



# UltraSPARC T1™ Supplement to the *UltraSPARC Architecture 2005*

---

*Draft D2.0, 17 Mar 2006*

*Privilege Levels: Privileged  
and Nonprivileged*

*Distribution: Public*

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

Part No: 819-3404-04  
Revision: Draft D2.0, 17 Mar 2006



Copyright 2002-2006 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, 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 2002–2006 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, 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>Preface</b> .....	<b>ix</b>
<b>1 UltraSPARC T1 Basics</b> .....	<b>1</b>
1.1 Background.....	1
1.2 UltraSPARC T1 Overview.....	3
1.3 UltraSPARC T1 Components .....	3
1.3.1 SPARC Physical Core .....	3
1.3.2 Floating-Point Unit (FPU) .....	4
1.3.3 L2 Cache.....	4
<b>2 Data Formats</b> .....	<b>5</b>
<b>3 Registers</b> .....	<b>7</b>
3.1 Ancillary State Registers (ASRs) .....	7
3.1.1 TICK Register.....	7
3.1.2 General Status Register (GSR) .....	7
3.1.3 Software Interrupt Register (SOFTINT).....	7
3.1.4 Tick Compare Register (TICK_CMPR).....	8
3.1.5 System Tick Register (STICK) .....	8
3.1.6 System Tick Compare Register (STICK_CMPR).....	8
3.1.7 PCR and PIC Registers.....	9
3.2 PR State Registers .....	9
3.2.1 Trap State (TSTATE) .....	9
3.2.2 Processor State Register (PSTATE) .....	9
3.2.3 Trap Level Register (TL).....	9
3.2.4 Global Level Register (GL).....	10
3.3 Floating-Point State Register (FSR).....	10
<b>4 Instruction Set Overview</b> .....	<b>11</b>
4.1 State Register Access.....	11

4.2	Floating-Point Operate (FPop) Instructions . . . . .	11
4.3	Reserved Opcodes and Instruction Fields . . . . .	12
4.4	Register Window Management . . . . .	12
<b>5</b>	<b>Instruction Definitions . . . . .</b>	<b>13</b>
5.1	Instruction Set Summary . . . . .	13
5.2	Prefetch and Prefetch from Alternate Space . . . . .	16
5.3	Trap on Integer Condition Codes (Tcc) . . . . .	16
5.4	VIS Instructions . . . . .	16
5.5	Partitioned Add/Subtract Instructions . . . . .	17
5.6	Align Data . . . . .	17
5.7	F Register Logical Operate Instructions . . . . .	17
5.8	Block Load and Store Instructions . . . . .	18
<b>6</b>	<b>Traps . . . . .</b>	<b>29</b>
6.1	Trap Levels . . . . .	29
<b>7</b>	<b>Interrupt Handling . . . . .</b>	<b>31</b>
7.0.1	Interrupt Queue Registers . . . . .	31
<b>8</b>	<b>Memory Models . . . . .</b>	<b>35</b>
8.1	Overview . . . . .	35
8.2	Supported Memory Models . . . . .	36
8.2.1	Total Store Order . . . . .	36
8.2.2	Relaxed Memory Order . . . . .	37
<b>9</b>	<b>Address Spaces and ASIs . . . . .</b>	<b>39</b>
9.1	Address Spaces . . . . .	39
9.1.1	Access to Nonexistent Memory or I/O . . . . .	39
9.1.2	48-bit Virtual Address Space . . . . .	39
9.2	Alternate Address Spaces . . . . .	41
9.2.1	ASI_REAL and ASI_REAL_LITTLE . . . . .	46
9.2.2	ASI_REAL_IO and ASI_REAL_IO_LITTLE . . . . .	46
9.2.3	ASI_SCRATCHPAD . . . . .	46
<b>10</b>	<b>Performance Instrumentation . . . . .</b>	<b>49</b>
10.1	Performance Control Register . . . . .	49
10.2	SPARC Performance Instrumentation Counter . . . . .	51
<b>11</b>	<b>Memory Management . . . . .</b>	<b>53</b>
11.1	Translation Table Entry (TTE) . . . . .	53
11.1.1	TTE Tag Format . . . . .	53
11.1.2	TTE Data Format . . . . .	53

11.2	Translation Storage Buffer	55
11.3	MMU-Related Faults and Traps	57
<b>12</b>	<b>Implementation Dependencies</b>	<b>59</b>
12.1	SPARC V9 General Information	59
12.1.1	Level-2 Compliance (Impl. Dep. #1)	59
12.1.2	Unimplemented Opcodes, ASIs, and ILLTRAP	59
12.1.3	Trap Levels (Impl. Dep. #37, 38, 39, 40, 101, 114, 115)	59
12.1.4	Trap Handling (Impl. Dep. #16, 32, 33, 35, 36, 44)	60
12.1.5	Population Count Instruction (POPC)	60
12.1.6	Secure Software	60
12.1.7	Address Masking (Impl. Dep. #125)	60
12.2	SPARC V9 Integer Operations	61
12.2.1	Integer Register File and Window Control Registers (Impl. Dep. #2)	61
12.2.2	SAVE Instruction	61
12.2.3	Clean Window Handling (Impl. Dep. #102)	62
12.2.4	Integer Multiply and Divide	62
12.2.5	MULSc	62
12.3	SPARC V9 Floating-Point Operations	62
12.3.1	Subnormal Operands and Results: Nonstandard Operation	62
12.3.2	Overflow, Underflow, and Inexact Traps (Impl. Dep. #3, 55)	63
12.3.3	Quad-Precision Floating-Point Operations (Impl. Dep. #3)	63
12.3.4	Floating-Point Square Root	64
12.3.5	Floating-Point Upper and Lower Dirty Bits in FPRS Register	64
12.3.6	Floating-Point State Register (FSR) (Impl. Dep. #13, 19, 22, 23, 24)	65
12.4	SPARC V9 Memory-Related Operations	67
12.4.1	Load/Store Alternate Address Space (Impl. Dep. #5, 29, 30)	67
12.4.2	Read/Write ASR (Impl. Dep. #6, 7, 8, 9, 47, 48)	67
12.4.3	FLUSH and Self-Modifying Code (Impl. Dep. #122)	67
12.4.4	PREFETCH{A} (Impl. Dep. #103, 117)	68
12.4.5	Instruction Prefetch	69
12.4.6	LDTW/STTW Handling (Impl. Dep. #107, 108)	69
12.4.7	Floating-Point <i>mem_address_not_aligned</i> (Impl. Dep. #109, 110, 111, 112)	69
12.4.8	Supported Memory Models (Impl. Dep. #113, 121)	69
12.4.9	Implicit ASI when TL > 0 (Impl. Dep. #124)	69
12.5	Non-SPARC V9 Extensions	70
12.5.1	Cache Subsystem	70
12.5.2	Block Memory Operations	70
12.5.3	Partial Stores	70

12.5.4	Short Floating-Point Loads and Stores . . . . .	70
12.5.5	Interrupt Vector Handling . . . . .	71
12.5.6	Power-Down Support . . . . .	71
12.5.7	UltraSPARC T1 Instruction Set Extensions (Impl. Dep. #106)	71
12.5.8	Performance Instrumentation . . . . .	71
<b>A</b>	<b>Assembly Language Syntax . . . . .</b>	<b>73</b>
<b>B</b>	<b>Programming Guidelines . . . . .</b>	<b>75</b>
B.1	Multithreading . . . . .	75
B.2	Pipeline Strand Flush . . . . .	76
B.3	Instruction Latencies . . . . .	76
B.4	Grouping Rules . . . . .	90
B.5	Floating-Point Operations . . . . .	90
B.6	Synchronization . . . . .	90
<b>C</b>	<b>Opcod Maps . . . . .</b>	<b>93</b>
<b>D</b>	<b>Instructions and Exceptions . . . . .</b>	<b>103</b>
<b>E</b>	<b>IEEE 754 Floating Point Support . . . . .</b>	<b>105</b>
E.1	Special Operand Handling . . . . .	105
E.1.1	Infinity Arithmetic . . . . .	106
E.1.2	Zero Arithmetic . . . . .	111
E.1.3	NaN Arithmetic . . . . .	112
E.1.4	Special Inexact Exceptions . . . . .	113
E.2	Subnormal Handling . . . . .	114
<b>F</b>	<b>Caches and Cache Coherency . . . . .</b>	<b>115</b>
F.1	Cache Flushing . . . . .	115
F.1.1	Displacement Flushing . . . . .	116
F.1.2	Memory Accesses and Cacheability . . . . .	116
F.1.3	Coherence Domains . . . . .	117
F.1.4	Memory Synchronization: MEMBAR and FLUSH . . . . .	119
F.1.5	Atomic Operations . . . . .	120
F.1.6	Nonfaulting Load . . . . .	121
<b>G</b>	<b>Glossary . . . . .</b>	<b>123</b>
	<b>Index . . . . .</b>	<b>1</b>

# Preface

---

Welcome to the UltraSPARC T1 Processor Supplement, D2.0. This document contains information about the processor-specific aspects of the architecture and programming of the UltraSPARC T1 processor, one of Sun Microsystems' family processors compliant with UltraSPARC Architecture™. It is intended to supplement the *UltraSPARC Architecture 2005* with processor-specific information.

---

## Target Audience

This User's Guide is mainly targeted for programmers who write software for the UltraSPARC T1 processor. This manual contains a depository of information that is useful to operating system programmers, application software programmers and logic designers, who are trying to understand the architecture and operation of the UltraSPARC T1 processor. This manual is both a guide and a reference manual for programming of the processor.

---

## 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. **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.

---

## 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<sub>2</sub>, FFFF 0000<sub>16</sub>). In some cases, numbers may be preceded by "0x" to indicate hexadecimal (base-16) notation (for example, 0xFFFF 0000). Long binary and hexadecimal numbers within the text may have spaces inserted every four characters to improve readability.

An en dash ( – ) with no spaces indicates a range, for example, 0001<sub>16</sub>–0000<sub>16</sub>.

Also see the colon ( : ) and double period ( .. ) notation described in the previous section.

---

## Informational Notes

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

**Note** | 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 processor might implement an architectural feature.
<b>V9 Compatibility Note</b>	Note containing information about possible differences between UltraSPARC Architecture and SPARC V9 implementations. Such information may not pertain to other SPARC V9 implementations.

# UltraSPARC T1 Basics

---

---

## 1.1 Background

UltraSPARC T1 is the first chip multiprocessor that fully implements Sun's Throughput Computing initiative. Throughput Computing is a technique that takes advantage of the thread-level parallelism that is present in most commercial workloads. Unlike desktop workloads, which often have a small number of threads concurrently running, most commercial workloads achieve their scalability by employing large pools of concurrent threads.

Historically, microprocessors have been designed to target desktop workloads, and as a result have focused on running a single thread as quickly as possible. Single thread performance is achieved in these microprocessors by a combination of extremely deep pipelines (over 20 stages in Pentium 4) and by executing multiple instructions in parallel (referred to as instruction-level parallelism, or ILP). The basic tenet behind Throughput Computing is that exploiting ILP and deep pipelining has reached the point of diminishing returns and as a result, current microprocessors do not utilize their underlying hardware very efficiently.

For many commercial workloads, the physical processor core will be idle most of the time waiting on memory, and even when it is executing it will often be able to only utilize a small fraction of its wide execution width. So rather than building a large and complex ILP processor that sits idle most of the time, a number of small, single-issue physical processor cores that employ multithreading are built in the same chip area. Combining multiple physical processors cores on a single chip with multiple hardware-supported threads (strands) per physical processor core, allows very high performance for highly threaded commercial applications. This approach is called thread-level parallelism (TLP). The difference between TLP and ILP is shown in FIGURE 1-1.



**FIGURE 1-1** Differences Between TLP and ILP

The memory stall time of one strand can often be overlapped with execution of other strands on the same physical processor core, and multiple physical processor cores run their strands in parallel. In the ideal case, shown in FIGURE 1-1, memory latency can be completely overlapped with execution of other strands. In contrast, instruction-level parallelism simply shortens the time to execute instructions, and does not help much in overlapping execution with memory latency.<sup>1</sup>

Given this ability to overlap execution with memory latency, why don't more processors utilize TLP? The answer is that designing processors is a mostly evolutionary process, and the ubiquitous deeply pipelined, wide ILP physical processor cores of today are the evolutionary outgrowth from a time when the CPU was the bottleneck in delivering good performance.

With physical processor cores capable of multiple-GHz clocking, the performance bottleneck has shifted to the memory and I/O subsystems and TLP has an obvious advantage over ILP for tolerating the large I/O and memory latency prevalent in commercial applications. Of course, every architectural technique has its advantages and disadvantages. The one disadvantage of employing TLP over ILP is that execution of a single strand may be slower on a TLP processor than an ILP processor. With physical processor cores running at frequencies well over one GHz, a strand capable of executing only a single instruction per cycle is fully capable of completing tasks in the time required by the application, making this disadvantage a non-issue for nearly all commercial applications.

<sup>1</sup>. Processors that employ out-of-order ILP can overlap some memory latency with execution. However, this overlap is typically limited to shorter memory latency events such as L1 cache misses that hit in the L2 cache. Longer memory latency events such as main memory accesses are rarely overlapped to a significant degree with execution by an out-of-order processor.

---

## 1.2 UltraSPARC T1 Overview

UltraSPARC T1 is a single-chip multiprocessor. UltraSPARC T1 contains eight SPARC® physical processor cores. Each SPARC physical processor core has full hardware support for four virtual processors (or “strands”). These four strands run simultaneously, with the instructions from each of the four strands executed round-robin by the single-issue pipeline. When a strand encounters a long-latency event, such as a cache miss, it is marked unavailable and instructions will not be issued from that strand until the long-latency event is resolved. Round-robin execution of the remaining available strands will continue while the long-latency event of the first strand is resolved.

Each SPARC physical core has a 16-Kbyte, 4-way associative instruction cache (32-byte lines), and 8K-byte, 4-way associative data cache (16-byte lines) that are shared by the four strands. The eight SPARC physical cores are connected through a crossbar to an on-chip unified 3-Mbyte, 12-way associative L2 cache (with 64-byte lines). The L2 cache is banked 4 ways to provide sufficient bandwidth for the eight SPARC physical cores.

---

## 1.3 UltraSPARC T1 Components

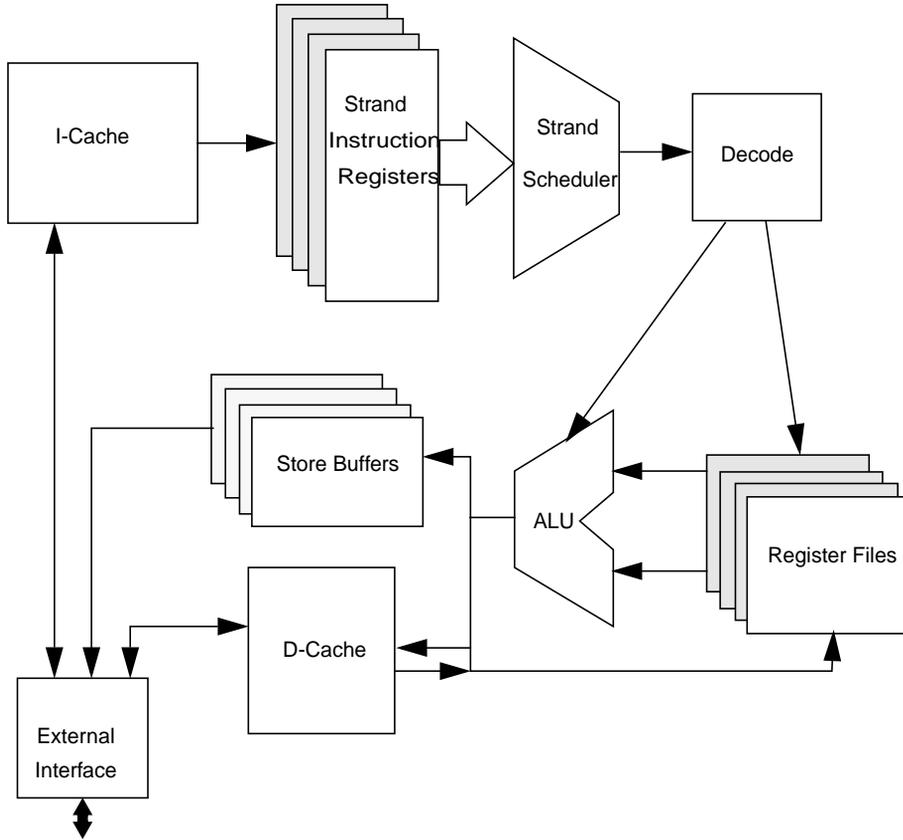
This section describes each component in UltraSPARC T1:

- SPARC physical core
- Floating-point unit
- L2 cache

### 1.3.1 SPARC Physical Core

Each SPARC physical core has hardware support for four strands. This support consists of a full register file (with eight register windows) per strand, with most of the ASI, ASR, and privileged registers replicated per strand. The four strands share the instruction and data caches.

FIGURE 1-2 illustrates SPARC physical core.



**FIGURE 1-2** SPARC Core Block Diagram

## 1.3.2 Floating-Point Unit (FPU)

A single floating-point unit is shared by all eight SPARC physical cores. The shared floating-point unit is sufficient for most commercial applications, in which fewer than 1% of instructions typically involve floating-point operations.

## 1.3.3 L2 Cache

The L2 cache is banked four ways, with the bank selection based on address bits 7:6. The cache is 3 Mbytes, 12-way set associative, and has a line size of 64 bytes.

# Data Formats

---

The UltraSPARC T1 processor supports all UltraSPARC Architecture 2005 data formats; see the Data Formats chapter of the *UltraSPARC Architecture 2005* for details.



# Registers

---

This chapter discusses the specifics of UltraSPARC T1 registers, as they differ from the register definitions in *UltraSPARC Architecture 2005*.

---

## 3.1 Ancillary State Registers (ASRs)

### 3.1.1 TICK Register

See the *UltraSPARC Architecture 2005* for a general description of this register.

The TICK register contains two fields: `npt` and `counter`. On an UltraSPARC T1 processor, the `npt` field is replicated per strand, while the `counter` field is shared by all four strands on a physical processor core. The counter increments each physical processor core clock but, on an UltraSPARC T1 processor, the least significant 2 bits of the `counter` field always read as 0.

### 3.1.2 General Status Register (GSR)

Each strand has a nonprivileged General Status register (GSR), as described in the *UltraSPARC Architecture 2005*.

All UltraSPARC Architecture 2005 GSR fields are supported in the UltraSPARC T1 implementation. However, the `mask` and `scale` fields are not directly written by VIS instructions; they are provided for use by software emulation.

### 3.1.3 Software Interrupt Register (SOFTINT)

Each strand has a privileged software interrupt register, as described in the *UltraSPARC Architecture 2005*.

The software interrupt register contains three fields: *sm*, *int\_level*, and *tm*. Setting any of *sm*, *tm*, or *SOFTINT*{14} generates an *interrupt\_level\_14* exception. However, these bits are considered completely independent of each other. Thus, a *Stick Compare* event will only set bit 16 and generate *interrupt\_level\_14* exception, not also set bit 14.

**UltraSPARC T1 Programming Note** | It is possible (but difficult) in UltraSPARC T1 for software to clear a *SOFTINT* bit between the setting of that bit and the generation of the interrupt from the bit being set because (there is a three-cycle window between the setting of the bit and the interrupt in UltraSPARC T1). If software were to do this, it would see an *interrupt\_level\_n* interrupt, but would find no bit set in the *SOFTINT* register. Note that normal software would only clear a bit in response to taking the *interrupt\_level\_n* exception, so this race condition should not occur in normal operation.

**UltraSPARC T1 Programming Note** | It is possible, but even more difficult than the above case, for software to zero a *SOFTINT* bit as it is getting set to 1, while another core is accessing its *SOFTINT* register, with timing such that hardware decides to take a *SOFTINT* trap, but the *SOFTINT* register is clear by the time it decides the trap number. In this case, hardware will take a trap  $40_{16}$ . Since software should only clear a bit that is known to be set, this should never happen in normal operation.

### 3.1.4 Tick Compare Register (TICK\_CMPR)

Each strand has a privileged Tick Compare (*TICK\_CMPR*) register, as described in the *UltraSPARC Architecture 2005*.

### 3.1.5 System Tick Register (STICK)

On an UltraSPARC T1 processor, the *STICK* register is an alias for the *TICK* register. Writes to *STICK* will be reflected in *TICK*, and vice versa. See the description of *TICK* above for the behavior of this register.

### 3.1.6 System Tick Compare Register (STICK\_CMPR)

Each strand has a privileged System Tick Compare (*STICK\_CMPR*) register, as described in the *UltraSPARC Architecture 2005*.

## 3.1.7 PCR and PIC Registers

TABLE 3-1 UltraSPARC T1-Specific Performance Instrumentation Registers

ASR Number	ASR Name	Access	priv	Replicated by Strand	Description
10 <sub>16</sub>	PCR	RW	Y <sup>2</sup>	Y	Performance counter control register
11 <sub>16</sub>	PIC	RW	Y <sup>1</sup>	Y	Performance Instrumentation Counter register

Notes:

1. Nonprivileged access with PCR.priv = 1 causes a *privileged\_action* exception.
2. Nonprivileged access causes a *privileged\_opcode* exception.

---

## 3.2 PR State Registers

### 3.2.1 Trap State (TSTATE)

Each virtual processor (strand) has *MAXPTL*(2) Trap State (TSTATE) registers, as described in the *UltraSPARC Architecture 2005*.

### 3.2.2 Processor State Register (PSTATE)

Each virtual processor (strand) has a Processor State register, as described in the *UltraSPARC Architecture 2005*.

### 3.2.3 Trap Level Register (TL)

Each virtual processor (strand) has a Trap Level register, as described in the *UltraSPARC Architecture 2005*.

The maximum trap level (*MAXPTL*) for UltraSPARC T1 is 2.

### 3.2.4 Global Level Register (GL)

Each virtual processor (strand) has a Global Level register, as described in the *UltraSPARC Architecture 2005*.

The maximum global level (*MAXPGL*) for UltraSPARC T1 is 2.

---

## 3.3 Floating-Point State Register (FSR)

Each virtual processor (strand) has a Floating-Point State register, FSR, as described in the *UltraSPARC Architecture 2005*.

UltraSPARC T1 does not provide a nonstandard floating-point mode, so the *ns* field of FSR is always 0.

On UltraSPARC T1, *FSR.ver* always reads as 0.

*FSR.qne* always reads as 0, because UltraSPARC T1 neither needs nor supports a floating-point queue (FQ).

# Instruction Set Overview

---

The UltraSPARC T1 processor implements the instruction set described in the *UltraSPARC Architecture 2005*. Additional UltraSPARC T1-specific details are described in this chapter.

---

## 4.1 State Register Access

UltraSPARC T1 supports the standard ASRs described in the *UltraSPARC Architecture 2005*.

---

## 4.2 Floating-Point Operate (FPop) Instructions

UltraSPARC T1 implements the floating-point instruction set described in the *UltraSPARC Architecture 2005*.

UltraSPARC T1 generates the correct IEEE Std 754-1985 results (impl. dep. #3).

All floating-point quad-precision operations cause an *fp\_exception\_other* trap with `FSR.ftt = unimplemented_FPop`, and system software must emulate those operations.

---

## 4.3 Reserved Opcodes and Instruction Fields

An attempt to execute an opcode to which no instruction is assigned causes a trap. Specifically:

- Attempting to execute a reserved FPop causes an *fp\_exception\_other* trap (with `FSR.ftt = unimplemented_FPop`).
- Attempting to execute any other reserved opcode causes an *illegal\_instruction* trap.
- Attempting to execute a Tcc instruction with a nonzero value in the reserved field (bits 10:8 and 6:5 when `i = 0` or bits 10:7 when `i = 1`) causes an *illegal\_instruction* trap. See *Trap on Integer Condition Codes (Tcc)* on page 16.

See Appendix C, *Opcode Maps*, for a complete enumeration of the opcode assignments.

---

## 4.4 Register Window Management

`N_REG_WINDOWS = 8` on UltraSPARC T1 (impl. dep. #2-V8). The state of the eight register windows is determined by the contents of the set of privileged registers described in the *UltraSPARC Architecture 2005*.

# Instruction Definitions

## 5.1 Instruction Set Summary

The UltraSPARC T1 CPU implements both the standard UltraSPARC Architecture 2005 instruction set and a number of implementation-dependent extended instructions. Standard UltraSPARC Architecture 2005 instructions are documented in the *UltraSPARC Architecture 2005*. UltraSPARC T1 extended instructions are documented in *VIS Instructions* on page 16.

The superscripts and their meanings are defined in TABLE 5-1.

**TABLE 5-1** Instruction Superscripts

Superscript	Meaning
D	Deprecated instruction
P	Privileged instruction

UltraSPARC T1 executes most UltraSPARC Architecture 2005 instructions in hardware. Those that trap and are emulated in software are listed in TABLE 5-2.

**TABLE 5-2** UltraSPARC Architecture 2005 Instructions Not Directly Implemented by UltraSPARC T1 Hardware (1 of 3)

Instruction	Description	Exception Caused by Attempted Execution
ALLCLEAN	Mark all windows as clean	<i>illegal_instruction</i>
ARRAY{8,16,32}	3-D address to blocked byte address conversion	<i>illegal_instruction</i>
BMASK	Write the GSR.mask field	<i>illegal_instruction</i>
BSHUFFLE	Permute bytes as specified by the GSR.mask field	<i>illegal_instruction</i>
EDGE{8,16,32}{L}{N}	Edge boundary processing {little-endian} {non-condition-code altering}	<i>illegal_instruction</i>
FABSq	Floating-point absolute value quad	<i>fp_exception_other</i> [unimplemented_FPop]

**TABLE 5-2** UltraSPARC Architecture 2005 Instructions Not Directly Implemented by UltraSPARC T1 Hardware (2 of 3)

<b>Instruction</b>	<b>Description</b>	<b>Exception Caused by Attempted Execution</b>
FADDq	Floating-point add quad	<i>fp_exception_other</i> [unimplemented_FPop]
FCMPq	Floating-point compare quad	<i>fp_exception_other</i> [unimplemented_FPop]
FCMPEq	Floating-point compare quad (exception if unordered)	<i>fp_exception_other</i> [unimplemented_FPop]
FCMPEQ{16,32}	Four 16-bit / two 32-bit compare: set integer dest if src1 = src2	<i>illegal_instruction</i>
FCMPGT{16,32}	Four 16-bit / two 32-bit compare: set integer dest if src1 > src2	<i>illegal_instruction</i>
FCMPLE{16,32}	Four 16-bit / two 32-bit compare: set integer dest if src1 ≤ src2	<i>illegal_instruction</i>
FCMPNE{16,32}	Four 16-bit / two 32-bit compare: set integer dest if src1 ≠ src2	<i>illegal_instruction</i>
FDIVq	Floating-point divide quad	<i>fp_exception_other</i> [unimplemented_FPop]
FdMULq	Floating-point multiply double to quad	<i>fp_exception_other</i> [unimplemented_FPop]
FEXPAND	Four 8-bit to 16-bit expand	<i>illegal_instruction</i>
FiTOq	Convert integer to quad floating-point	<i>fp_exception_other</i> [unimplemented_FPop]
FMOVq	Floating-point move quad	<i>fp_exception_other</i> [unimplemented_FPop]
FMOVqcc	Move quad floating-point register if condition is satisfied	<i>fp_exception_other</i> [unimplemented_FPop]
FMOVqr	Move quad floating-point register if integer register contents satisfy condition	<i>fp_exception_other</i> [unimplemented_FPop]
FMULq	Floating-point multiply quad	<i>fp_exception_other</i> [unimplemented_FPop]
FMUL8SUx16	Signed upper 8- x 16-bit partitioned product of corresponding components	<i>illegal_instruction</i>
FMUL8ULx16	Unsigned lower 8-bit x 16-bit partitioned product of corresponding components	<i>illegal_instruction</i>
FMUL8x16	8- x 16-bit partitioned product of corresponding components	<i>illegal_instruction</i>
FMUL8x16AL	Signed lower 8-bit x 16-bit lower α partitioned product of four components	<i>illegal_instruction</i>
FMUL8x16AU	Signed upper 8-bit x 16-bit lower α partitioned product of four components	<i>illegal_instruction</i>
FMULD8SUx16	Signed upper 8-bit x 16-bit multiply ← 32-bit partitioned product of components	<i>illegal_instruction</i>
FMULD8ULx16	Unsigned lower 8-bit x 16-bit multiply ← 32-bit partitioned product of components	<i>illegal_instruction</i>
FNEGq	Floating-point negate quad	<i>fp_exception_other</i> [unimplemented_FPop]

**TABLE 5-2** UltraSPARC Architecture 2005 Instructions Not Directly Implemented by UltraSPARC T1 Hardware (3 of 3)

<b>Instruction</b>	<b>Description</b>	<b>Exception Caused by Attempted Execution</b>
FPACKFIX	Two 32-bit to 16-bit fixed pack	<i>illegal_instruction</i>
FPACK{16,32}	Four 16-bit/two 32-bit pixel pack	<i>illegal_instruction</i>
FPMERGE	Two 32-bit to 64-bit fixed merge	<i>illegal_instruction</i>
FSQRT(s,d,q)	Floating-point square root	<i>fp_exception_other</i> [unimplemented_FPop]
F(s,d,q)TO(q)	Convert between floating-point formats to quad	<i>fp_exception_other</i> [unimplemented_FPop]
FqTOi	Convert quad floating point to integer	<i>fp_exception_other</i> [unimplemented_FPop]
FqTOx	Convert quad floating point to 64-bit integer	<i>fp_exception_other</i> [unimplemented_FPop]
FSUBq	Floating-point subtract quad	<i>fp_exception_other</i> [unimplemented_FPop]
FxTOq	Convert 64-bit integer to floating-point	<i>fp_exception_other</i> [unimplemented_FPop]
IMPDEP1	Implementation-dependent instruction	<i>illegal_instruction</i>
IMPDEP2	Implementation-dependent instruction	<i>illegal_instruction</i>
INVALW <sup>P</sup>	Mark all windows as CANSAVE	<i>illegal_instruction</i>
LDQF	Load quad floating-point	<i>illegal_instruction</i>
LDQFA	Load quad floating-point into alternate space	<i>illegal_instruction</i>
LDSHORTF	Short FP load, zero-extend 8/16-bit load to a double-precision floating-point register	<i>data_access_exception</i>
NORMALW	Mark other windows as restorable	<i>illegal_instruction</i>
OTHERW	Mark restorable windows as other	<i>illegal_instruction</i>
PDIST	Distance between eight 8-bit components	<i>illegal_instruction</i>
POPC	Population count	<i>illegal_instruction</i>
PST	Eight 8-bit/four 16-bit/two 32-bit partial stores	<i>data_access_exception</i>
SHUTDOWN <sup>D,P</sup>	Shut down	<i>illegal_instruction</i>
STBLOCKF	64-byte block store with commit	<i>data_access_exception</i>
STQF	Store quad floating-point	<i>illegal_instruction</i>
STQFA	Store quad floating-point into alternate space	<i>illegal_instruction</i>
STSHORTF	Short FP store, 8-/16-bit store from a double-precision floating-point register	<i>data_access_exception</i>

---

## 5.2 Prefetch and Prefetch from Alternate Space

PREFETCH and PREFETCHA with fcn codes of 0–3 and 16–23 ( $10_{16-17_{16}}$ ) are implemented; all map to the same operation that brings the cache line into the L2 cache. On an MMU miss, the prefetch is dropped (weak prefetching).

Prefetch fcn codes  $5_{16}-F_{16}$  cause an *illegal\_instruction* trap. These operations are all “weak” prefetches; in some cases the prefetch operation is dropped.

---

## 5.3 Trap on Integer Condition Codes (Tcc)

See the *UltraSPARC Architecture 2005* for a complete description of the Tcc instruction.

<b>UltraSPARC T1 Implementation Note</b>	For the $i = 0$ variant of Tcc, UltraSPARC T1 does not check that reserved instruction bit 7 is 0. If bit 7 is set to 1 with $i = 0$ , UltraSPARC T1 treats it as a valid Tcc instruction.
--	--

---

## 5.4 VIS Instructions

UltraSPARC T1 supports in hardware the VIS 2 SIAM instruction and a subset of the VIS 1 instructions.

All other VIS 1 and VIS 2 instructions (see TABLE 5-2 for a list) cause an *illegal\_instruction* exception on UltraSPARC T1 and are emulated in software.

<b>UltraSPARC T1 Programming Note</b>	The use of VIS instructions on UltraSPARC T1 is strongly discouraged; the performance of even the implemented VIS instructions will often be below that of a comparable set of non-VIS instructions. This includes the block load and block store instructions. An UltraSPARC T1 physical processor core (four virtual processors) can only have a single outstanding floating-point operation (including block load, block store, and VIS instructions) in progress at any given time.
---	---

---

## 5.5 Partitioned Add/Subtract Instructions

See the *UltraSPARC Architecture 2005* for detailed descriptions of the FPADD and FPSUB instructions.

<b>UltraSPARC T1 Programming Note</b>	For good performance on UltraSPARC T1, the result of a single FPADD should not be used as part of a 64-bit graphics instruction source operand in the next instruction group.  Similarly, the result of a standard FPADD should not be used as a 32-bit graphics instruction source operand in the next instruction group.
---	--

---

## 5.6 Align Data

See the *UltraSPARC Architecture 2005* for detailed descriptions of the FALIGNDATA instruction.

<b>UltraSPARC T1 Programming Note</b>	For good performance on UltraSPARC T1, the result of FALIGNDATA should not be used as the source operand of a 32-bit SIMD instruction in the next instruction group.
---	--

---

## 5.7 F Register Logical Operate Instructions

See the *UltraSPARC Architecture 2005* for a description of the F register logical operate instructions (1-, 2-, and 3-operand).

<b>UltraSPARC T1 Programming Note</b>	For good performance on UltraSPARC T1, the result of a single logical operate instruction should not be used as part of the source operand of a 64-bit SIMD instruction in the next instruction group.  Similarly, the result of a standard logical operate instruction should not be used as the source operand of a 32-bit SIMD instruction source operand in the next instruction group.
---	---

---

## 5.8 Block Load and Store Instructions

For architectural descriptions of the LDBLOCKF and STBLOCKF instructions, see the *UltraSPARC Architecture 2005*.

**UltraSPARC T1 Implementation Note** On UltraSPARC T1, a block load forces a miss in the primary cache and will *not* allocate a line in the primary cache, but does allocate in the L2 cache. On UltraSPARC T1, block loads and stores from multiple virtual processors are not overlapped.

**Compatibility Note** These instructions were intended for use in transferring large blocks of data (more than 256 bytes); for example, in BCOPY and BFILL operations.

The use of block loads and stores on UltraSPARC T1 is deprecated; they are provided primarily for compatibility with existing software. UltraSPARC T1 provides a separate set of ASIs for high performance BCOPY and BFILL, as described in TABLE 9-1 on page 41. The performance of parallel BCOPY using appropriate ASIs (from among  $22_{16}$ ,  $23_{16}$ ,  $E2_{16}$ ,  $E3_{16}$ ,  $EA_{16}$ , and  $EB_{16}$ ) will be 2.5 to 3.5 times that of a BCOPY using block loads and stores. The performance of a single-threaded BCOPY using these ASIs will be 15% to 50% better than that of a BCOPY using block loads and stores.

On UltraSPARC T1, to order an LDBLOCKF with respect to earlier stores, an intervening MEMBAR #Sync must be executed.

Similarly on UltraSPARC T1, STBLOCKF source data registers are not interlocked against completion of previous load instructions (even if a second LDBLOCKF has been performed). The previous load data must be referenced by some other intervening instruction, or an intervening MEMBAR #Sync must be performed. If the programmer violates these rules, data from before or after the load may be used. UltraSPARC T1 continues execution before all of the store data has been transferred. If store data registers are overwritten before the next block store or MEMBAR #Sync instruction, then the following rule must be observed. The first register can be overwritten in the same instruction group as the STBLOCKF, the second register can be overwritten in the instruction group following the block store and so on. If this rule is violated, the store may store correct data or the overwritten data. Block stores always operate under the relaxed memory order (RMO) memory model, regardless of the PSTATE.mm setting, and require a subsequent MEMBAR #Sync to order them with respect to following loads.

After an STBLOCKF instruction but before executing a DONE, RETRY, or WRPR to PSTATE instruction, there must be an intervening MEMBAR #Sync or a trap. If this rule is violated, instructions after the DONE, RETRY, or WRPR to PSTATE may not see the effects of the updated PSTATE.

On UltraSPARC T1, LDBLOCKF does not follow memory model ordering with respect to stores. In particular, read-after-write and write-after-read hazards to overlapping addresses are not detected. The side-effects bit associated with the access is ignored (see *Translation Table Entry (TTE)* on page 53). If ordering with respect to earlier stores is important (for example, a block load that overlaps previous stores), then there must be an intervening MEMBAR #StoreLoad (or stronger MEMBAR). If ordering with respect to later stores is important (for example, a block load that overlaps a subsequent store), then there must be an intervening MEMBAR #LoadStore or reference to the block load data. This restriction does not apply when a trap is taken, so the trap handler need not consider pending block loads. If the LDBLOCKF overlaps a previous or later store and there is no intervening MEMBAR, trap, or data reference, the LDBLOCKF may return data from before or after the store.

<b>Compatibility Note</b>	Prior UltraSPARC machines may have written loaded data into the first two registers at the same time. Software that depends on this unsupported behavior must be modified for UltraSPARC T1.
---------------------------	--

STBLOCKF does not follow memory model ordering with respect to loads, stores or flushes. In particular, read-after-write, write-after-write, flush-after-write, and write-after-read hazards to overlapping addresses are not detected. The side-effects bit associated with the access is ignored. If ordering with respect to earlier or later loads or stores is important, then there must be an intervening reference to the load data (for earlier loads), or appropriate MEMBAR instruction. This restriction does not apply when a trap is taken, so the trap handler does not have to worry about pending block stores. If the STBLOCKF overlaps a previous load and there is no intervening load data reference or MEMBAR #LoadStore instruction, the load may return data from before or after the store and the contents of the block are undefined. If the STBLOCKF overlaps a later load and there is no intervening trap or MEMBAR #StoreLoad instruction, the contents of the block are undefined. If the STBLOCKF overlaps a later store or flush and there is no intervening trap or MEMBAR #StoreStore instruction, the contents of the block are undefined.

Block load and store operations do not obey the ordering restrictions of the currently selected virtual processor memory model (always TSO in UltraSPARC T1); block operations always execute under an RMO memory ordering model. Explicit MEMBAR instructions are required to order block operations among themselves or with respect to normal loads and stores. In addition, block operations do not conform to dependence order on the issuing strand; that is, no read-after-write or writer-after-read checking occurs between block loads and stores. Explicit MEMBARs must be used to enforce dependence ordering between block operations that reference the same address.

Typically, LDBLOCKF and STBLOCKF are used in loops where software can ensure that there is no overlap between the data being loaded and the data being stored. The loop must be preceded and followed by the appropriate MEMBARs to ensure that there are no hazards with loads and stores outside the loops. CODE EXAMPLE 5-1 illustrates the inner loop of a byte-aligned block copy operation.

Note that the loop must be unrolled twice to achieve maximum performance. All FP register references in this code example are to 64-bit registers. Eight versions of this loop are needed to handle all the cases of double word misalignment between the source and destination.

**CODE EXAMPLE 5-1** Byte-Aligned Block Copy Inner Loop

```

loop:
    faligndata    %f0, %f2, %f34
    faligndata    %f2, %f4, %f36
    faligndata    %f4, %f6, %f38
    faligndata    %f6, %f8, %f40
    faligndata    %f8, %f10, %f42
    faligndata    %f10, %f12, %f44
    faligndata    %f12, %f14, %f46
    addcc         %l0, -1, %l0
    bg,pt        ll
    fmovd         %f14, %f48
    end of loop handling
ll: ldda         [regaddr] #ASI_BLK_P, %f0
    stda         %f32, [regaddr] #ASI_BLK_P
    faligndata    %f48, %f16, %f32
    faligndata    %f16, %f18, %f34
    faligndata    %f18, %f20, %f36
    faligndata    %f20, %f22, %f38
    faligndata    %f22, %f24, %f40
    faligndata    %f24, %f26, %f42
    faligndata    %f26, %f28, %f44
    faligndata    %f28, %f30, %f46
    addcc         %l0, -1, %l0
    be,pnt       done
    fmovd         %f30, %f48
    ldda         [regaddr] #ASI_BLK_P, %f16
    stda         %f32, [regaddr] #ASI_BLK_P
    ba           loop
    faligndata    %f48, %f0, %f32
done:  end of loop processing

```

## 5.9 Block Initializing Store ASIs

The Block Initializing Store ASIs are specific to the UltraSPARC T1 implementation and are not guaranteed to be portable to other UltraSPARC Architecture implementations. They should only appear in platform-specific dynamically-linked libraries or in code generated at runtime by software (for example, a just-in-time compiler) that is aware of the specific implementation upon which it is executing.

Instruction	imm_asi	ASI Value	Operation	Assembly Language Syntax
ST{B,H,W,X,D}A	ASI_STBI_AIUP	22 <sub>16</sub>	64-byte block initialing store to primary address space, user privilege	<code>st {b,h,w,x,d}a reg<sub>rd</sub>, [reg_addr] imm_asi</code> <code>st {b,h,w,x,d}a reg<sub>rd</sub>, [reg_plus_imm] %asi</code>
ST{B,H,W,X,D}A	ASI_STBI_AIUS	23 <sub>16</sub>	64-byte block initialing store to secondary address space, user privilege	
ST{B,H,W,X,D}A	ASI_STBI_N	27 <sub>16</sub>	64-byte block initialing store to nucleus address space	
ST{B,H,W,X,D}A	ASI_STBI_AIUPL_L	2A <sub>16</sub>	64-byte block initialing store to primary address space, user privilege, little-endian	
ST{B,H,W,X,D}A	ASI_STBI_AIUSL	2B <sub>16</sub>	64-byte block initialing store to secondary address space, user privilege, little-endian	
ST{B,H,W,X,D}A	ASI_STBI_NL	2F <sub>16</sub>	64-byte block initialing store to nucleus address space, little-endian	
ST{B,H,W,X,D}A	ASI_STBI_P	E2 <sub>16</sub>	64-byte block initialing store to primary address space	
ST{B,H,W,X,D}A	ASI_STBI_S	E3 <sub>16</sub>	64-byte block initialing store to secondary address space	
ST{B,H,W,X,D}A	ASI_STBI_PL	EA <sub>16</sub>	64-byte block initialing store to primary address space, little-endian	
ST{B,H,W,X,D}A	ASI_STBI_SL	EB <sub>16</sub>	64-byte block initialing store to secondary address space, little-endian	

## Description

The UltraSPARC T1-specific block initializing store instructions are selected by using one of the block-initializing ASIs with integer store alternate instructions. These ASIs allow block-initializing stores to be performed to the same address spaces as normal stores. Little-endian ASIs access data in little-endian format; otherwise, the access is assumed to be big-endian.

Integer stores of all sizes are allowed with these ASIs, and STDA behaves as a standard store doubleword. All stores to these ASIs operate under relaxed memory ordering (RMO), regardless of the value of PSTATE.mm. Software must follow a sequence of these stores with a MEMBAR #Sync to ensure ordering with respect to subsequent loads and stores.

A store to one of these ASIs where the least-significant 6 bits of the address are nonzero (that is, not the first word in the cache line) behaves the same as a normal store (with RMO ordering).

A store to one of these ASIs where the least-significant 6 bits of the address are zero will load a cache line in the L2 cache with either all zeros or the existing memory data, and then update the beginning of the cache line with the new store data. This special store behavior ensures that the line maintains coherency when it is loaded into the cache, but will not generally fetch the line from memory (instead, initializing it with zeroes).

A store using one of these ASIs to a noncacheable location behaves the same as a normal store.

<b>UltraSPARC T1 Implementation Note</b>	On UltraSPARC T1, a noncacheable address is identified by .
--	---

<b>Programming Note</b>	These instructions are particularly useful in combination with load twin extended word instructions for transferring large blocks (more than 256 bytes) of data; for example, in implementing <code>bcopy()</code> and <code>bfill()</code> operations.
-----------------------------	---

<b>UltraSPARC T1 Implementation Note</b>	On UltraSPARC T1, block initializing stores and load twin doublewords from multiple strands are fully overlapped.
--	---

Attempted use of any of these ASIs by a floating-point store alternate instruction (STFA, STDFA) causes a *data\_access\_exception* exception.

Access to any of these ASIs by an instruction with misaligned address causes a *mem\_address\_not\_aligned* exception.

**Programming  
Note**

The following pseudocode shows how these ASIs can be used to do a quadword-aligned (on both source and destination) copy of  $N$  quadwords from  $A$  to  $B$  (where  $N > 3$ ). Note that the final 64 bytes of the copy is performed using normal stores, to guarantee that all initial zeros in a cache line are overwritten with copy data.

```
%10 ← [A]; %11 ← [B]
prefetch [%10]
for (i = 0; i < N-4; i++) {
    if (!(i % 4)) { prefetch [%10+64] }
    ldda [%10] #ASI_BLK_INIT_ST_P, %12
    add %10, 16, %10
    stxa %12, [%11] #ASI_BLK_INIT_ST_P
    add %11, 8, %11
    stxa %13, [%11+8] #ASI_BLK_INIT_ST_P
    add %11, 8, %11
}
for (i = 0; i < 4; i++) {
    ldda [%10] #ASI_BLK_INIT_ST_P, %12
    add %10, 16, %10
    stx %12, [%11]
    stx %13, [%11+8]
    add %11, 16, %11
}
membar #Sync
```

An overlapped copy operation must avoid issuing a block-init store to a line before all loads from that line have been issued. Otherwise, one or more of the loads may see the interim "zero" side-effect value. This typically means that  $\text{abs}(A-B)$  must be 64.

**UltraSPARC T1  
Programming  
Notes**

(1) These ASIs are specific to UltraSPARC T1, to provide a high-performance mechanism for BCOPY operations, as an alternative to legacy block load and block store instructions (which rely on the floating-point register file and thus are limited by the single register file port). These ASIs are only allowed in platform-specific dynamically linked libraries and in code generated at runtime by software (for example, a just-in-time compiler) that is aware of the implementation upon which it is executing.

(2) These ASIs provide a higher performance `bcopy()` or `bfill()` than the block loads and stores described in Section 5.8, due to their ability to overlap multiple loads and stores between strands and to avoid the unnecessary fetch from memory of the data that is overwritten by the store. The performance of parallel `bcopy()` using these ASIs will be 2.5 to 3.5 times that of a `bcopy()` using block loads and stores. The performance of a single-threaded `bcopy()` using these ASIs will be 15% to 50% better than that of a `bcopy()` using block loads and stores.

*Exceptions*

*VA\_watchpoint*  
*mem\_address\_not\_aligned*  
*data\_access\_exception*

---

## 5.10 Load Twin Extended Word Instructions (nonprivileged)

The Load Twin Extended Word Instructions are not guaranteed to be portable to other UltraSPARC Architecture implementations. They should only appear in platform-specific dynamically-linked libraries or in code generated at runtime by software (for example, a just-in-time compiler) that is aware of the specific implementation upon which it is executing.

*Description* Load Twin Extended Word instructions are new in the UltraSPARC Architecture 2005; they are used to atomically read a 128-bit data item into a pair of integer registers.

See the *UltraSPARC Architecture 2005* for details.

**Programming Note** These instructions are particularly useful in combination with block-initializing stores for transferring large blocks of data (more than 256 bytes); for example, in implementing `bcopy()` and `bfill()` operations. See the description of Block Initializing Stores for an example of how Load Twin Extended Word can be used in combination with those instructions.

**UltraSPARC T1 Implementation Note** On UltraSPARC T1, a load twin extended word forces a miss in the primary cache and will *not* allocate a line in the primary cache, but does allocate in L2. On UltraSPARC T1, block initializing stores and load twin doublewords from multiple strands are fully overlapped.

See the description of Block Initializing Stores for an example of how Load Twin Doubleword can be used in combination with those instructions.

**UltraSPARC T1  
Programming  
Notes**

(1) These instructions, combined with store instructions using the UltraSPARC T1-specific Block Initializing Store ASIs, provide a high-performance mechanism for BCOPY operations, as an alternative to legacy block load and store (which rely on the floating-point register file and thus are limited by the single register file port). These ASIs are only allowed in platform-specific dynamically linked libraries and in code generated at runtime by software (for example, a just-in-time compiler) that is aware of the implementation upon which it is executing.

(2) These ASIs provide a higher performance `bcopy()` or `bfill()` than the block loads and stores described in Section 5.8, due to their ability to overlap multiple loads and stores between strands and to avoid the unnecessary fetch from memory of the data that is overwritten by the store. The performance of parallel `bcopy()` using these ASIs will be 2.5 to 3.5 times that of a `bcopy()` using block loads and stores. The performance of a single-threaded `bcopy()` using these ASIs will be 15% to 50% better than that of a `bcopy()` using block loads and stores.

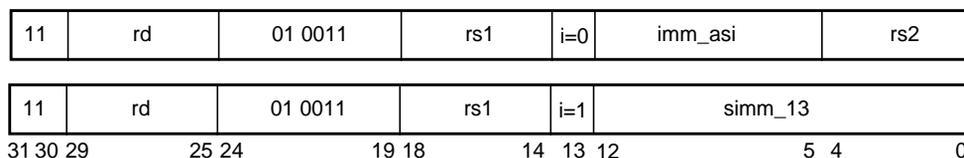
*See Also*      Block Initializing Store ASIs on page 21.

## 5.11 Load Twin Extended Word Instructions (privileged)

Instruction	imm_asi	ASI Value	Operation	Assembly Language Syntax
LDTX	ASI_LDTX_N	27 <sub>16</sub> <sup>†</sup>	128-bit atomic load	<code>ldda [reg_addr] imm_asi, reg_rd</code> <code>ldda [reg_plus_imm] %asi, reg_rd</code>
LDTX	ASI_LDTX_REAL	26 <sub>16</sub>	128-bit atomic load, real addressing (RA{63:0} set to VA{63:0})	
LDTX	ASI_LDTX_NL	2F <sub>16</sub> <sup>‡</sup>	128-bit atomic load, little endian	
LDTX	ASI_LDTX_REAL_L	2E <sub>16</sub>	128-bit atomic load, real addressing (RA{63:0} set to VA{63:0}), little endian	

<sup>†</sup> ASI 24<sub>16</sub> (deprecated) is aliased to ASI 27<sub>16</sub> in UltraSPARC T1.

<sup>‡</sup> ASI 2C<sub>16</sub> (deprecated) is aliased to ASI 2F<sub>16</sub> in UltraSPARC T1.



**Compatibility Note** | In previous UltraSPARC documents, these instructions were (loosely) referred to as "Quad LDD" instructions.

*Description* These instructions atomically read a 128-bit data item into two 64-bit integer registers. They are intended to be used to access TSB entries without requiring locks. The data is placed in an even/odd pair of 64-bit integer registers. The lowest address 64 bits is placed in the even-numbered register; the highest address 64-bits is placed in the odd-numbered register.

ASI\_LDTX\_REAL[\_L] bypasses the virtual-to-real portion of the translation, setting RA{63:0} = VA{63:0}.

In addition to the usual exceptions for LDTX using a privileged ASI, a *data\_access\_exception* trap occurs if these ASIs are used with any instruction other than LDTX or LDDA (which share an opcode). A *mem\_address\_not\_aligned* trap is taken if the access is not aligned on a 128-bit boundary.

*Exceptions*

*VA\_watchpoint*

*mem\_address\_not\_aligned* (Checked for opcode implied alignment if the opcode is not LDDA)

*data\_access\_exception*

# Traps

---

The UltraSPARC T1 processor implements the trap model described in the *UltraSPARC Architecture 2005*.

Additional UltraSPARC T1-specific details are described in this chapter.

---

## 6.1 Trap Levels

Each UltraSPARC T1 virtual processor supports two trap levels ( $MAXPTL = 2$ ).



# Interrupt Handling

---

## 7.0.1 Interrupt Queue Registers

Each strand has eight `ASI_QUEUE` registers at `ASI 2516`, `VA{63:0} = 3C016–3F816` that are used for communicating interrupts to the privileged mode operating system. These registers contain the head and tail pointers for four supervisor interrupt queues: `cpu_mondo`, `dev_mondo`, `resumable_error`, and `nonresumable_error`.

The tail registers are read-only. An attempted write to a tail register by privileged software generate a `data_access_exception` trap. The head registers are read/write.

Whenever the contents of the `CPU_MONDO_HEAD` and `CPU_MONDO_TAIL` registers are unequal, a `cpu_mondo` trap is generated. Whenever the contents of the `DEV_MONDO_HEAD` and `DEV_MONDO_TAIL` registers are unequal, a `dev_mondo` trap is generated. Whenever the contents of the `RESUMABLE_ERROR_HEAD` and `RESUMABLE_ERROR_TAIL` registers are unequal, a `resumable_error` trap is generated.

Unlike the other queue register pairs, the `nonresumable_error` trap is *not* automatically generated by hardware whenever the contents of the `NONRESUMABLE_ERROR_HEAD` and `NONRESUMABLE_ERROR_TAIL` registers are unequal; instead, hyperprivileged software must make it appear to privileged software as if a `nonresumable_error` trap has occurred.

**Warning** | There is a known “feature” in UltraSPARC T1 that affects `LDXA/STXA` by supervisor code to these ASI registers. If an immediately preceding instruction is a store that takes certain traps, an `LDXA` can corrupt an unrelated IRF (integer register file) register, or a `STXA` may complete in spite of the trap. To prevent this, it is *required* to have a non-store or NOP instruction before any `LDXA/STXA` to these ASIs. If the `LDXA/STXA` is at a branch target, there must be a non-store in the delay slot. Nonprivileged software is not affected by this.

**Programming Note**

These registers are intended to be used as head and tail pointers into a queue in memory storing the mondo or error interrupt data. When an interrupt is taken, the interrupt data are stored into the end of the appropriate queue. Then the corresponding tail register is updated to point beyond the new data, which causes a trap to be generated to privileged software (the operating system). Privileged software then processes the interrupt data from the head of the queue, updating the head register when the interrupt processing is completed.

While the first interrupt is being serviced, more interrupts may be placed on the queue. The operating system can read the tail pointer to service multiple interrupts at a time, or it can simply update the head pointer after each interrupt has been serviced and take a trap for each interrupt.

When all pending interrupts of the appropriate type have been serviced, the head and tail pointers will be equal again, and no further traps will be generated until new interrupt data is placed on the queue.

TABLE 7-1 through TABLE 7-8 define the format of the eight interrupt queue registers.

**TABLE 7-1** CPU Mondo Head Pointer – QUEUE\_CPU\_MONDO\_HEAD (ASI 25<sub>16</sub>, VA 3C0<sub>16</sub>)

Bit	Field	R/W	Description
63:14	—	R	<i>Reserved</i>
13:6	head	RW	Head pointer for CPU Mondo Interrupt Queue.
5:0	—	R	<i>Reserved</i>

**TABLE 7-2** CPU Mondo Tail Pointer – QUEUE\_CPU\_MONDO\_TAIL (ASI 25<sub>16</sub>, VA 3C8<sub>16</sub>)

Bit	Field	R/W	Description
63:14	—	R	<i>Reserved</i>
13:6	tail	RW	Tail pointer for CPU Mondo Interrupt Queue.
5:0	—	R	<i>Reserved</i>

**TABLE 7-3** Device Mondo Head Pointer – QUEUE\_DEV\_MONDO\_HEAD (ASI 25<sub>16</sub>, VA 3D0<sub>16</sub>)

Bit	Field	R/W	Description
63:14	—	R	<i>Reserved</i>
13:6	head	RW	Head pointer for Device Mondo Interrupt Queue.
5:0	—	R	<i>Reserved</i>

**TABLE 7-4** Device Mondo Tail Pointer – QUEUE\_DEV\_MONDO\_TAIL (ASI 25<sub>16</sub>, VA 3D8<sub>16</sub>)

Bit	Field	R/W	Description
63:14	—	R	<i>Reserved</i>
13:6	tail	RW	Tail pointer for Device Mondo Interrupt Queue.
5:0	—	R	<i>Reserved</i>

**TABLE 7-5** Resumable Error Head Pointer – QUEUE\_RESUMABLE\_HEAD (ASI 25<sub>16</sub>, VA 3E0<sub>16</sub>)

Bit	Field	R/W	Description
63:14	—	R	<i>Reserved</i>
13:6	head	RW	Head pointer for Resumable Error Queue.
5:0	—	R	<i>Reserved</i>

**TABLE 7-6** Resumable Error Tail Pointer – QUEUE\_RESUMABLE\_TAIL (ASI 25<sub>16</sub>, VA 3E8<sub>16</sub>)

Bit	Field	Initial Value	R/W	Description
63:14	—		R	<i>Reserved</i>
13:6	tail		RW	Tail pointer for Resumable Error Queue.
5:0	—		R	<i>Reserved</i>

**TABLE 7-7** Nonresumable Error Head Pointer – QUEUE\_NONRESUMABLE\_HEAD (ASI 25<sub>16</sub>, VA 3F0<sub>16</sub>)

Bit	Field	Initial Value	R/W	Description
63:14	—		R	<i>Reserved</i>
13:6	head		RW	Head pointer for NonResumable Error Queue.
5:0	—		R	<i>Reserved</i>

**TABLE 7-8** Nonresumable Error Tail Pointer – QUEUE\_NONRESUMABLE\_TAIL (ASI 25<sub>16</sub>, VA 3F8<sub>16</sub>)

Bit	Field	R/W	Description
63:14	—	R	<i>Reserved</i>
13:6	tail	RW	Tail pointer for NonResumable Error Queue.
5:0	—	R	<i>Reserved</i>



# Memory Models

---

## 8.1 Overview

SPARC V9 defines the semantics of memory operations for three memory models. From strongest to weakest, they are Total Store Order (TSO), Partial Store Order (PSO), and Relaxed Memory Order (RMO). The differences in these models lie in the freedom an implementation is allowed in order to obtain higher performance during program execution. The purpose of the memory models is to specify any constraints placed on the ordering of memory operations in uniprocessor and shared-memory multiprocessor environments.

For a full description of the TSO memory model, see the *UltraSPARC Architecture 2005*.

UltraSPARC T1 supports only TSO, with the exception that accesses using certain ASIs (notably, block loads and block stores) may operate under RMO (impl. dep. #113-V9-Ms10).

Although a program written for a weaker memory model potentially benefits from higher execution rates, it may require explicit memory synchronization instructions to function correctly if data is shared. MEMBAR is a memory synchronization primitive that enables a programmer to control explicitly the ordering in a sequence of memory operations. Processor consistency is guaranteed in all memory models.

The current memory model is indicated in the `PSTATE.mm` field. Its value is always 0 on UltraSPARC T1. An UltraSPARC T1 virtual processor always operates under the TSO memory model.

Memory is logically divided into real memory (cached) and I/O memory (noncached, with and without side effects) spaces (impl. dep. #118-V9). Real memory spaces may be cached and can be accessed without side effects. For example, a read (load) from real memory space returns the information most recently written. In

addition, an access to real memory space does not result in program-visible side effects. In contrast, a read from I/O space may not return the most recently written information and may result in program-visible side effects.

---

## 8.2 Supported Memory Models

The following sections contain brief descriptions of the two memory models supported by UltraSPARC T1. These definitions are for general illustration. Detailed definitions of these models can be found in *UltraSPARC Architecture 2005*. The definitions in the following sections apply to system behavior as seen by the programmer. A description of MEMBAR can be found in Section 8.3.2, “Memory Synchronization: MEMBAR and FLUSH” on page 72.

- Notes**
- (1) Stores to UltraSPARC T1 Internal ASIs, block loads, and block stores are outside the memory model; that is, they need MEMBARs to control ordering. See Section 8.3.8, “Instruction Prefetch to Side-Effect Locations” on page 79 and Section 13.5.3, “Block Load and Store Instructions” on page 172.
  - (2) Atomic load-stores are treated as both a load and a store and can only be applied to cacheable address spaces.

### 8.2.1 Total Store Order

UltraSPARC T1 implements the following programmer-visible properties in Total Store Order (TSO) mode:

- Loads are processed in program order; that is, there is an implicit MEMBAR #LoadLoad between them.
- Loads may bypass earlier stores. Any such load that bypasses such earlier stores must check (snoop) the store buffer for the most recent store to that address. A MEMBAR #Lookaside is not needed between a store and a subsequent load at the same noncacheable address.
- A MEMBAR #StoreLoad must be used to prevent a load from bypassing a prior store, if Strong Sequential Order is desired.
- Stores are processed in program order.
- Stores cannot bypass earlier loads.
- An L2 cache update is delayed on a store hit until all outstanding stores reach global visibility. For example, a cacheable store following a noncacheable store is not globally visible until the noncacheable store has reached global visibility; there is an implicit MEMBAR #MemIssue between them.

## 8.2.2 Relaxed Memory Order

UltraSPARC T1 implements the following programmer-visible properties for accesses through special ASIs that operate under the Relaxed Memory Order (RMO) model:

- There is no implicit order between any two memory references, either cacheable or noncacheable, except that noncacheable accesses to I/O space are all strongly ordered with respect to each other.
- A MEMBAR must be used between cacheable memory references if stronger order is desired. A MEMBAR #MemIssue is needed for ordering of cacheable after non-cacheable accesses. A MEMBAR #StoreLoad should be used between a store and a subsequent load at the same noncacheable address.



# Address Spaces and ASIs

---

---

## 9.1 Address Spaces

UltraSPARC T1 supports a 48-bit virtual address space.

### 9.1.1 Access to Nonexistent Memory or I/O

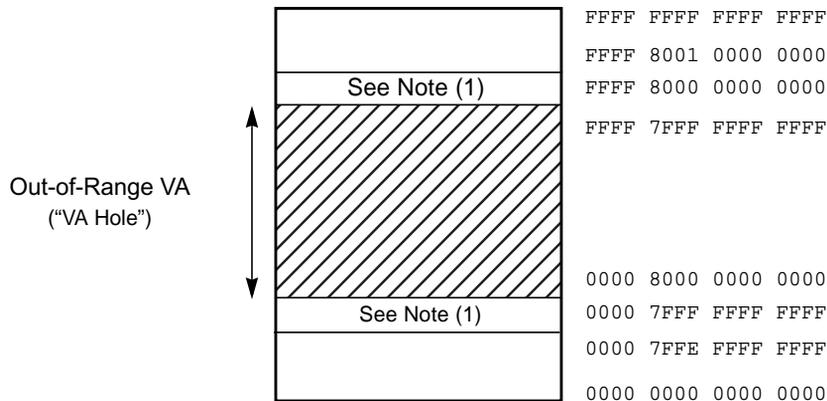
Accesses to nonexistent memory or I/O locations are treated as follows:

- A load access from a nonexistent memory or I/O location causes an exception
- An instruction fetch from a nonexistent memory or I/O location causes an exception
- A store access to a nonexistent memory or I/O location will be silently discarded by the system

### 9.1.2 48-bit Virtual Address Space

UltraSPARC T1 supports a 48-bit subset of the full 64-bit virtual address space (see FIGURE 9-1). Although the full 64 bits are generated and stored in integer registers, legal addresses are restricted to two equal halves at the extreme lower and upper portions of the full virtual address space. Virtual addresses between  $0000\ 8000\ 0000\ 0000_{16}$  and  $FFFF\ 7FFF\ FFFF\ FFFF_{16}$ , inclusive, lie within a “VA Hole”, are termed “out of range,” and are illegal.

Prior UltraSPARC implementations introduced the additional restriction on software to not use pages within 4 Gbytes of the VA hole as instruction pages, to avoid problems with prefetching into the VA hole. UltraSPARC T1 assumes that this convention is followed, for similar reasons. Note that there are no trap mechanisms to detect a violation of this convention.



Note (1): Prior implementations restricted use of this region to data only.

**FIGURE 9-1** UltraSPARC T1's 48-bit Virtual Address Space, With Hole

**Note** Throughout this document, when virtual address fields are specified as 64-bit quantities, bits 63:48 are assumed to be sign-extended from bit 47.

A number of state registers are affected by the reduced virtual address space. TBA, TPC, and TNPC registers are 48 bits wide, sign-extended to 64 bits on read accesses. VA watchpointing is 48 bits, zero-extended to 64-bits on read accesses. No checks are done when these registers are written by software. It is the responsibility of privileged software to properly update these registers.

An out of range address during an instruction access causes an *instruction\_access\_exception* trap if `PSTATE.am = 0`.

If the target address of a JMWL or RETURN instruction is an out-of-range address and `PSTATE.am` is not set, a trap is generated with `TPC[TL]` set to the address of the JMWL or RETURN instruction. This *instruction\_access\_exception* trap is lower priority than other traps on the JMWL or RETURN (*illegal\_instruction* due to nonzero reserved fields in the JMWL or RETURN, *mem\_address\_not\_aligned* trap, or *window\_fill* trap), because it really applies to the target. The trap handler can determine the out-of-range address by decoding the JMWL instruction from the code.

When any other control transfer instruction traps, it sets `TPC[TL]` to the address of the target instruction. Because the PC is sign-extended to 64 bits, the trap handler must adjust the PC value to compute the faulting address by **xoring** ones into the most significant 16 bits.

When a trap occurs on the delay slot of a taken branch or call whose target is out-of-range or is the last instruction below the VA hole, UltraSPARC T1 records the fact that NPC points to an out-of-range instruction in TNPC. If the trap handler executes a DONE or RETRY without saving TNPC, the *instruction\_access\_exception* trap is taken when the instruction at TNPC is executed. If TNPC is saved and subsequently restored by the trap handler, the fact that TNPC points to an out-of-range instruction is lost.

To guarantee that all out of range instruction accesses cause traps, software should not map addresses within  $2^{31}$  bytes of either side of the VA hole as executable.

An out-of-range address during a data access results in a *data\_access\_exception* trap if PSTATE.am is not set.

## 9.2 Alternate Address Spaces

TABLE 9-1 summarizes the ASI usage in UltraSPARC T1. The Section column lists where the operation of the ASI is explained. For several internal ASIs, a range of legal VAs is listed. An access outside the legal VA range will be aliased to a legal VA by ignoring the upper address bits.

- Notes**
- (1) All internal, nontranslating ASIs in UltraSPARC T1 can only be accessed using LDXA and STXA. This is different than UltraSPARC I/II, where LDDFA and STDFA can also be used to access internal ASIs. Using LDDFA and STDFA to access an internal ASI in UltraSPARC T1 results in a *data\_access\_exception* trap.
  - (2) ASIs  $80_{16}$ – $FF_{16}$  are unrestricted (nonprivileged and privileged software may access). ASIs  $00_{16}$ – $2F_{16}$  are restricted to privileged software.

**TABLE 9-1** UltraSPARC T1 ASI Usage (1 of 5)

ASI	ASI NAME	R/W	VA	Copy per strand	Description	Section
$00_{16}$ – $03_{16}$			Any	—	<i>data_access_exception</i>	
$04_{16}$	ASI_NUCLEUS	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	
$05_{16}$ – $0B_{16}$			Any	—	<i>data_access_exception</i>	
$0C_{16}$	ASI_NUCLEUS_LITTLE	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	

**TABLE 9-1** UltraSPARC T1 ASI Usage (2 of 5)

ASI	ASI NAME	R/W	VA	Copy per strand	Description	Section
0D <sub>16</sub> – 0F <sub>16</sub>			Any	—	<i>data_access_exception</i>	
10 <sub>16</sub>	ASI_AS_IF_USER_PRIMARY	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	
11 <sub>16</sub>	ASI_AS_IF_USER_SECONDARY	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	
12 <sub>16</sub> – 13 <sub>16</sub>			Any	—	<i>data_access_exception</i>	
14 <sub>16</sub>	ASI_REAL	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	9.2.1
15 <sub>16</sub>	ASI_REAL_IO	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	9.2.2
16 <sub>16</sub>	ASI_BLOCK_AS_IF_USER_PRIMARY	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	5.8
17 <sub>16</sub>	ASI_BLOCK_AS_IF_USER_SECONDARY	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	5.8
18 <sub>16</sub>	ASI_AS_IF_USER_PRIMARY_LITTLE	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	
19 <sub>16</sub>	ASI_AS_IF_USER_SECONDARY_LITTLE	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	
1A <sub>16</sub> – 1B <sub>16</sub>			Any	—	<i>data_access_exception</i>	
1C <sub>16</sub>	ASI_REAL_LITTLE	RW	Any	—	Nonallocating in L1 cache, same as ASI_REAL_IO_LITTLE for I/O addresses	9.2.1
1D <sub>16</sub>	ASI_REAL_IO_LITTLE	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	9.2.2
1E <sub>16</sub>	ASI_BLOCK_AS_IF_USER_PRIMARY_LITTLE	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	5.8
1F <sub>16</sub>	ASI_BLOCK_AS_IF_USER_SECONDARY_LITTLE	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	5.8
20 <sub>16</sub>	ASI_SCRATCHPAD	RW	0 <sub>16</sub> – 18 <sub>16</sub>	Y	Scratchpad registers 0–3	9.2.3
		RW	20 <sub>16</sub> – 28 <sub>16</sub>	—	<i>data_access_exception</i>	9.2.3
		RW	30 <sub>16</sub> – 38 <sub>16</sub>	Y	Scratchpad registers 6–7	9.2.3
21 <sub>16</sub>	ASI_MMU_CONTEXTID	RW	0 <sub>16</sub> – F8 <sub>16</sub>	—	(See <i>UltraSPARC Architecture 2005</i> )	

**TABLE 9-1** UltraSPARC T1 ASI Usage (3 of 5)

ASI	ASI NAME	R/W	VA	Copy per strand	Description	Section
22 <sub>16</sub>	ASI_LDTX_AIUP, ASI_STBI_AIUP	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> ) ASI_STBI_AIUP is used for Block-Initializing stores, As If User, Primary Context	5.10
23 <sub>16</sub>	ASI_LDTX_AIUS, ASI_STBI_AIUS	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> ) ASI_STBI_AIUS is used for Block-Initializing stores, As If User, Secondary Context	5.10
24 <sub>16</sub>	ASI_TWINK (ASI_LDTX) , ASI_QUAD_LDD <sup>D†</sup> , ASI_NUCLEUS_QUAD_LDD <sup>D†</sup>	R	Any	—	128-bit atomic Load Twin Doubleword (deprecated; superseded by ASI 27 <sub>16</sub> )	
25 <sub>16</sub>	ASI_QUEUE	RW	0 <sub>16</sub> – 3B8 <sub>16</sub>	—	Load/store does NOP	
		RW	3C0 <sub>16</sub> – 3F8 <sub>16</sub>	Y	(See <i>UltraSPARC Architecture 2005</i> )	
26 <sub>16</sub>	ASI_LDTX_REAL	R	Any	—	128-bit atomic LDTX, real address (see <i>UltraSPARC Architecture 2005</i> )	5.11
27 <sub>16</sub>	ASI_LDTX_N, ASI_STBI_N	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )  ASI_STBI_N is used for Block-Initializing stores, Nucleus Context	5.10
28 <sub>16</sub> – 29 <sub>16</sub>			Any	—	<i>data_access_exception</i>	
2A <sub>16</sub>	ASI_LDTX_AIUP_L, ASI_STBI_AIUP_L	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> ) ASI_STBI_AIUP_L is used for Block-Initializing stores, As If User, Primary Context, Little Endian	5.10
2B <sub>16</sub>	ASI_LDTX_AIUS_L, ASI_STBI_AIUS_L	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> ) ASI_STBI_AIUS_L is used for Block-Initializing stores, As If User, Secondary Context, Little Endian	5.10
2C <sub>16</sub>	ASI_TWINK_LITTLE (ASI_LDTX_L) , ASI_QUAD_LDD_LITTLE <sup>D†</sup> , ASI_NUCLEUS_QUAD_LDD_LITTLE <sup>D†</sup>	R	Any	—	128-bit atomic Load Twin Doubleword, little endian (deprecated; superseded by ASI 2F <sub>16</sub> )	

**TABLE 9-1** UltraSPARC T1 ASI Usage (4 of 5)

ASI	ASI NAME	R/W	VA	Copy per strand	Description	Section
2D <sub>16</sub>			Any	—	<i>data_access_exception</i>	
2E <sub>16</sub>	ASI_LDTX_REAL_L	R	Any	—	128-bit atomic LDTX, real address, little endian (see <i>UltraSPARC Architecture 2005</i> )	5.11
2F <sub>16</sub>	ASI_LDTX_NL, ASI_STBI_NL	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> ) ASI_STBI_NL is used for Block-Initializing stores, Nucleus context, Little-Endian	5.10
80 <sub>16</sub>	ASI_PRIMARY	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	
81 <sub>16</sub>	ASI_SECONDARY	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	
82 <sub>16</sub>	ASI_PRIMARY_NO_FAULT	R	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	
83 <sub>16</sub>	ASI_SECONDARY_NO_FAULT	R	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	
84 <sub>16</sub> – 87 <sub>16</sub>			Any	—	<i>data_access_exception</i>	
88 <sub>16</sub>	ASI_PRIMARY_LITTLE	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	
89 <sub>16</sub>	ASI_SECONDARY_LITTLE	RW	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	
8A <sub>16</sub>	ASI_PRIMARY_NO_FAULT_LITTLE	R	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	
8B <sub>16</sub>	ASI_SECONDARY_NO_FAULT_LITTLE	R	Any	—	(See <i>UltraSPARC Architecture 2005</i> )	
8C <sub>16</sub> – BF <sub>16</sub>			Any	—	<i>data_access_exception</i>	
C0 <sub>16</sub>	ASI_PST8_P		Any	—	<i>data_access_exception</i> <sup>1</sup>	
C1 <sub>16</sub>	ASI_PST8_S		Any	—	<i>data_access_exception</i> <sup>1</sup>	
C2 <sub>16</sub>	ASI_PST16_P		Any	—	<i>data_access_exception</i> <sup>1</sup>	
C3 <sub>16</sub>	ASI_PST16_S		Any	—	<i>data_access_exception</i> <sup>1</sup>	
C4 <sub>16</sub>	ASI_PST32_P		Any	—	<i>data_access_exception</i> <sup>1</sup>	
C5 <sub>16</sub>	ASI_PST32_S		Any	—	<i>data_access_exception</i> <sup>1</sup>	
C6 <sub>16</sub> – C7 <sub>16</sub>			Any	—	<i>data_access_exception</i>	
C8 <sub>16</sub>	ASI_PST8_PL		Any	—	<i>data_access_exception</i> <sup>1</sup>	
C9 <sub>16</sub>	ASI_PST8_SL		Any	—	<i>data_access_exception</i> <sup>1</sup>	
CA <sub>16</sub>	ASI_PST16_PL		Any	—	<i>data_access_exception</i> <sup>1</sup>	
CB <sub>16</sub>	ASI_PST16_SL		Any	—	<i>data_access_exception</i> <sup>1</sup>	

**TABLE 9-1** UltraSPARC T1 ASI Usage (5 of 5)

ASI	ASI NAME	R/W	VA	Copy per strand	Description	Section
CC <sub>16</sub>	ASI_PST32_PL		Any	—	<i>data_access_exception</i> <sup>1</sup>	
CD <sub>16</sub>	ASI_PST32_SL		Any	—	<i>data_access_exception</i> <sup>1</sup>	
CE <sub>16</sub> – CF <sub>16</sub>			Any	—	<i>data_access_exception</i>	
D0 <sub>16</sub>	ASI_FL8_P		Any	—	<i>data_access_exception</i> <sup>2</sup>	
D1 <sub>16</sub>	ASI_FL8_S		Any	—	<i>data_access_exception</i> <sup>2</sup>	
D2 <sub>16</sub>	ASI_FL16_P		Any	—	<i>data_access_exception</i> <sup>2</sup>	
D3 <sub>16</sub>	ASI_FL16_S		Any	—	<i>data_access_exception</i> <sup>2</sup>	
D4 <sub>16</sub> – D7 <sub>16</sub>			Any	—	<i>data_access_exception</i>	
D8 <sub>16</sub>	ASI_FL8_PL		Any	—	<i>data_access_exception</i> <sup>2</sup>	
D9 <sub>16</sub>	ASI_FL8_SL		Any	—	<i>data_access_exception</i> <sup>2</sup>	
DA <sub>16</sub>	ASI_FL16_PL		Any	—	<i>data_access_exception</i> <sup>2</sup>	
DB <sub>16</sub>	ASI_FL16_SL		Any	—	<i>data_access_exception</i> <sup>2</sup>	
DC <sub>16</sub> – DF <sub>16</sub>			Any	—	<i>data_access_exception</i>	
E0 <sub>16</sub>	ASI_BLK_COMMIT_P	RW	Any	—	<i>data_access_exception</i> <sup>3</sup>	
E1 <sub>16</sub>	ASI_BLK_COMMIT_S	RW	Any	—	<i>data_access_exception</i> <sup>3</sup>	
E4 <sub>16</sub> – E9 <sub>16</sub>			Any	—	<i>data_access_exception</i>	
EC <sub>16</sub> – EF <sub>16</sub>			Any	—	<i>data_access_exception</i>	
F0 <sub>16</sub>	ASI_BLK_P	RW	Any	—	64-byte block load/store, primary address	5.8
F1 <sub>16</sub>	ASI_BLK_S	RW	Any	—	64-byte block load/store, secondary address	5.8
F2 <sub>16</sub> – F7 <sub>16</sub>			Any	—	<i>data_access_exception</i>	
F8 <sub>16</sub>	ASI_BLK_PL	RW	Any	—	64-byte block load/store, primary address, little endian	5.8
F9 <sub>16</sub>	ASI_BLK_SL	RW	Any	—	64-byte block load/store, secondary address, little endian	5.8
FA <sub>16</sub> – FF <sub>16</sub>			Any	—	<i>data_access_exception</i>	

† This ASI name has been changed, for consistency; although use of this name is deprecated and software should use the new name, the old name is listed here for compatibility.

1. ASIs C0<sub>16</sub>–C5<sub>16</sub>, C8<sub>16</sub>–CD<sub>16</sub>, D0<sub>16</sub>–D3<sub>16</sub>, D8<sub>16</sub>–DB<sub>16</sub>, and E0<sub>16</sub>–E1<sub>16</sub> are checked for a VA watchpoint and will generate a *VA\_Watchpoint* trap if the watchpoint conditions are met. They are also checked for word-alignment and doubleword-alignment on STDFA, and will generate a *mem\_address\_not\_aligned* trap if the effective address (R[rs1] + R[rs2]; note that R[rs2] is not used as a mask) is not word-aligned or a *stdf\_mem\_address\_not\_aligned* trap if the address is word-aligned, but not doubleword-aligned.
2. ASIs D0<sub>16</sub>–D3<sub>16</sub> and D8<sub>16</sub>–DB<sub>16</sub> are checked for a VA watchpoint and will generate a *VA\_Watchpoint* trap if the watchpoint conditions are met. They are also checked for word-alignment and doubleword-alignment on STDFA and LDDFA, and will generate a *mem\_address\_not\_aligned* trap if the address is not word-aligned or a *stdf\_mem\_address\_not\_aligned/lddf\_mem\_address\_not\_aligned* trap if the address is word-aligned, but not doubleword-aligned.
3. ASIs E0<sub>16</sub>–E1<sub>16</sub> are checked for a VA watchpoint and will generate a *VA\_Watchpoint* trap if the watchpoint conditions are met. They are also checked for word-alignment and doubleword-alignment on STDFA, and will generate a *mem\_address\_not\_aligned* trap if the address is not word-aligned or a *stdf\_mem\_address\_not\_aligned* trap if the address is word-aligned, but not doubleword-aligned.

## 9.2.1 ASI\_REAL and ASI\_REAL\_LITTLE

This ASI is used to bypass the data MMU for memory addresses. Since the *cp* page attribute bit is clear, load accesses using this ASI will always fetch their data from the L2 cache. Using this ASI for an I/O address is permitted, and will follow the same page attributes (*w* = 1, all other attributes 0).

<b>Programming Note</b>	Although it is permitted to use <i>ASI_REAL</i> (or <i>ASI_REAL_LITTLE</i> ) for an I/O access, it is not recommended to do so because the <i>e</i> bit is not set for the access. <i>ASI_REAL_IO</i> and <i>ASI_REAL_IO_LITTLE</i> should be used instead.
-------------------------	---

## 9.2.2 ASI\_REAL\_IO and ASI\_REAL\_IO\_LITTLE

This ASI is used to bypass the data MMU for I/O addresses. The physical page attributes *e* and *w* are set to 1 and all other attribute bits are set to 0 for accesses to this ASI. Using this ASI for a memory address is permitted, and will follow the same page attributes (*e* = 1, *w* = 1, all other attributes 0).

<b>Note</b>	An atomic load-store operation is not permitted to these ASIs; an attempt to execute one will result in a <i>data_access_exception</i> exception.
-------------	---

## 9.2.3 ASI\_SCRATCHPAD

Each strand has a set of six privileged *ASI\_SCRATCHPAD* registers, accessed through *ASI* 20<sub>16</sub> with VA{63:0} = 0<sub>16</sub>, 8<sub>16</sub>, 10<sub>16</sub>, 18<sub>16</sub>, 30<sub>16</sub>, or 38<sub>16</sub>. These registers are for scratchpad use by privileged software. VA 20<sub>16</sub> and 28<sub>16</sub> may be used to access two additional scratchpad registers. However, access to those two scratchpad registers will be much slower than to the other six (because accesses to them will cause a trap and the access will be emulated).

TABLE 9-2 defines the format of these registers.

**TABLE 9-2** Scratchpad – ASI\_SCRATCHPAD (ASI 20<sub>16</sub>; VA 0<sub>16</sub>, 8<sub>16</sub>, 10<sub>16</sub>, 18<sub>16</sub>, 30<sub>16</sub>, or 38<sub>16</sub>)

Bit	Field	R/W	Description
63:0	scratchpad	RW	Scratchpad.

**Warning** | There is a known “feature” in UltraSPARC T1 that affects LDXA/STXA by privileged code to these ASI registers. If an immediately preceding instruction is a store that takes a trap, an LDXA can corrupt an unrelated IRF (integer register file) register, or a STXA may complete in spite of the trap. To prevent this, it is *required* to have a non-store or NOP instruction before any LDXA/STXA to this ASI. If the LDXA/STXA is at a branch target, there must be a non-store in the delay slot. Nonprivileged software is not affected by this.



# Performance Instrumentation

## 10.1 Performance Control Register

Each virtual processor has a privileged Performance Control register (PCR). Nonprivileged accesses to this register cause a *privileged\_opcode* trap. The performance control register contains six fields: *ovfh*, *ovfl*, *sl*, *ut*, *st*, and *priv*.

- *ovfh* and *ovfl* are state bits associated with the PIC.h and PIC.l overflow traps and are provided in this register to allow swapping out of a process that is in the state between the counter overflowing and the overflow trap being generated.
- *sl* controls which events are counted in PIC.l.
- *ut* controls whether user-level (nonprivileged) events are counted.
- *st* controls whether supervisor-level (privileged) events are counted.
- *priv* controls whether the PIC register can be read or written by nonprivileged software.

The format of this register is shown in TABLE 10-1. Note that changing the fields in PCR does not affect the PIC values. To change the events monitored, software needs to disable counting via PCR, reset the PIC, and then enable the new event via the PCR.

**TABLE 10-1** Performance Control Register – PCR (ASR 10<sub>16</sub>)

Bit	Field	Initial Value	R/W	Description
63:10	—	0	R	<i>Reserved</i>
9	<i>ovfh</i>	0	RW	If 1, PIC.h has overflowed, and the next count event will cause a disrupting trap to hyperprivileged software. The trap will appear to be precise to the instruction following the event.
8	<i>ovfl</i>	0	RW	If 1, PIC.l has overflowed
7	—	0	R	<i>Reserved</i>

**TABLE 10-1** Performance Control Register – PCR (ASR 10<sub>16</sub>)

Bit	Field	Initial Value	R/W	Description
6:4	sl	0	RW	Selects one of eight events to be counted for PIC.l, per TABLE 10-2.
3	—	0	R	<i>Reserved</i>
2	ut	0	RW	If ut = 1, count events in user mode; otherwise, ignore user mode events.
1	st	0	RW	If st = 1, count events in supervisor mode; otherwise, ignore supervisor mode events.
0	priv	0	RW	If priv = 1, prevent access to PIC by user-level code. If priv = 0, allow access to PIC by user-level code.

TABLE 10-2 contains the settings for the sl field.

**TABLE 10-2** sl Field Settings

Event Names	Encoding	PIC	Description
Instr_cnt	sl = XXX	H	Number of completed instructions. Annulled, mispredicted, or trapped instructions are not counted. <sup>1</sup>
SB_full	sl = 000	L	Number of store buffer full cycles. <sup>2</sup>
FP_instr_cnt	sl = 001	L	Number of completed floating-point instructions. <sup>3</sup> Annulled or trapped instructions are not counted.
IC_miss	sl = 010	L	Number of instruction cache (L1) misses.
DC_miss	sl = 011	L	Number of data cache (L1) misses for loads (store misses are not included as the cache is write-through, non-allocating).
ITLB_miss	sl = 100	L	Number of instruction TLB miss trap taken (includes real_translation misses).
DTLB_miss	sl = 101	L	Number of data TLB miss trap taken (includes real_translation misses).
L2_imiss	sl = 110	L	Number of secondary cache (L2) misses due to instruction cache requests.
L2_dmiss_ld	sl = 111	L	Number of secondary cache (L2) misses due to data cache load requests. <sup>4</sup>

1. Tcc instructions that are cancelled due to encountering a higher-priority trap are still counted.

2. SB\_full increments every cycle a strand (virtual processor) is stalled due to a full store buffer, regardless of whether other strands are able to keep the processor busy. The overflow trap for SB\_full is not precise to the instruction following the event that occurs when ovfl is set (the trap may occur on either the instruction following the event that occurs when ovfl is set, or on either of the next two instructions).

3. Only floating point instructions which execute in the shared FPU are counted. The following instructions are executed in the shared FPU: FADDs, FADDD, FSUBS, FSUBD, FMULS, FMULD, FDIVS, FDIVD, FSMULD, FSTOX, FDTOX, FXTOS, FXTOD, FITOS, FDTOS, FITOD, FSTOD, FSTOI, FDTOI, FCMPS, FCMPSD, FCMPE, FCMPEd.

4. L2 misses due to stores cannot be counted by the performance instrumentation logic.

---

## 10.2 SPARC Performance Instrumentation Counter

Each strand (virtual processor) has a Performance Instrumentation Counter register (PIC). Access privilege to PIC is controlled by the setting of PCR.priv. When PCR.priv = 1, a nonprivileged access to this register causes a *privileged\_action* trap. The PIC counter contains two fields, h and l. The PIC.h field always counts the number of completed instructions. The PIC.l field counts the event selected by PCR.sl.

The ut and st fields for PCR control whether events from user (nonprivileged) mode, supervisor (privileged) mode, both, or neither are counted. Whenever PCR.ovfh is set (which normally occurs when the PIC.h counter overflows, but may also be set via a write to the PCR), a disrupting trap is generated on the next event that increments the counter. This trap will appear to be precise to the instruction following the one that caused the event

**Programming Note** | A WRAsr to PCR that modifies the ovfh or ovfl bit behaves as if the ovfh or ovfl bit was modified before the WRAsr is executed.

This implies that if all of the following conditions are true, a performance counter overflow trap will be taken (to hyperprivileged software) on the instruction following the WRAsr:

- a WRAsr is executed in privileged mode that sets ovfh or ovfl to 1
- PCR.st = 1
- the WRAsr generates the event being counted

The format of the PIC register is shown in TABLE 10-3.

**TABLE 10-3** Performance Instrumentation Counter Register – PIC (ASR 11<sub>16</sub>)

Bit	Field	Initial Value	R/W	Description
63:32	h	0	RW	Instruction counter.
31:0	l	0	RW	Programmable event counter, event controlled by PCR.sl.



# Memory Management

---

---

## 11.1 Translation Table Entry (TTE)

The Translation Table Entry (TTE) holds information for a single page mapping. The TTE is broken 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 is fetched by software.

### 11.1.1 TTE Tag Format

UltraSPARC T1 supports both the UltraSPARC Architecture 2005 TTE tag format (as described in the *UltraSPARC Architecture 2005* specification; also known as the "sun4v" TTE format) and the older sun4u TTE tag format.

Note that UltraSPARC T1 only supports 13-bit context IDs; therefore, the most significant 3 bits of the (16-bit) context field are always zero.

UltraSPARC T1 supports 48-bit virtual addresses in hardware. When hardware writes a 48-bit virtual address into a 64-bit register, it sign-extends (copies) the most significant address bit (bit 47) into bits 63:48 of the register.

### 11.1.2 TTE Data Format

For the data portion of the TTE, both the sun4v and sun4u formats are supported by UltraSPARC T1. The sun4v TTE data format is described in the *UltraSPARC Architecture 2005* specification.

UltraSPARC Architecture 2005 specifies a 4-bit size field for TTE entries. Since UltraSPARC T1 only supports a 3-bit size field, the most significant bit of TTE (bit 3) is ignored when written.

In the sun4u TTE virtual address tag, bits 63:22 are used. Bits 21 through 13 are not maintained in the tag, since these bits are used to index the smallest direct-mapped TSB of 512 entries.

The sun4u TTE data format is shown in TABLE 11-1.

**TABLE 11-1** Format 16 Sun4u TTE Data Format

<b>Bit</b>	<b>Field</b>	<b>Description</b>
63	v	Valid
62:61	szl	size{1:0}
60	nfo	No-fault-only
59	ie	Invert endianness
58:49	soft2	Soft2
48	szh	size{2}
47:40	diag	Diagnostic
39:13	pa	PA{39:13}
12:8	soft	Soft
7	—	Reserved
6	l	Locked
5	cp	Cacheable in physically indexed cache
4	cv	Cacheable in virtually indexed cache
3	e	Side effect
2	p	Privileged
1	w	Writable
0	—	Reserved

TABLE 11-2 provides UltraSPARC T1-specific information regarding sun4u TTE data fields.

**TABLE 11-2** TTE Field Description

<b>Field</b>	<b>Description</b>
nfo	No-Fault-Only..
ie	Invert Endianness.
soft, soft2	Software-defined fields, provided for use by the operating system. Software fields are not implemented in UltraSPARC T1 hardware.
diag	Used by diagnostics
pa	The physical page number.
l	Lock. If this bit is set, the TTE entry will be “locked down”.
w	Writable.

---

## 11.2 Translation Storage Buffer

A Translation Storage Buffer (TSB) is an array of TTEs managed entirely by software. It serves as a cache of the Software Translation Table. The discussion in this section assumes the use of hardware support for TSB access, although the operating system is not required to make use of this support hardware.

Inclusion of the TLB entries in a TSB is not required; that is, translation information may exist in the TLB that is not present in the TSB.

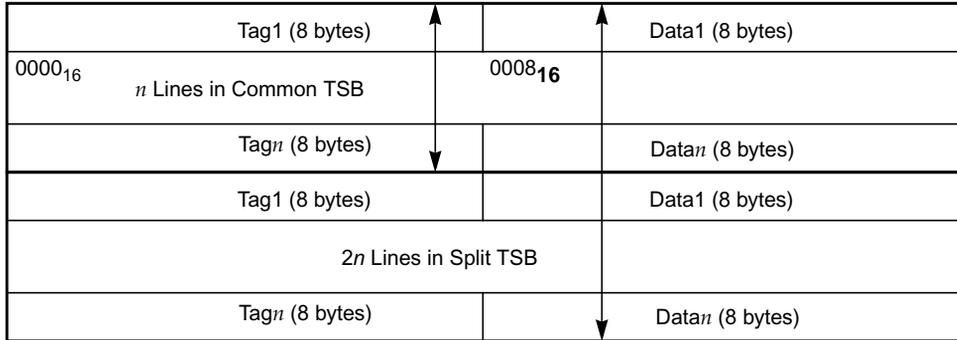
A TSB is arranged as a direct-mapped cache of TTEs. The  $n$  least significant bits of a virtual page number is used as the offset from the respective TSB base address, where  $n$  equals  $\log_2$  of the number of TTEs in the TSB.

A bit in the TSB register allows the PS0 and PS1 pointers to be computed for the case of separate or split PS0/PS1 TSB(s).

No hardware TSB indexing support is provided for TTEs of pages other than PS0 and PS1. Since the TSB is entirely software managed, however, the operating system may choose to place these different page TTEs in the TSB by forming the appropriate pointers. In addition, simple modifications to the PS0 and PS1 index pointers provided by the hardware allow formation of an M-way set-associative TSB, multiple TSBs per page size, and multiple TSBs per process.

The TSB exists as a normal data structure in memory, and therefore may 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 should provide a better overall solution than that provided by a fixed partitioning.

FIGURE 13-1 shows both the common and shared TSB organization. The constant  $n$  is determined by the size field in the TSB register; it may range from 512 entries to 16 M entries.



**FIGURE 11-1** TSB Organization

---

## 11.3 MMU-Related Faults and Traps

MMU traps are described in TABLE 11-3.

**TABLE 11-3** MMU Trap Description

Trap	Description
<i>data_access_exception</i>	Occurs when one of the following events (the D-MMU does not prioritize these and may set multiple bits) occurs: <ul style="list-style-type: none"><li>• The D-MMU detects a privilege violation for a data access; that is, an attempted access to a privileged page when <code>PSTATE.priv = 0</code>.</li><li>• A speculative (nonfaulting) load instruction issued to a page marked with the side-effect (<code>e</code>) bit = 1.</li><li>• An atomic instruction issued to an I/O address (that is, <code>VA{39} = 1</code>).</li><li>• An invalid LDA/STA ASI value, invalid virtual address, read to write-only register, or write to read-only register, but not for an attempted user access to a restricted ASI (see the <i>privileged_action</i> trap described below)</li><li>• An access with an ASI other than <code>ASI_&lt;PRIMARY,SECONDARY&gt;_NO_FAULT[_LITTLE]</code> to a page marked with the <code>nfo</code> (no-fault-only) bit.</li><li>• Virtual address out of range and <code>PSTATE.am</code> is not set. See <i>48-bit Virtual Address Space</i> on page 39 for details.</li></ul>
<i>instruction_access_exception</i>	Occurs when the I-MMU is enabled and one of the following happens: <ul style="list-style-type: none"><li>• The I-MMU detects a privilege violation for an instruction fetch; that is, an attempted access to a privileged page when <code>PSTATE.priv = 0</code>.</li><li>• Virtual address out of range and <code>PSTATE.am</code> is not set. See <i>48-bit Virtual Address Space</i> on page 39. Note that the case of JMWPL/RETURN and branch-CALL-sequential are handled differently.</li></ul>
<i>mem_address_not_aligned</i>	Occurs when a load, store, atomic, or JMWPL/RETURN instruction with a misaligned address is executed.
<i>privileged_action</i>	Occurs when an access is attempted using a <i>restricted</i> ASI while in nonprivileged mode ( <code>PSTATE.priv = 0</code> ).
<i>VA_watchpoint</i>	Occurs when virtual watchpoints are enabled and the D-MMU detects a load or store to the virtual address specified by the VA Data Watchpoint register.



# Implementation Dependencies

---

---

## 12.1 SPARC V9 General Information

### 12.1.1 Level-2 Compliance (Impl. Dep. #1)

UltraSPARC T1 is designed to meet Level-2 SPARC V9 compliance. It does the following:

- Correctly interprets all nonprivileged operations, and
- Correctly interprets all privileged elements of the architecture.

**Note** | System emulation routines (for example, for quad-precision floating-point operations) shipped with UltraSPARC T1 also must be Level-2 compliant.

### 12.1.2 Unimplemented Opcodes, ASIs, and ILLTRAP

SPARC V9 unimplemented instructions, *reserved* instructions, ILLTRAP opcodes, and instructions with invalid values in *reserved* fields (other than *reserved* FPOps and the *reserved* field in the Tcc instruction) encountered during execution cause an *illegal\_instruction* trap. Reserved FPOps cause an *fp\_exception\_other* (with `FSR.ftt = unimplemented_FPop`) trap. Unimplemented and *reserved* ASI values cause a *data\_access\_exception* trap.

### 12.1.3 Trap Levels (Imp. Dep. #37, 38, 39, 40, 101, 114, 115)

UltraSPARC T1 supports two trap levels; that is, `MAXPTL = 2`. Normal execution is at `TL = 0`.

A strand normally executes at trap level 0 (`execute_state`, `TL = 0`). In SPARC V9, a trap makes the CPU enter the next higher trap level, which is a fast and efficient process because there is one set of trap state registers for each trap level. After saving the most important machine states (`PC`, `NPC`, `PSTATE`) on the trap stack at this level, the trap (or error) condition is processed.

## 12.1.4 Trap Handling (Imp. Dep. #16, 32, 33, 35, 36, 44)

UltraSPARC T1 supports precise trap handling for all operations except for disrupting traps from hardware failures and interrupts. UltraSPARC T1 implements precise traps, interrupts, and exceptions for all instructions, including long latency floating-point operations.

UltraSPARC T1 can efficiently execute kernel code even in the event of multiple nested traps, promoting processor efficiency while dramatically reducing the system overhead needed for trap handling. Three sets of global registers are provided (`MAXPGL = 2`), for use at `TL = 0`, `TL = 1`, and `TL = 2`.

This further increases OS performance, providing fast trap execution by avoiding the need to save and restore registers while processing exceptions.

All traps supported in UltraSPARC T1 are listed in the “Traps” chapter of this document.

## 12.1.5 Population Count Instruction (POPC)

The population count instruction, `POPC`, generates an *illegal\_instruction* exception and is emulated in software rather than being executed in hardware.

## 12.1.6 Secure Software

To establish an enhanced security environment, it may be necessary to initialize certain strand states between contexts. Examples of such states are the contents of integer and floating-point register files, condition codes, and state registers. See also *Clean Window Handling (Impl. Dep. #102)*.

## 12.1.7 Address Masking (Impl. Dep. #125)

When `PSTATE.am = 1`, the `CALL`, `JMPL`, and `RDPC` instructions and all traps transmit zero in the high-order 32-bits of the `PC` to their specified destination registers. Traps also transmit zero in the high-order 32-bits of the `NPC` to the `TNPC`. Branch target addresses sent to the `NPC` and the updating of `NPC` with `NPC+4` for a

non-control-transferring instruction do not zero the high-order 32-bits. Restoration of PC and NPC from TPC and TNPC on a DONE or RETRY instruction do not mask the high-order 32-bits.

**Note** | When `PSTATE.am = 1`, address masking applies to all VAs, even those that immediately do a VA-to-RA bypass. This implies that with `PSTATE.am = 1`, `RA{63:32}` will be zeros after a VA-to-RA bypass.

---

## 12.2 SPARC V9 Integer Operations

### 12.2.1 Integer Register File and Window Control Registers (Impl. Dep. #2)

UltraSPARC T1 implements an eight-window 64-bit integer register file; that is, `N_REG_WINDOWS = 8`. UltraSPARC T1 truncates values stored in the `CWP`, `CANSAVE`, `CANRESTORE`, `CLEANWIN`, and `OTHERWIN` registers to three bits. This includes implicit updates to these registers by `SAVE(D)` and `RESTORE(D)` instructions. The most-significant two bits of these registers read as zero.

### 12.2.2 SAVE Instruction

Upon a `SAVE` instruction, UltraSPARC T1 initializes the values of the local registers in the new window to the same values as the local registers in the old window and initializes the values of the *out* registers in the new window to the same values as the *in* registers in the old window (that is, the new window matches the old window with the *ins* and *outs* swapped). Since this implies that they contain values from the executing process, V9 compliance is maintained. In this sense, the behavior of the `SAVE` instruction on UltraSPARC T1 differs from most other SPARC V9 implementations.<sup>1</sup>

<sup>1</sup> Most SPARC V9 processors do not initialize the *local* and *out* registers on a save instruction; instead, the values in the *local* and *out* registers are those left there from the last time the window was used.

## 12.2.3 Clean Window Handling (Impl. Dep. #102)

SPARC V9 introduced the concept of “clean window” to enhance security and integrity during program execution. A clean window is defined to be a register window that contains either all zeroes or addresses and data that belong to the current context. The CLEANWIN register records the number of available clean windows.

When a SAVE instruction requests a window and there are no more clean windows, a *clean\_window* trap is generated. Note that the behavior on a *clean\_window* trap for UltraSPARC T1 is the same as for a SAVE instruction, namely, the *local* registers for the new window remain the same as the *local* registers from the old window, while the *out* registers in the new window contain the contents of the *in* registers from the old window. Thus, while UltraSPARC T1 generates a *clean\_window* trap, the new window is automatically cleaned by hardware. System software only needs to increment CLEANWIN before returning to the requesting context.

## 12.2.4 Integer Multiply and Divide

Integer multiplications (MULSc, SMUL{cc}, MULX) and divisions (SDIV{cc}, UDIV{cc}, UDIVX) are executed directly in hardware.

## 12.2.5 MULSc

SPARC V9 does not define the value of xcc and R[rd]{63:32} for MULSc. UltraSPARC T1 sets xcc and rd based on the results of adding either (32 copies of R[rs1]{63}:: CCR.icc.n xor CCR.icc.v, R[rs1]{31:1}) or 0 (depending on Y{0}) to either R[rs2]{63:0} or the immediate operand.

---

## 12.3 SPARC V9 Floating-Point Operations

### 12.3.1 Subnormal Operands and Results: Nonstandard Operation

UltraSPARC T1 handles all cases of subnormal operands or results directly in hardware.

Because there is no trapping on subnormal operands, UltraSPARC T1 does not support the nonstandard result option of the SPARC V9 architecture, and the FSR.ns bit ignores any value written to it and always returns zero on a read.

## 12.3.2 Overflow, Underflow, and Inexact Traps (Impl. Dep. #3, 55)

UltraSPARC T1 implements precise floating-point exception handling. Underflow is detected before rounding.

**Note** | Major performance degradation may be observed while running with the inexact exception enabled.

## 12.3.3 Quad-Precision Floating-Point Operations (Impl. Dep. #3)

All quad-precision floating-point instructions, listed in TABLE 12-1, cause an *fp\_exception\_other* (with FSR.ftt = 3, unimplemented\_FPop) trap. These operations are emulated in system software.

**TABLE 12-1** Unimplemented Quad-Precision Floating-Point Instructions

Instruction	Description
F{s,d}TOq	Convert single-/double- to quad-precision floating-point
F{i,x}TOq	Convert 32-/64-bit integer to quad-precision floating-point
FqTO{s,d}	Convert quad- to single-/double-precision floating-point
FqTO{i,x}	Convert quad-precision floating-point to 32-/64-bit integer
FCMP{E}q	Quad-precision floating-point compares
FMOVq	Quad-precision floating-point move
FMOVqcc	Quad-precision floating-point move, if condition is satisfied
FMOVqr	Quad-precision floating-point move if register match condition
FABSq	Quad-precision floating-point absolute value
FADDq	Quad-precision floating-point addition
FDIVq	Quad-precision floating-point division
FdMULq	Double- to quad-precision floating-point multiply
FMULq	Quad-precision floating-point multiply

**TABLE 12-1** Unimplemented Quad-Precision Floating-Point Instructions (*Continued*)

<b>Instruction</b>	<b>Description</b>
FNEGq	Quad-precision floating-point negation
FSQRTq	Quad-precision floating-point square root
FSUBq	Quad-precision floating-point subtraction

## 12.3.4 Floating-Point Square Root

The three floating-point square root instructions: FSQRTS, FSQRTD, FSQRTQ are unimplemented. Execution of any of these instructions results in an *fp\_exception\_other* exception, with FSR.ftt= unimplemented\_FPop.

## 12.3.5 Floating-Point Upper and Lower Dirty Bits in FPRS Register

The FPRS\_dirty\_upper (du) and FPRS\_dirty\_lower (dl) bits in the Floating-Point Registers State (FPRS) register are set when an instruction that modifies the corresponding upper and lower half of the floating-point register file is issued. Floating-point register file modifying instructions include floating-point operate, graphics, floating-point loads, and block load instructions.

The FPRS.du and FPRS.dl may be set pessimistically, even though the instruction that modified the floating-point register file is nullified due to a trap. This includes the case where the floating-point instruction itself takes a *fp\_disabled* trap.

## 12.3.6 Floating-Point State Register (FSR) (Impl. Dep. #13, 19, 22, 23, 24)

UltraSPARC T1 supports precise-traps and implements all three exception fields (*tem*, *cexc*, and *aexc*) conforming to IEEE Standard 754-1985. TABLE 12-2 defines the register bits.

**TABLE 12-2** Floating-Point Status Register Format

Bits	Field	RW	Description										
63:38	—	R	<i>Reserved</i>										
37:36	<i>fcc3</i>	RW	Floating-point condition code (set 3). One of four sets of 2-bit floating-point condition codes, which are modified by the FCMP{E} (and LD{X}FSR) instructions. The FBfcc, FMOVcc, and MOVcc instructions use one of these condition code sets to determine conditional control transfers and conditional register moves.										
35:34	<i>fcc2</i>	RW	Floating-point condition code (set 2). See <i>fcc3</i> .										
33:32	<i>fcc1</i>	RW	Floating-point condition code (set 1). See <i>fcc3</i> .										
31:30	<i>rd</i>	RW	IEEE Std. 754-1985 rounding direction. Rounding modes are shown below.										
			<table border="1"> <thead> <tr> <th><i>rd</i></th> <th>Round Toward</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Nearest (even if tie)</td> </tr> <tr> <td>1</td> <td>0</td> </tr> <tr> <td>2</td> <td><math>+\infty</math></td> </tr> <tr> <td>3</td> <td><math>-\infty</math></td> </tr> </tbody> </table>	<i>rd</i>	Round Toward	0	Nearest (even if tie)	1	0	2	$+\infty$	3	$-\infty$
<i>rd</i>	Round Toward												
0	Nearest (even if tie)												
1	0												
2	$+\infty$												
3	$-\infty$												
29:28	—	R	<i>Reserved</i>										
27:23	<i>tem</i>	RW	5-bit trap enable mask for the IEEE-754 floating-point exceptions. If a floating-point operate instruction produces one or more exceptions, the corresponding <i>cexc/aexc</i> bits are set and an <i>fp_exception_ieee_754</i> (with FSR.ftt = 1, IEEE_754_exception) exception is generated.										
22	<i>ns</i>	R	Nonstandard floating-point results. Always 0: UltraSPARC T1 produces IEEE-754 compatible results.										
21:20	—	R	<i>Reserved</i>										
19:17	<i>ver</i>	R	FPU version number. Identifies a particular implementation of the UltraSPARC T1 FPU architecture.										
16:14	<i>ftt</i>	R	Floating-point trap type. The 3-bit floating point trap type field is set whenever an floating-point instruction causes the <i>fp_exception_ieee_754</i> or <i>fp_exception_other</i> traps. Trap types are listed in TABLE 12-4, below.										
13:	<i>qne</i>	R	Floating-point deferred-trap queue (FQ) not empty. Not used, because UltraSPARC T1 implements precise floating-point exceptions.										
12	—	R	<i>Reserved</i>										

**TABLE 12-2** Floating-Point Status Register Format (Continued)

Bits	Field	RW	Description
11:10	fcc0	RW	Floating-point condition code (set 0). See fcc3. Note: fcc0 is the same as fcc in SPARC V8.
9:5	aexc	RW	5-bit accrued exception field. Accumulates IEEE 754 exceptions while floating-point exception traps are disabled (that is, FSR.tem = 0).
4:0	cexc	RW	5-bit current exception field indicates the most recently generated IEEE 754 exceptions.

**Note** | fcc0 is the same as the fcc in SPARC V8.

**TABLE 12-3** Floating-Point Rounding Modes

rd	Round Toward
0	Nearest (even if tie)
1	0
2	$+\infty$
3	$-\infty$

**TABLE 12-4** Floating-Point Trap Type Values

ftt	Floating-Point Trap Type	Trap Signaled
0	None	—
1	IEEE_754_exception	<i>fp_exception_ieee_754</i>
2	unfinished_FPop	—
3	unimplemented_FPop	<i>fp_exception_other</i>
4	sequence_error	—
5	hardware_error	—
6	invalid_fp_register	—
7	reserved	—

- Notes**
- (1) UltraSPARC T1 neither detects nor generates the *unfinished\_FPop*, *sequence\_error*, *hardware\_error* or *invalid\_fp\_register* trap types directly in hardware.
  - (2) UltraSPARC T1 does not contain an FQ. An attempt to read the FQ with a RDPR instruction causes an *illegal\_instruction* trap.

---

## 12.4 SPARC V9 Memory-Related Operations

### 12.4.1 Load/Store Alternate Address Space (Impl. Dep. #5, 29, 30)

Supported ASI accesses are listed in *Alternate Address Spaces* on page 41.

### 12.4.2 Read/Write ASR (Impl. Dep. #6, 7, 8, 9, 47, 48)

Supported ASRs are discussed in *Ancillary State Registers (ASRs)* on page 7.

### 12.4.3 FLUSH and Self-Modifying Code (Impl. Dep. #122)

FLUSH is needed to synchronize code and data spaces after code space is modified during program execution. FLUSH is described in *Supported Memory Models* on page 36. On UltraSPARC T1, the FLUSH effective address is ignored, and as a result, FLUSH can not cause a *data\_access\_exception* trap.

**Note** SPARC V9 specifies that the FLUSH instruction has no latency on the issuing strand. In other words, a store to instruction space prior to the FLUSH instruction is visible immediately after the completion of FLUSH. MEMBAR #StoreStore is required to ensure proper ordering in multiprocessing system when the memory model is not TSO. When a MEMBAR #StoreStore, FLUSH sequence is performed, UltraSPARC T1 guarantees that earlier code modifications will be visible across the whole system.

## 12.4.4 PREFETCH{A} (Impl. Dep. #103, 117)

For UltraSPARC T1, PREFETCH{A} instructions follow TABLE 12-5 based on the fcn value. All prefetches in UltraSPARC T1 are weak (on an MMU miss or when the MMU is bypassed the prefetch is dropped). The only trap that a prefetch can generate on UltraSPARC T1 is *illegal\_instruction* (for fcn = 5<sub>16</sub>-F<sub>16</sub>).

**TABLE 12-5** PREFETCH{A} Variants

fcn	Prefetch Function	Action
0 <sub>16</sub>	Weak prefetch for several reads	Weak prefetch into Level 2 cache
1 <sub>16</sub>	Weak prefetch for one read	
2 <sub>16</sub>	Weak prefetch for several writes	
3 <sub>16</sub>	Weak prefetch for one write	
4 <sub>16</sub>	Prefetch Page	No operation
5 <sub>16</sub> -F <sub>16</sub>	-	<i>Illegal_instruction</i> trap.
10 <sub>16</sub>	Invalidate read-once prefetch	Weak prefetch into Level 2 cache
11 <sub>16</sub>	Prefetch for read to nearest unified cache	Weak prefetch into Level 2 cache
12 <sub>16</sub> -13 <sub>16</sub>	Strong prefetches	Weak prefetch into Level 2 cache
14 <sub>16</sub>	Strong prefetch for several reads	Weak prefetch into Level 2 cache
15 <sub>16</sub>	Strong prefetch for one read	
16 <sub>16</sub>	Strong prefetch for several writes	
17 <sub>16</sub>	Strong prefetch for one write	
18 <sub>16</sub> -1F <sub>16</sub>	—	No operation

PREFETCHA is legal for all implemented ASIs in UltraSPARC T1 and will prefetch into the Level 2 cache from memory using the context listed in TABLE 12-6. Prefetching is done regardless of privilege level (for example, user mode can use ASI 10<sub>16</sub> to prefetch into the L2 cache).

**TABLE 12-6** PREFETCH{A} ASIs

Context	ASIs (hexadecimal)
Primary	10, 16, 18, 1E, 22, 2A, 80, 82, 88, 8A, C0, C2, C4, C8, CA, CC, D0, D2, D8, DA, E0, E2, EA, F0, F8
Secondary	11, 17, 19, 1F, 23, 2B, 81, 83, 89, 8B, C1, C3, C5, C9, CB, CD, D1, D3, D9, DB, E1, E3, EB, F1, F9
Nucleus	04, 0C, 14, 15, 1C, 1D, 20, 21, 24, 25, 26, 27, 2C, 2E, 2F, 31, 32, 33, 35, 36, 37, 39, 3A, 3B, 3D, 3E, 3F, 40, 42, 43, 44, 45, 46, 47, 4B, 4C, 4D, 4F, 50, 51, 52, 54, 55, 56, 57, 58, 59, 5A, 5B, 5C, 5D, 5E, 5F, 60, 66, 67, 72, 73, 74

<b>Implementation Note</b>	Although it would have been desirable to treat PREFETCHA to restricted ASIs by underprivileged code as NOPs, PREFETCH only moves data between main memory and the L2 cache, so UltraSPARC T1's implementation causes no security issues.
----------------------------	--

## 12.4.5 Instruction Prefetch

UltraSPARC T1 does not implement an instruction prefetch. No prefetching is performed from the effective address of the BPN instruction.

## 12.4.6 LDTW/STTW Handling (Impl. Dep. #107, 108)

LDTW and STTW instructions are directly executed in hardware.

<b>Note</b>	LDTW/STTW were deprecated in SPARC V9. In UltraSPARC T1, it is more efficient to use LDX/STX for accessing 64-bit data. LDTW/STTW take longer to execute than two 32-/64-bit loads/stores.
-------------	--

## 12.4.7 Floating-Point *mem\_address\_not\_aligned* (Impl. Dep. #109, 110, 111, 112)

LDDF{A}/STDF{A} cause an LDDF/STDF *mem\_address\_not\_aligned* trap if the effective address is 32-bit aligned but not 64-bit (doubleword) aligned.

LDQF{A}/STQF{A} are not directly executed in hardware; they cause an *illegal\_instruction* trap.

## 12.4.8 Supported Memory Models (Impl. Dep. #113, 121)

UltraSPARC T1 supports only the TSO memory model, although certain specific operations such as block loads and stores operate under the RMO memory model. See *Supported Memory Models* on page 36.

## 12.4.9 Implicit ASI when TL > 0 (Impl. Dep. #124)

UltraSPARC T1 matches all UltraSPARC Architecture implementations, and makes the implicit ASI for instruction fetching `ASI_NUCLEUS` when `TL > 0`, while the implicit ASI for loads and stores when `TL > 0` is `ASI_NUCLEUS` if `PSTATE.cle = 0` or `ASI_NUCLEUS_LITTLE` if `PSTATE.cle = 1`.

**Compatibility Note** | With an implicit ASI for instruction fetching of `ASI_NUCLEUS`, if software was to set the strand in a state where `PSTATE.priv = 0` but `TL > 0`, an instruction fetch will generate an *instruction\_access\_exception*, because user-level code is accessing `ASI_NUCLEUS`. UltraSPARC I/II overrides this *instruction\_access\_exception* and allows instruction fetching when `PSTATE.priv = 0` and `TL > 0`. UltraSPARC T1 is compatible with UltraSPARC I/II and does the same override of *instruction\_access\_exception* when `PSTATE.priv = 0` and `TL > 0`.

---

## 12.5 Non-SPARC V9 Extensions

### 12.5.1 Cache Subsystem

UltraSPARC T1 contains one or more levels of cache. The cache subsystem architecture is described in Appendix F, *Caches and Cache Coherency*.

### 12.5.2 Block Memory Operations

UltraSPARC T1 supports 64-byte block memory operations utilizing a block of eight double-precision floating point registers as a temporary buffer. See *Block Load and Store Instructions* on page 18.

### 12.5.3 Partial Stores

UltraSPARC T1 does not support 8-/16-/32-bit partial stores to memory.

### 12.5.4 Short Floating-Point Loads and Stores

UltraSPARC T1 does not supports 8-/16-bit loads and stores to the floating-point registers.

## 12.5.5 Interrupt Vector Handling

CPUs and I/O devices can interrupt a selected CPU by assembling and sending an interrupt packet. This allows hardware interrupts and cross calls to have the same hardware mechanism and to share a common software interface for processing. Interrupt vectors are described in Chapter 7, *Interrupt Handling*.

## 12.5.6 Power-Down Support

UltraSPARC T1 supports the ability to power down virtual processors and I/O devices to reduce power requirements during idle periods.

## 12.5.7 UltraSPARC T1 Instruction Set Extensions (Impl. Dep. #106)

The UltraSPARC T1 CPU supports a subset of the VIS 1.0 and 2.0 instructions; see *UltraSPARC Architecture 2005 Instructions Not Directly Implemented by UltraSPARC T1 Hardware* on page 13.

Unimplemented IMPDEP1 and IMPDEP2 opcodes encountered during execution cause an *illegal\_instruction* trap.

## 12.5.8 Performance Instrumentation

UltraSPARC T1 performance instrumentation is described in *Performance Control Register* on page 49 and *SPARC Performance Instrumentation Counter* on page 51.



# Assembly Language Syntax

---

The assembly language syntax used in this document follows that described in the "Assembly Language Syntax" appendix of the *UltraSPARC Architecture 2005* specification.



# Programming Guidelines

---

## B.1 Multithreading

In UltraSPARC T1, execution is switched in round-robin fashion every cycle among the strands that are ready to issue another instruction. Context switching is built into the UltraSPARC T1 pipeline and takes place during the SWITCH stage, thus contexts are switched each cycle with no pipeline stall penalty.

The following instructions change a strand from a ready-to-issue state to a not-ready-to-issue state, until hardware determines that their input/execution requirements can be satisfied:

- All branches (including CALL, JMPL, etc.)
- All VIS instructions
- All floating point (FPops)
- All WRPR, WR
- All RDPR, RD
- SAVE(D), RESTORE(D), RETURN, FLUSHW (all register management)
- All MUL and DIV
- MULSCC
- MEMBAR #Sync, MEMBAR #StoreLoad, MEMBAR #MemIssue
- FLUSH
- All loads
- All floating-point memory operations
- All memory operations to alternate space
- All atomics load-store operations
- Prefetch

---

## B.2 Pipeline Strand Flush

The front end of the UltraSPARC T1 pipeline prevents instructions from being issued to the rest of the pipeline unless there is a high probability (for most instructions, a probability of 1.0) of the instruction having all its input dependencies satisfied. For certain instructions, the input dependencies cannot be determined by the front end, and the instruction (and any subsequent instructions issued from that strand) need to be flushed from the pipeline and replayed. TABLE B-1 lists instructions that may end up causing a strand flush.

**TABLE B-1** Pipeline Strand Flush Events

Event	Strand Flush Description
Loads	The strand will be flushed if the load encounters a cache miss while executing with STRAND_STS_REG.spec_en = 1.
multiply/divide/ floating-point operate	Resource contention can cause a strand flush.
store buffer full	The strand will be flushed until space is available in the store buffer.
trap	Instruction and any subsequent instructions in the pipeline from that strand are flushed, and fetching restarts at the trap vector.

---

## B.3 Instruction Latencies

TABLE B-2 lists the single-strand instruction latencies for UltraSPARC T1. When multiple strands are executing, much of the additional latency for multicycle instructions will be overlapped with execution of the additional strands.

In this table, certain opcodes are marked with mnemonic superscripts. These superscripts and their meanings are defined in TABLE 5-1 on page 13.

**TABLE B-2** Instruction Latencies (1 of 14)

Instruction	Latency	Comments
ADD	1	
ADD <sup>c</sup>	1	
ADD <sup>CC</sup>	1	
ADD <sup>Ncc</sup>	1	

**TABLE B-2** Instruction Latencies (2 of 14)

<b>Instruction</b>	<b>Latency</b>	<b>Comments</b>
AND	1	
ANDcc	1	
ANDN	1	
ANDNcc	1	
BA	3?	
BA_A	4	
BA_A_PN	4?	
BA_PN	3	
BA_XCC	3	
BA_XCC_A	4?	
BA_XCC_A_PN	4	
BA_XCC_PN	3?	
BCC	3	
BCC_A	3	
BCC_A_PN	4	
BCC_PN	3	
BCC_XCC	3	
BCC_XCC_A	3	
BCC_XCC_A_PN	4	
BCC_XCC_PN	3?	
BCS	3	
BCS_A	4	
BCS_A_PN	4?	
BCS_PN	3	
BCS_XCC	3	
BCS_XCC_A	4	
BCS_XCC_A_PN	4	
BCS_XCC_PN	4?	
BE	3	
BE_A	4?	

**TABLE B-2** Instruction Latencies (3 of 14)

<b>Instruction</b>	<b>Latency</b>	<b>Comments</b>
BE_A_PN	4	
BE_PN	3	
BE_XCC	3	
BE_XCC_A	4	
BE_XCC_A_PN	4	
BE_XCC_PN	3	
BG	3	
BG_A	3	
BG_A_PN	4	
BG_PN	3	
BG_XCC	3	
BG_XCC_A	3	
BG_XCC_A_PN	3?	
BG_XCC_PN	3	
BGE	3	
BGE_A	3	
BGE_A_PN	4?	
BGE_PN	3	
BGE_XCC	3	
BGE_XCC_A	4	
BGE_XCC_A_PN	4?	
BGE_XCC_PN	3	
BGU	3	
BGU_A	3	
BGU_A_PN	3?	
BGU_PN	3	
BGU_XCC	3	
BGU_XCC_A	4	
BGU_XCC_A_PN	4	
BGU_XCC_PN	3	

**TABLE B-2** Instruction Latencies (4 of 14)

<b>Instruction</b>	<b>Latency</b>	<b>Comments</b>
BL	3	
BL_A	3	
BL_A_PN	3?	
BL_PN	3	
BL_XCC	3	
BL_XCC_A	4?	
BL_XCC_A_PN	4	
BL_XCC_PN	3	
BLE	3	
BLE_A	3	
BLE_A_PN	3?	
BLE_PN	3	
BLE_XCC	3	
BLE_XCC_A	4	
BLE_XCC_A_PN	4?	
B LE_XCC_PN	3	
BLEU	3	
BLEU_A	4?	
BLEU_A_PN	4	
BLEU_PN	3?	
BLEU_XCC	3	
BLEU_XCC_A	4	
BLEU_XCC_A_PN	4	
BLEU_XCC_PN	3	
BN	3	
BN_A	4	
BN_A_PN	4	
BN_PN	3	
BN_XCC	3	
BN_XCC_A	4?	

**TABLE B-2** Instruction Latencies (5 of 14)

<b>Instruction</b>	<b>Latency</b>	<b>Comments</b>
BN_XCC_A_PN	4	
BN_XCC_PN	3	
BNE	3	
BNE_A	3	
BNE_A_PN	3?	
BNE_PN	3	
BNE_XCC	3?	
BNE_XCC_A	4	
BNE_XCC_A_PN	4	
BNE_XCC_PN	3?	
BNEG	3	
BNEG_A	4	
BNEG_A_PN	4	
BNEG_PN	3	
BNEG_XCC	3	
BNEG_XCC_A	4	
BNEG_XCC_A_PN	4	
BNEG_XCC_PN	3	
BPOS	3	
BPOS_A	4?	
BPOS_A_PN	4	
BPOS_PN	3	
BPOS_XCC	3	
BPOS_XCC_A	4?	
BPOS_XCC_A_PN	4	
BPOS_XCC_PN	3	
BRGEZ	3	
BRGEZ_A	4?	
BRGEZ_A_PN	4	
BRGEZ_PN	3	

**TABLE B-2** Instruction Latencies (6 of 14)

<b>Instruction</b>	<b>Latency</b>	<b>Comments</b>
BRGZ	3	
BRGZ_A	4?	
BRGZ_A_PN	4	
BRGZ_PN	3	
BRLEZ	3	
BRLEZ_A	4	
BRLEZ_A_PN	4	
BRLEZ_PN	3	
BRLZ	3	
BRLZ_A	3	
BRLZ_A_PN	3?	
BRLZ_PN	3	
BRNZ	3	
BRNZ_A	4	
BRNZ_A_PN	4	
BRNZ_PN	3	
BRZ	3	
BRZ_A	4	
BRZ_A_PN	4	
BRZ_PN	3	
BVC	3	
BVC_A	4?	
BVC_A_PN	4	
BVC_PN	3	
BVC_XCC	3	
BVC_XCC_A	4	
BVC_XCC_A_PN	4?	
BVC_XCC_PN	3	
BVS	3	
BVS_A	4	

**TABLE B-2** Instruction Latencies (7 of 14)

Instruction	Latency	Comments
BVS_A_PN	4	
BVS_PN	3	
BVS_XCC	3	
BVS_XCC_A	4	
BVS_XCC_A_PN	4	
BVS_XCC_PN	3	
CASA <sup>PASI</sup>	39	performed in L2
CASXA <sup>PASI</sup>	39	performed in L2
FABSd	8	
FABSs	21	
FADDd	26	
FADDs	26	
FBA	3	
FBA_A	4?	
FBA_A_PN	4	
FBA_PN	3	
FBE	3	
FBE_A	4	
FBE_A_PN	4	
FBE_PN	3	
FBG	3	
FBG_A	4	
FBG_A_PN	4?	
FBG_PN	3	
FBGE	3	
FBGE_A	4	
FBGE_A_PN	4	
FBGE_PN	3	
FBL	3	
FBL_A	4?	

**TABLE B-2** Instruction Latencies (8 of 14)

<b>Instruction</b>	<b>Latency</b>	<b>Comments</b>
FBL_A_PN	4	
FBL_PN	3	
FBLE	3	
FBLE_A	4	
FBLE_A_PN	4	
FBLE_PN	3	
FBLG	3	
FBLG_A	4	
FBLG_A_PN	4	
FBLG_PN	3	
FBN	3?	
FBN_A	4	
FBN_A_PN	4	
FBN_PN	3	
FBNE	3	
FBNE_A	4	
FBNE_A_PN	4?	
FBNE_PN	3	
FBUE	3	
FBUE_A	4	
FBUE_A_PN	4	
FBUE_PN	3	
FBUG	3	
FBUG_A	4	
FBUG_A_PN	4	
FBUG_PN	3	
FBUGE	3	
FBUGE_A	4	
FBUGE_A_PN	4	
FBUGE_PN	3	

**TABLE B-2** Instruction Latencies (9 of 14)

<b>Instruction</b>	<b>Latency</b>	<b>Comments</b>
FBUL	3	
FBUL_A	4	
FBUL_A_PN	4	
FBUL_PN	3	
FBULE	3	
FBULE_A	4	
FBULE_A_PN	4	
FBULE_PN	3	
FDIVd	83	
FDIVs	54	
FdTOi	25	
FdTOs	25	
FdTOx	25	
FiTOd	25	
FiTOs	26	
FMOVd	8	
FMOVDA		
FMOVDE	8	
FMOV DG	8	
FMOV DGE	8	
FMOV DL	8	
FMOV DLE	8	
FMOV DLG	8	
FMOV DN	8	
FMOV DNE	8	
FMOV DO	8	
FMOV DU	8	
FMOV DUE	8	
FMOV DUG	8	
FMOV DUGE	8	



**TABLE B-2** Instruction Latencies (11 of 14)

<b>Instruction</b>	<b>Latency</b>	<b>Comments</b>
FxTOs	26	
LD_FP	9	
LDFSR	9	
LDD_FP	9?	
LDD_FPD	9	
LDDFA <sup>P<sub>AsI</sub></sup>	9	
LDSB	22	performed in L2
LDSBA	21	performed in L2
LDSH	3	
LDSHA <sup>P<sub>AsI</sub></sup>	3	
LDSTUB	37	performed in L2
LDSTUBA <sup>P<sub>AsI</sub></sup>	37	performed in L2
LDSW	3	
LDSWA	3	
LDUB	3	
LDUBA	3	
LDUH	3	
LDUHA <sup>P<sub>AsI</sub></sup>	3	
LDUW	3	
LDUWA <sup>P<sub>AsI</sub></sup>	3	
LDX	3	
LDX_FSR	27	
LDXA <sup>P<sub>AsI</sub></sup>	3	
MOVA	1	
MOVA_FCC	1	
MOVCC	1	
MOVCS	1	
MOVE	1	
MOVE_FCC	1	
MOVG	1?	

**TABLE B-2** Instruction Latencies (12 of 14)

<b>Instruction</b>	<b>Latency</b>	<b>Comments</b>
MOVG_FCC	1	
MOVGE	1	
MOVGE_FCC	1	
MOVGU	1?	
MOVL	1	
MOVL_FCC	1	
MOVLE	1	
MOVLE_FCC	1	
MOVLEU	1	
MOVLG_FCC	1	
MOVN	1	
MOVN_FCC	1	
MOVNE	1	
MOVNE_FCC	1	
MOVNEG	1?	
MOVNO_FCC	1	
MOVPOS	1	
MOVRE	1	
MOVRGEZ	1	
MOVRGZ	1	
MOVRLEZ	1	
MOVRLZ	1	
MOVRNE	1	
MOVU_FCC	1	
MOVUE_FCC	1	
MOVUG_FCC	1	
MOVUGE_FCC	1	
MOVUL_FCC	1	
MOVULE_FCC	1	
MOVVC	1	

**TABLE B-2** Instruction Latencies (13 of 14)

Instruction	Latency	Comments
MOVVS	1	
MULSCC	7	
MULX	11	
OR	1	
ORcc	1	
ORN	1	
ORNcc	1	
RD_CCR	4	
RDASI	4	
RD_FPRS	4	
RD_Y	4	
SDIV <sup>D</sup>	72	
SDIVcc <sup>D</sup>	72	
SDIVX	72	
SETHI	1	
SLL	1	
SLLX	1	
SMUL <sup>D</sup>	11	
SMULcc <sup>D</sup>	11	
SRA	1?	
SRAX	1	
SRL	1	
SRLX	1	
STFA <sup>PASI</sup>	8	
ST_FSR	8	
STFA <sup>PASI</sup>	8	
STB	1	
STBA	4	
STDF	8	
STDA_FP	8?	

**TABLE B-2** Instruction Latencies (14 of 14)

<b>Instruction</b>	<b>Latency</b>	<b>Comments</b>
STDA_FP_ASI	8	
STH	1	
STHA	4	
STW	1?	
STWA	1?	
STX	1	
STX_FSR	8	
STXA	1-? (4-?)?	varies, depending on ASI
SUB	1	
SUBC	1	
SUBcc	1	
SUBCcc	1	
SWAP <sup>D</sup>	49	performed in L2
SWAPA <sup>D, P<sub>ASI</sub></sup>	37	performed in L2
TADDcc	1?	
TADDccTV <sup>D</sup>	1	
TSUBcc	1	
TSUBccTV <sup>D</sup>	1	
UDIV <sup>D</sup>	72	
UDIVcc <sup>D</sup>	72?	
UDIVX	72	
UMUL <sup>D</sup>	11	
UMULcc <sup>D</sup>	11	
WR_CCR	9	
WRASI	9	
WR_FPRS	9	
XNOR	1	
XNORcc	1	
XOR	1	
XORcc	1	

---

## B.4 Grouping Rules

Each physical cores in UltraSPARC T1 are single-issue, so there are no grouping rules for UltraSPARC T1.

---

## B.5 Floating-Point Operations

UltraSPARC T1 supports hardware floating-point operations, but since one floating-point unit (FPU) is shared among 8 physical cores (32 strands), there are limitations on dispatch of floating-point instructions to the FPU. Each physical processor core (four strands) can have a single floating-point instruction outstanding at any given time. For the purpose of this restriction, floating-point instructions include floating-point operations, VIS floating-point operations, floating-point loads and stores, and block loads and stores.

---

## B.6 Synchronization

UltraSPARC T1 has two varieties of instructions for synchronization: memory barriers and flush. The following memory barrier instructions ensure that any load, store, or atomic memory operation issued after it take effect after all memory operations issued before it:

- MEMBAR with `mmask{1} = 1` (MEMBAR #StoreLoad)
- MEMBAR with `cmask{1} = 1` (MEMBAR #MemIssue)
- MEMBAR with `cmask{2} = 1` MEMBAR #Sync)

All other types of membar instructions are treated as NOPs, since they are implied by the TSO memory ordering protocol followed by UltraSPARC T1.

However, the memory barriers do not guarantee that the instruction caches on UltraSPARC T1 have become consistent with the preceding memory operations. A FLUSH instruction guarantees that in addition to the preceding memory operations taking effect in the global memory system, all the instruction caches on UltraSPARC T1 are consistent with these operations. It also ensures that the instruction fetch buffer for the strand issuing the flush has become consistent with the preceding memory operations.

Thus, when one strand is modifying the instructions of another, the “producer” strand should

1. Complete all necessary modifications
2. Issue a FLUSH
3. Signal completion to the “consumer” strand

Completion may be signalled by a store/atomic instruction which modifies a predetermined location, or by issuing an interrupt to the consumer strand.

The consumer strand at this point should make sure that its instruction fetch buffer (of size 2 entries) becomes consistent with the global memory system. This can be done by either

1. Issuing a branch to the modified location; or
2. Issuing a flush instruction; or
3. Waiting for a two-instruction gap, to allow the two instructions in the fetch buffer to drain.

In the case of a branch, the delay slot is not guaranteed to be consistent with global memory. The branch is a better option than flush for high performance.

Note that the HALT instruction is not meant to be a synchronization instruction and should *not* be used as such. For example, the following code, which uses halt to make sure func\_A executes before func\_B, may cause T0 to hang:

<u>T0</u>	<u>T1</u>
st X	ld X
halt	do_func_A
do_func_B	intr T0



# Opcode Maps

This appendix contains the UltraSPARC T1 instruction opcode maps.

Opcodes marked with a dash (—) are reserved; an attempt to execute a reserved opcode causes a trap unless the opcode is an implementation-specific extension to the instruction set.

In this appendix, certain opcodes are marked with mnemonic superscripts. These superscripts and their meanings are defined in TABLE 5-1 on page 13.

In the tables in this appendix, *reserved* (—) and shaded entries indicate opcodes that are not implemented in the UltraSPARC T1 processor.

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).

**TABLE C-1** op{1:0}

op {1:0}			
0	1	2	3
Branches and SETHI <i>See</i> TABLE C-2.	CALL	Arithmetic and Miscellaneous <i>See</i> TABLE C-3	Loads/Stores <i>See</i> TABLE C-4

**TABLE C-2** op2{2:0} (op = 0)

op2 {2:0}							
0	1	2	3	4	5	6	7
ILLTRAP	BPcc – <i>See</i> TABLE C-7	Bicc <sup>D</sup> – <i>See</i> TABLE C-7	BPr – <i>See</i> TABLE C-8	SETHI NOP <sup>†</sup>	FBPfcc – <i>See</i> TABLE C-7	FBfcc <sup>D</sup> – <i>See</i> TABLE C-7	—

<sup>†</sup>rd = 0, imm22 = 0

The ILLTRAP and *reserved* (—) encodings generate an *illegal\_instruction* trap.

**TABLE C-3** op3{5:0} (op = 2) (1 of 3)

		op3 {5:4}			
		0	1	2	3
<b>op3 {3:0}</b>	<b>0</b>	ADD	ADDcc	TADDcc	WRY <sup>D</sup> (rd = 0) — (rd = 1) WRCCR (rd = 2) WRASI (rd = 3) — (rd = 4, 5) — (rd = 5, rd = 0, rd = 1) WRFPRS (rd = 6) — (rd = 7-14) — (rd = 15 and (rs1 > 0 or i = 0)) WRPCR <sup>P</sup> (rd = 16, rd = 0) — (rd = 16, rd = 1) WRPIC (rd = 17, rd = 0) — (rd = 17, rd = 1) — (rd = 18) WRGSR (rd = 19) WRSOFTINT_SET <sup>P</sup> (rd = 0) WRSOFTINT_CLR <sup>P</sup> (rd = 21) WRSOFTINT <sup>P</sup> (rd = 22) WRTICK_CMPR <sup>P</sup> (rd = 23) WRSTICK_CMPR (rd = 25) WR %asr26 (rd = 26, rd = 1) — (rd = 26, rd = 0) — (rd = 27-31))
	<b>1</b>	AND	ANDcc	TSUBcc	SAVED <sup>P</sup> (fcn = 0), RESTORED <sup>P</sup> (fcn = 1) — (fcn > 1)
	<b>2</b>	OR	ORcc	TADDccTV <sup>D</sup>	WRPR <sup>P</sup> — (rd = 15, 17-31)
	<b>3</b>	XOR	XORcc	TSUBccTV <sup>D</sup>	— (rs1 = 2, 4, 7-30)
	<b>4</b>	SUB	SUBcc	MULScc <sup>D</sup>	FPop1 – See TABLE C-5
	<b>5</b>	ANDN	ANDNcc	SLL (x = 0), SLLX (x = 1)	FPop2 – See TABLE C-6
	<b>6</b>	ORN	ORNcc	SRL (x = 0), SRLX (x = 1)	IMPDEP1 (VIS) – See TABLE C-12
	<b>7</b>	XNOR	XNORcc	SRA (x = 0), SRAX (x = 1)	IMPDEP2

**TABLE C-3** op3{5:0} (op = 2) (2 of 3)

		op3 {5:4}			
		0	1	2	3
<b>op3 {3:0}</b>	<b>8</b>	ADDC	ADDCcc	RDY <sup>D</sup> (rs1 = 0, i = 0) — (rs1 = 0, i = 1) — (rs1 = 1) RDCCR (rs1 = 2, i = 0) — (rs1 = 2, i = 1) RDASI (rs1 = 3, i = 0) — (rs1 = 3, i = 1) RTICK <sup>P<sub>npt</sub></sup> (rs1 = 4, i = 0) — (rs1 = 4, i = 1) RDPIC (rs1 = 5, i = 0) — (rs1 = 5, i = 1) RDPIC (rs1 = 6, i = 0) — (rs1 = 6, i = 1) — (rs1 = 7-14) MEMBAR (rs1 = 15, rd = 0, i = 1) STBAR <sup>D</sup> (rs1 = 15, rd = 0, i = 0) — (rs1 = 15, rd > 0) RDPIC <sup>P</sup> (rs1 = 16) RDPIC (rs1 = 17) — (rs1 = 18) RDGSR (rs1 = 19, i = 0) — (rs1 = 19, i = 1) — (rs1 = 20, 21) RDSOFTINT <sup>P</sup> (rs1 = 22, i = 0) — (rs1 = 22, i = 1) RTICK_CMPR <sup>P</sup> (rs1 = 23, i = 0) — (rs1 = 23, i = 1) RDSTICK <sup>P</sup> (rs1 = 24, i = 0) — (rs1 = 24, i = 1) RDSTICK_CMPR <sup>P</sup> (rs1 = 25, i = 0) — (rs1 = 25, i = 1) rd %asr26 (rs1 = 26) — (rs1 = 27 - 31)	JMPL

**TABLE C-3** op3{5:0} (op = 2) (3 of 3)

		op3 {5:4}			
		0	1	2	3
<b>op3 {3:0}</b>	<b>9</b>	MULX	—	— (rs1 = 2, 4, 7 - 30)	RETURN
	<b>A</b>	UMUL <sup>D</sup>	UMULcc <sup>D</sup>	RDPR <sup>P</sup> — (rs1 = 15, 17 - 30)	Tcc {(i = 0 and inst{10:5} = 0) or i = 1 and inst{10:8} = 0)}—See TABLE C-7 and TABLE C-11. — {(i = 0 and inst{10:5} > 0) or i = 1 and inst{10:8} > 0)}
	<b>B</b>	SMUL <sup>D</sup>	SMULcc <sup>D</sup>	FLUSHW	FLUSH
	<b>C</b>	SUBC	SUBCcc	MOVcc See TABLE C-9	SAVE
	<b>D</b>	UDIVX	—	SDIVX	RESTORE
	<b>E</b>	UDIV <sup>D</sup>	UDIVcc <sup>D</sup>	POPC (rs1 = 0) — (rs1 > 0)	DONE <sup>P</sup> (fcn = 0) RETRY <sup>P</sup> (fcn = 1) — (fcn > 1)
	<b>F</b>	SDIV <sup>D</sup>	SDIVcc <sup>D</sup>	MOVr See TABLE C-8	—

Shaded and the reserved (—) opcodes cause an illegal\_instruction trap.

TABLE C-4 op3{5:0} (op = 3)

		op3{5:4}			
		0	1	2	3
<b>op3 {3:0}</b>	0	LDUW	LDUWA <sup>P<sub>ASI</sub></sup>	LDF	LDFA <sup>P<sub>ASI</sub></sup>
	1	LDUB	LDUBA <sup>P<sub>ASI</sub></sup>	LDFSR <sup>D</sup> , LDXFSR — (rd > 1)	—
	2	LDUH	LDUHA <sup>P<sub>ASI</sub></sup>	LDQF	LDQFA <sup>P<sub>ASI</sub></sup>
	3	LDD <sup>D</sup> — (rd odd)	LDDA <sup>D, P<sub>ASI</sub></sup> — (rd odd)	LDDF	LDDFA <sup>P<sub>ASI</sub></sup> See 8.6.4 XREF
	4	STW	STWA <sup>P<sub>ASI</sub></sup>	STF	STFA <sup>P<sub>ASI</sub></sup>
	5	STB	STBA <sup>P<sub>ASI</sub></sup>	STFSR <sup>D</sup> , STXFSR — (rd > 1)	—
	6	STH	STHA <sup>P<sub>ASI</sub></sup>	STQF	STQFA <sup>P<sub>ASI</sub></sup>
	7	STD <sup>D</sup> — (rd odd)	STDA <sup>P<sub>ASI</sub></sup> — (rd odd)	STDF	STDFA <sup>P<sub>ASI</sub></sup> See 8.6.4 XREF
	8	LDSW	LDSWA <sup>P<sub>ASI</sub></sup>	—	—
	9	LDSB	LDSBA <sup>P<sub>ASI</sub></sup>	—	—
	A	LDSH	LDSHA <sup>P<sub>ASI</sub></sup>	—	—
	B	LDX	LDXA <sup>P<sub>ASI</sub></sup>	—	—
	C	—	—	—	CASA <sup>P<sub>ASI</sub></sup>
	D	LDSTUB	LDSTUBA <sup>P<sub>ASI</sub></sup>	PREFETCH — (fcn = 5–15)	PREFETCHA <sup>P<sub>ASI</sub></sup> — (fcn = 5–15)
	E	STX	STXA <sup>P<sub>ASI</sub></sup>	—	CASXA <sup>P<sub>ASI</sub></sup>
	F	SWAP <sup>D</sup>	SWAPA <sup>D, P<sub>ASI</sub></sup>	—	—

LDQF, LDQFA, STQF, STQFA, and the *reserved* (—) opcodes cause an *illegal\_instruction* trap.

TABLE C-5 opf{8:0} (op = 2, op3 = 34<sub>16</sub> = FPop1)

opf{8:3}	opf{2:0}							
	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>	—	FABS <sub>s</sub>	FABS <sub>d</sub>	FABS <sub>q</sub>	—	—	—	—
02 <sub>16</sub>	—	—	—	—	—	—	—	—
03 <sub>16</sub>	—	—	—	—	—	—	—	—
04 <sub>16</sub>	—	—	—	—	—	—	—	—
05 <sub>16</sub>	—	FSQRT <sub>s</sub>	FSQRT <sub>d</sub>	FSQRT <sub>q</sub>	—	—	—	—
06 <sub>16</sub>	—	—	—	—	—	—	—	—
07 <sub>16</sub>	—	—	—	—	—	—	—	—
08 <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>
09 <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>
0A <sub>16</sub>	—	—	—	—	—	—	—	—
0B <sub>16</sub>	—	—	—	—	—	—	—	—
0C <sub>16</sub>	—	—	—	—	—	—	—	—
0D <sub>16</sub>	—	FsMUL <sub>d</sub>	—	—	—	—	FdMUL <sub>q</sub>	—
0E <sub>16</sub>	—	—	—	—	—	—	—	—
0F <sub>16</sub>	—	—	—	—	—	—	—	—
10 <sub>16</sub>	—	FsTO <sub>x</sub>	FdTO <sub>x</sub>	FqTO <sub>x</sub>	FxTO <sub>s</sub>	—	—	—
11 <sub>16</sub>	FxTO <sub>d</sub>	—	—	—	FxTO <sub>q</sub>	—	—	—
12 <sub>16</sub>	—	—	—	—	—	—	—	—
13 <sub>16</sub>	—	—	—	—	—	—	—	—
14 <sub>16</sub>	—	—	—	—	—	—	—	—
15 <sub>16</sub>	—	—	—	—	—	—	—	—
16 <sub>16</sub>	—	—	—	—	—	—	—	—
17 <sub>16</sub>	—	—	—	—	—	—	—	—
18 <sub>16</sub>	—	—	—	—	FiTO <sub>s</sub>	—	FdTO <sub>s</sub>	FqTO <sub>s</sub>
19 <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>	—
1A <sub>16</sub>	—	FsTO <sub>i</sub>	FdTO <sub>i</sub>	FqTO <sub>i</sub>	—	—	—	—
1B <sub>16</sub> –3F <sub>16</sub>	—	—	—	—	—	—	—	—

Shaded and *reserved* (—) opcodes cause an *fp\_exception\_other* trap with FSR.ftt = 3 (unimplemented\_FPop).

TABLE C-6 opf{8:0} (op = 2, op3 = 35<sub>16</sub> = FPop2)

opf{8:4}	opf{3:0}								
	0	1	2	3	4	5	6	7	8-F
00	—	FMOV <sub>s</sub> (fcc0)	FMOV <sub>d</sub> (fcc0)	FMOV <sub>q</sub> (fcc0)	—	†	†	†	—
01	—	—	—	—	—	—	—	—	—
02	—	—	—	—	—	FMOV <sub>s</sub> Z	FMOV <sub>d</sub> Z	FMOV <sub>q</sub> Z	—
03	—	—	—	—	—	—	—	—	—
04	—	FMOV <sub>s</sub> (fcc1)	FMOV <sub>d</sub> (fcc1)	FMOV <sub>q</sub> (fcc1)	—	FMOV <sub>s</sub> LEZ	FMOV <sub>d</sub> LEZ	FMOV <sub>q</sub> LEZ	—
05	—	FCMP <sub>s</sub>	FCMP <sub>d</sub>	FCMP <sub>q</sub>	—	FCMP <sub>e</sub> <sub>s</sub>	FCMP <sub>e</sub> <sub>d</sub>	FCMP <sub>e</sub> <sub>q</sub>	—
06	—	—	—	—	—	FMOV <sub>s</sub> LZ	FMOV <sub>d</sub> LZ	FMOV <sub>q</sub> LZ	—
07	—	—	—	—	—	—	—	—	—
08	—	FMOV <sub>s</sub> (fcc2)	FMOV <sub>d</sub> (fcc2)	FMOV <sub>q</sub> (fcc2)	—	†	†	†	—
09	—	—	—	—	—	—	—	—	—
0A	—	—	—	—	—	FMOV <sub>s</sub> NZ	FMOV <sub>d</sub> NZ	FMOV <sub>q</sub> NZ	—
0B	—	—	—	—	—	—	—	—	—
0C	—	FMOV <sub>s</sub> (fcc3)	FMOV <sub>d</sub> (fcc3)	FMOV <sub>q</sub> (fcc3)	—	FMOV <sub>s</sub> GZ	FMOV <sub>d</sub> GZ	FMOV <sub>q</sub> GZ	—
0D	—	—	—	—	—	—	—	—	—
0E	—	—	—	—	—	FMOV <sub>s</sub> GEZ	FMOV <sub>d</sub> GEZ	FMOV <sub>q</sub> GEZ	—
0F	—	—	—	—	—	—	—	—	—
10	—	FMOV <sub>s</sub> (icc)	FMOV <sub>d</sub> (icc)	FMOV <sub>q</sub> (icc)	—	—	—	—	—
11–17	—	—	—	—	—	—	—	—	—
18	—	FMOV <sub>s</sub> (xcc)	FMOV <sub>d</sub> (xcc)	FMOV <sub>q</sub> (xcc)	—	—	—	—	—
19–1F	—	—	—	—	—	—	—	—	—

† Reserved variation of FMOV<sub>r</sub>

Shaded and *reserved* (—) opcodes cause an *fp\_exception\_other* trap with FSR.ftt = 3 (unimplemented\_FPop).

TABLE C-7 cond{3:0}

		BPcc	Bicc <sup>D</sup>	FBPfcc	FBfcc <sup>D</sup>	Tcc
		op = 0 op2 = 1	op = 0 op2 = 2	op = 0 op2 = 5	op = 0 op2 = 6	op = 2 op3 = 3a <sub>16</sub>
<b>cond {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 C-8 Encoding of rcond{2:0} Instruction Field

		BPr	MOVr	FMOVr
		op = 0 op2 = 3	op = 2 op3 = 2f <sub>16</sub>	op = 2 op3 = 35 <sub>16</sub>
<b>rcond {2:0}</b>	<b>0</b>	—	—	—
	<b>1</b>	BRZ	MOVRZ	FMOVRZ
	<b>2</b>	BRLEZ	MOVRLEZ	FMOVRLEZ
	<b>3</b>	BRLZ	MOVRLZ	FMOVRLZ
	<b>4</b>	—	—	—
	<b>5</b>	BRNZ	MOVRNZ	FMOVRNZ
	<b>6</b>	BRGZ	MOVRGZ	FMOVRGZ
	<b>7</b>	BRGEZ	MOVRGEZ	FMOVRGEZ

**TABLE C-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 C-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 C-11** cc Fields (BPcc and Tcc)

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

TABLE C-12 VIS Opcodes op = 2, op3 = 36<sub>16</sub> = IMPDEP1

		opf {3:0}							
		0	1	2	3	4	5	6	7
opf {8:4}	00	EDGE8	EDGE8N	EDGE8L	EDGE8LN	EDGE16	EDGE16N	EDGE16L	EDGE16LN
	01	ARRAY8		ARRAY16		ARRAY32			
	02	FCMPLE16		FCMPNE16		FCMPLE32		FCMPNE32	
	03		FMUL8X16		FMUL8X16AU		FMUL8X16AL	FMUL8SUX16	FMUL8ULX16
	04								
	05	FPADD16	FPADD16S	FPADD32	FPADD32S	FPSUB16	FPSUB16S	FPSUB32	FPSUB32S
	06	FZERO	FZEROS	FNOR	FNORS	FANDNOT2	FANDNOT2S	FNOT2	FNOT2S
	07	FAND	FANDS	FXNOR	FXNORS	FSRC1	FSRC1S	FORNOT2	FORNOT2S
	08	SHUTDOWN	SIAM						
	09..1F								

		opf {3:0}							
		8	9	A	B	C	D	E	F
opf {8:4}	00	EDGE32	EDGE32N	EDGE32L	EDGE32LN				
	01	ALIGN ADDRESS	BMASK	ALIGNADDRESS_LITTLE					
	02	FCMPGT16		FCMPEQ16		FCMPGT32		FCMPEQ32	
	03	FMULD8SUX16	FMULD8ULX16	FPACK32	FPACK16		FPACKFIX	PDIST	
	04	FALIGNDATA			FPMERGE	BSHUFFLE	FEXPAND		
	05								
	06	FANDNOT1	FANDNOT1S	FNOT1	FNOT1S	FXOR	FXORS	FNAND	FNANDS
	07	FSRC2	FSRC2S	FORNOT1	FORNOT1S	FOR	FORS	FONE	FONES
	08								
	09..1F								

**Note** | An *illegal\_instruction* exception is generated if the undefined or shaded opcodes in the IMPDEP1 space are used.

# Instructions and Exceptions

---

The instructions supported by UltraSPARC T1 and the exceptions they generate are listed in the UltraSPARC Architecture 2005 specification.



# IEEE 754 Floating Point Support

---

UltraSPARC T1 conforms to the SPARC V9 Appendix B (IEEE Std 754-1985 Requirements for SPARC-V9) recommendations.

**Note** | UltraSPARC T1 detects tininess before rounding.

---

## E.1 Special Operand Handling

The UltraSPARC T1 FPU provides full hardware support for subnormal operands and results. Unlike UltraSPARC I/II and UltraSPARC III, UltraSPARC T1 will never generate an unfinished\_FPop trap type. Also, unlike UltraSPARC I/II and UltraSPARC III, UltraSPARC T1 does not implement a nonstandard floating-point mode. The ns bit of the FSR is always read as 0, and writes to it are ignored.

The FPU generates +inf, -inf, +largest number, -largest number (depending on round mode) for overflow cases for multiply, divide, and add operations.

For higher-to-lower precision conversion instructions {FDTOS}:

- overflow, underflow, and inexact exceptions can be raised.
- overflow is treated the same way as an unrounded add result; depending on the round mode, we will either generate the properly signed infinity or largest number.
- underflow will produce a signed zero, smallest number, or subnormal result.

For conversion to integer instructions {F(s,d)TOi, F(s,d)TOx}: UltraSPARC T1 follows SPARC V9 appendix B.5, pg 246.

For NaN's: UltraSPARC T1 Follows SPARC V9 appendix B.2 (particularly Table 27) and B.5, pg 244-246.

- Please note that Appendix B applies to those instructions listed in IEEE 754 section 5: "All conforming implementations of this standard shall provide operations to add, subtract, multiply, divide, extract the sqrt, find the remainder, round to integer in fp format, convert between different fp formats, convert

between fp and integer formats, convert binary<->decimal, and compare. Whether copying without change of format is considered an operation is an implementation option.”

- The instructions involving copying/moving of fp data (FMOV, FABS, and FNEG) will follow earlier UltraSPARC implementations by doing the appropriate sign bit transformation but will not cause an invalid exception nor do a rs2 = SNaN to rd = QNaN transformation.
- Following UltraSPARC I/II implementations, all Fpops as defined in V9 will update cexc. All other instructions will leave cexc unchanged.
- Following SPARC V9 Manual 5.1.7.6, 5.1.7.8, 5.1.7.9, and figures in 5.1.7.10 Overflow Result is defined as:

If the appropriate trap enable masks are not set (FSR.ofm = 0 and FSR.nxm = 0), then set aexc and cexc overflow and inexact flags: FSR.ofa = 1, FSR.nxa = 1, FSR.ofc = 1, FSR.nxc = 1. No trap is generated.

If any or both of the appropriate trap enable masks are set (FSR.ofm = 1 or FSR.nxm = 1), then only an IEEE overflow trap is generated: FSR.ftt = 1. The particular cexc bit that is set diverges from UltraSPARC I/II to follow the SPARC V9 section 5.1.7.9 errata:

If FSR.ofm = 0 and FSR.nxm = 1, then FSR.nxc = 1.

If FSR.ofm = 1, independent of FSR.nxm, then FSR.ofc = 1 and FSR.nxc = 0.

- Following SPARC V9 Manual 5.1.7.6, 5.1.7.8, 5.1.7.9, and figures in 5.1.7.10 Underflow Result is defined as:

If the appropriate trap enable masks are not set (FSR.ufm = 0 and FSR.nxm = 0), then set aexc and cexc underflow and inexact flags: FSR.ufa = 1, FSR.nxa = 1, FSR.ufc = 1, FSR.nxc = 1. No trap is generated.

If any or both of the appropriate trap enable masks are set (FSR.ufm = 1 or FSR.nxm = 1), then only an IEEE underflow trap is generated: FSR.ftt = 1. The particular cexc bit that is set diverges from UltraSPARC I/II to follow the SPARC V9 section 5.1.7.9 errata:

If FSR.ufm = 0 and FSR.nxm = 1, then FSR.nxc = 1.

If FSR.ufm = 1, independent of FSR.nxm, then FSR.ufc = 1 and FSR.nxc = 0.

The remainder of this section gives examples of special cases to be aware of that could generate various exceptions.

## E.1.1 Infinity Arithmetic

Let “num” be defined as unsigned in the following tables.

## E.1.1.1 One Infinity Operand Arithmetic

- Do not generate exceptions

**TABLE E-1** One Infinity Operations That Do Not Generate Exceptions

---

**Cases**

---

+inf plus +num = +inf  
+inf plus -num = +inf  
-inf plus +num = -inf  
-inf plus -num = -inf

+inf minus +num = +inf  
+inf minus -num = +inf  
-inf minus +num = -inf  
-inf minus -num = -inf

+inf multiplied by +num = +inf  
+inf multiplied by -num = -inf  
-inf multiplied by +num = -inf  
-inf multiplied by -num = +inf

+inf divided by +num = +inf  
+inf divided by -num = -inf  
-inf divided by +num = -inf  
-inf divided by -num = +inf

+num divided by +inf = +0  
+num divided by -inf = -0  
-num divided by +inf = -0  
-num divided by -inf = +0

fstod, fdtos (+inf) = +inf  
fstod, fdtos (-inf) = -inF

+inf divided by +0 = +inf  
+inf divided by -0 = -inf  
-inf divided by +0 = -inf  
-inf divided by -0 = +inf

---

**TABLE E-1** One Infinity Operations That Do Not Generate Exceptions (*Continued*)

Cases
Any arithmetic operation involving infinity as 1 operand and a QNaN as the other operand: SPARC V9 B.2.2 Table 27 (+/- inf) OPERATOR (QNaN2) = QNaN2 (QNaN1) OPERATOR (+/- inf) = QNaN1
Compares when other operand is not a NaN treat infinity just like a regular number: +inf = +inf, +inf > anything else; -inf = -inf, -inf < anything else. Affects following instructions: V9 fp compares (rs1 and/or rs2 could be +/- inf): * FCMPE * FCMP
Compares when other operand is a QNaN, SPARC V9 A.13, B.2.1; fcc value = unordered = 2'b11 fcmp(s/d) (+/- inf) with (QNaN2) - no invalid exception fcmp(s/d) (QNaN1) with (+/- inf) - no invalid exception

- Could generate exceptions

**TABLE E-2** One Infinity Operations That Could Generate Exceptions

Cases	Possible Exception	Result (in addition to accrued exception) if tem is cleared
SPARC V9 Appendix B.5 <sup>1</sup>	IEEE_754 7.1	
F{s,d}TOi (+inf) = invalid	IEEE_754 invalid	2 <sup>31</sup> -1
F{s,d}TOx (+inf) = invalid	IEEE_754 invalid	2 <sup>63</sup> -1
F{s,d}TOi (-inf) = invalid	IEEE_754 invalid	-2 <sup>31</sup>
F{s,d}TOx (-inf) = invalid	IEEE_754 invalid	-2 <sup>63</sup>
SPARC V9 B.2.2	IEEE_754 7.1	(No NaN operand result)
+inf multiplied by +0 = invalid	IEEE_754 invalid	QNaN
+inf multiplied by -0 = invalid	IEEE_754 invalid	QNaN
-inf multiplied by +0 = invalid	IEEE_754 invalid	QNaN
-inf multiplied by -0 = invalid	IEEE_754 invalid	QNaN

**TABLE E-2** One Infinity Operations That Could Generate Exceptions (*Continued*)

<b>Cases</b>	<b>Possible Exception</b>	<b>Result (in addition to accrued exception) if tem is cleared</b>
SPARC V9 B.2.2 Table 27 <sup>2</sup> Any arithmetic operation involving infinity as 1 operand and a SNaN as the other operand except copying/moving data	IEEE_754 7.1	(One operand, a SNaN)
(+/- inf) OPERATOR (SNaN2)	IEEE_754 invalid	QNaN2
(SNaN1) OPERATOR (+/- inf)	IEEE_754 invalid	QNaN1
SPARC V9 A.13, B.2.1 <sup>2</sup> Any compare operation involving infinity as 1 operand and a SNaN as the other operand:	IEEE_754 7.1	
FCMP(s/d) (+/- inf) with (SNaN2)	IEEE_754 invalid	fcc value = unordered = 2'b11
FCMP(s/d) (SNaN1) with (+/- inf)	IEEE_754 invalid	fcc value = unordered = 2'b11
FCMPE(s/d) (+/- inf) with (SNaN2)	IEEE_754 invalid	fcc value = unordered = 2'b11
FCMPE(s/d) (SNaN1) with (+/- inf)	IEEE_754 invalid	fcc value = unordered = 2'b11
SPARC V9 A.13 <sup>2</sup> Any compare & generate exception operation involving infinity as 1 operand and a QNaN as the other operand:	IEEE_754 7.1	
FCMPE(s/d) (+/- inf) with (QNaN2)	IEEE_754 invalid	fcc value = unordered = 2'b11
FCMPE(s/d) (QNaN1) with (+/- inf)	IEEE_754 invalid	fcc value = unordered = 2'b11

1. Similar invalid exceptions also included in SPARC V9 B.5 are generated when the source operand is a NaN(QNaN or SNaN) or a resulting number that cannot fit in 32b[64b] integer format: (large positive argument  $\geq 2^{31}[2^{63}]$  or large negative argument  $\leq -(2^{31} + 1)[-(2^{63} + 1)]$ )
2. Note that in the IEEE 754 standard, infinity is an exact number; so this exception could also apply to non-infinity operands as well. Also note that the invalid exception and SNaN to QNaN transformation does not apply to copying/moving frops (fmov,fabs,fneg).

## E.1.1.2 Two Infinity Operand Arithmetic

- Do not generate exceptions

**TABLE E-3** Two Infinity Operations That Do Not Generate Exceptions

Cases
+inf plus +inf = +inf
-inf plus -inf = -inf
+inf minus -inf = +inf
-inf minus +inf = -inf
+inf multiplied by +inf = +inf
+inf multiplied by -inf = -inf
-inf multiplied by +inf = -inf
-inf multiplied by -inf = +inf
Compares treat infinity just like a regular number: +inf = +inf, +inf > anything else; -inf = -inf, -inf < anything else. Affects following instructions: V9 fp compares (rs1 and/or rs2 could be +/- inf): * FCMPE * FCMP

- Could generate exceptions

**TABLE E-4** Two Infinity Operations That Generate Exceptions

Cases	Possible Exception	Result (in addition to accrued exception) if tem is cleared
SPARC V9 B.2.2	IEEE_754 7.1	(No NaN operand result)
+inf plus -inf = invalid	IEEE_754 invalid	QNaN
-inf plus +inf = invalid	IEEE_754 invalid	QNaN
+inf minus +inf = invalid	IEEE_754 invalid	QNaN
-inf minus -inf = invalid	IEEE_754 invalid	QNaN
SPARC V9 B.2.2	IEEE_754 7.1	(No NaN operand result)
+inf divided by +inf = invalid	IEEE_754 invalid	QNaN
+inf divided by -inf = invalid	IEEE_754 invalid	QNaN
-inf divided by +inf = invalid	IEEE_754 invalid	QNaN
-inf divided by -inf = invalid	IEEE_754 invalid	QNaN

## E.1.2 Zero Arithmetic

**TABLE E-5** Zero Arithmetic Operations That Generate Exceptions

Cases	Possible Exception	Result (in addition to accrued exception) if tem is cleared
SPARC V9 B.2.2 & 5.1.7.10.4 +0 divided by +0 = invalid +0 divided by -0 = invalid -0 divided by +0 = invalid -0 divided by -0 = invalid	IEEE_754 7.1 IEEE_754 invalid IEEE_754 invalid IEEE_754 invalid IEEE_754 invalid	(No NaN operand result) QNaN QNaN QNaN QNaN
SPARC V9 5.1.7.10.4 +num divided by +0 = divide by zero +num divided by -0 = divide by zero -num divided by +0 = divide by zero -num divided by -0 = divide by zero	IEEE_754 7.2 IEEE_754 div_by_zero IEEE_754 div_by_zero IEEE_754 div_by_zero IEEE_754 div_by_zero	+inf -inf -inf +inf
SPARC V9 B.2.2 Table 27 <sup>1</sup> Any arithmetic operation involving zero as 1 operand and a SNaN as the other operand except copying/moving data (+/- 0) OPERATOR (SNaN2) (SNaN1) OPERATOR (+/- 0)	IEEE_754 7.1  IEEE_754 invalid IEEE_754 invalid	(One operand, a SNaN)  QNaN2 QNaN1

1. In this context, 0 is again another exact number; so this exception could also apply to non-zero operands as well. Also note that the invalid exception and SNaN to QNaN transformation does not apply to copying/moving data instructions (FMOV, FABS, FNEG)

**TABLE E-6** Interesting Zero Arithmetic Sign Result Case

Cases
+0 plus -0 = +0 for all round modes except round to -infinity where the result is -0.

## E.1.3 NaN Arithmetic

- Do not generate exceptions

**TABLE E-7** NaN Arithmetic Operations that do not generate exceptions

---

**Cases**

---

SPARC V9 B.2.1: Fp convert to wider NaN transformation

FsTOd (QNaN2) = QNaN2 widened

FsTOd(0x7fd10000) = 0x7ffa2000 8'h0

FsTOd(0xffd10000) = 0xfffa2000 8'h0

SPARC V9 B.2.1: Fp convert to narrower NaN transformation

FdTOs (QNaN2) = QNaN2 narrowed

FdTOs(0x7ffa2000 8'h0) = 0x7fd1000

FdTOs(0xfffa2000 8'h0) = 0xffd1000

SPARC V9 B.2.2 Table 27

Any non-compare arithmetic operations --result takes sign of QNaN pass through operand.

(+/- num) OPERATOR (QNaN2) = QNaN2

(QNaN1) OPERATOR (+/- num) = QNaN1

(QNaN1) OPERATOR (QNaN2) = QNaN2

---

- Could generate exceptions

**TABLE E-8** NaN Arithmetic Operations That Could Generate Exceptions

Cases	Possible Exception	Result (in addition to accrued exception) if tem is cleared
SPARC V9 B.2.1: Fp convert to wider NaN transformation FsTOd (SNaN2) = QNaN2 widened FsTOd(0x7f910000) = 0x7ffa2000 8'h0 FsTOd(0xff910000) = 0xfffa2000 8'h0	IEEE_754 7.1  IEEE_754 invalid	  QNaN2 widened
SPARC V9 B.2.1: Fp convert to narrower NaN transformation FdTos (SNaN2) = QNaN2 narrowed FdTos(0x7ff22000 8'h0) = 0x7fd1000 FdTos(0xfff22000 8'h0) = 0xffd1000	IEEE_754 7.1  IEEE_754 invalid	  QNaN2 narrowed
SPARC V9 B.2.2 Table 27 Any non-compare arithmetic operations except copying/ moving (fmov, fabs, fneg) (+/- num) OPERATOR (SNaN2)	IEEE_754 7.1  IEEE_754 invalid	  QNaN2
(SNaN1) OPERATOR (+/- num)	IEEE_754 invalid	QNaN1
(SNaN1) OPERATOR (SNaN2)	IEEE_754 invalid	QNaN2
(QNaN1) OPERATOR (SNaN2)	IEEE_754 invalid	QNaN2
(SNaN1) OPERATOR (QNaN2)	IEEE_754 invalid	QNaN1
SPARC V9 Appendix B.5 F{s,d}TOi (+QNaN) = invalid F{s,d}TOi (+SNaN) = invalid F{s,d}TOx (+QNaN) = invalid F{s,d}TOx (+SNaN) = invalid	IEEE_754 7.1 IEEE_754 invalid IEEE_754 invalid IEEE_754 invalid IEEE_754 invalid	 $2^{31}-1$ $2^{31}-1$ $2^{63}-1$ $2^{63}-1$
F{s,d}TOi (-QNaN) = invalid F{s,d}TOi (-SNaN) = invalid F{s,d}TOx (-QNaN) = invalid F{s,d}TOx (-SNaN) = invalid	IEEE_754 invalid IEEE_754 invalid IEEE_754 invalid IEEE_754 invalid	 $-2^{31}$ $-2^{31}$ $-2^{63}$ $-2^{63}$

## E.1.4 Special Inexact Exceptions

UltraSPARC T1 Follows SPARC V9 5.1.7.10.5 (IEEE\_754 Section 7.5) and sets FSR\_inexact whenever the rounded result of an operation differs from the infinitely precise unrounded result.

Additionally, there are a few special cases to be aware of:

**TABLE E-9** Fp <-> Int Conversions With Inexact Exceptions

Cases	Possible Exception	Result (in addition to accrued exception) if tem is cleared
SPARC V9 A.14: Fp convert to 32b integer when source operand lies between $-(2^{31}-1)$ and $2^{31}$ , but is not exactly an integer FsTOi, FdTOi	IEEE_754 7.5	
	IEEE_754 inexact	An integer number
SPARC V9 A.14: Fp convert to 64b integer when source operand lies between $-(2^{63}-1)$ and $2^{63}$ , but is not exactly an integer FsTOx, FdTOx	IEEE_754 7.5	
	IEEE_754 inexact	An integer number
SPARC V9 A.15: Convert integer to fp format when 32b integer source operand magnitude is not exactly representable in single precision (23b mantissa). Note, even if the operand is $> 2^{24}-1$ , if enough of its trailing bits are zeros, it may still be exactly representable. FiTOs	IEEE_754 7.5	
	IEEE_754 inexact	A SP number
SPARC V9 A.15: Convert integer to fp format when 64b integer source operand magnitude is not exactly representable in single precision (23b mantissa). Note, even if the operand is $> 2^{24}-1$ , if enough of its trailing bits are zeros, it may still be exactly representable. FxTOs	IEEE_754 7.5	
	IEEE_754 inexact	A SP number
SPARC V9 A.15: Convert integer to fp format when 64b integer source operand magnitude is not exactly representable in double precision (52b mantissa). Note, even if the operand is $> 2^{53}-1$ , if enough of its trailing bits are zeros, it may still be exactly representable. FxTOD	IEEE_754 7.5	
	IEEE_754 inexact	A DP number

## E.2 Subnormal Handling

The UltraSPARC T1 FPU provides full hardware support for subnormal operands and results. Unlike UltraSPARC I/II and UltraSPARC III, UltraSPARC T1 will never generate an unfinished\_FPop trap type.

# Caches and Cache Coherency

## Chapter Revision History

Date	By	Comment
28 Aug 03	Bill Bryg	Started outline of appendix.
22 Sep 03	Bill Croxton	Copied and reformatted Chapter 8, Cache and Memory Interactions, from USLi User Manual.
25 Nov 03	J. Laudon	Initial changes for UltraSPARC T1
4 Jun 04	J. Laudon	More information on unit of coherence.
2 Aug 04	J. Laudon	Add cache index information.
20 Dec 04	J. Laudon	Better documentation of how to flush L2.
11 Feb 05	M. L. Nohr	Converted to UltraSPARC Architecture format

This appendix describes various interactions between the caches and memory, and the management processes that an operating system must perform to maintain data integrity in these cases. In particular, it discusses the following subjects:

- Invalidation of one or more cache entries – when and how to do it
- Differences between cacheable and noncacheable accesses
- Ordering and synchronization of memory accesses
- Accesses to addresses that cause side effects (I/O accesses)
- Nonfaulting loads
- Cache sizes, associativity, replacement policy, etc.

---

## F.1 Cache Flushing

Data in the level-1 (read-only or write-through) caches can be flushed by invalidating the entry in the cache. Modified data in the level-2 (writeback) cache must be written back to memory when flushed.

Cache flushing is required in the following cases:

- **I-cache:** Flush is needed before executing code that is modified by a local store instruction other than block commit store, see Section 3.1.1.1, “Instruction Cache (I-cache).” This is done with the FLUSH instruction or by using ASI accesses. When ASI accesses are used, software must ensure that the flush is done on the same virtual core as the stores that modified the code space.
- **D-cache:** Flush is needed when a physical page is changed from (virtually) cacheable to (virtually) noncacheable. This is done with a displacement flush (*Displacement Flushing* on page 116).
- **L2 cache:** Flush is needed for stable storage. Examples of stable storage include battery-backed memory and transaction logs. This is done with a displacement flush (see *Displacement Flushing* on page 116). Flushing the L2 cache flushes the corresponding blocks from the I- and D-caches because UltraSPARC T1 maintains inclusion between the L2 and L1 caches.

## F.1.1 Displacement Flushing

Cache flushing can be accomplished by a displacement flush. This is done by placing the cache in direct-map mode, and reading a range of read-only addresses that map to the corresponding cache line being flushed, forcing out modified entries in the local cache. Care must be taken to ensure that the range of read-only addresses is mapped in the MMU before starting a displacement flush; otherwise, the TLB miss handler may put new data into the caches. In addition, the range of addresses used to force lines out of the cache must not be present in the cache when starting the displacement flush (if any of the displacing lines are present before starting the displacement flush, fetching the already present line will *not* cause the proper way in the direct-mapped mode L2 to be loaded, instead the already present line will stay at its current location in the cache.)

**Note** | Diagnostic ASI accesses to the L2 cache can be used to invalidate a line, but they are generally not an alternative to displacement flushing. Modified data in the L2 cache will not be written back to memory using these ASI accesses.

## F.1.2 Memory Accesses and Cacheability

**Note** | Atomic load-store instructions are treated as both a load and a store; they can be performed only in cacheable address spaces.

## F.1.3 Coherence Domains

Two types of memory operations are supported in UltraSPARC T1: cacheable and noncacheable accesses, as indicated by the page translation. Cacheable accesses are inside the coherence domain; noncacheable accesses are outside the coherence domain.

SPARC V9 does not specify memory ordering between cacheable and noncacheable accesses. In TSO mode, UltraSPARC T1 maintains TSO ordering, regardless of the cacheability of the accesses. For SPARC V9 compatibility while in PSO or RMO mode, a MEMBAR #Lookaside should be used between a store and a subsequent load to the same noncacheable address. See the *SPARC Architecture Manual, Version 9* for more information about the SPARC V9 memory models.

On UltraSPARC T1, a MEMBAR #Lookaside executes more efficiently than a MEMBAR #StoreLoad.

### F.1.3.1 Cacheable Accesses

Accesses that fall within the coherence domain are called cacheable accesses. They are implemented in UltraSPARC T1 with the following properties:

- Data resides in real memory locations.
- They observe supported cache coherence protocol.
- The unit of coherence is 64 bytes at the system level (coherence between the virtual processors and I/O), enforced by the L2 cache.
- The unit of coherence for the primary caches (coherence between multiple virtual processors) is the primary cache line size (16 bytes for the data cache, 32 bytes for the instruction cache), enforced by the L2 cache directories.

### F.1.3.2 Noncacheable and Side-Effect Accesses

Accesses that are outside the coherence domain are called noncacheable accesses. Accesses of some of these memory (or memory mapped) locations may result in side effects. Noncacheable accesses are implemented in UltraSPARC T1 with the following properties:

- Data may or may not reside in real memory locations.
- Accesses may result in program-visible side effects; for example, memory-mapped I/O control registers in a UART may change state when read.
- Accesses may not observe supported cache coherence protocol.
- The smallest unit in each transaction is a single byte.

Noncacheable accesses with the `e` bit set (that is, those having side-effects) are all strongly ordered with respect to other noncacheable accesses with the `e` bit set. Speculative loads with the `e` bit set cause a *data\_access\_exception* trap.

**Note** | The side-effect attribute does not imply noncacheability.

### F.1.3.3 Global Visibility and Memory Ordering

To ensure the correct ordering between the cacheable and noncacheable domains, explicit memory synchronization is needed in the form of MEMBARs or atomic instructions. CODE EXAMPLE F-1 illustrates the issues involved in mixing cacheable and noncacheable accesses.

#### CODE EXAMPLE F-1 Memory Ordering and MEMBAR Examples

```
Assume that all accesses go to non-side-effect memory locations.
Process A:
While (1)
{

    Store D1:data produced
1 MEMBAR #StoreStore (needed in PSO, RMO)
    Store F1:set flag
    While F1 is set (spin on flag)
    Load F1
2 MEMBAR #LoadLoad | #LoadStore (needed in RMO)

    Load D2
}

Process B:
While (1)
{

    While F1 is cleared (spin on flag)

    Load F1
2 MEMBAR #LoadLoad | #LoadStore (needed in RMO)

    Load D1

    Store D2
1 MEMBAR #StoreStore (needed in PSO, RMO)

    Store F1:clear flag
}
```

**Note** | A MEMBAR #MemIssue or MEMBAR #Sync is needed if ordering of cacheable accesses following noncacheable accesses must be maintained in PSO or RMO.

In TSO mode, loads and stores (except block stores) cannot pass earlier loads, and stores cannot pass earlier stores; therefore, no MEMBAR is needed.

In PSO mode, loads are completed in program order, but stores are allowed to pass earlier stores; therefore, only the MEMBAR at #1 is needed between updating data and the flag.

In RMO mode, there is no implicit ordering between memory accesses; therefore, the MEMBARs at both #1 and #2 are needed.

## F.1.4 Memory Synchronization: MEMBAR and FLUSH

The MEMBAR (STBAR in SPARC V8) and FLUSH instructions are provide for explicit control of memory ordering in program execution. MEMBAR has several variations; their implementations in UltraSPARC T1 are described below.

- **MEMBAR #LoadLoad** — Forces all loads after the MEMBAR to wait until all loads before the MEMBAR have reached global visibility.
- **MEMBAR #StoreLoad** — Forces all loads after the MEMBAR to wait until all stores before the MEMBAR have reached global visibility.
- **MEMBAR #LoadStore** — Forces all stores after the MEMBAR to wait until all loads before the MEMBAR have reached global visibility.
- **MEMBAR #StoreStore** and **STBAR** — Forces all stores after the MEMBAR to wait until all stores before the MEMBAR have reached global visibility.

**Notes** | (1) STBAR has the same semantics as MEMBAR #StoreStore; it is included for SPARC V8 compatibility.

(2) The above four MEMBARs do not guarantee ordering between cacheable accesses after noncacheable accesses.

- **MEMBAR #Lookaside** — SPARC V9 provides this variation for implementations having virtually tagged store buffers that do not contain information for snooping.

**Note** | For SPARC V9 compatibility, this variation should be used before issuing a load to an address space that cannot be snooped.

- **MEMBAR #MemIssue** — Forces all outstanding memory accesses to be *completed* before any memory access instruction after the MEMBAR is issued. It must be used to guarantee ordering of cacheable accesses following noncacheable

accesses. For example, I/O accesses must be followed by a MEMBAR #MemIssue before subsequent cacheable stores; this ensures that the I/O accesses reach global visibility before the cacheable stores after the MEMBAR.

MEMBAR #MemIssue is different from the combination of MEMBAR #LoadLoad | #LoadStore | #StoreLoad | #StoreStore. MEMBAR #MemIssue orders cacheable and noncacheable domains; it prevents memory accesses after it from issuing until it completes.

- MEMBAR #Sync (Issue Barrier) — Forces all outstanding instructions and all deferred errors to be completed before any instructions after the MEMBAR are issued.

**Note** | MEMBAR #Sync is a costly instruction; unnecessary usage may result in substantial performance degradation.

See the references to “Memory Barrier,” “The MEMBAR Instruction,” and “Programming With the Memory Models,” in *The SPARC Architecture Manual, Version 9* for more information.

### F.1.4.1 Self-Modifying Code (FLUSH)

The SPARC V9 instruction set architecture does not guarantee consistency between code and data spaces. A problem arises when code space is dynamically modified by a program writing to memory locations containing instructions. LISP programs and dynamic linking require this behavior. SPARC V9 provides the FLUSH instruction to synchronize instruction and data memory after code space has been modified.

In UltraSPARC T1, a FLUSH behaves like a store instruction for the purpose of memory ordering. In addition, all instruction fetch (or prefetch) buffers are invalidated. The issue of the FLUSH instruction is delayed until previous (cacheable) stores are completed. Instruction fetch (or prefetch) resumes at the instruction immediately after the FLUSH.

## F.1.5 Atomic Operations

SPARC V9 provides three atomic instructions to support mutual exclusion. These instructions behave like both a load and a store but the operations are carried out indivisibly. Atomic instructions may be used only in the cacheable domain.

An atomic access with a restricted ASI in nonprivileged mode (PSTATE.priv = 0) causes a *privileged\_action* trap. An atomic access with a noncacheable address causes a *data\_access\_exception* trap. An atomic access with an unsupported ASI causes a *data\_access\_exception* trap. TABLE F-1 lists the ASIs that support atomic accesses.

**TABLE F-1** ASIs that Support SWAP, LDSTUB, and CAS

ASI Name	Access
ASI_NUCLEUS{ <small>_LITTLE</small> }	Restricted
ASI_AS_IF_USER_PRIMARY{ <small>_LITTLE</small> }	Restricted
ASI_AS_IF_USER_SECONDARY{ <small>_LITTLE</small> }	Restricted
ASI_PRIMARY{ <small>_LITTLE</small> }	Unrestricted
ASI_SECONDARY{ <small>_LITTLE</small> }	Unrestricted
ASI_REAL{ <small>_LITTLE</small> }	Unrestricted

**Note** | Atomic accesses with nonfaulting ASIs are not allowed, because these ASIs have the load-only attribute.

### F.1.5.1 SWAP Instruction

SWAP atomically exchanges the lower 32 bits in an integer register with a word in memory. This instruction is issued only after store buffers are empty. Subsequent loads interlock on earlier SWAPs. A cache miss allocates the corresponding line.

**Note** | If a page is marked as virtually noncacheable but physically cacheable, allocation is done to the L2 cache only.

### F.1.5.2 LDSTUB Instruction

LDSTUB behaves like SWAP, except that it loads a byte from memory into an integer register and atomically writes all ones ( $FF_{16}$ ) into the addressed byte.

### F.1.5.3 Compare and Swap (CASX) Instruction

Compare-and-swap combines a load, compare, and store into a single atomic instruction. It compares the value in an integer register to a value in memory; if they are equal, the value in memory is swapped with the contents of a second integer register. All of these operations are carried out atomically; in other words, no other memory operation may be applied to the addressed memory location until the entire compare-and-swap sequence is completed.

## F.1.6 Nonfaulting Load

A nonfaulting load behaves like a normal load, except as follows:

- It does not allow side-effect access. An access with the **e** bit set causes a *data\_access\_exception* trap.
- It can be applied to a page with the **nfo** bit set; other types of accesses will cause a *data\_access\_exception* trap.

# Glossary

---

This chapter defines concepts and terminology common to all implementations of SPARC V9. It also includes terms that are unique to the UltraSPARC T1 implementation.

- ALU** Arithmetic Logical Unit
- address space identifier (ASI)** An 8-bit value that identifies an address space. For each instruction or data access, the integer unit appends an ASI to the address. *See also* **implicit ASI**.
- application program** A program executed with the 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* processor state (for example, as stored in a memory-image dump).
- architectural state** Software-visible registers and memory (including caches).
- ARF** Architectural Register File.
- ASI** Address Space Identifier.
- ASR** Ancillary State Register.
- big-endian** An addressing convention. Within a multiple-byte integer, the byte with the smallest address is the most significant; a byte's significance decreases as its address increases.
- BLD** (Obsolete) abbreviation for Block Load instruction; replaced by LDBLOCKF.
- blocking ASI** An ASI that will access its ASI register/array location once all older instructions in that strand have retired, there are no instructions in the other strand which can issue, and the store queue, TSW, and LMB are all empty. Additionally, the snoop pipeline is stalled before accessing the ASI register/array location. *See* **nonblocking ASI**.
- branch outcome** Refers to whether or not a branch instruction will alter the flow of execution from the sequential path. A taken branch outcome results in execution proceeding with the instruction at the branch target; a not-taken branch outcome results in execution proceeding with the instruction along the sequential path after the branch.

<b>branch resolution</b>	A branch is said to be resolved when the result (that is, the branch outcome and branch target address) has been computed and is known for certain. Branch resolution can take place late in the pipeline.
<b>branch target address</b>	The address of the instruction to be executed if the branch is taken.
<b>BST</b>	(Obsolete) abbreviation for Block Store instruction; replaced by STBLOCKF.
<b>bypass ASI</b>	An ASI that refers to memory and for which the MMU does not perform virtual-to-real address translation (that is, memory is accessed using a direct real address).
<b>byte</b>	Eight consecutive bits of data.
<b>CCR</b>	Condition Codes register
<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>CMP</b>	Chip multiprocessor. A single chip processor that contains more than one virtual processor. <i>See also</i> <b>processor</b> and <b>virtual processor</b> .
<b>coherence</b>	A set of protocols guaranteeing that all memory accesses are globally visible to all caches in a shared-memory system.
<b>commit</b>	An instruction commits when it modifies architectural state.
<b>completed</b>	A memory transaction is said to be completed 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>complex instruction</b>	An instruction that requires the creation of secondary “helper” instructions for normal operation, excluding trap conditions such as spill/fill traps (which use helpers).
<b>consistency</b>	<i>See coherence.</i>
<b>context</b>	A set of translations that supports a particular address space. <i>See also</i> <b>Memory Management Unit (MMU)</b> .
<b>CWP</b>	Current Window Pointer
<b>CPI</b>	Cycles per instruction. The number of clock cycles it takes to execute an instruction.
<b>CPU</b>	Central Processing Unit. A synonym for <b>virtual processor</b> .
<b>cross-call</b>	An interprocessor call in a multiprocessor system.
<b>CSR</b>	Control Status Register.
<b>CTI</b>	Control transfer instruction

<b>current window</b>	The block of 24 R registers that is currently in use. The Current Window Pointer (CWP) register points to the current window.
<b>DCTI</b>	Delayed control transfer instruction.
<b>DDR</b>	Double Data Rate.
<b>deprecated</b>	The term applied to an architectural feature (such as an instruction or register) for which a SPARC V9 implementation provides support only 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 SPARC V9 software and may not be supported in future versions of the architecture.
<b>DFT</b>	Designed for test.
<b>doublet</b>	Two bytes (16 bits) of data.
<b>doubleword</b>	An aligned octlet. <b>Note:</b> The definition of this term is architecture dependent and may differ from that used in other processor architectures.
<b>DTLB</b>	Data Cache Translation lookaside buffer.
<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 set bits.
<b>exception</b>	A condition that makes it impossible for the processor to continue executing the current instruction stream without software intervention. <i>See also</i> <b>trap</b> .
<b>extended word</b>	An aligned octlet, 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>EXU</b>	Execution Unit
<b>F register</b>	A floating-point register. SPARC V9 includes single-, double-, and quad-precision F registers.
<b><i>fccn</i></b>	One of the floating-point condition code fields <i>fcc0</i> , <i>fcc1</i> , <i>fcc2</i> , or <i>fcc3</i> .
<b>FGU</b>	Floating-point and Graphics Unit.
<b>floating-point exception</b>	An exception that occurs during the execution of an FPop instruction as defined by the <i>Fpop1</i> , <i>Fpop2</i> , <i>IMPDEP1</i> , and <i>IMPDEP2</i> opcodes. 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>floating-point IEEE-754 exception</b>	A floating-point exception, as specified by IEEE Std 754-1985. Listed within this specification as <i>IEEE_754_exception</i> .

<b>floating-point operate (FPop) instructions</b>	Instructions that perform floating-point and graphics calculations, as defined by the FPop1, FPop2, and IMPDEP1 opcodes. FPop instructions do not include FBfcc instructions, loads and stores between memory and the floating-point unit, or instructions defined by the IMPDEP2 opcodes.
<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>FP</b>	Floating Point
<b>FPRS</b>	Floating Point Register State (register).
<b>FRF</b>	Floating-point register file.
<b>FSR</b>	Floating-Point Status register.
<b>GL</b>	Global-Level register.
<b>GSR</b>	General Status register.
<b>halfword</b>	An aligned doublet. <b>Note:</b> The definition of this term is architecture dependent and may differ from that used in other processor architectures.
<b>helper</b>	An instruction generated by the IRU in response to a complex instruction. Helper instructions are not visible to software. Refer to <i>Instruction Latencies</i> on page 76 for a complete list of all complex instructions and their helper sequences.
<b>hexlet</b>	Sixteen bytes (128 bits) of data.
<b>hyperprivileged</b>	An adjective that describes the state of the processor when it is executing in <i>hyperprivileged mode</i>
<b>IFU</b>	Instruction Fetch Unit.
<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 architecture that can legitimately vary among implementations. In many cases, the permitted range of variation is specified in the SPARC V9 standard. When a range is specified, compliant implementations must not deviate from that range.
<b>implicit ASI</b>	The address space identifier that is supplied by the hardware on all instruction accesses and on data accesses that do not contain an explicit ASI or a reference to the contents of the ASI register.

<b>informative appendix</b>	An appendix containing information that is useful but not required to create an implementation that conforms to the SPARC V9 specification. <i>See also normative appendix.</i>
<b>initiated</b>	<i>Synonym: issued.</i>
<b>instruction field</b>	A bit field within an instruction word.
<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. The bulk of the ISA implemented by UltraSPARC T1 is defined in the <i>UltraSPARC Architecture 2005</i> ; a few extensions are described in this document.
<b>integer unit (IU)</b>	A processing unit that performs integer and control-flow operations and contains general-purpose integer registers and processor state registers, as defined by this specification.
<b>interrupt request</b>	A request for service presented to the processor by an external device.
<b>IRF</b>	Integer register file.
<b>IRU</b>	
<b>ISA</b>	Instruction set architecture.
<b>issue</b>	Used to describe the act of conveying an instruction from the instruction fetch unit for execution on the pipeline.
<b>L2C</b>	Level 2 Cache.
<b>leaf procedure</b>	A procedure that is a leaf in the program's call graph; that is, one that does not call (by using CALL or JMPL) any other procedures.
<b>little-endian</b>	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.
<b>load</b>	An instruction that reads (but does not write) memory or reads (but does not write) location(s) in an alternate address space. <i>Load</i> includes loads into integer or floating-point registers, block loads, Load Quadword Atomic, and alternate address space variants of those instructions. <i>See also load-store</i> and <i>store</i> , the definitions of which are mutually exclusive with <i>load</i> .
<b>load-store</b>	An instruction that explicitly both reads and writes memory or explicitly reads and writes location(s) in an alternate address space. <i>Load-store</i> includes instructions such as CASA, CASXA, LDSTUB, LDSTUBA and the deprecated SWAP and SWAPA instructions. <i>See also load</i> and <i>store</i> , the definitions of which are mutually exclusive with <i>load-store</i> .

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

**Memory Management Unit (MMU)**

The address translation hardware that translates 64-bit virtual address into real addresses. *See also* **context**, **physical address**, and **virtual address**.

**must** *Synonym:* **shall**.

**next program counter (NPC)**

A register that contains the address of the instruction to be executed next if a trap does not occur.

**NFO** Nonfault access only.

**nonblocking ASI**

An ASI that will access its ASI register/array location once all older instructions in that strand have retired and there are no instructions in the other strand which can issue. *See* **blocking ASI**.

**nonfaulting load**

A load operation that, in the absence of faults or in the presence of a recoverable fault, completes correctly, and in the presence of a nonrecoverable fault returns (with the assistance of system software) a known data value (nominally zero). *See* **speculative load**.

**nonprivileged**

An adjective that describes:  
(1) the state of the processor when `PSTATE.priv = 0`, that is, nonprivileged mode;  
(2) processor state information that is accessible to software while the processor is in either privileged mode or nonprivileged mode; for example, nonprivileged registers, nonprivileged ASRs, or, in general, nonprivileged state;  
(3) an instruction that can be executed when the processor is in either privileged mode or nonprivileged mode.

**nonprivileged mode**

The mode in which a processor is operating when `PSTATE.priv = 0`. *See also* **privileged**.

**normative appendix**

An appendix containing specifications that must be met by an implementation conforming to the SPARC V9 specification. *See also* **informative appendix**.

**nontranslating ASI**

An ASI that does not refer to memory (for example, refers to control/status register(s)) and for which the MMU does not perform address translation.

**NPC** Next program counter.

**npt** Nonprivileged trap.

***N\_REG\_WINDOWS***

The number of register windows present in a particular implementation.

**octlet**

Eight bytes (64 bits) of data. Not to be confused with “octet,” which has been commonly used to describe eight bits of data. In this document, the term *byte*, 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 contains an odd number of set bits.
<b>older instruction</b>	Refers to the relative fetch order of instructions. Instruction <i>i</i> is older than instruction <i>j</i> if instruction <i>i</i> was fetched before instruction <i>j</i> . Data dependencies flow from older instructions to younger instructions, and an instruction can only be dependent upon older instructions. <i>See</i> <b>younger instruction</b> .
<b>one-hot</b>	An <i>n</i> -bit binary signal is one-hot if, and only if, <i>n</i> – 1 of the bits are each 0 and a single bit is 1.
<b>opcode</b>	A bit pattern that identifies a particular instruction.
<b>optional</b>	A feature not required for compliance to an architecture specification (such as UltraSPARC Architecture 2005 or SPARC V9).
<b>PC</b>	Program counter.
<b>PCR</b>	Performance Control Register.
<b>PIC</b>	Performance Instrumentation Counter.
<b>PIL</b>	Processor Interrupt Level.
<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</i> <b>side effect</b> .
<b>privileged</b>	An adjective that describes (1) the state of the processor when it is executing in <i>privileged mode</i> ( PSTATE.priv = 1); (2) processor state that is only accessible to software while the processor is in <i>privileged mode</i> ; for example, privileged registers, privileged ASRs, privileged ASIs, 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. <i>See also</i> <b>nonprivileged</b> .

**processor** 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 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, like an UltraSPARC I processor, that can only execute one thread at a time would be a processor with a single physical core that is single-stranded. A processor that follows the academic model of simultaneous multithreading (SMT) would be a processor with a single physical core, where that physical core supports multiple strands in order to execute multiple threads at the same time (multi-stranded physical core). A processor that follows the academic model of a CMP would be a processor with multiple physical cores, each only supporting a single strand. One can also have multiple physical cores where each physical core is multi-stranded. UltraSPARC T1 is an example of the latter, where each UltraSPARC T1 processor contains eight physical cores, each of which contains four strands.

**program counter**

**(PC)** A register that contains the address of the instruction currently being executed by the IU.

**PSO** Partial store order.

**quadlet** Four bytes (32 bits) of data.

**quadword** Aligned hexlet. **Note:** The definition of this term is architecture dependent and may be different from that used in other processor architectures.

**RA** Real address.

**RAS** Return Address Stack;  
also Reliability, Availability and Serviceability.

**RAW** Read After Write

**R register** An integer register. Also called a general-purpose register or working register.

**rd** Rounding direction.

**RDPR** Read Privileged Register instruction

**real address** An address used by privileged mode code to describe the underlying physical memory. Real address are usually translated by a combination of hyperprivileged hardware and software to physical addresses which can be used to access real physical memory or I/O device space.

- reserved** Describing an instruction field, certain bit combinations within an instruction field, or a register field that is reserved for definition by future versions of the architecture.
- Reserved instruction fields* shall read as 0, unless the implementation supports extended instructions within the field. The behavior of SPARC V9 processors when they encounter nonzero values in reserved instruction fields is undefined.
- Reserved bit combinations within instruction fields* are defined in Chapter 5, *Instruction Definitions*. In all cases, SPARC V9 processors shall decode and trap on these reserved combinations.
- Reserved register fields* should always be written by software with values of those fields previously read from that register or with zeroes; they should read as zero in hardware. Software intended to run on future versions of SPARC V9 should not assume that these fields will read as 0 or any other particular value. Throughout this specification, figures and tables illustrating registers and instruction encodings indicate reserved fields and combinations with an em dash (—).
- restricted** Describing an address space identifier (ASI) that may be accessed only while the processor is operating in privileged mode.
- rs1, rs2, rd** The integer or floating-point register operands of an instruction. *rs1* and *rs2* are the source registers; *rd* is the destination register.
- RMO** Relaxed memory order.
- RTO** Read to own (cache line state).
- RTS** Read to share (cache line state).
- shall** A keyword indicating a mandatory requirement. Designers shall implement all such mandatory requirements to ensure interoperability with other SPARC V9-compliant products. *Synonym: must.*
- should** A keyword indicating flexibility of choice with a strongly preferred implementation. *Synonym: it is recommended.*
- SIAM** Set interval arithmetic mode instruction.
- side effect** The result of a memory location having additional actions beyond the reading or writing of data. A side effect can occur when a memory operation on that location is allowed to succeed. Locations with side effects include those that, when accessed, change state or cause external events to occur. For example, some I/O devices contain registers that clear on read; others have registers that initiate operations when read. *See also prefetchable.*
- SIMD** Single instruction stream, multiple data stream.

- store** An instruction that writes (but does not explicitly read) memory or writes (but does not explicitly read) location(s) in an alternate address space. *Store* includes stores from either integer or floating-point registers, block stores, partial store, and alternate address space variants of those instructions. *See also load and load-store*, the definitions of which are mutually exclusive with *store*.
- strand** A term for thread-specific hardware support that identifies the hardware state used to hold a software thread in order to execute it. Strand is specifically the software visible architected state (PC, 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. “Strand” replaces the ambiguous term “hardware thread”. The number of strands in a processor defines the number of threads that the operating system can schedule on that processor at any given time. *See also thread, and virtual processor.*
- strand identifier (SID)** An  $n$ -bit value, in a processor implementing  $2^n$  strands, that uniquely identifies each strand. The strand identifier in UltraSPARC T1 is five bits wide.
- superscalar** An implementation that allows several instructions to be issued, executed, and committed in one clock cycle.
- supervisor software** Software that executes when the processor is in privileged mode.
- thread** An executing process or lightweight process (LWP). Historically, the term thread is overused and ambiguous. Software and hardware have historically used it differently. From software’s (operating system) perspective, the term thread refers to an entity that can be run on hardware, it is something that is scheduled and may or may not be actively running on hardware at any given time, and may migrate around the hardware of a system. From hardware’s perspective, the term multithreaded processor refers to a processor that run multiple software threads simultaneously. To avoid confusion the term thread is used exclusively in the manner in which it is used by software and, specifically, the operating system. A thread can be viewed in a practical sense as a Solaris™ process or lightweight process (LWP). *See also strand, and virtual processor.*
- TICK** Hardware clock—TICK counter register.
- TL** Trap Level
- TPC** Trap-saved PC.
- trap** The action taken by the processor when it changes the instruction flow in response to the presence of an exception, a Tcc instruction, or an interrupt. The action is a vectored transfer of control to privileged or hyperprivileged software through a table. *See also exception.*
- TSB** Translation storage buffer. A table of the address translations that is maintained by software in system memory and that serves as a cache of 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>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>	<p>An aspect of the architecture that has deliberately been left unspecified. Software should have no expectation of, or 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.</p> <p>Notwithstanding any of the above, undefined aspects of the architecture shall not cause security holes (such as allowing user software to access privileged state), put the processor into supervisor mode, or put the processor into an unrecoverable state.</p>
<b>unimplemented</b>	An architectural feature that is not directly executed in hardware because it is optional or is emulated in software.
<b>unpredictable</b>	<i>Synonym: undefined.</i>
<b>unrestricted</b>	Describing an address space identifier (ASI) that can be used regardless of the processor mode; that is, regardless of the value of PSTATE.priv.
<b>user application program</b>	<i>Synonym: application program.</i>
<b>VA</b>	<i>Virtual address.</i>
<b>virtual address</b>	An address produced by a processor that maps all systemwide, program-visible memory. Virtual addresses usually are translated by a combination of hardware and software to real addresses.
<b>virtual processor</b>	The term <b>virtual processor</b> , is used to identify each strand in a processor. Each virtual processor corresponds to a specific strand on a specific physical core where there may be multiple physical cores each with multiple strands. In most respects, a virtual processor appears to the system, and to the operating system software, as a processing unit equivalent to a traditional single-stranded microprocessor (as in UltraSPARC I). Each virtual processor has its own interrupt ID and the operating system can schedule independent threads on each virtual processor. How multiple virtual processors are achieved within a processor is an implementation issue, and as much as possible the software interface is independent of how multiple virtual processors are implemented. The term virtual processor is used in place of strand because of the common association of the term strand with multi-stranded physical cores. <i>See also strand, and thread.</i>

**VIS™** Visual instruction set.

**word** An aligned quadlet. Note: The definition of this term is architecture dependent and may differ from that used in other processor architectures.

**younger instruction** *See older instruction.*

**writeback** The process of writing a dirty cache line back to memory before it is refilled.

**WRPR** Write Privileged Register.

# Index

---

## A

Accumulated Exception (aexc) field of FSR register, 65  
Address Mask (am), 60  
Address Mask (am)  
  field of PSTATE register, 40, 41, 57  
address space identifier (ASI)  
  bypass, **124**  
  definition, 123  
  nontranslating, **128**  
application program, **123**  
ASI\_PRIMARY\_NO\_FAULT, 57  
ASI\_PRIMARY\_NO\_FAULT\_LITTLE, 57  
ASI\_SECONDARY\_NO\_FAULT, 57  
ASI\_SECONDARY\_NO\_FAULT\_LITTLE, 57  
atomic quad load instructions (deprecated), 27

## B

BA instruction, 99  
BCC instruction, 99  
BCS instruction, 99  
BE instruction, 99  
BG instruction, 99  
BGE instruction, 99  
BGU instruction, 99  
Bicc instructions, 99  
BL instruction, 99  
BLE instruction, 99  
BLEU instruction, 99  
block  
  copy, inner loop pseudo-code, **20**  
  load instructions, 21

  memory operations, 70  
block-transfer ASIs, 22  
BN instruction, 99  
BNE instruction, 99  
BNEG instruction, 99  
BPA instruction, 100  
BPCC instruction, 100  
BPCS instruction, 100  
bpe instruction, 100  
BPG instruction, 100  
BPGE instruction, 100  
BPGU instruction, 100  
BPL instruction, 100  
BPLE instruction, 100  
BPLEU instruction, 100  
BPN instruction, 100  
BPNE instruction, 100  
BPNEG instruction, 100  
BPOS instruction, 99  
BPPOS instruction, 100  
BPr instructions, 100  
BPVC instruction, 100  
BPVS instruction, 100  
BVC instruction, 99  
BVS instruction, 99  
bypass ASI, **124**

## C

caching  
  TSB, 55  
canrestore Register, 61  
cansave Register, 61  
clean window, 62

*clean\_window* trap, 62  
cleanwin Register, 61  
CLEANWIN register, 62  
compatibility with SPARC V9  
    terminology and concepts, 123  
conventions  
    font, ix  
    notational, x  
cross call, 71  
Current Exception (*cexc*) field of FSR register, 65  
current window pointer (CWP) register  
    definition, 125  
cwp Register, 61

## D

D superscript on instruction name, 13  
data watchpoint  
    virtual address, 57  
*data\_access\_exception* trap, 24, 27, 28, 41, 57, 59, 67  
deferred  
    trap, 60  
Diagnostic (*diag*) field of TTE, 54  
Dirty Lower (*dl*) field of FPRS register, 64  
Dirty Upper (*du*) field of FPRS register, 64  
D-MMU, 57  
doublet, 125  
doubleword  
    definition, 125

## E

enhanced security environment, 60  
exceptions  
    *fp\_exception\_other*, 12  
    *illegal\_instruction*, 12  
extended  
    instructions, 71

## F

FABSd instruction, 98, 99  
FABSq instruction, 98, 99  
FBA instruction, 99  
FBE instruction, 99  
FBfcc instructions, 99  
FBG instruction, 99  
FBGE instruction, 99

FBL instruction, 99  
FBLE instruction, 99  
FBLG instruction, 99  
FBN instruction, 99  
FBNE instruction, 99  
FBO instruction, 99  
FBPA instruction, 100  
FBPE instruction, 100  
FBPfcc instructions, 99  
FBPG instruction, 100  
FBPGE instruction, 100  
FBPL instruction, 100  
FBPLE instruction, 100  
FBPLG instruction, 100  
FBPN instruction, 100  
FBPNE instruction, 100  
FBPO instruction, 100  
FBPU instruction, 100  
FBPUE instruction, 100  
FBPUG instruction, 100  
FBPUGE instruction, 100  
FBPUL instruction, 100  
FBPULE instruction, 100  
FBU instruction, 99  
FBUE instruction, 99  
FBUG instruction, 99  
FBUGE instruction, 99  
FBUL instruction, 99  
FBULE instruction, 99  
FCMPd instruction, 99  
FCMPEd instruction, 99  
FCMPEq instruction, 99  
FCMPEs instruction, 99  
FCMPq instruction, 99  
FCMPs instruction, 99  
FdTOx instruction, 98, 99  
floating point  
    deferred trap queue (*fq*), 66  
    exception handling, 63  
    trap type (*ftt*) field of FSR register, 65  
Floating Point Condition Code (*fcc*)  
    0 (*fcc0*) field of FSR register, 65, 66  
    1 (*fcc1*) field of FSR register, 65  
    2 (*fcc2*) field of FSR register, 65  
    3 (*fcc3*) field of FSR register, 65  
    field of FSR register in SPARC-V8, 66  
Floating Point Registers State (FPRS) Register, 64  
floating-point trap type (*ftt*) field of FSR register, 12  
floating-point trap types

- unimplemented\_FPop*, 12
- FLUSH instruction, 67
- FMOVcc instructions, 100
- FMOVccd instruction, 99
- FMOVccq instruction, 99
- FMOVccs instruction, 99
- FMOVd instruction, 98, 99
- FMOVq instruction, 98, 99
- FNEGd instruction, 98, 99
- FNEGq instruction, 98, 99
- fp\_exception\_ieee\_754* trap, 65, 66
- fp\_exception\_other* exception, 12
- fp\_exception\_other* trap, 59, 63, 65, 66
- fq*, see *floating-point deferred trap queue (fq)*
- FqTOx instruction, 98, 99
- FRF, 126
- FsTOx instruction, 98, 99
- FxTOd instruction, 98, 99
- FxTOq instruction, 98, 99
- FxTOs instruction, 98, 99

## H

- hardware
  - interrupts, 71
- hardware\_error floating-point trap type, 66

## I

- IEEE Std 754-1985, 65, 125
- IEEE support
  - inexact exceptions, 113
  - infinity arithmetic, 106
  - NaN arithmetic, 112
  - one infinity operand arithmetic, 107
  - two infinity operand arithmetic, 110
  - zero arithmetic, 111
- IEEE\_754\_exception* floating-point trap type, 66
- IEEE\_754\_exception* floating-point trap type, 125
- illegal\_instruction* exception, 12
- illegal\_instruction* trap, 40, 59, 66, 69, 71
- ILLTRAP instructions, 59
- implementation note, xii
- initiated, 127
- instruction fields
  - definition, 127
- instruction set architecture (ISA), 126, 127
- instruction\_access\_exception* trap, 40, 41, 57
- instructions

- reserved, 12
- integer
  - division, 62
  - multiplication, 62
  - register file, 61
- interrupt
  - packet, 71
  - request, 127
- invalid\_fp\_register* floating-point trap type, 66
- Invert Endianness
  - (ie) field of TTE, 54
- IRF, 127

## L

- LDDF\_mem\_address\_not\_aligned* trap, 69
- LDQF instruction, 69
- LDQFA instruction, 69
- LDTW instruction, 69
- little-endian
  - byte ordering, 127
- load instructions, 127
- load twin extended word instructions, 25
- load twin extended word instructions
  - (deprecated), 27
- load-store instructions
  - definition, 127
- Lock (l) field of TTE, 54

## M

- may (keyword), 128
- mem\_address\_not\_aligned* trap, 24, 27, 28, 40, 57
- MEMBAR
  - #LoadStore, 19
  - #StoreLoad, 19
  - #StoreStore, 19, 67
  - #Sync, 18, 19
- memory
  - model, 19
- MOVcc instructions, 100
- must (keyword), 128
- M-way set-associative TSB, 55

## N

- N\_REG\_WINDOWS*, 61
- nested traps
  - in SPARC-V9, 60

- No-Fault Only (nfo) field of TTE, 54
- nonfaulting load, 57
- nonfaulting loads
  - definition of, 128
- nonprivileged
  - mode, 123
- Non-Standard (ns) field of FSR register, 65
- nontranslating ASI, 128
- note
  - implementation, **xii**
  - programming, **xii**
- NPC register, 41

**O**

- opcode
  - definition, 129
- otherwin Register, 61
- out of range
  - virtual address, 39
  - virtual address, as target of JMWPL or RETURN, 40

**P**

- P superscript on instruction name, 13
- partial store
  - instruction, 70
- physical address (pa)
  - field of TTE, 54
- population count (POPC) instruction, 60
- power down mode, 71
- precise traps, 60
- PREFETCHA instruction, 68
- privileged
  - (priv) field of PSTATE register, 57
- privileged\_action* trap, 57
- programming note, **xii**
- pstate, 19
- PSTATE
  - priv field, 128, 129

**Q**

- quad-precision floating-point instructions, 63
- quadword
  - definition, 130

**R**

- reserved
  - fields in opcodes, 59
  - instructions, 12, 59
- Rounding Direction (rd) field of FSR register, 66

**S**

- SAVE instruction, 62
- secure environment, 60
- self-modifying code, 67
- sequence\_error floating-point trap type, 66
- shall (keyword), 131
- short floating point
  - load instruction, 70
  - store instruction, 70
- should (keyword), 131
- software
  - defined (soft) field of TTE, 54
  - defined (soft2) field of TTE, 54
  - Translation Table, 55
- SPARC
  - V9 compliance, 59
- SPARC V9
  - concepts and terminology, 123
- speculative load, 57
- STDF\_mem\_address\_not\_aligned* trap, 69
- store instructions, 132
- STQF instruction, 69
- STQFA instruction, 69
- STTW instruction, 69

**T**

- TA instruction, 99
- Tcc instruction, reserved fields, 59
- Tcc instructions, 99, 100
- TCS instruction, 99
- TE instruction, 99
- terminology for SPARC V9, definition of, 123
- TG instruction, 99
- TGE instruction, 99
- TGU instruction, 99
- tl instruction, 99
- tle instruction, 99
- TLEU instruction, 99
- TN instruction, 99
- TNE instruction, 99
- TNEG instruction, 99

TPOS instruction, 99  
Translation Table Entry *see* TTE  
trap

- stack, 60
- state registers, 60

Trap Enable Mask (*tem*) field of FSR register, 65, 65, 66  
TSB, 27

- caching, 55
- organization, 55
- Register, 55

TTE, 53  
TVC instruction, 99  
TVS instruction, 99

## U

UltraSPARC-I

- extended instructions, 71

unfinished\_FPop floating-point trap type, 66  
unimplemented

- instructions, 59

unimplemented\_FPop floating-point trap type, 63, 66  
*unimplemented\_FPop* floating-point trap type, 12

## V

VA Data Watchpoint register, 57  
*VA\_watchpoint* trap, 24, 28  
Version (*ver*) field of FSR register, 65  
virtual address

- space *illustrated*, 40

## W

*watchpoint* trap, 57  
*window\_fill* trap, 40  
Writable (*w*) field of TTE, 54

