

SPARCclassic Engine OEM Technical Manual

PART TWO:

Appendices K & L

microSPARC Reference Material



Sun Microsystems, Inc.
2550 Garcia Avenue
Mountain View, CA 94043
U.S.A.

Part No. 801-3137-10
Revision A, April, 1993

SPARCclassic Engine OEM Technical Manual

PART TWO:

Appendices K & L

microSPARC Reference Material



Sun Microsystems, Inc.
2550 Garcia Avenue
Mountain View, CA 94043
U.S.A.

Part No.: 801-3137-10
Revision A, April, 1993

microSPARC Reference Manual



Use the Texas Instruments microSPARC Reference Manual to define the operating parameters for the microprocessor, and for any firmware/software development that you require.

The TI documentation included here is accurate as of the date of release of the SPARCengine EC OEM Technical Manual. Please call Texas Instruments to ensure that you have the most current documentation. Please use caution in developing plans on this information until you confirm it is the latest information available.



TMS390S10

***micro*SPARC™**

Reference Guide

***Highly Integrated
SPARC
for Low-Cost
Desktop
Solutions***



Table of Contents

Overview	3
Integer Unit	5
Overview	5
Instruction Pipeline.....	7
Memory Operations.....	8
ALU/Shift Operations.....	10
Integer Multiply.....	10
Integer Divide.....	11
CTI's.....	12
Instruction Cache Interface.....	13
Data Cache Interface	13
Interlocks	14
Traps and Interrupts.....	14
Floating Point Interface	16
Special Features.....	17
Divergence from SPARC Version 8.....	17
Floating Point Unit.....	19
Overview	19
Deltas from SPARC Version 8.....	21
Implementation Specific Features	22
Software Considerations.....	23
FPU Instruction Timings	25
Memory Management Unit.....	27
Overview	27
Translation Lookaside Buffer	29
CPU TLB Lookup	35
CPU TLB Flush and Probe Operations	35
Processor MMU Registers.....	37
IO MMU Registers	46
IO MMU Bypass Mode.....	54
Physical Address Register	54
TLB Table Walk.....	55
Instruction Translation Buffer Register.....	57
Arbitration	58
Translation Modes	59
Page Mode Detection.....	59
Errors and Exceptions.....	59
Diagnostic Features	60

Data Cache	67
Overview	67
Data Cache Data Array	69
Data Cache Tags	69
Write Buffers	70
Data Cache Fill	70
Internal Memory Bus Interface	71
IU Data Bus Interface	71
RENU Register	71
Data Cache Flushing	71
Cacheability of Memory Accesses	71
Diagnostic Strategy	72
Instruction Cache	73
Overview	73
Instruction Cache Data Array	75
Instruction Cache Tags	75
Instruction Cache Fill	76
Internal Memory Bus Interface	77
IU Instruction Bus Interface	77
Instruction Cache Flushing	77
Cacheability of Memory Accesses	77
Diagnostic Strategy	77
Memory Interface	79
Overview	79
Memory Subsystem	79
Memory Control Block (MCB)	81
Data aligner and Parity Check/generate logic (DPC)	93
RAM Refresh Control (RFR)	96
SBus Controller	99
Overview	99
CPU Interface	104
Address Steering	107
SBus Arbiter	107
Main Control	110
Data Transfer	113
Slave Control Cycle	114
Slave Target Control	116
Data Path	117
Data Control	119
Error Handling	119
Diagnostic Testing	119

Additional Work	120
Reset, Clock Control, JTAG	121
Reset Controller	121
Reset Controller State Machine Operation	124
Clock Controller	125
Clock Signals	127
Stopping Clocks	127
Starting Clocks	127
Single-Step	127
Stop Clocks on Internal Event	128
External Cycle Counter	128
Counting Clocks	129
Issuing N Clocks	129
Count Clocks After Internal Event	133
Stop Clocks After N Internal Events	139
CCR Bits	140
JTAG	141
Board Level Architecture	141
TAP	141
Data Registers	142
JTAG Instructions	143
JTAG Interface to MISC	144
JTAG Operation	146
Error Handling	151
ASI Map	153
Overview	153
References	159

List of Figures

microSPARC Block Diagram	3
microSPARC IU Block Diagram	6
Meiko FPU Block Diagram.....	20
Untrapped FP Result in Same Format as Operands	21
Untrapped FP Result in Different Format	22
FPU Operation Modes	23
MMU Address and Data Path Block Diagram.....	28
TLB Replacement.....	29
TLB Entry.....	29
Page Table Entry in Page Table	31
Page Table Entry in TLB	32
Page Table Pointer in Page Table.....	32
Page Table Pointer in TLB	33
IO Page Table Entry in Page Table	34
IO Page Table Entry in TLB.....	34
CPU TLB Flush or Probe Address Format.....	36
Processor Control Register	38
Context Table Pointer Register.....	41
Context Register	41
Synchronous Fault Status Register	42
Synchronous Fault Address Register.....	46
TLB Replacement Control Register	46
IO Control Register	47
IO MMU Base Address Register.....	48
IOPTe Address Based Flush Format.....	49
Asynchronous Fault Status Register	50
Asynchronous Fault Address Register	51
SBUS Slot Configuration Register	51
Memory Fault Status Register	52
Memory Fault Address Register.....	53
MID Register	54
Physical Address Register	55
CPU Address Translation Using Table Walk.....	56
ITBR Page Table Entry	57
CPU Diagnostic TLB and ITLB Tag Access Format.....	60
Data Cache Block Diagram	68
Data Cache Tag Entry.....	69
Instruction Cache Block Diagram	74
Instruction Cache Tag Entry.....	75
MCB block diagram.....	82
MMU I-fetch beginning in page-mode	86

MMU page-mode write after a read	87
Non-paged write cycle, shown following a read	88
Non-paged read cycle, shown following a read	89
Paged Byte/Halfword (8/16 bit) write cycle.....	90
Non-paged Byte/Halfword (8/16 bit) write cycle.....	91
Datapath and Parity Control (DPC) block diagram.....	95
RAM Refresh Control block diagram.	96
Data Get and Data Put.....	101
SBus Controller Block Diagram.....	103
CPU State Machine	106
Arbitor State Machine	109
Main State Machine.....	112
D_ctl State Machine	113
S_ctl State Machine	115
t_ctl State Machine	116
SBC Data Path.....	118
microSPARC Reset State Machine.....	123
Clock Controller State Machine	126
Single Step with sbus_1st_half = 1.	128
Single Step with sbus_1st_half = 0.	128
With stop_on_ext_event.....	130
With stop_even_on_ext_event	130
N=2, stopped with sbus_1st_half=1.	132
N=7, stopped with sbus_1st_half=0.	132
N=7, stopped with sbus_1st_half=1.	133
Event in First half of bus cycle, N=8.....	135
Event in First half of bus cycle, N=8.....	136
Event in First half of bus cycle, N=9.....	137
Event in Second half of bus cycle, N=9.	138
Event in first half of bus cycle, N=2. Latency=6 cycles.	139
Event in second half of bus cycle, N=2. Latency=5 cycles.....	140
JTAG ID Reg Contents.....	142
JTAG LOGIC BLOCK DIAGRAM.....	148
microSPARC JTAG Data & Instruction Registers.....	149
TLB Flush or Probe Address Format	154
Instruction Cache Tag Entry.....	156
Data Cache Tag Entry.....	157

List of Tables

Cycles per Instruction	7
FSR Summary	24
FPU Instruction Timings	25
Virtual Tag Match Criteria.....	30
Page Table Access Permissions	31
Page Table Entry Types	32
Sizes of Page Tables	33
Page Table Entry Types	33
Virtual Tag Match Criteria	35
TLB Entry Flushing	36
CPU TLB Entry Probing	37
Address Map for MMU Registers	38
Memory Refresher Control Definition	39
Parity Control Definition	40
SFSR Level Field	43
SFSR Access Type Field	43
SFSR Fault Type Field	43
Setting of SFSR Fault Type Code	44
Overwrite Operations	45
Priority of Fault Types on Single Access	45
SBUS and IO MMU Control Space	47
IO MMU Page Table Address Generation	48
Memory Request Type	53
TLB Reference Priority	59
Translation Modes	59
TLB Entry Address Mapping	61
Virtual Address Match Conditions	62
Virtual Address Field Enable Decode	63
Memory Request Type	64
Data Cache Fill Ordering	70
Address Map for Data Cache Registers	72
Instruction Cache Fill Ordering	76
Memory operations performed by MCB.....	84
Physical Address decode for System Memory	93
Parity Control Definition	94
Refresh Rate Control bits	96
Clock Control and Scan	131
JTAG INSTRUCTIONS	143
Error Summary.....	151
ASI's Supported by microSPARC.....	153
TLB Entry Flushing	154

CPU TLB Entry Probing155
Address Map for MMU Registers155
Address Map for Data Cache Registers158

Preface

This guide contains application information for the highly integrated SPARC processor, named the TMS390S10. Throughout this guide the term microSPARC is used to describe the TMS390S10 chip.

This User's Guide should be used in conjunction with the TMS390S10 data sheet. Where conflicts between these documents exist, particularly in reference to exact timings and frequency information, the TMS390S10 data sheet has precedence.

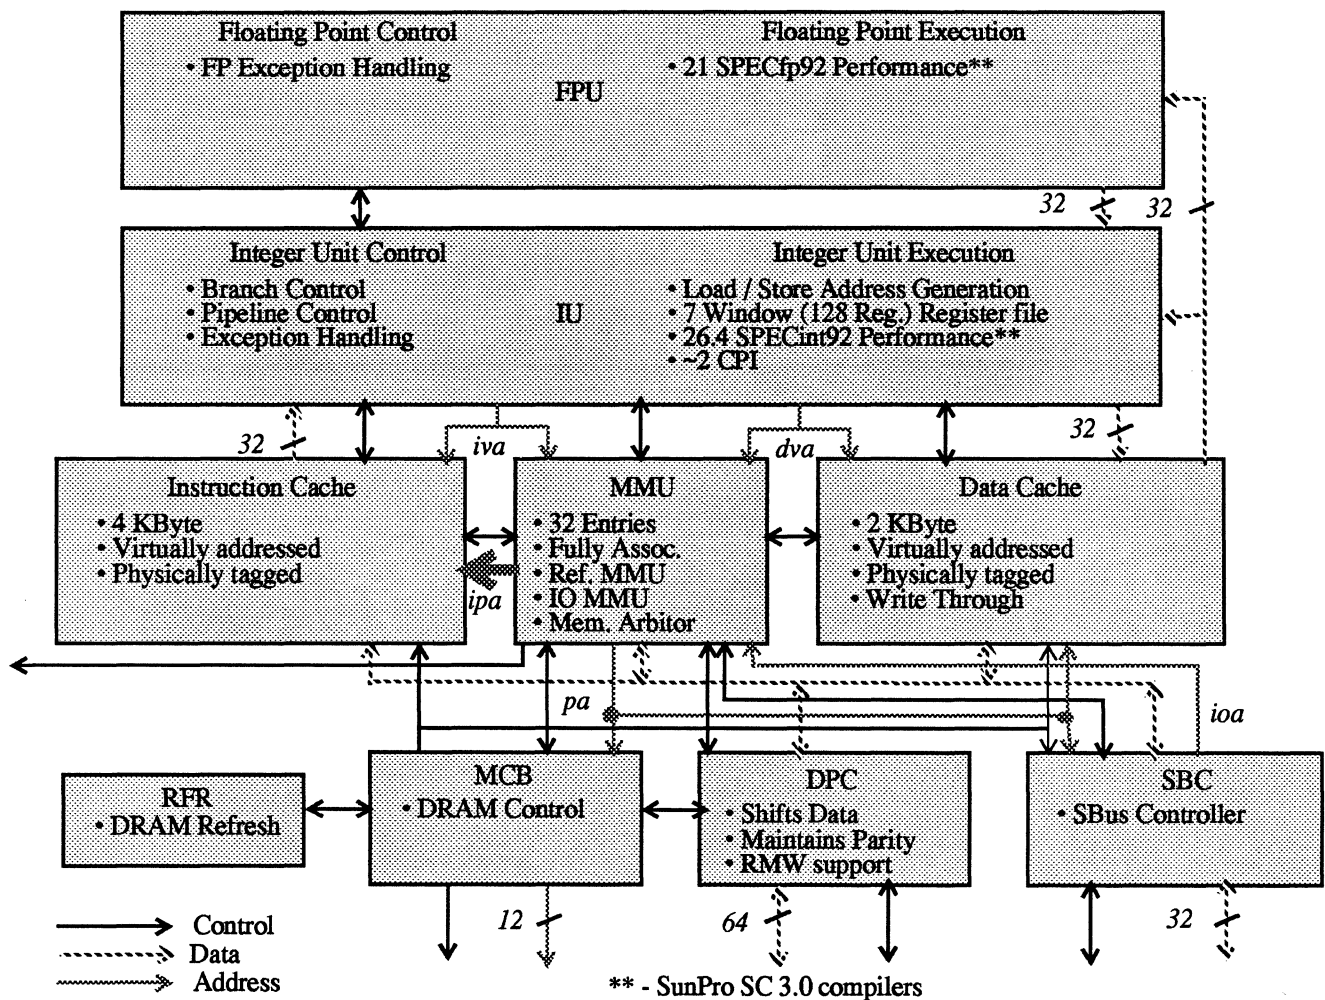
1.0 Overview

The microSPARC CPU is a highly integrated, low-cost implementation of the SPARC RISC architecture. High performance is achieved by the high level of integration including on chip instruction and data caches and the close coupling of the CPU with main memory. A full custom implementation allows for a target frequency of 50 MHz providing sustained performance of 26.4 SPECint92 with SunPro SC 3.0 compilers. The design is highly testable with the use of the full JTAG scan support. The microSPARC chip will support up to 128MB of DRAM and 4 SBus slots.

Integrated within microSPARC are a SPARC V8 Integer Unit core, a SPARC Reference Memory Management Unit, a Floating Point Unit, Instruction and Data Caches, DRAM controller, and an SBus Controller.

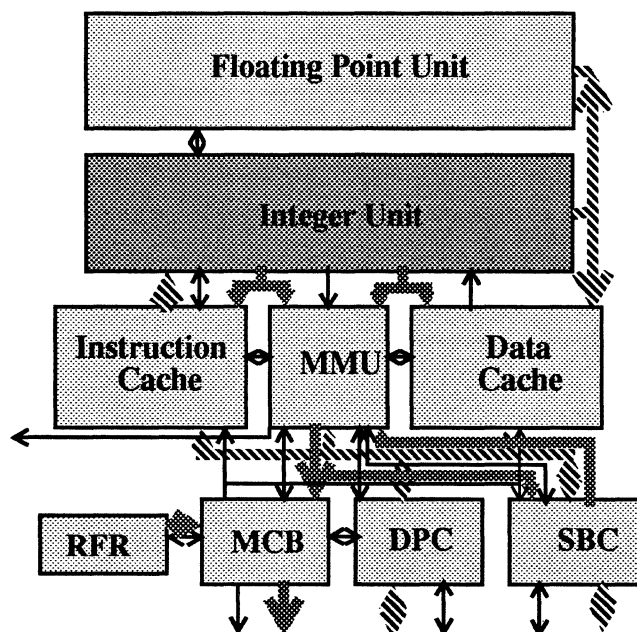
A simple block diagram follows.

Figure 1.0 - microSPARC Block Diagram



2.0 Integer Unit

The microSPARC integer unit is a SPARC integer unit as defined in the SPARC Architecture Manual (Version 8). The IU design goal is to maximize performance, given a constrained die size, using a predefined software architecture. The emphasis is on software compatibility, since the greatest cost impact would be on any software (i.e. kernel, compilers) that would need rewriting.

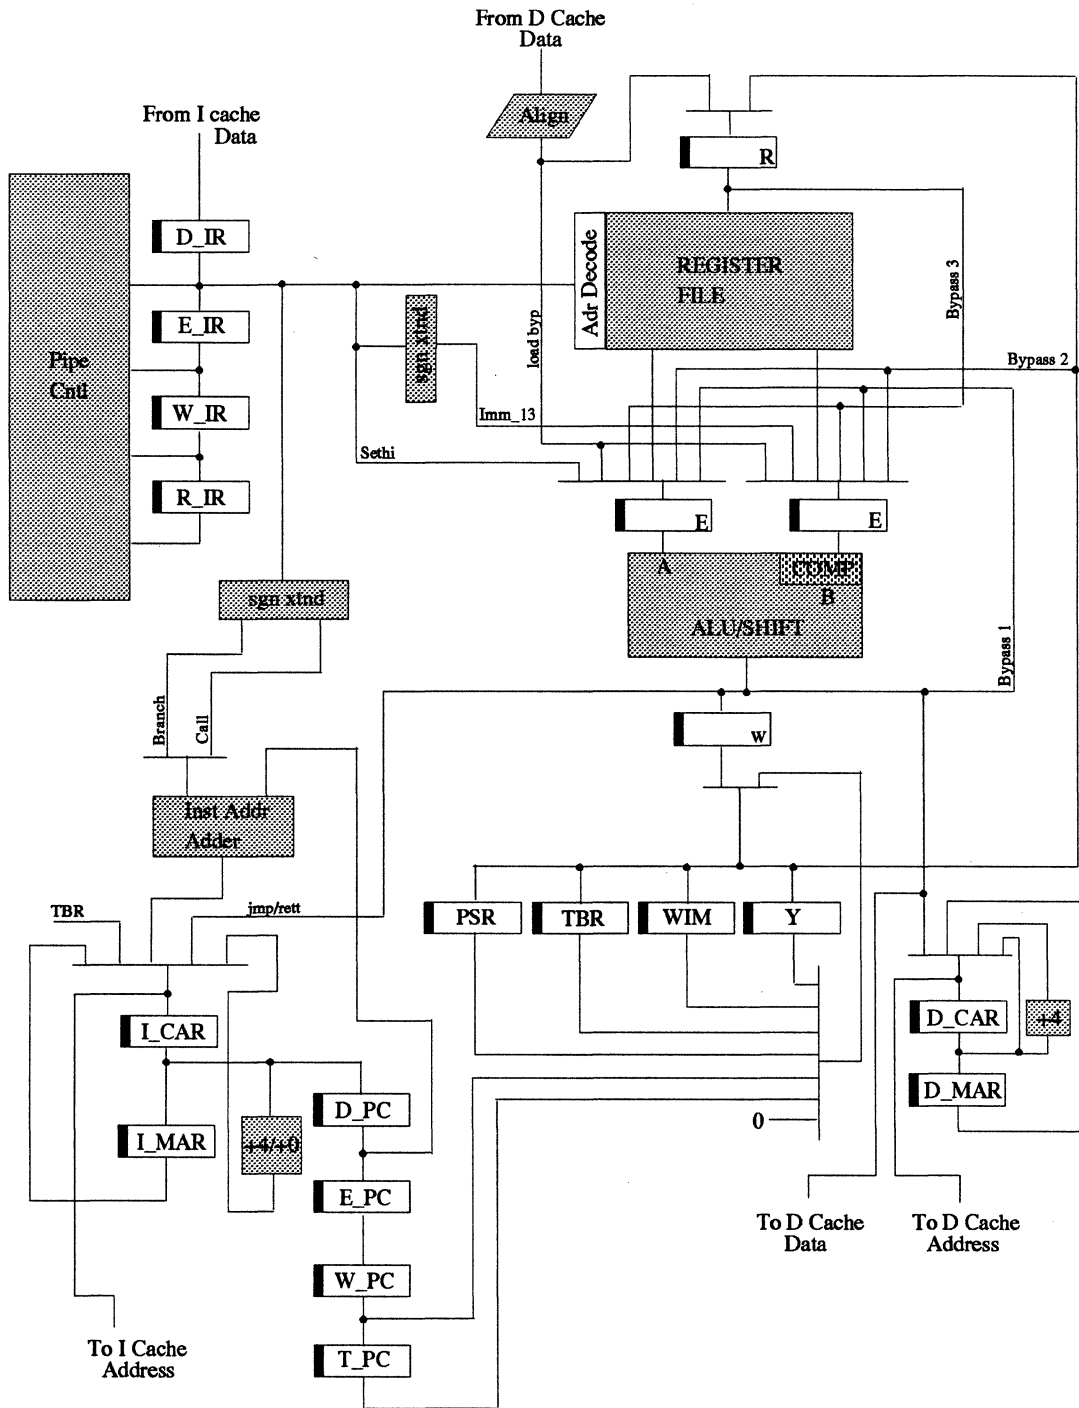


2.0.1 Overview

The microSPARC integer unit is a CMOS implementation of the SPARC 32-bit (Version 8) RISC architecture. Some important features of this design are:

- Single issue, 5 stage pipeline
- Harvard architecture
- Instruction and Data cache streaming support
- IMUL and IDIV implemented as integer operations
- 0 cycle branch penalty
- 120-register register file (7 register windows)

Figure 2.0 - microSPARC IU Block Diagram



2.0.2 Instruction Pipeline

The microSPARC IU uses a single instruction issue pipeline with 5 stages.

- F (Instruction Fetch): Instruction cache access occurs in this cycle based on the address generated in the previous cycle. The instruction is valid on the pins of the IU at the end of this cycle and are registered inside of the IU.
- D (Decode): The decode stage is used to decode the instruction and to read the necessary operands. Operands may come from the register file or from internal data bypasses. The register file has 2 independent read ports. For situations where the necessary operand is in the pipeline and has not yet been written to the register file, internal bypasses are supplied to prevent pipeline interlocks. In addition, addresses are computed for CALL and Branch in this cycle in the address adder.
- E (Execute): The execute stage is used to perform ALU, logical, and shift operations. For memory operations (e.g.: LD) and for JMPL/RETT the address is computed in this cycle.
- W (Write): This stage is used to access the data cache. For cache reads, the data will be valid by the end of this cycle, at which point it is aligned as appropriate. For cache writes, the data is presented to the data cache in this cycle.
- R (Result): This stage writes the result of any ALU, logical, shift, or cache read operation into the register file.

Table 2.0 - Cycles per Instruction

Instruction	Cycles	Words
Call	1	1
Single Loads	1	1
Jump/Rett	2	1
Double Loads	2	1
Single Stores	2	1
Double Stores	3	1
Taken Trap	3	1
Atomic Load/Store	2	1
SWAP	2	1
All Others	1	1

2.0.3 Memory Operations

2.0.3.1 Loads

All load operations take 1 cycle in the microSPARC IU except for LDD which takes 2. For LD, LDB, and LDH the pipeline does the following:

D - Register operands are read from the register file or are bypassed from instructions still in the pipe. An immediate operand is sign extended.

E - Address operands are added to yield the memory address. This address is passed to the cache in this cycle.

W - Address is registered in the cache and access is started. Data is expected at the end of this cycle. Any necessary alignment and sign extension is done in the IU prior to being registered.

R - Data is registered in the IU and is written into the register file.

In the event of a cache miss, the miss indication is given to the IU in the R cycle. It is flagged early enough to prevent writing bad data to the register file. The pipe is held and the miss address is resent to the cache to service the miss. The cache indicates when the miss data is available - the IU can then register it into the appropriate R cycle register and write it into the register file.

An LDD takes 2 cycles to complete because of the 32 bit datapaths. The pipeline does the following:

D - Register operands are read from the register file or are bypassed from instructions still in the pipe. An immediate operand is sign extended.

E - Address operands are added to yield the even memory address. This address is passed to the data cache in this cycle.

W (E2) - Even memory address is registered in the cache and access is started. This data is sent to the IU. At the same time, the odd address is generated by the IU and sent to the cache.

R (W2) - Even word is registered in the IU and written to the register file. The odd word address is registered in the cache and its access is started. This data gets sent to the IU.

R2 - Odd word is registered in the IU and written to the register file.

In the event of a cache miss, the miss indication is in the R cycle of the LDD (the same as the W cycle of the LDD's help cycle). The miss is indicated early enough to prevent writing bad data into the even register. The pipe is held and the even address is resent to the cache. When the cache sends the correct data, the R register is written with the correct data and the odd address is resent to get the odd word.

2.0.3.2 Stores

The microSPARC IU register file has only two independent read ports. As a result, store operations take 2 cycles, except STD which takes 3. For ST, STB, and STH the pipeline does the following:

D - Register operands are read from the register file or are bypassed from instructions still in the pipe. An immediate operand is sign extended.

E (D2) - The address operands are added to compute the memory address. This address will be registered within the IU to provide the data cache with the address in the correct cycle. At the same time, the store data is read from the register file or bypassed from instructions still in the pipe.

W (E2) - The store address is sent to the data cache.

R (W2) - The store data is sent to the data cache in this cycle along with the appropriate byte marks.

R2 - Store is complete.

For STD the pipeline does the following:

D - Register operands are read from the register file or are bypassed from instructions still in the pipe. An immediate operand is sign extended.

E (D2) - The address operands are added to compute the even memory address. This address will be registered within the IU to provide the data cache with the address in the correct cycle. At the same time, the even store data is read from the register file or bypassed from instructions still in the pipe.

W (E2/D3) - Even address is sent to the data cache. Odd word is read from register file.

R (W2/E3) - Even store data is sent to the data cache. Odd address is sent to the data cache.

R2 (W3) - Odd data is sent to the data cache.

R3 - STD complete.

2.0.3.3 Atomics

SWAP and LDSTUB each take two cycles to complete. The pipeline does the following on the SWAP instruction:

D - Register operands are read from the register file or are bypassed from instructions still in the pipe. An immediate operand is sign extended.

E (D2) - The address operands are added to compute the swap address. This address is sent to the data cache to start the cache read portion of the operation. The address is also registered inside of the IU to provide the data cache with the same address

for the store in the next cycle. The register to be swapped is read out in this cycle.

W (E2) - The data cache returns the memory location accessed. The store address is sent to the data cache.

R (W2) - The IU registers the read data and writes it to the register file. Also the store data is sent to the data cache.

R2 - SWAP complete.

The pipeline does the following on the LDSTUB instruction:

D - Register operands are read from the register file or are bypassed from instructions still in the pipe. An immediate operand is sign extended.

E (D2) - The address operands are added to compute the ldst address. This address is sent to the data cache to start the cache read portion of the operation. The address is also registered inside of the IU to provide the data cache with the same address for the store in the next cycle.

W (E2) - The data cache returns the memory location accessed and it is shifted appropriately inside the IU. The store address is sent to the data cache.

R (W2) - The IU registers the read data and writes it to the register file. Also 0xffffffff is sent to the data cache along with the appropriate byte marks to complete the store.

R2 - LDSTUB complete.

2.0.4 ALU/Shift Operations

Most ALU and shift operations take a single cycle to complete. The exceptions are Integer Multiply and Integer Divide. On Add, Subtract, Boolean, and Shift operations the pipeline does the following:

D - Read operands from register file or bypass from instructions still in the pipe.

E - Do appropriate operation in ALU or shifter. There is a selective inverter on the B input of the ALU to allow for subtracts and certain Boolean operation (e.g. ANDN).

W - Pipe result into R.

R - Write register file with result.

2.0.5 Integer Multiply

Integer multiply takes 19 cycles to complete. The algorithm implemented in the microSPARC IU is a modified Booth's (2-bit) multiply. The multiply process can be broken up into 4 distinct steps:

Initialization	1 cycle
Booth's iteration	16 cycles

Correction (ala Booth)	1 cycle
Writeback	1 cycle

The first cycle is used to set up the registers used in the multiply. The rs1 and rs2 registers initialize to the operands of the multiply. The W stage result register and the rs2 register are used as accumulators. At the completion of the multiply, the W register contains the most significant 32 bits of the result and the rs2 register contains the least significant 32 bits of the result. The W register contents are then written to the Y register and the rs2 contents to the destination register in the register file.

2.0.6 Integer Divide

Integer divide takes 39 cycles to complete. If an overflow is detected, the instruction completes in 6 cycles. The algorithm implemented in the microSPARC IU is non-restoring binary division (add and shift). The divide process can be broken into 5 distinct steps:

Divide by zero detection	1 cycle
Initialization/Ovf detection	3 cycles
Non-restoring division iteration	33 cycles
Correction (for non-restoring)	1 cycle
Writeback	1 cycle

Because the microSPARC IU does not allow traps to be taken by help instructions, the first step is to determine if we have a divide by 0 condition.

The high order bits of the dividend are in the Y register. The low order bits are in the rs1 operand. The divisor is in the rs2 operand. In the initialization step, the Y register is read out and put into the rs1 register in the datapath. The rs1 operand is passed through to the W register. The rs2 operand is passed to the rs2 register (surprise!). The W and rs1 registers are used as accumulators. At the completion of the divide, the W register contains the final quotient.

There are two overflow options for signed divide with a negative result defined in the SPARC Version 8 manual. The microSPARC IU implements:

$$\text{result} < -2^{31} \text{ with remainder} = 0.$$

If an overflow condition is detected, the divide terminates early with the appropriate result being written to the destination register.

If no overflow is detected, the non-restoring (add then shift) divide stage is started. A correction step is provided to correct the quotient (necessary for this algorithm). After the correction step, the quotient is written to the correct destination register.

2.0.7 CTI's

2.0.7.1 Branches

All branches take a single cycle to execute. There is no penalty for taken vs. untaken branches, even in the event that the instruction previous to the branch sets the condition codes.

In the Decode stage, the IU evaluates the condition codes and branch condition to determine taken or untaken. The IU outputs the correct instruction address for either the target or fall through paths in time to be registered by the instruction cache for the fetch occurring in the next cycle.

2.0.7.2 JMPL

JMPL is a two cycle instruction in the microSPARC IU. This is done somewhat uniquely in that there are no help cycles for the JMPL. Instead, there is an interlock that always occurs following the JMPL. This is done to force the IU to fetch the JMPL's delay instruction. In this way, the IU can evaluate whether an RETT is in the JMPL's delay slot and evaluate user/supervisor accesses correctly.

D - Read operands from register file or bypass from instructions still in the pipe. Sign extend immediate operands. The delay slot instruction is fetched in this cycle.

E - Compute target address and send this to the instruction cache.

W - Not much happens.

R - Write the PC of the JMPL instruction into the destination register.

2.0.7.3 RETT

RETT is a two cycle instruction in the microSPARC IU. Unlike JMPL, the RETT utilizes a help cycle. However, since it must follow an JMPL, the first cycle is always interlocked. This cycle allows the IU to determine that the RETT enters the pipe and can force the correct user/supervisor mode (contained in the PSR.PS bit) for subsequent instruction fetches.

D - Read operands from register file or bypass from instruction still in the pipe. Sign extend immediate operands.

E - Compute target address and send this to the instruction cache.

W - Not much happens.

R - Set PSR.ET to 1, move PSR.PS to PSR.S, and PSR.CWP++.

2.0.7.4 CALL

CALL is a single cycle instruction in the microSPARC IU.

D - Add PC and disp30 to form target address. Send this address to instruction cache. The delay slot fetch starts this cycle.

E - The CALL target is fetched.

W - Not much happens.

R - The PC of the CALL is written to r[15].

2.0.8 Instruction Cache Interface

In the event of an instruction cache miss, the IU will recirculate the missed address to the address bus to the instruction cache and hold the pipeline. Since the miss indication cannot be generated in time to prevent the missed instruction from moving from F to D, the missed instruction is physically in the Decode stage of the pipe.

The instruction cache is implemented so that the missed word of the cache line is returned first. This instruction word is strobed into the Decode stage. The IU is now free to stream instructions from the instruction cache as the cache is doing its line fill. This means that the IU is not held for the entire duration of the cache fill, but can use the instructions as soon as the instruction cache receives it. To do this, the IU is told when the instruction addressed by the IU is available to be strobed in. The IU can then selectively hold and release the pipe. One caveat is that the IU must correctly select the address to be sent to the instruction cache (determined by hold).

If one of the instructions encountered during the instruction streaming is a taken CTI whose target is outside of the cache line being filled, the IU can detect this condition (the instruction cache cannot) and hold the pipe.

2.0.9 Data Cache Interface

The data cache interface is roughly similar to the instruction cache interface. In the event of a data cache miss, the IU will recirculate the missed address to the data cache address bus and hold the pipeline. Since the data miss indication is not generated in time to prevent the instruction from moving from W to R, the instruction that caused the miss is in its R cycle. Any expected load data must then be directly strobed into the R stage and if the instruction in the E stage expects to get load data (via the load bypass), the load data must also be strobed into the correct E stage register(s).

The data cache is also implemented to return the missed word first. When the data cache indicates that the data is available, the data is passed through the load aligner (for any necessary alignment) and then strobed into the R cycle (and appropriate E cycle) register prior to being written to the register file.

The IU is then free to continue. To limit the complexity of the MMU, however, while the data cache is filling the line, no additional memory operations may be started until the line fill is complete. The exception to this is LDD, as the second word is allowed to be strobed in after the first.

2.0.10 Interlocks

2.0.10.1 Load Interlock

There is a single cycle load usage interlock in the microSPARC IU when a load instruction is followed by an instruction that uses the destination register of the load as a source operand.

2.0.10.2 Floating Point Interlocks

There are two types. The first is when the FPU is busy and a new floating point instruction is read into Decode. If the FPU detects a conflict, it will assert the FHOLD signal to prevent dispatch of that instruction until such time that the conflict is resolved.

The second is when a floating point branch enters decode and the FCCV bit from the FPU is deasserted. The interlock persists until the FPU asserts the FCCV bit.

2.0.10.3 Special Register Interlocks

Because of the execute datapath design, the microSPARC IU is unable to bypass special register read data to the instruction immediately following it in the pipeline. A single cycle interlock occurs.

2.0.11 Traps and Interrupts

2.0.11.1 Traps

The microSPARC IU implements all Version 8 traps except the following optional traps:

- data store error

- r register access error

- unimplemented FLUSH

- watchpoint detected

- coprocessor exception

Trap priorities are as defined in SPARC Version 8. If multiple traps occur during one instruction, only the highest priority trap is taken. Lower priority traps are ignored since it is assumed that lower priority traps will persist, recur, or are meaningless due to the presence of the higher priority trap.

In the pipeline, the trap indication usually occurs when the trapping instruction reaches the W stage of the pipeline. The exception to this are the exceptions detected by the MMU (e.g.: a LD which causes a data access exception trap) which occur in the MEMOP's R cycle. The reason for this difference is to allow the MMU an additional cycle to determine memory exceptions. Note that traps may be detected as early as the D cycle of the instruction. The trap indication is then piped to the W stage of that instruction.

After the assertion of the TRAP signal, instructions following the trapped instruction in the pipeline are flushed out. The PSR.ET <- 0, PSR.PS <- PSR.S, PSR.S <- 1, and PSR.CWP--. TBR.TT <- trapcode. The PC and nPC are written to r17 and r18. Instruction fetches then transfer operation to the trap vector as defined in the TBR.

The microSPARC IU does not allow help instructions to take a trap. There are no deferred integer traps. The IU will detect and act on deferred floating point traps.

2.0.11.2 Interrupts

The microSPARC IU is interrupted via the Interrupt Request Level bus. The IU depends on external logic to select the highest priority interrupting device and provide the appropriate IRL level. To discard glitches on the IRL lines, the IU must see at least two cycles where the level on the IRL are the same. Only then does it initiate an interrupt request to the processor. This request is pipelined by one cycle. The interrupt will be taken by the instruction currently in the W cycle of the pipeline (or, if that instruction is a help instruction, by the next non-help W cycle) if the IRL level is greater than the current PIL and there are no higher priority traps that take precedence.

Because there is a one cycle delay between when the IRL and PIL are compared and when the trap priorities are checked, this could cause a problem where back to back PSR writes could cause an interrupt to occur when the existing value in PSR.PIL is greater than the IRL. The microSPARC IU can prevent this from happening in hardware, so we avoid the difficulties encountered with previous designs.

2.0.11.3 Reset Trap

On reset, the following things occur:

- Traps are disabled (PSR.ET <- 0).
- If power-up reset, PSR.PS is undefined, else PSR.PS is unchanged.
- Enter supervisor mode (PSR.S <- 1).
- If power-up reset, PSR.CWP is undefined, else PSR.CWP is unchanged.
- If power-up reset, r[17] and r[18] are undefined, else are unchanged.
- If power-up reset, TBR.TT is undefined, else is unchanged.
- Execution begins at location PC=0 and nPC=4.

2.0.11.4 Error Mode

Error mode is entered when a trap occurs and PSR.ET = 0. Entry into error mode causes the following to happen:

- PSR.S <- 1.
- PSR.PS is unchanged.

- PSR.CWP --
- PC and nPC written to r[17] and r[18].
- PC <- 0, nPC <- 4.
- Assertion of the IU_ERROR signal.

In addition, the TBR.TT may be changed if the trapping instruction is an RETT. The TBR.TT will hold:

- With PSR.S = 0, TBR.TT will reflect privileged instruction
- With a window underflow, TBR.TT will reflect the underflow
- With a misaligned target address, TBR.TT will reflect the misaligned trap.

The IU will remain in error mode until it is reset.

2.0.12 Floating Point Interface

The microSPARC IU controls the addresses for all instructions and for floating point memory operations. Within the SingleSparc chip, the floating point unit has its own bus to the instruction cache. The IU provides the necessary strobe to load the FP's instruction register. This includes handling around instruction misses and instruction exceptions. In addition, the IU informs the FPU if the instruction just loaded is valid and should be continued down the pipe.

For floating point loads, the IU starts the cache access and the FPU reads the data. If the FPload causes a data cache miss, the IU will strobe the FPU's data register to pick up the data when it is available. For floating point stores, the IU starts the cache access and picks up the store data from the FPU. The IU then registers it and provides it to the data cache in the correct cycle(s).

When the FPU detects a usage conflict with the instruction just fetched in Decode, it asserts the FHOLD signal, which causes the IU to *interlock* the pipeline. The interlock is released when the FPU's internal status allows for the new FP instruction to start in the FPU.

FCC and FCCV are used by the IU to determine taken and untaken cases for floating point branches. If a floating point branch is detected in Decode and FCCV is not asserted, the IU will interlock until FCCV is asserted.

The FPU asserts the FEXC line when it detects a floating point exception. The IU will acknowledge the floating point exception (FXACK) when a floating point instruction is in the W stage of the pipe and the IU takes a floating point exception trap.

FPOps take one cycle in the IU, plus additional cycles in the FPU. For the number of cycles in the FPU, please refer to the FPU section in this document.

2.0.13 Special Features

The microSPARC IU has some build in features to make debug and bringup easier.

The IU is fully scanned, with all registers connected into the microSPARC IU scan chain (JTAG).

Certain registers of the scan chain are accessible only through the scan chain. These enable certain features useful for bringup and debug.

RF bypass - each read port has a bypass enable that causes the write data to be bypassed to the read port. Two registers in the scan chain can be set to enable this. These registers will be zeroed immediately on the next clock (when scan mode is off), disabling this feature.

Illegal opcode event - when this feature is enabled through the scan chain, the IU will assert the `iu_event` signal when a certain illegal opcode is decoded in the pipeline and the instruction causes an illegal instruction trap. The opcode in question is `op=10` binary and `op3=111111` binary. Once enabled, this feature can only be cleared through the scan chain.

IU error event - when this feature is enabled through the scan chain, the IU will assert the `iu_event` signal when the IU enters error mode. Once enabled, this feature can only be cleared through the scan chain.

2.0.14 Divergence from SPARC Version 8

The microSPARC IU has been designed to SPARC Version 8 compatible (as currently defined in the SPARC Architecture Manual, Version 8, Review-2) including hardware integer multiply and divide. microSPARC IU does deviate from full support of Version 8 features due to system design criteria. The deviations are as follows.

The microSPARC IU PSR is as implemented in the SPARC Version 8 manual. In early specifications of the microSPARC IU, it was stated that the EC bit of the PSR is not writeable. To maintain compatibility with Version 8 and IEEE 1754, the PSR.EC bit *is* writeable. Version 8 states that Coprocessor disabled traps occur when a coprocessor instruction is decoded and `PSR.EC=0` or *a coprocessor is not present*.

Alternate space memory operations proceed normally, however with a single caveat. Rather than the 8 bits of ASI, the microSPARC MMU only decodes 6 bits. The IU was directed to drop these bits, so out of bound ASI encodings are not detected.

The microSPARC IU does not implement STBAR since there is no need to force store ordering in this system. It will pass through the pipe as a Read Y Register operation with destination being the bit bucket (`%g0`).


The microSPARC IU also does not support reads and writes to the any Modes or Ancillary State Registers. We have no need for these. All read

cases will act like a Read Y Register operation. All write cases will act like a NOP.

The value read from the implementation field (IMPL) of the PSR for Tsunami will be (hexadecimal) 4. The value read from the version field of the PSR will be (hexadecimal) 1.

3.0 Floating Point Unit

The microSPARC Floating Point Unit is based on the Meiko FPU design. The Meiko FPU has been tailored for low-cost, and matched to the SPARC IU to balance performance. The FPU performance is more a result of the data cache hit rate, than the peak performance provided by the FPU design. The performance is therefore based on system level modeling, including the appropriate cache hit rates.

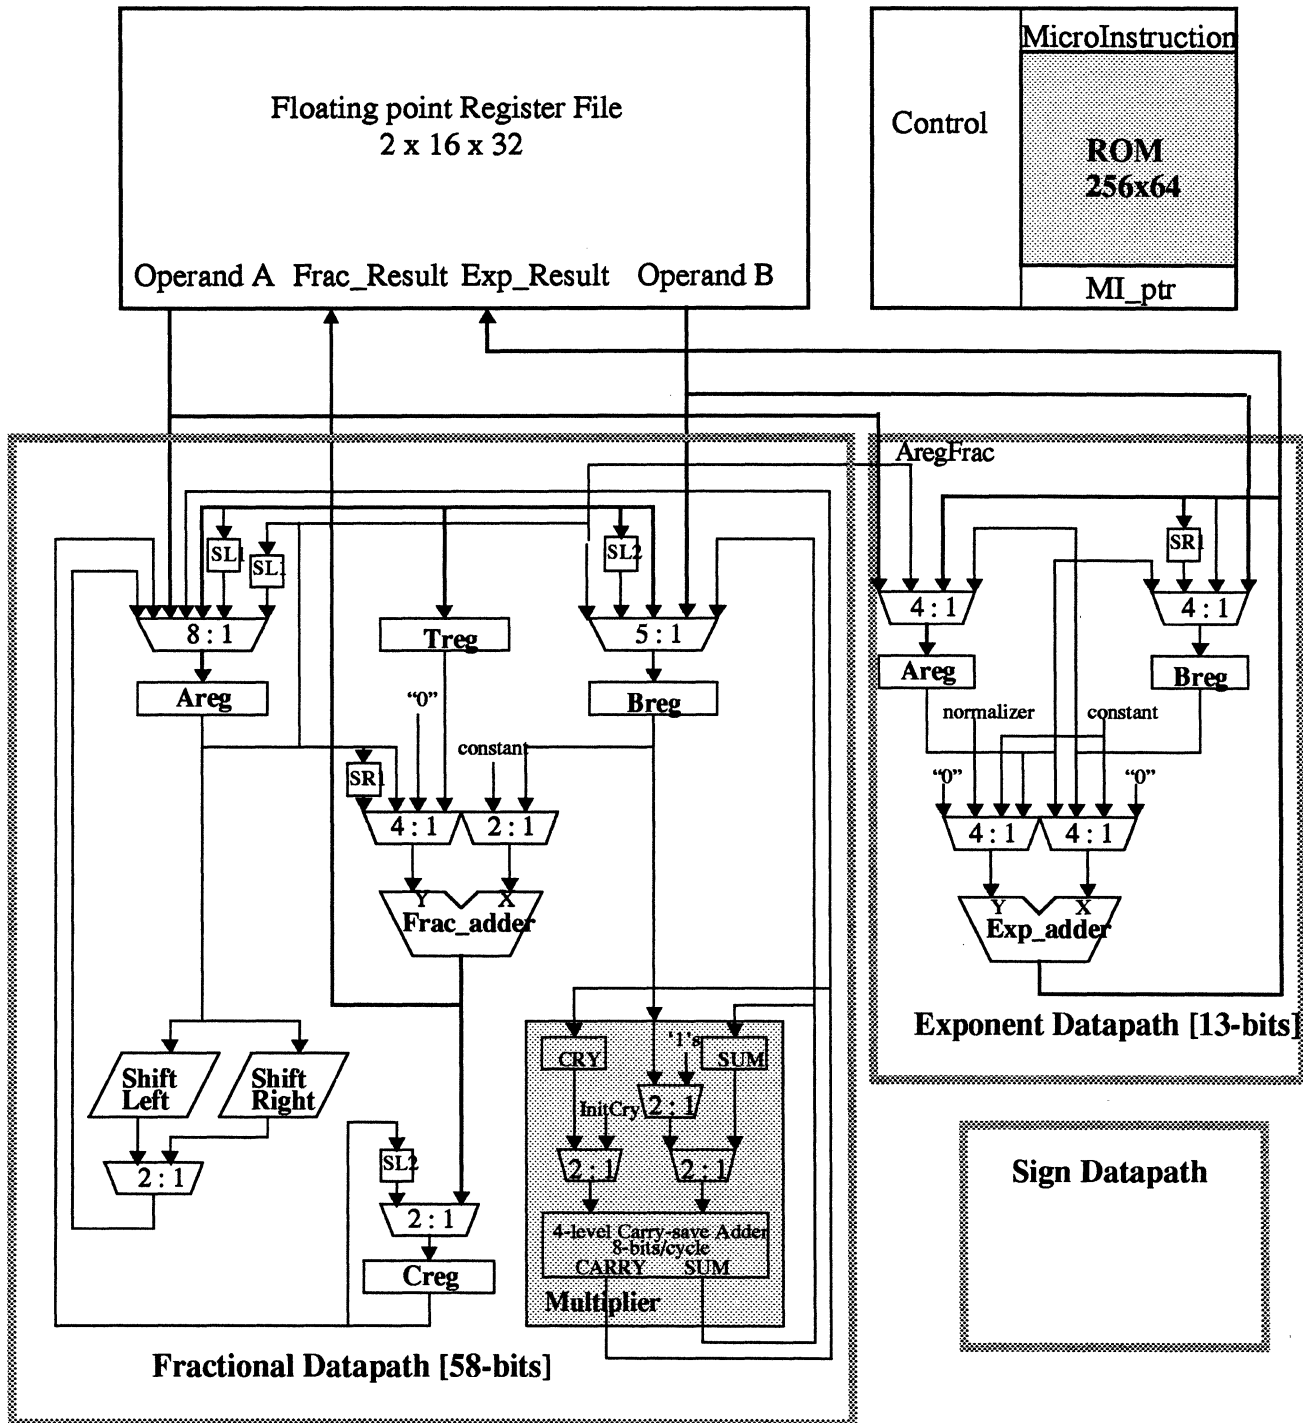


3.0.1 Overview

The Meiko FPU design is based on matching performance with the SPARC integer unit. The match is achieved by examining the maximum bandwidth of the integer unit in starting floating point operations and executing FPU LOADs and STOREs.

The Meiko FPU fully executes all single and double precision FP instructions as defined in the SPARC Architecture Manual (Version 8), except fsmuld. All other FP instructions (including fsmuld) trap to unimplemented. All implemented instructions will complete in hardware, therefore this FPU will never generate an unfinished exception. A block diagram follows:

Figure 3.0 - Meiko FPU Block Diagram



The bandwidth of the caches and main memory, and the integer unit's ability to fetch operands and schedule floating point instructions is the bottom line in performance. Through simulation, it has been determined

that the IU cannot provide data and schedule FP instructions at a rate faster than about 6 cycles per flop. The Meiko FPU can sustain floating operation times of about 5 cycles (as seen in LINPACK traces), and therefore will hardly impact overall operation time compared to an infinitely fast FPU.

The above conclusions allow a FPU implementation using multiple cycles to complete complex operations. The following algorithms were chosen for their positive trade-off in contributing to the final size and speed of the FPU.

- 8-bit multiply step
- 2-bit division step
- 1-bit square root step
- short distance (0-15 bits) shifter/normalizer
- separate single cycle round step
- microcode state machine to control FPU and decode operation

3.0.2 Deltas from SPARC Version 8

The microSPARC FPU deviates from SPARC Version 8 by not supporting the fsmuld instruction or quad-precision floating-point operations, and traps to unimplemented when these instructions are encountered. The microSPARC FPU also differs from the Appendix N, "SPARC IEEE 754 Implementation Recommendations" NaN format. The following figure shows the value returned for an untrapped floating-point result in the same format as the operands:

Figure 3.1 - Untrapped FP Result in Same Format as Operands

		rs2 operand		
		number	QNaN2	SNaN2
rs1 operand	none	IEEE 754	QNaN2	ME_NaN
	number	IEEE 754	QNaN2	ME_NaN
	QNaN1	QNaN1	QNaN1	ME_NaN
	SNaN1	ME_NaN	ME_NaN	ME_NaN

In the figure above, all QNaN results will have their sign bit set to 0. ME_NaN is 0x7ff0000 (single-precision) or 0x7ffe0000000000 (double-precision).

Figure 3.2 - Untrapped FP Result in Different Format operand (rs2)

operation	operand (rs2)			
	+QNaN	-QNaN	+SNaN	-SNaN
fstoi	ME_NaN	-imax	+imax	-imax
fstod	(QNaN2)	(QNaN2)	ME_NaN	ME_NaN
fdtos	ME_NaN	ME_NaN	ME_NaN	ME_NaN
fdtoi	ME_NaN	-imax	+imax	-imax

In the figure above, +imax = 0x7fffffff, and -imax = 0x80000000. (QNaN2) is a copy of the mantissa bits of the operand, with the extra low order bits zeroed, and the sign bit zeroed.

3.0.3 Implementation Specific Features

The microSPARC FPU implements a 1-entry floating-point deferred trap queue. When a floating-point instruction generates an fp_exception, microSPARC will delay the taking of an fp_exception trap until the next floating-point instruction is encountered in the instruction stream. The microSPARC FPU implementation can be modeled as having 3 states: fp_execute, fp_exception_pending, and fp_exception. These are shown in the figure below.

Normally the FPU is in fp_execute state. It moves from fp_execute to fp_exception_pending when an FPop generates a floating-point exception.

The FPU moves from fp_exception_pending to fp_exception, when the IU attempts to execute any floating-point instruction (including fbcc's). This transition (FXACK) generates an fp_exception trap. At this time the FQ contains the instruction and address of the FPop which originally caused the fp_exception.

An fp_exception trap can only be caused while the FPU is moving from the fp_exception_pending state to the fp_exception state (or by executing a STDFQ instruction when FSR.qne == 0, as described below). While in fp_exception state, only floating-point store instructions may be executed (particularly STDFQ and STFSR) and they can not cause an fp_exception trap.

The FPU remains in the fp_exception state until a STDFQ instruction is executed and the FQ becomes empty. At that time, the FPU returns to the fp_execute state.

If an FPop, or a floating-point load instruction (excluding fbcc's and all store instructions) is executed while the FPU is in fp_exception state, the FPU returns to fp_exception_pending state and also sets the FSR.ftt

field to `sequence_error` (0x4). The instruction that caused the `sequence_error` is not entered into the FQ.

If a STDFQ instruction is executed when the FQ is empty (`FSR.qne == 0`, FPU is in `fp_execute` state), the FPU will generate an immediate `fp_exception` trap (not deferred) and set the `FSR.ftt` field to `sequence_error` (0x4), but the FPU will remain in `fp_execute` state.

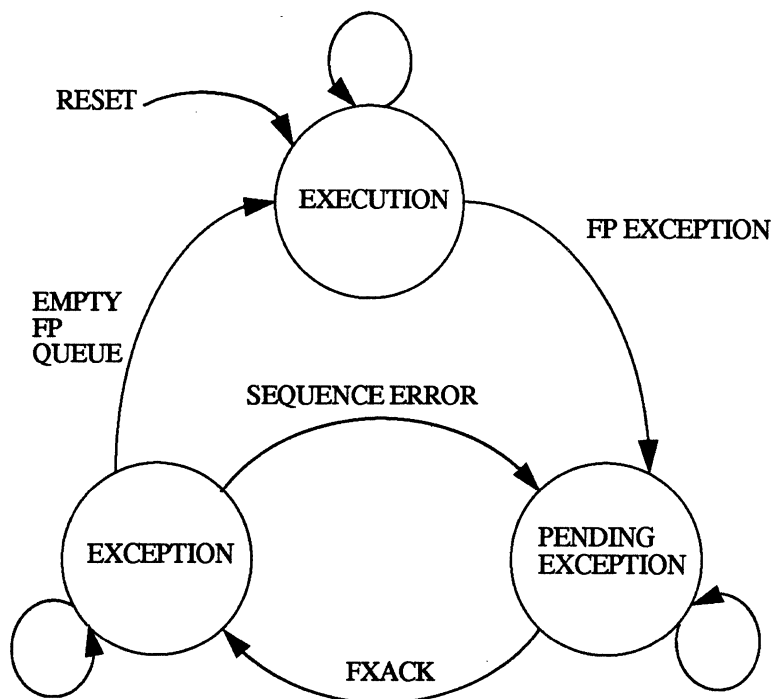


Figure 3.3 - FPU Operation Modes

The STDFQ instruction will store the address from the FQ to the effective address, and the instruction from the FQ to the effective address + 4.

3.0.4 Software Considerations

This section describes the software visible features of the microSPARC FPU/FPC.

The `FSR.ftt` field is set whenever an FPop completes or causes an exception. This field will remain unchanged until another FPop completes (or causes a sequence error). The `FSR.ftt` field may be cleared by executing a non-trapping FPop, such as `fmovs%f0,%f0`.

The following table describes the bits in the Floating-Point Status Register (FSR):

Table 3.0 - FSR Summary

FSR Bits	field	values	Description	writeable by LDFSR
31:30	RD	0 - Round to nearest (tie-even) 1 - Round to zero 2 - Round to +infinity 3 - Round to -infinity	Rounding Direction	Yes
29:28	res	always 0	reserved	No
27:23	TEM	0 - disables corresponding trap 1 - enables corresponding trap	Trap Enable Mask	Yes
22	NS	always 0	Nonstandard FP	No
21:20	res	always 0	reserved	No
19:17	ver	always 4	FPU version number	No
16:14	FTT	0 - None 1 - IEEE Exception 2 - Unfinished FPop 3 - Unimplemented FPop 4 - sequence error	FP trap type	No
13	QNE	0 - queue empty 1 - queue not empty	Queue Not Empty	No
12	res	always 0	reserved	No
11:10	FCC	0 - == 1 - < 2 - > 3 - ? (unordered)	FP Condition Codes	Yes
9:5	AEXC	0 - no corresponding exception 1 - corresponding exception	Accrued Exception Bits	Yes
4:0	CEXC	0 - no corresponding exception 1 - corresponding exception	Current Exception Bits	Yes

3.0.5 FPU Instruction Timings

The instruction timings, as quoted by Meiko, are provided in the following table. The timings are in CPU cycles.

Table 3.1 - FPU Instruction Timings

Instruction	Min	Typ	Max
fadds	4	4	17
faddd	4	4	17
fsubs	4	4	17
fsubd	4	4	17
fmuls	5	5	25
fmuld	7	9	32
fdivs	6	20	38
fdivd	6	35	56
fsqrts	6	37	51
fsqrtd	6	65	80
fnegs	2	2	2
fmovs	2	2	2
fabss	2	2	2
fstod	2	2	14
fdtos	3	3	16
fitos	5	6	13
fitod	4	6	13
fstoi	6	6	13
fdtoi	7	7	14
fcmps	4	4	15
fcmpd	4	4	15
fcmpes	4	4	15
fcmped	4	4	15
unimplemented	3	3	3

These cycle counts assume that the operands are available in the register file. A load-use interlock (fp load followed by an FPop which uses the destination register of the load as an operand) may add up to 2 cycles to the typical cycle count.

Because of the limited shifter size (0-15 bits was chosen to save hardware), the fpu instruction cycle counts are data dependent. There are 5 ways in which operations may take longer than the typical cycle count:


1. Exceptional operands (such as NaN, etc.) may add several cycles to the typical cycle count. In a normal environment, these are rare events probably caused by ill-conditioned data and will be trapped (if traps are enabled).
2. Possible exceptional results (results which are very close to underflow or overflow) may add up to 5 cycles to the typical cycle count. In a normal environment these are rare events, probably caused by ill-conditioned data.
3. Denormalized operands will add 1 extra cycle for each 15 bit shift required to normalize before the operation, and 1 extra cycle for each 15 bit shift required to denormalize the result after the operation (if necessary). Because operations on denormalized numbers will always complete in hardware (this fpu will never generate an unfinished exception), the overall performance will be greater than for an fpu which traps on denormalized operands.
4. Add or Subtract which require an initial alignment of more than 15 bits will add 1 extra cycle for each 15 bit shift. Also, a Subtract result which requires a shift of more than 15 bits to normalize will add 1 extra cycle for each 15 bit shift.
5. Non-standard rounding modes (RZ and RN are the typical operating modes) may require up to 3 additional cycles for some corner cases and exceptions.

Statistical analysis shows that, on average, 90% of fpu instructions will complete with the typical cycle count.

For a more detailed description of the Meiko floating point unit, please refer to the Meiko FPU specification, provided by Meiko Limited of Bristol, England.

4.0 Memory Management Unit

The microSPARC MMU provides the functionality of both a reference MMU as specified by the SPARC Reference MMU Architecture and a Sun 4M IO MMU. Additionally, much of the memory arbitration logic is contained within the MMU block.



4.0.1 Overview

This MMU provides four primary functions. First, the MMU translates virtual addresses of each running process to physical addresses in memory. More specifically, the MMU provides translation from a 32 bit virtual address to a 31 bit physical address by using a translation lookaside buffer (TLB). The 3 high order bits of Physical Address are maintained to support memory mapping into 8 different address spaces. The MMU supports the use of 64 contexts. Second, the MMU provides memory protection so that a process can be prohibited from reading or writing the address space of another process. Page protection and usage information is fully supported. Third, the MMU implements virtual memory. The page tables are maintained in main memory. When a miss occurs in the TLB the table walk is handled in hardware and a new virtual to physical address translation is loaded into the TLB. Finally, the MMU performs the arbitration function between IO, Data Cache, Instruction Cache, and TLB references to memory.

The reference MMU contains a 32 entry fully associative TLB and uses a pseudo random algorithm for the replacement of TLB entries. An address and data path block diagram follows.

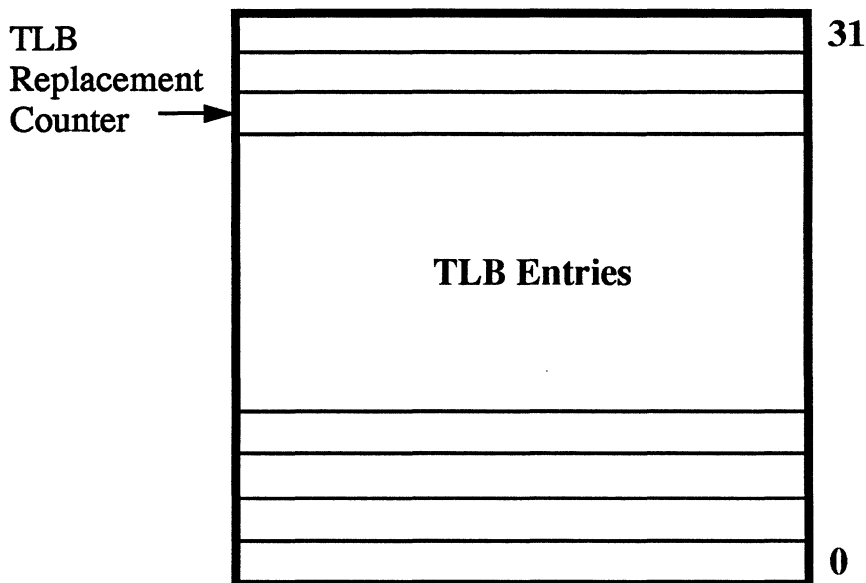
4.0.2 Translation Lookaside Buffer

The TLB is a 32 entry, fully associative cache of page descriptors. It caches virtual to physical address translations and the associated page protection and usage information. The pseudo random replacement algorithm determines which of the 32 entries should be replaced when needed. In the descriptions that follow the terms VA and PA are used to generically describe any virtual address (sb_ioa, iu_iva or iu_dva) or physical address (mm_pa, mm_dpa or mm_ipa) respectively.

4.0.2.1 TLB Replacement

The TLB uses a pseudo random replacement scheme. There is a 5 bit counter in the TLB Replacement Control Register (TRCR) which is incremented by one during each CPU clock cycle to address one of the TLB entries. When a TLB miss occurs, the counter value is used to address the TLB entry to be replaced. On reset the counter is initialized to zero. There is also a bit in the TRCR which is used to disable the counting function. A simple diagram follows.


Figure 4.1 - TLB Replacement



4.0.2.2 TLB Entry

An entry in the TLB has the following fields: a virtual address tag, a context tag, a PTE level field, and a page table field.

Figure 4.2 - TLB Entry



Field Definitions:

Virtual Address Tag - The 20 bit virtual address tag represents the most significant 20 bits (VA[31:12] the page address) of the virtual address being used when referencing PTEs and IOPTes. VA[11:00] is the byte within a page. The address in this field is physical when referencing PTPs with the least significant 19 bits containing PA[26:08].

Context Tag - The 6 bit context tag comes from the value in the context register as written by memory management software when referencing PTEs. Both it and the virtual address tag must match the CXR and VA[31:12] in order to have a TLB hit. This field contains a physical address (PA[07:02]) when referencing PTPs. This field is not used when referencing IOPTes.

Level - The 3 bit level field is used to enable the proper virtual tag match of region, and segment PTE's. IOPTes and PTP's will have this field set to use Index 1, 2 and 3 (b'111'). The most significant bit also serves as the TLB Valid bit because it is set for any valid PTE, IOPTe, or PTP. The following table defines the level field:

Table 4.1 - Virtual Tag Match Criteria

Level	Match Criteria
000	None
100	Index 1 (VA[31:24])
110	Index 1, 2 (VA[31:18])
111	Index 1, 2, 3 (VA[31:12])

Supervisor (S) - This bit is used to disable the matching of the context field indicating that a page is a supervisor level (ACC=6 or 7).

IO Page Table Entry (IO) - This bit indicates that an IOPTe resides in this entry of the TLB.

Page Table Pointer (PTP) - This bit indicates that a PTP resides in this entry of the TLB. Note that all SRMMU flush types (except page) will flush all PTPs from the TLB.

Page Table Field - The page table field can either be a Page Table Entry (PTE), a Page Table Pointer (PTP), or an IO Page Table Entry (IOPTe). This field can be read and written using ASI 0x06.

4.0.2.3 Page Table Entry

A Page Table Entry (PTE) defines both the physical address of a page and its access permissions. A PTE is defined for SPARC reference MMUs as follows.

Figure 4.3 - Page Table Entry in Page Table

Rsvd	PPN	C	M	R	ACC	ET
31 27 26		08 07	06	05	04 02 01	00

Field definitions:

Reserved (Rsvd) - Bits [31:27] should be written as zero, and will be read as zero.

Physical Page Number (PPN) - This field is the high order 19 bits ([30:12]) of the 31 bit physical address of the page. The PPN appears on PA[30:12] when a translation completes.

Cacheable (C) - When this bit is set to a one the page is cacheable by an instruction and/or data cache.

Modified (M) - This bit is set to a one when the page is written to.

Referenced (R) - This bit is set to a one when the page is accessed. All PTEs in the TLB have this bit set when the entry is loaded.

Access Permissions (ACC) - These bits indicate whether access to this page is allowed for the transaction being attempted. The Address Space Identifier (ASI) determines whether a given access is a data access or an instruction access, and whether the access is being done by the user or supervisor. The field is defined as follows.

Table 4.2 - Page Table Access Permissions

ACC	Permissions	
	User	Supervisor
0	Read only	Read only
1	Read/Write	Read/Write
2	Read/Execute	Read/Execute
3	Rd/Wrt/Exec	Rd/Wrt/Exec
4	Execute only	Execute only
5	Read only	Read/Write
6	No access	Read/Execute
7	No access	Rd/Wrt/Exec

Entry Type (ET) - This field differentiates the entry types in the TLB. Note that the entry type is not kept in the TLB RAM. On a probe operation the ET field is derived from a combination of other bits. The bit definitions of the ET field follows:

Table 4.3 - Page Table Entry Types

ET	Entry Type
0	Invalid
1	Page Table Pointer
2	Page Table Entry
3	Reserved

“Invalid” means that the corresponding range of virtual addresses is not currently mapped to a physical address.

In the TLB RAM the PTE has the following format:

Figure 4.4 - Page Table Entry in TLB

Rsvd	PPN	C	M	1	ACC	10
31 27 26		08 07	06	05	04 02	01 00

Bits [31:27] are not implemented, should be written as zero, and will be read as zero.

Bit [05] is set to one by hardware indicating that every PTE in the TLB has been referenced.

Bits [01:00] are set to one:zero by hardware indicating the entry type (ET) of a PTE. These bits are not actually stored in the TLB rather are derived as a function of the PTP bit of the tag.

4.0.2.4 Page Table Pointer

A Page Table Pointer (PTP) contains the physical address of a page table and may be found in the Context Table, in a Level 1 Page Table, or in a Level 2 Page Table. Page Table Pointers are put into the TLB during tablewalks and removed from the TLB either by natural replacement (also during tablewalks) or by flushing the entire TLB. Note that the Level field in a PTP tag is always set to 0x7. A PTP is defined as follows:

Figure 4.5 - Page Table Pointer in Page Table

Rsvd	PTP	Rsvd	ET
31 27 26		04 03 02 01	00

Field definitions:

Reserved (Rsvd) - Bits[31:27,03:02] should be written as zero, and will be read as zero.

Page Table Pointer (PTP) - The physical address of the base of a next level page table. The PTP appears on PA[30:08] during miss processing. The page table pointed to by a PTP must be aligned on a boundary equal to the size of the page table. Note that this is true of the context table at the root level also. The sizes of the tables are summarized as follows.

Table 4.4 - Sizes of Page Tables

Level	Size (Bytes)
Root	256
1	1024
2	256
3	256

Entry Type (ET) - This field differentiates the entry types in the TLB. Note that the entry type is not kept in the TLB RAM. On a probe operation the ET field is derived from a combination of other bits. The bit definitions of the ET field follows:


Table 4.5 - Page Table Entry Types

ET	Entry Type
0	Invalid
1	Page Table Pointer
2	Page Table Entry
3	Reserved

“Invalid” means that the corresponding range of virtual addresses is not currently mapped to a physical address.

In the TLB a PTP has the following format:

Figure 4.6 - Page Table Pointer in TLB



Bits [31:27] are not implemented, should be written as zero, and will be read as zero.

Bits [03:02] are set to zero by hardware and are unused.

Bits [01:00] are set to zero:one by hardware indicating the entry type (ET) of a PTP. These bits are not actually stored in the TLB rather are derived as a function of the PTP bit of the tag.

4.0.2.5 IO MMU Page Table Entry

An IO Page Table Entry (IOPTE) defines both the physical address of a page and its access permissions. Note that the Level field in a IOPTE tag is always set to 0x7 and the Supervisor bit is set to 0x0. An IOPTE is defined as follows.

Figure 4.7 - IO Page Table Entry in Page Table

Rsvd	PPN	Rsvd	W	V	WAZ
31 27 26		08 07 03	02	01	00

Field definitions:

Reserved (Rsvd) - Bits [31:27] are not implemented, should be written as zero, and will be read as zero. Bits [07:03] should also be written as zero, and will be read as zero.

Physical Page Number (PPN) - This field is the high order 19 bits of the 31 bit physical address of the page. The PPN appears on PA[30:12] when a translation completes. This address is concatenated with VA[11:00] to provide the entire translated address.

Writeable (W) - When this bit is set to a one both reads and writes to the page are allowed. When this bit is zero only reads are allowed.

Valid (V) - This bit is set to a one when the IOPTE is valid.

Write As Zero (WAZ) - This bit is to be written as zero in the memory io pagetable by software.

In the TLB an IOPTE has the following format:

Figure 4.8 - IO Page Table Entry in TLB

Rsvd	PPN	0	W	10
31 27 26		08 07	03 02	01 00

Bits [31:27] are not implemented, should be written as zero, and will be read as zero.

Bits [07:03] are set to zero by hardware. Bit[05] is used to distinguish between PTEs (set to 1) and IOPTEs (set to 0). Bits[07:06,04:03] are unused.

Bits [01:00] are set to one:zero by hardware indicating a valid IOPTE. These bits are not actually stored in the TLB.

4.0.3 CPU TLB Lookup

A virtual address to be translated by the MMU is compared to each entry in the TLB. During the TLB lookup the value of the Level field specifies which index fields are required to match the TLB virtual tag as follows:

Table 4.6 - Virtual Tag Match Criteria

Level	Match Criteria
000	None
100	Index 1 (VA[31:24])
110	Index 1, 2 (VA[31:18])
111	Index 1, 2, 3 (VA[31:12])


In addition to the virtual tag match, context matching of a PTE is required for all user page references (ACC is 0 to 5) when made by either user or supervisor (ASI = 0x8-0xB). Context matching is not required for a supervisor page reference (ACC is 6 or 7) when made by a supervisor (ASI = 0x9 or 0xB). This case takes advantage of the Supervisor bit in the TLB tag. Note that user references (ASI = 0x8 or 0xA) to supervisor pages (ACC is 6 or 7) result in address exceptions.

Note that the TLB ignores access level checking during probe operations. The most significant Level field bit is used as a Valid bit for the TLB. **This means that root level PTEs are not supported.**

4.0.4 CPU TLB Flush and Probe Operations

The flush operation allows software invalidation of TLB entries. TLB entries are flushed by using a store alternate instruction. The probe operation allows testing the TLB and page tables for a PTE corresponding to a virtual address. TLB entries are probed by using a load alternate instruction. The ASI value 0x3 is used to invalidate or probe entries in the TLB. In an alternate address space used for probing and flushing the address is composed as follows:

Figure 4.9 - CPU TLB Flush or Probe Address Format



Field Definitions:

Virtual Flush or Probe Address (VFPA) - This field is the address that is used to index into TLB. Depending on the type of flush or probe not all 20 bits are significant.

Type - This field specifies the extent of the flush or the level of the entry probed.

Reserved - These bits are ignored. They should be set to zero.

4.0.4.1 CPU TLB Flush

The flush operation must remove the PTEs and PTPs from the TLB that match the type criteria as follows:

Table 4.7 - TLB Entry Flushing

Type	Flush	PTE Match Criteria
0	Page	(Level 3) AND (Context match OR ACC=6-7) AND VA[31:12] match
1 to 4	Entire	None (Entire TLB Flush)
5 to F	Reserved	

4.0.4.2 CPU TLB Probe

The probe operation returns either a PTE from a page table in main memory or the TLB or it returns a zero if there is an invalid address or translation error while searching for the entry implied by the probe. If there is an error, a zero is returned for data. The reserved probe types (0x5-0xF) return an undefined value. A type 4 probe (entire) brings the accessed PTE and any PTPs that were needed into the TLB. If the PTE was not already there the referenced bit is updated. Probe type 0 affects

one entry of the TLB which is invalidated at the end of the probe operation.

Probe types 1-3 should be preceded by a TLB Flush Entire to ensure correct operation.

Table 4.8 - CPU TLB Entry Probing

Type	Probe	Returned Data
0	Page	Level 3 PTE or 0
* 1	Segment	Level 2 PTE or 0
* 2	Region	Level 1 PTE or 0
* 3	Context	Level 0 PTE or 0
4	Entire	PTE from Table Walk or 0
5 to F	Reserved	

* - Must be Preceded by TLB Flush Entire

4.0.5 Processor MMU Registers

The Processor Control Register (CR) contains general CPU control and status flags. The current context identifier is stored in the Context Register (CXR), and a pointer to the base of the context table in memory is stored in the Context Table Pointer Register (CTPR). If an MMU fault occurs on a CPU initiated transaction the address causing the fault is placed in the Synchronous Fault Address Register (SFAR) and the cause of the fault can be determined from the contents of the Synchronous Fault Status Register (SFSR). The TLB Replacement Control Register is used to control which TLB and Entries are to be replaced next. All of these internal MMU registers can be accessed directly by the processor

Implementation (IMPL) - The implementation number of this SPARC Reference MMU. This field is hardwired to 0x4 and read only.

Version (VER) - The version number of this SPARC Reference MMU. This field is hardwired to 0x1 read only.

Software Tablewalk enable (STW) - This bit enables the `instruction_access_MMU_miss` and `data_access_MMU_miss` traps for instruction and data tablewalking respectively for tablewalks to be done by software.

Address View (AV) - This bit is used for diagnostic purposes. Any address from the MMU Physical Address Register (PAR) is displayed on the SBus Address pins (SBADDR[27:00 = mm_pa[27:00]). This is a debug and test feature. During debug this can be monitored while running non io diagnostics. You cannot use the sbus while this bit is set.

Data View (DV) - This bit is used for diagnostic purposes. Any Data on the internal memory data bus will appear on the external SBus Data pins (SBDATA[31:00]). This is a debug and test feature. During debug this can be monitored while running non io diagnostics. You cannot use the sbus while this bit is set.

Memory Data View (MV) - This bit is used for diagnostic purposes. Any Data on the internal memory data bus (mdata[31:00]) will appear on the external memory data pins. This is useful for monitoring ASI and control space accesses (from/to both the IU and SBus). You cannot get to memory when this bit is set for either load or store operations.

Refresh Control (RC) - These 2 bits control the DRAM refresh rate of the system. Normal 40MHz operation would require a 0x2 value. The RC field is defined as follows:

Table 4.10 - Memory Refresher Control Definition

RFR_CNTL	Refresh Interval
0	Every 128 clocks (to 8.6 MHz)
1	No Refresh
2	Every 512 clocks (to 35 MHz)
3	Every 768 clocks (to 52 MHz)

Parity Control (PC) - This bit controls the generation of parity (and checking on memory reads) in the memory interface as follows:

Table 4.11 - Parity Control Definition

PC	Meaning
0	Even Parity
1	Odd Parity

ITBR Disable bit (ID) - This bit **disables** the use of the Instruction Translation Buffer Register when set.

Alternate Cacheability (AC) - This bit specifies that the caches are enabled by the IE and DE bits even with the mmu disabled when set. When not set, the caches are disabled when the mmu is disabled. This should not be used during boot mode accesses (or other instruction accesses to an sbus device).

Boot Mode (BM) - This bit is set by both sbus reset and watchdog reset and must be cleared for normal operation.

Parity Enable (PE) - When set to one this bit enables word parity checking for all non video data entering the processor over the memory bus.

Instruction Cache Enable (IE) - The instruction cache is enabled when this bit is set to a one. When zero, all references miss the cache. This bit is reset by both sbus reset and watchdog reset.

Data Cache Enable (DE) - The data cache is enabled when this bit is set to a one. When zero, all references miss the cache. This bit is reset by both sbus reset and watchdog reset.

No Fault bit (NF) - When set the supervisor accesses which cause exceptions will not be signaled to the processor (will be captured in the SFSR). Normal operation occurs while this bit is cleared.

MMU Enable (EN) - When this bit is set to a one the MMU is enabled and translation occurs normally. When this bit is not set the physical address is forced to the 31 least significant bits of the virtual address. This bit is reset by both sbus reset and watchdog reset.

4.0.5.2 Context Table Pointer Register

The Context Table Pointer Register (CTPR) contains the base of the Context table. It is defined as follows.

using ASI 0x4 and type 0x03 clears it. Using type 0x13 to read the SFSR does not clear it. Writes to the SFSR using ASI 0x4 and VA[12:08]=0x03 have no effect while writes using VA[12:08]=0x13 update the register. The SFSR is only guaranteed to be valid after an exception is actually signalled. In other words, it may not be valid if there is no exception.

Figure 4.13 - Synchronous Fault Status Register

Rsvd	CS	Rsv	PERR	Rsv	TO	BE	L	AT	FT	FAV	OW					
31	17	16	15	14	13	12	11	10	09	08	07	05	04	02	01	00

Field Definitions:

Reserved (Rsvd) - Bits [31:17,15,12] are not implemented, should be written as zero, and read as zero.

Control Space Error (CS) - This bit is asserted on the following conditions: [1] invalid ASI space, [2] invalid ASI size, [3] invalid VA field in valid ASI space and [4] invalid ASI operation (for example a swap instruction to an asi other than 0x8-0xB,0x20). Note that the AT field is not valid on Control Space Errors.

Parity Error (PERR) - The Parity Error[1:0] bits are set for external memory bus parity errors on the even and odd words respectively from memory.

Sbus Time Out (TO) - An Sbus Time Out resulted from a CPU initiated read transaction. No Sbus slave responded with an acknowledge within 256 Sbus cycles (12.8 us).

Sbus Bus Error (BE) - An error indication was returned from an Sbus slave on a CPU initiated read transaction. This may have been either an error acknowledgment or a late error.

Level (L) - The Level field is set to the page table level of the entry which caused the fault. If an error occurs while fetching a page table (either a PTP or PTE) this field records the page table level for the entry. The level field is defined as follows.

Table 4.12 - SFSR Level Field

L	Level
0	Entry in Context Table
1	Entry in Level 1 Page Table
2	Entry in Level 2 Page Table
3	Entry in Level 3 Page Table

Access Type (AT) - The Access Type field defines the type of access which caused the fault. Loads and Stores to user/supervisor instruction space can be caused by load/store alternate instructions with ASI = 0x8-0xB. The AT field is defined as follows. Note that this field is not valid on Control Space Errors.

Table 4.13 - SFSR Access Type Field

AT	Access Type
0	Load from User Data Space
1	Load from Supervisor Data Space
2	Load/Execute from User Instruction Space
3	Load/Execute from Supervisor Instruction Space
4	Store to User Data Space
5	Store to Supervisor Data Space
6	Store to User Instruction Space
7	Store to Supervisor Instruction Space

Fault Type (FT) - The Fault Type field defines the type of the current fault. The FT field is defined as follows.

Table 4.14 - SFSR Fault Type Field

FT	Fault Type
0	None
1	Invalid Address Error
2	Protection Error
3	Privilege Violation Error
4	Translation Error
5	Access Bus Error
6	Internal Error
7	Reserved

Invalid address errors, protection errors, and privilege violation errors depend on the AT field of the SFSR and the ACC field of the corresponding PTE. The errors are set as follows.

Table 4.15 - Setting of SFSR Fault Type Code

AT	FT Code								
	PTE[V]=0	PTE[V]=1							
		0	1	2	3	4	5	6	7
0	1	-	-	-	-	2	-	3	3
1	1	-	-	-	-	2	-	-	-
2	1	2	2	-	-	-	2	3	3
3	1	2	2	-	-	-	2	-	-
4	1	2	-	2	-	2	2	3	3
5	1	2	-	2	-	2	-	2	-
6	1	2	2	2	-	2	2	3	3
7	1	2	2	2	-	2	2	2	-

An invalid address error code (FT=1) is set when an invalid PTE or PTP is found while fetching an entry from the page table for a regular table walk or a probe entire operation. A translation error code (FT=4) is set when a SFSR PE type error occurs while the MMU is fetching an entry from a page table, a PTP is found in a level 3 page table, or a PTE has ET=3. The L field records the page table level at which the error occurred. The PE field records the word(s) having a parity error, if any. The protection error code (FT=2) is set if an access is attempted that is inconsistent with the protection attributes of the corresponding PTE. The privilege error code (FT=3) is set when a user program attempts to access a supervisor only page. An access bus error code (FT=5) is set when the SFSR PE field gets set on a memory operation that was not a table walk, or on a synchronously generated SBus error acknowledge or time out. Additionally, this error code is also set on an alternate space access to an unimplemented or reserved ASI or the memory access is using a size prohibited by the particular type of ASI. If multiple errors occur on a single access the highest priority fault is recorded in the FT field (see below).

Fault Address Valid (FAV) - The Fault Address Valid bit is set if the contents of the Synchronous Fault Address Register (SFAR) are valid. The SFAR is valid for data faults and data translation errors.

Overwrite (OW) - The Overwrite bit is set if the SFSR has been written more than once to indicate that previous status has been lost since the last time it was read.

Table 4.16 - Overwrite Operations

Pending Error	New Error	OW Status	Action Signalled
Translation Error	Translation Error	Set	Translation Error
Translation Error	Data Access Exception	Unchanged	Data Access Exception
Translation Error	Instruction Access Exception	Unchanged	Instruction Access Exception
Data Access Exception	Translation Error	Clear	Translation Error
Data Access Exception	Data Access Exception	Set	Data Access Exception
Data Access Exception	Instruction Access Exception	Unchanged	Instruction Access Exception
Instruction Access Exception	Translation Error	Clear	Translation Error
Instruction Access Exception	Data Access Exception	Clear	Data Access Exception
Instruction Access Exception	Instruction Access Exception	Set	Instruction Access Exception


If a single access causes multiple errors, the fault type is recognized in the following priority.

Table 4.17 - Priority of Fault Types on Single Access

Priority	Fault Type
1	Internal Error
2	Translation Error
3	Invalid Address Error
4	Privilege Violation Error
5	Protection Error


4.0.5.5 Synchronous Fault Address Register

The Synchronous Fault Address Register (SFAR) records the 32 bit virtual address of any data fault reported in the SFSR. The SFAR is overwritten according to the same policy as the SFSR on data faults. Reading the SFAR using ASI 0x4 and VA[12:08] 0x04 clears it. Using VA[12:08] 0x14 to read the SFSR does not clear it. Writes to the SFAR using ASI 0x4 and VA[12:08] 0x04 have no effect while writes using VA[12:08] 0x14 update the register. **Note that the SFAR should always be read before the SFSR to insure that a valid address is returned.** The structure of this register is as follows.

Figure 4.14 - Synchronous Fault Address Register

4.0.5.6 TLB Replacement Control Register

The TLB Replacement Control Register (TRCR) contains the TLB Replacement Counter and counter disable bit. The TRCR can be read and written using alternate load/store (LDA and STA) at ASI 0x4 with VA[12:08]=0x10. It is defined as follows.

Figure 4.15 - TLB Replacement Control Register

Field Definitions:

Reserved - Bits [31:06] are unimplemented, should be written as zero and will be read as zero.

TLB Replacement Counter Disable (TCD) - The TLBRC will not increment when this bit is set.

TLB Replacement Counter (TRC) - This is a 5 bit modulo 32 counter which is incremented by one during each CPU clock cycle to point to one of the TLB entries unless the TCD bit is set. When a TLB miss occurs, the counter value is used to address the entry to be replaced.

4.0.6 IO MMU Registers

The IO MMU Control Register (IOCR) contains IO MMU control and status flags. The IO MMU Base Address Register (IOBAR) defines the base address of the IO PTE Table in memory. The SBus Slot Configuration Registers (SSCR[0:3]) provides information about the slave device in the spare SBus slots. If a parity error occurs on an IO initiated transaction the physical address causing the fault is placed in the Asynchronous Fault Address Register (AFAR) and the cause of the fault can be determined from the contents of the Asynchronous Fault Status Register (AFSR). A DMA parity error will result in asserting the level 15 interrupt output (to be fed back to the IU externally as an interrupt) and the assertion of an error acknowledge to the SBC so it can return an SBus error acknowledge to the device that initiated the

transaction. IOPTE entries may be flushed from the TLB by doing writes to the Address Flush Register (AFR). This register is write only. All of these internal MMU registers can be accessed directly by software using SBus and IO MMU Control Space accesses with PA[30:24]=0x10. Also, the Entire TLB can be flushed using a control space access. The SBus and IOMMU Control Space address map follows.

Table 4.18 - SBUS and IO MMU Control Space

PA<30:00>	Device	R/W
1000 0000	IO MMU Control Register	R/W
1000 0004	IO MMU Base Address Register	R/W
1000 0014	Flush All TLB Entries	W
1000 0018	Address Flush Register	W
1000 1000	Asynchronous Fault Status Reg.	R/W
1000 1004	Asynchronous Fault Address Reg.	R/W
1000 1010	SBUS Slot Configuration Register0	R/W
1000 1014	SBUS Slot Configuration Register1	R/W
1000 1018	SBUS Slot Configuration Register2	R/W
1000 101C	SBUS Slot Configuration Register3	R/W
1000 1020	Memory Fault Status Register	R/W
1000 1024	Memory Fault Address Register	R/W
1000 2000	MID Register	R/W

4.0.6.1 IO MMU Control Register

The IO MMU Control Register (IOCR) contains control and status bits for the IO MMU. This register can be accessed using Sbus and IO MMU Control Space (0x10000000).

NOTE: Control space loads should not be executed while DMA is enabled (see MID register). A possible deadlock condition may occur if a DMA atomic or quad-word write coincides with the control space load.

The IOCR is defined as follows:

Figure 4.16 - IO Control Register

IMPL	VER	Rsvd	RANGE	Rsvd	ME
31 28 27 24 23			05 04 02 01 00		

Field definitions:

Implementation (IMPL) - The implementation number of this IO MMU. This field is hardwired to 0x4 and read only.

Version (VER) - The version number of this IO MMU. This field is hardwired to 0x1 and read only.

Reserved (Rsvd) - Bits [23:05,01] are not implemented, should be written as zero, and will be read as zero.

RANGE - This field defines the virtual address range for DVMA. Specifically, the translatable limit is defined to be $16\text{MB} * 2^{**} <\text{RANGE}>$. All VA bits above this limit must be set to one for an address to be valid. For example, if RANGE=2 then 64MB of virtual address are supported, and valid DVMA virtual addresses range from 0xFC000000 to 0xFFFFFFFF. Any access using a DVMA virtual address that is out of that range will receive an SBus error acknowledge. The only exception involves slots that have Bypass Enabled. The following table shows how the physical address of an IO MMU page table entry is generated:

Table 4.19 - IO MMU Page Table Address Generation


Range	Limit	Physical Address[30:00]
0	16MB	IBAR[26:10], IOVA[23:12],b'00'
1	32MB	IBAR[26:11], IOVA[24:12],b'00'
2	64MB	IBAR[26:12], IOVA[25:12],b'00'
3	128MB	IBAR[26:13], IOVA[26:12],b'00'
4	256MB	IBAR[26:14], IOVA[27:12],b'00'
5	512MB	IBAR[26:15], IOVA[28:12],b'00'
6	1GB	IBAR[26:16], IOVA[29:12],b'00'
7	2GB	IBAR[26:17], IOVA[30:12],b'00'

IO MMU Enable (ME) - IO MMU translation is enabled when this bit is set.

4.0.6.2 IO MMU Base Address Register

The IO MMU Base Address Register (IBAR) defines base address of the IO Reference Table. This register can be accessed using Sbus and IO MMU Control Space (0x10000004). The IBAR is defined as follows.

Figure 4.17 - IO MMU Base Address Register



Field definitions:

Reserved (Rsvd) - Bits [31:27,09:00] are not implemented, should be written as zero, and will be read as zero.

IO MMU Base Address (IBA) - When the IO MMU is enabled and the access translation misses the TLB, IBA is used as the base address for the (<RANGE/1024>)byte-aligned IO MMU Reference Table.


4.0.6.3 IOMMU Flush All TLB Entries

All TLB entries are flushed by writing to control space address PA=0x10000014. This address should not be read since the output of the TLB is unknown during a flash clear operation.

4.0.6.4 IOMMU Address Flush Register

The IOPTTE entries may be flushed from the TLB by doing writes to the Address Flush Register at PA=0x10000018 with the following format. The Address Flush Register is defined as follows.

Figure 4.18 - IOPTTE Address Based Flush Format



Field definitions:

Reserved (Rsv) - Bits [11:00] are not implemented and should be written as zero.

Flush Address (FA) - The virtual page address of the IOPTTE entry to be flushed.

Note that a register is not actually implemented to perform this function. Also note that to flush all IOMMU entries all TLB entries must be flushed (see section on CPU TLB Flush for details).

4.0.6.5 Asynchronous Fault Status Register

The Asynchronous Fault Status Register (AFSR) provides information on asynchronous faults during IO initiated transactions and CPU write operations. This register is used only for PIO operations, and is accessed using Sbus and IO MMU Control Space (0x10001000). A hardware lock is used to ensure that this register does not change while being read. Reading this register clears it. Multiple errors set the ME bit, but do not change any other states. The AFSR always reflects the status of the first error. Refer to the Sun 4M specification.

Note: The AFSR.size field is invalid when a late error (AFSR.le) is detected.

Note: Due to the pipelined nature of Processor I/O space writes, it is possible to receive a late error (AFSR.le) and no longer have the correct address stored in the AFAR. When this occurs, the AFSR.fav bit will not be asserted, indicating that the AFAR contains an invalid address.

Figure 4.19 - Asynchronous Fault Status Register

ERR	LE	TO	BE	SIZE	S	1000	ME	RD	FAV	Rsvd			
31	30	29	28	27	25	24	23	20	19	18	17	16	00

Field Definitions:


- Reserved (Rsvd) - Bits [23:20,16:00]. Bits [23:20] are forced to '1000'. Bits [16:00] are not implemented, should be written as zero, and read as zero.
- Summary Error Bit (ERR) - One or more of LE, TO, or BE is asserted.
- Late Error (LE) - The SBus reported an error after the transaction was done.
- Time Out (TO) - An SBus write access timed out.
- Bus Error (BE) - An SBus write access received an error acknowledge.
- Size (SIZE) - SBus size of error transaction.
- Supervisor (S) - CPU was in Supervisor mode when error occurred.
- Multiple Error (ME) - At least one other error was detected after the one shown.
- Read Operation (RD) - The error occurred during a read operation.
- Fault Address Valid (FAV) - The address contained in the AFAR is accurate and can be used in conjunction with the status in AFSR. The only time the AFAR will be invalid is on an SBus late error in which the second processor IO operation has already been requested and is queued up in the SBC.

4.0.6.6 Asynchronous Fault Address Register

The Asynchronous Fault Address Register (AFAR) records the 31 bit physical address that caused the fault. This register is accessed using Sbus and IO MMU Control Space (0x10001004). Bit [31] should be written as zero and will be read as zero. A hardware lock is used to

insure that this register does not change while being read. Writing the AFSR unlocks the AFAR. The structure of this register is as follows.

Figure 4.20 - Asynchronous Fault Address Register




Note that bit 31 is unimplemented, should be written as zero, and will be read as zero. Also, this register is only held when an error is reflected in the AFSR.

4.0.6.7 SBUS Slot Configuration Registers

The SBus Slot Configuration Registers (SSCR[0:3]) provide information about the slave device in sbus slots, and is also used for IO MMU bypass management for that slot. These registers can be accessed using Sbus and IO MMU Control Space (0x10001010, 0x10001014, 0x10001018 and 0x1000101C respectively). The SSCR is defined as follows:

Figure 4.21 - SBUS Slot Configuration Register



Field definitions:

Reserved - Bits [31:17,15:03] are not implemented, should be written as zero, and will be read as zero.

Segment Address Bit 30 (SA30) - This bit provides PA[30] when IO MMU bypass is used.

BA16 - Slave supports 16 byte bursts.

BA8 - Slave supports 8 byte bursts.

IO MMU Bypass (BY) - When this bit is set the MMU is bypassed and the virtual addresses from this slave are treated as physical when sb_ioa[31:30]=00. mm_pa[30] is given by the SA30 field and mm_pa[29:00] is defined as sb_ioa[29:00].

4.0.6.8 Memory Fault Status Register

The Memory Fault Status Register (MFSR) provides information on parity faults. This register is accessed using Sbus and MMU Control Space (0x10001020). This register is loaded on every request to memory

unless it is locked. A hardware lock is used to ensure that this register does not change while being read if there was an error condition. Reading this register allows it to begin loading once again.

When multiple memory errors occur, the MFSR will hold the status reflecting the operation in which the first error occurred, and also set the multiple error bit (MFSR.me). The MFSR will maintain the error status until cleared, which can be done by reading the MFSR.

Figure 4.22 - Memory Fault Status Register

ERR	Rsvd	S	CP	Rsvd	ME	Rsvd	PERR	BM	C	Rsvd	Type	Rsvd							
31	30	25	24	23	22	20	19	18	15	14	13	12	11	10	08	07	04	03	00

Field Definitions:

Reserved (Rsvd) - Bits [30:25,22:20,18:15,10:08,03:00] are not implemented, should be written as zero, and read as zero.

Summary Error Bit (ERR) - One or more of PERR[1] or PERR[0] is asserted.

Supervisor (S) - CPU was in Supervisor mode when error occurred.

CPU Transaction(CP) - CPU initiate the transaction that resulted in the parity error.

Multiple Error (ME) - At least one other error was detected after the one shown.

Parity Error[1:0] (PERR) - These bits are set on external memory parity errors for the even and odd words (respectively) from memory. Parity errors can result from CPU or IO initiated memory reads and byte or halfword (8 or 16 bit) write operations (which result in read-modify-writes).

Boot Mode (BM) - This bit indicates that the error occurred while the PCR was indicating that we were in Boot Mode.

Cacheable (C) - Address of error was mapped cacheable. On a CPU initiated transaction this bit is from the C bit of the PTE, otherwise it is set to zero.

Memory Request Type (Type[3:0]) - This field records the type of request that generated the parity error as follows:


Table 4.20 - Memory Request Type

Value (Hex)	Name	Meaning
0	NOP	No memory operation
1	RD64	Read of 64 bits (2 words)
2	RD128	Read of 128 bits (4 words)
3		Reserved
4	RD256	Read of 256 bits (8 words)
5		Reserved
6		Reserved
7		Reserved
8		Reserved
9	WR8	Write of 8 bits (1 byte)
A	WR16	Write of 16 bits (2 bytes)
B	WR32	Write of 32 bits (1 word)
C	WR64	Write of 64 bits (2 words)
D		Reserved
E		Reserved
F		Reserved

4.0.6.9 Memory Fault Address Register

The Memory Fault Address Register (MFAR) records the 31 bit physical address that caused the fault. This register is accessed using Sbus and IO MMU Control Space (0x10001024). This register is loaded on every request to memory unless it is locked. A hardware lock is used to ensure that this register does not change while being read if there was an error condition. Reading this register allows it to begin loading once again. Bit [31] should be written as zero and will be read as zero. The structure of this register is as follows.

Figure 4.23 - Memory Fault Address Register



Note that bit 31 is unimplemented, should be written as zero, and will be read as zero. Also, this register is only held when an error is reflected in the MFSR.

4.0.6.10 MID Register

The MID Register contains two fields. The MID field (Bits[3:0]) contain a constant value of 0x8 and the SBAE field which controls the ability

of SBus devices to arbitrate for the bus. This register can be accessed using Sbus and IO MMU Control Space (0x10002000). The SBAE bits are both readable and writeable while the MID field is read only. The MID is defined as follows:

Figure 4.24 - MID Register

Reserved	SBAE[4:0]	Reserved	'0x8'
31	21 20 16 15	04 03	00

Field definitions:

Reserved - Bits [31:21,15:04] are not implemented, should be written as zero, and will be read as zero.

SBus Arbitration Enable[4:0] (SBAE) - These bits control the ability for devices on the SBus to arbitrate for the bus. The most significant bit (SBAE[4]) controls arbitration for the SCSI/Ethernet master. The other bits (SBAE[3:0]) control arbitration for SBus devices 3:0 corresponding to SSCR[3:0]. These bits are R/W.

MID - This field is a constant 0x8 and is read only (writes to these bits are ignored).

4.0.7 IO MMU Bypass Mode


Bypass mode is provided to allow intelligent SBus masters to do their own memory management with assistance from the kernel. This facility is enabled by having the Bypass Enable bit set in that device's slot configuration register. It is assumed that such a master will have its own MMU. In order to bypass the IO MMU the DVMA master must issue a virtual address with sb_ioa[31:30]=0. In this case the Physical Address bus will have the Virtual Address bus put on it. The PA is checked to verify that it is in the valid main memory range and an error is issued to the master if it is not.

4.0.8 Physical Address Register

The Physical Address Register (PAR) is used to hold translated physical addresses before they are used for either memory requests or for Sbus

operations. This register cannot directly be read or written. The structure of this register is as follows.

Figure 4.25 - Physical Address Register

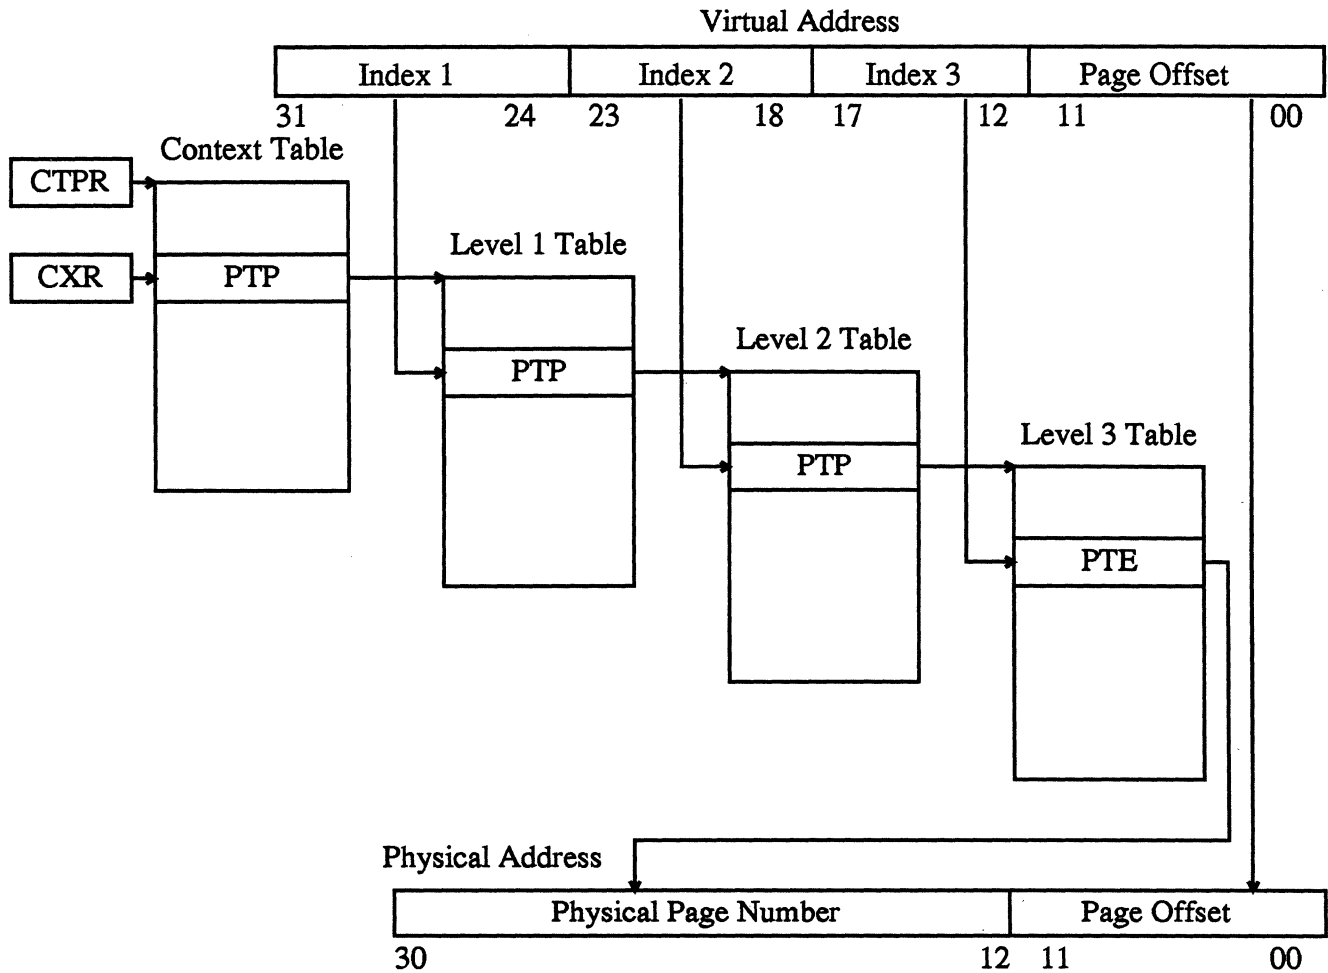


4.0.9 TLB Table Walk

On a translation miss the table walk hardware translates the virtual address to a physical address by “walking” through a context table and from 1 to 3 levels of page tables. The first and second levels of these tables typically (not necessarily) contain page table pointers (PTP) to the next level tables when accesses are due to CPU instruction or data addresses. IO accesses only the first level page table. A third level table entry should always be a page table entry (PTE) pointing to a physical page or else a translation fault occurs.

The table walk for a CPU generated virtual address uses the context table pointer register (CTPR) as a base register and the context number contained in the context register (CXR) as an offset to point to an entry in the context table. The context table entry is then used as a PTP into the first level page table. At any address the table walk hardware finds either a PTE which terminates its search or a PTP. A PTP is used in conjunction with a field in the virtual address to select an entry in the next level of tables. The table walk continues searching through levels of tables as long as PTPs are found pointing to the next table. The table walk terminates when either a PTE is found or an exception is generated if a PTE is not found after accessing the 3rd level page table (or if an invalid or reserved entry is found). Note that PTPs and PTEs encountered during a table walk are not cached in the data cache. A full table walk is shown in the following figure.

Figure 4.26 - CPU Address Translation Using Table Walk



When the PTE is found it is stored in an available TLB entry and used to complete the original virtual to physical address translation. A table walk which was forced by a store operation to an unmodified region of memory causes the M bit in the PTE to be set. Any "entire" probe or normal tablewalk operation causes the R bit of the PTE to be set if it had not been already.

The table walk for an IO generated virtual address uses the IO Base Address Register (IOBAR) as a base register and part of the DVMA virtual address as an index into an IOPTe table in memory. Specifically the IO MMU page table size and corresponding DVMA virtual address range are configured in the IOCR RANGE field. The table consists of 4 byte entries. The virtual address used for this mapping is VA[X:0] where "X" is the highest VA bit in the translatable range. VA[31:X+1] must be all "1"s in order for translation to take place; otherwise an error is signalled to the DVMA master. The bits VA[X:12] provide a virtual

page number which is used as an index into the IOMMU table in memory. These bits are placed on $mm_pa[X-10:2]$. The rest of the physical address is $mm_pa[1:0] = 00$, and $mm_pa[30:X-9] = IBA[30:X-9]$. This is the PA used for the one level IO walk.

4.0.10 Instruction Translation Buffer Register

Since instruction fetches occur every time the pipeline moves and there is only one TLB for translating instruction references, data references and DVMA requests, a method for dealing with conflicts between instruction references and data or IO references to the TLB was needed. A registered version of the last instruction translated TLB line is kept in the Instruction Translation Buffer Register (ITBR). When the TLB arbiter determines there is a conflict the iu_iva goes to the ITBR and the two translations occur simultaneously. When the iu_iva misses in the ITBR the translation is done in the TLB the next available cycle. Note that the default is to translate instruction addresses in the TLB and the ITBR is used only for conflict cases. This maximizes the hit rate of instruction address lookups. Each time an iu_iva is successfully translated in the TLB the ITBR is updated. The ITBR is logically split into a PTE and Tag section. Both the PTE and tag portions of the ITBR are read and written like other TLB PTE and tags using ASI 0x6. See diagnostic section for details.

An I-cache miss will require a translation using the TLB, as there is no datapath from the ITBR to PAR. Therefore, the ITBR is only useful for cached pages.

Any access error detected by the ITBR is seen as an ITBR miss, without updating any Fault Status logic. Normal execution will retry the translation using the TLB, and set the Fault Status logic accordingly.

The ITBR is invalidated whenever the TLB is written or flushed to maintain consistency. The ITBR is always a copy of the TLB entry, not an additional entry.

4.0.10.1 ITBR Page Table Entry

An ITBR Page Table Entry (ITBR/PTE) defines both the physical address of a page and its access permissions. A ITBR/PTE is defined as follows.

Figure 4.27 - ITBR Page Table Entry

Rsvd	PPN	C	Rsvd	Rsvd R	ACC	Rsvd ET
31 23 22	08 07	06	05	04	02 01	00

Field definitions:

Reserved (Rsvd) - Bits [31:23,06:05,01:00] are not implemented, should be written as zero, and will be read as zero except for bits [05 and 01] which are read as one. This was done to make the ITBR appear as a valid PTE when read. Bit [06] is the M bit (=0), bit [05] is the R bit (=1) and bits [01:00] is the ET field (=10 for PTE).

PPN, C, ACC - these fields are defined the same as they for TLB PTEs. Note that the 4 most significant PPN bits are not kept in the ITBR since instruction references must be made to main memory (limit 128MB in address space 0).

4.0.10.2 ITBR Tag

An ITBR Tag is defined in the section on MMU diagnostic strategy. Briefly, the tag consists of the Level field, the Instruction Virtual Address Tag, and the Context Tag.

4.0.11 Arbitration

The MMU block performs the primary memory arbitration function on the CPU. This is due to the central nature of the MMU in the address flow of the machine. The different sources of memory activity are the instruction cache block (for instruction fetches), the data cache block (for loads and stores), the TLB (during tablewalks and to keep the referenced and modified bits in the main memory page tables up to date), and IO DMA activity.

The other entity needing main memory is the DRAM refresh logic. This function is folded into the arbitration scheme by the Memory Controller which must arbitrate between it and a request out of the MMU.

The arbitrating requirements can be broken down into several different resource arbiters. The TLB (and ITLB) arbitration and the internal memory bus arbitration.

4.0.11.1 TLB Arbitration

The current priority scheme places TLB references as highest priority, followed by IO references, data references, and finally instruction references. Note that the TLB is referenced during every CPU clock in normal operation. Tablewalks and updates to the memory PTEs due to changes to the Referenced and Modified bits are the highest priority. They imply that some other operation is in progress.

Table 4.21 - TLB Reference Priority

Operation Pending			Result
IO DMA	IU Data Ref.	Instr. Fetch	
YES	X	X	Xlate for IO, Tablewalk if miss, use ITBR for IFetch Xlate
NO	YES	X	Xlate for IU Data Reference, Tablewalk if miss, use ITBR for IFetch Xlate
NO	NO	YES	Xlate for Instruction Fetch, Tablewalk if miss, load ITBR with Xlate output

Note: X=Don't Care, Xlate=Translate

4.0.12 Translation Modes

Translation of virtual addresses to physical addresses is done in the following modes:

Table 4.22 - Translation Modes

Name	ASI	Boot Mode	MMU En.	PA[30:00]
Boot IFetch	0x8, 0x9	Yes	Off	PA[30:28]=0x7, PA[27:00]=VA[27:00]
Pass Through	0x8, 0x9	No	Off	PA[30:00]=VA[30:00]
Translate	0x8, 0x9	No	On	PA[30:12]=PTE[26:08], PA[11:00]=VA[11:00]
Pass Through	0xA, 0xB	X	Off	PA[30:00]=VA[30:00]
Translate	0xA, 0xB	X	On	PA[30:12]=PTE[26:08], PA[11:00]=VA[11:00]
Bypass	0x20	X	X	PA[30:00]=VA[30:00]

4.0.13 Page Mode Detection

The MMU is responsible for generating a signal to the memory controller indicating whether or not the current memory request can use page mode of the DRAMs or not. This is done by comparing the contents of the MFAR (at the time of the last request) with the current physical address (mm_pa) the cycle before a request is ready. Specifically, bits [26:12] have to match between MFAR and the PA. If these bits match then the MMU will assert PAGE. The memory controller then has the option of using a page mode DRAM access or not. If mm_page is not asserted then a page mode access cannot be used to fulfill the request.

4.0.14 Errors and Exceptions

The MMU generates: instruction access error, instruction access exception, data access error, and data access exception for the SPARC IU. Also, an external interrupt is driven for asynchronous faults. In a Sun4M system, this would indicate a level 15 interrupt.

4.0.15 Diagnostic Features

All registers and RAM (and CAM) are accessible directly through alternate virtual address space loads and stores. In addition to this control is provided for putting the internal memory data bus onto the external memory data or SBus data pins. Also, any generated physical address can be seen at the SBus address pins.

There is also the ability to breakpoint on certain conditions. This is set up through use of the scan chain. More details follow.

4.0.15.1 Diagnostic Access of TLB, ITBR

Diagnostic reads and writes to the 32 TLB entries and the ITBR are performed by using load and store alternate instructions in ASI 0x6 and the virtual address to explicitly select a particular TLB entry. The access must be a word access, all other data sizes will result in an internal error. Depending on the virtual address specified either the TLB Tag, TLB PTE, ITLB Tag or ITLB PTE will be referenced. The format for the TLB PTE is as described earlier. The format of the Tag is shown below: (Note that bits [02:00] are not valid for an itbr tag and are read as zero)

Figure 4.28 - CPU Diagnostic TLB and ITLB Tag Access Format

Virtual Address Tag	Context Tag	V	Level	S	IO	PTP
31	12 11	06 05	04 03	02	01	00

Field Definitions:

Virtual Address Tag - The 20 bit virtual address tag represents the most significant 20 bits (VA[31:12] the page address) of the virtual address being used. VA[11:00] is the byte within a page. The address in this field is physical when referencing PTPs with the least significant 19 bits containing PA[26:08].

Context Tag - The 6 bit context tag comes from the value in the context register as written by memory management software. Both it and the virtual address tag must match the CXR and VA[31:12] in order to have a TLB hit. This field contains a physical address (PA[07:02]) when referencing PTPs.

Valid bit, Level bits - These 3 bits are used to enable the proper virtual tag match of root, region, and segment PTE's. The Valid bit indicates a valid entry.

Supervisor (S) - This bit is used to disable the matching of the context field indicating that a page is a supervisor level (ACC=6 or 7). This bit is non meaningful for an ITLB Tag and is read as 0.

IO Page Table Entry (IO) - This bit indicates that an IO PTE resides in this entry of the TLB. This bit is non meaningful for an ITLB Tag and is read as 0.

Page Table Pointer (PTP) - This bit indicates that a PTP resides in this entry of the TLB. Note that all SRMMU flush types (except page) will flush all PTPs from the TLB. This bit is non meaningful for an ITLB Tag and is read as 0.

Note that when loading TLB entries under software control (using alternate space accesses) care should be taken to ensure that multiple TLB entries cannot map to the same virtual address. This may inadvertently occur when combining TLB entries that map different sizes of addressing regions. A level 3 PTE could be included in a TLB region for a level 1 or 2 PTE for example. The TLB output is not valid when this occurs.

Note: Any sta to the TLB tag or data must be followed by 3 nops. This is to allow the pipelined TLB write sufficient time to complete.

The virtual address is used to select the TLB entries as follows:

Table 4.23 - TLB Entry Address Mapping

Virtual Address	TLB Entry
0x0	Entry 0 PTE
0x4	Entry 0 Tag
0x8	Entry 1 PTE
0xC	Entry 1 Tag
0x10	Entry 2 PTE
0x14	Entry 2 Tag
0x18	Entry 3 PTE
0x1C	Entry 3 Tag
0x20	Entry 4 PTE
0x24	Entry 4 Tag
0x28	Entry 5 PTE
.	.
.	.
.	.
0xF0	Entry 30 PTE
0xF4	Entry 30 Tag
0xF8	Entry 31 PTE
0xFC	Entry 31 Tag
0x100-0x7FC	Reserved
0x800	ITBR PTE
0x804	ITBR Tag
0x808-FFFFFFFC	Reserved

**4.0.15.2 MMU
Breakpoint
Debug Logic**

The MMU breakpoint debug logic is intended for use in lab debug only since it requires setup through a scan facility. The basic idea is to stop the clocks when certain conditions occur. This facility is general purpose in that there is a large matrixed selection of conditions to choose from. The breakpoints which can be enabled are virtual address matching, virtual address source matching, virtual address type matching, memory request matching, tablewalk detection (includes type), and tablewalk level matching. A more detailed description and suggested pairings of these conditions follows.

We have the ability to breakpoint on portions of the virtual address (the output of the virtual address muxing logic). The ITBR must be turned off to guarantee matches on instruction addresses. These portions of the virtual address can be combined with other conditions to make their match conditions more case specific as follows:

Table 4.24 - Virtual Address Match Conditions

Virtual Address Conditions	Conditions to be Paired With
VA[31:00]	Any address translation: io_tlb (DMA read, write or translate) dc_tlb (iu load, store, or atomic op) ic_tlb (instruction translation) or The following cycle types: read_w (iu load in w stage) write_w (iu store in w stage) ldsto_w (iu atomic in w stage) iu_fetch_f (instr. fetch in f stage) sb_read (DMA read op) sb_write (DMA write op) sb_translate (DMA translate - before DMA write op)
VA[31:01]	
VA[31:02]	
VA[31:03]	
VA[31:12]	
VA[31:18]	
VA[31:24]	
VA[10:02]	
VA[11:02]	
!VA[31:12] & VA[11:02]	
!VA[31:11] & VA[10:02]	

The virtual address breakpoint control register enables specific address bits for comparison. The details of the register are listed below:

Table 4.25 - Virtual Address Field Enable Decode

10	9	8	7	6	5	4	3	2	1	0
I 31:24	H 23:18	G 17:12	F 11	E 10:04	D 03	C 02	B 01	A 00	N11 N11	N NOT

Enables A-I enable their respective fields for comparison. The N11 and N bits are used to decode the 'compare not' function. The N11 bit only affects the F field (VA [11]), and the N bit affects the range of VA [31:12].

When N=1, normal comparisons are made. When N=0, the compare result is inverted; so, a 'hit' occurs when the addresses mismatch. The same control applies to N11. As an example, to enable the address bits VA [31:00], as listed in table 2.4.24, a value of 0x7FF is required in the virtual address breakpoint control register.

We also have the ability to breakpoint on the particular type of memory request being sent from the MMU to the MEMIF. This is sampled when a memory request is actually being issued (mm_issue_req = 1). This can be paired with two other fields indicating the type of tablewalk

occurring and the tablewalk level to match (if memory request indicates a tablewalk) as follows:

Table 4.26 - Memory Request Type

Memory Request	
NOP	No memory operation
RD64	Read of 64 bits (2 words)
RD128	Read of 128 bits (4 words)
RD256	Read of 256 bits (8 words)
WR8	Write of 8 bits (1 byte)
WR16	Write of 16 bits (2 bytes)
WR32	Write of 32 bits (1 word)
WR64	Write of 64 bits (2 words)
Tablewalk Type	
None	No tablewalk in progress
ic_tlb_tw	Tablewalk from instruction fetch
dc_tlb_tw	Tablewalk from data reference
io_tlb_tw	Tablewalk from DVMA
Tablewalk Level	
Root Level	
Level 1	
Level 2	
Level 3	

4.0.15.3 Additional Features

There are other features which can be used for microSPARC debug.

Some of these features are enabled using Processor Control Register bits. Software tablewalks can be enabled by asserting PCR[23], the STW bit. When in this mode the mmu will cause the instruction_access_MMU_miss and data_access_MMU_miss traps for instruction and data tablewalking respectively for tablewalks to be done by software.

The view modes are also very useful features for both debug and vector generation. There are three view modes: Address View, Data View, and Memory Data View which are enabled by PCR[22:20] respectively.


Address View mode is useful for non io testing allowing the Physical Address Register (PAR) to be viewed (1 cpu cycle later) on the SBus Address lines (bits [27:00] only).

Data View mode is useful for non io testing allowing the internal mc_mdata tristate bus to be viewed (1 cpu cycle later) on the SBus Data lines (bits [31:00]).

Memory Data View is useful for non memory sequences allowing the internal mc_mdata tristate bus to be viewed (1 cpu cycle later) on the Memory Data lines (bits [31:00]).

Alternate Cacheability is a diagnostic feature that allows the caches to be enabled by the IE and DE bits even with the mmu disabled. When not set, the caches are disabled when the mmu is disabled. This should not be used during boot mode accesses (or other instruction accesses to an sbus device). Specifically, having the mmu off, instruction cache on, alternate cacheability on and an sbus instruction access can cause indeterminate data to be put into the instruction cache. Instruction accesses work fine with alternate cacheability when the accesses are to main memory space.

5.0 Data Cache

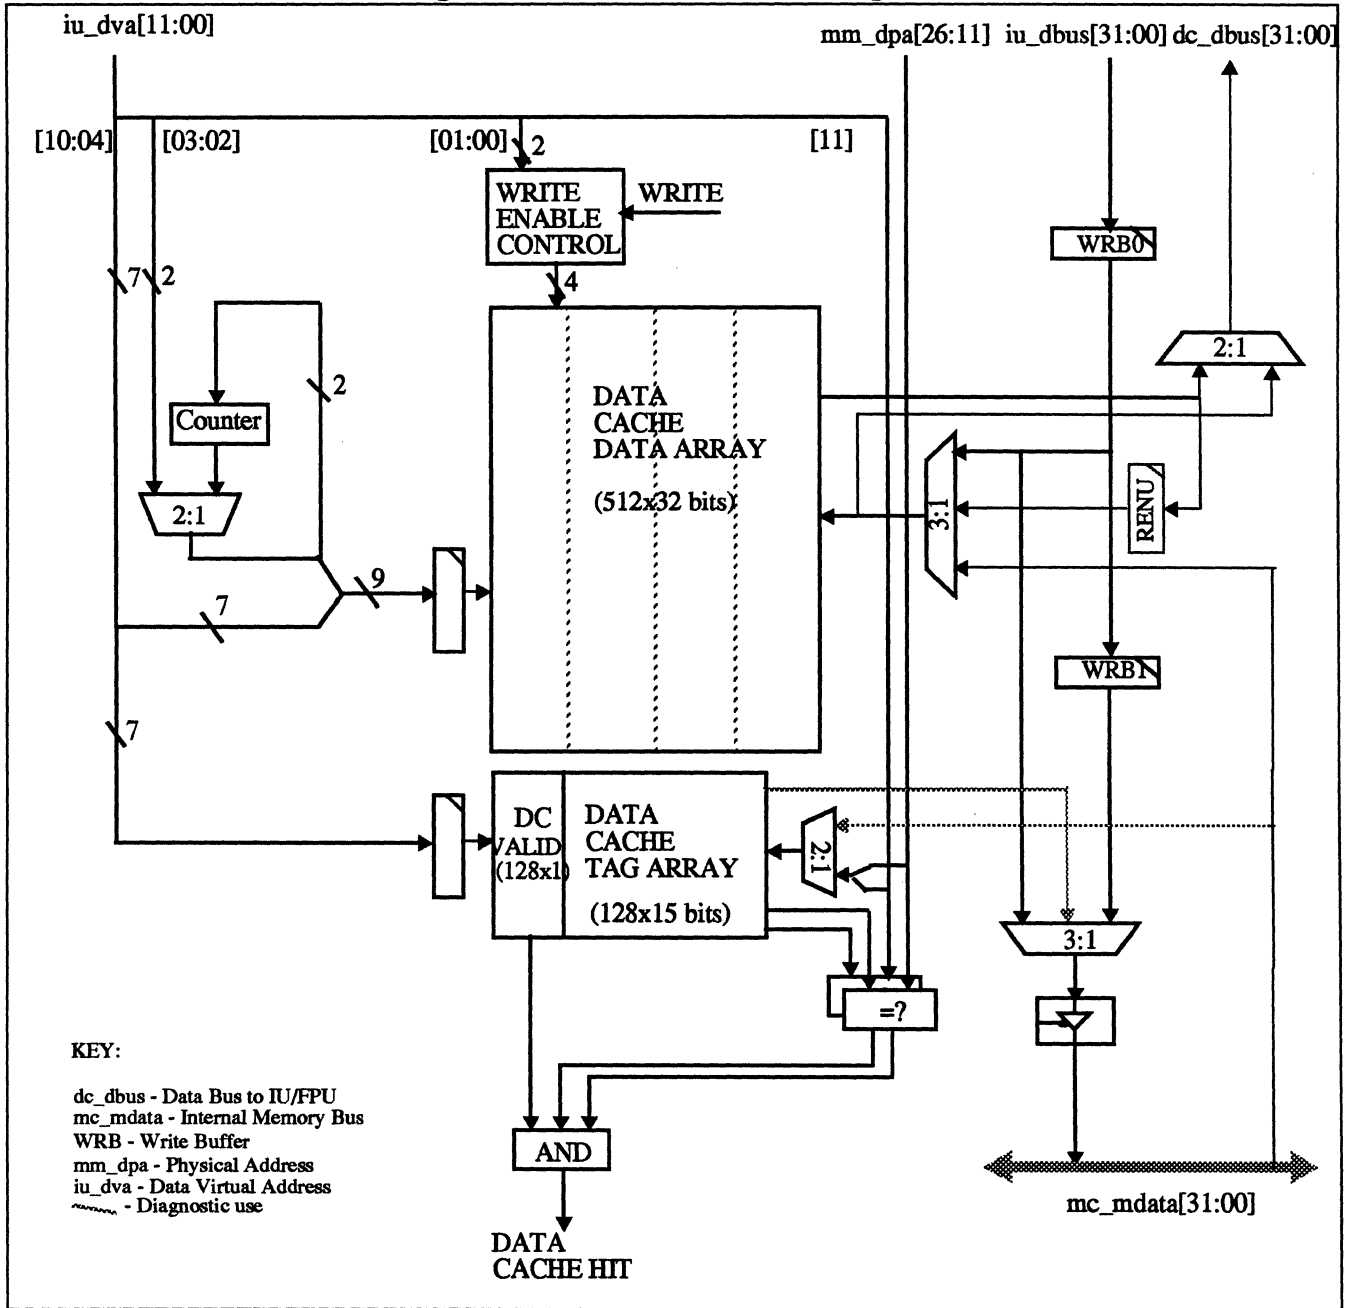


5.0.1 Overview

The microSPARC Data Cache is a 2K-Byte, direct mapped cache, used on load or store accesses from the CPU to cacheable pages of main memory. It is virtually addressed but physically tagged. Stores are write-through with no write allocate. The data cache is addressed by `iu_dva[10:0]`. The data cache is organized as 128 lines of 16 bytes of data. Each line has a cache tag store entry associated with it. On a data cache miss to a cacheable location, 16 bytes of data are written into the cache from main memory.

Within the data cache block there are also cache bypass paths. These paths are used for noncached load references, and for streaming data into the IU or FPU on cache miss. A simple block diagram follows.

Figure 5.1 - Data Cache Block Diagram



5.0.2 Data Cache Data Array

All IU write operations to cached locations write the data through to main memory, i.e. on a write hit, both the data cache and main memory are updated. There is, however, no write allocate, i.e. no cache fill is done on a write miss.

System software may read and write the data cache directly by executing load or store alternate space instructions, of any size, in ASI 0xF. Virtual address bits [10:0] will be used to address the data cache in this mode; all other virtual address bits are ignored during these operations.

There are three input sources to the data cache data array. The IU data_out bus (iu_dbus) is used when the data cache is updated on an integer or floating-point store operation. The internal memory data bus (mc_mdata) is used as input for fills on data cache misses. The RENU register is used in cancelling writes on stores which miss the cache.

5.0.3 Data Cache Tags

A data cache tag entry consists of several fields as follows.

Figure 5.2 - Data Cache Tag Entry

Reserved	PA Tag[26:11]	Reserved	Valid
31	27 26	11 10	01 00

Field Definitions:

Reserved (Rsvd) - Bits [31:27,10:01] are not implemented, should be written as 0 and will be read as 0.

Physical Address Tag - This field contains the physical address of the data held in the cache line. The Data Cache Controller writes this field from bits [26:11] of the physical address (mm_dpa) of the line.

Valid - This bit indicates that the line contains data. This bit is set when a cache line is filled due to a successful cache miss; a cache fill which results in a memory parity error will leave the Valid bit unset. An alternate address space data cache flash clear operation will clear the valid bits of all of the data cache tag entries.

There are two input sources to the data cache tag array. The Physical Address bits needed for the tag are used for cache updates due to data cache misses. The internal memory data bus (mc_mdata) is used as input for alternate store operations.

System software can read and write the data cache tags by executing word-length LDA and STA (Load and Store Alternate) instructions in ASI 0xE. The Virtual address bits [10:4] will select one of the 128 tags; all other address bits are ignored.

5.0.4 Write Buffers

The Write Buffers (WRB0,WRB1) are 32-bit registers in the data cache block used to hold data being stored from the IU or FPU to memory or other physical devices. On a store operation of a word or less, WRB0 holds the store data until it has been sent over the mc_mdata bus to the destination device. For halfword or byte stores, this data is left-shifted (with zero-fill) into proper byte alignment for writing to a word-addressed device before being loaded into WRB0. On a doubleword store the even word is first placed into WRB0. The next cycle the data from WRB0 is moved to WRB1 and WRB0 is loaded with the odd word. These registers can be read using a word-length LdA in ASI 0x39; for this operation, bit 8 of the Virtual address selects between the two registers (0 for WRB0, 1 for WRB1).

5.0.5 Data Cache Fill

The memory block size of data fetched from memory on data cache misses is 16 bytes. Memory will always return 16 bytes of data starting with the requested word first followed by the other word of the first doubleword and continuing with another doubleword (even word, then odd) which will wrap around a 16 byte boundary until the entire 16-byte block has been returned. The transfer rate is two words every three cycles from memory (two words of a doubleword, then a dead cycle). The Cache array is loaded the cycle that each word appears on the mc_mdata bus. The following table illustrates the fill operation showing the order that words are written into the cache:

Table 5.1 - Data Cache Fill Ordering

Requested Word	Order of fill (modulo 16B)
0	0, 1, dead cycle, 2, 3
1	1, 0, dead cycle, 2, 3
2	2, 3, dead cycle, 0, 1
3	3, 2, dead cycle, 0, 1

During cache fill, data is bypassed (or “streamed”) into the IU or FPU as it is written into the cache data array. For misses on word, halfword, or byte loads, the requested word is bypassed to the IU or FPU in the

same cycle that it appears on the mc_mdata bus; for LdD misses, each of the two requested words is bypassed to the IU or FPU in the same cycle that it appears on the mc_mdata bus.

5.0.6 Internal Memory Bus Interface

The data cache block interfaces to the internal memory bus (mc_mdata). Data from the data cache block to mc_mdata comes from either WRB0 or WRB1. WRB1 is used only for StD and ASI reads of WRB1. There are control signals from the MMU and Memory Controller to indicate when data is on mc_mdata to be loaded into the Data Cache and when data from WRB0 or WRB1 is to be put onto mc_mdata.

5.0.7 IU Data Bus Interface

The data cache block interfaces to an input and output IU data bus (iu_dbus and dc_dbus). Data to the IU or FPU is sourced from either the mc_mdata bus (for streamed data on data cache misses, and for non-cached loads) or the data cache (for data cache hits). Data from the IU or FPU on store operations is always loaded into WRB0.

5.0.8 RENU Register

In the event of a data cache miss on a Store instruction, the cache miss indication is not available until sometime into the cycle in which the store data is being written into the data array. This is too late to inhibit the write operation, so, to prevent the cache line from being corrupted by this write, we use the miss indication to MUX onto the cache array data-in bus a copy of the previous contents of the cache data array location being written. The previous contents of each stored-to location is captured in a special 32-bit register during the tag check access cycle which immediately precedes the write cycle of each store instruction. This register is known as the “REstore if Not Updated” (RENU) Register.

5.0.9 Data Cache Flushing

The data cache is implemented with a flash clear mechanism that is activated by any type of alternate store instruction to ASI 0x37. All data cache valid bits are reset (to zero) by this operation. Note that the data cache is not flushed by the FLUSH instruction (the instruction cache is).

5.0.10 Cacheability of Memory Accesses

Pages that are declared as non-cacheable (C=0 in the PTE) are not cached in the data cache. For data consistency and implementation reasons, the following operations are not cached.

Accesses when the MMU is disabled and alternate cacheability is disabled (EN, AC bits of the MMU CR=0).

Accesses while the data cache is disabled (DE bit of the MMU CR=0).

Accesses while using the MMU bypass ASI (ASI=0x20) and alternate cacheability is disabled (AC bits of the MMU CR=0).

Accesses while in Boot Mode.

Accesses to sources in physical address spaces 1-7.

Accesses by the MMU during tablewalks.

5.0.11 Diagnostic Strategy

Sublines and cache tags may be both read and written using ASI 0xF and 0xE respectively as previously discussed. The data cache will be structurally tested via the JTAG controller test ports. All register bits within the data cache and data cache tag are accessible via scan; on the chip level, all locations of these RAMs may be read or written by appropriate sequences of scan operations.

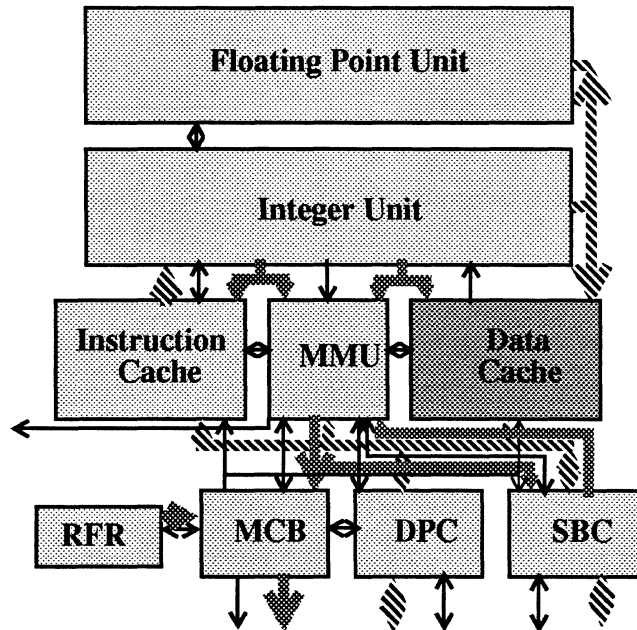
The internal Data Cache Registers may be read using ASI 0x39 and the Virtual Address to reference them. Single word accesses only should be used, others result in an internal error. The Virtual Address map to these registers:

Table 5.2 - Address Map for Data Cache Registers

VA[08]	Register
0	Write Buffer 0
1	Write Buffer 1

iu_dva bits [31:09,07:00] are ignored and should be set to zero by software.

6.0 Instruction Cache



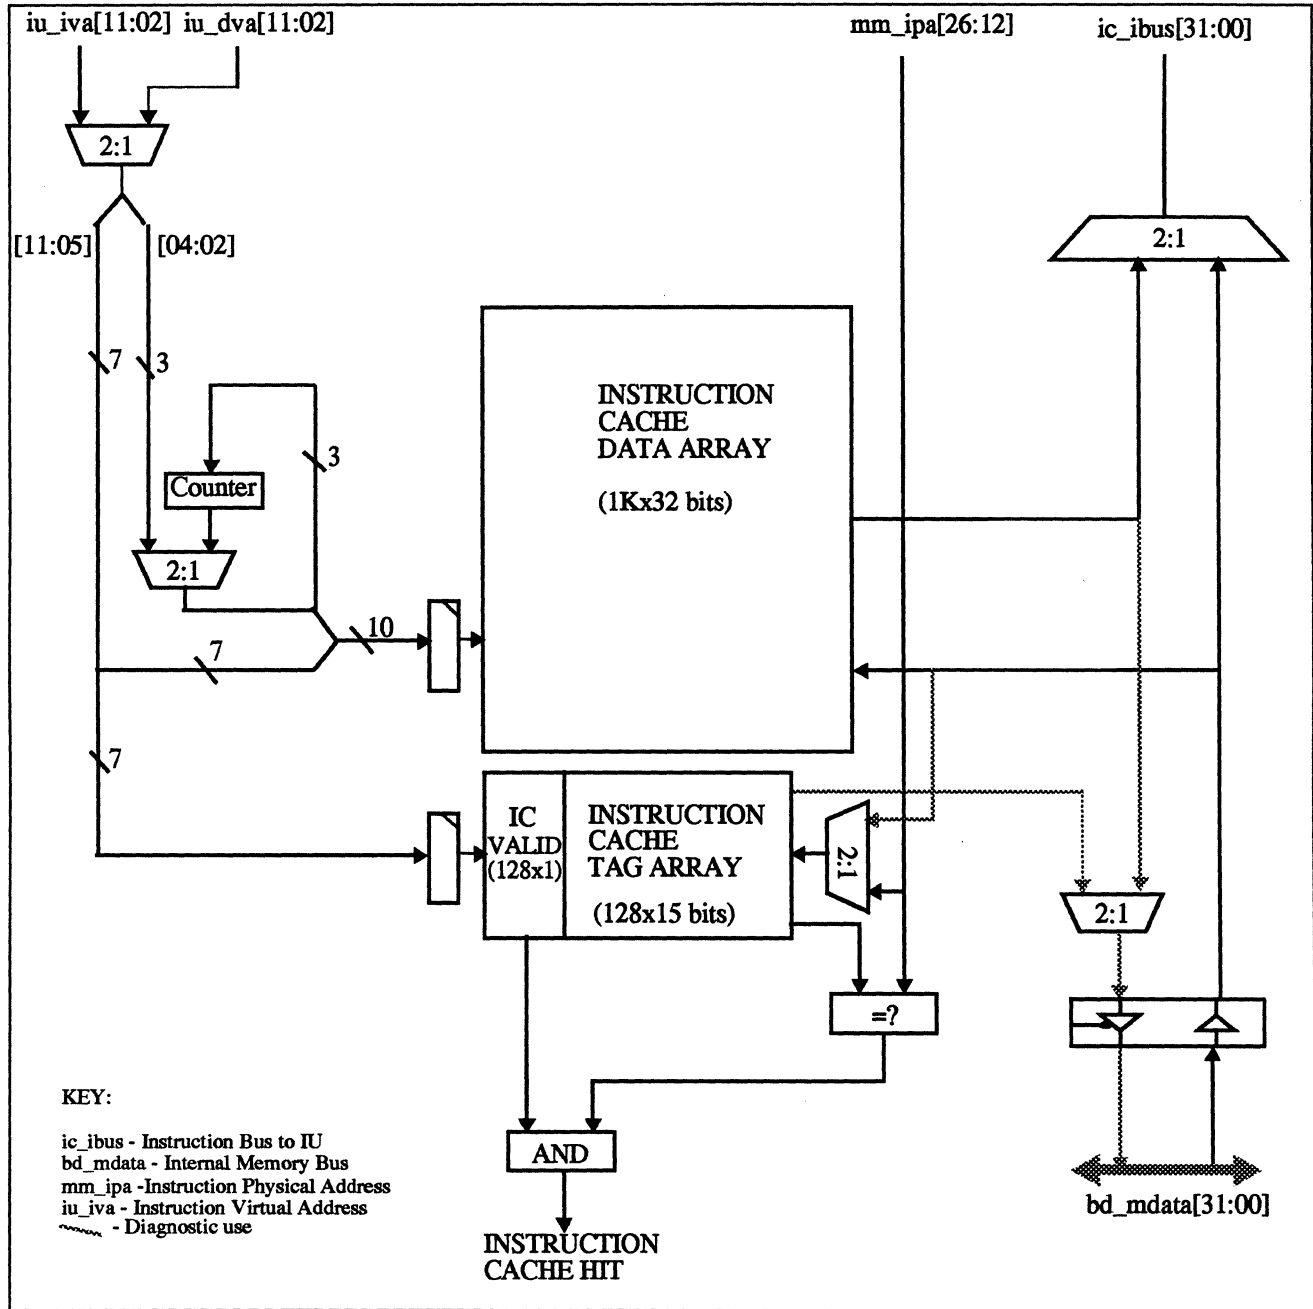
6.0.1 Overview

The microSPARC Instruction Cache is a 4K-Byte, direct mapped cache, used on instruction fetch accesses from the CPU to cacheable pages of main memory. It is virtually addressed but physically tagged. The instruction cache is normally addressed by `iu_iva[11:0]`. The instruction cache is organized as 128 lines of 32 bytes of data. Each line has a cache tag store entry associated with it. On a instruction cache miss to a cacheable location, 32 bytes of data are written into the cache from main memory.

Within the instruction cache block there are also cache bypass paths. These paths are used for noncached instruction fetches, and for

streaming instructions into the IU on cache miss. A simple block diagram follows.

Figure 6.1 - Instruction Cache Block Diagram



6.0.2 Instruction Cache Data Array

System software may read and write the instruction cache directly by executing load or store word alternate space instructions in ASI 0xD. Virtual address bits `iu_dva[11:2]` will be used to address the instruction cache in this mode; all other virtual address bits are ignored during these operations.

The internal memory data bus (`mc_mdata`) is used as input for fills on instruction cache misses, and as input for StA to ASI 0xD.

6.0.3 Instruction Cache Tags

A instruction cache tag entry consists of several fields as follows.

Figure 6.2 - Instruction Cache Tag Entry

Rsvd	IPA Tag[26:12]	Rsvd	Valid
31 27 26	12 11		01 00

Field Definitions:

Reserved (Rsvd) - Bits [31:27,11:01] are not implemented, should be written as 0 and will be read as 0.

Physical Address Tag - This field contains the physical address of the data held in the cache line. The Instruction Cache Controller writes this field from bits [26:12] of the physical address (`mm_ipa`) of the line.

Valid - This bit indicates that the line contains data. This bit is set when a cache line is filled due to a successful cache miss; a cache fill which results in a memory parity error will leave the Valid bit unset. An alternate address space instruction cache flash clear operation will clear the valid bits of all of the instruction cache tag entries. A Flush instruction will clear the valid bit of the single line which is addressed by `iu_dva[11:05]` (regardless of the contents of that line).

There are two input sources to the instruction cache tag array. The Physical Address bits needed for the tag are used for cache updates due to instruction cache misses. The internal memory instruction bus (`mc_mdata`) is used as input for alternate store operations.

System software can read and write the instruction cache tags by executing word-length LDA and STA (Load and Store Alternate)

6.0.4 Instruction Cache Fill

instructions in ASI 0xC.; dva bits [11:5] will select one of the 128 tags; all other address bits are ignored.

The memory block size of data fetched from memory on instruction cache misses is 32 bytes. Memory will always return 32 bytes of data, starting with the requested word first followed by the other word of the first doubleword and continuing with the three remaining doublewords (even word, then odd) which will wrap around a 32 byte boundary until the entire 32-byte block has been returned. The transfer rate is two words every three cycles from memory (two words of a doubleword, then a dead cycle). The Cache array is written during the cycle that each word appears on the mc_mdata bus. The following table illustrates the fill operation showing the order that words are written into the cache; 'D' represents a dead cycle in which no word is written:

Table 6.1 - Instruction Cache Fill Ordering

Requested Word	Order of fill
0	0, 1, D, 2, 3, D, 4, 5, D, 6, 7
1	1, 0, D, 2, 3, D, 4, 5, D, 6, 7
2	2, 3, D, 4, 5, D, 6, 7, D, 0, 1
3	3, 2, D, 4, 5, D, 6, 7, D, 0, 1
4	4, 5, D, 6, 7, D, 0, 1, D, 2, 3
5	5, 4, D, 6, 7, D, 0, 1, D, 2, 3
6	6, 7, D, 0, 1, D, 2, 3, D, 4, 5
7	7, 6, D, 0, 1, D, 2, 3, D, 4, 5

During an instruction cache fill, instructions from the missing line can be supplied to the IU or FPU by two separate mechanisms; these mechanisms are collectively called "streaming". In the first type of streaming ("bypass streaming"), instructions are bypassed around the cache data array to the IU/FPU in the same cycle that the array is being written - this can occur in all cycles of the fill sequence except the three dead cycles. The second form of streaming ("dead-cycle streaming") occurs only during the three dead cycles; any instruction word which has already been written into the RAM array can be accessed by reading the array. In a given cycle, the IU is only able to accept the instruction word which it is requesting; in some cycles, the IU may not be requesting any instruction at all, due to interlocks, multi-cycle instructions, or pipeline holds. If, in a given cycle, the IU is requesting a word which is available via streaming, then that word is supplied to the IU and the pipeline can advance.

6.0.5 Internal Memory Bus Interface

The instruction cache block interfaces to the internal memory bus (mc_mdata). Data for LdA from ASI 0x0d is driven onto mc_mdata by the Instruction Cache, under control of an enable signal from the MMU.

6.0.6 IU Instruction Bus Interface

The instruction cache block drives the IU instruction bus (ic_ibus). Instructions to the IU or FPU are sourced from either the mc_mdata bus (for bypass-streamed instructions on instruction cache misses, and for non-cached instruction fetches) or the instruction cache data array (for instruction cache hits, and for dead-cycle streamed instructions on instruction cache misses).

6.0.7 Instruction Cache Flushing

The instruction cache is implemented with a flash clear mechanism that is activated by any type of alternate store instruction to ASI 0x36. All instruction cache valid bits are reset (to zero) by this operation. Also, the FLUSH instruction always clears the single valid bit that is addressed by iu_dva[11:05], regardless of the contents of this tag entry.

6.0.8 Cacheability of Memory Accesses

Pages that are declared as non-cacheable (C=0 in the PTE) are not cached in the instruction cache. For data consistency and implementation reasons, the following instruction fetch operations are not cached.

Accesses when the MMU is disabled and alternate cacheability is disabled (EN, AC bits of the MMU CR=0).

Accesses while the instruction cache is disabled (IE bit of the MMU CR=0).


Accesses while in Boot Mode.

Accesses to sources in physical address spaces 1-7.

6.0.9 Diagnostic Strategy

Sublines and cache tags may be both read and written using ASI 0xD and 0xC respectively as previously discussed. The instruction cache will be structurally tested via the JTAG controller test ports. All register bits within the instruction cache and instruction cache tag are accessible via scan; on the chip level, all locations of these RAMs may be read or written by appropriate sequences of scan operations.

7.0 Memory Interface



7.0.1 Overview

The microSPARC architecture allocates 256MB of space for the system memory (Physical address space '0', defined by `mm_pa[30:28]`), while the actual memory interface and the memory management unit can only support up to 128MB.

The following sections describe the general memory layout for a microSPARC-based system and then explains each of the logical blocks within the Memory Interface block.

The microSPARC Memory Interface block is logically divided into three subsections, the Memory Control Block (MCB), The Data aligner and Parity check/generation logic (DPC) and the Ram Refresh control (RFR).

7.0.2 Memory Subsystem

microSPARC Memory Interface is designed to primarily satisfy the basic system requirements, while providing sufficient capabilities to support future expansion.

The interface is designed with the following criteria in mind:

- 64 bit Data bus to increase memory bandwidth.
- 1 bit parity per word (32 bits) for reduced cost.
- Memory divided into blocks which can support different density devices. This will allow relatively small memory increments with a small number of blocks.
- Allow for future higher memory requirements by supporting next generation of DRAM devices.

Typically a carefully laid out system board using the microSPARC chip would require 60ns DRAMs at 50MHz and 80ns DRAMs at 40MHz clock speeds. The designer however, should use the memory interface AC specifications in the microSPARC datasheet, to select the appropriate DRAM speed for a specific system and clock speed.

7.0.2.1 Memory Organization

microSPARC architecture defines a 28-bit physical address space for memory (PAS 0). This means a 256MB block for system DRAM. Electrically however, microSPARC uses only 27 bits of this address space, limiting the maximum memory for a microSPARC-based system to 128MB.

This 128MB is divided into 4 banks, each capable of addressing up to 32MB. The banks are defined as follows:

- Each bank is selected by a separate RAS line. There are a total of 4 RASes for DRAM banks (`c_mc_ras_1[3:0]`).
- The banks have a 64bit data path to microSPARC.
- All the banks use the same 2-bit CAS lines (`c_mc_cas_1[1:0]`), to select the upper or lower 32 bits (high or low word).
- All the banks use the same write signal (`c_mc_mwe_1`).
- All the banks use the same 22-bit multiplexed Row/Column address bus. At the time of finalizing the microSPARC memory interface, DRAM manufacturers were proposing 2 addressing schemes for 4Mx4 devices, an 11-row/11-column and a 12-row/10-column. MicroSPARC's memory interface will support DRAMs with an 11x11 matrix and DRAMs with a 12x10 matrix.

The memory interface is designed with the 4bit wide DRAM devices in mind. Using 16 such devices (or 2 SIMMs with eight devices on each) will provide the required 64bit wide data bus. In addition, each bank will require two 1bit wide devices of the same depth (If using SIMMs, one on each SIMM) to store the 2 parity bits.

Hence, each bank can be populated using one of the following configurations:

- 2MB (256Kx64) of data, using 16 of 256Kx4 devices for data and 2 of 256Kx1 for parity, or using 2 of 256Kx33 SIMMs.
- 8MB (1Mx64) of data, using 16 of 1Mx4 devices for data and 2 of 1MKx1 for parity, or using 2 of 1Mx33 SIMMs.
- 32MB (4Mx64) of data, using 16 of 4Mx4 devices for data and 2 of 4Mx1 for parity, or using 2 of 4Mx33 SIMMs.

Note that a pair of double-density (e.g. 512Kx33 or 16Mx33) SIMMs will occupy 2 banks (Need 2 RASes).

7.0.2.2 Access to Unused or Unpopulated Memory regions

Any access to a location in the upper 128MB will be mirrored to its corresponding location in the lower 128MB and no errors will be generated.

Similarly, if a bank contains less than the defined maximum of 32MB, the real memory will be mirrored on the higher unused portions and an access from any of the unused sections will be mirrored to the corresponding location in the lowest block and no errors will be generated. For example, if a bank contains 2MB of real memory, this will be mirrored on the remaining 15 empty portions.

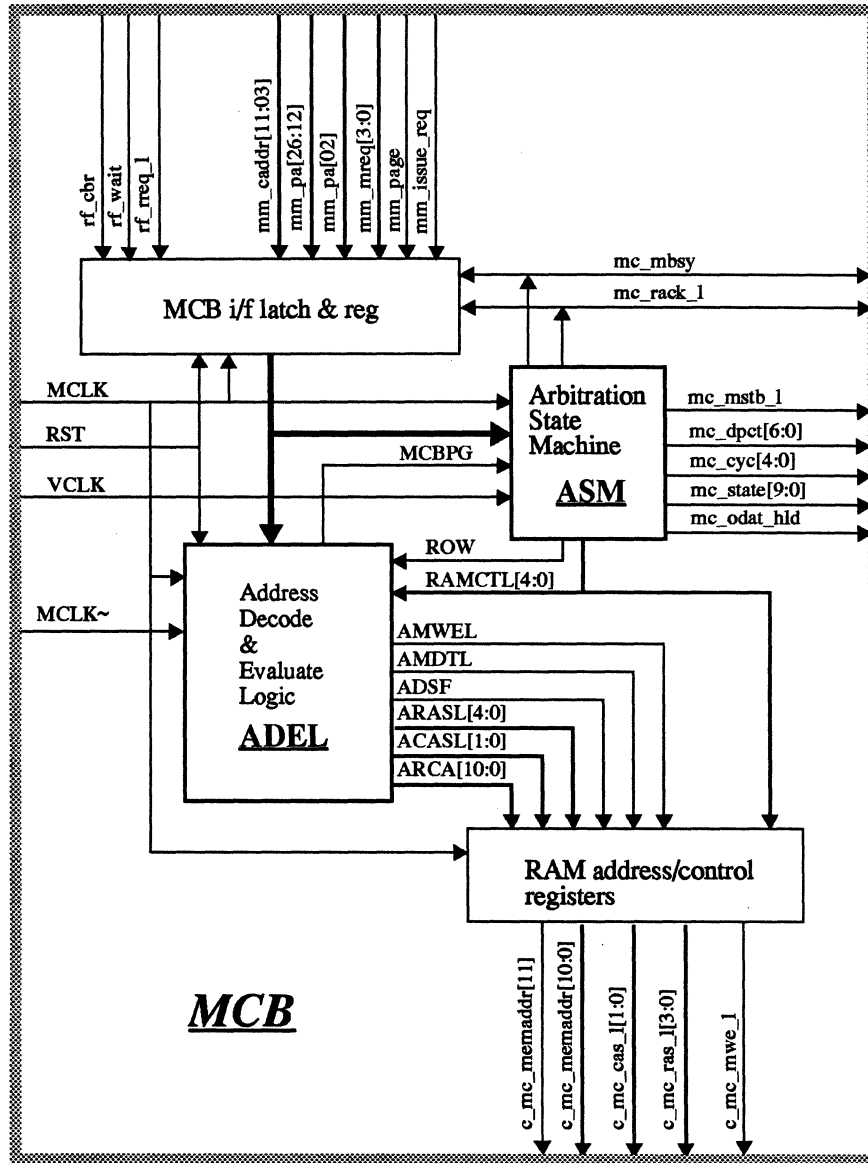
However, an access from a fully unused (empty) bank will complete, but it's result will be unknown and may cause a parity error.

7.0.3 Memory Control Block (MCB)

The operations that occur on the memory bus are data reads, writes, and read-modified-writes required for cpu execution, instruction fetches and prefetches, translation buffer accesses during table walks, reads and writes by IO devices, and all RAM refresh. The Memory Control Block (MCB) keeps track of the priorities of memory operations and completely controls the DRAM based main memory.

As shown in the following diagram, MCB contains 2 major logic blocks, labeled "ASM" and "ADEL" which perform the memory arbitration and address mapping functions respectively. This blocks will be described in the following subsections. MCB also includes some input and output register blocks, which provide the synchronization among input and output signals.

Figure 7.1 - MCB block diagram.



7.0.3.1 Arbitration State Machine (ASM)

ASM is responsible for detecting the requests from MMU and Refresh blocks, arbitrate between them if necessary and grant the appropriate request. Once a request is granted, the MCB will carry out the requested memory operation which will consist of one or more memory cycles. The following table lists all the types of memory operations performed by the MCB, the possible request sources and the type and number of cycles involved.

Table 7.1 - Memory operations performed by MCB

Operation	Source	Memory Cycles produced
<i>d.rd.32b</i>	MMU. Used to fill one line of I-cache (Inst-Fetch).	32 bytes are read from DRAM in a single operation, using 4 longword (64bit) read cycles. The first read is paged or non-paged, from the address given on PA. The following 3 reads are paged. ADEL will supply the address for the next 3 reads, incrementing or wrapping it as necessary, in order to read a 32 byte aligned block and fill a whole I-cache line.
<i>d.rd.16b</i>	MMU. Used to fill one line of D-cache or do an SBus burst read of 16bytes.	16 bytes are read from DRAM in a single operation, using 2 longword (64bit) read cycles. First read is a paged or non-paged cycle, using the address supplied on PA. The next cycle is a paged read, where ADEL will increment or wrap the address in order to read a 16byte aligned block from memory.
<i>d.rd.8b</i>	MMU. Used for IU and SBus longword reads.	8 bytes are read from DRAM, using a paged or non-paged longword read from the address supplied by PA.
<i>d.wr.8b</i>	MMU. Used for IU and SBus longword writes.	8 bytes are written to DRAM, using a paged or non-paged longword write to the address supplied by PA.
<i>d.wr.4b</i>	MMU. Used for IU and SBus word writes.	4 bytes are written to DRAM, using a paged or non-paged word write to the address supplied by PA.
<i>d.rmw.2b</i>	MMU. Used for IU and SBus halfword writes.	a halfword (16bit) write to DRAM in a single operation, using a paged or non-paged word read followed by a paged word write, using the same address supplied by PA. MCB will perform the read and write cycles and will instruct DPC to latch the 16bit write-data from the source, insert it in the appropriate halfword of the word read from memory and then gate it back on memory data-bus as the write data.
<i>d.rmw.b</i>	MMU. Used for IU and SBus byte writes.	a byte (8bit) write to DRAM in a single operation, using a paged or non-paged word read followed by a paged word write, using the same address supplied by PA. MCB will perform the read and write cycles and will instruct DPC to latch the 8bit write-data from the source, insert it in the appropriate byte of the word read from memory and then gate it back on memory data-bus as the write data.
<i>cbr.ref</i>	RFR. Used to do a refresh cycle on all DRAM/VRAM.	Will force a Cas-before-Ras refresh cycle to be performed on all DRAM and VRAM banks.

7.0.3.2 Arbitration for Memory Access and ASM Priority Scheme

ASM arbitration scheme is based on the following rules:

- All requests are checked at the end of each operation (for multi cycle operations, this means the end of last memory cycle) and:

- If: no requests are pending, ASM will enter the idle state and will remain there until a request is detected.
- If: only one request is pending, it will be granted and the operation will begin.
- If: More than one request is pending, the one with the highest priority will be granted and the operation will begin. The priorities are as follows:
 - g) MMU is the highest priority, except when the current cycle is also an MMU request, in which case it will be considered the lowest priority. This is to prevent bus locking as a result of back to back MMU requests.
 - h) RFR has the lowest priority, except when the current cycle is an MMU request, in which case it will have higher priority.
- If: While in idle, an RFR request is detected, the state machine will advance to a "Check" state, where it'll look to see if an MMU request occurred just as RFR request was accepted. If there are no MMU requests, ASM will continue to acknowledge the RFR request and do the cycle, else, it will do the MMU cycle.

7.0.3.3 Memory Operation Timing


Following pages contain the waveform diagrams for some of the memory operations requested by MMU and carried out by the memory interface. Each operation-type is defined using the operation-name given in table-1 of this section.

The diagrams are functional and do not represent actual delays. Synchronous signals are clocked with the positive edge of the MCLK (derived from system clock, running at same frequency and assumed to have negligible skew) and are shown to be valid about half a clock period later. In case of the falling edge of the RAS signals only, the transition occurs after the negative-edge of the system-clock.

In addition, the `mm_mreq[3:0]` is shown valid for 1 clock periods. This indicates the clock-cycle during which MMU asserts the `mm_issue_req` signal.

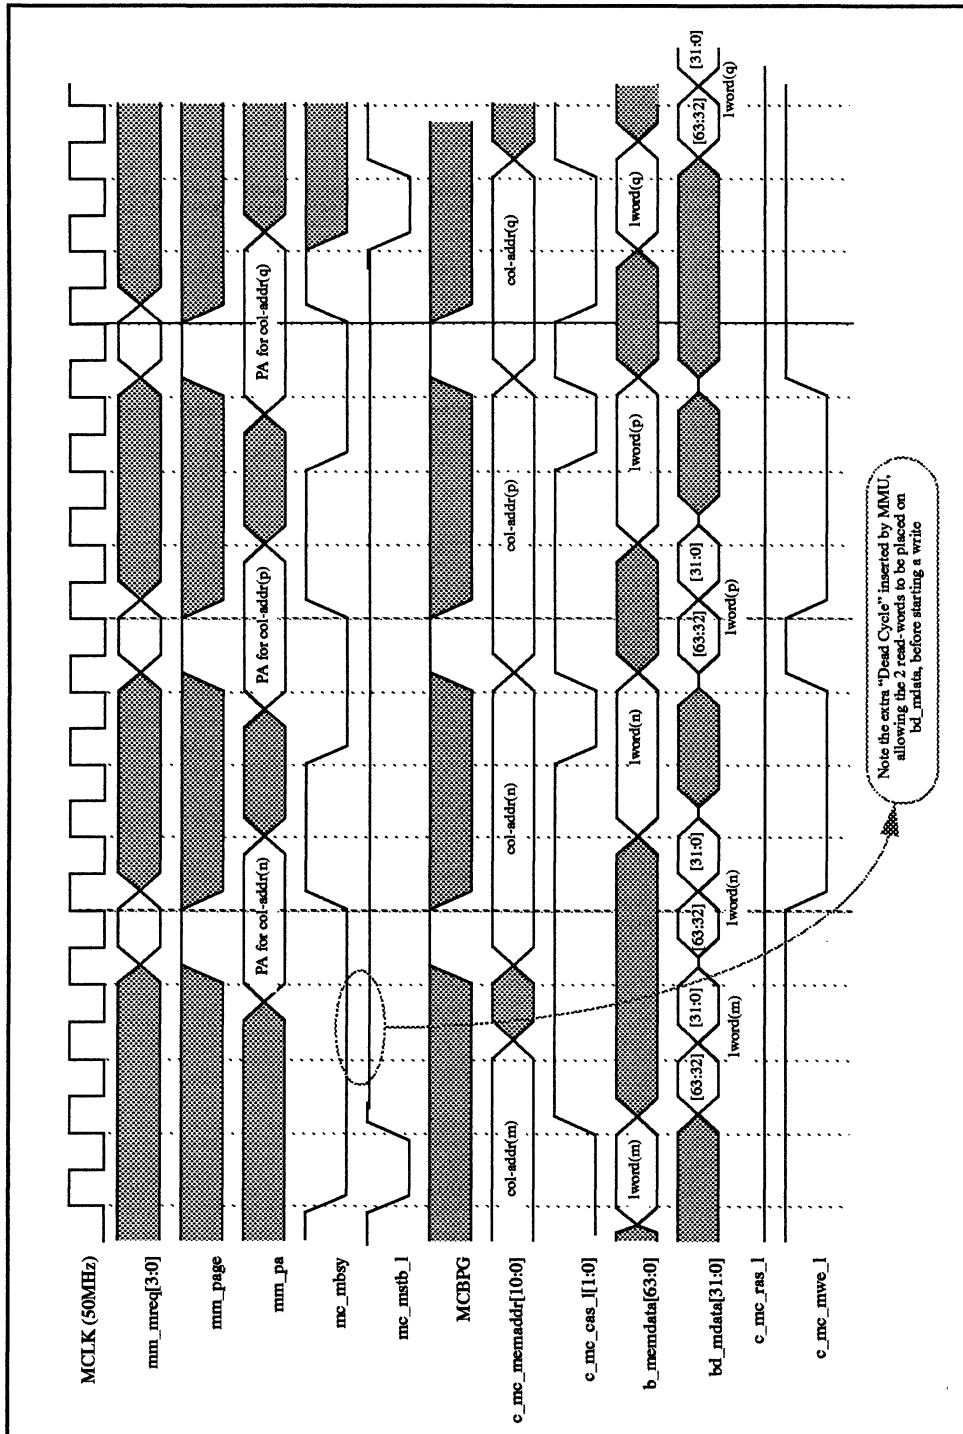
The waveforms are provided only as a general reference and do not reflect details such as the word/hword/byte order relative to the address or the MMU request type etc.

Figure 7.2 - MMU I-fetch beginning in page-mode



MMU I-fetch beginning in page-mode, followed by another MMU read request from same DRAM-page.

Figure 7.3 - MMU page-mode write after a read



MMU page-mode write after a read, followed by another write, followed by another read, all from same page.

Figure 7.4 - Non-paged write cycle, shown following a read




Figure 7.5 - Non-paged read cycle, shown following a read


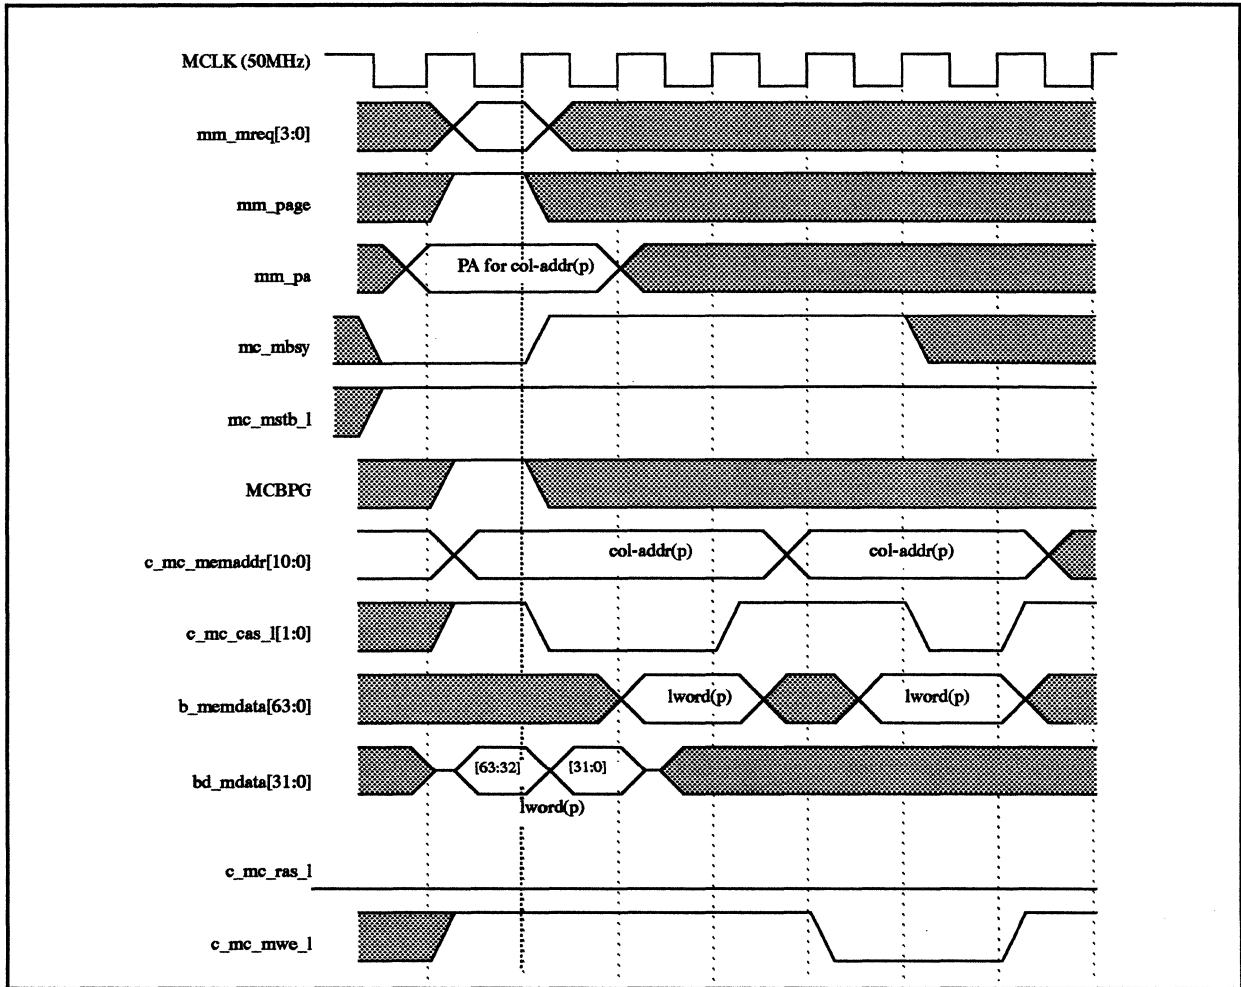



Figure 7.6 - Paged Byte/Halfword (8/16 bit) write cycle.



Paged Byte/Halfword (8/16 bit) write cycle, generating a hardware controlled Read-Modify-Write sequence.

Figure 7.7 - Non-paged Byte/Halfword (8/16 bit) write cycle.



Non-paged Byte/Halfword (8/16 bit) write cycle, generating a hardware controlled Read-Modify-Write sequence.

7.0.3.4 Address Decode & Evaluate Logic (ADEL)

This block primarily monitors the address and function-select signals coming from MMU and RFR and performs the necessary decode and re-mapping of the memory address and control lines. Based on commands received from ASM, ADEL gates the row/column address and memory control signals required for the current operation out to memory.

The mapping of system memory is discussed in the following section.

7.0.3.5 Address Mapping For System DRAM

From the 31 bits of the physical address bus driven by MMU block (`mm_pa[30:0]`), the three MSBs (`mm_pa[30:28]`) represent 1 of the 8 physical address spaces (PAS) as defined in microSPARC architecture. From these, only PAS0 is of concern to MCB, since an MMU request from MCB will only be made if an access to system memory is required. Hence ADEL ignores the `mm_pa[30:28]` bits.

When a memory cycle request is detected, ADEL uses the `mm_pa[26:02]` address bits to complete its decode. The following table describes the decode scheme used for system memory.

A maximum of 512 memory cycles can be made from a contiguous block, while remaining within a DRAM page. This gives a maximum of 4K (512x64) block size which can theoretically be accessed using page mode cycles only.

A point to note from the table below, are the staggered decoding of `mm_pa[24:21]` for `c_mc_memaddr[10:9]`. This was necessary in order to allow different size devices (256Kx4, 1Mx4 and 4Mx4) to be used while maintaining the largest common contiguous block, which is dictated by the least dense device.

Also, as shown in the table, `mm_pa[23]` is used as both `c_mc_memaddr[10]` for column address and `c_mc_memaddr[11]` for

row address. This is to cater for the 2 different 4Mx4 DRAM architectures, 11x11 matrix and 12x10 matrix.

Table 7.2 - Physical Address decode for System Memory

PA	Decode
30-27	Not Used. System memory limit is 128MB
26-25	Decode to select 1 of 4 RASes: 00 RASL0 1st 32MB bank. 01 RASL1 2nd 32MB bank. 10 RASL2 3rd 32MB bank. 11 RASL3 4th 32MB bank.
24	Decoded as row address bit 10 (c_mc_memaddr10). Required for 16MBit DRAMs.
23	Decoded as column address bit 10 (c_mc_memaddr10) and row address bit 11 (c_mc_memaddr11). Required for 16MBit DRAMs. See text for more information.
22	Decoded as row address bit 9 (c_mc_memaddr9). Required for 4MBit DRAMs.
21	Decoded as column address bit 9 (c_mc_memaddr9). Required for 4MBit DRAMs and up.
20-12	Decoded as row address bits 8 to 0 (c_mc_memaddr[8:0]). Required for 1MBit DRAMs and up.
11-3	Decoded as column address bits 8 to 0 (c_mc_memaddr[8:0]). Required for 1MBit DRAMs and up.
2	Decoded to select one of 2 CASes: 0 CASL0 Lower data word (bd_mdata[31:0]) 1 CASL1 Higher data word (bd_mdata[63:32])
1-0	Not used for external decode. Byte and halfword writes are achieved by MCB and DPC doing a read, update, write sequence. This bits are used then, to select the appropriate data fields.

7.0.4 Data aligner and Parity Check/generate logic (DPC)

DPC is responsible for transferring data between external memory data bus and the internal data path as well as generating and checking of parity for system main memory (DRAM).

During any read, write or hardware controlled read-modify-write cycle, DPC will perform the necessary data alignment and byte/halfword placement. It will also provide temporary storage for hardware controlled read-modify-write cycles, resulting from byte/halfword write cycles to memory.

DPC also contains the parity generation and checking logic. The parity is composed of 1 bit per word (32 bits) and is used for system DRAM only.

Type of parity operation for the system DRAM is determined by the state of the Parity Control bit (PC) in the Processor Control Register as described in the following table:

Table 7.3 - Parity Control Definition

PC	Description
0	Check/Generate even Parity.
1	Check/Generate odd Parity.

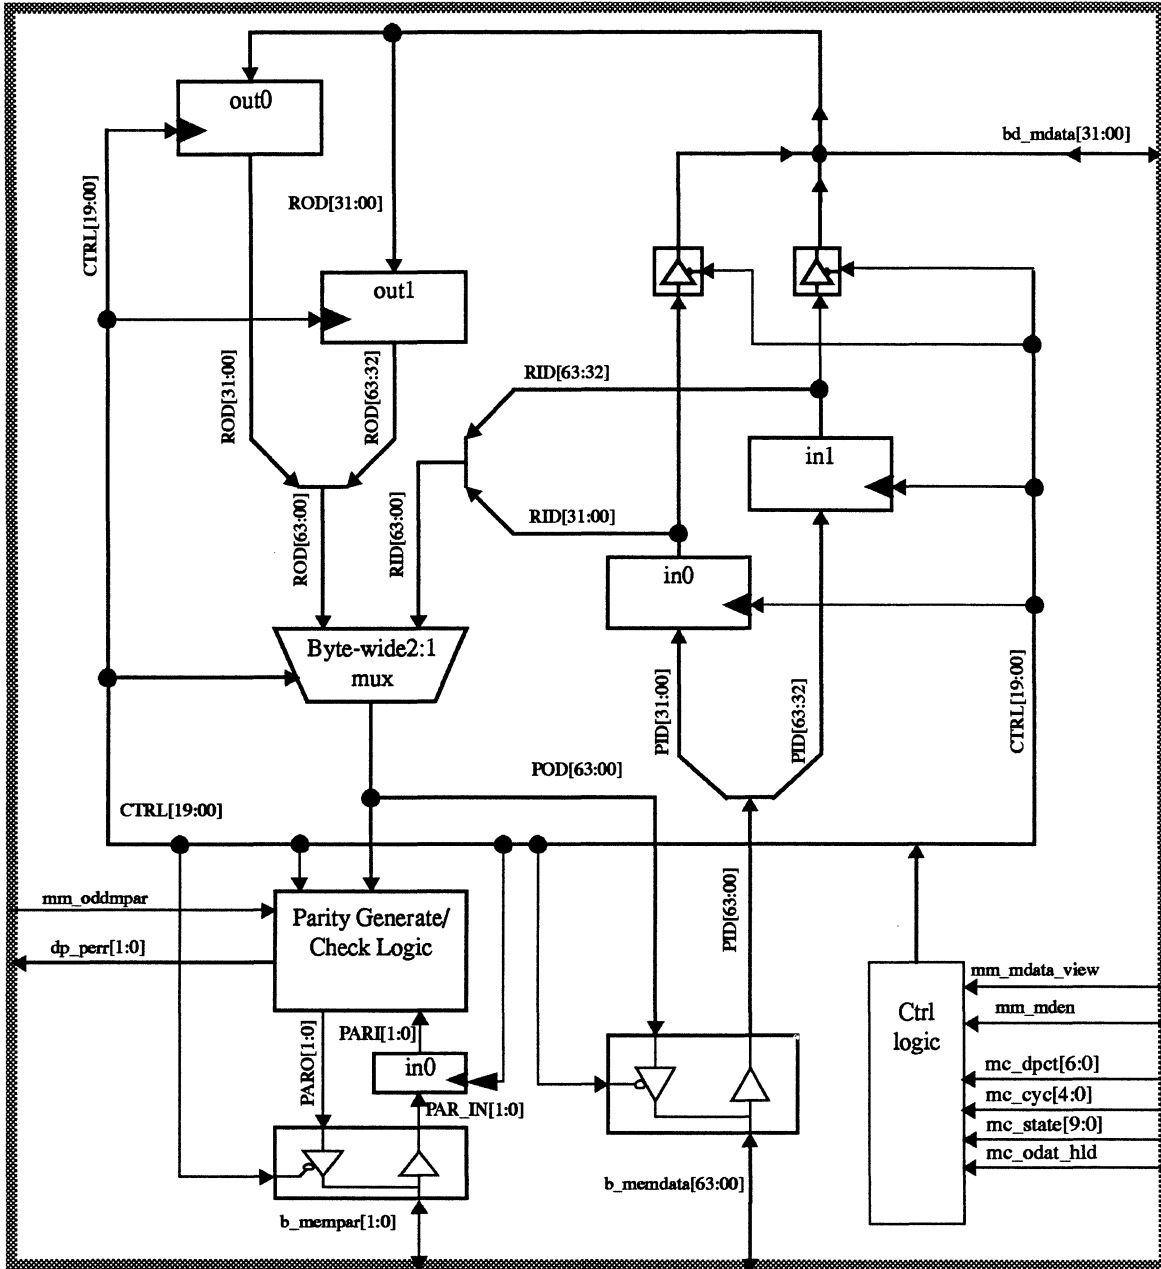
Since system parity is 1-bit per word, any byte or halfword store operation, will result in a hardware controlled read-modify-write cycle. During the read part of such operation, the word parity will be checked and if an error is detected, a parity error will be generated. After the word has been updated to contain the new byte/halfword, a write operation will be performed, which will also update the parity.

The flow of data and type of operations performed by DPC are governed by the Memory Control Block and the commands it receives from MCB.

DPC block diagram, given below, shows the basic data paths connecting the 64bit external memory bus (b_memdata[63:00]) to the 32bit internal one (bd_mdata[31:00]). The parity check/generation logic is shown to be on the output path, but for input data, parity is checked after it is clocked into the registers and gated through the alignment mux.

The alignment mux is also used to combine and produce the output data during a read modify write sequence. The complexity of this mux is reduced by having the byte or forward data which is to be written to memory, already in the correct position. This is done by the block sourcing the data on c_dp_mdata (D or I cache, IU, SBus controller).


Figure 7.8 - Datapath and Parity Control (DPC) block diagram



7.0.5 RAM Refresh Control (RFR)

The refresh control logic (RFR) is a simple request generator, asserting a request to MCB at fixed intervals. MCB will service this low priority request by performing a Cas-before-Ras type refresh cycle on all system RAM.

Figure 7.9 - RAM Refresh Control block diagram.



RFR refresh rate can be selected by programming 2 bits of the Processor Control Register according to the following table. These bits are then passed to RFR as mm_rf_cntl[1:0] input bits, which controls the rf_req_1 rate.

Table 7.4 - Refresh Rate Control bits.

mm_rf_cntl [1:0]	Refresh interval
00	Assert a refresh request once every 128 MCLK periods. With this setting, adequate refresh is guaranteed for MCLK values of down to 8.6MHz. This is the default after power up.
01	No Refresh!
10	Assert a refresh request once every 512 MCLK periods. With this setting, adequate refresh is guaranteed for MCLK values of down to 35MHz.
11	Assert a refresh request once every 768 MCLK periods. With this setting, adequate refresh is guaranteed for MCLK values of down to 52MHz.


The RFR is also responsible for initializing the DRAMs on power-up.

After power-up and before they can be reliably used, DRAMs require a 200us "Wait" period followed by 8 Cas-before-Ras refresh cycles.

For systems built around microSPARC, the reset must remain active for at least 200us after power-up, to satisfy the "Wait" period. In systems using the NCR 89C105 chip, the reset is supplied by the NCR 89C105 chip. On power-up the NCR 89C105 chip guarantees an active reset duration of ~200ms and for subsequent software initiated resets it will force reset active for ~1024 SBus clocks (~50us).

After an active reset, the "mm_rf_cntl" bits which reside in the MMU's PCR register are set to "00" (See table 2.7.4), setting RFR to generate a refresh request every 128 clocks. In addition, RFR itself, asserts its "rf_cbr" and "rf_rreq_l" signals, forcing MCB to enter a "cbr" state, where it will perform 8 CbR refresh cycles, completing the DRAM initialization cycle. After that, RFR will negate both "rf_cbr" and "rf_rreq_l" signals, allowing MCB to proceed to it's normal operation state.

8.0 SBus Controller



8.0.1 Overview

The SBus Controller (SBC) refers to the I/O subsystem that handles input and output between local resources, including the CPU, system memory, and control space, and all external system resources. The SBC is implemented as an SBus in accordance with the SBus Specification Rev. A.2. The SBC supports:

- Programmed Input/Output (PIO) transactions between the CPU and SBus slave devices
- Direct Virtual Memory Access (DVMA) transactions between SBus masters and local resources. (referred to as local DVMA)
- Direct Virtual Memory Access (DVMA) transactions between SBus masters and other SBus slave devices. (referred to as bypass DVMA)

Standard SBus features, such as dynamic bus sizing, reruns, atomic transactions, bus arbitration, burst transfers (up to 16 bytes), watchdog timer, and error reporting are fully supported. Interrupts and SBus Reset are not implemented in the SBC; these functions are handled elsewhere.

The SBC plays many SBus roles. It serves as an SBus controller by arbitrating bus requests, translating virtual addresses, enabling slave cycles, etc. In addition, the SBC may act as either an SBus master or an

SBus slave. For PIO transactions, the SBC acts as an SBus master. For DVMA transactions, the SBC can act as either a slave or have no role at all, depending on the target of the DVMA transaction as indicated by the physical address. For local DVMA, the SBC has a role as both a bus controller and a slave device. For bypass DVMA, the SBC has a role as a bus controller only, not as a slave.

PIO transactions consist of an SBus slave cycle only; the address translation is done in advance of the bus acquisition.

PIO transactions occur when the CPU executes loads or stores to I/O (SBus) space. In the case of a PIO write transaction, the write is posted. Processing in the CPU continues while the SBus transaction completes in the SBC. A stall will occur only if another PIO transaction is attempted before the previous PIO write transaction completes. In the case of a PIO read transaction, processing is always stalled until the data becomes valid at the end of the SBus transaction.

DVMA transactions occur when an SBus master has acquired the bus in order to execute a transaction to a slave. A DVMA transaction consists of an address translation cycle and a slave cycle. The target of the slave cycle is determined once the translation cycle completes. The slave target can be either a local resource, defined as locations in either system memory or system control space, or another SBus device.

During the address translation cycle, the SBC obtains a virtual address from the DVMA master and is submitted to the MMU for translation. The MMU returns a physical address. The type of DVMA slave cycle, either local or bypass, is determined from the physical address.

A significant distinction concerning memory data transfers is that since system memory is a local resource, it is necessary for memory data to pass through the SBC; "fly-by" memory data transfers are inappropriate. Local DVMA slave cycles have two distinct, sequential operations in the SBC: a data get followed by a data put operation. A data get operation loads up to 16 bytes of data into an internal data store. A data put operation transfers the data from the internal data store to a destination. The data get operation can either be an internal data transfer

or an SBus cycle, depending on the read/write direction; the data put operation will be the other.

Figure 8.1 - Data Get and Data Put

	Data Get	Data Put
Read	Internal Data Transfer	Sbus Cycle
Write	Sbus Cycle	Internal Data Transfer

For local DVMA slave read cycles, an internal data transfer occurs during the data get stage, and an SBus slave cycle occurs during the data put stage. In this case, the data get operation shows up as a pause between the SBus translation cycle and the SBus slave cycle.

For local DVMA slave write cycles, an SBus slave cycle is during the the data get stage, and an internal data transfer is during the data put stage. In this case, the DVMA transaction is finished after the slave cycle completes in the data get stage. The current cycle is not held up during the internal data transfer, but data put stage may show up as bus latency before the next translation cycle occurs.

Bypass DVMA slave cycles do not involve the SBC as a slave target. The data transfer is between an SBus master and another SBus slave. There is no data get and data put operations in this case.

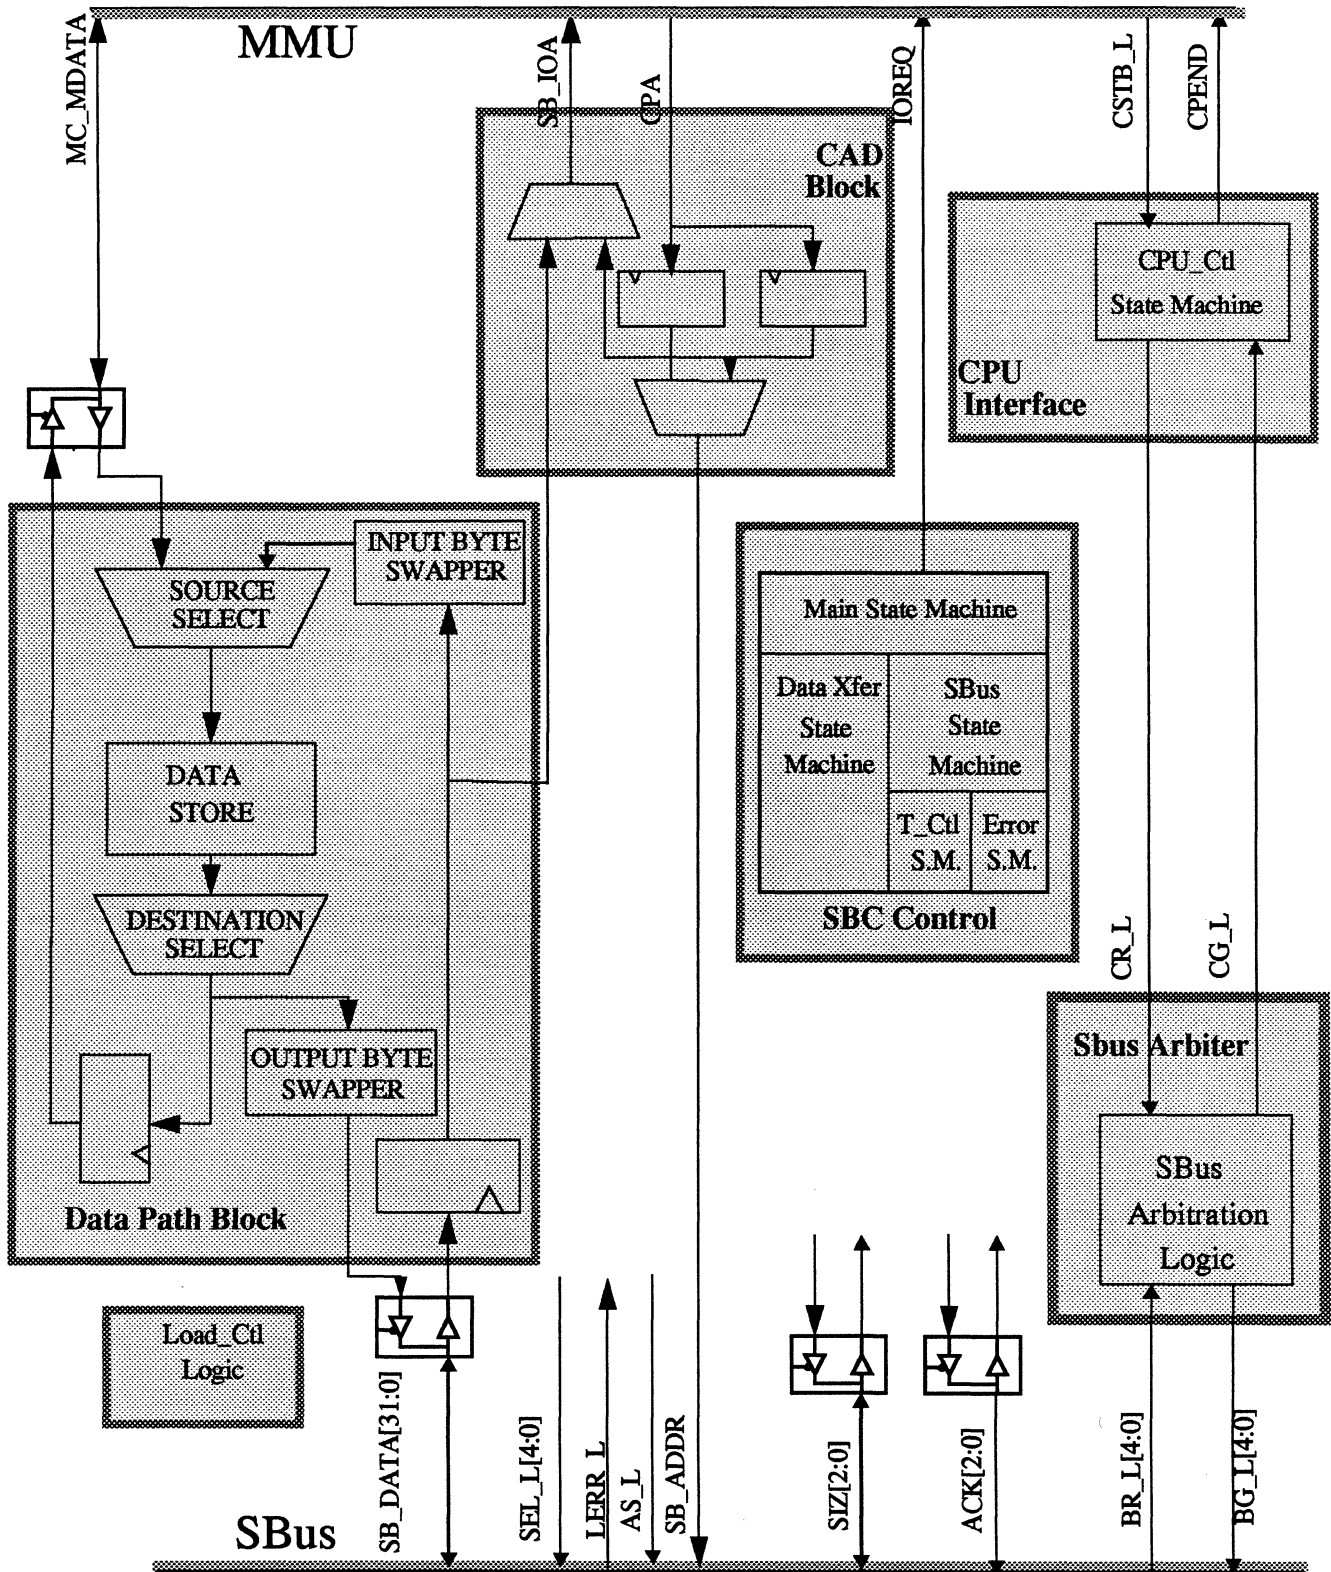
As a bus controller, the SBC has to handle bus errors and watchdog timeouts. Bus errors that occur during PIO cycles are handled by making the current state of the bus cycle available to the MMU. Bus errors that occur during DVMA cause the SBC to intercept the slave cycle from the intended slave target and, itself become the slave target in order to terminate the cycle with an error. Watchdog timeouts occur when an internal timer expires and the SBC terminates the slave cycle with an error.

The subcomponents of the SBC are the CPU Interface, Address Steering, SBus Arbiter, Main Control, Data Transfer Control, SBus Slave and Target Control, Data Path and Control, and Error Control blocks. These subcomponents are schematically shown in the following block diagram. A further description of these blocks is given in the following sections.

If it is appropriate, a state diagram is provided, along with a narrative walk-through. The state diagrams are provided for heuristic purposes. The intent is to purposely omit descriptions of some logic that tends to

cloud the understanding of the general functionality. This logic, used for such implementation-specific purposes as logic synthesis and timing aids, would detract from the overall comprehensibility of the SBC block.

Figure 8.2 - SBus Controller Block Diagram



8.0.2 CPU Interface

The CPU interface block handles the MMU/SBC handshake protocol, arbitrates for the SBus, catabolizes double word PIO into single word PIO, if appropriate, and supports dynamic bus sizing and bus cycle reruns. The data sizes supported for PIO cycles are byte, half byte, word and double word. There is also a high-performance feature that allows for very fast PIO writes to occur, which is especially important for certain operations that require fast output, such as graphics.

The CPU interface is double buffered, meaning that a copy of the entire state of the current cycle is retained for both PIO reads and PIO writes. The double buffering is necessary in the event of dynamic bus sizing, catabolic double word transactions or bus cycle reruns. The buffering also permits a DMA address translation to occur concurrent with a PIO transaction. This is important in deadlock avoidance.

The deadlock could occur when simultaneous PIO and DVMA transactions occur. The deadlock is avoided by buffering the entire state of the PIO transaction, and allow the DVMA transaction to proceed. Upon completion of the DVMA transaction, the PIO transaction, which had been retained in the SBC, would proceed.

The PIO buffers effectively provide a single element of a write buffer, since the CPU continues to execute instructions without waiting for a PIO write to complete.

A walk-through of the CPU State Machine (CSM) is given below: A set of signals, CSTB_L, CPEND, and IOREQ form a handshake between the MMU and the SBC. A PIO transaction is issued from the CPU through the MMU to the SBC by the assertion of CSTB_L. This occurs only if the SBC is not busy, as indicated by CPEND. De-assertion of CPEND indicates that the SBC is not busy and free to receive a PIO cycle. In the case of PIO reads, IOREQ is used to signal that data is ready. (IOREQ is also used for other various MMU/SBC communications).

The CSM remains in idle until CSTB_L is received. For a simple PIO transaction, control transitions into a bus request state where it remains until the bus is acquired. Once the bus is acquired, the state changes, and holds until the bus is relinquished. Next, in the case of PIO reads, control passes to a housekeeping state before returning to idle; in the case of PIO writes, control returns directly to idle.

Special states in CSM support catabolic double word transactions. Whenever a double word PIO transaction is attempted to an SBus device that does not support bursts, the SBC automatically catabolizes the double word burst transaction into single word transactions. (Status bits

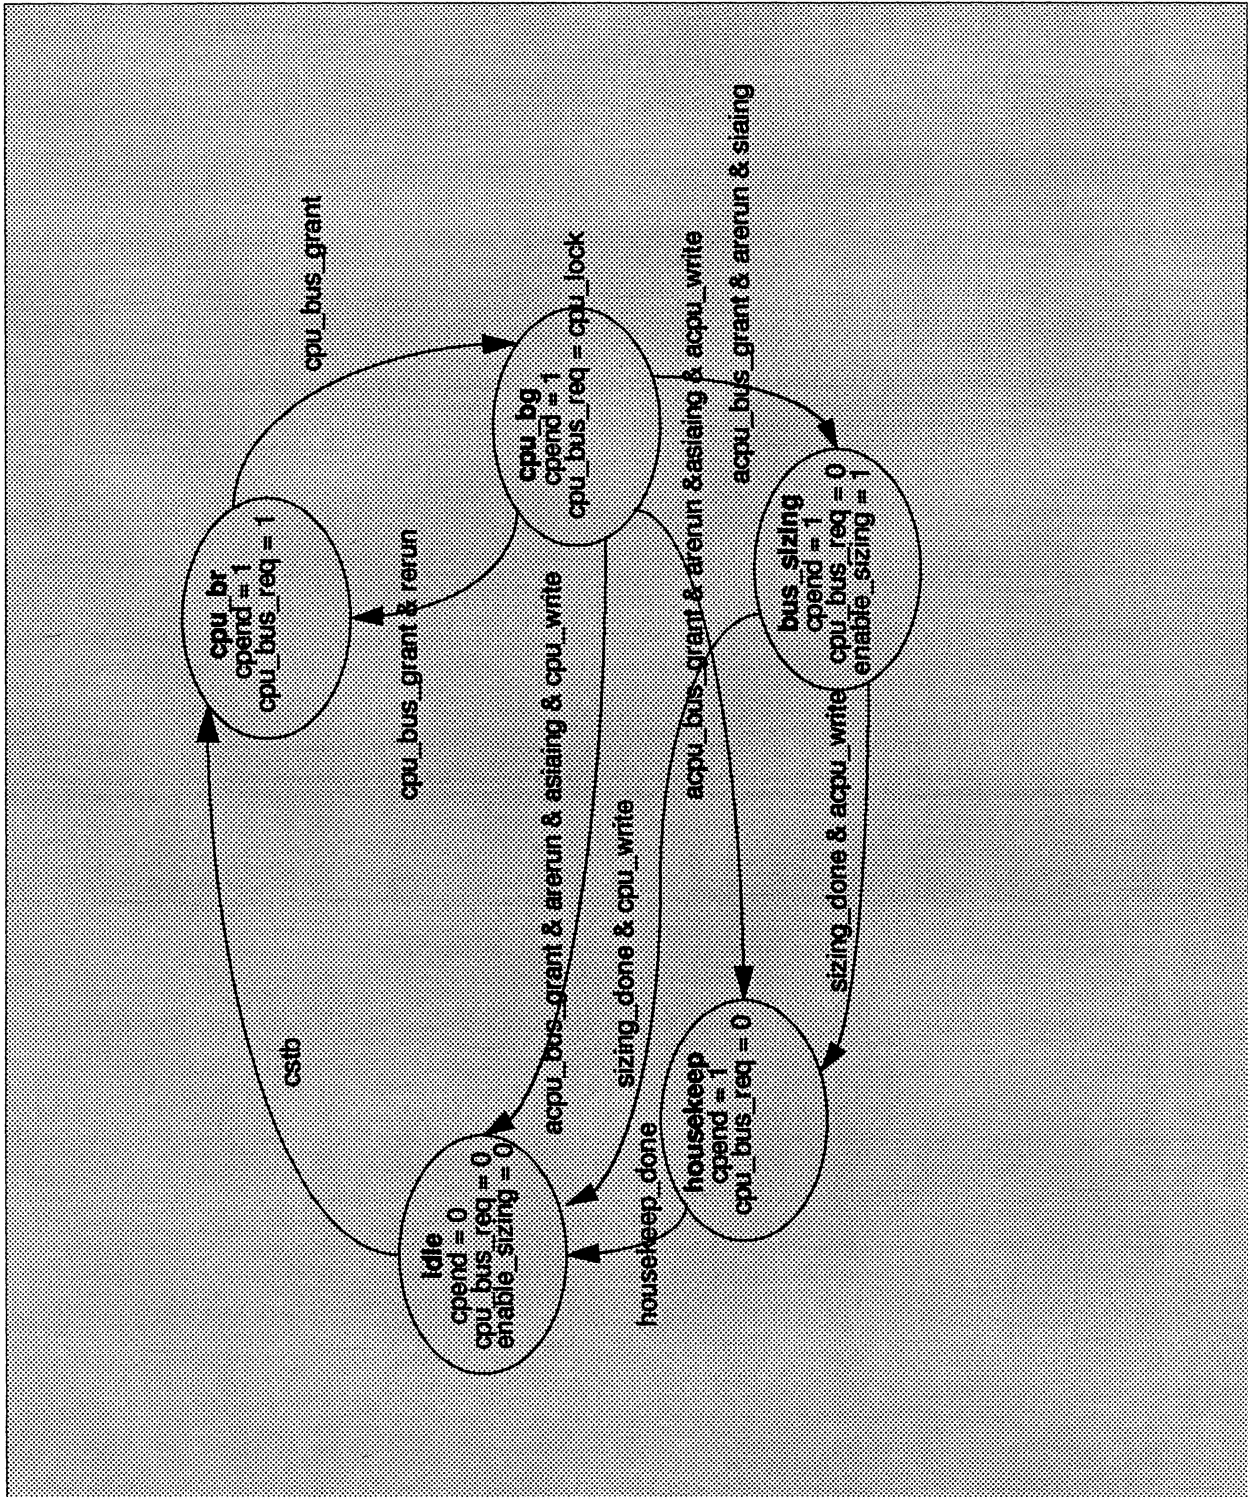
from the MMU Slot Configuration Registers indicate SBus device burst-handling capabilities) The MMU is held off during this time by CPEND. This operation is transparent to the MMU. Dynamic bus sizing and bus cycle reruns can occur in either portion of the catabolized transaction.

While a PIO cycle is in progress, as indicated by the CSM grant state, a dynamic bus sizing operation may occur which would cause control to branch to a special holding state. Simultaneously, the dynamic bus sizing state machine transitions from idle to handle this operation. When finished the CSM is signaled and control continues as if a simple PIO transaction had occurred. To improve latency during dynamic bus sizing, an attempt is made to keep the follow-on cycles atomic. If a rerun occurs, however, other DVMA masters are given a higher bus arbitration priority and the atomicity of the follow-on cycles will be broken. Reruns are supported whenever they occur.

Reruns can occur during any phase of a PIO transaction; during simple PIO transactions, dynamic bus sizing, atomic cycles, or catabolized transactions. When a rerun occurs, the transaction is ended, the bus is relinquished, and the cycle begins anew. Provisions are made the bus arbitor to allow any requesting DVMA masters onto the bus, before the PIO cycle is retried. For this reason atomic transfers may not work properly when the SBus slave recipient is capable of reruns.

A special speed path is built into the CSM to allow fast PIO writes. A prerequisite for this operation is that the data size must be a word (or double word) and must not be atomic. Another necessary condition is that the SBus slave device must respond with word acknowledge. If this criteria is met, then PIO word writes can sustain a bandwidth of 33 Mbytes/sec at 25 MHz. (PIO double word writes can sustain 50 MBytes/sec.).

Figure 8.3 - CPU State Machine



8.0.3 Address Steering

The Address Steering Block handles the address generation function for SBus transactions and local resources data transfers. The block insures that the proper SBus physical address is valid and stable whenever address strobe is asserted. In addition, this block generates the address used during a request for local resources.

There are two sources for an SBus physical address; the CPU generates a virtual address during PIO transactions and the SBus master generates a virtual address during DVMA transactions. In both cases the MMU translates the virtual address to a physical address. Since PIO and DVMA transactions can overlap, both physical addresses must be retained by the Address Steering Block. The only time the CAD block manipulates the SBus physical address is during double word catabolism and dynamic bus sizing.

In order to deal with such implementation-specific processes as memory burst order and local resource transfer sizes, the CAD block manipulates some low-order address bits to simplify data transfer control. In either case, data is transferred properly and control logic is simplified.

8.0.4 SBus Arbiter

The SBus Arbiter handles bus requests from the CPU and as many as five DVMA masters. The SBus Arbiter employs all fairness, and arbitration protocol as outlined in the SBus Specification A.2.

The fairness algorithm utilizes a token, which is passed round robin style. All six masters are given tokens which are prioritized based on the last master to have owned the bus. The requesting master with the highest priority is granted the bus. Once that master is finished with the bus, new tokens are assigned. The last owner is given the lowest priority.

The CPU is treated as one of the six masters. In this regard, the CPU master is indistinguishable from any other DVMA master. In addition to this, there are two ways in which the CPU is given special treatment. If the bus is free and is not about to be granted, the CPU has the ability to anticipate that its request will be granted. In this fast-bus-access case, the CPU will forego waiting for the bus grant in order to begin the bus cycle.

Another special case is made during times when a PIO transaction is dynamic bus sized. An attempt is made to keep the follow-on cycles atomic with the first cycle (although a rerun will cause the atomicity to be broken) in order to help the latency of that cycle.

A walk-through of the Arbiter State Machine (ASM) is given below: Control begins in Idle where it remains until a sampled version of at

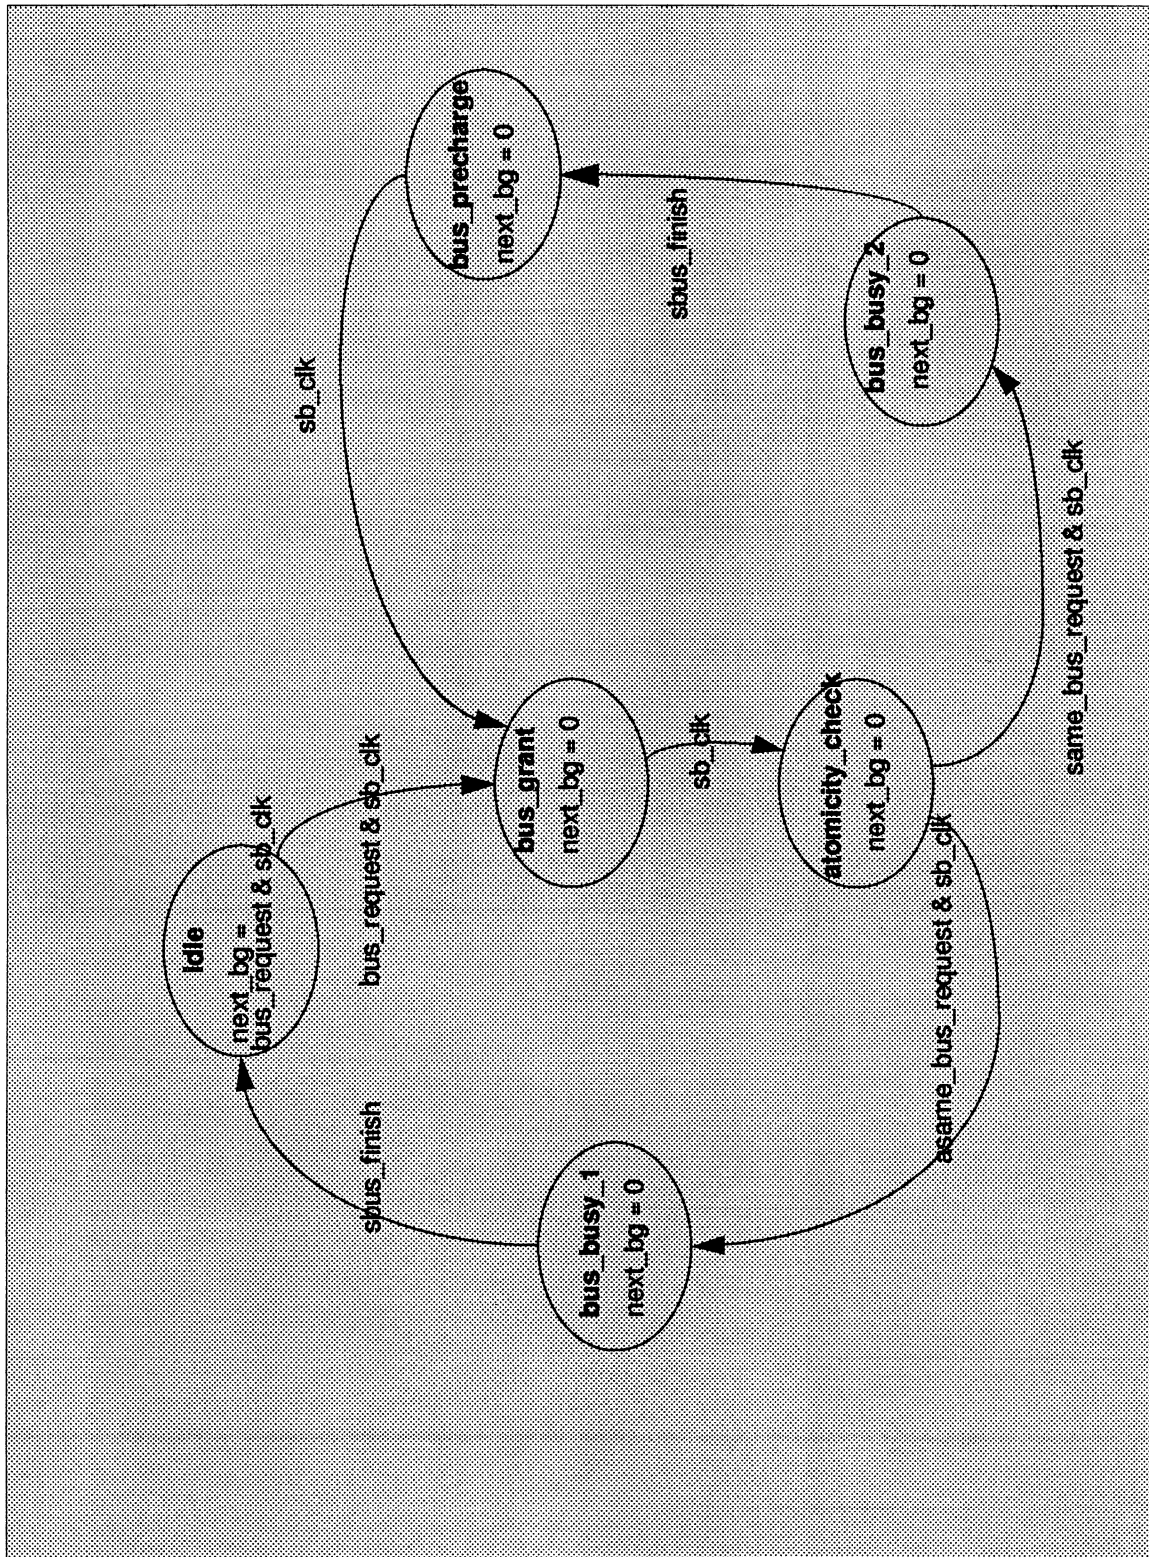
least one bus request is detected. Control transitions into the Bus_Grant state at the same time that the requesting master with the highest priority token is granted the bus.

On the proper phase of SBus clock, control moves into the Atomicity_Check state, where the current bus owner's request line is sampled for atomicity. For DVMA masters, the virtual address is latched during Atomicity_Check and a translation request is issued. If the request is still active, control branches to a special atomicity loop; otherwise control passes into a Bus_Busy state. Here it remains until the bus cycle is finished and then returns to Idle.

If the atomicity loop was taken, a different Bus_Busy state is entered. This Bus_Busy state is very similar to the first, except instead of entering Idle upon completion a Bus_Precharge state is entered. On the proper phase of SBus clock, control transitions into the Bus_Grant state. Other masters, however are not given an opportunity to compete for the bus and another bus cycle is granted to the previous master.

Once in the Bus_Grant state, control cannot distinguish how it arrived in that state (from either the Idle or the Bus_Precharge state). This means that each cycle is separate onto itself, and the atomicity check is made once during each bus cycle. It is possible for a master to retain the bus by constantly requesting it. (although good SBus citizens would never do this).

Figure 8.4 - Arbitor State Machine



8.0.5 Main Control

The Main State Machine (MSM) controls the internal data store, issues data transfer and translation requests to the MMU, and generally acts to coordinate the other state machines in the SBC. One of the major functions of the SBC is to move data between local resources and SBus devices. The SBC has to fill the internal data store by a get operation and then to empty the data store with a put operation. The MSM controls the above-mentioned get and put operations.

A walk-through of the MSM is given below: When in the Idle state, MSM monitors the state of ASM through two signals, CG, which indicates that the CPU has been granted the bus and XLAT, which indicates that a valid virtual address has been received and is ready to be translated. Control remains in the Idle state until one of these two signals is detected.

If CG is detected and it is a PIO write, then control transitions to the SputW state. In the case of PIO writes, the data store was filled concurrently with the issuance of the PIO cycle. All that remains to be done is to put the data to SBus space. An SBus cycle is issued. Control remains in SputW for the duration of the SBus cycle and then returns to Idle upon completion.

If CG is detected and it is a PIO read, then control transitions to the SgetR state. In the case of PIO reads, the data store must be first filled by a get operation from SBus space. An SBus cycle is issued from SgetR. Control remains in SgetR for the duration of the SBus cycle. Under certain conditions, such as reruns, dynamic bus sizing and catabolic double word cycles, control returns to Idle in anticipation of another CG. Upon completion, control passes to DputR. In DputR valid data is indicated by assertion of IOREQ. The data is then put to the CPU.

If XLAT is detected, a DVMA transaction is in progress and control passes to the Xlate state. A translation cycle is requested from the MMU through IOREQ. PA_VAL, from the MMU, indicates that the translation cycle has completed and the physical address is available. Upon receipt of PA_VAL, a 3-way branch occurs. The target of the DVMA cycle is determined from the physical address; it is not possible to know the target of the DVMA transaction until the transaction is complete. If the target is not system memory or control space, then Sbyp state is entered. If the target is memory or control space, then control branches to DgetR for DVMA reads; SgetW for DVMA writes. If any error had occurred at any time during the translation cycle, a special 4th branch, the Pet state, is entered.

From the SbyP state, the only function of the SBC is to act as an SBus controller; the SBC has neither a master or a slave role. An SBus cycle is simply issued and upon completion of the cycle control returns to Idle.

From the DgetR state, a get from memory or control space is required in accordance with a DVMA read transaction. Implicit to the read translation request is a data request if the target is determined to be system memory or control space; a separate request for data is not required. Once the data store is filled, control moves to SputR. If a parity error occurs, the Pet state is entered.

From the SputR state, the DVMA read transaction is completed by issuing an SBus cycle. Control remains in SputR state until completion and then transitions back to Idle.

From the SgetW state, a get from SBus space is required in accordance with a DVMA write transaction. An SBus cycle is issued, the data store is filled, and control moves to DputW upon completion.

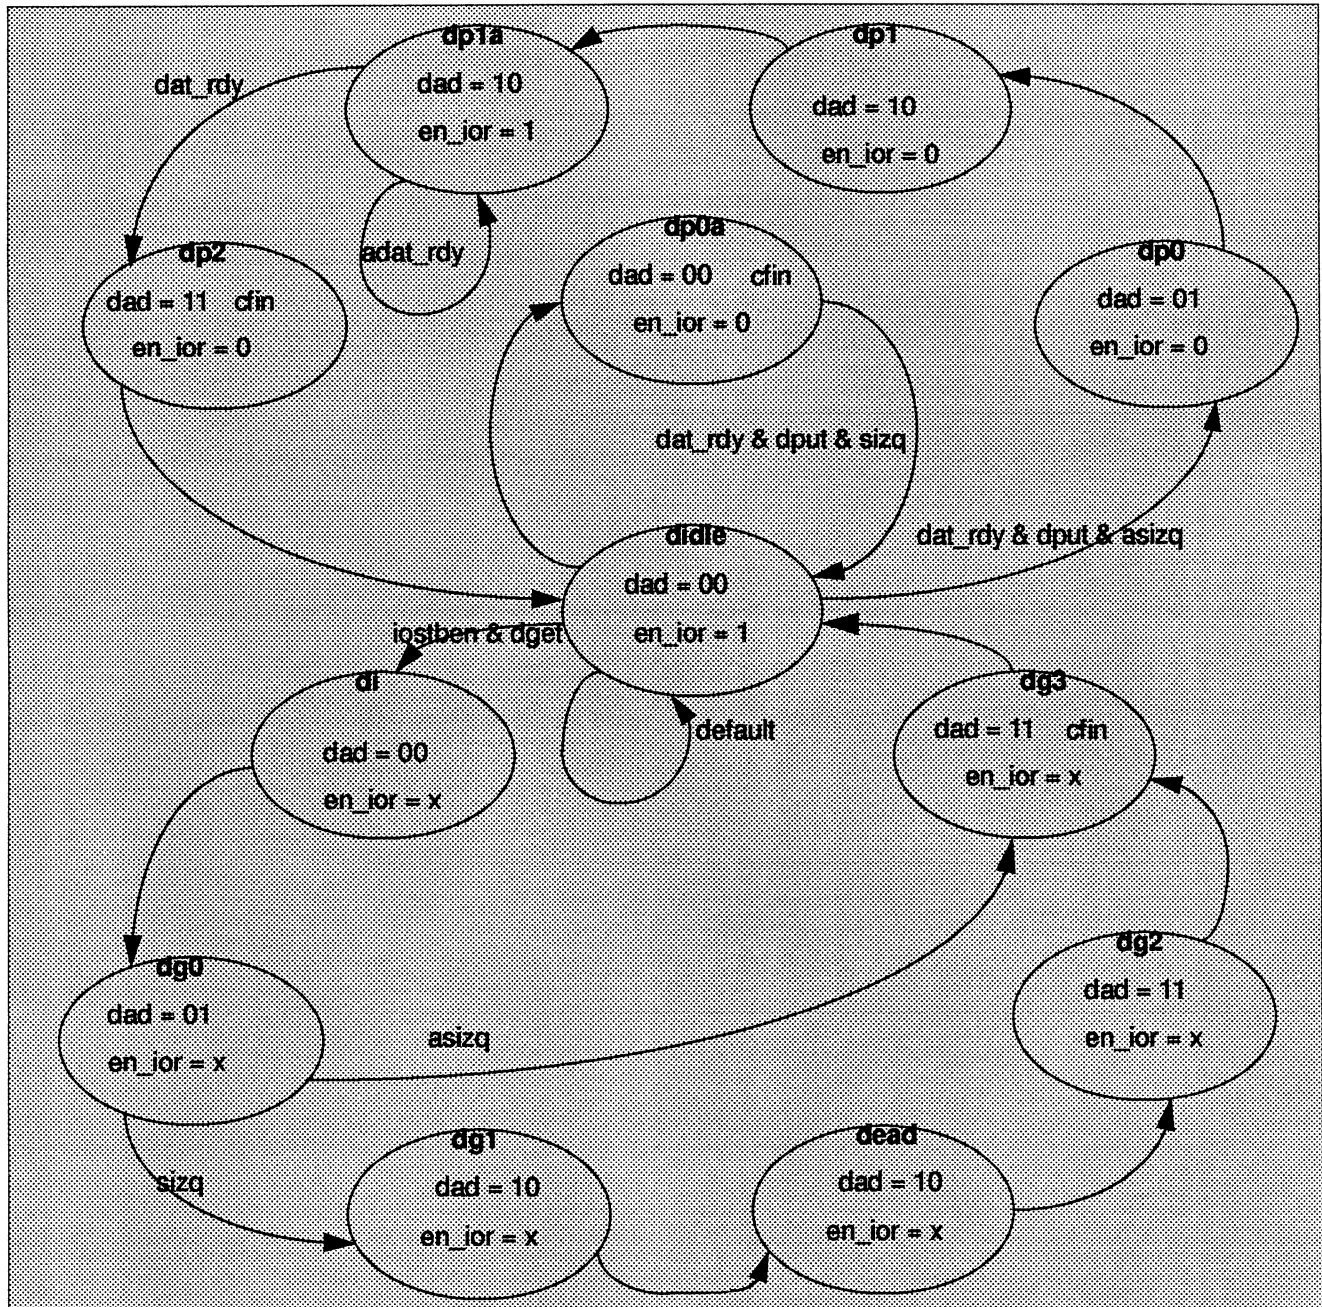
From the DputW state, the DVMA write transaction is completed by a put of the data to system memory or control space. IOREQ is asserted to request a write cycle. After the put operation is complete control returns back to Idle.

From the Pet state, an SBus cycle is issued, but the SBus controller must intervene since an error has occurred. The slave select must be suppressed in order for the SBC to become the target. A special signal is sent to the Target State Machine which has the responsibility of driving SB_ACK[2:0] to indicate an error in this case. Control remains in Pet until completion of the SBus cycle and then returns to Idle.

8.0.6 Data Transfer

The data control state machine (DSM) controls the movement of data between system memory or control space and the internal data store. The DSM monitors transaction size information and data transfer signals from the MMU. The data is counted and a signal, CFIN, is asserted upon completion.

Figure 8.6 - D_ctl State Machine



8.0.7 Slave Control Cycle

The SBus control state machine (SSM) is charged with tracking the progress of the current SBus cycle by monitoring the Transfer Acknowledgment (ACK) and terminating the cycle once completed or upon an error detection. The SSM does not differentiate between the ACKs from the TSM and other external ACKs; it treats the TSM as any other slave capable of responding with an ACK.

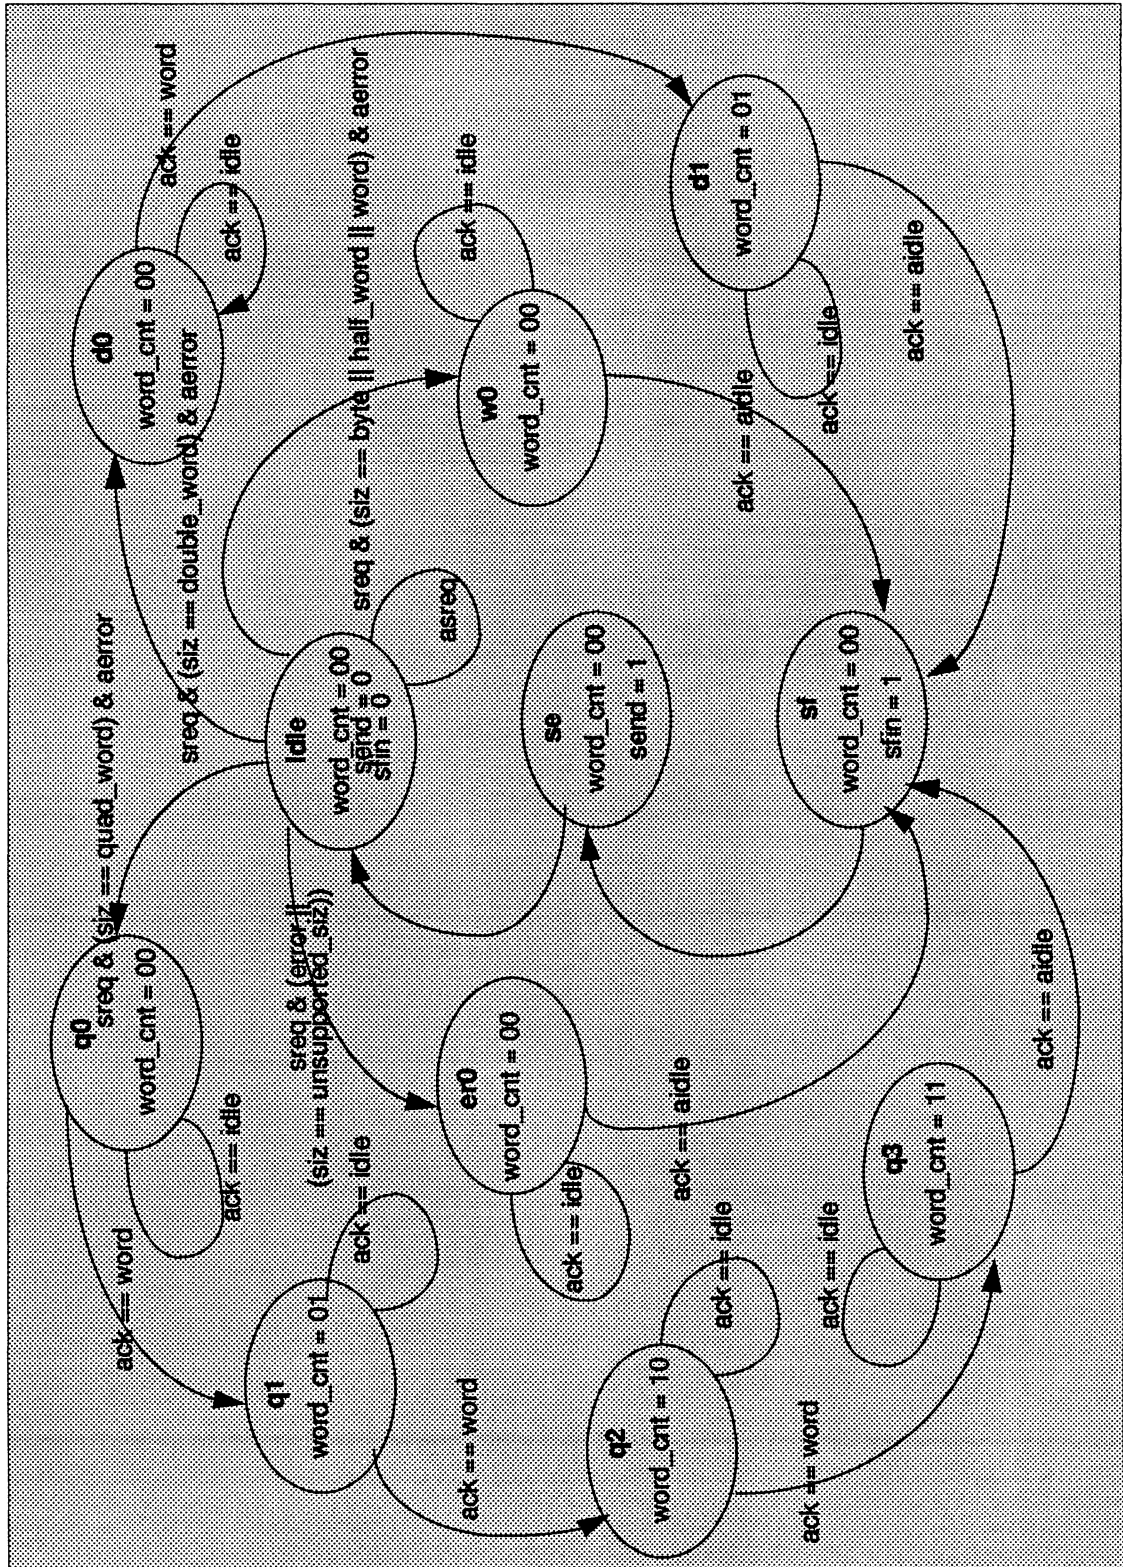
A walk-through of the SSM begins in Idle, where the SBus request line is sampled. Once a request is detected, control transitions to the appropriate state as a function of the bus size and the error signal; W0 if the size is a word or smaller, D0 if the size is a double word, Q0 if the size is a quad word, or Er0 if the size is unsupported or an error was detected. Once in either W0 or Er0, the ACK lines are monitored and any ACK code other than Idle/Wait will cause a transition to Sfin.

Once in D0, the ACK lines are monitored and the SSM is effectively enabled to count Word ACKs. Word ACK will cause a transition to D1. Idle/Wait ACK will keep control in the D0 state. D1 is similar to Er0 and W0. Any ACK code other than Idle/Wait will cause a transition to Sfin.

Once in Q0, (or Q1, or Q2) the ACK lines are monitored Word ACK will cause a transition to D1. Word ACK will bump control to the next higher word count stage until Q3 is reached. Idle/Wait ACK will retain state in Q0 (or Q1, or Q2). Q3 is similar to D1, Er0 and W0. Any ACK code other than Idle/Wait will cause a transition to Sfin.

Once in Sfin, the SBus cycle is nearly complete, except for some amount of housekeeping. Control transitions to Send and then returns to Idle.

Figure 8.7 - S_ctl State Machine

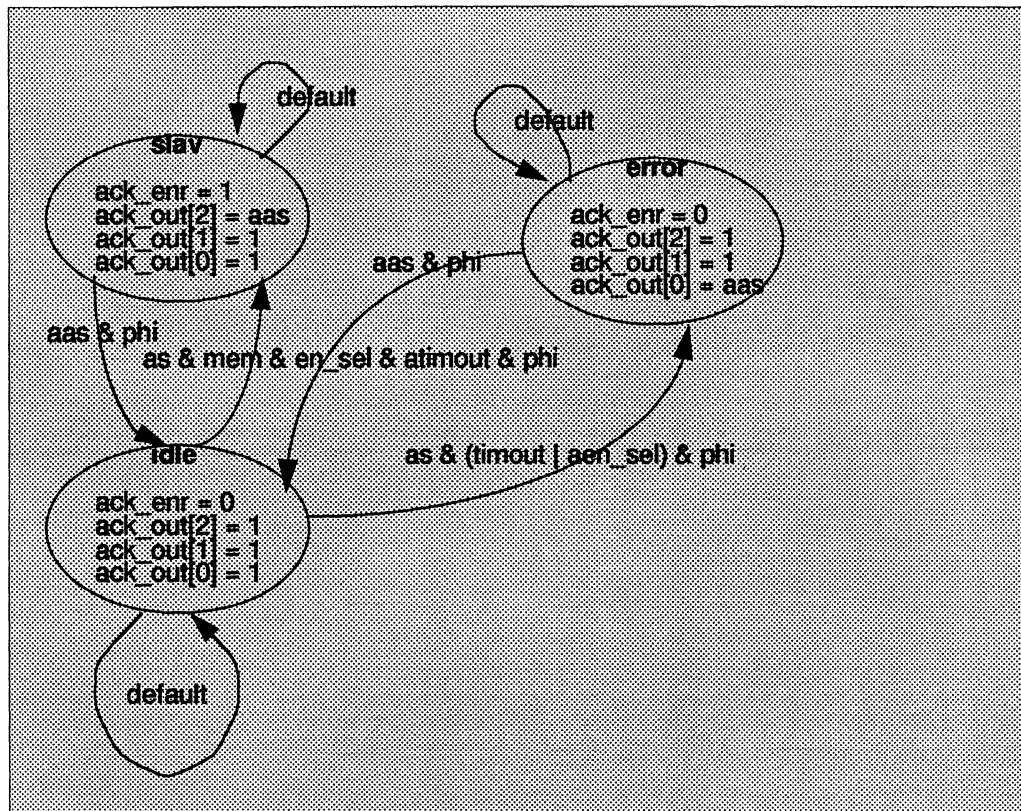


8.0.8 Slave Target Control

The Target control State Machine (TSM) controls the Transfer Acknowledgment (ACK) during local DVMA transactions or error conditions, when it is appropriate for the SBC to drive these signals.

A walk-through of the TSM begins in Idle, where control remains until it recognizes itself as the target of the current slave cycle. Since the TSM is clocked at twice the frequency as the SBus, the phase of SB_CLK is important. If the TSM is the target and either the memory_select or the error signal is detected, then control moves out of Idle to either the Error or Slave state. ACK is enabled and the proper code is asserted. When finished control transitions to the Precharge state where ACK is precharged and then control returns to Idle.

Figure 8.8 - t_ctl State Machine




8.0.9 Data Path

The SBC data path consists of a series of multiplexers and registers necessary to transfer data between the SBus devices and local resources. There are two sources of data: the internal data bus, MC_MDATA, which connects local resources to the SBC and the SBus data bus. Data from the SBus is buffered, then passes through a byte swapper, which is necessary to align SBus byte or half-word ports, passes through a source select mux on its way to the internal data store. Data from the internal data bus passes through the source select mux to the internal data store.

The heart of the data path is the internal data store, which provides temporary storage for up to 24 bytes of data. DVMA has exclusive use of 16 bytes of internal storage and 8 bytes are exclusively used for PIO. Each byte-sized register corresponds to an address location. This means that data from a given address location will always be loaded into the same internal data store location, regardless of the order in which the data arrives. Data from either the internal data bus or SBus can go to either the DVMA register bank or the PIO register bank.

Data destined for the internal data bus goes from the internal data store, passes through destination select muxes, into an output buffer and is enabled onto the internal data bus by a tristate driver. Data destined for the SBus passes through destination select muxes, through an output byte swapper, necessary to support dynamic bus sizing and is enabled onto the SBus by a tristate driver.

Figure 8.9 - SBC Data Path



8.0.10 Data Control

The SBC data path control logic steers the data through the source and destination multiplexers and loads the data into the internal data store. The source and destination multiplexers are straightforward to control. Since the get and put operations are serial, the data steering is almost static. The main state machine controls the get and put operations.

The internal data store load control works by the application of a mask to the load enables of each byte-sized register. Data can arrive in sizes of as small as a byte. As each piece of data arrives the load enable of its register is masked, thereby preserving the data until the put operation.

8.0.11 Error Handling

The Error Control Block handles errors that occur during both PIO and DVMA transactions. There are three possible sources of error for PIO transactions. There are PIO transactions terminated by timeouts, error acknowledge, and late error. A two-bit error status field, `ERR_TYPE`, is used to indicate to the CPU the source of error during PIO transactions. This bus is sampled and any code other than that indicating no error signifies that an error has occurred. During this time, the entire state of the current PIO transaction is made available to the CPU for error reporting.

The sources of error for DVMA transactions are translation, parity, timeout and SBus protocol errors. Parity can occur either during address translation or a get operation from local resources. In all cases the SBC becomes the slave target and drives `ACK` to indicate an error to the DVMA master. Errors during DVMA are transparent to the CPU. The SBC does not use the SBus late error signal to indicate errors.

8.0.12 Diagnostic Testing

The SBC employs JTAG and therefore allows all registers to be scanned during JTAG scan mode. Tristate enables for SBus signals that are bidirectional are disabled during scan mode.

A testing feature allows the internal data bus and address bus to be observed when the system is placed in a special diagnostic view mode. When placed in view mode, the SBC steers the internal data bus onto the SBus data bus and the internal address bus onto the SBus address bus. In both cases the address and the data bus information is delayed by one system clock. Of course, the proper SBus address and data is not available during view mode; the SBus is used exclusively for testing at this time.

8.0.13 Additional Work

This section is included for future work. There are two architectural aspects that can potentially yield a better SBus controller design: use a separate, non-unified I/O MMU and a complete handshake protocol between the memory controller and the SBus controller.

PIO and DVMA transactions are distinctly orthogonal operations and as such they can potentially occur in parallel at any time. Future designs should look carefully at the costs and benefits of sharing the resources involved (such as unified MMU/IOMMU which precludes some parallelism).

Another optimization that could be made on a future implementation is a full handshake for data transfer between system memory and SBus. This could reduce the amount of internal data storage in the SBC and at the same time increase the transfer sizes that can be supported. If the SBC could request a memory cycle and then signal when data is ready, then the internal data store need only be as big as the data bus size. This is particularly important during DVMA write transactions. With a better handshake mechanism, the SBC could request a memory transfer while the first data word is available. Latency could be avoided and storage elements larger than the size of one data word saved.

9.0 Reset, Clock Control, JTAG

This section will describe the Reset logic, Clock Control logic and the JTAG architecture. The JTAG, reset control and clock start/stop control logic are part of the Misc block, while the clock controller is a design block by itself.

All registers in the microSPARC CPU reset to zero except where otherwise noted. All RAMs including the IU and FPU register files, the data and instruction cache rams, the and TLB remain unchanged by the assertion of Reset.

State and pipeline registers internal to the IU are established on reset via reset logic in the IU, not via explicit reset to the flip-flop. This is to support clearing and setting certain bits (e.g.: S bit of the PSR).

The JTAG logic controls all the scan operation within the chip and in conjunction with the clock start/stop logic, enables the single step operation of the chip for debug purposes. All of the registers in the chip are scannable and are configured as one single internal scan chain for testing as well as debugging the chip.

9.0.1 Reset Controller

The microSPARC Reset Controller performs the simple task of driving microSPARC's internal reset lines, and inhibiting clocks during transitions on those lines to avoid timing violations on the flip-flops being reset.

microSPARC has two reset operations: General Reset (sometimes called SBus Reset) and Watchdog Reset. General Reset is done in response to assertion of the input_reset_1 microSPARC input pin; this happens on powerup and on any externally-triggered reset. Watchdog Reset is performed when the IU enters error state due to a taking a trap while the PSR ET bit is deasserted. General Reset will cause assertion of both Reset Controller output signals: reset_any and reset_nonwd; Watchdog Reset will cause only reset_any to be asserted. Reset_any resets the IU and any other logic which must be reset only on Watchdog Reset; reset_nonwd resets everything else except the clock and reset logic and the TAP controller.

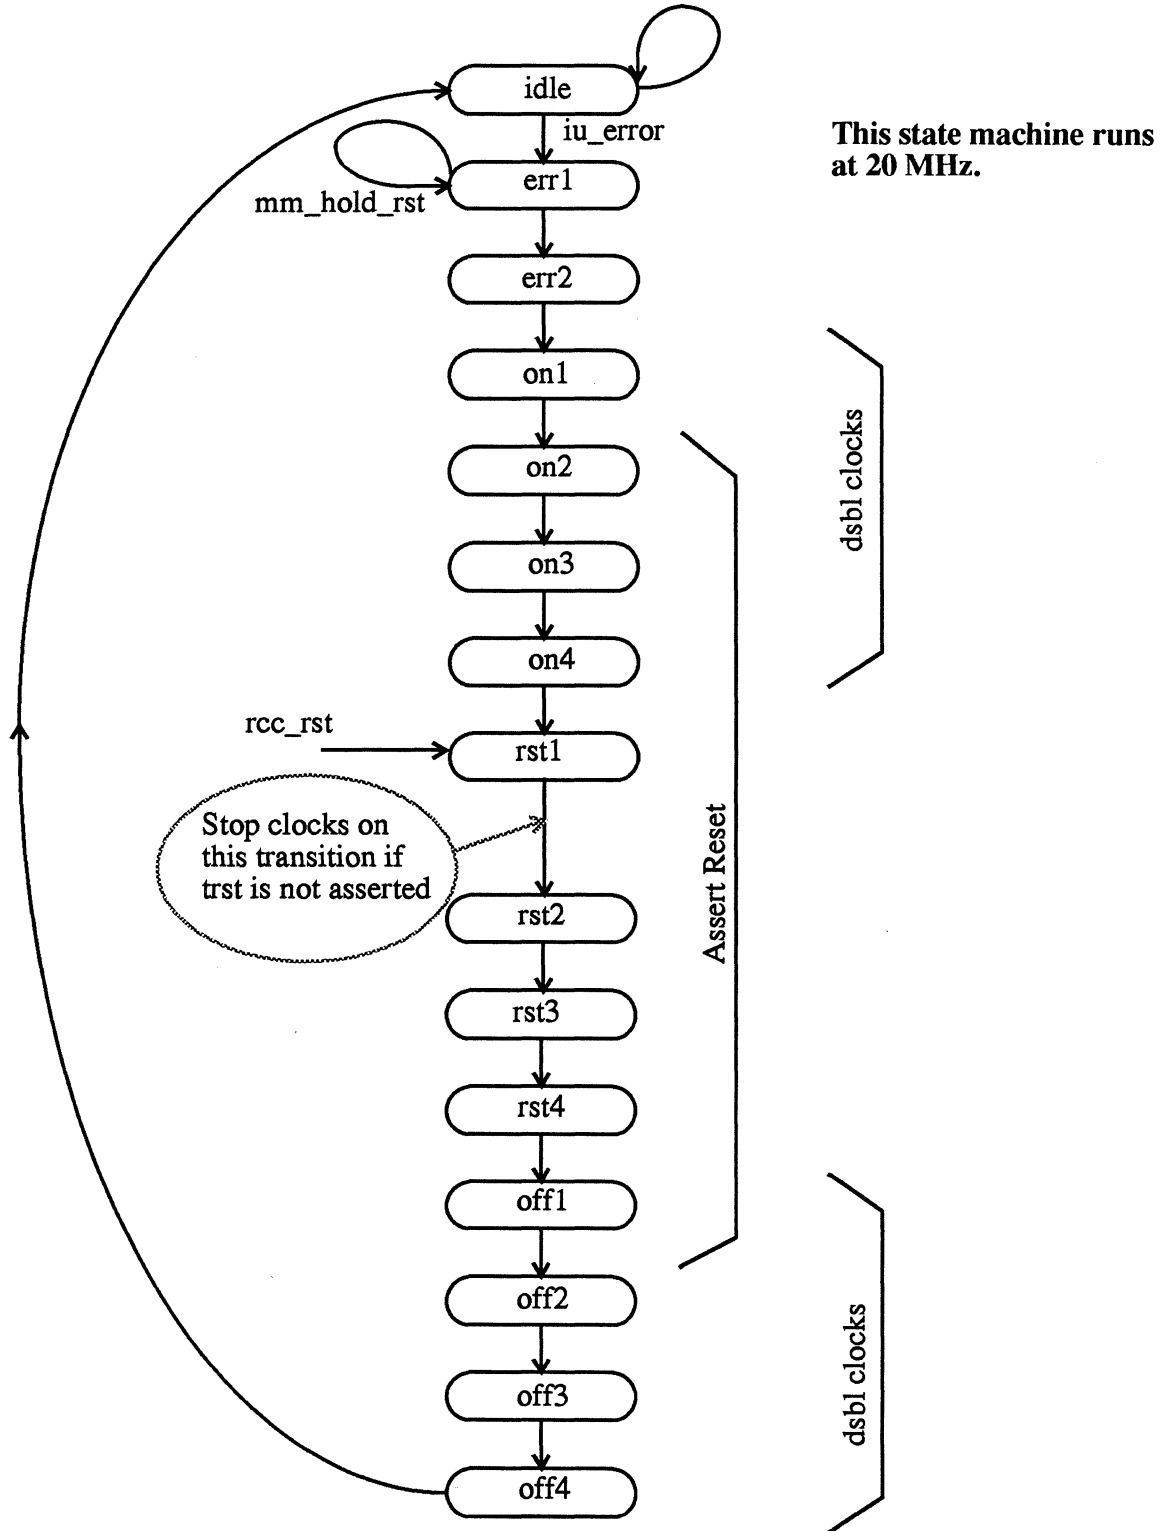
In addition to reset_any and reset_nonwd, the reset controller has another output, rs_dsbl_clocks, which is used to disable the outputs of the clock controller during transitions on the reset lines. This allows the heavily-loaded reset signals time to propagate throughout the chip completely between clocks, to avoid setup and hold time violations. All three of these outputs are controlled by the reset state machine.

However, `input_reset_1` is combinatorially ORed into both `reset_any` and `reset_nonwd`, and `rcc_rst` forces clocks to be running; taken together, these assure that any circuitry which must (for physical reasons) see reset asserted immediately on powerup will see it (assuming that `input_reset_1` is asserted, and `input_clock` is oscillating, immediately on powerup). As a consequence, timing violations may occur on the first clock after assertion of `input_reset_1`; presumably, the ensuing General Reset will eventually clean up any illegal states caused by these violations.

Inputs which affect operation of the reset state machine are: `rcc_rst`, a 20-MHz¹ - synchronized version of microSPARC's `input_reset_1` pin; `iu_error`, the error state indication from the IU which initiates a Watchdog Reset; and `mm_hold_rst`, a signal from the MMU which delays the start of a Watchdog Reset sequence until there are no loads, stores, or instruction fetches in progress. `Rcc_rst` is inhibited during scan shift operations, to prevent loss of non-resettable state if `input_reset_1` should happen to be asserted during a scan shift.

1. Throughout this section of the document, waveform frequencies and periods will be given as if the frequency of `input_clock` were 80 MHz, even though this logic will run correctly at any speed from the design frequency (100 MHz) down to DC.

Figure 9.1 - microSPARC Reset State Machine



9.0.2 Reset Controller State Machine Operation

The reset state machine is clocked at 20 MHz. Assertion of `rcc_rst` synchronously resets the state machine into the `rst1` state from any other state. The state machine will thus stay in state `rst1` for as long as `rcc_rst` is asserted. After completing a reset sequence, the state machine hangs in the idle state until either `iu_error` or `rcc_rst` is asserted. If `iu_error` is asserted while in the idle state, the state machine goes to state `err1`, waits there until `mm_hold_rst` is deasserted, and then completes the reset sequence and returns to idle. `Reset_any` and/or `reset_nonwd` are asserted in states `on2`, `on3`, `on4`, `rst1`, `rst2`, `rst3`, `rst4`, and `off1`: if the reset sequence was initiated by `iu_error`, only `reset_any` is asserted; if initiated by `rcc_rst`, both `reset_any` and `reset_nonwd` are asserted. Clocks are disabled in states `on1`, `on2`, `on3`, and `on4` as the reset signal is turned on; they are disabled again in states `off1`, `off2`, `off3`, and `off4` as reset is turned off again. This clock disabling does not put the clock state machine into the stopped state; it merely gates off the clock outputs. Note that the reset lines transition from 1 to 0 only during a clocks-disabled period, and, for Watchdog Reset, they transition from 0 to 1 only during a clocks-disabled period.

To facilitate scan-based debugging, the reset state machine will assert `rs_stop_even` upon exiting the `rst1` state during a General Reset sequence. If microSPARC's `jtag_trst_1` input is deasserted at that time, this will cause the clock control state machine to enter the stopped state. The reset sequence will continue as clocks are issued under scan control. It is thus possible to single-step through the remaining states of the reset state machine, and, more importantly, to reset the machine to a known, deterministic state during scan-based debug.

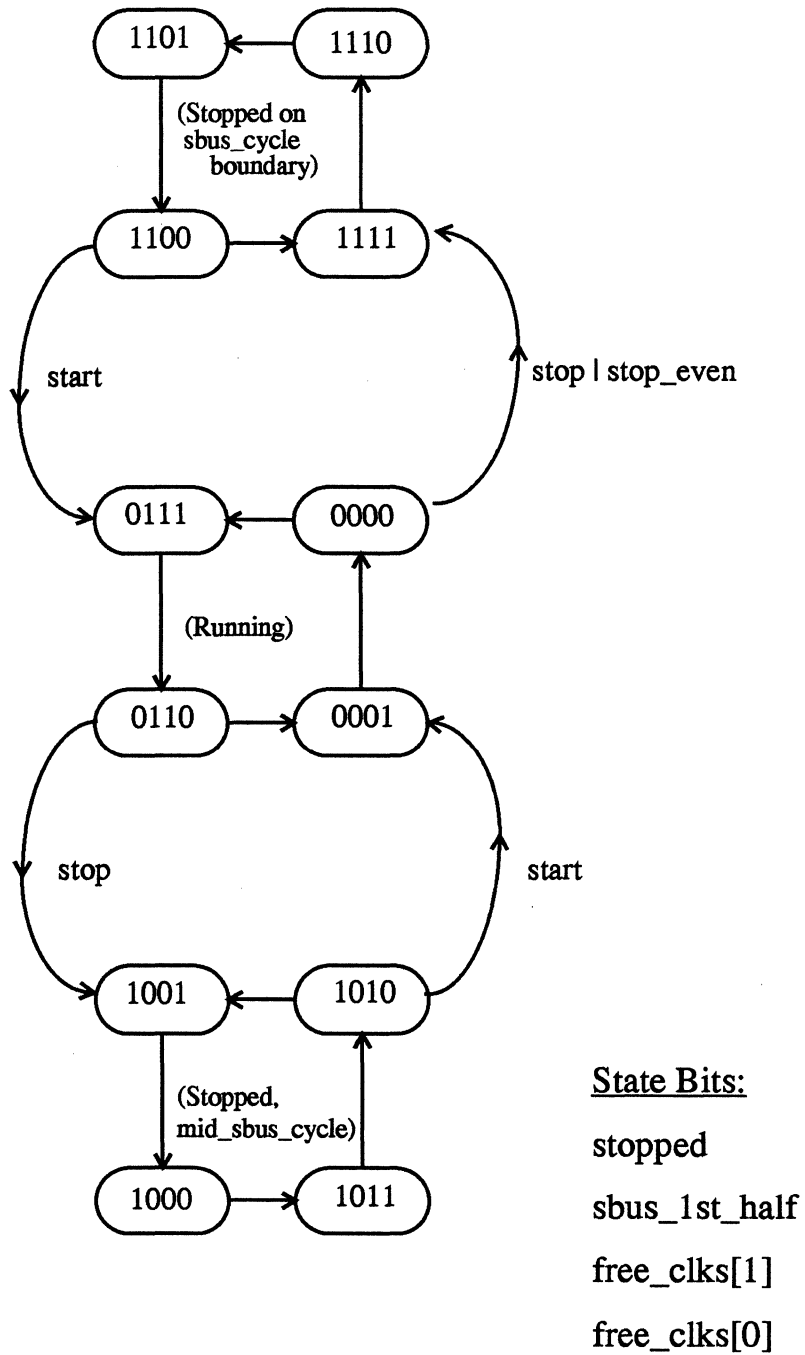
9.0.3 Clock Controller

The microSPARC Clock Controller generates the clock signals used by all of microSPARC (except the TAP controller), as well as the sbclk used by external SBus interface devices. Its operation is controlled by the Clock Control Register (CCR), a collection of internal register bits which are writable only by scan. On Reset, the CCR is cleared. Subsequent scan shift operations can be used to set bits of the CCR in order to alter the operation of clock state machine, as described below. However, the CCR has no effect on the operation of the clock state machine if the jtag_trst_1 microSPARC input pin is asserted (low).

At the heart of the clock controller is a 4-bit state machine, clocked by the 80-MHz input_clock. The low-order two bits of this state machine are a free-running two-bit down counter (free_clks[1:0]). The MSB (stopped) indicates whether clocks are stopped or running. The remaining bit (sbus_1st_half) indicates which half of the 20-MHz cycle the state machine is in, even when clocks are stopped. The two main clock outputs, ss_clock (40-MHz) and sbclk (20-MHz), are effectively equal to (free_clks[0] | stopped) and (free_clks[1] | stopped), respectively, although in the actual implementation these and all other clock outputs are driven by the Q outputs of 80-MHz-clocked flip-flops. There are three inputs to the clock state machine: start, stop, and stop_even; these are generated in the clk_stop submodule of the misc module. When stop is asserted while the stopped state bit is 0, the clock state machine will take one of these two transitions: 0000->1111 or 0110->1001, whichever comes first. When stop_even is asserted while the stopped state bit is 0, the clock state machine will take the 0000->1111 transition. When start is asserted while the stopped state bit is 1, the clock state machine will take one of these two transitions: 1100->0111 or 1010->0001, whichever comes first. The start input is actually a bit of the CCR, and it will reset itself on the first ss_clock positive edge, to facilitate the single-step operation.

The stopped and sbus_1st_half state bits are readable, but not writable, via scan. Synchronized copies of these two bits form a special two-bit scan chain which may be accessed via the sel_ccr TAP operation. This TAP operation, unlike sel_dbg_scan, does not interfere with the operation of the clock state machine, so the states of these bits may be polled at any time without affecting clocking. Note that 'sel_ccr' is a misnomer, since these two bits are not part of the CCR.

Figure 9.2 - Clock Controller State Machine



9.0.4 Clock Signals

Four distinct clock signals are generated by the clock controller. These are: `ss_clock`, the 40-MHz signal which clocks most of microSPARC; `sbclk`, the 20-MHz signal which is driven off-chip to clock the SBus interface logic and the external clock counter; `di_val`, a half-period-delayed version of `ss_clock`, used by the cache RAM megacells and other logic which requires a delayed clock; and `rcc_clock`, a 40-MHz signal which clocks the reset state machine and the CCR logic. All four of these signals will cleanly transition to the high state when the stopped bit of the clock state machine is high. During scan data shift and capture operations, all four clocks are disabled (i.e. forced high) by a synchronized version of the `testclken` signal sourced by the TAP controller; the clock state machine does not need to be in the stopped state for this to occur. All except `sbclk` are combinatorially ANDed with `testclk` (an active-low pulse train generated in the TAP controller by gating `jtag_ck`) during these disabled periods, so that flip-flops driven by all three of these clocks can be connected together in a single scan chain. All except `rcc_clk` are disabled by the `rs_dsbl_clocks` signal sourced by state decodes of the reset state machine, so that slow transitions on the internal reset lines will not cause setup violations. As with the `testclken` disable, the clock state machine need not be stopped when `rs_dsbl_clks` is asserted.

9.0.5 Stopping Clocks

To stop clocks, set the `stop_clocks` CCR bit. This will assert the stop input to the clock state machine, stopping clocks on the next 40-MHz rising edge.

9.0.6 Starting Clocks

To start clocks from a `stopped=1` state, set the start bit of the CCR.


9.0.7 Single-Step

From a `stopped=1` state, set both the `stop_clocks` and start bits of the CCR. A single 12.5-ns active-low `sys_clk` pulse will be issued; if `sbus_1st_half` was 0, a single 25-ns active-low `sbclk` pulse will also be issued (its rising edge will coincide with the rising edge of `sys_clk`).

Figure 9.3 - Single Step with sbus_1st_half = 1.



Figure 9.4 - Single Step with sbus_1st_half = 0.



9.0.8 Stop Clocks on Internal Event

To stop clocks on detection of an internal event, set the `stop_on_int_event` bit of the CCR and enable the desired internal event detection logic. Clocks will stop at the end of the `sys_clk` cycle in which the input to the `int_event` flip-flop is asserted. Internal events are detected by special logic in the IU and the MMU - see documentation on those units for more details.

9.0.9 External Cycle Counter

The microSPARC clock controller is designed to interface to a simple external cycle counter (XCC) for precise, at-speed control of system clocking. The interface consists of three microSPARC I/O pins:

- * `sbclk` (output) - the 20-MHz SBus clock output, which is gated off when system clocks are turned off. This output is used to clock the external SBus logic as well as the XCC.
- * `ext_event` (input) - this input is immediately registered in a 20-MHz-clocked flip-flop. Under control of some Clock Control Register (CCR) bits (which are writable only by scan), a logic 1

in this flip-flop will cause clocks to stop either at the next `ss_clock` rising edge or the next `sbclk` rising edge. This input should be driven by the `terminal_count` output of the XCC, perhaps ORed with other externally-detected clock stop signals. In a standard up-counter, the terminal count output is asserted when the counter contains all 1's (i.e. -1).

* `int_event` (output) - this is the output of a 20-MHz-clocked flip-flop. It is asserted whenever an internally-detected 'event' occurs (e.g. virtual address match). These events can, under control of some CCR bits, stop clocks; however, whether or not they stop clocks, they always cause assertion of the `int_event` output. This output can be used to trigger a logic analyzer; in addition, it can be used in conjunction with the XCC as described below to implement the 'stop N cycles after internal event' function.

Note that this interface runs at the 20-MHz SBus clock rate, and the signal I/O connect directly to inputs or outputs of flip-flops within microSPARC; thus, the XCC logic has nearly a full 50-ns cycle in which to set up its output to the `ext_event` input.

9.0.10 Counting Clocks

When the XCC is enabled, it increments on every `sbclk` positive edge. Since the states of the XCC and the CCR are accessible via scan, we can calculate how many 40-MHz system clocks have been issued between any two points in time by scanning out this state information before clocks are started and again after they have been stopped. The following formula can be used. `XCC.before` and `XCC.after` are the respective values of the clock counter before and after clocks have been issued; `sb1h.before` and `sb1h.after` are the corresponding values of the `sbus_1st_half` bit of the CCR.

$$N = 2*(XCC.after - XCC.before) - \sim sb1h.before + \sim sb1h.after$$

This formula of course assumes that XCC has not wrapped around; the XCC control logic should contain a wraparound detector that can be read by scan.

9.0.11 Issuing N Clocks

The XCC can be used to issue exactly N 40-MHz system clocks, at full speed. N can be any number from 1 to approximately $2^{*(X+1)}$, where X is the number of bits in XCC; for example, a 32-bit XCC lets us control clocks over a 200-second range at 40-MHz operation. This function does not require the use of the `int_event` output.

Several CCR bits are used for this function. When, while clocks are stopped, a 1 is scanned into `stop_on_ext_event` and 1 is scanned into `start_clocks`, clocks will start up and then stop on the next `ss_clock` rising edge after the `ext_event` FF goes active; `stop_even_on_ext_event` is similar to `stop_on_ext_event`, but it causes clocks to stop on the next `sbclk` rising edge after the `ext_event` FF goes active. Thus, clocks will stop either one or two 40-MHz cycles, respectively, after a logic 1 is clocked in on the `ext_event` input. Scan software can scan out the `clocks_stopped` and `sbus_1st_half` CCR bits to determine whether clocks are stopped, and if they are stopped in the first or second half of the 20-MHz `sbclk` cycle.

Figure 9.5 - With `stop_on_ext_event`





Figure 9.6 - With `stop_even_on_ext_event`



Scan software can, by scanning appropriate values into the CCR, XCC, and ext_event_ff while clocks are stopped, cause any number of clock pulses to be issued when clocks are restarted, from 1 on up to the maximum. In the table below, 'tc' is the terminal count value of XCC, and M is any integer greater than 1


Table 9.1 - Clock Control and Scan

N	sbus_1st_half	stop	stop_even	ext_event_ff	XCC
----	---- (scanout) ----	----- (scanin) -----		-----	-----
1	1	1	0	1	tc+1
1	0	0	1	1	tc+1
2	1	0	1	1	tc+1
2	0	1	0	0	tc
3	1	1	0	0	tc
3	0	0	1	0	tc
2*M	1	0	1	0	tc+2-M
2*M	0	1	0	0	tc+1-M
2*M+1	1	1	0	0	tc+1-M
2*M+1	0	0	1	0	tc+1-M

Examples:


N=2, stopped with sbus_1st_half=1: the table tells us to load (tc+1) into the counter, load 1 into ext_event_ff, and to assert stop_even_on_ext_event. Note that (tc+1)=0.

Figure 9.7 - N=2, stopped with sbus_1st_half=1.



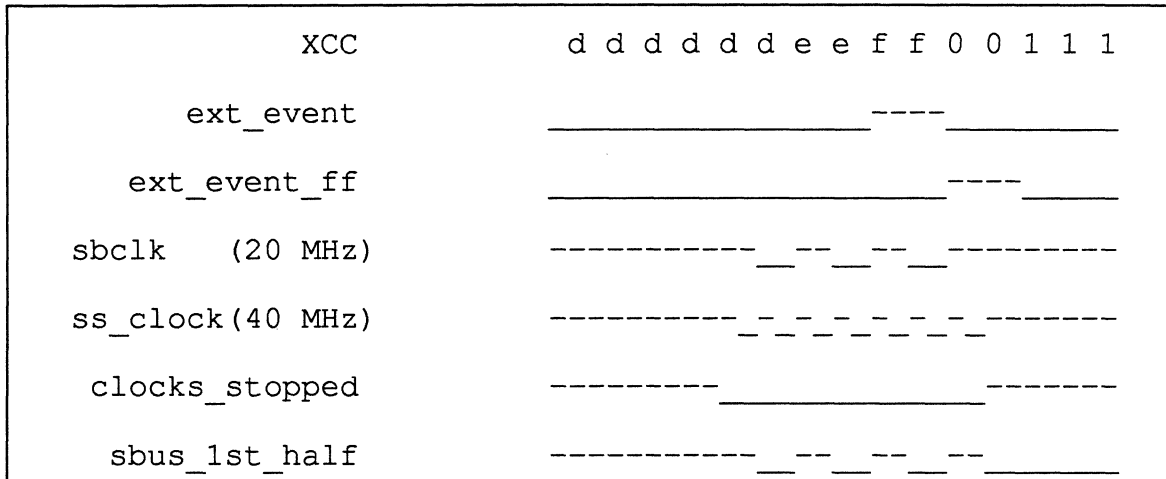
N=7, stopped with sbus_1st_half=0: $7=2*3+1$, so $M=3$. The table tells us to load $(tc+1-3)=-3$ into the counter, load 0 into ext_event_ff, and to assert stop_even_on_ext_event. I'll show -3 as 'd', which is the last hex digit of its 2's-complement representation.

Figure 9.8 - N=7, stopped with sbus_1st_half=0.



$N=7$, stopped with `sbus_1st_half=1`: $7=2*3+1$, so $M=3$. The table tells us to load $(tc+1-3)=-3$ into the counter, load 0 into `ext_event_ff`, and to assert `stop_on_ext_event`. I'll show -3 as 'd', which is the last hex digit of its 2's-complement representation.

Figure 9.9 - $N=7$, stopped with `sbus_1st_half=1`.



Note that the last two examples, taken together in sequence, cause 14 positive edges on `sys_clk` and 7 positive edges on `sbclk`.

9.0.12 Count Clocks After Internal Event

In this mode, the XCC is held until an internal event occurs. The internal event does not stop clocks, but does cause assertion of the `int_event` output; the `int_event` output will remain asserted until it is cleared by scan. The XCC is enabled to count whenever `int_event` is asserted, so clocks will continue to run until `ext_event` is asserted, either by XCC or by another external event detector. The intent of this mode is to issue exactly N clocks after the internal event has occurred. Logic in the clock controller records whether the internal event occurred in the first or second half of the bus cycle, and this information is factored into the subsequent clock stop on external event, so that N can be any even or odd integer. Due to latencies in the logic, N must be greater than or equal to 4.

To support this mode, the XCC must have logic which, under scan control, holds the count when `int_event` is not asserted.

The CCR also needs some logic. The signal `int_event_1st_half` records whether the internal event which caused the assertion of the `int_event`

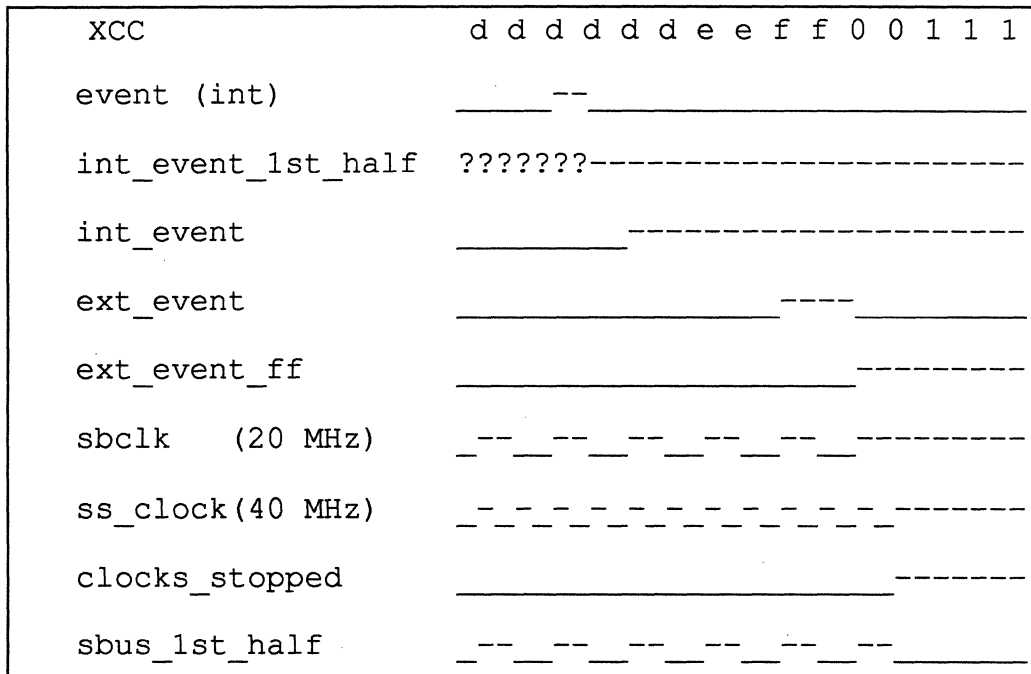
output happened in the first or second half of the SBus cycle. The CCR bit `stop_int_to_ext` will cause an even or odd number of `sys_clk` positive edges to occur after the internal event is detected, depending on whether `int_to_ext_odd` is 0 or 1, respectively. The actual number of clocks issued is $(2*(tc-XCC.before) + 4)$ with `int_to_ext_odd=0`, and $(2*(tc-XCC.before) + 5)$ with `int_to_ext_odd=1`. Logic in the clock controller works as follows when `stop_int_to_ext` is set:

- * if `int_to_ext_odd=0` and `int_event_1st_half=0`, then stop clocks at the end of the SBus cycle in which `ext_event_ff` is asserted (as described above for `stop_even_on_ext_event`);
- * if `int_to_ext_odd=0` and `int_event_1st_half=1`, then stop clocks midway through the SBus cycle in which `ext_event_ff` is asserted (as described above for `stop_on_ext_event`);
- * if `int_to_ext_odd=1` and `int_event_1st_half=1`, then stop clocks at the end of the SBus cycle in which `ext_event_ff` is asserted;
- * if `int_to_ext_odd=1` and `int_event_1st_half=0`, then stop clocks midway through the *next* SBus cycle *after* `ext_event_ff` is asserted.

If, in addition to setting `stop_int_to_ext`, we also set `stop_on_int_event`, then a special mode is enabled. In this mode, as with the simple `stop_int_to_ext` mode described above, the XCC starts counting clocks after the first internal event, and stops clocks when the count is exhausted. In addition, clocks will stop on an internal event as described in section 2.9.2.5, but only if the internal event occurs while the `int_event` microSPARC output pin is asserted. In other words, while in this mode, clocks will stop on the first internal event which occurs while the XCC is counting; if no such internal event occurs, clocks will stop when the count is exhausted.

N=8, event occurs in first half of bus cycle. stop_int_to_ext must be set, int_to_ext_odd must be cleared, and XCC.before must be set to (tc-2), here represented by 'd'. Clocks are stopped in the middle of the sbclk cycle in which ext_event_ff is active. We get eight more ss_clock rising edges than we would have gotten if clocks had been stopped immediately on the internal event.

Figure 9.11 - Event in First half of bus cycle, N=8.



N=9, event occurs in second half of bus cycle. stop_int_to_ext must be set, int_to_ext_odd must be set, and XCC.before must be set to (tc-2), here represented by 'd'. Clocks are stopped in the middle of the *next* sbclk cycle *after* ext_event_ff is active. We get nine more sys_clk rising edges than we would have gotten if clocks had been stopped immediately on the internal event.

Figure 9.13 - Event in Second half of bus cycle, N=9.



9.0.13 Stop Clocks After N Internal Events

In this mode clocks are stopped after the Nth sbclk cycle in which the int_event output is asserted. It is controlled by the stop_on_ext_event CCR bit, and XCC needs a scannable control bit which enables it to count only while int_event is active. To use this mode, we must load XCC with (tc-N) and turn on stop_on_ext_event. Latency will be six 40-MHz cycles if the final internal event occurs in the first half of the sbclk cycle, and five cycles if it occurs in the second half. Note that we are not able to handle more than one event per sbclk cycle.

Figure 9.14 - Event in first half of bus cycle, N=2. Latency=6 cycles.

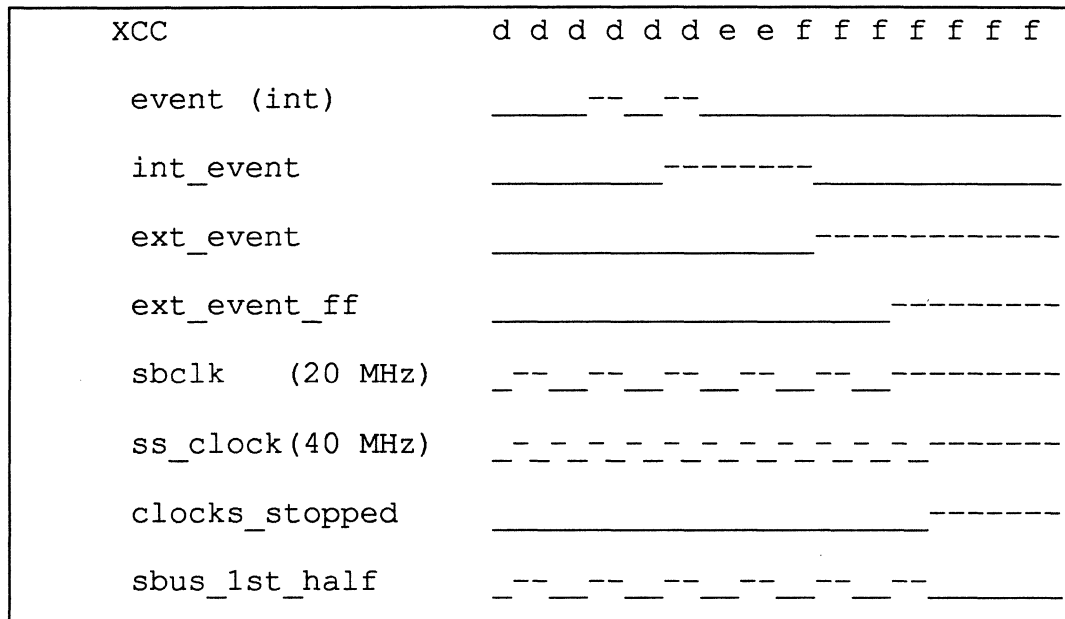



Figure 9.15 - Event in second half of bus cycle, N=2. Latency=5 cycles.



9.0.14 CCR Bits

Here is a list of the Clock Control Register bits. These are accessible by scan only, and their functionality is described above.

- *stop_on_ext_event (Issue N Clocks, Stop Clocks after N Internal Events)
- *stop_even_on_ext_event (Issue N Clocks)
- *stop_int_to_ext (Count Clocks after Internal Event)
- *int_to_ext_odd (Count Clocks after Internal Event)
- *stop_on_int_event (Stop Clocks on Internal Event)
- *stop_clocks (Stopping Clocks, Single-Step)
- *start (Starting Clocks, Single-Step)

9.0.15 JTAG

A variety of microSPARC test and diagnostic functions, including internal scan, boundary scan and clock control, are controlled through an IEEE 1149.1 (JTAG) Standard Test Access Port (TAP).

Commands and data are sent as serial data between the JTAG master and the microSPARC chip (a JTAG slave), via a 4 wire serial testability bus (JTAG bus). The TAP interfaces to the JTAG bus via 5 dedicated pins on the microSPARC chip. These pins are:

TCK - input - test clock

TMS - input - test mode select

TDI - input - test data input

TRST_L - input - JTAG TAP reset (asynchronous)

TDO - output - test data output

For more details on the IEEE protocol, please refer to the IEEE document "IEEE Standard Test Access Port and Boundary-Scan Architecture", published by IEEE.

9.0.16 Board Level Architecture

Typical microSPARC systems will contain several JTAG-compatible chips. These are connected using the minimum (single TMS signal) configuration as described in the 1149.1 specification (Figure 3-1, IEEE 1149.1 standards manual). This configuration contains three broadcast signals (TMS, TCK, and TRST,) which are fed from the JTAG master to all JTAG slaves in parallel, and a serial path formed by a daisy-chain connection of the serial test data pins (TDI and TDO) of all slaves.

The TAP supports a BYPASS instruction which places a minimum shift path (1 bit) between the chip's TDI and TDO pins. This allows efficient access to any single chip in the daisy-chain without board-level muxing.

9.0.17 TAP

The TAP consists of a TAP controller, plus a number of shift registers including an instruction register (IR) and multiple "data" registers.

The TAP controller is a synchronous FSM which controls the sequence of operations of the JTAG test circuitry, in response to changes at the JTAG bus. (Specifically, in response to changes at the TMS input with respect to the TCK input.) Note that the TAP controller is asynchronous with respect to the system clock(s), and can therefore be used to control

the clock control logic. The TAP FSM implements the state (16 states) diagram as detailed in the 1149.1 protocol.

The IR is a 6-bit register which allows a test instruction to be shifted into microSPARC. The instruction is used to select the test to be performed and/or the test data register to be accessed. The supported instructions are listed in a later section.

9.0.18 Data Registers

Although any number of loops may be supported by the TAP, the FSM in the TAP controller only distinguishes between the IR and a data register. The specific data register is decoded from the instruction in the IR.

The following data registers are supported in the microSPARC TAP:

- * Bypass Register - a single bit shift register for efficient board-level scan.
- * Device I.D. Register - a 32-bit register with the following field.

Figure 9.16 - JTAG ID Reg Contents

Ver	Part ID	Manufacturer's ID	Const
31 28 27	12 11		01 00

Field Definitions:

- Version - Bits[31:28] represent the version number which is 0x0 for this version
- Part ID - Bits[27:12] represent part number as assigned by TI, which is 0x0004
- Miff ID - Bits[11:01] represent manufacturer's ID as per JEDEC, which is 0x17
- Const - Bit[00] is tied to a constant logic '1'

Value in ID Register: 32'h0000202f

- * Data registers - A two bit clock control register to sample outputs from the clock controller(CCR)
- * Boundary Scan Register - a single scan chain consisting of all of the boundary scan cells (input, output and inout cells).

* Internal Scan Registers - a single scan chain of all the internal scan f/fs

9.0.19 JTAG Instructions

The following instructions are supported by the microSPARC TAP. The table contains the bit-value and mnemonic, as well as which data register is selected by that instruction. The encodings followed by an "*" are fixed by the IEEE JTAG protocol.

Table 9.2 - JTAG INSTRUCTIONS

value	Name of Instrn	Data register(s)	Scan Chains Accessed
000000*	EXTEST	Boundary Scan Register	Boundary Scan Chain
000001*	SAMPLE	Boundary Scan Register	Boundary Scan Chain
000010	INTEST	Boundary Scan Register	Boundary Scan Chain
000011	ATEINTEST	Boundary Scan Register	Boundary Scan Chain
100000	IDCODE	JTAG ID Register	ID Register Scan Chain
111111 *	BYPASS	Bypass Register	Bypass Register
011110	SEL_CCR	Clock Control Register	Clock Control Register Chain
010000	SEL_INT_SCAN	Internal Scan Register	Internal Scan Chain
011111	SEL_DBG_SCAN	Internal Scan Register	Internal Scan Chain

Note: 1. The two internal scan chain instructions differ with respect to the scan chain clocking during CAPTURE_DR state of the tap fsm. Sel_int_scan will be used for ATPG tests, where a clock pulse is needed to capture the next state when scan_mode signal is in the inactive state between shift cycles. The other scan instruction, Sel_dbg_scan is used during debug to read and write the scan chain. No pulse is generated during the transition from "shift --> capture ---> shift" states. In other words, the scan state is preserved during the shift, capture, shift cycle.

2. The TDO output becomes valid at the falling edge of TCK, per the 1149.1 protocol. This is so, that the TDI input (which is connected TDO

of the preceding component) of the component is stable to be clocked in during the rising edge of TCK.

3. The ATEINTEST operation is used to load the boundary scan f/fs after which, if it enters the 'run_test_idle' state, the JTAG controller will generate a single TCK pulse.

Although, we have the capability to single step the chip thru another mechanism (using sys_clock itself), ATEINTEST option provides the capability to perform ICT on the ATE, perhaps at slow speed.

4. The INTEST operation has been added so that it can be used in conjunction with the SEL_INT_SCAN instruction to perform the ATPG test using scan tool. This instruction will not generate any extra clock pulse in run_test_idle state. This is used primarily to load the boundary scan chain.

5. The Sel_CCR is used to sample two bits (stopped, sbus_1st_half) from the clock controller block. These two bits are synchronized (2 stage synchronizer using TCK) before being sampled during the shift-DR state.

9.0.20 JTAG Interface to MISC

The JTAG block provides two key signals to the clock controller section, two signals directly to the microSPARC core and a five wire control signal to the boundary scan f/fs.

Clock Controller Interface:

Testclk and Testclken are the two signal that are generated in the JTAG block and sent to the clock controller.

Testclken is an active high signal that switches the ss_clock (the 40MHz) to the core from the normal 40MHz clock to the Testclk. This happens only for certain JTAG instructions. They are:

sel_int_scan, sel_dbg_scan, intest, ateintest

For all other instructions (extest, sample, bypass, idcode, sel_ccr) testclken remains inactive thus enabling the normal 40 MHz clock to microSPARC core. The Testclken signal is synchronized inside the clock controller using the free_20MHz clock. By design Testclken is generated to be active at least three TCK cycles before the Testclk signal becomes active. Testclken signal changes state only after transition

from update_IR of the instruction scan cycle, on the positive edge of TCK. Testclken signal becomes inactive after transition to tap_logic_reset state on the falling edge of TCK.

Testclk is a gated version of TCK and the gating signals are sel_instruction and shift (function of shift_DR) and capture (capture-DR) states. Testclk toggles only during sel_int_scan and sel_dbg_scan instructions.

microSPARC Core Interface:

Sys_sen (ss_scan_mode) and tg_strobe are two signals that go directly to the core of microSPARC. Scan_mode signal is active high whenever the Tap enters any of the four DR states, shift,exit1,pause and exit2. During the last three state, Testclk will not toggle and the state of the f/f remains the same as the last bit scanned in during the shift state. It is necessary to activate the scan_mode signal during these three states, so that tri-states would remain disabled during repeat scan after going thru exit1, pause, exit2 states. Sys_sen is a registered signal that is clocked on the falling TCK. This has been done to avoid race conditions between the scan_mode signal and the shift clock(testclk) during the shortest tap state traversal from select-DR to shift-DR.

Since the Sys_sen is a heavily loaded (goes to all f/fs in the chip) signal, it may have a longer rise time and not meet the setup time requirement for the shortest tap state traversal from select-DR to shift-DR. In such a case, the TCK should not be run at greater than 5 MHz.

The tg_strobe signal is low going pulse that is used as a self-timing trigger for the megacells. It is generated during the update-DR state and adheres to the timing specified in the megacell document.

Boundary Control Interface:

The five wire boundary control signal corresponds to: bin_cap, bout_cap, b_sen, b_uen, b_mode.

bin_cap and bout_cap are generated during the capture-DR state and are used to load the value on the pins or the output of the core to the boundary scan f/f. b_sen is generated on the falling edge of the tck (to avoid race conditions) and is used as a scan_en signal for the boundary scan f/f. b_uen is an update signal for the boundary scan update latch and it happens at the falling edge of tck.

b_mode is a mux control signal that selects between the direct pin input and the value in the update latch. This signal will change

during the update-IR state and when the tap goes back to test-logic-reset state on the falling edge of TCK.

RESET Mechanism:

We also have a independent TRST_L signal which when active low would set the TAP into the tap_logic_reset state. This signal will asynchronously set the tap state machine to the tap_logic_reset state. It adheres to the 1149.1 IEEE protocol with respect to the initialization thru reset mechanism. There is no minimum active time requirement on this reset signal. If the board is not going to have an extra oscillator for TCK, then the JTAG reset pin (TRST_L) can be tied to an active low signal thus disabling JTAG operations in the chip.

The TDI and TMS inputs have pullups on the pad and when left unconnected will be equivalent to a signal value '1' on these pins. With a free running TCK, it would guarantee that the TAP would get into the tap_logic_reset state at the end of five TCKs.

9.0.21 JTAG Operation

The following are some of the basic operations which, when combined together will enable the user to run any of the JTAG instructions specified above. They are provided here just for understanding the TAP state transitions during various JTAG operations.

We will only be concerned with JTAG I/O, i.e. TCK, TMS, TDI, TRST and TDO. The first four are inputs and the last one is the output. All five are chip I/O. The other inputs to the chip are either in a don't care state or in a predetermined state. They shouldn't affect the operation of the JTAG controller. It should be noted, that, for a more robust operation of the chip, we should follow a proper procedure with regard to getting in and out and back to JTAG operations. (for instance resetting the system before and after JTAG operations. Once we are in the tap_logic_reset state, all outputs from JTAG become inactive and the chip should be back to normal functional mode.)

The tap state encodings (in hex) are as follows:

f-test-logic-reset, c-run-test-idle, 7-select-DR, 6-capture-DR, 2-shift-DR, 1-exit1-DR, 3-pause-DR, 0-exit2-DR, 5-update-DR, 4-select-IR, e-capture-IR, a-shift-IR, 9-exit1-IR, b-pause-IR, 8-exit2-DR, d-update-IR.

In order to run the JTAG instructions, we do the following tap state traversal for the various sub tasks:

Instruction Scan:

f --> c --> 7 --> 4 --> e --> 9 --> b --> 8 --> a (for 6 clocks) --> 9

(the opcode is shifted thru tdi while in the shift-IR state)

Data Scan:

9 --> b --> 8 --> d --> c --> 7 --> 6 --> 1 --> 3 --> 0 --> 2 (# of shifts equal to length of scan chain) --> 1

(At state 'd' the decode instruction is latched on the falling edge of tck. Data is shifted into appropriate data register during shift cycle and at the end of shift exit to exit1-DR(1) state.

Return to new instruction:

2 --> 1 --> 3 --> 0 --> 5 --> c

(we will wait in state c (run-test-idle) and go back to instruction scan as shown above.)

Figure 9.17 - JTAG LOGIC BLOCK DIAGRAM

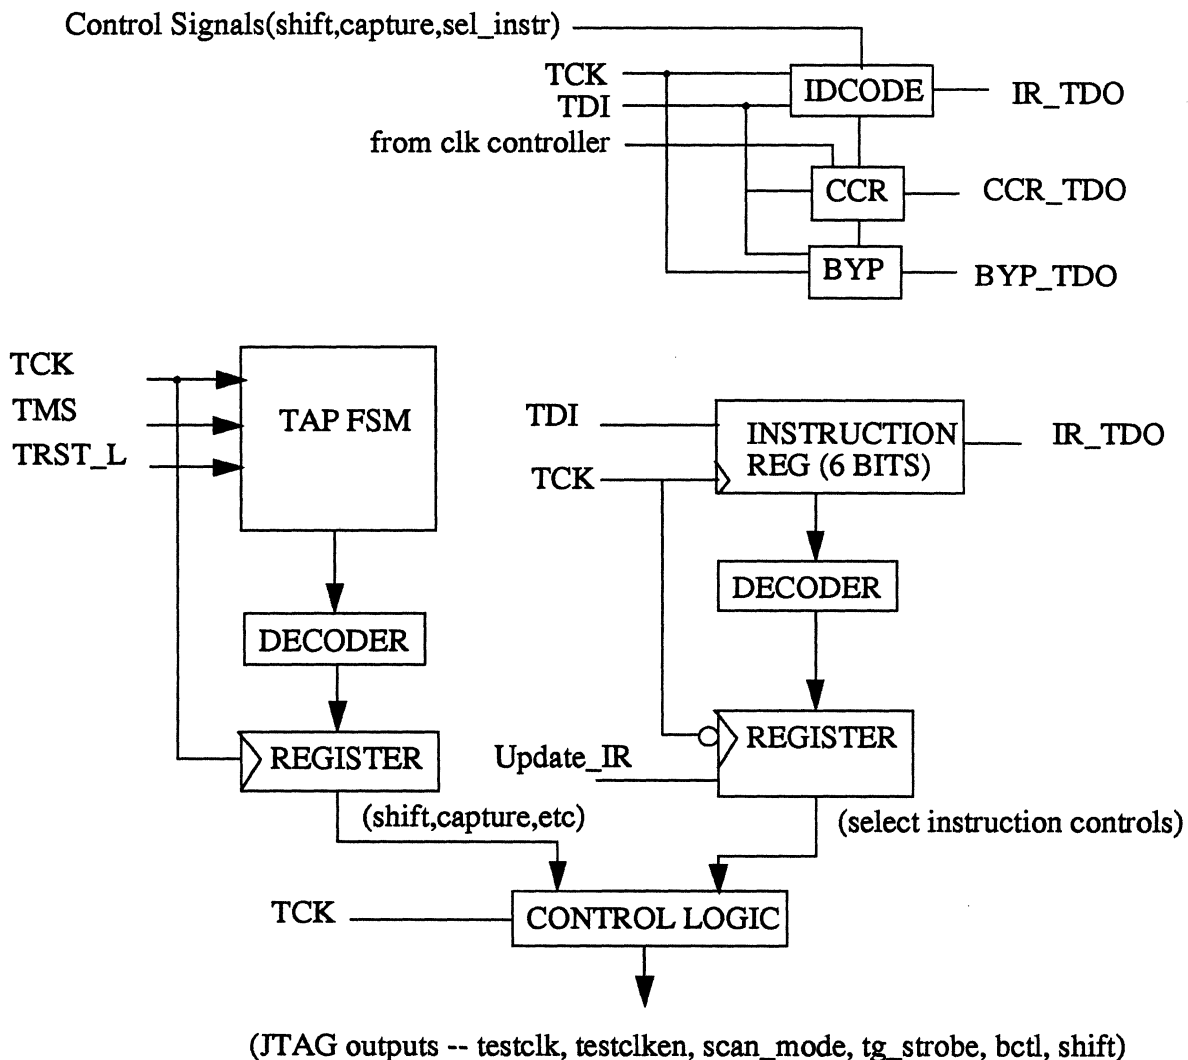



Figure 9.18 - microSPARC JTAG DATA & INSTRUCTION REGISTERS



10.0 Error Handling

The microSPARC CPU must detect and handle many kinds of errors and exceptions. The SPARC IU is interrupted by some type of trap in all CPU error cases. DMA masters other than the CPU should cause their own IU trap via the SBus interrupt mechanism. Physical address references to nonexistent addresses in any address space will either return garbage or cause timeouts. The following preliminary list attempts to describe what happens under various circumstances.

Table 10.1 - Error Summary

Error	Initiator	Result Summary	
Memory Parity Error	Instruction Memory Access	set PE, FT=5, L, AT in SFSR cause Instruction Access Error trap (D stage + 1)	
	IU, FPU Read Memory Access	set PE, ERR, CP, TYPE in MFSR save PA in MFAR cause L15 interrupt	
	IU, FPU Write Byte, Half- word Memory Access (Read-modify-write)	set PE, ERR, CP, TYPE in MFSR save PA in MFAR cause L15 interrupt	
	(Translation Error)	Tablewalk on Instruction Memory Access	set PE, FT=4, L, AT in SFSR cause Instruction Access Error trap (D stage)
	(Translation Error)	Tablewalk on IU, FPU Data Memory Access	set PE, FT=4, L, AT, FAV in SFSR save iu_dva in SFAR cause Data Access Error trap (R stage)
		IO DMA Read Memory Access	return SBus Error Acknowledge set PE, ERR in MFSR save PA in MFAR, cause L15 interrupt
		IO DMA Write Byte, Half- word Memory Access (Read-modify-write)	return SBus Error Acknowledge set PE, ERR in MFSR save PA in MFAR, cause L15 interrupt
SBus Controller Time Out	Tablewalk on IO DMA Memory Access	return SBus Error Acknowledge set PE, ERR in MFSR save PA in MFAR, cause L15 interrupt	
	CPU SBus Read Access	set TO, FT=5, FAV in SFSR save iu_dva in SFAR cause Data Access Error trap (R stage)	
	CPU SBus Write Access	set TO, ERR, SIZE, ~RD, FAV in AFSR save PA in AFAR cause L15 interrupt	
SBus Late Error (Ack)	IO DMA Access	return SBus Error Acknowledge	
	CPU SBus Read Access	set LE, ERR, SIZE, RD, FAV in AFSR	
	CPU SBus Write Access	set LE, ERR, SIZE, FAV(sometimes) in AFSR	

Error	Initiator	Result Summary
SBus Error Acknowledge	CPU Read Access	set BE, FT=5, FAV in SFSR save iu_dva in SFAR cause Data Access Error trap (R stage)
Invalid Address Error	CPU Write Access	set BE, ERR, SIZE, ~RD, FAV in AFSR, save PA in AFAR cause L15 interrupt using CP_STAT
Translation Error	IO DMA PTE Access (IO PTE V bit = 0)	return SBus Error Acknowledge
	ET=0 during Tablewalk on Instruction Memory Access	set FT=1, L, AT in SFSR cause Instruction Access Exception trap (D stage)
	ET=0 during Tablewalk on IU, FPU Data Memory Access	set FT=1, L, AT, FAV in SFSR save iu_dva in SFAR cause Data Access Exception trap (R stage)
	ET=3 during Tablewalk on Instruction Memory Access	set FT=4, L, AT in SFSR cause Instruction Access Error trap (D stage)
	ET=3 during Tablewalk on IU, FPU Data Memory Access	set FT=4, L, AT, FAV in SFSR save iu_dva in SFAR cause Data Access Error trap (R stage)
Control Space Error	CPU Invalid ASI Access	set FT=5, L, FAV, CS in SFSR save iu_dva in SFAR cause Data Access Exception trap (R stage)
	CPU Invalid Size of Access	set FT=5, L, FAV, CS in SFSR save iu_dva in SFAR cause Data Access Exception trap (R stage)
	CPU Invalid Virtual Address during ASI requiring VA	set FT=5, L, FAV, CS in SFSR save iu_dva in SFAR cause Data Access Exception trap (R stage)
Privilege Violation Error (S bit and not ACC 6,7)	IU Instruction Memory Access	set FT=3, L, AT in SFSR cause Instruction Access Exception trap (D stage)
Privilege Violation Error (ACC and ASI checked)	IU, FPU Data Memory Access	set FT=3, AT, FAV in SFSR save iu_dva in SFAR cause Data Access Exception trap (R stage)
Protection Error (Memory page ACC and the ASI are checked)	IU, FPU Data Memory Access	set FT=2, L, AT, FAV in SFSR save iu_dva in SFAR cause Data Access Exception trap (R stage)
Protection Error (Memory page ACC is checked)	IU, FPU Data Memory Access	set FT=2, L, AT, FAV in SFSR cause Instruction Access Exception trap (D stage)
Protection Error (Write to read only page)	IO DMA Write	return SBus Error Acknowledge

11.0 ASI Map

This chapter describes the microSPARC ASI map. The Address Space Identifier (ASI) is appended to the virtual address by the SPARC IU when it accesses memory. The ASI encodes whether the processor is in supervisor or user mode, whether an access is to instruction or data memory, and is used to perform other internal cpu functions.

11.0.1 Overview

The table below lists all of the ASI values supported in a microSPARC system. Only the least significant 6 bits of the ASI are decoded.

Table 11.1 - ASI's Supported by microSPARC

ASI	Function	Access	Size
00	Reserved	-	-
01-02	Unassigned	-	-
03	Ref MMU Flush/Probe	Read/Write	Single
04	MMU Registers	Read/Write	Single
05	Unassigned	-	-
06	Ref MMU Diagnostics	Read/Write	Single
07	Unassigned	-	-
08	User Instruction	Read/Write	All
09	Supervisor Instruction	Read/Write	All
0A	User Data	Read/Write	All
0B	Supervisor Data	Read/Write	All
0C	Instruction Cache Tag	Read/Write	Single
0D	Instruction Cache Data	Read/Write	Single
0E	Data Cache Tag	Read/Write	Single
0F	Data Cache Data	Read/Write	Single
10-14	Unassigned	-	-
15-16	Reserved	-	-
17-1C	Unassigned	-	-
1D-1E	Reserved	-	-
1F	Unassigned	-	-
20	Ref MMU Bypass	Read/Write	All
21-2F	Reserved	-	-
30-35	Unassigned	-	-
36	Instruction Cache Flash Clear	Write	Single
37	Data Cache Flash Clear	Write	Single
38	Unassigned	-	-
39	Data Cache Diagnostic Register Access	Read/Write	Single
3A-3F	Unassigned	-	-
40-FF	Reserved	-	-

ASI Descriptions:

ASI=0x00

Reserved - This space is architecturally reserved.


ASI=0x01-0x02

Unassigned - This space is unassigned and may be used in the future.

ASI=0x03

Ref MMU Flush/Probe - This space is used for a flush or probe operation. The Virtual Address is decoded as follows.

Figure 11.1 - TLB Flush or Probe Address Format



Field Definitions:

Virtual Flush or Probe Address (VFPA) - This field is the address that is used to index into TLB. Depending on the type of flush or probe not all 20 bits are significant.

Type - This field specifies the extent of the flush or the level of the entry probed.

Reserved - These bits are ignored. They should be set to zero.

A flush is caused by a single STA instruction and a probe by a single LDA instruction.

Flushes are used to maintain TLB consistency by conditionally removing one or more page descriptors. These conditions vary as shown. Note that any TLB flush also flushes the ITBR automatically.

Table 11.2 - TLB Entry Flushing

VA[11:08]	Flush	PTE Match Criteria
0	Page	(Level 3) AND (Context match OR ACC=6-7) AND VA[31:12] match
1 to 4	Entire	None (Entire TLB Flush)
5 to F	Reserved	

Probes cause the MMU to perform a table walk stopping when a PTE has been reached as shown.

Table 11.3 - CPU TLB Entry Probing

VA[11:08]	Probe	Returned Data
0	Page	Level 3 PTE or 0
1	Segment	Level 2 PTE or 0
2	Region	Level 1 PTE or 0
3	Context	Level 0 PTE or 0
4	Entire	PTE from Table Walk or 0
5 to F	Reserved	

ASI=0x04

MMU Registers - This space is used to read and write internal MMU registers using the Virtual Address to reference them. Single word accesses only should be used, others result in an error.

Table 11.4 - Address Map for MMU Registers

VA[12:08]	Register
00	Control Register
01	Context Table Pointer Register
02	Context Register
03	Synchronous Fault Status Register
04	Synchronous Fault Address Register
05-0F	Reserved
10	TLB Replacement Control Register
11-12	Reserved
13	Synchronous Fault Status Register**
14	Synchronous Fault Address Register**
15-1F	Reserved
	**Writeable for diagnostic purposes

VA bits [31:13] are zero. VA bits [07:00] are ignored and should be set to zero by software.

ASI=0x05

Unassigned - This space is unassigned and may be used in the future.

- ASI=0x06** **Ref MMU Diagnostics** - Diagnostic reads and writes can be made to the 32 TLB entries and the Instruction Translation Buffer Register using the virtual address to specify which entry and whether the PTE or Tag section is to be referenced.
- ASI=0x07** **Unassigned** - This space is unassigned and may be used in the future.
- ASI=0x08** **User Instruction** - This space is defined and reserved by SPARC for user instructions.
- ASI=0x09** **Supervisor Instruction** - This space is defined and reserved by SPARC for supervisor instructions.
- ASI=0x0A** **User Data** - This space is defined and reserved by SPARC for user data.
- ASI=0x0B** **Supervisor Data** - This space is defined and reserved by SPARC for supervisor data.
- ASI=0x0C** **Instruction Cache Tag** - This space is used for reading and writing instruction cache tags by using the LDA and STA instructions at virtual addresses in the range of 0x0 to 0x0FFF on modulo-32 boundaries.

Figure 11.2 - Instruction Cache Tag Entry

Rsvd	IPA Tag[26:12]	Rsvd	Valid
31 27 26		12 11	01 00

Bits [31:27,11:01] are not implemented, should be written as 0 and will be read as 0.

- ASI=0x0D** **Instruction Cache Data** - This space is used for reading and writing instruction cache data by using the LDA and STA instructions at virtual addresses in the range of 0x0 to 0x0FFF.

ASI=0x0E

Data Cache Tag - This space is used for reading and writing data cache tags by using the LDA and STA instructions at virtual addresses in the range of 0x0 to 0x03FF on modulo-16 boundaries.

Figure 11.3 - Data Cache Tag Entry

Reserved	PA Tag[26:11]	Reserved	Valid
31	27 26	11 10	01 00

Bits [31:27,10:01] are not implemented, should be written as 0 and will be read as 0.

ASI=0x0F

Data Cache Data - This space is used for reading and writing data cache data by using the LDA and STA instructions in ASI 0xF at virtual addresses in the range of 0x0 to 0x03FF.

ASI=0x10-0x14

Unassigned - This space is unassigned and may be used in the future.

ASI=0x15-0x16

Reserved - This space is architecturally reserved.

ASI=0x17-0x1C

Unassigned - This space is unassigned and may be used in the future.

ASI=0x1D-0x1E

Reserved - This space is architecturally reserved.

ASI=0x1F

Unassigned - This space is unassigned and may be used in the future.

ASI=0x20

Ref MMU Bypass - This space can be used to access an arbitrary physical address. It is particularly useful before the MMU or main memory have been initialized. The MMU does not perform an address translation rather a physical address is formed from the least significant 31 bits of the Virtual Address (PA[30:00] := VA[30:00]). Accesses in bypass mode are not cacheable.

ASI=0x21-0x2F

Reserved - This space is architecturally reserved.

- ASI=0x30-0x35** **Unassigned** - This space is unassigned and may be used in the future.
- ASI=0x36** **Instruction Cache Flash Clear** - The instruction cache is completely flushed by any type of alternate store instruction to this ASI. All instruction cache valid bits are reset (to zero) by this operation. Note that the pipeline is NOT flushed by this sta as it would be on a SPARC FLUSH instruction.
- ASI=0x37** **Data Cache Flash Clear** - The data cache is completely flushed by any type of alternate store instruction to this ASI. All data cache valid bits are reset (to zero) by this operation.
- ASI=0x38** **Unassigned** - This space is unassigned and may be used in the future.
- ASI=0x39** **Data Cache Diagnostic Register Access** - This space is used to read and write the internal Data Cache Registers. `iu_dva[08]` is also used to select from between WRB0 and WRB1. Single word accesses only should be used, others result in an internal error. The Virtual Address map to these registers:
- Table 11.5 - Address Map for Data Cache Registers**
- | VA[08] | Register |
|--------|----------------|
| 0 | Write Buffer 0 |
| 1 | Write Buffer 1 |
- VA bits [31:09] are zero. VA bits [07:00] are ignored and should be set to zero by software.
- ASI=0x3A-0x3F** **Unassigned** - This space is unassigned and may be used in the future.
- ASI=0x40-0xFF** **Reserved** - Since the 2 high order bits are not decoded these encodings should not be used.

2.16 References

1. The SPARC Architecture Manual, Version 8, of December 11, 1990
2. SBus Specification, Rev. B.0, 1990
3. Texas Instruments TMS390Z50 SuperSPARC Datasheet. SPKS011 Version, August, 1992.
4. Texas Instruments TMS390Z50 SuperSPARC User Guide.
5. Texas Instruments TMS390S10 microSPARC Datasheet. Version August 21, 1992.

TEXAS INSTRUMENTS

**ERRATA FOR TMS390S10 DATASHEET
(VERSION 21 AUG 92)**

CREATED : 8th Dec 1992

LAST UPDATED : 6th Jan 1993

TEXAS INSTRUMENTS

(8th Dec 1992)

Page 8 should read :-

"MPAR0 is for MDATA0 - MDATA31"
"MPAR1 is for MDATA32 - MDATA64"

(8th Dec 1992)

Page 11 The Text describing the N/C signals on pins 227,229,230 & 238 should be deleted. These pins should not be tied low. They are correctly described on page 8 as no connects.

(6 Jan 93)

Page 31 The description of the CAS lines in TABLE 13 should read :-

"CAS0 data word[63:32]"
"CAS1 data word [31:0]"

(8th Dec 1992)

Page 34 A block diagram showing the SBus address mapping should be added to this part of the datasheet. The diagram should contain the following data :-

TMS390S10 Physical Memory Map.

0x70000000	SLOT4
0x60000000	SLOT3
0x50000000	SLOT2
0x40000000	SLOT1
0x30000000	SLOT0
0x20000000	UNUSED
0x10000000	CONTROL SPACE
0x08000000	UNUSED
0x00000000	SYSTEM MEMORY

(8th Dec 1992)

Page 91 Add a line the Recommended Operating Conditions Table.

"VIH_IN_CLK High Level Input Voltage on IN_CLK 2.4V(MIN)

TEXAS INSTRUMENTS

(8th Dec 1992)

Page 92 Replace the description of Parameter 1 which currently says :-

"tac(MDATA) Access Time from ~MCAS(going low)" with "*tsu(MDATA) Setup time to REF_CLK(going high)"

Replace the min value for parameter 1 which is "30.0" with "8.0"

Replace the description of Parameter 2 which currently says :-

"tac(MPAR) Access Time from ~MCAS(going low)" with "*tsu(MPAR) Setup time to REF_CLK(going high)"

Replace the min value for parameter 2 which is "30.0" with "8.0"

Parameters 3 and 4 have "~MCAS(going low)" in the description, replace this with "~MCAS(going high)".

Change the minimum value for parameters 3 and 4 from "0.0" to "2.0"

Add the following text under the box containing parameters 1-4:-

"(*) These parameters imply (i) tac(MDATA), Access time from ~MCAS(falling edge) is a minimum of 30.0nS and (ii) tac(MPAR), Access time from ~MCAS(falling edge) is a minimum of 30.0nS, but this is not explicitly tested."

(6th Jan 1993)

Page 94 Change the min value of parameter 49 from 3.5 to 2.0 nS.

Change the min value of parameter 50 from 3.5 to 0.5 nS.

Change the min value of parameter 51 from 3.5 to 0.5 nS.


Add a "*" in the parameter description column of parameters 49, 50 & 51.

Under the "SBus interface switching table" add the lines :-

"(*) The SBus specification calls for a minimum hold time of 2.5nS for these parameters. Test equipment constraints result in a guaranteed minimum hold time which is less than 2.5nS"

(8th Dec 1992)

Page 97 In both parts of figure 23 replace all occurrences of "0.4V" with "0.2V". Replace all occurrences of "2.4V" with "4.0V".





microSPARC Data Sheet



Use the Texas Instruments microSPARC Data Sheet to define the operating parameters for the microprocessor, and for any firmware/software development that you require.

The TI documentation included here is accurate as of the date of release of the SPARCengine EC OEM Technical Manual. Please call Texas Instruments to ensure that you have the most current documentation. Please use caution in developing plans on this information until you confirm it is the latest information available.



- **High Performance SPARC Processor**
 - 50 MHz Operating Frequency
 - High Speed Floating Operation (>10 MFLOPS peak at 50MHz)
 - High Speed Integer Operation (>36 MIPS at 50MHz)
- **High Integration**
 - SPARC Integer Unit
 - SPARC Reference Memory Management Unit
 - MEIKO Floating Point Unit
 - Instruction Cache (4 KBytes)
 - Data Cache (2 KBytes)
 - SBus Controller supports up to five external SBus devices (SBus specification A.2)
 - Memory Controller up to 128MBytes of DRAM
- Full JTAG Testability (IEEE 1149.1)
- SPARC Version 8 Compatible
- TAB packaging for Low Cost and High Density Board Assembly
- Single 5-V Supply
- 0.8µm CMOS Technology
- Operating Temperature Range ...0°C to 70°C

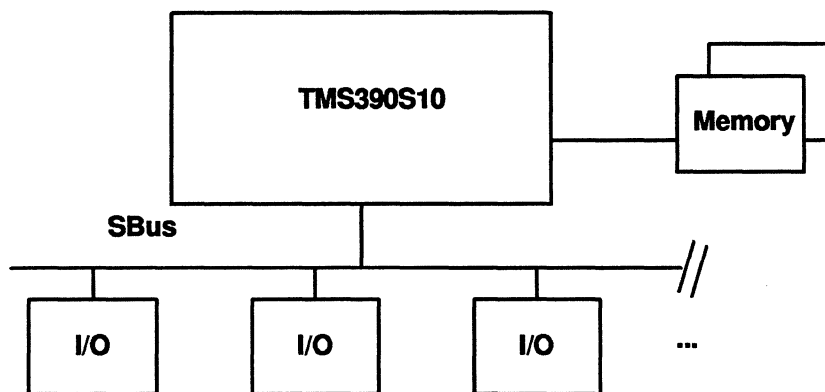



Figure 1. Typical System Configurations

PRODUCT PREVIEW

Table 1. Pin Assignments

TAB	SIGNAL	TAB	SIGNAL	TAB	SIGNAL	TAB	SIGNAL	TAB	SIGNAL	TAB	SIGNAL
1	VCC	49	MDATA[33]	97	MDATA[05]	145	GNDC	193	GNDP	241	JTDI
2	MPAR[1]	50	GNDP	98	MDATA[04]	146	SBDATA[16]	194	SBACK[1]	242	GNDP
3	MDATA[63]	51	MDATA[32]	99	MDATA[03]	147	SBADDR[16]	195	SBACK[0]	243	JTDO
4	MDATA[62]	52	VCCP	100	VCCP	148	SBDATA[15]	196	SBRQ[4]	244	MADDR[11]
5	VCCP	53	VCCC	101	MDATA[02]	149	SBADDR[15]	197	SBRQ[3]	245	Reserved
6	MDATA[61]	54	MPAR[0]	102	GNDP	150	SBDATA[14]	198	GNDC	246	JTMS
7	GNDP	55	GNDC	103	MDATA[01]	151	SBADDR[14]	199	SBRQ[2]	247	JTRST
8	MDATA[60]	56	MDATA[31]	104	MDATA[00]	152	GNDP	200	VCC	248	RESET
9	MDATA[59]	57	MDATA[30]	105	SBDATA[31]	153	SBDATA[13]	201	SBRQ[1]	249	JTCK
10	MDATA[58]	58	GNDP	106	VCC	154	SBADDR[13]	202	SBRQ[0]	250	VCC
11	MDATA[57]	59	MDATA[29]	107	SBDATA[30]	155	SBDATA[12]	203	SBLEERR	251	IN_CLK
12	GNDP	60	MDATA[28]	108	GNDC	156	SBADDR[12]	204	SBSEL[4]	252	GNDC
13	MDATA[56]	61	MDATA[27]	109	GNDP	157	SBDATA[11]	205	SBSEL[3]	253	GNDP
14	VCCP	62	VCCP	110	SBDATA[29]	158	SBADDR[11]	206	GNDP	254	SBCLK
15	MDATA[55]	63	MDATA[26]	111	VCCP	159	VCCP	207	SBSEL[2]	255	VCCP
16	MDATA[54]	64	GNDP	112	SBDATA[28]	160	SBDATA[10]	208	VCCP	256	IRQ[3]
17	VCC	65	MDATA[25]	113	SBDATA[27]	161	VCC	209	SBSEL[1]	257	IRQ[2]
18	MDATA[53]	66	MDATA[24]	114	VCC	162	SBADDR[10]	210	SBSEL[0]	258	IRQ[1]
19	GNDC	67	MDATA[23]	115	SBADDR[27]	163	GNDC	211	SBSIZE[2]	259	IRQ[0]
20	GNDP	68	MDATA[22]	116	GNDC	164	GNDP	212	GNDP	260	MADDR[10]
21	MDATA[52]	69	GNDP	117	SBDATA[26]	165	SBDATA[09]	213	SBSIZE[1]	261	MADDR[09]
22	MDATA[51]	70	MDATA[21]	118	SBADDR[26]	166	SBADDR[09]	214	SBSIZE[0]	262	GNDP
23	MDATA[50]	71	VCCP	119	GNDP	167	SBDATA[08]	215	SBREAD	263	MADDR[08]
24	VCCP	72	VCC	120	SBDATA[25]	168	SBADDR[08]	216	GNDC	264	VCCP
25	MDATA[49]	73	GNDC	121	SBADDR[25]	169	SBDATA[07]	217	VCC	265	MADDR[07]
26	GNDP	74	MDATA[20]	122	VCC	170	SBADDR[07]	218	VCCP	266	MADDR[06]
27	MDATA[48]	75	MDATA[19]	123	SBDATA[24]	171	SBDATA[06]	219	SBGR[4]	267	MADDR[05]
28	MDATA[47]	76	MDATA[18]	124	GNDC	172	SBADDR[06]	220	GNDP	268	GNDP
29	MDATA[46]	77	GNDP	125	SBADDR[24]	173	GNDP	221	SBGR[3]	269	MADDR[04]
30	VCC	78	MDATA[17]	126	SBDATA[23]	174	GNDC	222	SBGR[2]	270	MADDR[03]
31	MDATA[45]	79	MDATA[16]	127	SBADDR[23]	175	SBDATA[05]	223	SBGR[1]	271	MADDR[02]
32	GNDC	80	MDATA[15]	128	VCCP	176	VCC	224	SBGR[0]	272	VCCP
33	GNDP	81	VCCP	129	SBDATA[22]	177	SBADDR[05]	225	GNDP	273	VCC
34	MDATA[44]	82	MDATA[14]	130	SBADDR[22]	178	VCCP	226	EXT_EVENT	274	MADDR[01]
35	VCCP	83	GNDP	131	GNDP	179	SBDATA[04]	227	N/C	275	GNDC
36	MDATA[43]	84	MDATA[13]	132	SBDATA[21]	180	SBADDR[04]	228	VCCP	276	GNDP
37	MDATA[42]	85	MDATA[12]	133	SBADDR[21]	181	SBDATA[03]	229	N/C	277	MADDR[00]
38	MDATA[41]	86	GNDC	134	SBDATA[20]	182	SBADDR[03]	230	N/C	278	MRAS[3]
39	GNDP	87	MDATA[11]	135	SBADDR[20]	183	SBDATA[02]	231	N/C	279	MRAS[2]
40	MDATA[40]	88	VCC	136	SBDATA[19]	184	SBADDR[02]	232	CP_STAT[1]	280	MRAS[1]
41	MDATA[39]	89	MDATA[10]	137	SBADDR[19]	185	GNDP	233	GNDP	281	GNDP
42	MDATA[38]	90	GNDP	138	SBDATA[18]	186	SBDATA[01]	234	CP_STAT[0]	282	MRAS[0]
43	VCCP	91	MDATA[09]	139	SBADDR[18]	187	SBADDR[01]	235	INT_EVENT	283	VCCP
44	MDATA[37]	92	VCCP	140	GNDP	188	SBDATA[00]	236	REF_CLK	284	MCAS[1]
45	GNDP	93	MDATA[08]	141	SBDATA[17]	189	SBADDR[00]	237	GNDC	285	MCAS[0]
46	MDATA[36]	94	MDATA[07]	142	SBADDR[17]	190	SBAS	238	N/C	286	MWE
47	MDATA[35]	95	MDATA[06]	143	VCCP	191	VCCP	239	VCC	287	GNDP
48	MDATA[34]	96	GNDP	144	VCC	192	SBACK[2]	240	VCCP	288	GNDC

PRODUCT PREVIEW

PRODUCT PREVIEW documents contain information on products in the formative or design phase of development. Characteristic data and other specifications are design goals. Texas Instruments reserves the right to change or discontinue these products without notice.



description

The TMS390S10 is a highly integrated SPARC RISC single chip processor. This high performance and low cost processor consists of Integer Unit (IU), MEIKO Floating Point Unit (FPU), Instruction and Data Caches, Memory Management Unit (MMU), and interface circuits. The TMS390S10 has a performance of over 36 Million Instructions Per Second (MIPS) and 10+ MFLOPS peak. The TMS390S10 IU contains 120 general purpose registers (*r registers*) (7-window register file and 8-global registers) and 32 32-bit floating point registers (*f registers*). The TMS390S10 employs Harvard-architecture style of individual bus access to each cache, implemented using a five-stage pipeline to support a single cycle access of the data cache (single-word LOAD).

The TMS390S10 MMU provides a SPARC reference MMU, as specified by the SPARC Reference MMU Architecture and an I/O MMU. Additionally, internal arbitration logic controls access by the processor, and by I/O DMA to memory and I/O devices.

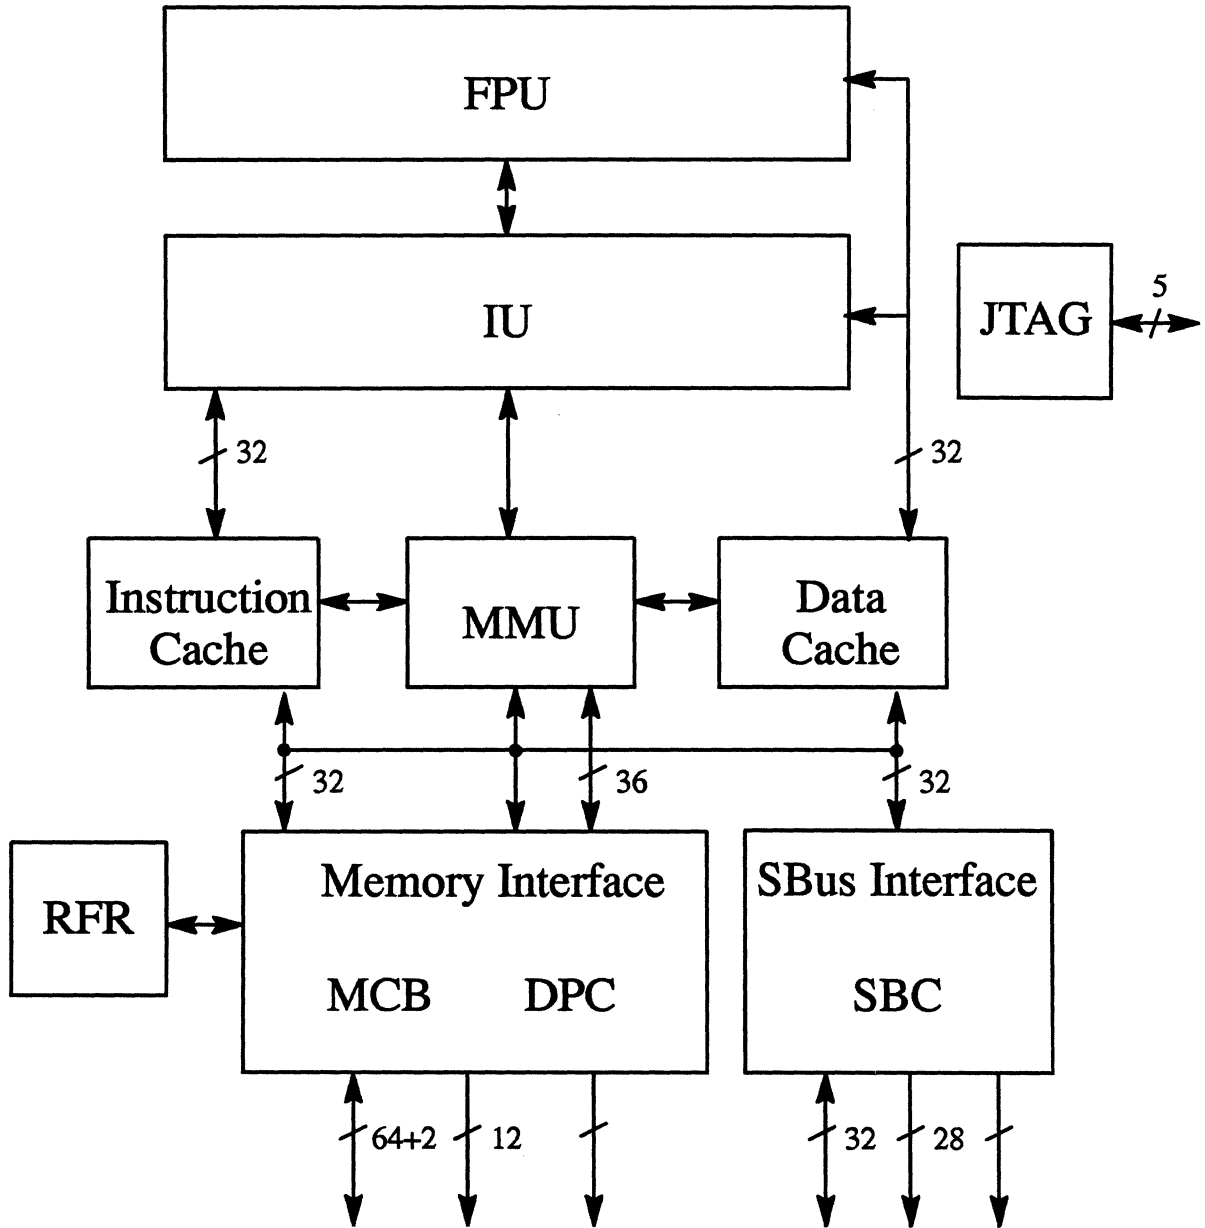
The TMS390S10 has two internal cache systems: 4-KBytes of Instruction Cache and 2-KBytes of Data Cache. The 2-KByte Data Cache is a direct-mapped, physically-addressed, write-through cache with no write allocate. This data store is organized as 128 lines of 16 bytes of data. The 4KByte instruction Cache is a physically-addressed cache. This Instruction store is organized as 128 lines of 32 bytes of data.

The system interfaces include an SBus interface and a Memory interface. The TMS390S10 SBus controller can support up to five SBus devices. The TMS390S10 memory interface can support a maximum of 128 MBytes in four banks for the system memory with using 16-Mbit DRAMs. When using 4-Mbit DRAMs, the TMS390S10 can support up to 32 MBytes of system memory. The TMS390S10 memory interface provides complete control/data signals for memory system with 64-bit data bus. Each 32-bit data word has a corresponding parity bit.

The TMS390S10 implements full JTAG (IEEE 1149.1) boundary scan using a JTAG test interface.

PRODUCT PREVIEW

Block Diagram



PRODUCT PREVIEW

Figure 2. TMS390S10 Block Diagram

Overall description

Integer Unit (IU)

The TMS390S10 Integer Unit executes SPARC integer instructions as defined in the SPARC Architecture Manual (Version 8). The IU contains 120 *r* registers (7-window registers and 8-global registers) and operates on instructions using a five-stage pipeline.

MEIKO Floating Point Unit (FPU)

The FPU fully executes all single and double precision FP instructions as defined in the SPARC Architecture Manual (Version 8). Quad-precision instructions and `fsmuld` instruction are not implemented. They are trapped to `unimplemented_FPop`. Since all implemented instructions are executed within hardware, this FPU never generates `unfinished_FPop` exception. The FPU contains 2x16x32 *f* registers.

Memory Management Unit (MMU)

The MMU is compatible with the SPARC Reference MMU, as defined by the SPARC Architecture Manual Version 8 appendix.H "SPARC Reference MMU Architecture". The TMS390S10 MMU also serves as an I/O MMU. This MMU translates 32-bit virtual addresses of each running process to 31-bit physical addresses in memory. The 3-high order bits of Physical address are maintained to support memory mapping into 8 different address spaces. The MMU supports 64 contexts.

The MMU also protects memory so that a process can be prohibited from reading or writing to the address space of another process.

The MMU controls arbitration between I/O, Data Cache, Instruction Cache, and TLB references to memory,

The MMU contains a 32-entry fully associative TLB and uses a pseudo random algorithm for the replacement of TLB entries.

Instruction Cache

The Instruction Cache is a 4-KByte, physically tagged cache. The Instruction Cache data is organized as 128 lines of 32 bytes.

Data Cache

The Data Cache is a 2-KByte, direct mapped, physically tagged, write through cache with no write allocate. The data store is organized as 128 lines of 16 bytes.

Data cache read and write hits take no extra pipe cycle except doubleword operations.

`LDD` takes 2 cycles to complete (1 extra cycle) and `STD` takes 3 cycles to complete (2 extra cycles).

There are two Store Buffers to hold data being stored from the IU or FPU to memory or other physical devices. The Store Buffers are 32-bit registers.

Memory Interface

The TMS390S10 provides a complete DRAM controller which generates all the signals necessary to support up to 128 MBytes of system memory.

The DRAM bus is 64 bits wide with two parity bits, one covering each 32 bits of data.

The system DRAM is organized as four banks, each of which may be 2 MBytes, 8 MBytes, or 32 MBytes depending upon the size of DRAM used.

Memory Control Block (MCB)

The Memory Control Block keeps track of the priorities of memory operations and completely provide all signals necessary to interface to DRAM based main memory. Memory operation can be:

- data read, data write, and read-modify-write for CPU execution
- instruction fetches
- translation buffer accesses during table walks
- reads and writes by I/O devices
- DRAM refresh.

Data aligner and Parity Check/generate logic (DPC)

Due to the minimum size of read/write to DRAM being 32 bits wide, 8 and 16-bit access requires read-modify-write operation and correct alignment to 32-bit boundary. The DPC aligns data and generates parity data for writing 32-bit data.

RAM Refresh Control (RFR)

The TMS390S10 RAM Refresh Control logic provides complete DRAM Refresh control. This Refresh Controller performs CAS-before-RAS refresh. Refresh interval is programmable.

SBus Interface

The SBC Interface performs all functions necessary to interface the TMS390S10 to the SBus, including dynamic bus sizing, cycle re-run control, burst cycle re-ordering, arbitration, and general SBus control. The SBC works with the MMU to arbitrate the system and memory resources and for I/O address translations.

SBus Controller (SBC)

The TMS390S10 SBC (SBus Controller) controls direct connected SBus devices. The SBC Control logic operates as the system SBus Controller. It supports following:

- Programmed Input/Output (PIO) transactions between the CPU and SBus devices
- Direct Virtual Memory Access (DVMA) transactions between SBus masters and local resources. (referred to as Local DVMA)
- Direct Virtual Memory Access (DVMA) transactions between SBus masters and other SBus slave devices. (referred to as Bypass DVMA)

For further details on SBus specification, please refer to the SBus Specification Rev A.2. available from Sun Microsystems. The TMS390S10 does not comply with SBus Rev B.0.

Table 2. Pin Function

SIGNAL	I/O	PIN NO	DESCRIPTION
CPSTAT1 CPSTAT0	O	232 234	11 — normal 10 — Level 15 interrupt 01 — Trap occurred, when trap disabled 00 — reserved
EXT_EVENT	I	226	External Event. To stop clocks during debug on an externally triggered event.
IN_CLK	I	251	CPU clock input
INT_EVENT	O	235	Internal Event. This is an internal counter overflow bit. For debug.
IRQ3 IRQ2 IRQ1 IRQ0	I	256 257 258 259	Interrupt request lines
JTCK	I	249	JTAG test clock input.
JTDI	I	241	JTAG test data input.
JTDO	O	243	JTAG test data output.
JTMS	I	246	JTAG test mode select input.
JTRST	I	247	JTAG test reset input.
MADDR11 MADDR10 MADDR9 MADDR8 MADDR7 MADDR6 MADDR5 MADDR4 MADDR3 MADDR2 MADDR1 MADDR0	O	244 260 261 263 265 266 267 269 270 271 274 277	Physical Address Bus. These signals are multiplexed address output. These pins need external buffer to provide the necessary drive for the DRAMs.
MCAST MCAS0	O	284 285	DRAM Column Address Strobe signals. These pins need external buffer to provide the necessary drive for the DRAMs.
MDATA63 MDATA62 MDATA61 MDATA60 MDATA59 MDATA58 MDATA57 MDATA56 MDATA55 MDATA54 MDATA53 MDATA52 MDATA51 MDATA50 MDATA49 MDATA48 MDATA47 MDATA46 MDATA45 MDATA44 MDATA43 MDATA42 MDATA41 MDATA40 MDATA39 MDATA38 MDATA37 MDATA36 MDATA35	I/O	3 4 6 8 9 10 11 13 15 16 18 21 22 23 25 27 28 29 31 34 36 37 38 40 41 42 44 46 47	Memory Data Bus

PRODUCT PREVIEW

Table 2.Pin Function (continued)

PRODUCT PREVIEW

SIGNAL	I/O	PIN NO	DESCRIPTION
MDATA34 MDATA33 MDATA32 MDATA31 MDATA30 MDATA29 MDATA28 MDATA27 MDATA26 MDATA25 MDATA24 MDATA23 MDATA22 MDATA21 MDATA20 MDATA19 MDATA18 MDATA17 MDATA16 MDATA15 MDATA14 MDATA13 MDATA12 MDATA11 MDATA10 MDATA9 MDATA8 MDATA7 MDATA6 MDATA5 MDATA4 MDATA3 MDATA2 MDATA1 MDATA0	I/O	48 49 51 56 57 59 60 61 63 65 66 67 68 70 74 75 76 78 79 80 82 84 85 87 89 91 93 94 95 97 98 99 101 103 104	Memory Data Bus
MPAR1 MPAR0	I/O	2 54	Memory parity bits. MPAR1 is for MDATA0 – MDATA31 MPAR0 is for MDATA32 – MDATA63
MRAS3 MRAS2 MRAS1 MRAS0	O	278 279 280 282	DRAM Row Address Strobe signals. These pins need external buffer to provide the necessary drive for the DRAMs. MRAS3 : 1st 32MB bank MRAS2 : 2nd 32MB bank MPAR1 : 3rd 32MB bank MPAR0 : 4th 32MB bank
MWE	O	286	DRAM Write Enable. This pin needs external buffer to provide the necessary drive for the DRAMs.
NO CONNECTION		238	This line must be open
NO CONNECTION		230	This line must be open
NO CONNECTION		229	This line must be open
NO CONNECTION		227	This line must be open
REF_CLK	O	236	Reference Clock output at 1/2 input clock (IN_CLK).
RESET	I	248	Power-up reset
RI_REQ	O	231	This output signal indicates that the MMU is doing either an operation to memory, to an ASI, control space, or an SBus device. This signal can be used to validate the physical address observation on the SBus address pins using view mode during debug.

PRODUCT PREVIEW documents contain information on products in the formative or design phase of development. Characteristic data and other specifications are design goals. Texas Instruments reserves the right to change or discontinue these products without notice.



Table 2.Pin Function (continued)

SIGNAL	I/O	PIN NO	DESCRIPTION																																				
SBACK2 SBACK1 SBACK0	I/O	192 194 195	<p>SBus Transfer Acknowledgement.</p> <p>These signals used to indicate the data has been transferred, or that the current bus cycle should be terminated, or both.</p> <table border="1"> <thead> <tr> <th>SBACK2</th> <th>SBACK1</th> <th>SBACK0</th> <th></th> </tr> </thead> <tbody> <tr> <td>1</td> <td>1</td> <td>1</td> <td>Idle/Wait</td> </tr> <tr> <td>1</td> <td>1</td> <td>0</td> <td>Error Acknowledgement</td> </tr> <tr> <td>1</td> <td>0</td> <td>1</td> <td>Byte (Data) Acknowledgement</td> </tr> <tr> <td>1</td> <td>0</td> <td>0</td> <td>Return Acknowledgement</td> </tr> <tr> <td>0</td> <td>1</td> <td>1</td> <td>Word (Data) Acknowledgement</td> </tr> <tr> <td>0</td> <td>1</td> <td>0</td> <td>Double-word (Data) Acknowledgement</td> </tr> <tr> <td>0</td> <td>0</td> <td>1</td> <td>Half-word (Data) Acknowledgement</td> </tr> <tr> <td>0</td> <td>0</td> <td>0</td> <td>Reserved</td> </tr> </tbody> </table>	SBACK2	SBACK1	SBACK0		1	1	1	Idle/Wait	1	1	0	Error Acknowledgement	1	0	1	Byte (Data) Acknowledgement	1	0	0	Return Acknowledgement	0	1	1	Word (Data) Acknowledgement	0	1	0	Double-word (Data) Acknowledgement	0	0	1	Half-word (Data) Acknowledgement	0	0	0	Reserved
SBACK2	SBACK1	SBACK0																																					
1	1	1	Idle/Wait																																				
1	1	0	Error Acknowledgement																																				
1	0	1	Byte (Data) Acknowledgement																																				
1	0	0	Return Acknowledgement																																				
0	1	1	Word (Data) Acknowledgement																																				
0	1	0	Double-word (Data) Acknowledgement																																				
0	0	1	Half-word (Data) Acknowledgement																																				
0	0	0	Reserved																																				
SBADDR27 SBADDR26 SBADDR25 SBADDR24 SBADDR23 SBADDR22 SBADDR21 SBADDR20 SBADDR19 SBADDR18 SBADDR17 SBADDR16 SBADDR15 SBADDR14 SBADDR13 SBADDR12 SBADDR11 SBADDR10 SBADDR9 SBADDR8 SBADDR7 SBADDR6 SBADDR5 SBADDR4 SBADDR3 SBADDR2 SBADDR1 SBADDR0	O	115 118 121 125 127 130 133 135 137 139 142 147 149 151 154 156 158 162 166 168 170 172 177 180 182 184 187 189	<p>SBus Address bus.</p> <p>The SBus signals used by the SBus controller to send a physical address to a slave.</p>																																				
SBAS	O	190	<p>SBus Address Strobe.</p> <p>This signal used by the SBus controller to indicate that a slave cycle is in progress.</p>																																				
SBCLK	O	254	<p>SBus Clock.</p> <p>This clock is the master clock reference for all SBus devices.</p>																																				
SBDATA31 SBDATA30 SBDATA29 SBDATA28 SBDATA27 SBDATA26 SBDATA25 SBDATA24 SBDATA23 SBDATA22	I/O	105 107 110 112 113 117 120 123 126 129	<p>SBus Data bus.</p> <p>The SBus signals used to transfer data between masters and slaves, and virtual address between masters and the SBus controller.</p>																																				

PRODUCT PREVIEW

Table 2.Pin Function(continued)

SIGNAL	I/O	PIN NO	DESCRIPTION			
SBDATA21	I/O	132	SBus Data bus. The SBus signals used to transfer data between masters and slaves, and virtual address between masters and the SBus controller.			
SBDATA20		134				
SBDATA19		136				
SBDATA18		138				
SBDATA17		141				
SBDATA16		146				
SBDATA15		148				
SBDATA14		150				
SBDATA13		153				
SBDATA12		155				
SBDATA11		157				
SBDATA10		160				
SBDATA9		165				
SBDATA8		167				
SBDATA7		169				
SBDATA6		171				
SBDATA5		175				
SBDATA4	179					
SBDATA3	181					
SBDATA2	183					
SBDATA1	186					
SBDATA0	188					
SBGR4	O	219	SBus Bus Grants (1/master)			
SBGR3		221				
SBGR2		222				
SBGR1		223				
SBGR0		224				
SBLERR	I	203	SBus Late Data Error. A special SBus signal to indicate that an error occurred during a preceding data transfer, even though the slave issued a byte, half-word, word, or double word-acknowledgement.			
SBREAD	I/O	215	Transfer Read/Write. H: read data from the slave L: write data to the slave			
SBRQ4	I	196	SBus Bus Requests (1/master). These signals used by the master to request the SBus controller to grant access to the bus.			
SBRQ3		197				
SBRQ2		199				
SBRQ1		201				
SBRQ0		202				
SBSEL4	O	204	SBus Slave Select (1/slave). These signals used to select which slave should be active during the current cycle. H: unselect L: select			
SBSEL3		205				
SBSEL2		207				
SBSEL1		209				
SBSEL0		210				
SBSIZE2	I/O	211	Transfer Size. SBSIZE2 SBSIZE1 SBSIZE0 FUNCTION			
SBSIZE1		213				
SBSIZE0		214				
		0		0	0	Word (four byte) transfer
		0		0	1	Byte transfer
		0		1	0	Half-word (two byte) transfer
		0		1	1	N/A
	1	0	0	Four Word Burst (16 bytes)		
	1	0	1	Eight Word Burst (32 bytes)		
	1	1	0	Sixteen Word Burst (64 bytes)		
	1	1	1	Two Word Burst (8 bytes)		

PRODUCT PREVIEW

PRODUCT PREVIEW documents contain information on products in the formative or design phase of development. Characteristic data and other specifications are design goals. Texas Instruments reserves the right to change or discontinue these products without notice.



Table 2. Pin Function (continued)
POWER CONNECTIONS

SIGNAL	I/O	PIN NO	DESCRIPTION
VCCC	I	1 17 30 53 72 88 106 114 122 144 161 176 200 217 239 250 273	+5 volts for core logic.
VCCP	I	5 14 52 62 71 81 92 100 111 128 143 159 178 191 208 218 228 24 240 255 264 272 283	+5 volts for peripheral logic.
N/C		227 229 230 238	These pins must be tied low.
Reserved	I	245	Reserved. This pin must be tied low.

PRODUCT PREVIEW

**Table 2.Pin Function(continued)
POWER CONNECTIONS**

SIGNAL	I/O	PIN NO	DESCRIPTION
GNDC	I	19	Ground for core logic.
		32	
		55	
		73	
		86	
		108	
		116	
		124	
		145	
		163	
		174	
		198	
		216	
		237	
		252	
275			
288			
GNDP	I	7	Ground for peripheral logic.
		12	
		20	
		26	
		33	
		39	
		45	
		50	
		58	
		64	
		69	
		77	
		83	
		90	
		96	
		253	
		102	
		109	
		119	
		131	
		140	
		152	
		164	
		173	
		185	
		193	
		206	
212			
220			
225			
233			
242			
262			
268			
276			
281			
287			

PRODUCT PREVIEW

PRODUCT PREVIEW documents contain information on products in the formative or design phase of development. Characteristic data and other specifications are design goals. Texas Instruments reserves the right to change or discontinue these products without notice.



Integer Unit (IU)

Overview

The TMS390S10 IU is a SPARC integer unit as defined in the SPARC Architecture Manual Version 8. It contains a 7-window register file and 8 global registers for a total of 120 registers.

TMS390S10 IU has a performance of >36 MIPS at 50 MHz operating frequency.

Pipeline

The IU has a five-stage pipeline. A typical instruction enters the pipeline, and completes five cycles later. The stages are fully overlapped to allow a peak execution rate of one instruction per cycle. Figure 3 shows IU pipeline.

The five stages are Fetch (F), Decode (D), Execute (E), Cache Access (C), and Write Back (W).

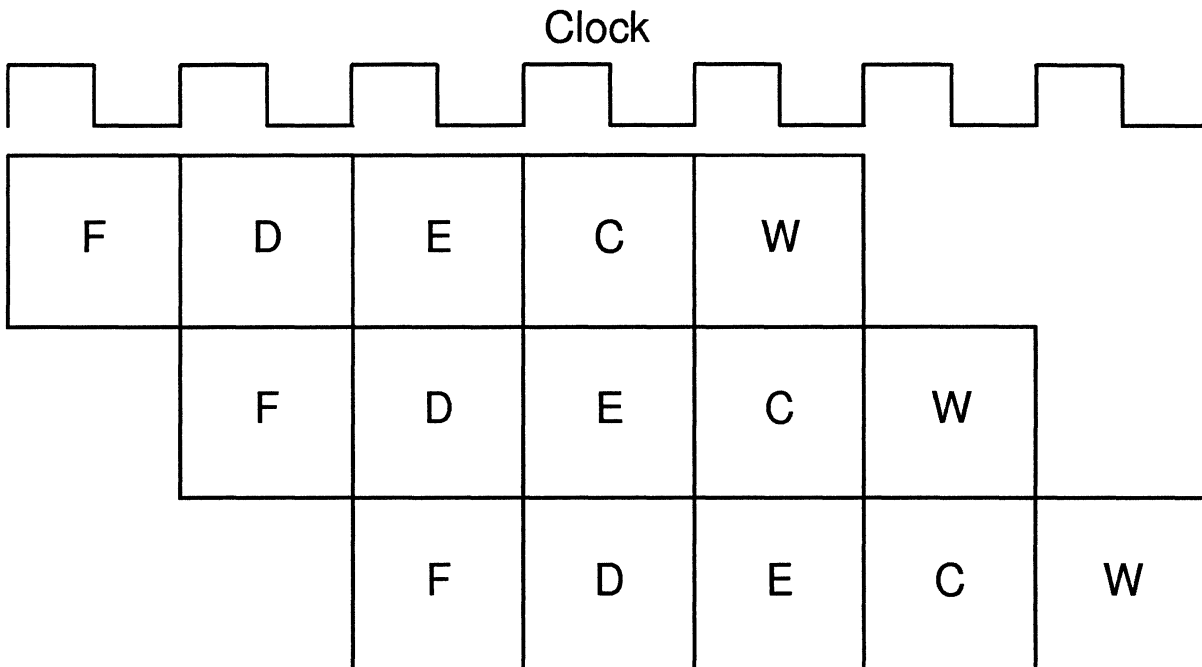


Figure 3. IU Pipeline

The IU has a dedicated 32-bit data bus interfacing directly with the Instruction Cache to support the fetching of one instruction per cycle. If the instruction is in the Instruction Cache, it is returned in the same cycle in which it was requested.

The Instruction Fetch (F) stage is used to access the instruction cache with the address generated in the previous cycle. The instruction is available and registered inside of the IU at the end of this cycle.

The Decode (D) stage of the pipeline is used to read operands from the register file as well as to decode instructions. During the Execution (E) stage, an ALU, SHIFT, MEMOP, JMP, RETT or address generation operation is performed.

During the Cache Access (C) stage of a Data Cache read, the Data Cache access is performed and data is registered inside of the IU. During the C-stage of a Data Cache write, the data registered in the IU is written to the Data Cache.

Finally, during the Write Back (W) stage, the loaded data, or computed result is written back into the register file.

PRODUCT PREVIEW

Instructions and Timings

The TMS390S10 IU execution times are shown below.

Table 3. IU Cycles per Instruction

INSTRUCTION	CYCLES
Call	1
Single Loads	1
Jump/Rett	2
Double Loads	2
Single Stores	2
Double Stores	3
Taken Trap	3
Atomic Load/Store	2
SWAP	2
Integer Multiply	19
Integer Divide	39
All Others	1

PRODUCT PREVIEW

MEIKO Floating Point Unit (FPU)**Overview**

The TMS390S10 FPU is designed to execute all of the single and double precision SPARC Version 8 compatible floating point instructions except `fsmuld`. All other FP instructions trap to `unimplemented_FPop`. Quad precision fp instructions are not supported. All implemented instructions are executed in hardware, therefore the TMS390S10 FPU will never generate an `unfinished_FPop` exception. This FPU contains a 32x32-bit register file.

The TMS390S10 FPU operating frequency is 50 MHz. The performance of this FPU reaches 10+ MFLOPS. With a well balanced combination of IU and FPU, the TMS390S10 can work well in a large number of range applications.

FPU Execution timings

The TMS390S10 FPU instruction execution times are shown in Table 4. The number of clock cycles used depends on the data processed. These cycle counts assume that the operands are available in the register file. A load-use interlock (fp load followed by an FPop which uses the destination register of the load as an operand) may add up to 2 cycles to the typical cycle count. The timings are in CPU cycles.

Table 4. FPU Cycles per Instruction

INSTRUCTION	MIN	TYP	MAX
<code>fadds</code>	4	4	17
<code>faddd</code>	4	4	17
<code>fsubs</code>	4	4	17
<code>fsubd</code>	4	4	17
<code>fmuls</code>	5	5	25
<code>fmuld</code>	7	9	32
<code>fdivs</code>	6	20	38
<code>fdivd</code>	6	35	56
<code>fsqrts</code>	6	37	51
<code>fsqrtd</code>	6	65	80
<code>fnegs</code>	2	2	2
<code>fmovs</code>	2	2	2
<code>fabss</code>	2	2	2
<code>fstod</code>	2	2	14
<code>fdtos</code>	3	3	16
<code>fitos</code>	5	6	13
<code>fitod</code>	4	6	13
<code>fstoi</code>	6	6	13
<code>fdtoi</code>	7	7	14
<code>fcmps</code>	4	4	15
<code>fcmpd</code>	4	4	15
<code>fcomes</code>	4	4	15
<code>fcmped</code>	4	4	15
<code>unimplemented</code>	3	3	3

FQ (Floating Point Deferred-Trap Queue)

The TMS390S10 has a 1-entry floating-point deferred-trap queue. When a floating-point instruction generates an fp_exception, the TMS390S10 will delay the taking of an fp_exception trap until the next floating-point instruction is encountered in the instruction stream. The TMS390S10 FPU implementation can be modeled as having 3 states: fp_execute, fp_exception_pending, and fp_exception. These are shown in Figure 4.

Normally the FPU is in fp_execute state. It moves from fp_execute to fp_exception_pending when an FPop generates a floating-point exception.

The FPU moves from fp_exception_pending to fp_exception, when the IU attempts to execute any floating-point instruction (including fbcc's). This transition (FXACK) generates an fp_exception trap. At this time the FQ contains the instruction and address of the FPop which originally caused the fp_exception.

An fp_exception trap can only be caused while the FPU is moving from the fp_exception_pending state to the fp_exception state (or by executing a STDFQ instruction when FSR.qne == 0, as described below). While in fp_exception state, only floating-point store instructions may be executed (particularly STDFQ and STFSR) and they can not cause an fp_exception trap.

The FPU remains in the fp_exception state until a STDFQ instruction is executed and the FQ becomes empty. At that time, the FPU returns to the fp_execute state.

If an FPop, or a floating-point load instruction (excluding fbcc's and all store instructions) is executed while the FPU is in fp_exception state, the FPU returns to fp_exception_pending state and also sets the FSR.ftt field to sequence_error (0x4). The instruction that caused the sequence_error is not entered into the FQ.

If a STDFQ instruction is executed when the FQ is empty (FSR.qne == 0, FPU is in fp_execute state), the FPU will generate an immediate fp_exception trap (not deferred) and set the FSR.ftt field to sequence_error (0x4), but the FPU will remain in fp_execute state.

The STDFQ instruction will store the address from the FQ to the effective address, and the instruction from the FQ to the effective address + 4.

PRODUCT PREVIEW

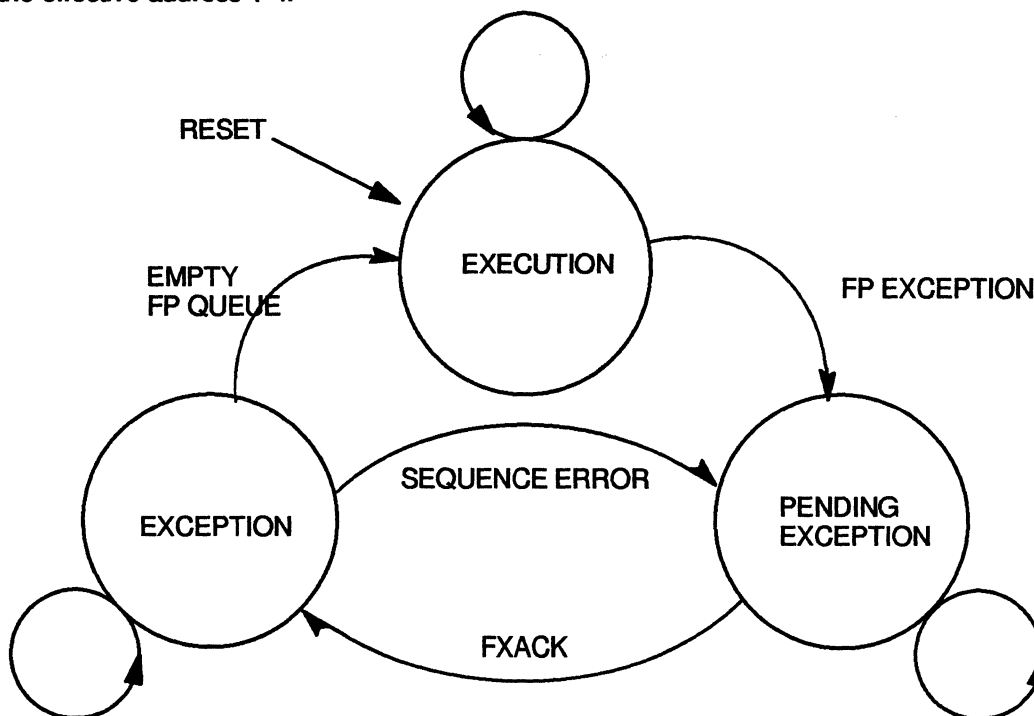


Figure 4. FPU Operation Modes

NaN Format

The TMS390S10 FPU uses different NaN format from the Appendix N, "SPARC IEEE 754 Implementation Recommendations" NaN.

The following shows the value returned for an untrapped floating-point result in the same format as the operands:

Table 5. Untrapped FP Result in Same Format as Operands

		rs2 operand		
		number	QNaN2	SNaN2
rs1 operand	none	IEEE 754	QNaN2	ME_NaN
	number	IEEE 754	QNaN2	ME_NaN
	QNaN1	QNaN1	QNaN1	ME_NaN
	SNaN1	ME_NaN	ME_NaN	ME_NaN

QNaN results will have their sign bit set to 0. ME_NaN is 0x7fff 0000 (Single-precision) or 0x7fff e000 0000 0000 (double-precision).

Table 6. Untrapped FP Result in Different Format

operation	operand (rs2)			
	+QNaN	-QNaN	+SNaN	-SNaN
fstoi	ME_NaN	-imax	+imax	-imax
fstod	(QNaN2)	(QNaN2)	ME_NaN	ME_NaN
fdtos	ME_NaN	ME_NaN	ME_NaN	ME_NaN
fdtoi	ME_NaN	-imax	+imax	-imax

+imax = 0x7fff ffff, and -imax = 0x8000 0000.

(QNaN2) is a copy of the mantissa bits of the operand, with the extra low order bits zeroed, and the sign bit zeroed.

Instruction Cache

The instruction cache is a 4-KByte, physically-tagged cache. An instruction cache data store is organized as 128 lines each of 32 bytes of data.

On a Cache miss

On instruction cache misses, 32 bytes of data are fetched from memory. Filling order is illustrated as the following:

REQUESTED WORD	ORDER OF FILL (MODULO 32B)
0	0, 1, dc, 2, 3, dc, 4, 5, dc, 6, 7
1	1, 0, dc, 2, 3, dc, 4, 5, dc, 6, 7
2	2, 3, dc, 4, 5, dc, 6, 7, dc, 0, 1
3	3, 2, dc, 4, 5, dc, 6, 7, dc, 0, 1
4	4, 5, dc, 6, 7, dc, 0, 1, dc, 2, 3
5	5, 4, dc, 6, 7, dc, 0, 1, dc, 2, 3
6	6, 7, dc, 0, 1, dc, 2, 3, dc, 4, 5
7	7, 6, dc, 0, 1, dc, 2, 3, dc, 4, 5

dc: dead cycle. During 'dead cycle', no data is written.

The ordering rule starts with requested word followed by the other word of the first doubleword and continuing with another doubleword (even word, odd word) which will wrap around a 32-byte boundary until the entire 32-byte block has been returned.

The transfer needs one dead cycle after every double word transfer.

Instruction Cache Flushing

All instruction cache valid bits can be reset (to zero) by any type of alternate store instruction to ASI 0x36. Also, the `FLUSH` instruction always clear the single valid bit that is addressed by `iu_dva[11:05]`, regardless of the contents of this tag entry.

Cacheability of Memory Accesses

Pages that are declared as non-cacheable (C=0 in the PTE) are not cached in the instruction cache. The following operations are not cached:

- Accesses when the MMU is disabled and alternate cacheability is disabled (EN, AC bits of the MMU CR =0).
- Accesses while the instruction cache is disabled (IE bit of the MMU CR=0).
- Accesses while in Boot Mode.
- Access to sources in physical address spaces 1-7.

PRODUCT PREVIEW

Data Cache

The 2-KByte data cache on the TMS390S10 supports a write-through, no write allocation cache protocol. All write operations write through data to main memory and on write hits both the data cache and main memory are updated.

System software may also read and write the data cache tags by executing word length LDA and STA (Load and Store Alternate) instructions in ASI 0xE. The Virtual address bits[10:0] will be used to address the data cache in this mode.

Cacheability

Pages that are declared as non-cacheable (C=0 in the PTE) are not cached in the data cache. The following operations are not cached:

- Accesses when the MMU is disabled and alternate cacheability is disabled (EN, AC bits of the MMU CR =0).
- Accesses while the data cache is disabled (DE bit of the MMU CR=0).
- Accesses while using the MMU bypass ASI (ASI=0x20) and alternate cacheability is disabled (AC bits of the MMU CR=0).
- Accesses while in Boot Mode.
- Accesses by sources in address spaces 0x01–0x07.
- Accesses by the MMU during tablewalks.

On a Cache miss

On a data cache miss to a cacheable location, 16 bytes of data are written into the cache from main memory. Filling is done as follow:

REQUESTED WORD	ORDER OF FILL (MODULO 16B)
0	0, 1, dc, 2, 3
1	1, 0, dc, 2, 3
2	2, 3, dc, 0, 1
3	3, 2, dc, 0, 1

dc: dead cycle. During 'dead cycle', no data is written.

The ordering rule starts with requested word followed by the other word of the first doubleword and continuing with another doubleword (even word, then odd word) which will wrap around a 16-byte boundary until the entire 16-byte block has been returned.

The transfer needs one dead cycle after two words of doubleword transfers.

Data Cache Flushing

The data cache valid bits can be reset to zero by any type of alternate store instruction to ASI 0x37.

Note that the data cache is not flushed by FLUSH instruction.

Memory Management Unit (MMU)

Overview

The MMU provides three primary functions:

- Translates the virtual addresses of instructions and data into physical addresses in memory. The MMU translates from a 32-bit virtual address to a 31-bit physical address by using a Translation Lookaside Buffer (TLB). This TLB contains 32 entries, is fully-associative, and is a pseudo-random algorithm for the replacement of TLB entries.
- Protects memory so that each process can be prohibited from reading or writing the address space of another process.
- Arbitrates between I/O, Data Cache, Instruction Cache, and TLB references to physical memory.

Translation Lookaside Buffer (TLB)

The TLB is a 32-entry, fully associative cache of page descriptors. It caches virtual to physical address translations and the associated page protections and usage information. The pseudo-random replacement algorithm determines which of the 32 entries should be replaced when a TLB miss occurs. In the descriptions that follow, the terms VA and PA are used to generically describe any virtual address or physical address, respectively.

TLB Replacement

The TLB uses a pseudo-random replacement scheme. There is a 5-bit modulo 32 counter in the TLB Replacement Control Register (TRCR) which is incremented by one during each CPU clock cycle to address one of the TLB entries. When a TLB miss occurs, the counter value is used to address the TLB entry to be replaced. On reset the counter is initialized to zero. There is also a bit in the TRCR which is used to disable the counting function. A simple diagram follows.

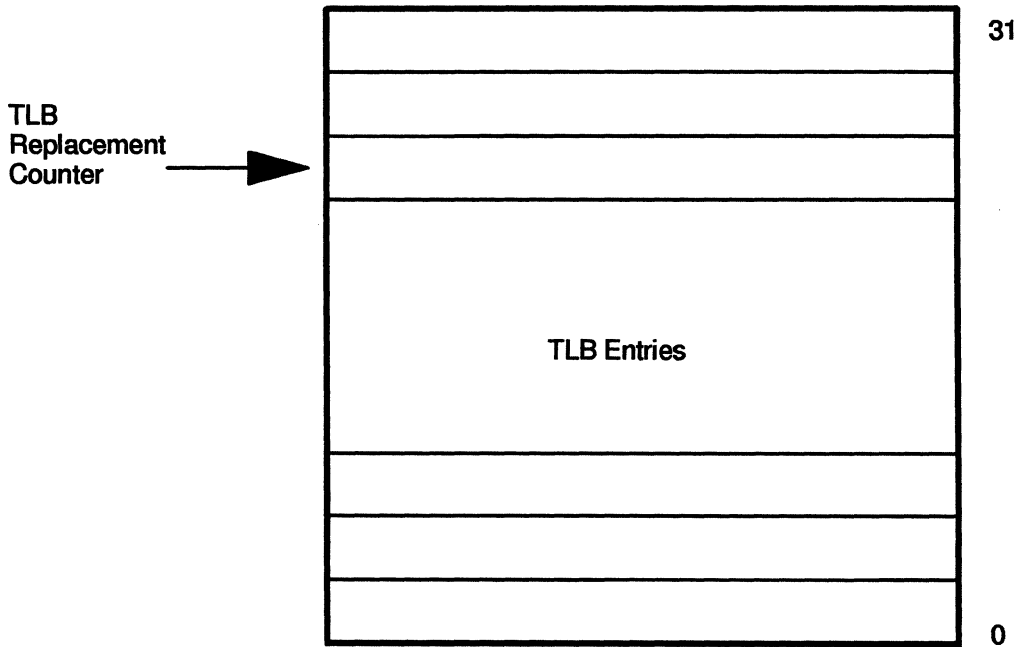



Figure 5. TLB Replacement

PRODUCT PREVIEW

TLB Entry

An entry in the TLB has the following fields.

TLB Entry



Virtual Address Tag

The 20-bit virtual address tag represents the most significant 20 bits (VA[31:12] the page address) of the virtual address being used when referencing Page Table Entries (PTEs) and I/OPTEs. VA[11:00] is the byte within a page. The address in this field is physical when referencing PTPs with the least significant 19 bits containing PA[26:08].

Context Tag

The 6-bit context tag comes from the value in the context register as written by memory management software when referencing PTEs. Both it and the virtual address tag must match the CXR and VA[31:12] in order to have a TLB hit. This field contains a physical address (PA[07:02]) when referencing PTPs. This is not used when referencing I/OPTEs.

Level

The 3-bit level field is used to enable the proper virtual tag match of region, and segment PTEs. I/OPTEs and PTPs will have this field set to use Index 1,2, and 3 (b'111). The most significant bit also serves as the TLB Valid Bit because it is set for any valid PTE, I/OPTE, or PTP. The following table defines the level field.

Table 7. Virtual Tag Match Criteria

LEVEL	MATCH CRITERIA
000	None
100	Index 1 (VA[31:24])
110	Index 1,2 (VA[31:18])
111	Index 1,2,3 (VA[31:12])

Supervisor (S)

This bit is used to disable the matching of the context field indicating that a page is a supervisor level (ACC=6 or 7).

I/O Page Table Entry (I/O)

This bit, if set to zero, indicates that an I/OPTE resides in this entry of the TLB.

Page Table Pointer (PTP)

This bit indicates that a PTP resides in this entry of the TLB.

All SRMMU flush types (except page) will flush all PTPs from the TLB.

Page Table Field

The page table field can either be a Page Table Entry (PTE), a Page Table Pointer (PTP), or an I/O Page Table Entry (I/OPTE). This field can be read and written using ASI 0x06.

PRODUCT PREVIEW

Page Table Entry (PTE)

A Page Table Entry (PTE) defines both the physical address of a page and its access permission. A PTE is defined for SPARC reference MMUs as follows.

Page Table Entry in Page Table

Reserved	PPN	C	M	R	ACC	ET
31 27 26	8	7	6	5	4 2 1	0

Reserved (Rsvd)

Bits[31:27] should be written as zero, and will be read as zero.

Physical Page Number (PPN)

This field is the high order 19 bits ([30:12]) of the 31-bit physical address of the page. The PPN appears on PA[30:12] when a translation completes.

Cacheable (C)

When this bit is set to a one the page is cacheable by the instruction and/or data cache.

Modified (M)

This bit is set to a one when the page is written to.

Referenced (R)

This bit is set to a one when the page is accessed. All PTEs in the TLB have this bit set when the entry is loaded.

Access Permissions (ACC)

These bits indicate whether access to this page is allowed for the transaction being attempted. The Address Space Identifier (ASI) determines whether a given access is a data access or an instruction access, and whether the access is being done by the user or supervisor. The field is defined as follows.

Table 8. Page Table Access Permissions

ACC	PERMISSIONS	
	USER	SUPERVISOR
000	Read only	Read only
001	Read/Write	Read/Write
010	Read/Execute	Read/Execute
011	Read/Write/Execute	Read/Write/Execute
100	Execute only	Execute only
101	Read only	Read/Write
110	No access	Read/Execute
111	No access	Read/Write/Execute

PRODUCT PREVIEW

Entry Type (ET)

This field differentiates the entry types in the TLB.

The entry type is not cached in the TLB. On a probe operation the ET field is derived from a combination of other bits. The bit definitions of the ET field are as follows:

Table 9. Page Table Entry Types

ET	ENTRY TYPE
00	Invalid
01	Page Table Pointer
10	Page Table Entry
11	Reserved

Invalid means that the corresponding range of virtual addresses is not currently mapped to a physical address.

In the TLB RAM the PTE has following fields:

Page Table Entry in TLB

Reserved	PPN	C	M	1	ACC	10
31 27 26	8	7	6	5	4 2 1 0	

Bit[31:27] are not implemented, should be written as zero, and will be read as zero.

Bit[05] is set to one by hardware to indicate that every PTE in the TLB has been referenced. This bit is used to distinguish between PTEs (set to a 1) and I/O PTEs (set to a 0).


Bits[1:0] are set to one:zero by hardware to indicate the entry type (ET) of a PTE. These bits are not actually stored in the TLB. Rather they are derived as a function of the PTP bit of the tag.

Page Table Pointer (PTP)

A PTP contains the physical address of a page table and may be found in the Context Table, in a Level 1 Page Table, or in a Level 2 Page Table. Page Table Pointers are put into the TLB during tablewalks and removed from the TLB either by natural replacement (also during tablewalks) or by flushing the entire TLB.

The Level field in a PTP tag is always set to 0x7. A PTP is defined as follows:

Page Table Pointer in Page Table



Reserved (Rsvd)

Bits[31:27, 03:02] should be written as zero, and will be read as zero.

Page Table Pointer (PTP)

The physical address of the base of a next level page table. The PTP appears on PA[30:08] during miss processing. The page table pointed to by a PTP must be aligned on a boundary equal to the size of the page table.

This is also true with the context table at the root level. The sizes of the tables are summarized as follows.

Table 10. Size of Page Tables

LEVEL	SIZE (BYTES)
Root	256
1	1024
2	256
3	256

Entry Type (ET)

This field differentiates the entry types in the TLB.

The entry type is not kept in the TLB RAM. On a probe operation the ET fields is derived from a combination of other bits. The bit definitions of the ET field follows:


ET	ENTRY TYPE
00	Invalid
01	Page Table Pointer
10	Page Table Entry
11	Reserved

Invalid means that the corresponding range of virtual addresses is not currently mapped to physical addresses.

PRODUCT PREVIEW

In the TLB, a PTP has the following format:

Page Table Pointer In TLB



Bit[31:27] are not implemented, should be written as zero, and will be read as zero.

Bit[03,02] are set to zero by hardware and are unused.

Bits[01:00] are set to zero:one by hardware to indicate the Entry Type(ET) of a PTP. These bits are not actually stored in the TLB rather are derived as a function of the PTP bit of the tag.


PRODUCT PREVIEW

I/O MMU PTE

An I/O Page Table Entry (I/O PTE) defines both the physical address of a page and its access permission.

The Level field in a I/O PTE tag is always set to 0x7 and the Supervisor bit is set to 0x0. An I/O PTE is defined as follows.

I/O Page Table Entry in Page Table



Reserved (Rsvd)

Bits[31:27] are not implemented, should be written as zero, and will be read as zero. Bits[07:03] should also be written as zero, and will be read as zero.

Physical Page Number (PPN)

This field is the high order 19 bits ([30:12]) of the 31-bit physical address of the page. The PPN appears on PA[30:12] when a translation completes. This address is concatenated with VA[11:00] to provide the entire translated address.

Writeable (W)

When this bit is set to a one both reads and writes to the page are allowed. When this bit is zero only reads are allowed.

Valid (V)


This bit is set to a one when the I/O PTE is valid.

Write As Zero (WAZ)

This bit is to be written as zero in the Memory I/O Page Table by software.

In the TLB a I/O PTE has the following format:

I/O Page Table Entry in TLB



Bits[31:27] are not implemented, should be written as zero, and will be read as zero.

Bits[07:03] are set to zero by hardware. Bit[05] is used to distinguish between PTEs (set to 1) and I/O PTEs (set to 0). Bits[07,06,04,03] are not used.

Bits[01:00] are set to one:zero by hardware to indicate a valid I/O PTE. These bits are not actually stored in the TLB.

PRODUCT PREVIEW

CPU TLB

A virtual address to be translated by the MMU is compared to each entry in the TLB. During the TLB lookup the value of the Level field specifies which index fields are required to match the TLB virtual tag as follows:

LEVEL	MATCH CRITERIA
000	None
100	Index 1 (VA[31:24])
110	Index 1,2 (VA[31:18])
111	Index 1,2,3 (VA[31:12])


In addition to the virtual tag match, context matching of a PTE is required for all user page references (ACC is 0 to 5) when made by either user or supervisor (ASI = 0x8—0xB). Context matching is not required for a supervisor page reference (ACC is 6 or 7) when made by a supervisor (ASI = 0x9 or 0xB). This case takes advantage of the Supervisor bit in the TLB tag.

User references (ASI = 0x8 or 0xA) to supervisor pages (ACC is 6 or 7) result in address exceptions.

The TLB ignores access level checking during probe operations. The most significant Level field bit is used as a Valid bit for the TLB. This means that root level PTEs are not supported.

CPU TLB Flush and Probe Operations

The flush operation allows software invalidation of TLB entries. TLB entries are flushed by using a store alternate instruction. The probe operation allows testing the TLB and page tables for a PTE corresponding to a virtual address. TLB entries are probed by using a load alternate instruction. The ASI value 0x3 is used to invalidate or probe entries in the TLB. In an alternate address space used for probing and flushing the address is composed as follows:

CPU TLB Flush or Probe Address Format*Virtual Flush or Probe Address (VFPA)*

This field is the address that is used to index into TLB. Depending on the type of flush or probe not all 20 bits are significant.

Type

This field specifies the extent of the flush or the level of the entry probed.

Reserved

These bits are ignored. They should be set to zero.

TLB Table Walk

On a translation miss the table walk hardware translates the virtual address to a physical address by "walking" through a context table and from 1 to 3 levels of page tables. The first and second levels of these tables typically (not necessarily) contain page table pointers (PTP) to the next level tables when accesses are due to CPU instruction or data addresses. I/O accesses only the first level page table. A third level table entry should always be a table entry (PTE) pointing to a physical page or else a translation fault occurs.

The table walk for a CPU generated virtual address uses the context table pointer register (CTPR, see Page 73) as a base register, and also the context number contained in the context register (CXR, see Page 73) as an offset to point to an entry in the context table. The context table entry is then used as a PTP into the first level page table. At any address, the table walk hardware finds either a PTP or a PTE which terminates its search. A PTP is used in conjunction with a field in the virtual address to select an entry in the next level of tables. The table walk continues searching through levels of tables as long as PTP pointing to the next table is found. The table walk terminates if either a PTE is found or an exception is generated when a PTE is not found after accessing the third level page table (or when an invalid or reserved entry is found).

Note that the PTPs and PTEs encountered during a table walk are not cached in the data cache.

A full table walk is shown in the Figure 6.

PRODUCT PREVIEW

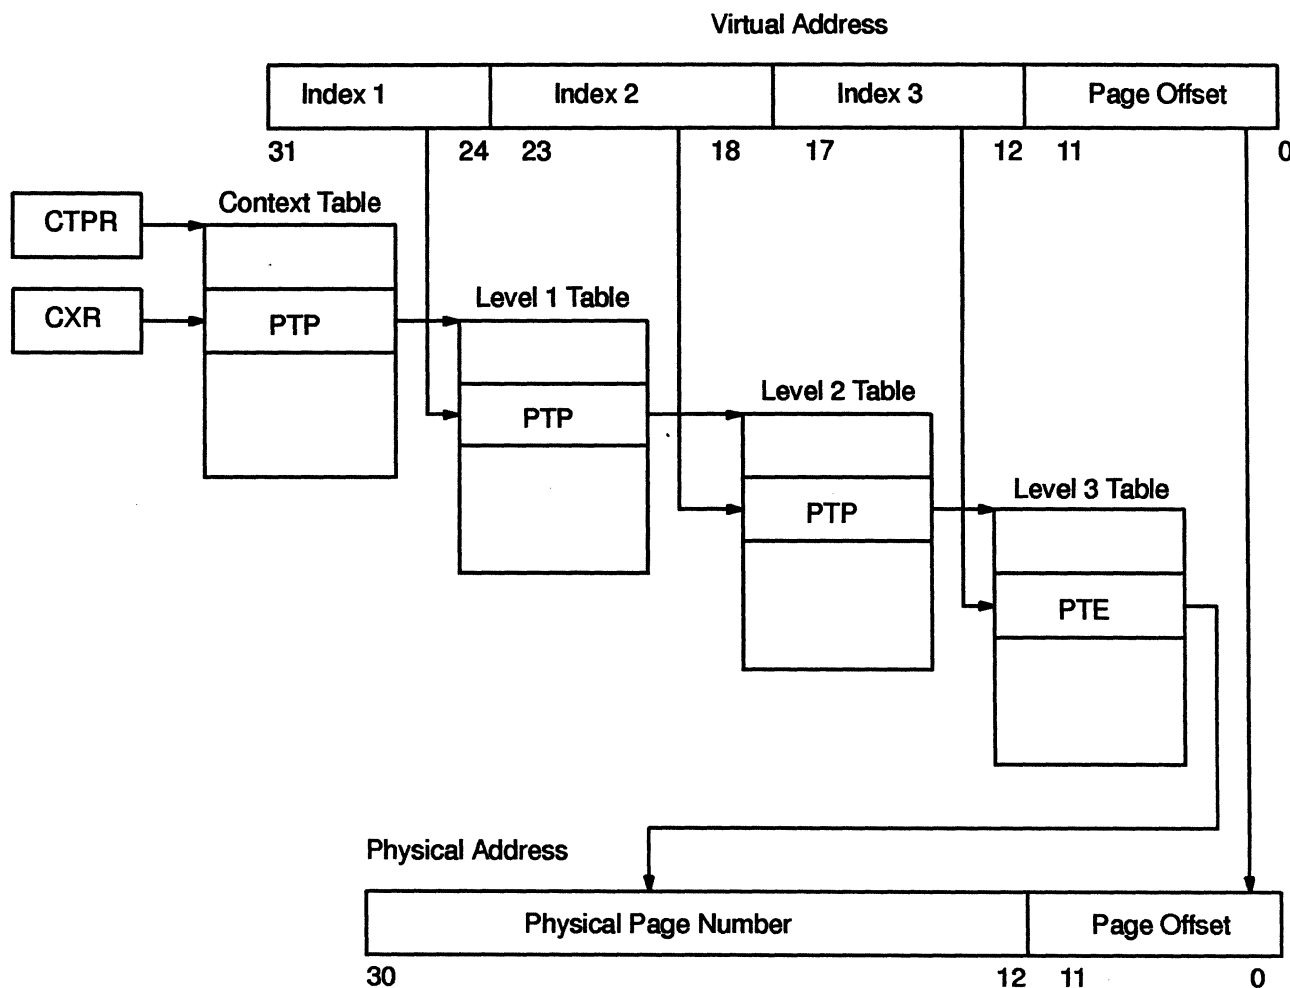


Figure 6. CPU Address Translation Using Table Walk

Translation modes

Translation of virtual addresses to physical addresses is done in the following modes:

Table 11. Translation Modes

NAME	ASI	BOOT MODE	MMU ENABLE	PA[30:00]
Boot IFetch	0x8, 0x9	Yes	N/A	PA[30:28]=0x7, PA[27:00]=VA[27:00]
Pass Through	0x8, 0x9	No	Off	PA[30:00]=VA[30:00]
Translate	0x8, 0x9	No	On	PA[30:12]=PTE[26:08], PA[11:00]=VA[11:00]
Pass Through	0xA, 0xB	N/A	Off	PA[30:00]=VA[30:00]
Translate	0xA, 0xB	N/A	On	PA[30:12]=PTE[26:08], PA[11:00]=VA[11:00]
Bypass	0x20	N/A	N/A	PA[30:00]=VA[30:00]

PA[bit range]: the bits of the physical address.

VA[bit range]: the bits of the virtual address.

PTE[bit range]: the bits of the PTE

N/A: Not Available

PRODUCT PREVIEW

Memory Interface

Overview

The TMS390S10 supports up to 128 MBytes of system memory.

This 128 MBytes of system memory is composed of four 32 MByte DRAM banks .

The banks can support different density devices. This allows the TMS390S10 to work with various sizes of system memory.

The TMS390S10 memory interface has a 64-bit data bus with 2-bit parity (1 bit parity per 32-bit word).

This Memory Interface also includes DRAM Refresh. Complete signals for CAS-before-RAS operation is supplied.

Memory Organization

The TMS390S10 has a 27-bit address space (128 MBytes) for its system. This 128 MByte space is divided into 4 banks. Each bank has capability of addressing up to 32 MBytes.

Each bank has a 64-bit data path to the TMS390S10. The TMS390S10 has 2 CAS lines to select upper or lower 32 bits (high or low word) of each bank. All the banks are controlled by the same write signal MWE. The TMS390S10 has 4 RAS lines to select DRAM banks.

The 22-bit multiplexed Row/Column address bus (11-row/11-column) are used for addressing all the banks. And the TMS390S10 also provides 12-row/10-column address signals to use 12x10 matrix DRAMs.

The TMS390S10 uses 16 pieces of 4-bit wide DRAM devices or 2 SIMMs with eight devices on each.

These configurations provide the 64-bit wide data bus. And, each bank needs two 1-bit wide devices of the same depth (if using SIMMs, one on each SIMM) to store the 2 parity bits.

Each bank can use one of following configurations:

Table 12. Memory Configuration Table

BANK MEMORY	PARITY	DRAM	SIMM†
2MB (256K x 64)	2 of 256K x 1	16 of 256K x 4	2 of 256K x 33
8MB (1M x 64)	2 of 1M x 1	16 of 1M x 4	2 of 1M x 33
32MB (4M x 64)	2 of 4M x 1	16 of 4M x 4	2 of 4M x 33

† A pair of double-density (e.g. 512Kx33 or 16Mx33) SIMMs will occupy 2 banks since they need 2 RASes.

Memory Access Speed and Refresh

Typically the TMS390S10 requires 60ns DRAMs at 50MHz, 80ns DRAMs at 40MHz and 33 MHz clock speed.

DRAM speed is highly system dependent. The designer must refer to the AC specification for the memory interface to select the proper DRAMs for the particular board design.

Refresh Conditions

The TMS390S10 supports CAS-before-RAS refresh.

The TMS390S10 can select one of three refresh cycles depending on the setting of the PC bits in the Processor Control Register (bits 10, 11). Please refer to Table 32 on page 72. These bits select a refresh of either 128, 512, or 768 cycles or no-refresh.

PRODUCT PREVIEW

Address Decode & Translation**Address Translation**

The TMS390S10 Memory Management Unit (MMU) translates 32-bit virtual addresses from the IU into 31-bit physical addresses. This translation is done by using look-up tables in the MMU. Normally, this translation is cached in a 32-entry fully associative TLB.

Table 13. Physical Address Decode

PA	DECODE
30—27	Not Used. System limit is 128MB
26—25	Decode to select 1 of 4 RASes 00 RAS0 1st 32 MB bank 01 RAS1 2nd 32 MB bank 02 RAS2 3rd 32 MB bank 03 RAS3 4th 32 MB bank
24	Decoded as row address bit 10. Required for 16 Mbit DRAMs
23†	Decoded as column address bit 10 and row address 11. Required for 16 Mbit DRAMs See note.
22	Decoded as row address bit 9. Required for 4 Mbit DRAMs
21	Decoded as column address bit 9. Required for 4 Mbit DRAMs
20—12	Decoded as row address bit 8 to 0. Required for 1 Mbit DRAMs and up
11—3	Decoded as column address bit 8 to 0 Required for 1 Mbit DRAMs and up
2	Decoded to select one of 2 CASes 0 CAS0 Lower data word [31:00] 1 CAS1 Higher data word [63:32]
1—0	Not Used. Byte and halfword are achieved by MCB and DPC doing a read-modify-write sequence

† Physical address 23 (PA[23]) is used as both column address 10 and row address 11. This is to provide for the 2 different 4Mx4 DRAM architecture, 11x11 matrix and 12x10 matrix.

As a result, the TMS390S10 can support up to 128 MByte system memory.

Access to unused or unpopulated memory regions

Any access to a location in the upper 128 MB will be mirrored to its corresponding location in the lower 128 MB and no errors will be generated. Similarly, if a bank contains less than the defined maximum of 32 MB, the real memory will be mirrored to the higher unused portions and an access from any of the unused sections will be mirrored to the corresponding location in the lowest block and no errors will be generated.

For example, if a bank contains 2 MB of real memory, this will be mirrored on the remaining 15 empty portions. However, although an access from a fully unused (empty) bank will complete, but its result will be unknown and may cause a parity error.

PRODUCT PREVIEW

Data Aligner and Parity Check/Generate Logic (DPC)

The DPC works to transfer data between the external memory data bus and the internal data path as well as to generate and check parity for system main memory.

Data Alignment

During any read, write, or hardware controlled read-modify-write cycle, DPC performs the necessary data alignment and byte/halfword placement. It also provides temporary storage for hardware controlled read-modify-write cycles, resulting from byte/halfword write cycles to memory.

Parity Check/Generate

The TMS390S10 uses 1 bit parity per 32-bit word. Therefore for 64-bit access two parity bits are always used. In the Processor Control Register, the PC bit (bit 17) designates whether even or odd parity is used.

PC	DESCRIPTION
0	Check/Generate Even Parity
1	Check/Generate Odd Parity

DRAM Interface Signals

Table 14 shows the signals which are associated with the system memory interface bus.

Table 14. Pin Assign Table

PIN NAME	IN/OUT	PINS	DESCRIPTION
MDATA[63:0]	I/O	64	Memory data bus
MPAR[1:0]	I/O	2	Memory parity (word)
MADDR[11:0]	O	12	DRAM address
MRAS[3:0]	O	4	DRAM RAS
MCAS[1:0]	O	2	DRAM CAS
MWE	O	1	DRAM WRITE

PRODUCT PREVIEW

Memory Access Latencies

Following information assumes no waiting due to other operations using the memory bus.

- For single-word operations, data cache read and write hits have no penalty.
- For double-word operations, LDD takes 2 cycles to complete (1 extra cycle) and STD takes 3 cycles to complete (2 extra cycles).
- Data cache read and instruction caches misses suffer a minimum of a 4 or 9 cycles latency depending on whether fast page mode was used.
- Typically, there is more of a penalty than those noted above due to interlock conditions that occur as the cache is being updated with the rest of the cache line.
- Data cache write misses take no extra time.

Translation misses cause a table walk to occur. There will be anywhere from 1 to 4 main memory operations depending on the number of tables that are walked through the hit rate of PTPs in the TLB. The following formula summarizes the number of cycles that table walks will take:

$$\begin{aligned} \text{Cycles} = & ((\# \text{ levels of page tables walked}) \times 2) \\ & + ((\# \text{ non fast page mode accesses}) \times 8) \\ & + ((\# \text{ fast page mode accesses}) \times 3) \end{aligned}$$

NOTE: # levels of page tables = 1, 2, or 3 (3 by default) 0 < # non fast page mode + fast page mode > 5.

SBus Interface

Overview

SBus is the user interface for the TMS390S10, used to access I/O devices. The SBus Controller Block (SBC), shown in NO TAG, controls access to the SBus and supports either PIO or DVMA.

For further details on SBus specification, please refer to SBus Specification Rev.A.2. SBus Specification is available from Sun Microsystems. Since the TMS390S10 is compatible to SBus Specification A.2, **do not** refer to SBus Specification B.0.

PIO

PIO transactions consist of an SBus slave cycle only; the address translation is done in advance of the bus acquisition. PIO transactions occur when the processor executes loads or stores to I/O (SBus) space. In the case of PIO write transaction, the write is posted to the SBus, so that processing by the processor can continue while the SBus transaction completes. A stall will occur only if another PIO transaction is attempted before the previous PIO write transaction completes. In the case of PIO read transaction, processing is always stalled until the data becomes valid at the end of the SBus transaction.

DVMA

DVMA transactions occur when an SBus master has acquired the bus in order to execute a transaction to a slave. A DVMA transaction consists of an address translation cycle and a slave cycle. The target of the slave cycle is determined once the translation cycle complete.

At the beginning of a bus cycle, a master places a virtual address on the data lines (SBDATA[31:00]). The TMS390S10's SBC translates into a physical address by using TMS390S10's MMU and places physical address on the address lines (SBADDR[27:00]).

DVMA transaction can be one the following two types:

- Local DVMA. Transactions between SBus masters and local resources.
- Bypass DVMA. Transactions between SBus masters and other SBus slave devices.

The TMS390S10 SBus Controller (SBC) Features

- A 32-bit data path.
- A 32-bit virtual address translation capability for Bus masters.
- A 28-bit physical address and 5 separate slave select signals.
- An SBus master clock equal to one half TMS390S10 system clock frequency.
- Completely synchronous operation, except for interrupt.
- Support up to 5 SBus devices.
- 2, 4, 8, and 16 byte burst transfers.
- Compatible with CMOS components.



SBus related signals

Table 15 shows all SBus related signals.

Table 15. SBus related signals

PIN NAME	IN/OUT	# OF PINS	DESCRIPTION
SBADDR[27:0]	out	28	SBus Address
SBDATA[31:0]	I/O	32	SBus Data bus
SBSEL[4:0]	out	5	SBus Slave Select (1/slave)
SBSIZE[2:0]	I/O	3	Transfer Size
SBREAD	I/O	1	Transfer Direction
SBCLK	out	1	SBus Clock . 0.5 x System Clock Frequency
SBAS	out	1	SBus Address Strobe
SBACK[2:0]	I/O	3	Transfer Acknowledgement
SBLERR	in	1	Late Data Error
SBRQ[4:0]	in	5	Bus Requests (1/master)
SBGR[4:0]	out	5	Bus Grants (1/master)

DVMA Cycles

The SBus DVMA cycle has two portions: translation cycle and slave cycle. But in host-based systems, no translation cycle occurs on the bus.

Translation cycle is the portion of a bus cycle between the assertion of grant and the placing of an address on the physical lines by the SBus controller. After receiving the grant, the designated master places a virtual address on the SBus data lines.

Slave cycle is the portion of a bus cycle that begins with placing an address on the physical address, and ends with Address Strobe being unasserted.

Reset

Reset is provided for the initialization of all the SBus devices.

$\overline{\text{Reset}}$ must be driven low at least 512 clock cycles before going high. For system power-up, power must be stable before these 512 clock cycles start.

The leading edge of $\overline{\text{Reset}}$ may or may not meet setup times with respect to Clock. The trailing edge of $\overline{\text{Reset}}$ must meet setup and hold time with respect to Clock. The SBus controller may keep $\overline{\text{Reset}}$ low for more than 512 clock cycles.

SBus Waveforms

In this section, SBus timing information is described by showing several waveforms. For the following waveforms, conventions like Figure 7 are used.

Detailed SBus timing information can be found in "SBus Specification A.2".

Note that the TMS390S10 SBus controller is compatible to "SBus Specification A.2". Please don't refer to "SBus Specification B.0" or later.

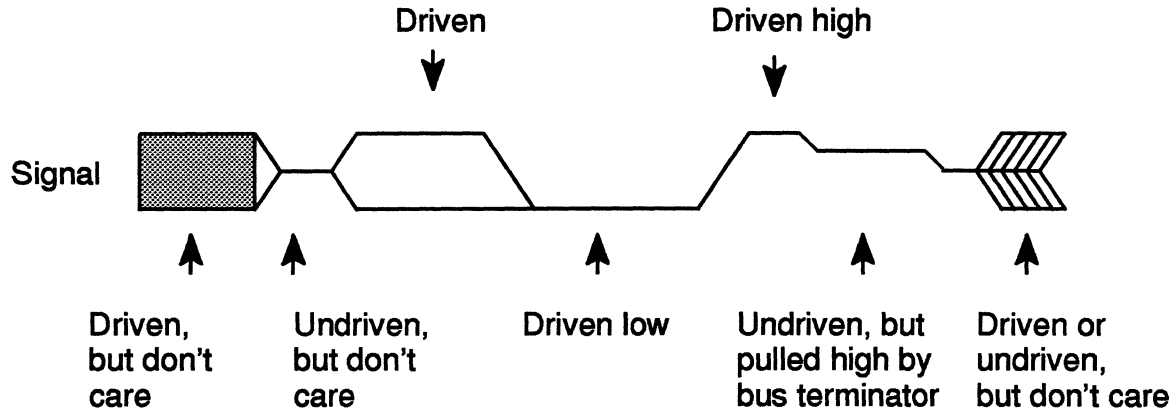



Figure 7. Timing Diagram Conventions

PRODUCT PREVIEW



PRODUCT PREVIEW

Figure 8. Atomic Transaction




Figure 9. Example of RESET Timing

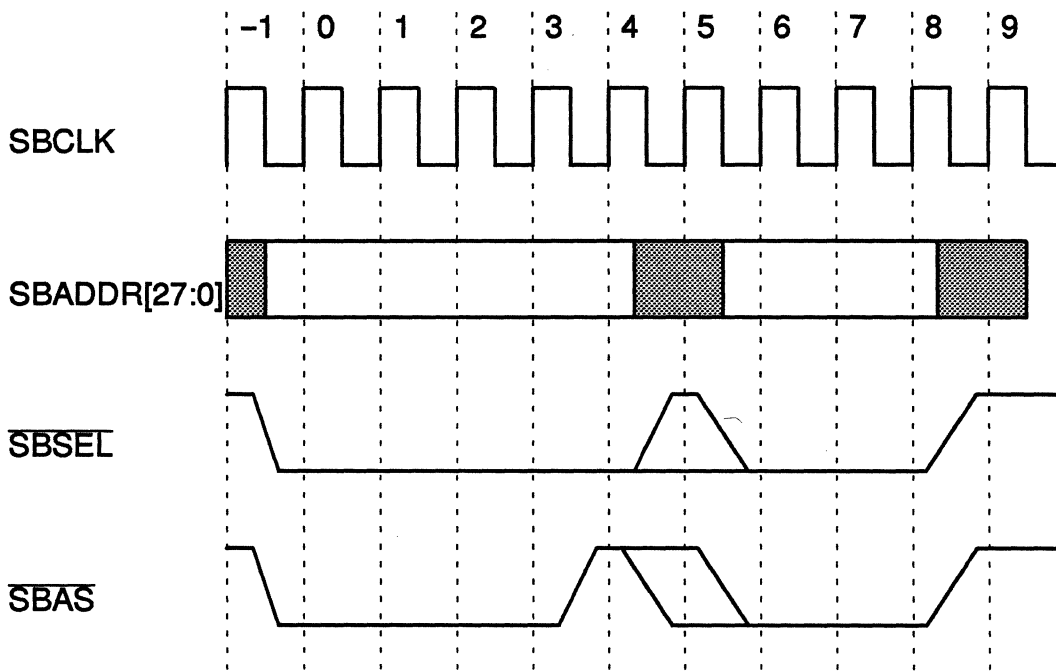


Figure 10. SBADDR[27:0], SBSEL, and SBAS

PRODUCT PREVIEW




Figure 11. Timing of SBBR, SBBG, and SBAS

PRODUCT PREVIEW

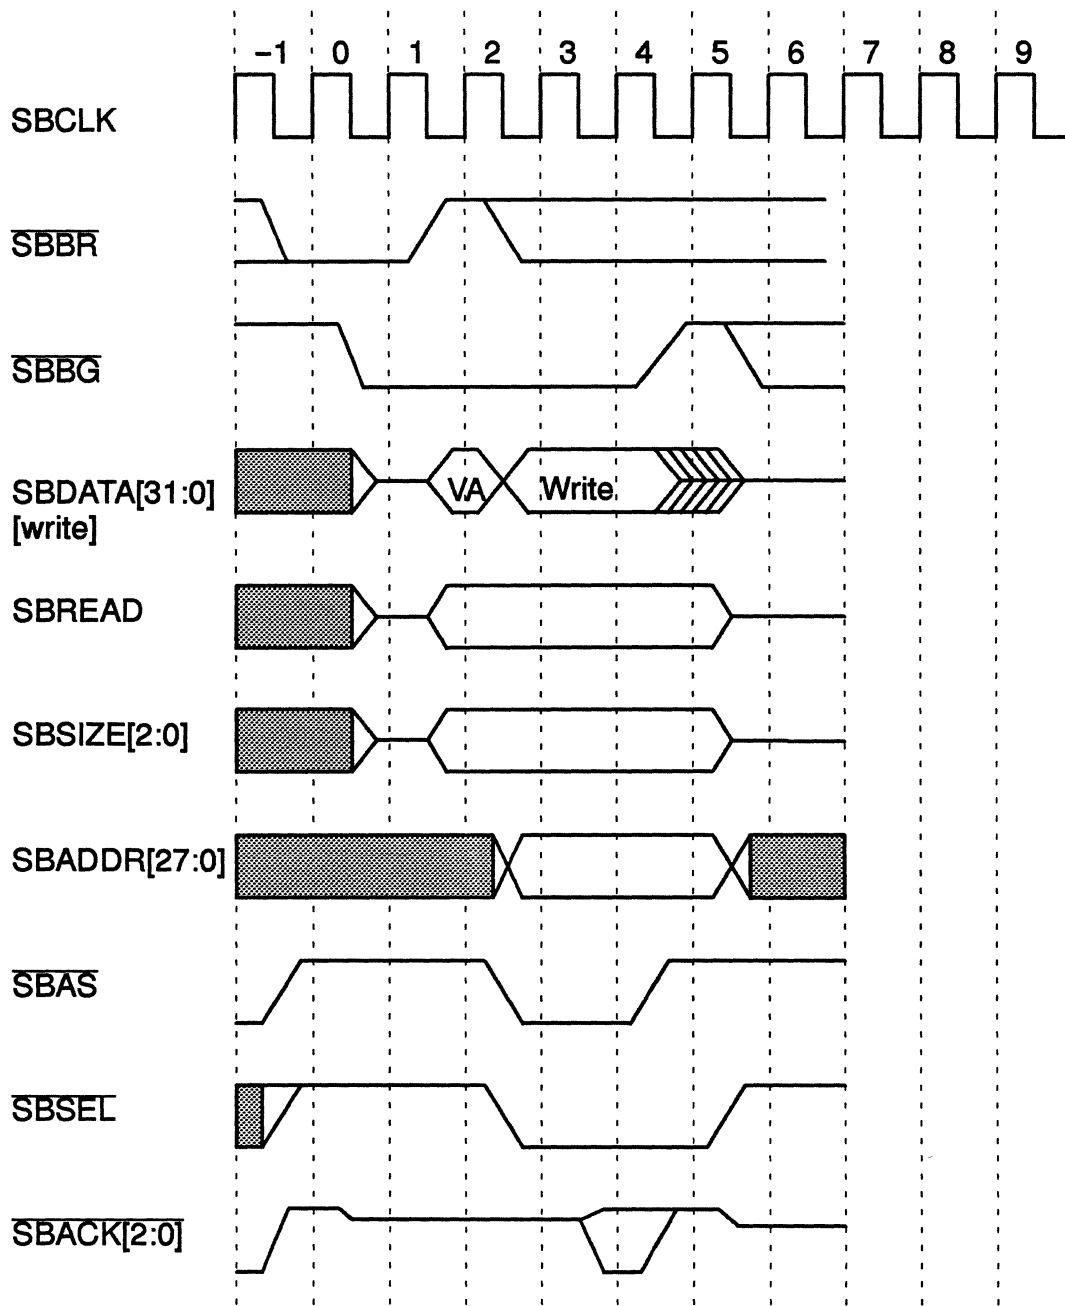


Figure 12. Translation Cycle and Slave Cycle




Figure 13. Basic Slave Cycle Timing

PRODUCT PREVIEW

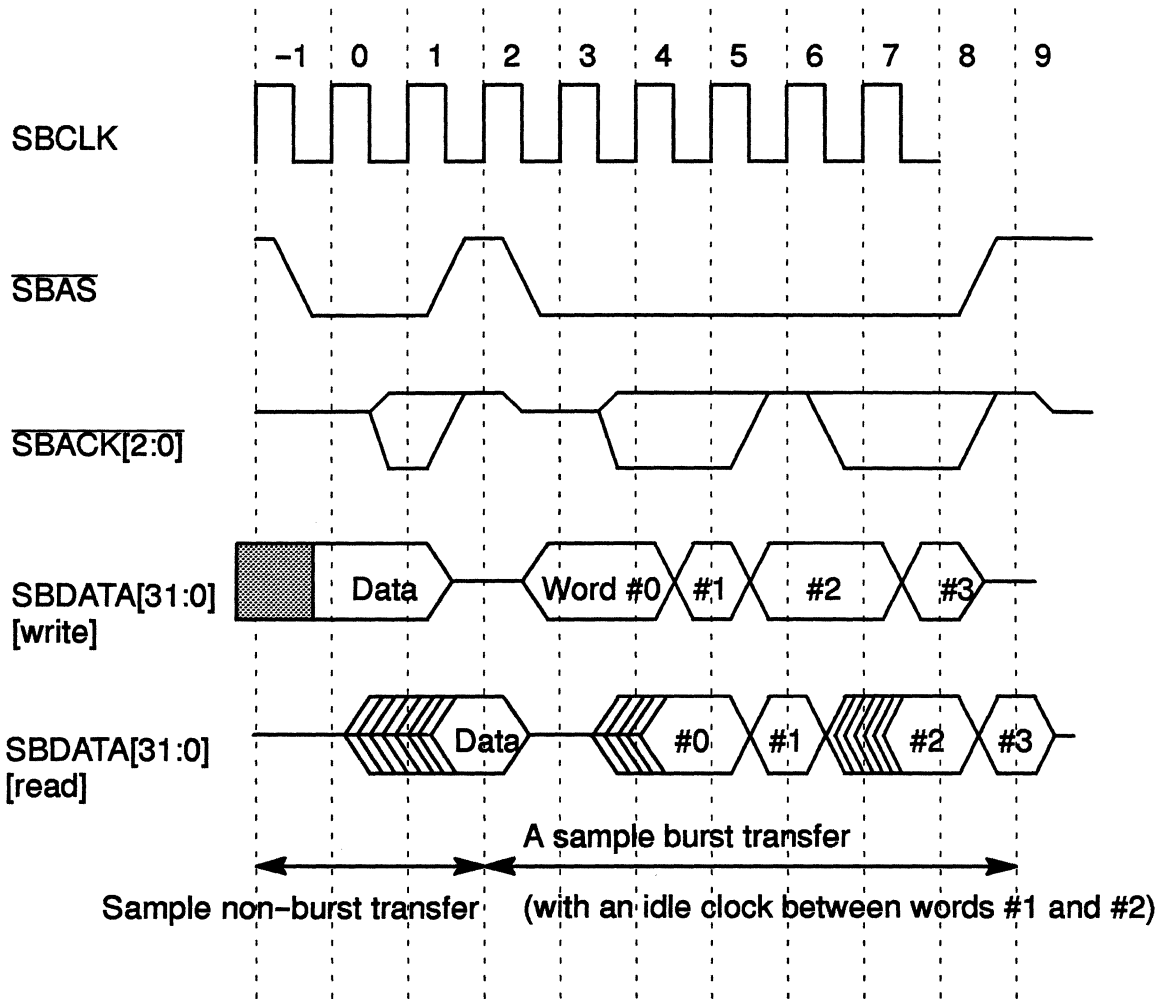


Figure 14. Sample Burst Transfer

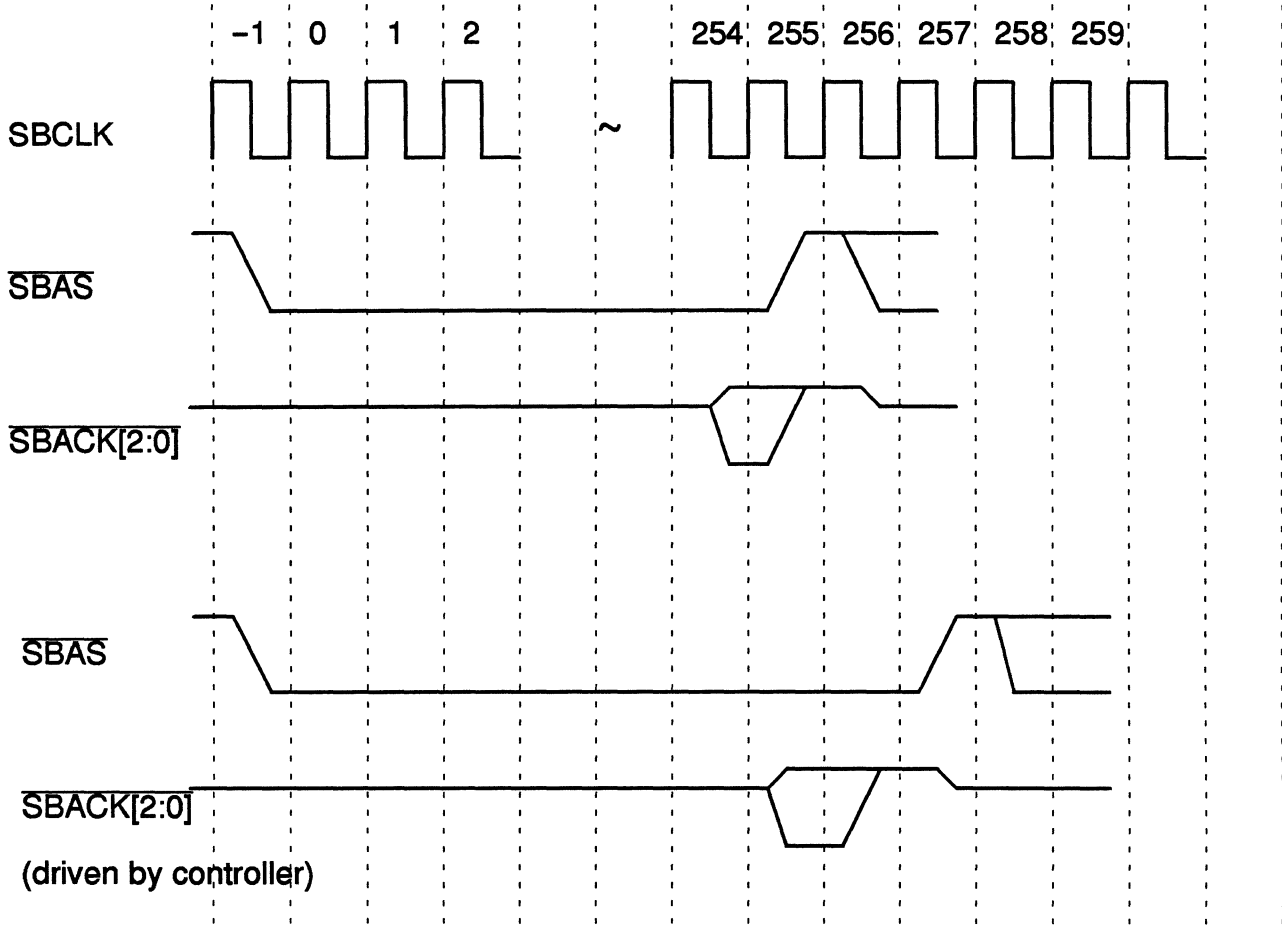


Figure 15. Bus Time-outs

PRODUCT PREVIEW




Figure 16. SBLERR Timing

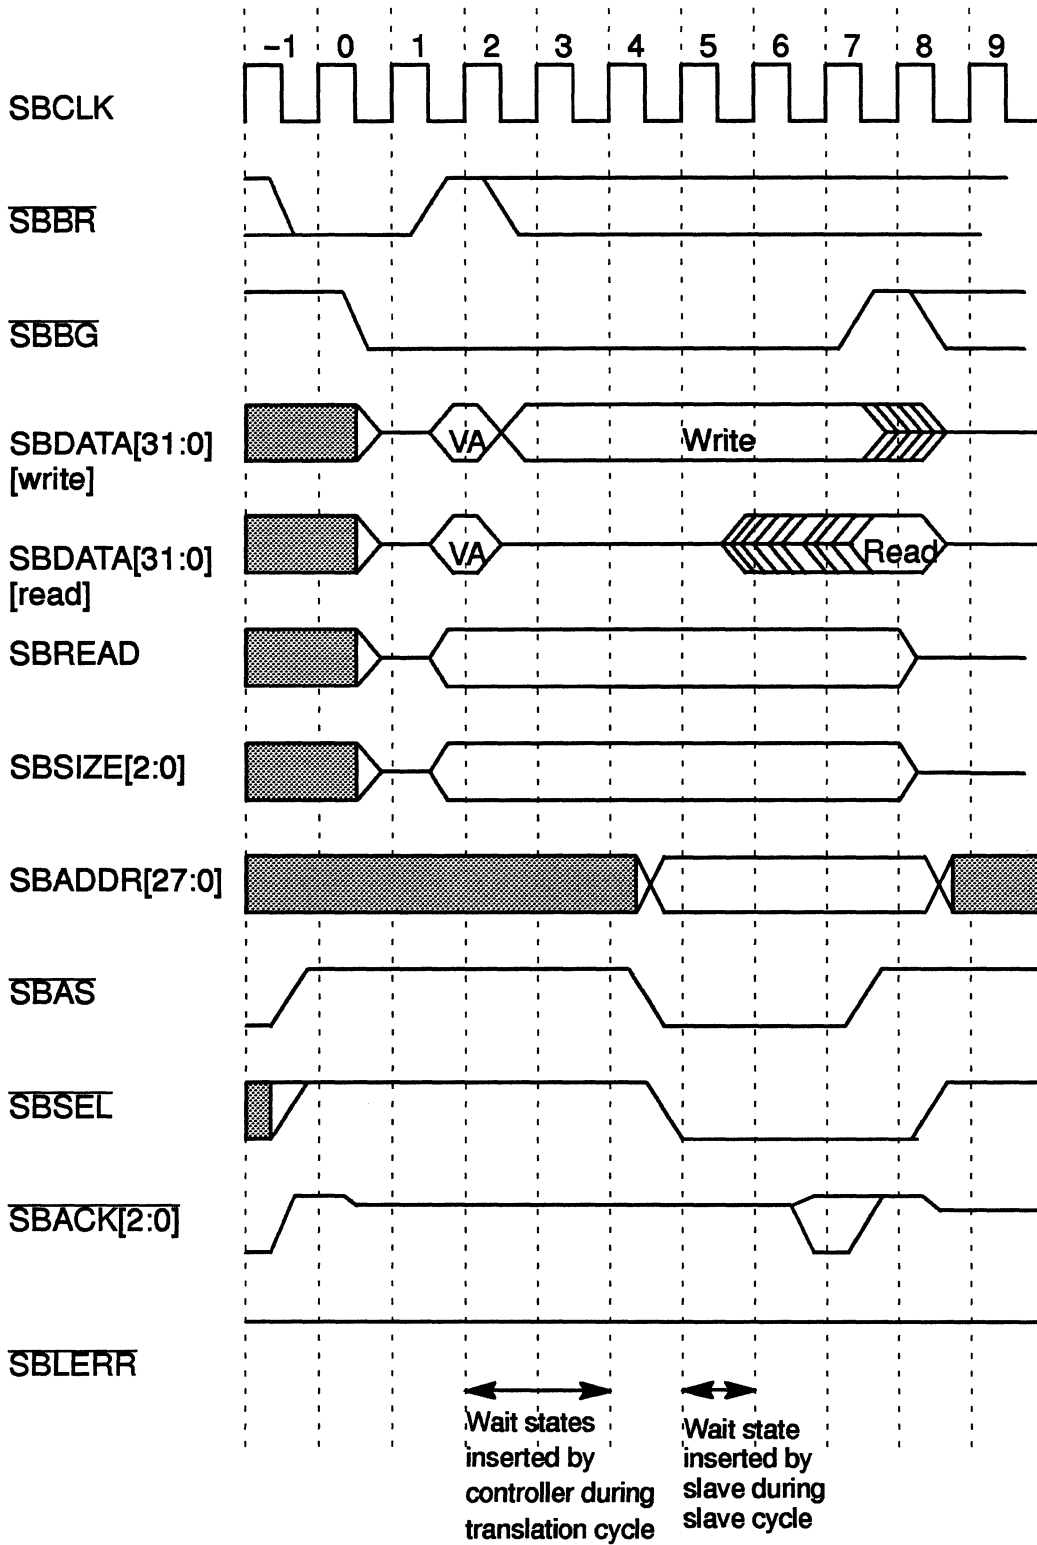


Figure 17. DVMA Cycle Wait States

PRODUCT PREVIEW

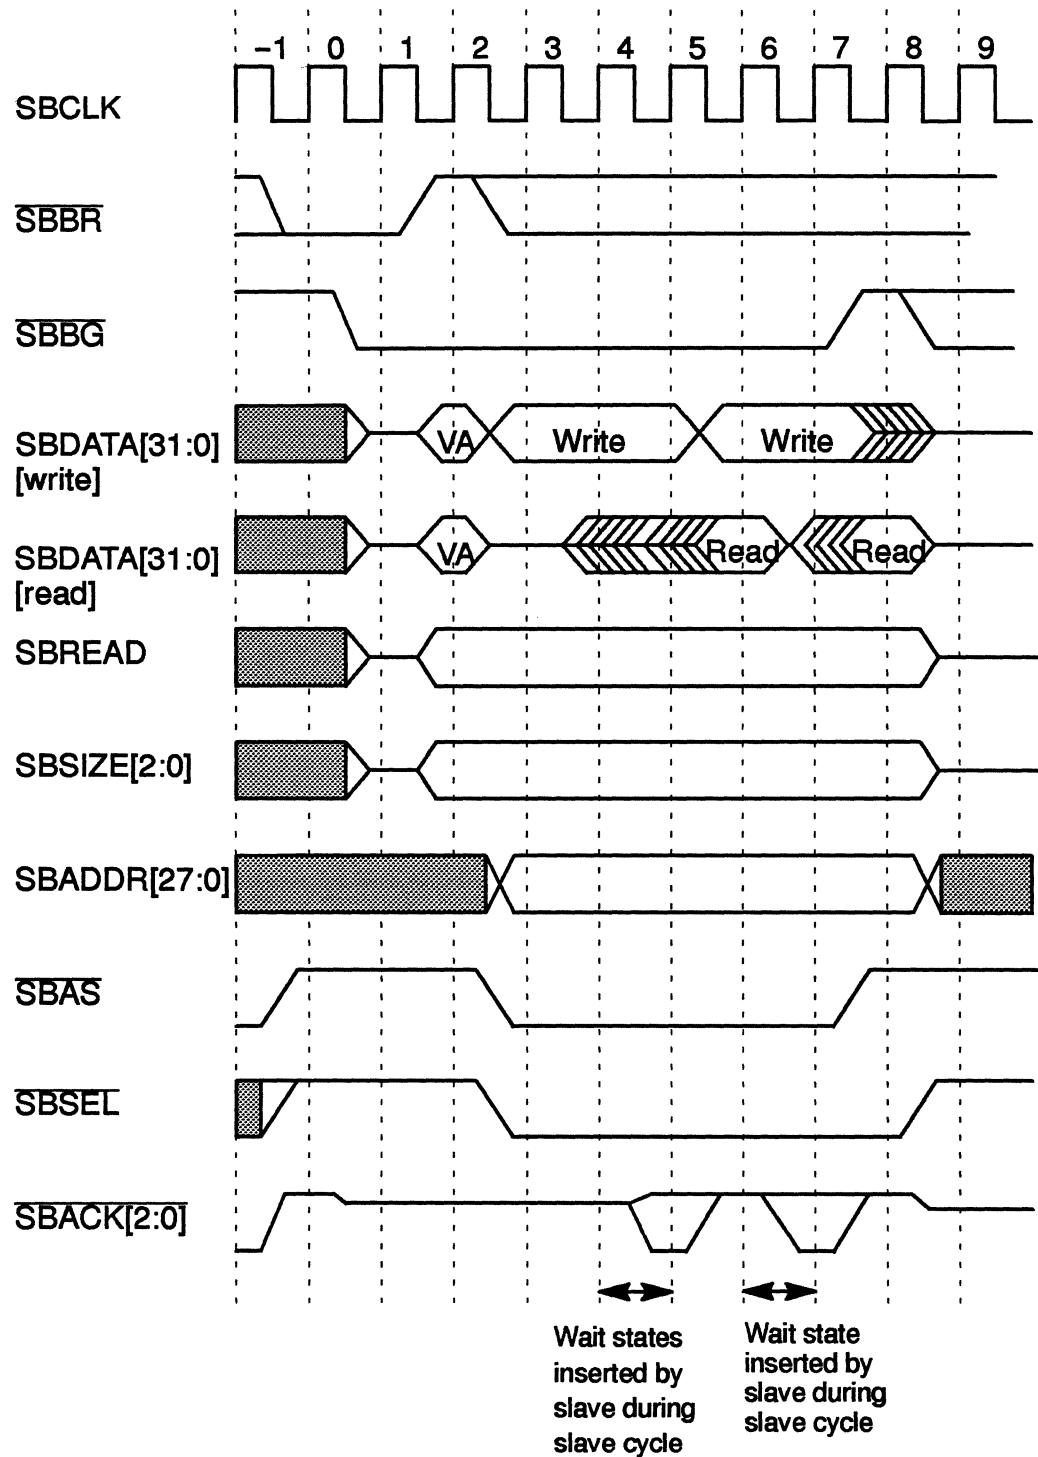


Figure 18. DVMA Burst Cycle Wait States

JTAG Test Bus Interface

The TMS390S10 has a five-wire Test Access Port(TAP) interface to support internal scan, boundary scan and clock control. This interface is compatible with the IEEE 1149.1 specification, "IEEE Standard Test Access Port and Boundary Scan Architecture". The TAP interfaces to the JTAG bus via five pins as per below.

Table 16. JTAG pins

PIN NAME	DIRECTION	DESCRIPTION
JTCK	in	Test Clock
JTMS	in	Test Mode Select
JTRST	in	Test TAP Reset (Asynchronous)
JTDI	in	Test Data In
JTDO	out	Test Data Out

The TAP supports a BYPASS instruction which places a minimum shift path(1 bit) between the chip's TDI and TDO pins. This allows efficient access to any single chip in the daisy-chain without board-level multiplexing.

TAP Controller

The TAP controller is a synchronous Finite State Machine (FSM) which controls the sequence of operations of the JTAG test circuitry, in response to changes at the JTAG bus. The TAP controller is asynchronous with respect to the system clock(s), and can therefore be used to control the clock control logic.

The TAP FSM implements the state (16 states) diagram as detailed in the 1149.1 protocol.

PRODUCT PREVIEW

Data Registers

The TMS390S10 TAP has following data registers:

- Bypass Register — a single bit shift register for efficient board level scan.
- Device ID Register — a 32-bit register. Details are shown in NO TAG.
- Clock Control Register (CCR) — a two bit clock control register to sample output from the clock controller.
- Boundary Scan Register — a single scan chain consisting of all of the boundary scan cells (input, output, and output enable cells).

TAP Instruction Register

The JTAG Instruction Register(IR) is a 6-bit serial register as shown in Figure 19. The instruction is used to determine what test is to be performed and/or what data register is to be accessed. The IR encoding is shown in Table 17. In accordance with IEEE 1149.1, the BYPASS, EXTEST, INTEST, SAMPLE and COMPONENT-ID instructions are public JTAG instructions. Sel-CCR is a private JTAG instruction.




Figure 19. Instruction Register Order of Scan

Table 17. JTAG Instruction Register encoding

IR VALUE	INSTRUCTION	SUMMARY	DATA REGISTERS
000000	EXTEST	Boundary Scan EXTEST	Boundary Scan Register
000011	ATEINTEST	Boundary Scan INTEST	Boundary Scan Register
000001	SAMPLE	Sample Boundary SAMPLE	Boundary Scan Register
100000	IDCODE	Get ID values	Device ID Register
111111	BYPASS	Bypass registers	Bypass Register
011110	Sel_CCR	Sample two bits from clock control logic.	Clock Control Register

BYPASS Instruction

The BYPASS instruction selects the BYPASS Register which is a single bit shift register. The BYPASS register can be used to reduce total length of scan chain when other devices on the same JTDI/JTDO chain are being accessed. When there are many chips on a board and in the scan chain, and when only a few chips need to be accessed, the other chips should have the Bypass register selected. The BYPASS Instruction can be scanned into all the other devices, thereby considerably reducing the length of the overall scan chain. When the Bypass test data register is selected, a single bit of logic zero will be sent to JTDO in the data shift state. Thereafter, JTDO will follow JTDI delayed by one JTCK.

PRODUCT PREVIEW

IDCODE Instruction

The IDCODE instruction selects the Device ID register, which is a 32-bit register. The Device ID register has the following format and contains the TMS390S10's assigned JTAG component identifier, 0x0000 202F. This value means version number of 0x0, a part number of 0x4, and a manufacturer ID of 0x17.

JTAG ID Register Contents

Version	Part Number	Manufacturer ID	1
31	28 27	12 11	1 0

FIELD NAME	DESCRIPTION	VALUE
Version	Version number. Incremented on component revisions.	0x0
Part Number	A component ID assigned by the manufacturer.	0x0004
Manufacturer ID	The identification number of the component's manufacturer.	0x17

SAMPLE Instruction

The SAMPLE instruction is used to sample data appearing at the TMS390S10 inputs and outputs without affecting normal device operation. The boundary scan register is selected in the scan path.

EXTEST Instruction

When the EXTEST instruction is selected, input data will be loaded into the boundary scan chain, and JTAG input data corresponding to output cells will force onto the TMS390S10 output pins on the falling edge of JTCK in the Update IR state. This data will change only on the falling edge of JTCK in the Update DR state. This instruction allows testing of board-level interconnections. All signals received at TMS390S10 input pins will be loaded into the boundary scan chain in the Capture Data Register state when this instruction is selected.

INTEST Instruction

Selection of the INTEST instruction allows testing of TMS390S10's internal logic. Data which has been loaded into the boundary scan chain and corresponding to input pins will be forced into TMS390S10's logic during the update IR state. This data will change only on the falling edge of JTCK in the Update DR state. Outputs of the logic will be captured into the boundary scan register in the Capture Data Register state. This data can be scanned out for analysis.

This instruction loads the boundary scan F/Fs after which, if it enters the "run_test_idle" state, the JTAG controller generates a JTCK pulse.

Sel_CCR Instruction

The Sel_CCR instruction selects CCR (Clock Controller). CCR is a 2-bit clock control register which is used to sample two internal state machine's signals.

Boundary Scan Register

The boundary scan register contains 287 bits as show in Table 18.

Table 18. Boundary Scan Registers

NO	REGISTER	DIR	NO	REGISTER	DIR	NO	REGISTER	DIR
1	INT_EVENT	O	49	SBDATA_out[3]	O	97	SBDATA_in[19]	I
2	CP_STAT[0]	O	50	SBADDR[4]	O	98	SBDATA_out[19]	O
3	CP_STAT[1]	O	51	SBDATA_in[4]	I	99	SBADDR[20]	O
4	EXT_EVENT	I	52	SBDATA_out[4]	O	100	SBDATA_in[20]	I
5	SBGR[0]	O	53	SBADDR[5]	O	101	SBDATA_out[20]	O
6	SBGR[1]	O	54	SBDATA_in[5]	I	102	SBADDR[21]	O
7	SBGR[2]	O	55	SBDATA_out[5]	O	103	SBDATA_in[21]	I
8	SBGR[3]	O	56	SBADDR[6]	O	104	SBDATA_out[21]	O
9	SBGR[4]	O	57	SBDATA_in[6]	I	105	SBADDR[22]	O
10	SBREAD_in	I	58	SBDATA_out[6]	O	106	SBDATA_in[22]	I
11	SBREAD_out	O	59	SBADDR[7]	O	107	SBDATA_out[22]	O
12	SBSIZE_in[0]	I	60	SBDATA_in[7]	I	108	SBADDR[23]	O
13	SBSIZE_out[0]	O	61	SBDATA_out[7]	O	109	SBDATA_in[23]	I
14	SBSIZE_in[1]	I	62	SBADDR[8]	O	110	SBDATA_out[23]	O
15	SBSIZE_out[1]	O	63	SBDATA_in[8]	I	111	SBADDR[24]	O
16	SBACK_OEN	E	64	SBDATA_out[8]	O	112	SBDATA_in[24]	I
17	SBSIZE_in[2]	I	65	SBADDR[9]	O	113	SBDATA_out[24]	O
18	SBSIZE_out[2]	O	66	SBDATA_in[9]	I	114	SBADDR[25]	O
19	SBSEL[0]	O	67	SBDATA_out[9]	O	115	SBDATA_in[25]	I
20	SBSEL[1]	O	68	SBADDR[10]	O	116	SBDATA_out[25]	O
21	SBSEL[2]	O	69	SBDATA_in[10]	I	117	SBADDR[26]	O
22	SBSEL[3]	O	70	SBDATA_out[10]	O	118	SBDATA_in[26]	I
23	SBSEL[4]	O	71	SBADDR[11]	O	119	SBDATA_out[26]	O
24	SBLERR	I	72	SBDATA_in[11]	I	120	SBADDR[27]	O
25	SBRO[0]	I	73	SBDATA_out[11]	O	121	SBDATA_in[27]	I
26	SBRO[1]	I	74	SBADDR[12]	O	122	SBDATA_out[27]	O
27	SBRO[2]	I	75	SBDATA_in[12]	I	123	SBDATA_in[28]	I
28	SBRO[3]	I	76	SBDATA_out[12]	O	124	SBDATA_out[28]	O
29	SBRO[4]	I	77	SBADDR[13]	O	125	SBDATA_in[29]	I
30	SBACK_in[0]	I	78	SBDATA_in[13]	I	126	SBDATA_out[29]	O
31	SBACK_out[0]	O	79	SBDATA_out[13]	O	127	SBDATA_in[30]	I
32	SBACK_in[1]	I	80	B_SIZE_OEN	E	128	SBDATA_out[30]	O
33	SBACK_out[1]	O	81	SBADDR[14]	O	129	SBDATA_in[31]	I
34	SBDATA_OEN	E	82	SBDATA_in[14]	I	130	SBDATA_out[31]	O
35	SBACK_in[2]	I	83	SBDATA_out[14]	O	131	MDATA_in[0]	I
36	SBACK_out[2]	O	84	SBADDR[15]	O	132	MDATA_out[0]	O
37	SBAS	O	85	SBDATA_in[15]	I	133	MDATA_in[1]	I
38	SBADDR[0]	O	86	SBDATA_out[15]	O	134	MDATA_out[1]	O
39	SBDATA_in[0]	I	87	SBADDR[16]	O	135	MDATA_in[2]	I
40	SBDATA_out[0]	O	88	SBDATA_in[16]	I	136	MDATA_out[2]	O
41	SBADDR[1]	O	89	SBDATA_out[16]	O	137	MDATA_in[3]	I
42	SBDATA_in[1]	I	90	SBADDR[17]	O	138	MDATA_out[3]	O
43	SBDATA_out[1]	O	91	SBDATA_in[17]	I	139	MDATA_in[4]	I
44	SBADDR[2]	O	92	SBDATA_out[17]	O	140	MDATA_out[4]	O
45	SBDATA_in[2]	I	93	SBADDR[18]	O	141	MDATA_in[5]	I
46	SBDATA_out[2]	O	94	SBDATA_in[18]	I	142	MDATA_out[5]	O
47	SBADDR[3]	O	95	SBDATA_out[18]	O	143	MDATA_in[6]	I
48	SBDATA_in[3]	I	96	SBADDR[19]	O	144	MDATA_out[6]	O

I: Input
O: Output
E: Output Enable

PRODUCT PREVIEW

PRODUCT PREVIEW documents contain information on products in the formative or design phase of development. Characteristic data and other specifications are design goals. Texas Instruments reserves the right to change or discontinue these products without notice.



Table 18. Boundary Scan Registers (continue)

NO	REGISTER	DIR	NO	REGISTER	DIR	NO	REGISTER	DIR
145	MDATA_in[7]	I	193	MDATA_in[31]	I	241	MDATA_out[53]	O
146	MDATA_out[7]	O	194	MDATA_out[31]	O	242	MDATA_in[54]	I
147	MDATA_in[8]	I	195	MPAR_in[0]	I	243	MDATA_out[54]	O
148	MDATA_out[8]	O	196	MPAR_out[0]	O	244	MDATA_in[55]	I
149	MDATA_in[9]	I	197	MDATA_in[32]	I	245	MDATA_out[55]	O
150	MDATA_out[9]	O	198	MDATA_out[32]	O	246	MDATA_in[56]	I
151	MDATA_in[10]	I	199	B_MEM_OEN	E	247	MDATA_out[56]	O
152	MDATA_out[10]	O	200	MDATA_in[33]	I	248	MDATA_in[57]	I
153	MDATA_in[11]	I	201	MDATA_out[33]	O	249	MDATA_out[57]	O
154	MDATA_out[11]	O	202	MDATA_in[34]	I	250	MDATA_in[58]	I
155	MDATA_in[12]	I	203	MDATA_out[34]	O	251	MDATA_out[58]	O
156	MDATA_out[12]	O	204	MDATA_in[35]	I	252	MDATA_in[59]	I
157	MDATA_in[13]	I	205	MDATA_out[35]	O	253	MDATA_out[59]	O
158	MDATA_out[13]	O	206	MDATA_in[36]	I	254	MDATA_in[60]	I
159	MDATA_in[14]	I	207	MDATA_out[36]	O	255	MDATA_out[60]	O
160	MDATA_out[14]	O	208	MDATA_in[37]	I	256	MDATA_in[61]	I
161	MDATA_in[15]	I	209	MDATA_out[37]	O	257	MDATA_out[61]	O
162	MDATA_out[15]	O	210	MDATA_in[38]	I	258	MDATA_in[62]	I
163	MDATA_in[16]	I	211	MDATA_out[38]	O	259	MDATA_out[62]	O
164	MDATA_out[16]	O	212	MDATA_in[39]	I	260	MDATA_in[63]	I
165	MDATA_in[17]	I	213	MDATA_out[39]	O	261	MDATA_out[63]	O
166	MDATA_out[17]	O	214	MDATA_in[40]	I	262	MPAR_in[1]	I
167	MDATA_in[18]	I	215	MDATA_out[40]	O	263	MPAR_out[1]	O
168	MDATA_out[18]	O	216	MDATA_in[41]	I	264	MWE	O
169	MDATA_in[19]	I	217	MDATA_out[41]	O	265	MCAS[0]	O
170	MDATA_out[19]	O	218	MDATA_in[42]	I	266	MCAS[1]	O
171	MDATA_in[20]	I	219	MDATA_out[42]	O	267	MRAS[0]	O
172	MDATA_out[20]	O	220	MDATA_in[43]	I	268	MRAS[1]	O
173	MDATA_in[21]	I	221	MDATA_out[43]	O	269	MRAS[2]	O
174	MDATA_out[21]	O	222	MDATA_in[44]	I	270	MRAS[3]	O
175	MDATA_in[22]	I	223	MDATA_out[44]	O	271	MADDR[0]	O
176	MDATA_out[22]	O	224	MDATA_in[45]	I	272	MADDR[1]	O
177	MDATA_in[23]	I	225	MDATA_out[45]	O	273	MADDR[2]	O
178	MDATA_out[23]	O	226	MDATA_in[46]	I	274	MADDR[3]	O
179	MDATA_in[24]	I	227	MDATA_out[46]	O	275	MADDR[4]	O
180	MDATA_out[24]	O	228	MDATA_in[47]	I	276	MADDR[5]	O
181	MDATA_in[25]	I	229	MDATA_out[47]	O	277	MADDR[6]	O
182	MDATA_out[25]	O	230	MDATA_in[48]	I	278	MADDR[7]	O
183	MDATA_in[26]	I	231	MDATA_out[48]	O	279	MADDR[8]	O
184	MDATA_out[26]	O	232	MDATA_in[49]	I	280	MADDR[9]	O
185	MDATA_in[27]	I	233	MDATA_out[49]	O	281	MADDR[10]	O
186	MDATA_out[27]	O	234	MDATA_in[50]	I	282	IRQ[0]	I
187	MDATA_in[28]	I	235	MDATA_out[50]	O	283	IRQ[1]	I
188	MDATA_out[28]	O	236	MDATA_in[51]	I	284	IRQ[2]	I
189	MDATA_in[29]	I	237	MDATA_out[51]	O	285	IRQ[3]	I
190	MDATA_out[29]	O	238	MDATA_in[52]	I	286	RESET	I
191	MDATA_in[30]	I	239	MDATA_out[52]	O	287	MADDR[11]	O
192	MDATA_out[30]	O	240	MDATA_in[53]	I			

I: Input
O: Output
E: Output Enable

PRODUCT PREVIEW

ASI Information

The Address Space Identifier (ASI) is appended to the virtual address by the SPARC IU when it accesses memory. The ASI encodes whether the processor is in supervisor or user mode, whether an access is to instruction or data memory, and is used to perform other internal CPU functions.

Table 19. ASIs Supported by the TMS390S10

ASI	FUNCTION	ACCESS	SIZE
00	Reserved	—	—
01—02	Unassigned	—	—
03	Ref MMU Flush/Probe	Read/Write	Single
04	MMU Registers	Read/Write	Single
05	Unassigned	—	—
06	Ref MMU Diagnostics	Read/Write	Single
07	Unassigned	—	—
08	User Instruction	Read/Write	All
09	Supervisor Instruction	Read/Write	All
0A	User Data	Read/Write	All
0B	Supervisor Data	Read/Write	All
0C	Instruction Cache Tag	Read/Write	Single
0D	Instruction Cache Data	Read/Write	Single
0E	Data Cache Tag	Read/Write	Single
0F	Data Cache Data	Read/Write	Single
10—14	Unassigned	—	—
15—16	Reserved	—	—
17—1C	Unassigned	—	—
1D—1E	Reserved	—	—
1F	Unassigned	—	—
20	Ref MMU Bypass	Read/Write	All
21—2F	Reserved	—	—
30—35	Unassigned	—	—
36	Instruction Cache Flash Clear	Write	Single
37	Data Cache Flash Clear	Write	Single
38	Unassigned	—	—
39	Data Cache Diagnostics Register Access	Read/Write	Single
3A—3F	Unassigned	—	—
40—FF	Reserved	—	—

PRODUCT PREVIEW

Table 21. CPU TLB Entry Probing

TYPE	PROBE	RETURNED DATA
0	Page	Level 3 PTE or 0
*1	Segment	Level 2 PTE or 0
*2	Region	Level 1 PTE or 0
*3	Context	Level 0 PTE or 0
4	Entire	PTE from Table Walk or 0
5—F	Reserved	

*: Must Proceed by TLB Flush Entire.

ASI = 0x04 MMU Registers

— This space is used to read and write internal MMU registers using the Virtual Address to reference them. Only single word accesses should be used, others will result in an error.

Table 22. Address Map for MMU Registers

VA[12:08]	REGISTER
00	Control Register
01	Context Table Pointer Register
02	Context Register
03	Synchronous Fault Status Register
04	Synchronous Fault Address Register
05—0F	Reserved
10	TLB Replacement Control Register
11—12	Reserved
13	Synchronous Fault Status Register
14	Synchronous Fault Address Register
15—1F	Reserved

VA[31:13] are zero. VA bits [07:00] are ignored and should be set to zero by software.

ASI = 0x05 Unassigned

— This space is unassigned and may be used in the future.

ASI = 0x06 Ref MMU Diagnostics

— Diagnostic reads and writes can be made to the 32 TLB entries and the Instruction Translation Buffer Register using the virtual address to specify which entry and whether the PTE or Tag section is to be referenced.

ASI = 0x07 Unassigned

— This space is unassigned and may be used in the future.

ASI = 0x08 User Instruction

— This space is defined and reserved by SPARC for user instruction.

PRODUCT PREVIEW

ASI = 0x09 **Supervisor Instruction**

— This space is defined and reserved by SPARC for supervisor instruction.

ASI = 0x0A **User Data**

— This space is defined and reserved by SPARC for user data.


ASI = 0x0B **Supervisor Data**

— This space is defined and reserved by SPARC for supervisor data.

ASI = 0x0C **Instruction Cache Tag**

— This space is used for reading and writing instruction cache tags by using the LDA and STA instructions at virtual addresses in the range of 0x0 to 0x0FFF on modulo-32 boundaries.

Instruction Cache Tag Entry



Bits [31:27,11:01] are not implemented, should be written 0 and will be read as 0.


ASI = 0x0D **Instruction Cache Data**

— This space is used for reading and writing instruction cache data by using the LDA and STA instructions at virtual addresses in the range of 0x0 to 0x0FFF.

ASI = 0x0E **Data Cache Tag**

— This space is used for reading and writing data cache tags by using the LDA and STA instructions at virtual addresses in the range of 0x0 to 0x03FF on modulo-16 boundaries.

Data Cache Tag Entry

**ASI = 0x0F** **Data Cache Data**

— This space is used for reading and writing data cache data by executing doubleword, word, halfword, or byte load or store alternate space instructions (respectively) in ASI 0xF at virtual addresses in the range of 0x0 to 0x03FF.

ASI = 0x10—0x14 **Unassigned**

— This space is unassigned and may be used in the future.

ASI = 0x15—0x16 **Reserved**

— This space is architecturally reserved.

ASI = 0x17—0x1C **Unassigned**

— This space is unassigned and may be used in the future.

ASI = 0x1D—0x1E *Reserved*

— This space is architecturally reserved.

ASI = 0x1F *Unassigned*

— This space is unassigned and may be used in the future.

ASI = 0x20 *Ref MMU Bypass*

— This space can be used to access an arbitrary physical address.

It is particularly useful before the MMU or main memory has been initialized. The MMU does not perform an address translation rather a physical address is formed from the least significant 31 bits of the Virtual Address (PA[30:00] := VA[30:00]).

ASI = 0x21—0x2F *Reserved*

— This space is architecturally reserved.

ASI = 0x30—0x35 *Unassigned*

— This space is unassigned and may be used in the future.

ASI = 0x36 *Instruction Cache Flush Clear*

— The instruction cache is completely flushed by any type of alternate store instruction to this ASI. All instruction cache valid bits are reset (to zero) by this operation.

The pipeline is not flushed by this STA as it would be on a SPARC FLUSH instruction.

ASI = 0x37 *Data Cache Flush Clear*

— The data cache is completely flushed by any type of alternate store instruction to this ASI. All data cache valid bits are reset (to zero) by this operation.

ASI = 0x38 *Unassigned*

— This space is unassigned and may be used in the future.

ASI = 0x39 *Data Cache Diagnostic Register Access*

— This space is used to read and write the internal Data Cache Registers. VA[08] is also used to select from between WRB0 and WRB1. Single word accesses only should be used, others result in an internal error. The Virtual Address maps to these registers:

Table 23. Address Map for Data Cache Registers

VA [08]	REGISTER
0	Write Buffer 0
1	Write Buffer 1

VA bits [31:09] are zero. VA bits [07:00] are ignored and should be set to zero by software.

ASI = 0x3A—0x3F *Unassigned*

— This space is unassigned and may be used in the future.

ASI = 0x40—0xFF *Reserved*

— Since the 2 high order bits are not decoded these encodings should not be used.

PRODUCT PREVIEW

Registers

Followings are descriptions of all TMS390S10 registers. There are four categories: IU, FPU, MMU and I/O MMU registers.

IU Registers

The TMS390S10 Integer Unit (IU) contains several control/status registers and 120 *r* registers (112 window registers and 8 global registers).

The 112 window registers are grouped into 7 sets of 24 *r* registers called windows.

Register Windows

At any given time, an instruction can access a total of 32 registers: 8 globals and a 24-register window. Register windows are shown in Figure 20.

Each register window has 3 groups of registers: 8 *ins*, 8 *locals*, and 8 *outs*. Registers are addressed as shown below.

Table 24. Register Windows

WINDOWED REGISTER ADDRESS	R REGISTER ADDRESS
in[0]—in[7]	r[24]—r[31]
local[0]—local[7]	r[16]—r[23]
out[0]—out[7]	r[8]—r[15]
global[0]—global[7]	r[0]—r[7]

The current window into the *r* registers is indicated by the Current Window Pointer (CWP) bits of Processor State Register (PSR).

The CWP is incremented by a RESTORE or RETT instruction and decremented by a SAVE instruction or a trap.

Window Overflow and Underflow

The TMS390S10 has seven register windows, and it is possible that the program will try to use more than seven windows. To prevent the overwriting of the oldest registers by subroutine calling, when the program tries to use more than seven windows, window overflow is detected and trapped. Window overflow and underflow are detected by using the window invalid mask (WIM) register which is controlled by the supervisor. The overflow and underflow traps are generated by SAVE, RESTORE, and RETT instructions.

Overlapping of Windows

Each window shares its *ins* and *outs* with the adjacent windows. See Figure 20.

The *outs* of the CWP+1 window can be accessed as *ins* of current window, and the *outs* in the current window are the *ins* of the CWP-1 window. The highest window 6 overlaps with window 0. The *outs* of window 6 are the *ins* of window 0. This mechanism provides very efficient method for parameter passing between two procedures.

The parameters can be passed by placing them into the *outs* registers. Also, executing a CALL instruction writes its own address (maybe usable as a return address) into register r[15] (out register 7) of the calling procedure's window. A SAVE instruction decrements CWP by one (modulo 7). Now, the calling procedure's *outs* become the called procedure's *ins* and the parameters can be accessed directly.

An *r* register with address *o*, where $8 \leq o \leq 15$, refers to exactly the same register as $(o+16)$ does after the CWP is decremented by 1 (modulo 7). Likewise, a 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 incremented by 1 (modulo 7).

PRODUCT PREVIEW

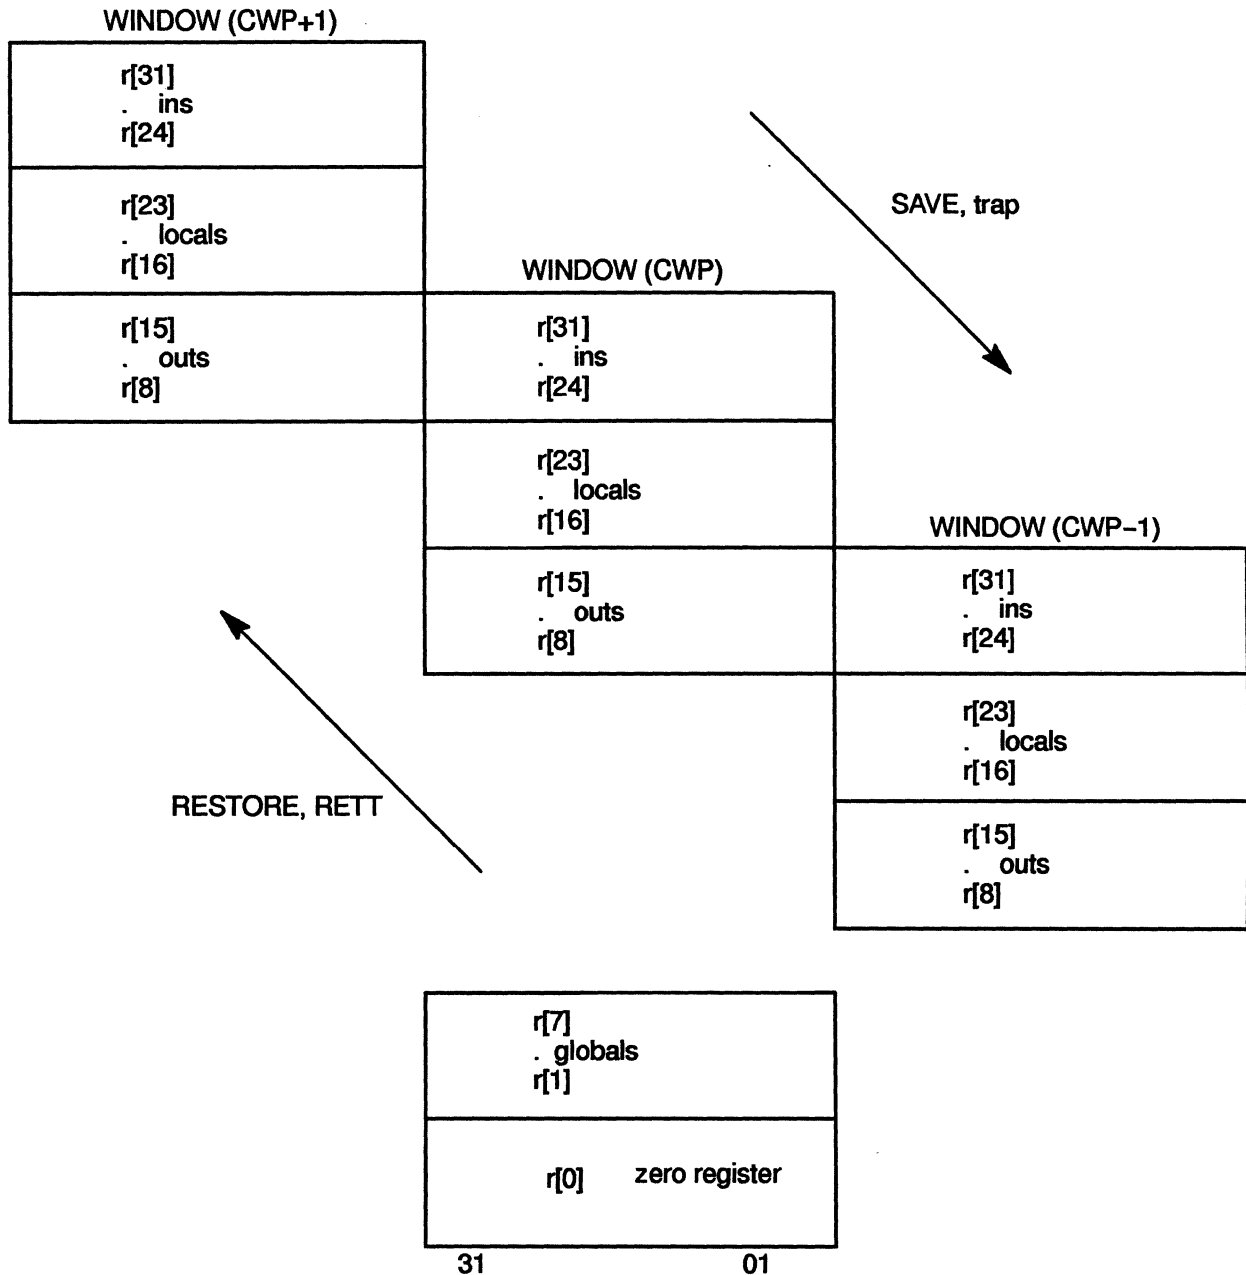



Figure 20. Three Overlapping Windows and the 8Global Registers

Doubleword Operands

The doubleword data is located with even-odd address pair, such as (%g0, %g1) in the *r* register file.



Special *r* registers

Generally, any register which is in the current register window can be used for any purpose, but the following three cases are exceptions:

- Register *r*[0] is hardwired to zero. If used as a source register, the value of zero is read. If used as a destination register, the data written is discarded.
- The CALL instruction writes its own address into register *r*[15] (out register7) of the calling procedure's window.
Its value can be used as a return address if a SAVE instruction follows it.
- When a trap occurs, registers *r*[17] and *r*[18] (*local* registers 1 and 2) of the trap's new register window are used to save the contents of program counters PC and nPC.

Control/Status Registers

The TMS390S10 has several control/status registers such as the Processor State Register (PSR), the Window Invalid Mask register (WIM), the Trap Base Register (TBR), and the multiply/divide (Y) register. All control/status registers are 32 bits wide and are located in IU.

Processor State Register(PSR)

The 32-bit PSR contains various fields that control the processor and hold status information.

It can be modified by the SAVE, RESTORE, Ticc, RETT instructions, all instructions that modify the condition codes, internally detected traps, or external interrupts. The privileged RDPSR and WRPSR instructions can read from and write to the PSR directly.

PSR fields

impl	ver	icc	reserved	EC	EF	PIL	S	PS	ET	CWP
31 28	27 24	23 20	19 14	13	12	11 8	7	6	5	4 0

Each field has following meaning.

PSR_implementation (impl) and PSR_version (Ver)

Bits 31 through 24 contain the information of implementation of the TMS390S10. The implementation number of the TMS390S10 is 0x4 (b'0100) and a version number is 0x1 (b'0001).

PSR_integer_cond_codes (icc)

Bits 23 through 20 are the IU's condition codes.

These bits are modified by the arithmetic and logical instructions whose names end with the letters cc (e.g., ANDcc), and by the WRPSR instruction. The Bicc and Ticc instructions cause a transfer of control based on the value of these bits, which are defined as follows:

Integer Condition Codes (icc) Fields of the PSR

n	z	v	c
23	22	21	20

PSR_negative (n)

Bits 23 indicates whether the 32-bit 2's complement ALU result was negative for the last instruction that modified the icc.

1	negative
0	not negative

PSR_zero (z)

Bit 22 indicates whether the 32-bit ALU result was zero for the last instruction that modified the icc field.

1	zero
0	nonzero

PRODUCT PREVIEW

PSR_overflow (v)

Bit 21 indicates whether the ALU result was representable in 32-bit 2's complement notation for the last instruction that modified the icc field.

1	overflow
0	no overflow

PSR_carry (c)

Bit 20 indicates whether a 2's complement carry out (or borrow) occurred for the last instruction that modified the icc field.

Carry is set on addition if there is a carry out of bit 31. Carry is set on subtraction if there is borrow into bit 31.

1	carry
0	no carry

PSR_reserved

Bits 19 through 14 are reserved.

When read by a RDPSR instruction, these bits are read as zeros. When the supervisor issues WRPSR, these bits must be set as zeros.

PSR_enable_coprocessor (EC)

Bit 13 determines whether the Coprocessor is enabled.

The TMS390S10 doesn't supports coprocessor, this bit is implemented to maintain compatibility with SPARC Architecture Version 8.

1	enabled
0	disable

PSR_enable_floating_point (EF)

Bit 12 determines whether the FPU is enabled. If disabled, a floating point instruction will trap.

1	enabled
0	disable

PSR_proc_interrupt_level (PIL)

Bits 11 (the most significant bit) through 8 (the least significant bit) identify the interrupt level above which the processor will accept an interrupt.

PSR_supervisor (S)

Bit 7 determines whether the processor is in supervisor or user mode.

1	supervisor mode
0	user mode

PSR_Previous_supervisor (PS)

Bit 6 contains the value of the S bit at the time of the most recent trap.

1	supervisor mode
0	user mode

PSR_enable_traps (ET)

Bit 5 determines whether traps are enabled.

A trap automatically resets ET to 0. When ET=0, an interrupt request is ignored and an exception trap causes the IU to halt execution, which typically results in a reset trap that resumes execution at address 0.

1	traps enabled
0	traps disabled


PSR_current_window_pointer (CWP)

Bits 4 (the most significant bit) through 0 (the least significant bit) comprise the current window pointer, a counter that identifies the current window into the *r* registers. The hardware decrements the CWP on traps and SAVE instructions, and increments it on RESTORE and RETT instructions (modulo 7).

Window Invalid Mask register (WIM)

The Window Invalid Mask register (WIM) is controlled by the supervisor and is used by hardware to determine whether a window overflow or underflow trap is to be generated by a SAVE, RESTORE, or RETT instruction.

WIM Fields



Each bit in the WIM represents active status of its corresponding window or register set.

If it is set to 1, the corresponding window is marked as invalid.

When a SAVE, RESTORE, or RETT instruction executes, the current value of the CWP is compared with the WIM.

If the SAVE, RESTORE, or RETT instruction causes the CWP to point a register set whose corresponding WIM bit equals 1, that is, WIM[CWP]=1, this would generate a window_overflow or window_underflow trap. The WIM can be accessed by the privileged instructions RDWIN to read and WRWIM to write.


Bits corresponding to unimplemented windows (W31—W7) read as zeros and values written to unimplemented bits are unused.

PRODUCT PREVIEW

Trap Base Register (TBR)

The Trap Base Register (TBR) contains three fields that together equal to the address to which control is transferred when a trap occurs.

TBR Fields

*TBR_trap_base_address (TBA)*

Bits 31 through 12 are the trap base address, which is established by supervisor software.

It contains the most-significant 20 bits of the trap table address. The TBA field is written by the privileged instruction WRTBR.

TBR_trap_type (tt)

Bits 11 through 4 comprise the trap type (tt) field.

This 8-bit field is written by the hardware when a trap occurs, and retains its value until the next trap. It provides an offset to the trap table. The WRTBR instruction does not affect the tt field.

TBR_zero (zero)

Bits 3 through 0 are zeros.

The WRTBR instruction does not affect this field. For future compatibility, the supervisor should only issue a WRTBR instruction with a zero value in this field.

Multiply/Divide Register (Y)

The 32-bit Y register contains the most significant word of the double-precision product of an integer multiplication, as a result of either an integer multiply (SMUL, SMULcc, UMUL, UMULcc) instruction, or of a routine that uses the integer multiply step (MULScc) instruction. The Y register also holds the most significant word of the double-precision dividend for an integer divide (SDIV, SDIVcc, UDIV, UDIVcc) instruction.

The Y register can be read and written with the RDY and WRY instructions, respectively.

FPU Registers

The TMS390S10 FPU contains several control/status registers and 32 *f* registers.

FPU *f* Registers

The TMS390S10 FPU contains 32 32-bit floating-point *f* registers, *f*[0] to *f*[31]. All registers can be accessed at any time. The *f* registers can be read and written with FPop (FPop1/FPop2 format) instructions and by load/store single/double floating-point instructions (LDF, LDDF, STF, STDF).

The *f* Registers

The *f* Registers

<i>f</i> [31]
<i>f</i> [30]
.
.
<i>f</i> [1]
<i>f</i> [0]

Double Operands

A single *f* register can hold one 32-bit single-precision operand.

A double-precision operand requires an aligned pair of *f* registers. So, at a given time, the *f* registers can hold a maximum of 32 single-precision or 16 double-precision operands.

The least-significant bit of a doubleword *f* register address specifier is reserved and should be set to zero by software.

Table 25. Floating-Point Doubles in Registers

SUB-FORMAT NAME	FORMAT FIELDS	<i>f</i> REGISTER ADDRESS
FD—0	s:exp[10:0]:fraction[51:32]	0 mod 2
FD—1	fraction[31:0]	1 mod 2

PRODUCT PREVIEW

FPU Control/Status Registers (FSR)

The FSR register fields contain FPU mode and status information. The FSR is read and written by the STFSR and LDFSR instruction.

FSR Fields

RD	res	TEM	NS	res	ver	flt	qne	res	fcc	aexc	cexc
31 30	29 28	27 23	22	21 20	19 17	16 14	13	12	11 10	9 5	4 0

FSR_rounding_direction (RD)

Bits 31 and 30 select the rounding direction of floating-point results according to ANSI/IEEE Standard 754-1985.

Table 26. Rounding Direction (RD) Field of FSR

RD	ROUND TOWARD:
0	Nearest (even, if tie)
1	0
2	+∞
3	-∞

FSR_trap_enable_mask (TEM)

Bits 27 through 23 are enable bits for each of the five floating-point exceptions that can be indicated in the current_exception field (cexc). Please refer to aexc and cexc from page NO TAG.

Trap Enable Mask Fields of FSR

NVM	OFM	UFM	DZM	NXM
27	26	25	24	23

FSR_nonstandard_fp (NS)

Bit 22 always 0.

FSR_reserved (res)

Bits 29, 28, 21, 20, and 12 are reserved.

When read by STFSR instruction, these bits are read as zeros. For future compatibility, software should only issue LDFSR instructions with zero values in these bits.

FSR_version (ver)

Bits 19 through 17 identify one or more particular implementations of the FPU architecture. For the TMS390S10 these bits are set to 4 (b'100).

PRODUCT PREVIEW

FSR_floating-point_trap_type (ftt)

Bits 16 through 14 identify floating-point exception trap types.

After an *fp_exception* trap occurs, the *ftt* field encodes the type of exception, until an STFSR or another FPop occurs.

Note that supervisor-mode software which handles floating-point traps must execute an STFSR to determine the floating-point trap type. The STFSR hardware implementation does not zero *ftt*, therefore the trap software must ensure that a subsequent STFSR from user mode shows only a zero FSR. LDFSR cannot be used for this purpose since it leaves *ftt* unchanged, although executing a non-trapping FPop such as "fmovs %f0, %f0" prior to user mode will zero the *ftt*.

Table 27. Floating-point trap type (*ftt*) Field of FSR

ftt	TRAP TYPE
0	None
1	IEEE 754 exception
2	unfinished FPop
3	unimplemented FPop
4	sequence error
5,6,7	reserved

The *sequence_error* and *hardware_error* are not expected to arise in the normal course of computation.

They are essentially unrecoverable from the view point of user applications.

In contrast, *IEEE_754_exception* and *unimplemented_FPop* are expected to arise occasionally in the normal course of computation and must be recoverable by supervisor software. Those traps facilitate recovery as follows:

- The value of *aexc* is unchanged.
- The value of *cexc* is unchanged, except that exactly one bit corresponding to the trapping exception will be set on an *IEEE_754_exception*.
- The source *f* registers are unchanged (usually implemented by leaving the destination *f* register unchanged).
- The value of *fcc* is normally unchanged, but when an *FCMP* or *FCMPE* instruction traps the value of *fcc* is undefined.

The foregoing describes the result seen by a user handler if an IEEE exception is signaled, either immediately from an *IEEE_754_exception* or after recovery from an *unimplemented_FPop*. In either case, *cexc* as seen by the trap handler will reflect the exception causing the trap.

In the case of *unimplemented_FPop* traps that don't subsequently generate IEEE exceptions, the recovery software is expected to define *cexc*, *aexc* and either the destination *f* register or *fcc*, as appropriate.

ftt = IEEE_754_exception

An *IEEE_754_exception* floating-point trap type indicates that a floating-point exception occurred that conforms to the ANSI/IEEE Standard 754-1985. The exception type is encoded in the *cexc* field.

Note that *aexc* and the destination *f* register are not affected by an *IEEE_754_exception* trap.

ftt = unimplemented_FPop

An *unimplemented_FPop* indicates that the an FPU decoded an FPop that it does not implement.

PRODUCT PREVIEW

PRODUCT PREVIEW

In this case, the *cexc* field is unchanged.

Note that in the case of unimplemented_FPop floating-point trap type, software should emulate or re-execute the exception-causing instruction, and update the FSR, destination *f* register(s), and *fcc*.

ftt = sequence_error

A *sequence_error* indicates an attempt to execute a non-store floating-point instruction while an FPU with a floating-point deferred-trap queue was in exception mode waiting for the FQ to be emptied by supervisor software.

ftt = hardware_error

A *hardware_error* indicates that the FPU detected a catastrophic internal error, such as an illegal state or a parity error on an *f* register access.

FSR_FQ_not_empty (qne)

Bit 13 indicates whether the optional floating-point deferred-trap queue (FQ) is empty after a deferred *fp_exception* trap or after a double floating-point queue (STDFQ) instruction has been executed. If *qne*=0, the queue is empty; if *qne*=1, the queue is not empty.

The *qne* bit can be read by the STFSR instruction.

The LDFSR instruction does not affect *qne*. However, executing successive STDFQ instructions will (eventually) cause the FQ to become empty (*qne*=0). If an implementation does not provide an FQ, this bit reads as zero.

This bit always reads as zero to user mode software.

FSR_fp_condition_codes (fcc)

Bits 11 and 10 contains the FPU condition codes.

These bits are updated by floating-point compare instructions (FCMP and FCMPE). They are read and written by the STFSR and LDFSR instruction, respectively. FBfcc bases its control transfer on this field.

In the following table, *frs1* and *frs2* correspond to the single or double values in the *f* registers specified by an instruction's *rs1* and *rs2* fields. The question mark (?) indicates an unordered relation, which is true if either *frs1* or *frs2* is a signaling NaN or quiet NaN.

The *fcc* is undefined if FCMPE generates an IEEE_exception trap.

Table 28. Floating-point Condition Codes (fcc) Field of FSR

FCC	RELATION
0	<i>frs1</i> = <i>frs2</i>
1	<i>frs1</i> < <i>frs2</i>
2	<i>frs1</i> > <i>frs2</i>
3	<i>frs1</i> ? <i>frs2</i> (unordered)

FSR_accrued_exception (aexc)

Bits 9 through 5 accumulate IEEE_754 floating-point exceptions while fp_exception traps are disabled using the TEM field.

Accrued Exception Bits (aexc) Fields of FSR

nva	ofa	ufa	dza	nxa
9	8	7	6	5

After an FPop completes, the TEM and cexc fields are logically and'd together. If the result is nonzero, an fp_exception trap is generated; otherwise, the new cexc field is or'd into the aexc field. Thus, while traps are masked, exceptions are accumulated in the aexc field.

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.

0	No Corresponding Exception.
1	Corresponding Exception

Current Exception Bits (cexc) Fields of FSR

nvc	ofc	ufc	dzc	nxc
4	3	2	1	0

The cexc bits are unchanged following an FPop that causes an fp_exception trap. Following a non-trapping FPop, the cexc bits are defined as described as followings.

FSR_invalid (nvc, nva)

An operand is improper for the operation to be performed. For example, $0 \div 0$, and $\infty - \infty$ are invalid.

1	invalid operand
0	valid operand

FSR_overflow (ofc, ofa)

The rounded result would be larger in magnitude than the largest normalized number in the specified format.

1	overflow
0	no overflow

PRODUCT PREVIEW

FSR_underflow (ufc, ufa)

The rounded result is inexact and would be smaller in magnitude than the smallest normalized number in the specified format.

1	underflow
0	no underflow

Underflow is never indicated when the correct unrounded result is zero.

Otherwise, if UFM=0: The ufc and ufa bits will be set if the correct unrounded result of an operation is less in magnitude than the smallest normalized number and the correctly-rounded result is inexact. These bits will be set if the correct unrounded result is less than the smallest normalized number, but the correct rounded result is the smallest normalized number. nxc and nxa are always set as well. if UFM=1: An IEEE_754_exception trap will occur if the correct unrounded result of an operation would be smaller than the smallest normalized number. A trap will occur if the correct unrounded result would be smaller than the smallest normalized number, but the correct rounded result would be the smallest normalized number.

FSR_division-by-zero (dzc, dza)

X÷0, where X is subnormal or normalized.

0÷0 does not set the dzc bit.

1	division-by-zero
0	no division-by-zero

FSR_inexact (nxc, nxa)

The rounded result of an operation differs from the infinitely precise correct result.

1	inexact result
0	exact result

Floating-Point Deferred-Trap Queue (FQ)

The TMS390S10 has one-entry floating-point deferred trap queue. Please refer to FPU description on page 16.

MMU Registers

The TMS390S10 MMU contains the following registers: Processor Control Register (CR), Context Register (CXR), Context Table Pointer Register (CTPR), Instruction Translation Buffer Register (ITBR), SBus Slot Configuration Register (SSCR), Synchronous Fault Address Register (SFAR), Synchronous Fault Status Register (SFSR), Asynchronous Fault Status Register (AFSR), Asynchronous Fault Address Register (AFAR), TLB Replacement Control Register (TRCR), I/OMMU Base Address Register (IBAR), I/O MMU Control Register (I/OCR), and Address Flush Register (AFR).

Processor Control Register (CR)

The Processor Control Register (CR) contains MMU control and status information.

Processor Control Register

IMPL	VER	STW	AV	DV	MV	Rsvd	PC	ID	AC	BM	Rsvd	PE	RC	IE	DE	Rsvd	NF	EN
31 28	27 24	23	22	21	20	19 18	17	16	15	14	13	12	11 10	9	8	7 2	1	0

The BM, IE, DE, and EN bits receive both the SBus reset (normal reset) and watchdog resets (BM is set, IE, ED, and EN are reset).

Note that it is highly recommended that STA's to the PCR are immediately followed by a SPARC FLUSH instruction to keep the machine in a very consistent state.

Reserved (Rsvd)

Bits [19:18, 13, 07:02] are unimplemented, should be read and written as zero.

Implementation (IMPL)

Bits 31 through 28 give the implementation number for TMS390S10. The implementation number is 0x4. This field is read only.

Version (VER)

Bits 27 through 24 specify the version number of the TMS390S10 MMU. This field is read as 0x1 (b'0001) and is read only

Software Tablewalk enable (STW)

Bit 23 is used to enable instruction_access_MMU_miss and data_access_MMU_miss traps for instruction and data tablewalking respectively for tablewalk.

Address View (AV)

Bit 22 is used for diagnostic purposes. This is a debug and test feature.

If set, any address from the MMU Physical Address Register (PAR) is displayed on the SBus Address pins (SBADDR[27:00]=mm_pa[27:00]). SBus cannot be used while this bit is set.

Data View (DV)

Bit 21 is used for diagnostic purposes. This is a debug and test feature.

When this bit is set, any data on the internal memory data bus will appear on the external SBus data pins SBDATA[31:00]. During debug this can be monitored while running non-I/O diagnostics. While this bit is set, SBus cannot be used.

PRODUCT PREVIEW

Memory Data View (MV)

Bit 20 is used for diagnostic purposes.

Any data on the internal memory data bus will appear on the external memory data pins.

This is useful for monitoring ASI and control space accesses (from/to both the IU and SBus). While this bit is set, *store* and *load* operation cannot be used.

Parity Control (PC)

Bit 17 defines the mode for parity checking and generating as follows:

Table 29. Parity Control Definition

PC	MEANING
0	Even Parity
1	Odd Parity

ITBR Disable bit (ID)

Bit 16 disables the use of the Instruction Translation Buffer Register.

Alternate Cacheability (AC)

If the AC bit is not set and the MMU is disabled (EN=0), then the instruction cache and data cache are automatically disabled also. If it is desired to enable the caches while the MMU is disabled, this can be accomplished by setting the AC bit. Setting the AC bit allows the enabling of the caches to be controlled by the setting of the IE and DE bits.

Table 30. MMU and Cache control by EN, IE, DE, and AC

EN	AC	IE	DE	MMU_EN	I-CACHE	D-CACHE
1	-	0	0	Yes	No	No
1	-	0	1	Yes	No	Yes
1	-	1	0	Yes	Yes	No
1	-	1	1	Yes	Yes	Yes
0	0	-	-	No	No	No
0	1	0	0	No	No	No
0	1	0	1	No	No	Yes
0	1	1	0	No	Yes	No
0	1	1	1	No	Yes	Yes

- : Don't care

Boot Mode (BM)

Bit 14 is set by both SBus reset and Watchdog reset and must be cleared for normal operation.

PRODUCT PREVIEW

Parity Enable (PE)

Bit 12 determines whether to enable the parity checking or not. If this bit is set to one, word parity checking is enabled.

Table 31. Parity Enable

PE	MEANING
0	Parity checking is enabled
1	Parity checking is not enabled

Refresh Control (RC)

Bits 11 and 10 control the DRAM refresh rate of the system. Normal 50 MHz operation would require a 0x2 value. The RC field is defined as follows:

Table 32. Memory Refresh Control Definition

RC	REFRESH INTERVAL
0 0	Every 128 clocks (operating frequency above 8.6 MHz) This is default after power up.
0 1	No Refresh
1 0	Every 512 clocks (operating frequency above 35 MHz)
1 1	Every 768 clocks (operating frequency above 52 MHz)

Instruction Cache Enable (IE)

Bit 9 controls the instruction cache. This bit is reset by both SBus reset and Watchdog reset.

Table 33. Instruction Cache Enable

IE	MEANING
1	enables instruction cache
0	all references miss the cache

Data Cache Enable (DE)

Bit 8 controls the data cache. This bit is reset by both SBus reset and Watchdog reset.

Table 34. Data Cache Enable

DE	MEANING
1	enables data cache
0	all references miss the cache

No Fault bit (NF)

when the NF is set, the processor does not act on any exceptions caused by supervisory accesses. Instead, occurrence of the exception is recorded in the SFCSR. For normal operation, the NF bit is not set.

MMU Enable (EN)

When bit 0 is set ,the MMU is enabled and translation occurs normally.


When this bit is not set, the physical address is forced to the 31 least significant bits of the virtual address. This bit is reset by both SBus reset and Watchdog reset.

PRODUCT PREVIEW

Context Register (CXR)

The Context Register (CXR) is used as an index into the Context Table. It is defined as follows.

Context Register Fields




The Context Register defines which virtual address space is considered the “current” address space through a 6-bit context number. The TMS390S10 supports 64 contexts.

Subsequent accesses to memory through the MMU are translated for the current address space. This continues until the CXR is changed. The physical address of the root pointer is obtained by taking bits [22:04] from the CTPR to form mm_pa[26:08] and bits [05:00] from the CXR to form mm_pa[07:02]. The bits mm_pa[30:27,01:00] are zero. Bits [31:06] of the CXR are unimplemented, should be read and written as a zero.

Context Table Pointer Register (CTPR)

The Context Table Pointer Register (CTPR) contains the base of the Context Table. It is defined as follows.

Context Table Pointer Register Fields



The Context Table Pointer is 19 bits wide.

The reserved fields are unimplemented, should be written as zero, and are read as a zero.

Synchronous Fault Status Register (SFSR)

The Synchronous Fault Status Register contains the information on exceptions (faults) issued by the MMU during CPU type transactions. There are three types of faults: instruction access faults, data access faults, and translation table access faults. If another instruction access fault occurs before the fault status of a previous instruction access fault has been read by the IU, the latest fault status is written into the SFSR and the OW bit is set. If multiple data access faults occur, only the status of the one taken by the IU is latched into the SFSR (and address in the SFAR). If data fault status overwrites previous instruction fault status, the OW bit is cleared since the fault status is represented correctly. An instruction access fault does not overwrite a data access fault. If a translation table access fault overwrites a previous instruction or data access fault, the OW bit is cleared. An instruction access or data fault does not overwrite a translation table access fault. Reading the SFSR using ASI 0x04 and type 0x03 clears it. Using type 0x13 to read the SFSR does not clear it. Writes to the SFSR using ASI 0x04 and VA[12:08]=0x03 have no effect while writes using VA[12:08]=0x13 update the register. The SFSR is only guaranteed to be valid after an exception is actually signalled. In other words, it may not be valid if there is no exception.

Each fields of SFSR are defined as follows:

Synchronous Fault Status Register

Rsvd	CS	Rsvd	PERR	Rsvd	TO	BE	L	AT	FT	FAV	OW
31 17	16	15	14 13	12	11	10	9 8	7 5	4 2	1	0

Reserved (Rsvd)

Bits [31:17,15,12] are not implemented. Each bit should be written and read as a zero.

Control Space Error (CE)

Bit 16 is asserted on the following conditions:

- Invalid ASI space
- Invalid ASI size
- Invalid VA field in valid ASI space
- Invalid ASI operation
example: a swap instruction to an ASI other than 0x08 — 0x0B, 0x20

The AT field is not valid on Control Space Errors.

Parity Error (PERR)

Bits 14 and 13 are set of external memory bus parity errors on the even and odd words respectively from memory.

SBus Time Out (TO)

Bit 11 is set as a result of an SBus Time Out from a CPU initiated read transaction.

If there is no SBus slave responding with an acknowledge within 256 SBus cycles (12.8 μs), time out will be reported.

SBus Bus Error (BE)

Bit 10 is set as an error indication which returned from an SBus slave on a CPU initiated read transaction.

This may have been either an error acknowledgement or a Late Error.

Level (L)

Bits 9 and 8 are set to the page table level of the entry which caused the fault.

If an error occurs while fetching a page table (either a PTP or PTE) this field records the page table level of the entry. The level field is defined as follows.

Table 35. SFSR Level Field

L	LEVEL
00	Entry in Context Table
01	Entry in Level 1 Page Table
10	Entry in Level 2 Page Table
11	Entry in Level 3 Page Table

Access Type (AT)

Bits 7 through 5 define the type of access which caused the fault.

Loads and Stores to user/supervisor instruction space can be caused by load/store alternate instructions with ASI=0x8—0xB. The AT field is defined as follows.

Table 36. SFSR Access Type Field

AT	ACCESS TYPE
0	Load from User Data Space
1	Load from Supervisor Data Space
2	Load/Execute from User Instruction Space
3	Load/Execute from Supervisor Instruction Space
4	Store to User Data Space
5	Store to Supervisor Data Space
6	Store to User Instruction Space
7	Store to Supervisor Instruction Space

Fault Type (FT)

Bits 4 through 2 define the type of the current fault. The FT field is defined as follows.

Table 37. SFSR Fault Type Field

FT	FAULT TYPE
0	None
1	Invalid Address Error
2	Protection Error
3	Privilege Violation Error
4	Transaction Error
5	Access Bus Error
6 — 7	Reserved

Invalid address errors, protection errors, and privilege violation errors depend on the AT field of the SFSR and the ACC field of the corresponding PTE. The errors are set as follows.

Table 38. Setting of SFSR Fault Type Code

AT	FT Code								
	PTE[V]=0	PTE[V]=1							
		ACC=0	ACC=1	ACC=2	ACC=3	ACC=4	ACC=5	ACC=6	ACC=7
0	1	—	—	—	—	2	—	3	3
1	1	—	—	—	—	2	—	—	—
2	1	2	2	—	—	—	2	3	3
3	1	2	2	—	—	—	2	—	—
4	1	2	—	2	—	2	2	3	3
5	1	2	—	2	—	2	—	2	—
6	1	2	2	2	—	2	2	3	3
7	1	2	2	2	—	2	2	2	—

An invalid address error code (FT=1) is set when an invalid PTE or PTP is found while fetching an entry from the page table for a regular table walk or a probe entire operation.

The protection error code (FT=2) is set if an access is attempted that is inconsistent with the protection attributes of the corresponding PTE.

The privilege error code (FT=3) is set when a user program attempts to access a supervisor only page.

The translation error code (FT=4) is set when a SFSR PE type error occurs while the MMU is fetching an entry from a page table, a PTP is found in a level 3 page table, or a PTE has ET=3. The L field records the table level at which the error occurred. The PE field records the word(s) having a parity error, if any.

The access bus error code (FT=5) is set when the SFSR PE field gets set on a memory operation that was not a table walk, or on a synchronously generated SBus error acknowledge or time out.

Also this code is set when the either on an alternate space access to an unimplemented or reserved ASI, or the memory access is using a size prohibited by the particular type of ASI. If multiple errors occur on a single access the highest priority fault is recorded in the FT field. (see Table 40)

Fault Address Valid (FAV)

Bit 1 is set if the contents of the Synchronous Fault Address Register (SFAR) are valid.

The SFAR is valid for data faults and translation errors.

PRODUCT PREVIEW

Overwrite (OW)

If the SFSR has been written to more than once, bit 0 is set to indicate that previous status has been lost since the last time it was read.

Table 39. Overwrite Operations

PENDING ERROR	NEW ERROR	OW STATUS	ACTION SIGNALLED
Translation Error	Translation Error	Set	Translation Error
Translation Error	Data Access Exception	Unchanged	Data Access Exception
Translation Error	Instruction Access Exception	Unchanged	Instruction Access Exception
Data Access Exception	Translation Error	Clear	Translation Error
Data Access Exception	Data Access Exception	Set	Data Access Exception
Data Access Exception	Instruction Access Exception	Unchanged	Instruction Access Exception
Instruction Access Exception	Translation Error	Clear	Translation Error
Instruction Access Exception	Data Access Exception	Clear	Data Access Exception
Instruction Access Exception	Instruction Access Exception	Set	Instruction Access Exception

If a single access causes multiple errors, the fault type is recognized in the following priority.

Table 40. Priority of Fault on Single Access

PRIORITY	FAULT TYPE
1	Translation Error
2	Invalid Address Error
3	Privilege Violation Error
4	Protection Error


Synchronous Fault Address Register (SFAR)

The Synchronous Fault Address Register (SFAR) records the 32-bit virtual address of any data fault reported in the SFSR.

The SFAR is overwritten according to the same policy as the SFSR on data faults. Reading the SFSR using ASI 0x04 and VA[12:08] 0x04 clears it. Using VA[12:08] 0x14 to read the SFSR does not clear it. Writes to the SFAR using ASI 0x04 and VA[12:08] 0x04 have no effect while writes using VA[12:08] 0x14 update the register.

The SFAR should always be read before the SFSR to insure that a valid address is returned. The structure of this register is as follows.

Synchronous Fault Address Register




TLB Replacement Control Register (TRCR)

The TLB Replacement Control Register (TRCR) contains the TLB Replacement Counter and counter disable bit.

The TRCR can be read and written using alternate load/store (LDA and STA) at ASI 0x04 with VA[12:8]=0x14. It is defined as follows.

TLB Replacement Control Register



Reserved

Bits 31 through 6 are unimplemented and should be written and read as zeros.

TLB Replacement Counter Disable (TCD)

The TLBRC will not increment when this bit is set.

TLB Replacement Counter (TLBRC)

This is a 5 bit modulo 32 counter which is incremented by one during each CPU clock cycle to point to one of the TLB entries unless the TCD bit is set. When a TLB miss occurs, the counter value is used to address the entry to be replaced.

PRODUCT PREVIEW

I/O MMU Registers

There are several I/O MMU related registers.

The registers are the I/O MMU Control Register (I/OCR), the I/O MMU Base Address Register (I/OBAR), the SBus Slot Configuration Registers (SSCR[0:3]), the Asynchronous Fault Address Register (AFAR), the Asynchronous Fault Status Register (AFSR), Memory Fault Address Register (MFAR), Memory Fault Status Register (MFSR), MID register, and the Address Flush Register (AFR).

All of these internal MMU registers can be accessed directly by software using SBus and I/O MMU Control Space access with PA[30:24]=0x10. Also, the Entire TLB can be flushed using a control space access. The SBus and I/O MMU Control Space address map follows.

Table 41. SBus and I/O MMU Control Space

PA <30:00>	DEVICE	R/W
1000 0000	I/O MMU Control Register	R/W
1000 0004	I/O MMU Base Address Register	R/W
1000 0014	Flush All TLB Entries	W
1000 0018	Address Flush Register	W
1000 1000	Asynchronous Fault Status Register	R/W
1000 1004	Asynchronous Fault Address Register	R/W
1000 1010	SBus Slot Configuration Register 0	R/W
1000 1014	SBus Slot Configuration Register 1	R/W
1000 1018	SBus Slot Configuration Register 2	R/W
1000 101C	SBus Slot Configuration Register 3	R/W
1000 1020	Memory Fault Status Register	R/W
1000 1024	Memory Fault Address Register	R/W
1000 2000	MID Register	R/W

I/O MMU Control Register (I/OCR)

The I/O MMU Control Register (I/OCR) contains control and status bits for the I/O MMU.

This register can be accessed using SBus and I/O MMU Control Space (0x1000 0000). The I/OCR is defined as follows:

I/O MMU Control Register

IMPL	VER	Rsvd	RANGE	Rsvd	ME
31 28	27 24	23	5 4 2	1	0

Implementation (IMPL)

Bits 31 through 28 are the implementation number of this I/O MMU. This field is hardwired to 0x4 (b'0100) and is read only.

Version (VER)

Bits 27 through 24 represent the version number of this I/O MMU. This field is hardwired to 0x1 (b'0001) and is read only.

Reserved (Rsvd)

Bits 23 through 5 and bit 1 are not implemented, should be written as zero, and will be read as zero.

RANGE

Bits 4 through 2 define the virtual address range for DVMA.

Specifically the limit is defined to be 16 MB x 2<RANGE>. All VA bits above this limit must be set to one for an address to be valid. For example, if RANGE=2 then 64 MB of virtual address are supported, and valid DVMA virtual addresses range from 0xFC000000 to 0xFFFFFFFF. Any access using a DVMA virtual address that is out of that range will receive an SBus error acknowledge. The only exception involves slots that have Bypass Enabled. The following table shows how the physical address of an I/O MMU page table entry is generated.

Table 42. I/O MMU Page Table Address Generation

RANGE	LIMIT	PHYSICAL ADDRESS[30:00]
0	16 MB	IBAR[26:10], I/OVA[23:12], b'00
1	32 MB	IBAR[26:11], I/OVA[24:12], b'00
2	64 MB	IBAR[26:12], I/OVA[25:12], b'00
3	128 MB	IBAR[26:13], I/OVA[26:12], b'00
4	256 MB	IBAR[26:14], I/OVA[27:12], b'00
5	512 MB	IBAR[26:15], I/OVA[28:12], b'00
6	1 GB	IBAR[26:16], I/OVA[29:12], b'00
7	2 GB	IBAR[26:17], I/OVA[30:12], b'00

I/O MMU Enable (ME)

When bit 0 is set, I/O MMU translation is enabled.

All TLB entries are flushed by writing to control space address PA=0x1000 0014. This address should not be read since the output of the TLB is unknown during a flash clear operation.

I/O MMU Base Address Register (IBAR)

The I/O MMU Base Address Register (IBAR) defines base address of the I/O Reference Table.

This register can be accessed using the SBus and I/O MMU Control Space (0x1000 0004). The IBAR is defined as follows.

I/O MMU Base Address Register

Rsvd	IBA[30:14]	Rsvd
31 27	26 10	9 0

Reserved (Rsvd)

Bits 31 through 27 and 9 through 0 are not implemented. They should be written and read as zero.

I/O MMU Base Address (IBA)


When the I/O MMU is enabled and the access translation misses the TLB, IBA is used as the base address for the (<RANGE/1024>) byte-aligned I/O MMU Reference Table.

PRODUCT PREVIEW

Address Flush Register (AFR)

The I/OPTe entries may be flushed from the TLB by doing writes to the Address Flush Register at PA=0x1000 0018 with the following format.

Address Flush Register



Reserved (Rsvd)

Bits 11 through 0 are not implemented, these bits should be written and read as zero.

Flush Address (FA)


The virtual page address of the I/OPTe entry to be flushed.

A register is not actually implemented to perform this function. Also to flush all I/O MMU entries, TLB entries must be flushed.

Asynchronous Fault Status Register (AFSR)

The Asynchronous Fault Status Register (AFSR) provides information on asynchronous faults during I/O initiated transactions and CPU write operations. This register is accessed using SBus and I/O MMU Control Space (0x10001000). A hardware lock is used to ensure that this register does not change while being read. Reading this register clears it.

Asynchronous Fault Status Register



Reserved (Rsvd)

Bits 23 through 20 are forced to 0x8 (b'1000) and bits 16 through 0 are not implemented, should be written as zero, and read as zero.

Summary Error Bit (ERR)

Bit 31 indicates one or more of LE, TO, or BE is asserted.

Late Error (LE)

Bit 30 indicates that the SBus reported an error after the transaction was done.

Time Out (TO)

Bit 29 reports the occurrence of an SBus write access time out. For the TMS390S10, this is 12.8 μs.

Bus Error (BE)

Bit 28 indicates that an SBus write access received an error.

Size (SIZE)

Bits 27 through 25 represent SBus size of error transaction.

Supervisor (S)

If bit 24 is set, CPU was in Supervisor mode when error occurred.

PRODUCT PREVIEW

I/O MMU Bypass (BY)

When bit 0 is set the MMU is bypassed and the virtual addresses from this slave are treated as physical when SBDATA[31:30]=00. Bit mm_pa[30] is given by the physical SA30 field and bits mm_pa[29:00] is defined as sb_ioa[29:00].

Memory Fault Status Register (MFSR)

The Memory Fault Status Register (MFSR) provides information on parity faults. This register is accessed using SBus and I/O MMU Control Space (0x1000 1020). This register is loaded on every request to memory unless it is locked. A hardware lock is used to ensure that this register does not change while being read if there was an error condition. Reading this register allows it to begin loading once again.

Memory Fault Status Register

ERR	Rsvd	S	CP	Rsvd	ME	Rsvd	PERR	BM	C	Rsvd	Type	Rsvd
31	30 25	24	23	22 20	19	18 15	14 13	12	11	10 8	7 4	3 0

Reserved (Rsvd)

Bits 30 through 25, 22 through 20, 18 through 15, 10 through 8, and 3 through 0 are not implemented, should be written as zero, and will be read as zero.

Summary Error Bit (ERR)

Bit 31 is set if one or more of PERR[1] or PERR[0] is asserted.

Supervisor (S)

Bit 24 is set if CPU was in Supervisor mode when error occurred.

CPU Transaction (CT)

Bit 23 is set if CPU initiated the translation that resulted in the parity error

Multiple Error (ME)

Bit 19 is set if there was at least one other error that occurred.

Parity Error[1:0] (PERR[1:0])

Bits 14 and 13 are set on external memory parity errors for the even and odd words respectively from memory. Parity errors can result from CPU or I/O initiated memory reads and byte or halfword (8- or 16-bit) write operations (which result in read-modify-writes).

Boot Mode (BM)

Bit 12 indicates that the error occurred while in Boot Mode.

Cacheable (C)

Bit 11 is set if the address of the error was mapped cacheable. On CPU initiated transactions this bit was from the C bit of the PTE, otherwise it is set to zero.

PRODUCT PREVIEW

Memory Request Type (Type[3:0])

Bits 7 through 4 record the type of request that generated the parity error as follows:

Table 43. Memory Request Type

VALUE (HEX)	NAME	MEANING
0	NOP	No memory operation
1	RD64	Read of 64 bits (2 words)
2	RD128	Read of 128 bits (4 words)
3		Reserved
4	RD256	Read of 256 bits (8 words)
5		Reserved
6		Reserved
7		Reserved
8		Reserved
9	WR8	Write of 8 bits (1 byte)
A	WR16	Write of 16 bits (2 bytes)
B	WR32	Write of 32 bits (1 word)
C	WR64	Write of 64 bits (2 words)
D		Reserved
E		Reserved
F		Reserved

Memory Fault Address Register (MFAR)

The Memory Fault Address Register (MFAR) records the 31-bit physical address that caused the fault. This register is accessed using the SBus and I/O MMU Control Space (0x1000 1024). This register is loaded on every request to memory unless it is locked. A hardware lock is used to ensure that this register does not change while being read if there is another error condition. Reading this register allows it to begin loading once again. Bit[31] should be written as zero and will be read as zero.

The structure of this register is as follows:

Memory Fault Address Register




Bit 31 is unimplemented, should be written as zero, and will be read as zero.

PRODUCT PREVIEW

MID Register

The MID Register contains two fields. The MID field (Bits[3:0]) contains a constant value of 0x8 and the SBAE field controls whether SBus devices can arbitrate the bus. This register can be accessed using SBus and I/O MMU Control Space (0x1000 2000). The SBAE bits are both readable and writeable while the MID field is read only. The MID is defined as follows:

MID Register**Reserved**

Bits 31 through 21 and 15 through 04 are not implemented and should be read and written as zero.

SBus Arbitration Enable[4:0] (SBAE)

Bits 20 through 16 control the ability for device on the SBus to arbitrate for the bus.

The most significant bit (SBAE[4]) controls arbitration for the preset 5th device.

The other bits (SBAE[3:0]) control arbitration for SBus device 3:0 corresponding to SSCR[3:0]. These bits are readable and writeable.

MID

Field 3 through 0 is constant 0x8 (b'1000) and is read only.

Writes to these bits are ignored.

I/O MMU Bypass Mode

Bypass mode is provided to allow intelligent SBus masters to do their own memory management with assistance from the kernel.

This facility is enabled by having the Bypass Enable bit set in that device's slot configuration register. It is assumed that such a master will have its own MMU. In order to bypass the I/O MMU the DVMA master must issue a virtual address with sb_ioa[31:30]=0. In this case the Physical Address bus will have the Virtual Address bus put on it. The PA is checked to verify that it is in the valid main memory range and an error is issued to the master if it is not.

Instruction Translation Buffer Register (ITBR)

The Instruction Translation Buffer Register (ITBR) records a registered version of the last instruction translated TLB line.

Table 44. TMS390S10 INSTRUCTIONS

INSTRUCTION	DESCRIPTION
LDSB (LDSBA) [†]	Load Signed Byte (from Alternate Space)
LDSH (LDSHA) [†]	Load Signed Halfword (from Alternate Space)
LDUB (LDUBA) [†]	Load Unsigned Byte (from Alternate Space)
LDUH (LDUHA) [†]	Load Unsigned Halfword (from Alternate Space)
LD (LDA) [†]	Load Word (from Alternate Space)
LDD (LDDA) [†]	Load Doubleword (from Alternate Space)
LDF	Load Floating-point
LDDF	Load Double Floating-point
LDFSR	Load Floating-point State Register
LDC [§]	Load Coprocessor
LDDC [§]	Load Double Coprocessor
LDCSR [§]	Load Coprocessor State Register
STB (STBA) [†]	Store Byte (into Alternate Space)
STH (STHA) [†]	Store Halfword (into Alternate Space)
ST (STA) [†]	Store Word (into Alternate Space)
STD (STDA) [†]	Store Doubleword (into Alternate Space)
STF	Store Floating-point
STDF	Store Double Floating-point
STFSR	Store Floating-point State Register
STDFQ [†]	Store Double Floating-point deferred-trap Queue
STC [§]	Store Floating-point
STDC [§]	Store Double Floating-point
STCSR [§]	Store Floating-point State Register
STDCQ ^{§†}	Store Double Floating-point deferred-trap Queue
LDSTUB (LDSTUBA) [†]	Atomic Load-Store Unsigned Byte (in Alternate Space)
SWAP (SWAPA) [†]	Swap r Register with Memory (in Alternate Space)
SETHI	Set High 22 Bits of r Register
NOP	No Operation
AND (ANDcc)	Logical And (and modify icc)
ANDN (ANDNcc)	Logical And Not (and modify icc)
OR (ORcc)	Logical Inclusive-Or (and modify icc)
ORN (ORNcc)	Logical Inclusive-Or Not (and modify icc)

[†] Privileged Instruction causes a privileged instruction trap if attempted with PSR.S clear.

[‡] Privileged instruction if the referenced register is a privileged register.

[§] Cannot have a coprocessor. These instruction always trap with a cp disabled trap.

[¶] The TMS390S10 does not support floating point quad format. These instruction will trap with an fp-exception trap and FSR.FTT of 3.

[#] unimplemented FPop.

PRODUCT PREVIEW

Table 44. TMS390S10 INSTRUCTIONS (Continued)

INSTRUCTION	DESCRIPTION
XOR (XORcc)	Logical Exclusive-Or (and modify icc)
XNOR (XNORcc)	Logical Exclusive-Nor (and modify icc)
SLL	Shift Left Logical
SRL	Shift Right Logical
SRA	Shift Right Arithmetic
ADD (ADDcc)	Add (and modify icc)
ADDX (ADDXcc)	Add with carry (and modify icc)
TADDcc (TADDccTV)	Tagged Add and modify icc (and Trap on overflow)
SUB (SUBcc)	Subtract (and modify icc)
SUBX (SUBXcc)	Subtract with Carry (and modify icc)
TSUBcc (TSUBccTV)	Tagged Subtract and modify icc (with Trap on overflow)
MULScc	Multiply Step and modify icc
UMUL (UMULcc)	Unsigned Integer Multiply (and modify icc)
SMUL (SMULcc)	Signed Integer Multiply (and modify icc)
UDIV (UDIVcc)	Unsigned Integer Divide (and modify icc)
SDIV (SDIVcc)	Signed Integer Divide (and modify icc)
SAVE	Save caller's window
RESTORE	Restore caller's window
Bicc	Branch on integer condition codes
FBfcc	Branch on Floating-point condition codes
CBccc [§]	Branch on coprocessor condition codes
CALL	Call and Link
JMPL	Jump and Link
RETT [†]	Return from Trap
Ticc	Trap on integer condition codes
RDASR [‡]	Read Ancillary State Register
RDY	Read Y Register
RDPSR [†]	Read Processor State Register
RDWIN [†]	Read Window Invalid Mask Register
RDTBR [†]	Read Trap Base Register
WRASR [‡]	Write Ancillary State Register

[†] Privileged Instruction causes a privileged instruction trap if attempted with PSR.S clear.

[‡] Privileged instruction if the referenced register is a privileged register.

[§] Cannot have a coprocessor. These instruction always trap with a cp disabled trap.

[¶] The TMS390S10 does not support floating point quad format. These instruction will trap with an fp-exception trap and FSR.FTT of 3.

[#] unimplemented FPop.

PRODUCT PREVIEW

Table 44. TMS390S10 INSTRUCTIONS (Continued)

INSTRUCTION	DESCRIPTION
WRY	Write Y Register
WRPSR†	Write Processor State Register
WRWIN†	Write Window Invalid Mask Register
WRTBR†	Write Trap Base Register
UNIMP#	Unimplemented
FLUSH	Flush Instruction Memory
FITOs	Convert Integer to Floating-point single
FITOd	Convert Integer to Floating-point double
FITOq‡	Convert Integer to Floating-point quad
FsTOi	Convert Floating-point single to Integer
FdTOi	Convert Floating-point double to Integer
FqTOi‡	Convert Floating-point quad to Integer
FsTOd	Convert Floating-point single to double
FsTOq‡	Convert Floating-point single to quad
FdTOs	Convert Floating-point double to single
FdTOq‡	Convert Floating-point double to quad
FqTOs‡	Convert Floating-point quad to single
FqTOd‡	Convert Floating-point quad to double
FMOVs	Move Floating-point single
FNEGs	Negate Floating-point single
FSQRTs	Floating-point Square root single
FSQRTd	Floating-point Square root double
FSQRTq‡	Floating-point Square root quad
FADDs	Floating-point Add single
FADDd	Floating-point Add double
FADDq‡	Floating-point Add quad
FSUBs	Floating-point Subtract single
FSUBd	Floating-point Subtract double
FSUBq‡	Floating-point Subtract quad
FMULs	Floating-point Multiply single

† Privileged Instruction causes a privileged instruction trap if attempted with PSR.S clear.

‡ Privileged instruction if the referenced register is a privileged register.

§ Cannot have a coprocessor. These instruction always trap with a cp disabled trap.

¶ The TMS390S10 does not support floating point quad format. These instruction will trap with an fp-exception trap and FSR.FTT of 3.

* unimplemented FPop.

PRODUCT PREVIEW

PRODUCT PREVIEW documents contain information on products in the formative or design phase of development. Characteristic data and other specifications are design goals. Texas Instruments reserves the right to change or discontinue these products without notice.



Table 44. TMS390S10 INSTRUCTIONS (Continued)

INSTRUCTION	DESCRIPTION
FMULd	Floating-point Multiply double
FMULq [†]	Floating-point Multiply quad
FDIVs	Floating-point Divide single
FDIVd	Floating-point Divide double
FDIVq [†]	Floating-point Divide quad
FsMULd [#]	Floating-point Multiply single produce double
FdMULq [†]	Floating-point Multiply double produce quad
FCMPs	Floating-point Compare single
FCMPd	Floating-point Compare double
FCMPq [†]	Floating-point Compare quad
FCMPEs	Floating-point Compare no Exceptions single
FCMPEd	Floating-point Compare no Exceptions double
FCMPEq [†]	Floating-point Compare no Exceptions quad

[†] Privileged Instruction causes a privileged instruction trap if attempted with PSR.S clear.

[‡] Privileged instruction if the referenced register is a privileged register.

[§] Cannot have a coprocessor. These instruction always trap with a cp disabled trap.

[¶] The TMS390S10 does not support floating point quad format. These instruction will trap with an fp-exception trap and FSR.FTT of 3.

[#] unimplemented FPop.

PRODUCT PREVIEW

Reset

The TMS390S10 has two Reset operations. One is General Reset (SBus reset) and another is Watchdog Reset.

Watchdog Reset

A Watchdog Reset is performed when the IU enters the error state taking a trap while the PSR ET bit is set to zero. Watchdog Reset only resets the IU, and the BOOT MODE bit in the MMU's Processor Control Register.

General Reset

General Reset is initiated by asserting **RESET**. This can happen on power-up and externally-triggered reset. During Power-on Reset, to get stable CPU state, **RESET** must remain active for at least 16 clock cycles after power-up.

To assure reliable DRAM operation, **RESET** must remain active for at least 200 μ s after power-up.)

After Power-on Reset, RC bits in MMU PCR are set to "00" (set refresh cycle to every 128 clocks).




Figure 21. Reset duration for CPU stabilization

PRODUCT PREVIEW



absolute maximum ratings over operating free-air temperature range (unless otherwise noted)†

Supply voltage range, V_{CC}	-0.5 V to 7 V
Input voltage range, V_I (see Note 1)	-0.5 V to 7 V
Output voltage range, V_O (see Note 1)	-0.5 V to $V_{CC} + 0.5$ V
Input clamp current, I_{IK} ($V_I < 0$ or $V_I > V_{CC}$)	± 20 mA
Output clamp current, I_{OK} ($V_O < 0$ or $V_O > V_{CC}$)	± 50 mA
Current into any output in the low state:	96 mA
Operating temperature range‡:	0°C to 70°C
Storage temperature range	-65°C to 150°C

† Stresses beyond those listed under "absolute maximum ratings" may cause permanent damage to the device. These are stress ratings only and functional operation of the device at these or any other conditions beyond those indicated under "recommended operating conditions" is not implied. Exposure to absolute-maximum-rated conditions for extended periods may affect device reliability.

‡ Devices are tested in order to guarantee operation over the junction temperature range.

recommended operating conditions

PARAMETER	MIN	NOM	MAX	UNIT
V_{CC} Supply Voltage	4.75	5.00	5.25	V
V_{IH} High-level input voltage	2.0			V
V_{IL} Low-level input voltage			0.8	V
I_{OH} High-level output current			1.0	mA
I_{OL} Low-level output current			2.0	mA
T_J Junction temperature §	15		95	°C

§ The maximum allowed operating temperature depends on system heat sinking and air flow velocity. This is not a guaranteed specification and should be determined for each system configuration.

electrical characteristics over recommended operating free-air temperature range (unless otherwise noted)

PARAMETER	TEST CONDITIONS	MIN	TYP	MAX	UNIT
V_{OH} High-level output voltage	$V_{CC}=5V \pm 5\%$	2.4			V
V_{OL} Low-level output voltage				0.4	V
I_{OZ} High-impedance output current §		-10		+10	μA
I_{IL} Low-level input current	Input signals except JTAG			-20	μA
	JTAG signals			-2	μA
I_{IH} High-level input current	Input signals except JTAG			20	μA
	JTAG signals			50	μA
I_{CC} Dynamic supply current			700	1000	mA
I_{CCQ} Quiescent supply current				2	mA
C_i Input capacitance ¶			10		pF
C_o Output capacitance ¶			10		pF

§ Except pins which cannot be hi-impedance: CP_STAT, INT_EVENT, JTDO, MADDR[11:0], MRAS, MCAS, MWE, REF_CLK, SBADDR[27:0], SBAS, SBCLK, SBGR[4:0], and SBSEL[4:0].

¶ This specification is provided as an aid to board design. This specification is not guaranteed during manufacturing testing.

VERSION 8/21/92

timing requirements Memory Interface input setup and hold times

NO.	PARAMETER	DESCRIPTION	SIGNALS	MIN	MAX	UNIT
1	$t_{ac}(MDATA)$	Access time from $MCAS\downarrow$	MDATA[63:0]	30.0		ns
2	$t_{ac}(MPAR)$	Access time from $MCAS\downarrow$	MPAR[1:0]	30.0		ns
3	$t_h(MDATA)$	hold from $MCAS\downarrow$	MDATA[63:0]	0.0		ns
4	$t_h(MPAR)$	hold from $MCAS\downarrow$	MPAR[1:0]	0.0		ns

timing requirements SBus input setup and hold times

NO.	PARAMETER	DESCRIPTION	DESCRIPTION	MIN	MAX	UNIT
5	$t_{su}(SBLERR)$	Setup to $SBCLK\uparrow$	SBLERR	10.5		ns
6	$t_{su}(SBRQ)$	Setup to $SBCLK\uparrow$	SBRQ[4:0]	10.5		ns
7	$t_{su}(SBDATA)$	Setup to $SBCLK\uparrow$	SBDATA[2:0]	10.5		ns
8	$t_{su}(SBSIZE)$	Setup to $SBCLK\uparrow$	SBSIZE[2:0]	10.5		ns
9	$t_{su}(SBREAD)$	Setup to $SBCLK\uparrow$	SBREAD	10.5		ns
10	$t_{su}(SBACK)$	Setup to $SBCLK\uparrow$	SBACK[2:0]	10.5		ns
11	$t_h(SBLERR)$	hold from $SBCLK\uparrow$	SBLERR	2.0		ns
12	$t_h(SBRQ)$	hold from $SBCLK\uparrow$	SBRQ[4:0]	2.0		ns
13	$t_h(SBDATA)$	hold from $SBCLK\uparrow$	SBDATA[2:0]	2.0		ns
14	$t_h(SBSIZE)$	hold from $SBCLK\uparrow$	SBSIZE[2:0]	2.0		ns
15	$t_h(SBREAD)$	hold from $SBCLK\uparrow$	SBREAD	2.0		ns
16	$t_h(SBACK)$	hold from $SBCLK\uparrow$	SBACK[2:0]	2.0		ns

PRODUCT PREVIEW

PRODUCT PREVIEW documents contain information on products in the formative or design phase of development. Characteristic data and other specifications are design goals. Texas Instruments reserves the right to change or discontinue these products without notice.



POST OFFICE BOX 655303 • DALLAS, TEXAS 75265

timing requirements JTAG input setup and hold times

NO.	PARAMETER	DESCRIPTION	SIGNALS	MIN	MAX	UNIT
17	$t_{su}(JTMS)$	Setup to JTCK↑	JTMS	20.0		ns
18	$t_{su}(JTDI)$	Setup to JTCK↑	JTDI	20.0		ns
19	$t_h(JTMS)$	hold from JTCK↑	JTMS	0.0		ns
20	$t_h(JTDI)$	hold from JTCK↑	JTDI	0.0		ns

JTRST is an asynchronous signal.

timing requirements JTAG Clock

NO.	PARAMETER	DESCRIPTION	SIGNALS	MIN	MAX	UNIT
21	$t_{su}(JTCK)$		JTCK	1.0	20.0	MHz

timing requirements Input Clock

NO.	PARAMETER	DESCRIPTION	SIGNALS	MIN	MIN	MAX	UNIT
22	$t_w(ICH)$	Pulse duration of IN_CLK high	IN_CLK		5.0		ns
23	$t_w(ICL)$	Pulse duration of IN_CLK low	IN_CLK		5.0		ns
24	$t_w(ICP)$	Period for IN_CLK	IN_CLK		10.0		ns

PRODUCT PREVIEW

Memory Interface switching characteristics over recommended operating conditions

NO.	PARAMETER	CONDITIONS	MIN	MAX	UNIT
25	$t_{su}(M\text{DATA})$ MDATA[63:0] to MCAS↓	$I_{OL} = \text{Max}$	5.0		ns
26	$t_{su}(M\text{PAR})$ MPAR[1:0] to MCAS↓	$I_{OH} = \text{Max}$	5.0		ns
27	$t_{su}(M\text{ADDR})$ MADDR[11:0] to MCAS↓	$V_{load} = 2.1\text{V}$	12.0		ns
28	$t_p(M\text{RAS})$ REF_CLK↑ to MRAS[3:0]↓	$C_L = 35\text{ pF}$		16.0	ns
29	$t_p(M\text{RAS})$ REF_CLK↑ to MRAS[3:0]↑	Figure 22		4.0	ns
30	$t_p(M\text{CAS})$ REF_CLK↑ to MCAS[1:0]			10.0	ns
31	$t_p(M\text{WE})$ REF_CLK↑ to MWE			10.0	ns
32	$t_{oh}(M\text{DATA})$ REF_CLK↑ to MDATA[63:0]		0.0		ns
33	$t_{oh}(M\text{PAR})$ REF_CLK↑ to MPAR[1:0]		0.0		ns
34	$t_{oh}(M\text{ADDR})$ MCAS↓ to MADDR[11:0]		16.0		ns
35	$t_{oh}(M\text{RAS})$ REF_CLK↑ to MRAS[3:0]↓		0.0		ns
36	$t_{oh}(M\text{RAS})$ REF_CLK↑ to MRAS[3:0]↑		0.0		ns
37	$t_{oh}(M\text{CAS})$ REF_CLK↑ to MCAS[1:0]		0.0		ns
38	$t_{oh}(M\text{WE})$ REF_CLK↑ to MWE		0.0		ns

SBus Interface switching characteristics over recommended operating conditions

NO.	PARAMETER	CONDITIONS	MIN	MAX	UNIT
39	$t_{so}(S\text{B}\text{DATA})$ SBCLK↑ to SBDATA[31:0]	$I_{OL} = \text{Max}$	18.0		ns
40	$t_{so}(S\text{B}\text{SIZE})$ SBCLK↑ to SBSIZE[2:0]	$I_{OH} = \text{Max}$	18.0		ns
41	$t_{so}(S\text{B}\text{READ})$ SBCLK↑ to SBREAD	$V_{load} = 2.1\text{V}$	16.0		ns
42	$t_{so}(S\text{B}\text{ACK})$ SBCLK↑ to SBACK[2:0]	$C_L = 35\text{ pF}$	18.0		ns
43	$t_{so}(S\text{B}\text{SEL})$ SBCLK↑ to SBSEL[4:0]	Figure 22	16.0		ns
44	$t_{so}(S\text{B}\text{ADDR})$ SBCLK↑ to SBADDR[27:0]		16.0		ns
45	$t_{so}(S\text{B}\text{AS})$ SBCLK↑ to SBAS		18.0		ns
46	$t_{so}(S\text{B}\text{GR})$ SBCLK↑ to SBGR[4:0]		16.0		ns
47	$t_p(S\text{B}\text{CLK}\uparrow)$ REF_CLK↑ to SBCLK↑		0	1	ns
48	$t_p(S\text{B}\text{CLK}\downarrow)$ REF_CLK↑ to SBCLK↓		0	1	ns
49	$t_{oh}(S\text{B}\text{DATA})$ SBCLK↑ to SBDATA[31:0]		3.5		ns
50	$t_{oh}(S\text{B}\text{SIZE})$ SBCLK↑ to SBSIZE[2:0]		3.5		ns
51	$t_{oh}(S\text{B}\text{READ})$ SBCLK↑ to SBREAD		3.5		ns
52	$t_{oh}(S\text{B}\text{ACK})$ SBCLK↑ to SBACK[2:0]		3.5		ns
53	$t_{oh}(S\text{B}\text{SEL})$ SBCLK↑ to SBSEL[4:0]		5.0		ns
54	$t_{oh}(S\text{B}\text{ADDR})$ SBCLK↑ to SBADDR[27:0]		5.0		ns
55	$t_{oh}(S\text{B}\text{AS})$ SBCLK↑ to SBAS		3.5		ns
56	$t_{oh}(S\text{B}\text{GR})$ SBCLK↑ to SBGR[4:0]		5.0		ns

These test values are designed to guarantee operation to SBus specification under SBus loading conditions.

PRODUCT PREVIEW

JTAG Interface switching characteristics over recommended operating conditions

NO.	PARAMETER	CONDITIONS	MIN	MAX	UNIT
57	$t_p(\text{JTDO})$	JTCK \uparrow to JTDO		20.0	ns
58	$t_{oh}(\text{JTDO})$	JTCK \uparrow to JTDO	$I_{OL} = \text{Max},$ $I_{OH} = \text{Max}$ $V_{load} = 2.1 \text{ V}$ $C_L = 35 \text{ pF}$ Figure 22		ns

Miscellaneous switching characteristics over recommended operating conditions

NO.	PARAMETER	CONDITIONS	MIN	MAX	UNIT
59	$t_p(\text{CP_STAT})$	REF_CLK \uparrow to CP_STAT †		—	
60	$t_{oh}(\text{CP_STAT})$	REF_CLK \uparrow to CP_STAT †	5		cycle
61	$t_{su}(\text{IRQ})$	Setup to REF_CLK \uparrow †	2		cycle
62	$t_h(\text{IRQ})$	hold from REF_CLK \uparrow †	—		
63	$t_{su}(\text{EXT_EVENT})$	Setup to REF_CLK \uparrow †	9.0		ns
64	$t_h(\text{EXT_EVENT})$	hold from REF_CLK \uparrow †	0.0		ns

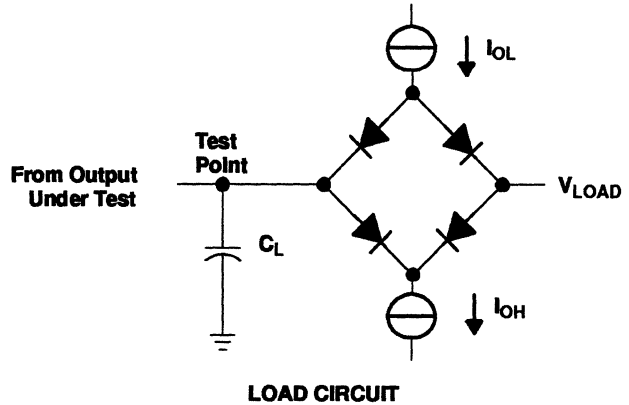
† For these signals, functionality is only tested.

PRODUCT PREVIEW

LOAD CIRCUIT PARAMETERS

TIMING PARAMETERS		C_{LOAD}^\dagger (pF)	I_{OL} (mA)	I_{OH} (μ A)	V_{LOAD} (V)
t_{en}	t_{pZH}	35	2.0	-370	2.1
	t_{pZL}				
t_{dis}	t_{PHZ}	35	2.0	-370	2.1
	t_{PLZ}				
t_{pd}		35	2.0	-370	2.1
t_{pd} MSH		35	6.0	-370	2.1

$\dagger C_{LOAD}$ includes probes and test fixture capacitance.



PRODUCT PREVIEW

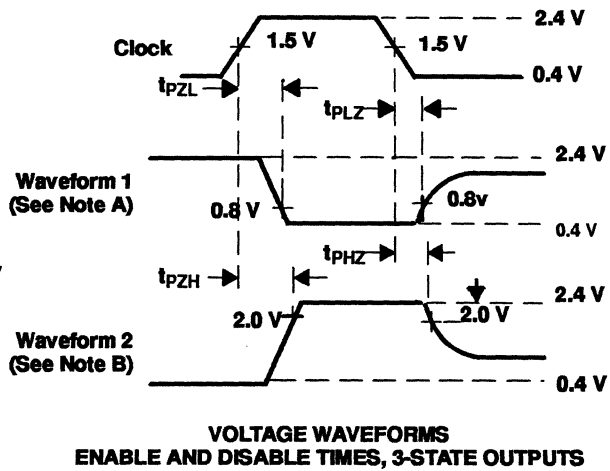
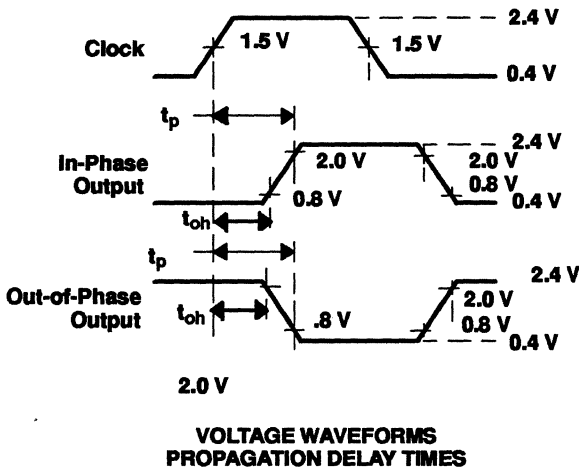


Figure 22.

- Notes: A. Waveform 1 is for an output with internal conditions such that the output is low except when disabled by the output control.
B. Waveform 2 is for an output with internal conditions such that the output is high except when disabled by the output control. For t_{pLZ} and t_{pHZ} , V_{OL} and V_{OH} are specified values.

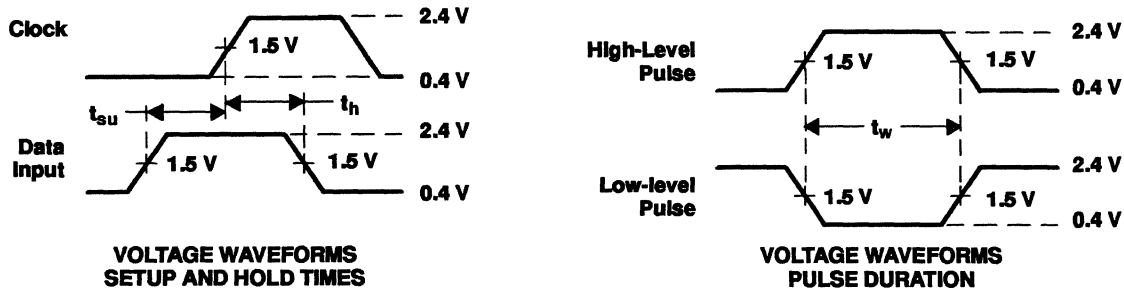


Figure 23.





Figure 24. Air Flow vs Maximum Ambient Temperature (T_A)

PRODUCT PREVIEW

Mechanical Data

PRODUCT PREVIEW



Tape Film Thickness:	125 mm (5mil)
Tape Format:	JEDEC Metric, BD-24
Outer Lead Pitch:	0.25 mm (9.8 mil)
Body Size:	20 mm square
Test PAD Pitch:	0.40 mm
Carrier Type:	JEDEC "Metric TAB Tape Carrier", Super 48mm
Shipping Method:	JEDEC "Metric TAB Magazine" 48mm
Die Metallization:	UP
Tape Metalization:	UP

PRODUCT PREVIEW documents contain information on products in the formative or design phase of development. Characteristic data and other specifications are design goals. Texas Instruments reserves the right to change or discontinue these products without notice.

TEXAS
INSTRUMENTS

