

SPARC M7™ Supplement to the *Oracle SPARC Architecture 2015*

Draft D1.0, 30 Jun 2016

*Privilege Levels: Privileged
and Nonprivileged*

Distribution: Public

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

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

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

Contents

1	SPARC M7 Basics	9
1.1	Background	9
1.2	SPARC M7 Overview	10
1.3	SPARC M7 Components.....	11
1.3.1	SPARC Physical Core.....	11
1.3.1.1	Single-threaded and multi-threaded performance ...	12
1.3.2	L3 Cache.....	12
2	Data Formats	15
3	Registers	17
3.1	Floating-Point State Register (FSR)	17
3.2	Ancillary State Registers (ASRs)	17
3.2.1	Tick Register (TICK).....	18
3.2.2	Program Counter (PC).....	18
3.2.3	Floating-Point Registers State Register (FPRS).....	19
3.2.4	General Status Register (GSR)	19
3.2.5	Software Interrupt Register (SOFTINT).....	19
3.2.6	System Tick Register (STICK)	19
3.2.7	System Tick Compare Register (STICK_CMPR).....	20
3.2.8	Compatibility Feature Register (CFR)	21
3.2.9	Pause (PAUSE)	23
3.2.10	MWAIT	23
3.3	Privileged PR State Registers	24
3.3.1	Trap State Register (TSTATE)	24
3.3.2	Processor State Register (PSTATE).....	25
3.3.3	Trap Level Register (TL).....	25
3.3.4	Current Window Pointer (CWP) Register.....	25
3.3.5	Global Level Register (GL)	26
4	Instruction Format	27
5	Instruction Definitions	29
5.1	Instruction Set Summary	29
5.2	PREFETCH/PREFETCHA	35
5.3	WRPAUSE	36
5.4	WRMWAIT.....	37
5.5	Block Load and Store Instructions	38
5.6	Block Initializing Store ASIs	41
6	Traps	45
6.1	Trap Levels.....	45
6.2	Trap Behavior.....	45
6.3	Trap Masking.....	45

7	Interrupt Handling	47
7.1	Interrupt Flow	47
7.1.1	Sources	47
7.1.2	States	47
7.2	CPU Interrupt Registers	47
7.2.1	Interrupt Queue Registers	47
8	Memory Models	51
8.1	Supported Memory Models	51
8.1.1	TSO	52
8.1.2	RMO	52
9	Address Spaces and ASIs	53
9.1	Address Spaces	53
9.1.1	54-bit Virtual and Real Address Spaces	53
9.2	Alternate Address Spaces	54
9.2.1	ASI_REAL, ASI_REAL_LITTLE, ASI_REAL_IO, and ASI_REAL_IO_LITTLE (ASIs 1416, 1C16, 1516, 1D16)	60
9.2.2	ASI_SCRATCHPAD (ASI 2016, VA 016-1816, 3016-3816)	60
10	Performance Instrumentation	61
10.1	Introduction	61
10.2	SPARC Performance Control Registers	61
10.3	SPARC Performance Instrumentation Counter	62
11	Implementation Dependencies	65
11.1	SPARC V9 General Information	65
11.1.1	Level-2 Compliance (Impdep #1)	65
11.1.2	Unimplemented Opcodes, ASIs, and ILLTRAP	65
11.1.3	Trap Levels (Impdep #37, 38, 39, 40, 114, 115)	65
11.1.4	Trap Handling (Impdep #16, 32, 33, 35, 36, 44)	65
11.1.5	Secure Software	66
11.1.6	Address Masking (Impdep #125)	66
11.2	Integer Operations	66
11.2.1	Integer Register File and Window Control Registers (Impdep #2)	66
11.2.2	Clean Window Handling (Impdep #102)	66
11.2.3	Integer Multiply and Divide	66
11.2.4	MULScc	67
11.3	SPARC V9 Floating-Point Operations	67
11.3.1	Overflow, Underflow, and Inexact Traps (Impdep #3, 55)	67
11.3.2	Quad-Precision Floating-Point Operations (Impdep #3)	67
11.3.3	Floating-Point Upper and Lower Dirty Bits in FPRS Register	68
11.3.4	Floating-Point Status Register (FSR) (Impdep #13, 19, 22, 23, 24)	68
11.4	SPARC V9 Memory-Related Operations	69
11.4.1	Load/Store Alternate Address Space (Impdep #5, 29, 30)	69
11.4.2	Read/Write ASR (Impdep #6, 7, 8, 9, 47, 48)	69
11.4.3	MMU Implementation (Impdep #41)	70
11.4.4	FLUSH and Self-Modifying Code (Impdep #122)	70
11.4.5	PREFETCH{A} (Impdep #103, 117)	70
11.4.6	LDD/STD Handling (Impdep #107, 108)	70
11.4.7	FP mem_address_not_aligned (Impdep #109, 110, 111, 112)	70
11.4.8	Supported Memory Models (Impdep #113, 121)	70
11.4.9	Implicit ASI When TL > 0 (Impdep #124)	70
11.5	Non-SPARC V9 Extensions	71
11.5.1	Cache Subsystem	71
11.5.2	Block Memory Operations	71
11.5.3	Partial Stores	71

11.5.4	Short Floating-Point Loads and Stores.	71
11.5.5	Load Twin Extended Word	71
11.5.6	SPARC M7 Instruction Set Extensions (Impdep #106)	71
11.5.7	Performance Instrumentation	71
11.5.8	ASI_MONITOR_AS_IF_USER_PRIMARY, ASI_MONITOR_AS_IF_USER_SECONDARY, ASI_MONITOR_PRIMARY, ASI_MONITOR_SECONDARY	72
12	Cryptographic Extensions	73
12.1	CFR Register.	73
12.2	Cryptographic Instructions	73
12.3	Cryptographic performance	73
12.4	SPARC M7 crypto coding guidance	74
13	Memory Management Unit	75
13.1	Translation Table Entry (TTE)	75
13.2	Translation Storage Buffer (TSB)	77
13.3	MMU-Related Faults and Traps.	78
13.3.1	<i>IAE_privilege_violation</i> Trap	78
13.3.2	<i>IAE_nfo_page</i> Trap	78
13.3.3	<i>DAE_privilege_violation</i> Trap	78
13.3.4	<i>DAE_side_effect_page</i> Trap	78
13.3.5	<i>DAE_nc_page</i> Trap	78
13.3.6	<i>DAE_invalid_asi</i> Trap	79
13.3.7	<i>DAE_nfo_page</i> Trap	79
13.3.8	<i>privileged_action</i> Trap	79
13.3.9	This trap occurs when an access is attempted using a restricted ASI while in non-privileged mode (PSTATE.priv = 0). <i>*mem_address_not_aligned</i> Traps	79
13.4	MMU Operation Summary.	79
13.5	Translation	81
13.5.1	Instruction Translation.	81
13.5.1.1	Instruction Prefetching	81
13.5.2	Data Translation	81
13.6	Compliance With the SPARC V9 Annex F	84
13.7	MMU Internal Registers and ASI Operations	85
13.7.1	Accessing MMU Registers.	85
13.7.2	Context Registers	85
A	Programming Guidelines	151
A.1	Multithreading	151
A.1.1	Instruction fetch	151
A.1.2	Select/Decode/Rename	151
A.1.3	Pick/Issue/Execute	152
A.1.4	Commit.	152
A.1.5	Context Switching Between Strands.	152
A.1.6	Synchronization	152
A.2	Optimizing for Single-Threaded Performance or Throughput	153
A.3	Instruction Latency	153
B	IEEE 754 Floating-Point Support	163
B.1	Special Operand and Result Handling	163
C	Differences Between SPARC M7 and SPARC M6.	165
C.1	Architectural and Microarchitectural Differences	166
C.2	Address Spaces and ASIs Differences	166
C.2.1	ASIs.	166

D	Cache Coherency and Ordering	167
D.1	Cache and Memory Interactions	167
D.2	Cache Flushing	167
D.2.1	Displacement Flushing	168
D.2.2	Memory Accesses and Cacheability	168
D.2.3	Coherence Domains	168
D.2.3.1	Cacheable Accesses	169
D.2.3.2	Noncacheable and Side-Effect Accesses	169
D.2.3.3	Global Visibility and Memory Ordering	169
D.2.4	Memory Synchronization: MEMBAR and FLUSH	170
D.2.4.1	MEMBAR #LoadLoad	170
D.2.4.2	MEMBAR #StoreLoad	170
D.2.4.3	MEMBAR #LoadStore	170
D.2.4.4	MEMBAR #StoreStore and STBAR	171
D.2.4.5	MEMBAR #Lookaside	171
D.2.4.6	MEMBAR #MemIssue	171
D.2.4.7	MEMBAR #Sync (Issue Barrier)	171
D.2.4.8	Self-Modifying Code (FLUSH)	171
D.2.5	Atomic Operations	172
D.2.5.1	SWAP Instruction	172
D.2.5.2	LDSTUB Instruction	172
D.2.5.3	Compare and Swap (CASX) Instruction	172
D.2.6	Nonfaulting Load	172
D.3	L1 I-Cache	173
D.3.1	LRU Replacement Algorithm	173
D.3.2	Direct-Mapped Mode	173
D.3.3	I-Cache Disable	173
D.4	L1 D-Cache	174
D.4.1	LRU Replacement Algorithm	174
D.4.2	Direct-Mapped Mode	174
D.4.3	D-Cache Disable	174
D.5	L2 Instruction Cache	174
D.5.1	NRU Replacement Algorithm	175
D.5.1.1	Mapping Out Lines	175
D.5.2	Directory Coherence	175
D.5.3	L2I Cache Disable	175
D.6	L2 Data Cache	176
D.6.1	NRU Replacement Algorithm	176
D.6.1.1	Mapping Out Lines	176
D.6.2	Directory Coherence	176
D.6.3	L2 Cache Disable	177
E	Glossary	179
F	Bibliography	181
	Index	183

SPARC M7 Basics

1.1 Background

SPARC M7 is the follow-on chip multi-threaded (CMT) processor to the SPARC M6 processor. SPARC M7 incorporates a new processor core (Core S4) a new L2 cache, and a new L3 cache structure. IO has moved off chip to an external ASIC, connected to SPARC M7 with a high speed serial link interface.


The SPARC M7 product line fully implements Oracle's Throughput Computing initiative for the horizontal system space. 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.

SPARC M7 supports up to a eight way glueless (without external hub chips) coherent system using 7 coherence link channels. SPARC M7 has 32 SPARC physical processor cores. Each core has full hardware support for eight strands, two integer execution pipelines, one floating-point execution pipeline, and one memory pipeline. The SPARC cores are connected to L2 and L3 caches. There are 8 L3 caches, each is 8 MB, 2-banked, 8 way associative per bank.

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 processors 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 processor is 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 processors that employ multithreading are built in the same chip area. Combining multiple processors on a single chip with multiple strands per

processor provides very high performance for highly threaded commercial applications. This approach is called thread-level parallelism (TLP), and the difference between TLP and ILP is shown in the 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 processor, and multiple processors 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.¹

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 processors of today are the evolutionary outgrowth from a time when the processor was the bottleneck in delivering good performance. With processors 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.

Unlike first-generation TLP processors, SPARC M7 seeks to provide the best of TLP and ILP processors. In particular, SPARC M7 provides a robust out-of-order, dual-issue processor core that is heavily threaded among eight strands. It has a 16-stage integer pipeline to achieve high operating frequencies, advanced branch prediction to mitigate the effect of a deep pipeline, and dynamic allocation of processor resources to threads. This allows SPARC M7 to achieve very high single-thread performance while still scaling to very high levels of throughput.

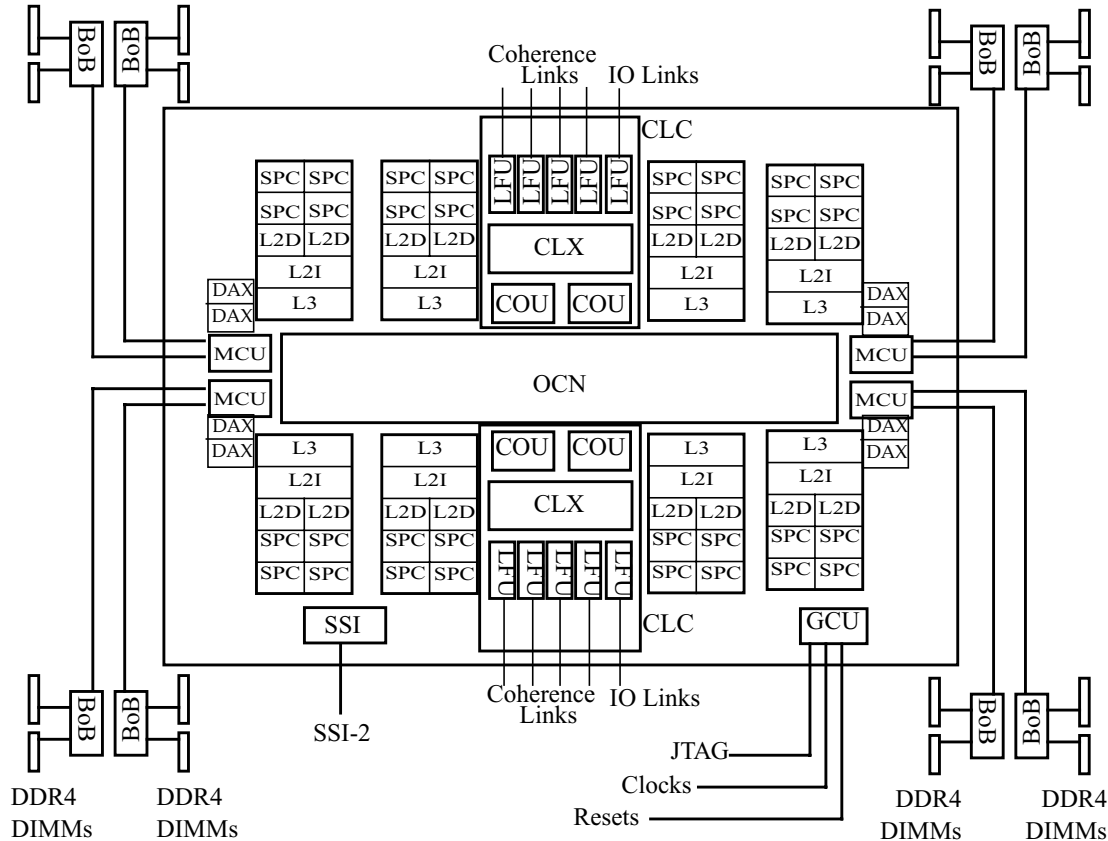
1.2 SPARC M7 Overview

SPARC M7 is a chip multi-threaded (CMT) processor which supports cache-coherent multi-socket systems. SPARC M7 contains 32 SPARC physical processor cores. Two SPARC physical cores connect to a single 256 KB L2 data cache of 2 banks and 8 ways, and four SPARC physical cores connect to a single 256 KB L2 instruction cache of 2 banks and 8 ways. The L2 instruction and data caches connect

¹. 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.

to an L3 cache. L2 and L3 caches have 64-byte lines. Each L3 cache is banked two ways and is local to the four SPARC cores, the two L2 data caches, and the L2 instruction cache that attach to each L3 cache. This collection of four cores, L2 caches, and L3 cache is called a SPARC Cache Cluster (SCC).

SPARC M7



1.3 SPARC M7 Components

This section describes each component in SPARC M7.

1.3.1 SPARC Physical Core

Each SPARC physical core has hardware support for eight strands. This support consists of a full integer register file with eight register windows per strand, a full floating-point register file per strand, and nearly all of the ASI, ASR, and privileged registers replicated per strand. The eight strands share the instruction and data caches. Each SPARC physical core has a 16 KB, 4-way set-associative instruction cache with 64-byte lines, a 16 KB, 4-way set-associative data cache (32-byte lines) that are shared by the eight strands. The L1 data cache is write-through and does not allocate on a write miss; the L2 is store-in and allocates on a write miss. All strands share a floating-point unit incorporating fused multiply-add and VIS3.0 instruction support.

Two physical cores share a 256KB, 8-way set-associative L2D cache with 64B lines. Four physical cores share a 256KB, 8-way set-associative L2I cache with 64B lines.

The strands share a dual-issue, out-of-order pipeline, divided into two "slots". One instruction can be issued each cycle to each slot. Slot 0 contains an integer unit and a load/store unit, while slot 1 contains an integer unit, a branch unit, and a floating-point and graphics unit. Up to two instructions can complete each cycle for a peak operation rate of two instructions per cycle. The pipeline is both horizontally and vertically threaded; various segments of the pipeline handle strands differently. The instruction fetch unit fetches instructions from a given strand each cycle. Strands are selected for fetching based upon a least-recently-fetched algorithm. Once fetched, strands are then selected for decoding in a least-recently-decoded fashion and are then renamed and supplied into an out-of-order processor core. Once inside the out-of-order core, strands are picked for issue independently between slots, and in an oldest-ready-first fashion within a slot. Instructions complete out-of-order and are committed in-order within a strand, but independently between strands. Up to 128 instructions can be in flight within the processor core, in any combination across the active strands. In certain circumstances, hardware may activate heuristics to avoid starvation or performance imbalances resulting from unfair access to hardware resources. The L1 cache load-use latency is 5 cycles, the L2 cache load-use latency is 19 cycles, and the L3 load-use latency is 41 cycles.

1.3.1.1 Single-threaded and multi-threaded performance

SPARC M7 is dynamically threaded. While software can activate up to 8 strands on each core at a time, hardware dynamically and seamlessly allocates core resources such as instruction, data, and L2 caches, and out-of-order execution resources such as the 128-entry re-order buffer in the core, among the active strands.

Since the core dynamically allocates resources among the active strands, there is no explicit "single-thread mode" or "multi-thread mode" for software to activate or deactivate. The extent to which strands compete for core resources depends upon their execution characteristics. These characteristics include cache footprints, inter-instruction dependencies in their execution streams, branch prediction effectiveness, and others. Consider one process which has a small cache footprint and a high correct branch prediction rate which, when running alone on a core, achieves 2 instructions per cycle (SPARC M7's peak rate of instruction execution). We term this a high-IPC process. If another process with similar characteristics is activated on a different strand on the same core, each of the strands will likely operate at approximately 1 instruction per cycle. In other words, the single-thread performance of each process has been cut in half. As a rule of thumb, activating N high-IPC strands will result in each strand executing at 1/N of its peak rate, assuming each strand is capable of executing close to 2 instructions per cycle.

Now consider a process which is largely memory-bound. Its native IPC will be small, perhaps 0.2. If this process runs on one strand on a core with another clone process running on a different strand, there is a good chance that both strands will suffer no noticeable performance loss, and the core throughput will improve to 0.4 IPC. If a low-IPC process runs on one strand with a high-IPC process running on another strand, it's likely that the IPC of either strand will not be greatly perturbed. The high-IPC strand may suffer a slight performance degradation (as long as the low-IPC strand does not cause a substantial increase in cache miss rates for the high-IPC strand).

The guidelines above are only general rules-of-thumb. The extent to which one strand affects another strand's performance depends upon many factors. Processes which run fine on their own but suffer from destructive cache interference when run with other strands may suffer unacceptable performance losses. Similarly, it is also possible for strands to cooperatively improve performance when run together. This may occur when the strands running on one core share code or data. In this case, one strand may prefetch instructions or data that other strands will use in the near future.

The same discussion can apply between cores running in the chip. Since the L3 cache and memory controllers are shared between the cores, activity on one core can influence the performance of strands on another core.

1.3.2 L3 Cache

Each of the eight L3 cache is banked two ways. It is inclusive of all chip-local L2 caches. Each L3 cache is 8 Mbytes, and each bank is 8-way set associative. The line size is 64 bytes.

Data Formats



Data formats supported by SPARC M7 are described in the Oracle SPARC Architecture 2015 specification.

Registers

3.1 Floating-Point State Register (FSR)

Each virtual processor has a Floating-Point State register. This register follows the Oracle SPARC Architecture 2015 specification, with the `ver` and `qne` fields permanently set to 0 (SPARC M7 does not support a FQ).

For more information on this register, see the Oracle SPARC Architecture 2015 specification.

3.2 Ancillary State Registers (ASRs)

This chapter discusses the SPARC M7 ancillary state registers. TABLE 3-1 summarizes and defines these registers.

TABLE 3-1 Summary of SPARC M7 Ancillary State Registers

ASR number	ASR Name	Access	priv	Description
0	Y	RW	N	Y Register
1	<i>Reserved</i>	—		Any access causes a <i>illegal_instruction</i> trap
2	CCR	RW	N	Condition Code register
3	ASI	RW	N	ASI register
4	TICK	RO	Y ¹	TICK register
5	PC	RO ²	N	Program counter
6	FPRS	RW	N	Floating-Point Registers Status register
07 - 13	<i>Reserved</i>	-		Any access causes an <i>illegal_instruction</i> trap
15	(MEMBAR, STBAR)	—	N	Instruction opcodes only, not an actual ASR.
16 - 18	<i>Reserved</i>	—		Any access causes an <i>illegal_instruction</i> trap
19	GSR	RW	N	General Status register
20	SOFTINT_SET	W	Y ⁴	Set bit in Soft Interrupt register
21	SOFTINT_CLR	W	Y ⁴	Clear bit in Soft Interrupt register

TABLE 3-1 Summary of SPARC M7 Ancillary State Registers (Continued)

ASR number	ASR Name	Access	priv	Description
22	SOFTINT	RW	Y ³	Soft Interrupt register
23	<i>Reserved</i>	-		Any access causes an <i>illegal_instruction</i> trap
24	STICK	RW	Y ⁵	System Tick register
25	STICK_CMPR	RW	Y ³	System TICK Compare register
26	CFR	RO ⁶	Y	Compatibility Feature Register
27	PAUSE	W	N	Any read causes an <i>illegal_instruction</i> trap; PAUSE is write-only
28	MWAIT	W	N	Any read causes an <i>illegal_instruction</i> trap; MWAIT is write-only
29 - 31	<i>Reserved</i>	—		Any access causes an <i>illegal_instruction</i> trap

Notes:

1. An attempted write by nonprivileged software to this register causes a *privileged_opcode* trap. An attempted write by privileged software to this register causes an *illegal_instruction* trap. See the Oracle SPARC Architecture 2015 specification for more detail.
2. A write to this register causes an *illegal_instruction* trap.
3. An attempted access in nonprivileged mode causes a *privileged_opcode* trap.
4. Read accesses cause an *illegal_instruction* trap. An attempted write access in nonprivileged mode causes a *privileged_opcode* trap.
5. A write by privileged or user software causes an *illegal_instruction* trap. See the Oracle SPARC Architecture 2015 specification for more detail.
6. Reads are nonprivileged. A write by privileged or user software causes an *illegal_instruction* trap.

3.2.1 Tick Register (TICK)

The TICK register contains one field: **counter**. The **counter** field is shared by the eight strands on a physical core. The **counter** increments each processor core clock. The format of this register is shown in TABLE 3-2.

For more information on this register, see the Oracle SPARC Architecture 2015 specification.

TABLE 3-2 TICK Register – TICK (ASR 04₁₆)

Bit	Field	RW	Description
63	—	RO	<i>Reserved</i>
62:0	counter	RW	Tick counter, increments each processor core clock cycle.

3.2.2 Program Counter (PC)

Each strand has a read-only program counter register. The PC contains a 54-bit virtual address and VA{63:54} is sign-extended from VA{53}. The format of this register is shown in TABLE 3-3.

TABLE 3-3 Program Counter – PC (ASR 05₁₆)

Bit	Field	R/W	Description
63:54	va_high	RO	Sign-extended from VA{53}.
53:2	va	RO	Virtual address contained in the program counter.
1:0	—	RO	The lower 2 bits of the program counter always read as 0.

3.2.3 Floating-Point Registers State Register (FPRS)

This register is described in Oracle SPARC Architecture 2015.

Implementation Note SPARC M7 sets FPRS.du or FPRS.dl when an instruction that updates the floating-point register file successfully completes, or when an FMOVcc or FMOVr instruction that does not satisfy the destination register update condition successfully completes.

3.2.4 General Status Register (GSR)

Each virtual processor has a nonprivileged general status register (GSR). When PSTATE.pef or FPRS.fef is zero, accesses to this register cause an *fp_disabled* trap.

For more information on this register, see the Oracle SPARC Architecture 2015 specification.

3.2.5 Software Interrupt Register (SOFTINT)

Each virtual processor has a privileged software interrupt register. Nonprivileged accesses to this register cause a *privileged_opcode* trap. The SOFTINT register contains two fields: *sm*, and *int_level*. Note that while setting the *sm* (bit 16) or SOFTINT{14} bits generate *interrupt_level_14*, these bits are considered completely independent of each other. Thus an STICK compare will only set bit 16 and generate *interrupt_level_14*, not also set bit 14.

TABLE 3-4 specifies how *interrupt_level_14* is shared between SOFTINT writes and STICK compares.

TABLE 3-4 Sharing of *interrupt_level_14*

Event	SOFTINT{14}	sm	Action
STICK compare when <i>sm</i> = 0	Unchanged	1	<i>interrupt_level_14</i> if PSTATE.ie = 1 and PIL < 14
Set <i>sm</i> = 1 when <i>sm</i> = 0	Unchanged	1	<i>interrupt_level_14</i> if PSTATE.ie = 1 and PIL < 14
Set SOFTINT{14} = 1 when SOFTINT{14} = 0.	1	Unchanged	<i>interrupt_level_14</i> if PSTATE.ie = 1 and PIL < 14

For more information on this register, see the Oracle SPARC Architecture 2015 specification.

3.2.6 System Tick Register (STICK)

Each SPARC M7 physical processor core implements an STICK register, shared by all strands of that core.

TABLE 3-5 System Tick Register – STICK (ASR 18₁₆)

Bit	Field	R/W	Description
63	—	RO	<i>Reserved.</i>
62:0	stick	RW	Elapsed time value, measured in increments of 1 nS.

Privileged software can read the STICK register with the RDSTICK instruction. Privileged software cannot write the STICK register; an attempt by privileged software to execute the WRSTICK instruction results in an *illegal_instruction* exception.

Nonprivileged software can read the STICK register with RDSTICK instruction. Nonprivileged software cannot write the STICK register; an attempt by nonprivileged software to execute the WRSTICK instruction results in an *illegal_instruction* exception.

In SPARC M7, the difference of the values of two different reads of the STICK register reflects the amount of time that has passed between the reads;

(value2 - value1) * 1 = the number of nanoseconds that passed between the reads.

For more information on this register, see the Oracle SPARC Architecture 2015 specification.

3.2.7 System Tick Compare Register (STICK_CMPR)

Each virtual processor has a privileged System Tick Compare (STICK_CMPR) register. Nonprivileged accesses to this register cause a *privileged_opcode* exception. STICK_CMPR contains two fields: `int_dis` and `stick_cmpr`. Only bits 62:9 of the `stick_cmpr` field are compared against the STICK counter field.

The `int_dis` bit controls whether a STICK *interrupt_level_14* interrupt is posted in the SOFTINT register when STICK_CMPR bits 62:9 match STICK bits 62:9. The format of this register is shown in TABLE 3-6.

TABLE 3-6 System Tick Compare Register – STICK_CMPR (ASR 19₁₆)

Bit	Field	R/W	Description
63	<code>int_dis</code>	RW	<code>stick_int</code> interrupt disable. If 1, <code>stick_int</code> interrupt generation is disabled.
62:9	<code>stick_cmpr</code>	RW	Compare value for <code>stick_int</code> interrupts.
8:0	—	RO	<i>Reserved.</i>

After a power-on reset trap, STICK_CMPR.int_dis is set to 1 and STICK_CMPR.cmpr is undefined.

An *stick_match* exception occurs in the cycle in which all of the following three conditions are met:

1. STICK_CMPR.int_dis == 0.

2. A transition occurs from

$$(\text{STICK.counter})[62:9] < \text{STICK_CMPR.cmpr}[62:9]$$

in one cycle, to

$$(\text{STICK.counter})[62:9] \geq \text{STICK_CMPR.cmpr}[62:9]$$

in the following cycle

3. This transition of state occurs due to incrementing STICK, and not due to writing STICK, or STICK_CMPR

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

The reason the comparison of STICK_CMPR and STICK ignore bits 8 to 0 (and the reason STICK_CMPR is not implemented below bit 9) is because at the minimum frequency of the processor core, the STICK register could reflect as much as 384 ns passing between cycles (due to acceleration of the STICK register increment by the 'drift fix').

For more information on this register, see the Oracle SPARC Architecture 2015 specification.

3.2.8 Compatibility Feature Register (CFR)

For general information on this register, see the Oracle SPARC Architecture 2015 specification.

Each virtual processor has a compatibility feature register (CFR). The CFR is read-only. The format of the CFR is shown in Table 3-7 .

TABLE 3-7 Compatibility Feature Register – CFR (ASR 1A₁₆)

Bit	Field	R/W	Description
63:15	—	RO	<i>Reserved</i>
14	xmontsq	RO	If set, the processor supports the XMONTSQR opcode. If not set, an attempt to execute an XMONTSQR instruction results in a <i>compatibility_feature</i> trap.
13	xmontmul	RO	If set, the processor supports the XMONTMUL opcode. If not set, an attempt to execute an XMONTMUL instruction results in a <i>compatibility_feature</i> trap.
12	xmpmul	RO	If set, the processor supports the XMPMUL opcode. If not set, an attempt to execute an XMPMUL instruction results in a <i>compatibility_feature</i> trap.
11	crc32c	RO	If set, the processor supports the CRC32C opcode. If not set, an attempt to execute a CRC32C instruction results in a <i>compatibility_feature</i> trap.
10	montsq	RO	If set, the processor supports the MONTSQR opcode. If not set, an attempt to execute a MONTSQR instruction results in a <i>compatibility_feature</i> trap.
9	montmul	RO	If set, the processor supports the MONTMUL opcode. If not set, an attempt to execute a MONTMUL instruction results in a <i>compatibility_feature</i> trap.
8	mpmul	RO	If set, the processor supports the MPMUL opcode. If not set, an attempt to execute an MPMUL instruction results in a <i>compatibility_feature</i> trap.
7	sha512	RO	If set, the processor supports the SHA512 opcode. If not set, an attempt to execute a SHA512 instruction results in a <i>compatibility_feature</i> trap.
6	sha256	RO	If set, the processor supports the SHA256 opcode. If not set, an attempt to execute a SHA256 instruction results in a <i>compatibility_feature</i> trap.
5	sha1	RO	If set, the processor supports the SHA1 opcode. If not set, an attempt to execute a SHA1 instruction results in a <i>compatibility_feature</i> trap.
4	md5	RO	If set, the processor supports the MD5 opcode. If not set, an attempt to execute an MD5 instruction results in a <i>compatibility_feature</i> trap.
3	camellia	RO	If set, the processor supports Camellia opcodes (CAMELLIA_F, CAMELLIA_FL, and CAMELLIA_FLI). If not set, an attempt to execute a Camellia instruction results in a <i>compatibility_feature</i> trap.

TABLE 3-7 Compatibility Feature Register – CFR (ASR 1A₁₆)

Bit	Field	R/W	Description
2	—	RO	<i>Reserved</i>
1	des	RO	If set, the processor supports DES opcodes (DES_ROUND, DES_IP, DES_IIP, and DES_KEXPAND). If not set, an attempt to execute a DES instruction results in a <i>compatibility_feature</i> trap.
0	aes	RO	If set, the processor supports AES opcodes (AES_EROUND01, AES_EROUND23, AES_DROUND01, AES_DROUND23, AES_EROUND_01_LAST, AES_EROUND_23_LAST, AES_DROUND_01_LAST, AES_DROUND_23_LAST, AES_KEXPAND0, AES_KEXPAND1, and AES_KEXPAND2). If not set, an attempt to execute an AES instruction results in a <i>compatibility_feature</i> trap.

The CFR enumerates the capabilities that SPARC M7 supports. While the current definition of the CFR only relates to cryptographic capability, additional capabilities may be added in future processors. Software can use the CFR to determine whether a set of cryptographic opcodes associated with a cryptographic function can be executed on an instance of SPARC M7. Hardware also uses the CFR to determine whether a cryptographic capability associated with an opcode is present. When SPARC M7 executes a cryptographic opcode, it associates a bit in the CFR with each opcode; the bit must be set, otherwise a *compatibility_feature* trap occurs.

Programming Note For optimal performance, prior to using instruction-level cryptographic functions, applications and libraries should first check the CFR to ensure that the desired algorithm is supported by the hardware.

The CFR allows software to construct an architecture that enables opcode reuse. A complete discussion is outside the scope of this document; however, a brief overview follows.

Consider the situation where a processor is introduced that supports three cryptographic opcodes: opA, opB, and opC. Cryptographic requirements could be such that opA=AES, opB=DES, and opC=Kasumi. Traditionally, for any derivative or next-generation processor for which different ciphers were of interest, it would be necessary to expend additional opcodes to achieve the necessary support: e.g. opD=Camellia, opE=MD5. OpA, OpB, and OpC would still be consumed in these follow-on processors, even if there was no longer any interest in the AES, DES, and Kasumi algorithms.

In conjunction with appropriate software architecture and infrastructure, the CFR enables opcode reuse by future processor generations when cryptographic algorithms become obsolete. Potential aliasing problems are disambiguated using the CFR. Each bit in the CFR is permanently assigned to a different cryptographic operation. For instance, bits 0, 1, and 2 are assigned to AES, DES, and Kasumi family opcodes, as shown above. The mapping in the CFR is fixed for all future and derivative processors. When an application wishes to perform an AES operation, it registers that request using the appropriate software architectural means, and uses opA in its binary. Prior to executing, system software or the application checks to make sure that the target processor binds the AES function to opA. It does so by examining the CFR to see if bit 0 is set. If so, the program executes using native AES instructions (opA); if not, system software and/or the application must support a non-native AES instruction implementation using standard instructions. It is expected that cryptographic libraries will contain the necessary checking, so hardware cryptographic support will be transparent to applications that perform cryptographic operations using cryptographic library calls. If the application does not use cryptographic libraries, it should check the CFR to make sure that hardware supports the appropriate function, otherwise it should emulate the function using standard instructions. Alternatively, if performance is not critical, it may rely on trap-and-emulate support provided by higher-level system software.

When the first generation of processor (G1) executes an AES opcode it checks that CFR bit 0 is set. If so, the hardware performs the requested AES operation. Accordingly, on G1, an application is free to perform AES operations using opA. Similar enforcement is applied to DES and Kasumi, respectively.

Now consider what happens if the application is moved to a future processor (G2) which has re-used opA to provide support for Camellia; i.e. opA=Camellia. When system software checks the capabilities for the program, or the program checks, it will see that G2 does not support AES using opA (CFR bit 0 will be 0). This allows system software or the application to emulate AES support using standard instructions. Note that if the application somehow runs without this check having been performed and issues opA, the G2 processor will examine the CFR bit for Camellia, and if set, the application will execute, and get erroneous results (Camellia instead of AES). A similar problem exists if the application is developed for G2 hardware, but somehow runs on a G1 processor. Thus it is vital that system software and/or the application appropriately register their intent and check hardware capability prior to executing cryptographic opcodes.

As a result, given appropriate software infrastructure, instruction set designers may reuse opcodes to perform a variety of different operations and applications will continue to see the expected results on different generation platforms.

3.2.9 Pause (PAUSE)

SPARC M7 physically implements a 16-bit PAUSE register. The value written to the PAUSE register via the WRPAUSE instruction is an unsigned 20-bit value that is then right-shifted by 4 bits (divided by 16) since hardware decrements the PAUSE register once every 16 ns. Thus the unsigned 16-bit value represents a count from 0 to a maximum of 1048576 ns. Writing to the non-privileged PAUSE register stalls a thread for the number of nanoseconds specified by the XOR of the source operands, except as follows:

1. Writing 0 to the PAUSE register stalls the thread for the minimum time (greater than zero since there is a minimum stall time due to internal pipeline delays).
2. Writing a value larger than $2^{20} - 1$ causes hardware to saturate the 16-bit PAUSE register; hardware sets PAUSE to F_FFF0_{16} prior to decrementing it.
3. If the STICK register is disabled (not measuring time), then writing any value to the PAUSE register behaves the same as writing 0 to the PAUSE register.

When the PAUSE register is written to a nonzero value, the strand is scheduled to be flushed and made inactive (i.e., all resources released by that strand, no core activity other than PAUSE register maintenance and monitoring for unmasked disrupting exceptions). No instructions are fetched by the strand while its PAUSE register is nonzero. An unmasked disrupting exception terminates the PAUSE. Once the PAUSE register reaches value 0 or an unmasked disrupting exception occurs, the virtual processor restarts fetch and execution of the strand. The reactivated strand restarts at either the instruction following the WRPAUSE or a disrupting trap handler.

For more information on this instruction, see the Oracle SPARC Architecture 2015 specification or Section 5.3, *WRPAUSE*, on page 36.

Note | Prior implementations implemented PAUSE based on cycles, not nanoseconds. Software written for prior implementations using PAUSE may need to adjust due to this change.

3.2.10 MWAIT

SPARC M7 shares the physical implementation of the PAUSE register for much of the MWAIT register functionality. That is, the counters for the PAUSE and MWAIT registers are physically the same counter. See the description of the PAUSE register in Section 3.2.9, *Pause (PAUSE)*, on page 23.

For more information on MWAIT, see the Oracle SPARC Architecture 2015 specification.

3.3 Privileged PR State Registers

TABLE 3-8 lists the privileged registers.

TABLE 3-8 Privileged Registers

Register	Register Name	Access	Description
0	TPC	RW	Trap PC ¹
1	TNPC	RW	Trap Next PC ¹
2	TSTATE	RW	Trap State
3	TT	RW	Trap Type
4	TICK	RW	Tick
5	TBA	RW	Trap Base Address ¹
6	PSTATE	RW	Process State
7	TL	RW	Trap Level
8	PIL	RW	Processor Interrupt Level
9	CWP	RW	Current Window Pointer
10	CANSAVE	RW	Savable Windows
11	CANRESTORE	RW	Restorable Windows
12	CLEANWIN	RW	Clean Windows
13	OTHERWIN	RW	Other Windows
14	WSTATE	RW	Window State
16	GL	RW	Global Level

1. SPARC M7 only implements bits 53:0 of the TPC, TNPC, and TBA registers. Bits 63:54 are always sign-extended from bit 53.

3.3.1 Trap State Register (TSTATE)

Each virtual processor has *MAXPTL* (2) Trap State registers. These registers hold the state values from the previous trap level. The format of one element the TSTATE register array (corresponding to one trap level) is shown in TABLE 3-9.

TABLE 3-9 Trap State Register

Bit	Field	R/W	Description
63:42	—	RO	<i>Reserved.</i>
41:40	gl	RW	Global level at previous trap level
39:32	ccr	RW	CCR at previous trap level
31:24	asi	RW	ASI at previous trap level
23:21	—	RO	<i>Reserved</i>
20	pstate tct	RW	PSTATE.tct at previous trap level
18	—	RO	<i>Reserved</i> (corresponds to bit 10 of PSTATE)
17	pstate cle	RW	PSTATE.cle at previous trap level
16	pstate tle	RW	PSTATE.tle at previous trap level
15:13	—	RO	<i>Reserved</i> (corresponds to bits 7:5 of PSTATE)

TABLE 3-9 Trap State Register (Continued)

Bit	Field	R/W	Description
12	pstate pef	RW	PSTATE.pef at previous trap level
11	pstate am	RW	PSTATE.am at previous trap level
10	pstate priv	RW	PSTATE.priv at previous trap level
9	pstate ie	RW	PSTATE.ie at previous trap level
8	—	RO	Reserved (corresponds to bit 0 of PSTATE)
7:3	—	RO	Reserved
2:0	cwp	RW	CWP from previous trap level

For more information on this register, see the Oracle SPARC Architecture 2015 specification.

3.3.2 Processor State Register (PSTATE)

Each virtual processor has a Processor State register. More details on PSTATE can be found in the Oracle SPARC Architecture 2015 specification. The format of this register is shown in TABLE 3-10; note that the memory model selection field (mm) mentioned in Oracle SPARC Architecture 2015 is not implemented in SPARC M7.

TABLE 3-10 Processor State Register

Bit	Field	R/W	Description
63:13	—	RO	Reserved
12	tct	RW	Trap on control transfer
10	—	RO	Reserved
9	cle	RW	Current little endian
8	tle	RW	Trap little endian
7:6	—	RO	Reserved (mm; not implemented in SPARC M7)
5	—	RO	Reserved
4	pef	RW	Enable floating-point
3	am	RW	Address mask
2	priv	RW	Privileged mode
1	ie	RW	Interrupt enable
0	—	RO	Reserved (was ag)

Programming Note | Hyperprivileged changes to translation in delay slots of delayed control transfer instructions should be avoided.

For more information on this register, see the Oracle SPARC Architecture 2015 specification.

3.3.3 Trap Level Register (TL)

Each virtual processor has a Trap Level register. Writes to this register saturate at *MAXPTL* (2). This saturation is based on bits 2:0 of the write data; bits 63:3 of the write data are ignored.

For more information on this register, see the Oracle SPARC Architecture 2015 specification.

3.3.4 Current Window Pointer (CWP) Register

Since *N_REG_WINDOWS* = 8 on SPARC M7, the CWP register in each virtual processor is implemented as a 3-bit register.

For more information on this register, see the Oracle SPARC Architecture 2015 specification.

3.3.5 Global Level Register (GL)

Each virtual processor has a Global Level register, which controls which set of global register windows is in use. The maximum global level (*MAXPGL*) for SPARC M7 is 2. GL is implemented as a 2-bit register on SPARC M7. On a trap, GL is set to $\min(\text{GL} + 1, \text{MAXPTL})$.

Writes to the GL register saturate at *MAXPTL*. This saturation is based on bits 1:0 of the write data; bits 63:2 of the write data are ignored.

The format of the GL register is shown in TABLE 3-11.

TABLE 3-11 Global Level Register

Bit	Field	R/W	Description
63:2	—	RO	<i>Reserved</i>
1:0	gl	RW	Global level.

For more information on this register, see the Oracle SPARC Architecture 2015 specification.

Instruction Format

Instruction formats are described in the Oracle SPARC Architecture 2015 specification.

Instruction Definitions

5.1 Instruction Set Summary

The SPARC M7 CPU implements the Oracle SPARC Architecture 2015 instruction set.

TABLE 5-1 lists the complete SPARC M7 instruction set supported in hardware. All instructions that are part of Oracle SPARC Architecture 2015 are documented in the Oracle SPARC Architecture 2015 specification; any instructions that are extensions to OSA 2011 are documented in this chapter.

TABLE 5-1 Complete SPARC M7 Hardware-Supported Instruction Set (1 of 6)

Opcode	Description
ADD (ADDcc)	Add (and modify condition codes)
ADDC (ADDCcc)	Add with carry (and modify condition codes)
ADDXC (ADDXCcc)	Add extended with carry (and modify condition codes)
AES_DROUND01	AES decrypt round, columns 0 & 1
AES_DROUND23	AES decrypt round, columns 2 & 3
AES_DROUND01_LAST	AES decrypt last round, columns 0 & 1
AES_DROUND23_LAST	AES decrypt last round, columns 2 & 3
AES_EROUND01	AES encrypt round, columns 0 & 1
AES_EROUND23	AES encrypt round, columns 2 & 3
AES_EROUND01_LAST	AES encrypt last round, columns 0 & 1
AES_EROUND23_LAST	AES encrypt last round, columns 2 & 3
AES_KEXPAND0	AES key expansion without round constant
AES_KEXPAND1	AES key expansion with round constant
AES_KEXPAND2	AES key expansion without SBOX
ALIGNADDRESS	Calculate address for misaligned data access
ALIGNADDRESS_LITTLE	Calculate address for misaligned data access (little-endian)
ALLCLEAN	Mark all windows as clean
AND (ANDcc)	And (and modify condition codes)
ANDN (ANDNcc)	And not (and modify condition codes)
ARRAY{8,16,32}	3-D address to blocked byte address conversion
Bicc	Branch on integer condition codes
BMASK	Writes the GSR.mask field
BPcc	Branch on integer condition codes with prediction
BPr	Branch on contents of integer register with prediction
BSHUFFLE	Permutates bytes as specified by the GSR.mask field
CALL ¹	Call and link
CAMELLIA_F	Camellia F operation
CAMELLIA_FL	Camellia FL operation

TABLE 5-1 Complete SPARC M7 Hardware-Supported Instruction Set (2 of 6)

Opcode	Description
CAMELLIA_FLI	Camellia FLI operation
CASA	Compare and swap word in alternate space
CASXA	Compare and swap doubleword in alternate space
CBcond	Fused 32 or 64 bit compare and conditional branch
CMASK{8,16,32}	Create GSR.maskfrom SIMD operation result
CRC32C	CRC32C polynomial instruction
DES_IP	DES initial permutation
DES_IIP	DES inverse initial permutation
DES_KEXPAND	DES key expansion
DES_ROUND	DES round
DONE	Return from trap
EDGE{8,16,32}{L}{N}	Edge boundary processing {little-endian} {non-condition-code altering}
FABS(s,d)	Floating-point absolute value
FADD(s,d)	Floating-point add
FALIGNDATAg	Perform data alignment for misaligned data
FALIGNDATAi	Perform data alignment for misaligned data using integer register
FANDNOT1{s,d}	Negated <i>src1</i> and <i>src2</i>
FANDNOT2{s,d}	<i>Src1</i> and negated <i>src2</i>
FAND{s,d}	Logical and
FBfcc	Branch on floating-point condition codes
FBPfcc	Branch on floating-point condition codes with prediction
FCHKSM16	16-bit partitioned checksum
FCMP(s,d)	Floating-point compare
FCMPE(s,d)	Floating-point compare (exception if unordered)
FDIV(s,d)	Floating-point divide
FEXPAND	Four 8-bit to 16-bit expand
FHADD{s,d}	Floating-point add and halve
FHSUB{s,d}	Floating-point subtract and halve
FiTO(s,d)	Convert integer to floating-point
FLCMP{s,d}	Lexicographic compare
FLUSH	Flush instruction memory
FLUSHW	Flush register windows
FMADD{s,d}	Floating-point multiply-add single/double (fused)
FMEAN16	16-bit partitioned average
FMOV(s,d)	Floating-point move
FMOV(s,d)cc	Move floating-point register if condition is satisfied
FMOV(s,d)R	Move floating-point register if integer register contents satisfy condition
FMSUB{s,d}	Floating-point multiply-subtract single/double (fused)
FMUL(s,d)	Floating-point multiply
FMUL8SUx16	Signed upper 8- x 16-bit partitioned product of corresponding components
FMUL8ULx16	Unsigned lower 8- x 16-bit partitioned product of corresponding components
FMUL8x16	8- x 16-bit partitioned product of corresponding components
FMUL8x16AL	Signed lower 8- x 16-bit lower α partitioned product of four components
FMUL8x16AU	Signed upper 8- x 16-bit lower α partitioned product of four components
FMULD8SUx16	Signed upper 8- x 16-bit multiply \rightarrow 32-bit partitioned product of components
FMULD8ULx16	Unsigned lower 8- x 16-bit multiply \rightarrow 32-bit partitioned product of components
FNADD(s,d)	Floating-point add and negate

TABLE 5-1 Complete SPARC M7 Hardware-Supported Instruction Set (3 of 6)

Opcode	Description
FNAND{s}	Logical nand (single precision)
FNEG(s,d)	Floating-point negate
FNHADD{s,d}	Floating-point add and halve, then negate
FNMADD{s,d}	Floating-point multiply-add and negate
FNMSUB{s,d}	Floating-point negative multiply-subtract single/double (fused)
FNMUL{s,d}	Floating-point multiply and negate
FNOR{s,d}	Logical nor
FNOT1{s,d}	Negate (1's complement) <i>src1</i>
FNOT2{s,d}	Negate (1's complement) <i>src2</i>
FNsMULd	Floating-point multiply and negate
FONE{s,d}	One fill
FORNOT1{s,d}	Negated <i>src1</i> or <i>src2</i>
FORNOT2{s,d}	<i>src1</i> or negated <i>src2</i>
FOR{s,d}	Logical or
FPACKFIX	Two 32-bit to 16-bit fixed pack
FPACK{16,32}	Four 16-bit/two 32-bit pixel pack
FPADD8	Eight 8-bit partitioned add
FPADD{16,32}{s}	Four 16-bit/two 32-bit partitioned add
FPADD64	Fixed-point partitioned add
FPADD{U}S8	Fixed-point partitioned add
FPADDS{16,32}{s}	Fixed-point partitioned add
FPADDUS16	Fixed-point partitioned add
FPCMPEQ{16,32}	Four 16-bit / two 32-bit compare: set integer dest if <i>src1</i> = <i>src2</i>
FPCMPGT{8,16,32}	Eight 8-bit / four 16-bit / two 32-bit compare: set integer dest if <i>src1</i> > <i>src2</i>
FPCMPLE{8,16,32}	Eight 8-bit / four 16-bit / two 32-bit compare: set integer dest if <i>src1</i> ≤ <i>src2</i>
FPCMPNE{16,32}	Four 16-bit / two 32-bit compare: set integer dest if <i>src1</i> ≠ <i>src2</i>
FPCMPU{GT,LE,NE,EQ}8	Compare 8-bit unsigned fixed-point values
FPCMPU{GT,LE}{16,32}	Compare four 16-bit/two 32-bit unsigned fixed-point values
FPMADDX	Unsigned integer multiply-add
FPMADDXHI	Unsigned integer multiply-add, return high-order 64 bits of result
FPMAX{U}{8,16,32}	Partitioned integer maximum
FPMERGE	Two 32-bit to 64-bit fixed merge
FPMIN{U}{8,16,32}	Partitioned integer minimum
FPSUB8	Eight 8-bit partitioned subtract
FPSUB{16,32}{s}	Four 16-bit/two 32-bit partitioned subtract (single precision)
FPSUB64	Fixed-point partitioned subtract, 64-bit
FPSUB{U}S8	Fixed-point partitioned subtract
FPSUBS{16,32}{s}	Fixed-point partitioned subtract
FPSUBUS16	Fixed-point partitioned subtract
FSL{16,32}	16- or 32-bit partitioned shift, left (old mnemonic FSHL)
FSLAS{16,32}	16- or 32-bit partitioned shift, left or right (old mnemonic FSHLAS)
FSRA{16,32}	16- or 32-bit partitioned shift, left or right (old mnemonic FSHRA)
FSRL{16,32}	16- or 32-bit partitioned shift, left or right (old mnemonic FSHRL)
FsMULd	Floating-point multiply single to double
FSQRT(s,d)	Floating-point square root
FSRC1{s,d}	Copy <i>src1</i>
FSRC2d	Copy <i>src2</i> (double precision)

TABLE 5-1 Complete SPARC M7 Hardware-Supported Instruction Set (4 of 6)

Opcode	Description
FSRC2s	Copy <i>src2</i> (single precision)
F(s,d)TO(s,d)	Convert between floating-point formats
F(s,d)TOi	Convert floating point to integer
F(s,d)TOx	Convert floating point to 64-bit integer
FSUB(s,d)	Floating-point subtract
FXNOR{s,d}	Logical xnor
FXOR{s,d}	Logical xor
FxTO(s,d)	Convert 64-bit integer to floating-point
FZERO{s}	Zero fill (single precision)
ILLTRAP	Illegal instruction
INVALW	Mark all windows as CANSAVE
JMPL	Jump and link
LDBLOCKF	64-byte block load
LDDF	Load double floating-point
LDDFA	Load double floating-point from alternate space
LDF	Load floating-point
LDFA	Load floating-point from alternate space
LDFSR	Load floating-point state register lower
LDSB	Load signed byte
LDSBA	Load signed byte from alternate space
LDSH	Load signed halfword
LDSHA	Load signed halfword from alternate space
LDSTUB	Load-store unsigned byte
LDSTUBA	Load-store unsigned byte in alternate space
LDSW	Load signed word
LDSWA	Load signed word from alternate space
LDTW	Load twin words
LDTWA	Load twin words from alternate space
LDUB	Load unsigned byte
LDUBA	Load unsigned byte from alternate space
LDUH	Load unsigned halfword
LDUHA	Load unsigned halfword from alternate space
LDUW	Load unsigned word
LDUWA	Load unsigned word from alternate space
LDX	Load extended
LDXA	Load extended from alternate space
LDXEFSR	Load extended floating-point state register
LDXFSR	Load extended floating-point state register
LZCNT	Leading zero count on 64-bit integer register
MD5	MD5 hash
MEMBAR	Memory barrier
MONTMUL	Montgomery multiplication
MONTSQR	Montgomery squaring
MOVcc	Move integer register if condition is satisfied
MOVdTOx	Move floating-point register to integer register
MOVr	Move integer register on contents of integer register
MOVsTO{u,s}w	Move floating-point register to integer register

TABLE 5-1 Complete SPARC M7 Hardware-Supported Instruction Set (5 of 6)

Opcode	Description
MOVwTOs	Move integer register to floating-point register
MOVxTOd	Move integer register to floating-point register
MPMUL	Multiple-precision multiplication
MULScC	Multiply step (and modify condition codes)
MULX	Multiply 64-bit integers
NOP	No operation
NORMALW	Mark other windows as restorable
OR (ORcc)	Inclusive-or (and modify condition codes)
ORN (ORNcc)	Inclusive-or not (and modify condition codes)
OTHERW	Mark restorable windows as other
PDIST	Distance between 8 8-bit components
PDISTN	Pixel component distance
POPC	Population count
PREFETCH	Prefetch data
PREFETCHA	Prefetch data from alternate space
PST	Eight 8-bit/4 16-bit/2 32-bit partial stores
RDASI	Read ASI register
RDASR	Read ancillary state register
RDCCR	Read condition codes register
RDCFR	Read compatibility feature register
RDFPRS	Read floating-point registers state register
RDPC	Read program counter
RDPR	Read privileged register
RTICK	Read TICK register
RESTORE	Restore caller's window
RESTORED	Window has been restored
RETRY	Return from trap and retry
RETURN	Return
SAVE	Save caller's window
SAVED	Window has been saved
SDIV (SDIVcc)	32-bit signed integer divide (and modify condition codes)
SDIVX	64-bit signed integer divide
SETHI	Set high 22 bits of low word of integer register
SHA1	SHA-1 hash
SHA256	SHA-256 hash
SHA512	SHA-512 hash
SIAM	Set interval arithmetic mode
SLL	Shift left logical
SLLX	Shift left logical, extended
SMUL (SMULcc)	Signed integer multiply (and modify condition codes)
SRA	Shift right arithmetic
SRAX	Shift right arithmetic, extended
SRL	Shift right logical
SRLX	Shift right logical, extended
STB	Store byte
STBA	Store byte into alternate space
STBAR	Store barrier

TABLE 5-1 Complete SPARC M7 Hardware-Supported Instruction Set (6 of 6)

Opcode	Description
STBLOCKF	64-byte block store
STD	Store doubleword
STDA	Store doubleword into alternate space
STDF	Store double floating-point
STDFA	Store double floating-point into alternate space
STF	Store floating-point
STFA	Store floating-point into alternate space
STFSR	Store floating-point state register
STH	Store halfword
STHA	Store halfword into alternate space
STPARTIALF	Eight 8-bit/4 16-bit/2 32-bit partial stores
STTW	Store twin words
STTWA	Store twin words into alternate space
STW	Store word
STWA	Store word into alternate space
STX	Store extended
STXA	Store extended into alternate space
STXFSR	Store extended floating-point state register
SUB (SUBcc)	Subtract (and modify condition codes)
SUBC (SUBCcc)	Subtract with carry (and modify condition codes)
SUBXC (SUBXCcc)	Subtract extended with carry (and modify condition codes)
SWAP	Swap integer register with memory
SWAPA	Swap integer register with memory in alternate space
TADDcc (TADDccTV)	Tagged add and modify condition codes (trap on overflow)
Tcc	Trap on integer condition codes (with 8-bit <code>sw_trap_number</code> , if bit 7 is set trap to hyperprivileged)
TSUBcc (TSUBccTV)	Tagged subtract and modify condition codes (trap on overflow)
UDIV (UDIVcc)	Unsigned integer divide (and modify condition codes)
UDIVX	64-bit unsigned integer divide
UMUL (UMULcc)	Unsigned integer multiply (and modify condition codes)
UMULXHI	Unsigned 64 x 64 multiply, returning upper 64 product bits
WRASI	Write ASI register
WRASR	Write ancillary state register
WRCCR	Write condition codes register
WRFPRS	Write floating-point registers state register
WRPR	Write privileged register
XMONTMUL	XOR Montgomery multiplication
XMONTSQR	XOR Montgomery squaring
XMPMUL	XOR multiple-precision multiplication
XMULX{HI}	XOR multiply
XNOR (XNORcc)	Exclusive-nor (and modify condition codes)
XOR (XORcc)	Exclusive-or (and modify condition codes)

1. The PC format saved by the CALL instruction is the same as the format of the PC register specified in Section 3.2.2, *Program Counter (PC)*, on page 18.

TABLE 5-2 lists the SPARC V9 and sun4v instructions that are not directly implemented in hardware by SPARC M7, and the exception that occurs when an attempt is made to execute them.

TABLE 5-2 Oracle SPARC Architecture 2015 Instructions Not Directly Implemented by SPARC M7 Hardware

Opcode	Description	Exception
FABSq	Floating-point absolute value quad	<i>illegal_instruction</i>
FADDq	Floating-point add quad	<i>illegal_instruction</i>
FCMPq	Floating-point compare quad	<i>illegal_instruction</i>
FCMPEq	Floating-point compare quad (exception if unordered)	<i>illegal_instruction</i>
FDIVq	Floating-point divide quad	<i>illegal_instruction</i>
FdMULq	Floating-point multiply double to quad	<i>illegal_instruction</i>
FiTOq	Convert integer to quad floating-point	<i>illegal_instruction</i>
FMOVq	Floating-point move quad	<i>illegal_instruction</i>
FMOVqcc	Move quad floating-point register if condition is satisfied	<i>illegal_instruction</i>
FMOVqr	Move quad floating-point register if integer register contents satisfy condition	<i>illegal_instruction</i>
FMULq	Floating-point multiply quad	<i>illegal_instruction</i>
FNEGq	Floating-point negate quad	<i>illegal_instruction</i>
FSQRTq	Floating-point square root quad	<i>illegal_instruction</i>
F(s,d,q)TO(q)	Convert between floating-point formats to quad	<i>illegal_instruction</i>
FQTOI	Convert quad floating point to integer	<i>illegal_instruction</i>
FQTOX	Convert quad floating point to 64-bit integer	<i>illegal_instruction</i>
FSUBq	Floating-point subtract quad	<i>illegal_instruction</i>
FxTOq	Convert 64-bit integer to floating-point	<i>illegal_instruction</i>
IMPDEP1 (not listed in TABLE 5-1)	Implementation-dependent instruction	<i>illegal_instruction</i>
IMPDEP2 (not listed in TABLE 5-1)	Implementation-dependent instruction	<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>
STQF	Store quad floating-point	<i>illegal_instruction</i>
STQFA	Store quad floating-point into alternate space	<i>illegal_instruction</i>

5.2 PREFETCH/PREFETCHA

See the PREFETCH and PREFETCHA instruction descriptions in the Oracle SPARC Architecture 2011 specification for the standard definitions of these instructions. This section describes how SPARC M7 handles PREFETCH instructions.

SPARC M7 interprets the function codes for prefetch variants as follows:

TABLE 5-3 SPARC M7 interpretation of prefetch variants

fcn	Prefetch Variant	Action
0	Weak prefetch for several reads	Prefetch to L1 data cache and L2 cache
1	Weak prefetch for one read	Prefetch to L2 cache
2	Weak prefetch for several writes	Prefetch to L2 cache (exclusive)
3	Weak prefetch for one write	Prefetch to L2 cache (exclusive)

TABLE 5-3 SPARC M7 interpretation of prefetch variants (*Continued*)

fcn	Prefetch Variant	Action
4	Prefetch Page	NOP - no action taken
5 - 15	<i>Reserved</i>	<i>illegal_instruction</i> trap
16	NOP	NOP - no action taken
17	Strong prefetch to nearest unified cache	Prefetch to L2 cache
18 - 19	NOP	NOP - no action taken
20	Strong prefetch for several reads	Prefetch to L1 data cache and L2 cache
21	Strong prefetch for one read	Prefetch to L2 cache
22	Strong prefetch for several writes	Prefetch to L2 cache (exclusive)
23	Strong prefetch for one write	Prefetch to L2 cache (exclusive)
24-31	NOP	NOP - no action taken

Programming Note SPARC M7 does not implement any prefetch functions that prefetch solely to the L3 cache.

Implementation Note On SPARC M7, prefetches can be dropped either at the L1 data cache, the L2 cache, or the L3 cache. Prefetches may be dropped regardless of whether they are strong or weak. Weak prefetches are dropped if they miss the DTLB, whereas strong prefetches are dropped if hardware tablewalk returns an error or is not enabled; otherwise, the following conditions apply to either type. Prefetches are dropped when:

1. The prefetch is to an I/O page, or a page marked as non-cacheable or with side-effects.
2. The miss buffer in the L1 data cache fills beyond a high-water mark (this only applies when more than one thread is unparked).
3. The prefetch is for a data cache miss which is already outstanding.
4. The prefetch is a read prefetch that hits in the L1 cache.
5. The prefetch is a read prefetch to L2 which hits in the L2 cache.
6. The prefetch is a write prefetch which exists in the L2 cache in the exclusive state.
7. The prefetch misses in the L2 cache, and the L2 miss buffer fills beyond a high water mark.
8. The prefetch misses in the L3 cache, and the L3 miss buffer fills beyond a high water mark.

5.3 WRPAUSE

WRPAUSE is a mnemonic for a WRASR to ASR 27, the PAUSE register.

Writing to the PAUSE register suspends a strand for a specified number of nanoseconds. The PAUSE register is write-only; the PAUSE register cannot be read. SPARC M7 implements a 16-bit PAUSE register as described below:

TABLE 5-4 PAUSE Register

Bit	Field	R/W	Description
63:20	—	WO	<i>Reserved.</i>
19:4	pause	WO	Pause value from 0..1048576 nanoseconds
3:0	—	WO	<i>Ignored.</i>

When WRPAUSE is executed, the following sequence of events occurs in SPARC M7:

1. Hardware places the strand in a paused state. Hardware flushes the strand, thereby making the strand inactive and releasing shared resources to the active strands.
2. Hardware checks the value that will be written to the PAUSE register. Hardware updates the strand's PAUSE register with the value of $((\min(2^{16} - 1, (R[rs1] \text{ xor } \text{sim}13)) \gg 4))$ or the value $((\min(2^{16} - 1, (R[rs1] \text{ xor } R[rs2])) \gg 4))$, depending upon the instruction format. If the value written to PAUSE is 0, hardware will pause the strand for a minimum delay. The value placed in the PAUSE register is divided by 16 since each strand's PAUSE register is decremented once every 16 ns. Thus the actual duration of a WRPAUSE ranges between a minimum delay of approximately 10 ns to a maximum of 1048576 ns.
3. Hardware decrements the PAUSE register every 16 ns. The strand remains in the paused state until either:
 - a. The PAUSE register decrements to zero, or
 - b. Any unmasked disrupting trap request, any deferred trap request, an XIR trap request, or a request to change the strand state from Running to Parked is received (See Section 15.1.2.6, *STRAND_RUNNING*, on page 458 for more details on Running and Parked states and transitions). Also see Oracle SPARC Architecture 2015 for more details on what terminates a WRPAUSE. These requests immediately force the PAUSE register to become 0.
4. When the PAUSE register becomes 0, SPARC M7 resumes instruction fetch and execution at the NPC of the WRPAUSE¹.

A masked trap request does not affect the PAUSE register or suspension of the strand.

Any disrupting trap request that is posted after WRPAUSE has updated the PAUSE register and the strand has suspended forward progress does not result in a trap being taken on the WRPAUSE instruction; the trap is taken on a later instruction. This ensures forward progress when the trap handler retries the instruction on which the trap was taken.

Programming Note | WRPAUSE is intended to be used as part of a progressive (exponential) backoff algorithm.

5.4 WRMWAIT

WRMWAIT is a mnemonic for a WRASR to ASR 28, the MWAIT register.

Writing to the MWAIT register suspends a strand up to a specified number of nanoseconds, or until the monitored memory address is modified. The MWAIT register is write-only; the MWAIT register cannot be read. SPARC M7 leverages the PAUSE register to implement the MWAIT register (specifically in respect to suspending a strand for a number of nanoseconds).

¹ Hardware releases the post-sync at the Select stage, enabling subsequent instructions to enter the pipeline.

See Oracle SPARC Architecture 2015 for details of MWAIT functionality, see Section 5.3 for implementation details relative to PAUSE, and see Section 11.5.8 for implementation dependencies.

5.5 Block Load and Store Instructions

See the LDBLOCKF and STBLOCKF instruction descriptions in the Oracle SPARC Architecture 2015 specification for the standard definitions of these instructions.

Block store commits in SPARC M7 do NOT force the data to be written to memory as specified in the Oracle SPARC Architecture 2015 specification. Block store commits are implemented the same as block stores in SPARC M7. As with all stores, block stores and block store commits maintain coherency with all I-caches, but will not flush any modified instructions executing down a pipeline. Flushing those instructions requires the pipeline to execute a FLUSH instruction.

Notes | If LDBLOCKF is used with an `ASI_BLK_COMMIT_{P,S}` and a destination register number `rd` is specified which is not a multiple of 8 (a misaligned `rd`), SPARC M7 generates an *illegal_instruction* exception (impl. dep. #255-U3-Cs10).

| If LDBLOCKF is used with an `ASI_BLK_COMMIT_{P,S}` and a memory address is specified with less than 64-byte alignment, SPARC M7 generates a *DAE_invalid_ASI* exception (impl. dep. #256-U3)

These instructions are used for transferring large blocks of data (more than 256 bytes); for example, `memcpy()` and `memset()`. On SPARC M7, a block load forces a miss in the primary cache and will not allocate a line in the primary cache, but does allocate in L2.

SPARC M7 breaks block load and store instructions into 8 individual "helper" instructions. Each helper is translated as an independent instruction. Thus, it is possible that any individual helper or set of helpers translates to a different memory page from other helpers from the same instruction, if the underlying memory mapping is changed by another process during the execution of the block instruction. Any individual helper or set of helpers may also trap if memory mapping attributes are changed by another process in the midst of a series of helper translations. In the event multiple helpers have exceptions, SPARC M7 commits the helpers in program order from the lowest virtual address to the highest virtual address. Thus, the helper with the lowest virtual address which experiences an exception determines which trap will be taken. SPARC M7 makes no guarantee about the atomicity of address translation for block operations.

Block stores execute differently on SPARC M7 than on prior UltraSPARC processors. On previous processors, such as UltraSPARC T2, UltraSPARC T2+, and SPARC T3, block stores fetched the data from memory prior to updating the line with the store data. On SPARC M7, the processor first establishes the line in the L2 cache and zeroes the data, prior to updating the line with the store source data. The block store is helperized into 8 individual block init stores. The first helper establishes the line in the L2 cache, zeroes the line out, then updates the first 8 bytes of the line with the first 8 bytes of the store source data. The remaining seven helpers collectively update the remaining 56 bytes with the remaining 56 bytes of store source data. As a result, it is possible for another process to see the old data, the new data, or a value of zero while the block store is being executed.

SPARC M7 treats LDBLOCKF as interlocked with respect to following instructions. All later instructions see the effect of the newly loaded values.

STBLOCKF source data registers are interlocked against completion of previous instructions, including block load instructions; STBLOCKF instructions don't commit until all previous instructions commit. Thus STBLOCKF instructions read the most recent value of the floating-point source register(s) when committing to memory. STBLOCKF instructions may or may not initialize the target memory locations to 0 prior to updating them with the source data. Thus another strand may observe these intermediate zero values prior to observing the final source data value.

LDBLOCKF does not follow memory model ordering with respect to stores. In particular, a read-after-write hazard to overlapping addresses is not detected. The side-effect bit associated with the access is ignored (see *Translation Table Entry (TTE)* on page 75). 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 the LDBLOCKF overlaps a previous store and there is no intervening MEMBAR or data reference, the LDBLOCKF may return data from before or after the store.

STBLOCKF instructions do not conform to TSO store-store ordering with respect to older non-overlapping stores. A subsequent load to the same address as a STBLOCKF may not read the results of the STBLOCKF. The side-effects bit associated with the access is ignored. If ordering with respect to later loads is important then there must be an intervening MEMBAR instruction. If the STBLOCKF overlaps a later load and there is no intervening MEMBAR #StoreLoad instruction, the result of the load is undefined.

Compatibility Notes	Block load and store operations do not obey the ordering restrictions of the currently selected processor memory model (TSO, PSO, or RMO). In general, explicit MEMBAR instructions are required to order block memory operations among themselves or with respect to normal loads and stores. In addition, block operations do not generally conform to dependence order on the issuing virtual processor; that is, no read-after-write or write-after-read checking occurs between block loads and stores. Explicit MEMBARs are required to enforce dependence ordering between block operations that reference the same address TABLE 5-5 describes the synchronization primitives required in SPARC M7, if any, to guarantee TSO ordering between various sequences of memory reference operations. The first column contains the reference type of the first or earlier instruction; the second column contains the reference type of the second or the later instruction. SPARC M7 orders loads and block loads against all subsequent instructions.
----------------------------	---

TABLE 5-5 SPARC M7 Synchronization Requirements for Memory Reference Operations

First reference	Second reference	Synchronization Required
Load	Load	—
	Block load	MEMBAR #LoadLoad
	Store	—
Block load	Block store	—
	Load	—
	Block load	MEMBAR #LoadLoad
Store	Store	—
	Block store	—
	Load	—
Block store	Block load	MEMBAR #StoreLoad or #Sync
	Store	—
	Block store	MEMBAR #StoreStore or stronger, if to non-overlapping addresses
Block store	Load	MEMBAR #StoreLoad or #Sync
	Block load	MEMBAR #StoreLoad or #Sync
	Store	MEMBAR #StoreStore or stronger, if to non-overlapping addresses
	Block store	MEMBAR #StoreStore or stronger, if to non-overlapping addresses

5.6 Block Initializing Store ASIs

Instruction	imm_asi	ASI Value	Operation
ST[B,H,W,TW,X]A, STFA, STDFA	ASI_ST_BLKINIT_AS_IF_USER_PRIMARY (ASI_STBI_AIUP)	22 ₁₆	64-byte block initializing store to primary address space, user privilege
	ASI_ST_BLKINIT_AS_IF_USER_SECONDARY (ASI_STBI_AIUS)	23 ₁₆	64-byte block initializing store to secondary address space, user privilege
	ASI_ST_BLKINIT_REAL (ASI_STBI_R)	26 ₁₆	64-byte block initializing store to real address
	ASI_ST_BLKINIT_NUCLEUS (ASI_STBI_N)	27 ₁₆	64-byte block initializing store to nucleus address space
	ASI_ST_BLKINIT_AS_IF_USER_PRIMARY_LITTLE (ASI_STBI_AIUPL)	2A ₁₆	64-byte block initializing store to primary address space, user privilege, little-endian
	ASI_ST_BLKINIT_AS_IF_USER_SECONDARY_LITTLE (ASI_STBI_AIUS_L)	2B ₁₆	64-byte block initializing store to secondary address space, user privilege, little-endian
	ASI_ST_BLKINIT_REAL_LITTLE (ASI_STBI_RL)	2E ₁₆	64-byte block initializing store to real address, little-endian
ST[B,H,W,TW,X]A, STFA, STDFA	ASI_ST_BLKINIT_NUCLEUS_LITTLE (ASI_STBI_NL)	2F ₁₆	64-byte block initializing store to nucleus address space, little-endian
	ASI_ST_BLKINIT_PRIMARY (ASI_STBI_P)	E2 ₁₆	64-byte block initializing store to primary address space
	ASI_ST_BLKINIT_SECONDARY (ASI_STBI_S)	E3 ₁₆	64-byte block initializing store to secondary address space
	ASI_ST_BLKINIT_PRIMARY_LITTLE (ASI_STBI_PL)	EA ₁₆	64-byte block initializing store to primary address space, little-endian
	ASI_ST_BLKINIT_SECONDARY_LITTLE (ASI_STBI_SL)	EB ₁₆	64-byte block initializing store to secondary address space, little-endian
	ASI_ST_BLKINIT_MRU_PRIMARY (ASI_STBIMRU_P)	F2 ₁₆	64-byte block initializing store to primary address space, install as MRU in L2 cache
	ASI_ST_BLKINIT_MRU_SECONDARY (ASI_STBIMRU_S)	F3 ₁₆	64-byte block initializing store to secondary address space, install as MRU in L2 cache
	ASI_ST_BLKINIT_MRU_PRIMARY_LITTLE (ASI_STBIMRU_PL)	FA ₁₆	64-byte block initializing store to primary address space, little-endian, install as MRU in L2 cache
	ASI_ST_BLKINIT_MRU_SECONDARY_LITTLE (ASI_STBIMRU_SL)	FB ₁₆	64-byte block initializing store to secondary address space, little-endian, install as MRU in L2 cache

Description Block initializing store instructions are selected by using one of the block initializing store ASIs with integer or floating-point store 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 and floating-point stores of all sizes (to alternate space) are allowed to use these ASIs.

All stores to these ASIs operate under relaxed memory ordering (RMO). To ensure ordering with respect to subsequent stores and loads, software must follow a sequence of these stores with a MEMBAR #StoreStore or #StoreLoad, respectively. To ensure ordering with respect to prior stores, software must precede these stores with a MEMBAR #StoreStore.

Stores to these ASIs where the least-significant 6 bits of the address are non-zero (that is, not the first word in the L2 cache line) behave the same as a normal RMO store. A store to these ASIs where the least-significant 6 bits are zero will load a 64 byte line in the L2 cache with all zeros, and then update that line with the new store data. The zeroing of the line and the storing of the new data are not atomic, therefore while a block-initializing store is being performed, another strand may observe any of the following: (1) the old data value, (2) zero, or (3) the new data value. When the operation is complete, only the new data value will be seen. This special store will make sure the 64B lines maintain coherency when they are loaded into the L2 cache, but will not fetch the line from memory (initializing it with zeros instead), except as noted above. Stores using these ASIs to a noncacheable address behave the same as a normal store.

The ASIs F2₁₆, F3₁₆, FA₁₆, and FB₁₆ establish the line in the L2 cache as recently-used, thereby helping to ensure they are not replaced shortly after being established. (The naming of the ASIs refers to MRU in this case, although the L2 does not use a true LRU policy.) This can aid in cases where the newly-established line is expected to be referenced in the near future from a process running on the same physical core. The remaining block initializing store ASIs establish the line in the L2 cache as not-used, increasing the likelihood of the line being replaced. This is useful when data is not expected to be used in the near future as it reduces the amount of cached data displaced by the copy routine.

One way the MRU and LRU variants can be used is in a copy routine. The first block initializing stores to the line can be of the MRU variety. This will reduce the chance that the line will be lost before all stores to the line are complete. The final block initializing store to the line would then be of the LRU variety, causing the line to then be favored for replacement and reducing the cache pollution associated with the copy operation.

Note | These instructions are used for transferring large blocks of data (more than 256 bytes); for example, `memcpy()` and `memset()`. On SPARC M7, a twin load forces a miss in the primary cache and will not allocate a line in the primary cache, but does allocate in L2.

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, guaranteeing that all initial zeros in a cache line are overwritten with copy data. This pseudocode may not be optimal for SPARC M7; it is provided as an example only.

```
%l0 ← [A]
%l1 ← [B]
prefetch [%l0]
for (i = 0; i < N-4; i++) {
    if ((i mod 4) ≠ 0) {
        prefetch [%l0+64]
    }
    ldtxa [%l0] #ASI_TWIX_P, %l2
    add %l0, 16, %l0
    stxa %l2, [%l1] #ASI_ST_BLKINIT_PRIMARY
    add %l1, 8, %l1
    stxa %l3, [%l1] #ASI_ST_BLKINIT_PRIMARY
    add %l1, 8, %l1
}
for (i = 0; i < 4; i++) {
    ldtxa [%l0] #ASI_TWIX_P, %l2
    add %l0, 16, %l0
    stx %l2, [%l1]
    stx %l3,d [%l1+8]
    add %l1, 16, %l1
}
membar #Sync
```

Programming Notes | The Block Initializing Store ASIs are of Class "N" and are only allowed in dynamically linked, platform-specific, OS-enabled libraries.

Traps

6.1 Trap Levels

Only SPARC M7 specific behavior is described in this chapter; refer to Oracle SPARC Architecture 2015 for more detail on trap handling.

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

6.2 Trap Behavior

TABLE 6-1 specifies the codes used in the tables below.

TABLE 6-1 Table Codes

Code	Meaning
H	Trap is taken in Hyperprivileged mode
P	Trap is taken via the Privileged trap table, in Privileged mode ($PSTATE.priv = 1$)
-x-	Not possible. Hardware cannot generate this trap in the indicated running mode. For example, all privileged instructions can be executed in privileged mode, therefore a <i>privileged_opcode</i> trap cannot occur in privileged mode.
—	This trap can only legitimately be generated by hyperprivileged software, not by the CPU hardware. So, for the purposes of sun4v, the trap vector has to be correct, but for a hardware CPU implementation these trap types are not generated by the hardware, therefore the resultant running mode is irrelevant.

6.3 Trap Masking

TABLE 6-2 specifies the codes used in TABLE 6-2.

TABLE 6-2 Codes

Code	Meaning
(nm)	Never Masked — when the condition occurs in this running mode, it is never masked out and the trap is always taken.
(ie)	When the outstanding disrupting trap condition occurs in this privilege mode, it may be conditioned (masked out) by PSTATE.ie = 0 (but remains pending).
PIL	Masked by PSTATE.ie and PIL
-x-	This trap can only legitimately be generated by hyperprivileged software, not by the CPU hardware. So, for the purposes of sun4v, the trap vector has to be correct, but for a hardware CPU implementation these trap types are not generated by the hardware, therefore the resultant running mode is irrelevant.

Interrupt Handling

The chapter describes the hardware interrupt delivery mechanism for the SPARC M7 chip.

Hyperprivileged code notifies privileged code about *sw_recoverable_error* traps through the *cpu_mondo*, *dev_mondo*, and *resumable_error* traps as described in *Interrupt Queue Registers* on page 47. Software interrupts are delivered to each virtual processor using the *interrupt_level_n* traps. Software interrupts are described in the Oracle SPARC Architecture 2015 specification.

7.1 Interrupt Flow

7.1.1 Sources

The processor SOFTINT_SET register can be written from sources external to the processor core.

7.1.2 States

7.2 CPU Interrupt Registers

7.2.1 Interrupt Queue Registers

Each virtual processor has eight ASI_QUEUE registers at ASI = 25₁₆, VA{63:0} = 3C0₁₆-3F8₁₆ that are used for communicating interrupts to the operating system. These registers contain the head and tail pointers for four supervisor interrupt queues: *cpu_mondo*, *dev_mondo*, *resumable_error*, *nonresumable_error*. The tail registers are read-only by supervisor, and read/write by hypervisor. Writes to the tail registers by the supervisor generate a *DAE_invalid_ASI* trap. The head registers are read/write by both supervisor and hypervisor.

Whenever the CPU_MONDO_HEAD register does not equal the CPU_MONDO_TAIL register, a *cpu_mondo* trap is generated. Whenever the DEV_MONDO_HEAD register does not equal the DEV_MONDO_TAIL register, a *dev_mondo* trap is generated. Whenever the RESUMABLE_ERROR_HEAD register does not equal the RESUMABLE_ERROR_TAIL register, a *resumable_error* trap is generated. Unlike the other queue register pairs, the *nonresumable_error* trap

is *not* automatically generated whenever the `NONRESUMABLE_ERROR_HEAD` register does not equal the `NONRESUMABLE_ERROR_TAIL` register; instead, the hypervisor will need to generate the `nonresumable_error` trap.

TABLE 7-1 through TABLE 7-8 define the format of the eight `ASI_QUEUE` registers.

TABLE 7-1 CPU Mondo Head Pointer – `ASI_QUEUE_CPU_MONDO_HEAD` (ASI 25₁₆, VA 3C0₁₆)

Bit	Field	Initial Value	Access	Description
63:31	—	0	RO	<i>Reserved</i>
30:6	head	X	RW	Head pointer for CPU mondo interrupt queue.
5:0	—	0	RO	<i>Reserved</i>

TABLE 7-2 CPU Mondo Tail Pointer – `ASI_QUEUE_CPU_MONDO_TAIL` (ASI 25₁₆, VA 3C8₁₆)

Bit	Field	Initial Value	Access	Description
63:31	—	0	RO	<i>Reserved</i>
30:6	tail	X	RW	Tail pointer for CPU mondo interrupt queue.
5:0	—	0	RO	<i>Reserved</i>

TABLE 7-3 Device Mondo Head Pointer – `ASI_QUEUE_DEV_MONDO_HEAD` (ASI 25₁₆, VA 3D0₁₆)

Bit	Field	Initial Value	Access	Description
63:31	—	0	RO	<i>Reserved</i>
30:6	head	X	RW	Head pointer for device mondo interrupt queue.
5:0	—	0	RO	<i>Reserved</i>

TABLE 7-4 Device Mondo Tail Pointer – `ASI_QUEUE_DEV_MONDO_TAIL` (ASI 25₁₆, VA 3D8₁₆)

Bit	Field	Initial Value	Access	Description
63:31	—	0	RO	<i>Reserved</i>
30:6	tail	X	RW	Tail pointer for device mondo interrupt queue.
5:0	—	0	RO	<i>Reserved</i>

TABLE 7-5 Resumable Error Head Pointer – `ASI_QUEUE_RESUMABLE_HEAD` (ASI 25₁₆, VA 3E0₁₆)

Bit	Field	Initial Value	Access	Description
63:31	—	0	RO	<i>Reserved</i> .
30:6	head	X	RW	Head pointer for resumable error queue.
5:0	—	0	RO	<i>Reserved</i>

TABLE 7-6 Resumable Error Tail Pointer – `ASI_QUEUE_RESUMABLE_TAIL` (ASI 25₁₆, VA 3E8₁₆)

Bit	Field	Initial Value	Access	Description
63:31	—	0	RO	<i>Reserved</i>
30:6	tail	X	RW	Tail pointer for resumable error queue.
5:0	—	0	RO	<i>Reserved</i>

TABLE 7-7 Nonresumable Error Head Pointer – ASI_QUEUE_NONRESUMABLE_HEAD (ASI 25₁₆, VA 3F0₁₆)

Bit	Field	Initial Value	Access	Description
63:31	—	0	RO	<i>Reserved</i>
30:6	head	X	RW	Head pointer for nonresumable error queue.
5:0	—	0	RO	<i>Reserved</i>

TABLE 7-8 Nonresumable Error Tail Pointer – ASI_QUEUE_NONRESUMABLE_TAIL (ASI 25₁₆, VA 3F8₁₆)

Bit	Field	Initial Value	Access	Description
63:31	—	0	RO	<i>Reserved</i>
30:6	tail	X	RW	Tail pointer for nonresumable error queue.
5:0	—	0	RO	<i>Reserved</i>

Memory Models

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. SPARC M7 supports only TSO, with the exception that certain ASI accesses (such as block loads and stores) may operate under RMO.

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 SPARC V9 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. It is unaffected by normal traps. SPARC M7 ignores the value set in this field and always operates under TSO.

A memory location is identified by an 8-bit address space identifier (ASI) and a 64-bit virtual address. The 8-bit ASI may be obtained from a ASI register or included in a memory access instruction. The ASI is used to distinguish between and provide an attribute for different 64-bit address spaces. For example, the ASI is used by the SPARC M7 MMU to control access to implementation-dependent control and data registers and for access protection. Attempts by nonprivileged software (PSTATE.priv = 0) to access restricted ASIs (ASI{7} = 0) cause a *privileged_action* trap.

Real memory spaces can be accessed without side effects. For example, a read 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.

8.1 Supported Memory Models

The following sections contain brief descriptions of the two memory models supported by SPARC M7. These definitions are for general illustration. Detailed definitions of these models can be found in *The SPARC Architecture Manual-Version 9*. The definitions in the following sections apply to system behavior as seen by the programmer.

Notes Stores to SPARC M7 internal ASIs, block loads, and block stores and block initializing stores are outside the memory model; that is, they need MEMBARs to control ordering.

Atomic load-stores are treated as both a load and a store and can only be applied to cacheable address spaces.

8.1.1 TSO

SPARC M7 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.
- Accesses to I/O space are all strongly ordered with respect to each other.
- 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.

8.1.2 RMO

SPARC M7 implements the following programmer-visible properties for special ASI accesses that operate under Relaxed Memory Order (RMO) mode:

- 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 noncacheable accesses.


Address Spaces and ASIs

9.1 Address Spaces

SPARC M7 supports a 54-bit virtual address space.

9.1.1 54-bit Virtual and Real Address Spaces

SPARC M7 supports a 54-bit subset of the full 64-bit virtual and real address spaces. 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 (real) address space. Virtual (real) addresses between $0020\ 0000\ 0000\ 0000_{16}$ and $FFDF\ FFFF\ FFFF\ FFFF_{16}$ inclusive lie within a “VA hole” (“RA hole”), are termed “out-of-range”¹, and are illegal. Prior UltraSPARC implementations introduced the additional restriction on software to not use pages within 4 Gbytes of the VA (RA) hole as instruction pages to avoid problems with prefetching into the VA (RA) hole. SPARC M7 implements a hardware check for instruction fetching near the VA (RA) hole and generates a trap when instructions are executed from a location in the address range $001F\ FFFF\ FFFF\ FFC0_{16}$ to $001F\ FFFF\ FFFF\ FFFF_{16}$, inclusive, if $PSTATE.am = 0$. However, even though SPARC M7 provides this hardware checking, it is still recommended that software should *not* use the 8-Kbyte page before the VA (RA) hole for instructions. Address translation and MMU related descriptions can be found in *Translation* on page 81.



Note (1): Use of this region restricted to data only.

FIGURE 9-1 SPARC M7’s 52-bit Virtual and Real Address Spaces, With Hole

¹. Another way to view an out-of-range address is as any address where bits {63:52} are not all equal to bit {51}.

Throughout this document, when virtual (real) address fields are specified as 64-bit quantities, they are assumed to be sign-extended based on VA{53} (RA{53}).

A number of state registers are affected by the reduced virtual and real address spaces. The PC register is 54 bits, sign-extended to 64-bits on read accesses. The TBA, TPC, and TNPC registers are 54 bits wide. No checks are done when these registers are written by software. It is the responsibility of privileged software to properly update these registers.

If the target virtual (real) address of a JMPL, RETURN, branch, or CALL instruction is an out-of-range address and PSTATE.am = 0, a trap is generated with TPC equal to the address of the JMPL, RETURN, branch, or CALL instruction.

An out-of-range virtual (real) address during a data access results in a trap if PSTATE.am = 0.

9.2 Alternate Address Spaces

The table below summarizes the ASI usage in SPARC M7. The Section/Page column contains a reference to the detailed explanation of the ASI (the page number refers to this chapter). For internal ASIs, the legal VAs are listed (or the field contains "Any" if all VAs are legal). An access outside the legal VA range generates a *DAE_invalid_asi* trap.

Notes All internal, nontranslating ASIs in SPARC M7 can only be accessed using LDXA and STXA.

ASIs 80₁₆–FF₁₆ are unrestricted (access allowed in all modes -- nonprivileged, privileged). ASIs 00₁₆–2F₁₆ are restricted to privileged and hyperprivileged modes.

TABLE 9-1 SPARC M7 ASI Usage (1 of 7)

ASI	ASI Name	R/W	VA	Copy per Strand	Description	Section/Page
00 ₁₆ –01 ₁₆			Any	—	<i>DAE_invalid_asi</i>	
03 ₁₆			Any	—	<i>DAE_invalid_asi</i>	
04 ₁₆	ASI_NUCLEUS	RW	Any	—	Implicit address space, nucleus context, TL > 0	(See UA-2015)
06 ₁₆ –0B ₁₆			Any	—	<i>DAE_invalid_asi</i>	
0C ₁₆	ASI_NUCLEUS_LITTLE	RW	Any	—	Implicit address space, nucleus context, TL > 0 (LE)	(See UA-2015)
0D ₁₆ –0F ₁₆			Any	—	<i>DAE_invalid_asi</i>	
10 ₁₆	ASI_AS_IF_USER_PRIMARY	RW	Any	—	Primary address space, user privilege	(See UA-2015)
11 ₁₆	ASI_AS_IF_USER_SECONDARY	RW	Any	—	Secondary address space, user privilege	(See UA-2015)
12 ₁₆	ASI_MONITOR_AS_IF_USER_PRIMARY	RO	Any	—	Primary address space, user privilege, set load monitor	(See UA-2015 and Section 11.5.8)
13 ₁₆	ASI_MONITOR_AS_IF_USER_SECONDARY	RO	Any	—	Secondary address space, user privilege, set load monitor	(See UA-2015 and Section 11.5.8)
14 ₁₆	ASI_REAL	RW	Any	—	Real address (normally used as cacheable)	9.2.1

TABLE 9-1 SPARC M7 ASI Usage (2 of 7)

ASI	ASI Name	R/W	VA	Copy per Strand	Description	Section/Page
15 ₁₆	ASI_REAL_IO	RW	Any	—	Real address (normally used as noncacheable, with side effect)	9.2.1
16 ₁₆	ASI_BLOCK_AS_IF_USER_PRIMARY	RW	Any	—	64-byte block load/store, primary address space, user privilege	5.5
17 ₁₆	ASI_BLOCK_AS_IF_USER_SECONDARY	RW	Any	—	64-byte block load/store, secondary address space, user privilege	5.5
18 ₁₆	ASI_AS_IF_USER_PRIMARY_LITTLE	RW	Any	—	Primary address space, user privilege (LE)	(See UA-2015)
19 ₁₆	ASI_AS_IF_USER_SECONDARY_LITTLE	RW	Any	—	Secondary address space, user privilege (LE)	(See UA-2015)
1A ₁₆ –1B ₁₆			Any	—	<i>DAE_invalid_asi</i>	
1C ₁₆	ASI_REAL_LITTLE	RW	Any	—	Real address (normally used as cacheable) (LE)	9.2.1
1D ₁₆	ASI_REAL_IO_LITTLE	RW	Any	—	Real address (normally used as noncacheable, with side effect) (LE)	9.2.1
1E ₁₆	ASI_BLOCK_AS_IF_USER_PRIMARY_LITTLE	RW	Any	—	64-byte block load/store, primary address space, user privilege (LE)	5.5
1F ₁₆	ASI_BLOCK_AS_IF_USER_SECONDARY_LITTLE	RW	Any	—	64-byte block load/store, secondary address space, user privilege (LE)	5.5
20 ₁₆	ASI_SCRATCHPAD	RW	0 ₁₆ –18 ₁₆	Y	Scratchpad registers	9.2.2
20 ₁₆	ASI_SCRATCHPAD		20 ₁₆ –28 ₁₆	—	<i>DAE_invalid_asi</i>	
20 ₁₆	ASI_SCRATCHPAD	RW	30 ₁₆ –38 ₁₆	Y	Scratchpad registers	9.2.2
21 ₁₆	ASI_MMU	RW	8 ₁₆	Y	I/DMMU Primary Context register 0	13.7.2
21 ₁₆	ASI_MMU	RW	10 ₁₆	Y	DMMU Secondary Context register 0	13.7.2
21 ₁₆	ASI_MMU	RW	28 ₁₆	Y	I/DMMU Primary Context register 0 (no Primary Context register 1 update)	13.7.2
21 ₁₆	ASI_MMU	RW	30 ₁₆	Y	DMMU Secondary Context register 0 (no Secondary Context register 1 update)	13.7.2
21 ₁₆	ASI_MMU	RW	108 ₁₆	Y	I/DMMU Primary Context register 1	13.7.2
21 ₁₆	ASI_MMU	RW	110 ₁₆	Y	DMMU Secondary Context register 1	13.7.2
22 ₁₆	ASI_TWIX_AIUP, ASI_STBI_AIUP	RW	Any	—	Load: 128-bit atomic load twin extended word, primary address space, user privilege Store: Block initializing store, primary address space, user privilege	5.7.4

TABLE 9-1 SPARC M7 ASI Usage (3 of 7)

ASI	ASI Name	R/W	VA	Copy per Strand	Description	Section/Page
23 ₁₆	ASI_TWIX_AIUS, ASI_STBI_AIUS	RW	Any	—	Load: 128-bit atomic load (See UA-2015) twin extended word, secondary address space, user privilege Store: Block initializing store	
24 ₁₆			Any	—	<i>DAE_invalid_asi</i>	
25 ₁₆	ASI_QUEUE	RW	3C0 ₁₆	Y	CPU Mondo Queue head pointer	7.2.1
25 ₁₆	ASI_QUEUE	RW (hyperpriv) RO (priv)	3C8	Y	CPU Mondo Queue tail pointer	7.2.1
25 ₁₆	ASI_QUEUE	RW	3D0 ₁₆	Y	Device Mondo Queue head pointer	7.2.1
25 ₁₆	ASI_QUEUE	RW (hyperpriv) RO (priv)	3D8 ₁₆	Y	Device Mondo Queue tail pointer	7.2.1
25 ₁₆	ASI_QUEUE	RW	3E0 ₁₆	Y	Resumable Error Queue head pointer	7.2.1
25 ₁₆	ASI_QUEUE	RW (hyperpriv) RO (priv)	3E8 ₁₆	Y	Resumable Error Queue tail pointer	7.2.1
25 ₁₆	ASI_QUEUE	RW	3F0 ₁₆	Y	Nonresumable Error Queue head pointer	7.2.1
25 ₁₆	ASI_QUEUE	RW (hyper-priv) RO (priv)	3F8 ₁₆	Y	Nonresumable Error Queue tail pointer	7.2.1
26 ₁₆	ASI_TWIX_REAL, ASI_STBI_REAL	RW	Any	—	Load: 128-bit atomic LDDA, real address Store: Block initializing store, real address	(See UA-2015)
27 ₁₆	ASI_TWIX_NUCLEUS, ASI_STBI_N	RW	Any	—	Load: 128-bit atomic load (See UA-2015) twin extended word from nucleus context Store: Block initializing store from nucleus context	
28 ₁₆ -29 ₁₆			Any	—	<i>DAE_invalid_asi</i>	
2A ₁₆	ASI_TWIX_AIUPL, ASI_STBI_AIUPL	RW	Any	—	Load: 128-bit atomic load (See UA-2015) twin extended word, primary address space, user privilege, little endian Store: Block initializing store, primary address space, user privilege, little endian	

TABLE 9-1 SPARC M7 ASI Usage (4 of 7)

ASI	ASI Name	R/W	VA	Copy per Strand	Description	Section/Page
2B ₁₆	ASI_TWIX_AIUSL, ASI_STBI_AIUSL	RW	Any	—	Load: 128-bit atomic load ((See UA-2015) twin extended word, secondary address space, user privilege, little endian Store: Block initializing store, secondary address space, user privilege, little endian	
2C ₁₆			Any	—	<i>DAE_invalid_asi</i>	
2D ₁₆			Any	—	<i>DAE_invalid_asi</i>	
2E ₁₆	ASI_TWIX_REAL_LITTLE, ASI_STBI_REAL_LITTLE	RW	Any	—	Load: 128-bit atomic LDDA, real address (LE) Store: Block initializing store, real address (LE)	(See UA-2015)
2F ₁₆	ASI_TWIX_NL, ASI_STBI_NL	RW	Any	—	Load: 128-bit atomic load from nucleus context, little endian Store: Block initializing store from nucleus context, little endian	(See UA-2015)
80 ₁₆	ASI_PRIMARY	RW	Any	—	Implicit primary address space	(See UA-2015)
81 ₁₆	ASI_SECONDARY	RW	Any	—	Implicit secondary address space	(See UA-2015)
82 ₁₆	ASI_PRIMARY_NO_FAULT	RO	Any	—	Primary address space, no fault	(See UA-2015)
83 ₁₆	ASI_SECONDARY_NO_FAULT	RO	Any	—	Secondary address space, no fault	(See UA-2015)
84 ₁₆	ASI_MONITOR_PRIMARY	RO	Any	—	Primary address space, set load monitor	(See UA-2015 and Section 11.5.8)
85 ₁₆	ASI_MONITOR_SECONDARY	RO	Any	—	Secondary address space, set load monitor	(See UA-2015 and Section 11.5.8)
86 ₁₆ –87 ₁₆			Any	—	<i>DAE_invalid_asi</i>	
88 ₁₆	ASI_PRIMARY_LITTLE	RW	Any	—	Implicit primary address space (LE)	(See UA-2015)
89 ₁₆	ASI_SECONDARY_LITTLE	RW	Any	—	Implicit secondary address space (LE)	((See UA-2015)
8A ₁₆	ASI_PRIMARY_NO_FAULT_LITTLE	RO	Any	—	Primary address space, no fault (LE)	(See UA-2015)
8B ₁₆	ASI_SECONDARY_NO_FAULT_LITTLE	RO	Any	—	Secondary address space, no fault (LE)	(See UA-2015)
8C ₁₆ –8F ₁₆			Any	—	<i>DAE_invalid_asi</i>	
91 ₁₆			Any	—	<i>DAE_invalid_asi</i>	
93 ₁₆ –AF ₁₆			Any	—	<i>DAE_invalid_asi</i>	
B0 ₁₆	ASI_PIC	RW	0 ₁₆	Y	Performance Instrumentation Counter 0	10.3

TABLE 9-1 SPARC M7 ASI Usage (5 of 7)

ASI	ASI Name	R/W	VA	Copy per Strand	Description	Section/Page
B0 ₁₆	ASI_PIC	RW	8 ₁₆	Y	Performance Instrumentation Counter 1	10.3
B0 ₁₆	ASI_PIC	RW	10 ₁₆	Y	Performance Instrumentation Counter 2	10.3
B0 ₁₆	ASI_PIC	RW	18 ₁₆	Y	Performance Instrumentation Counter 3	10.3
B1 ₁₆ –BF ₁₆			Any	—	<i>DAE_invalid_asi</i>	
C0 ₁₆	ASI_PST8_P	WO	Any	—	Eight 8-bit conditional stores, primary address	(See UA-2015)
C1 ₁₆	ASI_PST8_S	WO	Any	—	Eight 8-bit conditional stores, secondary address	(See UA-2015)
C2 ₁₆	ASI_PST16_P	WO	Any	—	Four 16-bit conditional stores, primary address	(See UA-2015)
C3 ₁₆	ASI_PST16_S	WO	Any	—	Four 16-bit conditional stores, secondary address	(See UA-2015)
C4 ₁₆	ASI_PST32_P	WO	Any	—	Two 32-bit conditional stores, primary address	(See UA-2015)
C5 ₁₆	ASI_PST32_S	WO	Any	—	Two 32-bit conditional stores, secondary address	(See UA-2015)
C6 ₁₆ –C7 ₁₆			Any	—	<i>DAE_invalid_asi</i>	
C8 ₁₆	ASI_PST8_PL	WO	Any	—	Eight 8-bit conditional stores, primary address, little endian	((See UA-2015)
C9 ₁₆	ASI_PST8_SL	WO	Any	—	Eight 8-bit conditional stores, secondary address, little endian	(See UA-2015)
CA ₁₆	ASI_PST16_PL	WO	Any	—	Four 16-bit conditional stores, primary address, little endian	(See UA-2015)
CB ₁₆	ASI_PST16_SL	WO	Any	—	Four 16-bit conditional stores, secondary address, little endian	(See UA-2015)
CC ₁₆	ASI_PST32_PL	WO	Any	—	Two 32-bit conditional stores, primary address, little endian	(See UA-2015)
CD ₁₆	ASI_PST32_SL	WO	Any	—	Two 32-bit conditional stores, secondary address, little endian	(See UA-2015)
CE ₁₆ –CF ₁₆			Any	—	<i>DAE_invalid_asi</i>	
D0 ₁₆	ASI_FL8_P	RW	Any	—	8-bit load/store, primary address	(See UA-2015)
D1 ₁₆	ASI_FL8_S	RW	Any	—	8-bit load/store, secondary address	(See UA-2015)
D2 ₁₆	ASI_FL16_P	RW	Any	—	16-bit load/store, primary address	(See UA-2015)
D3 ₁₆	ASI_FL16_S	RW	Any	—	16-bit load/store, secondary address	(See UA-2015)
D4 ₁₆ –D7 ₁₆			Any	—	<i>DAE_invalid_asi</i>	
D8 ₁₆	ASI_FL8_PL	RW	Any	—	8-bit load/store, primary address, little endian	(See UA-2015)

TABLE 9-1 SPARC M7 ASI Usage (6 of 7)

ASI	ASI Name	R/W	VA	Copy per Strand	Description	Section/Page
D9 ₁₆	ASI_FL8_SL	RW	Any	—	8-bit load/store, secondary address, little endian	(See UA-2015)
DA ₁₆	ASI_FL16_PL	RW	Any	—	16-bit load/store, primary address, little endian	(See UA-2015)
DB ₁₆	ASI_FL16_SL	RW	Any	—	16-bit load/store, secondary address, little endian	(See UA-2015)
DC ₁₆ –DF ₁₆			Any	—	<i>DAE_invalid_asi</i>	
E0 ₁₆	ASI_BLK_COMMIT_PRIMARY	WO	Any	—	64-byte block commit store, primary address	5.5
E1 ₁₆	ASI_BLK_COMMIT_SECONDARY	WO	Any	—	64-byte block commit store, secondary address	5.5
E2 ₁₆	ASI_TWIX_P, ASI_STBI_P	RW	Any	—	Load: 128-bit atomic load twin extended word, primary address space Store: Block initializing store, primary address space	(See UA-2015)
E3 ₁₆	ASI_TWIX_S, ASI_STBI_S	RW	Any	—	Load: 128-bit atomic load twin extended word, secondary address space Store: Block initializing store, secondary address space	(See UA-2015)
E4 ₁₆ –E9 ₁₆			Any	—	<i>DAE_invalid_asi</i>	
EA ₁₆	ASI_TWIX_PL, ASI_STBI_PL	RW	Any	—	Load: 128-bit atomic load twin extended word, primary address space, little endian Store: Block initializing store, primary address space, little endian	(See UA-2015)
EB ₁₆	ASI_TWIX_SL, ASI_STBI_SL	RW	Any	—	Load: 128-bit atomic load twin extended word, secondary address space, little endian Store: Block initializing store, secondary address space, little endian	(See UA-2015)
EC ₁₆ –EF ₁₆			Any	—	<i>DAE_invalid_asi</i>	
F0 ₁₆	ASI_BLK_P	RW	Any	—	64-byte block load/store, primary address	5.5
F1 ₁₆	ASI_BLK_S	RW	Any	—	64-byte block load/store, secondary address	5.5
F2 ₁₆	ASI_STBIMRU_PRIMARY	WO	Any		Block initializing store to primary, install as MRU in L2 cache	5.6
F3 ₁₆	ASI_STBIMRU_SECONDARY	WO	Any		Block initializing store to secondary, install as MRU in L2 cache	5.6
F4 ₁₆ –F7 ₁₆			Any	—	<i>DAE_invalid_asi</i>	

TABLE 9-1 SPARC M7 ASI Usage (7 of 7)

ASI	ASI Name	R/W	VA	Copy per Strand	Description	Section/Page
F8 ₁₆	ASI_BLK_PL	RW	Any	—	64-byte block load/store, 5.5 primary address (LE)	
F9 ₁₆	ASI_BLK_SL	RW	Any	—	64-byte block load/store, 5.5 secondary address (LE)	
FA ₁₆	ASI_STBIMRU_PRIMARY_LITTLE	WO	Any		Block initializing store to 5.6 primary little-endian, install as MRU in L2 cache	
FB ₁₆	ASI_STBIMRU_SECONDARY_LITTLE	WO	Any		Block initializing store to 5.6 secondary little-endian, install as MRU in L2 cache	
FC ₁₆ –FF ₁₆			Any	—	DAE_invalid_asi	

9.2.1 ASI_REAL, ASI_REAL_LITTLE, ASI_REAL_IO, and ASI_REAL_IO_LITTLE (ASIs 14₁₆, 1C₁₆, 15₁₆, 1D₁₆)

These ASIs are used to bypass the VA-to-RA translation. For these ASIs, the real address is set equal to the truncated virtual address (that is, RA{51:0} ← VA{51:0}), and the attributes used are those present in the matching TTE. The hypervisor will normally set the TTE attributes for ASI_REAL and ASI_REAL_LITTLE to cacheable (cp = 1) and for ASI_REAL_IO and ASI_REAL_IO_LITTLE to noncacheable, with side effect (cp = 0, e = 1). The hardware, however, does not require this, i.e. it allows an ASI_REAL/ASI_REAL_LITTLE to be issued to a noncacheable address (PA{49} = 1) or an ASI_REAL_IO/ASI_REAL_IO_LITTLE to be issued to a cacheable address (PA{49} = 0); no error is flagged in this case.

Future Implementation Note | Future implementations should explore generating an exception for the case of ASI_REAL_IO or ASI_REAL_IO_LITTLE used with TTE.cp = 1.

9.2.2 ASI_SCRATCHPAD (ASI 20₁₆, VA 0₁₆–18₁₆, 30₁₆–38₁₆)

Each virtual processor has a set of privileged ASI_SCRATCHPAD registers at ASI 20₁₆ with VA{63:0} = 0₁₆–18₁₆, 30₁₆–38₁₆. These registers are for scratchpad use by privileged software.

M7 Implementation Note | Accesses to VA 20₁₆ and 28₁₆ are much slower than to the other six scratchpad registers.

Performance Instrumentation

10.1 Introduction

As in previous UltraSPARC CMT processors, SPARC M7 supports monitoring processor performance by virtue of a set of performance counters. Significant differences from SPARC M5 are as follows:

1. SPARC M7 has a new cache hierarchy. All events associated with instruction cache hit/miss, instruction prefetch hit/miss/drop, data cache hit/miss and software/hardware prefetch hit/miss/drop have been updated to reflect the new cache hierarchy
2. SPARC M7 supports an additional 16GB page size. TLB fill events have been updated to include this page size
3. support of pipeline flush events
4. SPARC M7 supports new L2I and L2D events
5. SPARC M7 supports a richer set of per functional unit performance counters in the SOC blocks.

10.2 SPARC Performance Control Registers

Each virtual processor has four hyperprivileged, read/write Performance Control registers: PCR0, PCR1, PCR2, and PCR3. Each PCR controls its corresponding PIC: PCR0 controls PIC0, PCR1 controls PIC1, PCR2 controls PIC2, and PCR3 controls PIC3.

Each Performance Control register contains ten fields: `ntc`, `picnht`, `picnpt`, `sl`, `mask`, `ht`, `ut`, `st`, `toe`, and `ov`. All bits except `ntc` and `ov` are always updated on a Performance Control register write. `ov` is a state bit associated with PIC overflow traps and is provided to allow software to determine whether a PIC counter has overflowed.

`ntc` and `ov` can be reset by software but can never be written to 1. `sl` controls which events are counted in a PIC. `mask` is used in conjunction with `sl` to determine which set of subevents are counted in a PIC. `toe` controls whether a trap is generated when the PIC counter overflows. `ut` controls whether user-level events are counted. `st` controls whether supervisor-level events are counted. `ht` controls whether hypervisor level events are counted. The format of this register is shown in TABLE 10-1. Note that changing a field in the PCR does not directly affect a PIC value. To reliably change the events being monitored, software should perform the following sequence:

1. Disable counting by writing zeroes to PCR.sl and clearing PCR.ut, PCR.ht, and PCR.st.
2. Reset the PIC.

3. Enable the new event via writing a non-zero value to PCR.sl and setting PCR.ut, PCR.ht, or PCR.st, as appropriate.

TABLE 10-1 Performance Control Registers – PCR0-3 (ASI 64₁₆, VA 00₁₆, 08₁₆, 10₁₆, 18₁₆)

Bit	Field	Initial Value	R/W	Description
63:19	—	0	RO	<i>Reserved</i>
18	ntc	0	RW	Set to 1 when PIC wraps from 2 ³² –1 to 0 on a next-to-commit (ntc) instruction ¹ . Once set, ntc remains set until reset by software. Hardware sets ntc whenever it sets ov on a next-to-commit instruction.
17	picnht	0	RW	PIC non-hyperprivileged trap. Privileged software can access the PIC only if picnht = 0, otherwise a <i>privileged_action</i> trap occurs. Non-privileged software can access PIC only when picnht = 0 and picnpt = 0, otherwise a <i>privileged_action</i> trap occurs.
16	picnpt	0	RW	PIC non-privileged trap. Non-privileged software can access PIC only when picnht = 0 and picnpt = 0, otherwise a <i>privileged_action</i> trap occurs.
15:11	sl	0	RW	Selects one of 32 events to be counted for PIC as per the following table.
10:5	mask	0	RW	Mask event for PIC as listed in TABLE 10-2.
4	ht	0	RW	If ht = 1, count events in hyperprivileged mode; otherwise, ignore hyperprivileged mode events.
3	st	0	RW	If st = 1, count events in privileged mode; otherwise, ignore privileged mode events.
2	ut	0	RW	If ut = 1, count events in user mode; otherwise, ignore user mode events.
1	toe	0	RW	Trap-on-Event: This field controls whether a trap to hyperprivileged software occurs if the corresponding PIC counter overflows. Hardware ANDs the value of toe with ov to produce a trap.
0	ov	0	RW	Set to 1 when PIC wraps from 2 ³² –1 to 0. Once set, ov remains set until reset by software.

1. The following instructions are next-to-commit instructions: MD5, SHA1, SHA256, SHA512, MPMUL, MONTMUL, MONTSQL, XMPMUL, XMONTMUL, XMONTSQL, loads and stores to I/O space, CAS{X}A, LDSTUB, SWAP, WRHPR, WRASR, WRPR, RDHPR, RDPR, RDASR instructions, and any non-translating load or store alternate instruction as defined in Table 9-1, “SPARC M7 ASI Usage,” on page 54.

10.3 SPARC Performance Instrumentation Counter

Each virtual processor has four Performance Instrumentation Counter registers: PIC0, PIC1, PIC2, and PIC3. PCR0 controls PIC0, PCR1 controls PIC1, PCR2 controls PIC2, and PCR3 controls PCR3. Access privilege is controlled by the settings of PCR.picnht and PCR.picnpt. When PCR.picnht = 1 an attempt to access this register in privileged or nonprivileged mode causes a *privileged_action* trap. When PCR.picnpt = 1 an attempt to access this register in nonprivileged mode causes a *privileged_action* trap.

The PIC counter contains a single 32-bit counter field. The field counts the event selected by PCR.sl. The ut, st, and ht fields for PCR control which combination of user, supervisor, and/or hypervisor events are counted.

The format of the PIC registers are shown in TABLE 10-2.

TABLE 10-2 Performance Instrumentation Counter Register – PIC0-3 (ASI B0₁₆, VA 00₁₆, 08₁₆, 10₁₆, 18₁₆)

Bit	Field	Initial Value	R/W	Description
63:32	—	0	RW	<i>Reserved</i>
31:0	counter	0	RW	Programmable event counter, event controlled by PCR.sl.

█

Implementation Dependencies

11.1 SPARC V9 General Information

SPARC M7 complies with Oracle SPARC Architecture 2015 except where specifically noted. Oracle SPARC Architecture 2015 is generally a superset of SPARC V9.

11.1.1 Level-2 Compliance (Impdep #1)

SPARC M7 is designed to meet Level-2 SPARC V9 compliance. It

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

Note | System emulation routines (for example, quad-precision floating-point operations) shipped with SPARC M7 also must be Level-2 compliant.

11.1.2 Unimplemented Opcodes, ASIs, and ILLTRAP

SPARC V9 unimplemented, *reserved*, ILLTRAP opcodes, and instructions with invalid values in *reserved* fields (other than *reserved* FPOps) encountered during execution cause an *illegal_instruction* trap. Unimplemented and *reserved* ASI values cause a *DAE_invalid_ASI* trap.

11.1.3 Trap Levels (Impdep #37, 38, 39, 40, 114, 115)

SPARC M7 supports two trap levels; that is, *MAXPTL* = 2. Normal execution is at TL = 0.

A virtual processor normally executes at trap level 0 (*execute_state*, TL = 0). Per SPARC V9, a trap causes the virtual processor to enter the next higher trap level, which is a very 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.

11.1.4 Trap Handling (Impdep #16, 32, 33, 35, 36, 44)

SPARC M7 supports precise trap handling for all operations except for deferred and disrupting traps from hardware failures and interrupts. SPARC M7 implements precise traps, interrupts, and exceptions for all instructions, including long-latency floating-point operations. Multiple traps levels

are supported, allowing graceful recovery from faults. SPARC M7 can efficiently execute kernel code even in the event of multiple nested traps, promoting strand efficiency while dramatically reducing the system overhead needed for trap handling.

Three sets of global registers are provided. 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 SPARC M7 are listed in TABLE 6-2 on page 102.

11.1.5 Secure Software

To establish an enhanced security environment, it may be necessary to initialize certain virtual processor 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 (Impdep #102)*.

11.1.6 Address Masking (Impdep #125)

SPARC M7 follows Oracle SPARC Architecture 2015 for PSTATE.am masking and for PSTATE.vme masking. Addresses to non-translating ASIs, *REAL* ASIs, and accesses that bypass translation are never masked.

11.2 Integer Operations

11.2.1 Integer Register File and Window Control Registers (Impdep #2)

SPARC M7 implements an eight-window 64-bit integer register file; that is, $N_REG_WINDOWS = 8$. SPARC M7 truncates values stored in the CWP, CANSAVE, CANRESTORE, CLEANWIN, and OTHERWIN registers to three bits. This includes implicit updates to these registers by SAVE, SAVED, RESTORE, and RESTORED instructions. The most significant two bits of these registers read as zero.

11.2.2 Clean Window Handling (Impdep #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. System software needs to clean one or more windows before returning to the requesting context.

11.2.3 Integer Multiply and Divide

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

11.2.4 MULScC

SPARC V9 does not define the value of `xcc` and `rd{63:32}` for MULScC. SPARC M7 sets `xcc.n` to 0, `xcc.z` to 1 if `rd{63:0}` is zero and to 0 if `rd{63:0}` is not zero, `xcc.v` to 0, and `xcc.c` to 0. SPARC M7 sets `rd{63:33}` to zeros, and sets `rd{32}` to `icc.c` (that is, `rd{32}` is set if there is a carry-out of `rd{31}`; otherwise, it is cleared).

11.3 SPARC V9 Floating-Point Operations

11.3.1 Overflow, Underflow, and Inexact Traps (Impdep #3, 55)

SPARC M7 implements precise floating-point exception handling. Tininess, as it pertains to underflow is detected before rounding.

11.3.2 Quad-Precision Floating-Point Operations (Impdep #3)

All quad-precision floating-point instructions, listed in TABLE 11-1, cause an *illegal_instruction* trap. These operations are then emulated by system software.

TABLE 11-1 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.
FNEGq	Quad-precision floating-point negation.
FSQRTq	Quad-precision floating-point square root.
FSUBq	Quad-precision floating-point subtraction.

11.3.3 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 or 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.

SPARC V9 allows FPRS.du and FPRS.dl to be set pessimistically. SPARC M7 sets FPRS.du or FPRS.dl either when an instruction that updates the floating-point register file successfully completes, or when an FMOVcc or FMOVr that does not meet the condition successfully completes.

11.3.4 Floating-Point Status Register (FSR) (Impdep #13, 19, 22, 23, 24)

SPARC M7 supports precise-traps and implements all three exception fields (tem, cexc, and aexc) conforming to IEEE Standard 754-1985.

SPARC M7 implements the FSR register according to the definition in Oracle SPARC Architecture 2015, with the following implementation-specific clarifications:

- SPARC M7 does not contain an FQ, therefore FSR.qne always reads as 0 and an attempt to read the FQ with an RDPR instruction causes an *illegal_instruction* trap.
- SPARC M7 does not detect the unimplemented_FPop, unfinished_FPop, sequence_error, hardware_error, or invalid_fp_register floating-point trap types directly in hardware, therefore does not generate a trap when those conditions occur.

TABLE 11-2 documents the fields of the FSR.

TABLE 11-2 Floating-Point Status Register Format

Bits	Field	RW	Description										
63:38	—	RO	<i>Reserved</i>										
37:36	fcc3	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. Note: fcc0 is the same as the FCC in SPARC V8.										
35:34	fcc2	RW	Floating-point condition code (set 2). See fcc3 description.										
33:32	fcc1	RW	Floating-point condition code (set 1) See fcc3 description.										
31:30	rd	RW	IEEE Std. 754-1985 rounding direction, as follows: <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>rd</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>+∞</td> </tr> <tr> <td>3</td> <td>-∞</td> </tr> </tbody> </table>	rd	Round Toward	0	Nearest (even if tie)	1	0	2	+∞	3	-∞
rd	Round Toward												
0	Nearest (even if tie)												
1	0												
2	+∞												
3	-∞												
29:28	—	RO	<i>Reserved</i>										
27:23	tem	RW	IEEE-754 trap enable mask. Five-bit trap enable mask for the IEEE-754 floating-point exceptions. If a floating-point operate instruction produces one or more exceptions, the corresponding cexc/aexc bits are set and an <i>fp_exception_ieee_754</i> (with FSR.ftt = 1, <i>IEEE_754_exception</i>) exception is generated.										

TABLE 11-2 Floating-Point Status Register Format (Continued)

Bits	Field	RW	Description																											
22	ns	RO	Nonstandard floating-point results. SPARC M7 does not implement a non-standard floating-point mode. FSR.ns always reads as 0, and writes to it are ignored.																											
21:20	—	RO	<i>Reserved</i>																											
19:17	ver	RO	FPU version number. This field identifies a particular implementation of the SPARC M7 FPU architecture.																											
16:14	ftt	RW	Floating-point trap type. Set whenever a floating-point instruction causes the <i>fp_exception_ieee_754</i> or <i>fp_exception_other</i> traps. Values are as follows: <table border="1" data-bbox="706 489 1328 810"> <thead> <tr> <th>ftt</th> <th>Floating-Point Trap Type</th> <th>Trap Signalled</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>None</td> <td>—</td> </tr> <tr> <td>1</td> <td>IEEE_754_exception</td> <td><i>fp_exception_ieee_754</i></td> </tr> <tr> <td>2</td> <td>reserved</td> <td>—</td> </tr> <tr> <td>3</td> <td>reserved</td> <td>—</td> </tr> <tr> <td>4</td> <td>reserved</td> <td>—</td> </tr> <tr> <td>5</td> <td>reserved</td> <td>—</td> </tr> <tr> <td>6</td> <td>invalid_fp_register</td> <td><i>fp_exception_other</i></td> </tr> <tr> <td>7</td> <td>reserved</td> <td>—</td> </tr> </tbody> </table>	ftt	Floating-Point Trap Type	Trap Signalled	0	None	—	1	IEEE_754_exception	<i>fp_exception_ieee_754</i>	2	reserved	—	3	reserved	—	4	reserved	—	5	reserved	—	6	invalid_fp_register	<i>fp_exception_other</i>	7	reserved	—
ftt	Floating-Point Trap Type	Trap Signalled																												
0	None	—																												
1	IEEE_754_exception	<i>fp_exception_ieee_754</i>																												
2	reserved	—																												
3	reserved	—																												
4	reserved	—																												
5	reserved	—																												
6	invalid_fp_register	<i>fp_exception_other</i>																												
7	reserved	—																												
			Note: SPARC M7 neither detects nor generates the <i>unimplemented_FPop</i> , <i>unfinished_FPop</i> , <i>sequence_error</i> , <i>hardware_error</i> or <i>invalid_fp_register</i> trap types directly in hardware.																											
			Note: SPARC M7 does not contain an FQ. An attempt to read the FQ with an RDPR instruction causes an <i>illegal_instruction</i> trap.																											
13:	qne	RW	Floating-point deferred-trap queue (FQ) not empty. Not used, because SPARC M7 implements precise floating-point exceptions.																											
12	—	RO	<i>Reserved</i>																											
11:10	fcc0	RW	Floating-point condition code (set 0). See <i>fcc3</i> description.																											
9:5	aexc	RW	Accumulated outstanding exceptions. Accumulates IEEE 754 exceptions while floating-point exception traps are disabled (that is, while corresponding bit in FSR.tem is zero)																											
4:0	cexc	RW	Current outstanding exceptions. Indicates the most recently generated IEEE 754 exceptions.																											

11.4 SPARC V9 Memory-Related Operations

11.4.1 Load/Store Alternate Address Space (Impdep #5, 29, 30)

Supported ASI accesses are listed in Section 9.2.

11.4.2 Read/Write ASR (Impdep #6, 7, 8, 9, 47, 48)

Supported ASRs are listed in Chapter 3, *Registers*.

11.4.3 MMU Implementation (Impdep #41)

SPARC M7 memory management is based on in-memory Translation Storage Buffers (TSBs) backed by a Software Translation Table. See Chapter 13, *Memory Management Unit* for more details.

11.4.4 FLUSH and Self-Modifying Code (Impdep #122)

FLUSH is needed to synchronize code and data spaces after code space is modified during program execution. FLUSH is described in Section D.2.4. On SPARC M7, the FLUSH effective address is ignored, and as a result, FLUSH cannot cause a *DAE_invalid_ASI* trap.

Note SPARC V9 specifies that the FLUSH instruction has no latency on the issuing virtual processor. In other words, a store to instruction space prior to the FLUSH instruction is visible immediately after the completion of FLUSH. When a flush is performed, SPARC M7 guarantees that earlier code modifications will be visible across the whole system.

11.4.5 PREFETCH{A} (Impdep #103, 117)

For SPARC M7 PREFETCH{A} instruction documentation, see Section 5.2, *PREFETCH/PREFETCHA*, on page 35.

11.4.6 LDD/STD Handling (Impdep #107, 108)

LDD and STD instructions are directly executed in hardware.

Note LDD/STD are deprecated in SPARC V9. In SPARC M7 it is more efficient to use LDX/STX for accessing 64-bit data. LDD/STD take longer to execute than two 32- or 64-bit loads/stores.

11.4.7 FP mem_address_not_aligned (Impdep #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.

11.4.8 Supported Memory Models (Impdep #113, 121)

SPARC M7 supports only the TSO memory model, although certain specific operations such as block loads and stores operate under the RMO memory model. See Chapter 8, Section 8.2. Supported Memory Models.”.

11.4.9 Implicit ASI When TL > 0 (Impdep #124)

SPARC M7 matches all Oracle SPARC 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*.

11.5 Non-SPARC V9 Extensions

11.5.1 Cache Subsystem

SPARC M7 contains one or more levels of cache. The cache subsystem architecture is described in Appendix D, *Cache Coherency and Ordering*.

11.5.2 Block Memory Operations

SPARC M7 supports 64-byte block memory operations utilizing a block of eight double-precision floating point registers as a temporary buffer. See Section 5.5.

11.5.3 Partial Stores

SPARC M7 supports 8-/16-/32-bit partial stores to memory. See Section 5.5.

11.5.4 Short Floating-Point Loads and Stores

SPARC M7 supports 8-/16-bit loads and stores to the floating-point registers.

11.5.5 Load Twin Extended Word

SPARC M7 supports 128-bit atomic load operations to a pair of integer registers.

11.5.6 SPARC M7 Instruction Set Extensions (Impdep #106)

The SPARC M7 processor supports VIS 3.0. VIS instructions are designed to enhance graphics functionality and improve the efficiency of memory accesses.

Unimplemented IMPDEP1 and IMPDEP2 opcodes encountered during execution cause an *illegal_instruction* trap.

Other instruction extensions are described in Chapter 3, *Registers*.

11.5.7 Performance Instrumentation

SPARC M7 performance instrumentation is described in Chapter 10, *Performance Instrumentation*.

11.5.8 ASI_MONITOR_AS_IF_USER_PRIMARY, ASI_MONITOR_AS_IF_USER_SECONDARY, ASI_MONITOR_PRIMARY, ASI_MONITOR_SECONDARY

When using these ASIs with a load, SPARC M7 monitors for stores to the L1 data cache line (32 bytes) that contains the address of the load. Upon detecting that a store has occurred to the L1 data cache line, the load monitor is invalidated. If a subsequent MWAIT has suspended the strand, then the strand resumes on detection of the store to the L1 data cache line.

Cryptographic Extensions

SPARC M7 provides cryptographic support via non-privileged instructions. The instructions accelerate bulk ciphers, secure hashes, and public-key algorithms. Since these instructions are non-privileged, they can be used directly by applications, or by commonly used open source cryptographic libraries such as OpenSSL.

In SPARC M7, symmetric ciphers are implemented such that a single instruction is capable of performing a significant portion of a round. Secure hashes are implemented such that a single instruction performs a single block of the hash operation (i.e. multiple rounds). Public-key operations are accelerated via instructions that perform large (up to 2048-bit) Montgomery multiplication operations. More details on these instructions can be found in Chapter 5, *Instruction Definitions*.

The SPARC M7 implements the Compatibility Feature Register (CFR), which allows future UltraSPARC processors to drop support for older, deprecated ciphers (and introduce support for new ones) by reclaiming opcodes previously reserved for old ciphers.

12.1 CFR Register

The CFR is described in Chapter 3, *Registers*.

12.2 Cryptographic Instructions

SPARC M7 introduces a number of new cryptographic opcodes, which are detailed in Chapter 5, *Instruction Definitions*.

12.3 Cryptographic performance

For a single-thread executing on a core, the basic low-level performance on SPARC M7 is detailed in the following tables.

TABLE 12-1 Symmetric-key performance

Algorithm	Block Size (Bytes)	Block Latency (Cycles)
DES-ECB	8	
3DES-ECB	8	
AES-128-ECB	16	
AES-192-ECB	16	
AES-256-ECB	16	
Camellia		

TABLE 12-2 Secure hash performance

Algorithm	Block Size (Bytes)	Block Latency (Cycles)
MD5	64	186
SHA-1	64	220
SHA-256	64	188
SHA-512	128	236

12.4 SPARC M7 crypto coding guidance

It is anticipated that the SPARC M7 cryptographic instructions will be widely deployed - not only in Solaris libraries, but also in Open Source libraries like OpenSSL. Implementation of key cryptographic algorithms using these instructions is very straight-forward, and example use is provided in the instructions chapter. It is important that software use the CFR as detailed in Section 3.2.8, *Compatibility Feature Register (CFR)*, on page 21, or software may perform sub-optimally on future processors.

Memory Management Unit

This chapter provides detailed information about the SPARC M7 Memory Management Unit. It describes the internal architecture of the MMU and how to program it.

13.1 Translation Table Entry (TTE)

The Translation Table Entry 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.

SPARC M7 supports the sun4v TTE format as shown in the MMU chapter of Oracle SPARC Architecture 2015, with the following notes:

- SPARC M7 supports a 16-bit TTE Tag context ID field, formed from bits 63:48 of the TTE Tag
- SPARC M7 only supports 54-bit Virtual Addresses
- On SPARC M7, bits 55:13 of TTE Data contains the real page¹ number. Bits {55:50} should always be zero.
- The meaning of the `cp` bit in TTE Data on SPARC M7 is:

TABLE 13-1 Cacheable Field Encoding (from TSB)

Cacheable (cp)	Meaning of TTE When Placed in:	
	iTLB (I-cache PA-Indexed)	dTLB (D-cache PA-Indexed)
0	Cacheable in L2 and L3 caches only	Cacheable in L2 and L3 caches only
1	Cacheable in L3 cache, L2 cache, and I-cache	Cacheable in L3 cache, L2 cache, and D-cache

- For the IMMU and DMMU, on SPARC M7 the `ep` bit in the TTE is not written into the TLB, and returns zero on a Data Access read.
- On SPARC M7, the value of the `w` bit written into the ITLB will be read out on an ITLB Data Access read (impl. dep. #__).
- The following page sizes are supported on SPARC M7 (in `TTE.sz`): 8 KB, 64 KB, 4 MB, 256 MB, 2 GB, and 16 GB. Other encodings of `TTE.sz` are reserved.

TABLE 13-2 shows the Oracle SPARC Architecture 2015 TSB TTE tag format as interpreted by SPARC M7.

¹ sun4v supports translation from virtual addresses (VA) to real addresses (RA). Privileged code manages the VA-to-RA translations.

TABLE 13-2 sun4v TSB TTE Tag Format

Bit	Field	Description
63:48	context	The 16-bit context identifier associated with the TTE.
47:42	0	Must be 0
41:0	va	Virtual Address Tag{63:22}. The virtual page number. Bits 21 through 13 are not maintained in the tag, since these bits are used to index the smallest TSB (512 entries). NOTE: SPARC M7 hardware only supports a 54-bit VA.

The sun4v TSB TTE data format is shown in TABLE 13-3.

TABLE 13-3 sun4v TSB TTE Data Format

Bit	Field	Description
63	v	Valid. If the Valid bit is set, the remaining fields of the TTE are meaningful.
62	nfo	No-fault-only. If this bit is set, loads with <code>ASI_PRIMARY_NO_FAULT{LITTLE}</code> , <code>ASI_SECONDARY_NO_FAULT{LITTLE}</code> are translated. Any other DMMU access will trap with a <code>DAE_nfo_page</code> trap. For the IMMU, if the nfo bit is set, an <code>iae_nfo_page</code> trap will be taken.
61:56	soft2	<code>soft2</code> and <code>soft</code> are software-defined fields, provided for use by the operating system. Software fields are not implemented in the SPARC M7 TLB. <code>soft</code> and <code>soft2</code> fields may be written with any value; hardware ignores these fields. The fields are not preserved in the TLBs.
55:13	ra	The real page ¹ number. For SPARC M7, a 50-bit real address range is supported by the hardware tablewalker, and bits {55:50} should always be zero.
12	ie	Invert endianness. If this bit is set, accesses to the associated page are processed with inverse endianness from what is specified by the instruction (big-for-little and little-for-big). For the IMMU, the <code>ie</code> bit in the TTE is written into the ITLB but ignored during ITLB operation. The value of the <code>ie</code> bit written into the ITLB will be read out on an ITLB Data Access read. Note: This bit is intended to be set primarily for noncacheable accesses.
11	e	Side effect. If this bit is set, noncacheable memory accesses other than block loads and stores are strongly ordered against other <code>e</code> bit accesses, and noncacheable stores are not merged. This bit should be set for pages that map I/O devices having side effects. Note, however, that the <code>e</code> bit does not prevent normal instruction prefetching. For the IMMU, the <code>e</code> bit in the TTE is written into the ITLB, but ignored during ITLB operation. The value of the <code>e</code> bit written into the ITLB will be read out on an ITLB Data Access read. NOTE: The <code>e</code> bit does not force an uncacheable access. It is expected, but not required, that the <code>cp</code> bit will be set to zero when the <code>e</code> bit is set.
10	cp	The cacheable-in-physically-indexed-cache (<code>cp</code>) bit determines the placement of data in SPARC M7 caches, according to TABLE 13-4. The MMU does not operate on the cacheable bit, but merely passes them through to the cache subsystem.

TABLE 13-4 Cacheable Field Encoding (from TSB)

Cacheable (cp)	Meaning of TTE When Placed in:	
	iTLB (I-cache PA-Indexed)	dTLB (D-cache PA-Indexed)
0	Cacheable in L2 and L3 caches only	Cacheable in L2 and L3 caches only
1	Cacheable in L3 cache, L2 cache, and I-cache	Cacheable in L3 cache, L2 cache, and D-cache

TABLE 13-3 sun4v TSB TTE Data Format (Continued)

Bit	Field	Description
8	p	Privileged. If the p bit is set, only privileged software can access the page mapped by the TTE. If the p bit is set and an access to the page is attempted when PSTATE.priv = 0, the MMU will signal an <i>IAE_privilege_violation</i> or <i>DAE_privilege_violation</i> trap.
7	ep	Executable. If the ep bit is set, the page mapped by this TTE has execute permission granted. Otherwise, execute permission is not granted and the hardware table-walker will not load the ITLB with a TTE with ep = 0. For the IMMU and DMMU, the ep bit in the TTE is not written into the TLB. It returns one on a Data Access read for the ITLB and zero on a Data Access read for the DTLB.
6	w	Writable. If the w bit is set, the page mapped by this TTE has write permission granted. Otherwise, write permission is not granted and the MMU will cause a trap if a write is attempted. For the IMMU, the w bit in the TTE is written into the ITLB, but ignored during ITLB operation. The value of the w bit written into the ITLB will be read out on an ITLB Data Access read.
5:4	soft	(see soft2, above)
3:0	size	The page size of this entry, encoded as shown in TABLE 13-5.

TABLE 13-5 Size Field Encoding (from TTE)

Size{3:0}	Page Size
0000	8 KB
0001	64 KB
0010	Reserved
0011	4 MB
0100	Reserved
0101	256 MB
0110	2 GB
0111	16 GB
1000-1111	Reserved

1. sun4v supports translation from virtual addresses (VA) to real addresses (RA) to physical addresses (PA). Privileged code manages the VA-to-RA translations.

13.2 Translation Storage Buffer (TSB)

A TSB is an array of TTEs managed entirely by software. It serves as a cache of the Software Translation table

A TSB is arranged as a direct-mapped cache of TTEs.

The TSB exists as a normal data structure in memory and therefore may be cached. 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 the TSB organization. The constant *N* is determined by the size field in the TSB register; it may range from 512 entries to 1T entries.




FIGURE 13-1 TSB Organization

13.3 MMU-Related Faults and Traps

13.3.1 *IAE_privilege_violation* Trap

The I-MMU detects a privilege violation for an instruction fetch; that is, an attempted access to a privileged page when `PSTATE.priv = 0`.

13.3.2 *IAE_nfo_page* Trap

During a hardware tablewalk, the I-MMU matches a TTE entry whose `nfo` (no-fault-only) bit is set.

Implementation Note | The `nfo` bit is only checked on I-MMU translations. It is not checked on hardware tablewalks.

13.3.3 *DAE_privilege_violation* Trap

This trap occurs when the D-MMU detects a privilege violation for a data access; that is, a load or store instruction attempts access to a privileged page when `PSTATE.priv = 0`.

13.3.4 *DAE_side_effect_page* Trap

This trap occurs when a (nonfaulting) load instruction is issued to a page marked with the side-effect (`e`) bit = 1.

13.3.5 *DAE_nc_page* Trap

This trap occurs when an atomic instruction (including a 128-bit atomic load) is issued to a memory address marked uncacheable; for example,, with `cp = 0`.

Implementation Note | For SPARC M7, `cp` only controls cacheability in the L1 cache, not the private L2 caches or the shared L3. SPARC M7 performs atomic operations in the L2 cache and supports the ability to complete an atomic operation for pages with the `cp` bit = 0 even if the L2 cache is disabled. However, to keep SPARC M7 compliant with the Oracle SPARC Architecture 2011 specification, a *DAE_nc_page* trap is generated when an atomic is issued to a memory address marked with `cp = 0`.

13.3.6 *DAE_invalid_asi* Trap

This trap occurs when an invalid LDA/STA ASI value, invalid virtual address, read to write-only register, or write to read-only register occurs, but not for an attempted user access to a restricted ASI (see the *privileged_action* trap described below).

13.3.7 *DAE_nfo_page* Trap

This trap occurs when an access occurs with an ASI other than `ASI_{PRIMARY,SECONDARY}_NO_FAULT{_LITTLE}` to a page marked with the nfo (no-fault-only) bit.

13.3.8 *privileged_action* Trap

13.3.9 This trap occurs when an access is attempted using a *restricted* ASI while in non-privileged mode (PSTATE.priv = 0). **mem_address_not_aligned* Traps

The *lddf_mem_address_not_aligned*, *stdf_mem_address_not_aligned*, and *mem_address_not_aligned* traps occur when a load, store, atomic, or JMPL/RETURN instruction with a misaligned address is executed.

13.4 MMU Operation Summary

TABLE 13-8 summarizes the behavior of the D-MMU for noninternal ASIs using tabulated abbreviations. TABLE 13-11 summarizes the behavior of the I-MMU. In each case, and for all conditions, the behavior of the MMU is given by one of the abbreviations in TABLE 13-6. TABLE 13-7 lists abbreviations for ASI types.

TABLE 13-6 Abbreviations for MMU Behavior

Abbreviation	Meaning
ok	Normal translation
dasi	<i>DAE_invalid_asi</i> trap
dpriv	<i>DAE_privilege_violation</i> trap
dse	<i>DAE_side_effect_page</i> trap
ipriv	<i>IAE_privilege_violation</i> trap

TABLE 13-7 Abbreviations for ASI Types

Abbreviation	Meaning
NUC	ASI_NUCLEUS*
PRIM	Any ASI with PRIMARY translation, except *NO_FAULT
SEC	Any ASI with SECONDARY translation, except *NO_FAULT
PRIM_NF	ASI_PRIMARY_NO_FAULT*
SEC_NF	ASI_SECONDARY_NO_FAULT*
U_PRIM	ASI_*_AS_IF_USER_PRIMARY*

TABLE 13-7 Abbreviations for ASI Types

Abbreviation	Meaning
U_SEC	ASI_*_AS_IF_USER_SECONDARY*
U_PRIV	ASI_*_AS_IF_PRIV_*
REAL	ASI_*REAL*

Note | The *_LITTLE versions of the ASIs behave the same as the big-endian versions with regard to the MMU table of operations.

Other abbreviations include “w” for the writable bit, “e” for the side-effect bit, and “p” for the privileged bit.

TABLE 13-8 and TABLE 13-11 do not cover the following cases:

- Invalid ASIs, ASIs that have no meaning for the opcodes listed, or nonexistent ASIs; for example, ASI_PRIMARY_NO_FAULT for a store or atomic; also, access to SPARC M7 internal registers other than LDXA, LDFA, STDFA or STXA; the MMU signals a *DAE_invalid_asi* trap for this case.
- Attempted access using a restricted ASI in nonprivileged mode; the MMU signals a *privileged_action* trap for this case. Attempted use of a hyperprivileged ASI in privileged mode; the MMU also signals *privileged_action* trap for this case.
- An atomic instruction (including 128-bit atomic load) issued to a memory address marked uncacheable in a physical cache (that is, with cp = 0 or pa{49} = 1); the MMU signals a *DAE_nc_page* trap for this case.
- A data access with an ASI other than ASI_{PRIMARY,SECONDARY}_NO_FAULT{LITTLE} to a page marked nfo; the MMU signals a *DAE_nfo_page* for this case.
- An instruction access to a page marked with the nfo (no-fault-only) bit. The MMU signals an *IAE_nfo_page* trap for this case.
- An instruction fetch to a memory address marked non-executable (ep = 0). This is checked when Hardware Tablewalk attempts to load the I-MMU, and an *IAE_unauth_access* trap is taken instead.
- Real address out of range; the MMU signals an *instruction_real_range* trap for this case.
- Virtual address out of range and PSTATE.am is not set; the MMU signals an *instruction_address_range* trap for this case.

TABLE 13-8 D-MMU Operations for Normal ASIs

Condition				Behavior			
Opcode	priv Mode	ASI	w	e = 0 p = 0	e = 0 p = 1	e = 1 p = 0	e = 1 p = 1
Load	non-privileged	PRIM, SEC	—	ok	dpriv	ok	dpriv
		PRIM_NF, SEC_NF	—	ok	dpriv	dse	dpriv
	privileged	PRIM, SEC, NUC	—	ok			
		PRIM_NF, SEC_NF	—	ok		dse	
		U_PRIV, U_SEC	—	ok	dpriv	ok	dpriv
	REAL	—	ok				
FLUSH	non-privileged		—	ok			
	privileged		—	ok			

TABLE 13-8 D-MMU Operations for Normal ASIs

Condition				Behavior				
Opcode	priv Mode	ASI	w	e = 0 p = 0	e = 0 p = 1	e = 1 p = 0	e = 1 p = 1	
Store or Atomic	non-privileged	PRIM, SEC	0			dpriv	dpriv	
			1		ok	dpriv	ok	dpriv
	privileged	PRIM, SEC, NUC	0		ok			
			1					
		U_PRIM, U_SEC	0			dpriv		dpriv
			1		ok	dpriv	ok	dpriv
		REAL	0		ok			
			1					

See Section 9.2, *Alternate Address Spaces*, on page 54 for a summary of the SPARC M7 ASI map.

13.5 Translation

13.5.1 Instruction Translation

13.5.1.1 Instruction Prefetching

SPARC M7 fetches instructions sequentially (including delay slots). SPARC M7 fetches delay slots before the branch is resolved (before whether the delay slot will be annulled is known). SPARC M7 also fetches the target of a DCTI before the delay slot executes.

13.5.2 Data Translation

TABLE 13-9 DMMU Translation (1 of 4)

ASI Value (hex)	ASI NAME	Translation		
		Nonprivileged	Privileged	Hypervisor
00 ₁₆ –01 ₁₆	<i>Reserved</i>	<i>privileged_action</i>	<i>DAE_invalid_asi</i>	
03 ₁₆	<i>Reserved</i>	<i>privileged_action</i>	<i>DAE_invalid_asi</i>	
04 ₁₆	ASI_NUCLEUS	<i>privileged_action</i>	VA → PA	
06 ₁₆ –0B ₁₆	<i>Reserved</i>	<i>privileged_action</i>	<i>DAE_invalid_asi</i>	
0C ₁₆	ASI_NUCLEUS_LITTLE	<i>privileged_action</i>	VA → PA	
0D ₁₆ –0F ₁₆	<i>Reserved</i>	<i>privileged_action</i>	<i>DAE_invalid_asi</i>	
10 ₁₆	ASI_AS_IF_USER_PRIMARY	<i>privileged_action</i>	VA → PA	
11 ₁₆	ASI_AS_IF_USER_SECONDARY	<i>privileged_action</i>	VA → PA	
12 ₁₆	ASI_MONITOR_AS_IF_USER_PRIMARY	<i>privileged_action</i>	VA → PA	
13 ₁₆	ASI_MONITOR_AS_IF_USER_SECONDARY	<i>privileged_action</i>	VA → PA	
14 ₁₆	ASI_REAL	<i>privileged_action</i>	RA → PA	
15 ₁₆	ASI_REAL_IO	<i>privileged_action</i>	RA → PA	
16 ₁₆	ASI_BLOCK_AS_IF_USER_PRIMARY	<i>privileged_action</i>	VA → PA	

TABLE 13-9 DMMU Translation (2 of 4)

ASI Value (hex)	ASI NAME	Translation		
		Nonprivileged	Privileged	Hypervisor
17 ₁₆	ASI_BLOCK_AS_IF_USER_SECONDARY	<i>privileged_action</i>	VA → PA	
18 ₁₆	ASI_AS_IF_USER_PRIMARY_LITTLE	<i>privileged_action</i>	VA → PA	
19 ₁₆	ASI_AS_IF_USER_SECONDARY_LITTLE	<i>privileged_action</i>	VA → PA	
1A ₁₆ – 1B ₁₆	<i>Reserved</i>	<i>privileged_action</i>	<i>DAE_invalid_asi</i>	
1C ₁₆	ASI_REAL_LITTLE	<i>privileged_action</i>	RA → PA	
1D ₁₆	ASI_REAL_IO_LITTLE	<i>privileged_action</i>	RA → PA	
1E ₁₆	ASI_BLOCK_AS_IF_USER_PRIMARY_LITTLE	<i>privileged_action</i>	VA → PA	
1F ₁₆	ASI_BLOCK_AS_IF_USER_SECONDARY_LITTLE	<i>privileged_action</i>	VA → PA	
20 ₁₆	ASI_SCRATCHPAD	<i>privileged_action</i>	nontranslating	
21 ₁₆	ASI_PRIMARY_CONTEXT_0_REG, ASI_PRIMARY_CONTEXT_1_REG, ASI_SECONDARY_CONTEXT_0_REG, ASI_SECONDARY_CONTEXT_1_REG	<i>privileged_action</i>	nontranslating	
22 ₁₆	ASI_TWINK_AIUP, ASI_STBI_AIUP	<i>privileged_action</i>	VA → PA	
23 ₁₆	ASI_TWINK_AIUS, ASI_STBI_AIUS	<i>privileged_action</i>	VA → PA	
24 ₁₆	<i>Reserved</i>	<i>privileged_action</i>	<i>DAE_invalid_asi</i>	
25 ₁₆	ASI_QUEUE	<i>privileged_action</i>	nontranslating	
26 ₁₆	ASI_TWINK_REAL, ASI_STBI_REAL	<i>privileged_action</i>	RA → PA	
27 ₁₆	ASI_TWINK_NUCLEUS, ASI_STBI_N	<i>privileged_action</i>	VA → PA	
28 ₁₆ – 29 ₁₆	<i>Reserved</i>	<i>privileged_action</i>	<i>DAE_invalid_asi</i>	
2A ₁₆	ASI_TWINK_AIUPL, ASI_STBI_AIUPL	<i>privileged_action</i>	VA → PA	
2B ₁₆	ASI_TWINK_AIUSL, ASI_STBI_AIUSL	<i>privileged_action</i>	VA → PA	
2C ₁₆	<i>Reserved</i>	<i>privileged_action</i>	<i>DAE_invalid_asi</i>	
2D ₁₆	<i>Reserved</i>	<i>privileged_action</i>	<i>DAE_invalid_asi</i>	
2E ₁₆	ASI_TWINK_REAL_LITTLE, ASI_STBI_REAL_LITTLE	<i>privileged_action</i>	RA → PA	
2F ₁₆	ASI_TWINK_NL, ASI_STBI_NL	<i>privileged_action</i>	VA → PA	
80 ₁₆	ASI_PRIMARY	VA → PA	VA → PA	
81 ₁₆	ASI_SECONDARY	VA → PA	VA → PA	
82 ₁₆	ASI_PRIMARY_NO_FAULT	VA → PA	VA → PA	
83 ₁₆	ASI_SECONDARY_NO_FAULT	VA → PA	VA → PA	
84 ₁₆	ASI_MONITOR_PRIMARY	VA → PA	VA → PA	
85 ₁₆	ASI_MONITOR_SECONDARY	VA → PA	VA → PA	
86 ₁₆ – 87 ₁₆	<i>Reserved</i>	<i>DAE_invalid_asi</i>	<i>DAE_invalid_asi</i>	
88 ₁₆	ASI_PRIMARY_LITTLE	VA → PA	VA → PA	

TABLE 13-9 DMMU Translation (3 of 4)

ASI Value (hex)	ASI NAME	Translation		
		Nonprivileged	Privileged	Hypervisor
89 ₁₆	ASI_SECONDARY_LITTLE	VA → PA	VA → PA	
8A ₁₆	ASI_PRIMARY_NO_FAULT_LITTLE	VA → PA	VA → PA	
8B ₁₆	ASI_SECONDARY_NO_FAULT_LITTLE	VA → PA	VA → PA	
8C ₁₆ – 8F ₁₆	<i>Reserved</i>	<i>DAE_invalid_asi</i>	<i>DAE_invalid_asi</i>	
91 ₁₆	<i>Reserved</i>	<i>DAE_invalid_asi</i>	<i>DAE_invalid_asi</i>	
93 ₁₆ – AF ₁₆	<i>Reserved</i>	<i>DAE_invalid_asi</i>	<i>DAE_invalid_asi</i>	
B0 ₁₆	ASI_PIC0, ASI_PIC1, ASI_PIC2, ASI_PIC3	nontranslating	nontranslating	
B1 ₁₆ – BF ₁₆	<i>Reserved</i>	<i>DAE_invalid_asi</i>	<i>DAE_invalid_asi</i>	
C0 ₁₆	ASI_PST8_P	VA → PA	VA → PA	
C1 ₁₆	ASI_PST8_S	VA → PA	VA → PA	
C2 ₁₆	ASI_PST16_P	VA → PA	VA → PA	
C3 ₁₆	ASI_PST16_S	VA → PA	VA → PA	
C4 ₁₆	ASI_PST32_P	VA → PA	VA → PA	
C5 ₁₆	ASI_PST32_S	VA → PA	VA → PA	
C6 ₁₆ – C7 ₁₆	<i>Reserved</i>	<i>DAE_invalid_asi</i>	<i>DAE_invalid_asi</i>	
C8 ₁₆	ASI_PST8_PL	VA → PA	VA → PA	
C9 ₁₆	ASI_PST8_SL	VA → PA	VA → PA	
CA ₁₆	ASI_PST16_PL	VA → PA	VA → PA	
CB ₁₆	ASI_PST16_SL	VA → PA	VA → PA	
CC ₁₆	ASI_PST32_PL	VA → PA	VA → PA	
CD ₁₆	ASI_PST32_SL	VA → PA	VA → PA	
CE ₁₆ – CF ₁₆	<i>Reserved</i>	<i>DAE_invalid_asi</i>	<i>DAE_invalid_asi</i>	
D0 ₁₆	ASI_FL8_P	VA → PA	VA → PA	
D1 ₁₆	ASI_FL8_S	VA → PA	VA → PA	
D2 ₁₆	ASI_FL16_P	VA → PA	VA → PA	
D3 ₁₆	ASI_FL16_S	VA → PA	VA → PA	
D4 ₁₆ – D7 ₁₆		<i>DAE_invalid_asi</i>	<i>DAE_invalid_asi</i>	
D8 ₁₆	ASI_FL8_PL	VA → PA	VA → PA	
D9 ₁₆	ASI_FL8_SL	VA → PA	VA → PA	
DA ₁₆	ASI_FL16_PL	VA → PA	VA → PA	
DB ₁₆	ASI_FL16_SL	VA → PA	VA → PA	
DC ₁₆ – DF ₁₆	<i>Reserved</i>	<i>DAE_invalid_asi</i>	<i>DAE_invalid_asi</i>	
E0 ₁₆	ASI_BLK_COMMIT_PRIMARY	VA → PA	VA → PA	
E1 ₁₆	ASI_BLK_COMMIT_SECONDARY	VA → PA	VA → PA	
E2 ₁₆	ASI_TWIX_P, ASI_STBI_P	VA → PA	VA → PA	

TABLE 13-9 DMMU Translation (4 of 4)

ASI Value (hex)	ASI NAME	Translation		
		Nonprivileged	Privileged	Hypervisor
E3 ₁₆	ASI_TWIX_S, ASI_STBI_S	VA → PA	VA → PA	
E4 ₁₆ – E9 ₁₆	<i>Reserved</i>	<i>DAE_invalid_asi</i>	<i>DAE_invalid_asi</i>	
EA ₁₆	ASI_TWIX_PL, ASI_STBI_PL	VA → PA	VA → PA	
EB ₁₆	ASI_TWIX_SL, ASI_STBI_SL	VA → PA	VA → PA	
EC ₁₆ – EF ₁₆	<i>Reserved</i>	<i>DAE_invalid_asi</i>	<i>DAE_invalid_asi</i>	
F0 ₁₆	ASI_BLK_PRIMARY	VA → PA	VA → PA	
F1 ₁₆	ASI_BLK_SECONDARY	VA → PA	VA → PA	
F2 ₁₆	ASI_STBI_MRU_PRIMARY	VA → PA	VA → PA	
F3 ₁₆	ASI_STBI_MRU_SECONDARY	VA → PA	VA → PA	
F4 ₁₆ – F7 ₁₆	<i>Reserved</i>	<i>DAE_invalid_asi</i>	<i>DAE_invalid_asi</i>	
F8 ₁₆	ASI_BLK_PL	VA → PA	VA → PA	
F9 ₁₆	ASI_BLK_SL	VA → PA	VA → PA	
FA ₁₆	ASI_STBI_MRU_PRIMARY_LITTLE	VA → PA	VA → PA	
FB ₁₆	ASI_STBI_MRU_SECONDARY_LITTLE	VA → PA	VA → PA	
FC ₁₆ – FF ₁₆	<i>Reserved</i>	<i>DAE_invalid_asi</i>	<i>DAE_invalid_asi</i>	

13.6 Compliance With the SPARC V9 Annex F

The SPARC M7 MMU complies completely with the SPARC V9 MMU Requirements described in Annex F of the *The SPARC Architecture Manual, Version 9*. TABLE 13-10 shows how various protection modes can be achieved, if necessary, through the presence or absence of a translation in the I- or D-MMU.

TABLE 13-10 MMU Compliance With SPARC V9 Annex F Protection Mode

Condition			Resultant Protection Mode
TTE in D-MMU	TTE in I-MMU	Writable Attribute Bit	
Yes	No	0	Read-only
No	Yes	Don't Care	Execute-only
Yes	No	1	Read/Write
Yes	Yes	0	Read-only/Execute
Yes	Yes	1	Read/Write/Execute

13.7 MMU Internal Registers and ASI Operations

13.7.1 Accessing MMU Registers

All internal MMU registers can be accessed directly by the virtual processor through ASIs defined by SPARC M7.

See Section 13.5 for details on the behavior of the MMU during all other SPARC M7 ASI accesses.

Note STXA to an MMU register *does not* require any subsequent instructions such as a MEMBAR #Sync, FLUSH, DONE, or RETRY before the register effect will be visible to load / store / atomic accesses. SPARC M7 resolves all MMU register hazards via an automatic synchronization on all MMU register writes.

If the low order three bits of the VA are nonzero in an LDXA/STXA to/from these registers, a *mem_address_not_aligned* trap occurs. Writes to read-only, reads to write-only, illegal ASI values, or illegal VA for a given ASI may cause a *DAE_invalid_asi* trap.

Caution SPARC M7 does not check for out-of-range virtual addresses during an STXA to any internal register; it simply sign-extends the virtual address based on VA{53}. Software must guarantee that the VA is within range.

TABLE 13-11 SPARC M7 MMU Internal Registers and ASI Operations

I-MMU ASI	D-MMU ASI	VA{63:0}	Access	Register or Operation Name
21 ₁₆		8 ₁₆	Read/Write	Primary Context 0 register
—	21 ₁₆	10 ₁₆	Read/Write	Secondary Context 0 register
21 ₁₆		28 ₁₆	Read/Write	Primary Context 0 register (no Context 1 update)
—	21 ₁₆	30 ₁₆	Read/Write	Secondary Context 0 register (no Context 1 update)
21 ₁₆		108 ₁₆	Read/Write	Primary Context 1 register
—	21 ₁₆	110 ₁₆	Read/Write	Secondary Context 1 register

13.7.2 Context Registers

SPARC M7 supports a pair of primary and a pair of secondary context registers per strand, which are shared by the I- and D-MMUs. Primary Context 0 and Primary Context 1 are the primary context registers, and a TLB entry for a translating primary ASI can match the context field with either Primary Context 0 or Primary Context 1 to produce a TLB hit. Secondary Context 0 and Secondary Context 1 are the secondary context registers, and a TLB entry for a translating secondary ASI can match the context field with either Secondary Context 0 or Secondary Context 1 to produce a TLB hit.

Compatibility Note To maintain backward compatibility with software designed for a single primary and single secondary context register, writes to Primary (Secondary) Context 0 Register also update Primary (Secondary) Context 1 Register when using the original ASI and VA for the context registers (ASI 21₁₆, VA 8₁₆ and 0x10₁₆).

The Primary Context 0 and Primary Context 1 registers are defined as shown in FIGURE 13-2, where `pcontext` is the context value for the primary address space. ASI 21₁₆, VA 0x8₁₆ provides backward compatibility for software that does not use Primary Context 1; this register updates both Primary Context 0 and Primary Context 1. ASI 21₁₆, VA 0x28₁₆ updates only Primary Context 0, and leaves Primary Context 1 unaltered.




FIGURE 13-2 Primary Context 0/1 registers: ASI 21₁₆, VA 8₁₆; ASI 21₁₆, VA 28₁₆; and ASI 21₁₆, VA 108₁₆

The Secondary Context 0 and Secondary Context 1 Registers are defined in FIGURE 13-3, where `scontext` is the context value for the secondary address space. ASI 21₁₆, VA 0x10₁₆ provides backward compatibility for software that does not use Secondary Context 1; this register updates both Secondary Context 0 and Primary Context 1. ASI 21₁₆, VA 0x30₁₆ updates only Secondary Context 0, and leaves Secondary Context 1 unaltered.




FIGURE 13-3 Secondary Context 0/1 Registers: ASI 21₁₆, VA 10₁₆; ASI 21₁₆, VA 30₁₆; and 21₁₆, VA 110₁₆

The contents of the Nucleus Context register are hardwired to the value zero:




FIGURE 13-4 Nucleus Context Register

Programming Guidelines

A.1 Multithreading

In SPARC M7, each physical core contains eight strands. Each strand has a full set of architected state registers and appears to software as a complete processor¹. In general, each of the 8 strands share the execution pipeline including the instruction, data, and L2 caches, branch predictor, out-of-order scheduling, execution pipelines, and retirement mechanisms. The pipeline is both horizontally and vertically threaded. It is vertically threaded since instructions from different strands can be in adjacent pipeline stages. It is horizontally threaded where parallelism allows. For example, each cycle the Pick unit may pick one instruction from one thread, and another instruction from a different thread, to be issued to independent execution units. SPARC M7 utilizes advanced branch prediction, dual instruction issue, out-of-order execution with up to 128 instructions in flight using a reorder buffer, hardware prefetching of instruction and data cache misses, and seamless hardware thread switching to provide high per-thread performance as well as high throughput. The pipeline is partitioned into several major subsections: instruction fetch, select/decode/rename, pick/issue/execute, and commit, each of which are mostly independent of one another.

A.1.1 Instruction fetch

Each cycle an arbiter chooses one strand for instruction fetching. The least-recently-fetched strand among the strands which are ready for fetching is the one chosen. A strand may not be ready for fetching due to instruction cache misses, instruction buffer full conditions, or other reasons. Once selected for fetch, up to four instructions may be fetched from the instruction cache and placed in per-strand instruction buffers. Instruction fetching occupies the first few stages of the pipeline. Instruction fetching is decoupled from the rest of the pipeline by the Select stage.

A.1.2 Select/Decode/Rename

In the same fashion that instruction fetch chooses a strand for fetching, Select chooses a strand for decoding, renaming, and transfer to the Pick unit. Each cycle, in parallel with and independent of instruction fetch, Select determines which strand, among the ready strands, is the least-recently selected. A strand may not be ready due to per-strand wait conditions, such as an empty instruction buffer or a post-synchronizing² instruction pending, or due to pipeline-wide resource constraints, such as a lack of reorder buffer entries.

Select then reads up to 2 instructions per cycle from that strand's instruction buffers, and decodes and renames the instructions. As it decodes the instructions it identifies any intra-strand dependencies upon prior instructions, and enforces these dependencies until the instructions are sent to the Pick

¹ Certain state registers are shared across strands to conserve hardware resources. These shared registers will (eventually) be listed in this Appendix.

² A post-synchronizing instruction stalls instruction issue for the strand after issuing the post-synchronizing instruction until the instruction commits. Instructions which are post-synchronizing are listed in Section TABLE A-1, *SPARC M7 Instruction Latencies*, on page 153 below.

unit and written into the pick queue. Decode also assigns instructions to “slots”. There are 2 primary slots. Slot 0 is reserved for integer and load/store instructions. Slot 1 is reserved for integer, floating-point, graphics, cryptographic, and control transfer instructions. There is a third auxiliary slot, slot 2, which is reserved for store data operations.

A.1.3 Pick/Issue/Execute

Pick selects up to 2 instructions per cycle (an additional store data operation may also be picked) without regard to strand ID from a 36-entry out-of-order scheduler termed the pick queue. Instructions are written with a relative age in mind, so the pick queue picks the oldest ready instruction within a slot. An instruction is ready when all of its source data is available. Only one instruction can be picked for each slot each cycle. There are never any inter-strand instruction dependencies. As Pick issues instructions, pick queue entries are reclaimed, and made available for use by subsequent instructions coming from Select/decode/rename.

As Pick issues instructions to the execution units, the instructions execute in one of several functional units. There are 2 integer units, a floating-point and graphics unit, and a load/store unit. Each of these units has independent pipelines and operates in parallel with other execution units. When instructions finish execution, they report their completion status to the Commit unit.

A.1.4 Commit

Commit utilizes a 128-entry reorder buffer to hold completion status and other per-instruction information. Instructions commit once their completion status is available. Instructions which cause an exception complete, but do not commit. Instead, they trap, and the thread begins fetching instructions from the trap handler. Similarly, if a branch misprediction occurs, instruction fetching resumes from the correct path once the branch predictor has been updated, and execution resumes once all instructions prior to the mispredicted branch commit.

Commit is threaded and each cycle attempts to commit instructions from the least-recently-committed thread among the threads which are ready-to-commit.

A.1.5 Context Switching Between Strands

Since context switching is built into the SPARC M7 pipeline (via the instruction fetch, select/decode/rename, pick/issue/execute, and commit blocks), strands can be switched each cycle with no pipeline stall penalty.

A.1.6 Synchronization

Certain instructions require the pipeline to synchronize. One type of synchronization, post-synchronizing or post-sync'ing, puts the strand in a wait state at Select. The strand remains in a wait state, and subsequent instructions are not selected, decoded, or renamed until the post-sync clears. This is resolved by the commit of the post-sync'ing instruction.

A.2 Optimizing for Single-Threaded Performance or Throughput

Section 1.3.1.1, *Single-threaded and multi-threaded performance*, on page 12 describes some aspects of optimizing for single-threaded and/or multi-threaded performance.

A.3 Instruction Latency

TABLE A-1 lists the minimum single-strand instruction latencies for SPARC M7. When multiple strands are executing, some or much of the additional latency for multicycle instructions will be overlapped with execution of the additional strands.

A pre-sync'ing instruction waits at Pick for all prior instructions from the strand to commit before being picked; therefore these instructions have a variable latency, whose minimum is listed in TABLE A-1. A post-sync'ing instruction causes a flush after the instruction commits. Loads have a 5-cycle load-use delay (4 cycles need to be filled but out-of-order execution covers much of this latency in many cases).

TABLE A-1 SPARC M7 Instruction Latencies (1 of 9)

Instruction	Description	Latency	Post-sync	Notes
ADD (ADDcc)	Add (and modify condition codes)	1		
ADDC (ADDCcc)	Add with carry (and modify condition codes)	1		
ADDXC (ADDXCcc)	Add extended with carry (and modify condition codes)	1		
AES_DROUND01	AES decrypt round, columns 0 & 1	3 or 11 ¹		
AES_DROUND23	AES decrypt round, columns 2 & 3	3 or 11 ¹		
AES_DROUND01_LAST	AES decrypt last round, columns 0 & 1	3 or 11 ¹		
AES_DROUND23_LAST	AES decrypt last round, columns 2 & 3	3 or 11 ¹		
AES_EROUND01	AES encrypt round, columns 0 & 1	3 or 11 ¹		
AES_EROUND23	AES encrypt round, columns 2 & 3	3 or 11 ¹		
AES_EROUND01_LAST	AES encrypt last round, columns 0 & 1	3 or 11 ¹		
AES_EROUND23_LAST	AES encrypt last round, columns 2 & 3	3 or 11 ¹		
AES_KEXPAND0	AES key expansion without round constant	3 or 11 ¹		
AES_KEXPAND1	AES key expansion with round constant	3 or 11 ¹		
AES_KEXPAND2	AES key expansion without SBOX	3 or 11 ¹		
ALIGNADDRESS	Calculate address for misaligned data access	3 or 12 ¹		
ALIGNADDRESS_LITTLE	Calculate address for misaligned data access (little-endian)	3 or 12 ¹		
ALLCLEAN	Mark all windows as clean	1		breaks decode group

TABLE A-1 SPARC M7 Instruction Latencies (2 of 9)

Instruction	Description	Latency	Post-sync	Notes
AND (ANDcc)	Logical and (and modify condition codes)	1		
ANDN (ANDNcc)	Logical and not (and modify condition codes)	1		
ARRAY{8,16,32}	3-D address to blocked byte address conversion	12		
Bicc	Branch on integer condition codes	1		
BMASK	Write the GSR.mask field	3 or 12 ¹		
BPcc	Branch on integer condition codes with prediction	1		
BPr	Branch on contents of integer register with prediction	1		
BSHUFFLE	Permute bytes as specified by the GSR.mask field	3 or 11 ¹		
CALL	Call and link	1		
CAMELLIA_F	Camellia F operation	3 or 11 ¹		
CAMELLIA_FL	Camellia FL operation	3 or 11 ¹		
CAMELLIA_FLI	Camellia FLI operation	3 or 11 ¹		
CASA	Compare and swap word in alternate space	20-30		Done in L2 cache
CASXA	Compare and swap doubleword in alternate space	20-30		Done in L2 cache
CBcond	Compare and branch	1		
CMASK{8,16,32}	Create GSR.mask from SIMD operation result	3 or 12 ¹		
CRC32C	Two CRC32c operations	3 or 11 ¹		
DES_IP	DES initial permutation	3 or 11 ¹		
DES_IIP	DES inverse initial permutation	3 or 11 ¹		
DES_KEXPAND	DES key expansion	3 or 11 ¹		
DES_ROUND	DES round	3 or 11 ¹		
DONE	Return from trap	23		Causes flush and redirect to TNPC (23 cycle bubble)
EDGE{8,16,32}{L}{N}	Edge boundary processing {little-endian} {non-condition-code altering}	12		
FABS(s,d)	Floating-point absolute value	11		
FADD(s,d)	Floating-point add	11		
FALIGNDATAg	Perform data alignment for misaligned data	3 or 11 ¹		
FALIGNDATAi	Perform data alignment for misaligned data using integer register	3 or 11 ¹		
FANDNOT1{s,d}	Negated src1 and src2	3 or 11 ¹		
FANDNOT2{s,d}	src1 and negated src2	3 or 11 ¹		
FAND{s,d}	Logical and	3 or 11 ¹		
FBfcc	Branch on floating-point condition codes	1		
FBPfcc	Branch on floating-point condition codes with prediction	1		
FCHKSM16	16-bit partitioned checksum	11		
FCMP(s,d)	Floating-point compare	11		
FCMPE(s,d)	Floating-point compare (exception if unordered)	11		

TABLE A-1 SPARC M7 Instruction Latencies (3 of 9)

Instruction	Description	Latency	Post-sync	Notes
FDIV(s,d)	Floating-point divide	24 SP, 37 DP		
FEXPAND	Four 8-bit to 16-bit expand	11		
FHADD{s,d}	Floating-point add and halve	11		
FHSUB{s,d}	Floating-point subtract and halve	11		
FiTO(s,d)	Convert integer to floating-point	11		
FLCMP{s,d}	Lexicographic compare	11		
FLUSH	Flush instruction memory	27	Y	Flushes pipeline, 27 cycle bubble minimum
FLUSHW	Flush register windows	1		breaks decode group
FMADD{s,d}	Floating-point multiply-add single/double (fused)	11		
FMEAN16	16-bit partitioned average	11		
FMOV(s,d)	Floating-point move	11		
FMOV(s,d)cc	Move floating-point register if condition is satisfied	11		
FMOV(s,d)R	Move floating-point register if integer register contents satisfy condition	11		Cracked into 2 ops, breaks decode group
FMSUB{s,d}	Floating-point multiply-subtract single/double (fused)	11		
FMUL(s,d)	Floating-point multiply	11		
FMUL8SUx16	Signed upper 8- x 16-bit partitioned product of corresponding components	11		
FMUL8ULx16	Unsigned lower 8- x 16-bit partitioned product of corresponding components	11		
FMUL8x16	8- x 16-bit partitioned product of corresponding components	11		
FMUL8x16AL	Signed lower 8- x 16-bit lower α partitioned product of 4 components	11		
FMUL8x16AU	Signed upper 8- x 16-bit lower α partitioned product of 4 components	11		
FMULD8SUx16	Signed upper 8- x 16-bit multiply \rightarrow 32-bit partitioned product of components	11		
FMULD8ULx16	Unsigned lower 8- x 16-bit multiply \rightarrow 32-bit partitioned product of components	11		
FNADD(s,d)	Floating-point add and negate	11		
FNAND{s,d}	Logical nand	3 or 11 ¹		
FNEG(s,d)	Floating-point negate	11		
FNHADD{s,d}	Floating-point add and halve, then negate	11		
FNMADD{s,d}	Floating-point add and negate	11		
FNMSUB{s,d}	Floating-point negative multiply-subtract single/double (fused)	11		
FNMUL{s,d}	Floating-point multiply and negate	11		
FNOR{s,d}	Logical nor	3 or 11 ¹		

TABLE A-1 SPARC M7 Instruction Latencies (4 of 9)

Instruction	Description	Latency	Post-sync	Notes
FNOT1{s,d}	Negate (1's complement) src1	3 or 11 ¹		
FNOT2{s,d}	Negate (1's complement) src2	3 or 11 ¹		
FNsMULd	Floating-point multiply and negate	11		
FONE{s,d}	One fill	3 or 11 ¹		
FORNOT1{s,d}	Negated src1 or src2	3 or 11 ¹		
FORNOT2{s,d}	src1 or negated src2	3 or 11 ¹		
FOR{s,d}	Logical or	3 or 11 ¹		
FPACKFIX	Two 32-bit to 16-bit fixed pack	11		
FPACK{16,32}	Four 16-bit/two 32-bit pixel pack	11		
FPADD8	Eight 8-bit partitioned add	3 or 11 ¹		
FPADD{16,32}{s}	Four 16-bit/two 32-bit partitioned add	3 or 11 ¹		
FPADD64	Fixed-point partitioned add	3 or 11 ¹		
FPADD{U}S8	Fixed-point partitioned add	3 or 11 ¹		
FPADDS{16,32}{s}	Fixed-point partitioned add	3 or 11 ¹		
FPADDUS16	Fixed-point partitioned add	3 or 11 ¹		
FPCMPEQ{16,32}	Four 16-bit / two 32-bit compare: set integer dest if src1 = src2	3 or 12 ¹		
FPCMPGT{8,16,32}	Eight 8-bit / four 16-bit / two 32-bit compare: set integer dest if src1 > src2	3 or 12 ¹		
FPCMPLE{8,16,32}	Eight 8-bit / four 16-bit / two 32-bit compare: set integer dest if src1 ≤ src2	3 or 12 ¹		
FPCMPNE{16,32}	Four 16-bit / two 32-bit compare: set integer dest if src1 ≠ src2	3 or 12 ¹		
FPCMPU {GT,LE,NE,EQ}8	Compare 8-bit unsigned fixed-point values	3 or 12 ¹		
FPCMPU {GT,LE}{16,32}	Compare four 16-bit/two 32-bit unsigned fixed-point values	3 or 12 ¹		
FPMADDX	Unsigned integer multiply-add	11		
FPMADDXHI	Unsigned integer multiply-add, return high-order 64 bits of result	11		
FPMAX{U}{8,16,32}	Partitioned integer maximum	3 or 11 ¹		
FPMERGE	Two 32-bit to 64-bit fixed merge	11		
FPMIN{U}{8,16,32}	Partitioned integer minimum	3 or 11 ¹		
FPSUB8	Eight 8-bit partitioned subtract	3 or 11 ¹		
FPSUB{16,32}{s}	Four 16-bit/two 32-bit partitioned subtract	3 or 11 ¹		
FPSUB64	Fixed-point partitioned subtract, 64-bit	3 or 11 ¹		
FPSUB{U}S8	Fixed-point partitioned subtract	3 or 11 ¹		
FPSUBS{16,32}{s}	Fixed-point partitioned subtract	3 or 11 ¹		
FPSUBUS16	Fixed-point partitioned subtract	3 or 11 ¹		
FSL{16,32}	16- or 32-bit partitioned shift, left (old mnemonic FSHL)	11		
FSLAS{16,32}	16- or 32-bit partitioned shift, left or right (old mnemonic FSHLAS)	11		

TABLE A-1 SPARC M7 Instruction Latencies (5 of 9)

Instruction	Description	Latency	Post-sync	Notes
FSRA{16,32}	16- or 32-bit partitioned shift, left or right (old mnemonic FSHRA)	11		
FSRL{16,32}	16- or 32-bit partitioned shift, left or right (old mnemonic FSHRL)	11		
FsMULd	Floating-point multiply single to double	11		
FSQRT{s,d}	Floating-point square root	24 SP, 37 DP		
FSRC1{s,d}	Copy src1	3 or 11 ¹		
FSRC2d	Copy src2 double precision	2		
FSRC2s	Copy src2 single precision	3 or 11 ¹		
F(s,d)TO(s,d)	Convert between floating-point formats	11		
F(s,d)TOi	Convert floating point to integer	11		
F(s,d)TOx	Convert floating point to 64-bit integer	11		
FSUB(s,d)	Floating-point subtract	11		
FXNOR{s,d}	Logical xnor	3 or 11 ¹		
FXOR{s,d}	Logical xor	3 or 11 ¹		
FxTO(s,d)	Convert 64-bit integer to floating-point	11		
FZERO{s}	Zero fill (single precision)	3 or 11 ¹		
ILLTRAP	Illegal instruction	23		
INVALW	Mark all windows as CANSAVE	1		breaks decode group
JMPL	Jump and link	1		
LDBLOCKF	64-byte block load	8		
LDD	Load doubleword	1		
LDDA	Load doubleword from alternate space	1		
LDDF	Load double floating-point	1		
LDDFA	Load double floating-point from alternate space	1		
LDF	Load floating-point	1		
LDFA	Load floating-point from alternate space	1		
LDFSR	Load floating-point state register lower	variable	Y	
LDSB	Load signed byte	1		
LDSBA	Load signed byte from alternate space	1		
LDSH	Load signed halfword	1		
LDSHA	Load signed halfword from alternate space	1		
LDSTUB	Load-store unsigned byte	20-30		Done in L2 cache
LDSTUBA	Load-store unsigned byte in alternate space	20-30		Done in L2 cache
LDSW	Load signed word	1		
LDSWA	Load signed word from alternate space	1		
LDTW	Load twin word	2		breaks decode group

TABLE A-1 SPARC M7 Instruction Latencies (6 of 9)

Instruction	Description	Latency	Post-sync	Notes
LDTWA	Load twin word	2		breaks decode group
LDUB	Load unsigned byte	1		
LDUBA	Load unsigned byte from alternate space	1		
LDUH	Load unsigned halfword	1		
LDUHA	Load unsigned halfword from alternate space	1		
LDUW	Load unsigned word	1		
LDUWA	Load unsigned word from alternate space	1		
LDX	Load extended	1		
LDXA	Load extended from alternate space	variable if from nontranslating ASI, else 1		
LDXEFSR	Load extended floating-point state register	variable	Y	
LDXFSR	Load extended floating-point state register	variable	Y	
LZCNT	Leading zero count on 64-bit integer register	12		
MD5	MD5 hash	192	Y	
MEMBAR	Memory barrier	variable		membar #sync is post-sync'ing; other membar forms are not
MONTMUL	Montgomery multiplication	variable	Y	pre-sync
MONTSQR	Montgomery squaring	variable	Y	pre-sync
MOVcc	Move integer register if condition is satisfied	1		
MOVdTOx	Move floating-point register to integer register	2		
MOVr	Move integer register on contents of integer register	1		breaks decode group
MOVsTO{u,s}w	Move floating-point register to integer register	3 or 12 ¹		
MOVwTOs	Move integer register to floating-point register	3 or 11 ¹		
MOVxTOd	Move integer register to floating-point register	1		
MPMUL	Multiple-precision multiplication	variable	Y	pre-sync
MULScc	Multiply step (and modify condition codes)	12		pre-sync
MULX	Multiply 64-bit integers	12		
NOP	No operation	1		
NORMALW	Mark other windows as restorable	1		breaks decode group
OR (ORcc)	Inclusive-or (and modify condition codes)	1		
ORN (ORNcc)	Inclusive-or not (and modify condition codes)	1		

TABLE A-1 SPARC M7 Instruction Latencies (7 of 9)

Instruction	Description	Latency	Post-sync	Notes
OTHERW	Mark restorable windows as other	1		breaks decode group
PDIST	Distance between eight 8-bit components	11		
PDISTN	Pixel component distance	12		
POPC	Population count	12		
PREFETCH	Prefetch data	1		
PREFETCHA	Prefetch data from alternate space	1		
RDASI	Read ASI register	variable	Y	
RDASR	Read ancillary state register	variable	Y	
RDCCR	Read condition codes register	variable	Y	
RDCFR	Read compatibility feature register	variable		
RDFPRS	Read floating-point registers state register	variable	Y	
RDPC	Read PROGRAM COUNTER	2		
RDPR	Read privileged register	variable	Y	
RTICK	Read TICK register	variable	Y	
RESTORE	Restore caller's window	1		breaks decode group
RESTORED	Window has been restored	1		breaks decode group
RETRY	Return from trap and retry	23		Causes flush and redirect to TPC (23 cycle bubble)
RETURN	Return	1		breaks decode group
SAVE	Save caller's window	1		breaks decode group
SAVED	Window has been saved	1		breaks decode group
SDIV (SDIVcc)	32-bit signed integer divide (and modify condition codes)	42-61		pre-sync
SDIVX{i}	64-bit signed integer divide	26-44		
SETHI	Set high 22 bits of low word of integer register	1		
SHA1	SHA-1 hash	226	Y	pre-sync
SHA256	SHA-256 hash	194	Y	pre-sync
SHA512	SHA-512 hash	242	Y	pre-sync
SIAM	Set interval arithmetic mode	1		
SLL	Shift left logical	1		
SLLX	Shift left logical, extended	1		
SMUL (SMULcc)	Signed integer multiply (and modify condition codes)	12		
SRA	Shift right arithmetic	1		
SRAX	Shift right arithmetic, extended	1		
SRL	Shift right logical	1		

TABLE A-1 SPARC M7 Instruction Latencies (8 of 9)

Instruction	Description	Latency	Post-sync	Notes
SRLX	Shift right logical, extended	1		
STB	Store byte	1		
STBA	Store byte into alternate space	1		
STBAR	Store barrier	variable		
STBLOCKF	64-byte block store	8		
STD	Store doubleword	1		
STDA	Store doubleword into alternate space	1		
STDF	Store double floating-point	1		
STDFA	Store double floating-point into alternate space	1		
STF	Store floating-point	1		
STFA	Store floating-point into alternate space	1		
STFSR	Store floating-point state register	variable	Y	
STH	Store halfword	1		
STHA	Store halfword into alternate space	1		
STPARTIALF	Eight 8-bit/4 16-bit/2 32-bit partial stores	1		
STTW	Store twin words	2		
STTWA	Store twin words into alternate space	2		
STW	Store word	1		
STWA	Store word into alternate space	1		
STX	Store extended	1		
STXA	Store extended into alternate space	variable if from nontrans lating ASI, else 1	depends upon ASI	
STXF SR	Store extended floating-point state register	variable	Y	pre-sync
SUB (SUBcc)	Subtract (and modify condition codes)	1		
SUBC (SUBCcc)	Subtract with carry (and modify condition codes)	1		
SUBXC (SUBXCcc)	Subtract extended with carry (and modify condition codes)	1		
SWAP	Swap integer register with memory	20-30		Done in L2 cache
SWAPA	Swap integer register with memory in alternate space	20-30		Done in L2 cache
TADDcc (TADDccTV)	Tagged add and modify condition codes (trap on overflow)	1		
Tcc	Trap on integer condition codes (with 8-bit sw_trap_number, if bit 7 is set, trap to hyperprivileged)	1 if no trap or 23 if trap taken		
TSUBcc (TSUBccTV)	Tagged subtract and modify condition codes (trap on overflow)	1		
UDIV (UDIVcc)	Unsigned integer divide (and modify condition codes)	42-61		pre-sync
UDIVX	64-bit unsigned integer divide	26-44		

TABLE A-1 SPARC M7 Instruction Latencies (9 of 9)

Instruction	Description	Latency	Post-sync	Notes
UMUL (UMULcc)	Unsigned integer multiply (and modify condition codes)	12		
UMULXHI	Unsigned 64 x 64 multiply, returning upper 64 product bits	12		
WRASI	Write ASI register	variable	Y	
WRASR	Write ancillary state register	variable	Y	
WRCCR	Write condition codes register	variable	Y	
WRFPRS	Write floating-point registers state register	variable	Y	
WRPR	Write privileged register	variable	Y	
XMONTMUL	XOR Montgomery multiplication	variable	Y	pre-sync
XMONTSQR	XOR Montgomery squaring	variable	Y	pre-sync
XMPMUL	XOR multiple-precision multiplication	variable	Y	pre-sync
XMULX{HI}	XOR multiply	12		
XNOR (XNORcc)	Exclusive- nor (and modify condition codes)	1		
XOR (XORcc)	Exclusive- or (and modify condition codes)	1		

1. Latency is 3 cycles only if the consumer of the operations result is also capable of 3 cycle latency.

IEEE 754 Floating-Point Support

SPARC M7 conforms to Oracle SPARC Architecture 2015 and the corresponding IEEE Std 754-1985 Requirements chapter.

Note | SPARC M7 detects tininess before rounding.

B.1 Special Operand and Result Handling

The SPARC M7 FGU provides full hardware support for subnormal operands and results for all instructions. SPARC M7 never generates an unfinished_FPop trap type. SPARC M7 does not implement a non-standard floating-point mode. The NS bit of the FSR is always read as 0, and writes to it are ignored.

Differences Between SPARC M7 and SPARC M6

This chapter describes the differences between the earlier SPARC M6 and SPARC M7. A summary of the differences is provided in the table below.

Area	vs SPARC M6	Description
Architecture and Microarchitecture	Different	Section C.1
Data Format	Same	
Registers	Different	Section C.2
Instruction Format	Same	
Instruction Definitions	Same	
Traps	Different	Section C.3
Interrupt Handling	Different	Section C.4
Memory Models	Same	
Address Spaces & ASIs	Different	Section C.2
Performance Measurement	Different	Section C.7
Crypto	Same	
MMU	Different	
Clocks & Reset	Different	Section C.8
CMT	Different	Section C.9
Error Handling	Different	Section C.10
Power Management	Different	
Configuration	Different	Section C.12
Diagnostic	Different	Section C.13
HW Debug	Different	Section C.14

C.1 Architectural and Microarchitectural Differences

SPARC M7 modifies the SPARC core from SPARC M6 and SPARC M6, the unified L3 cache is shared among 16 cores (vs six in SPARC M6), and all the SPARC M7 SOC components are either re-designed or modified from SPARC M6. <under construction>

The architectural differences in the core are:

- Virtual, real, addresses is increased by two bits (VA is 54 bits, RA are 50 bits)
- Kasumi cipher instructions are no longer supported
- XOR versions of MPMUL, MONTMUL, and MONTSQR are now supported
- The core implements VA masking support for data accesses

The microarchitectural changes in the core are

- Instruction cache line size is increased to 64B

SPARC M7 is capable of supporting up to 8 processors in a glue-less fashion and provides scalability ports for scaling beyond 8 processors.

For details, refer to the following chapters:

- For details of overall architectural and microarchitectural differences, see VT Basic Chapter.

C.2 Address Spaces and ASIs Differences

C.2.1 ASIs

Addressing of all preexisting ASIs in SPARC core (including L2) does not change from SPARC M6 to SPARC M7. SPARC M7 does add several new ASIs

See Address Spaces and ASIs Chapter for details.

Cache Coherency and Ordering

D.1 Cache and Memory Interactions

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:

- 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.

D.2 Cache Flushing

Data in the level-1 (read-only or writethrough) caches can be flushed by invalidating the entry in the cache (in a way that also leaves the L2 directory in a consistent state). Modified data in the level-2 and level-3 (writeback) caches 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. This is done with the FLUSH instruction, which just forces previous stores to complete to all affected caches.. Flushing the I-cache with ASI accesses (Section 20.6, *L1 I-Cache Diagnostic Access*, on page 1006) also works, because the L2 directory correctly handles the cases where the directory thinks the line is in the L1, but the L1 doesn't.
- **D-cache:** Flush is needed when a physical page is changed from (physically) cacheable to (physically) noncacheable. This is done with a displacement flush (*Displacement Flushing*, below), or with ASI accesses (see Section 20.7, *L1 D-Cache Diagnostic Access*, on page 1009), which work for similar reasons as for the I-cache.
- **L2 cache:** Flush is needed for stable storage. Examples of stable storage include battery-backed memory and transaction logs. The recommended way to perform this is by using PIO line flushes to L2I and L2D CSR space to flush given index/ways (see Section 20.16.2, *L2I Line Flush with Optional Retire* and Section 20.16.5, *L2D Line Flush with Optional Retire*). This mechanism can also be used to "retire" cache lines that have persistent errors. Flushing the L2 caches flushes the corresponding blocks from the I- and D-caches, because SPARC M7 maintains inclusion between the L2 and L1 caches

- **L3 cache:** Flush is needed for stable storage. Examples of stable storage include battery-backed memory and transaction logs. The recommended way to perform this is by using diagnostic writes to L3 CSR space to flush given index/ways (see Section 20.22, *L3 diagnostic access and CSRs, L3 Off*, on page 1055). This mechanism can also be used to "retire" cache lines that have persistent errors. Alternatively, this can be done by a displacement flush (see the next section). Flushing the L3 cache flushes the corresponding blocks from the I- and D-caches, and L2I and D cache because SPARC M7 maintains inclusion between the L3, L2 and L1 caches
- **Errors:** Flush is needed for error processing. Examples include (1) forcing UE data from a cache to memory, in order to convert it to NotData, or (2) using flushes to force memory (not cache) reads and writes, to diagnose a memory error, or (3) writing a line of good data and flushing it to memory, to overwrite a memory soft error.

D.2.1 Displacement Flushing

Cache flushing of the L3 cache or the D-cache 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.

The L2 caches do not support a direct mapped mode. Flushing of the L2 caches should be done with PIO line flushes.

D.2.2 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.

In SPARC M7, all memory accesses are cached in the L2 and L3 caches (as long as the caches are enabled). The `cp` bit in the TTE corresponding to the access controls whether the memory access will be cached in the primary caches (if `cp = 1`, the access is cached in the primary caches; if `cp = 0` the access is not cached in the primary caches). Atomic operations are always performed at the L2 cache.

D.2.3 Coherence Domains

Two types of memory operations are supported in SPARC M7: 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. SPARC M7 maintains TSO ordering, regardless of the cacheability of the accesses, relative to other access by processors.

See the *The SPARC Architecture Manual-Version 9* for more information about the SPARC V9 memory models.

On SPARC M7, a `MEMBAR #Lookaside` is effectively a NOP and is not needed for forcing order of stores vs. loads to noncacheable addresses.

D.2.3.1 Cacheable Accesses

Accesses that fall within the coherence domain are called cacheable accesses. They are implemented in SPARC M7 with the following properties:

- Data resides in real memory locations.
- They observe the 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 and L3 caches.
- The unit of coherence for the primary caches (coherence between multiple virtual processors) is the primary cache line size (32 bytes for the data cache, 64 bytes for the instruction cache), enforced by the L2 cache directories.

D.2.3.2 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 SPARC M7 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 are all strongly ordered with respect to other noncacheable accesses (regardless of the *e* bit). Speculative loads with the *e* bit set cause a *DAE_so_page* trap.

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

D.2.3.3 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 D-1 illustrates the issues involved in mixing cacheable and noncacheable accesses.

CODE EXAMPLE D-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 for RMO cacheable accesses.

Due to load and store buffers implemented in SPARC M7, CODE EXAMPLE D-1 may not work for RMO accesses without the MEMBARs shown in the program segment.

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

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

D.2.4 Memory Synchronization: MEMBAR and FLUSH

The MEMBAR (STBAR in SPARC V8) and FLUSH instructions provide for explicit control of memory ordering in program execution. MEMBAR has several variations; their implementations in SPARC M7 are described below. 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.

D.2.4.1 MEMBAR #LoadLoad

All TSO loads on SPARC M7 are implicitly ordered so no MEMBAR #LoadLoad is required. Block loads are RMO and require an intervening MEMBAR #LoadLoad to ensure ordering with respect to prior or subsequent loads.

D.2.4.2 MEMBAR #StoreLoad

MEMBAR #StoreLoad forces all loads after the MEMBAR to wait until all stores before the MEMBAR have reached global visibility. All TSO loads and stores on SPARC M7 are implicitly ordered so no MEMBAR #StoreLoad is required. Block loads, block stores, and block initializing stores are RMO and require MEMBAR #StoreLoad to guarantee ordering.

D.2.4.3 MEMBAR #LoadStore

All loads and stores on SPARC M7 commit in order. Thus, MEMBAR #LoadStore is treated as a NOP on SPARC M7

D.2.4.4 MEMBAR #StoreStore and STBAR

TSO stores on SPARC M7 are implicitly ordered and no Membar #StoreStore is required. Block stores and block initializing stores are not implicitly ordered and require Membar #StoreStore (or stronger) to guarantee ordering on SPARC M7.

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

D.2.4.5 MEMBAR #Lookaside

Loads and stores to noncacheable addresses are “self-synchronizing” on SPARC M7. Thus MEMBAR #Lookaside is treated as a NOP on SPARC M7.

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

D.2.4.6 MEMBAR #MemIssue

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 noncacheable loads following cacheable stores. For example, a cacheable store must be followed by a MEMBAR #MemIssue before subsequent noncacheable loads; this ensures that the store reaches global visibility (as viewed by other strands) before the noncacheable load after the MEMBAR. All other ordering cases of noncacheable vs. cacheable accesses are implicitly order by the hardware.

SPARC M7 implements Membar #MemIssue identically to Membar #StoreLoad.

D.2.4.7 MEMBAR #Sync (Issue Barrier)

Membar #Sync 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.

D.2.4.8 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. Dynamic optimizers, 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 SPARC M7, 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.

SPARC M7 implements FLUSH identically to Membar #Sync.

D.2.5 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 *DAE_invalid_ASI* trap. TABLE D-1 lists the ASIs that support atomic accesses.

TABLE D-1 ASIs That Support SWAP, LDSTUB, and CAS

ASI Name
ASI_NUCLEUS{ _LITTLE }
ASI_AS_IF_USER_PRIMARY{ _LITTLE }
ASI_AS_IF_USER_SECONDARY{ _LITTLE }
ASI_PRIMARY{ _LITTLE }
ASI_SECONDARY{ _LITTLE }
ASI_REAL{ _LITTLE }

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

| For all atomics, allocation is done to the L2 cache only and will invalidate the L1s.

D.2.5.1 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.

D.2.5.2 LDSTUB Instruction

LDSTUB behaves like SWAP, except that it loads a byte from memory into an integer register and atomically writes all 1's (`FF16`) into the addressed byte.

D.2.5.3 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.

D.2.6 Nonfaulting Load

A nonfaulting load behaves like a normal load, except that

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

Nonfaulting loads are issued with `ASI_PRIMARY_NO_FAULT{ _LITTLE}` or `ASI_SECONDARY_NO_FAULT{ _LITTLE}`. A store with a `NO_FAULT` ASI causes a *DAE_invalid_ASI* trap.

When a nonfaulting load encounters a TLB miss, the operating system should attempt to translate the page. If the translation results in an error (for example, address out of range), a 0 is returned and the load completes silently.

Typically, optimizers use nonfaulting loads to move loads before conditional control structures that guard their use. This technique potentially increases the distance between a load of data and the first use of that data, to hide latency; it allows for more flexibility in code scheduling. It also allows for improved performance in certain algorithms by removing address checking from the critical code path.

For example, when following a linked list, nonfaulting loads allow the null pointer to be accessed safely in a read-ahead fashion if the operating system can ensure that the page at virtual address `016` is accessed with no penalty. The `nfo` (nonfault access only) bit in the MMU marks pages that are mapped for safe access by nonfaulting loads but can still cause a trap by other, normal accesses. This allows programmers to trap on wild pointer references (many programmers count on an exception being generated when accessing address `016` to debug code) while benefitting from the acceleration of nonfaulting access in debugged library routines.

D.3 L1 I-Cache

The L1 Instruction cache is 16 Kbytes, physically tagged and indexed, with 64-byte lines, and 4-way associative with true LRU replacement. The format used to index the cache is shown in TABLE D-2.

TABLE D-2 L1 Instruction Cache Addressing

Bit	Field	Description
49:12	tag	Tag for cache line.
11:6	set	Selects cache set containing the cache line.
5:2	instr	Selects 32-bit instruction in cache line.
1:0	—	Always 0 for access to 32-bit instructions.

D.3.1 LRU Replacement Algorithm

The I-cache replacement algorithm is true least-recently-used (LRU). Six bits are maintained for each cache index.

D.3.2 Direct-Mapped Mode

The I-cache direct-mapped mode works by forcing all replacements to the “way” identified by bits [13:11] of the virtual address. Since lines already present are not affected but only new lines brought into the cache are affected, it is safe to turn on (or off) the direct-mapped mode at any time.

D.3.3 I-Cache Disable

Clearing the I-cache enable bit stops all accesses to the I-cache for that strand. All fetches will miss, and the returned data will not fill the I-cache. Invalidates will still be serviced while the I-cache is disabled.

D.4 L1 D-Cache

The L1 Data cache is 16 Kbytes, writethrough, physically tagged and indexed, with 32-byte lines, and 4-way associative with true LRU replacement. The format used to index the cache is shown in TABLE D-3.

TABLE D-3 L1 Data Cache Addressing

Bit	Field	Description
49:12	tag	Tag for cache line.
11:5	set	Selects cache set containing the cache line.
4:0	data	Selects data byte(s) in cache line.

D.4.1 LRU Replacement Algorithm

The D-cache replacement algorithm is true least-recently-used (LRU). Six bits are maintained for each cache index.

D.4.2 Direct-Mapped Mode

The D-cache direct-mapped mode works by changing the replacement algorithm from LRU to instead use two bits of index (address[12:11]) to select the “way.” Since lines already present are not affected but only new lines brought into the cache are affected, it is safe to turn on (or off) the direct-mapped mode at any time.

Note that if the D-cache is in direct-mapped mode, and a parity error occurs, the way replaced will be the way which experienced the parity error. This overrides the index selected by the address in direct-mapped mode.

D.4.3 D-Cache Disable

The D-cache enable bit works by forcing all accesses to miss in the D-cache, and all misses are nonallocating. Stores that hit in the L1 will be performed in the L2, then update the L1 (as normal).

D.5 L2 Instruction Cache

The L2 Instruction cache is 256K, and is shared among 4 physical cores. It is 2 way banked (on PA[6]), and is 8 way set associative. The L2 instruction cache is physically tagged and indexed, with 64B lines. Replacement is NRU (Not Recently Used) replacement. The format used to index the full cache is shown in TABLE D-4.

TABLE D-4 L2 Cache Addressing (8 banks) Not Updated for M7

Bit	Field	Description
48:15	tag	Tag for cache line.
14:6	index	Selects cache set containing the cache line. Bit 6 selects the cache bank.
5:0	data	Selects data byte(s) in the cache line.

D.5.1 NRU Replacement Algorithm

A used-bit scheme is used to implement an NRU (Not Recently Used) replacement. The used bit is set each time a cache line is accessed or when initially fetched from memory. If setting the used-bit causes all used bits (at an index) to be set, the remaining used bits are cleared instead.

In addition, each line has a lock bit, which is set while a line is allocated for replacement as a result of a cache miss. The lock bit gets cleared when the location is filled with memory data. Any line that has the lock bit set is ineligible for replacement.

The next replacement way is computed by first checking the Valid and Lock bits. Starting from way 0 and searching to way 7, the first invalid and not locked way will be selected as the next replacement way. If all ways are valid, then starting from way 0 and searching to way 7, the first way with Used=0 and Lock=0 will be selected as the next replacement way. If all Used bits are set (which means all Locked bits are set), then a replacement way cannot be determined and any access which misses and requires a replacement way will be marked non-allocating when sent to the L3. The result from the L3 will be bypassed to the core, but the line will not fill in the L2I.

D.5.1.1 Mapping Out Lines

It is possible to "map out" individual cache lines that have gone bad, e.g. get too many errors, by setting the line's Lock bit to 1, and clearing its Valid bit to 0. This marks the line as "busy", but in a state where it will never become unbusy. These lines are never considered for replacement. This can be accomplished with the appropriate variant of line flush.

D.5.2 Directory Coherence

The L2 instruction cache has a directory of all L1 instruction cache lines, implemented as reverse directory. This means that for each location in the L1 I-cache, the directory knows the corresponding location in the L2 instruction cache. (Because SPARC M7 maintains inclusion between the L2 and L1 caches, a line which exists in an L1 cache will always exist in its connected L2 cache.) When the L1 requests a line from the L2, the virtual processor specifies whether the line will be allocated (put into the cache), and which "way" it will go into.

The L2 I-cache can issue invalidates to any/all of the cores simultaneously. An invalidation is issued to the L1 any time a line is invalidated or locked in the L2I. The invalidate transaction includes only index and way; it does not include the address.

For special cases, primarily parity errors, the directory will get "conservatively" out-of-sync, which means the directory thinks a line exists in the L1 but it doesn't. This is not a problem, as the only consequence is a possible invalidation for a line which is already invalid.

Since the L2 directory can handle the above cases, just invalidating an L1 line is safe, and can be used to flush out L1 lines.

D.5.3 L2I Cache Disable

The L2I cache disable is described in Section 20.16.1.1, *L2I Off Mode*.

D.6 L2 Data Cache

The L2 data cache is 256K, and is shared among 2 physical cores. It is 2 way banked (on PA[6]), and is 8 way set associative. The L2D cache is writeback, physically tagged and indexed, with 64B lines. Replacement is NRU (Not Recently Used) replacement. The format used to index the full cache is shown in TABLE D-5.

TABLE D-5 L2 Cache Addressing (8 banks) Not Updated for M7

Bit	Field	Description
48:15	tag	Tag for cache line.
14:6	index	Selects cache set containing the cache line. Bit 6 selects the cache bank.
5:0	data	Selects data byte(s) in the cache line.

D.6.1 NRU Replacement Algorithm

A used-bit scheme is used to implement an NRU (Not Recently Used) replacement. The used bit is set each time a cache line is accessed or when initially fetched from memory. If setting the used-bit causes all used bits (at an index) to be set, the remaining used bits are cleared instead.

In addition, each line has a lock bit, which is set while a line is allocated for replacement as a result of a cache miss. The lock bit gets cleared when the location is filled with memory data. Any line that has the lock bit set is ineligible for replacement.

The next replacement way is computed by first checking the Valid and Lock bits. Starting from way 0 and searching to way 7, the first invalid and not locked way will be selected as the next replacement way. If all ways are valid, then starting from way 0 and searching to way 7, the first way with Used=0 and Lock=0 will be selected as the next replacement way. If all Used bits are set (which means all Locked bits are set), then a replacement way cannot be determined and any access which misses and requires a replacement way will be inserted into the miss buffer to be replayed once a line at that index is unlocked.

D.6.1.1 Mapping Out Lines

It is possible to "map out" individual cache lines that have gone bad, e.g. get too many errors, by setting the line's Lock bit to 1, and clearing its Valid bit to 0. This marks the line as "busy", but in a state where it will never become unbusy. These lines are never considered for replacement.

D.6.2 Directory Coherence

The L2 data cache has a directory of all L1 data cache lines, implemented as reverse directory. This means that for each location in the L1 D-cache, the directory knows the corresponding location in the L2 data cache. (Because SPARC M7 maintains inclusion between the L2 and L1 caches, a line which exists in an L1 cache will always exist in its connected L2 cache.) When the L1 requests a line from the L2, the virtual processor specifies whether the line will be allocated (put into the cache), and which "way" it will go into.

The L2D can issue invalidates to any/all of the cores simultaneously. An invalidation is issued to the L1 any time a line is invalidated or locked in the L2D. In addition, if one of the cores connected to the L2D stores to a line, that corresponding line must be invalidated in the other core's D-cache. The invalidate transaction includes only index and way; it does not include the address.

■ For special cases, the directory will become "conservatively" out-of-sync, which means the directory thinks a line exists in the L1 but it doesn't. This is not a problem, as the only consequence is a possible invalidation for a line which is already invalid.

Since the L2 directory can handle the above cases, just invalidating an L1 line is safe, and can be used to flush out L1 lines.

D.6.3 L2 Cache Disable

The L2D cache disable is described in Section 20.16.7, *L2D Off Mode*.

Glossary

This chapter defines concepts and terminology unique to the SPARC M7 implementation. Definitions of terms common to all Oracle SPARC Architecture implementations may be found in the Definitions chapter of *Oracle SPARC Architecture 2015*.

- ALU** Arithmetic Logical Unit
- architectural state** Software-visible registers and memory (including caches).
- ARF** Architectural register file.
- blocking ASI** An ASI access that accesses its ASI register or array location once all older instructions in that strand have retired, no instructions in the other strand can issue, and the store queue, TSW, and LMB are all empty.
- branch outcome** A reference as 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.
- branch resolution** 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.
- branch target address** The address of the instruction to be executed if the branch is taken.
- commit** An instruction commits when it modifies architectural state.
- complex instruction** A complex instruction is an instruction that requires the creation of secondary “helper” instructions for normal operation, excluding trap conditions such as spill/fill traps (which use helpers). Refer to *Instruction Latency* on page 153 for a complete list of all complex instructions and their helper sequences.
- consistency** See **coherence**.
- CPU** Central Processing Unit. A synonym for **virtual processor**.
- CSR** Control Status register.
- FP** Floating point.
- L2C (or L2\$)** Level 2 cache.
- leaf procedure** 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.
- nonblocking ASI** A nonblocking ASI access 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.
- older instruction** Refers to the relative fetch order of instructions. Instruction *i* is older than instruction *j* if instruction *i* was fetched before instruction *j*. Data dependencies flow from older instructions to younger instructions, and an instruction can only be dependent upon older instructions.

one hot An n -bit binary signal is one hot if and only if $n - 1$ of the bits are each zero and a single bit is a 1.

quadlet

SIAM Set interval arithmetic mode instruction.

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

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

Bibliography

[contents of this appendix are TBD]

Index

A

Accumulated Exception (**aexc**) field of FSR register, 68, 69
Address Mask (**am**)
 field of PSTATE register, 53, 54, 80
address space identifier (ASI)
 identifying memory location, 51
ASI
 restricted, 80
 support for atomic instructions, 172
 usage, 54–60
ASI, *See* **address space identifier (ASI)**
ASI_AS_IF_USER_PRIMARY, 79
ASI_AS_IF_USER_SECONDARY, 80
ASI_NUCLEUS, 79
ASI_PRIMARY_NO_FAULT, 76, 79, 80
ASI_PRIMARY_NO_FAULT_LITTLE, 76, 79, 80
ASI_QUEUE registers, 47–49
ASI_REAL, 60
ASI_REAL_IO, 60
ASI_REAL_IO_LITTLE, 60
ASI_REAL_LITTLE, 60
ASI_SCRATCHPAD, 60
ASI_SECONDARY_NO_FAULT, 76, 79, 80
ASI_SECONDARY_NO_FAULT_LITTLE, 76, 79, 80
ASI_ST_BLK_INIT_PRIMARY, 41
ASI_ST_BLK_INIT_PRIMARY_LITTLE, 41
ASI_ST_BLK_INIT_SECONDARY, 41
ASI_ST_BLK_INIT_SECONDARY_LITTLE, 41
ASI_ST_BLKINIT_AS_IF_USER_PRIMARY, 41
ASI_ST_BLKINIT_AS_IF_USER_PRIMARY_LITTLE, 41
ASI_ST_BLKINIT_AS_IF_USER_SECONDARY, 41
ASI_ST_BLKINIT_AS_IF_USER_SECONDARY_LITTLE, 41
 1
ASI_ST_BLKINIT_NUCLEUS, 41
ASI_ST_BLKINIT_NUCLEUS_LITTLE, 41
ASI_STBI_AIUP, 41
ASI_STBI_AIUPL, 41
ASI_STBI_AIUS, 41
ASI_STBI_AIUS_L, 41
ASI_STBI_N, 41
ASI_STBI_NL, 41
ASI_STBI_P, 41
ASI_STBI_PL, 41
ASI_STBI_S, 41
ASI_STBI_SL, 41

atomic instructions, 172

B

block
 load instructions, 38, 41
 memory operations, 71
 store instructions, 38
block-initializing ASIs, 41
branch instruction, 54

C

cache flushing, when required, 167
cacheable in physically-indexed cache (**cp**) field of TTE, 76
caching
 TSB, 77
CALL instruction, 54
CANRESTORE register, 66
CANSAVE register, 66
clean window, 66
clean_window exception, 66
CLEANWIN register, 66
compatibility with SPARC V9
 terminology and concepts, 179
context
 field of TTE, 76
Current Exception (**cexc**) field of FSR register, 68, 69
CWP register, 66

D

DAE_invalid_ASI exception, 70, 85
DAE_invalid_asi exception, 54
DAE_privilege_violation exception, 77
DAE_so_page, 169
Dcache
 direct-mapped mode, 174
 disabling, 174
 displacement flush, 168
 flushing, 167
deferred
 trap, 65
Dirty Lower (**dl**) field of FPRS register, 68
Dirty Upper (**du**) field of FPRS register, 68

D-MMU, 79

E

endianness, 76

enhanced security environment, 66

errors

See also individual error entries

extended

instructions, 71

Extended instructions, 71

F

floating point

deferred trap queue (fq), 68, 69

exception handling, 67

trap type (ftt) field of FSR register, 69

Floating Point Condition Code (fcc)

0 (fcc0) field of FSR register, 68

3 (fcc3) field of FSR register, 68

field of FSR register in SPARC-V8, 68

Floating Point Registers State (FPRS) register, 68

FLUSH instruction, 70

fp_exception_ieee_754 exception, 68, 69

fp_exception_other exception, 69

G

global level register, *See* GL register

Graphics Status register, *See* GSR

H

hardware_error floating-point trap type, 68, 69

I

IAE_privilege_violation exception, 77

Icache

direct-mapped mode, 173

disabling, 173

flushing, 167

IEEE Std 754-1985, 68

IEEE support

infinity arithmetic, 163

normal operands/subnormal result, 163

IEEE_754_exception floating-point trap type, 69

illegal_instruction exception, 65, 68, 69, 70, 71

ILLTRAP instructions, 65

implementation-dependent instructions, *See* IMPDEP2A

instructions

instruction fetching

near VA (RA) hole, 53

instruction latencies, 153–161

instruction-level parallelism

history, 9

instruction-level parallelism, *See* ILP

integer

division, 66

multiplication, 66

register file, 66

internal registers, 80

interrupt

hardware delivery mechanism, 47

invalid_fp_register floating-point trap type, 68, 69

invert endianness, (ie) field of TTE, 76

ISA, *See* **instruction set architecture**

J

JMPL instruction, 54

jump and link, *See* JMPL instruction

L

L2 cache

configuration, 12

directory coherence, 175, 176–177

displacement flush, 168

flushing, 167, 168

instruction/data registers, 174–??, 176–??

latencies for instructions, 153–161

LDBLOCKF instruction, 38

LDD instruction, 70

LDDF_mem_address_not_aligned exception, 70

LDQF instruction, 70

LDQFA instruction, 70

LDXA instruction, 54

load

block, *See* **block load instructions**

short floating-point, *See* short floating-point load instructions

M

mem_address_not_aligned exception, 79, 85

MEMBAR #LoadLoad, 52

MEMBAR #Lookaside, 52

MEMBAR #MemIssue, 52, 170

MEMBAR #StoreLoad, 39, 40, 52

MEMBAR #StoreStore, 70

MEMBAR #Sync, 85

MEMBAR #Sync, 170

memory

cacheable and noncacheable accesses, 168

location identification, 51

model, 40

noncacheable accesses, 169

order between references, 52

ordering in program execution, 170–171

memory model supported, 51

memory models, 51

minimum single-strand instruction latencies, 153–161

MMU

requirements, compliance with SPARC V9, 84

N

N_REG_WINDOWS, 66

- nested traps
 - in SPARC-V9, 65
- No-Fault Only (nfo) field of TTE, 76, 80
- nonfaulting loads, 172
 - speculative, 78
- Non-Standard (ns) field of FSR register, 69
- Nucleus Context register, 86

O

- OTHERWIN register, 66
- out of range
 - virtual address, 53
 - virtual address, as target of JMWPL or RETURN, 54
 - virtual addresses, during STXA, 85

P

- page
 - size field of TTE, 77
 - size, encoding in TTE, 77
- partial store
 - instruction, 71
- Partial Store Order (PSO), 51
- pcontext field, 86
- PCR register
 - fields, 62
- performance instrumentation counter register, *See* PIC register
- physical core
 - components, 11
 - UltraSPARC T2 microarchitecture, 11
- PIC register
 - field description, 62
- precise traps, 65
- PREFETCHA instruction, 70
- Primary Context register, 86
- privileged
 - (p) field of TTE, 77
 - (priv) field of PSTATE register, 77, 78, 79
- privileged_action* exception
 - attempting access with restricted ASI, 51, 79, 80
- processor
 - memory model, 40
- processor interrupt level register, *See* PIL register
- processor state register, *See* PSTATE register
- processor states, *See* `execute_state`
- PSTATE register fields
 - ie
 - masking disrupting trap, 46
 - pef
 - See also* pef field of PSTATE register
- PTE (page table entry), *See* **translation table entry** (TTE)

Q

- quad-precision floating-point instructions, 67
- queue
 - Not Empty (qne) field of FSR register, 69

R

- RA hole, 53
- real page number (ra) field of TTE, 76
- Relaxed Memory Order (RMO), 51, 52
- reserved
 - fields in opcodes, 65
 - instructions, 65
- resumable_error* exception, 47
- RETURN instruction, 54
- RMO, *See* **relaxed memory order (RMO) memory model**
- Rounding Direction (rd) field of FSR register, 68

S

- SAVE instruction, 66
- scontext field, 86
- Secondary Context register, 86
- secure environment, 66
- self-modifying code, 70
- short floating point
 - load instruction, 71
 - store instruction, 71
- side effect
 - field of TTE, 76
- software
 - defined fields of TTE, 76
 - Translation Table, 70, 77
- software-defined field (soft) of TTE, 76
- SPARC V9
 - compliance with, 65
 - speculative load, 78
- STBLOCKF instruction, 38
- STD instruction, 70
- STDF_mem_address_not_aligned* exception, 70
- STQF instruction, 70
- STQFA instruction, 70
- STXA instruction, 54
- supervisor interrupt queues, 47

T

- TBA register, 54
- terminology for SPARC V9, definition of, 179
- thread-level parallelism
 - advantages, 10
 - background, 10
 - differences from instruction-level parallelism, 10
- thread-level parallelism, *See* **TLP**
- Throughput Computing, 9
- TNPC register, 54
- Total Store Order (TSO), 51, 52
- TPC register, 54
- Translation Table Entry *see* TTE
- Translation Table Entry, *See* TTE
- trap
 - mask behavior, ??-46
 - stack, 65
 - state registers, 65
- Trap Enable Mask (tem) field of FSR register, 68, 68, 69
- trap level register, *See* TL register

trap next program counter register, *See* TNPC register
trap program counter register, *See* TPC register
trap stack array, *See* **TSA**
trap state register, *See* TSTATE register
trap type register, *See* TT register
Trap-on-Event (toe) field of PCR register, 62
traps

See also exceptions *and* individual trap names

TSB
caching, 77
index to smallest, 76
in-memory, 70
organization, 77

TSO, *See* total store order (TSO) memory model

tstate, *See* **trap state** (TSTATE) register

TTE, 75

U

unimplemented instructions, 65

V

VA hole, 53

VA_tag field of TTE, 76

Valid (v) field of TTE, 76

Version (ver) field of FSR register, 69

virtual address

space *illustrated*, 53

Visual Instruction Set, *See* **VIS instructions**

W

window fill exception, *See also* *fill_n_normal* exception

window spill exception, *See also* *spill_n_normal* exception

writable (w) field of TTE, 77