

Preface

Thank you for your new or continued patronage of Toshiba semiconductor products. This is the 1998 edition of the user's manual for the TX49 Family of 64-bit RISC microprocessors, entitled 64-Bit TX System RISC TX49 Family Architecture.

This manual is written so as to be accessible to engineers who may be designing a Toshiba microprocessor into their products for the first time. No prior knowledge of these devices is assumed. The manual includes a review of the architecture of the processor family, a description of the TX49 instruction set, and sections dedicated to various other relevant topics, such as the Memory Management System (MMU) and CPU exceptions.

Toshiba continually updates its technical information. Your comments and suggestions concerning this and other Toshiba documents are sincerely appreciated and may be used in subsequent editions. For updates to this document or for additional information about the product, please contact your nearest Toshiba office or authorized Toshiba dealer.

September 1998

I□TX49 Processor Core Specification

1. Introduction

The TX49 Processor Core is a high performance and low-power 64-bit RISC microprocessor core developed by Toshiba which is well-suited to embedded applications such as networking, laser printer, STB (Set Top Box) and 3-D graphic.

2. Feature

- 64-bit operation
- 32 of 64 bit integer general purpose registers
- 32 of 32-bit floating point general purpose registers: Optional
- 64-GB physical address space
- Instruction Set
 - Upward compatible with MIPS I, MIPS II, and MIPS III ISA
 - MAC (Multiply and Accumulate) instructions %&
 - PREF (Prefetch) instruction %
- Optimized 5 stage pipeline
- Instruction Cache
 - 8-KB/16-KB/32KB : Selectable %
 - Four-way set associative %
 - Lock function support %
- Data cache
 - 8-KB/16-KB/32-KB: Selectable %
 - Four-way set associative %
 - Lock function support %&
 - Write Strategy
 - Write-back-Snoop
 - Write-through-No-Write-Allocate-Snoop %
 - Write-through-Write-Allocate-Snoop %
- MMU
 - 48-double-entry (even/odd) Joint TLB
 - 2-entry Instruction TLB
 - 4-entry Data TLB
- IEEE754 compatible single and double precision FPU: Optional ('99 available)
- Single and double precision FPU in hardware %
- Debug support (EJTAG): Optional %&
- Debug instructions %&
- Real time debugging is supported by debug module logic %&
- Power management modes (halt, doze) %&

3. TX49 Block Diagram

Figure 3-1 shows the block diagram of TX49 Pure Core, MPU Core and MCU. TX49 Pure Core includes an instruction cache and a data cache. These cache are selectable by user system from among a variety of possible configurations.

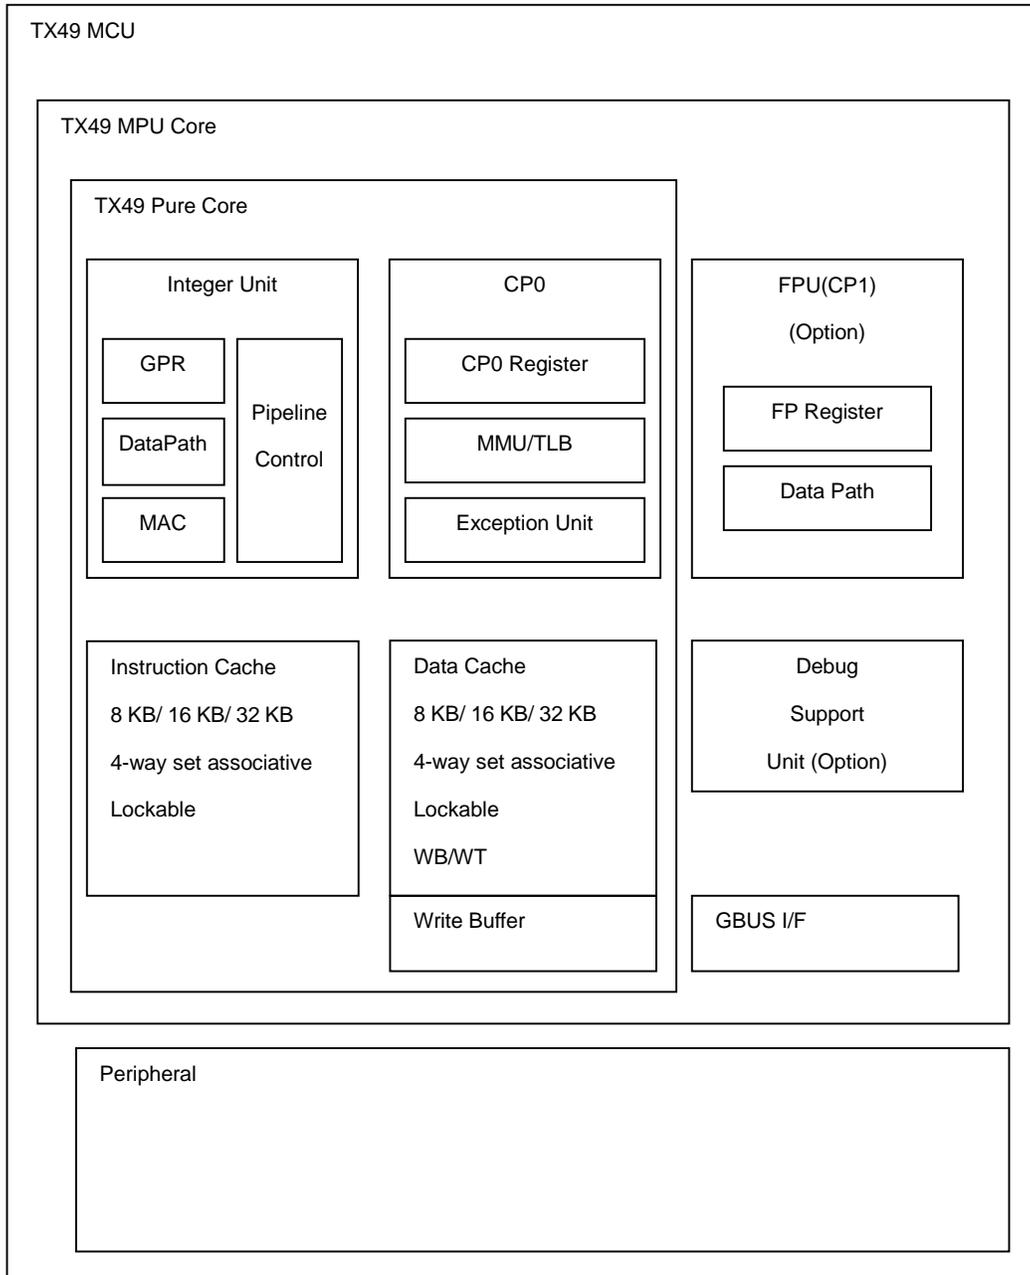


Figure 3-1 Block Diagram of the TX49

4. CPU Registers Overview

4.1 Introduction

The TX49 has the CPU registers for integer operation or address calculation and the CP0 registers for memory system or exception handling.

4.2 CPU Registers

The TX49 has the 64-bit CPU registers.

- 32 general-purpose registers
- 64-bit program counters
- HI/LO register for storing the result of multiply and divide operations

Figure 4-1 shows the configuration of these registers.

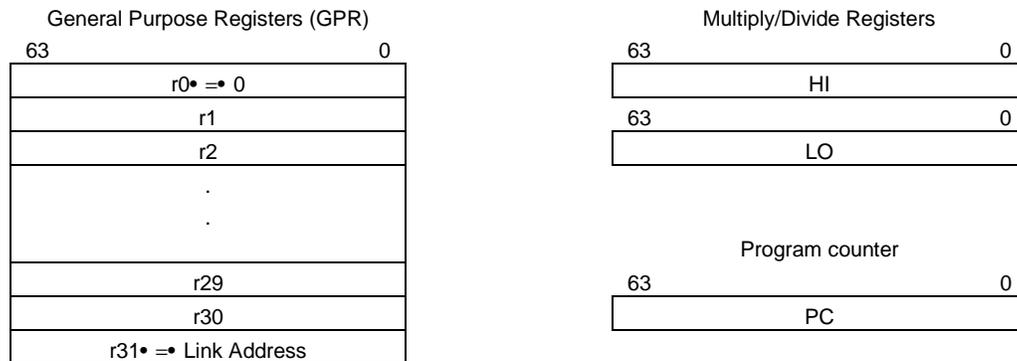


Figure 4-1 TX49 CPU registers

The r0 and r31 registers of GPR have special functions as follows.

- Register r0 always contains the value 0. It can be a target register of an instruction whose operation result is not needed. Or, it can be a source register of an instruction that requires a value of 0.
- Register r31 is the link register for the Jump and Link instruction. The address of the instruction after the delay slot is placed in r31.

The TX49 has the following some special registers that are used or modified implicitly by certain instructions.

- HI - Higher result for Multiply and Divide
- LO - Lower result for Multiply and Divide

These two registers are used to store that result of an integer multiplication or division. In multiplication, the 64 high-order bits of a 128-bit result are stored in the HI, and the 64 low-order bits are stored in the LO. In division, the resulting quotient is stored in the LO, and the remainder is stored in the HI.

- PC - Program Counters

The register contains the address of the currently executed instruction.

4.3 CP0 Registers

The TX49 has the 32-bit or 64-bit System control coprocessor(CP0) registers. These registers are used for memory system or exception handling. Table 4-1 lists the CP0 registers built into the TX49. The more detail information are described in Chapter 7.

Table 4-1 CP0 Registers

Register Name	Reg. No.	Register Name	Reg. No.
Index	Reg#0	Config	Reg#16
Random	Reg#1	LLAddr	Reg#17
EntryLo0	Reg#2	(Reserved)	Reg#18
EntryLo1	Reg#3	(Reserved)	Reg#19
Context	Reg#4	XContext	Reg#20
PageMask	Reg#5	(Reserved)	Reg#21
Wired	Reg#6	(Reserved)	Reg#22
(Reserved)	Reg#7	Debug %&	Reg#23
BadVAddr	Reg#8	DEPC %&	Reg#24
Count	Reg#9	(Reserved)	Reg#25
EntryHi	Reg#10	(Reserved)	Reg#26
Compare	Reg#11	(Reserved)	Reg#27
SR	Reg#12	TagLo	Reg#28
Cause	Reg#13	TagHi	Reg#29
EPC	Reg#14	ErrorEPC	Reg#30
PRId	Reg#15	DESAVE %	Reg#31

5. CPU Instruction Set Summary

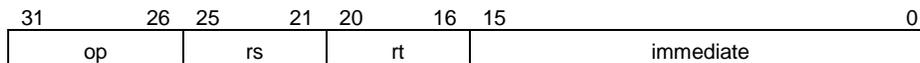
5.1 Introduction

Each instruction is 32 bits long. These instructions are upward compatible with the MIPS I, II and III instruction set architecture and the TX39's instructions.

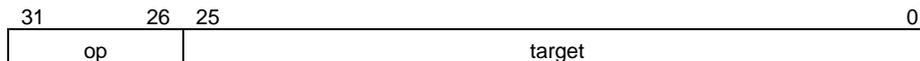
5.2 Instruction Format

There are three instruction formats: Immediate (I-type), Jump (J-type) and Register (R-type), as shown in Figure 5-1. Having just three instruction formats simplifies instruction decoding. If more complex functions or addressing modes are required, they can be produced with the compiler using combinations of the instructions.

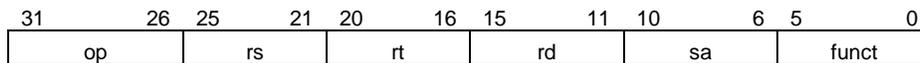
Immediate (I-type)



Jump (J-type)



Register (R-type)



op	Operation code (6• bits)
rs	Source register (5• bits)
rt	Target (source or destination) register, or branch condition (5• bits)
rd	Destination register (5• bits)
immediate	Immediate, branch displacement, address displacement (16• bits)
target	Branch target address (26• bits)
sa	Shift amount (5• bits)
funct	Function (6• bits)

Figure 5-1 Instruction formats and subfield mnemonics

5.3 Instruction Set Overview

5.3.1 Load and Store Instructions (Table 5-1)

Load and Store instructions move data between memory and general purpose registers, and are all I-type instructions. The only directly supported addressing mode is “base register plus 16-bit signed immediate offset”.

Table 5-1 CPU Instruction Set: Load and Store Instructions

Instruction	Description	Note
LB	Load Byte	MIPS I
LBU	Load Byte Unsigned	MIPS I
LH	Load Halfword	MIPS I
LHU	Load Halfword Unsigned	MIPS I
LW	Load Word	MIPS I
LWL	Load Word Left	MIPS I
LWR	Load Word Right	MIPS I
SB	Store Byte	MIPS I
SH	Store Halfword	MIPS I
SW	Store Word	MIPS I
SWL	Store Word Left	MIPS I
SWR	Store Word Right	MIPS I
LD	Load Doubleword	MIPS III
LDL	Load Doubleword Left	MIPS III
LDR	Load Doubleword Right	MIPS III
LL	Load Linked	MIPS II
LLD	Load Linked Doubleword	MIPS III
LWU	Load Word Unsigned	MIPS III
SC	Store Conditional	MIPS II
SCD	Store Conditional Doubleword	MIPS III
SD	Store Doubleword	MIPS III
SDL	Store Doubleword Left	MIPS III
SDR	Store Doubleword Right	MIPS III
SYNC	Sync	MIPS II

5.3.2 Computational Instructions (Table 5-1)

Computational instructions perform arithmetic, logical or shift operations on values in registers. This instruction format can be R-type or I-type. With R-type instructions, the two/three operands and the result are register values. With I-type instructions, one of the operands is 16-bit immediate data. Computational instructions can be classified as follows.

- ALU immediate
- Three-operand register-type
- Shift
- Multiply/Divide

Table 5-1 CPU Instruction Set: Computational Instructions

Instruction	Description	Note
	(ALU Immediate)	
ADDI	Add Immediate	MIPS I
ADDIU	Add Immediate Unsigned	MIPS I
SLTI	Set on Less Than Immediate	MIPS I
SLTIU	Set on Less Than Immediate Unsigned	MIPS I
ANDI	AND Immediate	MIPS I
ORI	OR Immediate	MIPS I
XORI	Exclusive OR Immediate	MIPS I
LUI	Load Upper Immediate	MIPS I
DADDI	Doubleword Add Immediate	MIPS III
DADDIU	Doubleword Add Immediate Unsigned	MIPS III
	(ALU 3-Operand, register type)	
ADD	Add	MIPS I
ADDU	Add Unsigned	MIPS I
SUB	Subtract	MIPS I
SUBU	Subtract Unsigned	MIPS I
SLT	Set on Less Than	MIPS I
SLTU	Set on Less Than Unsigned	MIPS I
AND	AND	MIPS I
OR	OR	MIPS I
XOR	Exclusive OR	MIPS I
NOR	NOR	MIPS I
DADD	Doubleword Add	MIPS III
DADDU	Doubleword Add Unsigned	MIPS III
DSUB	Doubleword Subtract	MIPS III
DSUBU	Doubleword Subtract Unsigned	MIPS III
	(Shift)	
SLL	Shift Left Logical	MIPS I
SRL	Shift Right Logical	MIPS I
SRA	Shift Right Arithmetic	MIPS I
SLLV	Shift Left Logical Variable	MIPS I
SRLV	Shift Right Logical Variable	MIPS I
SRAV	Shift Right Arithmetic Variable	MIPS I
DSLL	Doubleword Shift Left Logical	MIPS III
DSRL	Doubleword Shift Right Logical	MIPS III
DSRA	Doubleword Shift Right Arithmetic	MIPS III
DSLLV	Doubleword Shift Left Logical Variable	MIPS III
DSRLV	Doubleword Shift Right Logical Variable	MIPS III

Instruction	Description	Note
DSRAV	Doubleword Shift Right Arithmetic Variable	MIPS III
DSL32	Doubleword Shift Left Logical +32	MIPS III
DSRL32	Doubleword Shift Right Logical +32	MIPS III
DSRA32	Doubleword Shift Right Arithmetic +32	MIPS III
	(Multiply and Divide)	
MULT	Multiply	MIPS I
MULTU	Multiply Unsigned	MIPS I
DIV	Divide	MIPS I
DIVU	Divide Unsigned	MIPS I
MFHI	Move From HI	MIPS I
MTHI	Move To HI	MIPS I
MFLO	Move From LO	MIPS I
MTLO	Move To LO	MIPS I
DMULT	Doubleword Multiply	MIPS III
DMULTU	Doubleword Multiply Unsigned	MIPS III
DDIV	Doubleword Divide	MIPS III
DDIVU	Doubleword Divide Unsigned	MIPS III

5.3.3 Jump and Branch Instructions (Table 5-1)

Jump/branch instructions change the program flow. A jump/branch instruction will delay the pipeline by one instruction cycle, however, an instruction inserted into the delay slot (immediately following a branch instruction) can be executed while the instruction at the branch target address is being fetched.

Table 5-1 CPU Instruction Set: Jump and Branch Instructions

Instruction	Description	Note
J	Jump	MIPS I
JAL	Jump And Link	MIPS I
JR	Jump Register	MIPS I
JALR	Jump And Link Register	MIPS I
BEQ	Branch on Equal	MIPS I
BNE	Branch on Not Equal	MIPS I
BLEZ	Branch on Less Than or Equal to Zero	MIPS I
BGTZ	Branch on Greater Than Zero	MIPS I
BLTZ	Branch on Less Than Zero	MIPS I
BGEZ	Branch on Greater than or Equal to Zero	MIPS I
BLTZAL	Branch on Less Than Zero And Link	MIPS I
BGEZAL	Branch on Greater than or Equal to Zero And Link	MIPS I
BEQL	Branch on Equal Likely	MIPS II
BNEL	Branch on Not Equal Likely	MIPS II
BLEZL	Branch on Less Than or Equal to Zero Likely	MIPS II
BGTZL	Branch on Greater Than Zero Likely	MIPS II
BLTZL	Branch on Less Than Zero Likely	MIPS II
BGEZL	Branch on Greater Than or Equal to Zero Likely	MIPS II
BLTZALL	Branch on Less Than Zero And Link Likely	MIPS II
BGEZALL	Branch on Greater Than or Equal to Zero And Link Likely	MIPS II

5.3.4 Special Instructions (Table 5-1)

There are special instructions used for software trap. The instruction format is R-type for all two.

Table 5-1 CPU Instruction Set: Special Instructions

Instruction	Description	Note
SYSCALL	System Call	MIPS I
BREAK	Break	MIPS I

5.3.5 Exception Instructions (Table 5-1)

These instructions (R-type or I-type) cause a branch to the general exception handling vector based upon the result of a comparison.

Table 5-1 CPU Instruction Set: Exception Instructions

Instruction	Description	Note
TGE	Trap if Greater Than or Equal	MIPS II
TGEU	Trap if Greater Than or Equal Unsigned	MIPS II
TLT	Trap if Less Than	MIPS II
TLTU	Trap if Less Than Unsigned	MIPS II
TEQ	Trap if Equal	MIPS II
TNE	Trap if Not Equal	MIPS II
TGEI	Trap if Greater Than or Equal Immediate	MIPS II
TGEIU	Trap if Greater Than or Equal Immediate Unsigned	MIPS II
TLTI	Trap if Less Than Immediate	MIPS II
TLTIU	Trap if Less Than Immediate Unsigned	MIPS II
TEQI	Trap if Equal Immediate	MIPS II
TNEI	Trap if Not Equal Immediate	MIPS II

5.3.6 Coprocessor Instructions (Table 5-1)

Coprocessor instructions invoke coprocessor operations. The format of these instructions depends on which coprocessor is used.

Table 5-1 CPU Instruction Set: Coprocessor Instructions

Instruction	Description	Note
LWCz	Load Word to Coprocessor z (z* ⇒ 1,2)	MIPS I
SWCz	Store Word from Coprocessor z (z* ⇒ 1,2)	MIPS I
MTCz	Move To Coprocessor z (z* ⇒ 1,2)	MIPS I
MFCz	Move From Coprocessor z (z* ⇒ 1,2)	MIPS I
CTCz	Move Control To Coprocessor z (z* ⇒ 1,2)	MIPS I
CFCz	Move Control From Coprocessor z (z* ⇒ 1,2)	MIPS I
COPz	Coprocessor Operation z (z* ⇒ 1,2)	MIPS I
BCzT	Branch on Coprocessor z True (z* ⇒ 0,1,2)	MIPS I
BCzF	Branch on Coprocessor z False (z* ⇒ 0,1,2)	MIPS I
BCzTL	Branch on Coprocessor z True Likely (z* ⇒ 0,1,2)	MIPS II
BCzFL	Branch on Coprocessor z False Likely (z* ⇒ 0,1,2)	MIPS II
LDCz	Load Double Coprocessor z (z* ⇒ 1,2)	MIPS III
SDCz	Store Double Coprocessor z (z* ⇒ 1,2)	MIPS III
DMTCz	Doubleword Move To Coprocessor z (z* ⇒ 1,2)	MIPS III
DMFCz	Doubleword Move From Coprocessor z (z* ⇒ 1,2)	MIPS III

5.3.7 CP0 Instructions (Table 5-1)

Coprocessor 0 instructions are used for operations involving the system control coprocessor (CP0) registers, processor memory management and exception handling.

Table 5-1 Instruction Set: CP0 Instructions

Instruction	Description	Note
MTC0	Move To CP0	MIPS I
MFC0	Move From CP0	MIPS I
DMTC0	Doubleword Move To CP0	MIPS III
DMFC0	Doubleword Move From CP0	MIPS III
TLBR	Read Indexed TLB Entry	MIPS I
TLBWI	Write Indexed TLB Entry	MIPS I
TLBWR	Write Random TLB Entry	MIPS I
TLBP	Probe TLB for Matching Entry	MIPS I
CACHE	Cache	MIPS III
ERET	Exception Return	MIPS III
WAIT	Enter power management mode	

5.3.8 Multiply and Divide Instructions (Table 5-1) %&

Table 5-1 Extensions to the ISA: Multiply and Divide Instructions

Instruction	Description	Note
MULT	Multiply (3-operand)	% &
MULTU	Multiply Unsigned (3-operand)	% &
MADD	Multiply and ADD (3-operand)	% &
MADDU	Multiply and ADD Unsigned (3-operand)	% &

5.3.9 Debug Instructions (Table 5-1) %&

Table 5-1 Extensions to the ISA: Debug Instructions

Instruction	Description	Note
CTC0	Move Control To Coprocessor 0	%
CFC0	Move Control From Coprocessor 0	%
SDBBP	Software Debug Breakpoint	% &
DERET	Debug Exception Return	% &

5.3.10 Other Instructions (Table 5-1) %

Table 5-1 Other Instructions

Instruction	Description	Note
PREF	Prefetch	%

5.4 Instruction Execution Cycles

Because the TX49 employs the high-speed Multiply and Add Calculator (MAC), multiply instructions, such as MULT, MULTU, DMULT and DMULTU are executed faster. And, TX49 is improved the execution of divide instructions, too.

Instruction	Latency	Repeat
MULT 2/3 operand	4/4	1/3
MADD 2/3 operand	4/4	1/3
DMULT 2/3 operand	7/7	6/6
DIV	37	36
DDIV	69	68

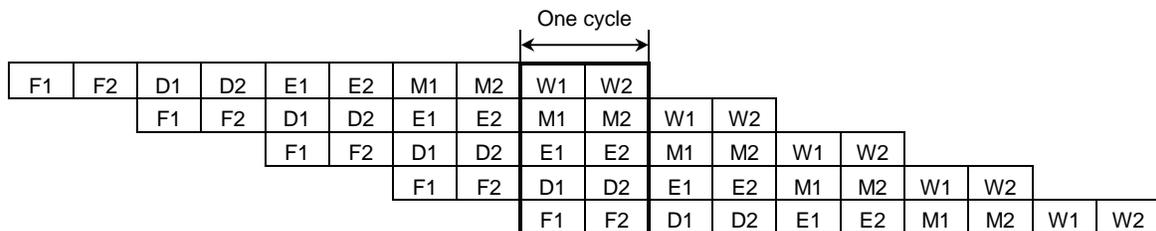
6. CPU Pipeline

6.1 Introduction

This chapter describes the operation of the TX49 pipeline. It explains the basic operation of the pipeline. And, it explains how the TX49 handled delay instructions; these are instructions that follow a branch or load instruction in the pipeline. A later section explains interruptions to the pipeline flow caused by interlocks and exceptions.

6.2 Basic Pipeline Operation

The TX49 executes instructions in an optimized 5 stage pipeline. Each pipeline stage is executed in one clock cycle. When the pipeline is fully utilized, five instructions are executed at the same time, resulting in an average instruction execution rate of one instruction per cycle as illustrated in Figure 6-1.



- F1 - Instruction Fetch, Phase one
- F2 - Instruction Fetch, Phase two
- D1 - Instruction Decode, Phase one
- D2 - Instruction Decode, Phase two
- E1 - Execution, Phase one
- E2 - Execution, Phase two
- M1 - Memory Access, Phase one
- M2 - Memory Access, Phase two
- W1 - Write Back, Phase one
- W2 - Write Back, Phase two

Figure 6-1 Pipeline stages for executing TX49 instructions

F1,F2 : Instruction Fetch

During the F1 phase the ITLB begins the virtual to physical address translation. And, during the F2 phase the instruction cache fetch and the virtual to physical address translation are completed.

D1,D2 : Instruction Decode

The instruction is decoded. Contents of the general-purpose registers are read. If the instruction involves a branch or jump, the target address is generated. The coprocessor condition signal is latched.

E1,E2 : Execution

Arithmetic, logical and shift operations are performed. The execution of multiple/divide instructions is begun.

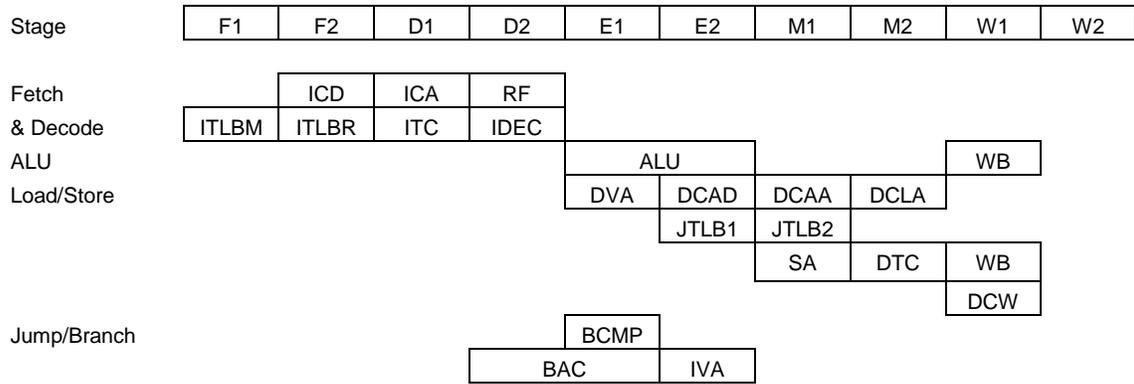
M1,M2 : Memory Access

The data cache is accessed in the case of load and store instructions.

W1,W2 : Write Back

The result is written to a general register.

6.3 TX49 Pipeline Activities



- ICD: Instruction cache address decode
- ICA: Instruction cache array access
- RF: Register fetch
- ITLBM: Instruction address translation match
- ITLBR: Instruction address translation read
- ITC: Instruction tag match
- IDEC: Instruction decode
- ALU: ALU operation
- WB: Write back to register file
- DVA: Data virtual address calculation
- DCAD: Data cache address decode
- DCAA: Data cache array access
- DCLA: Data cache load align
- JTLB1: Address translation in JTLB stage1
- JTLB2: Address translation in JTLB stage2
- SA: Store align
- DTC: Data cache tag check
- DCW: Data cache write
- BCMP: Branch compare
- BAC: Branch address calculation
- IVA: Generate instruction virtual address

6.4 Branch and Load Delay

Some TX49 instructions are executed with a delay of one instruction cycle. The cycle in which an instruction is delayed is called a delay slot. A delay occurs with load instruction and branch/jump instructions.

6.4.1 Delayed load

With load instructions, a one-cycle delay occurs while waiting for the data being loaded to become available for use by another instruction. The TX49 checks the instruction in the delay slot (the instruction immediately following the load instruction) to see if that instruction needs to use the load result; if so, it stalls the pipeline (see Figure 6-1).

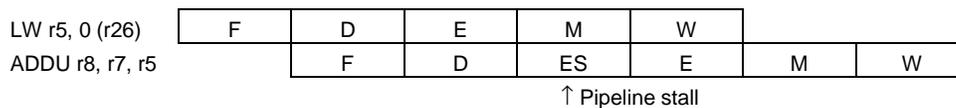


Figure 6-1 CPU Pipeline Load Delay

6.4.2 Delayed branching

Figure 6-1 shows the pipeline flow for jump/branch instructions. The branch target address that must be generated for these type of instructions does not become available until the E stage - too late to be used by the instruction in the branch delay slot. The branch target instruction is fetched immediately after the branch delay slot cycle.

It is, however, possible to fetch a different instruction that would normally be executed prior to the branch instruction.

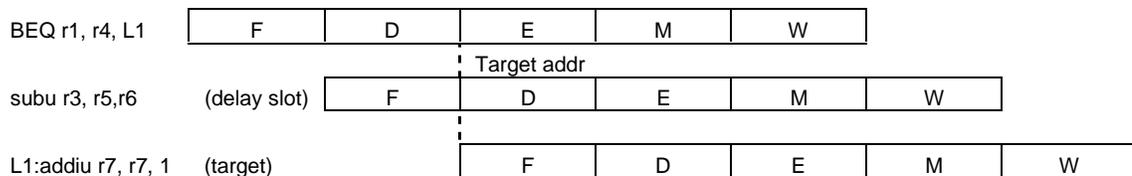


Figure 6-1 CPU Pipeline Branch Delay

You can make effective use of the branch delay slot as follows.

- Since the instruction immediately following a branch instruction will be executed just prior to the branch, you can therefore place an instruction (that logically should be executed just before the branch) into delay slot following the branch instruction.
- The TX49 provides Branch Likely instructions in addition to the normal Branch instructions that allow the instruction at the target branch address to be placed in the delay slot. If the branch condition of the Branch Likely instruction is met, the instruction in the delay slot is executed and the branch is taken. If the branch is not taken, the instruction in the delay is treated as a NOP.
- If no instruction is placed in the delay slot, a NOP is placed just after the branch instruction.

6.5 Non-blocking Load Function

The non-blocking load function prevents the pipeline from stalling when a cache miss occurs and a refill cycle is required to refill the data cache. Instructions after the load instruction that do not use registers affected by the load will continue to be executed. An example is shown in Figure 6-1. Here a cache miss occurs with the first load instruction. The two instructions following are executed prior to the load. The fourth instruction (ADD) must use a register that will be loaded by the load instruction, therefore the pipeline is stalled until the cache data becomes valid.

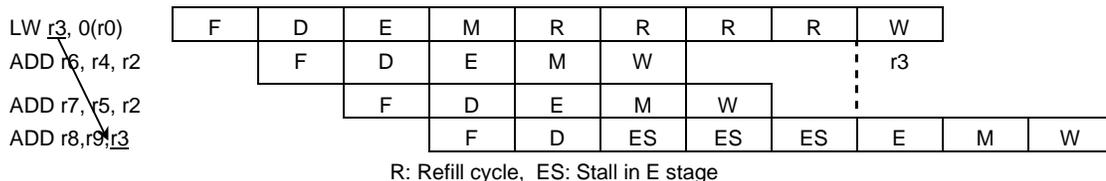


Figure 6-1 Non-blocking load function

6.6 Interlock and Exception Handling

6.6.1 Overview of Interlock and Exception Handling

Smooth pipeline flow is interrupted when cache misses or exceptions occur, or when data dependencies are detected. Interruptions handled using hardware, such as cache misses, are referred to as interlocks, while those that are handled using software are called exceptions.

As shown in Figure 6-1, all interlock and exception conditions are collectively referred to as faults.

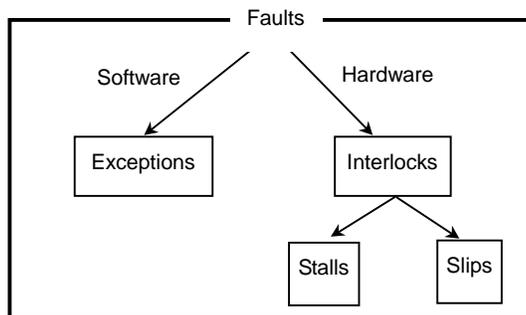


Figure 6-1 Interlocks, Exceptions, and Faults

These are two types of interlocks:

- stalls, which are resolved by halting the pipeline
- slips, which require one part of the pipeline to advance while another part of the pipeline is held static

At each cycle, exception and interlock condition corresponds to a particular pipeline stage, a condition can be traced to the particular instruction in the exception/interlock stage, as shown in Figure 6-2. For instance, an Illegal Instruction (II) exception is raised in the exception (EX) stage.

Table 6-1 and Table 6-2 describe the pipeline interlocks and exceptions listed in Figure 6-2.

State	Pipeline Stage				
	F	D	E	M	W
Stall	ITM	ICM		DCM	
				CPE	
Slip		LDI			
		MDS _t			
		(FCBSy)			
Exception	ITLB	IBE	RI	DBE	
			Cun	NMI	
			BP	Reset	
			SC	OVF	
			DTLB	Trap	
			DTMod		
			Intr		

Figure 6-2 Correspondence of pipeline stage to interlock condition

Table 6-1 Pipeline Interlocks

Interlock	Description
ITM	Instruction TLB Miss
ICM	Instruction Cache Miss
CPE	Coprocessor Possible Exception
DCM	Data Cache Miss
LDI	Load Interlock
MDS _t	Multiply / Divide Start
FCBSy	FP Coprocessor Busy

Table 6-2 Pipeline Exceptions

Exception	Description
ITLB	Instruction Translation or Address Exception
Intr	External Interrupt
IBE	Instruction Bus Error
RI	Reserved Instruction
BP	Breakpoint
SC	System Call
Cun	Coprocessor Unusable
OVF	Integer Overflow
FPE	FP Interrupt
ExTrap	EX Stage Traps
DTLB	Data Translation or Address Exception
TLBMod	TLB Modified
DBE	Data Bus Error
NMI	Nom-maskable Interrupt (or Soft Reset)
Reset	Reset

6.6.2 Exception Conditions

When an exception condition occurs, the relevant instruction and all those that follow it in the pipeline are cancelled. Accordingly, any stall conditions and any later exception conditions that may have referenced this instruction are inhibited; there is no benefit in servicing stalls for a cancelled instruction.

After instruction cancellation, a new instruction stream begins, starting execution at a predefined exception vector. System Control Coprocessor registers are loaded with information that identifies the type of exception and auxiliary information such as the virtual address at which translation exceptions occur.

6.6.3 Stall Conditions

Often, a stall condition is only detected after parts of the pipeline have advanced using incorrect data; this is called a pipeline overrun. When a stall condition is detected, all eight instructions—each different stage of the pipeline—are frozen at once. In this stalled state, no pipeline stages can advance until the interlock condition is resolved.

Once the interlock is removed, the restart sequence begins two cycles before the pipeline resumes execution. The restart sequence reverses the pipeline overrun by inserting the correct information into the pipeline.

6.6.4 External Stalls

External stall is another class of interlocks. An external stall originates outside the processor and is not referenced to a particular pipeline stage. This interlock is not affected by exceptions.

6.6.5 Interlock and Exception Timing

To prevent interlock and exception handling from adversely affecting the processor cycle time, the TX49 processor uses both logic and circuit pipeline techniques to reduce critical timing paths. Interlock and exception handling have the following effects on the pipeline:

- In some cases, the processor pipeline must be backed up (reversed and started over again from a prior stage) to recover from interlocks.
- In some cases, interlocks are serviced for instructions that will be aborted, due to an exception.

6.7 Multiply and Multiply/Add Instructions (MULT, MULTU, MADD, MADDU)

The TX49 can execute 32-bit multiply and multiply/add instructions of 2-operand continuously, and can use the results in the HI/LO registers in immediately following instructions, without pipeline stall as shown Figure 6-1. The TX49 requires three cycles to use the results of a general-purpose register as shown Figure 6-2.

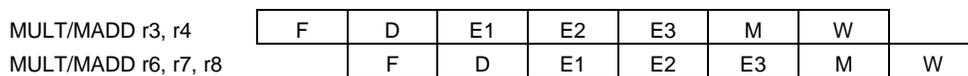


Figure 6-1 MULT and MADD Instructions w/o data dependency (32-bit and 2-operand)

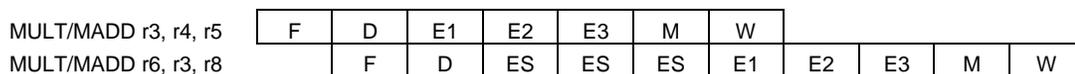


Figure 6-2 MULT and MADD Instructions with data dependency (32-bit and 3-operand)

6.8 Divide Instructions (DIV, DIVU)

Division starts from the pipeline E stage and takes 36 cycles.

Figure 6-1 shows an example of a divide instruction.

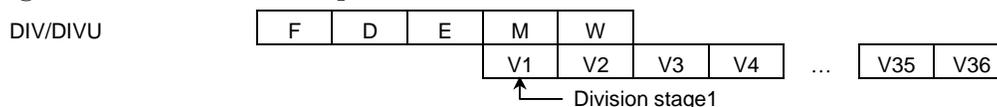


Figure 6-1 DIV and DIVU Instructions

6.9 Streaming

During a cache refill operation, the TX49 can resume execution immediately after arrival of necessary data or instruction in cache even though cache refill is not

completed. This is referred to as “streaming”.

7. System Control Coprocessor, CP0

7.1 Introduction

The TX49 has a System Control Co-Processor (CP0). CP0 translates virtual addresses to physical addresses. CP0 manages exceptions and transitions between kernel, supervisor, and user states. CP0 also controls the cache sub-system, as well as providing diagnostic control and error recovery facilities.

7.2 CP0 Registers

This section is described about the bit field of each register. The term “coldreset” of tables shows the value of each bit when GCOLDRESET* signal is asserted. The reserved bits in description must be written the same value in reset, and return the same value when read.

7.2.1 Index register (Reg#0)

The Index register is a 32-bit read/write register containing six bits to index an entry in the TLB. The P bit of the register shows the success/failure of a TLB Probe (TLBP) instruction.

The Index register also specifies the TLB entry affected by TLB Read (TLBR) or TLB Write Index (TLBWI) instructions.

Figure 7-1 shows the format of the Index register and Table 7-1 describes the Index register fields.



Figure 7-1 Index Register Format

Table 7-1 Index Register Field Descriptions

Bit	Field	Description	coldreset	Read/Write
31	P	Probe failure. Set to 1 when the previous TLB Probe (TLBP) instruction was unsuccessful.	Undefined	Read/Write
30~6	0	Reserved	0x0	Read
5~0	Index	Index to the TLB entry affected by the TLB Read and TLB Write Index instructions.	Undefined	Read/Write

7.2.2 Random register (Reg#1)

The Random register is a read only register containing six bits to index an entry in the TLB. This register decrements as each instruction executes. The values are as follows.

- A lower bound is set by the number of TLB entries reserved for exclusive use by the operating system (the contents of the Wired register).
- An upper bound is set by the total number of TLB entries (47 maximum).

The Random register specifies the TLB entry affected by TLB Write Random (TLBWR) instruction. However the register doesn't need to be read for this purpose, it is readable to verify proper operation of the processor.

To simplify testing, the Random register is set to the value of the upper bound upon system reset. This register is also set to the upper bound when the Wired register is written.

Figure 7-1 shows the format of the Random register and Table 7-1 describes the Random register fields.



Figure 7-1 Random Register Format

Table 7-1 Random Register Field Descriptions

Bit	Field	Description	coldreset	Read/Write
31~6	0	Reserved.	0x0	Read
5~0	Random	TLB random index.	Upper bound (47)	Read

7.2.4 Context register (Reg#4)

The Context register is a read/write register containing the pointer to an entry in the page table entry (PTE) array. This array is an operating system data structure that stores virtual to physical address translations. When there is a TLB miss, the CPU loads the TLB with the missing translation from the PTE array. Normally, the operating system uses the Context register to address the current page map which resides in the kernel mapped segment, kseg3. However the contents of this register duplicates some information of the BadVAddr register, it is arranged in a form that is more useful for TLB exception handler by a software.

Figure 7-1 shows the formats of the Context register and Table 7-1 describes the Context register fields.

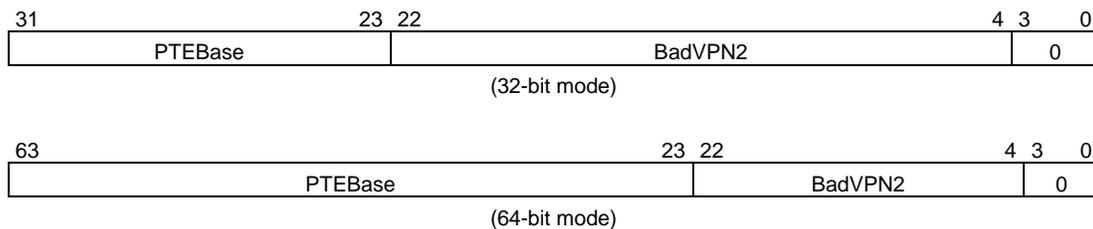


Figure 7-1 Context Register Formats

Table 7-1 Context Register Field Descriptions

32-bit mode

Bit	Field	Description	coldreset	Read/Write
31~23	PTEBase	Page table entry base pointer	Undefined	Read/Write
22~4	BadVPN2	Bad virtual address bits 31~13	Undefined	Read
3~0	0	Reserved	0x0	Read

64-bit mode

Bit	Field	Description	coldreset	Read/Write
63~23	PTEBase	Page table entry base pointer	Undefined	Read/Write
22~4	BadVPN2	Bad virtual address bits 31~13	Undefined	Read
3~0	0	Reserved	0x0	Read

7.2.5 PageMask register (Reg#5)

The PageMask register is a read/write register used for reading from/writing to the TLB. This register holds a comparison mask that sets the variable page size for each TLB entry.

TLB read and write operations use this register as either a source or a destination. When virtual addresses are presented for translation into physical address, the corresponding bits in the TLB identify which virtual address bits among bits 24~13 are used in the comparison. When the Mask field is not one of the values shown in Table 7-1, the operation of the TLB is undefined.

Figure 7-1 shows the format of the PageMask register and Table 7-1 describes the PageMask register fields.

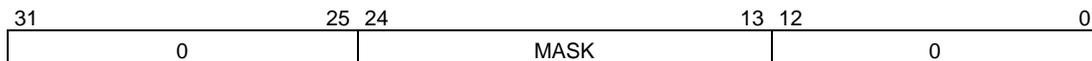


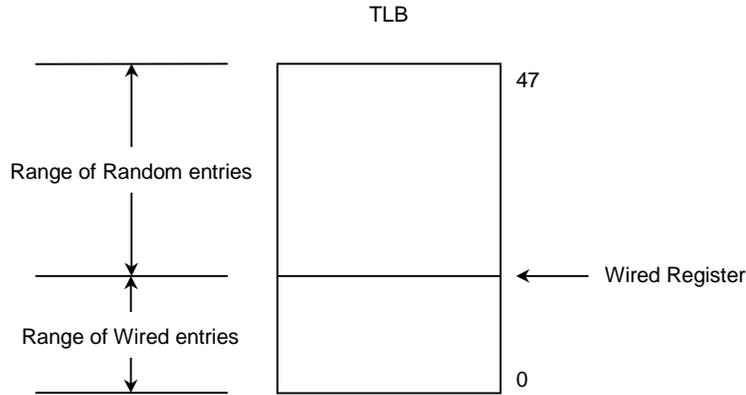
Figure 7-1 PageMask Register Format

Table 7-1 PageMask Register Field Descriptions

Bit	Field	Description	coldreset	Read/Write
31~25	0	Reserved	0x0	Read
24~13	MASK	Page comparison mask 000000000000: page size = 4• Kbytes 000000000011: page size = 16• Kbytes 000000001111: page size = 64• Kbytes 000000111111: page size = 256• Kbytes 000011111111: page size = 1• Mbytes 001111111111: page size = 4• Mbytes 111111111111: page size = 16• Mbytes	0x0	Read/Write
12~0	0	Reserved	0x0	Read

7.2.6 Wired register (Reg#6)

The Wired register is a read/write register specifies the boundary between the wired and random entries of the TLB as follows. Wired entries are non-replaceable entries, which can not be overwritten by a TLB write random operation. Random entries can be overwritten.



The Wired register is set to 0 upon system reset. Writing this register also sets the Random register to the value of its upper bound.

Figure 7-1 shows the format of the Wired register and Table 7-1 describes the Wired register fields.

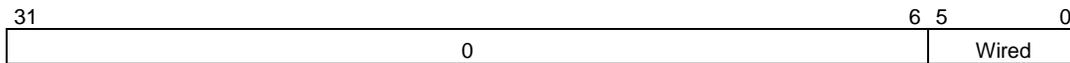


Figure 7-1 Wired Register

Table 7-1 Wired Register Filed Descriptions

Bit	Field	Description	coldreset	Read/Write
31-6	0	Reserved (Must be written as zeroes, and returns zeroes when read.)	0x0	Read
5-0	Wired	TLB Wired boundary.	0x0	Read/Write

7.2.7 BadVAddr register (Reg#8)

The Bad Virtual Address (BadVAddr) register is a read only register that displays the most recent virtual address that cause one of the following exceptions; Address Error, TLB Invalid, TLB Modified and TLB Refill exceptions.

The processor does not write to this register when the EXL bit in the Status register is set to a 1.

Figure 7-1 shows the formats of the BadVAddr register and Table 7-1 describes the BadVAddr register fields.

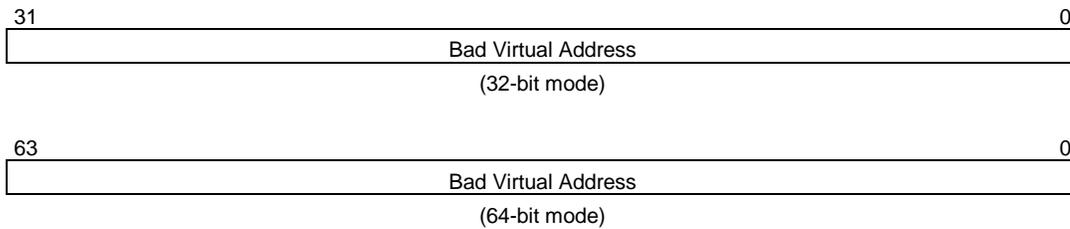


Figure 7-1 BadVAddr Register Formats

Table 7-1 BadVAddr Register Field Descriptions

32-bit mode

Bit	Field	Description	coldreset	Read/Write
31-0	BadVAddr	Bad Virtual address	Undefined	Read

64-bit mode

Bit	Field	description	coldreset	Read/Write
63-0	BadVAddr	Bad Virtual address	Undefined	Read

7.2.8 Count register (Reg#9)

The Count register is a read/write register. This register acts as a timer, incrementing at a constant rate whether or not an instruction is executed, retired, or any forward progress is made through the pipeline.

This register can be also written for diagnostic purpose or system initialization.

Figure 7-1 shows the format of the Count register and Table 7-1 describes the Count register field.

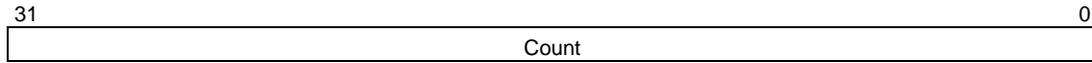


Figure 7-1 Count Register Format

Table 7-1 Count Register Field Description

Bit	Field	Description	coldreset	Read/Write
31-0	Count	32-bit timer, incrementing at half the maximum instruction issue rate.	0x0	Read/Write

7.2.9 EntryHi register (Reg#10)

The EntryHi is a read/write register, and holds the high-order bits of a TLB entry for TLB read and write operations.

This register is accessed by the TLB Probe (TLBP), TLB Write Ransom (TLBWR), TLB Write Indexed (TLBWI), and TLB Read Indexed (TLBR) instructions.

When either a TLB refill, TLB invalid, or TLB modified exception occurs, this register is loaded with the virtual page number (VPN2) and the ASID of the virtual address that did not have a matching TLB entry.

Figure 7-1 shows the formats of the EntryHi register and Table 7-1 describes the EntryHi register fields.

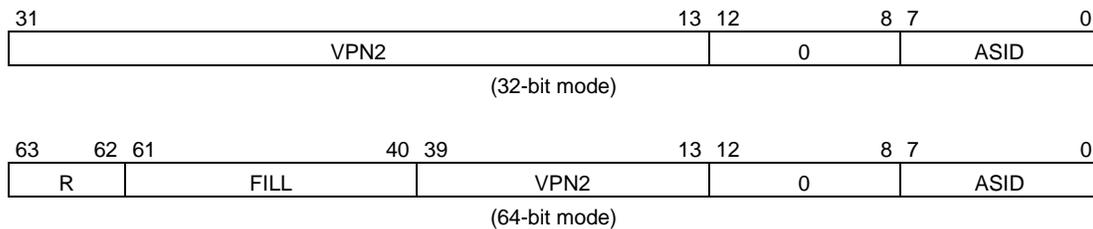


Figure 7-1 EntryHi Register Formats

Table 7-1 EntryHi Register Field Descriptions

32-bit mode

Bit	Field	Description	coldreset	Read/Write
31~1	VPN2	Virtual page number divided by two	Undefined	Read/Write
12~8	0	Reserved	0x0	Read
7~0	ASID	Address space ID field.	Undefined	Read/Write

64-bit mode

Bit	Field	Description	coldreset	Read/Write
63~62	R	Region.	Undefined	Read/Write
61~40	Fill	Reserved. 0 on read. Ignored on write.	Undefined	Read
39~13	VPN2	Virtual page number divided by two	Undefined	Read/Write
12~8	0	Reserved	0x0	Read
7~0	ASID	Address space ID field.	Undefined	Read/Write

7.2.10 Compare register (Reg#11)

The Compare register acts as a timer. When value of the Count register equals the value of the Compare register, interrupt bit IP (7) in the Cause register is set. This causes an interrupt exception as soon as the interrupt is enabled.

Writing a value to this register, as a side effect, clears the timer interrupt.

For diagnostic purpose, this register is a read/write register. However, in normal operation this register is write only.

Figure 7-1 shows the format of the Compare register and Table 7-1 describes the Compare register field.



Figure 7-1 Compare Register Format

Table 7-1 Compare Register Field Description

Bit	Field	Description	coldreset	Read/Write
31~0	Compare	Acts as a timer; it maintains a stable value that does not change on its own.	0x0	Read/Write

7.2.11 Status register (Reg#12)

The Status register is a read/write register that contains the operating mode, interrupt enabling, and diagnostic states of the processor. The more important Status register fields are as following;

- The Interrupt Mask (IM) field of 8 bits controls the enabling of eight interrupt conditions. Interrupt must be enabled before they can be asserted, and the corresponding bits are set in both the IM field of this register and the Interrupt Pending field of the Cause register.
- The Coprocessor Usability (CU) field of 4 bits controls the usability of four possible coprocessors. Regardless of the CU0 bit setting, CP0 is always usable in Kernel mode.
- The Diagnostic Status (DS) field of 9 bits is used for self-testing, and checks the cache and virtual memory system.
- The Reverse Endian (RE) bit reverses the endianness. The processor can be configured as either little/big-endian at reset; reverse-endian selection is used in Kernel and Supervisor modes, and in the User mode when the RE bit is 0. Setting the RE bit to 1 inverts the User mode endianness.

Figure 7-1 shows the format of the Status register and Table 7-1 describes the Status register field.

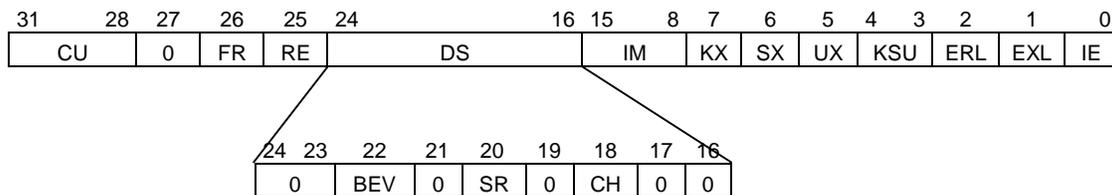


Figure 7-1 Status Register Format

Table 7-1 Status Register Field Descriptions

Bit	Field	Description	coldreset	Read/Write
31~28	CU (3,2,1,0)	Controls the usability of each of the four coprocessor unit numbers. 0: unusable, 1: usable.	0000	Read/Write
27	0	Reserved	0	Read
26	FR	Enables additional floating-point registers. 0: 16 registers, 1: 32 registers.	0	Read/Write
25	RE	Reverse-Endian bit, valid in User mode.	0	Read/Write
24~23	0	Reserved	0x0	Read
22	BEV	Controls the location of TLB refill and general exception vectors. 0: normal, 1: bootstrap	1	Read/Write

Bit	Field	Description	coldreset	Read/Write
21	0	Reserved	0	Read
20	SR	1: Indicates a soft reset or NMI has occurred.	0	Read/Write
19	0	Reserved	0	Read
18	CH	“Hit” or “miss” indication for last CACHE Hit Invalidate, Hit Write Back Invalidate, Hit Write Back for a primary cache. 0: miss, 1: hit.	0	Read/Write
17~16	0	Reserved	0x0	Read
15~8	IM	Interrupt Mask. 0: disabled, 1: enabled.	0x0	Read/Write
7	KX	Enables 64-bit addressing in kernel mode. 0: 32-bit, 1: 64-bit.	0	Read/Write
6	SX	Enables 64-bit addressing and operations in supervisor mode. 0: 32-bit, 1: 64-bit.	0	Read/Write
5	UX	Enables 64-bit addressing and operations in user mode. 0: 32-bit, 1: 64-bit.	0	Read/Write
4~3	KSU	Mode. 10: user, 01: supervisor, 00: kernel.	0x0	Read/Write
2	ERL	Error Level. 0: normal, 1: error.	1	Read/Write
1	EXL	Exception Level. 0: normal, 1: exception.	0	Read/Write
0	IE	Interrupt Enable. 0: disable, 1: enable.	0	Read/Write

7.2.12 Cause register (Reg#13)

The Cause register holds the cause of the most recent exception. This register is read-only, except for the IP[1~0] bits.

Figure 7-1 shows the format of the Cause register and Table 7-1 describes the Cause register field.

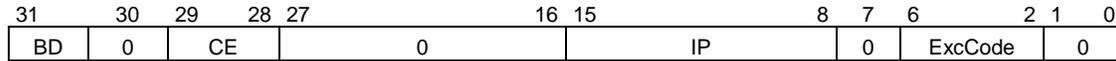


Figure 7-1 Cause Register Format

Table 7-1 Cause Register Field Descriptions

Bit	Field	Description	coldreset	Read/Write
31	BD	Indicates whether or not the last exception was taken while executing in a branch delay slot. 0: normal, 1: delay slot.	0	Read
30	0	Reserved	0	Read
29~28	CE	Indicates the coprocessor unit number referenced when a coprocessor unusable exception is taken. 00: coprocessor 0, 01: coprocessor 1, 10: coprocessor 2, 11: coprocessor 3.	0x0	Read
27~16	0	Reserved	0x0	Read
15~10	IP [7~2]	Indicates whether an interrupt is pending. 0: not pending, 1: pending.	INT[5:0]	Read
9~8	IP [1~0]	Software interrupts. 0: reset, 1: set.	0x0	Read/Write
7	0	Reserved	0	Read
6~2	ExcCode	Exception Code field. 0: Int: Interrupt. 1: Mod: TLB modification exception. 2: TLBL: TLB exception (load or instruction fetch) 3: TLBS: TLB exception (Store) 4: AdEL: Address error exception (load or instruction fetch) 5: AdES: Address error exception (store) 6: IBE: Bus error exception (instruction fetch) 7: DBE: Bus error exception (data reference: load or Store) 8: Sys: Syscall exception 9: Bp: Breakpoint exception 10: RI: Reserved instruction exception 11: CpU: Coprocessor Unusable exception 12: Ov: Arithmetic Overflow exception 13: Tr: Trap exception 14: Reserved: 15: FPE: Floating-Point exception 16-31: Reserved :	0x0	Read
1~0	0	Reserved	0x0	Read

7.2.13 EPC register (Reg#14)

The Exception Program Counter (EPC) register is a read/write register. This register contents the address at which processing resumes after an exception has been serviced.

For synchronous exceptions, this register contains either;

- the virtual address of the instruction that was the direct cause of the exception.
- the virtual address of the immediately preceding branch or jump instruction (when the instruction is in a branch delay slot, and the Branch Delay bit in the Cause register is set).

The processor does not write to the EPC register when EXL bit in the Status register is set to 1.

Figure 7-1 shows the formats of the EPC register and Table 7-1 describes the EPC register field.

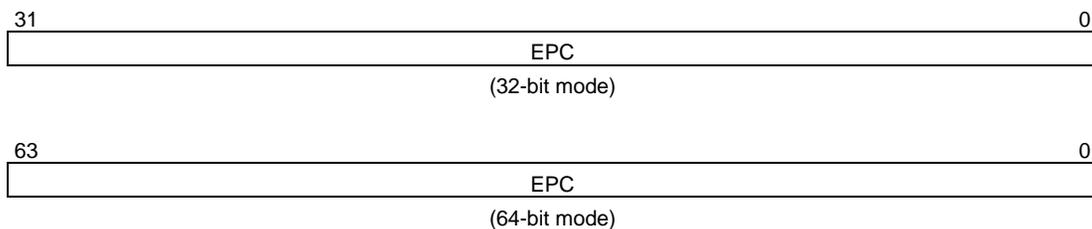


Figure 7-1 EPC Register Formats

Table 7-1 EPC Register Field Description

32-bit mode

Bit	Field	Description	coldreset	Read/Write
31~0	EPC	Exception program counter	Undefined	Read/Write

64-bit mode

Bit	Field	Description	coldreset	Read/Write
63~0	EPC	Exception program counter	Undefined	Read/Write

7.2.14 PRId register (Reg#15)

The Processor Revision Identifier (PRId) register is a read-only register. This register contents information identifying the implementation and revision level of the CPU and CP0.

Figure 7-1 shows the format of the PRId register and Table 7-1 describes the PRId register field.

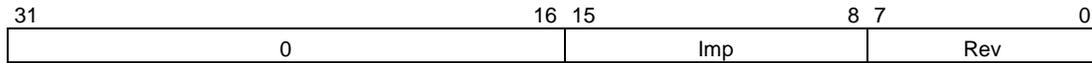


Figure 7-1 PRId Register Format

Table 7-1 PRId Register Field Descriptions

Bit	Field	Description	coldreset	Read/Write
31~16	0	Reserved	0x0	Read
15~8	Imp	Implementation number	0x2d	Read
7~0	Rev	Revision number +.	+	Read

+ Value is shown in product sheet

7.2.15 Config Register (Reg#16)

The Config register is a read-only register; except for HALT, ICE#, DCE# and K0 fields. This register specifies various configuration options selected on the TX49.

EC, BE, IC, DC, IB and DB fields are set by the hardware during reset and are included in this register as read-only status bits for the software to access.

Figure 7-1 shows the format of the Config register and Table 7-1 describes the Config register field.

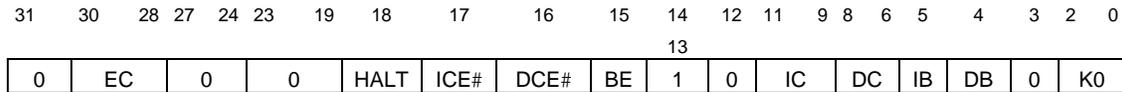


Figure 7-1 Config Register Format

Table 7-1 Config Register Field Descriptions

Bit	Field	Description	coldreset	Read/Write
31	0	Reserved	0	Read
30~28	EC	GBUS clock rate: 0: processor clock frequency divided by 2 1: processor clock frequency divided by 3 2: processor clock frequency divided by 4 7: processor clock frequency divided by 1 % 3, 4, 5, 6 : reserved	pin	Read
27~19	0	Reserved	0x0	Read
18	HALT %&	Wait mode. 0: Halt 1: Doze Indicates the power-down behavior of the TX49 when WAIT instruction is executed. The TX49 stalls the pipeline both in halt and doze mode. Cache snoops are possible during Doze mode but not possible during Halt mode. Halt mode reduces power consumption to a greater extent than Doze mode.	0	Read/Write
17	ICE# %&	Instruction Cache Enable 0: Instruction cache enable 1: Instruction cache disable	0	Read/Write
16	DCE# %&	Data Cache Enable 0: Data cache enable 1: Data cache disable	0	Read/Write
15	BE	Big Endian 0: Little Endian 1: Big Endian	pin	Read
14~13	1	Reserved	11	Read
12	0	Reserved	0	Read

Bit	Field	Description	coldreset	Read/Write
11~9	IC	Instruction cache size. In the TX49, this is set to 8• KB (001), 16• KB (010) or 32• KB (011).	001, 010 or 011	Read
8~6	DC	Data cache size. In the TX49, this is set to 8• KB (001), 16• KB (010) or 32• KB (011).	001, 010 or 011	Read
5	IB	Primary I-Cache line Size 1:32• bytes (8 words)	1	Read
4	DB	Primary D-cache line Size 1:32• bytes (8 words)	1	Read
3	0	Reserved	0	Read
2~0	K0	kseg0 coherency algorithm 0: Cacheable, noncoherent, write-through, no-WA 1: Cacheable, noncoherent, write-through, WA 2: Uncached 3: Cacheable, noncoherent, write-back, WA 4-7: Reserved	0x0	Read/Write

7.2.16 LLAddr register (Reg#17)

The Load Linked Address (LLAddr) register is a read/write register, and contains the physical address read by the most recent Load Linked (LL/LLD) instruction.

This register is for diagnostic purposes only, and serves no function during normal operation.

Figure 7-1 shows the format of the LLAddr register and Table 7-1 describes the LLAddr register field.

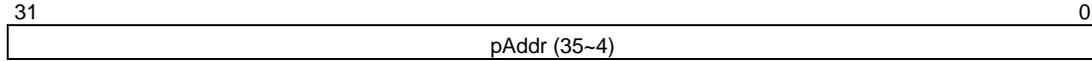


Figure 7-1 LLAddr Register Format

Table 7-1 LLAddr Register Field Description

Bit	Field	Description	coldreset	Read/Write
31~0	pAddr	Physical address bits 35~4	0x0	Read/Write

7.2.17 XContext register (Reg#20)

The XContext register is a read/write register, and contains a pointer to an entry in the page table entry (PTE) array, an operating system data structure that stores virtual to physical address translations. When there is a TLB miss, the operating system software loads the TLB with the missing translation from the PTE array. However the contents of this register duplicates some information of the BadVAddr register, it is arranged in a form that is more useful for TLB exception handler by a software. This register is for use with the XTLB refill handler, which loads TLB entries for references to a 64-bit address space, and is included solely for operating system use. The operating system sets the PTE base field in the register, as needed. Normally, the operating system uses this register to address the current page map which resides in the Kernel mapped segment, kseg3.

The BadVPN2 field of 27 bits has bit [39~13] of the virtual address that caused the TLB miss; bit 12 is excluded because a single TLB entry maps to an even-odd page pair. For a 4 KByte page size, this format may be used directly to access the pair-table of 8 Byte PTEs. For other page sizes and PTE sizes, shifting and masking this value produces the appropriate address.

Figure 7-1 shows the format of the XContext register and Table 7-1 describes the XContext register field.

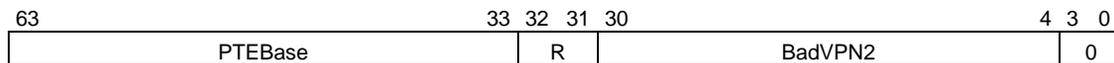


Figure 7-1 XContext Register Format

Table 7-1 XContext Register Field Description

Bit	Field	Description	coldreset	Read/Write
63~33	PTEBase	Page table entry base pointer	Undefined	Read/Write
32~31	R	00: user, 01: supervisor, 11: kernel.	Undefined	Read/Write
30~4	BadVPN2	Bad virtual page number divided by two.	Undefined	Read
3~0	0	Reserved	0x0	Read

7.2.18 Debug register (Reg#23)

The Debug register is a read-only; except for TLF, BsF, SSt and JtagRst fields.

This register holds the information for debug handler.

Figure 7-1 shows the format of the Debug register and Table 7-1 describes the Debug register field.

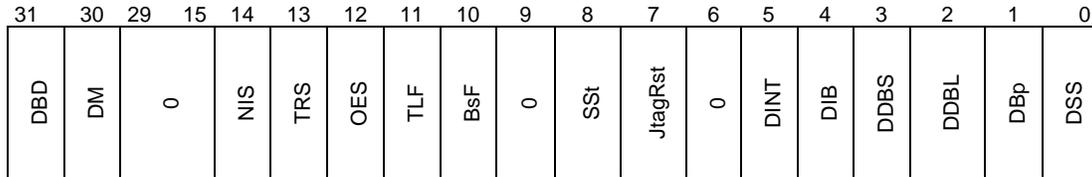


Figure 7-1 Debug Register Format

Table 7-1 Debug Register Field Descriptions

Bit	Field	Description	coldreset	Read/Write
31	DBD	Debug Branch Delay; When a debug exception occurs while an instruction in the branch delay slot is executing, this bit is set to 1.	0	Read
30	DM	Debug Mode; It indicates that a debug exception has taken place. This bit is set when a debug exception is taken, and is cleared upon return from the exception (DERET). While this bit is set all interrupts, including NMI, TLB exception, BUS error exception, and debug exception are masked and cache line locking function is disabled. 0: Debug handler not running. 1: Debug handler running.	0	Read
29~15	0	Reserved	0x0	Read
14	NIS	Non-maskable Interrupt Status; When this bit is set indicating that a non-maskable interrupt has occurred at the same time as a debug exception. In this case the Status, Cause, EPC, and BadVAddr registers assumes the usual status after occurrence of a non-maskable interrupt, but the address in DEPC is not the non-maskable exception vector address (0xbfc0 0000). Instead, 0xbfc0 0000 is put in DEPC by the debug handler software after which processing returns directly from the debug exception to the non-maskable interrupt handler.	0	Read
13	TRS	TLB Miss Status; When this bit is set indicating the Debug Exception and TLB/XTLB refill exception has occurred at the same time. In this case the Status, Cause, EPC, and BadVAddr registers assumes the usual status after occurrence of TLB/XTLB refill. The address in the DEPC is not the other exception vector address. Instead, 0xbfc0 0200 (if BEV• ⇒ 1) in case of TLB refill exception and 0xbfc0 0280 (if BEV• ⇒ 1) in case of XTLB refill exception or 0x8000 0000 (if BEV• ⇒ 0) in case of TLB refill exception and 0x8000 0080 (if BEV• ⇒ 0) in case of XTLB refill exception is put in DEPC by the debug exception handler software, after which processing returns directly from the debug exception to the other exception handler.	0	Read

Bit	Field	Description	coldreset	Read/Write
12	OES	Other Exception Status; When this bit is set indicates exception other than reset, NMI, or TLB/XTLB refill has occurred at the same time as a debug exception. In this case the Status, Cause, EPC, and BadVAddr registers assume the usual status after occurrence of such an exception, but the addressing the DEPC is not the other exception Vector address. Instead, 0xbfc0 0380 (if BEV• =• 1) or 0x8000 0180 (if BEV• =• 0) is put in DEPC by the debug exception handler software, after which processing returns directly from the other exception handler.	0	Read
11	TLF	TLB Exception Flag; This bit is set to 1 when TLB related exception occurs for immediately preceding load or store instruction while a debug exception handler is running (DM• =• 1). TLB exception will set this bit to 1 regardless of writing zero. It is cleared by writing 0 and writing 1 is ignored.	0	Read/Write
10	BsF	Bus Error Exception Flag; This bit is set to 1 when a bus error exception occurs for a load or store instruction while a debug exception handler is running (DM• =• 1). Bus error exception will set this bit to 1 regardless of writing zero. It is cleared by writing 0 and writing 1 is ignored.	0	Read/Write
9	0	Reserved	0	Read
8	SSt	Single Step; Set to 1 indicates the single step debug function is enable or disabled (0). The function is disable when the DM bit is set to 1 while the debug exception is running.	0	Read/Write
7	JtagRst	JTAG Reset; When this bit is set to 1 the processor reset the JTAG unit.	0	Read/Write
6	0	Reserved	0	Read
5	DINT	Debug Interrupt Break Exception Status; set to 1 when debug interrupts occurs.	0	Read
4	DIB	Debug Instruction Break Exception Status; Set to 1 on instruction address break.	0	Read
3	DDBS	Debug Data Break Store Exception Status; Set to 1 on data address break at store operation.	0	Read
2	DDBL	Debug Data Break Load Exception Status; Set to 1 on data address break at load operation.	0	Read
1	DBp	Debug Breakpoint Exception Status; This bit is set when executing SDBBP instruction.	0	Read
0	DSS	Debug Single Step Exception Status; Set to 1 indicate Single Step Exception.	0	Read

7.2.19 DEPC register (Reg#24)

The DEPC register holds the address where processing resumes after the debug exception routine has finished. The address that has been loaded in the DEPC register is the virtual address of the instruction that caused the debug exception. If the instruction is in the branch delay slot, the virtual address of the immediately preceding branch or jump instruction is placed in this register. Execution of the DERET instruction causes a jump to the address in the DEPC. If the DEPC is both written from software (by MTC0) and by hardware (debug exception) then the DEPC is loaded by the value generated by the hardware.

Figure 7-1 shows the formats of the DEPC register and Table 7-1 describes the DEPC register field.

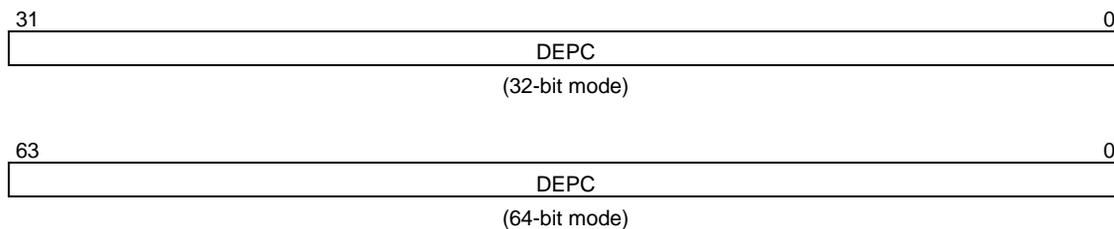


Figure 7-1 DEPC Register Formats

Table 7-1 DEPC Register Field Description

32-bit mode

Bit	Field	Description	coldreset	Read/Write
31-0	DEPC	Debug exception program counter.	Undefined	Read/Write

64-bit mode

Bit	Field	Description	coldreset	Read/Write
63-0	DEPC	Debug exception program counter.	Undefined	Read/Write

7.2.20 TagLo register (Reg#28) and TagHi register (Reg#29)

The TagLo and TagHi registers are a read/write registers. These registers hold the primary cache tag for cache lock function or cache diagnostics. These registers are written by the CACHE/MTC0 instruction.

Figure 7-1 shows the formats of the TagLo and TagHi registers and Table 7-1 describes the TagLo and TagHi registers field.

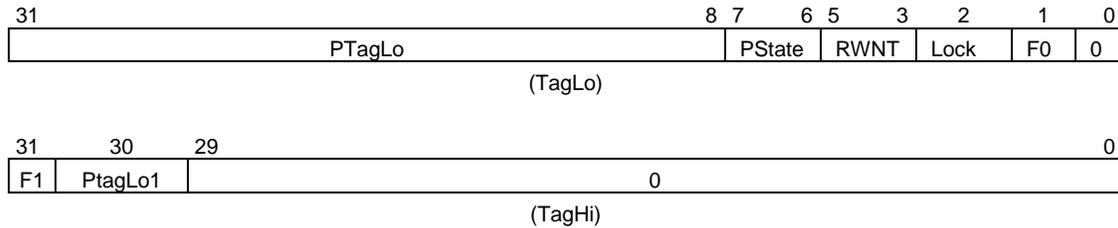


Figure 7-1 TagLo and TagHi Register Formats

Table 7-1 TagLo and TagHi Register Field Descriptions

TagLo

Bit	Field	Description	coldreset	Read/Write
31~8	PTagLo	Bits 35~12 of the physical address	0x0	Read/Write
7~6	PState	Specifies the primary cache state 0: Invalid 1: Reserved 2: Reserved 3: Valid	0x0	Read/Write
5~3	RWNT	Read/Write bits required for Windows NT	0x0	Read/Write
2	Lock	Lock bit (0: not locked, 1: locked)	0	Read/Write
1	F0	FIFO Replace bit 0 (indicates the set to be replaced)	0	Read/Write
0	0	Reserved	0	Read

TagHi

Bit	Field	Description	coldreset	Read/Write
31	F1	FIFO Replace bit 1 (indicates the set to be replaced)	0	Read/Write
30	PtagLo1	Bit 11 of the physical address	0	Read/Write
29~0	0	Reserved	0x0	Read

F1 and F0 are concatenated and indicate the set to be replaced.

F1 || F0

0 0 : way0

0 1 : way1

1 0 : way2

1 1 : way3

7.2.21 ErrorEPC register (Reg#30)

The ErrorEPC is a read/write register, and is similar to the EPC register. This register is used to store the program counter (PC) on ColdReset, SoftReset and NMI exceptions.

This register contains the virtual address at which instruction processing can resume after servicing an error. This address can be;

- The virtual address of the instruction that caused the exception
- The virtual address of the immediately preceding branch or jump instruction, when this address is in a branch delay slot.

There is no branch delay slot indication for this register.

Figure 7-1 shows the formats of the ErrorEPC register and Table 7-1 describes the ErrorEPC register field.

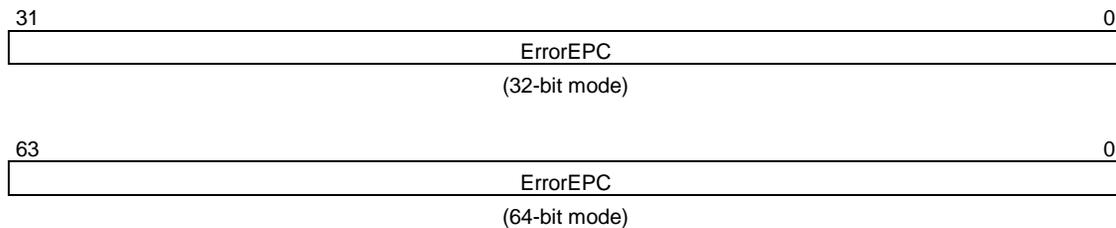


Figure 7-1 ErrorEPC Register Formats

Table 7-1 ErrorEPC Register Field Descriptions

32-bit mode

Bit	Field	Description	coldreset	Read/Write
31~0	ErrorEPC	Error Exception Program Counter.	Undefined	Read/Write

64-bit mode

Bit	Field	Description	coldreset	Read/Write
63~0	ErrorEPC	Error Exception Program Counter.	Undefined	Read/Write

7.2.22 DESAVE register (Reg#31)

This register is used by the debug exception handler to save one of the GPRs, that is then used to save the rest of the context to a pre-determined memory are, e.g. in the processor probe. This register allows the safe debugging of exception handlers and other types of code where the existence of a valid stack for context saving cannot be assumed.

Figure 7-1 shows the formats of the DESAVE register and Table 7-1 describes the DESAVE register field.

Note: This register can use for ICE system only.



Figure 7-1 DESAVE Register Format

Table 7-1 DESAVE register Field Description

32-bit mode

Bit	Field	Description	coldreset	Read/Write
63~0	DESAVE	Save one of the GPRs	Undefined	Read/Write

7.2.23 The initialization of CP0 registers in SoftReset exception

Table 7-1 shows the values of the registers that be initialized by SoftReset exception.

Table 7-1 The initial value by SoftReset Exception

Register	Bit	Field	SoftRest	Description
Status (Reg#12)	22	BEV	1	Same value as ColdReset
	20	SR	1	ColdReset has priority over SoftReset
	2	ERL	1	Same value as ColdReset

8. Memory Management System

8.1 Introduction

The TX49 provides a full-featured memory management unit (MMU) which uses an on-chip translation look aside buffer (TLB) to translate virtual addresses into physical addresses.

8.2 Address space overview

The TX49 physical address space is 64 Gbyte using a 36-bit address. The virtual address is either 64 or 32 bits wide depending on whether the processor is operating in 64- or 32-bit mode. In 32-bit mode, addresses are 32-bits wide and the maximum user process size is 2 Gbyte (2^{31}). In 64-bit mode, addresses are 64-bit wide and the maximum user process is 1 Tbyte (2^{40}). The virtual address is extended with an Address Space Identifier (ASID) to reduce the frequency of TLB flushing when switching context. The size of the ASID field is 8 bits. The ASID is contained in the CPO EntryHi register.

8.2.1 Virtual Address Space

The processor virtual address can be either 32 or 64 bits wide, depending on whether the processor is operating in 32-bit or 64-bit mode.

- In 32-bit mode, addresses are 32 bits wide.
The maximum user process size is 2 gigabytes (2^{31}).
- In 64-bit mode, addresses are 64 bits wide.
The maximum user process size is 1 terabyte (2^{40}).

Figure 8-1 shows the translation of a virtual address into a physical address.

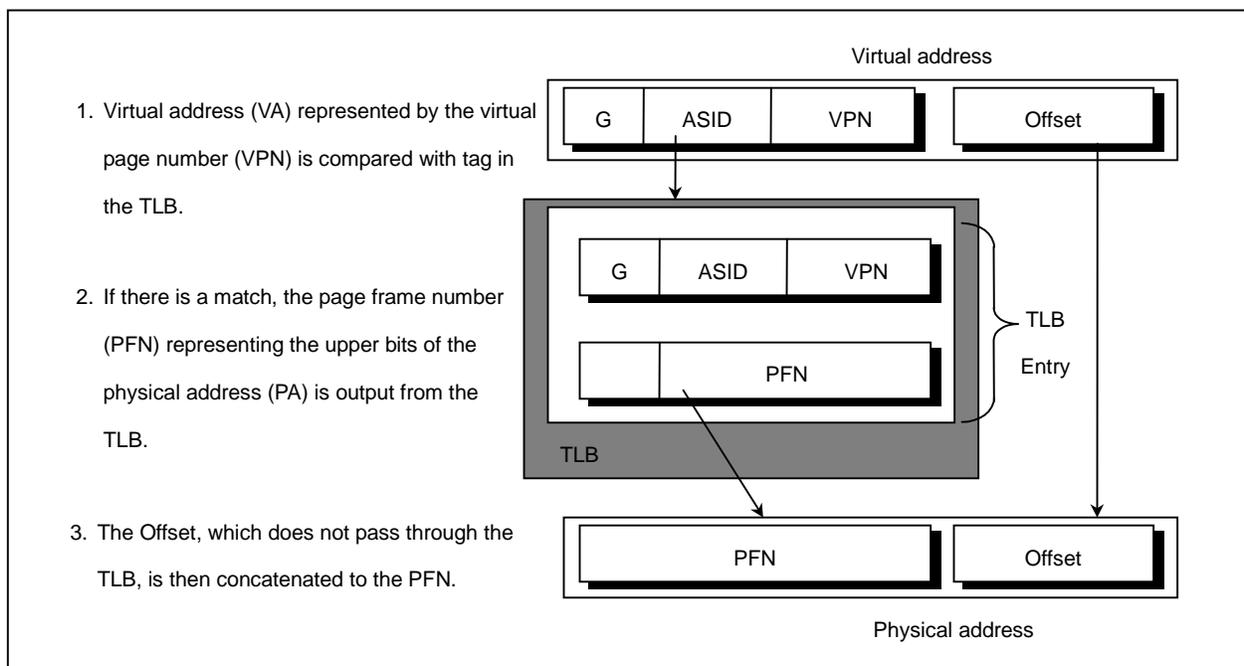


Figure 8-1 Overview of a Virtual-to-Physical Address Translation

As shown in Figure 8-1 and Figure 8-1, the virtual address is extended with an 8-bit address space identifier (ASID), which reduces the frequency of TLB flushing when switching contexts. This 8-bit ASID is in the CPO *EntryHi* register, described later in this chapter. The *Global* bit (G) is in the *EntryLo0* and *EntryLo1* registers, described later in this chapter.

8.2.2 Physical Address Space

Using a 36-bit address, the processor physical address space encompasses 64 Gbytes. The section following describes the translation of a virtual address to a physical address.

8.2.3 Virtual-to-Physical Address Translation

Converting a virtual address to a physical address begins by comparing the virtual address from the processor with the virtual addresses in the TLB; there is a match when the virtual page number (VPN) of the address is the same as the VPN field of the entry, and either:

- the Global (G) bit of the TLB entry is set, or
- the ASID field of the virtual address is the same as the ASID field of the TLB entry.

This match is referred to as a *TLB hit*. If there is no match, a TLB Miss exception is taken by the processor and software is allowed to refill the TLB from a page table of virtual/physical addresses in memory.

If there is a virtual address match in the TLB, the physical address is output from the TLB and concatenated with the *Offset*, which represents an address within the page frame space. The *Offset* does not pass through the TLB.

Virtual-to-physical translation is described in greater detail throughout the remainder of this chapter; Figure 8-1 is a flow diagram of the process shown at the end of this chapter. The next two sections describe the 32-bit and 64-bit address translations.

8.2.4 32-bit Mode Address Translation

Figure 8-1 shows the virtual-to-physical-address translation of a 32-bit mode address. This figure illustrates two of the possible page sizes: a 4-Kbyte page (12 bits) and a 16-Mbyte page (24 bits).

- The top portion of Figure 8-1 shows a virtual address with a 12-bit, or 4-Kbyte, page size, labeled *Offset*. The remaining 20 bits of the address represent the VPN, and Index the 1M-entry page table.
- The bottom portion of Figure 8-1 shows a virtual address with a 24-bit, or 16-Mbyte, page size, labeled *Offset*. The remaining 8 bits of the address represent the VPN, and index the 256-entry page table.

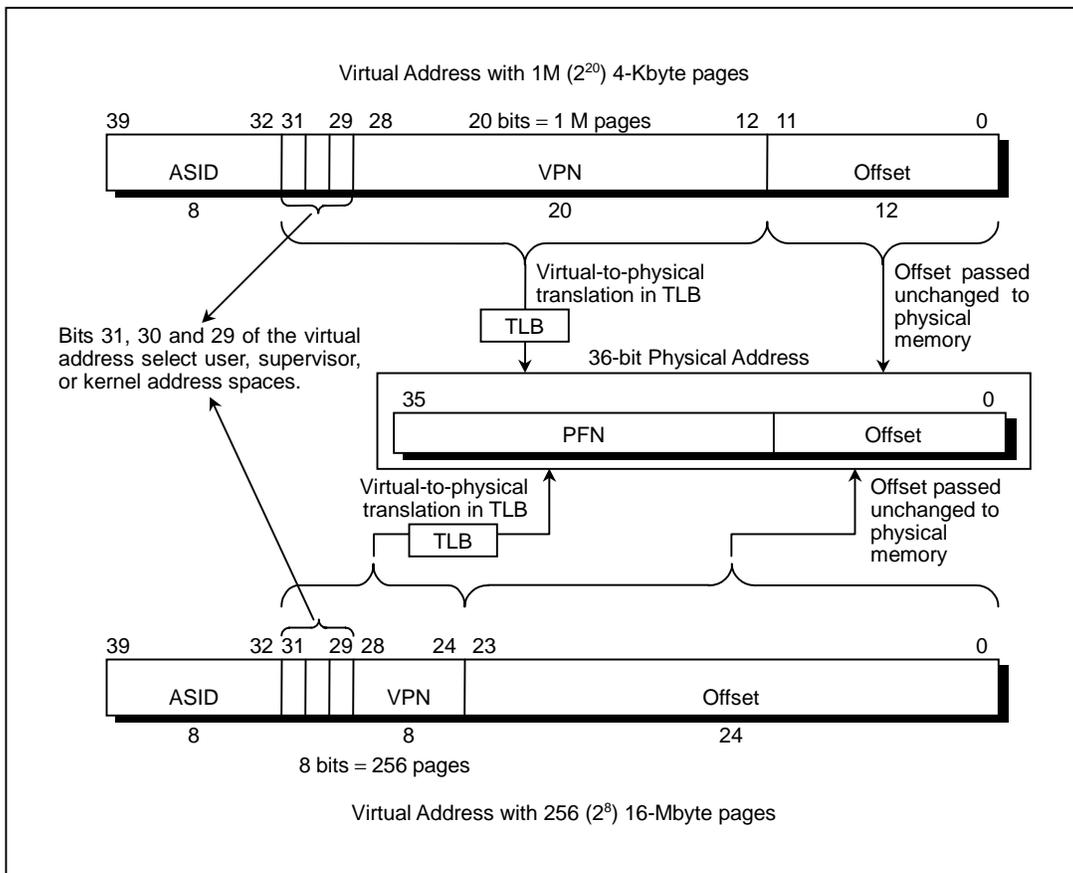


Figure 8-1 32-bit Mode Virtual Address Translation

8.2.5 64-bit Mode Address Translation

Figure 8-1 shows the virtual-to-physical-address translation of a 64-bit mode address. This figure illustrates two of the possible page sizes: a 4-Kbyte page (12 bits) and a 16-Mbyte page (24 bits).

- The top portion of Figure 8-1 shows a virtual address with a 12-bit, or 4-Kbyte, page size, labelled *Offset*. The remaining 28 bits of the address represent the VPN, and index the 256M-entry page table.
- The bottom portion of Figure 8-1 shows a virtual address with a 24-bit, or 16-Mbyte, page size, labelled *Offset*. The remaining 16 bits of the address represent the VPN, and index the 64K-entry page table.

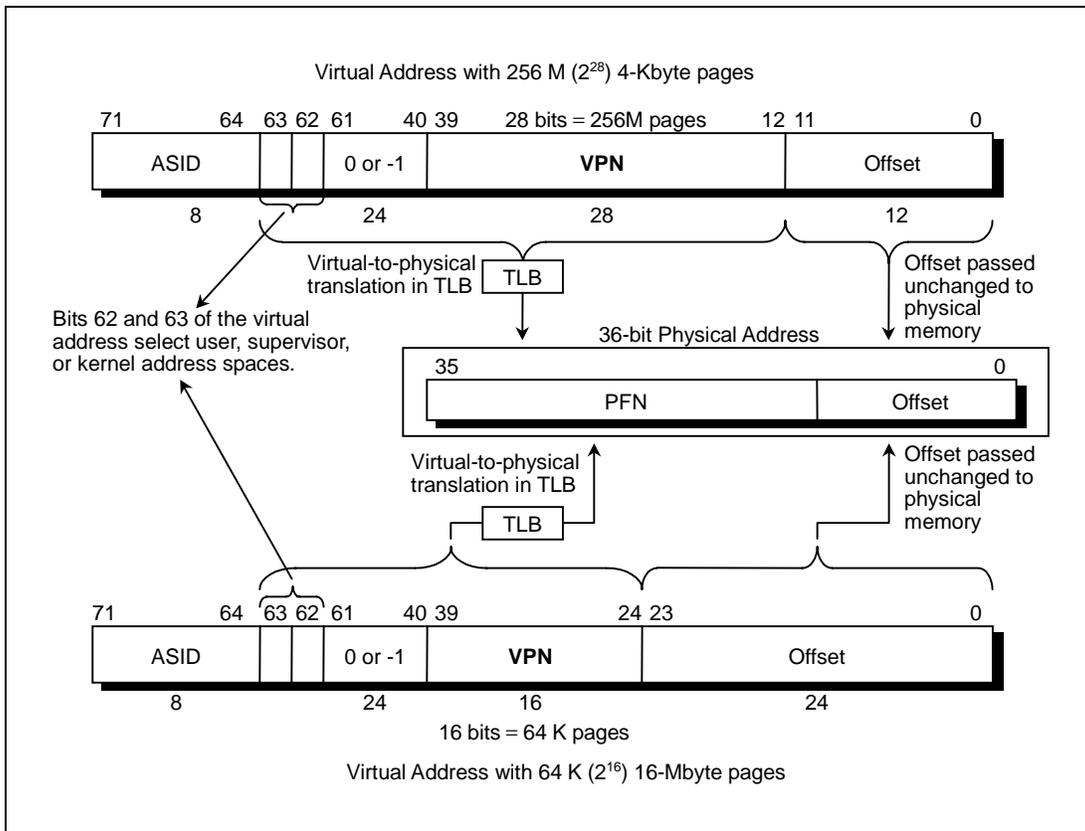


Figure 8-1 64-bit Mode Virtual Address Translation

8.3 Operating Modes

The TX49 has the three operating modes, User mode, Supervisor mode and Kernel mode, for 32- and 64-bit operation. The KSU, EXL and ERL bit in the Status register select User, Supervisor or Kernel mode. The UX, SX and KX bit in the Status register select 32- or 64-bit addressing in user, supervisor and kernel mode respectively.

KSU	EXL	ERL	UX	SX	KX	Mode
10	0	0	0	-	-	32-bit addressing in user mode
10	0	0	1	-	-	64-bit addressing in user mode
01	0	0	-	0	-	32-bit addressing in supervisor mode
01	0	0	-	1	-	64-bit addressing in supervisor mode
00	-	-	-	-	0	32-bit addressing in kernel mode
-	1	-	-	-	0	32-bit addressing in kernel mode
-	-	1	-	-	0	32-bit addressing in kernel mode
00	-	-	-	-	1	64-bit addressing in kernel mode
-	1	-	-	-	1	64-bit addressing in kernel mode
-	-	1	-	-	1	64-bit addressing in kernel mode

8.3.1 User Mode Operations

In User mode, a single, uniform virtual address space-labelled User segment-is available; its size is:

- 2 Gbytes (2^{31} bytes) in 32-bit mode (*useg*)
- 1 Tbyte (2^{40} bytes) in 64-bit mode (*xuseg*)

Figure 8-1 shows User mode virtual address space.

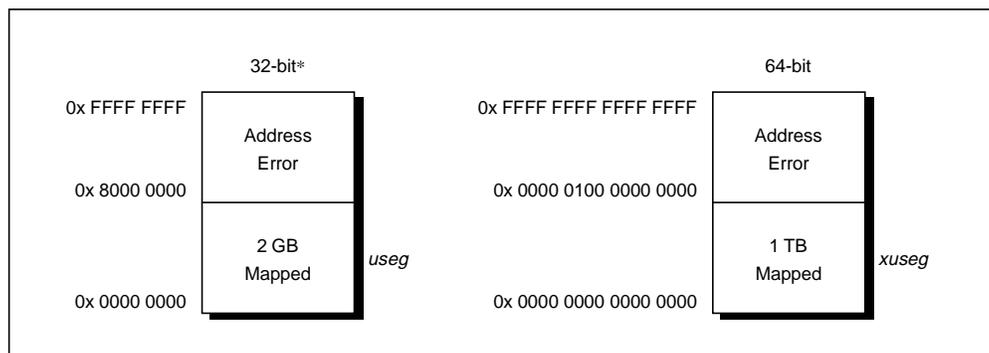


Figure 8-1 User Mode Virtual Address Space

*NOTE: In 32-bit mode, bit31 is sign-extended through bits 63~32. Failure results in an address error exception.

The User segment starts at address 0 and the current active user process resides in either *useg* (in 32-bit mode) or *xuseg* (in 64-bit mode). The TLB identically maps all references to *useg/xuseg* from all modes, and controls cache accessibility.

The processor operates in User mode when the Status register contains the following bit-values:

- KSU bits = 102
- EXL = 0
- ERL = 0

In conjunction with these bits, the *UX* bit in the Status register selects between 32- or 64-bit User mode addressing as follows:

- when *UX* = 0, 32-bit useg space is selected and TLB misses are handled by the 32-bit TLB refill exception handler
- when *UX* = 1, 64-bit xuseg space is selected and TLB misses are handled by the 64-bit TLB refill exception handler

Table 8-1 lists the characteristics of the two user mode segments, useg and xuseg.

Table 8-1 32-bit and 64-bit User Mode Segments

Address Bit Values	Status Register Bit Values				Segment Name	Address Range	Segment Size
	KSU	EXL	ERL	UX			
32-bit A (31) = 0	10 ₂	0	0	0	<i>useg</i>	0x0000 0000 through 0x7FFF FFFF	2• Gbyte (2 ³¹ • bytes)
64-bit A (63~40) = 0	10 ₂	0	0	1	<i>xuseg</i>	0x0000 0000 0000 0000 through 0x0000 00FF FFFF FFFF	1• Tbyte (2 ⁴⁰ • bytes)

32-bit User Mode (*useg*)

In User mode, when *UX* = 0 in the Status register, User mode addressing is compatible with the 32-bit addressing model shown in Figure 8-1, and a 2-Gbyte user address space is available, labelled useg.

All valid User mode virtual addresses have their most-significant bit cleared to 0; any attempt to reference an address with the most-significant bit set while in User mode causes an Address Error exception.

The system maps all references to useg through the TLB, and bit settings within the TLB entry for the page determine the cacheability of a reference.

64-bit User Mode (*xuseg*)

In User mode, when *UX* = 1 in the Status register, User mode addressing is extended to the 64-bit model shown in Figure 8-1. In 64-bit User mode, the processor provides a single, uniform address space of 2⁴⁰ bytes, labelled xuseg.

All valid User mode virtual addresses have bits 63~40 equal to 0; an attempt to reference an address with bits 63~40 not equal to 0 causes an Address Error exception.

8.3.2 Supervisor Mode Operations

Supervisor mode is designed for layered operating systems in which a true kernel runs in TX49 Kernel mode, and the rest of the operating system runs in Supervisor mode.

The processor operates in Supervisor mode when the Status register contains the following bit-values:

- *KSU* = 01₂
- *EXL* = 0
- *ERL* = 0

In conjunction with these bits, the *SX* bit in the Status register selects between 32- or 64-bit Supervisor mode addressing:

- when *SX* = 0, 32-bit supervisor space is selected and TLB misses are handled by the 32-bit TLB refill exception handler
- when *SX* = 1, 64-bit supervisor space is selected and TLB misses are handled by the 64-bit XTLB refill exception handler

Figure 8-1 shows Supervisor mode address mapping. Table 8-1 lists the characteristics of the supervisor mode segments; descriptions of the address spaces follow.

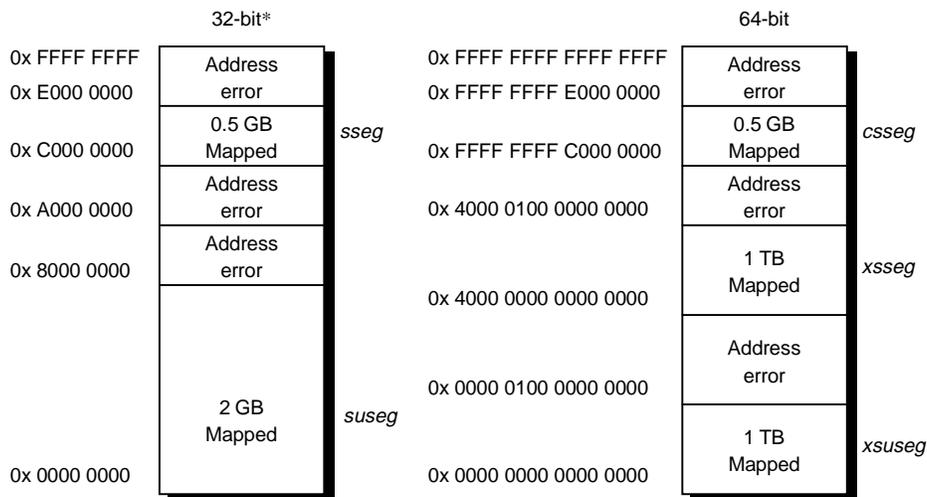


Figure 8-1 Supervisor Mode Address Space

*NOTE: In 32-bit mode, bit31 is sign-extended through bits 63~32. Failure results in an address error exception.

Table 8-1 32-bit and 64-bit Supervisor Mode Segments

Address Bit Values	Status Register Bit Values				Segment Name	Address Range	Segment Size
	KSU	EXL	ERL	SX			
32-bit A (31) = 0	0 ₁₂	0	0	0	<i>suseg</i>	0x0000 0000 through 0x7FFF FFFF	2• Gbyte (2 ³¹ • bytes)
32-bit A (31~29) = 110 ₂	0 ₁₂	0	0	0	<i>ssseg</i>	0xC000 0000 through 0xDFFF FFFF	512• Mbytes (2 ²⁹ • bytes)
64-bit A (63~62) = 00 ₂	0 ₁₂	0	0	1	<i>xsuseg</i>	0x0000 0000 0000 0000 through 0x0000 00FF FFFF FFFF	1• Tbyte (2 ⁴⁰ • bytes)
64-bit A (63~62) = 01 ₂	0 ₁₂	0	0	1	<i>xsseg</i>	0x4000 0000 0000 0000 through 0x4000 00FF FFFF FFFF	1• Tbyte (2 ⁴⁰ • bytes)
64-bit A (63~62) = 11 ₂	0 ₁₂	0	0	1	<i>csseg</i>	0xFFFF FFFF C000 0000 through 0xFFFF FFFF DFFF FFFF	512• Mbytes (2 ²⁹ • bytes)

32-bit Supervisor Mode, User Space (*suseg*)

In Supervisor mode, when $SX = 0$ in the *Status* register and the most-significant bit of the 32-bit virtual address is set to 0, the *suseg* virtual address space is selected; it covers the full 2³¹ bytes (2 Gbytes) of the current user address space. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address. This mapped space starts at virtual address 0x0000 0000 and runs through 0x7FFF FFFF.

32-bit Supervisor Mode, Supervisor Space (*sseg*)

In Supervisor mode, when $SX = 0$ in the *Status* register and the three most-significant bits of the 32-bit virtual address are 110₂, the *sseg* virtual address space is selected; it covers 2²⁹ bytes (512 Mbytes) of the current supervisor address space. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address. This mapped space begins at virtual address 0xC000 0000 and runs through 0xDFFF FFFF.

64-bit Supervisor Mode, User Space (*xsuseg*)

In Supervisor mode, when $SX = 1$ in the *Status* register and bits 63:62 of the virtual address are set to 00₂, the *xsuseg* virtual address space is selected; it covers the full 2⁴⁰ bytes (1 Tbyte) of the current user address space. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address. This mapped space starts at virtual address 0x0000 0000 0000 0000 and runs through 0x0000 00FF FFFF FFFF.

64-bit Supervisor Mode, Current Supervisor Space (*xsseg*)

In Supervisor mode, when $SX = 1$ in the *Status* register and bits 63~62 of the virtual address are set to 012, the *xsseg* current supervisor virtual address space is selected. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address. This mapped space begins at virtual address 0x4000 0000 0000 0000 and runs through 0x4000 00FF FFFF FFFF.

64-bit Supervisor Mode, Separate Supervisor Space (*csseg*)

In Supervisor mode, when $SX = 1$ in the *Status* register and bits 63~62 of the virtual address are set to 112, the *csseg* separate supervisor virtual address space is selected. Addressing of the *csseg* is compatible with addressing *sseg* in 32-bit mode. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address. This mapped space begins at virtual address 0xFFFF FFFF C000 0000 and runs through 0xFFFF FFFF DFFF FFFF.

8.3.3 Kernel Mode Operations

The processor operates in Kernel mode when the Status register contains one or more of the following values:

- $KSU = 00_2$
- $EXL = 1$
- $ERL = 1$

In conjunction with these bits, the KX bit in the Status register selects between 32- or 64-bit Kernel mode addressing:

- when $KX = 0$, 32-bit kernel space is selected and all TLB misses are handled by the 32-bit TLB refill exception handler
- when $KX = 1$, 64-bit kernel space is selected and all TLB misses are handled by the 64-bit XTLB refill exception handler

The processor enters Kernel mode whenever an exception is detected and it remains in Kernel mode until an Exception Return (ERET) instruction is executed and results in ERL and/or $EXL = 0$. The ERET instruction restores the processor to the mode existing prior to the exception.

Kernel mode virtual address space is divided into regions differentiated by the high-order bits of the virtual address, as shown in Figure 8-1. Table 8-1 lists the characteristics of the 32-bit kernel mode segments, and Table 8-2 lists the characteristics of the 64-bit kernel mode segments.

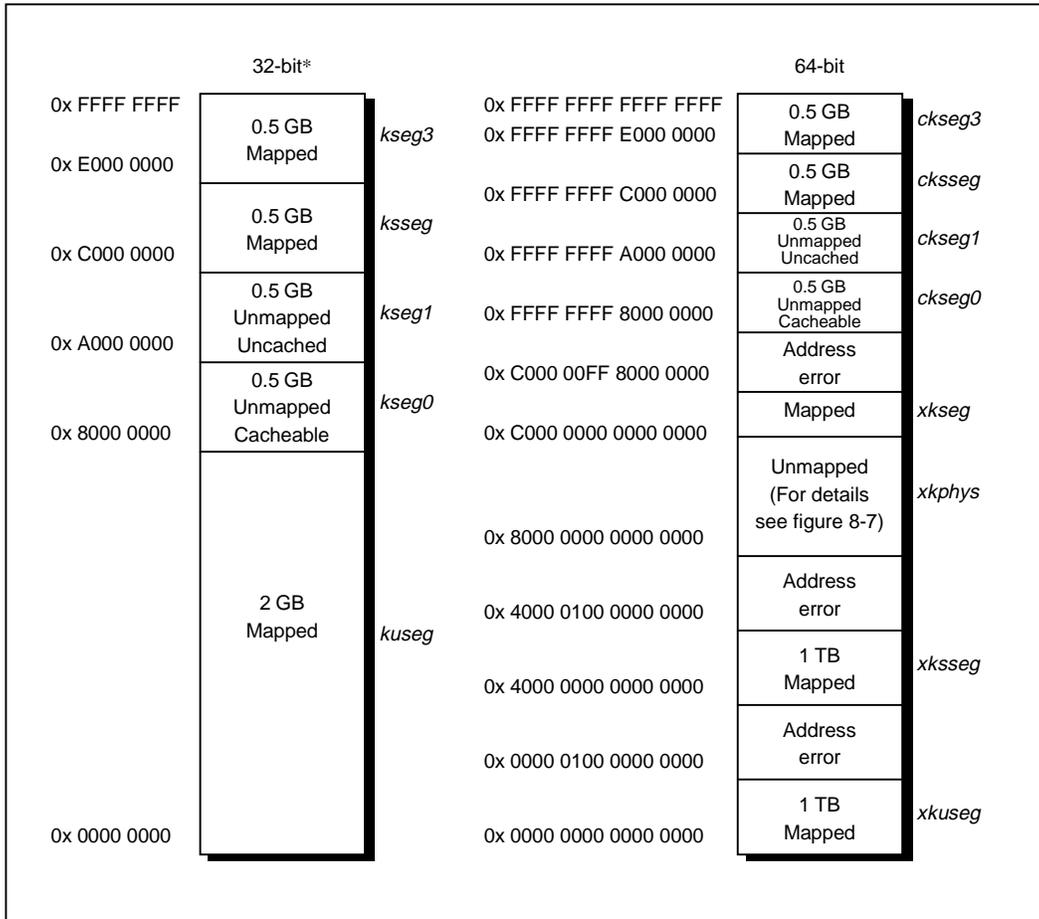


Figure 8-1 Kernel Mode Address Space

*NOTE: In 32-bit mode, bit31 is sign-extended through bits 63~32. Failure results in an address error exception.

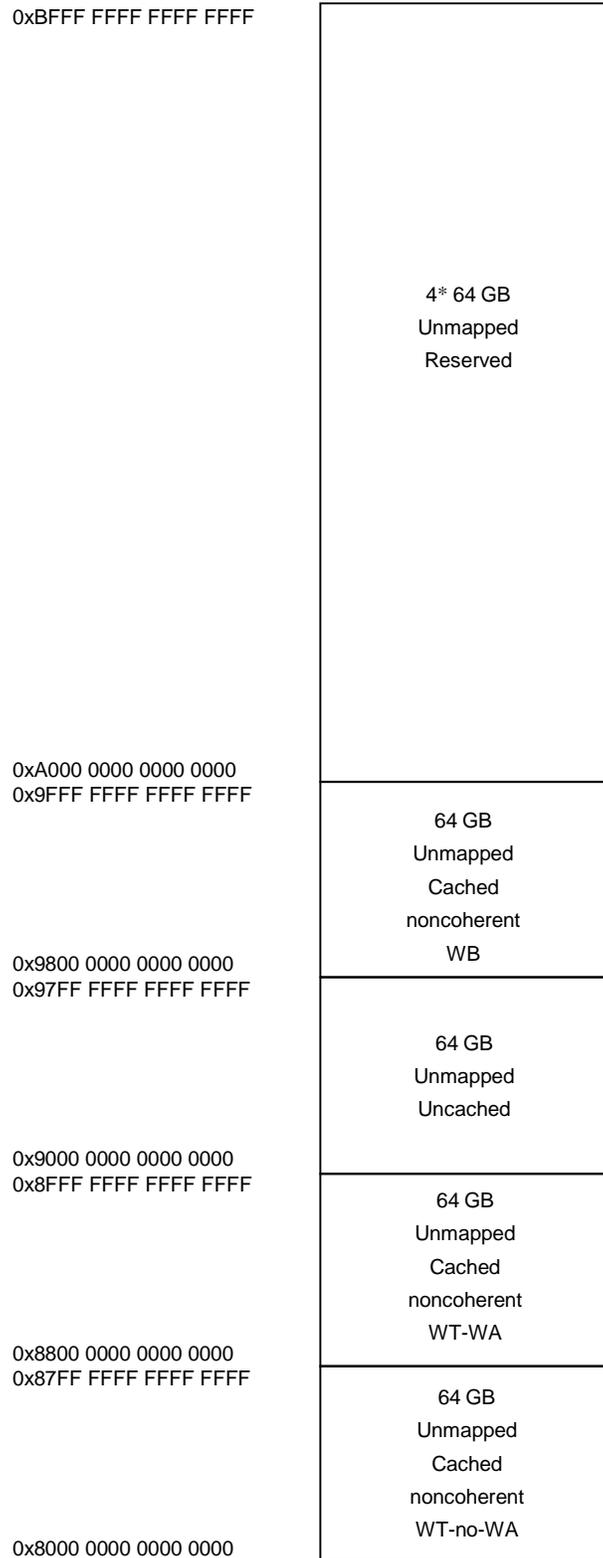


Figure 8-2 xkphys Address Space

Table 8-1 32-bit Kernel Mode Segments

Address Bit Values	Status Register Is One Of These Values				Segment Name	Address Range	Segment Size
	KSU	EXL	ERL	KX			
A (31) = 0	KSU = 00 ₂ or EXL = 1 or ERL = 1			0	<i>Kuseg</i>	0x0000 0000 through 0x7FFF FFFF	2• Gbyte (2 ³¹ • bytes)
A (31~29) = 100 ₂				0	<i>Kseg0</i>	0x8000 0000 through 0x9FFF FFFF	512• Mbytes (2 ²⁹ • bytes)
A (31~29) = 101 ₂				0	<i>Kseg1</i>	0xA000 0000 through 0xBFFF FFFF	512• Mbytes (2 ²⁹ • bytes)
A (31~29) = 110 ₂				0	<i>Ksseg</i>	0xC000 0000 through 0xDFFF FFFF	512• Mbytes (2 ²⁹ • bytes)
A (31~29) = 111 ₂				0	<i>Kseg3</i>	0xE000 0000 through 0xFFFF FFFF	512• Mbytes-4• Mbytes (2 ²⁹ • bytes)
				0	(Reserved)	0xFF00 0000 through 0xFF3F FFFF	4• Mbytes

32-bit Kernel Mode, User Space (*kuseg*)

In Kernel mode, when $KX = 0$ in the *Status* register, and the most-significant bit of the virtual address, A₃₁, is cleared, the 32-bit *kuseg* virtual address space is selected; it covers the full 2³¹ bytes (2 Gbytes) of the current user address space. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address. When ERL = 1 in the *Status* register, the user address region becomes a 2³¹ bytes unmapped (that is, mapped directly to physical addresses) uncached address space.

32-bit Kernel Mode, Kernel Space 0 (*kseg0*)

In Kernel mode, when $KX = 0$ in the *Status* register and the most-significant three bits of the virtual address are 100₂, 32-bit *kseg0* virtual address space is selected; it is the 2²⁹ bytes (512 Mbyte) kernel physical space. References to *kseg0* are not mapped through the TLB; the physical address selected is defined by subtracting 0x8000 0000 from the virtual address. The *K0* field of the *Config* register, described in this chapter, controls cacheability and coherency.

32-bit Kernel Mode, Kernel Space 1 (*kseg1*)

In Kernel mode, when $KX = 0$ in the *Status* register and the most-significant three bits of the 32-bit virtual address are 1012, 32-bit *kseg1* virtual address space is selected; it is the 2^{29} bytes (512 Mbyte) kernel physical space. References to *kseg1* are not mapped through the TLB; the physical address selected is defined by subtracting 0xA000 0000 from the virtual address. Caches are disabled for accesses to these addresses, and physical memory (or memory-mapped I/O device registers) are accessed directly.

32-bit Kernel Mode, Supervisor Space (*ksseg*)

In Kernel mode, when $KX = 0$ in the *Status* register and the most-significant three bits of the 32-bit virtual address are 1102, the *ksseg* virtual address space is selected; it is the current 2^{29} bytes (512 Mbyte) supervisor virtual space. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address.

32-bit Kernel Mode, Kernel Space 3 (*kseg3*)

In Kernel mode, when $KX = 0$ in the *Status* register and the most-significant three bits of the 32-bit virtual address are 1112, the *kseg3* virtual address space is selected; it is the current 2^{29} bytes (512 Mbyte-4 Mbyte) kernel virtual space. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address.

Note: These is the 4 Mbytes Reserved area, begin at virtual address 0xFF00_0000 and runs through 0xFF3F_FFFF.

Table 8-2 64-bit Kernel Mode Segments

Address Bit Values	Status Register Is One Of These Values				Segment Name	Address Range	Segment Size
	KSU	EXL	ERL	KX			
A (63-62) = 00 ₂	KSU = 00 ₂ or EXL = 1 or ERL = 1			1	<i>xkuseg</i>	0x0000 0000 0000 0000 through 0x0000 00FF FFFF FFFF	1• Tbytes (2 ⁴⁰ • bytes)
A (63-62) = 01 ₂				1	<i>xksseg</i>	0x4000 0000 0000 0000 through 0x4000 00FF FFFF FFFF	1• Tbytes (2 ⁴⁰ • bytes)
A (63-62) = 10 ₂				1	<i>xkphys</i>	0x8000 0000 0000 0000 through 0xBFFF FFFF FFFF FFFF	8*2 ³² • bytes
A (63-62) = 11 ₂				1	<i>xkseg</i>	0xC000 0000 0000 0000 through 0xC000 00FF 7FFF FFFF	2 ⁴⁰ -2 ³¹ • bytes
A (63-62) = 11 ₂ A (61-31) = -1				1	<i>ckseg0</i>	0xFFFF FFFF 8000 0000 through 0xFFFF FFFF 9FFF FFFF	512• Mbytes (2 ²⁹ • bytes)
A (63-62) = 11 ₂ A (61-31) = -1				1	<i>ckseg1</i>	0xFFFF FFFF A000 0000 through 0xFFFF FFFF BFFF FFFF	512• Mbytes (2 ²⁹ • bytes)
A (63-62) = 11 ₂ A (61-31) = -1				1	<i>cksseg</i>	0xFFFF FFFF C000 0000 through 0xFFFF FFFF DFFF FFFF	512• Mbytes (2 ²⁹ • bytes)
A (63-62) = 11 ₂ A (61-31) = -1				1	<i>ckseg3</i>	0xFFFF FFFF E000 0000 through 0xFFFF FFFF FFFF FFFF	512• Mbytes -4• Mbyte
				1	(Reserved)	0xFFFF FFFF FF00 0000 through 0xFFFF FFFF FF3F FFFF	4• Mbytes

64-bit Kernel Mode, User Space (*xkuseg*)

In Kernel mode, when $KX = 1$ in the *Status* register and bits 63~62 of the 64-bit virtual address are 00₂, the *xkuseg* virtual address space is selected; it covers the current user address space. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address.

When $ERL = 1$ in the *Status* register, the user address region becomes a 2^{31} bytes unmapped (that is, mapped directly to physical addresses) uncached address space.

64-bit Kernel Mode, Current Supervisor Space (*xksseg*)

In Kernel mode, when $KX = 1$ in the *Status* register and bits 63~62 of the 64-bit virtual address are 01₂, the *xksseg* virtual address space is selected; it is the current supervisor virtual space. The virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address.

64-bit Kernel Mode, Physical Spaces (*xkphys*)

In Kernel mode, when $KX = 1$ in the *Status* register and bits 63~62 of the 64-bit virtual address are 10₂, one of the two unmapped *xkphys* address spaces are selected, either cached or uncached. Accesses with address bits 58~36 not equal to 0 cause an address error.

References to this space are not mapped; the physical address selected is taken from bits 35~0 of the virtual address. Bits 61~59 of the virtual address specify the cacheability and coherency attributes, as shown in Table 8-3.

Table 8-3 Cacheability and Coherency Attributes

Value(61~59)	Cacheability and Coherency Attributes	Starting Address
0	Cacheable, non-coherent, write-through, no write allocate	0x8000 0000 0000 0000
1	Cacheable, non-coherent, write-through, no write allocate	0x8800 0000 0000 0000
2	Uncached	0x9000 0000 0000 0000
3	Cacheable, non-coherent	0x9800 0000 0000 0000
4-7	Reserved	0xA000 0000 0000 0000

64-bit Kernel Mode, Kernel Space (*xkseg*)

In Kernel mode, when $KX = 1$ in the *Status* register and bits 63~62 of the 64-bit virtual address are 11₂, the address space selected is one of the following:

- kernel virtual space, *xkseg*, the current kernel virtual space; the virtual address is extended with the contents of the 8-bit ASID field to form a unique virtual address
- one of the four 32-bit kernel compatibility spaces, as described in the next section.

64-bit Kernel Mode, Compatibility Spaces (*ckseg1~0, cksseg, ckseg3*)

In Kernel mode, when *KX* = 1 in the *Status* register, bits 63~62 of the 64-bit virtual address are 112, and bits 61~31 of the virtual address equal-1, the lower two bytes of address, as shown in Figure 8-1, select one of the following 512-Mbytes compatibility spaces.

- *ckseg0*. This 64-bit virtual address space is an unmapped region, compatible with the 32-bit address model
- *kseg0*. The *K0* field of the *Config* register, described in this chapter, controls cacheability and coherency.
- *ckseg1*. This 64-bit virtual address space is an unmapped and uncached region, compatible with the 32-bit address model *kseg1*.
- *cksseg*. This 64-bit virtual address space is the current supervisor virtual space, compatible with the 32-bit address model *ksseg*.
- *ckseg3*. This 64-bit virtual address space is kernel virtual space, compatible with the 32-bit address model *kseg3*.

8.4 Translation Look aside Buffer

8.4.1 Joint TLB

The TX49 has a fully associative TLB which maps 48 pairs(odd/even entry) of virtual pages to their corresponding physical addresses.

8.4.2 TLB Entry format

32-bit addressing

127	121	120	109	108	96				
0	MASK			0					
95	77 76 75 72 71				64				
VPN2				G	0	ASID			
63	62	61	38 37			35	34	33	32
0	PFN					C	D	V	0
31	30	29	6 5			3	2	1	0
0	PFN					C	D	V	0

64-bit addressing

255	217	216	205	204	96					
0	MASK			0						
191	190	189	168	167	141	140	139	136	135	128
R	0	VPN2			G	0	ASID			
127	94 93		70 69			67	66	65	64	
0	PFN					C	D	V	0	
63	30 29		6 5			3	2	1	0	
0	PFN					C	D	V	0	

MASK : Page comparison mask

VPN2 : Virtual page number divided by two (maps to two pages)

ASID : Address space ID field.

R : Region. (00: user, 01: supervisor, 11: kernel) used to match Vaddr63~62.

PFN : Page frame number.

C : Specifies the cache algorithm to be used (see the "C" field of the EntryLo0, 1).

D : Dirty. If this bit is set, the page is marked as dirty and therefore, writable.

V : Valid. If this bit is set, it indicates that the TLB entry is valid.

G : Global. If this bit is set in both Lo0 and Lo1, then ignore the ASID during TLB lookup.

0 : Reserved. Returns zeroes when read.

8.4.3 Instruction-TLB

The TX49 has a 2-entry instruction TLB (ITLB). Each ITLB entry is a subset of any single JTLB entry. The ITLB is completely invisible to software.

8.4.4 Data-TLB

The TX49 has a 4-entry data TLB (DTLB). Each DTLB entry is a subset of any single JTLB entry. The DTLB is completely invisible to software.

8.5 Virtual-to-Physical Address Translation Process

During virtual-to-physical address translation, the CPU compares the 8-bit ASID (if the Global bit, G, is not set) of the virtual address to the ASID of the TLB entry to see if there is a match. One of the following comparisons are also made:

- In 32-bit mode, the highest 7 to 19 bits (depending upon the page size) of the virtual address are compared to the contents of the TLB VPN2 (virtual page number divided by two).
- In 64-bit mode, the highest 15 to 27 bits (depending upon the page size) of the virtual address are compared to the contents of the TLB VPN2 (virtual page number divided by two).

If a TLB entry matches, the physical address and access control bits (C, D, and V) are retrieved from the matching TLB entry. While the V bit of the entry must be set for a valid translation to take place, it is not involved in the determination of a matching TLB entry.

Figure 8-1 illustrates the TLB address translation process.

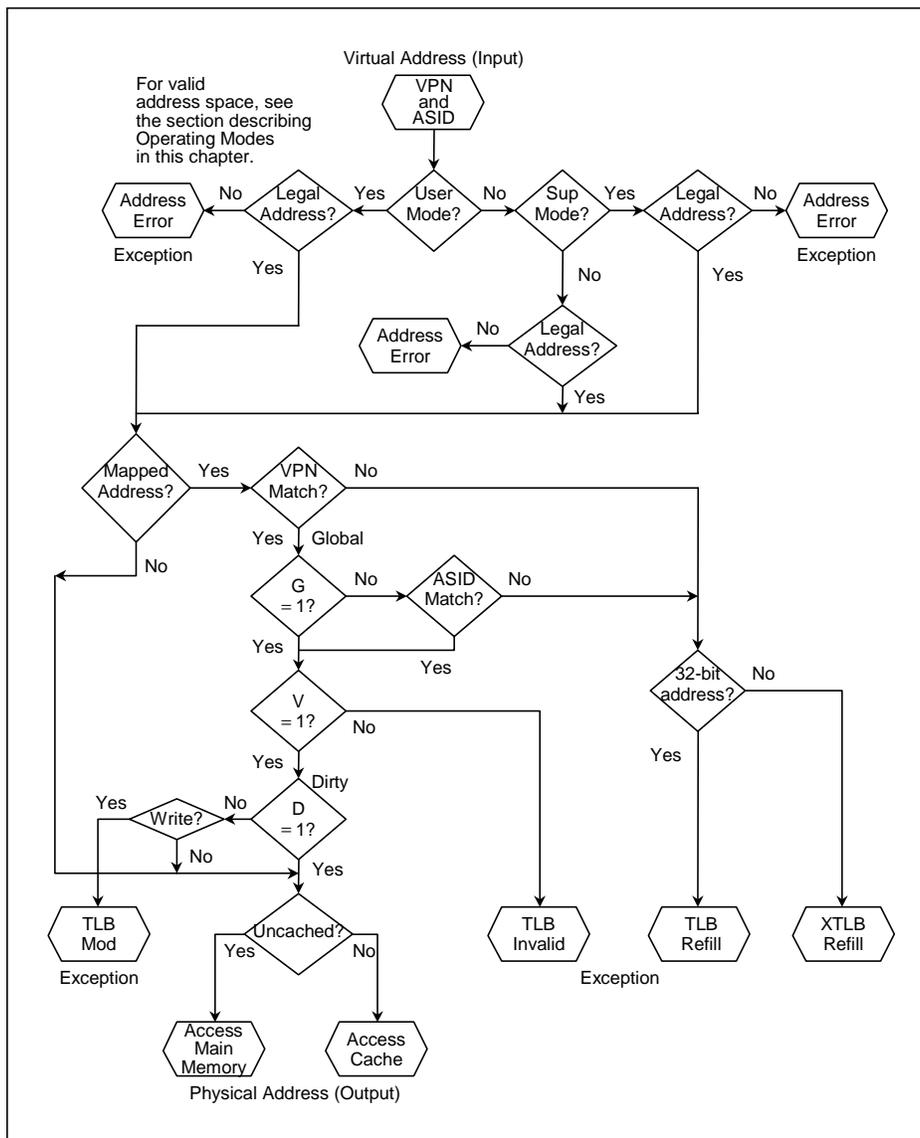


Figure 8-1 TLB Address Translation

9. Cache Organization

9.1 Introduction

This chapter describes the cache memory of TX49. This processor has two on-chip primary caches for instruction and data. Both caches are configured as either 8K-byte, 16K-byte or 32K-byte in size.

9.2 Instruction Cache (I-Cache)

The TX49 primary I-cache has the following characteristics:

- Cache size: 8K/16K/32K (Selectable)
- Four-way set associative
- Indexed with a virtual address
- Checked with a physical tag
- Block (line) size: 8 words (32 bytes)
- Burst refill size: 8 words (32 bytes)
- Lockable on a per-line basis
- All valid bits, lock and FIFO bits are cleared by a Reset exception

9.2.1 Instruction cache address field

Figure 9-1 shows the instruction cache address field.

When 4-KB page size is used in 32K Instruction cache, the bit 12 of the Physical Address and the Virtual Address must be same value.

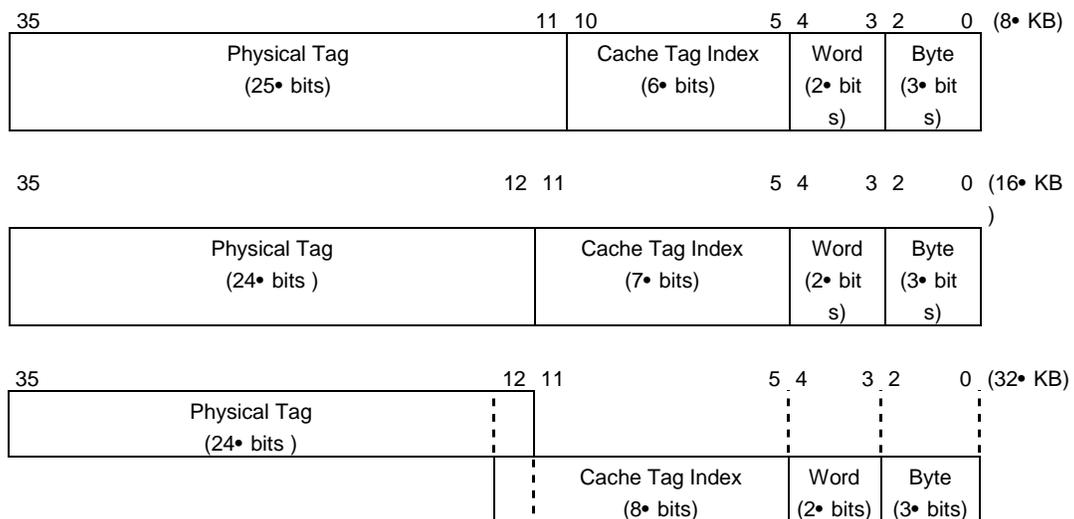


Figure 9-1 Instruction cache address field

9.2.2 Instruction cache configuration

Each line in the 4 ways of the instruction cache share F1, F0 replacement bits. Figure 9-1 shows the format of replacement bits. These bits are shared by way0, way1, way2 and way3 for 8KB/16KB/32KB cache, and indicate next set to which replacement will be directed; when lock bit is set to 1, indicate this set is not locked.

Each line of instruction cache data has an associated 27-bit/26-bit tag that contains a 25-bit/24-bit physical address, a single Lock bit and a single valid bit, except for the line in way0, which has an 26-bit/25-bit tag that excludes a lock bit. Figure 9-2 shows the formats of tag and data pair.



Figure 9-1 Format of replacement bits

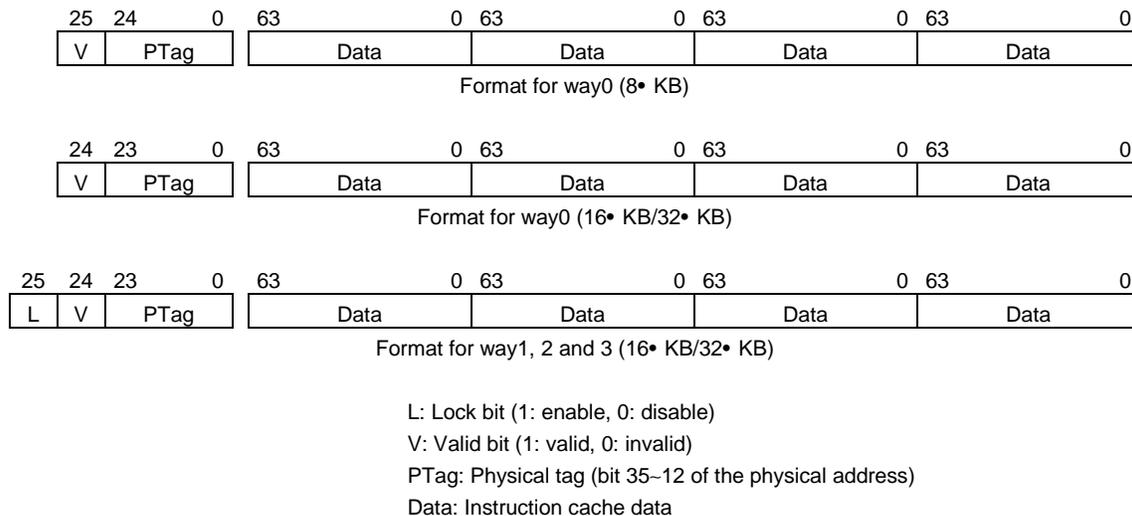


Figure 9-2 Format of tag and data pair for I-cache

9.3 Data Cache

The TX49 primary D-cache has the following characteristics:

- Cache size: 8KB/16KB/32KB (Selectable)
- Four-way set associative
- Indexed with a virtual address
- Checked with a physical tag
- Block (line) size: 8 words (32bytes)
- Burst size: 8 words (32bytes)
- Snoop function
- Store buffer
- Lockable on a per-line basis
- Write-back or write-through on a per-page basis

- All write-back, CS, FIFO and lock bits are cleared by a Reset exception

9.3.1 Data cache address field

Figure 9-1 shows the data cache address field.

When 4-KB page size is used in 32-KB Instruction cache, the bit 12 of the Physical Address and the Virtual Address must be same value.

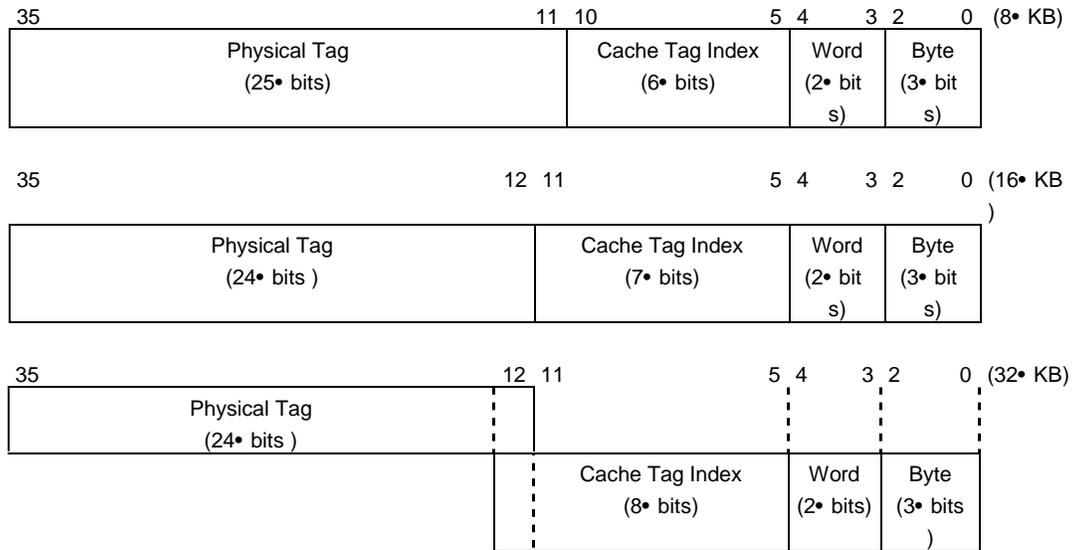
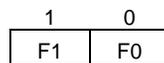


Figure 9-1 Data cache address field

9.3.2 Data cache configuration

Each line in the 4 ways of the data cache share F1, F0 replacement bits. Figure 9-1 shows the format of replacement bits. These bits are shared by way0, way1, way2 and way3 for 8-KB/16-KB/32-KB cache, and indicate next set to which replacement will be directed; when lock bit is set to 1, indicate this set is not locked.

Each line of data cache data has an associated 29-bit/28-bit tag that contains a 25-bit/24-bit physical address, a single Lock bit, a single write-back bit and a 2-bit cache state, except for the line in way0, which has an 28-bit/27-bit tag that excludes a Lock bit. Figure 9-2 shows the formats of tag and data pair.



F0: FIFO replace bit 0

F1: FIFO replace bit 1

Figure 9-1 Format of replacement bits

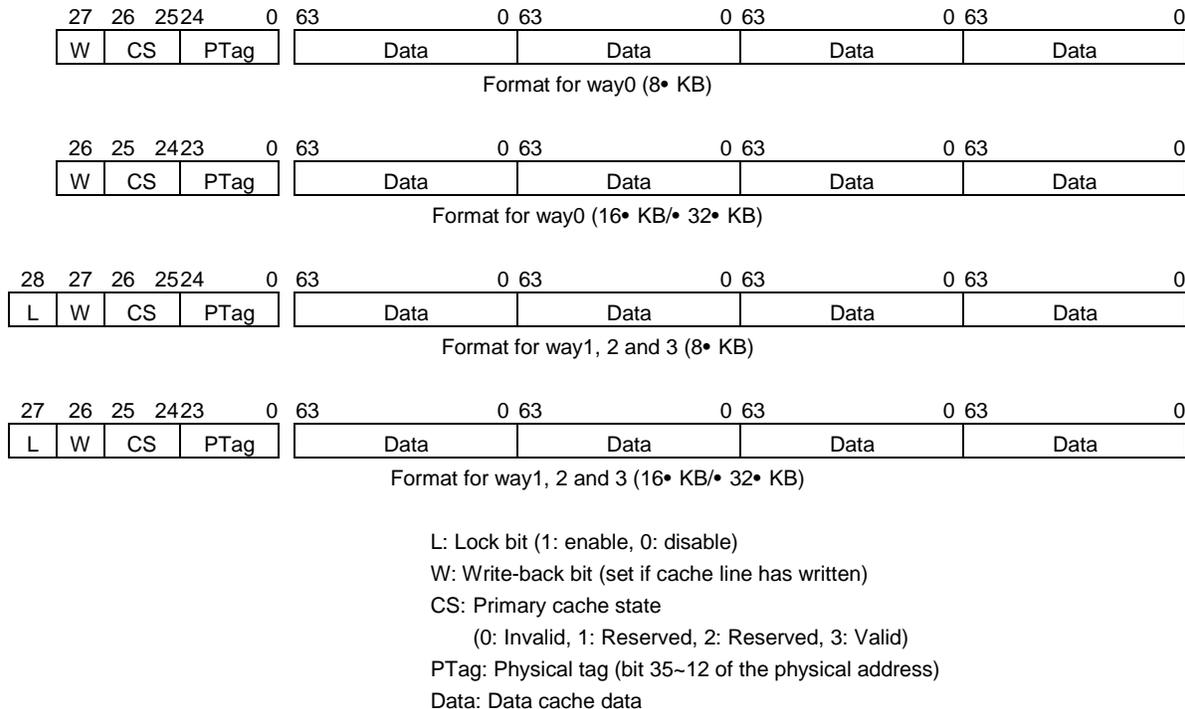


Figure 9-2 Format of tag and data pair for D-cache

In the TX49, the W (write-back) bit, not the cache state, indicates when the primary cache contents modified data that must be written back to memory. The states Invalid and Valid are used to describe the cache line. That is, there is no hardware support for cache coherency.

9.4 Lock function

The lock function can be used to locate critical instruction/data in one instruction/data cache set and they are not replaced when the lock bit is set.

9.4.1 Lock bit setting and clearing

Setting the Lock bit in each line cache enable the instruction/data cache lock function. When the lock function is enabled, the instruction/data in the valid line is locked and never be replaced. The set to be locked is pointed by FIFO bit. Refilled instruction/data during the lock function is enabled is locked. When a store miss occurs for the write-through data cache without write allocate, the store data is not written to the cache and will therefore not be locked.

The lock function is disabled by clearing the Lock bit in each line.

In order to clear or set the Lock bit in the cache, Cache instructions (Index store I-cache /D-cache Tag) can be used, and in order to load the instruction/data to cache from memory, another Cache instructions (Fill I-cache/D-cache) can be used (refer to Cache instruction).

Clear the lock bit as follows when data written to a locked line should be stored in main memory.

- (1) Read the locked data from cache memory
- (2) Clear the lock bit
- (3) Store the data that was read

9.4.2 Operation during lock

After the lock bit is set for a line, the line can be replaced only when its line state is invalid. The locked valid line can never be replaced. FIFO bit should point only to the set of locked invalid line or unlocked line.

A write access to a locked valid line takes place only to the cache not to the memory at Write Back mode. Both of the cache and the memory are replaced at Write Through mode.

9.4.3 Example of Data cache locking

During the load operation to the locked line of the cache, any interrupt should be disabled in order to avoid to lock the wrong data.

To lock data cache lines, the following sequence of codes could be used.

```

.....          /* Disable the interrupt */
mtc0  t0, TagLo  /* Load data into TagLo reg */
cache  2 (D), offset (base) /* Invalidate and lock line in desired set using
                           Index_Store_Tag cache instruction */
cache  7 (D), offset (base) /* Fill the cache line from desired memory location */
.....          / Enable the interrupt */

```

9.4.4 Example of Instruction cache locking

To lock instruction cache lines, the following sequence of codes could be used:

```

.....          /* Disable the interrupt */
mtc0  t0, TagLo  /* Load data into TagLo reg */
cache  2 (I), offset (base) /* Invalidate and lock line in desired set using
                           Index_Store_Tag cache instruction */
cache  5 (I), offset (base) /* Fill the cache line from desired memory location */
.....          /* Enable the interrupt */

```

9.5 The primary cache accessing

Figure 9-1 shows the virtual address (VA) index to the primary cache. Each instruction and data cache size is 8KB, 16KB or 32KB. The virtual address bits be used to index into the primary cache decided by the cache size.

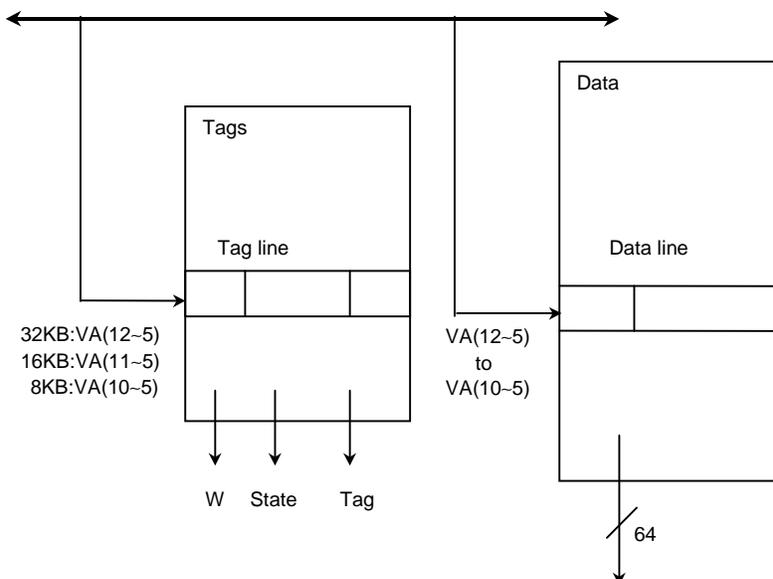


Figure 9-1 Primary Cache Data and Tag Organization

9.6 Cache States

The section describes about the state of a cache line. The cache line in the TX49 is in one of states described in Table 9-1.

The I-Cache line is in one of the following states:

- invalid
- valid

The D-Cache line is in one of the following states:

- invalid
- valid

Table 9-1 Cache States

Cache line State	Description
Invalid	A cache line that does not contain valid information must be marked invalid, and cannot be used. A cache line in any other state than invalid is assumed to contain valid information.
Valid	A Valid cache line contains valid information. The cache line may or not be consistent with memory and is owned by the processor (see Cache Line Ownership in this chapter).

9.7 Cache Line Ownership

The TX49 becomes the owner of a cache line after it writes to that cache line (that is, by entering the Valid), and is responsible for providing the contents of that line on a read request. There can only be one owner for each cache line.

9.8 Cache Multi-Hit Operation

The TX49 is not guaranteed the operation for the multi-hit of primary cache.

Thus, in case of locking the specified program/data in the primary cache, the program/data must be used after locked in the cache by Fill instruction.

Such as the previous description the cache multi hit does not guarantee in the TX49.

9.9 Cache Snoop

The TX49 has a bus arbitration function that releases bus mastership to an external bus master. Consistency between cache memory and main memory could deteriorate when an external bus master has write access to main memory. The purpose of the cache snoop function is to maintain this data consistency.

When TX49 releases the bus, the bus cycle is snooped by an external bus master. If an address access by the external bus master matches an address stored in the on-chip data cache (cache hit), the Cache State bit (CS) for that cache data is cleared to 0, invalidating it.

Locked data cannot be invalidated in a snoop, the FIFO bits point to the invalidated cache set. The lock bit is not changed as the result of a snoop.

Note: A snoop is possible even when the data cache is disabled.

The snooping of an external bus master is done by ST concurrency when the TX49 is in the doze mode. For the bus that is released by the assertion of the GSGNT* signal or the GHPSGNT* signal, snooping of the data cache can be performed by the GSNOOP* signal and the GA[35~0] signal. When an external bus master deasserts the GSREQ* signal or the GHPSREQ* signal, the TX49 deasserts the GSGNT* signal or the GHPSGNT* signal.

9.10 Cache Test Function

9.10.1 Cache disabling

The config register bits ICE# (Instruction Cache Enable) and DCE# (Data Cache Enable) are used to enable and disable the instruction and data cache, respectively.

When a cache is disabled, all cache accesses are misses and there is no refill (nor is there any burst bus cycle; this is the same as accessing a non-cacheable area). The valid bit (V) or Cache State bit (CS) for each entry cannot be modified.

9.10.2 Cache flushing

Both the instruction and data cache are flushed when a ColdReset/SoftReset exception is raised (all valid bits are cleared to 0).

The instruction cache is flushed by the CACHE instruction Index_Invalidate/Hit_Invalidate. The data cache is flushed by the CACHE instruction Hit_Invalidate.

10. Write Buffer

The TX49 contains a write buffer to improve the performance of writes to the external memory. Every write to external memory uses this on-chip write buffer. The write buffer holds up to four 64-bit address and data pairs.

For a cache miss write-back, the entire buffer is used for the write-back data and allows the processor to proceed in parallel with the memory update. For uncached and write-through stores, the write buffer uncouples the CPU from the write to memory. If the write buffer is full, additional stores will stall until there is room for them in the write buffer.

11. CPU Exception

11.1 Introduction

This chapter describes the explanation of CPU exception processing. The chapter concludes with a description of each exception's cause, together with the manner in which the CPU processes and services these exceptions.

11.2 Exception Vector Locations

Exception vector addresses are stored in an area of kseg0 or kseg1 except for Debug exception vector. The vector address of the ColdReset, SoftReset and NMI exception is always in a non-cacheable area of kseg1. Vector addresses of the other exceptions depend on the BEV bit of Status register. When BEV is 0, these exceptions are vectored to a cacheable area of kseg0. When BEV is 1, all vector addresses are in a non-cacheable area of kseg1.

Table 11-1 shows the list of the exception vector locations.

Table 11-1 Exception Vector Locations

Exception	TX49 Vector Address	
	(BEV• = 0)	(BEV• = 1)
ColdReset, SoftReset, NMI	0xffff_fff_bfc0_0000	0xffff_fff_bfc0_0000
TLB refill, EXL• =• 0	0xffff_fff_8000_0000	0xffff_fff_bfc0_0200
XTLB refill, EXL• =• 0 (X• =• 64• bit TLB)	0xffff_fff_8000_0080	0xffff_fff_bfc0_0280
Others	0xffff_fff_8000_0180	0xffff_fff_bfc0_0380

The cache error exception is not occurred because the TX49 does not have the parity bit into the primary cache.

Debug exception needs the care, it has the special address. (See 14.9.5)

Table 11-2 shows the list of the debug exception vector locations.

Table 11-2 Debug Exception Vector Locations

Exception	TX49 Debug Exception Vector Address	
	(ProbEnb• = 0)	(ProbEnb• = 1)
Debug	0xffff_fff_bfc0_0400	0xffff_fff_ff20_0200

11.3 Priority of Exception

More than one exception may be raised for the same instruction, in which case only the exception with the highest priority is reported. The TX49 Processor Core instruction exception priority is shown in Table 11-1.

Table 11-1 Priority of Exception

Priority	Exception		Mnemonic
High  ↓ Low	Cold Reset		
	Soft Reset		
	NMI		
	Address error	Inst. Fetch	AdEL
	TLB refill	Inst. Fetch	TLBL
	TLB invalid	Inst. Fetch	TLBL
	Bus error	Inst. Fetch	IBE
	Integer overflow, Trap, System Call, Breakpoint, Reserved Instruction, Coprocessor Unusable, or Floating-Point Exception		Ov, Tr, Sys, Bp, RI, CpU, FPE
	Address error	Data access	AdEL/AdES
	TLB refill	Data access	TLBL/TLBS
	TLB invalid	Data access	TLBL/TLBS
	TLB modified	Data write	Mod
	Bus error	Data access	DBE
	Interrupt		Int

11.4 ColdReset Exception

11.4.1 Cause

This ColdReset exception occurs when the GCOLDRESET* signal is asserted and then deasserted. This exception is not maskable.

11.4.2 Processing

A special interrupt vector that resides in an unmapped and uncached area is used. It is therefore not necessary for hardware to initialize TLB and cache memory in order to process this exception. The vector location of this exception is;

- In 32-bit mode, 0xbfc0 0000
- In 64-bit mode, 0xffff ffff bfc0 0000

The most register's contents are set/cleared when this exception occurs. The values of these bits are listed into the table of Section 7.

11.4.3 Servicing

The ColdReset exception is serviced by;

- initializing all registers, coprocessor registers, caches and the memory system
- performing diagnostic tests
- bootstrapping the operating system

11.5 SoftReset Exception

11.5.1 Cause

This SoftReset exception occurs when the GRESET* signal is asserted and then deasserted. This exception is not maskable.

11.5.2 Processing

A special interrupt vector that resides in an unmapped and uncached area is used. It is therefore not necessary for hardware to initialize TLB and cache memory in order to process this exception. The vector location of this exception is;

- In 32-bit mode, 0xbfc0 0000
- In 64-bit mode, 0xffff ffff bfc0 0000

All register contents are retained except for the following.

- ErrorEPC register, which contains the restart PC
- ERL, SR and BEV bits of Status register, which are set to “1”

Because Soft-reset exception can abort cache and bus operations, cache and memory state is undefined when this exception occurs.

11.5.3 Servicing

The SoftReset exception is serviced by saving the current processor state for diagnostic purposes, and reinitializing for the ColdReset exception.

11.6 NMI (Non-maskable Interrupt) Exception

11.6.1 Cause

The NMI (Non-maskable Interrupt) exception occurs at the falling edge of the GNMI* signal.

This interrupt is not maskable, and occurs regardless of the EXL, ERL and IE bits of the Status register.

11.6.2 Processing

The same special interrupt vector as for Cold-reset/Soft-reset exception (0xbfc0_0000/ 0xffff_fff_bfc0_0000). This vector is located within unmapped and uncached area so that the cache and TLB need not be initialized to process this exception. When this exception occurs, the SR bit of Status register is set.

Because NMI exception can occur in the midst of another exception, it is not normally possible to continue program execution after servicing NMI exception.

Unlike the Cold-reset/Soft-reset exception, but like other exceptions, this exception occurs at an instruction boundary. The state of the primary cache and memory system are preserved by this exception.

All register contents are retained except for the following.

- ErrorEPC register, which contains the restart PC
- ERL, SR and BEV bits of the Status register, which is set to 1.

11.6.3 Servicing

The NMI exception is serviced by saving the current processor state for diagnostic purposes, and reinitializing the system for the ColdReset exception.

11.7 Address Error Exception

11.7.1 Cause

The Address Error exception occurs when an attempt is made to execute one of the following.

- load or store a doubleword that is not aligned on a doubleword boundary
- load, fetch or store a word that is not aligned on a word boundary
- load or store a halfword that is not aligned on a halfword boundary
- reference Kernel mode address while in User or Supervisor mode
- reference Supervisor mode address while in User mode

This exception is not maskable.

11.7.2 Processing

The common exception vector is used. ExcCode AdEL or AdES in Cause register is set depending on whether the memory access attempt was a load or store. When this exception is raised, the misalign virtual address causing the exception, or the protected virtual address that was illegally referenced, is placed in BadVAddr register. The contents of the VPN field of Context and EntryHi registers are undefined, as are the contents of EntryLo register.

If EXL bit of Status register is only set to 0, the following operation is executed. EPC register points to the address of the instruction causing the exception. If, however, the affected instruction was in the branch delay slot (for execution during a branch), the immediately preceding branch instruction address is retained in EPC register and BD bit of Cause register is set to "1".

11.7.3 Servicing

The process executing at the time is handed a segmentation violation signal. This error is usually fatal to the process incurring the exception.

11.8 TLB Refill Exception

11.8.1 Cause

The TLB refill exception occurs when there is no TLB entry to match a reference to a mapped address. This exception is not maskable.

11.8.2 Processing

There are two special exception vectors for this exception; one for references to 32-bit virtual address, and one for references to 64-bit virtual address. The KX, SX and UX bits of Status register determine whether the User, Supervisor or Kernel address referenced are 32-bit mode or 64-bit mode. When EXL bit of Status register is set to "0", all references use these vectors. When this exception occurs, TLBL or TLBS code is set in the ExcCode field of Cause register. This code indicates whether the instruction, as shown by EPC register and BD bit of Cause register, caused the miss by an instruction reference, load operation, or store operation.

When this exception occurs;

- BadVAddr, Context, XContext and EntryHi registers hold the virtual address failed address translation
- EntryHi register contains ASID from which the translation fault occurred, too
- A valid address in which to place the replacement TLB entry is contained into Random register
- The contents of EntryLo register are undefined

If EXL bit of Status register is only set to 0, the following operation is executed. EPC register points to the address of the instruction causing the exception. If, however, the affected instruction was in the branch delay slot (for execution during a branch), the immediately preceding branch instruction address is retained in EPC register and BD bit of Cause register is set to "1".

11.8.3 Servicing

To service this exception, the contents of the Context or XContext register are used as a virtual address to fetch memory locations containing the physical page frame and access control bits for a pair of TLB entries. The two entries are placed into the EntryLo0/EntryLo1 register; the EntryHi and EntryLo registers are written into the TLB.

It is possible that the virtual address used to obtain the physical address and access control information is on a page that is not resident in the TLB. This condition is processed by allowing a TLB refill exception in the TLB refill handler. This second exception goes to the common exception vector because the EXL bit of the Status register is set.

11.9 TLB Invalid Exception

11.9.1 Cause

The TLB Invalid exception occurs when a virtual address reference matches a TLB entry that is marked invalid (TLB valid bit cleared). This exception is not maskable.

11.9.2 Processing

The common exception vector is used for this exception. When this exception occurs, TLBL or TLBS code is set in the ExcCode field of Cause register. This code indicates whether the instruction, as shown by EPC register and BD bit of Cause register, caused the miss by an instruction reference, load operation, or store operation.

When this exception occurs;

- BadVAddr, Context, XContext and EntryHi registers hold the virtual address failed address translation
- EntryHi register contains ASID from which the translation fault occurred, too
- A valid address in which to place the replacement TLB entry is contained into Random register
- The contents of EntryLo register are undefined

If EXL bit of Status register is only set to 0, the following operation is executed. EPC register points to the address of the instruction causing the exception. If, however, the affected instruction was in the branch delay slot (for execution during a branch), the immediately preceding branch instruction address is retained in EPC register and BD bit of Cause register is set to "1".

11.9.3 Servicing

A TLB entry is typically marked invalid when one of the following is true;

- a virtual address does not exist
- the virtual address exists, but is not in main memory (a page fault)
- a trap is desired on any reference to the page (for example, to maintain a reference bit or during debug)

After servicing the cause of a TLB Invalid exception, the TLB entry is located with TLB Probe (TLBP) instruction, and replaced by an entry with that entry's Valid bit set.

11.10 TLB Modified Exception

11.10.1 Cause

The TLB Modified exception occurs when a store operation virtual address reference to memory matches a TLB entry that is marked valid but is not dirty and therefore is not writable. This exception is not maskable.

11.10.2 Processing

The common exception vector is used for this exception, and Mod code in Cause register is set.

When this exception occurs;

- BadVAddr, Context, XContext and EntryHi registers hold the virtual address failed address translation
- EntryHi register contains ASID from which the translation fault occurred, too
- The contents of EntryLo register are undefined

If EXL bit of Status register is only set to 0, the following operation is executed. EPC register points to the address of the instruction causing the exception. If, however, the affected instruction was in the branch delay slot (for execution during a branch), the immediately preceding branch instruction address is retained in EPC register and BD bit of Cause register is set to 1.

11.10.3 Servicing

The kernel uses the failed virtual address or virtual page number to identify the corresponding access control information. The page identified may or may not permit write accesses; if writes are not permitted, a write protection violation occurs.

If write accesses are permitted, the page frame is marked dirty/writable by the kernel in its own data structures. The TLB Probe (TLBP) instruction places the index of the TLB entry that must be altered into the Index register. The EntryLo register is loaded with a word containing the physical page frame and access control bits (with the D bit set), and the EntryHi and EntryLo registers are written into the TLB.

11.11 Bus Error Exception

11.11.1 Cause

The Bus Error exception occurs when GBUSERR* signal is asserted during a memory read bus cycle. This exception is raised by board-level circuitry for events such as bus time-out, backplane bus parity errors, and invalid physical memory addresses or access types. This occurs during execution of the instruction causing the bus error. The memory bus cycle ends upon notification of a bus error. When a bus error is raised during a burst refill, the following refill is not performed. A bus error request made by asserting GBUSERR* signal will be ignored if TX49 is executing a cycle other than a bus cycle. It is therefore not possible to raise a Bus Error exception in a write access using a write buffer. A general interrupt must be used instead. This exception is not maskable.

11.11.2 Processing

The common interrupt vector is used for a Bus Error exception. The IBE or DBE code in the ExcCode field of the Cause register is set, signifying whether the instruction (as indicated by the EPC register and BD bit in the Cause register) caused the exception by an instruction reference, load operation, or store operation.

The EPC register contains the address of the instruction that caused the exception, unless it is in a branch delay slot, in which case the EPC register contains the address of the preceding branch instruction and the BD bit of the Cause register is set.

11.11.3 Servicing

The physical address at which the fault occurred can be computed from information available in the CPO registers.

- If the IBE code in the Cause register is set (indicating an instruction fetch reference), the virtual address is contained in the EPC register.
- If the DBE code is set (indicating a load or store reference), the instruction that caused the exception is located at the virtual address contained in the EPC register (or 4+ the contents of the EPC register if the BD bit of the Cause register is set).

The virtual address of the load and store reference can then be obtained by interpreting the instruction. The physical address can be obtained by using the TLB Probe (TLBP) instruction and reading the EntryLo register to compute the physical page number.

The process executing at the time of this exception is handed a bus error signal, which is usually fatal.

11.12 Integer Overflow Exception

11.12.1 Cause

The Integer Overflow exception occurs when ADD, ADDI, SUB, DADD, DADDI or DSUB instruction results in a 2's complement overflow. This exception is not maskable.

11.12.2 Processing

The common exception vector is used for this exception, and the Ov code in Cause register is set.

If EXL bit of Status register is only set to 0, the following operation is executed. EPC register points to the address of the instruction causing the exception. If, however, the affected instruction was in the branch delay slot (for execution during a branch), the immediately preceding branch instruction address is retained in EPC register and BD bit of Cause register is set to 1.

11.12.3 Servicing

The process executing at the time of the exception is handed a floating-point exception/integer overflow signal. This error is usually fatal to the current process.

11.13 Trap Exception

11.13.1 Cause

The Trap exception occurs when TGE, TGEU, TLT, TLTU, TEQ, TNE, TGEI, TGEIU, TLTI, TLTIU, TEQI or TNEI instruction results in a TRUE condition. This exception is not maskable.

11.13.2 Processing

The common exception vector is used for this exception, and the Tr code in Cause register is set.

If EXL bit of Status register is only set to 0, the following operation is executed. EPC register points to the address of the instruction causing the exception. If, however, the affected instruction was in the branch delay slot (for execution during a branch), the immediately preceding branch instruction address is retained in EPC register and BD bit of Cause register is set to 1.

11.13.3 Servicing

The process executing at the time of a Trap exception is handed a floating-point exception/integer overflow signal. This error is usually fatal.

11.14 System Call Exception

11.14.1 Cause

The System Call exception occurs during an attempt to execute the SYSCALL instruction. This exception is not maskable.

11.14.2 Processing

The common exception vector is used for this exception, and the Sys code in Cause register is set.

If EXL bit of Status register is only set to 0, the following operation is executed. EPC register points to the address of the SYSCALL instruction. If, however, the affected instruction was in the branch delay slot (for execution during a branch), the immediately preceding branch instruction address is retained in EPC register.

If the SYSCALL instruction is in a branch delay slot, BD bit of Status register is set, otherwise this bit is cleared.

11.14.3 Servicing

When this exception occurs, control is transferred to the applicable system routine.

To resume execution, the EPC register must be altered so that the SYSCALL instruction does not re-execute; this is accomplished by adding a value of 4 to the EPC register (EPC register $\square + \square 4$) before returning.

If a SYSCALL instruction is in a branch delay slot, a more complicated algorithm, beyond the scope of this description, may be required.

11.15 Breakpoint Exception

11.15.1 Cause

The Breakpoint exception occurs when an attempt is made to execute the BREAK instruction. This exception is not maskable.

11.15.2 Processing

The common exception vector is used for this exception, and the Bp code in Cause register is set.

If EXL bit of Status register is only set to 0, the following operation is executed. EPC register points to the address of the BREAK instruction. If, however, the affected instruction was in the branch delay slot (for execution during a branch), the immediately preceding branch instruction address is retained in EPC register.

If the BREAK instruction is in a branch delay slot, BD bit of Status register is set, otherwise this bit is cleared.

11.15.3 Servicing

When the Breakpoint exception occurs, control is transferred to the applicable system routine. Additional distinctions can be made by analyzing the unused bits of the BREAK instruction (bits 25~6), and loading the contents of the instruction whose address the EPC register contains. A value of 4 must be added to the contents of the EPC register (EPC register + 4) to locate the instruction if it resides in a branch delay slot.

To resume execution, the EPC register must be altered so that the BREAK instruction does not re-execute; this is accomplished by adding a value of 4 to the EPC register (EPC register + 4) before returning.

If a BREAK instruction is in a branch delay slot, interpretation of the branch instruction is required to resume execution.

11.16 Reserved Instruction Exception

11.16.1 Cause

The Reserved Instruction exception occurs when one of the following condition occurs:

- an attempt is made to execute an instruction with an undefined major opcode (bit 31~26)
- an attempt is made to execute a SPECIAL instruction with an undefined minor opcode (bit 5~0)
- an attempt is made to execute a REGIMM instruction with an undefined minor opcode (bit20~16)
- an attempt is made to execute 64-bit operations in 32-bit mode when in User or Supervisor modes
- an attempt is made to execute a COPz rs instruction with an undefined minor opcode (bit25~21)
- an attempt is made to execute a COPz rt instruction with an undefined minor opcode (bit20~16)
- 64-bit operations are always valid in Kernel mode regardless of the value of the KX bit in Status register. This exception is not maskable.

11.16.2 Processing

The common exception vector is used for this exception, and the RI code in Cause register is set.

If EXL bit of Status register is only set to 0, the following operation is executed. EPC register points to the address of the instruction causing the exception. If, however, the affected instruction was in the branch delay slot (for execution during a branch), the immediately preceding branch instruction address is retained in EPC register.

11.16.3 Servicing

No instruction in the MIPS ISA are currently interpreted. The process executing at the time of this exception is handed an illegal instruction/reserved operand fault signal. This error is usually fatal.

11.17 Coprocessor Unusable Exception

11.17.1 Cause

The Coprocessor Unusable exception occurs when an attempt is made to execute a coprocessor instruction for either

- attempting to execute a coprocessor CPz instruction when its corresponding CUz bit in Status register.
- in User or Supervisor mode attempting to execute a CP0 instruction when CU0 bit is cleared to “0”. (In Kernel mode, an exception is not raised when a CP0 instruction is issued , regardless of the CU0 bit setting)
- an attempt is made to execute a FPU instruction in TX49 without FPU

11.17.2 Processing

The common exception vector is used for this exception, and the CpU code in Cause register is set.

The coprocessor number referred to at the time of the exception is stored in Cause register CE (Coprocessor Error) field.

If EXL bit of Status register is only set to 0, the following operation is executed. EPC register points to the address of the instruction causing the exception. If, however, the affected instruction was in the branch delay slot (for execution during a branch), the immediately preceding branch instruction address is retained in EPC register and BD bit of Cause register is set to 1.

11.17.3 Servicing

The coprocessor unit to which an attempted reference was made is identified by the Coprocessor Usage Error field, which results in one of the following situations:

- If the process is entitled access to the coprocessor, the coprocessor is marked usable and the corresponding user state is restored to the coprocessor.
- If the process is entitled access to the coprocessor, but the coprocessor does not exist or has failed, interpretation of the coprocessor instruction is possible.
- If the BD bit is set in the Cause register, the branch instruction must be interpreted; then the coprocessor instruction can be emulated and execution resumed with the EPC register advanced past the coprocessor instruction.
- If the process is not entitled access to the coprocessor, the process executing at the time is handed an illegal instruction/privileged instruction fault signal. This error is usually fatal.

11.18 Floating-Point Exception

11.18.1 Cause

The Floating-Point exception is used by the floating-point coprocessor. This exception is not maskable.

11.18.2 Processing

The common exception vector is used for this exception, and the FPE code in Cause register is set.

The contents of the Floating-Point Control/Status register indicate the cause of this exception.

11.18.3 Servicing

This exception is cleared by clearing the appropriate bit in the Floating-Point Control/Status register.

For an unimplemented instruction exception, the kernel should emulate the instruction; for other exceptions, the kernel should pass the exception to the user program that caused the exception.

11.19 Interrupt Exception

11.19.1 Cause

The Interrupt exception is raised by any of eight interrupts (two software and six hardware). A hardware interrupt is raised when GINT* signal goes active. A software interrupt is raised by setting the IP[1]/IP[0] bit in Cause register. The significance of these interrupts is dependent upon the specific system implementation.

Each of the eight interrupts can be masked individually by clearing its corresponding bit in the IM(Interrupt Mask) field of Status register, and all interrupts can be masked at once by clearing IE bit of Status register to "0".

11.19.2 Processing

The common exception vector is used as following;

- In 32 bit mode, 0x8000 0180 (BEV \square = \square 0)
0xbfc0 0380 (BEV \square = \square 1)
- In 64 bit mode, 0xffff ffff 8000 0180 (BEV \square = \square 0)
0xffff ffff bfc0 0380 (BEV \square = \square 1)

11.19.3 Servicing

If the interrupt is caused by one of the two software-generated exceptions (SW1 or SW0), the interrupt condition is cleared by setting the corresponding Cause register bit to 0.

If the interrupt is hardware-generated, the interrupt condition is cleared by correcting the condition causing the interrupt pin to be asserted.

If the timer interrupt is caused, the interrupt condition is cleared by changing the value of the Compare register or setting the corresponding Cause register bit (IP[7]) to 0.

NOTE: due to the write buffer, a store to an external device will not necessary occur until after other instructions in the pipeline finish. Thus, the user must ensure that the store will occur before the return from exception instruction (ERET) is executed otherwise the interrupt may be serviced again even though there should be no interrupt pending.

11.20 Exception Handling and Servicing Flowcharts

The remainder of this chapter contains flowcharts for the following exceptions and guidelines for their handlers:

- general exceptions and their exception handler
- TLB/XTLB miss exception and their exception handler
- Cold Reset, Soft Reset and NMI exceptions, and a guideline to their handler.

Generally speaking, the exceptions are handled by hardware (HW); the exceptions are then serviced by software (SW).

Exceptions other than Reset, Soft Reset, NMI or first-level miss

Note: Interrupts can be masked by IE or IMs

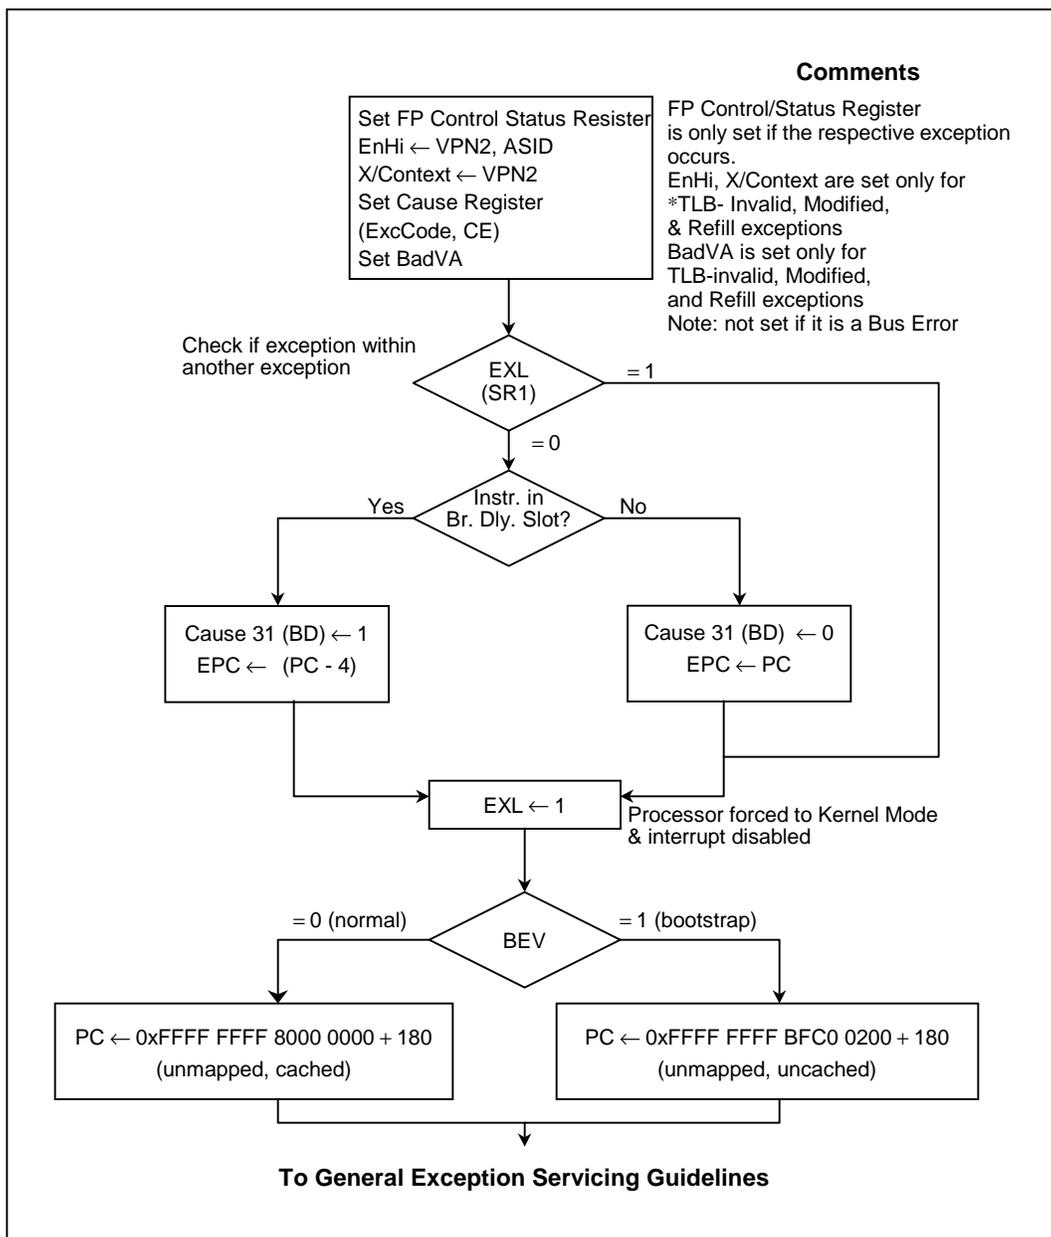


Figure 11-1 General Exception Handler (HW)

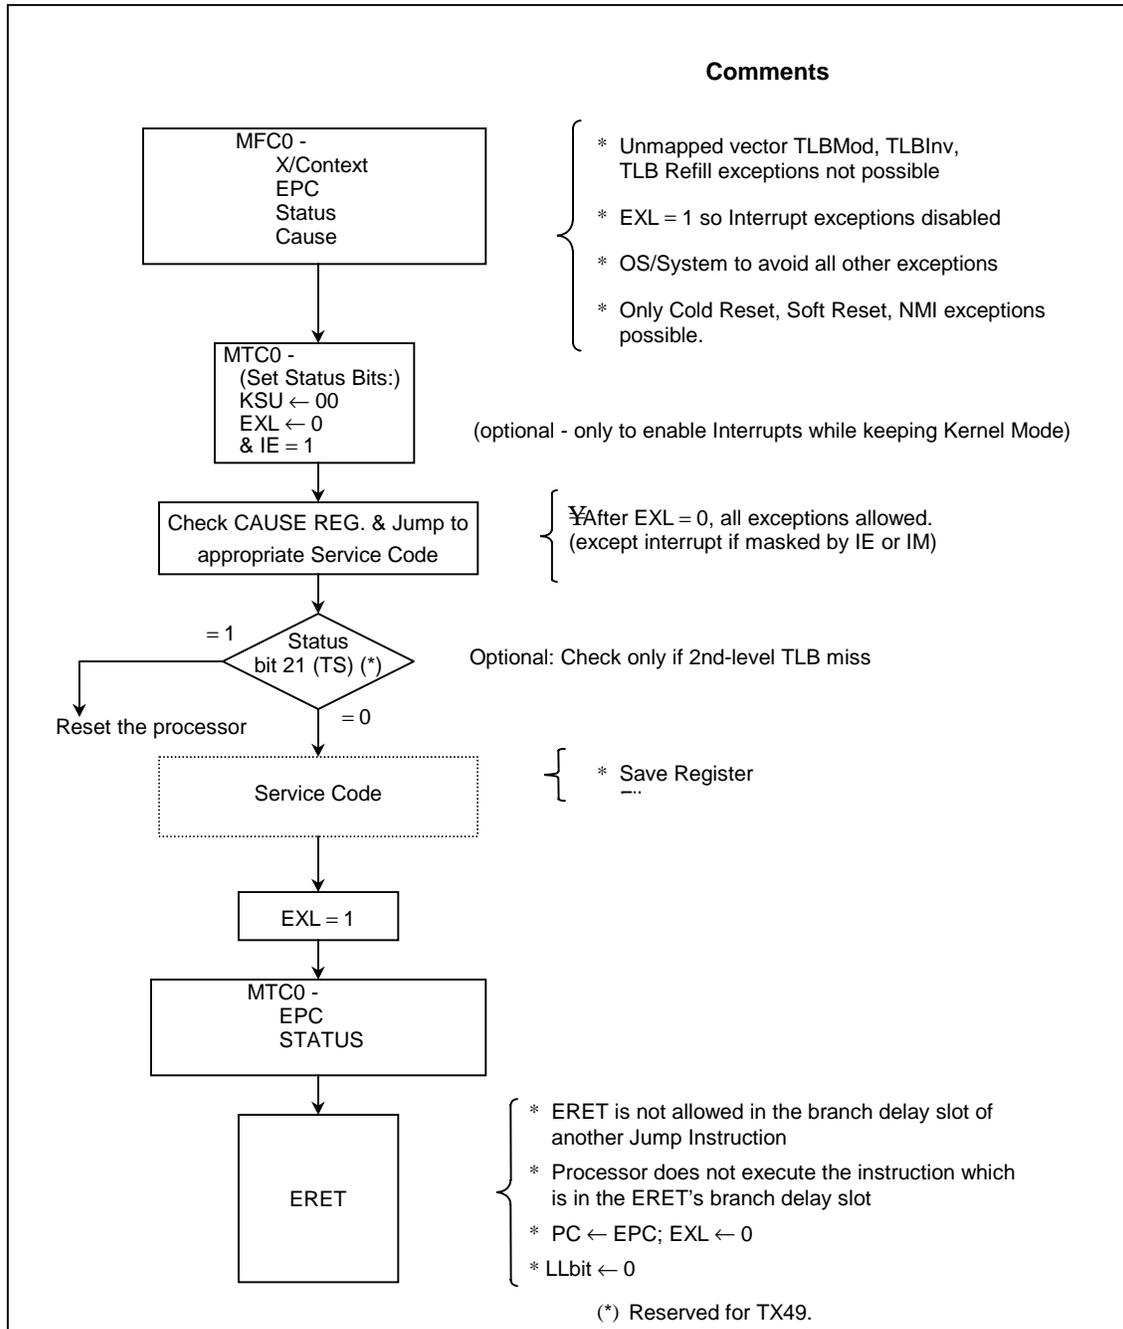


Figure 11-2 General Exception Servicing Guidelines (SW)

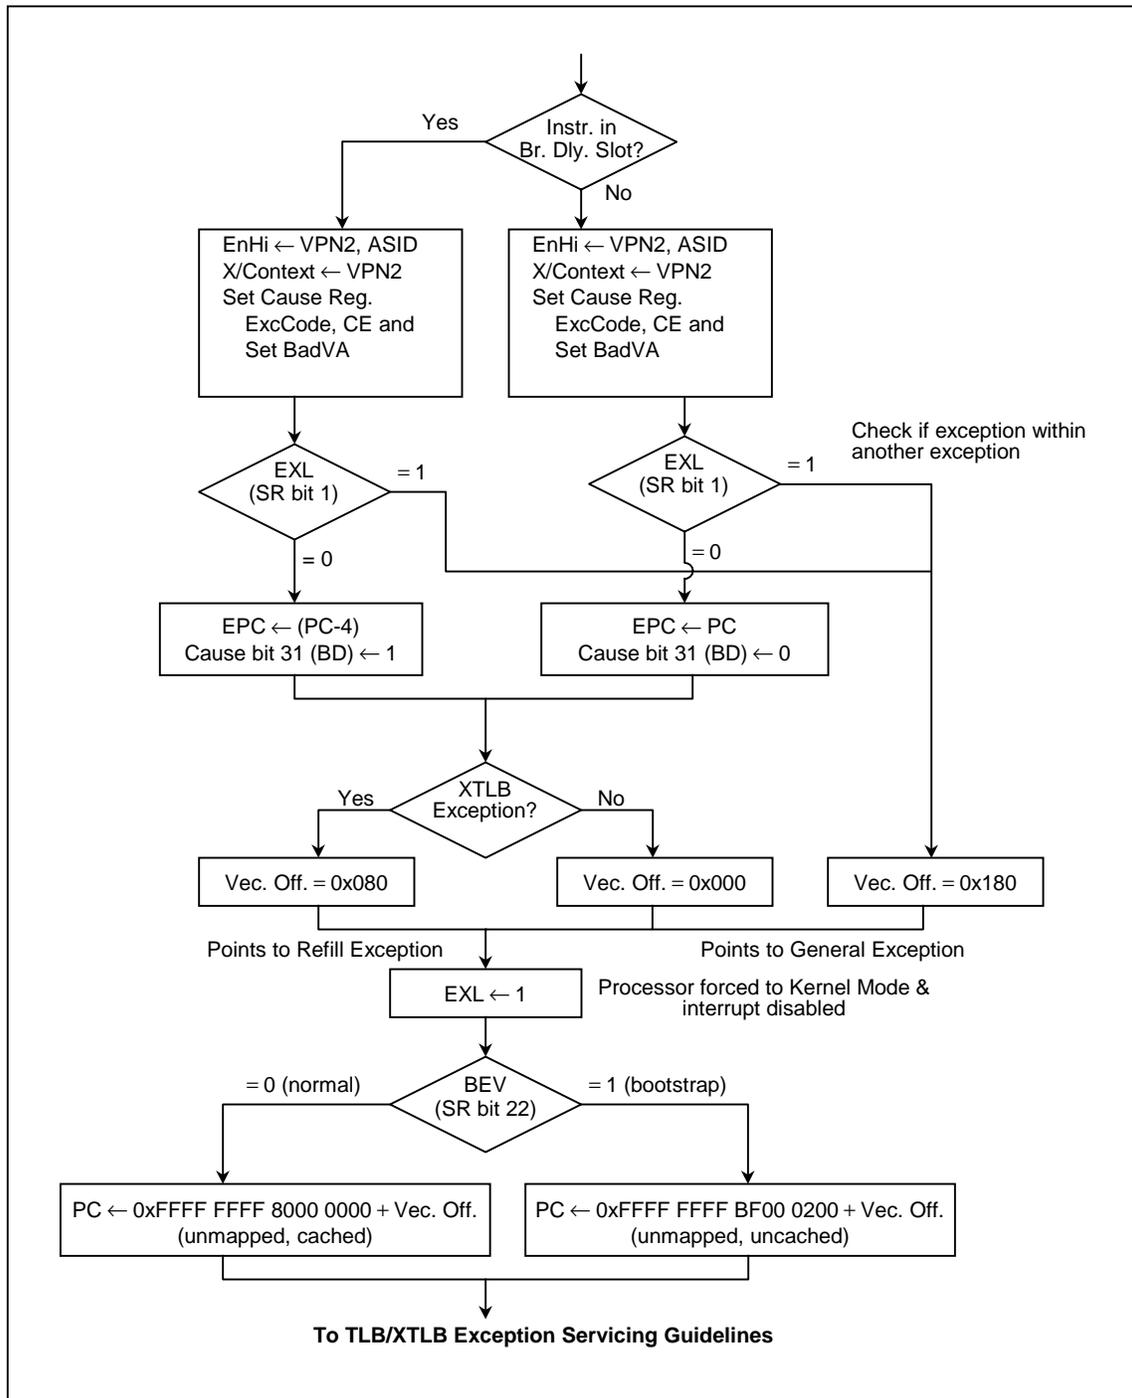


Figure 11-3 TLB/XTLB Miss Exception Handler (HW)

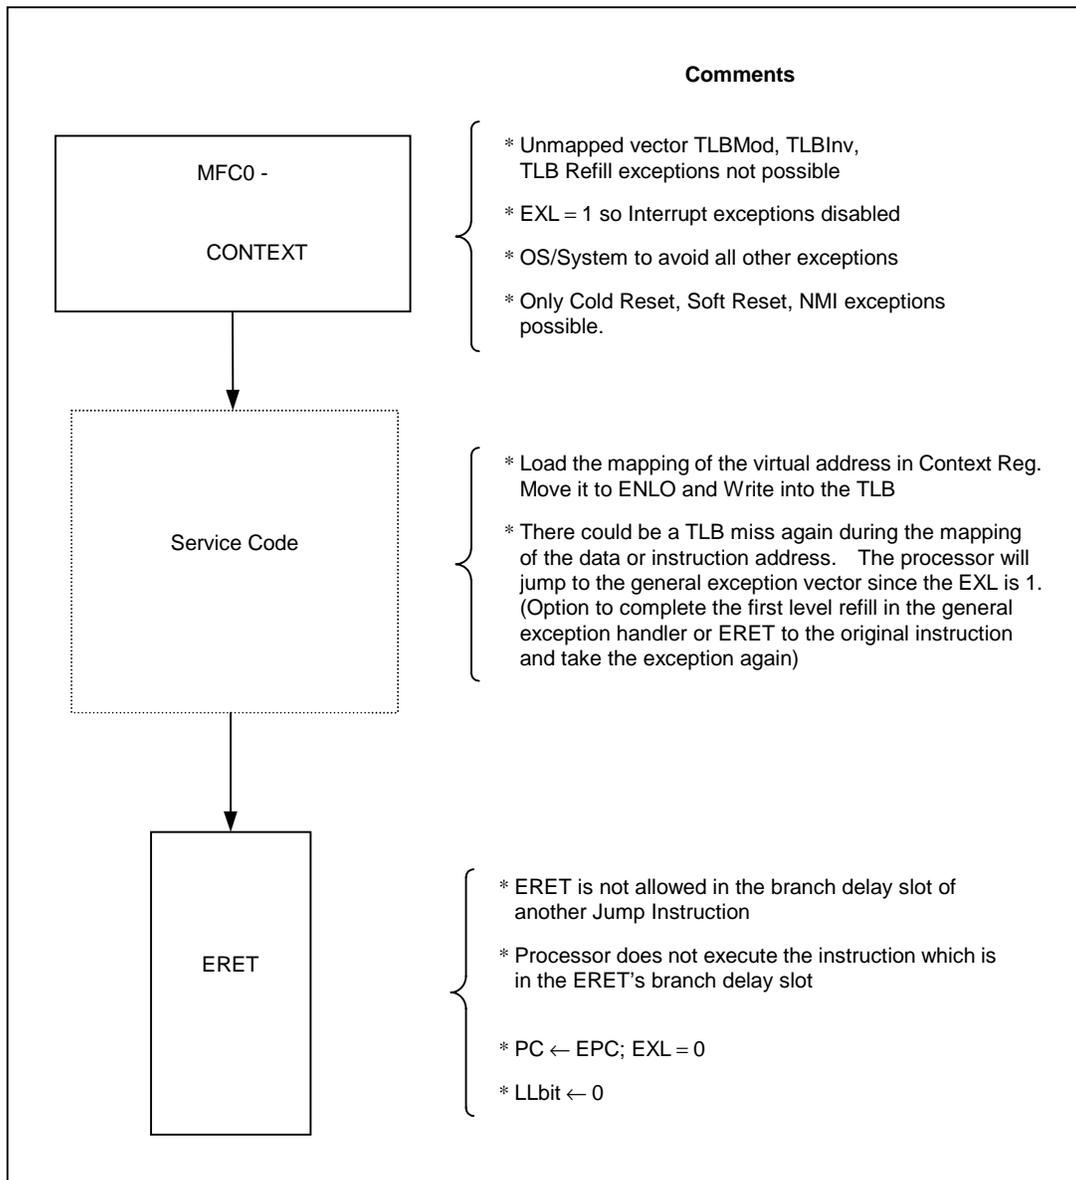


Figure 11-4 TLB/XTLB Exception Servicing Guidelines (SW)

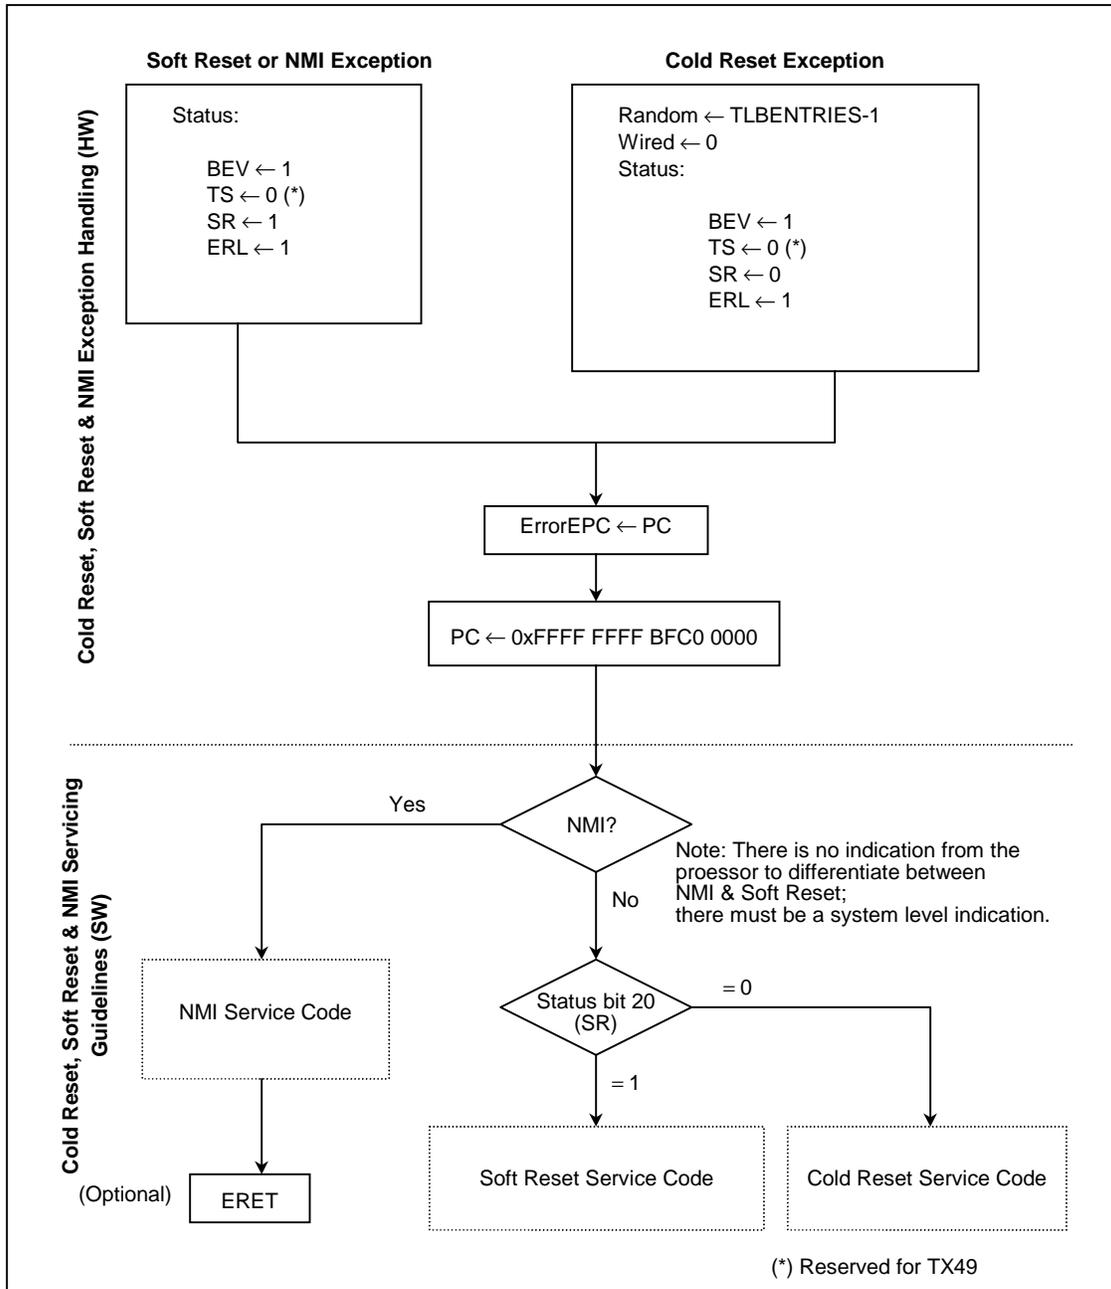


Figure 11-5 Cold Reset, Soft Reset & NMI Exception Handling (HW) and Servicing Guidelines (SW)

12. Floating-Point Unit, CP1 (Option)

This chapter describes the floating-point operations, including the programming model, instruction set and formats.

The floating-point operations fully conform to the requirements of ANSI/IEEE Standard 754-1985, *IEEE Standard for Binary Floating-Point Arithmetic*.

12.1 Overview

All floating-point instructions, as defined in the MIPS ISA for the floating-point coprocessor, CP1, are processed by the other hardware unit that executes integer instructions. Logically, the FPU exists as an individual coprocessor; however, unlike previous implementations, the TX49 FPU is physically integrated into the CPU.

The execution of floating-point instructions can be disabled by the coprocessor usability *CU* bit defined in the CP0 *Status* register.

12.2 Floating Point Register

12.2.1 Floating-Point General Registers (FGRs)

CP1 has a set of *Floating-Point General Purpose registers (FGRs)* that can be accessed in the following ways:

- As 32 general purpose registers (32 FGRs), each of which is 32 bits wide when the *FR* bit in the *CPU Status* register equals 0; or as 32 general purpose registers (32 FGRs), each of which is 64-bits wide when *FR* equals 1. The CPU accesses these registers through move, load, and store instructions.
- As 16 floating-point registers (see the next section for a description of FPRs), each of which is 64-bits wide, when the *FR* bit in the *CPU Status* register equals 0. The FPRs hold values in either single- or double-precision floating-point format. Each FPR corresponds to adjacently numbered FGRs as shown in Figure 12-1.
- As 32 floating-point registers (see the next section for a description of FPRs), each of which is 64-bits wide, when the *FR* bit in the *CPU Status* register equals 1. The FPRs hold values in either single- or double-precision floating-point format. Each FPR corresponds to an FGR as shown in Figure 12-1.

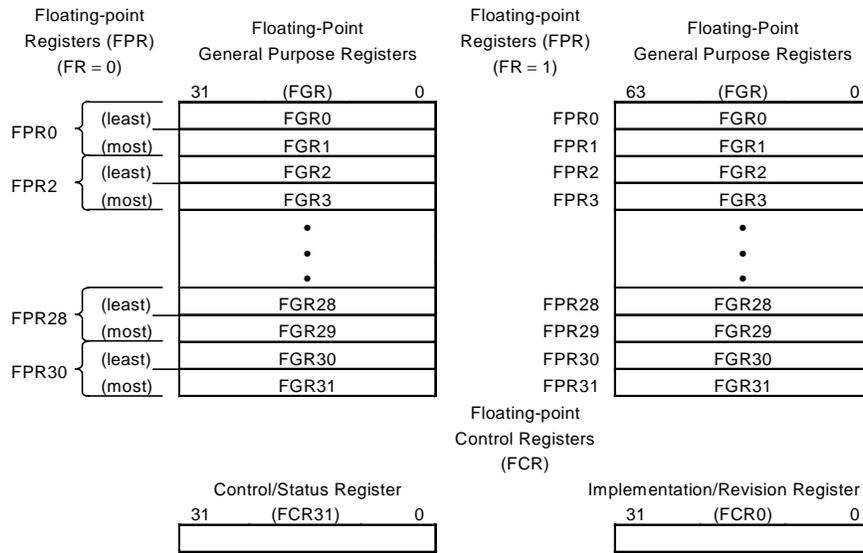


Figure 12-1 FP Registers

12.2.2 Floating-Point Control Registers

The MIPS RISC architecture defines 32 floating-point control registers (*FCRs*); the TX49 processor implements two of these registers: *FCR0* and *FCR31*. These *FCRs* are described below:

- The *Implementation/Revision* register (*FCR0*) holds revision information.
- The *Control/Status* register (*FCR31*) controls and monitors exceptions, holds the result of compare operations, and establishes rounding modes.
- *FCR1* to *FCR30* are reserved.

Table 12-1 lists the assignments of the *FCRs*.

Table 12-1 Floating-Point Control Register Assignments

FCR Number	Use
FCR0	Coprocessor implementation and revision register
FCR1 to FCR30	Reserved
FCR31	Rounding mode, cause, trap enables, and flags

Implementation and Revision Register, (*FCR0*)

The read-only *Implementation and Revision* register (*FCR0*) specifies the implementation and revision number of CP1. This information can determine the coprocessor revision and performance level, and can also be used by diagnostic software.

Figure 12-1 shows the layout of the register; Table 12-2 describes the *Implementation and Revision* register (*FCR0*) fields.

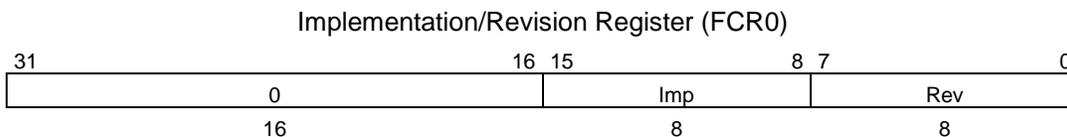


Figure 12-1 Implementation/Revision Register

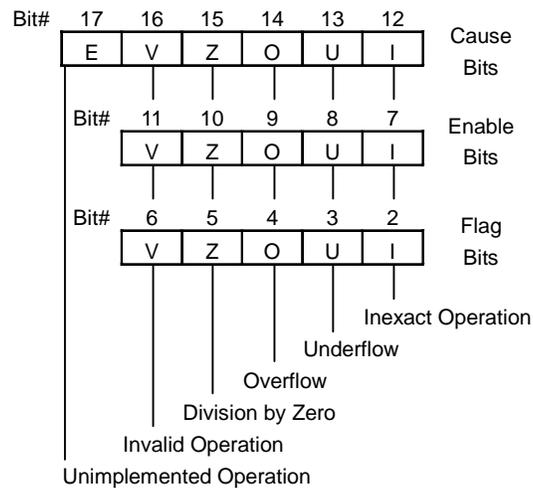


Figure 12-3 Control/Status Register Cause, Flag, and Enable Fields

Control/Status Register FS Bit

The *FS* bit enables the flushing of denormalized values. When the *FS* bit is set and the Underflow and Inexact *Enable* bits are not set, denormalized results are flushed instead of causing an Unimplemented Operation exception. Results are flushed either to 0 or the minimum normalized value, depending upon the rounding mode (see Table 12-4 below), and the Underflow and Inexact *Flag* and *Cause* bits are set.

Table 12-4 Flush Values of Denormalized Results

Denormalized Result	Flushed Result Rounding Mode			
	RN	RZ	RP	RM
Positive	+0	+0	$+2^{E_{min}}$	+0
Negative	-0	-0	-0	$-2^{E_{min}}$

Control/Status Register Condition Bit

When a floating-point Compare operation takes place, the result is stored at bit 23, the *Condition* bit. The *C* bit is set to 1 if the condition is true; the bit is cleared to 0 if the condition is false. Bit 23 is affected only by compare and CTC1 instructions.

Control/Status Register Cause, Flag, and Enable Fields

Figure 12-3 illustrates the *Cause*, *Flag*, and *Enable* fields of the *Control/Status* register. The *Cause* and *Flag* fields are updated by all conversion, computational (except MOV. fmt), CTC1, reserved, and unimplemented instructions. All other instructions have no affect on these fields.

Cause Bits

Bits 17:12 in the *Control/Status* register contain *Cause* bits, as shown in Figure 12-3, which reflect the results of the most recently executed floating-point instruction. The *Cause* bits are a logical extension of the CP0 *Cause* register; they identify the exceptions raised by the last floating-point operation. If the corresponding *Enable* bit is set at the time of the exception a floating-point exception and interrupt is raised. If more than one exception occurs on a single instruction, each appropriate bit is set.

The *Cause* bits are updated by most floating-point operations. The Unimplemented Operation (*E*) bit is set to 1 if software emulation is required, otherwise it remains 0. The other bits are set to 0 or 1 to indicate the occurrence or non-occurrence (respectively) of an IEEE 754 exception. Within the set of floating-point instructions that update the *Cause* bits, the *Cause* field indicates the exceptions raised by the most-recently-executed instruction.

When a floating-point exception is taken, no results are stored, and the only state affected is the *Cause* bit.

Enable Bits

A floating-point exception is generated any time a *Cause* bit and the corresponding *Enable* bit are set. A floating-point operation that sets an enabled *Cause* bit forces an immediate floating-point exception, as does setting both *Cause* and *Enable* bits with CTC1.

There is no enable for Unimplemented Operation (*E*). An Unimplemented exception always generates a floating-point exception.

Before returning from a floating-point exception, software must first clear the enabled *Cause* bits with a CTC1 instruction to prevent a repeat of the interrupt. Thus, User mode programs can never observe enabled *Cause* bits set; if this information is required in a User mode handler, it must be passed somewhere other than the *Status* register.

For a floating-point operation that sets only unenabled *Cause* bits, no floating-point exception occurs and the default result defined by IEEE 754 is stored. In this case, the exceptions that were caused by the immediately previous floating-point operation can be determined by reading the *Cause* field.

Flag Bits

The *Flag* bits are cumulative and indicate the exceptions that were raised by the operations that were executed since the bits were explicitly reset. *Flag* bits are set to 1 if an IEEE 754 exception is raised, otherwise they remain unchanged. The *Flag* bits are never cleared as a side effect of floating-point operations; however, they can be set or cleared by writing a new value into the *Status* register, using a CTC1 instruction.

When a floating-point exception is taken, the flag bits are not set by the hardware; floating-point exception software is responsible for setting these bits before invoking a user handler.

Control/Status Register Rounding Mode Control Bits

Bits 1 and 0 in the *Control/Status* register constitute the *Rounding Mode (RM)* field.

As shown in Table 12-5, these bits specify the rounding mode that CP1 uses for all floating-point operations.

Table 12-5 Rounding Mode Bit Decoding

Rounding Mode RM (1:0)	Mnemonic	Description
0	RN	Round result to nearest representable value; round to value with least-significant bit 0 when the two nearest representable values are equally near.
1	RZ	Round toward 0: round to value closest to and not greater in magnitude than the infinitely precise result.
2	RP	Round toward $+\infty$: round to value closest to and not less than the infinitely precise result.
3	RM	Round toward $-\infty$: round to value closest to and not greater than the infinitely precise result.

12.2.3 Accessing the FP Control and Implementation/Revision Registers

The *Control/Status* and the *Implementation/Revision* registers are read by a Move Control From Coprocessor 1 (CFC1) instruction.

The bits in the *Control/Status* register can be set or cleared by writing to the register using a Move Control To Coprocessor 1 (CTC1) instruction. The *Implementation/Revision* register is a read-only register. There are no pipeline hazards (between any instructions) associated with floating-point control registers.

12.3 Floating-Point Formats

CP1 performs both 32-bit (single-precision) and 64-bit (double-precision) IEEE standard floating-point operations. The 32-bit single-precision format has a 24-bit signed-magnitude fraction field ($f+s$) and an 8-bit exponent (e), as shown in Figure 12-1.

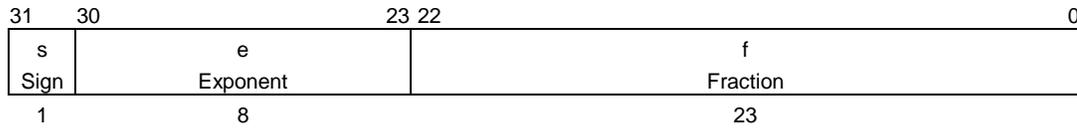


Figure 12-1 Single-Precision Floating-Point Format

The 64-bit double-precision format has a 53-bit signed-magnitude fraction field ($f+s$) and an 11-bit exponent, as shown in Figure 12-2.

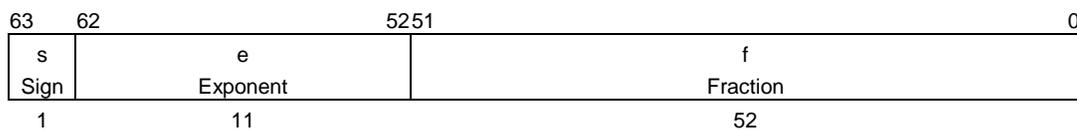


Figure 12-2 Double-Precision Floating-Point Format

As shown in the above figures, numbers in floating-point format are composed of three fields:

- sign field, s
- biased exponent, $e = E + bias$
- fraction, $f = b_1b_2\dots b_{p-1}$

The range of the unbiased exponent E includes every integer between the two values E_{min} and E_{max} inclusive, together with two other reserved values:

- $E_{min} - 1$ (to encode 0 and denormalized numbers)
- $E_{max} + 1$ (to encode ∞ and NaNs [Not a Number])

For single- and double-precision formats, each representable nonzero numerical value has just one encoding.

For single- and double-precision formats, the value of a number, v , is determined by the equations shown in Table 12-1.

Table 12-1 Equations for Calculating Values in Single and Double-Precision Floating-Point Format

No.	Equation
(1)	if $E = E_{max} + 1$ and $f \neq 0$, then v is NaN, regardless of s
(2)	if $E = E_{max} + 1$ and $f = 0$, then $v = (-1)^s \infty$
(3)	if $E_{min} \leq E \leq E_{max}$, then $v = (-1)^s 2^E (1.f)$
(4)	if $E = E_{min} - 1$ and $f \neq 0$, then $v = (-1)^s 2^{E_{min}} (0.f)$
(5)	if $E = E_{min} - 1$ and $f = 0$, then $v = (-1)^s 0$

For all floating-point formats, if v is NaN, the most-significant bit of f determines whether the value is a signaling or quiet NaN: v is a signaling NaN if the most-significant bit of f is set, otherwise, v is a quiet NaN.

Table 12-2 defines the values for the format parameters; minimum and maximum floating-point values are given in Table 12-3.

Table 12-2 Floating-Point Format Parameter Values

Parameter	Format	
	Single	Double
E_{\max}	+127	+1023
E_{\min}	126	1022
Exponent <i>bias</i>	127	+1023
Exponent width in bits	8	11
Integer bit	hidden	hidden
Fraction width in bits	23†	52†
Format width in bits	32	64

† Excluding the sign bit.

Table 12-3 Minimum and Maximum Floating-Point Values

Type	Value
Float Minimum	$1.40129846e^{-45}$
Float Minimum Norm	$1.17549435e^{-38}$
Float Maximum	$3.40282347e^{+38}$
Double Minimum	$4.9406564584124654e^{-324}$
Double Minimum Norm	$2.2250738585072014e^{-308}$
Double Maximum	$1.7976931348623157e^{+308}$

12.4 Binary Fixed-Point Format

Binary fixed-point values are held in 2's complement format. Unsigned fixed-point values are not directly provided by the floating-point instruction set. Figure 12-1 illustrates binary single fixed-point format and Figure 12-2 illustrates binary long fixed-point format; Table 12-1 lists the binary fixed-point format fields.



Figure 12-1 Binary Single Fixed-Point Format

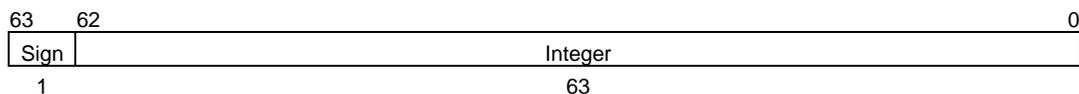


Figure 12-2 Binary Long Fixed-Point Format

Field assignments of the binary fixed-point format are:

Table 12-1 Binary Fixed-Point Format Fields

Field	Description
sign	sign bit
integer	integer value (2's complement)

12.5 Floating-Point Instruction Set Summary

Each instruction is 32 bits long, and aligned on a word boundary. This section describes the overview of instructions for floating-point unit. A detailed description of each instruction is provided in Appendix B.

12.5.1 Load, Move and Store Instructions (Table 12-1)

Load and Store instructions move data between memory and FPU general purpose registers, and Move instructions move data directly between CPU and FPU general purpose registers. These instructions are not perform format conversions and therefore never cause floating-point exceptions. The instruction immediately following a load can use the contents of the loaded register. However, in such case the hardware interlocks, requiring additional real cycles. Thus, the scheduling of load delay slots is required to avoid the interlocking.

Table 12-1 FPU Instruction Set (Optional): Load, Move and Store Instruction

Instruction	Description	Note
LWC1	Load Word to FPU (coprocessor 1)	MIPS I
SWC1	Store Word from FPU (coprocessor 1)	MIPS I
MTC1	Move Word to FPU (coprocessor 1)	MIPS I
CTC1	Move Control Word to FPU (coprocessor 1)	MIPS I
MFC1	Move Word from FPU (coprocessor 1)	MIPS I
CFC1	Move Control Word from FPU (coprocessor 1)	MIPS I

12.5.2 Conversion Instructions (Table 12-1)

Conversion instructions perform conversion operations between the various data formats.

Table 12-1 FPU Instruction Set(Optional): Conversion Instruction

Instruction	Description	Note
CVT.S.fmt	Floating-Point Convert to Single FP Format	MIPS I
CVT.W.fmt	Floating-Point Convert to Single Fixed-Point Format	MIPS I
ROUND.W.fmt	Floating-point Round	MIPS II
TRUNC.W.fmt	Floating-point Truncate	MIPS II
CEIL.W.fmt	Floating-point Ceiling	MIPS II
FLOOR.W.fmt	Floating-point Floor	MIPS II

12.5.3 Computational Instructions (Table 12-1)

Computational instructions perform arithmetic operations on floating-point values in the FPU registers. These are two categories of computational instructions:

- 3-Operand Register-Type instructions, which perform floating-point addition, multiplication, division, and square root operations
- 2-Operand

Table 12-1 FPU Instruction Set(Optional): Computational Instruction

Instruction	Description	Note
ADD.fmt	Floating-point Add	MIPS I
SUB.fmt	Floating-point Subtract	MIPS I
MUL.fmt	Floating-point Multiply	MIPS I
DIV.fmt	Floating-point Divide	MIPS I
ABS.fmt	Floating-point Absolute Value	MIPS I
MOV.fmt	Floating-point Move	MIPS I
NEG.fmt	Floating-point Negate	MIPS I
SQRT.fmt	Floating-point Square root	MIPS II

12.5.4 Compare and Branch Instructions (Table 12-1)

Compare instructions perform comparisons of the contents of registers and set a conditional bit based on the results. Branch on FPU Condition instructions perform a branch to the specified target if the specified coprocessor condition is met.

Table 12-1 FPU Instruction Set(Optional): Compare and Branch Instruction

Instruction	Description	Note
C.cond.fmt	Floating-point Compare	MIPS I
BC1T	Branch on FPU True	MIPS I
BC1F	Branch on FPU False	MIPS I
BC1TL	Branch on FPU True Likely	MIPS II
BC1FL	Branch on FPU False Likely	MIPS II

13. Floating-Point Exception(Option)

13.1 Introduction

This chapter describes floating-point exceptions, including FPU exception type, exception trap processing, exception flags, saving and restoring state when handling an exception, and trap handlers for IEEE Standard 754 exceptions.

13.2 Exception Types

The FP Control/Status register described in Chapter 12 contains an Enable bit for each exception type; exception Enable bits determine whether an exception will cause the FPU to initiate a trap or set a status flag.

- If a trap is taken, the FPU remains in the state found at the beginning of the operation and a software exception handling routine executes.
- If no trap is taken, an appropriate value is written into the FPU destination register and execution continues.

The FPU supports the five IEEE Standard 754 exceptions:

- Inexact (I)
- Underflow (U)
- Overflow (O)
- Division by Zero (Z)
- Invalid Operation (V)

Cause bits, Enables, and Flag bits (status flags) are used.

The FPU adds a sixth exception type, Unimplemented Operation (E). This exception indicates the use of a software implementation. The Unimplemented Operation exception has no Enable or Flag bit; whenever this exception occurs, an unimplemented exception trap is taken.

Figure 13-1 shows the Control/Status register bits that support exceptions.

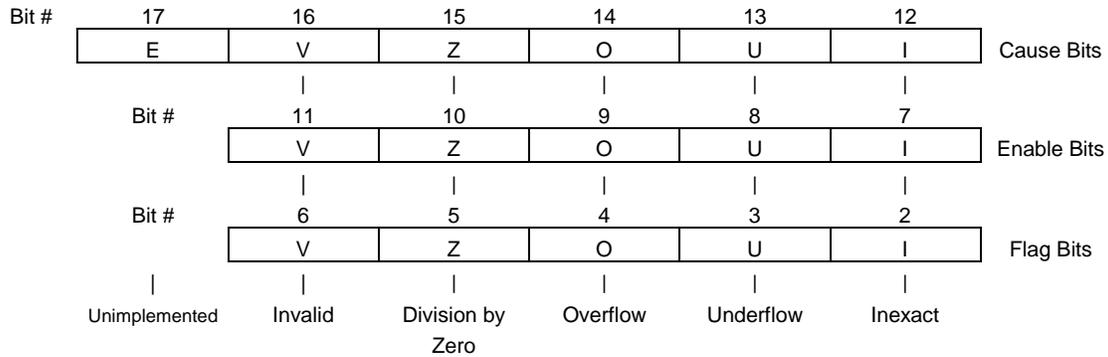


Figure 13-1 Control/Status Register Exception/Flag/Trap/Enable Bits

13.3 Exception Trap Processing

When a floating-point exception trap is taken, the Cause register indicates the floating-point coprocessor is the cause of the exception trap.

The Floating-Point Exception (FPE) code is used, and the Cause bits of the floating-point Control/Status register indicate the reason for the floating-point exception. These bits are, in effect, an extension of the system coprocessor Cause register.

13.4 Flags

A Flag bit is provided for each IEEE exception. This Flag bit is set to a 1 on the assertion of its corresponding exception, with no corresponding exception trap signaled.

13.5 FPU Exceptions

The following sections describe the conditions that cause the FPU to generate each of its exceptions, and details the FPU response to each exception-causing condition.

Inexact Exception (I)

The FPU generates the Inexact exception if one of the following occurs:

- the rounded result of an operation is not exact, or
- the rounded result of an operation overflows, or
- the rounded result of an operation underflows and both the Underflow and Inexact Enable bits are not set and the *FS* bit is set.

Trap Enabled Results: If Inexact exception traps are enabled, the result register is not modified and the source registers are preserved.

Trap Disabled Results: The rounded or overflowed result is delivered to the destination register if no other software trap occurs.

Invalid Operation Exception (V)

The Invalid Operation exception is signaled if one or both of the operands are invalid for an implemented operation. When the exception occurs without a trap, the MIPS ISA defines the result as a quiet Not a Number (qNaN). The invalid operations are:

- Addition or subtraction: magnitude subtraction of infinities, such as: $(\square + \square \infty) - \square + \square (-\infty)$ or $(-\infty) - \square - \square (-\infty)$
- Multiplication: 0 times ∞ , with any signs
- Division: 0/0, or ∞/∞ , with any signs
- Comparison of predicates involving '<' or '>' without '?', when the operands are unordered
- Any arithmetic operation, when one or both operands is a signaling NaN. A move (MOV) operation is not considered to be an arithmetic operation, but absolute value (ABS) and negate (NEG) are.
- Comparison or a Convert From Floating-point Operation on a signaling NaN.
- Square root: \sqrt{x} , where x is less than zero.

Software can simulate the Invalid Operation exception for other operations that are invalid for the given source operands. Examples of these operations include IEEE Standard 754-specified functions implemented in software, such as Remainder: $x \text{ REM } y$, where y is 0 or x is infinite; conversion of a floating-point number to a decimal format whose value causes an overflow, is infinity, or is NaN; and transcendental functions, such as $\ln(-5)$ or $\cos^{-1}(3)$. Refer to Appendix B for examples or for routines to handle these cases.

Trap Enabled Results: The result register is not modified, and the source registers are preserved.

Trap Disabled Results: A quiet NaN is delivered to the destination register if no other software trap occurs.

Divide-by-Zero Exception (Z)

The Division-by-Zero exception is signaled on an implemented divide operation if the divisor is zero and the dividend is a finite nonzero number. Software can simulate this exception for other operations that produce a signed infinity, such as $\ln(0)$, $\sec(\pi/2)$, $\csc(0)$, or 0^{-1} .

Trap Enabled Results: The result register is not modified, and the source registers are preserved.

Trap Disabled Results: The result, when no trap occurs, is a correctly signed infinity.

Overflow Exception (O)

The Overflow exception is signaled when the magnitude of the rounded floating-point result, with an unbounded exponent range, is larger than the largest finite number of the destination format. (This exception also signals an Inexact exception.)

Trap Enabled Results: The result register is not modified, and the source registers are preserved.

Trap Disabled Results: The result, when no trap occurs, is determined by the rounding mode and the sign of the intermediate result (as listed in Table 12-1).

Underflow Exception (U)

Two related events contribute to the Underflow exception:

- creation of a tiny nonzero result between $\pm 2^{E_{\min}}$ which can cause some later exception because it is so tiny
- extraordinary loss of accuracy during the approximation of such tiny numbers by denormalized numbers.

IEEE Standard 754 allows a variety of ways to detect these events, but requires they be detected the same way for all operations.

Tinness can be detected by one of the following methods:

- after rounding (when a nonzero result, computed as though the exponent range were unbounded, would lie strictly between $\pm 2^{E_{\min}}$)
- before rounding (when a nonzero result, computed as though the exponent range and the precision were unbounded, would lie strictly between $\pm 2^{E_{\min}}$).

The MIPS architecture requires that tininess be detected after rounding.

Loss of accuracy can be detected by one of the following methods:

- denormalization loss (when the delivered result differs from what would have been computed if the exponent range were unbounded)
- inexact result (when the delivered result differs from what would have been computed if the exponent range and precision were both unbounded).

The MIPS architecture requires that loss of accuracy be detected as an inexact result.

Trap Enabled Results: If Underflow or Inexact traps are enabled, or if the *FS* bit is not set, then an Unimplemented exception (E) is generated, and the result register is not modified.

Trap Disabled Results: If Underflow and Inexact traps are not enabled and the *FS* bit is set, the result is determined by the rounding mode and the sign of the intermediate result (as listed in Table 12-1).

Unimplemented Instruction Exception (E)

Any attempt to execute an instruction with an operation code or format code that has been reserved for future definition sets the *Unimplemented* bit in the *Cause* field in the FPU *Control/Status* register and traps. The operand and destination registers remain undisturbed and the instruction is emulated in software. Any of the IEEE Standard 754 exceptions can arise from the emulated operation, and these exceptions in turn are simulated.

The Unimplemented Instruction exception can also be signaled when unusual operands or result conditions are detected that the implemented hardware cannot handle properly. These include:

- Denormalized operand, except for Compare instruction
- Quiet Not a Number operand, except for Compare instruction
- Denormalized result or Underflow, when either Underflow or Inexact *Enable* bits are set or the *FS* bit is not set.
- Reserved opcodes
- Unimplemented formats
- Operations which are invalid for their format (for instance, CVT.S.S)

NOTE: Denormalized and NaN operands are only trapped if the instruction is a convert or computational operation. Moves do not trap if their operands are either denormalized or NaNs.

The use of this exception for such conditions is optional; most of these conditions are newly developed and are not expected to be widely used in early implementations. Loopholes are provided in the architecture so that these conditions can be implemented with assistance provided by software, maintaining full compatibility with the IEEE Standard 754.

Trap Enabled Results: The result register is not modified, and the source registers are preserved.

Trap Disabled Results: This trap cannot be disabled.

13.6 Saving and Restoring State

Sixteen doubleword[†] coprocessor load or store operations save or restore the coprocessor floating-point register state in memory. The remainder of control and status information can be saved or restored through CFC1/CTC1 instructions, and saving and restoring the processor registers. Normally, the *Control/Status* register is saved first and restored last.

When state is restored, state information in the *Control/Status* register indicates the exceptions that are pending. Writing a zero value to the *Cause* field of *Control/Status* register clears all pending exceptions, permitting normal processing to restart after the floating-point register state is restored.

13.7 Trap Handlers for IEEE Standard 754 Exceptions

The IEEE Standard 754 strongly recommends that users be allowed to specify a trap handler for any of the five standard exceptions that can compute; the trap handler can either compute or specify a substitute result to be placed in the destination register of the operation.

By retrieving an instruction using the processor *Exception Program Counter (EPC)* register, the trap handler determines:

- exceptions occurring during the operation
- the operation being performed
- the destination format

On Overflow or Underflow exceptions (except for conversions), and on Inexact exceptions, the trap handler gains access to the correctly rounded result by examining source registers and simulating the operation in software.

On Overflow or Underflow exceptions encountered on floating-point conversions, and on Invalid Operation and Divide-by-Zero exceptions,

the trap handler gains access to the operand values by examining the source registers of the instruction.

The IEEE Standard 754 recommends that, if enabled, the overflow and underflow traps take precedence over a separate inexact trap. This prioritization is accomplished in software; hardware sets the bits for both the Inexact exception and the Overflow or Underflow exception.

[†] 32 doublewords if the FR bit is set to 1.

14. Debug Support Unit (Option)

14.1 Features

1. Utilizes JTAG interface compatible with IEEE Std. 1149.1.
2. Additional Status pins and debug clock in conjunction with JTAG pins provide Real-Time Trace information.
3. Processor access to external processor probe to execute from the external trace memory during debug exception and boot time. This is to eliminate system memory for debugging purpose.
4. Supports DMA access through JTAG interface to internal processor bus to access internal registers, host system peripherals and system memory.
5. Debug functions
 - Instruction Address Break
 - Data Bus break
 - Processor Bus Break
 - Hardware Debug Interrupt
 - Reset, NMI, Interrupt Mask
6. Instructions for Debug
 - SDBBP, DERET, CTC0, CFC0
7. CP0 Registers for Debug
 - Debug, DEPC, DESAVE

14.2 EJTAG interface

This interface consists of two modes of operation a Run Time Mode and a Real Time Mode. The Run Time mode provides functions such as processor Run, Stop, Single Step, and access to internal registers and system memory. The Real Time mode provides additional status pins used in conjunction with JTAG pins for Real Time Trace information.

Pins	In/Out	Description
GTCK	I	Test Clock Input
GDCLK	O	Debug Clock (1/3 CPU Clock)
GTDI/GDINT	I	Test Data Input (GTDI) at Run Time mode /Debug Interrupt Input (GDINT) at Real Time mode
GTDO/GTPC[0]	O	Test Data Output (GTDO) /PC Output (GTPC)
GTMS	I	Test Mode Select Input
GTRST*	I	Reset
GPCST[8-0]	O	PC Trace Status Information
GTPC[3-1]	O	PC Output

14.3 JTAG Interface

Standard JTAG interface is used for on chip debugging during Run Time mode. The TX49 Debug Support Unit has following registers.

- Instruction Register
- Bypass Register
- Boundary-Scan Register
- Device Identification Register
- Implementation Register
- JTAG_Data_register
- JTAG_Address_Register
- JTAG_Control_Register

14.4 Processor Access Overview

The core processor can access external processor probe for reading and writing to external monitor memory, registers and other external resources.

In addition the processor can execute from the external monitor memory located from 0xf_ff20 0000 to 0xf_ff2f ffff when the ProbEnb bit is set and the processor probe is turned ON. Any access to the monitor location from 0xf_ff20 0000 to 0xf_ff3f ffff are only allowed when the processor is in the debug mode (DM \square = \square 1).

14.5 Instruction

The instruction is a 8 bit field. Instructions for the TX49 Debug Support Unit are encoded between 0x80 and 0x9f and other codes are reserved for Toshiba Standard JTAG instructions (Includes EXTEST, SAMPLE/PRELOAD, INTEST, IDCODE and HI-Z) and so on. Instructions are decoded as follows.

Hex Value	Instruction	Description
0x83	EJTAG_ImpCode	Select Implementation Register
0x88	JTAG_ADDRESS_IR	Select JTAG_Address Register
0x89	JTAG_DATA_IR	Select JTAG_Data Register
0x8A	JTAG_CONTROL_IR	Select JTAG_Control Register
0x8B	JTAG_ALL_IR	Select JTAG_All Register
0x90	PCTRACE	PCTRACE Instruction

Any unused instruction between 0x80 and 0x9f defaulted to BYPASS instruction.

14.6 Debug Unit

14.6.1 Extended Instructions

- SDBBP
- DERET
- CTC0
- CFC0

14.6.2 Extended Debug Registers in CP0

- Debug Register
- Debug Exception PC (DEPC)
- Debug SAVE

14.7 Register Map

Address	Mnemonic	Description
0xf ff30 0000	DCR	Debug Control Register
0xf ff30 0008	IBS	Instruction Break Status
0xf ff30 0010	DBS	Data Break Status
0xf ff30 0018	PBS	Processor Break Status
0xf ff30 0100	IBA0	Instruction Break Address 0
0xf ff30 0108	IBC0	Instruction Break Control 0
0xf ff30 0110	IBM0	Instruction Break Address Mask 0
0xf ff30 0300	DBA0	Data Break Address 0
0xf ff30 0308	DBC0	Data Break Control 0
0xf ff30 0310	DBM0	Data Break Address Mask 0
0xf ff30 0318	DB0	Data Break Value 0
0xf ff30 0600	PBA0	Processor Bus Break Address 0
0xf ff30 0608	PBD0	Processor Bus Break Data 0
0xf ff30 0610	PBM0	Processor Bus Break Mask 0
0xf ff30 0618	PBC0	Processor Bus Break Control 0

14.8 Processor Bus Break Function

This function is to monitor the interface to core and provide debug interruption or trace trigger for a given physical address and data.

14.9 Debug Exception

Three kinds of debug exception are supported.

- Debug Single Step (DSS bit)
- Debug Breakpoint Exception (SDBBP Instruction)
- JTAG Break Exception (Jtagbrk bit in JTAG_Control_Register)

Note: During real time debugging, first two functions are disabled.

14.9.1 Debug Single Step (DSS)

When the debug register DSS bit is set, this exception has been raised each time one instruction is executed.

14.9.2 Debug Breakpoint exception (Dbp)

This exception is raised when SDBBP instruction is executed.

14.9.3 JTAG Break Exception

This exception is raised when JTAG unit set the Jtagbrk in JTAG_Control_Register.

14.9.4 Debug Exception Handling

Updates DEPC and Debug register.

Registers other than DEPC and Debug register retain their values.

14.9.5 Branching to debug handler

If the ProbEnb bit in JTAG_Control_Register[15] is set, the debug exception vector is located at

PC: 0xffff ffff ff20 0200.

If the ProbEnb bit in JTAG_Conctrol_Register[15] is cleared, the debug exception vector is located at

PC: 0xffff ffff bfc0 0400.

14.9.6 Exception handling when in Debug Mode (DM bit is set)

All interrupts including NMI are masked. When the NMI interrupt has occurred during Debug mode, it is stored internally and the NMI interrupt is taken after debug handler is finished (DM is clear).

14.10 Real Time PC TRACE Output

In real time mode non-sequential Program Counter and trace information are outputted on GTPC[3~0] and GPCST[8~0]. at 1/3 of the processor clock speed.

14.11 Example of the PC trace output

15. TX49 MPU Core Signal Description %&

The TX49 MPU core has a 64-bit BUS I/F upward compatible to the TX39 G-Bus I/F.

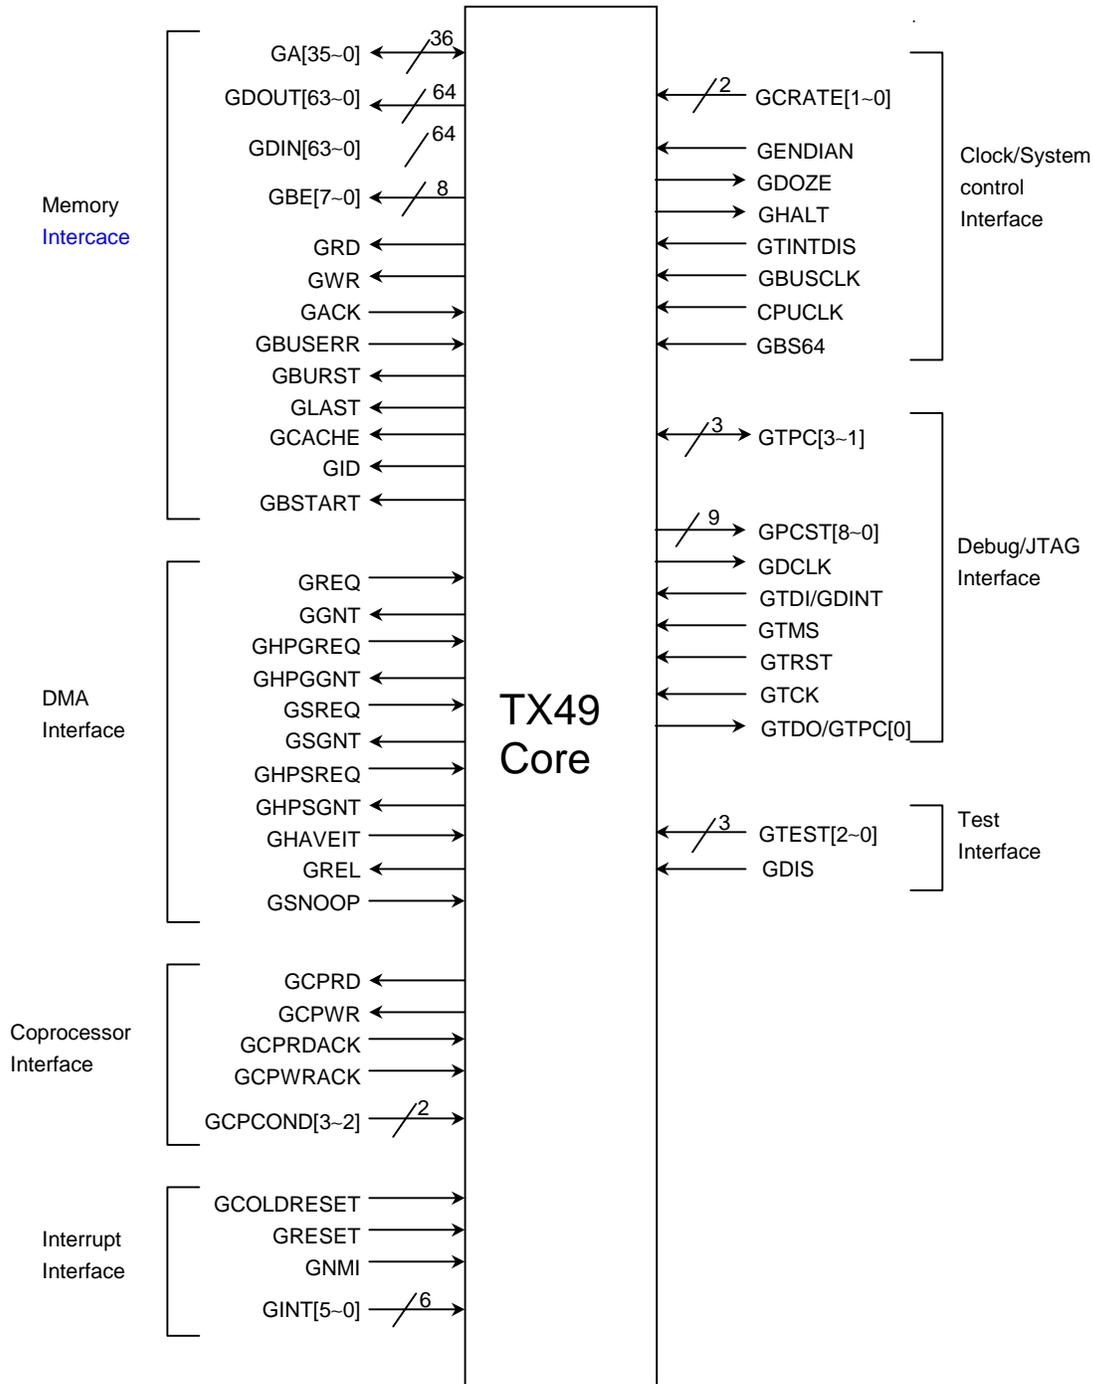


Figure 15-1 Interface signal of TX49 MPU Core

15.1 Signal description

15.1.1 Memory interface

Table 15-1 lists a memory interface signals.

Table 15-1 Signal description of a memory interface

Name	I/O	Description
GA[35~0]	I/O	Address Input/Output A 36-bit address input/output signal. It can also be used as a 36-bit address input signal to perform snooping of the data cache.
GDOUT[63~0]	O	Data Output A 64-bit data output signal.
GDIN[63~0]	I	Data Input A 64-bit data input signal
GBE[7~0] *	O	Byte Enable Indicates the byte that is being accessed in a bus operation. The correspondence with the data bus is. GBE[7]* GDOUT[63~56] and GDIN[63~56] GBE[6]* GDOUT[55~48] and GDIN[55~48] GBE[5]* GDOUT[47~40] and GDIN[47~40] GBE[4]* GDOUT[39~32] and GDIN[39~32] GBE[3]* GDOUT[31~24] and GDIN[31~24] GBE[2]* GDOUT[23~16] and GDIN[23~16] GBE[1]* GDOUT[15~8] and GDIN[15~8] GBE[0]* GDOUT[7~0] and GDIN[7~0]
GRD*	O	Read Indicates that a read operation is being executed.
GWR*	O	Write Indicates that a write operation is being executed.
GACK*	I	Read/Write Acknowledge Informs that TX49 that there are valid data the data bus. The TX49 ends the bus operation in the next cycle after acknowledging the assertion of this signal.

Name	I/O	Description
GBUSERR*	I	<p>Bus Error</p> <p>Indicates that an error has occurred in the TX49's bus operation. This bus error is valid only in a read bus operation. The read bus operation ends in the cycle where the assertion of the GBUSERR* signal is acknowledged. If a read bus operation is ended by the GBUSERR* signal, the TX49 generates a bus error exception. If a bus error occurred in a write bus operation, instead of the GBUSERR* signal being asserted, the write bus operation should be terminated by asserting the GACK* signal and an exception should be generated using the GINT[5-0]* signal.</p>
GBURST*	O	<p>Burst</p> <p>Indicates that a burst read operation is being executed.</p>
GLAST*	O	<p>Last</p> <p>Indicates the final data transfer.</p>
GCACHE*	O	<p>Cacheable</p> <p>Indicates that the bus operation being executed is either an operation to the cache area or an operation to the non-cache.</p> <p>H: Non-cache area L: Cache area</p>
GID	O	<p>Instruction or Data</p> <p>Indicates that the bus operation being executed is an operation performed on either an instruction or on the data.</p> <p>H: Instruction L: Data</p>
GBSTART*	O	<p>Bus Start</p> <p>A signal that is to be asserted for the first clock cycle when a bus operation starts. This signal can acknowledge that a bus operation has started.</p>

15.1.2 DMA interface

Table 15-1 lists a DMA interface signals.

Table 15-1 Signal description of a DMA interface

Name	I/O	Description
GREQ*	I	Normal Bus Request A signal with which the external bus master requests the TX49 for the G-Bus ownership based on ET concurrency. (Low priority)
GGNT*	O	Normal Bus Grant A signal which indicates that the TX49 granted the G-Bus ownership request (GREQ*) made by the external bus master.
GHPGREQ*	I	High Priority Normal Bus Request A signal with which the external bus master requests the TX49 for the G-Bus ownership based in ET concurrency. (High priority)
GHPGGNT*	O	High Priority Normal Bus Grant A signal which indicates that the TX49 granted the G-Bus ownership request (GHPGREQ*) made by the external bus master.
GSREQ*	I	Snoop Bus Request A signal with which the external bus master requests the TX49 for the X-Bus ownership based on ST concurrency. (Low priority)
GSGNT*	O	Snoop Bus Grant A signal which indicates that the TX49 granted the X-Bus ownership request (GSREQ*) made by the external bus master. Snoop operation can be performed while this signal is being asserted.
GHPGRESQ*	I	High Priority Snoop Bus Request A signal with which the external bus master requests the TX49 for the X-Bus ownership based on ST concurrency. (High priority)
GHPSGNT*	O	High Priority Snoop Bus Grant A signal which indicates that the TX49 granted the X-Bus ownership request (GHPGRESQ*) made by the external bus master. Snoop operation can be performed while this signal is being asserted.
GHAVEIT*	I	Have IT This is an input signal for confirming the bus ownership that is output from the external bus master which has obtained ownership of the G-Bus or the X-Bus.
GREL*	O	Release Request This is an output for the TX49 to request the external bus master to release the bus ownership it currently retains.
GSNOOP*	I	Snoop Indicates that the address on the GA[35-0] signal is valid for the snoop operation. The snoop operation is invalid while GGNT* or GHPGGNT* is being asserted.

15.1.3 Coprocessor interface

Table 15-1 lists a coprocessor interface signals.

Table 15-1 Signal description of a coprocessor interface

Name	I/O	Description
GCPRD*	O	Coprocessor Read A read signal to a coprocessor
GCPWR*	O	Coprocessor Write A write signal to a coprocessor
GCPRDACK*	I	Coprocessor Read Acknowledge A read acknowledge signal from a coprocessor
GCPWRACK*	I	Coprocessor Write Acknowledge A write acknowledge signal from a coprocessor
GCPCOND[3-2]	I	Coprocessor Condition An input signal that is to be a branch condition of a coprocessor condition branch instruction. The GCPCOND[3] signal, the GCPCOND[2] signal and the GCPCOND[1] signal are each used in the coprocessor condition branch instructions that support CP3, CP2 and CP1 respectively.

15.1.4 Interrupt interface

Table 15-1 lists a interrupt interface signals.

Table 15-1 Signal description of a interrupt interface

Name	I/O	Description
GCOLDRESET*	I	Coldreset This signal must be asserted for a power on reset or a cold reset. The clocks begin to cycle and are synchronized with the deasserted edge of GCOLDRESET*.
GRESET*	I	Reset Resets the TX49 and generates a reset exception.
GNMI*	I	Non-maskable Interrupt A signal for generating an non-maskable interrupt exception.
GINT[5-0] *	I	Interrupt A signal for generating an interrupt exception.

15.1.5 Clock/system control interface

Table 15-1 lists a clock/system control interface signals.

Table 15-1 Signal description of a clock/system control interface

Name	I/O	Description
CPUCLK	I	CPU Clock Input Establishes the processor operation frequency.
GCRATE [1-0]	I	GBUS Clock Rate Input from External Pin 00 divided by 2 01 divided by 3 10 divided by 4 11 divided by 1
GBUSCLK	I	GBUS Clock Input
GTINTDIS	I	Timer interrupt disable Input from External Pin H: disabled timer interrupt L: enabled timer interrupt
GHALT	O	Halt This signal output the status of the Halt bit of Config register. This signal indicates that the TX49 is in the halt mode when this signal is "H".
GDOZE	O	Doze This signal output the status of the Halt bit of Config register. This signal indicates that the TX49 is in the doze mode when this signal is "H".
GENDIAN	I	Endianess Input from External Pin Indicates the initial setting of the endian during a reset. H: Big Endian L: Little Endian
GBS64*	I	System bus size. H: 32• bit G-bus mode L: 64• bit G-bus mode

15.1.6 Debug interface

Table 15-1 lists a debug and JTAG interface signals.

Table 15-1 Signal description of Debug and JTAG interface

Name	I/O	Description
GDCLK	O	Debug Clock (1/3 CPU Clock) A clock output for a real-time debug system. The timing of a serial monitor bus and PC trace interface signals are all defined by this debug clock GDCLK. The operation clock of the TX49 Processor Core is divided by 3 at the time of a serial monitor bus operation.
GPCST[8-0]	O	PC Trace Status Outputs PC trace status information and the mode of the serial monitor bus.

Name	I/O	Description
GTPC[3-1]	I/O	Trace PC Output. Outputs a non-sequential program counter at DCLK.
GTCK	I	Test Clock Input Input clock used to shift data into or out of the EJTAG Instruction or Data register. The GTCK maybe independent of the CPUCLK.
GTDI/GDINT	I	Test Data Input / Debug Interrupt Run time mode: input serial data to data/instruction register of EJTAG Real time mode: interrupt line to change the debug unit state from real time mode to run time mode.
GTMS	I	Test Mode Select Input A signal select EJTAG mode.
GTRST*	I	Test Reset Input A reset input for a real-time debug system. When GTRST* is asserted, the debug support unit (DSU) is initialized.
GTDO	O	Test Data Output Run time mode: outputs serial data from data/instruction register of EJTAG Real time mode: outputs non-sequential program
GTDOE	O	Test Data Output Enable

15.1.7 Test interface

Table 15-1 lists a test interface signals.

Table 15-1 Signal description of a test interface

Name	I/O	Description
GTEST[2-0]	I	Test Sets test modes in order to test the TX49. Under normal operating conditions, all of these pins must be fixed to "L". The circuit must be configured so that external signals can be applied to these pins when testing the TX49.
GDIS	I	Disable output A signal for changing the TX49 output to Hi-Z (High impedance). Output can be changed to Hi-Z by setting this pin to "L". Under normal operating conditions, fix this pin to "H".

16. Bus Operations

All bus operations of the TX49 Megacell synchronize with the rising of GBUSCLK. In this section, consecutive serial numbers indicating timing are put on each rise of GBUSCLK in order to simplify the explanations.

The memory operations of the TX49 Megacell are divided into single read operations, the burst read operations, single write operations, and burst write operations.

16.1 Single read operations

Single read operations are bus operations for reading data of 8 bytes or less. These operations occur in the following situations:

- Instruction fetch or data load from the non-cache area
- Instruction fetch during instruction cache Disable
- Data load in the case of data cache Disable

In the single read operation, the GA [35~0] signal, the GDIN[63~0] signal, the GBE[7~0]* signal, the GBSTART* signal, the GRD* signal, the GLAST* signal, the GCACHE* signal, the GID signal, the GACK* signal, and the GBUSERR* signal are used. Figure 16-1 shows the basic single read operation.

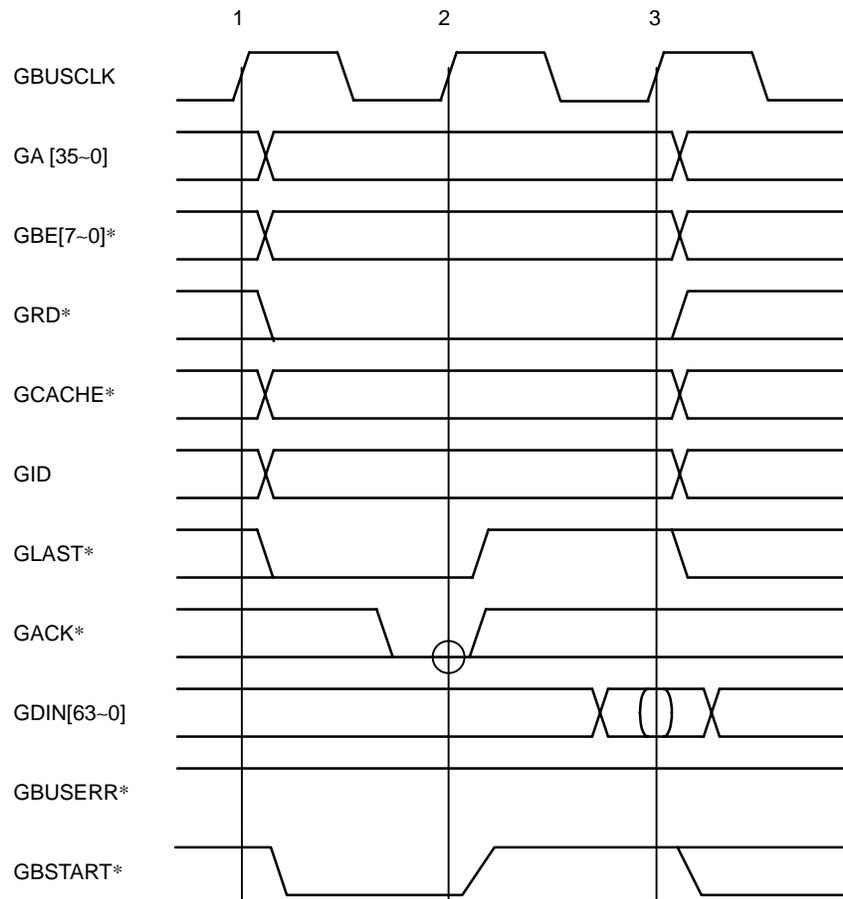


Figure 16-1 Basic Single Read Operation

- 1 The GBSTART* signal, the GRD* signal and the GLAST* signal are asserted, and a single read operation starts. At the same time, valid values are output to the GA [35~0] signal, the GBE[7~0]* signal, the GCACHE* signal, and the GID signal.
- 2 The GBSTART* signal is deasserted. After acknowledging that the GACK* signal is "L," the GLAST* signal is deasserted.
- 3 Data on the GDIN[63~0] signal are collected. Also, the GRD* signal is deasserted. This completes the single read operation. Up to this point in time, valid values are output to the GA[35~0] signal, the GBE[7~0]* signal, the GCACHE* signal, and the GID signal.

If a read operation follows immediately after this single read operation, the GRD* signal is not deasserted, and the read operation starts. A single read operation without a wait cycle completes in two cycles as shown in Figure 16-1.

A peripheral circuit can insert a wait cycle by not asserting the GACK* signal. Figure 16-2 shows a single read operation with one cycle wait inserted.

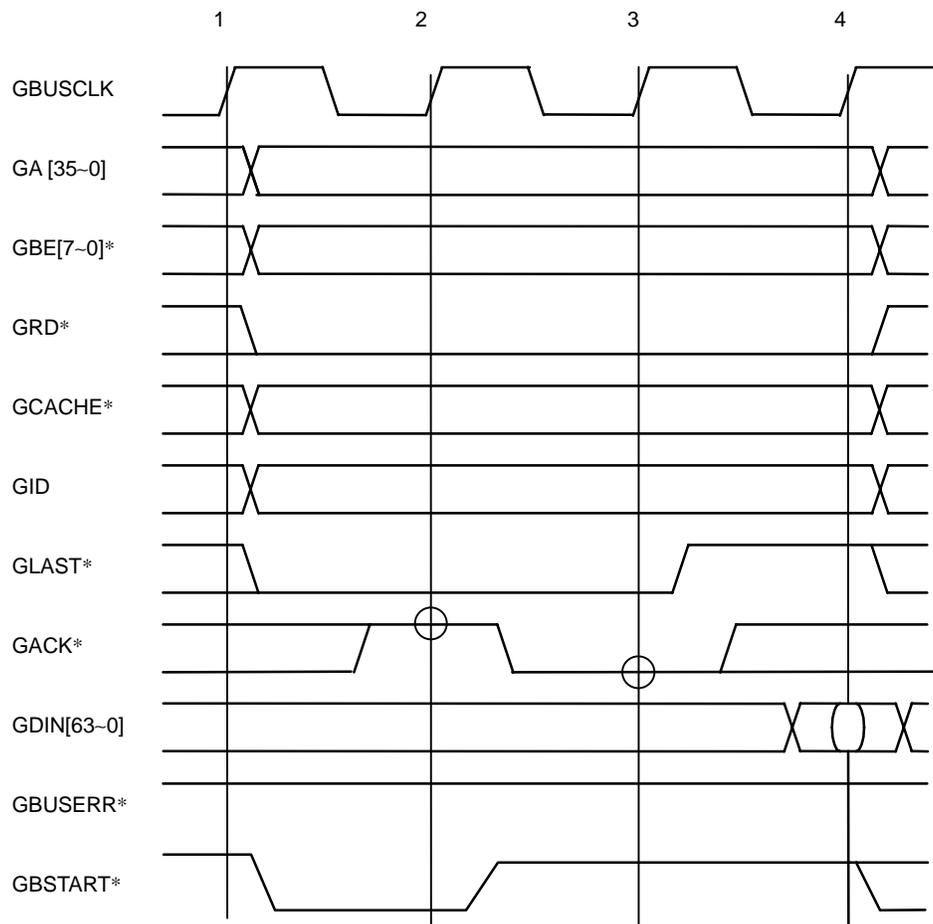


Figure 16-2 Single Read Operation with a Wait Cycle

- 1 The GBSTART* signal, the GRD* signal, and the GLAST* signal are asserted, and a single read operation starts. At the same time, valid values are output to the GA [35~0] signal, the GBE[7~0]* signal, the GCACHE* signal, and the GID signal.
- 2 The GBSTART* signal is deasserted. Since the GACK* signal is "H," data cannot be collected at the timing of Item 3.
- 3 After acknowledging that the GACK* signal is "L," the GLAST* signal is deasserted.
- 4 The data on the GDIN[63~0] signal are collected. Also, the GRD* signal is deasserted. This completes the single read operation. Up to this timing, valid values are output to the GA[35~0] signal, the GBE[7~0]* signal, the GCACHE* signal, and the GID signal.

If the single read operation causes an error in the system, peripheral circuits can inform the TX49 Megacell of the error through the GBUSERR* signal. The TX49 Megacell ends the single read operation by deasserting the GRD* signal at the rise of GBUSCLK where the assertion of the GBUSERR* signal is acknowledged. The GBUSERR* signal is sampled at all rise of GBUSCLK up to the timing when the TX49 Megacell collects data.

Figure 16-3 shows the status where the assertion of the GBUSERR* signal is acknowledged in the timing when the TX49 Megacell collects data. In this case, the data on the GDIN[63~0] signal are not read.

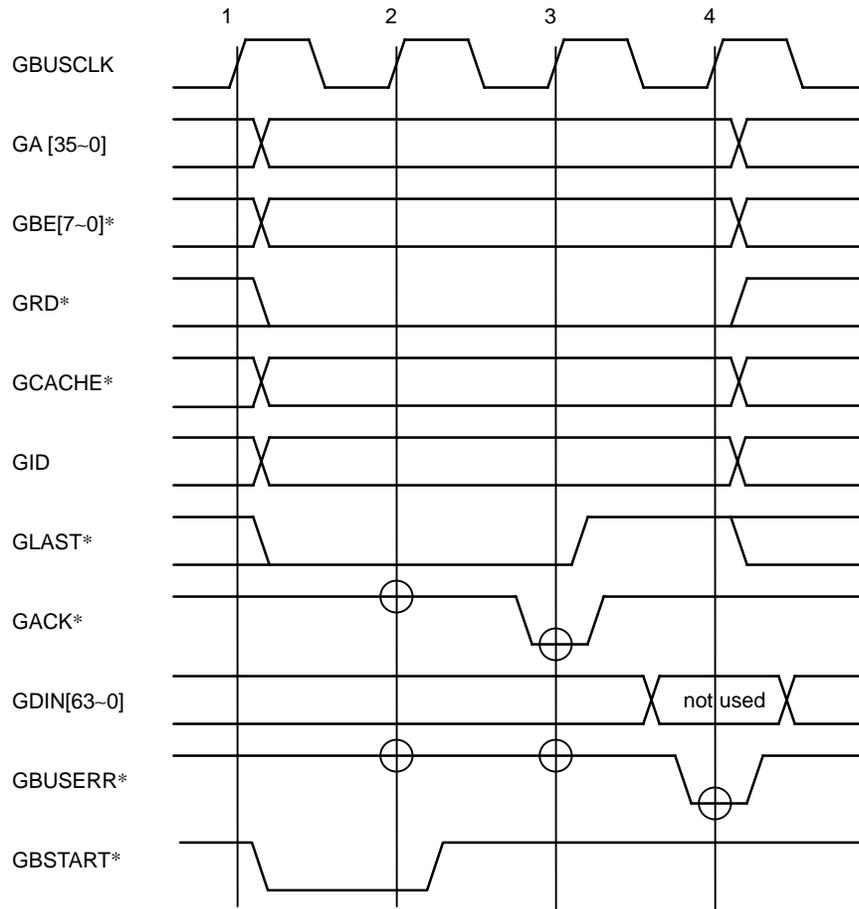


Figure 16-3 Single Read Operation That is Ended by GBUSERR* Signal Assertion

- 1 The GBSTART* signal, the GRD* signal, and the GLAST* signal are asserted, and a single read operation starts. At the same time, valid values are output to the GA[35~0] signal, the GBE[7~0]* signal, the GCACHE* signal, and the GID signal.
- 2 The GBSTART* signal is deasserted. Since the GACK* signal is "L," data cannot be collected at the timing of Item 3.
- 3 After acknowledging that the GACK* signal is "L," the GLAST* signal is deasserted.
- 4 After acknowledging that the GBUSERR* signal is "L," the GRD* signal is deasserted. The data on the GDIN[63~0] signal is not read. This completes the single read operation. Up to this timing, valid values are output to the GA[35~0] signal, the GBE[7~0]* signal, the GCACHE* signal, and the GID signal.

If a single read operation is ended by the assertion of the GBUSERR* signal, the TX49 Megacell generates a bus error exception.

The TX49 Megacell ends a single read operation when the GACK* signal or the GBUSERR* signal is asserted.

16.2 Burst read operations

The burst read operation is a bus operation used for quickly refilling multiple words to the cache. This operation enables second and subsequent data reads to be completed in one cycle. Therefore, the memory burst transfer mode can be used to read multiple words.

The burst read operation occurs due to an instruction cache miss or a data cache miss (Load data, Store data at write-through-write-allocate or write-back). The burst refill size of the data cache and instruction cache are 32 bytes.

In the burst read operation, the $GBURST^*$ signal is used in addition to the $GA[35\sim 0]$ signal, the $GDIN[63\sim 0]$ signal, the $GBE[7\sim 0]^*$ signal, the GRD^* signal, the $GLAST^*$ signal, the $GCACHE^*$ signal, the GID signal, the $GACK^*$ signal, the $GBUSERR^*$ signal. "L" is always output to the $GBE[7\sim 0]^*$ signal and the $GCACHE^*$ signal. Figure 16-1 shows a burst read operation with no wait cycle.

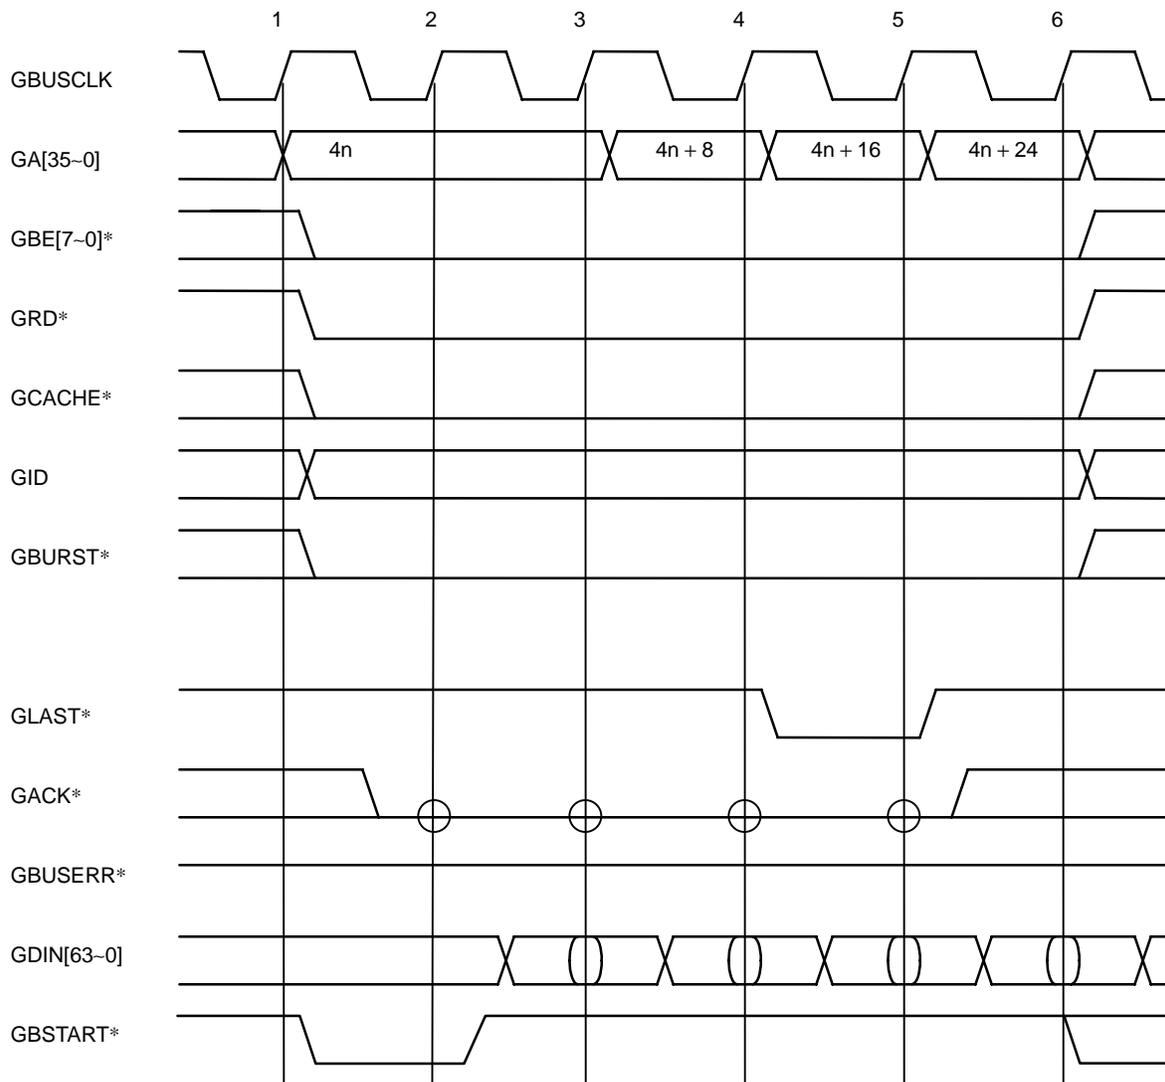


Figure 16-1 Burst Read Operation

- 1 The GRD* signal and the GBURST* signal are asserted, and a burst read operation starts. At the same time, valid values are output to the GA[35~0] signal, the GBE[7~0]* signal, the GCACHE* signal, the GID signal.
- 2 Since it is acknowledged that the GACK* signal is "L," data are collected at timing of Item 3.
- 3 The first data are collected from the GDIN[63~0] signal. The address values that are being output to the GA[35~0] signal are output in increments of eight. Since it is acknowledged that the GACK* signal is "L," data are collected at the timing of Item 4.
- 4 The second data are collected from the GDIN[63~0] signal. The address values that are being output to the GA[35~0] signal are output in increments of eight. Since it is acknowledged that the GACK* signal is "L," data are collected at the timing of Item 5. Also, the GLAST* signal is asserted.
- 5 The third data are collected from the GDIN[63~0] signal. The address values that are being output to the GA[35~0] signal are output in increments of eight. Since it is acknowledged that the GACK* signal is "L," data are collected at the timing of Item 6. Also, the GLAST* signal is deasserted.
- 6 The final data are collected from the GDIN[63~0] signal. The GRD* signal and the GBURST* signal are deasserted. This completes the burst read operation. Up to this timing, valid values are output to the GA[35~0] signal, the GBE[7~0] signal, the GCACHE* signal, and the GID signal.

If a single read operation follows immediately after the burst read operation, the GRD* signal is not deasserted, so the single read operation starts. Also, if another burst read operation follows immediately after this burst read operation, neither the GBURST* signal nor the GRD* signal is deasserted, so the next burst read operation starts.

Therefore, by checking the GACK* signal for the data to be collected next at the same timing as when collecting data, the second and subsequent data can be read in one cycle. The addresses on the GA[35~0] signal always indicate the address of the data that are being read. Since the TX49 Megacell starts a cache refill from the word of the lowest-order address in the refill boundary, the address in a burst read operation is simply output in increments of eight. The timing to assert the GLAST* signal is when the assertion of the GACK* signal for the data read second from the last is acknowledged. The timing to deassert the GLAST* signal is when the assertion of the GACK* signal for the last data read is acknowledged.

Similar to the single read operation, wait cycles can be inserted by deasserting the GACK* signal. Figure 16-2 shows a burst read operation with wait cycles. This is an example where wait cycles are inserted at the first and the third data read in a burst read operation.

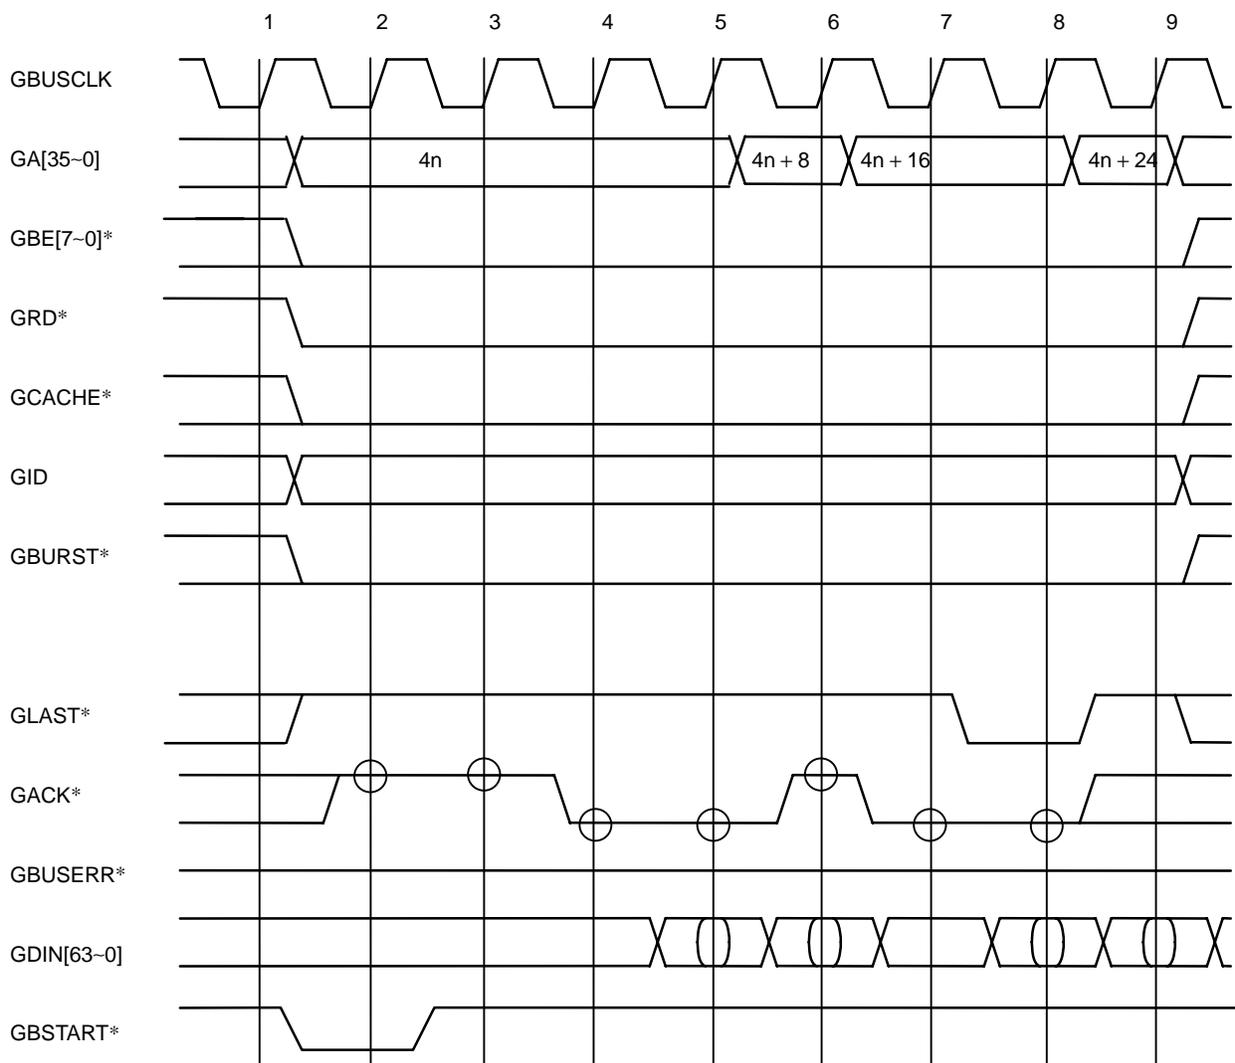


Figure 16-2 Burst Read Operation with Wait Cycles

- 1 The GBSTART* signal, the GRD* signal, and the GBURST* signal are asserted, and a burst read operation starts. At the same time, valid values are output to the GA[35~0] signal, the GBE[7~0]* signal, the GCACHE* signal, the GID signal.
- 2 The GBSTART* signal is deasserted. Since the GACK* signal is "H," data are not collected at timing of Item 3.
- 3 Since the GACK* signal is "H," data are not collected at the timing of Item 4.
- 4 Since it is acknowledged that the GACK* signal is "L," data are collected at the timing of Item 5.
- 5 The first data are collected from the GDIN[63~0] signal. The address values that are output on the GA[35~0] signal are output in increments of eight. Since it is acknowledged that the GACK* signal is "L," data are collected at the timing of Item 6.
- 6 The second data are collected from the GDIN[63~0] signal. The address values that are output on the GA[35~0] signal are output in increments of eight. Since it is acknowledged that the GACK* signal is "H," data are collected at the timing of Item 7.
- 7 Since it is acknowledged that the GACK* signal is "L," data are collected at the timing of Item 8. Also, the GLAST* signal is asserted.
- 8 The third data are collected from the GDIN[63~0] signal. The address values that are output on the GA[35~0] signal are output in increments of eight. Since it is acknowledged that the GACK* signal is "L," data are collected at the timing of Item 9. Also, the GLAST* signal is deasserted.
- 9 The final data are collected from the GDIN[63~0] signal. Also, the GRD* signal and the GBURST* signal are deasserted. This completes the burst read operation. Up to this timing, valid values are output to the GA[35~0] signal, GBE[7~0]* signal, the GCACHE* signal, and the GID signal.

Similar to the single read operation, if the burst read operation causes an error in the system, peripheral circuits can inform the TX49 Megacell of the error through the GBUSERR* signal. The TX49 Megacell ends the burst read operation by deasserting the GRD* signal and the GBURST* signal at the rise of the GBUSCLK where the assertion of the GBUSERR* signal is acknowledged.

Figure 16-3 shows the status where the GBUSERR* signal is asserted prior to or simultaneously with the assertion of the GACK* signal for the fourth data read.

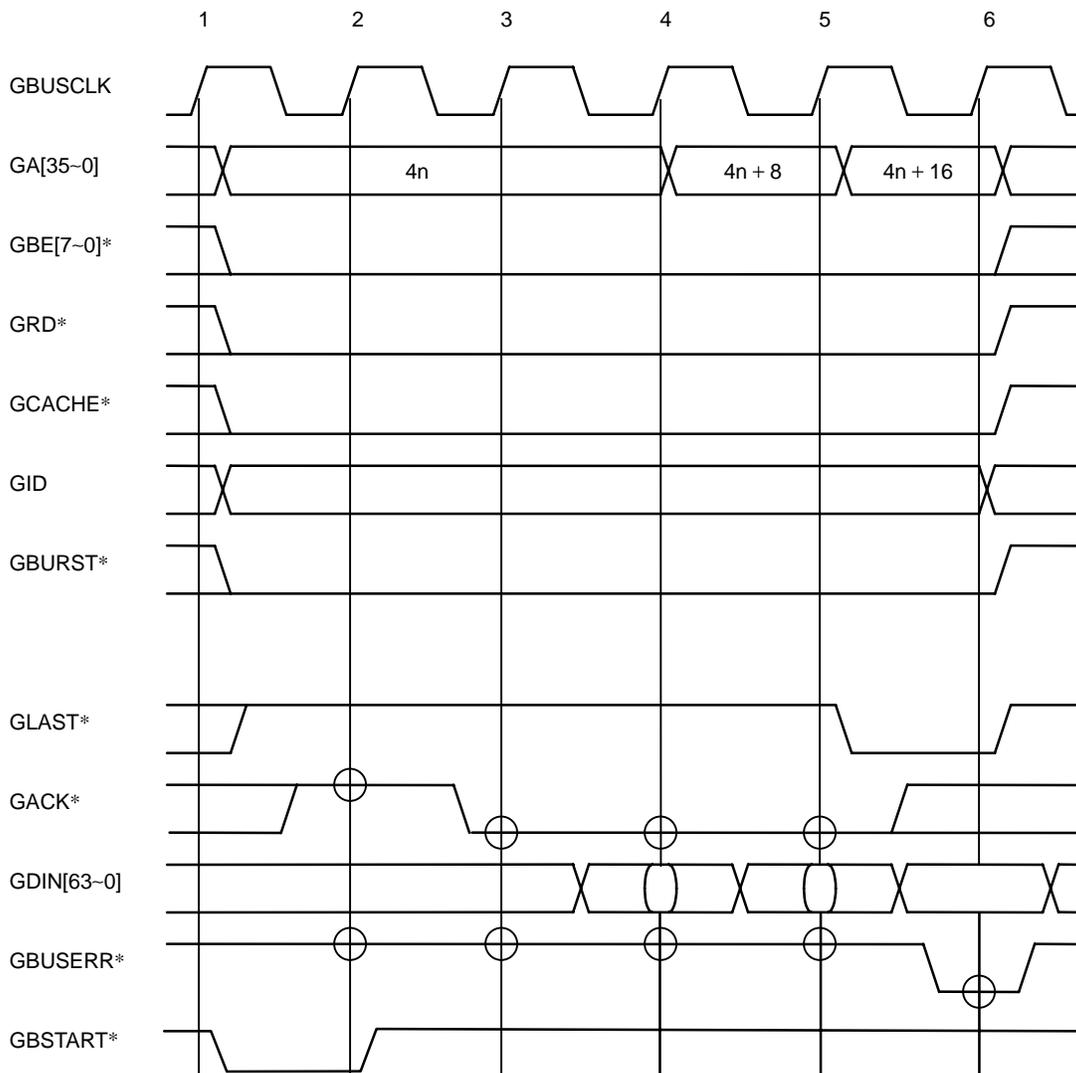


Figure 16-3 Burst Read Operation That is Ended by Assertion of GBUSERR* Signal

- 1 The GBSTART* signal, the GRD* signal, and the GBURST* signal are asserted, and a burst read operation starts. At the same time, valid values are output to the GA[35~0] signal, the GBE[7~0]* signal, the GCACHE* signal, and the GID signal.
- 2 The GBSTART* signal is deasserted. Since the GACK* signal is "H," data are not collected at the timing of Item 3.
- 3 Since it is acknowledged that the GACK* signal is "L," data are collected at the timing of Item 4 .
- 4 The first data are collected from the GDIN[63~0] signal. The address values that are output on the GA[35~0] signal are output in increments of eight. Since it is acknowledged that the GACK* signal is "L," data are collected at the timing of Item 5.
- 5 The second data are collected from the GDIN[63~0] signal. The address values that are output on the GA[35~0] signal are output in increments of eight. Since it is acknowledged that the GACK* signal is "L," data are collected at the timing of Item 6. Also, the GLAST* signal is asserted.

- 6 After acknowledging that the GBUSERR* signal is "L," the GRD* signal and the GBURST* signal are deasserted. This completes the burst read operation. Up to this timing, valid values are output to the GA[35~0] signal, the GBE[7~0]* signal, the GCACHE* signal, and the GID signal.

If a burst read operation is ended due to a bus error, the cache line is invalidated. The TX49 Megacell then generates a bus error exception.

16.3 Single write operations

Single write operations are bus operations for writing data of 8 bytes or less. These operations occur in the following situations:

- Data store to the non-cache area
- Data store in the case of the data cache uses write-through method (no relation with lock bit)
- Data store in the case of the data cache disable

A write operation from the TX49 Processor Core can be completed in one cycle without stalling the pipeline since the TX49 Megacell has an on-chip write buffer. However, if the write buffer is full, a write operation from the TX49 Processor Core stalls until an entry in the write buffer becomes empty. A write operation from the write buffer to the outside is a bus operation which requires a minimum of two clock cycles.

The write operation uses the GA[35~0] signal, the GDOUT[63~0] signal, the GBE[7~0]* signal, the GWR* signal, the GLAST* signal, the GCACHE* signal, the GID signal, and the GACK* signal. Figure 16-1 shows the basic write operation.

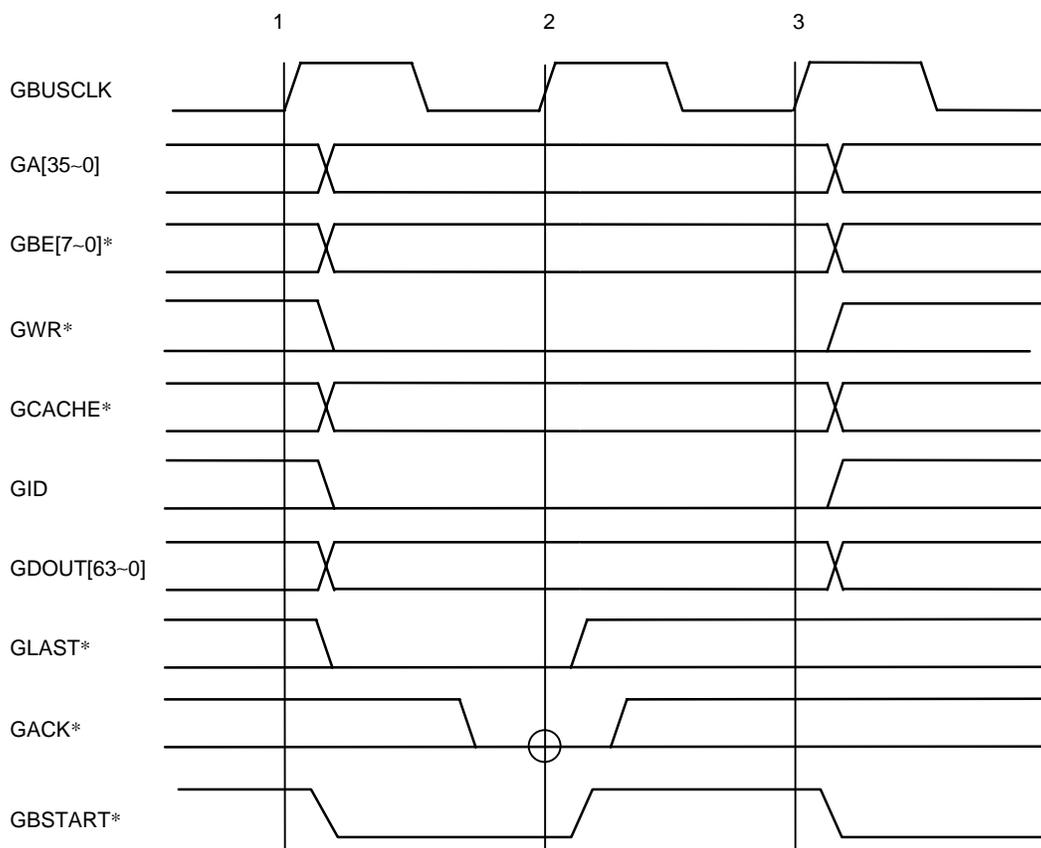


Figure 16-1 Basic Single Write Operation

- 1 The GBSTART* signal, the GWR* signal, and the GLAST* signal are asserted, and a write operation starts. At the same time, valid values are output to the GA[35~0] signal, the GDOUT[63~0] signal, the GBE[7~0]* signal, the GCACHE* signal, and the GID signal. The GID signal is "L" since a write operation always handles data.
- 2 The GBSTART* signal is deasserted. After acknowledging that the GACK* signal is "L," the GLAST* signal is deasserted.
- 3 In the next clock cycle after the clock cycle in which the GACK* signal is acknowledged, the GWR* signal is deasserted. This completes the write operation. Up to this timing, valid values are output to the GA[35~0] signal, the GDOUT[63~0] signal, the GBE[7~0]* signal, the GCACHE* signal, and the GID signal.

If another write operation immediately follows this write operation, the GWR* signal is not deasserted, and the write operation starts.

Similar to the read operations, the bus operation is completed in the timing of the clock rise after the clock cycle in which the assertion of the GACK* signal is acknowledged.

The peripheral circuits can insert wait cycles by not asserting the GACK* signal. Figure 16-2 shows a write operation with two wait cycles.

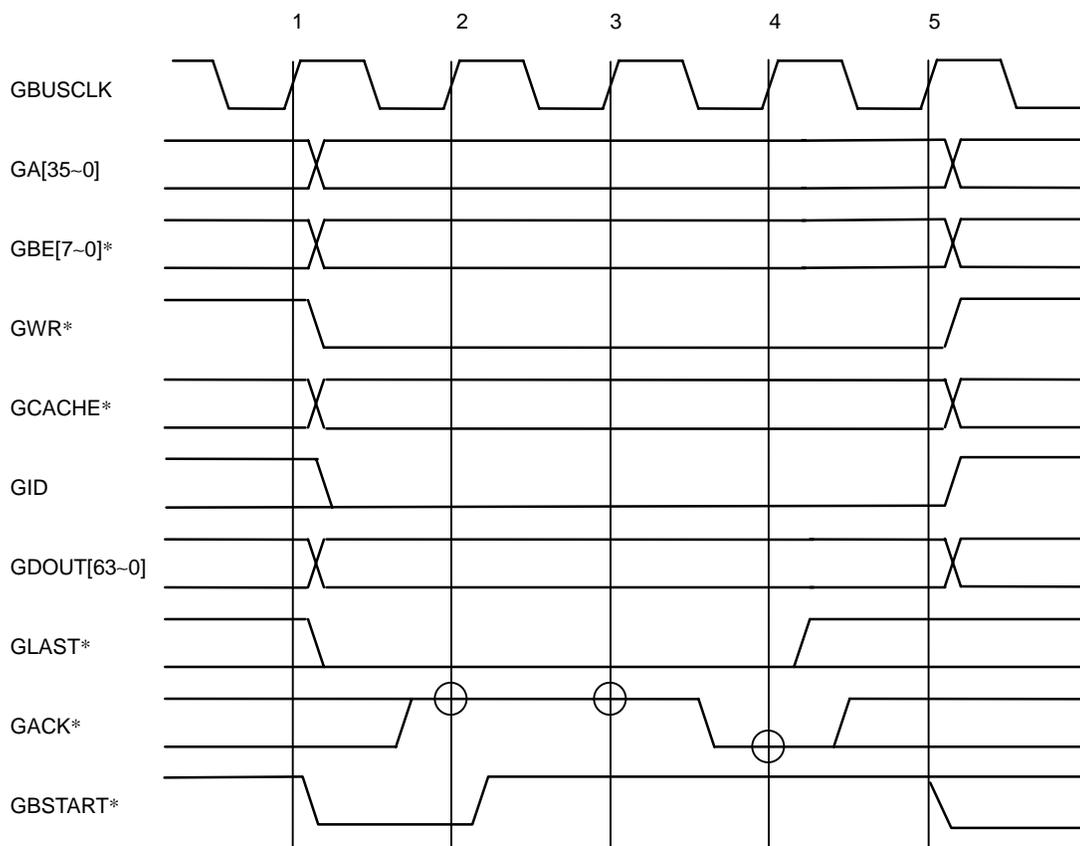


Figure 16-2 Single Write Operation with Wait Cycles

- 1 The GBSTART* signal, the GWR* signal, and the GLAST* signal are asserted, and a write operation starts. At the same time, valid values are output to the GA[35~0] signal, the GDOUT[63~0] signal, the GBE[7~0]* signal, the GCACHE* signal, and the GID signal. The GID signal is "L" since a write operation always handles data.
- 2 The GBSTART* signal is deasserted. Since the GACK* signal is "H," it becomes a wait cycle.
- 3 Since the GACK* signal is "H," it becomes a wait cycle.
- 4 After acknowledging that the GACK* signal is "L," the GLAST* signal is deasserted.
- 5 The GWR* signal is deasserted. This completes the write operation. Up to this timing, valid values are output to the GA[35~0] signal, the GDOUT[63~0] signal, the GBE[7~0]* signal, the GCACHE* signal, and the GID signal.

16.4 Burst write operations

The burst write operation is a bus operation used for quickly write-back multiple words to the memory. This operation enables second and subsequent data writes to be completed in one cycle. Therefore, the memory burst transfer mode can be used to write multiple words.

The burst write operation occurs due to a data cache miss and the cache line is dirty. The burst write size of the data cache is 32bytes.

In the burst write operation, the GBURST* signal is used in addition to the GA[35~0] signal, the GDOUT[63~0] signal, the GBE[7~0]* signal, the GWR* signal, the GLAST* signal, the GCACHE* signal, the GID signal and the GACK* signal. "L" is always output to the GBE[7~0]* signal and the GCACHE* signal.

Figure 16-1 shows a burst write operation with no wait cycle.

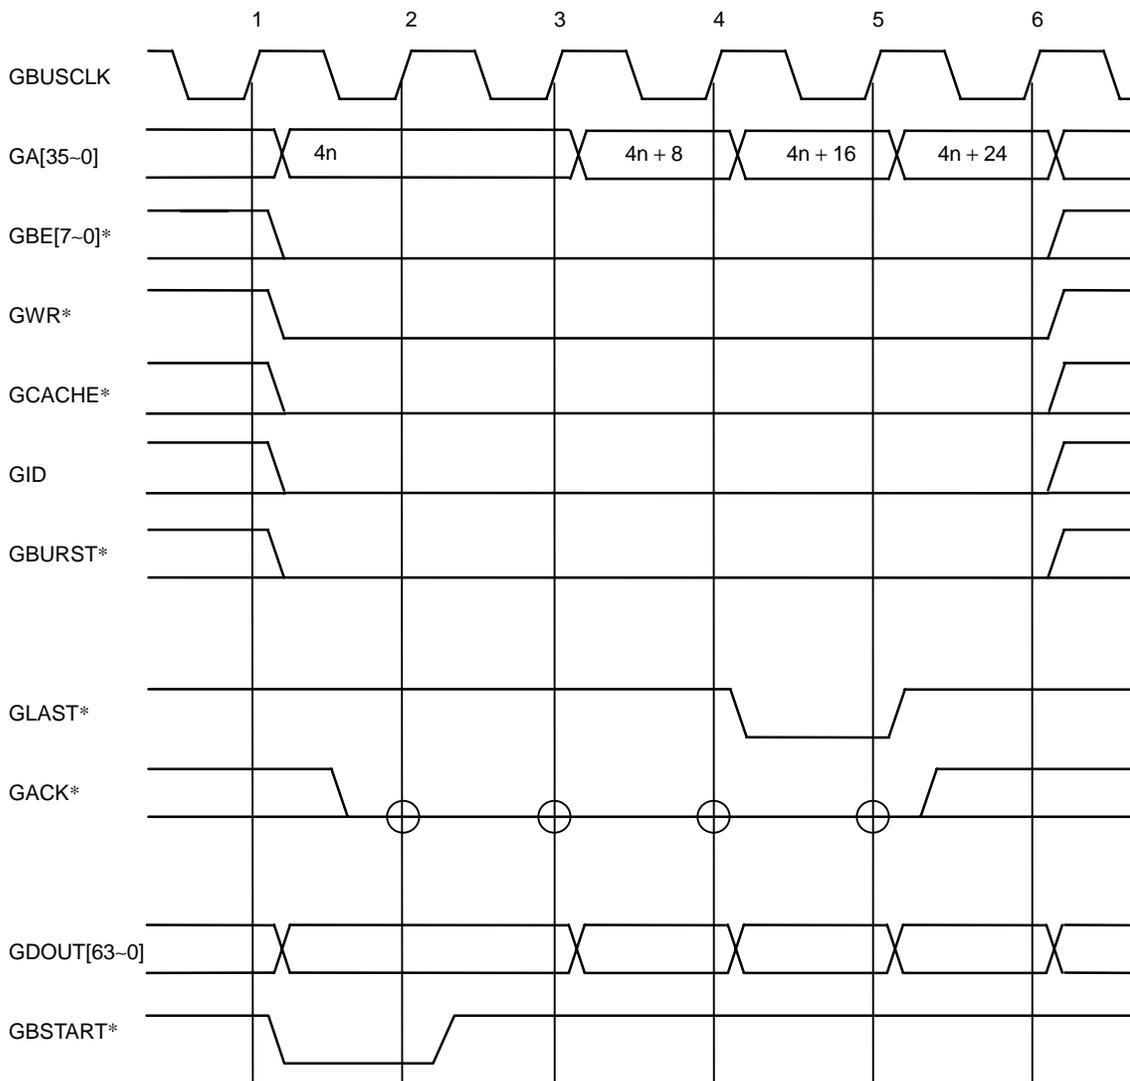


Figure 16-1 Burst write Operation (no wait)

- 1 The GBSTART* signal, GWR* signal and the GBURST* signal are asserted, and a burst write operation starts. At the same time, valid values are output to the GA[35~0] signal, the GDOUT[63~0] signal, the GBE[7~0]* signal, the GID signal.
- 2 The GBSTART* signal is deasserted. Since it is acknowledged that the GACK* signal is "L," next data are output at the timing of Item 3.
- 3 The second valid values are output to the GDOUT[63~0] signal. The address values that are being output to the GA[35~0] signal are output in increments of eight. Since it is acknowledged that the GACK* signal is "L," next data are output at the timing of Item 4.
- 4 The third valid values are output to the GDOUT[63~0] signal. The address values that are being output to the GA[35~0] signal are output in increments of eight. Since it is acknowledged that the GACK* signal is "L," next data are output at the timing of Item 5. Also, the GLAST* signal is asserted.
- 5 The final valid values are output to the GDOUT[63~0] signal. The address values that are being output to the GA[35~0] signal are output in increments of eight. Since it is acknowledged that the GACK* signal is "L," the GLAST* signal is deasserted.
- 6 The GWR* signal and the GBURST* signal are deasserted. This completes the burst write operation. Up to this timing, valid values are output to the GA[35~0] signal, the GBE[7~0]* signal, the GDOUT[63~0] signal, and the GID signal.

If a single write operation follows immediately after the burst write operation, the GWR* signal is not deasserted, so the single write operation starts. Also, if another burst write operation follows immediately after this burst write operation, neither the GBURST* signal nor the GWR* signal is deasserted, so the next burst write operation starts.

Therefore, by checking the GACK* signal for the data to be written next, the second and subsequent data can be written in one cycle. The addresses on the GA[35~0] signal always indicate the address of the data that are being write. Since the TX49 Megacell starts a burst write from the word of the lowest-order address in the refill boundary, the address in a burst write operation is simply output in increments of eight. The timing to assert the GLAST* signal is when the assertion of the GACK* signal for the data write second from the last is acknowledged. The timing to deassert the GLAST* signal is when the assertion of the GACK* signal for the last data write is acknowledged.

Similar to the single write operation, wait cycles can be inserted by deasserting the GACK* signal. Figure 16-2 shows a burst write operation with wait cycles. This is an example where wait cycles are inserted at the first and the third data write in a burst write operation.

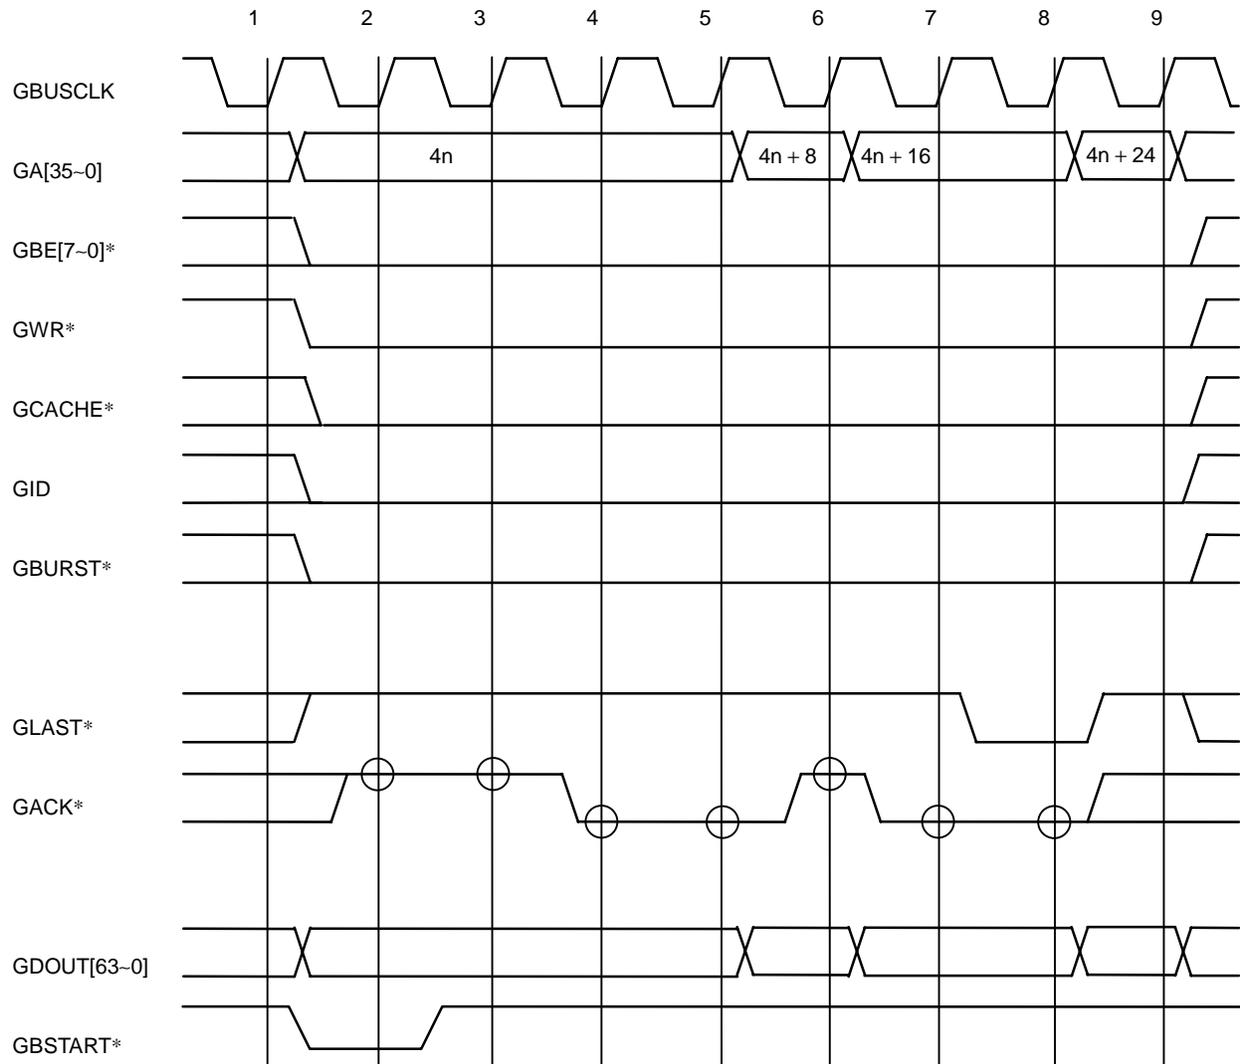


Figure 16-2 Burst write Operation with Wait Cycles

- 1 The GBSTART* signal, GWR* signal and the GBURST* signal are asserted, and a burst write operation starts. At the same time, valid values are output to the GA[35~0] signal, the GDOUT[63~0] signal, the GBE[7~0]* signal, the GCACHE* signal, the GID signal.
- 2 The GBSTART* signal is deasserted. Since the GACK* signal is "H," next data are not output at the timing of Item 3.
- 3 Since the GACK* signal is "H," data are not output at the timing of Item 4.
- 4 Since it is acknowledged that the GACK* signal is "L," next data are output at the timing of Item 5.
- 5 The second valid values are output to the GDOUT[63~0] signal. The address values that are being output to the GA[35~0] signal are output in increments of eight. Since it is acknowledged that the GACK* signal is "L," next data are output at the timing of Item 4.
- 6 The third valid values are output to the GDOUT[63~0] signal. The address values that are output on the GA[35~0] signal are output in increments of eight. Since it is acknowledged that the GACK* signal is "H," data are not output at the timing of Item 7.
- 7 Since it is acknowledged that the GACK* signal is "L," next data are output at the timing of Item 8. Also, the GLAST* signal is asserted.
- 8 The final valid values are output to the GDOUT[63~0] signal. The address values that are being output to the GA[35~0] signal are output in increments of eight. Since it is acknowledged that the GACK* signal is "L," the GLAST* signal is deasserted.
- 9 The GWR* signal and the GBURST* signal are deasserted. This completes the burst write operation. Up to this timing, valid values are output to the GA[35~0] signal, the GBE[7~0]* signal, the GDOUT[63~0] signal, and the GID signal.

The TX49 Megacell has an on-chip write buffer so that errors cannot be notified by the GBUSERR* signal.

When an error is caused in the system by the write operation of the TX49 Megacell, a peripheral circuit asserts the GACK* signal to suspend the write operation. Then, use an interrupt signal such as GINT[5~0]* to inform the TX49 Megacell of the error.

16.5 Signal status while no bus operation is executed

When no bus operation is being performed, the TX49 Megacell sets the memory interface output signals to the following states:

GA[35~0]	Indefinite
GDOUT[63~0]	Indefinite
GBE[7~0]*	All H
GRD*	H
GWR*	H
GCACHE*	H
GID	H
GBURST*	H
GBSTART*	H
GLAST*	H

The memory interface output signals during a reset are also in these states.

17. Bus Arbitration and Snoop Operation

17.1 Signals to be used in bus arbitration

The bus ownership transferring protocol between the TX49 Megacell and the external bus master is performed by ten signals: four bus request signals which are the GREQ* signal, the GSREQ* signal, the GHPGREQ* signal and the GHPSREQ* signal; four bus release signals that are the GGNT* signal, the GSGNT* signal, the GHPGGNT* signal and the GHPSGNT* signal; the bus ownership signal that is the GHAVEIT* signal; and the bus release request signal that is the GREL* signal. The external bus master can request the bus ownership using the GREQ* signal or the GHPGREQ* signal when in ET concurrency, and can request ownership using the GSREQ* signal or the GHPSREQ* signal when in ST concurrency. (For ET concurrency and ST concurrency, please see “Appendix E : Internal Bus Architecture”) In addition to the bus request signals, the TX49 Megacell has a cycle steal signal inside to retrieve the bus ownership that is surrendered to the external bus master.

The priority of each bus request signal is shown below in Table 17-1.

Table 17-1 Priority of Bus Request Signals

Bus Request Signal	Priority	Generation of Cycle Steal
GHPSREQ*	High	No
GHPGREQ*		Yes
GSREQ*		No
GREQ*	Low	Yes

The arbitration concerning the release and granting of the bus ownership is controlled by the G-Bus arbiter inside the TX49 Megacell.

When the TX49 Megacell releases the bus ownership, a bus release signal (one from among the GGNT* signal, the GSGNT* signal, the GHPGGNT* signal, or the GHPSGNT* signal) is asserted, and the memory interface output signals (the GA[35~0] signal, the GDOUT[63~0] signal, the GBE[7~0]* signal, the GRD* signal, the GWR* signal, the GLAST* signal, the GBURST* signal, and the GBSTART* signal) are set to Hi-Z (high impedance). After the priority is determined, the bus master with the highest priority obtains the G-Bus ownership, and asserts the GHAVEIT* signal to start driving the memory interface output signals. When the bus master completes the data transfer, it deasserts the GHAVEIT* signal, informs the TX49 Megacell of the completion, and changes its own memory interface output signal to Hi-Z (high impedance). The TX49 Megacell deasserts the bus release signal to resume its own bus operation. (See Figure 17-1)

The above is summarized as follows:

- The TX49 Megacell deasserts a bus release signal to drive the memory interface output signals. (Cycles 1, 7 and 8)
- When the GHAVEIT* signal is asserted, the bus master can drive the memory interface output signals. (Cycles 3, 4 and 5)
- When the current bus ownership changes to another, at least one dead cycle is inserted. The dead cycles are to prevent transient signal conflicts on the bus between the bus drivers while transferring. (Cycles 2 and 6)

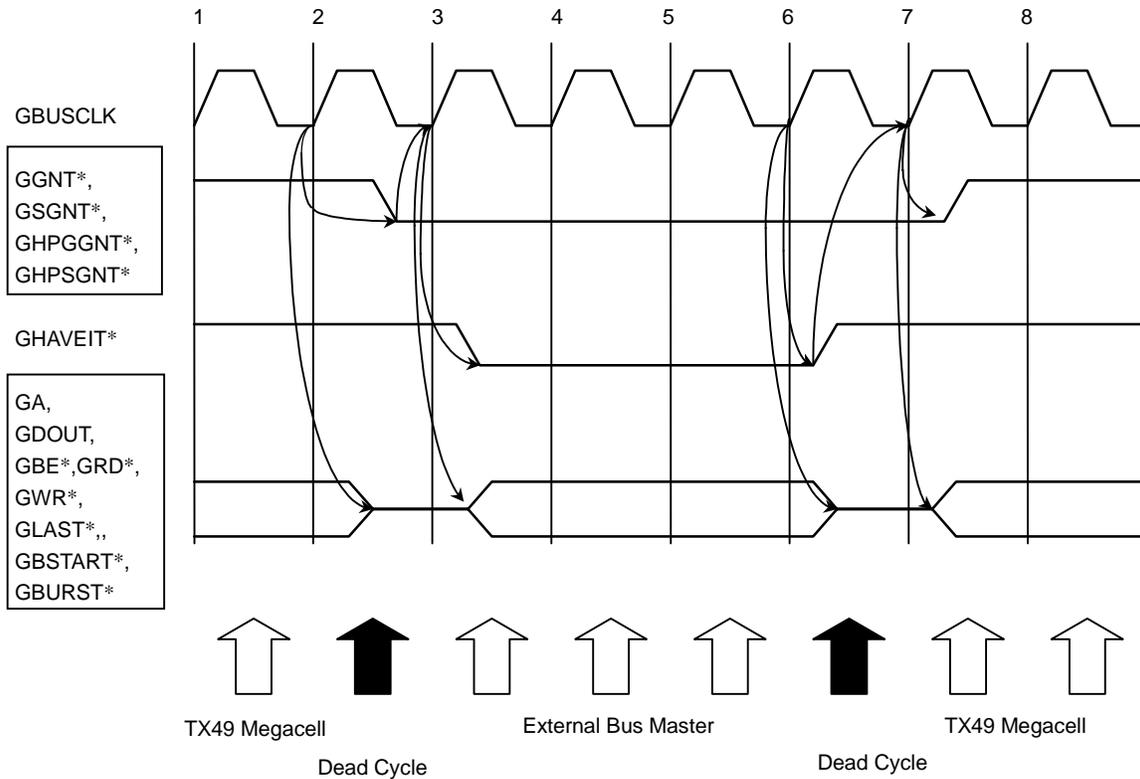


Figure 17-1 Hi-Z Timing of G-Bus

17.1.1 Bus request signals and bus release signals

The bus request signals are for the external bus master to request TX49 Megacell for the release of the bus ownership. The bus release signals are for the TX49 Megacell to inform the external bus master of the release of the bus ownership. There are four kinds of bus request signals and bus release signals with low/high priority for ET concurrency and ST concurrency as shown below.

Bus Request Signal	Bus Release Signal		
GREQ*	GGNT*	(Low priority)] ET Concurrency
GHPGREQ*	GHPGGNT*	(High priority)	
GSREQ*	SGNT*	(Low priority)] ST Concurrency
GHPSREQ*	GHPSGNT*	(High priority)	

Figure 17-1 shows an example where the TX49 Megacell releases the bus ownership using the GREQ* signal (ET Concurrency).

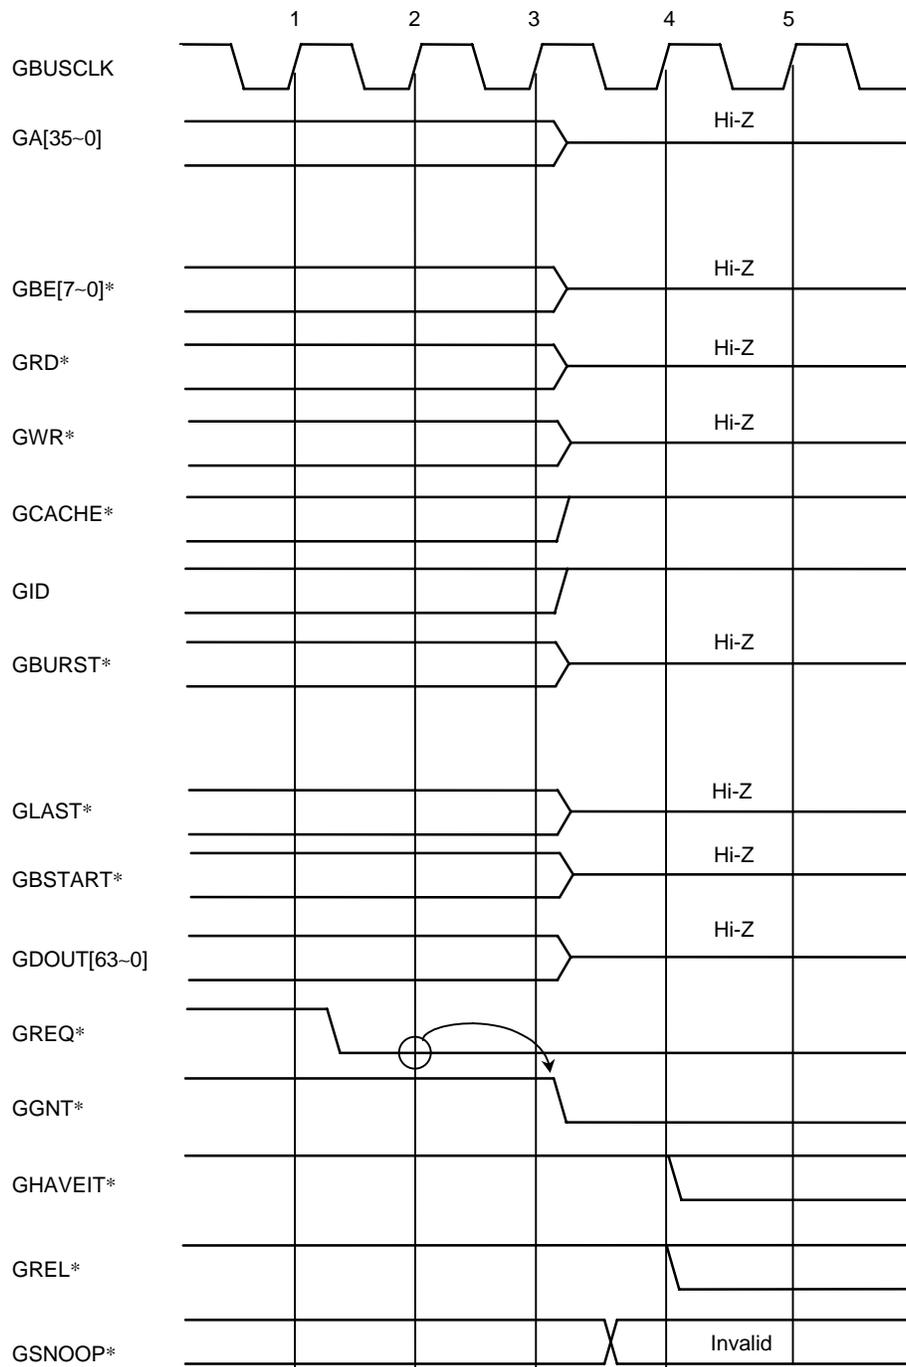


Figure 17-1 Bus Ownership Release by GREQ*

- 1 The deassertion of the GREQ* signal is acknowledged.
- 2 The assertion of the GREQ* signal is acknowledged.
- 3 The GGNT* signal is asserted, and bus ownership is released. This cycle is a dead cycle in which neither the TX49 Megacell nor the external bus master drives the memory interface output signals.
- 4 The external bus master asserts the GHAVEIT* signal so that it can drive the memory interface output signals. (In the above example, the external bus master does not drive the memory interface signals and stays at Hi-Z.) Note that the GSNOOP* signal is ignored if control of the bus is relinquished by GREQ*.

However, in order for the TX49 Megacell to release the bus ownership, it is necessary to make sure that a bus operation is not being executed in the timing when the GREQ* signal is acknowledged, and that a bus request signal with higher priority is not asserted.

If the TX49 Megacell is executing a bus operation, the assertion of the GREQ* signal can be acknowledged from the cycle prior to the completion of the bus operation. In this case, the TX49 Megacell asserts the GGNT* signal at the GBUSCLK rise where the bus operation ends. The TX49 Megacell sets the output signals to the following states after releasing the bus ownership.

Caution: The TX49 Megacell does not have the bus holder internally.

Table 17-1 States at Release of Bus Ownership

Name of Pin	I/O	At Bus Release
GA[35~0]	I/O	Hi-Z
GDOUT[63~0]	O	Hi-Z
GDIN[63~0]	I	□
GBE[7~0]*	O	Hi-Z
GRD*	O	Hi-Z
GWR*	O	Hi-Z
GACK*	I	□
GCACHE*	O	"H"
GID	O	"H"
GLAST*	O	Hi-Z
GBUSERR*	I	□
GBSTART*	O	Hi-Z
GBURST*	O	Hi-Z

Figure 17-2 shows an example where the external bus master releases the bus ownership and the TX49 Megacell regains it.

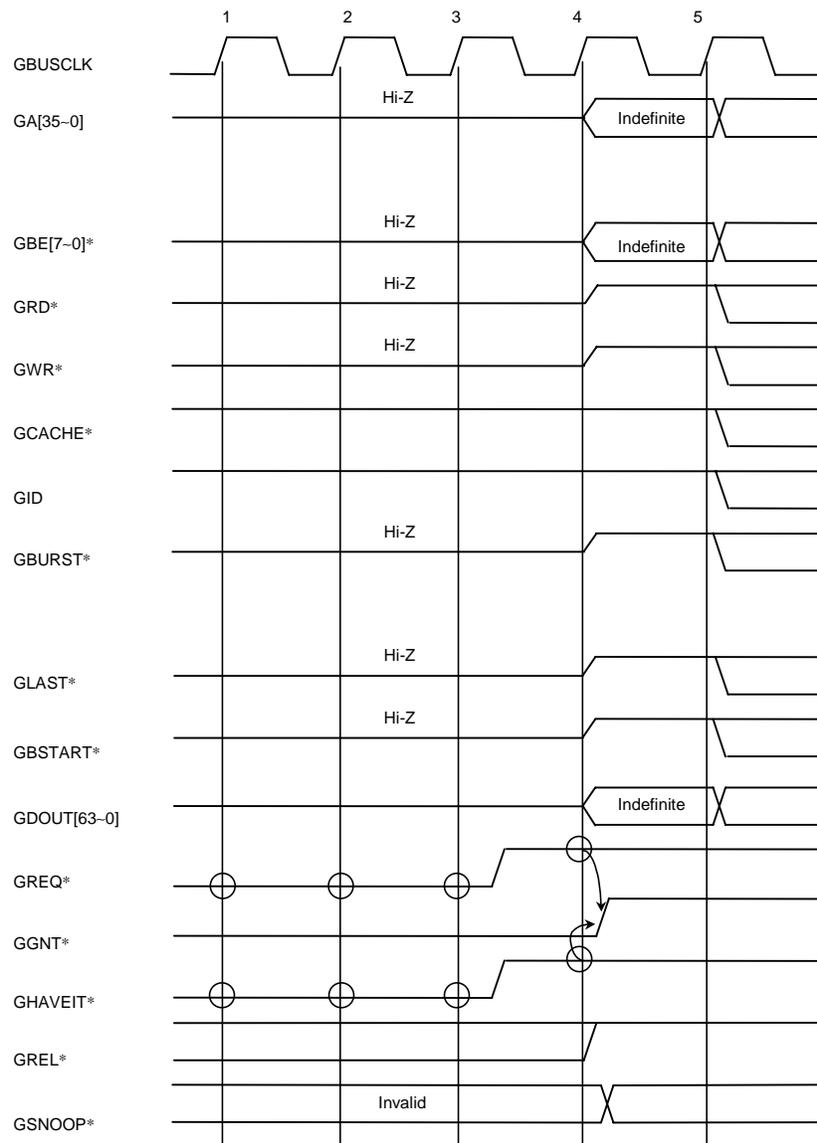


Figure 17-2 Bus Ownership Obtained from GREQ*

- 1 The assertion of the GREQ* signal is acknowledged.
- 2 The assertion of the GREQ* signal is acknowledged.
- 3 The assertion of the GREQ* signal is acknowledged.
- 4 The deassertion of the GREQ* signal and the GHAVEIT* signal are acknowledged. Having deasserted the GGNT* signal, the TX49 Megacell obtains the bus ownership.
- 5 Bus operations by the TX49 Megacell become available.

The TX49 Megacell in ET concurrency (GREQ* and GHPGREQ*) can write to the write buffer even when the bus ownership is released since only the apparent bus ownership of the G-Bus is released in ET concurrency. The internal pipeline also keeps operating, but no snoop operation can be performed. When a bus operation becomes necessary due to a reason such as a cache miss, the pipeline of the TX49 Processor Core is stalled, a cycle steal is generated, and the bus release request GREL* signal is asserted. The TX49 waits to obtain the bus ownership again.

On the other hand, the TX49 Megacell in ST concurrency (GSREQ* and GHPSREQ*) releases the bus ownership to enable a snoop function. It releases not only the ownership on the G-Bus, but also the X-Bus, to the external bus master.

If there is a bus ownership request by GSREQ* or GHPSREQ*, the TX49 Megacell writes all the data from the write buffer to external memory, and it asserts the GSGNT* or GHPSGNT* to release the bus ownership.

If the GCOLDRESET* signal is asserted while the TX49 Megacell uses a bus request signal to release the bus ownership, the TX49 Megacell forcibly deasserts a bus release signal and performs a reset operation.

When the TX49 Megacell is in the doze mode, it replies to a bus ownership request from the outside. In such a case, the TX49 Megacell asserts a bus release signal in the timing after the assertion of the bus request signal is acknowledged. On the other hand, when the TX49 Megacell is in the halt mode, it cannot reply to any bus ownership requests. It does not release the bus ownership until it has left the halt mode by the assertion of an external interrupt signal (the GNMI* signal or the GINT[5~0]* signal).

17.1.2 Cycle stealing

If a situation occurs where the TX49 Megacell needs to execute a bus operation (such as a read operation), during ET concurrency in which bus ownership was gained by the external master on the G-Bus using the GREQ* signal or GHPGREQ* signal, or if a situation occurs in which the write buffer is too full and must be flushed, the G-Bus arbiter determines that the G-Bus is in the cycle steal state. Then, a GREL* signal is asserted to get G-Bus ownership back from the external bus master.

Having sensed the assertion of the GREL* signal, the external bus master deasserts the GHAVEIT* signal and releases the G-Bus ownership.

It is possible to make an external bus master which does not permit cycle stealing by ignoring the GREL* signal. In such a case, the external bus master can halt the bus operation of the TX49 Megacell until releasing the G-Bus. After the G-Bus arbiter gets the G-Bus ownership back, the TX49 Megacell executes its bus operation.

When the bus operation by the TX49 Megacell is completed, the G-Bus arbiter starts receiving a new bus release request signal for the G-Bus. Then, the bus ownership is released to the bus master which asserted the bus request signal with the highest priority at that point.

A cycle steal does not occur in ST concurrency.

17.1.3 GREL* signal

The bus release request GREL* signal is a signal which prompts the bus master that currently owns the G-Bus to release the bus. This signal is asserted in the following cases:

If a bus request signal with higher priority is asserted while the bus ownership is released to the external bus master.

If a cycle steal request (Cycle Steal) occurs in the TX49 Megacell during an ET concurrency (while the GGNT* signal or the GHPGGNT* signal is being asserted).

When the external bus master acknowledges the assertion of the GREL* signal, the bus can be released to the TX49 Megacell or another bus master by deasserting the GHAVEIT* signal and a bus request signal. When the bus master which used to own the G-Bus ownership needs to retrieve it, the bus master can keep asserting a bus request signal.

Table 17-1 summarizes the current owner of the bus ownership (indicated by the bus release signals) and the corresponding bus request signals which causes the assertion of the GREL* signal.

Table 17-1 Conditions of GREL* Signal Assertion

Signal Currently Asserted	Signal to Cause GREL* Signal Assertion
GGNT*	GSREQ*
GHPGREQ*	GHPGREQ*
(Cycle Steal)	GSGNT*
GHPGREQ*	GHPGREQ*
GHPGGNT*	GHPGREQ* (Cycle Steal)
GHPGGNT*	□

17.1.4 GHAVEIT* signal

The bus ownership GHAVEIT* signal indicates that the external bus master possesses the bus ownership (drives G-bus). When the external bus master asserts a bus request signal (xxREQ* signal), the TX49 Megacell asserts a bus releases signal (xxGNT* signal), and release the bus ownership, the external bus master asserts the GHAVEIT* signal so that it can drive the memory interface signals. The external bus master then stops driving the memory interface signals after the bus operation is complete, and can return the bus ownership to the TX49 Megacell by deasserting the GHAVEIT* signal. Having acknowledged the deassertion of the GHAVEIT* signal at a rise of the GBUSCLK signal, the TX49 Megacell deasserts the bus release signal (xxGNT* signal).

If the external bus master that possesses the current bus ownership acknowledges the assertion of the bus release request signal (GREL* signal) by a cycle steal request from another external bus master with higher priority or from the TX49 Megacell, it can surrender the bus ownership to the bus master with higher priority by deasserting the GHAVEIT* signal. In this case, if the external bus master with lower priority needs to regain the bus ownership, it can continue asserting a bus request signal (xxREQ* signal).

For a system that consists of multiple bus masters, the bus masters can know if another bus master possesses the bus ownership by mutually observing the GHAVEIT* signal. By doing so, it is possible to know, by an external bus master (which has asserted a bus request signal) observing the assertion of the GHAVEIT* signal, whether another external bus master took bus ownership. It is then possible to realize functions such as a bus master withdrawing its own bus request or asserting another bus request signal with higher priority.

Note on the GHAVEIT* signal

The GHAVEIT* signal indicates that an external bus master is driving the G-Bus. Make sure it is asserted as long as the external bus master is driving the G-Bus. Therefore, we recommend using an output enable signal of the external bus master as the GHAVEIT* signal.

When deasserting the GREQ* signal, always be sure to deassert the GHAVEIT* signal also.

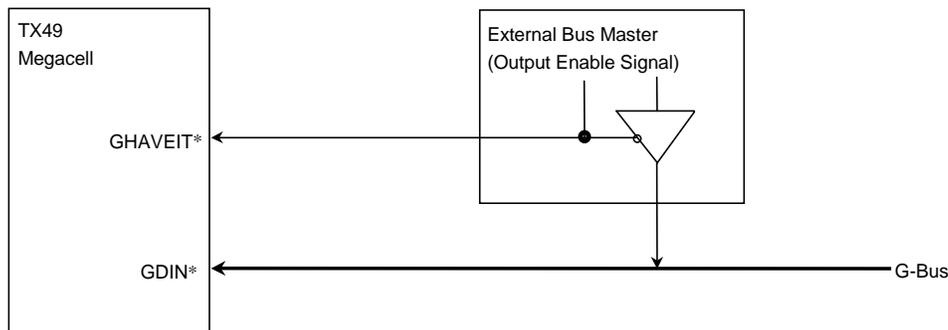


Figure 17-1 GHAVEIT* Usage Example

17.2 Operation timing of bus arbitration

17.2.1 Single bus master

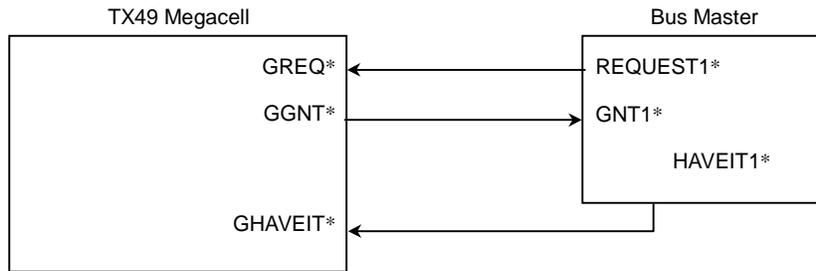


Figure 17-1 Single Bus Master Connection

Figure 17-1 is a connection example of the TX49 Megacell and one external bus master.

This example shows a case where the external bus master is requesting G-Bus ownership using the GREQ* signal. The REQUEST1* signal is the bus request signal of the external bus master. Figure 17-2 is the timing diagram.

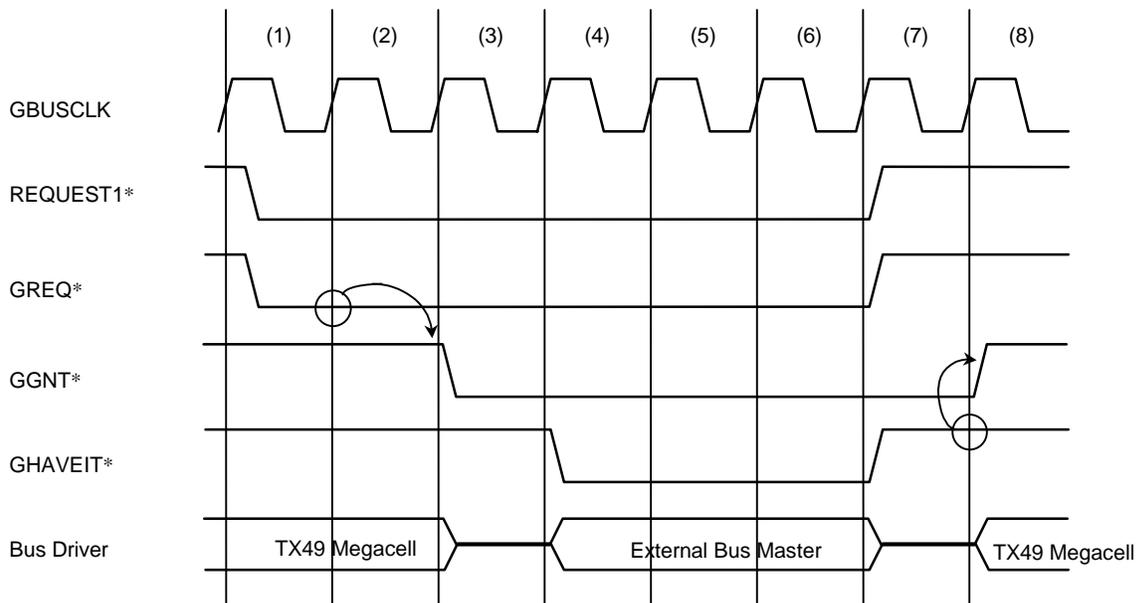


Figure 17-2 Single Bus Master Timing

- Cycle 1: The external bus master asserts $REQUEST1^*$ to request G-Bus ownership.
- Cycle 2: Having acknowledged the assertion of the $GREQ^*$ signal at a rise of $GBUSCLK$, the G-Bus arbiter of the TX49 Megacell implements bus arbitration.
- Cycle 3: The TX49 Megacell asserts the $GGNT^*$ signal to release G-Bus ownership. (The $GGNT^*$ signal is asserted in the cycle immediately after it is acknowledged that the $GREQ^*$ signal is "L." However, assertion of the $GGNT^*$ signal is delayed while a bus cycle is being executed by the TX49 Megacell.)
- Cycle 4: The external bus master asserts the $GHAVEIT^*$ signal, and processing starts. It is transfer cycle 1 of the external bus master.
- Cycle 5: It is transfer cycle 2 of the external bus master.
- Cycle 6: It is transfer cycle 3 of the external bus master.
- Cycle 7: The external bus master deasserts the $GHAVEIT^*$ signal. The $REQUEST1^*$ signal and the $GREQ^*$ signal are similarly deasserted in the same cycle. During this period, neither the external bus master nor the TX49 Megacell outputs any signal, so it becomes a dead cycle.
- Cycle 8: The TX49 Megacell deasserts the $GGNT^*$ signal and it becomes the owner of the G-Bus. (The $GGNT^*$ signal is deasserted in the cycle where it is acknowledged that the $GHAVEIT^*$ signal or the $GREQ^*$ signal is "H.")

17.2.2 Daisy chain

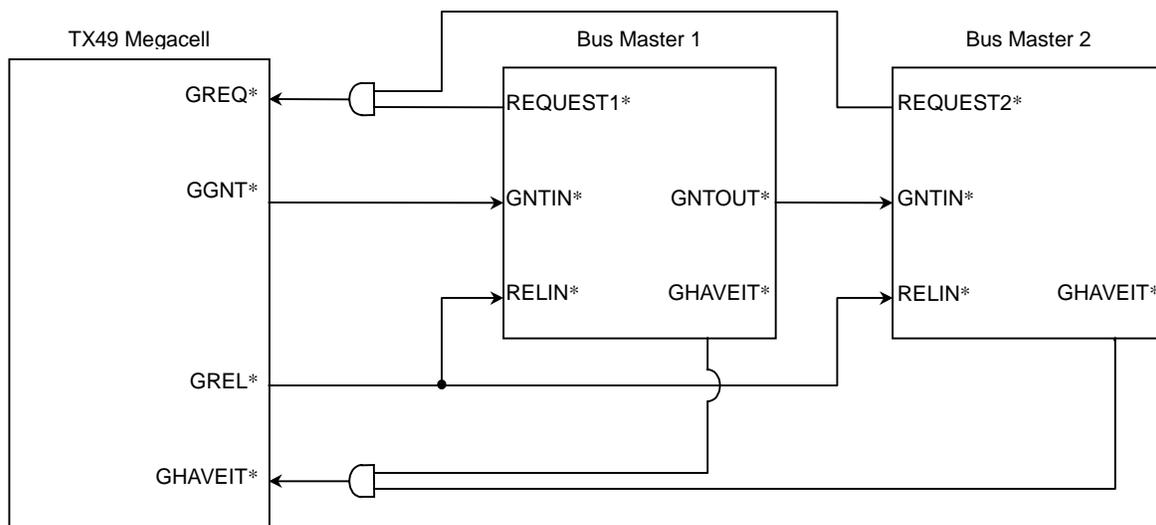


Figure 17-1 Daisy Chain Connection

Figure 17-1 is an example connection where a system is configured with two daisy-chained external bus masters: --Bus Master 1 and Bus Master 2. The $REQUEST1^*$ signal and the $REQUEST2^*$ signal are the bus request signals of the two bus masters. Bus Master 1 is on the upstream side of the daisy-chain and Bus Master 2 is on its downstream side. Figure 17-2 shows the timing diagram.

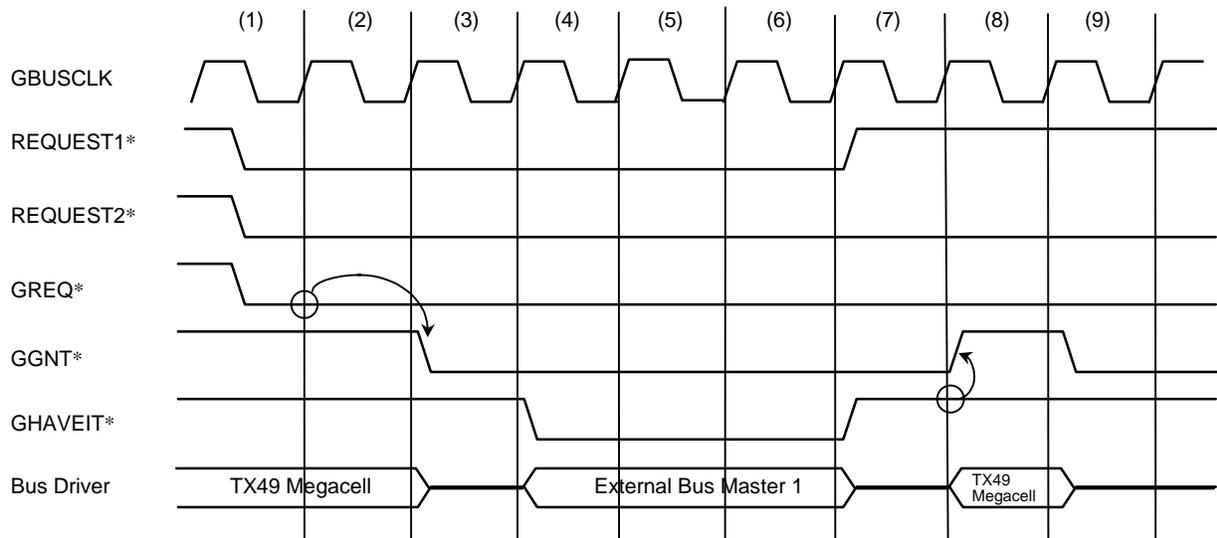


Figure 17-2 Daisy-Chain Timing

- Cycle 1: Both bus masters simultaneously assert the REQUEST1* signal and the REQUEST2* signal, and request G-Bus ownership.
- Cycle 2: Assertion of GREQ* is acknowledged at a rise of GBUSCLK. The G-Bus arbiter of the TX49 Megacell implements a bus arbitration.
- Cycle 3: The TX49 Megacell asserts the GGNT* signal, and releases G-Bus ownership. No signal is being output from the TX49 Megacell or Bus Master 1 during this period, so this becomes a dead cycle.
- Cycle 4: Bus Master 1 obtains G-Bus ownership, asserts the GHAVEIT* signal, and starts processing. Bus Master 2 continues asserting the REQUEST2* signal to obtain G-Bus ownership. This is transfer cycle 1 of Bus Master 1.
- Cycle 5: This is transfer cycle 2 of Bus Master 1.
- Cycle 6: This is transfer cycle 3 of Bus Master 1.
- Cycle 7: Bus Master 1 deasserts the GHAVEIT* signal and the REQUEST1* signal. No signal is being output from Bus Master 1 or the TX49 Megacell during this period, so it becomes a dead cycle. The GREQ* signal is the logical sum of the REQUEST1* signal and the REQUEST2* signal. Therefore, the GREQ* signal is being asserted even when the REQUEST1* signal has been deasserted.
- Cycle 8: The G-Bus arbiter acknowledges the GHAVEIT* signal as being "H" at a rise of GBUSCLK and deasserts the GGNT* signal. In addition, the TX49 Megacell possesses G-Bus ownership for the time being.
- Cycle 9: Since the GREQ* signal stays at "L" due to the REQUEST2* signal of the bus master, the G-Bus arbiter of the TX49 Megacell implements an arbitration, and asserts the GGNT* signals, releasing G-Bus ownership. No signal is output from the TX49 Megacell or the bus masters during this period, so it becomes a dead cycle.

17.2.3 Bus ownership release by GREL* signal

Figure 17-1 shows a situation where the external bus master abandons G-Bus ownership as a result of the TX49 Megacell asserting the GREL* signal in the system of the daisy-chain connection in example Figure 17-1. The GREL* signal is a bus release request signal of the TX49 Megacell.

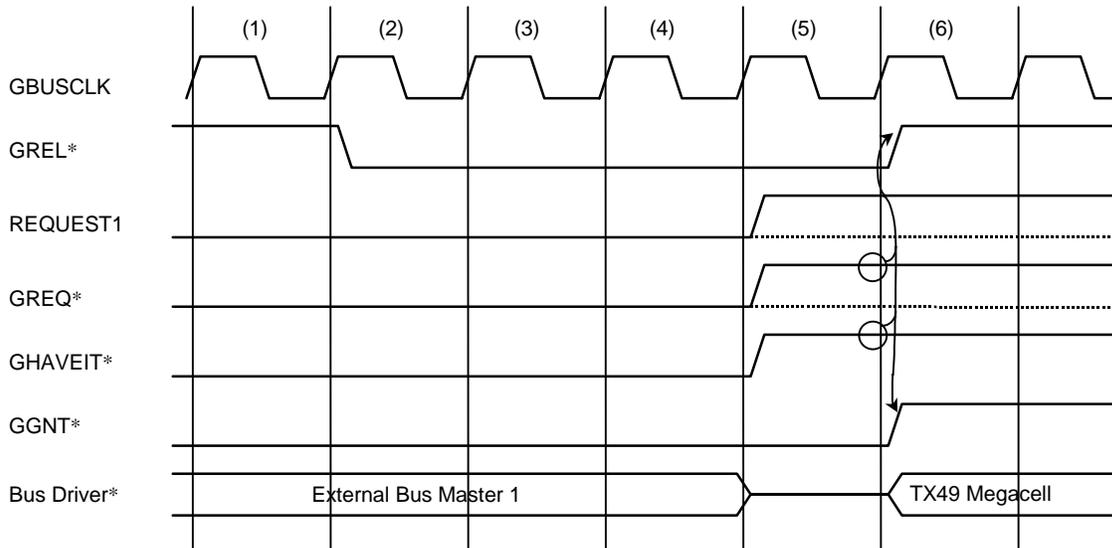


Figure 17-1 GREL* Signal Timing

- Cycle 1: The G-Bus ownership is surrendered to Bus Master 1 by the REQUEST1* signal from Bus Master 1.
- Cycle 2: The GREL* signal is asserted from the TX49 Megacell.
- Cycle 3: Having acknowledged the GREL* signal, Bus Master 1 implements completion operation of the data transfer.
- Cycle 4: Bus Master 1 is continuing the completion operation of the data transfer.
- Cycle 5: Bus Master 1 deasserts the GHAVEIT* signal. The REQUEST1* signal (GREL* signal) is similarly deasserted in the same cycle. If Bus Master 1 needs to regain the G-Bus which was released by the GREL* signal, it can keep asserting the REQUEST1* signal (GREQ* signal) instead of deasserting it. (See the broken lines in the diagram.)
- Cycle 6: The G-Bus arbiter deasserts the GGNT* signal and the GREL* signal. Since G-Bus ownership was granted to the TX49 Megacell by the GREL* signal, a bus cycle of the TX49 Megacell starts.

In this case, the GREL* signal is asserted in two situations: when the TX49 Megacell requests a bus cycle in an ET concurrency, or when a bus request signal is asserted by another external bus master with higher priority.

In Figure 17-1, the GHAVEIT* signal is deasserted by the GREL* signal, and G-Bus ownership returns to the TX49 Megacell. However, the G-Bus arbiter gives G-Bus ownership to either the TX49 Megacell which was the source of the assertion of the GREL* signal or to an external bus master with higher priority in the next arbitration.

17.2.4 Shift of bus ownership (Low priority to High priority)

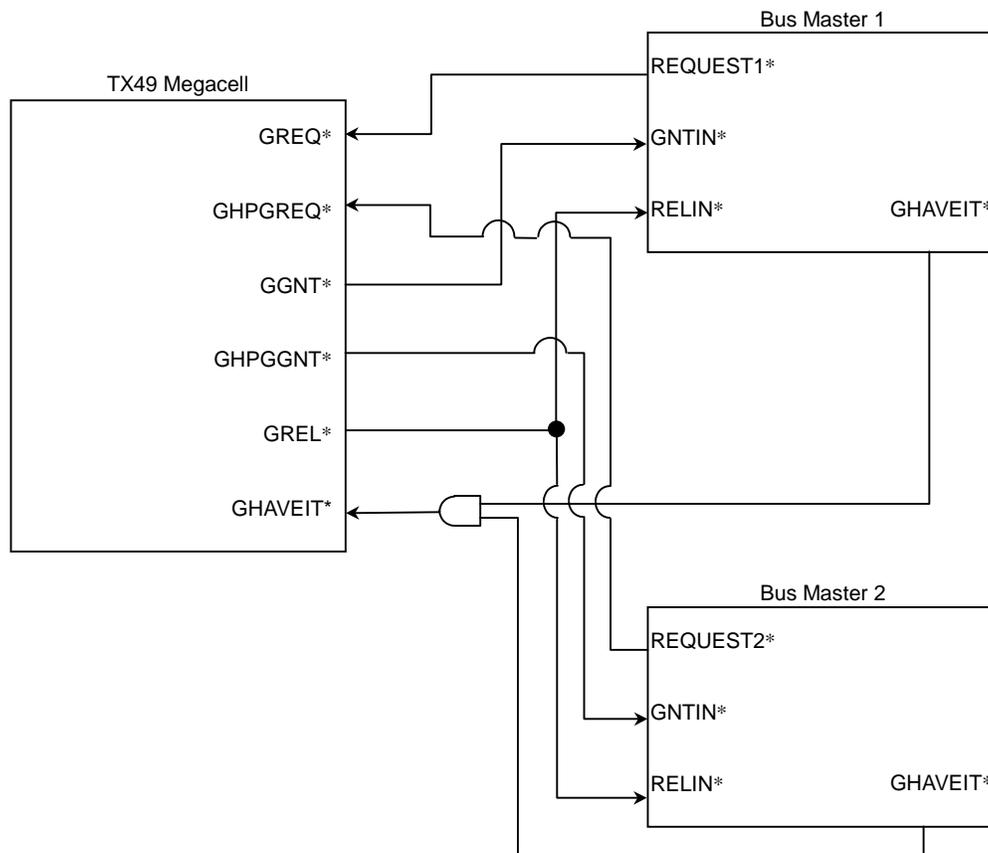


Figure 17-1 Multi-Bus-Master Connection

Figure 17-1 is a connection example of a system configured with two external bus masters: Bus Master 1 and Bus Master 2. Bus Master 1 uses the GREQ* signal and the GGNT* signal, and Bus Master 2 uses the GHPGREQ* signal and the GHPGGNT* signal to perform transfers of G-Bus ownership.

When a bus request signal (GHPGREQ* signal) is asserted from Bus Master 2, whose priority is higher while Bus Master 1 has G-Bus ownership, Bus Master 1 releases G-Bus ownership in response to the GREL* signal the Bus Master 2 then obtains G-Bus ownership. This operation is shown in Figure 17-2. See Table 17-1 for the priorities.

- Cycle 1: A bus request signal (GHPGREQ* signal) is asserted by Bus Master 2 during a bus arbitration by Bus Master 1.
- Cycle 2: The “L” of the GHPGREQ* signal is acknowledged at a rise of GBUSCLK and the GREL* signal is asserted.
- Cycle 3: Bus Master 1 acknowledges the GREL* signal and releases the G-Bus ownership by deasserting the GHAVEIT* signal and the GREQ* signal.
- Cycle 4: “H” of the GHAVEIT* signal is acknowledged at a rise of GBUSCLK and the GGNT* signal and the GREL* signal are deasserted. In this cycle, G-Bus ownership is momentarily returned to the TX49 Megacell and the TX49 Megacell drives the G-Bus.

- Cycle 5: "L" of the GHPGREQ* signal is acknowledged and the GHPGGNT* signal is asserted. This cycle becomes a dead cycle.
- Cycle 6: Bus Master 2 asserts the GHAVEIT* signal, and the processing starts process. This is transfer cycle 1 of Bus Master 2.
- Cycle 7: This is the transfer cycle 2 of Bus Master 2.

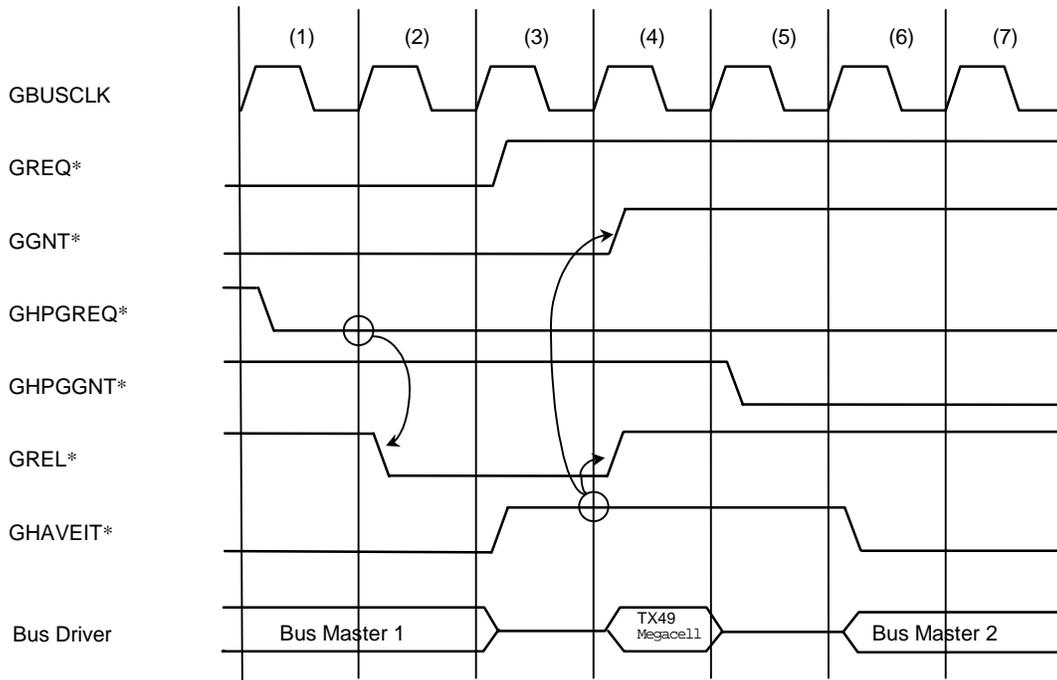


Figure 17-2 Bus Arbitration 4

In the example shown in Figure 17-2, it takes three cycles including a dead cycle to shift from the status where the G-Bus ownership is released by the GREQ* signal to the status where the G-Bus ownership is released by the GHPGREQ* signal with higher priority. Table 17-1 summarizes the minimum number of cycles for the G-Bus ownership to shift from a lower priority to a higher priority.

Table 17-1 Number of Cycles Necessary for Bus Ownership Shift

Shift of G-Bus Ownership	Min. Number of Cycles
GREQ* → GSREQ*	3• Cycles
GREQ* → GHPGREQ*	3• Cycles
GREQ* → GHPSREQ*	3• Cycles
GSREQ* → GHPGREQ*	3• Cycles
GSREQ* → GHPSREQ*	3• Cycles
GHPGREQ* → GHPSREQ*	3• Cycles

17.2.5 General rules

Following are general rules of bus arbitration.

- 1• An external bus master which needs to obtain ownership of the G-Bus or the X-Bus must assert an appropriate bus request signal.
2. When multiple bus request signals are asserted while the G-Bus is not being released to any of the external bus masters (the GHAVEIT* is being deasserted), the G-Bus arbiter of the TX49 Megacell surrenders G-Bus ownership to the bus master which is asserting the bus request signal with the highest priority.
3. The bus master which has the G-Bus ownership deasserts the GHAVEIT* signal in the final processing cycle and deasserts the appropriate bus request signal as well.

Note, however, that the GHAVEIT* signal must always be deasserted when the GREQ* signal is deasserted. Conversely, the GREQ* signal does not always have to be deasserted when the GHAVEIT* signal is deasserted.

4. In the cycle following the deassertion of the GHAVEIT* signal, the G-Bus arbiter of the TX49 Megacell deasserts the bus release signal which is being asserted. If there is no other request signal being asserted by another bus by the beginning of that cycle, G-Bus ownership is returned to the TX49 Processor Core which has the G-Bus arbiter. On the other hand, if bus request signals are asserted, the G-Bus ownership is released to the bus master which is asserting the bus request signal with the highest priority.

17.2.6 State machine

Table 17-1 below indicates situations in which G-Bus ownership is transferred from the current bus master to a new bus master by the assertion of a request signal. Bus masters are represented by the bus release signals. Current bus masters can be found in the “Current Bus Master” column. New bus masters are in the “Bus Master after Shift.” column.

For example, the top row of “GGNT*” can be interpreted as follows: if the GHPSREQ* signal is asserted by another bus master while the current bus master owns the G-Bus ownership according to GGNT*, the GREL* signal is asserted to shift ownership to the bus master with the GHPSREQ*.

Table 17-1 State Machine of Bus Arbitration

Bus Request Signal to be Asserted Current Bus Master	GHPSREQ*	Cycle Steal	GHPGREQ*	GSREQ*	GREQ*	GREL*	Bus Master After Shift
GGNT*	0	X	X	X	0	0	GHPGNT*
GGNT*	1	0	X	X	0	0	TX49 Megacell
GGNT*	1	1	0	X	0	0	GHPGGNT*
GGNT*	1	1	1	0	0	0	GSGNT*
GGNT*	1	1	1	1	0	1	GGNT*
GSGNT*	0	1	X	0	X	0	GHPGNT*
GSGNT*	1	1	0	0	X	0	GHPGGNT*
GSGNT*	1	1	1	0	X	1	GSGNT*
GHPGGNT*	0	X	0	X	X	0	GHPGNT*
GHPGGNT*	1	0	0	X	X	0	TX49 Megacell
GHPGGNT*	1	1	0	X	X	1	GHPGGNT*
GHPGNT*	0	1	X	X	X	1	GHPGNT*

X : Don't care 0: Assert 1: Deassert

17.3 Snoop operation

The TX49 Megacell has a snoop function. This function is used to maintain coherency between the data cache and the memory in the TX49 Processor Core during ST concurrency. This function uses the GSNOOP* signal and the GA[35~0] signal of the TX49 Megacell. The TX49 Megacell detects the GSNOOP* signal on a rising edge of the GBUSCLK signal. When either the TX49 Megacell has released both the G-Bus and the X-Bus; namely, when the GSGNT* signal or the GHPSGNT* signal has been asserted, the GSNOOP* signal and the GA[35~0] signal become valid. If the GSNOOP* signal is asserted at this time, the TX49 Megacell executes a snoop operation. If GA[35~5] hits an entry of the data cache, the entry becomes invalid.

In order to execute a snoop operation, all address output of the external bus master should be connected to GA[35~0]. The GSNOOP* signal should be asserted when the external bus master that has obtained the G-Bus ownership executes a write cycle. The address(es) to be input to GA[35~0] alone determine the address for the snoop execution, so it is necessary to supply correct write address(es) to the TX49 Megacell. Figure 17-1 shows the snoop operation.

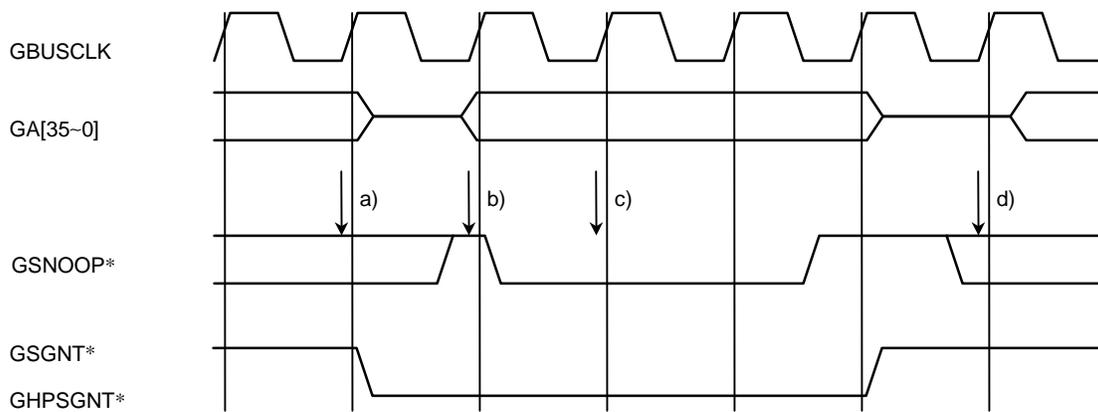


Figure 17-1 Snoop Operation

- Since the GSGNT* signal or the GHPSGNT* signal is deasserted, the GSNOOP* signal and the GA[35~0] signal are invalid.
- Since the GSNOOP* signal is deasserted, the snoop operation is not executed.
- The GSNOOP* signal and the GA[35~5] signal become valid and the snoop operation is executed. If the GA[35~0] signal hits the data cache, the entry becomes invalid.
- Since the GSGNT* signal or the GHPSGNT* signal is deasserted, the GSNOOP* signal and the GA[35~0] signal are invalid.

18. Coprocessor Operations

In the MIPS architecture, four coprocessors can be installed: CP0, CP1, CP2, and CP3. CP0 is reserved as the system control coprocessor which realizes the exception processing function, the system configuration function, and the memory management function. The TX49 Megacell has an on-chip CP0 coprocessor and CP1 is reserved for the internal Floating-Point Unit. TX49 does not support CP3 except coprocessor condition branch instruction operation.

The coprocessor operation uses the following memory interface signals in addition to the coprocessor interface signals: the GA[35~0] signal, the GDOUT[63~0] signal and the GDIN[63~0] signal. The output signals other than the above maintain the status from when no bus operation was executed. The GACK* signal and the GBUSERR* signal are ignored.

18.1 Coprocessor read operation

The coprocessor read operation is generated by the execution of the MFCz instruction, the CFCz instruction, and the COPz instruction. The timing of the coprocessor read operation is the same as the single read operation for the memory. The coprocessor read operation uses the GA[35~0] signal, the GDIN[63~0] signal, the GCPRD* signal, and the GCPRDACK* signal. Figure 18-1 shows the coprocessor read operation.

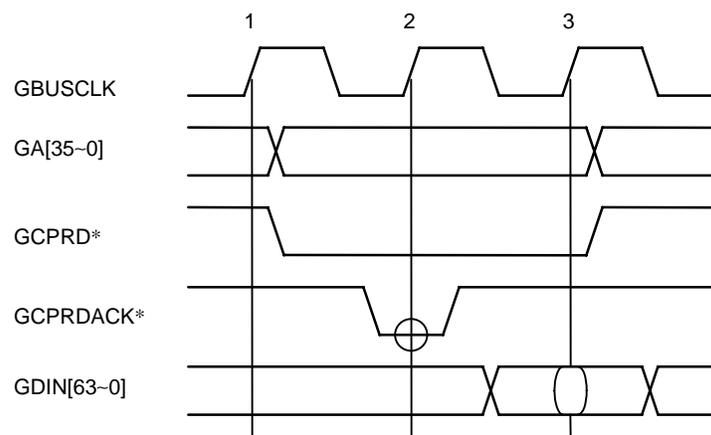


Figure 18-1 Coprocessor Read Operation

- 1 The GCPRD* signal is asserted, starting a coprocessor read operation. At the same time, coprocessor instruction codes are output to the GA[35~0] signal.
- 2 It is acknowledged that the GCPRDACK* signal is "L."
- 3 The data on the GDIN[63~0] signal are collected. Also, the GCPRD* signal is deasserted. This completes the coprocessor read operation. Up to this timing, valid values are retained in the GA[35~0] signal.

However, data transfer between the TX49 Megacell and the coprocessor do not occur with the COPz instruction, so the TX49 Megacell does not collect the data from the GDIN[63~0] signal.

If another coprocessor read operation immediately follows this coprocessor read operation, the coprocessor read operation starts without deasserting the GCPRD* signal.

When the GCPRDACK* signal is asserted, the TX49 Megacell ends the coprocessor read operation.

18.2 Coprocessor write operation

The coprocessor write operation is generated by the execution of the MTCz instruction and the CTCz instruction. The timing of the coprocessor write operation is the same as the write operation for the memory. The coprocessor write operation uses the GA[35~0] signal, the GDOUT[63~0] signal, the GCPWR* signal, and the GCPWRACK* signal. Figure 18-1 shows the coprocessor write operation.

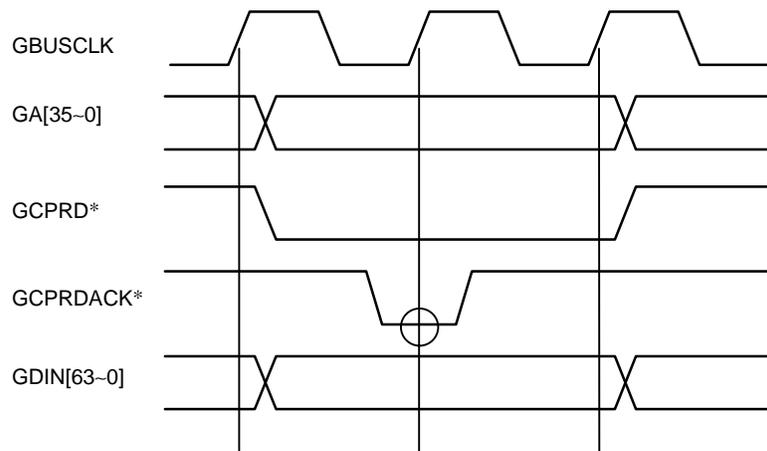


Figure 18-1 Coprocessor Write Operation

- 1 The GCPWR* signal is asserted, and a coprocessor write operation starts. At the same time, each coprocessor instruction code and the source register contents are output to the GA[35~0] signal, and the GDOUT[63~0] signal.
- 2 Having acknowledged that the GCPWRACK* signal is "L," the GCPWR* signal is deasserted. This completes the coprocessor write operation. Up to this timing, valid values are retained on the GA[35~0] signal, and the GDOUT[63~0] signal.

If another coprocessor write operation immediately follows this coprocessor write operation, the coprocessor write operation starts without deasserting the GCPWR* signal.

The TX49 Megacell continues the coprocessor write operation until the GCPWRACK* signal is asserted.

18.3 Coprocessor condition branch instruction operation

Figure 18-1 shows the timing for acquiring the GCPCOND[3~2] signal in a coprocessor condition branch instruction. The GCPCOND[3~2] signal is acquired at the GBUSCLK rise at the end of the instruction decoding stage of the TX49 Processor Core.

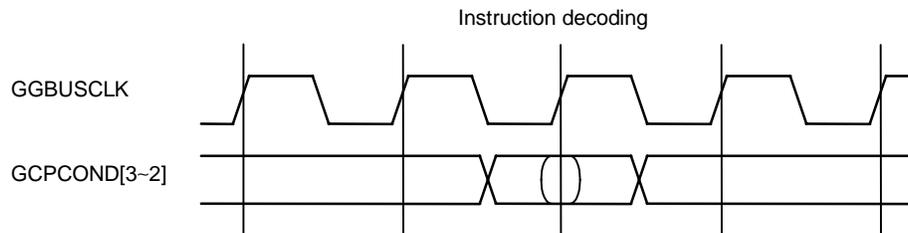


Figure 18-1 Coprocessor Condition Branch Instruction Operation

For pipeline operations, please Chapter 6 of this manual.

19. Interrupts

As for interrupt methods, the TX49 Megacell supports six hardware interrupts and two software interrupts. In addition to them, it supports the non-maskable interrupts. Interrupt exceptions can be generated by the GINT[5~0]* signal. Non-maskable interrupt exceptions can be generated by the GNMI* signal. These interrupt signals are active-low signals and are sampled at the positive-going transitions of GBUSCLK.

This section describes the timings of these external interrupt signals.

19.1 GNMI* Signal

As shown in Figure 19-1, the TX49 Megacell acquires the GNMI* signal at a rise of the GBUSCLK signal.

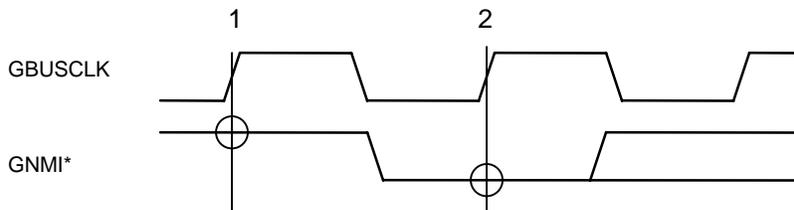


Figure 19-1 Non-Maskable Interrupt

- 1 The “H” of the GNMI* signal is acknowledged.
- 2 The transition from “H” to “L” of the GNMI* signal is acknowledged, so a non-maskable interrupt exception occurs.

A non-maskable interrupt exception occurs when the transition from “H” to “L” of the GNMI* signal is acknowledged. The TX49 Megacell retains this transition internally. A peripheral circuit can generate a non-maskable interrupt exception by asserting the GNMI* signal for one cycle or more. To generate the next non-maskable interrupt exception, it is necessary for the peripheral circuit to turn the GNMI* signal to “H” and then to “L” since the GNMI* signal is valid when transiting from “H” to “L.”

If the TX49 Megacell acknowledges the transition of the GNMI* signal from “H” to “L” during a bus operation, a non-maskable interrupt exception occurs when the bus operation ends.

If the TX49 Megacell acknowledges the transition of the GNMI* signal from “H” to “L” during release of the bus ownership, a non-maskable interrupt exception occurs immediately after the TX49 Megacell obtains the G-Bus ownership.

19.2 GINT[5:0]* Signal

An interrupt exception is generated by the GINT[5~0]* signal. Interrupts by the GINT[5~0]* signal can be masked by the bit which the IntMask bit of the Status register corresponds to and the IE bit. As shown in Figure 19-1, the TX49 Megacell acquires the GINT[5~0]* signal at a rise of the GBUSCLK signal.

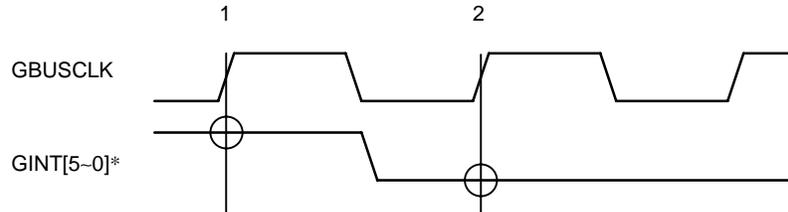


Figure 19-1 Interrupt

- 1 The "H" of the GINT* signal is acknowledged.
- 2 An interrupt exception occurs since the "L" of the GINT* signal is acknowledged.

As shown in Figure 19-2, the TX49 Megacell acknowledges the GINT[5~0]* signal at a rise of the GBUSCLK signal. To make sure that an interrupt request is acknowledged, the GINT[5~0]* signal must be asserted until an interrupt exception occurs. If it is deasserted before that, the occurrence of an interrupt exception is not guaranteed.

The IP bit of the Cause register reflects the GINT[5~0]* signal at the time when the Cause register was read. To correctly evaluate the cause using the interrupt exception handler, the GINT[5~0]* signal must be asserted until the handler reads the Cause register. Then, the GINT[5~0]* signal should be deasserted by the interrupt exception handler. When an interrupt becomes valid before deasserting the GINT[5~0]* signal, another interrupt exception occurs.

20. Reset and Endian Initial Setting

The TX49 Megacell initializes its internal status by acknowledging the assertion of the GCOLDRESET* signal. Then, a reset exception is generated by acknowledging a deassertion of the GCOLDRESET* signal. This section describes the reset of the TX49 Megacell and the endian initial setting during a reset.

20.1 Reset

Figure 20-1 shows the timing of the TX49 Megacell GCOLDRESET* signal.

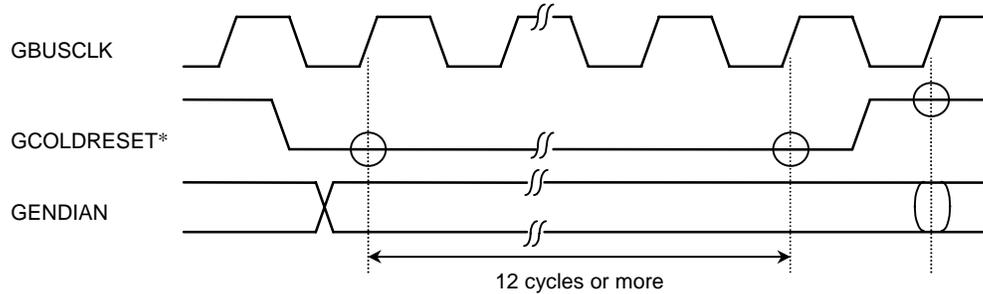


Figure 20-1 Reset

The TX49 Megacell acquires the GCOLDRESET* signal at a rise of the GBUSCLK signal. Having acknowledged the assertion of the GCOLDRESET* signal for twelve rises or more of the GBUSCLK signal, the TX49 Megacell will enter the reset status. The external signals of the TX49 Megacell in the reset status will be as follows.

GA[35~0]	Indefinite
GDOUT[63~0]	Indefinite
GBE[7~0]*	All H
GRD*	H
GWR*	H
GCACHE *	H
GID	H
GBURST*	H
GBSTART*	H
GLAST*	H
GBUSGNT*	H
GHALT	L
GDOZE	L

If an assertion of the GCOLDRESET* signal is acknowledged while the TX49 Megacell is executing a bus operation, the bus operation is forcibly suspended. Also, if an assertion of the GCOLDRESET* signal is acknowledged while the bus ownership is being released, the GGNT* signal, the GSGNT* signal, the GHPGGNT* signal, or the GHPSGNT* signal is forcibly deasserted.

Having acknowledged an assertion of the GCOLDRESET* signal, the TX49 Megacell performs an internal reset operation. The following are included in the reset operation of the TX49 Megacell.

- Control registers are initialized
- Valid bits of the instruction and data caches, and the lock bit are cleared.
- The internal pipeline is initialized.

When a deassertion of the GCOLDRESET* signal is acknowledged, the TX49 Megacell generates a reset exception. Figure 20-2 shows the power-on reset.

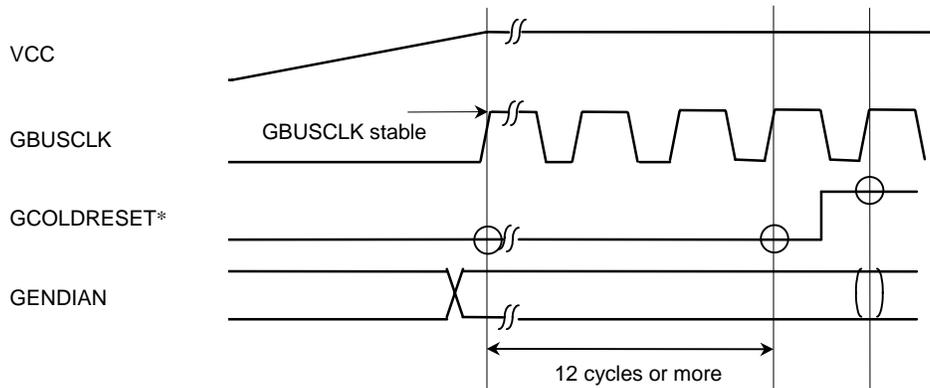


Figure 20-2 Power-On Reset

When the GBUSCLK is stabilized after the power has reached the specified voltage, an assertion of the GCOLDRESET* signal should be maintained for twelve cycles or more.

20.2 Endian Initial Setting

As shown in Figures 20-1 and 20-2, the endian in the kernel mode and the default endian in the user mode are set by the value of the GENDIAN signal at the first rise of the GBUSCLK signal after the GCOLDRESET* signal was deasserted. If the GENDIAN signal is "H," the TX49 Megacell is set as the big endian; and if it is "L," it is set as the little endian. The GENDIAN signal is valid only at the time of a reset.

20.3 Clocks

The TX49 Megacell operation clocks are generated by the GCPUCLK signal inside the TX49 Megacell. The GBUSCLK is the timing reference for the TX49 Megacell interface signals.

The GCPUCLK of the TX49 Megacell can be halted. This halts the TX49 Megacell operations, lowering the power consumption in the system.

21. Low Power Consumption Modes

The TX49 Megacell can reduce its power consumption compared to the normal mode by controlling its internal clocks. The following two operation modes function as low power consumption modes of the TX49 Megacell:

- Halt mode
- Doze mode

21.1 Halt mode

The halt mode reduces power consumption by halting TX49 Megacell operation. By setting the Wait mode bit of the Config register to 0 by the software and executing WAIT instruction, the TX49 Megacell mode shifts from the normal operation mode to the halt mode.

Therefore, as for bus control requests in the halt mode, a bus release request is responded to in cases of ET concurrency such as the GREQ* signal or the GHPGREQ* signal. However, the request is not responded to in cases of ST concurrency such as the GSREQ* signal or the GHPSREQ* signal. On the other hand, if WAIT instruction is executed while the bus is being released, the halt mode starts in cases of ET concurrency, but in cases of ST concurrency starts after bus ownership is granted and the GHALT signal is asserted.

If WAIT instruction is executed during a bus operation, the GHALT signal is asserted after the bus operation is completed.

If data remain in the write buffer, the write operation is executed even after shifting to the halt mode.

The internal halt bit is cleared by the assertion of the GINT[5~0]* signal, the GNMI* signal, the GRESET* signal or the GCOLDRESET* signal, and the TX49 Megacell return from the halt mode. If this is caused by the assertion of the GINT[5:0]* signal, the TX49 Megacell is released from the halt mode irrespective of the value in the IntMask field of the Status register. If the TX49 Megacell is brought back from the halt mode by the GCOLDRESET* signal, the GRESET* signal, the GNMI* signal, or a non-masked GINT[5~0]* signal, the initial instruction in the corresponding exception handler is executed. At this time, the EPC register is pointing to the instruction following the WAIT instruction. If it is recovered by a masked GINT[5~0]* signal, execution resumes from the instruction following the instruction that was being executed when it shifted to the halt mode.

As shown in Figure 21-1 the TX49 Megacell outputs the status of the internal halt bit on the GHALT signal. The memory interface output signals in the halt mode are maintained in the same status as when no bus operation was being executed.

Note: When the condition is brought back from the Power Consumption Modes are satisfied and WAIT instruction is executed, the TX49 Megacell does not shift to the mode.

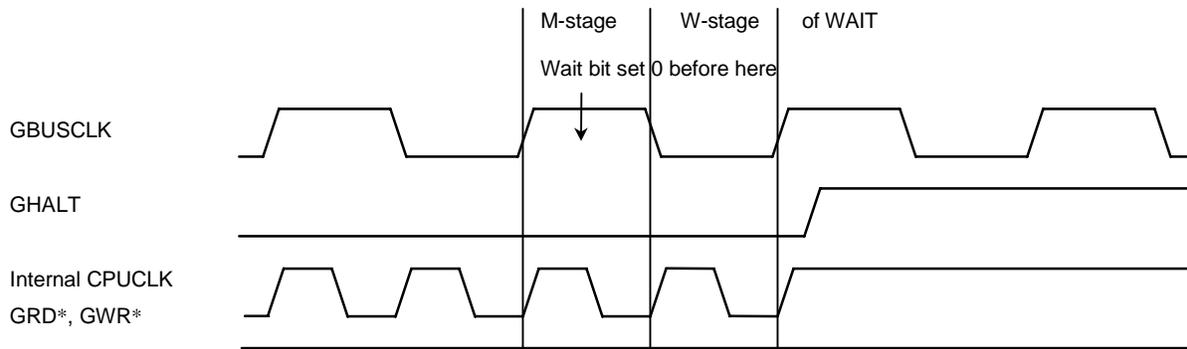


Figure 21-1 Halt Mode

21.2 Doze mode

The doze mode is also a mode which halts TX49 Megacell operation in order to lower power-consumption. However, the difference from the halt mode is that bus control requests (both ST concurrency and ET concurrency) from an external bus master can be responded to. Snooping operation of the data cache can also be performed in ST concurrency. By setting the Wait mode bit of the Config register to 1 by the software and executing WAIT instruction, the TX49 Megacell mode shifts from the normal operation mode to the doze mode. Then, the TX49 Processor Core that is built into the TX49 Megacell halts operation while retaining the pipeline status.

As mentioned above, bus control requests are responded to while in the doze mode in cases of ET concurrency such as the GREQ* signal and the GHPGREQ* signal, and in cases of ST concurrency such as the GSREQ* signal and the GHPSREQ* signal. On the other hand, if WAIT instruction is executed while the bus is being released, the doze mode starts in cases of ET concurrency, but in cases of ST concurrency starts after bus ownership is granted and the GDOZE signal is asserted. If WAIT instruction is executed during a bus operation, the GDOZE signal is asserted after the bus operation is completed. The snooping of an external bus master is done by ST concurrency when the TX49 Megacell is in the doze mode. For the bus that is released by the assertion of the SGNT* signal or the GHPSGNT* signal, snooping of the data cache can be performed by the GSNOOP* signal and the GA[35~0] signal. When an external bus master deasserts the GSREQ* signal or the GHPSREQ* signal, the TX49 Megacell deasserts the GSGNT* signal or the GHPSGNT* signal.

By asserting the GINT[5~0]* signal, the GNMI* signal, the GRESET* signal or the GCOLDRESET* signal, the internal doze bit is cleared and the TX49 Megacell returns from the doze mode. If this is caused by the assertion of the GINT[5~0]* signal, the TX49 Megacell is released from the doze mode irrespective of the value in the IntMask field of the Status register. If the TX49 Megacell is brought back from the doze mode by the GCOLDRESET* signal, the GNMI* signal, or a non-masked GINT[5~0]* signal, the top instruction in the corresponding exception handler is executed. At this time, the EPC is pointing to the instruction following the WAIT instruction. If it is recovered by a masked GINT[5~0]* signal, execution resumes from the instruction following the instruction that was being executed when it shifted to the doze mode.

As shown in Figure 21-1, the TX49 Megacell outputs the status of the internal doze bit on the GDOZE signal. The memory interface output signals in the doze mode are maintained in the same status as when no bus operation was executed.

Note: When the condition is brought back from the Power Consumption Modes are satisfied and WAIT instruction is executed, the TX49 Megacell does not shift to the mode.

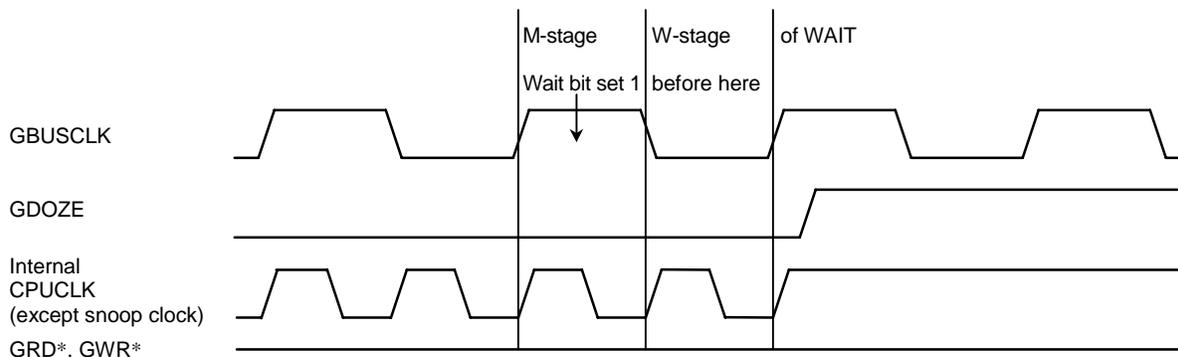


Figure 21-1 Doze Mode

21.3 Status Shifts

Figure 21-1 shows the status shifts in the operation mode of the TX49 Megacell.

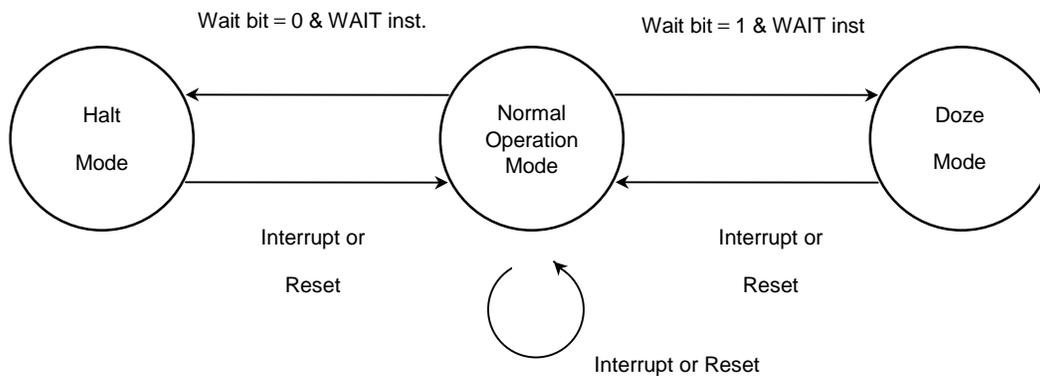


Figure 21-1 Status Shift Among Normal Operation Mode and Low Power Consumption Modes

When operation status shifts from the normal operation mode to the halt mode, it is returned to the normal operation mode by an interrupt or a reset. Similarly, when it shifts from the normal operation mode to the doze mode, it is returned to the normal operation mode by an interrupt or a reset. After a reset, the TX49 Megacell is initialized to the normal operation mode.

22. TX49 MCU Concept

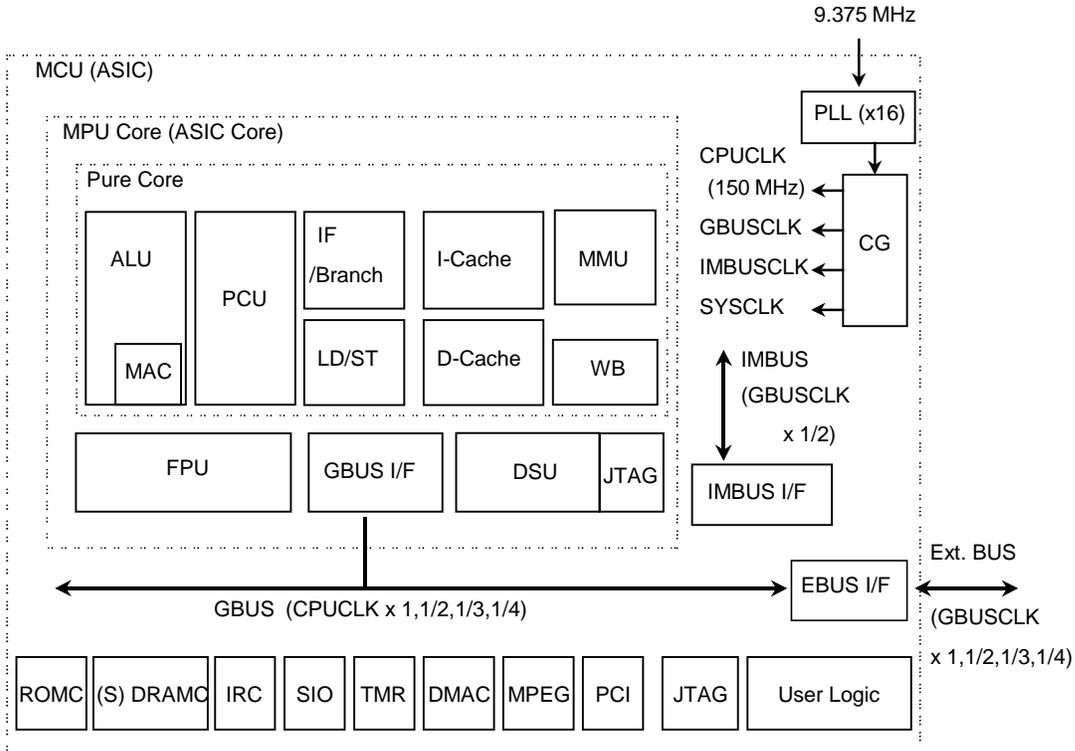


Figure 22-1 MCU concept

I. Using Toshiba Semiconductors Safely

TOSHIBA is continually working to improve the quality and the reliability of its products. Nevertheless, semiconductor devices in general can malfunction or fail due to their inherent electrical sensitivity and vulnerability to physical stress. It is the responsibility of the buyer, when utilizing TOSHIBA products, to observe standards of safety, and to avoid situations in which a malfunction or failure of a TOSHIBA product could cause loss of human life, bodily injury or damage to property.

In developing your designs, please ensure that TOSHIBA products are used within specified operating ranges as set forth in the most recent products specifications. Also, please keep in mind the precautions and conditions set forth in the TOSHIBA Semiconductor Reliability Handbook.

I. Safety Precautions

This section lists important precautions which users of semiconductor devices (and anyone else) should observe in order to avoid injury and damage to property, and to ensure safe and correct use of devices. Please be sure that you understand the meanings of the labels and the graphic symbol described below before you move on to the detailed descriptions of the precautions.

[Explanation of labels]

	Indicates an imminently hazardous situation which will most likely result in death or serious injury if you do not follow instructions.
	Indicates a potentially hazardous situation which could result in death or serious injury if you do not follow instructions.
	Indicates a potentially hazardous situation which, if not remedied, may result in minor injury or worse.

[Explanation of graphic symbol]

Graphic symbol	Meaning
	Indicates that caution is required (laser beam is dangerous to eyes).

A. General Precautions Regarding Semiconductor Devices

Do not use devices under conditions exceeding their absolute maximum ratings (e.g. current, voltage, power dissipation or temperature).

This may cause the device to break down, may degrade its performance, or cause it to catch fire or explode, resulting in injury.

Do not insert devices in the wrong orientation.

Make sure that the positive and negative terminals of power supplies are connected correctly.

Otherwise, the rated maximum current or power dissipation may be exceeded and the device may break down or undergo performance degradation, causing it to catch fire or explode and resulting in injury.

When power to a device is on, do not touch the device's heat sink.

Heat sinks become hot, so you may burn your hand.

Do not touch the tips of device leads.

Because some types of device have leads with pointed tips, you may prick your finger.

When conducting any kind of evaluation, inspection or testing, be sure to connect the testing equipment's electrodes or probes to the pins of the device under test before powering it on.

Otherwise, you may receive an electric shock, causing injury.

Before grounding an item of measuring equipment or a soldering iron, check that there is no electrical leakage from it.

Electrical leakage may cause the device which you are testing or soldering to break down, or could give you an electric shock.

Always wear protective glasses when cutting the leads of a device with clippers or a similar tool.

If you do not, small bits of metal flying off the cut ends may damage your eyes.

A. Precautions Specific to Each Product Group

1. Optical Semiconductor Devices

<p>When a visible semiconductor laser is operating, do not look directly into the laser beam or look through the optical system.</p> <p>This is highly likely to impair vision, and in the worst case, may cause blindness.</p> <p>If it is necessary to examine the laser apparatus, for example, to inspect its optical characteristics, always wear the appropriate type of laser protective glasses according to JIS standard JISC6802.</p>
<p>Ensure that the current flowing in an LED device does not exceed the device's maximum rated current. This is particularly important for resin-packaged LED devices, as excessive current may cause the package resin to blow up, scattering resin fragments and causing injury.</p> <p>When testing the dielectric strength of a photocoupler, use testing equipment which can shut off the supply voltage to the photocoupler. If you detect a leakage current of more than 100 μA, use the testing equipment to shut off the photocoupler's supply voltage; otherwise a large short-circuit current will flow continuously, and the device may break down or burst into flames, resulting in fire or injury.</p> <p>When incorporating a visible semiconductor laser into a design, use the device's internal photodetector or a separate photodetector to stabilize the laser's radiant power so as to ensure that laser beams exceeding the laser's rated radiant power cannot be emitted.</p> <p>If this stabilizing mechanism does not work and the rated radiant power is exceeded, the device may break down or the excessively powerful laser beams may cause injury.</p>

1. Power Devices (IGBTs)

<p>Never touch a power device while it is powered on. Also, after turning off a power device, do not touch it until it has thoroughly discharged all remaining electrical charge.</p> <p>Touching a power device while it is powered on or still charged could cause a severe electric shock, resulting in death or serious injury.</p> <p>When conducting any kind of evaluation, inspection or testing, be sure to connect the testing equipment's electrodes or probes to the device under test before powering it on.</p> <p>When you have finished, discharge any electrical charge remaining in the device.</p> <p>Connecting the electrodes or probes of testing equipment to a device while it is powered on may result in electric shock, causing injury.</p>
--

Do not use devices under conditions which exceed their absolute maximum ratings (current, voltage, power dissipation, temperature etc.).

This may cause the device to break down, causing a large short-circuit current to flow, which may in turn cause it to catch fire or explode, resulting in fire or injury.

Use a unit which can detect short-circuit currents and which will shut off the power supply if a short-circuit occurs.

If the power supply is not shut off, a large short-circuit current will flow continuously, which may in turn cause the device to catch fire or explode, resulting in fire or injury.

When designing a case for enclosing your system, consider how best to prevent the user from shrapnel in the event of the device catching fire or exploding.

Flying shrapnel can cause injury.

When conducting any kind of evaluation, inspection or testing, always use protective safety tools such as a cover for the device. Otherwise you may sustain injury caused by the device catching fire or exploding.

Make sure that all metal casings in your design are grounded to earth.

Even in modules where a device's electrodes and metal casing are insulated, capacitance in the module may cause the electrostatic potential in the casing to rise.

Dielectric breakdown may cause a high voltage to be applied to the casing, causing electric shock and injury to anyone touching it.

When designing the heat radiation and safety features of a system incorporating high-speed rectifiers, remember to take the device's forward and reverse losses into account.

The leakage current in these devices is greater than that in ordinary rectifiers; as a result, if a high-speed rectifier is used in an extreme environment (e.g. at high temperature or high voltage), its reverse loss may increase, causing thermal runaway to occur. This may in turn cause the device to explode and scatter shrapnel, resulting in injury to the user.

When conducting any kind of evaluation, inspection or testing, either wear protective gloves or wait until the device has cooled properly before handling it.

Devices become hot when they are operated. Even after the power has been turned off, the device will retain residual heat which may cause a burn to anyone touching it.

1. Bipolar ICs (for automotive use)

If your design includes an inductive load such as a motor coil, incorporate diodes or similar devices into your design to prevent negative current from flowing in.

The load current generated by powering the device on and off may cause it to function erratically or to break down, which could in turn cause injury.

Ensure that the power supply to any device which incorporates protective functions is stable.

If the power supply is unstable, the device may operate erratically, preventing the protective functions from working correctly. If protective functions fail, the device may break down, causing injury to the user.

I. General Safety Precautions and Usage Considerations

This section is designed to help you gain a better understanding of semiconductor devices, so as to ensure the safety, quality and reliability of the devices which you incorporate into your designs.

A. From Incoming to Shipping

1. Electrostatic Discharge (ESD)

When handling individual devices (which are not yet mounted on a printed circuit board), be sure that the environment is protected against electrostatic electricity. Operators should wear anti-static clothing, and containers and other objects that come into direct contact with devices should be made of anti-static materials and should be grounded to earth via an 0.5- to 1.0-M Ω protective resistor.

Please follow the precautions described below; this is particularly important for devices which are marked “Be careful of static.”.



a) Work Environment

- (1) When humidity in the working environment decreases, the human body and other insulators can easily become charged with static electricity due to friction. Maintain the recommended humidity of 40% to 60% in the work environment, while also taking into account the fact that moisture-proof-packed products may absorb moisture after unpacking.
- (2) Be sure that all equipment, jigs and tools in the working area are grounded to earth.
- (3) Place a conductive mat over the floor of the work area, or take other appropriate measures, so that the floor surface is protected against static electricity and is grounded to earth. The surface resistivity should be 10^4 to $10^8 \Omega/\text{sq}$ and the resistance between surface and ground, 7.5×10^5 to $10^8 \Omega$.
- (4) Cover the workbench surface also with a conductive mat (with a surface resistivity of 10^4 to $10^8 \Omega/\text{sq}$, for a resistance between surface and ground of 7.5×10^5 to $10^8 \Omega$). The purpose of this is to disperse static electricity on the surface (through resistive components) and ground it to earth. Workbench surfaces must not be constructed of low-resistance metallic materials that allow rapid static discharge when a charged device touches them directly.
- (5) Pay attention to the following points when using automatic equipment in your workplace:
 - (a) When picking up ICs with a vacuum unit, use a conductive rubber fitting at the end of the pick-up wand to protect against electrostatic charge.
 - (b) Minimize friction on IC package surfaces. If some rubbing is unavoidable due to the device's mechanical structure, minimize the friction plane or

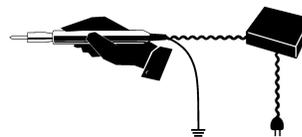
3 General Safety Precautions and Usage Considerations

use material with a small friction coefficient and low electrical resistance. Also consider the use of an ionizer.

- (c) In sections that come into contact with device lead terminals, use a material that dissipates static electricity.
- (d) Ensure that no statically charged bodies (such as work clothes or the human body) touch the devices.
- (e) Make sure that sections of the tape carrier which come into contact with installation devices or other electrical machinery are made of a low-resistance material.
- (f) Make sure that jigs and tools used in the assembly process do not touch devices.
- (g) In processes in which packages may retain an electrostatic charge, use an ionizer to neutralize the ions.
- (6) Make sure that CRT displays in the working area are protected against static charge, for example by a VDT filter. As much as possible, avoid turning displays on and off. Doing so can cause electrostatic induction in devices.
- (7) Keep track of charged potential in the working area by taking periodic measurements.
- (8) Ensure that work chairs are protected by an anti-static textile cover and are grounded to the floor surface by a grounding chain. (Suggested resistance between the seat surface and grounding chain is 7.5×10^5 to $10^{12} \Omega/\text{sq.}$)
- (9) Install anti-static mats on storage shelf surfaces. (Suggested surface resistivity is 10^4 to $10^8 \Omega/\text{sq.}$; suggested resistance between surface and ground is 7.5×10^5 to $10^8 \Omega/\text{sq.}$)
- (10) For transport and temporary storage of devices, use containers (boxes, jigs, bags) that are made of anti-static materials or of materials that dissipate electrostatic charge.
- (11) Make sure that cart surfaces which come into contact with device packaging are made of materials that will conduct static electricity, and verify that they are grounded to the floor surface with a grounding chain. (The suggested resistance between the cart surface and grounding chain is 7.5×10^5 to $10^{10} \Omega/\text{sq.}$)
- (12) In any location where the level of static electricity is to be closely controlled, the ground resistance level should be Class 3 or above. Use different ground wires for all items of equipment which may come into physical contact with devices.

a) Operating Environment

- (1) Operators must wear anti-static clothing and conductive shoes (or a leg or heel strap).
- (2) Operators must wear a wrist strap grounded to earth via a resistor of about $1M\Omega$.
- (3) Soldering irons must be grounded from iron tip to earth, and must be used only at low voltages (6V to 24V).



3 General Safety Precautions and Usage Considerations

- (4) If the tweezers you use are likely to touch the device terminals, use anti-static tweezers and in particular avoid metallic tweezers. If a charged device touches a low-resistance tool, rapid discharge can occur. When using vacuum tweezers, attach a conductive chucking pat to the tip, and connect it to a dedicated ground used especially for anti-static purposes (suggested resistance value: 10^4 to $10^8\Omega$).
- (5) Do not place devices or their containers near sources of strong electrical fields (such as above a CRT).
- (6) When storing printed circuit boards which have devices mounted on them, use a board container or bag that is protected against static charge. To avoid the occurrence of static charge or discharge due to friction, keep the boards separate from one other and do not stack them directly on top of one another.
- (7) Ensure, if possible, that any articles (such as clipboards) which are brought to any location where the level of static electricity must be closely controlled are constructed of anti-static materials.
- (8) In cases where the human body comes into direct contact with a device, be sure to wear anti-static finger covers or gloves (suggested resistance value: $10^8\Omega$ or less).
- (9) Equipment safety covers installed near devices should have resistance ratings of $10^9\Omega$ or less.
- (10) If a wrist strap cannot be used for some reason, and there is a possibility of imparting friction to devices, use an ionizer.

- (11) The transport film used in TCP products is manufactured from materials in which static charges tend to build up. When using these products, install an ionizer to prevent the film from being charged with static electricity. Also, ensure that no static electricity will be applied to the product's copper foils by taking measures to prevent static occurring in the peripheral equipment.

1. Vibration, Impact and Stress

Handle devices and packaging materials with care. To avoid damage to devices, do not toss or drop packages. Ensure that devices are not subjected to mechanical vibration or shock during transportation. Ceramic package devices, and devices in canister-type packages which have empty space inside them, are subject to damage from vibration and shock because bonding wires are secured only at their ends. Plastic molded devices, on the other hand, have a relatively high level of resistance to vibration and mechanical shock because their bonding wires are enveloped and fixed in resin. However, when any device or package type is installed in target equipment, it is to some extent susceptible to wiring disconnections and other damage from vibration, shock and stressed solder junctions. Therefore when incorporating devices into the design of vibration-prone equipment, the structural design of the equipment must be thought out carefully.

If a device is subjected to especially strong vibration, mechanical shock or stress, the package or the chip itself may crack. In products such as CCDs which incorporate window glass, this could cause surface flaws in the glass or cause the glass connection to separate.

Furthermore, it is generally known that stress applied to a semiconductor device through the package changes the resistance characteristics of the chip because of piezoelectric effects. In analog circuit design, attention must be paid to the problem of package stress as well as to the dangers of vibration and shock as described above.

A. Storage

1. General Storage

- (1) Avoid storage locations where devices will be exposed to moisture or direct sunlight. (Be especially careful during periods of rain or snow.)
- (2) Do not place device cartons upside down. Stack cartons on top of one another in an upright position only; do not place cartons on their sides.
- (3) The storage area temperature should be kept within a temperature range of 5°C to 35°C, and relative humidity should be maintained at between 45% and 75%.
- (4) Do not store devices in the presence of harmful (especially corrosive) gases, or in dusty conditions.
- (5) Use storage areas where there is minimal temperature fluctuation. Rapid temperature changes can cause moisture to form on stored devices, resulting in lead oxidation or corrosion. As a result, the solderability of the leads will be degraded.

3 General Safety Precautions and Usage Considerations

- (6) When repacking devices, use anti-static containers.
- (7) Do not allow external forces or loads to be applied to devices while they are in storage.
- (8) If devices have been stored for more than two years, their electrical characteristics should be tested and their leads should be tested for ease of soldering before they are used.

1. Moisture-Proof Packing

- (1) Do not drop or toss device packing. The laminated aluminum material in it can be rendered ineffective by rough handling.
- (2) Ensure that packing materials are stored in a 30°C, 90% RH environment. Use devices within 12 months.
- (3) If the 30% humidity indicator shown in Figure 1.1 is pink when the packing is opened, depending on the device and packing types, it may be advisable to bake the devices at high temperature to remove any moisture. See Section 3.2.2(4) below. It may also be advisable to bake the devices if the effective usage period for the indicator has expired. After the pack is opened, use the devices in a 30°C, 60% RH environment, and within the effective usage period listed on the moisture-proof package.
- (4) The following describes high-temperature treatments for the various packing types. Contact Toshiba or a Toshiba distributor for more information.
 - (a) Tray: If the tray is heat-proof, bake at 125°C for 20 hours (heat-proof trays bear a “Heat-Proof” marking). Bake non-heat-proof trays at 70°C for 168 hours.
 - (b) Tube: Tubes are not heat-proof. Transfer devices to heat-proof trays or aluminum tubes before baking at 125°C for 20 hours.
 - (c) Tape: Packing that includes adhesive or embossed tape cannot be baked. Devices packed on tape must be used within the permitted time limit after unpacking, as specified on the packing.



For types (a) and (b), be careful not to bend device leads when baking.

Figure 1.1 Humidity Indicator

A. Design

Care must be exercised in the design of electronic equipment to achieve the desired reliability. It is important not only to adhere to specifications concerning absolute maximum ratings and recommended operating conditions, it is also important to consider the overall environment in which equipment will be used, including factors such as the ambient temperature, transient noise, voltage and current surges, as well as mounting conditions that affect device reliability. This section describes some general precautions that you should observe when designing circuits and when mounting devices on printed circuit boards.

For more detailed information about each product family, refer to the relevant individual databooks, available from Toshiba.

1. Absolute Maximum Ratings

Do not use devices under conditions in which their absolute maximum ratings (e.g. current, voltage or power dissipation) will be exceeded. A device may break down or its performance may be degraded, causing it to catch fire or explode resulting in injury to the user.

The absolute maximum ratings are rated values which must not be exceeded during operation, even for an instant. Although absolute maximum ratings differ from product to product, they essentially concern the voltage and current at each pin, the allowable power dissipation, and the junction and storage temperatures. If the voltage or current on any pin exceeds the absolute maximum rating, the device's internal circuitry can become degraded. In the worst case, heat generated in internal circuitry can fuse wiring or cause the semiconductor chip to break down.

If storage or operating temperatures exceed rated values, the package seal can deteriorate or the wires can become disconnected due to the differences between the thermal expansion coefficients of the materials from which the device is constructed.

1. Recommended Operating Conditions

The recommended operating conditions for each device are those necessary to guarantee that the device will operate as specified in the datasheet.

If greater reliability is required, derate the device's absolute maximum ratings for voltage, current, power and temperature before using it.

1. Derating

When incorporating a device into your design, reduce its rated maximum voltage, current, power dissipation and operating temperature in order to ensure high reliability.

Since derating differs from application to application, refer to the technical datasheets available for the various devices used in your design.

1. **Unused Pins**

Some devices can exhibit input instability problems if unused pins are left open. Similarly, care must be taken not to connect the output pins of a device to the power supply (Vcc or Vdd) pin or to other output pins. For details concerning the handling of unused pins, follow the procedures described in the relevant technical datasheet or databook for the device being used. CMOS logic IC inputs, for example, have extremely high impedance. If an input pin is left open, it can easily pick up extraneous noise and become unstable. In this case, if the input voltage level reaches an intermediate level, both the P-channel and N-channel

transistors may be turned on, allowing unwanted supply current to flow. Therefore, ensure that the unused input pins of a device are connected to the power supply (Vcc) pin or ground (GND) pin of the same device. For details of what to do with the pins of heat sinks, refer to the relevant technical datasheet or databook.

1. Latch-up

Latch-up is an abnormal condition inherent in CMOS devices, in which Vcc gets shorted to ground. This happens when a parasitic PN-PN junction (thyristor structure) internal to the CMOS chip is turned on, causing a large current of the order of several hundred mA or more to flow between Vcc and GND, eventually causing the device to break down.

Latch-up occurs when the input/output voltage exceeds the rated value, causing a large current to flow in the internal chip, or when the voltage on the Vcc (Vdd) pin exceeds its rated value, forcing the internal chip into a breakdown condition. Once the chip falls into the latch-up state, even though the excess voltage may have been applied only for an instant, the large current continues to flow between Vcc (Vdd) and GND (Vss). This causes the device to heat up and, in extreme cases, to emit gas fumes as well. To avoid this problem, observe the following precautions:

- (1) Do not allow voltage levels on the input/output pins either to rise above Vcc (Vdd) or to fall below GND (Vss). Also, follow any prescribed power-on sequence, so that power is applied gradually or in steps rather than abruptly.
- (2) Allow no abnormal noise signals to be applied to the device.
- (3) Set the voltage levels of unused input pins to Vcc (Vdd) or (GND) Vss.
- (4) Do not connect outputs to one another.

1. Input/Output Protection

Wired-AND configurations, in which outputs are connected together, cannot be used, since this short-circuits the outputs. Outputs should, of course, never be connected to Vcc (Vdd) or GND (Vss).

Furthermore, ICs with tri-state outputs can undergo performance degradation if a shorted output current is allowed to flow for an extended period of time.

Therefore, when designing circuits, make sure that tri-state outputs will not be enabled simultaneously.

1. Load Capacitance

Some devices display increased delay times if the load capacitance is large. Also, large charging and discharging currents will flow in the device, causing noise. Furthermore, since outputs are shorted for a relatively long time, wiring can become fused.

Consult the technical information for the device being used to determine the recommended load capacitance.

1. Thermal Design

The failure rate of semiconductor devices is greatly increased as operating

3 General Safety Precautions and Usage Considerations

temperatures increase. As shown in Figure 1.1, the internal thermal stress on a device is the sum of the ambient temperature and the temperature rise due to power dissipation in the device. Therefore, to achieve optimum reliability, observe the following precautions concerning thermal design:

- (1) Keep the ambient temperature (T_a) as low as possible.
- (2) If the device's dynamic power dissipation is relatively large, select the most appropriate circuit board material, and consider the use of heat sinks or of forced air cooling. Such measures will help lower the thermal resistance of the package.

- (3) Derate the device's absolute maximum ratings to minimize thermal stress from power dissipation.
- $$\theta_{ja} = \theta_{jc} + \theta_{ca}$$
- $$\theta_{ja} = (T_j - T_a) / W$$
- $$\theta_{jc} = (T_j - T_c) / W$$
- $$\theta_{ca} = (T_c - T_a) / W$$
- in which θ_{ja} = thermal resistance between junction and surrounding air ($^{\circ}\text{C}/\text{W}$)
- θ_{jc} = thermal resistance between junction and package surface, or internal thermal resistance ($^{\circ}\text{C}/\text{W}$)
- θ_{ca} = thermal resistance between package surface and surrounding air,
or external thermal resistance ($^{\circ}\text{C}/\text{W}$)
- T_j = junction temperature or chip temperature ($^{\circ}\text{C}$)
- T_c = package surface temperature or case temperature ($^{\circ}\text{C}$)
- T_a = ambient temperature ($^{\circ}\text{C}$)
- W = power dissipation (W)

Figure 1.1 Thermal Resistance of Package

1. Interfacing

When connecting inputs and outputs between devices, make sure input voltage (V_{IL}/V_{IH}) and output voltage (V_{OL}/V_{OH}) levels are matched. Otherwise, the devices may malfunction. When connecting devices operating at different supply voltages, such as in a dual-power-supply system, be aware that erroneous power-on and power-off sequences can result in device breakdown. For details of how to interface particular devices, consult the relevant technical datasheets and databooks. If you have any questions or doubts about interfacing, contact your nearest Toshiba office or distributor.

1. Decoupling

Spike currents generated during switching can cause V_{cc} (V_{dd}) and GND (V_{ss}) voltage levels to fluctuate, causing ringing in the output waveform or a delay in response speed. (The power supply and GND wiring impedance is normally $50\ \Omega$ to $100\ \Omega$.) For this reason, the impedance of power supply lines with respect to high frequencies must be kept low. This can be accomplished by using thick and short wiring for the V_{cc} (V_{dd}) and GND (V_{ss}) lines and by installing decoupling capacitors (of approximately 0.01 to $1\ \mu\text{F}$ capacitance) as high-frequency filters between V_{cc} (V_{dd}) and GND (V_{ss}) at strategic locations on the printed circuit board.

For low-frequency filtering, it is a good idea to install a 10 - to $100\text{-}\mu\text{F}$ capacitor on the printed circuit board (one capacitor will suffice). If the capacitance is excessively large, however, (e.g. several thousand μF) latch-up can be a problem. Be sure to choose an appropriate capacitance value.

An important point about wiring is that, in the case of high-speed logic ICs, noise is caused

mainly by reflection and crosstalk, or by the power supply impedance. Reflections cause increased signal delay, ringing, overshoot and undershoot, thereby reducing the device's safety margins with respect to noise. To prevent reflections, reduce the wiring length by increasing the device mounting density so as to lower the inductance (L) and capacitance (C) in the wiring. Extreme care must be taken, however, when taking this corrective measure, since it tends to cause crosstalk between the wires. In practice, there must be a trade-off between these two factors.

1. External Noise

Printed circuit boards with long I/O or signal pattern lines are vulnerable to induced noise or surges from outside sources. Consequently, malfunctions or breakdowns can result from overcurrent or overvoltage, depending on the types of device used. To protect against noise, lower the impedance of the pattern line or insert a noise-canceling circuit. Protective measures must also be taken against surges. For details of the appropriate protective measures for a particular device, consult the relevant databook.

1. Electromagnetic Interference

Widespread use of electrical and electronic equipment in recent years has brought with it radio and TV reception problems due to electromagnetic interference. To use the radio spectrum effectively and to maintain radio communications quality, each country has formulated regulations limiting the amount of electromagnetic interference which can be generated by individual products.

Electromagnetic interference includes conduction noise propagated through power supply and telephone lines, and noise from direct electromagnetic waves radiated by equipment. Different measurement methods and corrective measures are used to assess and counteract each specific type of noise. Difficulties in controlling electromagnetic interference derive from the fact that there is no method available which allows designers to calculate, at the design stage, the strength of the electromagnetic waves which will emanate from each component in a piece of equipment. For this reason, it is only after the prototype equipment has been completed that the designer can take measurements using a dedicated instrument to determine the strength of electromagnetic interference waves.

Yet it is possible during system design to incorporate some measures for the prevention of electromagnetic interference, which can facilitate taking corrective measures once the design has been completed. These include installing shields and noise filters, and increasing the thickness of the power supply wiring patterns on the printed circuit board. One effective method, for example, is to devise several shielding options during design, and then select the most suitable shielding method based on the results of measurements taken after the prototype has been completed.

1. Peripheral Circuits

In most cases semiconductor devices are used with peripheral circuits and components. The input/output signal voltages and currents of these circuits must be chosen to match the semiconductor device's specifications. The

following factors must be taken into account.

- (1) Inappropriate voltages or currents applied to a device's input pins may cause it to operate erratically. Some devices contain pull-up/pull-down resistors. When designing your system, remember to take the effect of this on the required voltage and current levels into account.
- (2) The output pins on a device have a predetermined external circuit drive capability. If this drive capability is greater than that required, either incorporate a compensating circuit into your design or carefully select suitable components for use in external circuits.

1. Safety Standards

Each country has safety standards which must be observed. For example, for devices that handle high voltages, it is often required that an appropriate insulation distance be maintained between the device proper and the conductor pattern on the printed circuit board. Such requirements must be fully taken into account to ensure that your design conforms to the applicable safety standards.

1. Other Precautions

- (1) When designing a system, be sure to incorporate fail-safe and other appropriate measures according to the intended purpose of your system. Also, be sure to debug your system under actual board-mounted conditions.
- (2) If a plastic-package device is placed in a strong electric field, surface leakage may occur due to the charge-up phenomenon, resulting in device malfunction. In such cases, take appropriate measures to prevent this problem, for example by protecting the package surface with a conductive shield.
- (3) With some microcomputers and MOS memory devices, caution is required when powering on or resetting the device. To ensure that your design does not violate device specifications, consult the relevant databook for each constituent device.
- (4) Ensure that no conductive material or object (such as a metal pin) can drop onto and short the leads of a device mounted on a printed circuit board.

A. Inspection, Testing and Evaluation

1. Grounding

Ground all measuring instruments, jigs, tools and soldering irons to earth. Electrical leakage may cause a device to break down or may result in electric shock.

1. Inspection Sequence

symbol 129 \f "Wingdings" \s 11①“ Do not insert devices in the wrong orientation. Make sure that the positive and negative electrodes of the power supply are correctly connected. Otherwise, the rated maximum current or maximum power dissipation may be exceeded and the device may break down or undergo performance degradation, causing it to catch fire or explode, resulting in injury to the user.

symbol 130 \f "Wingdings" \s 11②“ When conducting any kind of evaluation, inspection or testing using AC power with a peak voltage of approximately 45 V or DC power exceeding 60 V, be sure to connect the electrodes or probes of the testing equipment to the device under test before powering it on. Connecting the electrodes or probes of testing equipment to a device while it is powered on may result in electric shock, causing injury.

- (1) Before beginning device inspection, make a final check to ensure that all associated equipment is properly grounded to earth and that there is no electrical leakage as described above. Apply voltage to the test jig only after inserting the device securely into it. (Do not power the test jig up or down abruptly; always apply or remove power gradually or in steps.)

- (2) Make sure that the voltage applied to the device is off before removing the device from the test jig. Otherwise, the device may undergo performance degradation or be destroyed.
- (3) Make sure that no surge voltages from the measuring equipment are applied to the device.
- (4) The chips housed in tape carrier packages (TCPs) are bare chips and are therefore exposed. During inspection take care not to crack the chip or cause any flaws in it.
Electrical contact may also cause a chip to become faulty. Therefore make sure that nothing comes into electrical contact with the chip.

A. Mounting

There are essentially two main types of semiconductor device package: lead insertion and surface mount. During mounting on printed circuit boards, devices can become contaminated by flux or damaged by thermal stress from the soldering process. Particularly with surface mount devices, the most significant problem is thermal stress from solder reflow, when the entire package is subjected to heat. This section describes a recommended temperature profile for each mounting method, as well as general precautions which you should take when mounting devices on printed circuit boards. Note, however, that even for devices with the same package type, the appropriate mounting method varies according to the size of the chip and the size and shape of the lead frame. Therefore, please consult the appropriate technical datasheet or databook.

1. Lead Forming

symbol 129 \f "Wingdings" \s 11①“ Always wear protective glasses when cutting the leads of a device with clippers or a similar tool. If you do not, small bits of metal flying off the cut ends may damage your eyes.

symbol 130 \f "Wingdings" \s 11②“ Because some types of device have leads with pointed tips, you may prick your finger.

Semiconductor devices must undergo a process in which the leads are cut and formed before the devices can be mounted on a printed circuit board. If undue stress is applied to the interior of a device during this process, mechanical breakdown or performance degradation can result. This is attributable primarily to differences between the stress on the device's external leads and the stress on the internal leads. If the relative difference is great enough, the device's internal leads, adhesive properties or sealant can be damaged. Observe these precautions during the lead forming process (this does not apply to surface mount devices):

- (1) Lead insertion hole intervals on the printed circuit board should match the lead pitch of the device precisely.
- (2) If lead insertion hole intervals on the printed circuit board do not precisely match the lead pitch of the device, do not attempt to forcibly insert devices by pressing on them or by pulling on their leads.
- (3) For the minimum clearance specification between a device and a printed circuit board, refer to the relevant device's datasheet or databook. If necessary, achieve the required clearance by forming the device's leads appropriately. Do not use the spacers which are used to raise

3 General Safety Precautions and Usage Considerations

devices above the surface of the printed circuit board during soldering to achieve clearance. These spacers normally continue to expand due to heat, even after the solder has begun to solidify; this applies severe stress to the device.

- (4) Do not repeatedly bend or stretch device leads.
- (5) Observe the following precautions when forming the leads of a device prior to mounting.
 - (a) Use a tool or jig to secure the lead at its base (where the lead meets the device package) while bending.
 - (b) Maintain a certain distance between the device package and the tool or jig.
 - (c) When forming a lead by bending it over a jig surface, be careful not to damage the lead on the edge of the jig surface.
 - (d) Follow any other precautions described in the individual datasheets and databooks for each device and package type.

1. Socket Mounting

- (1) When socket mounting devices on a printed circuit board, use sockets which match the inserted device's package.
- (2) Use sockets whose contacts have the appropriate contact pressure. If the contact pressure is insufficient, the socket may not make a perfect contact when the device is repeatedly inserted and removed; if the pressure is excessively high, the device leads may be bent or damaged when they are inserted into or removed from the socket.
- (3) When soldering sockets to the printed circuit board, use sockets whose construction prevents flux from penetrating into the contacts or which allows flux to be completely cleaned off.
- (4) Make sure the coating agent applied to the printed circuit board for moisture-proofing purposes does not stick to the socket contacts.
- (5) If the device leads are severely bent by a socket as it is inserted or removed and you wish to repair the leads so as to continue using the device, make sure that this lead correction is only performed once. Do not use devices whose leads have been corrected more than once.
- (6) If the printed circuit board with the devices mounted on it will be subjected to vibration from external sources, use sockets which have a strong contact pressure so as to prevent the sockets and devices from vibrating relative to one another.

1. Soldering Temperature Profile

The soldering temperature and heating time vary from device to device. Therefore, when specifying the mounting conditions, refer to the individual datasheets and databooks for the devices used.

a) Using a Soldering Iron

Complete soldering within ten seconds for lead temperatures of up to 260°C, or within three seconds for lead temperatures up to 350°C.

a) Using Medium Infrared Ray Reflow

- (a) Heating top and bottom with long or medium infrared rays is recommended
(see ref _Ref420308942 * Mergeformat Figure 3.5.3.1).

Figure 3.5.3.1 Heating Top and Bottom with Long or Medium Infrared Rays

3 General Safety Precautions and Usage Considerations

- (b) Complete the infrared ray reflow process within 30 seconds at a package surface temperature of between 210°C and 240°C.
- (c) Refer to ref _Ref420309387 * Mergeformat Figure 3.5.3.2 for an example of a good temperature profile for infrared or hot air reflow.

Figure styleref 3 \n 3.5.3.seq Figure * Arabic 2 Sample Temperature Profile for Infrared or Hot Air Reflow

a) Using Hot Air Reflow

- (a) Complete hot air reflow within 30 seconds at a package surface temperature of between 210°C and 240°C.
- (b) For an example of a recommended temperature profile, refer to Figure 3.5.3.2 above.

a) Using Vapor Phase Reflow Soldering (VPS)

- (a) The recommended solvent is Fluorinate FC-70 or equivalent.
- (b) Complete hot air reflow within 30 seconds at an ambient atmospheric temperature of 215°C, or within 60 seconds at an ambient atmospheric temperature of 200°C.
- (c) Refer to ref _Ref420309564 * Mergeformat Figure 3.5.3.3 for an example of a good temperature profile for vapor phase reflow soldering.

Figure styleref 3 \n 3.5.3.seq Figure * Arabic 3 Example Temperature Profile for Vapor Phase Reflow

a) Using Solder Flow

- (1) Apply preheating for 60 to 120 seconds at a temperature of 150°C.
- (2) For insertion-type packages, complete solder flow within 10 seconds with the temperature at the stopper, or at a location more than 1.5 mm from the body if there is no stopper, which does not exceed 260°C.
- (3) For surface mount packages, complete soldering within 5 seconds at a temperature of 250°C or less in order to prevent thermal stress in the device.
- (4) Figure 3.5.3.4 shows an example of a recommended temperature profile for surface mount packages using solder flow.

Figure styleref 3 \n 3.5.3.seq Figure * Arabic 4 Sample Temperature Profile for Solder Flow

1. Flux Cleaning and Ultrasonic Cleaning

- (1) When cleaning circuit boards to remove flux, make sure that no residual reactive ions such as Na or Cl remain. Note that organic solvents react with water to generate hydrogen chloride and other corrosive gases that can degrade device performance.
- (2) Washing devices with water will not cause any problems. However, make sure that no reactive ions such as sodium and chlorine are left as residues. Also, be sure to dry devices sufficiently after washing.
- (3) Do not rub device markings with a brush or with your hand during cleaning or while the devices are still wet from the cleaning agent. Doing so can rub off the markings.
- (4) The dip cleaning, shower cleaning and steam cleaning processes all involve the chemical action of a solvent. Use only recommended solvents for these cleaning methods. When immersing devices in a solvent or steam bath, make sure that the temperature of the liquid is 50°C or below, and that the circuit board is removed from the bath within one minute.
- (5) Ultrasonic cleaning can clean circuit boards efficiently in a short period of time. However, it should not be used with hermetically-sealed ceramic packages such as a leadless chip carrier (LCC), charge-coupled device (CCD) or pin grid array (PGA), because the bonding wires can become disconnected due to resonance during the cleaning process. Plastic packages do not have this problem. However, limit the duration of ultrasonic cleaning to as short a time as possible, since long hours of ultrasonic cleaning degrade the adhesion between the mold resin and the frame material. The following ultrasonic cleaning conditions are recommended:

Frequency: 27 to 29_kHz

Ultrasonic output power: 300_W or less (0.25_W/cm² or less)

Cleaning time: 30 seconds or less

Suspend the circuit board in the solvent bath during ultrasonic cleaning in such a way that the ultrasonic vibrator does not come into direct contact with the circuit board or the device.

Conventional cleaning solvents that contain freon are not recommended due to the danger that they pose to the earth's ozone layer. Alternative products listed below are available on the market. Some alternative cleaning agents that do not contain freon include:

- FRW-1, 17; FRV-100 from Toshiba Corporation
- AK-225AES from Asahi Glass Co., Ltd.
- 750H from Kao Co., Ltd.
- ST-100 from Arakawa Chemical Co., Ltd.

Contact Toshiba or a Toshiba distributor regarding cleaning conditions and other relevant information for each product type.

1. No Cleaning

If analog devices or high-speed devices are used without being cleaned, flux

residues may cause minute amounts of leakage between pins. Similarly, dew condensation, which occurs in environments containing residual chlorine when power to the device is on, may cause between-lead leakage or migration. Therefore, Toshiba recommends that these devices be cleaned.

However, if the flux used contains only a small amount of halogen (0.05% or less), the devices may be used without cleaning without any problems.

For details of individual devices' cleaning conditions, please contact Toshiba or a Toshiba distributor.

1. Mounting Tape Carrier Packages (TCPs)

- (1) When tape carrier packages (TCPs) are mounted, measures must be taken to prevent electrostatic breakdown of the devices.
- (2) If devices are being picked up from tape, or outer lead bonding (OLB) mounting is being carried out, consult the manufacturer of the insertion machine which is being used, in order to establish the optimum mounting conditions in advance and to avoid any possible hazards.
- (3) The base film, which is made of polyimide, is hard and thin. Be careful not to cut or scratch your hands or any objects while handling the tape.
- (4) When punching tape, try not to scatter broken pieces of tape too much.
- (5) Treat the extra film, reels and spacers left after punching as industrial waste, taking care not to destroy or pollute the environment.
- (6) Chips housed in tape carrier packages (TCPs) are bare chips and therefore have their reverse side exposed. To ensure that the chip will not be cracked during mounting, ensure that no mechanical shock is applied to the reverse side of the chip. Electrical contact may also cause a chip to fail. Therefore, when mounting devices, make sure that nothing comes into electrical contact with the reverse side of the chip.

If your design requires connecting the reverse side of the chip to the circuit board, please consult Toshiba or a Toshiba distributor beforehand.

1. Mounting Chips

Devices delivered in chip form tend to degrade or break under external forces much more easily than plastic-packaged devices. Therefore, caution is required when handling this type of device.

- (1) Mount devices in a well-prepared environment so that chip surfaces will not be exposed to polluted ambient air or other polluted substances.
- (2) When handling chips, be careful not to expose them to static electricity. In particular, measures must be taken to prevent static damage during the mounting of chips. With this in mind, Toshiba recommends mounting all peripheral parts first and then mounting chips last (after all other components have been mounted).
- (3) Make sure that circuit boards (e.g. PCBs) on which chips are being mounted do not have any chemical residues on them (such as the chemicals which were used for etching the boards).
- (4) When mounting chips on a board, use the method of assembly that is most

suitable for maintaining the appropriate electrical, thermal and mechanical properties of semiconductor devices.

* For details of devices in chip form, refer to the relevant devices' individual datasheets.

1. **Circuit Board Coating**

When devices are to be used in equipment requiring a high degree of reliability or in extreme environments (where moisture, corrosive gas or dust is present), circuit boards may be coated for protection. However, before doing so, you must carefully consider the possible stress and contamination effects that may result and choose the coating resin which applies the minimum level of stress to the device.

1. **Heat Sinks**

- (1) When attaching a heat sink to a device, be careful not to apply excessive force to the device in the process.
- (2) When attaching a device to a heat sink by fixing it at two or more locations, evenly tighten all the screws in stages (i.e. do not fully tighten one screw while the rest are still only loosely tightened). Finally, fully tighten all the screws up to the specified torque.
- (3) Drill holes for screws in the heat sink exactly as specified. Smooth the surface by removing burrs and protrusions or indentations which might interfere with the installation of any part of the device.
- (4) A coating of silicone compound can be applied between the heat sink and the device to improve heat conductivity. Be sure to apply the coating thinly and evenly; do not use too much. Also, be sure to use a nonvolatile compound, as volatile compounds can crack after a time, causing the heat radiation properties to deteriorate.
- (5) If the device is housed in a plastic package, use caution when selecting the type of silicone compound to be applied between the heat sink and the device. With some types, the base oil separates and penetrates the plastic package, significantly reducing the useful life of the device.

Two recommended silicone compounds in which base oil separation is not a problem are YG6260 from Toshiba Silicone and G746 from Shinetsu Chemical Industries.

- (6) Heat-sink-equipped devices can become very hot during operation. Do not touch them, or you may sustain a burn.

1. **Tightening Torque**

- (1) Make sure the screws are tightened with fastening torques not exceeding the torque values stipulated in individual datasheets and databooks for the devices used.
- (2) Do not allow a power screwdriver (electrical or air-driven) to touch devices.

A. Protecting Devices in the Field

1. Temperature

Semiconductor devices are generally more sensitive to temperature than are other electronic components. The various electrical characteristics of a semiconductor device are dependent on the ambient temperature at which the device is used. It is therefore necessary to understand the temperature characteristics of a device and to incorporate device derating into circuit design. Note also that if a device is used above its maximum temperature rating, device deterioration is more rapid and it will reach the end of its usable life sooner than expected.

1. Humidity

- (1) Resin-molded devices are sometimes improperly sealed. When these devices are used for an extended period of time in a high-humidity environment, moisture can penetrate into the device and cause chip degradation or malfunction. Furthermore, when devices are mounted on a regular printed circuit board, the impedance between wiring components can decrease under high-humidity conditions. In systems that require a high signal-source impedance, circuit board leakage or leakage between device lead pins can cause malfunctions. The application of a moisture-proof treatment to the device surface should be considered in this case. On the other hand, operation under low-humidity conditions can damage a device due to the occurrence of electrostatic discharge. Unless damp-proofing measures have been specifically carried out, use devices only in environments with appropriate ambient moisture levels (i.e. within a relative humidity range of 40% to 60%).
- (2) When semiconductor devices are to be used in equipment requiring a high degree of reliability or in extreme environments (where humidity is high, or where corrosive gas or dust is present), devices may be coated in order to moisture-proof them. In such cases, choose the coating resin which applies the minimum level of stress to the device.

1. Corrosive Gases

Corrosive gases can cause chemical reactions in devices, degrading device characteristics.

For example, sulphur-bearing corrosive gases emanating from rubber placed near a device (accompanied by condensation under high-humidity conditions) can corrode a device's leads. The resulting chemical reaction between leads forms foreign particles which can cause electrical leakage.

1. Radioactive and Cosmic Rays

Most industrial and consumer semiconductor devices are not designed with protection against radioactive and cosmic rays. Devices used in aerospace equipment or in radioactive environments must therefore be shielded.

1. Strong Electrical and Magnetic Fields

Devices exposed to strong magnetic fields can undergo a polarization phenomenon in

3 General Safety Precautions and Usage Considerations

plastic material, or within the chip, which gives rise to abnormal symptoms such as impedance changes or increased leakage current. Failures have been reported in LSIs mounted near malfunctioning deflection yokes in TV sets. In such cases, the device's installation location must be changed or the device must be shielded against the electrical or magnetic field. Shielding against magnetism is especially necessary for devices used in an alternating magnetic field, because of the electromotive forces generated in this type of environment.

1. Interference from Light (ultraviolet rays, sunlight, fluorescent lamps and incandescent lamps)

Light striking a semiconductor device generates electromotive force due to photoelectric effects. In some cases the device can malfunction. This is especially true for devices in which the internal chip is exposed. When designing circuits, make sure devices are protected against incident light from external sources. This problem is not limited to optical semiconductors and EPROMs. All types of device can be affected by light.

1. Dust and Oil

Just like corrosive gases, dust and oil can cause chemical reactions in devices, which will adversely affect a device's electrical characteristics. To avoid this problem, do not use devices in dusty or oily environments. This is especially important for optical devices because dust and oil can affect a device's optical characteristics as well as its physical integrity and the electrical performance factors mentioned above.

1. Fire

Semiconductor devices may be combustible; they can emit smoke and catch fire if heated sufficiently. When this happens, some devices may generate poisonous gases. Devices should therefore never be used in close proximity to an open flame or a heat-generating body, or near flammable or combustible materials.

A. Disposal of Devices and Packing Materials

When discarding unused devices and packing materials, follow all procedures specified by local regulations in order to protect the environment against contamination.

3 General Safety Precautions and Usage Considerations

I. Product-Specific Precautions and Usage Considerations

This section describes matters specific to each product group that need to be taken into consideration when using devices. If the same item is described in Sections 3 and 4, the description in Section 4 takes precedence.

A. Microcontrollers

1. Design

a) Using Resonators not Listed Under “Recommended Types”

Resonators recommended for use with Toshiba products in microcontroller oscillator applications are listed in Toshiba databooks along with information about oscillation conditions. If you use a resonator not included in this list, please consult Toshiba or the resonator manufacturer concerning the suitability of the device for your application.

a) Undefined Functions

In some microcontrollers certain instruction code values do not constitute valid processor instructions. Also, it is possible that the values of bits in registers will become undefined. Take care in your applications not to use invalid instructions or to let register bit values become undefined.

Appendix A: CPU Instruction Set Details

This appendix provides a detailed description of the operation of each TX49 instruction in both 32- and 64-bit modes. The instructions are listed in alphabetical order.

The exceptions that may occur due to the execution of each instruction are listed after the description of each instruction. The description of the immediate causes and manner of handling exceptions is omitted from the instruction descriptions in this chapter.

Figures at the end of this appendix list the bit encoding for the constant fields of each instruction, and the bit encoding for each individual instruction is included with that instruction.

Instruction Classes

The TX49 has some classes of CPU instructions, as follows.

- Load and Store
- Computational
- Jump and Branch
- Coprocessor
- Special
- Exception
- Multiply and Divide
- Debug
- Others

Instruction Formats

Every instruction consists of a single word (32 bits) aligned on a word boundary. The main instruction formats are shown in Figure A-1.

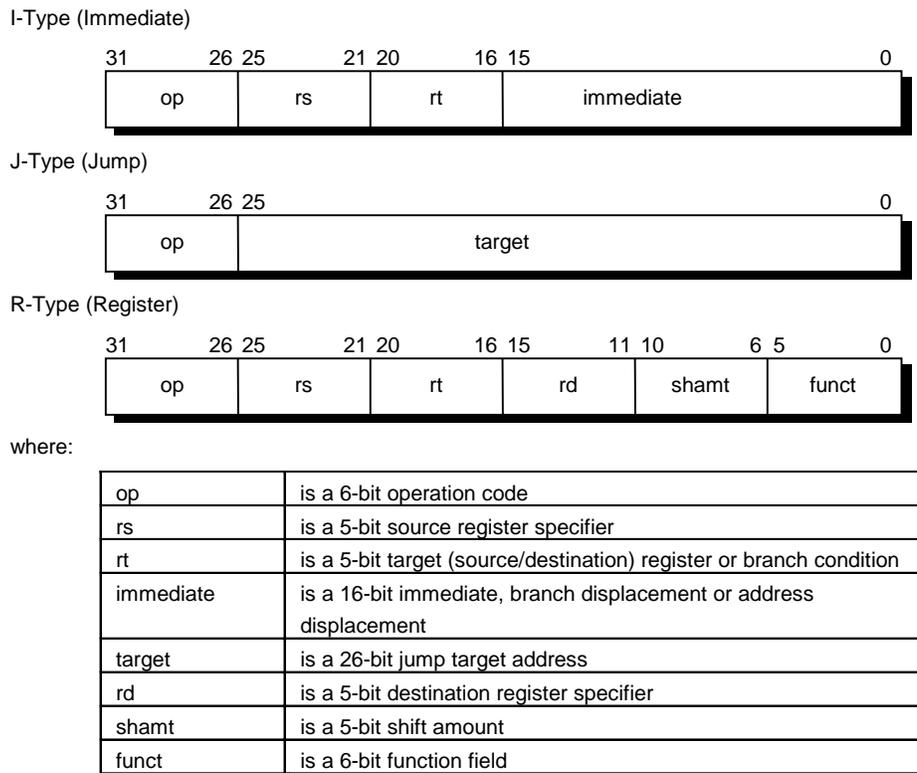


Figure A-1. CPU Instruction Formats

Instruction Notation Conventions

In this appendix, all variable subfields in an instruction format (such as *rs*, *rt* *immediate*, etc.) are shown in lowercase names.

For the sake of clarity, we sometimes use an alias for a variable subfield in the formats of specific instructions. For example, we use *rs* = *base* in the format for load and store instructions. Such an alias is always lower case, since it refers to a variable subfield.

Figures with the actual bit encoding for all the mnemonics are located at the end of this Appendix, and the bit encoding also accompanies each instruction.

In the instruction descriptions that follow, the *Operation* section describes the operation performed by each instruction using a high-level language notation. The TX49 can operate as either a 32- or 64-bit microprocessor. The operation for both modes is included with the instruction description. Special symbols used in the notation are described in Table A-1.

Table A-1 CPU Instruction Operation Notations

Symbol	Meaning
\leftarrow	Assignment.
$\ $	Bit string concatenation.
x^y	Replication of bit value x into a y -bit string. Note: x is always a single-bit value.
$x_{y...z}$	Selection of bits y through z of bit string x . Little-endian bit notation is always used. If y is less than z , this expression is an empty (zero length) bit string.
$+$	Two's complement or floating-point addition.
$-$	Two's complement or floating-point subtraction.
$*$	Two's complement or floating-point multiplication.
Div	Two's complement integer division.
Mod	Two's complement modulo.
$/$	Floating-point division.
$<$	Two's complement less than comparison.
And	Bitwise logic AND.
Or	Bitwise logic OR.
Xor	Bitwise logic XOR.
Nor	Bitwise logic NOR.
GPR[x]	General-Register x . The content of GPR[0] is always zero. Attempts to alter the content of GPR[0] have no effect.
CPR[z,x]	Coprocessor unit z , general register x .
CCR[z,x]	Coprocessor unit z , control register x .
COC[z]	Coprocessor unit z condition signal.
BigEndianMem	Big-endian mode as configured at reset (0 \rightarrow Little, 1 \rightarrow Big). Specifies the endianness of the memory interface (see LoadMemory and StoreMemory), and the endianness of Kernel and Supervisor mode execution.
ReverseEndian	Signal to reverse the endianness of load and store instructions. This feature is available in User mode only, and is effected by setting the <i>RE</i> bit of the <i>Status</i> register. Thus, ReverseEndian may be computed as (SR ₂₅ and User mode)
BigEndianCPU	The endianness for load and store instructions (0 \rightarrow Little, 1 \rightarrow Big). In User mode, this endianness may be reversed by setting SR ₂₅ . Thus, BigEndianCPU may be computed as BigEndianMem XOR ReverseEndian.
Llbit	Bit of state to specify synchronization instructions. Set by LL, cleared by ERET and Invalidate and read by SC.
$T \bullet + \bullet i$	Indicates the time steps between operations. Each of the statements within a time step are defined to be executed in sequential order (as modified by conditional and loop constructs). Operations which are marked $T \bullet + \bullet i$ are executed at instruction cycle i relative to the start of execution of the instruction. Thus, an instruction which starts at time j executes operations marked $T \bullet + \bullet i$ at time $j + \bullet i$. The interpretation of the order of execution between two instructions or two operations which execute at the same time should be pessimistic; the order is not defined.

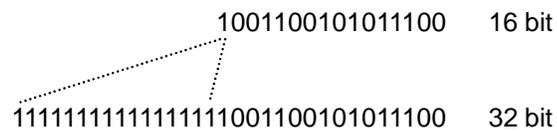
Sign Extension and Zero Extension

With some instructions the bit length may be extended; for example, a 16-bit offset may be extended to 32 bits. This extension can take the form of either a sign extension or zero extension.

- Sign extension

The extended part is filled with the value of the most significant bit.

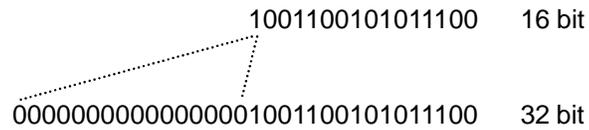
(example)



- Zero extension

The extended part is filled with zeros.

(example)



Instruction Notation Examples

The Following examples illustrate the application of some of the instruction notation conventions:

Example #1:
$GPR[rt] \leftarrow \text{immediate} \parallel 0^{16}$
Sixteen zero bits are concatenated with an immediate value (typically 16• bits), and the 32-bit string (with the lower 16• bits set to zero) is assigned to General-Purpose Register <i>rt</i> .
Example #2:
$(\text{immediate}_{15})^{16} \parallel \text{immediate}_{15-0}$
Bit 15 (the sign• bit) of an immediate value is extended for 16• bit positions, and the result is concatenated with bits 15 through 0 of the immediate value to form a 32-bit sign extended value.

Load and Store Instructions

In the TX49 implementation, the instruction immediately following a load may use the contents of the register loaded. In such cases, the hardware interlocks, requiring additional real cycles, so scheduling load delay slots is still desirable, although not required for functional code.

Two special instructions are provided in the TX49 implementation of the MIPS ISA, Load Linked and Store Conditional. These instructions are used in carefully coded sequences to provide one of several synchronization primitives, including test-and-set, bit-level locks, semaphores, and sequencers / event counts.

In the load and store operation descriptions, the functions listed in Table A-2 are used to summarize the handling of virtual addresses and physical memory.

Table A-2 Load and Store Common Functions

Function	Meaning
AddressTranslation	Uses the TLB to find the physical address given the virtual address. The function fails and an exception is taken if the required translation is not present in the TLB.
LoadMemory	Uses the cache and main memory to find the contents of the word containing the specified physical address. The low-order two bits of the address and the access type field indicates which of each of the four bytes within the data word need to be returned. If the cache is enabled for this access, the entire word is returned and loaded into the cache.
StoreMemory	Uses the cache, write buffer, and main memory to store the word or part of word specified as data in the word containing the specified physical address. The low-order two bits of the address and the access type field indicates which of each of the four bytes within the data word should be stored.

The access type field indicates the size of the data item to be loaded or stored as shown in Table A-3. Regardless of access type or byte-numbering order (endianness), the address specifies the byte which has the smallest byte address of the bytes in the addressed field. For a Big-endian machine, this is the leftmost byte and contains the sign for a 2's-complement number; for a Little-endian machine, this is the rightmost byte and contains the lowest precision byte.

Table A-3 Access Type Specifications for Loads/Stores

Access Type Mnemonic	Value	Meaning
DOUBLEWORD	7	doubleword (64• bits)
SEPTIBYTE	6	seven bytes (56• bits)
SEXTIBYTE	5	six bytes (48• bits)
QUINTIBYTE	4	five bytes (40• bits)
WORD	3	word (32• bits)
TRIPLEBYTE	2	triple-byte (24• bits)
HALFWORD	1	halfword (16• bits)
BYTE	0	byte (8• bits)

The bytes within the addressed doubleword which are used can be determined directly from the access type and the three low-order bits of the address, as shown in Chapter 2.

Jump and Branch Instructions

All jump and branch instructions have an architectural delay of exactly one instruction. That is, the instruction immediately following a jump or branch (i.e., occupying the delay slot) is always executed while the target instruction is being fetched from storage. It is not valid for a delay slot to be occupied itself by a jump or branch instruction; however, this error is not detected, and the results of such an operation are undefined.

If an exception or interrupt prevents the completion of a legal instruction during a delay slot, the hardware sets the *EPC* register to point at the jump or branch instruction which precedes it. When the code is restarted, both the jump or branch instructions and the instruction in the delay slot are reexecuted.

Because jump and branch instructions may be restarted after exceptions or interrupts, they must be restartable. Therefore, when a jump or branch instruction stores a return link value, register 31 (the register in which the link is stored) may not be used as a source register.

Since instructions must be word-aligned, a Jump Register or Jump and Link Register instruction must use a register whose two low-order bits are zero. If these low-order bits are not zero, an address exception will occur when the jump target instruction is subsequently fetched.

Coprocessor Instructions

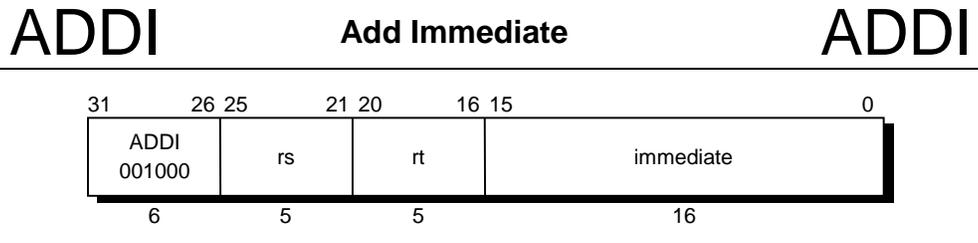
The MIPS architecture provides four coprocessor units, or classes. Coprocessors are alternate execution units, which have separate register files from the CPU. R-Series coprocessors have 2 register spaces, each with thirty-two 32-bit registers. The first space, *coprocessor general* registers, may be directly loaded from memory and stored into memory, and their contents may be transferred between the coprocessor and processor. The second, *coprocessor control* registers, may only have their contents transferred directly between the coprocessor and processor. Coprocessor instructions may alter registers in either space. Normally, by convention, *Coprocessor Control Register 0* is interpreted as a *Coprocessor Implementation And Revision* register. However, the system control coprocessor (CP0) uses *Coprocessor General Register 15* for the processor / coprocessor revision register. The register's low-order byte (bits 7-0) is interpreted as a coprocessor unit revision number. The second byte (bits 15-8) is interpreted as a coprocessor unit implementation descriptor. The revision number is a value of the form *y.x* where *y* is a major revision number in bits 7-4 and *x* is a minor revision number in bits 3-0.

The contents of the high-order halfword of the register are not defined (currently read as 0 and should be 0 when written).

System Control Coprocessor (CP0) Instructions

There are some special limitations imposed on operations involving CP0 that is incorporated within the CPU. Although load and store instructions to transfer data to and from coprocessors and move control to/from coprocessor instructions are generally permitted by the MIPS architecture, CP0 is given a somewhat protected status since it has responsibility for exception handling and memory management. Therefore, the move to/from coprocessor instructions are the only valid mechanism for reading from and writing to the CP0 registers.

Several coprocessor operation instructions are defined for CP0 to directly read, write, and probe TLB entries and to modify the operating modes in preparation for returning to User mode or interrupt-enabled states.



Format:

ADDI *rt*, *rs*, *immediate*

Description:

The 16-bit *immediate* is sign-extended and added to the contents of general register *rs* to form the result. The result is placed into general register *rt*. In 64-bit mode, the operand must be valid sign-extended, 32-bit values.

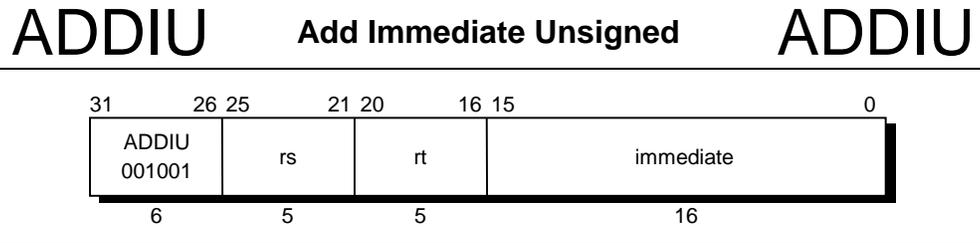
An overflow exception occurs if carries out of bits 30 and 31 differ (2's-complement overflow). The destination register *rt* is not modified when an integer overflow exception occurs.

Operation:

32	T:	$GPR[rt] \leftarrow GPR[rs] + (immediate_{15})^{16} \parallel immediate_{15-0}$
64	T:	$temp \leftarrow GPR[rs] + (immediate_{15})^{48} \parallel immediate_{15-0}$ $GPR[rt] \leftarrow (temp_{31})^{32} \parallel temp_{31-0}$

Exceptions:

Integer overflow exception



Format:

ADDIU *rt*, *rs*, *immediate*

Description:

The 16-bit *immediate* is sign-extended and added to the contents of general register *rs* to form the result. The result is placed into general register *rt*. No integer overflow exception occurs under any circumstances. In 64-bit mode, the operand must be valid sign-extended, 32-bit values.

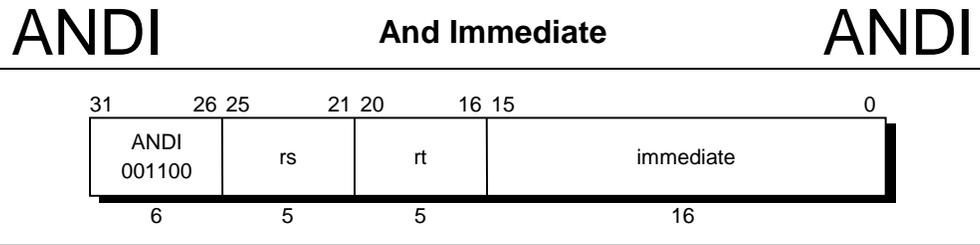
The only difference between this instruction and the ADDI instruction is that ADDIU never causes an overflow exception.

Operation :

32	T:	$GPR[rt] \leftarrow GPR[rs] + (immediate_{15})^{16} \parallel immediate_{15-0}$
64	T:	$temp \leftarrow GPR[rs] + (immediate_{15})^{48} \parallel immediate_{15-0}$ $GPR[rt] \leftarrow (temp_{31})^{32} \parallel temp_{31-0}$

Exceptions:

None



Format:

ANDI rt, rs, immediate

Description:

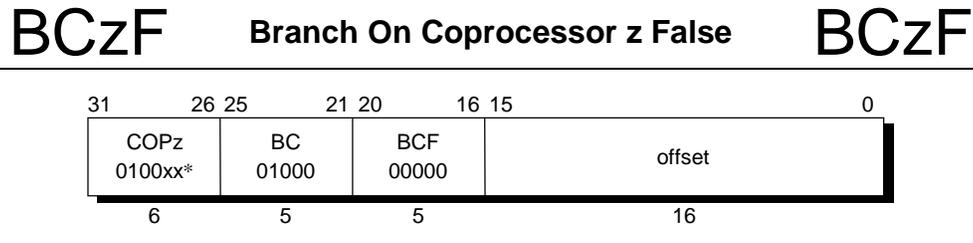
The 16-bit *immediate* is zero-extended and combined with the contents of general register *rs* in a bit-wise logical AND operation. The result is placed into general register *rt*.

Operation:

32	T:	$GPR[rt] \leftarrow 0^{16} \parallel (\text{immediate and } GPR[rs]_{15-0})$
64	T:	$GPR[rt] \leftarrow 0^{48} \parallel (\text{immediate and } GPR[rs]_{15-0})$

Exceptions:

None



Format:

BCzF offset

Description:

A branch target address is computed from the sum of the address of the instruction in the delay slot and the 16-bit *offset*, shifted left two bits and sign-extended. If coprocessor z's condition signal (CpCond), as sampled during the previous instruction, is false, then the program branches to the target address with a delay of one instruction.

Because the condition line is sampled during the previous instruction, there must be at least one instruction line between this instruction and a coprocessor instruction that changes the condition line.

Operation:

32	T-1:	condition \leftarrow <i>not</i> COC[z]
	T:	target \leftarrow (offset ₁₅) ¹⁴ offset 0 ²
	T • +• 1:	if condition then PC \leftarrow PC • +• target endif
64	T-1	condition \leftarrow <i>not</i> COC[z]
	T:	target \leftarrow (offset ₁₅) ⁴⁶ offset 0 ²
	T • +• 1:	if condition then PC \leftarrow PC • +• target endif

*See the table "Opcode Bit Encoding" on next page, or "CPU Instruction Opcode Bit Encoding" at the end of Appendix A.

BCzF

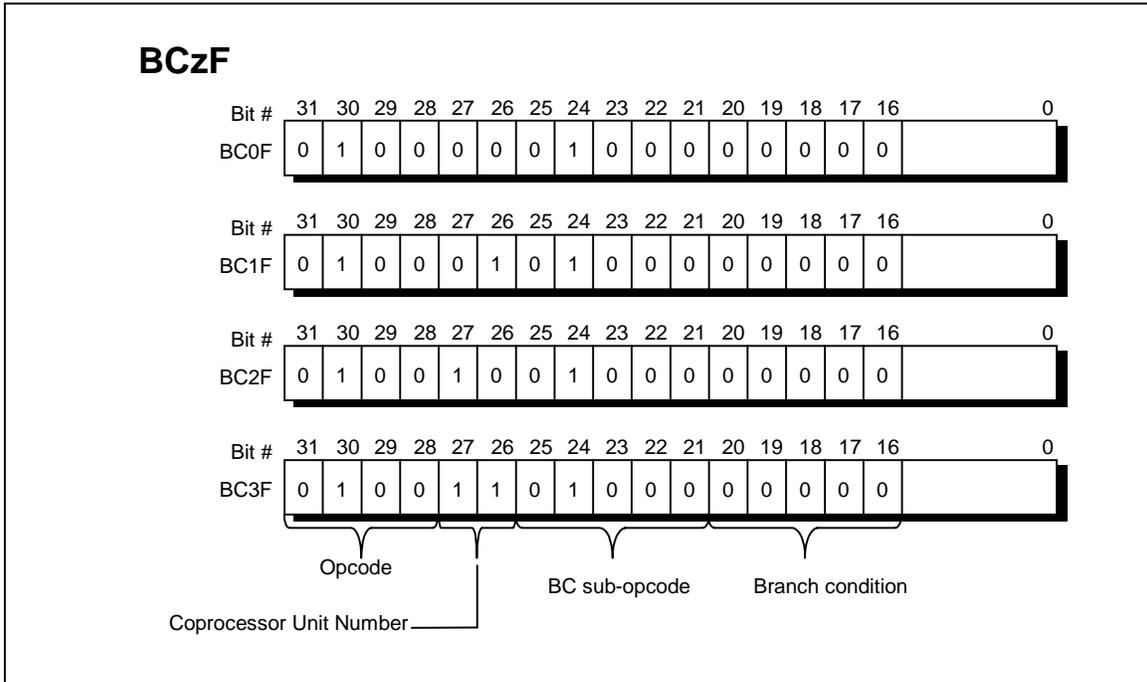
Branch On Coprocessor z False (continued)

BCzF

Exceptions:

Coprocessor unusable exception

Opcode Bit Encoding



Note:

CpCond0 = Write Buffer Empty

(Empty → true (1), Not empty → false (0))

CpCond1 = FPU (See the Appendix B)

CpCond2 = External Pin condition (GCPCOND2)

CpCond3 = External Pin condition (GCPCOND3)

Branch On Coprocessor Z False Likely (continued)

BCzFL

BCzFL

Operation:

```

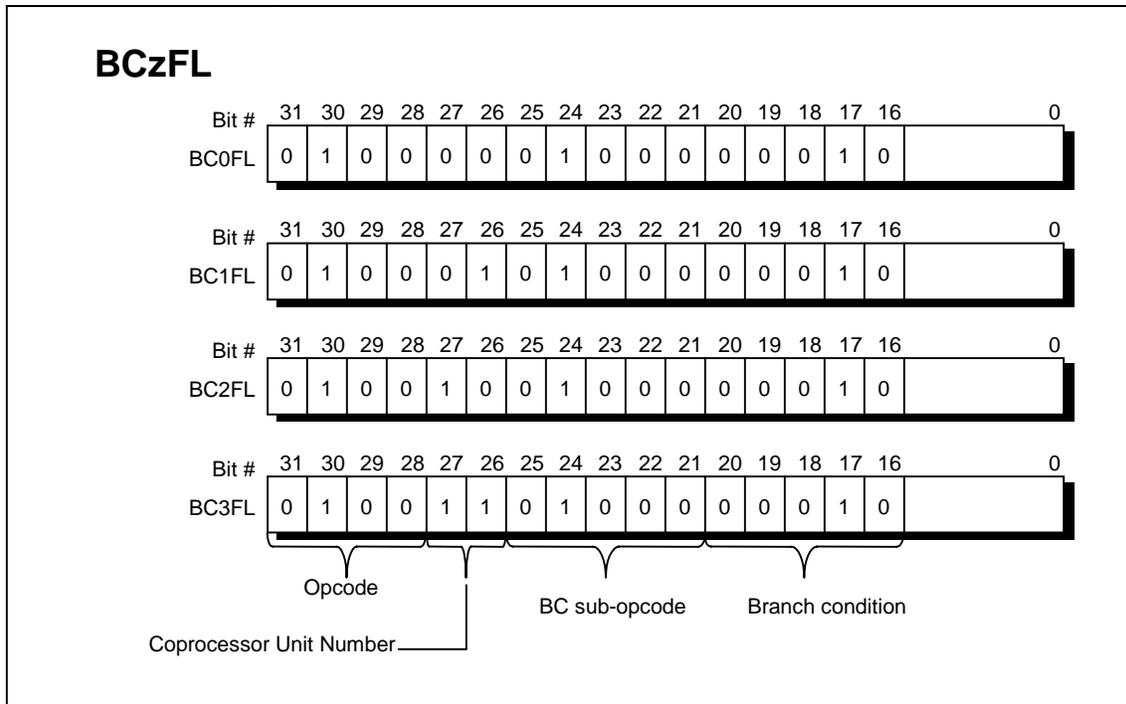
32 T-1:   condition ← not COC[z]
   T:     target ← (offset15)14 || offset || 02
   T• +• 1: if condition then
             PC ← PC• +• target
             else
             NullifyCurrentInstruction
             endif

64 T-1:   condition ← not COC[z]
   T:     target ← (offset15)46 || offset || 02
   T• +• 1: if condition then
             PC ← PC• +• target
             else
             NullifyCurrentInstruction
             endif
    
```

Exceptions:

Coprocessor unusable exception

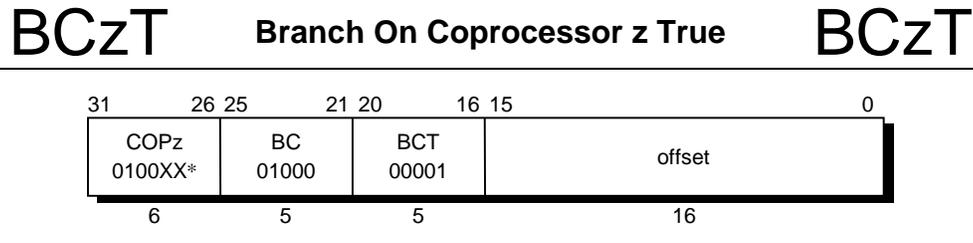
Opcode Bit Encoding:



Note:

- CpCond0 = Write Buffer Empty
(Empty → true (1), Not empty → false (0))
- CpCond1 = FPU (See the Appendix B)
- CpCond2 = External Pin condition (GCPCOND2)

CpCond3 = External Pin condition (GCPCOND3)



Format:

BCzT offset

Description :

A branch target address is computed from the sum of the address of the instruction in the delay slot and the 16-bit *offset*, shifted left two bits and sign-extended. If the coprocessor z's condition signal (CpCond) is true, then the program branches to the target address, with a delay of one instruction.

Because the condition line is sampled during the previous instruction, there must be at least one instruction between this instruction and a coprocessor instruction that changes the condition line.

Operation :

32	T-1:	condition \leftarrow COC[z]
	T:	target \leftarrow (offset ₁₅) ¹⁴ offset 0 ²
	T • +• 1:	if condition then
		PC \leftarrow PC • +• target
		endif
64	T-1	condition \leftarrow COC[z]
	T:	target \leftarrow (offset ₁₅) ⁴⁶ offset 0 ²
	T • +• 1:	if condition then
		PC \leftarrow PC • +• target
		Endif

*See the table "Opcode Bit Encoding" on next page, or "CPU Instruction Opcode Bit Encoding" at the end of Appendix A.

BCzT

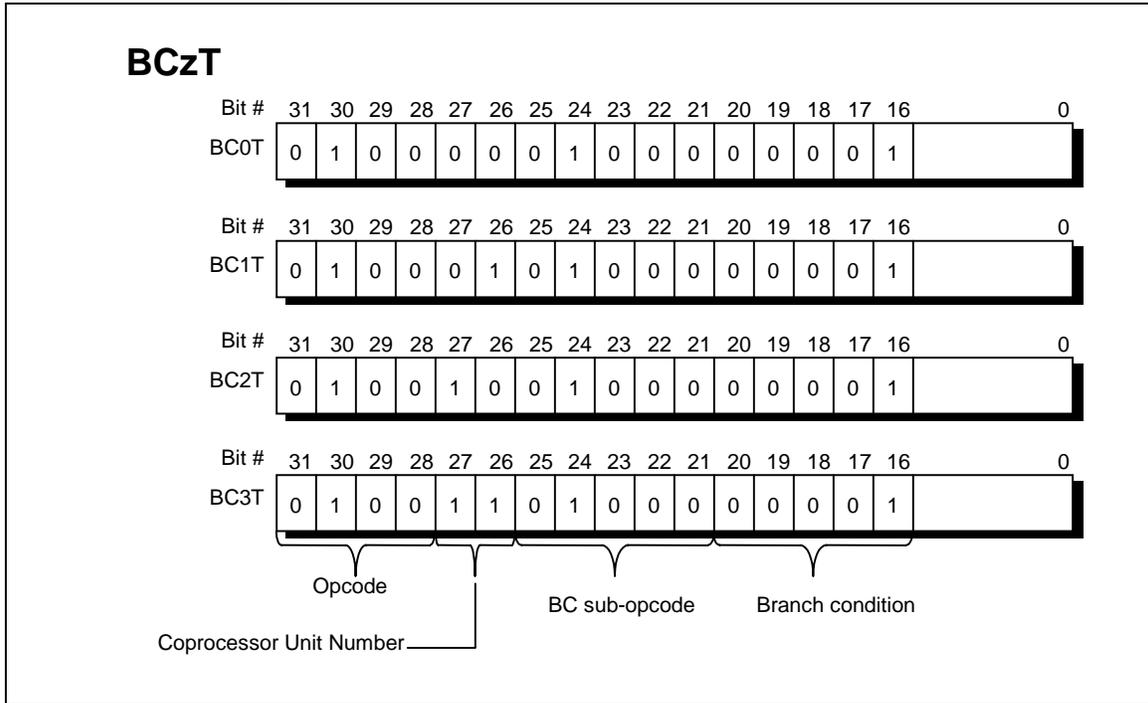
Branch On Coprocessor z True (continued)

BCzT

Exceptions:

Coprocessor unusable exception

Opcode Bit Encoding:



Note:

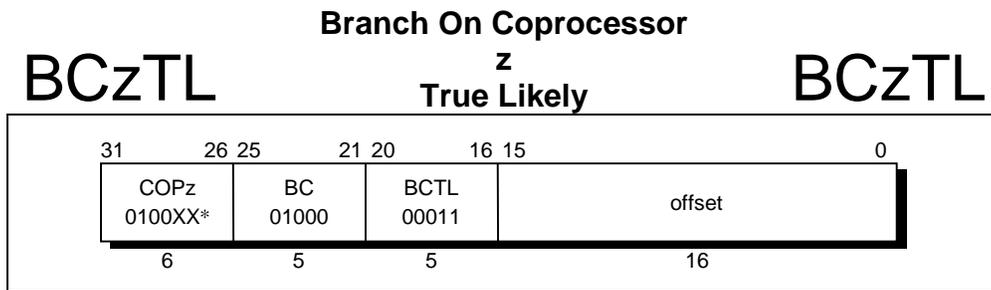
CpCond0 = Write Buffer Empty

(Empty → true (1), Not empty → false (0))

CpCond1 = FPU (See the Appendix B)

CpCond2 = External Pin condition (GCPCOND2)

CpCond3 = External Pin condition (GCPCOND3)



Format:

BCzTL offset

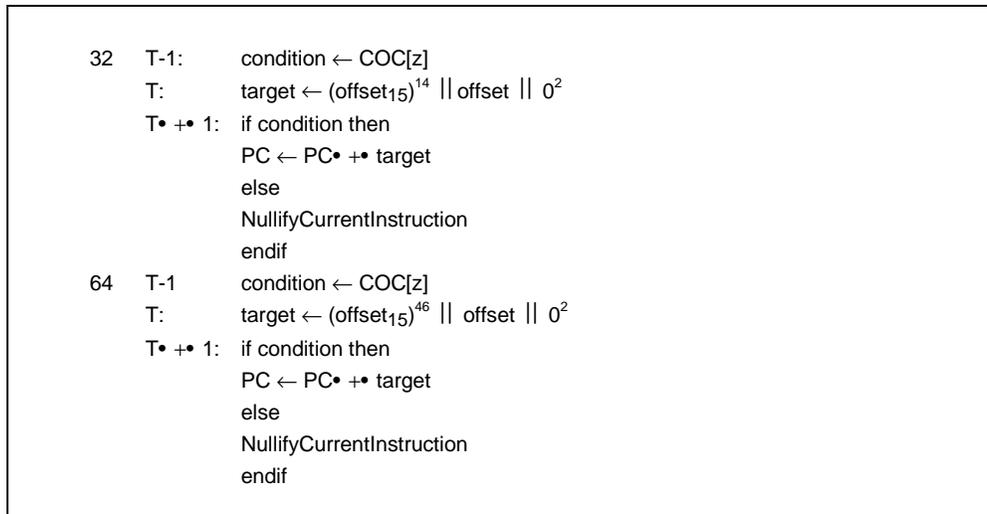
Description:

A branch target address is computed from the sum of the address of the instruction in the delay slot and the 16-bit *offset*, shifted left two bits and sign-extended. If the contents of coprocessor *z*'s condition line, as sampled during the previous instruction, is true, the target address is branched to with a delay of one instruction.

If the conditional branch is *not* taken, the instruction in the branch delay slot is nullified.

Because the condition line is sampled during the previous instruction, there must be at least one instruction between this instruction and a coprocessor instruction that changes the condition line.

Operation:



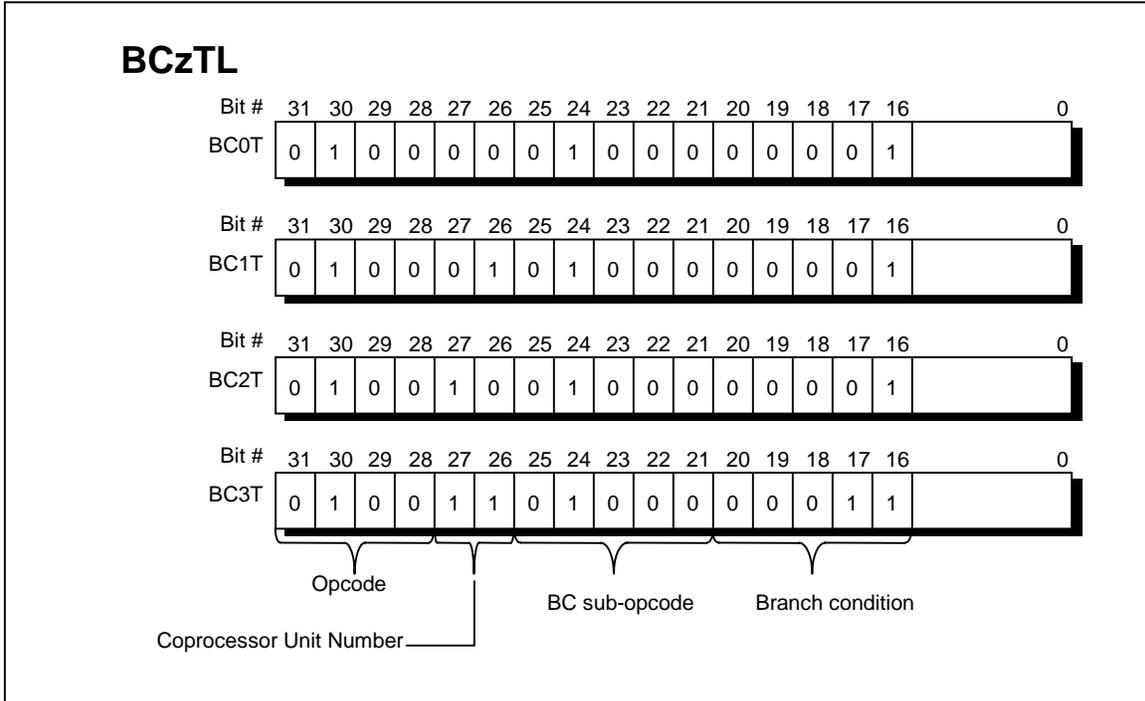
*See the table "Opcode Bit Encoding" on next page, or "CPU Instruction Opcode Bit Encoding" at the end of Appendix A.

BCzTL Branch On Coprocessor z True Likely (continued) BCzTL

Exceptions:

Coprocessor unusable exception

Opcode Bit Encoding:



Note:

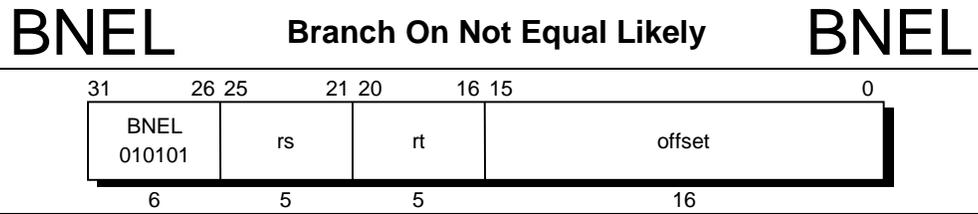
CpCond0 = Write Buffer Empty

(Empty → true (1), Not empty → false (0))

CpCond1 = FPU (See the Appendix B)

CpCond2 = External Pin condition (GCPCOND2)

CpCond3 = External Pin condition (GCPCOND3)

**Format:**

BNEL rs, rt, offset

Description:

A branch target address is computed from the sum of the address of the instruction in the delay slot and the 16-bit *offset*, shifted left two bits and sign-extended. The contents of general register *rs* and the contents of general register *rt* are compared. If the two registers are not equal, then the program branches to the target address, with a delay of one instruction.

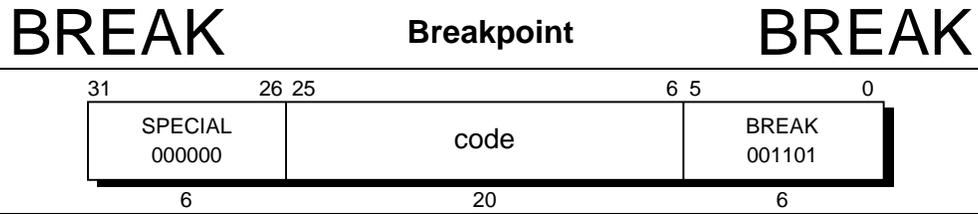
If the conditional branch is *not* taken, the instruction in the branch delay slot is nullified.

Operation:

32	T:	target \leftarrow (offset ₁₅) ¹⁴ offset 0 ² condition \leftarrow (GPR[rs] \neq GPR[rt])
	T + 1:	if condition then PC \leftarrow PC + target else NullifyCurrentInstruction endif
64	T:	target \leftarrow (offset ₁₅) ⁴⁶ offset 0 ² condition \leftarrow (GPR[rs] \neq GPR[rt])
	T + 1:	if condition then PC \leftarrow PC + target else NullifyCurrentInstruction endif

Exceptions:

None

**Format:**

BREAK

Description:

A breakpoint trap occurs, immediately and unconditionally transferring control to the exception handler.

The code field is available for use as software parameters, but is retrieved by the exception handler only by loading the contents of the memory word containing the instruction.

Operation:**Exceptions:**

Breakpoint exception

CACHE

Cache
(continued)

CACHE

Bits 20~18 of the instruction specify the operation as follows:

Code	Caches	Name	Operation
0	I	Index Invalidate	Set the cache state of the indexed block to invalid.
0	D	Index WriteBack Invalidate	Examine the cache state and W bit of the primary data cache block at the invalidate index specified by the virtual address. If the state is not invalid and the W bit is set, then write back the block to memory. The address to write is taken from the primary cache tag. Set cache state of primary cache block to invalid. LSB of VA select the way.
1	I / D	Index Load Tag	Read the tag for the cache block at the specified index and place it into the TagLo and TagHi CP0 registers. LSB of VA select the way.
2	I / D	Index Store Tag	Write the tag for the cache block at the specified index from the TagLo and TagHi CP0 registers. LSB of VA select the way.
3	I	Undefined	Undefined
3	D	Create Dirty Exclusive	This operation is used to avoid loading data needlessly from memory when writing new contents into an entire cache block. If the cacheblock does not contain the specified address, and the block is dirty, write it back to the memory. In all cases, set the cache block tag to the specified physical address, set the cache state to Dirty Exclusive.
4	I / D	Hit Invalidate	If the cache block contains the specified address, mark the cache block invalid. In case of multi-hit, lock bits of the specified line become ineffective and all way are invalidated.
5	I	Fill	Fill the primary instruction cache block from memory. LSB of VA select the way.
5	D	Hit WriteBack Invalidate	If the cache block contains the specified address, write back the data if it is dirty, and mark the cache block invalid.
6	I	Undefined	Undefined
6	D	Hit WriteBack	If the cache block contains the specified address, and the W bit is set, write back the data to memory, and clear the W bit.
7	I	Undefined	Undefined
7	D	Fill	Fill the primary data cache block from memory. LSB of VA select the way.

CACHE

Cache
(continued)

CACHE

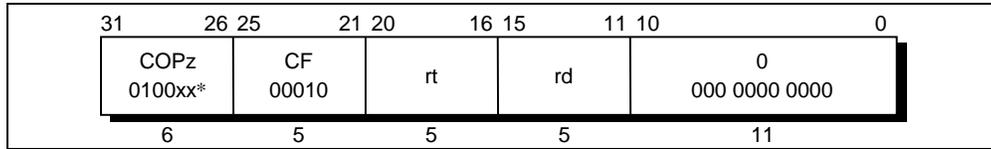
Operation:

32, 64	T:	$vAddr \leftarrow ((offset_{15})^{48} \parallel offset_{15-0}) + GPR[base]$ $(pAddr, uncached) \leftarrow AddressTranslation(vAddr, DATA)$ CacheOp(op, cAddr, pAddr)
--------	----	--

Exceptions:

Coprocessor unusable exception

CFCz Move Control From Coprocessor CFCz



Format:
CFCz rt, rd

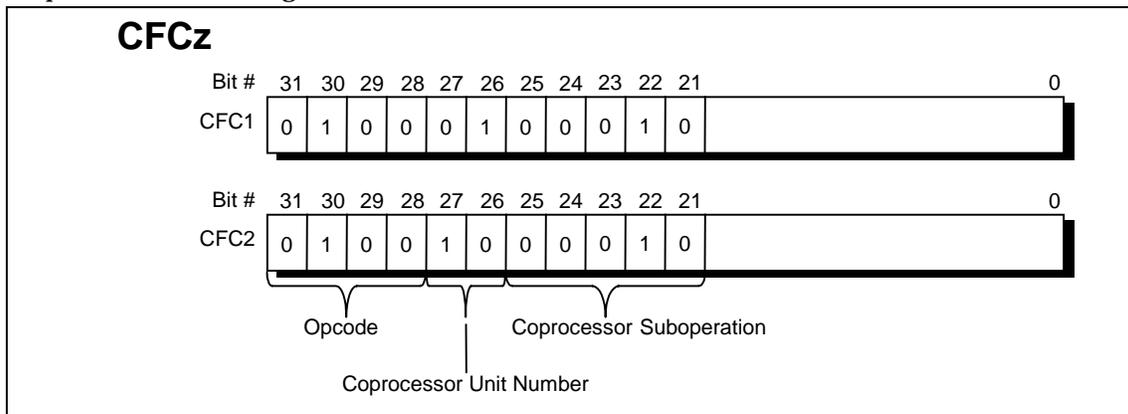
Description:
The contents of coprocessor control register *rd* of coprocessor unit *z* are loaded into general register *rt*.

Operation:

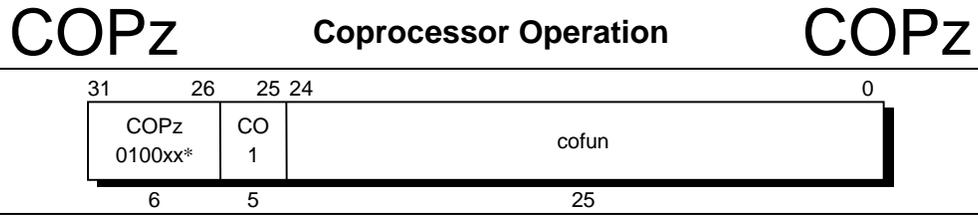
32	T:	data ← CCR[z,rd]
	T + 1:	GPR[rt] ← data
64	T:	data ← (CCR[z,rd] ₃₁) ³² CCR[z, rd]
	T + 1:	GPR[rt] ← data

Exceptions:
Coprocessor unusable exception
Reserved Instruction exception (CFC3)

***Opcode Bit Encoding:**



Note:
CFC1 for FPU (See the Appendix B)
CFC2 for Coprocessor 2 (user define)



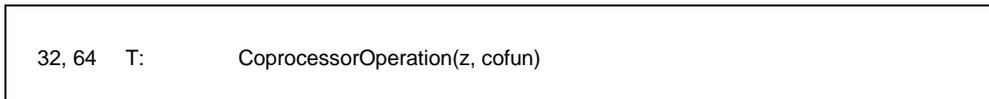
Format:

COPz cofun.

Description:

A coprocessor operation is performed. The operation may specify and reference internal coprocessor registers, and may change the state of the coprocessor condition line, but does not modify state within the processor or the cache / memory system. Details of coprocessor 1 operations are contained in Appendix B.

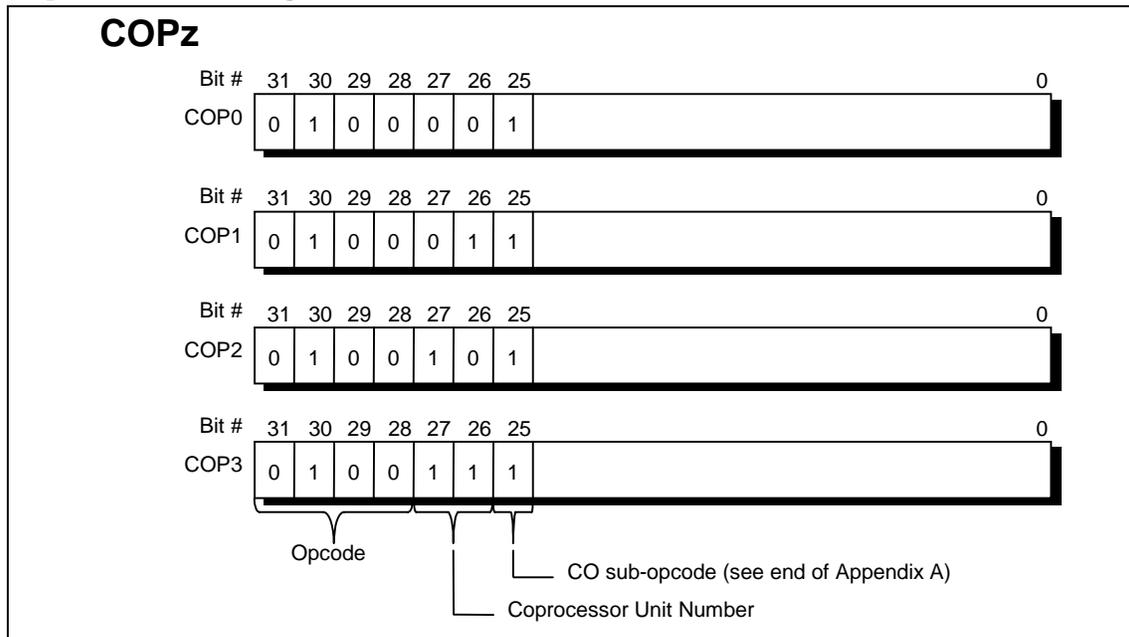
Operation:



Exceptions:

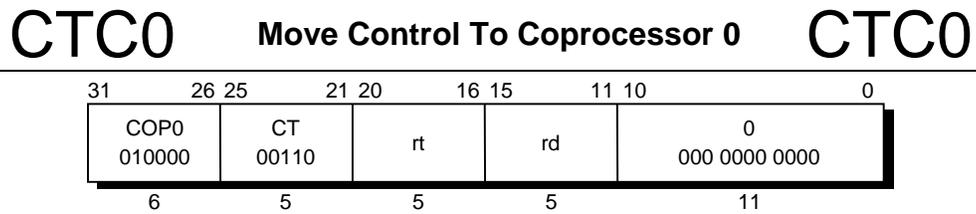
- Coprocessor unusable exception
- Coprocessor interrupt or Floating-Point Exception (TX49 CP1 only)
- Reserved Instruction exception (COP3)

***Opcode Bit Encoding:**



Note:

- COP0 for ICE system
- COP1 for FPU (See the Appendix B)
- COP2 for Coprocessor 2 (user define)



Format :

CTC0 rt, rd

Description :

For ICE system only.

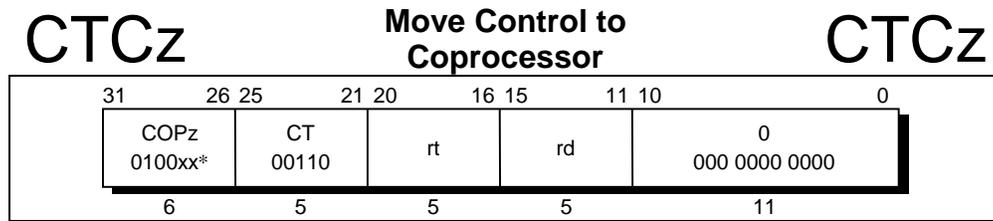
Loads the contents of general-purpose register rt into the Monitor memory.

Operation:

32, 64	T:	data ← GPR[rt]
	T + 1:	CCR[0,rd] ← data

Exceptions :

Coprocessor Unusable exception

**Format:**

CTCz rt, rd

Description:

The contents of general register *rt* are loaded into control register *rd* of coprocessor unit *z*.

Operation:

32, 64	T: data ← GPR[rt] T + 1: CCR[z,rd] ← data
--------	--

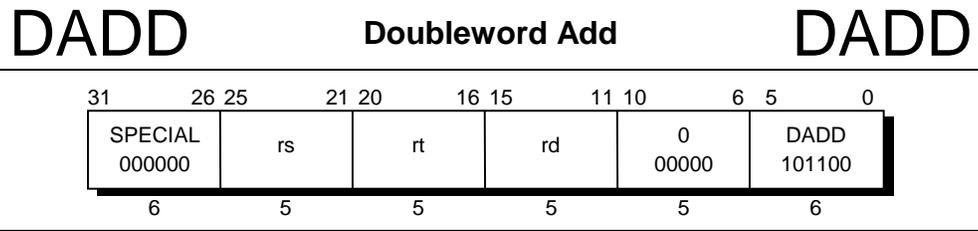
Exceptions:

Coprocessor unusable
Reserved Instruction exception (CTC3)

Note:

CTC1 for FPU (See the Appendix B)
CTC2 for Coprocessor 2 (user define)

*See "CPU Instruction Opcode Bit Encoding" at the end of Appendix A.

**Format:**

DADD rd, rs, rt

Description:

The contents of general register *rs* and the contents of general register *rt* are added to form the result. The result is placed into general register *rd*.

An overflow exception occurs if the carries out of bits 62 and 63 differ (2's-complement overflow). The destination register *rd* is not modified when an integer overflow exception occurs.

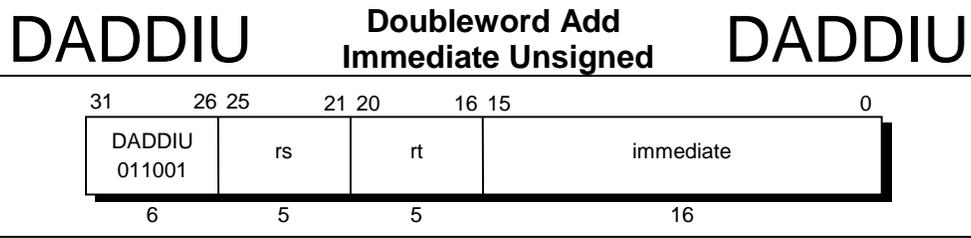
Operation:

64 T: $GPR[rd] \leftarrow GPR[rs] + GPR[rt]$

Exceptions:

Integer overflow exception

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**

DADDIU rt, rs, immediate

Description:

The 16-bit *immediate* is sign-extended and added to the contents of general register *rs* to form the result. The result is placed into general register *rt*. No integer overflow exception occurs under any circumstances.

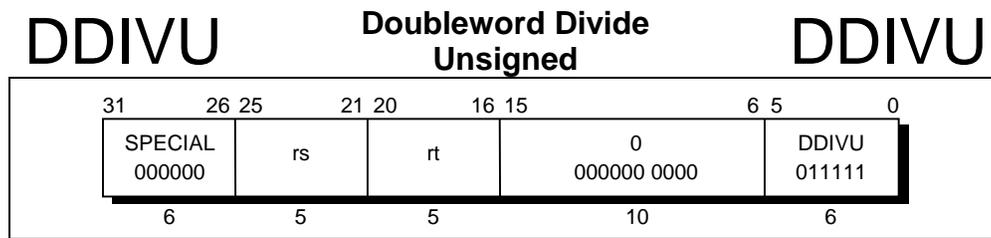
The only difference between this instruction and the DADDI instruction is that DADDIU never causes an overflow exception.

Operation:

64 T: $GPR[rt] \leftarrow GPR[rs] + (immediate_{15})^{48} \parallel immediate_{15-0}$

Exceptions:

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**

DDIVU rs, rt

Description:

The contents of general register *rs* are divided by the contents of general register *rt*, treating both operands as unsigned values. No integer overflow exception occurs under any circumstances, and the result of this operation is undefined when the divisor is zero.

This instruction is typically followed by additional instructions to check for a zero divisor.

When the operation completes, the quotient word of the double result is loaded into special register *LO*, and the remainder word of the double result is loaded into special register *HI*.

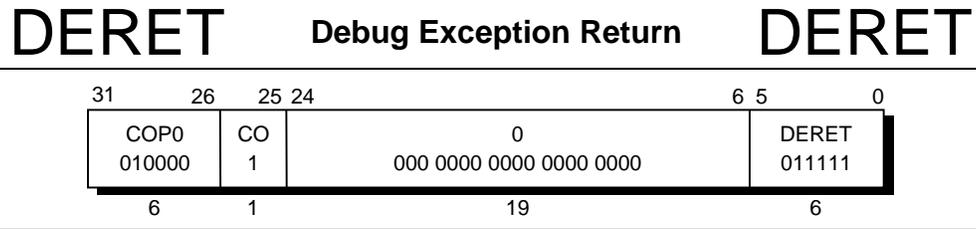
If either of the two preceding instructions is MFHI or MFLO, the results of those instructions are undefined. Correct operation requires separating reads of *HI* or *LO* from writes by two or more instructions.

Operation:

64	T-2:	LO	← undefined
		HI	← undefined
	T-1:	LO	← undefined
		HI	← undefined
	T:	LO	← (0 GPR[rs]) div (0 GPR[rt])
		HI	← (0 GPR[rs]) mod (0 GPR[rt])

Exceptions:

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**

DERET

Description:

Execute a return a self-debug interrupt or exception. This instruction requires a branch delay slot like that of the branch or jump instructions, and executes with a delay of one instruction cycle. The DERET instruction itself cannot be put in the delay slot.

The return address stored in the DEPC register is copied to the PC, and processing returns to the original program.

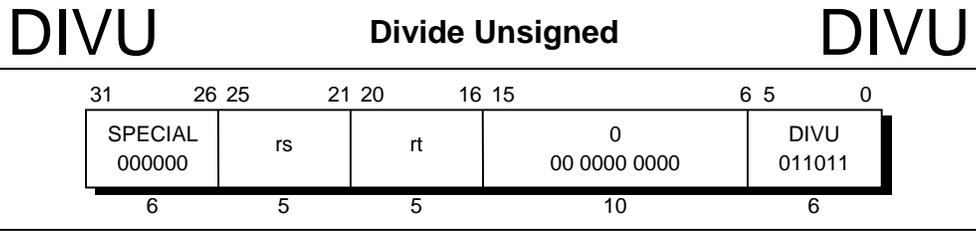
Note: If a MTC0 instruction was used to set the return address in the DEPC register, a minimum of two instructions must be executed before executing DERET.

Operation:

32, 64	T:	temp ← DEPC
	T-1:	PC ← temp
		Debug ₃₀ ← 0

Exceptions:

Coprocessor unusable exception

**Format:**

DIVU rs, rt

Description:

The contents of general register *rs* are divided by the contents of general register *rt*, treating both operands as unsigned values. No integer overflow exception occurs under any circumstances, and the result of this operation is undefined when the divisor is zero. In 64-bit mode, the operands must be valid sign-extended, 32-bit values. In 64-bit mode, the operands must be valid sign-extended, 32-bit values.

This instruction is typically followed by additional instructions to check for a zero divisor.

When the operation completes, the quotient word of the double result is loaded into special register *LO*, and the remainder word of the double result is loaded into special register *HI*.

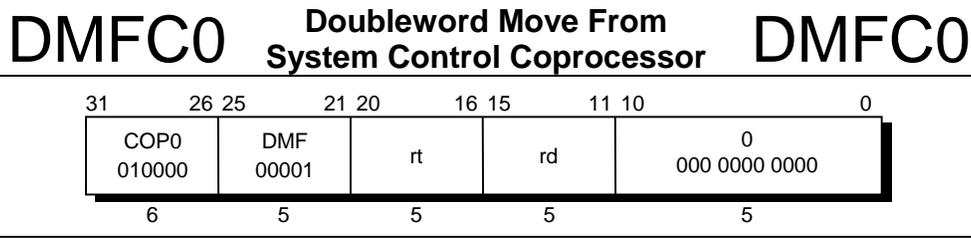
If either of the two preceding instructions is MFHI or MFLO, the results of those instructions are undefined. Correct operation requires separating reads of *HI* or *LO* from writes by two or more instructions.

Operation:

32	T-2:	LO	← undefined
		HI	← undefined
	T-1:	LO	← undefined
		HI	← undefined
T:	LO	← (0 GPR[rs]) div (0 GPR[rt])	
	HI	← (0 GPR[rs]) mod (0 GPR[rt])	
64	T-2:	LO	← undefined
		HI	← undefined
	T-1:	LO	← undefined
		HI	← undefined
T:	q	← (0 GPR[rs] ₃₁₋₀) div (0 GPR[rt] ₃₁₋₀)	
	r	← (0 GPR[rs] ₃₁₋₀) mod (0 GPR[rt] ₃₁₋₀)	
	LO	← (q ₃₁) ³² q ₃₁₋₀	
	HI	← (r ₃₁) ³² r ₃₁₋₀	

Exceptions:

None

**Format:**

DMFC0 rt, rd

Description:

The contents of coprocessor register rd of the CP0 are loaded into general register rt.

This operation is defined in kernel mode regardless of the setting of the Status. KX bit. Execution of this instruction with in supervisor mode with Status. SX = 0 or in user mode with UX = 0, causes a reserved instruction exception. All 64-bits of the general register destination are written from the coprocessor register source. The operation of DMFC0 on a 32-bit coprocessor 0 register is undefined.

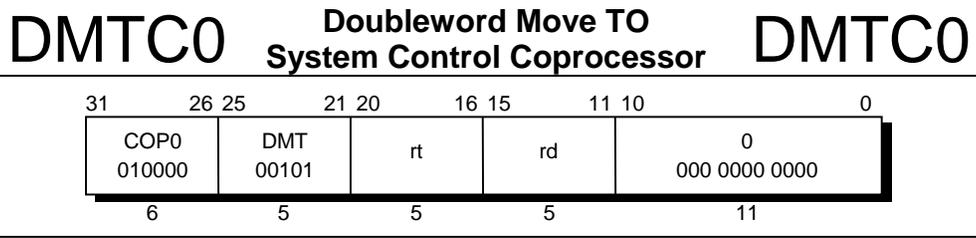
Operation:

64	T:	data ← CPR[0,rd]
	T + 1:	GPR[rt] ← data

Exceptions:

Coprocessor unusable exception

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**DMTC0 *rt*, *rd***Description:**

The contents of general register *rt* are loaded into coprocessor register *rd* of the CP0.

This operation is defined for the R4000 operating in 64-bit mode or in 32-bit kernel mode. Execution of this instruction in 32-bit user or supervisor mode causes a reserved instruction exception. All 64-bits of the coprocessor 0 register are written from the general register source. The operation of DMTC0 on a 32-bit coprocessor 0 register is undefined.

Because the state of the virtual address translation system may be altered by this instruction, the operation of load, store instructions and TLB operations immediately prior to and after this instruction are undefined.

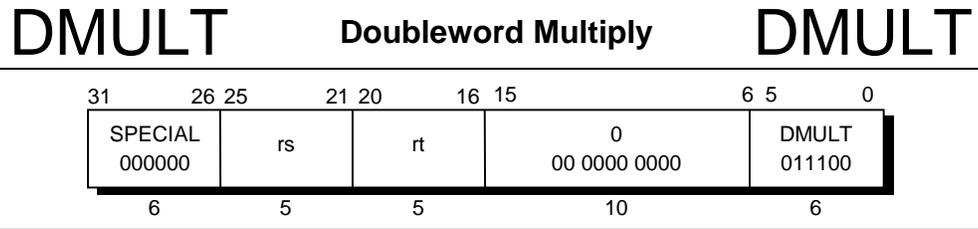
Operation:

64	T:	$data \leftarrow GPR[rt]$
	T + 1:	$CPR[0,rd] \leftarrow data$

Exceptions:

Coprocessor unusable exception

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**

DMULT rs, rt

Description:

The contents of general registers *rs* and *rt* are multiplied, treating both operands as 2's-complement values. No integer overflow exception occurs under any circumstances.

When the operation completes, the low-order word of the double result is loaded into special register *LO*, and the high-order word of the double result is loaded into special register *HI*.

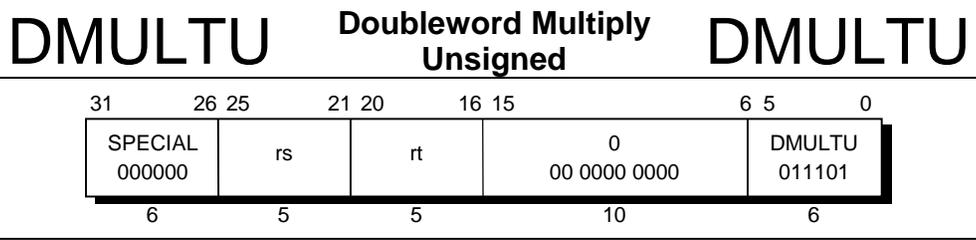
If either of the two preceding instructions is MFHI or MFLO, the results of these instructions are undefined. Correct operation requires separating reads of *HI* or *LO* from writes by a minimum of two other instructions.

Operation:

64	T-2:	LO	← undefined
		HI	← undefined
	T-1:	LO	← undefined
		HI	← undefined
	T:	t	← GPR[rs]* GPR[rt]
		LO	← t ₆₃₋₀
		HI	← t ₁₂₇₋₆₄

Exceptions:

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**

DMULTU rs, rt

Description:

The contents of general register *rs* and the contents of general register *rt* are multiplied, treating both operands as unsigned values. No over-flow exception occurs under any circumstances.

When the operation completes, the low-order word of the double re-sult is loaded into special register *LO*, and the high-order word of the double result is loaded into special register *HI*.

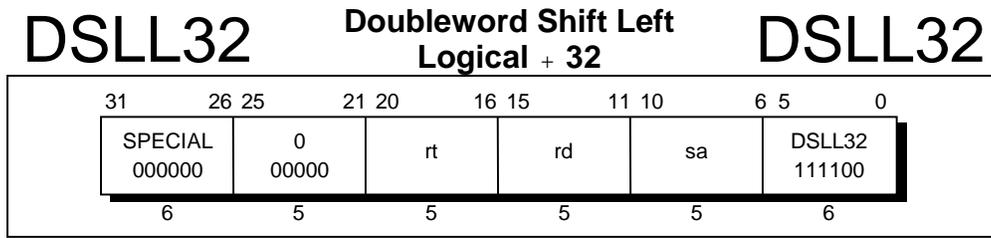
If either of the two preceding instructions is MFHI or MFLO, the re-suits of these instructions are undefined. Correct operation requires separating reads of *HI* or *LO* from writes by a minimum of two instructions.

Operation:

64	T-2:	LO	← undefined
		HI	← undefined
	T-1:	LO	← undefined
		HI	← undefined
	T:	t	← (0 GPR[rs]) * (0 GPR[rt])
		LO	← t ₆₃₋₀
		HI	← t ₁₂₇₋₆₄

Exceptions:

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**

DSLL32 rd, rt, sa

Description:

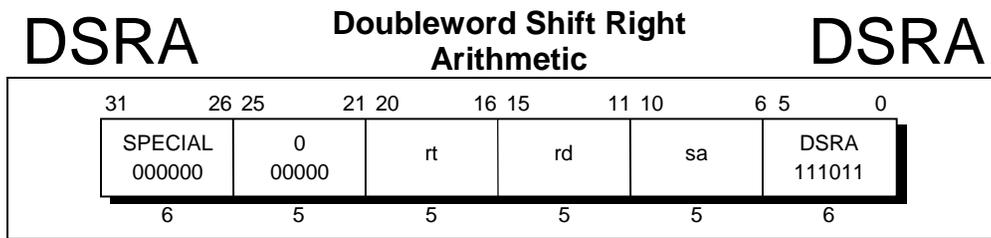
The contents of general register *rt* are shifted left by $32 + sa$ bits, inserting zeros into the low-order bits. The result is placed in register *rd*.

Operation:

<p>64 T: $s \leftarrow 1 \parallel sa$ $GPR[rd] \leftarrow GPR[rt]_{(63-s)-0} \parallel 0^s$</p>

Exceptions:

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**

DSRA rd, rt, sa

Description:

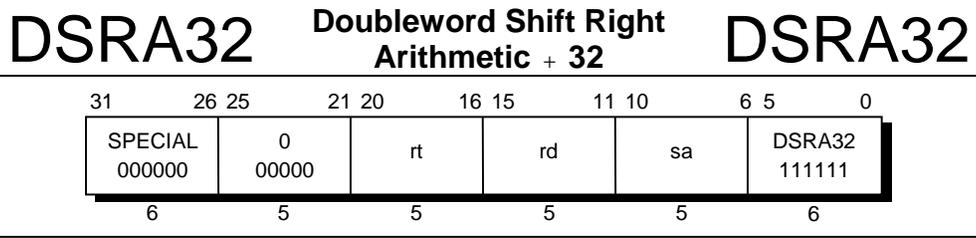
The contents of general register *rt* are shifted right by *sa* bits, sign-ex-tending the high-order bits. The result is placed in register *rd*.

Operation:

64 T: $s \leftarrow 0 \parallel sa$ $GPR[rd] \leftarrow (GPR[rt]_{63}^s) \parallel GPR[rt]_{63-s}$
--

Exceptions:

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**

DSRA32 rd, rt,sa

Description:

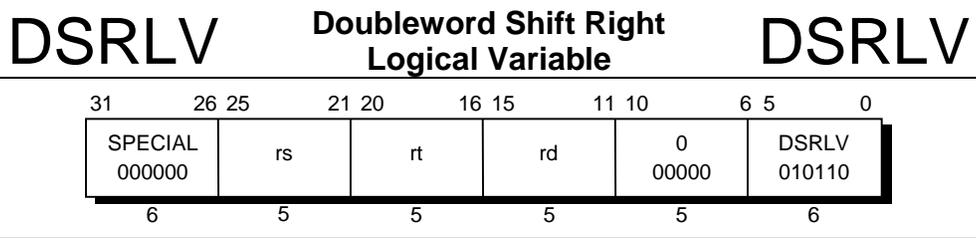
The contents of general register *rt* are shifted right by $32 + sa$ bits, sign-extending the high-order bits. The result is placed in register *rd*.

Operation:

64	T:	$s \leftarrow 1 \parallel sa$
		$GPR[rd] \leftarrow (GPR[rt]_{63}^s) \parallel GPR[rt]_{63-s}$

Exceptions:

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**

DSRLV rd, rt, rs

Description:

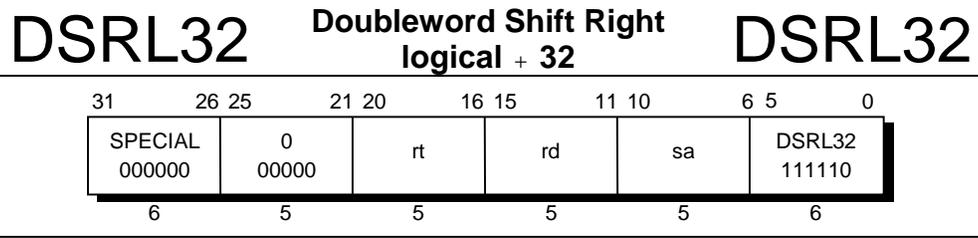
The contents of general register *rt* are shifted right by the number of bits specified by the low-order six bits of general register *rs*, inserting zeros unto the high-order bits. The result is placed in register *rd*.

Operation:

64 T: $s \leftarrow \text{GPR}[rs]_{5-0}$ $\text{GPR}[rd] \leftarrow 0^s \parallel \text{GPR}[rt]_{63-s}$

Exceptions:

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**

DSRL32 rd, rt, sa

Description:

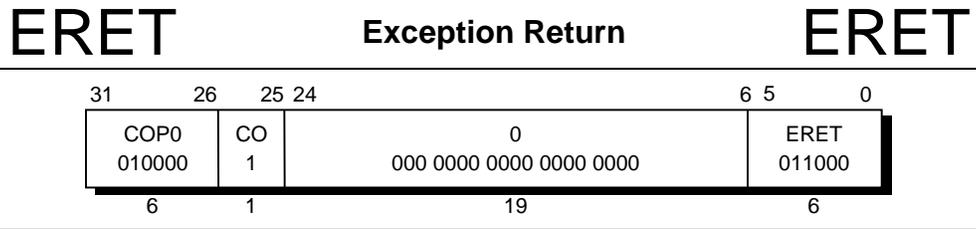
The contents of general register *rt* are shifted right by $32 + sa$ bits, inserting zeros into the high-order bits. The result is placed in register *rd*.

Operation:

64	T:	$s \leftarrow 1 \parallel sa$ $GPR[rd] \leftarrow 0^s \parallel GPR[rt]_{63-s}$
----	----	--

Exceptions:

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**

ERET

Description:

ERET is the TX49 instruction for returning from an interrupt, exception, or error trap. Unlike a branch or jump instruction, ERET does not execute the next instruction.

ERET must not itself be placed in a branch delay slot.

If the processor is servicing an error trap ($SR_2 = 1$), then load the PC from the *ErrorEPC* and clear the *ERL* bit of the *Status* register (SR_2). Otherwise ($SR_2 = 0$), load the PC from the *EPC*, and clear the *EXL* bit of the *Status* register (SR_1).

An ERET executed between a LL and SC also causes the SC to fail.

In case of this instruction is placed in the boundary of memory, it is necessary to keep the branch delay slot into same memory area.

Operation:

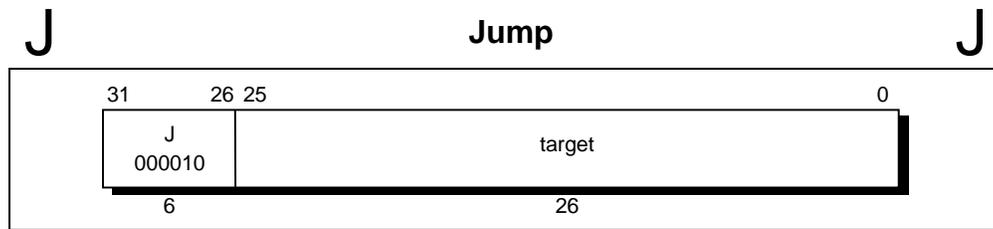
```

32, 64 T:   if SR2 = 1 then
             PC ← ErrorEPC
             SR ← SR31-3 || 0 || SR1-0
             else
             PC ← EPC
             SR ← SR31-2 || 0 || SR0
             endif
             LLbit ← 0

```

Exceptions:

Coprocessor unusable exception

**Format:**

J target

Description:

The 26-bit target address is shifted left two bits and combined with the high-order bits of the address of the delay slot. The program unconditionally jumps to this calculated address with a delay of one instruction.

Operation:

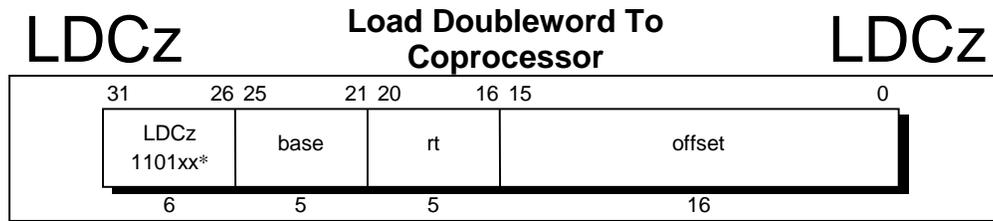
```

32  T:    temp ← target
    T + 1: PC ← PC31-28 || temp || 02
64  T:    temp ← target
    T + 1: PC ← PC63-28 || temp || 02

```

Exceptions:

None

**Format:**

LDCz rt, offset (base)

Description :

The 16-bit *offset* is sign-extended and added to the contents of general register *base* to form a virtual address. The processor reads a double-word from the addressed memory location and makes the data available to coprocessor unit *z*. The manner in which each coprocessor uses the data is defined by the individual coprocessor specifications.

If any of the three least-significant bits of the effective address are non-zero, an address error exception takes place.

This instruction is not valid for use with CP0.

This instruction is undefined when the least-significant bit of the *rt-field* is non-zero.

*See the table "Opcode Bit Encoding" on next page, or "CPU Instruction Opcode Bit Encoding" at the end of Appendix A.

LDCz Load Doubleword To Coprocessor (continued) LDCz

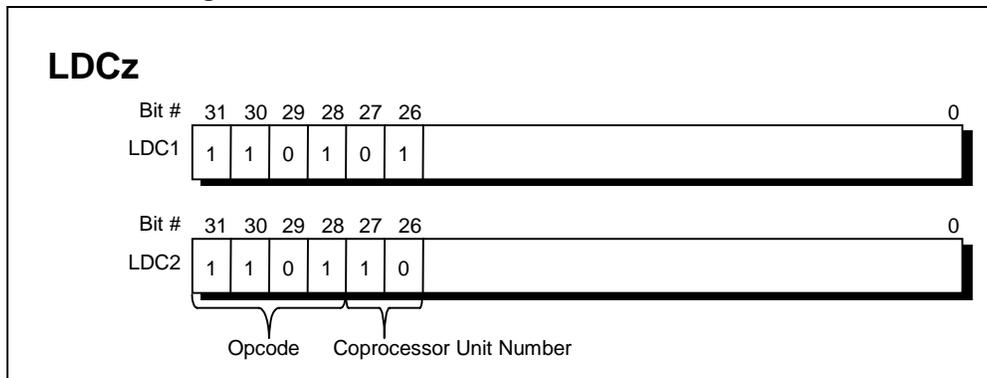
Operation:

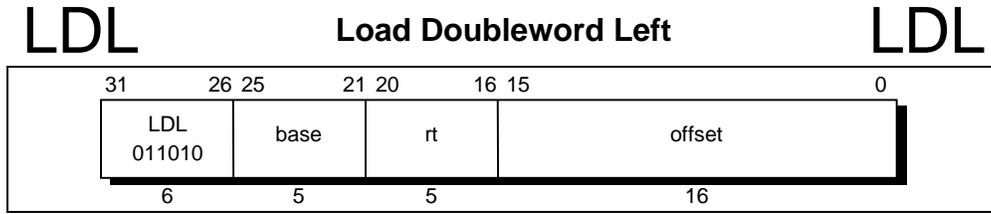
32	T:	$vAddr \leftarrow ((offset_{15})^{16} \parallel offset_{15-0}) \bullet + \bullet GPR[base]$ $(pAddr, uncached) \leftarrow AddressTranslation(vAddr, DATA)$ $mem \leftarrow LoadMemory(uncached, DOUBLEWORD, pAddr, vAddr, DATA)$ COPzLD(rt, mem)
64	T:	$vAddr \leftarrow ((offset_{15})^{48} \parallel offset_{15-0}) \bullet + \bullet GPR[base]$ $(pAddr, uncached) \leftarrow AddressTranslation(vAddr, DATA)$ $mem \leftarrow LoadMemory(uncached, DOUBLEWORD, pAddr, vAddr, DATA)$ COPzLD(rt, mem)

Exceptions:

- TLB refill exception
- TLB invalid exception
- Bus error exception
- Address error exception
- Coprocessor unusable exception

Opcode Bit Encoding:





Format:

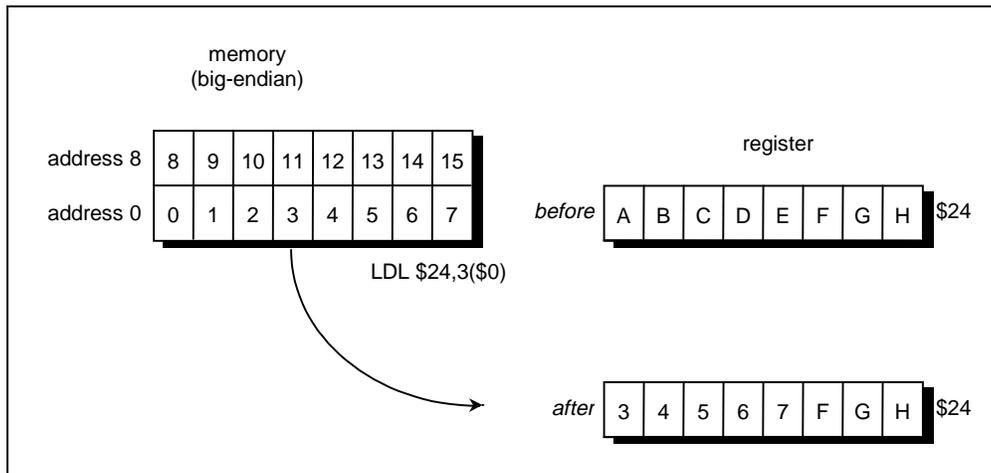
LDL rt, offset (base)

Description:

This instruction can be used in combination with the LDR instruction to load a register with eight consecutive bytes from memory, when the bytes cross a boundary between two doublewords. LDL loads the left portion of the register from the appropriate part of the high-order doubleword; LDR loads the right portion of the register from the appropriate part of the low-order doubleword.

The LDL instruction adds its sign-extended 16-bit *offset* to the contents of general register *base* to form a virtual address which can specify an arbitrary byte. It reads bytes only from the doubleword in memory which contains the specified starting byte. From one to eight bytes will be loaded, depending on the starting byte specified.

Conceptually, it starts at the specified byte in memory and loads that byte into the high-order (left-most) byte of the register; then it proceeds toward the low-order byte of the doubleword in memory and the low-order byte of the register, loading bytes from memory into the register until it reaches the low-order byte of the doubleword in memory. The least-significant (right-most) byte(s) of the register will not be changed.



LDL

Load Doubleword Left
(continued)

LDL

The contents of general register *rt* are internally bypassed within the processor so that no NOP is needed between an immediately preceding load instruction which specifies register *rt* and a following LDL (or LDR) instruction which also specifies register *rt*.

No address exceptions due to alignment are possible.

Operation:

```

64  T:    vAddr ← ((offset15)48 || offset15-0) • +• GPR[base]
        (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
        pAddr ← pAddrPSIZE-1-3 || (pAddr2-0 xor ReverseEncian3)
        if BigEndianMem• = 0 then
        pAddr ← pAddrPSIZE-1-3 || 03
        endif
        byte ← vAddr2-0 xor BigEndianCPU3
        mem ← LoadMemory (uncached, byte, pAddr, vAddr, DATA)
        GPR[rt] ← mem7• + 8*byte-0 || GPR[rt]55• - 8*byte-0

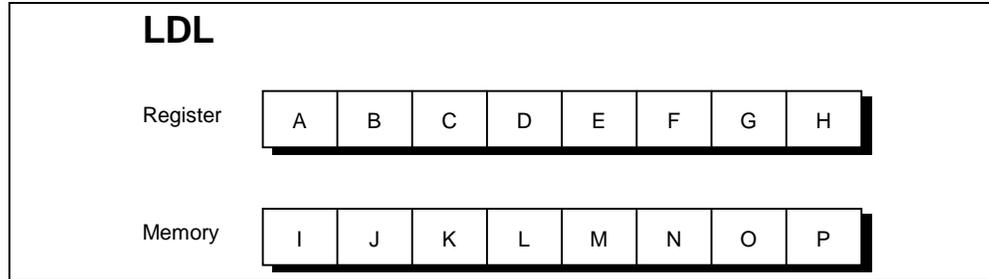
```

LDL

Load Doubleword Left
(continued)

LDL

Given a doubleword in a register and a doubleword in memory, the operation of LDL us as follows:



vAddr ₂₋₀	BigEndianCPU = 0					BigEndianCPU = 1				
	Destination	type	offset		Destination	type	offset			
			LEM	BEM			LEM	BEM		
0	P B C D E F G H	0	0	7	I J K L M N O P	7	0	0		
1	O P C D E F G H	2	0	5	K L M N O P G H	5	0	2		
3	M N O P E F G H	3	0	4	L M N O P F G H	4	0	3		
4	L M N O P F G H	4	0	3	M N O P E F G H	3	0	4		
5	K L M N O P G H	5	0	2	N O P D E F G H	2	0	5		
6	J K L M N O P H	6	0	1	O P C D E F G H	1	0	6		
7	I J K L M N O P	7	0	0	P B C D E F G H	0	0	7		

LEM BigEndianMem = 0

BEM BigEndianMem = 1

Type AccessType sent to memory

Offset Addr₂₋₀ sent to memory

Exceptions:

TLB refill exception

TLB invalid exception

Bus error exception

Address error exception

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

LDR

Load Doubleword Right
(continued)

LDR

The contents of general register *rt* are internally bypassed within the processor so that no NOP is needed between an immediately preceding load instruction which specifies register *rt* and a following LDR (or LDL) instruction which also specifies register *rt*.

No address exceptions due to alignment are possible.

Operation:

```

64  T:    vAddr ← ((offset15)48 || offset15-0) • +• GPR[base]
        (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
        pAddr ← pAddrPSIZE-1-3 || (pAddr2-0 xor ReverseEncian3)
        if BigEndianMem • ⇒ 1 then
        pAddr ← pAddr31-3 || 03
        endif
        byte ← vAddr2-0 xor BigEndianCPU3
        mem ← LoadMemory (uncached, byte, pAddr, vAddr, DATA)
        GPR[rt] ← GPR[rt]63-64 • - 8*byte || mem63-8*byte

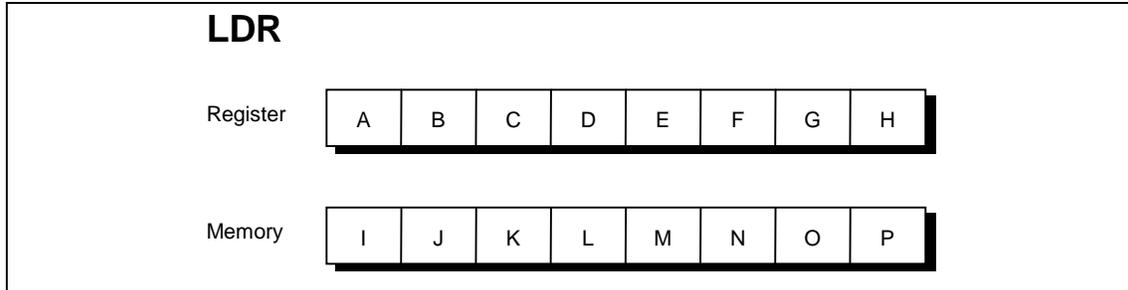
```

LDR

Load Doubleword Right (continued)

LDR

Given a doubleword in a register and a doubleword in memory, the operation of LDR is as follows:



vAddr ₂₋₀	BigEndianCPU = 0				BigEndianCPU = 1			
	Destination	type	offset		Destination	type	offset	
			LEM	BEM			LEM	BEM
0	I J K L M N O P	7	0	0	A B C D E F G I	0	7	0
1	A I J K L M N O	6	1	0	A B C D E F I J	1	6	0
2	A B I J K L M N	5	2	0	A B C D E I J K	2	5	0
3	A B I I J K L M	4	3	0	A B C D E I J K	3	4	0
4	A B C D I J K L	3	4	0	A B C I J K L M	4	3	0
5	A L C D E I J K	2	5	0	A B I J K L M K	5	2	0
6	A B C D E F I J	1	6	0	A I J K L M N O	6	1	0
7	A B C D E F G I	0	7	0	I J K L M N O P	7	0	0

LEM BigEndianMem = 0

BEM BigEndianMem = 1

Type AccessType sent to memory

Offset Addr₂₋₀ sent to memory

Exceptions:

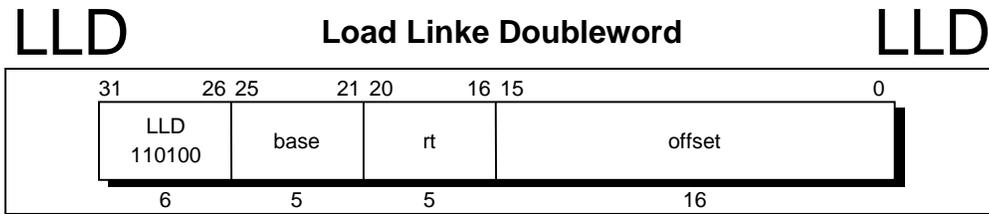
TLB refill exception

TLB invalid exception

Bus error exception

Address error exception

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**

LLD *rt*, *offset* (*base*)

Description:

The 16-bit *offset* is sign-extended and added to the contents of general register *base* to form a virtual address. The contents of the doubleword at the memory location specified by the effective address are loaded into general register *rt*.

The processor begins checking the accessed doubleword for modification by other processors and devices.

Load Linked Doubleword and Store Conditional Doubleword can be used to atomically update memory locations:

<pre> L1: LLD T1, (T0) ADD T2, T1, 1 SCD T2, (T0) BEQ T2, 0, L1 NOP </pre>
--

This atomically increments the word addressed by T0. Changing the ADD to an OR changes this to an atomic bit set.

LLD**Load LinkedDoubleword
(continued)****LLD**

The operation of LLD is undefined if the addressed location is uncached and, for synchronization between multiple processors, the operation of LLD is undefined if the addressed location is noncoherent.

A cache miss that occurs between LLD and SCD may cause SCD to fail, so no load or store instruction should occur between LLD and SCD. Exceptions also cause SCD to fail, so persistent exceptions must be avoided.

This instruction is available in User mode, and it is not necessary for CP0 to be enabled.

If any of the three least-significant bits of the effective address are non-zero, an address error exception takes place.

Operation:

<pre> 64 T: vAddr ← ((offset₁₅)⁴⁸ offset₁₅₋₀) • +• GPR[base] (pAddr, uncached) ← AddressTranslation (vAddr, DATA) mem ← LoadMemory (uncached, DOUBLE WORD, pAddr, vAddr, DATA) GPR[rt] ← mem LLbit ← 1 </pre>

Exceptions:

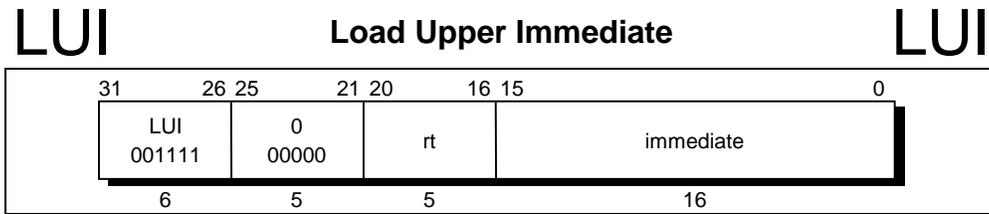
TLB refill exception

TLB invalid exception

Bus error exception

Address error exception

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**LUI *rt*, *immediate***Description:**

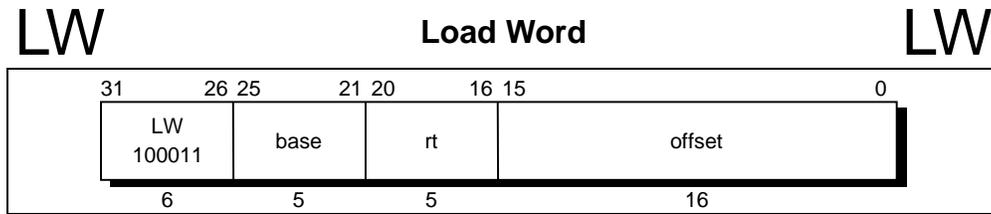
The 16-bit *immediate* is shifted left 16 bits and concatenated to 16 bits of zeros. The result is placed into general register *rt*. In 64-bit mode, the loaded word is sign-extended.

Operation:

32	T:	$GPR[rt] \leftarrow \text{immediate} \parallel 0^{16}$
64	T:	$GPR[rs] \leftarrow (\text{immediate}_{15})^{32} \parallel \text{immediate} \parallel 0^{16}$

Exceptions:

None

**Format:**

LW rt, offset (base)

Description:

The 16-bit *offset* is sign-extended and added to the contents of general register *base* to form a virtual address. The contents of the word at the memory location specified by the effective address are loaded into general register *rt*. In 64-bit mode, the loaded word is sign-extended.

If either of the two least-significant bits of the effective address is non-zero, an address error exception occurs.

Operation:

32	T:	$vAddr \leftarrow ((offset_{15})^{16} \parallel offset_{15-0}) \bullet \bullet GPR[base]$ $(pAddr, uncached) \leftarrow AddressTranslation(vAddr, DATA)$ $pAddr \leftarrow pAddr_{PSIZE-1-3} \parallel (pAddr_{2-0} \text{ xor } (ReverseEndian \parallel 0^2))$ $mem \leftarrow LoadMemory(uncached, WORD, pAddr, vAddr, DATA)$ $byte \leftarrow vAddr_{2-0} \text{ xor } (BigEndianCPU \parallel 0^2)$ $GPR[rt] \leftarrow mem_{31 \bullet + 8 \bullet \text{byte} - 8 \bullet \text{byte}}$
64	T:	$vAddr \leftarrow ((offset_{15})^{48} \parallel offset_{15-0}) \bullet \bullet GPR[base]$ $(pAddr, uncached) \leftarrow AddressTranslation(vAddr, DATA)$ $pAddr \leftarrow pAddr_{PSIZE-1-3} \parallel (pAddr_{2-0} \text{ xor } (ReverseEndian \parallel 0^2))$ $mem \leftarrow LoadMemory(uncached, WORD, pAddr, vAddr, DATA)$ $byte \leftarrow vAddr_{2-0} \text{ xor } (BigEndianCPU \parallel 0^2)$ $GPR[rt] \leftarrow (mem_{31 \bullet + 8 \bullet \text{byte}})^{32} \parallel mem_{31 \bullet + 8 \bullet \text{byte} - 8 \bullet \text{byte}}$

Exceptions:

- TLB refill exception
- TLB invalid exception
- Bus error exception
- Address error exception

LWCz

Load Word To Coprocessor
(continued)

LWCz

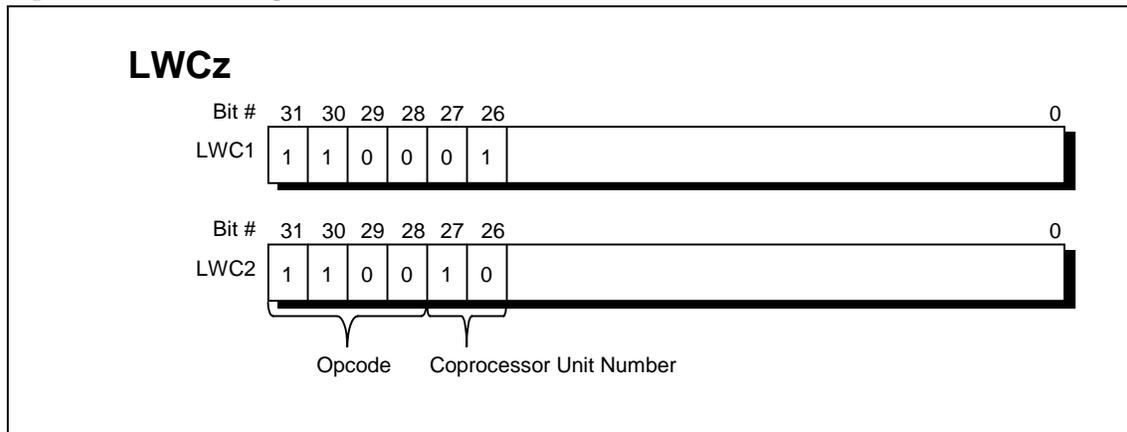
Operation:

32	T:	$vAddr \leftarrow ((offset_{15})^{16} \parallel offset_{15-0}) \bullet + \bullet GPR[base]$ $(pAddr, uncached) \leftarrow AddressTranslation(vAddr, DATA)$ $pAddr \leftarrow pAddr_{PSIZE-1-3} \parallel (pAddr_{2-0} \text{ xor } (ReverseEndian \parallel 0^2))$ $mem \leftarrow LoadMemory(uncached, WORD, pAddr, vAddr, DATA)$ $byte \leftarrow vAddr_{2-0} \text{ xor } (BigEndianCPU \parallel 0^2)$ COPzLW(byte, rt, mem)
64	T:	$vAddr \leftarrow ((offset_{15})^{48} \parallel offset_{15-0}) \bullet + \bullet GPR[base]$ $(pAddr, uncached) \leftarrow AddressTranslation(vAddr, DATA)$ $pAddr \leftarrow pAddr_{PSIZE-1-3} \parallel (pAddr_{2-0} \text{ xor } (ReverseEndian \parallel 0^2))$ $mem \leftarrow LoadMemory(uncached, WORD, pAddr, vAddr, DATA)$ $byte \leftarrow vAddr_{2-0} \text{ xor } (BigEndianCPU \parallel 0^2)$ COPzLW(byte, rt, mem)

Exceptions:

- TLB refill exception
- TLB invalid exception
- Bus error exception
- Address error exception
- Coprocessor unusable exception

Opcode Bit Encoding:



LWL

Load Word Left
(continued)

LWL

The contents of general register *rt* are internally bypassed within the processor so that no NOP is needed between an immediately preceding load instruction which specifies register *rt* and a following LWL (or LWR) instruction which also specifies register *rt*.

No address exceptions due to alignment are possible.

Operation:

```

32  T:    vAddr ← ((offset15)16 || offset15-0) • +• GPR[base]
        (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
        pAddr ← pAddrPSIZE-1-3 || (pAddr2-0 xor (ReverseEncian3))
        if BigEndianMem • =• 0 then
        pAddr ← pAddrPSIZE-1-2 || 02
        endif
        byte ← vAddr1-0 xor BigEndianCPU2
        word ← vAddr2 xor BigEndianCPU
        mem ← LoadMemory (uncached, 0 || byte, pAddr, vAddr, DATA)
        temp ← mem32*word • + 8*byte • + 7-32*word || GPR[rt]23 • - 8*byte-0
        GPR[rt] ← temp

64  T:    vAddr ← ((offset15)48 || offset15-0) • +• GPR[base]
        (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
        pAddr ← pAddrPSIZE-1-3 || (pAddr2-0 xor (ReverseEncian3))
        if BigEndianMem • =• 0 then
        pAddr ← pAddrPSIZE-1-2 || 02
        endif
        byte ← vAddr1-0 xor BigEndianCPU2
        word ← vAddr2 xor BigEndianCPU
        mem ← LoadMemory (uncached, 0 || byte, pAddr, vAddr, DATA)
        temp ← mem32*word • + 8*byte • + 7-32*word || GPR[rt]23 • - 8*byte-0
        GPR[rt] ← (temp31)32 || temp

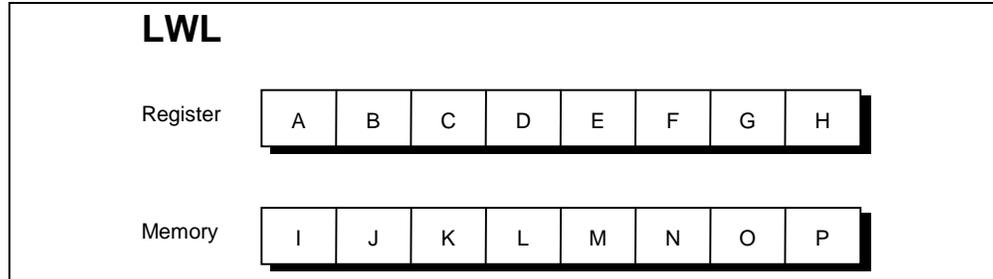
```

LWL

Load Word Left (continued)

LWL

Given a doubleword in a register and a doubleword in memory, the operation of LWL is as follows:



vAddr ₂₋₀	BigEndianCPU = 0					BigEndianCPU = 1				
	Destination	type	offset		Destination	type	offset			
			LEM	BEM			LEM	BEM		
0	S S S S P F G H	0	0	7	S S S S I J K L	3	4	0		
1	S S S S O P G H	1	0	6	S S S S J K L H	2	4	1		
2	S S S S N O P H	2	0	5	S S S S K L G H	1	4	2		
3	S S S S M N O P	3	0	4	S S S S L F G H	0	4	3		
4	S S S S L F G H	0	4	3	S S S S M N O P	3	0	4		
5	S S S S K L G H	1	4	2	S S S S N O P H	2	0	5		
6	S S S S J K L H	2	4	1	S S S S O P G H	1	0	6		
7	S S S S I J K L	3	4	0	S S S S P F G H	0	0	7		

LEM BigEndianMem = 0

BEM BigEndianMem = 1

Type AccessType (see Figure 2-2) sent to memory

Offset pAddr₂₋₀ sent to memory

S sign-extend of destination₃₁

Exception:

- TLB refill exception
- TLB invalid exception
- Bus error exception
- Address error exception

LWR

Load Word Right
(continued)

LWR

The contents of general register *rt* are internally bypassed within the processor so that no NOP is needed between an immediately preceding load instruction which specifies register *rt* and a following LWR (or LWL) instruction which also specifies register *rt*.

No address exceptions due to alignment are possible.

Operation:

```

32  T:    vAddr ← ((offset15)16 || offset15-0) • +• GPR[base]
        (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
        pAddr ← pAddrPSIZE-1-3 || (pAddr2-0 xor ReverseEndian3)
        if BigEndianMem • =• 1 then
        pAddr ← pAddrPSIZE-31-3 || 03
        endif
        byte ← vAddr1-0 xor BigEndianCPU2
        word ← vAddr2 xor BigEndianCPU
        mem ← LoadMemory (uncached, 0 || byte, pAddr, vAddr, DATA)
        temp ← GPR[rt]31-32 • - 8*byte || mem31 • + 32*word-32*word • + 8*byte
        GPR[rt] ← temp

64  T:    vAddr ← ((offset15)48 || offset15-0) • +• GPR[base]
        (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
        pAddr ← pAddrPSIZE-1-3 || (pAddr2-0 xor ReverseEndian3)
        if BigEndianMem • =• 1 then
        pAddr ← pAddrPSIZE-31-3 || 03
        endif
        byte ← vAddr1-0 xor BigEndianCPU2
        word ← vAddr2 xor BigEndianCPU
        mem ← LoadMemory (uncached, 0 || byte, pAddr, vAddr, DATA)
        temp ← GPR[rt]31-32 • - 8*byte || mem31 • + 32*word-32*word • + 8*byte
        GPR[rt] ← (temp31)32 || temp

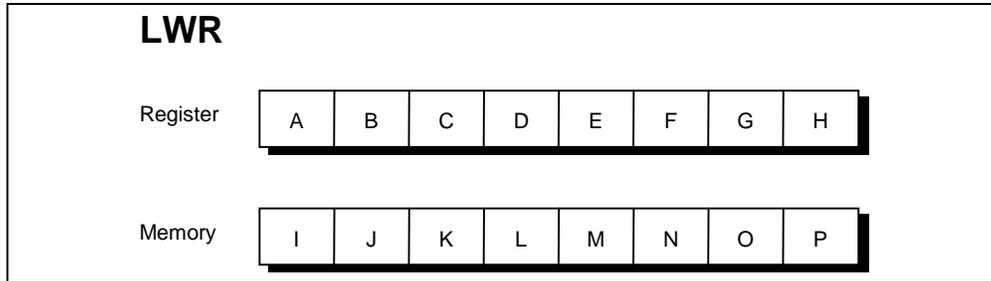
```

LWR

Load Word Right (continued)

LWR

Given a word in a register and a word in memory, the operation of LWR is as follows:



vAddr ₂₋₀	BigEndianCPU = 0					BigEndianCPU = 1				
	destination		type	offset		Destination		type	offset	
	S	S		LEM	BEM	S	S		LEM	BEM
0	S	S	0	0	4	X	X	0	7	0
1	X	X	1	1	4	X	X	1	6	0
2	X	X	2	2	4	X	X	2	5	0
3	X	X	3	3	4	S	S	3	4	0
4	S	S	0	4	0	X	X	4	3	4
5	X	X	1	5	0	X	X	5	2	4
6	X	X	2	6	0	X	X	6	1	4
7	X	X	3	7	0	S	S	7	0	4

LEM BigEndianMem = 0

BEM BigEndianMem = 1

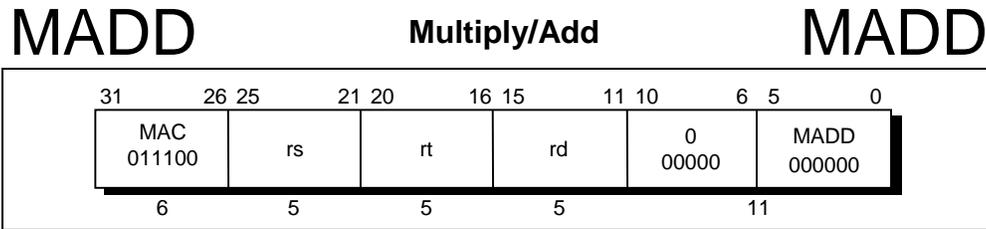
Type AccessType (see Figure 2-2) sent to memory

Offset pAddr₂₋₀ sent to memory

S sign-extend of destination₃₁

Exceptions:

- TLB refill exception
- TLB invalid exception
- Bus error exception
- Address error exception

**Format:**

- MADD rs, rt
- MADD rd, rs, rt

Description:

Multiplies the contents of general registers *rs* and *rt*, treating both values as two's complement, and puts the double-word result in special registers HI and LO. An overflow exception is never raised. The low-order word of the multiplication result is put in general register *rd* and in special register LO, whereas the high-order word of the result is put in special register HI.

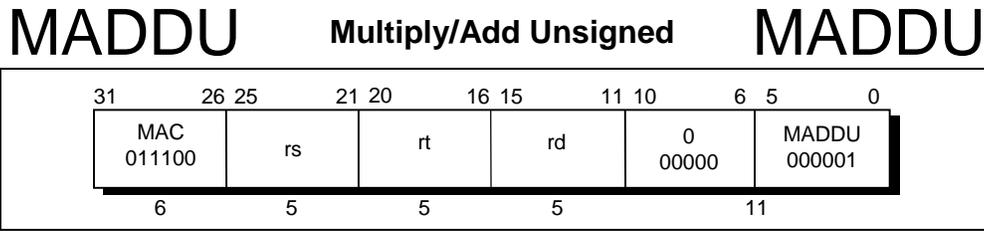
If *rd* is omitted in assembly language, 0 is used as the default value. To guarantee correct operation even if an interrupt occurs, neither of the two instructions following MADD should be DIV or DIVU instructions which modify the HI and LO register contents.

Operation:

<p>32, 64 T: $t \leftarrow (HI \parallel LO) \bullet + \bullet GPR[rs] * GPR[rt]$ $LO \leftarrow t_{31-0}$ $HI \leftarrow t_{63-32}$ $GPR[rd] \leftarrow t_{31-0}$</p>

Exception:

None

**Format:**

MADDU rs, rt

MADDU rd, rs, rt

Description:

Multiplies the contents of general registers rs and rt, treating both values as unsigned, and puts the double-word result in special registers HI and LO. An overflow exception is never raised. The low-order word of the multiplication result is put in general register rd and in special register LO, whereas the high-order word of the result is put in special register HI.

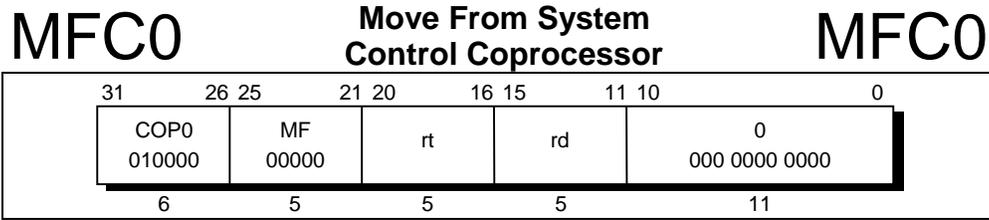
If rd is omitted in assembly language, 0 is used as the default value. To guarantee correct operation even if an interrupt occurs, neither of the two instructions following MADDU should be DIV or DIVU instructions which modify the HI and LO register contents.

Operation:

<p>32, 64 T: $t \leftarrow (HI \parallel LO) + (0 \parallel GPR[rs]) \cdot (0 \parallel GPR[rt])$ $LO \leftarrow t_{31-0}$ $HI \leftarrow t_{63-32}$ $GPR[rd] \leftarrow t_{31-0}$</p>

Exception:

None

**Format:**MFC0 *rt*, *rd***Description:**

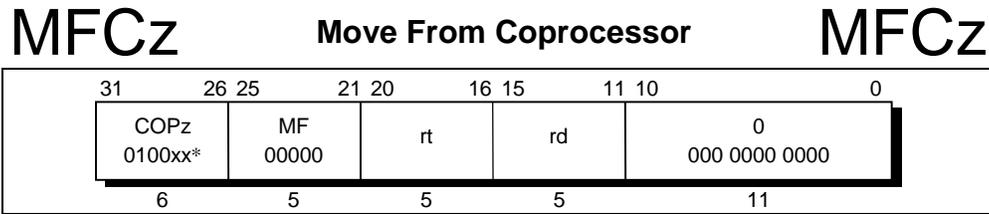
The contents of coprocessor register *rd* of the CP0 are loaded into general register *rt*.
May be used on both 32-bit and 64-bit CP0 registers.

Operation:

32	T:	$data \leftarrow CPR[0,rd]$
	T + 1:	$GPR[rt] \leftarrow data$
64	T:	$data \leftarrow CPR[0,rd]$
	T + 1:	$GPR[rt] \leftarrow (data_{31})^{32} data_{31-0}$

Exceptions:

Coprocessor unusable exception

**Format:**

MFCz rt, rd

Description:

The contents of coprocessor register *rd* of coprocessor *z* are loaded into general register *rt*.

Execution of the instruction referencing coprocessor 3 causes a reserved instruction exception, not a coprocessor unusable exception.

Operation:

```

32 T:   data ← CPR[z,rd]
      T + 1: GPR[rt] ← data
64 T:   if rd0 = 0
          data ← CPR[z,rd4-1 || 0]31-0
          else
          data ← CPR[z,rd4-1 || 0]63-32
          endif
      T + 1: GPR[rt] ← (data31)32 || data

```

Exceptions:

Coprocessor unusable exception

Reserved instruction exception (coprocessor 3)

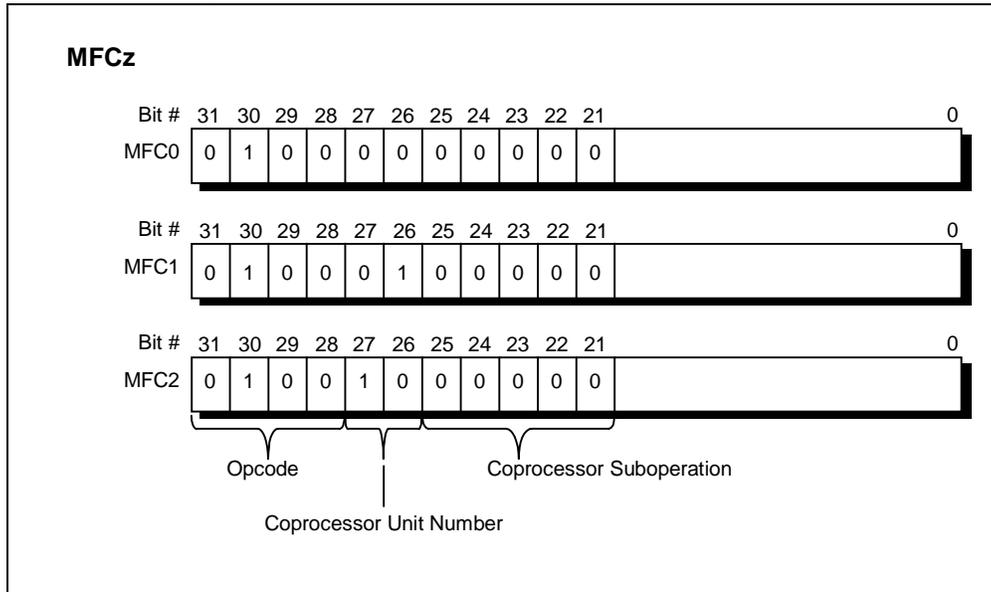
*See the table “Opcode Bit Encoding” on next page, or “CPU Instruction Opcode Bit Encoding” at the end of Appendix A.

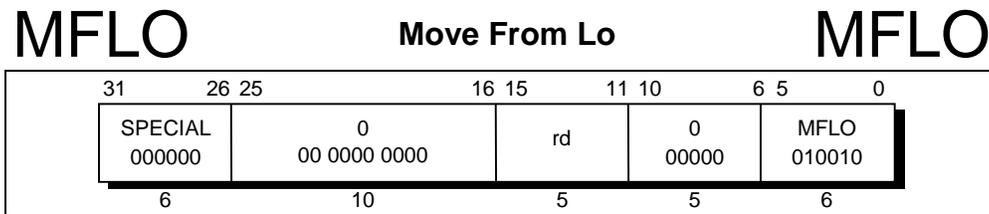
MFCz

Move From Coprocessor (continued)

MFCz

Opcode Bit Encoding:



**Format:**

MFLO rd

Description:

The contents of special register *LO* are loaded into general register *rd*.

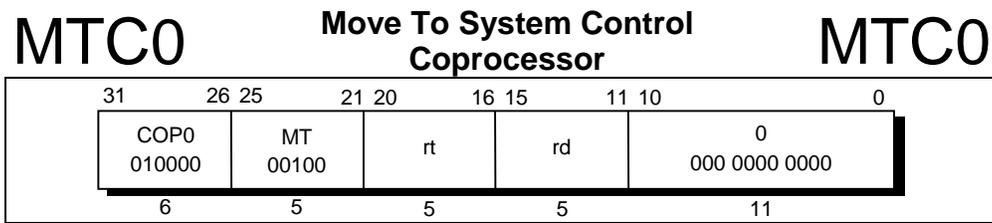
To ensure proper operation in the event of interruptions, the two instructions which follow a MFLO instruction may not be any of the instructions which modify the *LO* register: MULT, MULTU, DIV, DIVU, MTLO, DMULT, DMULTU, DDIV, DDIVU.

Operation:

32, 64 T: GPR[rd] ← LO

Exceptions:

None

**Format:**MTC0 *rt*, *rd***Description:**

The contents of general register *rt* are loaded into coprocessor register *rd* of the CP0.

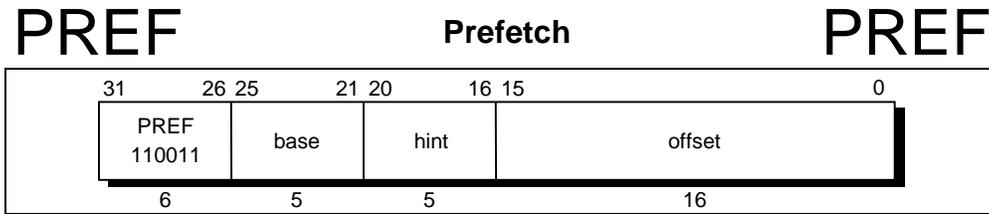
Because the state of the virtual address translation system may be altered by this instruction, the operation of load, store instructions and TLB operations immediately prior to and after this instruction are undefined.

Operation:

$32, 64 \text{ T: } \quad \text{data} \leftarrow \text{GPR}[\text{rt}]$ $\text{T} + 1: \text{CPR}[0, \text{rd}] \leftarrow \text{data}$

Exceptions:

Coprocessor unusable exception



Format :

PREF hint, offset (base)

Description :

PREF adds the 16-bit signed offset to the contents of GPR base to form an effective byte address. It advises that data at the effective address may be used in the near future.

If the hint field is 00000₂, this instruction prefetches a block of data from main memory into cache.

PREF is an advisory instruction. It may change the performance of the program. For all hint values and all effective addresses, it neither changes architecturally-visible state nor alters the meaning of the program.

PREF does not cause addressing-related exceptions. If it raises an exception condition, the exception conditions ignored. If an addressing-related exception is raised and ignored, no data will be prefetched, even if no data is prefetched in such a case, some action that is not architecturally-visible, such as writeback of a dirty cache line, might take place.

PREF will never generate a memory operation for a location with an uncached memory access type.

The defined hint values are shown in the table below. The TX49 only supports hint□=□0. The hint table may be extended in future implementations.

hint field: Value

Value	Name	Data use and desired prefetch action
0	Load	Data is expected to be loaded (not modified). Fetch data as if for a load.
1-31	Reserved	Reserved

PREF

Prefetch
(continued)

PREF

Programming Notes:

Prefetch can not prefetch data from a mapped location unless the translation for that location is present in the TLB. Locations in memory pages that have not been accessed recently may not have translations in the TLB, so prefetch may not be effective for such locations.

Prefetch does not cause addressing exceptions. It will not cause an exception to prefetch using an address pointer value before the validity of a pointer determined.

Operation :

32, 64	T:	$vAddr \leftarrow GPR[base] \bullet \Rightarrow \text{sign_extend}(\text{offset})$ $(pAddr, \text{uncached}) \leftarrow \text{Address Translation}(vAddr, \text{DATA}, \text{LOAD})$ Prefetch (uncached, pAddr, vAddr, DATA, hint)
--------	----	---

Exception :

None

SC

Store Conditional
(continued)

SC

If this instruction should both fail and take an exception, the exception takes precedence.

Operation:

```

32  T:    vAddr ← ((offset15)16 || offset15-0) + GPR[base]
        (pAddr, unchached) ← AddressTranslation (vAddr, DATA)
        pAddr ← pAddrPSIZE.1-3 || (pAddr2-0 xor ReverseEndian || 02)
        data ← GPR[rt]63-8*byte-0 || 08*byte
        if LLbit then
            StoreMemory (uncached, WORD, data, pAddr, vAddr, DATA)
        endif
        GPR[rt] ← 031 || LLbit

64  T:    vAddr ← ((offset15)48 || offset15-0) + GPR[base]
        (pAddr, unchached) ← AddressTranslation (vAddr, DATA)
        pAddr ← pAddrPSIZE.1-3 || (pAddr2-0 xor ReverseEndian || 02)
        data ← GPR[rt]63-8*byte-0 || 08*byte
        if LLbit then
            StoreMemory (uncached, WORD, data, pAddr, vAddr, DATA)
        endif
        GPR[rt] ← 063 || LLbit

```

Exceptions:

- TLB refill exception
- TLB invalid exception
- TLB modification exception
- Bus error exception
- Address error exception

SCD**Store Conditional
Doubleword
(continued)****SCD****Exceptions:**

TLB refill exception

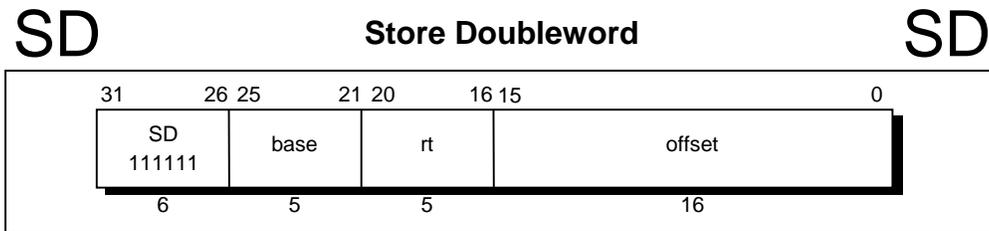
TLB invalid exception

TLB modification exception

Bus error exception

Address error exception

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**SD *rt*, *offset* (*base*)**Description:**

The 16-bit *offset* is sign-extended and added to the contents of general register *base* to form a virtual address. The contents of general register *rt* are stored at the memory location specified by the effective address.

If either of the three least-significant bits of the effective address are non-zero, an address error exception occurs.

Operation:

```

64 T:    vAddr ← ((offset15)48 || offset15-0) + GPR[base]
          (pAddr, uncached) ← AddressTranslation(vAddr, DATA)
          data ← GPR[rt]
          StoreMemory(uncached, DOUBLEWORD, data, pAddr, vAddr, DATA)

```

Exceptions:

TLB refill exception

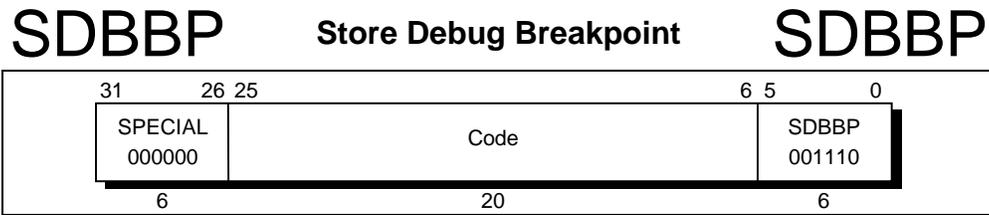
TLB invalid exception

TLB modification exception

Bus error exception

Address error exception

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**

SDBBP code

Description:

Raises a Debug Breakpoint exception, passing control to an exception handler. The code field can be used for passing information to the exception handler, but the only way to have the code field retrieved by the exception handler is to load the contents of the memory word containing this instruction using the DEPC register.

Operation:

32, 64	T:	Software Debug Breakpoint Exception
--------	----	-------------------------------------

Exception:

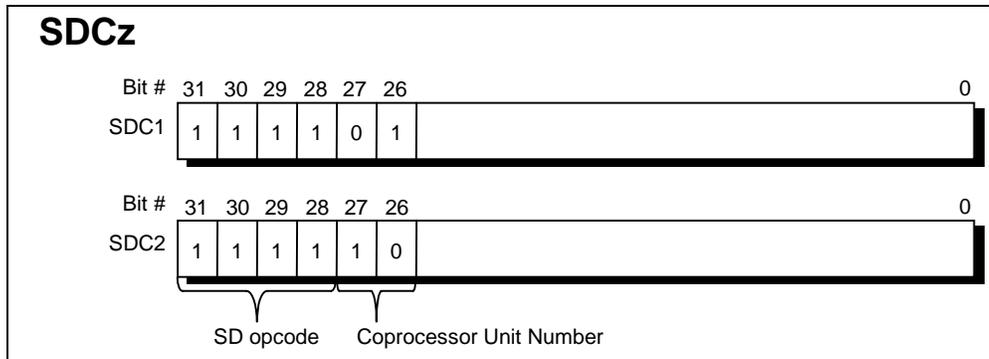
Debug Breakpoint exception

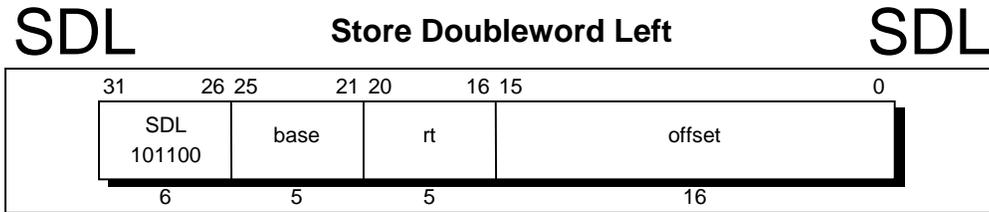
SDCz**Store Doubleword From Coprocessor (continued)****SDCz****Operation:**

32	T:	$vAddr \leftarrow ((offset_{15})^{16} offset_{15-0}) + GPR[base]$ $(pAddr, uncached) \leftarrow AddressTranslation(vAddr, DATA)$ $data \leftarrow COPzSD(rt),$ StoreMemory(uncached, DOUBLEWORD, data, pAddr, vAddr, DATA)
64	T:	$vAddr \leftarrow ((offset_{15})^{48} offset_{15-0}) + GPR[base]$ $(pAddr, uncached) \leftarrow AddressTranslation(vAddr, DATA)$ $data \leftarrow COPzSD(rt),$ StoreMemory(uncached, DOUBLEWORD, data, pAddr, vAddr, DATA)

Exceptions:

- TLB refill exception
- TLB invalid exception
- TLB modification exception
- Bus error exception
- Address error exception
- Coprocessor unusable exception

Opcode Bit Encoding:

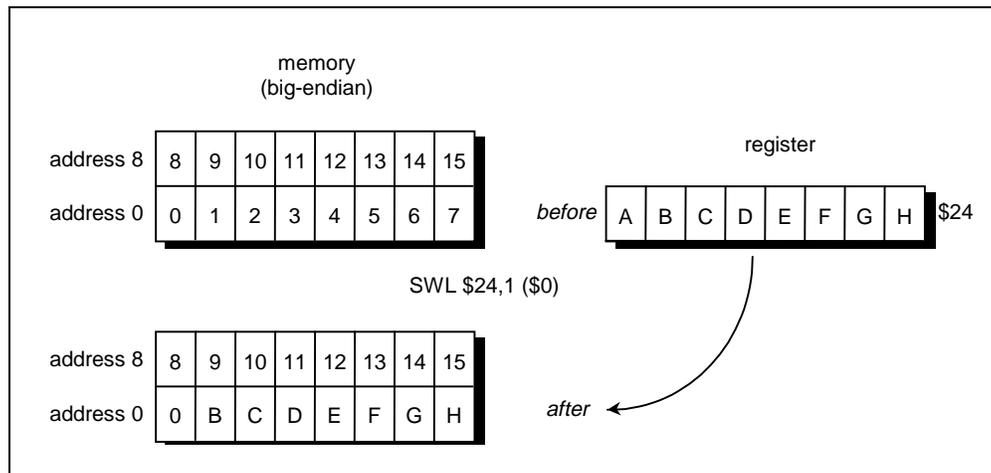
**Format:**SDL *rt*, *offset* (*base*)**Description:**

This instruction can be used with the SDR instruction to store the contents of a register into eight consecutive bytes of memory, when the bytes cross a boundary between two doublewords. SDL stores the left portion of the register into the appropriate part of the high-order doubleword of memory; SDR stores the right portion of the register into the appropriate part of the low-order doubleword.

The SDL instruction adds its sign-extended 16-bit *offset* to the contents of general register *base* to form a virtual address which may specify an arbitrary byte. It alters only the word in memory which contains that byte. From one to four bytes will be stored, depending on the starting byte specified.

Conceptually, it starts at the most-significant byte of the register and copies it to the specified byte in memory; then it proceeds toward the low-order byte of the register and the low-order byte of the word in memory, copying bytes from register to memory until it reaches the low-order byte of the word in memory.

No address exceptions due to alignment are possible.



SDL

Store Doubleword Left
(continued)

SDL

This operation is only defined for the TX4300 operating in 64-bit mode and 32-bit kernel mode.

Execution of this instruction in 32-bit user or supervisor mode causes a reserved instruction exception.

Operation:

```

64  T:      vAddr ← ((offset15)48 || offset15-0) + GPR[base]
          (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
          pAddr ← pAddrPSIZE.1-3 || (pAddr2-0 xor ReverseEndian3)
          If BigEndianMem = 0 then
              pAddr ← pAddr31-3 || 03
          endif
          byte ← vAddr2-0 xor BigEndianCPU3
          data ← 056-8*byte || GPR[rt]63-56-8*byte
          StoreMemory (uncached, byte, data, pAddr, vAddr, DATA)

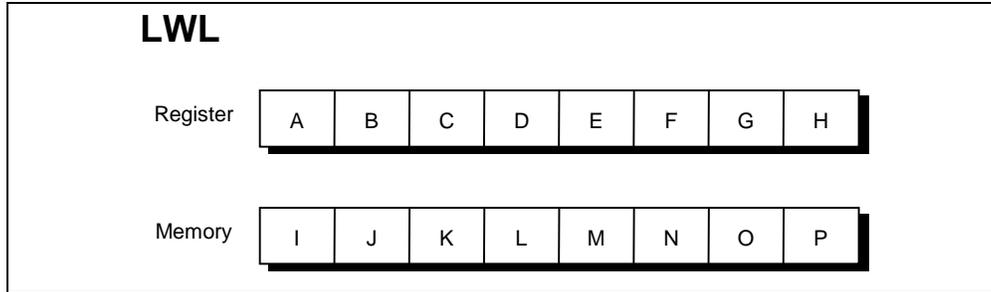
```

SDL

Store Doubleword Left (continued)

SDL

Given a doubleword in a register and a doubleword in memory, the operation of SWL is as follows:



vAddr ₂₋₀	BigEndianCPU ₀ ⇒ 0					BigEndianCPU ₀ ⇒ 1				
	destination	type	offset		destination	type	offset			
			LEM	BEM			LEM	BEM		
0	I J K L M N O A	□	0	□	A □ □ □ □ □ □ H	7	0	0		
1	I J K L M N □ □	□	□	□	□ □ □ □ □ □ □ □	6	0	1		
2	I J K L M A B C	2	0	5	I J A B C D E F	5	0	2		
3	I J K L A B C D	3	0	4	I J K A B C D E	4	0	3		
4	I J K A B C D E	4	0	3	I J K L A B C D	3	0	4		
5	I J A B C D E F	5	0	2	I J K L M A B C	2	0	5		
6	I A B C D E F G	6	0	1	I J K L M N A B	1	0	6		
7	A B C D E F G H	7	0	0	I J K L M N O A	0	0	7		

LEM BigEndianMem₀ = 0

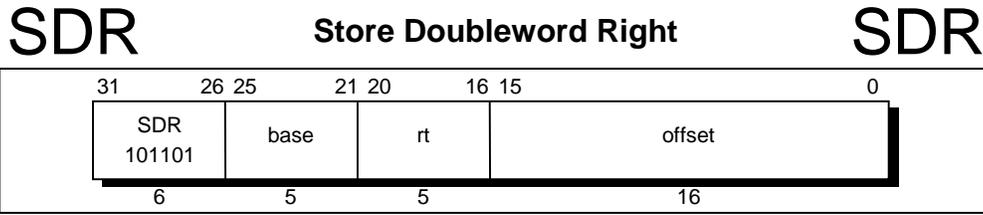
BEM BigEndianMem₀ = 1

Type Access Type (see Figure 2-2) sent to memory

Offset pAddr₂₋₀ sent to memory

Exceptions:

- TLB refill exception
- TLB invalid exception
- TLB modification exception
- Bus error exception
- Address error exception
- Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)



Format:

SDR rt, offset (base)

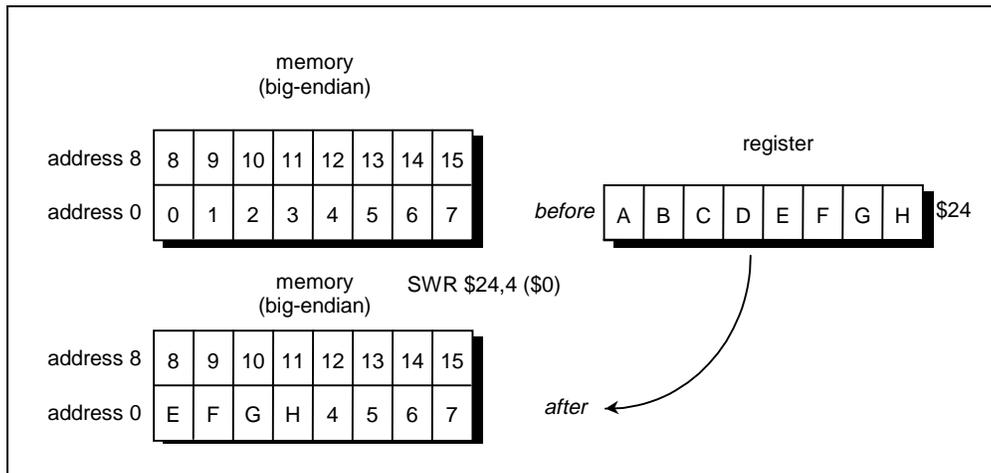
Description:

This instruction can be used with the SDL instruction to store the contents of a register into eight consecutive bytes of memory, when the bytes cross a boundary between two doublewords. SDR stores the right portion of the register into the appropriate part of the low-order doubleword; SDL stores the left portion of the register into the appropriate part of the low-order doubleword of memory.

The SDR instruction adds its sign-extended 16-bit *offset* to the contents of general register *base* to form a virtual address which may specify an arbitrary byte. It alters only the word in memory which contains that byte. From one to eight bytes will be stored, depending on the starting byte specified.

Conceptually, it starts at the least-significant (rightmost) byte of the register and copies it to the specified byte in memory; then it proceeds toward the high-order byte of the register and the high-order byte of the word in memory, copying bytes from register to memory until it reaches the high-order byte of the word in memory.

No address exceptions due to alignment are possible.



SDR

Store Doubleword Right
(continued)

SDR

This operation is only defined for the TX4300 operating in 64-bit mode and 32-bit kernel mode.

Execution of this instruction in 32-bit user or supervisor mode causes a reserved instruction exception

Operation:

```

64  T:      vAddr ← ((offset15)48 || offset15-0) + GPR[base]
          (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
          pAddr ← pAddrPSIZE.1-3 || (pAddr2-0 xor ReverseEndian3)
          if BigEndianMem = 0 then
              pAddr ← pAddrPSIZE.31-3 || 03
          endif
          byte ← vAddr1-0 xor BigEndianCPU3
          data ← GPR[rt]63-8*byte-0 || 08*byte
          StoreMemory (uncached, DOUBLEWORD-byte, data, pAddr, vAddr,

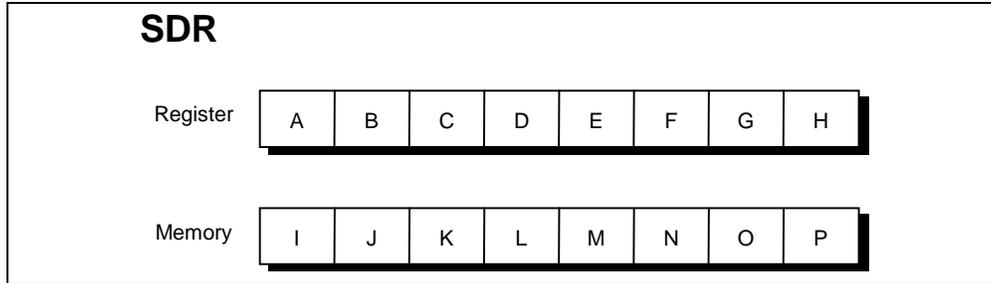
```

SDR

Store Doubleword Right (continued)

SDR

Given a doubleword in a register and a doubleword in memory, the operation of SDR is as follows:



vAddr ₂₋₀	BigEndianCPU ₀ ⇒ 0				BigEndianCPU ₁ ⇒ 1			
	destination	type	offset		destination	type	offset	
			LEM	BEM			LEM	BEM
0	A B C D E F G H	7	0	0	H J K L M N O P	0	7	0
1	B C D E F G H P	6	1	0	G H K L M N O P	1	6	0
2	C D E F G H O P	5	2	0	F G H L M N O P	2	5	0
3	D E F G H N O P	4	3	0	E F G H M N O P	3	4	0
4	E F G H M N O P	3	4	0	D E F G H N O P	4	3	0
5	F G H L M N O P	2	5	0	C D E F G H O P	5	2	0
6	G H K L M N O P	1	6	0	B C D E F G H P	6	1	0
7	H J K L M N O P	0	7	0	A B C D E F G H	7	0	0

LEM BigEndianMem = 0

BEM BigEndianMem = 1

Type Access Type (see Figure 2-2) sent to memory

Offset pAddr₂₋₀ sent to memory

Exceptions:

TLB refill exception

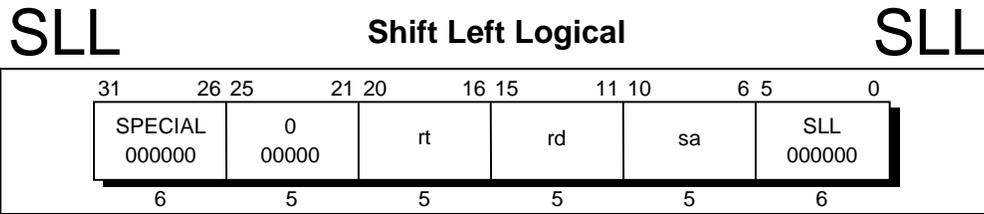
TLB invalid exception

TLB modification exception

Bus error exception

Address error exception

Reserved Instruction exception (in the 32-bit user or 32-bit supervisor mode)

**Format:**

SLL rd, rt, sa

Description:

The contents of general register *rt* are shifted left by *sa* bits, inserting zeros into the low-order bits. The result is placed in register *rd*. In 64-bit mode, the 32-bit result is sign extended when placed in the destination register. It is sign-extended for all shift amounts, including zero; SLL with a zero shift amount truncates a 64-bit value to 32-bits and sign extends this 32-bit value. SLL, unlike nearly all other word operations, does not require an operand to be a properly sign-extended word value to produce a valid sign-extended word result.

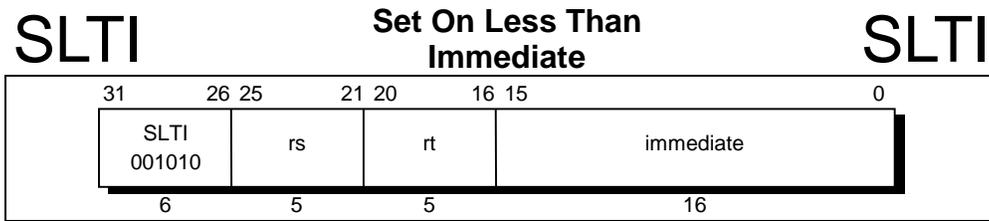
Note: SLL with a shift amount of zero may be treated as a NOP by some assemblers at some optimization levels. If using SLL with zero shift to truncate 64-bit values, check the assembler being used.

Operation:

32	T:	$GPR[rd] \leftarrow GPR[rt]_{31-sa-0} \parallel 0^{sa}$
64	T:	$s \leftarrow 0 \parallel sa$ $temp \leftarrow GPR[rt]_{31-s-0} \parallel 0^s$ $GPR[rd] \leftarrow (temp_{31})^{32} \parallel temp$

Exceptions:

None

**Format:**SLTI *rt*, *rs*, *immediate***Description:**

The 16-bit *immediate* is sign-extended and subtracted from the contents of general register *rs*. Considering both quantities as signed integers, if *rs* is less than the sign-extended *immediate*, the result is set to one, otherwise the result is set to zero. The result is placed into general register *rt*.

No integer overflow exception occurs under any circumstances. The comparison is valid even if the subtraction used during the comparison overflows.

Operation:

```

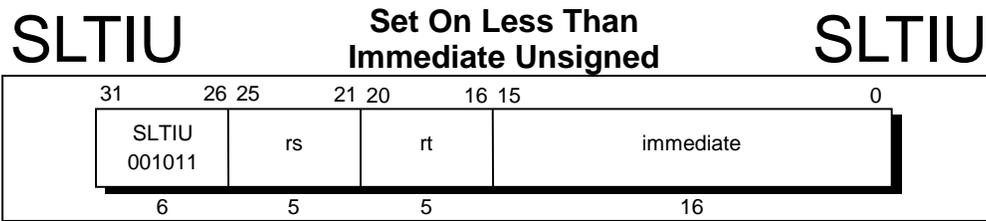
32  T:   if GPR[rs] <• (immediate15)16 || immediate15-0 then
        GPR[rt] ← 031 || 1
        else
        GPR[rt] ← 032
        endif

64  T:   if GPR[rs] <• (immediate15)48 || immediate15-0 then
        GPR[rt] ← 063 || 1
        else
        GPR[rt] ← 064
        endif

```

Exceptions:

None

**Format:**

SLTIU rt, rs, immediate

Description:

The 16-bit *immediate* is sign-extended and subtracted from the contents of general register *rs*. Considering both quantities as unsigned integers, if *rs* is less than the sign-extended immediate, the result is set to one, otherwise the result is set to zero. The result is placed into general register *rt*.

No integer overflow exception occurs under any circumstances. The comparison is valid even if the subtraction used during the comparison overflows.

Operation:

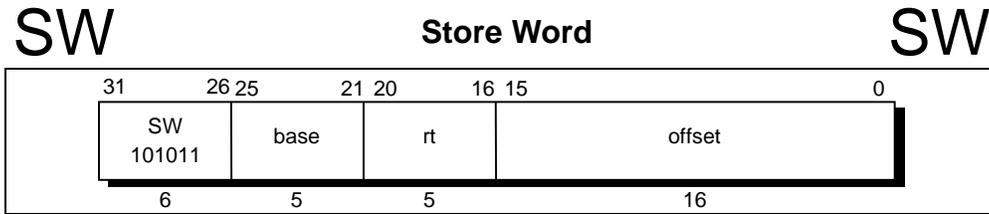
```

32  T:   if (0 || GPR[rs]) <• (immediate15)16 || immediate15-0 then
        GPR[rt] ← 031 || 1
        else
        GPR[rt] ← 032
        endif
64  T:   if (0 || GPR[rs]) <• (immediate15)48 || immediate15-0 then
        GPR[rt] ← 063 || 1
        else
        GPR[rt] ← 064
        endif

```

Exceptions:

None

**Format:**

SW rt, offset (base)

Description:

The 16-bit *offset* is sign-extended and added to the contents of general register *base* to form a virtual address. The contents of general register *rt* are stored at the memory location specified by the effective address.

If either of the two least-significant bits of the effective address are non-zero, an address error exception occurs.

Operation:

32	T:	$vAddr \leftarrow ((offset_{15})^{16} \parallel offset_{15-0}) \bullet + \bullet GPR[base]$ $(pAddr, uncached) \leftarrow AddressTranslation(vAddr, DATA)$ $pAddr \leftarrow pAddr_{PSIZE-1-3} \parallel (pAddr_{2-0} \text{ xor } (ReverseEndian \parallel 0^2))$ $byte \leftarrow vAddr_{2-0} \text{ xor } (BigEndianCPU \parallel 0^2)$ $data \leftarrow GPR[rt]_{63-8*byte} \parallel 0^{8*byte}$ $StoreMemory(uncached, WORD, data, pAddr, vAddr, DATA)$
64	T:	$vAddr \leftarrow ((offset_{15})^{48} \parallel offset_{15-0}) \bullet + \bullet GPR[base]$ $(pAddr, uncached) \leftarrow AddressTranslation(vAddr, DATA)$ $pAddr \leftarrow pAddr_{PSIZE-1-3} \parallel (pAddr_{2-0} \text{ xor } (ReverseEndian \parallel 0^2))$ $byte \leftarrow vAddr_{2-0} \text{ xor } (BigEndianCPU \parallel 0^2)$ $data \leftarrow GPR[rt]_{63-8*byte} \parallel 0^{8*byte}$ $StoreMemory(uncached, WORD, data, pAddr, vAddr, DATA)$

Exceptions:

- TLB refill exception
- TUB invalid exception
- TLB modification exception
- Bus error exception
- Address error exception

SWCz

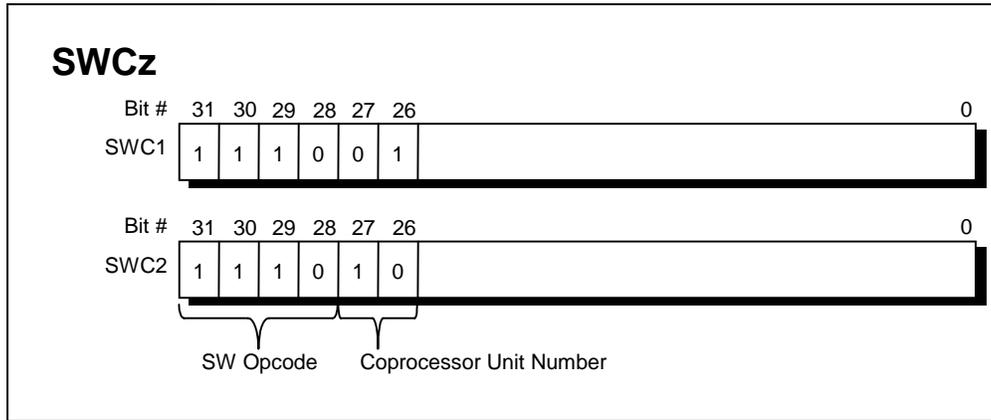
Store Word From Coprocessor (Continued)

SWCz

Exceptions:

- TLB refill exception
- TLB invalid exception
- TLB modification exception
- Bus error exception
- Address error exception
- Coprocessor unusable exception

Opcode Bit Encoding:



SWL

Store Word Left
(Continued)

SWL

Operation:

```

32  T:    vAddr ← ((offset15)16 || offset15-0) • +• GPR[base]
        (pAddr, unchached) ← AddressTranslation (vAddr, DATA)
        pAddr ← pAddrPSIZE-1-3 || (pAddr2-0 xor ReverseEndian3)
        if BigEndianMem • ⇒ 0 then
        pAddr ← pAddr31-2 || 02
        endif
        byte ← vAddr1-0 xor BigEndianCPU2
        if (vAddr2 xor BigEndianCPU) • ⇒ 0 then
        data ← 032 || 024-8*byte || GPR[rt]31-24-8*byte
        else
        data ← 024-8*byte || GPR[rt]31-24-8*byte || 032
        endif
        StoreMemory (uncached, byte, data, pAddr, vAddr, DATA)

64  T:    vAddr ← ((offset15)48 || offset15-0) • +• GPR[base]
        (pAddr, unchached) ← AddressTranslation (vAddr, DATA)
        pAddr ← pAddrPSIZE-1-3 || (pAddr2-0 xor ReverseEndian3)
        if BigEndianMem • ⇒ 0 then
        pAddr ← pAddr31-2 || 02
        endif
        byte ← vAddr1-0 xor BigEndianCPU2
        if (vAddr2 xor BigEndianCPU) • ⇒ 0 then
        data ← 032 || 024-8*byte || GPR[rt]31-24-8*byte
        else
        data ← 024-8*byte || GPR[rt]31-24-8*byte || 032
        endif
        StoreMemory (uncached, byte, data, pAddr, vAddr, DATA)

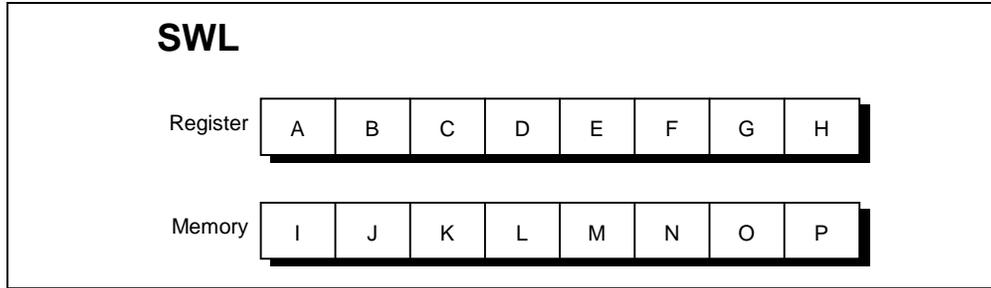
```

SWL

Store Word Left (Continued)

SWL

Given a doubleword in a register and a doubleword in memory, the operation of SWL is as follows:



vAddr2_0	BigEndianCPU = 0				BigEndianCPU = 1			
	Destination	type	offset		Destination	type	offset	
			LEM	BEM			LEM	BEM
0	I J K L M N O E	0	0	7	E F G H M N O P	3	4	0
1	I J K L M N E F	1	0	6	I E F G M N O P	2	4	1
2	I J K L M E F G	2	0	5	I J E F M N O P	1	4	2
3	I J K L E F G H	3	0	4	I J K E M N O P	0	4	3
4	I J K E M N O P	0	4	3	I J K L E F G H	3	0	4
5	I J E F M N O P	1	4	2	I J K L M E F G	2	0	5
6	I E F G M N O P	2	4	1	I J K L M N E F	1	0	6
7	E F G H M N O P	3	4	0	I J K L M N O E	0	0	7

LEM BigEndianMem = 0

BEM BigEndianMem = 1

Type AccessType (see Figure 2-2) sent to memory

Offset pAddr2_0 sent to memory

Exceptions:

- TLB refill exception
- TLB invalid exception
- TLB modification exception
- Bus error exception
- Address error exception

SWR

Store Word Right
(Continued)

SWR

Operation:

```

32  T:    vAddr ← ((offset15)16 || offset15-0) +• GPR[base]
        (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
        pAddr ← pAddrPSIZE-1-3 || (pAddr2-0 xor ReverseEndian3)
        if BigEndianMem• =• 0 then
        pAddr ← pAddr31-2 || 02
        endif
        byte ← vAddr1-0 xor BigEndianCPU2
        if (vAddr2 xor BigEndianCPU)• =• 0 then
        data ← 032 || GPR[rt]31-8*byte-0 || 08*byte
        else
        data ← GPR[rt]31-8*byte-0 || 08*byte || 032
        endif
        StoreMemory (uncached, WORD-byte, data, pAddr, vAddr, DATA)

64  T:    vAddr ← ((offset15)48 || offset15-0)• +• GPR[base]
        (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
        pAddr ← pAddrPSIZE-1-3 || (pAddr2-0 xor ReverseEndian3)
        if BigEndianMem• =• 0 then
        pAddr ← pAddr31-2 || 02
        endif
        byte ← vAddr1-0 xor BigEndianCPU2
        if (vAddr2 xor BigEndianCPU)• =• 0 then
        data ← 032 || GPR[rt]31-8*byte-0 || 08*byte
        else
        data ← GPR[rt]31-8*byte-0 || 08*byte || 032
        endif
        StoreMemory (uncached, WORD-byte, data, pAddr, vAddr, DATA)

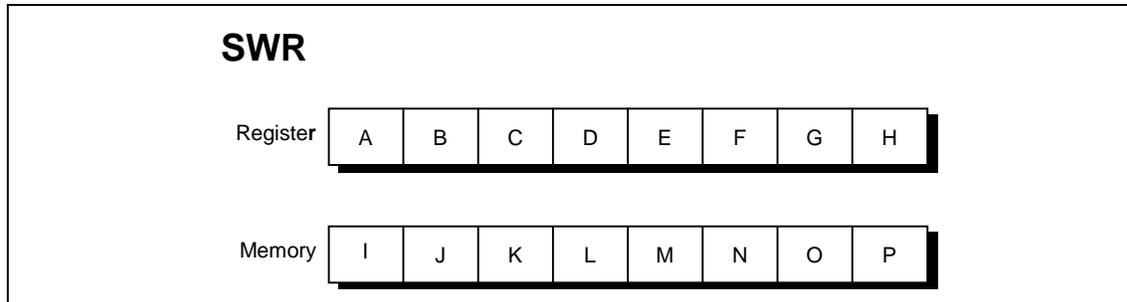
```

SWR

Store Word Right
(Continued)

SWR

Given a doubleword in a register and a doubleword in memory, the operation of SWR is as follows:



vAddr ₂₋₀	BigEndianCPU = 0				BigEndianCPU = 1			
	Destination	type	offset		Destination	type	offset	
			LEM	BEM			LEM	BEM
0	I J K L E F G H	3	0	4	E J K L M N O P	0	7	0
1	I J K L F G H P	2	1	4	G H K L M N O P	1	6	0
2	I J K L G H O P	1	2	4	F G H L M N O P	2	5	0
3	I J K L H N O P	0	3	4	E F G H M N O P	3	4	0
4	E F G H M N O P	3	4	0	I J K L H N O P	0	3	4
5	F G H L M N O P	2	5	0	I J K L G H O P	1	2	4
6	G H K L M N O P	1	6	0	I J K L F G H P	2	1	4
7	H J K L M N O P	0	7	0	I J K L E F G H	3	0	4

LEM BigEndianMem = 0

BEM BigEndianMem = 1

Type AccessType (see Figure 2-2) sent to memory

Offset pAddr₂₋₀ sent to memory

Exceptions:

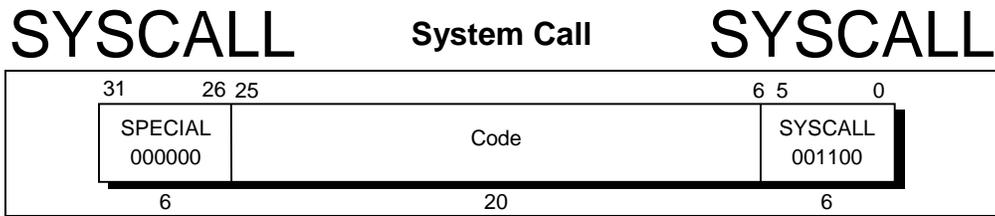
TLB refill exception

TLB invalid exception

TLB modification exception

BUS error exception

Address error exception

**Format:**

SYSCALL

Description:

A system call exception occurs, immediately and unconditionally transferring control to the exception handler.

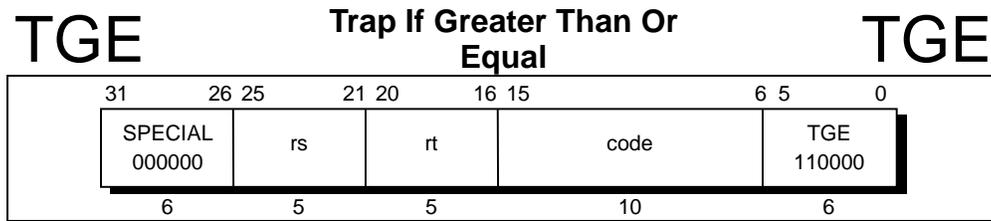
The code field is available for use as software parameters, but is retrieved by the exception handler only by loading the contents of the memory word containing the instruction.

Operation:

32, 64 T:	SystemCallException
-----------	---------------------

Exceptions:

System Call exception

**Format:**

TGE rs, rt

Description:

The contents of general register *rt* are compared to the contents of general register *rs*. Considering both quantities as signed integers, if the contents of general register *rs* are greater than or equal to the contents of general register *rt*, a trap exception occurs.

The code field is available for use as software parameters, but is retrieved by the exception handler only by loading the contents of the memory word containing the instruction.

Operation:

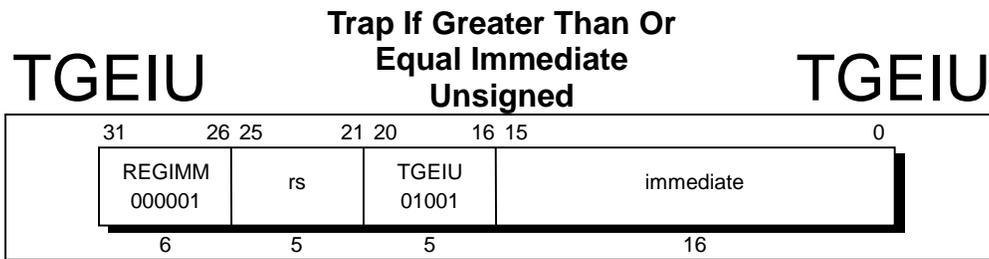
```

32, 64 T:   if GPR[rs] ≥ GPR[rt] then
            TrapException
            endif

```

Exceptions:

Trap exception

**Format:**

TGEIU rs, immediate

Description:

The 16-bit *immediate* is sign-extended and compared to the contents of general register *rs*. Considering both quantities as unsigned integers, if the contents of general register *rs* are greater than or equal to the sign-extended *immediate*, a trap exception occurs.

Operation:

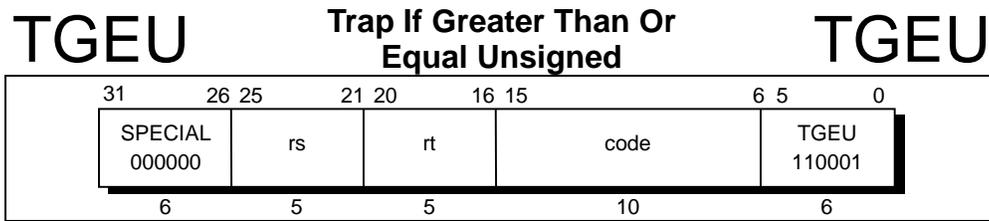
```

32  T:    if (0 || GPR[rs]) • ≥ • (0 || (immediate15)16 || immediate15-0) then
        TrapException
        endif
64  T:    if (0 || GPR[rs]) • ≥ • (0 || (immediate15)48 || immediate15-0) then
        TrapException
        endif

```

Exceptions:

Trap exception

**Format:**TGEU *rs*, *rt***Description:**

The contents of general register *rt* are compared to the contents of general register *rs*. Considering both quantities as unsigned integers, if the contents of general register *rs* are greater than or equal to the contents of general register *rt*, a trap exception occurs.

The code field is available for use as software parameters, but is retrieved by the exception handler only by loading the contents of the memory word containing the instruction.

Operation:

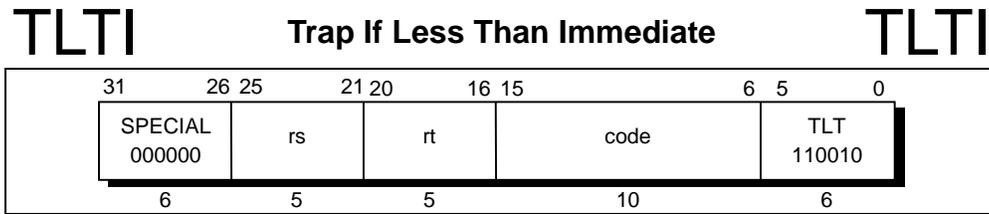
```

32, 64   T:   if (0 || GPR[rs]) • ≥ • (0 || GPR[rt]) then
           TrapException
           endif

```

Exceptions:

Trap exception

**Format:**

TLTI rs, immediate

Description:

The 16-bit *immediate* is sign-extended and compared to the contents of general register *rs*. Considering both quantities as signed integers, if the contents of general register *rs* are less than the sign-extended *immediate*, a trap exception occurs.

Operation:

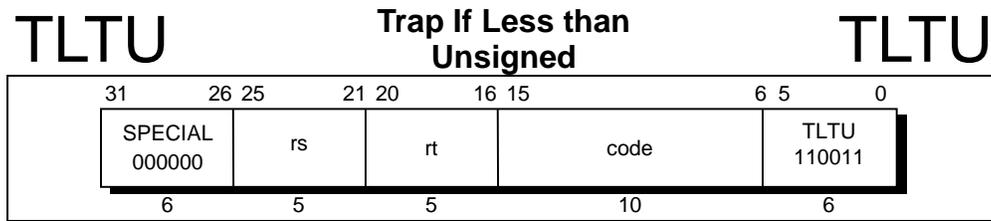
```

32  T:    if GPR[rs] <• (immediate15)16 || immediate15-0 then
        TrapException
        endif
64  T:    if GPR[rs] <• (immediate15)48 || immediate15-0 then
        TrapException
        endif

```

Exceptions:

Trap exception

**Format:**

TLTU rs, rt

Description:

The contents of general register *rt* are compared to general register *rs*. Considering both quantities as unsigned integers, if the contents of general register *rs* are less than the contents of general register *rt*, a trap exception occurs.

The code field is available for use as software parameters, but is retrieved by the exception handler only by loading the contents of the memory word containing the instruction.

Operation:

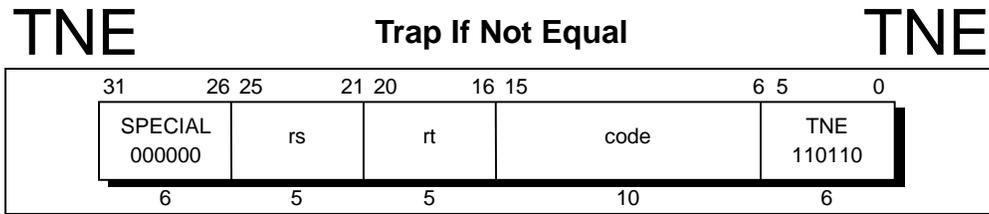
```

32, 64 T:   if (0 || GPR [rs]) <• (0 || GPR [rt]) then
            TrapException
            endif

```

Exceptions:

Trap exception

**Format:**TNE *rs*, *rt***Description:**

The contents of general register *rt* are compared to general register *rs*. If the contents of general register *rs* are not equal to the contents of general register *rt*, a trap exception occurs.

The code field is available for use as software parameters, but is retrieved by the exception handler only by loading the contents of the memory word containing the instruction.

Operation:

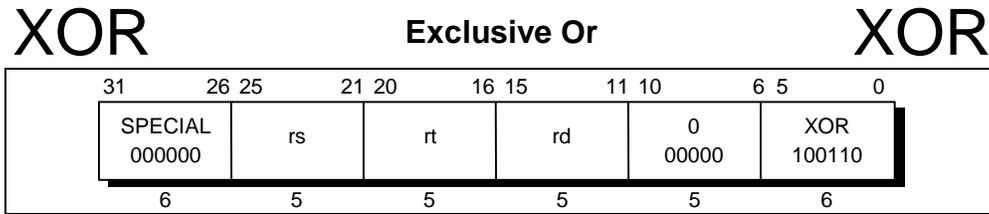
```

32, 64 T:   if GPR [rs] ≠ GPR [rt] then
            TrapException
            endif

```

Exceptions:

Trap exception

**Format:**

XOR rd, rs, rt

Description:

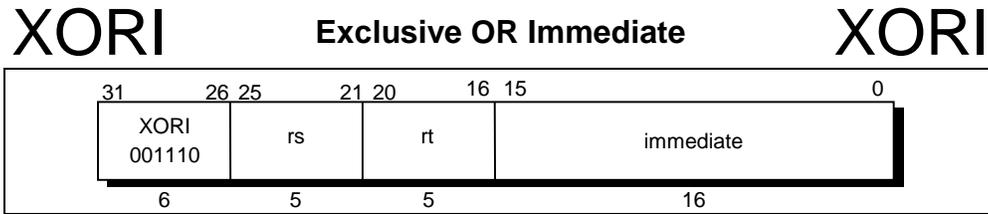
The contents of general register *rs* are combined with the contents of general register *rt* in a bit-wise logical exclusive OR operation. The result is placed into general register *rd*.

Operation:

32, 64 T: GPR [rd] ← GPR [rs] <i>xor</i> GPR [rt]

Exceptions:

None

**Format:**

XORI rt, rs, immediate

Description:

The 16-bit *immediate* is zero-extended and combined with the contents of general register *rs* in a bit-wise logical exclusive OR operation. The result is placed into general register *rt*.

Operation:

32	T:	$GPR [rt] \leftarrow GPR [rs] \text{ xor } (0^{16} \text{immediate})$
64	T:	$GPR [rt] \leftarrow GPR [rs] \text{ xor } (0^{48} \text{immediate})$

Exceptions:

None

Bit Encoding of CPU Instruction OPcodes

The Table A-2 shows the bit codes for all TX49 CPU instructions(ISA and extended ISA)

Table A-2 CPU Operation Code Bit Encoding

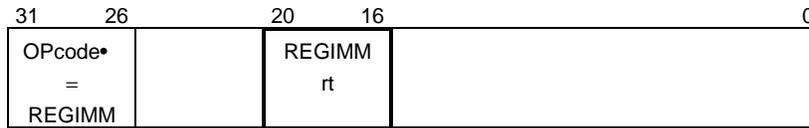
OPcode

31
26
0

	OPcode							
--	--------	--	--	--	--	--	--	--

		28-26						

REGIMM rt



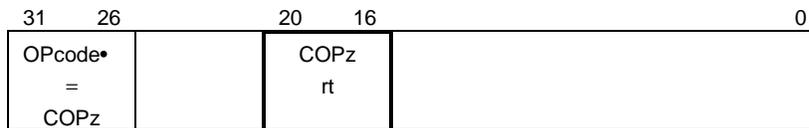
		18-16							
		0	1	2	3	4	5	6	7
20-19	0	BLTZ	BGEZ	BLTZL	BGEZL	*	*	*	*
1	1	TGEI	TGEIU	TLTI	TLTIU	TEQI	*	TNEI	*
2	2	BLTZAL	BGEZAL	BLTZALL	BGEZALL	*	*	*	*
3	3	*	*	*	*	*	*	*	*

COPz rs



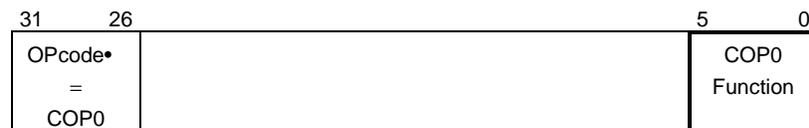
		23-21							
		0	1	2	3	4	5	6	7
25-24	0	MF	DMF ϵ	CF	γ	MT	DMT ϵ	CT	γ
1	1	BC	γ	γ	γ	γ	γ	γ	γ
2	2	CO							
3	3								

COPz rt



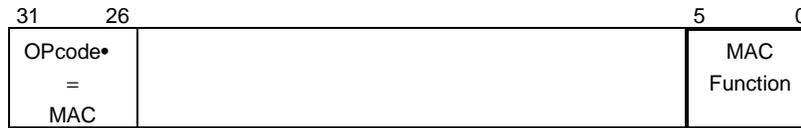
		18-16							
		0	1	2	3	4	5	6	7
20-19	0	BCF	BCT	BCFL	BCTL	γ	γ	γ	γ
1	1	γ							
2	2	γ							
3	3	γ							

COP0 Function



		2-0							
		0	1	2	3	4	5	6	7
5-3	0	ϕ	TLBR	TLBWI	ϕ	ϕ	ϕ	TLBWR	ϕ
1	1	TLBP	ϕ						
2	2	ϕ							
3	3	ERET	ϕ	ϕ	ϕ	ϕ	ϕ	ϕ	DERET
4	4	WAIT	ϕ						
5	5	ϕ							
6	6	ϕ							
7	7	ϕ							

MAC Function



	2-0								
5-3	0	1	2	3	4	5	6	7	
0	MADD	MADDU	γ	γ	γ	γ	γ	γ	γ
1	γ	γ	γ	γ	γ	γ	γ	γ	γ
2	γ	γ	γ	γ	γ	γ	γ	γ	γ
3	γ	γ	γ	γ	γ	γ	γ	γ	γ
4	γ	γ	γ	γ	γ	γ	γ	γ	γ
5	γ	γ	γ	γ	γ	γ	γ	γ	γ
6	γ	γ	γ	γ	γ	γ	γ	γ	γ
7	γ	γ	γ	γ	γ	γ	γ	γ	γ

Key :

- *: This opcode is reserved for future use. An attempt to execute it causes a Reserved Instruction exception.
- γ: This opcode is reserved for future use. An attempt to execute it causes a Reserved Instruction exception.
- λ: This opcode indicates an instruction class. The instruction word must be further decoded by examining additional tables that show the values for another instruction field.
- α: This opcode is a coprocessor operation, not a CPU operation. If the processor state does not allow access to the specified coprocessor, the instruction causes a Coprocessor Unusable exception. It is included in the table because it uses a primary opcode in the instruction encoding map.
- φ: This opcode is reserved for future use, but does not cause a Reserved Instruction exception in TX49 implementations. It is treated as "NOP".
- θ: This opcode is valid when BC is only selected in COPz rs; In other case, it causes a Reserved Instruction exception .
- ε: This opcode is valid when the processor is operating either in the Kernel mode or in the 64-bit non-Kernel (User or Supervisor) mode; In other case, it causes a Reserved Instruction exception .

Appendix B: FPU Instruction Set Details

This appendix provides a detailed description of the operation of each Floating-Point (FPU) instruction. The instructions are listed alphabetically. The exceptions that may occur due to the execution of each instruction are listed after the description of each instruction. The description of the immediate causes and the manner of handling exceptions is omitted from the instruction descriptions in this chapter. Refer to Chapter 6 for detailed descriptions of floating-point exceptions and handling.

Figure B-5 lists the entire bit encoding for the constant fields of the Floating-Point instruction set; the bit encoding for each instruction is included with that individual instruction.

Instruction Formats

There are three basic instruction format types:

- I-Type, or Immediate instructions, which include load and store operations,
- M-Type, or Move instructions, and
- R-Type, or Register instructions, which include the two- and three-register Floating-Point operations.

The instruction description subsections that follow show how the three basic instruction formats are used by:

- Load and store instructions,
- Move instructions, and
- Floating-Point Computational instructions.

A fourth instruction description subsection describes the special instruction format used by floating-point branch instructions.

Floating-point instructions are mapped onto the MIPS coprocessor instructions, defining coprocessor unit number one (CP1) as the floating-point unit.

Each operation is valid only for certain formats. Implementations may support some of these formats and operations only through emulation, but only need support combinations that are valid, which are marked with a V in Table B-1 below. Those combinations marked with a “R” are not currently specified by this architecture, causing an unimplemented instruction trap, to maintain compatibility with future architecture extensions.

Table B-1 Valid FPU Instruction Formats

Operation	Source Format			
	Single	Double	Word	Longword
ADD	V	V	R	R
SUB	V	V	R	R
MUL	V	V	R	R
DIV	V	V	R	R
SQRT	V	V	R	R
ABS	V	V	R	R
MOV	V	V		
NEG	V	V	R	R
TRUNC.L	V	V		
ROUND.L	V	V		
CEIL.L	V	V		
LOOR.L	V	V		
TRUNC.W	V	V		
ROUND.W	V	V		
CEIL.W	V	V		
FLOOR.W	V	V		
CVT.S		V	V	V
CVT.D	V		V	V
CVT.W	V	V		
CVT.L	V	V		
C	V	V	R	R

The coprocessor branch on condition true/false instructions can be used to logically negate any predicate. Thus, the 32 possible conditions require only 16 distinct comparisons, as shown in Table B-2 below.

Table B-2 Logical Negation of Predicates by Condition True/False

Condition		Code	Relations				Invalid Operation exception if unordered
Mnemonic			Greater Than	Less Than	Equal	Unordered	
True	False						
F	T	0	F	F	F	F	No
UN	OR	1	F	F	F	T	No
EQ	NEQ	2	F	F	T	F	No
UEQ	OGL	3	F	F	T	T	No
OLT	UGE	4	F	T	F	F	No
ULT	OGE	5	F	T	F	T	No
OLE	UGT	6	F	T	T	F	No
ULE	OGT	7	F	T	T	T	No
SF	ST	8	F	F	F	F	Yes
NGLE	GLE	9	F	F	F	T	Yes
SEQ	SNE	10	F	F	T	F	Yes
NGL	GL	11	F	F	T	T	Yes
LT	NLT	12	F	T	F	F	Yes
NGE	GE	13	F	T	F	T	Yes
LE	NLE	14	F	T	T	F	Yes
NGT	GT	15	F	T	T	T	Yes

Floating-Point Loads, Stores, and Moves

All movement of data between the floating-point coprocessor and memory is accomplished by coprocessor load and store operations, which reference the floating-point coprocessor's *General-Purpose Registers*. These operations are unformatted; no format conversions are performed and, therefore, no floating-point exceptions occur due to these operations.

Data may also be directly moved between the floating-point coprocessor and the processor by move to coprocessor and move from coprocessor instructions. Like the floating-point load and store operations, move to/from operations perform no format conversions and never cause floating-point exceptions.

An additional pair of coprocessor registers are available, called *Floating-Point Control* registers for which the only data movement operations supported are moves to and from processor *General-Purpose Registers*.

Floating-Point Operations

The floating-point unit's operation set includes floating-point add, subtract, multiply, divide, square root, convert between fixed-point and floating-point format, convert between floating-point formats, and floating-point compare. These operations satisfy IEEE Standard 754's requirements for accuracy. Specifically, these operations obtain a result which is identical to performing the result with infinite precision and then rounding to the specified format, using the current rounding mode.

Instructions must specify the format of their operands. Except for conversion functions, mixed-format operations are not provided.

Instruction Notational Conventions

In this appendix, all variable sub fields in an instruction format (such as *fs*, *ft*, immediate, and so on) are shown with lower-case names. The instruction name (such as ADD, SUB, and so on) is shown in upper-case.

For the sake of clarity, an alias is sometimes substituted for a variable subfield in the formats of specific instructions. For example, we use *rs = base* in the format for load and store instructions. Such an alias is always lower case, since it refers to a variable subfield.

In some instructions, however, the two instruction subfields *op* and *function* have constant 6-bit values. When reference is made to these instructions, upper-case mnemonics are used. In the floating-point instruction, for example, we use *op = COP1* and *function = FADD*. In some cases, a single field has both fixed and variable subfields, so the name contains both upper and lower case characters. Actual bit encoding for mnemonics is shown in Figure B-5 at the end of this appendix, and are also included with each individual instruction.

In the instruction description examples that follow, the *Operation* section describes the operation performed by each instruction using a high-level language notation.

Instruction Notation Examples

Example #1:

$$\boxed{\text{GPR}[ft] \leftarrow \text{immediate} || 0^{16}}$$

Sixteen zero bits are concatenated with an immediate value (typically 16-bit), and the 32-bit string (with the lower 16-bit set to zero) is assigned to GPR register *ft*.

Example #2:

$$\boxed{(\text{immediate}_{15})^{16} || \text{immediate}_{15-0}}$$

Bit 15 (the sign bit) of an immediate value is extended for 16-bit positions, and the result is concatenated with bits 15 through 0 of the immediate value to form a 32-bit sign-extended value.

Load and Store Instructions

In the MIPS ISA, all load operations have a delay of at least one instruction. That is, the instruction immediately following a load cannot use the contents of the register that will be loaded with the data being fetched from storage.

In the TX49, the instruction immediately following a load may use the contents of the register loaded. In such cases, the hardware will interlock, requiring additional real cycles, so scheduling load delay slots is still desirable, although not absolutely required for functional code.

When the FR bit in the *Status* register equals zero, the *Floating-Point General Registers (FGR)* are 32-bits wide. When the FR bit in the *Status* register equals one, the *Floating-Point General Registers (FGR)* are 64-bits wide. The behavior of the load store instructions is dependent on the width of the *FGRs*.

In the load/store operation descriptions, the functions listed in Table B-3 are used to summarize the handling of virtual addresses and physical memory.

Table B-3 Load/Store Common Functions

Function	Meaning
AddressTranslation	Uses the TLB to find the physical address given the virtual address. The function fails and an exception is taken if the required translation is not present in the TLB.
LoadMemory	Uses the cache and main memory to find the contents of the word containing the specified physical address. The low-order two bits of the address and the access type field indicates which of each of the four bytes within the data word need to be returned. If the cache is enabled for this access, the entire word is returned and loaded into the cache.
StoreMemory	Uses the cache, write buffer and main memory to store the word or part of word specified as data in the word containing the specified physical address. The low-order two bits of the address and the access type field indicates which of each of the four bytes within the data word should be stored.

Figure B-1 shows the I-Type instruction format used by load and store operations.

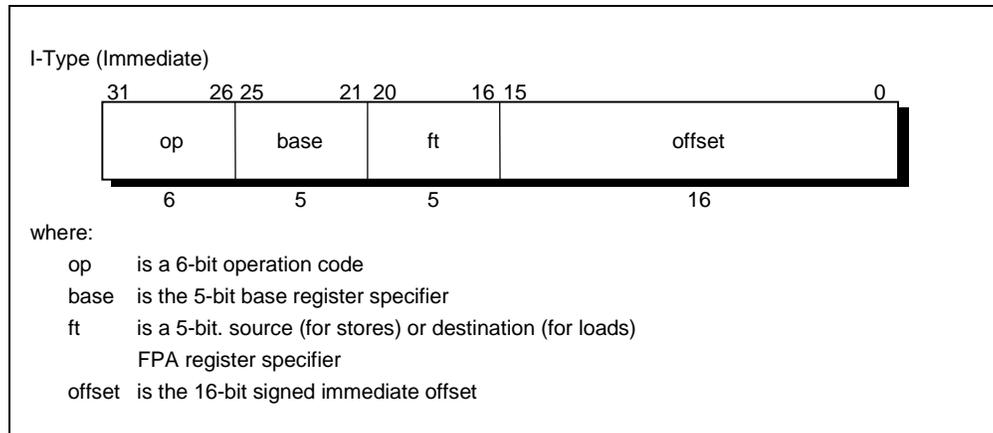


Figure B-1 Load and Store Instruction Format

All coprocessor loads and stores reference aligned word data items. Thus, for word loads and stores, the access type field is always WORD, and the low-order two bits of the address must always be zero.

For double word loads and stores, the access type field is always DOUBLEWORD, and the low-order three bits of the address must always be zero.

Regardless of byte-numbering order (endianness), the address specifies that byte which has the smallest byte-address of all of the bytes in the addressed field. For a Big-endian machine, this is the leftmost byte; for a Little-endian machine, this is the rightmost byte.

Computational Instructions

Computational instructions include all of the arithmetic floating-point operations performed by the FPU.

Figure B-2 shows the R-Type instruction format used for computational operations.

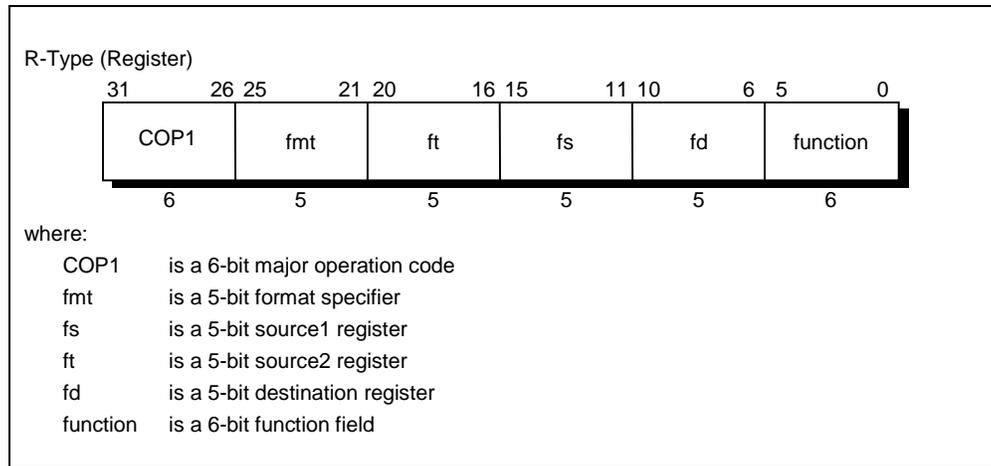


Figure B-2 Computational Instruction Format

Each floating-point instruction can be applied to a number of operand formats. The operand format for an instruction is specified by the 4-bit *Format* field; decoding for this field is shown in Table B-4.

Table B-4 Format Field Decoding

Code	Mnemonic	Size	Format
16	S	single	Binary floating-point
17	D	double	Binary floating-point
18	Reserved		
19	Reserved		
20	W	single	Binary fixed-point
21	L	longword	64-bit binary fixed-point
22-31	-	-	Reserved

The *function* indicates which floating-point operation is to be performed. Table B-5 lists all floating-point instructions.

Table B-5 Floating-Point Instructions and Operations

Code (5-0)	Mnemonic	Operation
0	ADD	Add
1	SUB	Subtract
2	MUL	multiply
3	DIV	Divide
4	SQRT	Square root
5	ABS	Absolute value
6	MOV	Move
7	NEG	Negate
8	ROUND.L	Convert to single fixed-point, rounded to nearest/even
9	TRUNC.L	Convert to single fixed-point, rounded toward zero
10	CEIL.L	Convert to single fixed-point, rounded to $+\infty$
11	FLOOR.L	Convert to single fixed-point, rounded to $-\infty$
12	ROUND.W	Convert to single fixed-point, rounded to nearest/even
13	TRUNC.W	Convert to single fixed-point, rounded toward zero
14	CEIL.W	Convert to single fixed-point, rounded to $+\infty$
15	FLOOR.W	Convert to single fixed-point, rounded to $-\infty$
16-31	-	Reserved
32	CVT.S	Convert to single floating-point
33	CVT.D	Convert to double floating-point
34	-	Reserved
35	-	Reserved
36	CVT.W	Convert to binary fixed-point
37	CVT.L	Convert to 64-bit binary fixed-point
38-47	-	Reserved
48-63	C	Floating-point compare

In the following pages, the notation FGR refers to the FPU's 32 General-Purpose Registers FGRO through FGR31, and FPR refers to the FPU's Floating-Point Registers. When the FR bit in the *Status* register (SR26) equals zero, only the even Floating-Point Registers are valid and the FPU's 32 General-Purpose Registers are 32-bits wide. When the FR bit in the *Status* register (SR26) equals one, both odd and even Floating-Point Registers may be used and the FPU's 32 General-Purpose Registers are 64-bits wide.

The following routines are used in the description of the floating-point operations to get the value of an FPR or to change the value of an FGR:

```

32• Bit Mode
value < - - ValueFPR (fpr, fmt)
/* undefined for odd fpr */
case fmt of
S, W:
    value < - - FGR[fpr + 0]
D:
    /* undefined for fpr not even */
    value < - - FGR[fpr + 1] || FGR[fpr + 0]
end

StoreFPR (fpr, fmt, value):
/* undefined for odd fpr */
case fmt of
S, W:
    FGR[fpr + 1] < - - undefined
    FGR[fpr + 0] < - - value
D:
    FGR[fpr + 1] < - - value63-32
    FGR[fpr + 0] < - - value31-0
end

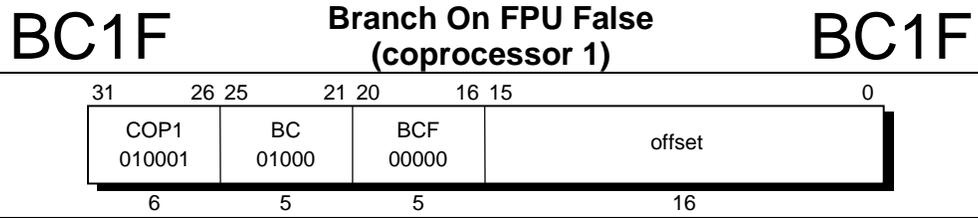
```

```

64• Bit Mode
value < - - ValueFPR (fpr, fmt)
case fmt of
S:
    value < - - FGR[fpr]31-0
D, L:
    value < - - FGR[fpr]
W:
    value < - - FGR[fpr]
end

StoreFPR (fpr, fmt, value):
case fmt of
S, W:
    FGR[fpr] < - - undefined32 || value
D, L:
    FGR[fpr] < - - value
end

```


**Format:**

BC1F offset

Description:

A branch target address is computed from the sum of the address of the instruction in the delay slot and the 16-bit *offset*, shifted left two bits and sign-extended. If the result of the last floating-point compare is false(zero), the program branches to the target address, with a delay of one instruction. There must be at least one instruction between C.cond. fmt and BC1F.

Operation:

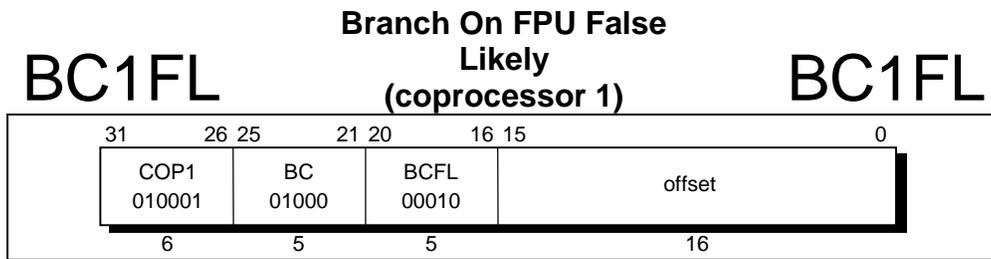
```

32  T• → 1: condition ← not COC[1]
      T:      target ← (offset15)14 || offset || 02
      T• +• 1: if condition then
                  PC ← PC • +• target
                  endif
64  T• → 1 condition ← not COC[1]
      T:      target ← (offset15)46 || offset || 02
      T• +• 1: if condition then
                  PC ← PC • +• target
                  endif

```

Exceptions:

Coprocessor unusable exception

**Format:**

BC1FL offset

Description:

A branch target address is computed from the sum of the address of the instruction in the delay slot and the 16-bit *offset*, shifted left two bits and sign-extended.

If the result of the last floating-point compare is false(zero), the program branches to the target address, with a delay of one instruction. If the conditional branch is *not* taken, the instruction in the branch delay slot is nullified. There must be at least one instruction between C.cond. fmt and BC1FL.

Operation:

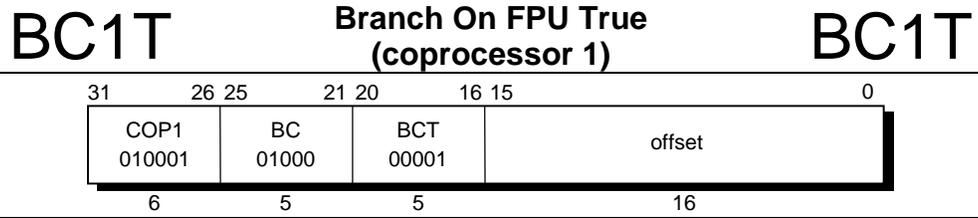
```

32 T• → 1: condition ← not COC[1]
   T:      target ← (offset15)14 || offset || 02
   T + 1:  if condition then
           PC ← PC + target
           Else
           NullifyCurrentInstruction
           Endif
64 T• → 1: condition ← not COC[1]
   T:      target ← (offset15)46 || offset || 02
   T + 1:  if condition then
           PC ← PC + target
           Else
           NullifyCurrentInstruction
           endif

```

Exceptions:

Coprocessor unusable exception

**Format:**

BC1T offset

Description:

A branch target address is computed from the sum of the address of the instruction in the delay slot and the 16-bit *offset*, shifted left two bits and sign-extended. If the result of the last floating-point compare is true(one), the program branches to the target address, with a delay of one instruction. There must be at least one instruction between C.cond. fmt and BC1T.

Operation:

```

32  T• → 1: condition ← COC[1]
      T:      target ← (offset15)14 || offset || 02
      T + 1:  if condition then
                PC ← PC + target
            endif

64  T• → 1: condition ← COC[1]
      T:      target ← (offset15)46 || offset || 02
      T + 1:  if condition then
                PC ← PC + target
            endif

```

Exceptions:

Coprocessor unusable exception

C.cond.fmt Floating-Point Compare (continued) C.cond.fmt

Operation:

```

32, 64  T:  if NaN (ValueFPR(is, fmt)) or NaN (ValueFPR(it, fmt)) then
           less ← false
           equal ← false
           unordered ← true
           if cond3 then
               signal InvalidOperationException
           endif
           else
               less ← ValueFPR (fs, fmt) < ValueFPR (It, fmt)
               equal ← ValueFPR (fs, fmt) = ValueFPR (it, fmt)
               unordered ← false
           endif
           condition ← (cond2 and less) or (cond1 and equal) or
                       (cond0 and unordered)
           FCR[31]23 ← condition
           COC[1] ← condition

```

Exceptions:

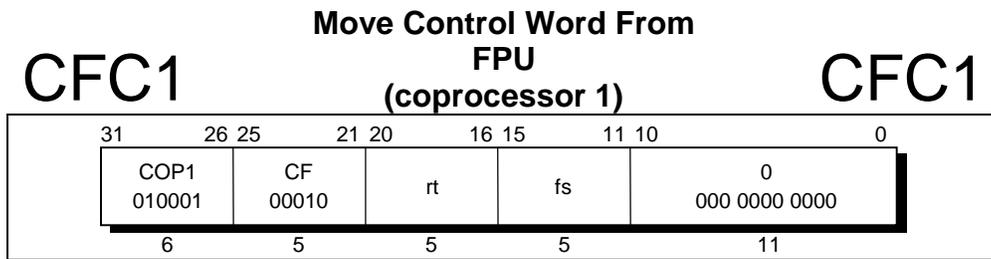
Coprocessor unusable

Floating-Point exception

Coprocessor Exceptions:

Unimplemented operation exception

Invalid operation exception

**Format:**CFC1 *rt*, *fs***Description:**

The contents of the FPU's control register *fs* are loaded into general register *rt*.

This operation is only defined when *fs* equals 0 or 31.

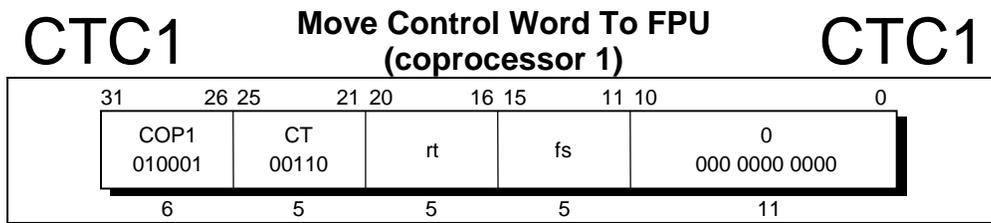
The contents of general register *rt* are undefined for the instruction immediately following CFC1.

Operation:

32	T:	$\text{temp} \leftarrow \text{FCR}[\text{fs}]$
	T• +• 1:	$\text{GPR}[\text{rt}] \leftarrow \text{temp}$
64	T:	$\text{temp} \leftarrow \text{FCR}[\text{fs}]$
	T• +• 1:	$\text{GPR}[\text{rt}] \leftarrow (\text{temp}_{31})^{32} \parallel \text{temp}$

Exceptions:

Coprocessor unusable exception

**Format:**

CTC1 rt, fs

Description:

The contents of general register *rt* are loaded into the FPU's control register *fs*. This operation is only defined when *fs* equals 0 or 31. Writing to *Control Register 31*, the floating-point *Control/Status* register, causes an interrupt or exception if any cause bit and its corresponding enable bit are both set. The register will be written before the exception occurs. The contents of floating-point control register *fs* are undefined for the instruction immediately following CTC1.

Operation:

<pre> 32 T: temp ← GPR[rt] T• +• 1: FCR[fs] ← temp COC[1] ← FCR[31]₂₃ 64 T: temp ← GPR[rt]_{31~0} T• +• 1: FCR[fs] ← temp COC[1] ← FCR[31]₂₃ </pre>

Exceptions:

Coprocessor unusable exception

Floating-Point exception

Coprocessor Exceptions:

Unimplemented operation exception

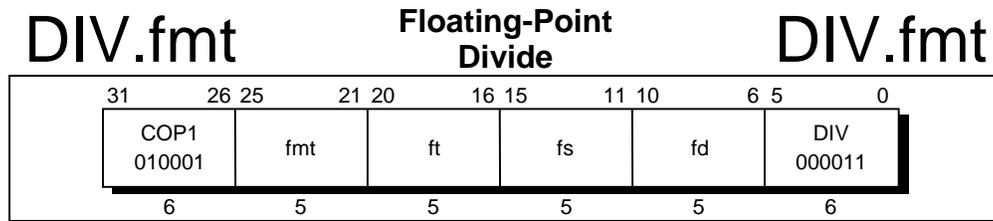
Invalid operation exception

Division by zero exception

Inexact exception

Overflow exception

Underflow exception

**Format:**

DIV.fmt fd, fs, ft

Description:

The contents of the floating-point registers specified by *fs* and *ft* are interpreted in the specified *format* and the value in *fs* is divided by the value in *ft*. The result is rounded as if calculated to infinite precision and then rounded to the specified format, according to the current rounding mode. The result is placed in the floating-point register specified by *fd*.

This instruction is valid for only single or double precision floating-point formats.

The operation is not defined if bit 0 of any register specification is set and the *FR* bit in the *Status* register equals zero, since the register numbers specify an even-odd pair of adjacent coprocessor general registers. When the *FR* bit in the *Status* register equals one, both even and odd register numbers are valid.

Operation:

32, 64 T: StoreFPR (fd, fmt, ValueFPR(fs, fmt)/ValueFPR(ft, fmt))

Exceptions:

Coprocessor unusable exception

Floating-Point exception

Coprocessor Exceptions:

Unimplemented operation exception

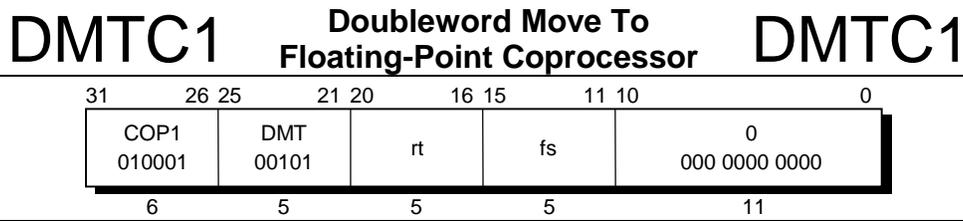
Invalid operation exception

Division-by-zero exception

Inexact exception

Overflow exception

Underflow exception

**Format:**DMTC1 *rt*, *fs***Description:**

The contents of general register *rt* are loaded into coprocessor register *fs* of the CP1.

The contents of floating-point register *fs* are undefined for the instruction immediately following DMTC1.

The *FR* bit in the *Status* register specifies whether all 32 register of the TX49 are addressable. When *FR* equals zero, this instruction is not defined when the least significant bit of *fs* is non-zero. When *FR* equals one, *fs* may specify either odd or even registers.

Operation:

```

64 T:    data ← GPR[rt]
T• +• 1: if SR26 = 1 then /*64-bit wide FGRs*/
          FGR[fs] ← data
          elseif fs0 = 0 then /*valid specifier, 32-bit wide valid FGRs*/
          FGR[fs• +• 1] ← data63-32
          FGR[fs] ← data31-0
          else /*undefined result for odd 32-bit reg #s */
          undefined_result
          endif

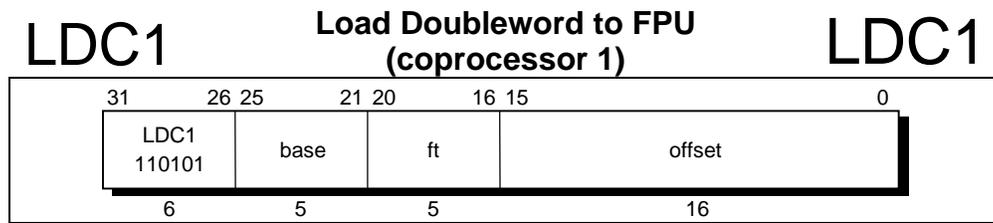
```

Exceptions:

Coprocessor unusable exception

Coprocessor Exceptions:

Unimplemented operation exception

**Format:**

LDC1 ft, offset (base)

Description:

The 16-bit *offset* is sign-extended and added to the contents of general register *base* to form an unsigned effective address. In 32-bit mode, the contents of the doubleword at the memory location specified by the effective address is loaded into registers *ft* and *ft + 1* of the floating-point coprocessor. This instruction is not valid, and is undefined, when the least significant bit of *ft* is non-zero. In 64-bit mode, the contents of the doubleword at the memory location specified by the effective address are loaded into the 64-bit register *ft* of the floating point coprocessor. The *FR* bit of the *Status* register (SR26) specifies whether all 32 registers of the TX49 are addressable. When *FR* = 0, this instruction is not defined when the least significant bit of *ft* is non-zero. When *FR* = 1, *ft* may specify either odd or even registers.

If any of the three least-significant bits of the effective address are non-zero, an address error exception takes place.

LDC1 Load Doubleword to FPU (coprocessor 1) (continued) LDC1

Operation:

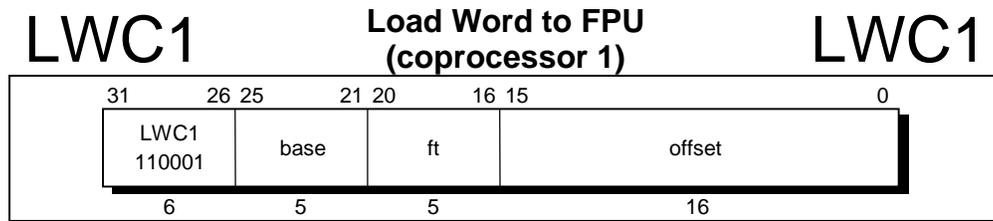
```

32  T:    vAddr ← ((offset15)16 || offset15-0) + GPR[base]
          (pAddr, uncached) ← Address Translation (vAddr, DATA)
          data ← LoadMemory (uncached, DLUBLEWORD, pAddr, vAddr, DATA)
          if SR26 = 1 then /*64-bit wide GFRs */
              FGR[ft] ← data
          elseif ft0 = 0 then /*valid specifier, 32-bit wide GFRs */
              FGR[ft + 1] ← data63-32
              FGR[ft] ← data31-0
          else /*undefined result if odd */
              undefined_result
          endif
64  T:    vAddr ← ((offset15)48 || offset15-0) + GPR[base]
          (pAddr, uncached) ← Address Translation (vAddr, DATA)
          data ← LoadMemory (uncached, DLUBLEWORD, pAddr, vAddr, DATA)
          if SR26 = 1 then /*64-bit wide GFRs */
              FGR[ft] ← data
          elseif ft0 = 0 then /*valid specifier, 32-bit wide GFRs */
              FGR[ft + 1] ← data63-32
              FGR[ft] ← data31-0
          else /*undefined result if odd */
              undefined_result
          endif

```

Exceptions:

- Coprocessor unusable**
- TLB refill exception**
- TLB invalid exception**
- Bus error exception**
- Address error exception**

**Format:**

LWC1 ft, offset (base)

Description:

The 16-bit *offset* is sign-extended and added to the contents of general register *base* to form an unsigned effective address. The contents of the word at the memory location specified by the effective address is loaded into register *ft* of the floating-point coprocessor.

The *FR* bit of the *Status* register specifies whether all 64-bit *Floating-Point Registers* are addressable. If *FR* equals zero, LWC1 loads either the high or low half of the 16 even *Floating-Point Registers*. If *FR* equals one, LWC1 loads the low 32-bits of both even and odd *Floating-Point Registers*.

If either of the two least-significant bits of the effective address is non-zero, an address error exception occurs.

LWC1 Load Word to FPU (coprocessor 1) (continued) LWC1

Operation:

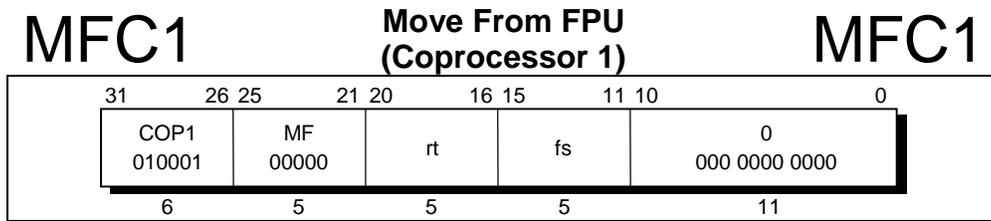
```

32  T:      vAddr ← ((offset15)16 || offset15-0) + GPR[base]
           (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
           pAddr ← pAddrPSIZE-1-3 || (pAddr2-0xor(ReverseEndian || 02))
           mem ← LoadMemory(uncached, WORD, pAddr, vAddr, DATA)
           byte ← vAddr2-0xor(BigEndianCPU || 02)
           /*"mem" is aligned 64-bits from memory. Pick out correct bytes. */
           if SR26 = 1 then /*64-bit wide FRGs */
               FGR[ft] ← undefined32 || mem31 + 8*byte-8*byte
           else /*32-bit wide FRGs */
               FGR[rf] ← mem31 + 8*byte-8*byte
           endif
64  T:      vAddr ← ((offset15)48 || offset15-0) + GPR[base]
           (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
           pAddr ← pAddrPSIZE-1-3 || (pAddr2-0xor(ReverseEndian || 02))
           mem ← LoadMemory(uncached, WORD, pAddr, vAddr, DATA)
           byte ← vAddr2-0xor(BigEndianCPU || 02)
           /*"mem" is aligned 64-bits from memory. Pick out correct bytes. */
           if SR26 = 1 then /*64-bit wide FRGs */
               FGR[ft] ← undefined32 || mem31 + 8*byte-8*byte
           else /*32-bit wide FRGs */
               FGR[rf] ← mem31 + 8*byte-8*byte
           endif

```

Exceptions:

- Coprocessor unusable**
- TLB-refill exception**
- TLB invalid exception**
- Bus error exception**
- Address error exception**

**Format:**

MFC1 rt, fs

Description:

The contents of register *fs* from the floating-point coprocessor are stored into processor register *rt*.

The contents of register *rt* are undefined for time *T* of the instruction immediately following this load instruction.

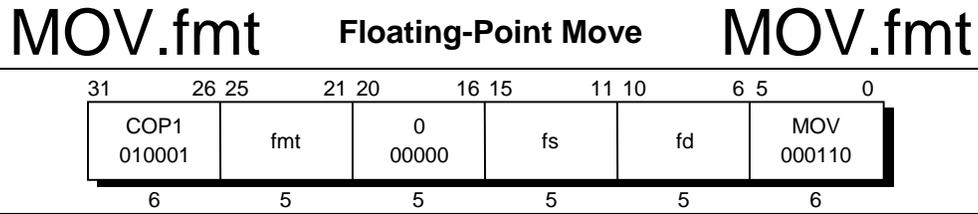
The *FR* bit of the Status register specifies whether all 32 registers of the TX49 are addressable. If *FR* equals zero, MFC1 stores either the high or low half of the 16 even *Floating-Point Registers*. If *FR* equals one, MFC1 stores the low 32-bits of both even and odd *Floating-Point Registers*.

Operation:

<pre> 32 T: data ← FGR[fs]₃₁₋₀ T•+• 1: GPR[rt] ← data 64 T: data ← FGR[fs]₃₁₋₀ T + 1: GPR[rt] ← (data₃₁)³² data </pre>
--

Exceptions:

Coprocessor unusable exception

**Format:**

MOV.fmt fd, fs

Description:

The contents of the FPU register specified by *fs* are interpreted in the specified *format* and are copied into the FPU register specified by *fd*. The move operation is non-arithmetic; no IEEE 754 exceptions occur as a result of the instruction.

This instruction is valid only for single- or double-precision floating-point formats.

The operation is not defined if bit 0 of any register specification is set and the *FR* bit in the *Status* register equals zero, since the register numbers specify an even-odd pair of adjacent coprocessor general registers. When the *FR* bit in the *Status* register equals one, both even and odd register numbers are valid.

Operation:

32, 64 T: StoreFPR (fd, fmt, ValueFPR (fs, fmt))
--

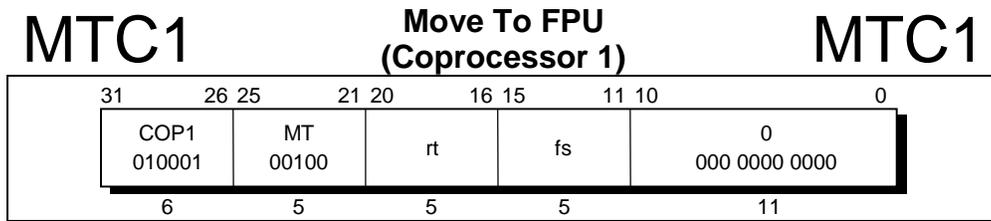
Exceptions:

Coprocessor unusable exception

Floating-Point exception

Coprocessor Exceptions:

Unimplemented operation exception

**Format:**

MTC1 rt, fs

Description:

The contents of register *rt* are loaded into the FPU's general register at location *fs*.

The contents of floating-point register *fs* is undefined for the instruction immediately following MTC1.

The *FR* bit of the *Status* register specifies whether all 32 registers of the TX49 are addressable. If *FR* equals zero, MTC1 loads either the high or low half of the 16 even *Floating-Point Registers*. If *FR* equals one, MTC1 loads the low 32-bits of both even and odd *Floating-Point Registers*.

Operation:

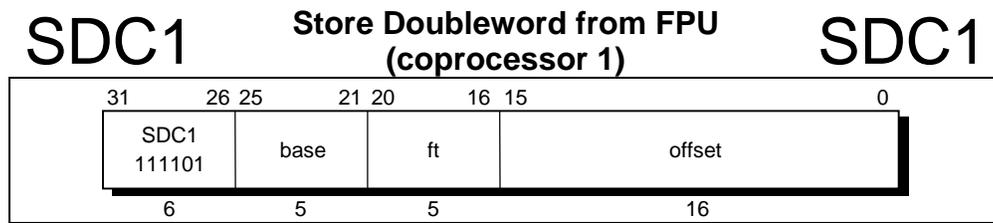
```

32, 64 T:   data ← GPR[rt]31-0
           T + 1: if SR26 = 1 then /* 64-bit wide FGRs */
                   FGR[fs] ← undefined32 || data
           else /* 32-bit wide FGRs */
           endif

```

Exceptions:

Coprocessor unusable exception

**Format:**

SDC1 ft, offset (base)

Description:

The 16-bit *offset* is sign-extended and added to the contents of general register *base* to form an unsigned effective address.

In 32-bit mode, the contents of registers *ft* and *ft + 1* from the floating-point coprocessor are stored at the memory location specified by the effective address. This instruction is not valid, and is undefined, when the least significant bit of *ft* is non-zero.

In 64-bit mode, the 64-bit register *ft* is stored to the contents of the doubleword at the memory location specified by the effective address. The *FR* bit of the *Status* register (SR26) specifies whether all 32 registers of the TX49 are addressable. When $FR = 0$, this instruction is not defined if the least significant bit of *ft* is non-zero. If $FR = 1$, *ft* may specify either odd or even registers.

If any of the three least-significant bits of the effective address are non-zero, an address error exception takes place.

Store Doubleword from FPU (coprocessor 1) (continued)

SDC1 SDC1

Operation:

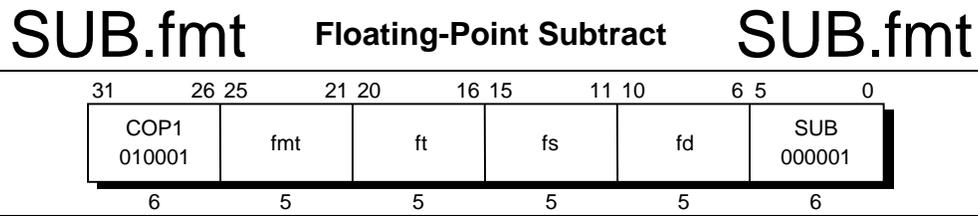
```

32  T:    vAddr ← ((offset15)16 || offset15-0) + GPR[base]
         (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
         if SR26 = 1 /*64-bit wide FGRs */
           data ← FGR[ft]
         elseif ft0 = then /* valid specifier, 32-bit wide FGRs */
           data ← FGR[ft + 1] || FGR[ft]
         else /*undefined for odd 32-bit reg #s */
           data ← undefined64
         endif
         StoreMemory (uncached, DOUBLEWORD, data, pAddr, vAddr, DATA)
64  T:    vAddr ← ((offset15)48 || offset15-0) + GPR[base]
         (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
         if SR26 = 1 /*64-bit wide FGRs */
           data ← FGR[ft]
         elseif ft0 = then /* valid specifier, 32-bit wide FGRs */
           data ← FGR[ft + 1] || FGR[ft]
         else /*undefined for odd 32-bit reg #s */
           data ← undefined64
         endif
         StoreMemory (uncached, DOUBLEWORD, data, pAddr, vAddr, DATA)

```

Exceptions:

- Coprocessor unusable**
- TLB refill exception**
- TLB invalid exception**
- TLB modification exception**
- Bus error exception**
- Address error exception**

**Format:**

SUB.fmt fd,fs, ft

Description:

The contents of the floating-point registers specified by *fs* and *ft* are interpreted in the specified *format* and the value in *ft* is subtracted from the value in *fs*. The result is rounded as if calculated to infinite precision and then rounded to the specified *format*, according to the current rounding mode. The result is placed in the floating-point register specified by *fd*.

This instruction is valid only for single- or double-precision floating-point formats.

The operation is not defined if bit 0 of any register specification is set and the *FR* bit in the *Status* register equals zero, since the register numbers specify an even-odd pair of adjacent coprocessor general registers. When the *FR* bit in the *Status* register equals one, both even and odd register numbers are valid.

Operation:

32, 64 T: StoreFPR (fd, fmt, ValueFPR (fs, fmt) ← ValueFPR (ft, fmt))

Exceptions:

Coprocessor unusable exception

Floating-Point exception

Coprocessor Exceptions:

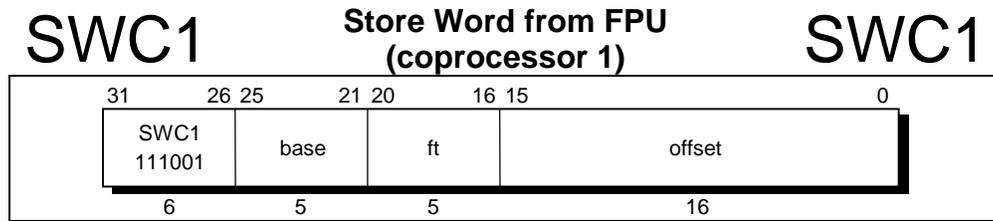
Unimplemented operation exception

Invalid operation exception

Inexact exception

Overflow exception

Underflow exception

**Format:**

SWC1 ft, offset (base)

Description:

The 16-bit *offset* is sign-extended and added to the contents of general register *base* to form an unsigned effective address. The contents of register *ft* from the floating-point coprocessor are stored at the memory location specified by the effective address.

The *FR* bit of the *Status* register specifies whether all 64-bit *Floating-Point Registers* are addressable. If *FR* equals zero, SWC1 stores either the high or low half of the 16 even *Floating-Point Registers*. If *FR* equals one, SWC1 stores the low 32-bits of both even and odd *Floating-Point Registers*.

If either of the two least-significant bits of the effective address are non-zero, an address error exception occurs.

SWC1 Store Word from FPU (coprocessor 1) (continued) SWC1

Operation:

```

32  T:    vAddr ← ((offset15)16 || offset15-0) + GPR[base]
          (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
          pAddr ← pAddrPSIZE-1-3 || (pAddr2-0xor (ReverseEndian || 02))
          byte ← vAddr2-0xor (BigEndianCPU || 02)
          /* the bytes of the word are put in the correct byte lanes in
           * "data" for a 64-bit path to memory */
          if SR26 = 1 then /*64-bit wide FGRs */
              data ← FGR[ft]63-8*byte-0 || 08*byte
          else /* 32-bit wide FGRs */
              data ← 032-8*byte || FGR[ft] || 08*byte
          endif
          StoreMemory (uncached, WORD, data, pAddr, vAddr, DATA)
64  T:    vAddr ← ((offset15)48 || offset15-0) + GPR[base]
          (pAddr, uncached) ← AddressTranslation (vAddr, DATA)
          pAddr ← pAddrPSIZE-1-3 || (pAddr2-0xor (ReverseEndian || 02))
          byte ← vAddr2-0xor (BigEndianCPU || 02)
          /* the bytes of the word are put in the correct byte lanes in
           * "data" for a 64-bit path to memory */
          if SR26 = 1 then /*64-bit wide FGRs */
              data ← FGR[ft]63-8*byte-0 || 08*byte
          else /* 32-bit wide FGRs */
              data ← 032-8*byte || FGR[ft] || 08*byte
          endif
          StoreMemory (uncached, WORD, data, pAddr, vAddr, DATA)

```

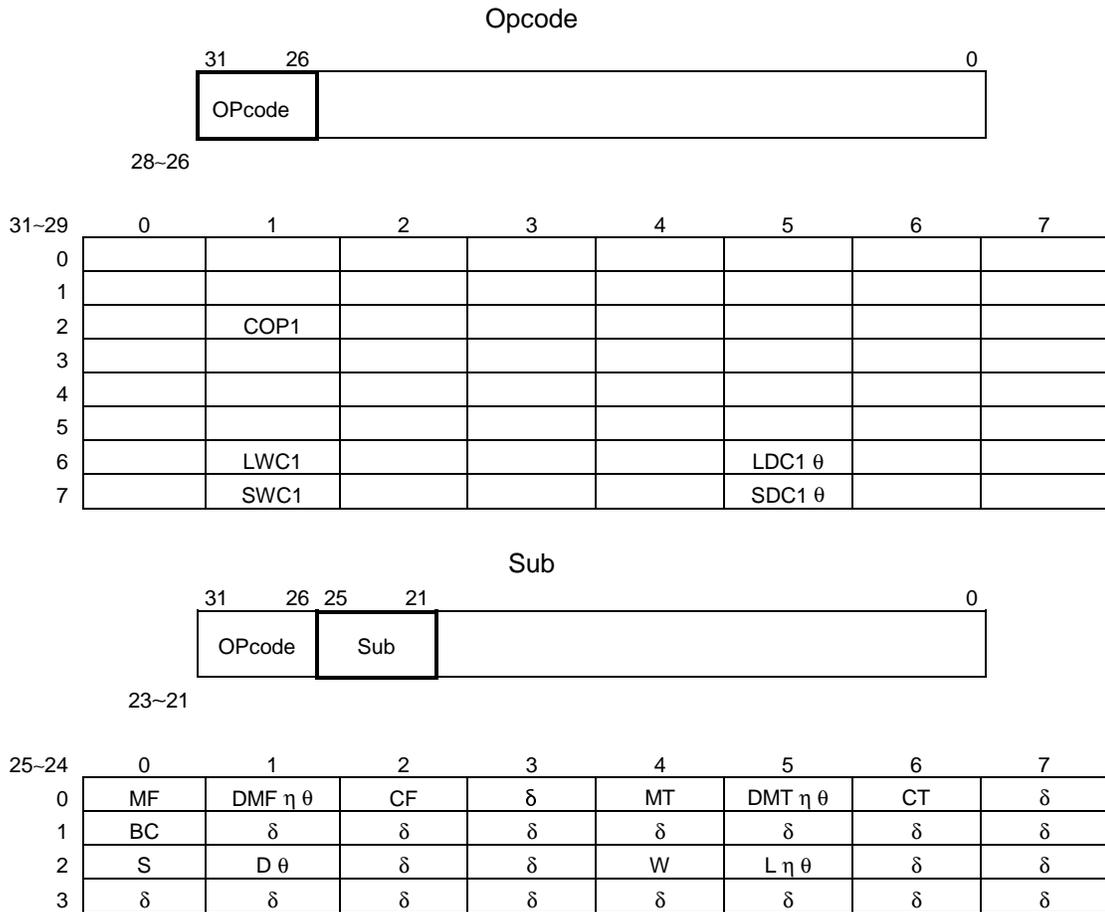
Exceptions:

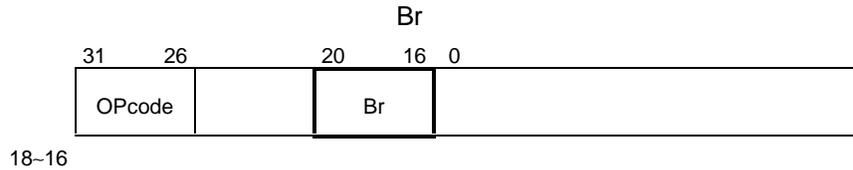
- Coprocessor unusable
- TLB refill exception
- TLB invalid exception
- TLB modification exception
- Bus error exception
- Address error exception

Bit Encoding of FPU Instruction Opcodes

The Table B-5 shows the bit codes for all TX49 FPU instructions (ISA and extended ISA)

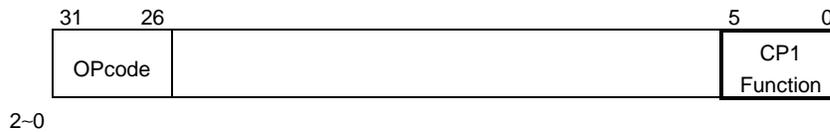
Table B-5 FPU Operation Code Bit Encoding





20-19	0	1	2	3	4	5	6	7
0	BCF	BCT	BCFL	BCTL	γ	γ	γ	γ
1	γ	γ	γ	γ	γ	γ	γ	γ
2	γ	γ	γ	γ	γ	γ	γ	γ
3	γ	γ	γ	γ	γ	γ	γ	γ

CP1 Function



5-3	0	1	2	3	4	5	6	7
0	ADD	SUB	MUL	DIV	SQRT	ABS	MOV	NEG
1	ROUND.L $\eta\theta$	TRUNC.L $\eta\theta$	CEIL.L $\eta\theta$	FLOOR.L $\eta\theta$	ROUND.W	TRUNC.W	CEIL.W	FLOORW
2	δ	δ	δ	δ	δ	δ	δ	δ
3	δ	δ	δ	δ	δ	δ	δ	δ
4	CVT.S	CVT.D θ	δ	δ	CVT.W	CVT.L $\eta\theta$	δ	δ
5	δ	δ	δ	δ	δ	δ	δ	δ
6	C.F	C.UN	C.EQ	C.UEQ	C.OLT	C.ULT	C.OLE	C.ULE
7	C.SF	C.NGLE	C.SEQ	C.NGL	C.LT	C.NGE	C.LE	C.NGT

Key :

- γ: This opcode is reserved for future use. An attempt to execute it causes a Reserved Instruction exception.
- δ: This opcode is reserved for future use. An attempt to execute it causes a Unimplemented operation exceptions in all current implementations.
- η: This opcode is valid only when MIPS III instructions are enabled. An attempt to execute these without MIPS III instruction enabled will cause an Unimplemented operation exception.
- θ: This opcode is valid only when the TX49 has a double precision FPU in hardware. An attempt to execute these without it will cause an Unimplemented operation exception.

Note :

□FPU Instructions are valid only when TX49 has with FPU(CP1). An attempt to execute these instruncions causes a Coprocessor Unusable exception, independent of C0_SR(bit 29)'s value.

Appendix C: Coprocessor 0 Hazards

Pipeline Interlock and Hazard in TX49

Interlock in Load Delay Slot

Pipeline control logic will interlock the pipeline when detecting a hazard condition and pipeline won't resume until the hazard is resolved.

An example is shown in Figure C-1. In this case, instruction in the load delay slot tries to read the destination register of the load instruction resulting in pipeline stall until the data is read from the cache.

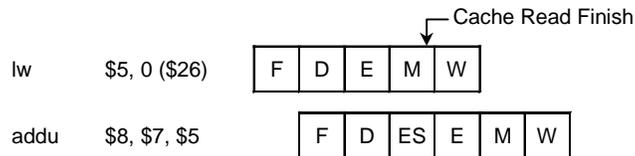


Figure C-1 Interlock in Load Delay Slot

Pipeline also interlocks when the cache miss occurs or when the data is loaded from uncached area (Figure C-2).

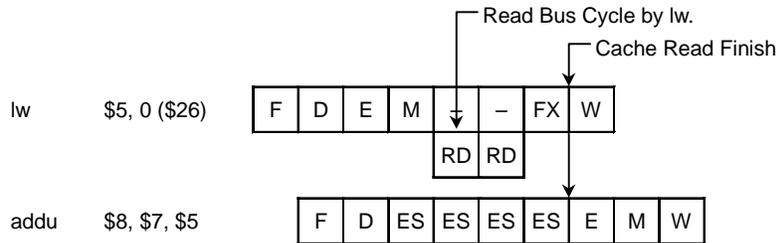


Figure C-2 Interlock in Cache Miss or in the Data Load from Non-cached Area

In this example where there is a register hazard between two consecutive instructions, ADDU will stall at E stage until the destination register of LW is written back.

However, if there is no data dependency between LW and ADDU, execution of ADDU will complete without stall before the destination register of LW is written back. Pipeline interlock occurs at the first instruction that has the data dependency with the preceding load instruction (Figure C-3).

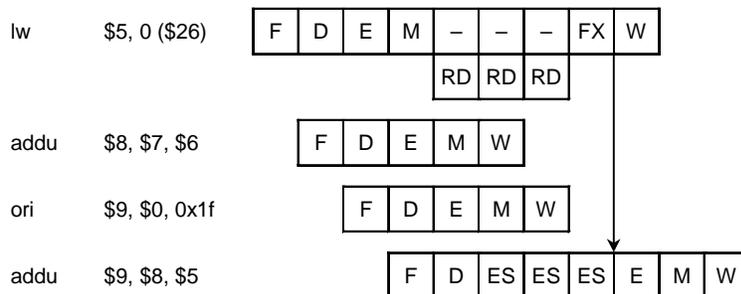


Figure C-3 Pipeline Interlock by Cache Miss

Pipeline also interlocks on write-after-write hazard which is illustrated in Figure C-4. Write-after-write hazard is detected when one of the instructions following a load has the destination register which is same as that of the load instruction. In this example, the ADDU instruction stalls at its E stage until the destination register (\$1) of the load is written back.

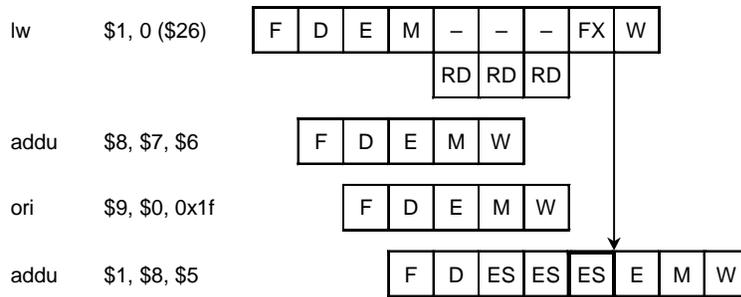


Figure C-4 Write-after-write Hazard by Load Instruction

A SYNC instruction may be placed right after a load instruction. This will cause pipeline stall until the bus cycle issued by the previous load instruction completes (Figure C-5). If the data is read from the cache, there is no bus cycle pending before the SYNC which results in no pipeline stall.

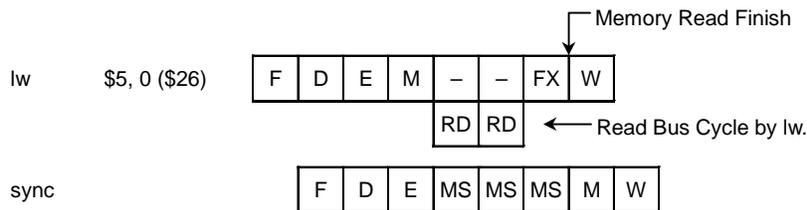


Figure C-5 SYNC Instruction After Load Instruction

Branch Delay Slot

Branch and jump instructions have a branch delay slot (Figure C-6). Also, DERET instruction has a branch delay slot. Note that the result is undefined when the branch/jump instruction is placed in the branch delay slot¹.

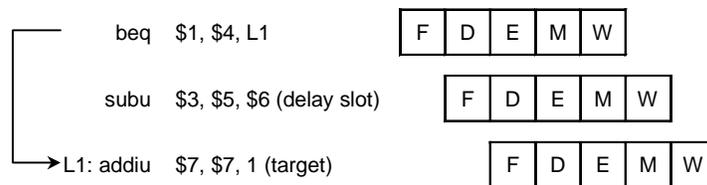


Figure C-6 Branch Delay Slot

¹ Instructions which cause exception, such as, SYSCALL, BREAK, and SDBBP may be placed in the branch delay slot.

Multiply, Multiply/Add and Division Instructions

This subsection explains the pipeline hazard/interlock caused by the combinations of multiply, multiply/add, division, and MTHI/MTLO/MFHI/MFLO instructions (Figure C-7). Basically, the pipeline hazard/interlock by these instructions can be summarized in this way:

- Pipeline interlocks when the data dependency exists.
- Pipeline interlocks when preceding 32-bit multiply or 32-bit multiply/add instruction has <rd> field.
- Pipeline interlocks when 32-bit instruction and 64-bit instruction are executed in sequence.
- HI/LO registers are in undefined state within two instructions before the division instruction, such as, DIV/DIVU/DDIV/DDIVU instruction².

		SUCCEEDING INSTRUCTION									
		MULT/ MULTU (2-operand)	MULT/ MULTU (3-operand)	MADD/ MADDU (2-operand)	MADD/ MADDU (3-operand)	MTHI/ MTLO	MFHI/ MFLO	DIV/ DIVU	DMULT/ DMULTU (2-operand)	DMULT/ DMULTU (3-operand)	DDIV/ DDIVU
PRECEDING INSTRUCTION	MULT/MULTU (2-operand)	NO STALL	NO STALL	NO STALL	NO STALL	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK
	MULT/MULTU (3-operand)	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK
	MADD/MADDU (2-operand)	NO STALL	NO STALL	NO STALL	NO STALL	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK
	MADD/MADDU (3-operand)	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK
	MTHI/MTLO	NO STALL	NO STALL	NO STALL	NO STALL	NO STALL	NO STALL	NO STALL	NO STALL	NO STALL	NO STALL
	MFHI/MFLO	NO STALL	NO STALL	NO STALL	NO STALL	NO STALL	NO STALL	*	NO STALL	NO STALL	*
	DIV/DIVU	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK
	DMULT/DMULTU (2-operand)	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK
	DMULT/DMULTU (3-operand)	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK
	DDIV/DDIVU	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK	INTERLOCK

*: HI/LO registers are in undefined state within two instructions before division instruction

Figure C-7 MAC pipeline hazard/interlock

In the following sections, the pipeline hazards/interlocks caused by the possible combinations of the instructions related multiply, multiply/add, division and both 32-bit and 64-bit operations are illustrated in detail. The Figures in the following sections classifies the cases in such a way that:

- The preceding instruction is immediately followed by 32-bit multiply or multiply/add instruction
- The preceding instruction is immediately followed by MFHI or MFLO instruction
- The preceding instruction is immediately followed by MTHI or MTLO instruction
- The preceding instruction is immediately followed by 32-bit division instruction
- The preceding instruction is immediately followed by 64-bit multiply instruction
- The preceding instruction is immediately followed by 64-bit division instruction

² In the original R3000, this can be applied to MULT, MULTU, MTHI, and MTLO instructions.

Case 1: Preceding Instruction Is 32-bit Multiply or 32-bit Multiply/Add Instruction

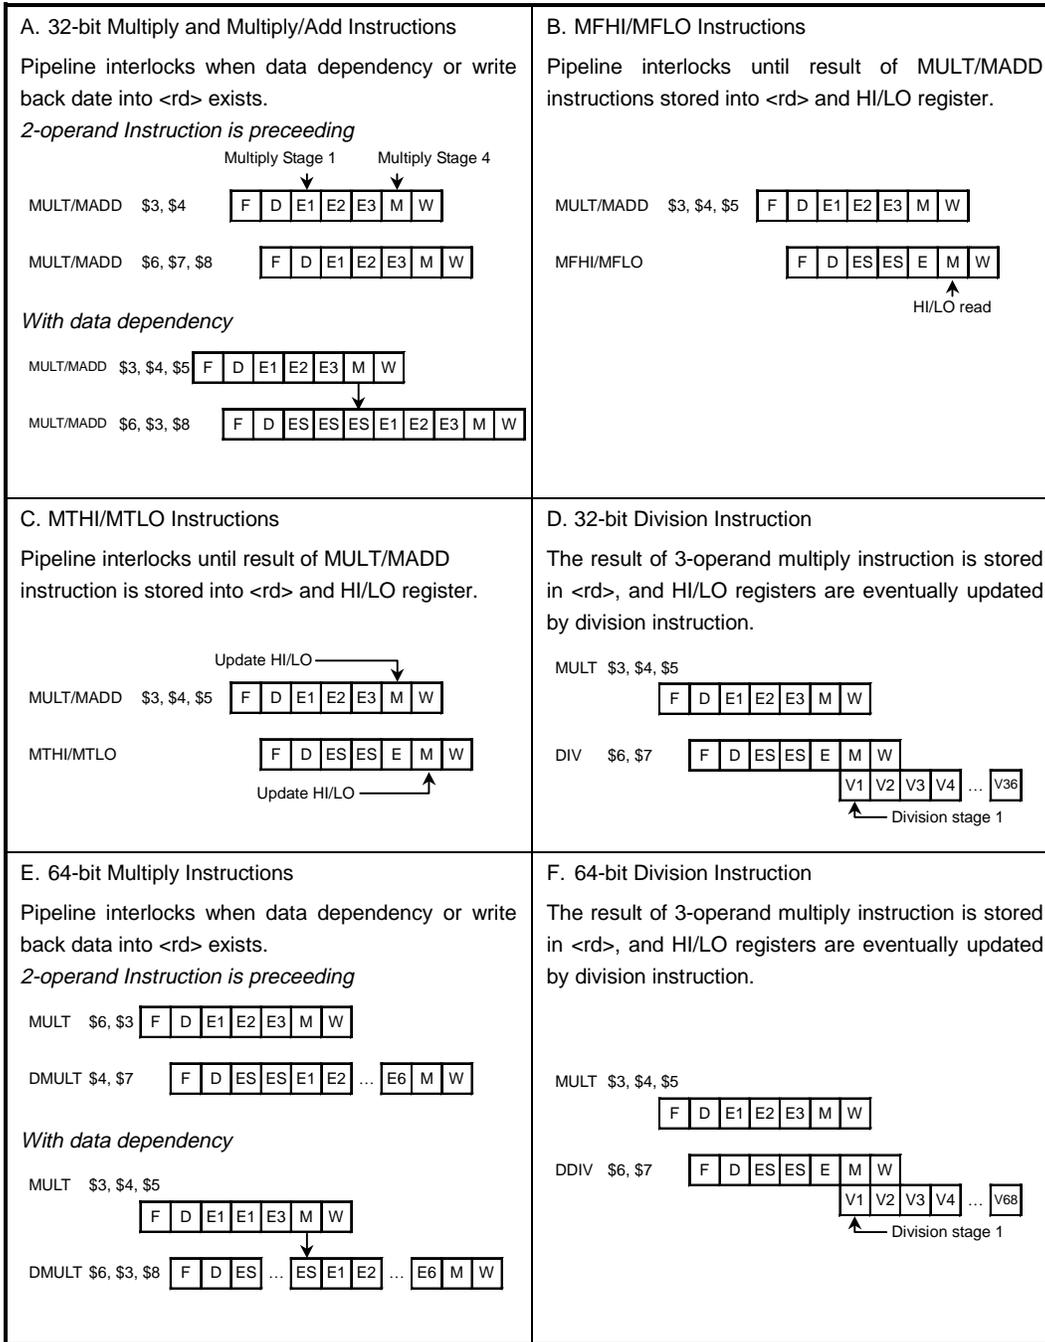


Figure C-8 Pipeline Hazard/Interlock by 32-bit Multiply or 32-bit Multiply/Add Instruction

Note that in the category A of the Figure C-8, pipeline interlocks for any instruction immediately after the multiply or multiply/add instruction when it has the data dependency regarding the general purpose registers. Thus, in the category D, the DIV instruction stalls at the E stage for three cycles when the division instruction has the data dependency with the preceding multiply instruction.

Also note that in the category D of the Figure C-8, Because the division instruction overwrites the HI/LO registers, the HI/LO registers as the result of the 2-operand multiply instruction is undefined. The result of the multiply instruction, as in this figure, is correctly stored in the <rd> register. If the preceding multiply or multiply/add instruction had a <rd> field, pipeline

interlocks due to the resource conflict.

Case 2: Preceding Instruction Is MFHI/MFLO Instruction

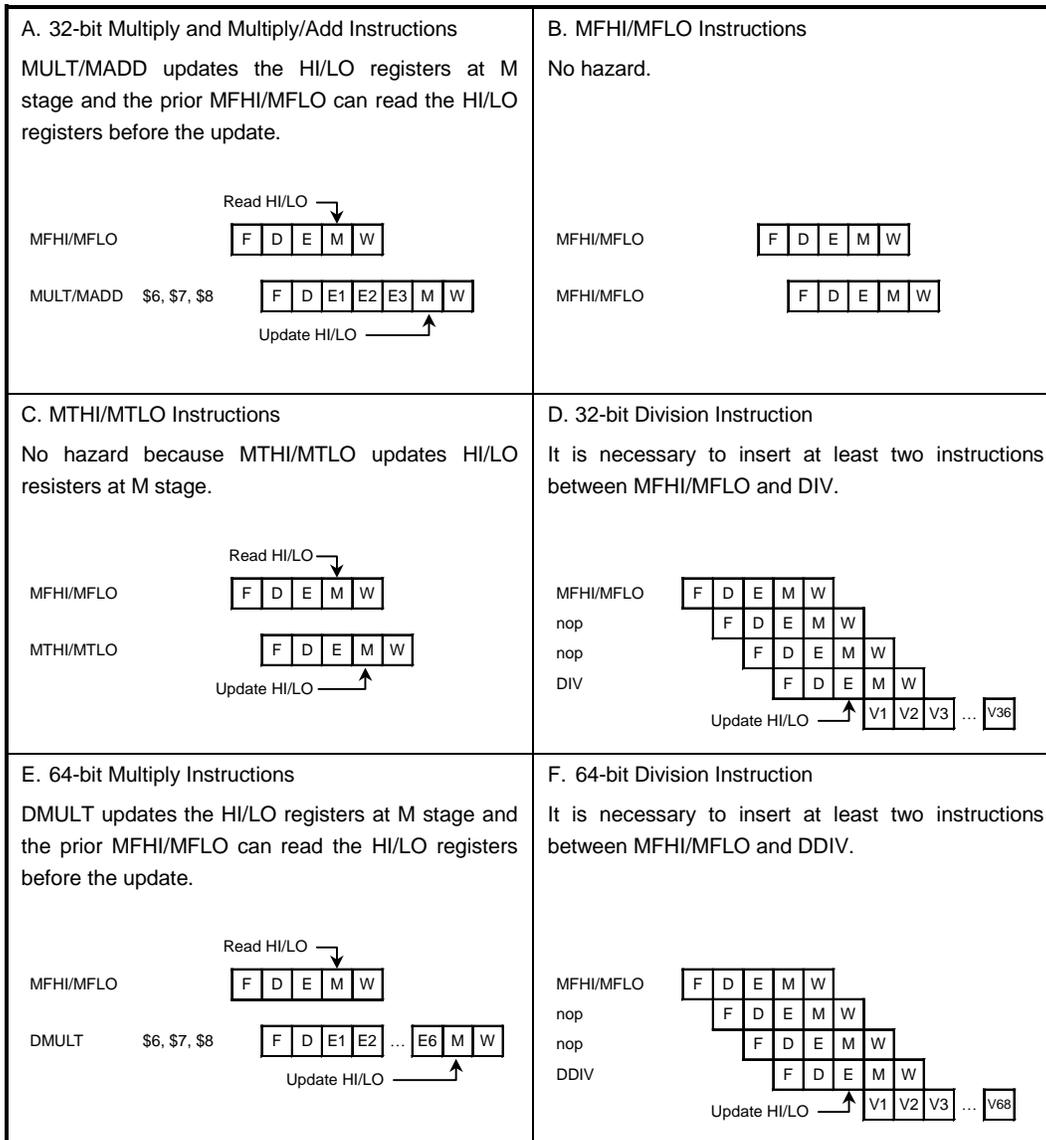


Figure C-9 Pipeline Hazard/Interlock by MFHI/MFLO Instructions

Case3: Preceding Instruction Is MTHI/MTLO Instruction

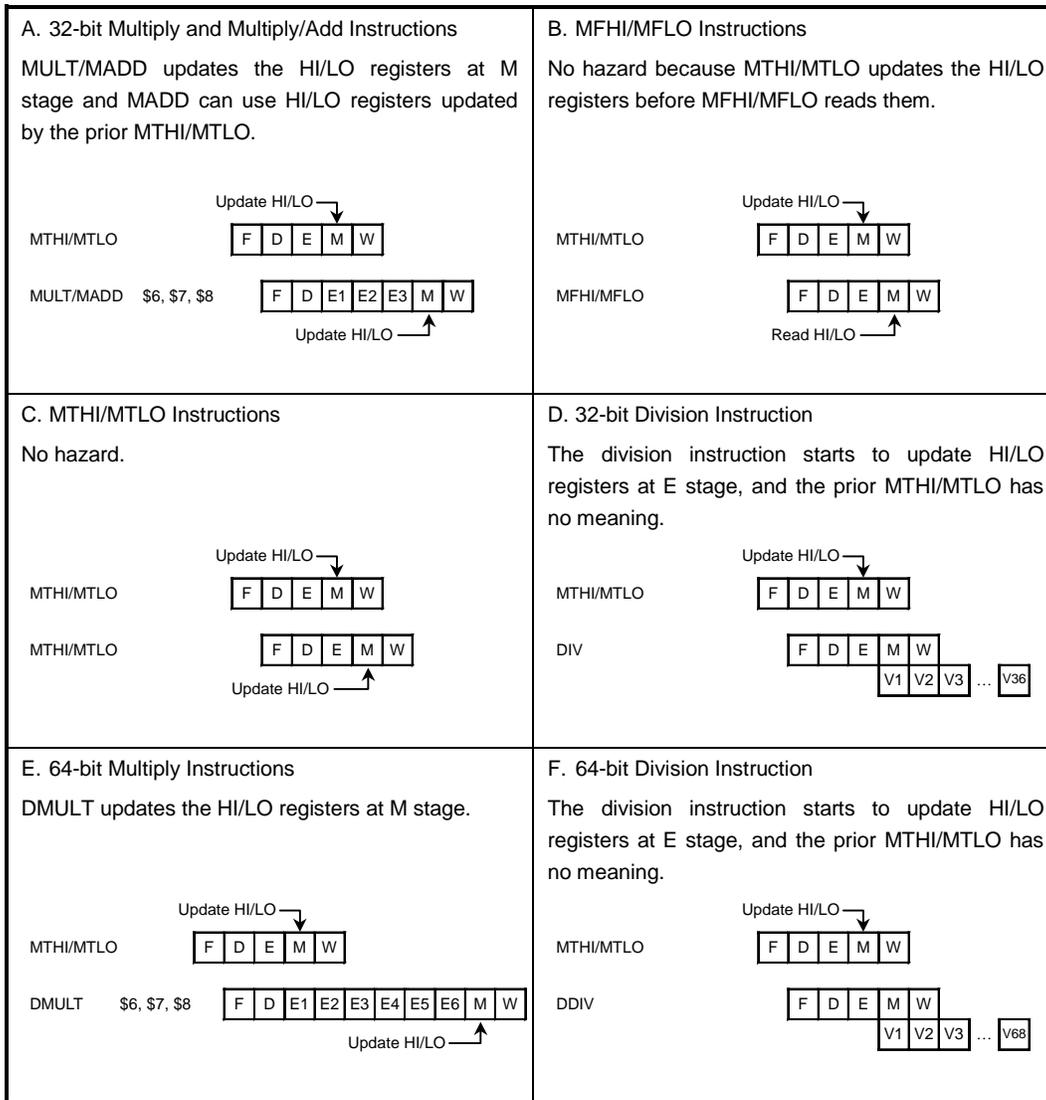


Figure C-10 Pipeline Hazard/Interlock by MTHI/MTLO Instructions

Case 4: Preceding Instruction Is 32-bit Division Instruction

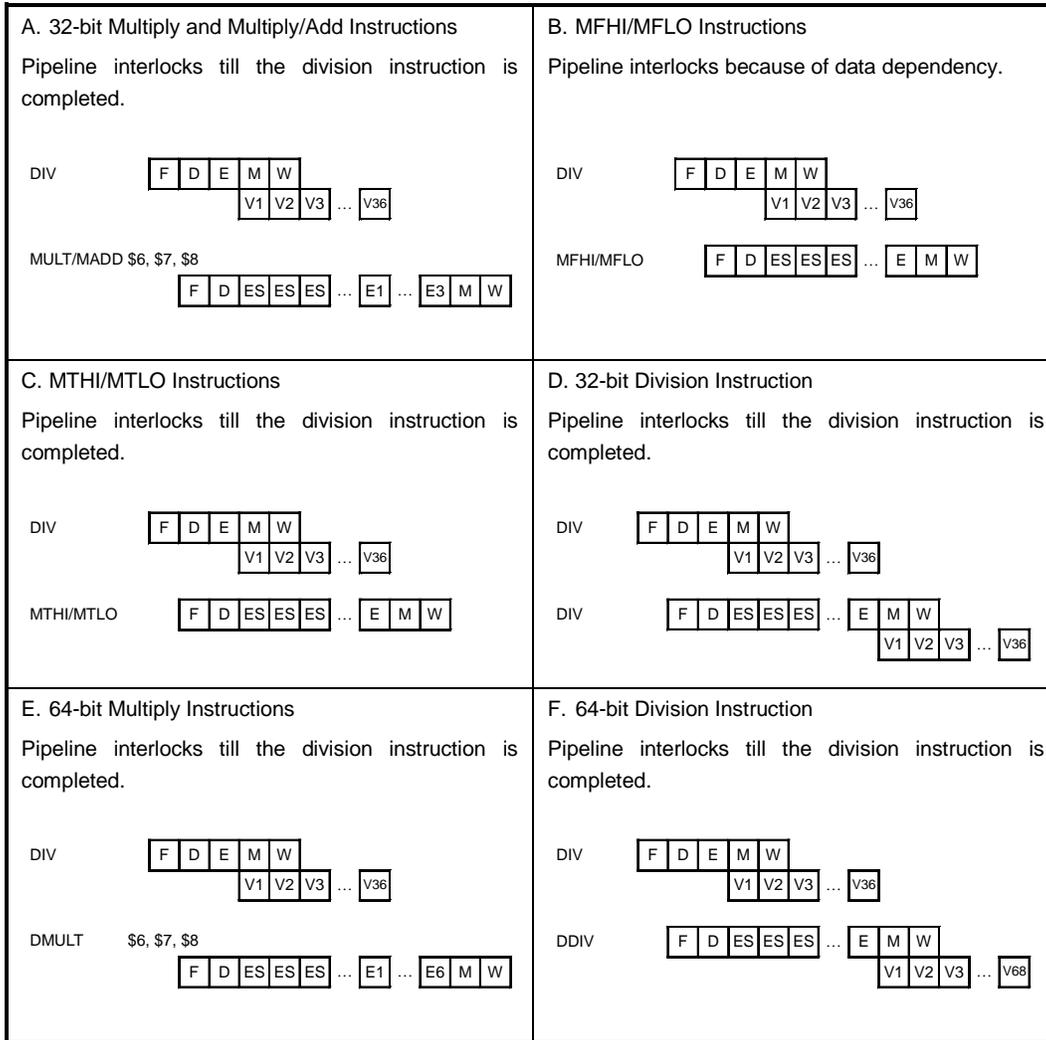


Figure C-11 Pipeline Hazard/Interlock by Division Instructions

Case 5: Preceding Instruction Is 64-bit Multiply Instruction

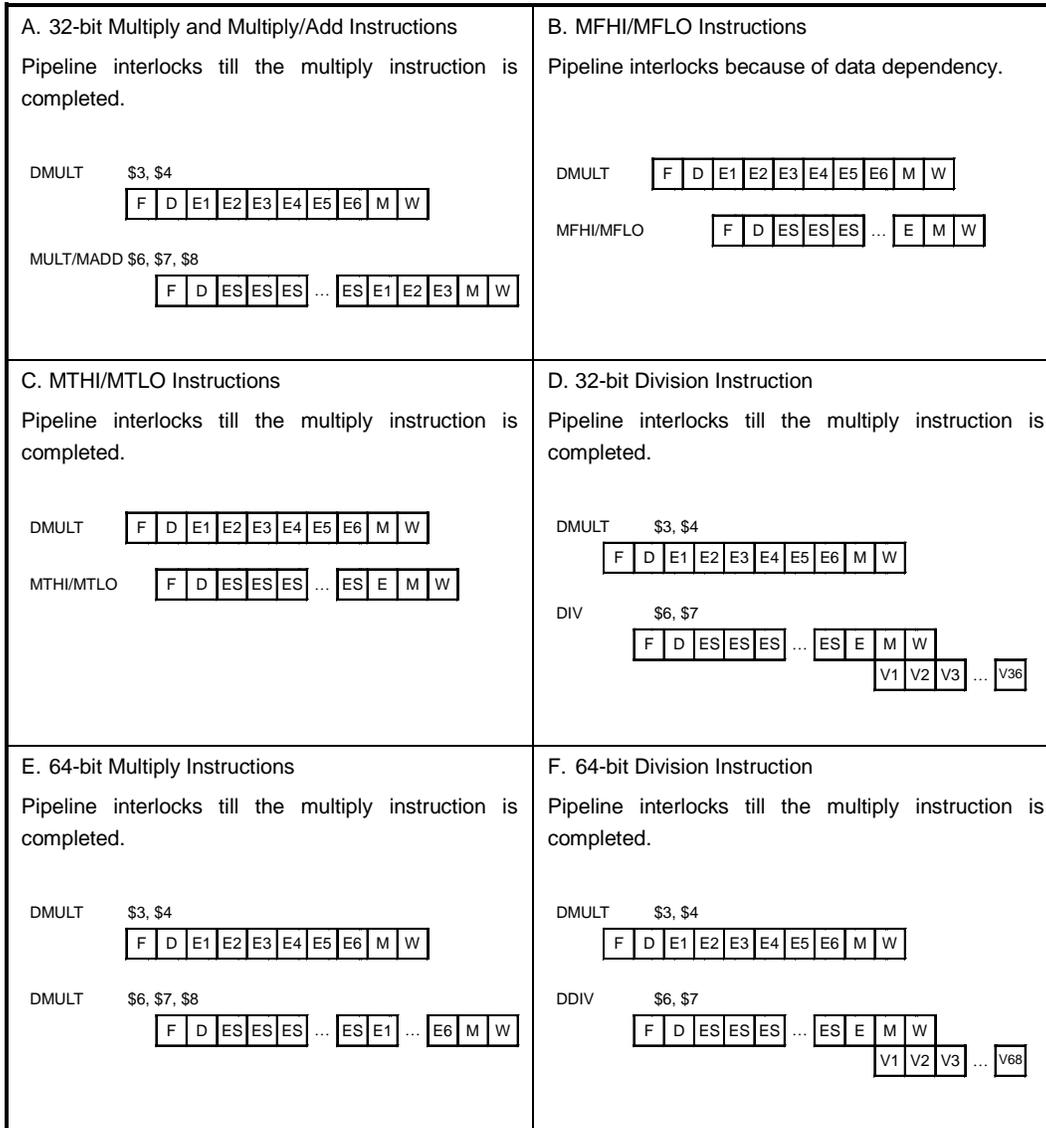


Figure C-12 Pipeline Hazard/Interlock by Division Instructions

Case 6: Preceding Instruction Is 64-bit Division Instruction

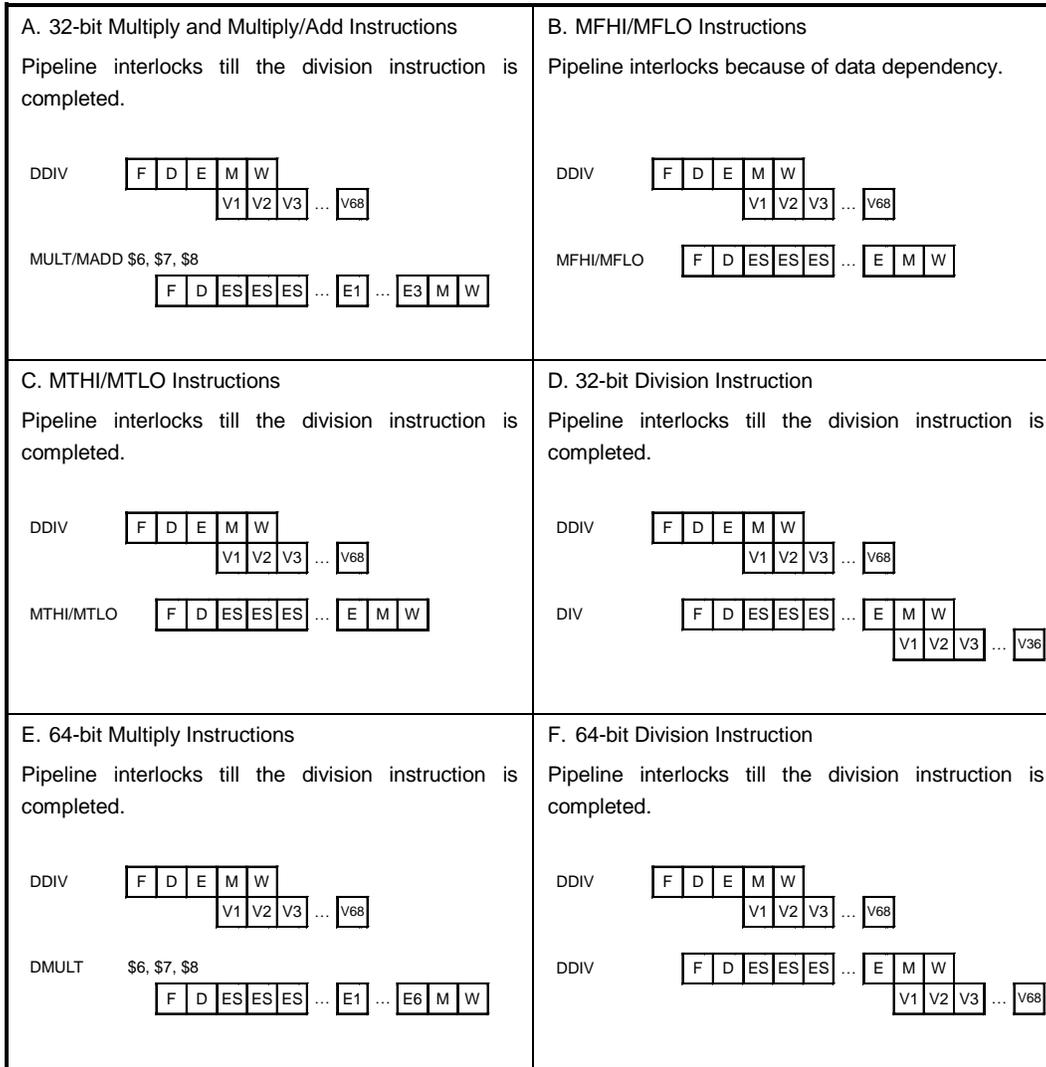


Figure C-13 Pipeline Hazard/Interlock by Division Instructions

Instructions regarding System Control Co-processor (CP0)

MFC0 and MTC0 Instructions

Pipeline interlocks when the MFC0 instruction is followed by the instruction that reads the destination register of MFC0 instruction (Figure C-14).

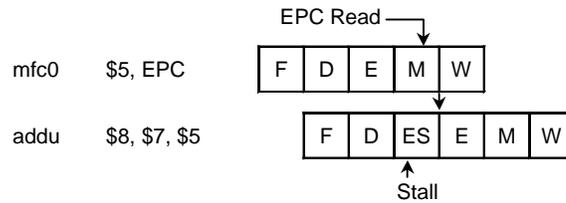


Figure C-14 Pipeline Interlock by MFC0 Instruction

No pipeline hazards occur when the MTC0 instruction is followed by MFC0 instruction because MTC0 writes the destination register in the M stage and MFC0 reads it also in the M stage (Figure C-15).

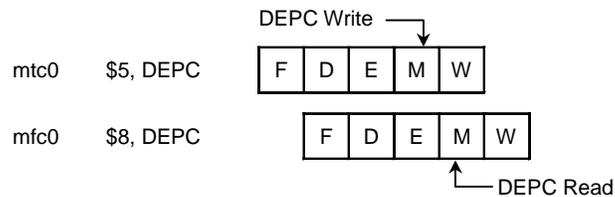


Figure C-15 MTC0 Instruction Followed by MFC0 Instruction

ERET Instruction

Unlike a branch or jump instruction, ERET does not execute the next instruction. The changed EPC becomes effective at the second instruction after the MTC0 instruction (Figure C-16).

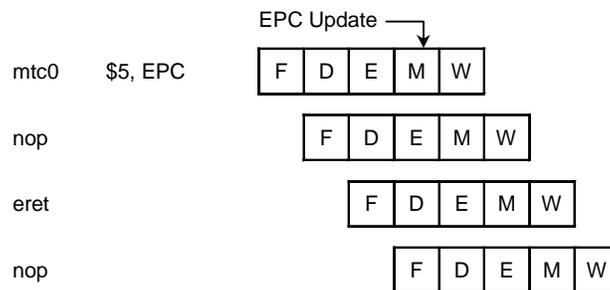


Figure C-16 MTC0 Instruction Followed by ERET Instruction

DERET Instruction

The DERET instruction has a branch delay slot, and the debug exception mode is effective till the delay slot instruction³. The instruction in the delay slot of DERET must be NOP instruction. Single step exception is disabled till the instruction to which DERET returns the control.

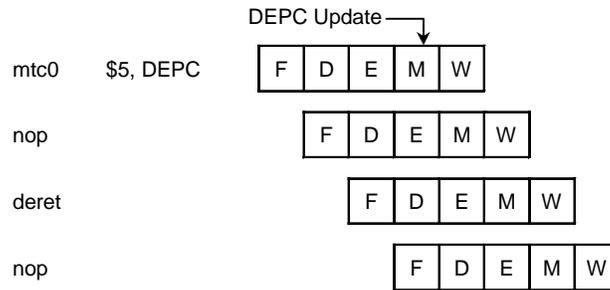


Figure C-17 MTC0 Instruction Followed by DERET Instruction

³ i.e. DM bit stays one (1) and interrupts and exceptions stay disabled.

Control Bits Change in CP0 Registers by MTC0 Instruction

The following sections describe the timings when the control bits change by the MTC0 instruction become effective.

Status Register

CU Bits: Because the co-processor instructions refer the CU bit in the D stage, if either of the two following instructions of the MTC0 instruction is the co-processor instruction, then its result is undefined because the CU bit is undefined (Figure C-18).

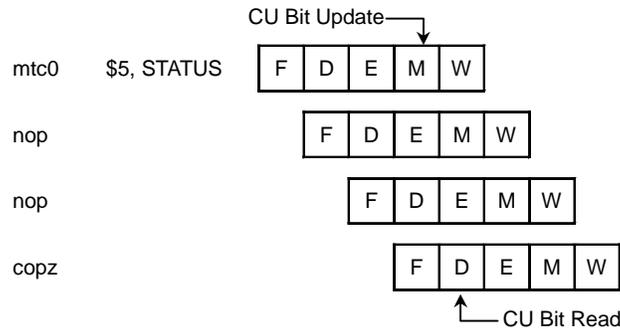


Figure C-18 Hazard regarding the CU Bits

Note that even if the CU bit is changed by the MTC0 instruction during the co-processor bus cycles of the preceding co-processor instruction, this gives no effect on the co-processor instruction currently being executed.

RE Bit: Because the load/store instructions refer the RE bit in the E stage, the change becomes effective at the second instruction after the MTC0 instruction. The result of the load/store instructions immediately after the MTC0 instruction is undefined (Figure C-19).

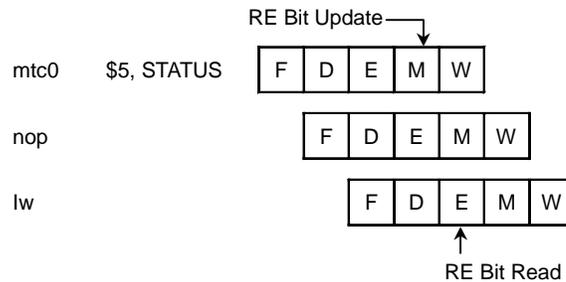


Figure C-19 Hazard regarding the RE Bits

Note that even if the RE bit is changed by the MTC0 instruction during the bus cycles of the preceding load/store instruction, this gives no effect on the load/store instruction currently being executed.

BEV Bit: For the exceptions that occur in the E stage, such as, the address error (AdEL) or the TLB miss (TLBL) exceptions which occurs in the instruction fetch stage, the exception vector base address designated by the changed BEV becomes effective at the second instruction after the MTC0 instruction. If these exceptions occur in the instruction immediately after the MTC0 instruction, the referred value of the BEV bit is undefined⁴ (Figure C-20).

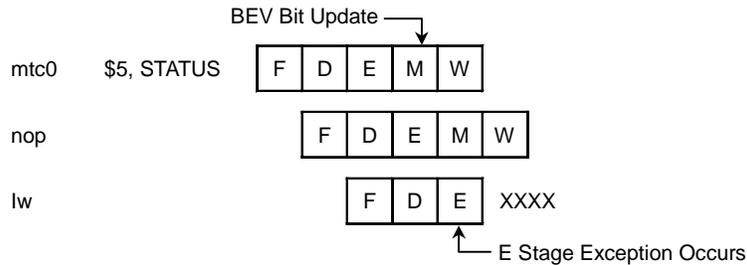


Figure C-20 Hazard regarding the BEV Bits (1)

For the exceptions that occur in the M stage, such as, IBE, DBE, NmI, CpU, Ov, Sys, Bp, RI, AdEL (data), TLBL (data), and TLBS, Mod, and Int, the exception vector base address designated by the changed BEV becomes effective at the instruction immediately after the MTC0 instruction (Figure C-21).

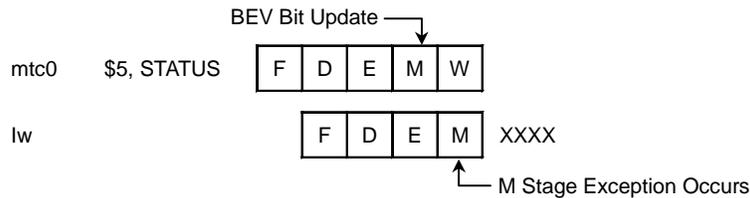


Figure C-21 Hazard regarding the BEV Bits (2)

Note that because the interrupts and the Bus Error exception occurs asynchronously with the instruction execution, the BEV bit value for them is the value which is hold in the BEV bit when they occurs.

IntMask Bits and IE Bit:

When the MTC0 instruction enables the interrupts by changing these bit, then the corresponding interrupts become enabled at the second instruction after the MTC0 instruction⁵ (Figure C-22).

On the other hand, when the MTC0 instruction disables the interrupts, the corresponding interrupts become disabled at the instruction immediately after the MTC0 instruction (Figure C-23).

FR Bit: Because the FR bit is changed in the M stage of the MTC0 instruction, new FR bit becomes effective at the third instruction after

⁴ The new exception vector base address may be effective because of pipeline stall.

⁵ They may become enable at the instruction immediately after the MTC0 instruction because of pipeline stall.

the MTC0 instruction (Figure C-24).

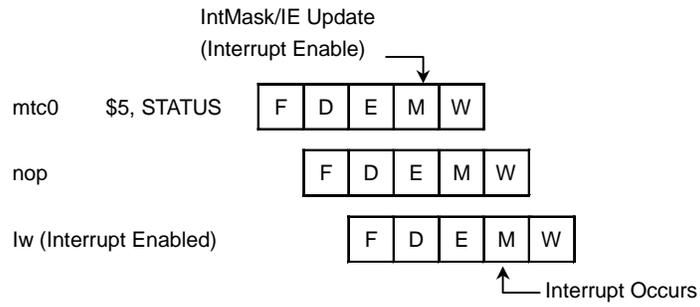


Figure C-22 Hazard regarding the IntMask Bits and IE Bit (1)

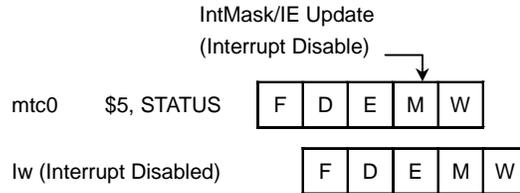


Figure C-23 Hazard regarding the IntMask Bits and IE Bit (2)

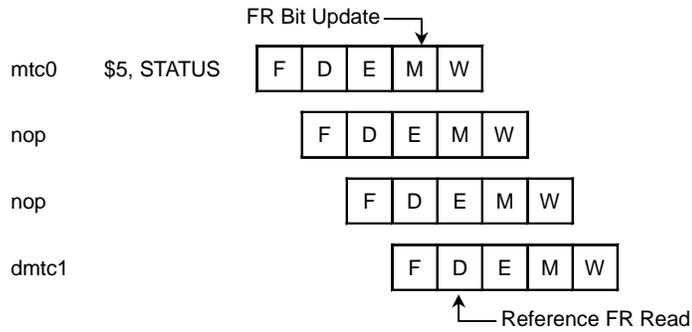


Figure C-24 Hazard regarding the FR Bit

EXL, ERL, KX, SX, UX, KSU Bit:

The modification of these bits become effective at the forth instruction after the MTC0 instruction. On the other hand, new addressing mode for a load/store instruction which is accessing the address in Kernel/Supervisor space or accessing in 64-bit addressing is effective at the second instruction after the MTC0 instruction. If either of the two instructions after the MTC0 instruction is co-processor instruction, result of the instruction is undefined (Figure C-25).

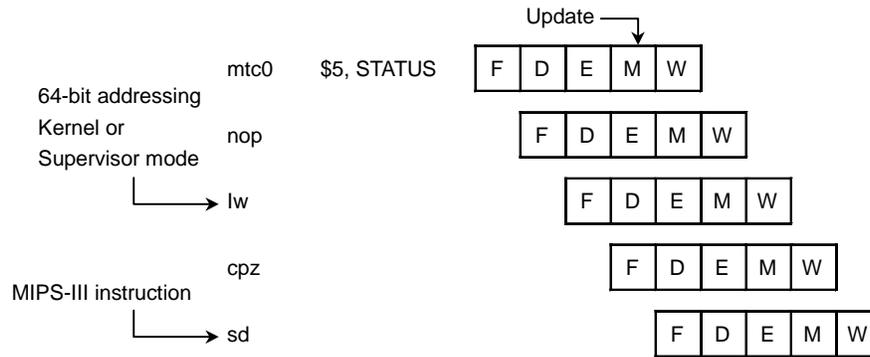


Figure C-25 EXL, ERL, KX, SX, UX, KSU Bit

Config Register

ICE# Bit: The MTC0 instruction may change the ICE# bit during the instruction cache streaming. In this case, the old ICE# bit are effective for the instructions during the streaming (Figure C-26).

```

mtc0 $5, Config      ; update ICE# bit
nop
beq $0, $0, L1      ; stop instruction streaming
nop
L1: lw $2, 0($0)    ; new ICE# bit is effective

```

Figure C-26 ICE# Bit update

DCE# Bit: The changed DCE# becomes effective at the second instruction after the MTC0 instruction. The DCE# bit is undefined at the instruction immediately after the MTC0 instruction. Note that the MTC0 instruction may change the DCE# bit during the data cache refill. In this case, the hardware interlock waits updating the DCE# bit till the data cache refill finishes.

K0 Bit: The modification of these bits becomes effective at the forth instruction after the MTC0 instruction, the result of the instruction in Kseg0 address space is undefined if they executed as first, second or third instruction after the MTC0 instruction. On the other hand, the modification of these bits are effective at the third instruction after MTC0 instruction. New addressing mode for a load/store instruction accessing the Kseg0 address space is undefined if the instruction executed as first or second instruction after MTC0 instruction.

Pipeline Behavior on Cache Miss

This section describes the pipeline behavior on cache miss.

Instruction Cache Miss

Instruction cache miss is detected in F stage and it is immediately followed by a cache refill cycle (Figure C-27).

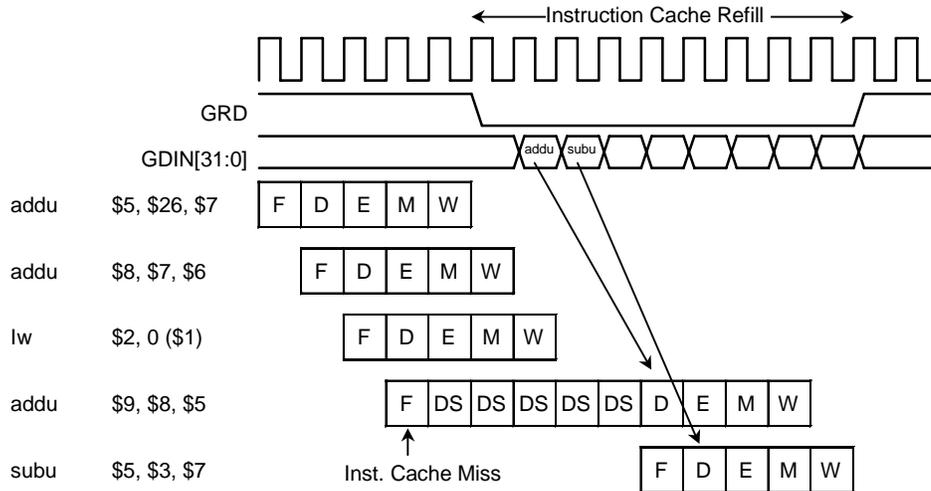


Figure C-27 Streaming on Instruction Cache Refill Cycle in 32-bit GBus mode

On cache miss, the fetched instructions are immediately decoded and executed before completion of refill cycle so that the pipeline resumes the execution of instruction stream as shown in Figure C-27. This is so called *streaming*⁶ and its refill cycle is called *stream cycle*.

When the branch or jump instruction is executed during the stream cycle, streaming will be terminated which means refill cycle will completed but the fetched instructions after the branch delay slot won't be executed. The pipeline will stall until the instruction at the branch or jump target is fetched. (Figure C-28).

⁶ No streaming in 64-bit GBus mode with 1:1 of GBus clock rate. TX49 executes one instruction per clock cycle even if two instructions are fetched in one cycle. In this case, fetched instruction won't be executed until the refill cycle completes.

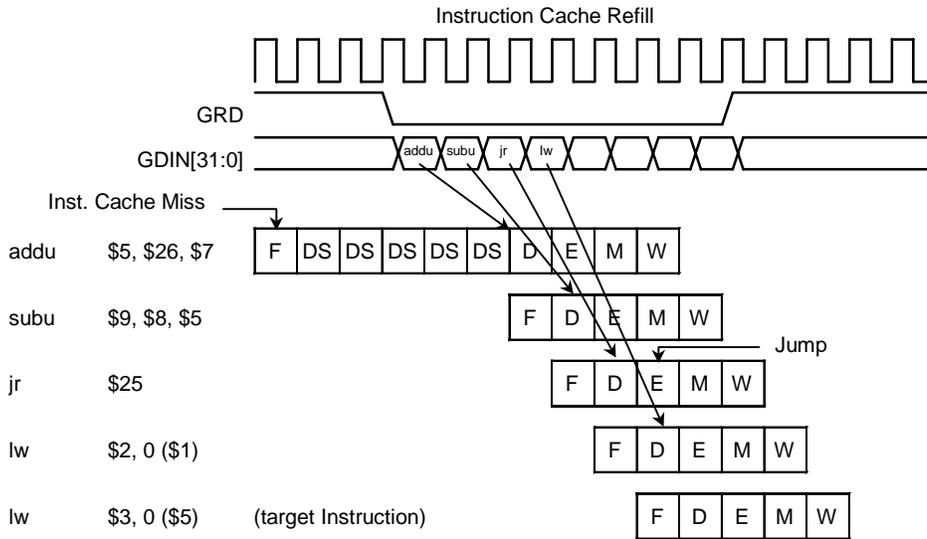


Figure C-28 Branch/Jump Instruction during Stream Cycle in GBus 32-bit Mode

Data Cache Miss

The data cache miss is detected in the M stage of load instruction and it is immediately followed by a cache refill cycle. Non-blocking load mechanism implemented in TX49 data cache allows the following instruction stream to be executed without waiting for the completion of data cache refill if there is no data dependency between the load and the following instructions.

The pipeline will stall at E-stage of the instruction which use the refilled data as its source until the data is loaded. (Figure C-29).

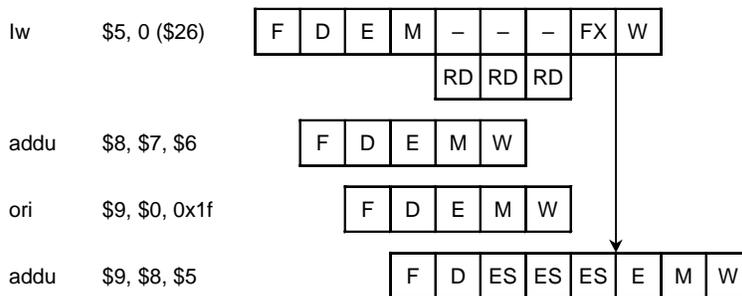


Figure C-29 Pipeline Interlock by Cache Miss

The pipeline also interlocks when a load/store instruction is issued during the data cache refill cycle because of the resource (i.e. data cache) conflict (Figure C-30).

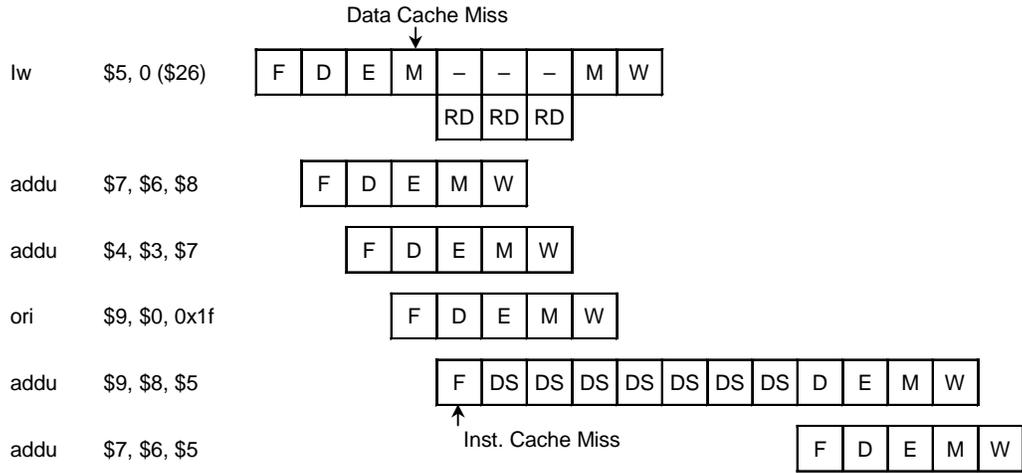


Figure C-32 Instruction Cache Miss during the Data Cache Refill Cycle

Pipeline Behavior in Uncached Area

The pipeline behavior regarding the memory access to an uncached area is similar to that of refill cycle sequence caused by the cache miss.

Data Read from Uncached Area

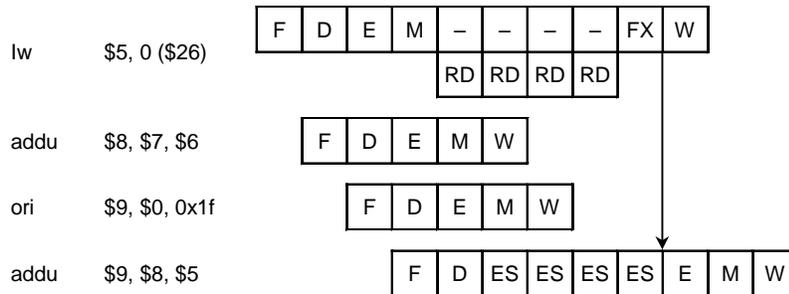


Figure C-33 Data Read from Uncached Area

Instruction Fetch from Uncached Area

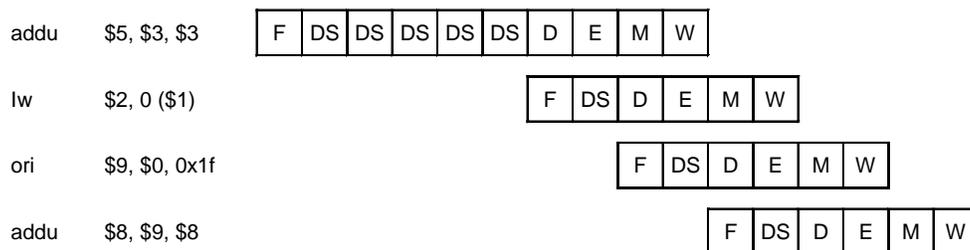


Figure C-34 Instruction Fetch from Uncached Area

Data Write to Uncached Area

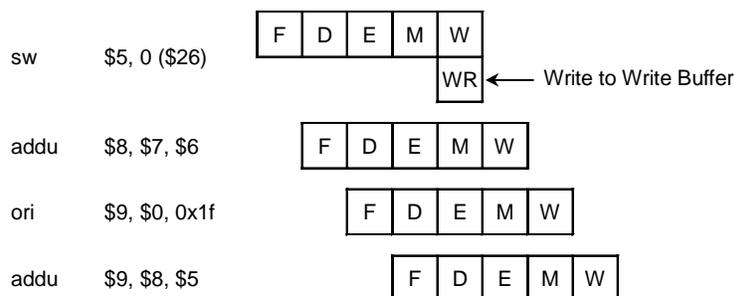


Figure C-35 Data Write to Uncached Area

Timings on the Exception Handling

This section describes the detail pipeline behavior on exception. When an exception takes place, the instruction on which the exception occurs is aborted. All instructions immediately after that instruction are also aborted and the processor passes the control to the exception handler.

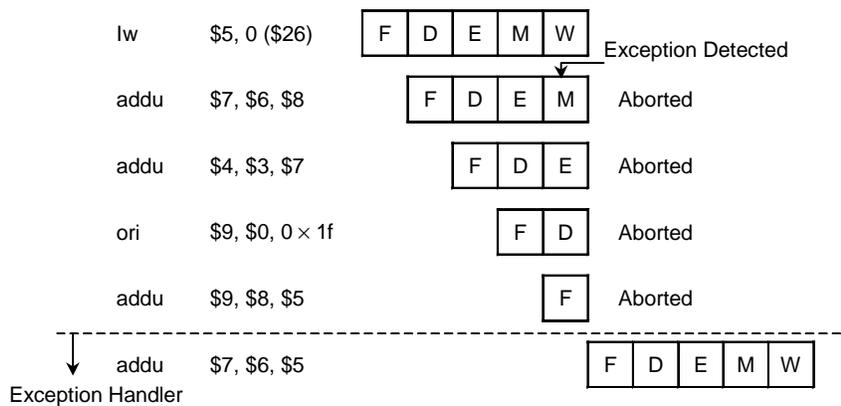
The exceptions normally occur in the M stage, but some of the exceptions occur in the E stage. The exceptions which occur in the E stage are:

- Debug Single Step (DSS)
- Debug Instruction Break (DIB)
- Address Error on Instruction Fetch (AdEL)
- TLB Refill/Invalid on Instruction Fetch (TLBL)

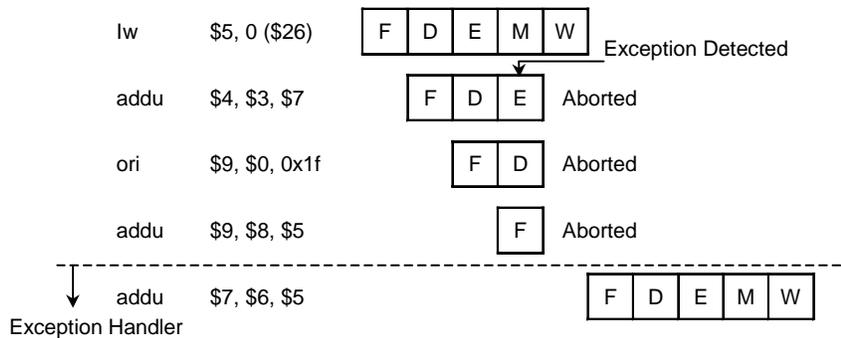
Note that the Reset/Soft Reset Exceptions occur in any stage.

Basic Pipeline Behavior When Exceptions Occur

The following Figure illustrates the pipeline behavior when an exception occurs.



(a) Exception Detected in the M Stage



(b) Exception Detected in the E Stage

Figure C-36 Pipeline Behavior in Case of Exception

Exceptions during the Execution of Multi-cycle Instructions

As described in the section entitle Multiply, Multiply/Add and Division Instructions, multi-cycle instructions which do not have a destination register file, such as DIV, and the following instructions will be executed in parallel if they do not have data dependency.

If an exception takes place at the instruction being executed in parallel with this type of multi-cycle instructions, the preceding multi-cycle instruction is completed while the instructions after the exception are aborted and the control is passed to the exception handler.

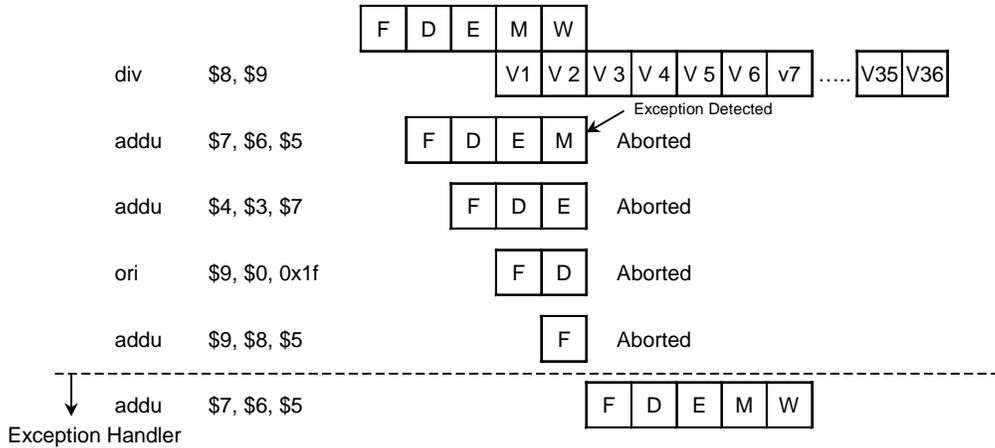


Figure C-37 Exception during the Execution of Division Instruction

Exceptions during the Data Cache Refill Cycle

When one of the exceptions occurs at the instruction which is being executed in parallel with data cache refill, the data cache refill cycle is completed while the instructions after the exception are aborted and the control is passed to the exception handler.

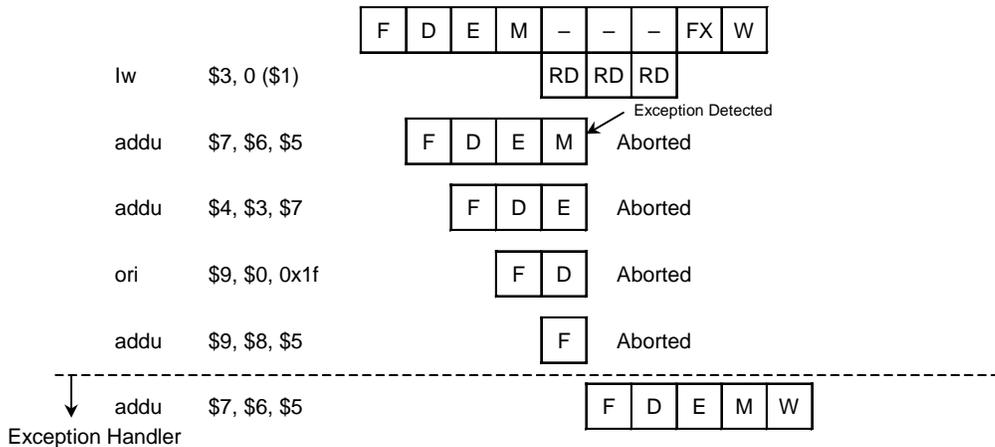


Figure C-38 Exceptions during the Data Cache Refill Cycle (1)

However, when one of the fatal exceptions, such as Bus Error or Reset occurs, the refill cycle is also aborted and the control is passed to the exception handler.

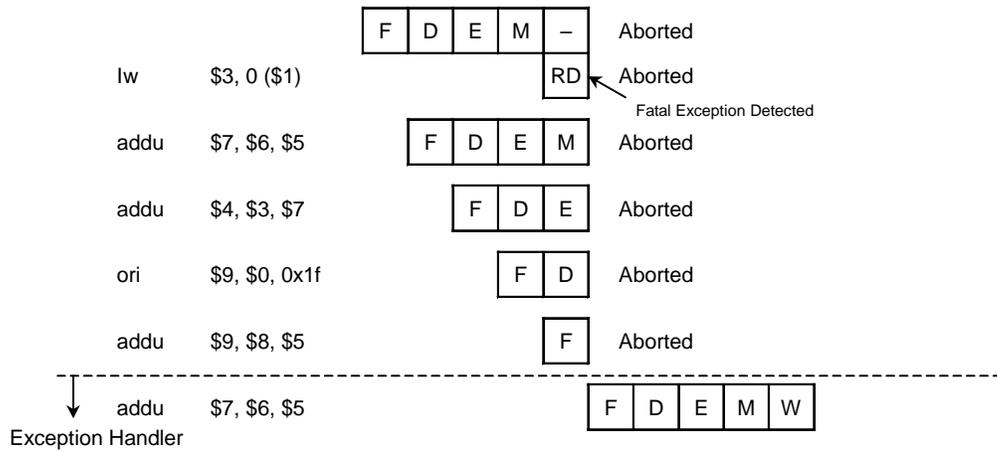


Figure C-39 Exception during Data Cache Refill Cycle (2)

Appendix D: Differences From TX3901, TX4300 and TX4600

Item	TX4901	TX3901	TX4300	TX4600
Datapath	64	32	64	64
ISA	MIPS I, II, III +MADD, +Debug +PREF	MIPS I, II +MADD, +Debug	MIPS I, II, III	MIPS I, II, III
Pipeline	5	5	5	5
MMU	TLB	No	TLB	TLB
JointTLB	48• double		32• double	48• double
I-TLB	2• entry		2• entry	2• entry
D-TLB	4• entry		No	4• entry
Page Size	4• K-16• MB		4• K-16• MB	4• K-16• MB
Shutdown	No-TS		Yes	No-TS
V.A. Size	40		40	40
P.A. Size	36		32	36
I-cache				
Size	16• KB	4• KB	16• KB	16• KB
Associate.	4-way	Dir.-map	Dir.-map	2-way
Lock	Yes	No	No	No
Snoop	No	No	No	No
Index	V	P	V	V
Tag	P	P	P	P
Line	32• B	16• B	32• B	32• B
Parity	No	No	No	Yes
D-cache				
Size	16• KB	1• KB	8• KB	16• KB
Associate.	4-way	2-way	Dir.-map	2-way
Lock	Yes	Yes	No	No
Write Policy	W.-back/-through	W.-through	W.-back	W.-back/-through
Snoop	Yes	Yes	No	No
Index	V	P	V	V
Tag	P	P	P	P
Line	32• B	4• B	16• B	32• B
Parity	No	No	No	Yes

Item	TX4901	TX3901	TX4300	TX4600
WriteBuffer	4A/D pairs	4A/D pairs	4A/D pairs	4A/D pairs
FPU (CP1)	No	No	Shared w/ IU Single Double	FPU Hard Shared w/ I-mul/div Single Double
Debug Support Unit	Yes	Yes	No	No
MPU Bus I/F	G-Bus I/F 36-bit/64-bit A/D separated	32-bit/32-bit A/D separated	SysAD 32-bit A/D multiplexed	SysAD 64-bit A/D multiplexed
Sys.Clock Ratio: 1:1 1.5:1 2:1 3:1 4:1 5:1 6:1 7:1 8:1	Yes No Yes Yes Yes No No No No	Yes No Yes No Yes No No No Yes	No Yes Yes Yes No No No No No	No No Yes Yes Yes Yes Yes Yes Yes
JTAG	No	No	Yes(No func.)	No
Power Sup.	3.3• V	3.3• V	3.3• V	3.3• V
Power down Mode	-WAIT Inst. (Halt/Doze)	-Config. Reg. (Halt/Doze)	-Status. Reg. (1/4 PClock)	-WAIT Inst. (Stand-by)
Package	PQFP-208	PQFP-160	PQFP-120	PGA-179 HSQFP-208

Appendix E: TX49 USER'S MANUAL Revision History

- (1) Rev1.1 Feb.14, 1997
- (2) Rev1.2 Feb.24, 1997
- (3) Rev1.3 Mar.31, 1997
- (4) Rev1.4 Apr.7, 1997
- (5) Rev1.5 Jun.15, 1997
- (6) Rev1.6 Jun.30, 1997
- (7) Rev1.7 Mar.2, 1998
- (8) Rev.2.0 Oct.1, 1998

