

# UltraSPARC<sup>®</sup> III Cu

User's Manual



Version 2.2.1

January 2004



Copyright © 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

This distribution may include materials developed by third parties.

Sun, Sun Microsystems, the Sun logo, Java, Solaris, Sun Fireplane Interconnect, VIS, OpenBoot PROM, UltraSPARC III Cu and SPARC are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

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

Use of any spare or replacement CPUs is limited to repair or one-for-one replacement of CPUs in products exported in compliance with U.S. export laws. Use of CPUs as product upgrades unless authorized by the U.S. Government is strictly prohibited.

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



# Table of Contents

---

**Preface xxxvii**

**Acronyms and Definitions xliii**

---

## **Section I: Processor Introduction**

---

|           |                                       |          |
|-----------|---------------------------------------|----------|
| <b>1.</b> | <b>Processor Introduction .....</b>   | <b>1</b> |
| 1.1       | Overview .....                        | 1        |
| 1.2       | CPU Features .....                    | 2        |
| 1.3       | Cache Features .....                  | 2        |
| 1.4       | Technology .....                      | 2        |
| 1.5       | UltraSPARC III Cu Differences .....   | 3        |
| 1.5.1     | Bootbus Limitations .....             | 3        |
| 1.5.2     | Instruction Set Extensions .....      | 4        |
| 1.5.3     | Instruction Differences .....         | 5        |
| 1.5.4     | Memory Subsystem .....                | 6        |
| 1.5.5     | Interrupts .....                      | 9        |
| 1.5.6     | Address Space Size .....              | 9        |
| 1.5.7     | Registers .....                       | 9        |
| 1.5.8     | Non-Cacheable Store Compression ..... | 10       |

|           |   |           |
|-----------|---|-----------|
| 1.5.9     | Error Correction .....  | 11        |
| 1.5.10    | SRAM Protection and RAS Features .....                            | 11        |
| <b>2.</b> | <b>System Introduction .....</b>                                  | <b>13</b> |
| 2.1       | System Configurations .....                                       | 13        |
| 2.1.1     | Two-Processor Configuration with UltraSPARC III Cu .....          | 13        |
| 2.1.2     | Four-Processor Configuration with UltraSPARC III Cu .....         | 14        |
| 2.1.3     | Multiprocessor System with the UltraSPARC III Cu .....            | 16        |
| 2.1.4     | Very Large Multiprocessor System with the UltraSPARC III Cu ..... | 17        |
| 2.2       | Cache Coherence .....   | 18        |
| 2.3       | System Interfaces .....   | 18        |
| 2.3.1     | Fireplane Interconnect .....                                      | 18        |
| 2.3.2     | SDRAM Interface .....   | 19        |
| 2.3.3     | DCDS Interface .....  | 19        |
| 2.3.4     | L2-Cache Interface .....  | 19        |
| 2.3.5     | Boot Bus Interface .....  | 19        |
| 2.3.6     | JTAG Interface .....  | 19        |

---

## Section II: Architecture and Functions

---

|           |   |           |
|-----------|---|-----------|
| <b>3.</b> | <b>CPU Architecture Basics .....</b>        | <b>23</b> |
| 3.1       | Component Overview .....                    | 23        |
| 3.1.1     | Instruction Fetch and Buffering .....       | 24        |
| 3.1.2     | Execution Pipelines .....                   | 25        |
| 3.1.3     | Load/Store Unit .....                       | 26        |
| 3.1.4     | Memory Management Units .....               | 27        |
| 3.1.5     | L2-cache Unit (Level-2 Unified Cache) ..... | 28        |
| 3.1.6     | System Interface Unit .....                 | 28        |
| 3.1.7     | Memory Controller Unit .....                | 29        |
| 3.2       | CPU Operating Modes .....                   | 29        |

|           |  |           |
|-----------|--|-----------|
| 3.2.1     | Privileged Mode .....  | 29        |
| 3.2.2     | Non-Privileged Mode .....  | 30        |
| 3.2.3     | Reset and RED_State .....  | 30        |
| 3.2.4     | Error Handling .....   | 32        |
| 3.2.5     | Debug and Diagnostics Mode .....                                       | 34        |
| <b>4.</b> | <b>Instruction Execution .....</b>                                     | <b>37</b> |
| 4.1       | Introduction .....   | 37        |
| 4.1.1     | NOP, Neutralized, and Helper Instructions .....                        | 38        |
| 4.2       | Processor Pipeline .....   | 38        |
| 4.2.1     | Grouping Dependencies .....  | 41        |
| 4.2.2     | Dispatch Dependencies .....  | 41        |
| 4.2.3     | Execution Dependencies .....   | 41        |
| 4.2.4     | Instruction-Fetch Stages .....   | 42        |
| 4.2.5     | Instruction Issue and Queue Stages .....                               | 43        |
| 4.2.6     | Execution Pipeline .....   | 44        |
| 4.2.7     | Trap and Done Stages .....   | 46        |
| 4.3       | Pipeline Recirculation .....   | 47        |
| 4.4       | Grouping Rules .....   | 47        |
| 4.4.1     | Execution Order .....  | 48        |
| 4.4.2     | Integer Register Dependencies to Instructions in the MS Pipeline ..... | 48        |
| 4.4.3     | Integer Instructions Within a Group .....                              | 49        |
| 4.4.4     | Same-Group Bypass .....  | 50        |
| 4.4.5     | Floating-Point Unit Operand Dependencies .....                         | 50        |
| 4.4.6     | Grouping Rules for Register-Window Management Instructions .....       | 51        |
| 4.4.7     | Grouping Rules for Reads and Writes of the ASRs .....                  | 52        |
| 4.4.8     | Grouping Rules for Other Instructions .....                            | 52        |
| 4.5       | Conditional Moves .....  | 53        |
| 4.6       | Instruction Latencies and Dispatching Properties .....                 | 54        |
| 4.6.1     | Latency .....  | 55        |

|       |                     |    |
|-------|---------------------|----|
| 4.6.2 | Blocking .....      | 55 |
| 4.6.3 | Pipeline .....      | 56 |
| 4.6.4 | Break and SIG ..... | 56 |

---

## Section III: Execution Environment

---

|           |  |           |
|-----------|--|-----------|
| <b>5.</b> | <b>Data Formats .....</b>                      | <b>65</b> |
| 5.1       | Integer Data Formats .....                     | 66        |
| 5.1.1     | Integer Data Value Range .....                 | 66        |
| 5.1.2     | Integer Data Alignment .....                   | 67        |
| 5.1.3     | Signed Integer Data Types .....                | 67        |
| 5.1.4     | Unsigned Integer Data Types .....              | 69        |
| 5.1.5     | Tagged Word .....                              | 70        |
| 5.2       | Floating-Point Data Formats .....              | 71        |
| 5.2.1     | Floating-Point Data Value Range .....          | 71        |
| 5.2.2     | Floating-Point Data Alignment .....            | 71        |
| 5.2.3     | Floating-Point, Single-Precision .....         | 72        |
| 5.2.4     | Floating-Point, Double-Precision .....         | 72        |
| 5.2.5     | Floating-Point, Quad-Precision .....           | 73        |
| 5.3       | VIS Execution Unit Data Formats .....          | 74        |
| 5.3.1     | Pixel Data Format .....                        | 75        |
| 5.3.2     | Fixed-Point Data Formats .....                 | 76        |
| <b>6.</b> | <b>Registers .....</b>                         | <b>79</b> |
| 6.1       | Introduction .....                             | 79        |
| 6.2       | Integer Unit General-Purpose r Registers ..... | 80        |
| 6.2.1     | Windowed (in/local/out) r Registers .....      | 82        |
| 6.2.2     | Global r Register Sets .....                   | 82        |
| 6.2.3     | 128-bit Operand Considerations .....           | 84        |
| 6.3       | Register Window Management .....               | 84        |



|       |   |     |
|-------|---|-----|
| 6.3.1 | CALL and JMPL Instructions .....                              | 86  |
| 6.3.2 | Circular Windowing .....                                      | 86  |
| 6.3.3 | Clean Window with RESTORE and SAVE Instructions .....         | 86  |
| 6.4   | Floating-Point General-Purpose Registers .....                | 86  |
| 6.4.1 | Floating-Point Register Number Encoding .....                 | 88  |
| 6.4.2 | Double and Quad Floating-Point Operands .....                 | 89  |
| 6.5   | Control and Status Register Summary .....                     | 91  |
| 6.5.1 | State and Ancillary State Register Summary .....              | 91  |
| 6.5.2 | Privileged Register Summary .....                             | 92  |
| 6.5.3 | ASI and Specially Accessed Register Summary .....             | 94  |
| 6.6   | State Registers .....   | 96  |
| 6.6.1 | 32-bit Multiply/Divide (YD) State Register 0 .....            | 96  |
| 6.6.2 | Integer Unit Condition Codes State Register 2 (CCR) .....     | 96  |
| 6.6.3 | Address Space Identifier (ASI) Register ASR 3 .....           | 98  |
| 6.6.4 | TICK Register (TICK) ASR4 .....                               | 98  |
| 6.6.5 | Program Counters State Register 5 .....                       | 99  |
| 6.6.6 | Floating-Point Registers State (FPRS) Register 6 .....        | 99  |
| 6.7   | Ancillary State Registers: ASRs 16-25 .....                   | 100 |
| 6.7.1 | Dispatch Control Register (DCR) ASR 18 .....                  | 100 |
| 6.7.2 | Graphics Status Register (GSR) ASR 19 .....                   | 103 |
| 6.7.3 | Software Interrupt State Registers: ASRs 20, 21, and 22 ..... | 104 |
| 6.7.4 | Timer State Registers: ASRs 4, 23, 24, 25 .....               | 106 |
| 6.8   | Privileged Registers .....                                    | 109 |
| 6.8.1 | Trap Stack Privileged Registers 0 through 3 .....             | 109 |
| 6.8.2 | Trap Base Address (TBA) Privileged Register 5 .....           | 111 |
| 6.8.3 | Processor State (PSTATE) Privileged Register 6 .....          | 112 |
| 6.8.4 | Trap Level (TL) Privileged Register 7 .....                   | 117 |
| 6.8.5 | Processor Interrupt Level (PIL) Privileged Register 8 .....   | 117 |
| 6.8.6 | Register-Window State Privileged Registers 9 through 13 ..... | 118 |
| 6.8.7 | Window State (WSTATE) Privileged Register 14 .....            | 120 |

|           |   |            |
|-----------|---|------------|
| 6.8.8     | Version (VER) Privileged Register 31 .....              | 120        |
| 6.9       | Special Access Register .....                           | 122        |
| 6.9.1     | Floating-Point Status Register (FSR) .....              | 122        |
| 6.10      | ASI Mapped Registers .....                              | 132        |
| 6.10.1    | Data Cache Unit Control Register (DCUCR) .....          | 132        |
| 6.10.2    | Data Watchpoint Registers .....                         | 136        |
| <b>7.</b> | <b>Instruction Types .....</b>                          | <b>139</b> |
| 7.1       | Introduction .....                                      | 140        |
| 7.2       | Memory Addressing for Load and Store Instructions ..... | 140        |
| 7.2.1     | Integer Unit Memory Alignment Requirements .....        | 141        |
| 7.2.2     | FP/VIS Memory Alignment Requirements .....              | 141        |
| 7.2.3     | Byte Order Addressing Conventions (Endianness) .....    | 141        |
| 7.2.4     | Address Space Identifiers .....                         | 142        |
| 7.2.5     | Maintaining Data Coherency .....                        | 143        |
| 7.3       | Integer Execution Environment .....                     | 143        |
| 7.3.1     | IU Data Access Instructions .....                       | 143        |
| 7.3.2     | IU Arithmetic Instructions .....                        | 147        |
| 7.3.3     | IU Logic Instructions .....                             | 148        |
| 7.3.4     | IU Compare Instructions .....                           | 148        |
| 7.3.5     | IU Miscellaneous Instructions .....                     | 149        |
| 7.4       | Floating-Point Execution Environment .....              | 150        |
| 7.4.1     | Floating-Point Operate Instructions .....               | 150        |
| 7.4.2     | FPU/VIS Data Access Instructions .....                  | 151        |
| 7.4.3     | FP Arithmetic Instructions .....                        | 152        |
| 7.4.4     | FP Conversion Instructions .....                        | 153        |
| 7.4.5     | FP Compare Instructions .....                           | 153        |
| 7.4.6     | FP Miscellaneous Instructions .....                     | 153        |
| 7.5       | VIS Execution Environment .....                         | 154        |
| 7.5.1     | VIS Pixel Data Instructions .....                       | 154        |

|           |   |            |
|-----------|---|------------|
| 7.5.2     | VIS Fixed-Point 16-bit and 32-bit Data Instructions ..... | 155        |
| 7.5.3     | VIS Logic Instructions .....                              | 156        |
| 7.6       | Data Coherency Instructions .....                         | 156        |
| 7.6.1     | FLUSH Instruction Cache Instruction .....                 | 157        |
| 7.6.2     | MEMBAR (Memory Synchronization) Instruction .....         | 157        |
| 7.6.3     | Store Barrier Instruction .....                           | 157        |
| 7.7       | Register Window Management Instructions .....             | 157        |
| 7.8       | Program Control Transfer Instructions .....               | 158        |
| 7.8.1     | Control Transfer Instructions (CTIs) .....                | 159        |
| 7.9       | Prefetch Instructions .....                               | 164        |
| 7.10      | Instruction Summary Table by Category .....               | 165        |
| 7.10.1    | Instruction Superscripts .....                            | 165        |
| 7.10.2    | Instruction Mnemonics Expansion .....                     | 165        |
| 7.10.3    | Instruction Grouping Rules .....                          | 165        |
| 7.10.4    | Table Organization .....                                  | 165        |
| 7.10.5    | Integer Execution Environment Instructions .....          | 166        |
| 7.10.6    | Floating-Point Execution Environment Instructions .....   | 170        |
| 7.10.7    | VIS Execution Environment Instructions .....              | 172        |
| 7.10.8    | Data Coherency Instructions .....                         | 173        |
| 7.10.9    | Register-Window Management Instructions .....             | 173        |
| 7.10.10   | Program Control Transfer Instructions .....               | 174        |
| 7.10.11   | Data Prefetch Instructions .....                          | 174        |
| 7.11      | Instruction Formats and Fields .....                      | 175        |
| 7.12      | Reserved Opcodes and Instruction Fields .....             | 179        |
| 7.12.1    | Summary of Unimplemented Instructions .....               | 179        |
| 7.13      | Big/Little-endian Addressing .....                        | 180        |
| 7.13.1    | Big-endian Addressing Convention .....                    | 180        |
| 7.13.2    | Little-endian Addressing Convention .....                 | 182        |
| <b>8.</b> | <b>Address Space Identifiers .....</b>                    | <b>185</b> |

|       |  |     |
|-------|--|-----|
| 8.1   | ASI Introduction .....                               | 185 |
| 8.1.1 | Load/Store Instructions .....                        | 186 |
| 8.1.2 | Processor State .....                                | 186 |
| 8.1.3 | Default ASIs .....                                   | 186 |
| 8.1.4 | Non-Translating and Bypassing ASIs .....             | 186 |
| 8.1.5 | Datapath .....                                       | 186 |
| 8.2   | ASI Heredity .....                                   | 188 |
| 8.2.1 | SPARC V9 ASIs .....                                  | 188 |
| 8.2.2 | UltraSPARC UltraSPARC III Family ASIs .....          | 188 |
| 8.2.3 | UltraSPARC III Cu Specific ASIs .....                | 189 |
| 8.3   | ASI Groups .....                                     | 189 |
| 8.4   | Instructions Associated with the ASIs .....          | 191 |
| 8.4.1 | Block Load and Block Store ASIs .....                | 191 |
| 8.4.2 | Partial Store ASIs .....                             | 191 |
| 8.4.3 | Short Floating-Point Load and Store ASIs .....       | 191 |
| 8.5   | Using ASIs .....                                     | 192 |
| 8.5.1 | Data Widths .....                                    | 192 |
| 8.5.2 | Operand Alignment .....                              | 192 |
| 8.5.3 | Common Exceptions .....                              | 192 |
| 8.6   | List of ASI Definitions .....                        | 193 |
| 8.7   | Special Memory Access ASIs .....                     | 204 |
| 8.7.1 | ASI 0x14 (ASI_PHYS_USE_EC) .....                     | 205 |
| 8.7.2 | ASI 0x15 (ASI_PHYS_BYPASS_EC_WITH_EBIT) .....        | 205 |
| 8.7.3 | ASI 0x1C (ASI_PHYS_USE_EC_LITTLE) .....              | 205 |
| 8.7.4 | ASI 0x1D (ASI_PHYS_BYPASS_EC_WITH_EBIT_LITTLE) ..... | 205 |
| 8.7.5 | ASIs 0x24 and 0x2C (Load Quadword ASIs) .....        | 206 |

---

## Section IV: Memory and Cache

---

|           |                            |            |
|-----------|----------------------------|------------|
| <b>9.</b> | <b>Memory Models .....</b> | <b>209</b> |
|-----------|----------------------------|------------|

|            |   |            |
|------------|---|------------|
| 9.1        | UltraSPARC III Cu TSO Behavior .....                  | 210        |
| 9.2        | Memory Location Identification .....                  | 210        |
| 9.3        | Memory Accesses and Cacheability .....                | 210        |
| 9.3.1      | Coherence Domains .....                               | 211        |
| 9.3.2      | Global Visibility .....                               | 212        |
| 9.3.3      | Memory Ordering .....                                 | 212        |
| 9.4        | Memory Synchronization .....                          | 213        |
| 9.4.1      | MEMBAR #Sync .....                                    | 214        |
| 9.4.2      | MEMBAR Rules .....                                    | 214        |
| 9.4.3      | FLUSH .....   | 216        |
| 9.5        | Atomic Operations .....                               | 216        |
| 9.6        | Non-Faulting Load .....                               | 218        |
| 9.7        | Prefetch Instructions .....                           | 218        |
| 9.8        | Block Loads and Stores .....                          | 219        |
| 9.9        | I/O and Accesses with Side-Effects .....              | 220        |
| 9.9.1      | Instruction Prefetch to Side-Effect Locations .....   | 220        |
| 9.9.2      | Instruction Prefetch Exiting Red State .....          | 221        |
| 9.10       | UltraSPARC III Cu Internal ASIs .....                 | 221        |
| 9.11       | Store Compression .....                               | 222        |
| 9.12       | Read-After-Write (RAW) Bypassing .....                | 222        |
| 9.12.1     | RAW Bypassing Algorithm .....                         | 223        |
| 9.12.2     | RAW Detection Algorithm .....                         | 224        |
| <b>10.</b> | <b>Caches and Cache Coherency .....</b>               | <b>225</b> |
| 10.1       | Cache Organization .....                              | 225        |
| 10.1.1     | Virtual Indexed, Physical Tagged Caches (VIPT) .....  | 225        |
| 10.1.2     | Physical Indexed, Physical Tagged Caches (PIPT) ..... | 226        |
| 10.2       | Cache Flushing .....                                  | 229        |
| 10.2.1     | Address Aliasing Flushing .....                       | 230        |
| 10.2.2     | Committing Block Store Flushing .....                 | 231        |

|            |  |            |
|------------|--|------------|
| 10.2.3     | Displacement Flushing .....  | 231        |
| 10.3       | Bypassing the D-Cache .....  | 231        |
| 10.4       | Controlling P-cache .....  | 232        |
| 10.5       | Coherence Tables .....   | 233        |
| 10.5.1     | Processor State Transition and the Generated Transaction .....     | 233        |
| 10.5.2     | Snoop Output and Input .....                                       | 236        |
| 10.5.3     | Transaction Handling .....   | 240        |
| <b>11.</b> | <b>Memory Management Unit .....</b>                                | <b>245</b> |
| 11.1       | Virtual Address Translation .....                                  | 245        |
| 11.2       | Translation Table Entry .....                                      | 248        |
| 11.3       | Translation Storage Buffer .....                                   | 251        |
| 11.3.1     | TSB Indexing Support .....   | 251        |
| 11.3.2     | TSB Cacheability .....   | 252        |
| 11.3.3     | TSB Organization .....   | 252        |
| 11.4       | Hardware Support for TSB Access .....                              | 253        |
| 11.4.1     | Typical TLB Miss/Refill Sequence .....                             | 253        |
| 11.4.2     | TSB Pointer Formation .....  | 254        |
| 11.4.3     | TSB Pointer Logic Hardware Description .....                       | 256        |
| 11.4.4     | Required TLB Conditions .....                                      | 258        |
| 11.4.5     | Required TSB Conditions .....                                      | 258        |
| 11.4.6     | MMU Global Registers Selection .....                               | 258        |
| 11.5       | Faults and Traps .....   | 259        |
| 11.6       | ASI Value, Context, and Endianness Selection for Translation ..... | 261        |
| 11.7       | Reset, Disable, and RED_state Behavior .....                       | 263        |
| 11.8       | SPARC V9 “MMU Requirements” Annex .....                            | 264        |
| 11.9       | Data Translation Lookaside Buffer .....                            | 265        |
| 11.9.1     | D-TLB Access Operation .....                                       | 266        |
| 11.9.2     | Same Page Size on Both dt512_0 and dt512_1 .....                   | 267        |
| 11.9.3     | D-TLB Automatic Replacement .....                                  | 268        |

|         |  |     |
|---------|--|-----|
| 11.9.4  | D-TLB Directed Data Read/Write .....           | 270 |
| 11.9.5  | D-TLB Tag Read Register .....                  | 271 |
| 11.9.6  | Demap Operation .....                          | 272 |
| 11.9.7  | D-TLB Access Summary .....                     | 273 |
| 11.9.8  | D-MMU Operation Summary .....                  | 274 |
| 11.9.9  | Internal Registers and ASI Operations .....    | 276 |
| 11.9.10 | D-MMU Bypass .....                             | 294 |
| 11.10   | Instruction Translation Lookaside Buffer ..... | 294 |
| 11.10.1 | I-TLB Access Operation .....                   | 294 |
| 11.10.2 | I-TLB Automatic Replacement .....              | 295 |
| 11.10.3 | I-TLB Access Summary .....                     | 296 |
| 11.10.4 | I-MMU Operation Summary .....                  | 298 |
| 11.10.5 | Internal Registers and ASI Operations .....    | 298 |
| 11.10.6 | I-MMU Bypass .....                             | 311 |

---

## Section V: Supervisor Programming

---

|            |  |            |
|------------|--|------------|
| <b>12.</b> | <b>Traps and Trap Handling .....</b>             | <b>315</b> |
| 12.1       | Processor States, Normal and Special Traps ..... | 316        |
| 12.1.1     | RED_state .....                                  | 317        |
| 12.1.2     | Error_state .....                                | 320        |
| 12.2       | Trap Categories .....                            | 320        |
| 12.2.1     | Precise Traps .....                              | 321        |
| 12.2.2     | Deferred Traps .....                             | 321        |
| 12.2.3     | Disrupting Traps .....                           | 322        |
| 12.2.4     | Reset Traps .....                                | 323        |
| 12.2.5     | Uses of the Trap Categories .....                | 323        |
| 12.3       | Trap Control .....                               | 324        |
| 12.3.1     | PIL Control .....                                | 324        |

|            |  |            |
|------------|--|------------|
| 12.3.2     | TEM Control .....                                  | 325        |
| 12.4       | Trap-Table Entry Addresses .....                   | 325        |
| 12.4.1     | Trap Table Organization .....                      | 325        |
| 12.4.2     | Trap Type (TT) .....                               | 326        |
| 12.4.3     | Trap Priorities .....                              | 331        |
| 12.4.4     | Details of Supported Traps .....                   | 331        |
| 12.5       | Trap Processing .....                              | 334        |
| 12.5.1     | Normal Trap Processing .....                       | 335        |
| 12.5.2     | Fast MMU Trap Processing .....                     | 337        |
| 12.5.3     | Interrupt Vector Trap Processing .....             | 338        |
| 12.5.4     | Special Trap Processing .....                      | 339        |
| 12.6       | Exception and Interrupt Descriptions .....         | 345        |
| <b>13.</b> | <b>Interrupt Handling .....</b>                    | <b>351</b> |
| 13.1       | Interrupt Vector Dispatch .....                    | 352        |
| 13.2       | Interrupt Vector Receive .....                     | 353        |
| 13.3       | Interrupt Global Registers .....                   | 354        |
| 13.4       | Interrupt ASI Registers .....                      | 354        |
| 13.4.1     | Outgoing Interrupt Vector Data<7:0> Register ..... | 354        |
| 13.4.2     | Interrupt Vector Dispatch Register .....           | 355        |
| 13.4.3     | Interrupt Vector Dispatch Status Register .....    | 356        |
| 13.4.4     | Incoming Interrupt Vector Data<7:0> .....          | 356        |
| 13.4.5     | Interrupt Vector Receive Register .....            | 357        |
| 13.5       | Software Interrupt Register (SOFTINT) .....        | 358        |
| 13.5.1     | Setting the Software Interrupt Register .....      | 358        |
| 13.5.2     | Clearing the Software Interrupt Register .....     | 358        |

---

## Section VI: Performance Programming

---

|            |  |            |
|------------|--|------------|
| <b>14.</b> | <b>Performance Instrumentation .....</b> | <b>363</b> |
|------------|--|------------|



|      |   |     |
|------|---|-----|
| 14.1 | Performance Control Register (PCR) .....                        | 364 |
| 14.2 | Performance Instrumentation Counter (PIC) Register .....        | 365 |
|      | 14.2.1 PIC Counter Overflow Trap Operation .....                | 367 |
| 14.3 | Performance Instrumentation Operation .....                     | 367 |
|      | 14.3.1 Performance Instrumentation Implementations .....        | 369 |
|      | 14.3.2 Performance Instrumentation Accuracy .....               | 369 |
| 14.4 | Pipeline Counters .....   | 369 |
|      | 14.4.1 Instruction Execution and CPU Clock Counts .....         | 369 |
|      | 14.4.2 IIU Statistics .....                                     | 370 |
|      | 14.4.3 IIU Stall Counts .....                                   | 370 |
|      | 14.4.4 R-stage Stall Counts .....                               | 372 |
|      | 14.4.5 Recirculation Counts .....                               | 372 |
| 14.5 | Cache Access Counters .....                                     | 373 |
|      | 14.5.1 Instruction Cache Events .....                           | 373 |
|      | 14.5.2 Data Cache Events .....                                  | 373 |
|      | 14.5.3 Write Cache Events .....                                 | 374 |
|      | 14.5.4 Prefetch Cache Events .....                              | 374 |
|      | 14.5.5 L2-Cache Events .....                                    | 375 |
|      | 14.5.6 Separating D-cache Stall Counts .....                    | 376 |
| 14.6 | Memory Controller Counters .....                                | 378 |
|      | 14.6.1 Memory Controller Read Request Events .....              | 378 |
| 14.7 | Data Locality Counters for Scalable Shared Memory Systems ..... | 379 |
|      | 14.7.1 Event Tree .....   | 380 |
|      | 14.7.2 Data Locality Event Matrix .....                         | 382 |
|      | 14.7.3 Synthesized Data Locality Events .....                   | 383 |
| 14.8 | Miscellaneous Counters .....                                    | 384 |
|      | 14.8.1 System Interface Events and Clock Cycles .....           | 384 |
|      | 14.8.2 Software Events .....                                    | 384 |
|      | 14.8.3 Floating-Point Operation Events .....                    | 385 |
| 14.9 | PCR.SL and PCR.SU Encodings .....                               | 385 |

|   |            |
|---|------------|
| <b>15. Processor Optimization .....</b>                         | <b>387</b> |
| 15.1 Introduction .....   | 387        |
| 15.2 Instruction Stream Issues .....                            | 387        |
| 15.2.1 Instruction Alignment .....                              | 388        |
| 15.2.2 Instruction Cache Timing .....                           | 390        |
| 15.2.3 Executing Code Out of the Level 2 Cache .....            | 390        |
| 15.2.4 TLB Misses .....   | 391        |
| 15.2.5 Conditional Moves vs. Conditional Branches .....         | 391        |
| 15.2.6 Instruction Cache Utilization .....                      | 392        |
| 15.2.7 Handling of CTI couples .....                            | 392        |
| 15.2.8 Mispredicted Branches .....                              | 392        |
| 15.2.9 Return Address Stack (RAS) .....                         | 392        |
| 15.3 Data Stream Issues .....                                   | 393        |
| 15.3.1 Data Cache Organization .....                            | 393        |
| 15.3.2 Data Cache Timing .....                                  | 394        |
| 15.3.3 Data Alignment .....                                     | 394        |
| 15.3.4 Store Considerations .....                               | 394        |
| 15.3.5 Read-After-Write Hazards .....                           | 395        |
| <b>16. Prefetch .....</b>                                       | <b>397</b> |
| 16.1 Prefetch Cache .....                                       | 397        |
| 16.1.1 P-Cache Data Flow .....                                  | 398        |
| 16.1.2 Fetched_Mode Tag Bit .....                               | 398        |
| 16.2 Software Prefetch .....                                    | 398        |
| 16.2.1 Software Prefetch Instructions .....                     | 399        |
| 16.2.2 Software Prefetch Instruction Uses the MS Pipeline ..... | 400        |
| 16.2.3 Cancelling Software Prefetch .....                       | 400        |
| 16.2.4 Prefetch Instruction Variants .....                      | 400        |
| 16.2.5 General Comments .....                                   | 401        |
| 16.2.6 Code Example .....                                       | 402        |

|            |   |            |
|------------|---|------------|
| 16.3       | Hardware Prefetch .....                           | 404        |
| 16.3.1     | Cancelling Hardware Prefetches .....              | 404        |
| 16.3.2     | FP/VIS Load Instruction Miss Fetch .....          | 405        |
| <b>17.</b> | <b>IEEE 754-1985 Standard .....</b>               | <b>407</b> |
| 17.1       | Introduction .....                                | 407        |
| 17.1.1     | Floating-point Operations .....                   | 407        |
| 17.1.2     | Rounding Mode .....                               | 408        |
| 17.1.3     | Nonstandard FP Operating Mode .....               | 408        |
| 17.1.4     | Memory and Register Data Images .....             | 408        |
| 17.1.5     | Subnormal Operations .....                        | 408        |
| 17.1.6     | FSR.CEXC and FSR.AEXC Updates .....               | 409        |
| 17.1.7     | Prediction Logic .....                            | 409        |
| 17.2       | Floating-point Numbers .....                      | 409        |
| 17.2.1     | Floating-point Number Line .....                  | 411        |
| 17.3       | IEEE Operations .....                             | 411        |
| 17.3.1     | Addition .....                                    | 412        |
| 17.3.2     | Subtraction .....                                 | 413        |
| 17.3.3     | Multiplication .....                              | 414        |
| 17.3.4     | Division .....                                    | 415        |
| 17.3.5     | Square Root .....                                 | 416        |
| 17.3.6     | Compare .....                                     | 416        |
| 17.3.7     | Precision Conversion .....                        | 417        |
| 17.3.8     | Floating-point to Integer Number Conversion ..... | 418        |
| 17.3.9     | Integer to Floating-point Number Conversion ..... | 419        |
| 17.3.10    | Copy/Move Operations .....                        | 419        |
| 17.3.11    | f Register Load/Store Operations .....            | 420        |
| 17.3.12    | VIS Operations .....                              | 420        |
| 17.4       | Traps and Exceptions .....                        | 420        |
| 17.4.1     | Summary of Exceptions .....                       | 421        |

|        |  |     |
|--------|--|-----|
| 17.4.2 | Trap Event .....                             | 422 |
| 17.4.3 | Trap Priority .....                          | 422 |
| 17.5   | IEEE Traps .....                             | 422 |
| 17.5.1 | IEEE Trap Enable Mask (TEM) .....            | 422 |
| 17.5.2 | IEEE Invalid (nv) Trap .....                 | 423 |
| 17.5.3 | IEEE Overflow (of) Trap .....                | 423 |
| 17.5.4 | IEEE Underflow (uf) Trap .....               | 423 |
| 17.5.5 | IEEE Divide-by-Zero (DZ) Trap .....          | 423 |
| 17.5.6 | IEEE Inexact (nx) Trap .....                 | 424 |
| 17.6   | Underflow Operation .....                    | 424 |
| 17.6.1 | Trapped Underflow .....                      | 425 |
| 17.6.2 | Untrapped Underflow .....                    | 425 |
| 17.7   | IEEE NaN Operations .....                    | 426 |
| 17.7.1 | Signaling and Quiet NaNs .....               | 426 |
| 17.7.2 | SNaN to QNaN Transformation .....            | 426 |
| 17.7.3 | Operations with NaN Operands .....           | 427 |
| 17.7.4 | NaN Results from Operands without NaNs ..... | 428 |
| 17.8   | Subnormal Operations .....                   | 428 |
| 17.8.1 | Response to Subnormal Operands .....         | 428 |
| 17.8.2 | Subnormal Number Generation .....            | 429 |

---

## Section VII: Special Topics

---

|            |  |            |
|------------|--|------------|
| <b>18.</b> | <b>Reset and RED_state .....</b>       | <b>435</b> |
| 18.1       | RED_state Characteristics .....        | 435        |
| 18.2       | Resets .....                           | 436        |
| 18.2.1     | Power-on Reset (POR) .....             | 436        |
| 18.2.2     | System Reset .....                     | 437        |
| 18.2.3     | Externally Initiated Reset (XIR) ..... | 437        |

|        |  |     |
|--------|--|-----|
| 18.2.4 | Watchdog Reset (WDR) and error_state ..... | 437 |
| 18.2.5 | Software-Initiated Reset (SIR) .....       | 438 |
| 18.3   | RED_state Trap Vector .....                | 438 |
| 18.4   | Machine States .....                       | 438 |

---

## Section VIII: Appendix

---

|           |   |            |
|-----------|---|------------|
| <b>A.</b> | <b>Instruction Definitions .....</b>                                    | <b>447</b> |
| A.1       | Add .....   | 454        |
| A.2       | Alignment Instructions (VIS I) .....                                    | 455        |
| A.3       | Three-Dimensional Array Addressing Instructions (VIS I) .....           | 457        |
| A.4       | Block Load and Block Store (VIS I) .....                                | 460        |
| A.5       | Byte Mask and Shuffle Instructions (VIS II) .....                       | 468        |
| A.6       | Branch on Integer Register with Prediction (BPr) .....                  | 469        |
| A.7       | Branch on Floating-Point Condition Codes with Prediction (FBPfcc) ..... | 471        |
| A.8       | Branch on Integer Condition Codes with Prediction (BPcc) .....          | 474        |
| A.9       | Call and Link .....   | 476        |
| A.10      | Compare and Swap .....  | 477        |
| A.11      | DONE and RETRY .....  | 479        |
| A.12      | Edge Handling Instructions (VIS I, VIS II) .....                        | 480        |
| A.13      | Floating-Point Add and Subtract .....                                   | 483        |
| A.14      | Floating-Point Compare .....  | 486        |
| A.15      | Convert Floating-Point to Integer .....                                 | 488        |
| A.16      | Convert Between Floating-Point Formats .....                            | 489        |
| A.17      | Convert Integer to Floating-Point .....                                 | 491        |
| A.18      | Floating-Point Move .....   | 493        |
| A.19      | Floating-Point Multiply and Divide .....                                | 494        |
| A.20      | Floating-Point Square Root .....  | 496        |
| A.21      | Flush Instruction Memory .....  | 497        |
| A.22      | Flush Register Windows .....  | 499        |

|        |  |     |
|--------|--|-----|
| A.23   | Illegal Instruction Trap .....   | 500 |
| A.24   | Jump and Link .....  | 501 |
| A.25   | Load Floating-Point .....  | 502 |
| A.26   | Load Floating-Point from Alternate Space .....                           | 504 |
| A.27   | Load Integer .....   | 506 |
| A.28   | Load Integer from Alternate Space .....                                  | 508 |
| A.29   | Load Quadword, Atomic (VIS I) .....                                      | 510 |
| A.30   | Load-Store Unsigned Byte .....   | 512 |
| A.31   | Load-Store Unsigned Byte to Alternate Space .....                        | 513 |
| A.32   | Logical Operate Instructions (VIS I) .....                               | 514 |
| A.33   | Logical Operations .....   | 517 |
| A.34   | Memory Barrier .....   | 519 |
| A.35   | Move Floating-Point Register on Condition (FMOVcc) .....                 | 524 |
| A.36   | Move Floating-Point Register on Integer Register Condition (FMOVr) ..... | 529 |
| A.37   | Move Integer Register on Condition (MOVcc) .....                         | 531 |
| A.38   | Move Integer Register on Register Condition (MOVr) .....                 | 536 |
| A.39   | Multiply and Divide (64-bit) .....                                       | 537 |
| A.40   | No Operation .....   | 538 |
| A.41   | Partial Store (VIS I) .....  | 540 |
| A.42   | Partitioned Add/Subtract Instructions (VIS I) .....                      | 542 |
| A.43   | Partitioned Multiply Instructions (VIS I) .....                          | 543 |
| A.43.1 | FMUL8x16 Instruction .....   | 545 |
| A.43.2 | FMUL8x16AU Instruction .....   | 545 |
| A.43.3 | FMUL8x16AL Instruction .....   | 546 |
| A.43.4 | FMUL8SUX16 Instruction .....   | 546 |
| A.43.5 | FMUL8ULX16 Instruction .....   | 547 |
| A.43.6 | FMULD8SUX16 Instruction .....  | 548 |
| A.43.7 | FMULD8ULX16 Instruction .....  | 549 |
| A.44   | Pixel Compare (VIS I) .....  | 550 |
| A.45   | Pixel Component Distance (PDIST) (VIS I) .....                           | 552 |

|        |   |     |
|--------|---|-----|
| A.46   | Pixel Formatting (VIS I)                    | 553 |
| A.46.1 | FPACK16                                     | 554 |
| A.46.2 | FPACK32                                     | 556 |
| A.46.3 | FPACKFIX                                    | 557 |
| A.46.4 | FEXPAND                                     | 558 |
| A.46.5 | FPMERGE                                     | 559 |
| A.47   | Population Count                            | 559 |
| A.48   | Prefetch Data                               | 560 |
| A.48.1 | Prefetch Instruction Variants               | 562 |
| A.48.2 | New Error Handling of Prefetches            | 563 |
| A.49   | Read Privileged Register                    | 566 |
| A.50   | Read State Register                         | 568 |
| A.51   | RETURN                                      | 570 |
| A.52   | SAVE and RESTORE                            | 572 |
| A.53   | SAVED and RESTORED                          | 574 |
| A.54   | Set Interval Arithmetic Mode (VIS II)       | 575 |
| A.55   | SETHI                                       | 577 |
| A.56   | Shift                                       | 578 |
| A.57   | Short Floating-Point Load and Store (VIS I) | 580 |
| A.58   | SHUTDOWN (VIS I)                            | 582 |
| A.59   | Software-Initiated Reset                    | 583 |
| A.60   | Store Floating-Point                        | 584 |
| A.61   | Store Floating-Point into Alternate Space   | 586 |
| A.62   | Store Integer                               | 588 |
| A.63   | Store Integer into Alternate Space          | 589 |
| A.64   | Subtract                                    | 591 |
| A.65   | Tagged Add                                  | 592 |
| A.66   | Tagged Subtract                             | 593 |
| A.67   | Trap on Integer Condition Codes (Tcc)       | 595 |
| A.68   | Write Privileged Register                   | 597 |

|           |  |            |
|-----------|--|------------|
| A.69      | Write State Register .....                             | 600        |
| A.70      | Deprecated Instructions .....                          | 603        |
| A.70.1    | Branch on Floating-Point Condition Codes (FBfcc) ..... | 603        |
| A.70.2    | Branch on Integer Condition Codes (Bicc) .....         | 605        |
| A.70.3    | Divide (64-bit/32-bit) .....                           | 608        |
| A.70.4    | Load Floating-Point Status Register .....              | 611        |
| A.70.5    | Load Integer Doubleword .....                          | 612        |
| A.70.6    | Load Integer Doubleword from Alternate Space .....     | 614        |
| A.70.7    | Multiply (32-bit) .....                                | 616        |
| A.70.8    | Multiply Step .....                                    | 617        |
| A.70.9    | Read Y Register .....                                  | 619        |
| A.70.10   | Store Barrier .....                                    | 619        |
| A.70.11   | Store Floating-Point Status Register Lower .....       | 621        |
| A.70.12   | Store Integer Doubleword .....                         | 622        |
| A.70.13   | Store Integer Doubleword into Alternate Space .....    | 623        |
| A.70.14   | Swap Register with Memory .....                        | 625        |
| A.70.15   | Swap Register with Alternate Space Memory .....        | 626        |
| A.70.16   | Tagged Add and Trap on Overflow .....                  | 628        |
| A.70.17   | Tagged Subtract and Trap on Overflow .....             | 629        |
| A.70.18   | Write Y Register .....                                 | 630        |
| <b>B.</b> | <b>Assembly Language Syntax .....</b>                  | <b>633</b> |
| B.1       | Notation Used .....                                    | 633        |
| B.1.1     | Register Names .....                                   | 634        |
| B.1.2     | Special Symbol Names .....                             | 634        |
| B.1.3     | Values .....   | 637        |
| B.1.4     | Labels .....   | 637        |
| B.1.5     | Other Operand Syntax .....                             | 637        |
| B.1.6     | Comments .....   | 638        |
| B.2       | Syntax Design .....                                    | 638        |



B.3 Synthetic Instructions ..... 638

---

**Section IX: Index**

---

**19. Revision History .....671**



# List of Figures

---

|             |   |    |
|-------------|---|----|
| FIGURE 2-1  | Two-Processor Configuration with the UltraSPARC III Cu .....      | 14 |
| FIGURE 2-2  | Multiprocessor System with the UltraSPARC III Cu .....            | 16 |
| FIGURE 2-3  | Very Large Multiprocessor System with the UltraSPARC III Cu ..... | 17 |
| FIGURE 3-1  | UltraSPARC III Cu Architecture .....                              | 24 |
| FIGURE 4-1  | Instruction Pipeline Diagram Instruction Dependencies .....       | 40 |
| FIGURE 5-1  | Signed Integer Byte Data Format .....                             | 68 |
| FIGURE 5-2  | Signed Integer Halfword Data Format .....                         | 68 |
| FIGURE 5-3  | Signed Integer Word Data Format .....                             | 68 |
| FIGURE 5-4  | Signed Integer Double Data Format .....                           | 68 |
| FIGURE 5-5  | Signed Extended Integer Data Format .....                         | 69 |
| FIGURE 5-6  | Unsigned Integer Byte Data Format .....                           | 69 |
| FIGURE 5-7  | Unsigned Integer Halfword Data Format .....                       | 69 |
| FIGURE 5-8  | Unsigned Integer Word Data Format .....                           | 70 |
| FIGURE 5-9  | Unsigned Integer Double Data Format .....                         | 70 |
| FIGURE 5-10 | Unsigned Extended Integer Data Format .....                       | 70 |
| FIGURE 5-11 | Tagged Word Data Format .....                                     | 70 |
| FIGURE 5-12 | Floating-Point Single-Precision Data Format .....                 | 72 |
| FIGURE 5-13 | Floating-Point Double-Precision Double Word Data Format .....     | 72 |
| FIGURE 5-14 | Floating-Point Double-Precision Extended Word Data Format .....   | 73 |
| FIGURE 5-15 | Floating-Point Quad-Precision Data Format .....                   | 74 |

|             |  |     |
|-------------|--|-----|
| FIGURE 5-16 | Pixel Data Format with Band Sequential Ordering Shown .....              | 76  |
| FIGURE 5-17 | Fixed16 VIS Data Format .....  | 76  |
| FIGURE 5-18 | Fixed32 VIS Data Format .....  | 77  |
| FIGURE 6-1  | Three Overlapping Windows and the Eight Global Registers .....           | 83  |
| FIGURE 6-2  | Windowed r Registers for NWINDOWS = 8 .....                              | 85  |
| FIGURE 6-3  | Integer Unit r Registers and Floating-Point Unit Working Registers ..... | 90  |
| FIGURE 6-4  | State and Ancillary State Registers .....                                | 91  |
| FIGURE 6-5  | Privileged Registers .....   | 93  |
| FIGURE 6-6  | ASI and Specially Accessed Registers .....                               | 95  |
| FIGURE 6-7  | Y Register .....   | 96  |
| FIGURE 6-8  | Condition Codes Register .....   | 97  |
| FIGURE 6-9  | Integer Condition Codes (CCR_icc and CCR_xcc) .....                      | 97  |
| FIGURE 6-10 | Address Space Identifier Register .....                                  | 98  |
| FIGURE 6-11 | Floating-Point Registers State Register .....                            | 99  |
| FIGURE 6-12 | Dispatch Control Register (ASR 18) .....                                 | 101 |
| FIGURE 6-13 | Graphic Status Register (ASR 19) .....                                   | 103 |
| FIGURE 6-14 | SOFTINT, SET_SOFTINT, and CLR_SOFTINT Register Formats .....             | 104 |
| FIGURE 6-15 | Timer State Registers .....  | 106 |
| FIGURE 6-16 | Trap State Register Format .....   | 110 |
| FIGURE 6-17 | Trap Stack and Event Example .....                                       | 111 |
| FIGURE 6-18 | Trap Base Address Register .....   | 112 |
| FIGURE 6-19 | Trap Vector Address Format .....   | 112 |
| FIGURE 6-20 | PSTATE Fields .....  | 112 |
| FIGURE 6-21 | Trap Level Register .....  | 117 |
| FIGURE 6-22 | Processor Interrupt Level Register .....                                 | 118 |
| FIGURE 6-23 | WSTATE Register .....  | 120 |
| FIGURE 6-24 | Version Register .....   | 121 |

|              |  |     |
|--------------|--|-----|
| FIGURE 6-25  | FSR Fields Reserved Bits .....   | 122 |
| FIGURE 6-26  | Trap Enable Mask (TEM) Fields of FSR .....                                   | 128 |
| FIGURE 6-27  | Accrued Exception Bits (aexc) Fields of FSR .....                            | 129 |
| FIGURE 6-28  | Current Exception Bits (cexc) Fields of FSR .....                            | 129 |
| FIGURE 6-29  | DCU Control Register Access Data Format (ASI 4516) .....                     | 133 |
| FIGURE 6-30  | VA Data Watchpoint Register Format .....                                     | 137 |
| FIGURE 6-31  | PA Data Watchpoint Register Format .....                                     | 137 |
| FIGURE 7-1   | Summary of Instruction Formats: Formats 1 and 2 .....                        | 175 |
| FIGURE 7-2   | Summary of Instruction Formats: Format 3 .....                               | 176 |
| FIGURE 7-3   | Summary of Instruction Formats: Format 4 .....                               | 177 |
| FIGURE 7-4   | Big-endian Addressing Conventions .....                                      | 181 |
| FIGURE 7-5   | Little-endian Addressing Conventions .....                                   | 182 |
| FIGURE 8-1   | ASI Source and Function Conceptual Diagram .....                             | 187 |
| FIGURE 8-2   | ASI Groups .....   | 188 |
| FIGURE 11-1  | Virtual-to-Physical Address Translation for All Four Page Sizes .....        | 246 |
| FIGURE 11-2  | Software View of the MMU .....   | 247 |
| FIGURE 11-3  | Translation Storage Buffer (TSB) Translation Table Entry (TTE) .....         | 248 |
| FIGURE 11-4  | TSB Organization, Illustrated for Both Common and Shared Cases .....         | 252 |
| FIGURE 11-5  | Formation of TSB Pointers for 8 KB and 64 KB TTE .....                       | 256 |
| FIGURE 11-6  | D-MMU Primary Context Register .....   | 265 |
| FIGURE 11-7  | D-MMU Secondary Context Register .....                                       | 266 |
| FIGURE 11-8  | Tag Access Extension Register Format for Saving Page Sizes Information ..... | 268 |
| FIGURE 11-9  | D-MMU Primary Context Register .....   | 278 |
| FIGURE 11-10 | D-MMU Secondary Context Register .....                                       | 278 |
| FIGURE 11-11 | D-MMU Nucleus Context Register .....   | 278 |
| FIGURE 11-12 | D-MMU TLB Tag Access Registers Format .....                                  | 279 |
| FIGURE 11-13 | D-MMU TLB Data Access Address Format .....                                   | 282 |

|              |   |     |
|--------------|---|-----|
| FIGURE 11-14 | D-MMU TLB Tag Read Registers .....  | 283 |
| FIGURE 11-15 | MMU Tag Target Registers .....  | 283 |
| FIGURE 11-16 | MMU Data TSB Registers .....  | 284 |
| FIGURE 11-17 | MMU Data TSB Extension Registers .....                                    | 285 |
| FIGURE 11-18 | D-MMU TSB 8 KB/64 KB Pointer and D-MMU Direct Pointer Register .....      | 286 |
| FIGURE 11-19 | MMU Data Synchronous Fault Status Registers (I-SFSR, D-SFSR) .....        | 287 |
| FIGURE 11-20 | MMU Data Synchronous Fault Address Register (D-SFAR) .....                | 289 |
| FIGURE 11-21 | MMU Demap Operation Address and Data Formats .....                        | 290 |
| FIGURE 11-22 | MMU TLB Diagnostic Access Virtual Address .....                           | 292 |
| FIGURE 11-23 | D-MMU TLB CAM Diagnostic Registers .....                                  | 293 |
| FIGURE 11-24 | I-MMU TLB Tag Access Registers .....                                      | 299 |
| FIGURE 11-25 | I-MMU TLB Data Access Address .....                                       | 301 |
| FIGURE 11-26 | I-MMU TLB Tag Read Registers .....  | 302 |
| FIGURE 11-27 | MMU Tag Target Registers .....  | 303 |
| FIGURE 11-28 | MMU Instruction TSB Registers .....                                       | 303 |
| FIGURE 11-29 | MMU Instruction TSB Extension Registers .....                             | 305 |
| FIGURE 11-30 | I-MMU TSB 8 KB/64 KB Pointer and I-MMU Direct Pointer Register .....      | 305 |
| FIGURE 11-31 | MMU Instruction Synchronous Fault Status Registers (I-SFSR, D-SFSR) ..... | 306 |
| FIGURE 11-32 | MMU Demap Operation Address and Data Formats .....                        | 309 |
| FIGURE 11-33 | MMU TLB Diagnostic Access Virtual Address Format .....                    | 310 |
| FIGURE 12-1  | Processor State Diagram .....   | 317 |
| FIGURE 12-2  | Trap Vector Address .....   | 325 |
| FIGURE 12-3  | Trap Table Layout .....   | 326 |
| FIGURE 12-4  | Trap Type Encoding for Spill/Fill Traps .....                             | 331 |
| FIGURE 14-1  | Performance Control Register .....  | 364 |
| FIGURE 14-2  | Performance Instrumentation Counter Register .....                        | 366 |
| FIGURE 14-3  | Operational Flow Diagram for Controlling Event Counters .....             | 368 |

|             |  |     |
|-------------|--|-----|
| FIGURE 14-4 | Dispatch Counters .....  | 371 |
| FIGURE 14-5 | D-Cache Load Miss Stall Regions .....                          | 377 |
| FIGURE 14-6 | Data Locality Event Tree for L2-cache Misses .....             | 381 |
| FIGURE 15-1 | Instruction Cache Organization .....                           | 388 |
| FIGURE 15-2 | Handling of Conditional Branches .....                         | 391 |
| FIGURE 15-3 | Handling of MOVCC .....  | 391 |
| FIGURE 15-4 | Logical Organization of Data Cache .....                       | 393 |
| FIGURE 17-1 | Floating-point Number Line .....                               | 411 |
| FIGURE A-1  | Three-Dimensional Array Fixed-Point Address Format .....       | 458 |
| FIGURE A-2  | Three-Dimensional Array Blocked-Address Format (Array8) .....  | 458 |
| FIGURE A-3  | Three-Dimensional Array Blocked-Address Format (Array16) ..... | 458 |
| FIGURE A-4  | Three-Dimensional Array Blocked-Address Format (Array32) ..... | 459 |
| FIGURE A-5  | FMUL8x16 Operation .....                                       | 545 |
| FIGURE A-6  | FMUL8x16AU Operation .....                                     | 546 |
| FIGURE A-7  | FMUL8x16AL Operation .....                                     | 546 |
| FIGURE A-8  | FMUL8SUx16 Operation .....                                     | 547 |
| FIGURE A-9  | FMUL8LUx16 Operation .....                                     | 548 |
| FIGURE A-10 | FMULD8SUx16 Operation .....                                    | 549 |
| FIGURE A-11 | FMULD8ULx16 Operation .....                                    | 549 |
| FIGURE A-12 | FPACK16 Operation .....  | 555 |
| FIGURE A-13 | FPACK32 Operation .....  | 556 |
| FIGURE A-14 | FPACKFIX Operation .....                                       | 557 |
| FIGURE A-15 | FEXPAND Operation .....  | 558 |
| FIGURE A-16 | FPMERGE Operation .....  | 559 |





# List of Tables

---

|           |  |    |
|-----------|--|----|
| TABLE 1-1 | Integer/Floating-Point <code>unfinished_FPop</code> Exception Conditions ..... | 5  |
| TABLE 1-2 | RAS Features .....   | 11 |
| TABLE 4-1 | Processor Pipeline Stages .....  | 38 |
| TABLE 4-2 | Execution Pipelines .....  | 43 |
| TABLE 4-3 | SPARC V9 Conditional Moves .....   | 54 |
| TABLE 4-4 | Execution Pipelines .....  | 54 |
| TABLE 4-4 | UltraSPARC III Cu Instruction Latencies and Dispatching Properties .....       | 56 |
| TABLE 5-1 | Signed Integer, Unsigned Integer, and Tagged Integer Format Ranges .....       | 66 |
| TABLE 5-2 | Integer Data Alignment .....   | 67 |
| TABLE 5-3 | Floating-Point Doubleword and Quadword Alignment .....                         | 71 |
| TABLE 5-4 | Floating-Point Single-Precision Format Definition .....                        | 72 |
| TABLE 5-5 | Floating-Point Double-Precision Format Definition .....                        | 73 |
| TABLE 5-6 | Floating-Point Quad-Precision Format Definition .....                          | 74 |
| TABLE 5-7 | Pixel, Fixed16, and Fixed32 Data Alignment .....                               | 75 |
| TABLE 6-1 | Integer Unit General-Purpose Registers .....                                   | 81 |
| TABLE 6-2 | 32-bit Floating-Point Registers with Aliasing .....                            | 87 |
| TABLE 6-3 | 64-bit Floating-Point Registers with Aliasing .....                            | 87 |
| TABLE 6-4 | 128-bit Floating-Point Registers with Aliasing .....                           | 88 |

|            |  |     |
|------------|--|-----|
| TABLE 6-5  | Floating-Point Register Number Encoding .....                          | 88  |
| TABLE 6-6  | State and Ancillary State Registers .....                              | 92  |
| TABLE 6-7  | Privileged Registers .....   | 94  |
| TABLE 6-8  | ASI and Specially Accessed Registers .....                             | 95  |
| TABLE 6-9  | DCR Bit Description .....  | 101 |
| TABLE 6-10 | GSR Bit Description .....  | 103 |
| TABLE 6-11 | Register-Window State Registers .....                                  | 104 |
| TABLE 6-12 | SOFTINT Bit Descriptions .....   | 105 |
| TABLE 6-13 | Timer State Registers .....  | 106 |
| TABLE 6-14 | Trap Stack Register Power-On and Normal Operation .....                | 111 |
| TABLE 6-15 | PSTATE Global Register Selection Events .....                          | 113 |
| TABLE 6-16 | MM Encodings .....   | 115 |
| TABLE 6-17 | Register-Window State Privileged Registers .....                       | 118 |
| TABLE 6-18 | Processor Implementation Codes .....                                   | 121 |
| TABLE 6-19 | UltraSPARC III Cu Processor Mask Version Codes .....                   | 121 |
| TABLE 6-20 | Floating-Point Condition Codes (f <sub>ccn</sub> ) Fields of FSR ..... | 123 |
| TABLE 6-21 | Rounding Direction (RD) Field of FSR .....                             | 124 |
| TABLE 6-22 | Floating-Point Trap Type (ft) Field of FSR .....                       | 125 |
| TABLE 6-23 | Standard Conditions in Which unfinished_FPop Trap Type Can Occur ..... | 127 |
| TABLE 6-24 | Setting of FSR.cexc Bits .....   | 130 |
| TABLE 6-25 | DCUCR Bit Field Descriptions .....                                     | 133 |
| TABLE 6-26 | ASIs Affected by Watchpoint Traps .....                                | 136 |
| TABLE 7-1  | MOV <sub>r</sub> and FMOV <sub>r</sub> Test Conditions .....           | 145 |
| TABLE 7-2  | Control Transfer Characteristics .....                                 | 161 |
| TABLE 7-3  | Instruction Summary for the Integer Execution Environment .....        | 166 |
| TABLE 7-4  | Instruction Summary for the Floating-Point Execution Environment ..... | 170 |
| TABLE 7-5  | Instruction Summary for the VIS Execution Environment .....            | 172 |

|            |  |     |
|------------|--|-----|
| TABLE 7-6  | Instruction Summary for Data Coherency .....                                     | 173 |
| TABLE 7-7  | Instruction Summary for Register-Window Management .....                         | 173 |
| TABLE 7-8  | Instruction Summary for Program Control Transfer .....                           | 174 |
| TABLE 7-9  | Instruction Summary Table .....  | 174 |
| TABLE 7-10 | Instruction Field Interpretation .....   | 177 |
| TABLE 7-11 | Processor Actions on Unimplemented Instructions .....                            | 179 |
| TABLE 8-1  | Processor Specific ASIs .....  | 189 |
| TABLE 8-2  | ASI Summary Table .....  | 190 |
| TABLE 8-3  | Operand Alignment .....  | 192 |
| TABLE 8-4  | ASI Definitions .....  | 193 |
| TABLE 9-1  | MEMBAR Semantics .....   | 213 |
| TABLE 9-2  | MEMBAR Rules for Column VA <12:5> ≠ Row VA <12:5> While Desiring Strong Ordering | 214 |
| TABLE 9-3  | MEMBAR Rules for Column VA<12:5> = Row VA<12:5> While Desiring Strong Ordering   | 215 |
| TABLE 9-4  | ASIs That Support SWAP, LDSTUB, and CAS .....                                    | 217 |
| TABLE 9-5  | Types of Software Prefetch Instructions .....                                    | 219 |
| TABLE 10-1 | External Cache Organizations .....   | 228 |
| TABLE 10-2 | Hit/Miss, State Change, and Transaction Generated for Processor Action .....     | 234 |
| TABLE 10-3 | Combined Tag/MTag States .....   | 236 |
| TABLE 10-4 | Deriving DTags, CTags, and MTags from Combined Tags .....                        | 236 |
| TABLE 10-5 | Snoop Output and DTag Transition .....   | 236 |
| TABLE 10-6 | Snoop Input and CIQ Operation Queued .....                                       | 239 |
| TABLE 10-7 | Transaction Handling at Head of CIQ .....  | 240 |
| TABLE 10-8 | No Snoop Transaction Handling .....  | 242 |
| TABLE 10-9 | Internal Transaction Handling .....  | 243 |
| TABLE 11-1 | TSB and TTE Bit Description .....  | 248 |
| TABLE 11-2 | MMU Trap Types, Causes, and Stored State Register Update Policy .....            | 259 |
| TABLE 11-3 | ASI Mapping for Instruction Access .....   | 261 |

|             |   |     |
|-------------|---|-----|
| TABLE 11-4  | ASI Mapping for Data Accesses .....                                   | 262 |
| TABLE 11-5  | I-MMU and D-MMU Context Register Usage .....                          | 262 |
| TABLE 11-6  | MMU SPARC V9 Annex G Protection Mode Compliance .....                 | 264 |
| TABLE 11-7  | D-MMU Primary Context Register .....                                  | 265 |
| TABLE 11-8  | D-MMU Primary Context Register .....                                  | 266 |
| TABLE 11-9  | Four-way Pseudo Random Selection .....                                | 267 |
| TABLE 11-10 | Tag Access Extension Register for Saving Page Sizes Information ..... | 269 |
| TABLE 11-11 | D-TLB Access Number .....   | 270 |
| TABLE 11-12 | D-MMU TLB Access Summary .....  | 273 |
| TABLE 11-13 | D-MMU Table of Operations for Normal ASIs .....                       | 276 |
| TABLE 11-14 | MMU Internal Registers and ASI Operations .....                       | 277 |
| TABLE 11-15 | D-MMU TLB Tag Access Registers .....                                  | 279 |
| TABLE 11-16 | MMU TLB Access Summary .....  | 280 |
| TABLE 11-17 | TLB Data Access Register .....  | 282 |
| TABLE 11-18 | D-MMU TLB Tag Read Registers .....                                    | 283 |
| TABLE 11-19 | MMU Tag Target Registers .....  | 284 |
| TABLE 11-20 | TSB Register Description .....  | 284 |
| TABLE 11-21 | D-SFSR Bit Description .....  | 287 |
| TABLE 11-22 | MMU Synchronous Fault Status Register FT (Fault Type) Field .....     | 289 |
| TABLE 11-23 | Demap Address Format .....  | 291 |
| TABLE 11-24 | TLB Diagnostic Register Virtual Address Format .....                  | 292 |
| TABLE 11-25 | CAM Diagnostic Register .....   | 293 |
| TABLE 11-26 | Bypass Attribute Bits .....   | 294 |
| TABLE 11-27 | I-MMU TLB Access Summary .....  | 296 |
| TABLE 11-28 | I-MMU Table of Operations for Normal ASIs .....                       | 298 |
| TABLE 11-29 | I-MMU Tag Access Register .....                                       | 299 |
| TABLE 11-30 | MMU TLB Access Summary .....  | 300 |

|             |   |     |
|-------------|---|-----|
| TABLE 11-31 | TLB Data Access Register .....                                    | 301 |
| TABLE 11-32 | I-MMU TLB Tag Read Register .....                                 | 302 |
| TABLE 11-33 | MMU Tag Target Register .....                                     | 303 |
| TABLE 11-34 | TSB Register Description .....                                    | 304 |
| TABLE 11-35 | SFSR Bit Description .....  | 306 |
| TABLE 11-36 | MMU Synchronous Fault Status Register FT (Fault Type) Field ..... | 307 |
| TABLE 11-37 | Demap Address Format .....  | 309 |
| TABLE 11-38 | MMU TLB Diagnostic Register Virtual Address .....                 | 311 |
| TABLE 12-1  | RED_state Trap Vector Layout .....                                | 318 |
| TABLE 12-2  | Typical Usage for Trap Levels .....                               | 320 |
| TABLE 12-3  | Exception and Interrupt Requests, by TT Value .....               | 327 |
| TABLE 12-4  | Exception and Interrupts Requests in Priority Order .....         | 329 |
| TABLE 12-5  | Trap Types for Spill/Fill Traps .....                             | 330 |
| TABLE 12-6  | Two New Traps Added for D-cache and I-cache Parity Error .....    | 332 |
| TABLE 12-7  | Trap Received While in execute_state .....                        | 334 |
| TABLE 12-8  | Trap Received While in RED_state .....                            | 335 |
| TABLE 12-9  | Reset Received While in error_state .....                         | 335 |
| TABLE 13-1  | Outgoing Interrupt Vector Data Register Format .....              | 354 |
| TABLE 13-2  | Interrupt Vector Dispatch Register Format .....                   | 355 |
| TABLE 13-3  | Interrupt Dispatch Status Register Format .....                   | 356 |
| TABLE 13-4  | Incoming Interrupt Vector Data Register Format .....              | 357 |
| TABLE 13-5  | Interrupt Receive Register Format .....                           | 357 |
| TABLE 13-6  | SOFTINT ASRs .....  | 359 |
| TABLE 14-1  | PCR Bit Description .....   | 365 |
| TABLE 14-2  | PIC Register Fields .....   | 367 |
| TABLE 14-3  | PIC Counter Overflow Processor Compatibility Comparison .....     | 367 |
| TABLE 14-4  | Instruction Execution Clock Cycles and Counts .....               | 369 |

|             |  |     |
|-------------|--|-----|
| TABLE 14-5  | Counters for Collecting IIU Statistics .....               | 370 |
| TABLE 14-6  | Counters for IIU Stalls .....                              | 371 |
| TABLE 14-7  | Counters for R-stage Stalls .....                          | 372 |
| TABLE 14-8  | Counters for Recirculation .....                           | 372 |
| TABLE 14-9  | Counters for Instruction Cache Events .....                | 373 |
| TABLE 14-10 | Counters for Data Cache Events .....                       | 373 |
| TABLE 14-11 | Counters for Write Cache Events .....                      | 374 |
| TABLE 14-12 | Counters for Prefetch Cache Events .....                   | 374 |
| TABLE 14-13 | Counters for L2-Cache Events .....                         | 375 |
| TABLE 14-14 | Re_DC_missovhd Stall Counter Processor Compatibility ..... | 376 |
| TABLE 14-15 | Counters for Memory Controller Read Events .....           | 378 |
| TABLE 14-16 | Counters for Memory Controller Write Events .....          | 378 |
| TABLE 14-17 | Counters for Memory Controller Stall Cycles .....          | 379 |
| TABLE 14-18 | Counters for Data Locality Events .....                    | 379 |
| TABLE 14-19 | Data Locality Events .....                                 | 382 |
| TABLE 14-20 | Synthesized Data Locality Events .....                     | 383 |
| TABLE 14-21 | Counters for System Interface Statistics .....             | 384 |
| TABLE 14-22 | Counters for Software Statistics .....                     | 384 |
| TABLE 14-23 | Counters for Floating-Point Operation Statistics .....     | 385 |
| TABLE 14-24 | PIC.SL and PIC.SU Selection Bit Field Encoding .....       | 385 |
| TABLE 16-1  | Types of Software Prefetch Instructions .....              | 399 |
| TABLE 17-1  | FSR.RD bit options .....                                   | 408 |
| TABLE 17-2  | Floating-point Numbers .....                               | 409 |
| TABLE 17-3  | Floating-point Addition .....                              | 412 |
| TABLE 17-4  | Floating-point Subtraction .....                           | 413 |
| TABLE 17-5  | Floating-point Multiplication .....                        | 414 |
| TABLE 17-6  | Floating-point Division .....                              | 415 |

|             |   |     |
|-------------|---|-----|
| TABLE 17-7  | Floating-point Square Root .....  | 416 |
| TABLE 17-8  | Number Compare .....  | 416 |
| TABLE 17-9  | Precision Conversion .....  | 417 |
| TABLE 17-10 | Floating-point to Integer Number Conversion .....                                     | 418 |
| TABLE 17-11 | Integer to Floating-point Number Conversion .....                                     | 419 |
| TABLE 17-12 | Floating-point Unit Exceptions .....  | 421 |
| TABLE 17-13 | Response to Traps .....   | 422 |
| TABLE 17-14 | FP $\leftrightarrow$ Integer Conversions that Generate Inexact Exceptions .....       | 424 |
| TABLE 17-15 | Underflow Exception Summary .....   | 425 |
| TABLE 17-16 | Results from NaN Operands .....   | 427 |
| TABLE 17-17 | Subnormal Handling Constants per Destination Register Precision .....                 | 430 |
| TABLE 18-1  | Machine State After Reset and in RED_state .....                                      | 438 |
| TABLE A-1   | Opcode Superscripts .....   | 448 |
| TABLE A-2   | Instruction Set .....   | 448 |
| TABLE A-3   | Three-Dimensional r[rs2] Array X/Y Dimensions .....                                   | 458 |
| TABLE A-4   | Edge Mask Specification .....   | 482 |
| TABLE A-5   | Edge Mask Specification (Little-Endian) .....   | 483 |
| TABLE A-6   | Floating-Point to Integer unfinished_FPop Exception Conditions .....                  | 489 |
| TABLE A-7   | Floating-Point/Floating-Point unfinished_FPop Exception Conditions .....              | 491 |
| TABLE A-8   | Integer/Floating-Point unfinished_FPop Exception Conditions .....                     | 492 |
| TABLE A-9   | MEMBAR mmask Encodings .....  | 520 |
| TABLE A-10  | MEMBAR cmask Encodings .....  | 520 |
| TABLE A-11  | MEMBAR Rules for Column VA <12:5> $\neq$ Row VA <12:5> While Desiring Strong Ordering | 522 |
| TABLE A-12  | MEMBAR Rules for Column VA <12:5> = Row VA <12:5> While Desiring Strong Ordering      | 523 |
| TABLE A-13  | Types of Software Prefetch Instructions .....   | 562 |
| TABLE A-14  | Error Handling of Prefetch Requests .....   | 564 |
| TABLE A-15  | Shift Count Encodings .....   | 579 |
| TABLE A-16  | UDIV / UDIVcc Overflow Detection and Value Returned .....                             | 609 |
| TABLE A-17  | SDIV / SDIVcc Overflow Detection and Value Returned .....                             | 610 |

|            |   |     |
|------------|---|-----|
| TABLE A-18 | UMULcc / SMULcc Condition Code Settings ..... | 617 |
| TABLE B-1  | Value Names and Values .....                  | 635 |
| TABLE B-2  | Prefetch Function Values .....                | 637 |



# Preface

---

Welcome to the UltraSPARC® III Cu User's Manual. This book contains information about the architecture and programming of the UltraSPARC III Cu processor, one of Sun Microsystems' family of SPARC® V9-compliant processors.

---

## Target Audience

This user's manual is mainly targeted for programmers who write software for the UltraSPARC III Cu processor. This user's manual contains a depository of information that is useful to operating system programmers, application software programmers, logic designers and third party vendors, who are trying to understand the architecture and operation of the UltraSPARC III Cu processor. This manual is both a guide and a reference manual for low-level programming of the processor.

---

## A Brief History of SPARC

SPARC stands for **Scalable Processor ARCHitecture**, which was first announced in 1987. Unlike more traditional processor architectures, SPARC is an open standard, freely available through license from SPARC International, Inc. Any company that obtains a license can manufacture and sell a SPARC-compliant processor.

By the early 1990s, SPARC processors were available from over a dozen different vendors, and over 8,000 SPARC-compliant applications had been certified.

In 1994, SPARC International, Inc. published *The SPARC Architecture Manual, Version 9*, which defined a powerful 64-bit enhancement to the SPARC architecture. SPARC V9 provided support for the following:

- 64-bit virtual addresses and 64-bit integer data

- Fault tolerance
- Fast trap handling and context switching
- Big-endian and little-endian byte orders

The UltraSPARC processor is the first family of SPARC V9-compliant processors available from Sun Microsystems, Inc.

---

## Prerequisites

This user's manual is a companion to *The SPARC Architecture Manual, Version 9*. The reader of this user's manual should be familiar with the contents of *The SPARC Architecture Manual, Version 9*, which is available from many technical bookstores or directly from its copyright holder:

SPARC International, Inc.  
2242 Camden Ave, Suite #105  
San Jose, CA 95124  
(408) 558-8111  
<http://www.sparc.org>

*The SPARC Architecture Manual, Version 9* provides a complete description of the SPARC V9 architecture. Since SPARC V9 is an open architecture, many of the implementation decisions have been left to the manufacturers of SPARC-compliant processors. These “implementation dependencies” are introduced in *The SPARC Architecture Manual, Version 9*.

---

## User's Manual Overview

This manual is focused on the treatment of the UltraSPARC III Cu processor. However, sometimes it refers to the UltraSPARC III family of processors to indicate generality of a certain feature. The term “UltraSPARC III family of processors” refers to processors that are similar to the UltraSPARC III Cu processor.

This manual is divided into multiple sections. The following sections are described:

## *Processor Introduction*

The processor introduction section describes the high level features of the UltraSPARC III Cu processor. This section also discusses how the UltraSPARC III Cu processor is used in a system.

## *Architecture and Functions*

This section discusses the details of the UltraSPARC III Cu architecture and the functions of various CPU units. An entire chapter is devoted to a discussion on the instruction execution pipeline.

## *Execution Environment*

This section describes the details necessary to understand the execution environment. Various topics such as memory models, data formats, registers and instruction types are discussed.

## *Supervisor Programming*

Supervisor software controls the processor and the instruction execution environment for itself and application programs. Chapters are devoted to trap and interrupt handling.

## *Performance Programming*

This section explores the opportunities to exploit the high-performance architecture of the processor. Chapters are devoted to performance instrumentation and prefetch, two special features of the UltraSPARC III Cu processor.

## *Instruction Definitions Appendix*

This section describes, in detail, each instruction for the UltraSPARC III Cu processor.

# SPARC V9 Architecture

*The SPARC Architecture Manual, Version 9* was used to implement the CPU in the processor to insure SPARC compatibility for user and application programs. The SPARC V9 manual provides important theoretical information for operating system programmers who write memory management software, compiler writers who write machine-specific optimizers, and

anyone who writes code to run on all SPARC V9 compatible machines. Book copies of the *The SPARC Architecture Manual, Version 9* are readily available at bookstores or from SPARC International, Inc.

Software that is intended to be portable across all SPARC V9 processors should adhere to *The SPARC Architecture Manual, Version 9*.

In this book, the word *architecture* refers to the machine details that are visible to an assembly language programmer or to the compiler code generator. It does not, necessarily, include details of the implementation that are not visible or easily observable by software. Where such details are provided, the intent is to enable faster and better programs.

---

## Textual Usage

### Fonts

Fonts are used as follows:

- *Italic* font is used for emphasis, book titles, the first instance of a word that is defined, and assembly language terms.
- *Italic sans serif* font is used for exception and trap names. “The *privileged\_action* exception...” is an example of how this font is used.
- *Typewriter* font (Courier) is used for register fields (named bits), instruction fields, and read-only register fields. “The `rs1` field contains...” is an example of how this font is used. It is also used for literals, instruction names, register names, and software examples.
- UPPERCASE items are acronyms, instruction names, or writable register fields. Some common acronyms are listed in Acronyms and Definitions.

---

**Note** – Names of some instructions contain both uppercase and lowercase letters.

---

- Underbar characters join words in register, register field, exception, and trap names.

---

**Note** – Such words can be split across lines at the underbar without an intervening hyphen. “This is true whenever the `integer_condition_code` field...” is an example of how the underbar characters are used.

---

# Notational Conventions

The following notational conventions are used:

- Square brackets, [ ], indicate a numbered register in a register file. For example,  $r[0]$  translates to register 0.
- Angle brackets, < >, indicate a bit number or colon-separated range of bit numbers within a field. “Bits  $FSR<29:28>$  and  $FSR<12>$  are...” is an example of how the angle brackets are used.
- Curly braces, { }, indicate textual substitution. For example, the string “PRIMARY{ \_LITTLE}” expands to “ASI\_PRIMARY” and “ASI\_PRIMARY\_LITTLE.”
- If the bar, |, is used with the curly braces, it represents multiple substitutions. For example, the string “ASI\_DMMU\_TSB\_{8KB|64KB|DIRECT}\_PTR\_REG” expands to “ASI\_DMMU\_TSB\_8KB\_PTR\_REG,” “ASI\_DMMU\_TSB\_64KB\_PTR\_REG,” and “ASI\_DMMU\_TSB\_DIRECT\_PTR\_REG.”
- The  $\square$  symbol designates concatenation of bit vectors. A comma (,) on the left side of an assignment separates quantities that are concatenated for the purpose of assignment. For example, if X, Y, and Z are 1-bit vectors and the 2-bit vector T equals  $11_2$ , then
$$(X, Y, Z) \leftarrow 0 \square T$$
results in  $X = 0$ ,  $Y = 1$ , and  $Z = 1$ .
- “A mod B” means “A modulus B,” where the calculated value is the remainder when A is divided by B.

## Notation for Numbers

Numbers throughout this specification are decimal (base-10) unless otherwise indicated. Numbers in other bases are followed by a numeric subscript indicating their base (for example,  $1001_2$ ,  $FFFF\ 0000_{16}$ ). In some cases, numbers may be preceded by “0x” to indicate hexadecimal (base-16) notation (for example,  $0xFFFF.0000$ ). Long binary and hexadecimal numbers within the text have spaces or periods inserted every four characters to improve readability.

The notation  $7h'1F$  indicates a hexadecimal number of  $1F_{16}$  with seven binary bits of width.

## Informational Notes

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

---

**Programming Note** – Programming notes contain incidental information about programming the UltraSPARC III Cu processor unless otherwise restricted to a particular processor in the family.

---

---

**Implementation Note** – Implementation notes contain information that contains implementation specific information the UltraSPARC III Cu processor compared to other UltraSPARC processors.

---

---

**Compatibility Note** – Compatibility notes contain information relevant to the previous SPARC V8 architecture.

---

---

**UltraSPARC Note** – UltraSPARC notes highlight the differences between the UltraSPARC I and UltraSPARC II processors and the UltraSPARC III family of processors. This note shows architectural and functional differences that may be generalized or applicable to one particular processor in one of the families. Check the appropriate User's Manual or section in this User's Manual to determine individual processor functionality as needed.

---

---

**Note** – This highlights a useful note regarding important and informative processor architecture or functional operation. This may be used for purposes not covered in one of the other notes.

---

# Acronyms and Definitions

---

This chapter defines concepts and terminology common to all implementations of SPARC V9.

- address space identifier** *See ASI.*
- AFAR** Asynchronous Fault Address Register.
- AFSR** Asynchronous Fault Status Register.
- aliased** Two virtual addresses that refer to the same physical address.
- application program** A program executed with the processor in non-privileged *mode*. **Note:** Statements made in this specification regarding application programs may not be applicable to programs (for example, debuggers) that have access to *privileged* processor state (for example, as stored in a memory-image dump).
- ASI** Address space identifier. An 8-bit value that identifies an address space. For each instruction or data access, the integer unit appends an ASI to the address. *See also implicit ASI.*
- ASR** Ancillary State Register.
- Ax** Either the A0 or A1 pipeline.
- BBC** Bootbus controller (UltraSPARC III Cu processors).
- big-endian** An addressing convention. Within a multiple-byte integer, the byte with the smallest address is the most significant; a byte's significance decreases as its address increases.
- BLD** Block load.
- BST** Block store.
- byte** Eight consecutive bits of data.
- CDS** Crossbar Data Switch. Data bus crossbars for the Sun<sup>TM</sup> Fireplane interconnect of the UltraSPARC III Cu Processor. Also known as Dual CPU Data Switch (**DCDS**).
- clean window** A register window in which all of the registers contain zero, a valid address from the current address space, or valid data from the current address space.

- coherence** A set of protocols guaranteeing that all memory accesses are globally visible to all caches on a shared-memory bus.
- completed** A memory transaction is completed when an idealized memory has executed the transaction with respect to all processors. A load is considered completed when no subsequent memory transaction can affect the value returned by the load. A store is considered completed when no subsequent load can return the value that was overwritten by the store.
- consistency** *See coherence.*
- context** A set of translations that supports a particular address space. *See also Memory Management Unit (MMU).*
- copyback** The process of copying back a dirty cache line in response to a cache hit while snooping.
- CPI** Cycles per instruction. The number of clock cycles it takes to execute an instruction.
- cross-call** An interprocessor call in a multiprocessor system.
- CSR** Control Status Register.
- current window** The block of 24  $x$  registers that is currently in use. The Current Window Pointer (CWP) register points to the current window.
- D-cache** Level-1 data memory cache.
- DCTI** Delayed control transfer instruction.
- DCU** Data Cache Unit. Includes controller and Tag and Data RAM arrays.
- demap** To invalidate a mapping in the MMU.
- deprecated** The term applied to an architectural feature (such as an instruction or register) for which a SPARC V9 implementation provides support only for compatibility with previous versions of the architecture. Use of a deprecated feature must generate correct results but may compromise software performance. Deprecated features should not be used in new SPARC V9 software and may not be supported in future versions of the architecture.
- DFT** Designed for test.
- DIMM** Dual In-line Memory Module. Provides a single or double bank of SDRAM devices 72 bits or 144 bits of data width.
- dispatch** To send a previously fetched instruction to one or more functional units for execution. Typically, the instruction is dispatched from a reservation station or other buffer of instructions waiting to be executed. *See also issued.*
- doublet** Two bytes (16 bits) of data.
- doubleword** An aligned octlet. **Note:** The definition of this term is architecture dependent and may differ from that used in other processor architectures.



|   |   |
|---|---|
| <b>DQM</b>  | Data input/output Mask. Q stands for either input or output.  |
| <b>D-TLB</b>                                      | Data Translation Lookaside Buffer.  |
| <b>ECU</b>  | External or embedded Cache Unit controller.   |
| <b>EMU</b>  | External Memory Unit. A combination of the ECU and the Memory Control Unit (MCU).   |
| <b>exception</b>                                  | A condition that makes it impossible for the processor to continue executing the current instruction stream without software intervention. <i>See also</i> <b>trap</b> .  |
| <b>extended word</b>                              | An aligned octlet, nominally containing integer data. <b>Note:</b> The definition of this term is architecture dependent and may differ from that used in other processor architectures.  |
| <b>f register</b>                                 | A floating-point register. SPARC V9 includes single-, double-, and quad-precision f registers.  |
| <b>fccN</b>                                       | One of the floating-point condition code fields <code>fcc0</code> , <code>fcc1</code> , <code>fcc2</code> , or <code>fcc3</code> .  |
| <b>FFA or FGA or FP1</b>                          | Floating-point/Graphics ALU pipeline.   |
| <b>FGM or FP0</b>                                 | Floating-point/Graphics Multiply pipeline.  |
| <b>FGU</b>  | Floating-point and Graphics Unit (FP0 and FP1).   |
| <b>floating-point exception</b>                   | An exception that occurs during the execution of a Floating-point operate (FPop) instruction while the corresponding bit in <code>FSR.TEM</code> is set to one. The exceptions are <i>unfinished_FPop</i> , <i>unimplemented_FPop</i> , <i>sequence_error</i> , <i>hardware_error</i> , <i>invalid_fp_register</i> , or <i>IEEE_754_exception</i> . |
| <b>floating-point IEEE-754 exception</b>          | A floating-point exception, as specified by IEEE Standard 754-1985. Listed within this specification as <i>IEEE_754_exception</i> .   |
| <b>floating-point operate (FPop) instructions</b> | Instructions that perform floating-point calculations, as defined by the <code>FPop1</code> and <code>FPop2</code> opcodes. FPop instructions do not include <code>Fbfc</code> instructions or loads and stores between memory and the floating-point unit.   |
| <b>floating-point trap type</b>                   | The specific type of a floating-point exception, encoded in the <code>FSR.ftt</code> field.   |
| <b>floating-point unit</b>                        | A processing unit that contains the floating-point registers and performs floating-point operations, as defined by this specification.  |
| <b>FPRS</b>                                       | Floating-point Register State.  |
| <b>FPU</b>  | Floating-point unit.  |
| <b>FRF</b>  | Floating-point Register File.   |
| <b>FSR</b>  | Floating-point Status Register.   |

|                                     |  |
|-------------------------------------|--|
| <b>halfword</b>                     | An aligned doublet. <b>Note:</b> The definition of this term is architecture dependent and may differ from that used in other processor architectures.   |
| <b>HBM</b>                          | Hierarchical Bus Mode.   |
| <b>hexlet</b>                       | Sixteen bytes (128 bits) of data.  |
| <b>HPE</b>                          | Hardware Prefetch Enable.  |
| <b>I-cache</b>                      | Level-2 Instruction memory cache.  |
| <b>IEU</b>                          | Instruction Execution Unit.  |
| <b>IIU</b>                          | Instruction Issue Unit.  |
| <b>implementation</b>               | Hardware or software that conforms to all of the specifications of an instruction set architecture (ISA).  |
| <b>implementation dependent</b>     | An aspect of the architecture that can legitimately vary among implementations. In many cases, the permitted range of variation is specified in the SPARC V9 standard. When a range is specified, compliant implementations must not deviate from that range.  |
| <b>implicit ASI</b>                 | The ASI that is supplied by the hardware on all instruction accesses and on data accesses that do not contain an explicit ASI or a reference to the contents of the ASI register.  |
| <b>informative appendix</b>         | An appendix containing information that is useful but not required to create an implementation that conforms to the SPARC V9 specification. <i>See also</i> <b>normative appendix</b> .  |
| <b>initiated</b>                    | <i>Synonym:</i> <b>issued</b> .  |
| <b>instruction field</b>            | A bit field within an instruction word.  |
| <b>instruction group</b>            | One or more independent instructions that can be dispatched for simultaneous execution.  |
| <b>instruction set architecture</b> | <i>See</i> <b>ISA</b> .  |
| <b>integer unit</b>                 | A processing unit that performs integer and control-flow operations and contains general-purpose integer registers and processor state registers, as defined by this specification.  |
| <b>interrupt request</b>            | A request for service presented to the processor by an external device.  |
| <b>ISA</b>                          | Instruction set architecture. A set that defines instructions, registers, instruction and data memory, the effect of executed instructions on the registers and memory, and an algorithm for controlling instruction execution. It does not define clock cycle times, cycles per instruction, data paths, etc. |

|                               |   |
|-------------------------------|---|
| <b>issued</b>                 | (1) A memory transaction (load, store, or atomic load-store) is “issued” when a processor has sent the transaction to the memory subsystem and the completion of the request is out of the processor’s control. <i>Synonym: initiated.</i><br>(2) An instruction (or sequence of instructions) is said to be <i>issued</i> when released from the processor's in-order instruction fetch unit. Typically, instructions are issued to a reservation station or other buffer of instructions waiting to be executed. (Other conventions for this term exist, but this document attempts to use “issue” consistently as defined here). <i>See also dispatched.</i> |
| <b>I-TLB</b>                  | Instruction Translation Lookaside Buffer.   |
| <b>I-TSB</b>                  | Instruction Translation Storage Buffer.   |
| <b>IU</b>                     | Integer Unit.   |
| <b>L2-cache</b>               | Second level cache.   |
| <b>leaf procedure</b>         | A procedure that is a leaf in the program’s call graph, that is, one that does not call (by using CALL or JMPL) any other procedures.   |
| <b>little-endian</b>          | An addressing convention. Within a multiple-byte integer, the byte with the smallest address is the least significant; a byte’s significance increases as its address increases.  |
| <b>load</b>                   | An instruction that reads (but does not write) memory or reads (but does not write) location(s) in an alternate address space. <i>Load</i> includes loads into integer or floating-point registers, block loads, Load Quadword Atomic, and alternate address space variants of those instructions. <i>See also load-store</i> and <b>store</b> , the definitions of which are mutually exclusive with <i>load</i> .   |
| <b>load-store</b>             | An instruction that explicitly both reads and writes memory or explicitly reads and writes location(s) in an alternate address space. <i>Load-store</i> includes instructions such as CASA, CASXA, LDSTUB, and the deprecated SWAP instruction. <i>See also load</i> and <b>store</b> , the definitions of which are mutually exclusive with <i>load-store</i> .  |
| <b>LPA</b>                    | Local Physical (or Processor) Address. Used in the context of Scalable Shared Memory (SSM) system architectures.  |
| <b>may</b>                    | A keyword indicating flexibility of choice with no implied preference. <b>Note:</b> “May” indicates that an action or operation is allowed; “can” indicates that it is possible.  |
| <b>MCU</b>                    | Memory Control Unit. Controls the SDRAM signals.  |
| <b>Memory Management Unit</b> | <i>See MMU.</i>   |
| <b>Microtag</b>               | A partial virtual address tag used for early way select of a virtually indexed, physically tagged set associative cache. Microtag is often referred to as utag or Utag in this documentation.   |

|                             |  |
|-----------------------------|--|
| <b>MMU</b>                  | Memory Management Unit. The address translation hardware in the UltraSPARC III Cu implementation that translates 64-bit virtual address into physical addresses. The MMU is composed of the translation lookaside buffers (TLBs), ASRs, and ASI registers used to manage address translation. <i>See also</i> <b>context</b> , <b>physical address</b> , and <b>virtual address</b> .  |
| <b>module</b>               | A master or slave device that attaches to the shared-memory bus.   |
| <b>MOESI</b>                | A cache-coherence protocol. Each of the letters stands for one of the states that a cache line can be in, as follows: <b>M</b> , modified, dirty data with no outstanding shared copy; <b>O</b> , owned, dirty data with outstanding shared copy(s); <b>E</b> , exclusive, clean data with no outstanding shared copy; <b>S</b> , shared, clean data with outstanding shared copy(s); <b>I</b> , invalid, invalid data.  |
| <b>must</b>                 | <i>Synonym:</i> <b>shall</b> .   |
| <b>NaN</b>                  | Not a Number.  |
| <b>NCPQ</b>                 | Non-coherent pending queue.  |
| <b>next program counter</b> | <i>See</i> <b>nPC</b> .  |
| <b>NFO</b>                  | Non-fault access only.   |
| <b>non-faulting load</b>    | A load operation that, in the absence of faults or in the presence of a recoverable fault, completes correctly, and in the presence of an unrecoverable fault returns (with the assistance of system software) a known data value (nominally zero). <i>See also</i> <b>speculative load</b> .  |
| <b>non-privileged</b>       | An adjective that describes: <ul style="list-style-type: none"> <li>(1) the state of the processor when <code>PSTATE.PRIV = 0</code>, that is, non-privileged mode;</li> <li>(2) processor state information that is accessible to software while the processor is in either privileged mode or non-privileged mode; for example, non-privileged registers, non-privileged ASRs, or, in general, non-privileged state;</li> <li>(3) an instruction that can be executed when the processor is in either privileged mode or non-privileged mode.</li> </ul> |
| <b>non-privileged mode</b>  | The mode in which a processor is operating when <code>PSTATE.PRIV = 0</code> . <i>See also</i> <b>privileged</b> .   |
| <b>normative appendix</b>   | An appendix containing specifications that must be met by an implementation conforming to the SPARC V9 specification. <i>See also</i> <b>informative appendix</b> .  |
| <b>nPC</b>                  | Next program counter. A register that contains the address of the next executed instruction if a trap does not occur.  |
| <b>NPT</b>                  | Non-privileged trap.   |
| <b>NWINDOWS</b>             | The number of register windows present in a particular implementation.   |
| <b>OBP</b>                  | OpenBoot™ PROM.  |

|                         |  |
|-------------------------|--|
| <b>octlet</b>           | Eight bytes (64 bits) of data. Not to be confused with “octet,” which has been commonly used to describe eight bits of data. In this document, the term <i>byte</i> , rather than octet, is used to describe eight bits of data.   |
| <b>opcode</b>           | A bit pattern that identifies a particular instruction.  |
| <b>optional</b>         | A feature not required for SPARC V9 compliance.  |
| <b>ORQ</b>              | Outgoing request queue.  |
| <b>PA</b>               | Physical address. An address that maps real physical memory or I/O device space. <i>See also virtual address.</i>  |
| <b>Page Table Entry</b> | <i>See PTE.</i>  |
| <b>PC</b>               | Program counter. A register that contains the address of the instruction currently being executed by the IU.   |
| <b>PCR</b>              | Performance Control Register.  |
| <b>physical address</b> | <i>See PA.</i>   |
| <b>PIC</b>              | Performance Instrumentation Counter.   |
| <b>PIO</b>              | Programmed I/O.  |
| <b>PIPT</b>             | Physically indexed, physically tagged.   |
| <b>PIVT</b>             | Physically indexed, virtually tagged.  |
| <b>POR</b>              | Power-on Reset. The most aggressive reset.   |
| <b>prefetchable</b>     | <p>(1) An attribute of a memory location that indicates to an MMU that PREFETCH operations to that location may be applied.</p> <p>(2) A memory location condition for which the system designer has determined that no undesirable effects will occur if a PREFETCH operation to that location is allowed to succeed. Typically, normal memory is prefetchable.</p> <p>Non-prefetchable locations include those that, when read, change state or cause external events to occur. For example, some I/O devices are designed with registers that clear on read; others have registers that initiate operations when read. <i>See also side-effect.</i></p> |
| <b>privileged</b>       | <p>An adjective that describes:</p> <p>(1) the state of the processor when <code>PSTATE.PRIV = 1</code>, that is, <i>privileged mode</i>;</p> <p>(2) processor state that is only accessible to software while the processor is in <i>privileged mode</i>; for example, privileged registers, privileged ASRs, or, in general, privileged state;</p> <p>(3) an instruction that can be executed only when the processor is in <i>privileged mode</i>.</p>  |
| <b>privileged mode</b>  | The mode in which a processor is operating when <code>PSTATE.PRIV = 1</code> . <i>See also non-privileged.</i>   |
| <b>processor</b>        | The combination of the integer unit and the floating-point unit.   |

|                        |  |
|------------------------|--|
| <b>program counter</b> | <i>See PC.</i>   |
| <b>PSO</b>             | Partial store order.   |
| <b>PTA</b>             | Pending tag array.   |
| <b>PTE</b>             | Page Table Entry. Describes the virtual-to-physical translation and page attributes for a specific page. A PTE generally means an entry in the page table or in the TLB; however, it is sometimes used as an entry in the translation storage buffer (TSB). In general, a PTE contains fewer fields than a TTE. <i>See also TLB and TSB.</i>   |
| <b>QNaN</b>            | Quiet Not a Number.  |
| <b>quadlet</b>         | Four bytes (32 bits) of data.  |
| <b>quadword</b>        | Aligned hexlet. <b>Note:</b> The definition of this term is architecture dependent and may be different from that used in other processor architectures.   |
| <b>r register</b>      | An integer register. Also called a general-purpose register or working register.   |
| <b>RAW</b>             | Read-After-Write.  |
| <b>RD</b>              | Rounding direction.  |
| <b>RDPR</b>            | Read Privileged Register.  |
| <b>RED_state</b>       | <b>Reset, Error, and Debug state.</b> The processor state when <code>PSTATE.RED = 1</code> . A restricted execution environment used to process resets and traps that occur when <code>TL = MAXTL - 1</code> .   |
| <b>reserved</b>        | Describes an instruction field, certain bit combinations within an instruction field, or a register field that is reserved for definition by future versions of the architecture.<br><br><i>Reserved instruction fields</i> shall read as zero, unless the implementation supports extended instructions within the field. The behavior of SPARC V9 processors when they encounter nonzero values in reserved instruction fields is undefined.<br><i>Reserved bit combinations within instruction fields</i> are defined in Appendix A, Instruction Definitions. In all cases, SPARC V9 processors shall decode and trap on these reserved combinations.<br><i>Reserved register fields</i> should always be written by software with values of those fields previously read from that register or with zeroes; they should read as zero in hardware. Software intended to run on future versions of SPARC V9 should not assume that these fields will read as zero or any other particular value. Throughout this specification, figures and tables illustrating registers and instruction encodings indicate reserved fields and combinations with an em dash (—). |
| <b>reset trap</b>      | A vectored transfer of control to privileged software through a fixed-address reset trap table. Reset traps cause entry into <code>RED_state</code> .  |
| <b>restricted</b>      | Describes an ASI that may be accessed only while the processor is operating in privileged mode.  |
| <b>RMO</b>             | Relaxed memory order.  |

|                     |   |
|---------------------|---|
| <b>rs1, rs2, rd</b> | The integer or floating-point register operands of an instruction. The source registers are rs1 and rs2; the destination register is rd.  |
| <b>RTO</b>          | Read to own.  |
| <b>RTOR</b>         | Read to own remote. A reissued RTO transaction.   |
| <b>RTS</b>          | Read to share.  |
| <b>RTSM</b>         | Read to share Mtag. An RTS to modify MTag transaction.  |
| <b>SAM</b>          | SPARC Architecture Manual, Version 9.   |
| <b>scrub</b>        | Writes data from the W-cache to the L2-cache.   |
| <b>SDRAM</b>        | Synchronous Dynamic Random Access Memory. May be prefaced with DDR, double data rate SDRAM.   |
| <b>SFAR</b>         | Synchronous Fault Address Register.   |
| <b>SFSR</b>         | Synchronous Fault Status Register.  |
| <b>shall</b>        | A keyword indicating a mandatory requirement. Designers shall implement all such mandatory requirements to ensure interoperability with other SPARC V9 compliant products. <i>Synonym: must.</i>  |
| <b>should</b>       | A keyword indicating flexibility of choice with a strongly preferred implementation. <i>Synonym: it is recommended.</i>   |
| <b>SIAM</b>         | Set interval arithmetic mode instruction.   |
| <b>side-effect</b>  | The result of a memory location having additional actions beyond the reading or writing of data. A side-effect can occur when a memory operation on that location is allowed to succeed. Locations with side-effects include those that, when accessed, change state or cause external events to occur. For example, some I/O devices contain registers that clear on read; others have registers that initiate operations when read. <i>See also prefetchable.</i> |
| <b>SIG</b>          | Single-Instruction Group. Sometimes shortened to “single-group.”  |
| <b>SIR</b>          | Software-initiated reset.   |
| <b>SIU</b>          | System Interface Unit (Sun Fireplane interconnect).   |
| <b>SNaN</b>         | Signalling Not a Number.  |
| <b>snooping</b>     | The process of maintaining coherency between caches in a shared-memory bus architecture. All cache controllers monitor (snoop) the bus to determine whether they have a copy of the shared cache block.   |
| <b>SPE</b>          | Software prefetch enable.   |

|                                     |   |
|-------------------------------------|---|
| <b>speculative load</b>             | A load operation that is issued by the processor speculatively, that is, before it is known whether the load will be executed in the flow of the program. Speculative accesses are used by hardware to speed program execution and are transparent to code. An implementation, through a combination of hardware and system software, must nullify speculative loads on memory locations that have side-effects; otherwise, such accesses produce unpredictable results. <i>Contrast with <b>non-faulting load</b></i> , which is an explicit load that always completes, even in the presence of recoverable faults. |
| <b>SSM</b>                          | Scalable shared memory. A directory based data coherency mechanism.   |
| <b>store</b>                        | An instruction that writes (but does not explicitly read) memory or writes (but does not explicitly read) location(s) in an alternate address space. <i>Store</i> includes stores from either integer or floating-point registers, block stores, partial store, and alternate address space variants of those instructions. <i>See also <b>load</b> and <b>load-store</b></i> , the definitions of which are mutually exclusive with <i>store</i> .   |
| <b>superscalar</b>                  | An implementation that allows several instructions to be issued, executed, and committed in one clock cycle.  |
| <b>supervisor software</b>          | Software that executes when the processor is in privileged mode.  |
| <b>TBA</b>                          | Trap base address.  |
| <b>TLB</b>                          | Translation Lookaside Buffer. A cache within an MMU that contains recent partial translations. TLBs speed up closely following translations by often eliminating the need to reread PTE from memory.  |
| <b>TLB hit</b>                      | The desired translation is present in the on-chip TLB.  |
| <b>TLB miss</b>                     | The desired translation is not present in the on-chip TLB.  |
| <b>TPC</b>                          | Trap-saved PC.  |
| <b>Translation Lookaside Buffer</b> | <i>See <b>TLB</b>.</i>  |
| <b>trap</b>                         | The action taken by the processor when it changes the instruction flow in response to the presence of an exception, a TCC instruction, or an interrupt. The action is a vectored transfer of control to supervisor software through a table, the address of which is specified by the privileged TBA register. <i>See also <b>exception</b>.</i>  |
| <b>TSB</b>                          | Translation storage buffer. A table of the address translations that is maintained by software in system memory and that serves as a cache of the address translations.   |
| <b>TSO</b>                          | Total store order.  |
| <b>TTE</b>                          | Translation table entry. Describes the virtual-to-physical translation and page attributes for a specific page in the Page Table. In some cases, the term is explicitly used for the entries in the TSB.  |
| <b>UE</b>                           | User process error.   |



|                                 |  |
|---------------------------------|--|
| <b>unassigned</b>               | A valued (for example, an ASI number) semantics which are not architecturally mandated and which may be determined independently by each implementation within any given guidelines.   |
| <b>undefined</b>                | An aspect of the architecture deliberately left unspecified. Software should have no expectation of, nor make any assumptions about, an undefined feature or behavior. Use of such a feature can deliver unexpected results, may or may not cause a trap, can vary among implementations, and can vary with time on a given implementation.<br><br>Notwithstanding any of the above, undefined aspects of the architecture shall not cause security holes (such as allowing user software to access privileged state), put the processor into supervisor mode or an unrecoverable state. |
| <b>unimplemented</b>            | An architectural feature that is not directly executed in hardware because it is optional or emulated in software.   |
| <b>unpredictable</b>            | <i>Synonym:</i> <b>undefined</b> .   |
| <b>unrestricted</b>             | Describes an ASI that can be used regardless of the processor mode, that is, regardless of the value of <code>PSTATE.PRIV</code> .   |
| <b>user application program</b> | <i>Synonym:</i> <b>application program</b> .   |
| <b>VA</b>                       | Virtual address. An address produced by a processor that maps all systemwide, program-visible memory. Virtual addresses usually are translated by a combination of hardware and software to physical addresses, which can be used to access physical memory.   |
| <b>victimize</b>                | [Error handling]   |
| <b>VIPT</b>                     | Virtually indexed, physically tagged.  |
| <b>virtual address</b>          | <i>See</i> <b>VA</b> .   |
| <b>VIS</b>                      | Visual Instruction Set. Performs partitioned integer arithmetic and other small integer operations.  |
| <b>VIVT</b>                     | Virtually indexed, virtually tagged (cache).   |
| <b>WAW</b>                      | Write-After-Write.   |
| <b>WDR</b>                      | Watchdog trap-level reset.   |
| <b>word</b>                     | An aligned quadlet. <b>Note:</b> The definition of this term is architecture dependent and may differ from that used in other processor architectures.   |
| <b>WRF</b>                      | Working Register File.   |
| <b>writeback</b>                | The process of writing a dirty cache line back to memory before it is refilled.  |
| <b>WRPR</b>                     | Write Privileged Register.   |
| <b>XIR</b>                      | Externally initiated reset.  |



**SECTION I**

---

---

---

**Processor Introduction**

---

---



# Processor Introduction

---

## 1.1 Overview

The UltraSPARC III Cu processor is a high-performance, highly integrated superscalar processor that implements the 64-bit SPARC V9 RISC architecture. It can sustain the execution of up to four instructions per cycle, even in the presence of conditional branches and cache misses, mainly because the units asynchronously feed instructions and data to the rest of the pipeline. Instructions that are predicted to be executed are issued in program order to multiple functional units, executed in parallel, and for added parallelism can be completed out-of-order. To further increase the number of instructions executed per cycle, instructions from two basic blocks can be issued in the same group.

The chip supports a 64-bit virtual address space and a 43-bit physical address space. The core instruction set has been extended to include graphics instructions that provide the most common operations related to two-dimensional image processing, two- and three-dimensional graphics and image compression algorithms, and parallel operations on pixel data with 8- and 16-bit components.

The processor is designed to offer very high clock speeds as well as wide superscalar issue to exploit instruction-level parallelism. The processor offers large Level-1 instruction and data caches, large flexible memory management units (MMUs), and support for large L2-cache. The processor was designed to work in systems ranging from single processor workstations through cache coherent servers with over a hundred processors. For building large systems, the processor has built-in support for both snooping-based cache coherency and directory-based cache coherency.

The architecture and implementation coupled with new compiler techniques make it possible to reduce each component while not degrading the other two.

---

## 1.2 CPU Features

The UltraSPARC III Cu processor is a richly featured processor. Features include:

- SPARC V9 Architecture with the VIS™ II Instruction Set, SPARC Binary code hardware compatible
- 4-way superscalar processor with nine execution units and six execution pipelines
- 14 stage, non-stalling pipeline
- Improved memory latency
- 64-bit data paths, 64-bit ALUs, 64-bit address arithmetic
- 64-bit virtual address and 43-bit physical address space
- Data prefetching mechanism
- L2-cache unit that supports a 2-way set associative cache
- Comprehensive error detection and recovery
- Data Memory Management Unit with 1040 Translation Lookaside Buffer (TLB) entries that can support up to 4 MB pages

---

## 1.3 Cache Features

The UltraSPARC III Cu processor cache features include:

- 32 KB, 4-way set associative primary instruction cache memory with parity
- 64 KB, 4-way set associative primary data cache memory with parity
- 2 KB, 4-way set associative Prefetch cache for software prefetch
- 2 KB, 4-way set associative Write cache reduces store bandwidth to Level 2 cache
- 8 MB, 2-way set associative external unified L2-cache with ECC protection (single bit correction, double bit detection)

---

## 1.4 Technology

The UltraSPARC III Cu processor technology features include:

- 0.18  $\mu$ , 7-layer Cu metal, CMOS process
- 1.6 V core and 1.5 V I/O power supplies

- 232 mm<sup>2</sup> die size
- 1368 pin ceramic LGA package
- 900 MHz and higher frequency
- 80 W power dissipation at 900 MHz

---

## 1.5 UltraSPARC III Cu Differences

The UltraSPARC III Cu processor differs from previous UltraSPARC processors in several key areas, including:

- Bootbus limitations
- Instruction set extensions
- Instruction differences
- Memory subsystem
- Interrupts
- Address space size
- Error correction
- Registers
- Non-cacheable store compression
- RAS Architecture

This section describes the UltraSPARC III Cu chip differences and includes a summary table of those differences. The section concludes with a discussion of the UltraSPARC III Cu processor performance enhancements and RAS architecture.

### 1.5.1 Bootbus Limitations

All bootbus addresses must be mapped as side-effect pages with the TTE.E bit set. In addition, programmers must not issue the following memory operations to any bootbus address:

- Prefetch instructions
- Block load and block store instructions
- Any memory operations with ASI\_PHYS\_USE\_EC or ASI\_PHYS\_USE\_EC\_LITTLE
- Partial store instructions

## 1.5.2 Instruction Set Extensions

The UltraSPARC III Cu processor has added Sun proprietary extensions to the SPARC V9 Instruction Set Architecture (ISA), in addition to those implemented in UltraSPARC I. The extensions are in the areas of VIS extensions, prefetch enhancement, and interval arithmetic support.

### 1.5.2.1 Visual Instruction Set (VIS) Extensions

Three new VIS instructions were added:

- **Byte Mask** — Sets the Graphics Status Register (GSR) for a following byte shuffle operation. One byte mask can be issued per instruction group as the last instruction of the group.

Byte Mask is a break-after instruction.

- **Byte Shuffle** — Allows any set of 8 bytes to be extracted from a pair of double-precision, floating-point registers and written to a destination double-precision, floating-point register. The 32-bit byte mask field of the GSR specifies the pattern of source bytes for the byte shuffle instruction.

- **Edge(ncc)** — Two variants: the original instruction sets the integer condition codes, and the new instruction does not set condition codes. Differences between the variants are as follows:

**Edge**

Sets integer condition codes  
Single instruction group

**Edge(ncc)**

Does not set integer condition codes  
Groupable

Because of implementation restrictions in the pipeline, all instructions that set condition codes and execute in the MS pipeline stage must be in a single instruction group.

### 1.5.2.2 Prefetch Enhancement

The processor supports an instruction to invalidate a prefetched line. It invalidates a prefetch cache line after prefetched data has been loaded into registers and on error conditions.

### 1.5.2.3 Interval Arithmetic Support

One new instruction was added to improve the efficiency of interval arithmetic computations. The Set Interval Arithmetic Mode (SIAM) instruction enables the rounding mode bits in the Floating-Point Status Register (FSR) to be overridden without the overhead of modifying the RD field of the FSR. Updates directly to FSR are expensive because they flush the pipeline.



## 1.5.3 Instruction Differences

Several instructions have changed relative to the previous UltraSPARC processors.

- **SHUTDOWN** — Low power mode compliance is achieved through a different mechanism than that used by the UltraSPARC I processor. For compatibility, the SHUTDOWN instruction in the UltraSPARC III Cu processor executes as a NOP.
- **FLUSH** — Since the processor maintains consistency between the instruction cache and all store and atomic instructions, the FLUSH instruction is used only to clear the pipeline. Unlike the case with the UltraSPARC I processor, the FLUSH address is ignored. It is not used for instruction cache flushing and is not propagated to the system.
- **Floating-point conversion instructions** — Because of implementation restrictions, the following integer to floating-point conversion instructions generate an *unfinished\_FPop* exception for certain ranges of integer operands, as shown in TABLE 1-1.

**TABLE 1-1** Integer/Floating-Point *unfinished\_FPop* Exception Conditions

| Instruction | Unfinished Trap Ranges   |
|-------------|--|
| FsTOi       | result < $-2^{31}$ , result $\geq 2^{31}$ , Inf, NaN   |
| FsTOx       | result  $\geq 2^{52}$ , Inf, NaN   |
| FdTOi       | result < $-2^{31}$ , result $\geq 2^{31}$ , Inf, NaN   |
| FdTOx       | result  $\geq 2^{52}$ , Inf, NaN   |
| FdTOS       | result  $\geq 2^{52}$ ,  result  < $2^{-31}$ , operand < $-2^{22}$ , operand $\geq 2^{22}$ , NaN |
| FiTOS       | operand < $-2^{22}$ , operand $\geq 2^{22}$  |
| FxTOS       | operand < $-2^{22}$ , operand $\geq 2^{22}$  |
| FxTOd       | operand < $-2^{51}$ , operand $\geq 2^{51}$  |

When the above instructions take an *unfinished\_FPop* trap, system software must properly emulate the instruction and resume execution.

- **NaN handling** — Because of implementation restrictions, the processor generates an *unfinished\_FPop* exception for operations that use the floating-point adder when one or more of the operands is NaN. Previous UltraSPARC processors would propagate the NaN in hardware.
- **Floating-point subnormal** — Because of implementation restrictions, the processor generates an *unfinished\_FPop* exception in nonstandard mode for floating-point addition and floating-point subtraction operations when the result is a subnormal value. Previous UltraSPARC processors handled these in hardware. When an *unfinished\_FPop* trap is generated, it is expected that system software will properly emulate the instruction and resume execution.
- **Ticc reserved field checking** — The processor checks the reserved field of the Ticc instruction for zero and generates an *illegal\_instruction* trap if the field is nonzero. Neither UltraSPARC I nor UltraSPARC II processors checked the Ticc reserved field for zero.

## 1.5.4 Memory Subsystem

The memory subsystem design is new. Differences include changes in the caches, cache flushing, and TLBs.

### 1.5.4.1 Caches

The UltraSPARC III Cu memory system comprises five caches: four on-chip and one external to the chip.

- **Data cache (D-cache)** — A 64 KB, 4-way associative, virtually indexed, physically tagged (VIPT) cache. The D-cache is write-through, no write-allocate, not included in the L2-cache. The line size is 32 bytes with no sub-blocking. The D-cache needs to be flushed only if an alias is created with virtual address (VA) bit 13. VA<13> is the only virtual bit used to index the D-cache.
- **Instruction cache (I-cache)** — A 32 KB, 4-way associative, VIPT cache. The I-cache is not included in the L2-cache.

The line size is 32 bytes - no sub-blocking. The I-cache is kept consistent with the store stream of the processor as well as with external stores from other processors.

You never need to flush the I-cache, not even for address aliases.

- **Prefetch cache (P-cache)** — A 2 KB, 4-way associative cache. It is virtually indexed, virtually tagged (VIVT) cache for lookup and install to the cache. It is physically indexed and physically tagged for snoop and invalidate operations. The P-cache is not included in the L2-cache. The line size is 64 bytes with 32-byte sub-blocks.

The P-cache is globally invalidated if any of the following conditions occur:

- If the context registers are written
- If there is a demap operation in the DMU
- When the DMU is turned on or off

Individual lines are invalidated on any of the following conditions:

- A store hits
- An external snoop hit
- Use of software prefetch invalidate function (PREFETCH with `fcn = 16`)

The P-cache is used for software prefetch instructions as well as for autonomous hardware prefetches from the L2-cache.

Software never needs to flush the P-cache, not even for address aliases.

- **Write cache (W-cache)** — A 2 KB, 4-way associative, PIPT cache. The line size is 64 bytes with 32-byte sub-blocks. The W-cache reduces bandwidth to the L2-cache by coalescing and bursting stores to the L2-cache.

The W-cache is included in the L2-cache; all lines in the W-cache have a corresponding line allocated in the L2-cache. The data state of the W-cache line always supersedes the state of the data in the corresponding L2-cache line.

It is necessary to flush the W-cache for stable storage. Flushing the L2-cache implicitly forces the flush of the W-cache.

- **L2-cache** — A 1 MB to 8 MB, direct mapped or 2-way set associative, PIPT cache. The L2-cache is write-allocate, write-back.

It is necessary to flush the L2-cache for stable storage.

## 1.5.4.2 Cache Flushing

The following are flushing requirements for specific caches:

- **Data cache** — The UltraSPARC III Cu D-cache differs in size and organization from the UltraSPARC I D-cache and so requires changes to the algorithms used to flush the cache.

The virtually indexed caches need to be flushed when a virtual address alias is created. Caches that contain modified data need to be flushed for stable storage.

The UltraSPARC III Cu D-cache is the only cache that needs to be flushed when a virtual address alias is created. Like the UltraSPARC I D-cache, the UltraSPARC III Cu D-cache uses one virtual address bit for indexing the cache and thus creates an alias boundary of 16 KB for the D-cache.

- **Instruction cache** — The processor maintains consistency of the on-chip I-cache with the stores from all processors so that a `FLUSH` instruction is needed only to ensure the pipeline is consistent. This means a single flush is sufficient at the end of a sequence of stores that updates the instruction stream to ensure correct operation.

Unlike the case with the UltraSPARC I processor, the `FLUSH` instruction does not propagate externally since all I-caches in an UltraSPARC III Cu multiprocessor system are maintained consistent. Since the I-cache is a PIPT cache, it does not have to be flushed for virtual address aliases. The I-cache never contains modified data; therefore, it does not need to be flushed for stable storage.

- **Prefetch cache** — The P-cache is physically indexed and tagged. It cannot contain modified data, so it never needs to be flushed.
- **L2-caches and write caches** — Since the L2-cache and W-cache can contain modified data, they must be flushed for stable storage. The W-cache is included in the L2-cache, so it is sufficient to flush a block from the L2-cache; if there is a corresponding block in the W-cache, it will also be flushed. The recommended procedure to flush modified data from the L2-cache back to memory is as follows:
  - Load the block (64 bytes) into the floating-point registers by using FP loads or Block Load.
  - Write the floating-point registers to memory with a Block Store Commit.
  - Issue `MEMBAR #Sync` to ensure completion.

The Block Store with Commit instruction will invalidate the block from both the L2-cache and the W-cache. Both of these caches are physically indexed, so they do not need to be flushed for address aliases.

### 1.5.4.3 Translation Lookaside Buffers (TLBs)

The implementation of instruction and data TLBs for the UltraSPARC III Cu processor is described in this section.

The following are two instruction TLBs that are accessed in parallel:

- A 16-entry, fully associative TLB to hold entries for 8 KB, 64 KB, 512 KB, and 4 MB page sizes. This TLB contains locked and unlocked pages of any size.
- A 128 entry, 2-way associative TLB used exclusively for 8 KB page entries. This TLB contains only unlocked pages.

The following are three data TLBs that are accessed in parallel.

- A 16-entry, fully associative TLB to hold entries for 8 KB, 64 KB, 512 KB, and 4 MB page sizes. This TLB contains locked and unlocked pages of any size.
- Two 512 entry, 2-way set associative that can each be programmed to support lookup of any one page size at a given time. However, multiple page sizes can be resident. This TLB contains only unlocked pages.

Other TLB differences are described below:

- **TLB flushing** — Both the instruction and data TLBs now have a demap-all operation that removes all unlocked Translation Table Entries (TTEs).
- **TTE format** — The UltraSPARC III Cu processor now has the additional elements in the TTE format:
  - Physical Address field: Expanded from 28 bits (PA<40:13>, TTE<40:13>) to 30 bits (PA<42:13>, TTE<42:13>).
- **Synchronous Fault Status Registers (SFSR) extensions** — A new fault type was added to the FT field of the SFSR to indicate an I/D-TLB miss, and one status bit was added to the D-TLB SFSR:
  - NF: Set to signify that the faulting operation was a speculative load instruction.
- **Instruction/Data Translation Storage Buffer (i/dTSB) Register** — Three new register extensions of the i/dTSB register were added to the UltraSPARC III Cu processor. These registers allow a different TSB virtual address base to be used for each of the three virtual address spaces (primary, secondary, nucleus) in the D-TLB and two virtual address spaces (primary, nucleus) in the I-TLB. On an I/D-TLB miss, the processor selects which TSB Extension Register to use to form the TSB base address, based on the virtual space accessed by the faulting instruction.
- **TLB Data Access Register** — The access address for the TLB Data Access Register has been expanded to enable access to three TLBs, each with up to 512 entries.
- **TLB Diagnostic Register** — A new register replaces the function of the diagnostic bits in the TTE.

## 1.5.5 Interrupts

The UltraSPARC III Cu processor extends the interrupt architecture previously implemented in the UltraSPARC I processor in these areas:

- **Module ID fields** — Extended from 5 bits to 10 bits.
- **Interrupt Transmit BUSY/NACK bits** — Extended from 1 pair to 32 pairs, enabling pipelining of outgoing interrupts.
- **Data Dispatch and Receive Registers** — Expanded from 3 to 8, enabling up to 64 bytes to be transmitted in an interrupt.
- **System Tick Interrupt bit** — Added to the soft interrupt register.

## 1.5.6 Address Space Size

The UltraSPARC III Cu processor extends both the virtual and physical address space previously implemented. It implements the full 64-bit virtual address range defined in the SPARC V9 architecture. There are no VA holes, compared to the UltraSPARC I. The physical address range has also been extended from 41 bits to 43 bits.

Address space with  $PA_{<42>} = 1$  is considered as the non-cacheable address space. Physical address  $40000000000_{16}$  to  $7FFFFFFF_{16}$  is in the non-cacheable area.

## 1.5.7 Registers

Differences in registers include enhancements to ASI registers and ASR registers.

### Address Space Identifier (ASI) Registers

Changes to the ASI registers include those to the following registers:

- **SRAM diagnostic registers** — Several new diagnostic ASI registers were added for the following on-chip SRAMs:
  - Prefetch cache
  - Write cache
  - Branch predict array
  - I/D-TLB CAM

Changes were made to fields of existing UltraSPARC II diagnostic ASI registers for the following on-chip SRAMs:

- Data cache
- L2-cache
- Instruction cache

The following ASI registers were removed:

- UDB Error Register
- UDB Control Register
- **Asynchronous Fault Status Register (AFSR)** — Several changes were made to add new fault types (L2-cache ECC errors) and remove old fault types (SDB errors). AFSR accumulates errors.
- **Asynchronous Fault Address Register (AFAR)** — The AFAR was extended to handle a 43-bit physical address. It is now updated on several errors that previously did not capture the address. AFAR overwrites for higher priority errors if a more severe error occurs.
- **Secondary Asynchronous Fault Status Register (AFSR)** — Secondary AFSR captures the first error.
- **Asynchronous Fault Address Register (AFAR)** — Secondary AFAR captures the address associated with the first error and it locks until it is explicitly cleared by software.
- **Software Interrupt Register (SOFTINT)** — The SOFTINT register has an additional bit added to signal SYSTEM TICK COMPARE interrupts.
- **System Interface Registers** — The UPA interface ASI has been reused for two new Fireplane Interconnect registers: a configuration register and an address register.

## Ancillary State Registers (ASRs)

Changes to the ASRs include changes to the following registers:

- **System Tick and Compare** — Two new ASRs were added to support a system clock: ASR 18<sub>16</sub>, a System Tick Register (analogous to the per-processor tick register ASR 4<sub>16</sub>), and ASR 19<sub>16</sub>, a System Tick Compare Register (analogous to the per-processor tick compare register 17<sub>16</sub>).
- **Graphics Status Register (GSR)** — New fields were added to the GSR:
  - 32-bit MASK field used by the BSHUFFLE instruction
  - 1-bit IM field to enable interval arithmetic round mode
  - 2-bit IRD field to specify round mode for interval arithmetic
- **Performance Control Register (PCR)** — The PCR has been extended to enable a larger number of performance events to be measured.
- **Dispatch Control Register (DCR)** — Many control fields were added to the DCR to aid in debugging first silicon.

## 1.5.8 Non-Cacheable Store Compression

Like previous implementations, the UltraSPARC III Cu processor uses a 16-byte buffer to merge adjacent non-cacheable stores into a single external data transaction. This merging greatly increases store bandwidth to the graphics frame buffer. A change in the algorithm for determining when to break merging improves store bandwidth to graphics devices.

## 1.5.9 Error Correction

Error correction differs from the UltraSPARC I processor and the UltraSPARC II processor handling, as follows:

- **L2-cache** — The processor uses ECC protection on the L2-cache instead of parity protection. It requires software correction and recovery for single bit L2-cache ECC read errors, which are signaled as a precise error.
- **System interface** — A new ECC code has been defined for ECC protection across 132 data bits (nine ECC bits) and three MTag bits (four ECC bits) on the system bus and on the data switch. The syndromes for these codes differ from the syndromes used previously. The processor requires software correction and recovery for single-bit system ECC errors, which are signaled as disrupting errors.

## 1.5.10 SRAM Protection and RAS Features

TABLE 1-2 lists all UltraSPARC III Cu on-chip SRAM protection and other RAS features.

**TABLE 1-2** RAS Features

| No. | Feature                                      | Feature Description  |
|-----|--|--|
| 1.  | D-cache Data array parity protection         | 2 parity bits per 64-bit data  |
| 2.  | D-cache Physical Tag array parity protection | 1-bit parity per tag entry   |
| 3.  | D-cache Snoop Tag array parity protection    | 1-bit parity per tag entry   |
| 4.  | I-cache Data array parity protection         | 1-bit parity per partially decoded instruction   |
| 5.  | I-cache Physical Tag array parity protection | 1-bit parity per tag entry   |
| 6.  | I-cache Snoop Tag array parity protection    | 1-bit parity per tag entry   |
| 7.  | L2-cache data array ECC protection           | ECC Protection with 1-bit correction and 2-bit detection   |
| 8.  | L2-cache Tag array ECC protection            | ECC Protection with 1-bit correction and 2-bit detection   |
| 9.  | Dual AFSR/AFAR                               | Secondary AFSR/AFAR captures on the first error event while primary AFSR/AFAR accumulates in case of multiple events |





## System Introduction

---

This chapter discusses how the UltraSPARC III Cu processor is used in systems.

---

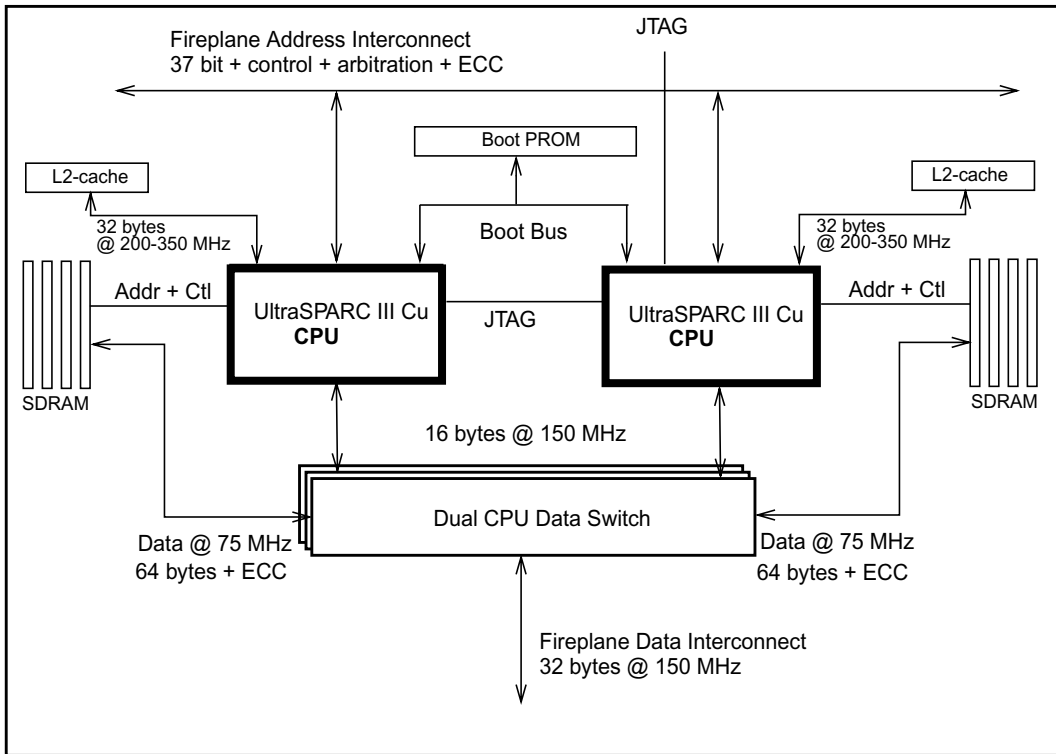
### 2.1 System Configurations

The UltraSPARC III Cu processor can be used in a variety of configurations from two-processor systems to very large, high-performance Symmetric Multiprocessor (SMP) Systems.

#### 2.1.1 Two-Processor Configuration with UltraSPARC III Cu

FIGURE 2-1 is an example of a two-processor UltraSPARC III Cu configuration. This configuration is the basic building block that can be used to build SMP systems.

In this configuration, both UltraSPARC III Cu processors are connected to a Dual CPU Data Switch (DCDS) via a 16-byte interface at 150 MHz. Each processor has an address and control interface to the SDRAMs. However, the data from the SDRAM is connected directly to the DCDS via a 64-byte interface running at 75 MHz. Note that the ranges for the SDRAM on different UltraSPARC III Cu CPUs should not overlap. Each UltraSPARC III Cu CPU connects to the Fireplane Interconnect Address bus. The UltraSPARC III Cu CPU also connects with an L2-cache via a 32-byte interface and with the Boot PROM via the Boot Bus interface. Both UltraSPARC III Cu CPUs provide a JTAG interface.



**FIGURE 2-1** Two-Processor Configuration with the UltraSPARC III Cu

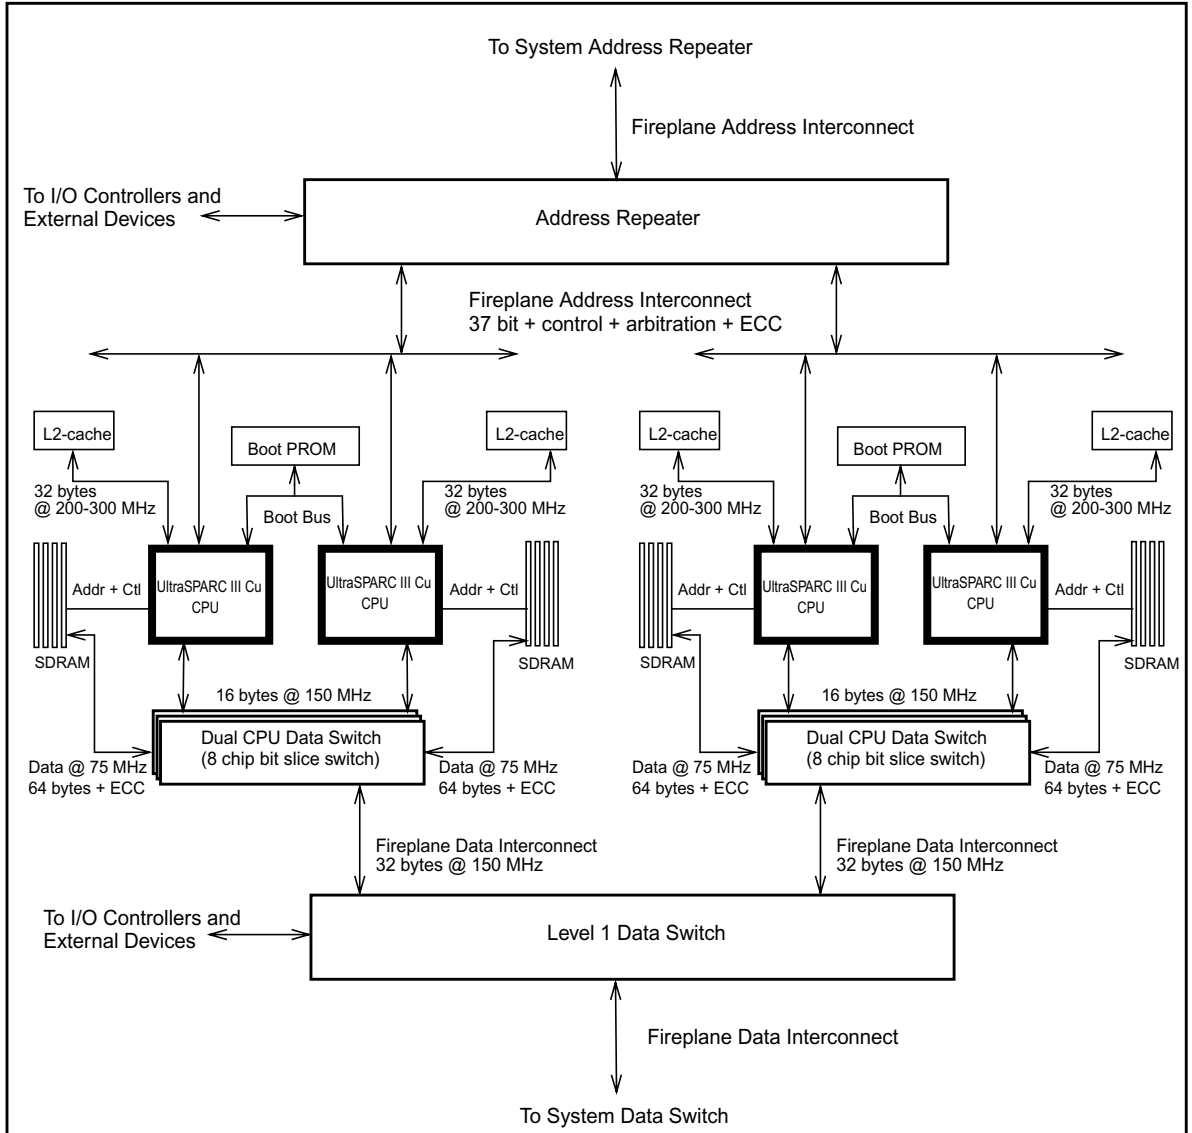
---

**Note** – Some processors or memories may not be present in the actual system.

---

## 2.1.2 Four-Processor Configuration with UltraSPARC III Cu

A four-processor configuration, shown in , is built from a two-processor configuration. This example also demonstrates how larger systems can be created by using an address repeater and a Level 1 data switch. I/O controllers, bridge chips, or external devices are connected to the system via the address repeater and the data switch. This configuration can be repeated to create a high performance SMP system.



Four-Processor Configuration with the UltraSPARC III Cu

## 2.1.3 Multiprocessor System with the UltraSPARC III Cu

A multiprocessor system, as shown in FIGURE 2-2, can be built using the two-processor configuration shown in FIGURE 2-1. A system address and data repeater are used to build such systems. Various devices can be connected to the system via PCI interfaces. Level 1 and Level 2 address repeaters and data switches are used to build systems that can accommodate up to six processor/memory boards and up to four I/O subsystems.

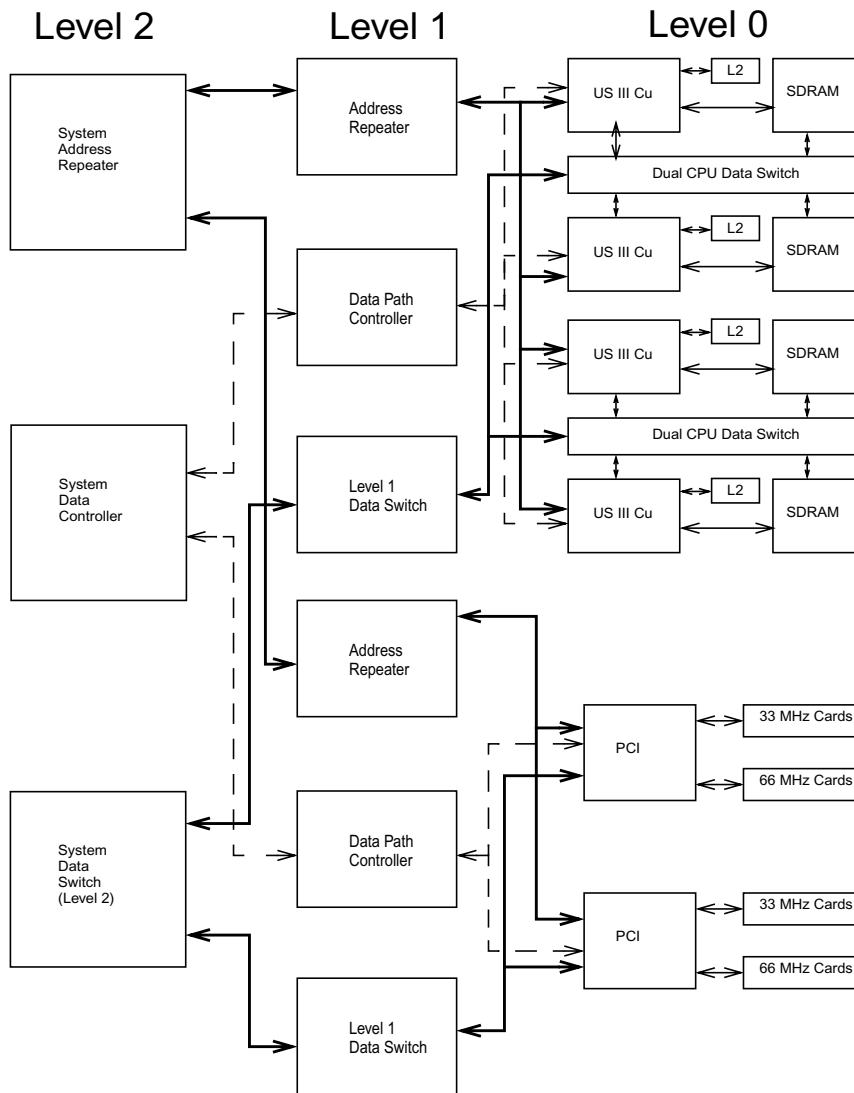
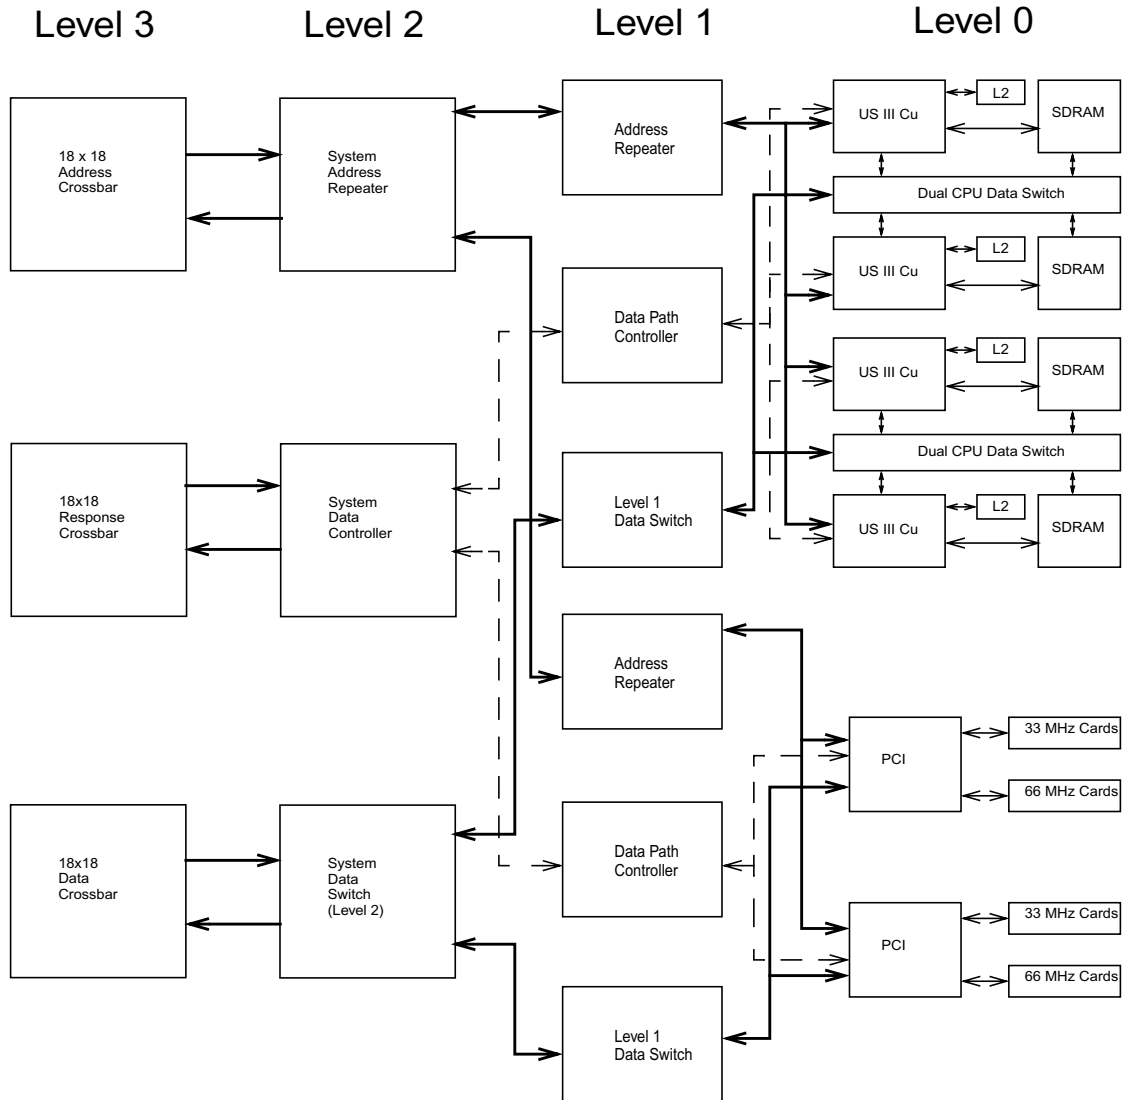


FIGURE 2-2 Multiprocessor System with the UltraSPARC III Cu

## 2.1.4 Very Large Multiprocessor System with the UltraSPARC III Cu

A very large multiprocessor system, as shown in FIGURE 2-3, can be built by interconnecting multiple configurations, shown in FIGURE 2-2, using Level 3, 18-port crossbar switches.



**FIGURE 2-3** Very Large Multiprocessor System with the UltraSPARC III Cu

---

## 2.2 Cache Coherence

UltraSPARC III Cu-based systems support a “snooping” based cache coherence protocol and a directory based cache coherence protocol (also known as Scalable Shared Memory (SSM)).

For small to medium systems, the UltraSPARC III Cu processor uses a MOESI snooping protocol. In this protocol, when a processor wants a line, it broadcasts the request to all other processors, which check their caches to see if they have the line.

For larger systems, the UltraSPARC III Cu processor has built-in support for a directory based coherence protocol (SSM).

In practice, there are small clusters of processors that are connected together with a snooping based coherence protocol. A directory based cache coherence protocol is used between clusters.

---

## 2.3 System Interfaces

There are multiple interfaces on the UltraSPARC III Cu processor. This section summarizes the various interfaces.

### 2.3.1 Fireplane Interconnect

The Fireplane Interconnect has two parts. Both typically run in the range of 150 MHz.

- A hierarchical bus for address and control.
- A point-to-point data interconnect.

The Fireplane Interconnect Address Bus has 37 bits of address plus control, arbitration and ECC bits. All UltraSPARC III Cu processors connect to this bus directly. To build larger multiprocessor systems, a repeater chip is used to create a hierarchical bus.

The Fireplane Interconnect Data Bus has 32 bytes of data along with ECC and routing information. All UltraSPARC III Cu processors connect to the Fireplane Interconnect Data Bus through a DCDS switch. To build larger multiprocessor systems, a point-to-point data network built with Level 1 data switches can be used.

## 2.3.2 SDRAM Interface

The UltraSPARC III Cu processor provides address and control bits to the SDRAM. The data from the SDRAM is directly driven to the DCDS via a 64-byte interface running at one-half of the Fireplane frequency.

## 2.3.3 DCDS Interface

The UltraSPARC III Cu processor connects to the Dual CPU Data Switch (DCDS) via a 16-byte interface running at 150 MHz and capable of delivering a peak bandwidth of 2.4 GB/s. DCDS is an eight chip, bit slice switch that converts it into a 32-byte Fireplane Interconnect Data Port running at 150 MHz and capable of delivering a peak bandwidth of 4.8 GB/s.

## 2.3.4 L2-Cache Interface

The UltraSPARC III Cu processor connects to an up to 8 MB, 2-way set associative external unified Level 2 (L2) cache with ECC protection via a 32-byte interface, running at 200 MHz or higher.

## 2.3.5 Boot Bus Interface

The UltraSPARC III Cu processor has a Boot Bus interface that connects to a Boot PROM, other booting mechanisms, and diagnostic and recovery mechanisms.

## 2.3.6 JTAG Interface

The UltraSPARC III Cu processor provides a standard 1149.1 compliant JTAG interface. This interface can be used to scan out the internal state of the processor for fault diagnosis. The full scan is a destructive operation that requires the CPU to go through a Power-on reset (POR) before being used again. In addition, the UltraSPARC III Cu processor also provides a shadow JTAG interface that allows a subset of the state to be scanned while the processor is running.





## SECTION II

---

---

# Architecture and Functions

---

---



## CPU Architecture Basics

---

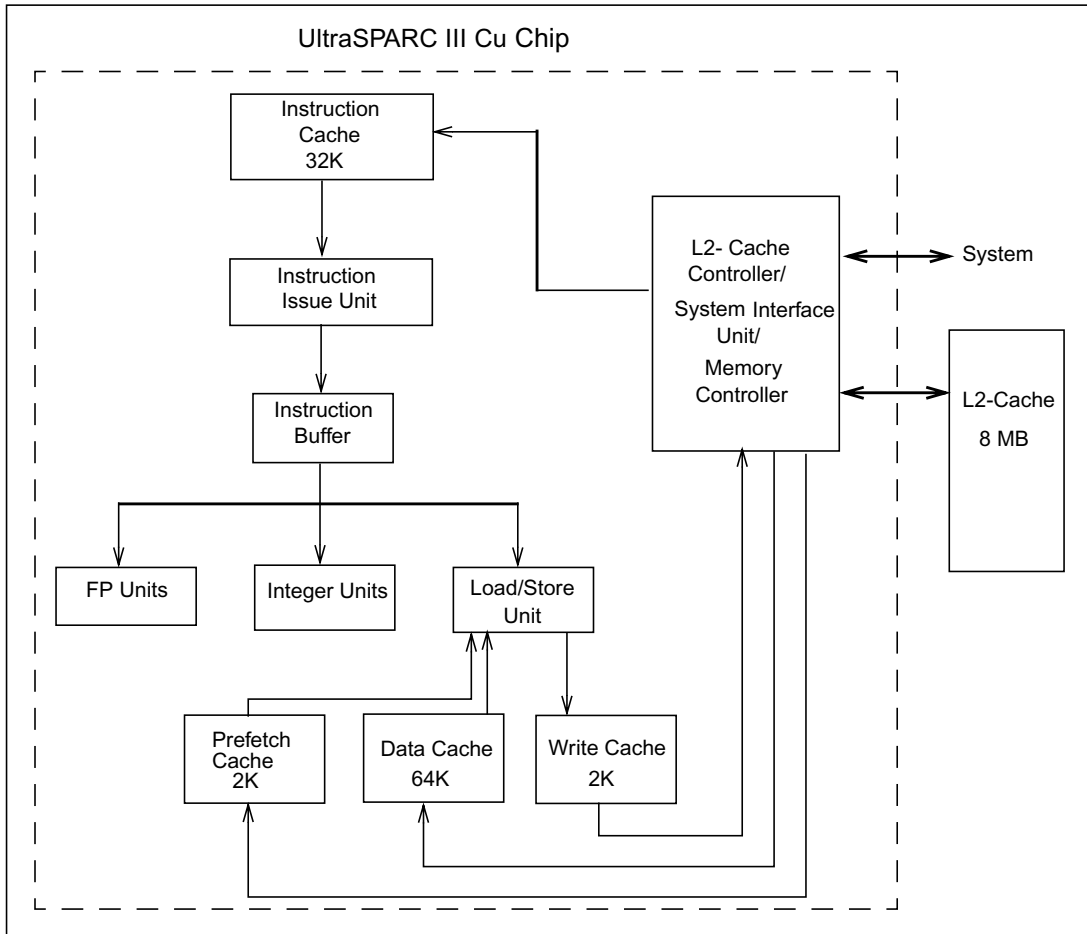
The UltraSPARC III Cu processor is a high-performance, highly integrated, superscalar processor. The UltraSPARC III Cu fully implements the 64-bit SPARC V9 architecture, supporting a 64-bit virtual address space and a 43-bit physical address space. The core instruction set is extended to include new SIMD operations. The processor was designed to offer very high clock speeds as well as wide superscalar issue to exploit instruction-level parallelism. The processor offers large Level-1 instruction and data caches, large flexible memory management units (MMUs), and support for a large L2-cache. The processor was designed to work in systems ranging from single processor workstations through cache-coherent servers with more than a thousand processors. For building a wide range of system configurations, the processors has built-in support for both snooping-based cache coherency and directory-based cache coherency.

The UltraSPARC III Cu processor also offers a number of performance enhancements over previous UltraSPARC processors. The processor incorporates a number of data prefetching mechanisms to exploit memory-level parallelism. The processor offers an enhanced data memory management unit (D-MMU) that has 1040 TLB entries and more support for flexibly using large pages, up to 4 MB pages, to more effectively map gigabytes of data. The processor supports a 2-way set associative L2-cache instead of a direct-mapped cache.

---

### 3.1 Component Overview

The processor consists of a high-performance, instruction fetch engine, called the instruction issue unit, that is decoupled from the rest of the pipeline by an instruction buffer. Instructions are steered to either floating-point execution units, integer execution units, or a load/store unit for execution. Integrated on the processor are controls for the L2-cache, interface to the Fireplane bus, and a memory controller. A simple block diagram of the UltraSPARC III Cu processor is shown in FIGURE 3-1.



**FIGURE 3-1** UltraSPARC III Cu Architecture

### 3.1.1 Instruction Fetch and Buffering

The instruction issue unit in the UltraSPARC III Cu processor is responsible for fetching, caching, and supplying groups of instructions to its execution pipelines. Instructions are fetched and decoded in groups of up to four instructions. Instruction fetching can be done in every cycle if the request hits instruction cache and other conditions are met. If the request misses instruction cache, a fill request is sent to the lower memory hierarchy and the requested 32-byte line is installed in instruction cache. If the requested line is the first half of a 64-byte boundary, the second half of the 64-byte boundary is prefetched and the line is filled in the Instruction Prefetch Buffer (IPB) for a potential reference in the future.

The UltraSPARC III Cu instruction cache is a 32 KB size, 32-byte line size (eight instructions), physically tagged, 4-way set associative cache. Its throughput is one cycle with two-cycle latency. In addition to data and tag array, it also has a microtag, predecode, Load Prediction Bit (LPB), and snoop tag array. The microtag uses eight bits of virtual address to enable way-select to be performed before the physical address translation is completed. The predecode bits include information about which pipeline each instruction will issue to and other information to optimize execution. The LPB is used to dynamically learn those load instructions that frequently see a read-after-write (RAW) hazard with preceding stores. The snoop tag is a copy of the tags dedicated for snoops caused by either stores from the same or different processors. The instruction cache in the UltraSPARC III Cu processor is kept completely coherent so the cache never needs to be flushed.

The instruction fetch engine is also dependent upon control transfer instructions, such as branches and jumps. The UltraSPARC III Cu processor uses a 16K entry branch predictor to predict fetch direction of conditional branches. The target must be determined for branches that are either predicted or known to redirect instruction fetches. For PC relative branches, the target of the branch is computed; this adds a one-cycle branch taken penalty, but avoids target misprediction. For predicting the target of return instructions an 8-entry Return Address Stack (RAS) is used. For other indirect branches (branches whose targets are determined by a register value), the software can provide a branch target prediction with a jump target preparation instruction.

Between the instruction fetch pipeline and the execution pipeline is an instruction buffer that can hold up to 16 instructions. The instruction buffer decouples the fetch and execute pipelines and buffers burstiness in each pipeline from each other. The buffer can effectively hide low latency issues like the taken branch penalty and even hides some of the penalty of instruction cache misses.

## 3.1.2 Execution Pipelines

The UltraSPARC III Cu processor has six execution pipelines and can issue up to four instructions per cycle. The six execution pipelines consist of the following:

- Two integer arithmetic and logic (ALU) pipelines
- Branch pipeline
- Load/store pipeline which also handles special instructions
- Floating-point multiply pipeline which also handles SIMD instructions
- Floating-point addition pipeline which also handles SIMD instructions

The integer ALU pipelines can issue integer addition, subtraction, logic operations, and shifts. These pipelines have single-cycle latency and throughput. The branch pipeline handles all branch instructions and can resolve one branch each cycle. The load/store pipeline can handle one load or store instruction each cycle and is discussed in more detail in

Section 3.1.3. Integer multiplication and division is performed by the load/store pipeline. Integer multiplication has a latency of 6 to 9 cycles depending on the size of the operands. Division is also iterative and requires 40 to 70 cycles.

The floating-point pipelines are each four-cycle latency pipelines but are fully pipelined (one instruction per cycle per pipeline). These pipelines handle single and double precision floating-point operations and a set of data parallel operations that operate on 8- or 16-bit fields. Floating-point division and square root operations use the floating-point multiplication pipeline and are iterative computations. Floating-point division requires 17 or 20 cycles for single and double precision, respectively. Floating-point square root requires 23 or 29 cycles for single and double precision, respectively.

### 3.1.3 Load/Store Unit

As stated earlier, a load or store instruction can be issued each cycle to the load/store pipeline. The load/store unit consists of the load/store pipeline, a store queue, a data cache, and a write cache.

Loads have either a two- or three-cycle latency. Integer loads that are for unsigned words or double words have a two-cycle latency. All other load instructions have a three-cycle latency. Data can be forwarded from earlier stores still in the store queue to subsequent loads if a RAW hazard is detected. Data forwarding requires a three-cycle latency. For those instructions that can have a two-cycle latency, there is a prediction bit in the instruction cache used to identify those loads that often require store forwarding, which will be issued as three-cycle loads. If a two-cycle load is not correctly predicted to have a RAW hazard, the load must be reissued.

There is an 8-entry store queue to buffer stores. Stores reside in the store queue from the time they are issued until they complete an update to the write cache. The store queue can effectively isolate the processor from the latency of completing stores. If the store queue fills up, the processor will block on a subsequent store. The store queue can coalesce stores to the same cache line. The store queue allows non-catchable stores (for example, stores to a graphics frame buffer) to be coalesced together such that the required bandwidth to the device is greatly reduced.

The data cache is a 64 KB, 4-way associative, two-cycle latency, one-cycle throughput, virtually indexed, physically tagged (VIPT) cache. The data cache, like the instruction cache, uses 8-bit microtags to do way-selection based on virtual addresses. The data cache is write-through, no write-allocate, and not included in the L2-cache. The line size is 32 bytes with no sub-blocking. The data cache needs to be flushed only if an alias is created using virtual address bit 13. VA[13] is the only virtual bit used to index the data cache.

The write cache is a write-back cache used to reduce the amount of store bandwidth required to the L2-cache. It exploits both temporal and spatial locality in the store stream. The small (2 KB) structure achieves a store bandwidth equivalent to a 64 KB write-back data cache

while maintaining SPARC V9 TSO compatibility. The write cache is kept fully coherent with both the processor pipeline and the system memory state. The write cache is 4-way set associative and has 64-byte lines. The write cache maintains per byte dirty bits.

### 3.1.3.1 Data Prefetching Support

The UltraSPARC III Cu processor makes use of an advanced data prefetching mechanism. This mechanism is used to both overlap load misses to increase memory-level parallelism and to hide load-miss latency. This mechanism allows software to explicitly expose the memory-level parallelism and to schedule memory operations. This mechanism is extremely important because the UltraSPARC III Cu processor has blocking loads; when the processor reaches a load instruction that misses in the cache, the processor waits for the load to complete before executing any other instructions. The processor supports software prefetching where the compiler (or Java™ JIT) can schedule prefetching of data to exploit memory-level parallelism. Some versions of the processor will also support hardware prefetching, where the processor observes common data sequences and attempts to prefetch the data automatically.

There are a number of variations of software prefetches. Software prefetches can specify if the data should be brought into the processor either for reading or both reading and writing. Software can also specify if the data should be installed into the L2-cache, for data that will be reused frequently, or only brought into the prefetch cache.

One of the main mechanisms for implementing prefetches is a special prefetch cache. The prefetch cache is a small (2 KB) cache that is accessed in parallel with the data cache for floating-point loads. Floating-point load misses, hardware prefetches, and software prefetches bring data into the prefetch cache. The prefetch cache is 4-way set associative and has 64-byte lines, which are broken into two 32-byte sub-blocks with separate valid bits. The prefetch cache is write invalidate.

## 3.1.4 Memory Management Units

There are separate Memory Management Units (MMUs) for instruction and data address translation. The MMUs consist of a set of translation lookaside buffers (TLBs) that are tables of translations from virtual to physical addresses. As long as a virtual address can be translated using one of the entries in a TLB, the operation proceeds without interruption. If there is no translation available for a virtual address, the processor traps to software to update the TLBs with a valid translation.

For the instruction address stream translation, there are two TLBs accessed in parallel. The first TLB is a 16-entry, fully associative TLB. This TLB can translate page sizes of 8K, 64K, 512K, and 4M, and locked page always reside in this TLB. The second TLB is a 64 set, 2-way set associative (128 entries) TLB. This large TLB is used to translate a page size of 8K.

The D-MMU of the UltraSPARC III Cu processor is enhanced to provide more translation entries and more support for using large pages for translation. For the data reference address stream translation there are three TLBs accessed in parallel. The first TLB is a 16-entry fully-associative TLB. This TLB can translate page sizes of 8K, 64K, 512K, and 4M. The second TLB is a 256-set, 2-way set-associative (512 entries) TLB. This TLB can translate at 8K, 64K, 512K, and 4M page sizes, but at any one time it is configured to only handle one of the page sizes. The third TLB is identical to the second. This TLB, like the second, can handle one of four page sizes and can be configured to the same or a different page size than the second TLB.

The two large TLBs is very important for general use of large pages for translation. One of the TLBs can be set for large pages (such as 4 MB pages) while the other can be set to the default page size (usually 8 KB pages). With this configuration, the processor provides robust support for large pages.

### 3.1.5 L2-cache Unit (Level-2 Unified Cache)

The UltraSPARC III Cu processor can support an L2-cache of 1 MB, 4 MB or 8 MB. The L2-cache is 2-way set associative, PIPT cache. The line size of the L2-cache depends on the cache size (64 bytes for 1 MB, up to 512 bytes for 8 MB). Regardless of the line size, the cache uses 64-byte sub-blocks that are the unit of fill and the unit of coherency. The L2-cache is write-allocate, write-back. The tags for the L2-cache are on the CPU chip.

For data, the L2-cache uses standard SRAM parts running at either one-third, one-fourth, or one-fifth of the processor speed. The interface between the processor and the SRAM is 32 bytes wide. The L2-cache for the UltraSPARC III Cu processor is fully protected with error correcting code (ECC). Single bit errors in the L2-cache are corrected and double bit errors are detected. These result in UltraSPARC III Cu state-of-the-art reliability.

### 3.1.6 System Interface Unit

The system interface unit (SIU) is the UltraSPARC III Cu processor's port to the external world. All data is transferred between the UltraSPARC III Cu processor and local DRAM, main memory associated with another CPU, or the system bus passes through the SIU. The SIU is the engine for the cache coherency protocol for multiprocessor systems.

The SIU supports clock divisors of 4, 5, and 6 between the system clock and the internal CPU clock. When the system reset becomes inactive, both the internal CPU clock and the system interface clock are synchronized at the rising edge.

The system interface allows for a low-cost interconnect of up to six agents. An agent may be another UltraSPARC III Cu processor, an I/O controller, bus repeater or an SSM controller. The bandwidth of the external interface buses allows the system interface to be implemented using a snooping coherence protocol. A snooping interface allows each agent to maintain



coherency without the need for an external coherency controller. An UltraSPARC III Cu processor snoops coherent transaction requests issued on the system interface buses and follows a write-invalidate MOESI cache coherence policy. Snooping-based systems can be built with tens of processors.

SSM is a directory-based protocol used for creating very large multiprocessor systems. Smaller groups of processors using a snooping interface can utilize SSM controllers to create systems with over a hundred processors. The SIU of the UltraSPARC III Cu processor has built-in support for working with an SSM controller to facilitate the creation of large systems.

### 3.1.7 Memory Controller Unit

The UltraSPARC III Cu processor has an on-chip Memory Controller Unit (MCU). The UltraSPARC III Cu memory system supports a minimum of 128 MB and a maximum of 16 GB of main memory. The MCU supports 75 MHz SDRAM and interfaces to various densities of DRAM and single or multiple banks of memory. The MCU only sends out control signals to the DRAM. The UltraSPARC III Cu SIU is responsible for delivering data to the data switch for write operations and retrieving data from the data switch for read operations.

---

## 3.2 CPU Operating Modes

The UltraSPARC III Cu processor operates in various modes.

### 3.2.1 Privileged Mode

This mode is a “**supervisor**” mode. In this mode, the software is allowed to access both privileged and non-privileged registers and ASIs. There are certain features that can be accessed only in privileged mode. Non-privileged software is not allowed to access these features.

Privileged mode execution is typically used by the kernel and operating system.

## 3.2.2 Non-Privileged Mode

This mode is a “**non-supervisor**” mode. In this mode, the software is allowed to access only non-privileged registers and ASIs. If non-privileged software tries to access privileged registers or ASIs, exceptions are generated and handled by the operating system.

Non-privileged mode execution is typically used by the application programmers.

## 3.2.3 Reset and RED\_State

The UltraSPARC III Cu processor can be reset using various mechanisms. This section deals with the reset and RED\_state for the UltraSPARC III Cu processor.

### 3.2.3.1 RED\_state Characteristics

A processor enters RED\_state by one of the two ways.

- First, by trapping when already at the maximum trap level.
- Second, by setting the PSTATE.RED.

When the processor enters the RED\_state, it will clear the DCU Control Register, including enable bits for I-cache, D-cache, I-MMU, D-MMU, and virtual and physical watchpoints.

---

**Note** – Exiting RED\_state by writing zero to PSTATE.RED in the delay slot of a JMPL is not recommended. A non-cacheable instruction prefetch can be made to the JMPL target, which may be in a cacheable memory area. This condition could result in a bus error on some systems and cause an *instruction\_access\_error* trap. You can mask the trap by setting the NCEEN bit in the ESTATE\_ERR\_EN register to zero, but this approach will mask all noncorrectable error checking. Exiting RED\_state with DONE or RETRY avoids the problem.

---

### 3.2.3.2 Resets

Reset priorities from highest to lowest are power-on resets (POR, hard or soft), externally initiated reset (XIR), watchdog reset (WDR), and software-initiated reset (SIR).

## *Power-on Reset (Hard Reset)*

A Power-on Reset (POR) occurs when the POK pin is activated and stays asserted until the processor is within its specified operating range. When the POK pin is active, all other resets and traps are ignored. POR has a trap type of 1 at physical address offset 20<sub>16</sub>. Any pending external transactions are canceled.

After POR, software must initialize values of certain registers and state that is unknown after POR. The following bits must be initialized before the caches are enabled:

- In the I-cache, valid bits must be cleared and microtag bits must be set so that each way within a set has a unique microtag value.
- In the D-cache, valid bits must be cleared and microtag bits must be set so that each way within a set has a unique microtag value.
- All L2-cache tags and data

The I-MMU and D-MMU TLBs must also be initialized. The P-cache valid bits must be initialized before any floating-point loads are executed.

The MCU refresh control register as well as the Fireplane configuration register must be initialized after a POR.

In SSM systems, the MTags contained in memory must be initialized before any Fireplane transactions are generated.

---

**Caution** – Executing a DONE or RETRY instruction when TSTATE is not initialized after a POR can damage the chip. The POR boot code should initialize TSTATE<3:0>, using wrpr writes, before any DONE or RETRY instructions are executed.

However, these operations can only be executed in privileged mode. Therefore, user code is not at the risk of damaging the chip.

---

## *System Reset (Soft Reset)*

A system reset occurs when the Reset pin is activated. When the Reset pin is active, all other resets and traps are ignored. System reset has a trap type of 1 at physical address offset 20<sub>16</sub>. Any pending external transactions are canceled.

---

**Note** – Memory refresh continues uninterrupted during a system reset. System interface, L2-cache configuration, and memory controller configuration are preserved across a system reset.

---

### *Externally Initiated Reset (XIR)*

An XIR is sent to the processor through an external hardware pin. It causes a SPARC V9 XIR, which has a trap type  $3_{16}$  at physical address offset  $60_{16}$ . XIR has higher priority than all other resets except Power-on Reset and System Reset.

XIR affects only one processor, rather than the entire system. Memory state, cache state, and most Control Status Register state are unchanged. System coherency is *not* guaranteed to be maintained through an XIR reset. The saved PC and nPC will only be approximate because the trap is not precise with respect to pipeline state.

### *Watchdog Reset (WDR) and error\_state*

The processor enters `error_state` when a trap occurs at `TL = MAXTL`.

The processor automatically exits `error_state` using WDR. The processor signals itself internally to take a WDR and sets `TT = 2`. The WDR traps to the address at `RSTVaddr + 0x4016`. WDR sets the processor in a state where it is prepared for diagnosis of failures.

WDR affects only one processor rather than the entire system. CWP updates due to window traps that cause watchdog traps are the same as the no watchdog trap case.

### *Software-Initiated Reset (SIR)*

An SIR is initiated by an SIR instruction within any processor. This per-processor reset has a trap type 4 at physical address offset  $80_{16}$ . SIR affects only one processor rather than the entire system.

### *RED\_state Trap Vector*

When the UltraSPARC III Cu processor processes a reset or trap that enters `RED_state`, it takes a trap at an offset relative to the `RED_state` trap vector base address (`RSTVaddr`); the base address is at virtual address `FFFF FFFF F000 000016`, which passes through to physical address `7FF F000 000016`.

## 3.2.4 Error Handling

The UltraSPARC III Cu processor provides extensive support for detecting and correcting errors. Note that some errors may still be uncorrectable.

## *Error Classes in Severity*

The classes of error in order of severity are as follows:

1. **Hardware-corrected errors.** Hardware tries to correct the error automatically. A trap is generated to log the error conditions when the error is corrected to enable the actions for preventive maintenance.
2. **Software-correctable errors.** Hardware does not correct the error automatically. Instead, it invokes a trap requesting the recovery software to correct the error. Corrective actions are expected from the recovery software. If recovery is successful, the system should continue the operation.
3. **Uncorrectable errors.** By its nature the error is uncorrectable, and hardware invokes a trap to signal the occurrence of the error to appropriate recovery software. Depending on the condition under which the error occurs, the system may be able to recover from the error and continue operation. If not, it may be able to isolate the error to a particular process and terminate it. Otherwise, the software should shutdown the system gracefully.
4. **Fatal errors.** By its nature, the error indicates either loss of system consistency or a system interconnect protocol error. It is dangerous to continue operation in this situation because of the impending threat of a failure to maintain data integrity. Therefore, upon the detection of the error, the processor generates an ERROR signal to its interconnect, expecting to be halted/reset by the system. System actions induced by the ERROR signal are system implementation dependent.

## *Errors Synchronous and Asynchronous to Instruction Execution*

Some errors can be detected asynchronously to instruction execution. Other errors are detected in the course of an instruction execution, that is, synchronous to instruction execution. Separate error recording mechanisms are used for synchronous and asynchronous errors.

An error asynchronous to instruction execution is signalled either through a disrupting trap to the processor or through an ERROR signal to system hardware to induce a system reset, depending on the severity of the error.

The errors signalled through a disrupting trap do not directly correspond to an instruction. Traps may or may not be recoverable.

Errors signalled with an ERROR are meant either to be loss of system consistency or a protocol error on system interconnect.

On the other hand, an error detected in the course of an instruction execution is signalled through an error trap to the instruction, with additional information recorded in hardware. The trap is either precise or deferred. The program (process) affected by the error should be given a corrected response, or if the error is uncorrectable, the process should be terminated appropriately. Precise traps are used wherever possible.

## *Corrective Actions*

Errors are handled by invocation of one of the following actions:

- **Reset-inducing ERROR signal.** The most severe fatal error generates an ERROR signal to induce a system reset. Both, an error detected in the course of instruction execution and an error asynchronous to instruction execution may generate an ERROR signal.
- **Precise traps.** Most errors detected in the course of an instruction execution generate a precise trap. If the error is hardware correctable, software just logs it. If the error is software correctable, software corrects it before continuing execution. If the error is uncorrectable, software takes appropriate action.
- **Deferred traps.** Some uncorrectable errors requiring immediate attention generate a deferred trap to request software intervention. The recovery software examines the recorded error information to determine the extent of the damage caused by the error. Depending on the observed effect, the system may need to be brought down, or it may continue to run when the effect is isolated within the user program. In any event, the error does not require immediate reset of the system.
- **Disrupting traps.** An error asynchronous to instruction execution generates a disrupting trap to request logging and clearing. The error may already be corrected by hardware and may only require logging. If the error is software correctable, software corrects it before continuing execution. If the error is uncorrectable, software takes appropriate action.

### 3.2.5 Debug and Diagnostics Mode

The UltraSPARC III Cu processor provides interfaces for diagnostic access to most internal state of the processor. This is important for diagnosing, and when possible recovering from failures. There are a couple of different diagnostic interfaces. All the diagnostic interfaces are accessible only from software running in privileged mode or from an external system controller in a server.

There are a number of diagnostic registers that are mapped to internal ASI registers. These registers are accessed by load and store alternate ASI instructions that specify certain configurations of ASI numbers and virtual addresses to access the register (all internal registers are 8 bytes and must be accessed as 8-byte units with 8-byte aligned addresses). Diagnostic registers are provided for recording various fault conditions as well as important information and state associated with the fault to help diagnosis and possibly recover.

For diagnostic and error recovery, large memories on chip, such as caches, can have each element of the memory array be directly read and written. These accesses are performed with load and store alternate ASIs that use specific ASIs that point to the memory array. These accesses can only be done by privileged software.

Special ASI numbers are used for diagnostic accesses to structures where the virtual address is used to specify the portion of the structure to be read (all internal state must be accessed in 8-byte units with 8-byte aligned addresses). Most structures can be directly read and many structures can also be directly written or quickly cleared.

The UltraSPARC III Cu processor also provides a serial JTAG interface that can be used by a system controller for diagnostics. A system controller can perform a shadow scan where various configuration and diagnostic information is scanned out of the processor without interfering with the operation of the processor. The system controller can also use the JTAG interface to scan in information to configure or control various aspects of the processor.

The JTAG interface can also be used to perform a full scan dump. When a full scan dump is performed, most of the flops in the processor are scanned out through a scan chain. A full scan dump is a destructive action and the processor must be reset after a full scan dump. The full scan provides an important tool for diagnosis of serious failures.

For controlling diagnostics mode, there is a range of configuration registers, which can enable and disable many features of the processor. The configuration registers are only accessible in privileged mode. Some of the configuration registers are implemented as ASRs. These registers are accessible from the RDASR/WRASR interface. Most of the configuration registers are mapped as internal ASI registers. These registers are accessed by load and store alternate ASI instructions that specify certain configurations of ASI numbers and virtual addresses to access the register (all internal registers are 8 bytes and must be accessed as 8-byte units with 8-byte aligned addresses).





# Instruction Execution

---

This chapter focuses on the needs of compiler writers and others who are interested in scheduling instructions to optimize program performance. The chapter discusses the following topics:

- Section 4.1, “Introduction”
- Section 4.2, “Processor Pipeline”
- Section 4.3, “Pipeline Recirculation”
- Section 4.4, “Grouping Rules”
- Section 4.5, “Conditional Moves”
- Section 4.6, “Instruction Latencies and Dispatching Properties”

---

## 4.1 Introduction

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

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

## 4.1.1 NOP, Neutralized, and Helper Instructions

The distinction between NOP and neutralized instructions is subtle.

### 4.1.1.1 NOP Instruction

The architected NOP instruction is coded as a SETHI instruction with destination register %g0. This instruction is groupable in the A0 or A1 pipeline.

### 4.1.1.2 Neutralized Instruction

Some instructions have no visible effects on the software. They have been de-implemented or assigned to not have an effect if the processor is in a certain mode. These instructions are often referred to as NOP instructions, but they are not the same as the NOP instruction in that they execute in the pipeline that is assigned to them. These are versions of instructions that have no effect because they only access the %g0 register and do not have any side-effects. Hence, these instructions are functionally neutral.

### 4.1.1.3 Helper Instructions

Helper instructions are generated by the hardware to help in the execution or re-execution of an instruction. The hardware partitions a single instruction into multiple instructions that flow through the pipeline consecutively. They have no software visibility and are part of the hardware function of the pipeline.

---

## 4.2 Processor Pipeline

The processor pipeline consists of fourteen stages plus an extra stage that is occasionally used by the hardware. The pipeline stages are referred to by the following mnemonic single letter names and are shown in TABLE 4-1.

**TABLE 4-1** Processor Pipeline Stages

| Pipeline Stage | Definition                      |
|----------------|---------------------------------|
| <b>A</b>       | Address generation              |
| <b>P</b>       | Preliminary Fetch               |
| <b>F</b>       | Fetch instructions from I-cache |

**TABLE 4-1** Processor Pipeline Stages (*Continued*)

| Pipeline Stage | Definition   |
|----------------|--|
| <b>B</b>       | Branch target computation                            |
| <b>I</b>       | Instruction group formation                          |
| <b>J</b>       | J: grouping  |
| <b>R</b>       | Register access (dispatch/dependency checking stage) |
| <b>E</b>       | Execute  |
| <b>C</b>       | Cache  |
| <b>M</b>       | Miss detect  |
| <b>W</b>       | Write  |
| <b>X</b>       | eXtend   |
| <b>T</b>       | Trap   |
| <b>D</b>       | Done   |

Rather than executing the instructions in a single pipeline, several separate pipelines are each dedicated to execution of a particular class of instructions. The execution pipelines start after the R-stage of the pipeline. Some instructions take a cycle or two to execute, others take a few cycles within the pipeline. As long as the execution fits within the fixed pipeline depth, execution can in general be fully pipelined. Some instructions have extended execution times that sometimes vary in duration depending on the state of the processor.

The following sections provide a stage-by-stage description of the pipeline. Chapter 3, “CPU Architecture Basics,” describes the functions of the various execution units. This chapter explains how the pipeline operates the execution units to process the instructions.

FIGURE 4-1 illustrates each pipeline stage in detail and the relationship between high level, large architectural structures.

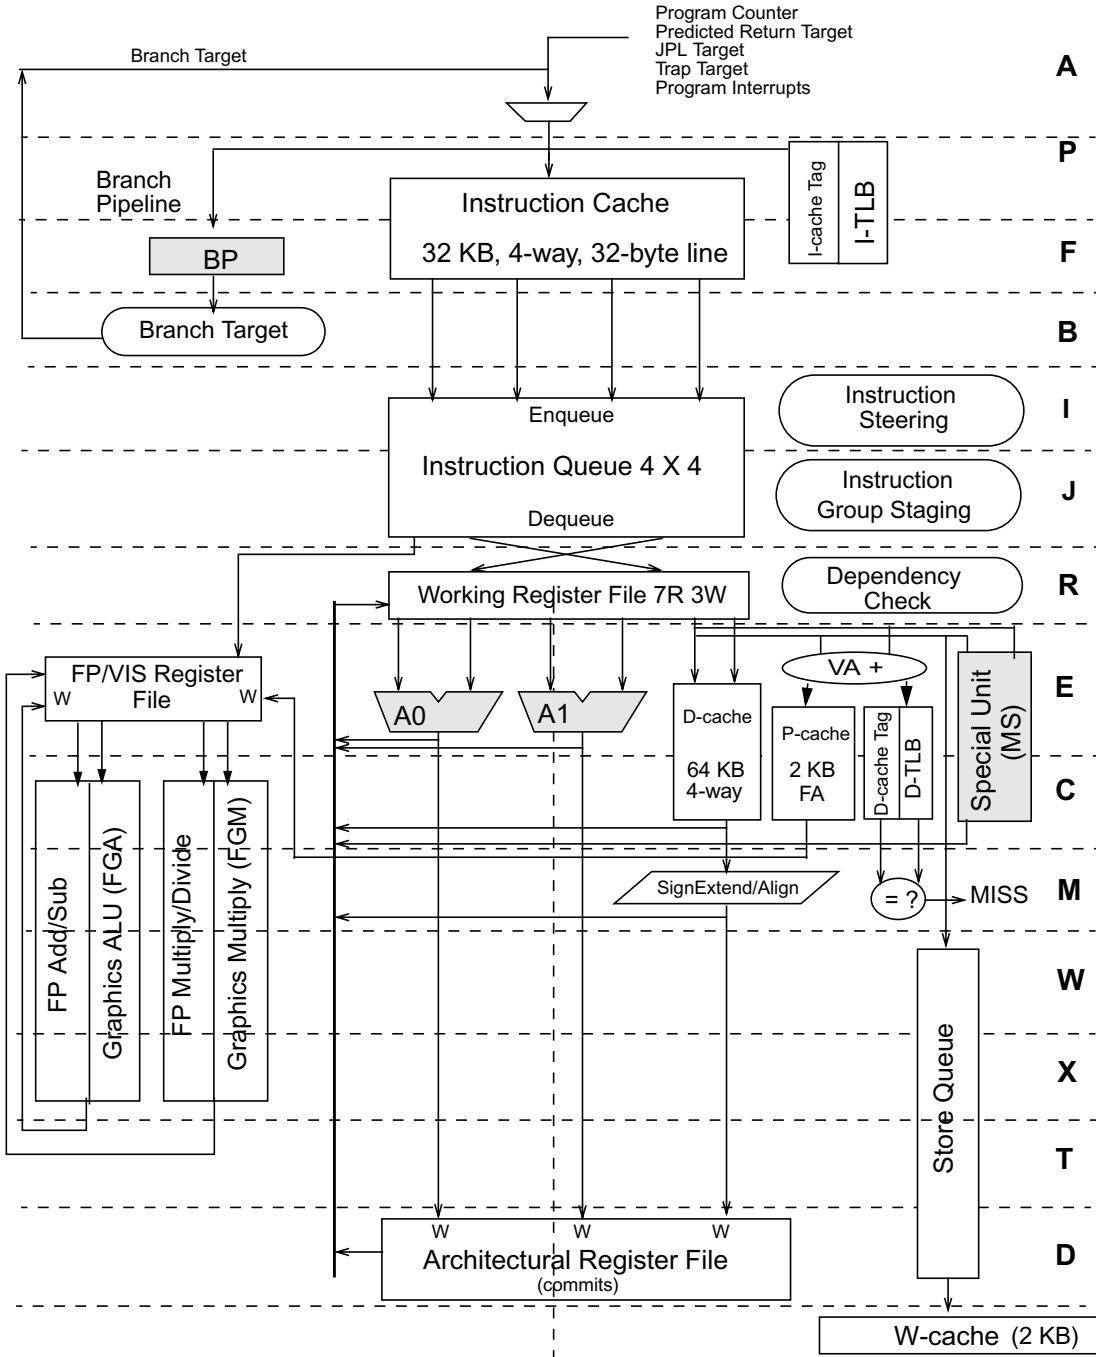


FIGURE 4-1 Instruction Pipeline Diagram Instruction Dependencies

Instruction dependencies exist in the grouping, dispatching, and execution of instructions.

## 4.2.1 Grouping Dependencies

Up to four instructions can be grouped together for simultaneous dispatch. The number of instructions that can be grouped together depends on the consecutive instructions that are present in the instruction fetch stream, the availability of execution resources (execution units), and the state of the system. Instructions are grouped together to provide superscalar execution of multiple instruction dispatches per clock cycle.

Some instructions are *single instruction group* instructions. These are dispatched by themselves one clock at a time as a single instruction in the group.

---

**Note – Pipeline Recirculation:** During recirculation, the recirculation invoking instruction is often re-executed as a single group instruction and often with helper instruction inserted into the pipeline by the hardware. Even groupable instructions are retried in a single instruction group. See Section 4.3, “Pipeline Recirculation” for details.

---

## 4.2.2 Dispatch Dependencies

Instructions can be held at the R-stage for many different reasons, including:

- Working register operand not available
- Functional Unit not available
- Store-load sequence in progress (atomic operation)

When instructions are held at the dispatch stage, the upper pipeline continues to operate until the instruction buffer is full. At that point, the upper pipeline stalls.

During recirculation, the recirculation invoking instruction is held at the dispatch stage until its execution dependency is resolved.

## 4.2.3 Execution Dependencies

The pipeline assumes all load instructions will hit in a primary cache, allowing the pipeline to operate at full speed. A cache miss will recirculate the pipeline.

- D-cache Miss
- Load requires data to be bypassed from an earlier store that has not completed and does not meet the criteria for read-after-write data bypassing.

## 4.2.4 Instruction-Fetch Stages

The instruction-fetch pipeline stages A, P, F, and B are described below.

### 4.2.4.1 A-stage (Address Generation)

The address stage generates and selects the fetch address to be used by the instruction cache in the next cycle. The address that can be selected in this stage for instruction fetching comes from several sources, including:

- Sequential PC
- Branch target (from B-stage)
- Trap target
- Interrupt
- Predicted return target
- Jmpl target
- Resolved branch/Jmpl target from execution pipeline

### 4.2.4.2 P-stage (Preliminary Fetch)

The preliminary fetch stage starts fetching four instructions from the instruction cache. Since the I-cache has a two-cycle latency, the P-stage and the F-stage are both used to complete an I-cache access. Although the I-cache has a two-cycle latency, it is pipelined and can access a new set of up to four instructions every cycle. The address used to start an I-cache access is generated in the previous cycle.

The P-stage also accesses the Branch Predictor (BP), which is a small, single-cycle access SRAM whose output is latched at the end of the P-stage. The BP predicts the direction of all conditional branches, based on the PC of the branch and the direction history of the most recent conditional branches.

### 4.2.4.3 F-stage (Fetch)

The F-stage is used for the second half of the I-cache access. At the end of this stage, up to four instructions from an I-cache line (32 bytes) are latched for decode. An I-cache fetch group is not permitted to cross an I-cache line (32-byte boundary).

#### 4.2.4.4 B-stage (Branch Target Computation)

The B-stage is the final stage of the instruction-fetch pipeline, A-P-F-B. In this stage, the four fetched instructions are first available in a register. The processor analyzes the instructions, looking for Delayed Control Transfer Instructions (DCTI) that can alter the path of execution. It finds the first DCTI, if any, among the four instructions and computes (if PC relative) or predicts (if register based) its target address. If this DCTI is predicted taken, the target address is passed to the A-stage to begin fetching from that stream; if predicted not taken, the target is passed on to the CTI queue for use in case of mispredict. Also in the B-stage, the computation of the hit or miss status of the instruction fetch is performed, so that the validity of the four instructions can be reported to the instruction queue.

In the case of an instruction cache miss, a request is issued to the L2-cache and all the way out to memory if needed to get the required line. The processor includes an optimization, where along with the line being fetched, the subsequent line (32 bytes) is also returned and placed into the instruction prefetch buffer. A subsequent miss that can get its instructions from the instruction prefetch buffer will behave like a fast miss.

### 4.2.5 Instruction Issue and Queue Stages

The I-stage and J-stage correspond to the enqueueing and dequeueing of instructions from the instruction queue. The R-stage is where instruction dependencies are resolved.

#### 4.2.5.1 I-stage (Instruction Group Formation)

In the I-stage, the instructions fetched from the I-cache are entered as a group into the instruction queue. The instruction queue is four instructions wide by four instruction groups deep. The instruction may wait in the queue for an arbitrary period of time until all earlier instructions are removed from the queue.

The instructions are grouped to use up to four of the execution pipelines, shown in TABLE 4-2.

**TABLE 4-2** Execution Pipelines

| Pipeline | Description  |
|----------|--|
| A0       | Integer ALU pipeline 0   |
| A1       | Integer ALU pipeline 1   |
| BR       | Branch pipeline  |
| MS       | Memory/Special pipeline  |
| FGM      | Floating Point/VIS multiply pipeline (with divide/square root pathway) |
| FGA      | Floating Point/VIS add ALU pipeline                                    |

## 4.2.5.2 J-stage (Instruction Group Staging)

In the J-stage, a group of instructions are dequeued from the instruction queue and prepared for being sent to the R-stage. If the R-stage is expected to be empty at the end of the current cycle, the group is sent to the R-stage.

## 4.2.5.3 R-stage (Dispatch and Register Access)

The integer working register file is accessed during the R-stage for the operands of the instructions (up to three) that have been steered to the A0, A1, and MS pipelines. At the end of the R-stage, results from previous instructions are bypassed in place of the register file operands, if required.

Up to two floating-point or VIS instructions are sent to the Floating Point/VIS Unit in this stage.

The register and pipeline dependencies between the instructions in the group and the instructions in the execution pipelines are calculated concurrently with the register file access. If a dependency is found, the dependent instruction and any older instruction in the group is held in the R-stage until the dependency is resolved.

## 4.2.6 Execution Pipeline

The execution pipeline contains the E, C, M, W and X stages.

### 4.2.6.1 Integer Instruction Execution: E-stage (Execute)

The E-stage is the first stage of the execution pipelines. Different actions are performed in each pipeline.

Integer instructions in the A0 and A1 pipelines compute their results in the E-stage. The instructions include most arithmetic, all shift, and all logical instructions. The results are available for bypassing to dependent instructions that are in the R-stage, resulting in single-cycle execution for most integer instructions. The A0 and A1 pipelines are the only two sources of bypass results in the E-stage.

Other integer instructions are steered to the MS pipeline and if necessary are sent with their operands to the special execution unit in this stage. They can start their execution during the E-stage, but will not produce any results to be bypassed until the C-stage or the M-stage.



Load instructions steered to the MS pipeline start accessing the D-cache during the E-stage. The D-cache features Sum Addressed Memory (SAM) decode logic that combines the arithmetic calculation for the virtual address with the row decode of the memory array to reduce lookup time. The virtual address is computed in the E-stage for translation lookaside buffer (TLB) access and possible access to the P-cache.

Floating-point and VIS instructions access the floating-point register file in the E-stage to obtain their operands. At the end of the E-stage, the results from previous completing floating-point/VIS instructions can be bypassed to the E-stage instructions.

Conditional branch instructions in the BR pipeline resolve their directions in the E-stage. Based on their original predicted direction, a *mispredict* signal is computed and sent to the A-stage for possible refetching of the correct instruction stream.

JMPL and RETURN instructions compute their target addresses in the E-stage of the MS pipeline. The results are sent to the A-stage to start fetching instructions from the target stream.

#### 4.2.6.2 C-stage (Cache)

The data cache delivers results for doubleword (64-bit) and unsigned word (32-bit) integer loads in the C-stage. The D-TLB access is initiated in the C-stage and proceeds in parallel with the D-cache access. For floating-point loads, the P-cache access is initiated in the C-stage. The results of the D-TLB access and P-cache access are available in the M-stage.

Special instruction unit results are produced at the end of this stage and can be bypassed to waiting dependent instructions in the R-stage — minimum two-cycle latency for SIU instructions. The integer pipelines, A0 and A1, write their results back to the working register file in the C-stage.

The C-stage is the first stage of execution for floating-point and VIS instructions in the FGA and FGM pipelines.

#### 4.2.6.3 M-stage (Miss)

Data cache misses are determined in the M-stage by a comparison of the physical address from the D-TLB to the physical address in the D-cache tags. If the load requires additional alignment or sign extension (such as signed word, all halfword, and all byte loads), it is carried out in this stage, resulting in a three-cycle latency for those load operations. This stage is used for the second execution cycle of floating-point and VIS instructions. Load data are available to the floating-point pipelines in the M-stage.

#### 4.2.6.4 W-stage (Write)

In the W-stage, the MS integer pipeline results are written into the working register file. The W-stage is also used as the third execution cycle for floating-point and VIS instructions. The results of the D-cache miss are available in this stage and the requests are sent to the L2-cache if needed.

#### 4.2.6.5 X-stage (Extend)

The X-stage is the last execution stage for most floating-point operations (except divide and square root) and for all VIS instructions. Floating-point results from this stage are available for bypass to dependent instructions that will be entering the C-stage in the next cycle.

### 4.2.7 Trap and Done Stages

This section describes the stages that interrupt or complete instruction execution.

The results of operations are bypassed and sent to the working register file. If no traps are generated, then they are successfully pipelined down to the architectural register file and committed. If a trap or recirculation occurs, then the architectural register file (contains committed data) is copied to the working register in preparation for the instructions to be re-executed.

#### 4.2.7.1 T-stage (Trap)

Traps, including floating-point and integer traps, are signalled in this stage. The trapping instruction, and all instructions younger than the trapping instruction must invalidate their results before reaching the D-stage to prevent their results from being erroneously written into the architectural or floating-point register files.

#### 4.2.7.2 D-stage (Done)

Integer results are written into the architectural register file in this stage. At this point, they are fully committed and are visible to any traps generated from younger instructions in the pipeline.

Floating-point results are written into the floating-point register file in this stage. These results are visible to any traps generated from younger instructions.

---

## 4.3 Pipeline Recirculation

When a dependency is encountered *in* or *before* the dispatch R-stage, then the pipeline is stalled. Most dependencies, like register or functional unit dependencies are resolved in the R-stage. When a dependency is encountered *after* the dispatch R-stage, then the pipeline is *recirculated*. Recirculation involves resetting the PC back to the recirculation invoking instruction. Instructions older than the dependent instruction continue to execute. The offending instructions and all younger instructions are recirculated. The offending instruction is re-fetched and goes through the entire pipeline again.

Upon recirculation, the instruction responsible for the recirculation becomes a single-group instruction that is held in the R-stage until the dependency is resolved.

### *Load Instruction Dependency*

In the case of a load instruction miss in a primary cache, the pipeline recirculates and the load instruction waits in the R-stage. When the data is returned in the D-cache *fill buffer*, the load instruction is dispatched again and the data is provided to the load instruction from the fill buffer. The pipeline logic inserts two helpers behind the load instruction to move the data in the fill buffer to the D-cache. The instruction in the instruction fetch stream, after the load instruction, follows the helpers and will regroup with younger instructions, if possible.

---

## 4.4 Grouping Rules

Grouping rules are made before going into R-stage. A *group* is a collection of instructions with no resource constraints that will limit them from being executed in parallel.

Instruction grouping rules are necessary for the following reasons:

- The instruction execution order is maintained.
- Each pipeline runs a subset of instructions.
- Resource dependencies, data dependencies, and multicycle instructions require helpers (NOPs) to maintain the pipelines.

Before continuing, we define a few terms that apply to instructions.

**break-before:** The instruction will always be the first instruction of a group.

**break-after:** The instruction will always be the last instruction of a group.

**single-instruction group (SIG):** The instruction will not be issued with any other instructions in the group. (SIG is sometimes shortened herein to “single-group.”)

**instruction latency:** The number of processor cycles after dispatching an instruction from the R-stage that a following data-dependent instruction can dispatch from the R-stage.

**blocking, multicycle:** The instruction reserves one or more of the execution pipelines for more than one cycle. The reserved pipelines are not available for other instructions to issue into until the blocking, multicycle instruction completes.

## 4.4.1 Execution Order

**Rule: Within the R-stage, some of the instructions can be dispatched and others cannot. If an instruction is younger than an instruction that is not able to dispatch, then the younger instruction will not be dispatched.**

“Younger” and “older” refer to instruction order within the program. The instruction that comes first in the program order is the older instruction.

## 4.4.2 Integer Register Dependencies to Instructions in the MS Pipeline

**Rule: If a source register operand of an instruction in the R-stage matches the destination register of an instruction in the MS pipeline’s E-stage, then the instruction in the R-stage may not proceed.**

The MS pipeline has no E-stage bypass.

If an operand of an instruction in the R-stage matches the destination register of an instruction in the MS pipeline’s C-stage, then the instruction in the R-stage may not proceed if the instruction in the MS pipeline’s C-stage does not generate its data until the M-stage. For example, LDSB does not have the load data until the M-stage, but LDX has its data in the C-stage. Thus, LDX would not cause an interlock, but LDSB would.

Most instructions in the MS pipeline have their data by the M-stage, so there is no dependency check on the MS pipeline’s M-stage destination register. In the case of multicycle MS instructions, the data is always available by the M-stage as the last of the instructions passes through the pipeline.

### 4.4.2.1 Helpers

Sometimes an instruction, as part of its operation, requires multiple flows in the pipeline. We call those extra flows after the initial instruction flow *helper cycles*. The only pipeline that executes such instructions is the MS pipeline. If an instruction requires a helper, that helper is generated in the R-stage. The help generation logic generates as many helpers as the instruction requires.

Most of the time the logic determines the number of helpers by examining the opcode. However, some recirculate cases run the recirculated instruction differently than the original flow down the pipeline, and some instructions, like integer multiply and divide, require variable numbers of helpers. Some helper counts are determined by I/O and memory controllers and system devices. For example, the D-cache unit requires helpers as it completes an atomic memory instruction.

**Rule: Instructions requiring helpers are always break-after.**

There can be no instruction in a group that is younger than an instruction that requires helpers. Another way of saying this is “an instruction that requires helpers will be the youngest in its group.” This rule preserves the in-order execution of the integer instructions.

**Rule: Helpers block the pipeline.**

Helpers block the pipeline from executing other instructions; thus, instructions with helpers are blocking.

**Rule: Helpers are always single-group.**

A helper cycle is always alone in a group. No other instruction will ever be dispatched from the R-stage if there is a helper cycle in the R-stage.

### 4.4.3 Integer Instructions Within a Group

**Rule: Integer instructions within a group are not allowed to write the same destination register.**

By not writing the same destination register at the same time, we simplify bypass logic and register file write-enable determination and potential Write-after-Write (WAW) errors. The instructions are break-before second destination is written.

This rule applies only to integer instructions writing integer registers. Floating-point instructions and floating-point loads (done in the integer A0, A1, and MS pipelines) can be grouped so that two or more instructions in the same group can write the same floating-point destination register. Instruction age is associated with each instruction. The write from an older instruction is not visible, but the execution of the instruction might still cause a trap and set condition codes.

**There are no special rules concerning integer instructions that set condition codes and integer branch instructions.**

Integer instructions that set condition codes can be grouped in any way with integer branches. In fact, any number instructions that set condition codes can be in any order relative to the branch are allowed, provided that they do not violate any other rules. No special rules apply to this specific case. Integer instructions that set condition codes in the A1 and A0 pipelines can compute a taken/not taken result in the E-stage, which is the same stage in which the branch is evaluating the correctness of its prediction. The control logic guarantees that the correct condition codes are used in the evaluation.

## 4.4.4 Same-Group Bypass

**Rule: Same-group bypass is disallowed, except store instructions.**

The group bypass rule states that no instruction can bypass its result to another instruction in the same group. The one exception to this rule is store. A store instruction can get its store data (`rd`), but not its address operands (`rs1`, `rs2`), from an instruction in the same group.

## 4.4.5 Floating-Point Unit Operand Dependencies

### 4.4.5.1 Latency and Destination Register Addresses

Floating-point operations have longer latencies than most integer instructions. Moreover, floating-point square root and divide instructions have varying latencies depending on whether the operands are single precision or double precision. All the floating-point instruction latencies are four clock cycles (except for floating-point divide and square root and `PDIST` → `PDIST`).

The operands for floating-point operations can either be single precision (32-bit) or double precision (64-bit). Sixteen of the double precision registers are each made up of two single precision registers. An operation using one of these double precision registers as a source operand may be dependent on an earlier single precision operation producing part of the register value. Similarly, an operation using one of the single precision registers as a source operand may be dependent on an earlier double precision operation, a part of which may produce the single precision register value.

### 4.4.5.2 Grouping Rules for Floating-Point Instructions

**Rule: Floating-point divide/square root is busy.**

The floating-point divide/square root unit is a non-pipelined unit. The Integer Execution Unit sets a busy bit for each of the two stages of the divide/square root and depends on the FGU to clear them. Only the first part of the divide/square root is considered to have a busy unit; therefore, once the first part is complete, a new floating-point divide/square root operation can be started.

**Rule: Floating-point divide/square root needs a write slot in FGM pipeline.**

In the stage in which a divide/square root is moved from the first part to the last part, we cannot issue any instructions to the FGM pipeline. This constraint provides the write slot in the FGM pipeline so the divide/square root can write the floating-point register file.

**Rule: Floating-point store is dependent on floating-point divide/square root.**

The floating-point divide/square root unit has a latency longer than the normal pipeline. As a result, if a floating-point store depend on the result of a floating-point divide/square root, then the floating-point store instruction may not be dispatched until the floating-point divide/square root instruction has completed.

### 4.4.5.3 Grouping Rules for VIS Instructions

**Rule: Graphics Status Register (GSR) Write instructions are break-after.**

The SIAM, BMASK, and FALIGNADDR instructions write the GSR. The BSHUFFLE and FALIGNDATA instructions read the GSR in their operation. Because of GSR write latency, a GSR reader cannot be in the same group as a GSR writer unless the GSR reader is older than the GSR writer. The simplest solution to this dependency is to make all GSR write instructions break-after.

---

**Note** – The WRGSR instruction is not included in this rule as a special case. The WRGSR instruction is already break-after by virtue of being a WRASR instruction.

---

### 4.4.5.4 PDIST Special Cases

PDIST-to-dependent-PDIST is handled as a special case with one-cycle latency. PDIST latency to any other dependent operation is four-cycle latency. In addition, a PDIST cannot be issued if there is ST, block store (BST), or partial store instruction in the M-stage of the pipeline. PDIST issue is delayed if there is a store type instruction two groups ahead of it.

## 4.4.6 Grouping Rules for Register-Window Management Instructions

**Rule: Window changing instructions are single-group.**

The window changing instructions `SAVE`, `RESTORE`, and `RETURN` are all single-group instructions. These instructions are never grouped with any other instruction. This rule greatly simplifies the tracking of register file addresses.

**Rule: Window changing instructions force bubbles after.**

The window changing instructions `SAVE`, `RESTORE`, and `RETURN` also force a subsequent pipeline bubble. A bubble is distinct from a helper cycle in that there is nothing valid in the pipeline within a bubble. During the bubble, control logic transfers the new window from the Architectural Register File (ARF) to the Working Register File (WRF).

**Rule: `FLUSHW` is single-group.**

To simplify the Integer Execution Unit's handling of the register file window flush, the `FLUSHW` instruction is single-group.

**Rule: `SAVED` and `RESTORED` are single-group.**

To simplify the Integer Execution Unit's window tracking, `SAVED` and `RESTORED` are single-group instructions.

## 4.4.7 Grouping Rules for Reads and Writes of the ASRs

**Rule: Write ASR and Write PR instructions are single-group.**

`WRASR` and `WRPR` are always the youngest instructions in a group. This case prevents problems with an instruction being dependent on the result of the write, which occurs late in the pipeline.

**Rule: Write ASR and Write PR force seven bubbles after.**

To guarantee that any instruction that starts in the R-stage is started with the most up-to-date status registers, `WRASR` and `WRPR` force bubbles after they are dispatched. Thus, if a `WRASR` or a `WRPR` instruction is in the pipeline anywhere from the E-stage to the T-stage, no instructions are dispatched from the R-stage (bubbles are forced in).

**Rule: Read ASR and Read PR force up to six bubbles before (break-before multicyle).**

Many instructions can update the ASRs and PRs. Therefore, if an `RDASR` or `RDPR` instruction is in the R-stage and any valid instruction is in the integer pipelines from the E-stage to the X-stage, the UltraSPARC III Cu processor does not allow the `RDASR` and `RDPR` instructions to be dispatched. Instead, we wait for all pipeline states to write the ASRs and privileged registers and then read them.

## 4.4.8 Grouping Rules for Other Instructions

**Rule: Block Load (BLD) and Block Store (BST) are single-group and multicyle.**



For simplicity in the Integer Execution Unit and memory system, BLD and BST are single-group instructions with helpers.

**Rule: FLUSH is single-group and seven bubbles after.**

To simplify the Instruction Issue Unit and Integer Execution Unit, the FLUSH instruction is single-group. This makes instruction cancellation and issue easier. FLUSH is held in the R-stage until the store queue and the pipeline from E-stage through D-stage is empty.

**Rule: MEMBAR (#Sync, #Lookaside, #StoreLoad, #Memissue) is single-group.**

To simplify the Integer Execution Unit and memory system, MEMBAR is a single-group instruction. MEMBAR will not dispatch until the memory system has completed necessary transactions.

**Rule: Software-initiated reset (SIR) is single-group.**

For simplicity, SIR is a single-group instruction.

**Rule: Load FSR (LDFSR) is single-group and forces seven bubbles after.**

For simplicity, LDFSR is a single-group instruction.

**Rule: DONE and RETRY are single-group.**

DONE and RETRY instructions are dispatched as a single-group.

**Rule: DONE and RETRY force seven bubbles after.**

DONE and RETRY are typically used to return from traps or interrupts and are known as trap exit instructions.

It takes a few cycles to properly restore the pre-trap state and the working register file from the architectural register file, so we force bubbles after the trap exit instructions to give us the cycles to do it all. We will not accept a new instruction until the trap exit instruction leaves the pipeline (also known as  $D + 1$ ).

---

## 4.5 Conditional Moves

The compiler needs to have a detailed model of the implementation of the various conditional moves so it can optimally schedule code. TABLE 4-3 describes the implementation of the five classes of SPARC V9 conditional moves in the pipeline. FADD and ADD instructions (shaded rows) are also described as a reference for comparison with the conditional move instructions.

**TABLE 4-3** SPARC V9 Conditional Moves

| Instruction | RD Latency | Pipelines Used | Busy Cycles | Groupable | Dependency |
|-------------|------------|----------------|-------------|-----------|------------|
| FMOVicc     | 3 cycles   | FGA and BR     | 1           | Yes       | icc - 0    |
| FMOVfcc     | 3 cycles   | FGA and BR     | 1           | Yes       | fcc - 0    |
| FMOVr       | 3 cycles   | FGA and MS     | 1           | Yes       | N/A        |
| FADD        | 4 cycles   | FGA            | 1           | Yes       | N/A        |
| ADD         | 1 cycle    | A0 or A1       | 1           | Yes       | N/A        |
| MOVcc       | 2 cycles   | MS and BR      | 1           | Yes       | icc - 0    |
| MOVR        | 2 cycles   | MS and BR      | 1           | Yes       | N/A        |

**Where:**

**RD Latency** — The number of processor cycles until the destination register is available for bypassing to a dependent instruction.

**Pipelines Used** — The pipeline that the instruction uses when it is issued. The pipelines are shown in TABLE 4-2.

**Busy Cycles** — The number of cycles that the pipelines are not available for other instructions to be issued. A value of one signifies a fully pipelined instruction.

**Groupable** — Whether instructions using pipelines, other than those used by the conditional move, can be issued in the same cycle as the conditional move.

**{i,f}CC Dependency** — The number of cycles that a CC setting instruction must be scheduled ahead of the conditional move in order to avoid incurring pipeline stall cycles.

---

## 4.6 Instruction Latencies and Dispatching Properties

In this section, a machine description is given in the form of a table (TABLE 4-4) dealing with dispatching properties and latencies of operations. The static or nominal properties are modelled in the following terms (columns in TABLE 4-4), which are discussed below.

- Latencies
- Blocking properties in dispatching
- Pipeline resources (A0, A1, FGA, FGM, MS, BR)

- Break rules in grouping (before, after, single-group)

The pipeline assumes the primary cache will be accessed. The dynamic properties, such as the effect of a cache miss and other conditions, are not described here.

## 4.6.1 Latency

In the Latency column, latencies are minimum cycles at which a dependent operation (consumer) can be dispatched relative to the producer operation without causing a dependency stall or instructions holding back in the R-stage to execute.

Operations like `ADDCC` produce two results, one in the destination register and another in the condition codes. For such operations, latencies are stated as a pair `x,y`, where `x` is for the destination register dependence and `y` is for the condition code.

A zero latency implies that the producer and consumer operations may be grouped together in a single group, as in `{SUBCC, BE %icc}`.

Operations like `UMUL` have different latencies, depending on operand values. These are given as a range, min–max, for example, 6–8 in `UMUL`. Operations like `LDFSR` involve waiting for a specified condition. Such cases are described by footnotes and a notation like 32+ for `CASA` (meaning at least 32 cycles).

Cycles for branch operations (like `BPCC`) give the dispatching cycle of the retiring target operation relative to the branch. A pair of numbers, for example 0,8, is given, depending on the outcome of a branch prediction, where 0 means a correct branch prediction and 8 means a mispredicted case.

Special cases, such as `FCMP(s,d)`, in which latencies depend on the type of consuming operations are described in footnotes (bracketed, for example, [1]).

## 4.6.2 Blocking

The Blocking column gives the number of clock cycles that the dispatch unit waits before issuing another group of instructions. Operations like `FDIVd` (MS pipeline) have limited blocking property; that is, the blocking is limited to the time before another instruction that uses MS pipeline can be dispatched. Such cases are noted with footnotes. All pipelines block instruction dispatch when an instruction is targeted to them, but they are not ready for another instruction to be pipelined-in.

## 4.6.3 Pipeline

The Pipeline column specifies the resource usage. Operations like `MOVcc` require more than one resource, as designated by the notation MS and BR. The operation `LDF` can dispatch to either MS, A0, or A1 as indicated.

## 4.6.4 Break and SIG

Grouping properties are given in columns Break and SIG (single-instruction group). In the Break column, an entry can be “Before,” meaning that this operation causes a break in a group so that the operation starts a new group. Operations like `RDCCR` require dispatching to be stalled until all operations in flight are completed (reach D-stage); in such cases, details are provided in a footnote reference in the Break column.

Operations like `ALIGNADDR` must be the last in an instruction group, causing a break in the group of type “After.”

Certain operations are not groupable and therefore are issued in single-instruction groups. A break “before” and “after” are implied for non-groupable instructions.

**TABLE 4-4** UltraSPARC III Cu Instruction Latencies and Dispatching Properties (*1 of 6*)

| Instruction                   | Latency  | Dispatch Blocking After | Pipeline  | Break | SIG |
|-------------------------------|----------|-------------------------|-----------|-------|-----|
| <code>ADD</code>              | 1        |                         | A0 or A1  |       |     |
| <code>ADDcc</code>            | 1,0 [1]  |                         | A0 or A1  |       |     |
| <code>ADDC</code>             | 5        | 4                       | MS        |       | Yes |
| <code>ADDCcc</code>           | 6,5 [2]  | 5                       | MS        |       | Yes |
| <code>ALIGNADDR</code>        | 2        |                         | MS        | After |     |
| <code>ALIGNADDRL</code>       | 2        |                         | MS        | After |     |
| <code>AND</code>              | 1        |                         | A0 or A1  |       |     |
| <code>ANDcc</code>            | 1,0 [1]  |                         | A0 or A1  |       |     |
| <code>ANDN</code>             | 1        |                         | A0 or A1  |       |     |
| <code>ANDNcc</code>           | 1,0 [1]  |                         | A0 or A1  |       |     |
| <code>ARRAY(8,16,32)</code>   | 2        |                         | MS        |       |     |
| <code>Bicc<sup>D</sup></code> | 0, 8 [3] | 0, 5 [4]                | BP        |       |     |
| <code>BMASK</code>            | 2        |                         | MS        | After |     |
| <code>BPcc</code>             | 0, 8 [3] | 0, 5 [4]                | BP        |       |     |
| <code>BPR</code>              | 0, 8 [3] | 0, 5 [4]                | BP and MS |       |     |
| <code>BSHUFFLE</code>         | 3        |                         | FGA       |       | Yes |

**TABLE 4-4** UltraSPARC III Cu Instruction Latencies and Dispatching Properties (2 of 6)

| Instruction          | Latency    | Dispatch Blocking After | Pipeline   | Break      | SIG |
|----------------------|------------|-------------------------|------------|------------|-----|
| CALL <i>label</i>    | 0-3 [5]    |                         | BP and MS  |            |     |
| CASA                 | 32+        | 31+                     | MS         | After      |     |
| CASXA                | 32+        | 31+                     | MS         | After      |     |
| DONE <sup>P</sup>    | 7          | Yes                     | BP and MS  |            | Yes |
| EDGE(8,16,32){L}     | 5          | 4                       | MS         |            | Yes |
| EDGE(8,16,32)N       | 2          |                         | MS         |            |     |
| EDGE(8,16,32)LN      | 2          |                         | MS         |            |     |
| FABS( <i>s,d</i> )   | 3          |                         | FGA        |            |     |
| FADD( <i>s,d</i> )   | 4          |                         | FGA        |            |     |
| FALIGNDATA           | 3          |                         | FGA        |            |     |
| FANDNOT1{ <i>s</i> } | 3          |                         | FGA        |            |     |
| FANDNOT2{ <i>s</i> } | 3          |                         | FGA        |            |     |
| FAND{ <i>s</i> }     | 3          |                         | FGA        |            |     |
| FBPfcc               |            |                         | BP         |            |     |
| FBfcc <sup>D</sup>   |            |                         | BP         |            |     |
| FCMP( <i>s,d</i> )   | 1,5 [6]    |                         | FGA        |            |     |
| FCMPE( <i>s,d</i> )  | 1,5 [6]    |                         | FGA        |            |     |
| FCMPEQ(16,32)        | 4          |                         | MS and FGA |            |     |
| FCMPGT(16,32)        | 4          |                         | MS and FGA |            |     |
| FCMPLE(16,32)        | 4          |                         | MS and FGA |            |     |
| FCMPNE(16,32)        | 4          |                         | MS and FGA |            |     |
| FDIVd                | 20(14) [6] | 17(11) [7]              | FGM        |            |     |
| FDIVs                | 17(14) [6] | 14(11) [7]              | FGM        |            |     |
| FEXPAND              | 3          |                         | FGA        |            |     |
| FiTO( <i>s,d</i> )   | 4          |                         | FGA        |            |     |
| FLUSH                | 8          | 7                       | BP and MS  | Before [8] | Yes |
| FLUSHW               |            | Yes                     | MS         |            | Yes |
| FMOV( <i>s,d</i> )   | 3          |                         | FGA        |            |     |
| FMOV( <i>s,d</i> )cc | 3          |                         | FGA and BP |            |     |
| FMOV( <i>s,d</i> )r  | 3          |                         | FGA and MS |            |     |
| FMUL( <i>s,d</i> )   | 4          |                         | FGM        |            |     |
| FMUL8(,SU,UL)x16     | 4          |                         | FGM        |            |     |
| FMUL8x16(AL,AU)      | 4          |                         | FGM        |            |     |
| FMULD8(SU,UL)x16     | 4          |                         | FGM        |            |     |

**TABLE 4-4** UltraSPARC III Cu Instruction Latencies and Dispatching Properties (3 of 6)

| Instruction             | Latency         | Dispatch<br>Blocking<br>After | Pipeline      | Break | SIG |
|-------------------------|-----------------|-------------------------------|---------------|-------|-----|
| FNAND{s}                | 3               |                               | FGA           |       |     |
| FNEG(s,d)               | 3               |                               | FGA           |       |     |
| FNOR{s}                 | 3               |                               | FGA           |       |     |
| FNOT(1,2){s}            | 3               |                               | FGA           |       |     |
| FONE{s}                 | 3               |                               | FGA           |       |     |
| FORNOT(1,2){s}          | 3               |                               | FGA           |       |     |
| FOR{s}                  | 3               |                               | FGA           |       |     |
| FPACK(FIX, 16,32)       | 4               |                               | FGM           |       |     |
| FPADD(16, 16s, 32, 32s) | 3               |                               | FGA           |       |     |
| FPMERGE                 | 3               |                               | FGA           |       |     |
| FPSUB(16, 16s, 32, 32s) | 3               |                               | FGA           |       |     |
| FsMULd                  | 4               |                               | FGM           |       |     |
| FSQRTd                  | 29(14) [6]      | 26(11) [7]                    | FGM           |       |     |
| FSQRTs                  | 23(14) [6]      | 20(11) [7]                    | FGM           |       |     |
| FSRC(1,2){s}            | 3               |                               | FGA           |       |     |
| F(s,d)TO(d,s)           | 4               |                               | FGA           |       |     |
| F(s,d)TOi               | 4               |                               | FGA           |       |     |
| F(s,d)TOx               | 4               |                               | FGA           |       |     |
| FSUB(s,d)               | 4               |                               | FGA           |       |     |
| FXNOR                   | 3               |                               | FGA           |       |     |
| FXOR{s}                 | 3               |                               | FGA           |       |     |
| FxTO(s,d)               | 4               |                               | FGA           |       |     |
| FZERO{s}                | 3               |                               | FGA           |       |     |
| ILLTRAP                 |                 |                               | MS            |       |     |
| JMPL reg,%o7            | 0-4, 9-10 [9]   | 0-3, 8-9                      | MS and BP     |       |     |
| JMPL %i7+8,%g0          | 3-5, 10-12 [10] | 2-4, 9-11                     | MS and BP     |       |     |
| JMPL %o7+8,%g0          | 0-4, 9 [11]     | 0-3, 8                        | MS and BP     |       |     |
| LDD <sup>D</sup>        | 2               | Yes                           | MS            | After |     |
| LDDA <sup>D</sup>       | 2               | Yes                           | MS            | After |     |
| LDDF{A}                 | 3               |                               | MS, A0, or A1 |       |     |
| LDF{A}                  | 3               |                               | MS, A0, or A1 |       |     |
| LDFSR <sup>D</sup>      | [22]            | Yes                           | MS            |       | Yes |
| LDSB{A}                 | 3               |                               | MS            |       |     |
| LDSH{A}                 | 3               |                               | MS            |       |     |

**TABLE 4-4** UltraSPARC III Cu Instruction Latencies and Dispatching Properties (4 of 6)

| Instruction        | Latency  | Dispatch Blocking After | Pipeline  | Break       | SIG |
|--------------------|----------|-------------------------|-----------|-------------|-----|
| LDSTUB{A}          | 31+      | 30+                     | MS        | After       |     |
| LDSW{A}            | 3        |                         | MS        |             |     |
| LDUB{A}            | 3        |                         | MS        |             |     |
| LDUH{A}            | 3        |                         | MS        |             |     |
| LDUW{A}            | 2        |                         | MS        |             |     |
| LDX{A}             | 2        |                         | MS        |             |     |
| LDXF <sub>SR</sub> | [22]     | Yes                     | MS        |             | Yes |
| MEMBAR #LoadLoad   |          | [12]                    | MS        |             | Yes |
| MEMBAR #LoadStore  |          | [12]                    | MS        |             | Yes |
| MEMBAR #Lookaside  |          | [13]                    | MS        |             | Yes |
| MEMBAR #MemIssue   |          | [13]                    | MS        |             | Yes |
| MEMBAR #StoreLoad  |          | [13]                    | MS        |             | Yes |
| MEMBAR #StoreStore |          | [12]                    | MS        |             | Yes |
| MEMBAR #Sync       |          | [14]                    | MS        |             | Yes |
| MOV <sub>cc</sub>  | 2        |                         | MS and BP |             |     |
| MOV <sub>fcc</sub> | 2        |                         | MS and BP |             |     |
| MOV <sub>r</sub>   | 2        |                         | MS        |             |     |
| MUL <sub>Scc</sub> | 6,5 [2]  | 5                       | MS        |             | Yes |
| MULX               | 6-9      | 5-8                     | MS        | After       |     |
| NOP                | N/A      |                         | MS        |             |     |
| OR                 | 1        |                         | A0 or A1  |             |     |
| OR <sub>cc</sub>   | 1,0 [1]  |                         | A0 or A1  |             |     |
| ORN                | 1        |                         | A0 or A1  |             |     |
| ORN <sub>cc</sub>  | 1,0 [1]  |                         | A0 or A1  |             |     |
| PDIST              | 4        |                         | FGM       |             |     |
| POPC               | Emulated |                         |           |             |     |
| PREFETCH{A}        |          |                         | MS        |             |     |
| RDASI              | 4        |                         | MS        | Before [15] |     |
| RDASR              | 4        |                         | MS        | Before [15] |     |
| RDCCR              | 4        |                         | MS        | Before [15] |     |
| RDDCR <sup>P</sup> |          |                         |           |             |     |
| RDFPRS             | 4        |                         | MS        | Before [15] |     |
| RDPC               | 4        |                         | MS        | Before [15] |     |
| RDPR               | 4        |                         | MS        | Before [15] |     |

**TABLE 4-4** UltraSPARC III Cu Instruction Latencies and Dispatching Properties (5 of 6)

| Instruction            | Latency       | Dispatch Blocking After | Pipeline  | Break       | SIG |
|------------------------|---------------|-------------------------|-----------|-------------|-----|
| RDSOFTINT <sup>P</sup> |               |                         |           |             |     |
| RDTICK                 | 4             |                         | MS        | Before [15] |     |
| RDY <sup>D</sup>       | 4             |                         | MS        | Before [15] |     |
| RESTORE                | 2             | 1                       | MS        | Before [16] | Yes |
| RESTORED <sup>P</sup>  |               |                         | MS        |             | Yes |
| RETRY <sup>P</sup>     | 2             | Yes                     | MS and BP | After       |     |
| RETURN                 | 2,9 [17]      | 1,8                     | MS and BP | Before [18] | Yes |
| SAVE                   | 2             | 1                       | MS        | Before [19] | Yes |
| SAVED <sup>P</sup>     | 2             | Yes                     | MS        |             | Yes |
| SDIV                   | 39            | 38                      | MS        | After       |     |
| SDIV{cc} <sup>D</sup>  | 40,39 [2]     | 39                      | MS        | After       |     |
| SDIVX                  | 71            | 70                      | MS        | After       |     |
| SETHI                  | 1             |                         | A0 or A1  |             |     |
| SHUTDOWN               | [23]          | NOP                     | MS        | NOP         |     |
| SIAM                   |               | Yes                     | MS        |             | Yes |
| SIR                    |               | Yes                     | BP and MS |             | Yes |
| SLL{X}                 | 1             |                         | A0 or A1  |             |     |
| SMUL <sup>D</sup>      | 6-7           | 5-6                     | MS        | After       |     |
| SMULcc <sup>D</sup>    | 7-8, -6-7 [2] | 6-8                     | MS        | After       |     |
| SRA{X}                 | 1             |                         | A0 or A1  |             |     |
| SRL{X}                 | 1             |                         | A0 or A1  |             |     |
| STB{A}                 |               |                         | MS        |             |     |
| STBAR <sup>D</sup>     | [20]          |                         | MS        |             | Yes |
| STD{A} <sup>D</sup>    |               | 2                       | MS        |             | Yes |
| STDF{A}                |               |                         | MS        |             |     |
| STF{A}                 |               |                         | MS        |             |     |
| STFSR <sup>P</sup>     |               | 9                       | MS        | Before [21] | Yes |
| ST(H,W,X){A}           |               |                         | MS        |             |     |
| STXF SR                |               | 9                       | MS        | Before [21] | Yes |
| SUB                    | 1             |                         | A0 or A1  |             |     |
| SUBcc                  | 1,0 [1]       |                         | A0 or A1  |             |     |
| SUBC                   | 5             | 4                       | MS        |             | Yes |
| SUBCcc                 | 6,5 [2]       | 5                       | MS        |             | Yes |
| SWAP{A}                | 31+           | 30+                     | MS        | After       |     |



**TABLE 4-4** UltraSPARC III Cu Instruction Latencies and Dispatching Properties (6 of 6)

| Instruction         | Latency      | Dispatch Blocking After | Pipeline  | Break | SIG |
|---------------------|--------------|-------------------------|-----------|-------|-----|
| TADDcc              | 5            | Yes                     | MS        |       | Yes |
| TSUBcc              | 5            | Yes                     | MS        |       | Yes |
| Tcc                 |              |                         | BR and MS |       |     |
| UDIV <sup>D</sup>   | 40           | 39                      | MS        | After |     |
| UDIVcc <sup>D</sup> | 41,40 [2]    | 40                      | MS        | After |     |
| UDIVX               | 71           | 70                      | MS        | After |     |
| UMUL <sup>D</sup>   | 6-8          | 5-7                     | MS        | After |     |
| UMULcc <sup>D</sup> | 7-8, 6-7 [2] | 6-8                     | MS        | After |     |
| WRASI               |              | 16                      | BR and MS |       | Yes |
| WRASR               |              | 7                       | BR and MS |       | Yes |
| WRCCR               |              | 7                       | BR and MS |       | Yes |
| WRFPRS              |              | 7                       | BR and MS |       | Yes |
| WRPR <sup>P</sup>   |              | 7                       | BR and MS |       | Yes |
| WRY <sup>D</sup>    |              | 7                       | BR and MS |       | Yes |
| XNOR                | 1            |                         | A0 or A1  |       |     |
| XNORcc              | 1,0 [1]      |                         | A0 or A1  |       |     |
| XOR                 | 1            |                         | A0 or A1  |       |     |
| XORcc               | 1,0 [1]      |                         | A0 or A1  |       |     |

1. These operations produce two results: destination register and condition code (%icc, %xcc). The latency is one in the former case and zero in the latter case. For example, SUBcc and BE %icc are grouped together (zero latency).
2. These operations produce two results: destination register and condition code (%icc, %xcc). The latency is given as a pair of numbers —*m,n*— for the register and condition code, respectively. When latencies vary in a range, such as in UMULcc, this range is indicated by pair – pair.
3. Latency is *x,y* for correct, incorrect branch prediction. It is measured as the difference in the dispatching cycle of the retiring target instruction and that of the branch.
4. Blocking cycles are *x,y* for correct, incorrect branch prediction. They are measured as the difference in the dispatching cycle of instruction in the delay slot (or target, if annulled) that retires and that of the branch.
5. Native Call and Link with immediate target address (label).
6. Latency in parentheses applies when operands involve IEEE special values (NaN, INF), including zero and illegal values.
7. Blocking is limited to another FD operation in succession; otherwise, it is unblocking. Blocking cycles in parentheses apply when operands involve special and illegal values.
8. Dispatching stall (7+ cycles) until all stores in flight retire.
9. 0–4 if predicted true; 9–10 if mispredicted.

10. Latency is taken to be the difference in dispatching cycles from `jmp1` to target operation, including the effect of an operation in the delay slot. Blocking cycles thus may include cycles due to restore in the delay slot. In a given pair `x,y`, `x` applies when predicted correctly and `y` when predicted incorrectly. Each `x` or `y` may be a range of values.
11. 0–4 if predicted true; 9 if mispredicted.
12. This `MEMBAR` has NOP semantics, since the ordering specified is implicitly done by processor (memory model is TSO).
13. All operations in flight complete as in `MEMBAR #Sync`.
14. All operations in flight complete.
15. Issue stalls a minimum of 7 cycles until all operations in flight are done (get to D-stage).
16. Dispatching stalls until previous save in flight, if any, reaches D-stage.
17. 2 if predicted correctly, 9 otherwise. Similarly for blocking cycles.
18. Dispatching stalls until previous restore in flight, if any, reaches D-stage.
19. Dispatching stall until previous restore in flight, if any, reaches D-stage.
20. Same as `MEMBAR #StoreStore`, which is NOP.
21. Dispatching stalls until all FP operations in flight are done.
22. Wait for completion of all FP operations in flight.
23. The Shutdown instruction is not implemented. The instruction is neutralized and appears as a NOP to software (no visible effects).

## SECTION III

---

---

# Execution Environment

---

---



## Data Formats

---

The processor recognizes the following fundamental data types:

- Signed integer: 8, 16, 32, and 64 bits
- Unsigned integer: 8, 16, 32, and 64 bits
- VIS Instruction data formats: pixel (32 bits), fixed16 (64 bits), and fixed32 (64 bits)
- Floating point: 32, 64, and 128 bits

The widths of the data types are as follows:

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

The signed integer values are stored as two's-complement numbers with a width commensurate with their range. The floating-point formats conform to the IEEE Standard for Binary Floating-point Arithmetic, IEEE Std 754-1985. In tagged words, the least significant two bits are treated as a tag; the remaining 30 bits are treated as a signed integer.

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

- Signed Integer Double
- Unsigned Integer Double
- Floating-Point, Double-Precision
- Floating-Point, Quad-Precision

---

## 5.1 Integer Data Formats

The processor supports the following integer data formats:

- Signed integer
- Unsigned integer
- Tagged integer word

### 5.1.1 Integer Data Value Range

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

**TABLE 5-1** Signed Integer, Unsigned Integer, and Tagged Integer Format Ranges

| Data Type                      | Width (bits) | Range     |              |
|--------------------------------|--------------|-----------|--------------|
|                                |              | Lower     | Upper        |
| Signed integer byte            | 8            | $-2^7$    | $2^7 - 1$    |
| Signed integer halfword        | 16           | $-2^{15}$ | $2^{15} - 1$ |
| Signed integer word            | 32           | $-2^{31}$ | $2^{31} - 1$ |
| Signed integer tagged word     | 32           | $-2^{29}$ | $2^{29} - 1$ |
| Signed integer double word     | 64           | $-2^{63}$ | $2^{63} - 1$ |
| Signed extended integer        | 64           | $-2^{63}$ | $2^{63} - 1$ |
| Unsigned integer byte          | 8            | 0         | $2^8 - 1$    |
| Unsigned integer halfword      | 16           | 0         | $2^{16} - 1$ |
| Unsigned integer word          | 32           | 0         | $2^{32} - 1$ |
| Unsigned integer tagged word   | 32           | 0         | $2^{30} - 1$ |
| Unsigned integer double word   | 64           | 0         | $2^{64} - 1$ |
| Unsigned extended integer      | 64           | 0         | $2^{64} - 1$ |
| (Unsigned) tagged integer word | 32           | 0         | $2^{30} - 1$ |

## 5.1.2 Integer Data Alignment

TABLE 5-2 describes the memory and register alignment for integer data.

**TABLE 5-2** Integer Data Alignment

| Subformat Type | Width                   | Subformat Field              | Required Address Alignment | Memory Address (Big-endian) | Register Number Alignment | Register Number |
|----------------|-------------------------|------------------------------|----------------------------|-----------------------------|---------------------------|-----------------|
| SB             | <i>B (byte)</i>         | signed_byte_integer<7:0>     | None                       | <i>n</i>                    | any                       | <i>r</i>        |
| UB             |                         | unsigned_byte_integer<7:0>   |                            |                             |                           |                 |
| SH             | <i>H (halfword)</i>     | signed_halfwd_integer<7:0>   | 0 mod 2                    | <i>n</i>                    | any                       | <i>r</i>        |
| UH             |                         | unsigned_halfwd_integer<7:0> |                            |                             |                           |                 |
| SW             | <i>W (word)</i>         | signed_word_integer<7:0>     | 0 mod 4                    | <i>n</i>                    | any                       | <i>r</i>        |
| UW             |                         | unsigned_word_integer<7:0>   |                            |                             |                           |                 |
| SD-0           | <i>D (double word)</i>  | signed_dbl_integer<63:32>    | 0 mod 8                    | <i>n</i>                    | 0 mod 2                   | <i>r</i>        |
| UD-0           |                         | unsigned_dbl_integer<63:32>  |                            |                             |                           |                 |
| SD-1           |                         | signed_dbl_integer<31:0>     | 4 mod 8                    | <i>n + 4</i>                | 1 mod 2                   | <i>r + 1</i>    |
| UD-1           |                         | unsigned_dbl_integer<31:0>   |                            |                             |                           |                 |
| SX             | <i>X (extendedword)</i> | signed_ext_integer<63:0>     | 0 mod 8                    | <i>n</i>                    | —                         | <i>r</i>        |
| UX             |                         | unsigned_ext_integer<63:0>   |                            |                             |                           |                 |

The data types are illustrated in the following subsections.


## 5.1.3 Signed Integer Data Types

Figures in this section illustrate the following signed data types:

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

### 5.1.3.1 Signed Integer Byte


FIGURE 5-1 illustrates the signed integer byte data format.



**FIGURE 5-1** Signed Integer Byte Data Format

### 5.1.3.2 Signed Integer Halfword


FIGURE 5-2 illustrates the signed integer halfword data format.



**FIGURE 5-2** Signed Integer Halfword Data Format

### 5.1.3.3 Signed Integer Word


FIGURE 5-3 illustrates the signed integer word data format.



**FIGURE 5-3** Signed Integer Word Data Format

### 5.1.3.4 Signed Integer Double

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




**FIGURE 5-4** Signed Integer Double Data Format



### 5.1.3.5 Signed Extended Integer

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



**FIGURE 5-5** Signed Extended Integer Data Format

## 5.1.4 Unsigned Integer Data Types

Figures in this section illustrate the following unsigned data types:

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

### 5.1.4.1 Unsigned Integer Byte


FIGURE 5-6 illustrates the unsigned integer byte data format.



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

### 5.1.4.2 Unsigned Integer Halfword


FIGURE 5-7 illustrates the unsigned integer halfword data format.



**FIGURE 5-7** Unsigned Integer Halfword Data Format

### 5.1.4.3 Unsigned Integer Word


FIGURE 5-8 illustrates the unsigned integer word data format.



**FIGURE 5-8** Unsigned Integer Word Data Format

### 5.1.4.4 Unsigned Integer Double


FIGURE 5-9 illustrates both components (UD-0 and UD-1) of the unsigned integer double data format.



**FIGURE 5-9** Unsigned Integer Double Data Format

### 5.1.4.5 Unsigned Extended Integer

FIGURE 5-10 illustrates the unsigned extended integer (UX) data format.




**FIGURE 5-10** Unsigned Extended Integer Data Format

### 5.1.5 Tagged Word

The Tagged word data format is similar to the unsigned word format except for a 2-bit field in the two least significant bit (LSB) positions. Bit 31 is the overflow bit.

FIGURE 5-11 illustrates the tagged word data format.



**FIGURE 5-11** Tagged Word Data Format

## 5.2 Floating-Point Data Formats

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

- Single-precision floating-point (32-bit)
- Double-precision floating-point (64-bit)
- Quad-precision floating-point (128-bit)

### 5.2.1 Floating-Point Data Value Range

The value range for each format is included with the format and description of each format.

### 5.2.2 Floating-Point Data Alignment

TABLE 5-3 describes the address and memory alignment for floating-point data.

**TABLE 5-3** Floating-Point Doubleword and Quadword Alignment

| Subformat Name | Subformat Field              | Required Address Alignment | Memory Address (Big-endian)* | Register Number Alignment | Available Registers     |
|----------------|------------------------------|----------------------------|------------------------------|---------------------------|-------------------------|
| FS             | s:exp<7:0>;fraction<22:0>    | 0 mod 4 †                  | <i>n</i>                     | any                       | <i>f0, f1, ..., f31</i> |
| FD-0           | s:exp<10:0>;fraction<51:32>  | 0 mod 4 †                  | <i>n</i>                     | 0 mod 2                   | <i>f0, f2, ..., f62</i> |
| FD-1           | fraction<31:0>               | 0 mod 4 †                  | <i>n</i> + 4                 | 1 mod 2                   | <i>f1, f3, ..., f63</i> |
| FX-0           |                              | 0 mod 4 †                  | <i>n</i>                     | 0 mod 4                   | <i>f0, f4, ..., f60</i> |
| FX-1           |                              | 0 mod 4 †                  | <i>n</i>                     | 0 mod 4                   | <i>f2, f6, ..., f62</i> |
| FQ-0           | s:exp<14:0>;fraction<111:96> | 0 mod 4 ‡                  | <i>n</i>                     | 0 mod 4                   | <i>f0, f4, ..., f60</i> |
| FQ-1           | fraction<95:64>              | 0 mod 4 ‡                  | <i>n</i> + 4                 | 1 mod 4                   | <i>f1, f5, ..., f61</i> |
| FQ-2           | fraction<63:32>              | 0 mod 4 ‡                  | <i>n</i> + 8                 | 2 mod 4                   | <i>f2, f6, ..., f62</i> |
| FQ-3           | fraction<31:0>               | 0 mod 4 ‡                  | <i>n</i> + 12                | 3 mod 4                   |                         |
| FX             |                              | 0 mod 4 †                  | <i>n</i>                     | 0 mod 4                   | <i>f3, f7, ..., f63</i> |


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

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

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

## 5.2.3 Floating-Point, Single-Precision

FIGURE 5-12 illustrates the floating-point single-precision data format, and TABLE 5-4 describes the formats.



**FIGURE 5-12** Floating-Point Single-Precision Data Format


**TABLE 5-4** Floating-Point Single-Precision Format Definition

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

## 5.2.4 Floating-Point, Double-Precision

FIGURE 5-13 illustrates both components (FD-0 and FD-1) of the floating-point double-precision data format when two 32-bit registers are used. FIGURE 5-14 illustrates a double-precision data format using one 64-bit register.

TABLE 5-5 describes the data formats.



**FIGURE 5-13** Floating-Point Double-Precision Double Word Data Format



**FIGURE 5-14** Floating-Point Double-Precision Extended Word Data Format

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

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


## 5.2.5 Floating-Point, Quad-Precision

FIGURE 5-15 illustrates all four components (FQ-0 through FQ-3) of the floating-point quad-precision data format, and TABLE 5-6 describes the formats.

---

**Compatibility Note** – Floating-point quad is not implemented in the processor. Quad-precision operations are emulated in the OS kernel.

---



**FIGURE 5-15** Floating-Point Quad-Precision Data Format

**TABLE 5-6** Floating-Point Quad-Precision Format Definition

|   |   |
|---|---|
| <i>s</i> = sign (1-bit)<br><i>e</i> = biased exponent (15 bits)<br><i>f</i> = fraction (112 bits)<br><i>u</i> = undefined |   |
| Normalized value ( $0 < e < 32767$ )  | $(-1)^s \times 2^{e-16383} \times 1.f$  |
| Subnormal value ( $e = 0$ )   | $(-1)^s \times 2^{-16382} \times 0.f$   |
| Zero ( $e = 0$ )  | $(-1)^s \times 0$   |
| Signalling NaN  | $s = u$ ; $e = 32767$ (max); $f = .0uu--uu$<br>(At least one bit of the fraction must be nonzero) |
| Quiet NaN   | $s = u$ ; $e = 32767$ (max); $f = .1uu--uu$   |
| $-\infty$ (negative infinity)   | $s = 1$ ; $e = 32767$ (max); $f = .000--00$   |
| $+\infty$ (positive infinity)   | $s = 0$ ; $e = 32767$ (max); $f = .000--00$   |

## 5.3 VIS Execution Unit Data Formats

VIS instructions are optimized for short integer arithmetic, where the overhead of converting to and from floating-point is significant. Data components can be 8 or 16 bits; intermediate results are 16 or 32 bits.

There are two VIS data formats:

- Pixel Data
- Fixed-point Data

## Data Conversions

Conversion from pixel data to fixed data occurs through pixel multiplications. Conversion from fixed data to pixel data is done with the pack instructions, which clip and truncate to an 8-bit unsigned value. Conversion from 32-bit fixed to 16-bit fixed is also supported with the FPACKFIX instruction.

## Rounding

Rounding can be performed by adding one to the round bit position. Complex calculations needing more dynamic range or precision should be performed using floating-point data.

## Range

The range of values that each format supports is described below.

## Data Alignment

The data in memory is expected to be aligned according to TABLE 5-7. If the address does not properly align, then an exception is generated and the load/store operation fails.

TABLE 5-7 Pixel, Fixed16, and Fixed32 Data Alignment

| VIS Data Format Type | Width | VIS Data Format Name | Required Address Alignment | Memory Address (Big-endian) | Register Number Alignment | Register Number |
|----------------------|-------|----------------------|----------------------------|-----------------------------|---------------------------|-----------------|
| Pixel 8              | 32    | Pixel Data Format    | $0 \bmod 4$                | $n$                         | $r$                       | $r$             |
| Fixed16              | 64    | Fixed16 Data Format  | $0 \bmod 8$                | $n$                         | $0 \bmod 2$               | $r$             |
| Fixed32              | 64    | Fixed32 Data Format  | $0 \bmod 8$                | $n$                         | $0 \bmod 2$               | $r$             |

### 5.3.1 Pixel Data Format

The Fixed 8-bit data format consists of four unsigned 8-bit integers contained in a 32-bit word.


One common use is to represent intensity values for the color components of an image. For example, R, G, B and  $\alpha$  are used as color components and are positioned as shown in FIGURE 5-16.





### 5.3.2.2 Fixed32 Data Format

The fixed 32-bit format consists of two 32-bit, signed, fixed-point values contained in a 64-bit word. FIGURE 5-18 illustrates the Fixed32 VIS data format.



**FIGURE 5-18** Fixed32 VIS Data Format



# Registers

---

This chapter discusses the following topics:

Section 6.1, “Introduction”

Section 6.2, “Integer Unit General-Purpose r Registers”

Section 6.3, “Register Window Management”

Section 6.4, “Floating-Point General-Purpose Registers”

Section 6.5, “Control and Status Register Summary”

Section 6.6, “State Registers”

Section 6.7, “Ancillary State Registers: ASRs 16-25”

Section 6.8, “Privileged Registers”

Section 6.9, “Special Access Register”

Section 6.10, “ASI Mapped Registers”

---

## 6.1 Introduction

The processor consists of many types of registers that serve various purposes and are accessed in many different ways.

There are separate working registers for the integer and floating-point units (FPUs). Both of these register sets have been expanded over the evolution of the SPARC processor. The integer unit registers are shadowed using windowing and selection methods. The registers in the floating-point register set (also used for VIS and block load store instructions) are combined in specific ways to support data sizes up to 128 bits. All integer registers and the upper floating-point registers are 64 bits wide.

The processor also has a vast array of control, status, state, and diagnostic registers that are used to setup, control, and operate the processor. The two main operating modes of the processor, privileged and non-privileged mode, have a profound effect on which of the control and status registers are available to the software.

The majority of the control and status registers are 64 bits wide and are accessed using the privileged register access instructions, state register access instructions, and load/store with ASI access instructions. For convenience, some registers in this chapter are illustrated as fewer than 64 bits wide. Any bits not shown are reserved for future extensions to the architecture. Such reserved bits are read as zeroes and, when written by software, should be written with the values of those bits previously read from that register or with zeroes.

- Integer Unit Working Registers (includes  $r$  and global)
- Floating-point Unit Working Registers
- Privileged Registers
- State and Ancillary State Registers (includes ASRs)
- Floating-point Status Register (FSR)
- ASI Mapped Registers (CSRs)

Some of the figures and tables in this chapter are reproduced from *The SPARC Architecture Manual, Version 9* and other sources. Many new diagrams and tables appear for the first time. Contents of this chapter applies to non-privileged mode unless stated otherwise.

---

## 6.2 Integer Unit General-Purpose $r$ Registers

An UltraSPARC III Cu processor contains 160 general-purpose 64-bit  $r$  registers. They are windowed into 32 registers addressable by Integer Unit Instructions.

The  $r$  registers are partitioned into eight addressable *global* registers and 24 addressable *windowed* registers. There are four *global* register sets: *normal*, *MMU*, *Interrupt*, and *Alternate*. The windowed registers point to eight working register sets that are windowed into  $r[8]$  to  $r[31]$ , as one full register set (eight *locals* and eight *ins*) and a half register set (eight *outs*) belonging to the next higher state.

In summary, the  $r$  registers consist of eight *in* registers, eight *local* registers, eight *out* registers, and the selected eight *global* registers.

The current window pointer (CWP) register selects the *in/local/out* windowed registers. SAVE and RESTORE instructions modify the CWP register.

The PSTATE.AG, .IG, and .MG fields select the *global* register set. Processor exceptions modify the PSTATE register fields to select the *global* register set.

PSTATE and CWP registers are accessible using privileged instructions.

At any moment, general-purpose registers appear in non-privileged mode as shown in TABLE 6-1.

**TABLE 6-1** Integer Unit General-Purpose Registers

| <b>Windowed Register Name</b> | <b>r Register Address</b> | <b>Source</b>                  | <b>Comments</b> |
|-------------------------------|---------------------------|--------------------------------|-----------------|
| <i>in</i> [7]                 | r[31]                     | Current Register Set           |                 |
| <i>in</i> [6]                 | r[30]                     | Current Register Set           |                 |
| <i>in</i> [5]                 | r[29]                     | Current Register Set           |                 |
| <i>in</i> [4]                 | r[28]                     | Current Register Set           |                 |
| <i>in</i> [3]                 | r[27]                     | Current Register Set           |                 |
| <i>in</i> [2]                 | r[26]                     | Current Register Set           |                 |
| <i>in</i> [1]                 | r[25]                     | Current Register Set           |                 |
| <i>in</i> [0]                 | r[24]                     | Current Register Set           |                 |
| <i>local</i> [7]              | r[23]                     | Current Register Set           |                 |
| <i>local</i> [6]              | r[22]                     | Current Register Set           |                 |
| <i>local</i> [5]              | r[21]                     | Current Register Set           |                 |
| <i>local</i> [4]              | r[20]                     | Current Register Set           |                 |
| <i>local</i> [3]              | r[19]                     | Current Register Set           |                 |
| <i>local</i> [2]              | r[18]                     | Current Register Set           |                 |
| <i>local</i> [1]              | r[17]                     | Current Register Set           |                 |
| <i>local</i> [0]              | r[16]                     | Current Register Set           |                 |
| <i>out</i> [7]                | r[15]                     | Next higher level Register Set | See footnote 1  |
| <i>out</i> [6]                | r[14]                     | Next higher level Register Set |                 |
| <i>out</i> [5]                | r[13]                     | Next higher level Register Set |                 |
| <i>out</i> [4]                | r[12]                     | Next higher level Register Set |                 |
| <i>out</i> [3]                | r[11]                     | Next higher level Register Set |                 |
| <i>out</i> [2]                | r[10]                     | Next higher level Register Set |                 |
| <i>out</i> [1]                | r[ 9]                     | Next higher level Register Set |                 |
| <i>out</i> [0]                | r[ 8]                     | Next higher level Register Set |                 |
| <i>global</i> [7]             | r[ 7]                     | Global[ 7]                     |                 |
| <i>global</i> [6]             | r[ 6]                     | Global[ 6]                     |                 |
| <i>global</i> [5]             | r[ 5]                     | Global[ 5]                     |                 |
| <i>global</i> [4]             | r[ 4]                     | Global[ 4]                     |                 |

**TABLE 6-1** Integer Unit General-Purpose Registers (*Continued*)

| Windowed Register Name | $\mathbf{r}$ Register Address | Source     | Comments                               |
|------------------------|-------------------------------|------------|--|
| <i>global</i> [3]      | $\mathbf{r}[ 3]$              | Global[ 3] |  |
| <i>global</i> [2]      | $\mathbf{r}[ 2]$              | Global[ 2] |  |
| <i>global</i> [1]      | $\mathbf{r}[ 1]$              | Global[ 1] |  |
| <i>global</i> [0]      | $\mathbf{r}[ 0]$              | Global[ 0] | Value ( $\mathbf{r}[ 0]$ ) always zero |

1. The CALL instruction writes its own address into the  $\mathbf{r}[15]$  register (*out*[7]).

## 6.2.1 Windowed (*in/local/out*) $\mathbf{r}$ Registers

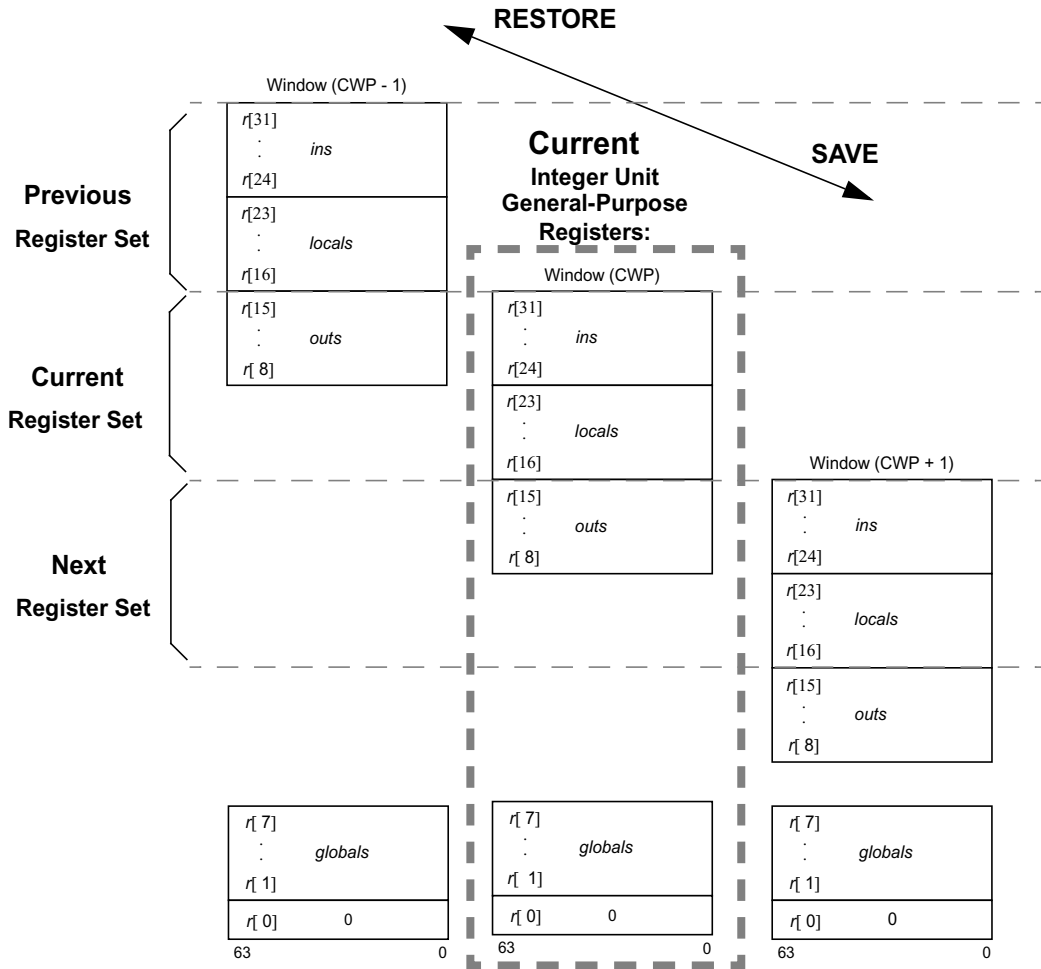
At any time, an integer unit instruction can access a 24-register *window* into the register sets. A register window comprises the eight *in* and eight *local* registers (a complete register set) together with the eight *in* registers (upper half of the next higher register set). The CALL instruction writes its own address into register  $\mathbf{r}[15]$  (*out* register 7).

## 6.2.2 Global $\mathbf{r}$ Register Sets

Registers  $\mathbf{r}[0] - \mathbf{r}[7]$  refer to a set of eight global registers ( $\mathbf{g}0 - \mathbf{g}7$ ). At any time, one of four sets of eight global register sets is selected and can be accessed as the current global register set. The currently enabled set of global registers is selected by the Alternate Global (AG), Interrupt Global (IG), and MMU Global (MG) fields in the PSTATE register. See “Processor State (PSTATE) Privileged Register 6” on page 112 for a description of the AG, IG, and MG fields.

Global register zero ( $\mathbf{g}0$ ) always reads as zero; writes to it have no program-visible effect.

An illustration showing the current IU registers is shown in FIGURE 6-1.



**FIGURE 6-1** Three Overlapping Windows and the Eight Global Registers

**Compatibility Note** – Since the `PSTATE` register is writable only by privileged software, existing non-privileged SPARC V8 software operates correctly on a processor if Supervisor Software ensures that User Software sees a consistent set of global registers.

In summary, the processor has eight windows or register sets ( $NWINDOWS = 8$ ). The total number of  $r$  registers in the processor is 160: 8 *normal* global registers, 8 *alternate* global registers, 8 *interrupt* global registers, 8 *MMU* global registers, plus the number of register sets (eight) times 16 registers/set.

### 6.2.2.1 Overlapping Windows

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

An *outs* register with address  $o$ , where  $8 \leq o \leq 15$ , refers to exactly the same register as  $(o+16)$  does after the  $CWP$  is incremented by one (modulo  $NWINDOWS$ ). Likewise, an *in* register with address  $i$ , where  $24 \leq i \leq 31$ , refers to exactly the same register as address  $(i - 16)$  does after the  $CWP$  is decremented by one (modulo  $NWINDOWS$ ). See FIGURE 6-1 and FIGURE 6-2.

Since  $CWP$  arithmetic is performed modulo  $NWINDOWS$ , the highest-numbered implemented window (window 7) overlaps with window 0. The *outs* of window  $NWINDOWS - 1$  are the *ins* of window 0. Implemented windows are numbered contiguously from 0 through  $NWINDOWS - 1$ .

### 6.2.3 128-bit Operand Considerations

LDD, LDDA, STD, and STDA instructions access 128-bit data associated with adjacent  $r$  registers and require even-odd register alignment. An attempt to execute an LDD, LDDA, STD, or STDA instruction that refers to a misaligned (odd) destination register number causes an *illegal\_instruction* trap.

---

## 6.3 Register Window Management

---

**Note** – Register window management is the responsibility of the operating system (supervisor code). The user code sees an unlimited stack of register windows and does not have to worry about register window management. The operating system provides support for underflow and overflow of the stack. This mechanism is transparent to the user code.

---

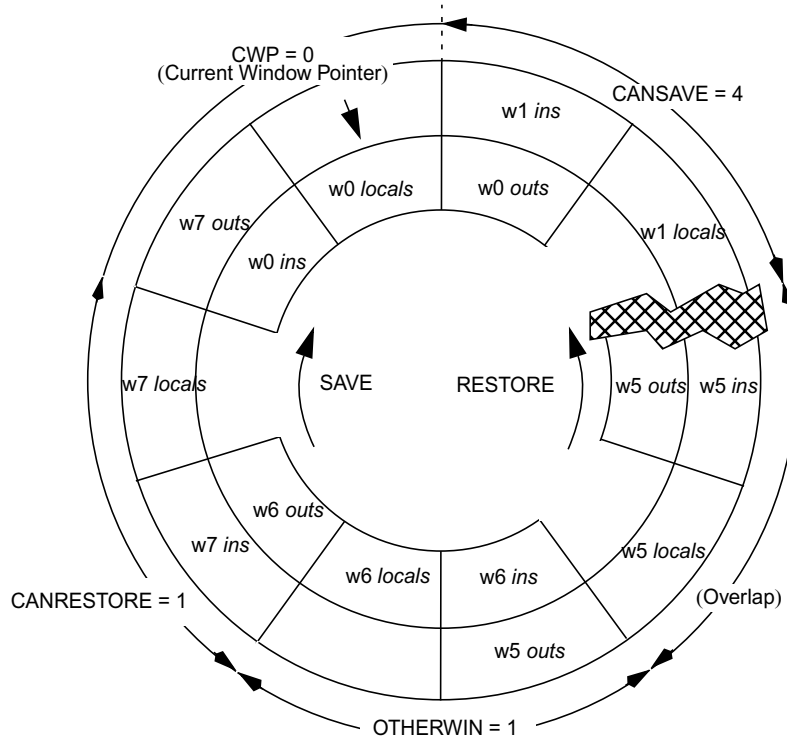
The current window in the windowed portion of  $r$  registers is given by the  $CWP$  register. The  $CWP$  is decremented by the `RESTORE` instruction and incremented by the `SAVE` instruction. Window overflow is detected by the `CANSAVE` register, and window underflow is detected by the `CANRESTORE` register, both of which are controlled by privileged software. A window overflow (underflow) condition causes a window spill (fill) trap.



---

**Programming Note** – Because the windows overlap, the number of windows available to software is one less than the number of implemented windows, that is, 7 ( $NWINDOWS - 1$ ).

---



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

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

$NWINDOWS = 8$ ,  $CWP = 0$ ,  $CANSERVE = 4$ ,  $OTHERWIN = 1$ , and  $CANRESTORE = 1$ . If the procedure using window  $w0$  executes a *RESTORE*, then window  $w7$  becomes the current window. If the procedure using window  $w0$  executes a *SAVE*, then window  $w1$  becomes the current window.

**FIGURE 6-2** Windowed r Registers for  $NWINDOWS = 8$

## 6.3.1 CALL and JMPL Instructions

---

**Programming Note** – Since the procedure call instructions (*CALL* and *JMPL*) do not change the *CWP*, a procedure can be called without changing the window.

---

## 6.3.2 Circular Windowing

---

**Programming Note** – When the register file is full, the *outs* of the newest window are the *ins* of the oldest window, which still contains valid data.

---

## 6.3.3 Clean Window with RESTORE and SAVE Instructions

---

**Programming Note** – The *local* and *out* registers of a register window are guaranteed to contain either zeroes or an old value that belongs to the current context upon reentering the window through a *SAVE* instruction. If a program executes a *RESTORE* followed by a *SAVE*, then the resulting window's *locals* and *outs* may not be valid after the *SAVE*, since a trap may have occurred between the *RESTORE* and the *SAVE*.

---

---

# 6.4 Floating-Point General-Purpose Registers

The Floating-point register file contains addressable registers for the following:

- Floating-point Instructions
- VIS Instructions
- Block load and store instructions
- FSR load and store instructions

The registers have various widths and assigned addresses as follows:

- 32 32-bit (single-precision) floating-point registers,  $f[0]$ ,  $f[1]$ , ...  $f[31]$
- 32 64-bit (double-precision) floating-point registers,  $f[0]$ ,  $f[2]$ , ...  $f[62]$
- 16 128-bit (quad-precision) floating-point registers,  $f[0]$ ,  $f[4]$ , ...  $f[60]$

The floating-point registers are arranged so that some of them overlap, that is, are aliased. The layout and numbering of the floating-point registers are shown in TABLE 6-2, TABLE 6-3, and TABLE 6-4. Unlike the windowed  $x$  registers, all of the floating-point registers are accessible at any time. The floating-point registers can be read and written by FPop (FPop1/FPop2 format) instructions, load/store single/double/quad floating-point instructions, and block load and block store instructions.

**TABLE 6-2** 32-bit Floating-Point Registers with Aliasing

| Operand Register and Field | From Register | Operand Register and Field | From Register |
|----------------------------|---------------|----------------------------|---------------|
| f31                        | <31:0>        | f15                        | <31:0>        |
| f30                        | <31:0>        | f14                        | <31:0>        |
| f29                        | <31:0>        | f13                        | <31:0>        |
| f28                        | <31:0>        | f12                        | <31:0>        |
| f27                        | <31:0>        | f11                        | <31:0>        |
| f26                        | <31:0>        | f10                        | <31:0>        |
| f25                        | <31:0>        | f9                         | <31:0>        |
| f24                        | <31:0>        | f8                         | <31:0>        |
| f23                        | <31:0>        | f7                         | <31:0>        |
| f22                        | <31:0>        | f6                         | <31:0>        |
| f21                        | <31:0>        | f5                         | <31:0>        |
| f20                        | <31:0>        | f4                         | <31:0>        |
| f19                        | <31:0>        | f3                         | <31:0>        |
| f18                        | <31:0>        | f2                         | <31:0>        |
| f17                        | <31:0>        | f1                         | <31:0>        |
| f16                        | <31:0>        | f0                         | <31:0>        |

**TABLE 6-3** 64-bit Floating-Point Registers with Aliasing

| Operand Register and Field | From Register | Operand Register and Field | From Register |
|----------------------------|---------------|----------------------------|---------------|
| f62                        | <63:0>        | f30                        | <63:0>        |
| f60                        | <63:0>        | f28                        | <63:0>        |
| f58                        | <63:0>        | f26                        | <63:0>        |
| f56                        | <63:0>        | f24                        | <63:0>        |
| f54                        | <63:0>        | f22                        | <63:0>        |
| f52                        | <63:0>        | f20                        | <63:0>        |
| f50                        | <63:0>        | f18                        | <63:0>        |
| f48                        | <63:0>        | f16                        | <63:0>        |
| f46                        | <63:0>        | f14                        | <63:0>        |
| f44                        | <63:0>        | f12                        | <63:0>        |
| f42                        | <63:0>        | f10                        | <63:0>        |
| f40                        | <63:0>        | f8                         | <63:0>        |
| f38                        | <63:0>        | f6                         | <63:0>        |
| f36                        | <63:0>        | f4                         | <63:0>        |
| f34                        | <63:0>        | f2                         | <63:0>        |
| f32                        | <63:0>        | f0                         | <63:0>        |

**TABLE 6-4** 128-bit Floating-Point Registers with Aliasing

| Operand Register and Field | From Register   |
|----------------------------|---|
| f60                        | <127:0> f 6 0<63:0>:f 6 2<63:0>                         |
| f56                        | <127:0> f 5 6<63:0>:f 5 8<63:0>                         |
| f52                        | <127:0> f 5 2<63:0>:f 5 4<63:0>                         |
| f48                        | <127:0> f 4 8<63:0>:f 5 0<63:0>                         |
| f44                        | <127:0> f 4 4<63:0>:f 4 6<63:0>                         |
| f40                        | <127:0> f 4 0<63:0>:f 4 2<63:0>                         |
| f36                        | <127:0> f 3 6<63:0>:f 3 8<63:0>                         |
| f32                        | <127:0> f 3 2<63:0>:f 3 4<63:0>                         |
| f28                        | <127:0> f 2 8<31:0>:f 2 9<31:0>:f 3 0<31:0>:f 3 1<31:0> |
| f24                        | <127:0> f 2 4<31:0>:f 2 5<31:0>:f 2 6<31:0>:f 2 7<31:0> |
| f20                        | <127:0> f 2 0<31:0>:f 2 1<31:0>:f 2 2<31:0>:f 2 3<31:0> |
| f16                        | <127:0> f 1 6<31:0>:f 1 7<31:0>:f 1 8<31:0>:f 1 9<31:0> |
| f12                        | <127:0> f 1 2<31:0>:f 1 3<31:0>:f 1 4<31:0>:f 1 5<31:0> |
| f8                         | <127:0> f 8<31:0>:f 9<31:0>:f 1 0<31:0>:f 1 1<31:0>     |
| f4                         | <127:0> f 4<31:0>:f 5<31:0>:f 6<31:0>:f 7<31:0>         |
| f0                         | <127:0> f 0<31:0>:f 1<31:0>:f 2<31:0>:f 3<31:0>         |

## 6.4.1 Floating-Point Register Number Encoding

The floating-point register number encoding in the instruction field depends on the width of register being addressed. The encoding for the 5-bit instruction field (labeled b<4>-b<0>, where b<4> is the most significant bit of the register number), is given in TABLE 6-5.

**TABLE 6-5** Floating-Point Register Number Encoding

| Register Operand Type | 6-bit Register Number, fn |      |      |      |      |      | Encoding in a 5-bit Register Field in an Instruction, rd/rs |      |      |      |      |
|-----------------------|---------------------------|------|------|------|------|------|---|------|------|------|------|
|                       | 0                         | b<4> | b<3> | b<2> | b<1> | b<0> | b<4>  | b<3> | b<2> | b<1> | b<0> |
| 32-bit (single)       | 0                         | b<4> | b<3> | b<2> | b<1> | b<0> | b<4>  | b<3> | b<2> | b<1> | b<0> |
| 64-bit (double)       | b<5>                      | b<4> | b<3> | b<2> | b<1> | 0    | b<4>  | b<3> | b<2> | b<1> | b<5> |
| 128-bit (quad)        | b<5>                      | b<4> | b<3> | b<2> | 0    | 0    | b<4>  | b<3> | b<2> | 0    | b<5> |

---

**Compatibility Note** – In the SPARC V8 architecture, bit 0 of 64- and 128-bit register numbers encoded in instruction fields was required to be zero. Therefore, all SPARC V8 floating-point instructions can run unchanged on a UltraSPARC III Cu processor, using the encoding in TABLE 6-5.

---

## 6.4.2 Double and Quad Floating-Point Operands

A 32-bit `f` register can hold one single-precision operand; a 64-bit (double-precision) operand requires an aligned pair of `f` registers, and a 128-bit (quad-precision) operand requires an aligned quadruple of `f` registers. At a given time, the floating-point registers can hold a maximum of 32 single-precision, 16 double-precision, or 8 quad-precision values in the lower half of the floating-point register file, plus an additional 16 double-precision or 8 quad-precision values in the upper half, or mixtures of the three sizes.

See FIGURE 6-3, TABLE 6-2, TABLE 6-3, and TABLE 6-4 for illustrative formats.

---

**Programming Note** – Data to be loaded into a floating-point double or quad register that is not doubleword aligned in memory must be loaded into the lower 16 double registers (8 quad registers) by means of single-precision LDF instructions. If desired, the data can then be copied into the upper 16 double registers (8 quad registers).

---

---

**Programming Note** – An attempt to execute an instruction that refers to a misaligned floating-point register operand (that is, a quad-precision operand in a register whose 6-bit register number is not `0 mod 4`) shall cause a `fp_exception_other` trap, with `FSR.fctt = 6` (`invalid_fp_register`).

---

---

**Programming Note** – Given the encoding in TABLE 6-5, it is impossible to specify a double-precision register with a misaligned register number.

---

---

**Note** – The processor does not implement quad-precision operations in hardware. All floating-point quad (including load and store) operations trap to the OS kernel and are emulated. Since the processor does not implement quad floating-point arithmetic operations in hardware, the `fp_exception_other` trap with `FSR.fctt = 6` (`invalid_fp_register`) does not occur in processors.

---

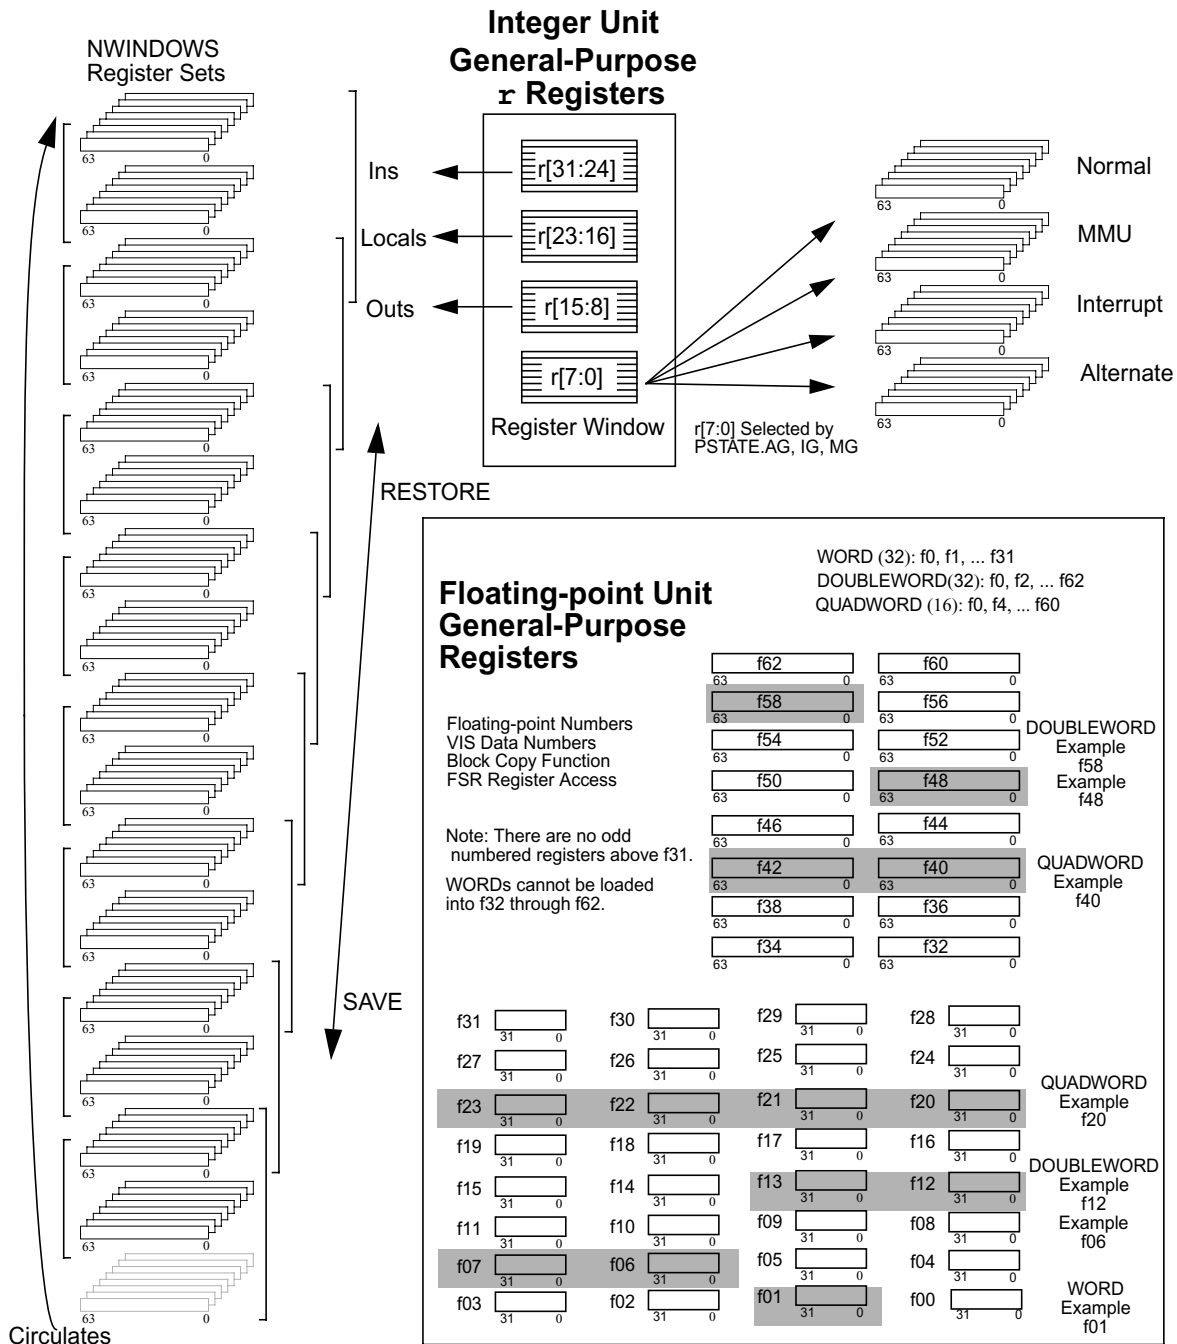


FIGURE 6-3 Integer Unit r Registers and Floating-Point Unit Working Registers

## 6.5 Control and Status Register Summary

This section presents a summary of control and status registers.

### 6.5.1 State and Ancillary State Register Summary

See FIGURE 6-4 and TABLE 6-6 for more information on state and ancillary state registers (ASRs).

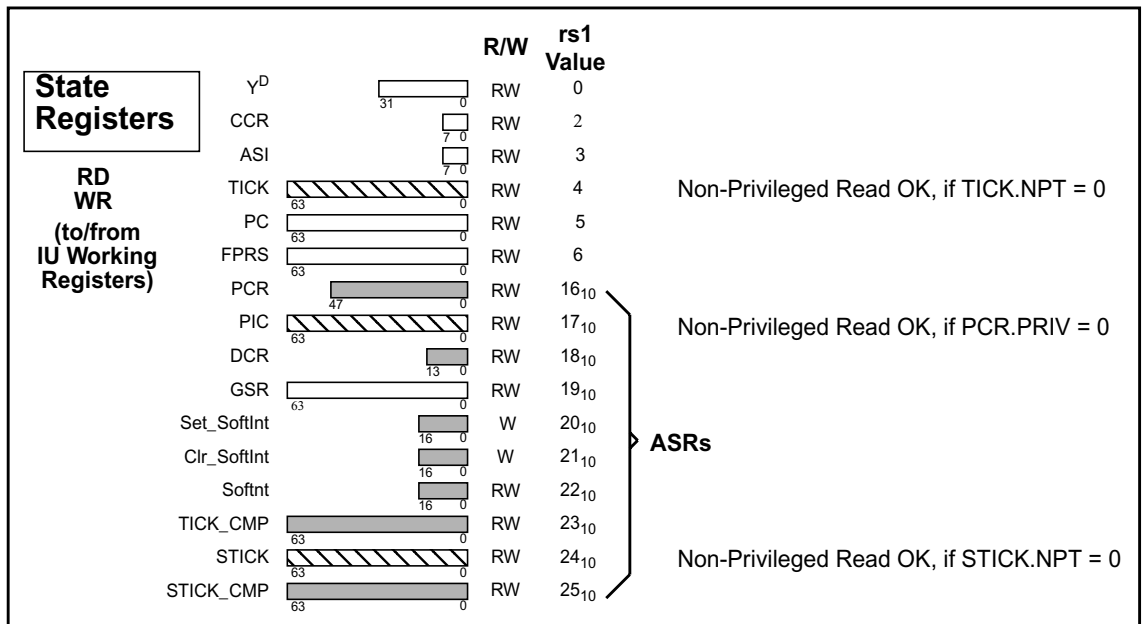


FIGURE 6-4 State and Ancillary State Registers

**TABLE 6-6** State and Ancillary State Registers

| State Register Number (base 10 used) | Access Restriction | R/W | Abbreviation            | Description   | Reference Section                         | Notes |
|--------------------------------------|--------------------|-----|-------------------------|---|---|-------|
| 0                                    | None               | RW  | Y <sup>D</sup> Register | 32-bit Multiply/Divide (deprecated)                                   |   |       |
| 1                                    |                    |     | Reserved                |   |   |       |
| 2                                    | None               | RW  | CCR                     | Condition Code  |   |       |
| 3                                    | None               | RW  | ASI                     | Address Space Identifier  | Section 6.6.3                             |       |
| 4                                    | Depends            | R   | TICK                    | TICK register for CPU Timer, also accessible as a privileged register | Section 6.7.4                             | 1     |
| 5                                    | None               | R   | PC                      | Program Counter   | Section 6.6.5                             |       |
| 6                                    | None               | RW  | FPRS                    | Floating-point Registers State  |   |       |
| ASR 7–15                             |                    |     | Reserved                | Reserved for future use; do not reference by software.                |   |       |
| ASR 16                               | Privileged         | RW  | PCR                     | Performance Instrumentation   | Chapter 14, “Performance Instrumentation” | 2     |
| ASR 17                               | Depends            | RW  | PIC                     |   |   | 3     |
| ASR 18                               | Privileged         | RW  | DCR                     | Dispatch Control Register   | Section 6.7.1                             |       |
| ASR 19                               | None               | RW  | GSR                     | Graphics (VIS) Status Register  | Section 6.7.2                             |       |
| ASR 20                               | Privileged         | W   | SET_SOFTINT             | Software Interrupts   | Section 6.7.3                             |       |
| ASR 21                               | Privileged         | W   | CLR_SOFTINT             |   |   |       |
| ASR 22                               | Privileged         | RW  | SOFTINT_REG             |   |   |       |
| ASR 23                               | Privileged         | RW  | TICK_CMP                | CPU and System Timer Registers  | Section 6.7.4                             |       |
| ASR 24                               | Depends            | RW  | STICK                   |   |   | 4     |
| ASR 25                               | Privileged         | RW  | STICK_CMP               |   |   |       |
| ASR 26–31                            |                    |     | Reserved                | Reserved for future use; do not reference by software.                |   |       |

1. Writes are always privileged; reads are privileged if `TICK.NPT = 1`; otherwise, reads are non-privileged.
2. If `PCR.NC = 0`, access is always privileged. If `PCR.NC ≠ 0` and `PCR.PRIV = 0`, access is non-privileged; otherwise, access is privileged.
3. All accesses are privileged if `PCR.PRIV = 1`; otherwise, all accesses are non-privileged.
4. Writes are always privileged; reads are privileged if `STICK.NPT = 1`. Otherwise, reads are non-privileged.

## 6.5.2 Privileged Register Summary

See FIGURE 6-5 and TABLE 6-7 for more information on privileged registers.



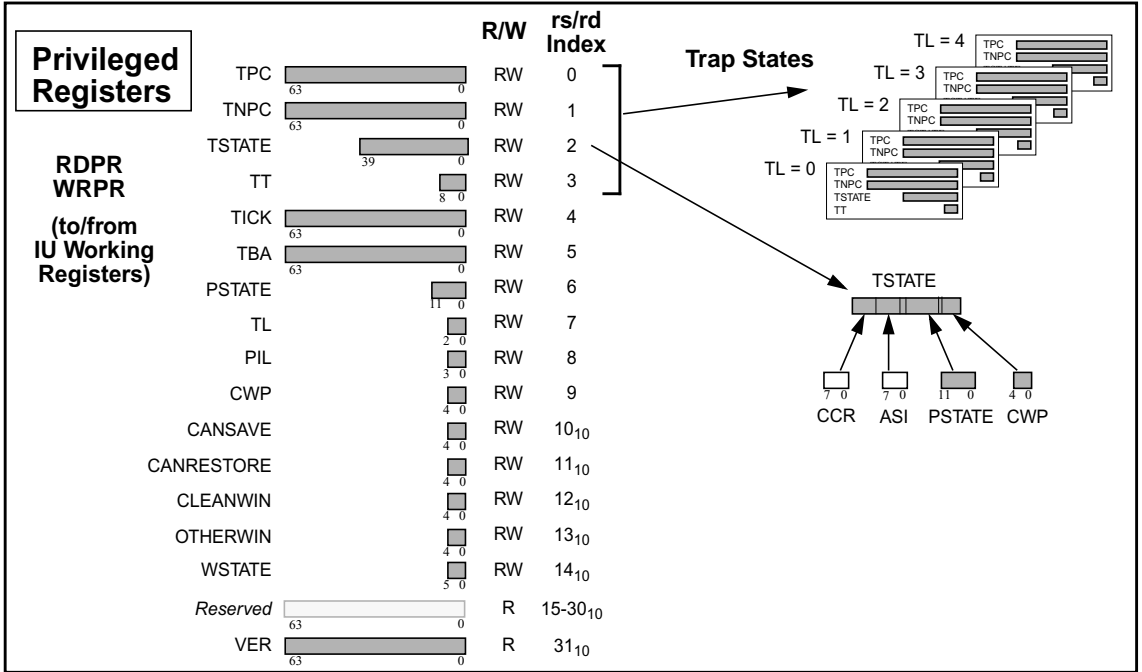


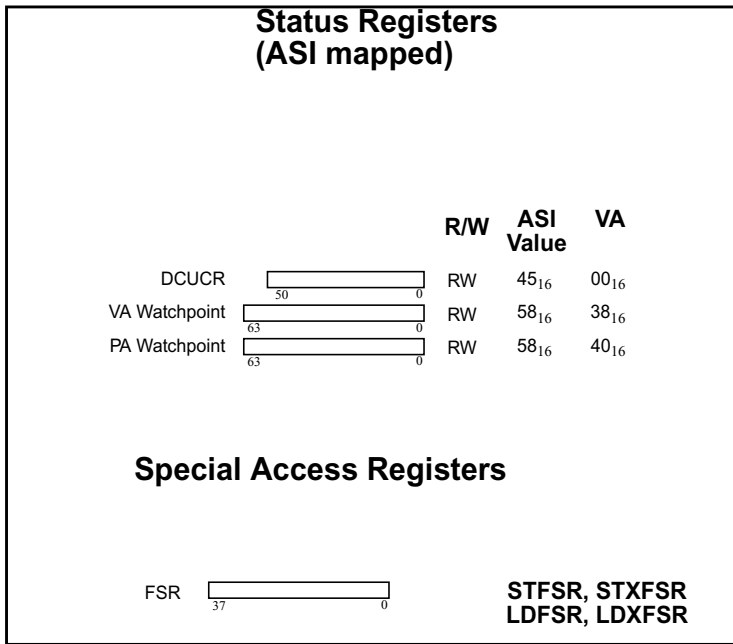
FIGURE 6-5 Privileged Registers

**TABLE 6-7** Privileged Registers

| Privileged Register Number (base 10 used) | Access Restriction | R/W | Abbreviation | Description  | Reference Section | Notes |
|---|--------------------|-----|--------------|--|-------------------|-------|
| 0   | Privileged         | RW  | TPC          | Trap stage program counter                                   | Section 6.8.1     |       |
| 1   | Privileged         | RW  | TNPC         | Trap state next program counter                              |                   |       |
| 2   | Privileged         | RW  | TSTATE       | Trap state register  |                   |       |
| 3   | Privileged         | RW  | TT           | Trap type register   |                   |       |
| 4   | Privileged         | RW  | TICK         | CPU TICK timer register, also accessible as a state register | Section 6.7.4     |       |
| 5   | Privileged         | RW  | TBA          | Trap base address register                                   | Section 6.8.2     |       |
| 6   | Privileged         | RW  | PSTATE       | Processor state register                                     | Section 6.8.3     |       |
| 7   | Privileged         | RW  | TL           | Trap level register  | Section 6.8.4     |       |
| 8   | Privileged         | RW  | PIL          | Processor Interrupt Level register                           | Section 6.8.5     |       |
| 9   | Privileged         | RW  | CWP          | Current window pointer                                       | Section 6.8.6     |       |
| 10  | Privileged         | RW  | CANSAVE      | Saveable register sets                                       |                   |       |
| 11  | Privileged         | RW  | CANRESTORE   | Restorable register sets                                     |                   |       |
| 12  | Privileged         | RW  | CLEANWIN     | Clean register sets  |                   |       |
| 13  | Privileged         | RW  | OTHERWIN     | Other register sets susceptible to spill/fill                |                   |       |
| 14  | Privileged         | RW  | WSTATE       | Window state register for traps due to spills and fills      | Section 6.8.7     |       |
| 15–30                                     | Privileged         |     | Reserved     |  |                   |       |
| 31  | Privileged         | R   | VER          | Processor version register                                   | Section 6.8.8     |       |

### 6.5.3 ASI and Specially Accessed Register Summary

See FIGURE 6-6 and TABLE 6-8 for more information on ASI and specially accessed registers.



**FIGURE 6-6** ASI and Specially Accessed Registers

**TABLE 6-8** ASI and Specially Accessed Registers

| Type                        | Abbreviation   | Description                               | Access Restriction | R/W | Reference Section | Notes |
|-----------------------------|----------------|---|--------------------|-----|-------------------|-------|
| ASI                         | DCUCR          | Data Cache Unit Control Register          |                    |     | Section 6.10.1    |       |
| ASI 58 <sub>16</sub>        | PA WATCHPOINT  | Watchpoint for physical addresses         |                    |     | Section 6.10.2    |       |
|                             | VA WATCHPOINT  | Watchpoint for virtual addresses          |                    |     |                   |       |
| LD/ST Floating-point Opcode | Load/Store FSR | Access the Floating-point Status Register |                    |     |                   |       |

## 6.6 State Registers

The state registers provide control and status to the Integer Execution Unit.

The type and accessibility of the registers (privileged vs. non-privileged mode) are summarized in FIGURE 6-4.

The SPARC V9 architecture provides for up to 31 state registers, 24 of which are classified as ASRs, numbered from 7 through 31. The eight State Registers, 0 through 7, are defined by the SPARC V9 architecture.

### 6.6.1 32-bit Multiply/Divide ( $Y^D$ ) State Register 0

The Y register is deprecated; it is provided only for compatibility with previous versions of the architecture. It should not be used in new SPARC V9 software. It is recommended that all instructions that reference the Y register (that is,  $SMUL^D$ ,  $SMULcc^D$ ,  $UMUL^D$ ,  $UMULcc^D$ ,  $MULScc^D$ ,  $SDIV^D$ ,  $SDIVcc^D$ ,  $UDIV^D$ ,  $UDIVcc^D$ ,  $RDY^D$ , and  $WRY^D$ ) be avoided.

The low-order 32 bits of the Y register, illustrated in FIGURE 6-7, contain the more significant word of the 64-bit product of an integer multiplication, as a result of either a 32-bit integer multiply ( $SMUL^D$ ,  $SMULcc^D$ ,  $UMUL^D$ ,  $UMULcc^D$ ) instruction or an integer multiply step ( $MULScc$ ) instruction. The Y register also holds the more significant word of the 64-bit dividend for a 32-bit integer divide ( $SDIV^D$ ,  $SDIVcc^D$ ,  $UDIV^D$ ,  $UDIVcc^D$ ) instruction.





FIGURE 6-7 Y Register

Although Y is a 64-bit register, its high-order 32 bits are reserved and always read as zero. The Y register is read and written with the  $RDY^D$  and  $WRY^D$  instructions, respectively.

### 6.6.2 Integer Unit Condition Codes State Register 2 (CCR)

The Condition Codes Register (CCR), shown in FIGURE 6-8, holds the integer condition codes.

The CCR is accessible using Read and Write State Register instructions ( $RDCCR$  and  $WRCCR$ ) in non-privileged or privileged mode.




**FIGURE 6-8** Condition Codes Register

### 6.6.2.1 CCR Condition Code Fields (`xcc` and `icc`)

All instructions that set integer condition codes set both the `xcc` and `icc` fields. The `xcc` condition codes indicate the result of an operation when viewed as a 64-bit operation. The `icc` condition codes indicate the result of an operation when viewed as a 32-bit operation. For example, if an operation results in the 64-bit value  $0000\ 0000\ FFFF\ FFFF_{16}$ , the 32-bit result is negative (`icc.N` is set to one) but the 64-bit result is nonnegative (`xcc.N` is set to zero).

Each of the 4-bit condition code fields is composed of four 1-bit subfields, as shown in FIGURE 6-9.



**FIGURE 6-9** Integer Condition Codes (`CCR_icc` and `CCR_xcc`)

The `n` bits indicate whether the two's-complement ALU result was negative for the last instruction that modified the integer condition codes; 1 = negative, 0 = non-negative.

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

The `v` bits signify whether the ALU result was within the range of (was representable in) 64-bit (`xcc`) or 32-bit (`icc`) two's-complement notation for the last instruction that modified the integer condition codes; 1 = overflow, 0 = no overflow.

The `c` bits indicate whether a two's-complement carry (or borrow) occurred during the last instruction that modified the integer condition codes. Carry is set on addition if there is a carry out of bit 63 (`xcc`) or bit 31 (`icc`). Carry is set on subtraction if there is a borrow into bit 63 (`xcc`) or bit 31 (`icc`); 1 = carry, 0 = no carry.

## Condition Codes

These bits are modified by the arithmetic and logical instructions, the names of which end with the letters “cc” (for example, ANDCC) and by the WRCCR instruction. They can be modified by a DONE or RETRY instruction, which replaces these bits with the CCR field of the TSTATE register. The BPCC and TCC instructions may cause a transfer of control based on the values of these bits. The MOVCC instruction can conditionally move the contents of an integer register based on the state of these bits. The FMOVCC instruction can conditionally move the contents of a floating-point register according to the state of these bits.

### CCR\_extended\_integer\_cond\_codes (xcc)

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

### CCR\_integer\_cond\_codes (icc)

Bits 3 through 0 are the IU condition codes, which indicate the results of an integer operation, with both of the operands and the result considered to be 32 bits wide. In addition to the BPCC and TCC instructions, the BiCC instruction may also cause a transfer of control based on the values of these bits.

## 6.6.3 Address Space Identifier (ASI) Register ASR 3

The ASI Register (FIGURE 6-10) specifies the ASI to be used for load and store alternate instructions that use the “rs1 + simm13” addressing form.

Non-privileged (user-mode) software may write any value into the ASI register; however, values with bit 7 equal to zero select restricted ASIs. When a non-privileged instruction makes an access that uses an ASI with bit 7 equal to zero, a *privileged\_action* exception is generated.




FIGURE 6-10 Address Space Identifier Register

## 6.6.4 TICK Register (TICK) ASR4

See Section 6.7.4, “Timer State Registers: ASRs 4, 23, 24, 25” for more details.

## 6.6.5 Program Counters State Register 5

The program counter (PC) contains the address of the instruction currently being executed. The next program counter (nPC) holds the address of the next instruction to be executed if a trap does not occur. The low-order two bits of PC and nPC always contain zero.

For a delayed control transfer, the instruction that immediately follows the transfer instruction is known as the delay instruction. This delay instruction is executed (unless the control transfer instruction annuls it) before control is transferred to the target. During execution of the delay instruction, the nPC points to the target of the control transfer instruction, and the PC points to the delay instruction. See Chapter 7, “Instruction Types” for more details.

The PC is used implicitly as a destination register by CALL, BiCC, BPCC, BPr, FBfCC, FBPFCC, JMPL, and RETURN instructions. It can be read directly by a RDPC instruction.

## 6.6.6 Floating-Point Registers State (FPRS) Register 6

The Floating-point Registers State (FPRS) Register, shown in FIGURE 6-11, holds control information for the floating-point register file. Mode and status information about the Floating-point unit (FPU) is presented in Section 6.9.1.

This register is readable and writable using the read and write state register instructions RDFPRS and WRFPRS when the processor is in non-privileged or privileged mode.




FIGURE 6-11 Floating-Point Registers State Register

### 6.6.6.1 FPRS\_enable\_fp (FEF)

Bit 2, FEF, determines whether the FPU is enabled. If it is disabled, executing a floating-point instruction causes a *fp\_disabled* trap. If this bit is set but the PSTATE.PEF bit is not set, then executing a floating-point instruction causes a *fp\_disabled* trap; that is, both FPRS.FEF and PSTATE.PEF must be set to enable floating-point operations.

### 6.6.6.2 FPRS\_dirty\_upper (DU)

Bit 1 is the “dirty” bit for the upper half of the floating-point registers; that is,  $f32$ – $f62$ . It is set whenever any of the upper floating-point registers is modified. The processor may set it pessimistically; it may be set whenever a floating-point instruction is issued, even though that instruction never completes and no output register is modified. The dirty bit may be set by instructions that the processor executes but does not complete due to wrong branch prediction. The DU bit is cleared only by software.

### 6.6.6.3 FPRS\_dirty\_lower (DL)

Bit 0 is the “dirty” bit for the lower 32 floating-point registers; that is,  $f0$ – $f31$ . It is set whenever any of the lower floating-point registers is modified. The processor may set it pessimistically; it may be set whenever a floating-point instruction is issued, even though that instruction never completes and no output register is modified. The DL bit is cleared only by software.

---

## 6.7 Ancillary State Registers: ASRs 16-25

The SPARC V9 architecture provides for optional *ancillary* state registers (ASRs) in addition to the six state registers defined for all SPARC V9 processors and already described.

An ASR is read and written with the RDASR and WRASR instructions, respectively. Access to a particular ASR may be privileged or non-privileged. A RDASR or WRASR instruction is privileged if the accessed register is privileged.

All the state and ancillary state registers are summarized in TABLE 6-6. Some of the registers descriptions are presented below.


### 6.7.1 Dispatch Control Register (DCR) ASR 18

The DCR provides control over the dispatch unit and branch prediction logic. The DCR also provides factory test equipment with access to internal logic states using the OBSDATA bus interface.

The DCR is a read/write register. Unused bits read as zero and should be written only with zero or values previously read from them. The DCR is a privileged register; attempted access by non-privileged (user) code causes a *privileged\_opcode* trap. POR value is  $xxxx.xx0x_2$ .

The DCR is illustrated in FIGURE 6-12 and described in TABLE 6-9.





**FIGURE 6-12** Dispatch Control Register (ASR 18)

**TABLE 6-9** DCR Bit Description

| Bit                                 | Field   | Type | Description  |
|-------------------------------------|---------|------|--|
| 63:14                               | —       |      | <i>Reserved.</i>   |
| 13:12                               | DPE     |      | Data Cache Parity Error Enable - If cleared, no parity checking at the Data Cache SRAM arrays (Data, Physical Tag, and Snoop Tag arrays) will be done. It also implies no Dcache_Parity_error trap (TT 0x071) will ever be generated. However, parity bits are still generated automatically and correctly by HW.  |
| 11:6                                | OBSDATA |      | These bits are used to select the set of signals driven on the OBSDATA<9:0> pins of the processor for factory test purposes.   |
| <b>Branch and Return Control</b>    |         |      |  |
| 5                                   | BPE     |      | <i>Branch Prediction Enable.</i> When BPE = 1, conditional branches are predicted through internal hardware. When BPE = 0, all branches are predicted not taken. After power-on reset initialization, this bit is set to zero. This bit is also automatically set to zero on any trap causing RED_state entry (but not cleared when privileged code enters RED_state by setting the RED bit in PSTATE).  |
| 4                                   | RPE     |      | <i>Return Address Prediction Enable.</i> When RPE = 0, the return address prediction stack is disabled. Even when encountering a JMPL instruction, instruction fetch will continue on a sequential path until the return address is generated and a mispredict is signalled. When RPE = 1, the processor may attempt to predict the target address of JMPL instructions and prefetch subsequent instructions accordingly. After power-on reset initialization, this bit is set to zero. This bit is also automatically set to zero on any trap causing a RED_state entry (but left unchanged when privileged code enters RED_state by setting PSTATE.RED). |
| <b>Instruction Dispatch Control</b> |         |      |  |
| 3                                   | SI      |      | <i>Single Issue Disable.</i> When SI = 0, only one instruction will be outstanding at a time. Superscalar instruction dispatch is disabled, and only one instruction is executed at a time. When SI = 1, normal pipelining is enabled. The processor can issue new instructions prior to the completion of previously issued instructions. After power-on reset initialization, this bit is set to zero. This bit is also automatically set to zero on any trap causing RED_state entry (but left unchanged when privileged code enters RED_state by setting PSTATE.RED).  |

**TABLE 6-9** DCR Bit Description (Continued)

| Bit | Field | Type | Description  |
|-----|-------|------|--|
| 2   | IPE   |      | Instruction Cache Parity Error Enable - If cleared, no parity checking at the Instruction Cache SRAM arrays (Data, Physical Tag, and Snoop Tag arrays) will be done. It also implies no <code>Icache_Parity_error</code> trap (TT 0x072) will ever be generated. However, parity bits are still generated automatically and correctly by HW.   |
| 1   | IFPOE |      | Interrupt Floating-point Operation Enable. The IFPOE bit enables system software to take interrupts on floating-point instructions. When set, the processor forces a <code>fp_disabled</code> trap when an interrupt occurs on floating-point code.  |
| 0   | MS    |      | <i>Multiscalar dispatch enable</i> . When <code>MS = 0</code> , the processor operates in scalar mode, issuing and executing one instruction at a time. Pipelined operation is still controlled by the <code>SI</code> bit. <code>MS = 1</code> enables superscalar (normal) instruction issue.<br><br>After power-on reset initialization, this bit is set to zero. The bit is also automatically set to zero on any trap causing <code>RED_state</code> entry (but left unchanged when privileged code enters <code>RED_state</code> by setting <code>PSTATE.RED</code> ). |

---

**Note** – Both IPE and DPE will default to 0 (disable) after power-on or system reset.

---

## Interrupt Floating-Point Operation Enable (Bit 1)

The IFPOE bit enables system software to take interrupts on floating-point instructions. This enable bit is cleared by hardware at power-on. System software must set the bit as needed. When this bit is enabled, the UltraSPARC III Cu processor forces a `fp_disabled` trap when an interrupt occurs on FP-only code. The trap handler is then responsible for checking whether the floating-point is indeed disabled. If it is not, the trap handler then enables interrupts to take the pending interrupt.

---

**Note** – This behavior deviates from SPARC V9 trap priorities in that interrupts are of lower priorities than the other two types of floating-point exceptions (*fp\_exception\_ieee\_754*, *fp\_exception\_other*).

---

- This mechanism is triggered for a floating-point instruction only if none of the approximately twelve preceding instructions across the two integer, load/store, and branch pipelines are valid, under the assumption that they are better suited to take the interrupt (only one trap entry/exit).
- Upon entry, the handler must check both `TSTATE.PEF` and `FPRF.FEF` bits. If `TSTATE.PEF = 1` and `FPRF.FEF = 1`, the handler has been entered because of an interrupt, either *interrupt\_vector* or *interrupt\_level*. In such a case:
  - The `fp_disabled` handler should enable interrupts (that is, set `PSTATE.IE = 1`); then, issue an integer instruction (for example, `add %g0, %g0, %g0`). An interrupt is triggered on this instruction.

- The processor then enters the appropriate interrupt handler (`PSTATE.IE` is turned off here) for the type of interrupt.
- At the end of the handler, the interrupted instruction is a `RETRY` after returning from the interrupt. The `add %g0, %g0, %g0` is a `RETRY`.
- The `fp_disabled` handler then returns to the original process with a `RETRY`.
- The “interrupted” FPop is then retried (taking a `fp_exception_ieee_754` or `fp_exception_other` at this time if needed).


## 6.7.2 Graphics Status Register (GSR) ASR 19

The GSR is used with the VIS Instruction Set.

The GSR is accessible in non-privileged mode. It can be read and written using the `RDASR` and `WRASR` state register instructions.

Accesses to the GSR cause a `fp_disabled` trap if `PSTATE.PEF` or `FPRS.FEF` is zero.

The GSR is illustrated in FIGURE 6-13 and described in TABLE 6-10.



**FIGURE 6-13** Graphic Status Register (ASR 19)

**TABLE 6-10** GSR Bit Description

| Bit   | Field                 | Description  |      |              |   |                       |   |   |   |    |   |    |
|-------|-----------------------|--|------|--------------|---|-----------------------|---|---|---|----|---|----|
| 63:32 | MASK<31:0>            | This field specifies the mask used by the <code>BSHUFFLE</code> instruction. The field contents are set by the <code>BMASK</code> instruction.   |      |              |   |                       |   |   |   |    |   |    |
| 31:28 | <i>Reserved.</i>      |  |      |              |   |                       |   |   |   |    |   |    |
| 27    | IM                    | Interval Mode: When <code>IM = 1</code> , the values in <code>FSR.RD</code> and <code>FSR.NS</code> are ignored; the processor operates as if <code>FSR.NS = 0</code> and rounds floating-point results according to <code>GSR.IRND</code> .   |      |              |   |                       |   |   |   |    |   |    |
| 26:25 | IRND<1:0>             | IEEE Std 754-1985 rounding direction to use in Interval Mode ( <code>GSR.IM = 1</code> ), as follows: <table border="1" style="margin-left: auto; margin-right: auto; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">IRND</th> <th style="text-align: left;">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> <p>When <code>GSR.IM = 1</code>, the value in <code>GSR.IRND</code> overrides the value in <code>FSR.RD</code>.</p> | IRND | Round toward | 0 | Nearest (even if tie) | 1 | 0 | 2 | +∞ | 3 | -∞ |
| IRND  | Round toward          |  |      |              |   |                       |   |   |   |    |   |    |
| 0     | Nearest (even if tie) |  |      |              |   |                       |   |   |   |    |   |    |
| 1     | 0                     |  |      |              |   |                       |   |   |   |    |   |    |
| 2     | +∞                    |  |      |              |   |                       |   |   |   |    |   |    |
| 3     | -∞                    |  |      |              |   |                       |   |   |   |    |   |    |

**TABLE 6-10** GSR Bit Description (Continued)

| Bit  | Field            | Description  |
|------|------------------|--|
| 24:8 | <i>Reserved.</i> |  |
| 7:3  | SCALE<4:0>       | Shift count in the range 0–31, used by the PACK instructions for formatting.   |
| 2:0  | ALIGN<2:0>       | Least three significant bits of the address computed by the last executed ALIGNADDRESS or ALIGNADDRESS_LITTLE instruction. |


## 6.7.3 Software Interrupt State Registers: ASRs 20, 21, and 22

Three registers are used to control software interrupts: SOFTINT, SET\_SOFTINT, and CLR\_SOFTINT. Bits written to the SOFTINT register will cause traps to the level the trap is enabled. The SOFTINT register can be written to directly using ASR 22, or indirectly using the SET\_SOFTINT and CLR\_SOFTINT registers as described in this section.

All three registers are accessible only in privileged mode. The SOFTINT register is accessed using the RD and WR state register access instructions. The SET\_SOFTINT and CLR\_SOFTINT registers are written using the WR state register access instruction. See TABLE 6-11 and FIGURE 6-14 for more details.

**TABLE 6-11** Register-Window State Registers

| Soft Interrupt Register | ASR # | Name and Description                    | Privileged Access Instructions |
|-------------------------|-------|---|--------------------------------|
| SOFTINT                 | 22    | Software Interrupt Register             | RDSOFTINIT<br>WRSOFTINT        |
| SET_SOFTINT             | 20    | Set Software Interrupt register bits.   | WRSOFTINIT_SET                 |
| CLR_SOFTINT             | 21    | Clear Software Interrupt register bits. | WRSOFTINIT_CLR                 |



**FIGURE 6-14** SOFTINT, SET\_SOFTINT, and CLR\_SOFTINT Register Formats

## *SOFTINT Register*

The operating system uses the SOFTINT to schedule interrupts. The field definitions are described in TABLE 6-12.

**TABLE 6-12** SOFTINT Bit Descriptions

| Bit  | Field             | Description   |
|------|-------------------|---|
| 16   | SM<br>(STICK_INT) | When the STICK_COMPARE.INT_DIS bit is zero (system tick compare is <i>enabled</i> ) and its STICK_CMPR field matches the value in the STICK register, then the SM field in SOFTINT is set to one and a Level-14 interrupt is generated. See Section 6.7.4, “Timer State Registers: ASRs 4, 23, 24, 25” for details. |
| 15:1 | INT_LEVEL         | When a bit is set within this field (bits 15:1), an interrupt is caused at the corresponding interrupt level. Note that INT_LEVEL<15> is shared by Level-15 interrupt and PIC overflow interrupt.   |
| 0    | TM<br>(TICK_INT)  | When the TICK_COMPARE.INT_DIS bit is zero (that is, tick compare is <i>enabled</i> ) and its TICK_CMPR field matches the value in the TICK register, then the TM field in the SOFTINT register is set to one and a Level-14 interrupt is generated. See “TICK_COMPARE Register” for details.                        |

## *SET\_SOFTINT Register*

The SET\_SOFTINT register is written to set bits in the SOFTINT register to set a bit in that register. When a bit in the SET\_SOFTINT register is set to a one, the corresponding bit in the SOFTINT is set.

## *CLR\_SOFTINT Register*

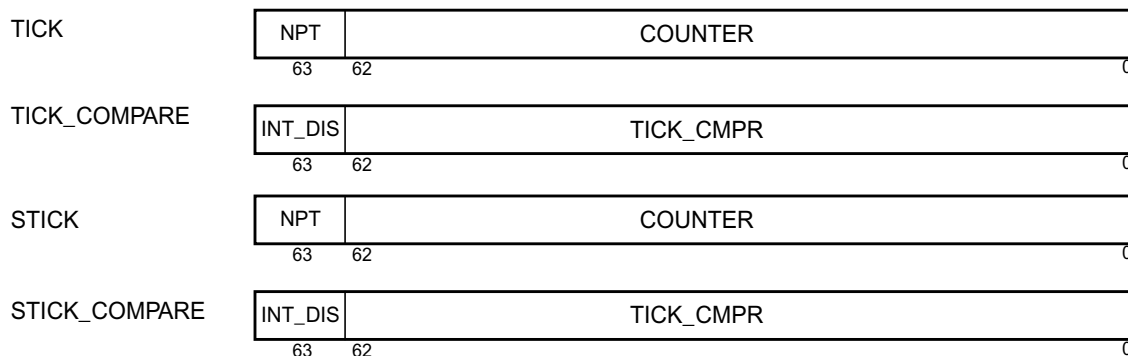
The CLR\_SOFTINT register is written in privileged mode using the WR write state register instruction to clear bits in the SOFTINT register. When a bit in the CLR\_SOFTINT register is set to a one, the corresponding bit in the SOFTINT register is cleared.

## 6.7.4 Timer State Registers: ASRs 4, 23, 24, 25

The processor has two timers. The TICK timer is driven by the CPU clock. The STICK timer is driven by the system clock. Four registers are used to implement the timer and support the timer interrupts.

**TABLE 6-13** Timer State Registers

| Soft Interrupt Register | ASR #<br>(base 10) | Name and Description   | Access Instructions                            |
|-------------------------|--------------------|------------------------|--|
| TICK                    | 4                  | TICK register          | Depends  |
| TICK_COMPARE            | 23                 | TICK Compare register  | State Register Instructions in privileged mode |
| STICK                   | 24                 | STICK register         | Depends  |
| STICK_COMPARE           | 25                 | STICK Compare register | State Register Instructions in privileged mode |



**FIGURE 6-15** Timer State Registers

### *TICK Register*

The TICK register is a 63-bit counter that counts processor clock cycle.

In privileged mode, the TICK register is always readable using either the RDPR (privileged read) or RDTICK (state register read) instructions. The TICK register is always writable in privileged mode using the WRPR (privileged write) instruction; there is no WRTICK (state register write) instruction.

The TICK.NPT bit (bit 63) selects the non-privileged mode readability. If TICK.NPT = 0, then the TICK register is readable in non-privileged mode using the RDTICK state register read instruction. When TICK.NPT = 1, an attempt by software to read the TICK register in non-privileged mode causes a *privileged\_action* exception. Software operating in non-privileged mode can never write to the TICK register.

TICK.NPT is set to one by a power-on reset trap. The value of TICK.COUNTER is reset after a power-on reset trap.

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

---

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

---

---

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

---

### *TICK\_COMPARE Register*

The TICK\_COMPARE register causes the processor to generate a trap when the TICK register reaches the value in the TICK\_COMPARE register and the INT\_DIS bit is zero. If the INT\_DIS bit is one, then no interrupt is generated.

When the TICK\_CMPR field exactly matches the TICK.COUNTER field and INT\_DIS = 0, then a TICK\_INT is posted in the SOFTINT register. This has the effect of posting a Level-14 interrupt to the processor when the processor has PIL register value less than fourteen and PSTATE.IE register field 1.

---

**Programming Note** – The Level-14 interrupt handler must check the SOFTINT<14>, TM (TICK\_INT), and SM (STICK\_INT) fields of the SOFTINT register to determine the source or sources of the Level-14 interrupt.

---

In privileged mode, the TICK\_COMPARE register is always accessible using the state register read and write instructions. The TICK\_COMPARE register is not accessible in non-privileged mode. Non-privileged accesses to this register causes a *privileged\_opcode* trap.

### *STICK Register*

The STICK register is a 63-bit counter that increments at a rate determined by the system clock.

The `STICK` register is always accessible in privileged mode using the `RDSTICK` and `WRSTICK` state register instructions.

The `STICK.NPT` bit (bit 63) selects the non-privileged mode readability. If `STICK.NPT = 0`, then the `STICK` register is readable in non-privileged mode using the `RDSTICK` state register read instruction. When `STICK.NPT = 1`, an attempt by software to read the `STICK` register in non-privileged mode causes a *privileged\_action* exception. Software operating in non-privileged mode can never write to the `STICK` register.

`STICK.NPT` bit is set to one by a power-on reset trap. The value of `STICK.COUNTER` is cleared after a power-on reset trap.

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

---

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

---

### *STICK\_COMPARE Register*

The `STICK_COMPARE` register causes the processor to generate a trap when the `STICK` register reaches the value in the `STICK_COMPARE` register and the `INT_DIS` bit is zero. If the `INT_DIS` bit is one, then no interrupt is generated.

The `STICK_COMPARE` is only accessible in privileged mode. Accesses to this register in non-privileged mode causes a *privileged\_opcode* trap.

When `STICK_CMPR` field exactly matches `STICK.COUNTER` field and `INT_DIS = 0`, then a `TICK_INT` is posted in the `SOFTINT` register. This has the effect of posting a Level-14 interrupt to the processor when the processor has `PIL` register value less than fourteen and `PSTATE.IE` register field 1.

---

**Programming Note** – The Level-14 interrupt handler must check `SOFTINT<14>`, `TICK_INT`, and `STICK_INT` to determine the source of the Level-14 interrupt.

---

After a power-on reset trap, the `INT_DIS` bit is set to one (disabling system tick compare interrupts), and the `STICK_CMPR` value is set to zero.



---

## 6.8 Privileged Registers

The privileged registers are described in this section. The privileged registers are visible only to software running in privileged mode (`PSTATE.PRIV = 1`). Privileged registers are written with the `WRPR` instruction and read with the `RDPR` instruction.

Refer to FIGURE 6-5 for more details.

### 6.8.1 Trap Stack Privileged Registers 0 through 3

The four trap stack registers (`TPC`, `TNPC`, `TSTATE`, and `TT`) form a group of registers that are shadowed for each of the five trap levels. Each instance of the registers save the state of key integer unit parameters at each trap level. FIGURE 6-16 shows the format for this register group. This figure is followed by a description of each register. FIGURE 6-17 shows how the register stack responds to an event example.

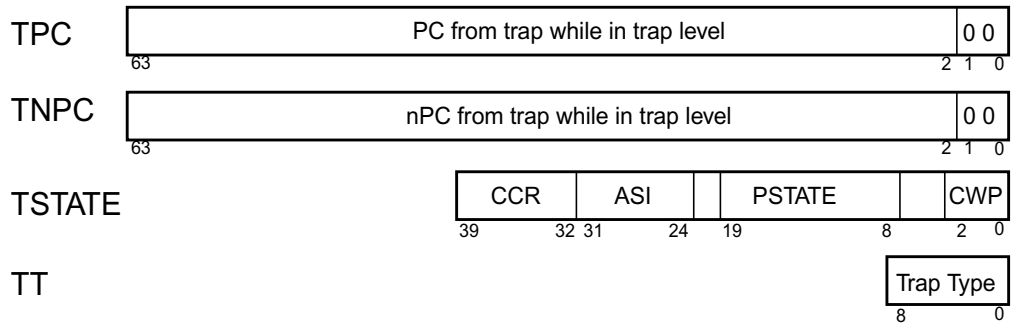
The group of trap stack registers contain state information from the previous trap level. The registers include values from the program counter (`PC`), the next program counter (`nPC`), the trap state (`TSTATE`) register (a group of fields comprising the contents of the `CCR`, `ASI`, `CWP`, and `PSTATE` registers), and the trap type (`TT`) register containing the value of the trap that caused entry into the current trap level.

#### 6.8.1.1 Common Attributes

There are `MAXTL = 5` instances of the trap control registers, but only one group is accessible at any time. The current value in the `TL` register determines which instance of the trap control registers are accessible.

All trap control registers are accessible in privileged mode. An attempt to read or write any of these registers in non-privileged mode causes a *privileged\_opcode* exception.

An attempt to read or write any of these registers when `TL = 0` causes an *illegal\_instruction* exception.



**FIGURE 6-16** Trap State Register Format

### *Trap Program Counter*

The Trap Program Counter (TPC) contains the PC from the previous trap level.

### *Trap Next Program Counter*

The Trap Next Program Counter (TNPC) register is the nPC from the previous trap level.

### *Trap State Register*

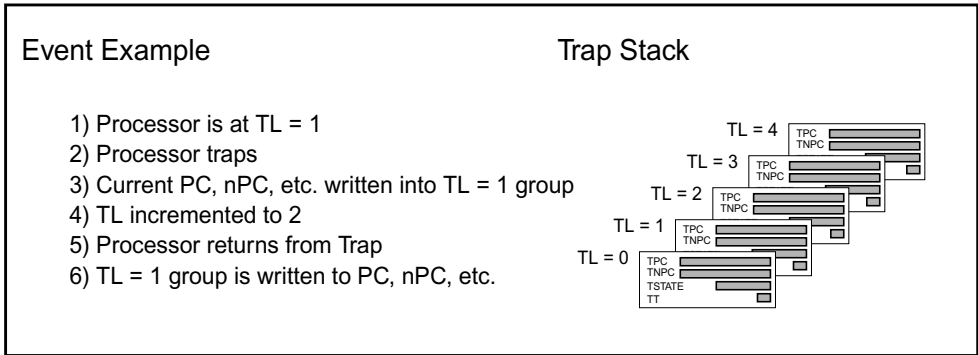
The Trap State (TSTATE) Register contains the state from the previous trap level, comprising the contents of the CCR, ASI, CWP, and PSTATE registers from the previous trap level.

### *Trap Type*

The Trap Type (TT) register normally contains the trap type of the trap that caused entry to the current trap level.

## 6.8.1.2 Trap Stack Operation

The trap stack and an event example are shown in FIGURE 6-17.



**FIGURE 6-17** Trap Stack and Event Example

### 6.8.1.3 Effects of Reset and Normal Operation


The effects of reset on each register are shown in TABLE 6-14. During normal operation, the trap stack register values defined for the trap levels above the current one are undefined.

**TABLE 6-14** Trap Stack Register Power-On and Normal Operation

| Trap Control Register | After Power-on Reset   | During Normal Operation, for <i>n</i> greater than the current trap level ( <i>n</i> > TL) |
|-----------------------|--|--|
| TPC                   | TPC[0] =<br>TPC[1] to TPC[5] are undefined   | TPC[ <i>n</i> ] is undefined   |
| TNPC                  | TNPC[0] =<br>TNPC[1] to TNPC[5] are undefined  | TNPC[ <i>n</i> ] is undefined  |
| TSTATE                | TSTATE[0] =<br>TSTATE[1] to TSTATE[5] are undefined                                  | TSTATE[ <i>n</i> ] is undefined  |
| TT                    | TT[0] = Reset Trap Type<br>TT[1] to TT[4] are undefined<br>TT[5] = 001 <sub>16</sub> | TT[ <i>n</i> ] is undefined  |


## 6.8.2 Trap Base Address (TBA) Privileged Register 5

The TBA register, shown in FIGURE 6-18, provides the upper 49 bits of the address used to select the trap vector for a trap. The TBA register is accessible using read and write privileged register instructions. The lower 15 bits of the TBA always read as zero, and writes to them are ignored.



**FIGURE 6-18** Trap Base Address Register

The full address for a trap vector is specified by the contents in the TBA, TL, and TT[TL] registers at the time the trap is taken, as shown in FIGURE 6-19.



**FIGURE 6-19** Trap Vector Address Format

**TL > 0 bit**


The “TL > 0” bit is zero if TL = 0 when the trap was taken, and one if TL > 0 when the trap was taken. This implies that there are two trap tables: one for traps from TL = 0 and one for traps from TL > 0. See Chapter 12, “Traps and Trap Handling” for more details on trap vectors.

**TT<sub>TL</sub> field**

The TT<sub>TL</sub> field is written with the contents of the TT register representing the new trap level that is being taken.

## 6.8.3 Processor State (PSTATE) Privileged Register 6

The PSTATE register (FIGURE 6-20) holds the current state of the processor. There is only one instance of the PSTATE register. The PSTATE register is copied to a 12-bit field in the TSTATE register of the trap stack. See Chapter 12, “Traps and Trap Handling” for more details.



**FIGURE 6-20** PSTATE Fields

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

Subsections on page 113 through page 114 describe the fields contained in the PSTATE register.

### 6.8.3.1 Global Register Set Selection - IG, MG, AG bits

The UltraSPARC III Cu processor provides Interrupt and MMU Global Register sets in addition to the two global register sets (normal and alternate) specified by SPARC V9. The currently active set of global registers is specified by the AG, IG, and MG bits and are set and cleared according to the events listed in TABLE 6-15.

---

**Note** – The IG, MG, and AG fields are saved on the trap stack along with the rest of the PSTATE Register.

---

**TABLE 6-15** PSTATE Global Register Selection Events

| Event  | Globals Selected for Use  | PSTATE Settings |    |    |
|--|---|-----------------|----|----|
|  |   | AG              | IG | MG |
| DONE, RETRY [1]  | Global Registers encoded in TSTATE register (Previous Global Registers before most recent trap) | 0               | 0  | 0  |
| <i>fast_instruction_access_MMU_miss,</i><br><i>fast_data_access_MMU_miss,</i><br><i>fast_data_access_protection,</i><br><i>data_access_exception,</i><br><i>instruction_access_exception</i> | MMU Global registers  | 0               | 0  | 1  |
| <i>interrupt_vector_trap</i>   | Interrupt Global registers  | 0               | 1  | 0  |
|  | Reserved [2]  | 0               | 1  | 1  |
| Write to privileged register (WPR) that modifies AG, IG or MG bits in PSTATE register  | Any Global Register   | x               | x  | x  |
| Any trap other than those listed above   | Alternate Global registers  | 1               | 0  | 0  |
|  | <i>Reserved.</i>  | 1               | 0  | 1  |
|  | <i>Reserved.</i>  | 1               | 1  | 0  |
|  | <i>Reserved.</i>  | 1               | 1  | 1  |

1. Since PSTATE is preserved in the TSTATE register when a trap occurs, the previous value of these bits are normally restored upon return from a trap (via DONE or RETRY instruction).
2. A WRPR to PSTATE, using a reserved combination of AG, IG, and MG bit values, causes an *illegal\_instruction* exception.

Executing a DONE or RETRY instruction restores the previous {AG, IG, MG} state before the trap is taken. Programmers can also set or clear these three bits by writing to the PSTATE register with a WRPR instruction.

---

**Note** – Attempting to use the “wrpr %pstate” instruction to set a reserved encoding for IG, MG, and AG (more than one of these bits set) results in an *illegal\_instruction* exception. However, the processor does not check for a reserved encoding when writing directly to the TSTATE register. Hence, executing a DONE or RETRY with an invalid AG, IG, MG bit combination may result in an undefined behavior of the processor.

---

---

**Compatibility Note** – UltraSPARC III Cu processors support two more sets (privileged only) of eight 64-bit global registers compared to the UltraSPARC II family: interrupt globals and MMU globals. These additional registers are called the *trap globals*. Two 1-bit fields, PSTATE.IG and PSTATE.MG, were added to the PSTATE register to select which set of global registers to use.

---

### *PSTATE\_interrupt\_globals (IG)*

When PSTATE.IG = 1, the processor interprets integer register numbers in the range 0–7 as referring to the interrupt global register set. See the **Note** on page 114. When an *interrupt\_vector* trap (trap type = 60<sub>16</sub>) is taken, processor sets IG and clears AG and MG.

### *PSTATE\_MMU\_globals (MG)*

When PSTATE.MG = 1, the processor interprets integer register numbers in the range 0–7 as referring to the MMU global register set.

The processor sets MG and clears IG and AG when any of the following traps are taken:

- *fast\_instruction\_access\_MMU\_miss* trap (trap type = 64<sub>16</sub>–67<sub>16</sub>)
- *fast\_data\_access\_MMU\_miss* trap (trap type = 68<sub>16</sub>–6B<sub>16</sub>)
- *fast\_data\_access\_protection* trap (trap type = 6C<sub>16</sub>–6F<sub>16</sub>)
- *data\_access\_exception* trap (trap type = 30<sub>16</sub>)
- *instruction\_access\_exception* trap (trap type = 08<sub>16</sub>)

### *PSTATE\_alternate\_globals (AG)*

When PSTATE.AG = 1, the processor interprets integer register numbers in the range 0–7 as referring to the alternate global register set.

If an exception is taken and it does not set another global bit, then the processor defaults to the Alternate Global register set by setting AG and clearing IG and MG.

### 6.8.3.2 PSTATE\_current\_little\_endian (CLE)

When `PSTATE.CLE = 1`, all data reads and writes using an implicit ASI are performed in little-endian byte order with an ASI of `ASI_PRIMARY_LITTLE`. When `PSTATE.CLE = 0`, all data reads and writes using an implicit ASI are performed in big-endian byte order with an ASI of `ASI_PRIMARY`. Instruction accesses are always big-endian.

### 6.8.3.3 PSTATE\_trap\_little\_endian (TLE)

When a trap is taken, the current `PSTATE` register is pushed onto the trap stack and the `PSTATE.TLE` bit is copied into `PSTATE.CLE` in the new `PSTATE` register. This behavior allows system software to have a different implicit byte ordering than the current process. Thus, if `PSTATE.TLE` is set to one, data accesses using an implicit ASI in the trap handler are little-endian. The original state of `PSTATE.CLE` is restored when the original `PSTATE` register is restored from the trap stack.

### 6.8.3.4 PSTATE\_mem\_model (MM)

The processor supports Total Store Order (TSO), only. The 2-bit field in the `PSTATE.MM` is hardwired to 00 indicating TSO mode. See TABLE 6-16 for MM Encodings.

TABLE 6-16 MM Encodings

| MM Value | SPARC V9                |
|----------|-------------------------|
| 00       | Total Store Order (TSO) |
| 01       | Reserved                |
| 10       | Reserved                |
| 11       | Reserved                |

- **Total Store Order (TSO)** — Loads are ordered with respect to earlier loads. Stores are ordered with respect to earlier loads and stores. Thus, loads can bypass earlier stores but cannot bypass earlier loads; stores cannot bypass earlier loads and stores. Programs that execute correctly in either PSO or RMO will execute correctly in the TSO model.

### 6.8.3.5 PSTATE\_RED\_state (RED)

PSTATE.RED (Reset, Error, and Debug state) is set whenever the UltraSPARC III Cu processor takes a RED state disrupting or nondisrupting trap. The IU sets PSTATE.RED when any hardware reset occurs. It also sets PSTATE.RED when a trap is taken while  $TL = (MAXTL - 1)$ . Software can exit RED\_state by the following method:

- Execute a DONE or RETRY instruction, which restores the stacked copy of PSTATE and clears PSTATE.RED if it was zero in the stacked copy.

---

**Note** – Software can also exit the RED\_state by writing a zero to PSTATE.RED with a WRPR instruction. However, this method is not recommended due to potential side-effects and unpredictable behavior.

---

### 6.8.3.6 PSTATE\_enable\_floating-point (PEF)

When set to one, the PEF bit enables the FPU, which allows privileged software to manage the FPU. For the FPU to be usable, both PSTATE.PEF and FPRS.FEF must be set. Otherwise, any floating-point instruction that tries to reference the FPU causes a *fp\_disabled* trap.

### 6.8.3.7 PSTATE\_address\_mask (AM)

When PSTATE.AM = 1, the high-order 32 bits of any virtual addresses for instruction and data are cleared to zero in the following cases:

- Before data addresses are sent out of the processor
- Before addresses are sent to the MMU
- For instruction accesses to all caches
- Before being stored to a general-purpose register for CALL, JMPL, and RDPC instructions
- Before being stored to TPC[n] and TNPC[n] when a trap occurs

When an ASI\_PHYS\_\* ASI is used in a load or store instruction, the setting of PSTATE.AM is ignored and the full 64-bit address is used. (See ASI 14<sub>16</sub>, ASI\_PHYS\_USE\_EC, for an example).

When PSTATE.AM = 1, the processor writes the full 64-bit program counter value (upper 32 bits are forced to be zero) to the destination register of a CALL, JMPL, or RDPC instruction.

When PSTATE.AM = 1 and a trap occurs, the processor writes the full 64-bit program counter value to TPC[TL].

When PSTATE.AM = 1 and a synchronous exception occurs, the processor writes the full 64-bit address to the Data Synchronous Fault Address Register (D-SFAR).



When `PSTATE.AM = 1` and an asynchronous exception occurs, the processor writes the full 64-bit address to the Data Asynchronous Fault Address Register (D-AFAR).

The `PSTATE.AM` bit must be set when 32-bit software is executed.

### 6.8.3.8 PSTATE\_privileged\_mode (PRIV)


When `PSTATE.PRIV = 1`, the processor is in privileged mode. This bit is controlled by events in the processor and can be explicitly set.

### 6.8.3.9 PSTATE\_interrupt\_enable (IE)

When `PSTATE.IE = 1`, the processor can accept interrupts.

## 6.8.4 Trap Level (TL) Privileged Register 7

The trap level register, shown in FIGURE 6-21, specifies the current trap level. `TL = 0` is the normal (nontrap) level of operation. `TL > 0` implies that one or more traps are being processed. The maximum valid value that the TL register may contain is `MAXTL = 5`, which is always equal to the number of supported trap levels beyond Level-0. See Chapter 12, “Traps and Trap Handling” for more details about the TL register.



**FIGURE 6-21** Trap Level Register


---

**Programming Note** – Writing the TL register with a value greater than `MAXTL` (five for UltraSPARC III Cu) causes the value `MAXTL` to be written. Writing the TL register with a `wrwr %t1` instruction does not alter any other processor state; that is, it is not equivalent to taking or returning from a trap.

---

## 6.8.5 Processor Interrupt Level (PIL) Privileged Register 8

The processor interrupt level (PIL), illustrated in FIGURE 6-22, is the interrupt level above which the processor will accept an interrupt. Interrupt priorities are mapped so that interrupt Level-2 has greater priority than interrupt Level-1, and so on.



**FIGURE 6-22** Processor Interrupt Level Register

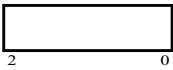
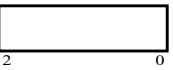
On SPARC V8 processors, the Level-15 interrupt is considered to be nonmaskable; therefore, it has different semantics from other interrupt levels. SPARC V9 processors do not treat Level-15 interrupts differently from other interrupt levels.

## 6.8.6 Register-Window State Privileged Registers 9 through 13

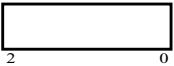
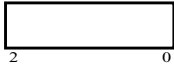
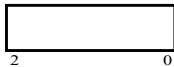
The state of the register window is determined by a set of privileged registers that are read and written by privileged mode software using the RDPR and WRPR instructions, respectively. In addition, these privileged registers are modified by instructions related to register windowing and are used to generate traps that allow supervisor software to spill, fill, and clean the register window sets.

Register-window management is described in a separate chapter.

**TABLE 6-17** Register-Window State Privileged Registers

| Register-Window State Registers   | Value Range | Description  |
|---|-------------|--|
| Current Window Pointer<br>CWP    | 0 to 7      | <b>State Register 9:</b> The CWP register is a counter that identifies the current window into the set of integer registers. See Chapter 12, “Traps and Trap Handling” for information on how hardware manipulates the CWP register. |
| Saveable Window Sets<br>CANSAVE  | 0 to 6      | <b>State Register 10:</b> The CANSAVE register contains the number of register sets following CWP that are not in use and are available to be allocated by a SAVE instruction without generating a window spill exception.           |

**TABLE 6-17** Register-Window State Privileged Registers (*Continued*)

| Register-Window State Registers  | Value Range | Description  |
|--|-------------|--|
| Restorable Window Sets<br><br>CANRESTORE  | 0 to 7      | <b>State Register 11:</b> The CANRESTORE register contains the number of register sets preceding CWP that are in use by the current program and can be restored (by the RESTORE instruction) without generating a window fill exception.   |
| Clean Window Sets<br><br>CLEANWIN         | 0 to 6      | <b>State Register 12:</b> The CLEANWIN register contains the number of windows that can be used by the SAVE instruction without causing a <i>clean_window</i> exception.   |
| Other Window Sets<br><br>OTHERWIN         | 0 to 7      | <b>State Register 13:</b> The OTHERWIN register contains the count of register sets that will be spilled/filled by a separate set of trap vectors based on the contents of WSTATE_OTHER. If OTHERWIN is zero, register sets are spilled/filled by use of trap vectors based on the contents of WSTATE_NORMAL. The OTHERWIN register can be used to split the register sets among different address spaces and handle spill/fill traps efficiently by use of separate spill/fill vectors. |

---

**Note** – The CWP, CANSAVE, CANRESTORE, OTHERWIN, and CLEANWIN registers contain values in the range 0 to 7 or 0 to 6 as indicated in TABLE 6-17. The effect of writing a value greater than indicated to any of these registers is undefined. The values programmed into these registers must combine into a consistent set of numbers that will work.

---



---

**Note** – The most significant 61 bits of all these registers are set to zero. When any are written, the most significant 61 bits are ignored.

---



---

**Compatibility Note** – The following differences between the SPARC V8 and SPARC V9 architectures are visible only to privileged software; they are invisible to non-privileged software.

1. In the SPARC V9 architecture, SAVE increments CWP and RESTORE decrements CWP. In the SPARC V8 architecture, the opposite is true: SAVE decrements PSR.CWP and RESTORE increments PSR.CWP.
  2. PSR.CWP in the SPARC V8 architecture is changed on each trap. In the SPARC V9 architecture, CWP is affected only by a trap caused by a window fill or spill exception.
-

## Clean Windows (CLEANWIN) Register Note

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

---

**Programming Note** – CLEANWIN must never be set to a value greater than six. Setting CLEANWIN greater than six would violate the register window state definition. Notice that the hardware does not enforce this restriction; it is up to Supervisor software to keep the window state consistent.

---

## 6.8.7 Window State (WSTATE) Privileged Register 14

The WSTATE register, shown in FIGURE 6-23, specifies bits that are inserted into  $TT_{TL}<4:2>$  on traps caused by window spill and fill exceptions.

This register is read/write by using the RDPR and WRPR privileged instructions.

These bits are used to select one of eight different window spill and fill handlers. If OTHERWIN = 0 at the time a trap is taken because of a window spill or window fill exception, then the WSTATE.NORMAL bits are inserted into  $TT[TL]$  field of the Trap Vector Address. Otherwise, the WSTATE.OTHER bits are inserted into  $TT[TL]$ .





FIGURE 6-23 WSTATE Register

## 6.8.8 Version (VER) Privileged Register 31

The version register, shown in FIGURE 6-24, specifies the fixed parameters pertaining to a particular processor implementation and mask set.

The VER register is read-only, readable by the RDPR privileged instruction.



**FIGURE 6-24** Version Register

## VER.manuf Field

The `VER.manuf` field contains our 16-bit manufacturer code, 003E<sub>16</sub>, which is our JEDEC semiconductor manufacturer code.

## VER.impl Field

The `VER.impl` field uniquely identifies the processor implementation or class of software-compatible implementations of the architecture. TABLE 6-18 shows the processor implementation codes.

**TABLE 6-18** Processor Implementation Codes

| Processor         | VER.impl           |
|-------------------|--------------------|
| UltraSPARC I      | 0010 <sub>16</sub> |
| UltraSPARC II     | 0011 <sub>16</sub> |
| UltraSPARC Ili    | 0012 <sub>16</sub> |
| UltraSPARC Iie    | 0013 <sub>16</sub> |
| UltraSPARC III Cu | 0015 <sub>16</sub> |

## VER.mask Field

The `VER.mask` specifies the current mask set revision and is chosen by the implementor. It generally increases numerically with successive releases of the processor but does not necessarily increase by one for consecutive releases. TABLE 6-19 lists the UltraSPARC III Cu Processor Mask Version.

**TABLE 6-19** UltraSPARC III Cu Processor Mask Version Codes

| Mask Version | VER.mask |
|--------------|----------|
| TO_1.x       | 4'h1     |
| TO_2.x       | 4'h2     |

## VER.maxt1 Field

The `VER.maxt1` value, 5, is the maximum number of trap levels supported by the processor.

## VER.maxwin Field

The `VER.maxwin` value, 7, is the maximum number of Integer Unit register windows that access the `NWINDOWS = 8` window register sets.

---

# 6.9 Special Access Register

## 6.9.1 Floating-Point Status Register (FSR)

The FSR register fields, illustrated in FIGURE 6-24, contain FPU mode and status information. Section 6.6.6, “Floating-Point Registers State (FPRS) Register 6” presents state information about the FPU.

The FSR is accessible using special load and store opcodes. They work in privileged and non-privileged mode. The lower 32 bits of the FSR are read and written by the `STFSRD` and `LDFSRD` floating-point instructions; all 64 bits of the FSR are read and written by the `STXFSR` and `LDXFSR` floating-point instructions, respectively. FIGURE 6-25 illustrates the FSR fields.

The `ver`, `ftt`, and `reserved (“—”)` fields are not modified by `LDFSR` or `LDXFSR`, which are read-only fields.




FIGURE 6-25 FSR Fields Reserved Bits

## Reserved Bits

Bits 63–38, 29–28, 21–20, and 12 are reserved. When read by an STXFSR instruction, these bits shall read as zero. Software should issue LDXFSR instructions only with zero values in these bits, unless the values of these bits are exactly those derived from a previous STXFSR.

The subsections on pages page 123 through page 131 describe the remaining fields in the FSR.

### 6.9.1.1 FSR\_fp\_condition\_codes (fcc0, fcc1, fcc2, fcc3)

The four sets of floating-point condition code fields are labeled `fcc0`, `fcc1`, `fcc2`, and `fcc3`.

---

**Compatibility Note** – SPARC V9 architecture’s `fcc0` is the same as SPARC V8 architecture’s `fcc`.

---

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

In TABLE 6-20,  $f_{rs1}$  and  $f_{rs2}$  correspond to the single, double, or quad values in the floating-point registers specified by a floating-point compare instruction’s `rs1` and `rs2` fields. The question mark (?) indicates an unordered relation, which is true if either  $f_{rs1}$  or  $f_{rs2}$  is a signalling NaN (SNaN) or a quiet NaN (QNaN). If FCMP or FCMPE generates a `fp_exception_ieee_754` exception, then `fccn` is unchanged. TABLE 6-20 shows the floating-point condition codes Fields of FSR.

**TABLE 6-20** Floating-Point Condition Codes (`fccn`) Fields of FSR

| Content of <code>fccn</code> | Indicated Relation                       |
|------------------------------|--|
| 0                            | $f_{rs1} = f_{rs2}$                      |
| 1                            | $f_{rs1} < f_{rs2}$                      |
| 2                            | $f_{rs1} > f_{rs2}$                      |
| 3                            | $f_{rs1} ? f_{rs2}$ ( <i>unordered</i> ) |

### 6.9.1.2 FSR\_rounding\_direction (RD)

Bits 31 and 30 select the rounding direction for floating-point results according to IEEE Std 754-1985. TABLE 6-21 shows the rounding direction fields.

**TABLE 6-21** Rounding Direction (RD) Field of FSR

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

If `GSR.IM = 1`, then the value of `FSR.RD` is ignored and floating-point results are instead rounded according to `GSR.IRND`.

### 6.9.1.3 FSR\_nonstandard\_fp (NS)

The NS bit allows the processor to flush a subnormal floating-point value to zero. If a floating-point add/subtract operation results in a subnormal value and `FSR.NS = 1`, the value is replaced by a floating-point zero value of the same sign. This replacement is usually performed in hardware. However, for the following cases when a subnormal value is generated in the course of the instruction and `FSR.NS = 1`, a *fp\_exception\_other* exception with `FSR.ftt = 2` (*unfinished\_FPop*) is taken and trap handler software is expected to replace the subnormal value with a zero value of the appropriate sign:

- FADD of numbers with opposite signs
- FSUB of numbers with the same signs
- FDTOS

The effects of `FSR.NS = 1` are as follows:

- If a floating-point source operand is subnormal, it is replaced by a floating-point zero value of the same sign (instead of causing an exception).
- If a floating-point operation generates a subnormal value, the value is replaced with a floating-point zero value of the same sign.
- This is implemented by performing the replacement in hardware, and sometimes cause a *fp\_exception\_other* exception with `FSR.ftt = 2` (*unfinished\_FPop*) so that trap handler software can perform the replacement.

If `GSR.IM = 1`, then the value of `FSR.NS` is ignored and the processor operates as if `FSR.NS = 0`.



### 6.9.1.4 FSR\_version (*ver*)

For the UltraSPARC III family of processors, the value in `FSR.ver` is zero.

Version number 7 is reserved to indicate that no hardware floating-point controller is present.

The `ver` field is read-only; it cannot be modified by the `LDFSR` and `LDXFSR` instructions.

### 6.9.1.5 FSR\_floating-point\_trap\_type (*ftt*)

When a floating-point exception trap occurs, `ftt` (bits 16 through 14 of the `FSR`) identifies the cause of the exception, the “floating-point trap type.” Several conditions can cause a floating-point exception trap. After a floating-point exception occurs, the `ftt` field encodes the type of the floating-point exception until an `STFSR` or `FPop` is executed.

The `ftt` field can be read by the `LDFSR` and `LDXFSR` instructions. The `STFSR` and `STXFSR` instructions do not affect `ftt` because this field is read-only.

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

---

**Programming Note** – Neither `LDFSR` nor `LDXFSR` can be used for the purpose of clearing `ftt`, since both leave `ftt` unchanged. However, executing a non-trapping `FPop`, such as “`fmovs %f0, %f0,`” prior to returning to non-privileged mode will zero `ftt`. The `ftt` remains valid until the next `FPop` instruction completes execution.

---

The `ftt` field encodes the floating-point trap type according to TABLE 6-22.

---

**Note** – The value “7” is reserved for future expansion.

---

**TABLE 6-22** Floating-Point Trap Type (`ftt`) Field of `FSR`

| <code>ftt</code> | Trap Type                 | Trap Vector                  |
|------------------|---------------------------|------------------------------|
| 0                | None                      | No trap taken                |
| 1                | <i>IEEE_754_exception</i> | <i>fp_exception_ieee_754</i> |
| 2                | <i>unfinished_FPop</i>    | <i>fp_exception_other</i>    |
| 3                | <i>unimplemented_FPop</i> | <i>fp_exception_other</i>    |
| 4                | <i>sequence_error</i>     | Reserved, unimplemented      |

**TABLE 6-22** Floating-Point Trap Type (*ftt*) Field of FSR (*Continued*)

| <b>ftt</b> | <b>Trap Type</b>           | <b>Trap Vector</b>      |
|------------|----------------------------|-------------------------|
| 5          | <i>hardware_error</i>      | Reserved, unimplemented |
| 6          | <i>invalid_fp_register</i> | Reserved, unimplemented |
| 7          | <i>Reserved</i>            | Reserved, unimplemented |

*IEEE\_754\_exception*, *unfinished\_FPop*, and *unimplemented\_FPop* will likely arise occasionally in the normal course of computation and must be recoverable by system software.

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

1. The value of *aexc* is unchanged. See Section 6.9.1.6 for details of *aexc*.
2. The value of *cexc* is unchanged, except for an *IEEE\_754\_exception*, where a bit corresponding to the trapping exception is set. The *unfinished\_FPop*, *unimplemented\_FPop*, *sequence\_error*, and *invalid\_fp\_register* floating-point trap types do not affect *cexc*. See Section 6.9.1.6 for details of *cexc*.
3. The source and destination registers are unchanged.
4. The value of *fccn* is unchanged.

The foregoing describes the result seen by a user trap handler if an IEEE exception is signalled, either immediately from an *IEEE\_754\_exception* or after recovery from an *unfinished\_FPop* or *unimplemented\_FPop*. In either case, *cexc* as seen by the trap handler reflects the exception causing the trap.

In the cases of *fp\_exception\_other* exceptions with *unfinished\_FPop* or *unimplemented\_FPop* trap types that do not subsequently generate IEEE traps, the recovery software should define *cexc*, *aexc*, and the destination registers or *fccs*, as appropriate.

***ftt = IEEE\_754\_exception.*** The *IEEE\_754\_exception* floating-point trap type indicates the occurrence of a floating-point exception conforming to IEEE Std 754-1985. The exception type is encoded in the *cexc* field.

The *aexc* and *fccs* fields and the destination *f* register are not affected by an *IEEE\_754\_exception* trap.

***ftt = unfinished\_FPop.*** The *unfinished\_FPop* floating-point trap type indicates that the processor was unable to generate correct results or that exceptions as defined by IEEE Std 754-1985 have occurred. Where exceptions have occurred, the *cexc* field is unchanged.

The conditions under which a *fp\_exception\_other* exception with floating-point trap type of *unfinished\_FPop* can occur are implementation dependent. The standard (recommended) set of conditions is listed in TABLE 6-23. An implementation may cause *fp\_exception\_other* with *unfinished\_FPop* under a different (but specified) set of conditions.

**TABLE 6-23** Standard Conditions in Which *unfinished\_FPop* Trap Type Can Occur

| FPU Operation | 1 Subnormal (SBN) Operand<br>IM = 1 or NS = 0 | 2 Subnormal (SBN) Operands<br>IM = 1 or NS = 0 | Result/Non-SBN Operand<br>IM = 1 or NS = 0           |
|---------------|---|--|--|
| <i>fadds</i>  | Unfinished trap                               | Unfinished trap                                | fi fv, fu, sbn (IM = NS = x)<br>NaN (either operand) |
| <i>fsubs</i>  | Unfinished trap                               | Unfinished trap                                | fi fv, fu, sbn (IM = NS = x)<br>NaN (either operand) |
| <i>faddd</i>  | Unfinished trap                               | Unfinished trap                                | fi fv, fu, sbn (IM = NS = x)<br>NaN (either operand) |
| <i>fsubd</i>  | Unfinished trap                               | Unfinished trap                                | fi fv, fu, sbn (IM = NS = x)<br>NaN (either operand) |
| <i>fmuls</i>  | Unfinished trap if<br>– result not zero       | Unfinished trap<br>– result not zero           | $-25 < Er \leq 1$                                    |
| <i>fdivs</i>  | Unfinished trap                               | Unfinished trap                                | $-25 < Er \leq 1$                                    |
| <i>fsmuld</i> | Unfinished trap                               | Unfinished trap                                | None   |
| <i>fmuld</i>  | Unfinished trap if<br>– result not zero       | Unfinished trap if<br>– result not zero        | $-54 < Er \leq 1$                                    |
| <i>fdivd</i>  | Unfinished trap                               | Unfinished trap                                | $-54 < Er \leq 1$                                    |
| <i>fsqrts</i> | Unfinished trap                               | N/A  | None   |
| <i>fsqrtd</i> | Unfinished trap                               | N/A  | None   |
| <i>fstoi</i>  | Unfinished trap                               | N/A  | $-2^{31} \leq \text{res} < 2^{31}$ , Infinity, NaN   |
| <i>fdtoi</i>  | Unfinished trap                               | N/A  | $-2^{31} \leq \text{res} < 2^{31}$ , Infinity, NaN   |
| <i>fstox</i>  | Unfinished trap                               | N/A  | $ \text{result}  \geq -2^{52}$ , Infinity, NaN       |
| <i>fdtox</i>  | Unfinished trap                               | N/A  | $ \text{result}  \geq -2^{52}$ , Infinity, NaN       |
| <i>fitos</i>  | N/A   | N/A  | $-2^{22} \leq \text{operand} < 2^{22}$               |
| <i>fxtos</i>  | N/A   | N/A  | $-2^{22} \leq \text{operand} < 2^{22}$               |
| <i>fitod</i>  | N/A   | N/A  | None   |
| <i>fxtod</i>  | N/A   | N/A  | $-2^{51} \leq \text{operand} < 2^{51}$               |

**TABLE 6-23** Standard Conditions in Which *unfinished\_FPop* Trap Type Can Occur (Continued)

| FPU Operation | 1 Subnormal (SBN) Operand<br>IM = 1 or NS = 0 | 2 Subnormal (SBN) Operands<br>IM = 1 or NS = 0 | Result/Non-SBN Operand<br>IM = 1 or NS = 0 |
|---------------|---|--|--|
| FSTOD         | Unfinished trap                               | N/A  | NaN  |
| FDTOS         | Unfinished trap                               | N/A  | fi fv, fu, sbn (IM = NS = x), NaN          |

Note:Er <- Biased Exponent of the result before rounding  
 Ei <- Biased Exponent of input operand  
 fi <- Invalid(Infinity - Infinity, Infinity\*0, 0/0, Infinity/Infinity)  
 fv <- OverflowEr >= 2047(DP) or 255(SP) but not exact infinity  
 fu <- Underflow0 < |result| < 2<sup>-1022</sup>(DP) or 2<sup>-126</sup>(SP)  
 sbnormal(sbn): |number| = 2<sup>-1022</sup> \* (significand x 2<sup>-52</sup>) (DP) or 2<sup>-126</sup> \* (significand x 2<sup>-23</sup>) (SP)  
 {-54 < Er < 1 (DP) or -25 < Er < 1 (SP)}

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


All quad FPOps variations set *ftt = unimplemented\_FPop*.

### 6.9.1.6 Floating-Point Exceptions Control and Status

There are three FSR register fields used to control and status the events associated with floating-point exceptions.

#### *FSR\_trap\_enable\_mask (TEM)*

Bits 27 through 23 are enable bits for each of the five IEEE-754 floating-point exceptions that can be indicated in the *current\_exception* field (*cexc*). See FIGURE 6-26 for an illustration. If a floating-point operate instruction generates one or more exceptions and the TEM bit corresponding to any of the exceptions is one, then this condition causes a *fp\_exception\_ieee\_754* trap. A TEM bit value of zero prevents the corresponding exception type from generating a trap.




**FIGURE 6-26** Trap Enable Mask (TEM) Fields of FSR

### *FSR\_accrued\_exception (aexc)*

Bits 9 through 5 accumulate IEEE\_754 floating-point exceptions as long as floating-point exception traps are disabled through the TEM field. See FIGURE 6-27 for an illustration. After an FPop completes with `ftt = 0`, the TEM and `cexc` fields are logically ANDed together. If the result is nonzero, `aexc` is left unchanged and a `fp_exception_ieee_754` trap is generated; otherwise, the new `cexc` field is ORED into the `aexc` field and no trap is generated. Thus, while (and only while) traps are masked, exceptions are accumulated in the `aexc` field.


This field is also written with the appropriate value when an LDFSR or LDXFSR instruction is executed.



**FIGURE 6-27** Accrued Exception Bits (`aexc`) Fields of FSR

### *FSR\_current\_exception (cexc)*

Bits 4 through 0 indicate that one or more IEEE\_754 floating-point exceptions were generated by the most recently executed FPop instruction. The absence of an exception causes the corresponding bit to be cleared. See FIGURE 6-28 for an illustration.



**FIGURE 6-28** Current Exception Bits (`cexc`) Fields of FSR

---

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

---

The `cexc` bits are set as described in Section 6.9.1.7, “Floating-Point Exception Fields” by the execution of an FPop that either does not cause a trap or causes a `fp_exception_ieee_754` exception with `FSR.ftt = IEEE_754_exception`. An `IEEE_754_exception` that traps shall cause exactly one bit in `FSR.cexc` to be set, corresponding to the detected IEEE Std 754-1985 exception.

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

- An IEEE 754 overflow condition (*of*) occurs:
  - If OFM = 0 and NXM = 0, the cexc.ofc and cexc.nxc bits are both set to one, the other three bits of cexc are set to zero, and a *fp\_exception\_ieee\_754* trap does *not* occur.
  - If OFM = 0 and NXM = 1, the cexc.nxc bit is set to one, the other four bits of cexc are set to zero, and a *fp\_exception\_ieee\_754* trap *does* occur.
  - If OFM = 1, the cexc.ofc bit is set to one, the other four bits of cexc are set to zero, and a *fp\_exception\_ieee\_754* trap *does* occur.
- An IEEE 754 underflow condition (*uf*) occurs:
  - If UFM = 0 and NXM = 0, the cexc.ufc and cexc.nxc bits are both set to one, the other three bits of cexc are set to zero, and a *fp\_exception\_ieee\_754* trap does *not* occur.
  - If UFM = 0 and NXM = 1, the cexc.nxc bit is set to one, the other four bits of cexc are set to zero, and a *fp\_exception\_ieee\_754* trap *does* occur.
  - If UFM = 1, the cexc.ufc bit is set to one, the other four bits of cexc are set to zero, and a *fp\_exception\_ieee\_754* trap *does* occur.

The behavior is summarized in TABLE 6-24 (where “x” indicates “don’t care”):

**TABLE 6-24** Setting of FSR.cexc Bits

| Exception(s) Detected in FP Operation |    |    | Trap Enable Mask bits (in FSR.TEM) |     |     | <i>fp_exception_ieee_754</i> Trap Occurs? | Current Exception bits (in FSR.cexc) |     |     | Notes |
|---------------------------------------|----|----|------------------------------------|-----|-----|---|--------------------------------------|-----|-----|-------|
| of                                    | uf | nx | OFM                                | UFM | NXM |   | ofc                                  | ufc | nxc |       |
| -                                     | -  | -  | x                                  | x   | x   | No  | 0                                    | 0   | 0   |       |
| -                                     | -  | 1  | x                                  | x   | 0   | No  | 0                                    | 0   | 1   |       |
| -                                     | 1  | 1  | x                                  | 0   | 0   | No  | 0                                    | 1   | 1   | (1)   |
| 1                                     | -  | 1  | 0                                  | x   | 0   | No  | 1                                    | 0   | 1   | (2)   |
| -                                     | -  | 1  | x                                  | x   | 1   | Yes                                       | 0                                    | 0   | 1   |       |
| -                                     | 1  | 1  | x                                  | 0   | 1   | Yes                                       | 0                                    | 0   | 1   |       |
| -                                     | 1  | -  | x                                  | 1   | x   | Yes                                       | 0                                    | 1   | 0   |       |
| -                                     | 1  | 1  | x                                  | 1   | x   | Yes                                       | 0                                    | 0   | 0   |       |
| 1                                     | -  | 1  | 1                                  | x   | x   | Yes                                       | 1                                    | 0   | 0   | (2)   |
| 1                                     | -  | 1  | 0                                  | x   | 1   | Yes                                       | 0                                    | 0   | 1   | (2)   |

Notes:  
 (1) When the underflow trap is disabled (UFM = 0), underflow is always accompanied by inexact.  
 (2) Overflow is always accompanied by inexact.

If the execution of an FPop causes a trap other than *fp\_exception\_ieee\_754*, *FSR.cexc* is left unchanged.

### 6.9.1.7 Floating-Point Exception Fields

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

#### *FSR\_invalid (nvc, nva)*

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

#### *FSR\_overflow (ofc, ofa)*

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

#### *FSR\_underflow (ufc, ufa)*

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

Underflow is never indicated when the correct unrounded result is zero. Otherwise:

- If UFM = 0, underflow occurs if a nonzero result is tiny and a loss of accuracy occurs.
- If UFM = 1, underflow occurs if a nonzero result is tiny.

SPARC V9 allows underflow to be detected either before or after rounding. The UltraSPARC III Cu processor detects underflow before rounding.

#### *FSR\_division-by-zero (dzc, dza)*

$X \div 0.0$ , where  $X$  is subnormal or normalized; 1 = division by zero, 0 = no division by zero.

---

**Note** –  $0.0 \div 0.0$  does not set the *dzc* or *dza* bits.

---

### *FSR\_inexact (nxc, nxa)*

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

---

**Programming Note** – Software must be capable of simulating the operation of the FPU in order to properly handle the *unimplemented\_FPop*, *unfinished\_FPop*, and *IEEE\_754\_exception* floating-point trap types. Thus, a user application program always sees an FSR that is fully compliant with IEEE Std 754-1985.

---

## 6.10 ASI Mapped Registers

In this section, we describe the Data Cache Unit Control Register and Data Watchpoint registers (virtual address data watchpoint and physical address data watchpoint).

### 6.10.1 Data Cache Unit Control Register (DCUCR)

ASI 45<sub>16</sub> (*ASI\_DCU\_CONTROL\_REGISTER*), VA = 0<sub>16</sub>

The DCUCR contains fields that control several memory related hardware functions. The functions include instruction, prefetch, write and data caches, MMUs, and watchpoint setting.

After a power-on reset (POR), all fields of DCUCR are set to zero. After a WDR, XIR, or SIR, all fields of DCUCR defined in this section are set to zero.

The DCUCR is illustrated in FIGURE 6-29 and described in TABLE 6-25. In the table, the field definitions and bits are grouped by function rather than by strict bit sequence.



|    |    |    |    |    |    |     |     |    |    |    |    |    |    |    |    |    |    |    |    |    |   |   |   |   |
|----|----|----|----|----|----|-----|-----|----|----|----|----|----|----|----|----|----|----|----|----|----|---|---|---|---|
| —  | CP | CV | ME | RE | PE | HPE | SPE | SL | WE | PM | VM | PR | PW | VR | VW | —  | DM | IM | DC | IC |   |   |   |   |
| 63 | 50 | 49 | 48 | 47 | 46 | 45  | 44  | 43 | 42 | 41 | 40 | 33 | 32 | 25 | 24 | 23 | 22 | 21 | 20 | 4  | 3 | 2 | 1 | 0 |

**FIGURE 6-29** DCU Control Register Access Data Format (ASI 45<sub>16</sub>)

**TABLE 6-25** DCUCR Bit Field Descriptions (1 of 3)

| Bits                       | Field    | Type | Description   | Note |
|----------------------------|----------|------|---|------|
| 63:50,<br>20:4             | reserved | RW   |   |      |
| <b>MMU Control</b>         |          |      |   |      |
| 49                         | CP       | RW   | <i>Cacheability of PA.</i> CP determines the physical cacheability of memory accesses when the I-MMU or D-MMU is disabled (IM = 0 or DM = 0). The TTE.E (side-effect) bit is set to the complement of CP when MMUs are enabled; 1 = cacheable, 0 = non-cacheable.                                 | 1    |
| 48                         | CV       | RW   | <i>Cacheability of VA.</i> CV determines the virtual cacheability of memory accesses when the D-MMU is disabled (DM = 0); 1 = cacheable, 0 = non-cacheable.   |      |
| 3                          | DM       |      | <i>D-MMU Enable.</i> If DM = 0, the D-MMU is disabled (pass-through mode). <b>Note:</b> When the MMU/TLB is disabled, a virtual address is passed through as a physical address.  |      |
| 2                          | IM       |      | <i>I-MMU Enable.</i> If IM = 0, the I-MMU is disabled (pass-through mode).  |      |
| <b>Store Queue Control</b> |          |      |   |      |
| 47                         | ME       | RW   | <i>Non-cacheable Store Merging Enable.</i> If cleared, no merging of non-cacheable, non-side-effect store data will occur. Each non-cacheable store will generate a system bus transaction.   |      |
| 46                         | RE       |      | <i>RAW Bypass Enable.</i> If cleared, no bypassing of data from the store queue to a dependent load instruction will occur. All load instructions will have their RAW predict field cleared.  |      |
| <b>Prefetch Control</b>    |          |      |   | 2    |
| 45                         | PE       |      | <i>Prefetch Cache Enable.</i> If prefetch is disabled by clearing the PE bit, all references to the P-cache are handled as P-cache misses. If cleared, the P-cache does not generate any hardware prefetch requests to the L2-cache. Software prefetch instructions are not affected by this bit. |      |
| 44                         | HPE      |      | <i>Prefetch Cache Hardware Prefetch Enable.</i>   | 3    |
| 43                         | SPE      |      | <i>Software Prefetch Enable.</i> Clear to disable prefetch instructions. When disabled, software prefetch instructions do not generate a request to the L2-cache or the system interface. They will continue to be issued to the pipeline, where they will be treated as NOPs.                    |      |

TABLE 6-25 DCUCR Bit Field Descriptions (2 of 3)

| Bits   | Field   | Type | Description  | Note |
|--|---------|------|--|------|
| <b>Second Load Control</b>                   |         |      |  |      |
| 42   | SL      |      | <i>Second Load Steering Enable.</i> If cleared, all load type instructions will be steered to the MS pipeline and no floating-point load type instructions will be issued to the A0 or A1 pipelines.   |      |
| <b>I-cache, D-cache, and W-cache Control</b> |         |      |  |      |
| 41   | WE      |      | <i>Write Cache Enable.</i> If zero, all W-cache references will be handled as W-cache misses. Each store queue entry will perform an RMW transaction to the L2-cache, and the W-cache will be maintained in a clean state. Software is required to flush the W-cache (force it to a clean state) before setting this bit to zero.  |      |
| 1  | DC      |      | <i>Data Cache Enable.</i> The DC is used to enable/disable the operation of the data cache closest to the processor (D-cache); DC = 1 enables the D-cache and DC = 0 disables it. When DC = 0, memory accesses (loads, stores, atomic load-stores) are satisfied by caches lower in the cache hierarchy.<br>When the data cache is disabled, its contents are not updated. When the D-cache is re-enabled, any D-cache lines still marked as “valid” may be inconsistent with the state of memory or other caches. In that case, software must handle any inconsistencies by flushing the inconsistent lines from the D-cache.     |      |
| 0  | IC      |      | <i>Instruction Cache Enable.</i> The IC is used to enable/disable the operation of the instruction cache closest to the processor (I-cache); IC = 1 enables the I-cache and IC = 0 disables it. When IC = 0, instruction fetches are satisfied by caches lower in the cache hierarchy.<br>When the instruction cache is disabled, its contents are not updated. When the I-cache is re-enabled, any I-cache lines still marked as “valid” may be inconsistent with the state of memory or other caches. In that case, software must handle any inconsistencies by invalidating the inconsistent lines in the I-cache.              |      |
| <b>Watchpoint Control</b>                    |         |      |  |      |
| 40:33  | PM<7:0> |      | <i>DCU Physical Address Data Watchpoint Mask.</i> The Physical Address Data Watchpoint Register contains the physical address of a 64-bit word to be watched. The 8-bit Physical Address Data Watch Point Mask controls which byte(s) within the 64-bit word should be watched. If all eight bits are cleared, the physical watchpoint is disabled. If the watchpoint is enabled and a data reference overlaps any of the watched bytes in the watchpoint mask, then a physical watchpoint trap is generated. Watchpoint behavior for a Partial Store instruction may differ.<br>Please see the VM field description in the table. | 4    |

TABLE 6-25 DCUCR Bit Field Descriptions (3 of 3)

| Bits                        | Field   | Type | Description   | Note                        |   |                  |                     |                  |           |                  |           |                  |           |   |
|-----------------------------|---|------|---|-----------------------------|---|------------------|---------------------|------------------|-----------|------------------|-----------|------------------|-----------|---|
| 32:25                       | VM<7:0>   |      | <p><i>DCU Virtual Address Data Watchpoint Mask.</i> The Virtual Address Data Watchpoint Register contains the virtual address of a 64-bit word to be watched. This 8-bit mask controls which byte(s) within the 64-bit word should be watched. If all eight bits are cleared, then the virtual watchpoint is disabled. If watchpoint is enabled and a data reference overlaps any of the watched bytes in the watchpoint mask, then a virtual watchpoint trap is generated.</p> <p>VA/PA data watchpoint byte mask examples are shown below.</p> <table border="1" data-bbox="565 482 1222 680"> <thead> <tr> <th>Watchpoint Mask (PM and VM)</th> <th>Least Significant 3 Bits of Address of Bytes Watched<br/>7654 3210</th> </tr> </thead> <tbody> <tr> <td>00<sub>16</sub></td> <td>Watchpoint disabled</td> </tr> <tr> <td>01<sub>16</sub></td> <td>0000 0001</td> </tr> <tr> <td>32<sub>16</sub></td> <td>0011 0010</td> </tr> <tr> <td>FF<sub>16</sub></td> <td>1111 1111</td> </tr> </tbody> </table> | Watchpoint Mask (PM and VM) | Least Significant 3 Bits of Address of Bytes Watched<br>7654 3210 | 00 <sub>16</sub> | Watchpoint disabled | 01 <sub>16</sub> | 0000 0001 | 32 <sub>16</sub> | 0011 0010 | FF <sub>16</sub> | 1111 1111 | 4 |
| Watchpoint Mask (PM and VM) | Least Significant 3 Bits of Address of Bytes Watched<br>7654 3210 |      |   |                             |   |                  |                     |                  |           |                  |           |                  |           |   |
| 00 <sub>16</sub>            | Watchpoint disabled   |      |   |                             |   |                  |                     |                  |           |                  |           |                  |           |   |
| 01 <sub>16</sub>            | 0000 0001   |      |   |                             |   |                  |                     |                  |           |                  |           |                  |           |   |
| 32 <sub>16</sub>            | 0011 0010   |      |   |                             |   |                  |                     |                  |           |                  |           |                  |           |   |
| FF <sub>16</sub>            | 1111 1111   |      |   |                             |   |                  |                     |                  |           |                  |           |                  |           |   |
| 24, 23                      | PR, PW  |      | <p><i>DCU Physical Address Data Watchpoint Enable.</i> If PR (PW) is one, then a data read (write) that matches the range of addresses in the Physical Watchpoint Register causes a watchpoint trap. If both PR and PW are set, a watchpoint trap will occur on either a read or write access.</p>  |                             |   |                  |                     |                  |           |                  |           |                  |           |   |
| 22, 21                      | VR, VW  |      | <p><i>DCU Virtual Address Data Watchpoint Enable.</i> If VR (VW) is one, then a data read (write) that matches the range of addresses in the Virtual Watchpoint Register causes a watchpoint trap. If both VR and VW are set, a watchpoint trap will occur on either a read or write access.</p>  |                             |   |                  |                     |                  |           |                  |           |                  |           |   |

1. The CP and CV bits of DCUCR must be changed with care. It is recommended that a MEMBAR #Sync be executed before and after CP or CV is changed. Also, software must manage cache states to be consistent before and after CP or CV is changed.
2. Prefetch is enabled in the UltraSPARC III Cu processor. Both hardware prefetch and software prefetch for data to the P-cache are valid only for floating-point load instructions and are not valid for integer load instructions.
3. Both hardware prefetch and second load unit may not be enabled at the same time. Enabling both may cause incorrect program behavior.
4. Watchpoint exceptions on Partial Store instruction occur conservatively. The DCUCR.VM masks are only checked for nonzero value (watchpoint disabled). The byte store mask (r[rs2]) in the Partial Store instruction is ignored, and a watchpoint exception can occur even if the mask is zero (that is, no store will take place).

## 6.10.2 Data Watchpoint Registers

UltraSPARC III Cu processors implement “break-before” watchpoint traps. When the address of a data access matches a preset physical or virtual watchpoint address, instruction execution is stopped immediately before the watched memory location is accessed.

TABLE 6-26 lists ASIs that are affected by the two watchpoint traps.

**TABLE 6-26** ASIs Affected by Watchpoint Traps

| ASI Type             | ASI Range  | Data MMU  | Watchpoint If Matching VA | Watchpoint If Matching PA |
|----------------------|--|-----------|---------------------------|---------------------------|
| Translating ASIs     | $04_{16}-11_{16}$ , $18_{16}-19_{16}$ , $24_{16}-2C_{16}$ ,<br>$70_{16}-71_{16}$ , $78_{16}-79_{16}$ , $80_{16}-FF_{16}$ | On<br>Off | Y<br>N                    | Y<br>Y                    |
| Bypass ASIs          | $14_{16}-15_{16}$ , $1C_{16}-1D_{16}$  | —         | N                         | Y                         |
| Non-translating ASIs | $30_{16}-6F_{16}$ , $72_{16}-77_{16}$ , $7A_{16}-7F_{16}$  | —         | N                         | N                         |

For 128-bit (quad) atomic load and 64-byte block load and store instructions, a watchpoint trap is generated only if the watchpoint overlaps the lowest address eight bytes of the access.

To avoid trapping infinitely, software should emulate the instruction that caused the trap and return from the trap by using a DONE instruction or turn off the watchpoint before returning from a watchpoint trap handler.

Two 64-bit data watchpoint registers provide the means to monitor data accesses during program execution. When Virtual/Physical Data Watchpoint is enabled, the virtual/physical addresses of all data references are compared against the content of the corresponding watchpoint register. If a match occurs, a *VA\_watchpoint* or *PA\_watchpoint* trap is signalled before the data reference instruction is completed. The virtual address watchpoint trap has higher priority than the physical address watchpoint trap.

Separate 8-bit byte masks allow watchpoints to be set for a range of addresses. Each zero bit in the byte mask causes the comparison to ignore the corresponding byte in the address. These watchpoint byte masks and the watchpoint enable bits reside in the DCUCR.


### *Virtual Address Data Watchpoint Register*

ASI  $58_{16}$ , VA =  $38_{16}$

Name: VA Data Watchpoint Register

FIGURE 6-30 illustrates the Virtual Address Watchpoint Register,

**where:** **DB\_VA** is the most significant 61 bits of the 64-bit virtual data watchpoint address.



**FIGURE 6-30** VA Data Watchpoint Register Format


### *Physical Address Data Watchpoint Register*

ASI  $58_{16}$ , VA= $40_{16}$

Name: PA Data Watchpoint Register

FIGURE 6-31 illustrates the PA Data Watchpoint Register,

**where:** **DB\_PA** is the most significant 61 bits of the physical data watchpoint address. The width of an UltraSPARC III Cu physical address is 43 bits.



**FIGURE 6-31** PA Data Watchpoint Register Format

---

**Compatibility Note** – The UltraSPARC III Cu processor supports a 43-bit physical address space. Software is responsible for writing a zero-extended 64-bit address into the PA Data Watchpoint register.

---

### *Data Watchpoint Reliability*

The processor supports watchpoint comparison on the MS (memory) pipeline; any second issue (Ax pipeline) floating-point loads will not trigger a watchpoint. For reliable use of the watchpoint mechanism, the second floating-point load feature must be disabled using DCUCR.SL.



# Instruction Types

---

Instructions are accessed by the processor from memory and are executed, annulled, or trapped. Instructions are discussed in seven general categories. The processor instructions are described in the following sections:

## Learning the Instructions

- Introduction
- Memory Addressing for Load and Store Instructions
- Integer Execution Environment
- Floating-Point Execution Environment
- VIS Execution Environment
- Data Coherency Instructions
- Register Window Management Instructions
- Program Control Transfer Instructions
- Trap Base Address (TBA) Register
- Prefetch Instructions

## Reference Section

- Instruction Summary Table by Category
- Integer Execution Environment Instructions
- Floating-Point Execution Environment Instructions
- VIS Execution Environment Instructions
- Data Coherency Instructions
- Register-Window Management Instructions
- Program Control Transfer Instructions

- Data Prefetch Instructions
- Instruction Formats and Fields
- Reserved Opcodes and Instruction Fields
- Big/Little-endian Addressing

---

## 7.1 Introduction

The processor's RISC architecture is defined primarily by the SPARC V9 architecture. The UltraSPARC II processors were the first to extend the SPARC V9 architecture with new instructions and additional logic units. The UltraSPARC III Cu processor further extends this instruction execution environment.

The UltraSPARC III Cu processor provides backward compatibility for SPARC application programs. Upgraded system software is required. Noteworthy enhancements to the processor include greater capability in the execution units to improved instruction scheduling, new VIS instructions to reduce the length of code sequences, and data prefetch instructions to provide the compiler with ways to improve cache hit rates.

Our compiler and other software development tools take advantage of the new instruction features to increase parallel execution, reduce code size, and achieve shorter instruction execution latencies.

---

## 7.2 Memory Addressing for Load and Store Instructions

The SPARC V9 architecture uses big-endian byte order by default; the address of a quadword, doubleword, word, or halfword is the address of its most significant byte. Increasing the address means decreasing the significance of the unit being accessed. All instruction accesses are performed using big-endian byte order. The SPARC V9 architecture also can support little-endian byte order for data accesses only; the address of a quadword, doubleword, word, or halfword is the address of its least significant byte. Increasing the address means increasing the significance of the unit being accessed.



## 7.2.1 Integer Unit Memory Alignment Requirements

Halfword accesses are aligned on 2-byte boundaries; word accesses (which include instruction fetches) are aligned on 4-byte boundaries; extended-word and doubleword accesses are aligned on 8-byte boundaries. An improperly aligned address in a load, store, or load-store instruction causes a trap to occur, with possible exceptions.

---

**Programming Note** – By setting `i = 1` and `rs1 = 0`, you can access any location in the lowest or highest 4 KB of an address space without using a register to hold part of the address.

---

## 7.2.2 FP/VIS Memory Alignment Requirements

Extended word and doubleword (64-bit) accesses must be aligned on 8-byte boundaries, quadword accesses must be aligned on 16-byte boundaries, and Block load (BLD) and Block store (BST) accesses must be aligned on 64-byte boundaries.

All references are 32, 64, or 128 bits. They must be naturally aligned to their data width in memory except for double-precision floating-point values, which may be aligned on word boundaries. However, if so aligned, doubleword loads/stores may not be used to access them, resulting in less efficient and nonatomic accesses.

An improperly aligned address in a load, store, or load-store instruction causes a *mem\_address\_not\_aligned* exception to occur, with the following exceptions:

- A `LDDF` or `LDDFA` instruction accessing an address that is word aligned but not doubleword aligned causes a *LDDF\_mem\_address\_not\_aligned* exception.
- A `STDF` or `STDFA` instruction accessing an address that is word aligned but not doubleword aligned causes a *STDF\_mem\_address\_not\_aligned* exception.

## 7.2.3 Byte Order Addressing Conventions (Endianess)

The processor uses big-endian byte order for all instruction accesses and, by default, for data accesses. It is possible to access data in little-endian format by using load and store alternate instructions that support little-endian data structures. It is also possible to change the default byte order for implicit data accesses.

See Section 7.13, “Big/Little-endian Addressing” for details.

## 7.2.4 Address Space Identifiers

Versions of load/store instructions, the *load and store alternate* instructions, can specify an 8-bit address space identifier (ASI) to go along with the load/store data instruction.

The load and store alternate instructions have three sources of ASIs:

- Explicit immediate of instruction
- ASI Register reference
- Hardcode to the instruction

Supervisor software (privileged mode) uses ASIs to access special, protected registers, such as MMU, cache control, and processor state registers, and other processor or system-dependent values.

ASIs are also used to modify the function of many instructions. This overloading of load/store instructions provide partial store, block load/store, and atomic memory access operations.

Chapter 8, “Address Space Identifiers” describes the ASIs in more detail. The chapter summary table associates ASI values to specific instructions.

### *Implicit ASI Value*

Load and store instructions provide an implicit ASI value of ASI\_PRIMARY, ASI\_PRIMARY\_LITTLE, ASI\_NUCLEUS, or ASI\_NUCLEUS\_LITTLE. Load and store alternate instructions provide an explicit ASI, specified by the `imm_asi` instruction field when `i = 0`, or the contents of the ASI register when `i = 1`.

### *Privileged and Non-privileged ASIs*

ASIs  $00_{16}$  through  $7F_{16}$  are restricted; only privileged software is allowed to access them. An attempt to access a restricted ASI by non-privileged software results in a *privileged\_action* exception. ASIs  $80_{16}$  through  $FF_{16}$  are unrestricted; software is allowed to access them whether the processor is operating in privileged or non-privileged mode.

---

**Compatibility Note** – The SPARC V9 architecture provides the basic framework and defines the required ASIs for the processor. Other ASIs are defined (and sometimes redefined) for a specific processor or family of processors as allowed by the SPARC V9 architecture.

---

---

**Implementation Note** – The processor decodes all eight bits of each ASI specifier. In addition, the processors redefine certain ASIs as appropriate for a specific processor.

---

## 7.2.5 Maintaining Data Coherency

The processor’s memory architecture requires some software intervention to provide data coherency during program execution. These requirements are discussed in Chapter 9, “Memory Models” using the FLUSH and Section 7.6, “Data Coherency Instructions” describes MEMBAR instructions.

The two types of data coherency instructions are needed to flush the cache for self-modifying code and to write data buffers out to memory.

---

## 7.3 Integer Execution Environment

### 7.3.1 IU Data Access Instructions

Load, store, and atomic instructions are the only instructions that access memory. All the IU data access instructions, except the compare and store (CASx) use either two  $r$  registers or SIMM13, a signed 13-bit immediate value, to calculate a 64-bit, byte-aligned memory address. Compare and Swap uses a single  $r$  register to specify a 64-bit memory address. Section 7.4.2, “FPU/VIS Data Access Instructions” discusses floating-point register load and store instructions.

The CPU appends an ASI to the 64-bit address used with all the data access instructions.

---

**Note** – In addition to the large physical main memory, the processor has many memory mapped control, status, and diagnostic registers that are accessed using load and store instructions with an appropriate ASI value.

---

The destination field of the data access instruction specifies an  $r$  or  $f$  (single, double/extended, or quadword) register that supplies the data for a store or that receives the data from a load.

### 7.3.1.1 Load and Store Instructions

Integer load and store instructions support byte, halfword (16-bit), word (32-bit), and doubleword (64-bit) accesses. Some versions of integer load instructions perform sign extension on 8-, 16-, and 32-bit values as they are loaded into a 64-bit destination register.

### 7.3.1.2 Move Instruction

There is no explicit integer move instruction. A move instruction can be easily synthesized by adding, subtracting or ORing a zero with a register and pointing the result to another register. The zero can come as a register input (such as `%r0` that has a value zero in SPARC V9) or as an immediate input to the instruction.

### 7.3.1.3 Conditional Move Instructions

#### *Based on Integer (icc/xcc) and Floating-point (fcc) Condition Codes*

This subsection describes two instructions that copy the contents of one register to another register within the same register file: one instruction for moving within the integer register file and another for moving within the floating-point register file.

- MOVcc Instruction

If a specified `icc/xcc` or `fcc` condition is satisfied, then the `MOVcc` instruction copies the contents of any integer to a destination integer register.

- FMOVcc Instruction

If a specified `icc/xcc` or `fcc` condition is satisfied, then the `FMOVcc` instruction copies the contents of any floating-point register to a destination floating-point register.

(A similar set of conditional move instructions are based on an integer register value. These conditional move instructions are described in the next section).

The condition code to test is specified in the instruction and may be any of the conditions allowed in conditional delayed control transfer instructions. This condition is tested against 1 of the 6 sets of condition codes (`icc`, `xcc`, `fcc0`, `fcc1`, `fcc2`, and `fcc3`), as specified by the instruction.

For example:

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

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

The MOVCC and FMOVCC instructions can be used to eliminate some branches in programs. In most situations, branches will take more clock cycles than the MOVCC or FMOVCC instructions.

For example, the following C statement:

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

can be coded as

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

which eliminates the need for a branch.

### *Based on Integer Register Value*

There are separate versions for the IU and floating-point unit (FPU) register files:

- MOVr Instruction

If the contents of an integer register satisfy a specified condition, then the MOVr instruction copies the contents of any integer register to a destination integer register.

- FMOVr Instruction

If the contents of an integer register satisfy a specified condition, then the FMOVr instruction copies the contents of any floating-point register to a destination floating-point register.

The conditions to test are enumerated in TABLE 7-1.

**TABLE 7-1** MOVr and FMOVr Test Conditions

| Condition | Symbol | Description                   |
|-----------|--------|-------------------------------|
| NZ        | ≠ 0    | Nonzero                       |
| Z         | = 0    | Zero                          |
| LZ        | < 0    | Less than zero                |
| LEZ       | ≤ 0    | Less than or equal to zero    |
| GZ        | > 0    | Greater than zero             |
| GEZ       | ≥ 0    | Greater than or equal to zero |

Any of the integer registers may be tested for one of the conditions, and the result used to control the move. For example,

```
movrnz %i2, %l4, %l6
```

moves integer register %l4 to integer register %l6 if integer register %i2 contains a nonzero value.

MOV<sub>r</sub> and FMOV<sub>r</sub> can be used to eliminate some branches in programs or can emulate multiple unsigned condition codes by using an integer register to hold the result of a comparison.

### 7.3.1.4 Atomic Instructions

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

The SWAP and LDSTUB instructions can optionally access alternate space. (The CASA instruction always accesses alternate memory spaces). If the ASI specified for any alternate form of these instructions is a privileged ASI (value  $80_{16}$ ), then the processor must be in privileged mode to access it.

#### *Atomic Quad Load Instruction (LDDA with ASI xx)*

The atomic quad load instruction supplies an indivisible quadword (16-byte) load that is important in system software programs.

#### *Compare and Swap Atomic Instruction (CASA)*

An *r* register specifies the value that is compared with the value in memory at the computed address. CASA accesses words, and CASXA accesses doublewords.

If the values are equal (memory location and *r* register), then the destination field specifies the *r* register that is to be exchanged atomically with the addressed memory location.

If the values are unequal, then the destination field specifies the *r* register that was to receive the value at the addressed memory location; in this case, the addressed memory location remains unchanged.

#### *Swap Atomic Instruction (SWAP<sup>D</sup>)*

The destination register identifies the *r* register to be exchanged atomically with the calculated memory location. SWAP accesses words.

## *Load-Store Unsigned Byte (LDSTUB)*

The LDSTUB instruction reads a byte from memory and writes ones to the location read. LDSTUB accesses bytes.

## 7.3.2 IU Arithmetic Instructions

The integer arithmetic instructions are generally triadic register address instructions that compute a result of a function of two source operands. They either write the result into the destination register  $r[rd]$  or discard it. One of the source operands is always  $r[rs1]$ . The other source operand depends on the  $i$  bit in the instruction. If  $i = 0$ , then the operand is  $r[rs2]$ . If  $i = 1$ , then the operand is the immediate constant  $simm10$ ,  $simm11$ , or  $simm13$  sign-extended to 64 bits.

The arithmetic/logical/shift instructions perform arithmetic, tagged arithmetic, logical, and shift operations. One exception is the `SETHI` instruction that can be used in combination with another arithmetic or logical instruction to create a 32-bit constant in an  $r$  register.

### *Condition Codes*

Most integer arithmetic instructions have two versions: one sets the integer condition codes (`icc` and `xcc`) as a side-effect; the other does not affect the condition codes.

### 7.3.2.1 Integer Add and Subtract Instructions

Sixty-four bit arithmetic is performed on two  $r$  registers to generate a 64-bit result. The `icc` and `xcc` condition codes can optionally be set.

### 7.3.2.2 Tagged Integer Add and Subtract Instructions

The tagged arithmetic instructions assume that the least significant two bits of each operand are a data-type tag. These instructions set the integer condition code (`icc`) and extended integer condition code (`xcc`) overflow bits on 32-bit (`icc`) or 64-bit (`xcc`) arithmetic overflow.

Appendix A “Instruction Definitions” describes the tagged instructions.

If either of the two operands has a nonzero tag or if 32-bit arithmetic overflow occurs, tag overflow is detected. If tag overflow occurs, then `TADDCC` and `TSUBCC` set the `CCR.icc.V` bit; if 64-bit arithmetic overflow occurs, then they set the `CCR.xcc.V` bit.

The `xcc` overflow bit is not affected by the tag bits.

The trapping versions (TADDccTV, TSUBccTV) are deprecated. See Section A.70.16, “Tagged Add and Trap on Overflow” and Section A.70.17, “Tagged Subtract and Trap on Overflow” for details.

### 7.3.2.3 Integer Multiply and Divide Instructions

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

Some versions of the 32-bit multiply and divide instructions set the condition codes.

### 7.3.2.4 Set High 22 Bits of Low Word

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

### 7.3.2.5 Integer Shift Instructions

Shift logical instructions (SLL, SRL) shift an *r* register left or right by an immediate constant in the instruction or by the amount pre-loaded in an *r* register.

## 7.3.3 IU Logic Instructions

### 7.3.3.1 ADD, ANDN, OR, ORN, XOR, XNOR Instructions

These are standard logic operations that work on all 64 bits of the register. The instructions can optionally set the integer condition codes (*icc/xcc*).

## 7.3.4 IU Compare Instructions

A special comparison instruction for integer values is not needed since it is easily synthesized with the “subtract and set condition codes” (SUBcc) instruction.



## 7.3.5 IU Miscellaneous Instructions

### 7.3.5.1 Interval Arithmetic Mode Instruction (SIAM) (VIS II)

The Set Interval Arithmetic Mode (SIAM) instruction sets the interval arithmetic mode fields in the GSR.

### 7.3.5.2 Align Address Instruction

The ALIGNADDR instruction takes two  $r$  registers and adds them together. The three least significant bits are forced to zero.

The ALIGNADDRL instruction supports little-endian data structures by taking the two  $r$  registers, adding them together, and placing the two's-complement of the three least significant bits of the result and storing them in the 3-bit GSR.ALIGN field.

### 7.3.5.3 Population of Ones Count

A population opcode is defined but not implemented in hardware; instead, a trap is generated.

### 7.3.5.4 Privileged Register Access Instructions

The privileged register access instructions read and write another group of state and status registers called privileged registers. These registers are visible only to privileged software. The read privileged register instruction moves the privileged register contents into an  $r$  register. The write privileged register instruction moves the contents of an  $r$  register into the selected privileged register.

### 7.3.5.5 State Register Access Instructions

The state register instructions access program-visible state and status registers. The read state register instruction moves the state register contents into an  $r$  register. The write state register instruction moves the contents of an  $r$  register into the selected state register.

Some state registers can only be accessed in privileged mode, others in either privileged or non-privileged mode. Some registers have access bits to restrict their availability as desired by the privileged software.

---

## 7.4 Floating-Point Execution Environment

The floating-point and VIS execution unit includes the floating-point register file for floating-point and fixed-point data formats and the execution pipelines for floating-point and VIS instructions.

This execution unit is a single unit that may be referred to any one of the following, depending on the textual context:

- Floating-point Unit (FPU)
- Floating-point and Graphics Unit (FGU)
- VIS Execution Unit (VIS)
- FPU/VIS

---

**Note** – The instructions associated with the FPU/VIS execution unit are divided between floating-point and VIS execution environments, but otherwise uses the same hardware pipelines.

---

### 7.4.1 Floating-Point Operate Instructions

Floating-point operate (FPop) instructions perform all floating-point calculations; they are register-to-register instructions that operate on the floating-point registers. Like arithmetic, logical, and shift instructions, FPops compute a result that is a function of one or two source operands. Specific floating-point operations are selected by a subfield of the `FPOP1/FPOP2` instruction formats.

FPops are generally triadic register address instructions. They compute a result that is a function of one or two source operands and place the result in one or more destination `f` registers, with two exceptions:

- Floating-point convert operations, which use one source and one destination operand.
- Floating-point compare operations, which do not write to an `f` register but update one of the `fccn` fields of the FSR instead.

The term “FPop” refers to those instructions encoded by the `FPOP1` and `FPOP2` opcodes and does *not* include branches based on the floating-point condition codes (`FBfccD` and `FBPfcc`) or the load/store floating-point instructions.

If `PSTATE.PEF = 0` or `FPRS.FEF = 0`, then any instruction, including an FPop instruction, that attempts to access an FPU register generates a *fp\_disabled* exception.

All FPop instructions clear the `ftt` field and set the `cexc` field unless they generate an exception. Floating-point compare instructions also write one of the `fccn` fields. All FPop instructions that can generate IEEE exceptions set the `cexc` and `aexc` fields unless they generate an exception. `FABS(s,d,q)`, `FMOV(s,d,q)`, `FMOVcc(s,d,q)`, `FMOVr(s,d,q)`, and `FNEG(s,d,q)` cannot generate IEEE exceptions; therefore, they clear `cexc` and leave `aexc` unchanged.

---

**Note** – The processor may indicate that a floating-point instruction did not produce a correct IEEE Std 754-1985 result by generating a `fp_exception_other` exception with `FSR.ftt = unfinished_FPop` or unimplemented FPop. In this case, privileged software must emulate any functionality not present in the hardware.

---

The processor does not implement quad-precision floating-point operations in hardware. Instead, these operations cause a `fp_exception_other` trap with `FSR.ftt = unimplemented_FPop`, and the system software emulates quad operations.

## 7.4.2 FPU/VIS Data Access Instructions

Floating-point load and store instructions support word, doubleword, and quadword memory accesses.

There are no move instructions to move data directly between the integer and floating-point register files.

### 7.4.2.1 Load Instructions

Byte, halfword, word, and double/extended word data widths are supported with access to alternate address spaces. Data loaded into a register that is not 64 bits is filled with zeroes in the high-order bits.

### 7.4.2.2 Store Instructions

Byte, halfword, word, and double/extended word data widths are supported with access to alternate address spaces.

### 7.4.2.3 Block Load and Store Instructions

Block load and store access eight consecutive doublewords. The LDDFA instruction is used with the various ASIs to specify a type of block transaction. The LDDFA instruction is specified with ASIs 70, 71, 78, 79, F0, F1, F8, F9, E0, and E1 to select between primary and secondary D-MMU contexts, little-endian and big-endian, privileged and non-privileged, and a set of block commit store ASIs.

### 7.4.2.4 Conditional Move Instructions

The FP/VIS conditional move instructions are described with the IU conditional move instructions, Section 7.3.1.3.

## 7.4.3 FP Arithmetic Instructions

Single-precision and double-precision FP is executed in hardware. Quad precision (128-bit) instructions are recognized by the CPU and trapped so they can be emulated in software.

### 7.4.3.1 Absolute Value and Negate Instructions

These instructions modify the sign of the floating-point operand.

### 7.4.3.2 Add and Subtract Instructions

These instruction use standard IEEE operation.

### 7.4.3.3 Multiply Instructions

These instructions use standard IEEE operation with some exceptions.

### 7.4.3.4 Square Root and Divide Instructions

The square root and divide instructions begin their execution in the FGM pipeline and block new instructions from entering until the result is nearly ready to leave the pipeline and be written to the register file.

## 7.4.4 FP Conversion Instructions

The following FP conversions are supported. Conversions do not generate `fcc` condition codes.

### 7.4.4.1 Floating-Point to Integer

All floating-point precision to word and double/extended word integer conversions are supported.

### 7.4.4.2 Integer to Floating-Point

Word and double/extended word integer to all floating-point precision number conversions are supported.

### 7.4.4.3 Floating-Point to Floating-Point

All floating-point precision to all floating-point precision number conversions are supported.

## 7.4.5 FP Compare Instructions

The same precision operands are compared and the `fcc` condition codes are set.

## 7.4.6 FP Miscellaneous Instructions

### 7.4.6.1 Load and Store FSR Register

The FSR register is accessed by load and store instructions into and out of the floating-point register file.

### 7.4.6.2 Data Alignment Instruction

The data alignment instruction `FALIGNDATA` concatenates two registers (16 bytes) and stores a contiguous block of eight of these bytes starting at the offset stored in the `GSR.ALIGN` field.

---

## 7.5 VIS Execution Environment

The floating-point and VIS execution unit includes the floating-point register file for floating-point and fixed-point data formats and the execution pipelines for floating-point and VIS instructions.

This execution unit is a single unit that may be referred to any one of the following, depending on the textual context:

- Floating-point Unit (FPU)
- Floating-point and Graphics Unit (FGU)
- VIS Execution Unit (VIS)
- FPU/VIS

---

**Note** – The instructions associated with the FPU/VIS execution unit are divided between floating-point and VIS execution environments, but otherwise uses the same hardware pipelines.

---

### 7.5.1 VIS Pixel Data Instructions

#### 7.5.1.1 Array Instruction

These instructions convert three-dimensional (3D) fixed-point addresses to a blocked byte address.

#### 7.5.1.2 Byte Mask and Shuffle Instructions

Byte Mask instruction adds two integer registers and stores the result in the integer register. The least significant 32 bits of the result are stored in a special field.

Byte Shuffle concatenates the two 64-bit floating-point registers to form a 16-byte value. Bytes in the concatenated value are numbered from most significant to least significant, with the most significant byte being byte 0.

### 7.5.1.3 Edge Handling Instructions

These instructions handle the boundary conditions for parallel pixel scan line loops, where the address of the next pixel to render and the address of the last pixel in the scan line is provided.

### 7.5.1.4 Pixel Packing Instructions

These instructions convert multiple values in a source register to a lower precision fixed or pixel format and store the resulting values in the destination register. Input values are clipped to the dynamic range of the output format. Packing applies a scale factor to allow flexible positioning of the binary point.

### 7.5.1.5 Expand and Merge Instructions

Expand takes four 8-bit unsigned integers, converts each integer to a 16-bit fixed-point value, and stores the four resulting 16-bit values in a 64-bit floating-point register.

Merge interleaves four corresponding 8-bit unsigned values to produce a 64-bit value in the 64-bit floating-point destination register. This instruction converts from packed to planar representation when it is applied twice in succession.

### 7.5.1.6 Pixel Distance Instruction

Eight unsigned 8-bit values are contained in the 64-bit floating-point source registers. The corresponding 8-bit values in the source registers are subtracted. The sum of the absolute value of each difference is added to the integer in the 64-bit floating-point destination register. The result is stored in the destination register. Typically, this instruction is used for motion estimation in video compression algorithms.

## 7.5.2 VIS Fixed-Point 16-bit and 32-bit Data Instructions

### 7.5.2.1 Partitioned Add and Subtract Instructions

The standard versions of these instructions perform four 16-bit or two 32-bit partitioned adds or subtracts between the corresponding fixed-point values contained in the source operands.

The single-precision versions of these instructions perform two 16-bit or one 32-bit partitioned add(s) or subtract(s); only the low 32 bits of the destination register are affected.

### 7.5.2.2 Partitioned Multiply Instructions

These instructions multiply signed and unsigned registers of different sizes and place the results in different types of destination registers.

### 7.5.2.3 Pixel Compare Instruction

Either four 16-bit or two 32-bit fixed-point values in the 64-bit floating-point source registers are compared. The 4-bit or 2-bit results are stored in the least significant bits in the integer destination register. Signed comparisons are used.

## 7.5.3 VIS Logic Instructions

### 7.5.3.1 Fill with Ones and Zeroes Instruction

These instructions perform a zero fill or a one fill.

### 7.5.3.2 Source Copy

These instructions perform a source copy.

### 7.5.3.3 AND, OR, NAND, NOR, and XNOR Instructions

These instructions perform the logical operations.

---

## 7.6 Data Coherency Instructions

The processor implements a Total Store Ordering (TSO) that provides the majority of data coherency support in hardware. Two instructions are used with this model to synchronize the data for memory operations to insure the latest data is accessed for load instructions and DMA activity.

Chapter 9, “Memory Models” discusses TSO in detail.



## 7.6.1 FLUSH Instruction Cache Instruction

The `FLUSH` instruction is used to flush the caches out to main memory. The `MEMBAR` instruction is used to flush the various data buffers in the CPU out to data coherent domain.

Self-modifying code (storable in the unified L2-cache) requires the use of the `FLUSH` instruction.

---

**Note** – The `FLUSHW` instruction flushes the window registers and is not related to the `FLUSH` command for the instruction cache.

---

## 7.6.2 MEMBAR (Memory Synchronization) Instruction

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

## 7.6.3 Store Barrier Instruction

---

**Note** – `STBARP` is also supported, but this instruction is deprecated and should not be used in newly developed software.

---

---

## 7.7 Register Window Management Instructions

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

The instructions that manage register windows include:

### *SAVE Instruction*

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

### *RESTORE Instruction*

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

### *SAVED<sup>P</sup> Instruction*

The SAVED instruction is used by a spill trap handler to indicate that a window spill has completed successfully. It increments CANSAVE.

### *RESTORED<sup>P</sup> Instruction*

The RESTORED instruction is used by a fill trap handler to indicate that a window has been filled successfully. It increments CANRESTORE.

### *Flush Register Windows Instruction*

The FLUSHW instruction cleans register windows of the data from other processes to insure a secure execution environment.

---

## 7.8 Program Control Transfer Instructions

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

The instruction following a delayed CTI is called a *delay* instruction. A bit in a delayed CTI (the *annul bit*) can cause the delay instruction to be annulled (that is, to have no effect) if the branch is not taken (or in the “branch always” case if the branch is taken).

---

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

---

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

---

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

---

## 7.8.1 Control Transfer Instructions (CTIs)

The following are the basic CTI types:

- Conditional branch (Bicc<sup>D</sup>, BPcc, BPr, FBfcc<sup>D</sup>, FBPfcc)
- Unconditional branch
- Call and link (CALL)
- Jump and link (JMPL, RETURN)
- Return from trap (DONE<sup>P</sup>, RETRY<sup>P</sup>)
- Trap (Tcc, ILLTRAP)
- No Operation (NOP, SIR when in non-privileged mode)

A CTI functions by changing the value of the next program counter (nPC) or by changing the value of both the program counter (PC) and the nPC. When only the nPC is changed, the effect of the transfer of control is delayed by one instruction. Most control transfers are of delayed variety. The instruction following a delayed CTI is said to be in the *delay slot* of the CTI. Some CTI (branches) can optionally annul, that is, not execute, the instruction in the delay slot, depending upon whether the transfer is taken or not taken. Annulled instructions have no effect upon the program-visible state, nor can they cause a trap.

---

**Programming Note** – The annul bit increases the likelihood that a compiler can find a useful instruction to fill the delay slot after a branch, thereby reducing the number of instructions executed by a program. For example, the annul bit can be used to move an instruction from within a loop to fill the delay slot of the branch that closes the loop.

Likewise, the annul bit can be used to move an instruction from either the “else” or “then” branch of an “if-then-else” program block to the delay slot of the branch that selects between them. Since a full set of conditions is provided, a compiler can arrange the code (possibly reversing the sense of the condition) so that an instruction from either the “else” branch or the “then” branch can be moved to the delay slot.

Use of annulled branches provided some benefit in older, single-issue SPARC implementations. The UltraSPARC III Cu processor is a superscalar SPARC implementation in which the only benefit of annulled branches might be a slight reduction in code size. Therefore, the use of annulled branch instructions is no longer encouraged.

---

TABLE 7-2 defines the value of the PC and the value of the nPC after execution of each instruction. Conditional branches have two forms: branches that test a condition (including branch-on-register), represented in the table by `BCC` (same as `BiCC`), and branches that are unconditional, that is, always or never taken, represented in the table by `B`. The effect of an annulled branch is shown in the table through explicit transfers of control, rather than fetching and annulling the instruction.

**TABLE 7-2** Control Transfer Characteristics

| Instruction Group | Address Form      | Delayed | Taken | Annul Bit | New PC   | New nPC      |
|-------------------|-------------------|---------|-------|-----------|----------|--------------|
| Non-CTIs          | —                 | —       | —     | —         | nPC      | nPC + 4      |
| Bcc               | PC-relative       | Yes     | Yes   | 0         | nPC      | EA           |
| Bcc               | PC-relative       | Yes     | No    | 0         | nPC      | nPC + 4      |
| Bcc               | PC-relative       | Yes     | Yes   | 1         | nPC      | EA           |
| Bcc               | PC-relative       | Yes     | No    | 1         | nPC + 4  | nPC + 8      |
| B                 | PC-relative       | Yes     | Yes   | 0         | nPC      | EA           |
| B                 | PC-relative       | Yes     | No    | 0         | nPC      | nPC + 4      |
| B                 | PC-relative       | Yes     | Yes   | 1         | EA       | EA + 4       |
| B                 | PC-relative       | Yes     | No    | 1         | nPC + 4  | nPC + 8      |
| CALL              | PC-relative       | Yes     | —     | —         | nPC      | EA           |
| JMPL, RETURN      | Register-indirect | Yes     | —     | —         | nPC      | EA           |
| DONE              | Trap state        | No      | —     | —         | TNPC[TL] | TNPC[TL] + 4 |
| RETRY             | Trap state        | No      | —     | —         | TPC[TL]  | TNPC[TL]     |
| Tcc               | Trap vector       | No      | Yes   | —         | EA       | EA + 4       |
| Tcc               | Trap vector       | No      | No    | —         | nPC      | nPC + 4      |

The effective address (EA) in TABLE 7-2, specifies the target of the control transfer instruction. The EA is computed in different ways, depending on the particular instruction.

- **PC-relative effective address** — A PC-relative EA is computed by sign extending the instruction’s immediate field to 64 bits, left-shifting the word displacement by two bits to create a byte displacement, and adding the result to the contents of the PC.
- **Register-indirect effective address** — A register-indirect EA computes its target address as either  $r[rs1] + r[rs2]$  if  $i = 0$ , or  $r[rs1] + \text{sign\_ext}(\text{simml3})$  if  $i = 1$ .
- **Trap vector effective address** — A trap vector EA first computes the software trap number as the least significant 7 bits of  $r[rs1] + r[rs2]$  if  $i = 0$ , or as the least significant 7 bits of  $r[rs1] + \text{sw\_trap\#}$  if  $i = 1$ . The trap level, TL, is incremented. The hardware trap type is computed as  $256 + \text{sw\_trap\#}$  and stored in TT[TL]. The EA is generated by concatenation of the contents of the TBA register, the “TL > 0” bit, and the contents of TT[TL].
- **Trap state effective address** — A trap state EA is not computed but is taken directly from either TPC[TL] or TNPC[TL].

---

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

The SPARC V8 architecture left undefined the result of executing a delayed conditional branch that had a delayed control transfer in its delay slot. For this reason, programmers should avoid such constructs when backward compatibility is an issue.

---

### 7.8.1.1 Conditional Branches

A conditional branch transfers control if the specified condition is true. If the annul bit is zero, the instruction in the delay slot is always executed. If the annul bit is one, the instruction in the delay slot is *not* executed *unless* the conditional branch is taken.

---

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

---

### 7.8.1.2 Unconditional Branches

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

---

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

---

### 7.8.1.3 CALL/JMPL and RETURN Instructions

#### *CALL*

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

When `PSTATE.AM = 1`, the value of the high-order 32 bits is transmitted to `r[15]` by the `CALL` instruction.

### *Jump and Link*

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

When `PSTATE.AM = 1`, the value of the high-order 32 bits transmitted to `r[rd]` by the `JMPL` instruction is zero.

### *RETURN*

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

## 7.8.1.4 DONE and RETRY Instructions

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

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

## 7.8.1.5 Trap Instruction (Tcc)

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

A `TCC` instruction can specify 1 of 128 software trap types. When a `TCC` is taken, 256 plus the seven least significant bits of the sum of the `TCC`'s source operands is written to `TT[TL]`. The only visible difference between a software trap generated by a `TCC` instruction and a hardware trap is the trap number in the `TT` register. See Chapter 12, “Traps and Trap Handling” for more information.

---

**Programming Note** – `TCC` can be used to implement breakpointing, tracing, and calls to supervisor software. `TCC` can also be used for runtime checks, such as out-of-range array index checks or integer overflow checks.

---

### 7.8.1.6 ILLTRAP

The ILLTRAP instruction causes an *illegal\_instruction* exception.

### 7.8.1.7 NOP

A NOP instruction occupies the entire (single) instruction group and performs no visible work.

- NOP Instruction

There are other instructions that also result in an operation that has no visible effect:

- SIR instruction executed in non-privileged mode
- SHUTDOWN instruction executed in privileged mode

There are other instructions that appear to be a NOP as long as they do not affect the condition codes.

---

## 7.9 Prefetch Instructions

The prefetch instruction is used to request that data be fetched from memory and put into the cache(s) if not already there for use in the floating-point and VIS execution environment. A subsequent load, if properly scheduled, can expect the data to more likely be in the cache, reducing the number of times the pipeline must recycle and thus improving performance.

The destination field of a `PREFETCH` instruction (`fcn`) is used to encode the prefetch type. The `PREFETCHA` instruction supports accesses to alternate space.

`PREFETCH` accesses at least 64 bytes. Refer to Appendix A, “Prefetch Data” on page 560 for further details.



---

## 7.10 Instruction Summary Table by Category

A summary of instructions are categorized in TABLE 7-3.

### 7.10.1 Instruction Superscripts

|                          |  |
|--------------------------|--|
| INSTRUCTION <sup>P</sup> | Instruction must execute in privileged mode.                 |
| INSTRUCTION              | Instruction can execute in privileged or non-privileged mode |

### 7.10.2 Instruction Mnemonics Expansion

|                                 |  |
|---------------------------------|--|
| INSTRUCTION{ <sub>A</sub> }     | means INSTRUCTION, INSTRUCTION <sub>A</sub>  |
| INSTRUCTION( <sub>A,B,C</sub> ) | means INSTRUCTION <sub>A</sub> , INSTRUCTION <sub>B</sub> , and INSTRUCTION <sub>C</sub> |

### 7.10.3 Instruction Grouping Rules

Chapter 4, “Instruction Execution” explains instruction grouping rules in detail.

#### Execution Latency

All instructions execute within the pipeline except the following:

- FSQRT (floating-point square root)
- FPDIVx (floating-point divide)

The latency of these instructions depend on the precision of the floating-point values. Some instructions execute early in the pipeline and have special bypass abilities. Chapter 4, “Instruction Execution” explains execution latencies in detail.

### 7.10.4 Table Organization

The Instruction Summary Table has the following main sections:

- Integer Execution Environment (TABLE 7-3)
  - Data access, Arithmetic, Logic, Compare, Miscellaneous instructions

- Floating-point Execution Environment (TABLE 7-4)
  - FP/VIS data access, FP arithmetic/logic/compare/miscellaneous
- VIS Execution Environment (TABLE 7-5)
  - VIS pixel and fixed-point arithmetic/logic
- Data Coherency Instructions (TABLE 7-6)
- Register-window Management Instructions (TABLE 7-7)
- Program Control Transfer Instructions (TABLE 7-8)
- Prefetch Instructions (TABLE 7-9)

Shaded areas indicate instructions that are completely deprecated (entire row) or always privileged (cell holding instruction name). Deprecated and privilege status is identified with a <sup>D</sup> or <sup>P</sup> superscript, respectively.

## 7.10.5 Integer Execution Environment Instructions

**TABLE 7-3** Instruction Summary for the Integer Execution Environment (*1 of 3*)

| Instruction   | Description   | ASI Load (hex)        | Notes |
|---|---|-----------------------|-------|
| <b>Integer Execution Environment</b>                                |   |                       |       |
| <b>IU Data Access Instructions</b><br>B= byte; H= halfword; W=word; |   | <b>ASI Load (hex)</b> |       |
| LDD <sup>D</sup>  | Load integer double word  | No                    |       |
| LDDA <sup>D, P<sup>ASI</sup></sup>                                  | Load integer double word from alternate space   |                       |       |
| LDDA <sup>P<sup>ASI</sup></sup>                                     | Atomic quad load  | 24, 2C                |       |
| LDS(B,H,W)  | Load signed extended byte, halfword, or word:<br>Memory → IU register   | No                    |       |
| LDX   | Load extended (double) word   | No                    |       |
| LDXA <sup>P<sup>ASI</sup></sup>                                     | Load extended (double) word from alternate space  |                       |       |
| LDS(B,H,W)A <sup>P<sup>ASI</sup></sup>                              | Load signed extended byte, halfword, or word from alternate space   |                       |       |
| LDSTUB  | Load-store (atomic) unsigned byte:<br>Memory → IU register & Compare logic;<br>IU register → Memory (conditional) | No                    |       |
| LDSTUBA <sup>P<sup>ASI</sup></sup>                                  | Load-store (atomic) unsigned byte (see LDSTUB) in alternate space   |                       |       |
| LDU(B,H,W)  | Load unsigned byte, halfword, word:<br>Memory → IU register   |                       |       |

**TABLE 7-3** Instruction Summary for the Integer Execution Environment (2 of 3)

| Instruction                                      | Description   | Notes   |
|--|---|---------|
| LDU(B,H,W)A <sup>PASI</sup>                      | Load unsigned byte, halfword, word from alternate space   |         |
| ST(B,H,W,D <sup>D</sup> ,X)                      | Store byte, halfword, word, double, or extended word:<br>IU register → Memory   |         |
| ST(B,H,W,D <sup>D</sup> ,X)A <sup>PASI</sup>     | Store byte, halfword, word, double, or extended word in alternate space   |         |
| MOV <sub>cc</sub>                                | Conditional move based on icc/fcc:<br>IU register → IU register   | 1       |
| MOV <sub>r</sub>                                 | Conditional move based on IU register value:<br>IU register → IU register   | 2       |
| CASA <sup>PASI</sup> , CASXA <sup>PASI</sup>     | Atomic Compare and Swap word/double word in alternate space:<br>Memory → Compare logic<br>Memory ↔ (conditional) Working register | 3, 4, 5 |
| SWAP <sup>D</sup> {A <sup>D</sup> , PASI}        | Atomically swap optionally with alternate space:<br>IU register ↔ Memory  |         |
| <b>IU Arithmetic Instructions</b>                |   |         |
| S= signed; U= unsigned; X= 64-bit (otherwise 32) |   |         |
| ADD{cc}  | Integer add   |         |
| ADD <sub>C</sub> {cc}                            | Integer add with carry  |         |
| SUB{cc}  | Integer subtract, optionally setting icc/xcc  |         |
| SUB <sub>C</sub> {cc}                            | Integer subtract with carry optionally setting icc/xcc  |         |
| MULX   | Signed or unsigned 64-bit multiply  |         |
| (S,U)MUL{cc} <sup>D</sup>                        | Signed/unsigned integer multiply optionally setting icc/xcc   |         |
| UDIVX  | Unsigned 64-bit integer divide  |         |
| SDIVX  | Signed 64-bit integer divide  |         |
| (S,U)DIV{cc} <sup>D</sup>                        | Signed/unsigned 32-bit integer divide optionally setting icc/xcc  |         |
| SETHI  | Modify highest 22 bits of low word in IU register:<br>Immediate → IU register (partial)   |         |
| SLL{X}   | Shift left logical (32/64-bit)  |         |
| SRL{X}   | Shift right logical (32/64-bit)   |         |
| SRA{X}   | Shift right arithmetic (32/64-bit)  |         |
| TADD <sub>cc</sub> {TV <sup>D</sup> }            | Tagged add and modify icc optionally trap on overflow   |         |

**TABLE 7-3** Instruction Summary for the Integer Execution Environment (3 of 3)

| Instruction  | Description   | Notes |
|--|---|-------|
| TSUBcc{TV <sup>D</sup> }   | Tagged subtract and modify icc optionally trap on overflow  |       |
| <b>IU Logic Instructions</b>   |   |       |
| AND{cc}  | Logical AND, optionally setting icc/xcc   |       |
| ANDN{cc}   | Logical AND-not, optionally setting icc/xcc   |       |
| OR{cc}   | Logical OR, optionally setting icc/xcc  |       |
| ORN{cc}  | Logical OR-not, optionally setting icc/xcc  |       |
| XOR{cc}  | Logical XOR, optionally setting icc/xcc   |       |
| XNOR{cc}   | Logical XNOR, optionally setting icc/xcc  |       |
| <b>IU Miscellaneous Instructions</b>   |   |       |
| SIAM   |   |       |
| ALIGNADDRESS{_LITTLE}  | Calculates aligned address  |       |
| POPC   | Defined to count the number of ones in register, unimplemented (causes an illegal instruction execution which traps to software for emulation)  |       |
| RDPR <sup>P</sup>  | Read privileged register  |       |
| WRPR <sup>P</sup>  | Write privileged register   |       |
| RDASR <sup>PASR</sup>  | Read ancillary state register (ASR) - see below. Privileged mode required for privileged ASRs.  |       |
| RDY <sup>D</sup> , RDCCR, RDASI, RDPC, RDFPRS, RDPCR <sup>P</sup> , RDPIC <sup>PPCR.PRIV</sup> , RDDCR <sup>P</sup> , RDGSR, RDSOFTINT <sup>P</sup> , RDTICK <sup>PNPT</sup> , RDSTICK <sup>PNPT</sup> , RDTICK_CMPR <sup>P</sup> , RDSTICK_CMPR <sup>P</sup>                            | Read state and ancillary state registers:<br><br>- If PCR.PRIV field is one, then PIC register access requires privileged mode.<br><br>- If {TICK STICK}.NPT field is zero, then TICK/STICK register reads require privileged mode. |       |
| WRASR <sup>PASR</sup>  | Write ancillary state register (ASR); Privileged mode required for privileged ASRs.   |       |
| WRY <sup>D</sup> , WRCCR, WRASI, WRFPRS, WRPCR <sup>P</sup> , WRPIC <sup>PPCR.PRIV</sup> , WRDCR <sup>P</sup> , WRGSR, WRSOFTINT <sup>P</sup> , WRSOFTINT_CLR <sup>P</sup> , WRSOFTINT_SET <sup>P</sup> , WRSTICK <sup>PNPT</sup> , WRTICK_CMPR <sup>P</sup> , WRSTICK_CMPR <sup>P</sup> | Read state and ancillary state registers:<br><br>- If PCR.PRIV field is one, then PIC register access requires privileged mode.<br><br>- If STICK.NPT field is zero, then STICK register writes require privileged mode.            |       |

1. A simple register-to-register move is accomplished by using the OR instruction with  $r[0]$ .

2. Load (LD) and store (ST) instructions are provided with many size formats (byte, word, double word, etc.) and most can be specified with an alternate space identifier (ASI).
3. The “r” refers to value in  $x$  registers.
4. The cc refers to settings of the integer condition codes.
5. The conditional move instructions (integer and floating-point) are influenced by the condition codes of either execution unit to facilitate moves in one type of execution unit based on the condition codes of the other or of those within the execution unit.

## 7.10.6 Floating-Point Execution Environment Instructions

**TABLE 7-4** Instruction Summary for the Floating-Point Execution Environment

| Instruction  | Description   | Reference Pages                              | Notes |
|--|---|--|-------|
| <b>FP/VIS Data Access Instruction</b><br>s= 32-bit; d= 64-bit; q= 128-bit (q is trapped) |   | <b>ASI Load (hex)</b>                        |       |
| LD <sub>{D}</sub> F  | Load word (or double word):<br>Memory → FPU register                        | No   |       |
| LD <sub>{D}</sub> FA <sup>PASI</sup>   | Load word (or double word) from alternate space:<br>Memory → FPU register   |  |       |
| LDDFA  | Block load 64 bytes:<br>Memory → FPU registers                              |  |       |
| LDDFA  | Load short:<br>Memory → FPU register  |  |       |
| LDQF   | Load quadword:<br>Memory → FPU register                                     | No   |       |
| LDQFA <sup>PASI</sup>  | Load quadword from alternate space:<br>Memory → FPU register                | No   |       |
| ST <sub>(F,DF,QF)</sub>  | Store word, double, or quad word to memory:<br>FPU register → Memory        | No   |       |
| ST <sub>(F,DF,QF)A</sub> <sup>PASI</sup>   | Store word, double, or quad word to memory using alternate memory space.    |  |       |
| STDFA  | Block store 64 bytes: uses ASIs   | 70, 71, 78, 79,<br>F0, F1, F8, F9,<br>E0, E1 |       |
| STDFA  | Short FP store: uses ASIs D(0:3) <sub>16</sub> ,<br>D(8:B) <sub>16</sub>    |  |       |
| STDFA  | Partial store FPU: uses ASIs C(0:5) <sub>16</sub> ,<br>C(8:D) <sub>16</sub> |  |       |
| FMOV <sub>(s,d,q)</sub>  | FPU → FPU register  | No   |       |
| FMOV <sub>(s,d,q)cc</sub>  | Conditional move, IU or FPU condition codes:<br>FPU → FPU register          | No   |       |
| FMOV <sub>(s,d,q)r</sub>   | Conditional move, IU or FPU register value: FPU → FPU register              | No   |       |
| <b>FP Arithmetic Instructions</b><br>s= 32-bit; d= 64-bit; q= 128-bit (q is trapped)     |   |  |       |
| FABS <sub>(s,d,q)</sub>  | FP absolute value   |  |       |
| FNEG <sub>(s,d,q)</sub>  | Change FP sign  |  |       |
| FADD <sub>(s,d,q)</sub>  | FP add  |  |       |
| FSUB <sub>(s,d,q)</sub>  | FP subtract   |  |       |

**TABLE 7-4** Instruction Summary for the Floating-Point Execution Environment (Continued)

| Instruction  | Description   | Reference Pages | Notes |
|--|---|-----------------|-------|
| FMUL <sub>(s,d,q)</sub>  | FP multiply   |                 |       |
| FdMULq   | FP multiple doubles to quadword                                 |                 |       |
| FsMULd   | FP multiple singles to doubleword                               |                 |       |
| FDIV <sub>(s,d,q)</sub>  | FP division   |                 |       |
| FSQRT <sub>(s,d,q)</sub>   | FP square root  |                 |       |
| <b>FP Conversion Instructions</b>  |   |                 |       |
| s= 32-bit; d= 64-bit; q= 128-bit (q is trapped); i= integer word; x= double (or extended) word |   |                 |       |
| F <sub>(s,d,q)</sub> TOi   | Floating-point to integer word                                  |                 |       |
| F <sub>(s,d,q)</sub> TOx   | Floating-point to integer double word                           |                 |       |
| F <sub>(s,d,q)</sub> TO <sub>(s,d,q)</sub>   | Floating-point to floating-point                                |                 |       |
| FiTO <sub>(s,d,q)</sub>  | Integer word to floating-point                                  |                 |       |
| FxTO <sub>(s,d,q)</sub>  | Integer double (or extended) word to floating-point             |                 |       |
| <b>FP Compare Instructions</b>   |   |                 |       |
| FCMP <sub>(s,d,q)</sub>  | FP compare of like precision, sets fcc condition codes          |                 |       |
| FCMPE <sub>(s,d,q)</sub>   | Same as FCMP, but an exception is generated if unordered        |                 |       |
| <b>FP Miscellaneous Instructions</b>   |   |                 |       |
| LDFSR <sup>D</sup>   | Load FSR into FP reg file:<br>FSR → FPU register (lower 32-bit) |                 |       |
| LDXFSR   | Load FSR into FP reg file:<br>FSR → FPU register (64-bit)       |                 |       |
| STFSR <sup>D</sup>   | Store FSR register:<br>FPU (lower 32-bit) → FSR register        |                 |       |
| STXFSR   | Store FSR register:<br>FPU → FSR register                       |                 |       |
| FALIGNDATA   | Concatenates two 64-bit registers into one based on GSR.ALIGN   |                 |       |

## 7.10.7 VIS Execution Environment Instructions

**TABLE 7-5** Instruction Summary for the VIS Execution Environment

| Instruction  | Description  | Reference Pages | Notes |
|--|--|-----------------|-------|
| <b>VIS Data Access Instructions</b>  |  |                 |       |
| Refer to Section 7.10.6, "Floating-Point Execution Environment Instructions" of the Instruction Summary Table. |  |                 |       |
| <b>VIS Pixel Data Instructions</b><br>L= little-endian; N= fcc not modified; S= 32-bit (otherwise 64-bit);     |  |                 |       |
| ARRAY(8,16,32)   | 3D-array addressing  |                 |       |
| BMASK  | Writes the GSR.MASK field  |                 |       |
| BSHUFFLE   | Permute bytes as specified by GSR.MASK field.                    |                 |       |
| EDGE(8,16,32)<br>( L,N,LN)   | Edge handling instructions                                       |                 |       |
| FEXPAND  | Pixel data expansion   |                 |       |
| FPMERGE  | Pixel merge  |                 |       |
| FPACK(16,32,FIX)   | Pixel packing  |                 |       |
| PDIST  | Pixel component distance   |                 |       |
| <b>VIS Fixed-point 16/32-bit Data Instructions</b>   |  |                 |       |
| FPADD(16,32){S}  | Fixed-point add, 16- or 32-bit operands, 32/64-bit register      |                 |       |
| FPSUB(16,32){S}  | Fixed-point subtract, 16- or 32-bit operands, 32/64-bit register |                 |       |
| FMUL8x16   | 8x16 partitioned multiply  |                 |       |
| FMUL8x16(AU,AL)  | 8x16 Upper/Lower $\alpha$ partitioned multiply                   |                 |       |
| FMUL8(SU,SL)x16  | 8x16 Upper/Lower partitioned multiply                            |                 |       |
| FMULD8(SU,SL)x16   | 8x16 Upper/Lower partitioned multiply                            |                 |       |
| FCMP(GT,LE,NE,EQ)(16,32)   | Fixed-point compare (also known as "pixel compare")              |                 |       |
| <b>VIS Logic Instructions</b><br>S= 32-bit (otherwise 64-bit)  |  |                 |       |
| FSRC(1,2){S}   | Copy source  |                 |       |
| FONE{S}  | Fill with ones (32/64-bit)                                       |                 |       |
| FZERO{S}   | Fill with zeroes (32/64-bit)                                     |                 |       |
| FAND{S}  | Logical AND (32/64-bit)  |                 |       |
| FANDNOT(1,2){S}  | Logical AND with a source inverted (32/64-bit)                   |                 |       |
| FOR{S}   | Logical OR (32/64-bit)   |                 |       |
| FNAND{S}   | Logical NAND (32/64-bit)   |                 |       |
| FNOR{S}  | Logical NOR (32/64-bit)  |                 |       |



**TABLE 7-5** Instruction Summary for the VIS Execution Environment (Continued)

| Instruction    | Description                                   | Reference Pages | Notes |
|----------------|---|-----------------|-------|
| FORNOT(1,2){S} | Logical OR with a source inverted (32/64-bit) |                 |       |
| FNOT(1,2){S}   | Logical inversion of source bits (32/64-bit)  |                 |       |
| FXNOR{S}       | Logical XNOR (32/64-bit)                      |                 |       |
| FXOR{S}        | Logical XOR (32/64-bit)                       |                 |       |

## 7.10.8 Data Coherency Instructions

**TABLE 7-6** Instruction Summary for Data Coherency

| Instruction                        | Description             | Reference Pages | Notes |
|------------------------------------|-------------------------|-----------------|-------|
| <b>Data Coherency Instructions</b> |                         |                 |       |
| FLUSH                              | Flush instruction cache |                 |       |
| MEMBAR                             | Memory barrier          |                 |       |
| STBAR <sup>D</sup>                 | Store barrier           |                 |       |

## 7.10.9 Register-Window Management Instructions

**TABLE 7-7** Instruction Summary for Register-Window Management

| Instruction                                    | Description              | Reference Pages | Notes |
|--|--------------------------|-----------------|-------|
| <b>Register-Window Management Instructions</b> |                          |                 |       |
| SAVE   | Save caller's window     |                 |       |
| SAVED <sup>P</sup>                             | Window has been saved    |                 |       |
| RESTORE  | Restore caller's window  |                 |       |
| RESTORED <sup>P</sup>                          | Window has been restored |                 |       |
| FLUSHW   | Flush register windows   |                 |       |

## 7.10.10 Program Control Transfer Instructions

**TABLE 7-8** Instruction Summary for Program Control Transfer

| Instruction   | Description  | Reference Pages | Notes |
|---|--|-----------------|-------|
| <b>Program Control Transfer Instructions</b>                          |  |                 |       |
| icc/xcc= integer condition codes (32/64-bit); fcc= FP condition codes |  |                 |       |
| Bicc <sup>D</sup>   | Conditional branch on icc/xcc  |                 |       |
| BPcc  | Conditional branch on icc/xcc with branch prediction                 |                 |       |
| BPr   | Conditional branch on IU reg value with branch prediction            |                 |       |
| CALL  | Call and link  |                 |       |
| DONE <sup>P</sup>   | Return from Trap   |                 |       |
| FBfcc <sup>D</sup>  | Conditional branch on fcc  |                 |       |
| FBPfcc  | Conditional branch on fcc with branch prediction                     |                 |       |
| ILLTRAP   | Causes <i>illegal_instruction</i> trap                               |                 |       |
| JMPL  | Jump and link  |                 |       |
| NOP   | No operation   |                 |       |
| RETRY <sup>P</sup>  | Return from trap entry   |                 |       |
| RETURN  | Return (jump and link)   |                 |       |
| SHUTDOWN <sup>P</sup>   | Intended for low power mode, but is a NOP in the processor           |                 |       |
| SIR <sup>P</sup> NOP  | Software initiated reset: a NOP when executed in non-privileged mode |                 |       |
| Tcc   | Trap on icc/xcc  |                 |       |

## 7.10.11 Data Prefetch Instructions


**TABLE 7-9** Instruction Summary Table

| Instruction                  | Description   | Reference Pages | Notes |
|------------------------------|---|-----------------|-------|
| <b>Prefetch Instructions</b> |   |                 |       |
| PREFETCH                     | Tells processor to fetch data                             |                 |       |
| PREFETCHA <sup>PASI</sup>    | Tells processor to fetch data from alternate memory space |                 |       |

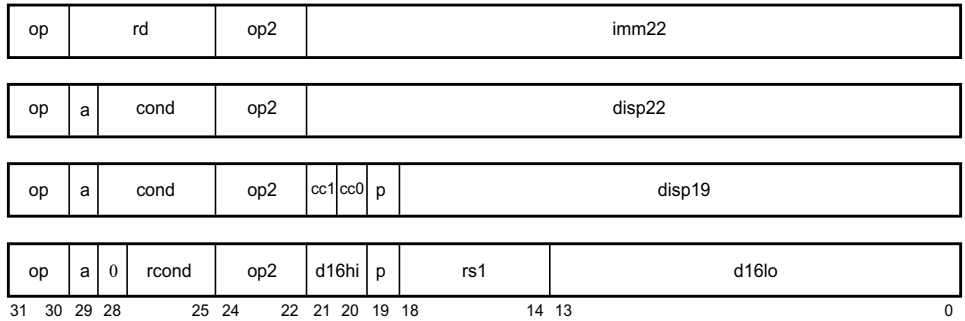
# 7.11 Instruction Formats and Fields

Instructions are encoded in four major 32-bit formats and several minor formats, as shown in FIGURE 7-1, FIGURE 7-2, and FIGURE 7-3.

*Format 1 (op = 1): CALL*



*Format 2 (op = 0): SETHI and Branches (Bicc, BPcc, BPr, FBfcc, FBPfcc)*



**FIGURE 7-1** Summary of Instruction Formats: Formats 1 and 2

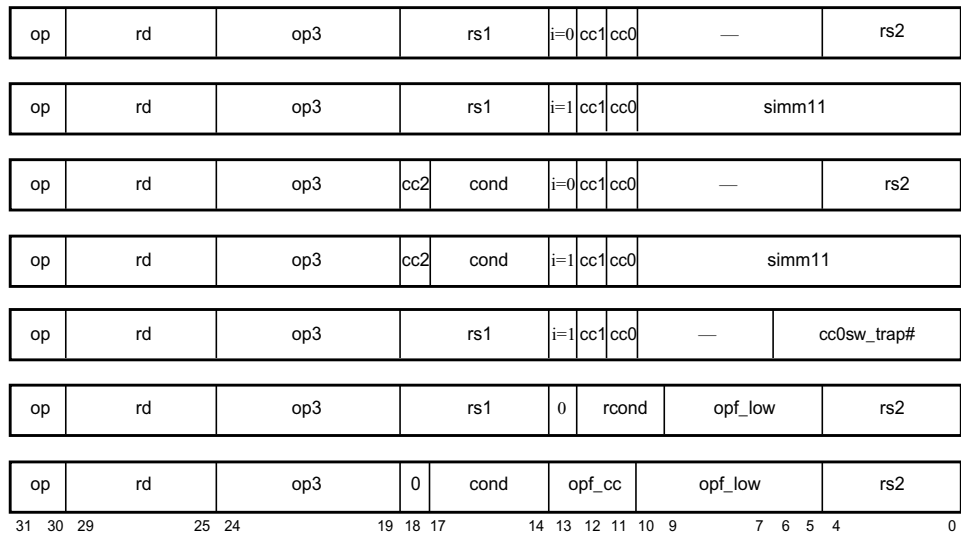
*Format 3 (op = 2 or 3): Arithmetic, Logical, MOV<sub>r</sub>, MEMBAR, Prefetch, Load, and Store*

|    |                 |         |                 |     |         |        |         |
|----|-----------------|---------|-----------------|-----|---------|--------|---------|
| op | rd              | op3     | rs1             | i=0 | —       | rs2    |         |
| op | rd              | op3     | rs1             | i=1 | simm13  |        |         |
| op | fcn             | op3     | rs1             | i=0 | —       | rs2    |         |
| op | fcn             | op3     | rs1             | i=1 | simm13  |        |         |
| op | —               | op3     | rs1             | i=0 | —       | rs2    |         |
| op | —               | op3     | rs1             | i=1 | simm13  |        |         |
| op | rd              | op3     | rs1             | i=0 | rcond   | —      | rs2     |
| op | rd              | op3     | rs1             | i=1 | rcond   | simm10 |         |
| op | rd              | op3     | rs1             | i=1 | —       | rs2    |         |
| op | rd              | op3     | rs1             | i=1 | —       | cmask  | mmask   |
| op | rd              | op3     | rs1             | i=0 | imm_asi |        | rs2     |
| op | <i>impl-dep</i> | op3     | <i>impl-dep</i> |     |         |        |         |
| op | rd              | op3     | rs1             | i=0 | x       | —      | rs2     |
| op | rd              | op3     | rs1             | i=1 | x=0     | —      | shcnt32 |
| op | rd              | op3     | rs1             | i=1 | x=1     | —      | shcnt64 |
| op | rd              | op3     | —               | opf |         | rs2    |         |
| op | 0 0 0           | cc1 cc0 | op3             | rs1 | opf     |        | rs2     |
| op | rd              | op3     | rs1             | opf |         | rs2    |         |
| op | rd              | op3     | rs1             | —   |         |        |         |
| op | fcn             | op3     | —               |     |         |        |         |
| op | fcn             | op3     | —               |     |         |        |         |

31 30 29 25 24 19 18 14 13 12 11 10 9 7 6 5 4 3 0

**FIGURE 7-2** Summary of Instruction Formats: Format 3

Format 4 (*op* = 2): MOVcc, FMOVr, FMOVcc, **and** Tcc



**FIGURE 7-3** Summary of Instruction Formats: Format 4

The instruction fields are interpreted as described in TABLE 7-10.

**TABLE 7-10** Instruction Field Interpretation (1 of 3)

| Field         | Description   |
|---------------|---|
| a             | The a bit annuls the execution of the following instruction if the branch is conditional and not taken, or if it is unconditional and taken.  |
| cc2, cc1, cc0 | cc2, cc1, and cc0 specify the condition codes (icc, xcc, fcc0, fcc1, fcc2, fcc3) to be used in the following instructions: <ul style="list-style-type: none"> <li>• Branch on Floating-point Condition Codes with Prediction Instructions (FBPfcc)</li> <li>• Branch on Integer Condition Codes with Prediction (BPcc)</li> <li>• Floating-point Compare Instructions (FCMP and FCMPE)</li> <li>• Move Integer Register If Condition Is Satisfied (MOVcc)</li> <li>• Move Floating-point Register If Condition Is Satisfied (FMOVcc)</li> <li>• Trap on Integer Condition Codes (Tcc)</li> </ul> In instructions such as Tcc that do not contain the cc2 bit, the missing cc2 bit takes on a default value. |
| cmask         | This 3-bit field specifies sequencing constraints on the order of memory references and the processing of instructions before and after a MEMBAR instruction.   |
| cond          | This 4-bit field selects the condition tested by a branch instruction.  |
| d16hi, d16lo  | These 2-bit and 14-bit fields together comprise a word-aligned, sign-extended, PC-relative displacement for a branch-on-register-contents with prediction (BPr) instruction.  |

**TABLE 7-10** Instruction Field Interpretation (2 of 3)

| Field          | Description  |   |                   |   |                                       |   |                                   |
|----------------|--|---|-------------------|---|---------------------------------------|---|-----------------------------------|
| disp19         | This 19-bit field is a word-aligned, sign-extended, PC-relative displacement for an integer branch-with-prediction (BPCC) instruction or a floating-point branch-with-prediction (FBPFC) instruction.  |   |                   |   |                                       |   |                                   |
| disp22, disp30 | These 22-bit and 30-bit fields are word-aligned, sign-extended, PC-relative displacements for a branch or call, respectively.  |   |                   |   |                                       |   |                                   |
| fcn            | This 5-bit field provides additional opcode bits to encode the DONE, RETRY, and PREFETCH(A) instructions.  |   |                   |   |                                       |   |                                   |
| i              | The i bit selects the second operand for integer arithmetic and load/store instructions. If i = 0, then the operand is r[rs2]. If i = 1, then the operand is simm10, simm11, or simm13, depending on the instruction, sign-extended to 64 bits.  |   |                   |   |                                       |   |                                   |
| imm22          | This 22-bit field is a constant that SETHI places in bits 31:10 of a destination register.   |   |                   |   |                                       |   |                                   |
| imm_asi        | This 8-bit field is the ASI in instructions that access alternate space.   |   |                   |   |                                       |   |                                   |
| mmask          | This 4-bit field imposes order constraints on memory references appearing before and after a MEMBAR instruction.   |   |                   |   |                                       |   |                                   |
| op, op2        | These 2-bit and 3-bit fields encode the three major formats and the Format 2 instructions.   |   |                   |   |                                       |   |                                   |
| op3            | This 6-bit field (together with one bit from op) encodes the Format 3 instructions.  |   |                   |   |                                       |   |                                   |
| opf            | This 9-bit field encodes the operation for a floating-point operate (FPop) instruction.  |   |                   |   |                                       |   |                                   |
| opf_cc         | Specifies the condition codes to be used in FMOVCC instructions. See field cc0, cc1, and cc2 for details.  |   |                   |   |                                       |   |                                   |
| opf_low        | This 6-bit field encodes the specific operation for a Move Floating-Point Register if condition is satisfied (FMOVCC) or Move Floating-Point Register if contents of integer register match condition (FMOVr) instruction.   |   |                   |   |                                       |   |                                   |
| p              | This 1-bit field encodes static prediction for BPCC and FBPFC instructions; branch prediction bit (p) encodings are shown below. <table border="1" data-bbox="458 999 1018 1107" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>p</th> <th>Branch Prediction</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Predict that branch will not be taken</td> </tr> <tr> <td>1</td> <td>Predict that branch will be taken</td> </tr> </tbody> </table> | p | Branch Prediction | 0 | Predict that branch will not be taken | 1 | Predict that branch will be taken |
| p              | Branch Prediction  |   |                   |   |                                       |   |                                   |
| 0              | Predict that branch will not be taken  |   |                   |   |                                       |   |                                   |
| 1              | Predict that branch will be taken  |   |                   |   |                                       |   |                                   |
| rcond          | This 3-bit field selects the register-contents condition to test for a move, based on register contents (MOVr or FMOVr) instruction or a Branch on Register Contents with Prediction (BPr) instruction.  |   |                   |   |                                       |   |                                   |
| rd             | This 5-bit field is the address of the destination (or source) r or f register(s) for a load, arithmetic, or store instruction.  |   |                   |   |                                       |   |                                   |
| rs1            | This 5-bit field is the address of the first r or f register(s) source operand.  |   |                   |   |                                       |   |                                   |
| rs2            | This 5-bit field is the address of the second r or f register(s) source operand with i = 0.  |   |                   |   |                                       |   |                                   |
| shcnt32        | This 5-bit field provides the shift count for 32-bit shift instructions.   |   |                   |   |                                       |   |                                   |
| shcnt64        | This 6-bit field provides the shift count for 64-bit shift instructions.   |   |                   |   |                                       |   |                                   |
| simm10         | This 10-bit field is an immediate value that is sign-extended to 64 bits and used as the second ALU operand for a MOVr instruction when i = 1.   |   |                   |   |                                       |   |                                   |
| simm11         | This 11-bit field is an immediate value that is sign-extended to 64 bits and used as the second ALU operand for a MOVCC instruction when i = 1.  |   |                   |   |                                       |   |                                   |

**TABLE 7-10** Instruction Field Interpretation (3 of 3)

| Field    | Description  |
|----------|--|
| simm13   | This 13-bit field is an immediate value that is sign-extended to 64 bits and used as the second ALU operand for an integer arithmetic instruction or for a load/store instruction when $i = 1$ . |
| sw_trap# | This 7-bit field is an immediate value that is used as the second ALU operand for a Trap on Condition Code instruction.  |
| x        | The x bit selects whether a 32-bit or 64-bit shift will be performed.  |

## 7.12 Reserved Opcodes and Instruction Fields

An attempt to execute an opcode to which no instruction is assigned causes a trap, specifically:

- Attempting to execute a reserved FPop (floating-point opcode) causes a *fp\_exception\_other* exception (with `FSR.ftt = unimplemented_FPop`).
- Attempting to execute any other reserved opcode causes an *illegal\_instruction* exception.
- Attempting to execute an FPop with a nonzero value in a reserved instruction field causes a *fp\_exception\_other* exception (with `FSR.ftt = unimplemented_FPop`).<sup>1</sup>
- Attempting to execute a TCC instruction with a nonzero value in a reserved instruction field causes an *illegal\_instruction* exception.
- Attempting to execute any other instruction with a nonzero value in a reserved instruction field causes an *illegal\_instruction* exception.<sup>1</sup>

### 7.12.1 Summary of Unimplemented Instructions

Certain SPARC V9 instructions are not implemented in hardware in the processor. Executing any of these instructions results in the behavior described in TABLE 7-11.

**TABLE 7-11** Processor Actions on Unimplemented Instructions

| Instructions                  | Trap Taken                 | Processor-Specific Behavior               | Operating System Response     |
|-------------------------------|----------------------------|---|-------------------------------|
| Quad FPOps (including FdMULq) | <i>fp_exception_other</i>  | <code>FSR.ftt = unimplemented_FPop</code> | Emulates Instruction          |
| POPC                          | <i>illegal_instruction</i> | None                                      | Emulates Instruction          |
| RDPR FQ                       | <i>illegal_instruction</i> | None                                      | Skips instruction and returns |
| LDQF                          | <i>illegal_instruction</i> | None                                      | Emulates Instruction          |
| STQF                          | <i>illegal_instruction</i> | None                                      | Emulates Instruction          |

1. Although it is recommended that this exception is generated, an UltraSPARC III Cu User's Manual implementation may ignore the contents of reserved instruction fields (in instructions other than TCC).

If a trap does not occur and the instruction is not a control transfer, the next program counter ( $nPC$ ) is copied into the  $PC$ , and the  $nPC$  is incremented by four (ignoring overflow, if any). If the instruction is a control transfer instruction, the  $nPC$  is copied into the  $PC$  and the target address is written to  $nPC$ . Thus, the two program counters provide for a delayed-branch execution model.

For each instruction access and each normal data access, the IU appends an 8-bit address space identifier (ASI) to the 64-bit memory address. Load/store alternate instructions (see Section 7.2.4, “Address Space Identifiers”) can provide an arbitrary ASI with their data addresses or can use the ASI value currently contained in the ASI register.

---

## 7.13 Big/Little-endian Addressing

The processor uses big-endian byte order for all instruction accesses and, by default, for data accesses.

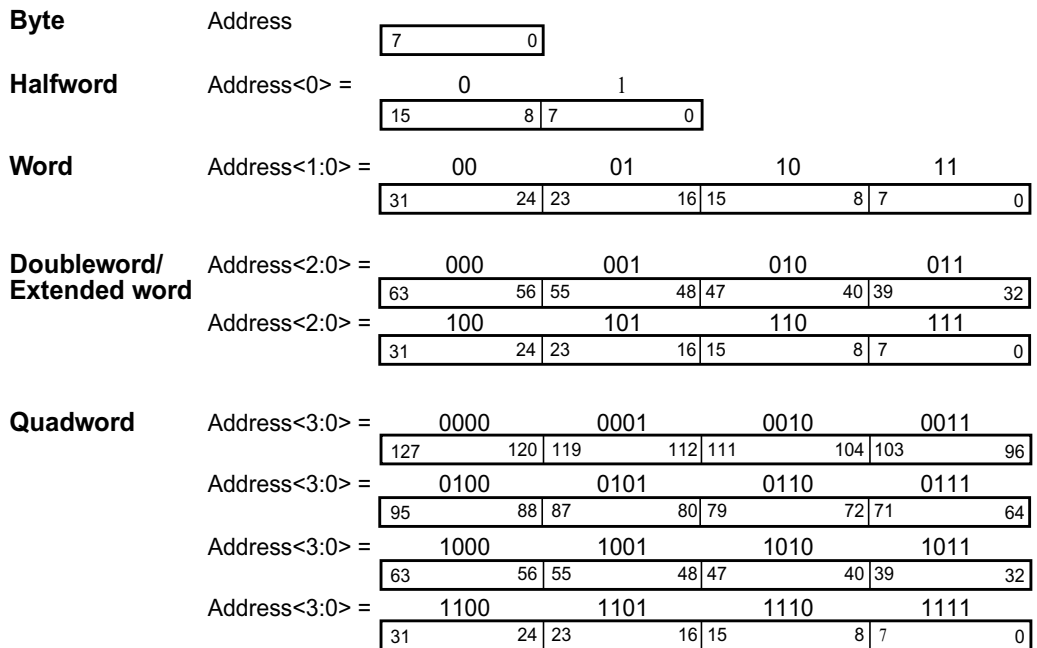
It is possible to access data in little-endian format by using selected ASIs. See Chapter 8, “Address Space Identifiers” for details.

It is also possible to change the default byte order for implicit data accesses.

### 7.13.1 Big-endian Addressing Convention

Within a multiple-byte integer, the byte with the smallest address is the most significant; a byte’s significance decreases as its address increases. The big-endian addressing conventions are illustrated in FIGURE 7-4 and described below the figure.





**FIGURE 7-4** Big-endian Addressing Conventions

- big-endian byte** A load/store byte instruction accesses the addressed byte in both big-endian and little-endian modes.
- big-endian halfword** For a load/store halfword instruction, 2 bytes are accessed. The most significant byte (bits 15–8) is accessed at the address specified in the instruction; the least significant byte (bits 7–0) is accessed at the address + 1.
- big-endian word** For a load/store word instruction, 4 bytes are accessed. The most significant byte (bits 31–24) is accessed at the address specified in the instruction; the least significant byte (bits 7–0) is accessed at the address + 3.
- big-endian doubleword or extended word** For a load/store extended or floating-point load/store double instruction, 8 bytes are accessed. The most significant byte (bits 63–56) is accessed at the address specified in the instruction; the least significant byte (bits 7–0) is accessed at the address + 7.

For the deprecated integer load/store double instructions (LDD/STD), two big-endian words are accessed. The word at the address specified in the instruction corresponds to the even register specified in the instruction; the word at address + 4 corresponds to the following odd-numbered register.

**big-endian quadword** For a load/store quadword instruction, 16 bytes are accessed. The most significant byte (bits 127–120) is accessed at the address specified in the instruction; the least significant byte (bits 7–0) is accessed at the address + 15.

## 7.13.2 Little-endian Addressing Convention

Within a multiple-byte integer, the byte with the smallest address is the least significant; a byte's significance increases as its address increases. The little-endian addressing conventions are illustrated in FIGURE 7-5 and defined below the figure.

|                                      |                |   |      |      |      |      |     |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
|--------------------------------------|----------------|---|------|------|------|------|-----|---|----|---|----|----|----|----|------|------|------|------|----|----|----|----|----|----|----|----|------|------|------|------|----|----|----|----|----|----|----|----|------|------|------|------|-----|----|-----|-----|-----|-----|-----|-----|
| <b>Byte</b>                          | Address        | <table border="1" style="display: inline-table; vertical-align: middle;"><tr><td style="width: 20px;">7</td><td style="width: 20px;">0</td></tr></table>  | 7    | 0    |      |      |     |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 7                                    | 0              |   |      |      |      |      |     |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| <b>Halfword</b>                      | Address<0> =   | <table style="display: inline-table; vertical-align: middle;"><tr><td style="width: 20px;">0</td><td style="width: 20px;">1</td></tr><tr><td style="border: 1px solid black; width: 20px;">7</td><td style="border: 1px solid black; width: 20px;">0</td><td style="border: 1px solid black; width: 20px;">15</td><td style="border: 1px solid black; width: 20px;">8</td></tr></table>   | 0    | 1    | 7    | 0    | 15  | 8 |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 0                                    | 1              |   |      |      |      |      |     |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 7                                    | 0              | 15  | 8    |      |      |      |     |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| <b>Word</b>                          | Address<1:0> = | <table style="display: inline-table; vertical-align: middle;"><tr><td style="width: 20px;">00</td><td style="width: 20px;">01</td><td style="width: 20px;">10</td><td style="width: 20px;">11</td></tr><tr><td style="border: 1px solid black; width: 20px;">7</td><td style="border: 1px solid black; width: 20px;">0</td><td style="border: 1px solid black; width: 20px;">15</td><td style="border: 1px solid black; width: 20px;">8</td><td style="border: 1px solid black; width: 20px;">23</td><td style="border: 1px solid black; width: 20px;">16</td><td style="border: 1px solid black; width: 20px;">31</td><td style="border: 1px solid black; width: 20px;">24</td></tr></table>   | 00   | 01   | 10   | 11   | 7   | 0 | 15 | 8 | 23 | 16 | 31 | 24 |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 00                                   | 01             | 10  | 11   |      |      |      |     |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 7                                    | 0              | 15  | 8    | 23   | 16   | 31   | 24  |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| <b>Doubleword/<br/>Extended word</b> | Address<2:0> = | <table style="display: inline-table; vertical-align: middle;"><tr><td style="width: 20px;">000</td><td style="width: 20px;">001</td><td style="width: 20px;">010</td><td style="width: 20px;">011</td></tr><tr><td style="border: 1px solid black; width: 20px;">7</td><td style="border: 1px solid black; width: 20px;">0</td><td style="border: 1px solid black; width: 20px;">15</td><td style="border: 1px solid black; width: 20px;">8</td><td style="border: 1px solid black; width: 20px;">23</td><td style="border: 1px solid black; width: 20px;">16</td><td style="border: 1px solid black; width: 20px;">31</td><td style="border: 1px solid black; width: 20px;">24</td></tr></table><br>Address<2:0> = <table style="display: inline-table; vertical-align: middle;"><tr><td style="width: 20px;">100</td><td style="width: 20px;">101</td><td style="width: 20px;">110</td><td style="width: 20px;">111</td></tr><tr><td style="border: 1px solid black; width: 20px;">39</td><td style="border: 1px solid black; width: 20px;">32</td><td style="border: 1px solid black; width: 20px;">47</td><td style="border: 1px solid black; width: 20px;">40</td><td style="border: 1px solid black; width: 20px;">55</td><td style="border: 1px solid black; width: 20px;">48</td><td style="border: 1px solid black; width: 20px;">63</td><td style="border: 1px solid black; width: 20px;">56</td></tr></table>  | 000  | 001  | 010  | 011  | 7   | 0 | 15 | 8 | 23 | 16 | 31 | 24 | 100  | 101  | 110  | 111  | 39 | 32 | 47 | 40 | 55 | 48 | 63 | 56 |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 000                                  | 001            | 010   | 011  |      |      |      |     |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 7                                    | 0              | 15  | 8    | 23   | 16   | 31   | 24  |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 100                                  | 101            | 110   | 111  |      |      |      |     |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 39                                   | 32             | 47  | 40   | 55   | 48   | 63   | 56  |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| <b>Quadword</b>                      | Address<3:0> = | <table style="display: inline-table; vertical-align: middle;"><tr><td style="width: 20px;">0000</td><td style="width: 20px;">0001</td><td style="width: 20px;">0010</td><td style="width: 20px;">0011</td></tr><tr><td style="border: 1px solid black; width: 20px;">7</td><td style="border: 1px solid black; width: 20px;">0</td><td style="border: 1px solid black; width: 20px;">15</td><td style="border: 1px solid black; width: 20px;">8</td><td style="border: 1px solid black; width: 20px;">23</td><td style="border: 1px solid black; width: 20px;">16</td><td style="border: 1px solid black; width: 20px;">31</td><td style="border: 1px solid black; width: 20px;">24</td></tr></table><br>Address<3:0> = <table style="display: inline-table; vertical-align: middle;"><tr><td style="width: 20px;">0100</td><td style="width: 20px;">0101</td><td style="width: 20px;">0110</td><td style="width: 20px;">0111</td></tr><tr><td style="border: 1px solid black; width: 20px;">39</td><td style="border: 1px solid black; width: 20px;">32</td><td style="border: 1px solid black; width: 20px;">47</td><td style="border: 1px solid black; width: 20px;">40</td><td style="border: 1px solid black; width: 20px;">55</td><td style="border: 1px solid black; width: 20px;">48</td><td style="border: 1px solid black; width: 20px;">63</td><td style="border: 1px solid black; width: 20px;">56</td></tr></table><br>Address<3:0> = <table style="display: inline-table; vertical-align: middle;"><tr><td style="width: 20px;">1000</td><td style="width: 20px;">1001</td><td style="width: 20px;">1010</td><td style="width: 20px;">1011</td></tr><tr><td style="border: 1px solid black; width: 20px;">71</td><td style="border: 1px solid black; width: 20px;">64</td><td style="border: 1px solid black; width: 20px;">79</td><td style="border: 1px solid black; width: 20px;">72</td><td style="border: 1px solid black; width: 20px;">87</td><td style="border: 1px solid black; width: 20px;">80</td><td style="border: 1px solid black; width: 20px;">95</td><td style="border: 1px solid black; width: 20px;">88</td></tr></table><br>Address<3:0> = <table style="display: inline-table; vertical-align: middle;"><tr><td style="width: 20px;">1100</td><td style="width: 20px;">1101</td><td style="width: 20px;">1110</td><td style="width: 20px;">1111</td></tr><tr><td style="border: 1px solid black; width: 20px;">103</td><td style="border: 1px solid black; width: 20px;">96</td><td style="border: 1px solid black; width: 20px;">111</td><td style="border: 1px solid black; width: 20px;">104</td><td style="border: 1px solid black; width: 20px;">119</td><td style="border: 1px solid black; width: 20px;">112</td><td style="border: 1px solid black; width: 20px;">127</td><td style="border: 1px solid black; width: 20px;">120</td></tr></table> | 0000 | 0001 | 0010 | 0011 | 7   | 0 | 15 | 8 | 23 | 16 | 31 | 24 | 0100 | 0101 | 0110 | 0111 | 39 | 32 | 47 | 40 | 55 | 48 | 63 | 56 | 1000 | 1001 | 1010 | 1011 | 71 | 64 | 79 | 72 | 87 | 80 | 95 | 88 | 1100 | 1101 | 1110 | 1111 | 103 | 96 | 111 | 104 | 119 | 112 | 127 | 120 |
| 0000                                 | 0001           | 0010  | 0011 |      |      |      |     |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 7                                    | 0              | 15  | 8    | 23   | 16   | 31   | 24  |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 0100                                 | 0101           | 0110  | 0111 |      |      |      |     |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 39                                   | 32             | 47  | 40   | 55   | 48   | 63   | 56  |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 1000                                 | 1001           | 1010  | 1011 |      |      |      |     |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 71                                   | 64             | 79  | 72   | 87   | 80   | 95   | 88  |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 1100                                 | 1101           | 1110  | 1111 |      |      |      |     |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |
| 103                                  | 96             | 111   | 104  | 119  | 112  | 127  | 120 |   |    |   |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |    |    |    |    |    |    |    |    |      |      |      |      |     |    |     |     |     |     |     |     |

**FIGURE 7-5** Little-endian Addressing Conventions

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

**little-endian halfword** For a load/store halfword instruction, 2 bytes are accessed. The least significant byte (bits 7–0) is accessed at the address specified in the instruction; the most significant byte (bits 15–8) is accessed at the address + 1.

**little-endian word** For a load/store word instruction, 4 bytes are accessed. The least significant byte (bits 7–0) is accessed at the address specified in the instruction; the most significant byte (bits 31–24) is accessed at the address + 3.

**little-endian doubleword  
or extended word**

For a load/store extended or floating-point load/store double instruction, 8 bytes are accessed. The least significant byte (bits 7–0) is accessed at the address specified in the instruction; the most significant byte (bits 63–56) is accessed at the address + 7.

For the deprecated integer load/store double instructions (LDD/STD), two little-endian words are accessed. The word at the address specified in the instruction corresponds to the even register in the instruction; the word at the address specified in the instruction plus four corresponds to the following odd-numbered register. With respect to little-endian memory, an LDD (STD) instruction behaves as if it is composed of two 32-bit loads (stores), each of which is byte-swapped independently before being written into each destination register (memory word).

**little-endian quadword**

For a load/store quadword instruction, 16 bytes are accessed. The least significant byte (bits 7–0) is accessed at the address specified in the instruction; the most significant byte (bits 127–120) is accessed at the address + 15.



---

# Address Space Identifiers

---

The address space identifiers (ASIs) are described in the following sections:

- ASI Introduction
- ASI Heredity
- ASI Groups
- Instructions Associated with the ASIs
- Using ASIs
- List of ASI Definitions
- Special Memory Access ASIs

---

## 8.1 ASI Introduction

Every instruction fetch, data load, or data store operation is specified by a 64-bit virtual address. In the SPARC architecture, there is always an ASI along with the virtual address. In most cases the ASI is implicit, but can be explicitly specified when appropriate. The ASI can provide rules for how the Memory Management Unit (MMU) should translate a virtual address to a physical address. The ASI can provide attributes for how an operation should be performed. The ASI can also be used to address internal state of the processor.

---

**SPARC Compatibility Note** – The SPARC V9 architecture has also extended the limit of virtual addresses from 32-bit (SPARC V8) to 64-bit for each address space. The SPARC V9 architecture supports 32-bit addressing through masking of the upper 32 bits to zero when the address mask (*AM*) bit in the *PSTATE* register is set.

---

Every instruction fetch, load or store address in the processor has an 8-bit ASI appended to the virtual address (*VA*). The *VA* plus the ASI fully specify the address. These address spaces map to main memory, the processor subsystems, and internal control, status, and

diagnostics registers (CSRs) within a processor. These ASIs are internal to the processor and are not visible outside. ASIs can create special transactions on internal processor busses and assert special internal control signals.

For instruction fetches and data loads or stores that do not use the load or store alternate instructions, the ASI is an implicit ASI generated by the hardware. If a load alternate or store alternate instruction is used, the value of the ASI can be specified in the `%asi` register or as an immediate value in the instruction.

## 8.1.1 Load/Store Instructions

Not all load/store instructions have explicit ASIs.

## 8.1.2 Processor State

Privileged versus non-privileged mode affects the way some ASIs are interpreted.

## 8.1.3 Default ASIs

Default ASIs are precoded. They are selected based on privileged/non-privileged mode.

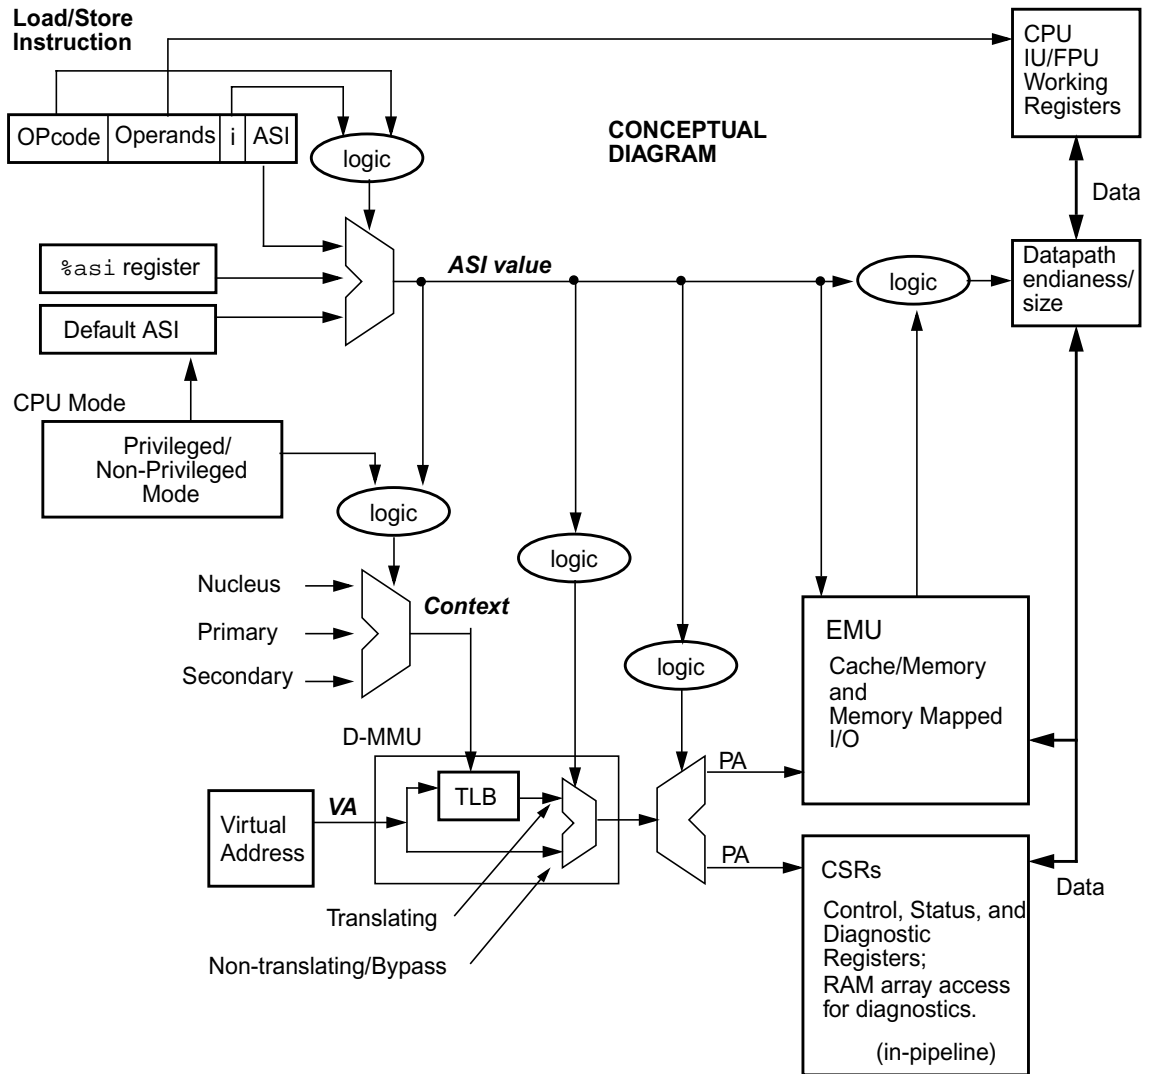
## 8.1.4 Non-Translating and Bypassing ASIs

Non-translating and bypassing ASIs use the VA for the PA. The bypassing ASIs select the processor control, status, and diagnostic registers and RAM arrays for diagnostic purposes. Non-translating ASIs essentially use the VA to address main memory, the system, and memory mapped control, status, and diagnostic registers.

Sometimes, bypassing ASIs are included when referring to non-translating ASIs.

## 8.1.5 Datapath

The datapath supports byte swapping for endianness and variable length data sizes from single bytes (partial store) to 64-byte blocks (block load/store). FIGURE 8-1 shows the datapath and logic for ASIs.

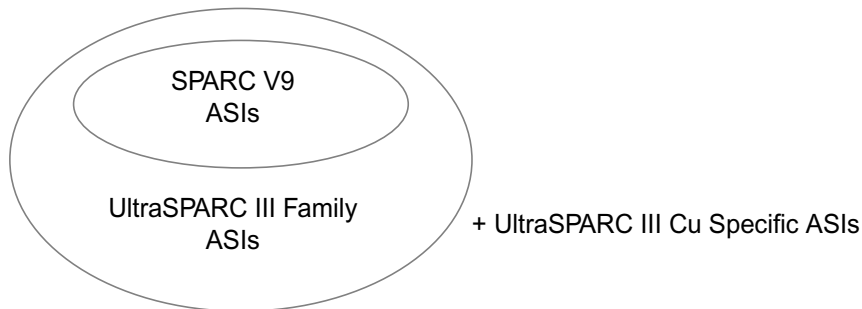


**FIGURE 8-1** ASI Source and Function Conceptual Diagram

---

## 8.2 ASI Heredity

FIGURE 8-2, illustrates how the ASIs are defined at different processor architectural levels. The most universal ASIs are defined by the SPARC V9 architecture. The next level is defined for all processors based on the UltraSPARC III Family of processors. The last level defines ASIs belonging to a particular processor.



**FIGURE 8-2** ASI Groups

ASIs defined by common architecture definitions should work across the processors within that family. All SPARC V9 architecture defined ASIs will work on all processors defined by the SPARC V9 architecture including all UltraSPARC processors, but an ASI defined for a specific processor will not necessarily work on other SPARC V9 processors.

### 8.2.1 SPARC V9 ASIs

The SPARC V9 architecture defines a set of required ASIs for SPARC V9 processors. These ASIs are supported in all UltraSPARC processors. Some of the ASIs have been deprecated in favor of newer ASIs that take advantage of the 64-bit architecture.

### 8.2.2 UltraSPARC UltraSPARC III Family ASIs

The UltraSPARC UltraSPARC III Family currently contains processors similar to the UltraSPARC III Cu processor. All SPARC V9 ASIs will work for the UltraSPARC UltraSPARC III Family of processors.



## 8.2.3 UltraSPARC III Cu Specific ASIs

The UltraSPARC III Cu processor specific ASIs are also defined. ASI  $4A_{16}$  is an example of how one ASI is defined differently for the Sun Fireplane interconnect in the UltraSPARC III Cu processor. All UltraSPARC UltraSPARC III Family ASIs will work with the UltraSPARC III Cu processor. TABLE 8-1 lists processor specific ASIs.

**TABLE 8-1** Processor Specific ASIs

| ASI Value | UltraSPARC III Cu           |
|-----------|-----------------------------|
| 4A        | Fireplane Interconnect CSRs |
| 72        | Memory Control Unit CSRs    |

---

## 8.3 ASI Groups

The ASI is evenly divided into restricted and unrestricted halves, defined by the SPARC V9. ASIs in the range  $00_{16}$ – $7F_{16}$  are restricted. ASIs in the range  $80_{16}$ – $FF_{16}$  are unrestricted. An attempt to access a restricted ASI in non-privileged mode causes a *privileged\_action* trap.

*Normal* or *translating* ASIs cause the CPU's VA to be translated to a physical one by the MMU. *Non-translating*, or *bypassing* ASIs, cause the CPU to not translate the VA; instead, the MMU passes the lower 43 bits of the CPU's virtual addresses as 43-bit physical addresses.

Access restrictions and translating abilities are summarized in TABLE 8-2.

ASIs can be grouped together to help understand the nature of the ASIs. The ASIs that map to a PA are targeted toward physical memory, memory mapped CSRs, and the processor's subsystem, depending on the VA and the ASI value.

The UltraSPARC III Cu processor implements the standard SPARC V9 ASIs and many processor specific ASIs for endian support and address CSR registers.

**Processor Compatibility Note** – In TABLE 8-2, text in **bold** means the ASI was not implemented in UltraSPARC I or UltraSPARC II. Text with ~~strike-through~~ means the ASI was implemented in UltraSPARC I or UltraSPARC II but not in the UltraSPARC III Family.

**TABLE 8-2** ASI Summary Table

| ASI Values   | Destination      | Translating VA to PA | Special Operations  | Architecture Definition            |
|--|------------------|----------------------|---------------------|------------------------------------|
| <b>Restricted, Accessible in Privileged Mode Only</b>                  |                  |                      |                     |                                    |
| 04h, 0Ch   | Physical Address | Translating          |                     | SPARC V9 and UltraSPARC III Family |
| 10h, 11h   |                  |                      |                     | SPARC V9                           |
| 14h, 15h   |                  | Bypassing            |                     | UltraSPARC III Family              |
| 18h, 19h   |                  | Translating          |                     | SPARC V9                           |
| 1Ch, 1Dh   |                  | Bypassing            |                     | UltraSPARC III Family              |
| 24h, 2Ch   |                  | Translating          |                     |                                    |
| <b>30-34h, 38-3Ch, 40-44h, 45-49h</b>                                  | CSR              | Non-translating      |                     | UltraSPARC III Family              |
| 4Ah  | CSR (Bus i/f)    |                      |                     | UltraSPARC III Cu                  |
| 4B-4Eh, 50-5Fh, <b>60h</b> , 66h, 67h, 68h, <del>6Eh</del> , 6Fh       | CSR              |                      |                     | UltraSPARC III Family              |
| 70h, 71h   | Physical Address | Translating          | Block Load/Store    |                                    |
| 72h  | CSR (MCU)        | Non-translating      |                     | UltraSPARC III Cu                  |
| <b>74-75h</b> , 76h, 77h   | CSR              | Non-translating      |                     |                                    |
| 78h, 79h   | Physical Address | Translating          | Block Load/Store    |                                    |
| 7Eh, 7Fh   | CSR              | Non-translating      |                     |                                    |
| <b>Non-Restricted, Accessible in Privileged or Non-privileged Mode</b> |                  |                      |                     |                                    |
| 80h, 81h, 82h, 83h, 88h, 89h, 8Ah, 8Bh                                 | Physical Address | Translating          |                     | SPARC V9                           |
| C0-C5h, C8-CDh   |                  |                      | Partial Store       | UltraSPARC III Family              |
| D0-D3h, D8-DBh   |                  |                      | Short FP Load/Store |                                    |
| E0h, E1h, F0h, F1h, F8h, F9h   |                  |                      | Block Load/Store    |                                    |

---

## 8.4 Instructions Associated with the ASIs

ASIs are used with load and store instructions. Their usage and restrictions are described in TABLE 8-4. Additional information is described in Section 8.5, “Using ASIs.”

### 8.4.1 Block Load and Block Store ASIs

Block load (BLD) and block store (BST) operations are generated by using the LDDFA and STDFA instructions with the  $70_{16}$ ,  $71_{16}$ ,  $78_{16}$ ,  $79_{16}$ ,  $E0_{16}$ ,  $E1_{16}$ ,  $F0_{16}$ ,  $F1_{16}$ ,  $F8_{16}$ , and  $F9_{16}$  ASIs.

If the operand address is not 64-byte aligned, then a *mem\_address\_not\_aligned* exception is generated.

If these ASIs are used with any other instruction, then a *data\_access\_exception* is generated and *mem\_address\_not\_aligned* is not generated.

### 8.4.2 Partial Store ASIs

Partial store operations are generated by using the STDFA instruction with the  $C0_{16}$ – $C5_{16}$  and  $C8_{16}$ – $CD_{16}$  ASIs.

If the operand address is not 8-byte aligned, then a *mem\_address\_not\_aligned* exception is generated and if  $i = 1$  in the instruction, then an *illegal\_instruction* exception is generated instead.

If these ASIs are used with any other instruction, then a *data\_access\_exception* exception is generated and neither a *mem\_address\_not\_aligned* nor a *illegal\_instruction* (for  $i = 1$ ) is generated.

### 8.4.3 Short Floating-Point Load and Store ASIs

Short floating-point load and store operations are generated by using the LDDFA and STDFA instructions with the  $D0_{16}$ – $D3_{16}$  and  $D8_{16}$ – $DB_{16}$  ASIs to load and store byte and halfword values in the floating-point registers.

If the data is not aligned, then a *mem\_address\_not\_aligned* is generated.

If these ASIs are used with any other instruction, a *data\_access\_exception* is generated and *mem\_address\_not\_aligned* will not be generated.

### 8.4.3.1 Halfword Alignment

If the operand address for a D2<sub>16</sub>, D3<sub>16</sub>, DA<sub>16</sub>, or DB<sub>16</sub> ASI (halfword) is not halfword aligned, then a *mem\_address\_not\_aligned* exception is generated.

---

## 8.5 Using ASIs

### 8.5.1 Data Widths

The ASIs for the UltraSPARC III Cu processor and the entire UltraSPARC UltraSPARC III Family of processors are accessible using 64-bit LDXA, STXA, LDDFA, and STDFA instructions, except where noted. SPARC V9 ASIs are accessible using aligned 8-, 16-, 32- and 64-bit load and store (read and write) instructions, except where noted. The UltraSPARC UltraSPARC III Family of processors and specific implementations require 64-bit aligned accesses, except where noted.

### 8.5.2 Operand Alignment

Addresses must align to the boundary of the data width. TABLE 8-3 shows the data size and address offset for operand alignment.

**TABLE 8-3** Operand Alignment

| Data Size             | Address Offset (binary) |
|-----------------------|-------------------------|
| Halfword              | xxxx.xxx0               |
| Word                  | xxxx.xx00               |
| Double word (8 bytes) | xxxx.x000               |
| Quad word (16 bytes)  | xxxx.0000               |
| Block (64 bytes)      | xx00.0000               |

### 8.5.3 Common Exceptions

Using ASIs improperly will generate one of the following exceptions in the CPU.

### 8.5.3.1 data\_access\_exception

When the wrong instruction is used with an ASI, a *data\_access\_exception* is generated. This is sometimes referred to as the *invalid\_ASI\_exception*.

### 8.5.3.2 mem\_address\_not\_aligned

When the address operand does not align to the boundary of the data size, a *mem\_address\_not\_aligned* exception is generated.

### 8.5.3.3 privileged\_action

If a restricted, privileged mode only ASI is accessed in non-privileged mode, then a *privileged\_mode\_exception* is generated.

---

## 8.6 List of ASI Definitions

TABLE 8-4 lists all the ASIs in the UltraSPARC III Cu processor.

**TABLE 8-4** ASI Definitions (1 of 12)

| Value (hex) | Processor   | ASI Name (Suggested Macro Syntax) | Type | VA (hex) | Description  | See Foot notes |
|-------------|---|-----------------------------------|------|----------|--|----------------|
| 00-7F       | <b>Restricted, Accessible in Privileged Mode only</b> |                                   |      |          |  |                |
| 00-03       | —   | —                                 |      |          | Implementation-dependent, Unassigned                           | 1              |
| 04          | SPARC V9 and UltraSPARC III Family                    | ASI_NUCLEUS (ASI_N)               | RW   |          | Implicit address space, nucleus privilege, TL>0                |                |
| 05-0B       | —   | —                                 |      |          | Implementation-dependent, Unassigned                           | 1              |
| 0C          | SPARC V9 and UltraSPARC III Family                    | ASI_NUCLEUS_LITTLE (ASI_NL)       | RW   |          | Implicit address space, nucleus privilege, TL>0, little-endian |                |
| 0D-0F       | —   | —                                 |      |          | Implementation-dependent, Unassigned                           | 1              |
| 10          | SPARC V9  | ASI_AS_IF_USER_PRIMARY (ASI_AIUP) | RW   |          | Primary address space, user privilege                          | 2, 3           |

**TABLE 8-4** ASI Definitions (2 of 12)

| Value (hex) | Processor             | ASI Name (Suggested Macro Syntax)                                    | Type | VA (hex) | Description  | See Foot notes |
|-------------|-----------------------|--|------|----------|--|----------------|
| 11          | SPARC V9              | ASI_AS_IF_USER_SECONDARY (ASI_AIUS)                                  | RW   |          | Secondary address space, user privilege                          | 2, 3           |
| 12–13       | —                     | —  |      |          | Implementation-dependent, Unassigned                             | 1              |
| 14          | UltraSPARC III Family | ASI_PHYS_USE_EC  | RW   |          | Physical address external cacheable only                         | 4, 5           |
| 15          | UltraSPARC III Family | ASI_PHYS_BYPASS_EC_WITH_EBIT   | RW   |          | Physical address, non-cacheable, with side-effect                | 4              |
| 16–17       | —                     | —  |      |          | Implementation-dependent, Unassigned                             | 1              |
| 18          | SPARC V9              | ASI_AS_IF_USER_PRIMARY_LITTLE (ASI_AIUPL)                            | RW   |          | Primary address space, user privilege, little-endian             | 3              |
| 19          | SPARC V9              | ASI_AS_IF_USER_SECONDARY_LITTLE (ASI_AIUSL)                          | RW   |          | Secondary address space, user privilege, little-endian           | 3              |
| 1A–1B       | —                     | —  |      |          | Implementation-dependent, Unassigned                             | 1              |
| 1C          | UltraSPARC III Family | ASI_PHYS_USE_EC_LITTLE (ASI_PHYS_USE_EC_L)                           | RW   |          | Physical address, external cacheable only, little-endian         | 4, 5           |
| 1D          | UltraSPARC III Family | ASI_PHYS_BYPASS_EC_WITH_EBIT_LITTLE (ASI_PHYS_BYPASS_EC_WITH_EBIT_L) | RW   |          | Physical address, non-cacheable, with side-effect, little-endian | 4              |
| 1E–23       | —                     | —  |      |          | Implementation-dependent, Unassigned                             | 1              |
| 24          | UltraSPARC III Family | ASI_NUCLEUS_QUAD_LDD   | R    |          | Cacheable, 128-bit atomic LDDA                                   | 6, 7           |
| 25–2B       | —                     | —  |      |          | Implementation-dependent, Unassigned                             | 1              |
| 2C          | UltraSPARC III Family | ASI_NUCLEUS_QUAD_LDD_LITTLE (ASI_NUCLEUS_QUAD_LDD_L)                 | R    |          | Cacheable, 128-bit atomic LDDA, little-endian                    | 6, 7           |
| 2D–2F       | —                     | —  |      |          | Implementation-dependent, Unassigned                             | 1              |
| 30          | UltraSPARC III Cu     | ASI_PCACHE_STATUS_DATA   | RW   |          | P-cache data status RAM diagnostic access                        |                |
| 31          | UltraSPARC III Cu     | ASI_PCACHE_DATA  | RW   |          | P-cache data RAM diagnostic access                               |                |

TABLE 8-4 ASI Definitions (3 of 12)

| Value (hex) | Processor             | ASI Name (Suggested Macro Syntax)    | Type | VA (hex) | Description                                      | See Foot notes |
|-------------|-----------------------|--------------------------------------|------|----------|--|----------------|
| 32          | UltraSPARC III Cu     | ASI_PCACHE_TAG                       | RW   |          | P-cache tag RAM diagnostic access                |                |
| 33          | UltraSPARC III Cu     | ASI_PCACHE_SNOOP_TAG                 | RW   |          | P-cache snoop tag RAM diagnostic access          |                |
| 34          | UltraSPARC III Cu     | ASI_ATOMIC_QUAD_LDD_PHYS             |      |          |  |                |
| 35–37       | —                     | —                                    |      |          | Implementation-dependent, Unassigned             | 1              |
| 38          | UltraSPARC III Cu     | ASI_WCACHE_VALID_BITS                |      |          | W-cache Valid Bits diagnostic access             |                |
| 39          | UltraSPARC III Cu     | ASI_WCACHE_DATA                      | RW   |          | W-cache data RAM diagnostic access               |                |
| 3A          | UltraSPARC III Cu     | ASI_WCACHE_TAG                       | RW   |          | W-cache tag RAM diagnostic access                |                |
| 3B          | UltraSPARC III Cu     | ASI_WCACHE_SNOOP_TAG                 | RW   |          | W-cache snoop tag RAM diagnostic access          |                |
| 3C          | UltraSPARC III Cu     | ASI_ATOMIC_QUAD_LDD_PHYS_L           |      |          |  | 1              |
| 3D–3F       | —                     | —                                    |      |          | Implementation-dependent, Unassigned             | 1              |
| 40          | UltraSPARC III Cu     | ASI_SRAM_FAST_INIT                   | W    |          | Interface to clean all major SRAM arrays on chip | 9              |
| 41          |                       | <i>Reserved.</i>                     |      |          |  |                |
| 42          | UltraSPARC III Cu     | ASI_DCACHE_INVALIDATE                | W    |          | D-cache Invalidate diagnostic access             |                |
| 43          | UltraSPARC III Cu     | ASI_DCACHE_UTAG                      | RW   |          | D-cache uTag diagnostic access                   |                |
| 44          | UltraSPARC III Cu     | ASI_DCACHE_SNOOP_TAG                 | RW   |          | D-cache snoop tag RAM diagnostic access          |                |
| 45          | UltraSPARC III Family | ASI_DCU_CONTROL_REGISTER (ASI_DCUCR) | RW   | 0        | D-cache Unit Control Register                    |                |
| 46          | UltraSPARC III Cu     | ASI_DCACHE_DATA                      | RW   |          | D-cache data RAM diagnostic access               |                |
| 47          | UltraSPARC III Cu     | ASI_DCACHE_TAG                       | RW   |          | D-cache tag/valid RAM diagnostic access          |                |

**TABLE 8-4** ASI Definitions (4 of 12)

| Value (hex) | Processor             | ASI Name (Suggested Macro Syntax)              | Type | VA (hex) | Description  | See Foot notes |
|-------------|-----------------------|--|------|----------|--|----------------|
| 48          | UltraSPARC III Family | ASI_INTR_DISPATCH_STATUS (ASI_MONDO_SEND_CTRL) | R    | 0        | Interrupt vector dispatch status   | 6              |
| 49          | UltraSPARC III Family | ASI_INTR_RECEIVE (ASI_MONDO_RECEIVE_CTRL)      | RW   | 0        | Interrupt vector receive status  |                |
| 4A          | UltraSPARC III Cu     | ASI_FIREPLANE_CONFIG_REG                       | RW   | 0        | Fireplane interconnect configuration register for the UltraSPARC III processor |                |
|             |                       | ASI_FIREPLANE_ADDRESS_REG                      | RW   | 08       | Fireplane interconnect configuration register for the UltraSPARC III processor |                |
| 4B          | UltraSPARC III Cu     | ASI_ESTATE_ERROR_EN_REG                        | RW   | 0        | E-state error enable register  |                |
| 4C          | UltraSPARC III Family | ASI_ASYNC_FAULT_STATUS (ASI_AFSR)              | RW   | 0        | Asynchronous fault status register   |                |
| 4D          | UltraSPARC III Family | ASI_ASYNC_FAULT_ADDR (ASI_AFAR)                | R    | 0        | Asynchronous fault address register  |                |
| 4E          | UltraSPARC III Cu     | ASI_ECACHE_TAG (ASI_EC_TAG)                    | RW   | <22:6 >  | L2-cache tag state RAM data diagnostic access                                  |                |
|             |                       | ASI_ECACHE_FLUSH                               | R    | [31] =1  |  |                |
| 4F          | —                     | —  |      |          | Implementation-dependent, Unassigned   | 1              |
| 50          | UltraSPARC III Family | ASI_IMMU_TAG_TARGET                            | R    | 0        | I-MMU tag target register  | 6              |
|             |                       | ASI_IMMU_SFSR                                  | RW   | 18       | I-MMU sync fault status register   |                |
|             |                       | ASI_IMMU_TSB_BASE                              | RW   | 28       | I-MMU TSB base register  |                |
|             |                       | ASI_IMMU_TAG_ACCESS                            | RW   | 30       | I-MMU TLB tag access register  |                |
|             |                       | ASI_IMMU_TSB_PEXT_REG                          | RW   | 48       | I-MMU TSB primary extension register   |                |
|             |                       | ASI_IMMU_TSB_SEXT_REG                          | R    | 50       | I-MMU TSB secondary extension register (hardwired to zero)                     |                |
|             |                       | ASI_IMMU_TSB_NEXT_REG                          | RW   | 58       | I-MMU TSB nucleus extension register   |                |
| 51          | UltraSPARC III Family | ASI_IMMU_TSB_8KB_PTR_REG                       | R    | 0        | I-MMU TSB 8 KB pointer register  | 6              |



**TABLE 8-4** ASI Definitions (5 of 12)

| Value (hex) | Processor             | ASI Name (Suggested Macro Syntax) | Type | VA (hex)          | Description   | See Foot notes |
|-------------|-----------------------|-----------------------------------|------|-------------------|---|----------------|
| 52          | UltraSPARC III Family | ASI_IMM_TSB_64KB_PTR_REG          | R    | 0                 | I-MMU TSB 64 KB pointer register  | 6              |
| 53          | UltraSPARC III Family | ASI_SERIAL_ID                     | R    |                   | Internal testability <sup>7</sup> (also known as ASI_DEVICE_ID+SERIAL_ID) | 6, 8           |
| 54          | UltraSPARC III Family | ASI_ITLB_DATA_IN_REG              | W    | 0                 | I-MMU TLB data in register  | 9              |
| 55          | UltraSPARC III Family | ASI_ITLB_DATA_ACCESS_REG          | RW   | 0 – 2.0F F8       | I-MMU TLB data access register  |                |
|             | UltraSPARC III Cu     | ASI_ITLB_CAM_ADDRESS_REG          | RW   | 4.000 0 – 6.0F F8 | I-MMU TLB CAM diagnostic register   |                |
| 56          | UltraSPARC III Family | ASI_ITLB_TAG_READ_REG             | R    | 0 – 2.0F F8       | I-MMU TLB tag read register; data access port to RAM array                | 6              |
| 57          | UltraSPARC III Family | ASI_IMMU_DEMAP                    | W    |                   | I-MMU TLB demap operations  | 9              |

**TABLE 8-4** ASI Definitions (6 of 12)

| Value (hex) | Processor             | ASI Name (Suggested Macro Syntax) | Type | VA (hex) | Description  | See Foot notes |
|-------------|-----------------------|-----------------------------------|------|----------|--|----------------|
| 58          | UltraSPARC III Family | ASI_DMMU_TAG_TARGET_REG           | R    | 0        | D-MMU TSB tag target register  |                |
|             |                       | ASI_PRIMARY_CONTEXT_REG           | RW   | 8        | I/D-MMU primary context register (this ASI register is shared between I-MMU and MMU) |                |
|             |                       | ASI_SECONDARY_CONTEXT_REG         | RW   | 10       | D-MMU secondary context register   |                |
|             |                       | ASI_DMMU_SFSR                     | RW   | 18       | D-MMU sync fault status register (D-SFSR)  |                |
|             |                       | ASI_DMMU_SFAR                     | R    | 20       | D-MMU sync fault address register (D-SFAR)   |                |
|             |                       | ASI_DMMU_TSB_BASE                 | RW   | 28       | D-MMU TSB base address register  |                |
|             |                       | ASI_DMMU_TAG_ACCESS               | RW   | 30       | D-MMU TLB tag access register  |                |
|             |                       | ASI_DMMU_VA_WATCHPOINT_REG        | RW   | 38       | D-MMU VA data watchpoint register  |                |
|             |                       | ASI_DMMU_PA_WATCHPOINT_REG        | RW   | 40       | D-MMU PA data watchpoint register  |                |
|             |                       | ASI_DMMU_TSB_PEXT_REG             | RW   | 48       | D-MMU TSB primary extension register   |                |
|             |                       | ASI_DMMU_TSB_SEXT_REG             | RW   | 50       | D-MMU TSB secondary extension register   |                |
|             |                       | ASI_DMMU_TSB_NEXT_REG             | RW   | 58       | D-MMU TSB nucleus extension register   |                |
| 59          | UltraSPARC III Family | ASI_DMMU_TSB_8KB_PTR_REG          | R    | 0        | D-MMU TSB 8 KB pointer register  | 6              |
| 5A          | UltraSPARC III Family | ASI_DMMU_TSB_64KB_PTR_REG         | R    | 0        | D-MMU TSB 64 KB pointer register   | 6              |
| 5B          | UltraSPARC III Family | ASI_DMMU_TSB_DIRECT_PTR_REG       | R    | 0        | D-MMU TSB direct pointer register  | 6              |
| 5C          | UltraSPARC III Family | ASI_DTLB_DATA_IN_REG              | W    | 0        | D-MMU TLB data in register   | 9              |

TABLE 8-4 ASI Definitions (7 of 12)

| Value (hex) | Processor             | ASI Name (Suggested Macro Syntax)             | Type | VA (hex)           | Description   | See Foot notes |
|-------------|-----------------------|---|------|--------------------|---|----------------|
| 5D          | UltraSPARC III Family | ASI_DTLB_DATA_ACCESS_REG                      | RW   | 0 to 2.0F F8       | D-MMU TLB data access register; Data access port to RAM array                           |                |
|             |                       | ASI_TLB_CAM_ACCESS_REG                        | RW   | 4.000 0 to 6.0F F8 | D-MMU TLB CAM diagnostic registers  |                |
| 5E          | UltraSPARC III Family | ASI_DTLB_TAG_READ_REG                         | R    | <17:0 >            | D-MMU TLB tag read register   | 6              |
| 5F          | UltraSPARC III Family | ASI_DMMU_DEMAP                                | W    |                    | D-MMU TLB demap operations  | 9              |
| 60          | UltraSPARC III Family | ASI_IIU_INST_TRAP                             | RW   | 0                  | Instruction breakpoint register   |                |
| 61–65       | —                     | —   |      |                    | Implementation-dependent, Unassigned  | 1              |
| 66          | UltraSPARC III Cu     | ASI_ICACHE_INSTR (ASI_IC_INSTR)               | RW   |                    | I-cache data RAM diagnostics access   |                |
| 67          | UltraSPARC III Cu     | ASI_ICACHE_TAG (ASI_IC_TAG)                   | RW   |                    | I-cache tag/valid RAM diagnostics access  |                |
| 68          | UltraSPARC III Cu     | ASI_ICACHE_SNOOP_TAG (ASI_IC_STAG)            | RW   |                    | I-cache snoop tag RAM diagnostics access  |                |
| 69–6E       | —                     | —   |      |                    | Implementation-dependent, Unassigned  | 1              |
| 6F          | Future                | (Reserved for ASI_BRANCH_PREDICTION_ARRAY)    | RW   |                    | Planned for future processor  |                |
| 70          | UltraSPARC III Family | ASI_BLOCK_AS_IF_USER_PRIMARY (ASI_BLK_AIUP)   | RW   |                    | Primary address space, block load/store, user privilege                                 | 3, 10          |
| 71          | UltraSPARC III Family | ASI_BLOCK_AS_IF_USER_SECONDARY (ASI_BLK_AIUS) | RW   |                    | Secondary address space, block load/store, user privilege                               | 3, 10          |
| 72          | UltraSPARC III Cu     | ASI_MCU_TIMING[1:4]_REG                       | RW   | 0, 8, 10, 18       | Memory Control Unit (MCU) programming registers; these registers are also memory mapped |                |
|             |                       | ASI_MCU_ADR_DEC[1:4]_REG                      | RW   | 20, 28, 30, 38     |   |                |
|             |                       | ASI_MCU_ADR_CNTL_REG                          | RW   | 40                 |   |                |

**TABLE 8-4** ASI Definitions (8 of 12)

| Value (hex) | Processor             | ASI Name (Suggested Macro Syntax)                     | Type | VA (hex) | Description  | See Foot notes |
|-------------|-----------------------|---|------|----------|--|----------------|
| 73          | —                     | —   |      |          | Implementation-dependent, Unassigned                                     |                |
| 74          | UltraSPARC III Cu     | ASI_ECACHE_DATA                                       | RW   |          | L2-cache data staging register   |                |
| 75          | UltraSPARC III Cu     | ASI_ECACHE_CONTROL (ASI_EC_CTRL)                      | RW   | 0        | L2-cache control register  |                |
| 76          | UltraSPARC III Cu     | ASI_ECACHE_W (ASI_EC_W)                               | W    |          | L2-cache data RAM diagnostic write access                                |                |
| 77          | UltraSPARC III Family | ASI_INTR_DATA0_W                                      | W    | 40       | Outgoing interrupt vector data Register 0 H                              | 9              |
|             |                       | ASI_INTR_DATA1_W                                      | W    | 48       | Outgoing interrupt vector data Register 0 L                              | 9              |
|             |                       | ASI_INTR_DATA2_W                                      | W    | 50       | Outgoing interrupt vector data Register 1 H                              | 9              |
|             |                       | ASI_INTR_DATA3_W                                      | W    | 58       | Outgoing interrupt vector data Register 1 L                              | 9              |
|             |                       | ASI_INTR_DATA4_W                                      | W    | 60       | Outgoing interrupt vector data Register 2 H                              | 9              |
|             |                       | ASI_INTR_DATA5_W                                      | W    | 68       | Outgoing interrupt vector data Register 2 L                              | 9              |
|             |                       | ASI_INTR_DISPATCH_W                                   | W    | 70       | Interrupt vector dispatch  | 9              |
|             |                       | ASI_INTR_DATA6_W                                      | W    | 80       | Outgoing interrupt vector data Register 3 H                              | 9              |
|             |                       | ASI_INTR_DATA7_W                                      | W    | 88       | Outgoing interrupt vector data Register 3 L                              | 9              |
|             |                       | ASI_INTR_DISPATCH_W                                   | W    |          | Interrupt vector dispatch: 01.0000.0070–8F.FFFF.C070                     | 9              |
| 78          | UltraSPARC III Family | ASI_BLOCK_AS_IF_USER_PRIMARY_LITTLE (ASI_BLK_AIUPL)   | RW   | 0        | Primary address space, block load/store, user privilege, little-endian   | 3, 10          |
| 79          | UltraSPARC III Family | ASI_BLOCK_AS_IF_USER_SECONDARY_LITTLE (ASI_BLK_AIUSL) | RW   | 0        | Secondary address space, block load/store, user privilege, little-endian | 3, 10          |
| 7A–7D       | —                     | —   |      |          | Implementation-dependent, Unassigned                                     | 1              |

TABLE 8-4 ASI Definitions (9 of 12)

| Value (hex) | Processor  | ASI Name (Suggested Macro Syntax)      | Type | VA (hex) | Description                                    | See Foot notes |
|-------------|--|--|------|----------|--|----------------|
| 7E          | UltraSPARC III Cu  | ASI_ECACHE_R (ASI_EC_R)                | R    | 0        | L2-cache data RAM diagnostic read access       |                |
| 7F          | UltraSPARC III Family  | ASI_INTR_DATA0_R                       | R    | 40       | Incoming interrupt vector data Register 0 H    | 6              |
|             |  | ASI_INTR_DATA1_R                       | R    | 48       | Incoming interrupt vector data Register 0 L    | 6              |
|             |  | ASI_INTR_DATA2_R                       | R    | 50       | Incoming interrupt vector data Register 1 H    | 6              |
|             |  | ASI_INTR_DATA3_R                       | R    | 58       | Incoming interrupt vector data Register 1 L    | 6              |
|             |  | ASI_INTR_DATA4_R                       | R    | 60       | Incoming interrupt vector data Register 2 H    | 6              |
|             |  | ASI_INTR_DATA5_R                       | R    | 68       | Incoming interrupt vector data Register 2 L    | 6              |
|             |  | ASI_INTR_DATA6_R                       | R    | 80       | Incoming interrupt vector data Register 3 H    | 6              |
|             |  | ASI_INTR_DATA7_R                       | R    | 88       | Incoming interrupt vector data Register 3 L    | 6              |
| 80–FF       | <b>Non-Restricted, Accessible in Privileged or Non-privileged Mode</b> |  |      |          |  |                |
| 80          | SPARC V9   | ASI_PRIMARY (ASI_P)                    | RW   |          | Implicit primary address space                 |                |
| 81          | SPARC V9   | ASI_SECONDARY (ASI_S)                  | RW   |          | Secondary address space                        |                |
| 82          | SPARC V9   | ASI_PRIMARY_NO_FAULT (ASI_PNF)         | R    |          | Primary address space, no fault                | 4, 6, 11       |
| 83          | SPARC V9   | ASI_SECONDARY_NO_FAULT (ASI_SNF)       | R    |          | Secondary address space, no fault              | 4, 6, 11       |
| 84–87       | —  | —                                      |      |          | Implementation-dependent, Unassigned           | 1              |
| 88          | SPARC V9   | ASI_PRIMARY_LITTLE (ASI_PL)            | RW   |          | Implicit primary address space, little-endian  |                |
| 89          | SPARC V9   | ASI_SECONDARY_LITTLE (ASI_SL)          | RW   |          | Secondary address space, little-endian         |                |
| 8A          | SPARC V9   | ASI_PRIMARY_NO_FAULT_LITTLE (ASI_PNFL) | R    |          | Primary address space, no fault, little-endian | 4, 6, 11       |

**TABLE 8-4** ASI Definitions (10 of 12)

| Value (hex) | Processor             | ASI Name (Suggested Macro Syntax)         | Type | VA (hex) | Description  | See Foot notes |
|-------------|-----------------------|---|------|----------|--|----------------|
| 8B          | SPARC V9              | ASI_SECONDARY_NO_FAULT_LITTLE (ASI_SNFL)  | R    |          | Secondary address space, no fault, little-endian               | 4, 6, 11       |
| 8C–BF       | —                     | —   |      |          | Implementation-dependent, Unassigned                           | 1              |
| C0          | UltraSPARC III Family | ASI_PST8_PRIMARY (ASI_PST8_P)             | W    |          | Primary address space, 8x8-bit partial store                   | 12             |
| C1          | UltraSPARC III Family | ASI_PST8_SECONDARY (ASI_PST8_S)           | W    |          | Secondary address space, 8x8-bit partial store                 | 12             |
| C2          | UltraSPARC III Family | ASI_PST16_PRIMARY (ASI_PST16_P)           | W    |          | Primary address space, 4x16-bit partial store                  | 12             |
| C3          | UltraSPARC III Family | ASI_PST16_SECONDARY (ASI_PST16_S)         | W    |          | Secondary address space, 4x16-bit partial store                | 12             |
| C4          | UltraSPARC III Family | ASI_PST32_PRIMARY (ASI_PST32_P)           | W    |          | Primary address space, 2x32-bit partial store                  | 12             |
| C5          | UltraSPARC III Family | ASI_PST32_SECONDARY (ASI_PST32_S)         | W    |          | Secondary address space, 2x32-bit partial store                | 12             |
| C6–C7       | —                     | —   |      |          | Implementation-dependent, Unassigned                           | 1              |
| C8          | UltraSPARC III Family | ASI_PST8_PRIMARY_LITTLE (ASI_PST8_PL)     | W    |          | Primary address space, 8x8-bit partial store, little-endian    | 12             |
| C9          | UltraSPARC III Family | ASI_PST8_SECONDARY_LITTLE (ASI_PST8_SL)   | W    |          | Secondary address space, 8x8-bit partial store, little-endian  | 12             |
| CA          | UltraSPARC III Family | ASI_PST16_PRIMARY_LITTLE (ASI_PST16_PL)   | W    |          | Primary address space, 4x16-bit partial store, little-endian   | 12             |
| CB          | UltraSPARC III Family | ASI_PST16_SECONDARY_LITTLE (ASI_PST16_SL) | W    |          | Secondary address space, 4x16-bit partial store, little-endian | 12             |
| CC          | UltraSPARC III Family | ASI_PST32_PRIMARY_LITTLE (ASI_PST32_PL)   | W    |          | Primary address space, 2x32-bit partial store, little-endian   | 12             |
| CD          | UltraSPARC III Family | ASI_PST32_SECONDARY_LITTLE (ASI_PST32_SL) | W    |          | Second address space, 2x32-bit partial store, little-endian    | 12             |

**TABLE 8-4** ASI Definitions (11 of 12)

| Value (hex) | Processor             | ASI Name (Suggested Macro Syntax)             | Type | VA (hex) | Description  | See Foot notes |
|-------------|-----------------------|---|------|----------|--|----------------|
| CE–CF       | —                     | —   |      |          | Implementation-dependent, Unassigned                                       | 1              |
| D0          | UltraSPARC III Family | ASI_FL8_PRIMARY (ASI_FL8_P)                   | RW   |          | Primary address space, one 8-bit floating-point load/store                 | 13             |
| D1          | UltraSPARC III Family | ASI_FL8_SECONDARY (ASI_FL8_S)                 | RW   |          | Second address space, one 8-bit floating-point load/store                  | 13             |
| D2          | UltraSPARC III Family | ASI_FL16_PRIMARY (ASI_FL16_P)                 | RW   |          | Primary address space, one 16-bit floating-point load/store                | 13             |
| D3          | UltraSPARC III Family | ASI_FL16_SECONDARY (ASI_FL16_S)               | RW   |          | Second address space, one 16-bit floating-point load/store                 | 13             |
| D4–D7       | —                     | —   |      |          | Implementation-dependent, Unassigned                                       | 1              |
| D8          | UltraSPARC III Family | ASI_FL8_PRIMARY_LITTLE (ASI_FL8_PL)           | RW   |          | Primary address space, one 8-bit floating-point load/store, little-endian  | 13             |
| D9          | UltraSPARC III Family | ASI_FL8_SECONDARY_LITTLE (ASI_FL8_SL)         | RW   |          | Second address space, one 8-bit floating-point load/store, little-endian   | 13             |
| DA          | UltraSPARC III Family | ASI_FL16_PRIMARY_LITTLE (ASI_FL16_PL)         | RW   |          | Primary address space, one 16-bit floating-point load/store, little-endian | 13             |
| DB          | UltraSPARC III Family | ASI_FL16_SECONDARY_LITTLE (ASI_FL16_SL)       | RW   |          | Second address space, one 16-bit floating-point load/store, little-endian  | 13             |
| DC–DF       | —                     | —   |      |          | Implementation-dependent, Unassigned                                       | 1              |
| E0          | UltraSPARC III Family | ASI_BLOCK_COMMIT_PRIMARY (ASI_BLK_COMMIT_P)   | W    |          | Primary address space, 8x8 - byte block store commit operation             | 9              |
| E1          | UltraSPARC III Family | ASI_BLOCK_COMMIT_SECONDARY (ASI_BLK_COMMIT_S) | W    |          | Secondary address space, 8x8 - byte block store commit operation           | 9              |
| E2–EE       | —                     | —   |      |          | Implementation-dependent, Unassigned                                       | 1              |
| EF          | Future                | (Reserved for ASI_BARRIER_SYNC)               |      |          | Planned for future processor   |                |

**TABLE 8-4** ASI Definitions (12 of 12)

| Value (hex) | Processor             | ASI Name (Suggested Macro Syntax)       | Type | VA (hex) | Description  | See Foot notes |
|-------------|-----------------------|---|------|----------|--|----------------|
| F0          | UltraSPARC III Family | ASI_BLOCK_PRIMARY (ASI_BLK_P)           | RW   |          | Primary address space, 8x8 - byte block load/store       | 9              |
| F1          | UltraSPARC III Family | ASI_BLOCK_SECONDARY (ASI_BLK_S)         | RW   |          | Secondary address space, block load/store                | 9              |
| F2–F7       | —                     | —                                       |      |          | Implementation-dependent, Unassigned                     | 1              |
| F8          | UltraSPARC III Family | ASI_BLOCK_PRIMARY_LITTLE (ASI_BLK_PL)   | RW   |          | Primary address space, block load/store, little endian   | 9              |
| F9          | UltraSPARC III Family | ASI_BLOCK_SECONDARY_LITTLE (ASI_BLK_SL) | RW   |          | Secondary address space, block load/store, little endian | 9              |
| FA–FF       | —                     | —                                       |      |          | Implementation-dependent, Unassigned                     | 1              |

1. Implementation-dependent, unassigned ASIs may be used in future architectures and processors.
2. Use of these ASIs causes access checks to be performed as if the memory access instruction were issued while PSTATE.PRIV = 0 (that is, in nonprivileged mode) and directed towards the corresponding address space.
3. If the memory page being accessed is privileged, then a *data\_access\_exception* occurs.
4. Accessible as 8-bit, 16-bit, 32-bit, and 64-bit.
5. Use in LDSTUBA, SWAPA, and CAS(X)A instructions.
6. Use in load instructions only; others cause a *data\_access\_exception*.
7. Use in LDDA instruction only; others cause a *data\_access\_exception*.
8. Do not use; reserved for factory use (part number, laser programming).
9. Use in store instructions only; others cause a *data\_access\_exception*.
10. Use in LDDFA and STDFA instructions only; others cause a *data\_access\_exception*. Align to 64-byte boundary. See Section 8.4.1, “Block Load and Block Store ASIs” for more details.
11. ASI\_PRIMARY\_NOFAULT{LITTLE} and ASI\_SECONDARY\_NOFAULT{LITTLE} refer to the same address spaces as ASI\_PRIMARY{LITTLE} and ASI\_SECONDARY{LITTLE}, respectively.
12. Use in STDFA instruction only; others cause a *data\_access\_exception*. Align to 8-byte boundary. See Section 8.4.2, “Partial Store ASIs” for more details.
13. Use in LDDFA and STDFA instructions only; others cause a *data\_access\_exception*. Align to 8-bit or 16-bit boundary as required by Section 8.4.3, “Short Floating-Point Load and Store ASIs.”

## 8.7 Special Memory Access ASIs

This section describes special memory access ASIs that are not specified in SPARC V9 and are not described in other sections.



## 8.7.1 ASI 0x14 (ASI\_PHYS\_USE\_EC)

When this ASI is specified in any memory access instructions, hardware does the following:

- VA is passed to PA, and CONTEXT values are disregarded.
- Address masking is ignored (PSTATE.AM) and the VA is used. The I-MMU passes the lower 43 bits of the VA through to create a 43-bit PA.
- Memory access behaves as if its byte order is big-endian.

Even if data address translation is disabled, the access with this ASI is still a cacheable access.

## 8.7.2 ASI 0x15 (ASI\_PHYS\_BYPASS\_EC\_WITH\_EBIT)

Accesses with this ASI bypass the L2-cache and behave as if the side-effect bit (E bit) is set. When this ASI is specified in any memory access instructions, hardware does the following:

- VA is passed to PA, and CONTEXT values are disregarded.
- Address masking is ignored (PSTATE.AM) and the VA is used. The I-MMU passes the lower 43 bits of the VA through to create a 43-bit PA.
- Memory access behaves as if its byte order is big-endian.

## 8.7.3 ASI 0x1C (ASI\_PHYS\_USE\_EC\_LITTLE)

Accesses with this ASI are cacheable. This ASI is a little-endian version of ASI 14<sub>16</sub>. When this ASI is specified in any memory access instructions, hardware does the following:

- VA is passed to PA, and CONTEXT values are disregarded.
- Address masking is ignored (PSTATE.AM) and the VA is used. The I-MMU passes the lower 43 bits of the VA through to create a 43-bit PA.
- Memory access behaves as if its byte order is little-endian.

## 8.7.4 ASI 0x1D (ASI\_PHYS\_BYPASS\_EC\_WITH\_EBIT\_LITTLE)

Accesses with this ASI bypass the L2-cache and behave as if the side-effect bit (E bit) is set. This ASI is a little-endian version of ASI 15<sub>16</sub>. When this ASI is specified in any memory access instructions, hardware does the following:

- VA is passed to PA, and CONTEXT values are disregarded.
- Address masking is ignored (PSTATE.AM) and the VA is used. The I-MMU passes the lower 43 bits of the VA through to create a 43-bit PA.

- Memory access behaves as if its byte order is little-endian.

## 8.7.5 ASIs 0x24 and 0x2C (Load Quadword ASIs)

ASIs  $24_{16}$  (ASI\_NUCLEUS\_QUAD\_LDD) and  $2C_{16}$  (ASI\_NUCLEUS\_QUAD\_LDD\_LITTLE) exist for use with the LDDA instruction as Load Quadword operations.

When these ASIs are used with LDDA for Load Quadword, a *mem\_address\_not\_aligned* exception is generated if the operand address is not 16-byte aligned.

If these ASIs are used with any other Load Alternate instruction or any Store Alternate instruction, a *data\_access\_exception* is always generated and *mem\_address\_not\_aligned* is not generated.

## SECTION IV

---

---

# Memory and Cache

---

---



## Memory Models

---

The SPARC V9 architecture is a *model* that specifies the behavior observable by software on SPARC V9 systems. Therefore, access to memory can be implemented in any manner, as long as the behavior observed by software conforms to that of the models described in the following:

- Chapter 8 of *The SPARC Architecture Manual, Version 9*
- Appendix D of *The SPARC Architecture Manual, Version 9*

The SPARC V9 architecture defines three different memory models: *Total Store Order (TSO)*, *Partial Store Order (PSO)*, and *Relaxed Memory Order (RMO)*. UltraSPARC III Cu processors implements TSO, the strongest of the memory models defined by SPARC V9. By implementing TSO, software written for any memory model (TSO, PSO, and RMO) executes correctly on the UltraSPARC III Cu processor.

This chapter departs from the organization of the memory models described in *The SPARC Architecture Manual, Version 9*. It describes the characteristics of the memory models for the UltraSPARC III Cu processor in sections organized as follows:

- UltraSPARC III Cu TSO Behavior
- Memory Location Identification
- Memory Accesses and Cacheability
- Memory Synchronization
- Atomic Operations
- Non-Faulting Load
- Prefetch Instructions
- Block Loads and Stores
- I/O and Accesses with Side-Effects
- UltraSPARC III Cu Internal ASIs
- Store Compression
- Read-After-Write (RAW) Bypassing

---

## 9.1 UltraSPARC III Cu TSO Behavior

The UltraSPARC III Cu processor implements the TSO memory model. The current memory model is indicated in the `PSTATE.MM` field and is set to TSO (`PSTATE.MM = 0`).

In some cases, the UltraSPARC III Cu processor implements stronger ordering than the TSO requirements. The significant cases are listed below:

- A `MEMBAR #Lookaside` is not needed between a store and a subsequent load to the same non-cacheable address.
- Accesses with the `TTE.E` bit set, such as those that have side-effects, are all strongly ordered with respect to each other.
- An L2-cache or write cache update is delayed on a store hit until all previous stores reach global visibility. For example, a cacheable store following a non-cacheable store will not appear globally visible until the non-cacheable store has become globally visible; there is an implicit `MEMBAR #MemIssue` between them.

---

## 9.2 Memory Location Identification

A memory location is identified by an 8-bit address space identifier (ASI) and a 64-bit (virtual) address. The 8-bit ASI can be obtained from an ASI register or included in a memory access instruction. The ASI distinguishes among and provides an attribute to different 64-bit address spaces. For example, the ASI is used by the MMU and memory access hardware for control of virtual-to-physical address translations, access to implementation-dependent control and data registers, and access protection. Attempts by non-privileged software (`PSTATE.PRIV = 0`) to access restricted ASIs (`ASI<7> = 0`) cause a *privileged\_action* exception. See Chapter 8 “Address Space Identifiers” for details on ASIs.

---

## 9.3 Memory Accesses and Cacheability

Memory is logically divided into real memory (cached) and I/O memory (noncached with and without side-effects) spaces. Real memory 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. In contrast, a read from I/O space may not return the most recently written information and may result in program-visible side-effects.

## 9.3.1 Coherence Domains

Two types of memory operations are supported in the UltraSPARC III Cu processor: cacheable and non-cacheable accesses, as indicated by the page translation (`TTE.CP`, `TTE.CV`) of the MMU or by an ASI override.

SPARC V9 does not specify memory ordering between cacheable and non-cacheable accesses. The UltraSPARC III Cu processor maintains TSO ordering between memory references regardless of their cacheability.

### 9.3.1.1 Cacheable Accesses

Accesses within the coherence domain are called cacheable accesses. They have the following properties:

- Data reside in real memory locations.
- Accesses observe supported cache coherency protocol(s).
- The unit of coherence is 64 bytes.

### 9.3.1.2 Non-Cacheable and Side-Effect Accesses

Accesses outside of the coherence domain are called non-cacheable accesses. Some of these memory-mapped locations may have side-effects when accessed. They have the following properties:

- Data might not reside in real memory locations. Accesses may result in programmer-visible side-effects. An example is memory-mapped I/O control registers, such as those in a UART.
- Accesses do not observe supported cache coherency protocol(s).
- The smallest unit in each transaction is a single byte.

Non-cacheable accesses with the `TTE.E` bit set (those having side-effects) are all strongly ordered with respect to other non-cacheable accesses with the `E` bit set. In addition, store compression is disabled for these accesses. Speculative loads with the `E` bit set cause a *data\_access\_exception* trap (with `SFSR.FT = 2`, speculative load to page marked with `E` bit).

---

**Note** – `TTE.E` bit comes from the page translation of the MMU or an ASI override.

---

Non-cacheable accesses with the `TTE.E` bit cleared (non-side-effect accesses) are processor consistent and obey TSO memory ordering. In particular, processor consistency ensures that a non-cacheable load that references the same location as a previous non-cacheable store will load the data of the previous store. Store compression is supported. See Section 9.11 “Store Compression” for details.

---

**Note** – Side-effect, as indicated in TTE.E, does not imply noncacheability.

---

## 9.3.2 Global Visibility

A memory access is considered globally visible when one of the following events occurs:

- Read or write permission is granted for a cacheable transaction in scalable shared memory (SSM) mode.
- The transaction request is issued for a non-cacheable transaction in SSM mode.
- The transaction request is issued when not in SSM mode.

## 9.3.3 Memory Ordering

To ensure the correct ordering between cacheable and non-cacheable domains, explicit memory synchronization is needed in the form of MEMBAR instructions. CODE EXAMPLE 9-1 illustrates the issues involved in mixing cacheable and non-cacheable accesses.

### CODE EXAMPLE 9-1 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 for SPARC V9 compliance)
    Store F1:set flag
    While F1 is set (spin on flag)
    Load F1
2  MEMBAR #LoadLoad, #LoadStore (needed in RMO for SPARC V9
compliance)
    Load D2
}
```

Process B:

```
While (1)
{
    While F1 is cleared (spin on flag)
    Load F1
2  MEMBAR #LoadLoad, #LoadStore (needed in RMO for SPARC-V9
compliance)
    Load D1
    Store D2
}
```



**CODE EXAMPLE 9-1** Memory Ordering and MEMBAR Examples (Continued)

```
1  MEMBAR #StoreStore (needed in PSO, RMO for SPARC-V9 compliance)
   Store Fl:clear flag
}
```

---

## 9.4 Memory Synchronization

Processors always see their own normal loads and stores performed in order. TSO defines how other processors may see the ordering of the loads and stores of a particular processor. Memory synchronizations are used to force the ordering that other processors see beyond the rules of TSO.

In some cases, memory synchronizations are required for deterministic behavior, even with respect to the program's own operations. This applies to memory operations outside of normal cacheable loads and stores.

The UltraSPARC III Cu processor achieves memory synchronization through MEMBAR and FLUSH. It provides MEMBAR (STBAR in SPARC V8) and FLUSH instructions for explicit control of memory ordering in program execution. MEMBAR has several variations. All MEMBARs are implemented in one of two ways in the UltraSPARC III Cu processor:

- As a NOP
- With MEMBAR #Sync semantics

Since the processor always executes with TSO memory ordering semantics, three of the ordering MEMBARs are implemented as NOPs. TABLE 9-1 lists the MEMBAR implementations.

**TABLE 9-1** MEMBAR Semantics

| MEMBAR             | Semantics   |
|--------------------|---|
| #LoadLoad          | NOP. All loads wait for completion of all previous loads.             |
| #LoadStore         | NOP. All stores wait for completion of all previous loads.            |
| #Lookaside         | #Sync. Wait until store buffer is empty.                              |
| #StoreStore, STBAR | NOP. All stores wait for completion of all previous stores.           |
| #StoreLoad         | #Sync. All loads wait for completion of all previous stores.          |
| #MemIssue          | #Sync. Wait until all outstanding memory accesses complete.           |
| #Sync              | #Sync. Wait for all outstanding instructions and all deferred errors. |

## 9.4.1 MEMBAR #Sync

Membar #Sync forces all outstanding instructions and all deferred errors to be completed before any instructions after the MEMBAR are issued.

## 9.4.2 MEMBAR Rules

TABLE 9-2 and TABLE 9-3 summarize the cases where the programmer must insert a MEMBAR to ensure *ordering* between two memory operations on the UltraSPARC III Cu processor. Use TABLE 9-2 and TABLE 9-3 for ordering purposes only. Be sure not to confuse memory operation ordering with processor consistency or deterministic operation; MEMBARs are required for deterministic operation of certain ASI register updates.

---

**Caution** – The MEMBAR requirements for the UltraSPARC III Cu processor are weaker than the requirements of SPARC V9. To ensure code portability across systems, use the stronger of the MEMBAR requirements of SPARC V9.

---

Read the tables as follows: Read from row to column; the first memory operation in program order in a row is followed by the memory operation found in the column. Two symbols are used as table entries:

- # — No intervening operation is required because Fireplane-compliant systems automatically order R before C.
- M — MEMBAR #Sync or MEMBAR #MemIssue or MEMBAR #StoreLoad

For VA<12:5> of a column operation not matching with VA<2:5> of a row operation while a strong ordering is desired, the MEMBAR rules summarized in TABLE 9-2 reflect the UltraSPARC III Cu processor's hardware implementation.

**TABLE 9-2** MEMBAR Rules for Column VA <12:5> ≠ Row VA <12:5> While Desiring Strong Ordering

| From Row<br>Operation R: | To Column Operation C: |                        |       |                       |        |           |            |            |             |       |        |               |          |           |
|--------------------------|------------------------|------------------------|-------|-----------------------|--------|-----------|------------|------------|-------------|-------|--------|---------------|----------|-----------|
|                          | load                   | load from internal ASI | store | store to internal ASI | atomic | load_nc_e | store_nc_e | load_nc_ne | store_nc_ne | blood | bstore | bstore_commit | blood_nc | bstore_nc |
| load                     | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M     | M      | #             | M        | M         |
| load from internal ASI   | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | #        | #         |
| store                    | M                      | #                      | #     | #                     | #      | M         | #          | M          | #           | M     | M      | #             | M        | M         |

**TABLE 9-2** MEMBAR Rules for Column VA <12:5> ≠ Row VA <12:5> While Desiring Strong Ordering (Continued)

| From Row Operation R: | To Column Operation C: |                        |       |                       |        |           |            |            |             |      |        |               |         |           |
|-----------------------|------------------------|------------------------|-------|-----------------------|--------|-----------|------------|------------|-------------|------|--------|---------------|---------|-----------|
|                       | load                   | load from internal ASI | store | store to internal ASI | atomic | load_nc_e | store_nc_e | load_nc_ne | store_nc_ne | load | bstore | bstore_commit | load_nc | bstore_nc |
| store to internal ASI | #                      | M                      | #     | #                     | #      | #         | #          | #          | #           | M    | #      | #             | M       | M         |
| atomic                | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M    | M      | #             | M       | M         |
| load_nc_e             | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M    | M      | #             | M       | M         |
| store_nc_e            | M                      | #                      | #     | #                     | #      | #         | #          | M          | #           | M    | M      | #             | M       | M         |
| load_nc_ne            | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M    | M      | #             | M       | M         |
| store_nc_ne           | M                      | #                      | #     | #                     | #      | M         | #          | M          | #           | M    | M      | #             | M       | M         |
| load                  | M                      | #                      | M     | #                     | M      | M         | M          | M          | M           | M    | M      | #             | M       | M         |
| bstore                | M                      | #                      | M     | #                     | M      | M         | M          | M          | M           | M    | M      | #             | M       | M         |
| bstore_commit         | M                      | #                      | M     | #                     | M      | M         | M          | M          | M           | M    | M      | #             | M       | M         |
| load_nc               | M                      | #                      | M     | #                     | M      | M         | M          | M          | M           | M    | M      | #             | M       | M         |
| bstore_nc             | M                      | #                      | M     | #                     | M      | M         | M          | M          | M           | M    | M      | #             | M       | M         |

When VA<12:5> of a column operation matches VA<12:5> of a row operation, the MEMBAR rules summarized in TABLE 9-3 reflect the UltraSPARC III Cu processor's hardware implementation.

**TABLE 9-3** MEMBAR Rules for Column VA<12:5> = Row VA<12:5> While Desiring Strong Ordering

| From Row Operation R:  | To Column Operation C: |                        |       |                       |        |           |            |            |             |      |        |               |         |           |
|------------------------|------------------------|------------------------|-------|-----------------------|--------|-----------|------------|------------|-------------|------|--------|---------------|---------|-----------|
|                        | load                   | load from internal ASI | store | store to internal ASI | atomic | load_nc_e | store_nc_e | load_nc_ne | store_nc_ne | load | bstore | bstore_commit | load_nc | bstore_nc |
| load                   | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #    | #      | #             | #       | #         |
| load from internal ASI | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #    | #      | #             | #       | #         |
| store                  | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M    | #      | #             | #       | #         |

**TABLE 9-3** MEMBAR Rules for Column VA<12:5> = Row VA<12:5> While Desiring Strong Ordering *(Continued)*

| From Row<br>Operation R: | To Column Operation C: |                        |       |                       |        |           |            |            |             |       |        |               |          |           |
|--------------------------|------------------------|------------------------|-------|-----------------------|--------|-----------|------------|------------|-------------|-------|--------|---------------|----------|-----------|
|                          | load                   | load from internal ASI | store | store to internal ASI | atomic | load_nc_e | store_nc_e | load_nc_ne | store_nc_ne | blast | bstore | bstore_commit | blast_nc | bstore_nc |
| store to internal ASI    | #                      | M                      | #     | #                     | #      | #         | #          | #          | #           | M     | #      | #             | M        | M         |
| atomic                   | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | #        | #         |
| load_nc_e                | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | #        | #         |
| store_nc_e               | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M     | #      | #             | M        | #         |
| load_nc_ne               | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | #        | #         |
| store_nc_ne              | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M     | #      | #             | M        | #         |
| blast                    | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | #        | #         |
| bstore                   | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M     | #      | #             | #        | #         |
| bstore_commit            | M                      | #                      | M     | #                     | M      | M         | M          | M          | M           | M     | M      | #             | M        | M         |
| blast_nc                 | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | #        | #         |
| bstore_nc                | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | M        | #         |

### 9.4.3 FLUSH

FLUSH behaves like a MEMBAR with further restrictions. MEMBAR blocks execution of subsequent instructions until all memory operations and errors are resolved. FLUSH is similar with further behavior that all instruction fetch and instruction buffering operations are also blocked.

---

## 9.5 Atomic Operations

The SPARC V9 architecture provides three atomic instructions to support mutual exclusion, including:

- **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.
- If a page is marked as virtually non-cacheable but physically cacheable (`TTE.CV = 0` and `TTE.CP = 1`), allocation is done to the L2-cache and W-cache only. This includes all of the atomic-access instructions.
- **LDSTUB** — Behaves like a SWAP except that it loads a byte from memory into an integer register and atomically writes all ones (`FF16`) into the addressed byte.
- **Compare and Swap (CAS(X)A)** — 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. If they are not equal, the value in memory is still swapped with the contents of the second integer register, but is not stored. The L2-cache will still go into M-state, even if there is no store.

All of these operations are carried out atomically; in other words, no other memory operation can be applied to the addressed memory location until the entire compare-and-swap sequence is completed.

These instructions behave like both a load and store access, but the operation is carried out indivisibly. These instructions can be used only in the cacheable domain (not in non-cacheable I/O addresses).

These atomic instructions can be used with the ASIs listed in TABLE 9-4. Access with a restricted ASI in unprivileged mode (`PSTATE.PRIV = 0`) results in a *privileged\_action* trap. Atomic accesses with non-cacheable addresses cause a *data\_access\_exception* trap (with `SFSR.FT = 4`, atomic to page marked non-cacheable). Atomic accesses with unsupported ASIs cause a *data\_access\_exception* trap (with `SFSR.FT = 8`, illegal ASI value or virtual address).

**TABLE 9-4** ASIs That Support SWAP, LDSTUB, and CAS

| ASI Name                          | Access       |
|-----------------------------------|--------------|
| ASI_NUCLEUS (LITTLE)              | Restricted   |
| ASI_AS_IF_USER_PRIMARY (LITTLE)   | Restricted   |
| ASI_AS_IF_USER_SECONDARY (LITTLE) | Restricted   |
| ASI_PRIMARY (LITTLE)              | Unrestricted |
| ASI_SECONDARY (LITTLE)            | Unrestricted |
| ASI_PHYS_USE_EC (LITTLE)          | Restricted   |

---

**Note** — Atomic accesses with non-faulting ASIs are not allowed, because the latter have the load-only attribute.

---

---

## 9.6 Non-Faulting Load

A non-faulting load behaves like a normal load, with the following exceptions:

- It does not allow side-effect access. An access with the TTE.E bit set causes a *data\_access\_exception* trap (with SFSR.FT = 2, speculative load to page marked E bit).
- It can be applied to a page with the TTE.NFO (non-fault access only) bit set; other types of accesses cause a *data\_access\_exception* trap (with SFSR.FT = 10<sub>16</sub>, normal access to page marked NFO).

These loads are issued with ASI\_PRIMARY\_NO\_FAULT{ \_LITTLE } or ASI\_SECONDARY\_NO\_FAULT{ \_LITTLE }. A store with a NO\_FAULT ASI causes a *data\_access\_exception* trap (with SFSR.FT = 8, illegal RW).

When a non-faulting load encounters a TLB miss, the operating system should attempt to translate the page. If the translation results in an error, then zero is returned and the load completes silently.

Typically, optimizers use non-faulting loads to move loads across conditional control structures that guard their use. This technique potentially increases the distance between a load of data and the first use of that data, in order to hide latency. The technique allows more flexibility in code scheduling and improves performance in certain algorithms by removing address checking from the critical code path.

For example, when following a linked list, non-faulting loads allow the null pointer to be accessed safely in a speculative, read-ahead fashion; the page at virtual address 0<sub>16</sub> can safely be accessed with no penalty. The NFO bit in the MMU marks pages that are mapped for safe access by non-faulting loads, but that can still cause a trap by other, normal accesses.

Thus, programmers can trap on wild pointer references — many programmers count on an exception being generated when accessing address 0<sub>16</sub> to debug code — while benefiting from the acceleration of non-faulting access in debugged library routines.

---

## 9.7 Prefetch Instructions

The UltraSPARC III Cu processor implements all SPARC V9 prefetch instructions except for prefetch page. All prefetches check the L2-cache before issuing a system request for the requested data. Prefetch instructions are a performance feature. Prefetch instructions do not change the underlying memory model and do not have any effect from an architectural standpoint.

TABLE 9-5 describes different types of software prefetch instructions.

**TABLE 9-5** Types of Software Prefetch Instructions

| fcn Value (hex) | Instruction Type                 | Prefetch (64 bytes of data) into:                  | Instruction Strength UltraSPARC III Cu | Request Exclusive Ownership |
|-----------------|----------------------------------|--|--|-----------------------------|
| 00              | Prefetch read many               | P-cache and L2-cache                               | Weak                                   | No                          |
| 01              | Prefetch read once               | P-cache only                                       | Weak                                   | No                          |
| 02              | Prefetch write many              | L2-cache only                                      | Weak                                   | Yes                         |
| 03              | Prefetch write once <sup>1</sup> | L2-cache only                                      | Weak                                   | No                          |
| 04              | <i>Reserved</i>                  | Undefined  |  |                             |
| 05 – 0F         | <i>Reserved</i>                  | Undefined  |  |                             |
| 10              | Prefetch invalidate              | Invalidates a P-cache line, no data is prefetched. |  | N/A                         |
| 11 – 13         | <i>Reserved</i>                  | Undefined  |  |                             |
| 14              | Same as fcn = 00                 |  | Weak <sup>2</sup>                      | No                          |
| 15              | Same as fcn = 01                 |  | Weak <sup>2</sup>                      | No                          |
| 16              | Same as fcn = 02                 |  | Weak <sup>2</sup>                      | Yes                         |
| 17              | Same as fcn = 03                 |  | Weak <sup>2</sup>                      | No                          |
| 18 – 1F         | <i>Reserved</i>                  | Undefined  |  |                             |

1. Although the name is “prefetch write once,” the actual use is prefetch to L2-cache for a future read.

2. These weak instructions may be implemented as strong in future implementations.

## 9.8 Block Loads and Stores

Block load and store instructions work like normal floating-point load and store instructions, except that the data size (granularity) is 64 bytes per transfer.

Block loads and stores do not obey TSO. They do not even obey the processor’s consistency rules without the correct use of MEMBARs. Section A.4 “Block Load and Block Store (VIS I)” on page A-460 discusses block loads and stores in detail.

---

## 9.9 I/O and Accesses with Side-Effects

I/O locations might not behave with memory semantics. Loads and stores could have side-effects; for example, a read access could clear a register or pop an entry off a FIFO. A write access could set a register address port so that the next access to that address will read or write a particular internal register. Such devices are considered order sensitive. Also, such devices may only allow accesses of a fixed size, so store merging of adjacent stores or stores within a 16-byte region would cause an error.

The UltraSPARC III Cu MMU includes an attribute bit in each page translation, TTE . E, which when set signifies that this page has side-effects. Accesses other than block loads or stores to pages that have this bit set exhibit the following behavior:

- Non-cacheable accesses are strongly ordered with respect to each other.
- Non-cacheable loads with the E bit set will not be issued to the system until all previous control transfers are resolved.
- Non-cacheable store compression is disabled for E bit accesses.
- Exactly those E-bit accesses implied by the program are made in program order.
- Non-faulting loads are not allowed and cause a *data\_access\_exception* (with SFSR . FT = 2, speculative load to page marked E bit).
- A MEMBAR may be needed between side-effect and non-side-effect accesses while in PSO and RMO modes, for portability across SPARC V9 processors, as well as in some cases of TSO.

### 9.9.1 Instruction Prefetch to Side-Effect Locations

The processor does instruction prefetching and follows branches that it predicts are taken. Addresses mapped by the I-MMU can be accessed even though they are not actually executed by the program. Normally, locations with side-effects or that generate timeouts or bus errors are not mapped by the I-MMU, so prefetching will not cause problems.

When running with the I-MMU disabled, software must avoid placing data in the path of a control transfer instruction target or sequentially following a trap or conditional branch instruction. Data can be placed sequentially following the delay slot of a BA, BPA(*p* = 1), CALL, or JMP.L instruction. Instructions should not be placed closer than 256 bytes to locations with side-effects.



## 9.9.2 Instruction Prefetch Exiting Red State

Exiting `RED_state` by writing zero to `PSTATE.RED` in the delay slot of a `JMPL` instruction is not recommended. A non-cacheable instruction prefetch may be made to the `JMPL` target, which may be in a cacheable memory area. This situation can result in a bus error on some systems and can cause an instruction access error trap. Programmers can mask the trap by setting the `NCEEN` bit in the L2-cache Error Enable Register to zero, but doing so will mask all uncorrectable error checking. Exiting `RED_state` with `DONE`, `RETRY`, or with the destination of the `JMPL` non-cacheable will avoid the problem.

---

## 9.10 UltraSPARC III Cu Internal ASIs

ASIs in the ranges  $30_{16}-6F_{16}$  and  $72_{16}-7F_{16}$  are used for accessing internal states. Stores to these ASIs do not follow the normal memory-model ordering rules. Correct operation can be assured by adhering to the following requirements:

- A `MEMBAR #Sync` is needed after a store to an internal ASI other than MMU ASIs before the point that side-effects must be visible. This `MEMBAR` must precede the next load or non-internal store. To avoid data corruption, the `MEMBAR` must also occur before the delay slot of a delayed control transfer instruction of any type.
- Alternatively, a `MEMBAR #Sync` could be inserted at the beginning of any vulnerable trap handler. “Vulnerable” trap handlers are those which contain one or more `LDXAs` from any internal ASI (ASIs  $0x30-0x6F$ ,  $0x72-0x77$ , and  $0x7A-0x7F$ ). However, this may cause unacceptable performance reduction in some trap handlers, so this is not the preferred alternative.
- A `FLUSH`, `DONE`, or `RETRY` is needed after a store to an internal I-MMU ASI (ASI  $50_{16}-52_{16}$ ,  $54_{16}-5F_{16}$ ), an I-cache ASI ( $66_{16}-6F_{16}$ ), or the `IC` bit in the DCU Control Register, prior to the point that side-effects must be visible. A store to D-MMU registers other than the context ASIs can use a `MEMBAR #Sync`. To avoid data corruption, the `MEMBAR` must also occur before the delay slot of a delayed control transfer instruction of any type.
- If the store is to an I-MMU state register (ASI =  $50_{16}$ , virtual address =  $18_{16}$ ), then the `FLUSH`, `DONE`, or `RETRY` must *immediately* follow the store. Furthermore, one of the following must be true, to prevent an intervening I-TLB miss from causing stale data to be stored:
  - The code must be locked down in the I-TLB, or
  - The store and the subsequent `FLUSH`, `DONE`, or `RETRY` should be kept on the same 8 KB page of instruction memory.

---

## 9.11 Store Compression

Consecutive non-side-effect, non-cacheable stores can be combined into aligned 16-byte entries in the store buffer to improve store bandwidth. Cacheable stores will naturally coalesce in the write cache rather than be compressed in the store buffer. Non-cacheable stores can be compressed only with adjacent non-cacheable stores. To maintain strong ordering for I/O accesses, stores with the side-effect attribute (E bit set) cannot be combined with any other stores.

A 16-byte non-cacheable merge buffer is used to coalesce adjacent non-cacheable stores. Non-cacheable stores will continue to coalesce into the 16-byte buffer until one of the following conditions occurs:

- The data is pulled from the non-cacheable merge buffer by the target device.
- The store would overwrite a previously written entry (a valid bit is kept for each of the 16 bytes).

---

**Caution** – This behavior is unique to the UltraSPARC III Cu processor and differs from previous UltraSPARC implementations.

---

- The store is not within the current address range of the merge buffer (within the 16-byte aligned merge region).
- The store is a cacheable store.
- The store is to a side-effect page.
- MEMBAR #Sync

---

## 9.12 Read-After-Write (RAW) Bypassing

Load data can be bypassed from previous stores before they become globally visible (data for load from the store queue). This is specifically allowed by the TSO memory model. Data for all types of loads cannot be bypassed from all types of stores.

All types of load instructions can get data from the store queue, except the following load instructions:

- Signed loads (ldsb, ldsh, ldsw)
- Atomics
- Load double to integer register file (ldd)
- Quad loads to integer register file

- Load from FSR register
- Block loads
- Short floating-point loads
- Loads from internal ASIs

All types of store instructions can give data to a load, except the following store instructions:

- Floating-point partial stores
- Store double from integer register file (`std`)
- Store part of atomic
- Short FP stores
- Stores to pages with side-effect bit set
- Stores to non-cacheable pages

## 9.12.1 RAW Bypassing Algorithm

The algorithm used in the UltraSPARC III Cu processor for RAW bypassing is as follows:

```

if ( (Load/store access the same physical address)      and
     (Load/store endianness is the same)                and
     (Load/store size is the same)                      and
     (Load data can get its data from store queue)      and
     (Store data in store can give its data to a load)  and
     (Load hits in either D-cache or P-cache)
)
then
    Load will get its data from store queue
else
    Load will get its data from the memory system
endif

```

## 9.12.2 RAW Detection Algorithm

When data for a load cannot be bypassed from previous stores before they become globally visible (store data is not yet retired from the store queue), the load is recirculated after the RAW hazard is removed. The following conditions can cause this recirculation:

- Load data can be bypassed from more than one store in the store queue.
- The load's VA<12:0> overlaps a store's VA<12:0> and store data cannot be bypassed from the store queue.
- The load's VA<12:5> matches a store's VA<12:5> and the load misses the D-cache.
- Load is from side-effect page (page attribute E = 1) when the store queue contains one or more stores to side-effect pages.

# Caches and Cache Coherency

---

This chapter describes the use of caches, and contains these sections:

- Cache Organization
- Cache Flushing
- Bypassing the D-Cache
- Controlling P-cache
- Coherence Tables

---

## 10.1 Cache Organization

In this section we describe two cache organizations: virtual indexed, physical tagged caches and physical indexed, physical tagged caches.

### 10.1.1 Virtual Indexed, Physical Tagged Caches (VIPT)

The Data Cache (D-cache) is virtual-indexed, physical-tagged (VIPT). Virtual addresses are used to index into the cache tag and data arrays while accessing the D-MMU, (i.e. D-TLBs). The resulting tag is compared against the translated physical address to determine cache hit.

A side-effect inherent in a virtual-indexed cache is *address aliasing*. This issue is addressed in Section 10.2.1 “Address Aliasing Flushing” on page 10-230.

#### 10.1.1.1 Data Cache (D-Cache)

The Data Cache (D-cache) is a write-through, non-allocating on write miss, 64 KB, four-way associative cache with a 32-byte line.

Data accesses bypass the D-cache when the D-cache enable bit in the DCU Control Register is clear. If the DM bit in the DCU Control Register is clear, the cacheability is determined by the CP and CV bits. If the access is mapped by the D-MMU as non-virtual-cacheable, then load misses will not allocate in the D-cache.

---

**Note** – A non-virtual-cacheable access may access data in the D-cache from an earlier cacheable access to the same physical block, unless the D-cache is disabled. Software must flush the D-cache when changing a physical page from cacheable to non-cacheable (see 10.2, “Cache Flushing”).

---

## 10.1.2 Physical Indexed, Physical Tagged Caches (PIPT)

### 10.1.2.1 Instruction Cache (I-Cache)

The Instruction Cache (I-cache) is a 32KB pseudo four-way set-associative, write-invalidate cache with 32-byte lines. Instruction fetches bypass the I-cache when any of the following occur:

- I-cache enable or I-MMU enable bits in the DCU Control Register are clear
- CP bit in the DCU Control Register is clear
- Processor is in RED mode, or
- Fetch is mapped by the I-MMU as not physically cacheable.

The I-cache snoops stores from other processors or DMA transfers, as well as stores in the same processor and block commit store.

The FLUSH instruction is not required to maintain coherency. Stores and block store commits invalidate the I-cache, but do not flush instructions that have already been prefetched into the pipeline. A FLUSH instruction should be used to flush the pipeline in the case of self modifying code.

If a program changes I-cache mode to I-cache-ON from I-cache-OFF, then the next instruction fetch always causes an I-cache miss even if it is supposed to hit. This rule applies even when the DONE instruction turns on the I-cache by changing its status from RED\_state to normal mode (see CODE EXAMPLE 10-1).

**CODE EXAMPLE 10-1** I-Cache Mode Example

```
(in RED_state)
setx 0x37e0000000007, %g1, %g2
stxa %g2, [%g0]0x45           // Turn on I-cache when processor
                               // returns normal mode.
done                          // Escape from RED_state.
```

(back to normal mode)

```
nop // 1st instruction; this always causes an I-cache miss.
```

### 10.1.2.2 Prefetch Cache (P-Cache)

The Prefetch Cache (P-cache) is a write-invalidate, 2KB, 4-way associative cache with a 64-byte line and two 32-byte sub-blocks. It is physically indexed and physically tagged and never contains modified data. The P-cache is filled by the following:

- Floating Point (FP) Load Miss of D-cache and P-cache
- Hardware or Software Prefetch

The `PREFETCH fcn=16` instruction can be used to invalidate, or flush a P-cache entry.

The cache line size is 64 bytes with 32-byte sub-blocks. Prefetches always generate 64-byte fills. An FP Load Miss generates 32-byte fills. The P-cache is globally invalidated on context changes and MMU updates. Individual lines are invalidated on store hits. There is a diagnostic interface (`ASI_PCACHE_TAG`) that allows privileged software to invalidate entries in P-cache.

The P-cache is globally invalidated if any of the following conditions occur:

- The context registers are written.
- There is a demap operation in the DMU.
- The DMU is turned on or off.

Individual lines are invalidated on any of the following conditions:

- A store hits
- An external snoop hit
- Use of software prefetch invalidate function (`PREFETCH` with `fcn = 16`)

The P-cache is used for software prefetch instructions as well as an autonomous hardware prefetch from the L2-cache. This cache never needs to be flushed (not even for address aliases).

---

**Note** – From a load’s perspective, the P-cache is a virtually indexed, virtually tagged (VIVT) cache.

---

### 10.1.2.3 Second Level and Write Caches (L2-Cache, W-Cache)

The level-2 caches — the L2-cache and the W-cache — are physical indexed, physical tagged (PIPT). These caches have no references to virtual address and context information. The operating system needs no knowledge of such caches after initialization, except for stable storage management and error handling.

Instruction fetches bypass the L2-cache in the following cases:

- The I-MMU is disabled *and* the CP bit in the DCU Control Register is not set.
- The processor is in RED\_state.
- The access is mapped by the I-MMU as not physically cacheable.

Data accesses bypass the L2-cache if the D-MMU enable bit in the DCU Control Register is clear or if the access is mapped by the D-MMU as not physically cacheable (unless ASI\_PHYS\_USE\_EC is used).

The system must provide a non-cacheable scratch memory for use by the booting code until the MMUs are enabled.

The L2-cache is a unified, write-back, allocating, one- or two-way set-associative cache. Its size ranges from 1 MB to 8 MB. Its line size varies from 64 bytes to 512 bytes with 64-byte sub-blocks. Fills are done at a sub-block/64-byte granularity. Allocations and evictions are done at a line granularity. See TABLE 10-1 for description on sub-block and line sizes. L2-cache uses a pseudo-random replacement policy. L2-cache cannot be disabled by software.

**TABLE 10-1** External Cache Organizations

| External Cache Size | Line Size | 64-Byte Sublines per line |
|---------------------|-----------|---------------------------|
| 1 MB                | 64 bytes  | 1                         |
| 4 MB                | 256 bytes | 4                         |
| 8 MB                | 512 bytes | 8                         |

Block loads and block stores, which load or store a 64-byte block of data between memory and the Floating Point Register file, do not allocate into the L2-cache, to avoid pollution. Block store commits invalidate the L2-cache. Prefetch Read Once instructions, which load a 64-byte block of data into the P-cache, do not allocate into the L2-cache. All other level-1 cache fills also allocate in L2-cache.

The W-cache is a 2-KB, 4-way associative cache, with 64 bytes per line and 32-byte sub-blocks. The W-cache is included in the L2-cache, and flushing the L2-cache ensures that the W-cache has also been flushed. The W-cache only contains those bytes of data written by stores and a per byte dirty (valid) bit is maintained. All cacheable stores allocate a 64-byte line in the W-cache. W-cache may contain newer data than in L2-cache. All data accesses to L2-cache must access the W-cache in parallel and merge at a byte granularity if the W-cache has newer data. On eviction, the W-cache writes back to the L2-cache.

#### 10.1.2.4 L2-cache Replacement Policy

CODE EXAMPLE 10-2 reflects the cache replacement algorithm when all 4 ways of the L2-cache are active.

**CODE EXAMPLE 10-2** L2-cache Replacement Policy



```

module lfsr (rand_out, event_in, reset, clk);

output [3:0] rand_out;
input event_in;
input reset;
input clk;

wire [4:0] lfsr_reg;

dfffe #(5) ff_lfsr (lfsr_reg, lfsr_in, ~reset, event_in, clk);

// 01010 is the non-reachable state for this implementation.
wire [4:0] lfsr_in = {~lfsr_reg[0],
                    lfsr_reg[0] ^ lfsr_reg[4],
                    lfsr_reg[3],
                    lfsr_reg[0] ^ lfsr_reg[2],
                    lfsr_reg[0] ^ lfsr_reg[1]};

// update on reads that miss the L2-cache
assign event_in = ec_lt_cs_r_d1 & ~ec_lt_we_r_d1 &
~lt_ec_hit_miss_d1;

dfffire #(5) f_lfsr (lfsr_reg, lfsr_in, reset, event_in, clk);

assign rand_out = { lfsr_reg[1] & lfsr_reg[0],
                  lfsr_reg[1] & ~lfsr_reg[0],
                  ~lfsr_reg[1] & lfsr_reg[0],
                  ~lfsr_reg[1] & ~lfsr_reg[0]};

endmodule

```

---

## 10.2 Cache Flushing

Data in the write-invalidate or write-through caches can be flushed by invalidating the entry in the cache. Modified data in the L2-cache and W-cache must be written back to memory when flushed.

Cache flushing is required in the following cases:

- A D-cache flush is needed when a physical page is changed from (virtually) cacheable to (virtually) non-cacheable, or an illegal address aliasing is created (see 10.2.1, “Address Aliasing Flushing”). This is done using ASI 0x42, ASI\_DCACHE\_INVALIDATE, which specifies a physical address to flush, like for a system bus snoop.
- L2-cache flush is needed for stable storage. This is done with either a displacement flush (using ASI\_ECACHE\_TAG) or a store with ASI\_BLK\_COMMIT. Flushing the L2-cache will flush the corresponding blocks from W-cache. (See 10.2.2, “Committing Block Store Flushing”).
- L2-cache, D-cache, P-cache, and I-cache flushes may be required when an ECC error occurs on a read from the memory or the L2-cache. When an ECC error occurs invalid data may be written into one of the caches, the cache lines must be flushed to prevent further corruption of data.

---

**Note** – When flushing a single 64-byte line, with a given PA, there are sixteen locations that must be flushed in the D-cache. This is because it has 32-byte lines (2 places), one VA index bit (2 places), and the PA can simultaneously exist in all 4 ways of a set (4 places).

---

## 10.2.1 Address Aliasing Flushing

A side-effect inherent in a virtual indexed cache is *illegal address aliasing*. Aliasing occurs when multiple virtual addresses map to the same physical address.

---

**Caution** – Since the D-cache is indexed with the virtual address bits (specifically bit 13 of the address) and is larger than the minimum page size, it is possible for the different aliased virtual addresses to end up in different cache blocks. Such aliases are illegal because updates to one cache block will not be reflected in aliased cache blocks. (There are corner cases where the same cache block can end up in different ways, within the same set (index). The hardware will update all ways within a set that have the line.)

---

Normally, software avoids illegal aliasing by forcing aliases to have the same address bits (*virtual color*) up to an *alias boundary*. The minimum alias boundary is 16KB. This size may increase in future designs.

When the alias boundary is violated, software must flush the D-cache if the page was virtual cacheable. In this case, only one mapping of the physical page can be allowed in the D-MMU at a time.

Alternatively, software can turn off virtual caching of illegally aliased pages. This allows multiple mapping of the alias to be in the D-MMU and avoids flushing the D-cache each time a different mapping is referenced.

---

**Note** – A change in virtual color when allocating a free page does not require a D-cache flush, since the D-cache is write-through.

---

## 10.2.2 Committing Block Store Flushing

Block store commit does a 64-byte store from Floating Point Registers directly to the memory and invalidates any caches in the system. Block store commit exists outside *Total Store Order (TSO)*.

For example, stable storage must be implemented by software cache flush. Examples of stable storage are battery-backed memory and a transaction log. Data which is present and modified in the L2-cache or the Write Cache must be written back to the stable storage.

Two ASIs (`ASI_BLK_COMMIT_PRIMARY` and `ASI_BLK_COMMIT_SECONDARY`) perform these writebacks efficiently when software can ensure exclusive write access to the block being flushed. These ASIs writeback the data from the Floating Point Registers to memory and invalidate the entry in the cache. The data in the Floating Point Registers must first be loaded by a block load instruction. A `MEMBAR #Sync` instruction can be used to ensure that the flush is complete.

## 10.2.3 Displacement Flushing

Cache flushing can also be accomplished by a displacement flush. This procedure reads a range of addresses that map to the corresponding cache line being flushed, forcing out modified entries in the local cache. 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.

---

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

---

---

## 10.3 Bypassing the D-Cache

D-cache can return stale data if `CP==1, CV==0` is used to bypass the cache, after use of `CP==1` and `CV==1`, for loads and stores to a particular address.

D-cache should be flushed, after mixing use of any CP/CV settings for a physical address, including cacheable (DRAM) and non-cacheable (I/O) physical addresses.

The term "virtually non-cacheable" refers to the "non-D-cacheable" CP==1, CV==0 case, as opposed to the more common use of "non-cacheable" to describe I/O or graphics related physical addresses

- CP == 1, CV == 1: Cacheable, Virtually-cacheable
- CP == 1, CV == 0: Cacheable, Virtually-non-cacheable (ASI\_PHYS\_USE\_EC has this effect)
- CP == 0, CV == 1: P-cache only
- CP == 0, CV == 0: Non-cacheable

Only two indexes in the D-cache need be flushed for each 32-byte aligned physical address:

- {VA[13] == 0, PA[12:5]}, and
- {VA[13] == 1, PA[12:5]}

---

**Note** – The behavior of a load or store instruction using ASI 0x14 (ASI\_PHYS\_USE\_EC) that hits on the D-cache is not defined.

---

---

## 10.4 Controlling P-cache

This section clarifies the use of DCR.PE, DCR.HPE, and DCR.SPE bits.

---

**Note** – Block loads do not cause installs into the P-cache. They are also not allowed to hit on the P-cache and, therefore, never trigger hardware prefetch.

Non-cacheable address space never installs in P-cache or L2-cache.

---

- PE = 0, HPE = x, SPE = x
  - No Hardware prefetches, no Software prefetches.
  - No FP load miss data (32 bytes) installed from L2-cache to P-cache.
  - All P-cache references treated as miss.
- PE = 1, HPE = 0, SPE = 0
  - No Hardware prefetch can work.
  - None of the software prefetches work.
  - No FP load miss data (32 bytes) installed in P-cache.
  - FP loads will be checked for P-cache hit/miss.

- PE = 1, HPE = 1, SPE = 0
  - Hardware prefetches can work
  - None of the software prefetches can work.
  - FP load miss data (32 bytes) can be installed in P-cache.
  - P-cache references (FP load) will be checked for hit/miss
- PE=1, HPE=0, SPE=1
  - No Hardware prefetches.
  - All Software prefetches can work.
  - No FP load miss data (32 bytes) installed in P-cache.
  - FP loads will be checked for P-cache hit/miss.
- PE = 1, HPE = 1, SPE = 1
  - Hardware prefetches can work.
  - All Software prefetches can work.
  - FP load miss data (32 bytes) will be installed in P-cache.
  - FP loads will checked for P-cache hit/miss.

---

**Note** – The configuration PE = 1, HPE = 1, and SPE=1 is not supported in the UltraSPARC III Cu processor. The use of this mode may cause data in the P-cache to become incoherent with the system.

---

## 10.5 Coherence Tables

The set of tables in this section describes the cache coherence protocol that governs the behavior of the processor on the interface bus.

### 10.5.1 Processor State Transition and the Generated Transaction

Tables in this section summarize the following:

- Hit/Miss, State Change, and Transaction Generated for Processor Action (TABLE 10-2)
- Combined Tag/MTag States (TABLE 10-3)
- Derivation of DTags, CTags, and MTags from Combined Tags (TABLE 10-4)

**TABLE 10-2** Hit/Miss, State Change, and Transaction Generated for Processor Action

| Combined State | MODE              | Processor action |                  |                 |             |                    |                  |
|----------------|-------------------|------------------|------------------|-----------------|-------------|--------------------|------------------|
|                |                   | Load             | Store/ Swap      | Block Load      | Block Store | Block Store Commit | Write Prefetch   |
| I              | ~SSM              | Miss: RTS        | Miss: RTO        | Miss : RS       | Miss: WS    | Miss: WS           | Miss: RTO        |
|                | SSM & LPA         | Miss: RTS        | Miss: RTO        | Miss: RS        | Miss: R_WS  | Miss: R_WS         | Miss: RTO        |
|                | SSM & LPA & retry | Mtag miss: R_RTS | Mtag miss: R_RTO | Mtag miss: R_RS | Invalid     | Invalid            | MTag Miss: R_RTO |
|                | SSM & ~LPA        | Miss: R_RTS      | Miss: R_RTO      | Miss: R_RS      | Miss: R_WS  | Miss: R_WS         | Miss: R_RTO      |
| E              | ~SSM              | Hit              | Hit: E→M         | Hit             | Hit: E→M    | Hiss: WS           | Hit              |
|                | SSM & LPA         | Hit              | Hit: E→M         | Hit             | Hit: E→M    | Miss: R_WS         | Hit              |
|                | SSM & LPA & retry | Invalid          |                  |                 |             |                    |                  |
|                | SSM & ~LPA        | Hit              | Hit: E→M         | Hit             | Hit: E→M    | Miss: R_WS         | Hit              |
| S              | ~SSM              | Hit              | Miss: RTO        | Hit             | Miss: WS    | Miss: WS           | Hit              |
|                | SSM & LPA         | Hit              | MTag miss: RTO   | Hit             | Miss: R_WS  | Miss: R_WS         | Hit              |
|                | SSM & LPA & retry | Invalid          | MTag Miss: R_RTO | Invalid         | Invalid     | Invalid            | Invalid          |
|                | SSM & ~LPA        | Hit              | MTag Miss: R_RTO | Hit             | Miss: R_WS  | Miss: R_WS         | Hit              |

**TABLE 10-2** Hit/Miss, State Change, and Transaction Generated for Processor Action

| Combined State                    | MODE                    | Processor action |                        |               |                |                          |                   |
|-----------------------------------|-------------------------|------------------|------------------------|---------------|----------------|--------------------------|-------------------|
|                                   |                         | Load             | Store/<br>Swap         | Block<br>Load | Block<br>Store | Block<br>Store<br>Commit | Write<br>Prefetch |
| O                                 | ~SSM                    | Hit              | Miss:<br>RTO           | Hit           | Miss:<br>WS    | Miss:<br>WS              | Hit               |
|                                   | SSM &<br>LPA            | Hit              | MTag<br>Miss:<br>RTO   | Hit           | Miss:<br>R_WS  | Miss:<br>R_WS            | Hit               |
|                                   | SSM &<br>LPA &<br>retry | Invalid          | MTag<br>Miss:<br>R_RTO | Invalid       | Invalid        | Invalid                  | Invalid           |
|                                   | SSM &<br>~LPA           | Hit              | MTag<br>Miss:<br>R_RTO | Hit           | Miss:<br>R_WS  | Miss:<br>R_WS            | Hit               |
| Os (Legal<br>only in SSM<br>mode) | ~SSM                    | Invalid          |                        |               |                |                          |                   |
|                                   | SSM &<br>LPA            | Hit              | MTag<br>Miss:<br>R_RTO | Hit           | Miss:<br>R_WS  | Miss:<br>R_WS            | Hit               |
|                                   | SSM &<br>LPA &<br>retry | Invalid          | MTag<br>Miss:<br>R_RTO | Invalid       | Invalid        | Invalid                  | Invalid           |
|                                   | SSM &<br>~LPA           | Hit              | MTag<br>Miss:<br>R_RTO | Hit           | Miss:<br>R_WS  | Miss:<br>R_WS            | Hit               |
| M                                 | ~SSM                    | Hit              | Hit                    | Hit           | Hit            | Miss:<br>WS              | Hit               |
|                                   | SSM &<br>LPA            | Hit              | Hit                    | Hit           | Hit            | Miss:<br>R_WS            | Hit               |
|                                   | SSM &<br>LPA &<br>retry | Invalid          |                        |               |                |                          |                   |
|                                   | SSM &<br>~LPA           | Hit              | Hit                    | Hit           | Hit            | Miss:<br>R_WS            | Hit               |

**TABLE 10-3** Combined Tag/MTag States

| <b>MTag State: CTag State</b> | <b>gI</b> | <b>gS</b> | <b>gM</b> |
|-------------------------------|-----------|-----------|-----------|
| cM                            | I         | Os        | M         |
| cO                            | I         | Os        | O         |
| cE                            | I         | S         | E         |
| cS                            | I         | S         | S         |
| cI                            | I         | I         | I         |

**TABLE 10-4** Deriving DTags, CTags, and MTags from Combined Tags

| <b>Combined Tags (CCTags)</b> | <b>DTag</b> | <b>CTag</b> | <b>MTag</b> |
|-------------------------------|-------------|-------------|-------------|
| I                             | dI          | cI          | gI          |
| E                             | dS          | cE          | gM          |
| S                             | dS          | cS          | gS          |
| O                             | dO          | cO          | gM          |
| Os                            | dO          | cO          | gS          |
| M                             | dO          | cM          | gM          |

## 10.5.2 Snoop Output and Input

TABLE 10-5 summarizes snoop output and DTag transition; TABLE 10-6 summarizes snoop input and CIQ operation queueing.

**TABLE 10-5** Snoop Output and DTag Transition (*1 of 4*)

| <b>Snooped Request</b>     | <b>DTag State</b> | <b>Shared Output</b> | <b>Owned Output</b> | <b>Error Output</b> | <b>Next DTag State</b> | <b>Action for Snoop Pipeline</b> |
|----------------------------|-------------------|----------------------|---------------------|---------------------|------------------------|----------------------------------|
| own RTS (for data)         | dI                | 0                    | 0                   | 0                   | dT                     | Own RTS wait data                |
|                            | dS                | 1                    | 0                   | 1                   | dS                     | Error                            |
|                            | dO                | 1                    | 0                   | 1                   | dO                     | Error                            |
|                            | dT                | 1                    | 0                   | 1                   | dT                     | Error                            |
| own RTS (for instructions) | dI                | 0                    | 0                   | 0                   | dS                     | Own RTS instruction wait data    |
|                            | dS                | 1                    | 0                   | 1                   | dS                     | Error                            |
|                            | dO                | 1                    | 0                   | 1                   | dO                     | Error                            |
|                            | dT                | 1                    | 0                   | 1                   | dT                     | Error                            |



**TABLE 10-5** Snoop Output and DTag Transition (2 of 4)

| Snooped Request | DTag State | Shared Output | Owned Output | Error Output | Next DTag State | Action for Snoop Pipeline       |
|-----------------|------------|---------------|--------------|--------------|-----------------|---------------------------------|
| foreign RTS     | dI         | 0             | 0            |              | dI              | None                            |
|                 | dS         | 1             | 0            |              | dS              | None                            |
|                 | dO         | 1             | 1            |              | dO              | Foreign RTS copyback            |
|                 | dT         | 1             | 0            |              | dS              | None                            |
| own RTO         | dI         | 0             | 0            |              | dO              | Own RTO wait data               |
|                 | dS & ~SSM  | 1             | 1            |              | dO              | Own RTO no data                 |
|                 | dS & SSM   | 0             | 0            |              | dO              | Own RTO wait data               |
|                 | dO         | 1             | 1            |              | dO              | Own RTO no data                 |
|                 | dT         | 1             | 1            | 1            | dO              | Error                           |
| foreign RTO     | dI         | 0             | 0            |              | dI              | None                            |
|                 | dS         | 0             | 0            |              | dI              | Foreign RTO invalidate          |
|                 | dO         | 0             | 1            |              | dI              | Foreign RTO copyback-invalidate |
|                 | dT         | 0             | 0            |              | dI              | Foreign RTO invalidate          |
| own RS          | dI         | 0             | 0            |              | dI              | Own RS wait data                |
|                 | dS         | 0             | 0            | 1            | dS              | Error                           |
|                 | dO         | 0             | 0            | 1            | dO              | Error                           |
|                 | dT         | 0             | 0            | 1            | dT              | Error                           |
| foreign RS      | dI         | 0             | 0            |              | dI              | None                            |
|                 | dS         | 0             | 0            |              | dS              | None                            |
|                 | dO         | 0             | 1            |              | dO              | Foreign RS copyback-discard     |
|                 | dT         | 0             | 0            |              | dT              | None                            |
| own WB          | dI         | 0             | 1            |              | dI              | Own WB (cancel)                 |
|                 | dS         | 0             | 1            |              | dI              | Own WB (cancel)                 |
|                 | dO         | 0             | 0            |              | dI              | Own WB                          |
|                 | dT         | 0             | 1            | 1            | dI              | Error                           |
| foreign WB      | dI         | 0             | 0            | 0            | dI              | None                            |
|                 | dS         | 0             | 0            | 0            | dS              | None                            |
|                 | dO         | 0             | 0            | 0            | dO              | None                            |
|                 | dT         | 0             | 0            | 0            | dT              | None                            |

**TABLE 10-5** Snoop Output and DTag Transition (3 of 4)

| Snooped Request                 | DTag State | Shared Output | Owned Output | Error Output | Next DTag State | Action for Snoop Pipeline                           |
|---------------------------------|------------|---------------|--------------|--------------|-----------------|---|
| foreign RTSM                    | dI         | 0             | 0            |              | dI              | None  |
|                                 | dS         | 0             | 0            |              | dS              | Foreign RTSM  |
|                                 | dO         | 0             | 1            |              | dS              | fRTSM copyback (if cacheline is not in the W-cache) |
|                                 | dO         | 0             | 1            |              | dI              | fRTSM copyback (if cacheline is in the W-cache)     |
|                                 | dT         | 0             | 0            |              | dS              | Foreign RTSM  |
| own RTSR (issued by SSM device) | dI         | 0             | 0            |              | dO              | Own RTSR wait data                                  |
|                                 | dS         | 0             | 0            | 0            | dO              | Own RTSR wait data                                  |
|                                 | dO         | 0             | 0            | 1            | dO              | Own RTSR wait data, Error                           |
|                                 | dT         | 0             | 0            | 1            | dT              | Own RTSR wait data, Error                           |
| foreign RTSR                    | dI         | 0             | 0            |              | dI              | None  |
|                                 | dS         | 1             | 0            |              | dS              | None  |
|                                 | dO         | 1             | 1            |              | dS              | Foreign RTSR  |
|                                 | dT         | 1             | 0            |              | dS              | None  |
| own RTOR (issued by SSM device) | dI         | 0             | 0            |              | dO              | Own RTOR wait data                                  |
|                                 | dS         | 0             | 0            |              | dO              | Own RTOR wait data                                  |
|                                 | dO         | 0             | 0            |              | dO              | Own RTOR wait data                                  |
|                                 | dT         | 0             | 0            | 1            | dO              | Error   |
| foreign RTOR                    | dI         | 0             | 0            |              | dI              | None  |
|                                 | dS         | 0             | 0            |              | dI              | Foreign RTOR invalidate                             |
|                                 | dO         | 0             | 0            |              | dI              | Foreign RTOR invalidate                             |
|                                 | dT         | 0             | 0            |              | dI              | Foreign RTOR invalidate                             |
| own RSR                         | dI         | 0             | 0            |              | dI              | Own RSR wait data                                   |
|                                 | dS         | 0             | 0            | 1            | dS              | Error   |
|                                 | dO         | 0             | 0            | 1            | dO              | Error   |
|                                 | dT         | 0             | 0            | 1            | dT              | Error   |
| foreign RSR                     | dI         | 0             | 0            |              | dI              | None  |
|                                 | dS         | 0             | 0            |              | dS              | None  |
|                                 | dO         | 0             | 0            |              | dO              | None  |
|                                 | dT         | 0             | 0            |              | dT              | None  |

**TABLE 10-5** Snoop Output and DTag Transition (4 of 4)

| Snooped Request | DTag State | Shared Output | Owned Output | Error Output | Next DTag State | Action for Snoop Pipeline |
|-----------------|------------|---------------|--------------|--------------|-----------------|---------------------------|
| own WS          | dI         | 0             | 0            |              | dI              | Own WS                    |
|                 | dS         | 0             | 0            |              | dI              | Own invalidate WS         |
|                 | dO         | 0             | 0            |              | dI              | Own invalidate WS         |
|                 | dT         | 0             | 0            |              | dI              | Own invalidate WS         |
| foreign WS      | dI         | 0             | 0            |              | dI              | None                      |
|                 | dS         | 0             | 0            |              | dI              | Invalidate                |
|                 | dO         | 0             | 0            |              | dI              | Invalidate                |

**TABLE 10-6** Snoop Input and CIQ Operation Queued

| Action from Snoop Pipeline      | Shared Input | Owned Input | Error (out) | Operation Queued in CIQ    |
|---------------------------------|--------------|-------------|-------------|----------------------------|
| own RTS wait data               | 1            | X           |             | RTS Shared                 |
|                                 | 0            | 0           |             | RTS ~Shared                |
|                                 | 0            | 1           | 1           | RTS Shared, Error          |
| own RTS inst wait data          | X            | X           |             | RTS Shared                 |
| foreign RTS copyback            | X            | 1           |             | Copyback                   |
|                                 | X            | 0           | 1           | Copyback, Error            |
| own RTO no data                 | 1            | X           |             | RTO nodata                 |
|                                 | 0            | X           | 1           | RTO nodata, error          |
| own RTO wait data               | 1            | X           | 1           | RTO data, error            |
|                                 | 0            | X           |             | RTO data                   |
| foreign RTO invalidate          | X            | X           |             | Invalidate                 |
| foreign RTO copyback-invalidate | X            | 0           | 1           | Copyback-invalidate, Error |
|                                 | 0            | 1           |             | Copyback-invalidate        |
|                                 | 1            | 1           |             | Invalidate                 |
| own RS wait data                | X            | X           |             | RS data                    |
| foreign RS copyback-discard     | X            | 0           | 1           | Error                      |
|                                 | X            | 1           |             | Copyback-discard           |
| foreign RTSM copyback           | X            | 0           | 1           | RTSM copyback, Error       |
|                                 | X            | 1           |             | RTSM copyback              |
| own RTSR wait data              | 1            | X           |             | RTSR shared                |

**TABLE 10-6** Snoop Input and CIQ Operation Queued (*Continued*)

| Action from Snoop Pipeline | Shared Input | Owned Input | Error (out) | Operation Queued in CIQ |
|----------------------------|--------------|-------------|-------------|-------------------------|
|                            | 0            | X           |             | RTSR~shared             |
| own RTOR wait data         | X            | X           |             | RTOR data               |
| foreign RTOR invalidate    | X            | X           |             | Invalidate              |
| own RSR                    | X            | X           |             | RS data                 |
| own WS                     | X            | X           |             | Own WS                  |
| own WB                     | X            | X           |             | Own WB                  |
| own invalidate WS          | X            | X           |             | Own invalidate WS       |
| invalidate                 | X            | X           |             | Invalidate              |

## 10.5.3 Transaction Handling

Tables in this section summarize handling of the following:

- Transactions at the head of CIQ (TABLE 10-7)
- No snoop transactions (TABLE 10-8)
- Transactions internal to the UltraSPARC III Cu processor (TABLE 10-9)

**TABLE 10-7** Transaction Handling at Head of CIQ (*1 of 3*)

| Operation at Head of CIQ | CCTag       | MTag (in/out) | Error | Retry | Next CCTag |
|--------------------------|-------------|---------------|-------|-------|------------|
| RTS Shared               | I           | gM (in)       |       |       | S          |
|                          |             | gS (in)       |       |       | S          |
|                          |             | gI (in)       |       | 1     | I          |
|                          | M,O, E,S,Os | X             | 1     |       |            |
| RTS ~ Shared             | I           | gM (in)       |       |       | E          |
|                          |             | gS (in)       |       |       | S          |
|                          |             | gI (in)       |       | 1     | I          |
|                          | M,O, E,S,Os | X (in)        | 1     |       |            |
| RTSR Shared              | I           | gM (in)       |       |       | O          |
|                          |             | gS (in)       |       |       | Os         |
|                          |             | gI (in)       | 1     | 1     | I          |
|                          | M,O, E,S,Os | X             | 1     |       |            |
| RTSR ~ Shared            | I           | gM (in)       |       |       | M          |
|                          |             | gS (in)       |       |       | Os         |

**TABLE 10-7** Transaction Handling at Head of CIQ (2 of 3)

| Operation at Head of CIQ | CCTag          | MTag (in/out) | Error | Retry | Next CCTag                                    |
|--------------------------|----------------|---------------|-------|-------|---|
|                          |                | gI (in)       | 1     | 1     | I   |
|                          | M, O, E, S, Os | X (in)        | 1     |       |   |
| RTO nodata               | I, M, E, Os    | None          | 1     |       |   |
|                          | O, S           | None          |       |       | M   |
| RTO data and SSM         | M, E, O, Os, S | X (in)        | 1     |       |   |
|                          | I              | gM (in)       |       |       | M   |
|                          |                | gS (in)       |       | 1     | Os  |
|                          |                | gI (in)       |       | 1     | I   |
| RTO data and SSM         | M, E, Os, O    | X(in)         | 1     |       |   |
|                          | I, S           | gM(in)        |       |       | M   |
|                          |                | gS(in)        |       | 1     | Os  |
|                          |                | gI(in)        |       | 1     | I   |
| RTOR data                | M, E           | X (in)        | 1     |       |   |
|                          | O              | gM (in)       | 1     |       | O   |
|                          | S, Os, I       | gM (in)       |       |       | M   |
|                          | S, O, Os, I    | gS (in)       | 1     | 1     | Os  |
|                          |                | gI (in)       | 1     | 1     | I   |
| Foreign RTSR             | I              | None          |       |       | I   |
|                          | M, O           | None          | 1     |       | No change                                     |
|                          | E, Os, S       | None          |       |       | S   |
| Foreign RTSM             | I              | X (in)        | 1     |       |   |
|                          | M, O, Os       | None          | 1     |       | S   |
|                          | E, S           | none          |       |       | S   |
| RTSM copyback            | M, O           | gM (out)      |       |       | S (if cacheline is <i>not</i> in the W-cache) |
|                          | M, O           | gM (out)      |       |       | I (if cacheline is in the W-cache)            |
|                          | Os             | gS (out)      |       |       | S   |
|                          | E, S, I        |               | 1     |       |   |
| Copyback                 | M, O           | gM (out)      |       |       | O   |
|                          | Os             | gS (out)      |       |       | Os  |
|                          | I              | gI (out)      |       |       | I   |
|                          | E, S           | gM (out)      | 1     |       | S   |
| Invalidate               | X              |               |       |       | I   |
| Copyback-invalidate      | M, O           | gM (out)      |       |       | I   |

**TABLE 10-7** Transaction Handling at Head of CIQ (3 of 3)

| Operation at Head of CIQ | CCTag       | MTag (in/out)              | Error | Retry | Next CCTag                          |
|--------------------------|-------------|----------------------------|-------|-------|-------------------------------------|
|                          | Os          | gS (out)                   |       |       | I                                   |
|                          | I           | gI (out)                   |       |       | I                                   |
|                          | E, S        | gM (out)                   | 1     |       | I                                   |
| Copyback-discard         | M, O        | gM (out)                   |       |       | No change                           |
|                          | Os          | gS (out)                   |       |       | Os                                  |
|                          | I           | gI (out)                   |       |       | I                                   |
|                          | E, S        | gM (out)                   | 1     |       | No change                           |
| RS data                  | X<br>X<br>X | gM(in)<br>gS(in)<br>gI(in) |       | 1     | No change<br>No change<br>No change |
| Own WS                   | X           | gM(out)                    |       |       | I                                   |
| Own WB                   | M,O         | gM (out)                   |       |       | I                                   |
|                          | Os          | gS (out)                   |       |       | I                                   |
|                          | I           | gI (out)                   |       |       | I                                   |
|                          | S, E        | gM (out)                   | 1     |       | I                                   |
| Own invalidate WS        | X           | gM(out)                    |       |       | I                                   |

**TABLE 10-8** No Snoop Transaction Handling

| Combined State | Mode     | Processor Action |                 |                |             |                    |              |
|----------------|----------|------------------|-----------------|----------------|-------------|--------------------|--------------|
|                |          | Load             | Store/ Swap     | Block Load     | Block Store | Block Store Commit | Dirty Victim |
| I              | No snoop | Miss:<br>RTS_ns  | Miss:<br>RTO_ns | Miss:<br>RS_ns | Miss:<br>WS | Miss:<br>WS        | None         |
| S              | No snoop | Error            |                 |                |             |                    |              |
| E              | No snoop | Hit              | Hit<br>E→M      | Hit            | Hit<br>E→M  | Miss:<br>WS        | None         |
| M              | No snoop | Hit              | Hit             | Hit            | Hit         | Miss:<br>WS        | WB           |
| O              | No snoop | Error            |                 |                |             |                    |              |
| Os             | No snoop | Error            |                 |                |             |                    |              |

RTS\_ns, RTO\_ns, and RS\_ns are transactions internal to the UltraSPARC III Cu processor and are not visible on the Sun™ Fireplane interconnect.

**TABLE 10-9** Internal Transaction Handling

| Operation at Head of CIQ | CCTag          | MTag (in/out)            | Error  | Next CCTag |
|--------------------------|----------------|--------------------------|--------|------------|
| RTS_ns                   | I              | gM (in)                  | 1      | E          |
|                          | S, E, M, O, Os | gS (in)<br>gI (in)       | 1<br>1 | S<br>I     |
| RTO_ns                   | I              | gM (in)                  | 1      | M          |
|                          | S, E, M, O, Os | gS (in)<br>gI (in)       | 1<br>1 | O<br>I     |
| RS_ns                    | I              | gM (in)                  | 1      | I          |
|                          | S, E, M, O, Os | gS (in)<br>gI (in)       | 1<br>1 | I<br>I     |
| WS                       | X              | gM (out)<br>gS, gI (out) | 1      | I<br>I     |





# Memory Management Unit

---

The Memory Management Unit (MMU) conforms to the requirements set forth in the *SPARC V9 Architecture Manual*. In particular, it supports a 64-bit virtual address space, software TLB miss processing only (no hardware page table walk), simplified protection encoding, and multiple page sizes.

This chapter describes the MMU, as seen by the operating system software, in these sections:

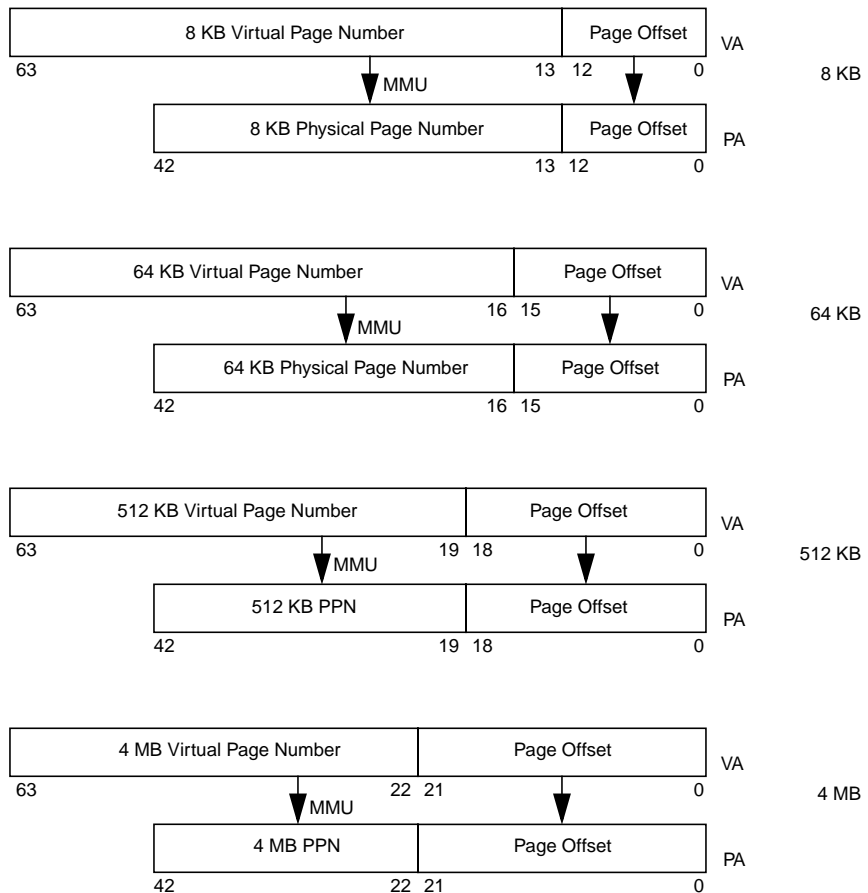
- “Virtual Address Translation” on page 245
- “Translation Table Entry” on page 248
- “Translation Storage Buffer” on page 251
- “Hardware Support for TSB Access” on page 253
- “Faults and Traps” on page 259
- “ASI Value, Context, and Endianness Selection for Translation” on page 261
- “Reset, Disable, and RED\_state Behavior” on page 263
- “SPARC V9 “MMU Requirements” Annex” on page 264
- “Data Translation Lookaside Buffer” on page 265
- “Instruction Translation Lookaside Buffer” on page 294

---

## 11.1 Virtual Address Translation

The MMUs support four page sizes: 8 KB, 64 KB, 512 KB, and 4 MB. Separate Instruction and Data MMUs (I-MMU and D-MMU, respectively) are provided to enable concurrent virtual-to-physical address translations for instruction and data. A 64-bit virtual address (VA) space is supported, with 43 bits of physical address (PA). In each translation, the virtual page number is replaced by a physical page number, which is concatenated with the page offset to form the full physical address, as illustrated in FIGURE 11-1.

Each data and I-MMU consists of multiple Translation Lookaside Buffers (TLBs) that can be accessed in parallel.



**FIGURE 11-1** Virtual-to-Physical Address Translation for All Four Page Sizes

The operating system maintains translation information in an arbitrary data structure, called the *software translation table* in this chapter. The I-MMU and D-MMU TLBs act as independent, software managed caches of the software translation table, providing appropriate concurrency for virtual-to-physical address translation.


In the D-MMU, there are two TLBs each of which have 512-entry two-way associative which can be programmed to hold unlocked 8 KB, 64 KB, 512 KB, and 4 MB pages for translations. Each TLB can hold any of the 4 page sizes, but are programmed to only one page size at any given time. Each TLB can be programmed to the either same or different page sizes. Also, a 16-entry fully associative TLB is used for 8 KB, 64 KB, 512 KB and

4 MB locked and unlocked pages. Specifically, it supports locked pages of the size supported in 512-entry, 2-way TLBs and unlocked and locked pages of the other page sizes. This TLB can hold any combination of page sizes that are locked/unlocked at a time.

In the I-MMU, a 128-entry, 2-way associative TLB is used exclusively for 8 KB page translations, and a 16-entry fully associative TLB is used for 64 KB, 512 KB, and 4 MB page translations and locked pages of all four sizes.

On a TLB miss, the MMU immediately traps to software for TLB miss processing. The TLB miss handler can fill the TLB by any available means, but it is likely to take advantage of the TLB miss support features provided by the MMU, since the TLB miss handler is time-critical code.

A general software view of the MMU is shown in FIGURE 11-2. The TLBs, which are part of the MMU hardware, are small and fast. The software translation table is likely to be large and complex. The translation storage buffer (TSB), which acts like a direct-mapped cache, is the interface between the two. The TSB can be shared by all processes running on a processor or can be process specific. The hardware does not require any particular scheme.



**FIGURE 11-2** Software View of the MMU

Aliasing between pages of different sizes (when multiple virtual addresses map to the same physical address) can occur. However, the reverse case of multiple mappings from one virtual address to multiple physical addresses producing a multiple TLB match is not necessarily detected in hardware and may produce undefined results.

---


**Note** – The hardware ensures the physical reliability of the TLB on multiple matches.

---

## 11.2 Translation Table Entry

The Translation Table Entry (TTE) is the equivalent of a SPARC V8 page table entry; it holds information for a single page mapping. The TTE is divided into two 64-bit words representing the tag and data of the translation. Just as in a hardware cache, the tag is used to determine whether there is a hit in the TSB; if there is a hit, the data is fetched by software.

The configuration of the TTE is illustrated in FIGURE 11-3 and described in TABLE 11-1.



**FIGURE 11-3** Translation Storage Buffer (TSB) Translation Table Entry (TTE)

**TABLE 11-1** TSB and TTE Bit Description (1 of 4)

| Bit         | Field             | Description   |
|-------------|-------------------|---|
| Tag – 63    | G                 | Global. If the Global bit is set, the Context field of the TLB entry is ignored during hit detection. This behavior allows any page to be shared among all (user or supervisor) contexts running in the same processor. The Global bit is duplicated in the TTE tag and data to optimize the software miss handler.   |
| Tag – 62:61 | —                 | <i>Reserved.</i>  |
| Tag – 60:48 | Context           | The 13-bit context identifier associated with the TTE.  |
| Tag – 47:42 | —                 | <i>Reserved.</i>  |
| Tag – 41:0  | VA_tag<br><63:22> | Virtual Address Tag. The virtual page number. Bits 21 through 13 are not maintained in the tag because these bits index the minimally sized, direct-mapped TSB of 512 entries.  |
| Data – 63   | V                 | Valid. If the Valid bit is set, then the remaining fields of the TTE are meaningful. Note that the explicit Valid bit is redundant with the software convention of encoding an invalid TTE with an unused context. The encoding of the context field is necessary to cause a failure in the TTE tag comparison, and the explicit Valid bit in the TTE data simplifies the TLB miss handler. |

**TABLE 11-1** TSB and TTE Bit Description (2 of 4)

| Bit          | Field     | Description  |           |           |    |      |    |       |    |        |    |      |
|--------------|-----------|--|-----------|-----------|----|------|----|-------|----|--------|----|------|
| Data – 62:61 | Size      | <p>The page size of this entry, encoded as shown below.</p> <table border="1"> <thead> <tr> <th>Size&lt;1:0&gt;</th> <th>Page Size</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>8 KB</td> </tr> <tr> <td>01</td> <td>64 KB</td> </tr> <tr> <td>10</td> <td>512 KB</td> </tr> <tr> <td>11</td> <td>4 MB</td> </tr> </tbody> </table>  | Size<1:0> | Page Size | 00 | 8 KB | 01 | 64 KB | 10 | 512 KB | 11 | 4 MB |
| Size<1:0>    | Page Size |  |           |           |    |      |    |       |    |        |    |      |
| 00           | 8 KB      |  |           |           |    |      |    |       |    |        |    |      |
| 01           | 64 KB     |  |           |           |    |      |    |       |    |        |    |      |
| 10           | 512 KB    |  |           |           |    |      |    |       |    |        |    |      |
| 11           | 4 MB      |  |           |           |    |      |    |       |    |        |    |      |
| Data – 60    | NFO       | No Fault Only. If the no-fault-only bit is set, loads with ASI_PRIMARY_NO_FAULT, ASI_SECONDARY_NO_FAULT, and their *_LITTLE variations are translated. Any other access will trap with a <i>data_access_exception</i> trap (FT = 10 <sub>16</sub> ). The NFO bit in the I-MMU is read as 0 and ignored when written. The I-TLB miss handler should generate an error if this bit is set before the TTE is loaded into the TLB.   |           |           |    |      |    |       |    |        |    |      |
| Data – 59    | IE        | Invert Endianness. If this bit is set for a page, accesses to the page are processed with inverse endianness from that specified by the instruction (big for little, little for big). See page 261 for details. The IE bit in the I-MMU is read as 0 and ignored when written.<br><b>Note:</b> This bit is intended to be set primarily for non-cacheable accesses. The performance of cacheable accesses will be degraded as if the access missed the D-cache.  |           |           |    |      |    |       |    |        |    |      |
| Data – 58:50 | Soft2     | Software-defined field, provided for use by the operating system. The Soft2 field can be written with any value in the TSB. Hardware is not required to maintain this field in the TLB; therefore, when it is read from the TLB, it may read as zero.  |           |           |    |      |    |       |    |        |    |      |
| Data – 49:43 | Reserved  | Reserved. Value on read is undefined.  |           |           |    |      |    |       |    |        |    |      |
| Data – 42:13 | PA        | The physical page number. Page offset bits for larger page sizes (PA<15:13>, PA<18:13>, and PA<21:13> for 64 KB, 512 KB, and 4 MB pages, respectively) are stored in the TLB and returned for a Data Access read but are ignored during normal translation. When page offset bits for larger page sizes (PA<15:13>, PA<18:13>, and PA<21:13> for 64 KB, 512 KB, and 4 MB pages, respectively) are stored in the TLB on the UltraSPARC III Cu processor, the data returned from those fields by a Data Access read are the data previously written to them. |           |           |    |      |    |       |    |        |    |      |
| Data – 12:7  | Soft      | Software-defined field, provided for use by the operating system. The Soft field can be written with any value in the TSB. Hardware is not required to maintain this field in the TLB; therefore, when it is read from the TLB, it may read as zero.   |           |           |    |      |    |       |    |        |    |      |

**TABLE 11-1** TSB and TTE Bit Description (3 of 4)

| Bit                            | Field                                | Description  |                                |  |  |                    |                                      |                               |        |               |               |    |                             |  |    |                             |   |
|--------------------------------|--------------------------------------|--|--------------------------------|--|--|--------------------|--------------------------------------|-------------------------------|--------|---------------|---------------|----|-----------------------------|--|----|-----------------------------|---|
| Data – 6                       | L                                    | <p>If the lock bit is set, then the TTE entry will be “locked down” when it is loaded into the TLB; that is, if this entry is valid, it will not be replaced by the automatic replacement algorithm invoked by an ASI store to the Data In Register. The lock bit has no meaning for an invalid entry. Arbitrary entries can be locked down in the TLB. Software must ensure that at least one entry is not locked when replacing a TLB entry; otherwise, a locked entry will be replaced. Since the 16-entry, fully associative TLB is shared for all locked entries as well as for 4 MB and 512 KB pages, the total number of locked pages is limited to less than or equal to 15.</p> <p>In the UltraSPARC III Cu processor, the TLB lock bit is implemented in the D-MMU 16-entry, fully associative TLB, and the I-MMU 16-entry, fully associative TLB. In the D-MMU 512-entry, 2-way associative TLB and I-MMU 128-entry, 2-way associative TLB, each TLB entry’s lock bit reads as 0 and writes to it are ignored.</p> <p>The lock bit set for 8 KB page translation in both I-MMU and D-MMU is read as 0 and ignored when written.</p> |                                |  |  |                    |                                      |                               |        |               |               |    |                             |  |    |                             |   |
| Data – 5<br>Data – 4           | CP,<br>CV                            | <p>The cacheable-in-physically-indexed-cache bit and cacheable-in-virtually-indexed-cache bit determine the placement of data in the caches. The UltraSPARC III Cu processor fully implements the CV bit. The following table describes how CP and CV control cacheability in specific UltraSPARC III Cu processor caches.</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th colspan="3">Meaning of TTE when placed in:</th> </tr> <tr> <th>Cacheable (CP, CV)</th> <th>I-TLB (Instruction Cache PA-indexed)</th> <th>D-TLB (Data Cache VA-indexed)</th> </tr> </thead> <tbody> <tr> <td>00, 01</td> <td>Non-cacheable</td> <td>Non-cacheable</td> </tr> <tr> <td>10</td> <td>Cacheable L2-cache, I-cache</td> <td>Cacheable L2-cache, W-cache, and P-cache</td> </tr> <tr> <td>11</td> <td>Cacheable L2-cache, I-cache</td> <td>Cacheable L2-cache, D-cache, W-cache, and P-cache</td> </tr> </tbody> </table> <p>The MMU does not operate on the cacheable bits but merely passes them through to the cache subsystem. The CV bit in the I-MMU is read as zero and ignored when written.</p>            | Meaning of TTE when placed in: |  |  | Cacheable (CP, CV) | I-TLB (Instruction Cache PA-indexed) | D-TLB (Data Cache VA-indexed) | 00, 01 | Non-cacheable | Non-cacheable | 10 | Cacheable L2-cache, I-cache | Cacheable L2-cache, W-cache, and P-cache | 11 | Cacheable L2-cache, I-cache | Cacheable L2-cache, D-cache, W-cache, and P-cache |
| Meaning of TTE when placed in: |                                      |  |                                |  |  |                    |                                      |                               |        |               |               |    |                             |  |    |                             |   |
| Cacheable (CP, CV)             | I-TLB (Instruction Cache PA-indexed) | D-TLB (Data Cache VA-indexed)  |                                |  |  |                    |                                      |                               |        |               |               |    |                             |  |    |                             |   |
| 00, 01                         | Non-cacheable                        | Non-cacheable  |                                |  |  |                    |                                      |                               |        |               |               |    |                             |  |    |                             |   |
| 10                             | Cacheable L2-cache, I-cache          | Cacheable L2-cache, W-cache, and P-cache   |                                |  |  |                    |                                      |                               |        |               |               |    |                             |  |    |                             |   |
| 11                             | Cacheable L2-cache, I-cache          | Cacheable L2-cache, D-cache, W-cache, and P-cache  |                                |  |  |                    |                                      |                               |        |               |               |    |                             |  |    |                             |   |
| Data – 3                       | E                                    | <p>Side-effect. If the side-effect bit is set, non-faulting loads will trap for addresses within the page, non-cacheable memory accesses other than block loads and stores are strongly ordered against other E-bit accesses, and non-cacheable stores are not merged. This bit should be set for pages that map I/O devices having side-effects. Note, however, that the E bit does not prevent normal instruction prefetching. The E bit in the I-MMU is read as 0 and ignored when written.</p> <p><b>Note:</b> The E bit does not force a non-cacheable access. It is expected, but not required, that the CP and CV bits will be set to 0 when the E bit is set. If both the CP and CV bits are set to 1 along with the E bit, the result is undefined.</p> <p><b>Note Also:</b> The E bit and the NFO bit are mutually exclusive; both bits should never be set in any TTE.</p>  |                                |  |  |                    |                                      |                               |        |               |               |    |                             |  |    |                             |   |

**TABLE 11-1** TSB and TTE Bit Description (4 of 4)

| Bit      | Field | Description  |
|----------|-------|--|
| Data – 2 | P     | Privileged. If the P bit is set, only the supervisor 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, then the MMU signals an <i>instruction_access_exception</i> or <i>data_access_exception</i> trap (FT = 1 <sub>16</sub> ). |
| Data – 1 | 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 causes a <i>fast_data_access_protection</i> trap if a write is attempted. The W bit in the I-MMU is read as 0 and ignored when written.               |
| Data – 0 | G     | Global. This bit must be identical to the Global bit in the TTE tag. Like the Valid bit, the Global bit in the TTE tag is necessary for the TSB hit comparison, and the Global bit in the TTE data facilitates the loading of a TLB entry.   |



## 11.3 Translation Storage Buffer

The Translation Storage Buffer (TSB) is an array of Translation Table Entries managed entirely by software. It serves as a cache of the software translation table, used to quickly reload the TLB in the event of a TLB miss. The discussion in this section assumes the use of the hardware support for TSB access described in “Hardware Support for TSB Access” on page 253, although the operating system is not required to make use of this support hardware.

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

A bit in the TSB register allows the TSB 64 KB pointer to be computed for the case of common or split 8 KB/64 KB TSBs.

### 11.3.1 TSB Indexing Support

No hardware TSB indexing support is provided for the 512 KB and 4 MB page TTEs. However, since the TSB is entirely software managed, the operating system may choose to place these larger page TTEs in the TSB by forming the appropriate pointers. In addition, simple modifications to the 8 KB and 64 KB index pointers provided by the hardware allow formation of an M-way, set-associative TSB, multiple TSBs per page size, and multiple TSBs per process.

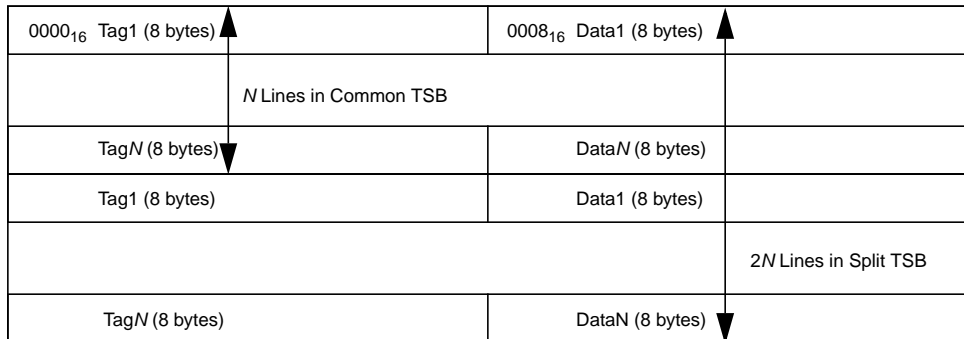
## 11.3.2 TSB Cacheability

The TSB exists as a normal data structure in memory and therefore can be cached. Indeed, the speed of the TLB miss handler relies on the TSB accesses hitting the level-2 cache at a substantial rate. This policy may result in some conflicts with normal instruction and data accesses, but it is hoped that the dynamic sharing of the level-2 cache resource will provide a better overall solution than that provided by a fixed partitioning.

## 11.3.3 TSB Organization

The TSB is arranged as a direct-mapped cache of TTEs. The MMU provides pre-computed pointers into the TSB for the 8 KB and 64 KB page TTEs. In each case,  $n$  least significant bits of the respective virtual page number are used as the offset from the TSB base address, with  $n$  equal to log base 2 of the number of TTEs in the TSB.

The TSB organization is illustrated in FIGURE 11-4. The constant  $N$  is determined by the `Size` field in the TSB Register. The `Size` Field is 3 bits and can have a value from 0 to 7. For the UltraSPARC III Cu processor, the `Size` Field is set to 7. The constant  $N$  is equal to  $512 \times 2^{\text{size}}$ . Therefore, for the UltraSPARC III Cu processor, the constant  $N$  is equal to 64K.



**FIGURE 11-4** TSB Organization, Illustrated for Both Common and Shared Cases



---

## 11.4 Hardware Support for TSB Access

The MMU hardware provides services to allow the TLB miss handler to efficiently reload a missing TLB entry for an 8 KB or 64 KB page. These services include:

- Formation of TSB Pointers, based on the missing virtual address and address space identifier
- Formation of the TTE Tag Target used for the TSB tag comparison
- Efficient atomic write of a TLB entry with a single store ASI operation
- Alternate globals on MMU-signalled traps

### 11.4.1 Typical TLB Miss/Refill Sequence

The following is a typical TLB miss and TLB refill sequence:

1. A TLB miss causes either a *fast\_instruction\_access\_MMU\_miss* or a *fast\_data\_access\_MMU\_miss* exception.
2. The appropriate TLB miss handler loads the TSB Pointers and the TTE Tag Target with loads from the MMU registers.
3. Using this information, the TLB miss handler checks to see if the desired TTE exists in the TSB. If so, the TTE data is loaded into the TLB Data In Register to initiate an atomic write of the TLB entry chosen by the replacement algorithm.
4. If the TTE does not exist in the TSB, then the TLB miss handler jumps to the more sophisticated, and slower, TSB miss handler.

The virtual address used in the formation of the pointer addresses comes from the Tag Access Register, which holds the virtual address and context of the load or store responsible for the MMU exception.

---

**Note** – There are no separate physical registers in hardware for the pointer registers; rather, they are implemented through a dynamic reordering of the data stored in the Tag Access and the TSB registers.

---

## 11.4.2 TSB Pointer Formation

Hardware provides pointers for the most common cases of 8 KB and 64 KB page miss processing. These pointers give the virtual addresses where the 8 KB and 64 KB TTEs are stored if either are present in the TSB.

### 11.4.2.1 Input Values for TSB Pointer Formation

The pointer to the TTE in the TSB is generated from the following parameters as inputs:

- TSB base address (`TSB_Base`)
- Virtual address (VA)
- `TSB_size`
- `TSB_split`
- `TSB_Hash`

The TSB base address is in either of I/D Primary/Secondary (provided only for data)/Nucleus TSB Extension Registers. Depending on the context that generated the TLB miss, an appropriate TSB Extension Register is selected (which may be combined with the `TSB_Base` field from the TSB Base Register; see Note). Note that the context with the TLB miss is logged in the I/D Synchronous Fault Status Register.

`TSB_size` and `TSB_split` are supplied also from the selected TSB Extension Register.

The virtual page number to be used for TSB pointer formation is in the I/D Tag Access Register.

---

**Note** – `TSB_Base` address may be generated by exclusive-ORing `TSB_Base` register and TSB Extension Register contents, for compatibility with the UltraSPARC I and UltraSPARC II processor TSB pointer formation. In this case, if the TSB Extension Registers hold 0 as `TSB_Base`, the value in `TSB_Base` register becomes the `TSB_Base` address, thereby maintaining compatibility with the UltraSPARC I and UltraSPARC II processor TLB miss handling software. In addition, `TSB_Base` may be taken directly from an appropriate TSB Extension Register. In that case, the implementation should provide the way to maintain compatibility with the UltraSPARC I and UltraSPARC II processor TLB miss handler software.

---

The UltraSPARC III Cu processor generates the `TSB_Base` address by taking the `TSB_Base` field directly from the TSB Extension Register.

## 11.4.2.2 TSB Pointer Formation

Hardware uses the following equations to form TSB pointers for TLB misses. In the equations,  $n$  is defined to be the `TSB_Size` field of the TSB register; it ranges from 0 to 7. Note that `TSB_Size` refers to the size of each TSB when the TSB is split. The  $\square$  symbol designates concatenation of bit vectors and  $\oplus$  indicates an exclusive-OR operation.

### *Exclusive-ORed TSB\_Base*

- **For a shared TSB (TSB Register split field = 0):**

$$8K\_POINTER = TSB\_Base[63:13+n] \oplus TSB\_Extension[63:13+n] \square VA[21+n:13] \square 0000$$

$$64K\_POINTER = TSB\_Base[63:13+n] \oplus TSB\_Extension[63:13+n] \square VA[24+n:16] \square 0000$$

- **For a split TSB (TSB Register split field = 1):**

$$8K\_POINTER = TSB\_Base[63:14+n] \oplus TSB\_Extension[63:14+n] \square 0 \square VA[21+n:13] \square 0000$$

$$64K\_POINTER = TSB\_Base[63:14+n] \oplus TSB\_Extension[63:14+n] \square 1 \square VA[24+n:16] \square 0000$$

### *TSB\_Base from TSB Extension Registers*

- **For a shared TSB (TSB Register split field = 0):**

$$8K\_POINTER = TSB\_Extension[63:13+n] \square (VA[21+n:13] \oplus TSB\_Hash) \square 0000$$

$$64K\_POINTER = TSB\_Extension[63:13+n] \square (VA[24+n:16] \oplus TSB\_Hash) \square 0000$$

- **For a split TSB (TSB Register split field = 1):**

$$8K\_POINTER = TSB\_Extension[63:14+n] \square 0 \square (VA[21+n:13] \oplus TSB\_Hash) \square 0000$$

$$64K\_POINTER = TSB\_Extension[63:14+n] \square 1 \square (VA[24+n:16] \oplus TSB\_Hash) \square 0000$$

The TSB Tag Target is formed by aligning the missing access VA (from the Tag Access Register) and the current context to positions found above in the description of the TTE tag, allowing a simple XOR instruction for TSB hit detection.

# 11.4.3 TSB Pointer Logic Hardware Description

FIGURE 11-5 illustrates the generation of the 8 KB and 64 KB pointers; CODE EXAMPLE 11-1 presents pseudocode for D-MMU pointer logic.

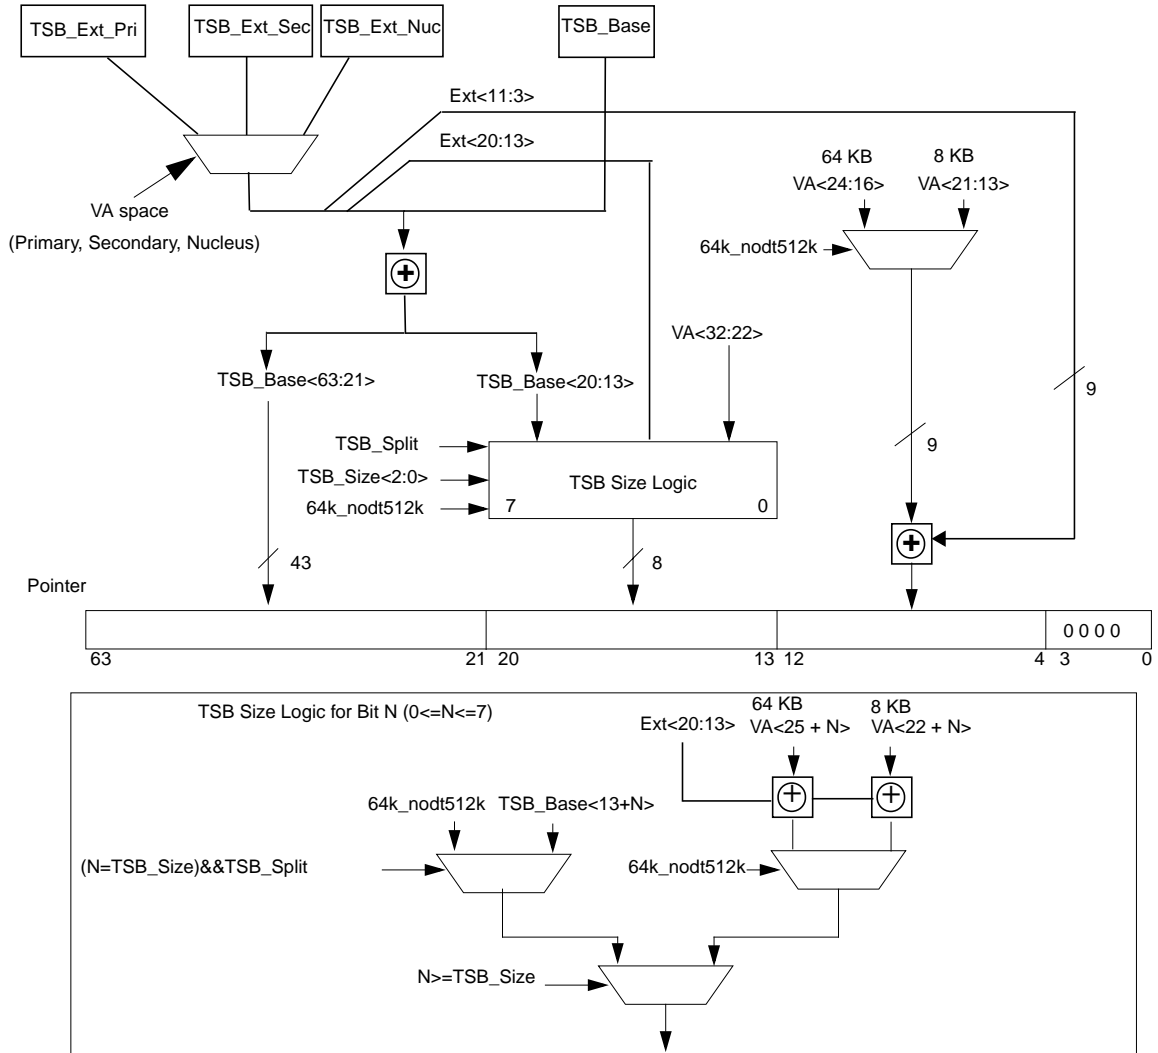


FIGURE 11-5 Formation of TSB Pointers for 8 KB and 64 KB TTE

**CODE EXAMPLE 11-1** Pseudocode for D-MMU Pointer Logic

```
int64 GenerateTSBPointer(  
    int64 va,          // Missing virtual address  
    PointerType type, // 8K_POINTER or 64K_POINTER  
    int64 TSBBase,    // TSB Register<63:13> << 13  
    Boolean split,    // TSB Register<12>  
    int TSBSize,      // TSB Register<2:0>  
    int SpaceType space)  
{  
    int64 vaPortion;  
    int64 TSBBaseMask;  
    int64 splitMask;  
  
    // Shift va towards lsb appropriately and  
    // zero out the original va page offset  
    vaPortion = (va >> ((type == 8K_POINTER)? 9: 12)) &  
                0xfffffffffffffff0;  
  
    switch (space) {  
    Primary:  
        TSBBASE ^=TSB_EXT_pri;  
        vaPortion ^= TSB_EXT_pri<<1 & 0x1ff0;  
        vaPortion ^= TSB_EXT_pri & 0x1fe000;  
        break;  
    Secondary:  
        TSBBASE ^=TSB_EXT_sec;  
        vaPortion ^= TSB_EXT_sec<<1 & 0x1ff0;  
        vaPortion ^= TSB_EXT_sec & 0x1fe000;  
        break;  
    Nucleus:  
        TSBBASE ^=TSB_EXT_nuc;  
        vaPortion ^= TSB_EXT_nuc<<1 & 0x1ff0;  
        vaPortion ^= TSB_EXT_nuc & 0x1fe000;  
        break;  
    }  
    // TSBBaseMask marks the bits from TSB Base Reg  
    TSBBaseMask = 0xffffffffffffe000 <<  
                (split? (TSBSize + 1) : TSBSize);  
  
    if (split) {  
        // There's only one bit in question for split  
        splitMask = 1 << (13 + TSBSize);  
        if (type == 8K_POINTER)  
            // Make sure we're in the lower half  
            vaPortion &= ~splitMask;  
        else  
            // Make sure we're in the upper half  
            vaPortion |= splitMask;  
    }  
    return (TSBBase & TSBBaseMask) | (vaPortion & ~TSBBaseMask);  
}
```

## 11.4.4 Required TLB Conditions

The following items must be locked in the TLB to avoid an error condition:

- TLB miss handler and data
- TSB and linked data
- Asynchronous trap handlers and data

## 11.4.5 Required TSB Conditions

The following items must be locked in the TSB (not necessarily the TLB) to avoid an error condition:

- TSB miss handler and data
- Interrupt-vector handler and data

## 11.4.6 MMU Global Registers Selection

In the SPARC V9 normal trap model, the software is presented with an alternate set of global registers in the Integer Register file. An UltraSPARC III Cu processor provides an additional feature to facilitate fast handling of TLB misses. For the following traps, the trap handler is presented with a special set of MMU globals:

- *fast\_instruction\_access\_MMU\_miss*
- *fast\_data\_access\_MMU\_miss*
- *instruction\_access\_exception*
- *data\_access\_exception*
- *fast\_data\_access\_protection*

Trap handlers for the *privileged\_action*, *mem\_address\_not\_aligned*, and *\*\_mem\_address\_not\_aligned* traps use the standard alternate global registers.

---

**Compatibility Note** – The MMU does not perform hardware tablewalking. The MMU hardware never directly reads or writes the TSB.

---

## 11.5 Faults and Traps

The traps recorded by the MMU are listed in TABLE 11-2 and described below the table, by the reference number. All listed traps are precise traps.

**TABLE 11-2** MMU Trap Types, Causes, and Stored State Register Update Policy

| Ref # | Trap Name   | Trap Cause                      | Registers Updated<br>(Stored State in MMU) |                        |                 |                        | Trap Type  |
|-------|---|---------------------------------|--|------------------------|-----------------|------------------------|--|
|       |   |                                 | I-SFSR                                     | I-MMU<br>Tag<br>Access | D-SFSR,<br>SFAR | D-MMU<br>Tag<br>Access |  |
| 1.    | <i>fast_instruction_access_MMU_miss</i>                               | I-TLB miss                      | X  | X                      |                 |                        | 64 <sub>16</sub> –<br>67 <sub>16</sub>   |
| 2.    | <i>instruction_access_exception</i>                                   | Privilege violation for I-fetch | X  | X                      |                 |                        | 08 <sub>16</sub>   |
| 3.    | <i>fast_data_access_MMU_miss</i>                                      | D-TLB miss                      |  |                        | X               | X                      | 68 <sub>16</sub> –<br>6B <sub>16</sub>   |
| 4.    | <i>data_access_exception</i>  | Several (see below)             |  |                        | X               | X <sup>†</sup>         | 30 <sub>16</sub>   |
| 5.    | <i>fast_data_access_protection</i>                                    | Protection violation            |  |                        | X               | X                      | 6C <sub>16</sub> –<br>6F <sub>16</sub>   |
| 6.    | <i>privileged_action</i>  | Use of privileged ASI           |  |                        | X               |                        | 37 <sub>16</sub>   |
| 7.    | watchpoints   | Watchpoint hit                  |  |                        | X               |                        | 61 <sub>16</sub> –<br>62 <sub>16</sub>   |
| 8.    | <i>mem_address_not_aligned</i> ,<br>* <i>_mem_address_not_aligned</i> | Misaligned mem op               |  |                        | X               |                        | 35 <sub>16</sub> ,<br>36 <sub>16</sub> ,<br>38 <sub>16</sub> ,<br>39 <sub>16</sub> |

<sup>†</sup> The contents of the context field of the D-MMU Tag Register are undefined after a *data\_access\_exception*.

---

**Note** – In an UltraSPARC III Cu processor, *fast\_instruction\_access\_MMU\_miss*, *fast\_data\_access\_MMU\_miss*, and *fast\_data\_access\_protection* traps are generated instead of SPARC V9 *instruction\_access\_MMU\_miss*, *data\_access\_MMU\_miss*, and *data\_access\_protection* traps, respectively.

---

1. ***fast\_instruction\_access\_MMU\_miss*** — Occurs when the MMU is unable to find a translation for an instruction access; that is, when the appropriate TTE is not in the I-TLB. In an UltraSPARC III Cu processor, the *fast\_instruction\_access\_MMU\_miss* exception is generated instead of the SPARC V8 *instruction\_access\_MMU\_miss*.

2. ***instruction\_access\_exception*** — Occurs when the I-MMU is enabled and detects a privilege violation for an instruction fetch; that is, an attempted access to a privileged page when `PSTATE.PRIV = 0`.
3. ***fast\_data\_access\_MMU\_miss*** — Occurs when the MMU is unable to find a translation for a data access; that is, when the appropriate TTE is not in the D-TLB.  
In an UltraSPARC III Cu processor, the *fast\_data\_access\_MMU\_miss* exception is generated instead of the SPARC V9 *data\_access\_MMU\_miss* trap.
4. ***data\_access\_exception*** — Signalled upon the detection of at least one of the following exceptional conditions:
  - The D-MMU detects a privilege violation for a data access; that is, an attempted access to a privileged page when `PSTATE.PRIV = 0`.
  - A speculative (non-faulting) load instruction issued to a page marked with the side-effect (E bit) set to 1, including cases in which the D-MMU is disabled.
  - An atomic instruction (including 128-bit atomic load) issued to a memory address marked non-cacheable in a physical cache; that is, with the CP bit set to 0, including cases in which the D-MMU is disabled.
  - An invalid LDA/STA ASI value, read to write-only register, or write to read-only register. Not for an attempted user access to a restricted ASI (see the *privileged\_action* trap described below).
  - An access with an ASI other than `“(PRIMARY, SECONDARY)_NO_FAULT(_LITTLE)”` to a page marked with the NFO (no-fault-only) bit.
5. ***fast\_data\_access\_protection*** — Occurs when the MMU detects a protection violation for a data access. A protection violation is defined to be an attempted store (including atomic load-store operations) to a page that does not have write permission.  
In an UltraSPARC III Cu processor, the *fast\_data\_access\_protection* exception is generated instead of the SPARC V9 *data\_access\_protection* trap.
6. ***privileged\_action*** — Occurs when an access is attempted using a *restricted* ASI while in non-privileged mode (`PSTATE.PRIV = 0`).
7. ***watchpoints*** — *PA\_watchpoint* and *VA\_watchpoint* traps are included in this category. Watchpoint traps occur when watchpoints are enabled and the D-MMU detects a load or store to the virtual or physical address specified by the watchpoint virtual or physical registers, respectively. See Section 6.10.2 “Data Watchpoint Registers” on page 6-136. The trap is precise and is signalled before the actual event, meaning that the contents of the location are not modified when the trap is invoked.
8. ***mem\_address\_not\_aligned*** — Occurs when a load, store, atomic, JMPL, or RETURN instruction with a misaligned address is executed.



On a *mem\_address\_not\_aligned* trap that occurs during a JMWPL or RETURN instruction, the UltraSPARC III Cu processor updates the D-SFAR and D-SFSR registers with the fault address and status, respectively.

---

## 11.6 ASI Value, Context, and Endianness Selection for Translation

The selection of the context for a translation is the result of a two-step process:

1. The ASI is determined (conceptually by the Integer Unit) from the instruction, ASI register, trap level, and the processor endian mode (`PSTATE.CLE`).
2. The Context Register is determined directly from the ASI. The context value is read by the Context Register selected by the ASI.

The ASI value and endianness (little or big) are determined for the I-MMU and D-MMU, respectively, according to TABLE 11-3 through TABLE 11-5. Note that the secondary context is never used to fetch instructions. The I-MMUs and D-MMUs, when using the Primary Context identifier, use the value stored in the shared Primary Context Register.

The endianness of a data access is specified by the following three conditions:

- ASI specified in the opcode or ASI register
- `PSTATE` current little-endian bit (`CLE`)
- D-MMU invert endianness bit

The D-MMU invert endianness bit does not affect the ASI value recorded in the `SFSR` but does invert the endianness that is otherwise specified for the access.

---

**Note** – The D-MMU invert endianness bit inverts the endianness for all accesses, including load/store/atomic alternates that have specified an ASI. That is, `LDXA [%g1]` `ASI_PRIMARY_LITTLE` will be `_BIG` if the `IE` bit is on.

---

**TABLE 11-3** ASI Mapping for Instruction Access

| Condition for Instruction Access | Resulting Action |                          |
|----------------------------------|------------------|--------------------------|
|                                  | Endianness       | ASI Value (in SFSR)      |
| <code>PSTATE.TL</code>           |                  |                          |
| 0                                | BIG              | <code>ASI_PRIMARY</code> |
| > 0                              | BIG              | <code>ASI_NUCLEUS</code> |

**TABLE 11-4** ASI Mapping for Data Accesses

| Condition for Data Access  |     |            |         | Access Processed with: |  |
|--|-----|------------|---------|------------------------|--|
| Opcode   | TL  | PSTATE.CLE | DMMU.IE | Endian                 | ASI Value (Recorded in SFSR)   |
| Load/Store/Atomic  | 0   | 0          | 0       | BIG                    | ASI_PRIMARY  |
|  |     |            | 1       | LITTLE                 | ASI_PRIMARY  |
|  |     | 1          | 0       | LITTLE                 | ASI_PRIMARY_LITTLE   |
|  |     |            | 1       | BIG                    | ASI_PRIMARY_LITTLE   |
|  | > 0 | 0          | 0       | BIG                    | ASI_NUCLEUS  |
|  |     |            | 1       | LITTLE                 | ASI_NUCLEUS  |
|  |     | 1          | 0       | LITTLE                 | ASI_NUCLEUS_LITTLE   |
|  |     |            | 1       | BIG                    | ASI_NUCLEUS_LITTLE   |
| Load/Store/Atomic<br>Alternate with specified ASI<br><i>not</i> ending in <code>_LITTLE</code> | x   | x          | 0       | BIG                    | Specified ASI value from<br>immediate field in opcode or<br>ASI Register |
|  |     |            | 1       | LITTLE                 |  |
| Load/Store/Atomic<br>Alternate with specified ASI<br>ending in <code>_LITTLE</code>            | x   | x          | 0       | LITTLE                 | Specified ASI value from<br>immediate field in opcode or<br>ASI Register |
|  |     |            | 1       | BIG                    |  |
| FLUSH  | 0   | x          | x       | —                      | ASI_PRIMARY_*  |
|  | > 0 | x          | x       | —                      | ASI_NUCLEUS  |

The Context Register used by the D-MMUs and I-MMUs is determined according to TABLE 11-5. The Context Register selection is not affected by the endianness of the access. For a comprehensive list of ASI values in the ASI map, see Chapter 8 “Address Space Identifiers.”

**TABLE 11-5** I-MMU and D-MMU Context Register Usage

| ASI Value  | Context Register                        |
|--|---|
| ASI_*NUCLEUS* (any ASI name containing the string “NUCLEUS”)     | Nucleus (0000 <sub>16</sub> hard-wired) |
| ASI_*PRIMARY* (any ASI name containing the string “PRIMARY”)     | Primary                                 |
| ASI_*SECONDARY* (any ASI name containing the string “SECONDARY”) | Secondary                               |
| All other ASI values   | (Not applicable, no translation)        |

---

## 11.7 Reset, Disable, and RED\_state Behavior

During global reset of the processor, the following statements apply:

- No change occurs in any block of the D-MMU.
- No change occurs in the datapath or TLB blocks of the I-MMU.
- The I-MMU resets its internal state machine to normal (nonsuspended) operation.
- The I-MMU and D-MMU Enable bits in the DCU Control Register are set to 0.

When the processor enters RED\_state, the following statement applies:

- The I-MMU and D-MMU Enable bits in the DCU Control Register are set to 0.

Either of the MMUs is defined to be *disabled* when its respective MMU Enable bit equals 0 or, for the I-MMU only, whenever the processor is in RED\_state. The D-MMU is enabled or disabled solely by the state of the D-MMU Enable bit.

When the D-MMU is disabled:

- The D-MMU passes all addresses through without translation ("bypasses" them); each address is truncated to the size of a physical address on the implementation, behaving as if the ASI\_PHYS\_\* ASI had been used for the access.
- The processor behaves as if the TTE bits were set as:
  - $TTE.IE \leftarrow 0$
  - $TTE.P \leftarrow 0$
  - $TTE.W \leftarrow 1$
  - $TTE.NFO \leftarrow 0$
  - If DCUCR.CP and DCUCR.CV are implemented:
    - $TTE.CP \leftarrow DCUCR.CP$
    - $TTE.CV \leftarrow DCUCR.CV$
    - $TTE.E \leftarrow \text{not } DCUCR.CP$
  - If DCUCR.CP and DCUCR.CV are not implemented:
    - $TTE.E \leftarrow \text{not } TTE.CP$

However, if a bypass ASI (ASI\_PHYS\_\*) is used while the D-MMU is disabled, the bypass operation behaves as it does when the D-MMU is enabled; that is, the access is processed with the E, CP, and CV bits as specified by the bypass ASI (see TABLE 11-26 on page 11-294).

When the I-MMU is disabled, it truncates all instruction accesses to the physical address size (43 bits) and passes the physically cacheable bit (Data Cache Unit Control Register CP bit) to the cache system. The access does not generate an *instruction\_access\_exception* trap.

---

**Note** – While the D-MMU is disabled and the default CV bit in the Data Cache Unit Control Register is set to 0, data in the D-cache can be accessed only through load and store alternates to the internal D-cache access ASI. Normal loads and stores bypass the D-cache. Data in the D-cache cannot be accessed by load or store alternates that use ASI\_PHYS\_\*. Other caches are physically indexed or are still accessible.

---

When disabled, both the I-MMU and D-MMU correctly perform all LDXA and STXA operations to internal registers, and traps are signalled just as if the MMU were enabled. For instance, if a non-faulting load is issued when the D-MMU is disabled and DCUCR.CP is set to 0 if the implementation has the bit, then the D-MMU signals a *data\_access\_exception* trap (FT = 02<sub>16</sub>), since E is set to 1.

---

**Note** – A reset of the TLB is not performed by a chip reset or by entry into RED\_state. Before the MMUs are enabled, the operating system software must explicitly write each entry with either a valid TLB entry or an entry with the valid bit set to 0. The operation of the I-MMU or D-MMU in enabled mode is undefined if the TLB valid bits have not been set explicitly beforehand.

---

## 11.8 SPARC V9 “MMU Requirements” Annex

The MMU complies completely with the SPARC V9 “MMU Requirements” Annex. TABLE 11-6 shows how various protection modes can be achieved, if necessary, through the presence or absence of a translation in the I-MMU or D-MMU. Note that this behavior requires specialized TLB miss handler code to guarantee these conditions.

**TABLE 11-6** MMU SPARC V9 Annex G Protection Mode Compliance

| Condition    |              |                        | Resultant Protection Mode |
|--------------|--------------|------------------------|---------------------------|
| TTE in D-MMU | TTE in I-MMU | Writable Attribute Bit |                           |
| Yes          | No           | 0                      | Read-only                 |
| No           | Yes          | N/A                    | Execute-only              |
| Yes          | No           | 1                      | Read/Write                |
| Yes          | Yes          | 0                      | Read-only/Execute         |
| Yes          | Yes          | 1                      | Read/Write/Execute        |

## 11.9 Data Translation Lookaside Buffer

In the D-MMU, there are two TLBs (dt512\_0 and dt\_512\_1) each of which have 512-entry two-way associative which can be programmed to hold unlocked 8 KB, 64 KB, 512 KB, and 4 MB pages for translations. Each TLB can hold any of the 4 page sizes, but are programmed to only one page size at any given time. Each TLB can be programmed to either the same or different page sizes. Also, a 16-entry fully associative TLB (dt16) is used for 8 KB, 64 KB, and 4 MB locked and unlocked pages. This TLB can hold any combination of page sizes that are locked/unlocked at a time.

---

**Note** – The UltraSPARC III Cu processor’s dt16 can support *unlocked* 8K pages. This is necessary since if dt512\_0 and dt512\_1 were programmed to non-8K page size, then a D-TLB fill of unlocked 8K page will not get dropped.

---




---

**Note** – When both large D-TLBs are configured with the same page size, then they behave like a single D-TLB with 1024-entry 4-way associative (256 entries per way).

---

Each dt512’s page size (PgSz) is programmable independently, one PgSz per context (Primary/Secondary/ Nucleus). Kernel can set the PgSz fields in ASI\_PRIMARY\_CONTEXT\_REG and ASI\_SECONDARY\_CONTEXT\_REG as illustrated in FIGURE 11-6 and FIGURE 11-7, respectively, and described in TABLE 11-7 and TABLE 11-8, respectively.



**FIGURE 11-6** D-MMU Primary Context Register

**TABLE 11-7** D-MMU Primary Context Register

| Bit   | Field   | Description  |
|-------|---------|--|
| 63:61 | N_pgsz0 | Nucleus context’s page size at the first large D-TLB (dt512_0).  |
| 60:58 | N_pgsz1 | Nucleus context’s page size at the second large D-TLB (dt512_1). |
| 57:22 | —       | <i>Reserved.</i>   |
| 21:19 | P_pgsz1 | Primary context’s page size at the second large D-TLB (dt512_0). |

**TABLE 11-7** D-MMU Primary Context Register (*Continued*)

| Bit   | Field    | Description   |
|-------|----------|---|
| 18:16 | P_pgsz0  | Primary context's page size at the first large D-TLB (dt512_1). |
| 15:13 | —        | <i>Reserved.</i>  |
| 12:0  | PContext | Context identifier for the primary address space.               |

|    |         |         |       |          |
|----|---------|---------|-------|----------|
| —  | S_pgsz1 | S_pgsz0 | —     | SContext |
| 63 | 22 21   | 19 18   | 16 15 | 13 12    |
|    |         |         |       | 0        |

**FIGURE 11-7** D-MMU Secondary Context Register

**TABLE 11-8** D-MMU Primary Context Register

| Bit   | Field    | Description  |
|-------|----------|--|
| 63:22 | —        | <i>Reserved.</i>   |
| 21:19 | S_pgsz1  | Secondary context's page size at the second large D-TLB (dt512_1). |
| 18:16 | S_pgsz0  | Secondary context's page size at the first large D-TLB (dt512_0).  |
| 15:13 | —        | <i>Reserved.</i>   |
| 12:0  | SContext | Context identifier for the secondary address space.                |

The following is the page size bit encoding (most significant bit is reserved to 0):

- 000 = 8 KB
- 001 = 64 KB
- 010 = 512 KB
- 011 = 4 MB

## 11.9.1 D-TLB Access Operation

When a memory access is issued, its VA, Context, and Page Size are presented to the D-MMU. All 3 D-TLBs (dt512\_0, dt512\_1, and dt16) are accessed in parallel. The fully associative dt16 only needs VA and Context to CAM-match and output an entry (1 out of 16). Proper VA bits are compared based on the page size bits of each dt16 entry (fast 3-bit encoding is used to define 8K, 64K, 512K, and 4M).

Since dt512's are not a fully associative structure, indexing the dt512 array requires knowledge of page size to properly select VA bits as index. The 2-way dt512\_0 and dt512\_1 needs Pgsz to mux select proper VA bits to index the dt512 arrays, as shown below:

- 8 KB page: index = VA[20:13]
- 64 KB page: index = VA[23:16]
- 512 KB page: index = VA[26:19]
- 4 MB page: index = VA[29:22]

Context bits are used later after the indexed entry comes out of each array bank/way, to qualify the context hit.

Three possible Context numbers are active in the CPU: Primary (PContext field in ASI\_PRIMARY\_CONTEXT\_REG), Secondary (SContext field in ASI\_SECONDARY\_CONTEXT\_REG), and Nucleus (default to Context = 0). The Context Register to send to D-MMU is determined based on the load/store's ASI encoding of Primary/Secondary/Nucleus.

Since all 3 D-TLBs are being accessed in parallel, software must guarantee that there are no duplicate (stale) entry hits. Most of this responsibility lies in the software (operating system) with the hardware providing some assistance to support full software control. A set of rules on D-TLB replacement, demap and context switch must be followed to maintain consistent and correct behavior.

## 11.9.2 Same Page Size on Both dt512\_0 and dt512\_1

When both dt512's are programmed to have identical page size, the behavior is a "single" 4-way 1024-entry dt512. During TTE fill, if there is no invalid TLB entry to take, then dt512 selection and way selection are determined based on a new 10-bit LFSR (pseudo random generator, the same one used on 2-way L2-cache). For 4-way pseudo random selection, LFSR[1:0] bits (two least significant bits of LFSR) is described in TABLE 11-9 will be used.

**TABLE 11-9** Four-way Pseudo Random Selection

| LFSR[1:0] | TTE Fill to:   |
|-----------|----------------|
| 00        | dt512_0, way 0 |
| 01        | dt512_0, way 1 |
| 10        | dt512_1, way 0 |
| 11        | dt512_1, way 1 |

---

**Note** – All LFSRs (in D-cache, I-cache, W-cache, L2-cache, and TLBs) get initialized on Power-on Reset (Hard\_POR) and System Reset (Soft\_POR).

---

This single LFSR is also shared among both dt512\_0 and dt512\_1 when they have different page sizes. The least significant bit (LFSR[0]) is used for entry replacement. It selects the same bank of both dt512's, but only one dt512's write-enable is eventually asserted at TTE fill time.

Demap Context is needed when the *same* context changes Page Size. During context-in, if the operating system decides to change any Page Size setting of dt512\_0 or dt512\_1 *differently* from the last context-out of the *same* Context (e.g., was both 8K at context-out, now 8K and 4M at context-in), then the operating system will perform Demap Context operation first. This avoids remnant entries in dt512\_0 or dt512\_1, which could cause duplicate, possibly stale, hit.


### 11.9.3 D-TLB Automatic Replacement

A D-TLB miss fast trap handler utilizes the automatic (hardware) replacement write using store ASI\_DTLB\_DATA\_IN\_REG. Section 11.9.4 will describe the D-TLB direct write using store ASI\_DTLB\_DATA\_ACCESS\_REG, which is the kernel uses for initialization (e.g., in OBP) and diagnostic.

When D-TLB miss, or DATA\_ACCESS\_EXCEPTION, or FAST\_DATA\_ACCESS\_PROTECTION is detected, hardware automatically saves the missing VA and context to the Tag Access Register (ASI\_DMMU\_TAG\_ACCESS). To ease indexing of the dt512's when the TTE data is presented (via STXA ASI\_DTLB\_DATA\_IN\_REG), the missing page size information of dt512\_0 and dt512\_1 is captured into a new Extension Register, called ASI\_DMMU\_TAG\_ACCESS\_EXT.

FIGURE 11-8 illustrates and TABLE 11-10 describes the format of the Tag Access Extension Register for saving page sizes

ASI 0x58, VA<63:0> == 0x60  
 Name: ASI\_DMMU\_TAG\_ACCESS\_EXT  
 Access: RW



**FIGURE 11-8** Tag Access Extension Register Format for Saving Page Sizes Information



**TABLE 11-10** Tag Access Extension Register for Saving Page Sizes Information

| Bit(s) | Field | Description   |
|--------|-------|---|
| 63:22  | —     | <i>Reserved.</i>  |
| 21:19  | pgsz1 | Page size of D-TLB miss' context (Primary/Secondary/Nucleus) in the second large D-TLB (dt512_0). |
| 18:16  | pgsz0 | Page size of D-TLB miss' context (Primary/Secondary/Nucleus) in the first large D-TLB (dt512_1).  |
| 15:0   | —     | <i>Reserved.</i>  |

---

**Note** – Bit 21 and 18 are hardwired to zero since the most significant bit of the page size fields in Primary and Secondary registers are reserved to 0.

---

---

**Note** – With the saved page sizes, hardware pre-computes in the background the index to dt512\_0 and dt512\_1 for TTE fill. When the TTE data arrives, only one write enable to dt512\_0, dt512\_1, and dt16 will get activated.

---

CODE EXAMPLE 11-2 shows the hardware D-TLB replacement algorithm.

---

**Note** – PgSz0 below is ASI\_DMMU\_TAG\_ACCESS\_EXT[18:16] bits; PgSz1 below is ASI\_DMMU\_TAG\_ACCESS\_EXT[21:19] bits.

---

**CODE EXAMPLE 11-2** D-TLB Replacement Algorithm

```
if (TTE to fill is a locked page, i.e., L bit is set) {
    fill TTE to dt16;
} else if (both dt512's have same page size, PgSz0 == PgSz1) {
    if (TTE's Size != PgSz0) {
        fill TTE to dt16;
    } else {
        if (one of the 4 same-index entries is invalid) {
            fill TTE to an invalid entry with selection order of
            (dt512_0 way0, dt512_0 way1, dt512_1 way0, dt512_1 way1)
        } else {
            case (LFSR[1:0]) {
                00: fill TTE to dt512_0 way0;
                01: fill TTE to dt512_0 way1;
                10: fill TTE to dt512_1 way0;
                11: fill TTE to dt512_1 way1;
            }
        }
    }
}
```

**CODE EXAMPLE 11-2** D-TLB Replacement Algorithm (Continued)

```
}
}
} else {
    if (TTE's Size == PgSz0) {
        if (one of the 2 same-index entries is invalid) {
            fill TTE to an invalid entry with selection order of
            (dt512_0 way0, dt512_0 way1)
        } else {
            case (LFSR[0]) {
                0: fill TTE to dt512_0 way0;
                1: fdt512ill TTE to dt512_0 way1;
            }
        }
    } else if (TTE's Size == PgSz1) {
        if (one of the 2 same-index entries is invalid) {
            fill TTE to an invalid entry with selection order of
            (dt512_1 way0, dt512_1 way1)
        } else {
            case (LFSR[0]) {
                0: fill TTE to dt512_1 way0;
                1: fill TTE to dt512_1 way1;
            }
        }
    } else {
        fill TTE to dt16;
    }
}
```

## 11.9.4 D-TLB Directed Data Read/Write

Each D-TLB can be directly written using the store `ASI_DTLB_DATA_ACCESS_REG` instruction. The kernel typically uses this method for initialization (e.g., in OBP) and diagnostic.

Due to the addition of a second large D-TLB, in the UltraSPARC III Cu processor, a new encoding value was added to the VA of `ASI_DTLB_DATA_ACCESS_REG` (0x5d), as listed in TABLE 11-11.

**TABLE 11-11** D-TLB Access Number

| TLB # | TLB Type   | Number of Entries |
|-------|--|-------------------|
| 0     | Small fully associative, all 4 page sizes (8K, 64K, 512K, 4M) and locked pages | 16                |



## 11.9.6 Demap Operation

For Demap Page in the large D-TLBs, the page size used to index the D-TLBs is derived based on the Context bits (Primary/Secondary/Nucleus). Hardware will automatically select proper P<sub>G</sub>Sz bits based on the “Context” field (Primary/Secondary/Nucleus) defined in ASI\_DMMU\_DEMAP (ASI 0x5f). These two P<sub>G</sub>Sz fields are used to properly index dt512\_0 and dt512\_1.

Demap Context and Demap All operations with dt512\_0/dt512\_1/dt16 are straight forward, except that one additional dt512 is demap in parallel.

---

**Note** – When global pages are used (G = 1), any active pages in a dt512 must have the *same* Page Size (but dt512\_0’s P<sub>G</sub>Sz and dt512\_1’s P<sub>G</sub>Sz can be different). When pages with G = 1 in a dt512 have variety of page sizes, the dt512 cannot index and locate the page correctly when trying to match the VA tag without the context number as a qualifier. Proper demap must be done, as well as proper Page Size settings in PContext and SContext registers.

---

## 11.9.7 D-TLB Access Summary

TABLE 11-12 lists the D-MMU TLB access summary.

**TABLE 11-12** D-MMU TLB Access Summary

| Software Operation |             | Effect on MMU Physical Registers                         |  |                   |                   |                   |
|--------------------|-------------|--|--|-------------------|-------------------|-------------------|
| Load/Store         | Register    | TLB Tag Array  | TLB Data Array   | Tag Access        | SFSR              | SFAR              |
| Load               | Tag Read    | Contents returned. From entry specified by LDXA's access | No effect  | No effect         | No effect         | No effect         |
|                    | Tag Access  | No effect  | No effect  | Contents returned | No effect         | No effect         |
|                    | Data In     | Trap with <i>data_access_exception</i>                   |  |                   |                   |                   |
|                    | Data Access | No effect  | Contents returned. From entry specified by LDXA's access | No effect         | No effect         | No effect         |
|                    | SFSR        | No effect  | No effect  | No effect         | Contents returned | No effect         |
|                    | SFAR        | No effect  | No effect  | No effect         | No effect         | Contents returned |

**TABLE 11-12** D-MMU TLB Access Summary (*Continued*)

| Software Operation |             | Effect on MMU Physical Registers  |  |                                       |  |  |
|--------------------|-------------|---|--|---------------------------------------|--|--|
| Load/Store         | Register    | TLB Tag Array   | TLB Data Array   | Tag Access                            | SFSR   | SFAR   |
| Store              | Tag Read    | Trap with <i>data_access_exception</i>  |  |                                       |  |  |
|                    | Tag Access  | No effect   | No effect  | Written with store data               | No effect  | No effect  |
|                    | Data In     | TLB entry determined by replacement policy written with contents of Tag Access Register | TLB entry determined by replacement policy written with store data | No effect                             | No effect  | No effect  |
|                    | Data Access | TLB entry specified by STXA address written with contents of Tag Access Register        | TLB entry specified by STXA address written with store data        | No effect                             | No effect  | No effect  |
|                    | SFSR        | No effect   | No effect  | No effect                             | Written with store data  | No effect  |
|                    | SFAR        | No effect   | No effect  | No effect                             | No effect  | Written with store data                              |
| TLB miss           |             | No effect   | No effect  | Written with VA and context of access | Written with fault status of faulting instruction and page sizes at faulting context for two 2-way set associative TLB | Written with virtual address of faulting instruction |

## 11.9.8 D-MMU Operation Summary

The behavior of the D-MMU is summarized in TABLE 11-13 on page 11-276. In each case and for all conditions, the behavior of each MMU is given by one of the following abbreviations:

| Abbreviation | Meaning   |
|--------------|---|
| OK           | Normal translation                                |
| Dmiss        | <i>fast_data_access_MMU_miss</i> exception        |
| Dexc         | <i>data_access_exception</i> exception            |
| Dprot        | <i>fast_data_access_protection</i> exception      |
| Imiss        | <i>fast_instruction_access_MMU_miss</i> exception |
| Iexc         | <i>instruction_access_exception</i> exception     |

The ASI is indicated by one the following abbreviations:

| 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*   |
| U_SEC        | ASI_AS_IF_USER_SECONDARY*   |
| BYPASS       | ASI_PHYS_* and also other ASIs that require the MMU to perform a bypass operation (such as D-cache access). |

---

**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.

The following cases are not covered in TABLE 11-13.

- Invalid ASIs or ASIs that have no meaning for the opcodes listed; for example, ASI\_PRIMARY\_NOFAULT for a store or atomic.
- Access to internal registers other than LDXA, LDDFA, STXA, or STDFA. See Chapter 8 “Address Space Identifiers.” The MMU signals a *data\_access\_exception* trap (FT = 08<sub>16</sub>) for these cases.
- Attempted access using a restricted ASI in non-privileged mode. The MMU signals a *privileged\_action* exception for this case.
- An atomic instruction (including 128-bit atomic load) issued to a memory address marked non-cacheable in a physical cache; that is, with the CP bit set to 0, including cases in which the D-MMU is disabled. The MMU signals a *data\_access\_exception* trap (FT = 04<sub>16</sub>) for this case.

- A data access with an ASI other than “(PRIMARY , SECONDARY )\_NO\_FAULT (\_LITTLE)” to a page marked with the NFO bit. The MMU signals a *data\_access\_exception* trap (FT = 10<sub>16</sub>) for this case.

**TABLE 11-13** D-MMU Table of Operations for Normal ASIs

| Condition       |           | ASI             | W | Behavior                 |                |                |                |                |
|-----------------|-----------|-----------------|---|--------------------------|----------------|----------------|----------------|----------------|
| Opcode          | PRIV mode |                 |   | TLB Miss                 | E = 0<br>P = 0 | E = 0<br>P = 1 | E = 1<br>P = 0 | E = 1<br>P = 1 |
| Load            | 0         | PRIM, SEC       | x | Dmiss                    | OK             | Dexc           | OK             | Dexc           |
|                 |           | PRIM_NF, SEC_NF | x | Dmiss                    | OK             | Dexc           | Dexc           | Dexc           |
|                 | 1         | PRIM, SEC, NUC  | x | Dmiss                    | OK             | OK             | OK             | OK             |
|                 |           | PRIM_NF, SEC_NF | x | Dmiss                    | OK             | OK             | Dexc           | Dexc           |
|                 |           | U_PRIM, U_SEC   | x | Dmiss                    | OK             | Dexc           | OK             | Dexc           |
| Store or Atomic | 0         | PRIM, SEC       | 0 | Dmiss                    | Dprot          | Dexc           | Dprot          | Dexc           |
|                 |           |                 | 1 | Dmiss                    | OK             | Dexc           | OK             | Dexc           |
|                 | 1         | PRIM, SEC, NUC  | 0 | Dmiss                    | Dprot          | Dprot          | Dprot          | Dprot          |
|                 |           |                 | 1 | Dmiss                    | OK             | OK             | OK             | OK             |
|                 |           | U_PRIM, U_SEC   | 0 | Dmiss                    | Dprot          | Dexc           | Dprot          | Dexc           |
|                 |           |                 | 1 | Dmiss                    | OK             | Dexc           | OK             | Dexc           |
| FLUSH           | 0         |                 | x | Dmiss                    | OK             | Dexc           | OK             | Dexc           |
|                 | 1         |                 | x | Dmiss                    | OK             | OK             | Dexc           | Desc           |
| x               | 0         | BYPASS          | x | <i>privileged_action</i> |                |                |                |                |
| x               | 1         | BYPASS          | x | Bypass                   |                |                |                |                |

## 11.9.9 Internal Registers and ASI Operations

In this section, how to access MMU registers is first described and following is the registers themselves, as follows:

- Context Registers
- D-MMU TLB Tag Access Registers
- D-TLB Data In, Data Access, and Tag Read Registers
- dTSB Tag Target Registers
- dTSB Base Registers
- dTSB Extension Registers
- dTSB 8 KB and 64 KB Pointer and Direct Pointer Registers
- Data Synchronous Fault Status Registers (D-SFSR)
- MMU Data Synchronous Fault Address Register

Following the register descriptions, the Data demap operation is described.



## 11.9.9.1 Accessing MMU Registers

All internal MMU registers can be accessed directly by the processor through defined ASIs. Several of the registers have been assigned their own ASI because these registers are crucial to the speed of the TLB miss handler. Allowing the use of %g0 for the address reduces the number of instructions required to perform the access to the alternate space (by eliminating address formation).

For instance, to facilitate an access to the D-cache, the MMU performs a bypass operation.

---

**Caution** – A store to an MMU register requires a MEMBAR #Sync, FLUSH, DONE, or RETRY before the point that the effect must be visible to load/store/atomic accesses. A FLUSH, DONE, or RETRY is needed before the point that the effect must be visible to instruction accesses, that is, MEMBAR #Sync is not sufficient. In either case, one of these instructions must be executed before the next non-internal store or load of any type and on or before the delay slot of a delayed-control transfer instruction of any type. This action is necessary to avoid data corruption.

---

If the low-order three bits of the VA are nonzero in an LDXA/STXA to or from these registers, then 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 can cause a *data\_access\_exception* trap (FT = 08<sub>16</sub>). (The hardware detects VA violations in only an unspecified lower portion of the virtual address.) TABLE 11-14 describes MMU registers.

**TABLE 11-14** MMU Internal Registers and ASI Operations

| D-MMU ASI        | VA<63:0>         | Access     | Register or Operation Name                       |
|------------------|------------------|------------|--|
| 58 <sub>16</sub> | 0 <sub>16</sub>  | Read-only  | dTSB Tag Target Registers                        |
| 58 <sub>16</sub> | 8 <sub>16</sub>  | Read/Write | Primary Context Register                         |
| 58 <sub>16</sub> | 10 <sub>16</sub> | Read/Write | Secondary Context Register                       |
| 58 <sub>16</sub> | 18 <sub>16</sub> | Read/Write | Data Synchronous Fault Status Registers (D-SFSR) |
| 58 <sub>16</sub> | 20 <sub>16</sub> | Read-only  | Data Synchronous Fault Address Register (D-SFAR) |
| 58 <sub>16</sub> | 28 <sub>16</sub> | Read/Write | dTSB Base Registers                              |
| 58 <sub>16</sub> | 30 <sub>16</sub> | Read/Write | D-TLB Tag Access Registers                       |
| 58 <sub>16</sub> | 38 <sub>16</sub> | Read/Write | Virtual Watchpoint Address                       |
| 58 <sub>16</sub> | 40 <sub>16</sub> | Read/Write | Physical Watchpoint Address                      |
| 58 <sub>16</sub> | 48 <sub>16</sub> | Read/Write | dTSB Primary Extension Registers                 |
| 58 <sub>16</sub> | 50 <sub>16</sub> | Read/Write | dTSB Secondary Extension Register                |
| 58 <sub>16</sub> | 58 <sub>16</sub> | Read/Write | dTSB Nucleus Extension Registers                 |
| 59 <sub>16</sub> | 0 <sub>16</sub>  | Read-only  | dTSB 8 KB Pointer Registers                      |

**TABLE 11-14** MMU Internal Registers and ASI Operations (Continued)

| D-MMU ASI        | VA<63:0>                                  | Access     | Register or Operation Name    |
|------------------|---|------------|-------------------------------|
| 5A <sub>16</sub> | 0 <sub>16</sub>                           | Read-only  | dTSB 64 KB Pointer Registers  |
| 5B <sub>16</sub> | 0 <sub>16</sub>                           | Read-only  | dTSB Direct Pointer Register  |
| 5C <sub>16</sub> | 0 <sub>16</sub>                           | Write-only | D-TLB Data In Registers       |
| 5D <sub>16</sub> | 0 <sub>16</sub> – 20FF8 <sub>16</sub>     | Read/Write | D-TLB Data Access Registers   |
| 5D <sub>16</sub> | 40000 <sub>16</sub> – 60FF8 <sub>16</sub> | Read/Write | D-TLB CAM Diagnostic Register |
| 5E <sub>16</sub> | 0 <sub>16</sub> – 20FF8 <sub>16</sub>     | Read-only  | D-TLB Tag Read Registers      |
| 5F <sub>16</sub> | See 11.9.9.14                             | Write-only | D-MMU Demap Operations        |


---

**Note** – TSB\_Hash field is implemented in TSB Extension Register.

---


## 11.9.9.2 Context Registers

The Primary Context Register is shared by the I-MMU and the D-MMU and resides in the MMU. The Primary Context Register is illustrated in FIGURE 11-9; PContext is the context identifier for the primary address space.




**FIGURE 11-9** D-MMU Primary Context Register

The Secondary Context Register is illustrated in FIGURE 11-10; SContext is the context identifier for the secondary address space.



**FIGURE 11-10** D-MMU Secondary Context Register

The Nucleus Context Register is hardwired to zero, as illustrated in FIGURE 11-11.



**FIGURE 11-11** D-MMU Nucleus Context Register

### 11.9.9.3 Data MMU TLB Tag Access Registers

In each MMU, the Tag Access Register is used as a temporary buffer for writing the TLB entry tag information. The Tag Access Register holds the tag portion, and the Data In or Data Access Register holds the data being accessed.

The Tag Access Register can be updated during either of the following operations:

- When the MMU signals a trap due to a miss, exception, or protection
 

The MMU hardware automatically writes the missing VA and the appropriate context into the Tag Access Register to facilitate formation of the TSB Tag Target Register. One exception is that after a *data\_access\_exception*, the contents of the Context field of the D-MMU Tag Access Register are undefined.
- An ASI write to the Tag Access Register
 

Before an ASI store to the TLB Data Access Registers, the operating system must set the Tag Access Register to the values desired in the TLB entry. Note that an ASI store to the TLB Data In Register for automatic replacement also uses the Tag Access Register, but typically the value written into the Tag Access Register by the MMU hardware is appropriate.

---


**Note** – Any update to the Tag Access Registers immediately affects the data that are returned from subsequent reads of the TSB Tag Target and TSB Pointer Registers.

---

The TLB Tag Access Register fields are defined in TABLE 11-15 and illustrated in FIGURE 11-12.

**TABLE 11-15** D-MMU TLB Tag Access Registers

| Bit(s) | Field   | Type | Description   |
|--------|---------|------|---|
| 63:13  | VA      | RW   | The 51-bit virtual page number.   |
| 12:0   | Context | RW   | The 13-bit context identifier. This field reads 0 when there is no associated context with the access. Its contents in the D-MMU are undefined after a <i>data_access_exception</i> . |



**FIGURE 11-12** D-MMU TLB Tag Access Registers Format

---

**Caution** – When the D-MMU causes a trap due to a protection violation or other exception, software should use the context number from `D-SFSR.CT` instead of the `Context` field of the D-TLB Tag Access Register.

---

### 11.9.9.4 Data TLB Data In, Data Access, and Tag Read Registers

Access to the TLB is complicated because of the need to provide an atomic write of a TLB entry data item (tag and data) that is larger than 64 bits, the need to replace entries automatically through the TLB entry replacement algorithm as well as to provide direct diagnostic access, and the need for hardware assist in the TLB miss handler.

TABLE 11-2 on page 11-259 shows when loads and stores update the Tag Access Registers.

TABLE 11-16 shows how the Tag Read, Tag Access, Data In, and Data Access Registers interact to provide atomic reads and writes to the TLBs.

**TABLE 11-16** MMU TLB Access Summary

| Software Operation |             | Effect on MMU Physical Registers  |  |                                       |
|--------------------|-------------|---|--|---------------------------------------|
| Load/Store         | Register    | TLB Tag Array   | TLB Data Array   | Tag Access Register                   |
| Load               | Tag Read    | Contents returned. Entry specified by <i>STXA</i> 's access.                            | No effect  | No effect                             |
|                    | Tag Access  | No effect   | No effect  | Contents returned                     |
|                    | Data In     | Trap with <i>data_access_exception</i> .  |  |                                       |
|                    | Data Access | No effect   | Contents returned. Entry specified by <i>STXA</i> 's access.       | No effect                             |
| Store              | Tag Read    | Trap with <i>data_access_exception</i> .  |  |                                       |
|                    | Tag Access  | No effect   | No effect  | Written with store data               |
|                    | Data In     | TLB entry determined by replacement policy written with contents of Tag Access Register | TLB entry determined by replacement policy written with store data | No effect                             |
|                    | Data Access | TLB entry specified by <i>STXA</i> address written with contents of Tag Access Register | TLB entry specified by <i>STXA</i> address written with store data | No effect                             |
| TLB miss           |             | No effect   | No effect  | Written with VA and context of access |

An ASI load from the TLB Tag Read Register initiates an internal read of the tag portion of the specified TLB entry.

## 11.9.9.5 Data In and Data Access Registers

The Data In and Data Access Registers are the means of reading and writing the TLB for all operations. The TLB Data In Register is used for TLB miss handler automatic replacement writes. The TLB Data Access Register is used for operating system and diagnostic directed writes (writes to a specific TLB entry).

An ASI load from the TLB Data Access Register initiates an internal read of the data portion of the specified TLB entry.

ASI loads from the TLB Data In Register are not supported.

An ASI store to the TLB Data In Register initiates an automatic atomic replacement of the TLB entry pointed to by an internal register that is updated by a proprietary replacement algorithm. The TLB data and tag are formed as in the case of an ASI store to the TLB Data Access Register.

---

**Caution** – Stores to the Data In Register are not guaranteed to replace the previous TLB entry, causing a fault. In particular, to change an entry’s attribute bits, software must explicitly demap the old entry before writing the new entry; otherwise, a multiple match error condition can result.

---


Both the TLB Data In Register and the TLB Data Access Register use the TTE format shown in FIGURE 11-3 on page 11-248. Refer to the description of the TTE data in Section 11.2 “Translation Table Entry” on page 11-248 for a complete description of the data fields.

Writes to the TLB Data In Register require the virtual address to be set to 0.

The format of the TLB Data Access Register virtual address is illustrated in FIGURE 11-13 and described in TABLE 11-17.

**TABLE 11-17** TLB Data Access Register

| Bit(s) | Field  | Type                  | Description   |       |          |              |   |  |    |   |                   |                       |   |                   |                       |
|--------|--|-----------------------|---|-------|----------|--------------|---|--|----|---|-------------------|-----------------------|---|-------------------|-----------------------|
| 63:19  | —  | —                     | <i>Reserved.</i>  |       |          |              |   |  |    |   |                   |                       |   |                   |                       |
| 18     | 0  | —                     | Set to zero.  |       |          |              |   |  |    |   |                   |                       |   |                   |                       |
| 17:16  | TLB #  | RW                    | The TLB to access, as defined below. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>TLB #</th> <th>TLB Type</th> <th># of Entries</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Fully associative 64 KB, 512 KB, and 4 MB page size and locked pages</td> <td>16</td> </tr> <tr> <td>2</td> <td>2-way associative</td> <td>512 (D-MMU) (dt512-0)</td> </tr> <tr> <td>3</td> <td>2-way associative</td> <td>512 (D-MMU) (dt512-1)</td> </tr> </tbody> </table> | TLB # | TLB Type | # of Entries | 0 | Fully associative 64 KB, 512 KB, and 4 MB page size and locked pages | 16 | 2 | 2-way associative | 512 (D-MMU) (dt512-0) | 3 | 2-way associative | 512 (D-MMU) (dt512-1) |
| TLB #  | TLB Type   | # of Entries          |   |       |          |              |   |  |    |   |                   |                       |   |                   |                       |
| 0      | Fully associative 64 KB, 512 KB, and 4 MB page size and locked pages | 16                    |   |       |          |              |   |  |    |   |                   |                       |   |                   |                       |
| 2      | 2-way associative  | 512 (D-MMU) (dt512-0) |   |       |          |              |   |  |    |   |                   |                       |   |                   |                       |
| 3      | 2-way associative  | 512 (D-MMU) (dt512-1) |   |       |          |              |   |  |    |   |                   |                       |   |                   |                       |
| 15:12  | —  | —                     | <i>Reserved.</i>  |       |          |              |   |  |    |   |                   |                       |   |                   |                       |
| 11:3   | TLB Entry  | RW                    | The TLB entry number to be accessed, in the range 0 – 511. Not all TLBs will have all 512 entries. All TLBs regardless of size are accessed from 0 to $N - 1$ , where $N$ is the number of entries in the TLB.  |       |          |              |   |  |    |   |                   |                       |   |                   |                       |
| 2:0    | 0  | —                     | Set to zero.  |       |          |              |   |  |    |   |                   |                       |   |                   |                       |




**FIGURE 11-13** D-MMU TLB Data Access Address Format

### 11.9.9.6 Data MMU TLB Tag Read Register

The format for the Tag Read Register virtual address is described in TABLE 11-18 and illustrated in FIGURE 11-14.

**TABLE 11-18** D-MMU TLB Tag Read Registers

| Bit(s) | Field        | Type | Description  |
|--------|--------------|------|--|
| 63:13  | VA           | RW   | The 51-bit virtual page number. In the fully associative TLB, page offset bits for larger page sizes are stored in the TLB; that is, VA<15:13>, VA<18:13>, and VA<21:13> for 64 KB, 512 KB, and 4 MB pages, respectively. These values are ignored during normal translation. When read, an implementation will return either 0 or the value previously written to them. |
| 12:0   | Data Context | RW   | The 13-bit context identifier.   |




**FIGURE 11-14** D-MMU TLB Tag Read Registers

### 11.9.9.7 Data MMU TLB Tag Access Register

An ASI store to the TLB Data Access or Data In Register initiates an internal atomic write to the specified TLB entry. The TLB entry data is obtained from the store data, and the TLB entry tag is obtained from the current contents of the TLB Tag Access Register.

### 11.9.9.8 Data TSB Tag Target Registers

The Data Translation Storage Buffer (dTSB) Tag Target Registers are simply bit-shifted versions of the data stored in the Data Tag Access Registers. Since the Data Tag Access Register is updated on an D-TLB miss, the Data Tag Target Registers appear to software to be updated on an D-TLB miss. The D-MMU Tag Target Register is described in TABLE 11-19 and illustrated in FIGURE 11-15.



**FIGURE 11-15** MMU Tag Target Registers


**TABLE 11-19** MMU Tag Target Registers

| Bit(s) | Field         | Type | Description   |
|--------|---------------|------|---|
| 63:61  | 000           | —    | Set to 000.   |
| 60:48  | Context<12:0> | RW   | The context associated with the missing virtual address.  |
| 47:42  | —             | —    | <i>Reserved.</i>  |
| 41:0   | VA<63:22>     | RW   | The most significant bits of the missing virtual address. |

### 11.9.9.9 Data TSB Base Registers

The TSB registers provide information for the hardware formation of TSB pointers and tag target, to assist software in quickly handling TLB misses. If the TSB concept is not employed in the software memory management strategy and therefore the Pointer and Tag Access Registers are not used, then the TSB registers need not contain valid data.

The TSB register is illustrated in FIGURE 11-16 and described in TABLE 11-20.



**FIGURE 11-16** MMU Data TSB Registers

**TABLE 11-20** TSB Register Description

| Bit(s) | Field         | Type | Description   |
|--------|---------------|------|---|
| 63:13  | Data TSB_Base | RW   | Provides the base virtual address of the TSB. Software must ensure that the TSB base is aligned on a boundary equal to the size of the TSB or both TSBs in the case of a split TSB. |



TABLE 11-20 TSB Register Description (Continued)

| Bit(s) | Field         | Type | Description  |
|--------|---------------|------|--|
| 12     | Split         | RW   | <p>When Split = 1, the TSB 64 KB pointer address is calculated assuming separate (but abutting and equally sized) TSB regions for the 8 KB and the 64 KB TTEs. In this case, TSB_Size refers to the size of each TSB. The TSB 8 KB pointer address calculation is not affected by the value of the Split bit. When Split = 0, the TSB 64 KB pointer address is calculated assuming that the same lines in the TSB are shared by 8 KB and 64 KB TTEs, called a “common TSB” configuration.</p> <p><b>Caution:</b> In the “common TSB” configuration (TSB.Split = 0), 8 KB and 64 KB page TTEs can conflict unless the TLB miss handler explicitly checks the TTE for page size. Therefore, do not use the common TSB mode in an optimized handler. For example, suppose an 8 KB page at VA = 2000<sub>16</sub> and a 64 KB page at VA = 10000<sub>16</sub> both exist — a legal situation. These both map to the second TSB line (line 1) and have the same VA tag of 0. Therefore, there is no way for the miss handler to distinguish these TTEs by the TTE tag alone, and unless the miss handler checks the TTE data, it may load an incorrect TTE.</p> |
| 11:3   | —             | —    | Reserved.  |
| 2:0    | Data TSB_Size | RW   | <p>The UltraSPARC III Cu processor implements a 3-bit TSB_Size field. The TSB_Size field provides the size of the TSB as follows:</p> <ul style="list-style-type: none"> <li>• The number of entries in the TSB (or each TSB if split) = <math>512 \times 2^{\text{TSB\_Size}}</math>.</li> <li>• The number of entries in the TSB ranges from 512 entries at TSB_Size = 0 (8 KB common TSB, 16 KB split TSB), to 64K entries at TSB_Size = 7 (1 MB common TSB, 2 MB split TSB).</li> </ul> <p><b>Note:</b> Any update to the TSB register immediately affects the data that are returned from later reads of the Tag Target and TSB Pointer Registers.</p>  |

### 11.9.9.10 Data TSB Extension Registers

The TSB Extension Registers provide information for the hardware formation of TSB pointers and tag target, to assist software in handling TLB misses quickly. If the TSB concept is not employed in the software memory management strategy and therefore the pointer and Tag Access Registers are not used, then the TSB Extension Registers need not contain valid data.

The TSB registers are defined as follows in FIGURE 11-17.




FIGURE 11-17 MMU Data TSB Extension Registers

In the UltraSPARC III Cu processor, `TSB_Hash` (bits 11:3 of the Extension Registers) are exclusive-ORed with the calculated TSB offset to provide a “hash” into the TSB. Changing the `TSB_Hash` field on a per-process basis minimizes the collision of TSB entries between different processes.

The register field definitions are the same as for Data TSB Base Registers. The field can be used either as a `TSB_Hash`, which is a representation of the context that generated the TLB miss, or as an extension to the `TSB_size` field, depending on the implementation. In the latter case, TSB pointer generation logic must incorporate the context ID into the process of TSB pointer generation, as described in “TSB Pointer Formation” on page 254.

There are three TSB Extension Registers, one for each of the virtual address spaces (Primary, Secondary, Nucleus); see TABLE 8-4 on page 8-193 for the ASI and VA of each register.


When a D-TLB miss occurs, an appropriate TSB Extension Register is selected and XORed either with the dTSB Register or with context ID, depending on the implementation. The result is then used to form a TSB pointer, as described in “TSB Pointer Formation” on page 254.

### 11.9.9.11 Data TSB 8 KB and 64 KB Pointer and Direct Pointer Registers

The dTSB 8 KB and 64 KB registers are provided as an aid to software in determining the location of the missing or trapping TTE in the software-maintained TSB. The TSB 8 KB and 64 KB Pointer Registers provide the possible locations of the 8 KB and 64 KB TTE, respectively.

As a fine point, the bit that controls selection of 8 KB or 64 KB address formation for the Direct Pointer Register is a state bit in the D-MMU that is updated during a *fast\_data\_access\_protection* exception. It records whether the page that hit in the TLB was a 64 KB page or a non-64 KB page, in which case, 8 KB is assumed.

The registers are illustrated in FIGURE 11-18, where `VA<63:4>` is the full virtual address of the TTE in the TSB, as determined by the MMU hardware, and is described in “Hardware Support for TSB Access” on page 253.



**FIGURE 11-18** D-MMU TSB 8 KB/64 KB Pointer and D-MMU Direct Pointer Register



**TABLE 11-21** D-SFSR Bit Description (Continued)

| Bit(s)     | Field         | Type          | Description   |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
|------------|---------------|---------------|---|------------|---------------|---------------|----|---------|---------|----|----------|-----------|----|---------|---------|----|----------|----------|
| 11:7       | FT            | RW            | <p>Specifies the exact condition that caused the recorded fault, according to TABLE 11-22 following this table. In the D-MMU, the Fault Type field is valid only for <i>data_access_exception</i> faults; there is no ambiguity in all other MMU trap cases. Note that the hardware does not priority-encode the bits set in the fault type (FT) field; that is, multiple bits can be set. In particular, the following ASI stores could set both the 01<sub>16</sub> and 08<sub>16</sub> Fault Type bits (the page is privileged, as well as storing to a read-only ASI):</p> <pre> stda          %g0,  [%g4]ASI_PRIMARY_NO_FAULT stda          %g0,  [%g4]ASI_SECONDARY_NO_FAULT stda          %g0,  [%g4]ASI_PRIMARY_NO_FAULT_LITTLE stda          %g0,  [%g4]ASI_SECONDARY_NO_FAULT_LITTLE                     </pre> |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 6          | E             | RW            | <p>Side-effect bit. Associated with the faulting data access or flush instruction. Set by translating ASI accesses (see Chapter 8 “Address Space Identifiers”) that are mapped by the TLB with the E bit set and bypass ASIs 15<sub>16</sub> and 1D<sub>16</sub>. Other cases that update the SFSR (including bypass or internal ASI accesses) set the E bit to 0.</p>  |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 5:4        | CT            | RW            | <p>Context Register selection, as described below. The context is set to 11<sub>2</sub> when the access does not have a translating ASI.</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>Context ID</th> <th>I-MMU Context</th> <th>D-MMU Context</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>Primary</td> <td>Primary</td> </tr> <tr> <td>01</td> <td>Reserved</td> <td>Secondary</td> </tr> <tr> <td>10</td> <td>Nucleus</td> <td>Nucleus</td> </tr> <tr> <td>11</td> <td>Reserved</td> <td>Reserved</td> </tr> </tbody> </table>   | Context ID | I-MMU Context | D-MMU Context | 00 | Primary | Primary | 01 | Reserved | Secondary | 10 | Nucleus | Nucleus | 11 | Reserved | Reserved |
| Context ID | I-MMU Context | D-MMU Context |   |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 00         | Primary       | Primary       |   |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 01         | Reserved      | Secondary     |   |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 10         | Nucleus       | Nucleus       |   |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 11         | Reserved      | Reserved      |   |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 3          | PR            | RW            | <p>Privilege bit. Set if the faulting access occurred while in privileged mode. This field is valid for all traps in which the FV bit is set.</p>   |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 2          | W             | RW            | <p>Write bit. Set if the faulting access indicated a data write operation (a store or atomic load/store instruction).</p>   |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 1          | OW            | RW            | <p>Overwrite bit. When the MMU detects a fault, the Overwrite bit is set to 1 if the Fault Valid bit has not been cleared from a previous fault; otherwise, it is set to 0.</p>   |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 0          | FV            | RW            | <p>Fault Valid bit. Set when the MMU detects a fault; it is cleared only on an explicit ASI write of 0 to the SFSR. This bit is not set on an MMU miss. Therefore, overwrites of MMU misses cannot be detected.</p> <p>When the Fault Valid bit is not set, the values of the remaining fields in the SFSR and SFAR are undefined for traps other than an MMU miss.</p>   |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |

TABLE 11-22 describes the SFSR fault type field (FT<11:7>).

**TABLE 11-22** MMU Synchronous Fault Status Register FT (Fault Type) Field

| Data | FT[6:0]          | Fault Type  |
|------|------------------|---|
| Data | 01 <sub>16</sub> | Privilege violation.  |
| D    | 02 <sub>16</sub> | Non-faulting load instruction to page marked with E bit. This bit is 0 for internal ASI accesses.                       |
| D    | 04 <sub>16</sub> | Atomic (including 128-bit atomic load) to page marked non-cacheable.  |
| D    | 08 <sub>16</sub> | Illegal LDA/STA ASI value, VA, RW, or size. Does not include cases where 02 <sub>16</sub> and 04 <sub>16</sub> are set. |
| D    | 10 <sub>16</sub> | Access other than non-faulting load to page marked NFO. This bit is 0 for internal ASI accesses.                        |

---

**Note** – A *fast\_data\_access\_MMU\_miss* trap causes the D-SFSR and the D-SFAR to be overwritten without setting either the OW or the FV bits.

---

The SFSR and the Tag Access Registers both maintain state concerning a previous translation causing an exception. The update policy for the SFSR and the Tag Access Registers is shown in TABLE 11-2 on page 11-259.


### 11.9.9.13 Synchronous Fault Addresses

This section describes how the D-MMU obtains a fault address.

#### *D-MMU Fault Address*

The Data Synchronous Fault Address Register contains the virtual memory address of the fault recorded in the D-MMU Synchronous Fault Status Register. The D-SFAR can be thought of as an additional field of the D-SFSR.

The D-SFAR register is illustrated in FIGURE 11-20, where *Fault Address* is the virtual address associated with the translation fault recorded in the D-SFSR; the field is set on an MMU miss fault or when the D-SFSR Fault Valid (FV) bit is set.



**FIGURE 11-20** MMU Data Synchronous Fault Address Register (D-SFAR)

## 11.9.9.14 Data MMU Demap

Demap is an MMU operation, not an MMU register. Demap removes selected entries from the TLBs.

---


**Note** — A store to a D-MMU Register requires a MEMBAR #Sync, FLUSH, DONE, or RETRY before the point that the effect must be visible to load/store/atomic accesses. A FLUSH, DONE, or RETRY is needed before the point that the effect must be visible to instruction accesses, that is, MEMBAR #Sync is not sufficient. In either case, one of these instructions must be executed before the next non-internal store or load of any type and on or before the delay slot of a delayed-control transfer instruction of any type. This action is necessary to avoid data corruption.

---

Three types of demap operations are provided:

- **Demap page** — Removes any TLB entry that matches exactly the specified virtual page and context number. It is illegal to have more than one TLB entry per page.  
Demap page may, in fact, remove more than one TLB entry in the condition of a multiple TLB match, but this is an error condition of the TLB and has undefined results.
- **Demap context** — Removes any TLB entries that match the specified context identifier.
- **Demap all** — Removes all of the TLB entries from the TLB except for locked entries.

Demap is initiated by an STXA with ASI 5F<sub>16</sub> for D-MMU demap. It removes TLB entries from an on-chip TLB. No bus-based demap is supported. The demap address format is illustrated in FIGURE 11-21 and described in TABLE 11-23.



**FIGURE 11-21** MMU Demap Operation Address and Data Formats

**TABLE 11-23** Demap Address Format

| Bit(s)           | Field                        | Type | Description  |                  |                       |    |                           |    |                              |    |                          |    |                    |
|------------------|------------------------------|------|--|------------------|-----------------------|----|---------------------------|----|------------------------------|----|--------------------------|----|--------------------|
| 63:13            | VA<63:13>                    | RW   | The virtual page number of the TTE to be removed from the TLB for Demap Page.  |                  |                       |    |                           |    |                              |    |                          |    |                    |
| 12:8             | <i>Reserved</i>              |      | This field is ignored by hardware.   |                  |                       |    |                           |    |                              |    |                          |    |                    |
| 7:6              | Type                         | RW   | The type of demap operation, as described below: <table border="1" data-bbox="572 449 1179 604"> <thead> <tr> <th>Type Field</th> <th>Demap Operation</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Demap page — see page 290</td> </tr> <tr> <td>1</td> <td>Demap context — see page 290</td> </tr> <tr> <td>2</td> <td>Demap all — see page 290</td> </tr> <tr> <td>3</td> <td>Reserved — Ignored</td> </tr> </tbody> </table>                       | Type Field       | Demap Operation       | 0  | Demap page — see page 290 | 1  | Demap context — see page 290 | 2  | Demap all — see page 290 | 3  | Reserved — Ignored |
| Type Field       | Demap Operation              |      |  |                  |                       |    |                           |    |                              |    |                          |    |                    |
| 0                | Demap page — see page 290    |      |  |                  |                       |    |                           |    |                              |    |                          |    |                    |
| 1                | Demap context — see page 290 |      |  |                  |                       |    |                           |    |                              |    |                          |    |                    |
| 2                | Demap all — see page 290     |      |  |                  |                       |    |                           |    |                              |    |                          |    |                    |
| 3                | Reserved — Ignored           |      |  |                  |                       |    |                           |    |                              |    |                          |    |                    |
| 5:4              | Context ID                   | RW   | Context Register selection, as described below. Use of the reserved value causes the demap to be ignored. <table border="1" data-bbox="572 711 1179 866"> <thead> <tr> <th>Context ID Field</th> <th>Context Used in Demap</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>Primary</td> </tr> <tr> <td>01</td> <td>Secondary (D-MMU only)</td> </tr> <tr> <td>10</td> <td>Nucleus</td> </tr> <tr> <td>11</td> <td>Reserved</td> </tr> </tbody> </table> | Context ID Field | Context Used in Demap | 00 | Primary                   | 01 | Secondary (D-MMU only)       | 10 | Nucleus                  | 11 | Reserved           |
| Context ID Field | Context Used in Demap        |      |  |                  |                       |    |                           |    |                              |    |                          |    |                    |
| 00               | Primary                      |      |  |                  |                       |    |                           |    |                              |    |                          |    |                    |
| 01               | Secondary (D-MMU only)       |      |  |                  |                       |    |                           |    |                              |    |                          |    |                    |
| 10               | Nucleus                      |      |  |                  |                       |    |                           |    |                              |    |                          |    |                    |
| 11               | Reserved                     |      |  |                  |                       |    |                           |    |                              |    |                          |    |                    |
| 3:0              | 0                            |      | Set to zero.   |                  |                       |    |                           |    |                              |    |                          |    |                    |

A demap operation does not invalidate the TSB in memory. Software must modify the appropriate TTEs in the TSB before initiating a demap operation.

Except for Demap All, the demap operation does not depend on the value of any entry's lock bit. A demap operation demaps both locked entries and unlocked entries.

The demap operation produces no output.

The following are Data demap page types:

- **Data Demap Page (Type = 0).** Demap Page removes the TTE (from the specified TLB), matching the specified virtual page number and Context Register. The match condition with regard to the global bit is the same as a normal TLB access; that is, if the global bit is set, the contexts do not need to match.

Virtual page offset bits 15:13, 18:13, and 21:13 for 64 KB, 512 MB, and 4 MB page TLB entries, respectively, do not participate in the match for that entry. This is the same condition as for a translation match.

---


**Note** – Each Demap Page operation removes only one TLB entry. A demap of a 64 KB, 512 KB, or 4 MB page does not demap any smaller page within the specified virtual address range.

---

- **Data Demap Context (Type = 1).** Demap Context removes from the TLB all TTEs having the specified context. If the TTE Global bit is set, then the TTE is not removed. VA is ignored for this operation.
- **Data Demap All (Type = 2).** Demap All removes all TTEs that do not have the lock bit set. VA and Context are ignored for this operation.

### 11.9.9.15 Data TLB CAM Diagnostic Register

Accesses to the TLB Diagnostic Register require the virtual address to be set to access a TLB and TLB entry. The virtual address format of the TLB Diagnostic Register virtual address is described in TABLE 11-24 and illustrated in FIGURE 11-22.



**FIGURE 11-22** MMU TLB Diagnostic Access Virtual Address

**TABLE 11-24** TLB Diagnostic Register Virtual Address Format


| Bit(s)    | Field   | Type          | Description   |     |          |         |   |  |           |   |   |               |   |   |               |
|-----------|---|---------------|---|-----|----------|---------|---|--|-----------|---|---|---------------|---|---|---------------|
| 63:19     | —   | —             | <i>Reserved.</i>  |     |          |         |   |  |           |   |   |               |   |   |               |
| 18        | 1   | —             | Set to one. Selects CAM Diagnostic mode for the t16 TLB only (the other two TLBs do not employ CAM registers).  |     |          |         |   |  |           |   |   |               |   |   |               |
| 17:16 (D) | TLB #   | RW            | The number of the TLB to access, as follows: <table border="1" style="width: 100%; border-collapse: collapse; margin-top: 5px;"> <thead> <tr> <th style="text-align: left;">TLB</th> <th style="text-align: left;">TLB Type</th> <th style="text-align: left;">Entries</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Fully associative 64 KB, 512 KB, and 4 MB page size and locked pages</td> <td>16 (dt16)</td> </tr> <tr> <td>2</td> <td>2-way associative 8 KB, 64 KB, 512 KB, and 4 MB locked/unlocked page size</td> <td>512 (dt512_0)</td> </tr> <tr> <td>3</td> <td>2-way associative 8 KB, 64 KB, 512 KB, and 4 MB locked/unlocked page size</td> <td>512 (dt512_1)</td> </tr> </tbody> </table> | TLB | TLB Type | Entries | 0 | Fully associative 64 KB, 512 KB, and 4 MB page size and locked pages | 16 (dt16) | 2 | 2-way associative 8 KB, 64 KB, 512 KB, and 4 MB locked/unlocked page size | 512 (dt512_0) | 3 | 2-way associative 8 KB, 64 KB, 512 KB, and 4 MB locked/unlocked page size | 512 (dt512_1) |
| TLB       | TLB Type  | Entries       |   |     |          |         |   |  |           |   |   |               |   |   |               |
| 0         | Fully associative 64 KB, 512 KB, and 4 MB page size and locked pages      | 16 (dt16)     |   |     |          |         |   |  |           |   |   |               |   |   |               |
| 2         | 2-way associative 8 KB, 64 KB, 512 KB, and 4 MB locked/unlocked page size | 512 (dt512_0) |   |     |          |         |   |  |           |   |   |               |   |   |               |
| 3         | 2-way associative 8 KB, 64 KB, 512 KB, and 4 MB locked/unlocked page size | 512 (dt512_1) |   |     |          |         |   |  |           |   |   |               |   |   |               |



**TABLE 11-24** TLB Diagnostic Register Virtual Address Format (*Continued*)

| Bit(s) | Field       | Type | Description   |
|--------|-------------|------|---|
| 15:12  | —           | —    | <i>Reserved.</i>  |
| 11:3   | TLB Entry # | RW   | The number of the TLB entry to be accessed, in the range 0 – 511. Not all TLBs will have all 512 entries. All TLBs regardless of size are accessed from 0 to $N - 1$ , where $N$ is the number of entries in the TLB. |
| 2:0    | 0           | —    | Set to zero.  |

The format for the CAM Diagnostic Register is described in TABLE 11-25 and illustrated in FIGURE 11-23.

**FIGURE 11-23** D-MMU TLB CAM Diagnostic Registers**TABLE 11-25** CAM Diagnostic Register

| Bit(s) | Field    | Type | Description  |
|--------|----------|------|--|
| 63:7   | —        | —    | <i>Reserved.</i>                                   |
| 6      | LRU      | RW   | The LRU bit in the CAM, read-write.                |
| 5:3    | RAM SIZE | R    | The 3-bit page size field from the RAM, read-only. |
| 2:0    | CAM SIZE | R    | The 3-bit page size field from the CAM, read-only. |

An ASI store to the TLB CAM Diagnostic Register initiates an internal atomic write to the specified TLB entry. The TLB RAM and CAM entry data are obtained from the store data.

An ASI load from the TLB CAM Diagnostic Register initiates an internal read of the data portion of the specified TLB RAM and CAM entry.

## 11.9.10 D-MMU Bypass

In a bypass access, the D-MMU sets the physical address equal to the truncated virtual address; that is, the low-order bits of the virtual address are passed through without translation as the physical address. The physical page attribute bits are set as shown in TABLE 11-26.

TABLE 11-26 Bypass Attribute Bits

| ASI   | Attribute Bits |    |    |   |   |   |     |      |
|---|----------------|----|----|---|---|---|-----|------|
|   | CP             | IE | CV | E | P | W | NFO | Size |
| ASI_PHYS_USE_EC<br>ASI_PHYS_USE_EC_LITTLE                           | 1              | 0  | 0  | 0 | 0 | 1 | 0   | 8 KB |
| ASI_PHYS_BYPASS_EC_WITH_EBIT<br>ASI_PHYS_BYPASS_EC_WITH_EBIT_LITTLE | 0              | 0  | 0  | 1 | 0 | 1 | 0   | 8 KB |

---

**Compatibility Note** – The virtual address is wider than the physical address; thus, there is no need to use multiple ASIs to fill in the high-order physical address bits, as is done in SPARC V8 machines.

---

---

## 11.10 Instruction Translation Lookaside Buffer

In the I-MMU, a 128-entry, 2-way associative TLB (it128) is used exclusively for 8 KB, unlocked page translations, and a 16-entry fully associative TLB is used for 64 KB, 512 KB, and 4 MB page translations and locked pages of all four sizes.

### 11.10.1 I-TLB Access Operation

When an instruction fetch access is issued, its VA and Context are presented to the I-MMU. Both I-TLBs (it128 and it16) are accessed in parallel. The two possible 8K page entries in the it128 are checked to determine a hit.

There are two possible Context numbers active in the CPU, Primary (PContext field in ASI\_PRIMARY\_CONTEXT\_REG), and Nucleus (default to Context = 0). When TL = 0 primary context number is used and when TL > 0, nucleus context number is used.

Since two I-TLBs are being accessed in parallel, software must guarantee that there are no duplicate (stale) entry hits. Most of this responsibility lies in software (operating system) with the hardware providing some assistance to support full software control. A set of rules on I-TLB replacement, demap and context switch must be followed to maintain consistent and correct behavior.

## 11.10.2 I-TLB Automatic Replacement

An I-TLB miss fast trap handler utilizes the automatic (hardware) replacement write using store `ASI_ITLB_DATA_IN_REG`.

When I-TLB miss, or `DATA_ACCESS_EXCEPTION`, or `FAST_DATA_ACCESS_PROTECTION` is detected, hardware automatically saves the missing VA and context to the Tag Access Register (`ASI_IMMU_TAG_ACCESS`). To ease indexing of the `it128` when the TTE data is presented (via `STXA ASI_ITLB_DATA_IN_REG`), the missing page size information of `it128` is captured into a new Extension Register, called `ASI_IMMU_TAG_ACCESS_EXT`.

CODE EXAMPLE 11-3 shows the hardware I-TLB replacement algorithm.

### CODE EXAMPLE 11-3 I-TLB Hardware Replacement Algorithm

```
if (TTE to fill is a locked page, i.e., L bit is set) {
    fill TTE to it16;
} else {
    if (TTE's Size == 8K) {
        if (one of the 2 same-index entries is invalid) {
            fill TTE to an invalid entry
        } else if (no entry is valid |
                   both entries are valid) {
            case (LFSR[0]) {
                0: fill TTE to it128 way0;
                1: fill TTE to it128 way1;
            }
        }
    } else {
        fill TTE to it16;
    }
}
```

## 11.10.3 I-TLB Access Summary

TABLE 11-27 lists the I-MMU TLB access summary.

**TABLE 11-27** I-MMU TLB Access Summary

| Software Operation |             | Effect on MMU Physical Registers                         |  |                    |                   |
|--------------------|-------------|--|--|--------------------|-------------------|
| Load/Store         | Register    | TLB Tag Array  | TLB Data Array   | Tag Access         | SFSR              |
| Load               | Tag Read    | Contents returned. From entry specified by LDXA's access | No effect  | No effect          | No effect         |
|                    | Tag Access  | No effect  | No effect  | Contents returned. | No effect         |
|                    | Data In     | Trap with <i>data_access_exception</i>                   |  |                    |                   |
|                    | Data Access | No effect  | Contents returned. From entry specified by LDXA's access | No effect          | No effect         |
|                    | SFSR        | No effect  | No effect  | No effect          | Contents returned |
|                    | SFAR        | No effect  | No effect  | No effect          | No effect         |

TABLE 11-27 I-MMU TLB Access Summary (Continued)

| Software Operation |             | Effect on MMU Physical Registers  |  |                                       |  |
|--------------------|-------------|---|--|---------------------------------------|--|
| Load/<br>Store     | Register    | TLB Tag Array   | TLB Data Array   | Tag Access                            | SFSR   |
| Store              | Tag Read    | Trap with <i>data_access_exception</i>  |  |                                       |  |
|                    | Tag Access  | No effect   | No effect  | Written with store data               | No effect  |
|                    | Data In     | TLB entry determined by replacement policy written with contents of Tag Access Register | TLB entry determined by replacement policy written with store data | No effect                             | No effect  |
|                    | Data Access | TLB entry specified by STXA address written with contents of Tag Access Register        | TLB entry specified by STXA address written with store data        | No effect                             | No effect  |
|                    | SFSR        | No effect   | No effect  | No effect                             | Written with store data  |
|                    | SFAR        | No effect   | No effect  | No effect                             | No effect  |
| TLB miss           |             | No effect   | No effect  | Written with VA and context of access | Written with fault status of faulting instruction and page sizes at faulting context for two 2-way set associative TLB |

## 11.10.4 I-MMU Operation Summary

The behavior of the I-MMU is summarized in TABLE 11-28.

**TABLE 11-28** I-MMU Table of Operations for Normal ASIs

| Condition | Behavior  |          |       |       |
|-----------|-----------|----------|-------|-------|
|           | PRIV mode | TLB Miss | P = 0 | P = 1 |
| 0         |           | Imiss    | OK    | Iexc  |
| 1         |           | Imiss    | OK    | OK    |

## 11.10.5 Internal Registers and ASI Operations

In this section, how to access MMU registers is first described followed by the registers themselves and are described as follows:

- I-MMU TLB Tag Access Registers
- I-TLB Data In, Data Access, and Tag Read Registers
- I-TSB Tag Target Registers
- I-TSB Base Registers
- I-TSB Extension Registers
- I-TSB 8 KB and 64 KB Pointer and Direct Pointer Registers
- Instruction Synchronous Fault Status Registers (I-SFSR, D-SFSR)
- MMU Synchronous Fault Address Register

Following the register descriptions, the Instruction demap operation is described.

### 11.10.5.1 Instruction MMU TLB Tag Access Registers

In each MMU, the Tag Access Register is used as a temporary buffer for writing the TLB Entry tag information. The Tag Access Register holds the tag portion, and the Data In or Data Access Register holds the data being accessed.

The Tag Access Register can be updated during either of the following operations:

- When the MMU signals a trap due to a miss, exception, or protection

The MMU hardware automatically writes the missing VA and the appropriate context into the Tag Access Register to facilitate formation of the TSB Tag Target Register. One exception is that after a *data\_access\_exception*, the contents of the Context field of the I-MMU Tag Access Register are undefined.

- An ASI write to the Tag Access Register


Before an ASI store to the TLB Data Access Registers, the operating system must set the Tag Access Register to the values desired in the TLB Entry. Note that an ASI store to the TLB Data In Register for automatic replacement also uses the Tag Access Register, but typically the value written into the Tag Access Register by the MMU hardware is appropriate.

---

**Note** – Any update to the Tag Access Registers immediately affects the data that is returned from subsequent reads of the TSB Tag Target and TSB Pointer Registers.

---

The TLB Tag Access Register fields are defined in TABLE 11-29 and illustrated in FIGURE 11-12.



**FIGURE 11-24** I-MMU TLB Tag Access Registers

**TABLE 11-29** I-MMU Tag Access Register

| Bit   | Field   | Type | Description   |
|-------|---------|------|---|
| 63:13 | VA      | RW   | The 51-bit virtual page number.   |
| 12:0  | Context | RW   | The 13-bit context identifier. This field reads 0 when there is no associated context with the access. Its contents in the I-MMU are undefined after a <i>data_access_exception</i> . |

---

**Caution** – When the I-MMU causes a trap due to a protection violation or other exception, software should use the context number from `I-SFSR.CT` instead of from the `Context` field of the I-TLB Tag Access Register.

---

### 11.10.5.2 Instruction TLB Data In, Data Access, and Tag Read Registers

Access to the TLB is complicated because of the need to provide an atomic write of a TLB entry data item (tag and data) that is larger than 64 bits, the need to replace entries automatically through the TLB entry replacement algorithm as well as to provide direct diagnostic access, and the need for hardware assist in the TLB miss handler.

TABLE 11-2 on page 11-259 shows when loads and stores update the Tag Access Registers.

TABLE 11-30 shows how the Tag Read, Tag Access, Data In, and Data Access Registers interact to provide atomic reads and writes to the TLBs.

**TABLE 11-30** MMU TLB Access Summary

| Software Operation |             | Effect on MMU Physical Registers  |  |                                       |
|--------------------|-------------|---|--|---------------------------------------|
| Load/Store         | Register    | TLB Tag Array   | TLB Data Array   | Tag Access Register                   |
| Load               | Tag Read    | Contents returned. Entry specified by STXA's access.                                    | No effect  | No effect                             |
|                    | Tag Access  | No effect   | No effect  | Contents returned                     |
|                    | Data In     | Trap with <i>data_access_exception</i> .  |  |                                       |
|                    | Data Access | No effect   | Contents returned. Entry specified by STXA's access.               | No effect                             |
| Store              | Tag Read    | Trap with <i>data_access_exception</i> .  |  |                                       |
|                    | Tag Access  | No effect   | No effect  | Written with store data               |
|                    | Data In     | TLB entry determined by replacement policy written with contents of Tag Access Register | TLB entry determined by replacement policy written with store data | No effect                             |
|                    | Data Access | TLB entry specified by STXA address written with contents of Tag Access Register        | TLB entry specified by STXA address written with store data        | No effect                             |
| TLB miss           |             | No effect   | No effect  | Written with VA and context of access |

An ASI load from the TLB Tag Read Register initiates an internal read of the tag portion of the specified TLB entry.

### 11.10.5.3 Data In and Data Access Registers

The Data In and Data Access Registers are the means of reading and writing the TLB for all operations. The TLB Data In Register is used for TLB miss handler automatic replacement writes. The TLB Data Access Register is used for operating system and diagnostic directed writes (writes to a specific TLB entry).

An ASI load from the TLB Data Access Register initiates an internal read of the data portion of the specified TLB entry.

ASI loads from the TLB Data In Register are not supported.



An ASI store to the TLB Data In Register initiates an automatic atomic replacement of the TLB Entry pointed to by an internal register that is updated by a proprietary replacement algorithm. The TLB data and tag are formed as in the case of an ASI store to the TLB Data Access Register.

---


**Caution** – Stores to the Data In Register are not guaranteed to replace the previous TLB entry, causing a fault. In particular, to change an entry’s attribute bits, software must explicitly demap the old entry before writing the new entry; otherwise, a multiple match error condition can result.

---

Both the TLB Data In Register and the TLB Data Access Register use the TTE format shown in FIGURE 11-3 on page 11-248. Refer to the description of the TTE data in “Translation Table Entry” on page 248 for a complete description of the data fields.

Writes to the TLB Data In Register require the virtual address to be set to 0.

The format of the TLB Data Access Register virtual address is illustrated in FIGURE 11-25 and described in TABLE 11-31.



**FIGURE 11-25** I-MMU TLB Data Access Address

**TABLE 11-31** TLB Data Access Register


| Bit(s) | Field  | Type         | Description   |       |          |              |   |  |    |   |                                   |             |
|--------|--|--------------|---|-------|----------|--------------|---|--|----|---|-----------------------------------|-------------|
| 63:19  | —  | —            | <i>Reserved.</i>  |       |          |              |   |  |    |   |                                   |             |
| 18     | 0  | —            | Set to zero.  |       |          |              |   |  |    |   |                                   |             |
| 17:16  | TLB #  | RW           | The TLB to access, as defined below. <table border="1" style="width: 100%; border-collapse: collapse; margin-top: 5px;"> <thead> <tr> <th style="text-align: left;">TLB #</th> <th style="text-align: left;">TLB Type</th> <th style="text-align: left;"># of Entries</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Fully associative 64 KB, 4 MB, and 512 KB page size and locked pages</td> <td>16</td> </tr> <tr> <td>2</td> <td>2-way associative, 8 KB page size</td> <td>128 (I-MMU)</td> </tr> </tbody> </table> | TLB # | TLB Type | # of Entries | 0 | Fully associative 64 KB, 4 MB, and 512 KB page size and locked pages | 16 | 2 | 2-way associative, 8 KB page size | 128 (I-MMU) |
| TLB #  | TLB Type   | # of Entries |   |       |          |              |   |  |    |   |                                   |             |
| 0      | Fully associative 64 KB, 4 MB, and 512 KB page size and locked pages | 16           |   |       |          |              |   |  |    |   |                                   |             |
| 2      | 2-way associative, 8 KB page size                                    | 128 (I-MMU)  |   |       |          |              |   |  |    |   |                                   |             |

**TABLE 11-31** TLB Data Access Register (Continued)

| Bit(s) | Field     | Type | Description  |
|--------|-----------|------|--|
| 15:10  | —         | —    | Reserved.  |
| 9:3    | TLB Entry | RW   | For it128, the TLB entry number to be accessed, in the range 0 – 127. Therefore, all 7 bits of the TLB entry are used.<br><br>For it16, the TLB entry number to be accessed, in the range 0 – 15. Therefore, all 4 bits of the TLB entry are used. |
| 2:0    | 0         | —    | Set to zero.   |

#### 11.10.5.4 Instruction MMU TLB Tag Read Register

The format for the Tag Read Register virtual address is described in TABLE 11-32 and illustrated in FIGURE 11-26.

**FIGURE 11-26** I-MMU TLB Tag Read Registers**TABLE 11-32** I-MMU TLB Tag Read Register


| Bit   | Field               | Type | Description  |
|-------|---------------------|------|--|
| 63:13 | VA                  | RW   | The 51-bit virtual page number. In the fully associative TLB, page offset bits for larger page sizes are stored in the TLB; that is, VA<15:13>, VA<18:13>, and VA<21:13> for 64 KB, 512 KB, and 4 MB pages, respectively. These values are ignored during normal translation. When read, the UltraSPARC III Cu processor will return either 0 or the value previously written to them. |
| 11:0  | Instruction Context | RW   | The 13-bit context identifier.   |

#### 11.10.5.5 Instruction MMU TLB Tag Access Register

An ASI store to the TLB Data Access or Data In Register initiates an internal atomic write to the specified TLB Entry. The TLB entry data are obtained from the store data, and the TLB entry tag is obtained from the current contents of the TLB Tag Access Register.

## 11.10.5.6 Instruction TSB Tag Target Registers

The I-TSB Tag Target Registers are simply bit-shifted versions of the data stored in the Instruction Tag Access Registers, respectively. Since the Instruction Tag Access Register is updated on an I-TLB miss, respectively, the Instruction Tag Target Registers appear to software to be updated on an I-TLB miss. The MMU Tag Target Register is described in TABLE 11-33 and illustrated in FIGURE 11-27.



**FIGURE 11-27** MMU Tag Target Registers


**TABLE 11-33** MMU Tag Target Register

| Bit(s) | Field         | Type | Description   |
|--------|---------------|------|---|
| 63:61  | 000           | —    | Set to 000.   |
| 60:48  | Context<12:0> | RW   | The context associated with the missing virtual address.  |
| 47:42  | —             | —    | <i>Reserved.</i>  |
| 41:0   | VA<63:22>     | RW   | The most significant bits of the missing virtual address. |

## 11.10.5.7 Instruction TSB Base Registers

The Translation Storage Buffer (TSB) registers provide information for the hardware formation of TSB pointers and tag target, to assist software in quickly handling TLB misses. If the TSB concept is not employed in the software memory management strategy and therefore the Pointer and Tag Access Registers are not used, then the TSB registers need not contain valid data.

The TSB register is illustrated in FIGURE 11-28 and described in TABLE 11-34.



**FIGURE 11-28** MMU Instruction TSB Registers


**TABLE 11-34** TSB Register Description

| Bit   | Field                   | Type | Description   |
|-------|-------------------------|------|---|
| 63:13 | Instruction<br>TSB_Base | RW   | Provides the base virtual address of the Translation Storage Buffer. Software must ensure that the TSB base is aligned on a boundary equal to the size of the TSB or both TSBs in the case of a split TSB.  |
| 12    | Split                   | RW   | When <code>Split = 1</code> , the TSB 64 KB pointer address is calculated assuming separate (but abutting and equally sized) TSB regions for the 8 KB and the 64 KB TTEs. In this case, <code>TSB_Size</code> refers to the size of each TSB. The TSB 8 KB pointer address calculation is not affected by the value of the <code>Split</code> bit. When <code>Split = 0</code> , the TSB 64 KB pointer address is calculated assuming that the same lines in the TSB are shared by 8 KB and 64 KB TTEs, called a “common TSB” configuration.<br><br><b>Caution:</b> In the “common TSB” configuration ( <code>TSB.Split = 0</code> ), 8 KB and 64 KB page TTEs can conflict unless the TLB miss handler explicitly checks the TTE for page size. Therefore, do not use the common TSB mode in an optimized handler. For example, suppose an 8 KB page at $VA = 2000_{16}$ and a 64 KB page at $VA = 10000_{16}$ both exist — a legal situation. These both map to the second TSB line (line 1) and have the same <code>VA</code> tag of 0. Therefore, there is no way for the miss handler to distinguish these TTEs by the TTE tag alone, and unless the miss handler checks the TTE data, it may load an incorrect TTE. |
| 11:3  | —                       | —    | Reserved.   |
| 2:0   | Instruction<br>TSB_Size | RW   | The UltraSPARC III Cu processor implements a 3-bit <code>TSB_Size</code> field. The <code>TSB_Size</code> field provides the size of the TSB as follows: <ul style="list-style-type: none"> <li>• The number of entries in the TSB (or each TSB if split) = <math>512 \times 2^{TSB\_Size}</math>.</li> <li>• The number of entries in the TSB ranges from 512 entries at <code>TSB_Size = 0</code> (8 KB common TSB, 16 KB split TSB), to 64K entries at <code>TSB_Size = 7</code> (1 MB common TSB, 2 MB split TSB).</li> </ul> <b>Note:</b> Any update to the TSB register immediately affects the data that are returned from later reads of the Tag Target and TSB Pointer Registers.  |

### 11.10.5.8 Instruction TSB Extension Registers

The TSB Extension Registers provide information for the hardware formation of TSB pointers and tag target, to assist software in handling TLB misses quickly. If the TSB concept is not employed in the software memory management strategy and therefore the pointer and Tag Access Registers are not used, then the TSB Extension Registers need not contain valid data.

The TSB registers are defined as follows in FIGURE 11-29.



**FIGURE 11-29** MMU Instruction TSB Extension Registers

In the UltraSPARC III Cu processor, TSB\_Hash (bits 11:3 of the Extension Registers) are exclusive-ORed with the calculated TSB offset to provide a “hash” into the TSB. Changing the TSB\_Hash field on a per-process basis minimizes the collision of TSB entries between different processes.

There are two TSB Extension Registers, one for each of the virtual address spaces (Primary, Nucleus).


When an I-TLB miss occurs, an appropriate TSB Extension Register is selected and XORed with the I-TSB Register. The result is then used to form a TSB pointer, as described in “TSB Pointer Formation” on page 254.

### 11.10.5.9 Instruction TSB 8 KB and 64 KB Pointer and Direct Pointer Registers

The I-TSB 8 KB and 64 KB registers are provided as an aid to software in determining the location of the missing or trapping TTE in the software-maintained TSB. The TSB 8 KB and 64 KB Pointer Registers provide the possible locations of the 8 KB and 64 KB TTE, respectively.

As a fine point, the bit that controls selection of 8 KB or 64 KB address formation for the Direct Pointer Register is a state bit in the I-MMU that is updated during a *fast\_data\_access\_protection* exception. It records whether the page that hit in the TLB was a 64 KB page or a non-64 KB page, in which case, 8 KB is assumed.

The registers are illustrated in FIGURE 11-30, where VA<63:4> is the full virtual address of the TTE in the TSB, as determined by the MMU hardware, and is described in “Hardware Support for TSB Access” on page 253.



**FIGURE 11-30** I-MMU TSB 8 KB/64 KB Pointer and I-MMU Direct Pointer Register

## TSB 8 KB and 64 KB Pointer Registers

The TSB Pointer Registers are implemented as a reorder of the current data stored in the Tag Access Register and the TSB Extension Register. If the Tag Access Register or TSB Extension Register is updated through a direct software write (through an STXA instruction), then the values in the Pointer Registers will be updated as well.

### Direct Pointer Register

The Direct Pointer Register is mapped by hardware to either the 8 KB or 64 KB Pointer Register in the case of a *fast\_data\_access\_protection* exception according to the known size of the trapping TTE. In the case of a 512 KB or 4 MB page miss, the Direct Pointer Register returns the pointer as if the fault were from an 8 KB page.

## 11.10.5.10 Instruction Synchronous Fault Status Registers (I-SFSR)

The I-MMU maintains its own SFSR Register. The SFSR is illustrated in FIGURE 11-31 and described in TABLE 11-35.




FIGURE 11-31 MMU Instruction Synchronous Fault Status Registers (I-SFSR, D-SFSR)

TABLE 11-35 SFSR Bit Description

| Bit(s) | Field | Type | Description  |
|--------|-------|------|--|
| 63:25  | —     |      | <i>Reserved.</i>   |
| 24     | NF    | RW   | Set in the I-MMU if the faulting instruction was a non-faulting load (a load to ASI_NOFAULT) (I-MMU = 0). NF is always 0 in I-SFSR.  |
| 23:16  | ASI   | RW   | Records the 8-bit ASI associated with the faulting instruction. This field is valid for both I-MMU and I-MMU SFSRs and for all traps in which the FV bit is set. A trapping alternate space load or store sets the ASI field to the ASI the instruction attempted to reference. A trapping non-alternate-space load or store sets ASI to ASI_PRIMARY if PSTATE.CLE = 0 or to ASI_PRIMARY_LITTLE if PSTATE.CLE = 1. A <i>mem_address_not_aligned</i> trap caused by a JMWPL or RETURN either does not set DSFSR.ASI or sets it as would a trapping non-alternate-space load or store. |
| 15     | TM    | RW   | I-TLB miss.  |
| 14:12  | —     |      | <i>Reserved.</i>   |

**TABLE 11-35** SFSR Bit Description (Continued)

| Bit(s)     | Field         | Type          | Description  |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
|------------|---------------|---------------|--|------------|---------------|---------------|----|---------|---------|----|----------|-----------|----|---------|---------|----|----------|----------|
| 11:7       | FT            | RW            | <p>Specifies the exact condition that caused the recorded fault, according to TABLE 11-35 following this table. In the I-MMU, the Fault Type field is valid only for <i>data_access_exception</i> faults; there is no ambiguity in all other MMU trap cases. Note that the hardware does not priority-encode the bits set in the fault type (FT) field; that is, multiple bits can be set. In particular, the following ASI stores could set both the 01<sub>16</sub> and 08<sub>16</sub> Fault Type bits (the page is privileged, as well as storing to a read-only ASI):</p> <pre>stda          %g0,  [ %g4 ]ASI_PRIMARY_NO_FAULT stda          %g0,  [ %g4 ]ASI_SECONDARY_NO_FAULT stda          %g0,  [ %g4 ]ASI_PRIMARY_NO_FAULT_LITTLE stda          %g0,  [ %g4 ]ASI_SECONDARY_NO_FAULT_LITTLE</pre> <p>The FT field in the I-MMU SFSR always reads 0 for <i>fast_instruction_access_MMU_miss</i> and reads 01<sub>16</sub> for <i>instruction_access_exception</i>, as all other fault types do not apply.</p> |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 6          | E             | RW            | Side-effect bit. It always reads as 0 in the I-MMU.  |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 5:4        | CT            | RW            | <p>Context Register selection, as described below. The context is set to 11<sub>2</sub> when the access does not have a translating ASI.</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>Context ID</th> <th>I-MMU Context</th> <th>D-MMU Context</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>Primary</td> <td>Primary</td> </tr> <tr> <td>01</td> <td>Reserved</td> <td>Secondary</td> </tr> <tr> <td>10</td> <td>Nucleus</td> <td>Nucleus</td> </tr> <tr> <td>11</td> <td>Reserved</td> <td>Reserved</td> </tr> </tbody> </table>  | Context ID | I-MMU Context | D-MMU Context | 00 | Primary | Primary | 01 | Reserved | Secondary | 10 | Nucleus | Nucleus | 11 | Reserved | Reserved |
| Context ID | I-MMU Context | D-MMU Context |  |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 00         | Primary       | Primary       |  |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 01         | Reserved      | Secondary     |  |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 10         | Nucleus       | Nucleus       |  |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 11         | Reserved      | Reserved      |  |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 3          | PR            | RW            | Privilege bit. Set if the faulting access occurred while in privileged mode. This field is valid for all traps in which the FV bit is set.   |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 2          | W             | RW            | Write bit. This bit always reads as 0 in the I-MMU SFSR.   |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 1          | OW            | RW            | Overwrite bit. When the MMU detects a fault, the Overwrite bit is set to 1 if the Fault Valid bit has not been cleared from a previous fault; otherwise, it is set to 0.   |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |
| 0          | FV            | RW            | <p>Fault Valid bit. Set when the MMU detects a fault; it is cleared only on an explicit ASI write of 0 to the SFSR. This bit is not set on an MMU miss. Therefore, overwrites of MMU misses cannot be detected.</p> <p>When the Fault Valid bit is not set, the values of the remaining fields in the SFSR and SFAR are undefined for traps other than an MMU miss.</p>  |            |               |               |    |         |         |    |          |           |    |         |         |    |          |          |

TABLE 11-36 describes the SFSR fault type field (FT<11:7>).

**TABLE 11-36** MMU Synchronous Fault Status Register FT (Fault Type) Field

| Instruction | FT[6:0]          | Fault Type           |
|-------------|------------------|----------------------|
| Instruction | 01 <sub>16</sub> | Privilege violation. |

---

**Note** – A *fast\_instruction\_MMU\_miss* trap causes the SFSR and the SFAR to be overwritten without setting either the OW or the FV bits.

---

The SFSR and the Tag Access Registers both maintain state concerning a previous translation causing an exception. The update policy for the SFSR and the Tag Access Registers is shown in TABLE 11-2 on page 11-259.

### 11.10.5.11 Synchronous Fault Addresses

This section describes how the I-MMU obtains a fault address.

#### *I-MMU Fault Address*

There is no I-MMU Synchronous Fault Address Register. Instead, software must read the TPC register appropriately as discussed here.

For *fast\_instruction\_access\_MMU\_miss* traps, TPC contains the virtual address that was not found in the I-MMU TLB.

For *instruction\_access\_exception* traps, “privilege violation” fault type, TPC contains the virtual address of the instruction in the privileged page that caused the exception.

### 11.10.5.12 Instruction MMU Demap

Demap is an MMU operation, not an MMU register. Demap removes selected entries from the TLBs.

---

**Note** – A store to an I-MMU Register requires a MEMBAR #Sync, FLUSH, DONE, or RETRY before the point that the effect must be visible to load/store/atomic accesses. A FLUSH, DONE, or RETRY is needed before the point that the effect must be visible to instruction accesses, that is, MEMBAR #Sync is not sufficient. In either case, one of these instructions must be executed before the next non-internal store or load of any type and on or before the delay slot of a delayed-control transfer instruction of any type. This action is necessary to avoid data corruption.

---

Three types of demap operations are provided:

- **Demap page** — Removes any TLB entry that matches exactly the specified virtual page and context number. It is illegal to have more than one TLB entry per page.  
Demap page may, in fact, remove more than one TLB entry in the condition of a multiple TLB match, but this is an error condition of the TLB and has undefined results.
- **Demap context** — Removes any TLB entries that match the specified context identifier.
- **Demap all** — Removes all of the TLB entries from the TLB except for locked entries.



Demap is initiated by an STXA with ASI 57<sub>16</sub> for I-MMU demap. It removes TLB entries from an on-chip TLB. No bus-based demap is supported. The demap address format is illustrated in FIGURE 11-32 and described in TABLE 11-37.




FIGURE 11-32 MMU Demap Operation Address and Data Formats

TABLE 11-37 Demap Address Format

| Bit(s)           | Field                        | Type | Description  |                  |                       |    |                           |    |                              |    |                          |    |                           |
|------------------|------------------------------|------|--|------------------|-----------------------|----|---------------------------|----|------------------------------|----|--------------------------|----|---------------------------|
| 63:13            | VA<63:13>                    | RW   | The virtual page number of the TTE to be removed from the TLB for Demap Page.  |                  |                       |    |                           |    |                              |    |                          |    |                           |
| 12:8             | Ignored                      |      | This field is ignored by hardware.   |                  |                       |    |                           |    |                              |    |                          |    |                           |
| 7:6              | Type                         | RW   | The type of demap operation, as described below: <table border="1" style="margin-left: 40px;"> <thead> <tr> <th>Type Field</th> <th>Demap Operation</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Demap page — see page 308</td> </tr> <tr> <td>1</td> <td>Demap context — see page 308</td> </tr> <tr> <td>2</td> <td>Demap all — see page 308</td> </tr> <tr> <td>3</td> <td><i>Reserved</i> — Ignored</td> </tr> </tbody> </table>                | Type Field       | Demap Operation       | 0  | Demap page — see page 308 | 1  | Demap context — see page 308 | 2  | Demap all — see page 308 | 3  | <i>Reserved</i> — Ignored |
| Type Field       | Demap Operation              |      |  |                  |                       |    |                           |    |                              |    |                          |    |                           |
| 0                | Demap page — see page 308    |      |  |                  |                       |    |                           |    |                              |    |                          |    |                           |
| 1                | Demap context — see page 308 |      |  |                  |                       |    |                           |    |                              |    |                          |    |                           |
| 2                | Demap all — see page 308     |      |  |                  |                       |    |                           |    |                              |    |                          |    |                           |
| 3                | <i>Reserved</i> — Ignored    |      |  |                  |                       |    |                           |    |                              |    |                          |    |                           |
| 5:4              | Context ID                   | RW   | Context Register selection, as described below. Use of the reserved value causes the demap to be ignored. <table border="1" style="margin-left: 40px;"> <thead> <tr> <th>Context ID Field</th> <th>Context Used in Demap</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>Primary</td> </tr> <tr> <td>01</td> <td><i>Reserved</i></td> </tr> <tr> <td>10</td> <td>Nucleus</td> </tr> <tr> <td>11</td> <td><i>Reserved</i></td> </tr> </tbody> </table> | Context ID Field | Context Used in Demap | 00 | Primary                   | 01 | <i>Reserved</i>              | 10 | Nucleus                  | 11 | <i>Reserved</i>           |
| Context ID Field | Context Used in Demap        |      |  |                  |                       |    |                           |    |                              |    |                          |    |                           |
| 00               | Primary                      |      |  |                  |                       |    |                           |    |                              |    |                          |    |                           |
| 01               | <i>Reserved</i>              |      |  |                  |                       |    |                           |    |                              |    |                          |    |                           |
| 10               | Nucleus                      |      |  |                  |                       |    |                           |    |                              |    |                          |    |                           |
| 11               | <i>Reserved</i>              |      |  |                  |                       |    |                           |    |                              |    |                          |    |                           |
| 3:0              | 0                            | —    | Set to zero.   |                  |                       |    |                           |    |                              |    |                          |    |                           |

A demap operation does not invalidate the TSB in memory. Software must modify the appropriate TTEs in the TSB before initiating a demap operation.

Except for Demap All, the demap operation does not depend on the value of any entry's lock bit. A demap operation demaps both locked entries and unlocked entries.

The demap operation produces no output.

The following are Instruction/Data demap page types:

- **Instruction Demap Page (Type = 0).** Demap Page removes the TTE (from the specified TLB), matching the specified virtual page number and Context Register. The match condition with regard to the global bit is the same as a normal TLB access; that is, if the global bit is set, the contexts do not need to match.

Virtual page offset bits 15:13, 18:13, and 21:13 for 64 KB, 512 KB, and 4 MB page TLB entries, respectively, do not participate in the match for that entry. This is the same condition as for a translation match.

---


**Note** – Each Demap Page operation removes only one TLB entry. A demap of a 64 KB, 512 KB, or 4 MB page does not demap any smaller page within the specified virtual address range.

---

- **Instruction Demap Context (Type = 1).** Demap Context removes from the TLB all TTEs having the specified context. If the TTE Global bit is set, then the TTE is not removed. VA is ignored for this operation.
- **Instruction Demap All (Type = 2).** Demap All removes all TTEs that do not have the lock bit set. VA and Context are ignored for this operation.

### 11.10.5.13 Instruction TLB Diagnostic Register

Accesses to the TLB Diagnostic Register require the virtual address to be set to access a TLB and TLB entry. The format of the TLB Diagnostic Register virtual address is described TABLE 11-38 and illustrated in FIGURE 11-33.



**FIGURE 11-33** MMU TLB Diagnostic Access Virtual Address Format

**TABLE 11-38** MMU TLB Diagnostic Register Virtual Address

| Bit   | Field  | Type        | Description  |     |          |         |   |  |           |   |                                  |             |
|-------|--|-------------|--|-----|----------|---------|---|--|-----------|---|----------------------------------|-------------|
| 63:19 | —  | —           | <i>Reserved.</i>   |     |          |         |   |  |           |   |                                  |             |
| 18    | 1  | —           | Set to one.  |     |          |         |   |  |           |   |                                  |             |
| 17:16 | TLB #  | RW          | The number of the table to access, as follows: <table border="1" data-bbox="803 418 1275 595"> <thead> <tr> <th>TLB</th> <th>TLB Type</th> <th>Entries</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Fully associative 64 KB, 4 MB, and 512 KB page size and locked pages (all sizes)</td> <td>16 (it16)</td> </tr> <tr> <td>2</td> <td>2-way associative 8 KB page size</td> <td>128 (it128)</td> </tr> </tbody> </table> | TLB | TLB Type | Entries | 0 | Fully associative 64 KB, 4 MB, and 512 KB page size and locked pages (all sizes) | 16 (it16) | 2 | 2-way associative 8 KB page size | 128 (it128) |
| TLB   | TLB Type   | Entries     |  |     |          |         |   |  |           |   |                                  |             |
| 0     | Fully associative 64 KB, 4 MB, and 512 KB page size and locked pages (all sizes) | 16 (it16)   |  |     |          |         |   |  |           |   |                                  |             |
| 2     | 2-way associative 8 KB page size   | 128 (it128) |  |     |          |         |   |  |           |   |                                  |             |
| 15:10 | —  | —           | <i>Reserved.</i>   |     |          |         |   |  |           |   |                                  |             |
| 9:3   | TLB Entry #  | RW          | For it128, the TLB entry number to be accessed, in the range 0 – 127. Therefore, all 7 bits of the TLB entry are used.<br><br>For it16, the TLB entry number to be accessed, in the range 0 – 15. Therefore, all 4 bits of the TLB entry are used.   |     |          |         |   |  |           |   |                                  |             |
| 2:0   | 0  | —           | Set to zero.   |     |          |         |   |  |           |   |                                  |             |

## 11.10.6 I-MMU Bypass

The I-MMU can only be bypassed by being disabled.



**SECTION V**

---

---

---

**Supervisor Programming**

---

---



## Traps and Trap Handling

---

A trap is a vectored transfer of control to supervisor software through a trap table that contains the first eight (32 for *clean\_window*, *spill*, *fill*, *fast\_instruction\_access\_MMU\_miss*, *fast\_data\_access\_MMU\_miss*, and *fast\_data\_access\_protection* traps) instructions of each trap handler. The base address of the table is established by supervisor software, by writing the Trap Base Address (TBA) register. The displacement within the table is determined by the trap type and the current trap level (TL). One-half of the table is reserved for hardware traps; one-quarter is reserved for software traps generated by TCC instructions; the remaining quarter is reserved for future use.

A trap behaves like an unexpected procedure call. It causes the hardware to do the following:

1. Save certain processors state (program counters, CWP, ASI, CCR, PSTATE, and the trap type) on a hardware register stack.
2. Enter privileged execution mode with a predefined PSTATE.
3. Begin executing trap handler code in the trap vector.

When the trap handler has finished, it uses either a DONE or RETRY instruction to return.

A trap may be caused by a TCC instruction, an instruction induced exception, a reset, an asynchronous error, or an interrupt request not directly related to a particular instruction. The processor must appear to behave as though, before executing each instruction, it determines if there are any pending exceptions or interrupt requests. If there are pending exceptions or interrupt requests, the processor selects the highest priority exception or interrupt request and causes a trap.

Thus, an *exception* is a condition that makes it impossible for the processor to continue executing the current instruction stream without software intervention. A *trap* is the action taken by the processor when it changes the instruction flow in response to the presence of an exception, interrupt, or TCC instruction.

A catastrophic error exception is due to the detection of a hardware malfunction, which due to the nature of the error, the state of the machine at the time of the exception cannot be restored. Since the machine state cannot be restored, execution after such an exception may not be resumable. An example of such an error is an uncorrectable bus parity error. Catastrophic errors cause deferred or disrupting traps.

Traps are described in the following sections:

- Processor States, Normal and Special Traps
- Trap Categories
- Trap Control
- Trap-Table Entry Addresses
- Trap Processing
- Exception and Interrupt Descriptions

---

## 12.1 Processor States, Normal and Special Traps

The processor is always in one of three discrete states:

- `execute_state`, which is the normal execution state of the processor
- `RED_state` (**R**eset, **E**rror, and **D**ebug state), which is a restricted execution state reserved for processing traps that occur when  $TL = MAXTL - 1$ , and for processing hardware-initiated and software-initiated resets
- `error_state`, which is a halted state that is entered as a result of a trap when  $TL = MAXTL$

Traps processed in `execute_state` are called *normal traps*. Traps processed in `RED_state` are called *special traps*.

FIGURE 12-1 shows the processor state diagram.



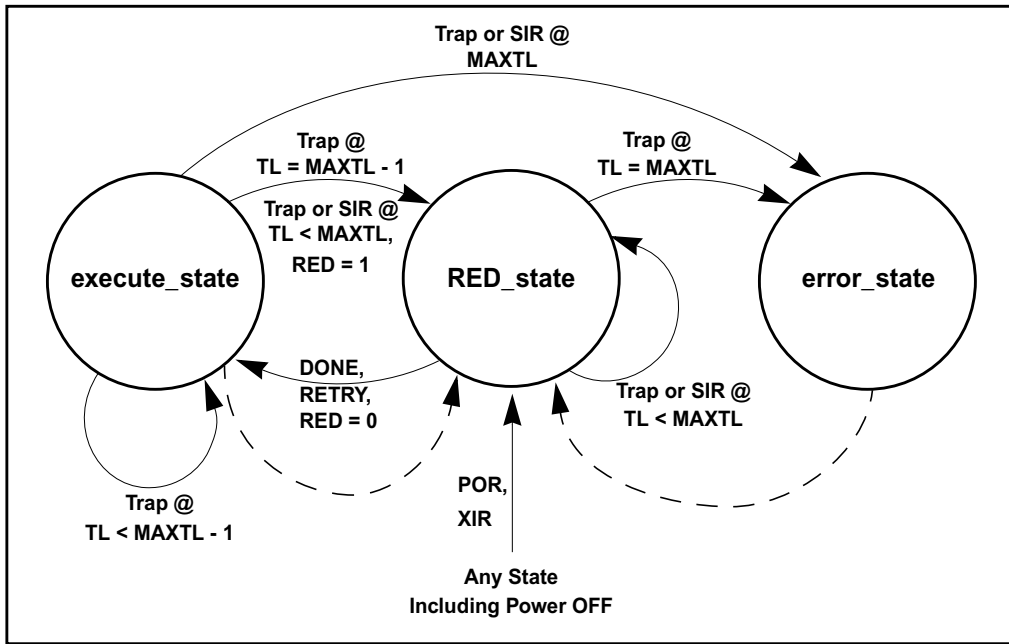


FIGURE 12-1 Processor State Diagram

### 12.1.1 RED\_state

`RED_state` is an acronym for **R**eset, **E**rror, and **D**ebug state. The processor enters `RED_state` under any one of the following conditions:

- A trap is taken when  $TL = MAXTL - 1$ .
- A POR, WDR, or XIR reset occurs.
- An SIR occurs when  $TL < MAXTL$ .
- System software sets `PSTATE.RED = 1`.

`RED_state` serves two mutually exclusive purposes:

- During trap processing, it indicates that no more trap levels are available; that is, if another nested trap is taken, the processor will enter `error_state` and halt. `RED_state` provides system software with a restricted execution environment.
- It provides the execution environment for all reset processing.

`RED_state` is indicated by `PSTATE.RED`. When this bit is set, the processor is in `RED_state`; when this bit is clear, the processor is not in `RED_state`, independent of the value of  $TL$ . Executing a `DONE` or `RETRY` instruction in `RED_state` restores the stacked copy of the `PSTATE` register, which clears the `PSTATE.RED` flag if the stacked copy had it

cleared. System software can also set or clear the `PSTATE.RED` flag with a `WRPR` instruction, which also forces the processor to enter or exit `RED_state`, respectively. In this case, the `WRPR` instruction should be placed in the delay slot of a jump so that the program counter (PC) can be changed in concert with the state change.

---

**Programming Notes** – Setting `TL = MAXTL` with a `WRPR` instruction does not also set `PSTATE.RED = 1` nor does it alter any other machine state. The values of `PSTATE.RED` and `TL` are independent.

Setting `PSTATE.RED` with a `WRPR` instruction causes the processor to execute in `RED_state`. However, it is different from a `RED_state` trap in the sense that there are no trap related changes in the machine state (for example, `TL` does not change).

---

### 12.1.1.1 RED\_state Trap Table

Traps occurring in `RED_state` or traps that cause the processor to enter `RED_state` use an abbreviated trap vector. The `RED_state` trap vector is constructed so that it can overlay the normal trap vector, if necessary. TABLE 12-1 illustrates the `RED_state` trap vector layout.

**TABLE 12-1** RED\_state Trap Vector Layout

| Offset           | TT             | Reason   |
|------------------|----------------|--|
| 00 <sub>16</sub> | 0              | Reserved (SPARC V8 reset)                      |
| 20 <sub>16</sub> | 1              | Power-on reset (POR)                           |
| 40 <sub>16</sub> | 2 <sup>†</sup> | Watchdog reset (WDR)                           |
| 60 <sub>16</sub> | 3 <sup>‡</sup> | Externally initiated reset (XIR)               |
| 80 <sub>16</sub> | 4              | Software-initiated reset (SIR)                 |
| A0 <sub>16</sub> | *              | All other exceptions in <code>RED_state</code> |

<sup>†</sup>TT = 2 if a WDR occurs while the processor is not in `error_state`; TT = trap type of the exception that caused entry into `error_state` if a WDR occurs in `error_state`.

<sup>‡</sup>TT = 3 if an *externally\_initiated\_reset* (XIR) occurs while the processor is not in `error_state`; TT = trap type of the exception that caused entry into `error_state` if the externally initiated reset occurs in `error_state`.

\*TT = trap type of the exception. See TABLE 12-3.

When the UltraSPARC III Cu processor processes a reset or trap that enters `RED_state`, it takes a trap at an offset relative to the `RED_state` trap vector base address (`RSTVaddr`); the base address is at virtual address `FFFF FFFF F000 000016`, which passes through to physical address `7FF F000 000016`.

### 12.1.1.2 RED\_state Execution Environment

In `RED_state`, the processor is forced to execute in a restricted environment by overriding the values of some processor controls and state registers.

The values are overridden, not set, allowing them to be switched atomically.

When `RED_state` is entered because of component failures, the handler should attempt to recover from potentially catastrophic error conditions or to disable the failing components. When `RED_state` is entered after a reset, the software should create the environment necessary to restore the system to a running state.

### 12.1.1.3 RED\_state Entry Traps

The following traps are processed in `RED_state` in all cases.

- **Power-on reset (POR)** — Implemented in hardware in UltraSPARC III Cu processors; not really a trap.
- **Watchdog reset (WDR)** — Implemented in hardware in UltraSPARC III Cu processors; this trap is used as a recovery mechanism from `error_state` in UltraSPARC III Cu. Upon an entry to `error_state`, the processor automatically invokes a WDR to enter `RED_state`.
- **Externally initiated reset (XIR)** — Implemented in hardware in UltraSPARC III Cu processors; typically used as a nonmaskable interrupt method for debug.

In addition, the following trap is processed in `RED_state` if  $TL < MAXTL$  when the trap is taken. Otherwise, it is processed in `error_state`.

- **Software-initiated reset (SIR)**

Traps that occur when  $TL = MAXTL - 1$  also set `PSTATE.RED = 1`; that is, any trap handler entered with  $TL = MAXTL$  runs in `RED_state`.

Any non-reset trap that sets `PSTATE.RED = 1` or that occurs when `PSTATE.RED = 1` branches to a special entry in the `RED_state` trap vector at `RSTVaddr + A016`.

### 12.1.1.4 RED\_state Software Considerations

In effect, `RED_state` reserves one level of the trap stack for recovery and reset processing. Software should be designed to require only  $MAXTL - 1$  trap levels for normal processing. That is, any trap that causes  $TL = MAXTL$  is an exceptional condition that should cause entry to `RED_state`.

The architected value for MAXTL in the UltraSPARC III Cu processor is five; typical usage of the trap levels is shown in TABLE 12-2.

**TABLE 12-2** Typical Usage for Trap Levels

| TL | Usage   |
|----|---|
| 0  | Normal execution  |
| 1  | System calls, interrupt handlers, instruction emulation |
| 2  | Window <i>spill/fill</i>                                |
| 3  | Page-fault handler                                      |
| 4  | Reserved for error handling                             |
| 5  | RED_state handler                                       |

---

**Programming Note** – To log the state of the processor, RED\_state handler software needs either a spare register or a preloaded pointer to a save area. To support recovery, the operating system might reserve one of the alternate global registers (for example, %a7) for use in RED\_state.

---

## 12.1.2 Error\_state

The processor enters `error_state` when a trap occurs while the processor is already at its maximum supported trap level, that is, when `TL = MAXTL`.

The processor automatically exits `error_state` using WDR. The processor signals itself internally to take a WDR and sets `TT = 2`. The WDR traps to the address at `RSTVaddr + 0x4016`. WDR sets the processor in a state where it is prepared to diagnose failures. A WDR behaves more like a trap than a reset.

WDR affects only one processor, rather than the entire system.

---

**Note** – When a window trap occurs at `TL = MAXTL`, the processor enters RED\_state and then performs a WDR. The current window pointer (CWP) is always updated as if the window trap was taken successfully.

---

---

## 12.2 Trap Categories

An exception or interrupt request can cause any of the following trap types:

- Precise trap
- Deferred trap
- Disrupting trap
- Reset trap

## 12.2.1 Precise Traps

A *precise trap* is induced by a particular instruction and occurs before any program-visible state has been changed by the trap-inducing instructions. When a precise trap occurs, several conditions must be true:

- The PC saved in TPC[TL] points to the instruction that induced the trap, and the next program counter (nPC) saved in TNPC[TL] points to the instruction that was to be executed next.
- All instructions issued before the one that induced the trap have completed execution.
- Any instructions issued after the one that induced the trap remain not executed.

Among the actions the trap handler software might take after a precise trap are the following:

- Return to the instruction that caused the trap and re-execute it by executing a RETRY instruction ( $PC \leftarrow \text{old } PC, nPC \leftarrow \text{old } nPC$ ).
- Emulate the instruction that caused the trap and return to the succeeding instruction by executing a DONE instruction ( $PC \leftarrow \text{old } nPC, nPC \leftarrow \text{old } nPC + 4$ ).
- Terminate the program or process associated with the trap.

## 12.2.2 Deferred Traps

A *deferred trap* is also induced by a particular instruction, but unlike a precise trap, a deferred trap may occur after program-visible state has been changed. Such state may have been changed by the execution of either the trap inducing instruction itself or by one or more other instructions.

Associated with a particular deferred trap implementation, the following must exist:

- An instruction that causes a potentially outstanding deferred trap exception to be taken as a trap.
- Privileged instructions that access the state information needed by the supervisor software to emulate the deferred trap inducing instruction and to resume execution of the trapped instruction stream.

---

**Programming Note** – Resuming execution may require the emulation of instructions that had not completed execution at the time of the deferred trap, that is, those instructions in the deferred trap queue.

---

Among the actions software can take after a deferred trap are the following:

- Emulate the instruction that caused the exception, emulate or cause to execute any other execution deferred instructions that were in an associated deferred trap state queue, and use `RETRY` to return control to the instruction at which the deferred trap was invoked.
- Terminate the program or process associated with the trap.

## 12.2.3 Disrupting Traps

A *disrupting trap* is neither a precise trap nor a deferred trap. A disrupting trap is caused by a *condition* (for example, an interrupt) rather than directly by a particular instruction; the cause distinguishes it from precise and deferred traps. When a disrupting trap has been serviced, trap handler software normally arranges for program execution to resume where it left off. This differentiates disrupting traps from reset traps, which trap to a unique reset address from which execution of the program that was running when the reset occurred is never expected to resume.

Disrupting traps are controlled by a combination of the Processor Interrupt Level (`PIL`) register and the Interrupt Enable (`IE`) field of `PSTATE`. A disrupting trap condition is ignored when interrupts are disabled (`PSTATE.IE = 0`) or when the condition's interrupt level is less than or equal to that specified in `PIL`.

A disrupting trap may be due either to an interrupt request not directly related to a previously executed instruction or an exception related to a previously executed instruction. Interrupt requests may be either internal or external. An interrupt request can be induced by the assertion of a signal not directly related to any particular processor or memory state, for example, the assertion of an "I/O done" signal.

A disrupting trap related to an earlier instruction causing an exception is similar to a deferred trap in that it occurs after instructions following the trap inducing instruction have modified the processor or memory state. The difference is that the condition that caused the instruction to induce the disrupting trap may lead to unrecoverable errors, since the implementation may not preserve the necessary state. An example is an ECC data access error reported after the corresponding load instruction has completed.

Disrupting trap conditions should persist until the corresponding trap is taken.

Among the actions that trap-handler software might take after a disrupting trap are the following:

- Use `RETRY` to return to the instruction at which the trap was invoked (`PC ← old PC, nPC ← old nPC`).

- Terminate the program or process associated with the trap.

## 12.2.4 Reset Traps

A *reset trap* occurs when supervisor software or the implementation's hardware determines that the machine must be reset to a known state. Reset traps differ from disrupting traps in that trap handler software for resets is never expected to resume execution of the program that was running when the reset trap occurred.

The following reset traps are defined for the SPARC V9 architecture:

- **Software-initiated reset (SIR)** — Initiated by software by executing the SIR instruction.
- **Power-on reset (POR)** — Initiated when power is applied (or reapplied) to the processor.
- **Watchdog reset (WDR)** — Initiated in response to entry into `error_state`.
- **Externally initiated reset (XIR)** — Initiated in response to an external signal. This reset trap is normally used for critical system events, such as power failure.

## 12.2.5 Uses of the Trap Categories

The SPARC V9 *trap model* makes the following stipulations:

1. Reset traps, except *software\_initiated\_reset* traps, occur asynchronously to program execution.
2. When recovery from an exception can affect the interpretation of subsequent instructions, such exceptions shall be precise. These exceptions include:
  - *software\_initiated\_reset*
  - *instruction\_access\_exception*
  - *privileged\_action*
  - *privileged\_opcode*
  - *trap\_instruction*
  - *instruction\_access\_error*
  - *clean\_window*
  - *fp\_disabled*
  - *LDDF\_mem\_address\_not\_aligned*
  - *STDF\_mem\_address\_not\_aligned*
  - *tag\_overflow*
  - *spill\_n\_normal*
  - *spill\_n\_other*

- *fill\_n\_normal*
  - *fill\_n\_other*
3. An exception caused after the initial access of a multiple access load or store instruction (load/store doubleword, block load, block store, LDSTUB, CASA, CASXA, or SWAP) that causes a catastrophic exception may be precise, deferred, or disrupting. Thus, a trap due to the second memory access can occur after the processor or memory state has been modified by the first access.
  4. Implementation-dependent catastrophic exceptions may cause precise, deferred, or disrupting traps.
  5. Exceptions caused by external events unrelated to the instruction stream, such as interrupts, are disrupting.

A deferred trap may occur one or more instructions after the trap inducing instruction is dispatched.

---

## 12.3 Trap Control

Several registers control how any given trap is processed:

- The IE field in PSTATE and the processor interrupt level PIL register control interrupt processing.
- The enable floating-point unit (FEF) field in FPRS, the floating-point unit enable (PEF) field in PSTATE, and the trap enable mask (TEM) in the FSR control floating-point traps.
- The TL register, which contains the current level of trap nesting, controls whether a trap causes entry to `execute_state`, `RED_state`, or `error_state`.
- PSTATE.TLE determines whether implicit data accesses in the trap routine will be performed with the big-endian or little-endian byte order.

### 12.3.1 PIL Control

Between the execution of instructions, the IU prioritizes the outstanding exceptions and interrupt requests. At any given time, only the highest priority exception or interrupt request is taken as a trap. When there are multiple outstanding exceptions or interrupt requests, the SPARC V9 architecture assumes that lower priority interrupt requests will persist and lower priority exceptions will recur if an exception causing instruction is re-executed.

For interrupt requests, the IU compares the interrupt request level against the PIL register. If the interrupt request level is greater than PIL, then the processor takes the interrupt request trap, assuming there are no higher-priority exceptions outstanding.



## 12.3.2 TEM Control

The occurrence of floating-point traps of type *IEEE\_754\_exception* can be controlled with the user accessible TEM field of the FSR. If a particular bit of TEM is one, the associated *IEEE\_754\_exception* can cause a *fp\_exception\_ieee\_754* trap.

If a particular bit of TEM is zero, the associated *IEEE\_754\_exception* does not cause a *fp\_exception\_ieee\_754* trap. Instead, the occurrence of the exception is recorded in the FSR's accrued exception field (aexc).

If an *IEEE\_754\_exception* results in a *fp\_exception\_ieee\_754* trap, then the destination f register, fccn, and aexc fields remain unchanged. However, if an *IEEE\_754\_exception* does not result in a trap, then the f register, fccn, and aexc fields are updated to their new values.

---

## 12.4 Trap-Table Entry Addresses

Privileged software initializes the trap base address (TBA) register to the upper 49 bits of the trap-table base address. Bit 14 of the vector address (the TL>0 field) is set based on the value of TL at the time the trap is taken, that is, to zero if TL = 0 and to one if TL > 0. Bits 13–5 of the trap vector address are the contents of the TT register. The lowest five bits of the trap address, bits 4–0, are always zero (hence, each trap-table entry is at least 2<sup>5</sup> or 32 bytes long). FIGURE 12-2 illustrates the trap vector address.




FIGURE 12-2 Trap Vector Address

### 12.4.1 Trap Table Organization

The trap table layout is illustrated in FIGURE 12-3.

| Value of TL<br>Before the Trap | Trap Table Contents     | Trap Type                            |
|--------------------------------|-------------------------|--------------------------------------|
| TL = 0                         | Hardware traps          | 000 <sub>16</sub> –07F <sub>16</sub> |
|                                | <i>Spill/fill</i> traps | 080 <sub>16</sub> –0FF <sub>16</sub> |
|                                | Software traps          | 100 <sub>16</sub> –7F <sub>16</sub>  |
|                                | <i>Reserved</i>         | 180 <sub>16</sub> –1FF <sub>16</sub> |
| TL > 0                         | Hardware traps          | 200 <sub>16</sub> –27F <sub>16</sub> |
|                                | <i>Spill/fill</i> traps | 280 <sub>16</sub> –2FF <sub>16</sub> |
|                                | Software traps          | 300 <sub>16</sub> –37F <sub>16</sub> |
|                                | <i>Reserved</i>         | 380 <sub>16</sub> –3FF <sub>16</sub> |

**FIGURE 12-3** Trap Table Layout

The trap table for TL = 0 comprises 512 32-byte entries; the trap table for TL > 0 comprises 512 more 32-byte entries. Therefore, the total size of a full trap table is 512 × 32 × 2, or 32 KB. However, if privileged software does not use software traps (TCC instructions) at TL > 0, the table can be made 24 KB long.

## 12.4.2 Trap Type (TT)

When a normal trap occurs, a value that uniquely identifies the trap is written into the current 9-bit TT register (TT[TL]) by hardware. Control is then transferred into the trap table to an address formed by the TBA register (TL > 0) and TT[TL]. The lowest five bits of the address are always zero; each entry in the trap table may contain the first eight instructions of the corresponding trap handler.

---

**Programming Notes** – The trap type for the *clean\_window* exception is 024<sub>16</sub>. Three subsequent trap vectors (025<sub>16</sub>–027<sub>16</sub>) are reserved to allow for an inline (branchless) trap handler. Three subsequent trap vectors are reserved for each *spill/fill* vector, to allow for an inline (branchless) trap handler.

The *spill/fill*, *clean\_window*, and MMU related traps (*fast\_instruction\_access\_MMU\_miss*, *fast\_data\_access\_MMU\_miss*, and *fast\_data\_access\_protection*) trap types are spaced such that their trap-table entries are 128 bytes (32 instructions) long in the UltraSPARC III Cu processor. This length allows the complete code for one *spill/fill* routine, a *clean\_window* routine, or a normal MMU miss handling routine to reside in one trap-table entry.

---

When a special trap occurs, the TT register is set as described in “RED\_state” on page 317. Control is then transferred into the RED\_state trap table to an address formed by the RSTVaddr and an offset depending on the condition.

TT values  $000_{16}$ – $0FF_{16}$  are reserved for hardware traps. TT values  $100_{16}$ – $17F_{16}$  are reserved for software traps (traps caused by execution of a TCC instruction). TT values  $180_{16}$ – $1FF_{16}$  are reserved for future uses.

TT values  $070_{16}$ ,  $071_{16}$  and  $072_{16}$  are reserved for fast\_ECC\_error, dcache\_parity\_error and icache\_parity\_error, respectively.

The assignment of TT values to traps is shown in TABLE 12-3. Hardware must detect and trap these exceptions and interrupts and must set the defined TT values. In the table, AG = alternate globals, MG = MMU globals, and IG = interrupt globals.

**TABLE 12-3** Exception and Interrupt Requests, by TT Value (1 of 3)

| Exception or Interrupt Request  | TT                      | Global Register Set | Priority  |
|---|-------------------------|---------------------|-----------|
| <i>Reserved</i>   | $000_{16}$              |                     | <i>NA</i> |
| <i>power_on_reset</i>   | $001_{16}$              | AG                  | 0         |
| <i>watchdog_reset</i>   | $002_{16}$              | AG                  | 1         |
| <i>externally_initiated_reset</i>   | $003_{16}$              | AG                  | 1         |
| <i>software_initiated_reset</i>   | $004_{16}$              | AG                  | 1         |
| <i>RED_state_exception</i>  | $005_{16}$              | AG                  | 1         |
| <i>Reserved</i>   | $006_{16}$ – $007_{16}$ |                     | <i>NA</i> |
| <i>instruction_access_exception</i>                                       | $008_{16}$              | MG                  | 5         |
| <i>instruction_access_MMU_miss (Not used. Replaced by newer version.)</i> | $008_{16}$              |                     | <i>NA</i> |
| <i>instruction_access_error</i>   | $00A_{16}$              | AG                  | 3         |
| <i>Reserved</i>   | $00B_{16}$ – $00F_{16}$ |                     | <i>NA</i> |
| <i>illegal_instruction</i>  | $010_{16}$              | AG                  | 7         |
| <i>privileged_opcode</i>  | $011_{16}$              | AG                  | 6         |
| <i>Reserved</i>   | $012_{16}$              |                     | <i>NA</i> |
| <i>Reserved</i>   | $013_{16}$              |                     | <i>NA</i> |
| <i>Reserved</i>   | $014_{16}$ – $01F_{16}$ |                     | <i>NA</i> |
| <i>fp_disabled</i>  | $020_{16}$              | AG                  | 8         |
| <i>fp_exception_ieee_754</i>  | $021_{16}$              | AG                  | 11        |
| <i>fp_exception_other</i>   | $022_{16}$              | AG                  | 11        |
| <i>tag_overflow</i>   | $023_{16}$              | AG                  | 14        |
| <i>clean_window</i>   | $024_{16}$ – $027_{16}$ | AG                  | 10        |
| <i>division_by_zero</i>   | $028_{16}$              | AG                  | 15        |

**TABLE 12-3** Exception and Interrupt Requests, by TT Value (2 of 3)

| Exception or Interrupt Request                                      | TT                                   | Global Register Set | Priority |
|---|--------------------------------------|---------------------|----------|
| <i>Reserved</i>   | 029 <sub>16</sub> –02F <sub>16</sub> |                     | NA       |
| <i>data_access_exception</i>  | 030 <sub>16</sub>                    | MG                  | 12       |
| <i>data_access_MMU_miss</i> (Not used; replaced by newer version)   | 031 <sub>16</sub>                    |                     | NA       |
| <i>data_access_error</i>  | 032 <sub>16</sub>                    | AG                  | 12       |
| <i>data_access_protection</i> (Not used; replaced by newer version) | 033 <sub>16</sub>                    |                     | 12       |
| <i>mem_address_not_aligned</i>                                      | 034 <sub>16</sub>                    | AG                  | 10       |
| <i>LDDF_mem_address_not_aligned</i>                                 | 035 <sub>16</sub>                    | AG                  | 10       |
| <i>STDF_mem_address_not_aligned</i>                                 | 036 <sub>16</sub>                    | AG                  | 10       |
| <i>privileged_action</i>  | 037 <sub>16</sub>                    | AG                  | 11       |
| <i>Reserved</i>   | 038 <sub>16</sub>                    |                     | NA       |
| <i>Reserved</i>   | 039 <sub>16</sub>                    |                     | NA       |
| <i>Reserved</i>   | 03A <sub>16</sub> –03F <sub>16</sub> |                     | NA       |
| <i>Reserved</i>   | 040 <sub>16</sub>                    |                     | NA       |
| <i>interrupt_level_n</i> (n = 1–15)                                 | 041 <sub>16</sub> –04F <sub>16</sub> | AG                  | 32-n     |
| <i>Reserved</i>   | 050 <sub>16</sub> –05F <sub>16</sub> |                     | NA       |
| <i>interrupt_vector</i>   | 060 <sub>16</sub>                    | IG                  | 16       |
| <i>PA_watchpoint</i>  | 061 <sub>16</sub>                    | AG                  | 12       |
| <i>VA_watchpoint</i>  | 062 <sub>16</sub>                    | AG                  | 11       |
| <i>ECC_error</i>  | 063 <sub>16</sub>                    | AG                  | 33       |
| <i>fast_instruction_access_MMU_miss</i>                             | 064 <sub>16</sub> –067 <sub>16</sub> | MG                  | 2        |
| <i>fast_data_access_MMU_miss</i>                                    | 068 <sub>16</sub> –06B <sub>16</sub> | MG                  | 12       |
| <i>fast_data_access_protection</i>                                  | 06C <sub>16</sub> –06F <sub>16</sub> | MG                  | 12       |
| <i>fast_ECC_error</i>   | 070 <sub>16</sub>                    | AG                  | 2        |
| <i>dcache_parity_error</i>  | 071 <sub>16</sub>                    | AG                  | 2        |
| <i>icache_parity_error</i>  | 072 <sub>16</sub>                    | AG                  | 2        |
| <i>spill_n_normal</i> (n = 0–7)                                     | 080 <sub>16</sub> –09F <sub>16</sub> | AG                  | 9        |
| <i>spill_n_other</i> (n = 0–7)                                      | 0A0 <sub>16</sub> –0BF <sub>16</sub> | AG                  | 9        |
| <i>fill_n_normal</i> (n = 0–7)                                      | 0C0 <sub>16</sub> –0DF <sub>16</sub> | AG                  | 9        |

**TABLE 12-3** Exception and Interrupt Requests, by TT Value (3 of 3)

| Exception or Interrupt Request    | TT                                   | Global Register Set | Priority |
|-----------------------------------|--------------------------------------|---------------------|----------|
| <i>fill_n_other</i> ( $n = 0-7$ ) | 0E0 <sub>16</sub> -0FF <sub>16</sub> | AG                  | 9        |
| <i>trap_instruction</i>           | 100 <sub>16</sub> -17F <sub>16</sub> | AG                  | 16       |
| <i>Reserved</i>                   | 180 <sub>16</sub> -1FF <sub>16</sub> |                     | NA       |

TABLE 12-4 lists the traps in priority order. Reserved and unused traps have been left out of this table.

**TABLE 12-4** Exception and Interrupts Requests in Priority Order

| Exception or Interrupt Request          | TT                                   | Global Register Set | Priority |
|---|--------------------------------------|---------------------|----------|
| <i>power_on_reset</i>                   | 001 <sub>16</sub>                    | AG                  | 0        |
| <i>watchdog_reset</i>                   | 002 <sub>16</sub>                    | AG                  | 1        |
| <i>externally_initiated_reset</i>       | 003 <sub>16</sub>                    | AG                  | 1        |
| <i>software_initiated_reset</i>         | 004 <sub>16</sub>                    | AG                  | 1        |
| <i>RED_state_exception</i>              | 005 <sub>16</sub>                    | AG                  | 1        |
| <i>fast_instruction_access_MMU_miss</i> | 064 <sub>16</sub> -067 <sub>16</sub> | MG                  | 2        |
| <i>fast_ECC_error</i>                   | 070 <sub>16</sub>                    | AG                  | 2        |
| <i>dcache_parity_error</i>              | 071 <sub>16</sub>                    | AG                  | 2        |
| <i>icache_parity_error</i>              | 072 <sub>16</sub>                    | AG                  | 2        |
| <i>instruction_access_error</i>         | 00A <sub>16</sub>                    | AG                  | 3        |
| <i>instruction_access_exception</i>     | 008 <sub>16</sub>                    | MG                  | 5        |
| <i>privileged_opcode</i>                | 011 <sub>16</sub>                    | AG                  | 6        |
| <i>illegal_instruction</i>              | 010 <sub>16</sub>                    | AG                  | 7        |
| <i>fp_disabled</i>                      | 020 <sub>16</sub>                    | AG                  | 8        |
| <i>spill_n_normal</i> ( $n = 0-7$ )     | 080 <sub>16</sub> -09F <sub>16</sub> | AG                  | 9        |
| <i>spill_n_other</i> ( $n = 0-7$ )      | 0A0 <sub>16</sub> -0BF <sub>16</sub> | AG                  | 9        |
| <i>fill_n_normal</i> ( $n = 0-7$ )      | 0C0 <sub>16</sub> -0DF <sub>16</sub> | AG                  | 9        |
| <i>fill_n_other</i> ( $n = 0-7$ )       | 0E0 <sub>16</sub> -0FF <sub>16</sub> | AG                  | 9        |
| <i>clean_window</i>                     | 024 <sub>16</sub> -027 <sub>16</sub> | AG                  | 10       |
| <i>mem_address_not_aligned</i>          | 034 <sub>16</sub>                    | AG                  | 10       |
| <i>LDDF_mem_address_not_aligned</i>     | 035 <sub>16</sub>                    | AG                  | 10       |

**TABLE 12-4** Exception and Interrupts Requests in Priority Order (Continued)

| Exception or Interrupt Request              | TT                                   | Global Register Set | Priority     |
|---|--------------------------------------|---------------------|--------------|
| <i>STDF_mem_address_not_aligned</i>         | 036 <sub>16</sub>                    | AG                  | 10           |
| <i>fp_exception_ieee_754</i>                | 021 <sub>16</sub>                    | AG                  | 11           |
| <i>fp_exception_other</i>                   | 022 <sub>16</sub>                    | AG                  | 11           |
| <i>privileged_action</i>                    | 037 <sub>16</sub>                    | AG                  | 11           |
| <i>VA_watchpoint</i>                        | 062 <sub>16</sub>                    | AG                  | 11           |
| <i>data_access_exception</i>                | 030 <sub>16</sub>                    | MG                  | 12           |
| <i>data_access_error</i>                    | 032 <sub>16</sub>                    | AG                  | 12           |
| <i>PA_watchpoint</i>                        | 061 <sub>16</sub>                    | AG                  | 12           |
| <i>fast_data_access_MMU_miss</i>            | 068 <sub>16</sub> –06B <sub>16</sub> | MG                  | 12           |
| <i>fast_data_access_protection</i>          | 06C <sub>16</sub> –06F <sub>16</sub> | MG                  | 12           |
| <i>tag_overflow</i>                         | 023 <sub>16</sub>                    | AG                  | 14           |
| <i>division_by_zero</i>                     | 028 <sub>16</sub>                    | AG                  | 15           |
| <i>trap_instruction</i>                     | 100 <sub>16</sub> –17F <sub>16</sub> | AG                  | 16           |
| <i>interrupt_vector</i>                     | 060 <sub>16</sub>                    | IG                  | 16           |
| <i>interrupt_level_n</i> ( <i>n</i> = 1–15) | 041 <sub>16</sub> –04F <sub>16</sub> | AG                  | 32- <i>n</i> |
| <i>ECC_error</i>                            | 063 <sub>16</sub>                    | AG                  | 33           |

### 12.4.2.1 Trap Type for Spill/Fill Traps

The trap type for window *spill/fill* traps is determined on the basis of the contents of the OTHERWIN and WSTATE registers as described in TABLE 12-5 and shown in FIGURE 12-4.

**TABLE 12-5** Trap Types for *Spill/Fill* Traps

| Bit | Field         | Description  |
|-----|---------------|--|
| 8:6 | SPILL_OR_FILL | 010 <sub>2</sub> for <i>spill</i> traps; 011 <sub>2</sub> for <i>fill</i> trap |
| 5   | OTHER         | (OTHERWIN ≠ 0)   |
| 4:2 | WTYPE         | If (OTHER) then WSTATE . OTHER; else WSTATE . NORMAL                           |




FIGURE 12-4 Trap Type Encoding for *Spill/Fill* Traps

## 12.4.3 Trap Priorities

TABLE 12-3 shows the assignment of traps to TT values and the relative priority of traps and interrupt requests. Priority 0 is highest, priority 31 is lowest; that is, if  $X < Y$ , a pending exception or interrupt request with priority  $X$  is taken instead of a pending exception or interrupt request with priority  $Y$ .

However, the TT values for the exceptions and interrupt requests shown in TABLE 12-3 must remain the same for every implementation.

The trap priorities given always need to be considered in light of how the processor actually issues and executes instructions. For example, if an *instruction\_access\_error* occurs (priority 3), it will be taken even if the instruction was an SIR (priority 1). This situation occurs because the processor gets the *instruction\_access\_error* during I-fetch and never actually issues or executes the instruction, so the SIR is never seen by the backend of the processor. This is an obvious case, but there are other more subtle cases.

In summary, the trap priorities are used to prioritize traps that occur in the same clock cycle. They do not take into consideration that an instruction may be alive for multiple cycles and that a trap may be detected and initiated early in the life of an instruction. Once the early trap is taken, any errors that might have occurred later in the instruction's life will not be seen.

## 12.4.4 Details of Supported Traps

### 12.4.4.1 MMU Traps

The UltraSPARC III Cu processor supports three 32-instruction traps for handling the most performance sensitive MMU traps:

- *fast\_instruction\_access\_MMU\_miss*
- *fast\_data\_access\_MMU\_miss*
- *fast\_data\_access\_protection*

The first two traps are taken when the TLBs miss on an instruction or data access. The third type of trap is taken when a protection violation occurs. The common case of this trap occurs when a write request is made to a page marked as clean in the TLB.

Each of these trap vectors takes up four slots in the trap table; this means that each trap handler can contain up to 32 instructions before a branch is needed.

## 12.4.4.2 Precise Correctable Data Corruption Error Traps

The UltraSPARC III Cu processor implements three error traps listed in TABLE 12-6.

**TABLE 12-6** Two New Traps Added for D-cache and I-cache Parity Error

| Exception or Interrupt Request | Globals | TT    | Priority |
|--------------------------------|---------|-------|----------|
| <i>fast_ECC_error</i>          | AG      | 0x070 | 2        |
| <i>dcache_parity_error</i>     | AG      | 0x071 | 2        |
| <i>icache_parity_error</i>     | AG      | 0x072 | 2        |

### *D-Cache and I-Cache Parity Errors*

Upon a *dcache\_parity\_error* trap taken, hardware will automatically turn off D-cache and I-cache by clearing the DC and IC bit in the DCU Control Register. The same action is done on *icache\_parity\_error*. The objective is to avoid any possible recursive/nested parity error of the same type, an event that is very complicated to handle by the trap handler and likely to reduce software ability to repair and recover from the parity error.

### *Fast ECC Errors*

On *fast\_ECC\_error* detection during D-cache load miss fill, D-cache will still install the uncorrected data. But since the *fast\_ECC\_error* trap is precise, hardware can rely on software to help clean up the bad data. I-cache is different. If I-cache is filled with errors, the line will not be installed in I-cache.

A D-cache or I-cache miss request may observe an ECC error in the line it reads from the L2-cache. When this occurs, a *fast\_ECC\_error* precise trap is generated for the instruction that detected the error. In the case of a D-cache request, the corrupted data will be installed in the D-cache, but the trap takes effect before the data can be used.

In case of an I-cache request, the data from the L2-cache will be corrected by hardware before being installed in the I-cache. When the *fast\_ECC\_error* trap is taken, the I-cache and D-cache are left enabled. Software is responsible to disable any caches as a part of the recovery. Software must flush the corrupted line from the D-cache if it was filled.



## *Prioritization of Errors*

All traps with priority 2 are of precise type. Miss/error traps with priority 2 that arrive at the same time are processed by hardware according to their “age” or program order. The oldest instruction with an error/miss will get the trap.

However, there are some cases where the same instruction (same PC) generate multiple traps. There are two cases:

1. **Case 1:** Singular trap type with highest priority.  
The processing order of which trap to take first follows the priority number (lowest number wins).
2. **Case 2:** Multiple traps having same highest priority.  
For trap priority 2, the only possible combination is simultaneous traps due to I-cache parity error and I-TLB miss. In this case, the hardware processing order is as follows:

*icache\_parity\_error > fast\_instruction\_access\_MMU\_miss*

There are no other “simultaneous traps” combinations of Case 2. The other priority 2 traps have later arrivals. D-cache access is further down the pipeline after instruction fetch from I-cache. Thus, D-cache parity error on a load instruction (if any) will be detected after I-cache parity error (if any) and I-TLB miss (if any). The other priority 2 trap, *fast\_ECC\_error* can only be caused by an I-cache miss or D-cache load miss; therefore, it arrives even later.

To summarize, precise traps are processed in the following order:

*program order > trap priority number > Hardware implementation order*

One exception to the trap priority number order in implementation is on a non-precise ECC error trap. If the victim instruction selected as the trap point *also* has a higher priority precise trap, then the ECC error trap is taken first. This should be okay since no trap is lost. When the ECC error trap handling finishes and the victim instruction is retried, the precise trap is encountered again.

### 12.4.4.3 Other Traps

The UltraSPARC III Cu processor supports the following trap types in addition to those in SPARC V9:

- *interrupt\_vector\_trap*
- *PA\_watchpoint*
- *VA\_watchpoint*
- *ECC\_error*
- *data\_access\_protection*

## 12.5 Trap Processing

The processor's action during trap processing depends on the trap type, the current level of trap nesting (given in the TL register), and the processor state. When a trap occurs, the global registers are replaced with one of three sets of trap global register — MMU globals, interrupt globals, or alternate globals — based on the type of trap.

All traps use normal trap processing, except those due to reset requests, catastrophic errors, traps taken when  $TL = MAXTL - 1$ , and traps taken when the processor is in `RED_state`. These traps use special `RED_state` trap processing.

During normal operation, the processor is in `execute_state`. It processes traps in `execute_state` and continues.

When a normal trap or software-initiated reset (SIR) occurs with  $TL = MAXTL$ , there are no more levels on the trap stack, so the processor enters `error_state` and halts. To avoid this catastrophic failure, SPARC V9 provides the `RED_state` processor state. Traps processed in `RED_state` use a special trap vector and a special trap-vectoring algorithm.

Traps that occur with  $TL = MAXTL - 1$  are processed in `RED_state`. In addition, reset traps are also processed in `RED_state`. Finally, supervisor software can force the processor into `RED_state` by setting the `PSTATE.RED` flag to one.

Once the processor has entered `RED_state`, no matter how it got there, all subsequent traps are processed in `RED_state` until software returns the processor to `execute_state` or a normal or SIR trap is taken when  $TL = MAXTL$ , which puts the processor in `error_state`. TABLE 12-7, TABLE 12-8, and TABLE 12-9 describe the processor mode and trap-level transitions involved in handling traps.

**TABLE 12-7** Trap Received While in `execute_state`

| Original State  | New State, After Receiving Trap Type                 |  |  |  |
|---|--|--|--|--|
|   | Normal Trap or Interrupt                             | POR                                    | XIR, WDR   | SIR  |
| <code>execute_state</code><br>$TL < MAXTL - 1$          | <code>execute_state</code><br>$TL \leftarrow TL + 1$ | <code>RED_state</code><br>$TL = MAXTL$ | <code>RED_state</code><br>$TL \leftarrow TL + 1$ | <code>RED_state</code><br>$TL \leftarrow TL + 1$ |
| <code>execute_state</code><br>$TL = MAXTL - 1$          | <code>RED_state</code><br>$TL = MAXTL$               | <code>RED_state</code><br>$TL = MAXTL$ | <code>RED_state</code><br>$TL = MAXTL$           | <code>RED_state</code><br>$TL = MAXTL$           |
| <code>execute_state</code> <sup>†</sup><br>$TL = MAXTL$ | <code>error_state</code><br>$TL = MAXTL$             | <code>RED_state</code><br>$TL = MAXTL$ | <code>RED_state</code><br>$TL = MAXTL$           | <code>error_state</code><br>$TL = MAXTL$         |

<sup>†</sup>This state occurs when software changes TL to MAXTL and does not set `PSTATE.RED`, or if it clears `PSTATE.RED` while at MAXTL.

**TABLE 12-8** Trap Received While in RED\_state

| Original State              | New State, After Receiving Trap Type |                         |                          |                           |
|-----------------------------|--------------------------------------|-------------------------|--------------------------|---------------------------|
|                             | Normal Trap or Interrupt             | POR                     | XIR, WDR                 | SIR                       |
| RED_state<br>TL < MAXTL - 1 | RED_state<br>TL ← TL + 1             | RED_state<br>TL = MAXTL | RED_state<br>TL ← TL + 1 | RED_state<br>TL ← TL + 1  |
| RED_state<br>TL = MAXTL - 1 | RED_state<br>TL = MAXTL              | RED_state<br>TL = MAXTL | RED_state<br>TL = MAXTL  | RED_state<br>TL = MAXTL   |
| RED_state<br>TL = MAXTL     | error_state<br>TL = MAXTL            | RED_state<br>TL = MAXTL | RED_state<br>TL = MAXTL  | error_state<br>TL = MAXTL |

**TABLE 12-9** Reset Received While in error\_state

| Original State                | New State, After Receiving Trap Type |                         |                          |     |
|-------------------------------|--------------------------------------|-------------------------|--------------------------|-----|
|                               | Normal Trap or Interrupt             | POR                     | XIR, WDR                 | SIR |
| error_state<br>TL < MAXTL - 1 | —                                    | RED_state<br>TL = MAXTL | RED_state<br>TL ← TL + 1 | —   |
| error_state<br>TL = MAXTL - 1 | —                                    | RED_state<br>TL = MAXTL | RED_state<br>TL = MAXTL  | —   |
| error_state<br>TL = MAXTL     | —                                    | RED_state<br>TL = MAXTL | RED_state<br>TL = MAXTL  | —   |

---

**Implementation Note** – The processor does not recognize interrupts while it is in error\_state.

---

## 12.5.1 Normal Trap Processing

A trap other than a fast MMU trap (see Section 12.5.2, “Fast MMU Trap Processing”) or an interrupt vector trap (see Section 12.5.3, “Interrupt Vector Trap Processing”) causes the following state changes to occur<sup>1</sup>:

- If the processor is already in RED\_state, the new trap is processed in RED\_state unless TL = MAXTL.

1. Please note that these state transitions include the minimum set covered in the SPARC V9 architecture. There are other non-SPARC V9 transitions that are implementation-dependent and are not described here.

- If the processor is in `execute_state` and the trap level is one less than its maximum value, that is,  $TL = MAXTL - 1$ , then the processor enters `RED_state`.
- If the processor is in either `execute_state` or `RED_state` and the trap level is already at its maximum value, that is,  $TL = MAXTL$ , then the processor enters `error_state`.

Otherwise, the trap uses normal trap processing, and the following state changes occur:

- The trap level is set. This provides access to a fresh set of privileged trap-state registers used to save the current state, in effect, pushing a frame on the trap stack.

$$TL \leftarrow TL + 1$$

- Existing state is preserved.

```
TSTATE[TL].CCR    ← CCR
TSTATE[TL].ASI    ← ASI
TSTATE[TL].PSTATE ← PSTATE
TSTATE[TL].CWP    ← CWP
TPC[TL]           ← PC
TNPC[TL]          ← nPC
```

- The trap type is preserved.

```
TT[TL]           ← the trap type
```

- The `PSTATE` register is updated to a predefined state.

```
PSTATE.MM        is unchanged
PSTATE.RED       ← 0
PSTATE.PEF       ← 1 (FPU is present)
PSTATE.AM        ← 0 (address masking is turned off)
PSTATE.PRIV      ← 1 (the processor enters privileged mode)
PSTATE.IE        ← 0 (interrupts are disabled)
PSTATE.AG        ← 1 (global registers are replaced with alternate globals)
PSTATE.MG        ← 0 (MMU globals are disabled)
PSTATE.IG        ← 0 (interrupt globals are disabled)
PSTATE.CLE       ← PSTATE.TLE (set endian mode for traps)
PSTATE.TLE       is unchanged
```

- For a register-window trap only, `CWP` is set to point to the register window that must be accessed by the trap-handler software, that is:
  - If  $TT[TL] = 024_{16}$  (a *clean\_window* trap), then  $CWP \leftarrow CWP + 1$ .
  - If  $080_{16} \leq TT[TL] \leq 0BF_{16}$  (window *spill* trap), then  $CWP \leftarrow CWP + CANSERVE + 2$ .
  - If  $0C0_{16} \leq TT[TL] \leq 0FF_{16}$  (window *fill* trap), then  $CWP \leftarrow CWP - 1$ .

For non-register-window traps, `CWP` is not changed.

- Control is transferred into the trap table:

```
PC    ← TBA<63:15> [ (TL > 0) ] TT[TL] [ 0 0000
nPC   ← TBA<63:15> [ (TL > 0) ] TT[TL] [ 0 0100
```

where “(TL > 0)” is zero if  $TL = 0$ , and one if  $TL > 0$ .

Interrupts are ignored as long as `PSTATE.IE = 0`.

---

**Programming Note** – State in `TPC[n]`, `TNPC[n]`, `TSTATE[n]`, and `TT[n]` is only changed autonomously by the processor when a trap is taken while `TL = n - 1`; however, software can change any of these values with a `WRPR` instruction when `TL = n`.

---

## 12.5.2 Fast MMU Trap Processing

Fast MMU traps (*fast\_instruction\_access\_MMU\_miss*, *fast\_data\_access\_MMU\_miss*, and *fast\_data\_access\_protection*) cause the following state changes to occur<sup>1</sup>:

- If the processor is already in `RED_state`, the new trap is processed in `RED_state` unless `TL = MAXTL`.
- If the processor is in `execute_state` and the trap level is one less than its maximum value, that is, `TL = MAXTL - 1`, then the processor enters `RED_state`.
- If the processor is in either `execute_state` or `RED_state` and the trap level is already at its maximum value, that is, `TL = MAXTL`, then the processor enters `error_state`.

Otherwise, the trap uses normal trap processing, and the following state changes occur:

- The trap level is set. This provides access to a fresh set of privileged trap-state registers used to save the current state, in effect, pushing a frame on the trap stack.

`TL ← TL + 1`

- Existing state is preserved:

`TSTATE[TL].CCR`      ← `CCR`  
`TSTATE[TL].ASI`      ← `ASI`  
`TSTATE[TL].PSTATE` ← `PSTATE`  
`TSTATE[TL].CWP`      ← `CWP`  
`TPC[TL]`              ← `PC`  
`TNPC[TL]`             ← `nPC`

- The trap type is preserved.

`TT[TL]`                ← the trap type

- The `PSTATE` register is updated to a predefined state.

`PSTATE.MM`            is unchanged  
`PSTATE.RED`           ← 0  
`PSTATE.PEF`          ← 1 (FPU is present)  
`PSTATE.AM`           ← 0 (address masking is turned off)  
`PSTATE.PRIV`        ← 1 (the processor enters privileged mode)  
`PSTATE.IE`           ← 0 (interrupts are disabled)  
`PSTATE.AG`           ← 0 (alternate globals are disabled)

<sup>1</sup> Please note that these state transitions include the minimum set covered in the SPARC V9 architecture. There are other non-SPARC V9 transitions that are implementation-dependent and are not described here.

|            |  |
|------------|--|
| PSTATE.MG  | ← 1 (global registers are replaced with MMU globals) |
| PSTATE.IG  | ← 0 (interrupt globals are disabled)                 |
| PSTATE.CLE | ← PSTATE.TLE (set endian mode for traps)             |
| PSTATE.TLE | is unchanged   |

- For non-register-window traps, CWP is not changed.

- Control is transferred into the trap table:

|     |              |              |        |          |
|-----|--------------|--------------|--------|----------|
| PC  | ← TBA<63:15> | [ (TL > 0) ] | TT[TL] | [ 0 0000 |
| nPC | ← TBA<63:15> | [ (TL > 0) ] | TT[TL] | [ 0 0100 |

where “(TL > 0)” is zero if TL = 0, and one if TL > 0.

Interrupts are ignored as long as PSTATE.IE = 0.

---

**Programming Note** – State in TPC[n], TNPC[n], TSTATE[n], and TT[n] is only changed autonomously by the processor when a trap is taken while TL = n – 1; however, software can change any of these values with a WRPR instruction when TL = n.

---

## 12.5.3 Interrupt Vector Trap Processing

An *interrupt\_vector* trap causes the following state changes to occur<sup>1</sup>:

- If the processor is already in RED\_state, the new trap is processed in RED\_state unless TL = MAXTL.
- If the processor is in execute\_state and the trap level is one less than its maximum value, that is, TL = MAXTL – 1, the processor enters RED\_state.
- If the processor is in either execute\_state or RED\_state and the trap level is already at its maximum value, that is, TL = MAXTL, then the processor enters error\_state.

Otherwise, the trap uses normal trap processing, and the following state changes occur:

- The trap level is set. This provides access to a fresh set of privileged trap-state registers used to save the current state, in effect, pushing a frame on the trap stack.

|    |          |
|----|----------|
| TL | ← TL + 1 |
|----|----------|

- Existing state is preserved.

|                   |          |
|-------------------|----------|
| TSTATE[TL].CCR    | ← CCR    |
| TSTATE[TL].ASI    | ← ASI    |
| TSTATE[TL].PSTATE | ← PSTATE |
| TSTATE[TL].CWP    | ← CWP    |
| TPC[TL]           | ← PC     |
| TNPC[TL]          | ← nPC    |

- The trap type is preserved.

<sup>1</sup> Please note that these state transitions include the minimum set covered in the SPARC V9 architecture. There are other non-SPARC V9 transitions that are implementation-dependent and are not described here.

TT[TL] ← the trap type

- The PSTATE register is updated to a predefined state.

PSTATE.MM is unchanged  
PSTATE.RED ← 0  
PSTATE.PEF ← 1 (FPU is present)  
PSTATE.AM ← 0 (address masking is turned off)  
PSTATE.PRIV ← 1 (the processor enters privileged mode)  
PSTATE.IE ← 0 (interrupts are disabled)  
PSTATE.AG ← 0 (alternate globals are disabled)  
PSTATE.MG ← 0 (MMU globals are disabled)  
PSTATE.IG ← 1 (global registers are replaced with interrupt globals)  
PSTATE.CLE ← PSTATE.TLE (set endian mode for traps)  
PSTATE.TLE is unchanged

- For non-register-window traps, CWP is not changed.

- Control is transferred into the trap table:

PC ← TBA<63:15> [ (TL > 0) ] TT[TL] [ 0 0000

nPC ← TBA<63:15> [ (TL > 0) ] TT[TL] [ 0 0100

where “(TL > 0)” is zero if TL = 0, and one if TL > 0.

Interrupts are ignored as long as PSTATE.IE = 0.

---

**Programming Note** – State in TPC[n], TNPC[n], TSTATE[n], and TT[n] is only changed autonomously by the processor when a trap is taken while TL = n – 1; however, software can change any of these values with a WRPR instruction when TL = n.

---

## 12.5.4 Special Trap Processing

The following conditions invoke special trap processing:

- Traps taken with TL = MAXTL – 1
- Power-on reset traps
- Watchdog reset traps
- Externally initiated reset traps
- Software-initiated reset traps
- Traps taken when the processor is already in RED\_state

### 12.5.4.1 Normal Traps with TL = MAXTL – 1

Normal traps that occur when TL = MAXTL – 1 are processed in RED\_state. The following state changes occur<sup>1</sup>:

- The trap level is advanced.

$$TL \leftarrow MAXTL$$

- Existing state is preserved.

|                   |          |
|-------------------|----------|
| TSTATE[TL].CCR    | ← CCR    |
| TSTATE[TL].ASI    | ← ASI    |
| TSTATE[TL].PSTATE | ← PSTATE |
| TSTATE[TL].CWP    | ← CWP    |
| TPC[TL]           | ← PC     |
| TNPC[TL]          | ← nPC    |

- The trap type is preserved.

$$TT[TL] \leftarrow \text{the trap type}$$

- The PSTATE register is set as follows:

|             |  |
|-------------|--|
| PSTATE.MM   | ← 00 <sub>2</sub> (TSO)                                    |
| PSTATE.RED  | ← 1 (enter RED_state)                                      |
| PSTATE.PEF  | ← 1 (FPU is present)                                       |
| PSTATE.AM   | ← 0 (address masking is turned off)                        |
| PSTATE.PRIV | ← 1 (the processor enters privileged mode)                 |
| PSTATE.IE   | ← 0 (interrupts are disabled)                              |
| PSTATE.AG   | ← 1 (global registers are replaced with alternate globals) |
| PSTATE.MG   | ← 0 (MMU globals are disabled)                             |
| PSTATE.IG   | ← 0 (interrupt globals are disabled)                       |
| PSTATE.CLE  | ← PSTATE.TLE (set endian mode for traps)                   |
| PSTATE.TLE  | ← undefined <sup>1</sup>                                   |

- For a register-window trap only, CWP is set to point to the register window that must be accessed by the trap-handler software, that is:
  - If  $TT[TL] = 024_{16}$  (a *clean\_window* trap), then  $CWP \leftarrow CWP + 1$ .
  - If  $(080_{16} \leq TT[TL] \leq 0BF_{16})$  (window *spill* trap), then  $CWP \leftarrow CWP + CANSERVE + 2$ .
  - If  $(0C0_{16} \leq TT[TL] \leq 0FF_{16})$  (window *fill* trap), then  $CWP \leftarrow CWP - 1$ .

For non-register-window traps, CWP is not changed.

- Implementation-specific state changes; for example, disabling an MMU.
- Control is transferred into the RED\_state trap table.

|     |                  |   |                        |
|-----|------------------|---|------------------------|
| PC  | ← RSTVaddr<63:8> | □ | 1010 0000 <sub>2</sub> |
| nPC | ← RSTVaddr<63:8> | □ | 1010 0100 <sub>2</sub> |

1. Please note that these state transitions include the minimum set covered in the SPARC V9 architecture. There are other non-SPARC V9 transitions that are implementation-dependent and are not described here.

1. Note that this differs from SPARC V9.



## 12.5.4.2 Power-On Reset (POR) Traps

POR traps occur when power is applied to the processor. If the processor is in `error_state`, a POR brings the processor out of `error_state` and places it in `RED_state`. Processor state is undefined after POR, except for the following<sup>1</sup>:

- The trap level is set.

$$TL \leftarrow MAXTL$$

- The trap type is set.

$$TT[TL] \leftarrow 001_{16}$$

- The `PSTATE` register is set as follows:

|                          |   |
|--------------------------|---|
| <code>PSTATE.MM</code>   | $\leftarrow 00_2$ (TSO)   |
| <code>PSTATE.RED</code>  | $\leftarrow 1$ (enter <code>RED_state</code> )                        |
| <code>PSTATE.PEF</code>  | $\leftarrow 1$ (FPU is present)                                       |
| <code>PSTATE.AM</code>   | $\leftarrow 0$ (address masking is turned off)                        |
| <code>PSTATE.PRIV</code> | $\leftarrow 1$ (the processor enters privileged mode)                 |
| <code>PSTATE.IE</code>   | $\leftarrow 0$ (interrupts are disabled)                              |
| <code>PSTATE.AG</code>   | $\leftarrow 1$ (global registers are replaced with alternate globals) |
| <code>PSTATE.MG</code>   | $\leftarrow 0$ (MMU globals are disabled)                             |
| <code>PSTATE.IG</code>   | $\leftarrow 0$ (interrupt globals are disabled)                       |
| <code>PSTATE.CLE</code>  | $\leftarrow 0$ (big-endian mode for nontraps)                         |
| <code>PSTATE.TLE</code>  | $\leftarrow 0$ (big-endian mode for traps)                            |

- The `TICK` register is protected.

$$TICK.NPT \leftarrow 1 \text{ (TICK unreadable by non-privileged software)}$$

- Implementation-specific state changes; for example, disabling an MMU.
- Control is transferred into the `RED_state` trap table.

$$PC \leftarrow RSTVaddr\langle 63:8 \rangle \square 0010\ 0000_2$$

$$nPC \leftarrow RSTVaddr\langle 63:8 \rangle \square 0010\ 0100_2$$

For any reset when  $TL = MAXTL$ , for all  $n < MAXTL$ , the values in `TPC[n]`, `TNPC[n]`, and `TSTATE[n]` are undefined.

## 12.5.4.3 Watchdog Reset (WDR) Traps

The WDR reset in the UltraSPARC III Cu processor provide automatic recovery from `error_state`.

Processor state is undefined after WDR, except for the following<sup>1</sup>:

- The trap level is set.

$$TL \leftarrow \mathbf{min} (TL + 1, MAXTL)$$

1. Please note that these state transitions include the minimum set covered in the SPARC V9 architecture. There are other non-SPARC V9 transitions that are implementation-dependent and are not described here.

- Existing state is preserved.

```
TSTATE[TL].CCR      ← CCR
TSTATE[TL].ASI      ← ASI
TSTATE[TL].PSTATE   ← PSTATE
TSTATE[TL].CWP      ← CWP
TPC[TL]              ← PC
TNPC[TL]             ← nPC
```

- The trap type is set.

```
TT[TL] ← 00216
```

- The PSTATE register is set as follows:

```
PSTATE.MM           ← 002 (TSO)
PSTATE.RED          ← 1 (enter RED_state)
PSTATE.PEF          ← 1 (FPU is present)
PSTATE.AM           ← 0 (address masking is turned off)
PSTATE.PRIV         ← 1 (the processor enters privileged mode)
PSTATE.IE           ← 0 (interrupts are disabled)
PSTATE.AG           ← 1 (global registers are replaced with alternate globals)
PSTATE.MG           ← 0 (MMU globals are disabled)
PSTATE.IG           ← 0 (interrupt globals are disabled)
PSTATE.CLE          ← PSTATE.TLE (set endian mode for traps)
PSTATE.TLE          ← undefined1
```

- Implementation-specific state changes; for example, disabling an MMU.

- Control is transferred into the RED\_state trap table.

```
PC      ← RSTVaddr<63:8> □ 0100 00002
nPC     ← RSTVaddr<63:8> □ 0100 01002
```

For any reset when  $TL = MAXTL$ , for all  $n < MAXTL$ , the values in  $TPC[n]$ ,  $TNPC[n]$ , and  $TSTATE[n]$  are undefined.

#### 12.5.4.4 Externally Initiated Reset (XIR) Traps

XIR traps are initiated by an external signal. They behave like an interrupt that cannot be masked by  $IE = 0$  or  $PILL$ . Typically, XIR is used for critical system events, such as power failure, reset button pressed, failure of external components, that does not require a WDR (which aborts operations), or systemwide reset in a multiprocessor. The following state changes occur<sup>2</sup>:

- Existing state is preserved.

1. Note that this differs from SPARC V9.

2. Please note that these state transitions include the minimum set covered in the SPARC V9 architecture. There are other non-SPARC V9 transitions that are implementation-dependent and are not described here.

```

TSTATE[TL].CCR      ← CCR
TSTATE[TL].ASI      ← ASI
TSTATE[TL].PSTATE   ← PSTATE
TSTATE[TL].CWP      ← CWP
TPC[TL]             ← PC
TNPC[TL]            ← nPC

```

- The trap type is set.

```
TT[TL] ← 00316
```

- The PSTATE register is set as follows:

```

PSTATE.MM           ← 002 (TSO)
PSTATE.RED          ← 1 (enter RED_state)
PSTATE.PEF          ← 1 (FPU is present)
PSTATE.AM           ← 0 (address masking is turned off)
PSTATE.PRIV         ← 1 (the processor enters privileged mode)
PSTATE.IE           ← 0 (interrupts are disabled)
PSTATE.AG           ← 1 (global registers are replaced with alternate globals)
PSTATE.MG           ← 0 (MMU globals are disabled)
PSTATE.IG           ← 0 (interrupt globals are disabled)
PSTATE.CLE          ← PSTATE.TLE (set endian mode for traps)
PSTATE.TLE          ← undefined1

```

- Implementation-specific state changes, for example, disabling an MMU.
- Control is transferred into the RED\_state trap table.

```

PC      ← RSTVaddr<63:8> □ 0110 00002
nPC     ← RSTVaddr<63:8> □ 0110 01002

```

For any reset when  $TL = MAXTL$ , for all  $n < MAXTL$ , the values in  $TPC[n]$ ,  $TNPC[n]$ , and  $TSTATE[n]$  are undefined.

### 12.5.4.5 Software-Initiated Reset (SIR) Traps

SIR traps are initiated by execution of an SIR instruction in privileged mode. Supervisor software uses the SIR trap as a panic operation or a metasupervisor trap.

The following state changes occur<sup>1</sup>:

- If  $TL = MAXTL$ , then enter error\_state. Otherwise, do the following:
- The trap level is set.

```
TL ← TL + 1
```

- Existing state is preserved.

1. Please note that these state transitions include the minimum set covered in the SPARC V9 architecture. There are other non-SPARC V9 transitions that are implementation-dependent and are not described here.

|                   |                          |
|-------------------|--------------------------|
| TSTATE[TL].CCR    | ← CCR                    |
| TSTATE[TL].ASI    | ← ASI                    |
| TSTATE[TL].PSTATE | ← PSTATE                 |
| TSTATE[TL].CWP    | ← CWP                    |
| TPC[TL]           | ← PC                     |
| TNPC[TL]          | ← undefined <sup>1</sup> |

- The trap type is set.

TT[TL] ← 04<sub>16</sub>

- The PSTATE register is set as follows:

|             |  |
|-------------|--|
| PSTATE.MM   | ← 00 <sub>2</sub> (TSO)                                    |
| PSTATE.RED  | ← 1 (enter RED_state)                                      |
| PSTATE.PEF  | ← 1 (FPU is present)                                       |
| PSTATE.AM   | ← 0 (address masking is turned off)                        |
| PSTATE.PRIV | ← 1 (the processor enters privileged mode)                 |
| PSTATE.IE   | ← 0 (interrupts are disabled)                              |
| PSTATE.AG   | ← 1 (global registers are replaced with alternate globals) |
| PSTATE.MG   | ← 0 (MMU globals are disabled)                             |
| PSTATE.IG   | ← 0 (interrupt globals are disabled)                       |
| PSTATE.CLE  | ← PSTATE.TLE (set endian mode for traps)                   |
| PSTATE.TLE  | ← undefined <sup>2</sup>                                   |

- Implementation-specific state changes, for example, disabling an MMU.
- Control is transferred into the RED\_state trap table.

|     |   |
|-----|---|
| PC  | ← RSTVaddr<63:8> □ 1000 0000 <sub>2</sub> |
| nPC | ← RSTVaddr<63:8> □ 1000 0100 <sub>2</sub> |

For any reset when TL = MAXTL, for all  $n < \text{MAXTL}$ , the values in TPC[n], TNPC[n], and TSTATE[n] are undefined.

## 12.5.4.6 Normal Traps When the Processor Is in RED\_state

Normal traps taken when the processor is already in RED\_state are also processed in RED\_state, unless TL = MAXTL, in which case the processor enters error\_state.

Assuming that TL < MAXTL, the processor state shall be set as follows<sup>1</sup>:

- The trap level is set.

TL ← TL + 1

- Existing state is preserved.

1. Please note that these state transitions include the minimum set covered in the SPARC V9 architecture. There are other non-SPARC V9 transitions that are implementation-dependent and are not described here.

2. Note that this differs from SPARC V9.

```

TSTATE[TL].CCR      ← CCR
TSTATE[TL].ASI      ← ASI
TSTATE[TL].PSTATE   ← undefined1
TSTATE[TL].CWP      ← CWP
TPC[TL]             ← PC
TNPC[TL]            ← nPC

```

- The trap type is preserved.

```
TT[TL] ← trap type
```

- The PSTATE register is set as follows:

```

PSTATE.MM           ← 002 (TSO)
PSTATE.RED          ← 1 (enter RED_state)
PSTATE.PEF          ← 1 (FPU is present)
PSTATE.AM           ← 0 (address masking is turned off)
PSTATE.PRIV         ← 1 (the processor enters privileged mode)
PSTATE.IE           ← 0 (interrupts are disabled)
PSTATE.AG           ← 1 (global registers are replaced with alternate globals)
PSTATE.MG           ← 0 (MMU globals are disabled)
PSTATE.IG           ← 0 (interrupt globals are disabled)
PSTATE.CLE          ← PSTATE.TLE (set endian mode for traps)
PSTATE.TLE          ← undefined1

```

- For a register-window trap only, CWP is set to point to the register window that must be accessed by the trap-handler software, that is:
  - If  $TT[TL] = 024_{16}$  (a *clean\_window* trap), then  $CWP \leftarrow CWP + 1$ .
  - If  $(080_{16} \leq TT[TL] \leq 0BF_{16})$  (window *spill* trap), then  $CWP \leftarrow CWP + CANSAVE + 2$ .
  - If  $(0C0_{16} \leq TT[TL] \leq 0FF_{16})$  (window *fill* trap), then  $CWP \leftarrow CWP - 1$ .
- For non-register-window traps, CWP is not changed.
- Implementation-specific state changes; for example, disabling an MMU.
- Control is transferred into the RED\_state trap table.

```

PC      ← RSTVaddr<63:8> □ 1010 00002
nPC     ← RSTVaddr<63:8> □ 1010 01002

```

---

## 12.6 Exception and Interrupt Descriptions

This section describes the various exceptions and interrupt requests and the conditions that cause them. Each exception and interrupt request describes the corresponding trap type as defined by the trap model. On the UltraSPARC III Cu processor, all traps are precise except for the deferred traps and disrupting traps.

<sup>1</sup>. Note that this differs from SPARC V9

- ***data\_access\_exception*** [ $tt = 030_{16}$ ] (Precise) — An exception occurred on an attempted data access. Detailed information regarding the error is logged into the `FTYPE` field of Data Synchronous Fault Status Register (ASI  $58_{16}$ , VA =  $18_{16}$ ). Below is a list of exceptions that cause a *data\_access\_exception* exception.
  - **Invalid ASI** — An attempt to perform a load or store with undefined or reserved ASI or a disallowed instruction/ASI combination.
  - **Illegal Access to Strongly Ordered Page** — An attempt to access a strongly ordered page by any type of load instruction with non-faulting ASI or by FLUSH instruction.
  - **Illegal Access to Non-Faulting-Only Page** — An attempt to access a non-faulting-only page by any type of load or store instruction or FLUSH instruction with ASI other than non-faulting ASI.
  - **Illegal Access to Non-cacheable Page** — An attempt to access a non-cacheable page by atomic instructions (CASA, CASXA, SWAP, SWAPA, LDSTUB, LDSTUBA), atomic quad load instructions (LDDA with ASI =  $24_{16}$ ,  $2C_{16}$ ), or by FLUSH instruction.
- ***division\_by\_zero*** [ $tt = 028_{16}$ ] (Precise) — An integer divide instruction attempted to divide by zero.
- ***fill\_n\_normal*** [ $tt = 0C0_{16}-0DF_{16}$ ] (Precise)
- ***fill\_n\_other*** [ $tt = 0E0_{16}-0FF_{16}$ ] (Precise)
 

A RESTORE or RETURN instruction has determined that the contents of a register window must be restored from memory.

---

**Compatibility Note** – The SPARC V9 *fill\_n\_\** exceptions supersede the SPARC V8 *window\_underflow* exception.

---

- ***fp\_disabled*** [ $tt = 020_{16}$ ] (Precise) — An attempt was made to execute a Floating-point operate (FPop), a floating-point branch, or a floating-point load/store instruction while an FPU was not present, `PSTATE.PEF = 0`, or `FPRS.FEF = 0`.
- ***illegal\_instruction*** [ $tt = 010_{16}$ ] (Precise) — An attempt was made to execute an instruction with an unimplemented opcode, an ILLTRAP instruction, an instruction with invalid field usage, instruction breakpoints, or an instruction that would result in illegal processor state.

---

**Note** – Unimplemented FPop instructions generate *fp\_exception\_other* traps.

---

An *illegal\_instruction* is generated in the following cases:

- An instruction encoding does not match any of the opcode map definitions.
- An instruction is not implemented in hardware (if the `op` and `op3` fields of the instruction decode as an FPop, then a *fp\_exception\_other* exception, with `ftt = 3`, will be generated instead of *illegal\_instruction*).
- An illegal value is present in an instruction `i` field.
- An illegal value is present in a field that is explicitly defined for an instruction, such as `cc2`, `cc1`, `cc0`, `fcn`, `impl`, `op2` (IMPDEP2A, IMPDEP2B), `rcond`, or `opf_cc`.

- Illegal register alignment (such as odd `rd` value in a doubleword load instruction).
- RDASR instruction with `rs1 = 1, 7–14, 20–21, or 26–31`.
- RDASR with `rs1 = 15` and nonzero `rd`.
- RDPR with `rs1 = 16–30`.
- RDPR with `rs1 ≤ 3` when `TL = 0`.
- WRPR with `rd = 15–31`.
- WRPR with `rd ≤ 3` when `TL = 0`.
- WRPR to `PSTATE` register that attempts to set more than one of bits `IG`, `MG`, and `AG`.
- Illegal `rd` value for `LDXFSR`, `STXFSR`, or the deprecated instructions `LDFSR` or `STFSR`.
- Illegal `rd` value for `WRPR`.
- Illegal `rs1` value for `RDPR`.
- WRASR instruction with `rd = 1, 4, 5, 7–14, 26–31`.
- WRASR with `rd = 15` and nonzero `rs1`.
- WRASR with `rd = 15` and `i = 0`.
- `DONE` or `RETRY` when `TL = 0`.
- `ILLTRAP` instruction.
- Instruction breakpoint occurred.
- A reserved instruction field in `TCC` instruction is nonzero.

If a reserved instruction field in an instruction other than `TCC` is nonzero, an *illegal\_instruction* exception should be generated.<sup>1</sup>

---

**Note** – If an instruction breakpoint triggers an *illegal\_instruction* trap, then the *illegal\_instruction* trap has a higher priority than does a *privileged\_opcode* trap.

---

- *instruction\_access\_exception* [`tt = 00816`] (Precise) — A protection exception occurred on an instruction access, typically as a result of an attempt to access a privileged page while the processor was executing in non-privileged mode.
- *interrupt\_level\_n* [`tt = 04116–04F16`] (Disrupting) — An interrupt request level of *n* was presented to the IU, while `PSTATE.IE = 1` and (interrupt request level > `PIL`).
- *mem\_address\_not\_aligned* [`tt = 03416`] (Precise) — A load/store instruction generated a memory address that was not properly aligned according to the instruction, or a `JMPL` or `RETURN` instruction generated a non-word-aligned address.
- *power\_on\_reset* (POR) [`tt = 00116`] (Reset) — An external signal was asserted. This trap is issued to bring a system reliably from the power-off to the power-on state.
- *privileged\_action* [`tt = 03716`] (Precise) — An action defined to be privileged has been attempted while `PSTATE.PRIV = 0`. Some examples include: a data access by non-privileged software using an ASI value with its most significant bit = 0 (a restricted ASI), or an attempt to read the `TICK` register by non-privileged software when `TICK.NPT = 1`.
- *privileged\_opcode* [`tt = 01116`] (Precise) — An attempt was made to execute a privileged instruction while `PSTATE.PRIV = 0`.

1. It is required that reserved fields have zero value.

---

**Compatibility Note** – The trap type for *privileged\_opcode* is identical to that of the SPARC V8 *privileged\_instruction* trap. The name was changed to distinguish it from the new *privileged\_action* trap type.

---

- **RED\_state\_exception** [tt = 005<sub>16</sub>] — Caused when TL = MAXTL – 1 and a trap occurs, an event that brings the processor into RED\_state.
- **software\_initiated\_reset** (SIR) [tt = 004<sub>16</sub>] (Precise/Reset) — Caused by the execution of the WRSIR, write to SIR register, instruction. It allows system software to reset the processor.
- **spill\_n\_normal** [tt = 080<sub>16</sub>–09F<sub>16</sub>] (Precise)
- **spill\_n\_other** [tt = 0A0<sub>16</sub>–0BF<sub>16</sub>] (Precise)  
A SAVE or FLUSHW instruction has determined that the contents of a register window must be saved to memory.

---

**Compatibility Note** – The SPARC V9 *spill\_n\_\** exceptions supersede the SPARC V8 *window\_overflow* exception.

---

- **tag\_overflow** [tt = 023<sub>16</sub>] (Precise) — A TADDccTV or TsubccTV instruction was executed, and either 32-bit arithmetic overflow occurred or at least one of the tag bits of the operands was nonzero.
- **trap\_instruction** [tt = 100<sub>16</sub>–17F<sub>16</sub>] (Precise) — A Tcc instruction was executed and the trap condition evaluated to TRUE.
- **clean\_window** [tt = 024<sub>16</sub>–027<sub>16</sub>] (Precise) — A SAVE instruction discovered that the window about to be used contains data from another address space; the window must be cleaned before it can be used.
- **data\_access\_error** [tt = 032<sub>16</sub>] (Precise or Deferred) — An error occurred on a data access.
- **externally\_initiated\_reset** (XIR) [tt = 003<sub>16</sub>] (Reset) — An external signal was asserted. This trap is used for catastrophic events, such as power failure, reset button pressed, and systemwide reset, in multiprocessor systems.
- **fp\_exception\_ieee\_754** [tt = 021<sub>16</sub>] (Precise) — An FPop instruction generated an *IEEE\_754\_exception* and its corresponding trap enable mask (TEM) bit was one. The floating-point exception type, *IEEE\_754\_exception*, is encoded in the FSR.ftt, and specific *IEEE\_754\_exception* information is encoded in FSR.cexc.
- **fp\_exception\_other** [tt = 022<sub>16</sub>] (Precise) — An FPop instruction generated an exception other than an *IEEE\_754\_exception*. Some examples include: the FPop is unimplemented, or there was a sequence or hardware error in the FPU. The floating-point exception type is encoded in the FSR's ftt field.
- **instruction\_access\_error** [tt = 00A<sub>16</sub>] (Precise) — An error occurred on an instruction access.



- *LDDF\_mem\_address\_not\_aligned* [ $tt = 035_{16}$ ] (Precise) — An attempt was made to execute an LDDF instruction and the effective address was not doubleword aligned.
- *STDF\_mem\_address\_not\_aligned* [ $tt = 036_{16}$ ] (Precise) — An attempt was made to execute an STDF instruction and the effective address was not doubleword aligned.
- *Watchdog reset* (WDR) [ $tt = 002_{16}$ ] (Reset) — This trap occurs as a transition from `error_state` to `RED_state`.
- *ECC\_error* [ $tt = 063_{16}$ ] (Disrupting) — The trap to signal the detection of hardware errors asynchronous to the instruction execution, or to request to save the information logged for the error that was detected and corrected by the processor.

---

**Implementation Note** – Some implementations may refer to this trap by the name “*corrected\_ECC\_error*.”

---

- *fast\_data\_access\_MMU\_miss* [ $tt = 068_{16} - 06B_{16}$ ] (Precise) — During an attempted data access, the MMU detected that a translation lookaside buffer (TLB) did not contain a translation for the virtual address (that is, a TLB miss occurred). Four trap vectors are allocated for this trap, allowing a TLB miss handler of up to 32 instructions to fit within the trap vector area.
- *fast\_data\_access\_protection* [ $tt = 06C_{16} - 06F_{16}$ ] (Precise) — During an attempted data write access (by a store or load-store instruction), the instruction had appropriate access privilege but the MMU signalled that the location was write-protected (write to a read-only location). Note that on an UltraSPARC III Cu processor, an attempt to read or write to a privileged location while in non-privileged mode causes the higher priority *data\_access\_exception* instead of this exception. Four trap vectors are allocated for this trap, allowing a trap handler of up to 32 instructions to fit within the trap vector area.
- *fast\_instruction\_access\_MMU\_miss* [ $tt = 064_{16} - 067_{16}$ ] (Precise) — During an attempted instruction access, the MMU detected a TLB miss. Four trap vectors are allocated for this trap, allowing a trap handler of up to 32 instructions to fit within the trap vector area.
- *interrupt\_vector\_trap* [ $tt = 060_{16}$ ] (Disrupting) — The processor has received an interrupt request.
- *PA\_watchpoint* [ $tt = 061_{16}$ ] (Precise) — The processor has detected a physical address breakpoint.
- *VA\_watchpoint* [ $tt = 062_{16}$ ] (Precise) — The processor has detected a virtual address breakpoint.
- *fast\_ECC\_error* [ $tt = 070_{16}$ ] (Precise) — A single-bit or multiple-bit ECC error is detected.

This trap is taken on ECC errors from the L2-cache. The trap handler is required to flush the cache line containing the error from both the D-cache and L2-cache since incorrect data would have already been written into the D-cache. The UltraSPARC III Cu hardware will automatically correct single-bit ECC errors on the L2-cache write back when the trap

handler performs the L2-cache flush. After the caches are flushed, the instruction that encountered the error should be retried; the corrected data will then be brought back in from memory and reinstalled in the D-cache and L2-cache.

On *fast\_ECC\_error* detection during D-cache load miss fill, D-cache installs the uncorrected data. But since the *fast\_ECC\_error* trap is precise, hardware can rely on software to help clean up the bad data. I-cache is different. If I-cache is filled with errors, the line will not be installed in I-cache.

A D-cache or I-cache miss request may observe an ECC error in the line it reads from the L2-cache. When this occurs, a *fast\_ECC\_error* precise trap is generated for the instruction that detected the error. In the case of a D-cache request, the corrupted data will be installed in the D-cache, but the trap takes effects before the data can be used.

In case of an I-cache request, the data from the L2-cache will be corrected by hardware before being installed in the I-cache. When the *fast\_ECC\_error* trap is taken, the I-cache and D-cache are left enabled. Software is responsible to disable any caches as a part of the recovery. Software must flush the corrupted line from the D-cache if it was filled.

## Interrupt Handling

---

Processors and I/O devices can interrupt a selected processor by assembling and sending an interrupt packet consisting of eight 64-bit words of interrupt vector data. The contents of these data are defined by software convention. Thus, hardware interrupts and cross-calls can have the same hardware mechanism for interrupt delivery and can share a common software interface for processing.

The interrupt requesting/receiving mechanism is a two-step process: the sending of an interrupt request on a vector data register to the target and the scheduling of the received interrupt request on the target upon receipt.

An interrupt request packet is sent by the interrupter through the interrupt vector dispatch mechanism and is received by the specified target through the interrupt vector receive mechanism. Upon receipt of an interrupt request packet, a special trap is invoked on the target processor. The trap handler software invoked in the target processor then schedules the interrupt request to itself by posting the interrupt into *SOFTINT* register at the desired interrupt level.

Note that the processor may not send an interrupt request packet to itself through the interrupt dispatch mechanism. Separate sets of dispatch (outgoing) and receive (incoming) interrupt data registers allow simultaneous interrupt dispatching and receiving.

In the following sections, we describe different aspects of interrupt handling:

- Interrupt Vector Dispatch
- Interrupt Vector Receive
- Interrupt Global Registers
- Interrupt ASI Registers
- Software Interrupt Register (*SOFTINT*)

---

## 13.1 Interrupt Vector Dispatch

To dispatch an interrupt or cross-call, a processor or I/O device first writes to the Outgoing Interrupt Vector Data Registers according to an established software convention, described below. A subsequent write to the Interrupt Vector Dispatch Register triggers the interrupt delivery. The status of the interrupt dispatch can be read by polling the `ASI_INTR_DISPATCH_STATUS`'s `BUSY` and `NACK` bits. A `MEMBAR #Sync` should be used before polling begins to ensure that earlier stores are completed. If both `NACK` and `BUSY` are cleared, the interrupt has been successfully delivered to the target processor. With the `NACK` bit cleared and `BUSY` bit set, the interrupt delivery is pending. Finally, if the delivery cannot be completed (if it is rejected by the target processor), the `NACK` bit is set. The pseudocode sequence in CODE EXAMPLE 13-1 sends an interrupt.

The `ASI_INTR_DISPATCH_STATUS` register contains 32 pairs of `BUSY/NACK` bit pairs enabling interrupts to be pipelined. Specifying a unique pair of `BUSY/NACK` bits to be used for each interrupt when writing the Interrupt Dispatch Register enables up to 32 interrupts to be outstanding at one time.

---

**Note** – The processor may not send an interrupt vector to itself through outgoing interrupt vector data registers. Doing so causes undefined interrupt vector data to be returned.

---

### CODE EXAMPLE 13-1 Code Sequence for Interrupt Dispatch

```
Read state of ASI_INTR_DISPATCH_STATUS; Error if BUSY
```

```
<no pending interrupt dispatch packet>
```

```
Repeat
```

```
    Begin atomic sequence(PSTATE.IE ← 0)
```

```
    Store to IV data reg 0 at ASI_INTR_W, VA=0x40 (optional)
```

```
    Store to IV data reg 1 at ASI_INTR_W, VA=0x48 (optional)
```

```
    Store to IV data reg 2 at ASI_INTR_W, VA=0x50 (optional)
```

```
    Store to IV data reg 3 at ASI_INTR_W, VA=0x58 (optional)
```

```
    Store to IV data reg 4 at ASI_INTR_W, VA=0x60 (optional)
```

```
    Store to IV data reg 5 at ASI_INTR_W, VA=0x68 (optional)
```

```
    Store to IV data reg 6 at ASI_INTR_W, VA=0x80 (optional)
```

```
    Store to IV data reg 7 at ASI_INTR_W, VA=0x88 (optional)
```

```
    Store to IV dispatch at ASI_INTR_W, VA<63:29>=0,
```

```
        VA<28:24>=BUSY/NACK bit #,VA<23:14>=ITID,
```

```
        VA<13:0>=0x70 initiates interrupt delivery
```

```
    Membar #Sync (wait for stores to finish)
```

**CODE EXAMPLE 13-1** Code Sequence for Interrupt Dispatch (Continued)

```
Poll state of ASI_INTR_DISPATCH_STATUS (BUSY, NACK)
  Loop if BUSY
End atomic sequence(PSTATE.IE ← 1)
```

---

**Note** – To avoid deadlocks, enable interrupts for some period before retrying the atomic sequence. Alternatively, implement the atomic sequence with locks without disabling interrupts.

---

## 13.2 Interrupt Vector Receive

When an interrupt is received, all eight Interrupt Data Registers are updated, regardless of which are being used by software. This update is done in conjunction with the setting of the BUSY bit in the ASI\_INTR\_RECEIVE register. At this point, the processor inhibits further interrupt packets from the system bus. If interrupts are enabled (PSTATE.IE = 1), then an interrupt trap (implementation-dependent trap type 60<sub>16</sub>) is generated. Software reads the ASI\_INTR\_RECEIVE register and Incoming Interrupt Data Registers to determine the entry point of the appropriate trap handler. All of the external interrupt packets are processed at the highest interrupt priority level and are then re-prioritized as lower priority interrupts in the software handler. CODE EXAMPLE 13-2 illustrates interrupt receive handling.

**CODE EXAMPLE 13-2** Code Sequence for an Interrupt Receive

```
Read state of ASI_INTR_RECEIVE; Error if !BUSY
Read from IV data reg 0 at ASI_SDB_INTR_R, VA=0x40 (optional)
Read from IV data reg 1 at ASI_SDB_INTR_R, VA=0x48 (optional)
Read from IV data reg 2 at ASI_SDB_INTR_R, VA=0x50 (optional)
Read from IV data reg 3 at ASI_SDB_INTR_R, VA=0x58 (optional)
Read from IV data reg 4 at ASI_SDB_INTR_R, VA=0x60 (optional)
Read from IV data reg 5 at ASI_SDB_INTR_R, VA=0x68 (optional)
Read from IV data reg 6 at ASI_SDB_INTR_R, VA=0x80 (optional)
Read from IV data reg 7 at ASI_SDB_INTR_R, VA=0x88 (optional)
Determine the appropriate handler
Handle interrupt or reprioritize this trap and
  set the SOFTINT register
```

---

## 13.3 Interrupt Global Registers

A separate set of global registers is implemented to expedite interrupt processing. As described in Section 13.2, “Interrupt Vector Receive,” the processor takes an implementation-dependent interrupt trap after receiving an interrupt packet. Software uses a number of scratch registers while determining the appropriate handler and constructing the interrupt state.

A separate set of eight Interrupt Global Registers (IGRs) replaces the eight programmer-visible global registers during interrupt processing. After an interrupt trap is dispatched, the hardware selects the interrupt global registers by setting the `PSTATE.IG` field. The previous value of `PSTATE` is restored from the trap stack by a `DONE` or `RETRY` instruction on exit from the interrupt handler.

---

## 13.4 Interrupt ASI Registers

`MEMBAR #Sync` is generally needed after stores to interrupt ASI registers, which avoids unnecessary effects caused by possible prefetches to the locations with side-effect.

### 13.4.1 Outgoing Interrupt Vector Data<7:0> Register

`ASI_INTR_W` (data 0):  $ASI = 77_{16}$ ,  $VA<63:0> = 40_{16}$   
`ASI_INTR_W` (data 1):  $ASI = 77_{16}$ ,  $VA<63:0> = 48_{16}$   
`ASI_INTR_W` (data 2):  $ASI = 77_{16}$ ,  $VA<63:0> = 50_{16}$   
`ASI_INTR_W` (data 3):  $ASI = 77_{16}$ ,  $VA<63:0> = 58_{16}$   
`ASI_INTR_W` (data 4):  $ASI = 77_{16}$ ,  $VA<63:0> = 60_{16}$   
`ASI_INTR_W` (data 5):  $ASI = 77_{16}$ ,  $VA<63:0> = 68_{16}$   
`ASI_INTR_W` (data 6):  $ASI = 77_{16}$ ,  $VA<63:0> = 80_{16}$   
`ASI_INTR_W` (data 7):  $ASI = 77_{16}$ ,  $VA<63:0> = 88_{16}$

**Name:** `ASI_INTR_W`: Outgoing Interrupt Vector Data Registers (privileged, write-only)

TABLE 13-1 describes the register field of the eight Outgoing Interrupt Vector Data Registers.

**TABLE 13-1** Outgoing Interrupt Vector Data Register Format

| Bits | Field | Type | Description    |
|------|-------|------|----------------|
| 63:0 | Data  | W    | Interrupt data |

A write to these eight registers modifies the outgoing Interrupt Dispatch Data Registers.

Non-privileged access to this register causes a *privileged\_action* trap. An attempt to read this register causes a *data\_access\_exception* trap.

## 13.4.2 Interrupt Vector Dispatch Register

ASI 77<sub>16</sub>

VA<63:39> = 0

VA<38:29> = SID<9:0>

VA<28:24> = BUSY/NACK bit pair # (BN)

VA<23:14> = interrupt target identifier (ITID)

VA<13:0> = 70<sub>16</sub>

**Name:** ASI\_INTR\_DISPATCH\_W (interrupt dispatch) (Privileged, write-only)

TABLE 13-2 describes the fields of the Interrupt Vector Dispatch Register.

**TABLE 13-2** Interrupt Vector Dispatch Register Format

| Bits      | Field | Type | Description   |
|-----------|-------|------|---|
| VA<28:24> | BN    | W    | Specifies which of the BUSY/NACK bit pairs to use for the interrupt. A 016 in this field (which current software is using) selects BUSY/NACK bits ASI_INTR_DISPATCH_STATUS<1:0> for backward compatibility. A 116 in this field selects BUSY/NACK bits ASI_INTR_DISPATCH_STATUS<3:2>. |
| VA<23:14> | ITID  | W    | Interrupt Target ID. Specifies the interrupt target CPU using the BUSY/NACK bit pair BN, along with the contents of the eight Interrupt Vector Data Registers.  |

A write to this ASI triggers an interrupt vector dispatch to the target CPU identified with Interrupt Target ID (ITID), using BUSY/NACK bit pair BN along with the contents of the eight Interrupt Vector Data Registers. Note that the write acts as a trigger; however, the data for the write is ignored. The upper bits VA<28:14> steer the interrupt.

The UltraSPARC III Cu processor interprets all ten bits of VA<38:39> when the Interrupt Vector Dispatch Register is written.

A read from the Interrupt Vector Dispatch Register causes a *data\_access\_exception* trap. Non-privileged access to this register causes a *privileged\_action* trap.

## 13.4.3 Interrupt Vector Dispatch Status Register

ASI 48<sub>16</sub>

VA<63:0> = 0

**Name:** ASI\_INTR\_DISPATCH\_STATUS (Privileged, read-only)

TABLE 13-3 describes the fields of the Interrupt Vector Dispatch Status Register.

**TABLE 13-3** Interrupt Dispatch Status Register Format

| Bits | Field | Type | Description   |
|------|-------|------|---|
| Odd  | NACK  | R    | Set if interrupt dispatch has failed. Cleared at the start of every interrupt dispatch attempt; set when a dispatch has failed. |
| Even | BUSY  | R    | Set when there is an outstanding dispatch.  |

In the UltraSPARC III Cu processor, 32 BUSY/NACK pairs are implemented in the Interrupt Vector Dispatch Status Register.

The status of up to 32 outgoing interrupts can be read from ASI\_INTR\_DISPATCH\_STATUS BUSY/NACK bits. This register contains up to 32 pairs of BUSY/NACK bit pairs: the pair at <1:0> is referred to as pair 0, <3:2> as pair 1, and so on up to pair 31 at bits <63:62>. The VA<28:24> field of the Interrupt Dispatch Register specifies which BUSY/NACK bit pair will be used for the interrupt.

Writes to this ASI cause a *data\_access\_exception* trap. Non-privileged access to this register causes a *privileged\_action* trap.

## 13.4.4 Incoming Interrupt Vector Data<7:0>

ASI\_INTR\_R (data 0): ASI = 7F<sub>16</sub>, VA<63:0> = 40<sub>16</sub>

ASI\_INTR\_R (data 1): ASI = 7F<sub>16</sub>, VA<63:0> = 48<sub>16</sub>

ASI\_INTR\_R (data 2): ASI = 7F<sub>16</sub>, VA<63:0> = 50<sub>16</sub>

ASI\_INTR\_R (data 3): ASI = 7F<sub>16</sub>, VA<63:0> = 58<sub>16</sub>

ASI\_INTR\_R (data 4): ASI = 7F<sub>16</sub>, VA<63:0> = 60<sub>16</sub>

ASI\_INTR\_R (data 5): ASI = 7F<sub>16</sub>, VA<63:0> = 68<sub>16</sub>

ASI\_INTR\_R (data 6): ASI = 7F<sub>16</sub>, VA<63:0> = 80<sub>16</sub>

ASI\_INTR\_R (data 7): ASI = 7F<sub>16</sub>, VA<63:0> = 88<sub>16</sub>

**Name:** ASI\_INTR\_R



TABLE 13-4 describes the register field of the eight Incoming Interrupt Vector Data Registers.

**TABLE 13-4** Incoming Interrupt Vector Data Register Format

| Bits | Field | Type | Description    |
|------|-------|------|----------------|
| 63:0 | Data  | R    | Interrupt data |

A read from these registers returns incoming interrupt information from the incoming Interrupt Receive Data Registers.

Non-privileged access to this register causes a *privileged\_action* trap.

## 13.4.5 Interrupt Vector Receive Register

ASI 49<sub>16</sub>

VA<63:0> = 0

**Name:** ASI\_INTR\_RECEIVE (Privileged)

TABLE 13-5 describes the fields of the Interrupt Receive Register.

**TABLE 13-5** Interrupt Receive Register Format

| Bits  | Field | Type | Description  |
|-------|-------|------|--|
| 63:11 |       | R    | Reserved.  |
| 10:6  | SID_U | R    | Most significant (upper) 5 bits of the physical module ID (MID) of the interrupter. Source ID bits <9:5> of interrupter. |
| 5     | BUSY  | RW   | Set when an interrupt vector is received. The BUSY bit must be cleared by software writing zero.                         |
| 4:0   | SID_L | R    | Least significant (lower) 5 bits of the physical module ID (MID) of the interrupter.                                     |

The status of an incoming interrupt can be read from ASI\_INTR\_RECEIVE. The BUSY bit is cleared by writing zero to this register.

The UltraSPARC III Cu processor sets all ten physical module ID (MID) bits in the SID\_U and SID\_L fields of the Interrupt Vector Receive Register. These bits correspond to the interrupt ID of the interrupter.

Non-privileged access to the Interrupt Vector Receive Register causes a *privileged\_action* trap.

---

## 13.5 Software Interrupt Register (SOFTINT)

To schedule interrupt vectors for processing at a later time, each processor can send itself signals by setting bits in the SOFTINT register.

The SOFTINT register (ASR 16<sub>16</sub>) is used for conication from nucleus (TL > 0) code to kernel (TL = 0) code. Interrupt packets and other service requests can be scheduled in queues or mailboxes in memory by the nucleus, which then sets SOFTINT<n> to cause an interrupt at level <n>.

Non-privileged access to this register causes a *privileged\_opcode* trap.

### 13.5.1 Setting the Software Interrupt Register

Setting SOFTINT<n> is done by a write to the SET\_SOFTINT register (ASR 14<sub>16</sub>), with bit *n* corresponding to the interrupt level set. The value written to the SET\_SOFTINT register is effectively ORed into the SOFTINT register. This approach allows the interrupt handler to set one or more bits in the SOFTINT register with a single instruction.

Read accesses to the SET\_SOFTINT register cause an *illegal\_instruction* trap. Non-privileged accesses to this register cause a *privileged\_opcode* trap.

When the nucleus returns, if (PSTATE.IE = 1) and (*n* > PIL), then the processor will receive the highest priority interrupt IRL<n> of the asserted bits in SOFTINT<16:0>. The processor then takes a trap for the interrupt request, and the nucleus sets the return state to the interrupt handler at that PIL and returns to TL = 0. In this manner, the nucleus can schedule services at various priorities and process them according to their priority.

### 13.5.2 Clearing the Software Interrupt Register

When all interrupts scheduled for service at level *n* have been serviced, the kernel writes to the CLEAR\_SOFTINT register (ASR 15<sub>16</sub>) with bit *n* set, to clear that interrupt. The complement of the value written to the CLEAR\_SOFTINT register is effectively ANDed with the SOFTINT register. This approach allows the interrupt handler to clear one or more bits in the SOFTINT register with a single instruction.

Read accesses to the CLEAR\_SOFTINT register cause an *illegal\_instruction* trap. Non-privileged write accesses to this register cause a *privileged\_opcode* trap.

The timer interrupt TICK\_INT and system timer interrupt STICK\_INT are equivalent to SOFTINT<14> and have the same effect.

---

**Note** – To avoid a race condition between the kernel clearing an interrupt and the nucleus setting it, the kernel should examine the queue for any valid entries again after clearing the interrupt bit.

---

TABLE 13-6 summarizes the SOFTINT ASRs.

**TABLE 13-6** SOFTINT ASRs

| ASR Value | ASR Name      | Type | Description                              |
|-----------|---------------|------|--|
| 1416      | SET_SOFTINT   | W    | Set bit(s) in Soft Interrupt Register.   |
| 1516      | CLEAR_SOFTINT | W    | Clear bit(s) in Soft Interrupt Register. |
| 1616      | SOFTINT       | RW   | Per-processor Soft Interrupt Register.   |



## SECTION VI

---

---

# Performance Programming

---

---



# Performance Instrumentation

---

Performance instrumentation consists of processor event counters that can be used to gather statistics during program execution and calls that start and stop the gathering process. Many events can be monitored, two at a time, to gain information about the performance of the processor. Memory access and stall times, for example, can be measured using two, 32-bit Performance Instrumentation Counters (PICs). The PCR and PIC are accessed through read /write Ancillary State Register instructions.

This chapter describes the performance instrumentation features in the following sections:

- Section 14.1, “Performance Control Register (PCR)”
- Section 14.2, “Performance Instrumentation Counter (PIC) Register”
- Section 14.3, “Performance Instrumentation Operation”
- Section 14.4, “Pipeline Counters”
- Section 14.5, “Cache Access Counters”
- Section 14.6, “Memory Controller Counters”
- Section 14.7, “Data Locality Counters for Scalable Shared Memory Systems”
- Section 14.8, “Miscellaneous Counters”
- Section 14.9, “PCR.SL and PCR.SU Encodings”

## *Supervisor/User Mode*

Access to the PCR is privileged. Non-privileged accesses cause a *privileged\_opcode* trap. Software can restrict non-privileged access to PICs by setting the PCR.PRIV field while in privileged mode. When PCR.PRIV = 1 (supervisor access only), an attempt by User Software to access the PIC register causes a *privileged\_action* trap. Software can control event measurements in non-privilege or privileged modes by setting the PCR.UT (user trace) and PCR.ST (system trace) fields.<sup>1</sup>

1. The PCR has mode bits to enable the counters in privileged mode, non-privilege mode, or to count when in either mode. The mode setting affects both counters.

# 14.1 Performance Control Register (PCR)

The Performance Control Register (PCR) is used to select the events to monitor and provides control for counting in privileged and/or non-privileged modes.

The 64-bit PCR is accessed through read/write Ancillary State Register (ASR) instructions (RDASR/WRASR). PCR is located at ASRs 16 ( $10_{16}$ ).

Two events can simultaneously be measured by setting the PIC\_SL and PIC\_SU fields. The counters can be enabled separately for Supervisor and User mode using UT and ST fields. The selected statistics are reflected during subsequent accesses to the PICs.

The PCR is a read/write register used to control the counting of performance monitoring events. FIGURE 14-1 shows the details of the PCR. TABLE 14-1 describes the various fields of the PCR. Counts are collected in the PIC register (see Section 14.2, “Performance Instrumentation Counter (PIC) Register”).

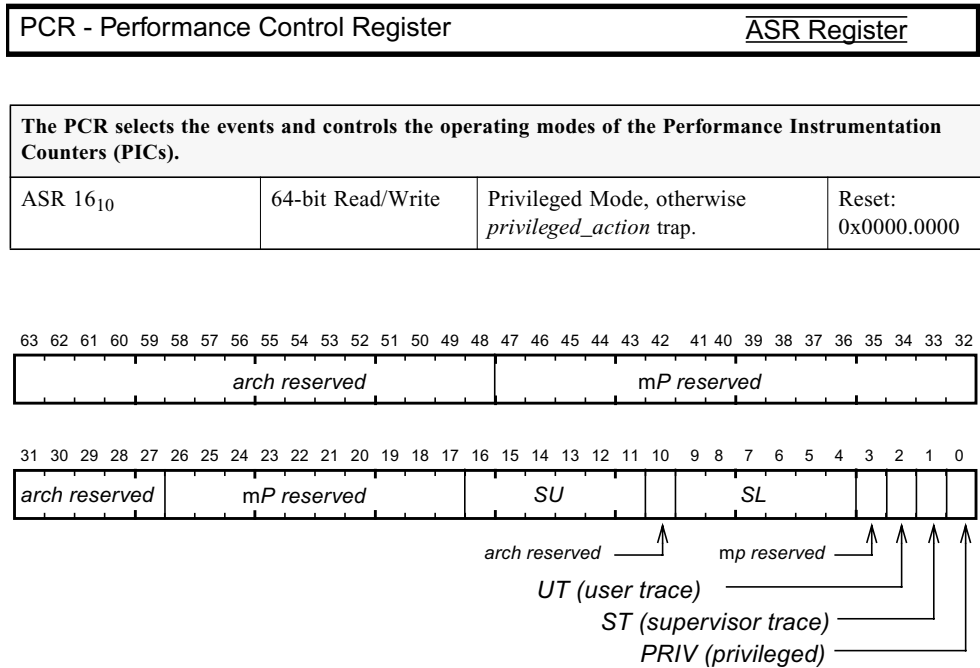


FIGURE 14-1 Performance Control Register



**TABLE 14-1** PCR Bit Description

| Bit   | Field | Description   |
|-------|-------|---|
| 63:48 | —     | Reserved by SPARC architecture.<br>Read zero, write zero, or write value read previously (read-modify-write).   |
| 47:32 | —     | Unused UltraSPARC III Cu processor.<br>Read zero, write zero, or write value read previously (read-modify-write).   |
| 31:27 | —     | Reserved by SPARC architecture.<br>Read zero, write zero, or write value read previously (read-modify-write).   |
| 26:17 | —     | Unused UltraSPARC III Cu processor.<br>Read zero, write zero, or write value read previously (read-modify-write).   |
| 16:11 | SU    | Selects 1 of up to 64 counters accessible in the upper half (bits <63:32>) of the PIC register.   |
| 10    | —     | Reserved by SPARC architecture.<br>Read zero, write zero, or write value read previously (read-modify-write).   |
| 9:4   | SL    | Selects 1 of up to 64 counters accessible in the lower half (bits <31:0>) of the PIC register.  |
| 3     | —     | Unused UltraSPARC III Cu processor.<br>Read zero, write zero, or write value read previously (read-modify-write).   |
| 2     | UT    | User Trace Enable.<br>If set to one, counts events in non-privileged mode (User).   |
| 1     | ST    | System Trace Enable.<br>If set to one, counts events in privileged mode (Supervisor).<br><b>Notes:</b><br>If both PCR.UT and PCR.ST are set to one, all selected events are counted.<br>If both PCR.UT and PCR.ST are zero, counting is disabled.<br>PCR.UT and PCR.ST are global fields which apply to both PIC pairs. |
| 0     | PRIV  | Privileged. If PCR.PRIV = 1, a non-privileged (PSTATE.PRIV = 0) attempt to access PIC (via a RDPIC or WRPIC instruction) will result in a <i>privileged_action</i> exception.   |

---

## 14.2 Performance Instrumentation Counter (PIC) Register

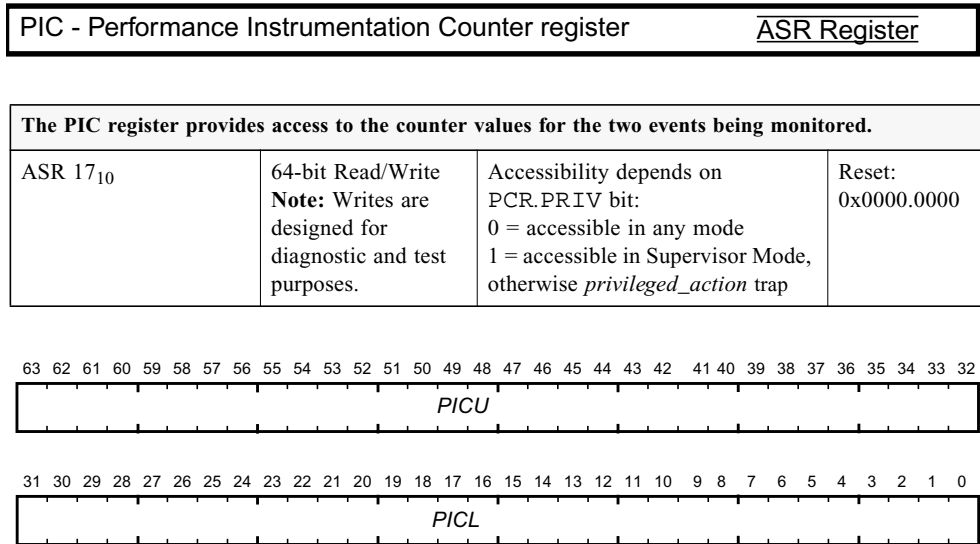
The 64-bit PIC is accessed through read/write Ancillary State Register (ASR) instructions (RDASR/WRASR). PIC is located at ASRs 17 ( $11_{16}$ ).

The PIC counters can be monitored during program execution to gather on-going statistics or reconfigure during steady-state program execution to gather statistics for more than two events. The pair of 32-bit counters can accumulate over four billion events each prior to

wrapping. Overflow of PICL or PICU causes a disrupting trap and SOFTINT. Active monitoring will allow the gathering software to extend the data range by periodically reading the contents of the PICs to detect and avoid overflow; an interrupt can be enabled on a counter overflow<sup>1</sup>.

Each of the two 32-bit PICs can accumulate over 4 billion events before wrapping around. Overflow of PICL or PICU causes a disrupting trap and SOFTINT register bit 15 to be set to 1; then, if ((PSTATE.IE = 1) and (PIL < 15)), causes an *interrupt\_level\_15* trap. Extended event logging can be accomplished by periodic reading of the contents of the PICs before each overflows. Additional statistics can be collected by use of the two PICs over multiple passes of program execution.<sup>2</sup>

The difference between the values read from the PIC on two reads reflects the number of events that occurred between register reads. Software can only rely on read-to-read PIC accesses to get an accurate count and not a write-to-read of the PIC counters. FIGURE 14-2 shows the details of the PIC. TABLE 14-2 describes the various fields of the PIC.



**FIGURE 14-2** Performance Instrumentation Counter Register

1. The point at which the interrupt due to a PIC overflow is delivered may be several instructions after the instruction responsible for the overflow event. This situation is known as a "skid". The degree of skid depends on the event that caused the overflow and the type of instructions being executed in the pipeline at the time the overflow occurred. It may not be possible to associate a counter overflow with the particular instruction that caused it due to the skid problem.
2. Two events can simultaneously be measured by setting PCR.SU/PCR.SL fields along with the PCR.UT and PCR.ST fields. The selected statistics are reflected during subsequent accesses to the PICs.

**TABLE 14-2** PIC Register Fields

| Bit   | Field | Description  |
|-------|-------|--|
| 63:32 | PICU  | 32-bit field representing the count of an event selected by the SU field of the Performance Control Register (PCR) |
| 31:0  | PICL  | 32-bit field representing the count of an event selected by the SL field of the Performance Control Register (PCR) |

## 14.2.1 PIC Counter Overflow Trap Operation

When a PIC counter overflows, an interrupt is generated as described in TABLE 14-3.

**TABLE 14-3** PIC Counter Overflow Processor Compatibility Comparison

| Function             | Description  |
|----------------------|--|
| PIC Counter Overflow | On overflow, a counter wraps to zero, SOFTINT register bit 15 is set to one, and an <i>interrupt_level_15</i> trap (a disrupting trap). The counter overflow trap is triggered on the transition from value FFFF FFFF <sub>16</sub> to value 0. The point at which the interrupt is delivered may be several instructions after the instruction responsible for the overflow event. This situation is known as a “skid.” |

---

## 14.3 Performance Instrumentation Operation

FIGURE 14-3 shows how an operating system might use the performance instrumentation features to provide event monitoring services.

Setup the PCR register as desired to select two events<sup>1</sup> and in which modes data should be collected. The monitoring must consider the real effects of the computer that includes calls to the system and interrupts. When used, the PCR register is considered part of a process state and must be saved and restored when switching process contexts.

Multiple data collection times can be done while the program executes to show on-going statistics.

1. When more than two events need to be monitored, the program, code sequence, or code loop need to be run again with the new events enabled. It is not possible to monitor more than two events at any given time.

FOR ILLUSTRATIVE  
PURPOSES ONLY

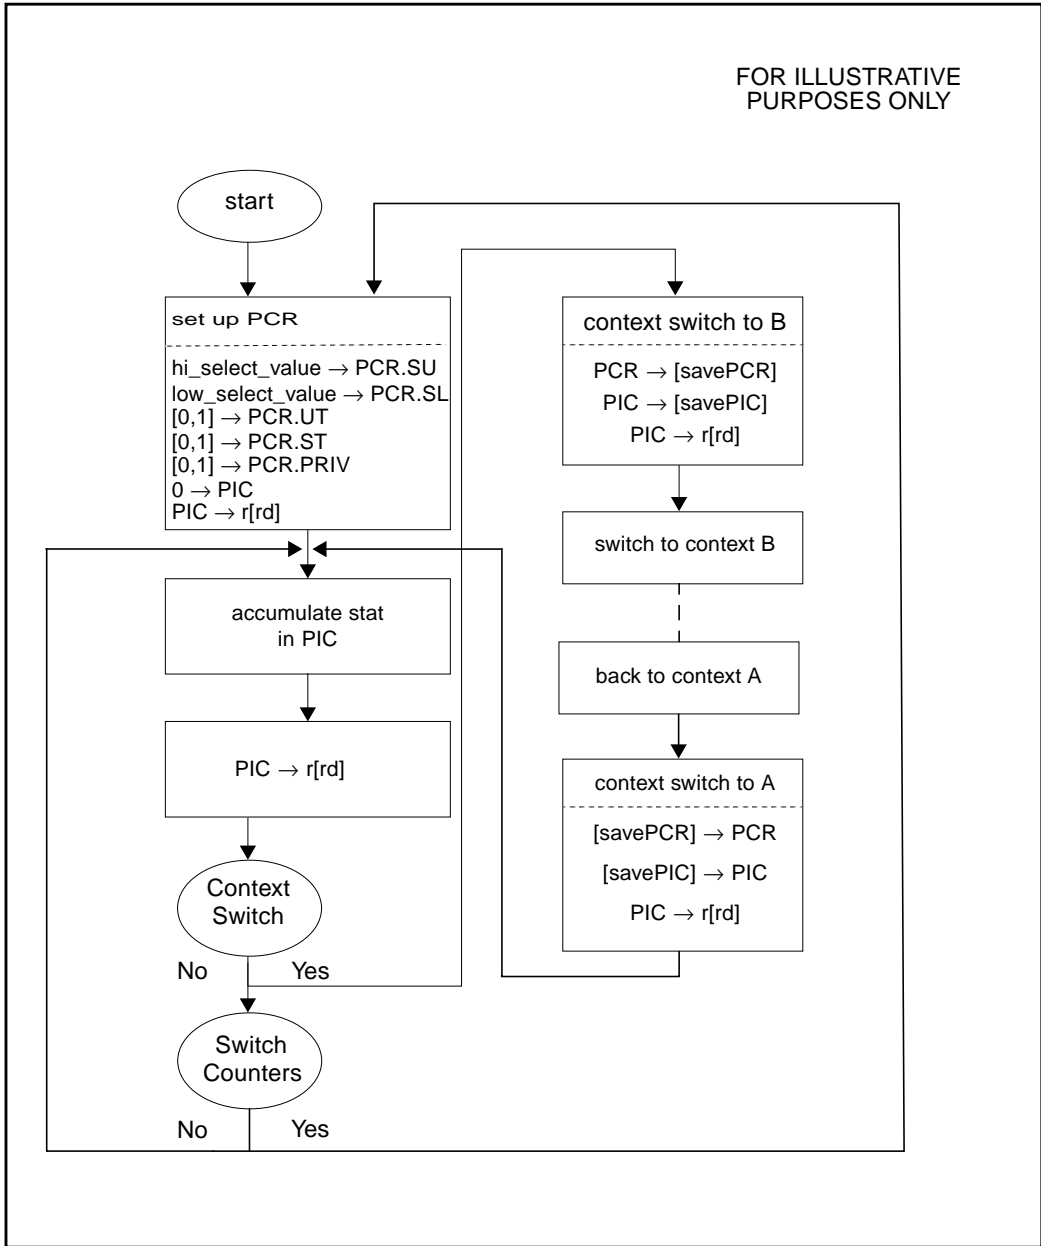


FIGURE 14-3 Operational Flow Diagram for Controlling Event Counters

## 14.3.1 Performance Instrumentation Implementations

Counting events and cycle stalls is sometimes complex because of the dynamic conditions and cancelled activities<sup>1</sup>.

## 14.3.2 Performance Instrumentation Accuracy

The performance instrumentation counters are designed to provide reasonable accuracy especially when used to count hundreds or thousands of events or stall cycles or when comparing the PIC counts that have recorded a similar number of events or stall cycles. Accuracy is most challenging when trying to associate an event to an instruction and when comparing PIC counts with one count rarely occurring.

When using the overflow trap, it is sometimes difficult to pin-point the instruction that is responsible for the overflow because of the way the pipeline is designed. A delay of several instructions is possible before the overflow is able to stop the current instruction flow and fetch the trap vector. This delay is referred to as skid and can occur for dozens of clock cycles. The skid for the load miss detection case is small. The skid value cannot be measured and its length depends on what event or stall cycle is being measured and what other instructions are in the pipeline.

---

# 14.4 Pipeline Counters

## 14.4.1 Instruction Execution and CPU Clock Counts

The instruction execution count monitors are described in TABLE 14-4 for clock and instruction execution counts.

**TABLE 14-4** Instruction Execution Clock Cycles and Counts

| Counter   | Description   |
|-----------|---|
| Cycle_cnt | [PICL 00.0000 and PICU 00.0000]<br>Counts clock cycles. This counter increments the same as the SPARC V9 TICK register, except that cycle counting is controlled by the PCR.UT and PCR.ST fields. |
| Instr_cnt | [PICL 00.0001 and PICU 00.0001]<br>Counts the number of instructions completed (retired) <sup>1</sup> .   |

1. This count does not include annulled, mispredicted, trapped, or helper instructions.

1. Specifics of each event and cycle stall are covered.

## *Synthesized Clocks Per Instruction (CPI)*

The cycle and instruction counts can be used to calculate the average number of instructions completed per cycle: Clock cycles per instruction,  $CPI = Cycle\_cnt / Instr\_cnt$ .

## 14.4.2 IIU Statistics

The counters listed in TABLE 14-5 record branch prediction event counts for taken and untaken branches in the Instruction Issue Unit (IIU). A retired branch in the following descriptions refers to a branch that reaches the D-stage without being invalidated.

**TABLE 14-5** Counters for Collecting IIU Statistics

| Counter                  | Description   |
|--------------------------|---|
| IU_Stat_Br_miss_taken    | [PICL 01.0101] Counts retired branches that were predicted to be taken, but in fact were not taken. |
| IU_Stat_Br_miss_untaken  | [PICU 01.1101] Counts retired branches that were predicted to be untaken, but in fact were taken.   |
| IU_Stat_Br_Count_taken   | [PICL 01.0110] Counts retired taken branches.   |
| IU_Stat_Br_Count_untaken | [PICU 01.1110] Counts retired untaken branches.   |

## 14.4.3 IIU Stall Counts

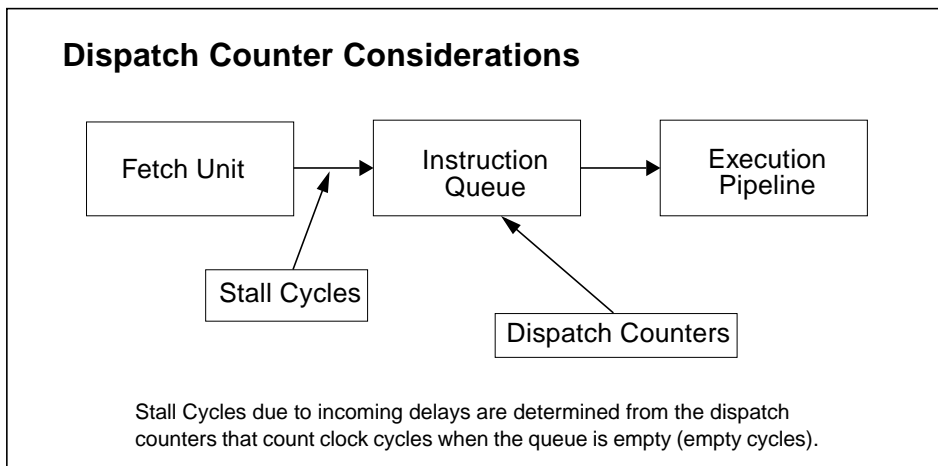
IIU stall counts, listed in TABLE 14-6 on page 14-371, are the major cause of pipeline stalls (bubbles) from the instruction fetch and decode pipeline. Stalls are counted for each clock cycle at which the associated condition is true.

FIGURE 14-4 on page 14-371 illustrates the first two considerations described below.

### 14.4.3.1 Dispatch Counter Considerations

1. Dispatch Counters count when the buffer is empty, regardless of whether the execution pipeline can accept more instructions from the instruction queue.
2. It is difficult to associate an empty queue to the reason it is empty. Multiple reasons together or separately can cause the instruction queue to be empty. The hardware picks the most recent disruptive event that is in the Fetch Unit to choose a counter to assign the empty queue cycles.

- Count accuracy is also subject to the conditions described in Section 14.3.2, “Performance Instrumentation Accuracy” for all counters.



**FIGURE 14-4** Dispatch Counters

**TABLE 14-6** Counters for IIU Stalls

| Counter             | Description  |
|---------------------|--|
| Dispatch0_IC_miss   | [PICL 00.0010] Counts the stall cycles due to the event that no instructions are issued because I-queue is empty from instruction cache miss. This count includes L2-cache miss processing if a L2-cache miss also occurs. <sup>1</sup>    |
| Dispatch0_mispred   | [PICU 00.0010] Counts the stall cycles due to the event that no instructions are issued because I-queue is empty due to branch misprediction. <sup>1</sup>   |
| Dispatch0_br_target | [PICL 00.0011] Counts the stall cycles due to the event that no instructions are issued because I-queue is empty due to a branch target address calculation. <sup>1</sup>  |
| Dispatch0_2nd_br    | [PICL 00.0100] Counts the stall cycles due to the event of having two branch instructions line-up in one 4-instruction group causing the second branch in the group to be re-fetched, delaying its entrance into the I-queue. <sup>1</sup> |
| Dispatch_rs_mispred | [PICL 01.0111] Counts the stall cycles due to the event that no instructions are issued because the I-queue is empty due to a Return Address Stack misprediction. <sup>1</sup>   |

1. See “Dispatch Counter Considerations” on page 370 for important information.

## 14.4.4 R-stage Stall Counts

Stalls are caused by dependency checks (data not ready for use by the instruction ready for dispatch) and by resources not being available (out-of-pipeline execution units needed, but are in-use).

The counters in TABLE 14-7 count the stall cycles at the R-stage of the pipeline. Stalls are counted for each clock at which the associated condition is true.

**TABLE 14-7** Counters for R-stage Stalls

| Counter       | Description   |
|---------------|---|
| Rstall_storeQ | [PICL 00.0101] Counts R-stage stall cycles for a store instruction which is the next instruction to be executed, but is stalled due to the store queue being full, that is, cannot hold additional stores. Up to eight entries can be in the store queue. |
| Rstall_FP_use | [PICU 00.1011] Counts R-stage stall cycles due to the event that the next instruction to be executed depends on the result of a preceding floating-point instruction in the pipeline that is not yet available.   |
| Rstall_IU_use | [PICL 00.0110] Counts R-stage stall cycles due to the event that the next instruction to be executed depends on the result of a preceding integer instruction in the pipeline that is not yet available.  |

## 14.4.5 Recirculation Counts

Recirculation instrumentation is implemented through the counters listed in TABLE 14-8.

**TABLE 14-8** Counters for Recirculation

| Counter       | Description  |
|---------------|--|
| Re_RAW_miss   | [PICU 10.0110] Counts stall cycles due to recirculation when there is a load in the E-stage which has a non-bypassable read-after-write (RAW) hazard with an earlier store instruction. This condition means that load data are being delayed by completion of an earlier store. See the Section 9.12, “Read-After-Write (RAW) Bypassing” for a description of the RAW hazard and causes of recirculation. |
| Re_FPU_bypass | [PICU 00.0101] Counts stall cycles due to recirculation when an FPU bypass condition that does not have a direct bypass path occurs.   |
| Re_DC_miss    | [PICU 00.0110] Counts stall cycles due to loads that miss D-cache and L2-cache and get recirculated. Includes cacheable loads only.  |
| Re_EC_miss    | [PICU 00.0111] Counts stall cycles due to loads that miss D-cache and L2-cache and get recirculated. Stall cycles from the point when L2-cache miss is detected to the D-stage of the recirculated flow are counted. Includes cacheable loads only.  |
| Re_PC_miss    | [PICU 01.0000] Counts stall cycles due to recirculation when a prefetch cache miss occurs on a prefetch predicted second load.   |



---

## 14.5 Cache Access Counters

Instruction, data, prefetch, write, and L2-cache access events can be collected through the counters listed in TABLE 14-9. Counts are updated by each cache access, regardless of whether the access will be used.

### 14.5.1 Instruction Cache Events

TABLE 14-9 describes the counters for instruction cache events.

**TABLE 14-9** Counters for Instruction Cache Events

| Counter           | Description  |
|-------------------|--|
| IC_ref            | [PICL 00.1000] Counts I-cache references. I-cache references are fetches (up to four instructions) from an aligned block of eight instructions. I-cache references are generally speculative and include instructions that are later cancelled due to mis-speculation. |
| IC_miss           | [PICU 00.1000] Counts I-cache misses. Includes fetches from mis-speculated execution paths which are later cancelled.  |
| IC_miss_cancelled | [PICU 00.0011] Counts I-cache misses cancelled due to mis-speculation, recycle, or other events.   |
| ITLB_miss         | [PICU 01.0001] Counts I-TLB miss traps taken.  |

### 14.5.2 Data Cache Events

TABLE 14-10 describes the counters for data cache events.

**TABLE 14-10** Counters for Data Cache Events

| Counter    | Description   |
|------------|---|
| DC_rd      | [PICL 00.1001] Counts D-cache read references (including accesses that subsequently trap). References to pages that are not virtually cacheable (TTE CV bit = 0) are not counted. |
| DC_rd_miss | [PICU 00.1001] Counts recirculated loads that miss the D-cache. Includes cacheable loads only.  |
| DC_wr      | [PICL 00.1010] Counts D-cache cacheable store accesses encountered (including cacheable stores that subsequently trap). Non-cacheable accesses are not counted.                   |

**TABLE 14-10** Counters for Data Cache Events *(Continued)*

| Counter    | Description   |
|------------|---|
| DC_wr_miss | [PICU 00.1010] Counts D-cache cacheable store accesses that miss D-cache. (There is no stall or recirculation on store miss). |
| DTLB_miss  | [PICU 01.0010] Counts memory reference instructions which trap due to D-TLB miss.   |

## 14.5.3 Write Cache Events

TABLE 14-11 describes the counters for write cache events.

**TABLE 14-11** Counters for Write Cache Events

| Counter       | Description   |
|---------------|---|
| WC_miss       | [PICU 01.0011] Counts W-cache misses.   |
| WC_snoop_cb   | [PICU 01.0100] Counts W-cache copybacks generated by a snoop from a remote processor. |
| WC_scrubbed   | [PICU 01.0101] Counts W-cache hits to clean lines.                                    |
| WC_wb_wo_read | [PICU 01.0110] Counts W-cache writebacks not requiring a read.                        |

## 14.5.4 Prefetch Cache Events

TABLE 14-12 describes the counters for prefetch cache events.

**TABLE 14-12** Counters for Prefetch Cache Events

| Counter      | Description   |
|--------------|---|
| PC_MS_miss   | [PICU 01.1111] Counts FP loads through the MS pipeline that miss P-cache.   |
| PC_soft_hit  | [PICU 01.1000] Counts FP loads that hit a P-cache line that was prefetched by a software-prefetch instruction.  |
| PC_hard_hit  | [PICU 01.1010] Counts FP loads that hit a P-cache line that was prefetched by a hardware prefetch.  |
| PC_snoop_inv | [PICU 01.1001] Counts P-cache invalidates generated by a snoop from a remote processor and stores by a local processor.                                   |
| PC_port0_rd  | [PICL 01.0000] Counts P-cache cacheable FP loads to the first port (general-purpose load path to D-cache and P-cache via MS pipeline).                    |
| PC_port1_rd  | [PICU 01.1011] Counts P-cache cacheable FP loads to the second port (memory and out-of-pipeline instruction execution loads via the A0 and A1 pipelines). |

## 14.5.5 L2-Cache Events

The L2-cache write hit count is determined by subtraction of the read hit and the instruction hit count from the total L2-cache hit count. The L2-cache write reference count is determined by subtraction of the D-cache read miss and I-cache misses from the total L2-cache references. Because of write caching, this is not the same as D-cache write misses. TABLE 14-13 describes the counter for L2-cache events.<sup>1</sup>

**TABLE 14-13** Counters for L2-Cache Events

| Counter          | Description   |
|------------------|---|
| EC_ref           | [PICL 00.1100] Counts L2-cache reference events. A 64-byte request is counted as one reference. Includes speculative D-cache load requests that turn out to be a D-cache hit. Count includes cacheable accesses only.                                   |
| EC_misses        | [PICU 00.1100] Counts L2-cache miss events sent to the System Interface Unit. Includes I-cache, D-cache, P-cache, W-cache exclusive (store), read stream (BLD), write stream (BST) requests that miss L2-cache. Count includes cacheable accesses only. |
| EC_write_hit_RTO | [PICL 00.1101] Counts W-cache exclusive requests that hit L2-cache in S, O, or O <sub>s</sub> state and thus, do a read-to-own (RTO) bus transaction.   |
| EC_wb            | [PICU 00.1101] Counts dirty sub-blocks that produce writebacks due to L2-cache miss events.   |
| EC_snoop_inv     | [PICL 00.1110] Counts L2-cache invalidates generated from a snoop by a remote processor.  |
| EC_snoop_cb      | [PICU 00.1110] Counts L2-cache copybacks generated from a snoop by a remote processor.  |
| EC_rd_miss       | [PICL 00.1111] Counts L2-cache miss events (including atomics) from D-cache requests. Cacheable D-cache loads only.   |
| EC_ic_miss       | [PICU 00.1110] Counts L2-cache read misses from I-cache requests. The counter counts all I-cache misses including those for instructions from the mis-speculated execution path. Cacheable requests only.   |

1. A block load or store access is counted as eight (8) references. For atomics, the read and write events are counted individually.

## 14.5.6 Separating D-cache Stall Counts

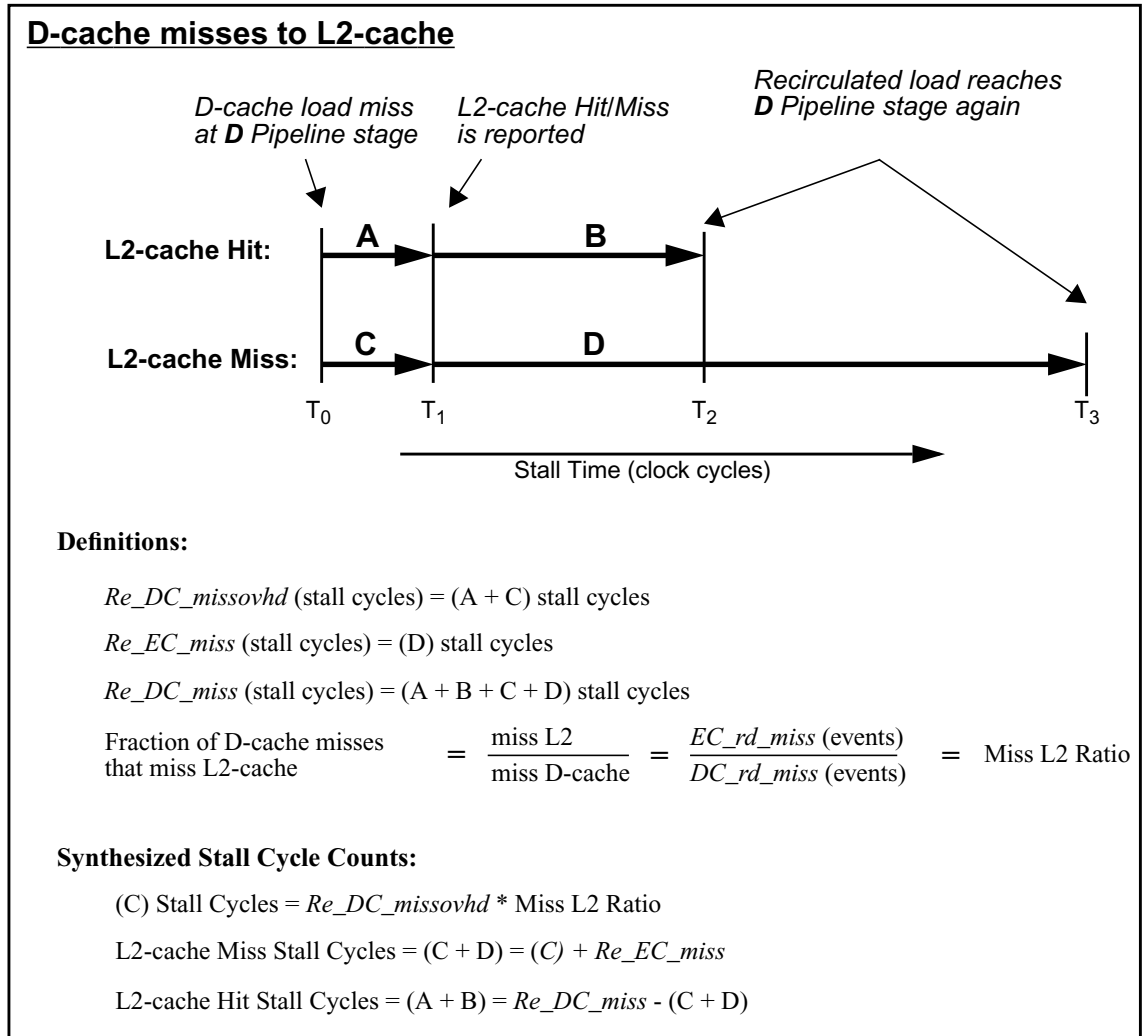
The D-cache stall cycle counts can be measured separately for L2-cache hits and misses by using the *Re\_DC\_missovhd* counter. The *Re\_DC\_missovhd* stall cycle counter is used with the recirculation and cache access events to separately calculate the D-cache loads that hit and miss the L2-cache. TABLE 14-14 describes the *Re\_DC\_missovhd* stall cycle counter processor compatibility.

**TABLE 14-14** *Re\_DC\_missovhd* Stall Counter Processor Compatibility

| Function              | Description  |
|-----------------------|--|
| <i>Re_DC_missovhd</i> | [PICL 00.0100]The <i>Re_DC_missovhd</i> cycle stall counter. |

## Synthesizing Individual Hit and Miss Stall Times

To explain the synthesis for L2-cache hit and miss stall times separately, consider the four stall regions A, B, C, and D shown in FIGURE 14-5 and the definitions and calculations that follow.



**FIGURE 14-5** D-Cache Load Miss Stall Regions

---

## 14.6 Memory Controller Counters

This section describes the memory controller counters in the UltraSPARC III Cu processor.

### 14.6.1 Memory Controller Read Request Events

If the snoop indicates that the line exists in some other device, the memory read request is cancelled. The cancelled read requests are not counted. TABLE 14-15 describes the counters for memory controller read events.

**TABLE 14-15** Counters for Memory Controller Read Events

| Counter    | Description   |
|------------|---|
| MC_reads_0 | [PICL 10.0000] Counts read requests completed to memory bank 0. |
| MC_reads_1 | [PICL 10.0001] Counts read requests completed to memory bank 1. |
| MC_reads_2 | [PICL 10.0010] Counts read requests completed to memory bank 2. |
| MC_reads_3 | [PICL 10.0011] Counts read requests completed to memory bank 3. |

#### 14.6.1.1 Memory Controller Write Request Events

TABLE 14-16 describes the counters for memory controller write events.

**TABLE 14-16** Counters for Memory Controller Write Events

| Counter     | Description  |
|-------------|--|
| MC_writes_0 | [PICU 10.0000] Counts write requests completed to memory bank 0. |
| MC_writes_1 | [PICU 10.0001] Counts write requests completed to memory bank 1. |
| MC_writes_2 | [PICU 10.0010] Counts write requests completed to memory bank 2. |
| MC_writes_3 | [PICU 10.0011] Counts write requests completed to memory bank 3. |

## 14.6.1.2 Memory Request Stall Cycles

The stall cycles may be generated due to bus contention, a bank being busy, data availability for a write, etc. TABLE 14-17 describes the counters for memory stall cycles.

**TABLE 14-17** Counters for Memory Controller Stall Cycles

| Counter     | Description  |
|-------------|--|
| MC_stalls_0 | [PICL 10.0100] Counts clock cycles that requests were stalled in the MCU queues because bank 0 was busy with a previous request. |
| MC_stalls_1 | [PICU 10.0100] Counts clock cycles that requests were stalled in the MCU queues because bank 1 was busy with a previous request. |
| MC_stalls_2 | [PICL 10.0101] Counts clock cycles that requests were stalled in the MCU queues because bank 2 was busy with a previous request. |
| MC_stalls_3 | [PICU 10.0101] Counts clock cycles that requests were stalled in the MCU queues because bank 3 was busy with a previous request. |

---

## 14.7 Data Locality Counters for Scalable Shared Memory Systems

There are four data locality performance event counters in the UltraSPARC III Cu processor. These event counters are provided to improve the ability to monitor and exploit performance in Scalable Shared Memory (SSM) systems where there are multiprocessor system clusters using Shared Memory Protocol (SMP) that are tied to other clusters using fabric interconnect utilizing the SSM architecture. TABLE 14-18 describes the counters for data locality events.

**TABLE 14-18** Counters for Data Locality Events

| Counter             | Description  |
|---------------------|--|
| EC_miss_local       | [PICL 01.1010] Counts any transaction to an LPA for which the processor issues an RTS/RTO/RS transaction.                                      |
| EC_miss_mtag_remote | [PICL 01.1011] and [PICU 10.1000] Counts any transaction to an LPA address in which the processor is required to generate a retry transaction. |

**TABLE 14-18** Counters for Data Locality Events *(Continued)*

| Counter        | Description   |
|----------------|---|
| EC_miss_remote | [PICU 10.1001] Counts the events triggered whenever the processor generates a remote (R_*) transaction and the address is to a non-LPA portion (remote) of the physical address space, or an R_WS transaction due to block store (BST)/block store commit (BSTC) to any address space (LPA or non-LPA), or an R_RTO due to Store/Swap request on O <sub>s</sub> state to LPA space. |
| EC_wb_remote   | [PICL 01.1001] Counts the retry event when any victimization for which the processor generates an R_WB transaction to non-LPA address region.<br><br>In practice, these are all NUMA cases, since for Coherence Memory Replication (CMR or COMA), the protocol insures that the processor only generates WB transactions.   |

### *SSM Systems*

Typically, four to six local processors are in a system cluster and have their own local memory subsystem(s). They use an SMP to maintain data coherency amongst themselves. Data coherency is maintained between system clusters using a directory based SSM data coherency mechanism to insure data coherency across systems with a large number of processors.

The data locality event counters are only valid for Shared System Memory architectures in SSM mode.

## 14.7.1 Event Tree

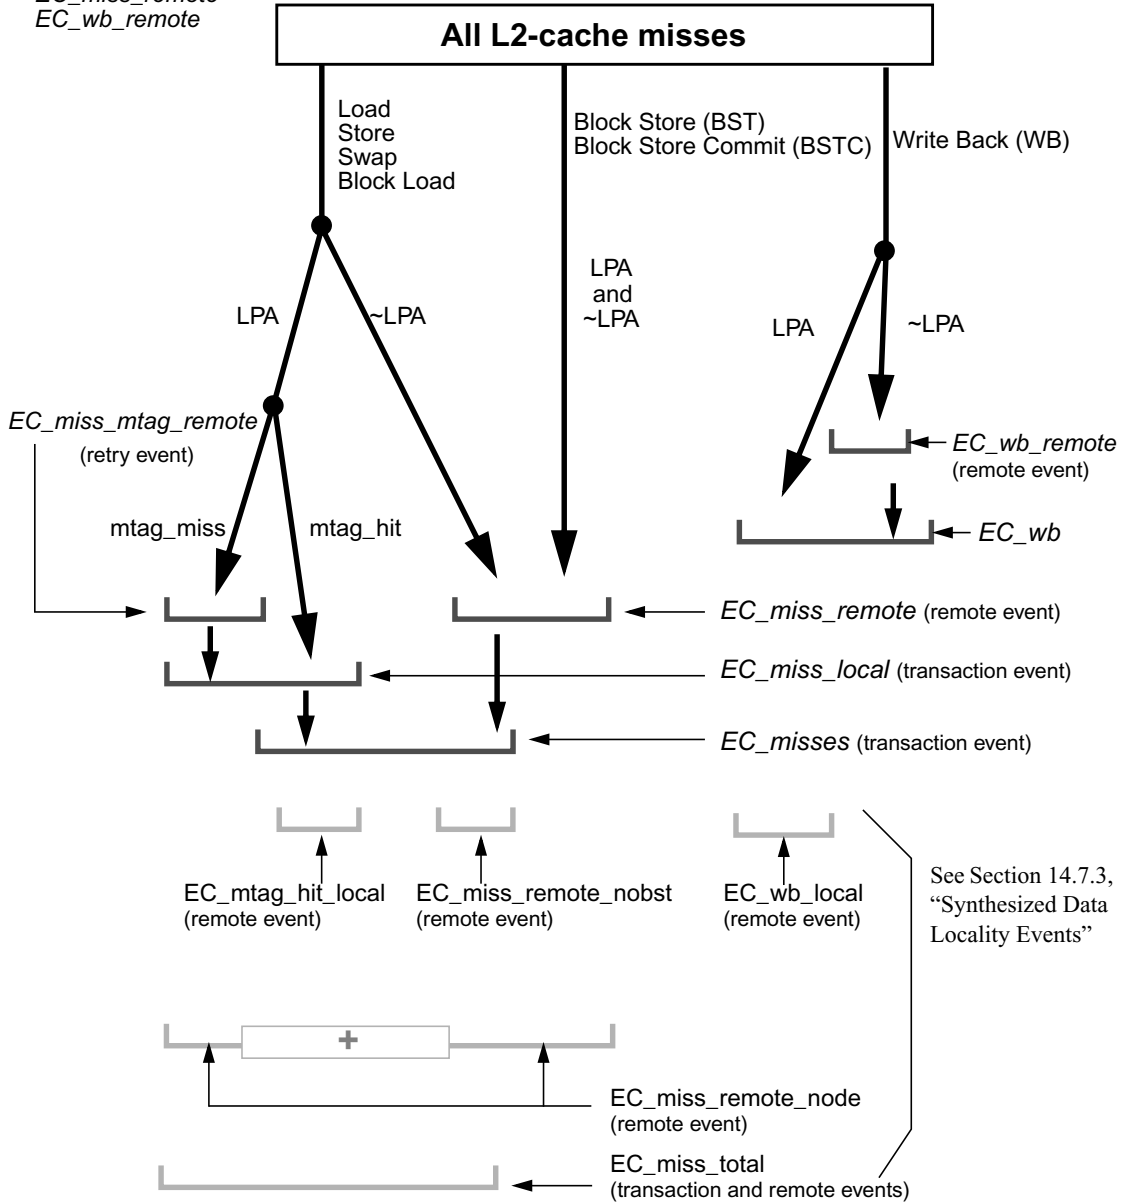
The event and cycle stall counters are illustrated in FIGURE 14-6 on page 14-381. The diagram includes actual counters and synthesized counts.



**Data Locality Events**

*EC\_miss\_local*  
*EC\_miss\_mtag\_remote*  
*EC\_miss\_remote*  
*EC\_wb\_remote*

LPA = Local Processor Physical Address  
 ~LPA = Remote Processor Physical Address



**FIGURE 14-6** Data Locality Event Tree for L2-cache Misses

## 14.7.2 Data Locality Event Matrix

TABLE 14-19 shows the data locality event matrix.

**TABLE 14-19** Data Locality Events

| MODE        | Combined State | Processor action    |                     |                    |               |                         |            |               |
|-------------|----------------|---------------------|---------------------|--------------------|---------------|-------------------------|------------|---------------|
|             |                | Load                | Store/ Swap         | Block Load         | Block Store   | Block Store with Commit | Write Back |               |
| LPA         | I              | miss:<br>RTS        | miss:<br>RTO        | miss:<br>RS        | miss:<br>R_WS | miss:<br>R_WS           | none       |               |
|             | E              | hit                 | hit:<br>E→M         | hit                | hit:<br>E→M   |                         |            |               |
|             | S              |                     | mtag miss:<br>RTO   |                    | miss:<br>R_WS |                         |            | WB            |
|             | O              |                     | mtag miss:<br>R_RTO |                    |               |                         |            |               |
|             | Os             |                     |                     |                    |               |                         |            |               |
|             | M              |                     | hit                 |                    |               |                         |            | WB            |
| LPA Retried | I              | mtag miss:<br>R_RTS | mtag miss:<br>R_RTO | mtag miss:<br>R_RS | invalid       |                         | none       |               |
|             | E              | invalid             |                     |                    |               |                         |            |               |
|             | S              | invalid             | mtag miss:<br>R_RTO | invalid            |               |                         | none       |               |
|             | O              |                     |                     |                    |               |                         |            |               |
|             | Os             |                     |                     |                    |               |                         |            |               |
|             | M              |                     |                     | invalid            |               |                         |            |               |
| ~LPA        | I              | miss:<br>R_RTS      | miss:<br>R_RTO      | miss:<br>R_RS      | miss:<br>R_WS | miss:<br>R_WS           | none       |               |
|             | E              | hit                 | hit:<br>E→M         | hit                | hit:<br>E→M   |                         |            |               |
|             | S              | hit                 | mtag miss:<br>R_RTO | hit                | miss:<br>R_WS |                         |            | miss:<br>R_WB |
|             | O              |                     |                     |                    |               |                         |            |               |
|             | Os             |                     |                     |                    |               |                         |            |               |
|             | M              |                     |                     |                    |               |                         |            |               |

### *LPA Retried Events*

Retry is to issue a R\_\* transaction for an RTS/RTO/RS transaction that gets unexpected MTag from the SSM system interconnect (for example, cache state = O and mtag state = gS). A retry takes place in LPA.

## 14.7.3 Synthesized Data Locality Events

The Data Locality event assignments allow the software to create synthetic events based on arithmetic combinations of events assigned to PICL and PICU, as shown in TABLE 14-20.

**TABLE 14-20** Synthesized Data Locality Events

| Synthesized Event  |   | PICL                     | PICU                     |
|--|---|--------------------------|--------------------------|
| <b>EC_wb_local</b>   | = | +<br>EC_wb               | -<br>EC_wb_remote        |
| <b>EC_miss_remote_nobst</b><br>(remote misses excluding<br>BST/BSTC) | = | +<br>EC_misses           | -<br>EC_miss_local       |
| <b>EC_miss_total</b><br>(L2-cache misses + retries)                  | = | +<br>EC_misses           | +<br>EC_miss_mtag_remote |
| <b>EC_miss_remote_node</b><br>(CMR, NUMA)                            | = | +<br>EC_miss_remote      | +<br>EC_miss_mtag_remote |
| <b>EC_mtag_hit_local</b>   | = | -<br>EC_miss_mtag_remote | +<br>EC_miss_local       |

### *EC\_miss\_total*

EC\_miss\_total counts all EC misses, which is EC\_misses plus all retries. Retry means you have two transactions for each miss (first it misses MTag and then reissued).

---

## 14.8 Miscellaneous Counters

### 14.8.1 System Interface Events and Clock Cycles

System interface statistics are collected through the counters listed in TABLE 14-21.

**TABLE 14-21** Counters for System Interface Statistics

| Counter     | Description  |
|-------------|--|
| SI_snoop    | [PICL 01.0001] Counts snoops from remote processor(s) including RTS, RTSR, RTO, RTOR, RS, RSR, RTSM, and WS. |
| SI_ciq_flow | [PICL 01.0010] Counts system clock cycles when the flow control (PauseOut) signal is asserted.               |
| SI_owned    | [PICL 010011] Counts events where owned_in is asserted on bus requests from the local processor.             |

### 14.8.2 Software Events

Software statistics are collected through the counters listed in TABLE 14-22.

**TABLE 14-22** Counters for Software Statistics

| Counter   | Description   |
|-----------|---|
| SW_count0 | [PICL 01.0100] Counts software-generated occurrences of sethi %hi(0xfc00), %g0 instruction. |
| SW_count1 | [PICU 01.1100] Counts software-generated occurrences of sethi %hi(0xfc00), %g0 instruction. |

---

**Note** – Both counters measure the same event; thus, the count can be programmed to be read from either the PICL or the PICU register.

---

## 14.8.3 Floating-Point Operation Events

Floating-point operation statistics are collected through the counters listed in TABLE 14-23.

**TABLE 14-23** Counters for Floating-Point Operation Statistics

| Event Counter      | Description   |
|--------------------|---|
| FA_pipe_completion | [PICL 01.1000] Counts instructions that complete execution on the Floating-point/Graphics ALU pipelines.      |
| FM_pipe_completion | [PICL 10.0111] Counts instructions that complete execution on the Floating-point/Graphics Multiply pipelines. |

---

## 14.9 PCR.SL and PCR.SU Encodings

TABLE 14-24 lists PCR . SL and PCR . SL selection bit field encoding. Shaded blocks show SL and SU field duplications with light shading.

**TABLE 14-24** PIC.SL and PIC.SU Selection Bit Field Encoding

| PCR.SL and PCR.SU Encodings | PICL Event Selection | PICU Event Selection |
|-----------------------------|----------------------|----------------------|
| 00.0000                     | Cycle_cnt            | Cycle_cnt            |
| 00.0001                     | Instr_cnt            | Instr_cnt            |
| 00.0010                     | Dispatch0_IC_miss    | Dispatch0_mispred    |
| 00.0011                     | Dispatch0_br_target  | IC_miss_cancelled    |
| 00.0100                     | Dispatch0_2nd_br     | Re_DC_missovhd       |
| 00.0101                     | Rstall_storeQ        | Re_FPU_bypass        |
| 00.0110                     | Rstall_IU_use        | Re_DC_miss           |
| 00.0111                     | <i>Reserved</i>      | Re_EC_miss           |
| 00.1000                     | IC_ref               | IC_miss              |
| 00.1001                     | DC_rd                | DC_rd_miss           |
| 00.1010                     | DC_wr                | DC_wr_miss           |
| 00.1011                     | <i>Reserved</i>      | Rstall_FP_use        |
| 00.1100                     | EC_ref               | EC_misses            |
| 00.1101                     | EC_write_hit_RTO     | EC_wb                |
| 00.1110                     | EC_snoop_inv         | EC_snoop_cb          |
| 00.1111                     | EC_rd_miss           | EC_ic_miss           |
| 01.0000                     | PC_port0_rd          | Re_PC_miss           |
| 01.0001                     | SI_snoop             | ITLB_miss            |
| 01.0010                     | SI_ciq_flow          | DTLB_miss            |

**TABLE 14-24** PIC.SL and PIC.SU Selection Bit Field Encoding (Continued)

| <b>PCR.SL and<br/>PCR.SU<br/>Encodings</b> | <b>PICL Event Selection</b> | <b>PICU Event Selection</b> |
|--|-----------------------------|-----------------------------|
| 01.0011                                    | SI_owned                    | WC_miss                     |
| 01.0100                                    | SW_count0                   | WC_snoop_cb                 |
| 01.0101                                    | IU_Stat_Br_miss_taken       | WC_scrubbed                 |
| 01.0110                                    | IU_Stat_Br_count_taken      | WC_wb_wo_read               |
| 01.0111                                    | Dispatch_rs_mispred         | <i>Reserved</i>             |
| 01.1000                                    | FA_pipe_completion          | PC_soft_hit                 |
| 01.1001                                    | EC_wb_remote                | PC_snoop_inv                |
| 01.1010                                    | EC_miss_local               | PC_hard_hit                 |
| 01.1011                                    | EC_miss_mtag_remote         | PC_port1_rd                 |
| 01.1100                                    | <i>Reserved</i>             | SW_count1                   |
| 01.1101                                    | <i>Reserved</i>             | IU_Stat_Br_miss_untaken     |
| 01.1110                                    | <i>Reserved</i>             | IU_Stat_Br_count_untaken    |
| 01.1111                                    | <i>Reserved</i>             | PC_MS_miss                  |
| 10.0000                                    | MC_reads_0                  | MC_writes_0                 |
| 10.0001                                    | MC_reads_1                  | MC_writes_1                 |
| 10.0010                                    | MC_reads_2                  | MC_writes_2                 |
| 10.0011                                    | MC_reads_3                  | MC_writes_3                 |
| 10.0100                                    | MC_stalls_0                 | MC_stalls_1                 |
| 10.0101                                    | MC_stalls_2                 | MC_stalls_3                 |
| 10.0110                                    | <i>Reserved</i>             | Re_RAW_miss                 |
| 10.0111                                    | <i>Reserved</i>             | FM_pipe_completion          |
| 10.1000                                    | <i>Reserved</i>             | EC_miss_mtag_remote         |
| 10.1001                                    | <i>Reserved</i>             | EC_miss_remote              |
| 10.1010 -<br>11.1111                       | <i>Reserved</i>             | <i>Reserved</i>             |

# Processor Optimization

---

---

## 15.1 Introduction

A significantly larger performance gain can be obtained if the code is re-compiled using a compiler specifically designed for a specific processor. Several features are provided on the UltraSPARC processors that can only be taken advantage of by using modern compiler technology. This technology was previously unavailable primarily due to insufficient hardware support to justify its development.

The front end of the processor consists of the following components:

- Prefetch Unit
- Instruction cache
- Next Field RAM
- The Branch and Set Prediction Logic
- Return address stack

The role of the front end in processor performance optimization is to supply as many valid instructions as possible to the grouping logic and eventually to the functional units, including the ALUs, the floating-point adder, the branch unit, the load/store pipe, to name a few.

Optimization is also about prefetching data, as cache misses are cycle intensive.

---

## 15.2 Instruction Stream Issues

The following section addresses the following issues:

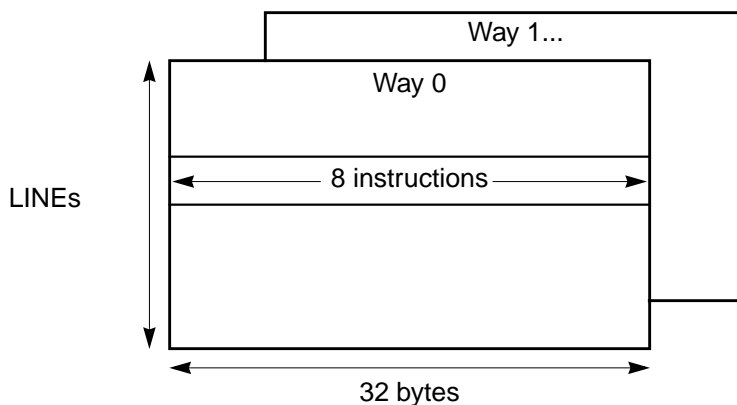
- Instruction Alignment

- Instruction Cache Timing
- Executing Code Out of the Level 2 Cache
- TLB Misses
- Instruction Cache Utilization
- Handling of CTI Couples
- Mispredicted Branches
- Return Address Stack

## 15.2.1 Instruction Alignment

### 15.2.1.1 Instruction Cache Organization

The I-cache is organized as a 4-way set associative cache, with each set containing a multiple of eight-instruction lines (see FIGURE 15-1). Depending on its address, for each line of 8 instructions, up to 4 instructions are sent to the instruction buffer. If the address points to one of the last three instructions in the line, only this last instruction and instructions (0-2) from the end of the line are selected. Consequently, on average for random accesses, 3.25 instructions are fetched from the I-cache. For sequential accesses, the fetching rate (4 instructions per cycle) matches the consuming rate of the pipeline (up to 4 instructions per cycle).



**FIGURE 15-1** Instruction Cache Organization



### 15.2.1.2 Branch Target Alignment

Given the restriction mentioned above regarding the number of instructions fetched from an I-cache access, it is desirable to align branch targets so that enough instructions are fetched to match the number of instructions issued in the first group of the branch target. The following examples highlight the logic behind branch target alignment:

- If the compiler scheduler indicates that the target can be grouped with only one more instruction, the target should be placed anywhere in the line except in the last slot, since only one instruction would be fetched in that case. It may be beneficial to fetch more instructions, if possible.
- If the target is accessed from more than one place, it should be aligned so that it accommodates the largest possible group (first 5 instructions of a line).
- If accesses to the I-cache are expected to miss, it may be desirable to align targets on a 32-byte boundary, or at least the front end of a block, so that 4 instructions are forwarded to the next stage. Such an alignment helps assure that the maximum number of instructions can be processed between cache misses, assuming that the code does not branch out of the sequence of instructions. In fact 64-byte alignment can help instruction prefetch.

In general, it is best to align for maximum fetching by aligning branch targets on 4 instruction (16-byte) boundaries. This can help ensure that the fetch bandwidth matches the issue width, which is a maximum of 4 instructions in the case of the UltraSPARC III processor.

### 15.2.1.3 Branch Optimization

The UltraSPARC processors favor branch not taken conditionals. Regardless of this preference, the instruction issue remains the same and the fetch is optimized.

### 15.2.1.4 Impact of the Delay Slot on Instruction Fetch

Most Control Transfer Instructions (CTIs) are actually *delayed* CTIs taking effect 1 instruction after the CTI, with the resulting lag known as the *delay slot*. That is the instruction following the branch, or after CTI, is always executed regardless of where the CTI directs execution (unless annulling is used). If the last instruction of a line is a branch, the next sequential line in the I-cache must be fetched even if the branch predicted is taken, since the delay slot must be sent to the grouping logic. This leads to inefficient fetches, since an entire L2-cache access must be dedicated to fetching the missing delay slot. Therefore one should take care not to place delayed CTIs at the end of a cache line.

### 15.2.1.5 Instruction Alignment for the Grouping Logic

See Chapter 4 “Grouping Rules” for a description of grouping logic.

### 15.2.1.6 Impact of Instruction Alignment on Instruction Dispatch

It is important that no two branches are in the same fetch group. If there are two branches in the same group the second branch will end the group and will cause a re-fetch taking 2 cycles. To guarantee this does not happen in the case of uncertain instruction alignment, ensure that no 2 branches are within 4 instructions of each other.

## 15.2.2 Instruction Cache Timing

While accesses to the I-cache hit successfully, the pipeline rarely starves for instructions. In rare cases however, the Instruction Dispatch is unable to provide a sufficient number of instructions to keep the functional units busy. For example, a taken branch to a taken branch sequence without any instructions between the branches (except for the delay slot) could only be executed at a peak rate of two instructions per cycle.

An I-cache miss does not necessarily result in bubbles being inserted into the pipeline. Part of the I-cache miss processing, or even all of it, can be overlapped with the execution of instructions that are already in the instruction buffer and are waiting to be grouped and executed. Moreover, since the operation of the Instruction Dispatch is somewhat separated from the rest of the pipeline, the I-cache miss may have occurred when the pipeline was already stalled (for example, due to a multi-cycle integer divide, floating-point divide dependency, dependency on load data that missed the D-cache, etc.). This means that the miss (or part of it) may be transparent to the pipeline.

Because of the possibility of stalling the processor when the pipeline is waiting for new instructions, it is desirable to try to make code routines fit in the I-cache and avoid cache misses. The UltraSPARC processor provides instrumentation to profile a program and detect if instruction accesses generate a cache miss or a cache hit. By checkpointing the counters before and after a large section of code, combined with profiling the section of code, one can determine if the frequently executed functions generally hit or miss the I-cache.

## 15.2.3 Executing Code Out of the Level 2 Cache

Executing out of the L2-cache is much better than executing from main memory. From main memory, hardware can prefetch the next 8 instructions if the initial fetch was from the lower 32 of a 64-byte aligned boundary.

## 15.2.4 TLB Misses

The TLB contains the virtual page number and the associated physical page number of the most recently accessed pages.

A TLB miss is handled by software through the use of the TSB, and takes a large number of cycles. In order to minimize the frequency of TLB misses, the UltraSPARC processor provides a large number of entries in the TLB. Nonetheless, techniques are encouraged to further decrease the TLB miss cost.

### 15.2.4.1 Impact of the Annulled Slot

Grouping rules in Chapter 4 “Grouping Rules” describes how the UltraSPARC processor handles instructions following an annulling branch. In connection with these instructions, pay regard to the rules:

- Avoid scheduling `WR(PR, ASR)`, `SAVE`, `SAVED`, `RESTORE`, `RESTORED`, `RETURN`, `RETRY`, and `DONE` in the delay slot and in the first three groups following an annulling branch.

## 15.2.5 Conditional Moves vs. Conditional Branches

The `MOVCC` and `MOVR` instructions provide an alternative to conditional branches for executing short code segments. The UltraSPARC processor differentiates the two as follows:

- Conditional branches: Distancing the `SETCC` from `BICC` does not gain any performance. The penalty for a mispredicted branch is always 8-cycles. `SETCC`, `BICC`, and the delay slot can be grouped together (FIGURE 15-2).

```
setcc G  E  C  N1 N2 N3 W
bicc  G  E  C  N1 N2 N3 W
delay G  E  C  N1 N2 N3 W
```

**FIGURE 15-2** Handling of Conditional Branches

- Conditional moves: A use of the destination register for the `MOVCC` follows the same rule as a load-use. FIGURE 15-3 shows a typical example.

```
setcc G  E  C  N1 N2 N3 W
movcc  G  E  C  N1 N2 N3 W
use           G  E  C  N1 N2 N3 W
```

**FIGURE 15-3** Handling of MOVCC

If a branch is correctly predicted, the issue rate can be higher than that of a branch that is replaced by a conditional move. If a branch is not predictable the mis-predicted penalty is significantly higher than the extra latency of a conditional move.

## 15.2.6 Instruction Cache Utilization

Grouping blocks that are executed frequently can effectively increase the apparent size of the I-cache. Case studies show that, often, half of the I-cache entries are never executed. Placing rarely executed code out of a line containing a frequently executed block (identified by profiling) achieves a better I-cache utilization.

## 15.2.7 Handling of CTI couples

Placing a CTI into the delay slot of another CTI will disrupt the fetch and cost many cycles and should be avoided.

## 15.2.8 Mispredicted Branches

Correctly predicted conditional branches allow the processor to group instructions from subsequent basic blocks and continue to progress speculatively until the branch is resolved. The capability of executing instructions speculatively is a significant performance boost for the UltraSPARC processor.

## 15.2.9 Return Address Stack (RAS)

In order to speed up returns from subroutines invoked through `CALL` instructions, UltraSPARC processor dedicates a 8-deep stack to store the return address. Each time a `CALL` is detected, the return address is pushed onto this RAS (Return Address Stack). Each time a return is encountered, the address is obtained from the top of the stack and the stack is popped. The UltraSPARC III Cu processor considers a return to be a `JMPL` or `RETURN` with `rs1` equal to `%o7` (normal subroutine) or `%i7` (leaf subroutine). The RAS provides a guess for the target address, so that prefetching can continue even though the address calculation has not yet been performed. `JMPL` or `RETURN` instructions using `rs1` values other than `%o7` or `%i7` use the value on the top of the RAS for continuing prefetching, but they do not pop the stack.

To take full advantage of the RAS, one should follow the standard call and return conventions so that hardware can correctly predict the return addresses.

---


## 15.3 Data Stream Issues

The following section addresses the following issues:

- Data Cache Organization
- Data Cache Timing
- Data Alignment
- Using LDDF to Load Two Single-Precision Operands/Cycle
- Store Considerations
- Read-After-Write Hazards

### 15.3.1 Data Cache Organization

The D-cache is a mapped, virtually indexed, physically tagged (VIPT), write-through, non-write-allocating cache. It is logically organized as lines of 32 bytes.



**FIGURE 15-4** Logical Organization of Data Cache

## 15.3.2 Data Cache Timing

The latency of a load to the D-cache depends on the opcode. LDX and LDUW have 2-cycle load-to-use latency while all other loads have 3. For instance, if the first two instructions in the instruction buffer are a load and an instruction dependent on that load, the grouping logic will break the group after the load and a bubble will be inserted in the pipeline. It is very important to separate loads from their use.

## 15.3.3 Data Alignment

SPARC V9 requires that all accesses be aligned on an address equal to the size of the access. Otherwise a *mem\_address\_not\_aligned* trap is generated. This is especially important for double precision floating-point loads, which should be aligned on an 8-byte boundary. If misalignment is determined to be possible at compile time, it is better to use two LDF (load floating-point, single precision) instructions and avoid the trap. The UltraSPARC processor supports single-precision loads mixed with double-precision operations, so that the case above can execute without penalty (except for the additional load). If a trap does occur, the UltraSPARC processor dedicates a trap vector for this specific misalignment, which reduces the overall penalty of the trap.

Grouping load data is desirable, since a D-cache sub-block can contain either four properly aligned single-precision operands or two properly aligned double-precision operands (eight and four respectively for a D-cache line). This is desirable not only for improving the D-cache hit rate (by increasing its utilization density), but also for D-cache misses where, for sequential accesses, one out of two requests to the L2-cache can be eliminated.

### 15.3.3.1 Using LDDF to Load Two Single-Precision Operands/Cycle

The UltraSPARC processor supports single cycle 8-byte data transfers into the floating-point register file for LDDF. Wherever possible, applications that use single-precision floating-point arithmetic heavily should organize their code and data to replace two LDFs with one LDDF. This reduces the load frequency by approximately one half, and cuts execution time considerably.

## 15.3.4 Store Considerations

The store on the UltraSPARC processor is designed so that stores can be issued even when the data is not ready. More specifically, a store can be issued in the same group as the instruction producing the result. The address of a store is buffered until the data is eventually available. Once in the store buffer, the store data is buffered until it can be completed.

The write cache can be used to exploit locality (both temporal and spatial) in the write stream. Try to organize data to exploit the locality by not having more than 4 data streams because of the 4-way set associative aliasing.

## 15.3.5 Read-After-Write Hazards

See Chapter 9 “Read-After-Write (RAW) Bypassing” for rules on RAW hazards.





# Prefetch

---

This chapter contains the following sections:

- Prefetch Cache
- Software Prefetch
- Hardware Prefetch
- FP/VIS Load Instruction Miss Fetch

Prefetches can be used to hide memory latency, increase memory-level parallelism (overlap memory accesses), and hide L2 latency for floating-point loads.

The Prefetch mechanism contains an eight entry prefetch request queue to hold up to eight outstanding prefetch requests to L2-cache and beyond. Prefetches are generated by software prefetch instructions and by an autonomous hardware prefetch engine. Prefetches may bring data from memory into L2-cache. Prefetches may also bring data (either from L2-cache or memory) all the way to the Level-1 P-cache. Only floating-point loads may get data from the P-cache.

---

## 16.1 Prefetch Cache

The P-cache is virtually indexed and virtually tagged for CPU cache reads and never contains modified data. It is physically indexed and physically tagged for system snooping of the cache. The prefetch cache is 2 KB, 4-way set associative with 64 bytes per line.

The P-cache is accessed in parallel with the D-cache for floating-point/VIS loads. If the P-cache contains the data and the D-cache does not, the data will be provided by the P-cache. Integer loads cannot get data for the P-cache.

## 16.1.1 P-Cache Data Flow

Data is put into the P-cache under the following conditions:

- Hardware Prefetch (64 bytes)
- FP Load Miss (32 bytes)
- Some variants of Software Prefetch (64 bytes)

The hardware prefetch mechanism and the software prefetch instructions fetch data into the P-cache but not the D-cache. The FP/VIS Load Instruction Miss Fetch puts data in both the D-cache and P-cache.

Block load data is never loaded into the P-cache. The P-cache can hold cacheable data that is not in the D-cache, and vice versa. The P-cache does not contain modified data. Stores that hit the P-cache invalidate the 64-byte P-cache line.

## 16.1.2 Fetched\_Mode Tag Bit

Each P-cache line contains a `fetched_mode` bit that indicates how the P-cache line entry was installed: software prefetch instruction or hardware prefetch mechanism.

Successful software prefetches set this bit to indicate the cache line was installed using a software prefetch instruction. This will inhibit the hardware prefetch mechanism from operating when it might otherwise operate using the cache line information. The hardware prefetch logic will not prefetch when the fetch mode bit of the current cache line is set to software prefetched data. This helps keep the prefetching from interfering with successful software prefetching.

---

## 16.2 Software Prefetch

The SPARC V9 instruction set contains data prefetch instructions. These instructions allow software to give warning to hardware that it will be using data in the future. This gives hardware the opportunity to potentially bring data closer to hide the latency of the memory subsystem when the data is actually needed. There are different types of prefetches that enable software to tell how the data will be used.

The UltraSPARC III Cu processor has extensive support for making use of prefetch directives from software. The prefetching can be used to either hide latency or memory operations or to overlap memory operations. In the UltraSPARC III Cu processor, different types of prefetches are handled differently. Prefetches may do one of the following:

- Bring data into L2-cache only

- Bring data into P-cache only
- Bring data into both P-cache and L2-cache
- Try to acquire ownership of a line for future writes

The software prefetch instructions are enabled by setting both the `pcache_enable` and the `software_prefetch_enable` (`SPE`) (bits 45 and 43, DCUCR register). If the P-cache is disabled or the `SPE` is zero, then all software prefetch instructions are ignored as they go through the pipeline.

The P-cache responds to software prefetches by fetching the requested data into the P-cache and/or L2-cache, as specified by the prefetch instruction.

## 16.2.1 Software Prefetch Instructions

Software prefetch instructions can prefetch data into the P-cache, L2-cache, or both, depending on the type of prefetch instruction executed. TABLE 16-1 lists the types of software prefetch instructions.

**TABLE 16-1** Types of Software Prefetch Instructions

| <b>fcn Value (hex)</b> | <b>Instruction Type</b>          | <b>Prefetch (64 bytes of data) into:</b>           | <b>Instruction Strength UltraSPARC III Cu</b> | <b>Request Exclusive Ownership</b> |
|------------------------|----------------------------------|--|---|------------------------------------|
| 00                     | Prefetch read many               | P-cache and L2-cache                               | Weak  | No                                 |
| 01                     | Prefetch read once               | P-cache only                                       | Weak  | No                                 |
| 02                     | Prefetch write many              | L2-cache only                                      | Weak  | Yes                                |
| 03                     | Prefetch write once <sup>1</sup> | L2-cache only                                      | Weak  | No                                 |
| 04                     | <i>Reserved</i>                  | Undefined  |   |                                    |
| 05 – 0F                | <i>Reserved</i>                  | Undefined  |   |                                    |
| 10                     | Prefetch invalidate              | Invalidates a P-cache line, no data is prefetched. |   | N/A                                |
| 11 – 13                | <i>Reserved</i>                  | Undefined  |   |                                    |
| 14                     | Same as fcn = 00                 |  | Weak <sup>2</sup>                             | No                                 |
| 15                     | Same as fcn = 01                 |  | Weak <sup>2</sup>                             | No                                 |
| 16                     | Same as fcn = 02                 |  | Weak <sup>2</sup>                             | Yes                                |
| 17                     | Same as fcn = 03                 |  | Weak <sup>2</sup>                             | No                                 |
| 18 – 1F                | <i>Reserved</i>                  | Undefined  |   |                                    |

1. Although the name is “prefetch write once,” the actual use is prefetch to L2-cache for a future read.

2. These Weak instructions may be implemented as strong in future implementations.

**Instruction Strength.** For each of the prefetch types, the software instruction can either be a weak or a strong software prefetch instruction. This allows software to tell hardware how confident it is that a prefetch will be helpful. Hardware may try harder to make sure that a strong prefetch gets executed. In the case of the UltraSPARC III Cu processors, all software prefetch instructions are weak, allowing the processor to more easily cancel their execution.

## 16.2.2 Software Prefetch Instruction Uses the MS Pipeline

All software prefetch instructions are executed by the MS pipeline because the MMU/TLB structure is required and it is dedicated to the load and store instructions that are executed in the MS pipeline.

The software prefetch instruction can be grouped with a second FP/VIS load instruction that executes in the A0 or A1 pipeline, but the software prefetch instruction will always execute in the MS pipeline. This means the processor never executes two software prefetch instructions within one instruction group.

## 16.2.3 Cancelling Software Prefetch

The processor will cancel a prefetch instruction under various conditions that depend on the instruction strength. The following four conditions can cancel a software prefetch instruction.

Regardless of software prefetch instruction strength, the prefetch is cancelled if:

- The data already exists in the appropriate cache, *or*
- The prefetch address is the same as one of the outstanding prefetches.

Weak software prefetch instruction strengths are also cancelled when:

- The prefetch instruction misses the D-TLB, *or*
- The prefetch queue is full (there are already too many outstanding prefetches).

## 16.2.4 Prefetch Instruction Variants

### *Prefetch for Several Reads*

The data installs in the P-Cache and L2-cache. In cases that a memory request is required to get the line, a Read to Share (RTS) is issued to the system. This is typically used when the data will be read more than once.

### *Prefetch for One Read*

The data installs in the P-Cache only. It does not install in the L2-cache if the line is not already there. In cases that a memory request is required to get the line, an RTS is issued to the system. This is typically used for streaming data that is only used for a short period.

---

**Note** – Prefetch for one read can be used to stream special non-coherent data. A Prefetch Invalidate should be used to flush such data to maintain correct behavior.

---

### *Prefetch for Several Writes*

The data installs a 64-byte line in L2-cache, if not present. In the case that a memory request is required, a Read to Own (RTO) is issued. This variant is typically used to prefetch for one or more writes.

If the prefetch line is already in the L2-cache, then the prefetch instruction will be cancelled even if the line is in a shared state. No RTO request will be sent out and no state change of the 64-byte sub-block in the L2-cache is made.

### *Prefetch for One Write*

The data installs a 64-byte line in L2-cache, if not present. In the case that a memory request is required, an RTS is issued. This is typically used for prefetching integer data from memory or for a two stage prefetch that separates prefetch to L2-cache and prefetch from L2-cache to P-cache.

## 16.2.5 General Comments

- *Block Load (BLD) data bypasses the P-cache and D-cache.*

The block load data is always supplied by the L2-cache and written into FP/VIS register file. The P-cache is used as a staging area utilizing eight 64-bit registers to hold the data that is forwarded to the FP/VIS register file. Using BLD can corrupt the P-cache. It is recommended that you use the Solaris<sup>TM</sup> bcopy routine for transferring large amounts of data.

- *P-cache data is not included in L2-cache.*

The prefetch-read-once software prefetch instruction puts data into the P-cache from the L2-cache or memory, but does not install it in L2-cache.

- *Stores invalidate prefetch cache lines.*

Stores invalidate P-cache lines that reference the same store address to maintain cache consistency. There is no mechanism to modify data in the P-cache, so code loops must be careful to place any required store instructions after the P-cache data line is no longer needed.

## 16.2.6 Code Example

CODE EXAMPLE 16-1 is an example of two memory operations per cycle, double-precision, vector-multiply code.

**CODE EXAMPLE 16-1** Two Memory Operations Per Cycle Code

```

! %i0 -> j
! %i1 -> c
! %i2 -> a
! %i3 -> b
! %i4 -> N
! %f32 -> K

! prefetch fcn=3 - one write, used to prefetch "c" into the E$
! prefetch fcn=1 - one read, used to prefetch "a" and "b" into
the P$

loop:
    prefetch [%i1+64],3                ! prefetch in first memory
slot
    lddf      [%i2],%f0                ! 2nd load from P$
    fmuld    %f32,%f24,%f24
    fadd     %f16,%f18,%f16

    stdf     %f8, [%i1+16]            ! stf in first memory slot
    add      %i1,64,%i1
    lddf     [%i3],%f2                ! 2nd load from P$

    prefetch [%i2+64],1
    lddf     [%i2+8],%f4
    fmuld    %f32,%f28,%f28
    fadd     %f20,%f22,%f20

    stdf     %f12, [%i1-40]
    lddf     [%i3+8],%f6

    prefetch [%i3+64],1
    lddf     [%i2+16],%f8
    fmuld    %f32,%f0,%f0

```

**CODE EXAMPLE 16-1** Two Memory Operations Per Cycle Code (*Continued*)

```
fadd    %f24,%f28,%f24

stdf    %f16, [%11-32]
lddf    [%13+16],%f10

lddf    [%12+24],%f12
fmul    %f32,%f4,%f4
fadd    %f28,%f30,%f28
stdf    %f20, [%11-24]

lddf    [%13+24],%f14
lddf    [%12+32],%f16
fmuld   %f32,%f8,%f8
fadd    %f0,%f2,%f0

stdf    %f24, [%11-16]
lddf    [%13+32],%f18

lddf    [%12+40],%f20
fmuld   %f32,%f12,%f12
fadd    %f4,%f6,%f4
stdf    %f28, [%11-8]

lddf    [%13+40],%f22
subcc   %10,8,%10

lddf    [%12+48],%f24
fmuld   %f32,%f16,%f16
fadd    %f8,%f10,%f8
stdf    %f0, [%11]

lddf    [%13+48],%f26
lddf    [%12+56],%f28

lddf    [%13+56],%f30
fmuld   %f32,%f20,%f20
fadd    %f14,%f12,%f12
stdf    %f4, [%11+8]

add     %13,64,%13
bgt     loop
add     %12,64,%12
```

---

## 16.3 Hardware Prefetch

---

**Note** – Some versions of the UltraSPARC III Cu processor may have a hardware prefetch mechanism as described in this section.

---

The hardware prefetch mechanism monitors the dispatching of floating-point load instructions. If there is a P-cache hit and the state of the line indicates that it was not brought in by a software prefetch and that a hardware prefetch for the next line has not already been generated, then the hardware attempts to prefetch the next 64-byte cache line from the L2-cache. If the request misses in the L2-cache, then the prefetch request is cancelled. The hardware prefetch mechanism triggers regardless of which word(s) position within the 64 bytes is loaded; they do not need to be sequentially accessed to trigger the prefetch mechanism.

The hardware prefetch mechanism is enabled when the `pcache_enable` and `hardware_prefetch_enable` (bits 45 and 44, DCUCR register) are both set.

The hardware prefetch from the L2-cache to the P-cache can get needed data to the P-cache well in advance of its use and hence reduce the L2-cache latency.

### 16.3.1 Cancelling Hardware Prefetches

Hardware prefetches are cancelled when they encounter one of the following conditions:

- The data already exists in the P-cache.
- The prefetch queue is full, *or*
- The hardware prefetch address is the same as one of the outstanding prefetches.
- The prefetch address is not in the same 8 KB boundary as the line that initiated the prefetch.
- The request misses the L2-cache.

Hardware prefetch will return 64 bytes of data to P-cache only if it hits the L2-cache. No more than one speculative prefetch is allowed to be generated from each prefetch cache line. Once a cache hit has caused a speculative prefetch, subsequent hits to that line will not cause redundant speculative prefetches.



## 16.3.2 FP/VIS Load Instruction Miss Fetch

When an FP/VIS load instruction misses in the D-cache and P-cache, the hardware fetches 32 bytes of data from the L2-cache and fills it into the D-Cache, regardless of the load instruction data size. If hardware prefetch is enabled, then the 32 bytes are also installed into the P-cache. This is how hardware prefetch can get initiated. If the request misses the L2-cache, then 64 bytes are fetched from main memory and installed in the L2-cache first.



## IEEE 754-1985 Standard

---

The implementation of the floating-point unit for standard and nonstandard operating modes are described in this chapter.

This chapter defines debug and diagnostics support in these sections:

- Introduction
- Floating-point Numbers
- IEEE Operations
- Traps and Exceptions
- IEEE Traps
- Underflow Operation
- IEEE NaN Operations
- Subnormal Operations

---

### 17.1 Introduction

#### 17.1.1 Floating-point Operations

Floating-point Operations (FPops) include the algebraic operations and usually do not include the specially treated FP Load/store, *FB/cc*, or the VIS instructions. The *FABS*, *FNEG*, and *FMOV* instructions are also treated separately from the algebraic operations.

## 17.1.2 Rounding Mode

The rounding mode of the FPU is determined either by the `FSR.RD` bit while in standard rounding mode or by the `GSR.IRND` bit when in interval arithmetic rounding mode. The rounding direction effects the result after any under or overflow condition is detected. Underflow is detected before rounding. The `FSR.RD` bit options are shown in TABLE 17-1.

TABLE 17-1 `FSR.RD` bit options

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

## 17.1.3 Nonstandard FP Operating Mode

The processor supports a nonstandard FP mode to facilitate in the handling of Subnormals by the hardware, avoiding a software trap to supervisor software. The FP operating mode is controlled by the `FSR.NS` bit. When `FSR.NS` = 1, nonstandard mode is selected. However, when `GSR.IM` = 1, interval arithmetic rounding mode is selected, then regardless of the `FSR.NS` bit the processor will be in standard mode.

## 17.1.4 Memory and Register Data Images

The floating-point values are represented in the `F` registers in the same way that they are represented in memory. Any conversions for ALU operations are completed within the FP execution unit. Load and store operations do not modify the register value.

VIS instructions (logical and move/copy operations) can be used with values generated by the FPU.

## 17.1.5 Subnormal Operations

Subnormal operations include operations with Subnormal number operands and situations where an operation without Subnormal number operands generate a Subnormal number result. The FPU response to Subnormal numbers is described in section 17.8, *Subnormal Operations*, on page 428.

## 17.1.6 FSR.CEXC and FSR.AEXC Updates

The current exception (*cexc*) and accrued exception (*aexc*) fields in the FSR register are described in section 17.5, *IEEE Traps*, on page 422.

In general:

- Only floating-point operations (FPops) will update *cexc* and only when an exceptional condition is detected. All other instructions will leave *cexc* unchanged.
- When an exception is detected, but the trap is masked, then the FPop will update the appropriate *aexc* field of the FSR register.

## 17.1.7 Prediction Logic

Prediction of overflow, underflow, and inexact traps is used in the hardware. Prediction always errors on the side of providing correct results when the hardware can and generating an exception when it cannot or is not sure.

Prediction of inexact occurs unless one of the operands is a Zero, NaN, or Infinity. When prediction occurs and the exception is enabled, system software will properly handle these cases and resume program execution. If the exception is not enabled, the result status is used to update the *FSR.aexc* and *FSR.cexc* bits of the FSR register.

---

# 17.2 Floating-point Numbers

The floating-point number types and their abbreviations are shown in TABLE 17-2. In general the IEEE 754-1985 Standard reserves exponent field values of all 0s and all 1s to represent special values in the standard's floating-point scheme.

TABLE 17-2 Floating-point Numbers

| Number Type    | Abbreviation | Data Representation |                        |                        |
|----------------|--------------|---------------------|------------------------|------------------------|
|                |              | Sign                | Exponent               | Fraction               |
| Zero           | 0            | 0 or 1              | 000...000              | 000...000              |
| Subnormal      | SbN          | 0 or 1              | 000...000              | 000...001 to 111...111 |
| Normal         | Normal       | 0 or 1              | 000...001 to 111...110 | 000...000 to 111...111 |
| Infinity       | Infinity     | 0 or 1              | 111...111              | 000...000              |
| Signalling NaN | SNaN         | 0 or 1              | 111...111              | 0xx...xxx              |
| Quiet NaN      | QNaN         | 0 or 1              | 111...111              | 1xx...xxx              |

## Zero

Zero is not directly representable if the straight format is followed, this is due to the assumption of a leading 1. To allow the number zero to yield a value of zero, the fraction (or mantissa) must be exactly zero. Therefore the number zero is special cased with exponent and fraction fields of zero. It is also important to note that -0 and +0 are considered to be distinct values, though they both compare as equal.

## SubNormal

If the exponent field is all 0s and the fraction field is non-zero then the value is a subnormal (denormalized) number. These numbers do not have an assumed leading 1 before the binary point. For single precision, these numbers are represented as  $(-1)^s \times 0.f \times 2^{-126}$ , in double precision the representation is  $(-1)^s \times 0.f \times 2^{-1022}$ . In both cases  $s$  is the sign bit and  $f$  is the fraction. Note that exponent and fraction fields of all 0s is the special representation of the number zero. From this point of view, the number zero can be considered a subnormal.

## Infinity

The values -infinity and +infinity are represented with an exponent field of all 1s and a fraction field of all 0s. The sign bit distinguishes between positive and negative infinities. The infinity representation is important as it allows operations to continue past overflow. Operations dealing with infinities are well defined by the IEEE 754-1985 Standard.

## Not a Number

The value NaN (Not a Number) is used to represent values that do not represent real numbers. The NaN exponent field is all 1s and the fraction field is non-zero. There are two categories of NaN; the QNaN (quiet NaN) and the SNaN (signalling NaN). A QNaN is a NaN with the most significant fraction field bit set. QNaN is allowed to freely propagate through most arithmetic operations; this NaN tends to appear when an operation produced mathematically undefined results. A SNaN fraction field significant bit is clear. The SNaN is used to signal an exception when it appears out of an operation being executed. Semantically QNaN can be considered to denote indeterminate operations, while SNaN indicates invalid operations.

## 17.2.1 Floating-point Number Line

The floating-point number line in FIGURE 17-1 represents the floating-point numbers used in the processor.

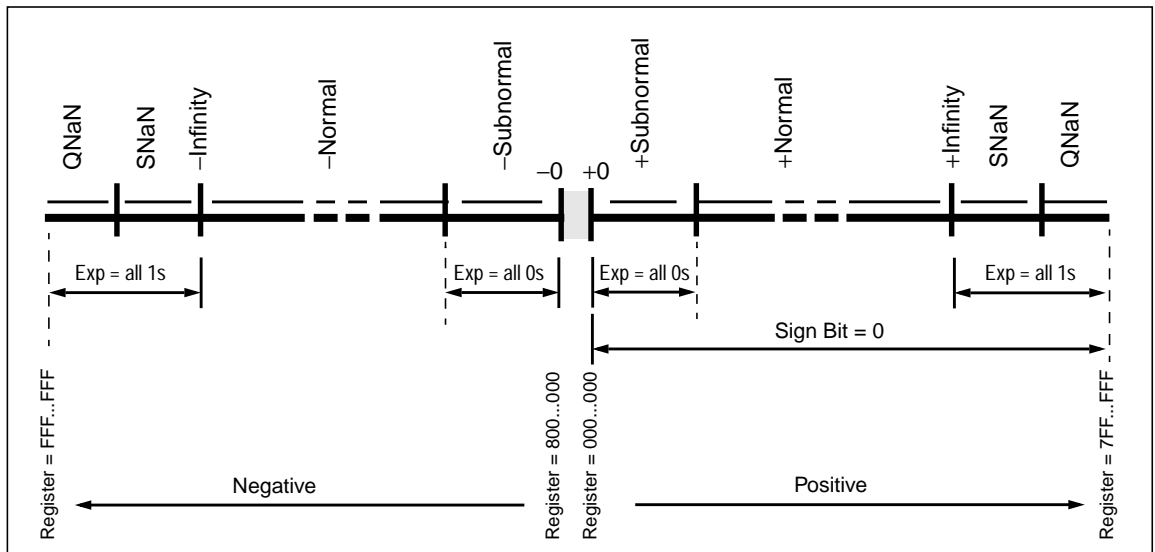


FIGURE 17-1 Floating-point Number Line

## 17.3 IEEE Operations

The response of each operation to operands with 0, Normal, Infinite, and NaN numbers are described in this section. The response to Subnormal numbers are described in section 17.8, *Subnormal Operations*, on page 428.

The result of each operation is concluded by one of the following:

- A number is written to the destination  $f$  register (rd).
- A number is written to the destination register **and** an IEEE flag is set.
- An IEEE flag is set **and** an IEEE trap is generated (rd is unchanged).

Each instruction is defined with one or more operands. Most instructions generate a result. The FCMP{E} instruction does not generate a result, instead it set the fccN bits.

## 17.3.1 Addition

The floating-point addition instruction is shown in TABLE 17-3.

**TABLE 17-3** Floating-point Addition

| <b>ADDITION Instruction</b><br><br><b>FADD</b> $rs_1, rs_2 [rs_2, rs_1] \rightarrow rd$ | <b>RESULT from the operation includes one or more of the following:</b> <ul style="list-style-type: none"> <li>• Number in f register, see <i>Trap Event</i> note, page 422.</li> <li>• Exception bit set, see TABLE 17-12.</li> <li>• Trap occurs, see abbreviations in TABLE 17-12.</li> <li>• Underflow/Overflow may occur.</li> </ul> |   |  |                                   |
|---|---|---|--|-----------------------------------|
|   | <b>Masked Exception, TEM=0</b>  |   | <b>Enabled Exception, TEM=1</b>          |                                   |
|   | <b>Destination register written (rd)</b>  | <b>Flag(s)</b>                              | <b>Destination register written (rd)</b> | <b>Flag(s), Trap</b>              |
| $+0, +0$  | +0  | no  | +0                                       | no                                |
| $+0, -0$  | +0 (FSR.RD=0,1,2)<br>-0 (FSR.RD=3)  | no  | +0 (FSR.RD=0,1,2)<br>-0 (FSR.RD=3)       | no                                |
| $-0, -0$  | -0  | no  | -0                                       | no                                |
| $0, +Normal$  | +Normal   | no  | +Normal                                  | no                                |
| $0, -Normal$  | -Normal   | no  | -Normal                                  | no                                |
| $0, +Infinity$  | +Infinity   | no  | +Infinity                                | no                                |
| $0, -Infinity$  | -Infinity   | no  | -Infinity                                | no                                |
| $Normal, +Infinity$   | +Infinity   | set ofc,<br>set ofa,<br>set nvc,<br>set nva | no                                       | set ofc,<br>set nvc,<br>ieee trap |
| $Normal, -Infinity$   | -Infinity   | set ofc,<br>set ofa,<br>set nvc,<br>set nva | no                                       | set ofc,<br>set nvc,<br>ieee trap |
| $+Normal, +Normal$  | May overflow, see 17.5.3  |   | May overflow, see 17.5.3                 |                                   |
| $+Normal, -Normal$  | Normal  |   | Normal                                   |                                   |
| $-Normal, +Normal$  | Normal  |   | Normal                                   |                                   |
| $-Normal, -Normal$  | May underflow, see 17.5.4   |   | May underflow, see 17.5.4                |                                   |
| $+Infinity, +Infinity$  | +Infinity   | no  | +Infinity                                | no                                |
| $+Infinity, -Infinity$  | QNaN  | set nvc,<br>set nva                         | no                                       | set nvc,<br>ieee trap             |
| $-Infinity, +Infinity$  | QNaN  | set nvc,<br>set nva                         | no                                       | set nvc,<br>ieee trap             |
| $-Infinity, -Infinity$  | -Infinity   | no  | -Infinity                                | no                                |



## 17.3.2 Subtraction

The floating-point subtraction instruction is shown in TABLE 17-4.

TABLE 17-4 Floating-point Subtraction

| SUBTRACTION Instruction<br><br><i>rs<sub>1</sub> - rs<sub>2</sub></i><br><br>FSUB <i>rs<sub>1</sub>, rs<sub>2</sub> → rd</i> | RESULT from the operation includes one or more of the following:  |   |                                   |                                   |
|--|---|---|-----------------------------------|-----------------------------------|
|  | <ul style="list-style-type: none"> <li>• Number in f register, see <i>Trap Event</i> note, page 422.</li> <li>• Exception bit set, see TABLE 17-12.</li> <li>• Trap occurs, see abbreviations in TABLE 17-12.</li> <li>• Underflow/Overflow may occur.</li> </ul> |   |                                   |                                   |
|  | Masked Exception, TEM=0   |   | Enabled Exception, TEM=1          |                                   |
|  | Destination register written (rd)   | Flag(s)                                     | Destination register written (rd) | Flag(s), Trap                     |
| <i>+0, +0</i>  | +0  | no  | +0                                | no                                |
| <i>+0, -0</i>  | -0  | no  | -0                                | no                                |
| <i>-0, +0</i>  | -0  | no  | -0                                | no                                |
| <i>-0, -0</i>  | +0  | no  | +0                                | no                                |
| <i>0, +Normal</i>  | -Normal   | no  | -Normal                           | no                                |
| <i>0, -Normal</i>  | +Normal   | no  | +Normal                           | no                                |
| <i>0, +Infinity</i>  | -Infinity   | no  | -Infinity                         | no                                |
| <i>0, -Infinity</i>  | +Infinity   | no  | +Infinity                         | no                                |
| <i>Normal, +Infinity</i>   | -Infinity   | set ufc,<br>set nvc,<br>set ufa,<br>set nva | no                                | set ufc,<br>set nvc,<br>ieee trap |
| <i>Normal, -Infinity</i>   | +Infinity   | set ufc,<br>set nvc,<br>set ufa,<br>set nva | no                                | set ofc,<br>set nvc,<br>ieee trap |
| <i>+Normal, -Normal</i>  | May overflow, see 17.5.3  |   | May overflow, see 17.5.3          |                                   |
| <i>Normal, +Normal</i>   | Normal  | no  | Normal                            | no                                |
| <i>-Normal, -Normal</i>  | May underflow, see 17.5.4   |   | May underflow, see 17.5.4         |                                   |
| <i>+Infinity, [0, Normal]</i>  | +Infinity   | no  | +Infinity                         | no                                |
| <i>-Infinity, [0, Normal]</i>  | -Infinity   | no  | -Infinity                         | no                                |
| <i>+Infinity, +Infinity</i>  | QNaN  | set nvc,<br>set nva                         | no                                | set nvc,<br>ieee trap             |
| <i>+Infinity, -Infinity</i>  | +Infinity   | no  | +Infinity                         | no                                |
| <i>-Infinity, +Infinity</i>  | -Infinity   | no  | -Infinity                         | no                                |
| <i>-Infinity, -Infinity</i>  | QNaN  | set nvc,<br>set nva                         | no                                | set nvc,<br>ieee trap             |

## 17.3.3 Multiplication

The floating-point multiplication instruction is shown in TABLE 17-5.

**TABLE 17-5** Floating-point Multiplication

| MULTIPLICATION Instruction<br><br>FMUL $rs_1, rs_2 [rs_2, rs_1] \rightarrow rd$ | RESULT from the operation includes one or more of the following: |                     |                                      |                       |
|---|--|---------------------|--------------------------------------|-----------------------|
|   | Masked Exception, TEM=0  |                     | Enabled Exception, TEM=1             |                       |
|   | Destination register written (rd)                                | Flag(s)             | Destination register written (rd)    | Flag(s), Trap         |
| +0, [+0/+Normal]  | +0   | no                  | +0                                   | no                    |
| +0, [-0/-Normal]  | -0   | no                  | -0                                   | no                    |
| -0, [+0/+Normal]  | -0   | no                  | -0                                   | no                    |
| -0, [-0/-Normal]  | +0   | no                  | +0                                   | no                    |
| +0, +Infinity   | QNaN   | set nvc,<br>set nva | no                                   | set nvc,<br>ieee trap |
| +0, -Infinity   | QNaN   | set nvc,<br>set nva | no                                   | set nvc,<br>ieee trap |
| -0, +Infinity   | QNaN   | set nvc,<br>set nva | no                                   | set nvc,<br>ieee trap |
| -0, -Infinity   | QNaN   | set nvc,<br>set nva | no                                   | set nvc,<br>ieee trap |
| Normal, Normal  | May underflow/<br>overflow, see 17.5                             |                     | May underflow/<br>overflow, see 17.5 |                       |
| [+Normal/+Infinity], +Infinity  | +Infinity  | no                  | +Infinity                            | no                    |
| [+Normal/+Infinity], -Infinity  | -Infinity  | no                  | -Infinity                            | no                    |
| [-Normal/-Infinity], +Infinity  | -Infinity  | no                  | -Infinity                            | no                    |
| [-Normal/-Infinity], -Infinity  | +Infinity  | no                  | +Infinity                            | no                    |

## 17.3.4 Division

The floating-point division instruction is shown in TABLE 17-6.

**TABLE 17-6** Floating-point Division

| <b>DIVISION Instruction</b><br><br>$rs_1 \ rs_2$<br><br><b>FDIV</b> $rs_1, rs_2 \rightarrow rd$ | <b>RESULT from the operation includes one or more of the following:</b> <ul style="list-style-type: none"> <li>• Number in f register, see <i>Trap Event</i> note, page 422.</li> <li>• Exception bit set, see TABLE 17-12.</li> <li>• Trap occurs, see abbreviations in TABLE 17-12.</li> <li>• Underflow/Overflow may occur.</li> </ul> |                     |  |                                   |
|---|---|---------------------|--|-----------------------------------|
|   | <b>Masked Exception, TEM=0</b>  |                     | <b>Enabled Exception, TEM=1</b>          |                                   |
|   | <b>Destination register written (rd)</b>  | <b>Flag(s)</b>      | <b>Destination register written (rd)</b> | <b>Flag(s), Trap</b>              |
| $0, 0$  | sign=0, expo=111...111, frac=111...111 (QNaN)   | set nvc,<br>set nva | no                                       | set nvc,<br>ieee trap             |
| $0, Normal$   | 0   | no                  | 0  | no                                |
| $0, Infinity$   | 0   | no                  | 0  | no                                |
| $+Normal, +0$   | +Infinity   | set nvc,<br>set nva | no                                       | set dzc,<br>set nvc,<br>ieee trap |
| $+Normal, -0$   | -Infinity   | set nvc,<br>set nva | no                                       | set dzc,<br>set nvc,<br>ieee trap |
| $-Normal, +0$   | -Infinity   | set nvc,<br>set nva | no                                       | set dzc,<br>set nvc,<br>ieee trap |
| $-Normal, -0$   | +Infinity   | set nvc,<br>set nva | no                                       | set dzc,<br>set nvc,<br>ieee trap |
| $Normal, Normal$  | May underflow/overflow,<br>see 17.5   |                     | May underflow/overflow,<br>see 17.5      |                                   |
| $Infinity, Infinity$  | QNaN  | set nvc,<br>set nva | no                                       | set nvc,<br>ieee trap             |
| $+Infinity, +Normal$  | +Infinity   | no                  | +Infinity                                | no                                |
| $+Infinity, -Normal$  | -Infinity   | no                  | -Infinity                                | no                                |
| $-Infinity, +Normal$  | -Infinity   | no                  | -Infinity                                | no                                |
| $-Infinity, -Normal$  | +Infinity   | no                  | +Infinity                                | no                                |

## 17.3.5 Square Root

The floating-point square root instruction is shown in TABLE 17-7.

**TABLE 17-7** Floating-point Square Root

| <b>SQUARE ROOT Instruction</b><br><br><i>sq root of rs<sub>2</sub></i><br><br><b>FSQRT</b> <i>rs<sub>2</sub> → rd</i> | <b>RESULT from the operation includes one or more of the following:</b> <ul style="list-style-type: none"> <li>• Number in f register, see <i>Trap Event</i> note, page 422.</li> <li>• Exception bit set, see TABLE 17-12.</li> <li>• Trap occurs, see abbreviations in TABLE 17-12.</li> <li>• Underflow/Overflow may occur.</li> </ul> |                     |  |                       |
|---|---|---------------------|--|-----------------------|
|   | <b>Masked Exception, TEM=0</b>  |                     | <b>Enabled Exception, TEM=1</b>          |                       |
|   | <b>Destination register written (rd)</b>  | <b>Flag(s)</b>      | <b>Destination register written (rd)</b> | <b>Flag(s), Trap</b>  |
| +0  | +0  | no                  | +0                                       | no                    |
| -0  | -0  | set nvc,<br>set nva | no                                       | set nvc,<br>ieec trap |
| +Normal   | May underflow/overflow,<br>see 17.5   |                     | May underflow/overflow,<br>see 17.5      |                       |
| [-Normal/-Infinity]   | QNaN<br>(sign=0, expo=111...111,<br>frac=111...111)   | set nvc,<br>set nva | no                                       | set nvc,<br>ieec trap |
| +Infinity   | + Infinity  | no                  | + Infinity                               | no                    |

## 17.3.6 Compare

Two f registers are compared. The result of the compare is reflected in the fccN bits of the FSR registers, shown in TABLE 17-8.

The FCMPE version of the instruction relates to Subnormal operations, see TABLE 17-16, *Results from NaN Operands*, on page 427.

**TABLE 17-8** Number Compare

| <b>FP NUMBER COMPARE Instruction</b><br><br><b>FCMP{E}</b> <i>rs<sub>1</sub>, rs<sub>2</sub></i> | <b>RESULT from the operation includes one or more of the following:</b> <ul style="list-style-type: none"> <li>• Exception bit set, see TABLE 17-12.</li> <li>• Trap occurs, see abbreviations in TABLE 17-12.</li> <li>• The fcc bit set.</li> </ul> |                |  |                      |
|--|---|----------------|--|----------------------|
|  | <b>Masked Exception, TEM=0</b>  |                | <b>Enabled Exception, TEM=1</b>            |                      |
|  | <b>FP Condition Code Setting (fccM)</b>   | <b>Flag(s)</b> | <b>FP Condition Code Setting (fccM)</b>    | <b>Flag(s), Trap</b> |
| +0, +0   | fcc=0 (rs <sub>1</sub> = rs <sub>2</sub> )  | no             | fcc=0 (rs <sub>1</sub> = rs <sub>2</sub> ) | no                   |
| -0, -0   | fcc=0 (rs <sub>1</sub> = rs <sub>2</sub> )  | no             | fcc=0 (rs <sub>1</sub> = rs <sub>2</sub> ) | no                   |
| +0, [+Normal/+Infinity]  | fcc=1 (rs <sub>1</sub> < rs <sub>2</sub> )  | no             | fcc=1 (rs <sub>1</sub> < rs <sub>2</sub> ) | no                   |
| -0, [-Normal/-Infinity]  | fcc=0 (rs <sub>1</sub> = rs <sub>2</sub> )  | no             | fcc=0 (rs <sub>1</sub> = rs <sub>2</sub> ) | no                   |
| -0, [+0/+Normal/+Infinity]   | fcc=1 (rs <sub>1</sub> < rs <sub>2</sub> )  | no             | fcc=1 (rs <sub>1</sub> < rs <sub>2</sub> ) | no                   |

TABLE 17-8 Number Compare (Continued)

| FP NUMBER<br>COMPARE Instruction | RESULT from the operation includes one or more of the following:  |         |                                  |               |
|----------------------------------|---|---------|----------------------------------|---------------|
|                                  | <ul style="list-style-type: none"> <li>Exception bit set, see TABLE 17-12.</li> <li>Trap occurs, see abbreviations in TABLE 17-12.</li> <li>The fcc bit set.</li> </ul> |         |                                  |               |
|                                  | Masked Exception, TEM=0   |         | Enabled Exception, TEM=1         |               |
| FCMP{E} $rs_1, rs_2$             | FP Condition Code Setting (fccM)  | Flag(s) | FP Condition Code Setting (fccM) | Flag(s), Trap |
| +0, [-0]/-Normal/-Infinity]      | fcc=2 ( $rs_1 > rs_2$ )   | no      | fcc=2 ( $rs_1 > rs_2$ )          | no            |
| Normal, Normal                   | =, >, or <  | no      | =, >, or <                       | no            |

## 17.3.7 Precision Conversion

Details of the precision conversion operations are shown in TABLE 17-9.

TABLE 17-9 Precision Conversion

| PRECISION CONVERSION<br>Operations                         | RESULT from the operation includes one or more of the following:  |         |                                       |               |
|--|---|---------|---------------------------------------|---------------|
|  | <ul style="list-style-type: none"> <li>Number in <math>\mathbb{F}</math> register, see <i>Trap Event</i> note, page 422.</li> <li>Exception bit set, see TABLE 17-12.</li> <li>Trap occurs, see abbreviations in TABLE 17-12.</li> <li>Underflow/Overflow may occur.</li> </ul> |         |                                       |               |
|  | Masked Exception, TEM=0   |         | Enabled Exception, TEM=1              |               |
| single operand   | Destination register written (rd)   | Flag(s) | Destination register written (rd)     | Flag(s), Trap |
| FsTOd $rs_2 \rightarrow rd$<br>FdTOs $rs_2 \rightarrow rd$ |   |         |                                       |               |
| FsTOd $\pm 0$<br>FdTOs $\pm 0$                             | $\pm 0$   | no      | $\pm 0$                               | no            |
| FsTOd $\pm Normal$   | $\pm Normal$  | no      | $\pm Normal$                          | no            |
| FdTOs $\pm Normal$   | May underflow/<br>overflow, see 17.4.   |         | May underflow/<br>overflow, see 17.4. |               |
| FsTOd $\pm Infinity$<br>FdTOs $\pm Infinity$               | $\pm Infinity$  | no      | $\pm Infinity$                        | no            |

Examples:

- FsTOd (7FD1.0000) = 7FFA.2000.0000.0000
- FsTOd (FDD1.0000) = FFFA.2000.0000.0000
- FdTOs (7FFA.2000.0000.0000) = 7FD1.0000
- FdTOs (FFFA.2000.0000.0000) = FFD1.0000

## 17.3.8 Floating-point to Integer Number Conversion

The floating-point to integer number conversion instruction is shown in TABLE 17-10.

**TABLE 17-10** Floating-point to Integer Number Conversion

| FP to Int NUMBER CONVERSION<br>Instruction |                                  | RESULT from the operation includes one or more of the following: |                     |  |                       |
|--|----------------------------------|--|---------------------|--|-----------------------|
|  |                                  | Masked Exception, TEM.NVM=0                                      |                     | Enabled Exception, TEM.NVM=1                   |                       |
| single operand                             |                                  | Destination register<br>written (rd)                             | Flag(s)             | Destination register<br>written (rd)           | Flag(s), Trap         |
| SP/DP<br>Int                               | +0                               | 000...000  | no                  | 000...000                                      | no                    |
|  | -0                               | 111...111  | no                  | 111...111                                      | no                    |
|  | +Infinity                        | 011...111  | no                  | no   | set nvc,<br>ieee trap |
|  | -Infinity                        | 100...000  | no                  | no   | set nvc,<br>ieee trap |
| SP Int                                     | +Normal < 2 <sup>31</sup>        | Integer representation of<br>the Normal number                   | no                  | Integer representation of<br>the Normal number | no                    |
|  | +Normal ≥ 2 <sup>31</sup>        | 011...111  | set nvc,<br>set nva | no   | set nvc,<br>ieee trap |
|  | -Normal > -[2 <sup>31</sup> + 1] | Integer representation of<br>the Normal number                   | no                  | Integer representation of<br>the Normal number | no                    |
|  | -Normal ≤ -[2 <sup>31</sup> + 1] | 100...000  | set nvc,<br>set nva | no   | set nvc,<br>ieee trap |
| DP Int                                     | +Normal < 2 <sup>63</sup>        | Integer representation of<br>the Normal number                   | no                  | Integer representation of<br>the Normal number | no                    |
|  | +Normal ≥ 2 <sup>63</sup>        | 011...111  | set nvc,<br>set nva | no   | set nvc,<br>ieee trap |
|  | -Normal > -[2 <sup>63</sup> + 1] | Integer representation of<br>the Normal number                   | no                  | Integer representation of<br>the Normal number | no                    |
|  | -Normal ≤ -[2 <sup>63</sup> + 1] | 100...000  | no                  | 100...000                                      | no                    |

## 17.3.9 Integer to Floating-point Number Conversion

The integer to floating-point number conversion instruction is shown in TABLE 17-11.

TABLE 17-11 Integer to Floating-point Number Conversion

| Int to FP NUMBER CONVERSION Instruction  |                               | RESULT from the operation includes one or more of the following:  |                  |                                   |                    |
|--|-------------------------------|---|------------------|-----------------------------------|--------------------|
| single operand   |                               | <ul style="list-style-type: none"> <li>Number in f register, see <i>Trap Event</i> note, page 422.</li> <li>Exception bit set, see TABLE 17-12.</li> <li>Trap occurs, see abbreviations in TABLE 17-12.</li> <li>Underflow/Overflow may occur.</li> </ul> |                  |                                   |                    |
| FiTOs $rs_2 \rightarrow rd$<br>FiTOd $rs_2 \rightarrow rd$<br>FxTOs $rs_2 \rightarrow rd$<br>FxTOd $rs_2 \rightarrow rd$ |                               | Masked Exception, TEM.NXM=0   | Flag(s)          | Enabled Exception, TEM.NXM=1      |                    |
|  |                               | Destination register written (rd)   | Flag(s)          | Destination register written (rd) | Flag(s), Trap      |
| SP/DP FP   | 0                             | $\pm 0$   | no               | $\pm 0$                           | no                 |
| SP FP  | $+Integer < 2^{23}$           | FP representation of +Normal  | no               | FP representation of +Normal      | no                 |
|  | $+Integer \geq 2^{23}$        | Integer is rounded to 23 msb and converted.   | set nvc, set nxc | no                                | set nvc, ieee trap |
|  | $-Integer > -[2^{23} + 1]$    | FP representation of +Normal  | no               | FP representation of +Normal      | no                 |
|  | $-Integer \leq -[2^{23} + 1]$ | Integer is rounded to 23 msb and converted.   | set nvc, set nxc | no                                | set nvc, ieee trap |
| DP FP  | $+Integer < 2^{52}$           | FP representation of +Normal  | no               | FP representation of +Normal      | no                 |
|  | $+Integer \geq 2^{52}$        | Integer is rounded to 52 msb and converted.   | set nvc, set nxc | no                                | set nvc, ieee trap |
|  | $-Integer > -[2^{52} + 1]$    | FP representation of +Normal  | no               | FP representation of +Normal      | no                 |
|  | $-Integer \leq -[2^{52} + 1]$ | Integer is rounded to 52 msb and converted.   | set nvc, set nxc | no                                | set nvc, ieee trap |

### 17.3.10 Copy/Move Operations

Floating-point numbers are not modified by the copy and move instructions: FMOV, FABS, and FNEG. The copy/move instructions will not generate an *unfinished\_FPop* or *unimplemented\_FPop* exception, but they will generate the *fp\_disabled* exception if the FPU is disabled.

The processor performs the appropriate sign bit transformation but will not cause an invalid exception and will not perform a QNaN to SNaN transformation.

These are single operand instructions that use the  $rs_2$  register as the source operand.

### *FMOV*

- *f* register-to- *f* register move.
  - No change to any bit, regardless of register content.
  - Useful with VIS instructions.

### *FABS*

- Changes the FP/Int sign bit to positive, if needed.
  - No change to any other bit, regardless of register content.

### *FNEG*

- Changes the FP/Int sign bit (If 0, then 1. If 1, then 0.)
- No change to any other bit, regardless of register content.

## 17.3.11 *f* Register Load/Store Operations

A load single floating-point (LDF) instruction writes to a 32-bit register. This must be converted to a 64-bit value (FsTOd) for use with double precision instructions.

A load double floating-point (LDDF) instruction writes to a pair of adjacent, 32-bit *f* registers aligned to an even boundary, and it can write to a 64-bit register. This must be converted to a 32-bit value (FdTOs) for use with single precision instructions.

Two LDF instructions can be used to load a 64-bit value when the memory address alignment to 64-bits is not guaranteed. Similarly, two STF instructions can be used to store a 64-bit value when the memory address alignment to 64-bits is not guaranteed.

## 17.3.12 VIS Operations

VIS instructions are unaffected by floating-point models. However, the FPU must be enabled. VIS instructions do not generate interrupts unless the FPU is disabled.

---

## 17.4 Traps and Exceptions

There are 3 trap vectors defined for floating-point operations:

- *fp\_disabled*
- *fp\_exception\_ieee\_754* (see section 17.5, *IEEE Traps*, on page 422)



- *fp\_exception\_other*

### *fp\_disabled* Trap

The floating-point unit can be enabled and disabled.

### *fp\_exception\_other* Trap

The *fp\_exception\_other* trap occurs when a floating-point operation cannot be completed by the processor (*unfinished\_FPop*) or an operation is requested that is not implemented by the processor (*unimplemented\_FPop*).

## 17.4.1 Summary of Exceptions

The floating-point unit exceptions are shown in TABLE 17-12.

TABLE 17-12 Floating-point Unit Exceptions

| Description                                 | IEEE Flag | Trap Abbreviation | Fault Trap Type                            | Exception/Trap Vector                                |
|---|-----------|-------------------|--|--|
| FPU disabled                                | none      | disable trap      | none                                       | <i>fp_disabled</i><br>(020 <sub>16</sub> )           |
| FP operation <b>invalid</b> (IEEE)          | <b>nv</b> | <b>ieee trap</b>  | <i>IEEE_745_exception</i><br>(FSR.FTT = 1) | <i>fp_exception_ieee_754</i><br>(021 <sub>16</sub> ) |
| FP operation <b>overflow</b> (IEEE)         | <b>of</b> |                   |  |  |
| FP operation <b>underflow</b> (IEEE)        | <b>uf</b> |                   |  |  |
| FP operation <b>division by zero</b> (IEEE) | <b>dz</b> |                   |  |  |
| FP operation <b>inexact</b> (IEEE)          | <b>nx</b> |                   |  |  |

## 17.4.2 Trap Event

When a floating-point exception causes a trap, the trap is precise. The response to traps is described in TABLE 17-13.

TABLE 17-13 Response to Traps

| Exception Event →   | fp_disabled | fp_exception_other |                 | fp_exception_ieee_754        |
|---|-------------|--------------------|-----------------|------------------------------|
|   |             | unimplemented_FPop | unfinished_FPop |                              |
| Resulting Action ↓  |             |                    |                 |                              |
| Address of instruction that caused the trap is put in the PC and pushed onto the trap stack.                                | ✓           | ✓                  | ✓               | ✓                            |
| The destination f register (rd) is unchanged from its state prior to the execution of the instruction that caused the trap. | ✓           | ✓                  | ✓               | ✓                            |
| The floating-point condition codes (fccN) are unchanged.  | ✓           | ✓                  | ✓               | ✓                            |
| The FSR.aexc field is unchanged.  | ✓           | ✓                  | ✓               | ✓                            |
| The FSR.cexc field is unchanged.  | ✓           | ✓                  | ✓               | Appropriate bit is set to 1. |
| The FSR.ftt field is set to:  | nc          | 3                  | 2               | 1                            |

## 17.4.3 Trap Priority

The traps generated by floating-point exceptions (*fp\_disabled*, *fp\_exception\_ieee\_754*, and *fp\_exception\_other*) are prioritized.

---

## 17.5 IEEE Traps

The Underflow, Overflow, Inexact, Division-by-zero, and Invalid IEEE traps are supported in standard and nonstandard modes. They are listed in TABLE 17-12, *Floating-point Unit Exceptions*, on page 421 and operate according to the IEEE 754-1985 Standard.

### 17.5.1 IEEE Trap Enable Mask (TEM)

Individual IEEE traps (*nv*, *of*, *uf*, *dz*, and *nx*) are masked by the FSR.TEM bits.

When a trap is masked and an exception is detected, then the appropriate `FSR.cexc` bit(s) are set and the destination register is written with data shown in TABLE 17-3, TABLE 17-4, TABLE 17-5, TABLE 17-6, TABLE 17-7, TABLE 17-8, and TABLE 17-9.

## 17.5.2 IEEE Invalid (nv) Trap

The IEEE invalid exception (nv) is generated when the source operand is a NaN (signalling or quiet), or the result cannot fit in the integer format.

The nv trap for an invalid case can be masked using the FSR register.

## 17.5.3 IEEE Overflow (of) Trap

When an overflow occurs the inexact flag is also set.

If an overflow occurs *and* the IEEE Overflow (of) and Invalid (nv) traps are enabled (`FSR.TEM.NVM = 1`), then a *fp\_exception\_IEEE\_754* is generated. If the Overflow trap is masked and the operation is valid, then the destination register (rd) receives Infinity.

The Overflow Trap is caused when the result of an arithmetic operation exceeds the range supported by the floating-point or integer number precision. This can happen in many different cases as listed in the tables of this section.

## 17.5.4 IEEE Underflow (uf) Trap

When a Normal number underflows the inexact flag is also set. Underflow is detected before rounding.

The Underflow condition leads to a Subnormal result unless gross underflow is detected. In that case the result is 0 and the inexact flag is raised.

Underflow is discussed in detail in section 17.6, *Underflow Operation*, on page 424.

## 17.5.5 IEEE Divide-by-Zero (DZ) Trap

When a number is divided by zero, the Divide-by-zero flag is asserted and an *ieee\_exception* is generated, if enabled. The dz flag and trap can only be generated by the FDIV instruction.

## 17.5.6 IEEE Inexact (nx) Trap

When an inexact condition occurs, the processor sets the `FSR.aexc.nxa` and/or the `FSR.cexc.nxc` bits whenever the rounded result of an operation differs from the precise result. The floating-point/integer conversions that generate inexact exceptions are shown in TABLE 17-14.

The Inexact (nx) flag is asserted for most an overflow or underflow conditions.

The Inexact trap is caused when the ideal result cannot fit into the destination format:

- most square root operations
- some add, subtract, multiply, and divide operations
- some number and precision conversion operations

**TABLE 17-14** FP ↔ Integer Conversions that Generate Inexact Exceptions

| Instruction    | Conversion Description  | Unmasked Exception, TEM=0   | Masked Exception, TEM=1 |
|----------------|---|-----------------------------|-------------------------|
| FsTOi<br>FdTOi | FP to 32-bit integer when the source operand is not between $-(2^{31} - 1)$ and $2^{31}$ , then the result is inexact.                          | Integer number, nx          | nx ieee trap            |
| FsTOx<br>FdTOx | FP to 64-bit integer when the source operand is not between $-(2^{63} - 1)$ and $2^{63}$ , then the result is inexact.                          | Integer number, nx          | nx ieee trap            |
| FiTOs          | Integer to FP when the 32-bit integer source operand magnitude is not exactly representable in single precision (23-bit fraction). <sup>1</sup> | Single Precision Normal, nx | nx ieee trap            |
| FxTOs          | Integer to FP when the 64-bit integer source operand magnitude is not exactly representable in single precision (23-bit fraction). <sup>1</sup> | Single Precision Normal, nx | nx ieee trap            |
| FxTOd          | Integer to FP when the 64-bit integer source operand magnitude is not exactly representable in double precision (52-bit fraction). <sup>2</sup> | Double Precision Normal, nx | nx ieee trap            |

1. Even if the operand is  $> 2^{24} - 1$ , if enough of its trailing bits are zeros, it may still be exactly representable.

2. Even if the operand is  $> 2^{53} - 1$ , if enough of its trailing bits are zeros, it may still be exactly representable.

## 17.6 Underflow Operation

Underflow occurs when the result of an operation (before rounding) is less than that representable by a Normal number.

After rounding, the tiny number (underflow) is usually represented by a Subnormal number, but may equal the smallest Normal number if the unrounded result is just below the range of Normal numbers and the rounding mode (specified in `FSR.RD`) moves it into the Normal number range. The underflow result will be zero, Subnormal, or the smallest Normal value.

---

**Compatibility Note** – The FPU does not support exponent wrapping for underflow or overflow.

---

## 17.6.1 Trapped Underflow

The FPU will trap on underflow if the `FSR.TEM.UFM` bit is set to 1. Since tininess is detected before rounding, trapped underflow occurs when the exact unrounded result has a magnitude between zero and the smallest representable Normal number in the precision of the destination format.

When underflow is trapped, the destination and other registers are left unchanged, see section 17.4.2, *Trap Event*, on page 422.

## 17.6.2 Untrapped Underflow

The FPU will not generate an underflow trap when an underflow occurs, if the `FSR.TEM.UFM` bit is set to 0.

If the result causes an underflow and the result after rounding is exact, then the FPU will not generate an inexact trap.

Tininess detection before rounding is summarized in TABLE 17-15.

Define a few terms:

- ***u*** is the unrounded (exact) value of the result.
- ***r*** is the rounded value of ***u*** (occurs when there is no trap generated)
- Underflow is when:  $0 < |u| < \textit{smallest Normal number}$ .

**TABLE 17-15** Underflow Exception Summary

| <b>Underflow :</b>                         |                                   | <b>enabled</b> (UFM = 1)    | <b>masked</b> (UFM = 0)  | <b>masked</b> (UFM = 0) |
|--|-----------------------------------|-----------------------------|--------------------------|-------------------------|
| <b>Inexact :</b>                           |                                   | <b>don't care</b> (NXM = x) | <b>enabled</b> (NXM = 1) | <b>masked</b> (NXM = 0) |
| <b><i>u = r</i></b><br><br>exact<br>result | <b><i>r</i></b> is minimum Normal | none                        | none                     | none                    |
|  | <b><i>r</i></b> is Subnormal      | set ufc, ieee trap          | none                     | none                    |
|  | <b><i>r</i></b> is Zero           | none                        | none                     | none                    |

**TABLE 17-15** Underflow Exception Summary (*Continued*)

| <b>Underflow :</b>   |                            | <b>enabled</b> (UFM = 1)    | <b>masked</b> (UFM = 0)  | <b>masked</b> (UFM = 0) |
|--|----------------------------|-----------------------------|--------------------------|-------------------------|
| <b>Inexact :</b>   |                            | <b>don't care</b> (NXM = x) | <b>enabled</b> (NXM = 1) | <b>masked</b> (NXM = 0) |
| <b>u r</b><br><br>inexact<br>result  | <b>r</b> is minimum Normal | set ufc, ieee trap          | set nxc, ieee trap       | set ufc, set ufa        |
|  | <b>r</b> is Subnormal      | set ufc, ieee trap          | set nxc, ieee trap       | set ufc, set ufa        |
|  | <b>r</b> is Zero           | set ufc, ieee trap          | set nxc, ieee trap       | set ufc, set ufa        |
| <p><b>set nxc</b> means FSR.cexc.nxc set to 1<br/> <b>set ufc</b> means FSR.cexc.ufc set to 1<br/> <b>set ufa</b> means FSR.aexc.ufa set to 1<br/> <b>ieee trap</b> means <i>fp_exception_ieee_754</i></p> |                            |                             |                          |                         |

## 17.7 IEEE NaN Operations

When a NaN operand appears or a NaN result is generated, and the invalid (*nv*) trap is enabled (FSR.TEM.NVM = 1), then the *fp\_exception\_ieee\_754* occurs.

If the invalid (*nv*) trap is masked (FSR.TEM.NVM = 0), then a signalling NaN operand is transformed into a quiet NaN. A quiet NaN operand will propagate to the destination register. Subnormals operations are described in TABLE 17-16, *Results from NaN Operands*, on page 427.

Whenever a NaN is created from non NaN operands, the *nv* flag is set.

### 17.7.1 Signaling and Quiet NaNs

SNaN and QNaN numbers are unsigned, the sign bit is an extension of the NaN's fraction field.

SNaN operands propagate to the destination register as a QNaN result when the *nv* exception is masked. All operations with NaN operands keep the sign bit unchanged including a FSQRT operation.

NaNs are generated for the conditions shown in section 17.7.4, *NaN Results from Operands without NaNs*, on page 428.

### 17.7.2 SNaN to QNaN Transformation

The signalling to quiet NaN transformation causes:

- The most significant bits of the operand fraction are copied to the most significant bits of the result's fraction. In conversion to a narrower format, excess low-order bits of the operand fraction are discarded. In conversion to a wider format, unwritten low-order bits of the result fraction are set to 0.
- The quiet bit (the most significant bit of the result fraction) is set to 1 (the NaN transformation produces a QNaN).
- The sign bit is copied from the operand to the result without modification.

## 17.7.3 Operations with NaN Operands

Operations with NaN operands may assert the IEEE invalid trap flag (*nv*). These operations are listed in TABLE 17-16.

If the Invalid Trap is enabled ( $\text{FSR.TEM.NVM} = 1$ ), then a trap event occurs as described in section 17.4.2, *Trap Event*, on page 422.

TABLE 17-16 Results from NaN Operands

| Operation   |  | RESULT from the operation includes one or more of the following: |                     |                                |                       |
|---|--|--|---------------------|--------------------------------|-----------------------|
|   |  | Masked Exception, TEM.NVM=0                                      |                     | Enabled Exception, TEM.NVM=1   |                       |
|   |  | rd or fcc register written                                       | flag set            | rd or fcc register written     | flag set              |
| <b>One Operand</b> $rs_2 \rightarrow rd$                    |  |  |                     |                                |                       |
| Any   | <i>QNaN</i>                                | QNaN,<br>see note <sup>1</sup>                                   | no                  | QNaN,<br>see note <sup>1</sup> | no                    |
| Any   | <i>SNaN</i>                                | SNaN $\rightarrow$ QNaN,<br>see note <sup>1</sup>                | set nvc,<br>set nva | no                             | set nvc,<br>ieee trap |
| <b>Two Operand</b> $rs_1, rs_2 [rs_2, rs_1] \rightarrow rd$ |  |  |                     |                                |                       |
| FADD, FSUB,<br>FMUL, FDIV                                   | <i>QNaN, QNaN</i>                          | QNaN <sub>rs2</sub>  | no                  | QNaN <sub>rs2</sub>            | no                    |
|   | <i>QNaN, anything except SNaN and QNaN</i> | QNaN   | no                  | QNaN                           | no                    |
|   | <i>SNaN, SNaN</i>                          | SNaN <sub>rs2</sub> $\rightarrow$ QNaN,<br>see note <sup>1</sup> | set nvc,<br>set nva | no                             | set nvc,<br>ieee trap |
|   | <i>SNaN, anything except SNaN</i>          | SNaN $\rightarrow$ QNaN,<br>see note <sup>1</sup>                | set nvc,<br>set nva | no                             | set nvc,<br>ieee trap |
| FCMPES,d  | <i>[SNaN or QNaN], anything</i>            | fcc=3 ( <i>unordered</i> )                                       | set nvc,<br>set nva | no                             | set nvc,<br>ieee trap |

**TABLE 17-16** Results from NaN Operands (Continued)

| Operation |                                   | RESULT from the operation includes one or more of the following: |                     |                                 |                       |
|-----------|-----------------------------------|--|---------------------|---------------------------------|-----------------------|
|           |                                   | Masked Exception,<br>TEM.NVM=0                                   |                     | Enabled Exception,<br>TEM.NVM=1 |                       |
|           |                                   | rd or fcc register written                                       | flag set            | rd or fcc register written      | flag set              |
| FCMPs,d   | <i>SNaN, anything</i>             | fcc=3 (unordered)  | set nvc,<br>set nva | no                              | set nvc,<br>ieee trap |
| FCMPs,d   | <i>QNaN, anything except SNaN</i> | fcc=3 (unordered)  | no                  | fcc=3<br>(unordered)            | no                    |

1. For the Fs,dTOs,d and other instructions, see section 17.7.2, *SNaN to QNaN Transformation*, on page 426.

---

**Note** – Notice from TABLE 17-16 that the *compare and cause exception if unordered* instruction (FCMPES,d) will cause an invalid (nv) exception if either operand is a quiet or signalling NaN. The FCMP instruction causes an exception for signalling NaNs only.

---

## 17.7.4 NaN Results from Operands without NaNs

The following operations generate NaNs, see section 17.3, *IEEE Operations*, on page 411, for details.

- FSQRT [*-Normal*, or *-0*]
- FDIV  $\pm 0$

---

## 17.8 Subnormal Operations

The handling of Subnormals is different for standard and nonstandard floating-point modes. The handling of operands and results are described separately in the following sections.

### 17.8.1 Response to Subnormal Operands

The FPU responds to Subnormal operands and results in either hardware or by generating an *fp\_exception\_other* (with FSR.fcc = 2, *unfinished\_FPop*).



The response of the FPU depends on the operating mode of the floating-point unit. This is controlled by the `FSR.NS` bit.

### *Standard Mode*

In Standard mode, the FPU generally traps when a Subnormal operand is detected or a Subnormal result is generated. In this situation, the system software must perform or complete the operation.

The FPU supports the following in Standard mode:

- Some cases of Subnormal operands are handled in hardware.
- Gross underflow results are supported in hardware for `FdTOs`, `FMULs,d`, and `FDIVs,d` instructions.

### *Nonstandard Mode*

In Nonstandard mode the FPU generally flushes Subnormal operands to 0 (with the same sign as the `SbN` number) and proceeds to use the value in the operation. Subnormal results (those that would otherwise cause an *unfinished\_FPop*) are also flushed to 0 in Nonstandard mode.

If the higher priority invalid operation (`nv`) or divide-by-zero (`dz`) condition occurs, then the corresponding condition(s) are flagged in the `FSR.cexc` register field. If the trap is enabled (`FSR.TEM` register), then an *fp\_exception\_ieee\_754* trap occurs. If the trap is disabled, then the corresponding condition(s) are also flagged in the `FSR.aexc` register field.

If neither the invalid nor divide-by-zero conditions occur, then an inexact condition plus any other detected floating-point exception conditions are flagged in the `FSR.cexc` register field. If an IEEE trap is enabled (`FSR.TEM` register), then an *fp\_exception\_ieee\_754* trap occurs. If the trap is disabled, then the corresponding condition(s) are also flagged in the `FSR.aexc` register field.

## 17.8.2 Subnormal Number Generation

Handling of the `FMULs`, `FMULd`, `FDIVs`, `FDIVd`, and `FdTOs` instructions requires further explanation.

Define:

- $\text{Sign}_r$  = sign of result,
- $\text{RT}_{\text{Eff}}$  = round nearest effective truncate or round truncate,
- $\text{RP}$  = round to +Infinity,
- $\text{RM}$  = round to -Infinity,
- $\text{RND} = \text{FSR.RD}$ ,

- $E_r$  = biased exponent result,
- $E_{rb}$  = the biased exponent result before rounding,
- $E(rs_1)$  = biased exponent of  $rs_1$  operand, and
- $P_{rs_1}$  = precision of the  $rs_1$  operand.

The value of the constants dependent on precision type, see TABLE 17-17.

**TABLE 17-17** Subnormal Handling Constants per Destination Register Precision

| Destination Register Precision (P) | Number of Bits in Exponent Field | Exponent Bias ( $E_{BIAS}$ ) | Exponent Max ( $E_{MAX}$ ) | Exponent Gross Underflow ( $E_{GUF}$ ) |
|------------------------------------|----------------------------------|------------------------------|----------------------------|--|
| Single                             | 8                                | 127                          | 255                        | -24                                    |
| Double                             | 11                               | 1023                         | 2047                       | -53                                    |

- For FMULs and FMULd:  $E_r = E(rs_1) + E(rs_2) - E_{BIAS}$ .
- For FDIVs and FDIVd:  $E_r = E(rs_1) - E(rs_2) + E_{BIAS} - 1$ .

When two Normal operands of FMULs,d and FDIVs,d generate a Subnormal result, the  $E_{rb}$  is calculated using the algorithm shown in code example 17-1.

**CODE EXAMPLE 17-1** Normal Operands Generating a Subnormal Result Pseudocode

```

If (fraction_msb overflows) // i.e., fraction_msb >= 1'd2
{
    Erb = Er + 1
}
ELSE
{
    Erb = Er
}

```

- For FdTOs,  $E_r = E(rs_2) - E_{BIAS}(P_{rs_2}) + E_{BIAS}(P_{rd})$ , where  $P_{rs_2}$  is the larger precision of the source and  $P_{rd}$  is the smaller precision of the destination.

Even though  $0 \leq [E(rs_1) \text{ or } E(rs_2)] \leq 255$  for each single precision biased operand exponent, the *computed* biased exponent result ( $E_r$ ) can be  $0 \leq E_r \leq 255$  or can even be negative. For example, for the FMULs instruction:

- If  $E(rs_1) = E(rs_2) = +127$ , then  $E_r = +127 (127 + 127 - 127)$
- If  $E(rs_1) = E(rs_2) = 0$ , then  $E_r = -127 (0 + 0 - 127)$

### *Overflow Result*

- If the appropriate trap enable masks are not set ( $FSR.OFM = 0$  and  $FSR.NXM = 0$ ), then set  $FSR.aexc$  and  $FSR.cexc$  overflow and inexact flags:  $FSR.ofa = 1$ ,  $FSR.nxa = 1$ ,  $FSR.ofc = 1$ , and  $FSR.nxc = 1$ . No trap is generated.
- If any or both of the appropriate trap enable masks are set ( $FSR.OFM = 1$  or  $FSR.NXM = 1$ ), then only an IEEE overflow trap is generated:  $FSR.ftt = 1$ . The particular  $FSR.cexc$  bit that is set follows the SPARC V9 architecture:
  - If  $FSR.OFM = 0$  and  $FSR.NXM = 1$ , then  $FSR.nxc = 1$ .
  - If  $FSR.OFM = 1$  (independent of  $FSR.NXM$ ), then  $FSR.ofc = 1$  and  $FSR.nxc = 0$ .

### *Gross Underflow Zero result*

- Result = 0 (with correct sign).
- If the appropriate trap enable masks are not set ( $FSR.UFM = 0$  and  $FSR.NXM = 0$ ), then set the  $FSR.aexc$  and  $FSR.cexc$  underflow and inexact flags:  $FSR.ufa = 1$ ,  $FSR.nxa = 1$ ,  $FSR.ufc = 1$ , and  $FSR.nxc = 1$ . A trap is not generated.
- If either or both of the appropriate trap enable masks are set ( $FSR.UFM = 1$  or  $FSR.NXM = 1$ ), then only an IEEE underflow trap is generated:  $FSR.ftt = 1$  and  $FSR.cexc.uf = 1$ . The particular  $FSR.cexc$  bit that is set diverges from previous UltraSPARC implementations to follow the SPARC V9 architecture:
  - If  $FSR.UFM = 0$  and  $FSR.NXM = 1$ , then  $FSR.nxc = 1$ .
  - If  $FSR.UFM = 1$ , independent of  $FSR.NXM$ , then  $FSR.ufc = 1$  and  $FSR.nxc = 0$ .

### *Subnormal Handling Override*

- Result is an QNaN or SNaN
  - Subnormal + SNaN = QNaN, invalid exception generated
    - Standard mode: No *unfinished\_FPop*
    - Nonstandard mode: No FSR.NX
  - Subnormal + QNaN = QNaN, no exception generated
    - Standard mode: No *unfinished\_FPop*
    - Nonstandard mode: No FSR.NX
- Result already generates an exception (Divide-by-zero or Invalid operation)
  - $FSQRT(\text{number less than zero}) = \text{invalid}$
- Result is Infinity:

- Subnormal + Infinity = Infinity, no exception generated
  - Standard mode: No *unfinished\_FPop*
  - Nonstandard mode: No FSR.nx
- Standard mode: Subnormal  $\times$  Infinity = Infinity
- Nonstandard mode: Subnormal  $\times$  Infinity = QNaN with nv exception (Subnormal is flushed to zero)
- Result is zero:
  - Subnormal  $\times$  0 = 0, no exception generated
    - Standard mode: No *unfinished\_FPop*
    - Nonstandard mode: No FSR.nx

## SECTION VII

---

---

# Special Topics

---

---



## Reset and RED\_state

---

This chapter examines RED\_state (Reset, Error, and Debug state) in the following sections:

- RED\_state Characteristics
- Resets
- RED\_state Trap Vector
- Machine States

---

### 18.1 RED\_state Characteristics

A reset or trap that sets PSTATE.RED (including a trap in RED\_state) will clear the DCU Control Register, including enable bits for I-cache, D-cache, I-MMU, D-MMU, and virtual and physical watchpoints. The characteristics of RED\_state include the following:

- The default access in RED\_state is non-cacheable; therefore, there must be non-cacheable scratch memory somewhere in the system.
- The D-cache, watchpoints, and D-MMU can be enabled by software in RED\_state, but any trap will disable them again.
- The I-MMU and consequently the I-cache are always disabled in RED\_state. Disabling overrides the enable bits in the DCU Control Register.

When PSTATE.RED is explicitly set by a software write, there are no side-effects other than the I-MMU is disabled. Software must create the appropriate state itself.

A trap when TL = MAXTL - 1 immediately brings the processor into RED\_state. In addition, a trap when TL = MAXTL immediately brings the processor into error\_state. Upon error\_state entry, the processor automatically recovers through watchdog reset (WDR) into RED\_state. A trap to error\_state immediately triggers WDR. A Signal

Monitor (SIGM) instruction generates a software-initiated reset (SIR) trap on the local processor. A trap to software-initiated reset causes an SIR trap on the processor and brings the processor into RED\_state.

During the RED\_state, the caches continue to snoop and maintain coherence if DVMA or other processors are still issuing cacheable accesses.

---

**Note** – A recommended way to exit RED\_state is with a DONE or RETRY. Exiting RED\_state by writing 0 to PSTATE.RED in the delay slot of a JMPL is not recommended and may result in an instruction access error.

---

---

## 18.2 Resets

Reset priorities from highest to lowest are Power-on Reset (POR), System Reset, externally initiated reset (XIR), watchdog reset (WDR), and software-initiated reset (SIR).

### 18.2.1 Power-on Reset (POR)

A Power-on Reset (POR) occurs when the POK pin is activated and stays asserted until the processor is within its specified operating range. When the POK pin is active, all other resets and traps are ignored. POR has a trap type of 1 at physical address offset  $20_{16}$ . Any pending external transactions are canceled.

After POR, software must initialize values of certain registers and state that is unknown. The following bits must be initialized before the caches are enabled:

- In the I-cache, valid bits must be cleared and Microtag (Utag) bits must be set so that each way within a set has a unique Utag value.
- In the D-cache, valid bits must be cleared and Utag bits must be set so that each way within a set has a unique Utag value.
- All L2-cache tags and data

The I-MMU and D-MMU TLBs must also be initialized. The P-cache valid bits must be initialized before any floating-point loads are executed.

The MCU Refresh Control Register, as well as the Fireplane Configuration Register, must be initialized after a POR.

In SSM systems, the Utags contained in memory must be initialized before any Fireplane transactions are generated.



---

**Caution** – Executing a `DONE` or `RETRY` instruction when `TSTATE` is not initialized after a POR can damage the chip. The POR boot code should initialize `TSTATE<3:0>`, using `wrpr` writes, before any `DONE` or `RETRY` instructions are executed.

However, these operations can only be executed in privileged mode. Therefore, user code is not at the risk of damaging the chip.

---

## 18.2.2 System Reset

A System Reset occurs when the Reset pin is activated. When the Reset pin is active, all other resets and traps are ignored. System Reset has a trap type of 1 at physical address offset  $20_{16}$ . Any pending external transactions are canceled.

---

**Note** – Memory refresh continues uninterrupted during a System Reset. System interface, L2-cache configuration, and memory controller configuration are preserved across a System Reset.

---

## 18.2.3 Externally Initiated Reset (XIR)

An externally initiated reset (XIR) is sent to the processor through an external hardware pin. It causes a SPARC V9 XIR, which has a trap type  $3_{16}$  at physical address offset  $60_{16}$ . XIR has higher priority than all other resets except POR and System Reset.

XIR affects only one processor rather than the entire system. Memory state, cache state, and most Control Status Register states are unchanged. System coherency is *not* guaranteed to be maintained through an XIR reset. The saved PC and nPC will only be approximate because the trap is not precise with respect to pipeline state.

## 18.2.4 Watchdog Reset (WDR) and `error_state`

The processor enters `error_state` when a trap occurs at `TL = MAXTL`.

The processor automatically exits `error_state` using WDR. The processor signals itself internally to take a WDR and sets `TT = 2`. The WDR traps to the address at `RSTVaddr + 0x4016`. WDR sets the processor in a state where it is prepared for diagnosis of failures.

WDR affects only one processor rather than the entire system. CWP updates due to window traps that cause watchdog traps are the same as the no watchdog trap case.

## 18.2.5 Software-Initiated Reset (SIR)

A software-initiated reset (SIR) is initiated by an SIR instruction within any processor. This per-processor reset has a trap type 4 at physical address offset  $80_{16}$ . SIR affects only one processor rather than the entire system.

## 18.3 RED\_state Trap Vector

When the UltraSPARC III Cu processor processes a reset or trap that enters RED\_state, it takes a trap at an offset relative to the RED\_state trap vector base address (RSTVaddr); the base address is at virtual address  $FFFF\ FFFF\ F000\ 0000_{16}$ , which passes through to physical address  $7FF\ F000\ 0000_{16}$ .

## 18.4 Machine States

TABLE 18-1 shows the machine states created as a result of any reset or after RED\_state is entered. RSTVaddr is often abbreviated as RSTV in the table.

**TABLE 18-1** Machine State After Reset and in RED\_state (1 of 6)

| Name                      | Fields | Hard_POR  | System Reset     | WDR              | XIR              | SIR              | RED_state <sup>†</sup> |
|---------------------------|--------|---|------------------|------------------|------------------|------------------|------------------------|
| Integer Registers         |        | Unknown   | Unchanged        | Unchanged        |                  |                  |                        |
| Floating-Point Registers  |        | Unknown   | Unchanged        | Unchanged        |                  |                  |                        |
| L2-cache Control Register |        | 0   | 0                | Unchanged        |                  |                  |                        |
| RSTVaddr value            |        | VA = $FFFF\ FFFF\ F000\ 0000_{16}$<br>PA = $7FF\ F000\ 0000_{16}$ |                  |                  |                  |                  |                        |
| PC                        |        | RSTV   $20_{16}$  | RSTV   $20_{16}$ | RSTV   $40_{16}$ | RSTV   $60_{16}$ | RSTV   $80_{16}$ | RSTV   $A0_{16}$       |
| nPC                       |        | RSTV   $24_{16}$  | RSTV   $24_{16}$ | RSTV   $44_{16}$ | RSTV   $64_{16}$ | RSTV   $84_{16}$ | RSTV   $A4_{16}$       |

**TABLE 18-1** Machine State After Reset and in RED\_state (2 of 6)

| Name       | Fields                       | Hard_POR                           | System Reset                       | WDR  | XIR                    | SIR | RED_state <sup>†</sup> |
|------------|------------------------------|------------------------------------|------------------------------------|--|------------------------|-----|------------------------|
| PSTATE     | MM                           | 0 (TSO)                            | 0 (TSO)                            | 0 (TSO)                                    |                        |     |                        |
|            | RED                          | 1<br>(RED_state)                   | 1<br>(RED_state)                   | 1 (RED_state)                              |                        |     |                        |
|            | PEF                          | 1 (FPU on)                         | 1 (FPU on)                         | 1 (FPU on)                                 |                        |     |                        |
|            | AM                           | 0 (Full 64-bit address)            | 0 (Full 64-bit address)            | 0 (Full 64-bit address)                    |                        |     |                        |
|            | PRIV                         | 1 (Privileged mode)                | 1 (Privileged mode)                | 1 (Privileged mode)                        |                        |     |                        |
|            | IE                           | 0 (Disable interrupts)             | 0 (Disable interrupts)             | 0 (Disable interrupts)                     |                        |     |                        |
|            | AG                           | 1 (Alternate globals selected)     | 1 (Alternate globals selected)     | 1 (Alternate globals selected)             |                        |     |                        |
|            | CLE                          | 0 (Current little-endian)          | 0 (Current little-endian)          | PSTATE.TLE                                 |                        |     |                        |
|            | TLE                          | 0 (Trap little-endian)             | 0 (Trap little-endian)             | Unchanged                                  |                        |     |                        |
|            | IG                           | 0 (Interrupt globals not selected) | 0 (Interrupt globals not selected) | 0 (Interrupt globals not selected)         |                        |     |                        |
| MG         | 0 (MMU globals not selected) | 0 (MMU globals not selected)       | 0 (MMU globals not selected)       |  |                        |     |                        |
| TBA<63:15> |                              | Unknown                            | Unchanged                          | Unchanged                                  |                        |     |                        |
| Y          |                              | Unknown                            | Unchanged                          | Unchanged                                  |                        |     |                        |
| PIL        |                              | Unknown                            | Unchanged                          | Unchanged                                  |                        |     |                        |
| CWP        |                              | Unknown                            | Unchanged                          | Unchanged except for register window traps |                        |     |                        |
| TT[TL]     |                              | 1                                  | 1                                  | Unchanged                                  | 3                      | 4   | Trap type              |
| CCR        |                              | Unknown                            | Unchanged                          | Unchanged                                  |                        |     |                        |
| ASI        |                              | Unknown                            | Unchanged                          | Unchanged                                  |                        |     |                        |
| TL         |                              | MAXTL                              | MAXTL                              | Min(TL + 1, MAXTL)                         |                        |     |                        |
| TPC[TL]    |                              | Unknown                            | Unchanged                          | PC   | PC & ~1F <sub>16</sub> |     | PC                     |
| TNPC[TL]   |                              | Unknown                            | Unchanged                          | nPC  | nPC = PC + 4           |     | nPC                    |

**TABLE 18-1** Machine State After Reset and in RED\_state (3 of 6)

| Name                     | Fields    | Hard_POR           | System Reset | WDR       | XIR          | SIR       | RED_state <sup>†</sup> |
|--------------------------|-----------|--------------------|--------------|-----------|--------------|-----------|------------------------|
| TSTATE                   | CCR       | Unknown            | Unchanged    | CCR       |              |           |                        |
|                          | ASI       | Unknown            | Unchanged    | ASI       |              |           |                        |
|                          | PSTATE    | Unknown            | Unchanged    | PSTATE    |              |           |                        |
|                          | CWP       | Unknown            | Unchanged    | CWP       |              |           |                        |
|                          | PC        | Unknown            | Unchanged    | PC        |              |           |                        |
|                          | nPC       | Unknown            | Unchanged    | nPC       |              |           |                        |
| TICK                     | NPT       | 1                  | 1            | Unchanged | Unchanged    | Unchanged |                        |
|                          | counter   | Restart at 0       | Restart at 0 | Count     | Restart at 0 | Count     |                        |
| CANSAVE                  |           | Unknown            | Unchanged    | Unchanged |              |           |                        |
| CANRESTORE               |           | Unknown            | Unchanged    | Unchanged |              |           |                        |
| OTHERWIN                 |           | Unknown            | Unchanged    | Unchanged |              |           |                        |
| CLEANWIN                 |           | Unknown            | Unchanged    | Unchanged |              |           |                        |
| WSTATE                   | OTHER     | Unknown            | Unchanged    | Unchanged |              |           |                        |
|                          | NORMAL    | Unknown            | Unchanged    | Unchanged |              |           |                        |
| VER                      | MANUF     | 003E <sub>16</sub> |              |           |              |           |                        |
|                          | IMPL      | 0015 <sub>16</sub> |              |           |              |           |                        |
|                          | MASK      | Mask dependent     |              |           |              |           |                        |
|                          | MAXTL     | 5                  |              |           |              |           |                        |
|                          | MAXWIN    | 7                  |              |           |              |           |                        |
| FSR                      | All       | 0                  | 0            | Unchanged |              |           |                        |
| FPRS                     | All       | Unknown            | Unchanged    | Unchanged |              |           |                        |
| <b>Non-SPARC V9 ASRs</b> |           |                    |              |           |              |           |                        |
| SOFTINT                  |           | Unknown            | Unchanged    | Unchanged |              |           |                        |
| TICK_COMPARE             | INT_DIS   | 1 (off)            | 1 (off)      | Unchanged |              |           |                        |
|                          | TICK_CMPR | 0                  | 0            | Unchanged |              |           |                        |
| STICK                    | NPT       | 1                  | 1            | Unchanged |              |           |                        |
|                          | counter   | 0                  | 0            | Count     |              |           |                        |
| STICK_COMPARE            | INT_DIS   | 1 (off)            | 1 (off)      | Unchanged |              |           |                        |
|                          | TICK_CMPR | 0                  | 0            | Unchanged |              |           |                        |

**TABLE 18-1** Machine State After Reset and in RED\_state (4 of 6)

| Name                     | Fields                   | Hard_POR | System Reset | WDR       | XIR | SIR | RED_state <sup>†</sup> |
|--------------------------|--------------------------|----------|--------------|-----------|-----|-----|------------------------|
| PCR                      | S1                       | Unknown  | Unchanged    | Unchanged |     |     |                        |
|                          | S0                       | Unknown  | Unchanged    | Unchanged |     |     |                        |
|                          | UT (trace user)          | Unknown  | Unchanged    | Unchanged |     |     |                        |
|                          | ST (trace system)        | Unknown  | Unchanged    | Unchanged |     |     |                        |
|                          | PRIV (privileged access) | Unknown  | Unchanged    | Unchanged |     |     |                        |
| PIC                      | All                      | Unknown  | Unknown      | Unknown   |     |     |                        |
| GSR                      | IM                       | 0        | 0            | Unchanged |     |     |                        |
|                          | others                   | Unknown  | Unchanged    | Unchanged |     |     |                        |
| DCR                      | MS                       | 0        | 0            | Unchanged |     |     |                        |
|                          | SI                       | 0        | 0            | Unchanged |     |     |                        |
|                          | RPE                      | 0        | 0            | Unchanged |     |     |                        |
|                          | BPE                      | 0        | 0            | Unchanged |     |     |                        |
|                          | OBS                      | 0        | 0            | Unchanged |     |     |                        |
|                          | IFPOE                    | 0        | 0            | Unchanged |     |     |                        |
|                          | IPE                      | 0        | 0            | Unchanged |     |     |                        |
|                          | DPE                      | 0        | 0            | Unchanged |     |     |                        |
| <b>Non-SPARC V9 ASIs</b> |                          |          |              |           |     |     |                        |
| Fireplane Information    |                          |          |              |           |     |     |                        |
| DCUCR                    | WE                       | 0 (off)  | 0 (off)      | 0         |     |     |                        |
|                          | All others               | 0 (off)  | 0 (off)      | 0 (off)   |     |     |                        |
| INSTRUCTION_TRAP         | All                      | 0 (off)  | 0 (off)      | Unchanged |     |     |                        |
| VA_WATCHPOINT            |                          | Unknown  | Unchanged    | Unchanged |     |     |                        |
| PA_WATCHPOINT            |                          | Unknown  | Unchanged    | Unchanged |     |     |                        |

**TABLE 18-1** Machine State After Reset and in RED\_state (5 of 6)

| Name   | Fields          | Hard_POR    | System Reset | WDR       | XIR | SIR | RED_state <sup>†</sup> |
|--|-----------------|-------------|--------------|-----------|-----|-----|------------------------|
| I-SFSR, D-SFSR   | ASI             | Unknown     | Unchanged    | Unchanged |     |     |                        |
|  | FT              | Unknown     | Unchanged    | Unchanged |     |     |                        |
|  | E               | Unknown     | Unchanged    | Unchanged |     |     |                        |
|  | CTXT            | Unknown     | Unchanged    | Unchanged |     |     |                        |
|  | PRIV            | Unknown     | Unchanged    | Unchanged |     |     |                        |
|  | W               | Unknown     | Unchanged    | Unchanged |     |     |                        |
|  | OW (overwrite)  | Unknown     | Unchanged    | Unchanged |     |     |                        |
|  | FV (SFSR valid) | 0           | 0            | Unchanged |     |     |                        |
|  | NF              | Unknown     | Unchanged    | Unchanged |     |     |                        |
| TM   | Unknown         | Unchanged   | Unchanged    |           |     |     |                        |
| DMMU_SFAR  |                 | Unknown     | Unchanged    | Unchanged |     |     |                        |
| INTR_DISPATCH  | All             | 0           | 0            | Unchanged |     |     |                        |
| INTR_RECEIVE   | BUSY            | 0           | 0            | Unchanged |     |     |                        |
|  | MID             | Unknown     | Unchanged    | Unchanged |     |     |                        |
| ESTATE_ERR_EN  | All             | 0 (all off) | 0 (all off)  | Unchanged |     |     |                        |
| AFAR   | PA              | Unknown     | Unchanged    | Unchanged |     |     |                        |
| AFSR   | All             | 0           | Unchanged    | Unchanged |     |     |                        |
| Rfr_CSR  | All             | Unknown     | Unchanged    | Unchanged |     |     |                        |
| Mem_Timing_CSR   | All             | Unknown     | Unchanged    | Unchanged |     |     |                        |
| Mem_Addr_Dec   | All             | Unknown     | Unchanged    | Unchanged |     |     |                        |
| Mem_Addr_Cnt1  | All             | Unknown     | Unchanged    | Unchanged |     |     |                        |
| <b>Other Processor-Specific States</b>   |                 |             |              |           |     |     |                        |
| Processor and L2-cache tags, Microtags and data (includes data, instruction, prefetch, and write caches) |                 | Unknown     | Unchanged    | Unchanged |     |     |                        |
| Cache snooping   |                 | Enabled     |              |           |     |     |                        |
| Instruction Queue  |                 | Empty       |              |           |     |     |                        |
| Store Queue  |                 | Empty       | Empty        | Unchanged |     |     |                        |

**TABLE 18-1** Machine State After Reset and in RED\_state (6 of 6)

| Name                      | Fields                                 | Hard_POR | System Reset        | WDR       | XIR | SIR | RED_state <sup>‡</sup> |
|---------------------------|--|----------|---------------------|-----------|-----|-----|------------------------|
| I-TLB, D-TLB <sup>1</sup> | Mappings in #2 (2-way set-associative) | Unknown  | Unchanged           | Unchanged |     |     |                        |
|                           | Mappings in #0 (fully set-associative) | Unknown  | Unknown and invalid | Unchanged |     |     |                        |
|                           | E (side-effect) bit                    | 1        | 1                   | 1         |     |     |                        |
|                           | NC (non-cacheable) bit                 | 1        | 1                   | 1         |     |     |                        |

1. The V, L, U, and G bits are cleared at reset.

<sup>‡</sup> Processor states are only updated according to the table if RED\_state is entered because of a reset or trap. If RED\_state is entered because the PSTATE.RED bit was explicitly set to 1, then software must create the appropriate states itself.





---

---

# Appendix

---

---



## Instruction Definitions

---

Related instructions are grouped into subsections. Each subsection consists of the following parts:

1. A table of the opcodes defined in the subsection with the values of the field(s) that uniquely identify the instruction(s).
2. An illustration of the applicable instruction format(s). In these illustrations a dash (—) indicates that the field is *reserved* for future versions of the architecture and shall be zero in any instance of the instruction. If the processor encounters nonzero values in these fields, its behavior is undefined.
3. A description of the features, restrictions, and exception-causing conditions.
4. A list of exceptions that can occur as a consequence of attempting to execute the instruction(s). Exceptions due to an *instruction\_access\_error*, *instruction\_access\_exception*, *fast\_instruction\_access\_MMU\_miss*, *fast\_ECC\_error*, *ECC\_error (corrected ECC\_error)*, *WDR*, and interrupts are not listed because they can occur on any instruction. Instructions not implemented in hardware shall generate an *illegal\_instruction* exception and therefore will not generate any of the other exceptions listed. The *illegal\_instruction* exception is not listed because it can occur on any instruction that triggers an instruction breakpoint or contains an invalid field.

Instruction Latencies and Execution rates are provided in Chapter 4 “Instruction Execution.”

TABLE A-2 summarizes the instruction set; the instruction definitions follow the table. Within TABLE A-2 and throughout this chapter, certain opcodes are marked with mnemonic superscripts. The superscripts and their meanings are defined in TABLE A-1.

**TABLE A-1** Opcode Superscripts

| Superscript      | Meaning  |
|------------------|--|
| D                | Deprecated instruction   |
| P                | Privileged opcode  |
| P <sub>ASI</sub> | Privileged action if bit 7 of the referenced ASI is zero       |
| P <sub>ASR</sub> | Privileged opcode if the referenced ASR register is privileged |
| P <sub>NPT</sub> | Privileged action if PSTATE.PRIV = 0 and (S)TICK.NPT = 1       |
| P <sub>PIC</sub> | Privileged action if PCR.PRIV = 1                              |

**TABLE A-2** Instruction Set (1 of 6)

| Operation                          | Name   | Page     | Ext. to V9? |
|------------------------------------|--|----------|-------------|
| ADD, ADD <sub>cc</sub>             | Add (and modify condition codes)   | page 454 |             |
| ADDC, ADDC <sub>cc</sub>           | Add with carry (and modify condition codes)                                | page 454 |             |
| ALIGNADDRESS{ <sub>LITTLE</sub> }  | Calculate address for misaligned data                                      | page 455 | 3           |
| AND, AND <sub>cc</sub>             | And (and modify condition codes)   | page 517 |             |
| ANDN, ANDN <sub>cc</sub>           | And not (and modify condition codes)                                       | page 517 |             |
| ARRAY(8,16,32)                     | Three-Dimensional array addressing instructions                            | page 457 | 3           |
| BP <sub>cc</sub>                   | Branch on integer condition codes with prediction                          | page 474 |             |
| Bi <sub>cc</sub> <sup>D</sup>      | Branch on integer condition codes  | page 605 |             |
| BMASK                              | Set the GSR.MASK field   | page 468 | 3           |
| BPr                                | Branch on contents of integer register with prediction (also known as BRr) | page 469 |             |
| BSHUFFLE                           | Permute bytes as specified by GSR.MASK                                     | page 468 | 3           |
| CALL                               | Call and link  | page 476 |             |
| CASA <sup>P<sub>ASI</sub></sup>    | Compare and swap word in alternate space                                   | page 477 |             |
| CASXA <sup>P<sub>ASI</sub></sup>   | Compare and swap doubleword in alternate space                             | page 477 |             |
| DONE <sup>P</sup>                  | Return from trap   | page 479 |             |
| EDGE(8,16,32){ <sub>L,N,LN</sub> } | Edge handling instructions   | page 480 | 3           |
| FABS(s,d,q)                        | Floating-point absolute value  | page 493 |             |
| FADD(s,d,q)                        | Floating-point add   | page 483 |             |
| FALIGNDATA                         | Perform data alignment for misaligned data                                 | page 455 | 3           |
| FAND{S}                            | Logical AND operation  | page 514 | 3           |
| FANDNOT(1,2){S}                    | Logical AND operation with one inverted source                             | page 514 | 3           |
| FBf <sub>cc</sub> <sup>D</sup>     | Branch on floating-point condition codes                                   | page 603 |             |
| FBP <sub>fcc</sub>                 | Branch on floating-point condition codes with prediction                   | page 471 |             |

**TABLE A-2** Instruction Set (2 of 6)

| Operation                | Name  | Page                               | Ext. to V9? |
|--------------------------|---|------------------------------------|-------------|
| FCMP(s,d,q)              | Floating-point compare  | page 486                           |             |
| FCMPE(s,d,q)             | Floating-point compare (exception if unordered)                             | page 486                           |             |
| FCMP(GT,LE,NE,EQ)(16,32) | Pixel compare operations  | page 550                           | 3           |
| FDIV(s,d,q)              | Floating-point divide   | page 494                           |             |
| FdMULq                   | Floating-point multiply double to quad                                      | page 494                           |             |
| FEXPAND                  | Pixel expansion   | page 558                           | 3           |
| FiTO(s,d,q)              | Convert integer to floating-point   | page 491                           |             |
| FLUSH                    | Flush instruction memory  | page 497                           |             |
| FLUSHW                   | Flush register windows  | page 499                           |             |
| FMOV(s,d,q)              | Floating-point move   | page 493                           |             |
| FMOV(s,d,q)cc            | Move floating-point register if condition is satisfied                      | page 524                           |             |
| FMOV(s,d,q)r             | Move floating-point register if integer register contents satisfy condition | page 529                           |             |
| FMUL(s,d,q)              | Floating-point multiply   | page 494                           |             |
| FMUL8x16                 | 8x16 partitioned product  | page 545                           | 3           |
| FMUL8x16(AU,AL)          | 8x16 upper/lower $\alpha$ partitioned product                               | page 545                           | 3           |
| FMUL8(SU,UL)x16          | 8x16 upper/lower partitioned product  | page 546                           | 3           |
| FMULD8(SU,UL)x16         | 8x16 upper/lower partitioned product  | page 548                           | 3           |
| FNAND{S}                 | Logical NAND operation  | page 514                           | 3           |
| FNEG(s,d,q)              | Floating-point negate   | page 493                           |             |
| FNOR{S}                  | Logical NOR operation   | page 514                           | 3           |
| FNOT(1,2){S}             | Copy negated source   | page 514                           | 3           |
| FONE{S}                  | One fill  | page 514                           | 3           |
| FOR{S}                   | Logical OR operation  | page 514                           | 3           |
| FORNOR(1,2){S}           | Logical OR operation with one inverted source                               | page 514                           | 3           |
| FPACK(16,32, FIX)        | Pixel packing   | page 554,<br>page 556,<br>page 557 | 3           |
| FPADD(16,32){S}          | Pixel add (single) 16- or 32-bit  | page 542                           | 3           |
| FPMERGE                  | Pixel merge   | page 559                           | 3           |
| FPSUB(16,32){S}          | Pixel subtract (single) 16- or 32-bit                                       | page 542                           | 3           |
| FsMULd                   | Floating-point multiply single to double                                    | page 494                           |             |
| FSQRT(s,d,q)             | Floating-point square root  | page 496                           |             |
| FSRC(1,2){S}             | Copy source   | page 514                           | 3           |
| F(s,d,q)TOi              | Convert floating-point to integer   | page 488                           |             |
| F(s,d,q)TO(s,d,q)        | Convert between floating-point formats                                      | page 489                           |             |
| F(s,d,q)TOx              | Convert floating-point to 64-bit integer                                    | page 488                           |             |

**TABLE A-2** Instruction Set (3 of 6)

| Operation                          | Name  | Page     | Ext. to V9? |
|------------------------------------|---|----------|-------------|
| FSUB(s,d,q)                        | Floating-point subtract                         | page 483 |             |
| FXNOR{S}                           | Logical XNOR operation                          | page 514 | 3           |
| FXOR{S}                            | Logical XOR operation                           | page 514 | 3           |
| FxTO(s,d,q)                        | Convert 64-bit integer to floating-point        | page 491 |             |
| FZERO{S}                           | Zero fill                                       | page 514 | 3           |
| ILLTRAP                            | Illegal instruction                             | page 500 |             |
| JMPL                               | Jump and link                                   | page 501 |             |
| LDD <sup>D</sup>                   | Load integer doubleword                         | page 612 |             |
| LDDA <sup>D, P<sub>ASI</sub></sup> | Load integer doubleword from alternate space    | page 614 |             |
| LDDA ASI_NUCLEUS_QUAD*             | Atomic quad load                                | page 510 | 3           |
| LDDF                               | Load double floating-point                      | page 502 |             |
| LDDFA <sup>P<sub>ASI</sub></sup>   | Load double floating-point from alternate space | page 460 |             |
| LDDFA ASI_BLK*                     | Block loads                                     | page 460 | 3           |
| LDDFA ASI_FL*                      | Short floating-point loads (VIS I)              | page 580 | 3           |
| LDF                                | Load floating-point                             | page 502 |             |
| LDFFA <sup>P<sub>ASI</sub></sup>   | Load floating-point from alternate space        | page 502 |             |
| LDFSR <sup>D</sup>                 | Load floating-point state register lower        | page 611 |             |
| LDQF                               | Load quad floating-point                        | page 502 |             |
| LDQFA <sup>P<sub>ASI</sub></sup>   | Load quad floating-point from alternate space   | page 502 |             |
| LDSB                               | Load signed byte                                | page 506 |             |
| LDSBA <sup>P<sub>ASI</sub></sup>   | Load signed byte from alternate space           | page 508 |             |
| LDSH                               | Load signed halfword                            | page 506 |             |
| LDSHA <sup>P<sub>ASI</sub></sup>   | Load signed halfword from alternate space       | page 508 |             |
| LDSTUB                             | Load-store unsigned byte                        | page 512 |             |
| LDSTUBA <sup>P<sub>ASI</sub></sup> | Load-store unsigned byte in alternate space     | page 513 |             |
| LDSW                               | Load signed word                                | page 506 |             |
| LDSWA <sup>P<sub>ASI</sub></sup>   | Load signed word from alternate space           | page 508 |             |
| LDUB                               | Load unsigned byte                              | page 506 |             |
| LDUBA <sup>P<sub>ASI</sub></sup>   | Load unsigned byte from alternate space         | page 508 |             |
| LDUH                               | Load unsigned halfword                          | page 506 |             |
| LDUHA <sup>P<sub>ASI</sub></sup>   | Load unsigned halfword from alternate space     | page 508 |             |
| LDUW                               | Load unsigned word                              | page 506 |             |
| LDUWA <sup>P<sub>ASI</sub></sup>   | Load unsigned word from alternate space         | page 508 |             |
| LDX                                | Load extended                                   | page 506 |             |
| LDXA <sup>P<sub>ASI</sub></sup>    | Load extended from alternate space              | page 508 |             |
| LDXFSR                             | Load floating-point state register              | page 502 |             |
| MEMBAR                             | Memory barrier                                  | page 519 |             |

**TABLE A-2** Instruction Set (4 of 6)

| Operation   | Name  | Page     | Ext. to V9? |
|---|---|----------|-------------|
| MOV <sub>cc</sub>                                   | Move integer register if condition is satisfied           | page 524 |             |
| MOV <sub>r</sub>                                    | Move integer register on contents of integer register     | page 536 |             |
| MUL <sub>sc</sub> <sup>D</sup>                      | Multiply step (and modify condition codes)                | page 616 |             |
| MULX  | Multiply 64-bit integers                                  | page 537 |             |
| NOP   | No operation  | page 538 |             |
| OR, OR <sub>cc</sub>                                | Inclusive OR (and modify condition codes)                 | page 517 |             |
| ORN, ORN <sub>cc</sub>                              | Inclusive OR not (and modify condition codes)             | page 517 |             |
| PDIST   | Pixel component distance                                  | page 552 | 3           |
| POPC  | Population Count  | page 559 |             |
| PREFETCH  | Prefetch data   | page 560 |             |
| PREFETCHA <sup>PASI</sup>                           | Prefetch data from alternate space                        | page 560 |             |
| RDASI   | Read ASI register   | page 568 |             |
| RDASR <sup>PASR</sup>                               | Read ancillary state register                             | page 568 |             |
| RDCCR   | Read condition codes register                             | page 568 |             |
| RDDCR <sup>P</sup>                                  | Read dispatch control register                            | page 568 |             |
| RDFPRS  | Read floating-point registers state register              | page 568 |             |
| RDGSR   | Read graphic status register                              | page 568 |             |
| RDPCL   | Read program counter                                      | page 568 |             |
| RDPCR <sup>P</sup>                                  | Read performance control register                         | page 568 |             |
| RDPIC <sup>PIC</sup>                                | Read performance instrumentation counters                 | page 568 |             |
| R DPR <sup>P</sup>                                  | Read privileged register                                  | page 566 |             |
| RDSOFTINT <sup>P</sup>                              | Read per-processor soft interrupt register                | page 568 |             |
| RDSTICK <sup>P<sub>NPT</sub></sup>                  | Read system TICK register                                 | page 568 |             |
| RDSTICK_CMPR  | Read system TICK compare register                         | page 568 |             |
| RDTICK <sup>P<sub>NPT</sub></sup>                   | Read TICK register  | page 568 |             |
| RDTICK_CMPR <sup>P</sup>                            | Read TICK compare register                                | page 568 |             |
| RDY <sup>D</sup>                                    | Read Y register   | page 619 |             |
| RESTORE   | Restore caller's window                                   | page 572 |             |
| RESTORED <sup>P</sup>                               | Window has been restored                                  | page 574 |             |
| RETRY <sup>P</sup>                                  | Return from trap and retry                                | page 479 |             |
| RETURN  | Return  | page 570 |             |
| SAVE  | Save caller's window                                      | page 572 |             |
| SAVED <sup>P</sup>                                  | Window has been saved                                     | page 574 |             |
| SDIV <sup>D</sup> , SDIV <sub>cc</sub> <sup>D</sup> | 32-bit signed integer divide (and modify condition codes) | page 608 |             |
| SDIVX   | 64-bit signed integer divide                              | page 537 |             |
| SETHI   | Set high 22 bits of low word of integer register          | page 577 |             |
| SHUTDOWN  | Shut down the processor                                   | page 582 | 3           |

**TABLE A-2** Instruction Set (5 of 6)

| Operation                               | Name   | Page               | Ext. to V9? |
|---|--|--------------------|-------------|
| SIAM                                    | Set Interval Arithmetic Mode (VIS II)                    |                    |             |
| SIR                                     | Software-initiated reset                                 | page 583           |             |
| SLL                                     | Shift left logical (IU)                                  | page 578           |             |
| SLLX                                    | Shift left logical, extended (IU)                        | page 578           |             |
| SMUL <sup>D</sup> , SMULcc <sup>D</sup> | Signed integer multiply (and modify condition codes)     | page 616           |             |
| SRA                                     | Shift right arithmetic (IU)                              | page 578           |             |
| SRAX                                    | Shift right arithmetic, extended (IU)                    | page 578           |             |
| SRL                                     | Shift right logical (IU)                                 | page 578           |             |
| SRLX                                    | Shift right logical, extended (IU)                       | page 578           |             |
| STB                                     | Store byte (IU)  | page 588           |             |
| STBA <sup>PASI</sup>                    | Store byte into alternate space (IU)                     | page 589           |             |
| STBAR <sup>D</sup>                      | Store barrier  | page 619           |             |
| STD <sup>D</sup>                        | Store doubleword   | page 622           |             |
| STDA <sup>D, PASI</sup>                 | Store doubleword into alternate space                    | page 623           |             |
| STDF                                    | Store double floating-point (FP)                         | page 584           |             |
| STDFA <sup>PASI</sup>                   | Store double floating-point into alternate space (FP)    | page 586           |             |
| STDFA ASI_BLK*                          | Block stores   | page 460           | 3           |
| STDFA ASI_FL*                           | Short floating-point stores (VIS I)                      | page 580           | 3           |
| STDFA ASI_PST*                          | Partial Store instructions                               | page 540           | 3           |
| STF                                     | Store floating-point (FP)                                | page 584           |             |
| STFA <sup>PASI</sup>                    | Store floating-point into alternate space (FP)           | page 586           |             |
| STFSR <sup>D</sup>                      | Store floating-point state register (FP)                 | page 621           |             |
| STH                                     | Store halfword (IU)                                      | page 588           |             |
| STHA <sup>PASI</sup>                    | Store halfword into alternate space (IU)                 | page 589           |             |
| STQF                                    | Store quad floating-point (FP)                           | page 584           |             |
| STQFA <sup>PASI</sup>                   | Store quad floating-point into alternate space (FP)      | page 586           |             |
| STW                                     | Store word (IU)  | page 588           |             |
| STWA <sup>PASI</sup>                    | Store word into alternate space (IU)                     | page 589           |             |
| STX                                     | Store extended (IU)                                      | page 588           |             |
| STXA <sup>PASI</sup>                    | Store extended into alternate space (IU)                 | page 589           |             |
| STXFSR                                  | Store extended floating-point state register (MS)        | page 584           |             |
| SUB, SUBcc                              | Subtract (and modify condition codes)                    | page 591           |             |
| SUBC, SUBCcc                            | Subtract with carry (and modify condition codes)         | page 591           |             |
| SWAP <sup>D</sup>                       | Swap integer register with memory                        | page 625           |             |
| SWAPA <sup>D, PASI</sup>                | Swap integer register with memory in alternate space     | page 626           |             |
| TADDcc, TADDccTV <sup>D</sup>           | Tagged add and modify condition codes (trap on overflow) | page 592, page 628 |             |




**TABLE A-2** Instruction Set (6 of 6)

| Operation                               | Name  | Page                  | Ext. to V9? |
|---|---|-----------------------|-------------|
| Tcc                                     | Trap on integer condition codes                               | page 595              |             |
| TSUBcc, TSUBccTV <sup>D</sup>           | Tagged subtract and modify condition codes (trap on overflow) | page 593,<br>page 629 |             |
| UDIV <sup>D</sup> , UDIVcc <sup>D</sup> | Unsigned integer divide (and modify condition codes)          | page 608              |             |
| UDIVX                                   | 64-bit unsigned integer divide                                | page 537              |             |
| UMUL <sup>D</sup> , UMULcc <sup>D</sup> | Unsigned integer multiply (and modify condition codes)        | page 616              |             |
| WRASI                                   | Write ASI register  | page 600              |             |
| WRASR <sup>PASR</sup>                   | Write ancillary state register                                | page 600              |             |
| WRCCR                                   | Write condition codes register                                | page 600              |             |
| WRDCR <sup>P</sup>                      | Write dispatch control register                               | page 600              |             |
| WRFPRS                                  | Write floating-point registers state register                 | page 600              |             |
| WRGSR                                   | Write graphic status register                                 | page 600              |             |
| WRPCR <sup>P</sup>                      | Write performance control register                            | page 600              |             |
| WRPIC <sup>Ppic</sup>                   | Write performance instrumentation counters register           | page 600              |             |
| WRPR <sup>P</sup>                       | Write privileged register                                     | page 597              |             |
| WRSOFTINT <sup>P</sup>                  | Write per-processor soft interrupt register                   | page 600              |             |
| WRSOFTINT_CLR <sup>P</sup>              | Clear bits of per-processor soft interrupt register           | page 600              |             |
| WRSOFTINT_SET <sup>P</sup>              | Set bits of per-processor soft interrupt register             | page 600              |             |
| WRTICK_CMPR <sup>P</sup>                | Write TICK compare register                                   | page 600              |             |
| WRSTICK <sup>P</sup>                    | Write System TICK register                                    | page 600              |             |
| WRSTICK_CMPR <sup>P</sup>               | Write System TICK compare register                            | page 600              |             |
| WRY <sup>D</sup>                        | Write Y register  | page 630              |             |
| XNOR, XNORcc                            | Exclusive NOR (and modify condition codes)                    | page 517              |             |
| XOR, XORcc                              | Exclusive OR (and modify condition codes)                     | page 517              |             |

## A.1 Add

| Opcode | Op3     | Operation                                 |
|--------|---------|---|
| ADD    | 00 0000 | Add                                       |
| ADDcc  | 01 0000 | Add and modify condition codes            |
| ADDC   | 00 1000 | Add with Carry                            |
| ADDCcc | 01 1000 | Add with Carry and modify condition codes |

### Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| add                      | <i>reg<sub>rs1</sub>, reg<sub>or_imm</sub>, reg<sub>rd</sub></i> |
| addcc                    | <i>reg<sub>rs1</sub>, reg<sub>or_imm</sub>, reg<sub>rd</sub></i> |
| addc                     | <i>reg<sub>rs1</sub>, reg<sub>or_imm</sub>, reg<sub>rd</sub></i> |
| addccc                   | <i>reg<sub>rs1</sub>, reg<sub>or_imm</sub>, reg<sub>rd</sub></i> |

### Description

ADD and ADDcc compute “ $r[rs1] + r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] + \text{sign\_ext}(simm13)$ ” if  $i = 1$ , and write the sum into  $r[rd]$ .

ADDC and ADDCcc (“ADD with carry”) also add the CCR register’s 32-bit carry ( $icc.c$ ) bit; that is, they compute “ $r[rs1] + r[rs2] + icc.c$ ” or “ $r[rs1] + \text{sign\_ext}(simm13) + icc.c$ ” and write the sum into  $r[rd]$ .

ADDcc and ADDCcc modify the integer condition codes ( $CCR.icc$  and  $CCR.xcc$ ). Overflow occurs on addition if both operands have the same sign and the sign of the sum is different.

---

**Programming Note** – ADDC and ADDC<sub>CC</sub> read the 32-bit condition codes’ carry bit (CCR.i<sub>cc.c</sub>), not the 64-bit condition codes’ carry bit (CCR.x<sub>cc.c</sub>).

---



---

**Compatibility Note** – ADDC and ADDC<sub>CC</sub> were named ADDX and ADDX<sub>CC</sub>, respectively, in the SPARC V8 architecture.

---

### Exceptions


None

---

## A.2 Alignment Instructions (VIS I)

| Opcode              | opf         | Operation  |
|---------------------|-------------|--|
| ALIGNADDRESS        | 0 0001 1000 | Calculate address for misaligned data access               |
| ALIGNADDRESS_LITTLE | 0 0001 1010 | Calculate address for misaligned data access little-endian |
| FALIGNDATA          | 0 0100 1000 | Perform data alignment for misaligned data                 |

### Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| alignaddr                | <i>reg<sub>rs1</sub>, reg<sub>rs2</sub>, reg<sub>rd</sub></i>    |
| alignaddr1               | <i>reg<sub>rs1</sub>, reg<sub>rs2</sub>, reg<sub>rd</sub></i>    |
| faligndata               | <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i> |

## Description

ALIGNADDRESS adds two integer values,  $r[rs1]$  and  $r[rs2]$ , and stores the result (with the least significant three bits forced to zero in the integer register  $r[rd]$ ). The least significant three bits of the result are stored in the `GSR.align` field.

ALIGNADDRESS\_LITTLE is the same as ALIGNADDRESS except that the two's-complement of the least significant 3 bits of the result is stored in `GSR.align`.

---

**Note** – ALIGNADDR\_LITTLE generates the opposite-endian byte ordering for a subsequent FALIGNDATA operation.

---

FALIGNDATA concatenates the two 64-bit floating-point registers specified by `rs1` and `rs2` to form a 128-bit (16-byte) intermediate value. The contents of the first source operand form the more-significant 8 bytes of the intermediate value, and the contents of the second source operand form the less-significant 8 bytes of the intermediate value. Bytes in the intermediate value are numbered from most significant (byte 0) to least significant (byte 15). Eight bytes are extracted from the intermediate value and stored in the 64-bit floating-point destination register specified by `rd`. `GSR.align`, specifying the number of the most significant byte to extract (therefore, the least significant byte extracted from the intermediate value is numbered `GSR.align+7`).

A byte-aligned 64-bit load can be performed as shown in CODE EXAMPLE A-1.

### CODE EXAMPLE A-1 Byte-Aligned 64-bit Load

```
alignaddr  Address, Offset, Address

ldd       [Address], %f0

ldd       [Address + 8], %f2

faligndata %f0, %f2, %f4
```

---

**Programming Note** – For good performance, the result of FALIGNDATA should not be used as a source operand for a 32-bit FP or VIS instruction in the next three instruction groups.

---


## Exceptions

*fp\_disabled*

## A.3 Three-Dimensional Array Addressing Instructions (VIS I)

| Opcode  | opf         | Operation   |
|---------|-------------|---|
| ARRAY8  | 0 0001 0000 | Convert 8-bit 3D address to blocked byte address  |
| ARRAY16 | 0 0001 0010 | Convert 16-bit 3D address to blocked byte address |
| ARRAY32 | 0 0001 0100 | Convert 32-bit 3D address to blocked byte address |

### Format (3)




| Assembly Language Syntax |                                  |
|--------------------------|----------------------------------|
| array8                   | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| array16                  | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| array32                  | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |

### Description

These instructions convert three-dimensional (3D) fixed-point addresses contained in  $r[rs1]$  to a blocked-byte address; they store the result in  $r[rd]$ . Fixed-point addresses typically are used for address interpolation for planar reformatting operations. Blocking is performed at the 64-byte level to maximize L2-cache block reuse, and at the 64 KB level to maximize TLB entry reuse, regardless of the orientation of the address interpolation. These instructions specify an element size of 8 bits (ARRAY8), 16 bits (ARRAY16), or 32 bits (ARRAY32). The second operand,  $r[rs2]$ , specifies the power-of-2 size of the X and Y dimensions of a 3D image array. The legal values for  $rs2$  and their meanings are shown in TABLE A-3. Illegal values produce undefined results in the destination register,  $r[rd]$ .


**TABLE A-3** Three-Dimensional  $r[rs2]$  Array X/Y Dimensions

| $r[rs2]$ value | Number of elements |
|----------------|--------------------|
| 0              | 64                 |
| 1              | 128                |
| 2              | 256                |
| 3              | 512                |
| 4              | 1024               |
| 5              | 2048               |




**FIGURE A-1** Three-Dimensional Array Fixed-Point Address Format

The integer parts of X, Y, and Z are converted to the following blocked-address formats illustrated in FIGURE A-2 through FIGURE A-4.



**FIGURE A-2** Three-Dimensional Array Blocked-Address Format (Array8)



**FIGURE A-3** Three-Dimensional Array Blocked-Address Format (Array16)



**FIGURE A-4** Three-Dimensional Array Blocked-Address Format (Array32)

The bits above Z upper are set to zero. The number of zeroes in the least significant bits is determined by the element size. An element size of 8 bits has no zeroes, an element size of 16 bits has one zero, and an element size of 32 bits has two zeroes. Bits in X and Y above the size specified by `r[rs2]` are ignored.

The code fragment in CODE EXAMPLE A-2 shows assembly of components along an interpolated line at the rate of one component per clock.

**CODE EXAMPLE A-2** Three-Dimensional Array Addressing Example

```

add      Addr, DeltaAddr, Addr

array8   Addr, %g0, bAddr

ldda     [bAddr] ASI_FL8_PRIMARY, data

faligndata data, accum, accum

```

---

**Note** – To maximize reuse of L2-cache and TLB data, software should block array references of a large image to the 64 KB level. This means processing elements within a 32x64x64 block.

---


### *Exceptions*

None

## A.4 Block Load and Block Store (VIS I)

| Opcode         | imm_asi          | ASI Value        | Operation   |
|----------------|------------------|------------------|---|
| LDDFA<br>STDFA | ASI_BLK_AIUP     | 70 <sub>16</sub> | 64-byte block load/store from/to primary address space, privilege mode access only                  |
| LDDFA<br>STDFA | ASI_BLK_AIUS     | 71 <sub>16</sub> | 64-byte block load/store from/to secondary address space, privilege mode access only                |
| LDDFA<br>STDFA | ASI_BLK_AIUPL    | 78 <sub>16</sub> | 64-byte block load/store from/to primary address space, little-endian, privilege mode access only   |
| LDDFA<br>STDFA | ASI_BLK_AIUSL    | 79 <sub>16</sub> | 64-byte block load/store from/to secondary address space, little-endian, privilege mode access only |
| LDDFA<br>STDFA | ASI_BLK_P        | F0 <sub>16</sub> | 64-byte block load/store from/to primary address space  |
| LDDFA<br>STDFA | ASI_BLK_S        | F1 <sub>16</sub> | 64-byte block load/store from/to secondary address space  |
| LDDFA<br>STDFA | ASI_BLK_PL       | F8 <sub>16</sub> | 64-byte block load/store from/to primary address space, little-endian                               |
| LDDFA<br>STDFA | ASI_BLK_SL       | F9 <sub>16</sub> | 64-byte block load/store from/to secondary address space, little-endian                             |
| STDFA          | ASI_BLK_COMMIT_P | E0 <sub>16</sub> | 64-byte block commit store to primary address space   |
| STDFA          | ASI_BLK_COMMIT_S | E1 <sub>16</sub> | 64-byte block commit store to secondary address space   |

*Format (3) LDDFA*





### Format (3) STDFA



| Assembly Language Syntax |   |
|--------------------------|---|
| ldda                     | <i>[reg_addr] imm_asi, freg<sub>rd</sub></i>  |
| ldda                     | <i>[reg_plus_imm] %asi, freg<sub>rd</sub></i> |
| stda                     | <i>freg<sub>rd</sub>, [reg_addr] imm_asi</i>  |
| stda                     | <i>freg<sub>rd</sub>, [reg_plus_imm] %asi</i> |

### Description

A block load (BLD) or block store (BST) instruction uses an LDDFA or STDFA instruction combined with a block transfer ASI. Block transfer ASIs allow BLDs and BSTs to be performed accessing the same address space as normal loads and stores. Little-endian ASIs (those with an ‘L’ suffix) access data in little-endian format; otherwise, the access is assumed to be big-endian. Byte swapping is performed separately for each of the eight double-precision registers used by the instruction. Endianness does not matter if these instructions are only being used for a block copy operation.

A BST with commit forces the data to be written to memory and invalidates copies in all caches present. As a result, a BST with commit maintains coherency with the I-cache.<sup>1</sup> It does not, however, flush instructions that have already been fetched into the pipeline before executing the modified code. If a BST with commit is used to write modified instructions, a FLUSH instruction must still be executed to guarantee that the instruction pipeline is flushed.

LDDFA with a block transfer ASI loads 64 bytes of data from a 64-byte aligned memory area into the eight double-precision floating-point registers specified by *rd*. The lowest addressed eight bytes in memory are loaded into the lowest numbered double-precision destination register. An *illegal\_instruction* exception occurs if the floating-point registers are not aligned on an eight double-precision register boundary. The least significant six bits of the memory address must be zero or a *mem\_address\_not\_aligned* exception occurs.

STDFA with a block transfer ASI stores data from the eight double-precision floating-point registers specified by *rs1* to a 64-byte-aligned memory area. The lowest addressed eight bytes in memory are stored from the lowest numbered double-precision *rd*. An

<sup>1</sup> All store instructions in the processor coherently update the instruction cache. In general SPARC V9 implementations, the store instructions (other than BST with Commit) do not maintain data coherency between instruction and data caches.

*illegal\_instruction* exception occurs if the floating-point registers are not aligned on an eight register boundary. The least significant six bits of the memory address must be zero or a *mem\_address\_not\_aligned* exception occurs.

ASIs  $E0_{16}$  and  $E1_{16}$  are only used for BST with commit operations; they are not used for BLD operations.

---

**Programming Note** – In the UltraSPARC III Cu processor, BLD does not offer a performance advantage over normal loads. For high performance, we recommend the use of prefetch instructions and 8-byte loads. BST and BST commit can offer performance advantage and are used in high performance UltraSPARC III Cu libraries.

---

---

**Programming Note** – BLD does not provide register dependency interlocks, as ordinary load instructions do.

---

Before BLD data can be referenced, a second BLD (to a different set of registers) or a `MEMBAR #Sync` must be performed. If a second BLD is used to synchronize against returning data, the processor will continue execution before all data has been returned. The programmer is then responsible for scheduling instructions so registers are only used when they become valid.

To determine when data is valid, the programmer must count instruction groups containing floating-point operate (FPop) instructions (not FP loads or stores). The lowest numbered destination register of the first BLD may be referenced in the first instruction group following the second BLD, using an FPop instruction only.

The second lowest numbered destination register of the first BLD may be referenced in the second instruction group containing an FPop instruction, and so on.

If this block-load/block-load synchronization mechanism is used, the initial reference to the BLD data must be an FPop instruction (not an FP store), and only instruction groups with FPop instructions are counted when determining BLD data availability.

If these rules are violated, data from before or after the BLD may be returned by a reference to any of the BLD's destination registers.

If a `MEMBAR #Sync` is used to synchronize on BLD data, there are no restrictions on data usage, although performance will be lower than if block-load/block-load synchronization is used. No other `MEMBARs` can be used to provide data synchronization for BLD.

FPop instructions can be issued in a single instruction group with FP stores. If block-load/block-load synchronization is used, FPOps and FP stores can be interlaced. This allows an FPop instruction, such as `FMOVD` or `FALIGNDATA`, to reference the returning data before using the data in any FP store (normal store or BST).

The processor also continues execution, without register interlocks, before all the store data for BSTs are transferred from the register file.

If store source registers are overwritten before the next `BST` or `MEMBAR #Sync` instruction, then the following rule must be observed: The first register can be overwritten in the same instruction group as the `BST`, the second register can be overwritten in the instruction group following the `BST`, and so on. If this rule is violated, the `BST` may use the old or the new (overwritten) data.

When determining correctness for a code sample, note that the processor may interlock more than what is required above. For example, there may be partial register interlocks, such as on the lowest number register.

Code that does not meet the above constraints may appear to work on a particular processor. However, to be portable across all processors similar to the UltraSPARC III Cu processor, all of the above rules should be followed.

## *Rules*

---

**Note** – These instructions are used for transferring large blocks of data (more than 256 bytes), for example, in C library routines `bcopy()` and `bfill()`. They do not allocate in the data cache or L2-cache on a miss. They update the L2-cache on a hit. One `BLD` and, in the most extreme cases, up to fifteen (maximum) `BST`s can be outstanding on the interconnect at one time.

---

To simplify the implementation, `BLD` destination registers may or may not interlock like ordinary load instructions. Before the `BLD` data is referenced, a second `BLD` (to a different set of registers) or a `MEMBAR #Sync` must be performed. If a second `BLD` is used to synchronize with returning data, then it continues execution before all data have been returned. The lowest number register being loaded can be referenced in the first instruction group following the second `BLD`, the second lowest number register can be referenced in the second group, and so on. If this rule is violated, data from before or after the load may be returned.

Similarly, `BST` source data registers are not interlocked against completion of previous load instructions (even if a second `BLD` has been performed). The previous load data must be referenced by some other intervening instruction, or an intervening `MEMBAR #Sync` must be performed. If the programmer violates these rules, data from before or after the load may be used. The load continues execution before all of the store data have been transferred. If store data registers are overwritten before the next `BST` or `MEMBAR #Sync` instruction, then the following rule must be observed. The first register can be overwritten in the same instruction group as the `BST`, the second register can be overwritten in the instruction group following the `BST`, and so on. If this rule is violated, the store may store correct data or the overwritten data.

There must be a `MEMBAR #Sync` or a trap following a `BST` before a `DONE`, `RETRY`, or `WRPR` to `PSTATE` instruction is executed. If this rule is violated, instructions after the `DONE`, `RETRY`, or `WRPR` to `PSTATE` may not see the effects of the updated `PSTATE` register.

BLD does not follow memory model ordering with respect to stores. In particular, read-after-write and write-after-read hazards to overlapping addresses are not detected. The side-effects bit (TTE.E) associated with the access is ignored. Some ordering considerations are as follows:

- If ordering with respect to earlier stores is important (for example, a BLD that overlaps previous stores), then there must be an intervening MEMBAR #StoreLoad or stronger MEMBAR.
- If ordering with respect to later stores is important (for example, a BLD that overlaps a subsequent store), then there must be an intervening MEMBAR #LoadStore or a reference to the BLD data. This restriction does not apply when a trap is taken, so the trap handler does not have to worry about pending BLDs.
- If the BLD overlaps a previous or later store and there is no intervening MEMBAR, then the trap or data referencing the BLD may return data from before or after the store.

BST does not follow memory model ordering with respect to loads, stores, or flushes. In particular, read-after-write, write-after-write, flush-after-write, and write-after-read hazards to overlapping addresses are not detected. The side-effects bit associated with the access is ignored. Some ordering considerations follow:

- If ordering with respect to earlier or later loads or stores is important, then there must be an intervening reference to the load data (for earlier loads) or an appropriate MEMBAR instruction. This restriction does not apply when a trap is taken, so the trap handler does not have to worry about pending BSTs.
- If the BST overlaps a previous load and there is no intervening load data reference or MEMBAR #StoreLoad instruction, then the load may return data from before or after the store and the contents of the block are undefined.
- If the BST overlaps a later load and there is no intervening trap or MEMBAR #LoadStore instruction, then the contents of the block are undefined.
- If the BST overlaps a later store or flush and there is no intervening trap or MEMBAR #Sync instruction, then the contents of the block are undefined.
- If the ordering of two successive BST instructions (overlapping or not) is required, then a MEMBAR #Sync must occur between the BST instructions.

Block operations do not obey the ordering restrictions of the currently selected processor memory model (TSO, PSO, RMO). Block operations always execute under an RMO memory ordering model. Explicit MEMBAR instructions are required to order block operations among themselves or with respect to normal memory operations. In addition, block operations do not conform to dependence order on the issuing processor; that is, no read-after-write, write-after-read, or write-after-write checking occurs between block operations. Explicit MEMBAR #Sync instructions are required to enforce dependence ordering between block operations that reference the same address.

Typically, BLD and BST will be used in loops where software can ensure that the data being loaded and the data being stored do not overlap. The loop will be preceded and followed by the appropriate MEMBARs to ensure that there are no hazards with loads and stores outside the loops. CODE EXAMPLE A-3 demonstrates the loop.

**CODE EXAMPLE A-3** Byte-Aligned Block Copy Inner Loop with BLD/BST

Note that the loop must be unrolled two times to achieve maximum performance. All FP registers are double-precision. Eight versions of this loop are needed to handle all the cases of doubleword misalignment between the source and destination.

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

## *Bcopy Code*

To achieve the highest Bcopy bandwidths, use prefetch instructions and floating-point loads instead of BLD instructions. Using prefetch instructions to bring memory data into the prefetch cache hides all of the latency to memory. This allows a Bcopy loop to run at maximum bandwidth. CODE EXAMPLE A-4 shows how to modify the standard UltraSPARC I bcopy() loop to use PREFETCH and floating-point load instructions instead of BLDs.

### **CODE EXAMPLE A-4** High-Performance UltraSPARC III Cu bcopy() Preamble Code (1 of 2)

preamble:

```
prefetch      [srcaddr],1
prefetch      [srcaddr+0x40],1
prefetch      [srcaddr+0x80],1
prefetch      [srcaddr+0xc0],1
lddf          [srcaddr],%f0
prefetch      [srcaddr+0x100],1
lddf          [srcaddr+0x8],%f2
lddf          [srcaddr+0x10],%f4
faligndata    %f0,%f2,%f32
lddf          [srcaddr+0x18],%f6
faligndata    %f2,%f4,%f34
lddf          [srcaddr+0x20],%f8
faligndata    %f4,%f6,%f36
lddf          [srcaddr+0x28],%f10
faligndata    %f6,%f8,%f38
lddf          [srcaddr+0x30],%f12
faligndata    %f8,%f10,%f40
lddf          [srcaddr+0x38],%f14
faligndata    %f10,%f12,%f42
lddf          [srcaddr+0x40],%f16
subcc         count,0x40,count
bpe          <exit>
add           srcaddr,0x40,srcaddr
```

loop:

```
1  fmovd      %f16,%f0
1  lddf       [srcaddr+0x8],%f2
2  faligndata %f12,%f14,%f44
2  lddf       [srcaddr+0x10],%f4
3  faligndata %f14,%f0,%f46
```

**CODE EXAMPLE A-4** High-Performance UltraSPARC III Cu `bcopy()` Preamble Code (2 of 2)

```
3  stda          %f32, [dstaddr] ASI_BLK_P
3  lddf          [srcaddr+0x18], %f6
4  faligndata   %f0, %f2, %f32
4  lddf          [srcaddr+0x20], %f8
5  faligndata   %f2, %f4, %f34
5  lddf          [srcaddr+0x28], %f10
6  faligndata   %f4, %f6, %f36
6  lddf          [srcaddr+0x30], %f12
7  faligndata   %f6, %f8, %f38
7  lddf          [srcaddr+0x38], %f14
8  faligndata   %f8, %f10, %f40
8  lddf          [srcaddr+0x40], %f16
8  prefetch     [srcaddr+0x100], 1
9  faligndata   %f10, %f12, %f42
9  subcc        count, 0x40, count
9  add          dstaddr, 0x40, dstaddr
9  bpg          loop
1  add          srcaddr, 0x40, srcaddr
```

### *Exceptions*

*fp\_disabled*

*PA\_watchpoint* (recognized on only the first 8 bytes of a transfer)

*VA\_watchpoint* (recognized on only the first 8 bytes of a transfer)

*illegal\_instruction* (misaligned rd)

*mem\_address\_not\_aligned*

*data\_access\_exception*

*data\_access\_error*

*fast\_data\_access\_MMU\_miss*

*fast\_data\_access\_protection*

## A.5 Byte Mask and Shuffle Instructions (VIS II)

| Opcode   | opf         | Operation  |
|----------|-------------|--|
| BMASK    | 0 0001 1001 | Set the GSR.MASK field in preparation for a following BSHUFFLE instruction |
| BSHUFFLE | 0 0100 1100 | Permute bytes as specified by GSR.MASK                                     |

### Format (3)

|          |       |        |       |     |     |
|----------|-------|--------|-------|-----|-----|
| 10       | rd    | 110110 | rs1   | opf | rs2 |
| 31 30 29 | 25 24 | 19 18  | 14 13 | 5 4 | 0   |

| Assembly Language Syntax |                                     |
|--------------------------|-------------------------------------|
| bmask                    | $reg_{rs1}, reg_{rs2}, reg_{rd}$    |
| bshuffle                 | $freg_{rs1}, freg_{rs2}, freg_{rd}$ |

### Description

BMASK adds two integer registers,  $r[rs1]$  and  $r[rs2]$ , and stores the result in the integer register  $r[rd]$ . The least significant 32 bits of the result are stored in the GSR.mask field.

BSHUFFLE concatenates the two 64-bit floating-point registers specified by  $rs1$  (more significant half) and  $rs2$  (less significant half) to form a 16-byte value. Bytes in the concatenated value are numbered from most significant to least significant, with the most significant byte being byte 0. BSHUFFLE extracts 8 of the 16 bytes and stores the result in the 64-bit floating-point register specified by  $rd$ . Bytes in the  $rd$  register are also numbered from most to least significant, with the most significant being byte 0. The following table indicates which source byte is extracted from the concatenated value for each byte in  $rd$ .

| Destination Byte (in $r[rd]$ ) | Source Byte                                  |
|--------------------------------|--|
| 0 (most significant)           | $(r[rs1] \parallel r[rs2])[GSR.mask<31:28>]$ |
| 1                              | $(r[rs1] \parallel r[rs2])[GSR.mask<27:24>]$ |
| 2                              | $(r[rs1] \parallel r[rs2])[GSR.mask<23:20>]$ |
| 3                              | $(r[rs1] \parallel r[rs2])[GSR.mask<19:16>]$ |
| 4                              | $(r[rs1] \parallel r[rs2])[GSR.mask<15:12>]$ |
| 5                              | $(r[rs1] \parallel r[rs2])[GSR.mask<11:8>]$  |
| 6                              | $(r[rs1] \parallel r[rs2])[GSR.mask<7:4>]$   |
| 7 (least significant)          | $(r[rs1] \parallel r[rs2])[GSR.mask<3:0>]$   |



---

**Note** – The BMASK instruction uses the MS pipeline; therefore it cannot be grouped with a store, non-prefetchable load, or a special instruction. The integer rd register result is available after a two-cycle latency. A younger BMASK can be grouped with an older BSHUFFLE (BMASK is “break-after”).

---

Results have a four-cycle latency to other dependent instructions executed in FGA and FGM pipelines. The FGA pipeline is used to execute BSHUFFLE. The GSR mask must be set at or before the instruction group previous to the BSHUFFLE (GSR.mask dependency). BSHUFFLE is fully pipelined (one per cycle).

### Exceptions


*fp\_disabled*

---

## A.6 Branch on Integer Register with Prediction (BPr)

| Opcode | rcond | Operation  | Register Contents Test |
|--------|-------|--|------------------------|
| —      | 000   | <i>Reserved</i>                                  | —                      |
| BRZ    | 001   | Branch on Register Zero                          | $r[rs1] = 0$           |
| BRLEZ  | 010   | Branch on Register Less Than or Equal to Zero    | $r[rs1] \leq 0$        |
| BRLZ   | 011   | Branch on Register Less Than Zero                | $r[rs1] < 0$           |
| —      | 100   | <i>Reserved</i>                                  | —                      |
| BRNZ   | 101   | Branch on Register Not Zero                      | $r[rs1] \neq 0$        |
| BRGZ   | 110   | Branch on Register Greater Than Zero             | $r[rs1] > 0$           |
| BRGEZ  | 111   | Branch on Register Greater Than or Equal to Zero | $r[rs1] \geq 0$        |

### Format (2)



| Assembly Language Syntax            |                                       |
|-------------------------------------|---------------------------------------|
| <code>brz{ ,a}{ ,pt   ,pn}</code>   | <code>reg<sub>rs1</sub>, label</code> |
| <code>brlez{ ,a}{ ,pt   ,pn}</code> | <code>reg<sub>rs1</sub>, label</code> |
| <code>brlz{ ,a}{ ,pt   ,pn}</code>  | <code>reg<sub>rs1</sub>, label</code> |
| <code>brnz{ ,a}{ ,pt   ,pn}</code>  | <code>reg<sub>rs1</sub>, label</code> |
| <code>brgz{ ,a}{ ,pt   ,pn}</code>  | <code>reg<sub>rs1</sub>, label</code> |
| <code>brgez{ ,a}{ ,pt   ,pn}</code> | <code>reg<sub>rs1</sub>, label</code> |

---

**Programming Note** – To set the annul bit for BPr instructions, append “, a” to the opcode mnemonic. For example, use “brz, a %i3, label.” In the preceding table, braces signify that the “, a” is optional. To set the branch prediction bit p, append either “, pt” for predict taken or “, pn” for predict not taken to the opcode mnemonic. If neither “, pt” nor “, pn” is specified, the assembler shall default to “, pt.”

---



---

**Programming Note** – Both BP and BR represent branch on integer register with prediction. They are, in fact, the same instruction.

---

### *Description*

These instructions branch based on the contents of `r[rs1]`. They treat the register contents as a signed integer value.

A BPr instruction examines all 64 bits of `r[rs1]` according to the `rcond` field of the instruction, producing either a TRUE or FALSE result. If TRUE, the branch is taken; that is, the instruction causes a PC-relative, delayed control transfer to the address “PC + (4 \* sign\_ext(d16hi □ d16lo)).” If FALSE, the branch is not taken.

If the branch is taken, the delay instruction is always executed, regardless of the value of the annul bit. If the branch is not taken and the annul bit (a) is one, the delay instruction is annulled (not executed).

The predict bit (p) gives the hardware a hint about whether the branch is expected to be taken. A one in the p bit indicates that the branch is expected to be taken; a zero indicates that the branch is expected not to be taken.

---

**Implementation Note** – The UltraSPARC III Cu processor does not implement this instruction by tagging each register value. The UltraSPARC III Cu processor looks at the full 64-bit register to determine a negative or zero.

---

## Exceptions


*illegal\_instruction* (if  $r_{cond} = 000_2$  or  $100_2$ )

---

# A.7 Branch on Floating-Point Condition Codes with Prediction (FBPfcc)

| Opcode | cond | Operation                               | fcc Test    |
|--------|------|---|-------------|
| FBPA   | 1000 | Branch Always                           | 1           |
| FBPN   | 0000 | Branch Never                            | 0           |
| FBPU   | 0111 | Branch on Unordered                     | U           |
| FBPG   | 0110 | Branch on Greater                       | G           |
| FBPUG  | 0101 | Branch on Unordered or Greater          | G or U      |
| FBPL   | 0100 | Branch on Less                          | L           |
| FBPUL  | 0011 | Branch on Unordered or Less             | L or U      |
| FBPLG  | 0010 | Branch on Less or Greater               | L or G      |
| FBPNE  | 0001 | Branch on Not Equal                     | L or G or U |
| FBPE   | 1001 | Branch on Equal                         | E           |
| FBPUE  | 1010 | Branch on Unordered or Equal            | E or U      |
| FBPGE  | 1011 | Branch on Greater or Equal              | E or G      |
| FBPUGE | 1100 | Branch on Unordered or Greater or Equal | E or G or U |
| FBPLE  | 1101 | Branch on Less or Equal                 | E or L      |
| FBPULE | 1110 | Branch on Unordered or Less or Equal    | E or L or U |
| FBPO   | 1111 | Branch on Ordered                       | E or L or G |

### Format (2)



| cc1 | cc0 | Condition Code |
|-----|-----|----------------|
| 00  |     | <i>fcc0</i>    |
| 01  |     | <i>fcc1</i>    |
| 10  |     | <i>fcc2</i>    |
| 11  |     | <i>fcc3</i>    |

| Assembly Language Syntax      |                     |                        |
|-------------------------------|---------------------|------------------------|
| <i>fba{ ,a}{ ,pt   ,pn}</i>   | <i>%fccn, label</i> |                        |
| <i>fbn{ ,a}{ ,pt   ,pn}</i>   | <i>%fccn, label</i> |                        |
| <i>fbu{ ,a}{ ,pt   ,pn}</i>   | <i>%fccn, label</i> |                        |
| <i>fbg{ ,a}{ ,pt   ,pn}</i>   | <i>%fccn, label</i> |                        |
| <i>fbug{ ,a}{ ,pt   ,pn}</i>  | <i>%fccn, label</i> |                        |
| <i>fb1{ ,a}{ ,pt   ,pn}</i>   | <i>%fccn, label</i> |                        |
| <i>fbul{ ,a}{ ,pt   ,pn}</i>  | <i>%fccn, label</i> |                        |
| <i>fb1g{ ,a}{ ,pt   ,pn}</i>  | <i>%fccn, label</i> |                        |
| <i>fbne{ ,a}{ ,pt   ,pn}</i>  | <i>%fccn, label</i> | <i>(synonym: fbnz)</i> |
| <i>fbe{ ,a}{ ,pt   ,pn}</i>   | <i>%fccn, label</i> | <i>(synonym: fbz)</i>  |
| <i>fbue{ ,a}{ ,pt   ,pn}</i>  | <i>%fccn, label</i> |                        |
| <i>fbge{ ,a}{ ,pt   ,pn}</i>  | <i>%fccn, label</i> |                        |
| <i>fbuge{ ,a}{ ,pt   ,pn}</i> | <i>%fccn, label</i> |                        |
| <i>fble{ ,a}{ ,pt   ,pn}</i>  | <i>%fccn, label</i> |                        |
| <i>fbule{ ,a}{ ,pt   ,pn}</i> | <i>%fccn, label</i> |                        |
| <i>fbo{ ,a}{ ,pt   ,pn}</i>   | <i>%fccn, label</i> |                        |

---

**Programming Note** – To set the annul bit for FBPfcc instructions, append “,a” to the opcode mnemonic. For example, use “fb1 ,a %fcc3 ,label.” In the preceding table, braces signify that the “,a” is optional. To set the branch prediction bit, append either “,pt” (for predict taken) or “,pn” (for predict not taken) to the opcode mnemonic. If neither “,pt” nor “,pn” is specified, the assembler shall default to “,pt.” To select the appropriate floating-point condition code, include “%fcc0”, “%fcc1”, “%fcc2”, or “%fcc3” before the label.

---

### *Description*

Unconditional branches and Fcc-conditional branches are described.

- **Unconditional branches (FBPA, FBPN)** — If its annul field is zero, an FBPN (Floating-Point Branch Never with Prediction) instruction acts like a NOP. If the Branch Never’s annul field is zero, the following (delay) instruction is executed; if the annul field is one, the following instruction is annulled (not executed). In no case does an FBPN cause a transfer of control to take place.

FBPA (Floating-Point Branch Always with Prediction) causes an unconditional PC-relative, delayed control transfer to the address “PC + (4 × sign\_ext (disp19)).” If the annul field of the branch instruction is one, the delay instruction is annulled (not executed). If the annul field is zero, the delay instruction is executed.

- **Fcc-conditional branches** — Conditional FBPFCC instructions (except FBPA and FBPN) evaluate one of the four floating-point condition codes (fcc0, fcc1, fcc2, fcc3) as selected by cc0 and cc1, according to the cond field of the instruction, producing either a TRUE or FALSE result. If TRUE, the branch is taken, that is, the instruction causes a PC-relative, delayed control transfer to the address “PC + (4 × sign\_ext (disp19)).” If FALSE, the branch is not taken.

If a conditional branch is taken, the delay instruction is always executed, regardless of the value of the annul field. If a conditional branch is not taken and the annul field (a) is one, the delay instruction is annulled (not executed). **Note:** The annul bit has a *different* effect on conditional branches than it does on unconditional branches.

The predict bit (p) gives the hardware a hint about whether the branch is expected to be taken. A one in the p bit indicates that the branch is expected to be taken. A zero indicates that the branch is expected not to be taken.

If FPRS.FEF = 0 or PSTATE.PEF = 0, or if an FPU is not present, an FBPFCC instruction is not executed and instead, a *fp\_disabled* exception is generated.

---

**Compatibility Note** – Unlike the SPARC V8 architecture, the SPARC V9 architecture does not require an instruction between a floating-point compare operation and a floating-point branch (FBFCC, FBPFCC).

---


## *Exceptions*

*fp\_disabled*

## A.8 Branch on Integer Condition Codes with Prediction (BPcc)

| Opcode | cond | Operation   | icc Test             |
|--------|------|---|----------------------|
| BPA    | 1000 | Branch Always   | 1                    |
| BPN    | 0000 | Branch Never  | 0                    |
| BPNE   | 1001 | Branch on Not Equal                                     | not Z                |
| BPE    | 0001 | Branch on Equal   | Z                    |
| BPG    | 1010 | Branch on Greater                                       | not (Z or (N xor V)) |
| BPLE   | 0010 | Branch on Less or Equal                                 | Z or (N xor V)       |
| BPGE   | 1011 | Branch on Greater or Equal                              | not (N xor V)        |
| BPL    | 0011 | Branch on Less  | N xor V              |
| BPGU   | 1100 | Branch on Greater Unsigned                              | not (C or Z)         |
| BPLEU  | 0100 | Branch on Less or Equal Unsigned                        | C or Z               |
| BPCC   | 1101 | Branch on Carry Clear (Greater Than or Equal, Unsigned) | not C                |
| BPCS   | 0101 | Branch on Carry Set (Less than, Unsigned)               | C                    |
| BPPOS  | 1110 | Branch on Positive                                      | not N                |
| BPNEG  | 0110 | Branch on Negative                                      | N                    |
| BPVC   | 1111 | Branch on Overflow Clear                                | not V                |
| BPVS   | 0111 | Branch on Overflow Set                                  | V                    |

*Format (2)*



| cc1 | cc0 | Condition Code |
|-----|-----|----------------|
| 00  |     | icc            |
| 01  |     | —              |
| 10  |     | xcc            |
| 11  |     | —              |

| Assembly Language Syntax |                         |                       |
|--------------------------|-------------------------|-----------------------|
| ba{ ,a }{ ,pt   ,pn }    | <i>i_or_x_cc, label</i> |                       |
| bn{ ,a }{ ,pt   ,pn }    | <i>i_or_x_cc, label</i> | (or: iprefetch label) |
| bne{ ,a }{ ,pt   ,pn }   | <i>i_or_x_cc, label</i> | (synonym: bnz)        |
| be{ ,a }{ ,pt   ,pn }    | <i>i_or_x_cc, label</i> | (synonym: bz)         |
| bg{ ,a }{ ,pt   ,pn }    | <i>i_or_x_cc, label</i> |                       |
| ble{ ,a }{ ,pt   ,pn }   | <i>i_or_x_cc, label</i> |                       |
| bge{ ,a }{ ,pt   ,pn }   | <i>i_or_x_cc, label</i> |                       |
| bl{ ,a }{ ,pt   ,pn }    | <i>i_or_x_cc, label</i> |                       |
| bgu{ ,a }{ ,pt   ,pn }   | <i>i_or_x_cc, label</i> |                       |
| bleu{ ,a }{ ,pt   ,pn }  | <i>i_or_x_cc, label</i> |                       |
| bcc{ ,a }{ ,pt   ,pn }   | <i>i_or_x_cc, label</i> | (synonym: bgeu)       |
| bcs{ ,a }{ ,pt   ,pn }   | <i>i_or_x_cc, label</i> | (synonym: blu)        |
| bpos{ ,a }{ ,pt   ,pn }  | <i>i_or_x_cc, label</i> |                       |
| bneg{ ,a }{ ,pt   ,pn }  | <i>i_or_x_cc, label</i> |                       |
| bvc{ ,a }{ ,pt   ,pn }   | <i>i_or_x_cc, label</i> |                       |
| bvs{ ,a }{ ,pt   ,pn }   | <i>i_or_x_cc, label</i> |                       |

---

**Programming Note** – To set the annul bit for BPCC instructions, append “, a” to the opcode mnemonic. For example, use “bgu, a %icc, label.” Braces in the preceding table signify that the “, a” is optional. To set the branch prediction bit, append to an opcode mnemonic either “, pt” for predict taken or “, pn” for predict not taken. If neither “, pt” nor “, pn” is specified, the assembler shall default to “, pt.” To select the appropriate integer condition code, include “%icc” or “%xcc” before the label.

---

### *Description*

Unconditional branches and conditional branches are described below:

- **Unconditional branches (BPA, BPN)** — A BPN (Branch Never with Prediction) instruction for this branch type (op2 = 1) is used in SPARC V9 as an instruction prefetch; that is, the effective address ( $PC + (4 \times \text{sign\_ext}(\text{disp19}))$ ) specifies an address of an instruction that is expected to be executed soon. If the Branch Never’s annul field is one, then the following (delay) instruction is annulled (not executed). If the annul field is zero, then the following instruction is executed. In no case does a Branch Never cause a transfer of control to take place.

BPA (Branch Always with Prediction) causes an unconditional PC-relative, delayed control transfer to the address “PC + (4 × sign\_ext (disp19)).” If the annul field of the branch instruction is one, then the delay instruction is annulled (not executed). If the annul field is zero, then the delay instruction is executed.

- **Conditional branches** — Conditional BPCC instructions (except BPA and BPN) evaluate one of the two integer condition codes (icc or xcc), as selected by cc0 and cc1, according to the cond field of the instruction, producing either a TRUE or FALSE result. If TRUE, the branch is taken; that is, the instruction causes a PC-relative, delayed control transfer to the address “PC + (4 × sign\_ext (disp19)).” If FALSE, the branch is not taken.

If a conditional branch is taken, the delay instruction is always executed regardless of the value of the annul field. If a conditional branch is not taken and the annul field (a) is one, the delay instruction is annulled (not executed). **Note:** The annul bit has a *different* effect for conditional branches than it does for unconditional branches.

The predict bit (p) is used to give the hardware a hint about whether the branch is expected to be taken. A one in the p bit indicates that the branch is expected to be taken; a zero indicates that the branch is expected not to be taken.

### Exceptions


*illegal\_instruction* (cc1 □ cc0 = 01<sub>2</sub> or 11<sub>2</sub>)

---

## A.9 Call and Link

| Opcode | op | Operation     |
|--------|----|---------------|
| CALL   | 01 | Call and Link |

### Format (1)



| Assembly Language Syntax |              |
|--------------------------|--------------|
| call                     | <i>label</i> |



## Description

The CALL instruction causes an unconditional, delayed, PC-relative control transfer to address  $PC + (4 \times \text{sign\_ext}(\text{disp30}))$ . Since the word displacement (`disp30`) field is 30 bits wide, the target address lies within a range of  $-2^{31}$  to  $+2^{31}$  minus four bytes. The PC-relative displacement is formed by sign-extending the 30-bit word displacement field to 62 bits and appending two low-order zeroes to obtain a 64-bit byte displacement.

The CALL instruction also writes the value of PC, which contains the address of the CALL, into `r[15]` (out register 7).

## Exceptions

None

---

# A.10 Compare and Swap

| Opcode                             | op3     | Operation                                      |
|------------------------------------|---------|--|
| <code>CASA</code> <sup>PASI</sup>  | 11 1100 | Compare and Swap Word from Alternate Space     |
| <code>CASXA</code> <sup>PASI</sup> | 11 1110 | Compare and Swap Extended from Alternate Space |

## Format (3)

|          |       |       |          |     |         |     |
|----------|-------|-------|----------|-----|---------|-----|
| 11       | rd    | op3   | rs1      | i=0 | imm_asi | rs2 |
| 11       | rd    | op3   | rs1      | i=1 | —       | rs2 |
| 31 30 29 | 25 24 | 19 18 | 14 13 12 |     | 5 4     | 0   |

| Assembly Language Syntax |   |
|--------------------------|---|
| <code>casa</code>        | <code>[reg<sub>rs1</sub>] imm_asi, reg<sub>rs2</sub>, reg<sub>rd</sub></code> |
| <code>casa</code>        | <code>[reg<sub>rs1</sub>] %asi, reg<sub>rs2</sub>, reg<sub>rd</sub></code>    |
| <code>casxa</code>       | <code>[reg<sub>rs1</sub>] imm_asi, reg<sub>rs2</sub>, reg<sub>rd</sub></code> |
| <code>casxa</code>       | <code>[reg<sub>rs1</sub>] %asi, reg<sub>rs2</sub>, reg<sub>rd</sub></code>    |

## Description

Concurrent processes use these instructions for synchronization and memory updates. Uses of compare-and-swap include spin-lock operations, updates of shared counters, and updates of linked-list pointers. The last two can use wait-free (nonlocking) protocols.

The CASXA instruction compares the value in register  $r[rs2]$  with the doubleword in memory pointed to by the doubleword address in  $r[rs1]$ . If the values are equal, the value in  $r[rd]$  is swapped with the doubleword pointed to by the doubleword address in  $r[rs1]$ . If the values are not equal, the contents of the doubleword pointed to by  $r[rs1]$  replaces the value in  $r[rd]$ , but the memory location remains unchanged.

The CASA instruction compares the low-order 32 bits of register  $r[rs2]$  with a word in memory pointed to by the word address in  $r[rs1]$ . If the values are equal, then the low-order 32 bits of register  $r[rd]$  are swapped with the contents of the memory word pointed to by the address in  $r[rs1]$  and the high-order 32 bits of register  $r[rd]$  are set to zero. If the values are not equal, the memory location remains unchanged, but the zero-extended contents of the memory word pointed to by  $r[rs1]$  replace the low-order 32 bits of  $r[rd]$  and the high-order 32 bits of register  $r[rd]$  are set to zero.

A compare-and-swap instruction comprises three operations: a load, compare, and swap. The overall instruction is atomic; that is, no intervening interrupts or deferred traps are recognized by the processor and no intervening update resulting from a compare-and-swap, swap, load, load-store unsigned byte, or store instruction to the doubleword containing the addressed location, or any portion of it, is performed by the memory system.

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

A compare-and-swap operation behaves as if it performs a store, either of a new value from  $r[rd]$  or of the previous value in memory. The addressed location must be writable, even if the values in memory and  $r[rs2]$  are not equal.

If  $i = 0$ , the address space of the memory location is specified in the `imm_asi` field; if  $i = 1$ , the address space is specified in the ASI register.

A *mem\_address\_not\_aligned* exception is generated if the address in  $r[rs1]$  is not properly aligned. CASXA and CASA cause a *privileged\_action* exception if `PSTATE.PRIV = 0` and bit 7 of the ASI is zero.

The coherence and atomicity of memory operations between processors and I/O DMA memory accesses is maintained for cacheable memory space.

---

**Programming Note** – Compare and Swap (CAS) and Compare and Swap Extended (CASX) synthetic instructions are available for “big-endian” memory accesses. Compare and Swap Little (CASL) and Compare and Swap Extended Little (CASXL) synthetic instructions are available for “little-endian” memory accesses.

---

The compare-and-swap instructions do not affect the condition codes.

### *Exceptions*


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

---

## A.11 DONE and RETRY

| Opcode             | op3     | fcn  | Operation                                    |
|--------------------|---------|------|--|
| DONE <sup>P</sup>  | 11 1110 | 0    | Return from Trap (skip trapped instruction)  |
| RETRY <sup>P</sup> | 11 1110 | 1    | Return from Trap (retry trapped instruction) |
| —                  | 11 1110 | 2–31 | <i>Reserved</i>                              |

### *Format (3)*



| Assembly Language Syntax |  |
|--------------------------|--|
| done                     |  |
| retry                    |  |

## Description

The DONE and RETRY instructions restore the saved state from TSTATE (CWP, ASI, CCR, and PSTATE), set PC and nPC, and decrement TL.

The RETRY instruction resumes execution with the trapped instruction by setting  $PC \leftarrow TPC[TL]$  (the saved value of PC on trap) and  $nPC \leftarrow TNPC[TL]$  (the saved value of nPC on trap).

The DONE instruction skips the trapped instruction by setting  $PC \leftarrow TNPC[TL]$  and  $nPC \leftarrow TNPC[TL] + 4$ .

Execution of a DONE or RETRY instruction in the delay slot of a control transfer instruction produces undefined results.

---

**Programming Note** – Use the DONE and RETRY instructions to return from privileged trap handlers.

---

## Exceptions

*privileged\_opcode*


*illegal\_instruction* (if TL = 0 or fcn = 2–31)

---

## A.12 Edge Handling Instructions (VIS I, VIS II)

| Opcode   | opf         | Operation  |
|----------|-------------|--|
| EDGE8    | 0 0000 0000 | Eight 8-bit edge boundary processing                       |
| EDGE8N   | 0 0000 0001 | Eight 8-bit edge boundary processing, no CC                |
| EDGE8L   | 0 0000 0010 | Eight 8-bit edge boundary processing little-endian         |
| EDGE8LN  | 0 0000 0011 | Eight 8-bit edge boundary processing, little-endian, no CC |
| EDGE16   | 0 0000 0100 | Four 16-bit edge boundary processing                       |
| EDGE16N  | 0 0000 0101 | Four 16-bit edge boundary processing, no CC                |
| EDGE16L  | 0 0000 0110 | Four 16-bit edge boundary processing little-endian         |
| EDGE16LN | 0 0000 0111 | Four 16-bit edge boundary processing, little-endian, no CC |
| EDGE32   | 0 0000 1000 | Two 32-bit edge boundary processing                        |
| EDGE32N  | 0 0000 1001 | Two 32-bit edge boundary processing, no CC                 |
| EDGE32L  | 0 0000 1010 | Two 32-bit edge boundary processing little-endian          |
| EDGE32LN | 0 0000 1011 | Two 32-bit edge boundary processing, little-endian, no CC  |

### Format (3)



| Assembly Language Syntax |                                  |
|--------------------------|----------------------------------|
| edge8                    | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| edge8n                   | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| edge8l                   | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| edge8ln                  | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| edge16                   | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| edge16n                  | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| edge16l                  | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| edge16ln                 | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| edge32                   | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| edge32n                  | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| edge32l                  | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| edge32ln                 | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |

### Description

These instructions handle the boundary conditions for parallel pixel scan line loops, where `src1` is the address of the next pixel to render and `src2` is the address of the last pixel in the scan line.

EDGE8L(N), EDGE16L(N), and EDGE32L(N) are little-endian versions of EDGE8(N), EDGE16(N), and EDGE32(N). They produce an edge mask that is bit reversed from their big-endian counterparts but are otherwise identical. This makes the mask consistent with the mask produced by the graphics compare operations (see Section A.44, “Pixel Compare (VIS I)”) and with the Partial Store instruction (see Section A.41, “Partial Store (VIS I)”) on little-endian data.

A 2-bit (EDGE32), 4-bit (EDGE16), or 8-bit (EDGE8) pixel mask is stored in the least significant bits of `r[rd]`. The mask is computed from left and right edge masks as follows:

1. The left edge mask is computed from the three least significant bits (LSBs) of `r[rs1]`, and the right edge mask is computed from the three LSBs of `r[rs2]`, according to TABLE A-4 (TABLE A-5 for little-endian byte ordering).

2. If 32-bit address masking is disabled ( $PSTATE.AM = 0$ , 64-bit addressing) and the upper 61 bits of  $r[rs1]$  are equal to the corresponding bits in  $r[rs2]$ ,  $r[rd]$  is set to the right edge mask ANDed with the left edge mask.
3. If 32-bit address masking is enabled ( $PSTATE.AM = 1$ , 32-bit addressing) and bits 31:3 of  $r[rs1]$  match bits 31:3 of  $r[rs2]$ ,  $r[rd]$  is set to the right edge mask ANDed with the left edge mask.
4. Otherwise,  $r[rd]$  is set to the left edge mask.

The integer condition codes are set per the rules of the SUBCC instruction with the same operands (see Section A.64, “Subtract”).

The EDGE(8,16,32)(L)N instructions do not set the integer condition codes.

### *Exceptions*

None

**TABLE A-4** Edge Mask Specification

| Edge Size | A2-A0 | Left Edge | Right Edge |
|-----------|-------|-----------|------------|
| 8         | 000   | 1111 1111 | 1000 0000  |
| 8         | 001   | 0111 1111 | 1100 0000  |
| 8         | 010   | 0011 1111 | 1110 0000  |
| 8         | 011   | 0001 1111 | 1111 0000  |
| 8         | 100   | 0000 1111 | 1111 1000  |
| 8         | 101   | 0000 0111 | 1111 1100  |
| 8         | 110   | 0000 0011 | 1111 1110  |
| 8         | 111   | 0000 0001 | 1111 1111  |
| 16        | 00x   | 1111      | 1000       |
| 16        | 01x   | 0111      | 1100       |
| 16        | 10x   | 0011      | 1110       |
| 16        | 11x   | 0001      | 1111       |
| 32        | 0xx   | 11        | 10         |
| 32        | 1xx   | 01        | 11         |

**TABLE A-5** Edge Mask Specification (Little-Endian)


| Edge Size | A2-A0 | Left Edge | Right Edge |
|-----------|-------|-----------|------------|
| 8         | 000   | 1111 1111 | 0000 0001  |
| 8         | 001   | 1111 1110 | 0000 0011  |
| 8         | 010   | 1111 1100 | 0000 0111  |
| 8         | 011   | 1111 1000 | 0000 1111  |
| 8         | 100   | 1111 0000 | 0001 1111  |
| 8         | 101   | 1110 0000 | 0011 1111  |
| 8         | 110   | 1100 0000 | 0111 1111  |
| 8         | 111   | 1000 0000 | 1111 1111  |
| 16        | 00x   | 1111      | 0001       |
| 16        | 01x   | 1110      | 0011       |
| 16        | 10x   | 1100      | 0111       |
| 16        | 11x   | 1000      | 1111       |
| 32        | 0xx   | 11        | 01         |
| 32        | 1xx   | 10        | 11         |

---

## A.13 Floating-Point Add and Subtract

| Opcode | op3     | opf         | Operation       |
|--------|---------|-------------|-----------------|
| FADDs  | 11 0100 | 0 0100 0001 | Add Single      |
| FADDd  | 11 0100 | 0 0100 0010 | Add Double      |
| FADDq  | 11 0100 | 0 0100 0011 | Add Quad        |
| FSUBs  | 11 0100 | 0 0100 0101 | Subtract Single |
| FSUBd  | 11 0100 | 0 0100 0110 | Subtract Double |
| FSUBq  | 11 0100 | 0 0100 0111 | Subtract Quad   |

### Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| fadds                    | <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fadd                     | <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| faddq                    | <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fsubs                    | <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fsubd                    | <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fsubq                    | <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i> |

### Description

The floating-point add instructions add the floating-point register(s) specified by the *rs1* field and the floating-point register(s) specified by the *rs2* field. The instructions then write the sum into the floating-point register(s) specified by the *rd* field.

The floating-point subtract instructions subtract the floating-point register(s) specified by the *rs2* field from the floating-point register(s) specified by the *rs1* field. The instructions then write the difference into the floating-point register(s) specified by the *rd* field.

Rounding is performed as specified by the *FSR.RD* field.



---

**Compatibility Note** – When `FSR.NS = 0`, the processor operates in standard floating-point mode. `FADD` or `FSUB` with a subnormal result causes a *fp\_exception\_other* exception with `FSR.fcc = unfinished_FPop`, system software emulates the instruction, and the correct numerical result is calculated. The UltraSPARC II and UltraSPARC III Cu class of processors operate identically in this case.

When `FSR.NS = 1`, the processor operates in “nonstandard” floating-point mode. When `FSR.NS = 1`, and `FADD` or `FSUB` produces a subnormal result on UltraSPARC II class of processors, the result is replaced by zero in hardware, without trapping. On the UltraSPARC III Cu processor, a *fp\_exception\_other* exception occurs with `FSR.fcc = unfinished_FPop` (even though the processor is operating in nonstandard floating-point mode), then system software emulates the instruction, and the correct numerical result is calculated (instead of replacing the result with zero).

So the processor may produce a different (albeit more accurate) result than in previous processors in the following situation:

`FADD` or `FSUB` produces a subnormal result  
`FSR.NS = 1`

---

### Notes –

1) The processor does not implement (in hardware) the instructions that refer to a quad floating-point register. Execution of such an instruction generates *fp\_exception\_other* (with `fcc = unimplemented_FPop`), which causes a trap. Supervisor software then emulates these instructions.

2) For `FADDs`, `FADDd`, `FSUBs`, `FSUBd`, a *fp\_exception\_other* with `fcc = unfinished_FPop` can occur if either operand is NaN.

---

### Exceptions

*fp\_disabled*

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

*fp\_exception\_other* (`fcc = unimplemented_FPop` (`FADDq` and `FSUBq` only))

*fp\_exception\_other* (`fcc = unfinished_FPop` (`FADDs`, `FADDd`, `FSUBs`, `FSUBd` only))



## Description

These instructions compare the floating-point register(s) specified by the `rs1` field with the floating-point register(s) specified by the `rs2` field, and set the selected floating-point condition code (`fccn`) as shown below.

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

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

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

FCMP causes an invalid (NV) exception if either operand is a signalling NaN.

---

**Compatibility Note** – Unlike the SPARC V8 architecture, the SPARC V9 architecture does not require an instruction between a floating-point compare operation and a floating-point branch (FBfcc, FBPfcc).

SPARC V8 floating-point compare instructions are required to have a zero in the `r[rd]` field. In SPARC V9, bits 26 and 25 of the `r[rd]` field specify the floating-point condition code to be set. Legal SPARC V8 code will work on SPARC V9 because the zeroes in the `r[rd]` field are interpreted as `fcc0` and the FBfcc instruction branches according to `fcc0`.

---

---

**Note** – The processor does not implement (in hardware) the instructions that refer to a quad floating-point register. Execution of such an instruction generates `fp_exception_other` (with `ftt = unimplemented_FPop`), which causes a trap. Supervisor software then emulates these instructions.

---

## Exceptions

*fp\_disabled*

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

*fp\_exception\_other* (`ftt = unimplemented_FPop` (FCMPq, FCMPEq only))

## A.15 Convert Floating-Point to Integer

| Opcode | op3     | opf         | Operation                        |
|--------|---------|-------------|----------------------------------|
| FsTOx  | 11 0100 | 0 1000 0001 | Convert Single to 64-bit Integer |
| FdTOx  | 11 0100 | 0 1000 0010 | Convert Double to 64-bit Integer |
| FqTOx  | 11 0100 | 0 1000 0011 | Convert Quad to 64-bit Integer   |
| FsTOi  | 11 0100 | 0 1101 0001 | Convert Single to 32-bit Integer |
| FdTOi  | 11 0100 | 0 1101 0010 | Convert Double to 32-bit Integer |
| FqTOi  | 11 0100 | 0 1101 0011 | Convert Quad to 32-bit Integer   |

### Format (3)



| Assembly Language Syntax |                         |
|--------------------------|-------------------------|
| fstox                    | $freg_{rs2}, freg_{rd}$ |
| fdtox                    | $freg_{rs2}, freg_{rd}$ |
| fqtox                    | $freg_{rs2}, freg_{rd}$ |
| fstoi                    | $freg_{rs2}, freg_{rd}$ |
| fdtoi                    | $freg_{rs2}, freg_{rd}$ |
| fqtoi                    | $freg_{rs2}, freg_{rd}$ |

### Description

FsTOx, FdTOx, and FqTOx convert the floating-point operand in the floating-point register(s) specified by `rs2` to a 64-bit integer in the floating-point register(s) specified by `rd`.

FsTOi, FdTOi, and FqTOi convert the floating-point operand in the floating-point register(s) specified by `rs2` to a 32-bit integer in the floating-point register specified by `rd`.

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

If the floating-point operand's value is too large to be converted to an integer of the specified size or is a NaN or infinity, then a *fp\_exception\_ieee\_754* “invalid” exception occurs.

---

**Note** – The processor does not implement (in hardware) the instructions that refer to a quad floating-point register. Execution of such an instruction generates *fp\_exception\_other* (with *fmt = unimplemented\_FPop*), which causes a trap. Supervisor software then emulates these instructions.

---

The following floating-point to integer conversion instructions generate an *unfinished\_FPop* exception for certain ranges of floating-point operands, as shown in TABLE A-6.

**TABLE A-6** Floating-Point to Integer *unfinished\_FPop* Exception Conditions

| Instruction | Unfinished Trap Ranges                               |
|-------------|--|
| FsTOi       | result < $-2^{31}$ , result $\geq 2^{31}$ , Inf, NaN |
| FsTOx       | result  $\geq 2^{52}$ , Inf, NaN                     |
| FdTOi       | result < $-2^{31}$ , result $\geq 2^{31}$ , Inf, NaN |
| FdTOx       | result  $\geq 2^{52}$ , Inf, NaN                     |

### Exceptions

*fp\_disabled*

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

*unfinished\_FPop*

*fp\_exception\_other* (*fmt = unimplemented\_FPop* (FqTOi, FqTOx only))

---

## A.16 Convert Between Floating-Point Formats

| Opcode | op3     | opf         | Operation                |
|--------|---------|-------------|--------------------------|
| FsTOd  | 11 0100 | 0 1100 1001 | Convert Single to Double |
| FsTOq  | 11 0100 | 0 1100 1101 | Convert Single to Quad   |
| FdTOs  | 11 0100 | 0 1100 0110 | Convert Double to Single |
| FdTOq  | 11 0100 | 0 1100 1110 | Convert Double to Quad   |
| FqTOs  | 11 0100 | 0 1100 0111 | Convert Quad to Single   |
| FqTOd  | 11 0100 | 0 1100 1011 | Convert Quad to Double   |

### Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| fstod                    | <i>freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fstoq                    | <i>freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fdtos                    | <i>freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fdtoq                    | <i>freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fqtos                    | <i>freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fqtod                    | <i>freg<sub>rs2</sub>, freg<sub>rd</sub></i> |

### Description

These instructions convert the floating-point operand in the floating-point register(s) specified by *rs2* to a floating-point number in the destination format. They write the result into the floating-point register(s) specified by *rd*.

Rounding is performed as specified by the `FSR.RD` field.

`FQTOD`, `FQTOS`, and `FDTOS` (the “narrowing” conversion instructions) can raise `OF`, `UF`, and `NX` exceptions. `FDTOQ`, `FSTOQ`, and `FSTOD` (the “widening” conversion instructions) cannot.

Any of these six instructions can trigger an `NV` exception if the source operand is a signalling NaN.

---

### Notes –

1) The UltraSPARC III Cu processor does not implement (in hardware) the instructions that refer to a quad floating-point register. Execution of such an instruction generates `fp_exception_other` (with `fmt = unimplemented_FPop`), which causes a trap. Supervisor software then emulates these instructions.

2) For `FDTOS` and `FSTOD`, a `fp_exception_other` with `fmt = unfinished_FPop` can occur if source operand is NaN or subnormal or out of range of the destination format.

---

The following floating-point to floating-point conversion instructions generate an *unfinished\_FPop* exception for certain ranges of floating-point operands, as shown in TABLE A-7.

**TABLE A-7** Floating-Point/Floating-Point *unfinished\_FPop* Exception Conditions

| Instruction | Unfinished Trap Ranges  |
|-------------|---|
| FdTOs       | $ \text{result}  \geq 2^{52}$ , $ \text{result}  < 2^{-31}$ , operand $< -2^{22}$ , operand $\geq 2^{22}$ , NaN |

### Exceptions

*fp\_disabled*

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

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

*unfinished\_FPop*

*fp\_exception\_other* (ftt = *unfinished\_FPop* (FdTOs and FsTOd only))

## A.17 Convert Integer to Floating-Point

| Opcode | op3     | opf         | Operation                        |
|--------|---------|-------------|----------------------------------|
| FxTOs  | 11 0100 | 0 1000 0100 | Convert 64-bit Integer to Single |
| FxTOd  | 11 0100 | 0 1000 1000 | Convert 64-bit Integer to Double |
| FxTOq  | 11 0100 | 0 1000 1100 | Convert 64-bit Integer to Quad   |
| FiTOs  | 11 0100 | 0 1100 0100 | Convert 32-bit Integer to Single |
| FiTOd  | 11 0100 | 0 1100 1000 | Convert 32-bit Integer to Double |
| FiTOq  | 11 0100 | 0 1100 1100 | Convert 32-bit Integer to Quad   |

### Format (3)

|          |       |       |       |     |     |
|----------|-------|-------|-------|-----|-----|
| 10       | rd    | op3   | —     | opf | rs2 |
| 31 30 29 | 25 24 | 19 18 | 14 13 | 5 4 | 0   |

| Assembly Language Syntax |  |
|--------------------------|--|
| <code>fxtos</code>       | <code>reg<sub>rs2</sub>, reg<sub>rd</sub></code> |
| <code>fxtod</code>       | <code>reg<sub>rs2</sub>, reg<sub>rd</sub></code> |
| <code>fxtoq</code>       | <code>reg<sub>rs2</sub>, reg<sub>rd</sub></code> |
| <code>fitos</code>       | <code>reg<sub>rs2</sub>, reg<sub>rd</sub></code> |
| <code>fitod</code>       | <code>reg<sub>rs2</sub>, reg<sub>rd</sub></code> |
| <code>fitoq</code>       | <code>reg<sub>rs2</sub>, reg<sub>rd</sub></code> |

### Description

`FXTOS`, `FXTOD`, and `FXTOQ` convert the 64-bit signed integer operand in the floating-point registers specified by `rs2` into a floating-point number in the destination format.

`FITOS`, `FITOD`, and `FITOQ` convert the 32-bit signed integer operand in floating-point register(s) specified by `rs2` into a floating-point number in the destination format. All write their result into the floating-point register(s) specified by `rd`.

`FITOS`, `FXTOS`, and `FXTOD` round as specified by the `FSR.RD` field.

---

**Note** – The UltraSPARC III Cu processor does not implement (in hardware) the instructions that refer to a quad floating-point register. Execution of such an instruction generates `fp_exception_other` (with `ftt = unimplemented_FPop`), which causes a trap. Supervisor software then emulates these instructions.

---

The following integer to floating-point conversion instructions generate an `unfinished_FPop` exception for certain ranges of integer operands, as shown in TABLE A-8.

**TABLE A-8** Integer/Floating-Point `unfinished_FPop` Exception Conditions

| Instruction        | Unfinished Trap Ranges                      |
|--------------------|---|
| <code>FITOS</code> | operand < $-2^{22}$ , operand $\geq 2^{22}$ |
| <code>FXTOS</code> | operand < $-2^{22}$ , operand $\geq 2^{22}$ |
| <code>FXTOD</code> | operand < $-2^{51}$ , operand $\geq 2^{51}$ |

### Exceptions

`fp_disabled`

`fp_exception_ieee_754` (NX (`FITOS`, `FXTOS`, `FXTOD` only))

`unfinished_FPop`

`fp_exception_other` (`ftt = unimplemented_FPop` (`FITOQ`, `FXTOQ` only))



## A.18 Floating-Point Move

| Opcode            | op3     | opf         | Operation             |
|-------------------|---------|-------------|-----------------------|
| FMOV <sub>s</sub> | 11 0100 | 0 0000 0001 | Move Single           |
| FMOV <sub>d</sub> | 11 0100 | 0 0000 0010 | Move Double           |
| FMOV <sub>q</sub> | 11 0100 | 0 0000 0011 | Move Quad             |
| FNEG <sub>s</sub> | 11 0100 | 0 0000 0101 | Negate Single         |
| FNEG <sub>d</sub> | 11 0100 | 0 0000 0110 | Negate Double         |
| FNEG <sub>q</sub> | 11 0100 | 0 0000 0111 | Negate Quad           |
| FABS <sub>s</sub> | 11 0100 | 0 0000 1001 | Absolute Value Single |
| FABS <sub>d</sub> | 11 0100 | 0 0000 1010 | Absolute Value Double |
| FABS <sub>q</sub> | 11 0100 | 0 0000 1011 | Absolute Value Quad   |

*Format (3)*



| Assembly Language Syntax |  |
|--------------------------|--|
| fmovs                    | <i>freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fmovd                    | <i>freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fmovq                    | <i>freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fnegs                    | <i>freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fnegd                    | <i>freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fnegq                    | <i>freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fabss                    | <i>freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fabsd                    | <i>freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fabsq                    | <i>freg<sub>rs2</sub>, freg<sub>rd</sub></i> |

## Description

The single-precision versions of these instructions copy the contents of a single-precision floating-point register to the destination. The double-precision versions copy the contents of a double-precision floating-point register to the destination. The quad-precision versions copy a quad-precision value in floating-point registers to the destination.

FMOV copies the source to the destination unaltered.

FNEG copies the source to the destination with the sign bit complemented.

FABS copies the source to the destination with the sign bit cleared.

These instructions do not round.

---

**Note** – The processor does not implement (in hardware) the instructions that refer to a quad floating-point register. Execution of such an instruction generates *fp\_exception\_other* (with `fmt = unimplemented_FPop`), which causes a trap. Supervisor software then emulates these instructions.

---

## Exceptions

*fp\_disabled*


*fp\_exception\_other* (`fmt = unimplemented_FPop` (FMOV<sub>q</sub>, FNEG<sub>q</sub>, FABS<sub>q</sub> only))

---

# A.19 Floating-Point Multiply and Divide

| Opcode | op3     | opf         | Operation                 |
|--------|---------|-------------|---------------------------|
| FMULs  | 11 0100 | 0 0100 1001 | Multiply Single           |
| FMULd  | 11 0100 | 0 0100 1010 | Multiply Double           |
| FMULq  | 11 0100 | 0 0100 1011 | Multiply Quad             |
| FsMULd | 11 0100 | 0 0110 1001 | Multiply Single to Double |
| FdMULq | 11 0100 | 0 0110 1110 | Multiply Double to Quad   |
| FDIVs  | 11 0100 | 0 0100 1101 | Divide Single             |
| FDIVd  | 11 0100 | 0 0100 1110 | Divide Double             |
| FDIVq  | 11 0100 | 0 0100 1111 | Divide Quad               |

### Format (3)



| <b>Assembly Language Syntax</b> |  |
|---------------------------------|--|
| fmuls                           | <i>freg<sub>rs1</sub></i> , <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i> |
| fmuld                           | <i>freg<sub>rs1</sub></i> , <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i> |
| fmulq                           | <i>freg<sub>rs1</sub></i> , <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i> |
| fsmuld                          | <i>freg<sub>rs1</sub></i> , <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i> |
| fdmulq                          | <i>freg<sub>rs1</sub></i> , <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i> |
| fdivs                           | <i>freg<sub>rs1</sub></i> , <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i> |
| fdivd                           | <i>freg<sub>rs1</sub></i> , <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i> |
| fdivq                           | <i>freg<sub>rs1</sub></i> , <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i> |

### Description

The floating-point multiply instructions multiply the contents of the floating-point register(s) specified by the `rs1` field by the contents of the floating-point register(s) specified by the `rs2` field. The instructions then write the product into the floating-point register(s) specified by the `rd` field.

The `FSMULD` instruction provides the exact double-precision product of two single-precision operands, without underflow, overflow, or rounding error. Similarly, `FDMULQ` provides the exact quad-precision product of two double-precision operands.

The floating-point divide instructions divide the contents of the floating-point register(s) specified by the `rs1` field by the contents of the floating-point register(s) specified by the `rs2` field. The instructions then write the quotient into the floating-point register(s) specified by the `rd` field.

Rounding is performed as specified by the `FSR.RD` field.

---

### Notes –

1) The processor does not implement (in hardware) the instructions that refer to a quad floating-point register. Execution of such an instruction generates *fp\_exception\_other* (with *ftt* = *unimplemented\_FPop*), which causes a trap. Supervisor software then emulates these instructions.

2) For *FDIVs* and *FDIVd*, a *fp\_exception\_other* with *ftt* = *unfinished\_FPop* can occur if the divide unit detects certain unusual conditions.

---

### Exceptions

*fp\_disabled*

*fp\_exception\_ieee\_754* (OF, UF, DZ (FDIV only), NV, NX)

*fp\_exception\_other* (*ftt* = *unimplemented\_FPop* (FMULq, FdMULq, FDIVq))

*fp\_exception\_other* (*ftt* = *unfinished\_FPop* (FMULs, FMULd, FSMULd, FDIVs, FDIV))

---

## A.20 Floating-Point Square Root

| Opcode | op3     | opf         | Operation          |
|--------|---------|-------------|--------------------|
| FSQRTs | 11 0100 | 0 0010 1001 | Square Root Single |
| FSQRTd | 11 0100 | 0 0010 1010 | Square Root Double |
| FSQRTq | 11 0100 | 0 0010 1011 | Square Root Quad   |

### Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| <i>fsqrts</i>            | <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i> |
| <i>fsqrtd</i>            | <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i> |
| <i>fsqrtq</i>            | <i>freg<sub>rs2</sub></i> , <i>freg<sub>rd</sub></i> |

## Description

These SPARC V9 instructions generate the square root of the floating-point operand in the floating-point register(s) specified by the `rs2` field and place the result in the destination floating-point register(s) specified by the `rd` field. Rounding is performed as specified by the `FSR.RD` field.

---

**Note** – The processor does not implement (in hardware) the instructions that refer to a quad floating-point register. Execution of such an instruction generates *fp\_exception\_other* (with `ftt = unimplemented_FPop`), which causes a trap. Supervisor software then emulates these instructions.

---

For `FSQRTs` and `FSQRTd` a *fp\_exception\_other* (with `ftt = unfinished_FPop`) can occur if the operand to the square root is positive denormalized.

## Exceptions

*fp\_disabled*

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

*fp\_exception\_other* (*unimplemented\_FPop*) (Quad forms)

*fp\_exception\_other* (*unfinished\_FPop*) (`FSQRTs`, `FSQRTd`)

---

## A.21 Flush Instruction Memory

| Opcode | op3     | Operation                |
|--------|---------|--------------------------|
| FLUSH  | 11 1011 | Flush Instruction Memory |

### Format (3)

|          |       |       |          |     |        |     |
|----------|-------|-------|----------|-----|--------|-----|
| 10       | —     | op3   | rs1      | i=0 | —      | rs2 |
| 10       | —     | op3   | rs1      | i=1 | simm13 |     |
| 31 30 29 | 25 24 | 19 18 | 14 13 12 | 5 4 | 0      |     |

| Assembly Language Syntax |                      |
|--------------------------|----------------------|
| <code>flush</code>       | <code>address</code> |

## Description

FLUSH ensures that the doubleword specified as the effective address is consistent across any local caches, and in a multiprocessor system, will eventually become consistent everywhere.

In the following discussion  $P_{\text{FLUSH}}$  refers to the processor that executed the FLUSH instruction.

FLUSH ensures that instruction fetches from the specified effective address by  $P_{\text{FLUSH}}$  appear to execute after any loads, stores, and atomic load-stores to that address issued by  $P_{\text{FLUSH}}$  prior to the FLUSH. In a multiprocessor system, FLUSH also ensures that these values will eventually become visible to the instruction fetches of all other processors. FLUSH behaves as if it were a store with respect to MEMBAR-induced orderings. See Section A.34, “Memory Barrier.”

The effective address operand for the FLUSH instruction is “ $r[\text{rs1}] + r[\text{rs2}]$ ” if  $i = 0$ , or “ $r[\text{rs1}] + \text{sign\_ext}(\text{simml3})$ ” if  $i = 1$ . The least significant two address bits of the effective address are unused and should be supplied as zeroes by software. Bit 2 of the address is ignored because FLUSH operates on at least a doubleword.

---

### Programming Note –

1. Typically, FLUSH is used in self-modifying code. The use of self-modifying code is discouraged.
  2. The order in which memory is modified can be controlled by means of FLUSH and MEMBAR instructions interspersed appropriately between stores and atomic load-stores. FLUSH is needed only between a store and a subsequent instruction fetch from the modified location. When multiple processes may concurrently modify live (that is, potentially executing) code, the programmer must ensure that the order of update maintains the program in a semantically correct form at all times.
  3. The memory model guarantees in a uniprocessor that *data* loads observe the results of the most recent store, even if there is no intervening FLUSH.
  4. FLUSH may be time consuming.
  5. In a multiprocessor system, the time it takes for a FLUSH to take effect is dependent on the system. No mechanism is provided to ensure or test completion.
  6. Because FLUSH is designed to act on a doubleword and on some implementations FLUSH may trap to system software, system software should provide a user-call service routine for flushing arbitrarily sized regions of memory. On some processor implementations, this routine would issue a series of FLUSH instructions; on others, it might issue a single trap to system software that would then flush the entire region.
- 

On an UltraSPARC III Cu processor:

- A FLUSH instruction flushes the processor pipeline and synchronizes the processor.
- The instruction cache is kept coherent so there is no need to perform any action on it.

- The address provided with the FLUSH instruction is ignored. However, for portability across all SPARC V9 implementations, software must supply the target effective address in FLUSH instructions.

FLUSH synchronizes code and data spaces after code space is modified during program execution. The FLUSH effective address is ignored. FLUSH does not access the data MMU and cannot generate a data MMU miss or exception.

SPARC V9 specifies that the FLUSH instruction has no latency on the issuing 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 operation is performed, the processor guarantees that earlier code modifications will be visible across the whole system.

### *Exceptions*

None

---

## A.22 Flush Register Windows

| Opcode | op3     | Operation              |
|--------|---------|------------------------|
| FLUSHW | 10 1011 | Flush Register Windows |

### *Format (3)*




---

#### Assembly Language Syntax

---

```
flushw
```

---

### *Description*

FLUSHW causes all active register windows except the current window to be flushed to memory at locations determined by privileged software. FLUSHW behaves as a NOP if there are no active windows other than the current window. At the completion of the FLUSHW instruction, the only active register window is the current one.

---

**Programming Note** – The `FLUSHW` instruction can be used by application software to switch memory stacks or to examine register contents for previous stack frames.

---

`FLUSHW` acts as a `NOP` if `CANSAVE = NWINDOWS - 2`. Otherwise, there is more than one active window, so `FLUSHW` causes a spill exception. The trap vector for the spill exception is based on the contents of `OTHERWIN` and `WSTATE`. The spill trap handler is invoked with the `CWP` set to the window to be spilled (that is,  $(CWP + CANSAVE + 2) \bmod NWINDOWS$ ).

---

**Programming Note** – Typically, the spill handler saves a window on a memory stack and returns to re-execute the `FLUSHW` instruction. Thus, `FLUSHW` traps and re-executes until all active windows other than the current window have been spilled.

---

### *Exceptions*

*spill\_n\_normal*


*spill\_n\_other*

---

## A.23 Illegal Instruction Trap

| Opcode  | op | op2 | Operation                       |
|---------|----|-----|---------------------------------|
| ILLTRAP | 00 | 000 | <i>illegal_instruction</i> trap |

### *Format (2)*



| Assembly Language Syntax |                      |
|--------------------------|----------------------|
| <code>illtrap</code>     | <code>const22</code> |



## Description

The ILLTRAP instruction causes an *illegal\_instruction* exception. The `const22` value is ignored by the hardware; specifically, this field is *not* reserved by the architecture for any future use.

---

**Compatibility Note** – Except for its name, this instruction is identical to the SPARC V8 UNIMP instruction.

---

## Exceptions


*illegal\_instruction*

---

# A.24 Jump and Link

| Opcode | op3     | Operation     |
|--------|---------|---------------|
| JMPL   | 11 1000 | Jump and Link |

## Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| <code>jmp1</code>        | <code>address, reg<sub>rd</sub></code> |

## Description

The JMPL instruction causes a register-indirect delayed control transfer to the address given by “`r[rs1] + r[rs2]`” if `i = 0`, or “`r[rs1] + sign_ext(simm13)`” if `i = 1`.

The JMPL instruction copies the PC, which contains the address of the JMPL instruction, into register `r[rd]`.

If either of the low-order two bits of the jump address is nonzero, a *mem\_address\_not\_aligned* exception occurs.

---

**Programming Note** – A `JMPL` instruction with `rd = 15` functions as a register-indirect call using the standard link register.

`JMPL` with `rd = 0` can be used to return from a subroutine. The typical return address is “`r[31] + 8`,” if a nonleaf routine (one that uses the `SAVE` instruction) is entered by a `CALL` instruction, or “`r[15] + 8`” if a leaf routine (one that does not use the `SAVE` instruction) is entered by a `CALL` instruction or by a `JMPL` instruction with `rd = 15`.

---

### Exceptions

*mem\_address\_not\_aligned*


---

## A.25 Load Floating-Point

| Opcode | op3     | rd   | Operation                           |
|--------|---------|------|-------------------------------------|
| LDF    | 10 0000 | 0–31 | Load Floating-Point Register        |
| LDDF   | 10 0011 | †    | Load Double Floating-Point Register |
| LDQF   | 10 0010 | †    | Load Quad Floating-Point Register   |
| LDXFSR | 10 0001 | 1    | Load Floating-Point State Register  |
| —      | 10 0001 | 2–31 | <i>Reserved</i>                     |

† Encoded floating-point register value.

### Format (3)



| Assembly Language Syntax |                                     |
|--------------------------|-------------------------------------|
| ld                       | [address], <i>freg<sub>rd</sub></i> |
| ldd                      | [address], <i>freg<sub>rd</sub></i> |
| ldq                      | [address], <i>freg<sub>rd</sub></i> |
| ldx                      | [address], %f <i>sr</i>             |

### Description

The load single floating-point instruction (LDF) copies a word from memory into  $f[rd]$ .

The load doubleword floating-point instruction (LDDF) copies a word-aligned doubleword from memory into a double-precision floating-point register.

The load quad floating-point instruction (LDQF) traps to software.

The load floating-point state register instruction (LDXFSR) waits for all FPop instructions that have not finished execution to complete and then loads a doubleword from memory into the FSR.

Load floating-point instructions access the primary address space ( $ASI = 80_{16}$ ). The effective address for these instructions is “ $r[rs1] + r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] + sign\_ext(simml3)$ ” if  $i = 1$ .

LDF causes a *mem\_address\_not\_aligned* exception if the effective memory address is not word aligned. LDXFSR causes a *mem\_address\_not\_aligned* exception if the address is not doubleword aligned. If the floating-point unit is not enabled (per `FPRS.FEF` and `PSTATE.PEF`) or if no FPU is present, then a load floating-point instruction causes a *fp\_disabled* exception.

LDDF requires doubleword aligned. If word alignment is used, then the LDDF causes a *LDDF\_mem\_address\_not\_aligned* exception. The trap handler software shall emulate the LDDF instruction and return.

---

**Programming Note** – In SPARC V8, some compilers issued sequences of single-precision loads when they could not determine that doubleword or quadword operands were properly aligned. For SPARC V9, since emulation of misaligned loads is expected to be fast, we recommend that compilers issue sets of single-precision loads only when they can determine that doubleword or quadword operands are *not* properly aligned.

---

If a load floating-point instruction traps with any type of access error, the contents of the destination floating-point register(s) is undefined.

In an UltraSPARC III Cu processor, an LDDF instruction causes a *LDDF\_mem\_address\_not\_aligned* trap if the effective address is 32-bit aligned but not 64-bit (doubleword) aligned.

### Exceptions

*illegal\_instruction* (op3=21<sub>16</sub> and rd = 2–31)  
*fp\_disabled*  
*LDDF\_mem\_address\_not\_aligned* (LDDF only)  
*mem\_address\_not\_aligned*  
*data\_access\_exception*  
*PA\_watchpoint*  
*VA\_watchpoint*  
*data\_access\_error*  
*fast\_data\_access\_MMU\_miss*  
*fast\_data\_access\_protection*

## A.26 Load Floating-Point from Alternate Space

| Opcode                | op3     | rd   | Operation  |
|-----------------------|---------|------|--|
| LDF <sup>PASI</sup>   | 11 0000 | 0–31 | Load Floating-Point Register from Alternate Space        |
| LDDFA <sup>PASI</sup> | 11 0011 | †    | Load Double Floating-Point Register from Alternate Space |
| LDQFA <sup>PASI</sup> | 11 0010 | †    | Load Quad Floating-Point Register from Alternate Space   |

† Encoded floating-point register value.

### Format (3)



| Assembly Language Syntax |   |
|--------------------------|---|
| lda                      | [regaddr] imm_asi, freg <sub>rd</sub>   |
| lda                      | [reg_plus_imm] %asi, freg <sub>rd</sub> |
| ldda                     | [regaddr] imm_asi, freg <sub>rd</sub>   |
| ldda                     | [reg_plus_imm] %asi, freg <sub>rd</sub> |
| ldqa                     | [regaddr] imm_asi, freg <sub>rd</sub>   |
| ldqa                     | [reg_plus_imm] %asi, freg <sub>rd</sub> |

### Description

The load single floating-point from alternate space instruction (LDFA) copies a word from memory into f[rd].

The load double floating-point from alternate space instruction (LDDFA) copies a word-aligned doubleword from memory into a double-precision floating-point register.

The load quad floating-point from alternate space instruction (LDQFA) traps to software.

Load floating-point from alternate space instructions contain the address space identifier (ASI) to be used for the load in the imm\_asi field if i = 0, or in the ASI register if i = 1. The access is privileged if bit 7 of the ASI is zero; otherwise, it is not privileged. The effective address for these instructions is “r[rs1] + r[rs2]” if i = 0, or “r[rs1] + sign\_ext(simm13)” if i = 1.

LDFA causes a *mem\_address\_not\_aligned* exception if the effective memory address is not word aligned. If the floating-point unit is not enabled (per FPRS.FEF and PSTATE.PEF) or if no FPU is present, then load floating-point from alternate space instructions cause a *fp\_disabled* exception.

LDDFA with certain target ASIs is defined to be a 64-byte block-load instruction. See Section A.4, “Block Load and Block Store (VIS I)” for details.

---

**Implementation Note** – LDFA and LDDFA cause a *privileged\_action* exception if PSTATE.PRIV = 0 and bit 7 of the ASI is zero.

---

LDDF requires doubleword alignment. If word alignment is used, then the LDDF causes a *LDDF\_mem\_address\_not\_aligned* exception. The trap handler software shall emulate the LDDF instruction and return.

---

**Programming Note** – In SPARC V8, some compilers issued sequences of single-precision loads when they could not determine that doubleword or quadword operands were properly aligned. For SPARC V9, since emulation of misaligned loads is expected to be fast, compilers should issue sets of single-precision loads only when they can determine that doubleword or quadword operands are *not* properly aligned.

---

If a load floating-point instruction traps with any type of access error, the contents of the destination floating-point register(s) is undefined.

In an UltraSPARC III Cu processor, an LDDFA instruction causes an *LDDF\_mem\_address\_not\_aligned* trap if the effective address is 32-bit aligned but not 64-bit (doubleword) aligned.

### *Exceptions*

*illegal\_instruction* (LDQFA only)  
*fp\_disabled*  
*LDDF\_mem\_address\_not\_aligned* (LDDFA only)  
*mem\_address\_not\_aligned*  
*privileged\_action*  
*data\_access\_exception*  
*data\_access\_error*  
*fast\_data\_access\_MMU\_miss*  
*fast\_data\_access\_protection*  
*VA\_watchpoint*  
*PA\_watchpoint*

---

## A.27 Load Integer

| Opcode | op3     | Operation              |
|--------|---------|------------------------|
| LDSB   | 00 1001 | Load Signed Byte       |
| LDSH   | 00 1010 | Load Signed Halfword   |
| LDSW   | 00 1000 | Load Signed Word       |
| LDUB   | 00 0001 | Load Unsigned Byte     |
| LDUH   | 00 0010 | Load Unsigned Halfword |
| LDUW   | 00 0000 | Load Unsigned Word     |
| LDX    | 00 1011 | Load Extended Word     |

### Format (3)



| Assembly Language Syntax |                                    |               |
|--------------------------|------------------------------------|---------------|
| ldsb                     | [address], <i>reg<sub>rd</sub></i> |               |
| ldsh                     | [address], <i>reg<sub>rd</sub></i> |               |
| ldsw                     | [address], <i>reg<sub>rd</sub></i> |               |
| ldub                     | [address], <i>reg<sub>rd</sub></i> |               |
| lduh                     | [address], <i>reg<sub>rd</sub></i> |               |
| lduw                     | [address], <i>reg<sub>rd</sub></i> | (synonym: ld) |
| ldx                      | [address], <i>reg<sub>rd</sub></i> |               |

### Description

The load integer instructions copy a byte, a halfword, a word, or an extended word from memory. All copy the fetched value into  $r[rd]$ . A fetched byte, halfword, or word is right-justified in the destination register  $r[rd]$ ; it is either sign-extended or zero-filled on the left, depending on whether the opcode specifies a signed or unsigned operation, respectively.

Load integer instructions access the primary address space (ASI = 80<sub>16</sub>). The effective address is “ $r[rs1] + r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] + \text{sign\_ext}(\text{simm13})$ ” if  $i = 1$ .

A successful load (notably, load extended) instruction operates atomically.

LDUH and LDSH cause a *mem\_address\_not\_aligned* exception if the address is not halfword aligned. LDUW and LDSW cause a *mem\_address\_not\_aligned* exception if the effective address is not word aligned. LDX causes a *mem\_address\_not\_aligned* exception if the address is not doubleword aligned.

---

**Compatibility Note** – The SPARC V8 LD instruction has been renamed LDUW in SPARC V9. The LDSW instruction is new in SPARC V9.

---

## Exceptions

*mem\_address\_not\_aligned* (all except LDSB, LDUB)

*data\_access\_exception*

*data\_access\_error*

*fast\_data\_access\_MMU\_miss*

*fast\_data\_access\_protection*

*VA\_watchpoint*

*PA\_watchpoint*

---

## A.28 Load Integer from Alternate Space

| Opcode                | op3     | Operation                                   |
|-----------------------|---------|---|
| LDSBA <sup>PASI</sup> | 01 1001 | Load Signed Byte from Alternate Space       |
| LDSHA <sup>PASI</sup> | 01 1010 | Load Signed Halfword from Alternate Space   |
| LDSWA <sup>PASI</sup> | 01 1000 | Load Signed Word from Alternate Space       |
| LDUBA <sup>PASI</sup> | 01 0001 | Load Unsigned Byte from Alternate Space     |
| LDUHA <sup>PASI</sup> | 01 0010 | Load Unsigned Halfword from Alternate Space |
| LDUWA <sup>PASI</sup> | 01 0000 | Load Unsigned Word from Alternate Space     |
| LDXA <sup>PASI</sup>  | 01 1011 | Load Extended Word from Alternate Space     |

### Format (3)

|          |       |       |          |     |         |     |
|----------|-------|-------|----------|-----|---------|-----|
| 11       | rd    | op3   | rs1      | i=0 | imm_asi | rs2 |
| 11       | rd    | op3   | rs1      | i=1 | simm13  |     |
| 31 30 29 | 25 24 | 19 18 | 14 13 12 | 5 4 | 0       |     |



| Assembly Language Syntax |  |                |
|--------------------------|--|----------------|
| ldsba                    | [regaddr] imm_asi, reg <sub>rd</sub>   |                |
| ldsha                    | [regaddr] imm_asi, reg <sub>rd</sub>   |                |
| ldswa                    | [regaddr] imm_asi, reg <sub>rd</sub>   |                |
| lduba                    | [regaddr] imm_asi, reg <sub>rd</sub>   |                |
| lduha                    | [regaddr] imm_asi, reg <sub>rd</sub>   |                |
| lduwa                    | [regaddr] imm_asi, reg <sub>rd</sub>   | (synonym: lda) |
| ldxa                     | [regaddr] imm_asi, reg <sub>rd</sub>   |                |
| ldsba                    | [reg_plus_imm] %asi, reg <sub>rd</sub> |                |
| ldsha                    | [reg_plus_imm] %asi, reg <sub>rd</sub> |                |
| ldswa                    | [reg_plus_imm] %asi, reg <sub>rd</sub> |                |
| lduba                    | [reg_plus_imm] %asi, reg <sub>rd</sub> |                |
| lduha                    | [reg_plus_imm] %asi, reg <sub>rd</sub> |                |
| lduwa                    | [reg_plus_imm] %asi, reg <sub>rd</sub> | (synonym: lda) |
| ldxa                     | [reg_plus_imm] %asi, reg <sub>rd</sub> |                |

### Description

The load integer from alternate space instructions copy a byte, halfword, word, or an extended word from memory. All copy the fetched value into  $r[rd]$ . A fetched byte, halfword, or word is right-justified in the destination register  $r[rd]$ ; it is either sign-extended or zero-filled on the left, depending on whether the opcode specifies a signed or unsigned operation, respectively.

The load integer from alternate space instructions contain the address space identifier (ASI) to be used for the load in the `imm_asi` field if  $i = 0$ , or in the ASI register if  $i = 1$ . The access is privileged if bit 7 of the ASI is zero; otherwise, it is not privileged. The effective address for these instructions is “ $r[rs1] + r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] + \text{sign\_ext}(\text{simml3})$ ” if  $i = 1$ .

A successful load (notably, load extended) instruction operates atomically.

LDUHA and LDSHA cause a *mem\_address\_not\_aligned* exception if the address is not halfword aligned. LDUWA and LDSWA cause a *mem\_address\_not\_aligned* exception if the effective address is not word aligned; LDXA causes a *mem\_address\_not\_aligned* exception if the address is not doubleword aligned.

These instructions cause a *privileged\_action* exception if `PSTATE.PRIV = 0` and bit 7 of the ASI is zero.

## Exceptions

*privileged\_action*

*mem\_address\_not\_aligned* (all except LDSBA and LDUBA)

*data\_access\_exception*

*PA\_watchpoint*

*VA\_watchpoint*

*fast\_data\_access\_MMU\_miss*

*fast\_data\_access\_protection*


*data\_access\_error*

---

## A.29 Load Quadword, Atomic (VIS I)

| Opcode | imm_asi                | ASI Value        | Operation                          |
|--------|------------------------|------------------|------------------------------------|
| LDDA   | ASI_NUCLEUS_QUAD_LDD   | 24 <sub>16</sub> | 128-bit atomic load                |
| LDDA   | ASI_NUCLEUS_QUAD_LDD_L | 2C <sub>16</sub> | 128-bit atomic load, little-endian |
| LDDA   | ASI_QUAD_LDD_PHYS      | 34 <sub>16</sub> | 128-bit atomic load                |
| LDDA   | ASI_QUAD_LDD_PHYS_L    | 3c <sub>16</sub> | 128-bit atomic load, little endian |

### Format (3) LDDA



| Assembly Language Syntax |                             |
|--------------------------|-----------------------------|
| ldda                     | [reg_addr] imm_asi, reg_rd  |
| ldda                     | [reg_plus_imm] %asi, reg_rd |

### Description

ASIs 24<sub>16</sub> and 2C<sub>16</sub> are used with the LDDA instruction to atomically read a 128-bit, virtually addressed data item. They are intended to be used by a TLB miss handler to access TSB entries without requiring locks. The data is placed in an even/odd pair of 64-bit registers. The lowest address 64 bits are placed in the even register; the highest-address 64 bits are placed

in the odd numbered register. The reference is made from the nucleus context. ASIs  $24_{16}$  and  $2C_{16}$  are translated by the MMU into physical addresses according to normal translation rules for the nucleus context.

To reduce the number of locked pages in D-TLB, a new ASI load instruction, atomic quad load physical ( $1d_{da}$  ASI\_QUAD\_LDD\_PHYS), was added. It allows a full TTE entry (128 bits, tag and data) in TSB to be read directly with PA, bypassing the VA-to-PA translation. In today's D-TLB miss handler, a TTE entry is read using two  $1d_x$  instructions. ASIs  $34_{16}$  and  $3C_{16}$  are not translated by the MMU and addresses provided are interpreted directly as physical addresses.

Since quad load with these ASIs bypasses the D-MMU, the physical address is set equal to the truncated virtual address, that is,  $PA[42:0]=VA[42:0]$ . Internally in hardware, the physical page attribute bits of these ASIs are hardcoded (not coming from DCU Control Register) as follows:

$CP = 1, CV = 0, IE = 0, E = 0, P = 0, W = 0, NFO = 0, Size = 8 K$

Note that  $(CP, CV) = 10$  means it is cacheable in L2-cache, W-cache, and P-cache, but not D-cache (since D-cache is VA-indexed). Therefore, this atomic quad load physical instruction can only be used with cacheable PA.

Semantically, ASI\_QUAD\_LDD\_PHYS is like a combination of ASI\_NUCLEUS\_QUAD\_LDD and ASI\_PHYS\_USE\_EC.

An *illegal\_instruction* occurs if an odd “rd” register number is used. If non-privileged software tries to use this ASI, a *privileged\_action* exception occurs. If the physical address of the data referenced matches the watchpoint register (ASI\_DMMU\_PA\_WATCHPOINT\_REG), the *PA\_watchpoint\_exception* occurs.

In addition to the usual traps for LDDA using a privileged ASI, a *data\_access\_exception* trap occurs for a non-cacheable access or if a quadword-load ASI is used with any instruction other than LDDA. A *mem\_address\_not\_aligned* trap is taken if the access is not aligned on a 128-byte boundary.

## Exceptions

### *privileged\_action*

*PA\_watchpoint* (recognized on only the first 8 bytes of an access)

*VA\_watchpoint* (recognized on only the first 8 bytes of an access)

*illegal\_instruction* (misaligned rd)

*mem\_address\_not\_aligned*

*data\_access\_exception* (an attempt to access a page marked as non-cacheable)

*data\_access\_error*


*fast\_data\_access\_MMU\_miss*

*fast\_data\_access\_protection*

## A.30 Load-Store Unsigned Byte

| Opcode | op3     | Operation                |
|--------|---------|--------------------------|
| LDSTUB | 00 1101 | Load-Store Unsigned Byte |

### Format (3)



| Assembly Language Syntax |                              |
|--------------------------|------------------------------|
| ldstub                   | [address], reg <sub>rd</sub> |

### Description

The load-store unsigned byte instruction copies a byte from memory into  $r[rd]$ , then rewrites the addressed byte in memory to all ones. The fetched byte is right-justified in the destination register  $r[rd]$  and zero-filled on the left.

The operation is performed atomically, that is, without allowing intervening interrupts or deferred traps. In a multiprocessor system, two or more processors executing LDSTUB, LDSTUBA, CASA, CASXA, SWAP, or SWAPA instructions addressing all or parts of the same doubleword simultaneously are guaranteed to execute them in an undefined, but serial order.

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

The coherence and atomicity of memory operations between processors and I/O DMA memory accesses is maintained for cacheable memory space.

### Exceptions

*data\_access\_exception*

*data\_access\_error*


*fast\_data\_access\_MMU\_miss*

*fast\_data\_access\_protection*  
*VA\_watchpoint*  
*PA\_watchpoint*

## A.31 Load-Store Unsigned Byte to Alternate Space

| Opcode                  | op3     | Operation                                     |
|-------------------------|---------|---|
| LDSTUBA <sup>PASI</sup> | 01 1101 | Load-Store Unsigned Byte into Alternate Space |

*Format (3)*



| Assembly Language Syntax |  |
|--------------------------|--|
| ldstuba                  | [regaddr] imm_asi, reg <sub>rd</sub>   |
| ldstuba                  | [reg_plus_imm] %asi, reg <sub>rd</sub> |

*Description*

The load-store unsigned byte into alternate space instruction copies a byte from memory into  $r[rd]$ , then rewrites the addressed byte in memory to all ones. The fetched byte is right-justified in the destination register  $r[rd]$  and zero-filled on the left.

The operation is performed atomically, that is, without allowing intervening interrupts or deferred traps. In a multiprocessor system, two or more processors executing LDSTUB, LDSTUBA, CASA, CASXA, SWAP, or SWAPA instructions addressing all or parts of the same doubleword simultaneously are guaranteed to execute them in an undefined, but serial order.

LDSTUBA contains the address space identifier (ASI) to be used for the load in the `imm_asi` field if `i = 0`, or in the ASI register if `i = 1`. The access is privileged if bit 7 of the ASI is zero; otherwise, it is not privileged. The effective address is “ $r[rs1] + r[rs2]$ ” if `i = 0`, or “ $r[rs1] + \text{sign\_ext}(simm13)$ ” if `i = 1`.

LDSTUBA causes a *privileged\_action* exception if `PSTATE.PRIV = 0` and bit 7 of the ASI is zero.

The coherence and atomicity of memory operations between processors and I/O DMA memory accesses is maintained for cacheable memory space.

### *Exceptions*

*privileged\_action*  
*data\_access\_exception*  
*data\_access\_error*  
*fast\_data\_access\_MMU\_miss*  
*fast\_data\_access\_protection*  
*VA\_watchpoint*  
*PA\_watchpoint*


---

## A.32 Logical Operate Instructions (VIS I)

| Opcode | opf         | Operation   |
|--------|-------------|---|
| FZERO  | 0 0110 0000 | Zero fill   |
| FZEROS | 0 0110 0001 | Zero fill, single-precision                                   |
| FONE   | 0 0111 1110 | One fill  |
| FONES  | 0 0111 1111 | One fill, single-precision                                    |
| FSRC1  | 0 0111 0100 | Copy <code>src1</code>  |
| FSRC1S | 0 0111 0101 | Copy <code>src1</code> , single-precision                     |
| FSRC2  | 0 0111 1000 | Copy <code>src2</code>  |
| FSRC2S | 0 0111 1001 | Copy <code>src2</code> , single-precision                     |
| FNOT1  | 0 0110 1010 | Negate (ones-complement) <code>src1</code>                    |
| FNOT1S | 0 0110 1011 | Negate (ones-complement) <code>src1</code> , single-precision |
| FNOT2  | 0 0110 0110 | Negate (ones-complement) <code>src2</code>                    |
| FNOT2S | 0 0110 0111 | Negate (ones-complement) <code>src2</code> , single-precision |
| FOR    | 0 0111 1100 | Logical OR  |
| FORS   | 0 0111 1101 | Logical OR, single-precision                                  |
| FNOR   | 0 0110 0010 | Logical NOR   |
| FNORS  | 0 0110 0011 | Logical NOR, single-precision                                 |
| FAND   | 0 0111 0000 | Logical AND   |

| Opcode    | opf         | Operation                               |
|-----------|-------------|---|
| FANDS     | 0 0111 0001 | Logical AND, single-precision           |
| FNAND     | 0 0110 1110 | Logical NAND                            |
| FNANDS    | 0 0110 1111 | Logical NAND, single-precision          |
| FXOR      | 0 0110 1100 | Logical XOR                             |
| FXORS     | 0 0110 1101 | Logical XOR, single-precision           |
| FXNOR     | 0 0111 0010 | Logical XNOR                            |
| FXNORS    | 0 0111 0011 | Logical XNOR, single-precision          |
| FORNOT1   | 0 0111 1010 | Negated src1 OR src2                    |
| FORNOT1S  | 0 0111 1011 | Negated src1 OR src2, single-precision  |
| FORNOT2   | 0 0111 0110 | src1 OR negated src2                    |
| FORNOT2S  | 0 0111 0111 | src1 OR negated src2, single-precision  |
| FANDNOT1  | 0 0110 1000 | Negated src1 AND src2                   |
| FANDNOT1S | 0 0110 1001 | Negated src1 AND src2, single-precision |
| FANDNOT2  | 0 0110 0100 | src1 AND negated src2                   |
| FANDNOT2S | 0 0110 0101 | src1 AND negated src2, single-precision |

### Format (3)



| Assembly Language Syntax |                         |
|--------------------------|-------------------------|
| fzero                    | $freq_{rd}$             |
| fzeros                   | $freq_{rd}$             |
| fone                     | $freq_{rd}$             |
| fores                    | $freq_{rd}$             |
| fsrc1                    | $freq_{rs1}, freq_{rd}$ |
| fsrc1s                   | $freq_{rs1}, freq_{rd}$ |
| fsrc2                    | $freq_{rs2}, freq_{rd}$ |
| fsrc2s                   | $freq_{rs2}, freq_{rd}$ |
| fnot1                    | $freq_{rs1}, freq_{rd}$ |
| fnot1s                   | $freq_{rs1}, freq_{rd}$ |
| fnot2                    | $freq_{rs2}, freq_{rd}$ |

| Assembly Language Syntax |  |
|--------------------------|--|
| <code>fnot2s</code>      | <code>freg<sub>rs2</sub>, freg<sub>rd</sub></code>                     |
| <code>for</code>         | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fors</code>        | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fnor</code>        | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fnors</code>       | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fand</code>        | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fand</code>        | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fnands</code>      | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fnands</code>      | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fxor</code>        | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fxors</code>       | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fxnor</code>       | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fxnors</code>      | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fornot1</code>     | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fornot1s</code>    | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fornot2</code>     | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fornot2s</code>    | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fandnot1</code>    | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fandnot1s</code>   | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fandnot2</code>    | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |
| <code>fandnot2s</code>   | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |

### Description

The standard 64-bit versions of these instructions perform 1 of 16 64-bit logical operations between the 64-bit floating-point registers specified by `rs1` and `rs2`. The result is stored in the 64-bit floating-point destination register specified by `rd`. The 32-bit (single-precision) version of these instructions perform 32-bit logical operations.

---

**Note** – For good performance, the result of a single logical should not be used as part of a 64-bit graphics instruction source operand in the next three instruction groups. Similarly, the result of a standard logical should not be used as a 32-bit graphics instruction source operand in the next three instruction groups.

---

### Exceptions


*fp\_disabled*



## A.33 Logical Operations

| Opcode             | op3     | Operation                        |
|--------------------|---------|----------------------------------|
| AND                | 00 0001 | AND                              |
| AND <sub>cc</sub>  | 01 0001 | AND and modify cc's              |
| ANDN               | 00 0101 | AND Not                          |
| ANDN <sub>cc</sub> | 01 0101 | AND Not and modify cc's          |
| OR                 | 00 0010 | Inclusive OR                     |
| OR <sub>cc</sub>   | 01 0010 | Inclusive OR and modify cc's     |
| ORN                | 00 0110 | Inclusive OR Not                 |
| ORN <sub>cc</sub>  | 01 0110 | Inclusive OR Not and modify cc's |
| XOR                | 00 0011 | Exclusive OR                     |
| XOR <sub>cc</sub>  | 01 0011 | Exclusive OR and modify cc's     |
| XNOR               | 00 0111 | Exclusive NOR                    |
| XNOR <sub>cc</sub> | 01 0111 | Exclusive NOR and modify cc's    |

*Format (3)*



| Assembly Language Syntax |  |
|--------------------------|--|
| and                      | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| andcc                    | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| andn                     | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| andncc                   | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| or                       | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| orcc                     | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| orn                      | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| orncc                    | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| xor                      | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| xorcc                    | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| xnor                     | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| xnorcc                   | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |

### Description

These instructions implement bitwise logical operations. They compute “*r[rs1] op r[rs2]*” if *i* = 0, or “*r[rs1] op sign\_ext(simml3)*” if *i* = 1, and write the result into *r[rd]*.

ANDCC, ANDNCC, ORCC, ORNCC, XORCC, and XNORCC modify the integer condition codes (*icc* and *xcc*). They set the condition codes as follows:

- *icc.v*, *icc.c*, *xcc.v*, and *xcc.c* to zero
  - *icc.n* to bit 31 of the result
  - *xcc.n* to bit 63 of the result
  - *icc.z* to one if bits 31:0 of the result are zero (otherwise to zero)
  - *xcc.z* to one if all 64 bits of the result are zero (otherwise to zero)

ANDN, ANDNCC, ORN, and ORNCC logically negate their second operand before applying the main (AND or OR) operation.

---

**Programming Note** – XNOR and XNORCC are identical to the XOR-Not and XOR-Not-cc logical operations, respectively.

---


### Exceptions

None

## A.34 Memory Barrier

| Opcode | op3     | Operation      |
|--------|---------|----------------|
| MEMBAR | 10 1000 | Memory Barrier |

*Format (3)*



| Assembly Language Syntax |                    |
|--------------------------|--------------------|
| MEMBAR                   | <i>membar_mask</i> |

*Description*

The memory barrier instruction, MEMBAR, has two complementary functions: to express order constraints between memory references and to provide explicit control of memory-reference completion. The *membar\_mask* field in the suggested assembly language is the concatenation of the *cmask* and *mmask* instruction fields.

MEMBAR introduces an order constraint between classes of memory references appearing before the MEMBAR and memory references following it in a program. The particular classes of memory references are specified by the *mmask* field. Memory references are classified as loads (including load instructions LDSTUB(A), SWAP(A), CASA, and CASXA) and stores (including store instructions LDSTUB(A), SWAP(A), CASA, CASXA, and FLUSH). The *mmask* field specifies the classes of memory references subject to ordering, as described. MEMBAR applies to all memory operations in all address spaces referenced by the issuing processor, but it has no effect on memory references by other processors. When the *cmask* field is nonzero, completion as well as order constraints are imposed, and the order imposed can be more stringent than that specifiable by the *mmask* field alone.

A load has been performed when the value loaded has been transmitted from memory and cannot be modified by another processor. A store has been performed when the value stored has become visible, that is, when the previous value can no longer be read by any processor. In specifying the effect of MEMBAR, instructions are considered to be executed as if they were processed in a strictly sequential fashion, with each instruction completed before the next has begun.

The `mmask` field is encoded in bits 3 through 0 of the instruction. TABLE A-9 specifies the order constraint that each bit of `mmask` (selected when set to one) imposes on memory references appearing before and after the MEMBAR. From zero to four, mask bits may be selected in the `mmask` field.

**TABLE A-9** MEMBAR `mmask` Encodings

| Mask Bit                    | Name                     | Description  |
|-----------------------------|--------------------------|--|
| <code>mmask&lt;3&gt;</code> | <code>#StoreStore</code> | The effects of all stores appearing prior to the MEMBAR instruction must be visible to all processors before the effect of any stores following the MEMBAR; equivalent to the deprecated <code>STBAR</code> instruction. |
| <code>mmask&lt;2&gt;</code> | <code>#LoadStore</code>  | All loads appearing prior to the MEMBAR instruction must have been performed before the effects of any stores following the MEMBAR are visible to any other processor.   |
| <code>mmask&lt;1&gt;</code> | <code>#StoreLoad</code>  | The effects of all stores appearing prior to the MEMBAR instruction must be visible to all processors before loads following the MEMBAR may be performed.  |
| <code>mmask&lt;0&gt;</code> | <code>#LoadLoad</code>   | All loads appearing prior to the MEMBAR instruction must have been performed before any loads following the MEMBAR may be performed.   |

The `cmask` field is encoded in bits 6 through 4 of the instruction. Bits in the `cmask` field, described in TABLE A-10, specify additional constraints on the order of memory references and the processing of instructions. If `cmask` is zero, then MEMBAR enforces the partial ordering specified by the `mmask` field; if `cmask` is nonzero, then completion and partial order constraints are applied.

**TABLE A-10** MEMBAR `cmask` Encodings

| Mask Bit              | Function                | Name                    | Description  |
|-----------------------|-------------------------|-------------------------|--|
| <code>cmask[2]</code> | Synchronization barrier | <code>#Sync</code>      | All operations (including non-memory reference operations) appearing prior to the MEMBAR must have been performed and the effects of any exceptions be visible before any instruction after the MEMBAR may be initiated. |
| <code>cmask[1]</code> | Memory issue barrier    | <code>#MemIssue</code>  | All memory reference operations appearing prior to the MEMBAR must have been performed before any memory operation after the MEMBAR may be initiated.  |
| <code>cmask[0]</code> | Lookaside barrier       | <code>#Lookaside</code> | A store appearing prior to the MEMBAR must complete before any load following the MEMBAR referencing the same address can be initiated.  |

The encoding of MEMBAR is identical to that of the `RDASR` instruction, except that `rs1 = 15`, `rd = 0`, and `i = 1`.

The coherence and atomicity of memory operations between processors and I/O DMA memory accesses is maintained for cacheable memory space.

---

**Compatibility Note** – MEMBAR with  $m_{mask} = 8_{16}$  and  $c_{mask} = 0_{16}$  (“MEMBAR #StoreStore”) is identical in function to the SPARC V8 STBAR instruction, which is deprecated.

---

The information included in this section should not be used for the decision as to when MEMBARs should be added to software that needs to be compliant across all UltraSPARC-based platforms. The operations of block load/block store (BLD/BST) on an UltraSPARC III Cu processor are generally more ordered with respect to other operations, compared to UltraSPARC I and UltraSPARC II processors. Code written and found to “work” on the UltraSPARC III Cu processor may not work on UltraSPARC I and UltraSPARC II processors if it does not follow the rules for BLD/BST specified for those processors. Code that happens to work on UltraSPARC I and UltraSPARC II processors may not work on the UltraSPARC III Cu processor if it did not meet the coding guidelines specified for those processors. In no case is the coding requirement for the UltraSPARC III Cu processor more restrictive than that of the UltraSPARC I and UltraSPARC II processors.

Software developers should not use the information in this section for determining the need for MEMBARs but instead should rely on the SPARC V9 MEMBAR rules. These UltraSPARC III Cu rules are less restrictive than SPARC V9, UltraSPARC I, and UltraSPARC II rules and are never more restrictive.

### MEMBAR Rules

The UltraSPARC III Cu hardware uses the following rules to guide the interlock implementation.

1. Non-cacheable load or store with side-effect bit on will always be blocked.
2. Cacheable or non-cacheable BLD will not be blocked.
3.  $VA<12:5>$  of a load (cacheable or non-cacheable) will be compared with the  $VA<12:5>$  of all entries in store queue. When a matching is detected, this load (cacheable or non-cacheable) will be blocked.
4. An insertion of MEMBAR is required if strong ordering is desired while not fitting rules 1 to 3.

TABLE A-11 and TABLE A-12 reflect the hardware interlocking mechanism implemented in the UltraSPARC III Cu processor. The tables are read from Row to Column, the first memory operation in program order being in Row followed by the memory operation found in Column. The following two symbols are used as table entries:

- # — No intervening operation required because Fireplane-compliant systems automatically order R before C.
- M — MEMBAR #Sync or MEMBAR #MemIssue or MEMBAR #StoreLoad required.

For VA<12:5> of a column operation not matching with VA<2:5> of a row operation while a strong ordering is desired, the MEMBAR rules summarized in TABLE A-11 reflect UltraSPARC III Cu's hardware implementation.

**TABLE A-11** MEMBAR Rules for Column VA <12:5> ≠ Row VA <12:5> While Desiring Strong Ordering

| From Row Operation R:  | To Column Operation C: |                        |       |                       |        |           |            |            |             |       |        |               |          |           |
|------------------------|------------------------|------------------------|-------|-----------------------|--------|-----------|------------|------------|-------------|-------|--------|---------------|----------|-----------|
|                        | load                   | load from internal ASI | store | store to internal ASI | atomic | load_nc_e | store_nc_e | load_nc_ne | store_nc_ne | blood | bstore | bstore_commit | blood_nc | bstore_nc |
| load                   | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M     | M      | #             | M        | M         |
| load from internal ASI | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | #        | #         |
| store                  | M                      | #                      | #     | #                     | #      | M         | #          | M          | #           | M     | M      | #             | M        | M         |
| store to internal ASI  | #                      | M                      | #     | #                     | #      | #         | #          | #          | #           | M     | #      | #             | M        | M         |
| atomic                 | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M     | M      | #             | M        | M         |
| load_nc_e              | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M     | M      | #             | M        | M         |
| store_nc_e             | M                      | #                      | #     | #                     | #      | #         | #          | M          | #           | M     | M      | #             | M        | M         |
| load_nc_ne             | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M     | M      | #             | M        | M         |
| store_nc_ne            | M                      | #                      | #     | #                     | #      | M         | #          | M          | #           | M     | M      | #             | M        | M         |
| blood                  | M                      | #                      | M     | #                     | M      | M         | M          | M          | M           | M     | M      | #             | M        | M         |
| bstore                 | M                      | #                      | M     | #                     | M      | M         | M          | M          | M           | M     | M      | #             | M        | M         |
| bstore_commit          | M                      | #                      | M     | #                     | M      | M         | M          | M          | M           | M     | M      | #             | M        | M         |
| blood_nc               | M                      | #                      | M     | #                     | M      | M         | M          | M          | M           | M     | M      | #             | M        | M         |
| bstore_nc              | M                      | #                      | M     | #                     | M      | M         | M          | M          | M           | M     | M      | #             | M        | M         |

When VA<12:5> of a column operation matches VA<12:5> of a row operation, the MEMBAR rules summarized in TABLE A-12 reflect the UltraSPARC III Cu processor's hardware implementation.

**TABLE A-12** MEMBAR Rules for Column VA<12:5> = Row VA<12:5> While Desiring Strong Ordering

| From Row<br>Operation R: | To Column Operation C: |                        |       |                       |        |           |            |            |             |       |        |               |          |           |
|--------------------------|------------------------|------------------------|-------|-----------------------|--------|-----------|------------|------------|-------------|-------|--------|---------------|----------|-----------|
|                          | load                   | load from internal ASI | store | store to internal ASI | atomic | load_nc_e | store_nc_e | load_nc_ne | store_nc_ne | bload | bstore | bstore_commit | bload_nc | bstore_nc |
| load                     | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | #        | #         |
| load from internal ASI   | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | #        | #         |
| store                    | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M     | #      | #             | #        | #         |
| store to internal ASI    | #                      | M                      | #     | #                     | #      | #         | #          | #          | #           | M     | #      | #             | M        | M         |
| atomic                   | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | #        | #         |
| load_nc_e                | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | #        | #         |
| store_nc_e               | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M     | #      | #             | M        | #         |
| load_nc_ne               | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | #        | #         |
| store_nc_ne              | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M     | #      | #             | M        | #         |
| bload                    | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | #        | #         |
| bstore                   | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | M     | #      | #             | #        | #         |
| bstore_commit            | M                      | #                      | M     | #                     | M      | M         | M          | M          | M           | M     | M      | #             | M        | M         |
| bload_nc                 | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | #        | #         |
| bstore_nc                | #                      | #                      | #     | #                     | #      | #         | #          | #          | #           | #     | #      | #             | M        | #         |

### Special Rules for Quad LDD (ASI 24<sub>16</sub> and ASI 2C<sub>16</sub>)

MEMBAR is only required before quad LDD if VA<12:5> of a preceding store to the same address space matches VA<12:5> of the quad LDD.

#### *Exceptions*

None

## A.35 Move Floating-Point Register on Condition (FMOVcc)

*For Integer Condition Codes*


| Opcode  | op3     | cond | Operation  | icc/xcc Test                              |
|---------|---------|------|--|---|
| FMOVA   | 11 0101 | 1000 | Move Always                                      | 1   |
| FMOVN   | 11 0101 | 0000 | Move Never                                       | 0   |
| FMOVNE  | 11 0101 | 1001 | Move if Not Equal                                | <b>not</b> Z                              |
| FMOVE   | 11 0101 | 0001 | Move if Equal                                    | Z   |
| FMOVG   | 11 0101 | 1010 | Move if Greater                                  | <b>not</b> (Z <b>or</b> (N <b>xor</b> V)) |
| FMOVLE  | 11 0101 | 0010 | Move if Less or Equal                            | Z <b>or</b> (N <b>xor</b> V)              |
| FMOVGE  | 11 0101 | 1011 | Move if Greater or Equal                         | <b>not</b> (N <b>xor</b> V)               |
| FMOVL   | 11 0101 | 0011 | Move if Less                                     | N <b>xor</b> V                            |
| FMOVGU  | 11 0101 | 1100 | Move if Greater Unsigned                         | <b>not</b> (C <b>or</b> Z)                |
| FMOVLEU | 11 0101 | 0100 | Move if Less or Equal Unsigned                   | (C <b>or</b> Z)                           |
| FMOVCC  | 11 0101 | 1101 | Move if Carry Clear (Greater or Equal, Unsigned) | <b>not</b> C                              |
| FMOVCS  | 11 0101 | 0101 | Move if Carry Set (Less than, Unsigned)          | C   |
| FMOVPOS | 11 0101 | 1110 | Move if Positive                                 | <b>not</b> N                              |
| FMOVNEG | 11 0101 | 0110 | Move if Negative                                 | N   |
| FMOVVC  | 11 0101 | 1111 | Move if Overflow Clear                           | <b>not</b> V                              |
| FMOVVS  | 11 0101 | 0111 | Move if Overflow Set                             | V   |



*For Floating-Point Condition Codes*

| Opcode   | op3     | cond | Operation                             | fcc Test    |
|----------|---------|------|---------------------------------------|-------------|
| FMOVFA   | 11 0101 | 1000 | Move Always                           | 1           |
| FMOVFN   | 11 0101 | 0000 | Move Never                            | 0           |
| FMOVFU   | 11 0101 | 0111 | Move if Unordered                     | U           |
| FMOVFG   | 11 0101 | 0110 | Move if Greater                       | G           |
| FMOVFUG  | 11 0101 | 0101 | Move if Unordered or Greater          | G or U      |
| FMOVFL   | 11 0101 | 0100 | Move if Less                          | L           |
| FMOVFUL  | 11 0101 | 0011 | Move if Unordered or Less             | L or U      |
| FMOVFLG  | 11 0101 | 0010 | Move if Less or Greater               | L or G      |
| FMOVFNE  | 11 0101 | 0001 | Move if Not Equal                     | L or G or U |
| FMOVFE   | 11 0101 | 1001 | Move if Equal                         | E           |
| FMOVFUE  | 11 0101 | 1010 | Move if Unordered or Equal            | E or U      |
| FMOVFGE  | 11 0101 | 1011 | Move if Greater or Equal              | E or G      |
| FMOVFUGE | 11 0101 | 1100 | Move if Unordered or Greater or Equal | E or G or U |
| FMOVFLE  | 11 0101 | 1101 | Move if Less or Equal                 | E or L      |
| FMOVFULE | 11 0101 | 1110 | Move if Unordered or Less or Equal    | E or L or U |
| FMOVFO   | 11 0101 | 1111 | Move if Ordered                       | E or L or G |

*Format (4)*



## Encoding of the *opf\_cc* Field

| <b>opf_cc</b> | <b>Condition Code</b> |
|---------------|-----------------------|
| 000           | <i>fcc0</i>           |
| 001           | <i>fcc1</i>           |
| 010           | <i>fcc2</i>           |
| 011           | <i>fcc3</i>           |
| 100           | <i>icc</i>            |
| 101           | —                     |
| 110           | <i>xcc</i>            |
| 111           | —                     |

## Encoding of *opf* Field (*opf\_cc* □ *opf\_low*)

| <b>Instruction Variation</b> |                     | <b>opf_cc</b> | <b>opf_low</b> | <b>opf</b>          |
|------------------------------|---------------------|---------------|----------------|---------------------|
| FMOVScC                      | <i>%fccn,rs2,rd</i> | <i>0nn</i>    | 00 0001        | 0 <i>nn</i> 00 0001 |
| FMOVDcC                      | <i>%fccn,rs2,rd</i> | <i>0nn</i>    | 00 0010        | 0 <i>nn</i> 00 0010 |
| FMOVQcC                      | <i>%fccn,rs2,rd</i> | <i>0nn</i>    | 00 0011        | 0 <i>nn</i> 00 0011 |
| FMOVScC                      | <i>%icc,rs2,rd</i>  | 100           | 00 0001        | 1 0000 0001         |
| FMOVDcC                      | <i>%icc,rs2,rd</i>  | 100           | 00 0010        | 1 0000 0010         |
| FMOVQcC                      | <i>%icc,rs2,rd</i>  | 100           | 00 0011        | 1 0000 0011         |
| FMOVScC                      | <i>%xcc,rs2,rd</i>  | 110           | 00 0001        | 1 1000 0001         |
| FMOVDcC                      | <i>%xcc,rs2,rd</i>  | 110           | 00 0010        | 1 1000 0010         |
| FMOVQcC                      | <i>%xcc,rs2,rd</i>  | 110           | 00 0011        | 1 1000 0011         |

## For Integer Condition Codes

| Assembly Language Syntax    |   |  |
|-----------------------------|---|--|
| <code>fmov{s,d,q}a</code>   | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |  |
| <code>fmov{s,d,q}n</code>   | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |  |
| <code>fmov{s,d,q}ne</code>  | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> | (synonyms: <code>fmov{s,d,q}nz</code> )  |
| <code>fmov{s,d,q}e</code>   | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> | (synonyms: <code>fmov{s,d,q}z</code> )   |
| <code>fmov{s,d,q}g</code>   | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |  |
| <code>fmov{s,d,q}le</code>  | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |  |
| <code>fmov{s,d,q}ge</code>  | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |  |
| <code>fmov{s,d,q}l</code>   | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |  |
| <code>fmov{s,d,q}gu</code>  | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |  |
| <code>fmov{s,d,q}leu</code> | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |  |
| <code>fmov{s,d,q}cc</code>  | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> | (synonyms: <code>fmov{s,d,q}geu</code> ) |
| <code>fmov{s,d,q}cs</code>  | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> | (synonyms: <code>fmov{s,d,q}lu</code> )  |
| <code>fmov{s,d,q}pos</code> | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |  |
| <code>fmov{s,d,q}neg</code> | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |  |
| <code>fmov{s,d,q}vc</code>  | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |  |
| <code>fmov{s,d,q}vs</code>  | <code>i_or_x_cc, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |  |

---

**Programming Note** – To select the appropriate condition code, include `%icc` or `%xcc` before the registers.

---

## For Floating-Point Condition Codes

| Assembly Language Syntax    |   |   |
|-----------------------------|---|---|
| <code>fmov{s,d,q}a</code>   | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |   |
| <code>fmov{s,d,q}n</code>   | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |   |
| <code>fmov{s,d,q}u</code>   | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |   |
| <code>fmov{s,d,q}g</code>   | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |   |
| <code>fmov{s,d,q}ug</code>  | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |   |
| <code>fmov{s,d,q}l</code>   | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |   |
| <code>fmov{s,d,q}ul</code>  | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |   |
| <code>fmov{s,d,q}lg</code>  | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |   |
| <code>fmov{s,d,q}ne</code>  | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> | (synonyms: <code>fmov{s,d,q}nz</code> ) |
| <code>fmov{s,d,q}e</code>   | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> | (synonyms: <code>fmov{s,d,q}z</code> )  |
| <code>fmov{s,d,q}ue</code>  | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |   |
| <code>fmov{s,d,q}ge</code>  | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |   |
| <code>fmov{s,d,q}uge</code> | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |   |
| <code>fmov{s,d,q}le</code>  | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |   |
| <code>fmov{s,d,q}ule</code> | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |   |
| <code>fmov{s,d,q}o</code>   | <code>%fccn, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |   |

### Description

These instructions copy the floating-point register(s) specified by `rs2` to the floating-point register(s) specified by `rd` if the condition indicated by the `cond` field is satisfied by the selected condition code. The condition code used is specified by the `opf_cc` field of the instruction. If the condition is FALSE, then the destination register(s) are not changed.

These instructions do not modify any condition codes.

---

**Programming Note** – In general, branches cause the processor’s performance to degrade. Frequently, the `MOVcc` and `FMOVcc` instructions can be used to avoid branches. For example, the following C language segment:

---

```
double A, B, X;
if (A > B) then X = 1.03; else X = 0.0;
```

can be coded as

```
! assume A is in %f0; B is in %f2; %xx points to constant area
ldd      [%xx+C_1.03],%f4    ! X = 1.03
fcmpd    %fcc3,%f0,%f2      ! A > B
```

```

fble ,a %fcc3,label
! following only executed if the branch is taken
fsubd %f4,%f4,%f4 ! X = 0.0
label:...

```

This code takes four instructions including a branch.

With `FMOVCC`, this could be coded as

```

ldd [%xx+C_1.03],%f4 ! X = 1.03
fsubd %f4,%f4,%f6 ! X' = 0.0
fcmpd %fcc3,%f0,%f2 ! A > B
fmovdle %fcc3,%f6,%f4 ! X = 0.0

```

This code also takes four instructions but requires no branches and may boost performance significantly. Use `MOVCC` and `FMOVCC` instead of branches wherever these instructions would improve performance.

### Exceptions

*fp\_disabled*

*fp\_exception\_other* (`ftt = unimplemented_FPop` (`opf_cc = 1012` or `1112` and quad forms))

## A.36 Move Floating-Point Register on Integer Register Condition (FMOVr)

| Opcode   | op3     | rcond | Operation                                      | Test            |
|----------|---------|-------|--|-----------------|
| —        | 11 0101 | 000   | <i>Reserved</i>                                | —               |
| FMOVRZ   | 11 0101 | 001   | Move if Register Zero                          | $r[rs1] = 0$    |
| FMOVRLEZ | 11 0101 | 010   | Move if Register Less Than or Equal to Zero    | $r[rs1] \leq 0$ |
| FMOVRLZ  | 11 0101 | 011   | Move if Register Less Than Zero                | $r[rs1] < 0$    |
| —        | 11 0101 | 100   | <i>Reserved</i>                                | —               |
| FMOVRNZ  | 11 0101 | 101   | Move if Register Not Zero                      | $r[rs1] \neq 0$ |
| FMOVRGZ  | 11 0101 | 110   | Move if Register Greater Than Zero             | $r[rs1] > 0$    |
| FMOVRGEZ | 11 0101 | 111   | Move if Register Greater Than or Equal to Zero | $r[rs1] \geq 0$ |

## Format (4)

|          |       |       |          |      |       |         |     |
|----------|-------|-------|----------|------|-------|---------|-----|
| 10       | rd    | op3   | rs1      | 0    | rcond | opf_low | rs2 |
| 31 30 29 | 25 24 | 19 18 | 14 13 12 | 10 9 | 5 4   | 0       |     |

### Encoding of opf\_low Field

| Instruction variation   | opf_low |
|-------------------------|---------|
| FMOVSRcond rs1, rs2, rd | 0 0101  |
| FMOVDRcond rs1, rs2, rd | 0 0110  |
| FMOVQRcond rs1, rs2, rd | 0 0111  |

| Assembly Language Syntax |   |                             |
|--------------------------|---|-----------------------------|
| fmovr{s, d, q}e          | reg <sub>rs1</sub> .freg <sub>rs2</sub> .freg <sub>rd</sub> | (synonym: fmovr{s, d, q}z)  |
| fmovr{s, d, q}lez        | reg <sub>rs1</sub> .freg <sub>rs2</sub> .freg <sub>rd</sub> |                             |
| fmovr{s, d, q}lze        | reg <sub>rs1</sub> .freg <sub>rs2</sub> .freg <sub>rd</sub> |                             |
| fmovr{s, d, q}ne         | reg <sub>rs1</sub> .freg <sub>rs2</sub> .freg <sub>rd</sub> | (synonym: fmovr{s, d, q}nz) |
| fmovr{s, d, q}gze        | reg <sub>rs1</sub> .freg <sub>rs2</sub> .freg <sub>rd</sub> |                             |
| fmovr{s, d, q}gez        | reg <sub>rs1</sub> .freg <sub>rs2</sub> .freg <sub>rd</sub> |                             |

### Description

If the contents of integer register  $r[rs1]$  satisfy the condition specified in the  $rcond$  field, these instructions copy the contents of the floating-point register(s) specified by the  $rs2$  field to the floating-point register(s) specified by the  $rd$  field. If the contents of  $r[rs1]$  do not satisfy the condition, the floating-point register(s) specified by the  $rd$  field are not modified.

These instructions treat the integer register contents as a signed integer value; they do not modify any condition codes.

---

**Implementation Note** – The UltraSPARC III Cu processor does not implement this instruction by tagging each register value; it looks at the full 64-bit register to determine a negative or zero.

---

## Exceptions

*fp\_disabled*

*fp\_exception\_other* (unimplemented\_FPop (xcond = 000<sub>2</sub> or 100<sub>2</sub> and quad forms))

---

# A.37 Move Integer Register on Condition (MOVcc)

*For Integer Condition Codes*

| Opcode | op3     | cond | Operation  | icc/xcc Test                              |
|--------|---------|------|--|---|
| MOVA   | 10 1100 | 1000 | Move Always                                      | 1   |
| MOVN   | 10 1100 | 0000 | Move Never                                       | 0   |
| MOVNE  | 10 1100 | 1001 | Move if Not Equal                                | <b>not</b> Z                              |
| MOVE   | 10 1100 | 0001 | Move if Equal                                    | Z   |
| MOVG   | 10 1100 | 1010 | Move if Greater                                  | <b>not</b> (Z <b>or</b> (N <b>xor</b> V)) |
| MOVLE  | 10 1100 | 0010 | Move if Less or Equal                            | Z <b>or</b> (N <b>xor</b> V)              |
| MOVGE  | 10 1100 | 1011 | Move if Greater or Equal                         | <b>not</b> (N <b>xor</b> V)               |
| MOVL   | 10 1100 | 0011 | Move if Less                                     | N <b>xor</b> V                            |
| MOVGU  | 10 1100 | 1100 | Move if Greater Unsigned                         | <b>not</b> (C <b>or</b> Z)                |
| MOVLEU | 10 1100 | 0100 | Move if Less or Equal Unsigned                   | (C <b>or</b> Z)                           |
| MOVCC  | 10 1100 | 1101 | Move if Carry Clear (Greater or Equal, Unsigned) | <b>not</b> C                              |
| MOVCS  | 10 1100 | 0101 | Move if Carry Set (Less than, Unsigned)          | C   |
| MOVPOS | 10 1100 | 1110 | Move if Positive                                 | <b>not</b> N                              |
| MOVNEG | 10 1100 | 0110 | Move if Negative                                 | N   |
| MOVVC  | 10 1100 | 1111 | Move if Overflow Clear                           | <b>not</b> V                              |
| MOVVS  | 10 1100 | 0111 | Move if Overflow Set                             | V   |





| cc2 | cc1 | cc0 | Condition Code  |
|-----|-----|-----|-----------------|
| 0   | 0   | 0   | fcc0            |
| 0   | 0   | 1   | fcc1            |
| 0   | 1   | 0   | fcc2            |
| 0   | 1   | 1   | fcc3            |
| 1   | 0   | 0   | icc             |
| 1   | 0   | 1   | <i>Reserved</i> |
| 1   | 1   | 0   | xcc             |
| 1   | 1   | 1   | <i>Reserved</i> |

### *For Integer Condition Codes*

| Assembly Language Syntax |  |                                     |
|--------------------------|--|-------------------------------------|
| mov <sub>a</sub>         | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> |                                     |
| mov <sub>n</sub>         | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> |                                     |
| mov <sub>ne</sub>        | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> | <i>(synonym: mov<sub>nz</sub>)</i>  |
| mov <sub>e</sub>         | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> | <i>(synonym: mov<sub>z</sub>)</i>   |
| mov <sub>g</sub>         | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> |                                     |
| mov <sub>le</sub>        | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> |                                     |
| mov <sub>ge</sub>        | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> |                                     |
| mov <sub>l</sub>         | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> |                                     |
| mov <sub>gu</sub>        | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> |                                     |
| mov <sub>leu</sub>       | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> |                                     |
| mov <sub>cc</sub>        | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> | <i>(synonym: mov<sub>geu</sub>)</i> |
| mov <sub>cs</sub>        | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> | <i>(synonym: mov<sub>lu</sub>)</i>  |
| mov <sub>pos</sub>       | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> |                                     |
| mov <sub>neg</sub>       | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> |                                     |
| mov <sub>vc</sub>        | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> |                                     |
| mov <sub>vs</sub>        | <i>i_or_x_cc, reg_or_imm11, reg<sub>rd</sub></i> |                                     |

---

**Programming Note** – To select the appropriate condition code, include `%icc` or `%xcc` before the register or immediate field.

---

### *For Floating-Point Condition Codes*

| <b>Assembly Language Syntax</b> |  |                           |
|---------------------------------|--|---------------------------|
| <code>movb</code>               | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> |                           |
| <code>movn</code>               | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> |                           |
| <code>movu</code>               | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> |                           |
| <code>movg</code>               | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> |                           |
| <code>movug</code>              | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> |                           |
| <code>movl</code>               | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> |                           |
| <code>movul</code>              | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> |                           |
| <code>movlg</code>              | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> |                           |
| <code>movne</code>              | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> | ( <i>synonym: movnz</i> ) |
| <code>move</code>               | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> | ( <i>synonym: movz</i> )  |
| <code>movue</code>              | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> |                           |
| <code>movge</code>              | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> |                           |
| <code>movuge</code>             | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> |                           |
| <code>movle</code>              | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> |                           |
| <code>movule</code>             | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> |                           |
| <code>movo</code>               | <code>%fccn, reg_or_imm11, reg<sub>rd</sub></code> |                           |

---

**Programming Note** – To select the appropriate condition code, include `%fcc0`, `%fcc1`, `%fcc2`, or `%fcc3` before the register or immediate field.

---

### *Description*

These instructions test to see if `cond` is TRUE for the selected condition codes. If so, they copy the value in `r[rs2]` if `i` field = 0, or “`sign_ext(sim11)`” if `i` = 1 into `r[rd]`. The condition code used is specified by the `cc2`, `cc1`, and `cc0` fields of the instruction. If the condition is FALSE, then `r[rd]` is not changed.

These instructions copy an integer register to another integer register if the condition is TRUE. The condition code that is used to determine whether the move will occur can be either integer condition code (`icc` or `xcc`) or any floating-point condition code (`fcc0`, `fcc1`, `fcc2`, or `fcc3`).

These instructions do not modify any condition codes.

---

**Programming Note** – In general, branches cause the processor performance to degrade. Frequently, the `MOVCC` and `FMOVCC` instructions can be used to avoid branches. For example, consider the C language if-then-else statement:

---

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

can be coded as

```
    cmp        %i0,%i2
    bg,a       %xcc,label
    or         %g0,1,%i3          ! X = 1
    or         %g0,0,%i3          ! X = 0
label:...
```

This takes four instructions including a branch. With `MOVCC`, this could be coded as

```
    cmp        %i0,%i2
    or         %g0,1,%i3          ! assume X = 1
    movle     %xcc,0,%i3          ! overwrite with X = 0
```

This approach takes only three instructions and no branches and may boost performance significantly. Use `MOVCC` and `FMOVCC` instead of branches wherever these instructions would increase performance.

---

### *Exceptions*

*illegal\_instruction* (cc2  $\square$  cc1  $\square$  cc0 = 101<sub>2</sub> or 111<sub>2</sub>)  
*fp\_disabled* (cc2  $\square$  cc1  $\square$  cc0 = 000<sub>2</sub>, 001<sub>2</sub>, 010<sub>2</sub>, or 011<sub>2</sub> and the FPU is disabled)

## A.38 Move Integer Register on Register Condition (MOVr)

| Opcode  | op3     | rcond | Operation                                      | Test            |
|---------|---------|-------|--|-----------------|
| —       | 10 1111 | 000   | <i>Reserved</i>                                | —               |
| MOVrZ   | 10 1111 | 001   | Move if Register Zero                          | $r[rs1] = 0$    |
| MOVrLEZ | 10 1111 | 010   | Move if Register Less Than or Equal to Zero    | $r[rs1] \leq 0$ |
| MOVrLZ  | 10 1111 | 011   | Move if Register Less Than Zero                | $r[rs1] < 0$    |
| —       | 10 1111 | 100   | <i>Reserved</i>                                | —               |
| MOVrNZ  | 10 1111 | 101   | Move if Register Not Zero                      | $r[rs1] \neq 0$ |
| MOVrGZ  | 10 1111 | 110   | Move if Register Greater Than Zero             | $r[rs1] > 0$    |
| MOVrGEZ | 10 1111 | 111   | Move if Register Greater Than or Equal to Zero | $r[rs1] \geq 0$ |

### Format (3)



| Assembly Language Syntax |                                       |                            |
|--------------------------|---------------------------------------|----------------------------|
| movrZ                    | $reg_{rs1}, reg\_or\_imm10, reg_{rd}$ | ( <i>synonym: movre</i> )  |
| movrLEZ                  | $reg_{rs1}, reg\_or\_imm10, reg_{rd}$ |                            |
| movrLZ                   | $reg_{rs1}, reg\_or\_imm10, reg_{rd}$ |                            |
| movrNZ                   | $reg_{rs1}, reg\_or\_imm10, reg_{rd}$ | ( <i>synonym: movrne</i> ) |
| movrGZ                   | $reg_{rs1}, reg\_or\_imm10, reg_{rd}$ |                            |
| movrGEZ                  | $reg_{rs1}, reg\_or\_imm10, reg_{rd}$ |                            |

## Description

If the contents of integer register  $r[rs1]$  satisfy the condition specified in the  $rcond$  field, these instructions copy  $r[rs2]$  (if  $i = 0$ ) or  $sign\_ext(simm10)$  (if  $i = 1$ ) into  $r[rd]$ . If the contents of  $r[rs1]$  do not satisfy the condition, then  $r[rd]$  is not modified. These instructions treat the register contents as a signed integer value; they do not modify any condition codes.

---

**Implementation Note** – The UltraSPARC III Cu processor does not implement this instruction by tagging each register value; it looks at the full 64-bit register to determine a negative or zero.

---

## Exceptions


*illegal\_instruction* ( $rcond = 000_2$  or  $100_2$ )

---

# A.39 Multiply and Divide (64-bit)

| Opcode | op3     | Operation                     |
|--------|---------|-------------------------------|
| MULX   | 00 1001 | Multiply (signed or unsigned) |
| SDIVX  | 10 1101 | Signed Divide                 |
| UDIVX  | 00 1101 | Unsigned Divide               |

## Format (3)



| Assembly Language Syntax |                                     |
|--------------------------|-------------------------------------|
| mulx                     | $reg_{rs1}, reg\_or\_imm, reg_{rd}$ |
| sdivx                    | $reg_{rs1}, reg\_or\_imm, reg_{rd}$ |
| udivx                    | $reg_{rs1}, reg\_or\_imm, reg_{rd}$ |

### Description

MULX computes “ $r[rs1] \times r[rs2]$ ” if  $i = 0$  or “ $r[rs1] \times \text{sign\_ext}(simm13)$ ” if  $i = 1$ , and writes the 64-bit product into  $r[rd]$ . MULX can be used to calculate the 64-bit product for signed or unsigned operands (the product is the same).

SDIVX and UDIVX compute “ $r[rs1] \div r[rs2]$ ” if  $i = 0$  or “ $r[rs1] \div \text{sign\_ext}(simm13)$ ” if  $i = 1$ , and write the 64-bit result into  $r[rd]$ . SDIVX operates on the operands as signed integers and produces a corresponding signed result. UDIVX operates on the operands as unsigned integers and produces a corresponding unsigned result.

For SDIVX, if the largest negative number is divided by  $-1$ , the result should be the largest negative number. That is:

$$8000\ 0000\ 0000\ 0000_{16} \div \text{FFFF FFFF FFFF FFFF}_{16} = 8000\ 0000\ 0000\ 0000_{16}$$

These instructions do not modify any condition codes.

### Exceptions


*division\_by\_zero*

---

## A.40 No Operation

| Opcode | op     | op2 | Operation    |
|--------|--------|-----|--------------|
| NOP    | 0 0000 | 100 | No Operation |

## Format (2)



| Assembly Language Syntax |  |
|--------------------------|--|
| nop                      |  |

### Description

The NOP instruction changes no program-visible state (except that of the PC and nPC).

NOP is a special case of the SETHI instruction, with  $imm22 = 0$  and  $rd = 0$ .

### Exceptions

None

## A.41 Partial Store (VIS I)

| Opcode | imm_asi      | ASI Value        | Operation  |
|--------|--------------|------------------|--|
| STDFA  | ASI_PST8_P   | C0 <sub>16</sub> | Eight 8-bit conditional stores to primary address space                  |
| STDFA  | ASI_PST8_S   | C1 <sub>16</sub> | Eight 8-bit conditional stores to secondary address space                |
| STDFA  | ASI_PST8_PL  | C8 <sub>16</sub> | Eight 8-bit conditional stores to primary address space, little-endian   |
| STDFA  | ASI_PST8_SL  | C9 <sub>16</sub> | Eight 8-bit conditional stores to secondary address space, little-endian |
| STDFA  | ASI_PST16_P  | C2 <sub>16</sub> | Four 16-bit conditional stores to primary address space                  |
| STDFA  | ASI_PST16_S  | C3 <sub>16</sub> | Four 16-bit conditional stores to secondary address space                |
| STDFA  | ASI_PST16_PL | CA <sub>16</sub> | Four 16-bit conditional stores to primary address space, little-endian   |
| STDFA  | ASI_PST16_SL | CB <sub>16</sub> | Four 16-bit conditional stores to secondary address space, little-endian |
| STDFA  | ASI_PST32_P  | C4 <sub>16</sub> | Two 32-bit conditional stores to primary address space                   |
| STDFA  | ASI_PST32_S  | C5 <sub>16</sub> | Two 32-bit conditional stores to secondary address space                 |
| STDFA  | ASI_PST32_PL | CC <sub>16</sub> | Two 32-bit conditional stores to primary address space, little-endian    |
| STDFA  | ASI_PST32_SL | CD <sub>16</sub> | Two 32-bit conditional stores to secondary address space, little-endian  |

### Format (3)

|          |       |        |       |     |         |     |
|----------|-------|--------|-------|-----|---------|-----|
| 11       | rd    | 110111 | rs1   | i=0 | imm_asi | rs2 |
| 31 30 29 | 25 24 | 19 18  | 14 13 | 5 4 | 0       |     |

#### Assembly Language Syntax<sup>1</sup>

|      |  |
|------|--|
| stda | <i>freg<sub>rd</sub>, reg<sub>rs2</sub>, [reg<sub>rs1</sub>] imm_asi</i> |
|------|--|

1. The original assembly language syntax for a partial store instruction (“stda *freg<sub>rd</sub> [reg<sub>rs1</sub>] reg<sub>rs2</sub> imm\_asi*”) has been deprecated because of inconsistency with the rest of the SPARC assembly language. Over time, assemblers will support the new syntax for this instruction. In the meantime, some assemblers may recognize only the original syntax.



## *Description*

The partial store instructions are selected by one of the partial store ASIs with the *STDFA* instruction.

Two 32-bit, four 16-bit, or eight 8-bit values from the 64-bit floating-point register specified by *rd* are conditionally stored at the address specified by *r[rs1]*, using the mask specified in *r[rs2]*. The value in *r[rs2]* has the same format as the result specified by the pixel compare instructions (see Section A.44, “Pixel Compare (VIS I)”). The most significant bit of the mask (not the entire register) corresponds to the most significant part of the floating-point register specified by *rd*. The data is stored in little-endian form in memory if the ASI name has an “L” suffix; otherwise, it is stored in big-endian format.

A partial store instruction can cause a virtual (or physical) watchpoint exception when the following conditions are met:

- The virtual (physical) address in *r[rs1]* matches the address in the VA (PA) Data Watchpoint Register.
- The byte store mask in *r[rs2]* indicates that a byte is to be stored.
- The Virtual (Physical) Data Watchpoint Mask in *DCUCR* indicates that one or more of the bytes to be stored at the watched address is being watched.

Watchpoint exceptions on partial store instructions behaves as if every partial store always stores all 8 bytes. The *DCUCR* Data Watchpoint masks are only checked for nonzero value (watchpoint enabled). The byte store mask (*r[rs2]*) in the partial store instruction is ignored, and a watchpoint exception can occur even if the mask is zero (that is, no store will take place).

ASIs *C0<sub>16</sub>-C5<sub>16</sub>* and *C8<sub>16</sub>-CD<sub>16</sub>* are only used for partial store operations. In particular, they should not be used with the *LDDFA* instruction.

---

**Note** – If the byte ordering is little-endian, the byte enables generated by this instruction are swapped with respect to big-endian.

---

## *Exceptions*

*fp\_disabled*

*illegal\_instruction* (When *i* = 1, no immediate mode is supported.)

*PA\_watchpoint*

*VA\_watchpoint*

*mem\_address\_not\_aligned*

*data\_access\_exception*

*data\_access\_error*


*fast\_data\_access\_MMU\_miss*

*fast\_data\_access\_protection*

## A.42 Partitioned Add/Subtract Instructions (VIS I)

| Opcode   | opf         | Operation            |
|----------|-------------|----------------------|
| FPADD16  | 0 0101 0000 | Four 16-bit Add      |
| FPADD16S | 0 0101 0001 | Two 16-bit Add       |
| FPADD32  | 0 0101 0010 | Two 32-bit Add       |
| FPADD32S | 0 0101 0011 | One 32-bit Add       |
| FPSUB16  | 0 0101 0100 | Four 16-bit Subtract |
| FPSUB16S | 0 0101 0101 | Two 16-bit Subtract  |
| FPSUB32  | 0 0101 0110 | Two 32-bit Subtract  |
| FPSUB32S | 0 0101 0111 | One 32-bit Subtract  |

*Format (3)*



| Assembly Language Syntax |  |
|--------------------------|--|
| fpadd16                  | <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fpadd16s                 | <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fpadd32                  | <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fpadd32s                 | <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fpsub16                  | <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fpsub16s                 | <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fpsub32                  | <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i> |
| fpsub32s                 | <i>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></i> |

### Description

The standard versions of these instructions perform four 16-bit or two 32-bit partitioned adds or subtracts between the corresponding fixed-point values contained in the source operands (the 64-bit floating-point registers specified by *rs1* and *rs2*). For subtraction, the second operand (*rs2*) is subtracted from the first operand(*rs1*). The result is placed in the 64-bit destination register specified by *rd*.

The single-precision versions of these instructions (FPADD16S, FPSUB16S, FPADD32S, FPSUB32S) perform two 16-bit or one 32-bit partitioned add(s) or subtract(s); only the low 32 bits of the destination register are affected.

---

**Note** – For good performance, the result of a single FPADD should not be used as part of a source operand of a 64-bit graphics instruction in the next instruction group. Similarly, the result of a standard FPADD should not be used as a 32-bit graphics instruction source operand in the next three instruction groups.

---

### Exceptions


*fp\_disabled*

---

## A.43 Partitioned Multiply Instructions (VIS I)

| Opcode      | opf         | Operation   |
|-------------|-------------|---|
| FMUL8x16    | 0 0011 0001 | 8-bit x 16-bit Partitioned Product                |
| FMUL8x16AU  | 0 0011 0011 | 8-bit x 16-bit Upper $\alpha$ Partitioned Product |
| FMUL8x16AL  | 0 0011 0101 | 8-bit x 16-bit Upper $\alpha$ Partitioned Product |
| FMUL8SUx16  | 0 0011 0110 | Upper 8-bit x 16-bit Partitioned Product          |
| FMUL8ULx16  | 0 0011 0111 | Lower Unsigned 8-bit x 16-bit Partitioned Product |
| FMULD8SUx16 | 0 0011 1000 | Upper 8-bit x 16-bit Partitioned Product          |
| FMULD8ULx16 | 0 0011 1001 | Lower Unsigned 8-bit x 16-bit Partitioned Product |

### Format (3)



| Assembly Language Syntax |                                  |
|--------------------------|----------------------------------|
| <code>fmul8x16</code>    | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| <code>fmul8x16au</code>  | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| <code>fmul8x16al</code>  | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| <code>fmul8sux16</code>  | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| <code>fmul8ulx16</code>  | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| <code>fmuld8sux16</code> | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |
| <code>fmuld8ulx16</code> | $reg_{rs1}, reg_{rs2}, reg_{rd}$ |

### Description

---

**Notes** – For good performance, the result of a partitioned multiply should not be used as a 32-bit graphics instruction source operand in the next three instruction groups.

---

---

**Programming Note** – When software emulates an 8-bit unsigned by 16-bit signed multiply, the unsigned value must be zero-extended and the 16-bit value sign-extended before the multiplication.

---

---

**Note** – For good performance, the result of a partitioned multiply should not be used as a source operand of a 32-bit graphics instruction in the next three instruction groups.

---

The following sections describe the versions of partitioned multiplies.

### Exceptions

*fp\_disabled*

## A.43.1 FMUL8x16 Instruction

FMUL8x16 multiplies each unsigned 8-bit value (that is, a pixel) in  $f[rs1]$  by the corresponding (signed) 16-bit fixed-point integer in the 64-bit floating-point register specified by  $rs2$ ; it rounds the 24-bit product (assuming binary point between bits 7 and 8) and stores the upper 16 bits of the result into the corresponding 16-bit field in the 64-bit floating-point destination register specified by  $rd$ . FIGURE A-5 illustrates the operation.

---

**Note** – This instruction treats the pixel values as fixed-point with the binary point to the left of the most significant bit. Typically, this operation is used with filter coefficients as the fixed-point  $rs2$  value and image data as the  $rs1$  pixel value. Appropriate scaling of the coefficient allows various fixed-point scaling to be realized.

---




FIGURE A-5 FMUL8x16 Operation

## A.43.2 FMUL8x16AU Instruction

FMUL8x16AU is the same as FMUL8x16, except that one 16-bit fixed-point value is used for all four multiplies. This value is the most significant 16 bits of the 32-bit register  $f[rs2]$ , which is typically a proportional value. FIGURE A-6 illustrates the operation.




FIGURE A-6 FMUL8x16AU Operation

### A.43.3 FMUL8x16AL Instruction

FMUL8x16AL is the same as FMUL8x16AU, except that the least significant 16 bits of the 32-bit register  $f[rs2]$  are used as a proportional value. FIGURE A-7 illustrates the operation.




FIGURE A-7 FMUL8x16AL Operation

### A.43.4 FMUL8SUx16 Instruction

FMUL8SUx16 multiplies the upper 8 bits of each 16-bit signed value in the 64-bit floating-point register specified by  $rs1$  by the corresponding signed, 16-bit fixed-point, signed integer in the 64-bit floating-point register specified by  $rs2$ . It rounds the 24-bit

product toward the nearest representable value and then stores the upper 16 bits of the result into the corresponding 16-bit field of the 64-bit floating-point destination register specified by `rd`. If the product is exactly halfway between two integers, the result is rounded toward positive infinity. FIGURE A-8 illustrates the operation.




FIGURE A-8 FMUL8SUx16 Operation

## A.43.5 FMUL8ULx16 Instruction

`FMUL8ULx16` multiplies the unsigned lower 8 bits of each 16-bit value in the 64-bit floating-point register specified by `rs1` by the corresponding fixed-point signed integer in the 64-bit floating-point register specified by `rs2`. Each 24-bit product is sign-extended to 32 bits. The upper 16 bits of the sign-extended value are rounded to nearest representable value and then stored in the corresponding 16-bit field of the 64-bit floating-point destination register specified by `rd`. If the result is exactly halfway between two integers, the result is rounded toward positive infinity. FIGURE A-9 illustrates the operation. CODE EXAMPLE A-5 shows an example.

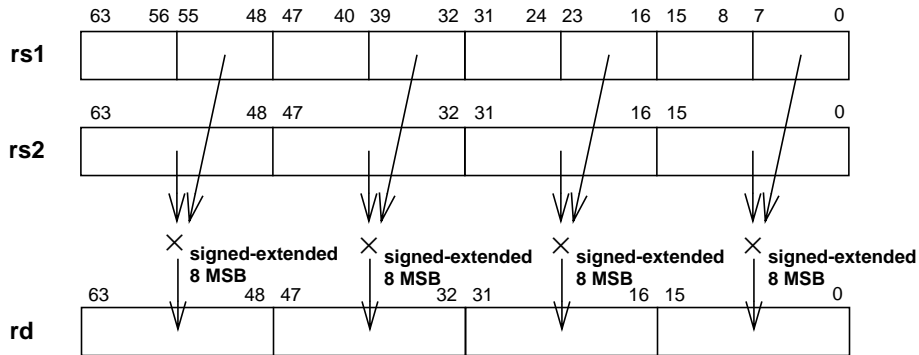


FIGURE A-9 FMUL8LUx16 Operation

CODE EXAMPLE A-5 FMUL8LUx16 Operation

```

fmul8sux16    %f0, %f1, %f2

fmul8ulx16    %f0, %f1, %f3

fpadd16       %f2, %f3, %f4

```

## A.43.6 FMULD8SUx16 Instruction

FMULD8SUx16 multiplies the upper 8 bits of each 16-bit signed value in  $f[rs1]$  by the corresponding signed 16-bit fixed-point signed integer in  $f[rs2]$ . Each 24-bit product is shifted left by 8 bits to make up a 32-bit result, which is then stored in the 64-bit floating-point register specified by  $rd$ . FIGURE A-10 illustrates the operation.






FIGURE A-10 FMULD8SUx16 Operation

## A.43.7 FMULD8ULx16 Instruction

FMULD8ULx16 multiplies the unsigned lower 8 bits of each 16-bit value in  $f[rs1]$  by the corresponding fixed-point signed integer in  $f[rs2]$ . Each 24-bit product is sign-extended to 32 bits and stored in the 64-bit floating-point register specified by  $rd$ . FIGURE A-11 illustrates the operation; CODE EXAMPLE A-6 exemplifies the operation.

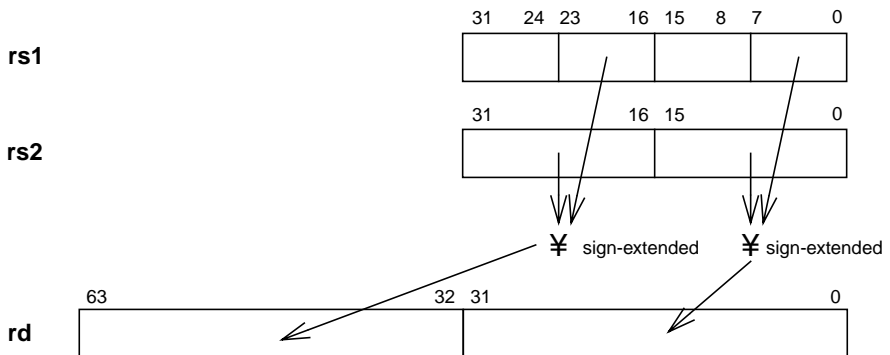


FIGURE A-11 FMULD8ULx16 Operation

**CODE EXAMPLE A-6** FMULD8ULx16 Operation

```

fmuld8sux16  %f0, %f1, %f2


fmuld8ulx16  %f0, %f1, %f3

fpadd32      %f2, %f3, %f4

```

## A.44 Pixel Compare (VIS I)

| Opcode   | opf         | Operation                                   |
|----------|-------------|---|
| FCMPGT16 | 0 0010 1000 | Four 16-bit Compares; set rd if src1 > src2 |
| FCMPGT32 | 0 0010 1100 | Two 32-bit Compares; set rd if src1 > src2  |
| FCMPLE16 | 0 0010 0000 | Four 16-bit Compares; set rd if src1 ≤ src2 |
| FCMPLE32 | 0 0010 0100 | Two 32-bit Compares; set rd if src1 ≤ src2  |
| FCMPNE16 | 0 0010 0010 | Four 16-bit Compares; set rd if src1 ≠ src2 |
| FCMPNE32 | 0 0010 0110 | Two 32-bit Compares; set rd if src1 ≠ src2  |
| FCMPEQ16 | 0 0010 1010 | Four 16-bit Compares; set rd if src1 = src2 |
| FCMPEQ32 | 0 0010 1110 | Two 32-bit Compares; set rd if src1 = src2  |

*Format (3)*

| Assembly Language Syntax |   |
|--------------------------|---|
| <code>fcmpgt16</code>    | <i>reg<sub>rs1</sub>, reg<sub>rs2</sub>, reg<sub>rd</sub></i> |
| <code>fcmpgt32</code>    | <i>reg<sub>rs1</sub>, reg<sub>rs2</sub>, reg<sub>rd</sub></i> |
| <code>fcmp1e16</code>    | <i>reg<sub>rs1</sub>, reg<sub>rs2</sub>, reg<sub>rd</sub></i> |
| <code>fcmp1e32</code>    | <i>reg<sub>rs1</sub>, reg<sub>rs2</sub>, reg<sub>rd</sub></i> |
| <code>fcmpne16</code>    | <i>reg<sub>rs1</sub>, reg<sub>rs2</sub>, reg<sub>rd</sub></i> |
| <code>fcmpne32</code>    | <i>reg<sub>rs1</sub>, reg<sub>rs2</sub>, reg<sub>rd</sub></i> |
| <code>fcmp16</code>      | <i>reg<sub>rs1</sub>, reg<sub>rs2</sub>, reg<sub>rd</sub></i> |
| <code>fcmp32</code>      | <i>reg<sub>rs1</sub>, reg<sub>rs2</sub>, reg<sub>rd</sub></i> |

### *Description*

Either four 16-bit or two 32-bit fixed-point values in the 64-bit floating-point source registers specified by `rs1` and `rs2` are compared. The 4-bit or 2-bit results are stored in the least significant bits in the integer destination register `r[rd]`. Signed comparisons are used. Bit 0 of `r[rd]` corresponds to the least significant 16-bit or 32-bit comparison.

For `FCMPGT`, each bit in the result is set if the corresponding value in the first source operand is greater than the value in the second source operand. Less-than comparisons are made by swapping the operands.

For `FCMPLE`, each bit in the result is set if the corresponding value in the first source operand is less than or equal to the value in the second source operand. Greater-than-or-equal comparisons are made by swapping the operands.

For `FCMP16`, each bit in the result is set if the corresponding value in the first source operand is equal to the value in the second source operand.

For `FCMPNE`, each bit in the result is set if the corresponding value in the first source operand is not equal to the value in the second source operand.


### *Exceptions*

*fp\_disabled*

## A.45 Pixel Component Distance (PDIST) (VIS I)

| Opcode | opf         | Operation                               |
|--------|-------------|---|
| PDIST  | 0 0011 1110 | Distance between eight 8-bit components |

### Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| <code>pdist</code>       | <code>freg<sub>rs1</sub>, freg<sub>rs2</sub>, freg<sub>rd</sub></code> |

### Description

Eight unsigned 8-bit values are contained in the 64-bit floating-point source registers specified by `rs1` and `rs2`. The corresponding 8-bit values in the source registers are subtracted (that is, the second source operand from the first source operand). The sum of the absolute value of each difference is added to the integer in the 64-bit floating-point destination register specified by `rd`. The result is stored in the destination register. Typically, this instruction is used for motion estimation in video compression algorithms.

---

**Note** – For good performance, the `rd` operand of `PDIST` should not reference the result of a non-`PDIST` instruction in the five previously executed instruction groups.

---


### Exceptions

*fp\_disabled*

## A.46 Pixel Formatting (VIS I)

| Opcode   | opf         | Operation                              |
|----------|-------------|--|
| FPACK16  | 0 0011 1011 | Four 16-bit packs into 8 unsigned bits |
| FPACK32  | 0 0011 1010 | Two 32-bit packs into 8 unsigned bit   |
| FPACKFIX | 0 0011 1101 | Four 16-bit packs into 16 signed bits  |
| FEXPAND  | 0 0100 1101 | Four 16-bit expands                    |
| FPMERGE  | 0 0100 1011 | Two 32-bit merges                      |

### Format (3)



| Assembly Language Syntax |                                     |
|--------------------------|-------------------------------------|
| <code>fpack16</code>     | $freg_{rs2}, freg_{rd}$             |
| <code>fpack32</code>     | $freg_{rs1}, freg_{rs2}, freg_{rd}$ |
| <code>fpackfix</code>    | $freg_{rs2}, freg_{rd}$             |
| <code>fexpand</code>     | $freg_{rs2}, freg_{rd}$             |
| <code>fpmerge</code>     | $freg_{rs1}, freg_{rs2}, freg_{rd}$ |

### Description

The FPACK instructions convert multiple values in a source register to a lower-precision fixed or pixel format and stores the resulting values in the destination register. Input values are clipped to the dynamic range of the output format. Packing applies a scale factor from `GSR.scale` to allow flexible positioning of the binary point.

---

**Programming Note** – For good performance, the result of an `FPACK` (including `FPACK32`) should not be used as part of a 64-bit graphics instruction source operand in the next three instruction groups.

---

`FEXPAND` performs the inverse of the `FPACK16` operation.

`FPMERGE` interleaves four 8-bit values from each of two 32-bit registers into a single 64-bit destination register.

---

**Programming Note** – The result of `FEXPAND` or `FPMERGE` should not be used as a 32-bit graphics instruction source operand in the next three instruction groups.

---

### *Exceptions*

*fp\_disabled*

## A.46.1 `FPACK16`

`FPACK16` takes four 16-bit fixed values from the 64-bit floating-point register specified by `rs2`, scales, truncates, and clips them into four 8-bit unsigned integers, and stores the results in the 32-bit destination register, `f[rd]`. FIGURE A-12 illustrates the `FPACK16` operation.

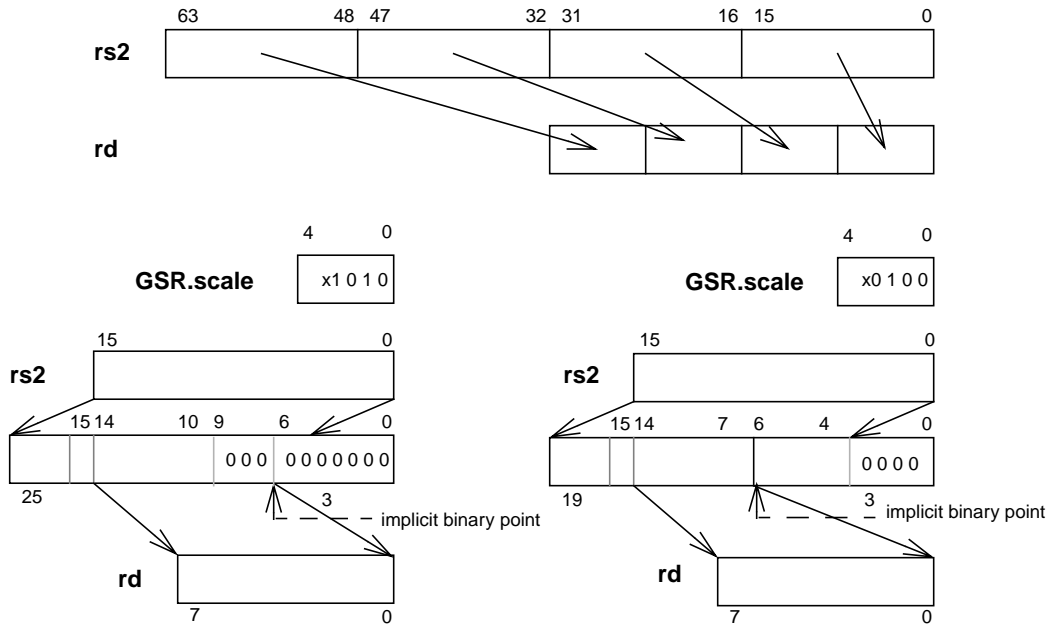


FIGURE A-12 FPACK16 Operation

---

**Note** – FPACK16 ignores the most significant bit of GSR.scale (GSR.scale<4>).

---

This operation is carried out as follows:

1. Left-shift the value from the 64-bit floating-point register specified by `rs2` by the number of bits specified in `GSR.scale` while maintaining clipping information.
2. Truncate and clip to an 8-bit unsigned integer starting at the bit immediately to the left of the implicit binary point (that is, between bits 7 and 6 for each 16-bit word). Truncation converts the scaled value into a signed integer (that is, round toward negative infinity). If the resulting value is negative (that is, its most significant bit is set), zero is returned as the clipped value. If the value is greater than 255, then 255 is delivered as the clipped value. Otherwise, the scaled value is returned as the result.
3. Store the result in the corresponding byte in the 32-bit destination register, `f[rd]`.

## A.46.2 FPACK32

FPACK32 takes two 32-bit fixed values from the second source operand (the 64-bit floating-point register specified by `rs2`) and scales, truncates, and clips them into two 8-bit unsigned integers. The two 8-bit integers are merged at the corresponding least significant byte positions with each 32-bit word in the 64-bit floating-point register specified by `rs1`, left-shifted by 8 bits. The 64-bit result is stored in the 64-bit floating-point register specified by `rd`. Thus, successive FPACK32 instructions can assemble two pixels by using three or four pairs of 32-bit fixed values. FIGURE A-13 illustrates the FPACK32 operation.




FIGURE A-13 FPACK32 Operation

This operation is carried out as follows:

1. Left-shift each 32-bit value from the second source operand by the number of bits specified in `GSR.scale`, while maintaining clipping information.
2. For each 32-bit value, truncate and clip to an 8-bit unsigned integer starting at the bit immediately to the left of the implicit binary point (that is, between bits 23 and 22 for each 32-bit word). Truncation converts the scaled value into a signed integer (that is,



round toward negative infinity). If the resulting value is negative (that is, MSB is set), then zero is returned as the clipped value. If the value is greater than 255, then 255 is delivered as the clipped value. Otherwise, the scaled value is returned as the result.

3. Left-shift each 32-bit value from the first source operand (the 64-bit floating-point register specified by `rs1`) by 8 bits.
4. Merge the two clipped 8-bit unsigned values into the corresponding least significant byte positions in the left-shifted value from the second source operand.
5. Store the result in the `rd` register.

### A.46.3 FPACKFIX

`FPACKFIX` takes two 32-bit fixed values from the 64-bit floating-point register specified by `rs2`, scales, truncates, and clips them into two 16-bit unsigned integers, and then stores the result in the 32-bit destination register `f[rd]`. FIGURE A-14 illustrates the `FPACKFIX` operation.




FIGURE A-14 `FPACKFIX` Operation

This operation is carried out as follows:

1. Left-shift each 32-bit value from the source operand (the 64-bit floating-point register specified by `rs2`) by the number of bits specified in `GSR.scale` while maintaining clipping information.
2. For each 32-bit value, truncate and clip to a 16-bit unsigned integer starting at the bit immediately to the left of the implicit binary point (that is, between bits 16 and 15 for each 32-bit word). Truncation converts the scaled value into a signed integer (that is, round toward negative infinity). If the resulting value is less than `-32768`, then `-32768` is returned as the clipped value. If the value is greater than `32767`, then `32767` is delivered as the clipped value. Otherwise, the scaled value is returned as the result.
3. Store the result in the 32-bit destination register `f[rd]`.

## A.46.4 FEXPAND

FEXPAND takes four 8-bit unsigned integers from `f[rs2]`, converts each integer to a 16-bit fixed-point value, and stores the four resulting 16-bit values in a 64-bit floating-point register specified by `rd`. FIGURE A-15 illustrates the operation.




FIGURE A-15 FEXPAND Operation

This operation is carried out as follows:

1. Left-shift each 8-bit value by four and zero-extend the results to a 16-bit fixed value.
2. Store the result in the destination register.

## A.46.5 FPMERGE

FPMERGE interleaves four corresponding 8-bit unsigned values in  $f[rs1]$  and  $f[rs2]$  to produce a 64-bit value in the 64-bit floating-point destination register specified by  $rd$ . This instruction converts from packed to planar representation when it is applied twice in succession, for example,

$R1G1B1A1, R3G3B3A3 \rightarrow R1R3G1G3A1A3 \rightarrow R1R2R3R4G1G2G3G4$ .

FPMERGE also converts from planar to packed when it is applied twice in succession, for example,  $R1R2R3R4, B1B2B3B4 \rightarrow R1B1R2B2R3B3R4B4 \rightarrow R1G1B1A1R2G2B2A2$ .

FIGURE A-16 illustrates the operation.




FIGURE A-16 FPMERGE Operation


Back-to-back FPMERGEs cannot be done on adjacent cycles.

---

## A.47 Population Count

| Opcode | op3     | Operation        |
|--------|---------|------------------|
| POPC   | 10 1110 | Population Count |

### Format (3)



| Assembly Language Syntax |                           |
|--------------------------|---------------------------|
| popc                     | <i>reg_or_imm, reg_rd</i> |

### Description

POPC counts the number of one bits in  $r[rs2]$  if  $i = 0$ , or the number of one bits in  $sign\_ext(simm13)$  if  $i = 1$ , and stores the count in  $r[rd]$ . This instruction does not modify the condition codes.

---

**Note** – The UltraSPARC III Cu processor does not implement this instruction in hardware; instead it traps to software. The instruction is emulated in supervisor software.

---

### Exceptions

*illegal\_instruction*

---

## A.48 Prefetch Data

| Opcode                   | op3     | Operation                          |
|--------------------------|---------|------------------------------------|
| PREFETCH                 | 10 1101 | Prefetch Data                      |
| PREFETCHA <sup>PAI</sup> | 11 1101 | Prefetch Data from Alternate Space |

---

**Implementation Note** – The PREFETCH{A} instructions are supported in the UltraSPARC III Cu processor.

---

### Format (3) PREFETCH{A}

|    |     |     |     |     |                                   |     |    |    |    |   |   |   |
|----|-----|-----|-----|-----|-----------------------------------|-----|----|----|----|---|---|---|
| 11 | fcn | op3 | rs1 | i=0 | PREFETCH: —<br>PREFETCHA: imm_asi | rs2 |    |    |    |   |   |   |
| 11 | fcn | op3 | rs1 | i=1 | simm13                            |     |    |    |    |   |   |   |
| 31 | 30  | 29  | 25  | 24  | 19                                | 18  | 14 | 13 | 12 | 5 | 4 | 0 |

See Chapter 16 “Prefetch” for complete details on the Prefetch Instructions and a definition of the `fcn` field.

| Assembly Language Syntax |  |
|--------------------------|--|
| <code>prefetch</code>    | <code>[address], prefetch_fcn</code>           |
| <code>prefetcha</code>   | <code>[regaddr] imm_asi, prefetch_fcn</code>   |
| <code>prefetcha</code>   | <code>[reg_plus_imm] %asi, prefetch_fcn</code> |

### Description

Prefetching is used to help manage data memory cache(s). A prefetch to a non-prefetchable location has no effect. Chapter 9 “Memory Models” describes non-prefetchable locations. Noncachable and non-prefetchable locations are not the same.

Variants of the prefetch instruction are used to prepare the memory system for different types of memory accesses. Chapter 16 “Prefetch” further discusses the prefetch instructions.

In non-privileged code, a prefetch instruction has no observable effect. Its execution is nonblocking and cannot cause an observable trap. In particular, a prefetch instruction shall not trap if it is applied to an illegal or nonexistent memory address.

---

**Programming Note** – When software needs to prefetch 64 bytes beginning at an arbitrary *address*, then issue two prefetch instructions to canvas all bytes:

```
prefetch[address], prefetch_fcn
prefetch[address + 63], prefetch_fcn
```

---

### PREFETCH A

Prefetch instructions that do *not* load from an alternate address space access the primary address space (ASI\_PRIMARY{`_LITTLE`}). Prefetch instructions that *do* load from an alternate address space contain the address space identifier (ASI) to be used for the load in

the `imm_as1` field if `i = 0`, or in the ASI register if `i = 1`. The access is privileged if bit 7 of the ASI is zero; otherwise, it is not privileged. The effective address for these instructions is “`r[rs1] + r[rs2]`” if `i = 0`, or “`r[rs1] + sign_ext(simm13)`” if `i = 1`.

See Chapter 8 “Address Space Identifiers” for acceptable ASI values used with the PREFETCH(A) instruction.

### Exceptions

*illegal\_instruction*

## A.48.1 Prefetch Instruction Variants

PREFETCH(A) instructions with `fcn = 0 – 3` are implemented.

Each prefetch variant reflects an intent on the part of the compiler or programmer. This is different from other instructions in SPARC V9 (except BFN), all of which specify specific actions.

The prefetch instruction variants are intended to provide scalability for future improvements in both hardware and compilers.

The prefetch variant is selected by the `fcn` field of the instruction. In accordance with SPARC V9, `fcn` values 4–15 cause an *illegal\_instruction* exception.

A prefetch with `fcn = 16` invalidates the P-cache line corresponding to the effective address of the prefetch. Use this characteristic to prefetch non-cacheable data after data are loaded into registers from the P-cache. A prefetch invalidate is issued to remove the data from the P-cache so it will not be found by a later reference. Prefetch with `fcn = 20, 21, 22, 23` map to `fcn 0–3` and are a new feature of the UltraSPARC III Cu processor.

TABLE A-13 lists the types of software prefetch instructions. Note that the table contains hexadecimal values for `fcn` unlike the decimal values in the explanation above.

**TABLE A-13** Types of Software Prefetch Instructions

| <b>fcn Value (hex)</b> | <b>Instruction Type</b>          | <b>Prefetch into:</b> | <b>Instruction Strength UltraSPARC III Cu</b> | <b>Request Exclusive Ownership</b> |
|------------------------|----------------------------------|-----------------------|---|------------------------------------|
| 00                     | Prefetch read many               | P-cache and L2-cache  | weak  | No                                 |
| 01                     | Prefetch read once               | P-cache only          | weak  | No                                 |
| 02                     | Prefetch write many              | L2-cache only         | weak  | Yes                                |
| 03                     | Prefetch write once <sup>1</sup> | L2-cache only         | weak  | No                                 |

**TABLE A-13** Types of Software Prefetch Instructions

| <b>fcn Value (hex)</b> | <b>Instruction Type</b> | <b>Prefetch into:</b>                              | <b>Instruction Strength UltraSPARC III Cu</b> | <b>Request Exclusive Ownership</b> |
|------------------------|-------------------------|--|---|------------------------------------|
| 04                     | <i>Reserved</i>         | Undefined  |   |                                    |
| 05 - 0F                | <i>Reserved</i>         | Undefined  |   |                                    |
| 10                     | Prefetch invalidate     | Invalidates a P-cache line, no data is prefetched. |   | N/A                                |
| 11 - 13                | <i>Reserved</i>         | Undefined  |   |                                    |
| 14                     | Same as fcn = 00        |  | weak <sup>2</sup>                             | No                                 |
| 15                     | Same as fcn = 01        |  | weak <sup>2</sup>                             | No                                 |
| 16                     | Same as fcn = 02        |  | weak <sup>2</sup>                             | Yes                                |
| 17                     | Same as fcn = 03        |  | weak <sup>2</sup>                             | No                                 |
| 18 - 1F                | <i>Reserved</i>         | Undefined  |   |                                    |

1. Although the name is “prefetch write once,” the actual use is prefetch to L2-cache for a future read.

2. These weak instructions may be implemented as strong in future implementations.

## A.48.2 New Error Handling of Prefetches

Since PREFETCH,<sub>2</sub> request for cache line ownership (RTO/R\_RTO), an error occurs while processing it will be handled differently compared to other prefetch requests with RTS/R\_RTS, as described in TABLE A-14.

**TABLE A-14** Error Handling of Prefetch Requests

| Prefetch Type             | L2-cache Hit/Miss    | Error Type                  | L2-cache Action                                | P-cache Action | Error Logging | Trap        |
|---------------------------|----------------------|-----------------------------|--|----------------|---------------|-------------|
| PREFETCH,2<br>(RTO/R_RTO) | Hit                  | Tag,<br>Hardware-corrected  | No state change                                | None           | THCE          | Disrupting  |
|                           | Miss                 | Tag,<br>Hardware-corrected  | Install data, state<br>change to M             | None           | THCE          | Disrupting  |
|                           | “Hit”<br>(tag error) | Tag, uncorrectable          | No data install,<br>no state change            | None           | TUE           | Fatal Error |
|                           | Hit                  | Data,<br>Hardware-corrected | No state change                                | None           | EDC           | Disrupting  |
|                           | Hit                  | Data,<br>uncorrectable      | No state change                                | None           | EDU           | Disrupting  |
|                           | Miss                 | Data,<br>Hardware-corrected | Install data, state<br>change to M             | None           | CE            | Disrupting  |
|                           | Miss                 | Data,<br>uncorrectable      | Install uncorrected<br>data, state change to M | None           | DUE           | Disrupting  |
|                           | Miss                 | MTag,<br>Hardware-corrected | Install data, state<br>change to M             | None           | EMC           | Disrupting  |
|                           | Miss                 | MTag,<br>uncorrectable      | Install data if L2-cache<br>state is M or Os   | None           | EMU           | Fatal Error |



**TABLE A-14** Error Handling of Prefetch Requests (*Continued*)

| Prefetch Type  | L2-cache Hit/Miss    | Error Type                  | L2-cache Action  | P-cache Action                             | Error Logging | Trap        |
|--|----------------------|-----------------------------|--|--|---------------|-------------|
| PREFETCH,0<br>PREFETCH,1<br>PREFETCH,3<br>Hardware prefetch<br>(RTS/R_RTS) | Hit                  | Tag,<br>Hardware-corrected  | No state change  | Install data<br>(except<br>PREFETCH,<br>3) | THCE          | Disrupting  |
|  | Miss                 | Tag,<br>Hardware-corrected  | Install data, state<br>change to S or E  | Install data<br>(except<br>PREFETCH,<br>3) | THCE          | Disrupting  |
|  | “Hit”<br>(tag error) | Tag,<br>uncorrectable       | No data install,<br>no state change  | Cancel<br>install                          | TUE           | Fatal Error |
|  | Hit                  | Data,<br>Hardware-corrected | No state change  | Install data<br>(except<br>PREFETCH,<br>3) | EDC           | Disrupting  |
|  | Hit                  | Data,<br>uncorrectable      | No state change  | Cancel<br>install                          | EDU           | Disrupting  |
|  | Miss                 | Data,<br>Hardware-corrected | Install data, state<br>change to S or E  | Install data<br>(except<br>PREFETCH,<br>3) | CE            | Disrupting  |
|  | Miss                 | Data,<br>uncorrectable      | - If RTS, cancel install,<br>no state change.<br>- If R_RTS, install<br>uncorrected data, state<br>change to Os. | Cancel<br>install                          | DUE           | Disrupting  |
|  | Miss                 | MTag,<br>Hardware-corrected | Install data, state<br>change to S or E  | None                                       | EMC           | Disrupting  |
|  | Miss                 | MTag,<br>uncorrectable      | Install data if L2-cache<br>state is M or Os   | None                                       | EMU           | Fatal Error |

### A.48.2.1 New Column in Coherence Table


A new column has been added to the UltraSPARC III Cu processor Coherence Table to describe the processor action on write prefetch RTO. Basically, the behavior of coherence state change is the following:

- On L2-cache hit: same as Load request (no state change)
- On L2-cache miss: same as Store request (send RTO/R\_RTO to get M state)

## A.49 Read Privileged Register

| Opcode            | op3     | Operation                |
|-------------------|---------|--------------------------|
| RDPR <sup>P</sup> | 10 1010 | Read Privileged Register |

*Format (3)*



| rs1   | Privileged Register |
|-------|---------------------|
| 0     | TPC                 |
| 1     | TNPC                |
| 2     | TSTATE              |
| 3     | TT                  |
| 4     | TICK                |
| 5     | TBA                 |
| 6     | PSTATE              |
| 7     | TL                  |
| 8     | PIL                 |
| 9     | CWP                 |
| 10    | CANSAVE             |
| 11    | CANRESTORE          |
| 12    | CLEANWIN            |
| 13    | OTHERWIN            |
| 14    | WSTATE              |
| 15    | FQ                  |
| 16–30 | —                   |
| 31    | VER                 |

| Assembly Language Syntax |                                      |
|--------------------------|--------------------------------------|
| rdpr                     | %tpc, <i>reg<sub>rd</sub></i>        |
| rdpr                     | %tnpc, <i>reg<sub>rd</sub></i>       |
| rdpr                     | %tstate, <i>reg<sub>rd</sub></i>     |
| rdpr                     | %tt, <i>reg<sub>rd</sub></i>         |
| rdpr                     | %tick, <i>reg<sub>rd</sub></i>       |
| rdpr                     | %tba, <i>reg<sub>rd</sub></i>        |
| rdpr                     | %pstate, <i>reg<sub>rd</sub></i>     |
| rdpr                     | %tl, <i>reg<sub>rd</sub></i>         |
| rdpr                     | %pil, <i>reg<sub>rd</sub></i>        |
| rdpr                     | %cwp, <i>reg<sub>rd</sub></i>        |
| rdpr                     | %cansave, <i>reg<sub>rd</sub></i>    |
| rdpr                     | %canrestore, <i>reg<sub>rd</sub></i> |
| rdpr                     | %cleanwin, <i>reg<sub>rd</sub></i>   |
| rdpr                     | %otherwin, <i>reg<sub>rd</sub></i>   |
| rdpr                     | %wstate, <i>reg<sub>rd</sub></i>     |
| rdpr                     | %fq, <i>reg<sub>rd</sub></i>         |
| rdpr                     | %ver, <i>reg<sub>rd</sub></i>        |

### *Description*

The `rs1` field in the instruction determines the privileged register that is read. There are MAXTL copies of the TPC, TNPC, TT, and TSTATE registers. A read from one of these registers returns the value in the register indexed by the current value in the trap level register (TL). A read of TPC, TNPC, TT, or TSTATE when the trap level is zero (TL = 0) causes an *illegal\_instruction* exception.

RDPR instructions with `rs1` in the range 16–30 are reserved; executing an RDPR instruction with `rs1` in that range causes an *illegal\_instruction* exception.

---

**Programming Note** – On this implementation with precise floating-point traps, the address of a trapping instruction will be in the TPC[TL] register when the trap code begins execution.

---

### Exceptions

*privileged\_opcode*


*illegal\_instruction* ((rs1 = 16–30) or ((rs1 ≤ 3) and (TL = 0)))

---

## A.50 Read State Register

| Opcode                             | op3     | rs1   | Operation  |
|------------------------------------|---------|-------|--|
| RDY <sup>D</sup>                   | 10 1000 | 0     | Read Y Register; deprecated (see Section A.70.9, “Read Y Register”)                        |
| —                                  | 10 1000 | 1     | <i>Reserved, do not access; attempt to access causes an illegal_instruction exception.</i> |
| RDCCR                              | 10 1000 | 2     | Read Condition Codes Register  |
| RDASI                              | 10 1000 | 3     | Read ASI Register  |
| RDTICK <sup>P<sub>NPT</sub></sup>  | 10 1000 | 4     | Read Tick Register   |
| RDPC                               | 10 1000 | 5     | Read Program Counter   |
| RDFPRS                             | 10 1000 | 6     | Read Floating-Point Registers Status Register  |
| —                                  | 10 1000 | 7–14  | <i>Reserved, do not access; attempt to access causes an illegal_instruction exception.</i> |
| <i>See description below</i>       | 10 1000 | 15    | STBAR, MEMBAR, or <i>Reserved</i> ; see description below.                                 |
| RDASR                              | 10 1000 | 16–31 | Read non-SPARC V9 ASRs   |
| RDPCR <sup>P<sub>PCR</sub></sup>   |         | 16    | Read Performance Control Registers (PCR)   |
| RDPIC <sup>P<sub>PIC</sub></sup>   |         | 17    | Read Performance Instrumentation Counters (PIC)  |
| RDDCR <sup>P</sup>                 |         | 18    | Read Dispatch Control Register (DCR)   |
| RDGSR                              |         | 19    | Read Graphic Status Register (GSR)   |
| —                                  |         | 20–21 | <i>Reserved, do not access; attempt to access causes an illegal_instruction exception.</i> |
| RDSOFTINT <sup>P</sup>             |         | 22    | Read per-processor Soft Interrupt Register   |
| RDTICK_CMPR <sup>P</sup>           |         | 23    | Read Tick Compare Register   |
| RDSTICK <sup>P<sub>NPT</sub></sup> |         | 24    | Read System TICK Register  |
| RDSTICK_CMPR <sup>P</sup>          |         | 25    | Read System TICK Compare Register  |
| —                                  |         | 26–31 | <i>Reserved, do not access; attempt to access causes an illegal_instruction exception.</i> |

### Format (3)



| <b>Assembly Language Syntax</b> |   |
|---------------------------------|---|
| rd                              | %ccr, <i>reg<sub>rd</sub></i>           |
| rd                              | %asi, <i>reg<sub>rd</sub></i>           |
| rd                              | %tick, <i>reg<sub>rd</sub></i>          |
| rd                              | %pc, <i>reg<sub>rd</sub></i>            |
| rd                              | %fprs, <i>reg<sub>rd</sub></i>          |
| rd                              | %pcr, <i>reg<sub>rd</sub></i>           |
| rd                              | %pic, <i>reg<sub>rd</sub></i>           |
| rd                              | %dcr, <i>reg<sub>rd</sub></i>           |
| rd                              | %gsr, <i>reg<sub>rd</sub></i>           |
| rd                              | %softint, <i>reg<sub>rd</sub></i>       |
| rd                              | %tick_cmpr, <i>reg<sub>rd</sub></i>     |
| rd                              | %sys_tick, <i>reg<sub>rd</sub></i>      |
| rd                              | %sys_tick_cmpr, <i>reg<sub>rd</sub></i> |

### Description

These instructions read the state register specified by *rs1* into *r[rd]*.

Values 7–14 of *rs1* are reserved for future versions of the architecture. A Read State Register instruction with *rs1* = 15, *rd* = 0, and *i* = 0 is defined to be a (deprecated) STBAR instruction (see Section A.70.10, “Store Barrier”). An RDASR instruction with *rs1* = 15, *rd* = 0, and *i* = 1 is defined to be a MEMBAR instruction. RDASR with *rs1* = 15 and *rd* ≠ 0 is reserved for future versions of the architecture; it causes an *illegal\_instruction* exception.

For RDPC, the processor writes the full 64-bit program counter value to the destination register of a CALL, JMPL, or RDPC instruction. When *PSTATE.AM* = 1 and a trap occurs, the processor writes the full 64-bit program counter value to *TPC[TL]*.

RDFPRS waits for all pending FPods and loads of floating-point registers to complete before reading the FPRS register.

RDGSR causes a *fp\_disabled* exception if *PSTATE.PEF* = 0 or *FPRS.FEF* = 0.

RDTICK causes a *privileged\_action* exception if `PSTATE.PRIV = 0` and `TICK.NPT = 1`.  
RDSTICK causes a *privileged\_action* exception if `PSTATE.PRIV = 0` and `STICK.NPT = 1`.

RDPIC causes a *privileged\_action* exception if `PSTATE.PRIV = 0` and `PCR.PRIV = 1`.

RDPCR causes a *privileged\_opcode* exception due to access privilege violation.

---

**Implementation Note** – Ancillary state registers include, for example, timer, counter, diagnostic, self-test, and trap-control registers.

---

---

**Compatibility Note** – The SPARC V8 RDPSR, RDWIM, and RDTBR instructions do not exist in SPARC V9 since the PSR, WIM, and TBR registers do not exist in SPARC V9.

---

### *Exceptions*

*privileged\_opcode*(RDDCR, RDSOFTINT, RDTICK\_CMPR, RDSTICK, RDSTICK\_CMPR,  
and RDPCR)

*illegal\_instruction*(RDASR with `rs1 = 1` or `7–14`;  
RDASR with `rs1 = 15` and `rd ≠ 0`;  
RDASR with `rs1 = 20–21, 26–31`)

*privileged\_action* (RDTICK with `PSTATE.PRIV = 0` and `TICK.NPT = 1`;  
RDPIC with `PSTATE.PRIV = 0` and `PCR.PRIV = 1`;  
RDSTICK with `PSTATE.PRIV = 0` and `STICK.NPT = 1`)


*fp\_disabled* (RDGSR with `PSTATE.PEF = 0` or `FPRS.FEF = 0`)

---

## A.51 RETURN

| Opcode | op3     | Operation |
|--------|---------|-----------|
| RETURN | 11 1001 | Return    |

### Format (3)



| Assembly Language Syntax |                |
|--------------------------|----------------|
| return                   | <i>address</i> |

### Description

The RETURN instruction causes a delayed transfer of control to the target address and has the window semantics of a RESTORE instruction; that is, it restores the register window prior to the last SAVE instruction. The target address is “ $r[rs1] + r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] + \text{sign\_ext}(simm13)$ ” if  $i = 1$ . Registers  $r[rs1]$  and  $r[rs2]$  come from the *old* window.

The RETURN instruction may cause an exception. It may cause a *window\_fill* exception as part of its RESTORE semantics, or it may cause a *mem\_address\_not\_aligned* exception if either of the two low-order bits of the target address is nonzero.

---

**Programming Note** – To re-execute the trapped instruction when returning from a user trap handler, use the RETURN instruction in the delay slot of a JMPL instruction, for example:

```

    jmpl      %l6,%g0   | Trapped PC supplied to user trap handler
    return   %l7       | Trapped nPC supplied to user trap handler

```

---

```

    save      %sp,-framesize, %sp
    . . .
    ret                               | Same as jmpl %i7+8, %g0
    restore                                 | Something useful like "restore
                                         | %o2,%l2,%o0"

```

or,

```

    save      %sp,-framesize, %sp
    . . .
    return   %i7+8
    nop                               | Could do some useful work in the caller's
                                         | window, for example, "or %o1, %o2,%o0"

```

---

---

**Programming Note** – A routine that uses a register window may be structured as either:

---

*Exceptions*


*mem\_address\_not\_aligned*  
*fill\_n\_normal* ( $n = 0 - 7$ )  
*fill\_n\_other* ( $n = 0 - 7$ )

---

## A.52 SAVE and RESTORE

| Opcode  | op3     | Operation               |
|---------|---------|-------------------------|
| SAVE    | 11 1100 | Save Caller's Window    |
| RESTORE | 11 1101 | Restore Caller's Window |

*Format (3)*



| Assembly Language Syntax |  |
|--------------------------|--|
| save                     | <i>reg<sub>rs1</sub>, reg<sub>or_imm</sub>, reg<sub>rd</sub></i> |
| restore                  | <i>reg<sub>rs1</sub>, reg<sub>or_imm</sub>, reg<sub>rd</sub></i> |

*Description (Effect on Non-privileged State)*

The SAVE instruction provides the routine executing it with a new register window. The *out* registers from the old window become the *in* registers of the new window. The contents of the *out* and the local registers in the new window are zero or contain values from the executing process; that is, the process sees a clean window.



The RESTORE instruction restores the register window saved by the last SAVE instruction executed by the current process. The *in* registers of the old window become the *out* registers of the new window. The *in* and local registers in the new window contain the previous values.

Furthermore, if and only if a spill or fill trap is not generated, SAVE and RESTORE behave like normal ADD instructions, except that the source operands  $r[rs1]$  or  $r[rs2]$  are read from the *old* window (that is, the window addressed by the original CWP) and the sum is written into  $r[rd]$  of the *new* window (that is, the window addressed by the new CWP).

**Note:** CWP arithmetic is performed modulo the number of windows, NWINDOWS.

---

**Programming Note** – Typically, if a SAVE (RESTORE) instruction traps, the spill (fill) trap handler returns to the trapped instruction to re-execute it. So, although the ADD operation is not performed the first time (when the instruction traps), it is performed the second time the instruction executes. The same applies to changing the CWP.

The SAVE instruction can be used to atomically allocate a new window in the register file and a new software stack frame in memory.

There is a performance trade-off to consider between using SAVE/RESTORE and saving and restoring selected registers explicitly.

---

### *Description (Effect on Privileged State)*

If the SAVE instruction does not trap, it increments the CWP (**mod** NWINDOWS) to provide a new register window and updates the state of the register windows by decrementing CANSAVE and incrementing CANRESTORE.

If the new register window is occupied (that is, CANSAVE = 0), a spill trap is generated. The trap vector for the spill trap is based on the value of OTHERWIN and WSTATE. The spill trap handler is invoked with the CWP set to point to the window to be spilled (that is, old CWP + 2).

If CANSAVE  $\neq$  0, the SAVE instruction checks whether the new window needs to be cleaned. It causes a *clean\_window* trap if the number of unused clean windows is zero, that is, (CLEANWIN – CANRESTORE) = 0. The *clean\_window* trap handler is invoked with the CWP set to point to the window to be cleaned (that is, old CWP + 1).

If the RESTORE instruction does not trap, it decrements the CWP (**mod** NWINDOWS) to restore the register window that was in use prior to the last SAVE instruction executed by the current process. It also updates the state of the register windows by decrementing CANRESTORE and incrementing CANSAVE.

If the register window to be restored has been spilled (CANRESTORE = 0), then a fill trap is generated. The trap vector for the fill trap is based on the values of OTHERWIN and WSTATE. The fill trap handler is invoked with CWP set to point to the window to be filled, that is, old CWP – 1.

---

**Programming Note** – The vectoring of spill and fill traps can be controlled by setting the value of the OTHERWIN and WSTATE registers appropriately.

The spill (fill) handler normally will end with a SAVED (RESTORED) instruction followed by a RETRY instruction.

---

### Exceptions

*clean\_window* (SAVE only)

*fill\_n\_normal* (RESTORE only,  $n = 0 - 7$ )

*fill\_n\_other* (RESTORE only,  $n = 0 - 7$ )

*spill\_n\_normal* (SAVE only,  $n = 0 - 7$ )


*spill\_n\_other* (SAVE only,  $n = 0 - 7$ )

---

## A.53 SAVED and RESTORED

| Opcode                | op3     | fcn  | Operation                |
|-----------------------|---------|------|--------------------------|
| SAVED <sup>P</sup>    | 11 0001 | 0    | Window has been saved    |
| RESTORED <sup>P</sup> | 11 0001 | 1    | Window has been restored |
| —                     | 11 0001 | 2–31 | <i>Reserved</i>          |

### Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| saved                    |  |
| restored                 |  |

### Description

SAVED and RESTORED adjust the state of the register-windows control registers.

SAVED increments CANSAVE. If OTHERWIN = 0, SAVED decrements CANRESTORE. If OTHERWIN ≠ 0, it decrements OTHERWIN.

RESTORED increments CANRESTORE. If CLEANWIN < (NWINDOWS-1), then RESTORED increments CLEANWIN. If OTHERWIN = 0, it decrements CANSAVE. If OTHERWIN ≠ 0, it decrements OTHERWIN.

---

**Programming Note** – The spill (fill) handlers use the SAVED (RESTORED) instruction to indicate that a window has been spilled (filled) successfully.

---

Normal privileged software would probably not do a SAVED or RESTORED from trap level zero (TL = 0). However, it is not illegal to do so and doing so does not cause a trap.

Executing a SAVED (RESTORED) instruction outside of a window spill (fill) trap handler is likely to create an inconsistent window state. Hardware will not signal an exception, however, since maintaining a consistent window state is the responsibility of privileged software.

---

### *Exceptions*

*privileged\_opcode*

*illegal\_instruction* (f.cn = 2-31)

---

## A.54 Set Interval Arithmetic Mode (VIS II)

| Opcode | opf         | Operation  |
|--------|-------------|--|
| SIAM   | 0 1000 0001 | Set the interval arithmetic mode fields in the GSR |

### *Format (3)*

|          |   |        |       |       |         |      |
|----------|---|--------|-------|-------|---------|------|
| 10       | — | 110110 | —     | opf   | —       | mode |
| 31 30 29 |   | 25 24  | 19 18 | 14 13 | 5 4 3 2 | 0    |

| Assembly Language Syntax |             |
|--------------------------|-------------|
| siam                     | <i>mode</i> |

### *Description*

The SIAM instruction sets the GSR.IM and GSR.IRND fields as follows:

GSR.IM = mode<2>

GSR.IRND = mode<1:0>

---

**Note** – SIAM is a groupable, break-after instruction. It enables the interval rounding mode to be changed every cycle without flushing the pipeline. FPops in the same instruction group as an SIAM instruction use the previous rounding mode.

---

### *Exceptions*


*fp\_disabled*

---

## A.55 SETHI

| Opcode | op2 | Operation                    |
|--------|-----|------------------------------|
| SETHI  | 100 | Set High 22 Bits of Low Word |

*Format (2)*



| Assembly Language Syntax |                                |
|--------------------------|--------------------------------|
| sethi                    | const22, reg <sub>rd</sub>     |
| sethi                    | %hi (value), reg <sub>rd</sub> |

*Description*

SETHI zeroes the least significant 10 bits and the most significant 32 bits of  $r[rd]$  and replaces bits 31 through 10 of  $r[rd]$  with the value from its imm22 field.

SETHI does not affect the condition codes.

Some SETHI instructions with  $rd = 0$  has a special use:

- $rd = 0$  and  $imm22 = 0$ : has no architectural effect and is defined to be a NOP instruction.
- $rd = 0$  and  $imm22 \neq 0$  is used to trigger hardware performance counters. See Chapter 14 “Performance Instrumentation” for details.

---

**Programming Note** – The most common form of 64-bit constant generation is creating stack offsets whose magnitude is less than  $2^{32}$ . The following code can be used to create the constant 0000 0000 ABCD 1234<sub>16</sub>:

---

```
sethi    %hi(0xabcd1234), %00
or       %00, 0x234, %00
```

The following code shows how to create a negative constant. **Note:** The immediate field of the xor instruction is sign extended and can be used to get ones in all of the upper 32 bits. For example, to set the negative constant FFFF FFFF ABCD 1234<sub>16</sub>:

```
sethi    %hi(0x5432edcb), %00 ! note 0x5432EDCB, not 0xABCD1234
xor      %00, 0x1e34, %00    ! part of imm. overlaps upper bits
```

---

## Exceptions


None

---

# A.56 Shift

| Opcode | op3     | x | Operation                        |
|--------|---------|---|----------------------------------|
| SLL    | 10 0101 | 0 | Shift Left Logical – 32 bits     |
| SRL    | 10 0110 | 0 | Shift Right Logical – 32 bits    |
| SRA    | 10 0111 | 0 | Shift Right Arithmetic – 32 bits |
| SLLX   | 10 0101 | 1 | Shift Left Logical – 64 bits     |
| SRLX   | 10 0110 | 1 | Shift Right Logical – 64 bits    |
| SRAX   | 10 0111 | 1 | Shift Right Arithmetic – 64 bits |

### Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| sll                      | <i>reg<sub>rs1</sub>, reg<sub>or_shcnt</sub>, reg<sub>rd</sub></i> |
| srl                      | <i>reg<sub>rs1</sub>, reg<sub>or_shcnt</sub>, reg<sub>rd</sub></i> |
| sra                      | <i>reg<sub>rs1</sub>, reg<sub>or_shcnt</sub>, reg<sub>rd</sub></i> |
| sllx                     | <i>reg<sub>rs1</sub>, reg<sub>or_shcnt</sub>, reg<sub>rd</sub></i> |
| srlx                     | <i>reg<sub>rs1</sub>, reg<sub>or_shcnt</sub>, reg<sub>rd</sub></i> |
| srax                     | <i>reg<sub>rs1</sub>, reg<sub>or_shcnt</sub>, reg<sub>rd</sub></i> |

### Description

When  $i = 0$  and  $x = 0$ , the shift count is the least significant five bits of  $r[rs2]$ . When  $i = 0$  and  $x = 1$ , the shift count is the least significant six bits of  $r[rs2]$ . When  $i = 1$  and  $x = 0$ , the shift count is the immediate value specified in bits 0 through 4 of the instruction. When  $i = 1$  and  $x = 1$ , the shift count is the immediate value specified in bits 0 through 5 of the instruction.

TABLE A-15 shows the shift count encodings for all values of  $i$  and  $x$ .

**TABLE A-15** Shift Count Encodings

| $i$ | $x$ | Shift Count             |
|-----|-----|-------------------------|
| 0   | 0   | bits 4–0 of $r[rs2]$    |
| 0   | 1   | bits 5–0 of $r[rs2]$    |
| 1   | 0   | bits 4–0 of instruction |
| 1   | 1   | bits 5–0 of instruction |

SLL and SLLX shift all 64 bits of the value in  $r[rs1]$  left by the number of bits specified by the shift count, replacing the vacated positions with zeroes, and write the shifted result to  $r[rd]$ .

SRL shifts the low 32 bits of the value in  $r[rs1]$  right by the number of bits specified by the shift count. Zeroes are shifted into bit 31. The upper 32 bits are set to zero, and the result is written to  $r[rd]$ .

SRLX shifts all 64 bits of the value in  $r[rs1]$  right by the number of bits specified by the shift count. Zeroes are shifted into the vacated high-order bit positions, and the shifted result is written to  $r[rd]$ .

SRA shifts the low 32 bits of the value in  $r[rs1]$  right by the number of bits specified by the shift count and replaces the vacated positions with bit 31 of  $r[rs1]$ . The high-order 32 bits of the result are all set with bit 31 of  $r[rs1]$ , and the result is written to  $r[rd]$ .

SRAX shifts all 64 bits of the value in  $r[rs1]$  right by the number of bits specified by the shift count and replaces the vacated positions with bit 63 of  $r[rs1]$ . The shifted result is written to  $r[rd]$ .

No shift occurs when the shift count is zero, but the high-order bits are affected by the 32-bit shifts as noted above.

These instructions do not modify the condition codes.

---

**Programming Note** – “Arithmetic left shift by 1 (and calculate overflow)” can be effected with the `ADDcc` instruction.

The instruction `sra rs1,0,rd` can be used to convert a 32-bit value to 64 bits, with sign extension into the upper word; `srl rs1,0,rd` can be used to clear the upper 32 bits of `r[rd]`.

---

### *Exceptions*

None

---

## A.57 Short Floating-Point Load and Store (VIS I)


| Opcode         | imm_asi     | ASI Value        | Operation  |
|----------------|-------------|------------------|--|
| LDDFA<br>STDFA | ASI_FL8_P   | D0 <sub>16</sub> | 8-bit load/store from/to primary address space                   |
| LDDFA<br>STDFA | ASI_FL8_S   | D1 <sub>16</sub> | 8-bit load/store from/to secondary address space                 |
| LDDFA<br>STDFA | ASI_FL8_PL  | D8 <sub>16</sub> | 8-bit load/store from/to primary address space, little-endian    |
| LDDFA<br>STDFA | ASI_FL8_SL  | D9 <sub>16</sub> | 8-bit load/store from/to secondary address space, little-endian  |
| LDDFA<br>STDFA | ASI_FL16_P  | D2 <sub>16</sub> | 16-bit load/store from/to primary address space                  |
| LDDFA<br>STDFA | ASI_FL16_S  | D3 <sub>16</sub> | 16-bit load/store from/to secondary address space                |
| LDDFA<br>STDFA | ASI_FL16_PL | DA <sub>16</sub> | 16-bit load/store from/to primary address space, little-endian   |
| LDDFA<br>STDFA | ASI_FL16_SL | DB <sub>16</sub> | 16-bit load/store from/to secondary address space, little-endian |



### Format (3) LDDFA



### Format (3) STDFA



| Assembly Language Syntax |   |
|--------------------------|---|
| ldda                     | <i>[reg_addr] imm_asi, freg<sub>rd</sub></i>  |
| ldda                     | <i>[reg_plus_imm] %asi, freg<sub>rd</sub></i> |
| stda                     | <i>freg<sub>rd</sub>, [reg_addr] imm_asi</i>  |
| stda                     | <i>freg<sub>rd</sub>, [reg_plus_imm] %asi</i> |

### Description

Short floating-point load and store instructions are selected by means of one of the short ASIs with the LDDFA and STDFA instructions.

These ASIs allow 8- and 16-bit loads or stores to be performed to/from the floating-point registers. Eight-bit loads can be performed to arbitrary byte addresses. For 16-bit loads, the least significant bit of the address must be zero or a *mem\_address\_not\_aligned* trap is taken. Short loads are zero-extended to the full floating-point register. Short stores access the low-order 8 or 16 bits of the register.

Little-endian ASIs transfer data in little-endian format in memory; otherwise, memory is assumed to be big-endian. Short loads and stores are typically used with the FALIGNEDDATA instruction (see Section A.2, “Alignment Instructions (VIS I)”) to assemble or store 64 bits on noncontiguous components.

### Exceptions

*fp\_disabled*

*PA\_watchpoint*

*VA\_watchpoint*

*mem\_address\_not\_aligned* (odd memory address for a 16-bit load or store)

*data\_access\_exception*

*data\_access\_error*

*fast\_data\_access\_MMU\_miss*

*fast\_data\_access\_protection*

---

## A.58 SHUTDOWN (VIS I)

| Opcode                | opf         | Operation                          |
|-----------------------|-------------|------------------------------------|
| SHUTDOWN <sup>P</sup> | 0 1000 0000 | Shut down to enter power-down mode |

### Format (3)

|          |       |        |       |     |   |
|----------|-------|--------|-------|-----|---|
| 10       | —     | 110110 | —     | opf | — |
| 31 30 29 | 25 24 | 19 18  | 14 13 | 5 4 | 0 |

| Assembly Language Syntax |  |
|--------------------------|--|
| shutdown                 |  |

### Description

SHUTDOWN is a privileged instruction.

The SHUTDOWN instruction executes as a NOP. An external system signal is used to enter and leave low power mode.

Because SHUTDOWN is a privileged instruction, an attempt to execute it while in non-privileged mode causes a *privileged\_opcode* trap.

## Exceptions


*privileged\_opcode*

---

# A.59 Software-Initiated Reset

| Opcode | op3     | rd | Operation                |
|--------|---------|----|--------------------------|
| SIR    | 11 0000 | 15 | Software-Initiated Reset |

*Format (3)*



| Assembly Language Syntax |                     |
|--------------------------|---------------------|
| <code>sir</code>         | <code>simm13</code> |

### Description

SIR is used to generate a software-initiated reset (SIR). As with other traps, a software-initiated reset performs different actions when  $TL = MAXTL$  than it does when  $TL < MAXTL$ .

When executed in non-privileged mode, SIR acts like a NOP with no visible effect.

### Exceptions


*software\_initiated\_reset*

## A.60 Store Floating-Point

| Opcode | op3     | rd   | Operation                            |
|--------|---------|------|--------------------------------------|
| STF    | 10 0100 | 0–31 | Store Floating-Point Register        |
| STDF   | 10 0111 | †    | Store Double Floating-Point Register |
| STQF   | 10 0110 | †    | Store Quad Floating-Point Register   |
| STXFSR | 10 0101 | 1    | Store Floating-Point State Register  |
| —      | 10 0101 | 2–31 | <i>Reserved</i>                      |

† Encoded floating-point register value.

### Format (3)



| Assembly Language Syntax |                        |
|--------------------------|------------------------|
| st                       | $freg_{rd}, [address]$ |
| std                      | $freg_{rd}, [address]$ |
| stq                      | $freg_{rd}, [address]$ |
| stx                      | $\%fsr, [address]$     |

### Description

The store single floating-point instruction (STF) copies  $f[rd]$  into memory.

The store double floating-point instruction (STDF) copies a doubleword from a double floating-point register into a word-aligned doubleword in memory.

The store quad floating-point instruction (STQF) traps to software.

The store floating-point state register instruction (STXFSR) waits for any currently executing FPop instructions to complete, and then it writes all 64 bits of the FSR into memory.

STXFSR zeroes FSR.ftt after writing the FSR to memory.

---

**Implementation Note** – FSR.ftt should not be zeroed until it is known that the store will not cause a precise trap.

---

The effective address for these instructions is “r[rs1] + r[rs2]” if *i* = 0, or “r[rs1] + sign\_ext(simml3)” if *i* = 1.

STF requires word alignment otherwise a *mem\_address\_not\_aligned* exception occurs.

STDF instruction causes a *STDF\_mem\_address\_not\_aligned* trap if the effective address is 32-bit aligned but not 64-bit (doubleword) aligned. In this case, the trap handler software shall emulate the STDF instruction and return.

STXFSR requires doubleword alignment; otherwise, it causes a *mem\_address\_not\_aligned* exception. In this case, the trap handler software shall emulate the STXFSR instruction and return.

If the floating-point unit (FPU) is not enabled for the source register *rd* (per FPRS.FEF and PSTATE.PEF) or if the FPU is not present, then a store floating-point instruction causes a *fp\_disabled* exception.

---

**Programming Note** – In SPARC V8, some compilers issued sets of single-precision stores when they could not determine that doubleword or quadword operands were properly aligned. For SPARC V9, since emulation of misaligned stores is expected to be fast, it is recommended that compilers issue sets of single-precision stores only when they can determine that doubleword or quadword operands are *not* properly aligned.

---

### *Exceptions*

*illegal\_instruction* (op3 = 25<sub>16</sub> and rd = 2–31)

*fp\_disabled*

*mem\_address\_not\_aligned*

*STDF\_mem\_address\_not\_aligned* (STDF only)

*data\_access\_exception*

*data\_access\_error*

*fast\_data\_access\_MMU\_miss*

*fast\_data\_access\_protection*

*PA\_watchpoint*


*VA\_watchpoint*

## A.61 Store Floating-Point into Alternate Space

| Opcode                | op3     | rd   | Operation   |
|-----------------------|---------|------|---|
| STFA <sup>PASI</sup>  | 11 0100 | 0–31 | Store Floating-Point Register to Alternate Space        |
| STDFA <sup>PASI</sup> | 11 0111 | †    | Store Double Floating-Point Register to Alternate Space |
| STQFA <sup>PASI</sup> | 11 0110 | †    | Store Quad Floating-Point Register to Alternate Space   |

† Encoded floating-point register value.

### Format (3)



| Assembly Language Syntax |   |
|--------------------------|---|
| sta                      | <i>f</i> reg <sub>rd</sub> , [ <i>regaddr</i> ] imm_asi   |
| sta                      | <i>f</i> reg <sub>rd</sub> , [ <i>reg_plus_imm</i> ] %asi |
| stda                     | <i>f</i> reg <sub>rd</sub> , [ <i>regaddr</i> ] imm_asi   |
| stda                     | <i>f</i> reg <sub>rd</sub> , [ <i>reg_plus_imm</i> ] %asi |
| stqa                     | <i>f</i> reg <sub>rd</sub> , [ <i>regaddr</i> ] imm_asi   |
| stqa                     | <i>f</i> reg <sub>rd</sub> , [ <i>reg_plus_imm</i> ] %asi |

### Description

The store single floating-point into alternate space instruction (STFA) copies *f*[rd] into memory.

The store double floating-point into alternate space instruction (STDFA) copies a doubleword from a double floating-point register into a word-aligned doubleword in memory.

The store quad floating-point into alternate space instruction (STQFA) traps to software.

Store floating-point into alternate space instructions contain the address space identifier (ASI) to be used for the load in the `imm_asi` field if `i = 0` or in the ASI register if `i = 1`. The access is privileged if bit 7 of the ASI is zero; otherwise, it is not privileged. The effective address for these instructions is “`r[rs1] + r[rs2]`” if `i = 0`, or “`r[rs1] + sign_ext(simm13)`” if `i = 1`.

STFA requires word alignment; otherwise, a *mem\_address\_not\_aligned* exception occurs.

STDFA instruction causes a *STDF\_mem\_address\_not\_aligned* trap if the effective address is 32-bit aligned but not 64-bit (doubleword) aligned. In this case, the trap handler software shall emulate the STDF instruction and return.

STDFA with certain target ASI is defined to be a 64-byte block-store instruction. See Section A.4, “Block Load and Block Store (VIS I)” for details.

If the floating-point unit is not enabled for the source register `rd` (per `FPRS.FEF` and `PSTATE.PEF`) or if the FPU is not present, store floating-point into alternate space instructions cause a *fp\_disabled* exception.

---

**Implementation Notes** – This check is not made for STQFA. STFA and STDFA cause a *privileged\_action* exception if `PSTATE.PRIV = 0` and bit 7 of the ASI is zero.

---

---

**Programming Note** – In SPARC V8, some compilers issued sets of single-precision stores when they could not determine that doubleword or quadword operands were properly aligned. For SPARC V9, since emulation of misaligned stores is expected to be fast, we recommend that compilers issue sets of single-precision stores only when they can determine that doubleword or quadword operands are *not* properly aligned.

---

## Exceptions

*illegal\_instruction*

*fp\_disabled*

*mem\_address\_not\_aligned*

*STDF\_mem\_address\_not\_aligned* (STDFA only)

*privileged\_action*

*data\_access\_exception*

*data\_access\_error*

*fast\_data\_access\_MMU\_miss*

*fast\_data\_access\_protection*


*PA\_watchpoint*

*VA\_watchpoint*

## A.62 Store Integer

| Opcode | op3     | Operation           |
|--------|---------|---------------------|
| STB    | 00 0101 | Store Byte          |
| STH    | 00 0110 | Store Halfword      |
| STW    | 00 0100 | Store Word          |
| STX    | 00 1110 | Store Extended Word |

*Format (3)*



| Assembly Language Syntax |                      |                            |
|--------------------------|----------------------|----------------------------|
| stb                      | $reg_{rd} [address]$ | (synonyms: stub, stsb)     |
| sth                      | $reg_{rd} [address]$ | (synonyms: stuh, stsh)     |
| stw                      | $reg_{rd} [address]$ | (synonyms: st, stuw, stsw) |
| stx                      | $reg_{rd} [address]$ |                            |

### Description

The store integer instructions (except store doubleword) copy the whole extended (64-bit) integer, the less significant word, the least significant halfword, or the least significant byte of  $r[rd]$  into memory.

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

A successful store (notably, store extended) instruction operates atomically.

STH causes a *mem\_address\_not\_aligned* exception if the effective address is not halfword aligned. STW causes a *mem\_address\_not\_aligned* exception if the effective address is not word aligned. STX causes a *mem\_address\_not\_aligned* exception if the effective address is not doubleword aligned.



## Exceptions

*mem\_address\_not\_aligned* (all except STB)

*data\_access\_exception*

*data\_access\_error*

*fast\_data\_access\_MMU\_miss*

*fast\_data\_access\_protection*

*PA\_watchpoint*


*VA\_watchpoint*

---

## A.63 Store Integer into Alternate Space

| Opcode               | op3     | Operation                                |
|----------------------|---------|--|
| STBA <sup>PASI</sup> | 01 0101 | Store Byte into Alternate Space          |
| STHA <sup>PASI</sup> | 01 0110 | Store Halfword into Alternate Space      |
| STWA <sup>PASI</sup> | 01 0100 | Store Word into Alternate Space          |
| STXA <sup>PASI</sup> | 01 1110 | Store Extended Word into Alternate Space |

### Format (3)



| Assembly Language Syntax |                                    |                               |
|--------------------------|------------------------------------|-------------------------------|
| stba                     | $reg_{rd}, [regaddr] imm_{asi}$    | (synonyms: stuba, stsba)      |
| stha                     | $reg_{rd}, [regaddr] imm_{asi}$    | (synonyms: stuha, stsha)      |
| stwa                     | $reg_{rd}, [regaddr] imm_{asi}$    | (synonyms: sta, stuwa, stswa) |
| stxa                     | $reg_{rd}, [regaddr] imm_{asi}$    |                               |
| stba                     | $reg_{rd}, [reg\_plus\_imm] \%asi$ | (synonyms: stuba, stsba)      |
| stha                     | $reg_{rd}, [reg\_plus\_imm] \%asi$ | (synonyms: stuha, stsha)      |
| stwa                     | $reg_{rd}, [reg\_plus\_imm] \%asi$ | (synonyms: sta, stuwa, stswa) |
| stxa                     | $reg_{rd}, [reg\_plus\_imm] \%asi$ |                               |

### Description

The store integer into alternate space instructions copy the whole extended (64-bit) integer, the less significant word, the least significant halfword, or the least significant byte of  $r[rd]$  into memory.

Store integer to alternate space instructions contain the address space identifier (ASI) to be used for the store in the  $imm_{asi}$  field if  $i = 0$ , or in the ASI register if  $i = 1$ . The access is privileged if bit 7 of the ASI is zero; otherwise, it is not privileged. The effective address for these instructions is “ $r[rs1] + r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] + sign\_ext(simm13)$ ” if  $i = 1$ .

A successful store (notably, store extended) instruction operates atomically.

STHA causes a *mem\_address\_not\_aligned* exception if the effective address is not halfword aligned. STWA causes a *mem\_address\_not\_aligned* exception if the effective address is not word aligned. STXA causes a *mem\_address\_not\_aligned* exception if the effective address is not doubleword aligned.

A store integer into alternate space instruction causes a *privileged\_action* exception if  $PSTATE.PRIV = 0$  and bit 7 of the ASI is zero.

---

**Compatibility Note** – The SPARC V8 STA instruction is renamed STWA in SPARC V9.

---

### Exceptions

*privileged\_action*  
*mem\_address\_not\_aligned* (all except STBA)  
*data\_access\_exception*  
*data\_access\_error*  
*fast\_data\_access\_MMU\_miss*

*fast\_data\_access\_protection*  
*PA\_watchpoint*  
*VA\_watchpoint*

## A.64 Subtract

| Opcode | op3     | Operation                           |
|--------|---------|-------------------------------------|
| SUB    | 00 0100 | Subtract                            |
| SUBcc  | 01 0100 | Subtract and modify cc's            |
| SUBC   | 00 1100 | Subtract with Carry                 |
| SUBCcc | 01 1100 | Subtract with Carry and modify cc's |

*Format (3)*



| Assembly Language Syntax |  |
|--------------------------|--|
| sub                      | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i> |
| subcc                    | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i> |
| subc                     | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i> |
| subccc                   | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i> |

*Description*

These instructions compute “ $r[rs1] - r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] - \text{sign\_ext}(simm13)$ ” if  $i = 1$ , and write the difference into  $r[rd]$ .

SUBC and SUBCcc (“subtract with carry”) also subtract the CCR register’s 32-bit carry ( $icc.c$ ) bit; that is, they compute “ $r[rs1] - r[rs2] - icc.c$ ” or “ $r[rs1] - \text{sign\_ext}(simm13) - icc.c$ ,” and write the difference into  $r[rd]$ .

SUBcc and SUBCcc modify the integer condition codes (CCR.icc and CCR.xcc). A 32-bit overflow (CCR.icc.v) occurs on subtraction if bit 31 (the sign) of the operands differs and bit 31 (the sign) of the difference differs from  $r[rs1] < 31 >$ . A 64-bit overflow (CCR.xcc.v) occurs on subtraction if bit 63 (the sign) of the operands differs and bit 63 (the sign) of the difference differs from  $r[rs1] < 63 >$ .

---

**Programming Note** – A SUBcc with  $rd = 0$  can be used to effect a signed or unsigned integer comparison.

SUBC and SUBCcc read the 32-bit condition codes' carry bit (CCR.icc.c), not the 64-bit condition codes' carry bit (CCR.xcc.c).

---

### Exceptions


None

---

## A.65 Tagged Add

| Opcode | op3     | Operation                  |
|--------|---------|----------------------------|
| TADDcc | 10 0000 | Tagged Add and modify cc's |

### Format (3)



| Assembly Language Syntax |                                     |
|--------------------------|-------------------------------------|
| taddcc                   | $reg_{rs1}, reg\_or\_imm, reg_{rd}$ |

### Description

This instruction computes a sum that is “ $r[rs1] + r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] + \text{sign\_ext}(\text{simm13})$ ” if  $i = 1$ .

TADDCC modifies the integer condition codes (*icc* and *xcc*).

A tag overflow condition occurs if bit 1 or bit 0 of either operand is nonzero or if the addition generates 32-bit arithmetic overflow (that is, both operands have the same value in bit 31 and the sum of bit 31 is different).

If a TADDCC causes a tag overflow, the 32-bit overflow bit (*CCR.icc.v*) is set to one; if TADDCC does not cause a tag overflow, *CCR.icc.v* is set to zero.

In either case, the remaining integer condition codes (both the other *CCR.icc* bits and all the *CCR.xcc* bits) are also updated as they would be for a normal ADD instruction. In particular, the setting of the *CCR.xcc.v* bit is not determined by the tag overflow condition (tag overflow is used only to set the 32-bit overflow bit). *CCR.xcc.v* is set only, based on the normal 64-bit arithmetic overflow condition, like a normal 64-bit add.

### Exceptions

None

---

## A.66 Tagged Subtract

| Opcode | op3     | Operation                       |
|--------|---------|---------------------------------|
| TSUBCC | 10 0001 | Tagged Subtract and modify cc's |

### Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| <code>tsubcc</code>      | <code>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></code> |

### *Description*

This instruction computes “ $r[rs1] - r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] - \text{sign\_ext}(\text{simml3})$ ” if  $i = 1$ .

TSUBCC modifies the integer condition codes (`icc` and `xcc`).

A tag overflow condition occurs if bit 1 or bit 0 of either operand is nonzero or if the subtraction generates 32-bit arithmetic overflow; that is, the operands have different values in bit 31 (the 32-bit sign bit) and the sign of the 32-bit difference in bit 31 differs from bit 31 of  $r[rs1]$ .

If a TSUBCC causes a tag overflow, the 32-bit overflow bit (`CCR.icc.v`) is set to one; if TSUBCC does not cause a tag overflow, `CCR.icc.v` is set to zero.

In either case, the remaining integer condition codes (both the other `CCR.icc` bits and all the `CCR.xcc` bits) are also updated as they would be for a normal subtract instruction. In particular, the setting of the `CCR.xcc.v` bit is not determined by the tag overflow condition (tag overflow is used only to set the 32-bit overflow bit). The `CCR.xcc.v` setting is based only on the normal 64-bit arithmetic overflow condition, like a normal 64-bit subtract.

### *Exceptions*

None

## A.67 Trap on Integer Condition Codes (Tcc)

| Opcode | op3     | cond | Operation   | icc Test                    |
|--------|---------|------|---|-----------------------------|
| TA     | 11 1010 | 1000 | Trap Always   | 1                           |
| TN     | 11 1010 | 0000 | Trap Never  | 0                           |
| TNE    | 11 1010 | 1001 | Trap on Not Equal                                     | not Z                       |
| TE     | 11 1010 | 0001 | Trap on Equal   | Z                           |
| TG     | 11 1010 | 1010 | Trap on Greater                                       | <b>not</b> (Z or (N xor V)) |
| TLE    | 11 1010 | 0010 | Trap on Less or Equal                                 | Z or (N xor V)              |
| TGE    | 11 1010 | 1011 | Trap on Greater or Equal                              | <b>not</b> (N xor V)        |
| TL     | 11 1010 | 0011 | Trap on Less  | N xor V                     |
| TGU    | 11 1010 | 1100 | Trap on Greater Unsigned                              | <b>not</b> (C or Z)         |
| TLEU   | 11 1010 | 0100 | Trap on Less or Equal Unsigned                        | (C or Z)                    |
| TCC    | 11 1010 | 1101 | Trap on Carry Clear (Greater than or Equal, Unsigned) | <b>not</b> C                |
| TCS    | 11 1010 | 0101 | Trap on Carry Set (Less Than, Unsigned)               | C                           |
| TPOS   | 11 1010 | 1110 | Trap on Positive or zero                              | <b>not</b> N                |
| TNEG   | 11 1010 | 0110 | Trap on Negative                                      | N                           |
| TVC    | 11 1010 | 1111 | Trap on Overflow Clear                                | <b>not</b> V                |
| TVS    | 11 1010 | 0111 | Trap on Overflow Set                                  | V                           |

*Format (4)*



| cc1 | cc0 | Condition Codes |
|-----|-----|-----------------|
| 00  |     | icc             |
| 01  |     | —               |
| 10  |     | xcc             |
| 11  |     | —               |

| Assembly Language Syntax |                                 |                 |
|--------------------------|---------------------------------|-----------------|
| ta                       | i_or_x_cc, software_trap_number |                 |
| tn                       | i_or_x_cc, software_trap_number |                 |
| tne                      | i_or_x_cc, software_trap_number | (synonym: tnz)  |
| te                       | i_or_x_cc, software_trap_number | (synonym: tz)   |
| tg                       | i_or_x_cc, software_trap_number |                 |
| tle                      | i_or_x_cc, software_trap_number |                 |
| tge                      | i_or_x_cc, software_trap_number |                 |
| tl                       | i_or_x_cc, software_trap_number |                 |
| tgu                      | i_or_x_cc, software_trap_number |                 |
| tleu                     | i_or_x_cc, software_trap_number |                 |
| tcc                      | i_or_x_cc, software_trap_number | (synonym: tgeu) |
| tcs                      | i_or_x_cc, software_trap_number | (synonym: tlu)  |
| tpos                     | i_or_x_cc, software_trap_number |                 |
| tneg                     | i_or_x_cc, software_trap_number |                 |
| tvc                      | i_or_x_cc, software_trap_number |                 |
| tvs                      | i_or_x_cc, software_trap_number |                 |

### Description

The Tcc instruction evaluates the selected integer condition codes (icc or xcc) according to the cond field of the instruction, producing either a TRUE or FALSE result. If TRUE and no higher-priority exceptions or interrupt requests are pending, then a *trap\_instruction* exception is generated. If FALSE, a *trap\_instruction* exception does not occur and the instruction behaves like a NOP.

The software trap number is specified by the least significant seven bits of “r[rs1] + r[rs2]” if i = 0, or the least significant seven bits of “r[rs1] + sw\_trap\_#” if i = 1.

When i = 1, bits 7 through 10 are reserved and should be supplied as zeroes by software. When i = 0, bits 5 through 10 are reserved, the most significant 57 bits of “r[rs1] + r[rs2]” are unused, and both should be supplied as zeroes by software.

### Description (Effect on Privileged State)

If a *trap\_instruction* traps, 256 plus the software trap number is written into TT[TL]. As described in Chapter 12 “Traps and Trap Handling” the trap is taken, and the processor performs the normal trap entry procedure.



---

**Programming Note** – `TCC` can be used to implement breakpointing, tracing, and calls to supervisor software. It can also be used for runtime checks, such as out-of-range array indexes, integer overflow, and so on.

---



---

**Compatibility Note** – `TCC` is upward compatible with the SPARC V8 `Ticc` instruction, with one qualification: a `Ticc` with `i = 1` and `simm13 < 0` may execute differently on a SPARC V9 processor. Use of the `i = 1` form of `Ticc` is believed to be rare in SPARC V8 software, and `simm13 < 0` is probably not used at all; therefore, it is believed in practice, that full software compatibility will be achieved.

---

### Exceptions

*trap\_instruction*

*illegal\_instruction* (`cc1`  $\square$  `cc0 = 012 or 112`, or reserved fields nonzero)

---

## A.68 Write Privileged Register

| Opcode                        | op3                  | Operation                 |
|-------------------------------|----------------------|---------------------------|
| <code>WRPR<sup>P</sup></code> | <code>11 0010</code> | Write Privileged Register |

### Format (3)

|          |    |       |       |          |        |       |
|----------|----|-------|-------|----------|--------|-------|
| 10       | rd | op3   | rs1   | i=0      | —      | rs2   |
| 10       | rd | op3   | rs1   | i=1      | simm13 |       |
| 31 30 29 |    | 25 24 | 19 18 | 14 13 12 |        | 5 4 0 |

| <b>rd</b> | <b>Privileged Register</b> |
|-----------|----------------------------|
| 0         | TPC                        |
| 1         | TNPC                       |
| 2         | TSTATE                     |
| 3         | TT                         |
| 4         | TICK                       |
| 5         | TBA                        |
| 6         | PSTATE                     |
| 7         | TL                         |
| 8         | PIL                        |
| 9         | CWP                        |
| 10        | CANSAVE                    |
| 11        | CANRESTORE                 |
| 12        | CLEANWIN                   |
| 13        | OTHERWIN                   |
| 14        | WSTATE                     |
| 15–31     | <i>Reserved</i>            |

| <b>Assembly Language Syntax</b> |   |
|---------------------------------|---|
| wrpr                            | <i>reg_rsl, reg_or_imm, %tpc</i>        |
| wrpr                            | <i>reg_rsl, reg_or_imm, %tnpc</i>       |
| wrpr                            | <i>reg_rsl, reg_or_imm, %tstate</i>     |
| wrpr                            | <i>reg_rsl, reg_or_imm, %tt</i>         |
| wrpr                            | <i>reg_rsl, reg_or_imm, %tick</i>       |
| wrpr                            | <i>reg_rsl, reg_or_imm, %tba</i>        |
| wrpr                            | <i>reg_rsl, reg_or_imm, %pstate</i>     |
| wrpr                            | <i>reg_rsl, reg_or_imm, %tl</i>         |
| wrpr                            | <i>reg_rsl, reg_or_imm, %pil</i>        |
| wrpr                            | <i>reg_rsl, reg_or_imm, %cwp</i>        |
| wrpr                            | <i>reg_rsl, reg_or_imm, %cansave</i>    |
| wrpr                            | <i>reg_rsl, reg_or_imm, %canrestore</i> |
| wrpr                            | <i>reg_rsl, reg_or_imm, %cleanwin</i>   |
| wrpr                            | <i>reg_rsl, reg_or_imm, %otherwin</i>   |
| wrpr                            | <i>reg_rsl, reg_or_imm, %wstate</i>     |

### *Description*

This instruction stores the value “ $r[rs1] \mathbf{xor} r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] \mathbf{xor} \mathit{sign\_ext}(simml3)$ ” if  $i = 1$ , to the writable fields of the specified privileged state register. **Note:** The operation is an exclusive OR.

The  $rd$  field in the instruction determines the privileged register that is written. There are at least four copies of the TPC, TNPC, TT, and TSTATE registers, one for each trap level. A write to one of these registers sets the register indexed by the current value in the trap level register (TL). A write to TPC, TNPC, TT, or TSTATE when the trap level is zero (TL = 0) causes an *illegal\_instruction* exception.

A WRPR of TL does not cause a trap or return from trap; it does not alter any other machine state.

---

**Programming Note** – A WRPR of TL can be used to read the values of TPC, TNPC, and TSTATE for any trap level; however, take care that traps do not occur while the TL register is modified.

---

The WRPR instruction is a *non*-delayed write instruction. The instruction immediately following the WRPR observes any changes made to processor state made by the WRPR.

WRPR instructions with  $rd$  in the range 15–31 are reserved for future versions of the architecture; executing a WRPR instruction with  $rd$  in that range causes an *illegal\_instruction* exception.

---

**Implementation Note** – Some WRPR instructions could serialize the processor in some implementations.

---

### *Exceptions*


*privileged\_opcode*

*illegal\_instruction* ( $(rd = 15-31)$  or  $((rd \leq 3) \text{ and } (TL = 0))$ )

## A.69 Write State Register

| Opcode                            | op3     | rd    | Operation  |
|-----------------------------------|---------|-------|--|
| WRY <sup>D</sup>                  | 11 0000 | 0     | Write Y register; deprecated (see Section A.70.18, “Write Y Register”).                    |
| —                                 | 11 0000 | 1     | <i>Reserved, do not access; attempt to access causes an illegal_instruction exception.</i> |
| WRCCR                             | 11 0000 | 2     | Write Condition Codes Register   |
| WRASI                             | 11 0000 | 3     | Write Graphics Status Register   |
| —                                 | 11 0000 | 4, 5  | <i>Reserved, do not access; attempt to access causes an illegal_instruction exception.</i> |
| WRFPRS                            | 11 0000 | 6     | Write Floating-Point Registers Status Register   |
| —                                 | 11 0000 | 7–14  | <i>Reserved, do not access; attempt to access causes an illegal_instruction exception.</i> |
| —                                 | 11 0000 | 15    | Software-initiated reset (see Section A.59, “Software-Initiated Reset”).                   |
| WRASR                             | 11 0000 | 16–31 | Write non-SPARC V9 ASRs  |
| WRPCR <sup>P</sup> <sub>PCR</sub> |         | 16    | Write Performance Control Registers (PCR)  |
| WRPIC <sup>P</sup> <sub>PIC</sub> |         | 17    | Write Performance Instrumentation Counters (PIC)   |
| WRDCR <sup>P</sup>                |         | 18    | Write Dispatch Control Register (DCR)  |
| WRGSR                             |         | 19    | Write Graphic Status Register (GSR)  |
| WRSOFTINT_SET <sup>P</sup>        |         | 20    | Set bits of per-processor Soft Interrupt Register  |
| WRSOFTINT_CLR <sup>P</sup>        |         | 21    | Clear bits of per-processor Soft Interrupt Register  |
| WRSOFTINT <sup>P</sup>            |         | 22    | Write per-processor Soft Interrupt Register  |
| WRTICK_CMPR <sup>P</sup>          |         | 23    | Write Tick Compare Register  |
| WRSTICK <sup>P</sup>              |         | 24    | Write System TICK Register   |
| WRSTICK_CMPR <sup>P</sup>         |         | 25    | Write System TICK Compare Register   |
| —                                 |         | 26–31 | <i>Reserved, do not access; attempt to access causes an illegal_instruction exception.</i> |

*Format (3)*



| Assembly Language Syntax |   |
|--------------------------|---|
| wr                       | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %ccr           |
| wr                       | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %asi           |
| wr                       | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %fprs          |
| wr                       | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %pcr           |
| wr                       | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %pic           |
| wr                       | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %dcr           |
| wr                       | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %gsr           |
| wr                       | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %set_softint   |
| wr                       | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %clear_softint |
| wr                       | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %softint       |
| wr                       | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %tick_cmpr     |
| wr                       | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %sys_tick      |
| wr                       | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , %sys_tick_cmpr |

### Description

These instructions store the value “*r[rs1] xor r[rs2]*” if *i* = 0, or “*r[rs1] xor sign\_ext(simm13)*” if *i* = 1, to the writable fields of the specified state register. **Note:** The operation is an exclusive OR.

WRASR writes a value to the ancillary state register (ASR) indicated by *rd*. The operation performed to generate the value written may be *rd* dependent or implementation dependent (see below). A WRASR instruction is indicated by *op* = 2, *rd* = ≥ 16, and *op3* = 30<sub>16</sub>.

The WRASR opcode for *rd* = 15, *rs1* = 0, and *i* = 1 is used for the software-initiated reset (SIR) instruction (see Section A.59, “Software-Initiated Reset”).

The WRCCR, WRFPRS, and WRASI instructions are *not* delayed-write instructions. The instruction immediately following a WRCCR, WRFPRS, or WRASIR observes the new value of the CCR, FPRS, or ASI register.

WRFPRS waits for any pending floating-point operations to complete before writing the FPRS register.

WRGSR causes a *fp\_disabled* trap if *PSTATE.PEF* = 0 or *FPRS.FEF* = 0.

WRPIC causes a *privileged\_action* exception if *PSTATE.PRIV* = 0 and *PCR.PRIV* = 1.

WRPCR causes a *privileged\_opcode* exception due to access privilege violation.

---

**Implementation Note** – Ancillary state registers may include, for example, timer, counter, diagnostic, self-test, and trap-control registers.

---

---

**Compatibility Note** – The SPARC V8 WRIER, WRPSR, WRWIM, and WRTBR instructions do not exist in SPARC V9 because the IER, PSR, TBR, and WIM registers do not exist in SPARC V9.

---

---

**Implementation Note** – Some WRASR instructions could serialize the processor in some implementations.

---

### *Exceptions*

*software\_initiated\_reset* ( $rd = 15$ ,  $rs1 = 0$ , and  $i = 1$  only)

*privileged\_opcode* (WRDCR, WRSOFTINT\_SET, WRSOFTINT\_CLR, WRSOFTINT, WRTICK\_CMPR, WRSTICK, WRSTICK\_CMPR, and WRPCR)

*illegal\_instruction* (WRASR with  $rd = 1, 4, 5, 7-14, 26-31$ ;

WRASR with  $rd = 15$  and  $rs1 \neq 0$  or  $i \neq 1$ )

*privileged\_action* (WRPIC with  $PSTATE.PRIV = 0$  and  $PCR.PRIV = 1$ )

*fp\_disabled* (WRGSR with  $PSTATE.PEF = 0$  or  $FPRS.FEF = 0$ )

---

## A.70 Deprecated Instructions

The following instructions are deprecated; they are provided only for compatibility with previous versions of the architecture. They should not be used in new SPARC V9 software. For each deprecated instruction, we recommend the instruction to be used instead. Please see TABLE A-2 for the page number at which you can find a description of the preferred instruction.

### A.70.1 Branch on Floating-Point Condition Codes (FBfcc)

The FBfcc instructions are deprecated. Use the FBPfcc instructions instead.

| Opcode             | cond | Operation                               | fcc Test    |
|--------------------|------|---|-------------|
| FBA <sup>D</sup>   | 1000 | Branch Always                           | 1           |
| FBN <sup>D</sup>   | 0000 | Branch Never                            | 0           |
| FBU <sup>D</sup>   | 0111 | Branch on Unordered                     | U           |
| FBG <sup>D</sup>   | 0110 | Branch on Greater                       | G           |
| FBUG <sup>D</sup>  | 0101 | Branch on Unordered or Greater          | G or U      |
| FBL <sup>D</sup>   | 0100 | Branch on Less                          | L           |
| FBUL <sup>D</sup>  | 0011 | Branch on Unordered or Less             | L or U      |
| FBLG <sup>D</sup>  | 0010 | Branch on Less or Greater               | L or G      |
| FBNE <sup>D</sup>  | 0001 | Branch on Not Equal                     | L or G or U |
| FBE <sup>D</sup>   | 1001 | Branch on Equal                         | E           |
| FBUE <sup>D</sup>  | 1010 | Branch on Unordered or Equal            | E or U      |
| FBGE <sup>D</sup>  | 1011 | Branch on Greater or Equal              | E or G      |
| FBUGE <sup>D</sup> | 1100 | Branch on Unordered or Greater or Equal | E or G or U |
| FBLE <sup>D</sup>  | 1101 | Branch on Less or Equal                 | E or L      |
| FBULE <sup>D</sup> | 1110 | Branch on Unordered or Less or Equal    | E or L or U |
| FBO <sup>D</sup>   | 1111 | Branch on Ordered                       | E or L or G |

## Format (2)

|    |    |      |     |        |    |    |    |   |
|----|----|------|-----|--------|----|----|----|---|
| 00 | a  | cond | 110 | disp22 |    |    |    |   |
| 31 | 30 | 29   | 28  | 25     | 24 | 22 | 21 | 0 |

| Assembly Language Syntax |       |                 |
|--------------------------|-------|-----------------|
| fba{ , a }               | label |                 |
| fbn{ , a }               | label |                 |
| fbu{ , a }               | label |                 |
| fbg{ , a }               | label |                 |
| fbug{ , a }              | label |                 |
| fb1{ , a }               | label |                 |
| fbu1{ , a }              | label |                 |
| fb1g{ , a }              | label |                 |
| fbne{ , a }              | label | (synonym: fbnz) |
| fbe{ , a }               | label | (synonym: fbz)  |
| fbue{ , a }              | label |                 |
| fbge{ , a }              | label |                 |
| fbuge{ , a }             | label |                 |
| fb1e{ , a }              | label |                 |
| fbule{ , a }             | label |                 |
| fbo{ , a }               | label |                 |

---

**Programming Note** – To set the annul bit for FBfCC instructions, append “ , a ” to the opcode mnemonic. For example, use “fb1 , a label.” In the preceding table, braces around “ , a ” signify that “ , a ” is optional.

---

### Description

Unconditional and FCC branches are described below:

- **Unconditional branches (FBA, FBN)** — If its annul field is zero, an FBN (Branch Never) instruction acts like a NOP. If its annul field is one, the following (delay) instruction is annulled (not executed) when the FBN is executed. In neither case does a transfer of control take place.



FBA (Branch Always) causes a PC-relative, delayed control transfer to the address “PC + (4 × sign\_ext(dispatch22)),” regardless of the value of the floating-point condition code bits. If the annul field of the branch instruction is one, the delay instruction is annulled (not executed). If the annul field is zero, the delay instruction is executed.

- **Fcc-conditional branches** — Conditional FBfcc instructions (except FBA and FBN) evaluate floating-point condition code zero (fcc0) according to the cond field of the instruction. Such evaluation produces either a TRUE or FALSE result. If TRUE, the branch is taken, that is, the instruction causes a PC-relative, delayed control transfer to the address “PC + (4 × sign\_ext(dispatch22)).” If FALSE, the branch is not taken.

If a conditional branch is taken, the delay instruction is always executed, regardless of the value of the annul field. If a conditional branch is not taken and the annul (a) field is one, the delay instruction is annulled (not executed).

---

**Note** – The annul bit has a *different* effect on conditional branches than it does on unconditional branches.

---

---

**Compatibility Note** – Unlike SPARC V8, SPARC V9 does not require an instruction between a floating-point compare operation and a floating-point branch (FBfcc, FBPFcc).

---

If FPRS.FEF = 0 or PSTATE.PEF = 0, or if an FPU is not present, the FBfcc instruction is not executed and instead generates a *fp\_disabled* exception.

*Exceptions*


*fp\_disabled*

## A.70.2 Branch on Integer Condition Codes (Bicc)

Use the BPcc instructions in place of BiCC instructions.

| Opcode            | cond | Operation   | icc Test                    |
|-------------------|------|---|-----------------------------|
| BA <sup>D</sup>   | 1000 | Branch Always   | 1                           |
| BN <sup>D</sup>   | 0000 | Branch Never  | 0                           |
| BNE <sup>D</sup>  | 1001 | Branch on Not Equal                                     | <b>not Z</b>                |
| BE <sup>D</sup>   | 0001 | Branch on Equal   | Z                           |
| BG <sup>D</sup>   | 1010 | Branch on Greater                                       | <b>not (Z or (N xor V))</b> |
| BLE <sup>D</sup>  | 0010 | Branch on Less or Equal                                 | Z or (N xor V)              |
| BGE <sup>D</sup>  | 1011 | Branch on Greater or Equal                              | <b>not (N xor V)</b>        |
| BL <sup>D</sup>   | 0011 | Branch on Less  | N xor V                     |
| BGU <sup>D</sup>  | 1100 | Branch on Greater Unsigned                              | <b>not (C or Z)</b>         |
| BLEU <sup>D</sup> | 0100 | Branch on Less or Equal Unsigned                        | C or Z                      |
| BCC <sup>D</sup>  | 1101 | Branch on Carry Clear (Greater Than or Equal, Unsigned) | <b>not C</b>                |
| BCS <sup>D</sup>  | 0101 | Branch on Carry Set (Less Than, Unsigned)               | C                           |
| BPOS <sup>D</sup> | 1110 | Branch on Positive                                      | <b>not N</b>                |
| BNEG <sup>D</sup> | 0110 | Branch on Negative                                      | N                           |
| BVC <sup>D</sup>  | 1111 | Branch on Overflow Clear                                | <b>not V</b>                |
| BVS <sup>D</sup>  | 0111 | Branch on Overflow Set                                  | V                           |

*Format (2)*



| Assembly Language Syntax |       |                 |
|--------------------------|-------|-----------------|
| ba{ , a }                | label |                 |
| bn{ , a }                | label |                 |
| bne{ , a }               | label | (synonym: bnz)  |
| be{ , a }                | label | (synonym: bz)   |
| bg{ , a }                | label |                 |
| ble{ , a }               | label |                 |
| bge{ , a }               | label |                 |
| bl{ , a }                | label |                 |
| bgu{ , a }               | label |                 |
| bleu{ , a }              | label |                 |
| bcc{ , a }               | label | (synonym: bgeu) |
| bcs{ , a }               | label | (synonym: blu)  |
| bpos{ , a }              | label |                 |
| bneg{ , a }              | label |                 |
| bvc{ , a }               | label |                 |
| bvs{ , a }               | label |                 |

---

**Programming Note** – To set the annul bit for `Bicc` instructions, append “, a” to the opcode mnemonic. For example, use “`bgu , a label`.” In the preceding table, braces signify that the “, a” is optional.

---

### *Description*

Unconditional branches and `icc`-conditional branches are described below:

- **Unconditional branches (BA, BN)** — If its annul field is zero, a BN (Branch Never) instruction is treated as a NOP. If its annul field is one, the following (delay) instruction is annulled (not executed). In neither case does a transfer of control take place.

BA (Branch Always) causes an unconditional PC-relative, delayed control transfer to the address “`PC + (4 × sign_ext (disp22))`.” If the annul field of the branch instruction is one, the delay instruction is annulled (not executed). If the annul field is zero, the delay instruction is executed.

- **Icc-conditional branches** — Conditional `Bicc` instructions (all except BA and BN) evaluate the 32-bit integer condition codes (`icc`), according to the `cond` field of the instruction, producing either a TRUE or FALSE result. If TRUE, the branch is taken, that is, the instruction causes a PC-relative, delayed control transfer to the address “`PC + (4 × sign_ext (disp22))`.” If FALSE, the branch is not taken.

If a conditional branch is taken, the delay instruction is always executed regardless of the value of the annul field. If a conditional branch is not taken and the annul (a) field is one, the delay instruction is annulled (not executed).

---

**Note** – The annul bit has a *different* effect on conditional branches than it does on unconditional branches.

---

### Exceptions


None

## A.70.3 Divide (64-bit/32-bit)

The UDIV, UDIV<sub>cc</sub>, SDIV, and SDIV<sub>cc</sub> instructions are deprecated. Use the UDIVX and SDIVX instructions instead.

| Opcode                          | op3     | Operation                               |
|---------------------------------|---------|---|
| UDIV <sup>D</sup>               | 00 1110 | Unsigned Integer Divide                 |
| SDIV <sup>D</sup>               | 00 1111 | Signed Integer Divide                   |
| UDIV <sub>cc</sub> <sup>D</sup> | 01 1110 | Unsigned Integer Divide and modify cc's |
| SDIV <sub>cc</sub> <sup>D</sup> | 01 1111 | Signed Integer Divide and modify cc's   |

### Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| udiv                     | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| sdiv                     | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| udiv <sub>cc</sub>       | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| sdiv <sub>cc</sub>       | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |

### Description

The divide instructions perform 64-bit by 32-bit division, producing a 32-bit result. If  $i = 0$ , they compute “ $(Y \llcorner r[rs1] \langle 31:0 \rangle) \div r[rs2] \langle 31:0 \rangle$ .” Otherwise (that is, if  $i = 1$ ), the divide instructions compute “ $(Y \llcorner r[rs1] \langle 31:0 \rangle) \div (\text{sign\_ext}(\text{simm13}) \langle 31:0 \rangle)$ .” In either case, if overflow does not occur, the less significant 32 bits of the integer quotient are sign-extended or zero-extended to 64 bits and are written into  $r[rd]$ .

The contents of the Y register are undefined after any 64-bit by 32-bit integer divide operation.

### Unsigned Divide

Unsigned divide (UDIV, UDIVCC) assumes an unsigned integer doubleword dividend ( $Y \llcorner r[rs1] \langle 31:0 \rangle$ ) and an unsigned integer word divisor  $r[rs2 \langle 31:0 \rangle]$  or ( $\text{sign\_ext}(\text{simm13}) \langle 31:0 \rangle$ ) and computes an unsigned integer word quotient ( $r[rd]$ ). Immediate values in `simm13` are in the ranges 0 to  $2^{12} - 1$  and  $2^{32} - 2^{12}$  to  $2^{32} - 1$  for unsigned divide instructions.

Unsigned division rounds an inexact rational quotient toward zero.

In the UltraSPARC III Cu processor, LDD is implemented in hardware.

---

**Programming Note** – The *rational quotient* is the infinitely precise result quotient. It includes both the integer part and the fractional part of the result. For example, the rational quotient of  $11/4 = 2.75$  (integer part = 2, fractional part = .75).

---

The result of an unsigned divide instruction can overflow the less significant 32 bits of the destination register  $r[rd]$  under certain conditions. When overflow occurs, the largest appropriate unsigned integer is returned as the quotient in  $r[rd]$ . The condition under which overflow occurs and the value returned in  $r[rd]$  under this condition are specified in TABLE A-16.

**TABLE A-16** UDIV / UDIVCC Overflow Detection and Value Returned

| Condition Under Which Overflow Occurs | Value Returned in $r[rd]$                            |
|---------------------------------------|--|
| Rational quotient $\geq 2^{32}$       | $2^{32} - 1$<br>(0000 0000 FFFF FFFF <sub>16</sub> ) |

When no overflow occurs, the 32-bit result is zero-extended to 64 bits and written into register  $r[rd]$ .

UDIV does not affect the condition code bits. UDIVCC writes the integer condition code bits as shown in the following table. Note that negative (N) and zero (Z) are set according to the value of  $r[rd]$  after it has been set to reflect overflow, if any.

| Bit   | UDIVcc                           |
|-------|----------------------------------|
| icc.N | Set if $r[rd] < 31 > = 1$        |
| icc.Z | Set if $r[rd] < 31:0 > = 0$      |
| icc.V | Set if overflow (per TABLE A-16) |
| icc.C | Zero                             |
| xcc.N | Set if $r[rd] < 63 > = 1$        |
| xcc.Z | Set if $r[rd] < 63:0 > = 0$      |
| xcc.V | Zero                             |
| xcc.C | Zero                             |

### Signed Divide

Signed divide (SDIV, SDIVCC) assumes a signed integer doubleword dividend ( $Y \square$  lower 32 bits of  $r[rs1]$ ) and a signed integer word divisor (lower 32 bits of  $r[rs2]$  or lower 32 bits of  $sign\_ext(signmm13)$ ) and computes a signed integer word quotient ( $r[rd]$ ).

Signed division rounds an inexact quotient toward zero. For example,  $-7 \div 4$  equals the rational quotient of  $-1.75$ , which rounds to  $-1$  (not  $-2$ ) when rounding toward zero.

The result of a signed divide can overflow the low-order 32 bits of the destination register  $r[rd]$  under certain conditions. When overflow occurs, the largest appropriate signed integer is returned as the quotient in  $r[rd]$ . The conditions under which overflow occurs and the value returned in  $r[rd]$  under those conditions are specified in TABLE A-17.

**TABLE A-17** SDIV / SDIVCC Overflow Detection and Value Returned

| Condition Under Which Overflow Occurs | Value Returned in $r[rd]$                            |
|---------------------------------------|--|
| Rational quotient $\geq 2^{31}$       | $2^{31} - 1$<br>(0000 0000 7FFF FFFF <sub>16</sub> ) |
| Rational quotient $\leq -2^{31} - 1$  | $-2^{31}$<br>(FFFF FFFF 8000 0000 <sub>16</sub> )    |

When no overflow occurs, the 32-bit result is sign-extended to 64 bits and written into register  $r[rd]$ .

SDIV does not affect the condition code bits. SDIVCC writes the integer condition code bits as shown in the following table. Note that negative (N) and zero (Z) are set according to the value of  $r[rd]$  after it has been set to reflect overflow, if any.

| Bit   | SDIVcc                           |
|-------|----------------------------------|
| icc.N | Set if $r[rd]<31> = 1$           |
| icc.Z | Set if $r[rd]<31:0> = 0$         |
| icc.V | Set if overflow (per TABLE A-17) |
| icc.C | Zero                             |
| xcc.N | Set if $r[rd]<63> = 1$           |
| xcc.Z | Set if $r[rd]<63:0> = 0$         |
| xcc.V | Zero                             |
| xcc.C | Zero                             |

### Exceptions

*division\_by\_zero*

## A.70.4 Load Floating-Point Status Register

The LDFSR instruction is deprecated. Use the LDXFSSR instruction instead.

| Opcode             | op3     | rd | Operation                                |
|--------------------|---------|----|--|
| LDFSR <sup>D</sup> | 10 0001 | 0  | Load Floating-Point State Register Lower |

### Format (3)

|          |       |       |          |     |        |     |
|----------|-------|-------|----------|-----|--------|-----|
| 11       | rd    | op3   | rs1      | i=0 | —      | rs2 |
| 11       | rd    | op3   | rs1      | i=1 | simm13 |     |
| 31 30 29 | 25 24 | 19 18 | 14 13 12 | 5 4 | 0      |     |

| Assembly Language Syntax |                  |
|--------------------------|------------------|
| ld                       | [address], %fscr |

### *Description*

The load floating-point state register lower instruction (LDFSR) waits for all FPop instructions that have not finished execution to complete and then loads a word from memory into the less significant 32 bits of the FSR. The upper 32 bits of FSR are unaffected by LDFSR.

LDFSR causes a *mem\_address\_not\_aligned* exception if the effective memory address is not word aligned.

---

**Compatibility Note** – SPARC V9 supports two different instructions to load the FSR: the SPARC V8 LDFSR instruction is defined to load only the less significant 32 bits of the FSR, whereas LDXFSR allows SPARC V9 programs to load all 64 bits of the FSR.

---

### *Exceptions*

*mem\_address\_not\_aligned*  
*data\_access\_exception*  
*data\_access\_error*  
*fast\_data\_access\_MMU\_miss*  
*fast\_data\_access\_protection*  
*PA\_watchpoint*  
*VA\_watchpoint*

## A.70.5 Load Integer Doubleword


The LDD instruction is deprecated; it is provided only for compatibility with previous versions of the architecture. It should not be used in new SPARC V9 software. Use the LDX instruction instead.

Please refer to Section A.27 “Load Integer” for the current load integer instructions.

| Opcode           | op3     | Operation       |
|------------------|---------|-----------------|
| LDD <sup>D</sup> | 00 0011 | Load doubleword |



### Format (3)



| Assembly Language Syntax |   |
|--------------------------|---|
| l <code>dd</code>        | [ <i>address</i> ], <i>reg<sub>rd</sub></i> |

#### Description

The load doubleword integer instruction (LDD) copies a doubleword from memory into an *r* register pair. The word at the effective memory address is copied into the even *r* register. The word at the effective memory address + 4 is copied into the following odd numbered *r* register. The upper 32 bits of both the even numbered and odd numbered *r* registers are zero-filled.

---

**Notes** – A load doubleword with *rd* = 0 modifies only *r*[1]. The least significant bit of the *rd* field in an LDD instruction is unused and should be set to zero by software. An attempt to execute a load doubleword instruction that refers to a misaligned (odd numbered) destination register causes an *illegal\_instruction* exception.

With respect to little-endian memory, an LDD instruction behaves as if it is composed of two 32-bit loads, each of which is byte swapped independently before being written into each destination register.

---

Load integer doubleword instructions access the primary address space (ASI = 80<sub>16</sub>). The effective address is “*r*[*rs1*] + *r*[*rs2*]” if *i* = 0, or “*r*[*rs1*] + *sign\_ext*(*simm13*)” if *i* = 1.

A successful load doubleword instruction operates atomically.

---

**Programming Note** – LDD is provided for compatibility with SPARC V8. It may execute slowly on SPARC V9 machines because of data path and register-access difficulties.

---

*Exceptions*

*illegal\_instruction* (LDD with odd rd)  
*mem\_address\_not\_aligned*  
*data\_access\_exception*  
*data\_access\_error*  
*fast\_data\_access\_MMU\_miss*  
*fast\_data\_access\_protection*  
*PA\_watchpoint*  
*VA\_watchpoint*

## A.70.6 Load Integer Doubleword from Alternate Space

The LDDA instruction is deprecated. Use the LDXA instruction in its place.

Please refer to Section A.28 “Load Integer from Alternate Space” for current load integer from alternate space instructions.

| Opcode                  | op3     | Operation                            |
|-------------------------|---------|--------------------------------------|
| LDDA <sup>D, PAsi</sup> | 01 0011 | Load Doubleword from Alternate Space |

*Format (3)*



| Assembly Language Syntax |  |
|--------------------------|--|
| ldda                     | [regaddr] imm_asi, reg <sub>rd</sub>   |
| ldda                     | [reg_plus_imm] %asi, reg <sub>rd</sub> |

## Description

The load doubleword integer from alternate space instruction (LDDA) copies a doubleword from memory into an  $r$  register pair. The word at the effective memory address is copied into the even  $r$  register. The word at the effective memory address + 4 is copied into the following odd numbered  $r$  register. The upper 32 bits of both the even numbered and odd numbered  $r$  registers are zero-filled.

---

**Notes** – A load doubleword with  $rd = 0$  modifies only  $r[1]$ . The least significant bit of the  $rd$  field in an LDDA instruction is unused and should be set to zero by software. An attempt to execute a load doubleword instruction that refers to a misaligned (odd numbered) destination register causes an *illegal\_instruction* exception.

With respect to little-endian memory, an LDDA instruction behaves as if it is composed of two 32-bit loads, each of which is byte-swapped independently before being written into each destination register.

---

The load integer doubleword from alternate space instructions contain the address space identifier (ASI) to be used for the load in the `imm_asi` field if  $i = 0$ , or in the ASI register if  $i = 1$ . The access is privileged if bit 7 of the ASI is zero; otherwise, it is not privileged. The effective address for these instructions is “ $r[rs1] + r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] + \text{sign\_ext}(simm13)$ ” if  $i = 1$ .

A successful load doubleword instruction operates atomically.

LDDA causes a *mem\_address\_not\_aligned* exception if the address is not doubleword aligned.

These instructions cause a *privileged\_action* exception if `PSTATE.PRIV = 0` and bit 7 of the ASI is zero.

In the UltraSPARC III Cu processor, LDDA is implemented in hardware.

LDDA with  $ASI=24_{16}$  or  $2C_{16}$  is defined to be a Load Quadword Atomic instruction. See Section A.29 “Load Quadword, Atomic (VIS I)” for details.

---

**Programming Note** – LDDA is provided for compatibility with SPARC V8. It may execute slowly on SPARC V9 machines because of data path and register-access difficulties.

If LDDA is emulated in software, an LDXA instruction should be used for the memory access in order to preserve atomicity.

---

## Exceptions

*privileged\_action*

*illegal\_instruction* (LDDA with odd  $rd$ )

*mem\_address\_not\_aligned*

*data\_access\_exception*  
*fast\_data\_access\_MMU\_miss*  
*fast\_data\_access\_protection*  
*PA\_watchpoint*  
*VA\_watchpoint*

## A.70.7 Multiply (32-bit)

The UMUL, UMULCC, SMUL, and SMULCC instructions are deprecated. Use the MULX instruction instead.

| Opcode              | op3     | Operation                                 |
|---------------------|---------|---|
| UMUL <sup>D</sup>   | 00 1010 | Unsigned Integer Multiply                 |
| SMUL <sup>D</sup>   | 00 1011 | Signed Integer Multiply                   |
| UMULCC <sup>D</sup> | 01 1010 | Unsigned Integer Multiply and modify cc's |
| SMULCC <sup>D</sup> | 01 1011 | Signed Integer Multiply and modify cc's   |

*Format (3)*



| Assembly Language Syntax |  |
|--------------------------|--|
| umul                     | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| smul                     | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| umulcc                   | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |
| smulcc                   | <i>reg<sub>rs1</sub>, reg_or_imm, reg<sub>rd</sub></i> |

*Description*

The multiply instructions perform 32-bit by 32-bit multiplications, producing 64-bit results. They compute “ $r[rs1]<31:0> \times r[rs2]<31:0>$ ” if  $i = 0$ , or “ $r[rs1]<31:0> \times \text{sign\_ext}(simm13)<31:0>$ ” if  $i = 1$ . They write the 32 most significant bits of the product into the Y register and all 64 bits of the product into  $r[rd]$ .

Unsigned multiply instructions (UMUL, UMULCC) operate on unsigned integer word operands and compute an unsigned integer doubleword product. Signed multiply instructions (SMUL, SMULCC) operate on signed integer word operands and compute a signed integer doubleword product.

UMUL and SMUL do not affect the condition code bits. UMULCC and SMULCC write the integer condition code bits, ICC and XCC, as shown in TABLE A-18. **Note:** 32-bit negative (ICC.N) and zero (ICC.Z) condition codes are set according to the *less* significant word of the product, not according to the full 64-bit result.

**TABLE A-18** UMULCC / SMULCC Condition Code Settings

| Bit   | UMULCC / SMULCC          |
|-------|--------------------------|
| ICC.N | Set if product<31> = 1   |
| ICC.Z | Set if product<31:0>= 0  |
| ICC.V | 0                        |
| ICC.C | 0                        |
| XCC.N | Set if product<63> = 1   |
| XCC.Z | Set if product<63:0> = 0 |
| XCC.V | 0                        |
| XCC.C | 0                        |

---

**Programming Note** – 32-bit overflow after UMUL/UMULCC is indicated by Y ≠ 0.

Thirty-two bit overflow after SMUL/SMULCC is indicated by Y ≠ (r[rd] >> 31), where “>>” indicates 32-bit arithmetic right-shift.

---

### *Exceptions*


None

## A.70.8 Multiply Step

The MULSCC instruction is deprecated. Use the MULX instruction instead.

| Opcode              | op3     | Operation                     |
|---------------------|---------|-------------------------------|
| MULSCC <sup>D</sup> | 10 0100 | Multiply Step and modify cc's |

### Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| mul <code>scc</code>     | <code>reg<sub>rs1</sub>, reg<sub>or_imm</sub>, reg<sub>rd</sub></code> |

### Description

MUL`scc` treats the less significant 32 bits of both `r[rs1]` and the `Y` register as a single 64-bit, right-shiftable doubleword register. The least significant bit of `r[rs1]` is treated as if it were adjacent to bit 31 of the `Y` register. The MUL`scc` instruction adds, based on the least significant bit of `Y`.

Multiplication assumes that the `Y` register initially contains the multiplier, `r[rs1]` contains the most significant bits of the product, and `r[rs2]` contains the multiplicand. Upon completion of the multiplication, the `Y` register contains the least significant bits of the product.

**Note:** A standard MUL`scc` instruction has `rs1 = rd`.

MUL`scc` operates as follows:

1. The multiplicand is `r[rs2]` if `i = 0`, or `sign_ext(simm13)` if `i = 1`.
2. A 32-bit value is computed by shifting `r[rs1]` right by one bit with “`CCR.icc.n xor CCR.icc.v`” replacing bit 31 of `r[rs1]`. (This is the proper sign for the previous partial product).
3. If the least significant bit of `Y = 1`, the shifted value from step (2) and the multiplicand are added. If the least significant bit of `Y = 0`, then zero is added to the shifted value from step (2).
4. The sum from step (3) is written into `r[rd]`. The upper 32 bits of `r[rd]` are undefined. The integer condition codes are updated according to the addition performed in step (3). The values of the extended condition codes are undefined.
5. The `Y` register is shifted right by one bit, with the least significant bit of the unshifted `r[rs1]` replacing bit 31 of `Y`.

### Exceptions


None

## A.70.9 Read Y Register

The RDY instruction from the Read State Register instructions (Section A.50 “Read State Register”) is deprecated. We recommend that all instructions that reference the Y register be avoided.

| Opcode           | op3     | rs1 | Operation       |
|------------------|---------|-----|-----------------|
| RDY <sup>D</sup> | 10 1000 | 0   | Read Y Register |

### Format (3)



| Assembly Language Syntax |                       |
|--------------------------|-----------------------|
| rd                       | %Y, reg <sub>rd</sub> |

### Description

This instruction reads the Y register into r[rd].

### Exceptions


None

## A.70.10 Store Barrier

The STBAR instruction is deprecated. Use the MEMBAR instruction instead.

| Opcode             | op3     | Operation     |
|--------------------|---------|---------------|
| STBAR <sup>D</sup> | 10 1000 | Store Barrier |

### Format (3)



| Assembly Language Syntax |
|--------------------------|
| stbar                    |

### Description

The store barrier instruction (STBAR) forces *all* store and atomic load-store operations issued by a processor prior to the STBAR to complete their effects on memory before *any* store or atomic load-store operations issued by that processor subsequent to the STBAR are executed by memory.

**Note:** The encoding of STBAR is identical to that of the RDASR instruction except that  $rs1 = 15$  and  $rd = 0$ , and it is identical to that of the MEMBAR instruction except that bit 13 ( $i$ ) = 0.

---

**Compatibility Note** – STBAR is identical in function to a MEMBAR instruction with  $mask = 8_{16}$ . STBAR is retained for compatibility with SPARC V8.

---



---

**Implementation Note** – For correctness, it is sufficient for a processor to stop issuing new store and atomic load-store operations when an STBAR is encountered and to resume after all stores have completed and are observed in memory by all processors. More efficient implementations may take advantage of the fact that the processor is allowed to issue store and load-store operations after the STBAR, as long as those operations are guaranteed not to become visible before all the earlier stores and atomic load-stores have become visible to all processors.

---

### Exceptions

None



## A.70.11 Store Floating-Point Status Register Lower

The STFSR instruction is deprecated. Use the STXFSR instruction instead.

| Opcode             | op3     | rd | Operation                                 |
|--------------------|---------|----|---|
| STFSR <sup>D</sup> | 10 0101 | 0  | Store Floating-Point State Register Lower |

*Format (3)*



| Assembly Language Syntax |                         |
|--------------------------|-------------------------|
| st                       | %f $sr$ , [ $address$ ] |

*Description*

The store floating-point state register lower instruction (STFSR) waits for any currently executing FPop instructions to complete, and then it writes the less significant 32 bits of the FSR into memory.

---

**Compatibility Note** – SPARC V9 needs two store-FSR instructions, since the SPARC V8 STFSR instruction is defined to store only 32 bits of the FSR into memory. STXFSR allows SPARC V9 programs to store all 64 bits of the FSR.

---

STFSR zeroes FSR.ftt after writing the FSR to memory.

---

**Implementation Note** – FSR.ftt should not be zeroed until it is known that the store will not cause a precise trap.

---

The effective address for this instruction is “ $r[rs1] + r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] + sign\_ext(simm13)$ ” if  $i = 1$ .

STFSR causes a *mem\_address\_not\_aligned* exception if the effective memory address is not word aligned.

## Exceptions

*illegal\_instruction* (op3 = 25<sub>16</sub> and rd = 2–31)

*fp\_disabled*

*mem\_address\_not\_aligned*

*data\_access\_exception*

*data\_access\_error*

*fast\_data\_access\_MMU\_miss*

*fast\_data\_access\_protection*

*PA\_watchpoint*


*VA\_watchpoint*

## A.70.12 Store Integer Doubleword

The STD instruction is deprecated. Use the STX instruction instead.

| Opcode           | op3     | Operation        |
|------------------|---------|------------------|
| STD <sup>D</sup> | 00 0111 | Store Doubleword |

### Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| std                      | <i>reg<sub>rd</sub></i> [ <i>address</i> ] |

### Description

The store doubleword integer instruction (STD) copies two words from an *r* register pair into memory. The least significant 32 bits of the even numbered *r* register are written into memory at the effective address, and the least significant 32 bits of the following odd numbered *r* register are written into memory at the “effective address + 4.” The least significant bit of the *rd* field of a store doubleword instruction is unused and should always be set to zero by software. An attempt to execute a store doubleword instruction that refers to a misaligned (odd numbered) *rd* causes an *illegal\_instruction* exception.

The effective address for this instruction is “ $r[rs1] + r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] + \text{sign\_ext}(\text{simml3})$ ” if  $i = 1$ .

A successful store doubleword instruction operates atomically.

STD causes a *mem\_address\_not\_aligned* exception if the effective address is not doubleword aligned.

In the UltraSPARC III Cu processor, STD is implemented in hardware.

---

**Programming Note** – STD is provided for compatibility with SPARC V8. It may execute slowly on SPARC V9 machines because of data path and register-access difficulties. Therefore, programmers should avoid using STD.

If STD is emulated in software, STX should be used to preserve atomicity.

---

With respect to little-endian memory, a STD instruction behaves as if it is composed of two 32-bit stores, each of which is byte-swapped independently before being written into each destination memory word.

### *Exceptions*

*illegal\_instruction* (STD with odd  $rd$ )  
*mem\_address\_not\_aligned* (all except STB)  
*data\_access\_exception*  
*data\_access\_error*  
*fast\_data\_access\_MMU\_miss*  
*fast\_data\_access\_protection*  
*PA\_watchpoint*  
*VA\_watchpoint*

## A.70.13 Store Integer Doubleword into Alternate Space

The STDA instruction is deprecated. Instead, use the STXA instruction.

| Opcode                   | op3     | Operation                             |
|--------------------------|---------|---------------------------------------|
| STDA <sup>D</sup> , PASI | 01 0111 | Store Doubleword into Alternate Space |

### Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| stda                     | <i>reg<sub>rd</sub></i> , [ <i>reg_plus_imm</i> ] %asi |

#### Description

The store doubleword integer instruction (STDA) copies two words from an *r* register pair into memory. The least significant 32 bits of the even numbered *r* register are written into memory at the effective address, and the least significant 32 bits of the following odd numbered *r* register are written into memory at the “effective address + 4.” The least significant bit of the *rd* field of a store doubleword instruction is unused and should always be set to zero by software. An attempt to execute a store doubleword instruction that refers to a misaligned (odd numbered) *rd* causes an *illegal\_instruction* exception.

Store integer doubleword to alternate space instructions contain the address space identifier (ASI) to be used for the store in the *imm\_asi* field if *i* = 0, or in the ASI register if *i* = 1. The access is privileged if bit 7 of the ASI is zero; otherwise, it is not privileged. The effective address for these instructions is “*r*[*rs1*] + *r*[*rs2*]” if *i* = 0, or “*r*[*rs1*] + *sign\_ext*(*simm13*)” if *i* = 1.

A successful store doubleword instruction operates atomically.

STDA causes a *mem\_address\_not\_aligned* exception if the effective address is not doubleword aligned.

A store integer into alternate space instruction causes a *privileged\_action* exception if *PSTATE.PRIV* = 0 and bit 7 of the ASI is zero.

In the UltraSPARC III Cu processor, STDA is implemented in hardware.

---

**Programming Note** – STDA is provided for compatibility with SPARC V8. It may execute slowly on SPARC V9 machines because of data path and register-access difficulties. Therefore, programmers should avoid using STDA.

---

*Exceptions*

*illegal\_instruction* (STDA with odd rd)  
*privileged\_action*  
*mem\_address\_not\_aligned*  
*data\_access\_exception*  
*data\_access\_error*  
*fast\_data\_access\_MMU\_miss*  
*fast\_data\_access\_protection*  
*PA\_watchpoint*  
*VA\_watchpoint*

## A.70.14 Swap Register with Memory

The SWAP instruction is deprecated. Use the CASA or CASXA instruction in its place.

| Opcode            | op3     | Operation                 |
|-------------------|---------|---------------------------|
| SWAP <sup>D</sup> | 00 1111 | Swap Register with Memory |

*Format (3)*



| Assembly Language Syntax |                              |
|--------------------------|------------------------------|
| swap                     | [address], reg <sub>rd</sub> |

## *Description*

SWAP exchanges the less significant 32 bits of  $r[rd]$  with the contents of the word at the addressed memory location. The upper 32 bits of  $r[rd]$  are set to zero. The operation is performed atomically, that is, without allowing intervening interrupts or deferred traps. In a multiprocessor system, two or more processors executing CASA, CASXA, SWAP, SWAPA, LDSTUB, or LDSTUBA instructions addressing any or all of the same doubleword simultaneously are guaranteed to execute them in an undefined, but serial order.

The effective address for these instructions is “ $r[rs1] + r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] + \text{sign\_ext}(\text{simml3})$ ” if  $i = 1$ . This instruction causes a *mem\_address\_not\_aligned* exception if the effective address is not word aligned.

The coherence and atomicity of memory operations between processors and I/O DMA memory accesses are implementation dependent.

---

**Implementation Note** – See *Implementation Characteristics of Current SPARC-V9-based Products, Revision 9.x*, a document available from SPARC International, for information on the presence of hardware support for these instructions in the various SPARC V9 implementations.

---

## *Exceptions*

*mem\_address\_not\_aligned*  
*data\_access\_exception*  
*data\_access\_error*  
*fast\_data\_access\_MMU\_miss*  
*fast\_data\_access\_protection*  
*PA\_watchpoint*  
*VA\_watchpoint*

## A.70.15 Swap Register with Alternate Space Memory

The SWAPA instruction is deprecated. Use the CASXA instruction instead.

| Opcode                   | op3     | Operation                                 |
|--------------------------|---------|---|
| SWAPA <sup>D, PAsI</sup> | 01 1111 | Swap register with Alternate Space Memory |

### Format (3)

|    |    |     |     |     |         |     |    |    |    |   |   |   |
|----|----|-----|-----|-----|---------|-----|----|----|----|---|---|---|
| 11 | rd | op3 | rs1 | i=0 | imm_asi | rs2 |    |    |    |   |   |   |
| 11 | rd | op3 | rs1 | i=1 | simm13  |     |    |    |    |   |   |   |
| 31 | 30 | 29  | 25  | 24  | 19      | 18  | 14 | 13 | 12 | 5 | 4 | 0 |

| Assembly Language Syntax |  |
|--------------------------|--|
| swapa                    | [regaddr] imm_asi, reg <sub>rd</sub>   |
| swapa                    | [reg_plus_imm] %asi, reg <sub>rd</sub> |

### Description

SWAPA exchanges the less significant 32 bits of  $r[rd]$  with the contents of the word at the addressed memory location. The upper 32 bits of  $r[rd]$  are set to zero. The operation is performed atomically, that is, without allowing intervening interrupts or deferred traps. In a multiprocessor system, two or more processors executing CASA, CASXA, SWAP, SWAPA, LDSTUB, or LDSTUBA instructions addressing any or all of the same doubleword simultaneously are guaranteed to execute them in an undefined, but serial order.

The SWAPA instruction contains the address space identifier (ASI) to be used for the load in the `imm_asi` field if  $i = 0$ , or in the ASI register if  $i = 1$ . The access is privileged if bit 7 of the ASI is zero; otherwise, it is not privileged. The effective address for this instruction is “ $r[rs1] + r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] + \text{sign\_ext}(simm13)$ ” if  $i = 1$ .

This instruction causes a *mem\_address\_not\_aligned* exception if the effective address is not word aligned. It causes a *privileged\_action* exception if `PSTATE.PRIV = 0` and bit 7 of the ASI is zero.

The coherence and atomicity of memory operations between processors and I/O DMA memory accesses are implementation dependent.

---

**Implementation Note** – See *Implementation Characteristics of Current SPARC-V9-based Products, Revision 9.x*, a document available from SPARC International, for information on the presence of hardware support for this instruction in the various SPARC V9 implementations.

---

## Exceptions

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

## A.70.16 Tagged Add and Trap on Overflow

The TADDccTV instruction is deprecated. Use the TADDcc followed by BPVS instruction (with instructions to save the pre-TADDcc integer condition codes if necessary).

| Opcode                | op3     | Operation                                       |
|-----------------------|---------|---|
| TADDccTV <sup>D</sup> | 10 0010 | Tagged Add and modify cc's, or Trap on Overflow |

### Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| taddcctv                 | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i> |

### Description

This instruction computes a sum that is “ $r[rs1] + r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] + \text{sign\_ext}(simm13)$ ” if  $i = 1$ .

TADDccTV modifies the integer condition codes if it does not trap.

A *tag\_overflow* exception occurs if bit 1 or bit 0 of either operand is nonzero or if the addition generates 32-bit arithmetic overflow (that is, both operands have the same value in bit 31 and the sum of bit 31 is different).



If TADDccTV causes a tag overflow, a *tag\_overflow* exception is generated and  $r[rd]$  and the integer condition codes remain unchanged. If a TADDccTV does not cause a tag overflow, the sum is written into  $r[rd]$  and the integer condition codes are updated. CCR.icc.v is set to zero to indicate no 32-bit overflow.

In either case, the remaining integer condition codes (both the other CCR.icc bits and all the CCR.xcc bits) are also updated as they would be for a normal ADD instruction. In particular, the setting of the CCR.xcc.v bit is not determined by the tag overflow condition (tag overflow is used only to set the 32-bit overflow bit). CCR.xcc.v is set, based only on the normal 64-bit arithmetic overflow condition, like a normal 64-bit add.

---

**Compatibility Note** – TADDccTV traps based on the 32-bit overflow condition, just as in SPARC V8. Although the tagged add instructions set the 64-bit condition codes CCR.xcc, there is no form of the instruction that traps the 64-bit overflow condition.

---

### Exceptions


*tag\_overflow*

## A.70.17 Tagged Subtract and Trap on Overflow

The TSubccTV instruction is deprecated. Use the TSubcc instruction followed by BPVS (with instructions to save the pre-TSubcc integer condition codes if necessary).

| Opcode                | op3     | Operation  |
|-----------------------|---------|--|
| TSubccTV <sup>D</sup> | 10 0011 | Tagged Subtract and modify cc's, or Trap on Overflow |

### Format (3)



| Assembly Language Syntax |  |
|--------------------------|--|
| tsubcctv                 | <i>reg<sub>rs1</sub></i> , <i>reg_or_imm</i> , <i>reg<sub>rd</sub></i> |

## Description

This instruction computes “ $r[rs1] - r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] - \text{sign\_ext}(\text{simml3})$ ” if  $i = 1$ .

TSUBccTV modifies the integer condition codes (`icc` and `xcc`) if it does not trap.

A tag overflow occurs if bit 1 or bit 0 of either operand is nonzero or if the subtraction generates 32-bit arithmetic overflow; that is, the operands have different values in bit 31 (the 32-bit sign bit) and the sign of the 32-bit difference in bit 31 differs from bit 31 of  $r[rs1]$ .

If TSUBccTV causes a tag overflow, then a *tag\_overflow* exception is generated and  $r[rd]$  and the integer condition codes remain unchanged. If a TSUBccTV does not cause a tag overflow condition, the difference is written into  $r[rd]$  and the integer condition codes are updated. `CCR.icc.v` is set to zero to indicate no 32-bit overflow.

In either case, the remaining integer condition codes (both the other `CCR.icc` bits and all the `CCR.xcc` bits) are also updated as they would be for a normal subtract instruction. In particular, the setting of the `CCR.xcc.v` bit is not determined by the tag overflow condition (tag overflow is used only to set the 32-bit overflow bit). `CCR.xcc.v` is set, based only on the normal 64-bit arithmetic overflow condition, like a normal 64-bit subtract.

---

**Compatibility Note** – TSUBccTV traps are based on the 32-bit overflow condition, just as in SPARC V8. Although the tagged-subtract instructions set the 64-bit condition codes `CCR.xcc`, there is no form of the instruction that traps on 64-bit overflow.

---

## Exceptions

*tag\_overflow*

## A.70.18 Write Y Register

The WRY instruction is deprecated. It is recommended that all instructions that reference the Y register be avoided.

| Opcode           | op3     | rd   | Operation                               |
|------------------|---------|------|---|
| WRY <sup>D</sup> | 11 0000 | 0    | Write Y register                        |
| —                | 11 0000 | 1–31 | See Section A.69 “Write State Register” |

### Format (3)

|          |    |       |       |          |        |       |
|----------|----|-------|-------|----------|--------|-------|
| 10       | rd | op3   | rs1   | i=0      | —      | rs2   |
| 10       | rd | op3   | rs1   | i=1      | simm13 |       |
| 31 30 29 |    | 25 24 | 19 18 | 14 13 12 |        | 5 4 0 |

#### Assembly Language Syntax

|    |                                |
|----|--------------------------------|
| wr | $reg_{rs1}, reg\_or\_imm, \%y$ |
|----|--------------------------------|

#### Description

This instruction stores the value “ $r[rs1] \mathbf{xor} r[rs2]$ ” if  $i = 0$ , or “ $r[rs1] \mathbf{xor} sign\_ext(simm13)$ ” if  $i = 1$ , to the writable fields of the Y register.

---

**Note** – The operation is an exclusive OR.

---

The WRY instruction is *not* a delayed-write instruction. The instruction immediately following a WRY observes the new value of the Y register.

#### Exceptions

None



# Assembly Language Syntax

---

This appendix supports Appendix A “Instruction Definitions.” Each instruction description in Appendix A includes a table that describes the suggested assembly language format for that instruction. This appendix describes the notation used in those assembly language syntax descriptions and lists some synthetic instructions provided by the SPARC V9 assemblers for the convenience of assembly language programmers.

The appendix contains these sections:

- Notation Used
- Syntax Design
- Synthetic Instructions

---

## B.1 Notation Used

The notations defined here are also used in the assembly language syntax descriptions in Appendix A “Instruction Definitions.”

Items in *typewriter font* are literals to be written exactly as they appear. Items in *italic font* are metasymbols that are to be replaced by numeric or symbolic values in actual SPARC V9 assembly language code. For example, “*imm\_asi*” would be replaced by a number in the range 0 to 255 (the value of the *imm\_asi* bits in the binary instruction) or by a symbol bound to such a number.

Subscripts on metasymbols further identify the placement of the operand in the generated binary instruction. For example, *reg<sub>rs2</sub>* is a *reg* (register name) whose binary value will be placed in the *rs2* field of the resulting instruction.

## B.1.1 Register Names

See *The SPARC Architecture Manual, Version 9* regarding notational conventions for register names.

## B.1.2 Special Symbol Names

Certain special symbols appear in the syntax table in typewriter font. They must be written exactly as they are shown, including the leading percent sign (%). See *The SPARC Architecture Manual, Version 9* regarding special symbol names.

Additional UltraSPARC III Cu symbol names and the registers or operators to which they refer are as follows:

|                             |   |
|-----------------------------|---|
| <code>%ccr</code>           | Condition Codes Register                      |
| <code>%clear_softint</code> | Soft Interrupt Register (clear selected bits) |
| <code>%dcr</code>           | Dispatch Control Register                     |
| <code>%fprs</code>          | Floating-Point Registers State Register       |
| <code>%gsr</code>           | Graphics Status Register                      |
| <code>%pccr</code>          | Performance Control Register                  |
| <code>%pic</code>           | Performance Instrumentation Counters          |
| <code>%set_softint</code>   | Soft Interrupt Register (set selected bits)   |
| <code>%softint</code>       | Soft Interrupt Register                       |
| <code>%sys_tick</code>      | System Timer (TICK) Register                  |
| <code>%sys_tick_cmpr</code> | System Timer (STICK) Compare Register         |
| <code>%tba</code>           | Trap Base Address Register                    |
| <code>%tick_cmpr</code>     | Timer (TICK) Compare Register                 |

The following special symbol names are unary operators that perform the functions described:

|                                       |   |
|---------------------------------------|---|
| <code>%uhi</code>                     | Extracts bits 63:42 (high 22 bits of upper word) of its operand           |
| <code>%ulo</code> or <code>%hm</code> | Extracts bits 41:32 (low-order 10 bits of upper word) of its operand      |
| <code>%hi</code> or <code>%lm</code>  | Extracts bits 31:10 (high-order 22 bits of low-order word) of its operand |
| <code>%lo</code>                      | Extracts bits 9:0 (low-order 10 bits) of its operand                      |

Certain predefined value names appear in the syntax table in typewriter font. They must be written exactly as they are shown, including the leading sharp sign (#). The value names and the values to which they refer are listed in TABLE B-1.

**TABLE B-1** Value Names and Values (1 of 2)

| Name                            | Value            | Description   |
|---------------------------------|------------------|---|
| #n_reads                        | 00 <sub>16</sub> | for PREFETCH instruction, Strong version = 14 <sub>16</sub> |
| #one_read                       | 01 <sub>16</sub> | for PREFETCH instruction, Strong version = 15 <sub>16</sub> |
| #n_writes                       | 02 <sub>16</sub> | for PREFETCH instruction, Strong version = 16 <sub>16</sub> |
| #one_write                      | 03 <sub>16</sub> | for PREFETCH instruction, Strong version = 17 <sub>16</sub> |
| #page                           | 04 <sub>16</sub> | for PREFETCH instruction                                    |
| #Sync                           | 40 <sub>16</sub> | for MEMBAR instruction cmask field                          |
| #MemIssue                       | 20 <sub>16</sub> | for MEMBAR instruction cmask field                          |
| #Lookaside                      | 10 <sub>16</sub> | for MEMBAR instruction cmask field                          |
| #StoreStore)                    | 08 <sub>16</sub> | for MEMBAR instruction cmask field                          |
| #LoadStore                      | 04 <sub>16</sub> | for MEMBAR instruction cmask field                          |
| #StoreLoad                      | 02 <sub>16</sub> | for MEMBAR instruction cmask field                          |
| #LoadLoad                       | 01 <sub>16</sub> | for MEMBAR instruction cmask field                          |
| #ASI_AIUP                       | 10 <sub>16</sub> | ASI_AS_IF_USER_PRIMARY                                      |
| #ASI_AIUS                       | 11 <sub>16</sub> | ASI_AS_IF_USER_SECONDARY                                    |
| #n_reads_strong                 | 14 <sub>16</sub> | Strong version for PREFETCH instruction (00 <sub>16</sub> ) |
| #one_read_strong                | 15 <sub>16</sub> | Strong version for PREFETCH instruction (01 <sub>16</sub> ) |
| #n_writes_strong                | 16 <sub>16</sub> | Strong version for PREFETCH instruction (02 <sub>16</sub> ) |
| #one_write_strong               | 17 <sub>16</sub> | Strong version for PREFETCH instruction (03 <sub>16</sub> ) |
| #ASI_AIUP_L                     | 18 <sub>16</sub> | ASI_AS_IF_USER_PRIMARY_LITTLE                               |
| #ASI_AIUS_L                     | 19 <sub>16</sub> | ASI_AS_IF_USER_SECONDARY_LITTLE                             |
| #ASI_PHYS_USE_EC_L              | 1C <sub>16</sub> | ASI_PHYS_USE_EC_LITTLE                                      |
| #ASI_PHYS_BYPASS_EC_WITH_EBIT_L | 1D <sub>16</sub> | ASI_PHYS_BYPASS_EC_WITH_EBIT_LITTLE                         |
| #ASI_NUCLEUS_QUAD_LDD_L         | 2C <sub>16</sub> | ASI_NUCLEUS_QUAD_LDD_LITTLE                                 |
| #ASI_MONDO_SEND_CTRL            | 48 <sub>16</sub> | ASI_INTR_DISPATCH_STATUS                                    |
| #ASI_MONDO_RECEIVE_CTRL         | 49 <sub>16</sub> | ASI_INTR_RECEIVE  |
| #ASI_AFSR                       | 4C <sub>16</sub> | ASI_ASYNC_FAULT_STATUS                                      |
| #ASI_AFAR                       | 4D <sub>16</sub> | ASI_ASYNC_FAULT_ADDR  |
| #ASI_BLK_AIUP                   | 70 <sub>16</sub> | ASI_BLOCK_AS_IF_USER_PRIMARY                                |
| #ASI_BLK_AIUS                   | 71 <sub>16</sub> | ASI_BLOCK_AS_IF_USER_SECONDARY                              |
| #ASI_BLK_AIUPL                  | 78 <sub>16</sub> | ASI_BLOCK_AS_IF_USER_PRIMARY_LITTLE                         |
| #ASI_BLK_AIUSL                  | 79 <sub>16</sub> | ASI_BLOCK_AS_IF_USER_SECONDARY_LITTLE                       |
| #ASI_P                          | 80 <sub>16</sub> | ASI_PRIMARY   |
| #ASI_S                          | 81 <sub>16</sub> | ASI_SECONDARY   |
| #ASI_PNF                        | 82 <sub>16</sub> | ASI_PRIMARY_NOFAULT   |

**TABLE B-1** Value Names and Values (2 of 2)

| Name              | Value            | Description                  |
|-------------------|------------------|------------------------------|
| #ASI_SNF          | 83 <sub>16</sub> | ASI_SECONDARY_NOFAULT        |
| #ASI_P_L          | 88 <sub>16</sub> | ASI_PRIMARY_LITTLE           |
| #ASI_S_L          | 89 <sub>16</sub> | ASI_SECONDARY_LITTLE         |
| #ASI_PNF_L        | 8A <sub>16</sub> | ASI_PRIMARY_NOFAULT_LITTLE   |
| #ASI_SNF_L        | 8B <sub>16</sub> | ASI_SECONDARY_NOFAULT_LITTLE |
| #ASI_PST8_P       | C0 <sub>16</sub> | ASI_PST8_PRIMARY             |
| #ASI_PST8_S       | C1 <sub>16</sub> | ASI_PST8_SECONDARY           |
| #ASI_PST16_P      | C2 <sub>16</sub> | ASI_PST16_PRIMARY            |
| #ASI_PST16_S      | C3 <sub>16</sub> | ASI_PST16_SECONDARY          |
| #ASI_PST32_P      | C4 <sub>16</sub> | ASI_PST32_PRIMARY            |
| #ASI_PST32_S      | C5 <sub>16</sub> | ASI_PST32_SECONDARY          |
| #ASI_PST8_PL      | C8 <sub>16</sub> | ASI_PST8_PRIMARY_LITTLE      |
| #ASI_PST8_SL      | C9 <sub>16</sub> | ASI_PST8_SECONDARY_LITTLE    |
| #ASI_PST16_PL     | CA <sub>6</sub>  | ASI_PST16_PRIMARY_LITTLE     |
| #ASI_PST16_SL     | CB <sub>16</sub> | ASI_PST16_SECONDARY_LITTLE   |
| #ASI_PST32_PL     | CC <sub>16</sub> | ASI_PST32_PRIMARY_LITTLE     |
| #ASI_PST32_SL     | CD <sub>16</sub> | ASI_PST32_SECONDARY_LITTLE   |
| #ASI_FL8_P        | D0 <sub>16</sub> | ASI_FL8_PRIMARY              |
| #ASI_FL8_S        | D1 <sub>16</sub> | ASI_FL8_SECONDARY            |
| #ASI_FL16_P       | D2 <sub>16</sub> | ASI_FL16_PRIMARY             |
| #ASI_FL16_S       | D3 <sub>16</sub> | ASI_FL16_SECONDARY           |
| #ASI_FL8_PL       | D8 <sub>16</sub> | ASI_FL8_PRIMARY_LITTLE       |
| #ASI_FL8_SL       | D9 <sub>16</sub> | ASI_FL8_SECONDARY_LITTLE     |
| #ASI_FL16_PL      | DA <sub>16</sub> | ASI_FL16_PRIMARY_LITTLE      |
| #ASI_FL16_SL      | DB <sub>16</sub> | ASI_FL16_SECONDARY_LITTLE    |
| #ASI_BLK_COMMIT_P | E0 <sub>16</sub> | ASI_BLOCK_COMMIT_PRIMARY     |
| #ASI_BLK_COMMIT_S | E1 <sub>16</sub> | ASI_BLOCK_COMMIT_SECONDARY   |
| #ASI_BLK_P        | F0 <sub>16</sub> | ASI_BLOCK_PRIMARY            |
| #ASI_BLK_S        | F1 <sub>16</sub> | ASI_BLOCK_SECONDARY          |
| #ASI_BLK_PL       | F8 <sub>16</sub> | ASI_BLOCK_PRIMARY_LITTLE     |
| #ASI_BLK_SL       | F9 <sub>16</sub> | ASI_BLOCK_SECONDARY_LITTLE   |

The full names of the ASIs, listed in the Description column of TABLE B-1 can also be defined.



## B.1.3 Values

See *The SPARC Architecture Manual, Version 9* regarding notational conventions for values.

## B.1.4 Labels

See *The SPARC Architecture Manual, Version 9* regarding notational conventions for labels.

## B.1.5 Other Operand Syntax

See *The SPARC Architecture Manual, Version 9* regarding notational conventions for general operand syntax. Additional operand syntax is listed below.

### *membar\_mask*

Operand syntax is as follows:

*const7* — A constant that can be represented in 7 bits. Typically, this is an expression involving the logical OR of some combination of #Lookaside, #MemIssue, #Sync, #StoreStore, #LoadStore, #StoreLoad, and #LoadLoad.

### *prefetch\_fcn (Prefetch Function)*

Operand syntax is summarized in the TABLE B-2.

**TABLE B-2** Prefetch Function Values

| Prefetch Function Value (Decimal) | Operand                                |
|-----------------------------------|--|
| 0                                 | #n_reads (Strong version = 20)         |
| 1                                 | #one_read (Strong version = 21)        |
| 2                                 | #n_writes (Strong version = 22)        |
| 3                                 | #one_write (Strong version = 23)       |
| 4                                 | #page                                  |
| 5 – 15                            | Reserved and Traps                     |
| 16                                | P-cache Invalidate                     |
| 17-19                             | NOP                                    |
| 20                                | #n_reads_strong (Strong version of 0)  |
| 21                                | #one_read_strong (Strong version of 1) |

**TABLE B-2** Prefetch Function Values (*Continued*)

| Prefetch Function Value<br>(Decimal) | Operand                                 |
|--------------------------------------|---|
| 22                                   | #n_writes_strong (Strong version of 2)  |
| 23                                   | #one_write_strong (Strong version of 3) |
| 24 – 31                              | NOP                                     |

## B.1.6 Comments

See *The SPARC Architecture Manual, Version 9* regarding notational conventions for comments.

---

## B.2 Syntax Design

See *The SPARC Architecture Manual, Version 9* regarding syntax design.

---

## B.3 Synthetic Instructions

See *The SPARC Architecture Manual, Version 9* regarding synthetic instructions.

---

---

# Index

---

---



## A

- a* field of instructions 177, 470, 473, 474, 477, 604, 607
  - A pipeline stage 42, 45
  - A0 pipeline stage 43
  - A1 pipeline stage 43
  - accesses
    - cacheable 211
    - I/O 220
    - noncacheable 211
    - nonfaulting ASIs 217
    - real memory space 210
    - restricted ASI 210
    - with side effects 210, 211, 220
  - accrued exception (*aexc*) field of FSR register 126, 129, 325
  - ADD instruction 454
  - ADDC instruction 454
  - ADDcc instruction 454, 580
  - ADDCcc instruction 454
  - address
    - 64-bit virtual data watchpoint 136
    - aliasing 225, 247
    - illegal address aliasing 230
    - physical address data watchpoint 137
    - space identifier (ASI) 210
    - space size, differences from UltraSPARC-I 9
    - virtual address
      - data watchpoint 136
      - watchpoint priority 136
    - virtual passed to physical 133
    - virtual-to-physical translation 210
  - address mask (AM) field of PSTATE register 116
  - address space identifier (ASI)
    - accessing MMU registers 277
    - affected by watchpoint traps 136
    - appended to memory address 143, 180
    - bit 7 setting for *privileged\_action* exception 587
    - bypass 189
    - definition xliii
    - explicit values 142
    - imm\_asi* instruction field 178
    - implicit values 142
    - load floating-point instructions 503
    - load from TLB Data Access register 281, 300
    - load from TLB Data In register 281, 300
    - load from TLB Tag Read register 280, 300
    - load integer doubleword instructions 613
    - load integer instructions 507
    - operations 276, 298
      - with prefetch instructions 561
      - restricted 275
      - restriction indicator 98
      - unrestricted 189
  - address space identifier (ASI) register
    - for load/store alternate instructions 98
    - and *imm\_asi* instruction field 142
    - and LDDA instruction 615
    - and LDSTUBA instruction 513
    - load floating-point from alternate space instructions 505
    - load integer from alternate space instructions 509
    - with prefetch instructions 562
  - restoring saved state 480
  - saving state 315
    - and STDA instruction 624
  - store floating-point into alternate space instructions 587
  - store integer to alternate space instructions 590
    - and SWAPA instruction 627
    - and TSTATE Register 110
    - and write state register instructions 601
- addressing conventions 141, 180
- ADDX instruction (SPARC V8) 455
- ADDXcc instruction (SPARC V8) 455
- aexc* field of FSR 409
- AFAR
  - extension 10
  - state after reset 442
- AFSR
  - changes to 10
  - state after reset 442
- alias
  - address 225
  - boundary 230
  - floating-point registers 87
- ALIGNADDRESS instruction 455
- ALIGNADDRESS\_LITTLE instruction 455
- aligning branch targets 389
- alignment
  - data (load/store) 141, 141
  - doubleword 141, 141
  - extended-word 141
  - halfword 141, 141
  - instructions 141, 141
  - integer registers 613, 615
  - memory 347
  - quadword 141, 141
  - word 141, 141

- alternate address space 561
- alternate global registers 82
- alternate globals enable (AG) field of PSTATE register 82, 114
- alternate space instructions 98
- ancillary state registers (ASRs) 96–??
  - access 100
  - number 96
  - possible registers included 570, 601
  - writing to 601
- AND instruction 517
- ANDcc instruction 517
- ANDN instruction 517
- ANDNcc instruction 517
- annul bit
  - in branch instructions 470
  - in conditional branches 605
  - in control transfer instruction 99
- annulled branches 470
- annulled slot 391
- application program xliii, 132
- Architectural Register File (ARF) 52
- architecture, meaning for SPARC V9 xl
- ARF (Architectural Register File) 52
- arithmetic overflow 97
- ARRAY16 instruction 457
- ARRAY32 instruction 457
- ARRAY8 instruction 457
- ASI
  - \_BLK\_COMMIT\_PRIMARY 231
  - \_BLK\_COMMIT\_SECONDARY 231
  - \_NUCLEUS\_QUAD\_LDD\_S 510
  - atomic access 217
  - load from TLB Data Access register 293
  - nonfaulting 217
  - registers, differences from UltraSPARC-I 9
  - registers, removed in UltraSPARC-III 10
  - registers, state after reset 439
  - restricted 210
  - UltraSPARC III internal 221
- ASI\_AFAR 196
- ASI\_AFSR 196
- ASI\_AIUP 193
- ASI\_AIUPL 194
- ASI\_AIUS 194
- ASI\_AIUSL 194
- ASI\_AS\_IF\_USER\_PRIMARY 193, 217
- ASI\_AS\_IF\_USER\_PRIMARY\_LITTLE 194, 217
- ASI\_AS\_IF\_USER\_SECONDARY 194, 217
- ASI\_AS\_IF\_USER\_SECONDARY\_LITTLE 194, 217
- ASI\_ASYNC\_FAULT\_ADDR 196
- ASI\_ASYNC\_FAULT\_STATUS 196
- ASI\_ATOMIC\_QUAD\_LDD\_PHYS 195
- ASI\_ATOMIC\_QUAD\_LDD\_PHYS\_LITTLE 199, 203
- ASI\_BARRIER\_SYNCH 203
- ASI\_BARRIER\_SYNCH\_P 199
- ASI\_BLK\_AIUP 199
- ASI\_BLK\_AIUPL 200
- ASI\_BLK\_AIUS 199
- ASI\_BLK\_AIUSL 200
- ASI\_BLK\_COMMIT\_P 203
- ASI\_BLK\_COMMIT\_S 203
- ASI\_BLK\_P 204
- ASI\_BLK\_PL 204
- ASI\_BLK\_S 204
- ASI\_BLK\_SL 204
- ASI\_BLOCK\_AS\_IF\_USER\_PRIMARY 199
- ASI\_BLOCK\_AS\_IF\_USER\_PRIMARY\_LITTLE 200
- ASI\_BLOCK\_AS\_IF\_USER\_SECONDARY 199
- ASI\_BLOCK\_AS\_IF\_USER\_SECONDARY\_LITTLE 200
- ASI\_BLOCK\_COMMIT\_PRIMARY 203
- ASI\_BLOCK\_COMMIT\_SECONDARY 203
- ASI\_BLOCK\_PRIMARY 204
- ASI\_BLOCK\_PRIMARY\_LITTLE 204
- ASI\_BLOCK\_SECONDARY 204
- ASI\_BLOCK\_SECONDARY\_LITTLE 204
- ASI\_DCACHE\_DATA 195
- ASI\_DCACHE\_INVALIDATE 195
- ASI\_DCACHE\_SNOOP\_TAG 195
- ASI\_DCACHE\_TAG 195
- ASI\_DCACHE\_UTAG 195
- ASI\_DCU\_CONTROL\_REGISTER 132, 195
- ASI\_DCUCR 195
- ASI\_DEVICE\_ID+SERIAL\_ID 197
- ASI\_DMMU\_DEMAP 199
- ASI\_DMMU\_PA\_WATCHPOINT\_REG 198
- ASI\_DMMU\_SFAR 198
- ASI\_DMMU\_SFSR 198
- ASI\_DMMU\_TAG\_ACCESS 198
- ASI\_DMMU\_TAG\_TARGET\_REG 198
- ASI\_DMMU\_TSB\_64KB\_PTR\_REG 198
- ASI\_DMMU\_TSB\_8KB\_PTR\_REG 198
- ASI\_DMMU\_TSB\_BASE 198
- ASI\_DMMU\_TSB\_DIRECT\_PTR\_REG 198
- ASI\_DMMU\_TSB\_NEXT\_REG 198
- ASI\_DMMU\_TSB\_PEXT\_REG 198
- ASI\_DMMU\_TSB\_SEXT\_REG 198

ASI\_DMMU\_VA\_WATCHPOINT\_REG 198  
 ASI\_DTLB\_DATA\_ACCESS\_REG 199  
 ASI\_DTLB\_DATA\_IN\_REG 198  
 ASI\_DTLB\_TAG\_READ\_REG 199  
 ASI\_FL16\_P 203  
 ASI\_FL16\_PL 203  
 ASI\_FL16\_PRIMARY 203  
 ASI\_FL16\_PRIMARY\_LITTLE 203  
 ASI\_FL16\_S 203  
 ASI\_FL16\_SECONDARY 203  
 ASI\_FL16\_SECONDARY\_LITTLE 203  
 ASI\_FL16\_SL 203  
 ASI\_FL8\_P 203  
 ASI\_FL8\_PL 203  
 ASI\_FL8\_PRIMARY 203  
 ASI\_FL8\_PRIMARY\_LITTLE 203  
 ASI\_FL8\_S 203  
 ASI\_FL8\_SECONDARY 203  
 ASI\_FL8\_SECONDARY\_LITTLE 203  
 ASI\_FL8\_SL 203  
 ASI\_IIU\_INST\_TRAP 199, 203  
 ASI\_IMMU\_DEMAP 197  
 ASI\_IMMU\_SFSR 196  
 ASI\_IMMU\_TAG\_TARGET 196  
 ASI\_IMMU\_TSB\_64KB\_PTR\_REG 197  
 ASI\_INTR\_DATA0\_R 201  
 ASI\_INTR\_DATA0\_W 200  
 ASI\_INTR\_DATA1\_R 201  
 ASI\_INTR\_DATA1\_W 200  
 ASI\_INTR\_DATA2\_R 201  
 ASI\_INTR\_DATA2\_W 200  
 ASI\_INTR\_DATA3\_R 201  
 ASI\_INTR\_DATA3\_W 200  
 ASI\_INTR\_DATA4\_R 201  
 ASI\_INTR\_DATA4\_W 200  
 ASI\_INTR\_DATA5\_R 201  
 ASI\_INTR\_DATA5\_W 200  
 ASI\_INTR\_DATA6\_R 201  
 ASI\_INTR\_DATA6\_W 200  
 ASI\_INTR\_DATA7\_R 201  
 ASI\_INTR\_DATA7\_W 200  
 ASI\_INTR\_DISPATCH\_STATUS 352, 355, 356  
 ASI\_INTR\_DISPATCH\_STATUS.BUSY bit 352  
 ASI\_INTR\_DISPATCH\_STATUS.NACK bit 352  
 ASI\_INTR\_DISPATCH\_W 200, 355  
 ASI\_INTR\_RECEIVE 196, 353, 357  
 ASI\_INTR\_W 352, 354  
 ASI\_ITLB\_DATA\_ACCESS\_REG 197  
 ASI\_ITLB\_DATA\_IN\_REG 197  
 ASI\_ITLB\_TAG\_READ\_REG 197  
 ASI\_MONDO\_RECEIVE\_CTRL 196  
 ASI\_MONDO\_SEND\_CTRL 196  
 ASI\_N 193  
 ASI\_NL 193  
 ASI\_NUCLEUS 193, 217, 261, 262  
 ASI\_NUCLEUS\_LITTLE 193, 217, 262  
 ASI\_NUCLEUS\_QUAD\_LDD 206  
 ASI\_NUCLEUS\_QUAD\_LDD\_L 194  
 ASI\_NUCLEUS\_QUAD\_LDD\_LITTLE 194, 206  
 ASI\_P 201  
 ASI\_PCACHE\_DATA 194  
 ASI\_PCACHE\_SNOOP\_TAG 195  
 ASI\_PCACHE\_STATUS\_DATA 194  
 ASI\_PCACHE\_TAG 195  
 ASI\_PHYS\_BYPASS\_EC\_WITH\_EBIT 194, 205, 263  
 ASI\_PHYS\_BYPASS\_EC\_WITH\_EBIT\_L 194  
 ASI\_PHYS\_BYPASS\_EC\_WITH\_EBIT\_LITTLE 194,  
 205  
 ASI\_PHYS\_USE\_EC 3, 194, 205, 217, 228  
 ASI\_PHYS\_USE\_EC\_L 194  
 ASI\_PHYS\_USE\_EC\_LITTLE 3, 194, 205, 217  
 ASI\_PL 201  
 ASI\_PNF 201  
 ASI\_PNFL 201  
 ASI\_PRIMARY 142, 201, 217, 261  
 ASI\_PRIMARY\_\* 262  
 ASI\_PRIMARY\_CONTEXT\_REG 198  
 ASI\_PRIMARY\_LITTLE 115, 201, 217, 261, 262  
 ASI\_PRIMARY\_NO\_FAULT 201, 218, 288, 307  
 ASI\_PRIMARY\_NO\_FAULT\_LITTLE 201, 218, 288,  
 307  
 ASI\_PST16\_P 202, 540  
 ASI\_PST16\_PL 202, 540  
 ASI\_PST16\_PRIMARY 202  
 ASI\_PST16\_PRIMARY\_LITTLE 202  
 ASI\_PST16\_S 202, 540  
 ASI\_PST16\_SECONDARY 202  
 ASI\_PST16\_SECONDARY\_LITTLE 202  
 ASI\_PST16\_SL 540  
 ASI\_PST32\_P 202, 540  
 ASI\_PST32\_PL 202, 540  
 ASI\_PST32\_PRIMARY 202  
 ASI\_PST32\_PRIMARY\_LITTLE 202  
 ASI\_PST32\_S 202, 540  
 ASI\_PST32\_SECONDARY 202  
 ASI\_PST32\_SECONDARY\_LITTLE 202  
 ASI\_PST32\_SL 202, 540  
 ASI\_PST8\_P 202, 540

- ASI\_PST8\_PL 202, 540
- ASI\_PST8\_PRIMARY 202
- ASI\_PST8\_PRIMARY\_LITTLE 202
- ASI\_PST8\_S 202, 540
- ASI\_PST8\_SECONDARY 202
- ASI\_PST8\_SECONDARY\_LITTLE 202
- ASI\_PST8\_SL 202, 540
- ASI\_S 201
- ASI\_SDB\_INTR 354, 356
- ASI\_SDB\_INTR\_R 353
- ASI\_SECONDARY 201, 217
- ASI\_SECONDARY\_CONTEXT\_REG 198
- ASI\_SECONDARY\_LITTLE 201, 217
- ASI\_SECONDARY\_NO\_FAULT 201, 218, 288, 307
- ASI\_SECONDARY\_NO\_FAULT\_LITTLE 202, 218, 288, 307
- ASI\_SL 201
- ASI\_SNF 201
- ASI\_SNFL 202
- ASI\_SRAM\_TEST\_INIT 195
- ASI\_WCACHE\_DATA 195
- ASI\_WCACHE\_SNOOP\_TAG 195
- ASI\_WCACHE\_TAG 195
- ASI\_WCACHE\_VALID\_BITS 195
- ASRs
  - differences from UltraSPARC-I 10
  - grouping rules 52
- assembler, synthetic instructions 638–??
- async\_data\_error* exception 504, 510, 512, 615
- atomic
  - load quadword 510
  - memory operations 510, 511
  - store doubleword instruction 623, 624
  - store instructions 588, 590
- atomic instructions
  - compare and swap 217
  - LDSTUB 217
  - mutual exclusion support 216
  - and store queue 222
  - SWAP 216
  - use with ASIs 217
- atomic load-store instructions 478
  - compare and swap 477
  - load-store unsigned byte 512, 626, 627
  - load-store unsigned byte to alternate space 513
  - swap *r* register with alternate space memory 626
  - swap *r* register with memory 478, 625

## B

- B pipeline stage 43
- BA instruction 606, 607
- BCC instruction 606
- BCS instruction 606
- BE instruction 606
- BG instruction 606
- BGE instruction 606
- BGU instruction 606
- Bicc instructions 98, 99, 605
- big-endian
  - swapping in partial store instructions 541
- big-endian byte order 115, 140, 141, 180
- bit vector concatenation xli
- bit vectors, concatenation 255
- BLE instruction 606
- BLEU instruction 606
- block
  - load and store instructions
    - compliance across UltraSPARC platforms 521
    - data size (granularity) 219
    - external cache allocation 228
  - load instruction 219
    - E-cache allocation 228
    - grouping 53
    - ordering 464
    - and prefetch cache 401
    - and store queue 223
  - load instructions 87, 231, 460
  - operations and memory model 464
  - overlapping stores 464
  - store instruction
    - data size (granularity) 219
    - E-cache allocation 228
    - grouping 53
    - ordering 464
    - and PDIST 51
    - use in loops 464
  - store instructions 87, 460
  - store with commit 7
  - use in loops 464
- block load instructions 191
- block store instructions 191
- BMASK instruction 468
  - and BSHUFFLE instruction 469
  - and MS pipeline 469
  - grouping rules 51
- BN instruction 606, 607
- BNE instruction 606



- BNEG instruction 606
- bootbus address
  - mapping 3
  - unallowed memory operations 3
- booting code 228
- BPA instruction 474
- BPCC instruction 474
- BPcc instructions 98, 99, 177, 178, 474
- BPCS instruction 474
- BPE instruction 474
- BPG instruction 474
- BPGE instruction 474
- BPGU instruction 474
- BPL instruction 474
- BPLE instruction 474
- BPLEU instruction 474
- BPN instruction 474
- BPNE instruction 474
- BPNEG instruction 474
- BPOS instruction 606
- BPPOS instruction 474
- BPr instructions 99, 177, 178, 469
- BPVC instruction 474
- BPVS instruction 474
- BR pipeline 43
- branch
  - annulled 470
  - delayed 180
  - elimination 145, 146
  - fcc*-conditional 473, 605
  - icc*-conditional 607
  - prediction bit 470
  - target alignment 389
  - unconditional 473, 475, 604, 607
- branch if contents of integer register match condition
  - instructions 469
- branch instructions, conditional 45
- branch on floating-point condition codes instructions 603
- branch on floating-point condition codes with prediction
  - instructions 471
- branch on integer condition codes instructions, *See Bicc instructions*
- branch prediction
  - in B pipeline stage 43
  - mispredict signal 45
  - statistics for taken/untaken 370
- Branch Predictor (BP) 42
- break-after, definition 47
- break-before, definition 47
- BRGEZ instruction 469
- BRGZ instruction 469
- BRLEZ instruction 469
- BRLZ instruction 469
- BRNZ instruction 469
- BRZ instruction 469
- BSHUFFLE instruction 468
  - and BMASK instruction 469
  - fully pipelined 469
  - grouping rules 51
- bubble, vs. helper 52
- bubbles 370
- BUSY/NACK pairs 356
- BVC instruction 606
- BVS instruction 606
- byte
  - addressing 182
  - data format 65
  - order 140, 141, 180
  - order, big-endian 115, 140
  - order, implicit 115
  - order, little-endian 115, 140
- byte mask
  - grouping 469
  - instruction 4
- byte ordering 541

## C

- C pipeline stage 45, 46
- cache
  - coherency protocol 211
  - differences from UltraSPARC-I 6
  - external 227
  - flushing 229
  - level 1 225
  - level 2 226, 227
  - organization 225
  - write 227
- cacheable accesses
  - indication 211
  - properties 211
- caching, TSB 252
- CALL instruction
  - description 476
  - destination register 99
  - displacement 159
  - does not change CWP 86

- and JMWPL instruction 502
  - writing address into r[15] 82
- CAM Diagnostic Register format 293
- CANRESTORE register 119, 440
- CANSAVE register 118, 440
- carry (C) bit of condition fields of CCR 97
- CAS(X)A instruction 217
- CASA instruction 146, 477, 512, 513, 626, 627
- CASXA instruction 146, 477, 512, 513, 626, 627
- catastrophic\_error* exception **316**, 346
- cc0* field of instructions 177, 473, 474, 486, 533
- cc1* field of instructions 177, 473, 474, 486, 533
- cc2* field of instructions 177, 533
- CCR register 439
- CCR, *See* condition codes (CCR) register
- cexc* field of FSR 409
- clean register window 120, 348, 572
- clean windows (CLEANWIN) register 119, 119, 566, 598
  - clean\_window* exception 119, 323, 326, 348, 573, 574
- CLEANWIN register 440
- CLEAR\_SOFTINT pseudo-register 358
- clock-tick register (TICK) 106, 107, 347, 566, 598
- code
  - kernel 358
  - nucleus 358
- coherence
  - domain 211
  - unit of 211
- compare and swap instructions 477
- comparison instruction 148, 592
- complex calculations, fixed data format 76
- compliance
  - SPARC V9 264
- concatenation of bit vectors xli
- cond* field of instructions 177, 473, 474, 526, 533, 604, 607
- condition code register 439
- condition codes 479
  - adding 593
  - extended integer (Xcc) 98
  - floating-point 605
  - icc field 97
  - integer 96
  - results of integer operation (icc) 98
  - subtracting 592, 594
  - trapping on 596
  - xcc field 97
- condition codes (CCR) register 96, 110, 315, 454, 480, 601, 618
- conditional branch instructions 45
- conditional branches 473, 605, 607
- conditional move instructions
  - grouping rules 53
- const22* field of instructions 500
- constants, generating 577
- Context field of the D-TLB Tag Access Register 280, 299
- context register
  - determination of 261
  - Nucleus 278
  - Primary 278
  - Secondary 278
- control and status registers 96
- control-transfer instructions (CTIs) 158, 480
- conventions
  - font xli
  - notational xli
- conversion
  - between floating-point formats instructions 489
  - floating-point to integer instructions 488
  - integer to floating-point instructions 491
  - planar to packed 559
- CTI couple 392
- CTI queue 43
- current exception (*cexc*) field of FSR register 126, 128, 129, 129, 131, 151, 348
- current window pointer (CWP) register
  - and CALL/JMWPL instructions 86
  - and clean windows 120
  - definition xliv
  - and FLUSHW instruction 500
  - function 118
  - incremented/decremented 84, 573
  - and overlapping windows 84
  - range of values 119
  - reading CWP with RDPR instruction 566
  - and RESTORE instruction 158, 573
  - restored during DONE or RETRY 480
  - and SAVE instruction 158, 573
  - saved during a trap 315
  - and TSTATE Register 110
  - writing CWP with WRPR instruction 598
- current\_little\_endian (CLE) field of PSTATE register 115, 115
- CWP register 439
- cycles accumulated, count 369

## D

D pipeline stage 46, 370

*d16hi* field of instructions 177, 470

*d16lo* field of instructions 177, 470

data

formats

byte 65

doubleword 65

extended word 65

halfword 65

quadword 65

tagged word 65

word 65

types

floating-point 65

signed integer 65

unsigned integer 65

width 65

watchpoint

behavior 136

exception 541

physical address 137

register format 137

virtual address 137

data alignment 394

Data Cache 225

flush 230

data cache

and RED\_state 435

and block load/store 463

flushing 7

Data Cache Unit Control Register, *See* DCUCR

*data\_access\_error* exception 346, 467, 479, 504, 508, 511, 512, 514, 541, 582, 585, 587, 589, 590, 614, 622, 623, 625

*data\_access\_exception* exception 114, 251, 258, 259, 260, 264, 277, 288, 307, 349, 355, 356, 479, 504, 506, 512, 514, 585, 587, 589, 590, 623, 625, 626, 628

*data\_access\_exception* exception 191, 206, 211, 217, 218, 220

*data\_access\_MMU\_miss* exception 259, 260

*data\_access\_protection* exception 259, 260, 287, 306, 467, 508, 510, 511, 541, 582, 614, 616

DB\_PA field of PA Data Watchpoint register 137

DC\_wr 373

DC\_wr\_miss 374

D-cache

hit rate 394

line 394

logical organization *illustrated* 393

misses 394

organization 393

sub-block 394

timing 394

DCR

branch and return control 101

changes to 10

fields

BPE (branch prediction enable) 101

MS (multiscalar dispatch enable) 102

RPE (return address prediction enable) 101

SI (single issue disable) 101

IFPOE field 102

instruction dispatch control 101

layout 100

state after reset 441

DCTI 43

DCU

control register 435

DCUCR

CP (cacheability) field 228, 263

CV (cacheability) field 264

DC (data cache enable) field 134

DM (DMMU enable) field 133

IC (I-cache enable) field 221

IC (instruction cache enable) field 134

IMI (IMMU enable) field 133

PM (PA data watchpoint mask) field 134

PR/PW (PA watchpoint enable) fields 135

RED\_state 263

SPE (software prefetch enable) field 399

state after reset 441

VM (VA data watchpoint mask) field 135

VR/VW (VA data watchpoint enable) fields 135

watchpoint byte masks/enable bits 136

deferred trap

occurrence 321

queue, floating-point (FQ) 566

software actions 322

vs. disrupting trap 322

delay instruction 99, 158, 470, 473, 476, 480, 571, 605

delay slot

and instruction fetch 389

delayed branch 180

delayed control transfer 99, 470

delayed control-transfer instruction (DCTI) 43

demap operation and output 291, 309

deprecated instructions

BA 606  
 BCC 606  
 BCS 606  
 BE 606  
 BG 606  
 BGE 606  
 BGU 606  
 Bicc 605  
 BLE 606  
 BLEU 606  
 BN 606  
 BNE 606  
 BNEG 606  
 BPOS 606  
 BVC 606  
 BVS 606  
 FBA 603  
 FBE 603  
 FBG 603  
 FBGE 603  
 FBL 603  
 FBLE 603  
 FBLG 603  
 FBN 603  
 FBNE 603  
 FBO 603  
 FBU 603  
 FBUE 603  
 FBUGE 603  
 FBUL 603  
 FBULE 603  
 LDD 612  
 LDDA 614  
 LDFSR 611  
 MULSc 96, 617  
 RDY 96, 568, 619  
 SDIV 96, 608  
 SDIVcc 96, 608  
 SMUL 96, 616  
 SMULcc 96, 616  
 STD 622  
 STDA 623  
 STFSR 621  
 SWAP 625  
 SWAPA 626  
 TSUBccTV 628, 629  
 UDIV 96, 608  
 UDIVcc 96, 608  
 UMUL 96, 616  
 UMULcc 96, 616  
 WRY 96, 600, 630  
 Direct Pointer Register 287, 306  
*disp19* field of instructions 178, 473, 474  
*disp22* field of instructions 178, 604, 607  
*disp30* field of instructions 178, 476  
 Dispatch\_rs\_mispred 371  
 Dispatch0\_2nd\_br 371  
 Dispatch0\_br\_target 371  
 displacement flush 231  
 disrupting traps 322, 323  
 divide instructions 537, 608  
 divide-by-zero mask (*DZM*) bit of TEM field of FSR  
   register 128  
*division\_by\_zero* exception 148, 346, 538  
 division-by-zero accrued (*dza*) bit of *aexc* field of FSR  
   register 131  
 division-by-zero current (*dzc*) bit of *cexc* field of FSR  
   register 131  
 D-MMU  
   and RED\_state 435  
   context # upon a trap 280, 299  
   context register usage 262  
   determining ASI value and endianness 261  
   Direct Pointer register 286, 305  
   disabled 264  
   disabled, effect on D-cache 264  
   Enable bit 263  
   enable bits 263  
   memory operation summary 274  
   Nucleus Context Register 278  
   pointer logic pseudocode 257  
   Registers:Primary, Secondary, Nucleus 278  
   Secondary Context Register 278  
   virtual address translation 246  
 DONE instruction 98, 114, 315, 317, 347, 479  
   after internal store to ASI 221  
   and BST 463  
   exiting RED\_state 30  
   grouping rules 53  
   restoring AG, IG, MG bits 113  
   target address 159  
   when TSTATE uninitialized 31, 437  
 doublet xliv  
 doubleword  
   addressing 183  
   alignment 141  
   data format 65  
   definition xliv

- in memory 84
- D-SFAR
  - defined 289
  - description 277
- D-SFAR register
  - exception address (64-bit) 116, 117
  - state after reset 442
  - update after trap 261
- D-SFSR
  - and ASI operations 277
  - bit description 287, 306
  - FTYPE field upon data\_access\_exception 346
- D-SFSR register
  - CT (context) field 280, 299
  - state after reset 442
  - update after trap 261
- D-TLB
  - access 45
  - state after reset 443
- DVMA 436

## E

- E pipeline stage 44
- EC\_ic\_miss 375
- EC\_misses 375
- ECC
  - error correction 349
  - protection, on external cache 11
- ECC\_error exception 349
- Edge instructions 4
- EDGE16 instruction 480
- EDGE16L instruction 480
- EDGE16LN instruction 480
- EDGE16N instruction 480
- EDGE32 instruction 480
- EDGE32L instruction 480
- EDGE32LN instruction 480
- EDGE32N instruction 480
- EDGE8 instruction 480
- EDGE8L instruction 480
- EDGE8LN instruction 480
- EDGE8N instruction 480
- Edgenc instruction 4
- emulating multiple unsigned condition codes 146
- enable floating-point (FEF) field of FPRS register 99,
  - 116, 150, 324, 346, 473, 503, 505, 585, 587, 605
- enable floating-point (PEF) field of PSTATE register 99,

- 116, 150, 324, 346, 473, 503, 505, 585, 587, 605
- enable RED\_state field (RED) of PSTATE register 317
- Energy Star
  - compliance with 5
- Error Enable Register
  - NCEEN field 221
  - state after reset 442
- error\_state 316, 317, 319, 320, 334, 335, 336, 337, 338, 344
- error\_state, and watchdog reset 32, 437
- errors
  - handling differences from UltraSPARC-I 11
- exceptions
  - async\_data\_error* 504, 510, 512, 615
  - catastrophic\_error* 316, 346
  - causing traps 315
  - clean\_window* 119, 323, 326, 348, 573, 574
  - data\_access\_error* 346, 467, 479, 504, 508, 511, 512, 514, 541, 582, 585, 587, 589, 590, 614, 622, 623, 625
  - data\_access\_exception* 479, 504, 506, 512, 514, 585, 587, 589, 590, 623, 625, 626, 628
  - data\_access\_protection* 467, 508, 510, 511, 541, 582, 614, 616
  - definition 315
  - division\_by\_zero* 148, 346, 538
  - fill\_n\_normal* 324, 346, 572, 574
  - fill\_n\_other* 324, 572, 574
  - fp\_disabled* 99, 150, 323, 346, 473, 485, 489, 491, 492, 494, 496, 503, 504, 505, 506, 529, 531, 535, 585, 587, 605, 612, 622
  - fp\_exception\_ieee\_754* 123, 129, 130, 131, 325, 348, 485, 489, 491, 492, 496
  - fp\_exception\_other* 89, 179, 348, 485, 487, 489, 491, 492, 494, 496, 497, 531
  - illegal\_instruction* 84, 109, 179, 346, 471, 476, 480, 501, 504, 535, 537, 560, 567, 568, 569, 575, 585, 587, 597, 599, 613, 614, 615, 622, 623, 624, 625
  - instruction\_access\_error* 323
  - instruction\_access\_exception* 323, 347
  - LDDF\_mem\_address\_not\_aligned* 141, 323, 349, 503, 505, 506
  - mem\_address\_not\_aligned* 141, 347, 479, 502, 503, 504, 506, 507, 508, 509, 510, 571, 572, 585, 587, 589, 590, 614, 615, 622, 623, 625, 626, 628
  - privileged\_action* 98, 142, 323, 347, 479, 505, 506, 509, 510, 514, 570, 587, 590, 615, 624, 625, 628
  - privileged\_instruction* (SPARC V8) 348
  - privileged\_opcode* 323, 347, 480, 568, 575, 599

- spill\_n\_normal* 323, 348, 500, 574
- spill\_n\_other* 323, 348, 500, 574
- STDF\_mem\_address\_not\_aligned* 141, 323, 349, 585, 587
- tag\_overflow* 148, 348, 593, 594, 628, 629, 630
- trap\_instruction* 323, 348, 596, 597
- window\_fill* 120, 571
- window\_spill* 120
- execute\_state 316, 334, 336, 337, 338
- extended word addressing 183
- extended word data format 65
- External Cache 227
  - flush 230
    - Tag Access Address Format 268
- external cache
  - and block load/store 463
  - bypassing by instruction fetches 228
  - description 227, 228
  - diagnostics accesses 231
  - error correction 11
  - flushing 7
  - invalidating a line 231
  - PIPT 227
  - update 210
- externally\_initiated\_reset* (XIR) 317, 318, 319, 323, 339, 342, 348

## F

- F pipeline stage 42
- f* registers 325
- FABSd instruction 493
- FABSq instruction 493
- FABSs instruction 493
- FADD instruction 485
- FADDd instruction 483
- FADDq instruction 483
- FADDs instruction 483
- FALIGNADDR instruction
  - grouping rules 51
- FALIGNDATA instruction 455
  - grouping rules 51
- FAND instruction 514
- FANDNOT1 instruction 515
- FANDNOT1S instruction 515
- FANDNOT2 instruction 515
- FANDNOT2S instruction 515
- FANDS instruction 515

- fast\_data\_access\_MMU\_miss* exception 114, 258, 259, 260, 289
- fast\_data\_access\_MMU\_miss* exception 253
- fast\_data\_access\_protection* exception 114, 251, 258, 259, 260
- fast\_instruction\_access\_MMU\_miss* exception 114, 258, 259, 307, 308
- fast\_instruction\_access\_MMU\_miss* exception 253
- fast\_instruction\_MMU\_miss* exception 289, 307
- FBA instruction 603, 605
- FBE instruction 603
- FBfcc instructions 99, 123, 150, 346, 603, 605
- FBG instruction 603
- FBGE instruction 603
- FBL instruction 603
- FBLE instruction 603
- FBLG instruction 603
- FBN instruction 603, 604
- FBNE instruction 603
- FBO instruction 603
- FBPA instruction 471, 473
- FBPcc instructions 178
- FBPE instruction 471
- FBPfcc instructions 99, 123, 150, 177, 178, 471, 605
- FBPG instruction 471
- FBPGE instruction 471
- FBPL instruction 471
- FBPLE instruction 471
- FBPLG instruction 471
- FBPN instruction 471, 473
- FBPNE instruction 471
- FBPO instruction 471
- FBPU instruction 471
- FBPUE instruction 471
- FBPUG instruction 471
- FBPUGE instruction 471
- FBPUL instruction 471
- FBPULE instruction 471
- FBU instruction 603
- FBUE instruction 603
- FBUG instruction 603
- FBUGE instruction 603
- FBUL instruction 603
- FBULE instruction 603
- fcc*-conditional branches 473, 605
- FCMP\* instructions 123, 486
- FCMPd instruction 486
- FCMPE\* instructions 123, 486
- FCMPEd instruction 486

- FCMPEQ instruction 551
- FCMPEq instruction 486
- FCMPEQ16 instruction 550
- FCMPEQ32 instruction 550
- FCMPEs instruction 486
- FCMPG instruction 551
- FCMPGT16 instruction 550
- FCMPGT32 instruction 550
- FCMPL instruction 551
- FCMPLE16 instruction 550
- FCMPLE32 instruction 550
- FCMPNE instruction 551
- FCMPNE16 instruction 550
- FCMPNE32 instruction 550
- FCMPq instruction 486
- FCMPs instruction 486
- fcn* field of instructions 480
- FDIVd instruction 494
- FDIVq instruction 494
- FDIVs instruction 494
- FdMULq instruction 494
- FdTOi instruction 5, 488, 489
- FdTOq instruction 489
- FdTOs instruction 489
- FdTOx instruction 5, 488, 489
- FEXPAND instruction 155, 553, 558
- FEXPAND instruction, pixel formatting 554
- FEXPAND operation 558
- FFA (f.p./Graphics ALU) pipeline 43
- FFA pipeline 385
- FGA pipeline xlv, 469
- FGM (F.p./Graphics multiply) pipeline 43
- FGM pipeline xlv, 385
- fill register window 84, 158, 346, 573, 575
- fill\_n\_normal* exception 324, 346, 346, 572, 574
- fill\_n\_other* exception 324, 346, 572, 574
- FiTOd instruction 491
- FiTOq instruction 491
- FiTOs instruction 5, 491, 492
- fixed-point scaling 545
- floating point
  - divide/square root 51
  - grouping rules ??–51
  - latencies 50
  - operation statistics 385
  - register file access 45
  - store instructions 51
  - subnormal value generation 124
- floating point complex calculations 76
- floating-point add and subtract instructions 483
- floating-point compare instructions 123, 486, 486
- floating-point condition code bits 605
- floating-point condition codes (*fcc*) fields of FSR register 123, 126, 325, 473, 487, 605
- floating-point data type 65
- floating-point deferred-trap queue (FQ) 566
- floating-point exception 125
- floating-point move instructions 493
- floating-point multiply and divide instructions 494
- floating-point operate (FPop) instructions 125, 129, 150, 178, 346, 348, 612
- floating-point registers 89
- floating-point registers state (FPRS) register 99, 569, 601
- floating-point square root instructions 496
- floating-point state (FSR) register 122, 129, 132, 585, 612, 621
- floating-point trap type (*ftt*) field of FSR register 129
- floating-point trap type (*ftt*) field of FSR register 122, 125, 129, 151, 348, 585, 621
- floating-point trap types
  - IEEE\_754\_exception* 126, 126, 129, 132, 325, 348
  - invalid\_fp\_register* 89, 126, 494, 497
  - numeric values 125
  - sequence\_error* 126
  - unfinished\_FPop* 126, 126, 132, 485, 496
  - unimplemented\_FPop* 126, 132, 485, 487, 489, 491, 492, 496, 529, 531
- floating-point traps
  - precise 568
- FLUSH instruction 497
  - after internal store 221
  - differences from UltraSPARC-I 5
  - grouping rule 53
  - I-cache 7
  - memory ordering control 213
  - self-modifying code 499
- flush register windows instruction 499
- flushing
  - data cache 7
  - differences from UltraSPARC-I 7
  - displacement 231
  - external cache 228
  - instruction cache 7
  - prefetch cache 7
  - TLB 8
  - write cache 7, 228
- FLUSHW instruction 157, 348, 499
- FLUSHW instruction, grouping rule 52

FMOVA instruction 524  
 FMOVcc instruction 524  
 FMOVcc instructions 98, 123, 144, 177, 178, 524, 528, 529, 535  
     grouping rules 54  
 FMOVCS instruction 524  
 FMOVd instruction 493  
 FMOVDec instruction 526  
 FMOVE instruction 524  
 FMOVFA instruction 525  
 FMOVFE instruction 525  
 FMOVFG instruction 525  
 FMOVFGE instruction 525  
 FMOVFL instruction 525  
 FMOVFLE instruction 525  
 FMOVFLG instruction 525  
 FMOVFN instruction 525  
 FMOVFNE instruction 525  
 FMOVFO instruction 525  
 FMOVFU instruction 525  
 FMOVFUE instruction 525  
 FMOVFUG instruction 525  
 FMOVFUGE instruction 525  
 FMOVFUL instruction 525  
 FMOVFULE instruction 525  
 FMOVG instruction 524  
 FMOVGE instruction 524  
 FMOVGU instruction 524  
 FMOVL instruction 524  
 FMOVLE instruction 524  
 FMOVLEU instruction 524  
 FMOVN instruction 524  
 FMOVNE instruction 524  
 FMOVNEG instruction 524  
 FMOVPOS instruction 524  
 FMOVq instruction 493  
 FMOVQcc instruction 526  
 FMOVr instructions 178, 529  
 FMOVRGEZ instruction 529  
 FMOVRGZ instruction 529  
 FMOVRLEZ instruction 529  
 FMOVRLZ instruction 529  
 FMOVRNZ instruction 529  
 FMOVRZ instruction 529  
 FMOV's instruction 493  
 FMOVSec instruction 526  
 FMOVVC instruction 524  
 FMOVVS instruction 524  
 FMUL8SUx16 instruction 543, 546  
 FMUL8ULx16 instruction 543, 547  
 FMUL8x16 instruction 156, 543, 545  
 FMUL8x16AL instruction 543, 546  
 FMUL8x16AU instruction 543, 545  
 FMULd instruction 494  
 FMULD8SUx16 instruction 543, 548  
 FMULD8ULx16 instruction 543, 549  
 FMULq instruction 494  
 FMULs instruction 494  
 FNAND instruction 515  
 FNANDS instruction 515  
 FNEGd instruction 493  
 FNEGq instruction 493  
 FNEGs instruction 493  
 FNOR instruction 514  
 FNORS instruction 514  
 FNOT1 instruction 514  
 FNOT1S instruction 514  
 FNOT2 instruction 514  
 FNOT2S instruction 514  
 FONE instruction 514  
 FONES instruction 514  
 FOR instruction 514  
 formats, instruction 175  
 FORNOT1 instruction 515  
 FORNOT1S instruction 515  
 FORNOT2 instruction 515  
 FORNOT2S instruction 515  
 FORS instruction 514  
*fp\_disabled* exception 99, 102, 150, 323, 346, 473, 485, 489, 491, 492, 494, 496, 503, 504, 505, 506, 529, 531, 535, 582, 585, 587, 605, 612, 622  
*fp\_exception* exception 129  
*fp\_exception\_ieee\_754* "invalid" exception 489  
*fp\_exception\_ieee\_754* exception 102, 103, 123, 129, 130, 131, 325, 348, 485, 489, 491, 492, 496  
*fp\_exception\_other* exception 89, 102, 103, 124, 127, 151, 179, 346, 348, 485, 487, 489, 491, 492, 494, 496, 497, 531  
*fp\_exception\_other* exception 409, 420, 421, 428  
 FPACK instructions 155–??, 553–558  
 FPACK, performance usage 554  
 FPACK16 instruction 155, 553, 554  
 FPACK16 operation 555  
 FPACK32 instruction 553, 556  
 FPACK32 operation 556  
 FPACKFIX instruction 553, 557  
 FPACKFIX operation 558  
 FPADD16 instruction 542



- FPADD16S instruction 542
- FPADD32 instruction 542
- FPADD32S instruction 542
- FPMERGE instruction 553, 559
- FPMERGE instruction, back-to-back execution 554
- FPRS register
  - description 99
  - FEF field 102, 601
  - state after reset 440
- FPSUB16 instruction 542
- FPSUB16S instruction 542
- FPSUB32 instruction 542
- FPSUB32S instruction 542
- FqTOd instruction 489
- FqTOi instruction 488
- FqTOs instruction 489
- FqTOx instruction 488
- FsMULd instruction 494
- FSQRTd instruction 496
- FSQRTq instruction 496
- FSQRTs instruction 496
- FSR
  - aexc field 409
  - cexc field 409
  - ftt field 124
  - nonstandard floating-point operation 124
  - NS field 124
    - = 1 124
    - =0 485
    - =1 485
  - RD field 4
    - state after reset 440
- FSRC1 instruction 514
- FSRC1S instruction 514
- FSRC2 instruction 514
- FSRC2S instruction 514
- FsTOD instruction 489
- FsTOi instruction 5, 488, 489
- FsTOq instruction 489
- FsTOx instruction 5, 488, 489
- FSUB instruction 485
- FSUBd instruction 483
- FSUBq instruction 483
- FSUBs instruction 483
- FXNOR instruction 515
- FXNORS instruction 515
- FXOR instruction 515
- FXORS instruction 515
- FxTOD instruction 5, 491, 492

- FxTOq instruction 491
- FxTOs instruction 5, 491, 492
- FZERO instruction 514
- FZEROS instruction 514

## G

- generating constants 577
- global registers
  - interrupt 114
  - trap 114
- global* registers 80, 82, 82
- global visibility 212
- graphics data format
  - fixed 16-bit 76
- grouping rules 47–51
  - BMASK and BSHUFFLE 469
  - SIAM instruction 576
- GSR
  - byte mask instruction 4
  - fields
    - ALIGN 104
    - IM (interval mode) field 103
    - IRND (rounding) 103
    - MASK 103
    - SCALE 104
  - mask, setting before BSHUFFLE 469
  - new fields 10
  - state after reset 441
  - write instruction latency 51

## H

- halfword
  - addressing 182
  - alignment 141
  - data format 65
- halt 334
- hardware
  - interlocking mechanism 521
  - table walking 258
  - traps 327
- helper
  - cycle 49
  - execution order 49
  - generation 49
  - in pipelines 49

**I**

- i* field of instructions 178, 454, 497, 499, 501, 503, 505, 507, 509, 512, 513, 518, 533, 536, 538, 560, 569, 571, 609, 611, 613, 614, 616, 618, 619
- I pipeline stage 43
- I/D
  - MMU Demap Operation 278
  - MMU TLB Tag Access Registers 279
  - MMU TSB Pointer register 286, 305
- I/D Translation Storage Buffer Register
  - differences from UltraSPARC-I 8
- I/D TSB Tag Target registers 277
- I/O
  - access 220, 222
  - memory 210
  - memory-mapped 211
  - noncacheable address 217
- IC\_miss 373
- IC\_miss\_cancelled 373
- I-cache
  - miss processing 390
  - organization 388
  - organization *illustrated* 388
  - utilization 392
- icc* field of CCR register 96, 98, 454, 476, 518, 534, 592, 593, 596, 607, 610, 611, 617, 618
- icc*-conditional branches 607
- IE, Invert Endianness bit 249
- IEEE Std 754-1985 xlv, 124, 126, 131, 132, 151
- IEEE\_754\_exception* floating-point trap type xlv, 126, 126, 129, 132, 325, 348
- IER register (SPARC V8) 602
- IIU
  - branch prediction statistics 370
  - stall counts 370
- illegal address aliasing 230
- illegal\_instruction* exception 84, 109, 179, 346, 447, 471, 476, 480, 501, 504, 535, 537, 560, 567, 568, 569, 575, 585, 587, 597, 599, 613, 614, 615, 622, 623, 624, 625
- illegal\_instruction* exception 191, 347, 562
- ILLTRAP instruction 346, 500
- images
  - band interleaved 76
  - band sequential 76
- imm\_asi* field of instructions 142, 178, 477, 503, 505, 507, 509, 512, 513, 611, 613, 614
- imm22* field of instructions 178
- I-MMU
  - bypassing E-cache 228
  - context register usage 262
  - disabled 220, 263
  - Enable bit 133, 263
  - enable bits 263
  - and instruction prefetching 220
  - memory operation summary 298
  - Registers: Primary, Secondary, Nucleus 278
  - virtual address translation 294
- IMPDEP2A instruction 346
- IMPDEP2B instruction 346
- implementation
  - dependency xxxviii
- implementation note xlii
- implementation number (*impl*) field of VER register **121**
- implicit
  - ASI 142
  - byte order 115
- in* registers 80, 84, 572
- inexact accrued (*nxa*) bit of *aexc* field of FSR register 132
- inexact current (*nxc*) bit of *cexc* field of FSR register 132
- inexact mask (*NXM*) bit of TEM field of FSR register 128
- inexact quotient 609, 610
- initiated xlvii
- instruction
  - breakpoint, trap priorities 347
  - buffer 390, 394
  - bypass 50
  - conditional branch 45
  - dependency check 48
  - dispatching properties 54
  - Edge 4
  - Edgenc 4
  - execution order 48
  - explicit synchronization 464
  - grouping rules 47–51
  - latency 48, 54
  - multicycle, blocking 48
  - number completed 369
  - prefetch 30, 220
  - SIAM 4
  - window-saving 52
  - with helpers 53
  - writing integer register 49
- Instruction Cache 226
  - physically indexed
  - physically tagged 226
- instruction cache
  - consistency 7

- disabled in RED\_state 435
- effect of mode change 226
- flushing 7
- reference counts 373
- instruction fields
  - a* 177, 470, 474, 477, 604, 607
  - cc0* 177, 473, 474, 486, 533
  - cc1* 177, 473, 474, 486, 533
  - cc2* 177, 533
  - cond* 177, 473, 474, 526, 533, 604, 607
  - const22* 500
  - d16hi* 177, 470
  - d16lo* 177, 470
  - definition xlvi
  - disp19* 178, 473, 474
  - disp22* 178, 604, 607
  - disp30* 178, 476
  - fcn* 480
  - i* 178, 454, 497, 499, 501, 503, 505, 507, 509, 512, 513, 518, 533, 536, 538, 560, 569, 571, 609, 611, 613, 614, 616, 618, 619
  - imm\_asi* 142, 178, 477, 503, 505, 507, 509, 611, 613, 614
  - imm22* 178
  - mmask* 178, 620
  - op3* 178, 454, 477, 480, 497, 499, 501, 503, 505, 507, 509, 512, 513, 518, 538, 566, 569, 571, 609, 611, 613, 614, 616, 618, 619
  - opf* 178, 484, 486, 488, 490, 492, 493, 495, 496
  - opf\_cc* 178, 526
  - opf\_low* 178, 526, 530
  - p* 178, 470, 473, 474
  - rcond* 178, 470, 530, 536
  - rd* 178, 454, 477, 484, 488, 490, 492, 493, 495, 496, 501, 503, 505, 507, 509, 512, 513, 518, 526, 530, 533, 536, 538, 560, 566, 569, 609, 611, 613, 614, 616, 618, 619
  - reserved* 447
  - rs1* 178, 454, 470, 477, 484, 486, 495, 497, 501, 503, 505, 507, 509, 512, 513, 518, 530, 536, 538, 566, 569, 571, 609, 611, 613, 614, 616, 618, 619
  - rs2* 178, 454, 477, 484, 486, 488, 490, 492, 493, 495, 496, 497, 501, 503, 505, 507, 509, 512, 513, 518, 526, 530, 533, 536, 538, 560, 571, 609, 611, 613, 614, 616, 618
  - shcnt32* 178
  - shcnt64* 178
  - simm10* 178, 536
  - simm11* 178, 533
  - simm13* 179, 454, 497, 501, 503, 505, 507, 508, 512, 513, 518, 538, 560, 571, 609, 611, 613, 614, 616, 618
  - sw\_trap#* 179
  - x* 179
- instruction queue, state after reset 442
- instruction set architecture (ISA) xlvi
- instruction TLB 8
- Instruction Translation Lookaside Buffer (iTLB)
  - misses 391
- instruction\_access\_error* exception 323
- instruction\_access\_error* exception 30
- instruction\_access\_exception* (ISA) exception 323
- instruction\_access\_exception* exception 114, 251, 258, 259, 307, 308, 347
- instruction\_access\_exception* exception 263
- instruction\_access\_MMU\_miss* exception 259
- INSTRUCTION\_TRAP register 441
- instructions
  - alignment 141, 141, 456
  - array addressing 154, 457
  - atomic 478
  - atomic load-store 477, 478, 512, 513, 625, 626
  - block load and store 461
  - branch if contents of integer register match condition 469
  - branch on floating-point condition codes 603
  - branch on floating-point condition codes with prediction 471
  - branch on integer condition codes 605
  - causing illegal instruction 501
  - compare and swap 477
  - comparison 148, 592
  - control-transfer (CTIs) 158, 480
  - convert between floating-point formats 489
  - convert floating-point to integer 488
  - convert integer to floating-point 491
  - count of number of bits 560
  - divide 537, 608
  - DONE 113, 479
  - edge handling 155, 481
  - floating-point add and subtract 483
  - floating-point compare 123, 486, 486
  - floating-point move 493
  - floating-point multiply and divide 494
  - floating-point operate (FPop) 125, 129, 150, 612
  - floating-point square root 496
  - flush instruction memory 497
  - flush register windows 499

- formats 175
- generate software-initiated reset 583
- jump and link 159, 501
- load floating-point 611
- load floating-point from alternate space 504
- load integer 506, 612
- load integer from alternate space 508, 614
- load quadword 510
- load-store unsigned byte 478, 512, 626, 627
- load-store unsigned byte to alternate space 513
- logical 517
- logical operate 516
- move floating-point register if condition is true 524
- move floating-point register if contents of integer register satisfy condition 529
- move integer register if contents of integer register satisfies condition 536
- multiply 537, 616, 616
- ordering MEMBAR 157
- partial store 541
- partitioned add/subtract 155, 543
- partitioned multiply 544
- permuting bytes specified by GSR.MASK 468
- pixel compare 156, 551
- pixel component distance 552
- pixel formatting (PACK) 155, 553
- prefetch data 560
- read privileged register 566
- read state register 568, 619
- register window management 157
- reserved 179
- reserved fields 447
- RETRY 113, 479
- RETURN vs. RESTORE 571
- sequencing MEMBAR 157
- set high bits of low word 577
- set interval arithmetic mode 576
- setting GSR.MASK field 154, 468
- shift 147, 578
- shift count 579
- short floating-point load/store 581
- shut down to enter power-down mode 582
- software-initiated reset 583
- store 588
- store floating point 584
- store floating-point into alternate space 586, 586
- store integer 588
- store integer into alternate space 590
- subtract 591, 591
- swap *r* register with alternate space memory 626
- swap *r* register with memory 625
- tagged addition 593
- tagged arithmetic 147
- tagged subtraction 594
- timing 447
- trap on condition codes 596
- trap on integer condition codes 595
- unimplemented 179
- write privileged register 597
- writing privileged register 599
- integer register file access 44
- integer unit (IU)
  - condition codes 98
- interrupt
  - differences from UltraSPARC-I 9
  - enable (IE) field of PSTATE register 117, 322, 324, 347
  - on floating-point instructions 102
  - global registers 114
  - level 117
  - request *xlvi*, 315
  - trap 353
  - vector dispatch 352
  - vector dispatch register 355
  - vector dispatch status register 356
  - vector receive 353
  - vector receive register 357
- Interrupt Vector Dispatch Register 355
- Interrupt Vector Dispatch Status Register 356
- interrupt\_vector* exception 102
- interrupt\_vector* trap 114
- interrupt\_vector\_trap* exception 349
- interval arithmetic
  - support 4
- INTR\_DISPATCH register 442
- INTR\_RECEIVE register 442
- invalid accrued (*nva*) bit of *axc* field of FSR register 131
- invalid current (*mc*) bit of *cexc* field of FSR register 131
- invalid mask (*NVM*) bit of TEM field of FSR register 128
- invalid\_exception* exception 489
- invalid\_fp\_register* floating-point trap type 89, 126, 494, 497
- invalidation
  - prefetch cache 562
- I-SFSR
  - and ASI operations 277
  - bit description 287, 306
  - NF field always 0 306

issued xlvii  
italic font, in assembly language syntax 633  
ITID field of Interrupt Vector Dispatch register 352, 355  
I-TLB  
    state after reset 443

## J

JMPL instruction 30, 45, 261, 436  
    computing target address 159  
    description 501  
    destination register 99  
    does not change CWP 86  
    *mem\_address\_not\_aligned* exception 347  
    reexecuting trapped instruction 571  
jump and link (JMPL) instruction 159, 501

## K

kernel code 358

## L

latency  
    BMASK and BSHUFFLE 469  
    floating-point operations 50  
    FPADD instruction 543  
    partitioned multiply 544  
LD instruction (SPARC V8) 507  
LDD instruction 222, 506, 612  
LDDA instruction 84, 206, 508, 510, 614  
LDDF instruction 141, 349, 502, 611  
*LDDF\_mem\_address\_not\_aligned* exception 141, 323, 349, 506  
LDDFA instruction 141, 191, 460, 504, 541, 580  
LDF instruction 502, 611  
LDFB instruction 504  
LDFSR instruction 53, 123, 125, 223, 347, 611  
LDQF instruction 179, 502, 611  
LDQFA instruction 504  
LDSB instruction 222, 506, 612  
LDSBA instruction 508, 614  
LDSH instruction 222, 506, 612  
LDSHA instruction 508, 614  
LDSTUB instruction 143, 217, 512, 513  
LDSTUBA instruction 512, 513

LDSW instruction 222, 506, 612  
LDSWA instruction 508, 614  
LDUB instruction 506, 612  
LDUBA instruction 508, 614  
LDUH instruction 506, 612  
LDUHA instruction 508, 614  
LDUW instruction 506, 612  
LDUWA instruction 508, 614  
LDX instruction 506, 612  
LDXA instruction 508, 614  
LDXFSR instruction 122, 123, 125, 223, 347, 502, 611  
leaf subroutine 392  
level 2 cache 227  
level-1 cache 225  
    flushing 229  
little-endian  
    ordering in partial store instructions 541  
little-endian byte order xlvii, 115, 140  
load floating-point from alternate space instructions 504  
load floating-point instructions 611  
load instructions xlvii  
load instructions, getting data from store queue 222  
load integer from alternate space instructions 508, 614  
load integer instructions 506, 612  
load quadword atomic 206, 510  
load recirculation 224  
LoadLoad MEMBAR relationship 520  
LoadLoad predefined constant 637  
loads  
    from alternate space 98, 142  
load-store alignment 141, 141  
load-store instructions 143  
    compare and swap 477  
    definition xlvii  
    and *fast\_data\_access\_protection* exception 349  
    load-store unsigned byte 478, 512, 626, 627  
    load-store unsigned byte to alternate space 513  
    swap *r* register with alternate space memory 626  
    swap *r* register with memory 478, 625  
LoadStore MEMBAR relationship 520  
LoadStore predefined constant 637  
local registers 80, 84, 573  
logical instructions 517  
Lookaside MEMBAR relationship 520  
Lookaside predefined constant 637  
lower registers dirty (DL) field of FPRS register 100

**M**

- M pipeline stage 45
- machine state
  - after reset 438
  - in RED\_state 438
- mask number (mask) field of VER register 121
- maximum trap levels (MAXTL) field of VER register 122
- MAXTL 117, 317, 319, 336, 337, 338, 435, 583
- may (keyword) xvii
- MCU 31, 436
- Mem\_Addr\_CTL register 442
- Mem\_Addr\_Dec register 442
- mem\_address\_not\_aligned* exception 141, 258, 259, 260, 277, 287, 306, 347, 479, 502, 503, 504, 506, 507, 508, 509, 510, 571, 572, 582, 585, 587, 589, 590, 614, 615, 622, 623, 625, 626, 628
- mem\_address\_not\_aligned* exception 191, 192, 206, 261
- mem\_address\_not\_aligned* trap 394
- Mem\_Timing\_CSR register 442
- MEMBAR
  - #LoadLoad 212, 520, 637
  - #LoadStore 212, 520, 637
  - #LoadStore and block store 464
  - #Lookaside 210
  - #MemIssue 210, 521
  - #StoreLoad 520, 637
    - and BLD 464
    - and BST 464
    - for strong ordering 521
  - #StoreStore 499, 520, 637
    - and BST 464
    - code example 212
  - #Sync 231, 277, 290, 308
    - after BST 463
    - after internal ASI store 221
    - BLD and BST 463
    - E-cache flushing 7
    - semantics 214
    - for strong ordering 521
- instruction 157, 178, 354, 498, 519, 569, 620
  - explicit synchronization 212
  - grouping rules 53
  - memory ordering 213
  - side-effect accesses 220
  - single group 53
- QUAD\_LDD requirement 523
- rules for interlock implementation 521
- UltraSPARC-III specifics 521
- membar\_mask* 637
- MemIssue MEMBAR relationship 520
- MemIssue predefined constant 637
- memory
  - access instructions 143
  - bank, access counts 378
  - cached 210
  - current model, indication 210
  - global visibility of memory accesses 212
  - location 210
  - models
    - and block operations 464
    - ordering and block store 464
    - partial store order (PSO) 209, 464
    - relaxed memory order (RMO) 464
    - strongly ordered 222, 521
    - total store order (TSO) 209
    - total store order (TSO)TSO 464
  - noncacheable, scratch 228
  - ordering 212
  - subsystem, differences from UltraSPARC-I 6
  - synchronization 213
- Memory Management Unit (MMU) 245
- memory\_model (MM) field of PSTATE register 115
- memory-mapped I/O 211
- merge buffer 222
- mispredict signal 45
- mmask* field of instructions 178, 620
- MMU ??–293
  - accessing registers 277
  - behavior during reset 263
  - bypass 294
  - D Synchronous Fault Address Register 277
  - D TSB Secondary Extension Registers 277
  - demap 290, 308
    - all 290, 308
    - context 290, 292, 308, 310
    - operation syntax 290, 309
    - page 290, 291, 308, 310
  - disable 263
  - global registers 114, 258
  - I/D Synchronous Fault Status Registers 277, 287, 306
  - I/D TLB Data Access Registers 278
  - I/D TLB Data In Registers 278
  - I/D TLB Tag Access register 277
  - I/D TLB Tag Read Register 278
  - I/D TSB 64K Pointer Registers 278
  - I/D TSB 8K Pointer Registers 277
  - I/D TSB Extension Registers 285, 305

- I/D TSB Nucleus Extension Register 277
- I/D TSB Primary Extension Register 277
- I/D TSB register 277
- I/D TSB Registers 284, 303
- page sizes 245
- Physical Watchpoint Address 277
- Primary Context Register 277
- Secondary Context Register 277
- SPARC V9 compliance 264
- Synchronous Fault Address Registers 289
- Synchronous Fault Status Register
  - fault types 289, 307
- Tag Target Registers 283, 303
- TLB Diagnostic Access Address 292, 310
- Virtual Watchpoint Address 277
- mode
  - privileged 109
  - user 98
- MOVA instruction 531
- MOVCC instruction 531
- MOVcc instructions 98, 123, 144, 177, 178, 528, 529, 535
  - grouping rules 54
- MOVCS instruction 531
- move floating-point register if condition is true 524
- move floating-point register if contents of integer register satisfy condition 529
- MOVE instruction 531
- move integer register if contents of integer register satisfies condition instructions 536
- MOVFA instruction 532
- MOVFE instruction 532
- MOVFG instruction 532
- MOVFGI instruction 532
- MOVFL instruction 532
- MOVFLE instruction 532
- MOVFLG instruction 532
- MOVFN instruction 532
- MOVFNE instruction 532
- MOVFO instruction 532
- MOVFU instruction 532
- MOVFUE instruction 532
- MOVFUG instruction 532
- MOVFUGE instruction 532
- MOVFUL instruction 532
- MOVFULE instruction 532
- MOVG instruction 531
- MOVGE instruction 531
- MOVGU instruction 531
- MOVL instruction 531
- MOVLE instruction 531
- MOVLEU instruction 531
- MOVN instruction 531
- MOVNE instruction 531
- MOVNEG instruction 531
- MOVPOS instruction 531
- MOVR instructions
  - grouping rules 54
- MOVr instructions 178, 536
- MOVRRGEZ instruction 536
- MOVRRGZ instruction 536
- MOVRRLEZ instruction 536
- MOVRRLZ instruction 536
- MOVRRNZ instruction 536
- MOVRRZ instruction 536
- MOVVC instruction 531
- MOVVS instruction 531
- MS pipeline
  - description 43
  - E-stage bypass 48
  - instruction requirements 4
  - integer instruction execution 44
  - and W-stage 46
- multiple unsigned condition codes, emulating 146
- multiply instructions 537, 616, 616
- multiprocessor synchronization instructions 478, 626, 627
- multiprocessor system 498, 626, 627
- MULX instruction 537
- must (keyword) xviii
- mutual exclusion, atomic instructions 216
- M-way set-associative TSB 251

**N**

- NaN (not-a-number)
  - converting floating-point to integer 489
  - quiet 487
  - signalling 123, 487, 490
- negative (*N*) bit of condition fields of CCR 97
- next program counter (nPC) 99, 110, 180, 480, 539
- nonallocating cache 393
- noncacheable
  - accesses 211
  - I/O address 217
  - instruction prefetch 30, 221
  - store compression 222

- store compression, differences from UltraSPARC-I 10
- store merging enable 133
- nonfaulting
  - ASIs and atomic accesses 217
  - load
    - and TLB miss 218
    - behavior 218
    - use by optimizer 218
- nonfaulting load xlvi, 250, 289
- nonleaf routine 502
- nonprivileged
  - mode xliii, 125
  - software 99
- nonprivileged trap (NPT) field of TICK register 570
- nonstandard floating-point operation 124
- NOP instruction 473, 538, 596, 604, 607
- normal traps 316, 326, 335, 335, 337, 337, 338, 339
- note
  - implementation xlii
  - programming xlii
- nPC register 438
- nPC register, *See* next program counter (nPC)
- NS field of FSR 124
- Nucleus code 358
- Nucleus Context Register 278
- NWINDOWS 83, 84, 573

## O

- op3* field of instructions 178, 454, 477, 480, 497, 499, 501, 503, 505, 507, 509, 512, 513, 518, 538, 566, 569, 571, 609, 611, 613, 614, 616, 618, 619
- opcode
  - definition xlix
- opf* field of instructions 178, 484, 486, 488, 490, 492, 493, 495, 496
- opf\_cc* field of instructions 178, 526
- opf\_low* field of instructions 178, 526, 530
- OR instruction 517
- ORcc instruction 517
- ordering
  - block load 464
  - block store 464
- ordering MEMBAR instructions 157
- ORN instruction 517
- ORNcc instruction 517
- other windows (OTHERWIN) register 119, 500, 566, 573, 598

- OTHERWIN register 440
- out* register #7 82
- out* registers 84, 572
- overflow (*V*) bit of condition fields of CCR 97, 147
- overflow accrued (*ofa*) bit of *aexc* field of FSR register 131
- overflow current (*ofc*) bit of *cexc* field of FSR register 131
- overflow mask (*OFM*) bit of TEM field of FSR register 128

## P

- p* field of instructions 178, 470, 473, 474
- PA Data Watchpoint Register
  - DB\_PA field 137
  - format 137
- PA\_watchpoint* exception 136, 349
- PA\_WATCHPOINT register 441
- packed-to-planar conversion 155, 559
- partial store instruction 51
- partial store instructions 540
- partitioned multiply instructions 544
- PC register 438
- PC register, *See* program counter (PC)
- PC, Instr\_cnt 369
- PC\_1st\_rd 374
- PC\_2nd\_rd 374
- PC\_counter\_inv 374
- PC\_hard\_hit 374
- PC\_MS\_misses 374
- PC\_soft\_hit 374
- PCR
  - access 364, 365
  - extension 10
  - fields
    - PRIV 365
    - ST(system trace enable) field 365
    - SU (select upper bits of PIC) field 365
    - UT (user trace enable) field 365
  - function
    - Cycle\_cnt 369
    - DC\_hit 373
    - Dispatch0\_2nd\_br 371
    - Dispatch0\_br\_target 371
    - Dispatch0\_IC\_miss 370
    - Dispatch0\_mispred 371
    - EC\_ref 375
    - EC\_snoop\_inv 375



- EC\_snoop\_wb 375
- EC\_wb 375
- EC\_write\_hit\_clean 375
- IC\_ref 373
- SI\_snoops 384
- ST field 369
- state after reset 441
- UT field 369
- PDIST instruction 552
- PDIST, instruction latency 51
- performance hints
  - FPACK usage 554
  - FPADD usage 543
  - logical operate instructions 516
  - partitioned multiply usage 544
- physical address
  - data watchpoint 137
- Physical Indexed Caches 226
- Physical Tagged Caches 226
- PIC register
  - and PCR 364
  - access 364, 365
  - event logging 366
  - PIC0 Events 385
  - PIC1 Events 385
  - PICL field 367
  - SL selection bit field encoding 385
  - state after reset 441
- PIL register 439
- pipeline
  - A0 43, 44
  - A1 43
  - BR 43
  - conditional moves 53
  - dependencies 44
  - FFA 43, 385
  - FGA xlv, 469
  - FGM xlv, 43, 385
  - MS 4, 43, 44, 46
  - stages
    - A 42, 45
    - B 43
    - C 45, 46
    - D 46, 370
    - E 44
    - F 42
    - I 43
    - M 45
    - mnemonics 38
    - R 44, 372
    - T 46
    - W 46
  - stalls, causes 370
- PIPT cache
  - E-cache 7
  - I-cache 7
  - level 2 caches 227
  - W-cache 6
- pixel instructions
  - comparison 156, 551
  - component distance 552
  - formatting 155, 553
- planar-to-packed conversion 559
- POK pin 31, 436
- POPC instruction 179, 559
- power failure 323, 342
- power\_on\_reset* (POR) 317, 319, 339
- power-on reset (POR) 107, 108
  - system reset when Reset pin activated 31, 437
- precise floating-point traps 568
- precise trap
  - conditions for 321
  - software actions 321
  - vs. disrupting trap 322
- predefined constants
  - LoadLoad 637
  - lookaside 637
  - MemIssue 637
  - StoreLoad 637
  - StoreStore 637
  - Sync 637
- predict bit 470
- prefetch
  - differences from UltraSPARC-I 4
  - instruction, noncacheable 30
  - instructions 220
  - invalidate prefetched line 4
  - noncacheable data 562
- Prefetch Cache
  - physically indexed
    - physically tagged 227
- prefetch cache
  - characteristics 6
  - corrupted by block load 401
  - flushing 7
  - invalidation 402, 562
  - valid bits 31, 436
- prefetch data instruction 560

PREFETCH instruction 164, 560  
     descriptions 218  
     execution 400  
     external cache allocation 228  
     P-cache/E-cache 399  
     types 399, 562  
*prefetch\_fcn* 637  
 PREFETCHA instruction 560  
 Primary Context Register 278  
 priority  
     traps 324, 331  
     VA vs. PA\_watchpoint 136  
 privileged  
     mode 109  
     registers 109  
     software 84, 116, 125, 142, 326, 499  
 privileged (PRIV) field of PSTATE register 117, 347, 478, 505, 514, 570, 587, 590, 624, 627  
 privileged mode (PRIV) field of PSTATE register 117  
 privileged registers 52  
*privileged\_action* exception 98, 142, 258, 259, 260, 275, 323, 347, 355, 356, 357, 479, 505, 506, 509, 510, 514, 570, 587, 590, 615, 624, 625, 628  
*privileged\_action* exception 189, 210, 217, 363, 364, 366  
     PIC access 365  
*privileged\_instruction* exception (SPARC V8) 348  
*privileged\_opcode* exception 323, 347, 358, 480, 568, 575, 599  
*privileged\_opcode* exception 347  
 processor  
     front end components 387  
     halt 334  
     state diagram 316  
 processor interrupt level (PIL) register 117, 322, 324, 347, 358, 566, 598  
 processor pipeline  
     address stage 42  
     branch target computation stage 43  
     cache stage 45  
     done stage 46  
     execute stage 44  
     fetch stage 42  
     instruction issue 43  
     register stage 44  
     trap stage 46  
 processor state (PSTATE) register 83, 110, 112, 115, 315, 317, 480, 566, 598  
 processor states  
     error\_state 317, 320, 334, 335, 336, 337, 338, 344  
     execute\_state 334, 336, 337, 338  
     RED\_state 317, 319, 326, 334, 335, 336, 337, 338, 339, 341, 344  
 program counter (PC) 99, 109, 180, 315, 477, 480, 501, 539  
 programming note xlii  
 PSO memory model 209, 212, 213, 220  
 PSR register (SPARC V8) 602  
 PSTATE  
     AM field 116, 117, 185  
     global register selection encodings 113  
     IE field 102, 103, 358  
     IG field 113, 114, 354  
     *illegal\_instruction* exception 347  
     MG field 113, 114  
     MM field 210  
     PEF field 601  
     PRIV field xlviii, xlix, 210, 217, 251, 260  
     RED field 102  
         clearing DCUCR 435  
         exiting RED\_state 30, 221, 436  
         explicitly set 435  
     register 114  
     state after reset 439  
     WRPR instruction and BST 463

**Q**  
 Quad FPop instructions 179  
 quad load instruction 222, 523  
 quadword  
     addressing 183  
     alignment 141  
     data format 65  
     definition 1  
 queue  
     instruction, state after reset 442  
     store, state after reset 442  
 quiet NaN (not-a-number) 123, 427, 487

**R**  
 R pipeline stage 44  
*r* register  
     #15 82  
     categories 81  
     alignment 613, 615

- rational quotient 609
- R-A-W
  - Bypass Enable bit in DCUCR 133
  - bypassing algorithm 222
  - bypassing data from store queue 133
  - detection algorithm 224
- rcond* field of instructions 178, 470, 530, 536
- RD field of FSR register 4
- rd* field of instructions 178, 454, 477, 484, 488, 490, 492, 493, 495, 496, 501, 503, 505, 507, 509, 512, 513, 518, 526, 530, 533, 536, 538, 560, 566, 569, 609, 611, 613, 614, 616, 618, 619
- RDASI instruction 568, 568, 619
- RDASR instruction 100, 347, 364, 365, 568, 568, 619, 620
  - dispatching 52
  - forcing bubbles before 52
- RDCCR instruction 56, 568, 568, 619
- RDDCR instruction 568
- RDFPRS instruction 568, 568, 619
- RDGSR instruction 568
- RDPC instruction 99, 568, 568, 619
- RDPCR instruction 568
- RDPIC instruction 365, 568
- RDPR FQ instruction 179
- RDPR instruction 109, 113, 118, 120, 347, 566, 570
  - dispatching 52
  - forcing bubbles before 52
- RDSOFTINT instruction 568
- RDSTICK instruction 568
- RDSTICK\_CMPR instruction 568
- RTICK instruction 568, 568, 570, 619
- RTICK\_CMPR instruction 568
- RDY instruction 96
- Re\_DC\_miss counter 372
- Re\_EC\_miss counter 372
- Re\_FPU\_bypass counter 372
- Re\_PC\_miss counter 372
- Re\_RAW\_miss counter 372
- read privileged register (RDPR) instruction 566
- read state register instructions 568, 619
- real memory 210
- recirculation instrumentation 372
- RED\_state 263, 264, 316, 317, 319, 326, 334, 335, 336, 337, 338, 339, 341, 344
  - exiting 221
  - MMU behavior 263
  - restricted environment 319
  - trap vector 32, 318, 438
- RED\_state (RED) field of PSTATE register 115, 317
- RED\_state trap table 326
- reference MMU 633
- register
  - access
    - floating-point 45
    - integer 44
  - Floating-Point Status (FSR) 124
  - global trap 114
  - PSTATE 114
  - SOFTINT 10
  - STICK 10
  - values after reset 438
- register window management instructions 157
- register windows 84
  - clean 120, 348
  - fill 84, 158, 346, 573, 575
  - spill 84, 158, 348, 573, 575
- registers
  - accessing MMU registers 277
  - address space identifier (ASI) 315, 480, 505, 509, 513, 562, 587, 590, 601, 615, 624, 627
  - alternate global* 82
  - ancillary state registers (ASRs) 96, 100
  - ASI 98, 110
  - CANRESTORE 119
  - CANSAVE 118
  - clean windows (CLEANWIN) 119, 119, 566, 598
  - CLEAR\_SOFTINT 358
  - clock-tick (TICK) 347
  - condition codes register (CCR) 110, 315, 454, 480, 601, 618
  - control and status 96
  - current window pointer (CWP) 84, 110, 118, 119, 120, 315, 480, 500, 566, 573, 598
  - dispatch control register (DCR) 100
  - f* (floating point) 325
  - floating-point 89
  - floating-point registers state (FPRS) 99, 569, 601
  - floating-point state (FSR) 122, 129, 132, 612, 621
  - global* 80, 82, 82
  - IER (SPARC V8) 602
  - in* 80, 84, 572
  - Interrupt Vector Dispatch register 355
  - Interrupt Vector Dispatch Status register 356
  - Interrupt Vector Receive register 357
  - local* 80, 84, 573
  - MMU Tag Target 283, 303
  - other windows (OTHERWIN) 119, 500, 566, 573,

- 598
- out* 84, 572
- out* #7 82
- PC 99
- performance control (PCR) 364
- privileged 109
- processor interrupt level (PIL) 117, 566, 598
- processor state (PSTATE) 83, 110, 112, 115, 315, 317, 480, 566, 598
- PSR (SPARC V8) 602
- r* 81
- r* register #15 82
- restorable windows (CANRESTORE) 84, 119, 120, 566, 573, 575, 598
- savable windows (CANSAVE) 84, 118, 119, 500, 566, 573, 575, 598
- SET\_SOFTINT 358
- SOFTINT 358
- TBR (SPARC V8) 602
- TICK 106, 107, 566, 598
- TICK\_COMPARE 107
- trap base address (TBA) 111, 315, 325, 566, 598
- trap level (TL) 109, 112, 117, 117, 120, 122, 315, 480, 566, 567, 575, 583, 598, 599
- trap next program counter (TNPC) 110, 566, 598
- Trap Program Counter 308
- trap program counter (TPC) 566, 568, 598
- trap state (TSTATE) 110, 113, 480, 566, 598
- trap type (TT) 110, 112, 120, 326, 343, 566, 596, 598
- update 289, 308
- version register (VER) 120, 566
- WIM (SPARC V8) 602
- window state (WSTATE) 118, 120, 500, 566, 573, 598
- Y 96, 96, 609, 616, 618, 631
- reserved
  - fields in instructions 447
  - instructions 179
- reset
  - externally\_initiated\_reset* (XIR) 317, 318, 319, 323, 339, 342, 348
  - global 263
  - power\_on\_reset* (POR) 317, 319, 339, 347, 347
  - power-on 107, 108
  - processing 317
  - PSTATE.RED 435
  - register values after reset 438
  - request 317, 347, 348
  - reset trap 107, 108, 322, 323
  - software Initiated Reset (SIR) 436
  - software\_initiated\_reset* (SIR) 317, 319, 323, 334, 343, 348
  - system 31, 437
  - watchdog\_reset* (WDR) 339, 342, 349
- Reset pin 31, 437
- restorable windows (CANRESTORE) register 84, 119, 120, 566, 573, 575, 598
- RESTORE instruction 572–574
  - actions 158
  - and current window 85
  - decrementing CWP register 84
  - fill trap 346
  - followed by SAVE instruction 86
  - managing register windows 157
  - operation 572
  - performance trade-off 573
  - and restorable windows (CANRESTORE) register 119
  - restoring register window 573
  - SPARC V9 vs. SPARC V8 119
- RESTORED instruction 158, 574, 574, 574
  - use by privileged software 157
- RESTORED instruction, single group 52
- restricted address space identifier 142
- restricted ASI 189, 210, 275
- RETRY instruction 98, 103, 114, 159, 315, 317, 322, 347, 479
  - after internal store to ASI 221
  - and BST 463
  - exiting RED\_state 30, 436
  - grouping rules 53
  - restoring AG, IG, MG bits 113
  - use with IFPOE 103
  - when TSTATE uninitialized 31, 437
- Return Address Stack (RAS) 392
- RETURN instruction 45, 261, 570–572
  - computing target address 159
  - destination register 99
  - fill trap 346
  - mem\_address\_not\_aligned* exception 347
  - operation 570
  - reexecuting trapped instruction 571
- Rfr\_CSR register 442
- RMO memory model 209, 212, 213, 220, 464
- rounding
  - behavior in GSR 103
  - for floating-point results 124
  - in signed division 610

rounding direction (RD) field of FSR register 124, 484, 488, 490, 492, 495, 497  
 routine, nonleaf 502  
*rs1* field of instructions 178, 454, 470, 477, 484, 486, 495, 497, 501, 503, 505, 507, 509, 512, 513, 518, 530, 536, 538, 566, 569, 571, 609, 611, 613, 614, 616, 618, 619  
*rs2* field of instructions 178, 454, 477, 484, 486, 488, 490, 492, 493, 495, 496, 497, 501, 503, 505, 507, 509, 518, 526, 530, 533, 536, 538, 560, 609, 611, 613, 614, 616, 618  
 R-stage stall counts 372  
 Rstall\_FP\_use counter 372  
 Rstall\_IU\_use counter 372  
 Rstall\_storeQ counter 372  
 RSTVaddr 32, 318, 326, 438

## S

savable windows (CANSERVE) register 84, 118, 119, 500, 566, 573, 575, 598  
 SAVE instruction 572–574  
   actions 158  
   after RESTORE instruction 571  
   *clean\_window* exception 348  
   and current window 85  
   decrementing CWP register 84  
   leaf procedure 502  
   and local/out registers of register window 86  
   managing register windows 157  
   no clean window available 120  
   number of usable windows 119  
   operation 572  
   performance trade-off 573  
   and savable windows (CANSERVE) register 118  
   SPARC V9 vs. SPARC V8 119  
   spill trap 348  
 SAVED instruction 157, 158, 574, 574, 574  
 SAVED instruction, single group 52  
 Scalable Processor Architecture *see* SPARC  
 scaling of the coefficient 545  
 SDIV instruction 96, 608  
 SDIVcc instruction 96, 608  
 SDIVX instruction 537  
 Secondary Context Register 278  
 self-modifying code 498  
*sequence\_error* floating-point trap type 126, 348  
 sequencing MEMBAR instructions 157

Set Interval Arithmetic Mode (SIAM) instruction 4  
 SET\_SOFTINT pseudo-register 358  
 SETCC instruction, grouping 49  
 SETHI instruction 147, 148, 178, 539, 577, 577  
 SFAR Fault Address field 289  
 SFSR  
   bit description 287, 306  
   extensions 8  
   FT field 8  
     extension: I/D TLB miss 8  
     FT = 10 218  
     FT = 2 211, 218, 220  
     FT = 4 217  
     FT = 8 217, 218  
   NF field 8  
   state after reset 442  
   update policy 289, 308  
 shall (keyword) li  
*shcnt32* field of instructions 178  
*shcnt64* field of instructions 178  
 shift count encodings 579  
 shift instructions 147, 148, 578  
 short floating-point load and store instructions 191, 580  
 short floating-point load instruction 223  
 should (keyword) li  
 SHUTDOWN instruction 582  
   differences from UltraSPARC-I 5  
 SIAM instruction 575  
   grouping rules 51  
   interval arithmetic support 4  
   rounding 576  
   setting GSR fields 576  
 side effect  
   accesses 211, 220  
   and block load 464  
   instruction placement 220  
   instruction prefetching 220  
   visible 211  
 SIGM instruction 435  
 signalling NaN (not-a-number) 123, 487, 490  
 signed integer data type 65  
 sign-extended 64-bit constant 179  
*simmm10* field of instructions 178, 536  
*simmm11* field of instructions 178, 533  
*simmm13* field of instructions 179, 454, 497, 501, 503, 505, 507, 509, 512, 513, 518, 538, 560, 571, 609, 611, 613, 614, 616, 618  
 single-instruction group 48, 49, 51, 52, 53, 56  
 SIR instruction 32, 323, 343, 348, 438, 583, 601

- grouping rule 53
- SLL instruction 578, 578
- SLLX instruction 578, 578
- SMUL instruction 96, 616
- SMULcc instruction 96, 616
- snooping
  - snoop counts 384
- SOFTINT register 10, 358, 440
- software interrupt (SOFTINT) register
  - clearing 358
  - in code sequence for Interrupt Receive 353
  - scheduling interrupt vectors 358
  - setting 358
- software prefetch 399
- software statistics, counters 384
- software translation table 246
- software trap 326, 327, 327, 596
- software\_initiated\_reset* (SIR) 32, 317, 319, 323, 334, 339, 343, 348, 438, 583
- Software-Initiated Reset (SIR) 53, 436
- SPARC xxxvii
  - Architecture Manual, Version 9* xxxvii
  - brief history xxxvii
  - International, address of xxxviii
  - V9, architecture xxxvii
- SPARC V8 compatibility
  - ADDC/ADDCcc renamed 455
  - current window pointer (CWP) register differences 119
  - delay instruction 159
  - delay instruction fetch 162
  - executing delayed conditional branch 162
  - existing nonprivileged SPARC V8 software 83
  - instruction between FBfcc /FBPfcc 473
  - LD, LDUW instructions 507
  - level 15 interrupt 118
  - read state register instructions 570
  - STA instruction renamed 590
  - STBAR instruction 521, 620
  - STD instruction 623
  - STDA instruction 625
  - STFSR instruction 621
  - tagged add instructions 629
  - tagged subtract instructions 630
  - Ticc instruction 597
  - UNIMP instruction renamed 501
  - window\_overflow* exception superseded 348
  - window\_underflow* exception superseded 346
  - write state register instructions 602
- SPARC V9
  - compliance xlix, 264
  - special traps 316, 326
  - speculative load 211
  - spill register window 84, 158, 348, 573, 575
  - spill windows 573
  - spill\_n\_normal* exception 348, 500, 574
  - spill\_n\_other* exception 348, 500, 574
  - SRA instruction 578, 578
  - SRAM
    - changes 9
    - new diagnostic registers 9
  - SRAX instruction 578, 578
  - SRL instruction 578, 578
  - SRLX instruction 578, 578
  - SSM mode
    - globally visible memory access 212
  - stable storage 231
  - stack frame 573
  - stalls
    - counted 370
    - pipeline 370
    - R Stage counts 372
  - STB instruction 588
  - STBA instruction 589
  - STBAR instruction 213, 521, 569
  - STDA instruction 84
  - STDF instruction 141, 349, 584
  - STDF\_mem\_address\_not\_aligned* exception 141, 323, 349, 585, 587
  - STDFA instruction 141, 191, 460, 540, 580, 586, 586
  - STF instruction 584
  - STFA instruction 586
  - STFSR instruction 122, 123, 125, 347
  - STH instruction 588
  - STHA instruction 589
  - STICK register 10, 440, 568
  - STICK\_COMPARE register 108, 440, 568
  - STICK\_INT 358
  - store
    - buffer
      - merging 220
    - compression 211, 222
    - instructions, giving data to a load 223
    - noncacheable, coalescing 222
    - queue
      - R-stage stall count 372
      - state after reset 442
  - store buffer 394

- store floating-point into alternate space instructions 586
- store instructions lii, 349
- StoreLoad MEMBAR relationship 520
- StoreLoad predefined constant 637
- stores to alternate space 98, 142
- StoreStore MEMBAR relationship 520
- StoreStore predefined constant 637
- STQF instruction 179, 584
- STQFA instruction 586, 586
- strongly ordered memory model 222, 521
- STW instruction 588
- STWA instruction 589
- STX instruction 588
- STXA instruction 589
- STXFSR instruction 122, 123, 125, 347, 584
- SUB instruction 591, 591
- SUBC instruction 591, 591
- SUBcc instruction 148, 591, 591
- SUBCcc instruction 591, 591
- subtract instructions 591
- supervisor software 83, 126, 142, 315, 334, 343
- supervisor-mode trap handler 326
- SW\_count\_0 384
- SW\_count\_1 384
- sw\_trap#* field of instructions 179
- SWAP instruction 217, 512, 513, 625
- swap *r* register with alternate space memory instructions 626
- swap *r* register with memory instructions 478, 625
- SWAPA instruction 512, 513, 626
- Sync MEMBAR relationship 520
- Sync predefined constant 637
- Synchronous Fault Address Register (SFAR) 289
- Synchronous Fault Status Register (SFSR)
  - fault types 289, 307
  - register bits 287, 306
- synthetic instructions in assembler 638–??
- system interface
  - registers 10
  - statistics, counters 384
- system interface unit (SIU) instructions 45
- system software 348, 498
- system timer interrupt, STICK\_INT 358

**T**

- T pipeline stage 46
- TA instruction 595
- TADDcc instruction 147, 592
- TADDccTV instruction 148, 348
- Tag Access Register 253, 279, 280, 298, 299
- tag overflow 147
- tag\_overflow* exception 148, 348, 593, 594, 628, 629, 630
- tagged arithmetic instructions 147
- tagged word data format 65
- tagged words 65
- TBA register 439
- TBR register (SPARC V8) 602
- TCC instruction 595
- Tcc instructions 98, 177, 179, 315, 326, 327, 347, 348, 595
- TCS instruction 595
- TE instruction 595
- TG instruction 595
- TGE instruction 595
- TGU instruction 595
- Ticc instruction 5
- Ticc instruction (SPARC V8) 597
- TICK register
  - state after reset 440
- TICK\_COMPARE register 107, 440
- TICK\_INT 358
- timer interrupt, TICK\_INT 358
- timing of instructions 447
- TL instruction 595
- TL register 439, 599
- TLB
  - and 3-dimensional arrays 459
  - and RED\_state 263
  - CAM Diagnostic Register 292, 310
  - data access 45
  - Data Access register 8, 281, 300, 301
  - Data In register 253, 281, 301
  - data in register 281, 281, 300, 301
  - Diagnostic Register 8
  - differences from UltraSPARC-I 8
  - D-TLB state after reset 443
  - flushing 8
  - hit lii
  - instruction 259
  - I-TLB state after reset 443
  - miss
    - fast handling 258
    - handler 247, 253
    - MMU behavior 247
    - reloading TLB 251

- miss and nonfaulting load 218
- miss counts 373
- miss handler 253
- miss/refill sequence 253
- missing entry 253
- required conditions 258
- specialized miss handler code 264
- Tag Access Registers 279, 298
- TLE instruction 595
- TLEU instruction 595
- TN instruction 595
- TNE instruction 595
- TNEG instruction 595
- TNPC register 321, 439
- total store order (TSO) memory model 115
- TPC register 439
- TPOS instruction 595
- Translation
  - Storage Buffer (TSB) 251, 284, 285, 303, 304
  - Table Entry (TTE) 248, 259
- Translation Table Entry, *See TTE*
- trap
  - See also exceptions and traps*
  - atomic accesses 217
  - atomic instructions 217
  - definition 315
  - ECC\_error* 349
  - fp\_disabled*
    - GSR access 601
  - fp\_disabled* 102
  - fp\_exception\_ieee\_754* 102
  - fp\_exception\_other* 102, 124
  - level 117
    - TL = MAXTL 435
    - TL = MAXTL - 1 435
  - model 323
  - noncacheable accesses 211
  - priority 324, 331
  - processing 334
  - stack 113, 336, 337, 338
  - VA\_/PA\_watchpoint* 136
  - vector, RED\_state 318
- trap base address (TBA) register 111, 315, 325, 566, 598
- trap categories
  - deferred 322
  - disrupting 322, 322, 323
  - precise 322
  - reset 323
- trap enable mask (TEM) field of FSR register 128, 128,
  - 129, 324, 325, 348
- trap globals 114
- trap handler 480
  - ECC errors 349
  - supervisor-mode 326
  - user 126, 425
- trap level (TL) register 109, 112, 117, 117, 120, 122,
  - 315, 480, 566, 567, 575, 583, 598, 599
- trap next program counter (TNPC) register 110, 566, 598
- trap on integer condition codes instructions 595
- trap program counter (TPC) register 308, 566, 568, 598
- trap state (TSTATE) register 110, 113, 480, 566, 598
- trap type (TT) register 110, 112, 120, 326, 343, 566, 596,
  - 598
- trap\_instruction* (ISA) exception 323, 348, 596, 597
- trap\_little\_endian* (TLE) field of PSTATE register 115,
  - 115
- traps
  - See also exceptions and trap*
  - deferred 321
  - disrupting 321
  - hardware 327
  - normal 316, 326, 335, 335, 337, 337, 338, 339
  - precise 321
  - reset 321, 322, 323, 334
  - software 327, 596
  - software\_initiated\_reset* (SIR) 339
  - special 316, 326
  - window fill 326
  - window spill 326
- TSB 391
  - cacheability 252
  - caching 252
  - demap operation 291, 309
  - Direct Pointer registers 286, 305
  - Extension Register 277, 285, 304
  - Extension Registers
    - new additions 8
    - TSB\_Hash field 286, 305
  - I/D Translation Storage Buffer Register 284, 303
  - indexing support 251
  - miss handler 253
  - organization 252
  - pointer logic 256
  - pointer logic hardware 256
  - Pointer register 287, 306
  - register, computing 64-Kbyte pointer 251
  - required conditions 258
  - SB\_Size field 255



- split 255
- Tag Target Register 255
- Tag Target register 283, 303
- TSO memory model 209, 210, 211, 213, 220
- TSTATE register
  - initializing 31, 437
  - PEF field 102
  - state after reset 440
- TSUBcc instruction 147, 593
- TSUBccTV instruction 148, 348
- TT register 439
- TTE
  - CP (cacheability) field 211, 217
  - CP field 250
  - CV (cacheability) field 211, 217
  - CV field 250
  - Diag field 249
  - E field 210, 211, 218, 220, 250
  - entry, locking in TSB 250
  - format 8
  - G field 248, 251
  - L field 250
  - NFO field 218, 249
  - P field 251
  - PA field 249
  - Size field 249
  - Soft2 field 249
  - V field 248
  - VA\_tag field 248
  - W field 251
- TVC instruction 595
- TVS instruction 595
- typewriter font, in assembly language syntax 633

## U

- UART 211
- UDIV instruction 96, 608
- UDIVcc instruction 96, 608
- UDIVX instruction 537
- UltraSPARC-I 3, 521
- UltraSPARC-II 3, 521
- UMUL instruction 96, 616
- UMULcc instruction 96, 616
- unconditional branches 473, 475, 604, 607
- underflow accrued (*ufa*) bit of *aexc* field of FSR register
  - 131
- underflow current (*ufe*) bit of *cexc* field of FSR register

- 131
- underflow mask (*UFM*) bit of TEM field of FSR register
  - 128, 131, 424, 425
- unfinished\_FPop* exception 124
- unfinished\_FPop* exception 5, 428, 489, 491, 492
- unfinished\_FPop* floating-point trap type 126, 126, 132, 496
- UNIMP instruction (SPARC V8) 501
- unimplemented instructions 179
- unimplemented\_FPop* floating-point trap type 126, 128, 132, 485, 487, 489, 491, 492, 496, 529, 531
- unsigned integer data type 65
- upper registers dirty (DU) field of FPRS register 100
- user
  - mode 98
  - trap handler 126, 425

## V

- VA Data Watchpoint Register
  - DB\_VA field 136
- VA\_watchpoint* exception 136, 349
- VA\_WATCHPOINT register 441
- VER
  - register 440
- version register (VER) 120, 566
- VIPT cache
  - D-cache 6
- virtual address 210
  - data watchpoint 136
- virtual address 0 218
- Virtual Indexed, Physical Tagged Caches 225
- virtual-indexed
  - physical-tagged (VIPT) cache 225
- virtually indexed, physically tagged (VIPT) 393
- virtual-to-physical address translation 210
- VIS extensions
  - byte mask 4
  - byte shuffle 4
  - differences from UltraSPARC-I 4
  - edge variants 4
- VIS instruction execution 45
- Visual Instruction Set (VIS) 103

## W

- W pipeline stage 46

- watchdog\_reset* (WDR) 32, 320, 339, 342, 349, 437
- watchpoint
  - and RED\_state 435
- watchpoints
  - data registers 136
  - PA/VA watchpoint traps 260
  - trap 259
- WC\_miss 374
- WC\_scrubbed 374
- WC\_snoop\_cb 374
- WC\_wb\_wo\_read 374
- WIM register (SPARC V8) 602
- window changing 52
- window fill trap handler 157
- window overflow 84
- window spill trap handler 157
- window state (WSTATE) register
  - description 120
  - overview 118
  - reading WSTATE with RDPR instruction 566
  - spill exception 500
  - spill trap 573
  - writing WSTATE with WRPR instruction 598
- window underflow 84
- window, clean 572
- window\_fill* exception 120, 326, 571
- window\_spill* exception 120, 326
- word
  - addressing 183
  - alignment 141
  - data format 65
- Working Register File (WRF) 52
- WRASI instruction 600
- WRASR instruction 100, 347, 364, 365, 600
  - forcing bubbles after 52
  - grouping rule 52
  - WRDCR instruction 600
  - WRGSR instruction 600
  - WRPCR instruction 600
  - WRPIC instruction 600
  - WRSOFTINT instruction 600
  - WRSOFTINT\_CLR instruction 600
  - WRSOFTINT\_SET instruction 600
  - WRSTICK instruction 600
  - WRSTICK\_CMPR instruction 600
  - WRTICK\_CMP instruction 600
- WRCCR instruction 98, 600
- WRF (Working Register File) 52
- WRFPRS instruction 600

- WRGSR instruction 51
- WRIER instruction (SPARC V8) 602
- Write Cache 227
- write cache
  - characteristics 6
  - description 227, 228
  - diagnostic accesses 195
  - flushing 7, 228
  - miss counts 374
- write privileged register instruction 597
- write-through cache 393
- WRPIC instruction 365
- WRPR instruction 106, 113, 118, 318, 347, 597, 597
  - forcing bubbles after 52
  - grouping rule 52
  - to PSTATE and BST 463
- WRPSR instruction (SPARC V8) 602
- WRTBR instruction (SPARC V8) 602
- WRWIM instruction (SPARC V8) 602
- WRY instruction 96, 600
- WSTATE register 440

## X

- x* field of instructions 179
- xcc* field of CCR register 98, 454, 476, 518, 534, 592, 593, 610, 611, 617, 618
- XNOR instruction 517
- XNORcc instruction 517
- XOR instruction 517
- XORcc instruction 517

## Y

- Y register 96, 96, 439, 609, 616, 618, 631

## Z

- zero (Z) bit of condition fields of CCR 97
- zero virtual address 218

# Revision History

---

| <b>Date</b>   | <b>Comments</b>   | <b>Version</b> |
|---------------|---|----------------|
| May 2002      | Initial release   | 1.0            |
| February 2003 | Additional chapters included  | 2.0            |
| April 2003    | Font modifications  | 2.1            |
| May 2003      | Modification to IEEE 754 and Performance Instrumentation chapters                             | 2.2            |
| December 2003 | Additional clarifications on pages 11-290, 11-291, and 18-441. Revised figures 6-12 and 6-29. | 2.2            |
| January 2004  | Updated table 6-9 and added note on page 6-102.   | 2.2.1          |

