Implementation** | RDasr shares an opcode with MEMBAR and STBAR<sup>D</sup>; it is  
**Note** | distinguished by  $rs1 = 15$  or  $rd = 0$  or  $(i = 0, \text{ and bit } 12 = 0)$ .

*Exceptions*

- illegal\_instruction*
- privileged\_opcode*
- fp\_disabled*
- privileged\_action*


*See Also*

- RDHPR on page 302
- RDPR on page 303
- WRasr on page 369

# RDHPR

## 7.74 Read Hyperprivileged Register

Instruction	op3	Operation	rs1	Assembly Language Syntax	Class
RDHPR <sup>H</sup>	10 1001	Read hyperprivileged register			<b>C2</b>
		HPSTATE	0	rdhpr %hpstate, reg <sub>rd</sub>	
		HTSTATE	1	rdhpr %htstate, reg <sub>rd</sub>	
		Reserved	2		
		HINTP	3	rdhpr %hintp, reg <sub>rd</sub>	
		Reserved	4		
		HTBA	5	rdhpr %htba, reg <sub>rd</sub>	
		HVER	6	rdhpr %hver, reg <sub>rd</sub>	
		Reserved	7–30		
		HSTICK_CMPR	31	rdhpr %hsys_tick_cmpr, reg <sub>rd</sub>	



**Description** This instruction reads the contents of the specified hyperprivileged state register into the destination register, R[rd]. The rs1 field in the RDHPR instruction determines which hyperprivileged register is read.

There are *MAXTL* copies of the HTSTATE register. A read from HTSTATE returns the value in the copy of HTSTATE indexed by the current value in the trap level register (TL).

An attempt to execute a RDHPR instruction when any of the following conditions exist causes an *illegal\_instruction* exception:

- instruction bits 13:0 are nonzero
- rs1 = 2, rs1 = 4, or 7 ≤ rs1 ≤ 30 (reserved rs1 values)
- HPSTATE.hpriv = 0 (the processor is not in hyperprivileged mode)
- rs1 = 1 (attempt to read the HTSTATE register) while TL = 0 (current trap level is zero)


**Exceptions** *illegal\_instruction*

**See Also** RDAsr on page 299  
RDPR on page 303  
WRHPR on page 372

# RDPR

## 7.75 Read Privileged Register

Instruction	op3	Operation	rs1	Assembly Language Syntax	Class
RDPR <sup>P</sup>	10 1010	Read Privileged register			<b>C2</b>
		TPC	0	rdpr %tpc, reg <sub>rd</sub>	
		TNPC	1	rdpr %tnpc, reg <sub>rd</sub>	
		TSTATE	2	rdpr %tstate, reg <sub>rd</sub>	
		TT	3	rdpr %tt, reg <sub>rd</sub>	
		TICK	4	rdpr %tick, reg <sub>rd</sub>	
		TBA	5	rdpr %tba, reg <sub>rd</sub>	
		PSTATE	6	rdpr %pstate, reg <sub>rd</sub>	
		TL	7	rdpr %tl, reg <sub>rd</sub>	
		PIL	8	rdpr %pil, reg <sub>rd</sub>	
		CWP	9	rdpr %cwp, reg <sub>rd</sub>	
		CANSAVE	10	rdpr %cansave, reg <sub>rd</sub>	
		CANRESTORE	11	rdpr %canrestore, reg <sub>rd</sub>	
		CLEANWIN	12	rdpr %cleanwin, reg <sub>rd</sub>	
		OTHERWIN	13	rdpr %otherwin, reg <sub>rd</sub>	
		WSTATE	14	rdpr %wstate, reg <sub>rd</sub>	
		<i>Reserved</i>	15		
		GL	16	rdpr %gl, reg <sub>rd</sub>	
		<i>Reserved</i>	17–31		



**Description** The rs1 field in the instruction determines the privileged register that is read. There are *MAXTL* 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 = 15, or 17 ≤ rs1 ≤ 31 (reserved rs1 values)
- 0 ≤ rs1 ≤ 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 or hyperprivileged mode.

**Implementation** In nonprivileged mode, *illegal\_instruction* exception due to  
**Note** 0 ≤ rs1 ≤ 3 and TL = 0 does not occur; the *privileged\_opcode* exception occurs instead.

# RDPR

An attempt to execute a RDPR instruction in nonprivileged mode (PSTATE.priv = 0 and HSTATE.hpriv = 0) causes a *privileged\_opcode* exception.

**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 UltraSPARC 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 299  
                      RDHPR on page 302  
                      WRPR on page 374

# RESTORE

## 7.76 RESTORE

Instruction	op3	Operation	Assembly Language Syntax	Class
RESTORE	11 1101	Restore Caller's Window	<code>restore reg<sub>rs1</sub>, reg<sub>or_imm</sub>, reg<sub>rd</sub></code>	<b>A1</b>



*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 CANSAVE.

# RESTORE

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 Spi ll/Fill Traps* on page 469. The fill trap handler is invoked with CWP set to point to the window to be filled, that is, old CWP – 1.

<b>Programming Note</b>	The vectoring of fill traps can be controlled by setting the value of the OTHERWIN and WSTATE registers appropriately. For details, see the section “Splitting the Register Windows” in <i>Software Considerations</i> , contained in the separate volume <i>UltraSPARC Architecture Application Notes</i> .  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.

<i>Exceptions</i>	<i>illegal_instruction</i> <i>fill_n_normal</i> ( $n = 0-7$ ) <i>fill_n_other</i> ( $n = 0-7$ )
-------------------	---

<i>See Also</i>	SAVE on page 313
-----------------	------------------

# RESTORED

## 7.77 RESTORED

Instruction	Operation	Assembly Language Syntax	Class
RESTORED <sup>P</sup>	Window has been restored	restored	C1



*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.

<b>Programming Notes</b>	Trap handler software for register window fills use the RESTORED instruction to indicate that a window has been filled successfully. For details, see the section “Example Code for Spill Handler” in <i>Software Considerations</i> , contained in the separate volume <i>UltraSPARC Architecture Application Notes</i> .  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 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 88) 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 and HSTATE.hpriv = 0) causes a *privileged\_opcode* exception.

# RESTORED

*Exceptions*     *illegal\_instruction*  
                  *privileged\_opcode*

*See Also*        ALLCLEAN on page 148  
                  INVALW on page 238  
                  NORMALW on page 285  
                  OTHERW on page 287  
                  SAVED on page 315

# RETRY

## 7.78 RETRY

Instruction	op3	Operation	Assembly Language Syntax	Class
RETRY <sup>P</sup>	11 1110	Return from Trap (retry trapped instruction)	retry	C1



*Description* The RETRY instruction restores the saved state from TSTATE[TL] (GL, CCR, ASI, PSTATE, and CWP), HTSTATE[TL] (HPSTATE), 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.

When a RETRY instruction is executed in *privileged* mode and HTSTATE[TL].hpstate.hpriv = 0 (which will cause the RETRY to return the virtual processor to nonprivileged or privileged mode), the value of GL restored from TSTATE[TL] saturates at MAXPGL. That is, if the value in TSTATE[TL].gl is greater than MAXPGL, then MAXPGL is substituted and written to GL. This protects against non-hyperprivileged software executing with  $GL > MAXPGL$ .

If software writes invalid or inconsistent state to TSTATE or HTSTATE before executing RETRY, virtual processor behavior during and after execution of the RETRY instruction is undefined.

The RETRY instruction does not provide an error barrier, as MEMBAR #Sync does (impl. dep. #215-U3).

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.

# RETRY

**Exceptions.** An attempt to execute the RETRY instruction when the following condition is true causes an *illegal\_instruction* exception:

- TL = 0 and the virtual processor is in privileged mode or hyperprivileged mode (PSTATE.priv = 1 or HPSTATE.hpriv = 1)

An attempt to execute a RETRY instruction in nonprivileged mode (PSTATE.priv = 0 and HPSTATE.hpriv = 0) 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).

A *trap\_level\_zero* disrupting trap can occur upon the *completion* of a RETRY instruction, if the following three conditions are true after RETRY has executed:

- *trap\_level\_zero* exceptions are enabled (HPSTATE.tlz = 1),
- the virtual processor is in nonprivileged or privileged mode (HPSTATE.hpriv = 0), and
- the trap level (TL) register's value is zero (TL = 0)

*Exceptions*     *illegal\_instruction*  
                  *privileged\_opcode*


*trap\_level\_zero*

*See Also*        DONE on page 166

# RETURN

## 7.79 RETURN

Instruction	op3	Operation	Assembly Language Syntax	Class
RETURN	11 1001	Return	return <i>address</i>	<b>A1</b>



*Description* The RETURN instruction causes a 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%16,%g0 !Trapped PC supplied to user trap handler
return%17 !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 !Same as jmp1 %i7 + 8, %g0
restore !Something useful like "restore
! %02,%12,%00"
```

or as:

```
save %sp, -framesize, %sp
. . .
return %i7 + 8
nop !Could do some useful work in the
! caller's window, e.g., "or %01, %02,%00"
```

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.

# RETURN

A RETURN instruction causes a *mem\_address\_not\_aligned* exception if either of the two least-significant bits of the target address is nonzero.

## *Exceptions*

*illegal\_instruction*

*fill\_n\_normal* ( $n = 0-7$ )


*fill\_n\_other* ( $n = 0-7$ )

*mem\_address\_not\_aligned*

# SAVE

## 7.80 SAVE

Instruction	op3	Operation	Assembly Language Syntax	Class
SAVE	11 1100	Save Caller's Window	save <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></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. For details, see the section “Leaf-Procedure Optimization” in Software Considerations, contained in the separate volume *UltraSPARC Architecture Application Notes*.

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.

# SAVE

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.

If `CANSAVE ≠ 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, `old CWP + 1`).

<b>Programming Note</b>	The vectoring of spill traps can be controlled by setting the value of the <code>OTHERWIN</code> and <code>WSTATE</code> registers appropriately. For details, see the section “Splitting the Register Windows” in <i>Software Considerations</i> , contained in the separate volume <i>UltraSPARC Architecture Application Notes</i> .  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 305

# SAVED

## 7.81 SAVED

Instruction	Operation	Assembly Language Syntax	Class
SAVED <sup>P</sup>	Window has been saved	saved	C1



*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. For details, see the section “Example Code for Spill Handler” in Software Considerations, contained in the separate volume *UltraSPARC Architecture Application Notes*.

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  $CANRESTORE = 0$  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 88) 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 and HSTATE.hpriv = 0) causes a *privileged\_opcode* exception.

*Exceptions* *illegal\_instruction*  
*privileged\_opcode*

# SAVED

*See Also*

ALLCLEAN on page 148  
INVALW on page 238  
NORMALW on page 285  
OTHERW on page 287  
RESTORED on page 307

# SETHI

## 7.82 SETHI

Instruction	op2	Operation	Assembly Language Syntax	Class
SETHI	100	Set High 22 Bits of Low Word	<code>sethi const22, reg<sub>rd</sub></code> <code>sethi %hi (value), reg<sub>rd</sub></code>	<b>A1</b>



*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 UltraSPARC 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  $2^{32}$ . The code below can be used to create the constant 0000 0000 ABCD 1234<sub>16</sub>:

```
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 1234<sub>16</sub>:

```
sethi %hi(0x5432edcb),%o0! note 0x5432EDCB, not 0xABCD1234
xor   %o0, 0x1e34, %o0! part of imm. overlaps upper bits
```

*Exceptions* None

# SHUTDOWN (Deprecated)

## 7.83 SHUTDOWN VIS 1

The SHUTDOWN instruction is deprecated and should not be used in new software.

Instruction	opf	Operation	Assembly Language Syntax	Class
SHUTDOWN <sup>D,P</sup>	0 1000 0000	Enter low-power mode	shutdown	D3



*Description* SHUTDOWN is a deprecated, privileged instruction that was used in early UltraSPARC implementations to bring the virtual processor or its containing system into a low-power state in an orderly manner. It had no effect on software-visible virtual processor state.

On an UltraSPARC Architecture implementation operating in privileged or hyperprivileged mode, SHUTDOWN behaves like a NOP (impl. dep. #206-U3-Cs10).

In an UltraSPARC Architecture 2005 implementation, this instruction is not implemented in hardware, causes an *illegal\_instruction* exception, and its effect is emulated in software.

*Exceptions* *illegal\_instruction* (instruction not implemented in hardware)

# SIAM

## 7.84 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	<i>siam</i> <i>siam_mode</i>	<b>B1</b>



*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*

# SIR

## 7.85 Software-Initiated Reset

Instruction	op3	rd	Operation	Assembly Language Syntax	Class
SIR <sup>H</sup>	11 0000	15	Software-Initiated Reset	<i>sir</i> <i>simm13</i>	<b>A2</b>



*Description*    SIR is a hyperprivileged instruction, used to generate a software-initiated reset (SIR). As with other traps, a software-initiated reset performs different actions when  $TL = MAXTL$  than it does when  $TL < MAXTL$ .

See *Software-Initiated Reset (SIR) Traps* on page 479 and *Software-Initiated Reset (SIR)* on page 548 for more information about software-initiated resets.

When executed in nonprivileged or privileged mode ( $HPSTATE.hpriv = 0$ ), SIR causes an *illegal\_instruction* exception (impl. dep. #116-V9).

<b>Implementation Notes</b>	The SIR instruction shares an opcode with WRasr; they are distinguished by the rd, rs1, and i fields ( $rd = 15, rs1 = 0$ , and $i = 1$ for SIR).  An instruction that uses the WRasr opcode ( $op1 = 10_2$ , $op3 = 11\ 0000_2$ ) with $i = 1$ is not an SIR instruction; see <i>Write Ancillary State Register</i> on page 369 for specification of its behavior.
-----------------------------	---


*Exceptions*    *software\_initiated\_reset*  
*illegal\_instruction*

*See Also*    WRasr on page 369

# SLL / SRL / SRA

## 7.86 Shift

Instruction	op3	x	Operation	Assembly Language Syntax	Class
SLL	10 0101	0	Shift Left Logical – 32 bits	sll <i>reg<sub>rs1</sub></i> , <i>reg_or_shcnt</i> , <i>reg<sub>rd</sub></i>	A1
SRL	10 0110	0	Shift Right Logical – 32 bits	srl <i>reg<sub>rs1</sub></i> , <i>reg_or_shcnt</i> , <i>reg<sub>rd</sub></i>	A1
SRA	10 0111	0	Shift Right Arithmetic– 32 bits	sra <i>reg<sub>rs1</sub></i> , <i>reg_or_shcnt</i> , <i>reg<sub>rd</sub></i>	A1
SLLX	10 0101	1	Shift Left Logical – 64 bits	sllx <i>reg<sub>rs1</sub></i> , <i>reg_or_shcnt</i> , <i>reg<sub>rd</sub></i>	A1
SRLX	10 0110	1	Shift Right Logical – 64 bits	srlx <i>reg<sub>rs1</sub></i> , <i>reg_or_shcnt</i> , <i>reg<sub>rd</sub></i>	A1
SRAX	10 0111	1	Shift Right Arithmetic – 64 bits	srax <i>reg<sub>rs1</sub></i> , <i>reg_or_shcnt</i> , <i>reg<sub>rd</sub></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 7-13 shows the shift count encodings for all values of  $i$  and  $x$ .

**TABLE 7-13** 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].

## SLL / SRL / SRA

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].

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 reg<sub>rs1</sub>, 0, reg<sub>rd</sub>” can be used to convert a 32-bit value to 64 bits, with sign extension into the upper word. “srl reg<sub>rs1</sub>, 0, reg<sub>rd</sub>” 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*

# STB / STH / STW / STX

## 7.87 Store Integer

Instruction	op3	Operation	Assembly Language Syntax	Class
STB	00 0101	Store Byte	stb <sup>†</sup> <i>reg<sub>rd</sub></i> [ <i>address</i> ]	<b>A1</b>
STH	00 0110	Store Halfword	sth <sup>‡</sup> <i>reg<sub>rd</sub></i> [ <i>address</i> ]	<b>A1</b>
STW	00 0100	Store Word	stw <sup>◇</sup> <i>reg<sub>rd</sub></i> [ <i>address</i> ]	<b>A1</b>
STX	00 1110	Store Extended Word	stx <i>reg<sub>rd</sub></i> [ <i>address</i> ]	<b>A1</b>

<sup>†</sup> *synonyms*: stub, stsb    <sup>‡</sup> *synonyms*: stuh, stsh    <sup>◇</sup> *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 104). 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.

STH causes a *mem\_address\_not\_aligned* exception if the effective address is not halfword-aligned. STW causes a *mem\_address\_not\_aligned* exception if the effective address is not word-aligned. STX causes a *mem\_address\_not\_aligned* exception if the effective address is not doubleword-aligned.

**Exceptions**

- illegal\_instruction*
- mem\_address\_not\_aligned*
- VA\_watchpoint*
- data\_access\_MMU\_error*

**See Also** STTW on page 346

# STBA / STHA / STWA / STXA

## 7.88 Store Integer into Alternate Space

Instruction	op3	Operation	Assembly Language Syntax	Class
STBA <sup>PASI</sup>	01 0101	Store Byte into Alternate Space	stba <sup>†</sup> <i>reg<sub>rd</sub></i> , [ <i>regaddr</i> ] <i>imm_asi</i> stba <i>reg<sub>rd</sub></i> , [ <i>reg_plus_imm</i> ] %asi	A1
STHA <sup>PASI</sup>	01 0110	Store Halfword into Alternate Space	stha <sup>‡</sup> <i>reg<sub>rd</sub></i> , [ <i>regaddr</i> ] <i>imm_asi</i> stha <i>reg<sub>rd</sub></i> , [ <i>reg_plus_imm</i> ] %asi	A1
STWA <sup>PASI</sup>	01 0100	Store Word into Alternate Space	stwa <sup>◇</sup> <i>reg<sub>rd</sub></i> , [ <i>regaddr</i> ] <i>imm_asi</i> stwa <i>reg<sub>rd</sub></i> , [ <i>reg_plus_imm</i> ] %asi	A1
STXA <sup>PASI</sup>	01 1110	Store Extended Word into Alternate Space	stxa <i>reg<sub>rd</sub></i> , [ <i>regaddr</i> ] <i>imm_asi</i> stxa <i>reg<sub>rd</sub></i> , [ <i>reg_plus_imm</i> ] %asi	A1

<sup>†</sup> synonyms: stuba, stsba    <sup>‡</sup> synonyms: stuha, stsha    <sup>◇</sup> synonyms: sta, stua, 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 access is privileged if bit 7 of the ASI is 0; otherwise, it is not privileged. 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 and HPSTATE.hpriv = 0), if bit 7 of the ASI is 0, these instructions cause a *privileged\_action* exception. In privileged mode (PSTATE.priv = 1 and HPSTATE.hpriv = 0), if the ASI is in the range 30<sub>16</sub> to 7F<sub>16</sub>, these instructions cause a *privileged\_action* exception.

STHA causes a *mem\_address\_not\_aligned* exception if the effective address is not halfword-aligned. STWA causes a *mem\_address\_not\_aligned* exception if the effective address is not word-aligned. STXA causes a *mem\_address\_not\_aligned* exception if the effective address is not doubleword-aligned.

# STBA / STHA / STWA / STXA

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 *data\_access\_exception* exception.

---

ASIs valid for STBA, STHA, and STWA	
ASI_NUCLEUS	ASI_NUCLEUS_LITTLE
ASI_AS_IF_USER_PRIMARY	ASI_AS_IF_USER_PRIMARY_LITTLE
ASI_AS_IF_USER_SECONDARY	ASI_AS_IF_USER_SECONDARY_LITTLE
ASI_REAL	ASI_REAL_LITTLE
ASI_REAL_IO	ASI_REAL_IO_LITTLE
ASI_PRIMARY	ASI_PRIMARY_LITTLE
ASI_SECONDARY	ASI_SECONDARY_LITTLE

STXA can be used with any ASI (including, but not limited to, the above list), 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 *data\_access\_exception* exception.

---

ASIs invalid for STXA (cause <i>data_access_exception</i> exception)	
24 <sub>16</sub> (aliased to 27 <sub>16</sub> , ASI_LDTX_N)	2C <sub>16</sub> (aliased to 2F <sub>16</sub> , ASI_LDTX_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
24 <sub>16</sub> (deprecated ASI_QUAD_LDD)	2C <sub>16</sub> (deprecated ASI_QUAD_LDD_L)
ASI_PST8_PRIMARY	ASI_PST8_PRIMARY_LITTLE
ASI_PST8_SECONDARY	ASI_PST8_SECONDARY_LITTLE
ASI_PRIMARY_NO_FAULT	ASI_PRIMARY_NO_FAULT_LITTLE
ASI_SECONDARY_NO_FAULT	ASI_SECONDARY_NO_FAULT_LITTLE
ASI_PST16_PRIMARY	ASI_PST16_PRIMARY_LITTLE
ASI_PST16_SECONDARY	ASI_PST16_SECONDARY_LITTLE
ASI_PST32_PRIMARY	ASI_PST32_PRIMARY_LITTLE
ASI_PST32_SECONDARY	ASI_PST32_SECONDARY_LITTLE
ASI_FL8_PRIMARY	ASI_FL8_PRIMARY_LITTLE
ASI_FL8_SECONDARY	ASI_FL8_SECONDARY_LITTLE
ASI_FL16_PRIMARY	ASI_FL16_PRIMARY_LITTLE
ASI_FL16_SECONDARY	ASI_FL16_SECONDARY_LITTLE
ASI_BLOCK_COMMIT_PRIMARY	ASI_BLOCK_COMMIT_SECONDARY
ASI_BLOCK_PRIMARY	ASI_BLOCK_PRIMARY_LITTLE
ASI_BLOCK_SECONDARY	ASI_BLOCK_SECONDARY_LITTLE

**V8 Compatibility** | The SPARC V8 STA instruction was renamed STWA in the  
**Note** | SPARC V9 architecture.

## STBA / STHA / STWA / STXA

*Exceptions*     *mem\_address\_not\_aligned* (all except STBA)  
                      *privileged\_action*  
                      *VA\_watchpoint*


*See Also*        LDA on page 242  
                      STTWA on page 348

# STBAR - Deprecated

## 7.89 Store Barrier

The STBAR instruction is deprecated. Use the MEMBAR instruction instead.

Opcode	op3	Operation	Assembly Language Syntax	Class
STBAR <sup>D</sup>	10 1000	Store Barrier	stbar	Y2



### Description

The store barrier instruction (STBAR) forces *all* store and atomic load-store operations issued by a virtual processor prior to the STBAR to complete their effects on memory before *any* store or atomic load-store operations issued by that virtual processor subsequent to the STBAR are executed by memory.

### V8 Compatibility Notes

STBAR is identical in function to a MEMBAR instruction with  $m\text{mask} = 8_{16}$ . STBAR is retained for compatibility with existing SPARC V8 software.

For correctness, it is sufficient for a virtual processor to stop issuing new store and atomic load-store operations when an STBAR is encountered and to resume after all stores have completed and are observed in memory by all virtual processors. More efficient implementations may take advantage of the fact that the virtual processor is allowed to issue store and load-store operations after the STBAR, as long as those operations are guaranteed not to become visible before all the earlier stores and atomic load-stores have become visible to all virtual processors.

An attempt to execute a STBAR instruction when instruction bits 12:0 are nonzero causes an *illegal\_instruction* exception.

### Implementation Note

STBAR shares an opcode with MEMBAR, and RDasr; it is distinguished by  $rs1 = 15$ ,  $rd = 0$ ,  $i = 0$ , and bit 12 = 0.


*Exceptions*     *illegal\_instruction*

# STBLOCKF

## 7.90 Block Store VIS 1

The STBLOCKF instruction is intended to be a processor-specific instruction, which may or may not be implemented in future UltraSPARC Architecture implementations. Therefore, it should only be used in platform-specific dynamically-linked libraries, in hyperprivileged software, or in software created by a runtime code generator that is aware of the specific virtual processor implementation on which it is executing.

Instruction	ASI		Operation	Assembly Language Syntax	Class
	Value				
STBLOCKF	16 <sub>16</sub>		64-byte block store to primary address space, user privilege	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_BLK_AIUP</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>A2</b>
STBLOCKF	17 <sub>16</sub>		64-byte block store to secondary address space, user privilege	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_BLK_AIUS</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>A2</b>
STBLOCKF	1E <sub>16</sub>		64-byte block store to primary address space, little-endian, user privilege	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_BLK_AIUPL</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>A2</b>
STBLOCKF	1F <sub>16</sub>		64-byte block store to secondary address space, little-endian, user privilege	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_BLK_AIUSL</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>A2</b>
STBLOCKF	F0 <sub>16</sub>		64-byte block store to primary address space	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_BLK_P</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>A2</b>
STBLOCKF	F1 <sub>16</sub>		64-byte block store to secondary address space	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_BLK_S</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>A2</b>
STBLOCKF	F8 <sub>16</sub>		64-byte block store to primary address space, little-endian	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_BLK_PL</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>A2</b>
STBLOCKF	F9 <sub>16</sub>		64-byte block store to secondary address space, little-endian	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_BLK_SL</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>A2</b>



*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

# STBLOCKF

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, and its companion, LDBLOCKF, were originally defined to provide a fast mechanism for block-copy operations.

STBLOCKF 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 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 instructions may or may not allow for case (2) above. Such software should be checked to verify that either it always waits for STBLOCKF 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 operation is in progress.

A Block Store only guarantees atomicity for each 64-bit (8-byte) portion of the 64 bytes that it stores.

Software should assume the following (where “load operation” includes load, load-store, and LDBLOCKF instructions and “store operation” includes store, load-store, and STBLOCKF instructions):

- A STBLOCKF 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 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 and subsequent load operations.
- A STBLOCKF 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 and subsequent store operations.
- STBLOCKFs do not follow register dependency interlocks, as do ordinary stores.

# STBLOCKF

**Programming Note** | STBLOCKF is intended to be a processor-specific instruction (see the warning at the top of page 328). If STBLOCKF *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 #411-S10, described below.

**IMPL. DEP. #411-S10:** The following aspects of the behavior of the block store (STBLOCKF) instruction are implementation dependent:

- The memory ordering model that STBLOCKF 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 (the recommended behavior), or only on accesses to the first eight bytes.
- Whether STBLOCKFs to non-cacheable (TTE.cp = 0) pages execute in strict program order or not. If not, a STBLOCKF to a non-cacheable page causes an *illegal\_instruction* exception.
- Whether STBLOCKF follows register dependency interlocks (as ordinary stores do).
- Whether a STBLOCKF forces the data to be written to memory and invalidates copies in all caches present.
- Any other restrictions on the behavior of STBLOCKF, 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 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 and HPSTATE.hpriv = 0), if bit 7 of the ASI is 0 (ASIs 16<sub>16</sub>, 17<sub>16</sub>, 1E<sub>16</sub>, and 1F<sub>16</sub>), STBLOCKF causes a *privileged\_action* exception.

An access caused by STBLOCKF may trigger a *VA\_watchpoint* exception (impl. dep. #411-S10).

**Implementation Note** | STBLOCKF shares an opcode with the STDFA, STPARTIALF, and STSHORTF instructions; it is distinguished by the ASI used.

# STBLOCKF

*Exceptions*     *illegal\_instruction*  
                  *mem\_address\_not\_aligned*  
                  *privileged\_action*  
                  *VA\_watchpoint* (impl. dep. #411-S10)

*See Also*        LDBLOCKF on page 245

# STF / STDF / STQF / STXFSR

## 7.91 Store Floating-Point

Instruction	op3	rd	Operation	Assembly Language		Class
STF	10 0100	0–31	Store Floating-Point register	st	<i>freg<sub>rd</sub></i> , [ <i>address</i> ]	<b>A1</b>
STDF	10 0111	†	Store Double Floating-Point register	std	<i>freg<sub>rd</sub></i> , [ <i>address</i> ]	<b>A1</b>
STQF	10 0110	†	Store Quad Floating-Point register	stq	<i>freg<sub>rd</sub></i> , [ <i>address</i> ]	<b>C3</b>
STXFSR	10 0101	1	Store Floating-Point State register	stx	<i>%fsr</i> , [ <i>address</i> ]	<b>A1</b>
—	10 0101	2–31	<i>Reserved</i>			

† 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).

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** |  $FSR.ftt$  should not be zeroed by STXFSR until it is known that the  
**Note** | store will not cause a precise trap.

These instruction access memory using the implicit ASI (see page 104). 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, STDF, or STXFSR instruction when  $i = 0$  and instruction bits 12:5 are nonzero causes an *illegal\_instruction* exception.

# STF / STDF / STQF / STXFSR

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, STDF, or STXFSR instruction causes an *fp\_disabled* exception.

STF causes a *mem\_address\_not\_aligned* exception if the effective memory address is not word-aligned. STXFSR causes a *mem\_address\_not\_aligned* exception if the address is not doubleword-aligned.

STDF requires only word alignment in memory. However, if the effective address is word-aligned but not doubleword-aligned, an attempt to execute an STDF instruction 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)).

STQF requires only word alignment in memory. If the effective address is word-aligned but not quadword-aligned, an attempt to execute an STQF instruction 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)).

<b>Programming Note</b>	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 <i>not</i> properly aligned.
-------------------------	--

An attempt to execute an STQF instruction when rd{1} ≠ 0 causes an *fp\_exception\_other* (FSR.ftt = invalid\_fp\_register) exception.

<b>Implementation Note</b>	Since UltraSPARC Architecture 2005 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 = invalid_fp_register) 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 UltraSPARC Architecture 2005)
- mem\_address\_not\_aligned*
- fp\_exception\_other* (FSR.ftt = invalid\_fp\_register (STQF only))
- VA\_watchpoint*

## **STF / STDF / STQF / STXFSR**

### *See Also*

*Load Floating-Point* on page 249

*Block Store* on page 328

*Store Floating-Point into Alternate Space* on page 335

*Store (Lower) Floating-Point Status Register* on page 339

*Store Short Floating-Point* on page 344


*Store Partial Floating-Point* on page 341

# STFA / STDFA / STQFA

## 7.92 Store Floating-Point into Alternate Space

Instruction	op3	rd	Operation	Assembly Language Syntax	Class
STFA <sup>P<sub>ASI</sub></sup>	11 0100	0–31	Store Floating-Point Register to Alternate Space	<code>sta freg<sub>rd</sub>, [regaddr] imm<sub>asi</sub></code> <code>sta freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>A1</b>
STDFA <sup>P<sub>ASI</sub></sup>	11 0111	†	Store Double Floating-Point Register to Alternate Space	<code>stda freg<sub>rd</sub>, [regaddr] imm<sub>asi</sub></code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>A1</b>
STQFA <sup>P<sub>ASI</sub></sup>	11 0110	†	Store Quad Floating-Point Register to Alternate Space	<code>stqa freg<sub>rd</sub>, [regaddr] imm<sub>asi</sub></code> <code>stqa freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>C3</b>

† 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 access is privileged if bit 7 of the ASI is 0; otherwise, it is not privileged. 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.** STFA causes a *mem\_address\_not\_aligned* exception if the effective memory address is not word-aligned.

# STFA / STDFA / STQFA

STDFA requires only word alignment in memory. However, if the effective address is word-aligned but not doubleword-aligned, an attempt to execute an STDFA instruction 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(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 *STQF\_mem\_address\_not\_aligned* exception. In this case, the trap handler software must emulate the STQFA instruction and return (impl. dep. #112-V9-Cs10(b)).

**Implementation Note** | STDFA shares an opcode with the STBLOCKF, STPARTIALF, and STSHORTF instructions; it is distinguished by the ASI used.

An attempt to execute an STQFA instruction when  $rd\{1\} \neq 0$  causes an *fp\_exception\_other* (FSR.ftt = *invalid\_fp\_register*) exception.

**Implementation Note** | Since UltraSPARC Architecture 2005 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 and HPSTATE.hpriv = 0), if bit 7 of the ASI is 0, this instruction causes a *privileged\_action* exception. In privileged mode (PSTATE.priv = 1 and HPSTATE.hpriv = 0), if the ASI is in the range  $30_{16}$  to  $7F_{16}$ , 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 *data\_access\_exception* exception.

---

## ASIs valid for STFA and STQFA

ASI_NUCLEUS	ASI_NUCLEUS_LITTLE
ASI_AS_IF_USER_PRIMARY	ASI_AS_IF_USER_PRIMARY_LITTLE
ASI_AS_IF_USER_SECONDARY	ASI_AS_IF_USER_SECONDARY_LITTLE
ASI_REAL	ASI_REAL_LITTLE
ASI_REAL_IO	ASI_REAL_IO_LITTLE
ASI_PRIMARY	ASI_PRIMARY_LITTLE
ASI_SECONDARY	ASI_SECONDARY_LITTLE

# STFA / STDFA / STQFA

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 *data\_access\_exception* exception.

---

ASIs valid for STDFA	
ASI_NUCLEUS	ASI_NUCLEUS_LITTLE
ASI_AS_IF_USER_PRIMARY	ASI_AS_IF_USER_PRIMARY_LITTLE
ASI_AS_IF_USER_SECONDARY	ASI_AS_IF_USER_SECONDARY_LITTLE
ASI_REAL	ASI_REAL_LITTLE
ASI_REAL_IO	ASI_REAL_IO_LITTLE
ASI_PRIMARY	ASI_PRIMARY_LITTLE
ASI_SECONDARY	ASI_SECONDARY_LITTLE
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 †
ASI_BLOCK_PRIMARY †	ASI_BLOCK_PRIMARY_LITTLE †
ASI_BLOCK_SECONDARY †	ASI_BLOCK_SECONDARY_LITTLE †
ASI_BLOCK_COMMIT_PRIMARY †	
ASI_BLOCK_COMMIT_SECONDARY †	
ASI_FL8_PRIMARY ‡	ASI_FL8_PRIMARY_LITTLE ‡
ASI_FL8_SECONDARY ‡	ASI_FL8_SECONDARY_LITTLE ‡
ASI_FL16_PRIMARY ‡	ASI_FL16_PRIMARY_LITTLE ‡
ASI_FL16_SECONDARY ‡	ASI_FL16_SECONDARY_LITTLE ‡
ASI_PST8_PRIMARY *	ASI_PST8_PRIMARY_LITTLE *
ASI_PST8_SECONDARY *	ASI_PST8_SECONDARY_LITTLE *
ASI_PST16_PRIMARY *	ASI_PST16_PRIMARY_LITTLE *
ASI_PST16_SECONDARY *	ASI_PST16_SECONDARY_LITTLE *
ASI_PST32_PRIMARY *	ASI_PST32_PRIMARY_LITTLE *
ASI_PST32_SECONDARY *	ASI_PST32_SECONDARY_LITTLE *

† If this ASI is used with the opcode for STDFA, the STBLOCKF instruction is executed instead of STFA. For behavior of STBLOCKF, see *Block Store* on page 328.

‡ 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 344.

\* If this ASI is used with the opcode for STDFA, the STPARTIALF instruction is executed instead of STDFA. For behavior of STPARTIALF, see *Store Partial Floating-Point* on page 341.

## Exceptions

*illegal\_instruction*

*fp\_disabled*

*STDF\_mem\_address\_not\_aligned*

*STQF\_mem\_address\_not\_aligned* (STQFA only) (not used in UA-2005)

*mem\_address\_not\_aligned*

*fp\_exception\_other* (FSR.ftt = invalid\_fp\_register (STQFA only))

# STFA / STDFA / STQFA

*privileged\_action*  
*VA\_watchpoint*

## *See Also*

*Load Floating-Point from Alternate Space* on page 252  
*Block Store* on page 328  
*Store Floating-Point* on page 332  
*Store Short Floating-Point* on page 344  
*Store Partial Floating-Point* on page 341

# STFSR (Deprecated)

## 7.93 Store (Lower) Floating-Point Status Register

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 <sup>D</sup>	10 0101	0	Store Floating-Point State Register Lower	st %fsr, [address]	D2



**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 the *FSR* into memory.

STFSR zeroes *FSR.ftt* after writing the *FSR* to memory.

**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 104). The effective address for this instruction is “*R*[*rs1*] + *R*[*rs2*]” if *i* = 0, or “*R*[*rs1*] + *sign\_ext*(*simm13*)” 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, UltraSPARC Architecture implementations continue to support it for compatibility with existing SPARC V8 software. The STFSR instruction is defined to store only 32 bits of the *FSR* into memory, while STXFSR allows SPARC V9 software to store all 64 bits of the *FSR*.

## STFSR (Deprecated)

*Exceptions*     *illegal\_instruction*  
                  *fp\_disabled*  
                  *mem\_address\_not\_aligned*  
                  *VA\_watchpoint*


*See Also*        *Store Floating-Point* on page 332



# STPARTIALF

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 locations.

The mask value in  $R[rs2]$  has the same format as the result specified by the pixel compare instructions (see *SIMD Signed Compare* on page 178). The most significant bit of the mask (not 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 7-29** Mask Format for Partial Store

In an UltraSPARC Architecture 2005 implementation, these instructions are not implemented in hardware, cause a *data\_access\_exception* exception, and are emulated in software.

**Exceptions.** An attempt to execute a STPARTIALF instruction when  $i = 1$  causes an *illegal\_instruction* exception.

# STPARTIALF

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.

STPARTIALF causes a *mem\_address\_not\_aligned* exception if the effective memory address is not word-aligned.

STPARTIALF requires only word alignment in memory for eight byte stores. If the effective address is word-aligned but not doubleword-aligned, it generates an *STDF\_mem\_address\_not\_aligned* exception. In this case, the trap handler software shall emulate the STDFA 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  $C0_{16}$ – $C5_{16}$  and  $C8_{16}$ – $CD_{16}$  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 254.

<b>Implementation</b>	STPARTIALF shares an opcode with the STBLOCKF, STDFA,
<b>Note</b>	and STSHORTF instructions; it is distinguished by the ASI used.

## Exceptions

*illegal\_instruction*

*fp\_disabled*


*data\_access\_exception* (not implemented in hardware in UA-2005)

*data\_access\_MMU\_error*

# STSHORTF

## 7.95 Store Short Floating-Point VIS 1

Instruction	ASI Value	Operation	Assembly Language Syntax	Class
STSHORTF	D0 <sub>16</sub>	8-bit store to primary address space	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_FL8_P</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>C3</b>
STSHORTF	D1 <sub>16</sub>	8-bit store to secondary address space	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_FL8_S</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>C3</b>
STSHORTF	D8 <sub>16</sub>	8-bit store to primary address space, little-endian	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_FL8_PL</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>C3</b>
STSHORTF	D9 <sub>16</sub>	8-bit store to secondary address space, little-endian	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_FL8_SL</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>C3</b>
STSHORTF	D2 <sub>16</sub>	16-bit store to primary address space	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_FL16_P</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>C3</b>
STSHORTF	D3 <sub>16</sub>	16-bit store to secondary address space	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_FL16_S</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>C3</b>
STSHORTF	DA <sub>16</sub>	16-bit store to primary address space, little-endian	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_FL16_PL</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>C3</b>
STSHORTF	DB <sub>16</sub>	16-bit store to secondary address space, little-endian	<code>stda freg<sub>rd</sub>, [regaddr] #ASI_FL16_SL</code> <code>stda freg<sub>rd</sub>, [reg_plus_imm] %asi</code>	<b>C3</b>



**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 stores are typically used with the `FALIGNDATA` instruction (see *Align Data* on page 173) to assemble or store 64 bits on noncontiguous components.

**Implementation** STSHORTF shares an opcode with the `STBLOCKF`, `STDFA`, and `STPARTIALF` instructions; it is distinguished by the ASI used.

In an UltraSPARC Architecture 2005 implementation, these instructions are not implemented in hardware, cause a `data_access_exception` exception, and are emulated in software.

# STSHORTF

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.

STSHORTF causes a *mem\_address\_not\_aligned* exception if the effective memory address is not halfword-aligned.

An 8-bit STSHORTF (using ASI D0<sub>16</sub>, D1<sub>16</sub>, D8<sub>16</sub>, or D9<sub>16</sub>) can be performed to an arbitrary memory address (no alignment requirement).

A 16-bit STSHORTF (using ASI D2<sub>16</sub>, D3<sub>16</sub>, DA<sub>16</sub>, or DB<sub>16</sub>) to an address that is not halfword-aligned (an odd address) causes a *mem\_address\_not\_aligned* exception.

## *Exceptions*

*VA\_watchpoint*  
*data\_access\_exception*  
*data\_access\_MMU\_error*


# STTW (Deprecated)

## 7.96 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 <sup>D</sup>	00 0111	Store Integer Twin Word	<i>sttw</i> <i>reg<sub>rd</sub></i> , [ <i>address</i> ]	<b>D2</b>

† 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 104). 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 UltraSPARC Architecture 2005 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 (Deprecated)

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.

<b>Programming Notes</b>	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.  If STTW is emulated in software, STX instruction should be used for the memory access in the emulation code to preserve atomicity.
--------------------------	--

### Exceptions

*unimplemented\_STTW*  
*illegal\_instruction*  
*mem\_address\_not\_aligned*  
*VA\_watchpoint*  
*fast\_data\_access\_MMU\_miss*  
*data\_access\_MMU\_miss*  
*data\_access\_MMU\_error*  
*fast\_data\_access\_protection*

### See Also

STW/STX on page 323  
STTWA on page 348

# STTWA (Deprecated)


## 7.97 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 <sup>D, P</sup> ASI	01 0111	Store Twin Word into Alternate Space	<code>sttwa reg<sub>rd</sub> [reg<sub>addr</sub>] imm<sub>asi</sub></code> <code>sttwa reg<sub>rd</sub> [reg<sub>plus_imm</sub>] %asi</code>	‡

† 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 (00<sub>16</sub>-7F<sub>16</sub>); **D2** for unrestricted ASIs (80<sub>16</sub>-FF<sub>16</sub>)



**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<sub>asi</sub> field if i = 0, or in the ASI register if i = 1. The access is privileged if bit 7 of the ASI is 0; otherwise, it is not privileged. 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.

# STTWA (Deprecated)

**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 UltraSPARC Architecture 2005 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 and HPSTATE.hpriv = 0), if bit 7 of the ASI is 0, this instruction causes a *privileged\_action* exception. In privileged mode (PSTATE.priv = 1 and HPSTATE.hpriv = 0), if the ASI is in the range 30<sub>16</sub> to 7F<sub>16</sub>, this instruction causes a *privileged\_action* exception.

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 *data\_access\_exception* exception (impl. dep. #300-U4-Cs10).

---

ASIs valid for STTWA	
ASI_NUCLEUS	ASI_NUCLEUS_LITTLE
ASI_AS_IF_USER_PRIMARY	ASI_AS_IF_USER_PRIMARY_LITTLE
ASI_AS_IF_USER_SECONDARY	ASI_AS_IF_USER_SECONDARY_LITTLE
ASI_REAL	ASI_REAL_LITTLE
ASI_REAL_IO	ASI_REAL_IO_LITTLE
ASI_PRIMARY	ASI_PRIMARY_LITTLE
ASI_SECONDARY	ASI_SECONDARY_LITTLE

**Programming Note** Nontranslating ASIs (see page 407) 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 *data\_access\_exception* exception (impl. dep. #300-U4-Cs10).

**Programming Note** STTWA 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. Therefore, software should avoid using STTWA.  
If STTWA is emulated in software, the STXA instruction should be used for the memory access in the emulation code to preserve atomicity.

*Exceptions*  
*unimplemented\_STTW*  
*illegal\_instruction*  
*mem\_address\_not\_aligned*

## STTWA (Deprecated)


*privileged\_action*  
*VA\_watchpoint*  
*fast\_data\_access\_MMU\_miss*  
*data\_access\_MMU\_miss*  
*data\_access\_MMU\_error*  
*fast\_data\_access\_protection*

*See Also*      STWA/STXA on page 324  
                  STTW on page 346

# SUB

## 7.98 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*

# SWAP (Deprecated)

## 7.99 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 <sup>D</sup>	00 1111	Swap Register with Memory	<code>swap [address], reg<sub>rd</sub></code>	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 104). 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*  
*fast\_data\_access\_MMU\_miss*  
*data\_access\_MMU\_miss*  
*data\_access\_MMU\_error*  
*fast\_data\_access\_protection*


# SWAPA (Deprecated)

## 7.100 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 <sup>D, P<sub>ASI</sub></sup>	01 1111	Swap register with Alternate Space Memory	swapa [regaddr] imm_asi, reg <sub>rd</sub> swapa [reg_plus_imm] %asi, reg <sub>rd</sub>	‡

‡ **Y3** for restricted ASIs (00<sub>16</sub>-7F<sub>16</sub>); **D2** for unrestricted ASIs (80<sub>16</sub>-FF<sub>16</sub>)



*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 access is privileged if bit 7 of the ASI is 0; otherwise, it is not privileged. 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.

# SWAPA (Deprecated)

In nonprivileged mode (PSTATE.priv = 0 and HPSTATE.hpriv = 0), if bit 7 of the ASI is 0, this instruction causes a *privileged\_action* exception. In privileged mode (PSTATE.priv = 1 and HPSTATE.hpriv = 0), if the ASI is in the range 30<sub>16</sub> to 7F<sub>16</sub>, 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 *data\_access\_exception* exception.

---

## ASIs valid for SWAPA

ASI_NUCLEUS	ASI_NUCLEUS_LITTLE
ASI_AS_IF_USER_PRIMARY	ASI_AS_IF_USER_PRIMARY_LITTLE
ASI_AS_IF_USER_SECONDARY	ASI_AS_IF_USER_SECONDARY_LITTLE
ASI_PRIMARY	ASI_PRIMARY_LITTLE
ASI_SECONDARY	ASI_SECONDARY_LITTLE
ASI_REAL	ASI_REAL_LITTLE


*Exceptions*

- mem\_address\_not\_aligned*
- privileged\_action*
- VA\_watchpoint*
- data\_access\_exception*
- fast\_data\_access\_MMU\_miss*
- data\_access\_MMU\_miss*
- data\_access\_MMU\_error*
- fast\_data\_access\_protection*

# TADDcc

## 7.101 Tagged Add

Instruction	op3	Operation	Assembly Language Syntax	Class
TADDcc	10 0000	Tagged Add and modify cc's	<code>taddcc reg_rs1, reg_or_imm, reg_rd</code>	<b>A1</b>



*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<sup>D</sup> on page 356  
TSUBcc on page 361

# TADDccTV (Deprecated)

## 7.102 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 <sup>D</sup>	10 0010	Tagged Add and modify cc's, or Trap on Overflow	taddcctv <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i>	<b>D2</b>



*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.

# TADDccTV (Deprecated)

<b>SPARC V8 Compatibility Note</b>	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 355  
                  TSUBccTV<sup>D</sup> on page 362

# Tcc

## 7.103 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	<b>not</b> Z	tn <sup>†</sup> <i>i_or_x_cc, software_trap_number</i>	A1
TE	11 1010	0001	Trap on Equal	Z	te <sup>‡</sup> <i>i_or_x_cc, software_trap_number</i>	A1
TG	11 1010	1010	Trap on Greater	<b>not</b> (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	<b>not</b> (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	<b>not</b> (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)	<b>not</b> C	tcc <sup>◇</sup> <i>i_or_x_cc, software_trap_number</i>	A1
TCS	11 1010	0101	Trap on Carry Set (Less Than, Unsigned)	C	tcs <sup>∇</sup> <i>i_or_x_cc, software_trap_number</i>	A1
TPOS	11 1010	1110	Trap on Positive or zero	<b>not</b> 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	<b>not</b> 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

<sup>†</sup> synonym: tnz    <sup>‡</sup> synonym: tz    <sup>◇</sup> synonym: tgeu    <sup>∇</sup> synonym: tlu



# Tcc

cc1 :: cc0	Condition Codes Evaluated
00	CCR.icc
01	— ( <i>illegal_instruction</i> )
10	CCR.xcc
11	— ( <i>illegal_instruction</i> )

## 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 and hyperprivileged modes, 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 and hyperprivileged modes 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:

# Tcc

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 and HSTATE.hpriv = 0)	<i>trap_instruction</i> exception (to privileged mode) ( $256 \leq \text{TT} \leq 383$ )	— (not possible)
Privileged (PSTATE.priv = 1 and HSTATE.hpriv = 0)	<i>trap_instruction</i> exception (to privileged mode) ( $256 \leq \text{TT} \leq 383$ )	<i>htrap_instruction</i> exception (to hyperprivileged mode) ( $384 \leq \text{TT} \leq 511$ )
Hyperprivileged (and HSTATE.hpriv = 1)	<i>htrap_instruction</i> exception (to hyperprivileged 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 12:5 are nonzero
- $i = 1$  and instruction bits 10:8 are nonzero
- $\text{cc0} = 1$

If a Tcc instruction causes a *trap\_instruction* or *htrap\_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 470.


*Exceptions*

- illegal\_instruction*
- trap\_instruction* ( $0 \leq \text{SWTN} \leq 127$ )
- htrap\_instruction* ( $128 \leq \text{SWTN} \leq 255$ )

# TSUBcc

## 7.104 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>	<b>A1</b>



**Description** This instruction computes “ $R[rs1] - R[rs2]$ ” if  $i = 0$ , or “ $R[rs1] - \text{sign\_ext}(\text{simm13})$ ” 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 355  
TSUBccTV<sup>D</sup> on page 362

# TSUBccTV (Deprecated)

## 7.105 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 <sup>D</sup>	10 0011	Tagged Subtract and modify cc's, or Trap on Overflow	<code>tsubcc tv reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></code>	<b>D2</b>



**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.

# TSUBccTV (Deprecated)

<b>SPARC V8 Compatibility Note</b>	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<sup>D</sup> on page 356  
                  TSUBcc on page 361

# UDIV / UDIVcc / SDIV / SDIVcc (Deprecated)

## 7.106 Divide (64-bit ÷ 32-bit)

The UDIV, UDIVcc, SDIV, and SDIVcc instructions are deprecated and should not be used in new software. The UDIVX and SDIVX instructions should be used instead.

Opcode	op3	Operation	Assembly Language Syntax	Class
UDIV <sup>D</sup>	00 1110	Unsigned Integer Divide	<code>udiv reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></code>	<b>C2</b>
SDIV <sup>D</sup>	00 1111	Signed Integer Divide	<code>sdiv reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></code>	<b>C2</b>
UDIVcc <sup>D</sup>	01 1110	Unsigned Integer Divide and modify cc's	<code>udivcc reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></code>	<b>C2</b>
SDIVcc <sup>D</sup>	01 1111	Signed Integer Divide and modify cc's	<code>sdivcc reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></code>	<b>C2</b>



*Description* The 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).

## UDIV / UDIVcc / SDIV / SDIVcc (Deprecated)

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 7-14.

**TABLE 7-14** 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 <sub>16</sub> )

When no overflow occurs, the 32-bit result is zero-extended to 64 bits and written into register R[rd].

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	UDIVcc
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 7-14)
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

*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 **sign\_ext**(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 7-15.

# UDIV / UDIVcc / SDIV / SDIVcc (Deprecated)

TABLE 7-15 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 <sub>16</sub> )
Rational quotient $\leq -2^{31} - 1$	$-2^{31}$ (FFFF FFFF 8000 0000 <sub>16</sub> )

When no overflow occurs, the 32-bit result is sign-extended to 64 bits and written into register R[rd].

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	SDIVcc
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 a UDIV, UDIVcc, 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*

# UMUL / UMULcc / SMUL / SMULcc (Deprecated)

## 7.107 Multiply (32-bit)

The UMUL, UMULcc, 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
UMUL <sup>D</sup>	00 1010	Unsigned Integer Multiply	umul <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i>	<b>C2</b>
SMUL <sup>D</sup>	00 1011	Signed Integer Multiply	smul <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i>	<b>C2</b>
UMULcc <sup>D</sup>	01 1010	Unsigned Integer Multiply and modify cc's	umulcc <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i>	<b>C2</b>
SMULcc <sup>D</sup>	01 1011	Signed Integer Multiply and modify cc's	smulcc <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i>	<b>C2</b>



*Description* The multiply instructions perform 32-bit by 32-bit multiplications, producing 64-bit results. They compute “ $R[rs1]\{31:0\} \times R[rs2]\{31:0\}$ ” if  $i = 0$ , or “ $R[rs1]\{31:0\} \times \text{sign\_ext}(\text{simm13})\{31:0\}$ ” 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. Signed multiply instructions (SMUL, SMULcc) operate on signed integer word operands and compute a signed integer doubleword product.

UMUL and SMUL do not affect the condition code bits. UMULcc and SMULcc write the integer condition code bits, icc and xcc, as shown below.

Bit	UMULcc / SMULcc
icc.n	Set to 1 if $\text{product}\{31\} = 1$ ; otherwise, set to 0
icc.z	Set to 1 if $\text{product}\{31:0\} = 0$ ; otherwise, set to 0
icc.v	Set to 0
icc.c	Set to 0
xcc.n	Set to 1 if $\text{product}\{63\} = 1$ ; otherwise, set to 0
xcc.z	Set to 1 if $\text{product}\{63:0\} = 0$ ; otherwise, set to 0
xcc.v	Set to 0
xcc.c	Set to 0

# UMUL / UMULcc / SMUL / SMULcc (Deprecated)

**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/UMULcc is indicated by  $Y \neq 0$ .  
32-bit overflow after SMUL/SMULcc is indicated by  $Y \neq (R[rd] \gg 31)$ , where “ $\gg$ ” indicates 32-bit arithmetic right-shift.


An attempt to execute a UMUL, UMULcc, SMUL, or SMULcc instruction when  $i = 0$  and instruction bits 12:5 are nonzero causes an *illegal\_instruction* exception.

*Exceptions*      *illegal\_instruction*

## 7.108 Write Ancillary State Register

Instruction	rd	Operation	Assembly Language Syntax	Class
WRY <sup>D</sup>	0	Write Y register ( <i>deprecated</i> )	<code>wr reg_rs1, reg_or_imm, %y</code>	<b>C1</b>
—	1	<i>Reserved</i>		
WRCCR	2	Write Condition Codes register	<code>wr reg_rs1, reg_or_imm, %ccr</code>	<b>A1</b>
WRASI	3	Write ASI register	<code>wr reg_rs1, reg_or_imm, %asi</code>	<b>A1</b>
—	4	<i>Reserved</i> (read-only ASR (TICK))		
—	5	<i>Reserved</i> (read-only ASR (PC))		
WRFPRS	6	Write Floating-Point Registers Status register	<code>wr reg_rs1, reg_or_imm, %fprs</code>	<b>A1</b>
—	7–14	<i>Reserved</i>		
—	15	Software-initiated reset (see <i>Software-Initiated Reset</i> on page 320)		
WRPCR <sup>P</sup>	16	Write Performance Control register (PCR)	<code>wr reg_rs1, reg_or_imm, %pcr</code>	<b>A1</b>
WRPIC <sup>P</sup> <sub>PIC</sub>	17	Write Performance Instrumentation Counters (PIC)	<code>wr reg_rs1, reg_or_imm, %pic</code>	<b>A1</b>
—	18	<i>Reserved</i> (impl. dep. #8-V8-Cs20, #9-V8-Cs20)		
WRGSR	19	Write General Status register (GSR)	<code>wr reg_rs1, reg_or_imm, %gsr</code>	<b>A1</b>
WRSOFTINT_SET <sup>P</sup>	20	Set bits of per-virtual processor Soft Interrupt register	<code>wr reg_rs1, reg_or_imm, %softint_set</code>	<b>N1</b>
WRSOFTINT_CLR <sup>P</sup>	21	Clear bits of per-virtual processor Soft Interrupt register	<code>wr reg_rs1, reg_or_imm, %softint_clr</code>	<b>N1</b>
WRSOFTINT <sup>P</sup>	22	Write per-virtual processor Soft Interrupt register	<code>wr reg_rs1, reg_or_imm, %softint</code>	<b>N1</b>
WRTICK_CMPR <sup>P</sup>	23	Write Tick Compare register	<code>wr reg_rs1, reg_or_imm, %tick_cmpr</code>	<b>N1</b>
WRSTICK <sup>H</sup>	24	Write System Tick register	<code>wr reg_rs1, reg_or_imm, %sys_tick</code>	<b>N1</b>
WRSTICK_CMPR <sup>P</sup>	25	Write System Tick Compare register	<code>wr reg_rs1, reg_or_imm, %sys_tick_cmpr</code>	<b>N1</b>
—	26–27	<i>Reserved</i> (impl. dep. #8-V8-Cs20, 9-V8-Cs20)		
—	28–31	Implementation dependent (impl. dep. #8-V8-Cs20, 9-V8-Cs20)		

# WRAsr



**Description** The WRAsr instructions each store a value to the writable fields of the ancillary state register (ASR) specified by rd.

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  $rs1 = 0$  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.

**IMPL. DEP. #48-V8-Cs20:** WRAsr instructions with rd in the range 26–31 are available for implementation-dependent uses (impl. dep. #8-V8-Cs20). For a WRAsr instruction with rd in the range 26–31, 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 or hyperprivileged (impl. dep. #9-V8-Cs20), and
- whether an attempt to execute the instruction causes an *illegal\_instruction* exception.

**Note** | See the section “Read/Write Ancillary State Registers (ASRs)” in *Extending the UltraSPARC Architecture*, contained in the separate volume *UltraSPARC Architecture Application Notes*, for a discussion of extending the SPARC V9 instruction set by means of read/write ASR instructions.

# WRAsr

<b>V9 Compatibility Notes</b>	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 UltraSPARC Architecture because the IER, PSR, TBR, and WIM registers do not exist in the UltraSPARC Architecture.

See *Ancillary State Registers* on page 70 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
- $rd = 1, 4, 5, 7-14, 18,$  or  $26-31$
- $rd = 15$  and  $((rs1 \neq 0)$  or  $(i = 0))$
- the instruction is WRSTICK and the virtual processor is not in hyperprivileged mode ( $HPSTATE.hpriv = 0$ )

An attempt to execute a WRPCR (impl. dep. #250-U3-Cs10), WRSOFTINT\_SET, WRSOFTINT\_CLR, WRSOFTINT, WRTICK\_CMPR, or WRSTICK\_CMPR instruction in nonprivileged mode ( $PSTATE.priv = 0$  and  $HPSTATE.hpriv = 0$ ) 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.

An attempt to execute a WRPIC instruction in nonprivileged mode ( $PSTATE.priv = 0$  and  $HPSTATE.hpriv = 0$ ) when  $PCR.priv = 1$  causes a *privileged\_action* exception.

<b>Implementation Note</b>	The SIR instruction shares an opcode with WRAsr; they are distinguished by the $rd$ , $rs1$ , and $i$ fields ( $rd = 15, rs1 = 0,$ and $i = 1$ for SIR). See <i>Software-Initiated Reset</i> on page 320.
--------------------------------	---

*Exceptions*

- illegal\_instruction*
- privileged\_opcode*
- fp\_disabled*
- privileged\_action*


*See Also*

- RDAsr on page 299
- WRHPR on page 372
- WRPR on page 374

# WRHPR

## 7.109 Write Hyperprivileged Register

Instruction	op3	Operation	rd	Assembly Language Syntax	Class
WRHPR <sup>H</sup>	11 0011	Write hyperprivileged register			<b>C1</b>
		HPSTATE	0	wrhpr <i>reg_rs1</i> , <i>reg_or_imm</i> , %hpstate	
		HTSTATE	1	wrhpr <i>reg_rs1</i> , <i>reg_or_imm</i> , %htstate	
		<i>Reserved</i>	2		
		HINTP	3	wrhpr <i>reg_rs1</i> , <i>reg_or_imm</i> , %hintp	
		<i>Reserved</i>	4		
		HTBA	5	wrhpr <i>reg_rs1</i> , <i>reg_or_imm</i> , %htba	
		<i>Reserved</i>	6–30		
		HSTICK_CMPR	31	wrhpr <i>reg_rs1</i> , <i>reg_or_imm</i> , %sys_tick_cmpr	



*Description* A WRHPR instruction stores the value “R[rs1] xor R[rs2]” if  $i = 0$ , or “R[rs1] xor sign\_ext(simm13)” if  $i = 1$  to the writable fields of the specified hyperprivileged state register.

**Note** | The operation is **exclusive-or**.

The *rd* field in the instruction determines the hyperprivileged register that is written. There are *MAXTL* copies of the HTSTATE register, one for each trap level. A write to one of these registers sets the copy of HTSTATE indexed by the current value in the trap-level register (TL).

The WRHPR instruction is a *non-delayed-write* instruction. The instruction immediately following the WRHPR observes any changes made to virtual processor state made by the WRHPR.

An attempt to execute a WRHPR instruction when any of the following conditions exist causes an *illegal\_instruction* exception:

- $i = 0$  and instruction bits 12:5 are nonzero
- $rd = 2, 4$ , or 6-30 (reserved for future versions of the architecture)
- $rd = 1$  and  $TL = 0$  (write to HTSTATE when the trap level is zero)
- virtual processor is in nonprivileged or privileged mode ( $HPSTATE.hpriv = 0$ )

A *trap\_level\_zero* disrupting trap can occur upon the *completion* of a WRHPR instruction to HPSTATE, if the following three conditions are true after WRHPR has executed:

- *trap\_level\_zero* exceptions are enabled ( $HPSTATE.tlz = 1$ ),

# WRHPR

- the virtual processor is in nonprivileged or privileged mode (HPSTATE.hpriv = 0), and
- the trap level (TL) register's value is zero (TL = 0)

**Programming Note** Execution of a WRHPR instruction that causes the value of HPSTATE.hpriv to change from 1 to 0 is not guaranteed to work if the WRHPR is in the delay slot of a DCTI instruction. Therefore, it is recommended that WRHPR not be executed in a delay slot, especially if it will toggle the value of HPSTATE.hpriv to 0.

**Programming Note** For historical reasons, the WRPR instruction, not WRHPR, is used to write to the hyperprivileged TICK register. See *Write Privileged Register* on page 374.

*Exceptions*      *illegal\_instruction*  
                      *trap\_level\_zero*

*See Also*        RDHPR on page 302  
                      WRAsr on page 369  
                      WRPR on page 374

# WRPR

## 7.110 Write Privileged Register

Instruction	op3	Operation	rd	Assembly Language Syntax	Class
WRPR <sup>P</sup>	11 0010	Write Privileged register			<b>C1</b>
		TPC	0	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %tpc	
		TNPC	1	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %tnpc	
		TSTATE	2	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %tstate	
		TT	3	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %tt	
		TICK	4	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %tick	
		TBA	5	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %tba	
		PSTATE	6	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %pstate	
		TL	7	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %tl	
		PIL	8	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %pil	
		CWP	9	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %cwp	
		CANSAVE	10	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %cansave	
		CANRESTORE	11	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %canrestore	
		CLEANWIN	12	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %cleanwin	
		OTHERWIN	13	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %otherwin	
		WSTATE	14	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %wstate	
		<i>Reserved</i>	15		
		GL	16	wrpr <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %gl	
		<i>Reserved</i>	17–31		



*Description* This instruction stores the value “R[rs1] xor R[rs2]” if *i* = 0, or “R[rs1] xor sign\_ext(simm13)” if *i* = 1 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 *MAXTL* 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).

# 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.

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.

In privileged mode, *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. In hyperprivileged mode, *MAXTL* is the maximum value that may be written by a WRPR to TL; an attempt to write a larger value results in *MAXTL* being written to TL. For details, see TABLE 6-22 on page 100.

In privileged mode, *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. In hyperprivileged mode, *MAXGL* is the maximum value that may be written by a WRPR to GL; an attempt to write a larger value results in *MAXGL* being written to GL. For details, see TABLE 6-23 on page 103.

**Programming Note** | For historical reasons, the WRPR instruction, not WRHPR, is used to write to the hyperprivileged TICK register.

**Exceptions.** An attempt to execute a WRPR instruction in nonprivileged mode (PSTATE.priv = 0 and HSTATE.hpriv = 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) and (PSTATE.priv = 1 and HSTATE.hpriv = 0)  
(an attempt to write to hyperprivileged register TICK while in privileged mode)
- rd = 15, or 17-31 (reserved for future versions of the architecture)
- $0 \leq rd \leq 3$  (attempt to write TPC, TNPC, TSTATE, or TT register) while TL = 0 (current trap level is zero) and the virtual processor is in privileged or hyperprivileged mode.

**Implementation Note** | In nonprivileged mode, *illegal\_instruction* exception due to  $0 \leq rd \leq 3$  and TL = 0 does not occur; the *privileged\_opcode* exception occurs instead.

A *trap\_level\_zero* disrupting trap can occur upon the *completion* of a WRPR instruction to TL, if the following three conditions are true after WRPR has executed:

- *trap\_level\_zero* exceptions are enabled (HPSTATE.tlz = 1)
- the virtual processor is in nonprivileged or privileged mode (HPSTATE.hpriv = 0), and

# WRPR

- the trap level (TL) register's value is zero (TL = 0)


*Exceptions*     *privileged\_opcode*  
                  *illegal\_instruction*  
                  *trap\_level\_zero*

*See Also*        RDPR on page 303  
                  WRAsr on page 369  
                  WRHPR on page 372

# XOR / XNOR

## 7.111 XOR Logical Operation

Instruction	op3	Operation	Assembly Language Syntax	Class
XOR	00 0011	Exclusive <b>or</b>	<code>xor <i>reg_rs1</i>, <i>reg_or_imm</i>, <i>reg_rd</i></code>	<b>A1</b>
XORcc	01 0011	Exclusive <b>or</b> and modify cc's	<code>xorcc <i>reg_rs1</i>, <i>reg_or_imm</i>, <i>reg_rd</i></code>	<b>A1</b>
XNOR	00 0111	Exclusive <b>nor</b>	<code>xnor <i>reg_rs1</i>, <i>reg_or_imm</i>, <i>reg_rd</i></code>	<b>A1</b>
XNORcc	01 0111	Exclusive <b>nor</b> and modify cc's	<code>xnorcc <i>reg_rs1</i>, <i>reg_or_imm</i>, <i>reg_rd</i></code>	<b>A1</b>



**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*

# XOR / XNOR

**Note: This chapter is undergoing final review; please check back later for a copy of UltraSPARC Architecture 2005 containing the final version of this chapter.**

## IEEE Std 754-1985 Requirements for UltraSPARC Architecture 2005

---

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 379.
- **NaN Operand and Result Definitions** on page 380.
- **Trapped Underflow Definition ( $ufm = 1$ )** on page 382.
- **Untrapped Underflow Definition ( $ufm = 0$ )** on page 383.
- **Integer Overflow Definition** on page 383.
- **Floating-Point Nonstandard Mode** on page 384.

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 61 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 or unimplemented FPOps execute as if by hardware; that is, such a trap is undetectable by application software, except that timing may be affected.

<b>Programming Note</b>	<p>A user-mode trap handler invoked for an <i>IEEE_754_exception</i>, 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 <i>FSR.ftt</i> = <i>unfinished_FPop</i> or <i>FSR.ftt</i> = <i>unimplemented_FPop</i>, can rely on the following behavior:</p> <ul style="list-style-type: none"> <li>■ The address of the instruction that caused the exception will be available.</li> <li>■ The destination floating-point register(s) are unchanged from their state prior to that instruction's execution.</li> <li>■ The floating-point condition codes (<i>fcc0</i>, <i>fcc1</i>, <i>fcc2</i>, and <i>fcc3</i>) are unchanged.</li> <li>■ The <i>FSR.aexc</i> field is unchanged.</li> <li>■ The <i>FSR.cexc</i> field contains exactly one bit set to 1, corresponding to the exception that caused the trap.</li> <li>■ The <i>FSR.ftt</i>, <i>FSR.qne</i>, and reserved fields of <i>FSR</i> are zero.</li> </ul>
-------------------------	---

---

## 8.2 NaN Operand and Result Definitions

An untrapped floating-point result can be in a format that is either the same as, or different from, the format of the source operands. These two cases are described separately below.

### 8.2.1 Untrapped Result in Different Format from Operands

- **F<sdq>TO<sdq> or F<sd>MUL<dq> with a quiet NaN operand** — No exception caused; result is a quiet NaN. The operand is transformed as follows:

**NaN transformation:** The most significant bits of the operand fraction are copied to the most significant bits of the result fraction. In conversion to a narrower format, excess low-order bits of the operand fraction are discarded (which is not considered a "rounding" operation). In conversion to a wider format, excess low-order bits of the result fraction are set to 0. The quiet bit (the most significant bit of the result fraction) is always set to 1, so the NaN transformation always produces a quiet NaN. The sign bit is copied from the operand to the result without modification.

- **F<sdq>TO<sdq> or F<sd>MUL<dq> with a signalling NaN operand** — Invalid exception; result is the signalling NaN operand processed by the NaN transformation above to produce a quiet NaN.
- **FCMPE<sdq> with any NaN operand** — Invalid exception; the selected floating-point condition code is set to unordered.
- **FCMP<sdq> with any signalling NaN operand** — Invalid exception; the selected floating-point condition code is set to unordered.
- **FCMP<sdq> with any quiet NaN operand but no signalling NaN operand** — No exception; the selected floating-point condition code is set to unordered.

## 8.2.2 Untrapped Result in Same Format as Operands

- **No NaN operand** — For an invalid operation such as `sqrt(-1.0)` or `0.0 ÷ 0.0`, the result is the quiet NaN with sign = zero, exponent = all 1's, and fraction = all ones. The sign is zero to distinguish such results from storage initialized to all ones.
- **One operand, a quiet NaN** — No exception; result is the quiet NaN operand.
- **One operand, a signalling NaN** — Invalid exception; result is the signalling NaN with its quiet bit (most significant bit of fraction field) set to 1.
- **Two operands, both quiet NaNs** — No exception; result is the rs2 (second source) operand.
- **Two operands, both signalling NaNs** — Invalid exception; result is the rs2 operand with the quiet bit set to 1.
- **Two operands, only one is a signalling NaN** — Invalid exception; result is the signalling NaN operand with the quiet bit set to 1.
- **Two operands, neither is a signalling NaN, only one is a quiet NaN** — No exception; result is the quiet NaN operand.

In TABLE 8-1, NaN<sub>*n*</sub> means that the NaN is in rs<sub>*n*</sub>, Q means quiet, S signalling.

**TABLE 8-1** Untrapped Floating-Point Results

		rs2 Operand		
		Number	QNaN2	SNaN2
rs1 operand	None	IEEE 754	QNaN2	QNaN2
	Number	IEEE 754	QNaN2	QNaN2
	QNaN1	QNaN1	QNaN2	QNaN2
	SNaN1	QNaN1	QNaN1	QNaN2

QNaN $n$  means a quiet NaN produced by the NaN transformation on a signalling NaN from rs $n$ ; the invalid exception is always indicated. The QNaN $n$  results in the table never generate an exception, but IEEE 754 specifies several cases of invalid exceptions, and QNaN results from operands that are both numbers.

## 8.3 Trapped Underflow Definition (ufm = 1)

An UltraSPARC Architecture virtual processor detects tininess before rounding occurs. (impl. dep. #55-V8-Cs10)

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 SPARC V9 architecture at the hardware, hyperprivileged, and privileged software levels. If they are created at all, it would be by user software in a nonprivileged-mode trap handler.

## 8.4 Untrapped Underflow Definition (ufm = 0)

On an implementation that detects tininess before rounding, 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.

TABLE 8-2 summarizes what happens on an implementation that detects tininess before rounding, 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 8-2** Untrapped Floating-Point Underflow (Tininess Detected Before Rounding)

	Underflow trap: Inexact trap:	ufm = 1 nxm = x	ufm = 0 nxm = 1	ufm = 0 nxm = 0
$u = r$	$r$ is minimum normal	None	None	None
	$r$ is subnormal	UF	None	None
	$r$ is zero	None	None	None
$u \neq r$	$r$ is minimum normal	UF	NX	uf nx
	$r$ is subnormal	UF	NX	uf nx
	$r$ is zero	UF	NX	uf nx
UF = <i>fp_exception_ieee_754</i> trap with <i>cexc.ufc</i> = 1 NX = <i>fp_exception_ieee_754</i> trap with <i>cexc.nxc</i> = 1 uf = <i>cexc.ufc</i> = 1, <i>aexc.ufa</i> = 1, no <i>fp_exception_ieee_754</i> trap nx = <i>cexc.nxc</i> = 1, <i>aexc.nxa</i> = 1, no <i>fp_exception_ieee_754</i> trap				

## 8.5 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* should be set and *fp\_exception\_IEEE\_754* should be 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` should be set and `fp_exception_IEEE_754` should be 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  $-2^{63}$ .

---

## 8.6 Floating-Point Nonstandard Mode

Please refer to *Nonstandard Floating-Point (ns)* on page 63 for information.

# Memory

---

The UltraSPARC 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. For additional information on the use of the models in programming real systems, see *Programming with the Memory Models*, contained in the separate volume *UltraSPARC Architecture Application Notes*.

This chapter contains a great deal of theoretical information so that the discussion of the UltraSPARC Architecture TSO memory model has sufficient background.

This chapter describes memory models in these sections:

- **Memory Location Identification** on page 386.
- **Memory Accesses and Cacheability** on page 386.
- **Memory Addressing and Alternate Address Spaces** on page 389.
- **SPARC V9 Memory Model** on page 393.
- **The UltraSPARC Architecture Memory Model — TSO** on page 396.
- **Nonfaulting Load** on page 405.
- **Store Coalescing** on page 405.

---

## 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, real, or physical address.

---

## 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 UltraSPARC Architecture: cacheable and noncacheable accesses. The manner in which addresses are differentiated is implementation dependent. In some implementations, it is indicated by the page translation (TTE.cp), while in other implementations, it is physical address bit specific.

Although SPARC V9 does not specify memory ordering between cacheable and noncacheable accesses, the UltraSPARC Architecture maintains TSO ordering between memory references regardless of their cacheability.

The UltraSPARC Architecture obeys the Sun-5 Ordering rules as documented in the “Sun-4u/Sun-5 Ordering with TSO” specification.

### 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 UltraSPARC 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 of 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 405).

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 402.

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 UltraSPARC 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 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.

<b>V9 Compatibility Note</b>	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 114, 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 doubleword<sup>1</sup>. 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 UltraSPARC 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 are translated by the MMU in order to locate data in physical memory. Virtual addresses can be presented in nonprivileged mode and privileged mode, or in hyperprivileged mode using the `ASI_AS_IF_USER*` ASI variants.

<sup>1</sup> Two exceptions to this are the special `ASI_TWIN_DW_NUCLEUS[_L]` and `ASI_LD_TWINK_REAL[_L]` which provide hardware support for an atomic quad load to be used for TTE loads from TSBs.

**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.

**Physical addresses (PA).** A physical address is one that appears on the system bus and is the same as the physical addresses in legacy architectures. Hyperprivileged software is responsible for managing the translation of real addresses into physical addresses.

Nonprivileged software only uses virtual addresses. Privileged software uses virtual and real addresses. Hyperprivileged software uses physical addresses, except when the explicit `ASI_AS_IF_USER*` or `ASI_*REAL*` ASI variants are used for load and store alternate instructions.

## 9.3.2 Memory Address Spaces

The UltraSPARC Architecture supports accessing memory using virtual, real, or physical addresses. Multiple virtual address spaces within the same real address space are distinguished by a *context identifier* (context ID). Multiple real address spaces within the same physical address space are distinguished by a *partition identifier* (partition 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.

Hyperprivileged software can create multiple real address spaces, using the partition register to associate a partition ID with every real address. Hyperprivileged software manages the allocation of partition IDs.

**IMPL. DEP. #\_\_** The number of bits in the partition register is implementation dependent.

The full representation of each type of address is as follows:

```
real_address = context_ID :: virtual_address
physical_address = partition ID :: real_address
or
physical_address = partition ID :: context ID :: 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
- 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.

<b>Programming Note</b>	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 the `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, privileged and hyperprivileged modes is indicated in TABLE 10-1 on page 409.

Attempts by nonprivileged software (`PSTATE.priv = 0` and `HPSTATE.hpriv = 0`) to access restricted ASIs (ASI bit 7 = 0) cause a *privileged\_action* exception. Attempts by privileged software (`PSTATE.priv = 1` and `HPSTATE.hpriv = 0`) 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).

In hyperprivileged mode, all instruction fetches and loads and stores with implicit ASIs use a physical address, regardless of the value of TL.

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, restricted-to-privileged, or restricted-to-hyperprivileged:

- A nonrestricted ASI (ASI range  $80_{16} - FF_{16}$ ) is one that may be used independently of the privilege level (PSTATE.priv and HPSTATE.hpriv) 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 or hyperprivileged mode for a legal access to occur.
- A restricted-to-hyperprivileged ASI (ASI range  $30_{16} - 7F_{16}$ ) requires that the virtual processor be in hyperprivileged 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	Result of ASI Access in HP Mode
$00_{16} - 2F_{16}$	Restricted-to-privileged	<i>privileged_action</i> exception	Valid Access	Valid Access
$30_{16} - 7F_{16}$	Restricted-to-hyperprivileged	<i>privileged_action</i> exception	<i>privileged_action</i> exception	Valid Access
$80_{16} - FF_{16}$	Nonrestricted	Valid Access	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 UltraSPARC 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.




FIGURE 9-1 Processor Model: Uniprocessor System

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*.

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).

<b>V9 Compatibility Note</b>	An implementation can avoid blocking instruction execution in 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 401 and *The UltraSPARC Architecture Memory Model — TSO* on page 396 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, *Formal Specification of the Memory Models*, for more information.

## 9.4.2 Virtual Processor/Memory Interface Model

Each UltraSPARC 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 to be consistent (coherent). Instruction caches need not be kept consistent with data caches and therefore require explicit program action to ensure consistency when a program modifies an executing instruction stream. See *Synchronizing Instruction and Data Memory* on page 403 for details. Memory is shared in terms of address space, but it may be nonhomogeneous and distributed in an implementation. Mapping and caches are ignored in the model, since their functions are transparent to the memory model<sup>1</sup>.

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. The ASI-address couples it generates are translated by a memory management unit (MMU), which associates attributes with the address and may, in some instances, abort the memory transaction and signal an exception to the virtual processor.

For example, a region of memory may be marked as nonprefetchable, noncacheable, read-only, or restricted. It is the MMU's responsibility, working in conjunction with system software, to ensure that memory attribute constraints are not violated. See implementation-specific MMU documentation for detailed information about how this is accomplished in each UltraSPARC Architecture implementation.

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

<sup>1</sup> The model described here is only a model; implementations of UltraSPARC Architecture systems are unconstrained as long as their observable behaviors match those of the model.

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.

---

## 9.5 The UltraSPARC Architecture Memory Model — TSO

The UltraSPARC Architecture is a *model* that specifies the behavior observable by software on UltraSPARC 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 UltraSPARC Architecture virtual processors implement TSO ordering. The PSO and RMO models from SPARC V9 are not described in this UltraSPARC Architecture specification. UltraSPARC Architecture 2005 processors do not implement the PSO memory model directly, but all software written to run under PSO will execute correctly on an UltraSPARC Architecture 2005 processor (using the TSO model).

Whether memory models represented by  $PSTATE.mm = 10_2$  or  $11_2$  are supported in an UltraSPARC Architecture processor is implementation dependent (impl. dep. #113-V9-Ms10). 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 *The SPARC Architecture Manual-Version 9*. If the 11<sub>2</sub> 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 having this model available.

## 9.5.1 Memory Model Selection

The active memory model is specified by the 2-bit value in `PSTATE.mm`. The value 00<sub>2</sub> 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 UltraSPARC 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 UltraSPARC 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 394), 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, and #StoreStore.

In addition to the above TSO rules, the following rules apply to UltraSPARC 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 UltraSPARC Architecture Memory Model — TSO* on page 396) is desired.
- Accesses that have side effects are all strongly ordered with respect to each other.
- A MEMBAR #Lookaside 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 UltraSPARC Architecture virtual processor running in TSO mode, to ensure that the operations appear to complete in a particular 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 UltraSPARC Architecture Ordering Rules (TSO Memory Model)

From Memory Operation R (row):	To Memory Operation C (column):										
	load	store	atomic	bload	bstore	load_nc_e	store_nc_e	load_nc_ne	store_nc_ne	bload_nc	bstore_nc
load	#	#	#	S	S	#	#	#	#	S	S
store	M <sup>2</sup>	#	#	M	S	M	#	M	#	M	S
atomic	#	#	#	M	S	#	#	#	#	M	S
bload	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	# <sup>1</sup>	# <sup>1</sup>	# <sup>1</sup>	# <sup>1</sup>	S	S
store_nc_e	S	#	#	S	S	# <sup>1</sup>	# <sup>1</sup>	M <sup>2</sup>	# <sup>1</sup>	M	S
load_nc_ne	#	#	#	S	S	# <sup>1</sup>	# <sup>1</sup>	# <sup>1</sup>	# <sup>1</sup>	S	S
store_nc_ne	S	#	#	S	S	M <sup>2</sup>	# <sup>1</sup>	M <sup>2</sup>	# <sup>1</sup>	M	S
bload_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.

## 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 UltraSPARC Architecture provides three 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 *data\_access\_exception* 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 *data\_access\_exception* 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. For examples, see *Programming with the Memory Models*, contained in the separate volume *UltraSPARC Architecture Application Notes*.

### 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 UltraSPARC 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.<sup>1</sup> 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 SPARCV9 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

<sup>1</sup>Sequencing 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. See *Programming with the Memory Models*, contained in the separate volume *UltraSPARC Architecture Application Notes*, for examples of the use of sequencing MEMBARs.

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 01<sub>16</sub>, 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 08<sub>16</sub> (#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, *Formal Specification 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 <sub>16</sub>	0
Store < <i>m</i> Load	#StoreLoad	#StoreLoad	02 <sub>16</sub>	1
Load < <i>m</i> Store	#LoadStore	nop	04 <sub>16</sub>	2
Store < <i>m</i> Store	#StoreStore	nop	08 <sub>16</sub>	3

<b>Implementation Note</b>	An UltraSPARC Architecture 2005 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.

- **Lookaside Barrier** — 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. See the example in “Control and Status Registers” in *Programming with the Memory Models*, contained in the separate volume *UltraSPARC Architecture Application Notes*.
- **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. See the example in “I/O Registers with Side Effects” in *Programming with the Memory Models*, contained in the separate volume *UltraSPARC Architecture Application Notes*.
- **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	#Lookaside	10 <sub>16</sub>	0
Memory Issue Barrier	#MemIssue	20 <sub>16</sub>	1
Synchronization Barrier	#Sync	40 <sub>16</sub>	2

**Implementation Note** | In UltraSPARC Architecture 2005 implementations, MEMBAR #Lookaside and MEMBAR #MemIssue are typically implemented as a MEMBAR #Sync.

For more details, see the MEMBAR instruction on page 271 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 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.

<b>Programming Note</b>	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 UltraSPARC 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 UltraSPARC Architecture implementation may ignore the address in the operands of a FLUSH instruction.

<b>Programming Note</b>	UltraSPARC 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 186 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 *data\_access\_exception* 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 *data\_access\_exception* 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 *data\_access\_exception* 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<sup>1</sup>. 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.

<sup>1</sup>Other than the impact of occupying TLB entries.

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 407.
- **ASI Values** on page 407.
- **ASI Assignments** on page 408.
- **Special Memory Access ASIs** on page 422.

---

## 10.1 Address Space Identifiers and Address Spaces

An UltraSPARC 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. This physical address space can be accessed through virtual-to-physical address mapping or through the MMU bypass mode.

---

## 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 *Normal* or *Translating* ASI is translated by the MMU.
- A *Nontranslating* ASI is not translated by the MMU; instead the address is passed through unchanged. Nontranslating ASIs are typically used for accessing internal registers.
- A *Bypass* ASI, like a nontranslating ASI, is not translated by the MMU and the address is passed through unchanged. However, unlike a nontranslating ASI, an access using a bypass ASI can cause exception(s) only visible in hyperprivileged mode (such as a *PA\_watchpoint* exception). Bypass ASIs are typically used by privileged or hyperprivileged software for directly accessing memory using real or physical (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 UltraSPARC 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 UltraSPARC Architecture implementations and some are only present in some implementations.

An ASI marked with a closed bullet (●) is required to be implemented on all UltraSPARC Architecture 2005 processors.

An ASI marked with an open bullet (○) is defined by the UltraSPARC Architecture 2005 but is not necessarily implemented in all UltraSPARC Architecture 2005 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.

ASIs marked "Reserved" are set aside for use in future revisions to the architecture and are not to be used by implementations. 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** UltraSPARC Architecture ASIs (1 of 13)

ASI Value	req'd(●) opt'l(○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	T/ Non-T/ Bypass	Shared /per strand	Description
00 <sub>16</sub> – 03 <sub>16</sub>	○	—	— <sup>2,12</sup>	—	—	—	Implementation dependent <sup>1</sup>
04 <sub>16</sub>	●	ASI_NUCLEUS (ASI_N)	RW <sup>2,4</sup>	(decoded)	T	—	Implicit address space, nucleus context, TL > 0
05 <sub>16</sub> – 0B <sub>16</sub>	○	—	— <sup>2,12</sup>	—	—	—	Implementation dependent <sup>1</sup>
0C <sub>16</sub>	●	ASI_NUCLEUS_LITTLE (ASI_NL)	RW <sup>2,4</sup>	(decoded)	T	—	Implicit address space, nucleus context, TL > 0, little-endian

**TABLE 10-1** UltraSPARC Architecture ASIs (2 of 13)

ASI Value	req'd(●) opt'l (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	T/ Non-T/ Bypass	Shared /per strand	Description
0D <sub>16</sub> – 0F <sub>16</sub>	○	—	— <sup>2,12</sup>	—	—	—	Implementation dependent <sup>I</sup>
10 <sub>16</sub>	●	ASI_AS_IF_USER_PRIMARY (ASI_AIUP)	RW <sup>2,4,18</sup>	(decoded)	T	—	Primary address space, as if user (nonprivileged)
11 <sub>16</sub>	●	ASI_AS_IF_USER_SECONDARY (ASI_AIUS)	RW <sup>2,4,18</sup>	(decoded)	T	—	Secondary address space, as if user (nonprivileged)
12 <sub>16</sub> – 13 <sub>16</sub>	○	—	— <sup>2,12</sup>	—	—	—	Implementation dependent <sup>I</sup>
14 <sub>16</sub>	○	ASI_REAL	RW <sup>2,4</sup>	(decoded)	B	—	Real address
15 <sub>16</sub>	○	ASI_REAL_IO <sup>D</sup>	RW <sup>2,5</sup>	(decoded)	B	—	Real address, noncacheable, with side effect (deprecated)
16 <sub>16</sub>	○	ASI_BLOCK_AS_IF_USER_PRIMARY (ASI_BLK_AIUP)	RW <sup>2,8,14,18</sup>	(decoded)	T	—	Primary address space, block load/store, as if user (nonprivileged)
17 <sub>16</sub>	○	ASI_BLOCK_AS_IF_USER_SECONDAR Y (ASI_BLK_AIUS)	RW <sup>2,8,14,18</sup>	(decoded)	T	—	Secondary address space, block load/store, as if user (nonprivileged)
18 <sub>16</sub>	●	ASI_AS_IF_USER_PRIMARY_LITTLE (ASI_AIUPL)	RW <sup>2,4,18</sup>	(decoded)	T	—	Primary address space, as if user (nonprivileged), little-endian
19 <sub>16</sub>	●	ASI_AS_IF_USER_SECONDARY_ LITTLE (ASI_AIUSL)	RW <sup>2,4,18</sup>	(decoded)	T	—	Secondary address space, as if user (nonprivileged), little-endian
1A <sub>16</sub> – 1B <sub>16</sub>	○	—	— <sup>2,12</sup>	—	—	—	Implementation dependent <sup>I</sup>
1C <sub>16</sub>	○	ASI_REAL_LITTLE (ASI_REAL_L)	RW <sup>2,4</sup>	(decoded)	B	—	Real address, little-endian
1D <sub>16</sub>	○	ASI_REAL_IO_LITTLE <sup>D</sup> (ASI_REAL_IO_L <sup>D</sup> )	RW <sup>2,5</sup>	(decoded)	B	—	Physical address, noncacheable, with side effect, little-endian (deprecated)
1E <sub>16</sub>	○	ASI_BLOCK_AS_IF_USER_PRIMARY_ LITTLE (ASI_BLK_AIUPL)	RW <sup>2,8,14,18</sup>	(decoded)	T	—	Primary address space, block load/store, as if user (nonprivileged), little-endian
1F <sub>16</sub>	○	ASI_BLOCK_AS_IF_USER_ SECONDARY_LITTLE (ASI_BLK_AIUS_L)	RW <sup>2,8,14,18</sup>	(decoded)	T	—	Secondary address space, block load/store, as if user (nonprivileged), little-endian

**TABLE 10-1** UltraSPARC Architecture ASIs (3 of 13)

ASI Value	req'd (●) opt'l (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	T/ Non-T/ Bypass	Shared /per strand	Description
20 <sub>16</sub>	○	ASI_SCRATCHPAD	RW <sup>2,6</sup>	(decoded; see below)	N	per strand	Privileged Scratchpad registers; implementation dependent <sup>1</sup>
	○		"	0 <sub>16</sub>	"	"	Scratchpad Register 0 <sup>1</sup>
	○		"	8 <sub>16</sub>	"	"	Scratchpad Register 1 <sup>1</sup>
	○		"	10 <sub>16</sub>	"	"	Scratchpad Register 2 <sup>1</sup>
	○		"	18 <sub>16</sub>	"	"	Scratchpad Register 3 <sup>1</sup>
	○		"	20 <sub>16</sub>	"	"	Scratchpad Register 4 <sup>1</sup>
	○		"	28 <sub>16</sub>	"	"	Scratchpad Register 5 <sup>1</sup>
	○		"	30 <sub>16</sub>	"	"	Scratchpad Register 6 <sup>1</sup>
	○		"	38 <sub>16</sub>	"	"	Scratchpad Register 7 <sup>1</sup>
21 <sub>16</sub>	○	ASI_MMU_CONTEXTID	RW <sup>2,6</sup>	(decoded; see below)	N	per strand	MMU context registers
	○		"	8 <sub>16</sub>	"	"	I/D MMU Primary Context ID register
	○		"	10 <sub>16</sub>	"	"	I/D MMU Secondary Context ID register
22 <sub>16</sub>	○	ASI_LD_TWIX_AS_IF_USER_PRIMARY (ASI_LDTX_AIUP)	R <sup>2,7,11</sup>	(decoded)	T	—	Primary address space, 128-bit atomic load twin extended word, as if user (nonprivileged)
23 <sub>16</sub>	○	ASI_LD_TWIX_AS_IF_USER_SECONDARY (ASI_LDTX_AIUS)	R <sup>2,7,11</sup>	(decoded)	T	—	Secondary address space, 128-bit atomic load twin extended word, as if user (nonprivileged)
24 <sub>16</sub>	○	—	—	—	—	—	Implementation dependent <sup>1</sup>

**TABLE 10-1** UltraSPARC Architecture ASIs (4 of 13)

ASI Value	req'd (●) opt'l (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	T/ Non-T/ Bypass	Shared /per strand	Description
25 <sub>16</sub>	○	ASI_QUEUE	(see below)	(decoded; see below)	N	per strand	
	○		RW <sup>2,6</sup>	3C0 <sub>16</sub>	"	"	CPU Mondo Queue Head Pointer
	○		RW <sup>2,6,17</sup>	3C8 <sub>16</sub>	"	"	CPU Mondo Queue Tail Pointer
	○		RW <sup>2,6</sup>	3D0 <sub>16</sub>	"	"	Device Mondo Queue Head Pointer
	○		RW <sup>2,6,17</sup>	3D8 <sub>16</sub>	"	"	Device Mondo Queue Tail Pointer
	○		RW <sup>2,6</sup>	3E0 <sub>16</sub>	"	"	Resumable Error Queue Head Pointer
	○		RW <sup>2,6,17</sup>	3E8 <sub>16</sub>	"	"	Resumable Error Queue Tail Pointer
	○		RW <sup>2,6</sup>	3F0 <sub>16</sub>	"	"	Nonresumable Error Queue Head Pointer
	○		RW <sup>2,6,17</sup>	3F8 <sub>16</sub>	"	"	Nonresumable Error Queue Tail Pointer
26 <sub>16</sub>	○	ASI_LD_TWINK_REAL (ASI_LDTX_REAL) ASI_QUAD_LDD_REAL <sup>D†</sup>	R <sup>2,7,11</sup>	(decoded)	B	—	128-bit atomic twin extended-word load from real address
27 <sub>16</sub>	○	ASI_LD_TWINK_NUCLEUS (ASI_LDTX_N)	R <sup>2,7,11</sup>	(decoded)	T	—	Nucleus context, 128-bit atomic load twin extended-word
28 <sub>16</sub> – 29 <sub>16</sub>	○	—	— <sup>2,12</sup>	—	—	—	Implementation dependent <sup>I</sup>
2A <sub>16</sub>	○	ASI_LD_TWINK_AS_IF_USER_PRIMARY_LITTLE (ASI_LDTX_AIUPL)	R <sup>2,7,11</sup>	(decoded)	T	—	Primary address space, 128-bit atomic load twin extended-word, as if user (nonprivileged), little-endian
2B <sub>16</sub>	○	ASI_LD_TWINK_AS_IF_USER_SECONDARY_LITTLE (ASI_LDTX_AIUS_L)	R <sup>2,7,11</sup>	(decoded)	T	—	Secondary address space, 128-bit atomic load twin extended-word, as if user (nonprivileged), little-endian
2C <sub>16</sub>	○	—	— <sup>2</sup>	—	—	—	Implementation dependent <sup>I</sup>
2D <sub>16</sub>	○	—	— <sup>2,12</sup>	—	—	—	Implementation dependent <sup>I</sup>
2E <sub>16</sub>	○	ASI_LD_TWINK_REAL_LITTLE (ASI_LDTX_REAL_L) ASI_QUAD_LDD_REAL_LITTLE <sup>D†</sup>	R <sup>2,7,11</sup>	(decoded)	B	—	128-bit atomic twin-extended-word load from real address, little-endian

**TABLE 10-1** UltraSPARC Architecture ASIs (5 of 13)

ASI Value	req'd(●) opt'l (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	T/ Non-T/ Bypass	Shared /per strand	Description
2F <sub>16</sub>	○	ASI_LD_TWIXX_NUCLEUS_LITTLE (ASI_LDTX_NL)	R <sup>2,7,11</sup>	(decoded)	T	—	Nucleus context, 128-bit atomic load twin extended-word, little-endian
30 <sub>16</sub> – 40 <sub>16</sub>	○	—	— <sub>3,13</sub>	—	—	—	Implementation dependent <sup>1</sup>
3D <sub>16</sub>	○	—	— <sub>3,13</sub>	—	—	—	Implementation dependent <sup>1</sup>
3E <sub>16</sub>	●	—	— <sub>3</sub>	—	—	—	<i>Reserved</i>
3F <sub>16</sub> – 40 <sub>16</sub>	○	—	— <sub>3,13</sub>	—	—	—	Implementation dependent <sup>1</sup>
41 <sub>16</sub>	○	ASI_CMT_SHARED	(see below)	(decoded; see below)	N	shared	CMT control/status (shared)
	○		R <sup>3,6,11</sup>	00 <sub>16</sub>	"	"	Virtual Processor (strand) Available Register
	○		R <sup>3,6,11</sup>	10 <sub>16</sub>	"	"	Virtual Processor (strand) Enable Status Register
	○		RW <sup>3,6</sup>	20 <sub>16</sub>	"	"	Virtual Processor (strand) Enable Register
	○		RW <sup>1,3,6</sup>	30 <sub>16</sub>	"	"	XIR Steering Register Implementation dependent <sup>1</sup> (impl. dep. #1105)
	○		RW <sup>3,6</sup>	50 <sub>16</sub>	"	"	Virtual Processor (strand) Running Register, general access
	○		R <sup>3,6,11</sup>	58 <sub>16</sub>	"	"	Virtual Processor (strand) Running Status Register
	○		W <sup>3,6,10</sup>	60 <sub>16</sub>	"	"	Virtual Processor (strand) Running Register, general access. Write '1' to set bit
	○		W <sup>3,6,10</sup>	68 <sub>16</sub>	"	"	Virtual Processor (strand) Running Register, general access. Write '1' to clear bit
42 <sub>16</sub> – 44 <sub>16</sub>	○	—	— <sub>3,13</sub>	—	—	—	Implementation dependent <sup>1</sup>
45 <sub>16</sub>	○	—	— <sub>3,13</sub>	—	—	—	Implementation dependent <sup>1</sup>
46 <sub>16</sub> – 48 <sub>16</sub>	○	—	— <sub>3,13</sub>	—	—	—	Implementation dependent <sup>1</sup>
49 <sub>16</sub>	○	—	— <sub>3,13</sub>	—	—	—	Implementation dependent <sup>1</sup>

**TABLE 10-1** UltraSPARC Architecture ASIs (6 of 13)

ASI Value	req'd(●) opt'l (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	T/ Non-T/ Bypass	Shared /per strand	Description
4A <sub>16</sub> – 4B <sub>16</sub>	○	—	— <sup>3,13</sup>	—	—	—	Implementation dependent <sup>1</sup>
4C <sub>16</sub>	○	Error Status and Enable Registers					Implementation dependent <sup>1</sup>
4D <sub>16</sub> – 4E <sub>16</sub>	○	—	— <sup>3,13</sup>	—			Implementation dependent <sup>1</sup>
4F <sub>16</sub>	○	ASI_HYP_SCRATCHPAD	RW <sup>3,6</sup>	(decoded; see below)	N	per strand	Hyperprivileged Scratchpad registers; implementation dependent <sup>1</sup>
	○			0 <sub>16</sub>			Hyperprivileged Scratchpad Register 0 <sup>1</sup>
	○			8 <sub>16</sub>			Hyperprivileged Scratchpad Register 1 <sup>1</sup>
	○			10 <sub>16</sub>			Hyperprivileged Scratchpad Register 2 <sup>1</sup>
	○			18 <sub>16</sub>			Hyperprivileged Scratchpad Register 3 <sup>1</sup>
	○			20 <sub>16</sub>			Hyperprivileged Scratchpad Register 4 <sup>1</sup>
	○			28 <sub>16</sub>			Hyperprivileged Scratchpad Register 5 <sup>1</sup>
	○			30 <sub>16</sub>			Hyperprivileged Scratchpad Register 6 <sup>1</sup>
	○			38 <sub>16</sub>			Hyperprivileged Scratchpad Register 7 <sup>1</sup>
50 <sub>16</sub>	○	ASI_IMMU	—	(decoded; see below)	N	per strand	IMMU registers
	○		R <sup>3,6,11</sup>	0 <sub>16</sub>	N	per strand	IMMU tag target register
	○		RW <sup>3,6</sup>	18 <sub>16</sub>	N	per strand	Instruction fault status register (SFSR)
	○		RW <sup>3,6</sup>	30 <sub>16</sub>	N	per strand	I TLB tag access register
52 <sub>16</sub>	○	ASI_MMU_REAL	RW <sup>3,6</sup>	(see below)	N	per strand	MMU registers
	○			108 <sub>16</sub>	"	"	MMU Real Range
	○			110 <sub>16</sub>	"	"	MMU Real Range
	○			118 <sub>16</sub>	"	"	MMU Real Range
	○			120 <sub>16</sub>	"	"	MMU Real Range

**TABLE 10-1** UltraSPARC Architecture ASIs (7 of 13)

ASI Value	req'd (●) opt'l (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	T/ Non-T/ Bypass	Shared /per strand	Description
	○			208 <sub>16</sub>	"	"	MMU Physical Address Offset Registers
	○			210 <sub>16</sub>	"	"	MMU Physical Address Offset Registers
	○			218 <sub>16</sub>	"	"	MMU Physical Address Offset Registers
	○			220 <sub>16</sub>	"	"	MMU Physical Address Offset Registers
53 <sub>16</sub>	○	—	— <sub>3,13</sub>	—	—	—	Implementation dependent <sup>1</sup>

**TABLE 10-1** UltraSPARC Architecture ASIs (8 of 13)

ASI Value	req'd(●) opt'l (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	T/ Non-T/ Bypass	Shared /per strand	Description
54 <sub>16</sub>	○	ASI_MMU	(see below)	(decoded; see below)	N	per strand	(more) MMU registers
	○		W <sup>3,6,10</sup>	0 <sub>16</sub>	"	"	I TLB data in register
	○		RW <sup>3,6</sup>	10 <sub>16</sub>	"	"	Context Zero TSB Configuration register 0
	○		RW <sup>3,6</sup>	18 <sub>16</sub>	"	"	Context Zero TSB Configuration register 1
	○		RW <sup>3,6</sup>	20 <sub>16</sub>	"	"	Context Zero TSB Configuration register 2
	○		RW <sup>3,6</sup>	28 <sub>16</sub>	"	"	Context Zero TSB Configuration register 3
	○		RW <sup>3,6</sup>	30 <sub>16</sub>	"	"	Context Nonzero TSB Configuration register 0
	○		RW <sup>3,6</sup>	38 <sub>16</sub>	"	"	Context Nonzero TSB Configuration register 1
	○		RW <sup>3,6</sup>	40 <sub>16</sub>	"	"	Context Nonzero TSB Configuration register 2
	○		RW <sup>3,6</sup>	48 <sub>16</sub>	"	"	Context Nonzero TSB Configuration register 3
	○		RW <sup>3,6</sup>	50 <sub>16</sub>	"	"	Instruction TSB Pointer register 0
	○		RW <sup>3,6</sup>	58 <sub>16</sub>	"	"	Instruction TSB Pointer register 1
	○		RW <sup>3,6</sup>	60 <sub>16</sub>	"	"	Instruction TSB Pointer register 2
	○		RW <sup>3,6</sup>	68 <sub>16</sub>	"	"	Instruction TSB Pointer register 3
	○		RW <sup>3,6</sup>	70 <sub>16</sub>	"	"	Data/Unified TSB Pointer register 0
	○		RW <sup>3,6</sup>	78 <sub>16</sub>	"	"	Data/Unified TSB Pointer register 1
	○		RW <sup>3,6</sup>	80 <sub>16</sub>	"	"	Data/Unified TSB Pointer register 2
	○		RW <sup>3,6</sup>	88 <sub>16</sub>	"	"	Data/Unified TSB Pointer register 3
	○		RW <sup>3,6</sup>	90 <sub>16</sub>	"	"	Tablewalk Pending Control register
	○		RW <sup>3,6</sup>	98 <sub>16</sub>	"	"	Tablewalk Pending Status register

**TABLE 10-1** UltraSPARC Architecture ASIs (9 of 13)

ASI Value	req'd(●) opt'l (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	T/ Non-T/ Bypass	Shared /per strand	Description
55 <sub>16</sub>	○	ASI_ITLB_DATA_ACCESS_REG	RW <sup>3,6</sup>	0 <sub>16</sub> –3F8 <sub>16</sub> , 800 <sub>16</sub> – 7FFF8 <sub>16</sub>	N	per strand	IMMU TLB data access register
56 <sub>16</sub>	○	ASI_ITLB_TAG_READ_REG	R <sup>3,6,11</sup>	0 <sub>16</sub> – FFFF8 <sub>16</sub>	N	per strand	IMMU TLB tag read register
57 <sub>16</sub>	○	ASI_IMMU_DEMAP	W <sup>3,6,10</sup>	0 <sub>16</sub>	N	per strand	IMMU TLB demap
58 <sub>16</sub>	○	ASI_DMMU /ASI_UMMU	(see below)	(decoded; see below)	N	—	Data or Unified MMU registers
	○		R <sup>3,6,11</sup>	0 <sub>16</sub>	"	per strand	D/U TSB tag target register
	○		RW <sup>3,6</sup>	18 <sub>16</sub>	"	per strand	Data error status register (DSFSR)
	○		R <sup>3,6,11</sup>	20 <sub>16</sub>	"	/core	Data error address register (DSFAR)
	○		RW <sup>3,6</sup>	30 <sub>16</sub>	"	/core	D/U TLB tag access register
	○		RW <sup>3,6</sup>	38 <sub>16</sub>	"	per strand	VA instruction, and PA/VA data watchpoint register
	○		RW <sup>3,6</sup>	40 <sub>16</sub>	"	per strand	I/D/U MMU hardware tablewalk configuration register
	○		RW <sup>3,6</sup>	80 <sub>16</sub>	"	per strand	I/D/U MMU partition ID register
59 <sub>16</sub> – 5B <sub>16</sub>	○	—	— <sup>3,13</sup>	—	—	—	Reserved
5C <sub>16</sub>	○	ASI_DTLB_DATA_IN_REG	W <sup>3,6,10</sup>	0 <sub>16</sub>	N	per strand	D/U TLB data in register
5D <sub>16</sub>	○	ASI_DTLB_DATA_ACCESS_REG	RW <sup>3,6</sup>	0 <sub>16</sub> –3F8 <sub>16</sub> , 800 <sub>16</sub> – 7FFF8 <sub>16</sub>	N	per strand	D/U TLB data access register
5E <sub>16</sub>	○	ASI_DTLB_TAG_READ_REG	R <sup>3,6,11</sup>	0 <sub>16</sub> – FFFF8 <sub>16</sub>	N	per strand	D/U TLB tag read register
5F <sub>16</sub>	○	ASI_DMMU_DEMAP	W <sup>3,6,10</sup>	0 <sub>16</sub>	N	per strand	D/U TLB demap
60 <sub>16</sub> – 62 <sub>16</sub>	○	—	— <sup>3,13</sup>	—	—	—	Implementation dependent <sup>1</sup>
61 <sub>16</sub> – 62 <sub>16</sub>	○	—	— <sup>3,13</sup>	—	—	—	Implementation dependent <sup>1</sup>

**TABLE 10-1** UltraSPARC Architecture ASIs (10 of 13)

ASI Value	req'd (●) opt'l (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	T/ Non-T/ Bypass	Shared /per strand	Description
63 <sub>16</sub>	○	ASI_CMT_PER_STRAND, ASI_CMT_PER_CORE <sup>†</sup>	(see below)	(decoded; see below)	N	per strand	CMT control/status (per strand)
	○		RW <sup>3,6</sup>	00 <sub>16</sub>	"	"	Virtual Processor (strand) Interrupt ID
	○		R <sup>3,6,11</sup>	10 <sub>16</sub>	"	"	Virtual Processor (strand) ID
64 <sub>16</sub> – 67 <sub>16</sub>	○	—	— <sub>3,13</sub>	—	—	—	Implementation dependent <sup>I</sup>
68 <sub>16</sub> – 7F <sub>16</sub>	●	—	— <sub>3,13</sub>	—	—	—	<i>Reserved</i>
80 <sub>16</sub>	●	ASI_PRIMARY (ASI_P)	RW <sup>4</sup>	(decoded)	T	—	Implicit primary address space
81 <sub>16</sub>	●	ASI_SECONDARY (ASI_S)	RW <sup>4</sup>	(decoded)	T	—	Secondary address space
82 <sub>16</sub>	●	ASI_PRIMARY_NO_FAULT (ASI_PNF)	R <sup>9,11</sup>	(decoded)	T	—	Primary address space, no fault
83 <sub>16</sub>	●	ASI_SECONDARY_NO_FAULT (ASI_SNF)	R <sup>9,11</sup>	(decoded)	T	—	Secondary address space, no fault
84 <sub>16</sub> – 87 <sub>16</sub>	●	—	— <sup>16</sup>	—	—	—	<i>Reserved</i>
88 <sub>16</sub>	●	ASI_PRIMARY_LITTLE (ASI_PL)	RW <sup>4</sup>	(decoded)	T	—	Implicit primary address space, little-endian
89 <sub>16</sub>	●	ASI_SECONDARY_LITTLE (ASI_SL)	RW <sup>4</sup>	(decoded)	T	—	Secondary address space, little-endian
8A <sub>16</sub>	●	ASI_PRIMARY_NO_FAULT_LITTLE (ASI_PNFL)	R <sup>9,11</sup>	(decoded)	T	—	Primary address space, no fault, little-endian
8B <sub>16</sub>	●	ASI_SECONDARY_NO_FAULT_LITTLE (ASI_SNFL)	R <sup>9,11</sup>	(decoded)	T	—	Physical address, noncacheable, with side effect, little-endian
8C <sub>16</sub> – BF <sub>16</sub>	●	—	— <sup>16</sup>	—	—	—	<i>Reserved</i>
C0 <sub>16</sub>	○	ASI_PST8_PRIMARY (ASI_PST8_P)	W <sup>8,10,14</sup>	(decoded)	T	—	Primary address space, 8×8-bit partial store
C1 <sub>16</sub>	○	ASI_PST8_SECONDARY (ASI_PST8_S)	W <sup>8,10,14</sup>	(decoded)	T	—	Secondary address space, 8×8-bit partial store
C2 <sub>16</sub>	○	ASI_PST16_PRIMARY (ASI_PST16_P)	W <sup>8,10,14</sup>	(decoded)	T	—	Primary address space, 4×16-bit partial store
C3 <sub>16</sub>	○	ASI_PST16_SECONDARY (ASI_PST16_S)	W <sup>8,10,14</sup>	(decoded)	T	—	Secondary address space, 4×16-bit partial store
C4 <sub>16</sub>	○	ASI_PST32_PRIMARY (ASI_PST32_P)	W <sup>8,10,14</sup>	(decoded)	T	—	Primary address space, 2×32-bit partial store

**TABLE 10-1** UltraSPARC Architecture ASIs (11 of 13)

ASI Value	req'd(●) opt'l (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	T/ Non-T/ Bypass	Shared /per strand	Description
C5 <sub>16</sub>	○	ASI_PST32_SECONDARY (ASI_PST32_S)	W <sup>8,10,14</sup>	(decoded)	T	—	Secondary address space, 2×32-bit partial store
C6 <sub>16</sub> <sup>-</sup> C7 <sub>16</sub>	●	—	— <sup>15</sup>	—	—	—	Implementation dependent <sup>1</sup>
C8 <sub>16</sub>	○	ASI_PST8_PRIMARY_LITTLE (ASI_PST8_PL)	W <sup>8,10,14</sup>	(decoded)	T	—	Primary address space, 8×8-bit partial store, little-endian
C9 <sub>16</sub>	○	ASI_PST8_SECONDARY_LITTLE (ASI_PST8_SL)	W <sup>8,10,14</sup>	(decoded)	T	—	Secondary address space, 8×8-bit partial store, little-endian
CA <sub>16</sub>	○	ASI_PST16_PRIMARY_LITTLE (ASI_PST16_PL)	W <sup>8,10,14</sup>	(decoded)	T	—	Primary address space, 4×16-bit partial store, little-endian
CB <sub>16</sub>	○	ASI_PST16_SECONDARY_LITTLE (ASI_PST16_SL)	W <sup>8,10,14</sup>	(decoded)	T	—	Secondary address space, 4×16-bit partial store, little-endian
CC <sub>16</sub>	○	ASI_PST32_PRIMARY_LITTLE (ASI_PST32_PL)	W <sup>8,10,14</sup>	(decoded)	T	—	Primary address space, 2×32-bit partial store, little-endian
CD <sub>16</sub>	○	ASI_PST32_SECONDARY_LITTLE (ASI_PST32_SL)	W <sup>8,10,14</sup>	(decoded)	T	—	Second address space, 2×32-bit partial store, little-endian
CE <sub>16</sub> <sup>-</sup> CF <sub>16</sub>	●	—	— <sup>15</sup>	—	—	—	Implementation dependent <sup>1</sup>
D0 <sub>16</sub>	○	ASI_FL8_PRIMARY (ASI_FL8_P)	RW <sup>8,14</sup>	(decoded)	T	—	Primary address space, one 8-bit floating-point load/store
D1 <sub>16</sub>	○	ASI_FL8_SECONDARY (ASI_FL8_S)	RW <sup>8,14</sup>	(decoded)	T	—	Second address space, one 8-bit floating-point load/store
D2 <sub>16</sub>	○	ASI_FL16_PRIMARY (ASI_FL16_P)	RW <sup>8,14</sup>	(decoded)	T	—	Primary address space, one 16-bit floating-point load/store
D3 <sub>16</sub>	○	ASI_FL16_SECONDARY (ASI_FL16_S)	RW <sup>8,14</sup>	(decoded)	T	—	Second address space, one 16-bit floating-point load/store
D4 <sub>16</sub> <sup>-</sup> D7 <sub>16</sub>	●	—	— <sup>15</sup>	—	—	—	Implementation dependent <sup>1</sup>
D8 <sub>16</sub>	○	ASI_FL8_PRIMARY_LITTLE (ASI_FL8_PL)	RW <sup>8,14</sup>	(decoded)	T	—	Primary address space, one 8-bit floating point load/store, little-endian
D9 <sub>16</sub>	○	ASI_FL8_SECONDARY_LITTLE (ASI_FL8_SL)	RW <sup>8,14</sup>	(decoded)	T	—	Second address space, one 8-bit floating point load/store, little-endian

**TABLE 10-1** UltraSPARC Architecture ASIs (12 of 13)

ASI Value	req'd(●) opt'l (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	T/ Non-T/ Bypass	Shared /per strand	Description
DA <sub>16</sub>	○	ASI_FL16_PRIMARY_LITTLE (ASI_FL16_PL)	RW <sup>8,14</sup>	(decoded)	T	—	Primary address space, one 16-bit floating-point load/store, little-endian
DB <sub>16</sub>	○	ASI_FL16_SECONDARY_LITTLE (ASI_FL16_SL)	RW <sup>8,14</sup>	(decoded)	T	—	Second address space, one 16-bit floating point load/store, little-endian
DC <sub>16</sub> -DF <sub>16</sub>	●	—	— <sup>15</sup>	—	—	—	Implementation dependent <sup>1</sup>
E0 <sub>16</sub> <sup>-</sup> E1 <sub>16</sub>	●	—	— <sup>15</sup>	—	—	—	<i>Reserved</i>
E2 <sub>16</sub>	○	ASI_LD_TWIXX_PRIMARY (ASI_LDTX_P)	R <sup>19</sup>	(decoded)	T	—	Primary address space, 128-bit atomic load twin extended word
E3 <sub>16</sub>	○	ASI_LD_TWIXX_SECONDARY (ASI_LDTX_S)	R <sup>19</sup>	(decoded)	T	—	Secondary address space, 128-bit atomic load twin extended-word
E4 <sub>16</sub> <sup>-</sup> E9 <sub>16</sub>	●	—	— <sup>15</sup>	—	—	—	Implementation dependent <sup>1</sup>
EA <sub>16</sub>	○	ASI_LD_TWIXX_PRIMARY_LITTLE (ASI_LDTX_PL)	R <sup>19</sup>	(decoded)	T	—	Primary address space, 128-bit atomic load twin extended word, little endian
EB <sub>16</sub>	○	ASI_LD_TWIXX_SECONDARY_LITTLE (ASI_LDTX_SL)	R <sup>19</sup>	(decoded)	T	—	Secondary address space, 128-bit atomic load twin extended word, little endian
EC <sub>16</sub> <sup>-</sup> EF <sub>16</sub>	○	—	— <sup>15</sup>	—	—	—	Implementation dependent <sup>1</sup>
F0 <sub>16</sub>	○	ASI_BLOCK_PRIMARY (ASI_BLK_P)	RW <sup>8,14</sup>	(decoded)	T	—	Primary address space, 8x8-byte block load/store
F1 <sub>16</sub>	○	ASI_BLOCK_SECONDARY (ASI_BLK_S)	RW <sup>8,14</sup>	(decoded)	T	—	Secondary address space, 8x8- byte block load/store
F2 <sub>16</sub> <sup>-</sup> F5 <sub>16</sub>	●	—	— <sup>15</sup>	—	—	—	Implementation dependent <sup>1</sup>
F6 <sub>16</sub> <sup>-</sup> F7 <sub>16</sub>	●	—	—	—	—	—	Implementation dependent <sup>1</sup>
F8 <sub>16</sub>	○	ASI_BLOCK_PRIMARY_LITTLE (ASI_BLK_PL)	RW <sup>8,14</sup>	(decoded)	T	—	Primary address space, 8x8-byte block load/store, little endian

**TABLE 10-1** UltraSPARC Architecture ASIs (13 of 13)

ASI Value	req'd(●) opt'l (○)	ASI Name (and Abbreviation)	Access Type(s)	Virtual Address (VA)	T/ Non-T/ Bypass	Shared /per strand	Description
F9 <sub>16</sub>	○	ASI_BLOCK_SECONDARY_LITTLE (ASI_BLK_SL)	RW <sup>8,14</sup>	(decoded)	T	—	Secondary address space, 8x8- byte block load/store, little endian
FA <sub>16</sub> <sup>-</sup> FD <sub>16</sub>	●	—	— <sup>15</sup>	—	—	—	Implementation dependent <sup>†</sup>
FE <sub>16</sub> <sup>-</sup> FF <sub>16</sub>	●	—	— <sup>15</sup>	—	—	—	Implementation dependent <sup>†</sup>

† 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.

‡ This ASI was named ASI\_DEVICE\_ID+SERIAL\_ID in older documents.

- 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 *data\_access\_exception* 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 *data\_access\_exception* 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 *data\_access\_exception* 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 *data\_access\_exception* 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 *data\_access\_exception* exception.
  - 10 Write(store)-only ASI; an attempted load alternate, atomic alternate, or prefetch alternate instruction to this ASI causes a *data\_access\_exception* exception.
  - 11 Read(load)-only ASI; an attempted store alternate or atomic alternate instruction to this ASI causes a *data\_access\_exception* exception.
  - 12 An attempted load alternate, store alternate, atomic alternate or prefetch alternate instruction to this ASI in privileged mode or hyperprivileged mode causes a *data\_access\_exception* exception.
  - 13 An attempted load alternate, store alternate, atomic alternate or prefetch alternate instruction to this ASI in hyperprivileged mode causes a *data\_access\_exception* exception if this ASI is not implemented by the specific implementation.
  - 14 An attempted access to this ASI may cause an exception (see *Special Memory Access ASIs* on page 422 for details).
  - 15 An attempted load alternate, store alternate, atomic alternate or prefetch alternate instruction to this ASI in any mode causes a *data\_access\_exception* 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 *data\_access\_exception* exception.
  - 17 The Queue Tail Registers (ASI 25<sub>16</sub>) are read-only by privileged software and read-write by hyperprivileged software. An attempted write to the Queue Tail Registers by privileged software causes a *data\_access\_exception* exception.
  - 18 An access to a privileged page (TTE.p = 1) using an ASI\_\*AS\_IF\_USER\* ASI causes a *data\_access\_exception* 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 10<sub>16</sub>, 11<sub>16</sub>, 16<sub>16</sub>, 17<sub>16</sub> and 18<sub>16</sub> (ASI\_\*AS\_IF\_USER\_\*)

These ASI are intended to be used in accesses from privileged and hyperprivileged 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 *data\_access\_exception* (privilege violation) exception occurs
  - Otherwise, the access occurs and its endianness is determined by the current privileged mode and the U/DMMU TTE.ie bit. In hyperprivileged mode, the access is always made in big-endian byte order. In privileged mode, 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 <sub>16</sub>	ASI_AS_IF_USER_PRIMARY (ASI_AIUP)	Virtual (Primary)	In nonprivileged or privileged mode: Big-endian when U/DMMU TTE.ie = 0; little-endian when U/DMMU TTE.ie = 1  In nyperprivileged mode: always big-endian.
11 <sub>16</sub>	ASI_AS_IF_USER_SECONDARY (ASI_AIUS)	Virtual (Secondary)	
16 <sub>16</sub>	ASI_BLOCK_AS_IF_USER_PRIMARY (ASI_BLK_AIUP)	Virtual (Primary)	
17 <sub>16</sub>	ASI_BLOCK_AS_IF_USER_SECONDARY (ASI_BLK_AIUS)	Virtual (Secondary)	

## 10.4.2 ASIs 18<sub>16</sub>, 19<sub>16</sub>, 1E<sub>16</sub>, and 1F<sub>16</sub> (ASI\_\*AS\_IF\_USER\_\*\_LITTLE)

These ASIs are little-endian versions of ASIs 10<sub>16</sub>, 11<sub>16</sub>, 16<sub>16</sub>, and 17<sub>16</sub> (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 and hyperprivileged 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 *data\_access\_exception* (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 <sub>16</sub>	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 <sub>16</sub>	ASI_AS_IF_USER_SECONDARY_LITTLE (ASI_AIUSL)	Virtual (Secondary)	
1E <sub>16</sub>	ASI_BLOCK_AS_IF_USER_PRIMARY_LITTLE (ASI_BLK_AIUP)	Virtual (Primary)	
1F <sub>16</sub>	ASI_BLOCK_AS_IF_USER_SECONDARY_LITTLE (ASI_BLK_AIUSL)	Virtual (Secondary)	

### 10.4.3 ASI 14<sub>16</sub> (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 15<sub>16</sub> (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
- If used with a CASA, CASXA, LDSTUBA, SWAPA, or PREFETCHA instruction, a *data\_access\_exception* exception occurs

- Used with any other load alternate or store alternate instruction, in privileged mode or hyperprivileged 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 1C<sub>16</sub> (ASI\_REAL\_LITTLE)

ASI\_REAL\_LITTLE is a little-endian version of ASI 14<sub>16</sub> (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 1D<sub>16</sub> (ASI\_REAL\_IO\_LITTLE)

ASI\_REAL\_IO\_LITTLE is a little-endian version of ASI 15<sub>16</sub> (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 22<sub>16</sub>, 23<sub>16</sub>, 27<sub>16</sub>, 2A<sub>16</sub>, 2B<sub>16</sub>, 2F<sub>16</sub> (Privileged Load Integer Twin Extended Word)

ASIs 22<sub>16</sub>, 23<sub>16</sub>, 27<sub>16</sub>, 2A<sub>16</sub>, 2B<sub>16</sub> and 2F<sub>16</sub> 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 263). 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 Word / Block Store Init ASIs

ASI	Names	Addressing (Context)	Endianness of Access
22 <sub>16</sub>	ASI_LD_TWINK_AS_IF_USER_PRIMARY (ASI_LDCTX_AIUP)	Virtual (Primary)	Big-endian when U/ DMMU
23 <sub>16</sub>	ASI_LD_TWINK_AS_IF_USER_SECONDARY (ASI_LDCTX_AIUS)	Virtual (Secondary)	TTE.ie = 0; little-endian
27 <sub>16</sub>	ASI_LD_TWINK_NUCLEUS (ASI_LDCTX_N)	Virtual‡ (Nucleus)	when U/ DMMU TTE.ie = 1
2A <sub>16</sub>	ASI_LD_TWINK_AS_IF_USER_PRIMARY_LITTLE (ASI_LDCTX_AIUP_L)	Virtual (Primary)	Little-endian when U/ DMMU
2B <sub>16</sub>	ASI_LD_TWINK_AS_IF_USER_SECONDARY_LITTLE (ASI_LDCTX_AIUS_L)	Virtual (Secondary)	TTE.ie = 0; big-endian
2F <sub>16</sub>	ASI_LD_TWINK_NUCLEUS_LITTLE (ASI_LDCTX_NL)	Virtual‡ (Nucleus)	when U/ DMMU TTE.ie = 1

‡ In hyperprivileged mode, this ASI uses Physical addressing

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 *data\_access\_exception* exception is always generated and *mem\_address\_not\_aligned* is not generated.

**Compatibility Note** | These ASIs replaced ASIs 24<sub>16</sub> and 2C<sub>16</sub> used in earlier UltraSPARC implementations; see the detailed Compatibility Note on page 433 for details.

## 10.4.8 ASIs 26<sub>16</sub> and 2E<sub>16</sub> (Privileged Load Integer Twin Extended Word, Real Addressing)

ASIs 26<sub>16</sub> and 2E<sub>16</sub> 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 263). 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 <sub>16</sub>	ASI_LD_TWIX_REAL (ASI_LDTX_REAL)	Real (—)	Big-endian when U/DMMU TTE.ie = 0; little-endian when U/ DMMU TTE.ie = 1
2E <sub>16</sub>	ASI_LD_TWIX_REAL_LITTLE (ASI_LDTX_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 *data\_access\_exception* exception is always generated and *mem\_address\_not\_aligned* is not generated.

**Compatibility Note** | These ASIs replaced ASIs 34<sub>16</sub> and 3C<sub>16</sub> used in earlier UltraSPARC implementations; see the Compatibility Note on page 433 for details.

## 10.4.9 ASIs 30<sub>16</sub>, 31<sub>16</sub>, 36<sub>16</sub>, 38<sub>16</sub>, 39<sub>16</sub>, 3E<sub>16</sub> (ASI\_AS\_IF\_PRIV\_\*)

These ASI are intended to be used in accesses from hyperprivileged mode, but are processed as if they were issued from privileged mode. These ASIs are distinguished by the context from which the access is made and the endianness of the access, as described in TABLE 10-6.

When one of these ASIs is specified in a load alternate or store alternate instruction, the virtual processor behaves as follows:

- In nonprivileged or privileged mode, a *privileged\_action* exception occurs
- In hyperprivileged mode:
  - 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.

**TABLE 10-6** Hyperprivileged AS\_IF\_PRIV\_\* ASIs

ASI	Names	Addressing (Context)	Endianness of Access
30 <sub>16</sub>	ASI_AS_IF_PRIV_PRIMARY (ASI_AIPP)	Virtual (Primary)	Big-endian when U/DMMU
31 <sub>16</sub>	ASI_AS_IF_PRIV_SECONDARY (ASI_AIPS)	Virtual (Secondary)	TTE.ie = 0; little-endian when U/DMMU
36 <sub>16</sub>	ASI_AS_IF_PRIV_NUCLEUS (ASI_AIPN)	Virtual (Nucleus)	TTE.ie = 1
38 <sub>16</sub>	ASI_AS_IF_PRIV_PRIMARY_LITTLE (ASI_AIPP_L)	Virtual (Primary)	Little-endian when U/DMMU
39 <sub>16</sub>	ASI_AS_IF_PRIV_SECONDARY_LITTLE (ASI_AIPS_L)	Virtual (Secondary)	TTE.ie = 0; big-endian when U/DMMU
3E <sub>16</sub>	ASI_AS_IF_PRIV_NUCLEUS_LITTLE (ASI_AIPN_L)	Virtual (Nucleus)	TTE.ie = 1

## 10.4.10 ASIs E2<sub>16</sub>, E3<sub>16</sub>, EA<sub>16</sub>, EB<sub>16</sub> (Nonprivileged Load Integer Twin Extended Word)

ASIs E2<sub>16</sub>, E3<sub>16</sub>, EA<sub>16</sub>, and EB<sub>16</sub> 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 263). These ASIs are distinguished by the address space accessed (Primary or Secondary) and the endianness of the access, as described in TABLE 10-7.

**TABLE 10-7** Load Integer Twin Extended Word ASIs

ASI	Names	Addressing (Context)	Endianness of Access
E2 <sub>16</sub>	ASI_LD_TWIXX_PRIMARY (ASI_LDTX_P)	Virtual (Primary)	Big-endian when U/DMMU
E3 <sub>16</sub>	ASI_LD_TWIXX_SECONDARY (ASI_LDTX_S)	Virtual (Secondary)	TTE.ie = 0, little-endian when U/DMMU TTE.ie = 1
EA <sub>16</sub>	ASI_LD_TWIXX_PRIMARY_LITTLE (ASI_LDTX_PL)	Virtual (Primary)	Little-endian when U/DMMU
EB <sub>16</sub>	ASI_LD_TWIXX_SECONDARY_LITTLE (ASI_LDTX_SL)	Virtual (Secondary)	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 *data\_access\_exception* exception is always generated and *mem\_address\_not\_aligned* is not generated.

## 10.4.11 Block Load and Store ASIs

ASIs 16<sub>16</sub>, 17<sub>16</sub>, 1E<sub>16</sub>, 1F<sub>16</sub>, F0<sub>16</sub>, F1<sub>16</sub>, F8<sub>16</sub>, and F9<sub>16</sub> exist for use with LDDFA and STDFA instructions as Block Load (LDBLOCKF) and Block Store (STBLOCKF) operations (see *Block Load* on page 245 and *Block Store* on page 328).

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.

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 *data\_access\_exception* exception is always generated and *mem\_address\_not\_aligned* is not generated.

## 10.4.12 Partial Store ASIs

ASIs C0<sub>16</sub>–C5<sub>16</sub> and C8<sub>16</sub>–CD<sub>16</sub> exist for use with the STDFA instruction as Partial Store (STPARTIALF) operations (see *Store Partial Floating-Point* on page 341).

When these ASIs are used with STDFA 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 STDFA, a Load Alternate, Store Alternate, Atomic Load-Store Alternate, or PREFETCHA instruction, a *data\_access\_exception* exception is generated and *mem\_address\_not\_aligned*, *LDDF\_mem\_address\_not\_aligned*, and *illegal\_instruction* (for *i* = 1) are not generated.

ASIs C0<sub>16</sub>–C5<sub>16</sub> and C8<sub>16</sub>–CD<sub>16</sub> are only defined for use in Partial Store operations (see page 341). 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 254.

## 10.4.13 Short Floating-Point Load and Store ASIs

ASIs D0<sub>16</sub>–D3<sub>16</sub> and D8<sub>16</sub>–DB<sub>16</sub> exist for use with the LDDFA and STDFA instructions as Short Floating-point Load and Store operations (see *Load Floating-Point* on page 249 and *Store Floating-Point* on page 332).

When ASI D2<sub>16</sub>, D3<sub>16</sub>, DA<sub>16</sub>, or DB<sub>16</sub> is used with LDDFA (STDFA) 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 *data\_access\_exception* exception is always generated and *mem\_address\_not\_aligned* is not generated.

---

## 10.5 ASI-Accessible Registers

In this section the Data Watchpoint registers, scratchpad registers, and CMT registers are described.

A list of UltraSPARC Architecture 2005 ASIs is shown in TABLE 10-1 on page 409.

## 10.5.1 Privileged Scratchpad Registers (ASI\_SCRATCHPAD) D1

An UltraSPARC 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.

For conventional uses of Scratchpad registers, see “Scratchpad Register Usage” in *Software Considerations*, contained in the separate volume *UltraSPARC Architecture Application Notes*.

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 10-8.

**TABLE 10-8** Scratchpad Registers

Assembly Language ASI Name	ASI #	Virtual Address	Privileged Scratchpad Register #
		00 <sub>16</sub>	0
		08 <sub>16</sub>	1
		10 <sub>16</sub>	2
		18 <sub>16</sub>	3
ASI_SCRATCHPAD	20 <sub>16</sub>	20 <sub>16</sub>	4
		28 <sub>16</sub>	5
		30 <sub>16</sub>	6
		38 <sub>16</sub>	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 (emulated by *data\_access\_exception*trap to hyperprivileged software), or
- (3) inaccessible (cause a *data\_access\_exception*).

**V9 Compatibility Note** | Privileged scratchpad registers are an UltraSPARC Architecture extension to SPARC V9.

## 10.5.2 Hyperprivileged Scratchpad Registers (ASI\_HYP\_SCRATCHPAD) (D2)

An UltraSPARC Architecture virtual processor includes eight hyperprivileged Scratchpad registers (64 bits each, read/write accessible). The use of the hyperprivileged Scratchpad registers is completely defined by software.

The hyperprivileged Scratchpad registers are intended to be used in hyperprivileged trap handler code.

The hyperprivileged Scratchpad registers are accessed with Load Alternate and Store Alternate instructions, using the ASIs and addresses listed in TABLE 10-9.

**IMPL. DEP. #407-S10:** It is implementation dependent whether any of the hyperprivileged Scratchpad registers are aliased to the corresponding privileged Scratchpad register or is an independent register.

**TABLE 10-9** Hyperprivileged Scratchpad Registers

Assembly Language	ASI Name	ASI #	Virtual Address	Hyperprivileged Scratchpad Register #
			00 <sub>16</sub>	0
			08 <sub>16</sub>	1
			10 <sub>16</sub>	2
			18 <sub>16</sub>	3
ASI_HYP_SCRATCHPAD		4F <sub>16</sub>	20 <sub>16</sub>	4
			28 <sub>16</sub>	5
			30 <sub>16</sub>	6
			38 <sub>16</sub>	7

**V9 Compatibility** | Hyperprivileged Scratchpad registers are an UltraSPARC  
**Note** | Architecture extension to SPARC V9.

## 10.5.3 CMT Registers Accessed Through ASIs (D2)

All chip-level multithreading (CMT) registers are accessed through ASIs. See *Accessing CMT Registers* on page 515, for descriptions of ASI registers used to control CMT functions.

## 10.5.4 ASI Changes in the UltraSPARC Architecture

The following Compatibility Notes summarize the UltraSPARC ASI changes in UltraSPARC Architecture.

**Compatibility Note** | The names of several ASIs used in earlier UltraSPARC implementations have changed in UltraSPARC Architecture. Their functions have not changed; just their names have changed.

<u>ASI#</u>	<u>Previous UltraSPARC</u>	<u>UltraSPARC Architecture</u>
14 <sub>16</sub>	ASI_PHYS_USE_EC	ASI_REAL
15 <sub>16</sub>	ASI_PHYS_BYPASS_EC_WITH_EBIT	ASI_REAL_IO
1C <sub>16</sub>	ASI_PHYS_USE_EC_LITTLE (ASI_PHYS_USE_EC_L)	ASI_REAL_LITTLE
1D <sub>16</sub>	ASI_PHYS_BYPASS_EC_WITH_EBIT_LITTLE (ASI_PHY_BYPASS_EC_WITH_EBIT_L)	ASI_REAL_IO_LITTLE

**Compatibility Note** | The names *and* ASI assignments (but not functions) changed between earlier UltraSPARC implementations and UltraSPARC Architecture, for the following ASIs:

<u>Previous UltraSPARC</u>		<u>UltraSPARC Architecture</u>	
<u>ASI#</u>	<u>Name</u>	<u>ASI#</u>	<u>Name</u>
24 <sub>16</sub>	ASI_NUCLEUS_QUAD_LDD	27 <sub>16</sub>	ASI_LD_TWIX_NUCLEUS (ASI_LDTX_N)
2C <sub>16</sub>	ASI_NUCLEUS_QUAD_LDD_LITTLE (ASI_NUCLEUS_QUAD_LDD_L)	2F <sub>16</sub>	ASI_LD_TWIX_NUCLEUS_LITTLE (ASI_LDTX_NL)
34 <sub>16</sub>	ASI_QUAD_LDD_PHYS	26 <sub>16</sub>	ASI_LD_TWIX_REAL (ASI_LDTX_REAL)
3C <sub>16</sub>	ASI_QUAD_LDD_LITTLE (ASI_QUAD_LDD_L)	2E <sub>16</sub>	ASI_LD_TWIX_REAL_LITTLE (ASI_LDTX_REAL_L)



# Performance Instrumentation

---

(contents to be supplied in a later revision)



# Traps

---

A *trap* is a vectored transfer of control to software running in a privilege mode (see page 438) 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. A trap that occurs while executing in hyperprivileged mode can only be delivered to hyperprivileged mode.

The actual transfer of control occurs through a trap table that contains the first eight instructions (32 instructions for *clean\_window*, *fast\_instruction\_access\_MMU\_miss*, *fast\_data\_access\_MMU\_miss*, *fast\_data\_access\_protection*, 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 physical base address of the trap table for traps to be delivered in hyperprivileged mode is specified in the Hyperprivileged Trap Base Address (HTBA) register. The displacement within either 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, or enter hyperprivileged mode with a predefined PSTATE and HPSTATE.
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.

<b>V9 Compatibility Note</b>	Exceptions referred to as “catastrophic error exceptions” in the SPARC V9 specification do not exist in the UltraSPARC 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 438.
- **Virtual Processor States, Normal Traps, and RED\_state Traps** on page 440.
- **Trap Categories** on page 445.
- **Trap Control** on page 451.
- **Trap-Table Entry Addresses** on page 452.
- **Trap Processing** on page 470.
- **Exception and Interrupt Descriptions** on page 481.
- **Register Window Traps** on page 491.

---

## 12.1 Virtual Processor Privilege Modes

An UltraSPARC 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, in which hypervisor software operates, serving as a layer between the supervisor software and the underlying virtual processor

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

HPSTATE.hpriv	PSTATE.priv	Virtual Processor Privilege Mode
0	0	Nonprivileged
0	1	Privileged
1	—	Hyperprivileged

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. Traps detected in hyperprivileged mode are either delivered to the virtual processor in hyperprivileged mode or may be masked. If masked, they are held pending.

TABLE 12-4 on page 459 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 453 for details. A trap delivered to hyperprivileged mode uses the hyperprivileged mode trap vector address, based upon the HTBA register. See *Trap-Table Entry Address to Hyperprivileged Mode* on page 454 for details.

The maximum trap level at which privileged software may execute is *MAXPTL* (which, on a virtual processor, is 2). Therefore, if  $TL \geq MAXPTL$  and a trap occurs that would normally be delivered in privileged mode, it is instead delivered in hyperprivileged mode; the trap table offset for *watchdog\_reset* ( $40_{16}$ ) is used, and the priority and trap type of the original exception is retained. This is referred to as a “*guest\_watchdog*” trap (so named because it uses *watchdog\_reset*'s trap table offset).

- Notes**
- Execution in nonprivileged or privileged mode with  $TL > MAXPTL$  is an invalid condition that hyperprivileged software should never allow to occur.
  - Execution in nonprivileged mode with  $TL > 0$  is an invalid condition that privileged and hyperprivileged 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 or HPSTATE.hpriv. In this figure, **PT** indicates a “trap destined for privileged mode” and **HT** indicates a “trap destined for hyperprivileged mode”.




FIGURE 12-1 Virtual Processor Privilege Mode Transition Diagram

## 12.2 Virtual Processor States, Normal Traps, and RED\_state Traps

An UltraSPARC Architecture virtual processor is always in one of three discrete states:

- `execute_state`, which is the normal execution state of the virtual processor
- `RED_state` (**R**eset, **E**rror, and **D**ebug state), which is a restricted execution state reserved for processing traps that occur when  $TL = MAXTL - 1$ , and for processing hardware- and software-initiated resets
- `error_state`, which is a transient state that is entered as a result of a non-reset trap, SIR, or XIR when  $TL = MAXTL$

The values of TL and HPSTATE.red affect the generated trap vector address. TL also determines where (that is, into which element of the TSTATE and HTSTATE arrays) the states are saved..

Traps processed in `execute_state` are called *normal traps*. Traps processed in `RED_state` are called *RED\_state traps*.

**V9 Compatibility Note** | `RED_state` traps were called “special traps” in the SPARC V9 specification. The name was changed to clarify the terminology.

FIGURE 12-2 shows the virtual processor state transition diagram.




FIGURE 12-2 Virtual Processor State Diagram (“NRT” = “non-reset trap”)

## 12.2.1 RED\_state

`RED_state` is an acronym for **R**eset, **E**rror, and **D**ebug state. The virtual processor enters `RED_state` under any one of the following conditions:

- A non-reset trap is taken when  $TL = MAXTL - 1$ .
- A POR or WDR reset occurs.
- An SIR reset occurs when  $TL < MAXTL$ .
- An XIR reset occurs when  $TL < MAXTL$ .
- System software sets `HPSTATE.red = 1`. For this condition, no other machine state or operation is modified as a side-effect of the write to `HPSTATE`; software must set the appropriate machine state.

`RED_state` serves two purposes:

- During trap processing, it indicates that no more trap levels are available; that is, while executing in `RED_state` with  $TL = MAXTL$ , if another nested non-reset trap, SIR, or XIR is taken, the virtual processor will enter `error_state`. `RED_state` provides system software with a restricted execution environment.

- It provides the execution environment for all reset processing.

`RED_state` is indicated by `HPSTATE.red`. When this bit is set to 1, the virtual processor is in `RED_state`; when this bit is zero, the virtual processor is not in `RED_state`, independent of the value of `TL`. Executing a `DONE` or `RETRY` instruction in `RED_state` restores the stacked copy of the `HPSTATE` register, which zeroes the `HPSTATE.red` flag if it was zero in the stacked copy. System software can also directly write 1 or 0 to `HPSTATE.red` with a `WRHPR` instruction, which forces the virtual processor to enter or exit `RED_state`, respectively. In this case, the `WRHPR` instruction should be placed in the delay slot of a jump instruction so that the `PC` can be changed in concert with the state change.

When `RED_state` is entered due to a reset or a trap, the execution environment is altered in four ways:

- Address translation is disabled in the MMU, for both instruction and data references.
- Watchpoints are disabled.
- The trap vector for the traps occurring in `RED_state` is based on the `RED_state` Trap Table.
- The virtual processor enters hyperprivileged mode (`HPSTATE.hpriv` ← 1).

<b>Programming Note</b>	Setting <code>TL</code> ← <code>MAXTL</code> with a <code>WRHPR</code> instruction does not also set <code>HPSTATE.red</code> ← 1, nor does it alter any other machine state. The values of <code>HPSTATE.red</code> and <code>TL</code> are independent.
	Setting <code>HPSTATE.red</code> with a <code>WRHPR</code> instruction causes the virtual processor to execute in <code>RED_state</code> . This results in the execution environment defined in <i>RED_state Execution Environment</i> on page 442. However, it is different from a <code>RED_state</code> trap in the sense that there are no trap-related changes in the machine state (for example, <code>TL</code> does not change).

The trap table organization for `RED_state` traps is described in *RED\_state Trap Table Organization* on page 456.

### 12.2.1.1 `RED_state` Execution Environment

In `RED_state`, the virtual processor is forced to execute in a restricted environment by overriding the values of some virtual processor control and state registers.

The values are overridden, not set, allowing them to be switched atomically.

Some of the characteristics of `RED_state` include:

- Memory accesses in `RED_state` are by default noncacheable, so there must be noncacheable scratch memory somewhere in the system.
- The D-cache watchpoints and DMMU/UMMU can be enabled by software in `RED_state`, but any trap will disable them again.

- The caches continue to snoop and maintain coherence in RED\_state if DMA or other virtual processors are still issuing cacheable accesses.

**IMPL. DEP. #115-V9:** A processor's behavior in RED\_state is implementation dependent.

<b>Programming Note</b>	When RED_state is entered because of component failures, trap handler software should attempt to recover from potentially fatal error conditions or to disable the failing components. When RED_state is entered after a reset, the software should create the environment necessary to restore the system to a running state.
-------------------------	--

### 12.2.1.2 RED\_state Entry Traps

The following reset traps are processed in RED\_state:

- **Power-on reset (POR)** — POR causes the virtual processor to start execution at this trap table entry.
- **Watchdog reset (WDR)** — While in error\_state, the virtual processor automatically invokes a watchdog reset to enter RED\_state (impl. dep. #254-U3-Cs10).
- **Externally initiated reset (XIR)** — This trap is typically used as a nonmaskable interrupt for debugging purposes. If  $TL < MAXTL$  when an XIR occurs, the reset trap is processed in RED\_state; if  $TL = MAXTL$  when an XIR occurs, the virtual processor transitions directly to error\_state.
- **Software-initiated reset (SIR)** — If  $TL < MAXTL$  when an SIR occurs, the reset trap is processed in RED\_state; if  $TL = MAXTL$  when an SIR occurs, the virtual processor transitions directly to error\_state.

Non-reset traps that occur when  $TL = MAXTL - 1$  also set `HPSTATE.red = 1`; that is, any non-reset trap handler entered with  $TL = MAXTL$  runs in RED\_state.

Any non-reset trap that sets `HPSTATE.red = 1` or that occurs when `HPSTATE.red = 1` branches to a special entry in the RED\_state trap vector at  $RSTVADDR + A0_{16}$ . Reset traps are described in *Reset Traps* on page 449.

### 12.2.1.3 RED\_state Software Considerations

In effect, `RED_state` reserves one level of the trap stack for recovery and reset processing. Hyperprivileged software should be designed to require only `MAXTL - 1` trap levels for normal processing. That is, any trap that causes `TL = MAXTL` is an exceptional condition that should cause entry to `RED_state`.

**Programming Note** To log the state of the virtual processor, `RED_state`-handler software needs either a spare register or a preloaded pointer to a save area. To support recovery, the operating system might reserve one of the hyperprivileged scratchpad registers for use in `RED_state`.

### 12.2.1.4 Usage of Trap Levels

If `MAXPTL = 2` and `MAXTL = 5` in an UltraSPARC 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
3	Hyperprivileged	Real address TLB miss handler
4	Hyperprivileged	Reserved for error handling
5	Hyperprivileged	<code>RED_state</code> handler

## 12.2.2 error\_state

The virtual processor enters `error_state` when a trap occurs while the virtual processor is already at its maximum supported trap level — that is, it enters `error_state` when a trap occurs while `TL = MAXTL`. No other conditions cause entry into `error_state` on an UltraSPARC Architecture virtual processor. (impl. dep. #39-V8-Cs10)

**IMPL. DEP. #40-V8:** Effects when `error_state` is entered are implementation-dependent, but it is recommended that as much processor state as possible be preserved upon entry to `error_state`. In addition, an UltraSPARC Architecture virtual processor may have other `error_state` entry traps that are implementation dependent.

Upon entering `error_state`, a virtual processor automatically generates a `watchdog_reset` (WDR) (impl. dep. #254-U3-Cs10), which causes entry into `RED_state`.

---

## 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.

<b>Programming Note</b>	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 UltraSPARC Architecture 2005 implementations) or in hyperprivileged mode if  $TL \geq MAXPTL$ .

**Outstanding in Hyperprivileged mode, destined for delivery in Privileged mode.** An outstanding disrupting trap condition detected while in hyperprivileged mode and destined for delivery in privileged mode is held pending while in hyperprivileged mode (HPSTATE.priv = 1), regardless of the values of TL and PSTATE.ie.

**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.

**Outstanding in Hyperprivileged mode, destined for delivery in Hyperprivileged mode.** An outstanding disrupting trap condition detected in hyperprivileged mode and destined to be delivered in hyperprivileged mode is masked and held pending while `PSTATE.ie = 0`.

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 <code>PSTATE.ie = 0</code> or <code>interrupt level ≤ PIL</code>	—
	Hyperprivileged	Held pending while <code>HPSTATE.hpriv = 1</code>	—
All other disrupting traps	Nonprivileged or Privileged	Held pending while <code>PSTATE.ie = 0</code>	Delivered immediately
	Hyperprivileged	Held pending while <code>HPSTATE.hpriv = 1</code>	Held pending while <code>PSTATE.ie = 0</code>

### 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 Reset Traps

A *reset trap* occurs when hyperprivileged software or the implementation's hardware determines that the machine must be reset to a known state. Reset traps differ from disrupting traps in that:

- They are not maskable.

- Trap handler software for resets is generally not expected to resume execution of the program that was running when the reset trap occurred. After an SIR or XIR, execution of the interrupted program may resume.

All *reset traps* are delivered to the virtual processor in hyperprivileged mode.

**IMPL. DEP. #37-V8:** Some of a virtual processor's behavior during a reset trap is implementation dependent. See *RED\_state Trap Processing* on page 474 for details.

The following reset traps are defined by the SPARC V9 architecture:

- **Power-on reset (POR)** — Used for initialization purposes (for example, when power is applied or reapplied to the virtual processor).
- **Watchdog reset (WDR)** — Initiated when the virtual processor enters *error\_state* (impl. dep. #254-U3-Cs10). The WDR reset trap is taken instead of the trap request that caused entry to *error\_state* at  $TL = MAXTL$ .  $TSTATE[MAXTL]$ ,  $TPC[MAXTL]$ ,  $TNPC[MAXTL]$  and  $TT[MAXTL]$  observed after a WDR reset trap are those associated with the trap request that caused entry to *error\_state*. The value of  $TT[MAXTL]$  indicates the trap type of this trap. Machine state is consistent; however, software should not resume normal instruction processing at the address in  $TPC[TL]$  after the WDR reset trap. The values in  $TSTATE[MAXTL]$ ,  $TPC[MAXTL]$ ,  $TNPC[MAXTL]$  and  $TT[MAXTL]$  are accurate and are intended for debug purposes.
- **Externally initiated reset (XIR)** — Initiated in response to a signal or event that is external to the virtual processor. This reset trap is normally used for critical system events, such as power failure. The XIR reset trap is treated as an interrupt and processed similarly to a disrupting trap (but without masking). Software can resume the interrupted program at the conclusion of trap handler execution.
- **Software-initiated reset (SIR)** — Initiated by software by executing the SIR instruction in hyperprivileged mode. In nonprivileged and privileged mode, the SIR instruction causes an *illegal\_instruction* exception (which results in a trap to hyperprivileged mode). The SIR reset trap is processed similar to a precise trap. The PC saved in  $TPC[TL]$  points to the SIR instruction. If the SIR reset is detected when  $TL =$ , the processor enters *error\_state* and triggers a WDR reset.

## 12.3.5 Uses of the Trap Categories

The SPARC V9 *trap model* stipulates the following:

1. Reset traps (except *software\_initiated\_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 481 for identification of which traps are precise.

3. In an UltraSPARC Architecture implementation, all exceptions that occur as the result of program execution, except for errors on store instructions that occur after the store instruction that has passed the retirement point, 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) 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 hyperprivileged mode bit (*hpriv*) field in the *HPSTATE* register, which can affect how a trap is delivered. See *Conditioning of Disrupting Traps* on page 448 for details.
- The *TL* register, which contains the current level of trap nesting, controls whether a trap causes entry to *execute\_state*, *RED\_state*, or *error\_state*. It also affects whether the trap is processed in privileged mode or hyperprivileged mode.
- For a trap delivered to the virtual processor in privileged mode, *PSTATE.tle* determines whether implicit data accesses in the trap handler routine will be performed using big-endian or little-endian byte order. A trap delivered to the virtual processor in hyperprivileged mode always uses a default byte order of big-endian.

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 469.

## 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 459 and TABLE 12-5 on page 465 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  $2^5$  or 32 bytes long. Each entry in the trap table may contain the first eight instructions of the corresponding trap handler.

FIGURE 12-3 illustrates the trap vector address for a trap delivered to privileged mode. In FIGURE 12-3, 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 12-3 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 12-4.

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 <sub>16</sub> –07F <sub>16</sub>	0 <sub>16</sub> – FE0 <sub>16</sub>	Hardware traps
	—	080 <sub>16</sub> –0FF <sub>16</sub>	1000 <sub>16</sub> –1FE0 <sub>16</sub>	<i>Spill / fill traps</i>
	0 <sub>16</sub> – 7F <sub>16</sub>	100 <sub>16</sub> –17F <sub>16</sub>	2000 <sub>16</sub> –2FE0 <sub>16</sub>	Software traps to Privileged level
	—	180 <sub>16</sub> –1FF <sub>16</sub>	3000 <sub>16</sub> –3FE0 <sub>16</sub>	<i>unassigned</i>
TL = 1 (TL = MAXPTL–1)	—	000 <sub>16</sub> –07F <sub>16</sub>	4000 <sub>16</sub> –4FE0 <sub>16</sub>	Hardware traps
	—	080 <sub>16</sub> –0FF <sub>16</sub>	5000 <sub>16</sub> –5FE0 <sub>16</sub>	<i>Spill / fill traps</i>
	0 <sub>16</sub> – 7F <sub>16</sub>	100 <sub>16</sub> –17F <sub>16</sub>	6000 <sub>16</sub> –6FE0 <sub>16</sub>	Software traps to Privileged level
	—	180 <sub>16</sub> –1FF <sub>16</sub>	7000 <sub>16</sub> –7FE0 <sub>16</sub>	<i>unassigned</i>

FIGURE 12-4 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-Table Entry Address to Hyperprivileged Mode

Hyperprivileged software initializes bits 63:14 of the Hyperprivileged Trap Base Address (HTBA) register (its most significant 50 bits) with bits 63:14 of the desired 64-bit hyperprivileged trap table base address.

At the time a trap to hyperprivileged mode is taken:

- Bits 63:14 of the trap vector address are taken from HTBA{63:14}.
- 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  $2^5$  or 32 bytes long. Each entry in the trap table may contain the first eight instructions of the corresponding trap handler.

FIGURE 12-5 illustrates the trap vector address used for a trap delivered to hyperprivileged mode.




FIGURE 12-5 Hyperprivileged Mode Trap Vector Address

## 12.5.4 Hyperprivileged Trap Table Organization

The layout of the hyperprivileged-mode trap table (which is accessed using physical addresses) is illustrated in FIGURE 12-6.

Software Trap Type	Hardware Trap Type (TT[TL])	Trap Table Offset (from HTBA)	Contents of Trap Table
—	$000_{16} - 07F_{16}$	$0_{16} - FE0_{16}$	Hardware traps
—	$080_{16} - 0FF_{16}$	$1000_{16} - 1FE0_{16}$	<i>Spill / fill traps</i>
$0_{16} - 7F_{16}$	$100_{16} - 17F_{16}$	$2000_{16} - 2FE0_{16}$	Software traps from hyperprivileged level to hyperprivileged level
$80_{16} - FF_{16}$	$180_{16} - 1FF_{16}$	$3000_{16} - 3FE0_{16}$	Software traps to hyperprivileged level

FIGURE 12-6 Hyperprivileged-mode Trap Table Layout

The hyperprivileged trap table comprises 512 thirty-two-byte entries. Therefore, the total size of a full hyperprivileged trap table is  $512 \times 32$  bytes (16 Kbytes).

## 12.5.5 Trap Table Entry Address to RED\_state

Traps occurring in RED\_state or traps that cause the virtual processor to enter RED\_state use an abbreviated trap vector, called the RED\_state trap vector.

The RED\_state trap vector is located at the following address, referred to as RSTVADDR (impl. dep. #114-V9-Cs10):

Physical Address       $RSTVADDR = FFFF\ FFFF\ F000\ 0000_{16}$   
(highest 256 MB of physical address space)

In an implementation that implements fewer than 64 bits of physical addressing, unimplemented high-order bits of the above RSTVADDR are ignored.

FIGURE 12-7 illustrates the trap vector address used for a trap delivered to RED\_state (in hyperprivileged mode).




FIGURE 12-7 RED\_state Trap Vector Address

## 12.5.6 RED\_state Trap Table Organization

The RED\_state trap table is constructed so that it can overlay the hyperprivileged trap table (see FIGURE 12-6) if necessary. For a trap to RED\_state, the trap table offset is added to the base address contained in RSTVADDR to yield the RED\_state trap vector. FIGURE 12-8 illustrates the layout of the RED\_state trap table.

Hardware Trap Type (TT[TL])	Trap Table Offset (from RSTVADDR)	Contents of Trap Table
0	00 <sub>16</sub>	<i>Reserved</i>
1	20 <sub>16</sub>	Power-on reset (POR)
TT <sup>†</sup>	40 <sub>16</sub>	Watchdog reset (WDR)
3 or TT <sup>‡</sup>	60 <sub>16</sub>	Externally initiated reset (XIR)
4	80 <sub>16</sub>	Software-initiated reset (SIR)
TT <sup>*</sup>	A0 <sub>16</sub>	All other exceptions in RED_state

- † TT = trap type of the exception that caused entry into error\_state  
 TT = 3 if an *externally\_initiated\_reset* (XIR) occurs while the virtual processor is not in error\_state; TT = trap type of the exception that caused entry into error\_state if the externally initiated reset occurs in error\_state.
- \* TT = trap type of the exception. See TABLE 12-4 on page 459.

FIGURE 12-8 RED\_state Trap Table Layout

## 12.5.7 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 one of the following, depending on the trap's destination privilege mode:

- The TBA register, (TL > 0), and TT[TL] (see *Trap-Table Entry Address to Privileged Mode* on page 453)
- The HTBA register and TT[TL] (see *Trap-Table Entry Address to Hyperprivileged Mode* on page 454)

**Programming Note** | The *spill\_n\_\**, *fill\_n\_\**, *clean\_window*, and MMU-related traps (*fast\_instruction\_access\_MMU\_miss*, *fast\_data\_access\_MMU\_miss*, and *fast\_data\_access\_protection*) are spaced such that their trap-table entries are 128 bytes (32 instructions) long in the UltraSPARC Architecture. This length allows the complete code for one spill/fill routine, a *clean\_window* routine, or a normal MMU miss handling routine to reside in one trap-table entry.

When a *RED\_state* trap occurs, the TT register is set as described in *RED\_state* on page 441. Control is then transferred into the *RED\_state* trap table at an address formed by *RSTVADDR* and an offset depending on the condition.

TT values  $000_{16}$ – $0FF_{16}$  are reserved for hardware traps. TT values  $100_{16}$ – $17F_{16}$  are reserved for software traps (caused by execution of a Tcc instruction) to privileged-mode trap handlers. TT values  $180_{16}$ – $1FF_{16}$  are used for software traps to trap handlers operating in hyperprivileged mode.

**IMPL. DEP. #35-V8-Cs20:** TT values  $060_{16}$  to  $07F_{16}$  were reserved for *implementation\_dependent\_exception\_n* exceptions in the SPARC V9 specification, but are now all defined as standard UltraSPARC 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 UltraSPARC Architecture 2005. 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 UltraSPARC Architecture 2005, but its implementation is optional.
P	Trap is taken via the Privileged trap table, in Privileged mode (PSTATE.priv = 1)
H	Trap is taken via the Hyperprivileged trap table, in Hyperprivileged mode (HSTATE.hpriv = 1)
H <sup>U</sup>	Trap is taken via the Hyperprivileged trap table, in Hyperprivileged mode (HSTATE.hpriv = 1). However, the trap is <i>unexpected</i> . While hardware can legitimately generate this trap, it should not do so unless there is a programming error or some other error. Therefore, occurrence of this trap indicates an actual error to hyperprivileged software.
-x-	Not possible. Hardware cannot generate this trap in the indicated running mode. For example, all privileged instructions can be executed in both privileged and hyperprivileged modes, therefore a <i>privileged_opcode</i> trap cannot occur in privileged or hyperprivileged mode.
—	This trap is reserved for future use.
(am)	Always Masked — when the condition occurs in this privilege mode, it is always masked out (but remains pending).
(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 occur in this running mode, but can't be serviced 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 6)

UA-2005 ●=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	HP
—	<i>Reserved</i>	000 <sub>16</sub>	—	—	—	—	—
●	<i>power_on_reset</i>	001 <sub>16</sub>	reset	0	H (nm)	H (nm)	H (nm)
●	<i>watchdog_reset</i>	TT <sup>▲</sup>	reset	1.2	H (nm)	H (nm)	H (nm)
●	<i>externally_initiated_reset</i>	003 <sub>16</sub>	reset	1.1	H (nm)	H (nm)	H (nm)
●	<i>software_initiated_reset</i>	004 <sub>16</sub>	reset	1.3	-x-	-x-	H (nm)
—	<i>Reserved</i>	005 <sub>16</sub>	—	—	—	—	—
●	<i>RED_state_exception</i>	TT <sup>♣</sup>	precise	♣	H (nm)	H (nm)	H (nm)
—	<i>implementation-dependent</i>	006 <sub>16</sub>	—	—	—	—	—
○	<i>store_error</i>	007 <sub>16</sub>	deferred	2.01	H (nm)	H (nm)	H (nm)
●	<i>instruction_access_exception</i>	008 <sub>16</sub>	precise	3	H (nm)	H (nm)	H <sup>U</sup> (nm)
●	<i>instruction_access_MMU_miss<sup>†</sup></i>	009 <sub>16</sub>	precise	2.08	H (nm)	H (nm)	-x-
●	<i>instruction_access_error</i>	00A <sub>16</sub>	precise	4	H (nm)	H (nm)	H (nm)
—	<i>Reserved</i>	00B <sub>16</sub> – 00D <sub>16</sub>	—	—	—	—	—
—	<i>Reserved</i>	00D <sub>16</sub> – 00E <sub>16</sub>	—	—	—	—	—
—	<i>Reserved</i>	00F <sub>16</sub>	—	—	—	—	—
●	<i>illegal_instruction</i>	010 <sub>16</sub>	precise	6.2	H (nm)	H (nm)	H (nm)
●	<i>privileged_opcode</i>	011 <sub>16</sub>	precise	7	P (nm)	-x-	-x-
—	<i>Reserved</i>	012 <sub>16</sub> – 013 <sub>16</sub>	—	—	—	—	—

**TABLE 12-4** Exception and Interrupt Requests, by TT Value (2 of 6)

UA-2005 ●=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	HP
—	<i>Reserved</i>	014 <sub>16</sub> – 017 <sub>16</sub>	—	—	—	—	—
—	<i>Reserved</i>	018 <sub>16</sub> – 01F <sub>16</sub>	—	—	—	—	—
●	<i>fp_disabled</i>	020 <sub>16</sub>	precise	8	P (nm)	P (nm)	H <sup>U</sup> (nm)
○	<i>fp_exception_ieee_754</i>	021 <sub>16</sub>	precise	11.1	P (nm)	P (nm)	H <sup>U</sup> (nm)
○	<i>fp_exception_other</i>	022 <sub>16</sub>	precise	11.1	P (nm)	P (nm)	H <sup>U</sup> (nm)
●	<i>tag_overflow</i> <sup>D</sup>	023 <sub>16</sub>	precise	14	P (nm)	P (nm)	H <sup>U</sup> (nm)
●	<i>clean_window</i>	024 <sub>16</sub> <sup>‡</sup>	precise	10.1	P (nm)	P (nm)	H <sup>U</sup> (nm)
—	<i>Reserved</i>	025 <sub>16</sub> – 027 <sub>16</sub>	—	—	—	—	—
●	<i>division_by_zero</i>	028 <sub>16</sub>	precise	15	P (nm)	P (nm)	H <sup>U</sup> (nm)
○	<i>internal_processor_error</i>	029 <sub>16</sub>	precise	◆	H (nm)	H (nm)	H (nm)
○	<i>instruction_invalid_TSB_entry</i>	02A <sub>16</sub>	precise	2.10	H (nm)	H (nm)	-x-
○	<i>data_invalid_TSB_entry</i>	02B <sub>16</sub>	precise	12.03	H (nm)	H (nm)	H (nm)
—	<i>Reserved</i>	02C <sub>16</sub>	—	—	—	—	—
—	<i>Reserved</i>	02D <sub>16</sub> – 02F <sub>16</sub>	—	—	—	—	—
●	<i>data_access_exception</i>	030 <sub>16</sub>	precise	12.01	H (nm)	H (nm)	H <sup>U</sup> (nm)
●	<i>data_access_MMU_miss</i> <sup>†</sup>	031 <sub>16</sub>	precise	12.03	H (nm)	H (nm)	H (nm)
○	<i>data_access_error</i>	032 <sub>16</sub>	precise	12.10	H (nm)	H (nm)	H (nm)

TABLE 12-4 Exception and Interrupt Requests, by TT Value (3 of 6)

UA-2005 ●=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	HP
—	<i>data_access_protection</i>	033 <sub>16</sub>	precise	12.07	H (nm)	H (nm)	H (nm)
●	<i>mem_address_not_aligned</i>	034 <sub>16</sub>	precise	10.2	H (nm)	H (nm)	H <sup>U</sup> (nm)
●	<i>LDDF_mem_address_not_aligned</i>	035 <sub>16</sub>	precise	10.1	H (nm)	H (nm)	H <sup>U</sup> (nm)
●	<i>STDF_mem_address_not_aligned</i>	036 <sub>16</sub>	precise	10.1	H (nm)	H (nm)	H <sup>U</sup> (nm)
●	<i>privileged_action</i>	037 <sub>16</sub>	precise	11.1	H (nm)	H (nm)	-x-
○	<i>LDQF_mem_address_not_aligned</i>	038 <sub>16</sub>	precise	10.1	H (nm)	H (nm)	H <sup>U</sup> (nm)
○	<i>STQF_mem_address_not_aligned</i>	039 <sub>16</sub>	precise	10.1	H (nm)	H (nm)	H <sup>U</sup> (nm)
—	<i>Reserved</i>	03A <sub>16</sub>	—	—	—	—	—
—	<i>Reserved</i>	03B <sub>16</sub>	—	—	—	—	—
—	<i>Reserved</i>	03B <sub>16</sub> – 03D <sub>16</sub>	—	—	—	—	—
●	<i>instruction_real_translation_miss</i>	03E <sub>16</sub>	precise	2.08	H (nm)	H (nm)	-x-
●	<i>data_real_translation_miss</i>	03F <sub>16</sub>	precise	12.03	H (nm)	H (nm)	H (nm)
—	<i>Reserved</i>	040 <sub>16</sub>	—	—	—	—	—
●	<i>interrupt_level_n</i> (n = 1–15)	041 <sub>16</sub> – 04F <sub>16</sub>	disrupting	32–n (31 to 17)	P (ie)	P (ie)	(pend)
—	<i>Reserved</i>	050 <sub>16</sub> – 05D <sub>16</sub>	—	—	—	—	—
●	<i>hstick_match</i>	05E <sub>16</sub>	disrupting	16.01	H (nm)	H (nm)	H (ie)
●	<i>trap_level_zero</i>	05F <sub>16</sub>	disrupting	2.02	H	H	-x-
—	<i>Reserved</i>	060 <sub>16</sub>	—	—	—	—	—

**TABLE 12-4** Exception and Interrupt Requests, by TT Value (4 of 6)

UA-2005 ●=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	HP
○	<i>PA_watchpoint (RA_watchpoint)</i>	061 <sub>16</sub>	precise	12.09	H (nm)	H (nm)	H (nm)
—	<i>Reserved</i>	062 <sub>16</sub>	—	—	—	—	—
○	<i>VA_watchpoint</i>	062 <sub>16</sub>	precise	11.2	P (nm)	P (nm)	-x-
●	<i>fast_instruction_access_MMU_miss</i>	064 <sub>16</sub> <sup>‡</sup>	precise	2.08	H (nm)	H (nm)	-x-
—	<i>Reserved</i>	065 <sub>16</sub> – 067 <sub>16</sub>	—	—	—	—	—
●	<i>fast_data_access_MMU_miss</i>	068 <sub>16</sub> <sup>‡</sup>	precise	12.03	H (nm)	H (nm)	H (nm)
—	<i>Reserved</i>	069 <sub>16</sub> – 06B <sub>16</sub>	—	—	—	—	—
●	<i>fast_data_access_protection</i>	06C <sub>16</sub> <sup>‡</sup>	precise	12.07	H (nm)	H (nm)	H (nm)
—	<i>Reserved</i>	06D <sub>16</sub> – 06F <sub>16</sub>	—	—	—	—	—
○	<i>implementation_dependent_exception_n</i> (impl. dep. #35-V8-Cs20)	070 <sub>16</sub> – 075 <sub>16</sub>	—	∇	—	—	—
●	<i>instruction_breakpoint</i>	076 <sub>16</sub>	precise	6.1	H	H	H
□	<i>implementation_dependent_exception_n</i> (impl. dep. #35-V8-Cs20)	077	—	∇	—	—	—
□	<i>implementation_dependent_exception_n</i> (impl. dep. #35-V8-Cs20)	079 <sub>16</sub> – 07B <sub>16</sub>	—	∇	—	—	—
—	<i>Reserved</i>	079 <sub>16</sub>	—	—	—	—	—
●	<i>cpu_mondo</i>	07C <sub>16</sub>	disrupting	16.08	P (ie)	P (ie)	(pend)
●	<i>dev_mondo</i>	07D <sub>16</sub>	disrupting	16.11	P (ie)	P (ie)	(pend)
●	<i>resumable_error</i>	07E <sub>16</sub>	disrupting	33.3	P (ie)	P (ie)	(pend)
□	<i>implementation_dependent_exception_15</i> (impl. dep. #35-V8-Cs20)	07F <sub>16</sub>	—	∇	—	—	—

TABLE 12-4 Exception and Interrupt Requests, by TT Value (5 of 6)

UA-2005 ●=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	HP
—	<i>nonresumable_error</i> (generated by hyperprivileged software, not by hardware)	07F <sub>16</sub>	—	—	—	—	—
●	<i>spill_n_normal</i> ( <i>n</i> = 0–7)	080 <sub>16</sub> <sup>‡</sup> – 09C <sub>16</sub> <sup>‡</sup>	precise	9	P (nm)	P (nm)	H <sup>U</sup> (nm)
●	(reserved for use by <i>spill_7_normal</i> ; see footnote for trap type 09C <sub>16</sub> )	09D <sub>16</sub> – 09F <sub>16</sub>	—	—	—	—	—
●	<i>spill_n_other</i> ( <i>n</i> = 0–7)	0A0 <sub>16</sub> <sup>‡</sup> – 0BC <sub>16</sub> <sup>‡</sup>	precise	9	P (nm)	P (nm)	H <sup>U</sup> (nm)
●	(reserved for use by <i>spill_7_other</i> see footnote for trap type 0BC <sub>16</sub> )	0BD <sub>16</sub> – 0BF <sub>16</sub>	—	—	—	—	—
●	<i>fill_n_normal</i> ( <i>n</i> = 0–7)	0C0 <sub>16</sub> <sup>‡</sup> – 0DC <sub>16</sub> <sup>‡</sup>	precise	9	P (nm)	P (nm)	H <sup>U</sup> (nm)
●	(reserved for use by <i>fill_7_normal</i> ; see footnote for trap type 0DC <sub>16</sub> )	0DD <sub>16</sub> – 0DF <sub>16</sub>	—	—	—	—	—
●	<i>fill_n_other</i> ( <i>n</i> = 0–7)	0E0 <sub>16</sub> <sup>‡</sup> – 0FC <sub>16</sub> <sup>‡</sup>	precise	9	P (nm)	P (nm)	H <sup>U</sup> (nm)
●	(reserved for use by <i>fill_7_other</i> see footnote for trap type 0FC <sub>16</sub> )	0FD <sub>16</sub> – 0FF <sub>16</sub>	—	—	—	—	—
●	<i>trap_instruction</i>	100 <sub>16</sub> – 17F <sub>16</sub>	precise	16.02	P (nm)	P (nm)	H <sup>U</sup> (nm)

TABLE 12-4 Exception and Interrupt Requests, by TT Value (6 of 6)

UA-2005 ●=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	HP
●	<i>htrap_instruction</i>	180 <sub>16</sub> – 1FF <sub>16</sub>	precise	16.02	-x-	H (nm)	H <sup>U</sup> (nm)
●	<i>guest_watchdog</i> <sup>◇</sup>	TT <sup>◇</sup>	precise or disrupting <sup>◇</sup>	◇	H (nm)	H (nm)	-x-

\* Although these trap priorities are recommended, all trap priorities are implementation dependent (impl. dep. #36-V8 on page 469), including relative priorities within a given priority level.

† This exception type is only used in UltraSPARC Architecture 2005 implementations that support hardware MMU table walking. See description of this exception in *Exception and Interrupt Descriptions* on page 481.

‡ 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.

◇ The *guest\_watchdog* trap is caused when  $TL \geq MAXPTL$  and any precise or disrupting trap occurs that is destined for privileged mode. *guest\_watchdog* shares a trap table offset with *watchdog\_reset* (40<sub>16</sub>), but retains the trap type (TT) value and priority of the exception that caused the trap.

♣ *watchdog\_reset* uses the trap vector entry for trap type 002<sub>16</sub> (trap table offset 40<sub>16</sub>), but retains the trap type (TT) value of the exception that caused entry into error\_state .

♣ *RED\_state\_exception* uses the trap vector entry for trap type 005<sub>16</sub> (trap table offset A0<sub>16</sub>), but retains the trap type (TT) value and priority of the exception that caused the trap.

◆ The priority of *internal\_processor\_error* is implementation dependent (impl. dep. # 402-S10)

∇ The priority of an *implementation\_dependent\_exception\_n* trap is implementation dependent (impl. dep. # 35-V8-Cs20)

Ⓓ 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 4)

UA-2005 ●=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	HP
●	<i>power_on_reset</i>	001 <sub>16</sub>	reset	0	H (nm)	H (nm)	H (nm)
●	<i>externally_initiated_reset</i>	003 <sub>16</sub>	reset	1.1	H (nm)	H (nm)	H (nm)
●	<i>watchdog_reset</i>	TT <sup>▲</sup>	reset	1.2	H (nm)	H (nm)	H (nm)
●	<i>software_initiated_reset</i>	004 <sub>16</sub>	reset	1.3	-x-	-x-	H (nm)
○	<i>store_error</i>	007 <sub>16</sub>	deferred	2.01	H (nm)	H (nm)	H (nm)
●	<i>trap_level_zero</i>	05F <sub>16</sub>	disrupting	2.02	H (nm)	H (nm)	-x-
●	<i>instruction_real_translation_miss</i>	03E <sub>16</sub>	precise	2.08	H (nm)	H (nm)	-x-
●	<i>instruction_access_MMU_miss<sup>†</sup></i>	009 <sub>16</sub>	precise		H (nm)	H (nm)	-x-
●	<i>fast_instruction_access_MMU_miss</i>	064 <sub>16</sub> <sup>‡</sup>	precise		H (nm)	H (nm)	-x-
○	<i>instruction_invalid_TSB_entry</i>	02A <sub>16</sub>	precise	2.10	H (nm)	H (nm)	-x-
●	<i>instruction_access_exception</i>	008 <sub>16</sub>	precise	3	H (nm)	H (nm)	H <sup>U</sup> (nm)
●	<i>instruction_access_error</i>	00A <sub>16</sub>	precise	4	H (nm)	H (nm)	H (nm)
●	<i>instruction_breakpoint</i>	076 <sub>16</sub>	precise	6.1	H (nm)	H (nm)	H (nm)
●	<i>illegal_instruction</i>	010 <sub>16</sub>	precise	6.2	H (nm)	H (nm)	H (nm)
●	<i>privileged_opcode</i>	011 <sub>16</sub>	precise	7	P (nm)	-x-	-x-
●	<i>fp_disabled</i>	020 <sub>16</sub>	precise	8	P (nm)	P (nm)	H <sup>U</sup> (nm)

TABLE 12-5 Exception and Interrupt Requests, by Priority (2 of 4)

UA-2005 ●=Req'd. ○=Opt'l □=Impl- Specific	Exception or Interrupt Request	TT (Trap Type)	Trap Category	Priority (0 = High- est)	Mode in which Trap is Delivered and (and Conditioning Applied), based on Current Privilege Mode		
					NP	Priv	HP
●	<i>spill_n_normal</i> ( <i>n</i> = 0–7)	080 <sub>16</sub> <sup>‡</sup> – 09C <sub>16</sub> <sup>‡</sup>	precise	9	P (nm)	P (nm)	H <sup>U</sup> (nm)
●	<i>spill_n_other</i> ( <i>n</i> = 0–7)	0A0 <sub>16</sub> <sup>‡</sup> – 0BC <sub>16</sub> <sup>‡</sup>	precise		P (nm)	P (nm)	H <sup>U</sup> (nm)
●	<i>fill_n_normal</i> ( <i>n</i> = 0–7)	0C0 <sub>16</sub> <sup>‡</sup> – 0DC <sub>16</sub> <sup>‡</sup>	precise		P (nm)	P (nm)	H <sup>U</sup> (nm)
●	<i>fill_n_other</i> ( <i>n</i> = 0–7)	0E0 <sub>16</sub> <sup>‡</sup> – 0FC <sub>16</sub> <sup>‡</sup>	precise		P (nm)	P (nm)	H <sup>U</sup> (nm)
●	<i>clean_window</i>	024 <sub>16</sub> <sup>‡</sup>	precise	10.1	P (nm)	P (nm)	H <sup>U</sup> (nm)
●	<i>LDDF_mem_address_not_aligned</i>	035 <sub>16</sub>	precise		H (nm)	H (nm)	H <sup>U</sup> (nm)
●	<i>STDF_mem_address_not_aligned</i>	036 <sub>16</sub>	precise		H (nm)	H (nm)	H <sup>U</sup> (nm)
○	<i>LDQF_mem_address_not_aligned</i>	038 <sub>16</sub>	precise		H (nm)	H (nm)	H <sup>U</sup> (nm)
○	<i>STQF_mem_address_not_aligned</i>	039 <sub>16</sub>	precise		H (nm)	H (nm)	H <sup>U</sup> (nm)
●	<i>mem_address_not_aligned</i>	034 <sub>16</sub>	precise	10.2	H (nm)	H (nm)	H <sup>U</sup> (nm)
○	<i>fp_exception_other</i>	022 <sub>16</sub>	precise	11.1	P (nm)	P (nm)	H <sup>U</sup> (nm)
○	<i>fp_exception_ieee_754</i>	021 <sub>16</sub>	precise		P (nm)	P (nm)	H <sup>U</sup> (nm)
●	<i>privileged_action</i>	037 <sub>16</sub>	precise		H (nm)	H (nm)	-x-
○	<i>VA_watchpoint</i>	062 <sub>16</sub>	precise	11.2	P (nm)	P (nm)	-x-
●	<i>data_access_exception</i>	030 <sub>16</sub>	precise	12.01	H (nm)	H (nm)	H <sup>U</sup> (nm)

TABLE 12-5 Exception and Interrupt Requests, by Priority (3 of 4)

UA-2005 ●=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	HP
●	<i>data_real_translation_miss</i>	03F <sub>16</sub>	precise	12.03	H (nm)	H (nm)	H (nm)
●	<i>data_access_MMU_miss</i> <sup>†</sup>	031 <sub>16</sub>	precise		H (nm)	H (nm)	H (nm)
●	<i>fast_data_access_MMU_miss</i>	068 <sub>16</sub> <sup>‡</sup>	precise		H (nm)	H (nm)	H (nm)
○	<i>data_invalid_TSB_entry</i>	02B <sub>16</sub>	precise		H (nm)	H (nm)	H (nm)
●	<i>fast_data_access_protection</i>	06C <sub>16</sub> <sup>‡</sup>	precise	12.07	H (nm)	H (nm)	H (nm)
—	<i>data_access_protection</i>	033 <sub>16</sub>	precise		H (nm)	H (nm)	H (nm)
○	<i>PA_watchpoint (RA_watchpoint)</i>	061 <sub>16</sub>	precise	12.09	H (nm)	H (nm)	H (nm)
○	<i>data_access_error</i>	032 <sub>16</sub>	precise	12.10	H (nm)	H (nm)	H (nm)
●	<i>tag_overflow</i> <sup>D</sup>	023 <sub>16</sub>	precise	14	P (nm)	P (nm)	H <sup>U</sup> (nm)
●	<i>division_by_zero</i>	028 <sub>16</sub>	precise	15	P (nm)	P (nm)	H <sup>U</sup> (nm)
●	<i>hstick_match</i>	05E <sub>16</sub>	disrupting	16.01	H (nm)	H (nm)	H (ie)
●	<i>trap_instruction</i>	100 <sub>16</sub> – 17F <sub>16</sub>	precise	16.02	P (nm)	P (nm)	H (nm)
●	<i>htrap_instruction</i>	180 <sub>16</sub> – 1FF <sub>16</sub>	precise		-x- (nm)	H (nm)	H <sup>U</sup> (nm)
●	<i>cpu_mondo</i>	07C <sub>16</sub>	disrupting	16.08	P (ie)	P (ie)	(pend)
●	<i>dev_mondo</i>	07D <sub>16</sub>	disrupting	16.11	P (ie)	P (ie)	(pend)

TABLE 12-5 Exception and Interrupt Requests, by Priority (4 of 4)

UA-2005 ●=Req'd. ○=Opt'l □.=Impl- Specific	Exception or Interrupt Request	TT (Trap Type)	Trap Category	Priority (0 = High- est)	Mode in which Trap is Delivered and (and Conditioning Applied), based on Current Privilege Mode		
					NP	Priv	HP
●	<i>interrupt_level_n</i> ( <i>n</i> = 1–15)	041 <sub>16</sub> – 04F <sub>16</sub>	disrupting	32- <i>n</i> (31 to 17)	P (ie)	P (ie)	(pend)
●	<i>resumable_error</i>	07E <sub>16</sub>	disrupting	33.3	P (ie)	P (ie)	(pend)
●	<i>guest_watchdog</i> <sup>◇</sup>	TT <sup>◇</sup>	precise or disrupting <sup>◇</sup>	◇	H (nm)	H (nm)	-x-
●	<i>RED_state_exception</i>	TT <sup>*</sup>	precise	♣	H (nm)	H (nm)	H (nm)
○	<i>internal_processor_error</i>	029 <sub>16</sub>	precise	◆	H (nm)	H (nm)	H (nm)
○	<i>implementation_dependent_exception_n</i> (impl. dep. #35-V8-Cs20)	070 <sub>16</sub> – 075 <sub>16</sub> ′ 077 <sub>16</sub> ′ 079 <sub>16</sub> – 07B <sub>16</sub> ′ 07F <sub>16</sub>	—	∇	—	—	—
—	<i>nonresumable_error</i> (generated by hyperprivileged software, not by hardware)	07F <sub>16</sub>	—	—	—	—	—

\* Although these trap priorities are recommended, all trap priorities are implementation dependent (impl. dep. #36-V8 on page 469), including relative priorities within a given priority level.

† This exception type is only used in UltraSPARC Architecture 2005 implementations that support hardware MMU table walking. See description of this exception in *Exception and Interrupt Descriptions* on page 481.

‡ 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.

◇ The *guest\_watchdog* trap is caused when  $TL \geq MAXPTL$  and any precise or disrupting trap occurs that is destined for privileged mode. *guest\_watchdog* shares a trap table offset with *watchdog\_reset* (40<sub>16</sub>), but retains the trap type (TT) value and priority of the exception that caused the trap.

♣ *watchdog\_reset* uses the trap vector entry for trap type 002<sub>16</sub> (trap table offset 40<sub>16</sub>), but retains the trap type (TT) value of the exception that caused entry into error\_state.

♣ *RED\_state\_exception* uses the trap vector entry for trap type 005<sub>16</sub> (trap table offset A0<sub>16</sub>), but retains the trap type (TT) value and priority of the exception that caused the trap.

◆ The priority of *internal\_processor\_error* is implementation dependent (impl. dep. # 402-S10)

∇ The priority of an *implementation\_dependent\_exception\_n* trap is implementation dependent (impl. dep. # 35-V8-Cs20)

□ This exception is deprecated, because the only instructions that can generate it have been deprecated.

### 12.5.7.1 Trap Type for Spi ll/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 12-9.

Bit	Field	Description
8:6	spill_or_fill	010 <sub>2</sub> for spill traps; 011 <sub>2</sub> for fill traps
5	other	(OTHERWIN ≠ 0)
4:2	wtype	If (other) then WSTATE.other; else WSTATE.normal




FIGURE 12-9 Trap Type Encoding for Spill/Fill Traps

## 12.5.8 Trap Priorities

TABLE 12-4 on page 459 and TABLE 12-5 on page 465 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 UltraSPARC 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. For example, if an *instruction\_access\_error* occurs (priority 3), it will be taken even if the instruction is an SIR (priority 1). This situation occurs because the virtual processor detects *instruction\_access\_error* during instruction fetch and never actually issues or executes the instruction, so the SIR instruction is never seen by the execution units of the virtual processor. This is an obvious case, but there are other more subtle cases.

---

## 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), HPSTATE, 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.

The following traps are processed in RED\_state:

- POR, XIR, and WDR reset requests
- SIR reset request when  $TL < MAXTL$
- Non-reset traps taken when  $TL = MAXTL - 1$
- Traps taken when the virtual processor is in RED\_state

All other traps are handled in execute\_state using normal trap processing.

During normal operation, the virtual processor is in execute\_state. It processes traps in execute\_state and continues.

When a nonreset trap or software-initiated reset (SIR) occurs with  $TL = MAXTL$ , there are no more levels on the trap stack, so the virtual processor enters the transitory state error\_state. The virtual processor remains in error\_state for an implementation-dependent duration, then generates a WDR reset (impl. dep. #254-U3-Cs10) to effect a change from error\_state to RED\_state.

Traps processed in RED\_state use a special trap vector and a special trap-vectoring algorithm. RED\_state vectoring and the setting of the TT value for RED\_state traps are described in *RED\_state Trap Table Organization* on page 456.

Traps that occur with  $TL = MAXTL - 1$  are processed in RED\_state. In addition, reset traps are also processed in RED\_state. Reset trap processing is described in *RED\_state Trap Processing* on page 474. Finally, software can force the processor into RED\_state by setting the HPSTATE.red bit to 1.

Once the virtual processor has entered RED\_state, no matter how it got there, all subsequent traps are processed in RED\_state until software returns the virtual processor to execute\_state or a normal or SIRtrap is taken with  $TL = MAXTL$ , which puts the virtual processor in error\_state.

TABLE 12-6, TABLE 12-7, and TABLE 12-8 describe 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 Type				
	Nonreset Trap or Interrupt	POR	XIR	WDR ‡	SIR
<code>execute_state</code> TL < MAXTL - 1	<code>execute_state</code> TL ← TL + 1	<code>RED_state</code> TL = MAXTL	<code>RED_state</code> TL ← TL + 1	‡	<code>RED_state</code> TL ← TL + 1
<code>execute_state</code> TL = MAXTL - 1	<code>RED_state</code> TL = MAXTL	<code>RED_state</code> TL = MAXTL	<code>RED_state</code> TL = MAXTL	‡	<code>RED_state</code> TL = MAXTL
<code>execute_state</code> <sup>†</sup> TL = MAXTL	<code>error_state</code> TL = MAXTL	<code>RED_state</code> TL = MAXTL	<code>RED_state</code> TL = MAXTL	‡	<code>error_state</code> TL = MAXTL

<sup>†</sup> This state occurs when software changes TL to MAXTL and leaves HPSTATE.red = 0, or if software sets HPSTATE.red ← 0 while TL = MAXTL.

‡ WDR can only be generated from `error_state`.

**TABLE 12-7** Trap Received While in `RED_state`

Original State	New State, After Receiving Trap Type				
	Nonreset Trap or Interrupt	POR	XIR	WDR ‡	SIR
<code>RED_state</code> TL < MAXTL - 1	<code>RED_state</code> TL ← TL + 1	<code>RED_state</code> TL = MAXTL	<code>RED_state</code> TL ← TL + 1	‡	<code>RED_state</code> TL ← TL + 1
<code>RED_state</code> TL = MAXTL - 1	<code>RED_state</code> TL = MAXTL	<code>RED_state</code> TL = MAXTL	<code>RED_state</code> TL = MAXTL	‡	<code>RED_state</code> TL = MAXTL
<code>RED_state</code> TL = MAXTL	<code>error_state</code> TL = MAXTL	<code>RED_state</code> TL = MAXTL	<code>RED_state</code> TL = MAXTL	‡	<code>error_state</code> TL = MAXTL

‡ WDR can only be generated from `error_state`.

**TABLE 12-8** Reset Received While in `error_state`

Original State	New State, After Receiving Trap Type				
	Nonreset Trap or Interrupt	POR	XIR	WDR	SIR
<code>error_state</code> TL = MAXTL	—	<code>RED_state</code> TL = MAXTL	<code>RED_state</code> TL = MAXTL	<code>RED_state</code> TL = MAXTL	—

The virtual processor does not recognize interrupts while it is in `error_state`.

A non-reset trap causes the following state changes to occur:

- If the virtual processor is already in `RED_state`, the new trap is processed in `RED_state` unless  $TL = MAXTL$ . See *Nonreset Traps When the Virtual Processor Is in RED\_state* on page 480.
- If the virtual processor is in `execute_state` and the trap level is one less than its maximum value, that is,  $TL = MAXTL - 1$ , then the virtual processor enters `RED_state`. See *RED\_state* on page 441 and *Nonreset Traps with  $TL = MAXTL - 1$*  on page 474.
- If the virtual processor is in either `execute_state` or `RED_state` and the trap level is already at its maximum value, that is,  $TL = MAXTL$ , then the virtual processor enters `error_state`. See *error\_state* on page 444.

Otherwise, the trap uses normal trap processing, described in the following section on *Normal Trap Processing*.

## 12.6.1 Normal Trap Processing

Normal traps comprise all traps processed in `execute_state`; that is, all non-`RED_state` and non-`error_state` traps.

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          ← TL + 1    // note that if TL = MAXTL - 1 before this trap,
                       // trap would have been processed in
                       // RED_state, not here using normal trap
                       // processing.

```

- Existing state is preserved.
 

```

TSTATE[TL].gl ← GL
TSTATE[TL].ccr ← CCR
TSTATE[TL].asi ← ASI
TSTATE[TL].pstate ← PSTATE
TSTATE[TL].cwp ← CWP
TPC[TL]       ← PC // (upper 32 bits zeroed if PSTATE.am = 1)
TNPC[TL]      ← NPC // (upper 32 bits zeroed if PSTATE.am = 1)
HTSTATE[TL].hpstate ← HPSTATE // even for traps to privileged mode

```
- The trap type is preserved.
 

```

TT[TL]       ← the trap type

```
- The Global Level register (GL) is updated. This normally provides access to a fresh set of global registers:

```

if (the trap is being delivered in privileged mode)
then GL      ← min (GL + 1, MAXPGL)
else (trap is being delivered in hyperprivileged mode)
      GL      ← min (GL + 1, MAXGL)
endif

```

- The PSTATE register is updated to a predefined state (even for traps to hyperprivileged mode):

```

PSTATE.mm   is unchanged
PSTATE.pef  ← 1 // if an FPU is present, it is enabled
PSTATE.am   ← 0 // address masking is turned off
if (the trap is being delivered in privileged mode)
then PSTATE.priv ← 1 // the virtual processor enters privileged mode
      PSTATE.cle ← PSTATE.tle //set endian mode for traps
else // trap is being delivered in hyperprivileged mode
      PSTATE.priv ← 0
      PSTATE.cle ← 0
endif
PSTATE.ie   ← 0 // interrupts are disabled
PSTATE.tle  is unchanged
PSTATE.tct  ← 0 // trap on CTI disabled

```

- The HPSTATE register is updated:

```

if (the trap is to hyperprivileged mode)
then HPSTATE.red ← 0
      HPSTATE.hpriv ← 1 // enter hyperprivileged mode
      HPSTATE.ibe   ← 0 disable instruction breakpoints
      HPSTATE.tlz is unchanged
endif

```

- 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] = 02416 // a clean_window trap
then CWP ← CWP + 1
endif
if (08016 ≤ TT[TL] ≤ 0BF16) // window spill trap
then CWP ← CWP + CANSERVE + 2
endif
if (0C016 ≤ TT[TL] ≤ 0FF16) // window fill trap
then CWP ← CWP - 1
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 if ( (trap is to privileged mode) and (TL > MAXPTL) )
then // this is the guest_watchdog case; the trap is handled in
  // hyperprivileged mode using trap table offset 4016.
  PC ← HTBA{63:14} :: 002 :: 04016
  NPC ← HTBA{63:14} :: 002 :: 04416
else { trap is handled in hyperprivileged mode }
  PC ← HTBA{63:14} :: TT[TL] :: 0 00002
  NPC ← HTBA{63:14} :: TT[TL] :: 0 01002
endif

```

Interrupts are ignored as long as PSTATE.ie = 0.

<b>Programming Note</b>	State in TPC[ <i>n</i> ], TNPC[ <i>n</i> ], TSTATE[ <i>n</i> ], and TT[ <i>n</i> ] is only changed autonomously by the processor when a trap is taken while TL = <i>n</i> - 1; however, software can change any of these values with a WRPR instruction when TL = <i>n</i> .
-------------------------	--

## 12.6.2 RED\_state Trap Processing

The following conditions invoke RED\_state trap processing, and cause the trap to be delivered in hyperprivileged mode:

- Traps taken with TL = MAXTL - 1
- Power-on reset traps
- Watchdog reset traps
- Externally initiated reset traps
- Software-initiated reset traps
- Traps taken when the virtual processor is already in RED\_state

**IMPL. DEP. #38-V8:** Implementation-dependent registers may or may not be affected by the various reset traps.

### 12.6.2.1 Nonreset Traps with TL = MAXTL - 1

Nonreset traps that occur when TL = MAXTL - 1 are processed in RED\_state.

The following state changes occur (conceptually, in this order) during a nonreset trap that occurs when  $TL = MAXTL - 1$ :

- The trap level is advanced.
  - $TL \leftarrow MAXTL$
- Existing state is preserved.
  - $TSTATE[TL].gl \leftarrow GL$
  - $TSTATE[TL].ccr \leftarrow CCR$
  - $TSTATE[TL].asi \leftarrow ASI$
  - $TSTATE[TL].pstate \leftarrow PSTATE$
  - $TSTATE[TL].cwp \leftarrow CWP$
  - $TPC[TL] \leftarrow PC$  // (upper 32 bits zeroed if  $PSTATE.am = 1$ )
  - $TNPC[TL] \leftarrow NPC$  // (upper 32 bits zeroed if  $PSTATE.am = 1$ )
  
  - $HTSTATE[TL].hpstate \leftarrow HPSTATE$
- The trap type is preserved.
  - $TT[TL] \leftarrow$  the trap type
- The Global Level register is updated.
  - $GL \leftarrow \min(GL + 1, MAXGL)$
- The PSTATE register is set as follows:
  - $PSTATE.mm \leftarrow 00_2$  // TSO
  - $PSTATE.pef \leftarrow 1$  // if an FPU is present, it is enabled
  - $PSTATE.am \leftarrow 0$  // address masking is turned off
  - $PSTATE.priv \leftarrow 0$  // entering hyperprivileged mode
  - $PSTATE.ie \leftarrow 0$  // interrupts are disabled
  - $PSTATE.cle \leftarrow 0$  // big-endian is default for hyperprivileged mode
  - $PSTATE.tle$  is unchanged // (*was unspecified in SPARC V9 specification*)
  - $PSTATE.tct \leftarrow 0$  // trap on CTI disabled
- The HPSTATE register is updated:
  - $HPSTATE.red \leftarrow 1$  // enter RED\_state
  - $HPSTATE.hpriv \leftarrow 1$  // enter hyperprivileged mode
  - $HPSTATE.ibe \leftarrow 0$  // disable instruction breakpoints
  - $HPSTATE.tlz \leftarrow 0$  // disable *trap\_level\_zero* exceptions
- For a register-window trap 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  $CWP \leftarrow CWP + 1$
  - endif
  - If  $(080_{16} \leq TT[TL] \leq 0BF_{16})$  // window *spill* trap
  - then  $CWP \leftarrow CWP + CANSAVE + 2$
  - endif

```

    If  $(0C0_{16} \leq TT[TL] \leq 0FF_{16})$  // window fill trap
    then CWP  $\leftarrow$  CWP - 1
    endif

```

For non-register-window traps, CWP is not changed.

- Implementation-specific state changes; for example, disabling an MMU.
- Control is transferred into the RED\_state trap table. See *Trap Table Entry Address to RED\_state* on page 455 for further details of RSTVADDR.
 

```

          PC           $\leftarrow$  RSTVADDR{63:8} :: 1010 00002
          NPC          $\leftarrow$  RSTVADDR{63:8} :: 1010 01002
      
```

### 12.6.2.2 Power-On Reset (POR) Traps

A POR trap occurs when power is applied to the virtual processor. If the virtual processor is in error\_state, a POR brings the virtual processor out of error\_state and places it in RED\_state. See Chapter 16, *Resets* for further details.

Virtual processor state is undefined after POR, except for the following:

- The trap level is set.
 

```

          TL           $\leftarrow$  MAXTL
      
```
- The trap type is set.
 

```

          TT[TL]      $\leftarrow$  00116
      
```
- The Global Level register is updated.
 

```

          GL          $\leftarrow$  MAXGL
      
```
- The PSTATE register is set as follows:
 

```

          PSTATE.mm   $\leftarrow$  002 // TSO
          PSTATE.pef  $\leftarrow$  1 // if an FPU is present, it is enabled
          PSTATE.am   $\leftarrow$  0 // address masking is turned off
          PSTATE.priv  $\leftarrow$  0 // entering hyperprivileged mode
          PSTATE.ie   $\leftarrow$  0 // interrupts are disabled
          PSTATE.cle  $\leftarrow$  0 // big-endian is default for hyperprivileged mode
          PSTATE.tle  $\leftarrow$  0 // big-endian mode for traps
          PSTATE.tct  $\leftarrow$  0 // trap on CTI disabled
      
```
- The HPSTATE register is updated:
 

```

          HPSTATE.red  $\leftarrow$  1 // enter RED_state
          HPSTATE.hpriv  $\leftarrow$  1 // enter hyperprivileged mode
          HPSTATE.ibe  $\leftarrow$  0 // disable instruction breakpoints
          HPSTATE.tlz  $\leftarrow$  0 // disable trap_level_zero exceptions
      
```
- The TICK register is protected.
 

```

          TICK.npt    $\leftarrow$  1 // TICK is unreadable by nonprivileged software
      
```
- Implementation-specific state changes; for example, disabling an MMU.

- Control is transferred into the `RED_state` trap table.
 

PC	$\leftarrow RSTVADDR\{63:8\} :: 0010\ 0000_2$
NPC	$\leftarrow RSTVADDR\{63:8\} :: 0010\ 0100_2$

### 12.6.2.3 Watchdog Reset (WDR) Traps

Entry to `error_state` is caused by occurrence of a trap when `TL = MAXTL` (impl. dep. #39-V8-Cs10). See `error_state` on page 444.

To recover from `error_state`, the UltraSPARC Architecture provides *watchdog\_reset* (WDR), which causes a transition from `error_state` to `RED_state` (impl. dep. #254-U3-Cs10).

The following virtual processor state changes occur during WDR (conceptually, in this order):

- The trap level is updated.
 

TL	$\leftarrow \min(TL + 1, MAXTL)$
----	----------------------------------
- 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]	$\leftarrow PC$ // (upper 32 bits zeroed if <code>PSTATE.am = 1</code> )
TNPC[TL]	$\leftarrow NPC$ // (upper 32 bits zeroed if <code>PSTATE.am = 1</code> )
HTSTATE[TL].hpstate	$\leftarrow HPSTATE$
- The trap type is set.
 

TT[TL]	$\leftarrow$ the trap type that caused the WDR
--------	--
- The Global Level register is updated.
 

GL	$\leftarrow \min(GL + 1, MAXGL)$
----	----------------------------------
- The PSTATE register is set as follows:
 

PSTATE.mm	$\leftarrow 00_2$ // TSO
PSTATE.pef	$\leftarrow 1$ // if an FPU is present, it is enabled
PSTATE.am	$\leftarrow 0$ // address masking is turned off
PSTATE.priv	$\leftarrow 0$ // entering hyperprivileged mode
PSTATE.ie	$\leftarrow 0$ // interrupts are disabled
PSTATE.cle	$\leftarrow 0$ // big-endian is default for hyperprivileged mode
PSTATE.tle	is unchanged // ( <i>was unspecified in SPARC V9 specification</i> )
PSTATE.tct	$\leftarrow 0$ // trap on CTI disabled
- The HPSTATE register is updated:

```

HPSTATE.red ← 1 // enter RED_state
HPSTATE.hpriv ← 1 // enter hyperprivileged mode
HPSTATE.ibe ← 0 // disable instruction breakpoints
HPSTATE.tlz ← 0 // disable trap_level_zero exceptions

```

- Implementation-specific state changes; for example, disabling an MMU.

- Control is transferred into the RED\_state trap table.

```

PC          ← RSTVADDR{63:8} :: 0100 00002
NPC         ← RSTVADDR{63:8} :: 0100 01002

```

### 12.6.2.4 Externally Initiated Reset (XIR) Traps

XIR traps are initiated by an external signal. They behave like an interrupt that cannot be masked by PSTATE.ie = 0 or PIL. Typically, XIR is used for critical system events such as power failure, reset button pressed, failure of external components that does not require a WDR (which aborts operations), or systemwide reset in a multiprocessor. See Chapter 16, *Resets* for further details.

If TL = MAXTL, then the virtual processor enters error\_state. The following virtual processor state changes occur during XIR (conceptually, in this order):

- The trap level is updated:

```

TL          ← min (TL + 1, MAXTL)

```

- Existing state is preserved.

```

TSTATE[TL].gl ← GL
TSTATE[TL].ccr ← CCR
TSTATE[TL].asi ← ASI
TSTATE[TL].pstate ← PSTATE
TSTATE[TL].cwp ← CWP
TPC[TL]       ← PC // (upper 32 bits zeroed if PSTATE.am = 1)
TNPC[TL]      ← NPC // (upper 32 bits zeroed if PSTATE.am = 1)
HTSTATE[TL].hpstate ← HPSTATE

```

- The trap type is set.

```

TT[TL]       ← 00316

```

- The Global Level register is updated.

```

GL          ← min (GL + 1, MAXGL)

```

- The PSTATE register is set as follows:

```

PSTATE.mm   ← 002 // TSO
PSTATE.pef  ← 1 // if an FPU is present, it is enabled
PSTATE.am   ← 0 // address masking is turned off
PSTATE.priv ← 0 // entering hyperprivileged mode
PSTATE.ie   ← 0 // interrupts are disabled
PSTATE.cle  ← 0 // big-endian is default for hyperprivileged mode
PSTATE.tle  is unchanged // (was unspecified in SPARC V9 specification)
PSTATE.tct  ← 0 // trap on CTI disabled

```

- The HPSTATE register is updated:
  - HPSTATE.red  $\leftarrow 1$  // enter RED\_state
  - HPSTATE.hpriv  $\leftarrow 1$  // enter hyperprivileged mode
  - HPSTATE.ibe  $\leftarrow 0$  // disable instruction breakpoints
  - HPSTATE.tlz  $\leftarrow 0$  // disable *trap\_level\_zero* exceptions
- Implementation-specific state changes; for example, disabling an MMU.
- Control is transferred into the RED\_state trap table.
  - PC  $\leftarrow RSTVADDR\{63:8\} :: 0110\ 0000_2$
  - NPC  $\leftarrow RSTVADDR\{63:8\} :: 0110\ 0100_2$

See *Externally Initiated Reset (XIR)* on page 547 and the documentation for specific processor implementations for more information.

### 12.6.2.5 Software-Initiated Reset (SIR) Traps

A software-initiated reset trap is initiated by execution of an SIR instruction in hyperprivileged mode. Hyperprivileged software uses the SIR trap as a panic operation or a metahypervisor trap. See Chapter 16, *Resets* for further details.

If  $TL = MAXTL$ , then the virtual processor enters *error\_state*.

Otherwise,  $TL < MAXTL$  as trap processing begins and the following virtual processor state changes occur (conceptually, in this order):

- The trap level is updated.
  - 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]  $\leftarrow PC$  // (upper 32 bits zeroed if PSTATE.am = 1)
  - TNPC[TL]  $\leftarrow NPC$  // (upper 32 bits zeroed if PSTATE.am = 1)
  - HTSTATE[TL].hpstate  $\leftarrow HPSTATE$
- The trap type is set.
  - TT[TL]  $\leftarrow 04_{16}$
- The Global Level register is updated.
  - GL  $\leftarrow \min(GL + 1, MAXGL)$
- The PSTATE register is set as follows:
  - PSTATE.mm  $\leftarrow 00_2$  // TSO
  - PSTATE.pef  $\leftarrow 1$  // if an FPU is present, it is enabled
  - PSTATE.am  $\leftarrow 0$  // address masking is turned off
  - PSTATE.priv  $\leftarrow 0$  // entering hyperprivileged mode

PSTATE.ie  $\leftarrow 0$  // interrupts are disabled  
 PSTATE.cle  $\leftarrow 0$  // big-endian is default for hyperprivileged mode  
 PSTATE.tle is unchanged // (*was unspecified in SPARC V9 specification*)  
 PSTATE.tct  $\leftarrow 0$  // trap on CTI disabled

- The HPSTATE register is updated:

HPSTATE.red  $\leftarrow 1$  // enter RED\_state  
 HPSTATE.hpriv  $\leftarrow 1$  // enter hyperprivileged mode  
 HPSTATE.ibe  $\leftarrow 0$  // disable instruction breakpoints  
 HPSTATE.tlz  $\leftarrow 0$  // disable *trap\_level\_zero* exceptions

- Implementation-specific state changes; for example, disabling an MMU.

- Control is transferred into the RED\_state trap table.

PC  $\leftarrow RSTVADDR\{63:8\} :: 1000\ 0000_2$   
 NPC  $\leftarrow RSTVADDR\{63:8\} :: 1000\ 0100_2$

See *Software-Initiated Reset (SIR)* on page 548 and the documentation for specific processor implementations for more information.

### 12.6.2.6 Nonreset Traps When the Virtual Processor Is in RED\_state

When a nonreset trap occurs while the virtual processor is in RED\_state, if  $TL = MAXTL$ , then the virtual processor enters error\_state.

Otherwise,  $TL < MAXTL$  as trap processing begins, the virtual processor remains in RED\_state, and the following virtual processor state changes occur (conceptually, in this order):

- The trap level is updated.

$TL \leftarrow TL + 1$

- 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]  $\leftarrow PC$  // (upper 32 bits zeroed if PSTATE.am = 1)  
 TNPC[TL]  $\leftarrow NPC$  // (upper 32 bits zeroed if PSTATE.am = 1)  
 HTSTATE[TL].hpstate  $\leftarrow HPSTATE$

- The trap type is preserved.

$TT[TL] \leftarrow$  trap type

- The Global Level register is updated.

$GL \leftarrow \min(GL + 1, MAXGL)$

- The PSTATE register is set as follows:
  - PSTATE.mm  $\leftarrow 00_2$  // TSO
  - PSTATE.pef  $\leftarrow 1$  // if an FPU is present, it is enabled
  - PSTATE.am  $\leftarrow 0$  // address masking is turned off
  - PSTATE.priv  $\leftarrow 0$  // entering hyperprivileged mode
  - PSTATE.ie  $\leftarrow 0$  // interrupts are disabled
  - PSTATE.cle  $\leftarrow 0$  // big-endian is default for hyperprivileged mode
  - PSTATE.tle  $\leftarrow$  unchanged // (*was unspecified in SPARC V9 specification*)
  - PSTATE.tct  $\leftarrow 0$  // trap on CTI disabled
- The HPSTATE register is updated:
  - HPSTATE.red  $\leftarrow 1$  // enter RED\_state
  - HPSTATE.hpriv  $\leftarrow 1$  // enter hyperprivileged mode
  - HPSTATE.ibe  $\leftarrow 0$  // disable instruction breakpoints
  - HPSTATE.tlz is unchanged
- For a register-window trap 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  $CWP \leftarrow CWP + 1$
  - endif
  - If  $(080_{16} \leq TT[TL] \leq 0BF_{16})$  // window spill trap
  - then  $CWP \leftarrow CWP + CANSAVE + 2$
  - endif
  - If  $(0C0_{16} \leq TT[TL] \leq 0FF_{16})$  // window fill trap
  - then  $CWP \leftarrow CWP - 1$
  - endif
- For non-register-window traps, CWP is not changed.
- Implementation-specific state changes; for example, disabling an MMU.
- Control is transferred into the RED\_state trap table.
  - PC  $\leftarrow RSTVADDR\{63:8\} :: 1010\ 0000_2$
  - NPC  $\leftarrow RSTVADDR\{63:8\} :: 1010\ 0100_2$

---

## 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 UltraSPARC 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 UltraSPARC Architecture 2005 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 = 024<sub>16</sub>-027<sub>16</sub>] (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.

- **cpu\_mondo** [TT = 07C<sub>16</sub>] (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, which is generated through software calls to hyperprivileged software. The standard software interface (API) to hyperprivileged software allows 64 bytes of data to be sent to one or more target virtual processors. When the CPU mondo queue has a valid entry, a *cpu\_mondo* exception is sent to the target virtual processor.
- **data\_access\_error** [TT = 032<sub>16</sub>] (Precise) — A hardware error occurred during a data access.
- **data\_access\_exception** [TT = 030<sub>16</sub>] (Precise) — An exception occurred on an attempted data access. Detailed information regarding the error is logged into the *ft* field of the DSFSR (Data Synchronous Fault Status register, ASI 58<sub>16</sub>, VA = 18<sub>16</sub>).

The conditions that may cause a *data\_access\_exception* exception are:

- **Privilege Violation** — 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.

- **Illegal Access to Noncacheable Page** — An access to a noncacheable page (TTE.cp = 0) (including cases with the TLB disabled) was attempted by an atomic load-store instruction (CASA, CASXA, SWAP, SWAPA, LDSTUB, or LDSTUBA) or an LDTXA instruction.
- **Illegal Access to Page That May Cause Side Effects** — An attempt was made to access a page which may cause side effects (TTE.e = 1) (including cases with the TLB disabled) by any type of load instruction with nonfaulting ASI.
- **Invalid ASI** — 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 *data\_access\_exception* 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 429 and *Partial Store ASIs* on page 430). This includes the following:
    - An attempt to use a Load Twin Extended Word (LDTXA) ASI (see *ASIs 10<sub>16</sub>, 11<sub>16</sub>, 16<sub>16</sub>, 17<sub>16</sub> and 18<sub>16</sub> (ASI\_\*AS\_IF\_USER\_\*)* on page 422) with any load alternate opcode other than LDTXA's (which is shared by LDTWA)
    - 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
    - An attempt to write to a read-only ASI-accessible register
- **Illegal Access to Non-Faulting-Only Page** — An attempt was made to access a non-faulting-only page (TTE.nfo = 1) by any type of load, store, or load-store instruction with an ASI other than a nonfaulting ASI (PRIMARY\_NO\_FAULT[\_LITTLE] or SECONDARY\_NO\_FAULT[\_LITTLE]).

<b>Forward Compatibility Note</b>	The next revision of the UltraSPARC Architecture is expected to replace <i>data_access_exception</i> with several more specific exceptions — one for each condition that currently can cause a <i>data_access_exception</i> . This will support slightly faster trap handling for these exceptions and allow elimination of the D-SFSR register.
---	--

- **data\_access\_MMU\_miss** [TT = 031<sub>16</sub>] (Precise) — During an attempted data access to memory,
  - (1) hardware tablewalk is enabled,
  - (2) the MMU detects that a translation lookaside buffer did not contain a translation for the data's virtual address, and
  - (3) the required TTE was not found in the configured TSBs.

- ***data\_invalid\_TSB\_entry*** [TT = 02B<sub>16</sub>] (Precise) — During an attempted data access,
  - (1) hardware tablewalk was enabled,
  - (2) the MMU detected that a translation lookaside buffer did not contain a translation for the virtual address, and
  - (3) the required TTE was found in the configured TSBs to be a real address, requiring real-to-physical address translation, and
  - (4) the real address cannot be translated to a physical address by hardware.
- ***data\_real\_translation\_miss*** [TT = 03F<sub>16</sub>] (Precise) — During an attempted real address data access, the MMU detected that a translation lookaside buffer (TLB) did not contain a translation for the real address (that is, a TLB miss occurred).
- ***dev\_mondo*** [TT = 07D<sub>16</sub>] (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.
- ***division\_by\_zero*** [TT = 028<sub>16</sub>] (Precise) — An integer divide instruction attempted to divide by zero.
- ***externally\_initiated\_reset*** (XIR) [TT = 003<sub>16</sub>] (Reset) — An external signal was asserted. This trap is used for catastrophic events such as power failure, reset button pressed, and system-wide reset in multiprocessor systems.
- ***fast\_data\_access\_MMU\_miss*** [TT = 068<sub>16</sub>] (Precise) — During an attempted data access to memory,
  - (1) hardware tablewalk was disabled (or is not implemented) and
  - (2) the MMU detected that a translation lookaside buffer did not contain a translation for the virtual address.
 Four trap vectors are allocated for this trap, allowing a TLB miss handler of up to 32 instructions to fit within the trap vector area.
- ***fast\_data\_access\_protection*** [TT = 06C<sub>16</sub>] (Precise) — During an attempted data write access (by a store or load-store instruction), the instruction had appropriate access privilege but the MMU signalled that the location was write-protected (write to a read-only location (TTE.w = 0)). Four trap vectors are allocated for this trap, allowing a trap handler of up to 32 instructions to fit within the trap vector area.  
 Note that on an UltraSPARC Architecture virtual processor, an attempt to read or write to a privileged location while in nonprivileged mode causes the higher-priority instead of this exception.
- ***fast\_instruction\_access\_MMU\_miss*** [TT = 064<sub>16</sub>] (Precise) — During an attempted instruction virtual address access,
  - (1) hardware tablewalk was disabled (or is not implemented) and
  - (2) the MMU detected a TLB miss.
 Four trap vectors are allocated for this trap, allowing a trap handler of up to 32 instructions to fit within the trap vector area.

- **fill\_n\_normal** [TT = 0C0<sub>16</sub>–0DF<sub>16</sub>] (Precise)
- **fill\_n\_other** [TT = 0E0<sub>16</sub>–0FF<sub>16</sub>] (Precise)
 

A RESTORE or RETURN instruction has determined that the contents of a register window must be restored from memory.
- **fp\_disabled** [TT = 020<sub>16</sub>] (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 = 021<sub>16</sub>] (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 = 022<sub>16</sub>] (Precise) — An FPop instruction generated an exception other than an IEEE\_754\_exception. Examples: the FPop is unimplemented or execution of an FPop requires software assistance to complete. The floating-point exception type is encoded in FSR.ftt.
- **guest\_watchdog** [TT = (see text)] (Precise, Disrupting) — The virtual processor was in nonprivileged or privileged mode, TL was  $\geq$  MAXPTL, and a precise or disrupting exception to privileged mode occurred. *guest\_watchdog* uses the same trap table entry (table offset 040<sub>16</sub>) as *watchdog\_reset*. When a *guest\_watchdog* trap occurs, the trap type (TT) value and priority of the exception that caused the trap are retained.
- **hstick\_match** [TT = 05E<sub>16</sub>] (Disrupting) — This interrupt indicates that a match between the System Tick (STICK) and the Hypervisor System Tick Compare (HSTICK\_CMPR) register has occurred. The event is recorded in the *hstick\_match\_pending* (hsp) bit of the Hypervisor Interrupt Pending (HINTP) register. The *hstick\_match* disrupting trap is recognized when HINTP.hsp = 1 and PSTATE.ie = 1 or HPSTATE.hpriv = 0; otherwise, it remains pending. HINTP.hsp provides a mechanism for hyperprivileged software to determine that an *hstick\_match* trap is pending while PSTATE.ie = 0 and to clear the condition without actually having to take the *hstick\_match* trap.
- **htrap\_instruction** [TT = 180<sub>16</sub>–1FF<sub>16</sub>] (Precise) — A Tcc instruction was executed in privileged or hyperprivileged mode, the trap condition evaluated to TRUE, and the software trap number was greater than 127. The trap is delivered in hyperprivileged mode, using the hyperprivileged mode trap base address (HTBA). See also *trap\_instruction* on page 489.
- **illegal\_instruction** [TT = 010<sub>16</sub>] (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.

**Note** | An unimplemented FPop instruction generates an *fp\_exception\_other* exception with ftt = 3, instead of an *illegal\_instruction* exception.

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*).
- A non-FPop 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 132.)

- 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, op2 (IMPDEP2A, IMPDEP2B), 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*.

- ***instruction\_access\_error*** [TT = 00A<sub>16</sub>] (Precise) — A hardware error occurred during an instruction access.
- ***instruction\_access\_MMU\_miss*** [TT = 009<sub>16</sub>] (Precise) — During an attempted instruction access (instruction fetch) from memory,
  - (1) hardware tablewalk was enabled,
  - (2) the MMU detected that a translation lookaside buffer did not contain a translation for the virtual address (that is, a TLB miss occurred), and
  - (3) the required TTE was not found in the configured TSBs.
- ***instruction\_access\_exception*** [TT = 008<sub>16</sub>] (Precise) — An exception occurred on an instruction access. The conditions that may cause an *instruction\_access\_exception* exception are:
  - **Privilege Violation** — An attempt to fetch an instruction from a privileged memory page (TTE.p = 1) while the virtual processor was executing in nonprivileged mode.
  - **Unauthorized Access** — An attempt to fetch an instruction from a memory page which was missing “execute” permission (TTE.ep = 0).
  - **No-Fault Only Access** — An attempt to fetch an instruction from a memory page which was marked for access only by nonfaulting loads (TTE.nfo = 1).
- ***instruction\_breakpoint*** [TT = 076<sub>16</sub>] (Precise) — This exception is generated if HPSTATE.ibc = 1 and the processor has detected a breakpoint condition based on the values in the Instruction Breakpoint Control register for the current instruction. As part of the trap, the HPSTATE.ibc bit is cleared (set to 0).

- ***instruction\_invalid\_TSB\_entry*** [TT = 02A<sub>16</sub>] (Precise) — During an attempted instruction access (instruction fetch), (1) hardware tablewalk was enabled, (2) the MMU detected that a translation lookaside buffer did not contain a translation for the virtual address, (3) the required TTE was found in the configured TSBs to be a real address, requiring real-to-physical address translation, and (4) the real address cannot be translated to a physical address by hardware.
- ***instruction\_real\_translation\_miss*** [TT = 03E<sub>16</sub>] (Precise) — During an attempted real address instruction access (instruction fetch), the MMU detected a TLB miss.
- ***internal\_processor\_error*** [TT = 029<sub>16</sub>] (Precise) — A serious internal error occurred in the virtual processor.
 

**IMPL. DEP. #402-S10:** The trap priority of the *internal\_processor\_error* exception is implementation dependent. Furthermore, its priority may vary within an implementation, based on the cause of the error being reported.
- ***interrupt\_level\_n*** [TT = 041<sub>16</sub>–04F<sub>16</sub>] (Disrupting) — SOFTINT{n} was set to 1 or an external interrupt request of level *n* was presented to the virtual processor and *n* > PIL.
 

<b>Implementation</b>	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>SOFTINT<sup>P</sup> Register (ASRs 20, 21, 22)</i> on page 80).
<b>Note</b>	
- ***LDDF\_mem\_address\_not\_aligned*** [TT = 035<sub>16</sub>] (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 = 034<sub>16</sub>] (Precise) — A load/store instruction generated a memory address that was not properly aligned according to the instruction, or a JMPL or RETURN instruction generated a non-word-aligned address. (See also *Special Memory Access ASIs* on page 422.)
- ***nonresumable\_error*** [TT = 07F<sub>16</sub>] (Disrupting) — 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.
- ***PA\_watchpoint*** [TT = 061<sub>16</sub>] (Precise) — The virtual processor has detected a load or store to a physical address specified by the PA Watchpoint register while PA watchpoints are enabled. Hyperprivileged software may reflect this trap back to privileged software as a synthetic *RA\_watchpoint* exception.
- ***pic\_overflow*** [TT = 04F<sub>16</sub>] (Disrupting) — A performance counter has overflowed and PIL < 15. Note that this exception shares a trap type, 04F<sub>16</sub>, with *interrupt\_level\_15*. The disrupting trap caused by *pic\_overflow* is conditioned by PSTATE.ie.
 

If PSTATE.ie = 1 and PIL < 15 when the possible counter overflow is detected and depending on the event being monitored by the counter, the disrupting trap may be reported prior to retirement of the instruction that incremented the counter to

cause the possible counter overflow. Upon entry to the trap handler, TPC points to an instruction that increments the performance counter and the counter is within some epsilon of overflow.

If `PSTATE.ie = 0` or `PIL = 15` when the possible overflow is detected, the trap remains pending and will be taken on the first instruction for which `PSTATE.ie = 1` and `PIL < 15`. In this case, TPC may not point to an instruction that increments the counter.

- **power\_on\_reset** (POR) [TT = 001<sub>16</sub>] (Reset) — An external signal was asserted. This trap is issued to bring a system reliably from the power-off to the power-on state.
- **privileged\_action** [TT = 037<sub>16</sub>] (Precise) — An action defined to be privileged has been attempted while in nonprivileged mode (`PSTATE.priv = 0` and `HPSTATE.hpriv = 0`), or an action defined to be hyperprivileged has been attempted while in nonprivileged or privileged mode (`HPSTATE.hpriv = 0`).  
Examples:
  - A data access by nonprivileged software using a restricted (privileged or hyperprivileged) ASI, that is, an ASI in the range 00<sub>16</sub> to 7F<sub>16</sub> (inclusively)
  - A data access by nonprivileged or privileged software using a hyperprivileged ASI, that is, an ASI in the range 30<sub>16</sub> to 7F<sub>16</sub> (inclusively)
  - Execution by nonprivileged software of an instruction with a privileged operand value
  - An attempt to read the TICK register by nonprivileged software when `TICK.npt = 1`
  - An attempt to access the PIC register (using RDPIC or WRPIC) while `PSTATE.priv = 0` and `PCR.priv = 1`
  - An attempt to execute a nonprivileged instruction with an operand value requiring more privilege than available in the current privilege mode.
- **privileged\_opcode** [TT = 011<sub>16</sub>] (Precise) — An attempt was made to execute a privileged instruction while `PSTATE.priv = 0`.
- **RED\_state\_exception** [TT = (see text)] (Precise) — Caused when `TL = MAXTL - 1` and a trap occurs, an event that brings the virtual processor into `RED_state`. Uses the trap vector entry reserved for trap type 005<sub>16</sub>, but the trap type recorded in TT is the trap type of the original exception that triggered *RED\_state\_exception*.
- **resumable\_error** [TT = 07E<sub>16</sub>] (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.
- **software\_initiated\_reset** (SIR) [TT = 004<sub>16</sub>] (Precise) — Caused by the execution of the SIR instruction. It allows system software to reset the virtual processor.
- **spill\_n\_normal** [TT = 080<sub>16</sub>–09F<sub>16</sub>] (Precise)
- **spill\_n\_other** [TT = 0A0<sub>16</sub>–0BF<sub>16</sub>] (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 = 036<sub>16</sub>] (Precise) — An attempt was made to execute an STDF or STDEFA instruction and the effective address was not doubleword aligned. (impl. dep. #110)
- **store\_error** [TT = 007<sub>16</sub>] (Deferred) — An error has been detected on a store instruction that prevents it from completing, but the error was detected after the store had passed its instruction retirement point. Since the store cannot be made globally visible, the software thread that issued the store must be terminated. Therefore, this is a termination deferred trap.
- **sw\_recoverable\_error** [TT = 040<sub>16</sub>] (Disrupting) — Indicates that one or more potentially recoverable errors have been detected in the virtual processor. A single *sw\_recoverable\_error* exception may indicate multiple errors and may occur asynchronously to instruction execution. When *sw\_recoverable\_error* causes a trap, the TPC and TNPC stacked by the trap do not necessarily indicate the instruction or data access that caused the error. (impl. dep. #31-V8-Cs10, #218-U3-Cs10)

**IMPL. DEP. #218-U3-Cs10:** Whether *sw\_recoverable\_error* exception is implemented is implementation dependent. If it does exist, it indicates that an error is detected in a processor core and its trap type is 40<sub>16</sub>.

<b>SPARC V9 Compatibility Note</b>	The <i>sw_recoverable_error</i> exception was called <i>async_data_error</i> in the SPARC V9 specification, which in turn superseded the less general SPARC V8 <i>data_store_error</i> exception.
--	---

- **tag\_overflow** [TT = 023<sub>16</sub>] (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 = 100<sub>16</sub>–17F<sub>16</sub>] (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.
- **trap\_level\_zero** [TT = 05F<sub>16</sub>] (Disrupting) — This exception indicates a simultaneous existence of three conditions:
  - *trap\_level\_zero* exceptions are enabled (HPSTATE.tlz = 1),
  - the virtual processor is in nonprivileged or privileged mode (HPSTATE.hpriv = 0), and
  - the trap level (TL) register's value is zero (TL = 0)

Upon entry to the trap handler for *trap\_level\_zero*, TPC points to the instruction that was about to be executed when all three of these conditions were met.

**Programming Note** | The purpose of this trap is to improve efficiency when descheduling a virtual processor. When a descheduling event occurs and the virtual processor is executing in privileged mode at  $TL > 0$ , hyperprivileged software can choose to enable the *trap\_level\_zero* exception (set  $HPSTATE.tlz \leftarrow 1$ ) and return to privileged mode, enabling privileged software to complete its  $TL > 0$  processing. When privileged code returns to  $TL = 0$ , this exception enables the hyperprivileged code to regain control and deschedule the virtual processor with low overhead.

- **unimplemented\_LDTW** [ $TT = 012_{16}$ ] (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** [ $TT = 013_{16}$ ] (Precise) — An attempt was made to execute an STTW instruction that is not implemented in hardware on this implementation (impl. dep. #108-V9).
- **watchdog\_reset** (WDR) [ $TT = 002_{16}$ ] (Reset) — This trap occurs in *error\_state* and causes a transition to *RED\_state* (impl. dep. #254-U3-Cs10).
- **VA\_watchpoint** [ $TT = 062_{16}$ ] (Precise) — The virtual processor has detected an attempt to access 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 physical 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. This exception is always masked in hyperprivileged mode; therefore, a *VA\_watchpoint* trap cannot occur in hyperprivileged mode (even if memory is accessed using *ASI\_AS\_IF\_USER\_PRIMARY* or *ASI\_AS\_IF\_USER\_SECONDARY*).

## 12.7.1 SPARC V9 Traps Not Used in UltraSPARC Architecture 2005

The following traps were optional in the SPARC V9 specification and are not used in UltraSPARC Architecture 2005:

- **async\_data\_error** [ $TT = 040_{16}$ ] (Disrupting) — This exception was superseded by the *sw\_recoverable\_error* exception.
- **data\_access\_protection** [ $TT = 033_{16}$ ] (Precise or Deferred) — This exception is generally superseded by *fast\_data\_access\_protection* (see page 484).

**IMPL. DEP. #fast\_ECC\_error** [TT = 070<sub>16</sub>] (Precise) — A single-bit or multiple-bit ECC error was detected. **202-U3:** Whether or not a *fast\_ECC\_error* trap exists is implementation dependent. If it does exist, it indicates that an ECC error was detected in an external cache and its trap type is 070<sub>16</sub>.

- **implementation\_dependent\_exception\_n** [TT = 077<sub>16</sub> - 07A<sub>16</sub>] 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 = 038<sub>16</sub>] (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* on page 249. (impl. dep. #111)
- **STQF\_mem\_address\_not\_aligned** [TT = 039<sub>16</sub>] (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 332. (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 452, 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 Note** | All spill traps can set CWP by using the calculation:  
 $CWP \leftarrow (CWP + CANSAVE + 2) \bmod N\_REG\_WINDOWS$   
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 RETRY 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:

$$CLEANWIN \leftarrow (CLEANWIN + 1)$$



**Note: This chapter is undergoing final review; please check back later for a copy of UltraSPARC Architecture 2005 containing the final version of this chapter.**

## 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 496).

<b>Programming Note</b>	An interrupt request packet is sent by an interrupt source (through an implementation-specific mechanism) 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 496.
- **Software Interrupt Register (SOFTINT)** on page 496.

- **Interrupt Queues** on page 497.
- **Interrupt Traps** on page 500.
- **Strand Interrupt ID Register (STRAND\_INTR\_ID)** on page 501.

---

## 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. Similarly, hyperprivileged software can schedule interrupt vectors for privileged software running on the same virtual processor by setting bits in SOFTINT.

See *SOFTINT<sup>P</sup> Register (ASRs 20, 21, 22)* on page 80 for a detailed description of the SOFTINT register.

- |                         |   |
|-------------------------|---|
| <b>Programming Note</b> | The SOFTINT register (ASR 16 <sub>16</sub> ) 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 <i>n</i> . |
| <b>Programming Note</b> | The SOFTINT mechanism is independent of the “mondo” interrupt mechanism mentioned in <i>Interrupt Queues</i> on page 497. 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<sup>P</sup> 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\_SET<sup>P</sup> Pseudo-Register (ASR 20)* on page 82 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<sup>P</sup> 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\_CLR<sup>P</sup> Pseudo-Register (ASR 21)* on page 82 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 mondos<sup>1</sup>
- CPU mondos
- Resumable errors
- Nonresumable errors

<sup>1</sup> “mondo” is a historical term, referring to the name of the original UltraSPARC 1 bus transaction in which these interrupts were introduced

New interrupt entries are appended to the tail of a queue (by hardware or by hyperprivileged software) 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.

**IMPL. DEP. #421-S10:** It is implementation dependent whether interrupt queue head and tail registers (a) are datatype-agnostic “scratch registers” used for communication between privileged and hyperprivileged software, in which case their contents are defined purely by software convention, or (b) are maintained to some degree by virtual processor hardware, imposing a fixed meaning on their contents.

**Programming Note** | If the contents of Queue Head and Tail registers are set only by software convention in a given implementation, software could place any type of data in them (such as addresses, address offsets, or index values).  
It is expected that Queue Head and Tail registers will typically contain a byte offset from the base of an appropriately-aligned queue region in memory.

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	Hyper-privileged mode Access
CPU Mondo Queue Head	25 <sub>16</sub> (ASI_QUEUE)	3C0 <sub>16</sub>	RW	R/W
CPU Mondo Queue Tail	25 <sub>16</sub> (ASI_QUEUE)	3C8 <sub>16</sub>	R or RW†	R/W
Device Mondo Queue Head	25 <sub>16</sub> (ASI_QUEUE)	3D0 <sub>16</sub>	RW	R/W
Device Mondo Queue Tail	25 <sub>16</sub> (ASI_QUEUE)	3D8 <sub>16</sub>	R or RW†	R/W
Resumable Error Queue Head	25 <sub>16</sub> (ASI_QUEUE)	3E0 <sub>16</sub>	RW	R/W
Resumable Error Queue Tail	25 <sub>16</sub> (ASI_QUEUE)	3E8 <sub>16</sub>	R or RW†	R/W
Nonresumable Error Queue Head	25 <sub>16</sub> (ASI_QUEUE)	3F0 <sub>16</sub>	RW	R/W
Nonresumable Error Queue Tail	25 <sub>16</sub> (ASI_QUEUE)	3F8 <sub>16</sub>	R or RW†	R/W

† see **IMPL. DEP.#422-S10**

**IMPL. DEP. #422-S10:** It is implementation dependent whether tail registers are writable in privileged mode. If a tail register is read-only in privileged mode, an attempt to write to it causes a *data\_access\_exception* exception. If a tail register is writable in privileged mode, an attempt to write to it results in undefined behavior.

**Implementation Note** | Although Queue Head and Tail registers behave as registers, they may or may not be implemented using actual *hardware* registers. For example, they may reside in memory, mapped by a mechanism visible only to hyperprivileged software. In any case, the means by which Queue Head and Tail registers are implemented is not visible to privileged software.

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.

**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 UltraSPARC Architecture 2005: *cpu\_mondo*, *dev\_mondo*, *resumable\_error*, and *nonresumable\_error*. The first three (*cpu\_mondo*, *dev\_mondo*, and *resumable\_error*) are all generated by hardware, while *nonresumable\_error* is generated by hyperprivileged software. See Chapter 12, *Traps*, for details.

UltraSPARC Architecture 2005 also supports the *interrupt\_level\_n* traps defined in the SPARC V9 specification.

How interrupts are delivered is implementation-specific; see the relevant implementation-specific Supplement to this specification for details.

---

## 13.5 Strand Interrupt ID Register (STRAND\_INTR\_ID)

The STRAND\_INTR\_ID per-virtual-processor register allows software to assign a 16-bit interrupt ID to a virtual processor that is unique within the system. This is important, to enable virtual processors to receive interrupts. See *Strand Interrupt ID Register (STRAND\_INTR\_ID)* on page 520 for details.



# Memory Management

---

An UltraSPARC 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.

In UltraSPARC Architecture 2005, memory management is implementation-specific. Basic concepts are described in this chapter, but see the relevant processor-specific Supplement to this specification for a detailed description of a particular processor's memory management facilities.

This appendix describes the Memory Management Unit, as observed by hyperprivileged software, in these sections:

- **Virtual Address Translation** on page 503.
- **TSB Translation Table Entry (TTE)** on page 505.
- **Translation Storage Buffer (TSB)** on page 508.
- **Faults and Traps** on page 509.

---

## 14.1 Virtual Address Translation

The MMUs may support up to four page sizes: 8 KBytes, 64 KBytes, 4 MBytes, and 256 MBytes. 8-KByte, 64-KByte and 4- MByte page sizes must be supported; other page sizes are optional.

Each MMU consists of one or more Translation Lookaside Buffers (TLBs), and may include micro-TLB structures. Separate Instruction and Data MMUs (IMMU and DMMU, respectively) may be provided to enable concurrent virtual-to-physical address translations for instruction and data.

**IMPL. DEP. #222-U3:** TLB organization is implementation dependent.

Privileged software manages virtual-to-real address translations. Hyperprivileged software manages real-to-physical 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.

The MMU TLBs act as independent caches of the software translation table, providing appropriate concurrency for virtual-to-physical address translation.


Hyperprivileged software maintains translation information for real-to-physical translations.

During a memory access, one or more TLBs are searched for a VA (or RA) translation. A TLB hit is indicated when the virtual address, context ID, and partition ID (or real address and partition ID) match an entry in the TLB.

A TLB miss is indicated when no such match occurs, and the MMU immediately traps to hyperprivileged software for TLB miss processing. The TLB miss handler can fill the TLB by any available means, but it is likely to take advantage of the TLB miss support features provided by the MMU, since the TLB miss handler is time-critical code.

A conceptual view of privileged-mode memory management the MMU is shown in FIGURE 14-1. The TLBs, which are part of the MMU hardware, are small and fast. 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.

The UltraSPARC Architecture provides a memory partitioning mechanism that allows for multiple partitions, each containing its own real address space. Hyperprivileged software provides real address to physical address translations.




**FIGURE 14-1** Conceptual View of the MMU

Aliasing of multiple virtual addresses to the same physical address is supported. However, the reverse case of multiple mappings from one virtual address to multiple physical addresses producing a multiple TLB match is detected in hardware as a multiple tag hit TLB error.

## 14.2 TSB Translation Table Entry (TTE)

The Translation Storage Buffer (TSB) Translation Table Entry (TTE) 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 14-2 and described in TABLE 14-1.



**FIGURE 14-2** Translation Storage Buffer (TSB) Translation Table Entry (TTE)

**TABLE 14-1** 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;"> <b>Programming Note</b> 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 context_id 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 ASI_PRIMARY_NO_FAULT{ _LITTLE } or ASI_SECONDARY_NO_FAULT{ _LITTLE } are translated. Any other data access with the D/UMMU TTE.nfo = 1 will trap with a <i>data_access_exception</i> (with SFSR.ft = 10 <sub>16</sub> ). An instruction fetch access to a page with the IMMU TTE.nfo = 1 results in an <i>instruction_access_exception</i> exception.
Data - 61:56	soft2	Software-defined field, provided for use by the operating system. The soft2 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	t_addr	Target address from TSB (Real Address {55:13} or Physical Address {55:13}). UltraSPARC Architecture TLBs store physical addresses, not real addresses. Hyperprivileged code is responsible for translation between real and physical addresses. Whether this field contains a Real or Physical address is determined by the bit in the corresponding MMU TSB Configuration register.  <b>IMPL. DEP. #224-U3:</b> Physical address width support by the MMU is implementation dependent in the UltraSPARC Architecture; minimum PA width is 40 bits. <b>IMPL. DEP. #238-U3:</b> 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.
Data - 12	ie	Invert Endianness. If ie = 1 for a page, accesses to the page are processed with inverse endianness from that specified by the instruction (big for little, little for big). <b>Note:</b> This bit is intended to be set to 1 primarily for noncacheable accesses. The performance of cacheable accesses may be degraded as if the access missed the D-cache. <b>IMPL. DEP. #__:</b> The ie bit in the IMMU is ignored during ITLB operation. It is implementation dependent if it is implemented and how it is read and written.

**TABLE 14-1** TSB TTE Bit Description (2 of 3)

Bit	Field	Description														
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 e-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 e bit causes the prefetch instruction to be treated as a nop, but does not prevent normal (hardware) instruction prefetching.</p> <p><b>Note:</b> The e bit does not force a noncacheable access. It is expected, but not required, that the cp and cv bits will be set to 0 when the e bit is set to 1. If both the cp and cv bits are set to 1 along with the e bit, the result is undefined.</p> <p><b>Note:</b> The e bit and the nfo bit are mutually exclusive; both bits should never be set to 1 in any TTE.</p>														
Data – 10 Data – 9	cp, cv	<p>The cacheable-in-physically-indexed-cache bit and cacheable-in-virtually-indexed-cache bit determine the cacheability of the page. Given an implementation with a physically indexed instruction cache, a virtually indexed data cache, and a physically indexed unified second-level cache, the following table illustrates how the cp and cv bits could be used:</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th rowspan="2">Cacheable (cp, cv)</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 VA-indexed)</th> </tr> </thead> <tbody> <tr> <td>00, 01</td> <td>Noncacheable</td> <td>Noncacheable</td> </tr> <tr> <td>10</td> <td>Cacheable L2-cache, I-cache</td> <td>Cacheable L2-cache</td> </tr> <tr> <td>11</td> <td>Cacheable L2-cache, I-cache</td> <td>Cacheable L2-cache, D-cache</td> </tr> </tbody> </table>	Cacheable (cp, cv)	Meaning of TTE when placed in:		I-TLB (Instruction Cache PA-indexed)	D-TLB (Data Cache VA-indexed)	00, 01	Noncacheable	Noncacheable	10	Cacheable L2-cache, I-cache	Cacheable L2-cache	11	Cacheable L2-cache, I-cache	Cacheable L2-cache, D-cache
Cacheable (cp, cv)	Meaning of TTE when placed in:															
	I-TLB (Instruction Cache PA-indexed)	D-TLB (Data Cache VA-indexed)														
00, 01	Noncacheable	Noncacheable														
10	Cacheable L2-cache, I-cache	Cacheable L2-cache														
11	Cacheable L2-cache, I-cache	Cacheable L2-cache, D-cache														
		<p>The MMU does not operate on the cacheable bits but merely passes them through to the cache subsystem. The cv bit in the IMMU is read as zero and ignored when written.</p> <p><b>IMPL. DEP. #226-U3:</b> Whether the cv bit is supported in hardware is implementation dependent in the UltraSPARC Architecture. The cv bit in hardware should be provided if the implementation has virtually indexed caches, and the implementation should support hardware unaliasing for the caches.</p>														
Data – 8	p	<p>Privileged. If p = 1, only privileged software can access the page mapped by the TTE. If p = 0 and an access to the page is attempted by nonprivileged mode (<code>PSTATE.priv = 0</code>), then the MMU signals an <i>instruction_access_exception</i> exception or <i>data_access_exception</i> exception.</p>														
Data – 7	ep	<p>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>instruction_access_exception</i> exception.</p> <p><b>IMPL. DEP. #___:</b> Some UltraSPARC Architecture ITLB implementations may not implement the ep bit, and present the <i>instruction_access_exception</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.</p>														

**TABLE 14-1** TSB TTE Bit Description (3 of 3)

Bit	Field	Description																				
Data – 6	w	Writable. If w = 1, the page mapped by this TTE has write permission granted. Otherwise, write permission is not granted, and the MMU causes a <i>fast_data_access_protection</i> trap if a write is attempted. <b>IMPL. DEP. #___</b> : The w bit in the IMMU is ignored during ITLB operation. It is model dependent if the bit is implemented and how it is written and read.																				
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" style="margin-left: 20px;"> <thead> <tr> <th><u>sz</u></th> <th><u>Page Size</u></th> </tr> </thead> <tbody> <tr> <td>0000</td> <td>8 Kbyte</td> </tr> <tr> <td>0001</td> <td>64 Kbyte</td> </tr> <tr> <td>0010</td> <td><i>Reserved</i></td> </tr> <tr> <td>0011</td> <td>4 Mbyte</td> </tr> <tr> <td>0100</td> <td><i>Reserved</i></td> </tr> <tr> <td>0101</td> <td>256 Mbyte</td> </tr> <tr> <td>0110</td> <td><i>Reserved</i></td> </tr> <tr> <td>0111</td> <td><i>Reserved</i></td> </tr> <tr> <td>1000-1111</td> <td><i>Reserved</i></td> </tr> </tbody> </table>	<u>sz</u>	<u>Page Size</u>	0000	8 Kbyte	0001	64 Kbyte	0010	<i>Reserved</i>	0011	4 Mbyte	0100	<i>Reserved</i>	0101	256 Mbyte	0110	<i>Reserved</i>	0111	<i>Reserved</i>	1000-1111	<i>Reserved</i>
<u>sz</u>	<u>Page Size</u>																					
0000	8 Kbyte																					
0001	64 Kbyte																					
0010	<i>Reserved</i>																					
0011	4 Mbyte																					
0100	<i>Reserved</i>																					
0101	256 Mbyte																					
0110	<i>Reserved</i>																					
0111	<i>Reserved</i>																					
1000-1111	<i>Reserved</i>																					

## 14.3 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 TLB in the event of a TLB miss.

Inclusion of the TLB entries in the TSB is not required; that is, translation information that is not present in the TSB can exist in the TLB.

### 14.3.1 TSB Indexing Support

Hardware TSB indexing support via TSB pointers should be provided for the TTEs. Hardware tablewalk uses the TSB pointers. If the hardware tablewalk is disabled, the TLB miss handler software can use the TSB pointers.

## 14.3.2 TSB Cacheability

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.

## 14.3.3 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 14-3. The constant  $n$  is determined by the size field in the TSB register; it can range from 512 to an implementation-dependent number.




FIGURE 14-3 TSB Organization

---

## 14.4 Faults and Traps

The traps recorded by the MMU are listed in TABLE 14-2. For a detailed description of each trap, see Chapter 12, *Traps*. All listed traps are precise traps.

**TABLE 14-2** MMU Trap Types, Causes, and Stored State Register Update Policy

Ref #	Trap Name	Trap Cause	Registers Updated (Stored State in MMU)			# of Trap Vectors Used	
			IMMU Tag Access	D-SFAR	D/ UMMU Tag Access		
1.	<i>fast_instruction_access_MMU_miss</i>	I-TLB miss	X			64 <sub>16</sub>	4
2.	<i>instruction_access_exception</i>	Several (see below)	X			08 <sub>16</sub>	1
3.	<i>fast_data_access_MMU_miss</i>	D-TLB miss		X	X	68 <sub>16</sub>	4
4.	<i>data_access_exception</i>	Several (see below)		X	X <sup>†</sup>	30 <sub>16</sub>	1
5.	<i>fast_data_access_protection</i>	Protection violation		X	X	6C <sub>16</sub>	4
6.	<i>privileged_action</i>	Use of privileged ASI		X		37 <sub>16</sub>	1
7.	<i>PA_watchpoint</i>	Watchpoint hit		X		61 <sub>16</sub>	1
7b.	<i>VA_watchpoint</i>	Watchpoint hit		X		62 <sub>16</sub>	1
8.	<i>mem_address_not_aligned,</i> <i>*_mem_address_not_aligned</i>	Misaligned memory operation		impl. dep. #237-U3		35 <sub>16</sub> , 36 <sub>16</sub> , 38 <sub>16</sub> , 39 <sub>16</sub>	1

<sup>†</sup>The contents of the context\_id field of the DMMU Tag Access register are undefined after a *data\_access\_exception*.

# Chip-Level Multithreading (CMT)

---

An UltraSPARC Architecture 2005 processor may include multiple virtual processors on the same processor module to provide a dense, high-throughput system. This may be achieved by having a combination of multiple physical processor cores and/or multiple strands (threads) per physical processor core.

This chapter specifies a common interface between hardware and software for such products, referred to here as chip-level multithreaded processors (CMTs). It addresses issues common to CMT processors, regardless of the microarchitecture of the individual physical processor cores, in the following sections:

- **Overview of CMT** on page 511.
- **Accessing CMT Registers** on page 515.
- **CMT Registers** on page 518.
- **Disabling and Parking Virtual Processors** on page 522.
- **Reset and Trap Handling** on page 532.
- **Error Handling in CMT Processors** on page 535.
- **Additional CMT Software Interfaces** on page 540.
- **Performance Issues for CMT Processors** on page 541.
- **Recommended Subset for Single-Strand Processors** on page 541.
- **Machine State Summary** on page 543.

---

## 15.1 Overview of CMT

A broad range of designs may fall under the definition of CMT. The interface specified here is intended to provide a set of common behaviors to enable operating system software and other privileged software to be common across UltraSPARC Architecture 2005 processors. This interface is not complete, as a range of implementation dependent features will exist to configure and control these processors.

The CMT Programming Model describes a set of privileged registers that are used for identification and configuration of CMT processors. Equally important, the CMT Programming Model describes certain behavior that is common across CMT implementations. The set of registers and the common behavior are covered in the following sections, grouped by topic.

UltraSPARC Architecture 2005 processors that are not CMT processors (are single-threaded) should implement a subset of the CMT interface. This enables those virtual processors to be more easily integrated into products that may also contain CMT processors and also enables more consistent software to be deployed across future products. See *Recommended Subset for Single-Strand Processors* on page 541 for additional information on non-CMT processor implementations.

## 15.1.1 CMT Definition

An UltraSPARC Architecture 2005 CMT processor is defined by its externally-visible nature and not by its internal organization. Note that this definition is not fully consistent with the common hardware definition of CMT. The following section gives some background terminology, followed by a description of the CMT definition.

### 15.1.1.1 Background Terminology

**Thread.** Historically, the term *thread* is overused and ambiguous; software and hardware have used it differently. From a software (operating system) perspective, the term “thread” refers to an entity that:

- Can be executed on underlying hardware
- Is scheduled
- May or may not be actively running on hardware at any given time
- May migrate around the hardware of a system.

From the hardware perspective, the term “multithreaded processor” refers to a processor that can run multiple software threads simultaneously.

To avoid confusion, the term “thread” in UltraSPARC Architecture 2005 is used exclusively in the manner that it is used by software (specifically, the operating system). A thread can be viewed in a practical sense as a Solaris™ process or lightweight process (LWP).

**Strand.** The term *strand* is used to identify the state that hardware must maintain in order to execute a software thread. Specifically, a “strand” is the software-visible architected state (PC, NPC, general-purpose registers, floating-point registers, condition codes, status registers, ASRs, etc.) of a thread plus any microarchitecture

state required by hardware for its execution. “Strand” replaces the ambiguous term “hardware thread.” The number of strands in a processor defines the number of threads that an operating system can schedule on that processor at any given time.

**Pipeline.** The term *pipeline* refers to an execution pipeline. It is a loose term for the basic collection of hardware needed to execute instructions. A pipeline may be used by one or more strands, in order to execute instruction from one or more threads. *Synonym: microcore.*

**Physical Core.** The term *physical processor core*, or just *physical core*, is similar to the term “pipeline” but represents a broader collection of hardware. A physical core includes one or more execution pipelines and associated structures, such as caches, that are required for executing instructions from one or more software threads. A physical core contains one or more strands. The physical core provides the necessary resources for the threads on each strand to make forward progress at a reasonable rate. A multistranded physical core can execute multiple software threads by time-multiplexing resources, partitioning resources, or any combination thereof.

The delineations among the terms strand, pipeline, and physical core are not precise. Among different microarchitecture organizations the scope of the terms may vary. In general, in a specific microarchitecture it will be apparent what constitutes a physical core. A physical core will be a highly integrated unit with a clearly defined interface to more distant levels of the memory hierarchy and the system interface unit. A physical core will contain a defined number of strands, that is, a maximum number of software threads that may be scheduled on it at any given time.

**Processor.** A *processor* is the unit on which a shared interface is provided to control the configuration and execution of a collection of strands. A processor contains one or more physical cores, each of which contains one or more strands. Physically, a *processor* is a physical module that plugs into a system. A processor is expected to appear logically as a single agent on the system interconnect fabric.

Therefore, a simple processor that can only execute one thread at a time (for example, an UltraSPARC I processor) would contain a single physical core which is single-stranded. A processor that follows the academic model of simultaneous multithreading (SMT) would contain a single physical core, where that physical core supports multiple strands in order to execute multiple simultaneous threads (multi-stranded physical core). A processor that follows the academic model of a chip multi-processor (CMP) would be a processor with multiple physical cores, each supporting only a single strand. A processor may also contain multiple physical cores, where each physical core is multi-stranded.

**Virtual Processor.** The term *virtual processor* is used to identify each strand in a processor. Each virtual processor corresponds to a specific strand on a specific physical core, where multiple physical cores, each with multiple strands, may exist. In most respects a virtual processor appears to the system and to operating system

software as a processing unit equivalent to a traditional single-stranded processor (as in UltraSPARC I). Each virtual processor is capable of having interrupts directed specifically to it. At any given time, an operating system can have a different thread scheduled on each virtual processor.

The UltraSPARC Architecture 2005 CMT architecture (software interface) described in this chapter is independent of the specific method by which multiple virtual processors are implemented. The term “virtual processor” is generally used instead of “strand” because “strand” is commonly associated with multistranded physical cores.

**CPU.** The term *CPU* is ambiguous in reference to processors with multiple virtual processors. The term could potentially refer to a virtual processor or to an entire processor. Therefore, the term “CPU” is considered ambiguous and will not be used in this document.

**CMT.** *CMT* is an abbreviation for “Chip MultiThreaded” or “Chip MultiThreading” (depending on context). A CMT processor is one containing one or more virtual processors.

### 15.1.1.2 CMT Definition

*CMT*, as defined in UltraSPARC Architecture 2005, applies to all SPARC virtual processors. A processor containing a single virtual processor (strand) is a special case, covered in *Recommended Subset for Single-Strand Processors* on page 541. The CMT interface is the same whether multiple strands are provided by multiple physical cores, a single physical core with multiple strands, or multiple physical cores each with multiple strands.

A *virtual processor* is a processing entity that can execute a software thread. A virtual processor has a number of key characteristics and includes all the architecturally visible state, as defined elsewhere in this specification, to execute a thread (general purpose registers, floating-point registers, process state, status registers, condition codes, etc.). A virtual processor is the smallest unit to which an interrupt can be delivered. The addressability of interrupts to individual virtual processors is a very important aspect of the CMT programming interface. An UltraSPARC Architecture 2005 implementation must provide sufficient resources so that every virtual processor within the processor makes forward progress at a reasonable rate.

Each virtual processor contains a separate instance of all user-visible architected state; that is, nonprivileged architected state is per-virtual processor.

The privileged and hyperprivileged architected state of a processor falls into four classes (described in *Classes of CMT Registers* on page 516), based on the degree of sharing among virtual processors.

<b>Implementation Note</b>	The UltraSPARC Architecture 2005 applies to a single physical processor chip. In a multiple-chip system, the UltraSPARC Architecture 2005 applies to each processor chip.
----------------------------	---

## 15.1.2 General CMT Behavior

In general, each virtual processor of a CMT processor behaves functionally as if it was an independent processor. This is an important aspect of CMT processors because user code running on a virtual processor does not need to know whether or not that virtual processor is part of a CMT processor. At a high level, most privileged code in an operating system can treat virtual processors of a CMT processor as if each was an independent processor. Some software (for example, boot, error, and diagnostic) must be aware that it is executing on a CMT processor. This chapter deals chiefly with the interface between this software and a CMT processor.

Each virtual processor of a CMT processor obeys the same memory model semantics as if it was an independent processor. All software designed to run in a multiprocessing environment, including thread libraries, must be able to operate on a CMT processor without modification.

There are significant performance implications of CMT processors, especially when shared resources (such as caches) exist within a CMT processor. The virtual processors' proximity will potentially mean drastically different costs for communicating between two virtual processors on the same CMT processor compared to communicating between two virtual processors on different CMT processors. This adds another degree of non-uniform memory access (NUMA) to a system. For high performance, the operating system, and even some user applications, will want to program specifically for the NUMA nature of CMT processors. There may also be resource contention issues between virtual processors on the same CMT processor. *Performance Issues for CMT Processors* on page 541 discusses some key performance issues related to CMT processors.

---

## 15.2 Accessing CMT Registers

A key part of the CMT Programming Model is a set of privileged registers. This section covers how these registers are organized and accessed. The registers can be accessed by software running on a virtual processor of the CMT processor.

CMT-specific registers can be accessed by privileged software running on a virtual processor, using Load and Store Alternate (notably, LDXAs and STXAs) instructions that provide an address space identifier value and a (virtual) address. The CMT Programming Model defines address space identifiers and associated virtual addresses (VAs) for accessing the CMT-specific registers.

## 15.2.1 Classes of CMT Registers

Nonprivileged architected state, including registers visible to nonprivileged software, is (or at least appears to be) per-virtual-processor.

Privileged architected state, including registers visible to privileged software, is (or at least appears to be) per-virtual-processor.

The hyperprivileged architected state of a processor falls into four categories:

- Per-virtual-processor (per-strand) registers, of which each virtual processor has a private (not shared) copy
- Subset-shared registers, where a copy of each register is shared by a non-overlapping subset of virtual processors<sup>1</sup>.
- Per-physical-core shared registers (a special case of subset-shared registers), where a copy of each register is shared by all virtual processors contained within a physical core.
- Processor-shared CMT registers, in which a single copy of each register is shared by all virtual processors in the processor

Registers that are read-only in privileged mode (for example, TICK) need not be strictly implemented as per-virtual-processor registers; they may be implemented in one of the “shared” categories above, such that their shared nature is not visible to privileged software.

CMT-specific registers of all classes can be accessed as ASI-mapped registers through hyperprivileged software running on a virtual processor. Software running on a given virtual processor can access:

- all the per-virtual processor registers belonging to the virtual processor on which it is running
- the per-physical-core shared registers belonging to the physical core on which it is running
- subset-shared registers for any group of virtual processors to which the virtual processor on which it is running belongs
- all processor-shared registers

<sup>1</sup> Currently, no architectural CMT registers fall into this category. It is defined here for completeness, because registers in this category may need to exist as implementation-specific registers

In nonprivileged or privileged mode, it is normally not possible for a virtual processor on one physical core to address (much less, read) the per-physical-core registers of another physical core. On some implementations it may be possible for a virtual processor on one physical core to address the per-physical-core registers of another physical core, but *only* in hyperprivileged mode or if hyperprivileged software grants such privileges to software running at a lower privilege level.

The semantics for accessing the CMT registers through the ASI interface are described in *Accessing CMT Registers Through ASIs* on page 517.

## 15.2.2 Accessing CMT Registers Through ASIs

Each CMT-specific register is accessible through a restricted ASI (accessible only in hyperprivileged software). The ASI number and virtual address corresponding to each CMT register are described later in this chapter.

Each virtual processor can access the per-physical-core CMT registers associated with that virtual processor. The implementation must guarantee that accesses to per-physical-core registers follow sequential semantics on the virtual processor with which they are associated.

Each virtual processor can access all the per-processor shared CMT registers on its processor. An update to a per-processor shared register from one virtual processor will be visible to all other virtual processors that share that register. The ordering of accesses to per-processor shared registers from different virtual processors is not defined, but an implementation must guarantee that:

- Accesses to a shared register from the same virtual processor follow sequential semantics.
- If multiple virtual processors attempt to store to a shared CMT register at the same time, the value observed in (readable from) the register will always be that written by one of those stores. That is, a store to a CMT register must be performed atomically on all bits of the register. In the case of the `STRAND_RUNNING` register, there is a third option — a write to the register may be dropped (ignored) entirely in certain situations (for details, see *Simultaneous Updates to the STRAND\_RUNNING Register* on page 528).

There may be additional implementation-enforced restrictions on updates to some CMT registers.

All CMT registers are 64-bit registers, although some of the bits of individual registers can be reserved or defined to contain a fixed value in a given implementation. *Reserved* register fields should always be written by software with values of those fields previously read from that register or with zeroes and they should read as zero in hardware (see *Reserved Opcodes and Instruction Fields* on page

132). Software intended to run on future versions of CMTs should not assume that these fields will read as 0 or any other particular value. This convention simplifies future expansion of the CMT interface.

A CMT register is accessed through load and store instructions, using a defined ASI number and virtual address. CMT registers can only be accessed in hyperprivileged mode. An attempt to access a CMT register in nonprivileged or privileged mode results in a *privileged\_action* exception.

Only the LDXA or LDDFA instruction can be used to read a CMT register. Only the STXA or STDFA instruction can be used to store to a CMT register. An attempt to access a CMT register with any other instruction results in a *data\_access\_exception* exception. An attempt to write to a read-only CMT register with a STXA instruction results in a *data\_access\_exception* (invalid ASI) exception.

---


## 15.3 CMT Registers

In this section, the registers used to control operation of a processor in a CMT implementation are described. For each register defined in this document, a six-column quick-reference table is provided that specifies the key attributes of the register, as follows:

Column Heading	Meaning of column contents
<b>Register Name</b>	The name of the CMT register
<b>ASI # (Name)</b>	The address space identifier number used for accessing the register from software running on the CMT processor (and the recommended ASI name for use in assembly-language hyperprivileged software)
<b>VA</b>	The virtual address used for accessing the register from software running on the CMT processor
<b>Scope</b>	The scope of sharing for the register — whether the register is a “per-virtual processor” (per-strand) register, or a single instance of a register that is “shared” among the virtual processors within a physical core (per-core), “shared” among a subset of virtual processors within a physical core (per-subset), or “shared” among all the virtual processors within a processor (per-proc).
<b>Access</b>	Whether software access to the register is read/write (RW), read-only (R only), write-only (W only), Write-1-to-Set (W1S), or Write-1-to-Clear (W1C)
<b>Note</b>	Any additional information

## 15.3.1 Strand ID Register (STRAND\_ID)

Register Name	ASI # (Name)	VA	Scope	Access	Note
STRAND_ID	63 <sub>16</sub> (ASI_CMT_PER_STRAND)	10 <sub>16</sub>	per-strand	R only	



**FIGURE 15-1** STRAND\_ID Register

STRAND\_ID is a read-only, per-virtual processor register that holds the ID value assigned by hardware to each implemented virtual processor. The ID value is unique within the CMT processor.

As shown above, the STRAND\_ID register has three fields:

1. `strand_id`, which represents this virtual processor's number, as assigned by hardware. The strand ID is encoded in 6 bits.
2. `max_strand_id`, which is the bit-position index (bit number) of the most significant '1' bit in the STRAND\_AVAILABLE register. This is the Strand ID of the highest-numbered implemented virtual processor in this CMT processor.
3. `max_strand_id_per_core`, which specifies the number of strands minus one that are implemented on each physical core. For a single-stranded processor, `max_strand_id_per_core` will be 0.

Many other CMT-specific registers provide a bit mask in which each bit corresponds to an individual virtual processor. For these registers, the `strand_id` field indicates which bit of a bit mask corresponds to this specific virtual processor.

**Strand Numbering Convention.** The numbering of virtual processors (strands) may or may not be contiguous; system software may only assume that each strand ID is unique within a CMT processor. In general, virtual processors should be numbered in a sequential, contiguous series starting with strand number 0. When numbering the virtual processors within a CMT processor, this convention appears straightforward. There are cases, however, where this might not be so simple. This numbering convention is recommended but not required.

In a CMT processor designed with many virtual processors, some physical cores in a manufactured CMT processor may fail to function correctly. It is likely that there would be a desire to salvage a partially good CMT processor (one where a subset of the virtual processors and all the common area function correctly) and use it as a CMT processor with fewer than the maximum number of functional virtual processors. In such a case, it would be possible that the functional strands be

numbered contiguously, starting from 0, and that the STRAND\_ID.max\_strand\_id field be set to the highest-numbered functional virtual processor. This requires some way to reassign the identity of individual virtual processors after manufacturing. If this is not practical, the functioning virtual processors may not be contiguously numbered.

### 15.3.1.1 Exposing Stranding

If a processor implements multiple strands per physical core, the stranding is exposed in STRAND\_ID.max\_strand\_id\_per\_core. This field encodes one less than the number of strands that are implemented on the physical processor core; for example, on a physical core with 4 strands, STRAND\_ID.max\_strand\_id\_per\_core = 3. Every virtual processor within the physical core must observe the same value of max\_strand\_id\_per\_core. An implementation defines and count strands and physical processor cores as appropriate for that implementation.

When STRAND\_ID.max\_strand\_id\_per\_core is nonzero, there are additional constraints on the numbering of virtual processors. virtual processors that correspond to strands on the same physical processor core must have contiguous STRAND\_ID.strand\_id values, with the lowest numbered virtual processor on a physical core having a strand\_id value that is a multiple of the number of strands on each physical core.

It is important to expose stranding to software. From a performance standpoint, stranding must be exposed for the operating system to understand resource sharing and contention issues and to optimally schedule software threads on the processor. From a power management perspective, knowledge of stranding enables the facility to park or disable all strands on a physical core to obtain significant power savings.

## 15.3.2 Strand Interrupt ID Register (STRAND\_INTR\_ID)

Register Name	ASI # (Name)	VA	Scope	Access	Note
STRAND_INTR_ID	63 <sub>16</sub> (ASI_CMT_PER_STRAND)	00 <sub>16</sub>	per-strand	RW	



FIGURE 15-2 STRAND\_INTR\_ID Register

The STRAND\_INTR\_ID register allows software to assign a 16-bit interrupt ID, unique within a system, to each virtual processor. This is necessary in order to enable virtual processors to receive interrupts. The identifier in this register is used

by other virtual processors (on the same and different CMT processors) and other bus agents to address interrupts to this specific virtual processor. It can also be used by this virtual processor to identify itself as the source of an interrupt it sends to other virtual processors and bus agents.

This register is Read/Write, accessible only in hyperprivileged mode (HPSTATE.hpriv = 1). It is expected that it will be modified only at boot or reconfiguration time. An attempt to access this register in privileged mode or nonprivileged mode results in a *privileged\_action* exception.

The STRAND\_INTR\_ID register has only one field, a 16-bit interrupt ID field, named *int\_id*.

If an implementation uses fewer than 16 bits for its interrupt ID, the unused bits read as zero and writes to them are ignored.

**IMPL. DEP. #:** It is implementation dependent whether any portion of the *int\_id* field of the STRAND\_INTR\_ID register is read-only (see following subsection, *Assigning an Interrupt ID*).

### 15.3.2.1 Assigning an Interrupt ID

When assigning the interrupt ID to a virtual processor, software must be aware of interrupt routing conventions used in the system. Some portion of the interrupt ID might be required to follow a hardware convention to enable the interrupt to be correctly routed through the system interconnect. In some implementations, a part of the interrupt ID can be fixed by the processor to correspond to the strand ID. This portion of the interrupt ID can be read-only in the STRAND\_INTR\_ID register. Such requirements are both processor- and system-platform-specific.

Each virtual processor in the CMT processor must have an interrupt ID that is unique within the system. If the interrupt ID of multiple virtual processors in the same system are set to the same value, the behavior of the processor is undefined when an interrupt specifying that ID is sent or received.

### 15.3.2.2 Dispatching and Receiving Interrupts

The mechanisms used to dispatch and receive interrupts must work with the interrupt ID register. A processor's interrupt dispatch mechanism must be able to specify the interrupt ID of the destination virtual processor to which the interrupt is to be delivered. When a destination interrupt ID is specified, the interrupt must be delivered to the virtual processor that has the matching ID in its STRAND\_INTR\_ID register.

### 15.3.2.3 Updating the Strand Interrupt ID Register

It is expected that the interrupt ID register of a virtual processor will be written once by software, when a virtual processor is initially booted. It is assumed that while a virtual processor is being booted, there will be no interrupt traffic in the system.

The latency from when software writes to `STRAND_INTR_ID` to when the write takes effect is implementation dependent. Use of a `MEMBAR #Sync` instruction after a write to `STRAND_INTR_ID` will cause the write to become visible before any instructions after the `MEMBAR` are executed on the virtual processor.

Updates to `STRAND_INTR_ID` are atomic: if `STRAND_INTR_ID` is written, the value observed at any time will be either the old value or the new value; no transient value will be observed. If an interrupt is issued to a virtual processor while its interrupt ID register is being updated (addressed either to its old or new interrupt ID), the interrupt may or may not be received by the virtual processor. Once a virtual processor acknowledges an interrupt using its new interrupt ID, it will not acknowledge any interrupts addressed to the old interrupt ID.

If an interrupt is issued to a system, addressed to an interrupt ID that does not match any virtual processors or other system agents, the interrupt will not be acknowledged and will be dropped.

---

## 15.4 Disabling and Parking Virtual Processors

The CMT programming model provides the ability to disable virtual processors and temporarily suspend (park) virtual processors. This section describes the interface for probing what virtual processors are available, enabled, and running (not parked). This section also describes the interface for enabling/disabling virtual processors and parking/unparking virtual processors.

## 15.4.1 Strand Available Register (STRAND\_AVAILABLE)

Register Name	ASI # (Name)	VA	Scope	Access	Note
STRAND_AVAILABLE	41 <sub>16</sub> (ASI_CMT_SHARED)	00 <sub>16</sub>	per-proc	R only	




FIGURE 15-3 STRAND\_AVAILABLE Register

The STRAND\_AVAILABLE register is a shared (one per processor) register that indicates which virtual processors are available for use (that is, are present and functional) in a CMT implementation.

The STRAND\_AVAILABLE register is read-only, comprising a single 64-bit field. As illustrated in FIGURE 15-3, bit  $n$  corresponds to virtual processor  $n$ ; therefore up to 64 virtual processors are supported per CMT. If a bit in the register is 1, the corresponding virtual processor is available for use in the CMT. If a bit in the register is 0, the corresponding virtual processor is not available for use. An “available” virtual processor is one that is present and functional, therefore can be enabled and used.

## 15.4.2 Enabling and Disabling Virtual Processors

The CMT programming model allows virtual processors to be enabled and disabled. Enabling or disabling a virtual processors is a heavyweight operation that in most cases requires either a *power\_on\_reset* (POR) or a *warm\_reset* (WRM) for updates. A disabled virtual processor produces no architectural effects observable by other virtual processors, and does not participate in cache coherency. The behavior of any transaction (such as an interrupt) issued to a disabled virtual processor is undefined.

**IMPL. DEP. #322-U4:** Whether disabling a virtual processor reduces the power used by a CMT is implementation dependent. It is recommended that a disabled virtual processor consume a minimal amount of power.

**IMPL. DEP. #423-S10:** Whether disabling a virtual processor increases the performance of other virtual processors in the CMT is implementation dependent.

## 15.4.2.1 Strand Enable Status Register (STRAND\_ENABLE\_STATUS)

Register Name	ASI # (Name)	VA	Scope	Access	Note
STRAND_ENABLE_STATUS	41 <sub>16</sub> (ASI_CMT_SHARED)	10 <sub>16</sub>	per-proc	R only	



FIGURE 15-4 STRAND\_ENABLE\_STATUS Register

The STRAND\_ENABLE\_STATUS register is a shared (one per processor) register that indicates which virtual processors are currently enabled. The register is a read-only register, in which each bit corresponds to a virtual processor.

As shown in FIGURE 15-4, bit  $n$  corresponds to virtual processor  $n$ . If a bit in the STRAND\_ENABLE\_STATUS register is 1, the corresponding virtual processor is available and enabled. A virtual processor indicated as “not available” in the STRAND\_AVAILABLE register cannot be enabled, and its corresponding enabled bit in this register will be 0. An available, enabled virtual processor that is parked is still considered enabled.

**Programming** | Hyperprivileged software should never set bit  
**Note** | STRAND\_ENABLE $\{n\}$  to 1 if STRAND\_AVAILABLE $\{n\} = 0$ .

**State After Reset.** The STRAND\_ENABLE\_STATUS register changes due to a *power\_on\_reset*. (POR) or a *warm\_reset*. (WRM). During a *power\_on\_reset*, the contents of its STRAND\_AVAILABLE register are copied to the STRAND\_ENABLE\_STATUS register. During a *warm\_reset* reset, the contents of the STRAND\_ENABLE register are copied to the STRAND\_ENABLE\_STATUS register.

## 15.4.2.2 Strand Enable Register (STRAND\_ENABLE)

Register Name	ASI # (Name)	VA	Scope	Access	Note
STRAND_ENABLE	41 <sub>16</sub> (ASI_CMT_SHARED)	20 <sub>16</sub>	per-proc	RW	Changes take effect during reset



FIGURE 15-5 STRAND\_ENABLE Register

The STRAND\_ENABLE register is a shared (one per processor) register, used by software to enable and disable a CMT's virtual processors. When disabled, a virtual processor and any structures private to that virtual processor behave as though they were not present.

<b>Programming Note</b>	When re-enabled, per-strand architectural state that existed when the virtual processor was previously enabled should be assumed to be lost. Therefore, hyperprivileged software must initialize any needed per-strand architectural state each time a virtual processor is enabled.
-------------------------	--

Changing a bit in the STRAND\_ENABLE register does *not* take effect (cause a virtual processor to be enabled/disabled) immediately. Instead, it indicates a *pending* change to the STRAND\_ENABLE\_STATUS register, which will not take effect until the next *warm\_reset*. (WRM) reset — at which time, the contents of the STRAND\_ENABLE register are copied to the STRAND\_ENABLE\_STATUS register. A change in the STRAND\_ENABLE register may also take place at some other implementation-dependent time (see *Dynamically Enabling/Disabling Virtual Processors* on page 526 (impl. dep. #\_\_\_)).

As shown in FIGURE 15-5, the STRAND\_ENABLE register contains one bit per possible virtual processor, with bit *n* corresponding to virtual processor *n*. If bit *n* is 1, then virtual processor *n* should be enabled after the next warm reset (if that virtual processor is available). If bit *n* is 0, then virtual processor *n* should be disabled after the next warm reset.

When bit *n* in the STRAND\_AVAILABLE register is 0 (the virtual processor is unavailable), the corresponding bit (bit *n*) in the STRAND\_ENABLE register is forced to 0 and attempts to write “1” to bit *n* in the STRAND\_ENABLE register are ignored.

### Restrictions on Updating the STRAND\_ENABLE Register.

**IMPL. DEP. #323-U4:** Whether an implementation provides a restriction that prevents software from writing a value of all zeroes (or zeroes corresponding to all available virtual processors) to the STRAND\_ENABLE register is implementation dependent. This restriction avoids the dangerous case where all virtual processors become disabled and the only way to enable any virtual processor is a hard *power\_on\_reset* (a warm reset would not suffice). If such a restriction is implemented and software running on any virtual processor attempts to write a value of all zeroes (or zeroes corresponding to all available virtual processors) to the STRAND\_ENABLE register, hardware forces the STRAND\_ENABLE register to an implementation-dependent value which enables at least one of the available virtual processors.

**State After Reset.** Upon assertion of *power\_on\_reset*, the value of the STRAND\_AVAILABLE register is copied to the STRAND\_ENABLE register. The STRAND\_ENABLE register does not change during any other reset, including system (or equivalent) resets.

### 15.4.2.3 Dynamically Enabling/Disabling Virtual Processors

**IMPL. DEP. #424-S10:** Whether a CMT implementation provides the ability to dynamically enable and disable virtual processors is implementation dependent. It is tightly coupled to the underlying microarchitecture of a specific CMT implementation. This feature is implementation dependent because any implementation-independent interface would be too inefficient on some implementations.

## 15.4.3 Parking and Unparking Virtual Processors

Parking is a way to temporarily suspend the operation of a virtual processor, intended for use by critical diagnostic and recovery code. A parked virtual processor can be later unparked to allow it to resume running. A virtual processor can be parked or unparked at arbitrary times using the STRAND\_RUNNING register and a WMR or POR reset is not required for parking/unparking to become effective. The STRAND\_RUNNING\_STATUS register can be used to determine whether a virtual processor that has been directed to park has completed the process of parking.

A parked virtual processor does not execute instructions and does not initiate any transactions on its own. If any portion of the memory system resides in a parked virtual processor, it will continue to be updated as necessary for it to remain coherent with the rest of the memory system while the virtual processor is parked.

When a virtual processor is unparked, it continues execution with the instruction that was next to be executed when the virtual processor was parked. It is transparent to software running on a virtual processor that it was ever parked (except for observable timing considerations).

While a virtual processor is parked, the STICK register continues to count.

**IMPL. DEP. #425-S10:** It is implementation dependent whether the TICK register continues to count while a virtual processor is parked.

Using the TICK or STICK counter to detect the parking of a virtual processor is not recommended.

An interrupt to a parked virtual processor behaves the same as if the virtual processor was too busy to accept the interrupt.

**IMPL. DEP. #324-U4:** It is implementation dependent whether parking a virtual processor reduces the power used by a CMT. It is recommended that a parked virtual processor use a reduced amount of power.

Parking a virtual processor should, when appropriate, reduce the contention for shared resources and enable other virtual processors to potentially run faster.

**IMPL. DEP. #426-S10:** The degree to which parking a virtual processor impacts the performance of other virtual processors is implementation dependent.

<b>Implementation Note</b>	One possible way to implement virtual processor parking is to disable instruction fetching in a parked virtual processor. In such an implementation, after a virtual processor is parked, it will execute the instructions currently in its pipeline, complete pending transactions (such as draining the store queue), and then become idle.
----------------------------	---

### 15.4.3.1 Strand Running Register (STRAND\_RUNNING)

Register Name	ASI # (Name)	VA	Scope	Access	Note
STRAND_RUNNING_RW	41 <sub>16</sub> (ASI_CMT_SHARED)	50 <sub>16</sub>	per-proc	RW	General RW access
STRAND_RUNNING_W1S	41 <sub>16</sub> (ASI_CMT_SHARED)	60 <sub>16</sub>	per-proc	W1S	Write 1s to set bits
STRAND_RUNNING_W1C	41 <sub>16</sub> (ASI_CMT_SHARED)	68 <sub>16</sub>	per-proc	W1C	Write 1s to clear bits



**FIGURE 15-6** STRAND\_RUNNING Register

STRAND\_RUNNING is a shared (one per processor) register, used by software to park and unpark selected virtual processors in a CMT implementation. When a virtual processor is parked, the virtual processor stops executing new instructions and will not initiate new transactions except in response to a coherency transaction initiated by another virtual processor.

**IMPL. DEP. #427-S10:** There may be an arbitrarily long, but bounded, delay (“skid”) from the time when a virtual processor is directed to park or unpark (via an update to the STRAND\_RUNNING register) until the corresponding virtual processor(s) actually park or unpark.

Multiple access methods are provided for writing bits in the STRAND\_RUNNING register, distinguished by the virtual address used (listed above):

- STRAND\_RUNNING\_RW, for normal reading and writing of the entire register
- STRAND\_RUNNING\_W1S (“Write 1 to Set”), where writing ‘1’ to a bit sets the destination bit to ‘1’ and writing ‘0’ to a bit leaves the destination bit unchanged

- **STRAND\_RUNNING\_W1C** (“Write 1 to Clear”), where writing ‘1’ to a bit sets the destination bit to ‘0’ (clears it) and writing ‘0’ to a bit leaves the destination bit unchanged

A specific value can be atomically written to all bits of the **STRAND\_RUNNING** register, using **STRAND\_RUNNING\_RW**, or bits can be individually modified, using **STRAND\_RUNNING\_W1S** or **STRAND\_RUNNING\_W1C**. When a virtual processor parks itself, software should write to **STRAND\_RUNNING\_W1C**. When a virtual processor wants to become the only active virtual processor (parking all other virtual processors in the CMT), it is more appropriate to write the desired value directly to **STRAND\_RUNNING\_RW**. A direct write eliminates the need to perform separate set and clear operations to write a specific value to the register.

As shown in **FIGURE 15-6**, the **STRAND\_RUNNING** register contains one bit per possible virtual processor, with bit  $n$  corresponding to virtual processor  $n$ . Writing a value of 1 to bit position  $n$  activates (unparks) virtual processor  $n$  for normal execution, while writing a value of 0 to bit  $n$  parks virtual processor  $n$ . If bit  $n$  in the **STRAND\_ENABLE\_STATUS** register is 0 (not enabled), hardware forces the corresponding bit in the **STRAND\_RUNNING** register to 0 and attempts to write to that bit are ignored.

**Updating the STRAND\_RUNNING Register.** When a virtual processor parks itself by updating the **STRAND\_RUNNING** register and follows the update with a **FLUSH** instruction, no instruction after the **FLUSH** instruction will be executed until the virtual processor is unparked. The virtual address specified in the **FLUSH** instruction is not important. The **FLUSH** instruction may be executed either before parking takes effect or after the virtual processor is unparked. The **FLUSH** can, therefore, enable software to bound when parking takes effect, in the case when a virtual processor parks itself.

**IMPL. DEP. #428-S10:** When a virtual processor writes to the **STRAND\_RUNNING** register to park itself, the method by which completion of parking is assured (instructions stop being issued) is implementation dependent.

**Simultaneous Updates to the STRAND\_RUNNING Register.** Hardware is not required to provide a mechanism for handling simultaneous updates from different strands to the **STRAND\_RUNNING** register.

<b>Programming Note</b>	<p>It is the responsibility of hyperprivileged software to insure that a livelock condition, resulting from simultaneous updates from different strands to the STRAND_RUNNING register, does not occur.</p> <p>After writing to STRAND_RUNNING with a STXA instruction, hyperprivileged software should check the STRAND_RUNNING_STATUS register to verify when the attempted parking/unparking of virtual processor(s) actually completed.</p>
-------------------------	---

**At Least One Virtual Processor Must Remain Unparked.** Hardware enforces the restriction that an update to the STRAND\_RUNNING register by software running on one of the virtual processors cannot cause all of the enabled virtual processors to become parked. This restriction is important to avoid the dangerous situation where all virtual processors become parked and there is no way to reactivate any of the virtual processors (without a warm reset or power-on reset).

**IMPL. DEP. #429-S10:** If an update to the STRAND\_RUNNING register would cause all enabled virtual processors to become parked, it is implementation dependent which virtual processor is automatically unparked by hardware. The preferred implementation is that when an update to the STRAND\_RUNNING register (STXA instruction) would cause all virtual processors to become parked, hardware silently ignores (discards) that STXA instruction.

<b>Implementation Note</b>	<p>It is important that when a virtual processor attempts to issue an update to the STRAND_RUNNING register that would cause all virtual processors to become parked, that virtual processor is <i>not</i> parked. A virtual processor updating the STRAND_RUNNING register will be executing a section of software (error diagnostic or other special code) that is aware of the behavior and implications of parking. When an attempt is made to park all virtual processors, automatically unparking an <i>arbitrary</i> virtual processor would be problematic, because a virtual processor in the midst of running nonprivileged code could become the only unparked virtual processor. If this were to happen, the only active virtual processor in the CMT would be unaware of the state of the CMT and would not know to check the running status of other virtual processors.</p>
----------------------------	--

**At Least One Virtual Processor Must Remain Unparked — Multiprocessor Configuration.** When there are multiple processors (chips) in the configuration, there is still a requirement to have at least one virtual processor unparked on each processor. However, from a testing point of view, it is desirable to be able to unpark all but one virtual processor in the entire multiprocessor configuration.

**IMPL. DEP. #430-S10:** In a multiprocessor configuration, whether all but one virtual processor can be parked is implementation dependent.

**State After Reset.** Upon power-on reset or warm reset, the STRAND\_RUNNING register by default is initialized such that all the virtual processors are parked except for the lowest-numbered enabled virtual processor. This provides a default on-chip “boot master” virtual processor, reducing BootBus contention.

**Note** | For systems that use a system reset pin, the value of the STRAND\_RUNNING register is updated upon assertion of the warm reset signal.

### 15.4.3.2 Strand Running Status Register (STRAND\_RUNNING\_STATUS)

Register Name	ASI # (Name)	VA	Scope	Access	Note
STRAND_RUNNING_STATUS	41 <sub>16</sub> (ASI_CMT_SHARED)	58 <sub>16</sub>	per-proc	R only	



**FIGURE 15-7** STRAND\_RUNNING\_STATUS Register

STRAND\_RUNNING\_STATUS is a shared (one per processor) register. It indicates whether a virtual processor is still active (running) or has actually become parked. It is needed because there may be a delay between the time when a virtual processor is directed to park (via the STRAND\_RUNNING register) and the time when it actually becomes parked. The STRAND\_RUNNING\_STATUS register is a shared, read-only register in which bit  $n$  indicates if strand  $n$  is active.

There is an implementation-dependent delay from the time virtual processor  $n$  is directed to park by writing 0 to bit  $n$  of the STRAND\_RUNNING register until it actually becomes parked (impl. dep. #427-S10).

As shown in FIGURE 15-7, the STRAND\_RUNNING\_STATUS register has one 64-bit field (one bit per possible virtual processor), with bit  $n$  corresponding to virtual processor  $n$ .

- If virtual processor  $n$  is enabled (STRAND\_ENABLE\_STATUS $\{n\}$  = 1):
  - a value of 0 in bit  $n$  of the STRAND\_RUNNING\_STATUS register indicates that virtual processor  $n$  is truly parked and will not execute any additional instructions or initiate new transactions until it is unparked.

- A value of 1 in bit  $n$  of the STRAND\_RUNNING\_STATUS register indicates that a virtual processor is active and can execute instructions and initiate transactions. All virtual processors that have a 1 in the STRAND\_RUNNING register must have a 1 in the STRAND\_RUNNING\_STATUS register.
- If virtual processor  $n$  is disabled (STRAND\_ENABLE\_STATUS $\{n\}$  = 0), bit  $n$  of the STRAND\_RUNNING\_STATUS register must be 0.

The STRAND\_RUNNING\_STATUS register indicates when a virtual processor that has been directed to park has actually parked, that is, is no longer executing instructions or initiating any transactions (except in response to coherency transactions generated by other virtual processors).

**IMPL. DEP. #431-S10:** The criteria used for determining whether a virtual processor is fully parked (corresponding bit set to '1' in the STRAND\_RUNNING\_STATUS register) are implementation dependent.

After bit  $n$  in the STRAND\_RUNNING register has been changed from 1 to 0, hardware must guarantee that only a single transition from 1 to 0 in bit  $n$  of the STRAND\_RUNNING\_STATUS register will be observed.

**State After Reset.** The value of the STRAND\_RUNNING\_STATUS register is the same as the value of the STRAND\_RUNNING register at the end of a system reset.

## 15.4.4 Virtual Processor Standby (or Wait) State

**IMPL. DEP. #432-S10:** Whether an implementation implements a Standby (or Wait) state for virtual processors, how that state is controlled, and how that state is observed are implementation-dependent.

In a Standby state, the virtual processor is suspended for a predetermined period of time and/or until an external interrupt is received. A Standby state may appear similar to a Parked state, but virtual processor Standby state (if implemented) must be completely orthogonal to parking. The details of the software interface to and implementation of Standby/Wait state is beyond the scope of this specification.

With respect to parking, the virtual processor is either Running or not running (Parked), as indicated in the STRAND\_RUNNING\_STATUS register. With respect to standby, the virtual processor is either in Standby or Normal state. Since these features are independent, the virtual processor can be in any of the four possible combinations of these states. A virtual processor is still considered running if it is in a Standby mode but is not Parked. If a virtual processor is in a Standby mode and becomes Parked, it will remain Parked even if an event causes it to change from Standby to Normal mode; it will not execute instructions until it is later unparked.

Implementing a Standby mode may provide performance and/or power-consumption benefits. A virtual processor in Standby mode may cause less resource contention with other running virtual processors and may consume less power.

---

## 15.5 Reset and Trap Handling

In a CMT, some resets apply globally to all virtual processors, some apply to an individual virtual processor, and some apply to an arbitrary subset of virtual processors. The following sections address how each type of reset affects the virtual processors in a CMT.

The reset nomenclature used in this section is generally consistent with that used for UltraSPARC Architecture 2005 processors. If future processors classify resets differently, this model should be extended appropriately to the new classifications.

Traps (as opposed to resets) apply to individual virtual processors and are discussed in *Traps* on page 437.

### 15.5.1 Per-Strand Resets (SIR and WDR Resets)

The only resets that affect only a single virtual processor are those that are internally generated by a virtual processor, such as software initiated reset (SIR) and watchdog reset (WDR). These resets are generated by an individual virtual processor and are not propagated to the other virtual processors in a CMT.

### 15.5.2 Full-Processor Resets (POR and WRM Resets)

There is a class of resets that are generated by an external agent and apply to all the virtual processors within a processor. This class includes all resets associated with fundamental CMT reconfigurations.

*power\_on\_reset* (POR) is one case of full-processor reset. Warm reset is another example of such a reset (warm reset may be either processor or physical strand-specific, depending on the implementation). Full-processor reset is required for certain reconfigurations of the processor.

Power-on reset and warm reset (or their equivalents in future processors) are global resets, sent to all strands in a CMT processor.

## 15.5.2.1 Boot Sequence

As discussed in *Strand Running Register (STRAND\_RUNNING)* on page 527, the default boot sequence is for all virtual processors except one (nominally, the lowest-numbered enabled virtual processor) to be set to `Parked` state at the beginning of full-processor reset. The single unparked virtual processor is the master virtual processor, which should arbitrate for the BootBus (if multiple CMT processors share the same BootBus). The master virtual processor (or service processor) should unpark the other virtual processors in the processor at the appropriate time in the booting process.

## 15.5.3 Partial Processor Resets (XIR Reset)

There is a class of resets, referred to here as “partial-processor resets,” that are generated by an external agent and affect an arbitrary subset of virtual processors within a processor. The subset may be anything from all virtual processors to no virtual processors (impl. dep. #433-S10).

Externally-initiated reset (XIR) is a partial-processor reset. XIR is intended to reset a specific virtual processor in a system, primarily for diagnostic and recovery purposes.

**IMPL. DEP. #433-S10:** A mechanism must exist to specify which subset of virtual processors in a processor should be reset when a partial-processor reset (for example, XIR) occurs. The specific mechanism is implementation-dependent.

Possible methods of specifying the subset include the following:

1. Before the partial-processor reset occurs, set up a steering register that specifies the subset of virtual processors that should be affected. For systems using an XIR reset, the XIR Steering register described in *XIR Steering Register (XIR\_STEERING)* on page 534 should be used.
2. Specify the subset of virtual processors concurrently with the reset request, across the same interface used for communicating the reset. This method would require that the interface used for communicating resets supports sending packets of information along with the resets.

In an implementation that replaces the XIR reset with a different set of resets, the following rules apply for extending this CMT programming interface:

- Each partial-processor reset may use an interface where the set of virtual processors to reset is communicated along with the reset request.
- For partial-processor resets for which the set of virtual processors to be reset is *not* communicated along with the reset request:
  - The highest priority virtual processor will use the `XIR_STEERING` register to determine the subset of virtual processors to be reset.

- Each subsequent lower-priority virtual processor can either use the XIR\_STEERING register or use an additional steering register (comparable to XIR\_STEERING), specifically associated with that reset. Each additional steering register will be accessed using the same ASI number (41<sub>16</sub>) as the XIR\_STEERING register but with a distinct virtual address.

### 15.5.3.1 XIR Steering Register (XIR\_STEERING)

Register Name	ASI # (Name)	VA	Scope	Access	Note
XIR_STEERING	41 <sub>16</sub> (ASI_CMT_SHARED)	30 <sub>16</sub>	per-proc	RW	General access




FIGURE 15-8 XIR\_STEERING Register

An externally initiated reset (XIR) can be steered to an arbitrary subset of virtual processors, using the XIR\_STEERING register. The XIR\_STEERING register is shared across virtual processors and is used by software to control which virtual processor(s) within a processor will receive the XIR reset signal when XIR is asserted for the processor module.

As shown in FIGURE 15-8, the XIR\_STEERING register has one 64-bit field (one bit per possible virtual processor), in which bit  $n$  corresponds to virtual processor  $n$ .

When an external reset is asserted for the CMT, if bit  $n$  in the XIR\_STEERING register is 1, virtual processor  $n$  receives an XIR reset; if bit  $n$  in the XIR\_STEERING register is 0, virtual processor  $n$  continues execution, unaware of the external reset asserted for the CMT.

A virtual processor that is parked when it receives an XIR reset remains parked and will handle the XIR reset immediately after being unparked.

**IMPL. DEP. #325-U4a:** Whether XIR\_STEERING $\{n\}$  is a read-only bit or a read/write bit is implementation dependent. If XIR\_STEERING $\{n\}$  is read-only, then (1) writes to XIR\_STEERING $\{n\}$  are ignored and (2) XIR\_STEERING $\{n\}$  is set to 1 if virtual processor  $n$  is available and to 0 if it is not available (that is, XIR\_STEERING $\{n\}$  reads the same as STRAND\_AVAILABLE $\{n\}$ ).

It may be desirable for an XIR to effectively unpark and reset all virtual processors in a CMT. If so, that effect can be generated by having the first action of software on virtual processor receiving an XIR to unpark all other virtual processors in the CMT.

### State After Reset.

During *power\_on\_reset*, the contents of the STRAND\_AVAILABLE register are copied to the XIR\_STEERING register. During a warm reset, the contents of the STRAND\_ENABLE register are copied to the XIR\_STEERING register. This provides for a default condition in which all enabled virtual processors receive an XIR reset when an external reset is asserted for the processor. (impl. dep. #325-U4b)

---

## 15.6 Error Handling in CMT Processors

Errors in a structure *private* to a virtual processor are considered virtual-processor(strand)-specific and are reported to that virtual processor using its error-reporting mechanism.

When an error in a structure *shared* among virtual processors occurs:

- If the virtual processor initiating the request that caused or detected the error can be identified, the error is considered virtual-processor-specific and is reported back to the originating virtual processor.
- If the virtual processor initiating the request that caused or detected the error cannot be identified, the error is considered non-virtual-processor-specific.
- All virtual processors that share a structure are considered to be part of the *error-handling group* for that structure. This implies that any virtual processor in the group can be assigned to handle error traps associated with the structure and have diagnostic access to the structure for error recovery.

The following sections describe how a CMT processor handles both virtual-processor-specific and non-virtual-processor-specific errors.

### 15.6.1 Virtual-Processor-Specific Error Reporting

Errors specific to a particular virtual processor are reported to the virtual processor associated with the error, using the virtual processor's error reporting mechanism. A virtual-processor-specific error can be either synchronous or asynchronous. It may be an error that occurred in a shared structure but is traceable to the originating virtual processor. It is the responsibility of error handling software to recognize the implication of errors in shared structures and take appropriate action.

## 15.6.2 Reporting Errors on Shared Structures

Errors in shared structures are more complicated than virtual-processor-specific errors. When a non-virtual-processor-specific error occurs, it must be recorded and an exception must be generated on one of the virtual processors within the CMP to deal with the error. More precisely, the virtual processor that reports the exception must be part of the error-handling group for the shared structure in which the error was detected. The following subsections describe where the error should be recorded and in which virtual processor the exception should be generated.

### 15.6.2.1 Error Steering

When an error occurs in a shared resource, the error must be reported to a virtual processor that shares that resource and is part of its error-handling group. That virtual processor has the capability of issuing diagnostic reads and writes to the structure for diagnosis, correction, and error-clearing purposes. Error steering registers are used to determine which virtual processor will handle the error. Software configures an error steering register to specify which virtual processor should handle the error(s) associated with that error steering register. That is, an error steering register defines in which virtual processor an exception will be generated, to report and handle the error.

A given CMT implementation may contain resources shared by all the virtual processors of the CMT processor or shared by a subset of two or more virtual processors.

**IMPL. DEP. #434-S10:** Because of the range of implementation, the number of, organization of, and ASI assignments for error steering registers in a CMT processor are implementation dependent.

Error steering registers may be provided per shared resource or per level of sharing. In the case that all shared resources are shared by all virtual processors, it is recommended that a single error steering register be used and that error steering register should follow the behavior of the `ERROR_STEERING` register defined in *Error Steering Register (ERROR\_STEERING)* on page 538. If a mechanism is used where error steering registers are used per level of sharing, it is recommended that the `ERROR_STEERING` register be used for the level at which all virtual processors share and provide error-handling groups.

**General Guidelines for Error Steering Registers.** An error steering register controls which virtual processor handles non-virtual-processor-specific errors. Such an error is recorded using the virtual processor's asynchronous error reporting mechanism (as relevant to the error) and generates an appropriate exception.

An error steering register is accessed through an ASI or a memory-mapped address. It must be accessible for both reading and writing by software (using load and store alternate instructions).

A processor contains one or more error steering registers. The number of error steering registers needed depends on how resources are shared and the ability of a virtual processor to diagnose errors in a resource it does not share.

An error steering register specifies a virtual processor by an encoded field, `target_id`, that corresponds to the `strand_id` of the targeted virtual processor. Use of an encoded representation guarantees that only one virtual processor can be specified. An error steering register should contain only one field, the `target_id` field, that encodes the `strand_id` of the virtual processor that should be informed of non-virtual-processor-specific errors in its sharing group.

**IMPL. DEP. #326-U4-Cs10a:** The number of implemented bits of `ERROR_STEERING.target_id` is nominally six, but is implementation dependent and must be sufficient to encode the highest implemented virtual processor ID.

It is the responsibility of software to ensure that an error steering register identifies an appropriate virtual processor for handling the error(s) assigned to it. If an error steering register identifies a virtual processor that is not available (per `STRAND_AVAILABLE`) or is disabled (per `STRAND_ENABLE_STATUS`), none of the enabled virtual processors in the error-handling group will be affected by the reporting of a non-virtual-processor-specific error to the disabled virtual processor. However, the behavior of the specified disabled virtual processor is undefined; for example, the error status register in the disabled virtual processor may or may not be observed to have been updated.

If an error steering register identifies a virtual processor that is not part of the error-handling group, operation is also undefined. An example would be if the error steering register identifies a virtual processor in another error-handling group for a virtual-processor-specific error. To avoid this case, an error steering register should be assigned on a core basis for core errors that are non-virtual-processor-specific.

If an error steering register identifies a virtual processor that is parked, the non-virtual-processor-specific error is reported to that virtual processor and the virtual processor will observe the appropriate exception, but not until after it is unparked.

When an error steering register is written by software, the update becomes visible after an unspecified delay. If a store to the register is followed by a `MEMBAR` synchronization barrier instruction, it is guaranteed that the write to the error steering register will complete by the time the execution of the `MEMBAR` instruction completes.

When a non-virtual-processor-specific error occurs, the corresponding error steering register is consulted. The error is reported to and an exception is generated in the virtual processor indicated by the error steering register.

If a non-virtual-processor-specific error occurs and at the same time `target_id` is being changed in the corresponding error steering register, the subsequent error report and the generated exception will occur together on the *same* virtual processor, either the virtual processor indicated by the old value in the error steering register or

the one indicated by the new value. That is, for non-virtual-processor-specific errors, the generation of an error report plus an exception is atomic with respect to changes to the contents of the error steering register.

### State of Error Steering Register After Reset.

The `target_id` field of an error steering register is initialized during a power-on-reset and warm reset. After a power-on-reset, the value in the `target_id` field of an error steering register should refer to the lowest-numbered available virtual processor (as indicated by the `STRAND_AVAILABLE` register) that corresponds to the resource(s) covered by the steering register. After a warm reset, the value in the `target_id` field of an error steering register should refer to the lowest-numbered enabled virtual processor (as indicated by the `STRAND_ENABLE` register) that corresponds to the resource(s) covered by the steering register.

**Error Steering Register (ERROR\_STEERING).** The `ERROR_STEERING` register is the recommended mechanism for specifying which virtual processor in an error-handling group should handle non-virtual-processor-specific errors in resources shared by all virtual processors of the error-handling group. `ERROR_STEERING` is a shared register, accessible from all virtual processors in the error-handling group.

When a non-virtual-processor-specific error occurs, the error is recorded using the asynchronous error reporting mechanism in the virtual processor indicated by `ERROR_STEERING`. The appropriate exception is generated in that same virtual processor.

Register Name	ASI # (Name)	VA	Scope	Access	Note
<code>ERROR_STEERING</code>			per-proc	RW	

The Error Steering register has only one field that encodes the strand ID of the strand that should be informed of non-virtual-processor-specific errors. When an error is detected that cannot be traced back to a specific virtual processor, the error is recorded in, and a trap is sent to, the virtual processor identified by the Error Steering register.




FIGURE 15-9 `ERROR_STEERING` Register

**IMPL. DEP. #435-S10:** Although the `ERROR_STEERING` register is the recommended mechanism for steering non-virtual-processor-specific errors to a virtual processor for handling, the actual mechanism used in a given implementation is implementation dependent.

The `ERROR_STEERING` register contains one field, `target_id`, that encodes the virtual processor ID of the virtual processor that should be informed of non-virtual-processor-specific errors (see FIGURE 15-9).

**IMPL. DEP. #436-S10:** The width of the `target_id` field of the `ERROR_STEERING` register is implementation dependent.

The `target_id` field (refer to FIGURE 15-9) must be wide enough to encode the strand ID of the highest-numbered implemented virtual processor. If  $n$  bits of this field are implemented, the unused most-significant bits numbered 5 to  $6-n$  read as zero and writes to those bits are ignored.

**IMPL. DEP. #437-S10:** An implementation may provide multiple `target_id` fields in an `ERROR_STEERING` register for different types of non-virtual-processor-specific errors.

### 15.6.2.2 Reporting Non-Virtual-Processor-Specific Errors

Before an exception can be generated for a non-virtual-processor-specific error, the error must be recorded. Non-virtual-processor-specific errors are recorded using the asynchronous error reporting mechanism of the virtual processor specified by the `ERROR_STEERING` register. The mechanism used is the same as that for reporting virtual processor-specific errors.

Each asynchronous error is defined as either virtual-processor-specific or non-virtual-processor-specific. If the same error can occur as either a virtual-processor-specific error or a non-virtual-processor-specific error, the two cases must be reported as two identifiably distinct errors.

**IMPL. DEP. #438-S10:** It is implementation dependent whether the error-reporting structures for errors in shared resources appear within a virtual processor in per-virtual-processor registers or are contained within shared registers associated with the shared structures in which the errors may occur.

**IMPL. DEP. #439-S10:** The type of exception generated in a virtual processor to handle each type of non-virtual-processor-specific error is implementation dependent. A virtual processor can choose to use the same exceptions used for corresponding virtual-processor-specific asynchronous errors or it can choose to generate different exceptions.

---

## 15.7 Additional CMT Software Interfaces

### 15.7.1 Diagnostic/RAS Registers

The CMT software interface defines how virtual processors are disabled or parked (for diagnostic and error recovery) and how errors are reported in a CMT processor. It is up to the implementation to provide appropriate diagnostic and recovery mechanisms, which are not specified here.

A future extension of the CMT Programming Model may include more common features for diagnostics and RAS. Increasing commonality without significantly limiting the implementation options is best.

### 15.7.2 Configuration Registers

Given the broad range of possible implementations, no common configuration interface is defined here.

At this time the CMT Programming Model does not specify any common configuration registers. A future extension of the CMT Programming Model may include some. Increasing commonality without significantly limiting the implementation options is best.

### 15.7.3 Performance Registers

At this time, no common performance registers are specified. A future extension of the CMT Programming Model may include some.

This is a specifically important area to have common features. A range of software tools rely on the performance registers and common features will enable software tools to be more quickly deployed on new architectures with less work.

### 15.7.4 Booting Support

Some of the registers previously described can be used by firmware for booting support. See *Strand Running Register (STRAND\_RUNNING)* on page 527 for an example of such a register.

During a power-on-reset, only one enabled virtual processor per processor will be unparked. Only this virtual processor will begin fetching instructions after the reset.

**IMPL. DEP. #440-S10:** Which virtual processor is unparked during POR and whether it is unparked by processor hardware or by a service processor is implementation dependent. Conventionally, the virtual processor with the lowest-numbered `strand_id` is unparked.

In a recommended booting sequence, software determines when virtual processors become unparked after reset. The default behavior is for only one virtual processor to be unparked when the system reset signal is removed. That virtual processor, in turn, configures common registers and then unparks other virtual processors one at a time. This is only one possible boot sequence; software is free to implement other boot sequences.

---

## 15.8 Performance Issues for CMT Processors

Which resources are shared among which virtual processors in a CMT processor is implementation-dependent. Resources such as caches, TLBs, and even execution pipelines may be shared by virtual processors. From a performance perspective, there are significant issues that result from this sharing. In this section, hyperprivileged software issues of thread scheduling and configuration of inactive virtual processors is discussed. Issues of how to develop algorithms and approaches to take advantage of the low communication latencies between virtual processors are not covered here.

To understand and take advantage of performance issues in a CMT processor requires some knowledge of the underlying implementation. The existence of implementation dependencies is unavoidable, but hopefully abstract representations and general approaches can reduce the degree of implementation dependence in hyperprivileged software.

---

## 15.9 Recommended Subset for Single-Strand Processors

It is recommended that single-strand UltraSPARC Architecture 2005 processors implement a subset of the CMT interface. This enables them to more easily integrate into systems that may also contain CMT processors and enables more consistent software to be deployed across those and other future systems.

Single-strand UltraSPARC Architecture 2005 processors should implement all of the CMT registers described in this chapter, as follows:

- The Strand Interrupt ID register (STRAND\_INTR\_ID) should be fully implemented.
- All other registers can be implemented as read-only registers containing fixed values, writes to which are ignored.

TABLE 15-1 summarizes the recommended implementation of CMT registers for a single-strand processor implementation:

**TABLE 15-1** Recommended CMT Register Set for Single-Strand Processors

<b>ASI</b>	<b>VA</b>	<b>Register Name</b>	<b>Type</b>	<b>Note</b>
41 <sub>16</sub>	00 <sub>16</sub>	STRAND_AVAILABLE	R only	Read value of 01 <sub>16</sub>
	10 <sub>16</sub>	STRAND_ENABLE_STATUS	R only	Read value of 01 <sub>16</sub>
	20 <sub>16</sub>	STRAND_ENABLE	R only	Read value of 01 <sub>16</sub>
	30 <sub>16</sub>	XIR_STEERING	R only	Read value of 01 <sub>16</sub>
	50 <sub>16</sub>	STRAND_RUNNING_RW	R only	Read value of 01 <sub>16</sub>
	58 <sub>16</sub>	STRAND_RUNNING_STATUS	R only	Read value of 01 <sub>16</sub>
	60 <sub>16</sub>	STRAND_RUNNING_W1S	W only (ignored)	Access (write) ignored
	68 <sub>16</sub>	STRAND_RUNNING_W1C	W only (ignored)	Access (write) ignored
63 <sub>16</sub>	00 <sub>16</sub>	STRAND_INTR_ID	RW	Software assigned unique interrupt ID for virtual processor (read/write)
	10 <sub>16</sub>	STRAND_ID	R only	Read value of 00 <sub>16</sub>

## 15.10 Machine State Summary

TABLE 15-2 describes the ASI extensions that support CMT registers. The states of CMT registers after resets are enumerated in TABLE 16-2 on page 552.

**TABLE 15-2** ASI Extensions

ASI	VA	Register Name	Scope	Type	Description
41 <sub>16</sub>	00 <sub>16</sub>	STRAND_AVAILABLE	per-proc	R	Bit mask of implemented virtual processors
	10 <sub>16</sub>	STRAND_ENABLE_STATUS	per-proc	R	Bit mask of enabled virtual processors
	20 <sub>16</sub>	STRAND_ENABLE	per-proc	RW	Bit mask of virtual processors to enable after next reset (read/write)
	30 <sub>16</sub>	XIR_STEERING	per-proc	RW	Bit mask of virtual processors to propagate XIR to (read/write)
	50 <sub>16</sub>	STRAND_RUNNING_RW	per-proc	RW	Bit mask to control which virtual processors are active and which are parked (read/write): 1 = active, 0 = parked
	58 <sub>16</sub>	STRAND_RUNNING_STATUS	per-proc	R	Bit mask of virtual processors that are currently active: 1 = active, 0 = parked
	60 <sub>16</sub>	STRAND_RUNNING_W1S	per-proc	W1S	Pseudo-register for write-one-to-set access to STRAND_RUNNING
	68 <sub>16</sub>	STRAND_RUNNING_W1C	per-proc	W1C	Pseudo-register for write-one-to-clear access to STRAND_RUNNING
63 <sub>16</sub>	00 <sub>16</sub>	STRAND_INTR_ID	per-strand	RW	Software assigned unique interrupt ID for virtual processor (read/write)
	10 <sub>16</sub>	STRAND_ID	per-strand	R	Hardware assigned ID for virtual processor (read-only)
	40 <sub>16</sub> and greater	<i>Reserved</i>	per-strand	Impl. Dep.	Reserved for implementation-specific per-strand registers



# Resets

---

## 16.1 Resets

The UltraSPARC Architecture 2005 defines 5 types of resets. Reset priorities, listed in order from highest to lowest, are as follows:

- power-on reset (POR)
- warm reset (WMR)
- externally initiated reset (XIR)
- watchdog reset (WDR), and
- software-initiated reset (SIR)

POR, WMR, and XIR resets are initiated external to the processor (chip). WDR and SIR resets are initiated by the virtual processor itself, in response to specific conditions.

POR resets are processor-wide (affect all virtual processors on the chip). WDR and SIR resets are directed to a specific virtual processor. XIR resets are directed to the virtual processor(s) indicated by the XIR\_STEERING register. WMR resets are implementation dependent and may be either processor-wide or directed to specific virtual processor(s).

Resets are used to initialize a virtual processor and place it in an operating state, to attempt recovery of a failing or stuck virtual processor, to attempt recovery of failing operating system privileged software, and for debug purposes. The defined states for each reset show an increasing amount of resource reset, such that, for example, a XIR, WDR or SIR reset will leave most architectural and memory resources unchanged, while a WMR reset will leave most memory resources unchanged but reset certain architectural resources, and a POR reset will initialize all processor resources.

All resets are processed as traps and place the virtual processor in *RED\_state*. *RED\_state* (Reset, Error, and Debug state) is a restricted execution state reserved for processing hardware- and software-initiated resets. Please refer to *Reset Traps* on page 449 and the subsections regarding reset traps in *RED\_state Trap Processing*, which begins on page 474.

## 16.1.1 Power-on Reset (POR)

A POR reset occurs when the assigned POR pin is asserted and deasserted. During this time, all other resets and traps are ignored. POR reset has the highest trap priority. POR causes any pending external transactions to be cancelled.

The POR reset is a processor-wide reset. It affects all virtual processors on the chip, as well as all IO, cache, and DRAM subsystems.

During a POR reset, hardware sets registers to a known state (see *Machine States* on page 548). All hardware-based initialization functions are performed, all logic (including the pipeline) is initialized, all architectural registers are placed in their reset state (as defined in TABLE 16-1 on page 549), and all entries in caches and TLBs are invalidated.

A service processor may also participate in the POR reset process. The POR reset functions provided by a service processor are documented in the relevant Service Processor specification.

After a POR reset is complete:

- The first available<sup>1</sup> virtual processor begins executing at physical address  $RSTVADDR + 20_{16}$  (*RED\_state* trap vector base address plus the POR offset of  $20_{16}$ ), with a trap type of  $01_{16}$ .
- All other virtual processors are in "parked" state (see *Parking and Unparking Virtual Processors* on page 526)

**Implementation Note** From the perspective of this specification, which describes a processor architecture, after a Power-On Reset (POR) execution begins on one strand of the processor. However, in a multiprocessor system, after POR a service processor might arrange for execution to initially occur on only one strand per system. If and how that that occurs is beyond the scope of this specification and would be described in system-level documentation.

<sup>1</sup>. per the Strand Available register (see *Strand Available Register (STRAND\_AVAILABLE)* on page 523)

**Programming Note** After a POR reset, software must initialize values that are specified as “undefined” in TABLE 16-1. In particular, I-cache tags, D-cache tags and L2 cache tags must be initialized before enabling the caches. The ITLB, DTLB and UTLB also must be initialized before enabling memory management. If a service processor participates in the reset, software should also reference the Service Processor Specification to determine which machine state has been reset by the service processor.

## 16.1.2 Warm Reset (WMR)

A Warm Reset (WMR) occurs when software writes into a particular implementation-dependent reset register or when an implementation-dependent reset input pin is asserted and then deasserted. When a WMR reset is received, all other resets and traps except POR are ignored.

The extent to which the processor is reset by a WMR reset is implementation dependent. A WMR reset may be chip-wide or it may be core-wide (resetting all virtual processors on the core, but allowing virtual processors on other cores to continue processing and maintaining cache coherency).

A WMR, even if it is chip-wide, will not alter the contents of external memory. It may, however, alter on-chip portions of the memory system (for example, store queues or cache(s)).

Warm reset has the same trap type ( $1_{16}$ ) and trap vector offset ( $20_{16}$ ) as a POR reset. By what means hyperprivileged software can distinguish between WMR and POR resets is implementation dependent.

**IMPL. DEP. #420-S10:** The following aspects of Warm Reset (WMR) are implementation dependent:

- (a) by what means WMR can be applied (for example, write to reset register or assertion/deassertion of an input pin)
- (b) the extent to which a processor is reset by WMR (for example, single physical core, entire processor (chip), and how the on-chip memory system is affected),
- (c) by what means hyperprivileged software can distinguish between WMR and POR resets

## 16.1.3 Externally Initiated Reset (XIR)

An externally initiated reset (XIR) is sent by asserting and deasserting an input pin, setting and clearing a bit in a reset register, or both.

An XIR reset is sent to all virtual processors specified in the XIR\_STEERING register (impl. dep. #304-U4-Cs10). It causes an XIR reset trap in each affected virtual processor. An XIR reset trap has trap type  $3_{16}$  and uses a trap vector with a physical address offset of  $60_{16}$ .

Memory state, cache state, and most architectural state (see TABLE 16-1) are unchanged by an XIR reset. System coherency is guaranteed to be maintained during an XIR reset. The PC (NPC) saved in TPC (TNPC) observed after an XIR will be mutually consistent, such that execution could resume using the saved PC and NPC. In effect, XIR behaves like a non-maskable interrupt.

## 16.1.4 Watchdog Reset (WDR)

An UltraSPARC Architecture virtual processor enters `error_state` when a non-reset trap or SIR reset occurs at  $TL = MAXTL$ .

The virtual processor signals itself internally to take a watchdog reset (WDR) and sets TT to the trap type of the trap that caused entry to `error_state`. The WDR causes a trap using a trap vector with a physical address offset of  $40_{16}$ . WDR only affects the virtual processor on which it occurs; no other virtual processors are affected.

On a watchdog reset trap caused by a register window-related trap, CWP register is updated the same as if a WDR had not occurred.

## 16.1.5 Software-Initiated Reset (SIR)

A software-initiated reset is initiated by an SIR instruction executing on a virtual processor. This virtual processor reset has a trap type 4 and uses a trap vector with a physical address offset of  $80_{16}$ . SIR affects only the virtual processor on which it executes; all other virtual processors are unaffected.

---

# 16.2 Machine States

Machine state changes when a trap is taken at  $TL = MAXTL - 1$  or when a reset occurs.

TABLE 16-1 specifies the machine states observed by software after a trap is taken at  $TL = MAXTL - 1$  or after a reset occurs. For details of how those machine states are set, see processor-specific documentation and/or relevant Service Processor documentation.

The UltraSPARC Architecture specifies the machine state that must be observed by software after reset.

<b>Implementation Note</b>	For the POR reset (and possibly the WMR reset), the change in machine state may be accomplished directly by processor hardware or with support from a service processor.
----------------------------	--

**Programming Note** Virtual processor states are *only* updated according to TABLE 16-1 if RED\_state is entered because of a trap at TL = MAXTL – 1 or a reset. If RED\_state is entered because the HPSTATE.red bit was explicitly set to 1 by software, then software is responsible for setting the appropriate machine state.

In the following tables, a value marked as "Undefined" may or may not be set to a known value by hardware (and/or a service processor) after reset.

**Programming Note** Values marked as "Undefined" after POR in the following tables should be initialized by software after the power-on reset.

**TABLE 16-1** Machine State After Reset or a Trap @ TL = MAXTL – 1 (1 of 3)

Name	Fields	POR	WMR	WDR	XIR	SIR	Traps taken @ TL=MAXTL-1
Integer registers		Undefined <sup>1</sup>	Unchanged				
Floating-point registers		Undefined					
RSTVADDR		VA = FFFF FFFF F000 0000 <sub>16</sub> PA = 0000 7FFF F000 0000 <sub>16</sub> (impl. dep. #114)					
PC		RSTVADDR   20 <sub>16</sub>	RSTVADDR   40 <sub>16</sub>	RSTVADDR   60 <sub>16</sub>	RSTVADDR   80 <sub>16</sub>	RSTVADDR   A0 <sub>16</sub>	
NPC		RSTVADDR   24 <sub>16</sub>	RSTVADDR   44 <sub>16</sub>	RSTVADDR   64 <sub>16</sub>	RSTVADDR   84 <sub>16</sub>	RSTVADDR   A4 <sub>16</sub>	
PSTATE	tct	0 (Trap on control transfer disabled)					
	mm	00 <sub>2</sub> (TSO)					
	pef	1 (FPU on)					
	am	0 (Full 64-bit address)					
	priv	0					
	ie	0 (Disable interrupts)					
	cle	0 (Current not little-endian)					
HPSTATE	tle	0 (Trap little-endian)	Unchanged				
	ibe	0 (Instruction breakpoint disabled)					
	red	1 (RED_state)					
	hpriv	1 (Hyperprivileged mode)					
	tlz	0 (trap_level_zero traps disabled)					
TBA<63:15>	tba_high49	Undefined	Unchanged				
HTBA<63:14>	htba_high50	Undefined	Unchanged				
Y		Undefined	Unchanged				
PIL		Undefined	Unchanged				
CWP		Undefined	Unchanged except for register window traps				
TT[TL]		1	1	trap type	3	4	trap type
CCR		Undefined	Unchanged				
ASI		Undefined	Unchanged				

**TABLE 16-1** Machine State After Reset or a Trap @ TL = MAXTL - 1 (2 of 3)

Name	Fields	POR	WMR	WDR	XIR	SIR	Traps taken @ TL=MAXTL-1
TL			MAXTL				min(TL+1, MAXTL)
GL			MAXGL				min(GL+1, MAXGL)
TPC[TL]		Undefined	(impl.dep.# 419-S10)‡				PC
TNPC[TL]		Undefined	(impl.dep.# 419-S10)‡				NPC
TSTATE[TL]	gl	Undefined	(impl.dep.# 419-S10)‡				GL
	ccr						CCR
	asi						ASI
	pstate						PSTATE
	cwp						CWP
HTSTATE[TL]	ibe	Undefined	(impl.dep.# 419-S10)‡				HPSTATE.ibe
	red						HPSTATE.red
	hpriv						HPSTATE.hpriv
	tlz						HPSTATE.tlz
TICK	npt	1	1				Unchanged
	counter	Undefined					Count
CANSAVE		Undefined					Unchanged
CANRESTORE		Undefined					Unchanged
OTHERWIN		Undefined					Unchanged
CLEANWIN		Undefined					Unchanged
WSTATE	other	Undefined					Unchanged
	normal						
HVER	manuf		<i>Implementation dependent (impl. dep. #104-V9)</i>				
	impl		<i>Implementation dependent (impl. dep. # 13-V8)</i>				
	mask		Mask dependent				
	maxgl		MAXGL				
	maxtl		MAXTL				
	maxwin		N_REG_WINDOWS - 1				
FSR	all	Undefined					Unchanged
GSR	all	Undefined					Unchanged
FPRS	fef	Undefined					Unchanged
	du						
	dl						
SOFTINT		Undefined					Unchanged
HINTP	hsp	Undefined					Unchanged
TICK_CMPR	int_dis	Undefined					Unchanged
	tick_cmpr						
STICK	npt	Undefined					Unchanged
	counter						

**TABLE 16-1** Machine State After Reset or a Trap @ TL = MAXTL – 1 (3 of 3)

Name	Fields	POR	WMR	WDR	XIR	SIR	Traps taken @ TL=MAXTL-1
STICK_CMPR	int_dis	Undefined				Unchanged	
	stick_cmpr					Unchanged	
HSTICK_CMPR	int_dis	Undefined				Unchanged	
	hstick_cmpr					Unchanged	
SCRATCHPAD_n		Undefined				Unchanged	
HYP_SCRATCHPAD_n		Undefined†	Unchanged†	Unchanged	Unchanged†	Unchanged	
I_SFSTR, D_SFSTR		Undefined				Unchanged	
D_SFAR		Undefined				Unchanged	
I-cache controls	enable(s)	0 (disable I\$)				Unchanged	
I-cache entries		Invalidated				Unchanged	
I-cache data		Undefined				Unchanged	
D-cache controls	enable(s)	0 (disable D\$)				Unchanged	
D-cache entries		Invalidated				Unchanged	
D-cache data		Undefined				Unchanged	
MMU controls /Demap	enable(s)	0 (disable MMU)				Unchanged	
MMU registers	all	Undefined				Unchanged	
ITLB/DTLB/UTLB entries		Invalidated				Unchanged	
store queue entries		Invalidated		Unchanged			
L2 cache controls	enable(s)	1 (enable L2\$)				Unchanged	
L2 cache entries		Invalidated				Unchanged	
L2 cache directory		Invalidated				Unchanged	
L2 cache data		Undefined				Unchanged	
Error enable registers		Undefined				Unchanged	
Error Trap enable registers		Undefined				Unchanged	
Error Status registers	error events	Undefined <sup>2</sup>				Unchanged	
Watchpoint Controls	enables	0 (disabled)				Unchanged	
Interrupt Queue pointers	all	Undefined				Unchanged	
Error Queue pointers	all	Undefined				Unchanged	

† If a service processor is present, it may change the value of hyperprivileged scratchpad register(s) before execution of the reset trap handler begins

‡ **IMPL. DEP. #419-S10:** It is implementation dependent whether, after a Warm Reset (WMR), the contents of TPC[TL], TNPC[TL], TSTATE[TL], and HTSTATE[TL] are unchanged from their values before the WMR, or contain the same values saved as during a WDR, XIR, or SIR reset. (The latter implementation is the preferred one.)

1. After POR, integer register R[0] must read as zero (with good ECC/parity). The value to which all other integer and all floating-point registers are set during a Power-On Reset (POR) is Undefined. For an implementation that protect these registers with ECC/parity, the registers must be initialized with good ECC/parity as part of a POR reset, either by hardware or software.
2. For a POR reset, the Error Status register(s) can be set either by hardware or by a service processor.

## 16.2.1 Machines States for CMT

TABLE 16-2 shows the CMT machine state set by hardware as a result of a trap taken at  $TL = MAXTL - 1$  or when a reset occurs.

**TABLE 16-2** Machine State After Reset and in RED\_state for CMT Registers

Name	Fields	POR	WMR	WDR	XIR	SIR	Traps taken @TL=MAXTL-1
<b>Registers shared among Virtual Processors (Strands)</b>							
STRAND_AVAILABLE		Unchanged (Predefined value, set at time of manufacture)					
STRAND_ENABLE_STATUS		Copied from STRAND_AVAILABLE† (but may be changed by a service processor during reset)	Copied from STRAND_ENABLE† (but may be changed by a service processor during reset)	Unchanged			
STRAND_ENABLE							
XIR_STEERING		Copied from STRAND_AVAILABLE† (impl. dep. #325-U4(b)) (but may be changed by a service processor during reset)	Copied from STRAND_ENABLE† (but may be changed by a service processor during reset)	Unchanged			
STRAND_RUNNING and STRAND_RUNNING_STATUS		Set to 0†, then during the reset either (1) virtual processor hardware sets to 1 the bit in the position corresponding to lowest-numbered <i>implemented and available</i> virtual processor (as specified by STRAND_AVAILABLE), or (2) this register is initialized by a service processor.	Set to 0†, then during the reset either (1) virtual processor hardware sets to 1 the bit in the position corresponding to lowest-numbered <i>enabled</i> virtual processor (as specified by the value of STRAND_ENABLE before the reset), or (2) this register is initialized by a service processor.	Unchanged			

**TABLE 16-2** Machine State After Reset and in RED\_state for CMT Registers (Continued)

Name	Fields	POR	WMR	WDR	XIR	SIR	Traps taken @TL=MAXTL-1
<b>Per-Strand Registers (not shared)</b>							
STRAND_ID	max_strand_id	max strand ID †	Unchanged				
	max_core_id	max core ID †					
	core_id	core ID † of this core					
CORE_INTR_ID	core_intr_id	interrupt ID † of this core					

† if the implementation is *always* paired with a service processor and the service processor *always* initializes this register during reset, processor hardware can leave this register unchanged (or set it to 0) and allow the service processor to perform the initialization



---



# Opcode Maps

This appendix contains the UltraSPARC Architecture 2005 instruction opcode maps. Also included are the optional UltraSPARC V instruction opcode maps; UltraSPARC V opcodes are highlighted in bold face.

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 136. 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 UltraSPARC Architecture 2005 strands.

Shading	Meaning
	An attempt to execute opcode will cause an <i>illegal_instruction</i> exception.
	An attempt to execute opcode will cause an <i>fp_exception_other</i> exception with FSR.ftt = 3 (unimplemented_FPop).

An attempt to execute a reserved opcode behaves as defined in *Reserved Opcodes and Instruction Fields* on page 132.

**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	BPcc (See TABLE A-7)	Bicc <sup>D</sup> (See TABLE A-7)	BPr (bit 28 = 0) (See TABLE A-8) — (bit 28 = 1) <sup>1</sup>	SETHI NOP <sup>2</sup>	FBPfcc (See TABLE A-7)	FBfcc <sup>D</sup> (See TABLE A-7)	—

1. See the footnote regarding bit 28 on page 160.

2. rd = 0, imm22 = 0

TABLE A-3 op3{5:0} (op = 10<sub>2</sub>) (1 of 2)

		op3{5:4}			
		0	1	2	3
op3 {3:0}	0	ADD	ADDcc	TADDcc	WRY <sup>D</sup> (rd = 0) — (rd = 1) WRCCR (rd = 2) WRASI (rd = 3) — (rd = 4, 5) SIR <sup>H</sup> (rd = 15, rs1 = 0, i = 1) — (rd = 15) <b>and</b> (rs1 ≠ 0 <b>or</b> i ≠ 1) — (rd = 7 – 14) WRFPRS (rd = 6) WRasr <sup>PASR</sup> (7 ≤ rd ≤ 14) WRPCR <sup>P</sup> (rd = 16) WRPIC (rd = 17) — (rd = 18) WRGSR (rd = 19) WRSOFTINT_SET <sup>P</sup> (rd = 20) WRSOFTINT_CLR <sup>P</sup> (rd = 21) WRSOFTINT <sup>P</sup> (rd = 22) WRTICK_CMPR <sup>P</sup> (rd = 23) WRSTICK <sup>H</sup> (rd = 24) WRSTICK_CMPR <sup>P</sup> (rd = 25) — (rd = 26 - 31)
	1	AND	ANDcc	TSUBcc	SAVED <sup>P</sup> (fcn = 0) RESTORED <sup>P</sup> (fcn = 1) ALLCLEAN <sup>P</sup> (fcn = 2) OTHERW <sup>P</sup> (fcn = 3) NORMALW <sup>P</sup> (fcn = 4) INVALW <sup>P</sup> (fcn = 5) — (fcn ≥ 6)
	2	OR	ORcc	TADDccTV <sup>D</sup>	—
	2	OR	ORcc	TADDccTV <sup>D</sup>	WRPR <sup>P</sup> (rd = 0-14 <b>or</b> 16) — (rd = 15 <b>or</b> 17–31)
	3	XOR	XORcc	TSUBccTV <sup>D</sup>	WRHPR <sup>H</sup>
	4	SUB	SUBcc	MULScc <sup>D</sup>	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)	IMPDEP1 (VIS) (See TABLE A-12)
	7	XNOR	XNORcc	SRA (x = 0), SRAX (x = 1)	IMPDEP2

TABLE A-3 op3{5:0} (op = 10<sub>2</sub>) (2 of 2)

		op3{5:4}			
		0	1	2	3
op3 {3:0}	<b>8</b>	ADDC	ADDC <sub>cc</sub>	RDY <sup>D</sup> (rs1 = 0, i = 0) — (rs1 = 1, i = 0) RDCCR (rs1 = 2, i = 0) RDASI (rs1 = 3, i = 0) RDTICK <sup>Pnpt</sup> (rs1 = 4, i = 0) RDPC (rs1 = 5, i = 0) RDFPRS (rs1 = 6, i = 0) RDasr <sup>PASR</sup> (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)) STBAR <sup>D</sup> (rs1 = 15, rd = 0, i = 0) — (rs1 = 15 and rd > 0 and i = 0) RDPCR <sup>P</sup> (rs1 = 16 and i = 0) RDPIC (rs1 = 17 and i = 0) — (rs1 = 18 and i = 0) RDGSR (rs1 = 19 and i = 0) — (rs1 = 20 or 21) and (i = 0)) RDSOFTINT <sup>P</sup> (rs1 = 22 and i = 0) RDTICK_CMPR <sup>P</sup> (rs1 = 23 and i = 0) RDSTICK (rs1 = 24 and i = 0) RDSTICK_CMPR <sup>P</sup> (rs1 = 25 and i = 0) — ((rs1 = 26 – 31) and (i = 0))	JMPL
	<b>9</b>	MULX	—	RDHPR <sup>H</sup>	RETURN
	<b>A</b>	UMUL <sup>D</sup>	UMUL <sub>cc</sub> <sup>D</sup>	RDPR <sup>P</sup> (rs1 = 1–14 or 16) — (rs1 = 15 or 17 – 30)	Tcc ((i = 0 and inst{10:5} = 0) or ((i = 1) and (inst{10:7} = 0))) (See TABLE A-7) — (bit 29 = 1) — ((i = 0 and (inst{10:5} ≠ 0)) or (i = 1 and (inst{10:7} ≠ 0)))
	<b>B</b>	SMUL <sup>D</sup>	SMUL <sub>cc</sub> <sup>D</sup>	FLUSHW	FLUSH
op3 {3:0}	<b>C</b>	SUBC	SUBC <sub>cc</sub>	MOV <sub>cc</sub>	SAVE
	<b>D</b>	UDIVX	—	SDIVX	RESTORE
	<b>E</b>	UDIV <sup>D</sup>	UDIV <sub>cc</sub> <sup>D</sup>	POPC (rs1 = 0) — (rs1 > 0)	DONE <sup>P</sup> (fcn = 0) RETRY <sup>P</sup> (fcn = 1) — (fcn = 2..15) — (fcn = 16..31)
	<b>F</b>	SDIV <sup>D</sup>	SDIV <sub>cc</sub> <sup>D</sup>	MOV <sub>r</sub> (See TABLE A-8)	—

TABLE A-4 op3{5:0} (op = 11<sub>2</sub>)

		op3{5:4}			
		0	1	2	3
<b>op3 {3:0}</b>	<b>0</b>	LDUW	LDUWA <sup>PASI</sup>	LDF	LDFA <sup>PASI</sup>
	<b>1</b>	LDUB	LDUBA <sup>PASI</sup>	LDFSR <sup>D</sup> , LDXFSR — (rd > 1)	—
	<b>2</b>	LDUH	LDUHA <sup>PASI</sup>	LDQF	LDQFA <sup>PASI</sup>
	<b>3</b>	LDTW <sup>D</sup> — (rd odd)	LDTWA <sup>D, PASI</sup> — (rd odd)	LDDEF	LDDEFA <sup>PASI</sup> LDBLOCKF LDSHORTF
	<b>4</b>	STW	STWA <sup>PASI</sup>	STF	STFA <sup>PASI</sup>
	<b>5</b>	STB	STBA <sup>PASI</sup>	STFSR <sup>D</sup> , STXFSR — (rd > 1)	—
	<b>6</b>	STH	STHA <sup>PASI</sup>	STQF	STQFA <sup>PASI</sup>
	<b>7</b>	STTW <sup>D</sup> — (rd odd)	STTWA <sup>PASI</sup> — (rd odd)	STDF	STDFA <sup>PASI</sup> STLBLOCKF STPARTIALF STSHORTF
	<b>8</b>	LDSW	LDSWA <sup>PASI</sup>	—	—
	<b>9</b>	LDSB	LDSBA <sup>PASI</sup>	—	—
	<b>A</b>	LDSH	LDSHA <sup>PASI</sup>	—	—
	<b>B</b>	LDX	LDXA <sup>PASI</sup>	—	—
	<b>C</b>	—	—	—	CASA <sup>PASI</sup>
	<b>D</b>	LDSTUB	LDSTUBA <sup>PASI</sup>	PREFETCH — (fcn = 5 – 15)	PREFETCHA <sup>PASI</sup> — (fcn = 5 – 15)
	<b>E</b>	STX	STXA <sup>PASI</sup>	—	CASXA <sup>PASI</sup>
	<b>F</b>	SWAP <sup>D</sup>	SWAPA <sup>D, PASI</sup>	—	—

TABLE A-5 opf{8:0} (op = 10<sub>2</sub>, op3 = 34<sub>16</sub> = FPop1)

opf{8:4}	opf{3:0}							
	0	1	2	3	4	5	6	7
00 <sub>16</sub>	—	FMOV <sub>s</sub>	FMOV <sub>d</sub>	FMOV <sub>q</sub>	—	FNEG <sub>s</sub>	FNEG <sub>d</sub>	FNEG <sub>q</sub>
01 <sub>16</sub>	—	—	—	—	—	—	—	—
02 <sub>16</sub>	—	—	—	—	—	—	—	—
03 <sub>16</sub>	—	—	—	—	—	—	—	—
04 <sub>16</sub>	—	FADD <sub>s</sub>	FADD <sub>d</sub>	FADD <sub>q</sub>	—	FSUB <sub>s</sub>	FSUB <sub>d</sub>	FSUB <sub>q</sub>
05 <sub>16</sub>	—	—	—	—	—	—	—	—
06 <sub>16</sub>	—	—	—	—	—	—	—	—
07 <sub>16</sub>	—	—	—	—	—	—	—	—
08 <sub>16</sub>	—	FsTO <sub>x</sub>	FdTO <sub>x</sub>	FqTO <sub>x</sub>	FxTO <sub>s</sub>	—	—	—
09 <sub>16</sub>	—	—	—	—	—	—	—	—
0A <sub>16</sub>	—	—	—	—	—	—	—	—
0B <sub>16</sub>	—	—	—	—	—	—	—	—
0C <sub>16</sub>	—	—	—	—	FiTO <sub>s</sub>	—	FdTO <sub>s</sub>	FqTO <sub>s</sub>
0D <sub>16</sub>	—	FsTO <sub>i</sub>	FdTO <sub>i</sub>	FqTO <sub>i</sub>	—	—	—	—
0E <sub>16</sub> –1F <sub>16</sub>	—	—	—	—	—	—	—	—
	8	9	A	B	C	D	E	F
00 <sub>16</sub>	—	FABS <sub>s</sub>	FABS <sub>d</sub>	FABS <sub>q</sub>	—	—	—	—
01 <sub>16</sub>	—	—	—	—	—	—	—	—
02 <sub>16</sub>	—	FSQRT <sub>s</sub>	FSQRT <sub>d</sub>	FSQRT <sub>q</sub>	—	—	—	—
03 <sub>16</sub>	—	—	—	—	—	—	—	—
04 <sub>16</sub>	—	FMUL <sub>s</sub>	FMUL <sub>d</sub>	FMUL <sub>q</sub>	—	FDIV <sub>s</sub>	FDIV <sub>d</sub>	FDIV <sub>q</sub>
05 <sub>16</sub>	—	—	—	—	—	—	—	—
06 <sub>16</sub>	—	FsMUL <sub>d</sub>	—	—	—	—	FdMUL <sub>q</sub>	—
07 <sub>16</sub>	—	—	—	—	—	—	—	—
08 <sub>16</sub>	FxTO <sub>d</sub>	—	—	—	FxTO <sub>q</sub>	—	—	—
09 <sub>16</sub>	—	—	—	—	—	—	—	—
0A <sub>16</sub>	—	—	—	—	—	—	—	—
0B <sub>16</sub>	—	—	—	—	—	—	—	—
0C <sub>16</sub>	FiTO <sub>d</sub>	FsTO <sub>d</sub>	—	FqTO <sub>d</sub>	FiTO <sub>q</sub>	FsTO <sub>q</sub>	FdTO <sub>q</sub>	—
0D <sub>16</sub>	—	—	—	—	—	—	—	—
0E <sub>16</sub> –1F <sub>16</sub>	—	—	—	—	—	—	—	—

**TABLE A-6** opf{8:0} (op = 10<sub>2</sub>, op3 = 35<sub>16</sub> = FPop2)

opf{8:4}	opf{3:0}								
	0	1	2	3	4	5	6	7	8-F
00 <sub>16</sub>	—	FMOV <sub>s</sub> (fcc0)	FMOV <sub>d</sub> (fcc0)	FMOV <sub>q</sub> (fcc0)	—	† ‡	† ‡	† ‡	—
01 <sub>16</sub>	—	—	—	—	—	—	—	—	—
02 <sub>16</sub>	—	—	—	—	—	FMOVR <sub>s</sub> Z ‡	FMOVR <sub>d</sub> Z ‡	FMOVR <sub>q</sub> Z ‡	—
03 <sub>16</sub>	—	—	—	—	—	—	—	—	—
04 <sub>16</sub>	—	FMOV <sub>s</sub> (fcc1)	FMOV <sub>d</sub> (fcc1)	FMOV <sub>q</sub> (fcc1)	—	FMOVR <sub>s</sub> LEZ ‡	FMOVR <sub>d</sub> LEZ ‡	FMOVR <sub>q</sub> LEZ ‡	—
05 <sub>16</sub>	—	FCMP <sub>s</sub>	FCMP <sub>d</sub>	FCMP <sub>q</sub>	—	FCMP <sub>E</sub> s ‡	FCMP <sub>E</sub> d ‡	FCMP <sub>E</sub> q ‡	—
06 <sub>16</sub>	—	—	—	—	—	FMOVR <sub>s</sub> LZ ‡	FMOVR <sub>d</sub> LZ ‡	FMOVR <sub>q</sub> LZ ‡	—
07 <sub>16</sub>	—	—	—	—	—	—	—	—	—
08 <sub>16</sub>	—	FMOV <sub>s</sub> (fcc2)	FMOV <sub>d</sub> (fcc2)	FMOV <sub>q</sub> (fcc2)	—	†	†	†	—
09 <sub>16</sub>	—	—	—	—	—	—	—	—	—
0A <sub>16</sub>	—	—	—	—	—	FMOVR <sub>s</sub> NZ ‡	FMOVR <sub>d</sub> NZ ‡	FMOVR <sub>q</sub> NZ ‡	—
0B <sub>16</sub>	—	—	—	—	—	—	—	—	—
0C <sub>16</sub>	—	FMOV <sub>s</sub> (fcc3)	FMOV <sub>d</sub> (fcc3)	FMOV <sub>q</sub> (fcc3)	—	FMOVR <sub>s</sub> GZ ‡	FMOVR <sub>d</sub> GZ ‡	FMOVR <sub>q</sub> GZ ‡	—
0D <sub>16</sub>	—	—	—	—	—	—	—	—	—
0E <sub>16</sub>	—	—	—	—	—	FMOVR <sub>s</sub> GEZ ‡	FMOVR <sub>d</sub> GEZ ‡	FMOVR <sub>q</sub> GEZ ‡	—
0F <sub>16</sub>	—	—	—	—	—	—	—	—	—
10 <sub>16</sub>	—	FMOV <sub>s</sub> (icc)	FMOV <sub>d</sub> (icc)	FMOV <sub>q</sub> (icc)	—	—	—	—	—
11 <sub>16</sub> –17 <sub>16</sub>	—	—	—	—	—	—	—	—	—
18 <sub>16</sub>	—	FMOV <sub>s</sub> (xcc)	FMOV <sub>d</sub> (xcc)	FMOV <sub>q</sub> (xcc)	—	—	—	—	—
19 <sub>16</sub> –1F <sub>16</sub>	—	—	—	—	—	—	—	—	—

† Reserved variation of FMOVR

‡ bit 13 of instruction = 0

TABLE A-7 cond{3:0}

		<b>BPcc</b> <b>op = 0</b> <b>op2 = 1</b>	<b>Bicc</b> <b>op = 0</b> <b>op2 = 2</b>	<b>FBPfcc</b> <b>op = 0</b> <b>op2 = 5</b>	<b>FBfcc<sup>D</sup></b> <b>op = 0</b> <b>op2 = 6</b>	<b>Tcc</b> <b>op = 2</b> <b>op3 = 3a<sub>16</sub></b>
<b>cond</b> <b>{3:0}</b>	<b>0</b>	BPN	BN <sup>D</sup>	FBPN	FBN <sup>D</sup>	TN
	<b>1</b>	BPE	BE <sup>D</sup>	FBPNE	FBNE <sup>D</sup>	TE
	<b>2</b>	BPLE	BLE <sup>D</sup>	FBPLG	FBLG <sup>D</sup>	TLE
	<b>3</b>	BPL	BL <sup>D</sup>	FBPUL	FBUL <sup>D</sup>	TL
	<b>4</b>	BPLEU	BLEU <sup>D</sup>	FBPL	FBL <sup>D</sup>	TLEU
	<b>5</b>	BPCS	BCS <sup>D</sup>	FBPUG	FBUG <sup>D</sup>	TCS
	<b>6</b>	BPNEG	BNEG <sup>D</sup>	FBPG	FBG <sup>D</sup>	TNEG
	<b>7</b>	BPVS	BVS <sup>D</sup>	FBPU	FBU <sup>D</sup>	TVS
	<b>8</b>	BPA	BA <sup>D</sup>	FBPA	FBA <sup>D</sup>	TA
	<b>9</b>	BPNE	BNE <sup>D</sup>	FBPE	FBE <sup>D</sup>	TNE
	<b>A</b>	BPG	BG <sup>D</sup>	FBPUE	FBUE <sup>D</sup>	TG
	<b>B</b>	BPGE	BGE <sup>D</sup>	FBPGE	FBGE <sup>D</sup>	TGE
	<b>C</b>	BPGU	BGU <sup>D</sup>	FBPUGE	FBUGE <sup>D</sup>	TGU
	<b>D</b>	BPCC	BCC <sup>D</sup>	FBPLE	FBLE <sup>D</sup>	TCC
	<b>E</b>	BPPOS	BPOS <sup>D</sup>	FBPULE	FBULE <sup>D</sup>	TPOS
<b>F</b>	BPVC	BVC <sup>D</sup>	FBPO	FBO <sup>D</sup>	TVC	

TABLE A-8 Encoding of rcond{2:0} Instruction Field

		<b>BPr</b> <b>op = 0</b> <b>op2 = 3</b>	<b>MOVr</b> <b>op = 2</b> <b>op3 = 2F<sub>16</sub></b>	<b>FMOVr</b> <b>op = 2</b> <b>op3 = 35<sub>16</sub></b>
<b>rcond</b> <b>{2:0}</b>	<b>0</b>	—	—	—
	<b>1</b>	BRZ	MOVZR	FMOVr<s d q>Z
	<b>2</b>	BRLEZ	MOVRLZ	FMOVr<s d q>LEZ
	<b>3</b>	BRLZ	MOVRLZ	FMOVr<s d q>LZ
	<b>4</b>	—	—	—
	<b>5</b>	BRNZ	MOVRNZ	FMOVr<s d q>NZ
	<b>6</b>	BRGZ	MOVRGZ	FMOVr<s d q>GZ
	<b>7</b>	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
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** IMPDEP1: opf{8:0} for VIS opcodes (op = 10<sub>2</sub>, op3 = 36<sub>16</sub>)

		opf {8:4}								
		00	01	02	03	04	05	06	07	08
opf	0	EDGE8	ARRAY8	FCMPLE16	—	—	FPADD16	FZERO	FAND	SHUT DOWN <sup>D,P</sup>
	1	EDGE8N	—	—	FMUL 8x16	—	FPADD16S	FZEROS	FANDS	SIAM
	2	EDGE8L	ARRAY16	FCMPNE16	—	—	FPADD32	FNOR	FXNOR	—
	3	EDGE8LN	—	—	FMUL 8x16AU	—	FPADD32S	FNORS	FXNORS	—
	4	EDGE16	ARRAY32	FCMPLE32	—	—	FPSUB16	FANDNOT2	FSRC1	—
	5	EDGE16N	—	—	FMUL 8x16AL	—	FPSUB16S	FANDNOT2S	FSRC1S	—
	6	EDGE16L	—	FCMPNE32	FMUL 8SUx16	—	FPSUB32	FNOT2	FORNOT2	—
	7	EDGE16LN	—	—	FMUL 8ULx16	—	FPSUB32S	FNOT2S	FORNOT2S	—
	8	EDGE32	ALIGN ADDRESS	FCMPGT16	FMULD 8SUx16	FALIGN DATA	—	FANDNOT1	FSRC2	—
	9	EDGE32N	BMASK	—	FMULD 8ULx16	—	—	FANDNOT1S	FSRC2S	—
	A	EDGE32L	ALIGN ADDRESS _LITTLE	FCMPEQ16	FPACK32	—	—	FNOT1	FORNOT1	—
	B	EDGE32LN	—	—	FPACK16	FPMERGE	—	FNOT1S	FORNOT1S	—
	C	—	—	FCMPGT32	—	BSHUFFLE	—	FXOR	FOR	—
	D	—	—	—	FPACKFIX	FEXPAND	—	FXORS	FORS	—
	E	—	—	FCMPEQ32	PDIST	—	—	FNAND	FONE	—
	F	—	—	—	—	—	—	FNANDS	FONES	—

**TABLE A-14** IMPDEP1:  $\text{opf}\{8:0\}$  for VIS opcodes ( $\text{op} = 10_2$ ,  $\text{op3} = 36_{16}$ ) (3 of 3)

		opf {8:4}							
		09-1F	10	11	12	13	14	15	16-1F
opf {3:0}	0	—	—	—		—	—	—	—
	1	—	—	—	—	—	—	—	—
	2	—	—	—		—	—	—	—
	3	—	—	—	—	—	—	—	—
	4	—	—	—	—	—	—	—	—
	5	—	—	—	—	—	—	—	—
	6	—	—	—	—	—	—	—	—
	7	—	—	—	—	—	—	—	—
	8	—	—	—		—	—	—	—
	9	—	—	—	—	—	—	—	—
	A	—	—	—		—	—	—	—
	B	—	—	—	—	—	—	—	—
	C	—	—	—	—	—	—	—	—
	D	—	—	—	—	—	—	—	—
	E	—	—	—	—	—	—	—	—
	F	—	—	—	—	—	—	—	—

**Note: This chapter is undergoing final review; please check back later for a copy of UltraSPARC Architecture 2005 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 569.

The appendix contains these sections:

- **Definition of an Implementation Dependency** on page 567.
- **Hardware Characteristics** on page 568.
- **Implementation Dependency Categories** on page 568.
- **List of Implementation Dependencies** on page 569.

---

### 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 *UltraSPARC Architecture 2005 Compliance with SPARC V9 Architecture* on page 25.

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 10)

<b>Nbr</b>	<b>Category</b>	<b>Description</b>	<b>Page</b>
<b>1-V8</b>	<b>f</b>	<b>Software emulation of instructions</b> Whether an instruction complies with UltraSPARC Architecture 2005 by being implemented directly by hardware, simulated by software, or emulated by firmware is implementation dependent.	<b>26</b>
<b>2-V8</b>	<b>v</b>	<b>Number of IU registers</b> An UltraSPARC Architecture implementation may contain from 72 to 640 general-purpose 64-bit R registers. This corresponds to a grouping of the registers into $MAXGL + 1$ 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).	<b>26, 51</b>
<b>3-V8</b>	<b>f</b>	<b>Incorrect IEEE Std 754-1985 results</b> 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 <code>FSR.ftt = unfinished_FPop</code> or <code>FSR.ftt = unimplemented_FPop</code> . In this case, software running in a higher privilege mode shall emulate any functionality not present in the hardware.	<b>132</b>
<b>4, 5</b>		<i>Reserved.</i>	

TABLE B-1 SPARC V9 Implementation Dependencies (2 of 10)

Nbr	Category	Description	Page
6-V8	f	<b>I/O registers privileged status</b> Whether I/O registers can be accessed by nonprivileged code is implementation dependent.	29
7-V8	t	<b>I/O register definitions</b> The contents and addresses of I/O registers are implementation dependent.	30
8-V8- Cs20	t	<b>RDAsr/WRAsr target registers</b> Ancillary state registers (ASRs) in the range 0–27 that are not defined in UltraSPARC Architecture 2005 are reserved for future architectural use. ASRs in the range 28–31 are available to be used for implementation-dependent purposes.	31, 70, 299, 369
9-V8- Cs20	f	<b>RDAsr/WRAsr privileged status</b> Whether each of the implementation-dependent read/write ancillary state register instructions (for ASRs 28–31) is privileged is implementation dependent.	31, 70, 299, 369
10-V8–12-V8		<i>Reserved.</i>	
13-V8	a	<b>HVER.impl</b> HVER.impl uniquely identifies an implementation or class of software-compatible implementations of the architecture. Values FFF0 <sub>16</sub> –FFFF <sub>16</sub> are reserved and are not available for assignment.	108
14-V8–15-V8		<i>Reserved.</i>	
16-V8-Cu3		<i>Reserved.</i>	
17-V8		<i>Reserved.</i>	
18- V8- Ms10	f	<b>Nonstandard IEEE 754-1985 results</b> UltraSPARC Architecture 2005 implementations do not implement a nonstandard floating-point mode. FSR.ns is a reserved bit; it always reads as 0 and writes to it are ignored.	63
19-V8	a	<b>FPU version, FSR.ver</b> Bits 19:17 of the FSR, FSR.ver, identify one or more implementations of the FPU architecture.	63
20-V8–21-V8		<i>Reserved.</i>	
22-V8	f	<b>FPU tem, cexc, and aexc</b> An UltraSPARC Architecture implementation implements the tem, cexc, and aexc fields in hardware, conformant to IEEE Std 754-1985.	70
23-V8		<i>Reserved.</i>	
24-V8		<i>Reserved.</i>	
25-V8	f	<b>RDPR of FQ with nonexistent FQ</b> An UltraSPARC 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.	66, 304
26-V8–28-V8		<i>Reserved.</i>	

TABLE B-1 SPARC V9 Implementation Dependencies (3 of 10)

Nbr	Category	Description	Page
29-V8	t	<p><b>Address space identifier (ASI) definitions</b></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 UltraSPARC Architecture. Others remain implementation dependent in the UltraSPARC Architecture. See <i>ASI Assignments</i> on page 408 and <i>Block Load and Store ASIs</i> on page 429 for details.</p>	121
30-V8-Cu3	f	<p><b>ASI address decoding</b></p> <p>In SPARC V9, an implementation could choose to decode only a subset of the 8-bit ASI specifier. In UltraSPARC 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>	121
31-V8-Cs10	f	<p>This implementation dependency is no longer used in the UltraSPARC Architecture, since “catastrophic” errors are now handled using normal error-reporting mechanisms.</p>	—
32-V8-Ms10	t	<p><b>Restartable deferred traps</b></p> <p>Whether any restartable deferred traps (and associated deferred-trap queues) are present is implementation dependent.</p>	447
33-V8-Cs10	f	<p><b>Trap precision</b></p> <p>In an UltraSPARC Architecture implementation, all exceptions that occur as the result of program execution, except for <i>store_error</i>, are precise.</p>	451
34-V8	f	<p><b>Interrupt clearing</b></p> <p><b>a:</b> The method by which an interrupt is removed is now defined in the UltraSPARC Architecture (see <i>Clearing the Software Interrupt Register</i> on page 497).</p> <p><b>b:</b> How quickly a virtual processor responds to an interrupt request, like all timing-related issues, is implementation dependent.</p>	497
35-V8-Cs20	t	<p><b>Implementation-dependent traps</b></p> <p>Trap type (TT) values <math>060_{16}</math>–<math>07F_{16}</math> were reserved for <i>implementation_dependent_exception_n</i> exceptions in SPARC V9 but are now all defined as standard UltraSPARC Architecture exceptions.</p>	457
36-V8	f	<p><b>Trap priorities</b></p> <p>The relative priorities of traps defined in the UltraSPARC 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.</p>	469
37-V8	f	<p><b>Reset trap</b></p> <p>Some of a virtual processor’s behavior during a reset trap is implementation dependent.</p>	449
38-V8	f	<p><b>Effect of reset trap on implementation-dependent registers</b></p> <p>Implementation-dependent registers may or may not be affected by the various reset traps.</p>	474

TABLE B-1 SPARC V9 Implementation Dependencies (4 of 10)

Nbr	Category	Description	Page
39-V8-Cs10	f	<b>Entering error_state on implementation-dependent errors</b> The virtual processor enters <code>error_state</code> when a trap occurs while the virtual processor is already at its maximum supported trap level — that is, it enters <code>error_state</code> when a trap occurs while <code>TL = MAXTL</code> . No other conditions cause entry into <code>error_state</code> on an UltraSPARC Architecture virtual processor.	444, 477
40-V8	f	<b>error_state virtual processor state</b> Effects when <code>error_state</code> is entered are implementation dependent, but it is recommended that as much virtual processor state as possible be preserved upon entry to <code>error_state</code> . In addition, an UltraSPARC Architecture virtual processor may have other <code>error_state</code> entry traps that are implementation dependent.	444
41-V8		<i>Reserved.</i>	
42-V8-Cs10	t, f, v	<b>FLUSH instruction</b> FLUSH is implemented in hardware in all UltraSPARC Architecture 2005 implementations, so never causes a trap as an unimplemented instruction.	
43-V8		<i>Reserved.</i>	
44-V8-Cs10	f	<b>Data access FPU trap</b> <b>a:</b> 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) are undefined or are guaranteed to remain unchanged. <b>b:</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.	251 254
45-V8–46-V8		<i>Reserved.</i>	
47-V8-Cs20	t	<b>RDAsr</b> RDAsr instructions with <code>rd</code> in the range 28–31 are available for implementation-dependent uses (impl. dep. #8-V8-Cs20). For an RDAsr instruction with <code>rs1</code> in the range 28–31, the following are implementation dependent: <ul style="list-style-type: none"> <li>the interpretation of bits 13:0 and 29:25 in the instruction</li> <li>whether the instruction is nonprivileged or privileged or hyperprivileged (impl. dep. #9-V8-Cs20)</li> <li>whether an attempt to execute the instruction causes an <i>illegal_instruction</i> exception</li> </ul>	300
48-V8-Cs20	t	<b>WRAsr</b> WRAsr instructions with <code>rd</code> in the range 26–31 are available for implementation-dependent uses (impl. dep. #8-V8-Cs20). For a WRAsr instruction with <code>rd</code> in the range 26–31, the following are implementation dependent: <ul style="list-style-type: none"> <li>the interpretation of bits 18:0 in the instruction</li> <li>the operation(s) performed (for example, <code>xor</code>) to generate the value written to the ASR</li> <li>whether the instruction is nonprivileged or privileged or hyperprivileged (impl. dep. #9-V8-Cs20)</li> <li>whether an attempt to execute the instruction causes an <i>illegal_instruction</i> exception</li> </ul>	370
49-V8–54-V8		<i>Reserved.</i>	

TABLE B-1 SPARC V9 Implementation Dependencies (5 of 10)

Nbr	Category	Description	Page
55- V8- Cs10	f	<p><b>Tininess detection</b></p> <p>In SPARC V9, it is implementation-dependent whether “tininess” (an IEEE 754 term) is detected before or after rounding. In all UltraSPARC Architecture implementations, tininess is detected before rounding.</p>	69
56–100		<i>Reserved.</i>	
101- V9- CS10		<p><b>Maximum trap level (MAXPTL, MAXTL)</b></p> <p>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 visible to privileged software). In a typical implementation <i>MAXPTL</i> = <i>MAXPGL</i> (see impl. dep. #401-S10).</p> <p>The architectural parameter <i>MAXTL</i> is a constant for each implementation; its legal values are from 3 to 7 (supporting from 3 to 7 levels of saved trap state).</p> <p>Architecturally, <i>MAXPTL</i> must be <math>\geq 2</math>, <i>MAXTL</i> must be <math>\geq 4</math>, and <i>MAXTL</i> must be <math>&gt; MAXPTL</math>.</p>	99, 101
102- V9	f	<p><b>Clean windows trap</b></p> <p>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.</p>	482
103- V9- Ms10	f	<p><b>Prefetch instructions</b></p> <p>The following aspects of the PREFETCH and PREFETCHA instructions are implementation dependent:</p> <p><b>a:</b> the attributes of the block of memory prefetched: its size (minimum = 64 bytes) and its alignment (minimum = 64-byte alignment)</p> <p><b>b:</b> whether each defined prefetch variant is implemented (1) as a NOP, (2) with its full semantics, or (3) with common-case prefetching semantics</p> <p><b>c:</b> whether and how variants 16, 18, 19 and 24–31 are implemented; if not implemented, a variant must execute as a NOP</p> <p>The following aspects of the PREFETCH and PREFETCHA instructions used to be (but are no longer) implementation dependent:</p> <p><b>d:</b> while in nonprivileged mode (<i>PSTATE.priv</i> = 0 and <i>HPSTATE.hpriv</i> = 0), an attempt to reference an ASI in the range <math>0_{16}..7F_{16}</math> by a PREFETCHA instruction executes as a NOP; specifically, it does not cause a <i>privileged_action</i> exception.</p> <p><b>e:</b> PREFETCH and PREFETCHA have no observable effect in privileged code</p> <p><b>f:</b> In UltraSPARC Architecture 2005, neither PREFETCH nor PREFETCHA can cause a <i>data_access_MMU_miss</i> exception (because a Strong prefetch is treated the same as a Weak prefetch)</p> <p><b>g:</b> while in privileged mode (<i>PSTATE.priv</i> = 1 and <i>HPSTATE.hpriv</i> = 0), an attempt to reference an ASI in the range <math>30_{16}..7F_{16}</math> by a PREFETCHA instruction executes as a NOP (specifically, it does not cause a <i>privileged_action</i> exception)</p>	292 292, 295 297C

**TABLE B-1** SPARC V9 Implementation Dependencies (6 of 10)

<b>Nbr</b>	<b>Category</b>	<b>Description</b>	<b>Page</b>
<b>104-V9</b>	a	<b>HVER.manuf</b> HVER.manuf contains a 16-bit semiconductor manufacturer code. This field is optional and, if not present, reads as zero. VER.manuf may indicate the original supplier of a second-sourced processor in cases involving mask-level second-sourcing. It is intended that the contents of HVER.manuf track the JEDEC semiconductor manufacturer code as closely as possible. If the manufacturer does not have a JEDEC semiconductor manufacturer code, then SPARC International will assign a HVER.manuf value.	<b>108</b>
<b>105-V9</b>	f	<b>TICK register</b> <b>a:</b> If an accurate count cannot always be returned when TICK is read, any inaccuracy should be small, bounded, and documented. <b>b:</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 0.	<b>75</b>
<b>106-V9</b>	f	<b>IMPDEP2A instructions</b> The IMPDEP2A instructions are completely implementation dependent. Implementation-dependent aspects include their operation, the interpretation of bits 29:25 and 18:0 in their encodings, and which (if any) exceptions they may cause.	<b>236</b>
<b>107-V9</b>	f	<b>Unimplemented LDTW(A) trap</b> <b>a:</b> 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>b:</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.	<b>266</b> <b>269</b>
<b>108-V9</b>	f	<b>Unimplemented STTW(A) trap</b> <b>a:</b> 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. <b>b:</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.	<b>346</b> <b>349</b>

TABLE B-1 SPARC V9 Implementation Dependencies (7 of 10)

Nbr	Category	Description	Page
109- V9- Cs10	f	<p><b><i>LDDF(A)_mem_address_not_aligned</i></b></p> <p><b>a:</b> 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 UltraSPARC 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>b:</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 UltraSPARC Architecture processor, the <i>LDDF_mem_address_not_aligned</i> exception occurs in this case and trap handler software emulates the LDDFA instruction)</p>	114, 114, 250, 487          253
110- V9- Cs10	f	<p><b><i>STDF(A)_mem_address_not_aligned</i></b></p> <p><b>a:</b> 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 UltraSPARC 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>b:</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 UltraSPARC Architecture processor, the <i>STDF_mem_address_not_aligned</i> exception occurs in this case and trap handler software emulates the STDFA instruction)</p>	114, 333, 489          336

TABLE B-1 SPARC V9 Implementation Dependencies (8 of 10)

Nbr	Category	Description	Page
111- V9- Cs10	f	<p><b><i>LDQF(A)_mem_address_not_aligned</i></b></p> <p><b>a:</b> 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 UltraSPARC 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 UltraSPARC Architecture 2005 implementations, because they do not implement the LDQF instruction in hardware)</p> <p><b>b:</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 UltraSPARC 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 UltraSPARC Architecture 2005 implementations, because they do not implement the LDQFA instruction in hardware)</p>	115, 114, 250, 491          253
112- V9- Cs10	f	<p><b><i>STQF(A)_mem_address_not_aligned</i></b></p> <p><b>a:</b> 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 UltraSPARC 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 UltraSPARC Architecture 2005 implementations, because they do not implement the STQF instruction in hardware)</p> <p><b>b:</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 UltraSPARC 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 UltraSPARC Architecture 2005 implementations, because they do not implement the STQFA instruction in hardware)</p>	115, 333, 491          336

TABLE B-1 SPARC V9 Implementation Dependencies (9 of 10)

Nbr	Category	Description	Page
113- V9- Ms10	f	<p><b>Implemented memory models</b></p> <p>Whether memory models represented by <code>PSTATE.mm = 10<sub>2</sub></code> or <code>11<sub>2</sub></code> are supported in an UltraSPARC Architecture processor is implementation dependent. If the <code>10<sub>2</sub></code> model is supported, then when <code>PSTATE.mm = 10<sub>2</sub></code> the implementation must correctly execute software that adheres to the RMO model described in <i>The SPARC Architecture Manual-Version 9</i>. If the <code>11<sub>2</sub></code> model is supported, its definition is implementation dependent.</p>	95, 396
114- V9- Cs10	f	<p><b>RED_state trap vector address (RSTVADDR)</b></p> <p>The <code>RED_state</code> trap vector is located at an address referred to as <code>RSTVADDR</code>. In the UltraSPARC Architecture, <code>RSTVADDR</code> is bound to the following address:</p> <p style="padding-left: 40px;">Physical Address      <code>FFFF FFFF F000 0000</code><sub>16</sub> (the highest 256MB of physical address space)</p> <p>In an implementation that implements fewer than 64 bits of physical addressing, unimplemented high-order bits of the above <code>RSTVADDR</code> are ignored.</p>	455, 549
115- V9	f	<p><b>RED_state</b></p> <p>What occurs after the processor enters <code>RED_state</code> is implementation dependent.</p>	443
116- V9	f	<p><b>SIR_enable control flag</b></p> <p>SPARC V9 states that the location of the <code>SIR_enable</code> control flag and the means by which it is accessed are implementation dependent. In UltraSPARC Architecture virtual processors, the <code>SIR_enable</code> control flag does not explicitly exist; the <code>SIR</code> instruction always generated an <i>illegal_instruction</i> exception in nonprivileged and privileged modes. <code>SIR</code> only causes a <i>software_initiated_reset</i> trap when executed in hyperprivileged mode.</p>	320
118- V9	f	<p><b>Identifying I/O locations</b></p> <p>The manner in which I/O locations are identified is implementation dependent.</p>	388
119- Ms10	f	<p><b>Unimplemented values for PSTATE.mm</b></p> <p>The effect of an attempt to write an unsupported memory model designation into <code>PSTATE.mm</code> is implementation dependent; however, it should never result in a value of <code>PSTATE.mm</code> value greater than the one that was written. In the case of an UltraSPARC Architecture implementation that only supports the TSO memory model, <code>PSTATE.mm</code> always reads as zero and attempts to write to it are ignored.</p>	96, 397
120- V9	f	<p><b>Coherence and atomicity of memory operations</b></p> <p>The coherence and atomicity of memory operations between virtual processors and I/O DMA memory accesses are implementation dependent.</p>	388
121- V9	f	<p><b>Implementation-dependent memory model</b></p> <p>An implementation may choose to identify certain addresses and use an implementation-dependent memory model for references to them.</p>	388
122- V9	f	<p><b>FLUSH latency</b></p> <p>The latency between the execution of <code>FLUSH</code> on one virtual processor and the point at which the modified instructions have replaced outdated instructions in a multiprocessor is implementation dependent.</p>	186, 404
123- V9	f	<p><b>Input/output (I/O) semantics</b></p> <p>The semantic effect of accessing I/O registers is implementation dependent.</p>	29

**TABLE B-1** SPARC V9 Implementation Dependencies (10 of 10)

Nbr	Category	Description	Page
<b>124-V9</b>	v	<b>Implicit ASI when TL &gt; 0</b> In SPARC V9, when TL > 0, the implicit ASI for instruction fetches, loads, and stores is implementation dependent. In all UltraSPARC Architecture 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.	391
<b>125-V9-Cs10</b>	f	<b>Address masking</b> (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.	97, 97, 162, 239, 300, 472
<b>126-V9-Ms10</b>		<b>Register Windows State registers width</b> Privileged registers CWP, CANSAVE, CANRESTORE, OTHERWIN, and CLEANWIN contain values in the range 0 to N_REG_WINDOWS - 1. An attempt to write a value greater than N_REG_WINDOWS - 1 to any of these registers causes an implementation-dependent value between 0 and N_REG_WINDOWS - 1 (inclusive) to be written to the register. Furthermore, an attempt to write a value greater than N_REG_WINDOWS - 2 violates the register window state definition in <i>Register Window Management Instructions</i> on page 129. Although the width of each of these five registers is architecturally 5 bits, the 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. All five registers should have the same width. For UltraSPARC Architecture 2005 processors, N_REG_WINDOWS = 8. Therefore, each register window state register is implemented with 3 bits, the maximum value for CWP and CLEANWIN is 7, and the maximum value for CANSAVE, CANRESTORE, and OTHERWIN is 6. When these registers are written by the WRPR instruction, bits 63:3 of the data written are ignored.	85
<b>127–199</b>		<i>Reserved.</i>	—

TABLE B-2 provides a list of implementation dependencies that, in addition to those in TABLE B-1, apply to UltraSPARC 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** UltraSPARC Architecture Implementation Dependencies (1 of 12)

Nbr	Description	Page
<b>200–201</b>	<i>Reserved.</i>	—
<b>202-U3</b>	<b>fast_ECC_error</b> trap Whether or not a <i>fast_ECC_error</i> trap exists is implementation dependent. If it does exist, it indicates that an ECC error was detected in an external cache and its trap type is 070 <sub>16</sub> .	491
<b>203-U3-Cs10</b>	<b>Dispatch Control register (DCR) bits 13:6 and 1</b> <i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	

TABLE B-2 UltraSPARC Architecture Implementation Dependencies (2 of 12)

Nbr	Description	Page
204-U3- CS10	<b>DCR bits 5:3 and 0</b> <i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	
205-U3- CS10	<b>Instruction Trap Register</b> <i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	
206-U3- CS10	<b>SHUTDOWN instruction</b> On an UltraSPARC Architecture implementation executing in privileged or hyperprivileged mode, SHUTDOWN behaves like a NOP.	318
207-U3	<b>PCR register bits 47:32, 26:17, and 3</b> The values and semantics of bits 47:32, 26:17, and bit 3 of the PCR register are implementation dependent.	78
208-U3	<b>Ordering of errors captured in instruction execution</b> The order in which errors are captured in instruction execution is implementation dependent. Ordering may be in program order or in order of detection.	—
209-U3	<b>Software intervention after instruction-induced error</b> Precision of the trap to signal an instruction-induced error of which recovery requires software intervention is implementation dependent.	—
210-U3	<b>ERROR output signal</b> The following aspects of the ERROR output signal are implementation dependent in the UltraSPARC Architecture: <ul style="list-style-type: none"> <li>• The causes of the ERROR signal</li> <li>• Whether each of the causes of the ERROR signal, when it generates the ERROR signal, halts the virtual processor or allows the virtual processor to continue running</li> <li>• The exact semantics of the ERROR signal</li> </ul>	—
211-U3	<b>Error logging registers' information</b> 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	<b>AFSR .priv</b> 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.	—
214-U3	<b>Enable/disable control for deferred traps</b> Whether an implementation provides an enable/disable control feature for deferred traps is implementation dependent.	—
215-U3	<b>Error barrier</b> DONE and RETRY instructions may implicitly provide an error barrier function as MEMBAR #Sync. Whether DONE and RETRY instructions provide an error barrier is implementation dependent.	—

TABLE B-2 UltraSPARC Architecture Implementation Dependencies (3 of 12)

Nbr	Description	Page
216-U3	<b>data_access_error trap precision</b> The precision of a <i>data_access_error</i> trap is implementation dependent.	—
217-U3	<b>instruction_access_error trap precision</b> The precision of an <i>instruction_access_error</i> trap is implementation dependent.	—
218-U3- Cs10	<b>async_data_error</b> The <i>async_data_error</i> exception has been superseded by <i>sw_recoverable_error</i> , so this implementation dependency no longer applies.	—
219-U3	<b>Asynchronous Fault Address register (AFAR) allocation</b> Allocation of Asynchronous Fault Address register (AFAR) is implementation dependent. There may be one instance or multiple instances of AFAR. Although the ASI for AFAR is defined as 4D <sub>16</sub> , the virtual address of AFAR if there are multiple AFARs is implementation dependent.	—
220-U3	<b>Addition of logging and control registers for error handling</b> Whether the implementation supports additional logging and control registers for error handling is implementation dependent.	—
221-U3	<b>Special/signalling ECCs</b> The method to generate “special” or “signalling” ECCs and whether a processor ID is embedded into the data associated with special/signalling ECCs is implementation dependent.	—
223-U3	<b>TLB multiple-hit detection</b> Whether TLB multiple-hit detection is supported in an UltraSPARC Architecture implementation is implementation dependent.	—
225-U3	<b>TLB locking of entries</b> The mechanism by which entries in TLB are locked is implementation dependent in UltraSPARC Architecture implementations.	—
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> <b>TSB Base address generation</b> Whether the implementation generates the TSB Base address by <b>exclusive-OR</b> ing the TSB Base register and a TSB register or by taking the <i>tsb_base</i> field directly from a TSB register is implementation dependent in UltraSPARC 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	<b>data_access_exception trap</b> The causes of a <i>data_access_exception</i> trap are implementation dependent in UltraSPARC Architecture 2005.	—
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>	—

TABLE B-2 UltraSPARC Architecture Implementation Dependencies (4 of 12)

Nbr	Description	Page
<b>235-U3-Cs10</b>	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
<b>236-U3-Cs10</b>	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
<b>237-U3</b>	<b>JMPL/RETURN <i>mem_address_not_aligned</i></b> Whether the fault status and/or address (D-SFSR/D-SFAR) are is captured when a <i>mem_address_not_aligned</i> trap occurs during a JMPL or RETURN instruction is implementation dependent.	—
<b>239-U3-Cs10</b>	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
<b>240-U3-Cs10</b>	<i>Reserved.</i>	—
<b>241-U3</b>	<b>Address Masking and DSFAR</b> When PSTATE.am = 1 and an exception occurs, the value written to the more-significant 32 bits of the Data Synchronous Fault Address Register (DSFAR) is implementation dependent.	97, 97
<b>243-U3</b>	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
<b>244-U3-Cs10</b>	<b>Data Watchpoint Reliability</b> Data Watchpoint traps are completely implementation-dependent in UltraSPARC Architecture processors.	—
<b>245-U3-Cs10</b>	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
<b>246-U3</b>	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
<b>247-U3</b>	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
<b>248-U3</b>	<b>Conditions for <i>fp_exception_other</i> with unfinished_FPop</b> 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.	65
<b>249-U3-Cs10</b>	<b>Data Watchpoint for Partial Store Instruction</b> 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 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.	343

TABLE B-2 UltraSPARC Architecture Implementation Dependencies (5 of 12)

Nbr	Description	Page
<b>250-U3-Cs10</b>	<b>PCR accessibility when PSTATE.priv = 0</b> In an UltraSPARC Architecture implementation, PCR is never accessible to nonprivileged software. Specifically, when a virtual processor is operating in nonprivileged mode (PSTATE.priv = 0 and HPSTATE.HPRIV = 0), an attempt to access PCR (using an RDPCR or a WRPCR instruction) results in a <i>privileged_opcode</i> exception.	78, 301, 371
<b>251</b>	<i>Reserved.</i>	
<b>252-U3-Cs10</b>	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
<b>253-U3-Cs10</b>	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
<b>254-U3-Cs10</b>	<b>Means of exiting error_state</b> A virtual processor, upon entering <i>error_state</i> , automatically generates a <i>watchdog_reset</i> (WDR).	444, 450, 477, 490, 548
<b>257-U3</b>	<b>LDDFA with ASI C0<sub>16</sub>–C5<sub>16</sub> or C8<sub>16</sub>–CD<sub>16</sub> and misaligned memory address</b> If an LDDFA opcode is used with an ASI of C0 <sub>16</sub> –C5 <sub>16</sub> or C8 <sub>16</sub> –CD <sub>16</sub> (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>data_access_exception</i> , <i>mem_address_not_aligned</i> , or <i>LDDF_mem_address_not_aligned</i> .	254
<b>258-U3-Cs10</b>	<b>ASI_SERIAL_ID</b> (This register is not defined in the UltraSPARC Architecture, so this implementation dependency does not apply to UltraSPARC Architecture 2005.)	—
<b>259–299</b>	<i>Reserved.</i>	—
<b>300-U4-Cs10</b>	<b>Attempted access to ASI registers with LDTWA</b> If an LDTWA instruction referencing a non-memory ASI is executed, it generates a <i>data_access_exception</i> exception.	270
<b>301-U4-Cs10</b>	<b>Attempted access to ASI registers with STTWA</b> If an STTWA instruction referencing a non-memory ASI is executed, it generates a <i>data_access_exception</i> exception.	349
<b>302-U4-Cs10</b>	<b>Scratchpad registers</b> An UltraSPARC Architecture processor includes eight privileged Scratchpad registers (64 bits each, read/write accessible).	431
<b>303-U4-CS10</b>	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
<b>304-U4-Cs10</b>	<b>XIR</b> XIR affects only the virtual processors identified in the XIR_STEERING register (not a whole system).	547
<b>305-U4-Cs10</b>	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—

**TABLE B-2** UltraSPARC Architecture Implementation Dependencies (6 of 12)

Nbr	Description	Page
<b>306-U4-Cs10</b>	<b>Trap type generated upon attempted access to noncacheable page with LDTXA</b> When an LDTXA instruction attempts access from an address that is not mapped to cacheable memory space, a <i>data_access_exception</i> exception is generated.	264
<b>307-U4-Cs10</b>	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
<b>308-U3-Cs10</b>	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
<b>309-U4-Cs10</b>	<i>Reserved.</i>	—
<b>311–319</b>	<i>Reserved.</i>	
	<b>Strand Interrupt ID register</b>	521
	Whether any portion of the <i>int_id</i> field of the Strand Interrupt ID register is read-only is implementation dependent.	
<b>321-U4</b>	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	
<b>322-U4</b>	<b>Power used by CMT</b>	523
	Whether disabling a virtual processor reduces the power used by a CMT processor is implementation dependent.	
<b>323-U4</b>	<b>Updating Strand Enable Register</b>	525
	Whether an implementation provides a restriction that prevents software from writing a value of all zeroes (or zeroes corresponding to all available virtual processors) to the STRAND_ENABLE register is implementation dependent. This restriction avoids the dangerous case where all virtual processors become disabled and the only way to enable any virtual processor is a hard <i>power_on_reset</i> (a warm reset would not suffice). If such a restriction is implemented and software running on any virtual processor attempts to write a value of all zeroes (or zeroes corresponding to all available virtual processors) to the STRAND_ENABLE register, hardware forces the STRAND_ENABLE register to an implementation-dependent value which enables at least one of the available virtual processors.	
<b>324-U4</b>	<b>Parking a virtual processor</b>	527
	Whether parking a virtual processor reduces the power used by a CMT processor is implementation dependent.	
<b>325-U4</b>	<b>XIR Steering register (XIR Reset)</b>	
	<b>a:</b> Whether XIR_STEERING{n} is a read-only bit or a read/write bit is implementation dependent. If XIR_STEERING{n} is read-only, then (1) writes to XIR_STEERING{n} are ignored and (2) XIR_STEERING{n} is set to 1 if virtual processor n is available and to 0 if it is not available (that is, XIR_STEERING{n} reads the same as STRAND_AVAILABLE{n}).	534
	<b>b:</b> If XIR_STEERING{n} is read/write, upon de-assertion of reset the value of STRAND_AVAILABLE{n} is copied to XIR_STEERING{n} for all UltraSPARC Architecture implementations.	535, 552
<b>326-U4-Cs10</b>	<i>This implementation dependency no longer applies, as of UltraSPARC Architecture 2005.</i>	—
<b>327–399</b>	<i>Reserved</i>	

TABLE B-2 UltraSPARC Architecture Implementation Dependencies (7 of 12)

Nbr	Description	Page
<b>400-S10</b>	<b>Global Level register (GL) implementation</b> 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>MAXGL</i> for that implementation. If any bits of GL are not implemented, they read as zero and writes to them are ignored.	102
<b>401-S10</b>	<b>Maximum Global Level (<i>MAXPGL</i>, <i>MAXGL</i>)</b> 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 visible to privileged software). In a typical implementation <i>MAXPGL</i> = <i>MAXPTL</i> (see impl. dep. #101-V9-CS10). The architectural parameter <i>MAXGL</i> is a constant for each implementation; its legal values are from 4 to 15 (supporting from 5 to 16 sets of global registers). Architecturally, <i>MAXPTL</i> must be $\geq 2$ and <i>MAXGL</i> must be $> \text{MAXPGL}$ .	99, 101
<b>402-S10</b>	<b>Priority of <i>internal_processor_error</i></b> The trap priority of the <i>internal_processor_error</i> exception is implementation dependent. Furthermore, its priority may vary within an implementation, based on the cause of the error being reported.	464, 468, 487
<b>403-S10</b>	<b>Setting of “dirty” bits in FPRS</b> 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. The specific conditions under which a dirty bit is set are implementation dependent.	77, 77
<b>404-S10</b>	<b>Privileged Scratchpad registers 4 through 7</b> 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(emulated by trap to hyperprivileged software), or (3) inaccessible (cause a <i>data_access_exception</i> exception).	431
<b>405-S10</b>	<b>Virtual address range</b> An UltraSPARC Architecture implementation may support a full 64-bit virtual address space or a more limited range of virtual addresses. In an implementation that does 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 $2^{64} - 1$ .	28
<b>406-S10</b>	<b>HTBA high-order bits</b> It is implementation dependent whether all 50 bits of HTBA{63:14} are implemented or if only bits <i>n</i> -1:0 are implemented. If the latter, writes to bits 63: <i>n</i> are ignored and when HTBA is read, bits 63: <i>n</i> read as sign-extended copies of the most significant implemented bit, HTBA{ <i>n</i> - 1}.	107
<b>407-S10</b>	<b>Hyperprivileged Scratchpad register aliasing</b> It is implementation dependent whether any of the hyperprivileged Scratchpad registers are aliased to the corresponding privileged Scratchpad register or is an independent register.	432
<b>408-S10</b>	<b>HPSTATE bit 11</b> The contents and semantics of HPSTATE{11} are implementation dependent.	104



TABLE B-2 UltraSPARC Architecture Implementation Dependencies (9 of 12)

Nbr	Description	Page
411-S10	<p><b>Block Store behavior</b></p> <p>The following aspects of the behavior of block store (STBLOCKF) instructions are implementation dependent:</p> <ul style="list-style-type: none"> <li>• The memory ordering model that STBLOCKF follows (other than as constrained by the rules outlined on page 330).</li> <li>• Whether <i>VA_watchpoint</i> exceptions are recognized on accesses to all 64 bytes of a STBLOCKF (the recommended behavior), or only on accesses to the first eight bytes.</li> <li>• Whether STBLOCKFs to non-cacheable (TTE.cp = 0) pages execute in strict program order or not. If not, a STBLOCKF to a non-cacheable page causes an <i>illegal_instruction</i> exception.</li> <li>• Whether STBLOCKF follows register dependency interlocks (as ordinary stores do).</li> <li>• Whether a non-Commit STBLOCKF forces the data to be written to memory and invalidates copies in all caches present (as the Commit variants of STBLOCKF do).</li> <li>• Whether the MMU ignores the side-effect bit (TTE.e) for STBLOCKF accesses (in which case, STBLOCKFs behave as if TTE.e = 0)</li> <li>• Any other restrictions on the behavior of STBLOCKF, as described in implementation-specific documentation.</li> </ul>	330, 330
412-S10	<p><b>MEMBAR behavior</b></p> <p>An UltraSPARC Architecture implementation may define the operation of each MEMBAR variant in any manner that provides the required semantics.</p>	273
413-S10	<p><b>Load Twin Extended Word behavior</b></p> <p>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>	264
414	<i>Reserved.</i>	—
417-S10	<p><b>Behavior of DONE and RETRY when TSTATE[TL].pstate.am = 1</b></p> <p>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>	98, 167309
418	— <i>unused</i> —	,
419-S10	<p><b>Contents of TPC[TL], TNPC[TL], TSTATE[TL], and HTSTATE[TL] after a Warm Reset (WMR)</b></p> <p>It is implementation dependent whether, after a Warm Reset (WMR), the contents of TPC[TL], TNPC[TL], TSTATE[TL], and HTSTATE[TL] are unchanged from their values before the WMR, or are contain the same values saved as during a WDR, XIR, or SIR reset. (The latter implementation is the preferred one.)</p>	551

**TABLE B-2** UltraSPARC Architecture Implementation Dependencies (10 of 12)

Nbr	Description	Page
<b>420-S10</b>	<p><b>Implementation Dependent Aspects of a Warm Reset (WMR)</b>                      The following aspects of Warm Reset (WMR) are implementation dependent:                      (a) by what means WMR can be applied (for example, write to reset register or assertion/deassertion of an input pin)                      (b) the extent to which a processor is reset by WMR (for example, single physical core, entire processor (chip), and how the on-chip memory system is affected),                      (c) by what means hyperprivileged software can distinguish between WMR and POR resets</p>	547
<b>421-S10</b>	<p><b>Interrupt Queue Head and Tail Register Contents</b>                      It is implementation dependent whether interrupt queue head and tail registers (a) are datatype-agnostic “scratch registers” used for communication between privileged and hyperprivileged software, in which case their contents are defined purely by software convention, or (b) are maintained to some degree by virtual processor hardware, imposing a fixed meaning on their contents.</p>	498
<b>422-S10</b>	<p><b>Interrupt Queue Tail Register Writability</b>                      It is implementation dependent whether tail registers are writable in privileged mode. If a tail register is read-only in privileged mode, an attempt to write to it causes a <i>data_access_exception</i> exception. If a tail register is writable in privileged mode, an attempt to write to it results in undefined behavior.</p>	499, 499
<b>423-S10</b>	<p><b>Performance Impact of Disabling a Virtual Processor</b>                      Whether disabling a virtual processor increases the performance of other virtual processors in the CMT is implementation dependent.</p>	523
<b>424-S10</b>	<p><b>Ability to Dynamically Enable/Disable a Virtual Processor</b>                      Whether a CMT implementation provides the ability to dynamically enable and disable virtual processors is implementation dependent. It is tightly coupled to the underlying microarchitecture of a specific CMT implementation. This feature is implementation dependent because any implementation-independent interface would be too inefficient on some implementations.</p>	526
<b>425-S10</b>	<p><b>TICK Register Counting While a Virtual Processor is Parked</b>                      It is implementation dependent whether the TICK register continues to count while a virtual processor is parked.</p>	526
<b>426-S10</b>	<p><b>Performance Impact of Parking a Virtual Processor</b>                      The degree to which parking a virtual processor impacts the performance of other virtual processors is implementation dependent.</p>	527
<b>427-S10</b>	<p><b>Latency to Park or Unpark a Virtual Processor</b>                      There may be an arbitrarily long, but bounded, delay (“skid”) from the time when a virtual processor is directed to park or unpark (via an update to the STRAND_RUNNING register) until the corresponding virtual processor(s) actually park or unpark.</p>	527, 530
<b>428-S10</b>	<p><b>Method by Which Self-Parking is Assured</b>                      When a virtual processor writes to the STRAND_RUNNING register to park itself, the method by which completion of parking is assured (instructions stop being issued) is implementation dependent.</p>	528

**TABLE B-2** UltraSPARC Architecture Implementation Dependencies (11 of 12)

Nbr	Description	Page
<b>429-S10</b>	<p><b>Which Virtual Processor is Automatically Unparked</b></p> <p>If an update to the STRAND_RUNNING register would cause all enabled virtual processors to become parked, it is implementation dependent which virtual processor is automatically unparked by hardware. The preferred implementation is that when an update to the STRAND_RUNNING register (STXA instruction) would cause all virtual processors to become parked, hardware silently ignores (discards) that STXA instruction.</p>	529
<b>430-S10</b>	<p><b>Parking All But One Virtual Processor in a Multiprocessor Configuration</b></p> <p>In a multi'io configuration, whether all but one virtual processor can be parked is implementation dependent.</p>	530
<b>431-S10</b>	<p><b>Criteria for Completion of Park/Unpark</b></p> <p>The criteria used for determining whether a virtual processor is fully parked (corresponding bit set to '1' in the STRAND_RUNNING_STATUS register) are implementation dependent.</p>	531
<b>432-S10</b>	<p><b>Standby/Wait state</b></p> <p>Whether an implementation implements a Standby (or Wait) state for virtual processors, how that state is controlled, and how that state is observed are implementation-dependent.</p>	531
<b>433-S10</b>	<p><b>Partial-Processor Reset Subsetting Mechanism</b></p> <p>A mechanism must exist to specify which subset of virtual processors in a processor should be reset when a partial-processor reset (for example, XIR) occurs. The specific mechanism is implementation-dependent.</p>	533
<b>434-S10</b>	<p><b>Error Steering Register(s)</b></p> <p>Because of the range of implementation, the number of, organization of, and ASI assignments for error steering registers in a CMT processor are implementation dependent.</p>	536
<b>435-S10</b>	<p><b>Error Steering Register Alternatives</b></p> <p>Although the ERROR_STEERING register is the recommended mechanism for steering non-virtual-processor-specific errors to a virtual processor for handling, the actual mechanism used in a given implementation is implementation dependent.</p>	538
<b>436-S10</b>	<p><b>Error Steering Register</b></p> <p>The width of the target_id field of the ERROR_STEERING register is implementation dependent.</p>	539
<b>437-S10</b>	<p><b>Error Steering Register targetid Field Plurality</b></p> <p>An implementation may provide multiple target_id fields in an ERROR_STEERING register for different types of non-virtual-processor-specific errors.</p>	539
<b>438-S10</b>	<p><b>Non-Virtual Processor-Specific Errors in Shared Resources</b></p> <p>It is implementation dependent whether the error-reporting structures for errors in shared resources appear within a virtual processor in per-virtual-processor registers or are contained within shared registers associated with the shared structures in which the errors may occur.</p>	539

**TABLE B-2** UltraSPARC Architecture Implementation Dependencies (12 of 12)

<b>Nbr</b>	<b>Description</b>	<b>Page</b>
<b>439-S10</b>	<b>Exception Generated for Each Non-Virtual Processor-Specific Error</b> The type of exception generated in a virtual processor to handle each type of non-virtual-processor-specific error is implementation dependent. A virtual processor can choose to use the same exceptions used for corresponding virtual-processor-specific asynchronous errors or it can choose to generate different exceptions.	<b>539</b>
<b>440-S10</b>	<b>Which Virtual Processor Unparked During Power-on-Reset (POR)</b> Which virtual processor is unparked during POR and whether it is unparked by processor hardware or by a service processor is implementation dependent. Conventionally, the virtual processor with the lowest-numbered <code>strand_id</code> is unparked	<b>541</b>
<b>441–449</b>	<i>Reserved for UltraSPARC Architecture 2005</i>	
<b>450-499</b>	<i>Reserved for UltraSPARC Architecture 2006</i>	
<b>451 and up</b>	<i>Reserved.</i>	



**Note: This chapter is undergoing final review; please check back later for a copy of UltraSPARC Architecture 2005 containing the final version of this chapter.**

## 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 UltraSPARC Architecture assemblers for the convenience of assembly language programmers.

The appendix contains these sections:

- **Notation Used** on page 591.
- **Syntax Design** on page 598.
- **Synthetic Instructions** on page 598.

---

### 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 metasymbols further identify the placement of the operand in the generated binary instruction. For example, *reg<sub>rs2</sub>* 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:<sup>1</sup>

- %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:

- reg<sub>rs1</sub>* (rs1 field)
- reg<sub>rs2</sub>* (rs2 field)
- reg<sub>rd</sub>* (rd field)

**freg.** An *freg* is a floating-point register name. It may have the following values:

- %f0, %f1, %f2–%f63*

See *Floating-Point Registers* on page 55.

Subscripts further identify the placement of the operand in the binary instruction as one of the following:

- freg<sub>rs1</sub>* (rs1 field)
- freg<sub>rs2</sub>* (rs2 field)
- freg<sub>rs3</sub>* (rs3 field)
- freg<sub>rd</sub>* (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<sub>rs1</sub>* (rs1 field)
- asr\_reg<sub>rd</sub>* (rd field)

<sup>1</sup> In actual usage, the *%sp*, *%fp*, *%gn*, *%on*, *%ln*, and *%in* forms are preferred over *%rn*.

*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:

<code>%asi</code>	Address Space Identifier (ASI) register
<code>%canrestore</code>	Restorable Windows register
<code>%cansave</code>	Savable Windows register
<code>%ccr</code>	Condition Codes register
<code>%cleanwin</code>	Clean Windows register
<code>%cwp</code>	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>%hintp</code>	Hyperprivileged Interrupt Pending (HINTP) register
<code>%hpstate</code>	Hyperprivileged State (HSTATE) register
<code>%hstick</code>	Hyperprivileged System Timer (HSTICK) register
<code>%hstick_cmpr</code>	Hyperprivileged System Tick Compare (HSTICK_CMPR) register
<code>%htba</code>	Hyperprivileged Trap Base Address (HTBA) register
<code>%htstate</code>	Hyperprivileged Trap State (HTSTATE) register
<code>%hver</code>	Hyperprivileged Version (HVER) register
<code>%otherwin</code>	Other Windows (OTHERWIN) register
<code>%pc</code>	Program Counter (PC) register
<code>%pcr</code>	Performance Control Register (PCR)
<code>%pic</code>	Performance Instrumentation Counters
<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>%sys_tick</code>	System Timer (STICK) register
<code>%sys_tick_cmpsr</code>	System Timer Compare (STICK_CMPR) register
<code>%tba</code>	Trap Base Address (TBA) register
<code>%tick</code>	Cycle count (TICK) register
<code>%tick_cmpsr</code>	Timer Compare (TICK_CMPR) 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 following special symbol names are unary operators that perform the functions described:

<code>%uhi</code>	Extracts bits 63:42 (high 22 bits of upper word) of its operand
<code>%ulo</code> or <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

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 "fcn" 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 <sub>16</sub> )	

**TABLE C-1** Value Names and Values (2 of 2)

Value Name in Assembly Language	Value	Comments
#n_reads_strong	20 (14 <sub>16</sub> )	
#one_read_strong	21 (15 <sub>16</sub> )	
#n_writes_strong	22 (16 <sub>16</sub> )	
#one_write_strong	23 (17 <sub>16</sub> )	
<b>for MEMBAR instruction “mmask” field</b>		
#LoadLoad	01 <sub>16</sub>	
#StoreLoad	02 <sub>16</sub>	
#LoadStore	04 <sub>16</sub>	
<b>for MEMBAR instruction “cmask” field</b>		
#StoreStore	08 <sub>16</sub>	
#Lookaside	10 <sub>16</sub>	
#MemIssue	20 <sub>16</sub>	
#Sync	40 <sub>16</sub>	

## 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:

*reg<sub>rs1</sub>* (equivalent to *reg<sub>rs1</sub>* + %g0)  
*reg<sub>rs1</sub>* + *simm13*  
*reg<sub>rs1</sub>* - *simm13*  
*simm13* (equivalent to %g0 + *simm13*)  
*simm13* + *reg<sub>rs1</sub>* (equivalent to *reg<sub>rs1</sub>* + *simm13*)

*address* Can be any of the following:

*reg<sub>rs1</sub>* (equivalent to *reg<sub>rs1</sub>* + %g0)  
*reg<sub>rs1</sub>* + *simm13*  
*reg<sub>rs1</sub>* - *simm13*  
*simm13* (equivalent to %g0 + *simm13*)  
*simm13* + *reg<sub>rs1</sub>* (equivalent to *reg<sub>rs1</sub>* + *simm13*)  
*reg<sub>rs1</sub>* + *reg<sub>rs2</sub>*

*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, #MemIssue, #Sync, #StoreStore, #LoadStore, #StoreLoad, and #LoadLoad.

*prefetch\_fcn* (*prefetch function*) Can be any of the following:

#n\_reads  
#one\_read  
#n\_writes  
#one\_write  
#page

*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:

$reg_{rs2}$   
 $simm13$

*reg\_or\_imm10* (*register or immediate value*) Can be either of:

$reg_{rs2}$   
 $simm10$

*reg\_or\_imm11* (*register or immediate value*) Can be either of:

$reg_{rs2}$   
 $simm11$

*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 (in a Load, Store, CASA, CASXA, LDSTUB[A], or SWAP[A] 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.

---

## 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<sub>rs1</sub></i> , <i>reg_or_imm</i>	subcc <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %g0	Compare.
jmp <i>address</i>	jmp1 <i>address</i> , %g0	
call <i>address</i>	jmp1 <i>address</i> , %o7	
iprefetch <i>label</i>	bn, a, pt %xcc, <i>label</i>	Instruction prefetch.
tst <i>reg<sub>rs1</sub></i>	orcc    %g0, <i>reg<sub>rs1</sub></i> , %g0	Test.
ret	jmp1    %i7+8, %g0	Return from subroutine.
retl	jmp1    %o7+8, %g0	Return from leaf subroutine.

**TABLE C-2** Mapping Synthetic to SPARC V9 Instructions (2 of 3)

Synthetic Instruction	SPARC V9 Instruction(s)	Comment	
restore	restore %g0, %g0, %g0	Trivial RESTORE.	
save	save %g0, %g0, %g0	Trivial SAVE. <b>(Warning:</b> trivial SAVE should only be used in kernel code!)	
setuw	<i>value, reg<sub>rd</sub></i>	sethi %hi( <i>value</i> ), <i>reg<sub>rd</sub></i> — or —	(When (( <i>value</i> &3FF <sub>16</sub> ) == 0).)
		or %g0, <i>value</i> , <i>reg<sub>rd</sub></i> — or —	(When 0 ≤ <i>value</i> ≤ 4095).
		sethi %hi( <i>value</i> ), <i>reg<sub>rd</sub></i> ;	(Otherwise)
		or <i>reg<sub>rd</sub></i> , %lo( <i>value</i> ), <i>reg<sub>rd</sub></i>	Warning: do not use setuw in the delay slot of a DCTI.
set	<i>value, reg<sub>rd</sub></i>	synonym for setuw.	
setsw	<i>value, reg<sub>rd</sub></i>	sethi %hi( <i>value</i> ), <i>reg<sub>rd</sub></i> — or —	(When ( <i>value</i> >= 0) and (( <i>value</i> & 3FF <sub>16</sub> ) == 0).)
		or %g0, <i>value</i> , <i>reg<sub>rd</sub></i> — or —	(When 4096 ≤ <i>value</i> ≤ 4095).
		sethi %hi( <i>value</i> ), <i>reg<sub>rd</sub></i>	(Otherwise, if ( <i>value</i> < 0) and (( <i>value</i> & 3FF <sub>16</sub> ) = = 0))
		sra <i>reg<sub>rd</sub></i> , %g0, <i>reg<sub>rd</sub></i> — or —	
		sethi %hi( <i>value</i> ), <i>reg<sub>rd</sub></i> ;	(Otherwise, if <i>value</i> 0)
		or <i>reg<sub>rd</sub></i> , %lo( <i>value</i> ), <i>reg<sub>rd</sub></i> — or —	
		sethi %hi( <i>value</i> ), <i>reg<sub>rd</sub></i> ;	(Otherwise, if <i>value</i> < 0)
		or <i>reg<sub>rd</sub></i> , %lo( <i>value</i> ), <i>reg<sub>rd</sub></i>	
		sra <i>reg<sub>rd</sub></i> , %g0, <i>reg<sub>rd</sub></i>	Warning: do not use setsw in the delay slot of a CTI.
			Create 64-bit constant.
setx	<i>value, reg, reg<sub>rd</sub></i>	sethi %uhi( <i>value</i> ), <i>reg</i> or <i>reg</i> , %ulo( <i>value</i> ), <i>reg</i>	(“ <i>reg</i> ” is used as a temporary register.)
		sllx <i>reg</i> , 32, <i>reg</i>	
		sethi %hi( <i>value</i> ), <i>reg<sub>rd</sub></i> or <i>reg<sub>rd</sub></i> , <i>reg</i> , <i>reg<sub>rd</sub></i>	Note: setx optimizations are possible but not enumerated here. The worst case is shown.
		or <i>reg<sub>rd</sub></i> , %lo( <i>value</i> ), <i>reg<sub>rd</sub></i>	Warning: do not use setx in the delay slot of a CTI.
signx	<i>reg<sub>rs1</sub>, reg<sub>rd</sub></i>	sra <i>reg<sub>rs1</sub></i> , %g0, <i>reg<sub>rd</sub></i>	Sign-extend 32-bit value to 64 bits.
signx	<i>reg<sub>rd</sub></i>	sra <i>reg<sub>rd</sub></i> , %g0, <i>reg<sub>rd</sub></i>	
not	<i>reg<sub>rs1</sub>, reg<sub>rd</sub></i>	xnor <i>reg<sub>rs1</sub></i> , %g0, <i>reg<sub>rd</sub></i>	One’s complement.
not	<i>reg<sub>rd</sub></i>	xnor <i>reg<sub>rd</sub></i> , %g0, <i>reg<sub>rd</sub></i>	One’s complement.

**TABLE C-2** Mapping Synthetic to SPARC V9 Instructions (3 of 3)

Synthetic Instruction		SPARC V9 Instruction(s)		Comment
neg	<i>reg<sub>rs2</sub>, reg<sub>rd</sub></i>	sub	<i>%g0, reg<sub>rs2</sub>, reg<sub>rd</sub></i>	Two's complement.
neg	<i>reg<sub>rd</sub></i>	sub	<i>%g0, reg<sub>rd</sub>, reg<sub>rd</sub></i>	Two's complement.
cas	<i>[reg<sub>rs1</sub>], reg<sub>rs2</sub>, reg<sub>rd</sub></i>	cas	<i>[reg<sub>rs1</sub>]#ASI_P, reg<sub>rs2</sub>, reg<sub>rd</sub></i>	Compare and swap.
casl	<i>[reg<sub>rs1</sub>], reg<sub>rs2</sub>, reg<sub>rd</sub></i>	cas	<i>[reg<sub>rs1</sub>]#ASI_P_L, reg<sub>rs2</sub>, reg<sub>rd</sub></i>	Compare and swap, little-endian.
casx	<i>[reg<sub>rs1</sub>], reg<sub>rs2</sub>, reg<sub>rd</sub></i>	casxa	<i>[reg<sub>rs1</sub>]#ASI_P, reg<sub>rs2</sub>, reg<sub>rd</sub></i>	Compare and swap extended.
casxl	<i>[reg<sub>rs1</sub>], reg<sub>rs2</sub>, reg<sub>rd</sub></i>	casxa	<i>[reg<sub>rs1</sub>]#ASI_P_L, reg<sub>rs2</sub>, reg<sub>rd</sub></i>	Compare and swap extended, little-endian.
inc	<i>reg<sub>rd</sub></i>	add	<i>reg<sub>rd</sub>, 1, reg<sub>rd</sub></i>	Increment by 1.
inc	<i>const13, reg<sub>rd</sub></i>	add	<i>reg<sub>rd</sub>, const13, reg<sub>rd</sub></i>	Increment by <i>const13</i> .
inccc	<i>reg<sub>rd</sub></i>	addcc	<i>reg<sub>rd</sub>, 1, reg<sub>rd</sub></i>	Increment by 1; set icc & xcc.
inccc	<i>const13, reg<sub>rd</sub></i>	addcc	<i>reg<sub>rd</sub>, const13, reg<sub>rd</sub></i>	Incr by <i>const13</i> ; set icc & xcc.
dec	<i>reg<sub>rd</sub></i>	sub	<i>reg<sub>rd</sub>, 1, reg<sub>rd</sub></i>	Decrement by 1.
dec	<i>const13, reg<sub>rd</sub></i>	sub	<i>reg<sub>rd</sub>, const13, reg<sub>rd</sub></i>	Decrement by <i>const13</i> .
decc	<i>reg<sub>rd</sub></i>	subcc	<i>reg<sub>rd</sub>, 1, reg<sub>rd</sub></i>	Decrement by 1; set icc & xcc.
deccc	<i>const13, reg<sub>rd</sub></i>	subcc	<i>reg<sub>rd</sub>, const13, reg<sub>rd</sub></i>	Decr by <i>const13</i> ; set icc & xcc.
btst	<i>reg_or_imm, reg<sub>rs1</sub></i>	andcc	<i>reg<sub>rs1</sub>, reg_or_imm, %g0</i>	Bit test.
bset	<i>reg_or_imm, reg<sub>rd</sub></i>	or	<i>reg<sub>rd</sub>, reg_or_imm, reg<sub>rd</sub></i>	Bit set.
bclr	<i>reg_or_imm, reg<sub>rd</sub></i>	andn	<i>reg<sub>rd</sub>, reg_or_imm, reg<sub>rd</sub></i>	Bit clear.
btog	<i>reg_or_imm, reg<sub>rd</sub></i>	xor	<i>reg<sub>rd</sub>, reg_or_imm, reg<sub>rd</sub></i>	Bit toggle.
clr	<i>reg<sub>rd</sub></i>	or	<i>%g0, %g0, reg<sub>rd</sub></i>	Clear (zero) register.
clrb	<i>[address]</i>	stb	<i>%g0, [address]</i>	Clear byte.
clrh	<i>[address]</i>	sth	<i>%g0, [address]</i>	Clear half-word.
clr	<i>[address]</i>	stw	<i>%g0, [address]</i>	Clear word.
clrx	<i>[address]</i>	stx	<i>%g0, [address]</i>	Clear extended word.
clruw	<i>reg<sub>rs1</sub>, reg<sub>rd</sub></i>	srl	<i>reg<sub>rs1</sub>, %g0, reg<sub>rd</sub></i>	Copy and clear upper word.
clruw	<i>reg<sub>rd</sub></i>	srl	<i>reg<sub>rd</sub>, %g0, reg<sub>rd</sub></i>	Clear upper word.
mov	<i>reg_or_imm, reg<sub>rd</sub></i>	or	<i>%g0, reg_or_imm, reg<sub>rd</sub></i>	
mov	<i>%y, reg<sub>rd</sub></i>	rd	<i>%y, reg<sub>rd</sub></i>	
mov	<i>%asrn, reg<sub>rd</sub></i>	rd	<i>%asrn, reg<sub>rd</sub></i>	
mov	<i>reg_or_imm, %y</i>	wr	<i>%g0, reg_or_imm, %y</i>	
mov	<i>reg_or_imm, %asrn</i>	wr	<i>%g0, reg_or_imm, %asrn</i>	

# Index

---

## A

a (annul) instruction field

branch instructions, 154, 155, 157, 160, 175, 177

accesses

cacheable, 387

I/O, 387

restricted ASI, 391

with side effects, 387, 398

accrued exception (aexc) field of FSR register, 66, 452, 570

ADD instruction, 146

ADDC instruction, 146

ADDcc instruction, 146, 322

ADDccc instruction, 146

address

aliasing, 505

operand syntax, 596

space identifier (ASI), 407

address mask (am) field of PSTATE register  
description, 96

address space, 22

address space identifier (ASI), 386

appended to memory address, 27, 112

architecturally specified, 391

bypass, 8, 408

bypassing, 97

changed in, 433

changed in UA

ASI\_LD\_TWINK\_NUCLEUS\_LITTLE, 433

ASI\_LD\_TWINK\_REAL\_LITTLE, 433

ASI\_LDTX\_N, 433

ASI\_LDTX\_NL, 433

ASI\_LDTX\_REAL, 433

ASI\_LDTX\_REAL\_L, 433

ASI\_QUAD\_LDD\_REAL, 433

ASI\_REAL, 433

ASI\_REAL\_IO, 433

ASI\_REAL\_IO\_LITTLE, 433

ASI\_REAL\_LITTLE, 433

definition, 7

encoding address space information, 113

explicit, 120

explicitly specified in instruction, 120

implicit, *See* implicit ASIs

nontranslating, 13, 270, 349

nontranslating ASIs, 408

with prefetch instructions, 292

restricted, 391, 407

hyperprivileged, 392

privileged, 392

restriction indicator, 74

SPARC V9 address, 389

translating ASIs, 408

unrestricted, 392, 407

address space identifier (ASI) register

for load/store alternate instructions, 74

address for explicit ASI, 120

and LDDA instruction, 252, 268

and LDSTUBA instruction, 261

load integer from alternate space  
instructions, 242

with prefetch instructions, 292

for register-immediate addressing, 392

restoring saved state, 166, 309

saving state, 437

and STDA instruction, 348

store floating-point into alternate space

- instructions, 335
- store integer to alternate space instructions, 324
- and SWAPA instruction, 353
- after trap, 32
- and TSTATE register, 92
- and write state register instructions, 370
- addressing modes, 22
- ADDX instruction (SPARC V8), 146
- ADDXcc instruction (SPARC V8), 146
- AFAR, *See* Asynchronous Fault Address register (AFAR)
- AFSR, *See* Asynchronous Fault Status register (AFSR)
- alias
  - floating-point registers, 55
- aliased, 7
- ALIGNADDRESS instruction, 147
- ALIGNADDRESS\_LITTLE instruction, 147
- alignment
  - data (load/store), 28, **114**, 389
  - doubleword, 28, **114**, 389
  - extended-word, **114**
  - halfword, 28, **114**, 389
  - instructions, 28, **114**, 389
  - integer registers, 267, 269
  - memory, 389, 487
  - quadword, 28, **114**, 389
  - word, 28, **114**, 389
- ALLCLEAN instruction, **148**
- alternate space instructions, 29, 74
- ancillary state registers (ASRs)
  - access, 70
  - assembly language syntax, 592
  - I/O register access, 29
  - possible registers included, 300, 371
  - privileged, 31, 570
  - reading/writing implementation-dependent
    - processor registers, 31, 570
  - writing to, 370
- AND instruction, **149**
- ANDcc instruction, **149**
- ANDN instruction, **149**
- ANDNcc instruction, **149**
- annul bit
  - in branch instructions, 160
  - in conditional branches, 175
- annulled branches, 160
- application program, 7, 70
- architectural direction note, 5
- architecture, meaning for SPARC V9, 21
- arithmetic overflow, 73
- ARRAY16 instruction, 150
- ARRAY32 instruction, 150
- ARRAY8 instruction, 150
- ASI
  - invalid, and *data\_access\_exception*, 483
- ASI register, 70
- ASI, *See* address space identifier (ASI)
- ASI\_\*REAL\* ASIs, 390
- ASI\_AIPN, 428
- ASI\_AIPN\_L, 428
- ASI\_AIPP, 428
- ASI\_AIPP\_L, 428
- ASI\_AIPS, 428
- ASI\_AIPS\_L, 428
- ASI\_AIUP, 410, 423
- ASI\_AIUPL, 410, 424
- ASI\_AIUS, 410, 423
- ASI\_AIUS\_L, 263
- ASI\_AIUSL, 410, 424
- ASI\_AS\_IF\_PRIV\_NUCLEUS, 428
- ASI\_AS\_IF\_PRIV\_NUCLEUS\_LITTLE, 428
- ASI\_AS\_IF\_PRIV\_PRIMARY, 428
- ASI\_AS\_IF\_PRIV\_PRIMARY\_LITTLE, 428
- ASI\_AS\_IF\_PRIV\_SECONDARY, 428
- ASI\_AS\_IF\_PRIV\_SECONDARY\_LITTLE, 428
- ASI\_AS\_IF\_USER\*, 97, 389, 390
- ASI\_AS\_IF\_USER\* ASIs, 389
- ASI\_AS\_IF\_USER\_NONFAULT\_LITTLE, 393
- ASI\_AS\_IF\_USER\_PRIMARY, 410, 423, 490
- ASI\_AS\_IF\_USER\_PRIMARY\_LITTLE, 392, 410, 424, 482
- ASI\_AS\_IF\_USER\_SECONDARY, 392, 410, 423, 482, 490
- ASI\_AS\_IF\_USER\_SECONDARY\_LITTLE, 392, 410, 424, 482
- ASI\_AS\_IF\_USER\_SECONDARY\_NOFAULT\_LITTLE, 393
- ASI\_BLK\_AIUP, 410, 423
- ASI\_BLK\_AIUPL, 410, 424
- ASI\_BLK\_AIUS, 410, 423
- ASI\_BLK\_AIUSL, 410, 424
- ASI\_BLK\_P, 420
- ASI\_BLK\_PL, 420
- ASI\_BLK\_S, 420
- ASI\_BLK\_SL, 421
- ASI\_BLOCK\_AS\_IF\_USER\_PRIMARY, 410, 423
- ASI\_BLOCK\_AS\_IF\_USER\_PRIMARY\_LITTLE, 4

10, 424  
 ASI\_BLOCK\_AS\_IF\_USER\_SECONDARY, 410, 423  
 ASI\_BLOCK\_AS\_IF\_USER\_SECONDARY\_LITTLE, 410, 424  
 ASI\_BLOCK\_PRIMARY, 420  
 ASI\_BLOCK\_PRIMARY\_LITTLE, 420  
 ASI\_BLOCK\_SECONDARY, 420  
 ASI\_BLOCK\_SECONDARY\_LITTLE, 421  
 ASI\_CMT\_PER\_CORE, 418  
 ASI\_CMT\_PER\_STRAND, 418, 519, 520  
 ASI\_CMT\_SHARED, 413, 523, 524, 527, 530, 534  
 ASI\_DEVICE\_ID+SERIAL\_ID, 421  
 ASI\_DMMU, 417  
 ASI\_DMMU\_DEMAP, 417  
 ASI\_DTLB\_DATA\_ACCESS\_REG, 417  
 ASI\_DTLB\_DATA\_IN\_REG, 417  
 ASI\_DTLB\_TAG\_READ\_REG, 417  
 ASI\_FL16\_P, 419  
 ASI\_FL16\_PL, 420  
 ASI\_FL16\_PRIMARY, 419  
 ASI\_FL16\_PRIMARY\_LITTLE, 420  
 ASI\_FL16\_S, 419  
 ASI\_FL16\_SECONDARY, 419  
 ASI\_FL16\_SECONDARY\_LITTLE, 420  
 ASI\_FL16\_SL, 420  
 ASI\_FL8\_P, 419  
 ASI\_FL8\_PL, 419  
 ASI\_FL8\_PRIMARY, 419  
 ASI\_FL8\_PRIMARY\_LITTLE, 419  
 ASI\_FL8\_S, 419  
 ASI\_FL8\_SECONDARY, 419  
 ASI\_FL8\_SECONDARY\_LITTLE, 419  
 ASI\_FL8\_SL, 419  
 ASI\_IMMU, 414  
 ASI\_IMMU\_DEMAP, 417  
 ASI\_ITLB\_DATA\_ACCESS\_REG, 417  
 ASI\_ITLB\_TAG\_READ\_REG, 417  
 ASI\_LD\_TWINK\_AS\_IF\_USER\_PRIMARY, 411, 426  
 ASI\_LD\_TWINK\_AS\_IF\_USER\_PRIMARY\_LITTLE, 412, 426  
 ASI\_LD\_TWINK\_AS\_IF\_USER\_SECONDARY, 411, 426  
 ASI\_LD\_TWINK\_AS\_IF\_USER\_SECONDARY\_LITTLE, 412, 426  
 ASI\_LD\_TWINK\_NUCLEUS, 412, 426, 433  
 ASI\_LD\_TWINK\_NUCLEUS[\_L], 389  
 ASI\_LD\_TWINK\_NUCLEUS\_LITTLE, 413, 426, 433  
 ASI\_LD\_TWINK\_PRIMARY, 420, 429  
 ASI\_LD\_TWINK\_PRIMARY\_LITTLE, 420, 429  
 ASI\_LD\_TWINK\_REAL, 412, 427  
 ASI\_LD\_TWINK\_REAL\_LITTLE, 412, 427, 433  
 ASI\_LD\_TWINK\_REAL\_REAL[\_]L, 389  
 ASI\_LD\_TWINK\_SECONDARY, 420, 429  
 ASI\_LD\_TWINK\_SECONDARY\_LITTLE, 420, 429  
 ASI\_LDTX\_AIUP, 263, 411, 426  
 ASI\_LDTX\_AIUP\_L, 263, 426  
 ASI\_LDTX\_AIUPL, 412  
 ASI\_LDTX\_AIUS, 263, 426  
 ASI\_LDTX\_AIUS\_L, 412, 426  
 ASI\_LDTX\_N, 263, 412, 433  
 ASI\_LDTX\_NL, 263, 413, 426, 433  
 ASI\_LDTX\_P, 263, 420  
 ASI\_LDTX\_PL, 263, 420  
 ASI\_LDTX\_R, 427  
 ASI\_LDTX\_REAL, 263, 412, 433  
 ASI\_LDTX\_REAL\_L, 412, 427, 433  
 ASI\_LDTX\_S, 263, 420  
 ASI\_LDTX\_SL, 263, 420  
 ASI\_MMU, 416  
 ASI\_MMU\_CONTEXTID, 411  
 ASI\_MMU\_REAL, 414  
 ASI\_N, 409  
 ASI\_NL, 409  
 ASI\_NUCLEUS, 120, 409  
 ASI\_NUCLEUS\_LITTLE, 120, 409  
 ASI\_NUCLEUS\_QUAD\_LDD, 433  
 ASI\_NUCLEUS\_QUAD\_LDD\_L, 433  
 ASI\_NUCLEUS\_QUAD\_LDD\_LITTLE, 433  
 ASI\_P, 418  
 ASI\_PHY\_BYPASS\_EC\_WITH\_EBIT\_L, 433  
 ASI\_PHYS\_BYPASS\_EC\_WITH\_EBIT, 433  
 ASI\_PHYS\_BYPASS\_EC\_WITH\_EBIT\_LITTLE, 433  
 ASI\_PHYS\_USE\_EC, 433  
 ASI\_PHYS\_USE\_EC\_L, 433  
 ASI\_PHYS\_USE\_EC\_LITTLE, 433  
 ASI\_PL, 418  
 ASI\_PNF, 418  
 ASI\_PNFL, 418  
 ASI\_PRIMARY, 120, 391, 392, 418  
 ASI\_PRIMARY\_LITTLE, 120, 391, 418  
 ASI\_PRIMARY\_NO\_FAULT, 388, 405, 418, 483  
 ASI\_PRIMARY\_NO\_FAULT\_LITTLE, 388, 405, 418, 483  
 ASI\_PRIMARY\_NOFAULT\_LITTLE, 393  
 ASI\_PST16\_P, 341, 418  
 ASI\_PST16\_PL, 341, 419

ASI\_PST16\_PRIMARY, 418  
 ASI\_PST16\_PRIMARY\_LITTLE, 419  
 ASI\_PST16\_S, 341, 418  
 ASI\_PST16\_SECONDARY, 418  
 ASI\_PST16\_SECONDARY\_LITTLE, 419  
 ASI\_PST16\_SL, 341  
 ASI\_PST32\_P, 341, 418  
 ASI\_PST32\_PL, 341, 419  
 ASI\_PST32\_PRIMARY, 418  
 ASI\_PST32\_PRIMARY\_LITTLE, 419  
 ASI\_PST32\_S, 341, 419  
 ASI\_PST32\_SECONDARY, 419  
 ASI\_PST32\_SECONDARY\_LITTLE, 419  
 ASI\_PST32\_SL, 341, 419  
 ASI\_PST8\_P, 418  
 ASI\_PST8\_PL, 419  
 ASI\_PST8\_PRIMARY, 418  
 ASI\_PST8\_PRIMARY\_LITTLE, 419  
 ASI\_PST8\_S, 418  
 ASI\_PST8\_SECONDARY, 418  
 ASI\_PST8\_SECONDARY\_LITTLE, 419  
 ASI\_PST8\_SL, 341, 419  
 ASI\_QUAD\_LDD\_L, 433  
 ASI\_QUAD\_LDD\_LITTLE, 433  
 ASI\_QUAD\_LDD\_PHYS, 433  
 ASI\_QUAD\_LDD\_REAL, 433  
 ASI\_QUAD\_LDD\_REAL\_LITTLE, 412  
 ASI\_QUAD\_LOAD\_REAL, 412  
 ASI\_REAL, 410, 424, 433  
 ASI\_REAL\_IO, 410, 424, 433  
 ASI\_REAL\_IO\_L, 410  
 ASI\_REAL\_IO\_LITTLE, 410, 425, 433  
 ASI\_REAL\_L, 410  
 ASI\_REAL\_LITTLE, 410, 425, 433  
 ASI\_S, 418  
 ASI\_SECONDARY, 418  
 ASI\_SECONDARY\_LITTLE, 418  
 ASI\_SECONDARY\_NO\_FAULT, 405, 418, 483  
 ASI\_SECONDARY\_NO\_FAULT\_LITTLE, 405, 418, 483  
 ASI\_SECONDARY\_NOFAULT, 393  
 ASI\_SL, 418  
 ASI\_SNF, 418  
 ASI\_SNFL, 418  
 ASI\_UMMU, 417  
*asr\_reg*, 592  
*async\_data\_error* exception, 490  
*async\_data\_error* exception (superseded), 489, 580  
 atomic

memory operations, 264, 400, 402  
 store doubleword instruction, 346, 348  
 store instructions, 323, 324  
 atomic load-store instructions  
   compare and swap, 163  
   load-store unsigned byte, 260, 353  
   load-store unsigned byte to alternate space, 261  
   simultaneously addressing doublewords, 352  
   swap R register with alternate space  
     memory, 353  
     swap R register with memory, 163, 352  
 atomicity, 388, 577  
 available (core), 7

## B

BA instruction, 154, 155, 563  
 BCC instruction, 154, 563  
 bclrg synthetic instruction, 600  
 BCS instruction, 154, 563  
 BE instruction, 154, 563  
 Berkeley RISCs, 24  
 BG instruction, 154, 563  
 BGE instruction, 154, 563  
 BGU instruction, 154, 563  
 Bicc instructions, 154, 557  
 big-endian, 7  
 big-endian byte order, 28, 95, 115  
   in hyperprivileged mode, 451  
 binary compatibility, 24  
 BL instruction, 563  
 BLD, *See* LDBLOCKF instruction  
 BLE instruction, 154, 563  
 BLEU instruction, 154, 563  
 block load instructions, 56, 245, 429  
 block store instructions, 56, 328, 429  
 blocked byte formatting, 151  
 BMASK instruction, 156  
 BN instruction, 154, 563  
 BNE instruction, 154, 563  
 BNEG instruction, 154, 563  
 BP instructions, 563  
 BPA instruction, 157, 563  
 BPCC instruction, 157, 563  
 BPcc instructions, 73, 74, 157, 564  
 BPCS instruction, 157, 563  
 BPE instruction, 157, 563  
 BPG instruction, 157, 563  
 BPGE instruction, 157, 563

- BPGU instruction, 157, 563
- BPL instruction, 157, 563
- BPLE instruction, 157, 563
- BPLEU instruction, 157, 563
- BPN instruction, 157, 563
- BPNE instruction, 157, 563
- BPNEG instruction, 157, 563
- BPOS instruction, 154, 563
- BPPOS instruction, 157, 563
- BPr instructions, **160**, 563
- BPVC instruction, 157, 563
- BPVS instruction, 157, 563
- branch
  - annulled, 160
  - delayed, 111
  - elimination, 128
  - fcc-conditional, 175, 177
  - icc-conditional, 155
  - instructions
    - on floating-point condition codes, **174**
    - on floating-point condition codes with prediction, **176**
    - on integer condition codes with prediction (BPcc), 157
    - on integer condition codes, *See* Bicc instructions
    - when contents of integer register match condition, **160**
  - prediction bit, 160
  - unconditional, 154, 158, 175, 177
  - with prediction, 22
- BRGEZ instruction, 160
- BRGZ instruction, 160
- BRLEZ instruction, 160
- BRLZ instruction, 160
- BRNZ instruction, 160
- BRZ instruction, 160
- bset synthetic instruction, 600
- BSHUFFLE instruction, 156
- BST, *See* STBLOCKF instruction
- btog synthetic instruction, 600
- btst synthetic instruction, 600
- BVC instruction, 154, 563
- BVS instruction, 154, 563
- bypass ASIs, **8**, 408
- byte, **8**
  - addressing, 120
  - data format, **35**
  - order, 28

- order, big-endian, 28
- order, little-endian, 28
- byte order
  - big-endian, 95
    - in hyperprivileged mode, 451
  - implicit, 95
  - in trap handlers, 451
  - little-endian, 95

## C

- cache
  - coherency protocol, 387
  - data, 395
  - instruction, 395
  - miss, 297
  - nonconsistent instruction cache, 395
- cacheable accesses, 386
- caching, TSB, 509
- CALL instruction
  - description, 162
  - displacement, 31
  - does not change CWP, 53
  - and JMPL instruction, 239
  - writing address into R[15], 55
- call synthetic instruction, 598
- CANRESTORE (restorable windows) register, **87**
  - and *clean\_window* exception, 129
  - and CLEANWIN register, 87, 89, 492
  - counting windows, 89
  - decremented by RESTORE instruction, 305
  - decremented by SAVED instruction, 315
  - detecting window underflow, 53
  - if registered window was spilled, 306
  - incremented by SAVE instruction, 313
  - modified by NORMALW instruction, 285
  - modified by OTHERW instruction, 287
  - range of values, 85, 86, 578
  - RESTORE instruction, 129
  - specification for RDPR instruction, 303
  - specification for WRPR instruction, 374
  - state after reset, 550
  - window underflow, 491
- CANSAVE (savable windows) register, **86**
  - decremented by SAVE instruction, 313
  - detecting window overflow, 53
  - FLUSHW instruction, 190
  - if equals zero, 129
  - incremented by RESTORE, 305

- incremented by *SAVED* instruction, 315
- range of values, 85, 86, 578
- SAVE* instruction, 493
- specification for *RDPR* instruction, 303
- specification for *WRPR* instruction, 374
- state after reset, 550
- window overflow, 491
- CAS synthetic instruction, 402
- CASA instruction, 163
  - 32-bit compare-and-swap, 400
  - alternate space addressing, 29
  - and *data\_access\_exception* (noncacheable page) exception, 483
  - atomic operation, 260
  - hardware primitives for mutual exclusion of *CASXA*, 400
  - in multiprocessor system, 261, 352, 353
  - R register use, 113
  - word access (memory), 114
- casn* synthetic instructions, 600
- CASX synthetic instruction, 400, 402
- CASXA instruction, 163
  - 64-bit compare-and-swap, 400
  - alternate space addressing, 29
  - and *data\_access\_exception* (noncacheable page) exception, 483
  - atomic operation, 261
  - doubleword access (memory), 114
  - hardware primitives for mutual exclusion of *CASA*, 400
  - in multiprocessor system, 260, 261, 352, 353
  - R register use, 113
- catastrophic error exception, 438
- cc0 instruction field
  - branch instructions, 157, 177
  - floating point compare instructions, 181
  - move instructions, 277, 564
- cc1 instruction field
  - branch instructions, 157, 177
  - floating point compare instructions, 181
  - move instructions, 277, 564
- cc2 instruction field
  - move instructions, 277, 564
- CCR (condition codes) register, 72
  - 32-bit operation (icc) bit of condition field, 73, 74
  - 64-bit operation (xcc) bit of condition field, 73, 74
  - ADD* instructions, 146
  - ASR for, 70
  - carry (c) bit of condition fields, 73
  - icc field, *See* CCR.icc field
  - MULScc* instruction, 281
  - negative (n) bit of condition fields, 73
  - overflow bit (v) in condition fields, 73
  - restored by *RETRY* instruction, 166, 309
  - saved after trap, 437
  - saving after trap, 32
  - state after reset, 549
  - TSTATE register, 92
  - write instructions, 370
  - xcc field, *See* CCR.xcc field
  - zero (z) bit of condition fields, 73
- CCR.icc field
  - add instructions, 146, 355
  - bit setting for signed division, 366
  - bit setting for signed/unsigned multiply, 367
  - bit setting for unsigned division, 365
  - branch instructions, 155, 158, 277
  - integer subtraction instructions, 351
  - logical operation instructions, 149, 286, 377
  - MULScc* instruction, 281
  - Tcc instruction, 359
- CCR.xcc field
  - add instructions, 146, 355
  - bit setting for signed/unsigned divide, 365, 366
  - bit setting for signed/unsigned multiply, 367
  - branch instructions, 158, 277
  - logical operation instructions, 149, 286, 377
  - subtract instructions, 351
  - Tcc instruction, 359
- clean register window, 313, 482
- clean window, 8
  - and window traps, 89, 491
  - CLEANWIN register, 89
  - definition, 492
  - number is zero, 129
  - trap handling, 493
- clean\_window* exception, 87, 129, 314, 457, 482, 492, 573
- CLEANWIN (clean windows) register, 87
  - CANSAVE* instruction, 129
  - clean window counting, 87
  - incremented by trap handler, 493
  - range of values, 85, 86, 578
  - specification for *RDPR* instruction, 303
  - specification for *WRPR* instruction, 374
  - specifying number of available clean windows, 492

- state after reset, 550
- value calculation, 89
- clock cycle, counts for virtual processor, 75
- clock tick registers, *See* TICK *and* STICK registers
- clock-tick register (TICK), 488
- `clr` synthetic instructions, 600
- CMP
  - disabling a core, 523
  - parking a core, 526
- `cmp` synthetic instruction, 351, 598
- CMT, 8, 511, 514
  - enabling a core, 523
  - ERROR\_STEERING register, 536, 538
  - Programming Model, 512
  - registers, 518
  - STRAND\_AVAILABLE register, 519, 523
  - STRAND\_ENABLE register, 524
  - STRAND\_ENABLE\_STATUS register, 524
  - STRAND\_ID register, 519
  - STRAND\_INTR\_ID register, 501, 520
  - STRAND\_RUNNING register
    - simultaneous updates, 528
  - STRAND\_RUNNING register, 527
  - STRAND\_RUNNING\_STATUS, 530
  - unparking a core, 526
  - XIR\_STEERING register, 534
- coherence, 8
  - between processors, 577
  - data cache, 395
  - domain, 387
  - memory, 388
  - unit, memory, 389
- compare and swap instructions, 163
- comparison instruction, 122, 351
- compatibility note, 4
- completed (memory operation), 8
- compliant SPARC V9 implementation, 25
- cond instruction field
  - branch instructions, 155, 157, 175, 177
  - floating point move instructions, 193
  - move instructions, 277
- condition codes
  - adding, 355
  - effect of compare-and-swap instructions, 164
  - extended integer (xcc), 74
  - floating-point, 175
  - icc field, 73
  - integer, 72
  - results of integer operation (icc), 74
  - subtracting, 351, 361
  - trapping on, 359
  - xcc field, 73
- condition codes register, *See* CCR register
- conditional branches, 155, 175, 177
- conditional move instructions, 32
- conforming SPARC V9 implementation, 25
- const22 instruction field of ILLTRAP
  - instruction, 235
- constants, generating, 317
- context, 8
  - nucleus, 188
- context identifier, 390
- control transfer
  - pseudo-control-transfer via WRPR to
    - PSTATE.am, 98
- control-transfer instructions (CTIs), 30, 166, 309
- conventions
  - font, 2
  - notational, 3
- conversion
  - between floating-point formats instructions, 231
  - floating-point to integer instructions, 229, 383
  - integer to floating-point instructions, 185, 234
  - planar to packed, 219
- copyback, 8
- core, 8
- CPI, 8
- CPU, pipeline draining, 86, 90
- `cpu_mondo` exception, 482
- cross-call, 8
- CTI, 8, 18
- current exception (cexc) field of FSR register, 67, 132, 570
- current window, 8
- current window pointer register, *See* CWP register
- current\_little\_endian (cle) field of PSTATE
  - register, 95, 391
- CWP (current window pointer) register
  - and instructions
    - CALL and JMPL instructions, 53
    - FLUSHW instruction, 190
    - RDPR instruction, 303
    - RESTORE instruction, 129, 305
    - SAVE instruction, 129, 305, 313
    - WRPR instruction, 374
  - and traps
    - after spill trap, 493
    - after spill/fill trap, 33

- on window trap, 493
- saved by hardware, 437

CWP (current window pointer) register, **86**

- clean windows, 87
- definition, 8
- incremented/decremented, 52, 305, 313
- overlapping windows, 52
- range of values, 85, 86, 578
- restored during RETRY, 166, 309
- specifying windows for use without cleaning, 492
- state after reset, 549
- and TSTATE register, 92
- updated during a WDR reset, 548

## D

D superscript on instruction name, **136**

d16hi instruction field

- branch instructions, 160

d16lo instruction field

- branch instructions, 160

data

- access, **8**
- cache coherence, 395
- conversion between SIMD formats, 43
- flow order constraints
  - memory reference instructions, **394**
  - register reference instructions, **394**
- formats
  - byte, **35**
  - doubleword, **35**
  - halfword, **35**
  - Int16 SIMD, **44**
  - Int32 SIMD, **44**
  - quadword, **35**
  - tagged word, **35**
  - Uint8 SIMD, **44**
  - word, **35**
- memory, 403
- types
  - floating-point, 35
  - signed integer, 35
  - unsigned integer, 35
  - width, 35

Data Cache Unit Control register, *See* DCUCR

Data Synchronous Fault Address register, *See* D-SFAR

Data Synchronous Fault Status register, *See* D-SFSR

*data\_access\_error* exception, **482**

- with load instructions, 265

*data\_access\_exception* (invalid ASI) exception

- with load alternate instructions, 243

*data\_access\_exception* exception, **482**

- register update policy, 510

- with compare-and-swap instructions, 165

- with LD instructions, 241

- with LDSHORTF instructions, 244, 247

- with LDTXA instructions, 265

- with load instructions, 251, 267, 270

- with load instructions and ASIs, 254, 426, 427, 429, 430, 431

- with store instructions and ASIs, 254, 426, 427, 429, 430, 431

- with STPARTIALF instructions, 343

- with SWAPA instruction, 354

*data\_access\_MMU\_error* exception

- on PREFETCH, 293, 298

- with CASA instruction, 165

- with load instructions, 241, 244, 248, 251, 255, 257, 260, 262, 265, 267, 270

- with store instructions, 323, 343, 345, 347, 350

- with SWAP instruction, 352, 354

*data\_access\_MMU\_miss* exception

- with integer load instructions, 241

- with load alternate instructions, 244

- with load instructions, 248

- with PREFETCH instruction, 293

*data\_access\_MMU\_miss* exception

- (superseded), **483**

*data\_access\_protection* exception

- (superseded), **490**

*data\_invalid\_TSB\_entry* exception, **484**

*data\_real\_translation\_miss* exception, **484**

*data\_store\_error* exception (SPARC V8), **489**

DCTI couple, **127**

DCTI instructions, **8**

- behavior, 111

- RETURN instruction effects, 311

dec synthetic instructions, 600

decccg synthetic instructions, 600

deferred trap, **445**

- distinguishing from disrupting trap, 447

- floating-point, 304

- restartable

- implementation dependency, 447

- software actions, 447

delay instruction

- and annul field of branch instruction, 175
- annulling, 31
- conditional branches, 177
- DONE instruction, 166
- executed after branch taken, 160
- following delayed control transfer, 31
- RETRY instruction, 309
- RETURN instruction, 311
- unconditional branches, 177
- with conditional branch, 158
- delayed branch, 111
- delayed control transfer, 160
- delayed CTL, *See* DCTI
- demap, 8
- denormalized number, 9
- deprecated, 9
- deprecated exceptions
  - tag\_overflow*, 489
- deprecated instructions
  - FBA, 174
  - FBE, 174
  - FBG, 174
  - FBGE, 174
  - FBL, 174
  - FBLE, 174
  - FBLG, 174
  - FBN, 174
  - FBNE, 174
  - FBO, 174
  - FBU, 174
  - FBUE, 174
  - FBUGE, 174
  - FBUL, 174
  - FBULE, 174
  - LDFSR, 256
  - LDTW, 266
  - LDTWA, 268
  - MULScC, 72, 281
  - RDY, 70, 72, 299
  - SDIV, 72, 364
  - SDIVcc, 72, 364
  - SMUL, 72, 367
  - SMULcc, 72, 367
  - STFSR, 339
  - STTW, 346
  - STTWA, 348
  - SWAP, 352
  - SWAPA, 353
  - TADDccTV, 356
  - TSUBccTV, 362
  - UDIV, 72, 364
  - UDIVcc, 72, 364
  - UMUL, 72, 367
  - UMULcc, 72, 367
  - WRY, 70, 72, 369
- dev\_mondo* exception, 484
- disable (core), 9
- disabled (core), 9
- disabling CMP core, 523
- disp19 instruction field
  - branch instructions, 157, 177
- disp22 instruction field
  - branch instructions, 154, 175
- disp30 instruction field
  - word displacement (CALL), 162
- dispatch, 9
- disrupting trap, 447
  - differences from reset trap, 449
- divide instructions, 30, 283, 364
- division\_by\_zero* exception, 123, 283, 484
- division-by-zero bits of FSR.aexc/FSR.cexc
  - fields, 69
- DMMU Tag Access register
  - context field after *data\_access\_exception*, 510
- DONE instruction, 166
  - effect on HTSTATE, 106
  - effect on TNPC register, 91
  - effect on TSTATE register, 93
  - executed in RED\_state, 442
  - generating *illegal\_instruction* exception, 486
  - modifying CCR.xcc condition codes, 73
  - return from trap, 437
  - return from trap handler with different GL
    - value, 103
  - target address, 31
- doublet, 9
- doubleword, 9
  - addressing, 118
  - alignment, 28, 114, 389
  - data format, 35
  - definition, 9
- D-SFAR
  - state after reset, 551
- D-SFSR register, *See* SFSR register
- DuTLB, disabled, 483

**E**

- EDGE16 instruction, 168
- EDGE16L instruction, 168
- EDGE16LN instruction, 170
- EDGE16N instruction, 170
- EDGE32 instruction, 168
- EDGE32L instruction, 168
- EDGE32LN instruction, 170
- EDGE32N instruction, 170
- EDGE8 instruction, 168
- EDGE8L instruction, 168
- EDGE8LN instruction, 170
- EDGE8N instruction, 170
- emulating multiple unsigned condition codes, 128
- enable (core), 9
- enable floating-point
  - See FPRS register, *fef* field
  - See PSTATE register, *pef* field
- enabled (core), 9
- enabling CMP core, 523
- error\_state*, 444
  - definition, 440
  - effects when entering, 572
  - entering, 470, 472, 478, 479, 480
  - exiting, 470
  - recognizing interrupts, 471
    - and *RED\_state*, 441
- ERROR\_STEERING register, 538
- even parity, 9
- exception, 9
- exceptions
  - See also individual exceptions
  - async\_data\_error*, 490
  - async\_data\_error* (superseded), 489
  - catastrophic error, 438
  - causing traps, 437
  - clean\_window*, 457, 482, 573
  - cpu\_mondo*, 482
  - data\_access\_error*, 482
  - data\_access\_exception*, 482
  - data\_access\_MMU\_error*
    - on *PREFETCH*, 293
    - on *PREFETCH*, 293
  - data\_access\_MMU\_error*, 165, 241, 244, 248, 251, 255, 257, 260, 262, 265, 267, 270, 298, 323, 343, 345, 347, 350, 352
  - data\_access\_MMU\_miss*, 293
  - data\_access\_MMU\_miss* (superseded), 483
  - data\_access\_protection* (superseded), 490
  - data\_invalid\_TSB\_entry*, 484
  - data\_real\_translation\_miss*, 484
  - data\_store\_error* (SPARC V8), 489
  - definition, 438
  - dev\_mondo*, 484
  - division\_by\_zero*, 484
  - fast\_data\_access\_MMU\_miss*, 293, 484
  - fast\_data\_access\_protection*, 484
  - fast\_ECC\_error*, 491
  - fill\_n\_normal*, 485
  - fill\_n\_other*, 485
  - fp\_disabled*
    - and GSR, 80
  - fp\_disabled*, 485
  - fp\_exception\_ieee\_754*, 485
  - fp\_exception\_other*, 485
  - guest\_watchdog*, 485
  - hstick\_match*, 106, 109, 485
  - htrap\_instruction*, 485
  - illegal\_instruction*
    - and SIR instruction, 450
  - illegal\_instruction*, 104, 485
  - instruction\_access\_error*, 486
  - instruction\_access\_exception*, 486, 486
  - instruction\_access\_MMU\_miss*, 486
  - instruction\_breakpoint*, 486
  - instruction\_invalid\_TSB\_entry*, 487
  - instruction\_real\_translation\_miss*, 487
  - internal\_processor\_error*, 487
  - interrupt\_level\_14*
    - and *SOFTINT.int\_level*, 81
    - and *STICK\_CMPR.stick\_cmpr*, 85
    - and *TICK\_CMPR.tick\_cmpr*, 83
  - interrupt\_level\_14*, 487
  - interrupt\_level\_15*
    - and *SOFTINT.int\_level*, 81
  - interrupt\_level\_15*, 487
  - interrupt\_level\_n*
    - and *SOFTINT* register, 80
    - and *SOFTINT.int\_level*, 81
  - interrupt\_level\_n*, 448, 487
  - LDDF\_mem\_address\_not\_aligned*, 487
  - LDQF\_mem\_address\_not\_aligned*, 491
  - mem\_address\_not\_aligned*, 487
  - nonresumable\_error*, 487
  - PA\_watchpoint*, 487
  - pending, 33
  - pic\_overflow*, 487
  - power\_on\_reset*, 488

- privileged\_action*, 488
- privileged\_opcode*
  - and access to register-window PR state registers, 85, 90, 100, 103
  - and access to SOFTINT, 81
  - and access to SOFTINT\_CLR, 82
  - and access to SOFTINT\_SET, 82
  - and access to STICK\_CMPR, 84
  - and access to TICK\_CMPR, 83
- privileged\_opcode*, 488
- RA\_watchpoint*, 462, 467
- RED\_state\_exception*, 488
- resumable\_error*, 488
- software\_initiated\_reset*, 488
- spill\_n\_normal*, 314, 489
- spill\_n\_other*, 314, 489
- STDF\_mem\_address\_not\_aligned*, 489
- store\_error*, 489
- STQF\_mem\_address\_not\_aligned*, 491
- sw\_recoverable\_error*, 489, 490
- tag\_overflow* (deprecated), 489
- trap\_instruction*, 489
- trap\_level\_zero*
  - state after reset, 549
- trap\_level\_zero*, 489
- unimplemented\_LDTW*, 490
- unimplemented\_STTW*, 490
- VA\_watchpoint*, 490
- watchdog\_reset*
  - and *guest\_watchdog*, 439
- watchdog\_reset*, 490
- window fill, 457
- window spill, 457
- execute unit, 393
- execute\_state*
  - and *error\_state*, 472
  - and *RED\_state*, 472
  - returning to, 470
  - trap processing, 440, 470
- explicit ASI, 9, 120, 409
- extended word, 10
  - addressing, 118
- externally\_initiated\_reset* (XIR), 478, 484, 547
  - causing entry into *RED\_state*, 441
  - entering *error\_state*, 440
  - and *error\_state*, 456
  - for critical system events, 450
  - for debugging, 443
  - partial-processor, 533

- RED\_state* trap processing, 474
- to virtual processor, 547

## F

- F registers, 10, 27, 131, 379, 452
- FABSd instruction, 171, 561, 562
- FABSq instruction, 171, 561, 562
- FABSs instruction, 171
- FADD, 172
- FADDd instruction, 172
- FADDq instruction, 172
- FADDs instruction, 172
- FALIGNDATA instruction, 173
- FAND instruction, 227
- FANDNOT1 instruction, 227
- FANDNOT1S instruction, 227
- FANDNOT2 instruction, 227
- FANDNOT2S instruction, 227
- FANDS instruction, 227
- fast\_data\_access\_MMU\_miss* exception, 484
  - register update policy, 510
  - with integer load instructions, 241
  - with load alternate instructions, 244
  - with PREFETCH instruction, 293
- fast\_data\_access\_protection* exception, 484
  - register update policy, 510
  - write permission not granted, 508
- fast\_ECC\_error* exception, 491
- fast\_instruction\_access\_MMU\_miss* exception
  - register update policy, 510
- FBA instruction, 174, 175, 563
- FBE instruction, 174, 563
- FBfcc instructions, 61, 174, 485, 557, 563
- FBG instruction, 174, 563
- FBGE instruction, 174, 563
- FBL instruction, 174, 563
- FBLE instruction, 174, 563
- FBLG instruction, 174, 563
- FBN instruction, 174, 175, 563
- FBNE instruction, 174, 563
- FBO instruction, 174, 563
- FBPA instruction, 176, 177, 563
- FBPE instruction, 176, 563
- FBPfcc instructions, 61, 176, 557, 563, 564
- FBPG instruction, 176, 563
- FBPGE instruction, 176, 563
- FBPL instruction, 176, 563
- FBPLE instruction, 176, 563

- FBPLG instruction, 176, 563
- FBPN instruction, 176, 177, 563
- FBPNE instruction, 176, 563
- FBPO instruction, 176, 563
- FBPU instruction, 176, 563
- FBPUE instruction, 176, 563
- FBPUG instruction, 176, 563
- FBPUGE instruction, 176, 563
- FBPUL instruction, 176, 563
- FBPULE instruction, 176, 563
- FBU instruction, 174, 563
- FBUE instruction, 174, 563
- FBUG instruction, 174, 563
- FBUGE instruction, 174, 563
- FBUL instruction, 174, 563
- FBULE instruction, 174, 563
- fcc-conditional branches, 175, 177
- fccn*, 10
- FCMP instructions, 564
- FCMP\* instructions, 61, 62, 181
- FCMPd instruction, 181, 381, 562
- FCMPE instructions, 564
- FCMPE\* instructions, 61, 62, 181
- FCMPEd instruction, 181, 381, 562
- FCMPEq instruction, 181, 381, 562
- FCMPEQ16 instruction, 178
- FCMPEQ32 instruction, 178
- FCMPEs instruction, 181, 381, 562
- FCMPGT instruction, 178
- FCMPGT16 instruction, 178
- FCMPGT32 instruction, 178
- FCMPLE16 instruction, 178
- FCMPLE16 instruction, 178
- FCMPLE32 instruction, 178
- FCMPLE32 instruction, 178
- FCMPNE16 instruction, 178, 179
- FCMPNE32 instruction, 178, 179
- FCMPq instruction, 181, 381, 562
- FCMPs instruction, 181, 381, 562
- fcn instruction field
  - DONE instruction, 166
  - PREFETCH, 291
  - RETRY instruction, 309
- FDIVd instruction, 183
- FDIVq instruction, 183
- FDIVs instructions, 183
- FdMULq instruction, 207
- FdTOi instruction, 229, 383
- FdTOq instruction, 231
- FdTOs instruction, 231
- FdTOx instruction, 229, 562
- fef field of FPRS register, 77
  - and access to GSR, 80
  - and *fp\_disabled* exception, 485
  - branch operations, 175, 177
  - byte permutation, 156
  - comparison operations, 179, 182
  - data movement operations, 278
  - enabling FPU, 96
  - floating-point operations, 171, 172, 183, 185, 191, 196, 199, 207, 209, 228, 229, 231, 233, 234, 250, 252, 256, 258
  - integer arithmetic operations, 218, 223
  - logical operations, 224, 225, 227
  - memory operations, 247
  - read operations, 301, 319, 330
  - special addressing operations, 147, 173, 333, 339, 343, 345, 371
- fef, *See* FPRS register, fef field
- FEXPAND instruction, 184
- FEXPAND operation, 184
- fill handler, 306
- fill register window, 485
  - overflow/underflow, 53
  - RESTORE instruction, 89, 305, 491
  - RESTORED instruction, 130, 307, 493
  - RETRY instruction, 493
  - selection of, 492
  - trap handling, 492, 493
  - trap vectors, 306
  - window state, 89
- fill\_n\_normal* exception, 306, 312, 485, 485
- fill\_n\_other* exception, 306, 312, 485
- FiTOd instruction, 185
- FiTOq instruction, 185
- FiTOs instruction, 185
- fixed values, 236
- fixed-point scaling, 202
- floating point
  - absolute value instructions, 171
  - add instructions, 172
  - compare instructions, 61, 62, 181, 181, 381
  - condition code bits, 175
  - condition codes (fcc) fields of FSR register, 64, 175, 177, 181
  - data type, 35
  - deferred-trap queue (FQ), 304
  - divide instructions, 183

- exception, **10**
- exception, encoding type, **63**
- FPRS register, **370**
- FSR condition codes, **62**
- move instructions, **191**
- multiply instructions, **207**
- negate instructions, **209**
- operate (FPop) instructions, **10, 32, 63, 67, 131, 256**
- registers
  - destination F, **379**
  - FPRS, *See* FPRS register
  - FSR, *See* FSR register
  - programming, **59**
- rounding direction, **62**
- square root instructions, **228**
- subtract instructions, **233**
- trap types, **10**
  - IEEE\_754\_exception, **64, 65, 67, 70, 380**
  - invalid\_fp\_register, **171, 172, 233**
  - unfinished\_FPop, **64, 65, 70, 172, 183, 208, 232, 233, 380**
    - results after recovery, **65**
  - unimplemented\_FPop, **65, 70, 171, 172, 182, 183, 185, 191, 197, 200, 208, 209, 230, 232, 233, 380**
- traps
  - deferred, **304**
  - precise, **304**
- floating-point condition codes (fcc) fields of FSR register, **452**
- floating-point operate (FPop) instructions, **485**
- floating-point trap types
  - IEEE\_754\_exception, **452, 485**
- floating-point unit (FPU), **10, 27**
- FLUSH instruction, **187**
  - memory ordering control, **273**
- FLUSH instruction
  - memory/instruction synchronization, **186**
- FLUSH instruction, **186, 404**
  - data access, **8**
  - immediacy of effect, **188**
  - in multiprocessor system, **186**
  - in self-modifying code, **187**
  - latency, **577**
- flush instruction memory, *See* FLUSH instruction
- flush register windows instruction, **190**
- FLUSHW instruction, **190, 489**
  - effect, **32**
  - management by window traps, **89, 491**
  - spill exception, **131, 190, 493**
- FMOVcc instructions
  - conditionally moving floating-point register contents, **74**
  - conditions for copying floating-point register contents, **127**
  - copying a register, **61**
  - encoding of opf<84> bits, **562**
  - encoding of opf\_cc instruction field, **564**
  - encoding of rcond instruction field, **563**
  - floating-point moves, **193**
  - FPop instruction, **131**
    - used to avoid branches, **197, 277**
- FMOVccd instruction, **562**
- FMOVccq instruction, **562**
- FMOVd instruction, **191, 561, 562**
- FMOVDfcc instructions, **193**
- FMOVDGEZ instruction, **198**
- FMOVDGZ instruction, **198**
- FMOVDicc instructions, **193**
- FMOVDLEZ instruction, **198**
- FMOVDLZ instruction, **198**
- FMOVDNZ instruction, **198**
- FMOVDZ instruction, **198**
- FMOVq instruction, **191, 561, 562**
- FMOVQfcc instructions, **193, 196**
- FMOVqGEZ instruction, **198**
- FMOVqGZ instruction, **198**
- FMOVQicc instructions, **193, 196**
- FMOVqLEZ instruction, **198**
- FMOVqLZ instruction, **198**
- FMOVqNZ instruction, **198**
- FMOVqZ instruction, **198**
- FMOVr instructions, **132, 563**
- FMOVRq instructions, **199**
- FMOVRsGZ instruction, **198**
- FMOVRsLEZ instruction, **198**
- FMOVRsLZ instruction, **198**
- FMOVRsNZ instruction, **198**
- FMOVRsZ instruction, **198**
- FMOVs instruction, **191**
- FMOVSc instructions, **195**
- FMOVsfcc instructions, **193**
- FMOVSGEZ instruction, **198**
- FMOVsicc instructions, **193**
- FMOVsxcc instructions, **193**
- FMOVxcc instructions, **193, 196**
- FMUL8SUX16 instruction, **201, 204**

FMUL8ULx16 instruction, 201, 204  
 FMUL8x16 instruction, 201, 202  
 FMUL8x16AL instruction, 201, 203  
 FMUL8x16AU instruction, 201, 203  
 FMULd instruction, 207  
 FMULD8SUx16 instruction, 201, 205  
 FMULD8ULx16 instruction, 201, 206  
 FMULq instruction, 207  
 FMULs instruction, 207  
 FNAND instruction, 227  
 FNANDS instruction, 227  
 FNEG instructions, 209  
 FNEGd instruction, 209, 561, 562  
 FNEGq instruction, 209, 561, 562  
 FNEGs instruction, 209  
 FNOR instruction, 227  
 FNORS instruction, 227  
 FNOT1 instruction, 225  
 FNOT1S instruction, 225  
 FNOT2 instruction, 225  
 FNOT2S instruction, 225  
 FONE instruction, 224  
 FONES instruction, 224  
 FOR instruction, 227  
 formats, instruction, 112  
 FORNOT1 instruction, 227  
 FORNOT1S instruction, 227  
 FORNOT2 instruction, 227  
 FORNOT2S instruction, 227  
 FORS instruction, 227  
*fp\_disabled* exception, 485
 

- absolute value instructions, 171, 172, 233
- and GSR, 80
- FPop instructions, 132
- FPRS.fef disabled, 77
- PSTATE.pef not set, 77
- with branch instructions, 175, 177
- with compare instructions, 180
- with conversion instructions, 185, 230, 232, 234
- with floating-point arithmetic instructions, 183, 208, 218, 223
- with FMOV instructions, 191
- with load instructions, 255
- with move instructions, 197, 200, 278
- with store instructions, 333, 337, 339, 340, 343, 345, 371

*fp\_exception* exception, 67  
*fp\_exception\_ieee\_754* "invalid" exception, 229  
*fp\_exception\_ieee\_754* exception, 485
 

- and tem bit of FSR, 63
- cause encoded in FSR.ftt, 64
- FSR.aexc, 67
- FSR.cexc, 68
- FSR.ftt, 67
- generated by FCMP or FCMPE, 62
- and IEEE 754 overflow / underflow conditions, 67, 68
- trap handler, 380
- when FSR.tem = 0, 452
- when FSR.tem = 1, 452
- with floating-point arithmetic instructions, 172, 183, 208, 233

*fp\_exception\_other* exception, 70, 485
 

- absolute value instructions, 171
- cause encoded in FSR.ftt, 64
- FADDq instruction, 172, 233
- FCMP{E}q instructions, 182
- FDIVq instruction, 183
- FdTOq, FqTOd instructions, 232
- FiTOq instruction, 185
- FMOVcc instruction, 197
- FMOVq instruction, 191
- FMOVRq instruction, 200
- FMULq, FdMULq instructions, 208
- FNEGq instruction, 209
- FqTOx, FqTOi instructions, 230
- FSQRT instructions, 228
- FxTOq instruction, 234
- incorrect IEEE Std 754-1985 result, 132, 569
- occurrence, 145
- supervisor handling, 380
- trap type of unfinished\_FPop, 65
- unimplemented\_FPop for quad FPOps, 60
- when quad FPop unimplemented in hardware, 66
- with floating-point arithmetic instructions, 183, 208

 FPACK instruction, 80  
 FPACK instructions, 210–214  
 FPACK16 instruction, 210, 211  
 FPACK16 operation, 211  
 FPACK32 instruction, 210, 212  
 FPACK32 operation, 212  
 FPACKFIX instruction, 210, 214  
 FPACKFIX operation, 214  
 FPADD16 instruction, 216  
 FPADD16S instruction, 216  
 FPADD32 instruction, 216

FPADD32S instruction, 216  
 FPMERGE instruction, 219  
 FPop, 10  
 FPop instruction  
     unimplemented, 485  
 FPop, *See* floating-point operate (FPop) instructions  
 FPRS register  
     *See also* floating-point registers state (FPRS)  
     register  
 FPRS register, 76  
     ASR summary, 71  
     definition, 10  
     fef field, 132, 451  
     RDFPRS instruction, 300  
     state after reset, 550  
 FPRS register fields  
     dl (dirty lower fp registers), 77  
     du (dirty upper fp registers), 77  
     fef, 77  
     fef, *See also* fef field of FPRS register  
 FPSUB16 instruction, 221  
 FPSUB16S instruction, 221  
 FPSUB32 instruction, 221  
 FPSUB32S instruction, 221  
 FPU, 10  
 FqTOd instruction, 231  
 FqTOi instruction, 229, 383  
 FqTOs instruction, 231  
 FqTOx instruction, 229, 561, 562  
*freq*, 592  
 FsMULd instruction, 207  
 FSQRTd instruction, 228  
 FSQRTq instruction, 228  
 FSQRTs instruction, 228  
 FSR (floating-point state) register  
     fields  
         aexc (accrued exception), 64, 65, 66, 67, 380  
         aexc (accrued exceptions)  
             in user-mode trap handler, 380  
         -- dza (division by zero) bit of aexc, 69  
         -- nxa (rounding) bit of aexc, 70  
         cexc (current exception), 62, 64, 65, 67, 67, 68,  
             380, 485  
         cexc (current exceptions)  
             in user-mode trap handler, 380  
         -- dzc (division by zero) bit of cexc, 69  
         -- nxc (rounding) bit of cexc, 70  
         fcc (condition codes), 61, 64, 65, 380, 593  
         fccn, 62  
         ftt (floating-point trap type), 63, 67, 132, 332,  
             339, 485  
             in user-mode trap handler, 380  
         not modified by LDFSR/LDXFSR  
             instructions, 61  
         qne (queue not empty), 66  
             in user-mode trap handler, 380  
         rd (rounding), 62  
         tem (trap enable mask), 62, 66, 68, 485  
         ver, 63  
     FSR (floating-point state) register, 61  
     after floating-point trap, 380  
     compliance with IEEE Std 754-1985, 70  
     LDFSR instruction, 256  
     reading/writing, 61  
     state after reset, 550  
     values in ftt field, 64  
     writing to memory, 332, 339  
 FSRC1 instruction, 225  
 FSRC1S instruction, 225  
 FSRC2 instruction, 225  
 FSRC2S instruction, 225  
 FsTOd instruction, 231  
 FsTOi instruction, 229, 383  
 FsTOq instruction, 231  
 FsTOx instruction, 229, 561, 562  
 FSUBd instruction, 233  
 FSUBq instruction, 233  
 FSUBs instruction, 233  
 functional choice, implementation-dependent, 569  
 FXNOR instruction, 227  
 FXNORS instruction, 227  
 FXOR instruction, 227  
 FXORS instruction, 227  
 FxTOd instruction, 234, 562  
 FxTOq instruction, 234, 562  
 FxTOs instruction, 234, 562  
 FZERO instruction, 224  
 FZEROS instruction, 224

## G

General Status register, *See* GSR  
 generating constants, 317  
 GL register, 101  
     access, 102  
     during resets, 103  
     during trap processing, 470  
     function, 101

- reading with RDPR instruction, 303, 374
- relationship to TL, 102
- restored during RETRY, 166, 309
- SPARC V9 compatibility, 99
- and TSTATE register, 92
- value restored from TSTATE[TL], 103
- value restored from TSTATE[TL], 102, 166, 309
- and VER.maxgl, 108
- writing to, 102
- global level register, *See* GL register
- global registers, 22, 26, 49, 49, 49, 569
- Graphics Status register, *See* GSR
- GSR (general status) register
  - fields
    - align, 80
    - im (interval mode) field, 80
    - irnd (rounding), 80
    - mask, 80
    - scale, 80
- GSR (general status) register
  - ASR summary, 71
- GSR (graphics status) register
  - state after reset, 550, 551
- guest\_watchdog* exception, 485

## H

- H superscript on instruction name, 136
- halfword, 10
  - alignment, 28, 114, 389
  - data format, 35
- hardware
  - dependency, 568
  - traps, 457
- hardware tablewalk
  - disabled, 508
  - and TSB pointers, 508
- hardware trap stack, 32
- HINTP register, 106
- HPR state registers (ASRs), 103–109
- hpriv field of HPSTATE register, 442
- HPSTATE register
  - fields
    - hpriv
      - and access to PCR, 78
- HPSTATE register, 104
  - entering hyperprivileged execution mode, 437
  - hpriv field, 105
  - hpriv field, *See also* hyperprivileged (hpriv) field

- of HPSTATE register
  - and HTSTATE register, 105
  - ibe field, 486
  - ibe field, 104
  - red field, 104, 442
  - state after reset, 549
  - tlz field, 105
  - tlz field, and *trap\_level\_zero* exception, 105, 489
- hsp (*hstick\_match* pending) field of HINTP register, 106, 107, 109
- HSTICK\_CMPR register, 109, 485
  - and HINTP, 106
- hstick\_match* exception, 106, 109, 485
- hstick\_match* pending (hsp) field of HINTP register, 106, 107, 109
- HTBA (hyperprivileged trap base address) register, 107, 439, 485
  - establishing table address, 437
  - initialization, 454
  - state after reset, 549
- htrap\_instruction* exception, 360, 485
- HTSTATE (hyperprivileged trap state) register, 105
  - number of copies for reading, 302
  - number of copies for writing, 372
  - reading, 302
  - writing to, 372
- HVER (version) register
  - fields
    - maxtl, 108
    - maxwin, 108
- HVER (version) register, 107
  - state after reset, 550
- HVER (version) register fields
  - maxwin, 108
- hyperprivileged
  - mode, 90
  - registers, 103
- hyperprivileged (hpriv) field of HPSTATE register, 349, 371
  - access to register-window PR state registers, 90
  - and trap control, 451
  - compare and swap instructions, 164, 354
  - disrupting trap condition detected, 448
  - load instructions, 243, 247, 253, 261, 269, 330
  - privileged\_action* exception, 391
  - store instructions, 324, 336, 349
  - trap\_level\_zero* exception, 167, 310, 372, 373, 375, 489
- hyperprivileged (software), 10

- hyperprivileged (state), **10**
- hyperprivileged mode
  - byte order, 451
- hyperprivileged scratchpad registers
  - state after reset, 551
- hypervisor (software), **11**

## I

- i (integer) instruction field
  - arithmetic instructions, 281, 283, 286, 364, 367
  - floating point load instructions, 249, 252, 256
  - flush memory instruction, 186
  - flush register instruction, 190
  - jump-and-link instruction, 239
  - load instructions, 240, 260, 261, 266, 268
  - logical operation instructions, 149, 286, 377
  - move instructions, 277, 279
  - POPC, 289
  - PREFETCH, 291
  - RETURN, 311
- I/O
  - access, 387
  - memory, 386
  - memory-mapped, 387
- IEEE 754, **11**
- IEEE Std 754-1985, 11, 21, 62, 65, 68, 70, 132, 379, 569
- IEEE\_754\_exception floating-point trap type, **11**, 64, 65, 67, 70, 380, 452, 485
- IEEE-754 exception, **11**
- IER register (SPARC V8), 371
- illegal\_instruction*
  - and OTHERW instruction, 318
- illegal\_instruction* exception, 190, **485**
  - and SIR instruction, 450
  - attempt to write in nonprivileged mode, 84
  - DONE/RETRY, 167, 310, 311
  - HTSTATE register, reading/writing, 104, 106
  - ILLTRAP, 235
  - instruction not specifically defined in architecture, 133
  - not implemented in hardware, 145
  - POPC, 290
  - PREFETCH, 298
  - RETURN, 312
  - with BPr instruction, 161
  - with branch instructions, 158, 161
  - with CASA and CASXA instructions, 164, 286
    - with CASXA instruction, 165
    - with DONE instruction, 167
    - with FMOV instructions, 191
    - with FMOVcc instructions, 197
    - with load instructions, 55, 247, 251, 267, 269, 430
    - with move instructions, 278, 280
    - with RDHPR instructions, 302
    - with read hyperprivileged register instructions, 302, 303
    - with read instructions, 300, 301, 302, 303, 375, 572
    - with store instructions, 333, 340, 346, 347, 349
    - with STQFA instruction, 337
    - with Tcc instructions, 360
    - with TPC register, 90
    - with TSTATE register, 92
    - with write instructions, 371, 372, 373, 376
    - write to ASR 5, 76
    - write to STICK register, 84
    - write to TICK register, 75
- ILLTRAP instruction, **235**, 485
- imm\_asi instruction field
  - explicit ASI, providing, 120
  - floating point load instructions, 252
  - load instructions, 261, 266, 268
  - PREFETCH, 291
- immediate CTI, 111
- I-MMU
  - and instruction prefetching, 388
- IMPDEP1 instruction, 237
- IMPDEP1 instructions, 236, 565, 566
- IMPDEP2A instructions, 236, 486, 574
- IMPDEP2B instructions, 132, 236, 486
- implementation, **11**
- implementation (impl) field of VER register, **108**
- implementation dependency, **567**
- implementation dependent, **11**
- implementation note, **4**, **5**
- implementation number (impl) field of HVER register, **108**
- implementation-dependent functional choice, 569
- implementation-dependent instructions, *See* IMPDEP2A instructions
- implicit ASI, **11**, **120**, 408
- implicit ASI memory access
  - LDFSR, 256
  - LDSTUB, 260
  - load fp instructions, 249
  - load integer doubleword instructions, 266

- load integer instructions, 240
  - STD, 346
  - STFSR, 339
  - store floating-point instructions, 332
  - store integer instructions, 323
  - SWAP, 352
- implicit byte order, 95
- in* registers, 49, 52, 313
- inccc* synthetic instructions, 600
- inexact accrued (*nx*a) bit of *aexc* field of FSR register, 383
- inexact current (*nx*c) bit of *cx* field of FSR register, 383
- inexact quotient, 364, 365
- infinity, 383, 384
- initiated, **11**
- input/output (I/O) locations
  - access by nonprivileged code, 570
  - behavior, 386
  - contents and addresses, 570
  - identifying, 577
  - order, 386
  - semantics, 577
  - value semantics, 386
- instruction fields, **11**
  - See also* individual instruction fields
  - definition, **11**
- instruction group, **11**
- instruction MMU, *See* I-MMU
- instruction prefetch buffer, invalidation, 187
- instruction set architecture (ISA), **11**, **11**, 23
- Instruction Synchronous Fault Status register, *See* I-SFSR
- instruction\_access\_error* exception, **486**
- instruction\_access\_exception* exception, **486**
  - register update policy, 510
- instruction\_access\_MMU\_miss* exception, **486**
- instruction\_breakpoint* exception, **486**
- instruction\_invalid\_TSB\_entry* exception, **487**
- instruction\_real\_translation\_miss* exception, **487**
- instructions
  - 32-bit wide, 22
  - alignment, **114**
  - alignment, 28, 147, 389
  - arithmetic, integer
    - addition, 146, 355
    - division, 30, 283, 364
    - multiplication, 30, 281, 283, 367
    - subtraction, 351, 361
  - tagged, 30
- array addressing, 150
- atomic
  - CASA/CASXA, 163
  - load twin extended word from alternate space, 263
  - load-store, 113, 163, 260, 261, 352, 353
  - load-store unsigned byte, 260, 261
  - successful loads, 240, 242, 267, 269
  - successful stores, 323, 324
- branch
  - branch if contents of integer register match condition, **160**
  - branch on floating-point condition codes, **174**, **176**
  - branch on integer condition codes, **154**, **157**
- cache, 395
- causing illegal instruction, 235
- compare and swap, **163**
- comparison, 122, 351
- conditional move, 32
- control-transfer (CTIs), 30, 166, 309
- conversion
  - convert between floating-point formats, **231**
  - convert floating-point to integer, **229**
  - convert integer to floating-point, **185**, **234**
  - floating-point to integer, 383
- count of number of bits, 289
- edge handling, 168
- fetches, 114
- floating point
  - compare, 61, 62, 181
  - floating-point add, **172**
  - floating-point compare, 381
  - floating-point divide, **183**
  - floating-point load, 113, 249
  - floating-point load from alternate space, 252
  - floating-point move, 191, **193**, 198
  - floating-point operate (FPop), 32, 256
  - floating-point square root, **228**
  - floating-point store, 113, 332
  - floating-point store to alternate space, 335
  - floating-point subtract, **233**
  - operate (FPop), 63, 67
  - short floating-point load, 258
  - short floating-point store, 344
  - status of floating-point load, 256
- flush instruction memory, **186**
- flush register windows, **190**

- formats, **112**
- generate software-initiated reset, **320**
- implementation-dependent, *See* IMPDEP2A
  - instructions
- jump and link, **31, 239**
- loads
  - block load, **245**
  - floating point, *See* instructions: floating point
  - integer, **113**
  - integer from alternate space, **516**
  - simultaneously addressing doublewords, **352**
  - unsigned byte, **163, 260**
  - unsigned byte to alternate space, **261**
- logical operations
  - 64-bit/32-bit, **225, 227**
  - AND, **149**
  - logical 1-operand ops on F registers, **224**
  - logical 2-operand ops on F registers, **225**
  - logical 3-operand ops on F registers, **227**
  - logical XOR, **377**
  - OR, **286**
- memory, **403**
- moves
  - floating point, *See* instructions: floating point
  - move integer register, **275, 279**
  - on condition, **22**
- ordering MEMBAR, **122**
- permuting bytes specified by GSR.mask, **156**
- pixel component distance, **288, 288**
- pixel formatting (PACK), **210**
- prefetch data, **291**
- read hyperprivileged register, **302**
- read privileged register, **303**
- read state register, **31, 299**
- register window management, **32**
- reordering, **394**
- reserved, **132**
- reserved* fields, **145**
- RETRY
  - and restartable deferred traps, **447**
- RETURN vs. RESTORE, **311**
- sequencing MEMBAR, **122**
- set high bits of low word, **317**
- set interval arithmetic mode, **319**
- setting GSR.mask field, **156**
- shift, **30**
- shift, **321**
- shift count, **321**
- shut down to enter power-down mode, **318**
- SIMD, **17**
- simultaneous addressing of doublewords, **353**
- SIR, **320**
- software-initiated reset, **320**
- stores
  - block store, **328**
  - floating point, *See* instructions: floating point
  - integer, **113, 323**
  - integer (except doubleword), **323**
  - integer into alternate space, **324, 516**
  - partial, **341**
  - unsigned byte, **163**
  - unsigned byte to alternate space, **261**
  - unsigned bytes, **260**
- swap R register, **352, 353**
- synthetic (for assembly language programmers), **598–600**
- tagged addition, **355**
- test-and-set, **401**
- timing, **145**
- trap on integer condition codes, **358**
- write hyperprivileged register, **372**
- write privileged register, **374**
- write state register, **370**
- integer unit (IU)
  - condition codes, **74**
  - definition, **11**
  - description, **26**
- internal\_processor\_error* exception, **487**
- interrupt
  - enable (ie) field of PSTATE register, **448, 451**
  - level, **101**
  - request, **11, 33, 437**
- interrupt\_level\_14* exception, **81, 487**
  - and SOFTINT.int\_level, **81**
  - and STICK\_CMPR.stick\_cmpr, **85**
  - and TICK\_CMPR.tick\_cmpr, **83**
- interrupt\_level\_15* exception, **487**
  - and SOFTINT.int\_level, **81**
- interrupt\_level\_n* exception, **448, 487**
  - and SOFTINT register, **80**
  - and SOFTINT.int\_level, **81**
- inter-strand operation, **11**
- intra-strand operation, **11**
- invalid accrued (nva) bit of aexc field of FSR register, **69**
- invalid ASI
  - and *data\_access\_exception*, **483**
- invalid current (nvc) bit of cexc field of FSR

- register, 69, 383, 384
- invalid\_exception* exception, 229
- invalid\_fp\_register floating-point trap type, 171, 172, 182, 183, 185, 191, 197, 200, 228, 233
- INVALW instruction, 238
- iprefetch synthetic instruction, 598
- ISA, 11
- ISA, *See* instruction set architecture
- I-SFSR register, *See* SFSR register
- issue unit, 393, 393
- issued, 12
- italic font, in assembly language syntax, 591
- IU, 12
- ixc synthetic instructions, 600
- IXX>data\_access\_exception (invalid ASI)
  - with load alternate instructions, 269

## J

- jmp synthetic instruction, 598
- JMPL instruction, 239
  - computing target address, 31
  - does not change CWP, 53
  - mem\_address\_not\_aligned* exception, 487
  - reexecuting trapped instruction, 311
- jump and link, *See* JMPL instruction

## L

- LD instruction (SPARC V8), 240
- LDBLOCKF instruction, 245, 429
- LDD instruction (SPARC V8 and V9), 267
- LDDA instruction, 428
- LDDA instruction (SPARC V8 and V9), 269
- LDDF instruction, 114, 249, 487
- LDDF\_mem\_address\_not\_aligned* exception, 487
  - address not doubleword aligned, 575
  - address not quadword aligned, 576
  - LDDF/LDDFA instruction, 114
  - load instruction with partial store ASI and misaligned address, 254
  - with load instructions, 250, 253, 430
  - with store instructions, 336, 430
- LDDF\_mem\_not\_aligned* exception, 60
- LDDFA instruction, 252, 343
  - alignment, 114
  - ASIs for fp load operations, 430
  - behavior with partial store ASIs, 250–??, 254, 254–??, 430–??

- causing *LDDF\_mem\_address\_not\_aligned* exception, 114, 487
- for block load operations, 429
- reading from a CMP register, 518
- used with ASIs, 429
- LDF instruction, 60, 249
- LDFA instruction, 60, 252
- LDFSR instruction, 61, 63, 64, 256, 486
- LDQF instruction, 249, 491
- LDQF\_mem\_address\_not\_aligned* exception, 491
  - address not quadword aligned, 576
  - LDQF/LDQFA instruction, 115
  - with load instructions, 253
- LDQFA instruction, 252
- LDSB instruction, 240
- LDSBA instruction, 242
- LDSH instruction, 240
- LDSHA instruction, 242
- LDSHORTF instruction, 258
- LDSTUB instruction, 113, 260, 261, 401, 402
  - and *data\_access\_exception* (noncacheable page) exception, 483
  - hardware primitives for mutual exclusion of LDSTUB, 400
- LDSTUBA instruction, 260, 261
  - alternate space addressing, 29
  - and *data\_access\_exception* exception, 483
  - hardware primitives for mutual exclusion of LDSTUBA, 400
- LDSW instruction, 240
- LDSWA instruction, 242
- LDTW instruction, 55, 114
- LDTW instruction (deprecated), 266
- LDTWA instruction, 55, 114
- LDTWA instruction (deprecated), 268
- LDTX instruction, 425
- LDTXA instruction, 116, 118, 263, 426
  - access alignment, 114
  - access size, 114
  - and *data\_access\_exception* (noncacheable page) exception, 483
- LDUB instruction, 240
- LDUBA instruction, 242
- LDUH instruction, 240
- LDUHA instruction, 242
- LDUW instruction, 240
- LDUWA instruction, 242
- LDX instruction, 240
- LDXA instruction, 242, 270, 398, 516

- reading from a CMP register, 518
- LDXFSR instruction, 61, 63, 64, **249**, 256, **315**, 486
- leaf procedure
  - modifying windowed registers, 130
- little-endian byte order, **12**, 28, 95
- load
  - block, *See* block load instructions
  - floating-point from alternate space
    - instructions, **252**
  - floating-point instructions, **249**, 256
  - from alternate space, 29, 74, 120, 516
  - instructions, **12**
  - instructions accessing memory, 113
  - nonfaulting, 393
  - short floating-point, *See* short floating-point load instructions
- LoadLoad MEMBAR relationship, 272
- LoadLoad MEMBAR relationship, 402
- LoadLoad predefined constant, 596
- loads
  - nonfaulting, 405
- load-store alignment, 28, **114**, 389
- load-store instructions
  - compare and swap, **163**
  - definition, **12**
  - and *fast\_data\_access\_protection* exception, 484
  - load-store unsigned byte, 163, **260**, 352, 353
  - load-store unsigned byte to alternate space, **261**
  - memory access, 27
  - swap R register with alternate space
    - memory, **353**
  - swap R register with memory, 163, **352**
- LoadStore MEMBAR relationship, 272, 402
- LoadStore predefined constant, 596
- local registers, 49, 52, 305
- logical XOR instructions, **377**
- Lookaside predefined constant, 596
- LSTPARTIALF instruction, 430

## M

- machine state
  - after reset, 548, 552
  - in RED\_state, 548, 552
- manufacturer (manuf) field of VER register, **108**, 574
- mask number (mask) field of VER register, **108**
- MAXGL, 26, 49, **99**, **101**, 102
- maximum global levels maxgl field of VER
  - register, **108**

- maximum trap levels maxtl field of HVER
  - register, **108**
- MAXPGL, **99**, **101**
- MAXTL
  - and error\_state, 472
  - and MAXGL, 102
  - and RED\_state, 472
  - instances of HTSTATE register, 105
  - instances of TNPC register, 91
  - instances of TPC register, 90
  - instances of TSTATE register, 92
  - instances of TT register, 93
  - non-reset trap, 441
- may (keyword), **12**
- mem\_address\_not\_aligned* exception, **487**
- JMPL instruction, 239
- LDTXA, 426, 427, 429
- load instruction with partial store ASI and
  - misaligned address, 254
- register update policy, 510
- RETURN, 312
- when recognized, 165
- with CASA instruction, 164
- with compare instructions, 165
- with load instructions, 114–115, 240, 241, 243, 250, 256, 267, 269, 270, 429, 430
- with store instructions, 114–115, 323, 324, 326, 337, 340, 347, 349, 429, 430
- with swap instructions (deprecated), 352, 354

- MEMBAR
  - #Sync
    - semantics, 274
- instruction
  - atomic operation ordering, 402
  - FLUSH instruction, 186, 404
  - functions, **271**, **401–403**
  - memory ordering, 273
  - memory synchronization, 122
  - side-effect accesses, 388
  - STBAR instruction, 273
  - write to error steering register, 537
- mask encodings
  - #LoadLoad, 272, 402
  - #LoadStore, 272, 402
  - #Lookaside, 272, 403
  - #MemIssue, 272, 403
  - #StoreLoad, 272, 402
  - #StoreStore, 272, 402
  - #Sync, 272, 403

- predefined constants
  - #LoadLoad, 596
  - #LoadStore, 596
  - #Lookaside, 596
  - #MemIssue, 596
  - #StoreLoad, 596
  - #StoreStore, 596
  - #Sync, 596
- MEMBAR
  - #Lookaside, 398
  - #StoreLoad, 398
- membar\_mask*, 596
- MemIssue predefined constant, 596
- memory
  - access instructions, 27, 113
  - alignment, 389
  - atomic operations, 400
  - atomicity, 577
  - cached, 386
  - coherence, 388, 577
  - coherency unit, 389
  - data, 403
  - instruction, 403
  - location, 386
  - models, 385
  - ordering unit, 389
  - real, 386
  - reference instructions, data flow order
    - constraints, 394
  - synchronization, 273
  - virtual address, 386
  - virtual address 0, 405
- Memory Management Unit
  - definition, 12
- Memory Management Unit, *See* MMU
- memory model
  - mode control, 397
  - partial store order (PSO), 396
  - relaxed memory order (RMO), 273, 396
  - sequential consistency, 397
  - strong, 397
  - total store order (TSO), 273, 396, 397
  - weak, 397
- memory model (mm) field of PSTATE register, 95
- memory order
  - pending transactions, 395
  - program order, 393
- memory\_model (mm) field of PSTATE register, 397
- memory-mapped I/O, 387
- mmask instruction field
  - store instructions, 327
- MMU
  - bypass, 407
  - definition, 12
  - page sizes, 503
- mode
  - hyperprivileged, 90, 392
  - MMU bypass, 407
  - nonprivileged, 24
  - privileged, 26, 90, 392
- motion estimation, 288
- MOVA instruction, 275
- MOVCC instruction, 275
- MOVcc instructions, 275
  - conditionally moving integer register
    - contents, 74
  - conditions for copying integer register
    - contents, 127
  - copying a register, 61
  - encoding of cond field, 563
  - encoding of opf\_cc instruction field, 564
  - used to avoid branches, 197, 277
- MOVCS instruction, 275
- move floating-point register if condition is true, 193
- move floating-point register if contents of integer register satisfy condition, 198
- MOVE instruction, 275
- move integer register if condition is satisfied
  - instructions, 275
- move integer register if contents of integer register satisfies condition instructions, 279
- move on condition instructions, 22
- MOVFA instruction, 276
- MOVFE instruction, 276
- MOVFG instruction, 276
- MOVFGE instruction, 276
- MOVFL instruction, 276
- MOVFLE instruction, 276
- MOVFLG instruction, 276
- MOVFN instruction, 276
- MOVFNE instruction, 276
- MOVFO instruction, 276
- MOVFU instruction, 276
- MOVFUE instruction, 276
- MOVFUG instruction, 276
- MOVFUGE instruction, 276
- MOVFUL instruction, 276
- MOVFULE instruction, 276



- 283, 364, 367
- floating point load instructions, 249, 252, 256
- flush instructions, 186, 190
- jump-and-link instruction, 239
- load instructions, 240, 260, 261, 266, 268
- logical operation instructions, 149, 286, 377
- PREFETCH, 291
- RETURN, 311
- opcode
  - definition, 13
  - format, 237
- opf instruction field
  - floating point arithmetic instructions, 172, 183, 207, 228
  - floating point compare instructions, 181
  - floating point conversion instructions, 229, 231, 234
  - floating point instructions, 171
  - floating point integer conversion, 185
  - floating point move instructions, 191
  - floating point negate instructions, 209
- opf\_cc instruction field
  - floating point move instructions, 193
  - move instructions, 564
- opf\_low instruction field, 193
- optional, 14
- OR instruction, 286
- ORcc instruction, 286
- ordering MEMBAR instructions, 122
- ordering unit, memory, 389
- ORN instruction, 286
- ORNcc instruction, 286
- OTHERW instruction, 287
- OTHERWIN (other windows) register, 87
  - FLUSHW instruction, 190
  - keeping consistent state, 89
  - modified by OTHERW instruction, 287
  - partitioned, 89
  - range of values, 85, 86, 578
  - rd designation for WRPR instruction, 374
  - rs1 designation for RDPR instruction, 303
  - SAVE instruction, 314
  - state after reset, 550
  - zeroed by INVALIDW instruction, 238
  - zeroed by NORMALW instruction, 285
- OTHERWIN register trap vectors
  - fill/spill traps, 492
  - handling spill/fill traps, 492
  - selecting spill/fill vectors, 492

- out register #7, 55
- out registers, 49, 52, 313
- overflow
  - bits
    - (v) in condition fields of CCR, 123
    - accrued (ofa) in aexc field of FSR register, 69
    - current (ofc) in cexc field of FSR register, 69
  - causing spill trap, 491
  - tagged add/subtract instructions, 123

## P

- p (predict) instruction field of branch
  - instructions, 157, 160, 161, 177
- P superscript on instruction name, 136
- PA\_watchpoint* exception, 408, 487
- packed-to-planar conversion, 219
- packing instructions, *See* FPACK instructions
- page fault, 297
- page table entry (PTE), *See* translation table entry (TTE)
- parity, even, 9
- parity, odd, 13
- park, 14
- parked, 14
- parking CMP core, 526
- partial store instructions, 341, 430
- partial store order (PSO) memory model, 396, 397
- partition identifier, 390
- partitioned
  - additions, 216
  - subtracts, 221
- P<sub>ASI</sub> superscript on instruction name, 136
- P<sub>ASR</sub> superscript on instruction name, 136
- PC (program counter) register, 15, 71, 76
  - after instruction execution, 111
  - CALL instruction, 162
  - changed by NOP instruction, 284
  - copied by JMPL instruction, 239
  - saving after trap, 32
  - set by DONE instruction, 166
  - set by RETRY instruction, 309
  - state after reset, 549
  - Trap Program Counter register, 90
- PCR
  - ASR summary, 71
- PCR register fields
  - priv, 78

- sl (select lower bits of PIC), 78
- st (system trace enable), 78
- su (select upper bits of PIC), 78
- ut (user trace enable), 78
- PDIST instruction, **288**
- pef field of PSTATE register
  - and access to GSR, 80
  - and *fp\_disabled* exception, 485
  - and FPop instructions, 132
  - branch operations, 175, 177
  - byte permutation, 156
  - comparison operations, 179, 182
  - data movement operations, 278
  - enabling FPU, 77
  - floating-point operations, 171, 172, 183, 185, 191, 196, 199, 207, 209, 228, 229, 231, 233, 234, 250, 252, 256, 258
  - integer arithmetic operations, 218, 223
  - logical operations, 224, 225, 227
  - memory operations, 247
  - read operations, 301, 319, 330
  - special addressing operations, 147, 173, 333, 339, 343, 345, 371
  - trap control, 451
- pef, *See* PSTATE, pef field
- Performance Control register, *See* PCR
- performance instrumentation counter register, *See* PIC register
- physical address, **14**
- physical core, **14**
- PIC (performance instrumentation counter)
  - register, **14, 79**
  - accessing, 488
  - ASR summary, 71
  - and PCR, 78
  - picl field, 79
  - picu field, 79
- pic\_overflow* exception, **487**
- PIL (processor interrupt level) register, **101**
  - interrupt conditioning, 448
  - interrupt request level, 452
  - interrupt\_level\_n*, 487
  - specification of register to read, 303
  - specification of register to write, 374
  - state after reset, 549
  - trap processing control, 451
- pipeline, **14**
- pipeline draining of CPU, 86, 90
- PIPT, **14**
- pixel instructions
  - compare, 178
  - component distance, **288, 288**
  - formatting, 210
- pixel registers for storing values, 236
- planar-to-packed conversion, 219
- P<sub>npt</sub> superscript on instruction name, **136**
- POPC instruction, **289**
- POR, **14**
- POR (*power\_on\_reset*), **546**
  - machine state changes, 548
- POR, *See* *power\_on\_reset* (POR)
- positive infinity, 383, 384
- power failure, 450, 478
- power\_on\_reset* (POR)
  - hard reset when POR pin activated, 546
- power\_on\_reset* (POR), **488, 546**
  - effect on HTSTATE, 106
  - effect on STICK register fields, 84
  - effect on TNPC register, 91
  - effect on TPC, 91
  - effect on TT register, 93
  - enabling/disabling virtual processors, 523, 524
  - full-processor reset, 532
  - hard reset, 525, 583
  - machine state changes, 548
  - and RED\_state, 441, 443, 474
  - setting TICK.npt, 75
  - STRAND\_ENABLE\_STATUS register, 535
  - system reset, 532
  - when initiated, 450
- P<sub>pic</sub> superscript on instruction name, **136**
- precise floating-point traps, 304
- precise trap, **445**
  - conditions for, 445
  - software actions, 445
  - vs. disrupting trap, 447
- predefined constants
  - LoadLoad, 596
  - lookaside, 596
  - MemIssue, 596
  - StoreLoad, 596
  - StoreStore, 596
  - Sync, 596
- predict bit, 161
- prefetch
  - for one read, 296
  - for one write, 297

- for several reads, 295
- for several writes, 296
- page, 297
- prefetch data instruction, **291**
- PREFETCH instruction, 113, **291**, 573
- prefetch\_fcn*, **596**
- PREFETCHA instruction, **291**, 573
  - and invalid ASI or VA, 483
- prefetchable, **15**
- priority of traps, 451, 469
- privilege violation
  - and *data\_access\_exception*, 482, 486
- privileged, **15**
  - mode, 26, 90, 392
  - registers, **90**
  - software, 25, 53, 64, 96, 121, 190, 454, 573
- privileged (priv) field of PCR register, 301
- privileged (priv) field of PSTATE register, **98**, **105**, 164, 167, 243, 247, 252, 253, 261, 269, 301, 324, 330, 336, 349, 353, 354, 371, 392, 488
- privileged mode, **15**
- privileged\_action* exception, **488**
  - accessing restricted ASIs, 391
  - PIC access, 79
  - register update policy, 510
  - restricted ASI access attempt, 121, 408
  - TICK register access attempt, 74
  - with CASA instruction, 164
  - with compare instructions, 165
  - with load alternate instructions, 243, 247, 253, 261, 269, 324, 330, 336, 349, 354, 371
  - with load instructions, 252
  - with RDasr instructions, 301
  - with read instructions, 301
  - with store instructions, 338
  - with swap instructions, 354
- privileged\_opcode* exception, **488**
  - DONE instruction, 167
  - RETRY instruction, 310
  - SAVED instruction, 315
  - with DONE instruction, 167, 303, 310, 375
  - with write instructions, 376
- processor, **15**
  - execute unit, 393
  - issue unit, 393, **393**
  - privilege-mode transition diagram, 440
  - reorder unit, 393
  - self-consistency, **394**
  - state diagram, 441
- processor cluster, *See* processor module
- processor interrupt level register, *See* PIL register
- processor state register, *See* PSTATE register
- processor states
  - error\_state*, 441, **444**, 470, 471, 472
    - entering, 478, 479, 480
  - execute\_state*, 470, 472
  - RED\_state*, 441, 442, 443, 457, 470, 472, 474, 476, 480
- processor states, *See* *error\_state*, *execute\_state*, and *RED\_state*
- program counter register, *See* PC register
- program counters, saving, 437
- program order, **393**, 394
- programming note, **4**
- PSO, *See* partial store order (PSO) memory model
- PSR register (SPARC V8), 371
- PSTATE register
  - fields
    - priv*
      - and access to PCR, 78
- PSTATE register
  - entering privileged execution mode, 437
  - restored by RETRY instruction, 166, 309
  - saved after trap, 437
  - saving after trap, 32
  - specification for RDPR instruction, 303
  - specification for WRPR instruction, 374
  - state after reset, 549
  - and TSTATE register, 92
- PSTATE register fields
  - ag*
    - unimplemented, 99
  - am*
    - CALL instruction, 162
    - description, **96**
    - masked/unmasked address, 166, 239, 309, 311
  - cle*
    - and implicit ASIs, 120
    - and PSTATE.tle, 95
    - description, **95**
  - ie*
    - description, 98, 99
    - enabling disrupting traps, 448
    - interrupt conditioning, 448
    - masking disrupting trap, 458
  - mm*
    - description, **95**

- implementation dependencies, 95, 96, 396, 577
- reserved values, 95
- pef**
  - and FPRS.fef, 96
  - description, 96
  - See also* pef field of PSTATE register
- priv**
  - access to register-window PR state registers, 90
  - accessing restricted ASIs, 391
  - description, 98
  - determining mode, 13, 15, 507
  - when processor in privileged mode, 105
- tle**
  - and PSTATE.cle, 95
  - description, 95
- PTE (page table entry), *See* translation table entry (TTE)

## Q

- quadword, 15**
  - alignment, 28, 114, 389
  - data format, 35
- quiet NaN (not-a-number), 62, 181, 182, 381

## R

- R register, 15**
  - #15, 55
  - special-purpose, 55
  - alignment, 267, 269
- RA\_watchpoint* exception, 462, 467
- rational quotient, 364
- R-A-W, *See* read-after-write memory hazard
- rcond instruction field**
  - branch instructions, 160
  - encoding of, 563
  - move instructions, 279
- rd (rounding), 15**
- rd instruction field, 17**
  - arithmetic instructions, 146, 158, 161, 163, 281, 283, 364, 367
  - floating point arithmetic, 172
  - floating point arithmetic instructions, 183, 207, 228
  - floating point conversion instructions, 229, 231, 234

- floating point integer conversion, 185
- floating point load instructions, 249, 252, 256
- floating point move instructions, 191, 193
- floating point negate instructions, 209
- floating-point instructions, 171
- jump-and-link instruction, 239
- load instructions, 240, 260, 261, 266, 268
- logical operation instructions, 149, 286, 377
- move instructions, 277, 279
- POPC, 289
- RDASI instruction, 70, 74, 299**
- RDAsr instruction, 299**
  - accessing I/O registers, 29
  - implementation dependencies, 300, 572
  - reading ASRs, 70
- RDCCR instruction, 70, 72, 299, 299**
- RDFPRS instruction, 71, 76, 299**
- RDGSR instruction, 71, 80, 299**
- RDHPR instruction, 103, 104, 106, 107, 302**
  - hyperprivileged registers read, 302
- RDPC instruction, 71, 299**
  - reading PC register, 76
- RDPCR instruction, 71, 299**
- RDPIC instruction, 71, 299, 488**
- RDPR instruction, 15, 71, 303**
  - accessing GL register, 102
  - accessing non-register-window PR state registers, 90
  - accessing register-window PR state registers, 85
  - and register-window PR state registers, 85
  - effect on TNPC register, 91
  - effect on TPC register, 91
  - effect on TSTATE register, 93
  - effect on TT register, 94
  - reading privileged registers, 90
  - reading PSTATE register, 94
  - reading the TICK register, 75
  - registers read, 303
- RDSOFTINT instruction, 71, 81, 299**
- RDSTICK instruction, 71, 84, 299**
- RDSTICK\_CMPR instruction, 71, 299**
- RDITICK instruction, 71, 75, 299**
- RDITICK\_CMPR instruction, 71, 299**
- RDY instruction, 72**
- read ancillary state register (RDAsr) instructions, 299**
- read state register instructions, 31**
- read-after-write memory hazard, 394**
- real memory, 386**

**RED\_state**, 16  
     catastrophic failure avoidance, 470  
     description, 440  
     entering, 443, 476, 577  
     entry conditions, 441  
     exiting, 105  
     red field of HPSTATE register, 441, 442, 443, 470, 471  
     restricted environment, 442  
     special trap processing, 474  
     trap processing, 443, 470, 472  
     trap table, 457  
     trap vector, 455, 577  
**RED\_state trap**, 16  
**RED\_state\_exception** exception, 488  
 reference MMU, 591  
*reg*, 592  
*reg\_or\_imm*, 597  
*reg\_plus\_imm*, 596  
*regaddr*, 597  
 register reference instructions, data flow order constraints, 394  
 register window, 49  
 register window management instructions, 32  
 register windows  
     clean, 87, 89, 129, 482, 491, 492, 493  
     fill, 53, 89, 129, 130, 306, 307, 315, 485, 491, 492, 493  
     management of, 24  
     overlapping, 52–54  
     spill, 53, 89, 129, 130, 131, 314, 315, 489, 491, 492, 493  
 registers  
     *See also* individual register (common) names  
     address space identifier (ASI), 392  
     ASI (address space identifier), 74  
     chip-level multithreading, *See* CMT  
     clean windows (CLEANWIN), 87  
     clock-tick (TICK), 488  
     current window pointer (CWP), 86  
     F (floating point), 379, 452  
     floating-point, 27  
         programming, 59  
     floating-point registers state (FPRS), 76  
     floating-point state (FSR), 61  
     general status (GSR), 80  
     GL (global level), 108  
     *global*, 22, 26, 49, 49, 49, 569  
     global level (GL), 101  
     HSTICK\_CMPR  
         and HINTP, 106  
     HSTICK\_CMPR, 109  
     HTSTATE (hyperprivileged trap state), 105  
     HVER (version register), 107  
     hyperprivileged, 103  
     IER (SPARC V8), 371  
     *in*, 49, 52, 313  
     *local*, 49, 52  
     next program counter (NPC), 76  
     other windows (OTHERWIN), 87  
     *out*, 49, 52, 313  
     *out* #7, 55  
     performance control (PCR), 78  
     performance instrumentation counter (PIC), 79  
     pixel storage registers, 236  
     processor interrupt level (PIL)  
         and PIC, 79  
         and PIC counter overflow, 79  
         and SOFTINT, 81  
         and STICK\_CMPR, 85  
         and TICK\_CMPR, 83  
     processor interrupt level (PIL), 101  
     program counter (PC), 76  
     PSR (SPARC V8), 371  
     R register #15, 55  
     renaming mechanism, 394  
     restorable windows (CANRESTORE), 87, 87  
     savable windows (CANSAVE), 86  
     scratchpad  
         hyperprivileged, 432  
         privileged, 431  
     SOFTINT, 71  
     SOFTINT\_CLR pseudo-register, 71, 82  
     SOFTINT\_SET pseudo-register, 71, 82  
     STICK, 83  
     STICK\_CMPR  
         and HINTP, 106  
         ASR summary, 71  
         *int\_dis* field, 81, 85  
         *stick\_cmpr* field, 85  
         and system software trapping, 84  
     TBR (SPARC V8), 371  
     TICK, 74  
     TICK\_CMPR  
         *int\_dis* field, 81, 83  
         *tick\_cmpr* field, 83  
     TICK\_CMPR, 71, 83  
     TL (trap level), 108

- trap base address (TBA), **94**
- trap base address, *See* registers: TBA
- trap level (TL), **99**
- trap level, *See* registers: TL
- trap next program counter (TNPC), **91**
- trap next program counter, *See* registers: TNPC
- trap program counter (TPC), **90**
- trap program counter, *See* registers: TPC
- trap state (TSTATE), **92**
- trap state, *See* registers: TSTATE
- trap type (TT), **93, 456**
- trap type, *See* registers: TT
- VA\_WATCHPOINT, **490**
- visible to software in privileged mode, **90–103**
- WIM (SPARC V8), **371**
- window state (WSTATE), **88**
- window state, *See* registers: WSTATE
- Y (32-bit multiply/divide), **72**
- relaxed memory order (RMO) memory model, **273, 396**
- renaming mechanism, register, **394**
- reorder unit, **393**
- reordering instruction, **394**
- reserved, **16**
  - fields in instructions, **145**
  - register field, **48**
- reset
  - externally\_initiated\_reset* (XIR), **440, 441, 443, 450, 456, 474, 478, 484, 533, 547**
  - power\_on\_reset* (POR)
    - enabling/disabling virtual processors, **523, 524**
    - machine state changes, **548**
    - STRAND\_ENABLE\_STATUS register, **535**
  - power\_on\_reset* (POR), **441, 443, 450, 474, 488, 532, 546**
  - power-on, **75**
  - processing, **442**
  - request, **488**
  - reset trap, **75, 93, 447, 449**
  - software\_initiated\_reset* (SIR), **440, 441, 443, 449, 450, 456, 470, 479, 488, 532, 548**
  - trap, **571**
  - trap vector address, *See* RSTVaddr
  - warm\_reset* (WMR)
    - and STRAND\_ENABLE register, **525**
    - enabling/disabling virtual processors, **523, 524**
    - machine state changes, **548**
    - warm\_reset* (WMR), **547**
    - watchdog* (WDR), **532**
    - watchdog\_reset* (POR), **443**
    - watchdog\_reset* (WDR)
      - and *guest\_watchdog*, **439**
    - watchdog\_reset* (WDR), **474, 478, 490, 532, 548**
    - XIR, **533**
  - reset trap, **16**
  - Reset, Error, and Debug state, *See* RED\_state
  - restartable deferred trap, **446**
  - restorable windows register, *See* CANRESTORE register
  - RESTORE instruction, **53, 305–306**
    - actions, **129**
    - and current window, **55**
    - decrementing CWP register, **52**
    - fill trap, **485, 491**
    - followed by SAVE instruction, **53**
    - managing register windows, **32**
    - operation, **305**
    - performance trade-off, **305, 313**
    - and restorable windows (CANRESTORE) register, **87**
    - restoring register window, **305**
    - role in register state partitioning, **89**
  - restore synthetic instruction, **599**
  - RESTORED instruction, **130, 307**
    - creating inconsistent window state, **307**
    - fill handler, **306**
    - fill trap handler, **130, 493**
    - register window management, **32**
  - restricted, **16**
  - restricted address space identifier, **121**
  - restricted ASI, **391, 407**
  - resumable\_error* exception, **488**
  - ret/retl synthetic instructions, **598**
  - RETRY instruction, **309**
    - and restartable deferred traps, **447**
    - effect on HTSTATE, **106**
    - effect on TNPC register, **91**
    - effect on TPC register, **91**
    - effect on TSTATE register, **93**
    - executed in RED\_state, **442**
    - generating *illegal\_instruction* exception, **486**
    - modifying CCR.xcc, **73**
    - reexecuting trapped instruction, **493**
    - restoring gl value in GL, **103**
    - return from trap, **437**
    - returning to instruction after trap, **449**

- target address, return from privileged traps, 31
- RETURN instruction, **311–312**
  - computing target address, 31
  - fill trap, 485
  - mem\_address\_not\_aligned* exception, 487
  - operation, 311
  - reexecuting trapped instruction, 311
- RETURN vs. RESTORE instructions, 311
- RMO, **16**
- RMO, *See* relaxed memory order (RMO) memory model
- rounding
  - for floating-point results, 62
  - in signed division, 365
- rounding direction (rd) field of FSR register, 172, 183, 207, 228, 229, 231, 233, 234
- routine, nonleaf, 239
- rs1 instruction field, **17**
  - arithmetic instructions, 146, 158, 161, 163, 281, 283, 364, 367
  - branch instructions, 160
  - floating point arithmetic instructions, 172, 183, 207
  - floating point compare instructions, 181
  - floating point load instructions, 249, 252, 256
  - flush memory instruction, 186
  - jump-and-link instruction, 239
  - load instructions, 240, 260, 261, 266, 268
  - logical operation instructions, 149, 286, 377
  - move instructions, 279
  - PREFETCH, 291
  - RETURN, 311
- rs2 instruction field, **17**
  - arithmetic instructions, 146, 158, 161, 163, 281, 283, 286, 364, 367
  - floating point arithmetic instructions, 172, 183, 207, 228
  - floating point compare instructions, 181
  - floating point conversion instructions, 229, 231, 234
  - floating point instructions, 171
  - floating point integer conversion, 185
  - floating point load instructions, 249, 252, 256
  - floating point move instructions, 191, 193
  - floating point negate instructions, 209
  - flush memory instruction, 186
  - jump-and-link instruction, 239
  - load instructions, 240, 266, 268
  - logical operation instructions, 149, 377

- move instructions, 277, 279
- POPC, 289
- PREFETCH, 291
- RSTVADDR, 443, 455, 456, 457, 476, 477, 478, 479, 480, 481, 549, 577
- RTO, **17**
- RTS, **17**

## S

- savable windows register, *See* CANSERVE register
- SAVE instruction, 52, **313**
  - actions, 129
  - after RESTORE instruction, 311
  - clean\_window* exception, 482, 492
  - and current window, 55
  - decrementing CWP register, 52
  - effect on privileged state, 314
  - leaf procedure, 239
  - and *local/out* registers of register window, 53
  - managing register windows, 32
  - no clean window available, 87
  - number of usable windows, 87
  - operation, 313
  - performance trade-off, 313
  - role in register state partitioning, 89
  - and savable windows (CANSERVE) register, 86
  - spill trap, 489, 491, 493
- save synthetic instruction, 599
- SAVED instruction, 130, **315**
  - creating inconsistent window state, 315
  - register window management, 32
  - spill handler, 314, 315
  - spill trap handler, 130, 493
- scaling of the coefficient, 202
- scratchpad registers
  - hyperprivileged, **432**
  - privileged, **431**
  - state after reset, 551
- SDIV instruction, 72, **364**
- SDIVcc instruction, 72, **364**
- SDIVX instruction, **283**
- self-consistency, processor, 394
- self-modifying code, 186, 187
- sequencing MEMBAR instructions, 122
- sequential consistency memory model, **397**
- service processor, **17**
- SETHI instruction, **123, 317**
  - creating 32-bit constant in R register, 30

- and NOP instruction, 284
- with rd = 0, 317
- set*n* synthetic instructions, 599
- SFSR register, 17
  - data\_access\_exception*, 482
  - fault type field (ft), 482
  - state after reset, 551
  - update policy, 510
- shall (keyword), 17
- shared memory, 385
- shift count encodings, 321
- shift instructions, 30
- shift instructions, 122, 321
- short floating-point load and store instructions, 430
- short floating-point load instructions, 258
- short floating-point store instructions, 344
- should (keyword), 17
- SHUTDOWN instruction, 318
- SIAM instruction, 319
- side effect
  - accesses, 387
  - definition, 17
  - I/O locations, 386
  - instruction prefetching, 388
  - real memory storage, 386
  - visible, 387
- signalling NaN (not-a-number), 62, 181, 182, 231, 381
- signed integer data type, 35
- signx synthetic instructions, 599
- SIMD, 17
  - instruction data formats, 43–45
- simm10 instruction field
  - move instructions, 279
- simm11 instruction field
  - move instructions, 277
- simm13 instruction field
  - floating point
    - load instructions, 249
- simm13 instruction field
  - arithmetic instructions, 281, 283, 286, 364, 367
  - floating point load instructions, 252, 256
  - flush memory instruction, 186
  - jump-and-link instruction, 239
  - load instructions, 240, 260, 261, 266, 268
  - logical operation instructions, 149, 377
  - POPC, 289
  - PREFETCH, 291
  - RETURN, 311
  - single instruction/multiple data, *See* SIMD
  - SIR, 17
  - SIR (*software\_initiated\_reset*), 548
  - SIR instruction, 320
    - affecting virtual processor, 548
    - causing *software\_initiated\_reset* exception, 450, 488
    - and trap priority, 469
    - use by supervisor software, 479
  - SIR, *See software\_initiated\_reset (SIR)*
  - SLL instruction, 321
  - SLLX instruction, 321
  - SMUL instruction, 72, 367
  - SMULcc instruction, 72, 367
  - snooping, 17
  - SOFTINT register, 71, 80
    - clearing, 497
    - clearing of selected bits, 82
    - communication from nucleus code to kernel
      - code, 496
    - scheduling interrupt vectors, 495, 496
    - setting, 497
    - state after reset, 550
  - SOFTINT register fields
    - int\_level, 81
    - sm (stick\_int), 81
    - tm (tick\_int), 81, 83
  - SOFTINT\_CLR pseudo-register, 71, 82
  - SOFTINT\_SET pseudo-register, 71, 82, 82
  - software
    - nucleus, 13
    - software translation table, 504
    - software trap, 359, 454, 457
    - software trap number (SWTN), 359
    - software, nonprivileged, 76
    - software\_initiated\_reset* (SIR), 479, 488, 548
      - entering error\_state, 440
      - entering RED\_state, 441
      - and MAXTL, 443
      - per-strand reset, 532
      - RED\_state trap processing, 474
      - RED\_state trap vector, 456
      - SIR instruction, 320, 449
      - and virtual processor, 548
      - virtual processor trap processing, 470
      - when TL = MAXTL, 470
    - software\_trap\_number*, 597
    - source operands, 216, 221
  - SPA

- ASI\_TWIN\_DW\_NUCLEUS, 433
- SPARC V8 compatibility
  - LD, LDUW instructions, 240
  - operations to I/O locations, 388
  - read state register instructions, 300
  - STA instruction renamed, 325
  - STBAR instruction, 273, 327
  - STD instruction, 347
  - STDA instruction, 349
  - tagged subtract instructions, 363
  - UNIMP instruction renamed, 235
  - window\_overflow* exception superseded, 485
  - write state register instructions, 371
- SPARC V9
  - compliance, 14
  - features, 22
- SPARC V9 Application Binary Interface (ABI), 24
- special trap, renamed, 441
- special traps, 441, 457
- speculative load, 17
- spill register window, 489
  - FLUSH instruction, 131
  - overflow/underflow, 53
  - RESTORE instruction, 129
  - SAVE instruction, 89, 129, 313, 491
  - SAVED instruction, 130, 315, 493
  - selection of, 492
  - trap handling, 493
  - trap vectors, 314, 493
  - window state, 89
- spill\_n\_normal* exception, 314, 489
  - and FLUSHW instruction, 190
- spill\_n\_other* exception, 314, 489
  - and FLUSHW instruction, 190
- SRA instruction, 321
- SRAX instruction, 321
- SRL instruction, 321
- SRLX instruction, 321
- stack frame, 313
- state registers (ASRs), 70–85
- STB instruction, 323
- STBA instruction, 324
- STBAR instruction, 273
- STBAR instruction, 300, 370, 394, 402
- STBLOCKF instruction, 328, 429
- STDF instruction, 114, 332, 489
- STDF\_mem\_address\_not\_aligned* exception, 489
  - and store instructions, 333, 337
  - STDF/STDFA instruction, 114
- STDFA instruction, 335
  - alignment, 114
  - ASIs for fp store operations, 430
  - causing *data\_access\_exception* exception, 430
  - causing *mem\_address\_not\_aligned* or *illegal\_instruction* exception, 430
  - causing *STDF\_mem\_address\_not\_aligned* exception, 114, 489
  - for block load operations, 429
  - for partial store operations, 430
  - used with ASIs, 429
  - writing to a CMP register, 518
- STF instruction, 332
- STFA instruction, 335
- STFSR instruction, 61, 63, 64, 486
- STH instruction, 323
- STHA instruction, 324
- STICK register, 71, 75, 83
  - and *hstick\_match* exception, 485
  - counter field, 83, 84
  - fields after power-on reset trap, 84
  - npt field, 75, 83
  - RDSTICK instruction, 299
  - state after reset, 550
  - while virtual processor is parked, 526
- STICK\_CMPR register, 71, 84
  - and HINTP, 106
  - int\_dis field, 81, 85
  - RDSTICK\_CMPR instruction, 299
  - state after reset, 551
  - stick\_cmpr field, 85
- store
  - block, *See* block store instructions
  - partial, *See* partial store instructions
  - short floating-point, *See* short floating-point store instructions
- store buffer
  - merging, 387
- store floating-point into alternate space
  - instructions, 335
- store instructions, 17, 113, 484
- store\_error* exception, 489
- StoreLoad MEMBAR relationship, 272, 402
- StoreLoad predefined constant, 596
- stores to alternate space, 29, 74, 120
- StoreStore MEMBAR relationship, 272, 402
- StoreStore predefined constant, 596
- STPARTIALF instruction, 341
- STQF instruction, 115, 332, 491

**STQF\_mem\_address\_not\_aligned** exception, **491**  
     STQF/STQFA instruction, 115  
 STQFA instruction, 115, **335**, 335  
 strand, **18**  
 STRAND\_AVAILABLE register, 519, **523**, 523, 525, 526  
     state after reset, 552  
 STRAND\_ENABLE register, **524**  
     state after reset, 552  
 STRAND\_ENABLE\_STATUS register, **524**  
     state after reset, 552  
 STRAND\_ID register, **519**  
     state after reset, 553  
 STRAND\_INTR\_ID register, 501, **520**, 542  
     state after reset, 553  
 STRAND\_RUNNING register, 526, **527**  
     simultaneous updates, 528  
     state after reset, 552  
 STRAND\_RUNNING\_RW pseudo-register, **527**, **528**  
 STRAND\_RUNNING\_STATUS register, 526, **530**  
     Parked or Unparked status, 531  
     state after reset, 552  
 STRAND\_RUNNING\_W1C pseudo-register, **527**, **528**  
 STRAND\_RUNNING\_W1S pseudo-register, **527**, **528**  
 strong consistency memory model, 397  
 strong ordering, 397  
 Strong Sequential Order, 398  
 strongly ordered page, illegal access to, 483  
 STSHORTF instruction, 344  
 STTW instruction, 55, 114  
 STTW instruction (deprecated), **346**  
 STTWA instruction, 55, 114  
 STTWA instruction (deprecated), **348**  
 STW instruction, 323  
 STWA instruction, 324  
 STX instruction, 323  
 STXA instruction, **324**  
     accessing CMP-specific registers, 516  
     accessing nontranslating ASIs, 349  
     *mem\_address\_not\_aligned* exception, 324  
     referencing internal ASIs, 398  
     writing to a CMP register, 518  
 STXFSR instruction, 61, 63, 64, 332, 486  
 SUB instruction, **351**, 351  
 SUBC instruction, **351**, 351  
 SUBcc instruction, 122, **351**, 351  
 SUBCcc instruction, **351**, 351  
 subnormal number, **18**  
 subtract instructions, **351**  
 superscalar, **18**  
 supervisor software  
     accessing special protected registers, 29  
     definition, **18**  
     forcing processing into RED\_state, 470  
     use of SIR trap, 479  
 suspend, **18**  
*sw\_recoverable\_error* exception, **489**, 490  
 SWAP instruction, **352**  
     accessing doubleword simultaneously with other  
         instructions, 353  
     and *data\_access\_exception* (noncacheable page)  
         exception, 483  
     hardware primitive for mutual exclusion, 400,  
         401  
     identification of R register to be exchanged, 113  
     in multiprocessor system, 260, 261  
     memory accessing, 352  
     ordering by MEMBAR, 402  
 swap R register  
     bit contents, 163  
     with alternate space memory instructions, **353**  
     with memory instructions, **352**  
 SWAPA instruction, **353**  
     accessing doubleword simultaneously with other  
         instructions, 353  
     alternate space addressing, 29  
     and *data\_access\_exception* (noncacheable page)  
         exception, 483  
     hardware primitive for mutual exclusion, 400  
     in multiprocessor system, 260, 261  
     ordering by MEMBAR, 402  
 SWTN (software trap number), **359**  
 Sync predefined constant, 596  
 synchronization, **274**  
 synchronization, **18**  
 Synchronous Fault Address register (SFAR), **17**  
 synchronous fault status register, *See* SFSR register  
 synthetic instructions  
     mapping to SPARC V9 instructions, **598–600**  
     for assembly language programmers, 598  
     mapping  
         **bclrg**, 600  
         **bset**, 600  
         **btog**, 600  
         **btst**, 600  
         **call**, 598  
         **casn**, 600  
         **clrn**, 600  
         **cmp**, 598

- dec, 600
- deccc, 600
- inc, 600
- inccc, 600
- iprefetch, 598
- jmp, 598
- movn, 600
- neg, 600
- not, 599
- restore, 599
- ret/retl, 598
- save, 599
- setn, 599
- signx, 599
- tst, 598
- vs. pseudo ops, 598
- system clock-tick register (STICK), 83
- system software, 488
  - accessing memory space by server program, 391
  - ASIs allowing access to memory space, 392
  - FLUSH instruction, 188, 404
  - processing exceptions, 391
  - trap types from which software must recover, 64
- System Tick Compare register, *See* STICK\_CMPR register
- System Tick register, *See* STICK register

## T

- TA instruction, 358, 563
- TADDcc instruction, 123, 355
- TADDccTV instruction, 123, 489
- tag overflow, 123
- tag\_overflow* exception, 123, 355, 356, 357, 361, 363
- tag\_overflow* exception (deprecated), 489
- tagged arithmetic, 123
- tagged arithmetic instructions, 30
- tagged word data format, 35
- tagged words, 35
- TBA (trap base address) register, 94, 439
  - establishing table address, 32, 437
  - initialization, 453
  - specification for RDPR instruction, 303
  - specification for WRPR instruction, 374
  - state after reset, 549
  - trap behavior, 19
- TBR register (SPARC V8), 371
- TCC instruction, 358
- Tcc instructions, 358

- at TL > 0, 454
- causing trap, 437
- causing trap to privileged trap handler, 457
- CCR register bits, 73
- generating *htrap\_instruction* exception, 485
- generating *illegal\_instruction* exception, 486
- generating *trap\_instruction* exception, 489
- opcode maps, 559, 563, 564
- programming uses, 360
- trap table space, 32
- vector through trap table, 437
- TCS instruction, 358, 563
- TE instruction, 358, 563
- termination deferred trap, 445
- test-and-set instruction, 401
- TG instruction, 358, 563
- TGE instruction, 358, 563
- TGU instruction, 358, 563
- thread, 18
- TICK register, 71
  - controlling access to timing information, 76
  - counter field, 75, 574
  - fields after power-on reset trap, 75
  - inaccuracies between two readings of, 574
  - npt field, 76
  - specification for RDPR instruction, 303
  - specification for WRPR instruction, 374
  - state after reset, 550
  - while virtual processor is parked, 526, 587
- TICK\_CMPR register, 71, 83
  - int\_dis field, 81, 83
  - state after reset, 550
  - tick\_cmpr field, 83
- timer registers, *See* TICK register *and* STICK register
- timing of instructions, 145
- tininess (floating-point), 69
- TL (trap level) register, 99, 439
  - affect on privilege level to which a trap is delivered, 452
  - and implicit ASIs, 120
  - displacement in trap table, 437
  - executing RESTORED instruction, 307
  - executing SAVED instruction, 315
  - indexing for WRHPR instruction, 372
  - indexing for WRPR instruction, 374
  - indexing hyperprivileged register after RDHPR, 302
  - indexing privileged register after RDPR, 303
  - setting register value after WRHPR, 372

- setting register value after WRPR, 374
- specification for RDPR instruction, 303
- specification for WRPR instruction, 374
- state after reset, 550
- and TBA register, 453
- and TPC register, 90
- and TSTATE register, 92, 105
- and TT register, 93
- use in calculating privileged trap vector
  - address, 453
  - and VER.maxtl, 108
  - and WSTATE register, 88
- TL instruction, 358, 563
- TLB, **18**
  - and 3-dimensional arrays, 153
  - definition, **18**
  - hit, **18**
  - miss, **18**
    - handler, 504
    - MMU behavior, 504
    - reloading TLB, 504, 508
- TLE instruction, 358, 563
- TLEU instruction, 358, 563
- TN instruction, 358, 563
- TNE instruction, 358, 563
- TNEG instruction, 358, 563
- TNPC (trap next program counter) register, **91**
  - after *async\_data\_error* disrupting trap, 489
  - saving NPC, 445
  - specification for RDPR instruction, 303
  - specification for WRPR instruction, 374
  - state after reset, 550
- total order, **396**
- total store order (TSO) memory model, 95, 273, 387, **396, 397, 397**
- TPC (trap program counter) register, **18, 90**
  - address of trapping instruction, 304
  - after *async\_data\_error* disrupting trap, 489
  - number of instances, 90
  - specification for RDPR instructions, 303
  - specification for WRPR instruction, 374
  - state after reset, 550
- TPOS instruction, 358, 563
- translating ASIs, 408
- Translation Lookaside Buffer, *See* TLB
- Translation Table Entry, *See* TTE
- trap
  - See also* exceptions *and* traps
  - noncacheable accesses, 388
  - when taken, 18
- trap enable mask (tem) field of FSR register, 451, 452, 570
- trap handler
  - for global registers, 103
  - hyperprivileged mode, 456
  - privileged mode, 456
  - regular/nonfaulting loads, 13
  - returning from, 166, 309
  - user, 65, 383
- trap level register, *See* TL register
- trap next program counter register, *See* TNPC register
- trap on integer condition codes instructions, **358**
- trap program counter register, *See* TPC register
- trap state register, *See* TSTATE register
- trap type (TT) register, **456**
- trap type register, *See* TT register
- trap\_instruction* (ISA) exception, 359, 360, **489**
- trap\_level\_zero* exception, 105, **489**
  - state after reset, 549
  - with WRHPR instructions, 373
  - with write instructions, 376
- trap\_little\_endian (tle) field of PSTATE register, **95**
- traps, **19**
  - See also* exceptions *and* individual trap names
  - categories
    - deferred, 445, **445**, 447
    - disrupting, 445, **447**, 449
    - precise, 445, **445**, 447
    - priority, 451, 469
    - reset, 93, 445, 447, **449**, 449, 470, 571
    - restartable
      - implementation dependency, 447
      - restartable deferred, **446**
      - termination deferred, **445**
  - caused by undefined feature/behavior, 19
  - causes, **33**, 33
  - definition, 32, 438
  - hardware, 457
  - hardware stack, 23
  - level specification, 99
  - model stipulations, **450**
  - nested, 23
  - normal, **13**, **441**, 456, 471, 474
  - processing, 470
  - software, 359, 454, 457
  - software\_initiated\_reset* (SIR), 474
  - special, **441**, 457
  - stack, 472

- vector address, specifying, 94, 107
- vector, *RED\_state*, 455
- TSB, **19**, **508**
  - cacheability, 509
  - caching, 509
  - indexing support, 508
  - organization, 509
- TSO, **19**
- TSO, *See* total store order (TSO) memory model
- tst* synthetic instruction, 598
- TSTATE (trap state) register, **92**
  - DONE instruction, 166, 309
  - registers saved after trap, 32
  - restoring GL value, 103
  - specification for RDPR instruction, 303
  - specification for WRPR instruction, 374
  - state after reset, 550
- tstate*, *See* trap state (TSTATE) register
- TSUBcc instruction, 123, 361
- TSUBccTV instruction, 123, 489
- TT (trap type) register, **93**
  - and privileged trap vector address, 453, 454
  - reserved values, 571
  - specification for RDPR instruction, 303
  - specification for WRPR instruction, 374
  - state after reset, 549
  - and Tcc instructions, 360
  - transferring trap control, 456
  - trap type recorded after
    - RED\_state\_exception*, 488
  - window spill/fill exceptions, 88
  - WRHPR instruction, 372
  - WRPR instruction, 374
- TTE, **19**
  - context ID field, **506**
  - cp (cacheability) field, 386
  - cp field, 483, 507, **507**
  - cv field, 507, **507**
  - e field, 387, 405, 483, **507**
  - ie field, **506**
  - indexing support, 508
  - nfo field, 405, 483, **506**, **507**
  - p field, 482, **507**
  - size field, **508**
  - soft2 field, **506**
  - SPARC V8 equivalence, 505
  - taddr field, **506**
  - v field, **506**
  - va\_tag field, **506**

- w field, **508**
- TVC instruction, 358, 563
- TVS instruction, 358, 563
- typewriter font, in assembly language syntax, 591

## U

- UDIV instruction, 72, **364**
- UDIVcc instruction, 72, **364**
- UDIVX instruction, **283**
- UltraSPARC, previous ASIs
  - ASI\_NUCLEUS\_QUAD\_LDD, 433
  - ASI\_NUCLEUS\_QUAD\_LDD\_L, 433
  - ASI\_NUCLEUS\_QUAD\_LDD\_LITTLE, 433
  - ASI\_PHY\_BYPASS\_EC\_WITH\_EBIT\_L, 433
  - ASI\_PHYS\_BYPASS\_EC\_WITH\_EBIT, 433
  - ASI\_PHYS\_BYPASS\_EC\_WITH\_EBIT\_LITTLE, 433
  - ASI\_PHYS\_USE\_EC, 433
  - ASI\_PHYS\_USE\_EC\_L, 433
  - ASI\_PHYS\_USE\_EC\_LITTLE, 433
  - ASI\_QUAD\_LDD\_L, 433
  - ASI\_QUAD\_LDD\_LITTLE, 433
  - ASI\_QUAD\_LDD\_PHYS, 433
- UMUL instruction, 72, **367**
- UMULcc instruction, 72, **367**
- unassigned, **19**
- unconditional branches, 154, 158, 175, 177
- undefined, **19**
- underflow
  - bits of FSR register
    - accrued (*ufa*) bit of *aexc* field, **69**, 383
    - current (*ufc*) bit of *cexc*, **69**
    - current (*ufc*) bit of *cexc* field, 383
    - mask (*ufm*) bit of FSR.tem, **69**
    - mask (*ufm*) bit of tem field, 382
  - detection, 53
  - occurrence, 491
- unfinished\_FPop floating-point trap type, **65**, 172, 183, 208, 232, 233, 380
  - handling, 70
  - in normal computation, 64
  - results after recovery, 65
- UNIMP instruction (SPARC V8), 235
- unimplemented, **19**
- unimplemented\_FPop floating-point trap type, **65**, 171, 172, 182, 183, 185, 191, 197, 200, 208, 209, 230, 232, 233, 380
  - handling, 70

- result after recovery, 65
- unimplemented\_LDTW* exception, 267, 490
- unimplemented\_STTW* exception, 347, 490
- uniprocessor system, 19
- unpark, 19
- unparking CMP core, 526
- unrestricted, 19
- unrestricted ASI, 407
- unsigned integer data type, 35
- user application program, 19
- user trap handler, 65, 383

## V

- VA, 20
- VA\_watchpoint* exception, 490
- VA\_WATCHPOINT register, 490
- value clipping, *See* FPACK instructions
- value semantics of input/output (I/O)
  - locations, 386
- VER (version) register (SPARC V9), 108
- VER (version) register fields
  - impl, 63, 108
  - manuf, 108
  - mask, 108
  - maxgl, 108
- virtual
  - address, 386
  - address 0, 405
- virtual address, 20
- virtual core, 20
- virtual memory, 297
- VIS, 20
- VIS instructions
  - encoding, 565, 566
  - implicitly referencing GSR register, 80
- Visual Instruction Set, *See* VIS instructions

## W

- W-A-R, *See* write-after-read memory hazard
- warm\_reset* (WMR), 547
  - and STRAND\_ENABLE register, 525
  - enabling/disabling virtual processors, 523, 524
  - machine state changes, 548
- watchdog\_reset* (POR)
  - and RED\_state, 443
- watchdog\_reset* (WDR), 490
  - entering error\_state, 444

- exiting error\_state, 548, 582
- full-processor reset, 532
- invoking RED\_state trap processing, 474
- per-strand reset, 532
  - and XIR traps, 478
- watchdog\_reset* (WDR), and *guest\_watchdog*, 439
- watchdog\_reset* (WMR), 548
- watchpoint comparator, 97
- W-A-W, *See* write-after-write memory hazard
- WDR, 20
- WDR (*watchdog\_reset*), 548
- WDR, *See* *watchdog\_reset* (WDR)
- WIM register (SPARC V8), 371
- window fill exception, *See also* *fill\_n\_normal*
  - exception
- window fill trap handler, 32
- window overflow, 53, 491
- window spill exception, *See also* *spill\_n\_normal*
  - exception
- window spill trap handler, 32
- window state register, *See* WSTATE register
- window underflow, 491
- window, clean, 313
- window\_fill* exception, 88, 129, 457
  - RETURN, 311
- window\_spill* exception, 88, 457
- WMR (*warm\_reset*), 547
  - machine state changes, 548
- word, 20
  - alignment, 28, 114, 389
  - data format, 35
- WRASI instruction, 70, 74, 369
- WRAsr instruction, 369
  - accessing I/O registers, 29
  - attempt to write to ASR 5 (PC), 76
  - cannot write to PC register, 76
  - implementation dependencies, 572
  - writing ASRs, 70
- WRCCR instruction, 70, 72, 73, 369
- WRFP RS instruction, 71, 76, 369
- WRGSR instruction, 71, 80, 369
- WRHPR instruction, 103, 104, 106, 372
- WRIER instruction (SPARC V8), 371
- write ancillary state register (WRAsr)
  - instructions, 369
- write ancillary state register instructions, *See* WRAsr
  - instruction
- write hyperprivileged register instruction, 372
- write privileged register instruction, 374

write-after-read memory hazard, 394  
write-after-write memory hazard, 394  
WRPCR instruction, 71, 369  
WRPIC instruction, 71, 369, 488  
WRPR instruction, 20, 442

- accessing non-register-window PR state registers, 90
- accessing register-window PR state registers, 85
- and register-window PR state registers, 85
- effect on TNPC register, 91
- effect on TPC register, 91
- effect on TSTATE register, 93
- effect on TT register, 94
- writing to GL register, 102
- writing to PSTATE register, 94
- writing to TICK register, 75

WRPSR instruction (SPARC V8), 371  
WRSOFTINT instruction, 71, 81, 369  
WRSOFTINT\_CLR instruction, 71, 81, 82, 369, 497  
WRSOFTINT\_SET instruction, 71, 81, 82, 369, 497  
WRSTICK instruction, 71, 84, 369  
WRSTICK\_CMPR instruction, 71, 369  
WRTBR instruction (SPARC V8), 371  
WRTICK\_CMP instruction, 71, 369  
WRWIM instruction (SPARC V8), 371  
WRY instruction, 70, 72, 369  
WSTATE (window state) register

- description, 88
- and fill/spill exceptions, 492
- normal field, 492
- other field, 492
- overview, 85
- reading with RDPR instruction, 303
- spill exception, 190
- spill trap, 314
- state after reset, 550
- writing with WRPR instruction, 374

## X

XIR, 20  
XIR (*externally\_initiated\_reset*), 547  
XIR reset, 533  
XIR, *See externally\_initiated\_reset* (XIR)  
XIR\_STEERING register, 534

- state after reset, 552

XNOR instruction, 377  
XNORcc instruction, 377  
XOR instruction, 377

XORcc instruction, 377

## Y

Y register, 70, 72

- after multiplication completed, 281
- content after divide operation, 364
- divide operation, 364
- multiplication, 281
- state after reset, 549
- unsigned multiply results, 367
- WRY instruction, 370

Y register (deprecated), 72

## Z

zero virtual address, 405