

6 Instructions

Instructions are accessed by the processor from memory and are executed, annulled, or trapped. Instructions are encoded in four major formats and partitioned into eleven general categories.

6.1 Instruction Execution

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

An instruction may generate an exception if it encounters some condition that makes it impossible to complete normal execution. Such an exception may in turn generate a precise trap. Other events may also cause traps: an exception caused by a previous instruction (a deferred trap), an interrupt or asynchronous error (a disrupting trap), or a reset request (a reset trap). If a trap occurs, control is vectored into a trap table. See Chapter 7, “Traps,” for a detailed description of exception and trap processing.

If a trap does not occur and the instruction is not a control transfer, the next program counter (nPC) is copied into the PC and the nPC is incremented by 4 (ignoring overflow, if any). If the instruction is a control-transfer instruction, the next program counter (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.


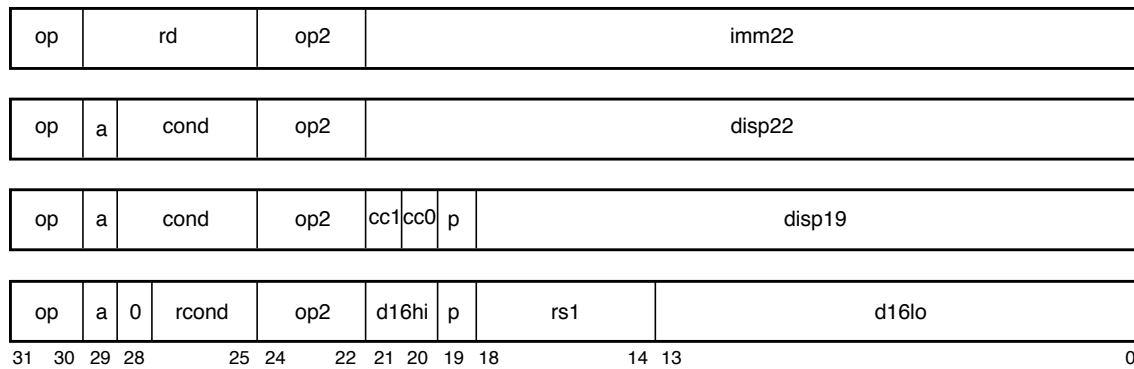
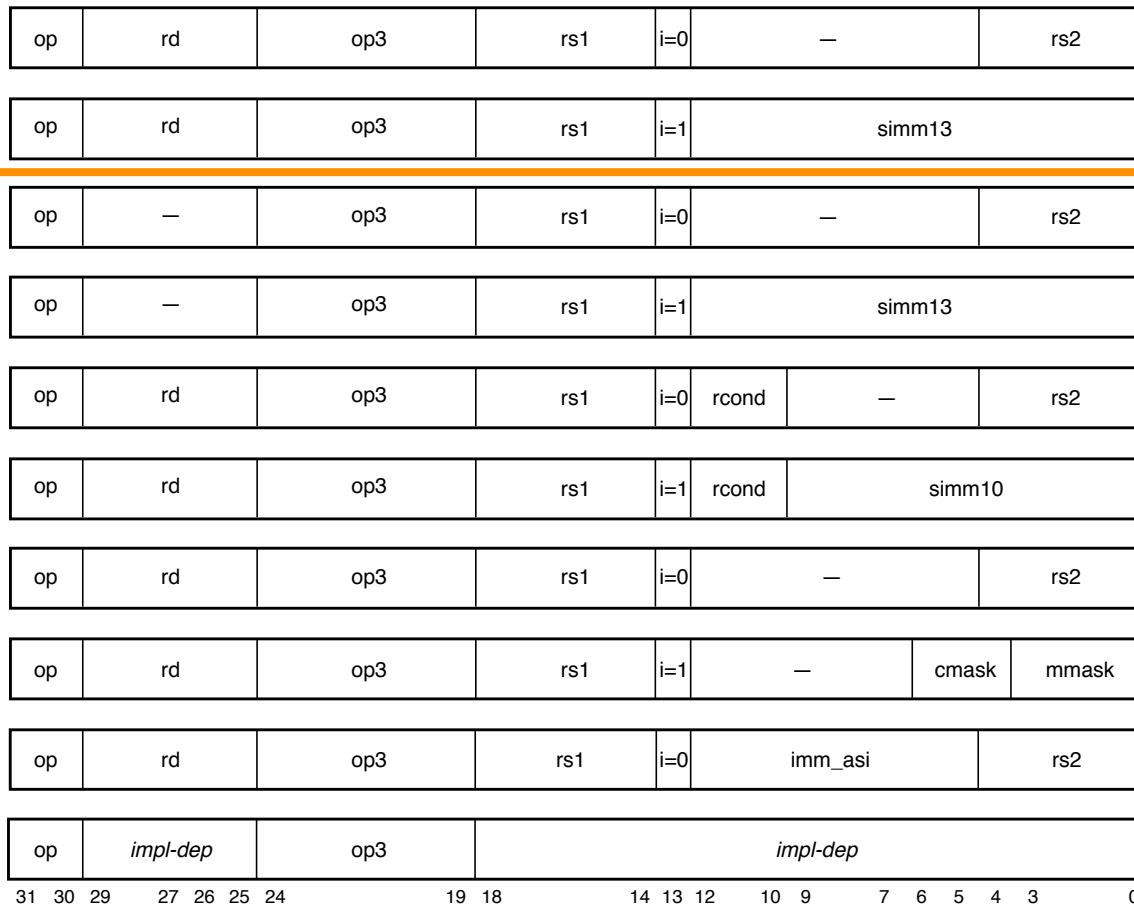
For each instruction access and each normal data access, the IU appends an 8-bit address space identifier, or ASI, to the 64-bit memory address. Load/store alternate instructions (see 6.3.1.3, “Address Space Identifiers (ASIs),”) can provide an arbitrary ASI with their data addresses, or use the ASI value currently contained in the ASI register.

Implementation Note:

The time required to execute an instruction is implementation-dependent, as is the degree of execution concurrency. In the absence of traps, an implementation should cause the same program-visible register and memory state changes as if a program had executed according to the sequential model implied in this document. See Chapter 7, “Traps,” for a definition of architectural compliance in the presence of traps.

6.2 Instruction Formats

Instructions are encoded in four major 32-bit formats and several minor formats, as shown in figures 33 and 34.

Format 1 ($op = 1$): CALL**Format 2** ($op = 0$): SETHI & Branches (Bicc, BPcc, BPr, FBfcc, FBPfcc)**Format 3** ($op = 2$ or 3): Arithmetic, Logical, MOV_r, MEMBAR, Load, and Store**Figure 33—Summary of Instruction Formats: Formats 1, 2, and 3**

Format 3 (op = 2 or 3): *Continued*

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		000		cc1		cc0		op3		rs1		opf				rs2		
op		rd		op3		rs1		opf						rs2				
op		rd		op3		rs1		—										
op		fcn		op3		—												
op		rd		op3		—												
31	30	29		25	24		19	18		14	13	12	11		6	5	4	0

Format 4 (op = 2): MOVcc, FMOVr, FMOVcc, and Tcc

op		rd		op3		rs1		i=0		cc1		cc0		—				rs2					
op		rd		op3		rs1		i=1		cc1		cc0		simm11									
op		rd		op3		cc2		cond		i=0		cc1		cc0		—				rs2			
op		rd		op3		cc2		cond		i=1		cc1		cc0		simm11							
op		rd		op3		rs1		i=1		cc1		cc0		—				sw_trap#					
op		rd		op3		rs1		0		rcond		opf_low				rs2							
op		rd		op3		0		cond		opf_cc		opf_low				rs2							
31	30	29	25		24	19		18	17	14		13	12	11	10	9	7		6	5	4	0	

E Opcode Maps

E.1 Overview

This appendix contains the SPARC-V9 instruction opcode maps.

Opcodes marked with a dash ‘—’ are reserved; an attempt to execute a reserved opcode shall cause a trap, unless it is an implementation-specific extension to the instruction set. See 6.3.11, “Reserved Opcodes and Instruction Fields,” for more information.

In this appendix and in Appendix A, “Instruction Definitions,” certain opcodes are marked with mnemonic superscripts. These superscripts and their meanings are defined in table 21 on page 133. For deprecated opcodes, see the appropriate instruction pages in Appendix A, “Instruction Definitions,” for preferred substitute instructions.

E.2 Tables

Table 30—*op*[1:0]

op [1:0]			
0	1	2	3
Branches & SETHI <i>See table 31</i>	CALL	Arithmetic & Misc. <i>See table 32</i>	Loads/Stores <i>See table 33</i>

Table 31—*op2*[2:0] (*op* = 0)

op2 [2:0]							
0	1	2	3	4	5	6	7
ILLTRAP	BPcc <i>See table 36</i>	Bicc ^D <i>See table 36</i>	BPr <i>See table 37</i>	SETHI NOP [†]	FBPfcc <i>See table 36</i>	FBfcc ^D <i>See table 36</i>	—

[†]*rd* = 0, *imm22* = 0

Table 32—*op3[5:0] (op = 2)*

		op3 [5:4]			
		0	1	2	3
op3 [3:0]	0	ADD	ADDcc	TADDcc	WRY ^D (<i>rd</i> = 0) — (<i>rd</i> = 1) WRCCR (<i>rd</i> = 2) WRASI (<i>rd</i> = 3) WRASR ^{PASR} (<i>see</i> A.63) WRFPRS (<i>rd</i> = 6) SIR (<i>rd</i> = 15, <i>rsI</i> = 0, <i>i</i> = 1)
	1	AND	ANDcc	TSUBcc	SAVED ^P (<i>fcn</i> = 0), RESTORED ^P (<i>fcn</i> = 1)
	2	OR	ORcc	TADDccTV ^D	WRPR ^P
	3	XOR	XORcc	TSUBccTV ^D	—
	4	SUB	SUBcc	MULScc ^D	FPop1 <i>See table 34</i>
	5	ANDN	ANDNcc	SLL (<i>x</i> = 0), SLLX (<i>x</i> = 1)	FPop2 <i>See table 35</i>
	6	ORN	ORNcc	SRL (<i>x</i> = 0), SRLX (<i>x</i> = 1)	IMPDEP1
	7	XNOR	XNORcc	SRA (<i>x</i> = 0), SRAX (<i>x</i> = 1)	IMPDEP2
	8	ADDC	ADDCcc	RDY ^D (<i>rsI</i> = 0) — (<i>rsI</i> = 1) RDCCR (<i>rsI</i> = 2) RDASI (<i>rsI</i> = 3) RDTICK ^{P_{NPT}} (<i>rsI</i> = 4) RDPC (<i>rsI</i> = 5) RDFPRS (<i>rsI</i> = 6) RDASR ^{PASR} (<i>see</i> A.44) MEMBAR (<i>rsI</i> = 15, <i>rd</i> = 0, <i>i</i> = 1) STBAR ^D (<i>rsI</i> = 15, <i>rd</i> = 0, <i>i</i> = 0)	JMPL
	9	MULX	—	—	RETURN
	A	UMUL ^D	UMULcc ^D	RDPR ^P	Tcc <i>See table 36</i>
	B	SMUL ^D	SMULcc ^D	FLUSHW	FLUSH
	C	SUBC	SUBCcc	MOVcc	SAVE
	D	UDIVX	—	SDIVX	RESTORE
	E	UDIV ^D	UDIVcc ^D	POPC (<i>rsI</i> = 0) — (<i>rsI</i> > 0)	DONE ^P (<i>fcn</i> = 0) RETRY ^P (<i>fcn</i> = 1)
	F	SDIV ^D	SDIVcc ^D	MOV _r <i>See table 37</i>	—

Table 33—*op3*[5:0] (*op* = 3)

		<i>op3</i> [5:4]			
		0	1	2	3
<i>op3</i> [3:0]	0	LDUW	LDUWA ^{PASI}	LDF	LDFA ^{PASI}
	1	LDUB	LDUBA ^{PASI}	LDFSR ^D , LDXFSR	—
	2	LDUH	LDUHA ^{PASI}	LDQF	LDQFA ^{PASI}
	3	LDD ^D	LDDA ^{D, PASI}	LDDF	LDDFA ^{PASI}
	4	STW	STWA ^{PASI}	STF	STFA ^{PASI}
	5	STB	STBA ^{PASI}	STFSR ^D , STXFSR	—
	6	STH	STHA ^{PASI}	STQF	STQFA ^{PASI}
	7	STD ^D	STDA ^{PASI}	STDF	STDFA ^{PASI}
	8	LDSW	LDSWA ^{PASI}	—	—
	9	LDSB	LDSBA ^{PASI}	—	—
	A	LDSH	LDSHA ^{PASI}	—	—
	B	LDX	LDXA ^{PASI}	—	—
	C	—	—	—	CASA ^{PASI}
	D	LDSTUB	LDSTUBA ^{PASI}	PREFETCH	PREFETCHA ^{PASI}
	E	STX	STXA ^{PASI}	—	CASXA ^{PASI}
	F	SWAP ^D	SWAPA ^{D, PASI}	—	—

Table 34—*opf*[8:0] (*op* = 2, *op3* = 34₁₆ = FPop1)

[illegible]

Table 35—*opf*[8:0] (*op* = 2, *op3* = 35₁₆ = FPop2)

opf [8:4]	opf[3:0]								8..F
	0	1	2	3	4	5	6	7	
00	—	FMOV _s (fcc0)	FMOV _d (fcc0)	FMOV _q (fcc0)	—	†	†	†	—
01	—	—	—	—	—	—	—	—	—
02	—	—	—	—	—	FMOVR _s Z	FMOVR _d Z	FMOVR _q Z	—
03	—	—	—	—	—	—	—	—	—
04	—	FMOV _s (fcc1)	FMOV _d (fcc1)	FMOV _q (fcc1)	—	FMOVR _s LEZ	FMOVR _d LEZ	FMOVR _q LEZ	—
05	—	FCMP _s	FCMP _d	FCMP _q	—	FCMP _E _s	FCMP _E _d	FCMP _E _q	—
06	—	—	—	—	—	FMOVR _s LZ	FMOVR _d LZ	FMOVR _q LZ	—
07	—	—	—	—	—	—	—	—	—
08	—	FMOV _s (fcc2)	FMOV _d (fcc2)	FMOV _q (fcc2)	—	†	†	†	—
09	—	—	—	—	—	—	—	—	—
0A	—	—	—	—	—	FMOVR _s NZ	FMOVR _d NZ	FMOVR _q NZ	—
0B	—	—	—	—	—	—	—	—	—
0C	—	FMOV _s (fcc3)	FMOV _d (fcc3)	FMOV _q (fcc3)	—	FMOVR _s GZ	FMOVR _d GZ	FMOVR _q GZ	—
0D	—	—	—	—	—	—	—	—	—
0E	—	—	—	—	—	FMOVR _s GEZ	FMOVR _d GEZ	FMOVR _q GEZ	—
0F	—	—	—	—	—	—	—	—	—
10	—	FMOV _s (icc)	FMOV _d (icc)	FMOV _q (icc)	—	—	—	—	—
11..17	—	—	—	—	—	—	—	—	—
18	—	FMOV _s (xcc)	FMOV _d (xcc)	FMOV _q (xcc)	—	—	—	—	—
19..1F	—	—	—	—	—	—	—	—	—

† Undefined variation of FMOVR

Table 36—*cond*[3:0]

		BPcc	Bicc^D	FBPfcc	FBfcc^D	Tcc
		op = 0 op2 = 1	op = 0 op2 = 2	op = 0 op2 = 5	op = 0 op2 = 6	op = 2 op3 = 3A₁₆
cond [3:0]	0	BPN	BN ^D	FBPN	FBN ^D	TN
	1	BPE	BE ^D	FBPNE	FBNE ^D	TE
	2	BPLE	BLE ^D	FBPLG	FBLG ^D	TLE
	3	BPL	BL ^D	FBPUL	FBUL ^D	TL
	4	BPLEU	BLEU ^D	FBPL	FBL ^D	TLEU
	5	BPCS	BCS ^D	FBPUG	FBUG ^D	TCS
	6	BPNEG	BNEG ^D	FBPG	FBG ^D	TNEG
	7	BPVS	BVS ^D	FBPU	FBU ^D	TVS
	8	BPA	BA ^D	FBPA	FBA ^D	TA
	9	BPNE	BNE ^D	FBPE	FBE ^D	TNE
	A	BPG	BG ^D	FBPUE	FBUE ^D	TG
	B	BPGE	BGE ^D	FBPGE	FBGE ^D	TGE
	C	BPGU	BGU ^D	FBPUGE	FBUGE ^D	TGU
	D	BPCC	BCC ^D	FBPLE	FBLE ^D	TCC
	E	BPPOS	BPOS ^D	FBPULE	FBULE ^D	TPOS
	F	BPVC	BVC ^D	FBPO	FBO ^D	TVC

Table 37—Encoding of *rcond*[2:0] Instruction Field

		BPr	MOVr	FMOVr
		op = 0 op2 = 3	op = 2 op3 = 2F₁₆	op = 2 op3 = 35₁₆
rcond [2:0]	0	—	—	—
	1	BRZ	MOVRZ	FMOVZ
	2	BRLEZ	MOVRLEZ	FMOVLEZ
	3	BRLZ	MOVRLZ	FMOVLZ
	4	—	—	—
	5	BRNZ	MOVRNZ	FMOVNZ
	6	BRGZ	MOVRGZ	FMOVGZ
	7	BRGEZ	MOVERGEZ	FMOVGEZ

Table 38—*cc/opf_cc* Fields (MOVcc and FMOVcc)

opf_cc			Condition code selected
cc2	cc1	cc0	
0	0	0	fcc0
0	0	1	fcc1
0	1	0	fcc2
0	1	1	fcc3
1	0	0	icc
1	0	1	—
1	1	0	xcc
1	1	1	—

Table 39—*cc* Fields (FBPfcc, FCMP and FCMPE)

cc1	cc0	Condition code selected
0	0	fcc0
0	1	fcc1
1	0	fcc2
1	1	fcc3

Table 40—*cc* Fields (BPcc and Tcc)

cc1	cc0	Condition code selected
0	0	icc
0	1	—
1	0	xcc
1	1	—

A.27 Load Integer

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. It is recommended that the LDX instruction be used in its place.

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
LDD ^D	00 0011	Load Doubleword

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	

Suggested Assembly Language Syntax	
ldsb	[<i>address</i>], <i>reg_{rd}</i>
ldsh	[<i>address</i>], <i>reg_{rd}</i>
ldsw	[<i>address</i>], <i>reg_{rd}</i>
ldub	[<i>address</i>], <i>reg_{rd}</i>
lduh	[<i>address</i>], <i>reg_{rd}</i>
lduw	[<i>address</i>], <i>reg_{rd}</i> (synonym: ld)
ldx	[<i>address</i>], <i>reg_{rd}</i>
ldd	[<i>address</i>], <i>reg_{rd}</i>

Description:

The load integer instructions copy a byte, a halfword, a word, an extended word, or a doubleword from memory. All except LDD 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 doubleword integer instructions (LDD) copy 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. Note that 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.

IMPL. DEP. #107(1): It is implementation-dependent whether LDD is implemented in hardware. If not, an attempt to execute it will cause an *unimplemented_ldd* exception.

Load integer instructions access the primary address space ($ASI = 80_{16}$). The effective address 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 and load doubleword) instruction operates atomically.

LDUH and LDSH cause a *mem_address_not_aligned* exception if the address is not half-word-aligned. LDW and LDSW cause a *mem_address_not_aligned* exception if the effective address is not word-aligned. LDX and LDD cause a *mem_address_not_aligned* exception if the address is not doubleword-aligned.

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. In some systems it may trap to emulation code. It is suggested that programmers and compilers avoid using these instructions.

If LDD is emulated in software, an LDX instruction should be used for the memory access in order to preserve atomicity.

Compatibility Note:

The SPARC-V8 LD instruction has been renamed LDW in SPARC-V9. The LDSW instruction is new in SPARC-V9.

Exceptions:

- async_data_error*
- unimplemented_LDD* (LDD only (impl. dep. #107))
- illegal_instruction* (LDD with odd rd)
- mem_address_not_aligned* (all except LDSB, LDUB)
- data_access_exception*
- data_access_protection*
- data_access_MMU_miss*
- data_access_error*

A.54 Store Integer

The STD 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. It is recommended that the STX instruction be used in its place.

Opcode	op3	Operation
STB	00 0101	Store Byte
STH	00 0110	Store Halfword
STW	00 0100	Store Word
STX	00 1110	Store Extended Word
STD ^D	00 0111	Store Doubleword

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	

Suggested Assembly Language Syntax		
stb	<i>reg_{rd}</i> , [<i>address</i>]	(synonyms: stub, stsb)
sth	<i>reg_{rd}</i> , [<i>address</i>]	(synonyms: stuh, stsh)
stw	<i>reg_{rd}</i> , [<i>address</i>]	(synonyms: st, stuw, stsw)
stx	<i>reg_{rd}</i> , [<i>address</i>]	
std	<i>reg_{rd}</i> , [<i>address</i>]	

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 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.

IMPL. DEP. #108(1): IT is implementation-dependent whether STD is implemented in hardware. If not, an attempt to execute it will cause an *unimplemented_STD* exception.

The effective address for these instructions is “ $r[rs1] + r[rs2]$ ” if $i = 0$, or “ $r[rs1] + \text{sign_ext}(\text{simm13})$ ” if $i = 1$.

A successful store (notably, store extended and store doubleword) 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 and STD causes a *mem_address_not_aligned* exception if the effective address is not doubleword-aligned.

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. In some SPARC-V9 systems it may cause a trap to emulation code; therefore, STD should be avoided.

If STD is emulated in software, STX should be used in order to preserve atomicity.

Compatibility Note:

The SPARC-V8 ST instruction has been renamed STW in SPARC-V9.

Exceptions:

- async_data_error*
- unimplemented_STD* (STD only) (impl. dep. #108)
- illegal_instruction* (STD with odd *rd*)
- mem_address_not_aligned* (all except STB)
- data_access_exception*
- data_access_error*
- data_access_protection*
- data_access_MMU_miss*