



IDT Assembler Software Reference Guide Volume 2

**Version 3.0
December 1998**

2975 Stender Way, Santa Clara, California 95054
Telephone: (800) 345-7015 • TWX: 910-338-2070 • FAX: (408) 492-8674
Printed in U.S.A.
© 1998 Integrated Device Technology, Inc.

Integrated Device Technology, Inc. reserves the right to make changes to its products or specifications at any time, without notice, in order to improve design or performance and to supply the best possible product. IDT does not assume any responsibility for use of any circuitry described other than the circuitry embodied in an IDT product. The Company makes no representations that circuitry described herein is free from patent infringement or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, patent rights or other rights, of Integrated Device Technology, Inc.

LIFE SUPPORT POLICY

Integrated Device Technology's products are not authorized for use as critical components in life support devices or systems unless a specific written agreement pertaining to such intended use is executed between the manufacturer and an officer of IDT.

1. Life support devices or systems are devices or systems which (a) are intended for surgical implant into the body or (b) support or sustain life and whose failure to perform, when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.
2. A critical component is any components of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.

The IDT logo is a registered trademark, and BiCameral, BurstRAM, BUSMUX, CacheRAM, DECnet, Double-Density, FASTX, Four-Port, FLEXI-CACHE, Flexi-PAK, Flow-thruEDC, IDT/c, IDTenvY, IDT/sae, IDT/sim, IDT/ux, MacStation, MICROSLICE, PalatteDAC, REAL8, RC3041, RC3051, RC3052, RC3081, RC36100, RC32364, RC4600, RC4640, RC4650, RC4700, RC5000, RC64474, RC64475, RISController, RISCORE, RISC Subsystem, RISC Windows, SARAM, SmartLogic, SyncFIFO, SyncBiFIFO, SPC, TargetSystem and WideBus are trademarks of Integrated Device Technology, Inc.

MIPS is a registered trademark, and RISCompiler, RISComponent, RISComputer, RISCware, RISC/os, R3000, and R3010 are trademarks of MIPS Computer Systems, Inc. Postscript is a registered trademark of Adobe Systems, Inc. AppleTalk, LocalTalk, and Macintosh are registered trademarks of Apple Computer, Inc. Centronics is a registered trademark of Genicom, Inc. Ethernet is a registered trademark of Digital Equipment Corp. PS2 is a registered trademark of IBM Corp.



About This Manual

Notes

This manual provides a reference for all real hardware (non-synthetic) assembler instructions. A sister publication of this manual provides an introduction and design overview as well as more detailed descriptions for the following IDT product families:

- ◆ *IDT79RC30xx family of 32-bit RISC controllers*
- ◆ *IDT79RC323xx family of 32-bit enhanced MIPS-2 embedded devices*
- ◆ *IDT79RC4xxx 64-BIT RISC CONTROLLER family of high-performance 64-bit CPUs*
- ◆ *IDT79RC5000 family of MIPS-4 ISA compatible CPU devices*

Summary of Contents

Chapter 1, “CPU Instructions Basics,” presents an overview and broad classification of the CPU instruction set of all IDT microprocessors and RISC controllers.

Chapter 2, “CPU Instructions Reference,” is the detailed reference material for each of the CPU instructions in alphabetical order. Each new instruction starts on a new page and the instruction mnemonic is easily locatable at the top of the page in large bold letters.

Chapter 3, “CPU instructions Encoding,” explains the format and encoding of all of the CPU instructions.

Chapter 4, “FPU Instructions Basics,” is similar to Chapter 1 except that it deals with the FPU (hardware floating point unit) instructions.

Chapter 5, “FPU Instructions Reference,” is similar to Chapter 2 except that it deals with the FPU (hardware floating point unit) instructions.

Chapter 6, “FPU instructions Encoding,” is similar to Chapter 3 except that it deals with the FPU (hardware floating point unit) instructions.



Table of Contents

Notes

1 About This Manual

2 CPU Instructions Basics

- Introduction 1-1
- Functional Instruction Groups 1-1
 - Load and Store Instructions 1-2
 - Delayed Loads 1-3
 - CPU Loads and Stores 1-3
 - Atomic Update Loads and Stores 1-4
 - Coprocessor Load and Store Instructions 1-4
- Computational Instructions 1-4
 - Arithmetic Logic Unit 1-4
 - Shift Instructions 1-5
 - Multiply and Divide Instructions 1-6
 - Jump and Branch Instructions 1-6
- Miscellaneous Instructions 1-8
 - Exception Instructions 1-8
 - Serialization Instructions 1-8
 - Conditional Move Instructions 1-8
 - Prefetch Instructions 1-9
- Coprocessor Instructions 1-9
 - Coprocessor Load and Store Instructions 1-10
 - Coprocessor Operations 1-10
- Memory Access Types 1-10
 - Uncached 1-10
 - Cached Noncoherent 1-10
 - Cached Coherent 1-10
 - Cached 1-10
- Mixing References with Different Access Types 1-10
- Cache Coherence Algorithms and Access Types 1-11
- Implementation-Specific Access Types 1-11
- Instruction Descriptions 1-12
 - Instruction Mnemonic and Name 1-12
 - Instruction Encoding Picture 1-12
 - Format 1-13
 - Purpose 1-13
 - Description 1-13
 - Restrictions 1-13
 - Operation 1-13
 - Exceptions 1-14
 - Programming and Implementation Notes 1-14
 - Operation Section Notation and Functions 1-14
 - Pseudocode Language 1-14
 - Pseudocode Symbols 1-14
- Pseudocode Functions 1-16
- Coprocessor General Register Access Functions 1-16
- Load and Store Memory Functions 1-17
- Access Functions for Floating-Point Registers 1-19
- Miscellaneous Functions 1-20

Individual CPU Instruction Descriptions	1-21
3 CPU Instruction Reference	
4 CPU Instructions Encoding	
CPU Instruction Encoding	3-2
Instruction Decode	3-2
SPECIAL Instruction Class	3-2
REGIMM Instruction Class	3-2
Instruction Subsets of MIPS III and MIPS IV Processors	3-2
Non-CPU Instructions in the Tables	3-2
Coprocessor 0 - COP0	3-2
Coprocessor 1 - COP1, COP1X, MOVCI, and CP1 load/store	3-3
Coprocessor 2 - COP2 and CP2 load/store	3-3
Coprocessor 3 - COP3 and CP3 load/store	3-3
5 FPU Instructions Basics	
FPU Instruction Set Details	4-1
FPU Instructions	4-1
Data Transfer Instructions	4-1
Arithmetic Instructions	4-2
Conversion Instructions	4-3
Formatted Operand Value Move Instructions	4-4
Conditional Branch Instructions	4-5
Miscellaneous Instructions	4-5
Valid Operands for FP Instructions	4-5
Description of an Instruction	4-6
Operation Notation Conventions and Functions	4-7
Individual FPU Instruction Descriptions	4-7
6 FPU Instructions Reference	
7 FPU Instructions Encoding	
FPU (CP1) Instruction Opcode Bit Encoding	6-3
Instruction Decode	6-3
COP1 Instruction Class	6-4
COP1X Instruction Class	6-4
SPECIAL Instruction Class	6-4
Instruction Subsets of MIPS III and MIPS IV Processors	6-4
Key to all FPU (CP1) instruction encoding tables:	6-25
8 Index	



List of Tables

Notes

Table 1.1	Load/Store Operations Using Register + Offset Addressing Mode	1-2
Table 1.2	Load/Store Operations Using Register + Register Addressing Mode.....	1-2
Table 1.3	Normal CPU Load/Store Instructions.....	1-3
Table 1.4	Unaligned CPU Load/Store Instructions	1-3
Table 1.5	Atomic Update CPU Load/Store Instructions.....	1-4
Table 1.6	Coprocessor Load/Store Instructions	1-4
Table 1.7	PFU Load/Store Instructions Using Register + Register Addressing.....	1-4
Table 1.8	ALU Instructions With an Immediate Operand	1-5
Table 1.9	Operand ALU Instructions	1-5
Table 1.10	Shift Instructions	1-5
Table 1.11	Multiply/Divide Instructions	1-6
Table 1.12	Jump Instructions Jumping Within a 256 Megabyte Region.....	1-7
Table 1.13	Jump Instructions to Absolute Address	1-7
Table 1.14	PC-Relative Conditional Branch Instructions, Comparing 2 Registers.....	1-7
Table 1.15	PC-Relative Conditional Branch Instructions, Comparing Against Zero.....	1-7
Table 1.16	System Call and Breakpoint Instructions	1-8
Table 1.17	Trap-on-Condition Instructions, Comparing Two Registers.....	1-8
Table 1.18	Trap-on-Condition Instructions, Comparing an Immediate	1-8
Table 1.19	Serialization Instructions.....	1-8
Table 1.20	CPU Conditional Move Instructions	1-9
Table 1.21	Prefetch Using Register + Offset Address Mode.....	1-9
Table 1.22	Prefetch Using Register + Register Address Mode	1-9
Table 1.23	Coprocessor Definition and Use in the MIPS Architecture	1-9
Table 1.24	Coprocessor Operation Instructions	1-10
Table 1.25	Symbols in Instruction Operation Statements (Page 1 of 2).....	1-14
Table 1.26	Coprocessor General Register Access Functions	1-16
Table 1.27	AccessLength Specifications for Loads/Stores.....	1-18
Table 2.28	64-bit RISController Family Primary Cache Indexing	2-25
Table 2.29	Values of Hint Field for Prefetch Instruction in RC32364	2-115
Table 2.30	Values of Hint Field for Prefetch Instruction in RC5000	2-116
Table 2.31	Bytes Stored by SDL Instruction.....	2-127
Table 2.32	Bytes Stored by SDR Instruction	2-129
Table 2.33	Unaligned Word Store using SWL and SWR.....	2-146
Table 2.34	Bytes Stored by SWL Instruction	2-147
Table 2.35	Bytes Stored by SWR Instruction	2-150
Table 3.1	CPU Instruction Formats	3-1
Table 3.2	CPU Instruction Encoding - MIPS I Architecture	3-4
Table 3.3	CPU Instruction Encoding - MIPS II Architecture	3-5
Table 3.4	CPU Instruction Encoding - MIPS III Architecture	3-6
Table 3.5	CPU Instruction Encoding - MIPS IV Architecture	3-7
Table 3.6	Architecture Level in Which CPU Instructions are Defined or Extended	3-8
Table 3.7	CPU Instruction Encoding Changes - MIPS II Revision	3-9
Table 3.8	CPU Instruction Encoding Changes - MIPS III Revision	3-10
Table 3.9	CPU Instruction Encoding Changes - MIPS IV Revision	3-11
Table 4.10	FPU Loads and Stores Using Register + Offset Address Mode.....	4-2
Table 4.11	FPU Loads and Stores Using Register + Register Address Mode	4-2
Table 4.12	FPU Move To/From Instructions.....	4-2
Table 4.13	FPU IEEE Arithmetic Operations.....	4-3

Table 4.14	FPU Approximate Arithmetic Operations	4-3
Table 4.15	FPU Multiply-Accumulate Arithmetic Operations	4-3
Table 4.16	FPU Conversion Operations Using a Directed Rounding Mode	4-4
Table 4.17	FPU Formatted Operand Move Instructions	4-4
Table 4.18	FPU Conditional Move on True/False Instructions	4-4
Table 4.19	FPU Conditional Move on Zero/Nonzero Instructions	4-4
Table 4.20	FPU Conditional Branch Instructions	4-5
Table 4.21	CPU Conditional Move on FPU True/False Instructions	4-5
Table 4.22	FPU Operand Format Field (fmt, fmt3) Decoding	4-5
Table 4.23	Valid Formats for FPU Operations	4-6
Table 5.24	FPU Comparisons Without Special Operand Exceptions	5-13
Table 5.25	FPU Comparisons With Special Operand Exceptions for QNaNs	5-14



List of Figures

Notes

Figure 1.1 MIPS Architecture Extensions 1-1

Figure 1.2 Example Instruction Description 1-12

Figure 2.3 Unaligned Doubleword Load using LDL and LDR 2-67

Figure 2.4 Bytes Loaded by LDL Instruction 2-68

Figure 2.5 Unaligned Doubleword Load using LDR and LDL 2-69

Figure 2.6 Bytes Loaded by LDR Instruction 2-70

Figure 2.7 Unaligned Doubleword Load using LDL and LDR 2-77

Figure 2.8 Bytes Loaded by LDL Instruction 2-78

Figure 2.9 Unaligned Doubleword Load using LDR and LDL 2-79

Figure 2.10 Bytes Loaded by LDR Instruction 2-80

Figure 2.11 Unaligned Word Load using LWL and LWR 2-90

Figure 2.12 Bytes Loaded by LWL Instruction 2-91

Figure 2.13 Unaligned Word Load using LWR and LWL 2-93

Figure 2.14 Bytes Loaded by LWR Instruction..... 2-94

Figure 2.15 Unaligned Doubleword Store with SDL and SDR 2-126

Figure 2.16 Unaligned Doubleword Store with SDR and SDL 2-128

Figure 2.17 Unaligned Word Store using SWR and SWL 2-149



CPU Instructions Basics

Notes

Introduction

This chapter describes the instruction set architecture (ISA) for the central processing unit (CPU) in the MIPS IV architecture. The CPU architecture defines the non-privileged instructions that execute in user mode. It does not define privileged instructions providing processor control executed by the implementation-specific System Control Processor. Instructions for the floating-point unit (FPU) are described in Chapters 4, 5 and 6.

The original MIPS I CPU ISA has been extended in a backward-compatible fashion three times. The ISA extensions are inclusive as the diagram illustrates; each new architecture level (or version) includes the former levels. The description of an architectural feature includes the architecture level in which the feature is (first) defined or extended. The feature is also available in all later (higher) levels of the architecture.

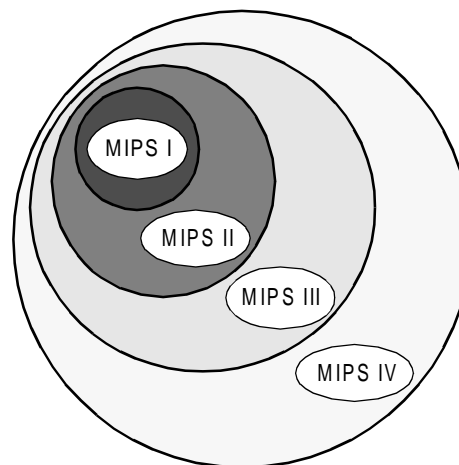


Figure 1.1 MIPS Architecture Extensions

The practical result is that a processor implementing MIPS IV is also able to run MIPS I, MIPS II, or MIPS III user-mode binary programs without change.

It should be noted that there may not always be a one-to-one relationship between an IDT microprocessor or RISController and a MIPS ISA level. Some IDT parts adhere strictly to a MIPS ISA level, some implement a specific MIPS ISA level and also implement additional special instructions (for example, in the case of RC4640, RC4650), while yet others implement a combination of different MIPS ISA levels and also additional special instructions (for example, in the case of RC32364).

The CPU instruction set is first summarized by functional group. In Chapter 2 each instruction is described separately in alphabetical order. Chapter 3 describes the organization of the individual instruction descriptions and the notation used in them (including FPU instructions). It concludes with the CPU instruction formats and opcode encoding tables.

Functional Instruction Groups

CPU instructions are divided into the following functional groups:

- ◆ *Load and Store*
- ◆ *Arithmetic Logic Unit*
- ◆ *Jump and Branch*
- ◆ *Miscellaneous*
- ◆ *Coprocessor*

Load and Store Instructions

Load and store instructions transfer data between the memory system and the general register sets in the CPU and the coprocessors. There are separate instructions for different purposes: transferring various sized fields, treating loaded data as signed or unsigned integers, accessing unaligned fields, selecting the addressing mode, and providing atomic memory update (read-modify-write).

Regardless of byte ordering (big- or little-endian), the address of a halfword, word, or doubleword is the smallest byte address among the bytes forming the object. For big-endian ordering this is the most-significant byte; for a little-endian ordering this is the least-significant byte.

Except for the few specialized instructions listed in Table 1.4, loads and stores must access naturally aligned objects. An attempt to load or store an object at an address that is not an even multiple of the size of the object will cause an Address Error exception.

Load and store operations have been added in each revision of the architecture:

- ◆ *MIPS II*
 - *64-bit coprocessor transfers¹*
 - *atomic update*
- ◆ *MIPS III*
 - *64-bit CPU transfers*
 - *unsigned word load for CPU*
- ◆ *MIPS IV*
 - *register + register addressing mode for FPU*

Table 1.1 and Table 1.2 tabulate the supported load and store operations and indicate the MIPS architecture level at which each operation was first supported. The instructions themselves are listed in the following sections.

Data Size	Load Signed	CPU		Coprocessor (except 0)	
		Load Unsigned	Store	Load	Store
byte	I	I	I		
halfword	I	I	I		
word	I	III	I	I	I
doubleword	III		III	II	II
unaligned word	I		I		
unaligned doubleword	III		III		
linked word (atomic modify)	II		II		
linked doubleword (atomic modify)	III		III		

Table 1.1 Load/Store Operations Using Register + Offset Addressing Mode

Data Size	floating-point coprocessor only	
	Load	Store
word	IV	IV
doubleword	IV	IV

Table 1.2 Load/Store Operations Using Register + Register Addressing Mode

¹ Even though the RISCORE32300 implements MIPS II, double word accesses will signal a trap.

Delayed Loads

The MIPS I architecture defines delayed loads; an instruction scheduling restriction requires that an instruction immediately following a load into register *Rn* cannot use *Rn* as a source register. The time between the load instruction and the time the data is available is the “load delay slot”. If no useful instruction can be put into the load delay slot, then a null operation (assembler mnemonic NOP) must be inserted.

In MIPS II, this instruction scheduling restriction is removed. Programs will execute correctly when the loaded data is used by the instruction following the load, but this may require extra real cycles. Most processors cannot actually load data quickly enough for immediate use and the processor will be forced to wait until the data is available. Scheduling load delay slots is desirable for performance reasons even when it is not necessary for correctness.

CPU Loads and Stores

There are instructions to transfer different amounts of data: bytes, halfwords, words, and doublewords. Signed and unsigned integers of different sizes are supported by loads that either sign-extend or zero-extend the data loaded into the register.

Mnemonic	Description	Defined in
LB	Load Byte	I
LBU	Load Byte Unsigned	I
SB	Store Byte	I
LH	Load Halfword	I
LHU	Load Halfword Unsigned	I
SH	Store Halfword	I
LW	Load Word	I
LWU	Load Word Unsigned	III
SW	Store Word	I
LD	Load Doubleword	III
SD	Store Doubleword	III

Table 1.3 Normal CPU Load/Store Instructions

Unaligned words and doublewords can be loaded or stored in only two instructions by using a pair of special instructions. The load instructions read the left-side or right-side bytes (left or right side of register) from an aligned word and merge them into the correct bytes of the destination register. MIPS I, though it prohibits other use of loaded data in the load delay slot, permits LWL and LWR instructions targeting the same destination register to be executed sequentially. Store instructions select the correct bytes from a source register and update only those bytes in an aligned memory word (or doubleword).

Mnemonic	Description	Defined in
LWL	Load Word Left	I
LWR	Load Word Right	I
SWL	Store Word Left	I
SWR	Store Word Right	I
LDL	Load Doubleword Left	III
LDR	Load Doubleword Right	III
SDL	Store Doubleword Left	III
SDR	Store Doubleword Right	III

Table 1.4 Unaligned CPU Load/Store Instructions

Atomic Update Loads and Stores

There are paired instructions, Load Linked and Store Conditional, that can be used to perform atomic read-modify-write of word and doubleword cached memory locations. 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. The individual instruction descriptions describe how to use them.

Mnemonic	Description	Defined in
LL	Load Linked Word	II
SC	Store Conditional Word	II
LLD	Load Linked Doubleword	III
SCD	Store Conditional Doubleword	III

Table 1.5 Atomic Update CPU Load/Store Instructions

Coprocessor Load and Store Instructions

These loads and stores are coprocessor instructions, however it seems more useful to summarize all load and store instructions in one place instead of listing them in the coprocessor instructions functional group.

If a particular coprocessor is not enabled, loads and stores to that processor cannot execute and will cause a Coprocessor Unusable exception. Enabling a coprocessor is a privileged operation provided by the System Control Coprocessor.

Mnemonic	Description	Defined in
LWCz	Load Word to Coprocessor-z	I
SWCz	Store Word from Coprocessor-z	I
LDCz	Load Doubleword to Coprocessor-z	II
SDCz	Store Doubleword from Coprocessor-z	II

Table 1.6 Coprocessor Load/Store Instructions

Mnemonic	Description	Defined in
LWXC1	Load Word Indexed to Floating Point	IV
SWXC1	Store Word Indexed from Floating Point	IV
LDXC1	Load Doubleword Indexed to Floating Point	IV
SDXC1	Store Doubleword Indexed from Floating Point	IV

Table 1.7 PFU Load/Store Instructions Using Register + Register Addressing

Computational Instructions

Computational instructions perform arithmetic, logical, shift, multiply, and divide operations on values in registers. Two's complement arithmetic is performed on integers represented in two's complement notation. There are signed versions of add, subtract, multiply, and divide. There are add and subtract operations, called "unsigned," that are actually modulo arithmetic without overflow detection. There are unsigned versions of multiply and divide. There is a full complement of shift and logical operations.

MIPS I provides 32-bit integers and 32-bit arithmetic. MIPS III adds 64-bit integers and provides separate arithmetic and shift instructions for 64-bit operands. Logical operations are not sensitive to the width of the register.

Arithmetic Logic Unit

Some arithmetic and logical instructions operate on one operand, from a register and the other from a 16-bit immediate value in the instruction word. The immediate operand is treated as signed for the arithmetic and compare instructions, and treated as logical (zero-extended to register length) for the logical instructions.

Mnemonic	Description	Defined in
ADDI	Add Immediate Word	I
ADDIU	Add Immediate Unsigned Word	I
SLTI	Set on Less Than Immediate	I
SLTIU	Set on Less Than Immediate Unsigned	I
ANDI	And Immediate	I
ORI	Or Immediate	I
XORI	Exclusive Or Immediate	I
LUI	Load Upper Immediate	I
DADDI	Doubleword Add Immediate	III
DADDIU	Doubleword Add Immediate Unsigned	III

Table 1.8 ALU Instructions With an Immediate Operand

Mnemonic	Description	Defined in
ADD	Add Word	I
ADDU	Add Unsigned Word	I
SUB	Subtract Word	I
SUBU	Subtract Unsigned Word	I
DADD	Doubleword Add	III
DADDU	Doubleword Add Unsigned	III
DSUB	Doubleword Subtract	III
DSUBU	Doubleword Subtract Unsigned	III
SLT	Set on Less Than	I
SLTU	Set on Less Than Unsigned	I
AND	And	I
OR	Or	I
XOR	Exclusive Or	I
NOR	Nor	I

Table 1.9 Operand ALU Instructions

Shift Instructions

There are shift instructions that take the shift amount from a 5-bit field in the instruction word and shift instructions that take a shift amount from the low-order bits of a general register. The instructions with a fixed shift amount are limited to a 5-bit shift count, so there are separate instructions for doubleword shifts of 0-31 bits and 32-63 bits.

Mnemonic	Description	Defined in
SLL	Shift Word Left Logical	I
SRL	Shift Word Right Logical	I
SRA	Shift Word Right Arithmetic	I
SLLV	Shift Word Left Logical Variable	I
SRLV	Shift Word Right Logical Variable	I
SRAV	Shift Word Right Arithmetic Variable	I
DSLL	Doubleword Shift Left Logical	III
DSRL	Doubleword Shift Right Logical	III

Table 1.10 Shift Instructions

Mnemonic	Description	Defined in
DSRA	Doubleword Shift Right Arithmetic	III
DSLL32	Doubleword Shift Left Logical + 32	III
DSRL32	Doubleword Shift Right Logical + 32	III
DSRA32	Doubleword Shift Right Arithmetic + 32	III
DSLLV	Doubleword Shift Left Logical Variable	III
DSRLV	Doubleword Shift Right Logical Variable	III
DSRAV	Doubleword Shift Right Arithmetic Variable	III

Table 1.10 Shift Instructions

Multiply and Divide Instructions

Multiply produces a full-width product twice the width of the input operands: the low half is placed in LO and the high half is placed in HI. Integer divides produce both a quotient in LO and a remainder in HI. These results are accessed by instructions that transfer data between these special purpose registers and the general registers.

The RC4650 adds the MAD or MADU instruction (multiply-accumulate or multiply-accumulate unsigned, with HI and LO as the accumulator) to the base MIPS-III ISA. The MAD or MADU instruction uses the HI and LO registers as a 64-bit accumulator. This process allows these instructions to compatibly operate in 32-bit processors.

The RC4650 also adds MUL, a 3-operand 32x32 → 32 multiply instruction that eliminates the need to explicitly move the multiply result from the LO register back to a general register.

Note: After executing the MUL instruction, the HI and LO registers are undefined.

Mnemonic	Description	Defined in
MAD	Multiply/Add	IDT extension
MADU	Multiply/Add Unsigned	IDT extension
MUL	Multiply	IDT extension
MULT	Multiply Word	MIPS I
MULTU	Multiply Unsigned Word	MIPS I
DIV	Divide Word	I
DIVU	Divide Unsigned Word	I
DMULT	Doubleword Multiply	III
DMULTU	Doubleword Multiply Unsigned	III
DDIV	Doubleword Divide	III
DDIVU	Doubleword Divide Unsigned	III
MFHI	Move From HI	I
MTHI	Move To HI	I
MFLO	Move From LO	I
MTLO	Move To LO	I

Table 1.11 Multiply/Divide Instructions

Jump and Branch Instructions

The architecture defines PC-relative conditional branches, a PC-region unconditional jump, an absolute (register) unconditional jump, and a similar set of procedure calls that record a return link address in a general register. For convenience this discussion refers to them all as branches.

All branches have an architectural delay of one instruction. When a branch is taken, the instruction immediately following the branch instruction, in the branch delay slot, is executed before the branch to the target instruction takes place. Conditional branches come in two versions that treat the instruction in the delay slot differently when the branch is not taken and execution falls through. The “branch” instructions execute the instruction in the delay slot, but the “branch likely” instructions do not (they are said to nullify it).

By convention, if an exception or interrupt prevents the completion of an instruction occupying a branch delay slot, the instruction stream is continued by re-executing the branch instruction. To permit this, branches must be restartable; procedure calls may not use the register in which the return link is stored (usually register 31) to determine the branch target address.

Mnemonic	Description	Defined in
J	Jump	I
JAL	Jump and Link	I

Table 1.12 Jump Instructions Jumping Within a 256 Megabyte Region

Mnemonic	Description	Defined in
JR	Jump Register	I
JALR	Jump and Link Register	I

Table 1.13 Jump Instructions to Absolute Address

Mnemonic	Description	Defined in
BEQ	Branch on Equal	I
BNE	Branch on Not Equal	I
BLEZ	Branch on Less Than or Equal to Zero	I
BGTZ	Branch on Greater Than Zero	I
BEQL	Branch on Equal Likely	II
BNEL	Branch on Not Equal Likely	II
BLEZL	Branch on Less Than or Equal to Zero Likely	II
BGTZL	Branch on Greater Than Zero Likely	II

Table 1.14 PC-Relative Conditional Branch Instructions, Comparing 2 Registers

Mnemonic	Description	Defined in
BLTZ	Branch on Less Than Zero	I
BGEZ	Branch on Greater Than or Equal to Zero	I
BLTZAL	Branch on Less Than Zero and Link	I
BGEZAL	Branch on Greater Than or Equal to Zero and Link	I
BLTZL	Branch on Less Than Zero Likely	II
BGEZL	Branch on Greater Than or Equal to Zero Likely	II
BLTZALL	Branch on Less Than Zero and Link Likely	II
BGEZALL	Branch on Greater Than or Equal to Zero and Link Likely	II

Table 1.15 PC-Relative Conditional Branch Instructions, Comparing Against Zero

Miscellaneous Instructions

Exception Instructions

Exception instructions have as their sole purpose causing an exception that will transfer control to a software exception handler in the kernel. System call and breakpoint instructions cause exceptions unconditionally. The trap instructions cause exceptions conditionally based upon the result of a comparison.

Mnemonic	Description	Defined in
SYSCALL	System Call	I
BREAK	Breakpoint	I

Table 1.16 System Call and Breakpoint Instructions

Mnemonic	Description	Defined in
TGE	Trap if Greater Than or Equal	II
TGEU	Trap if Greater Than or Equal Unsigned	II
TLT	Trap if Less Than	II
TLTU	Trap if Less Than Unsigned	II
TEQ	Trap if Equal	II
TNE	Trap if Not Equal	II

Table 1.17 Trap-on-Condition Instructions, Comparing Two Registers

Mnemonic	Description	Defined in
TGEI	Trap if Greater Than or Equal Immediate	II
TGEIU	Trap if Greater Than or Equal Unsigned Immediate	II
TLTI	Trap if Less Than Immediate	II
TLTIU	Trap if Less Than Unsigned Immediate	II
TEQI	Trap if Equal Immediate	II
TNEI	Trap if Not Equal Immediate	II

Table 1.18 Trap-on-Condition Instructions, Comparing an Immediate

Serialization Instructions

The order in which memory accesses from load and store instruction appears **outside** the processor executing them, such as in a multiprocessor system, is not specified by the architecture. The SYNC instruction creates a point in the executing instruction stream at which the relative order of some loads and stores is known. Loads and stores executed before the SYNC are completed before loads and stores after the SYNC can start.

Mnemonic	Description	Defined in
SYNC	Synchronize Shared Memory	II

Table 1.19 Serialization Instructions

Conditional Move Instructions

Instructions were added in MIPS IV to conditionally move one CPU general register to another, based on the value in a third general register.

Mnemonic	Description	Defined in
MOVN	Move Conditional on Not Zero	IV
MOVZ	Move Conditional on Zero	IV

Table 1.20 CPU Conditional Move Instructions

Prefetch Instructions

There are two prefetch advisory instructions: one with register+offset addressing (PREF) and the other with register+register addressing (PREFX). These instructions advise that memory is likely to be used in a particular way in the near future and should be prefetched into the cache. The PREFX instruction using register+register addressing mode is coded in the FPU opcode space, along with the other operations using register+register addressing. The RC32364 implements PREF instruction.

Mnemonic	Description	Defined in
PREF	Prefetch Indexed	IV

Table 1.21 Prefetch Using Register + Offset Address Mode

Mnemonic	Description	Defined in
PREFX	Prefetch Indexed	IV

Table 1.22 Prefetch Using Register + Register Address Mode

Coprocessor Instructions

Coprocessors are alternate execution units, with register files separate from the CPU. The MIPS architecture provides an abstraction for up to 4 coprocessor units, numbered 0 to 3. Each architecture level defines some of these coprocessors, as shown in Table 1.23.

Coprocessor 0 is always used for system control and coprocessor 1 is used for the floating-point unit. Other coprocessors are architecturally valid, but do not have a reserved use. Some coprocessors are not defined and their opcodes are either reserved or used for other purposes.

MIPS Architecture Level				
coprocessor	I	II	III	IV
0	Sys Control	Sys Control	Sys Control	Sys Control
1	FPU	FPU	FPU	FPU
2	unused	unused	unused	unused
3	unused	unused	not defined	FPU (COP 1X)

Table 1.23 Coprocessor Definition and Use in the MIPS Architecture

The coprocessors may have two register sets—coprocessor general registers and coprocessor control registers—each set containing up to thirty two registers. Coprocessor computational instructions may alter registers in either set.

System control for all MIPS processors is implemented as coprocessor 0 (CP0), the System Control Coprocessor. It provides the processor control, memory management, and exception handling functions. The CP0 instructions are specific to each CPU and are documented with the CPU-specific information.

If a system includes a floating-point unit, it is implemented as coprocessor 1 (CP1). In MIPS IV, the FPU also uses the computation opcode space for coprocessor unit 3, renamed COP1X. The FPU instructions are documented in Chapters 4, 5 and 6.

The coprocessor instructions are divided into these two main groups:

- ◆ *Load and store instructions that are reserved in the main opcode space.*
- ◆ *Coprocessor-specific operations that are defined entirely by the coprocessor.*

Coprocessor Load and Store Instructions

Load and store instructions are not defined for CP0; the move to/from coprocessor instructions are the only way to write and read the CP0 registers. The loads and stores for coprocessors are summarized on page 1-1.

Coprocessor Operations

There are up to four coprocessors and the instructions are shown generically for coprocessor-z. Within the operation main opcode, the coprocessor has further coprocessor-specific instructions encoded.

Mnemonic	Description	Defined in
COPz	Coprocessor-z Operation	I

Table 1.24 Coprocessor Operation Instructions

Memory Access Types

MIPS processors provide a few *memory access types* that are characteristic ways to use physical memory and caches to perform a memory access. The memory access type is specified as a cache coherence algorithm (CCA) in the CP0 descriptions of a virtual address. The access type used for a location is associated with the virtual address, not the physical address or the instruction making the reference. Implementations without multiprocessor (MP) support provide uncached and cached accesses. Implementations with MP support provide uncached, cached noncoherent and cached coherent accesses. The memory access types use the memory hierarchy as follows:

Uncached

Physical memory is used to resolve the access. Each reference causes a read or write to physical memory. Caches are neither examined nor modified.

Cached Noncoherent

Physical memory and the caches of the processor performing the access are used to resolve the access. Other caches are neither examined nor modified.

Cached Coherent

Physical memory and all caches in the system containing a coherent copy of the physical location are used to resolve the access. A copy of a location is coherent (noncoherent) if the copy was placed in the cache by a cached coherent (cached noncoherent) access. Caches containing a coherent copy of the location are examined and/or modified to keep the contents of the location coherent. It is unpredictable whether caches holding a noncoherent copy of the location are examined and/or modified during a cached coherent access.

Cached

For early 32-bit processors without MP support, cached is equivalent to cached noncoherent. If an instruction description mentions the cached noncoherent access type, the comment applies equally to the cached access type in a processor that has the cached access type.

For processors with MP support, cached is a collective term, e.g. "cached memory" or "cached access", that includes both cached noncoherent and cached coherent. Such a collective use does not imply that cached is an access type, it means that the statement applies equally to cached noncoherent and cached coherent access types.

Mixing References with Different Access Types

It is possible to have more than one virtual location simultaneously mapped to the same physical location. The memory access type that is used for virtual mappings may be different.

For all accesses to virtual locations with the **same** memory access type, a processor executing load and store instructions must observe the effect of those instructions to a physical location in the order that they occur in the instruction stream (such as program order).

If a processor executes a load or store using one access type to a physical location, the behavior of a subsequent load or store to the same location, using a different memory access type, is undefined unless a privileged instruction sequence is executed between the two accesses. Each implementation has a privileged implementation-specific mechanism that must be used to change the access type being used to access a location.

The 64-bit RISController family allows physical memory to be described simultaneously with different access characteristics, such as write-back and write-through. The caches are physically tagged, and provide sufficient state bites, to ensure memory coherency in a uniprocessor system.

The memory access type of a location affects the behavior of I-fetch, load, store, and prefetch operations to the location. In addition, memory access types affect some instruction descriptions. Load linked (LL, LLD) and store conditional (SC, SCD) have defined operation only for locations with cached memory access type. SYNC affects only load and stores made to locations with uncached or cached coherent memory access types.

Cache Coherence Algorithms and Access Types

The memory access types are specified by implementation-specific cache coherence algorithms (CCAs) in TLB entries. Slightly different cache coherence algorithms such as “cached coherent, update on write” and “cached coherent, exclusive on write” can map to the same memory access type, in this case they both map to cached coherent.

To map to the same access type, the fundamental mechanism of both CCAs must be the same. When it affects the operation of the instruction, the instructions are described in terms of the memory access types. The load and store operations in a processor proceeds according to the specific CCA of the reference, however, and the pseudocode for load and store common functions in the section “Load and Store Memory Functions” on page 1-18 use the CCA value rather than the corresponding memory access type.

Implementation-Specific Access Types

An implementation may provide memory access types other than uncached, cached noncoherent, or cached coherent. Implementation-specific documentation will define the properties of the new access types and their effect on all memory-related operations.

Instruction Descriptions

The CPU instructions are described in alphabetic order. Each description contains several sections that contain specific information about the instruction. The content of the section is described in detail below. An example description is shown in Figure 1.2.




Figure 1.2 Example Instruction Description

Instruction Mnemonic and Name

The instruction mnemonic and name are printed as page headings for each page in the instruction description.

Instruction Encoding Picture

The instruction word encoding is shown in pictorial form at the top of the instruction description. This picture shows the values of all constant fields and the opcode names for opcode fields in upper-case. It labels all variable fields with lower-case names that are used in the instruction description. Fields that contain zeroes but are not named are unused fields that are required to be zero. A summary of the instruction formats and a definition of the terms used to describe the contents can be found in **CPU Instruction Formats**.

Format

The assembler formats for the instruction and the architecture level at which the instruction was originally defined are shown. If the instruction definition was later extended, the architecture levels at which it was extended and the assembler formats for the extended definition are shown in order of extension. The MIPS architecture levels are inclusive; higher architecture levels include all instructions in previous levels. Extensions to instructions are backwards compatible. The original assembler formats are valid for the extended architecture.

The assembler format is shown with literal parts of the assembler instruction in upper-case characters. The variable parts, the operands, are shown as the lower-case names of the appropriate fields in the instruction encoding picture. The architecture level at which the instruction was first defined, e.g. "MIPS I", is shown at the right side of the page.

There can be more than one assembler format per architecture level. This is sometimes an alternate form of the instruction. Floating-point operations on formatted data show an assembly format with the actual assembler mnemonic for each valid value of the "fmt" field. For example the ADD.fmt instruction shows ADD.S and ADD.D.

The assembler format lines sometimes have comments to the right in parentheses to help explain variations in the formats. The comments are not a part of the assembler format.

Purpose

This section provides a short statement on the purpose of the instruction.

Description

If a one-line symbolic description of the instruction is feasible, it will appear immediately to the right of the *Description* heading. The main purpose is to show how fields in the instruction are used in the arithmetic or logical operation.

The body of the section is a description of the operation of the instruction in text, tables, and figures. This description complements the high-level language description in the *Operation* section.

This section uses acronyms for register descriptions. "GPR *rt*" is CPU General Purpose Register specified by the instruction field *rt*. "FPR *fs*" is the Floating Point Operand Register specified by the instruction field *fs*. "CP1 register *fd*" is the coprocessor 1 General Register specified by the instruction field *fd*. "FCSR" is the floating-point control and status register.

Restrictions

This section documents the restrictions on the instruction. Most restrictions fall into one of the following six categories:

- ◆ *The valid values for instruction fields (see floating-point ADD.fmt).*
- ◆ *The alignment requirements for memory addresses (see LW).*
- ◆ *The valid values of operands (see DADD).*
- ◆ *The valid operand formats (see floating-point ADD.fmt).*
- ◆ *The order of instructions necessary to guarantee correct execution.*
- ◆ *The valid memory access types (see LL/SC).*

These ordering constraints avoid pipeline hazards for which some processors do not have hardware interlocks (see MUL).

Operation

This section describes the operation of the instruction as pseudocode in a high-level language notation resembling Pascal. The purpose of this section is to describe the operation of the instruction clearly in a form with less ambiguity than prose. This formal description complements the *Description* section; it is not complete in itself because many of the restrictions are either difficult to include in the pseudocode or omitted for readability.

There will be separate *Operation* sections for 32-bit and 64-bit processors if the operation is different. This is usually necessary because the path to memory is a different size on these processors.

See "Operation Section Notation and Functions" on page 1-15 for more information on the formal notation.

Exceptions

This section lists the exceptions that can be caused by **operation** of the instruction. It omits exceptions that can be caused by instruction fetch, e.g. TLB Refill. It omits exceptions that can be caused by asynchronous external events, e.g. Interrupt. Although the Bus Error exception may be caused by the operation of a load or store instruction this section does not list Bus Error for load and store instructions because the relationship between load and store instructions and external error indications, like Bus Error, are implementation dependent.

Reserved Instruction is listed for every instruction not in MIPS I because the instruction will cause this exception on a MIPS I processor. To execute a MIPS II, MIPS III, or MIPS IV instruction, the processor must both support the architecture level and have it enabled. The mechanism to do this is implementation specific.

The mechanism used to signal a floating-point unit (FPU) exception is implementation specific. Some implementations use the exception named "Floating Point". Others use external interrupts (the Interrupt exception). This section lists Floating Point to represent all such mechanisms. The specific FPU traps are listed, indented, under the Floating Point entry.

An instruction may cause implementation-dependent exceptions that are not present in the *Exceptions* section.

Programming and Implementation Notes

These sections contain material that is useful for programmers and implementors respectively but that is not necessary to describe the instruction and does not belong in the description sections.

Operation Section Notation and Functions

In an instruction description, the *Operation* section describes the operation performed by each instruction using a high-level language notation. The contents of the *Operation* section are described here. The special symbols and functions used are documented here.

Pseudocode Language

Each of the high-level language statements is executed in sequential order (as modified by conditional and loop constructs).

Pseudocode Symbols

Special symbols used in the notation are described in Table 1.25.

Symbol	Meaning
..	Assignment.
=, ≠	Tests for equality and inequality.
	Bit string concatenation.
x^y	A y -bit string formed by y copies of the single-bit value x .
$x_{y..z}$	Selection of bits y through z of bit string x . Little-endian bit notation (rightmost bit is 0) is used. If y is less than z , this expression is an empty (zero length) bit string.
+, -	2's complement or floating-point arithmetic: addition, subtraction.
*, ¥	2's complement or floating-point multiplication (both used for either).
div	2's complement integer division.
mod	2's complement modulo.
/	Floating-point division.
<	2's complement less than comparison.
nor	Bit-wise logical NOR.
xor	Bit-wise logical XOR.
and	Bit-wise logical AND.
or	Bit-wise logical OR.

Table 1.25 Symbols in Instruction Operation Statements (Page 1 of 2)

Symbol	Meaning
GPRLen	The length in bits (32 or 64), of the CPU General Purpose Registers.
GPR[x]	CPU General Purpose Register x. The content of GPR[0] is always zero.
FPR[x]	Floating-Point operand register x.
FCC[cc]	Floating-Point condition code cc. FCC[0] has the same value as COC[1].
FGR[x]	Floating-Point (Coprocessor unit1), general register x.
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	Endian mode as configured at chip reset (0 → Little, 1 → 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 RE bit of the Status register. Thus, ReverseEndian may be computed as (SR _{RE} and User mode).
BigEndianCPU	The endianness for load and store instructions (0 → Little, 1 → Big). In User mode, this endianness may be switched by setting the RE bit in the Status Register. Thus, BigEndianCPU may be computed as (BigEndianMem XOR ReverseEndian).
LLbit	Bit of virtual state used to specify operation for instructions that provide atomic read-modify-write. It is set when a linked load occurs. It is tested and cleared by the conditional store. It is cleared, during other CPU operation, when a store to the location would no longer be atomic. In particular, it is cleared by exception return instructions.
I, I+n, I-n:	<p>This occurs as a prefix to operation description lines and functions as a label. It indicates the instruction time during which the effects of the pseudocode lines appears to occur (i.e. when the pseudocode is “executed”). Unless otherwise indicated, all effects of the current instruction appear to occur during the instruction time of the current instruction. No label is equivalent to a time label of “I:”. Sometimes effects of an instruction appear to occur either earlier or later – during the instruction time of another instruction. When that happens, the instruction operation is written in sections labelled with the instruction time, relative to the current instruction I, in which the effect of that pseudocode appears to occur. For example, an instruction may have a result that is not available until after the next instruction. Such an instruction will have the portion of the instruction operation description that writes the result register in a section labelled “I+1:”.</p> <p>The effect of pseudocode statements for the current instruction labelled “I+1:” appears to occur “at the same time” as the effect of pseudocode statements labelled “I:” for the following instruction. Within one pseudocode sequence the effects of the statements takes place in order. However, between sequences of statements for different instructions that occur “at the same time”, there is no order defined. Programs must not depend on a particular order of evaluation between such sections.</p>
PC	The Program Counter value. During the instruction time of an instruction this is the address of the instruction word. The address of the instruction that occurs during the next instruction time is determined by assigning a value to PC during an instruction time. If no value is assigned to PC during an instruction time by any pseudocode statement, it is automatically incremented by 4 before the next instruction time. A taken branch assigns the target address to PC during the instruction time of the instruction in the branch delay slot.
PSIZE	The SIZE, number of bits, of Physical address in an implementation.

Table 1.25 Symbols in Instruction Operation Statements (Page 2 of 2)

Pseudocode Functions

There are several functions used in the pseudocode descriptions. These are used either to make the pseudocode more readable, to abstract implementation specific behavior, or both. The functions are defined in this section.

Coprocessor General Register Access Functions

Defined coprocessors, except for CP0, have instructions to exchange words and doublewords between coprocessor general registers and the rest of the system. What a coprocessor does with a word or doubleword supplied to it and how a coprocessor supplies a word or doubleword is defined by the coprocessor itself. This behavior is abstracted into the following functions:

<p>COP_LW (z, rt, memword)</p> <p>z: The coprocessor unit number. rt: Coprocessor general register specifier. memword: A 32-bit word value supplied to the coprocessor.</p> <p>This is the action taken by coprocessor z when supplied with a word from memory during a load word operation. The action is coprocessor specific. The typical action would be to store the contents of <i>memword</i> in coprocessor general register <i>rt</i>.</p>
<p>COP_LD (z, rt, memdouble)</p> <p>z: The coprocessor unit number. rt: Coprocessor general register specifier. memdouble: 64-bit doubleword value supplied to the coprocessor.</p> <p>This is the action taken by coprocessor z when supplied with a doubleword from memory during a load doubleword operation. The action is coprocessor specific. The typical action would be to store the contents of <i>memdouble</i> in coprocessor general register <i>rt</i>.</p>
<p>dataword " COP_SW (z, rt)</p> <p>z: The coprocessor unit number. rt: Coprocessor general register specifier. dataword: 32-bit word value.</p> <p>This defines the action taken by coprocessor z to supply a word of data during a store word operation. The action is coprocessor specific. The typical action would be to supply the contents of the low-order word in coprocessor general register <i>rt</i>.</p>
<p>datadouble " COP_SD (z, rt)</p> <p>z: The coprocessor unit number. rt: Coprocessor general register specifier. datadouble: 64-bit doubleword value.</p> <p>This defines the action taken by coprocessor z to supply a doubleword of data during a store doubleword operation. The action is coprocessor specific. The typical action would be to supply the contents of the doubleword in coprocessor general register <i>rt</i>.</p>

Table 1.26 Coprocessor General Register Access Functions

Load and Store Memory Functions

Regardless of byte ordering (big- or little-endian), the address of a halfword, word, or doubleword is the smallest byte address among the bytes forming the object. For big-endian ordering this is the most-significant byte; for a little-endian ordering this is the least-significant byte.

In the operation description pseudocode for load and store operations, the functions shown below are used to summarize the handling of virtual addresses and accessing physical memory. The size of the data item to be loaded or stored is passed in the *AccessLength* field.

The valid constant names and values are shown in Table 1.27. The bytes within the addressed unit of memory (word for 32-bit processors or doubleword for 64-bit processors) can be determined directly from the *AccessLength* and the two or three low-order bits of the address.

(pAddr, CCA) ~ AddressTranslation (vAddr, lorD, LorS)

- pAddr: Physical Address.
- CCA: Cache Coherence Algorithm: the method used to access caches and memory and resolve the reference.
- vAddr: Virtual Address.
- lorD: Indicates whether access is for INSTRUCTION or DATA.
- LorS: Indicates whether access is for LOAD or STORE.

Translate a virtual address to a physical address and a cache coherence algorithm describing the mechanism used to resolve the memory reference.

Given the virtual address *vAddr*, and whether the reference is to Instructions or Data (*lorD*), find the corresponding physical address (*pAddr*) and the cache coherence algorithm (*CCA*) used to resolve the reference. If the virtual address is in one of the unmapped address spaces the physical address and *CCA* are determined directly by the virtual address. If the virtual address is in one of the mapped address spaces then the TLB is used to determine the physical address and access type; if the required translation is not present in the TLB or the desired access is not permitted the function fails and an exception is taken.

MemElem ~ LoadMemory (CCA, AccessLength, pAddr, vAddr, lorD)

- MemElem: Data is returned in a fixed width with a natural alignment. The width is the same size as the CPU general purpose register, 32 or 64 bits, aligned on a 32 or 64-bit boundary respectively.
- CCA: Cache Coherence Algorithm: the method used to access caches and memory and resolve the reference.
- AccessLength: Length, in bytes, of access.
- pAddr: Physical Address.
- vAddr: Virtual Address.
- lorD: Indicates whether access is for Instructions or Data.

Load a value from memory.

Uses the cache and main memory as specified in the Cache Coherence Algorithm (*CCA*) and the sort of access (*lorD*) to find the contents of *AccessLength* memory bytes starting at physical location *pAddr*. The data is returned in the fixed width naturally-aligned memory element (*MemElem*). The low-order two (or three) bits of the address and the *AccessLength* indicate which of the bytes within *MemElem* needs to be given to the processor. If the memory access type of the reference is uncached then only the referenced bytes are read from memory and valid within the memory element. If the access type is cached, and the data is not present in cache, an implementation specific size and alignment block of memory is read and loaded into the cache to satisfy a load reference. At a minimum, the block is the entire memory element.

StoreMemory (CCA, AccessLength, MemElem, pAddr, vAddr)

- CCA: Cache Coherence Algorithm: the method used to access caches and memory and resolve the reference.
- AccessLength: Length, in bytes, of access.
- MemElem: Data in the width and alignment of a memory element. The width is the same size as the CPU general purpose register, 4 or 8 bytes, aligned on a 4 or 8-byte boundary. For a partial-memory-element store, only the bytes that will be stored must be valid.
- pAddr: Physical Address.
- vAddr: Virtual Address.

Store a value to memory.

The specified data is stored into the physical location *pAddr* using the memory hierarchy (data caches and main memory) as specified by the Cache Coherence Algorithm (CCA). The *MemElem* contains the data for an aligned, fixed-width memory element (word for 32-bit processors, doubleword for 64-bit processors), though only the bytes that will actually be stored to memory need to be valid. The low-order two (or three) bits of *pAddr* and the *AccessLength* field indicates which of the bytes within the *MemElem* data should actually be stored; only these bytes in memory will be changed.

Prefetch (CCA, pAddr, vAddr, DATA, hint)

- CCA: Cache Coherence Algorithm: the method used to access caches and memory and resolve the reference.
- pAddr: physical Address.
- vAddr: Virtual Address.
- DATA: Indicates that access is for DATA.
- hint: hint that indicates the possible use of the data.

Prefetch data from memory.

Prefetch is an advisory instruction for which an implementation specific action is taken. The action taken may increase performance but must not change the meaning of the program or alter architecturally-visible state.

AccessLength Name	Value	Meaning
DOUBLEWORD	7	8 bytes (64 bits)
SEPTIBYTE	6	7 bytes (56 bits)
SEXTIBYTE	5	6 bytes (48 bits)
QUINTIBYTE	4	5 bytes (40 bits)
WORD	3	4 bytes (32 bits)
TRIPLEBYTE	2	3 bytes (24 bits)
HALFWORD	1	2 bytes (16 bits)
BYTE	0	1 byte (8 bits)

Table 1.27 AccessLength Specifications for Loads/Stores

Access Functions for Floating-Point Registers

The details of the relationship between CP1 general registers and floating-point operand registers is encapsulated in the functions included in this section. See **Valid Operands for FP Instructions** in the Chapter Titled “FPU Instruction Set” for more information.

This function returns the current logical width, in bits, of the CP1 general registers. All 32-bit processors will return “32”. 64-bit processors will return “32” when in 32-bit-CP1-register emulation mode and “64” when in native 64-bit mode.

The following pseudocode referring to the Status_{FR} bit is valid for all existing MIPS 64-bit processors at the time of this writing, however this is a privileged processor-specific mechanism and it may be different in some future processor.

```

SizeFGR() -- current size, in bits, of the CP1 general registers
size ``SizeFGR()
  if 32_bit_processor then
    size ``32
  else
    /* 64-bit processor */
    if StatusFR = 1 then
      size ``64
    else
      size ``32
    endif
  endif
endif

```

This pseudocode specifies how the unformatted contents loaded or moved-to CP1 registers are interpreted to form a formatted value. If an FPR contains a value in some format, rather than unformatted contents from a load (uninterpreted), it is valid to interpret the value in that format, but not to interpret it in a different format.

```

ValueFPR() -- Get a formatted value from an FPR.
value ``ValueFPR (fpr, fmt) /* get a formatted value from an FPR */
  if SizeFGR() = 64 then
    case fmt of
      S, W:
        value ``FGR[fpr]31..0
      D, L:
        value ``FGR[fpr]
    endcase
  elseif fpr0 = 0 then /* fpr is valid (even), 32-bit wide FGRs */
    case fmt of
      S, W:
        value ``FGR[fpr]
      D, L:
        value ``FGR[fpr+1] || FGR[fpr]
    endcase
  else /* undefined for odd 32-bit FGRs */
    UndefinedResult
  endif

```

This pseudocode specifies the way that a binary encoding representing a formatted value is stored into CP1 registers by a computational or move operation. This binary representation is visible to store or move-from instructions. Once an FPR contains a value via StoreFPR(), it is not valid to interpret the value with ValueFPR() in a different format.

```

StoreFPR() -- store a formatted value into an FPR.
StoreFPR(fpr, fmt, value):          /* place a formatted value into an FPR */
  if SizeFGR() = 64 then           /* 64-bit wide FGRs */
    case fmt of
      S, W:
        FGR[fpr] `` undefined32 || value
      D, L:
        FGR[fpr] `` value
    endcase
  elseif fpr0 = 0 then           /* fpr is valid (even), 32-bit wide FGRs */
    case fmt of
      S, W:
        FGR[fpr+1] `` undefined32
        FGR[fpr] `` value
      D, L:
        FGR[fpr+1] `` value63..32
        FGR[fpr] `` value31..0
    endcase
  else                             /* undefined for odd 32-bit FGRs */
    UndefinedResult
  endif
    
```

Miscellaneous Functions

<p>SyncOperation(stype) stype: Type of load/store ordering to perform. order loads and stores to synchronize shared memory. Perform the action necessary to make the effects of groups synchronizable loads and stores indicated by <i>stype</i> occur in the same order for all processors.</p>
<p>SignalException(Exception) Exception The exception condition that exists. Signal an exception condition. This will result in an exception that aborts the instruction. The instruction operation pseudocode will never see a return from this function call.</p>
<p>UndefinedResult() This function indicates that the result of the operation is undefined.</p>
<p>NullifyCurrentInstruction() Nullify the current instruction. This occurs during the instruction time for some instruction and that instruction is not executed further. This appears for branch-likely instructions during the execution of the instruction in the delay slot and it kills the instruction in the delay slot.</p>
<p>CoprocessorOperation (z, cop_fun) z Coprocessor unit number cop_fun Coprocessor function from function field of instruction Perform the specified Coprocessor operation.</p>

Individual CPU Instruction Descriptions


The user-mode CPU instructions are described in alphabetic order. See "Instruction Descriptions" on page 1-13 for a description of the information in each instruction description.



CPU Instruction Reference

Notes

This chapter contains the detailed reference material for each of the CPU instructions in alphabetical order. Each new instruction starts on a new page and the instruction mnemonic is easily locatable at the top of the page in large bold letters.



Format: ADD rd, rs, rt **MIPS I**

Purpose: To add 32-bit integers. If overflow occurs, then trap.

Description: $rd \leftarrow rs + rt$

The 32-bit word value in GPR *rt* is added to the 32-bit value in GPR *rs* to produce a 32-bit result. If the addition results in 32-bit 2's complement arithmetic overflow then the destination register is not modified and an Integer Overflow exception occurs. If it does not overflow, the 32-bit result is placed into GPR *rd*.

Restrictions:

On 64-bit processors, if either GPR *rt* or GPR *rs* do not contain sign-extended 32-bit values (bits 63..31 equal), then the result of the operation is undefined.

Operation:

```


if (NotWordValue(GPR[rs]) or NotWordValue(GPR[rt])) then UndefinedResult() endif
temp ← GPR[rs] + GPR[rt]
if (32_bit_arithmetic_overflow) then
    SignalException(IntegerOverflow)
else
    GPR[rd] ← sign_extend(temp31..0)
endif
    
```

Exceptions:

Integer Overflow

Programming Notes:

ADDU performs the same arithmetic operation but, does not trap on overflow.



Format: ADDI rt, rs, immediate **MIPS I**

Purpose: To add a constant to a 32-bit integer. If overflow occurs, then trap.

Description: $rt \leftarrow rs + \text{immediate}$

The 16-bit signed *immediate* is added to the 32-bit value in GPR *rs* to produce a 32-bit result. If the addition results in 32-bit 2's complement arithmetic overflow then the destination register is not modified and an Integer Overflow exception occurs. If it does not overflow, the 32-bit result is placed into GPR *rt*.

Restrictions:

On 64-bit processors, if GPR *rs* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is undefined.

Operation:

```


if (NotWordValue(GPR[rs])) then UndefinedResult() endif
temp ← GPR[rs] + sign_extend(immediate)
if (32_bit_arithmetic_overflow) then
    SignalException(IntegerOverflow)
else
    GPR[rt] ← sign_extend(temp31..0)
endif
    
```

Exceptions:

Integer Overflow

Programming Notes:

ADDIU performs the same arithmetic operation but, does not trap on overflow.



Format: ADDIU rt, rs, immediate MIPS I

Purpose: To add a constant to a 32-bit integer.

Description: $rt \leftarrow rs + \text{immediate}$

The 16-bit signed *immediate* is added to the 32-bit value in GPR *rs* and the 32-bit arithmetic result is placed into GPR *rt*.

No Integer Overflow exception occurs under any circumstances.

Restrictions:

On 64-bit processors, if GPR *rs* does not contain a sign-extended 32-bit value (bits 63..31 equal), then the result of the operation is undefined.

Operation:

```


if (NotWordValue(GPR[rs])) then UndefinedResult() endif
temp ← GPR[rs] + sign_extend(immediate)
GPR[rt] ← sign_extend(temp31..0)
    
```

Exceptions:

None

Programming Notes:

The term "unsigned" in the instruction name is a misnomer; this operation is 32-bit modulo arithmetic that does not trap on overflow. It is appropriate for arithmetic which is not signed, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as "C" language arithmetic.



Format: ADDU rd, rs, rt MIPS I

Purpose: To add 32-bit integers.

Description: $rd \leftarrow rs + rt$

The 32-bit word value in GPR *rt* is added to the 32-bit value in GPR *rs* and the 32-bit arithmetic result is placed into GPR *rd*.

No Integer Overflow exception occurs under any circumstances.

Restrictions:

On 64-bit processors, if either GPR *rt* or GPR *rs* do not contain sign-extended 32-bit values (bits 63..31 equal), then the result of the operation is undefined.

Operation:

```


if (NotWordValue(GPR[rs]) or NotWordValue(GPR[rt])) then UndefinedResult() endif
temp ← GPR[rs] + GPR[rt]
GPR[rd] ← sign_extend(temp31..0)
    
```

Exceptions:

None

Programming Notes:

The term “unsigned” in the instruction name is a misnomer; this operation is 32-bit modulo arithmetic that does not trap on overflow. It is appropriate for arithmetic which is not signed, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as “C” language arithmetic.



Format: AND rd, rs, rt MIPS I

Purpose: To do a bitwise logical AND.

Description: $rd \leftarrow rs \text{ AND } rt$

The contents of GPR *rs* are combined with the contents of GPR *rt* in a bitwise logical AND operation. The result is placed into GPR *rd*.

Restrictions:


None

Operation:

$GPR[rd] \leftarrow GPR[rs] \text{ and } GPR[rt]$

Exceptions:

None



Format: ANDI rt, rs, immediate MIPS I

Purpose: To do a bitwise logical AND with a constant.

Description: $rt \leftarrow rs \text{ AND } \text{immediate}$

The 16-bit *immediate* is zero-extended to the left and combined with the contents of GPR *rs* in a bitwise logical AND operation. The result is placed into GPR *rt*.

Restrictions:


None

Operation:

$GPR[rt] \leftarrow \text{zero_extend}(\text{immediate}) \text{ and } GPR[rs]$

Exceptions:

None



Format: BEQ rs, rt, offset MIPS I

Purpose: To compare GPRs then do a PC-relative conditional branch.

Description: if (rs = rt) then branch

An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.

If the contents of GPR *rs* and GPR *rt* are equal, branch to the effective target address after the instruction in the delay slot is executed.

Restrictions:

None

Operation:


```
I:  tgt_offset ← sign_extend(offset || 02)
      condition ← (GPR[rs] = GPR[rt])
I+1: if condition then
      PC ← PC + tgt_offset
      endif
```

Exceptions:

None

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.



Format: BEQL rs, rt, offset **MIPS II**
Purpose: To compare GPRs then do a PC-relative conditional branch; execute the delay slot only if the branch is taken.

Description: if (rs = rt) then branch_likely
 An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.


If the contents of GPR *rs* and GPR *rt* are equal, branch to the target address after the instruction in the delay slot is executed. If the branch is not taken, the instruction in the delay slot is not executed.

Restrictions:
 None

Operation:
 I: $\text{tgt_offset} \leftarrow \text{sign_extend}(\text{offset} \ll 2)$
 $\text{condition} \leftarrow (\text{GPR}[\text{rs}] = \text{GPR}[\text{rt}])$
 I+1: if condition then
 $\text{PC} \leftarrow \text{PC} + \text{tgt_offset}$
 else
 NullifyCurrentInstruction()
 endif

Exceptions:
 Reserved Instruction

Programming Notes:
 With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.



Format: BGEZ rs, offset MIPS I

Purpose: To test a GPR then do a PC-relative conditional branch.

Description: if ($rs \geq 0$) then branch

An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.

If the contents of GPR *rs* are greater than or equal to zero (sign bit is 0), branch to the effective target address after the instruction in the delay slot is executed.

Restrictions:

None

Operation:


I: $tgt_offset \leftarrow sign_extend(offset \ll 2)$
 $condition \leftarrow GPR[rs] \geq 0^{GPREN}$
 I+1: if condition then
 $PC \leftarrow PC + tgt_offset$
 endif

Exceptions:

None

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.



Format: BGEZAL rs, offset MIPS I

Purpose: To test a GPR then do a PC-relative conditional procedure call.

Description: if (rs ≥ 0) then procedure_call

Place the return address link in GPR 31. The return link is the address of the second instruction following the branch, where execution would continue after a procedure call.

An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.

If the contents of GPR *rs* are greater than or equal to zero (sign bit is 0), branch to the effective target address after the instruction in the delay slot is executed.

Restrictions:

GPR 31 must not be used for the source register *rs*, because such an instruction does not have the same effect when re-executed. The result of executing such an instruction is undefined. This restriction permits an exception handler to resume execution by re-executing the branch when an exception occurs in the branch delay slot.

Operation:

```


I:  tgt_offset ← sign_extend(offset || 02)
    condition ← GPR[rs] ≥ 0GPRLEN
    GPR[31] ← PC + 8
I+1: if condition then
      PC ← PC + tgt_offset
    endif
    
```

Exceptions:

None

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump and link (JAL) or jump and link register (JALR) instructions for procedure calls to more distant addresses.



Format: BGEZALL rs, offset **MIPS II**

Purpose: To test a GPR then do a PC-relative conditional procedure call; execute the delay slot only if the branch is taken.

Description: if (rs ≥ 0) then procedure_call_likely

Place the return address link in GPR 31. The return link is the address of the second instruction following the branch, where execution would continue after a procedure call.

An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.

If the contents of GPR *rs* are greater than or equal to zero (sign bit is 0), branch to the effective target address after the instruction in the delay slot is executed. If the branch is not taken, the instruction in the delay slot is not executed.

Restrictions:

GPR 31 must not be used for the source register *rs*, because such an instruction does not have the same effect when re-executed. The result of executing such an instruction is undefined. This restriction permits an exception handler to resume execution by re-executing the branch when an exception occurs in the branch delay slot.

Operation:


```
I: tgt_offset ← sign_extend(offset || 02)
   condition ← GPR[rs] ≥ 0GPRLEN
   GPR[31] ← PC + 8
I+1: if condition then
      PC ← PC + tgt_offset
   else
      NullifyCurrentInstruction()
   endif
```

Exceptions:

Reserved Instruction

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump and link (JAL) or jump and link register (JALR) instructions for procedure calls to more distant addresses.



Format: BGEZL rs, offset **MIPS II**
Purpose: To test a GPR then do a PC-relative conditional branch; execute the delay slot only if the branch is taken.

Description: if (rs ≥ 0) then branch_likely
 An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.


If the contents of GPR *rs* are greater than or equal to zero (sign bit is 0), branch to the effective target address after the instruction in the delay slot is executed. If the branch is not taken, the instruction in the delay slot is not executed.

Restrictions:
 None

Operation:
 I: $\text{tgt_offset} \leftarrow \text{sign_extend}(\text{offset} \ll 2)$
 $\text{condition} \leftarrow \text{GPR}[\text{rs}] \geq 0^{\text{GPRLEN}}$
 I+1: if condition then
 PC ← PC + tgt_offset
 else
 NullifyCurrentInstruction()
 endif

Exceptions:
 Reserved Instruction

Programming Notes:
 With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.



Format: BGTZ rs, offset **MIPS I**

Purpose: To test a GPR then do a PC-relative conditional branch.

Description: if (rs > 0) then branch

An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.

If the contents of GPR *rs* are greater than zero (sign bit is 0 but value not zero), branch to the effective target address after the instruction in the delay slot is executed.

Restrictions:

None

Operation:


```
I:  tgt_offset ← sign_extend(offset || 02)
      condition ← GPR[rs] > 0GPREN
I+1: if condition then
      PC ← PC + tgt_offset
      endif
```

Exceptions:

None

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.



Format: BGTZL rs, offset **MIPS II**
Purpose: To test a GPR then do a PC-relative conditional branch; execute the delay slot only if the branch is taken.

Description: if (rs > 0) then branch_likely
 An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.

If the contents of GPR *rs* are greater than zero (sign bit is 0 but value not zero), branch to the effective target address after the instruction in the delay slot is executed. If the branch is not taken, the instruction in the delay slot is not executed.


Restrictions:
 None

Operation:
 I: $\text{tgt_offset} \leftarrow \text{sign_extend}(\text{offset} \ll 0^2)$
 $\text{condition} \leftarrow \text{GPR}[\text{rs}] > 0^{\text{GPREN}}$
 I+1: if condition then
 PC \leftarrow PC + tgt_offset
 else
 NullifyCurrentInstruction()
 endif

Exceptions:
 Reserved Instruction

Programming Notes:
 With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.

BLEZ Branch on Less Than or Equal to Zero



Format: BLEZ rs, offset **MIPS I**

Purpose: To test a GPR then do a PC-relative conditional branch.

Description: if (rs ≤ 0) then branch

An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.

If the contents of GPR *rs* are less than or equal to zero (sign bit is 1 or value is zero), branch to the effective target address after the instruction in the delay slot is executed.

Restrictions:

None

Operation:

I: tgt_offset ← sign_extend(offset || 0²)
 condition ← GPR[rs] ≤ 0^{GPRLEN}

I+1: if condition then
 PC ← PC + tgt_offset


endif

Exceptions:

None

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.



Format: BLEZL rs, offset **MIPS II**
Purpose: To test a GPR then do a PC-relative conditional branch; execute the delay slot only if the branch is taken.

Description: if (rs ≤ 0) then branch_likely
 An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.


If the contents of GPR *rs* are less than or equal to zero (sign bit is 1 or value is zero), branch to the effective target address after the instruction in the delay slot is executed. If the branch is not taken, the instruction in the delay slot is not executed.

Restrictions:
 None

Operation:
 I: $\text{tgt_offset} \leftarrow \text{sign_extend}(\text{offset} \ll 2)$
 $\text{condition} \leftarrow \text{GPR}[\text{rs}] \leq 0^{\text{GPRLEN}}$
 I+1: if condition then
 PC ← PC + tgt_offset
 else
 NullifyCurrentInstruction()
 endif

Exceptions:
 Reserved Instruction

Programming Notes:
 With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.



Format: BLTZ rs, offset MIPS I

Purpose: To test a GPR then do a PC-relative conditional branch.

Description: if (rs < 0) then branch

An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.

If the contents of GPR *rs* are less than zero (sign bit is 1), branch to the effective target address after the instruction in the delay slot is executed.

Restrictions:

None

Operation:


I: $\text{tgt_offset} \leftarrow \text{sign_extend}(\text{offset} \ll 2)$
 $\text{condition} \leftarrow \text{GPR}[\text{rs}] < 0^{\text{GPREN}}$
 I+1: if condition then
 $\text{PC} \leftarrow \text{PC} + \text{tgt_offset}$
 endif

Exceptions:

None

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.



Format: BLTZAL rs, offset MIPS I

Purpose: To test a GPR then do a PC-relative conditional procedure call.

Description: if (rs < 0) then procedure_call

Place the return address link in GPR 31. The return link is the address of the second instruction following the branch (**not** the branch itself), where execution would continue after a procedure call.

An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch, in the branch delay slot, to form a PC-relative effective target address.

If the contents of GPR *rs* are less than zero (sign bit is 1), branch to the effective target address after the instruction in the delay slot is executed.

Restrictions:

GPR 31 must not be used for the source register *rs*, because such an instruction does not have the same effect when re-executed. The result of executing such an instruction is undefined. This restriction permits an exception handler to resume execution by re-executing the branch when an exception occurs in the branch delay slot.

Operation:

```


I:  tgt_offset ← sign_extend(offset || 02)
      condition ← GPR[rs] < 0GPRLEN
      GPR[31] ← PC + 8
I+1: if condition then
      PC ← PC + tgt_offset
      endif
    
```

Exceptions:

None

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump and link (JAL) or jump and link register (JALR) instructions for procedure calls to more distant addresses.



Format: BLTZALL rs, offset **MIPS II**

Purpose: To test a GPR then do a PC-relative conditional procedure call; execute the delay slot only if the branch is taken.

Description: if (rs < 0) then procedure_call_likely

Place the return address link in GPR 31. The return link is the address of the second instruction following the branch (**not** the branch itself), where execution would continue after a procedure call.

An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch, in the branch delay slot, to form a PC-relative effective target address.

If the contents of GPR *rs* are less than zero (sign bit is 1), branch to the effective target address after the instruction in the delay slot is executed. If the branch is not taken, the instruction in the delay slot is not executed.

Restrictions:

GPR 31 must not be used for the source register *rs*, because such an instruction does not have the same effect when re-executed. The result of executing such an instruction is undefined. This restriction permits an exception handler to resume execution by re-executing the branch when an exception occurs in the branch delay slot.

Operation:


```
I: tgt_offset ← sign_extend(offset || 02)
   condition ← GPR[rs] < 0GPREN
   GPR[31] ← PC + 8
I+1: if condition then
      PC ← PC + tgt_offset
   else
      NullifyCurrentInstruction()
   endif
```

Exceptions:

Reserved Instruction

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump and link (JAL) or jump and link register (JALR) instructions for procedure calls to more distant addresses.



Format: BLTZ rs, offset **MIPS II**

Purpose: To test a GPR then do a PC-relative conditional branch; execute the delay slot only if the branch is taken.

Description: if (rs < 0) then branch_likely

An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.

If the contents of GPR *rs* are less than zero (sign bit is 1), branch to the effective target address after the instruction in the delay slot is executed. If the branch is not taken, the instruction in the delay slot is not executed.

Restrictions:

None

Operation:

```
I:   tgt_offset ← sign_extend(offset || 02)
      condition ← GPR[rs] < 0GPRLEN
```


```
I+1: if condition then
       PC ← PC + tgt_offset
     else
       NullifyCurrentInstruction()
     endif
```

Exceptions:

Reserved Instruction

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.



Format: BNE rs, rt, offset **MIPS I**

Purpose: To compare GPRs then do a PC-relative conditional branch.

Description: if (rs ≠ rt) then branch

An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.

If the contents of GPR *rs* and GPR *rt* are not equal, branch to the effective target address after the instruction in the delay slot is executed.

Restrictions:

None

Operation:

```


I:  tgt_offset ← sign_extend(offset || 0²)
    condition ← (GPR[rs] ≠ GPR[rt])
I+1: if condition then
      PC ← PC + tgt_offset
    endif
  
```

Exceptions:

None

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.



Format: BNEL rs, rt, offset **MIPS II**

Purpose: To compare GPRs then do a PC-relative conditional branch; execute the delay slot only if the branch is taken.

Description: if (rs ≠ rt) then branch_likely

An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.

If the contents of GPR *rs* and GPR *rt* are not equal, branch to the effective target address after the instruction in the delay slot is executed. If the branch is not taken, the instruction in the delay slot is not executed.

Restrictions:

None

Operation:


```
I: tgt_offset ← sign_extend(offset || 02)
    condition ← (GPR[rs] ≠ GPR[rt])
I+1: if condition then
        PC ← PC + tgt_offset
    else
        NullifyCurrentInstruction()
    endif
```

Exceptions:

Reserved Instruction

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.



Format: BREAK MIPS I

Purpose: To cause a Breakpoint exception.

Description:

A breakpoint 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.

Restrictions:

None


Operation:

SignalException(Breakpoint)

Exceptions:

Breakpoint

CACHE op, offset(base)



Format: CACHE op, 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 virtual address is translated to a physical address, and the 5-bit sub-opcode specifies a cache operation for that address.

If CP0 is not usable (User or Supervisor mode) the CP0 enable bit in the *Status* register is clear, and a coprocessor unusable exception is taken. The operation of this instruction on any operation/cache combination not listed below is undefined. The operation of this instruction on uncached addresses is also undefined.

The 64-bit RISController family uses only the tag comparisons, not the valid bits, to choose which data it supplies to the instruction unit. This makes it important that the tags of the A and B sets are never the same.

The Index operation uses part of the virtual address to specify a cache block and set access as shown in Table 2.28

64-bit RISController-Family Processor	Set Selection	Block Selection
RC4640/RC4650	VAddr ₁₂	11..5
RC4700	VAddr ₁₃	12..5
RC5000	VAddr ₁₄	13..5
Note: TagLo[12] is the valid bit and TagLo[31:15] is the tag for all secondary cache operations.		

Table 2.28 64-bit RISController Family Primary Cache Indexing

Index Load Tag also uses vAddr_{4,3} to select the doubleword for reading parity. When the CE bit of the Status register is set, Hit WriteBack, Hit WriteBack Invalidate, Index WriteBack Invalidate, and Fill also use vAddr_{4,3} to select the doubleword that has its parity modified. This operation is performed unconditionally.

The Hit operation accesses the specified cache as normal data references, and performs the specified operation if the cache block contains valid data with the specified physical address (a hit). If both sets are invalid or contain different addresses (a miss), no operation is performed.

Write back from a primary cache goes to memory. The address to be written is specified by the cache tag and not the translated physical address.

For Index operations (where the physical address is used to index the cache but need not match the cache tag), unmapped addresses may be used to avoid exceptions.

This operation will never cause Virtual Coherency exceptions.

Bits 17..16 of the instruction specify the cache as follows:

Code	Name	Cache
0	I	Primary instruction
1	D	Primary data
2	NA	Undefined
3	SC	Secondary Cache (RV5000)

Bits 20..18 (this value is listed under the **Code** column) of the instruction specify the operation as follows:

Code	Caches	Name	Operation
0	I	Index Invalidate	Set the cache state of the cache block to Invalid. Index_Invalidate_I writes the physical address of the cache op into the tag when it clears the valid bit, which is different from the RC4000.
0	D	Index WriteBack Invalidate	Examine the cache state and W bit of the primary data cache block at the 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.
0	SC	Cache Clear	Generate a valid clear sequence to flush the entire cache in one operation.
1	I, D	Index Load Tag	Read the tag for the cache block at the specified index and place it into the TagLo CPO registers, ignoring parity errors. Also load the data parity bits into the ECC register.
1	SC	Index Load Tag	Read the secondary cache for the specified index and places it into the TagLo CPO register.
2	I, D	Index Store Tag	Write the tag for the cache block at the specified index from the TagLo and TagHi CPO registers.
2	SC	Index Store Tag	Write the secondary cache for the specified index from the physical address generated by the CACHE instruction.
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 cache block 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.
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.
5	I	Fill	Fill the primary instruction cache block from memory. If the CE bit of the Status register is set, the contents of the ECC register is used instead of the computed parity bits for addressed doubleword when written to the instruction cache.
5	SC	Cache Page Invalidate	Flush 128 lines of the cache in one operation with the tag value from the TagLo CPO register. The index for the cache page invalidate must be page aligned. Interrupts are deferred until a cache page invalidate instruction completes (up to 512 processor clocks for a SysClock ratio of 4).


Code	Caches	Name	Operation
6	D	<i>Hit WriteBack</i>	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.
6	I	<i>Hit WriteBack</i>	If the cache block contains the specified address, write back the data unconditionally.

Operation:

T: $vAddr \leftarrow ((offset_{15})^{48} \parallel offset_{15..0}) + GPR[base]$
 $(pAddr, uncached) \leftarrow AddressTranslation(vAddr, DATA)$
 CacheOp (op, vAddr, pAddr)

Exceptions:

Coprocessor unusable exception



Format: CFC1 rt, fs MIPS I

Purpose: To copy a word from an FPU control register to a GPR.

Description: $rt \leftarrow FP_Control[fs]$

Copy the 32-bit word from FP (coprocessor 1) control register *fs* into GPR *rt*, sign-extending it if the GPR is 64 bits.

Restrictions:

There are only a couple control registers defined for the floating-point unit. The result is not defined if *fs* specifies a register that does not exist.

For MIPS I, MIPS II, and MIPS III, the contents of GPR *rt* are undefined for the instruction immediately following CFC1.

Operation: MIPS I - III


- I: temp $\leftarrow FCR[fs]$
- I+1: GPR[rt] $\leftarrow sign_extend(temp)$

Operation: MIPS IV

- temp $\leftarrow FCR[fs]$
- GPR[rt] $\leftarrow sign_extend(temp)$

Exceptions:

Coprocessor Unusable



Format: CLO rt, rs **RC32364**

Description:

The RC32364 adds this new instruction. The content of general register rs is scanned from most significant bit to least significant bit, the number of leading ones is written into general register rt. If no bits were cleared in general register rs, i.e. rs=0xffffffff, the content of general register rt is 32.

Operation:


T: rt <-- Leading_ones(rs)

Exceptions:

None

Programming Notes:

This is an IDT proprietary extension.



Format: CLZ rt, rs **RC32364**

Description:

The RC32364 adds this new instruction. The content of general register rs is scanned from most significant bit to least significant bit, the number of leading zeros is written into general register rt. If no bits were set in general register rs, i.e. rs=0x0, the content of general register rt is 32.

Operation:


T: rt ← Leading_zeros(rs)

Exceptions:

None

Programming Notes:

This is an IDT proprietary extension.



Format: COP0 cop_fun MIPS I
 COP1 cop_fun
 COP2 cop_fun
 COP3 cop_fun

Purpose: To execute a coprocessor instruction.

Description:

The coprocessor operation specified by *cop_fun* is performed by coprocessor unit *zz*. Details of coprocessor operations must be found in the specification for each coprocessor.

Each MIPS architecture level defines up to 4 coprocessor units, numbered 0 to 3. The opcodes corresponding to coprocessors that are not defined by an architecture level may be used for other instructions.

Restrictions:

Access to the coprocessors is controlled by system software. Each coprocessor has a "coprocessor usable" bit in the System Control coprocessor. The usable bit must be set for a user program to execute a coprocessor instruction. If the usable bit is not set, an attempt to execute the instruction will result in a Coprocessor Unusable exception. An unimplemented coprocessor must never be enabled. The result of executing this instruction for an unimplemented coprocessor when the usable bit is set, is undefined.


See specification for the specific coprocessor being programmed.

Operation:

CoprocessorOperation (z, cop_fun)

Exceptions:

- Reserved Instruction
- Coprocessor Unusable
- Coprocessor interrupt or Floating-Point Exception (CP1 only for some processors)



Format: CTC1 rt, fs MIPS I

Purpose: To copy a word from a GPR to an FPU control register.

Description: FP_Control[fs] ← rt

Copy the low word from GPR *rt* into FP (coprocessor 1) control register *fs*.

Writing to control register 31, the *Floating-Point Control and Status Register* or FCSR, causes the appropriate exception if any cause bit and its corresponding enable bit are both set. The register will be written before the exception occurs.

Restrictions:

There are only a couple control registers defined for the floating-point unit. The result is not defined if *fs* specifies a register that does not exist.

For MIPS I, MIPS II, and MIPS III, the contents of floating-point control register *fs* are undefined for the instruction immediately following CTC1.

Operation: MIPS I - III


- I: temp ← GPR[rt]_{31..0}
- I+1: FCR[fs] ← temp
- COC[1] ← FCR[31]₂₃

Operation: MIPS IV

- temp ← GPR[rt]_{31..0}
- FCR[fs] ← temp
- COC[1] ← FCR[31]₂₃

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Unimplemented Operation
 - Invalid Operation
 - Division-by-zero
 - Inexact
 - Overflow
 - Underflow



Format: DADD rd, rs, rt MIPS III

Purpose: To add 64-bit integers. If overflow occurs, then trap.

Description: $rd \leftarrow rs + rt$

The 64-bit doubleword value in GPR *rt* is added to the 64-bit value in GPR *rs* to produce a 64-bit result. If the addition results in 64-bit 2's complement arithmetic overflow then the destination register is not modified and an Integer Overflow exception occurs. If it does not overflow, the 64-bit result is placed into GPR *rd*.

Restrictions:

None

Operation: 64-bit processors


```
temp ← GPR[rs] + GPR[rt]
if (64_bit_arithmetic_overflow) then
    SignalException(IntegerOverflow)
else
    GPR[rd] ← temp
endif
```

Exceptions:

- Integer Overflow
- Reserved Instruction

Programming Notes:

DADDU performs the same arithmetic operation but, does not trap on overflow.



Format: DADDI rt, rs, immediate MIPS III

Purpose: To add a constant to a 64-bit integer. If overflow occurs, then trap.

Description: $rt \leftarrow rs + \text{immediate}$

The 16-bit signed *immediate* is added to the 64-bit value in GPR *rs* to produce a 64-bit result. If the addition results in 64-bit 2's complement arithmetic overflow then the destination register is not modified and an Integer Overflow exception occurs. If it does not overflow, the 64-bit result is placed into GPR *rt*.

Restrictions:

None

Operation: 64-bit processors


```
temp ← GPR[rs] + sign_extend(immediate)
if (64_bit_arithmetic_overflow) then
    SignalException(IntegerOverflow)
else
    GPR[rt] ← temp
endif
```

Exceptions:

- Integer Overflow
- Reserved Instruction

Programming Notes:

DADDIU performs the same arithmetic operation but, does not trap on overflow.



Format: DADDIU rt, rs, immediate MIPS III

Purpose: To add a constant to a 64-bit integer.

Description: $rt \leftarrow rs + \text{immediate}$

The 16-bit signed *immediate* is added to the 64-bit value in GPR *rs* and the 64-bit arithmetic result is placed into GPR *rt*.

No Integer Overflow exception occurs under any circumstances.

Restrictions:

None

Operation: 64-bit processors


$$GPR[rt] \leftarrow GPR[rs] + \text{sign_extend}(\text{immediate})$$

Exceptions:

Reserved Instruction

Programming Notes:

The term "unsigned" in the instruction name is a misnomer; this operation is 64-bit modulo arithmetic that does not trap on overflow. It is appropriate for arithmetic which is not signed, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as "C" language arithmetic.



Format: DADDU rd, rs, rt

MIPS III

Purpose: To add 64-bit integers.

Description: $rd \leftarrow rs + rt$

The 64-bit doubleword value in GPR *rt* is added to the 64-bit value in GPR *rs* and the 64-bit arithmetic result is placed into GPR *rd*.

No Integer Overflow exception occurs under any circumstances.

Restrictions:

None

Operation: 64-bit processors


$$GPR[rd] \leftarrow GPR[rs] + GPR[rt]$$

Exceptions:

Reserved Instruction

Programming Notes:

The term “unsigned” in the instruction name is a misnomer; this operation is 64-bit modulo arithmetic that does not trap on overflow. It is appropriate for arithmetic which is not signed, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as “C” language arithmetic.



Format: DDIV rs, rt **MIPS III**

Purpose: To divide 64-bit signed integers.

Description: (LO, HI) ← rs / rt

The 64-bit doubleword in GPR *rs* is divided by the 64-bit doubleword in GPR *rt*, treating both operands as signed values. The 64-bit quotient is placed into special register *LO* and the 64-bit remainder is placed into special register *HI*.

No arithmetic exception occurs under any circumstances.

Restrictions:

If either of the two preceding instructions is MFHI or MFLO, the result of the MFHI or MFLO is undefined. Reads of the *HI* or *LO* special registers must be separated from subsequent instructions that write to them by two or more other instructions.

If the divisor in GPR *rt* is zero, the arithmetic result value is undefined.

Operation: 64-bit processors


- I-2, I-1: LO, HI ← undefined
- I: LO ← GPR[rs] div GPR[rt]
- HI ← GPR[rs] mod GPR[rt]

Exceptions:

Reserved Instruction

Programming Notes:

See the Programming Notes for the DIV instruction.



Format: DDIVU rs, rt MIPS III

Purpose: To divide 64-bit unsigned integers.

Description: (LO, HI) ← rs / rt

The 64-bit doubleword in GPR *rs* is divided by the 64-bit doubleword in GPR *rt*, treating both operands as unsigned values. The 64-bit quotient is placed into special register *LO* and the 64-bit remainder is placed into special register *HI*.

No arithmetic exception occurs under any circumstances.

Restrictions:

If either of the two preceding instructions is MFHI or MFLO, the result of the MFHI or MFLO is undefined. Reads of the *HI* or *LO* special registers must be separated from subsequent instructions that write to them by two or more other instructions.

If the divisor in GPR *rt* is zero, the arithmetic result value is undefined.

Operation: 64-bit processors


I-2, I-1: LO, HI ← undefined
 I: LO ← (0 || GPR[rs]) div (0 || GPR[rt])
 HI ← (0 || GPR[rs]) mod (0 || GPR[rt])

Exceptions:

Reserved instruction

Programming Notes:

See the Programming Notes for the DIV instruction.



Format: DIV rs, rt **MIPS I**

Purpose: To divide 32-bit signed integers.

Description: (LO, HI) ← rs / rt

The 32-bit word value in GPR *rs* is divided by the 32-bit value in GPR *rt*, treating both operands as signed values. The 32-bit quotient is placed into special register *LO* and the 32-bit remainder is placed into special register *HI*.

No arithmetic exception occurs under any circumstances.

Restrictions:

On 64-bit processors, if either GPR *rt* or GPR *rs* do not contain sign-extended 32-bit values (bits 63..31 equal), then the result of the operation is undefined.

If either of the two preceding instructions is MFHI or MFLO, the result of the MFHI or MFLO is undefined. Reads of the *HI* or *LO* special registers must be separated from subsequent instructions that write to them by two or more other instructions.

If the divisor in GPR *rt* is zero, the arithmetic result value is undefined.

Operation:

if (NotWordValue(GPR[rs]) or NotWordValue(GPR[rt])) then UndefinedResult() endif

I-2, I-1: LO, HI ← undefined

I: q ← GPR[rs]_{31..0} div GPR[rt]_{31..0}

LO ← sign_extend(q_{31..0})

r ← GPR[rs]_{31..0} mod GPR[rt]_{31..0}

HI ← sign_extend(r_{31..0})

Exceptions:


None

Programming Notes:

In some processors the integer divide operation may proceed asynchronously and allow other CPU instructions to execute before it is complete. An attempt to read *LO* or *HI* before the results are written will wait (interlock) until the results are ready. Asynchronous execution does not affect the program result, but offers an opportunity for performance improvement by scheduling the divide so that other instructions can execute in parallel.

No arithmetic exception occurs under any circumstances. If divide-by-zero or overflow conditions should be detected and some action taken, then the divide instruction is typically followed by additional instructions to check for a zero divisor and/or for overflow. If the divide is asynchronous then the zero-divisor check can execute in parallel with the divide. The action taken on either divide-by-zero or overflow is either a convention within the program itself or more typically, the system software; one possibility is to take a BREAK exception with a code field value to signal the problem to the system software.

As an example, the C programming language in a UNIX environment expects division by zero to either terminate the program or execute a program-specified signal handler. C does not expect overflow to cause any exceptional condition. If the C compiler uses a divide instruction, it also emits code to test for a zero divisor and execute a BREAK instruction to inform the operating system if one is detected.



Format: DIVU rs, rt

MIPS I

Purpose: To divide 32-bit unsigned integers.

Description: (LO, HI) \leftarrow rs / rt

The 32-bit word value in GPR *rs* is divided by the 32-bit value in GPR *rt*, treating both operands as unsigned values. The 32-bit quotient is placed into special register *LO* and the 32-bit remainder is placed into special register *HI*.

No arithmetic exception occurs under any circumstances.

Restrictions:

On 64-bit processors, if either GPR *rt* or GPR *rs* do not contain sign-extended 32-bit values (bits 63..31 equal), then the result of the operation is undefined.

If either of the two preceding instructions is MFHI or MFLO, the result of the MFHI or MFLO is undefined. Reads of the *HI* or *LO* special registers must be separated from subsequent instructions that write to them, like this one, by two or more other instructions.

If the divisor in GPR *rt* is zero, the arithmetic result is undefined.

Operation:

if (NotWordValue(GPR[rs]) or NotWordValue(GPR[rt])) then UndefinedResult() endif

I-2, I-1: LO, HI \leftarrow undefined

I: q \leftarrow (0 || GPR[rs]_{31..0}) div (0 || GPR[rt]_{31..0})

LO \leftarrow sign_extend(q_{31..0})

r \leftarrow (0 || GPR[rs]_{31..0}) mod (0 || GPR[rt]_{31..0})


HI \leftarrow sign_extend(r_{31..0})

Exceptions:

None

Exceptions:

See the programming Notes for the DIV instruction.



Format: DMFCO *rt*, *rd*

RC5000

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 DMFCO on a 32-bit coprocessor 0 register is undefined.


Operation:

T: data ← CPR[0,*rd*]
T+1: GPR[*rt*] ← data

Exceptions:

Coprocessor unusable exception

Reserved instruction exception for supervisor mode with Status.SX = 0 or user mode with Status.UX = 0.



Format: DMTC0 rt, rd

RC5000

Description:

The contents of general register *rt* are loaded into coprocessor register *rd* of the CP0.

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 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 instructions and store instructions immediately prior to and after this instruction are undefined.

Operation:

T: data ← GPR[rt]
 T+1: CPR[0,rd] ← data

Exceptions:

Reserved instruction exception for supervisor mode with Status.SX = 0 or user mode with Status.UX = 0.



Format: DMFC1 *rt*, *fs* **MIPS III**

Purpose: To copy a doubleword from an FPR to a GPR.

Description: $rt \leftarrow fs$

The doubleword contents of FPR *fs* are placed into GPR *rt*.

If the coprocessor 1 general registers are 32-bits wide (a native 32-bit processor or 32-bit register emulation mode in a 64-bit processor), FPR *fs* is held in an even/odd register pair. The low word is taken from the even register *fs* and the high word is from *fs+1*.

Restrictions:

If *fs* does not specify an FPR that can contain a doubleword, the result is undefined; see the Chapter titled "FPU Instruction Set" for information on FP registers, etc..

For MIPS III, the contents of GPR *rt* are undefined for the instruction immediately following DMFC1.

Operation: MIPS I - III

```

I:  if SizeFGR() = 64 then          /* 64-bit wide FGRs */
    data ← FGR[fs]
    elseif fs0 = 0 then          /* valid specifier, 32-bit wide FGRs */
    data ← FGR[fs+1] || FGR[fs]
    else                            /* undefined for odd 32-bit FGRs */
    UndefinedResult()
    endif
I+1: GPR[rt] ← data
    
```


Operation: MIPS IV

```

if SizeFGR() = 64 then          /* 64-bit wide FGRs */
    data ← FGR[fs]
elseif fs0 = 0 then          /* valid specifier, 32-bit wide FGRs */
    data ← FGR[fs+1] || FGR[fs]
else                            /* undefined for odd 32-bit FGRs */
    UndefinedResult()
endif
GPR[rt] ← data
    
```

Exceptions:

- Reserved Instruction
- Coprocessor Unusable



Format: DMTC1 *rt*, *fs* MIPS III

Purpose: To copy a doubleword from a GPR to an FPR.

Description: $fs \leftarrow rt$

The doubleword contents of GPR *rt* are placed into FPR *fs*.

If coprocessor 1 general registers are 32-bits wide (a native 32-bit processor or 32-bit register emulation mode in a 64-bit processor), FPR *fs* is held in an even/odd register pair. The low word is placed in the even register *fs* and the high word is placed in *fs*+1.

Restrictions:

If *fs* does not specify an FPR that can contain a doubleword, the result is undefined; see the Chapter Titled "FPU Instruction Set" for information about FP registers, etc.

For MIPS III, the contents of FPR *fs* are undefined for the instruction immediately following DMTC1.

Operation: MIPS I - III

```

I: data ← GPR[rt]
I+1: if SizeFGR() = 64 then /* 64-bit wide FGRs */
      FGR[fs] ← data
    elseif fs0 = 0 then /* valid specifier, 32-bit wide FGRs */
      FGR[fs+1] ← data63..32
      FGR[fs] ← data31..0
    else /* undefined result for odd 32-bit FGRs */
      UndefinedResult()
    endif
  
```


Operation: MIPS IV

```

data ← GPR[rt]
if SizeFGR() = 64 then /* 64-bit wide FGRs */
  FGR[fs] ← data
elseif fs0 = 0 then /* valid specifier, 32-bit wide FGRs */
  FGR[fs+1] ← data63..32
  FGR[fs] ← data31..0
else /* undefined result for odd 32-bit FGRs */
  UndefinedResult()
endif
  
```

Exceptions:

- Reserved Instruction
- Coprocessor Unusable



Format: DMULT rs, rt MIPS III

Purpose: To multiply 64-bit signed integers.

Description: (LO, HI) ← rs × rt

The 64-bit doubleword value in GPR *rt* is multiplied by the 64-bit value in GPR *rs*, treating both operands as signed values, to produce a 128-bit result. The low-order 64-bit doubleword of the result is placed into special register *LO*, and the high-order 64-bit doubleword is placed into special register *HI*.

No arithmetic exception occurs under any circumstances.

Restrictions:

If either of the two preceding instructions is MFHI or MFLO, the result of the MFHI or MFLO is undefined. Reads of the *HI* or *LO* special registers must be separated from subsequent instructions that write to them by two or more other instructions.

Operation: 64-bit processors

- I-2, I-1: LO, HI ← undefined
- I: prod ← GPR[rs] * GPR[rt]
- LO ← prod_{63..0}
- HI ← prod_{127..64}


Exceptions:

Reserved Instruction

Programming Notes:

In some processors the integer multiply operation may proceed asynchronously and allow other CPU instructions to execute before it is complete. An attempt to read *LO* or *HI* before the results are written will wait (interlock) until the results are ready. Asynchronous execution does not affect the program result, but offers an opportunity for performance improvement by scheduling the multiply so that other instructions can execute in parallel.

Programs that require overflow detection must check for it explicitly.



Format: DMULTU rs, rt MIPS III

Purpose: To multiply 64-bit unsigned integers.

Description: (LO, HI) ← rs × rt

The 64-bit doubleword value in GPR *rt* is multiplied by the 64-bit value in GPR *rs*, treating both operands as unsigned values, to produce a 128-bit result. The low-order 64-bit doubleword of the result is placed into special register *LO*, and the high-order 64-bit doubleword is placed into special register *HI*.

No arithmetic exception occurs under any circumstances.

Restrictions:


If either of the two preceding instructions is MFHI or MFLO, the result of the MFHI or MFLO is undefined. Reads of the *HI* or *LO* special registers must be separated from subsequent instructions that write to them by two or more other instructions.

Operation: 64-bit processors

I-2, I-1: LO, HI ← undefined
 I: prod ← (0 || GPR[rs]) * (0 || GPR[rt])
 LO ← prod_{63..0}
 HI ← prod_{127..64}

Exceptions:

Reserved Instruction



Format: DSLL rd, rt, sa MIPS III

Purpose: To left shift a doubleword by a fixed amount — 0 to 31 bits.

Description: $rd \leftarrow rt \ll sa$

The 64-bit doubleword contents of GPR *rt* are shifted left, inserting zeros into the emptied bits; the result is placed in GPR *rd*. The bit shift count in the range 0 to 31 is specified by *sa*.

Restrictions:

None


Operation: 64-bit processors

$$s \leftarrow 0 \parallel sa$$

$$GPR[rd] \leftarrow GPR[rt]_{(63-s)..0} \parallel 0^s$$

Exceptions:

Reserved Instruction



Format: DSLL32 rd, rt, sa MIPS III

Purpose: To left shift a doubleword by a fixed amount — 32 to 63 bits.

Description: $rd \leftarrow rt \ll (sa+32)$

The 64-bit doubleword contents of GPR *rt* are shifted left, inserting zeros into the emptied bits; the result is placed in GPR *rd*. The bit shift count in the range 32 to 63 is specified by *sa*+32.

Restrictions:

None


Operation: 64-bit processors

$$s \leftarrow 1 \parallel sa \quad /* 32+sa */$$

$$GPR[rd] \leftarrow GPR[rt]_{(63-s)..0} \parallel 0^s$$

Exceptions:

Reserved Instruction



Format: DSLLV rd, rt, rs MIPS III

Purpose: To left shift a doubleword by a variable number of bits.

Description: $rd \leftarrow rt \ll rs$

The 64-bit doubleword contents of GPR *rt* are shifted left, inserting zeros into the emptied bits; the result is placed in GPR *rd*. The bit shift count in the range 0 to 63 is specified by the low-order six bits in GPR *rs*.

Restrictions:

None


Operation: 64-bit processors

$$s \leftarrow 0 \parallel \text{GPR}[rs]_{5..0}$$

$$\text{GPR}[rd] \leftarrow \text{GPR}[rt]_{(63-s)..0} \parallel 0^s$$

Exceptions:

Reserved Instruction



Format: DSRA rd, rt, sa MIPS III

Purpose: To arithmetic right shift a doubleword by a fixed amount — 0 to 31 bits.

Description: $rd \leftarrow rt \gg sa$ (arithmetic)

The 64-bit doubleword contents of GPR *rt* are shifted right, duplicating the sign bit (63) into the emptied bits; the result is placed in GPR *rd*. The bit shift count in the range 0 to 31 is specified by *sa*.

Restrictions:

None


Operation: 64-bit processors

$$s \leftarrow 0 \parallel sa$$

$$GPR[rd] \leftarrow (GPR[rt]_{63})^s \parallel GPR[rt]_{63..s}$$

Exceptions:

Reserved Instruction



Format: DSRA32 rd, rt, sa **MIPS III**

Purpose: To arithmetic right shift a doubleword by a fixed amount — 32-63 bits.

Description: $rd \leftarrow rt \gg (sa+32)$ (arithmetic)

The doubleword contents of GPR *rt* are shifted right, duplicating the sign bit (63) into the emptied bits; the result is placed in GPR *rd*. The bit shift count in the range 32 to 63 is specified by *sa*+32.

Restrictions:

None


Operation: 64-bit processors

$$s \leftarrow 1 \parallel sa \quad /* 32+sa */$$

$$GPR[rd] \leftarrow (GPR[rt]_{63})^s \parallel GPR[rt]_{63..s}$$

Exceptions:

Reserved Instruction



Format: DSRAV rd, rt, rs MIPS III

Purpose: To arithmetic right shift a doubleword by a variable number of bits.

Description: $rd \leftarrow rt \gg rs$ (arithmetic)

The doubleword contents of GPR *rt* are shifted right, duplicating the sign bit (63) into the emptied bits; the result is placed in GPR *rd*. The bit shift count in the range 0 to 63 is specified by the low-order six bits in GPR *rs*.

Restrictions:

None


Operation: 64-bit processors

$$s \leftarrow \text{GPR}[rs]_{5..0}$$

$$\text{GPR}[rd] \leftarrow (\text{GPR}[rt]_{63})^s \parallel \text{GPR}[rt]_{63..s}$$

Exceptions:

Reserved Instruction



Format: DSRL rd, rt, sa **MIPS III**

Purpose: To logical right shift a doubleword by a fixed amount — 0 to 31 bits.

Description: $rd \leftarrow rt \gg sa$ (logical)

The doubleword contents of GPR *rt* are shifted right, inserting zeros into the emptied bits; the result is placed in GPR *rd*. The bit shift count in the range 0 to 31 is specified by *sa*.

Restrictions:

None


Operation: 64-bit processors

$$s \leftarrow 0 \parallel sa$$

$$GPR[rd] \leftarrow 0^s \parallel GPR[rt]_{63..s}$$

Exceptions:

Reserved Instruction



Format: DSRL32 rd, rt, sa MIPS III

Purpose: To logical right shift a doubleword by a fixed amount — 32 to 63 bits.

Description: $rd \leftarrow rt \gg (sa+32)$ (logical)

The 64-bit doubleword contents of GPR *rt* are shifted right, inserting zeros into the emptied bits; the result is placed in GPR *rd*. The bit shift count in the range 32 to 63 is specified by *sa*+32.

Restrictions:

None


Operation: 64-bit processors

$$s \leftarrow 1 \parallel sa \quad /* 32+sa */$$

$$GPR[rd] \leftarrow 0^s \parallel GPR[rt]_{63..s}$$

Exceptions:

Reserved Instruction



Format: DSRLV rd, rt, rs MIPS III

Purpose: To logical right shift a doubleword by a variable number of bits.

Description: $rd \leftarrow rt \gg rs$ (logical)

The 64-bit doubleword contents of GPR *rt* are shifted right, inserting zeros into the emptied bits; the result is placed in GPR *rd*. The bit shift count in the range 0 to 63 is specified by the low-order six bits in GPR *rs*.

Restrictions:

None


Operation: 64-bit processors

$$s \leftarrow GPR[rs]_{5..0}$$

$$GPR[rd] \leftarrow 0^s \parallel GPR[rt]_{63..s}$$

Exceptions:

Reserved Instruction



Format: DSUB rd, rs, rt MIPS III

Purpose: To subtract 64-bit integers; trap if overflow.

Description: $rd \leftarrow rs - rt$

The 64-bit doubleword value in GPR *rt* is subtracted from the 64-bit value in GPR *rs* to produce a 64-bit result. If the subtraction results in 64-bit 2's complement arithmetic overflow then the destination register is not modified and an Integer Overflow exception occurs. If it does not overflow, the 64-bit result is placed into GPR *rd*.

Restrictions:

None

Operation: 64-bit processors


```
temp ← GPR[rs] - GPR[rt]
if (64_bit_arithmetic_overflow) then
    SignalException(IntegerOverflow)
else
    GPR[rd] ← temp
endif
```

Exceptions:

- Integer Overflow
- Reserved Instruction

Programming Notes:

DSUBU performs the same arithmetic operation but, does not trap on overflow.



Format: DSUBU rd, rs, rt MIPS III

Purpose: To subtract 64-bit integers.

Description: $rd \leftarrow rs - rt$

The 64-bit doubleword value in GPR *rt* is subtracted from the 64-bit value in GPR *rs* and the 64-bit arithmetic result is placed into GPR *rd*.

No Integer Overflow exception occurs under any circumstances.

Restrictions:

None

Operation: 64-bit processors


$$GPR[rd] \leftarrow GPR[rs] - GPR[rt]$$

Exceptions:

Reserved Instruction

Programming Notes:

The term "unsigned" in the instruction name is a misnomer; this operation is 64-bit modulo arithmetic that does not trap on overflow. It is appropriate for arithmetic which is not signed, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as "C" language arithmetic.



Format: ERET

Description:

ERET is the RC4650 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.


Operation:

```

T:  if  $SR_2 = 1$  then
      PC ← ErrorEPC
      SR ←  $SR_{31..3} || 0 || SR_{1..0}$ 
    else
      PC ← EPC
      SR ←  $SR_{31..2} || 0 || SR_0$ 
    endif
    LLbit ← 0
  
```

Exceptions:

Coprocessor unusable exception



Format: J target MIPS I

Purpose: To branch within the current 256 MB aligned region.

Description:

This is a PC-region branch (not PC-relative); the effective target address is in the "current" 256 MB aligned region. The low 28 bits of the target address is the *instr_index* field shifted left 2 bits. The remaining upper bits are the corresponding bits of the address of the instruction in the delay slot (**not** the branch itself).

Jump to the effective target address. Execute the instruction following the jump, in the branch delay slot, before jumping.

Restrictions:

None

Operation:

I:

$$I+1: PC \leftarrow PC_{GPREN..28} \parallel instr_index \parallel 0^2$$


Exceptions:

None

Programming Notes:

Forming the branch target address by concatenating PC and index bits rather than adding a signed offset to the PC is an advantage if all program code addresses fit into a 256 MB region aligned on a 256 MB boundary. It allows a branch to anywhere in the region from anywhere in the region which a signed relative offset would not allow.

This definition creates the boundary case where the branch instruction is in the last word of a 256 MB region and can therefore only branch to the following 256 MB region containing the branch delay slot



Format: JAL target MIPS I

Purpose: To procedure call within the current 256 MB aligned region.

Description:

Place the return address link in GPR 31. The return link is the address of the second instruction following the branch, where execution would continue after a procedure call.

This is a PC-region branch (not PC-relative); the effective target address is in the "current" 256 MB aligned region. The low 28 bits of the target address is the *instr_index* field shifted left 2 bits. The remaining upper bits are the corresponding bits of the address of the instruction in the delay slot (**not** the branch itself).

Jump to the effective target address. Execute the instruction following the jump, in the branch delay slot, before jumping.

Restrictions:

None

Operation:

$$I: \text{GPR}[31] \leftarrow \text{PC} + 8$$

$$I+1: \text{PC} \leftarrow \text{PC}_{\text{GPREN}..28} \parallel \text{instr_index} \parallel 0^2$$


Exceptions:

None

Programming Notes:

Forming the branch target address by concatenating PC and index bits rather than adding a signed offset to the PC is an advantage if all program code addresses fit into a 256 MB region aligned on a 256 MB boundary. It allows a branch to anywhere in the region from anywhere in the region which a signed relative offset would not allow.

This definition creates the boundary case where the branch instruction is in the last word of a 256 MB region and can therefore only branch to the following 256 MB region containing the branch delay slot.



Format: JALR *rs* (rd = 31 implied) MIPS I
JALR *rd*, *rs*

Purpose: To procedure call to an instruction address in a register.

Description: $rd \leftarrow return_addr, PC \leftarrow rs$

Place the return address link in GPR *rd*. The return link is the address of the second instruction following the branch, where execution would continue after a procedure call.

Jump to the effective target address in GPR *rs*. Execute the instruction following the jump, in the branch delay slot, before jumping.

Restrictions:

Register specifiers *rs* and *rd* must not be equal, because such an instruction does not have the same effect when re-executed. The result of executing such an instruction is undefined. This restriction permits an exception handler to resume execution by re-executing the branch when an exception occurs in the branch delay slot.

The effective target address in GPR *rs* must be naturally aligned. If either of the two least-significant bits are not -zero, then an Address Error exception occurs, not for the jump instruction, but when the branch target is subsequently fetched as an instruction.

Operation:


I: temp \leftarrow GPR[*rs*]
GPR[*rd*] \leftarrow PC + 8
I+1: PC \leftarrow temp

Exceptions:

None

Programming Notes:

This is the only branch-and-link instruction that can select a register for the return link; all other link instructions use GPR 31. The default register for GPR *rd*, if omitted in the assembly language instruction, is GPR 31.



Format: JR rs MIPS I

Purpose: To branch to an instruction address in a register.

Description: PC ← rs

Jump to the effective target address in GPR *rs*. Execute the instruction following the jump, in the branch delay slot, before jumping.

Restrictions:

The effective target address in GPR *rs* must be naturally aligned. If either of the two least-significant bits are not -zero, then an Address Error exception occurs, not for the jump instruction, but when the branch target is subsequently fetched as an instruction.


Operation:

I: temp ← GPR[rs]

I+1: PC ← temp

Exceptions:

None



Format: LB rt, offset(base) **MIPS I**

Purpose: To load a byte from memory as a signed value.

Description: $rt \leftarrow \text{memory}[\text{base}+\text{offset}]$

The contents of the 8-bit byte at the memory location specified by the effective address are fetched, sign-extended, and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

None

Operation: 32-bit processors


$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $pAddr \leftarrow pAddr_{(PSIZE-1)..2} \parallel (pAddr_{1..0} \text{ xor ReverseEndian}^2)$
 $\text{memword} \leftarrow \text{LoadMemory}(\text{uncached}, \text{BYTE}, pAddr, vAddr, \text{DATA})$
 $\text{byte} \leftarrow vAddr_{1..0} \text{ xor BigEndianCPU}^2$
 $\text{GPR}[rt] \leftarrow \text{sign_extend}(\text{memword}_{7+8*\text{byte}..8*\text{byte}})$

Operation: 64-bit processors

$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $pAddr \leftarrow pAddr_{PSIZE-1..3} \parallel (pAddr_{2..0} \text{ xor ReverseEndian}^3)$
 $\text{memdouble} \leftarrow \text{LoadMemory}(\text{uncached}, \text{BYTE}, pAddr, vAddr, \text{DATA})$
 $\text{byte} \leftarrow vAddr_{2..0} \text{ xor BigEndianCPU}^3$
 $\text{GPR}[rt] \leftarrow \text{sign_extend}(\text{memdouble}_{7+8*\text{byte}..8*\text{byte}})$

Exceptions:

TLB Refill, TLB Invalid
 Address Error



Format: LBU rt, offset(base) **MIPS I**

Purpose: To load a byte from memory as an unsigned value.

Description: $rt \leftarrow \text{memory}[\text{base} + \text{offset}]$

The contents of the 8-bit byte at the memory location specified by the effective address are fetched, zero-extended, and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

None

Operation: 32-bit processors


$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $pAddr \leftarrow pAddr_{\text{PSIZE}-1..2} \parallel (pAddr_{1..0} \text{ xor ReverseEndian}^2)$
 $\text{memword} \leftarrow \text{LoadMemory}(\text{uncached}, \text{BYTE}, pAddr, vAddr, \text{DATA})$
 $\text{byte} \leftarrow vAddr_{1..0} \text{ xor BigEndianCPU}^2$
 $\text{GPR}[\text{rt}] \leftarrow \text{zero_extend}(\text{memword}_{7+8* \text{byte}..8* \text{byte}})$

Operation: 64-bit processors

$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $pAddr \leftarrow pAddr_{\text{PSIZE}-1..3} \parallel (pAddr_{2..0} \text{ xor ReverseEndian}^3)$
 $\text{memdouble} \leftarrow \text{LoadMemory}(\text{uncached}, \text{BYTE}, pAddr, vAddr, \text{DATA})$
 $\text{byte} \leftarrow vAddr_{2..0} \text{ xor BigEndianCPU}^3$
 $\text{GPR}[\text{rt}] \leftarrow \text{zero_extend}(\text{memdouble}_{7+8* \text{byte}..8* \text{byte}})$

Exceptions:

TLB Refill, TLB Invalid
 Address Error



Format: LD *rt*, *offset*(*base*) **MIPS III**

Purpose: To load a doubleword from memory.

Description: $rt \leftarrow \text{memory}[\text{base} + \text{offset}]$

The contents of the 64-bit doubleword at the memory location specified by the aligned effective address are fetched and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

The effective address must be naturally aligned. If any of the three least-significant bits of the address are non-zero, an Address Error exception occurs.

MIPS IV: The low-order 3 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.


Operation: 64-bit processors

```

vAddr ← sign_extend(offset) + GPR[base]
if (vAddr2..0) ≠ 03 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
memdouble ← LoadMemory (uncached, DOUBLEWORD, pAddr, vAddr, DATA)
GPR[rt] ← memdouble
    
```

Exceptions:

- TLB Refill, TLB Invalid
- Bus Error
- Address Error
- Reserved Instruction



Format: LDC1 rt, offset(base) MIPS II
 LDC2 rt, offset(base)

Purpose: To load a doubleword from memory to a coprocessor general register.

Description: $rt \leftarrow \text{memory}[\text{base}+\text{offset}]$

The contents of the 64-bit doubleword at the memory location specified by the aligned effective address are fetched and made available to coprocessor unit zz. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The manner in which each coprocessor uses the data is defined by the individual coprocessor specifications. The usual operation would place the data into coprocessor general register *rt*.

Each MIPS architecture level defines up to 4 coprocessor units, numbered 0 to 3. The opcodes corresponding to coprocessors that are not defined by an architecture level may be used for other instructions.

Restrictions:

Access to the coprocessors is controlled by system software. Each coprocessor has a "coprocessor usable" bit in the System Control coprocessor. The usable bit must be set for a user program to execute a coprocessor instruction. If the usable bit is not set, an attempt to execute the instruction will result in a Coprocessor Unusable exception. An unimplemented coprocessor must never be enabled. The result of executing this instruction for an unimplemented coprocessor when the usable bit is set, is undefined.

This instruction is not available for coprocessor 0, the System Control coprocessor, and the opcode may be used for other instructions.

The effective address must be naturally aligned. If any of the three least-significant bits of the effective address are non-zero, an Address Error exception occurs.

MIPS IV: The low-order 3 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation: 32-bit processors


$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 if $(vAddr_{2..0}) \neq 0^3$ then $\text{SignalException}(\text{AddressError})$ endif
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $\text{memdouble} \leftarrow \text{LoadMemory}(\text{uncached}, \text{DOUBLEWORD}, pAddr, vAddr, \text{DATA})$
 COP_LD (z, rt, memdouble)

Operation: 64-bit processors

$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 if $(vAddr_{2..0}) \neq 0^3$ then $\text{SignalException}(\text{AddressError})$ endif
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $\text{memdouble} \leftarrow \text{LoadMemory}(\text{uncached}, \text{DOUBLEWORD}, pAddr, vAddr, \text{DATA})$
 COP_LD (z, rt, memdouble)

Exceptions:

- TLB Refill, TLB Invalid
- Bus Error
- Address Error
- Reserved Instruction
- Coprocessor Unusable



Format: LDL rt, offset(base) MIPS III

Purpose: To load the most-significant part of a doubleword from an unaligned memory address.

Description: $rt \leftarrow rt \text{ MERGE memory}[\text{base}+\text{offset}]$

The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address (*EffAddr*). *EffAddr* is the address of the most-significant of eight consecutive bytes forming a doubleword in memory (*DW*) starting at an arbitrary byte boundary. A part of *DW*, the most-significant one to eight bytes, is in the aligned doubleword containing *EffAddr*. This part of *DW* is loaded appropriately into the most-significant (left) part of GPR *rt* leaving the remainder of GPR *rt* unchanged.

The figure below illustrates this operation for big-endian byte ordering. The eight consecutive bytes in 2..9 form an unaligned doubleword starting at location 2. A part of *DW*, six bytes, is contained in the aligned doubleword containing the most-significant byte at 2. First, LDL loads these six bytes into the left part of the destination register and leaves the remainder of the destination unchanged. Next, the complementary LDR loads the remainder of the unaligned doubleword.

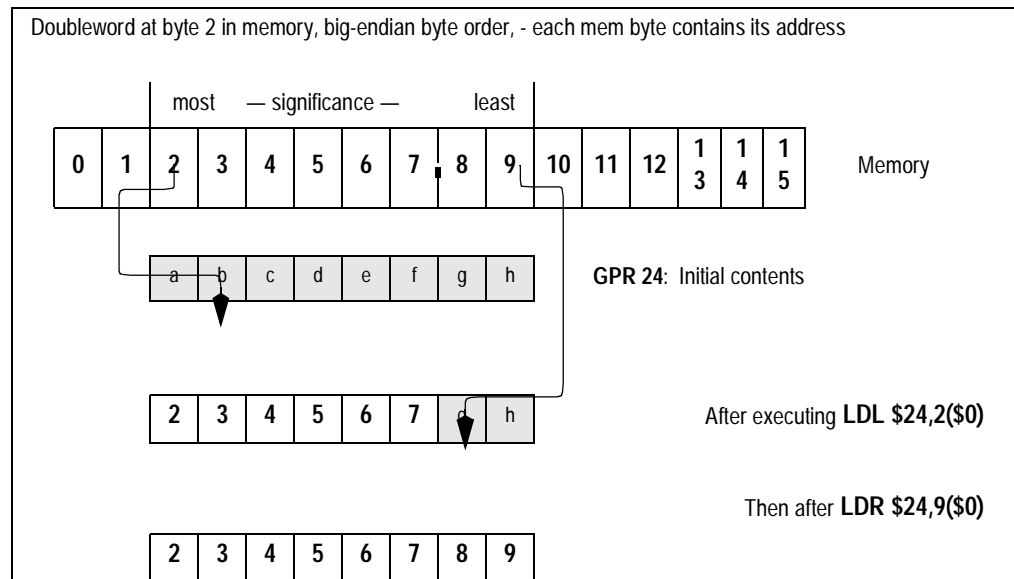


Figure 2.3 Unaligned Doubleword Load using LDL and LDR

The bytes loaded from memory to the destination register depend on both the offset of the effective address within an aligned doubleword, i.e. the low three bits of the address ($vAddr_{2..0}$), and the current byte ordering mode of the processor (big- or little-endian). The table below shows the bytes loaded for every combination of offset and byte ordering.

Memory contents and byte offsets ($vAddr_{2..0}$)								Initial contents of Destination Register								
most – significance – least								most – significance – least								
0	1	2	3	4	5	6	7	big-								
I	J	K	L	M	N	O	P	a	b	c	d	e	f	g	h	
7	6	5	4	3	2	1	0	little-endian offset								
Destination register contents after instruction (shaded is unchanged)																
Big-endian byte ordering								Little-endian byte ordering								
I	J	K	L	M	N	O	P	0	P	b	c	d	e	f	g	h
J	K	L	M	N	O	P	h	1	O	P	c	d	e	f	g	h
K	L	M	N	O	P	g	h	2	N	O	P	d	e	f	g	h
L	M	N	O	P	f	g	h	3	M	N	O	P	e	f	g	h
M	N	O	P	e	f	g	h	4	L	M	N	O	P	f	g	h
N	O	P	d	e	f	g	h	5	K	L	M	N	O	P	g	h
O	P	c	d	e	f	g	h	6	J	K	L	M	N	O	P	h
P	b	c	d	e	f	g	h	7	I	J	K	L	M	N	O	P

Figure 2.4 Bytes Loaded by LDL Instruction

Restrictions:

None


Operation: 64-bit processors

```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddr(PSIZE-1)..3 || (pAddr2..0 xor ReverseEndian3)
if BigEndianMem = 0 then
    pAddr ← pAddr(PSIZE-1)..3 || 03
endif
byte ← vAddr2..0 xor BigEndianCPU3
memdouble ← LoadMemory (uncached, byte, pAddr, vAddr, DATA)
GPR[rt] ← memdouble7+8*byte..0 || GPR[rt]55-8*byte..0
    
```

Exceptions:

- TLB Refill, TLB Invalid
- Bus Error
- Address Error
- Reserved Instruction



Format: LDR *rt*, offset(*base*)

MIPS III

Purpose: To load the least-significant part of a doubleword from an unaligned memory address.

Description: $rt \leftarrow rt \text{ MERGE } \text{memory}[\text{base}+\text{offset}]$

The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address (*EffAddr*). *EffAddr* is the address of the least-significant of eight consecutive bytes forming a doubleword in memory (*DW*) starting at an arbitrary byte boundary. A part of *DW*, the least-significant one to eight bytes, is in the aligned doubleword containing *EffAddr*. This part of *DW* is loaded appropriately into the least-significant (right) part of GPR *rt* leaving the remainder of GPR *rt* unchanged.

The figure below illustrates this operation for big-endian byte ordering. The eight consecutive bytes in 2..9 form an unaligned doubleword starting at location 2. A part of *DW*, two bytes, is contained in the aligned doubleword containing the least-significant byte at 9. First, LDR loads these two bytes into the right part of the destination register and leaves the remainder of the destination unchanged. Next, the complementary LDL loads the remainder of the unaligned doubleword.

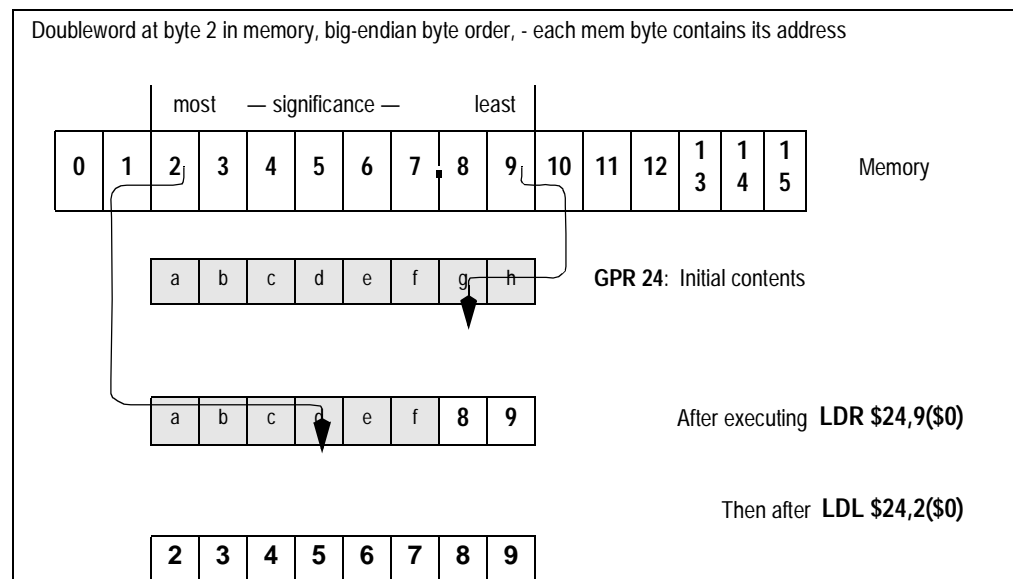


Figure 2.5 Unaligned Doubleword Load using LDR and LDL

The bytes loaded from memory to the destination register depend on both the offset of the effective address within an aligned doubleword, i.e. the low three bits of the address ($vAddr_{2..0}$), and the current byte ordering mode of the processor (big- or little-endian). The table below shows the bytes loaded for every combination of offset and byte ordering.




Figure 2.6 Bytes Loaded by LDR Instruction

Restrictions:


None

Operation: 64-bit processors

$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $pAddr \leftarrow pAddr_{(P\text{SIZE}-1)..3} \parallel (pAddr_{2..0} \text{ xor ReverseEndian}^3)$
 if BigEndianMem = 1 then
 $pAddr \leftarrow pAddr_{(P\text{SIZE}-1)..3} \parallel 0^3$
 endif
 $\text{byte} \leftarrow vAddr_{2..0} \text{ xor BigEndianCPU}^3$
 $\text{memdouble} \leftarrow \text{LoadMemory}(\text{uncached}, \text{byte}, pAddr, vAddr, \text{DATA})$
 $\text{GPR}[\text{rt}] \leftarrow \text{GPR}[\text{rt}]_{63..64-8*\text{byte}} \parallel \text{memdouble}_{63..8*\text{byte}}$

Exceptions:

- TLB Refill, TLB Invalid
- Bus Error
- Address Error
- Reserved Instruction



Format: JALR rs (rd = 31 implied) MIPS I
JALR rd, rs

Purpose: To procedure call to an instruction address in a register.

Description: rd ← return_addr, PC ← rs

Place the return address link in GPR *rd*. The return link is the address of the second instruction following the branch, where execution would continue after a procedure call.

Jump to the effective target address in GPR *rs*. Execute the instruction following the jump, in the branch delay slot, before jumping.

Restrictions:

Register specifiers *rs* and *rd* must not be equal, because such an instruction does not have the same effect when re-executed. The result of executing such an instruction is undefined. This restriction permits an exception handler to resume execution by re-executing the branch when an exception occurs in the branch delay slot.

The effective target address in GPR *rs* must be naturally aligned. If either of the two least-significant bits are not -zero, then an Address Error exception occurs, not for the jump instruction, but when the branch target is subsequently fetched as an instruction.

Operation:


I: temp ← GPR[rs]
GPR[rd] ← PC + 8
I+1: PC ← temp

Exceptions:

None

Programming Notes:

This is the only branch-and-link instruction that can select a register for the return link; all other link instructions use GPR 31. The default register for *GPR rd*, if omitted in the assembly language instruction, is GPR 31.



Format: JR rs MIPS I

Purpose: To branch to an instruction address in a register.

Description: PC ← rs

Jump to the effective target address in GPR *rs*. Execute the instruction following the jump, in the branch delay slot, before jumping.

Restrictions:

The effective target address in GPR *rs* must be naturally aligned. If either of the two least-significant bits are not -zero, then an Address Error exception occurs, not for the jump instruction, but when the branch target is subsequently fetched as an instruction.


Operation:

I: temp ← GPR[rs]

I+1: PC ← temp

Exceptions:

None



Format: LB rt, offset(base) **MIPS I**

Purpose: To load a byte from memory as a signed value.

Description: $rt \leftarrow \text{memory}[\text{base} + \text{offset}]$

The contents of the 8-bit byte at the memory location specified by the effective address are fetched, sign-extended, and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

None

Operation: 32-bit processors

$$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$$

$$(\text{pAddr}, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$$

$$\text{pAddr} \leftarrow \text{pAddr}_{(\text{PSIZE}-1)..2} \parallel (\text{pAddr}_{1..0} \text{ xor ReverseEndian}^2)$$

$$\text{memword} \leftarrow \text{LoadMemory}(\text{uncached}, \text{BYTE}, \text{pAddr}, vAddr, \text{DATA})$$

$$\text{byte} \leftarrow vAddr_{1..0} \text{ xor BigEndianCPU}^2$$

$$\text{GPR}[\text{rt}] \leftarrow \text{sign_extend}(\text{memword}_{7+8*\text{byte}..8*\text{byte}})$$

Operation: 64-bit processors

$$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$$

$$(\text{pAddr}, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$$

$$\text{pAddr} \leftarrow \text{pAddr}_{\text{PSIZE}-1..3} \parallel (\text{pAddr}_{2..0} \text{ xor ReverseEndian}^3)$$


$$\text{memdouble} \leftarrow \text{LoadMemory}(\text{uncached}, \text{BYTE}, \text{pAddr}, vAddr, \text{DATA})$$

$$\text{byte} \leftarrow vAddr_{2..0} \text{ xor BigEndianCPU}^3$$

$$\text{GPR}[\text{rt}] \leftarrow \text{sign_extend}(\text{memdouble}_{7+8*\text{byte}..8*\text{byte}})$$

Exceptions:

TLB Refill, TLB Invalid
Address Error



Format: LBU rt, offset(base) MIPS I

Purpose: To load a byte from memory as an unsigned value.

Description: $rt \leftarrow \text{memory}[\text{base} + \text{offset}]$

The contents of the 8-bit byte at the memory location specified by the effective address are fetched, zero-extended, and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

None

Operation: 32-bit processors


$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $pAddr \leftarrow pAddr_{\text{PSIZE}-1..2} \parallel (pAddr_{1..0} \text{ xor ReverseEndian}^2)$
 $\text{memword} \leftarrow \text{LoadMemory}(\text{uncached}, \text{BYTE}, pAddr, vAddr, \text{DATA})$
 $\text{byte} \leftarrow vAddr_{1..0} \text{ xor BigEndianCPU}^2$
 $\text{GPR}[rt] \leftarrow \text{zero_extend}(\text{memword}_{7+8* \text{byte}..8* \text{byte}})$

Operation: 64-bit processors

$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $pAddr \leftarrow pAddr_{\text{PSIZE}-1..3} \parallel (pAddr_{2..0} \text{ xor ReverseEndian}^3)$
 $\text{memdouble} \leftarrow \text{LoadMemory}(\text{uncached}, \text{BYTE}, pAddr, vAddr, \text{DATA})$
 $\text{byte} \leftarrow vAddr_{2..0} \text{ xor BigEndianCPU}^3$
 $\text{GPR}[rt] \leftarrow \text{zero_extend}(\text{memdouble}_{7+8* \text{byte}..8* \text{byte}})$

Exceptions:

TLB Refill, TLB Invalid
 Address Error



Format: LD *rt*, offset(*base*) MIPS III

Purpose: To load a doubleword from memory.

Description: $rt \leftarrow \text{memory}[\text{base} + \text{offset}]$

The contents of the 64-bit doubleword at the memory location specified by the aligned effective address are fetched and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

The effective address must be naturally aligned. If any of the three least-significant bits of the address are non-zero, an Address Error exception occurs.


MIPS IV: The low-order 3 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation: 64-bit processors

```
vAddr ← sign_extend(offset) + GPR[base]
if (vAddr2..0) ≠ 03 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation(vAddr, DATA, LOAD)
memdouble ← LoadMemory(uncached, DOUBLEWORD, pAddr, vAddr, DATA)
GPR[rt] ← memdouble
```

Exceptions:

- TLB Refill, TLB Invalid
- Bus Error
- Address Error
- Reserved Instruction



Format: LDC1 rt, offset(base) MIPS II
 LDC2 rt, offset(base)

Purpose: To load a doubleword from memory to a coprocessor general register.

Description: $rt \leftarrow \text{memory}[\text{base}+\text{offset}]$

The contents of the 64-bit doubleword at the memory location specified by the aligned effective address are fetched and made available to coprocessor unit zz. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The manner in which each coprocessor uses the data is defined by the individual coprocessor specifications. The usual operation would place the data into coprocessor general register *rt*.

Each MIPS architecture level defines up to 4 coprocessor units, numbered 0 to 3. The opcodes corresponding to coprocessors that are not defined by an architecture level may be used for other instructions.

Restrictions:

Access to the coprocessors is controlled by system software. Each coprocessor has a “coprocessor usable” bit in the System Control coprocessor. The usable bit must be set for a user program to execute a coprocessor instruction. If the usable bit is not set, an attempt to execute the instruction will result in a Coprocessor Unusable exception. An unimplemented coprocessor must never be enabled. The result of executing this instruction for an unimplemented coprocessor when the usable bit is set, is undefined.

This instruction is not available for coprocessor 0, the System Control coprocessor, and the opcode may be used for other instructions.

The effective address must be naturally aligned. If any of the three least-significant bits of the effective address are non-zero, an Address Error exception occurs.

MIPS IV: The low-order 3 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation: 32-bit processors


$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 if $(vAddr_{2..0}) \neq 0^3$ then SignalException(AddressError) endif
 (pAddr, uncached) \leftarrow AddressTranslation (vAddr, DATA, LOAD)
 memdouble \leftarrow LoadMemory (uncached, DOUBLEWORD, pAddr, vAddr, DATA)
 COP_LD (z, rt, memdouble)

Operation: 64-bit processors

$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 if $(vAddr_{2..0}) \neq 0^3$ then SignalException(AddressError) endif
 (pAddr, uncached) \leftarrow AddressTranslation (vAddr, DATA, LOAD)
 memdouble \leftarrow LoadMemory (uncached, DOUBLEWORD, pAddr, vAddr, DATA)
 COP_LD (z, rt, memdouble)

Exceptions:

- TLB Refill, TLB Invalid
- Bus Error
- Address Error
- Reserved Instruction
- Coprocessor Unusable



Format: LDL *rt*, offset(*base*) MIPS III

Purpose: To load the most-significant part of a doubleword from an unaligned memory address.

Description: $rt \leftarrow rt \text{ MERGE memory}[\text{base}+\text{offset}]$

The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address (*EffAddr*). *EffAddr* is the address of the most-significant of eight consecutive bytes forming a doubleword in memory (*DW*) starting at an arbitrary byte boundary. A part of *DW*, the most-significant one to eight bytes, is in the aligned doubleword containing *EffAddr*. This part of *DW* is loaded appropriately into the most-significant (left) part of GPR *rt* leaving the remainder of GPR *rt* unchanged.

The figure below illustrates this operation for big-endian byte ordering. The eight consecutive bytes in 2..9 form an unaligned doubleword starting at location 2. A part of *DW*, six bytes, is contained in the aligned doubleword containing the most-significant byte at 2. First, LDL loads these six bytes into the left part of the destination register and leaves the remainder of the destination unchanged. Next, the complementary LDR loads the remainder of the unaligned doubleword.

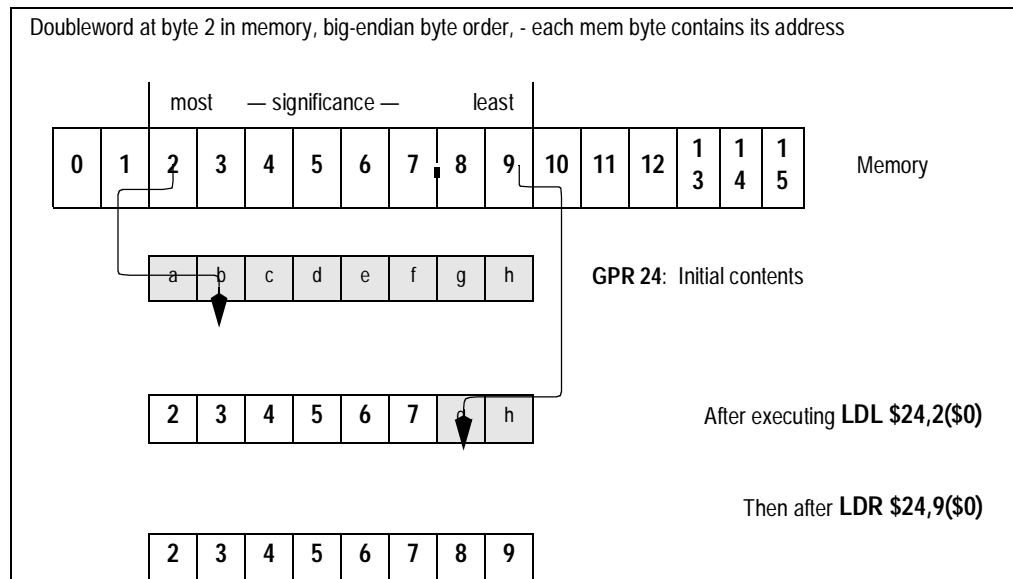


Figure 2.7 Unaligned Doubleword Load using LDL and LDR

The bytes loaded from memory to the destination register depend on both the offset of the effective address within an aligned doubleword, i.e. the low three bits of the address ($vAddr_{2..0}$), and the current byte ordering mode of the processor (big- or little-endian). The table below shows the bytes loaded for every combination of offset and byte ordering.

Memory contents and byte offsets ($vAddr_{2..0}$)								Initial contents of Destination Register								
most —significance— least								most —significance— least								
0	1	2	3	4	5	6	7	big-								
I	J	K	L	M	N	O	P	a	b	c	d	e	f	g	h	
7	6	5	4	3	2	1	0	little-endian offset								
Destination register contents after instruction (shaded is unchanged)																
Big-endian byte ordering								Little-endian byte ordering								
I	J	K	L	M	N	O	P	0	P	b	c	d	e	f	g	h
J	K	L	M	N	O	P	h	1	O	P	c	d	e	f	g	h
K	L	M	N	O	P	g	h	2	N	O	P	d	e	f	g	h
L	M	N	O	P	f	g	h	3	M	N	O	P	e	f	g	h
M	N	O	P	e	f	g	h	4	L	M	N	O	P	f	g	h
N	O	P	d	e	f	g	h	5	K	L	M	N	O	P	g	h
O	P	c	d	e	f	g	h	6	J	K	L	M	N	O	P	h
P	b	c	d	e	f	g	h	7	I	J	K	L	M	N	O	P

Figure 2.8 Bytes Loaded by LDL Instruction

Restrictions:


None

Operation: 64-bit processors

$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $pAddr \leftarrow pAddr_{(P\text{SIZE}-1)..3} \parallel (pAddr_{2..0} \text{ xor ReverseEndian}^3)$
 if BigEndianMem = 0 then
 $pAddr \leftarrow pAddr_{(P\text{SIZE}-1)..3} \parallel 0^3$
 endif
 $\text{byte} \leftarrow vAddr_{2..0} \text{ xor BigEndianCPU}^3$
 $\text{memdouble} \leftarrow \text{LoadMemory}(\text{uncached}, \text{byte}, pAddr, vAddr, \text{DATA})$
 $\text{GPR}[\text{rt}] \leftarrow \text{memdouble}_{7+8*\text{byte}..0} \parallel \text{GPR}[\text{rt}]_{55-8*\text{byte}..0}$

Exceptions:

- TLB Refill, TLB Invalid
- Bus Error
- Address Error
- Reserved Instruction



Format: LDR rt, offset(base) MIPS III

Purpose: To load the least-significant part of a doubleword from an unaligned memory address.

Description: rt ← rt MERGE memory[base+offset]

The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address (*EffAddr*). *EffAddr* is the address of the least-significant of eight consecutive bytes forming a doubleword in memory (*DW*) starting at an arbitrary byte boundary. A part of *DW*, the least-significant one to eight bytes, is in the aligned doubleword containing *EffAddr*. This part of *DW* is loaded appropriately into the least-significant (right) part of GPR *rt* leaving the remainder of GPR *rt* unchanged.

The figure below illustrates this operation for big-endian byte ordering. The eight consecutive bytes in 2..9 form an unaligned doubleword starting at location 2. A part of *DW*, two bytes, is contained in the aligned doubleword containing the least-significant byte at 9. First, LDR loads these two bytes into the right part of the destination register and leaves the remainder of the destination unchanged. Next, the complementary LDL loads the remainder of the unaligned doubleword.

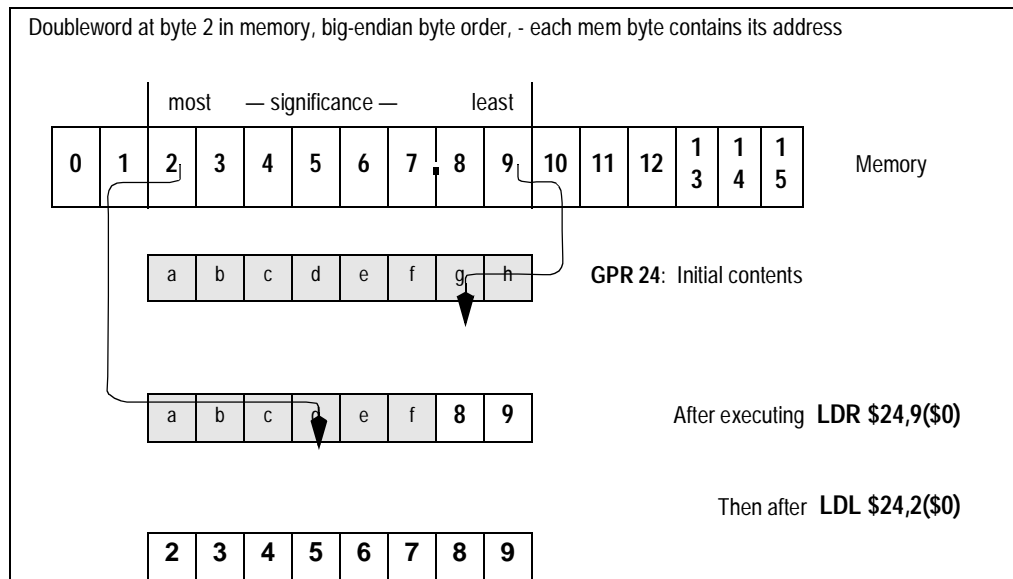


Figure 2.9 Unaligned Doubleword Load using LDR and LDL

The bytes loaded from memory to the destination register depend on both the offset of the effective address within an aligned doubleword, i.e. the low three bits of the address ($vAddr_{2..0}$), and the current byte ordering mode of the processor (big- or little-endian). The table below shows the bytes loaded for every combination of offset and byte ordering.

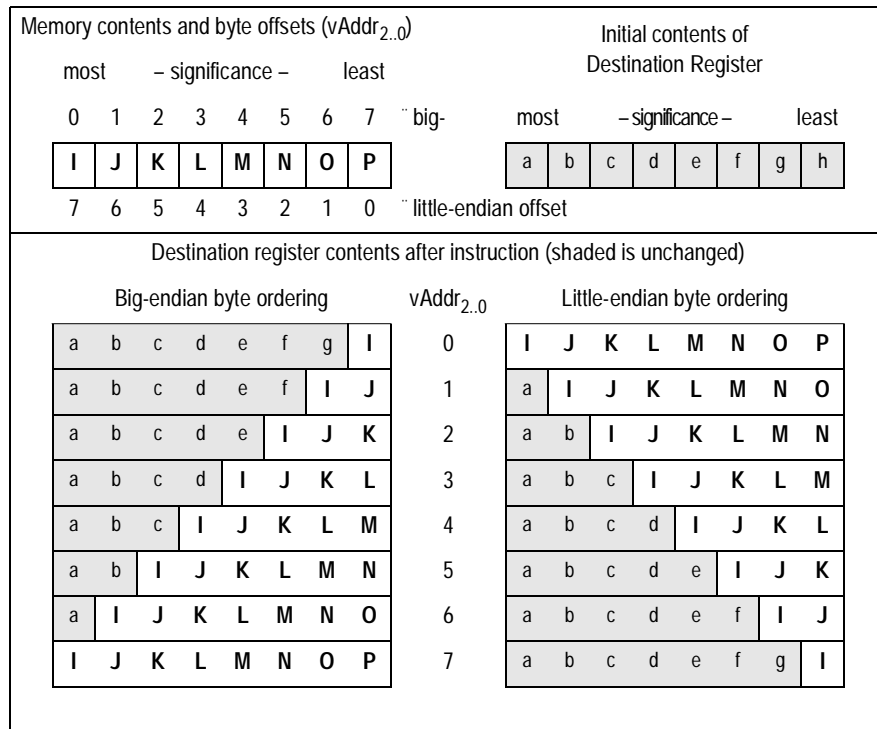


Figure 2.10 Bytes Loaded by LDR Instruction

Restrictions:


None

Operation: 64-bit processors

$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $pAddr \leftarrow pAddr_{(P\text{SIZE}-1)..3} \parallel (pAddr_{2..0} \text{ xor ReverseEndian}^3)$
 if BigEndianMem = 1 then
 $pAddr \leftarrow pAddr_{(P\text{SIZE}-1)..3} \parallel 0^3$
 endif
 $\text{byte} \leftarrow vAddr_{2..0} \text{ xor BigEndianCPU}^3$
 $\text{memdouble} \leftarrow \text{LoadMemory}(\text{uncached}, \text{byte}, pAddr, vAddr, \text{DATA})$
 $\text{GPR}[\text{rt}] \leftarrow \text{GPR}[\text{rt}]_{63..64-8*\text{byte}} \parallel \text{memdouble}_{63..8*\text{byte}}$

Exceptions:

- TLB Refill, TLB Invalid
- Bus Error
- Address Error
- Reserved Instruction



Format: LH *rt*, offset(*base*) MIPS I

Purpose: To load a halfword from memory as a signed value.

Description: $rt \leftarrow \text{memory}[\text{base}+\text{offset}]$

The contents of the 16-bit halfword at the memory location specified by the aligned effective address are fetched, sign-extended, and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

The effective address must be naturally aligned. If the least-significant bit of the address is non-zero, an Address Error exception occurs.

MIPS IV: The low-order bit of the *offset* field must be zero. If it is not, the result of the instruction is undefined.

Operation: 32-bit processors

```

vAddr ← sign_extend(offset) + GPR[base]
if (vAddr0) ≠ 0 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE - 1..2 || (pAddr1..0 xor (ReverseEndian || 0))
memword ← LoadMemory (uncached, HALFWORD, pAddr, vAddr, DATA)
byte ← vAddr1..0 xor (BigEndianCPU || 0)
GPR[rt] ← sign_extend(memword15+8*byte..8* byte)
    
```


Operation: 64-bit processors

```

vAddr ← sign_extend(offset) + GPR[base]
if (vAddr0) ≠ 0 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE - 1..3 || (pAddr2..0 xor (ReverseEndian || 0))
memdouble ← LoadMemory (uncached, HALFWORD, pAddr, vAddr, DATA)
byte ← vAddr2..0 xor (BigEndianCPU2 || 0)
GPR[rt] ← sign_extend(memdouble15+8*byte..8* byte)
    
```

Exceptions:

- TLB Refill, TLB Invalid
- Bus Error
- Address Error



Format: LHU *rt*, *offset*(*base*) **MIPS I**

Purpose: To load a halfword from memory as an unsigned value.

Description: $rt \leftarrow \text{memory}[\text{base} + \text{offset}]$

The contents of the 16-bit halfword at the memory location specified by the aligned effective address are fetched, zero-extended, and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

The effective address must be naturally aligned. If the least-significant bit of the address is non-zero, an Address Error exception occurs.

MIPS IV: The low-order bit of the *offset* field must be zero. If it is not, the result of the instruction is undefined.

Operation: 32-bit processors

```

vAddr ← sign_extend(offset) + GPR[base]
if (vAddr0) ≠ 0 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..2 || (pAddr1..0 xor (ReverseEndian || 0))
memword ← LoadMemory (uncached, HALFWORD, pAddr, vAddr, DATA)
byte ← vAddr1..0 xor (BigEndianCPU || 0)
GPR[rt] ← zero_extend(memword15+8*byte..8*byte)
    
```


Operation: 64-bit processors

```

vAddr ← sign_extend(offset) + GPR[base]
if (vAddr0) ≠ 0 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian2 || 0))
memdouble ← LoadMemory (uncached, HALFWORD, pAddr, vAddr, DATA)
byte ← vAddr2..0 xor (BigEndianCPU2 || 0)
GPR[rt] ← zero_extend(memdouble15+8*byte..8*byte)
    
```

Exceptions:

- TLB Refill, TLB Invalid
- Address Error



Format: LL rt, offset(base) **MIPS II**

Purpose: To load a word from memory for an atomic read-modify-write.

Description: $rt \leftarrow \text{memory}[\text{base}+\text{offset}]$

The LL and SC instructions provide primitives to implement atomic Read-Modify-Write (RMW) operations for cached memory locations.

The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address.

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched, sign-extended to the GPR register length if necessary, and written into GPR *rt*. This begins a RMW sequence on the current processor.

There is one active RMW sequence per processor. When an LL is executed it starts the active RMW sequence replacing any other sequence that was active.

The RMW sequence is completed by a subsequent SC instruction that either completes the RMW sequence atomically and succeeds, or does not and fails. See the description of SC for a list of events and conditions that cause the SC to fail and an example instruction sequence using LL and SC.

Executing LL on one processor does not cause an action that, by itself, would cause an SC for the same block to fail on another processor.

An execution of LL does not have to be followed by execution of SC; a program is free to abandon the RMW sequence without attempting a write.

Restrictions:

The addressed location must be cached; if it is not, the result is undefined.

The effective address must be naturally aligned. If either of the two least-significant bits of the effective address are non-zero an Address Error exception occurs.

MIPS IV: The low-order 2 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation: 32-bit processors

```

vAddr ← sign_extend(offset) + GPR[base]
if (vAddr1..0) ≠ 02 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
memword ← LoadMemory (uncached, WORD, pAddr, vAddr, DATA)
GPR[rt] ← memword
LLbit ← 1
    
```

Operation: 64-bit processors

```

vAddr ← sign_extend(offset) + GPR[base]
if (vAddr1..0) ≠ 02 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
memdouble ← LoadMemory (uncached, WORD, pAddr, vAddr, DATA)
byte ← vAddr2..0 xor (BigEndianCPU || 02)
GPR[rt] ← sign_extend(memdouble31+8*byte..8*byte)
LLbit ← 1
    
```

Exceptions:


TLB Refill, TLB Invalid
Address Error
Reserved Instruction

Programming Notes:

There is no Load Linked Word Unsigned operation corresponding to Load Word Unsigned.

Implementation Notes:

An LL on one processor must not take action that, by itself, would cause an SC for the same block on another processor to fail. If an implementation depends on retaining the data in cache during the RMW sequence, cache misses caused by LL must not fetch data in the exclusive state, thus removing it from the cache, if it is present in another cache.



Format: LLD rt, offset(base) MIPS III

Purpose: To load a doubleword from memory for an atomic read-modify-write.

Description: $rt \leftarrow \text{memory}[\text{base}+\text{offset}]$

The LLD and SCD instructions provide primitives to implement atomic Read-Modify-Write (RMW) operations for cached memory locations.

The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address.

The contents of the 64-bit doubleword at the memory location specified by the aligned effective address are fetched and written into GPR *rt*. This begins a RMW sequence on the current processor.

There is one active RMW sequence per processor. When an LLD is executed it starts the active RMW sequence replacing any other sequence that was active.

The RMW sequence is completed by a subsequent SCD instruction that either completes the RMW sequence atomically and succeeds, or does not and fails. See the description of SCD for a list of events and conditions that cause the SCD to fail and an example instruction sequence using LLD and SCD.

Executing LLD on one processor does not cause an action that, by itself, would cause an SCD for the same block to fail on another processor.

An execution of LLD does not have to be followed by execution of SCD; a program is free to abandon the RMW sequence without attempting a write.

Restrictions:

The addressed location must be cached; if it is not, the result is undefined.

The effective address must be naturally aligned. If either of the three least-significant bits of the effective address are non-zero an Address Error exception occurs.

MIPS IV: The low-order 3 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation: 64-bit processors

```

vAddr ← sign_extend(offset) + GPR[base]
if (vAddr2..0) ≠ 03 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
memdouble ← LoadMemory (uncached, DOUBLEWORD, pAddr, vAddr, DATA)
GPR[rt] ← memdouble
LLbit ← 1
    
```


Exceptions:

- TLB Refill, TLB Invalid
- Address Error
- Reserved Instruction

Programming Notes:

Implementation Notes:

An LLD on one processor must not take action that, by itself, would cause an SCD for the same block on another processor to fail. If an implementation depends on retaining the data in cache during the RMW sequence, cache misses caused by LLD must not fetch data in the exclusive state, thus removing it from the cache, if it is present in another cache.



Format: LUI rt, immediate MIPS I

Purpose: To load a constant into the upper half of a word.

Description: $rt \leftarrow \text{immediate} \parallel 0^{16}$

The 16-bit *immediate* is shifted left 16 bits and concatenated with 16 bits of low-order zeros. The 32-bit result is sign-extended and placed into GPR *rt*.

Restrictions:


None

Operation:

$GPR[rt] \leftarrow \text{sign_extend}(\text{immediate} \parallel 0^{16})$

Exceptions:

None



Format: LW *rt*, offset(*base*) **MIPS I**

Purpose: To load a word from memory as a signed value.

Description: $rt \leftarrow \text{memory}[\text{base}+\text{offset}]$

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched, sign-extended to the GPR register length if necessary, and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

The effective address must be naturally aligned. If either of the two least-significant bits of the address are non-zero, an Address Error exception occurs.

MIPS IV: The low-order 2 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation: 32-bit processors


$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 if $(vAddr_{1..0}) \neq 0^2$ then SignalException(AddressError) endif
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $\text{memword} \leftarrow \text{LoadMemory}(\text{uncached}, \text{WORD}, pAddr, vAddr, \text{DATA})$
 $\text{GPR}[rt] \leftarrow \text{memword}$

Operation: 64-bit processors

$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 if $(vAddr_{1..0}) \neq 0^2$ then SignalException(AddressError) endif
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $pAddr \leftarrow pAddr_{\text{PSIZE}-1..3} \parallel (pAddr_{2..0} \text{ xor } (\text{ReverseEndian} \parallel 0^2))$
 $\text{memdouble} \leftarrow \text{LoadMemory}(\text{uncached}, \text{WORD}, pAddr, vAddr, \text{DATA})$
 $\text{byte} \leftarrow vAddr_{2..0} \text{ xor } (\text{BigEndianCPU} \parallel 0^2)$
 $\text{GPR}[rt] \leftarrow \text{sign_extend}(\text{memdouble}_{31+8*\text{byte}..8*\text{byte}})$

Exceptions:

- TLB Refill, TLB Invalid
- Bus Error
- Address Error



Format: LWC1 rt, offset(base) MIPS I
 LWC2 rt, offset(base)
 LWC3 rt, offset(base)

Purpose: To load a word from memory to a coprocessor general register.

Description: $rt \leftarrow \text{memory}[\text{base}+\text{offset}]$

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched and made available to coprocessor unit *zz*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The manner in which each coprocessor uses the data is defined by the individual coprocessor specification. The usual operation would place the data into coprocessor general register *rt*.

Each MIPS architecture level defines up to 4 coprocessor units, numbered 0 to 3 (see Section titled "Coprocessor Instructions" earlier in this Chapter). The opcodes corresponding to coprocessors that are not defined by an architecture level may be used for other instructions.

Restrictions:

Access to the coprocessors is controlled by system software. Each coprocessor has a "coprocessor usable" bit in the System Control coprocessor. The usable bit must be set for a user program to execute a coprocessor instruction. If the usable bit is not set, an attempt to execute the instruction will result in a Coprocessor Unusable exception. An unimplemented coprocessor must never be enabled. The result of executing this instruction for an unimplemented coprocessor when the usable bit is set, is undefined.

This instruction is not available for coprocessor 0, the System Control coprocessor, and the opcode may be used for other instructions.

The effective address must be naturally aligned. If either of the two least-significant bits of the address are non-zero, an Address Error exception occurs.

MIPS IV: The low-order 2 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation: 32-bit processors


I: $vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 if $(vAddr_{1..0}) \neq 0^2$ then $\text{SignalException}(\text{AddressError})$ endif
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $\text{memword} \leftarrow \text{LoadMemory}(\text{uncached}, \text{WORD}, pAddr, vAddr, \text{DATA})$
 I+1: COP_LW (*z*, *rt*, *memword*)

Operation: 64-bit processors

$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 if $(vAddr_{1..0}) \neq 0^2$ then $\text{SignalException}(\text{AddressError})$ endif
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $pAddr \leftarrow pAddr_{\text{PSIZE}-1..3} \parallel (pAddr_{2..0} \text{ xor } (\text{ReverseEndian} \parallel 0^2))$
 $\text{memdouble} \leftarrow \text{LoadMemory}(\text{uncached}, \text{DOUBLEWORD}, pAddr, vAddr, \text{DATA})$
 $\text{byte} \leftarrow vAddr_{2..0} \text{ xor } (\text{BigEndianCPU} \parallel 0^2)$
 $\text{memword} \leftarrow \text{memdouble}_{31+8*\text{byte}..8*\text{byte}}$
 COP_LW (*z*, *rt*, *memdouble*)

Exceptions:

- TLB Refill, TLB Invalid
- Bus Error
- Address Error
- Coprocessor Unusable



Format: LWL rt, offset(base) MIPS I

Purpose: To load the most-significant part of a word as a signed value from an unaligned memory address.

Description: $rt \leftarrow rt \text{ MERGE } \text{memory}[\text{base}+\text{offset}]$

The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address (*EffAddr*). *EffAddr* is the address of the most-significant of four consecutive bytes forming a word in memory (*W*) starting at an arbitrary byte boundary. A part of *W*, the most-significant one to four bytes, is in the aligned word containing *EffAddr*. This part of *W* is loaded into the most-significant (left) part of the word in GPR *rt*. The remaining least-significant part of the word in GPR *rt* is unchanged.

If GPR *rt* is a 64-bit register, the destination word is the low-order word of the register. The loaded value is treated as a signed value; the word sign bit (bit 31) is always loaded from memory and the new sign bit value is copied into bits 63..32.

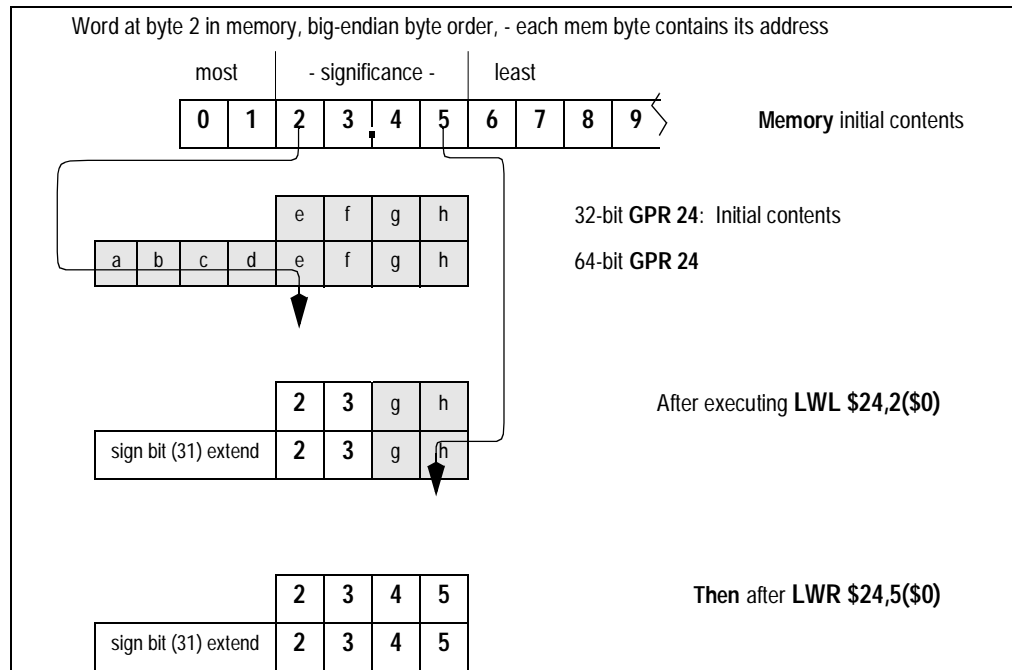


Figure 2.11 Unaligned Word Load using LWL and LWR

The figure above illustrates this operation for big-endian byte ordering for 32-bit and 64-bit registers. The four consecutive bytes in 2..5 form an unaligned word starting at location 2. A part of *W*, two bytes, is in the aligned word containing the most-significant byte at 2. First, LWL loads these two bytes into the left part of the destination register word and leaves the right part of the destination word unchanged. Next, the complementary LWR loads the remainder of the unaligned word.

The bytes loaded from memory to the destination register depend on both the offset of the effective address within an aligned word, i.e. the low two bits of the address ($vAddr_{1..0}$), and the current byte ordering mode of the processor (big- or little-endian). The table below shows the bytes loaded for every combination of offset and byte ordering.

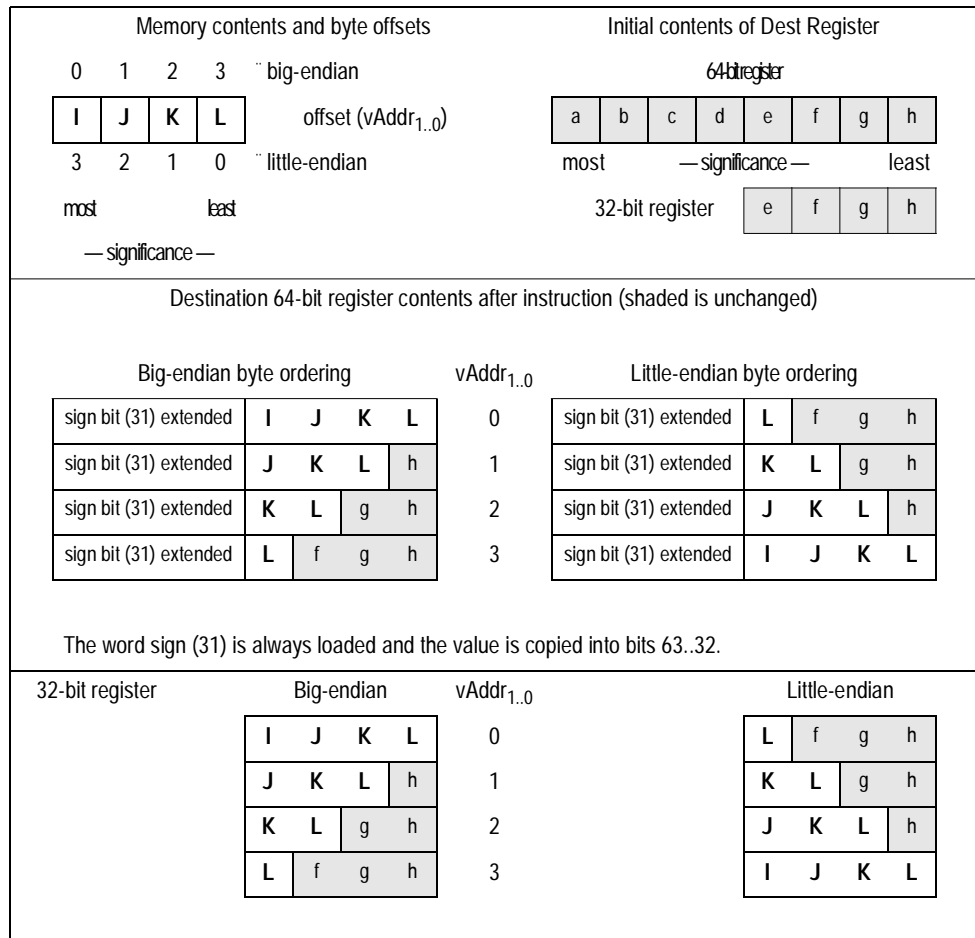


Figure 2.12 Bytes Loaded by LWL Instruction

The unaligned loads, LWL and LWR, are exceptions to the load-delay scheduling restriction in the MIPS I architecture. An unaligned load instruction to GPR rt that immediately follows another load to GPR rt can “read” the loaded data. It will correctly merge the 1 to 4 loaded bytes with the data loaded by the previous instruction.

Restrictions:

MIPS I scheduling restriction: The loaded data is not available for use by the following instruction. The instruction immediately following this one, unless it is an unaligned load (LWL, LWR), may not use GPR rt as a source register. If this restriction is violated, the result of the operation is undefined.

Operation: 32-bit processors

```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddr(PSIZE-1)..2 || (pAddr1..0 xor ReverseEndian2)
if BigEndianMem = 0 then
    pAddr ← pAddr(PSIZE-1)..2 || 02
endif
byte ← vAddr1..0 xor BigEndianCPU2
memword ← LoadMemory (uncached, byte, pAddr, vAddr, DATA)
GPR[rt] ← memword7+8*byte..0 || GPR[rt]23-8*byte..0
    
```

Operation: 64-bit processors

```


vAddr ← sign_extend(offset) + GPR[base]
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddr(PSIZE-1)..3 || (pAddr2..0 xor ReverseEndian3)
if BigEndianMem = 0 then
    pAddr ← pAddr(PSIZE-1)..3 || 03
endif
byte ← 0 || (vAddr1..0 xor BigEndianCPU2)
word ← vAddr2 xor BigEndianCPU
memdouble ← LoadMemory (uncached, byte, pAddr, vAddr, DATA)
temp ← memdouble31+32*word-8*byte..32*word || GPR[rt]23-8*byte..0
GPR[rt] ← (temp31)32 || temp
    
```

Exceptions:

- TLB Refill, TLB Invalid
- Bus Error
- Address Error

Programming Notes:

The architecture provides no direct support for treating unaligned words as unsigned values, i.e. zeroing bits 63..32 of the destination register when bit 31 is loaded. See SLL or SLLV for a single-instruction method of propagating the word sign bit in a register into the upper half of a 64-bit register.



Format: LWR rt, offset(base) MIPS I

Purpose: To load the least-significant part of a word from an unaligned memory address as a signed value.

Description: $rt \leftarrow rt \text{ MERGE memory}[\text{base} + \text{offset}]$

The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address (*EffAddr*). *EffAddr* is the address of the least-significant of four consecutive bytes forming a word in memory (*W*) starting at an arbitrary byte boundary. A part of *W*, the least-significant one to four bytes, is in the aligned word containing *EffAddr*. This part of *W* is loaded into the least-significant (right) part of the word in GPR *rt*. The remaining most-significant part of the word in GPR *rt* is unchanged.

If GPR *rt* is a 64-bit register, the destination word is the low-order word of the register. The loaded value is treated as a signed value; if the word sign bit (bit 31) is loaded (i.e. when all four bytes are loaded) then the new sign bit value is copied into bits 63..32. If bit 31 is not loaded then the value of bits 63..32 is implementation dependent; the value is either unchanged or a copy of the current value of bit 31. Executing both LWR and LWL, in either order, delivers in a sign-extended word value in the destination register.

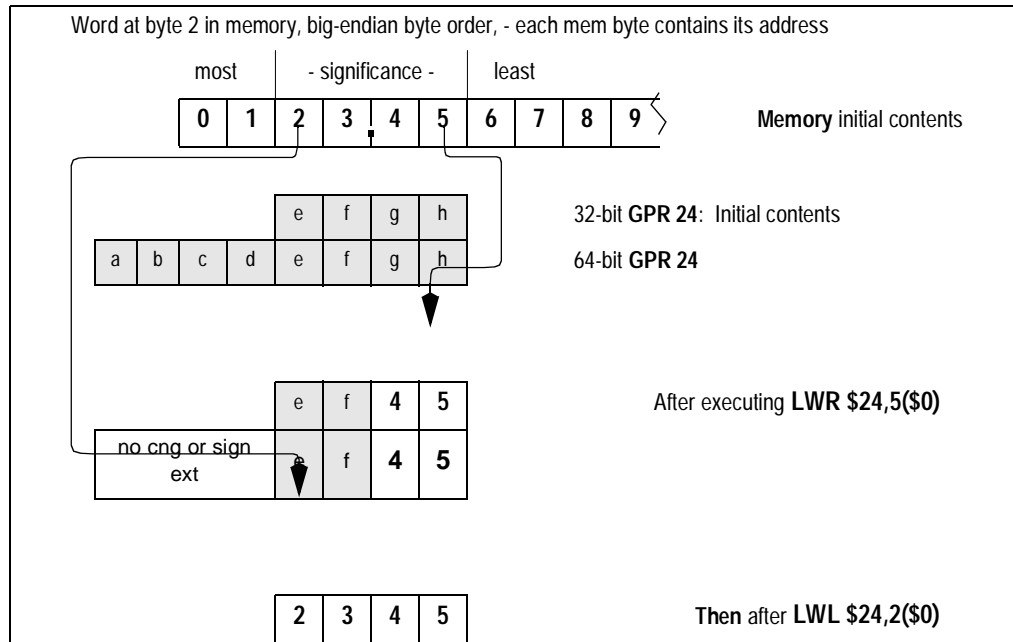


Figure 2.13 Unaligned Word Load using LWR and LWL

The figure above illustrates this operation for big-endian byte ordering for 32-bit and 64-bit registers. The four consecutive bytes in 2..5 form an unaligned word starting at location 2. A part of *W*, two bytes, is in the aligned word containing the least-significant byte at 5. First, LWR loads these two bytes into the right part of the destination register. Next, the complementary LWL loads the remainder of the unaligned word.

The bytes loaded from memory to the destination register depend on both the offset of the effective address within an aligned word, i.e. the low two bits of the address ($vAddr_{1..0}$), and the current byte ordering mode of the processor (big- or little-endian). The table below shows the bytes loaded for every combination of offset and byte ordering.

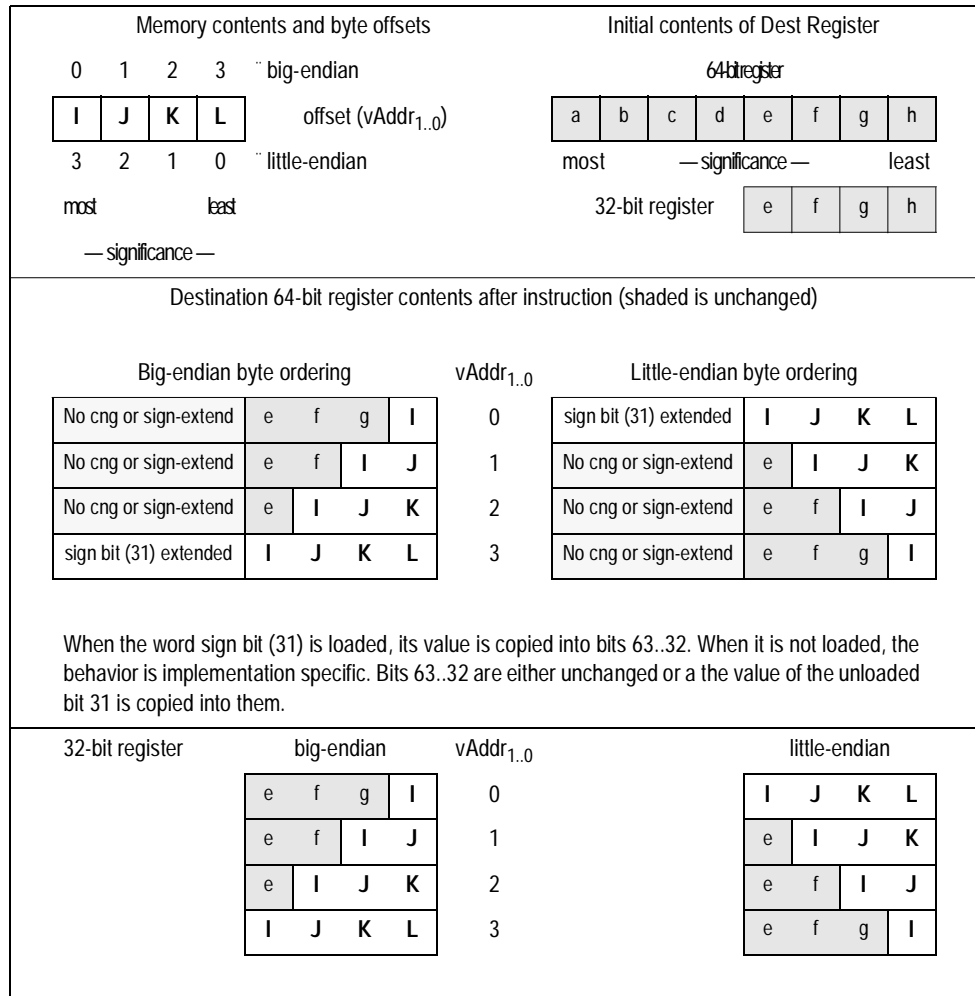


Figure 2.14 Bytes Loaded by LWR Instruction

The unaligned loads, LWL and LWR, are exceptions to the load-delay scheduling restriction in the MIPS I architecture. An unaligned load to GPR *rt* that immediately follows another load to GPR *rt* can “read” the loaded data. It will correctly merge the 1 to 4 loaded bytes with the data loaded by the previous instruction.

Restrictions:

MIPS I scheduling restriction: The loaded data is not available for use by the following instruction. The instruction immediately following this one, unless it is an unaligned load (LWL, LWR), may not use GPR *rt* as a source register. If this restriction is violated, the result of the operation is undefined.

Restrictions:

None

Operation: 32-bit processors

```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddr(PSIZE-1)..2 || (pAddr1..0 xor ReverseEndian2)
if BigEndianMem = 0 then
    pAddr ← pAddr(PSIZE-1)..2 || 02
endif
byte ← vAddr1..0 xor BigEndianCPU2
memword ← LoadMemory (uncached, byte, pAddr, vAddr, DATA)
GPR[rt] ← memword31..32-8*byte || GPR[rt]31-8*byte..0
    
```

Operation: 64-bit processors

```


vAddr ← sign_extend(offset) + GPR[base]
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddr(PSIZE-1)..3 || (pAddr2..0 xor ReverseEndian3)
if BigEndianMem = 1 then
    pAddr ← pAddr(PSIZE-1)..3 || 03
endif
byte ← vAddr1..0 xor BigEndianCPU2
word ← vAddr2 xor BigEndianCPU
memdouble ← LoadMemory (uncached, 0 || byte, pAddr, vAddr, DATA)
temp ← GPR[rt]31..32-8*byte || memdouble31+32*word..32*word+8*byte
if byte = 4 then
    utemp ← (temp31)32 /* loaded bit 31, must sign extend */
else
    one of the following two behaviors:
        utemp ← GPR[rt]63..32 /* leave what was there alone */
        utemp ← (GPR[rt]31)32 /* sign-extend bit 31 */
endif
GPR[rt] ← utemp || temp
    
```

Exceptions:

- TLB Refill, TLB Invalid
- Bus Error
- Address Error

Programming Notes:

The architecture provides no direct support for treating unaligned words as unsigned values, i.e. zeroing bits 63..32 of the destination register when bit 31 is loaded. See SLL or SLLV for a single-instruction method of propagating the word sign bit in a register into the upper half of a 64-bit register.



Format: LWU rt, offset(base) **MIPS III**

Purpose: To load a word from memory as an unsigned value.

Description: $rt \leftarrow \text{memory}[\text{base}+\text{offset}]$

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched, zero-extended, and placed in GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

The effective address must be naturally aligned. If either of the two least-significant bits of the address are non-zero, an Address Error exception occurs.


MIPS IV: The low-order 2 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation: 64-bit processors

$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 if $(vAddr_{1..0}) \neq 0^2$ then SignalException(AddressError) endif
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 $pAddr \leftarrow pAddr_{PSIZE-1..3} \parallel (pAddr_{2..0} \text{ xor } (\text{ReverseEndian} \parallel 0^2))$
 $\text{memdouble} \leftarrow \text{LoadMemory}(\text{uncached}, \text{WORD}, pAddr, vAddr, \text{DATA})$
 $\text{byte} \leftarrow vAddr_{2..0} \text{ xor } (\text{BigEndianCPU} \parallel 0^2)$
 $\text{GPR}[rt] \leftarrow 0^{32} \parallel \text{memdouble}_{31+8*\text{byte}..8*\text{byte}}$

Exceptions:

- TLB Refill, TLB Invalid
- Bus Error
- Address Error
- Reserved Instruction



Format: MAD rs, rt

Description:

The RC4650 and RC32364 add a MAD instruction (multiply-accumulate, with HI and LO as the accumulator) to the base MIPS-III ISA. The MAD instruction is defined as:

$$HI,LO \leftarrow HI,LO + rs*rt$$

The lower 32-bits of the accumulator are stored in the lower 32 bits of LO, while the upper 32 bits of the result are stored in the lower 32 bits of HI. This is done to allow this instruction to operate compatibly in 32-bit processors.

The actual repeat rate and latency of this operation are dependent on the size of the operands.

Operation:

$$T: \quad temp \leftarrow (HI_{31..0} \parallel LO_{31..0}) + ((rs_{31..0})^{32} \parallel rs_{31..0}) \times ((rt_{31..0})^{32} \parallel rt_{31..0})$$

$$Hi \leftarrow (temp_{63})^{32} \parallel temp_{63..32}$$


$$Lo \leftarrow (temp_{31})^{32} \parallel temp_{31..0}$$

Exceptions:

None

Programming Notes:

This is an IDT proprietary extension.



Format: MADU rs, rt

Description:

The RC4650 and RC32364 add a MAD instruction (multiply-accumulate, with HI and LO as the accumulator) to the base MIPS-III ISA. The MAD instruction is defined as:

$$HI,LO \leftarrow HI,LO + rs*rt$$

The lower 32-bits of the accumulator are stored in the lower 32 bits of LO, while the upper 32 bits of the result are stored in the lower 32 bits of HI. This is done to allow this instruction to operate compatibly in 32-bit processors. The actual repeat rate and latency of this operation are dependent on the size of the operands.

Operation:

$$T: \text{temp} \leftarrow (HI_{31..0} \parallel LO_{31..0}) + (0^{32} \parallel rs_{31..0}) \times (0^{32} \parallel rt_{31..0})$$

$$Hi \leftarrow (\text{temp}_{63})^{32} \parallel \text{temp}_{63..32}$$


$$Lo \leftarrow (\text{temp}_{31})^{32} \parallel \text{temp}_{31..0}$$

Exceptions:

None

Programming Notes:

This is an IDT proprietary extension.



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.

Note: *See "Opcode Bit Encoding" on page 120, or "CPU InstructionEncoding" at the end of Appendix A.

Operation:

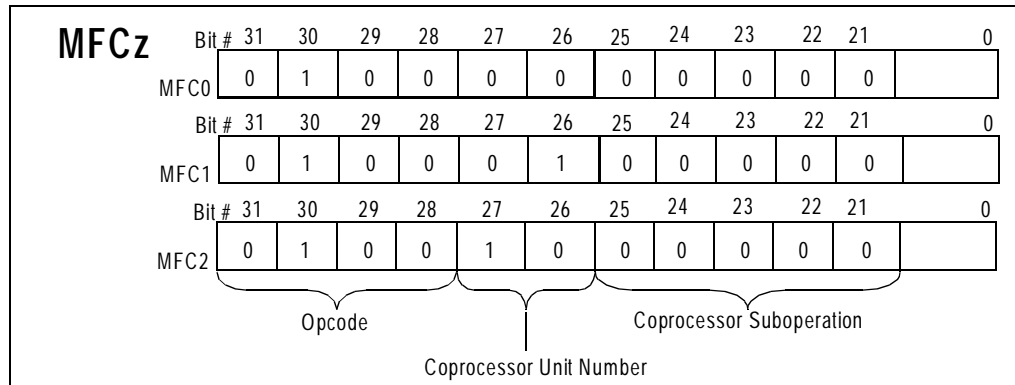
```


T:   if rd0 = 0 then
        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)

Opcode Bit Encoding:





Format: MFHI rd MIPS I

Purpose: To copy the special purpose HI register to a GPR.

Description: rd ← HI

The contents of special register *HI* are loaded into GPR *rd*.

Restrictions:


The two instructions that follow an MFHI instruction must not be instructions that modify the *HI* register: DDIV, DDIVU, DIV, DIVU, DMULT, DMULTU, MTHI, MULT, MULTU. If this restriction is violated, the result of the MFHI is undefined.

Operation:

GPR[rd] ← HI

Exceptions:

None



Format: MFLO rd MIPS I

Purpose: To copy the special purpose LO register to a GPR.

Description: rd ← LO

The contents of special register *LO* are loaded into GPR *rd*.

Restrictions:


The two instructions that follow an MFLO instruction must not be instructions that modify the *LO* register: DDIV, DDIVU, DIV, DIVU, DMULT, DMULTU, MTLO, MULT, MULTU. If this restriction is violated, the result of the MFLO is undefined.

Operation:

GPR[rd] ← LO

Exceptions:

None



Format: MOVN rd, rs, rt MIPS IV, RC32364

Purpose: To conditionally move a GPR after testing a GPR value.

Description: if (rt ≠ 0) then rd ← rs

If the value in GPR *rt* is not equal to zero, then the contents of GPR *rs* are placed into GPR *rd*.

Restrictions:

None

Operation:

```


if GPR[rt] ≠ 0 then
    GPR[rd] ← GPR[rs]
endif
    
```

Exceptions:

Reserved Instruction

Programming Notes:

The nonzero value tested here is the "condition true" result from the SLT, SLTI, SLTU, and SLTIU comparison instructions.



Format: MOVZ rd, rs, rt MIPS IV, RC32364

Purpose: To conditionally move a GPR after testing a GPR value.

Description: if (rt = 0) then rd ← rs

If the value in GPR *rt* is equal to zero, then the contents of GPR *rs* are placed into GPR *rd*.

Restrictions:

None

Operation:

```


if GPR[rt] = 0 then
    GPR[rd] ← GPR[rs]
endif
    
```

Exceptions:

Reserved Instruction

Programming Notes:

The zero value tested here is the "condition false" result from the SLT, SLTI, SLTU, and SLTIU comparison instructions.



Format: MSUB rs, rt RC32364

Description:

The RC32364 adds this new instruction. The content of general registers rs and rt are multiplied, treating both operands as 32-bit two's complement values, and the result is subtracted from HI/LO. No overflow exception occur under any circumstances.

When the operation is complete, the low-order word of the double result is loaded in LO, and the high-order word of the double result is loaded into HI. The instruction is not interlocked so any attempt to read HI/LO before the operation completes returns undefined value.

Operation:


T: temp <-- (HI || LO) - GPR[rs] * GPR[rt]
 LO <-- temp
 HI <-- temp

Exceptions:

None

Programming Notes:

This is an IDT proprietary extension.



Format: MSUB rs, rt RC32364

Description:

The RC32364 adds this new instruction. The content of general registers rs and rt are multiplied, treating both operand as 32-bit unsigned values, and the result is subtracted from HI/LO. No overflow exception occur under any circumstances.

When the operation completes, the low-order word of the double result is loaded in LO, and the high-order word of the double result is loaded in HI. The instruction is not interlocked so any attempt to read HI/LO before the operation completes returns undefined value.

Operation:


T: temp <-- (HI || LO) - (0||GPR[rs]) * (0||GPR[rt])
 LO <-- temp
 HI <-- temp

Exceptions:

None

Programming Notes:

This is an IDT proprietary extension.



Format: MTCz rt, rd

Note: *See "Opcode Bit Encoding" on this page, or "CPU Instruction Encoding" at the end of Appendix A.

Description:

The contents of general register *rt* are loaded into coprocessor register *rd* of coprocessor *z*. Execution of the instruction referencing coprocessor 3 causes a reserved instruction exception, not a coprocessor unusable exception.

Operation:

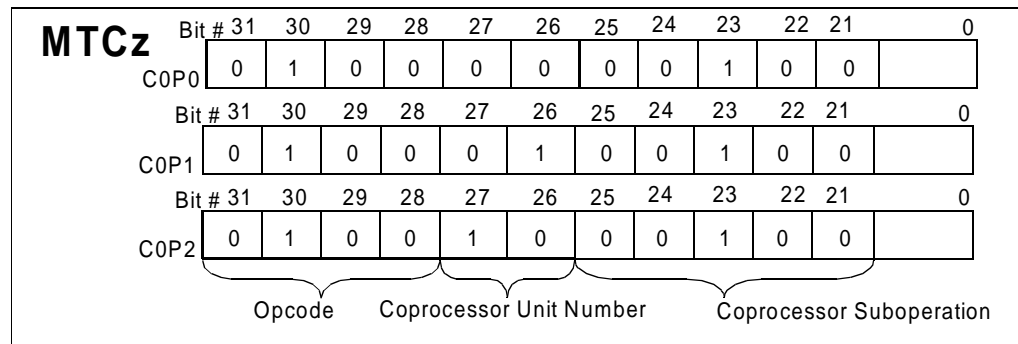
```


T:  data ← GPR[rt]31..0
T+1: if rd0 = 0
      CPR[z,rd4..1 || 0] ← CPR[z, rd4..1 || 0]63..32 || data
      else
      CPR[z,rd4..1 || 0] ← data || CPR[z,rd4..1 || 0]31..0
      endif
    
```

Exceptions:

- Coprocessor unusable exception
- Reserved instruction exception (coprocessor 3)

Opcode Bit Encoding:





Format: MTHI rs MIPS I

Purpose: To copy a GPR to the special purpose HI register.

Description: HI ← rs

The contents of GPR *rs* are loaded into special register *HI*.

Restrictions:

If either of the two preceding instructions is MFHI, the result of that MFHI is undefined. Reads of the *HI* or *LO* special registers must be separated from subsequent instructions that write to them by two or more other instructions.

A computed result written to the *HI/LO* pair by DDIV, DDIVU, DIV, DIVU, DMULT, DMULTU, MULT, or MULTU must be read by MFHI or MFLO before another result is written into either *HI* or *LO*. If an MTHI instruction is executed following one of these arithmetic instructions, but before a MFLO or MFHI instruction, the contents of *LO* are undefined. The following example shows this illegal situation:


```
MUL    r2,r4    # start operation that will eventually write to HI,LO
...          # code not containing mfhi or mflo
MTHI   r6
...          # code not containing mflo
MFLO   r3      # this mflo would get an undefined value
```

Operation:

```
I-2, I-1: HI ← undefined
I:         HI ← GPR[rs]
```

Exceptions:

None



Format: MTLO rs MIPS I

Purpose: To copy a GPR to the special purpose LO register.

Description: LO ← rs

The contents of GPR *rs* are loaded into special register *LO*.

Restrictions:

If either of the two preceding instructions is MFLO, the result of that MFLO is undefined. Reads of the *HI* or *LO* special registers must be separated from subsequent instructions that write to them by two or more other instructions.

A computed result written to the *HI/LO* pair by DDIV, DDIVU, DIV, DIVU, DMULT, DMULTU, MULT, or MULTU must be read by MFHI or MFLO before another result is written into either *HI* or *LO*. If an MTLO instruction is executed following one of these arithmetic instructions, but before a MFLO or MFHI instruction, the contents of *HI* are undefined. The following example shows this illegal situation:


```
MUL    r2,r4    # start operation that will eventually write to HI,LO
...    # code not containing mfhi or mflo
MTLO   r6
...    # code not containing mfhi
MFHI   r3      # this mfhi would get an undefined value
```

Operation:

```
I-2, I-1: LO ← undefined
I:        LO ← GPR[rs]
```

Exceptions:

None



Format: MUL rd, rs, rt

Description:

The RC4650 and RC32364 add a true 3-operand 32x32→32 multiply instruction to the MIPS-III ISA, where by rd = rs*rt. This instruction eliminates the need to explicitly move the multiply result from the LO register back to a general register. The execution time of this operation is operand size dependent.

The HI and LO registers are undefined after executing this instruction. For 16-bit operands, the latency of MUL is 3 cycles, with a repeat rate of 2 cycles. In addition, the MUL instruction will unconditionally slip or stall for all but 2 cycles of its latency.

Operation:


```
T: Temp " rs 31 . . . 0 x rt 31 . . 0
rd " (temp31)32 || temp 31. . . 0
HI " undefined
LO " undefined
```

Exceptions:

None

Programming Notes:

This instruction is an IDT proprietary extension.



Format: MULT rs, rt MIPS I

Purpose: To multiply 32-bit signed integers.

Description: (LO, HI) ← rs × rt

The 32-bit word value in GPR *rt* is multiplied by the 32-bit value in GPR *rs*, treating both operands as signed values, to produce a 64-bit result. The low-order 32-bit word of the result is placed into special register *LO*, and the high-order 32-bit word is placed into special register *HI*.

No arithmetic exception occurs under any circumstances.

Restrictions:

On 64-bit processors, if either GPR *rt* or GPR *rs* do not contain sign-extended 32-bit values (bits 63..31 equal), then the result of the operation is undefined.

If either of the two preceding instructions is MFHI or MFLO, the result of the MFHI or MFLO is undefined. Reads of the *HI* or *LO* special registers must be separated from subsequent instructions that write to them by two or more other instructions.

Operation:

```

if (NotWordValue(GPR[rs]) or NotWordValue(GPR[rt])) then UndefinedResult() endif
I-2, I-1:    LO, HI ← undefined
I:          prod ← GPR[rs]31..0 * GPR[rt]31..0
            LO ← sign_extend(prod31..0)
            HI ← sign_extend(prod63..32)
    
```


Exceptions:

None

Programming Notes:

In some processors the integer multiply operation may proceed asynchronously and allow other CPU instructions to execute before it is complete. An attempt to read *LO* or *HI* before the results are written will wait (interlock) until the results are ready. Asynchronous execution does not affect the program result, but offers an opportunity for performance improvement by scheduling the multiply so that other instructions can execute in parallel.

Programs that require overflow detection must check for it explicitly.



Format: MULTU rs, rt MIPS I

Purpose: To multiply 32-bit unsigned integers.

Description: (LO, HI) ← rs × rt

The 32-bit word value in GPR *rt* is multiplied by the 32-bit value in GPR *rs*, treating both operands as unsigned values, to produce a 64-bit result. The low-order 32-bit word of the result is placed into special register *LO*, and the high-order 32-bit word is placed into special register *HI*.

No arithmetic exception occurs under any circumstances.

Restrictions:

On 64-bit processors, if either GPR *rt* or GPR *rs* do not contain sign-extended 32-bit values (bits 63..31 equal), then the result of the operation is undefined.

If either of the two preceding instructions is MFHI or MFLO, the result of the MFHI or MFLO is undefined. Reads of the *HI* or *LO* special registers must be separated from subsequent instructions that write to them by two or more other instructions.

Operation:

```

if (NotWordValue(GPR[rs]) or NotWordValue(GPR[rt])) then UndefinedResult() endif
I-2, I-1:    LO, HI    ← undefined
I:          prod    ← (0 || GPR[rs]31..0) * (0 || GPR[rt]31..0)
              LO     ← sign_extend(prod31..0)
              HI     ← sign_extend(prod63..32)
    
```


Exceptions:

None

Programming Notes:

In some processors the integer multiply operation may proceed asynchronously and allow other CPU instructions to execute before it is complete. An attempt to read *LO* or *HI* before the results are written will wait (interlock) until the results are ready. Asynchronous execution does not affect the program result, but offers an opportunity for performance improvement by scheduling the multiply so that other instructions can execute in parallel.

Programs that require overflow detection must check for it explicitly.



Format: NOR rd, rs, rt MIPS I

Purpose: To do a bitwise logical NOT OR.

Description: $rd \leftarrow rs \text{ NOR } rt$

The contents of GPR *rs* are combined with the contents of GPR *rt* in a bitwise logical NOR operation. The result is placed into GPR *rd*.

Restrictions:


None

Operation:

$GPR[rd] \leftarrow GPR[rs] \text{ nor } GPR[rt]$

Exceptions:

None



Format: OR rd, rs, rt MIPS I

Purpose: To do a bitwise logical OR.

Description: $rd \leftarrow rs \text{ OR } rt$

The contents of GPR *rs* are combined with the contents of GPR *rt* in a bitwise logical OR operation. The result is placed into GPR *rd*.

Restrictions:


None

Operation:

$GPR[rd] \leftarrow GPR[rs] \text{ or } GPR[rt]$

Exceptions:

None



Format: ORI rt, rs, immediate MIPS I

Purpose: To do a bitwise logical OR with a constant.

Description: $rd \leftarrow rs \text{ OR } \text{immediate}$

The 16-bit *immediate* is zero-extended to the left and combined with the contents of GPR *rs* in a bitwise logical OR operation. The result is placed into GPR *rt*.

Restrictions:


None

Operation:

$GPR[rt] \leftarrow \text{zero_extend}(\text{immediate}) \text{ or } GPR[rs]$

Exceptions:

None



Format: PREF hint, offset(base) **MIPS IV, RC32364**

Purpose: To prefetch data from memory.

Description: prefetch_memory(base+offset)

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. The *hint* field supplies information about the way that the data is expected to be used.

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. An implementation may do nothing when executing a PREF instruction.

If MIPS IV instructions are supported and enabled, PREF does not cause addressing-related exceptions. If it raises an exception condition, the exception condition is ignored. If an addressing-related exception condition 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 (see the section titled “**Memory Access Types**” earlier in this Chapter).

If PREF results in a memory operation, the memory access type used for the operation is determined by the memory access type of the effective address, just as it would be if the memory operation had been caused by a load or store to the effective address.

PREF enables the processor to take some action, typically prefetching the data into cache, to improve program performance. The action taken for a specific PREF instruction is both system and context dependent. Any action, including doing nothing, is permitted that does not change architecturally-visible state or alter the meaning of a program. It is expected that implementations will either do nothing or take an action that will increase the performance of the program.

For a cached location, the expected, and useful, action is for the processor to prefetch a block of data that includes the effective address. The size of the block, and the level of the memory hierarchy it is fetched into are implementation specific.

The *hint* field supplies information about the way the data is expected to be used. No *hint* value causes an action that modifies architecturally-visible state. A processor may use a *hint* value to improve the effectiveness of the prefetch action. The defined *hint* values and the recommended prefetch action are shown in the table below. The *hint* table may be extended in future implementations.

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	store	Data is expected to be stored or modified. Fetch data as if for a store.
31	Ignore hit (Kernel Mode only)	Invalidate the cache line and bring in the new data from memory regardless of the state of the valid bit.

Table 2.29 Values of Hint Field for Prefetch Instruction in RC32364

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	store	Data is expected to be stored or modified. Fetch data as if for a store.
2-3		Not yet defined.
4	load_streamed	Data is expected to be loaded (not modified) but not reused extensively; it will "stream" through cache. Fetch data as if for a load and place it in the cache so that it will not displace data prefetched as "retained".
5	store_streamed	Data is expected to be stored or modified but not reused extensively; it will "stream" through cache. Fetch data as if for a store and place it in the cache so that it will not displace data prefetched as "retained".
6	load_retained	Data is expected to be loaded (not modified) and reused extensively; it should be "retained" in the cache. Fetch data as if for a load and place it in the cache so that it will not be displaced by data prefetched as "streamed".
7	store_retained	Data is expected to be stored or modified and reused extensively; it should be "retained" in the cache. Fetch data as if for a store and place it in the cache so that will not be displaced by data prefetched as "streamed".
8-31		Not yet defined.

Table 2.30 Values of Hint Field for Prefetch Instruction in RC5000

Restrictions:

None

Operation:

$vAddr \leftarrow GPR[base] + sign_extend(offset)$
 $(pAddr, uncached) \leftarrow AddressTranslation(vAddr, DATA, LOAD)$
 Prefetch(uncached, pAddr, vAddr, DATA, hint)

Exceptions:

Reserved Instruction


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 is determined.

Implementation Notes:

It is recommended that a reserved *hint* field value either cause a default prefetch action that is expected to be useful for most cases of data use, such as the "load" *hint*, or cause the instruction to be treated as a NOP.



Format: RFE

Description:

This instruction is not implemented on RC4000 and RISCORE32300 processors; use ERET instead.

RFE restores the previous interrupt mask and Kernel/User-mode bits (IEp and KUp) of the Status register (SR) into the corresponding current status bits (IEc and KUp) and restores the old status bits (IEo and KUo) into the corresponding previous status bits (IEp and KUp). The old status bits remain unchanged.

The architecture does not specify the operation of memory references associated with load/store instructions immediately prior to an RFE instruction. Normally, the RFE instruction follows in the delay slot of a JR (jump register) instruction to restore the PC.


R2000/RC3000/RC6000

Operation:

$$T: SR \leftarrow SR31..4 || SR5..2 \\ LLbit \leftarrow 0$$

Exceptions:

- Coprocessor unusable exception
- Reserved instruction exception (RC4000)



Format: SB rt, offset(base) **MIPS I**

Note: It is recommended that a reserved *hint* field value either cause a default prefetch action that is expected to be useful for most cases of data use, such as the “load” *hint*, or cause the instruction to be treated as a NOP.

Purpose: To store a byte to memory.

Description: memory[base+offset] ← rt

The least-significant 8-bit byte of GPR *rt* is stored in memory at the location specified by the effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

None

Operation: 32-bit processors


vAddr ← sign_extend(offset) + GPR[base]
 (pAddr, uncached) ← AddressTranslation (vAddr, DATA, STORE)
 pAddr ← pAddr_{PSIZE-1..2} || (pAddr_{1..0} xor ReverseEndian²)
 byte ← vAddr_{1..0} xor BigEndianCPU²
 dataword ← GPR[rt]_{31-8*byte..0} || 0^{8*byte}
 StoreMemory (uncached, BYTE, dataword, pAddr, vAddr, DATA)

Operation: 64-bit processors

vAddr ← sign_extend(offset) + GPR[base]
 (pAddr, uncached) ← AddressTranslation (vAddr, DATA, STORE)
 pAddr ← pAddr_{PSIZE-1..3} || (pAddr_{2..0} xor ReverseEndian³)
 byte ← vAddr_{2..0} xor BigEndianCPU³
 datadouble ← GPR[rt]_{63-8*byte..0} || 0^{8*byte}
 StoreMemory (uncached, BYTE, datadouble, pAddr, vAddr, DATA)

Exceptions:

TLB Refill, TLB Invalid
 TLB Modified
 Bus Error
 Address Error



Format: SC rt, offset(base) **MIPS II**

Purpose: To store a word to memory to complete an atomic read-modify-write.

Description: if (atomic_update) then memory[base+offset] ← rt, rt ← 1 else rt ← 0

The LL and SC instructions provide primitives to implement atomic Read-Modify-Write (RMW) operations for cached memory locations.

The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address.

The SC completes the RMW sequence begun by the preceding LL instruction executed on the processor. If it would complete the RMW sequence atomically, then the least-significant 32-bit word of GPR *rt* is stored into memory at the location specified by the aligned effective address and a one, indicating success, is written into GPR *rt*. Otherwise, memory is not modified and a zero, indicating failure, is written into GPR *rt*.

If any of the following events occurs between the execution of LL and SC, the SC will fail:

- ◆ *A coherent store is completed by another processor or coherent I/O module into the block of physical memory containing the word. The size and alignment of the block is implementation dependent. It is at least one word and is at most the minimum page size.*
- ◆ *An exception occurs on the processor executing the LL/SC. An implementation may detect "an exception" in one of three ways:*
 - 1) *Detect exceptions and fail when an exception occurs.*
 - 2) *Fail after the return-from-interrupt instruction (RFE or ERET) is executed.*
 - 3) *Do both 1 and 2.*

If any of the following events occurs between the execution of LL and SC, the SC may succeed or it may fail; the success or failure is unpredictable. Portable programs should not cause one of these events:

- ◆ *A load, store, or prefetch is executed on the processor executing the LL/SC.*
- ◆ *The instructions executed starting with the LL and ending with the SC do not lie in a 2048-byte contiguous region of virtual memory. The region does not have to be aligned, other than the alignment required for instruction words.*

The following conditions must be true or the result of the SC will be undefined:

- ◆ *Execution of SC must have been preceded by execution of an LL instruction.*
- ◆ *A RMW sequence executed without intervening exceptions must use the same address in the LL and SC. The address is the same if the virtual address, physical address, and cache-coherence algorithm are identical.*

Atomic RMW is provided only for cached memory locations. The extent to which the detection of atomicity operates correctly depends on the system implementation and the memory access type used for the location. See the section titled "Memory Access Types" earlier in this Chapter.

MP atomicity: To provide atomic RMW among multiple processors, all accesses to the location must be made with a memory access type of cached coherent.

Uniprocessor atomicity: To provide atomic RMW on a single processor, all accesses to the location must be made with memory access type of either cached noncoherent or cached coherent. All accesses must be to one or the other access type, they may not be mixed.

I/O System: To provide atomic RMW with a coherent I/O system, all accesses to the location must be made with a memory access type of cached coherent. If the I/O system does not use coherent memory operations, then atomic RMW cannot be provided with respect to the I/O reads and writes.

The definition above applies to user-mode operation on all MIPS processors that support the MIPS II architecture. There may be other implementation-specific events, such as privileged CPO instructions, that will cause an SC instruction to fail in some cases. System programmers using LL/SC should consult implementation-specific documentation.

Restrictions:

The addressed location must have a memory access type of cached noncoherent or cached coherent; if it does not, the result is undefined (see the section titled “Memory Access Types” earlier in this Chapter).

The effective address must be naturally aligned. If either of the two least-significant bits of the address are non-zero, an Address Error exception occurs.

MIPS IV: The low-order 2 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation: 32-bit processors

```
vAddr ← sign_extend(offset) + GPR[base]
if (vAddr1..0) ≠ 02 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, STORE)
dataword ← GPR[rt]
if LLbit then
    StoreMemory (uncached, WORD, dataword, pAddr, vAddr, DATA)
endif
GPR[rt] ← 031 || LLbit
```

Operation: 64-bit processors

```
vAddr ← sign_extend(offset) + GPR[base]
if (vAddr1..0) ≠ 02 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
byte ← vAddr2..0 xor (BigEndianCPU || 02)
datadouble ← GPR[rt]63-8*byte..0 || 08*byte
if LLbit then
    StoreMemory (uncached, WORD, datadouble, pAddr, vAddr, DATA)
endif
GPR[rt] ← 063 || LLbit
```

Exceptions:

- TLB Refill, TLB Invalid
- TLB Modified
- Address Error
- Reserved Instruction

Programming Notes:

LL and SC are used to atomically update memory locations as shown in the example atomic increment operation below.


```
L1:
    LL      T1, (T0)    # load counter
    ADDI   T2, T1, 1   # increment
    SC     T2, (T0)    # try to store, checking for atomicity
    BEQ    T2, 0, L1   # if not atomic (0), try again
    NOP
```

Exceptions between the LL and SC cause SC to fail, so persistent exceptions must be avoided. Some examples of these are arithmetic operations that trap, system calls, floating-point operations that trap or require software emulation assistance.

LL and SC function on a single processor for cached noncoherent memory so that parallel programs can be run on uniprocessor systems that do not support cached coherent memory access types.

Implementation Notes:

The block of memory that is "locked" for LL/SC is typically the largest cache line in use.



Format: SCD *rt*, offset(*base*) **MIPS III**

Purpose: To store a doubleword to memory to complete an atomic read-modify-write.

Description: if (*atomic_update*) then memory[*base+offset*] ← *rt*, *rt* ← 1 else *rt* ← 0

The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address.

The SCD completes the RMW sequence begun by the preceding LLD instruction executed on the processor. If it would complete the RMW sequence atomically, then the 64-bit doubleword of GPR *rt* is stored into memory at the location specified by the aligned effective address and a one, indicating success, is written into GPR *rt*. Otherwise, memory is not modified and a zero, indicating failure, is written into GPR *rt*.

If any of the following events occurs between the execution of LLD and SCD, the SCD will fail:

- ◆ A coherent store is completed by another processor or coherent I/O module into the block of physical memory containing the word. The size and alignment of the block is implementation dependent. It is at least one doubleword and is at most the minimum page size.
- ◆ An exception occurs on the processor executing the LLD/SCD. An implementation may detect "an exception" in one of three ways:
 - 1) Detect exceptions and fail when an exception occurs.
 - 2) Fail after the return-from-interrupt instruction (RFE or ERET) is executed.
 - 3) Do both 1 and 2.

If any of the following events occurs between the execution of LLD and SCD, the SCD may succeed or it may fail; the success or failure is unpredictable. Portable programs should not cause one of these events.

- ◆ A memory access instruction (load, store, or prefetch) is executed on the processor executing the LLD/SCD.
- ◆ The instructions executed starting with the LLD and ending with the SCD do not lie in a 2048-byte contiguous region of virtual memory. The region does not have to be aligned, other than the alignment required for instruction words.

The following conditions must be true or the result of the SCD will be undefined:

- ◆ Execution of SCD must have been preceded by execution of an LLD instruction.
- ◆ A RMW sequence executed without intervening exceptions must use the same address in the LLD and SCD. The address is the same if the virtual address, physical address, and cache-coherence algorithm are identical.

Atomic RMW is provided only for memory locations with cached noncoherent or cached coherent memory access types. The extent to which the detection of atomicity operates correctly depends on the system implementation and the memory access type used for the location. See the section titled "Memory Access Types" earlier in this Chapter.

MP atomicity: To provide atomic RMW among multiple processors, all accesses to the location must be made with a memory access type of cached coherent.

Uniprocessor atomicity: To provide atomic RMW on a single processor, all accesses to the location must be made with memory access type of either cached noncoherent or cached coherent. All accesses must be to one or the other access type, they may not be mixed.

I/O System: To provide atomic RMW with a coherent I/O system, all accesses to the location must be made with a memory access type of cached coherent. If the I/O system does not use coherent memory operations, then atomic RMW cannot be provided with respect to the I/O reads and writes.

The defemination above applies to user-mode operation on all MIPS processors that support the MIPS III architecture. There may be other implementation-specific events, such as privileged CPO instructions, that will cause an SCD instruction to fail in some cases. System programmers using LLD/SCD should consult implementation-specific documentation.

Restrictions:

The addressed location must have a memory access type of cached noncoherent or cached coherent; if it does not, the result is undefined (see the section titled “**Memory Access Types**” earlier in this Chapter. The 64-bit doubleword of register *rt* is conditionally stored in memory at the location specified by the aligned effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The effective address must be naturally aligned. If any of the three least-significant bits of the address are non-zero, an Address Error exception occurs.

MIPS IV: The low-order 3 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation: 64-bit processors

```
vAddr ← sign_extend(offset) + GPR[base]
if (vAddr2..0) ≠ 03 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, STORE)
datadouble ← GPR[rt]
if LLbit then
    StoreMemory (uncached, DOUBLEWORD, datadouble, pAddr, vAddr, DATA)
endif
GPR[rt] ← 063 || LLbit
```

Exceptions:

- TLB Refill, TLB Invalid
- TLB Modified
- Address Error
- Reserved Instruction

Programming Notes:

LLD and SCD are used to atomically update memory locations as shown in the example atomic increment operation below.


```
L1:
    LLD    T1, (T0)    # load counter
    ADDI   T2, T1, 1   # increment
    SCD   T2, (T0)    # try to store, checking for atomicity
    BEQ   T2, 0, L1   # if not atomic (0), try again
    NOP                                # branch-delay slot
```

Exceptions between the LLD and SCD cause SCD to fail, so persistent exceptions must be avoided. Some examples of these are arithmetic operations that trap, system calls, floating-point operations that trap or require software emulation assistance.

LLD and SCD function on a single processor for cached noncoherent memory so that parallel programs can be run on uniprocessor systems that do not support cached coherent memory access types.

Implementation Notes:

The block of memory that is “locked” for LLD/SCD is typically the largest cache line in use.



Format: SD rt, offset(base) MIPS III

Purpose: To store a doubleword to memory.

Description: memory[base+offset] ← rt

The 64-bit doubleword in GPR *rt* is stored in memory at the location specified by the aligned effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

The effective address must be naturally aligned. If any of the three least-significant bits of the effective address are non-zero, an Address Error exception occurs.

MIPS IV: The low-order 3 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.


Operation: 64-bit processors

```

vAddr ← sign_extend(offset) + GPR[base]
if (vAddr2..0) ≠ 03 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, STORE)
datadouble ← GPR[rt]
StoreMemory (uncached, DOUBLEWORD, datadouble, pAddr, vAddr, DATA)
    
```

Exceptions:

- TLB Refill, TLB Invalid
- TLB Modified
- Address Error
- Reserved Instruction



Format: SDC1 rt, offset(base) MIPS II
 SDC2 rt, offset(base)

Purpose: To store a doubleword from a coprocessor general register to memory.

Description: memory[base+offset] ← rt

Coprocessor unit zz supplies a 64-bit doubleword which is stored at the memory location specified by the aligned effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The data supplied by each coprocessor is defined by the individual coprocessor specifications. The usual operation would read the data from coprocessor general register *rt*.

Each MIPS architecture level defines up to 4 coprocessor units, numbered 0 to 3. The opcodes corresponding to coprocessors that are not defined by an architecture level may be used for other instructions.

Restrictions:

Access to the coprocessors is controlled by system software. Each coprocessor has a “coprocessor usable” bit in the System Control coprocessor. The usable bit must be set for a user program to execute a coprocessor instruction. If the usable bit is not set, an attempt to execute the instruction will result in a Coprocessor Unusable exception. An unimplemented coprocessor must never be enabled. The result of executing this instruction for an unimplemented coprocessor when the usable bit is set, is undefined.

This instruction is not defined for coprocessor 0, the System Control coprocessor, and the opcode may be used for other instructions.


The effective address must be naturally aligned. If any of the three least-significant bits of the effective address are non-zero, an Address Error exception occurs.

MIPS IV: The low-order 3 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation: 32-bit processors
 $vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 if $(vAddr_{2..0}) \neq 0^3$ then SignalException(AddressError) endif
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{STORE})$
 $\text{datadouble} \leftarrow \text{COP_SD}(z, rt)$
 StoreMemory(uncached, DOUBLEWORD, datadouble, pAddr, vAddr, DATA)

Operation: 64-bit processors
 $vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 if $(vAddr_{2..0}) \neq 0^3$ then SignalException(AddressError) endif
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{STORE})$
 $\text{datadouble} \leftarrow \text{COP_SD}(z, rt)$
 StoreMemory(uncached, DOUBLEWORD, datadouble, pAddr, vAddr, DATA)

Exceptions:
 TLB Refill, TLB Invalid
 TLB Modified
 Address Error
 Reserved Instruction
 Coprocessor Unusable



Format: SDL rt, offset(base) **MIPS III**

Purpose: To store the most-significant part of a doubleword to an unaligned memory address.

Description: memory[base+offset] ← Some_Bytes_From rt

The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address (*EffAddr*). *EffAddr* is the address of the most-significant of eight consecutive bytes forming a doubleword in memory (*DW*) starting at an arbitrary byte boundary. A part of *DW*, the most-significant one to eight bytes, is in the aligned doubleword containing *EffAddr*. The same number of most-significant (left) bytes of GPR *rt* are stored into these bytes of *DW*.

The figure below illustrates this operation for big-endian byte ordering. The eight consecutive bytes in 2..9 form an unaligned doubleword starting at location 2. A part of *DW*, six bytes, is contained in the aligned doubleword containing the most-significant byte at 2. First, SDL stores the six most-significant bytes of the source register into these bytes in memory. Next, the complementary SDR instruction stores the remainder of *DW*.

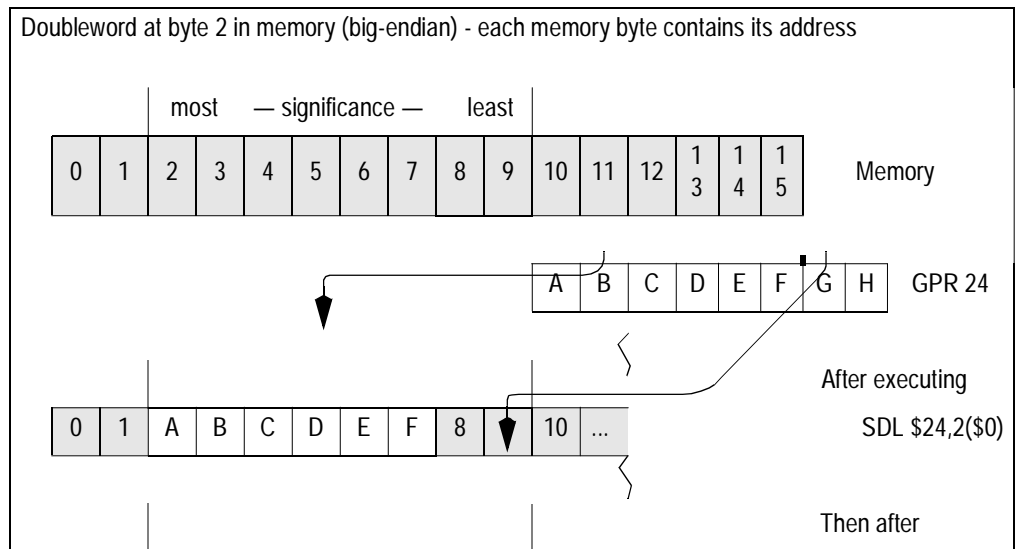


Figure 2.15 Unaligned Doubleword Store with SDL and SDR

The bytes stored from the source register to memory depend on both the offset of the effective address within an aligned doubleword, i.e. the low three bits of the address ($vAddr_{2..0}$), and the current byte ordering mode of the processor (big- or little-endian). The table below shows the bytes stored for every combination of offset and byte ordering.

Initial Memory contents and byte offsets								Contents of Source Register								
most	— significance —						least	most	— significance —						least	
0	1	2	3	4	5	6	7	big-	most	— significance —						least
i	j	k	l	m	n	o	p		A	B	C	D	E	F	G	H
7	6	5	4	3	2	1	0	little-endian								

Memory contents after instruction (shaded is unchanged)																
Big-endian byte ordering								$vAddr_{2..0}$	Little-endian byte ordering							
A	B	C	D	E	F	G	H	0	i	j	k	l	m	n	o	A
i	A	B	C	D	E	F	G	1	i	j	k	l	m	n	A	B
i	j	A	B	C	D	E	F	2	i	j	k	l	m	A	B	C
i	j	k	A	B	C	D	E	3	i	j	k	l	A	B	C	D
i	j	k	l	A	B	C	D	4	i	j	k	A	B	C	D	E
i	j	k	l	m	A	B	C	5	i	j	A	B	C	D	E	F
i	j	k	l	m	n	A	B	6	i	A	B	C	D	E	F	G
i	j	k	l	m	n	o	A	7	A	B	C	D	E	F	G	H

Table 2.31 Bytes Stored by SDL Instruction

Restrictions:

None

Operation: 64-bit processors

$$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$$

$$(\text{pAddr}, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{STORE})$$

$$\text{pAddr} \leftarrow \text{pAddr}_{(\text{PSIZE}-1)..3} \parallel (\text{pAddr}_{2..0} \text{ xor ReverseEndian}^3)$$

If BigEndianMem = 0 then

$$\text{pAddr} \leftarrow \text{pAddr}_{(\text{PSIZE}-1)..3} \parallel 0^3$$

endif


$$\text{byte} \leftarrow vAddr_{2..0} \text{ xor BigEndianCPU}^3$$

$$\text{datadouble} \leftarrow 0^{56-8*\text{byte}} \parallel \text{GPR}[\text{rt}]_{63..56-8*\text{byte}}$$

StoreMemory (uncached, byte, datadouble, pAddr, vAddr, DATA)

Exceptions:

- TLB Refill, TLB Invalid
- TLB Modified
- Bus Error
- Address Error
- Reserved Instruction



Format: SDR rt, offset(base) **MIPS III**

Purpose: To store the least-significant part of a doubleword to an unaligned memory address.

Description: memory[base+offset] ← Some_Bytes_From rt

The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address (*EffAddr*). *EffAddr* is the address of the least-significant of eight consecutive bytes forming a doubleword in memory (*DW*) starting at an arbitrary byte boundary. A part of *DW*, the least-significant one to eight bytes, is in the aligned doubleword containing *EffAddr*. The same number of least-significant (right) bytes of GPR *rt* are stored into these bytes of *DW*.

The figure below illustrates this operation for big-endian byte ordering. The eight consecutive bytes in 2..9 form an unaligned doubleword starting at location 2. A part of *DW*, two bytes, is contained in the aligned doubleword containing the least-significant byte at 9. First, SDR stores the two least-significant bytes of the source register into these bytes in memory. Next, the complementary SDL stores the remainder of *DW*.




Figure 2.16 Unaligned Doubleword Store with SDR and SDL

The bytes stored from the source register to memory depend on both the offset of the effective address within an aligned doubleword, i.e. the low three bits of the address ($vAddr_{2..0}$), and the current byte ordering mode of the processor (big- or little-endian). The table below shows the bytes stored for every combination of offset and byte ordering.

Initial Memory contents and byte offsets								Contents of Source Register															
most				— significance				least				most				— significance				least			
0	1	2	3	4	5	6	7	big-								0	1	2	3	4	5	6	7
i	j	k	l	m	n	o	p									A	B	C	D	E	F	G	H
7	6	5	4	3	2	1	0	little-endian															

Memory contents after instruction (shaded is unchanged)																
Big-endian byte ordering								$vAddr_{2..0}$	Little-endian byte ordering							
H	j	k	l	m	n	o	p	0	A	B	C	D	E	F	G	H
G	H	k	l	m	n	o	p	1	B	C	D	E	F	G	H	p
F	G	H	l	m	n	o	p	2	C	D	E	F	G	H	o	p
E	F	G	H	m	n	o	p	3	D	E	F	G	H	n	o	p
D	E	F	G	H	n	o	p	4	E	F	G	H	m	n	o	p
C	D	E	F	G	H	o	p	5	F	G	H	l	m	n	o	p
B	C	D	E	F	G	H	p	6	G	H	k	l	m	n	o	p
A	B	C	D	E	F	G	H	7	H	j	k	l	m	n	o	p

Table 2.32 Bytes Stored by SDR Instruction

Restrictions:

None

Operation: 64-bit processors

$$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$$

$$(\text{pAddr}, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{STORE})$$

$$\text{pAddr} \leftarrow \text{pAddr}_{(\text{PSIZE}-1)..3} \parallel (\text{pAddr}_{2..0} \text{ xor ReverseEndian}^3)$$

If BigEndianMem = 0 then

$$\text{pAddr} \leftarrow \text{pAddr}_{(\text{PSIZE}-1)..3} \parallel 0^3$$

endif

$$\text{byte} \leftarrow vAddr_{1..0} \text{ xor BigEndianCPU}^3$$

$$\text{datadouble} \leftarrow \text{GPR}[\text{rt}]_{63-8*\text{byte}} \parallel 0^{8*\text{byte}}$$

StoreMemory (uncached, DOUBLEWORD-byte, datadouble, pAddr, vAddr, DATA)

Exceptions:


TLB Refill, TLB Invalid

TLB Modified

Bus Error

Address Error

Reserved Instruction



Format: SH rt, offset(base) MIPS I

Purpose: To store a halfword to memory.

Description: memory[base+offset] ← rt

The least-significant 16-bit halfword of register *rt* is stored in memory at the location specified by the aligned effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

The effective address must be naturally aligned. If the least-significant bit of the address is non-zero, an Address Error exception occurs.

MIPS IV: The low-order bit of the *offset* field must be zero. If it is not, the result of the instruction is undefined.

Operation: 32-bit processors

$$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$$

$$\text{if } (vAddr_0) \neq 0 \text{ then SignalException(AddressError) endif}$$

$$(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{STORE})$$

$$pAddr \leftarrow pAddr_{PSIZE-1..2} \parallel (pAddr_{1..0} \text{ xor } (\text{ReverseEndian} \parallel 0))$$

$$\text{byte} \leftarrow vAddr_{1..0} \text{ xor } (\text{BigEndianCPU} \parallel 0)$$

$$\text{dataword} \leftarrow \text{GPR}[\text{rt}]_{31-8*\text{byte}..0} \parallel 0^{8*\text{byte}}$$

$$\text{StoreMemory}(\text{uncached}, \text{HALFWORD}, \text{dataword}, pAddr, vAddr, \text{DATA})$$

Operation: 64-bit processors

$$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$$

$$\text{if } (vAddr_0) \neq 0 \text{ then SignalException(AddressError) endif}$$

$$(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{STORE})$$

$$pAddr \leftarrow pAddr_{PSIZE-1..3} \parallel (pAddr_{2..0} \text{ xor } (\text{ReverseEndian}^2 \parallel 0))$$


$$\text{byte} \leftarrow vAddr_{2..0} \text{ xor } (\text{BigEndianCPU}^2 \parallel 0)$$

$$\text{datadouble} \leftarrow \text{GPR}[\text{rt}]_{63-8*\text{byte}..0} \parallel 0^{8*\text{byte}}$$

$$\text{StoreMemory}(\text{uncached}, \text{HALFWORD}, \text{datadouble}, pAddr, vAddr, \text{DATA})$$

Exceptions:

- TLB Refill, TLB Invalid
- TLB Modified
- Address Error



Format: SLL rd, rt, sa MIPS I

Purpose: To left shift a word by a fixed number of bits.

Description: $rd \leftarrow rt \ll sa$

The contents of the low-order 32-bit word of GPR *rt* are shifted left, inserting zeroes into the emptied bits; the word result is placed in GPR *rd*. The bit shift count is specified by *sa*. If *rd* is a 64-bit register, the result word is sign-extended.

Restrictions:

None

Operation:

$s \leftarrow sa$
 $temp \leftarrow GPR[rt]_{(31-s).0} \parallel 0^s$
 $GPR[rd] \leftarrow sign_extend(temp)$


Exceptions:

None

Programming Notes:

Unlike nearly all other word operations the input operand does not have to be a properly sign-extended word value to produce a valid sign-extended 32-bit result. The result word is always sign extended into a 64-bit destination register; this instruction with a zero shift amount truncates a 64-bit value to 32 bits and sign extends it.

Some assemblers, particularly 32-bit assemblers, treat this instruction with a shift amount of zero as a NOP and either delete it or replace it with an actual NOP.



Format: SLLV rd, rt, rs MIPS I

Purpose: To left shift a word by a variable number of bits.

Description: $rd \leftarrow rt \ll rs$

The contents of the low-order 32-bit word of GPR *rt* are shifted left, inserting zeroes into the emptied bits; the result word is placed in GPR *rd*. The bit shift count is specified by the low-order five bits of GPR *rs*. If *rd* is a 64-bit register, the result word is sign-extended.

Restrictions:

None

Operation:

$$s \leftarrow GP[rs]_{4..0}$$

$$temp \leftarrow GPR[rt]_{(31-s)..0} \parallel 0^s$$

$$GPR[rd] \leftarrow sign_extend(temp)$$


Exceptions:

None

Programming Notes:

Unlike nearly all other word operations the input operand does not have to be a properly sign-extended word value to produce a valid sign-extended 32-bit result. The result word is always sign extended into a 64-bit destination register; this instruction with a zero shift amount truncates a 64-bit value to 32 bits and sign extends it.

Some assemblers, particularly 32-bit assemblers, treat this instruction with a shift amount of zero as a NOP and either delete it or replace it with an actual NOP.



Format: SLT rd, rs, rt MIPS I

Purpose: To record the result of a less-than comparison.

Description: $rd \leftarrow (rs < rt)$

Compare the contents of GPR *rs* and GPR *rt* as signed integers and record the Boolean result of the comparison in GPR *rd*. If GPR *rs* is less than GPR *rt* the result is 1 (true), otherwise 0 (false).

The arithmetic comparison does not cause an Integer Overflow exception.

Restrictions:

None


Operation:

```

if GPR[rs] < GPR[rt] then
    GPR[rd] ← 0GPRLEN-1 || 1
else
    GPR[rd] ← 0GPRLEN
endif
    
```

Exceptions:

None



Format: SLTI *rt*, *rs*, *immediate* MIPS I

Purpose: To record the result of a less-than comparison with a constant.

Description: $rt \leftarrow (rs < immediate)$

Compare the contents of GPR *rs* and the 16-bit signed *immediate* as signed integers and record the Boolean result of the comparison in GPR *rt*. If GPR *rs* is less than *immediate* the result is 1 (true), otherwise 0 (false).

The arithmetic comparison does not cause an Integer Overflow exception.

Restrictions:

None


Operation:

```

if GPR[rs] < sign_extend(immediate) then
    GPR[rd] ← 0GPREN-1 || 1
else
    GPR[rd] ← 0GPREN
endif
    
```

Exceptions:

None



Format: SLTIU rt, rs, immediate MIPS I

Purpose: To record the result of an unsigned less-than comparison with a constant.

Description: $rt \leftarrow (rs < immediate)$

Compare the contents of GPR *rs* and the sign-extended 16-bit *immediate* as unsigned integers and record the Boolean result of the comparison in GPR *rt*. If GPR *rs* is less than *immediate* the result is 1 (true), otherwise 0 (false).

Because the 16-bit *immediate* is sign-extended before comparison, the instruction is able to represent the smallest or largest unsigned numbers. The representable values are at the minimum [0, 32767] or maximum [max_unsigned-32767, max_unsigned] end of the unsigned range.

The arithmetic comparison does not cause an Integer Overflow exception.

Restrictions:

None


Operation:

```

if (0 || GPR[rs]) < (0 || sign_extend(immediate)) then
    GPR[rd] ← 0GPREN-1 || 1
else
    GPR[rd] ← 0GPREN
endif
    
```

Exceptions:

None



Format: SLTU rd, rs, rt MIPS I

Purpose: To record the result of an unsigned less-than comparison.

Description: $rd \leftarrow (rs < rt)$

Compare the contents of GPR *rs* and GPR *rt* as unsigned integers and record the Boolean result of the comparison in GPR *rd*. If GPR *rs* is less than GPR *rt* the result is 1 (true), otherwise 0 (false).

The arithmetic comparison does not cause an Integer Overflow exception.

Restrictions:

None


Operation:

```

if (0 || GPR[rs]) < (0 || GPR[rt]) then
  GPR[rd] ← 0GPRLEN-1 || 1
else
  GPR[rd] ← 0GPRLEN
endif
    
```

Exceptions:

None



Format: SRA rd, rt, sa MIPS I

Purpose: To arithmetic right shift a word by a fixed number of bits.

Description: $rd \leftarrow rt \gg sa$ (arithmetic)

The contents of the low-order 32-bit word of GPR *rt* are shifted right, duplicating the sign-bit (bit 31) in the emptied bits; the word result is placed in GPR *rd*. The bit shift count is specified by *sa*. If *rd* is a 64-bit register, the result word is sign-extended.

Restrictions:

On 64-bit processors, if GPR *rt* does not contain a sign-extended 32-bit value (bits 63..31 equal) then the result of the operation is undefined.


Operation:

```

if (NotWordValue(GPR[rt])) then UndefinedResult() endif
s ← sa
temp ← (GPR[rt]31)s || GPR[rt]31..s
GPR[rd] ← sign_extend(temp)
    
```

Exceptions:

None



Format: SRAV rd, rt, rs MIPS I

Purpose: To arithmetic right shift a word by a variable number of bits.

Description: $rd \leftarrow rt \gg rs$ (arithmetic)

The contents of the low-order 32-bit word of GPR *rt* are shifted right, duplicating the sign-bit (bit 31) in the emptied bits; the word result is placed in GPR *rd*. The bit shift count is specified by the low-order five bits of GPR *rs*. If *rd* is a 64-bit register, the result word is sign-extended.

Restrictions:

On 64-bit processors, if GPR *rt* does not contain a sign-extended 32-bit value (bits 63..31 equal) then the result of the operation is undefined.


Operation:

```

if (NotWordValue(GPR[rt])) then UndefinedResult() endif
s ← GPR[rs]4..0
temp ← (GPR[rt]31)s || GPR[rt]31..s
GPR[rd] ← sign_extend(temp)
    
```

Exceptions:

None



Format: SRL rd, rt, sa MIPS I

Purpose: To logical right shift a word by a fixed number of bits.

Description: $rd \leftarrow rt \gg sa$ (logical)

The contents of the low-order 32-bit word of GPR *rt* are shifted right, inserting zeros into the emptied bits; the word result is placed in GPR *rd*. The bit shift count is specified by *sa*. If *rd* is a 64-bit register, the result word is sign-extended.

Restrictions:

On 64-bit processors, if GPR *rt* does not contain a sign-extended 32-bit value (bits 63..31 equal) then the result of the operation is undefined.


Operation:

```

if (NotWordValue(GPR[rt])) then UndefinedResult() endif
s      ← sa
temp   ← 0s || GPR[rt]31..s
GPR[rd] ← sign_extend(temp)
    
```

Exceptions:

None



Format: SRLV rd, rt, rs MIPS I

Purpose: To logical right shift a word by a variable number of bits.

Description: $rd \leftarrow rt \gg rs$ (logical)

The contents of the low-order 32-bit word of GPR *rt* are shifted right, inserting zeros into the emptied bits; the word result is placed in GPR *rd*. The bit shift count is specified by the low-order five bits of GPR *rs*. If *rd* is a 64-bit register, the result word is sign-extended.

Restrictions:

On 64-bit processors, if GPR *rt* does not contain a sign-extended 32-bit value (bits 63..31 equal) then the result of the operation is undefined.


Operation:

```

if (NotWordValue(GPR[rt])) then UndefinedResult() endif
s ← GPR[rs]4..0
temp ← 0s || GPR[rt]31..s
GPR[rd] ← sign_extend(temp)
    
```

Exceptions:

None



Format: SUB rd, rs, rt **MIPS I**

Purpose: To subtract 32-bit integers. If overflow occurs, then trap.

Description: $rd \leftarrow rs - rt$

The 32-bit word value in GPR *rt* is subtracted from the 32-bit value in GPR *rs* to produce a 32-bit result. If the subtraction results in 32-bit 2's complement arithmetic overflow then the destination register is not modified and an Integer Overflow exception occurs. If it does not overflow, the 32-bit result is placed into GPR *rd*.

Restrictions:

On 64-bit processors, if either GPR *rt* or GPR *rs* do not contain sign-extended 32-bit values (bits 63..31 equal), then the result of the operation is undefined.

Operation:

```


if (NotWordValue(GPR[rs]) or NotWordValue(GPR[rt])) then UndefinedResult() endif
temp ← GPR[rs] - GPR[rt]
if (32_bit_arithmetic_overflow) then
    SignalException(IntegerOverflow)
else
    GPR[rd] ← temp
endif
    
```

Exceptions:

Integer Overflow

Programming Notes:

SUBU performs the same arithmetic operation but, does not trap on overflow.



Format: SUBU rd, rs, rt MIPS I

Purpose: To subtract 32-bit integers.

Description: $rd \leftarrow rs - rt$

The 32-bit word value in GPR *rt* is subtracted from the 32-bit value in GPR *rs* and the 32-bit arithmetic result is placed into GPR *rd*.

No integer overflow exception occurs under any circumstances.

Restrictions:

On 64-bit processors, if either GPR *rt* or GPR *rs* do not contain sign-extended 32-bit values (bits 63..31 equal), then the result of the operation is undefined.

Operation:


```
if (NotWordValue(GPR[rs]) or NotWordValue(GPR[rt])) then UndefinedResult() endif
temp ← GPR[rs] - GPR[rt]
GPR[rd] ← temp
```

Exceptions:

None

Programming Notes:

The term “unsigned” in the instruction name is a misnomer; this operation is 32-bit modulo arithmetic that does not trap on overflow. It is appropriate for arithmetic which is not signed, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as “C” language arithmetic.



Format: SW rt, offset(base) **MIPS I**

Purpose: To store a word to memory.

Description: memory[base+offset] ← rt

The least-significant 32-bit word of register *rt* is stored in memory at the location specified by the aligned effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

The effective address must be naturally aligned. If either of the two least-significant bits of the address are non-zero, an Address Error exception occurs.

MIPS IV: The low-order 2 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation: 32-bit Processors


```
vAddr ← sign_extend(offset) + GPR[base]
if (vAddr1..0) ≠ 02 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, STORE)
dataword ← GPR[rt]
StoreMemory (uncached, WORD, dataword, pAddr, vAddr, DATA)
```

Operation: 64-bit Processors

```
vAddr ← sign_extend(offset) + GPR[base]
if (vAddr1..0) ≠ 02 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
byte ← vAddr2..0 xor (BigEndianCPU || 02)
datadouble ← GPR[rt]63-8*byte || 08*byte
StoreMemory (uncached, WORD, datadouble, pAddr, vAddr, DATA)
```

Exceptions:

- TLB Refill, TLB Invalid
- TLB Modified
- Address Error



Format: SWC1 rt, offset(base) MIPS I
 SWC2 rt, offset(base)
 SWC3 rt, offset(base)

Purpose: To store a word from a coprocessor general register to memory.

Description: memory[base+offset] ← rt

Coprocessor unit zz supplies a 32-bit word which is stored at the memory location specified by the aligned effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

The data supplied by each coprocessor is defined by the individual coprocessor specifications. The usual operation would read the data from coprocessor general register *rt*.

Each MIPS architecture level defines up to 4 coprocessor units, numbered 0 to 3. The opcodes corresponding to coprocessors that are not defined by an architecture level may be used for other instructions.

Restrictions:

Access to the coprocessors is controlled by system software. Each coprocessor has a “coprocessor usable” bit in the System Control coprocessor. The usable bit must be set for a user program to execute a coprocessor instruction. If the usable bit is not set, an attempt to execute the instruction will result in a Coprocessor Unusable exception. An unimplemented coprocessor must never be enabled. The result of executing this instruction for an unimplemented coprocessor when the usable bit is set, is undefined.

This instruction is not available for coprocessor 0, the System Control coprocessor, and the opcode may be used for other instructions.

The effective address must be naturally aligned. If either of the two least-significant bits of the address are non-zero, an Address Error exception occurs.


MIPS IV: The low-order 2 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation: 32-bit processors
 $vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 if $(vAddr_{1..0}) \neq 0^2$ then SignalException(AddressError) endif
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{STORE})$
 $\text{dataword} \leftarrow \text{COP_SW}(z, rt)$
 StoreMemory(uncached, WORD, dataword, pAddr, vAddr, DATA)

Operation: 64-bit processors
 $vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 if $(vAddr_{1..0}) \neq 0^2$ then SignalException(AddressError) endif
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{STORE})$
 $pAddr \leftarrow pAddr_{\text{PSIZE}-1..3} \parallel (pAddr_{2..0} \text{ xor } (\text{ReverseEndian} \parallel 0^2))$
 $\text{byte} \leftarrow vAddr_{2..0} \text{ xor } (\text{BigEndianCPU} \parallel 0^2)$
 $\text{dataword} \leftarrow \text{COP_SW}(z, rt)$
 $\text{datadouble} \leftarrow 0^{32-8*\text{byte}} \parallel \text{dataword} \parallel 0^{8*\text{byte}}$
 StoreMemory(uncached, WORD, datadouble, pAddr, vAddr DATA)

Exceptions:

- TLB Refill, TLB Invalid
- TLB Modified
- Address Error
- Reserved Instruction
- Coprocessor Unusable



Format: SWL rt, offset(base) **MIPS I**

Purpose: To store the most-significant part of a word to an unaligned memory address.

Description: memory[base+offset] ← rt

The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address (*EffAddr*). *EffAddr* is the address of the most-significant of four consecutive bytes forming a word in memory (*W*) starting at an arbitrary byte boundary. A part of *W*, the most-significant one to four bytes, is in the aligned word containing *EffAddr*. The same number of the most-significant (left) bytes from the word in GPR *rt* are stored into these bytes of *W*.

If GPR *rt* is a 64-bit register, the source word is the low word of the register.

Figure 1.2.11 illustrates this operation for big-endian byte ordering for 32-bit and 64-bit registers. The four consecutive bytes in 2..5 form an unaligned word starting at location 2. A part of *W*, two bytes, is contained in the aligned word containing the most-significant byte at 2. First, SWL stores the most-significant two bytes of the low-word from the source register into these two bytes in memory. Next, the complementary SWR stores the remainder of the unaligned word.

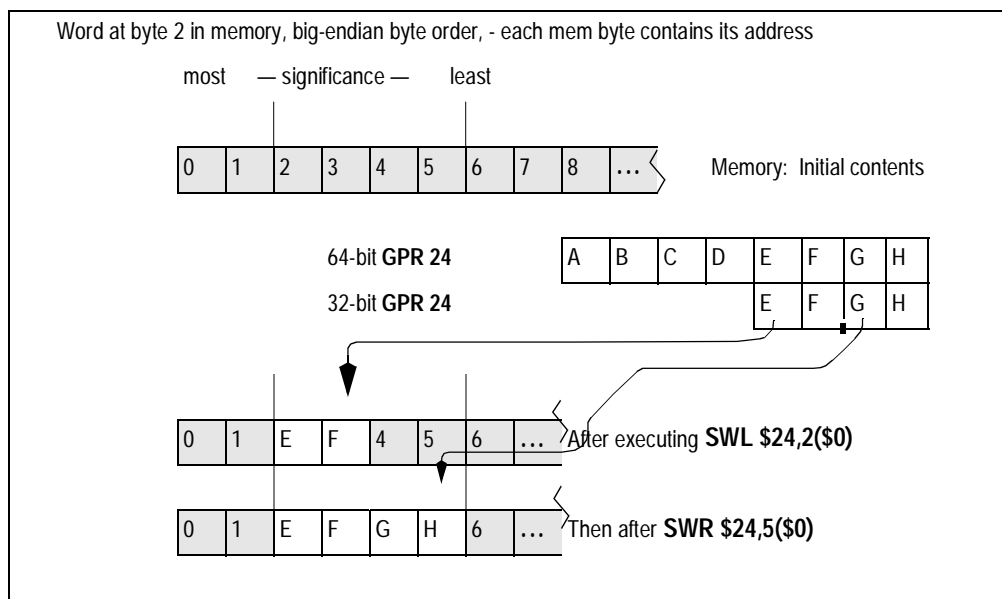


Table 2.33 Unaligned Word Store using SWL and SWR

The bytes stored from the source register to memory depend on both the offset of the effective address within an aligned word, i.e. the low two bits of the address ($vAddr_{1..0}$), and the current byte ordering mode of the processor (big- or little-endian). The table below shows the bytes stored for every combination of offset and byte ordering.

Memory contents and byte offsets				Initial contents of Dest Register							
0	1	2	3	big-endian							
i	j	k	l	offset ($vAddr_{1..0}$)							
3	2	1	0	little-endian							
most				least							
— significance —				64-bit register							
				A	B	C	D	E	F	G	H
				most — significance — least							
				32-bit register							
				E	F	G	H				

Memory contents after instruction (shaded is unchanged)											
Big-endian byte ordering				$vAddr_{1..0}$	Little-endian byte ordering						
E	F	G	H	0	i	j	k	E			
i	E	F	G	1	i	j	E	F			
i	j	E	F	2	i	E	F	G			
i	j	k	E	3	E	F	G	H			

Table 2.34 Bytes Stored by SWL Instruction

Operation: 32-bit Processors


$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{STORE})$
 $pAddr \leftarrow pAddr_{(P\text{SIZE}-1)..2} \parallel (pAddr_{1..0} \text{ xor ReverseEndian}^2)$
 If BigEndianMem = 0 then
 $pAddr \leftarrow pAddr_{(P\text{SIZE}-1)..2} \parallel 0^2$
 endif
 $\text{byte} \leftarrow vAddr_{1..0} \text{ xor BigEndianCPU}^2$
 $\text{dataword} \leftarrow 0^{24-8*\text{byte}} \parallel \text{GPR}[\text{rt}]_{31..24-8*\text{byte}}$
 StoreMemory(uncached, byte, dataword, pAddr, vAddr, DATA)

Operation: 64-bit Processors

$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{STORE})$
 $pAddr \leftarrow pAddr_{(P\text{SIZE}-1)..3} \parallel (pAddr_{2..0} \text{ xor ReverseEndian}^3)$
 If BigEndianMem = 0 then
 $pAddr \leftarrow pAddr_{(P\text{SIZE}-1)..2} \parallel 0^2$
 endif
 $\text{byte} \leftarrow vAddr_{1..0} \text{ xor BigEndianCPU}^2$
 if $(vAddr_2 \text{ xor BigEndianCPU}) = 0$ then
 $\text{datadouble} \leftarrow 0^{32} \parallel 0^{24-8*\text{byte}} \parallel \text{GPR}[\text{rt}]_{31..24-8*\text{byte}}$
 else
 $\text{datadouble} \leftarrow 0^{24-8*\text{byte}} \parallel \text{GPR}[\text{rt}]_{31..24-8*\text{byte}} \parallel 0^{32}$
 endif
 StoreMemory(uncached, byte, datadouble, pAddr, vAddr, DATA)

Exceptions:

- TLB Refill, TLB Invalid
- TLB Modified
- Bus Error
- Address Error



Format: SWR rt, offset(base) **MIPS I**

Purpose: To store the least-significant part of a word to an unaligned memory address.

Description: memory[base+offset] ← rt

The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address (*EffAddr*). *EffAddr* is the address of the least-significant of four consecutive bytes forming a word in memory (*W*) starting at an arbitrary byte boundary. A part of *W*, the least-significant one to four bytes, is in the aligned word containing *EffAddr*. The same number of the least-significant (right) bytes from the word in GPR *rt* are stored into these bytes of *W*.

If GPR *rt* is a 64-bit register, the source word is the low word of the register.

Figure 1.2.11 illustrates this operation for big-endian byte ordering for 32-bit and 64-bit registers. The four consecutive bytes in 2..5 form an unaligned word starting at location 2. A part of *W*, two bytes, is contained in the aligned word containing the least-significant byte at 5. First, SWR stores the least-significant two bytes of the low-word from the source register into these two bytes in memory. Next, the complementary SWL stores the remainder of the unaligned word.

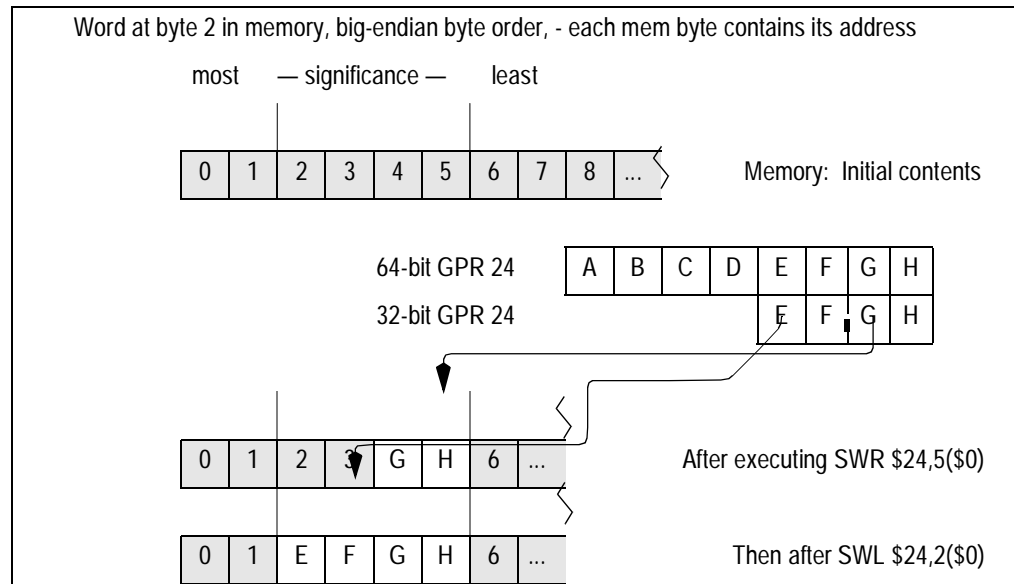


Figure 2.17 Unaligned Word Store using SWR and SWL

The bytes stored from the source register to memory depend on both the offset of the effective address within an aligned word, i.e. the low two bits of the address ($vAddr_{1..0}$), and the current byte ordering mode of the processor (big- or little-endian). The table below shows the bytes stored for every combination of offset and byte ordering.

Memory contents and byte offsets				Initial contents of Dest Register							
0	1	2	3	← big-endian							
i	j	k	l	offset ($vAddr_{1..0}$)							
3	2	1	0	← little-endian							
most		least		—significance—							
—significance—				64-bit register							
				A	B	C	D	E	F	G	H
				32-bit register				E	F	G	H
Memory contents after instruction (shaded is unchanged)											
Big-endian byte ordering				$vAddr_{1..0}$	Little-endian byte ordering						
				0							
H	j	k	l	0	E	F	G	H			
G	H	k	l	1	F	G	H	l			
F	G	H	l	2	G	H	k	l			
E	F	G	H	3	H	j	k	l			

Table 2.35 Bytes Stored by SWR Instruction

Restrictions:

None

Operation: 32-bit Processors

$vAddr \leftarrow \text{sign_extend}(\text{offset}) + \text{GPR}[\text{base}]$
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{STORE})$
 $pAddr \leftarrow pAddr_{(P\text{SIZE}-1)..2} \parallel (pAddr_{1..0} \text{ xor ReverseEndian}^2)$
 BigEndianMem = 0 then
 $pAddr \leftarrow pAddr_{(P\text{SIZE}-1)..2} \parallel 0^2$
 endif
 $\text{byte} \leftarrow vAddr_{1..0} \text{ xor BigEndianCPU}^2$
 $\text{dataword} \leftarrow \text{GPR}[\text{rt}]_{31-8*\text{byte}} \parallel 0^{8*\text{byte}}$
 StoreMemory (uncached, WORD-byte, dataword, pAddr, vAddr, DATA)


Operation: 64-bit Processors

```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, uncached) ← AddressTranslation(vAddr, DATA, STORE)
pAddr ← pAddr(PSIZE-1)..3 || (pAddr2..0 xor ReverseEndian3)
If BigEndianMem = 0 then
    pAddr ← pAddr(PSIZE-1)..2 || 02
endif
byte ← vAddr1..0 xor BigEndianCPU2
if (vAddr2 xor BigEndianCPU) = 0 then
    datadouble ← 032 || GPR[rt]31-8*byte..0 || 08*byte
else
    datadouble ← GPR[rt]31-8*byte..0 || 08*byte || 032
endif
StoreMemory(uncached, WORD-byte, datadouble, pAddr, vAddr, DATA)
    
```

Exceptions:

- TLB Refill, TLB Invalid
- TLB Modified
- Bus Error
- Address Error



Format: SYNC (stype = 0 implied) MIPS II

Purpose: To order loads and stores to shared memory in a multiprocessor system.

Description:

To serve a broad audience, two descriptions are given. A simple description of SYNC that appeals to intuition is followed by a precise and detailed description.

A Simple Description:

SYNC affects only uncached and cached coherent loads and stores. The loads and stores that occur prior to the SYNC must be completed before the loads and stores after the SYNC are allowed to start.

Loads are completed when the destination register is written. Stores are completed when the stored value is visible to every other processor in the system.

A Precise Description:

If the *stype* field has a value of zero, every synchronizable load and store that occurs in the instruction stream prior to the SYNC instruction must be globally performed before any synchronizable load or store that occurs after the SYNC may be performed with respect to any other processor or coherent I/O module.

Sync does not guarantee the order in which instruction fetches are performed.

The *stype* values 1-31 are reserved; they produce the same result as the value zero.

Synchronizable: A load or store instruction is *synchronizable* if the load or store occurs to a physical location in shared memory using a virtual location with a memory access type of either uncached or cached coherent. *Shared memory* is memory that can be accessed by more than one processor or by a coherent I/O system module.

Performed load: A load instruction is *performed* when the value returned by the load has been determined. The result of a load on processor A has been *determined* with respect to processor or coherent I/O module B when a subsequent store to the location by B cannot affect the value returned by the load. The store by B must use the same memory access type as the load.

Performed store: A store instruction is *performed* when the store is observable. A store on processor A is *observable* with respect to processor or coherent I/O module B when a subsequent load of the location by B returns the value written by the store. The load by B must use the same memory access type as the store.

Globally performed load: A load instruction is *globally performed* when it is performed with respect to all processors and coherent I/O modules capable of storing to the location.

Globally performed store: A store instruction is *globally performed* when it is globally observable. It is *globally observable* when it is observable by all processors and I/O modules capable of loading from the location.

Coherent I/O module: A *coherent I/O module* is an Input/Output system component that performs coherent Direct Memory Access (DMA). It reads and writes memory independently as though it were a processor doing loads and stores to locations with a memory access type of cached coherent.

Restrictions:

The effect of SYNC on the global order of the effects of loads and stores for memory access types other than uncached and cached coherent is not defined.

Operation:

SyncOperation(stype)

Exceptions:

Reserved Instruction

Programming Notes:

A processor executing load and store instructions observes the effects of the loads and stores that use the same memory access type in the order that they occur in the instruction stream; this is known as *program order*. A *parallel program* has multiple instruction streams that can execute at the same time on different processors. In multiprocessor (MP) systems, the order in which the effects of loads and stores are observed by other processors, the *global order* of the loads and stores, determines the actions necessary to reliably share data in parallel programs.

When all processors observe the effects of loads and stores in program order, the system is *strongly ordered*. On such systems, parallel programs can reliably share data without explicit actions in the programs. For such a system, SYNC has the same effect as a NOP. Executing SYNC on such a system is not necessary, but is also not an error.

If a multiprocessor system is not strongly ordered, the effects of load and store instructions executed by one processor may be observed out of program order by other processors. On such systems, parallel programs must take explicit actions in order to reliably share data. At critical points in the program, the effects of loads and stores from an instruction stream must occur in the same order for all processors. SYNC separates the loads and stores executed on the processor into two groups and the effects of these **groups** are seen in program order by all processors. The effect of all loads and stores in one group is seen by all processors before the effect of any load or store in the other group. In effect, SYNC causes the system to be strongly ordered for the executing processor at the instant that the SYNC is executed.

Many MIPS-based multiprocessor systems are strongly ordered or have a mode in which they operate as strongly ordered for at least one memory access type. The MIPS architecture also permits MP systems that are not strongly ordered. SYNC enables the reliable use of shared memory on such systems. A parallel program that does not use SYNC will generally not operate on a system that is not strongly ordered, however a program that does use SYNC will work on both types of systems. System-specific documentation will describe the actions necessary to reliably share data in parallel programs for that system.

The behavior of a load or store using one memory access type is undefined if a load or store was previously made to the same physical location using a different memory access type. The presence of a SYNC between the references does not alter this behavior.

SYNC affects the order in which the effects of load and store instructions appears to all processors; it not generally affect the **physical** memory-system ordering or synchronization issues that arise in system programming. The effect of SYNC on implementation specific aspects of the cached memory system, such as writeback buffers, is not defined. The effect of SYNC on reads or writes to memory caused by privileged implementation-specific instructions, such as CACHE, is not defined.

Prefetch operations have no effects detectable by user-mode programs so ordering the effects of prefetch operations is not meaningful.


EXAMPLE: These code fragments show how SYNC can be used to coordinate the use of shared data between separate writer and reader instruction streams in a multiprocessor environment. The FLAG location is used by the instruction streams to determine whether the shared data item DATA is valid. The SYNC executed by processor A forces the store of DATA to be performed globally before the store to FLAG is performed. The SYNC executed by processor B ensures that DATA is not read until after the FLAG value indicates that the shared data is valid.

Processor A (writer)			
# Conditions at entry:			
# The value 0 has been stored in FLAG and that value is observable by B.			
SW	R1, DATA		# change shared DATA value
LI	R2, 1		
SYNC			# perform DATA store before performing FLAG store
SW	R2, FLAG		# say that the shared DATA value is valid

Processor B (reader)			
	LI	R2, 1	
1:	LW	R1, FLAG	# get FLAG
	BNE	R2, R1, 1B	# if it says that DATA is not valid, poll again
	NOP		
	SYNC		# FLAG value checked before doing DATA reads
	LW	R1, DATA	# read (valid) shared DATA values

Implementation Notes:

There may be side effects of uncached loads and stores that affect cached coherent load and store operations. To permit the reliable use of such side effects, buffered uncached stores that occur before the SYNC must be written to memory before cached coherent loads and stores after the SYNC may be performed.



Format: SYSCALL MIPS I

Purpose: To cause a System Call exception.

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.

Restrictions:


None

Operation:

SignalException(SystemCall)

Exceptions:

System Call



Format: TEQ rs, rt MIPS II

Purpose: To compare GPRs and do a conditional Trap.

Description: if (rs = rt) then Trap

Compare the contents of GPR *rs* and GPR *rt* as signed integers; if GPR *rs* is equal to GPR *rt* then take a Trap exception.

The contents of the *code* field are ignored by hardware and may be used to encode information for system software. To retrieve the information, system software must load the instruction word from memory.

Restrictions:

None


Operation:

```

if GPR[rs] = GPR[rt] then
    SignalException(Trap)
endif
    
```

Exceptions:

Reserved Instruction
Trap



Format: TEQI rs, immediate MIPS II

Purpose: To compare a GPR to a constant and do a conditional Trap.

Description: if (rs = immediate) then Trap

Compare the contents of GPR *rs* and the 16-bit signed *immediate* as signed integers; if GPR *rs* is equal to *immediate* then take a Trap exception.

Restrictions:

None


Operation:

```

if GPR[rs] = sign_extend(immediate) then
    SignalException(Trap)
endif
    
```

Exceptions:

Reserved Instruction
Trap



Format: TGE rs, rt MIPS II

Purpose: To compare GPRs and do a conditional Trap.

Description: if (rs ≥ rt) then Trap

Compare the contents of GPR *rs* and GPR *rt* as signed integers; if GPR *rs* is greater than or equal to GPR *rt* then take a Trap exception.

The contents of the *code* field are ignored by hardware and may be used to encode information for system software. To retrieve the information, system software must load the instruction word from memory.

Restrictions:

None


Operation:

```

if GPR[rs] ≥ GPR[rt] then
    SignalException(Trap)
endif
    
```

Exceptions:

Reserved Instruction
Trap



Format: TGEI rs, immediate MIPS II

Purpose: To compare a GPR to a constant and do a conditional Trap.

Description: if (rs ≥ immediate) then Trap

Compare the contents of GPR *rs* and the 16-bit signed *immediate* as signed integers; if GPR *rs* is greater than or equal to *immediate* then take a Trap exception.

Restrictions:


None

Operation:

```
if GPR[rs] ≥ sign_extend(immediate) then
    SignalException(Trap)
endif
```

Exceptions:

Reserved Instruction
Trap



Format: TGEIU rs, immediate MIPS II

Purpose: To compare a GPR to a constant and do a conditional Trap.

Description: if (rs ≥ immediate) then Trap

Compare the contents of GPR *rs* and the 16-bit sign-extended *immediate* as unsigned integers; if GPR *rs* is greater than or equal to *immediate* then take a Trap exception.

Because the 16-bit *immediate* is sign-extended before comparison, the instruction is able to represent the smallest or largest unsigned numbers. The representable values are at the minimum [0, 32767] or maximum [max_unsigned-32767, max_unsigned] end of the unsigned range.

Restrictions:


None

Operation:

```
if (0 || GPR[rs]) ≥ (0 || sign_extend(immediate)) then
    SignalException(Trap)
endif
```

Exceptions:

- Reserved Instruction
- Trap



Format: TGEU rs, rt MIPS II

Purpose: To compare GPRs and do a conditional Trap.

Description: if (rs ≥ rt) then Trap

Compare the contents of GPR *rs* and GPR *rt* as unsigned integers; if GPR *rs* is greater than or equal to GPR *rt* then take a Trap exception.

The contents of the *code* field are ignored by hardware and may be used to encode information for system software. To retrieve the information, system software must load the instruction word from memory.

Restrictions:


None

Operation:

```
if (0 || GPR[rs]) ≥ (0 || GPR[rt]) then
    SignalException(Trap)
endif
```

Exceptions:

Reserved Instruction
Trap



Format: TLT rs, rt MIPS II

Purpose: To compare GPRs and do a conditional Trap.

Description: if (rs < rt) then Trap

Compare the contents of GPR *rs* and GPR *rt* as signed integers; if GPR *rs* is less than GPR *rt* then take a Trap exception.

The contents of the *code* field are ignored by hardware and may be used to encode information for system software. To retrieve the information, system software must load the instruction word from memory.

Restrictions:


None

Operation:

```
if GPR[rs] < GPR[rt] then
    SignalException(Trap)
endif
```

Exceptions:

Reserved Instruction
Trap



Format: TLTI rs, immediate MIPS II

Purpose: To compare a GPR to a constant and do a conditional Trap.

Description: if (rs < immediate) then Trap

Compare the contents of GPR *rs* and the 16-bit signed *immediate* as signed integers; if GPR *rs* is less than *immediate* then take a Trap exception.

Restrictions:


None

Operation:

```
if GPR[rs] < sign_extend(immediate) then
    SignalException(Trap)
endif
```

Exceptions:

Reserved Instruction
Trap



Format: TLTIU rs, immediate MIPS II

Purpose: To compare a GPR to a constant and do a conditional Trap.

Description: if (rs < immediate) then Trap

Compare the contents of GPR *rs* and the 16-bit sign-extended *immediate* as unsigned integers; if GPR *rs* is less than *immediate* then take a Trap exception.

Because the 16-bit *immediate* is sign-extended before comparison, the instruction is able to represent the smallest or largest unsigned numbers. The representable values are at the minimum [0, 32767] or maximum [max_unsigned-32767, max_unsigned] end of the unsigned range.

Restrictions:


None

Operation:

```
if (0 || GPR[rs]) < (0 || sign_extend(immediate)) then
    SignalException(Trap)
endif
```

Exceptions:

Reserved Instruction
Trap



Format: TLTU rs, rt MIPS II

Purpose: To compare GPRs and do a conditional Trap.

Description: if (rs < rt) then Trap

Compare the contents of GPR *rs* and GPR *rt* as unsigned integers; if GPR *rs* is less than GPR *rt* then take a Trap exception.

The contents of the *code* field are ignored by hardware and may be used to encode information for system software. To retrieve the information, system software must load the instruction word from memory.

Restrictions:


None

Operation:

```
if (0 || GPR[rs]) < (0 || GPR[rt]) then
    SignalException(Trap)
endif
```

Exceptions:

Reserved Instruction
Trap



Format: TNE rs, rt MIPS II

Purpose: To compare GPRs and do a conditional Trap.

Description: if (rs ≠ rt) then Trap

Compare the contents of GPR *rs* and GPR *rt* as signed integers; if GPR *rs* is not equal to GPR *rt* then take a Trap exception.

The contents of the *code* field are ignored by hardware and may be used to encode information for system software. To retrieve the information, system software must load the instruction word from memory.

Restrictions:


None

Operation:

```
if GPR[rs] ≠ GPR[rt] then
    SignalException(Trap)
endif
```

Exceptions:

Reserved Instruction
Trap



Format: TNEI rs, immediate MIPS II

Purpose: To compare a GPR to a constant and do a conditional Trap.

Description: if (rs ≠ immediate) then Trap

Compare the contents of GPR *rs* and the 16-bit signed *immediate* as signed integers; if GPR *rs* is not equal to *immediate* then take a Trap exception.

Restrictions:


None

Operation:

```
if GPR[rs] ≠ sign_extend(immediate) then
    SignalException(Trap)
endif
```

Exceptions:

Reserved Instruction
Trap



Format: TLBP

Description:

The *Index* register is loaded with the address of the TLB entry whose contents match the contents of the *EntryHi* register. If no TLB entry matches, the high-order bit of the *Index* register is set.


The architecture does not specify the operation of memory references associated with the instruction immediately after a TLBP instruction, nor is the operation specified if more than one TLB entry matches.

Operation:

```
T:  Index ← 1 || 0 31
    for i in 0..TLBEntries-1
        if (TLB[i]167..141 and not (015 || TLB[i]216..205))
            = EntryHi39..13) and not (015 || TLB[i]216..205)) and
            (TLB[i]140 or (TLB[i]135..128 = EntryHi7..0)) then
                Index ← 026 || i 5..0
            endif
        endfor
```

Exceptions:

Coprocessor unusable exception



Format: TLBR

Description:

The *G* bit (which controls ASID matching) read from the TLB is written into both of the *EntryLo0* and *EntryLo1* registers.


The *EntryHi* and *EntryLo* registers are loaded with the contents of the TLB entry pointed at by the contents of the TLB *Index* register. The operation is invalid (and the results are unspecified) if the contents of the TLB *Index* register are greater than the number of TLB entries in the processor.

Operation:

T: PageMask ~ TLB[Index5..0]255..192
 EntryHi ~ TLB[Index5..0]191..128 and not TLB[Index5..0]255..192
 EntryLo1 ~ TLB[Index5..0]127..65 || TLB[Index5..0]140
 EntryLo0 ~ TLB[Index5..0]63..1 || TLB[Index5..0]140

Exceptions:

Coprocessor unusable exception



Format: TLBWI

Description:

The *G* bit of the TLB is written with the logical AND of the *G* bits in the *EntryLo0* and *EntryLo1* registers.

The TLB entry pointed at by the contents of the TLB *Index* register is loaded with the contents of the *EntryHi* and *EntryLo* registers.

The operation is invalid (and the results are unspecified) if the contents of the TLB *Index* register are greater than the number of TLB entries in the processor.

Operation:

$$T: TLB[Index5..0] \leftarrow PageMask \parallel (EntryHi \text{ and not } PageMask) \parallel EntryLo1 \parallel EntryLo0$$

Exceptions:

Coprocessor unusable exception



Format: TLBWR

Description:

The *G* bit of the TLB is written with the logical AND of the *G* bits in the *EntryLo0* and *EntryLo1* registers.


The TLB entry pointed at by the contents of the TLB *Random* register is loaded with the contents of the *EntryHi* and *EntryLo* registers.


Operation:

$$T: \text{TLB}[\text{Random}5..0] = \text{PageMask} \parallel (\text{EntryHi and not PageMask}) \parallel \text{EntryLo1} \parallel \text{EntryLo0}$$

Exceptions:

Coprocessor unusable exception





Format: WAIT

Purpose: To stop the internal pipeline and reduce power used by the CPU.

Description:


The WAIT instruction is used to halt the internal pipeline and thus reduce the power consumption of the CPU.

Operation:

```
T:  if SysAD bus is idle then
      StopPipeline
    endif
```

Exceptions:

Coprocessor unusable exception



Format: XOR rd, rs, rt MIPS I

Purpose: To do a bitwise logical EXCLUSIVE OR.

Description: $rd \leftarrow rs \text{ XOR } rt$

Combine the contents of GPR *rs* and GPR *rt* in a bitwise logical exclusive OR operation and place the result into GPR *rd*.

Restrictions:


None

Operation:

$GPR[rd] \leftarrow GPR[rs] \text{ xor } GPR[rt]$

Exceptions:

None



Format: XORI *rt*, *rs*, *immediate* **MIPS I**

Purpose: To do a bitwise logical EXCLUSIVE OR with a constant.

Description: $rt \leftarrow rs \text{ XOR } immediate$

Combine the contents of GPR *rs* and the 16-bit zero-extended *immediate* in a bitwise logical exclusive OR operation and place the result into GPR *rt*.

Restrictions:

None

Operation:

$$GPR[rt] \leftarrow GPR[rs] \text{ xor } \text{zero_extend}(immediate)$$

Exceptions:

None




CPU Instructions Encoding


Notes

A CPU instruction is a single 32-bit aligned word. The major instruction formats are shown in Table 3.1.


I-Type (Immediate).



J-Type (Jump).



R-Type (Register).



opcode	6-bit primary operation code
rd	5-bit destination register specifier
rs	5-bit source register specifier
rt	5-bit target (source/destination) register specifier or used to specify functions within the primary opcode value <i>REGIMM</i>
immediate	16-bit signed immediate used for: logical operands, arithmetic signed operands, load/store address byte offsets, PC-relative branch signed instruction displacement
instr_index	26-bit index shifted left two bits to supply the low-order 28 bits of the jump target address.
sa	5-bit shift amount
function	6-bit function field used to specify functions within the primary operation code value <i>SPECIAL</i> .

Table 3.1 CPU Instruction Formats

CPU Instruction Encoding

This section describes the encoding of user-level, i.e. non-privileged, CPU instructions for the four levels of the MIPS architecture, MIPS I through MIPS IV. Each architecture level includes the instructions in the previous level;¹ MIPS IV includes all instructions in MIPS I, MIPS II, and MIPS III. This section presents eight different views of the instruction encoding.

- ◆ *Separate encoding tables for each architecture level.*
- ◆ *A MIPS IV encoding table showing the architecture level at which each opcode was originally defined and subsequently modified (if modified).*
- ◆ *Separate encoding tables for each architecture revision showing the changes made during that revision.*

Instruction Decode

Instruction field names are printed in **bold** in this section.

The primary **opcode** field is decoded first. Most **opcode** values completely specify an instruction that has an immediate value or offset. **Opcode** values that do not specify an instruction specify an instruction class. Instructions within a class are further specified by values in other fields. The **opcode** values *SPECIAL* and *REGIMM* specify instruction classes. The *COP0*, *COP1*, *COP2*, *COP3*, and *COP1X* instruction classes are not CPU instructions; See “Non-CPU Instructions in the Tables” below.

SPECIAL Instruction Class

The **opcode**=*SPECIAL* instruction class encodes 3-register computational instructions, jump register, and some special purpose instructions. The class is further decoded by examining the **format** field. The **format** values fully specify the CPU instructions; the *MOVCI* instruction class is not a CPU instruction class.

REGIMM Instruction Class

The **opcode**=*REGIMM* instruction class encodes conditional branch and trap immediate instructions. The class is further decoded, and the instructions fully specified, by examining the **rt** field.

Instruction Subsets of MIPS III and MIPS IV Processors

MIPS III processors, such as the RC4000, RC4200, RC4300, RC4400, and RC4600, have a processor mode in which only the MIPS II instructions are valid. The MIPS II encoding table describes the MIPS II-only mode except that the Coprocessor 3 instructions (*COP3*, *LWC3*, *SWC3*, *LDC3*, *SDC3*) are not available and cause a Reserved Instruction exception.

MIPS IV processors, such as the R8000 and R10000, have processor modes in which only the MIPS II or MIPS III instructions are valid. The MIPS II encoding table describes the MIPS II-only mode except that the Coprocessor 3 instructions (*COP3*, *LWC3*, *SWC3*, *LDC3*, *SDC3*) are not available and cause a Reserved Instruction exception. The MIPS III encoding table describes the MIPS III-only mode.

Non-CPU Instructions in the Tables

The encoding tables show all values for the field they describe and by doing this they include some entries that are not user-level CPU instructions. The primary opcode table includes coprocessor instruction classes (*COP0*, *COP1*, *COP2*, *COP3/COP1X*) and coprocessor load/store instructions (*LWCx*, *SWCx*, *LDCx*, *SDCx* for *x*=1, 2, or 3). The **opcode**=*SPECIAL* + **function**=*MOVCI* instruction class is an FPU instruction.

Coprocessor 0 - COP0

COP0 encodes privileged instructions for Coprocessor 0, the System Control Coprocessor. The definition of the System Control Coprocessor is processor-specific and further information on these instructions are not included in this document.

¹ An exception to this rule is that the reserved, but never implemented, Coprocessor 3 instructions were removed or changed to another use starting in MIPS III.

Coprocessor 1 - COP1, COP1X, MOVCI, and CP1 load/store


Coprocessor 1 is the floating-point unit in the MIPS architecture. *COP1*, *COP1X*, and the (*opcode=SPECIAL + function=MOVCI*) instruction classes encode floating-point instructions. LWC1, SWC1, LDC1, and SDC1 are floating-point loads and stores. The FPU instruction encoding is documented in section "FPU" (CP1) Instruction Opcode Bit Encoding in the Chapter "FPU Instruction Set"..

Coprocessor 2 - COP2 and CP2 load/store

Coprocessor 2 is optional and implementation-specific. No standard processor from MIPS has implemented coprocessor 2, but MIPS' semiconductor licensees may have implemented it in a product based on one of the standard MIPS processors. At this time the standard processors are: RC2000, RC3000, RC4000, RC4200, RC4300, RC4400, RC4600, RC6000, R8000, and R10000.

Coprocessor 3 - COP3 and CP3 load/store


Coprocessor 3 is optional and implementation-specific in the MIPS I and MIPS II architecture levels. It was removed from MIPS III and later architecture levels. Note that in MIPS IV the *COP3* primary opcode was reused for the *COP1X* instruction class. No standard processor from MIPS has implemented coprocessor 2, but MIPS' semiconductor licensees may have implemented it in a product based on one of the standard MIPS processors. At this time the standard processors are: RC2000, RC3000, RC4000, RC4200, RC4300, RC4400, RC4600, RC6000, R8000, and R10000.



opcode bits 28..26		Instructions encoded by opcode field.							
bits 31..29	0	1	2	3	4	5	6	7	
	000	001	010	011	100	101	110	111	
0	000	<i>SPECIAL</i> d	<i>REGIMM</i> d	J	JAL	BEQ	BNE	BLEZ	BGTZ
1	001	ADDI	ADDIU	SLTI	SLTIU	ANDI	ORI	XORI	LUI
2	010	<i>COP0</i> d,p	<i>COP1</i> d,p	<i>COP2</i> d,p	<i>COP3</i> d,p,k	*	*	*	*
3	011	*	*	*	*	*	*	*	*
4	100	LB	LH	LWL	LW	LBU	LHU	LWR	*
5	101	SB	SH	SWL	SW	*	*	SWR	*
6	110	*	LWC1 p	LWC2 p	LWC3 p,k	*	*	*	*
7	111	*	SWC1 p	SWC2 p	SWC3 p,k	*	*	*	*




function bits 2..0		Instructions encoded by function field when opcode field = SPECIAL.							
bits 5..3	0	1	2	3	4	5	6	7	
	000	001	010	011	100	101	110	111	
0	000	SLL	*	SRL	SRA	SLLV	*	SRLV	SRAV
1	001	JR	JALR	*	*	SYSCALL	BREAK	*	*
2	010	MFHI	MTHI	MFLO	MTLO	*	*	*	*
3	011	MULT	MULTU	DIV	DIVU	*	*	*	*
4	100	ADD	ADDU	SUB	SUBU	AND	OR	XOR	NOR
5	101	*	*	SLT	SLTU	*	*	*	*
6	110	*	*	*	*	*	*	*	*
7	111	*	*	*	*	*	*	*	*




rt bits 18..16		Instructions encoded by the rt field when opcode field = REGIMM.							
bits 20..19	0	1	2	3	4	5	6	7	
	000	001	010	011	100	101	110	111	
0	00	BLTZ	BGEZ	=	=	=	=	=	=
1	01	=	=	=	=	=	=	=	=
2	10	BLTZAL	BGEZAL	=	=	=	=	=	=
3	11	=	=	=	=	=	=	=	=


Table 3.2 CPU Instruction Encoding - MIPS I Architecture



opcode bits 28..26		Instructions encoded by opcode field.							
bits 31..29	0	1	2	3	4	5	6	7	
	000	001	010	011	100	101	110	111	
0	000	<i>SPECIAL</i> d	<i>REGIMM</i> d	J	JAL	BEQ	BNE	BLEZ	BGTZ
1	001	ADDI	ADDIU	SLTI	SLTIU	ANDI	ORI	XORI	LUI
2	010	<i>COP0</i> d,p	<i>COP1</i> d,p	<i>COP2</i> d,p	<i>COP3</i> d,p,k	BEQL	BNEL	BLEZL	BGTZL
3	011	*	*	*	*	*	*	*	*
4	100	LB	LH	LWL	LW	LBU	LHU	LWR	*
5	101	SB	SH	SWL	SW	*	*	SWR	r
6	110	LL	LWC1 p	LWC2 p	LWC3 p,k	*	LDC1 p	LDC2 p	LDC3 p,k
7	111	SC	SWC1 p	SWC2 p	SWC3 p,k	*	SDC1 p	SDC2 p	SDC3 p,k




function bits 2..0		Instructions encoded by function field when opcode field = <i>SPECIAL</i> .							
bits 5..3	0	1	2	3	4	5	6	7	
	000	001	010	011	100	101	110	111	
0	000	SLL	*	SRL	SRA	SLLV	*	SRLV	SRAV
1	001	JR	JALR	*	*	SYSCALL	BREAK	*	SYNC
2	010	MFHI	MTHI	MFLO	MTLO	*	*	*	*
3	011	MULT	MULTU	DIV	DIVU	*	*	*	*
4	100	ADD	ADDU	SUB	SUBU	AND	OR	XOR	NOR
5	101	*	*	SLT	SLTU	*	*	*	*
6	110	TGE	TGEU	TLT	TLTU	TEQ	*	TNE	*
7	111	*	*	*	*	*	*	*	*



rt bits 18..16		Instructions encoded by the rt field when opcode field = <i>REGIMM</i> .							
bits 20..19	0	1	2	3	4	5	6	7	
	000	001	010	011	100	101	110	111	
0	00	BLTZ	BGEZ	BLTZL	BGEZL	*	*	*	*
1	01	TGEI	TGEIU	TLTI	TLTIU	TEQI	*	TNEI	*
2	10	BLTZAL	BGEZAL	BLTZALL	BGEZALL	*	*	*	*
3	11	*	*	*	*	*	*	*	*


Table 3.3 CPU Instruction Encoding - MIPS II Architecture



opcode		Instructions encoded by opcode field.							
bits 31..29	bits 28..26	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000	<i>SPECIAL</i> d	<i>REGIMM</i> d	J	JAL	BEQ	BNE	BLEZ	BGTZ
1	001	ADDI	ADDIU	SLTI	SLTIU	ANDI	ORI	XORI	LUI
2	010	<i>COPO</i> d,p	<i>COPI</i> d,p	<i>COPI2</i> d,p	*	BEQL	BNEL	BLEZL	BGTZL
3	011	DADDI	DADDIU	LDL	LDR	*	*	*	*
4	100	LB	LH	LWL	LW	LBU	LHU	LWR	LWU
5	101	SB	SH	SWL	SW	SDL	SDR	SWR	r
6	110	LL	LWC1 p	LWC2 p	*	LLD	LDC1 p	LDC2 p	LD
7	111	SC	SWC1 p	SWC2 p	*	SCD	SDC1 p	SDC2 p	SD




function		Instructions encoded by function field when opcode field = SPECIAL.							
bits 5..3	bits 2..0	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000	SLL	*	SRL	SRA	SLLV	*	SRLV	SRAV
1	001	JR	JALR	*	*	SYSCALL	BREAK	*	SYNC
2	010	MFHI	MTHI	MFLO	MTLO	DSLLV	*	DSRLV	DSRAV
3	011	MULT	MULTU	DIV	DIVU	DMULT	DMULTU	DDIV	DDIVU
4	100	ADD	ADDU	SUB	SUBU	AND	OR	XOR	NOR
5	101	*	*	SLT	SLTU	DADD	DADDU	DSUB	DSUBU
6	110	TGE	TGEU	TLT	TLTU	TEQ	*	TNE	*
7	111	DSLL	*	DSRL	DSRA	DSLL32	*	DSRL32	DSRA32



rt		Instructions encoded by the rt field when opcode field = REGIMM.							
bits 20..19	bits 18..16	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	00	BLTZ	BGEZ	BLTZL	BGEZL	*	*	*	*
1	01	TGEI	TGEIU	TLTI	TLTIU	TEQI	*	TNEI	*
2	10	BLTZAL	BGEZAL	BLTZALL	BGEZALL	*	*	*	*
3	11	*	*	*	*	*	*	*	*


Table 3.4 CPU Instruction Encoding - MIPS III Architecture



opcode		Instructions encoded by opcode field.							
bits 31..29	bits 28..26	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000	<i>SPECIAL</i> d	<i>REGIMM</i> d	J	JAL	BEQ	BNE	BLEZ	BGTZ
1	001	ADDI	ADDIU	SLTI	SLTIU	ANDI	ORI	XORI	LUI
2	010	<i>COP0</i> d,p	<i>COP1</i> d,p	<i>COP2</i> d,p	<i>COP1X</i> d,p	BEQL	BNEL	BLEZL	BGTZL
3	011	DADDI	DADDIU	LDL	LDR		*	*	*
4	100	LB	LH	LWL	LW	LBU	LHU	LWR	LWU
5	101	SB	SH	SWL	SW	SDL	SDR	SWR	r
6	110	LL	LWC1 p	LWC2 p	PREF	LLD	LDC1 p	LDC2 p	LD
7	111	SC	SWC1 p	SWC2 p	*	SCD	SDC1 p	SDC2 p	SD



function		Instructions encoded by function field when opcode field = <i>SPECIAL</i> .							
bits 5..3	bits 2..0	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000	SLL	<i>MOVCI</i> d,m	SRL	SRA	SLLV	*	SRLV	SRAV
1	001	JR	JALR	MOVZ	MOVN	SYSCALL	BREAK	*	SYNC
2	010	MFHI	MTHI	MFLO	MTLO	DSLLV	*	DSRLV	DSRAV
3	011	MULT	MULTU	DIV	DIVU	DMULT	DMULTU	DDIV	DDIVU
4	100	ADD	ADDU	SUB	SUBU	AND	OR	XOR	NOR
5	101	*	*	SLT	SLTU	DADD	DADDU	DSUB	DSUBU
6	110	TGE	TGEU	TLT	TLTU	TEQ	*	TNE	*
7	111	DSLL	*	DSRL	DSRA	DSLL32	*	DSRL32	DSRA32



rt		Instructions encoded by the rt field when opcode field = <i>REGIMM</i> .							
bits 20..19	bits 18..16	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	00	BLTZ	BGEZ	BLTZL	BGEZL	*	*	*	*
1	01	TGEI	TGEIU	TLTI	TLTIU	TEQI	*	TNEI	*
2	10	BLTZAL	BGEZAL	BLTZALL	BGEZALL	*	*	*	*
3	11	*	*	*	*	*	*	*	*

Table 3.5 CPU Instruction Encoding - MIPS IV Architecture


The architecture level in which each MIPS IV encoding was defined is indicated by a subscript 1, 2, 3, or 4 (for architecture level I, II, III, or IV). If an instruction or instruction class was later extended, the extending level is indicated after the defining level.



opcode		Instructions encoded by opcode field.							
bits 31..29	bits 28..26	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000	<i>SPECIAL</i> _{1,4}	<i>REGIMM</i> _{1,2}	<i>J</i> ₁	<i>JAL</i> ₁	<i>BEQ</i> ₁	<i>BNE</i> ₁	<i>BLEZ</i> ₁	<i>BGTZ</i> ₁
1	001	<i>ADDI</i> ₁	<i>ADDIU</i> ₁	<i>SLTI</i> ₁	<i>SLTIU</i> ₁	<i>ANDI</i> ₁	<i>ORI</i> ₁	<i>XORI</i> ₁	<i>LUI</i> ₁
2	010	<i>COP0</i> ₁	<i>COP1</i> _{1,2,3,4}	<i>COP2</i> ₁	<i>COPIX</i> ₄	<i>BEQL</i> ₂	<i>BNEL</i> ₂	<i>BLEZL</i> ₂	<i>BGTZL</i> ₂
3	011	<i>DADDI</i> ₃	<i>DADDIU</i> ₃	<i>LDL</i> ₃	<i>LDR</i> ₃	* ₁	* ₁	* ₁	* ₁
4	100	<i>LB</i> ₁	<i>LH</i> ₁	<i>LWL</i> ₁	<i>LW</i> ₁	<i>LBU</i> ₁	<i>LHU</i> ₁	<i>LWR</i> ₁	<i>LWU</i> ₃
5	101	<i>SB</i> ₁	<i>SH</i> ₁	<i>SWL</i> ₁	<i>SW</i> ₁	<i>SDL</i> ₃	<i>SDR</i> ₃	<i>SWR</i> ₁	<i>r</i> ₂
6	110	<i>LL</i> ₂	<i>LWC1</i> ₁	<i>LWC2</i> ₁	<i>PREF</i> ₄	<i>LLD</i> ₃	<i>LDC1</i> ₂	<i>LDC2</i> ₂	<i>LD</i> ₃
7	111	<i>SC</i> ₂	<i>SWC1</i> ₁	<i>SWC2</i> ₁	* ₃	<i>SCD</i> ₃	<i>SDC1</i> ₂	<i>SDC2</i> ₂	<i>SD</i> ₃




function		Instructions encoded by function field when opcode field = SPECIAL.							
bits 5..3	bits 2..0	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000	<i>SLL</i> ₁	<i>MOVCI</i> ₄	<i>SRL</i> ₁	<i>SRA</i> ₁	<i>SLLV</i> ₁	* ₁	<i>SRLV</i> ₁	<i>SRV</i> ₁
1	001	<i>JR</i> ₁	<i>JALR</i> ₁	<i>MOVZ</i> ₄	<i>MOVN</i> ₄	<i>SYSCALL</i> ₁	<i>BREAK</i> ₁	* ₁	<i>SYNC</i> ₂
2	010	<i>MFHI</i> ₁	<i>MTHI</i> ₁	<i>MFLO</i> ₁	<i>MTLO</i> ₁	<i>DSLIV</i> ₃	* ₁	<i>DSRLV</i> ₃	<i>DSRAV</i> ₃
3	011	<i>MULT</i> ₁	<i>MULTU</i> ₁	<i>DIV</i> ₁	<i>DIVU</i> ₁	<i>DMULT</i> ₃	<i>DMULTU</i> ₃	<i>DDIV</i> ₃	<i>DDIVU</i> ₃
4	100	<i>ADD</i> ₁	<i>ADDU</i> ₁	<i>SUB</i> ₁	<i>SUBU</i> ₁	<i>AND</i> ₁	<i>OR</i> ₁	<i>XOR</i> ₁	<i>NOR</i> ₁
5	101	* ₁	* ₁	<i>SLT</i> ₁	<i>SLTU</i> ₁	<i>DADD</i> ₃	<i>DADDU</i> ₃	<i>DSUB</i> ₃	<i>DSUBU</i> ₃
6	110	<i>TGE</i> ₂	<i>TGEU</i> ₂	<i>TLT</i> ₂	<i>TLTU</i> ₂	<i>TEQ</i> ₂	* ₁	<i>TNE</i> ₂	* ₁
7	111	<i>DSLL</i> ₃	* ₁	<i>DSRL</i> ₃	<i>DSRA</i> ₃	<i>DSLL32</i> ₃	* ₁	<i>DSRL32</i> ₃	<i>DSRA32</i> ₃



rt		Instructions encoded by the rt field when opcode field = REGIMM.							
bits 20..19	bits 18..16	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	00	<i>BLTZ</i> ₁	<i>BGEZ</i> ₁	<i>BLTZL</i> ₂	<i>BGEZL</i> ₂	* ₁	* ₁	* ₁	* ₁
1	01	<i>TGEI</i> ₂	<i>TGEIU</i> ₂	<i>TLTI</i> ₂	<i>TLTIU</i> ₂	<i>TEQI</i> ₂	* ₁	<i>TNEI</i> ₂	* ₁
2	10	<i>BLTZAL</i> ₁	<i>BGEZAL</i> ₁	<i>BLTZALL</i> ₂	<i>BGEZALL</i> ₂	* ₁	* ₁	* ₁	* ₁
3	11	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁

Table 3.6 Architecture Level in Which CPU Instructions are Defined or Extended




An instruction encoding is shown if the instruction is added in this revision.

opcode		Instructions encoded by opcode field.							
bits 31..29	bits 28..26	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000								
1	001								
2	010					BEQL	BNEL	BLEZL	BGTZL
3	011								
4	100								
5	101								r
6	110	LL					LDC1 _p	LDC2 _p	LDC3 _p
7	111	SC					SDC1 _p	SDC2 _p	SDC3 _p




function		Instructions encoded by function field when opcode field = SPECIAL.							
bits 5..3	bits 2..0	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000								
1	001								SYNC
2	010								
3	011								
4	100								
5	101								
6	110	TGE	TGEU	TLT	TLTU	TEQ		TNE	
7	111								



rt		Instructions encoded by the rt field when opcode field = REGIMM.							
bits 20..19	bits 18..16	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	00			BLTZL	BGEZL				
1	01	TGEI	TGEIU	TLTI	TLTIU	TEQI		TNEI	
2	10			BLTZALL	BGEZALL				
3	11								

Table 3.7 CPU Instruction Encoding Changes - MIPS II Revision




An instruction encoding is shown if the instruction is added or modified in this revision.

opcode		Instructions encoded by opcode field.							
bits 31..29	bits 28..26	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000								
1	001								
2	010				*				
					(was COP3)				
3	011	DADDI	DADDIU	LDL	LDR				
4	100								LWU
5	101					SDL	SDR		
6	110				*	LLD			LD
					(was LWC3)				(was LDC3)
7	111				*	SCD			SD
					(was SWC3)				(was SDC3)




function		Instructions encoded by function field when opcode field = SPECIAL.							
bits 5..3	bits 2..0	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000								
1	001								
2	010					DSLLV		DSRLV	DSRAV
3	011					DMULT	DMULTU	DDIV	DDIVU
4	100								
5	101					DADD	DADDU	DSUB	DSUBU
6	110								
7	111	DSLL		DSRL	DSRA	DSLL32		DSRL32	DSRA32




rt		Instructions encoded by the rt field when opcode field = REGIMM.							
bits 20..19	bits 18..16	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	00								
1	01								
2	10								
3	11								

Table 3.8 CPU Instruction Encoding Changes - MIPS III Revision




An instruction encoding is shown if the instruction is added or modified in this revision.

opcode		Instructions encoded by opcode field.							
bits 31..29	bits 28..26	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000								
1	001								
2	010				COPIX d,p				
3	011								
4	100								
5	101								
6	110				PREF				
7	111								



function		Instructions encoded by function field when opcode field = SPECIAL.							
bits 5..3	bits 2..0	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000		MOVCl d,m						
1	001			MOVZ	MOVN				
2	010								
3	011								
4	100								
5	101								
6	110								
7	111								



rt		Instructions encoded by the rt field when opcode field = REGIMM.							
bits 20..19	bits 18..16	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	00								
1	01								
2	10								
3	11								

Table 3.9 CPU Instruction Encoding Changes - MIPS IV Revision

Key to notes in CPU instruction encoding tables:

- * This opcode is reserved for future use. An attempt to execute it causes a Reserved Instruction exception.
- = This opcode is reserved for future use. An attempt to execute it produces an undefined result. The result may be a Reserved Instruction exception but this is not guaranteed.
- δ (also *italic* opcode name) This opcode indicates an instruction class. The instruction word must be further decoded by examining additional tables that show values for another instruction field.
- π 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.
- κ This opcode is removed in a later revision of the architecture. If a MIPS III or MIPS IV processor is operated in MIPS II-only mode this opcode will cause a Reserved Instruction exception.
- μ This opcode indicates a class of coprocessor 1 instructions. If the processor state does not allow access to coprocessor 1, the opcode causes a Coprocessor Unusable exception. It is included in the table because the encoding uses a location in what is otherwise a CPU instruction encoding map. Further encoding information for this instruction class is in the FPU Instruction Encoding tables.
- ρ This opcode is reserved for Coprocessor 0 (System Control Coprocessor) instructions that require base+offset addressing. If the instruction is used for COP0 in an implementation, an attempt to execute it without Coprocessor 0 access privilege will cause a Coprocessor Unusable exception. If the instruction is not used in an implementation, it will cause a Reserved Instruction exception.



FPU Instructions Basics

Notes

FPU Instruction Set Details

This appendix documents the instructions for the floating-point unit (FPU) in MIPS processors. It contains some descriptive material at the beginning, a detailed description for each instruction in alphabetic order, and an instruction opcode encoding table at the end of the section.

The descriptive material describes the FPU instruction categories, the instruction encoding formats, the valid operands for FPU computational instructions, compare and condition values, FPU use of the coprocessor registers, and a description of the notation used for the detailed instruction description.

This section does not describe the operation of floating-point arithmetic, the exception conditions within FP arithmetic, the exception mechanism of the FPU, or the handling of these FP exceptions.

FPU Instructions

The floating-point unit (FPU) is implemented as Coprocessor unit 1 (CP1) within the MIPS architecture. A floating-point instruction needs access to coprocessor 1 to execute; if CP1 is not enabled, an FP instruction will cause a Coprocessor Unusable exception. The FPU has a load/store architecture. All computations are done on data held in registers, and data is transferred between registers and the rest of the system with dedicated load, store, and move instructions.

The FPU instructions fall into the following categories:

- ◆ *Data Transfer*
- ◆ *Arithmetic*
- ◆ *Conversion*
- ◆ *Formatted Operand Value Move*
- ◆ *Conditional Branch*
- ◆ *Miscellaneous*

Data Transfer Instructions

The FPU has two separate register sets: coprocessor general registers and coprocessor control registers. The FPU has a load/store architecture; all computations are done on data held in coprocessor general registers. The control registers are used to control FPU operation. Data is transferred between registers and the rest of the system with dedicated load, store, and move instructions. The transferred data is treated as unformatted binary data; no format conversions are performed and, therefore, no IEEE floating-point exceptions can occur.

The supported transfer operations are:

- FPU general reg ↔ memory (word/doubleword load/store)
- FPU general reg ↔ CPU general reg (word/doubleword move)
- FPU control reg ↔ CPU general reg (word move)

All coprocessor loads and stores operate on naturally-aligned data items. An attempt to load or store to an address that is not naturally aligned for the data item will cause an Address Error exception. Regardless of byte-numbering order (endianness), the address of a word or doubleword is the smallest byte address among the bytes in the object. For a big-endian machine this is the most-significant byte; for a little-endian machine this is the least-significant byte.

The FPU has loads and stores using the usual register+offset addressing. In MIPS IV, for the FPU only, there are also load and store instructions using register+register addressing.

MIPS I specifies that loads are delayed by one instruction and that proper execution must be insured by observing an instruction scheduling restriction. The instruction immediately following a load into an FPU register *F_n* must not use *F_n* as a source register. The time between the load instruction and the time the data is available is the “load delay slot”. If no useful instruction can be put into the load delay slot, then a null operation (NOP) must be inserted.

In MIPS II, this instruction scheduling restriction is removed. Programs will execute correctly when the loaded data is used by the instruction following the load, but this may require extra real cycles. Most processors cannot actually load data quickly enough for immediate use and the processor will be forced to wait until the data is available. Scheduling load delay slots is desirable for performance reasons even when it is not necessary for correctness.

Mnemonic	Description	Defined in
LWC1	Load Word to Floating-Point	I
SWC1	Store Word to Floating-Point	I
LDC1	Load Doubleword to Floating-Point	III
SDC1	Store Doubleword to Floating-Point	III

Table 4.10 FPU Loads and Stores Using Register + Offset Address Mode

Mnemonic	Description	Defined in
LWXC1	Load Word Indexed to Floating-Point	IV
SWXC1	Store Word Indexed to Floating-Point	IV
LDXC1	Load Doubleword Indexed to Floating-Point	IV
SDXC1	Store Doubleword Indexed to Floating-Point	IV

Table 4.11 FPU Loads and Stores Using Register + Register Address Mode

Mnemonic	Description	Defined in
MTC1	Move Word To Floating-Point	I
MFC1	Move Word From Floating-Point	I
DMTC1	Doubleword Move To Floating-Point	III
DMFC1	Doubleword Move From Floating-Point	III
CTC1	Move Control Word To Floating-Point	I
CFC1	Move Control Word From Floating-Point	I

Table 4.12 FPU Move To/From Instructions

Arithmetic Instructions

The arithmetic instructions operate on formatted data values. The result of most floating-point arithmetic operations meets the IEEE standard specification for accuracy; a result which is identical to an infinite-precision result rounded to the specified format, using the current rounding mode. The rounded result differs from the exact result by less than one unit in the least-significant place (ulp).

Mnemonic	Description	Defined in
ADD. <i>fmt</i>	Floating-Point Add	I
SUB. <i>fmt</i>	Floating-Point Subtract	I
MUL. <i>fmt</i>	Floating-Point Multiply	I
DIV. <i>fmt</i>	Floating-Point Divide	I
ABS. <i>fmt</i>	Floating-Point Absolute Value	I
NEG. <i>fmt</i>	Floating-Point Negate	I
SQRT. <i>fmt</i>	Floating-Point Square Root	II
C. <i>cond.fmt</i>	Floating-Point Compare	I

Table 4.13 FPU IEEE Arithmetic Operations

Two operations, Reciprocal Approximation (RECIP) and Reciprocal Square Root Approximation (RSQRT), may be less accurate than the IEEE specification. The result of RECIP differs from the exact reciprocal by no more than one ulp. The result of RSQRT differs by no more than two ulp. Within these error limits, the result of these instructions is implementation specific.

Mnemonic	Description	Defined in
RECIP. <i>fmt</i>	Floating-Point Reciprocal Approximation	IV
RSQRT. <i>fmt</i>	Floating-Point Reciprocal Square Root Approximation	IV

Table 4.14 FPU Approximate Arithmetic Operations

There are four compound-operation instructions that perform variations of multiply-accumulate: multiply two operands and accumulate to a third operand to produce a result. The accuracy of the result depends which of two alternative arithmetic models is used for the computation. The unrounded model is more accurate than a pair of IEEE operations and the rounded model meets the IEEE specification.

Mnemonic	Description	Defined in
MADD. <i>fmt</i>	Floating-Point Multiply Add	IV
MSUB. <i>fmt</i>	Floating-Point Multiply Subtract	IV
NMADD. <i>fmt</i>	Floating-Point Negative Multiply Add	IV
NMSUB. <i>fmt</i>	Floating-Point Negative Multiply Subtract	IV

Table 4.15 FPU Multiply-Accumulate Arithmetic Operations

The RC5000 uses the rounded model which meets the specification.

- ◆ *Rounded or non-fused*
 - *The product is rounded according to the current rounding mode prior to the accumulation. This model meets the IEEE accuracy specification; the result is numerically identical to the equivalent computation using multiply, add, subtract, and negate instructions.*
- ◆ *Unrounded or fused (R8000 implementation)*
 - *The product is not rounded and all bits take part in the accumulation. This model does not match the IEEE accuracy requirements; the result is more accurate than the equivalent computation using IEEE multiply, add, subtract, and negate instructions.*

Conversion Instructions

There are instructions to perform conversions among the floating-point and fixed-point data types. Each instruction converts values from a number of operand formats to a particular result format. Some convert instructions use the rounding mode specified in the Floating Control and Status Register (FCSR), others specify the rounding mode directly.

Mnemonic	Description	Defined in
CVT.S.fmt	Floating-Point Convert to Single Floating-Point	I
CVT.D.fmt	Floating-Point Convert to Double Floating-Point	I
CVT.W.fmt	Floating-Point Convert to Word Fixed-Point	I
CVT.L.fmt	Floating-Point Convert to Long Fixed-Point	I

Table 4.1 FPU Conversion Operations Using the FCSR Rounding Mode

Mnemonic	Description	Defined in
ROUND.W.fmt	Floating-Point Round to Word Fixed-Point	II
ROUND.L.fmt	Floating-Point Round to Long Fixed-Point	III
TRUNC.W.fmt	Floating-Point Truncate to Word Fixed-Point	II
TRUNC.L.fmt	Floating-Point Truncate to Long Fixed-Point	III
CEIL.W.fmt	Floating-Point Ceiling to Word Fixed-Point	II
CEIL.L.fmt	Floating-Point Ceiling to Long Fixed-Point	III
FLOOR.W.fmt	Floating-Point Floor to Word Fixed-Point	II
FLOOR.L.fmt	Floating-Point Floor to Long Fixed-Point	III

Table 4.16 FPU Conversion Operations Using a Directed Rounding Mode

Formatted Operand Value Move Instructions

These instructions all move formatted operand values among FPU general registers. A particular operand type must be moved by the instruction that handles that type. There are three kinds of move instructions:

- ◆ *Unconditional move*
- ◆ *Conditional move that tests an FPU condition code*
- ◆ *Conditional move that tests a CPU general register value against zero*

The conditional move instructions operate in a way that may be unexpected. They always force the value in the destination register to become a value of the format specified in the instruction. If the destination register does not contain an operand of the specified format, before the conditional move is executed, the contents become undefined.

Mnemonic	Description	Defined in
MOV.fmt	Floating-Point Move	I

Table 4.17 FPU Formatted Operand Move Instructions

Mnemonic	Description	Defined in
MOVT.fmt	Floating-Point Move Conditional on FP True	IV
MOVF.fmt	Floating-Point Move Conditional on FP False	IV

Table 4.18 FPU Conditional Move on True/False Instructions

Mnemonic	Description	Defined in
MOVZ.fmt	Floating-Point Move Conditional on Zero	IV
MOVN.fmt	Floating-Point Move Conditional on Nonzero	IV

Table 4.19 FPU Conditional Move on Zero/Nonzero Instructions

Conditional Branch Instructions

The FPU has PC-relative conditional branch instructions that test condition codes set by FPU compare instructions (*C.cond.fmt*).

All branches have an architectural delay of one instruction. When a branch is taken, the instruction immediately following the branch instruction, in the branch delay slot, is executed before the branch to the target instruction takes place. Conditional branches come in two versions that treat the instruction in the delay slot differently when the branch is not taken and execution falls through. The “branch” instructions execute the instruction in the delay slot, but the “branch likely” instructions do not (they are said to nullify it).

MIPS I defines a single condition code which is implicit in the compare and branch instructions. MIPS IV defines seven additional condition codes and includes the condition code number in the compare and branch instructions. The MIPS IV extension keeps the original condition bit as condition code zero and the extended encoding is compatible with the MIPS I encoding.

Mnemonic	Description	Defined in
BC1T	Branch on FP True	I
BC1F	Branch on FP False	I
BC1TL	Branch on FP True Likely	II
BC1FL	Branch on FP False Likely	II

Table 4.20 FPU Conditional Branch Instructions

Miscellaneous Instructions

CPU Conditional Move

There are instructions to conditionally move one CPU general register to another based on an FPU condition code as shown in Table 4.21.

Mnemonic	Description	Defined in
MOVZ	Move Conditional on FP True	IV
MOVN	Move Conditional on FP False	IV

Table 4.21 CPU Conditional Move on FPU True/False Instructions

Valid Operands for FP Instructions

The floating-point unit arithmetic, conversion, and operand move instructions operate on formatted values with different precision and range limits and produce formatted values for results. Each representable value in each format has a binary encoding that is read from or stored to memory. The *fmt* or *fmt3* field of the instruction encodes the operand format required for the instruction. A conversion instruction specifies the result type in the *function* field; the result of other operations is the same format as the operands. The encoding of the *fmt* and *fmt3* fields is shown in Table 4.22.

fmt	fmt3	Instruction Mnemonic	Size		data type
			name	bits	
0-15	-	Reserved			
16	0	S	single	32	floating-point
17	1	D	double	64	floating-point
18-19	2-3	Reserved			
20	4	W	word	32	fixed-point
21	5	L	long	64	fixed-point
22-31	6-7	Reserved			

Table 4.22 FPU Operand Format Field (fmt, fmt3) Decoding

Each type of arithmetic or conversion instruction is valid for operands of selected formats. A summary of the computational and operand move instructions and the formats valid for each of them is listed in Table 4.23. Implementations must support combinations that are valid either directly in hardware or through emulation in an exception handler.

The result of an instruction using operand formats marked “U” is not currently specified by this architecture and will cause an exception. They are available for future extensions to the architecture. The exact exception mechanism used is processor specific. Most implementations report this as an Unimplemented Operation for a Floating Point exception. Other implementations report these combinations as Reserved Instruction exceptions.

Mnemonic	Operation	operand fmt			
		float		fixed	
		S	D	W	L
ABS	Absolute value	2	2	U	U
ADD	Add	2	2	U	U
<i>C.cond</i>	Floating-point compare	2	2	U	U
CEIL.L	Convert to word/longword fixed-point, round toward $+\infty$	2	2	U	U
CEIL.W					
CVT.D	Convert to double floating-point	2	U	2	2
CVT.L	Convert to longword fixed-point	2	2	U	U
CVT.S	Convert to single floating-point	U	2	2	2
CVT.W	Convert to 32-bit fixed-point	2	2	U	U
DIV	Divide	2	2	U	U
FLOOR.L	Convert to word/longword fixed-point, round toward $-\infty$	2	2	U	U
FLOOR.W					
MOV	Move Register	2	2	U	U
MOVF	FP Move Conditional on condition	2	2	U	U
MOVT					
MOVN	FP Move Conditional on GPR \neq zero	2	2	U	U
MOVZ	FP Move Conditional on GPR = zero	2	2	U	U
NEG	Negate	2	2	U	U
RECIP	Reciprocal approximation	2	2	U	U
ROUND.L	Convert to word/longword fixed-point, round to nearest/even	2	2	U	U
ROUND.W					
RSQRT	Reciprocal square root approximation	2	2	U	U
SQRT	Square root	2	2	U	U
SUB	Subtract	2	2	U	U
TRUNC.L	Convert to word/longword fixed-point, round toward zero	2	2	U	U
TRUNC.W					
Key:	• - Valid. U - Causes unimplemented exception traps.				

Table 4.23 Valid Formats for FPU Operations

Description of an Instruction

For the FPU instruction detail documentation, all variable subfields in an instruction format (such as *fs*, *ft*, *immediate*, and so on) are shown in lower-case. The instruction name (such as ADD, SUB, and so on) is shown in upper-case.

For 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.

In some instructions, the instruction subfields *op* and *function* can have constant 6-bit values. When reference is made to these instructions, upper-case mnemonics are used. For instance, in the floating-point ADD instruction we use *op* = COP1 and *function* = ADD. In other cases, a single field has both fixed and variable subfields, so the name contains both upper and lower case characters. Bit encodings for mnemonics are shown at the end of this section, and are also included with each individual instruction.

Operation Notation Conventions and Functions

The instruction description includes an *Operation* section that describes the operation of the instruction in a pseudocode. The pseudocode and terms used in the description are described in “Operation Section Notation and Functions” on page 15 of Chapter 1.

Individual FPU Instruction Descriptions


The FP instructions are described in alphabetic order. For a description of the information in each instruction, see “Instruction Descriptions” on page 13 of Chapter 1.



FPU Instructions Reference

Notes

This chapter is similar to Chapter 2 except that it deals with the FPU (hardware floating point unit) instructions.



Format: ABS.S *fd, fs* **MIPS I**
 ABS.D *fd, fs*

Purpose: To compute the absolute value of an FP value.

Description: $fd \leftarrow \text{absolute}(fs)$

The absolute value of the value in FPR *fs* is placed in FPR *fd*. The operand and result are values in format *fmt*.

This operation is arithmetic; a NaN operand signals invalid operation.

Restrictions:

The fields *fs* and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.

The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR(*fd, fmt, AbsoluteValue(ValueFPR(*fs, fmt*)))*

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Unimplemented Operation
 - Invalid Operation

Operation:

MIPS I, II, and III define a single condition code; MIPS IV adds 7 more condition codes. This operation specification is for the general "Branch On Condition" operation with the *tf* (true/false) and *nd* (nullify delay slot) fields as variables. The individual instructions BC1F, BC1FL, BC1T, and BC1TL have specific values for *tf* and *nd*.

MIPS I

```

I-1: condition ← COC[1] = tf
I: target_offset ← (offset15)GPRLEN-(16+2) || offset || 02
I+1: if condition then
    PC ← PC + target
endif
    
```

MIPS II and MIPS III:

```

I-1: condition ← COC[1] = tf
I: target_offset ← (offset15)GPRLEN-(16+2) || offset || 02
I+1: if condition then
    PC ← PC + target
else if nd then
    NullifyCurrentInstruction()
endif
    
```

MIPS IV:

```


I: condition ← FCC[cc] = tf
    target_offset ← (offset15)GPRLEN-(16+2) || offset || 02
I+1: if condition then
    PC ← PC + target
else if nd then
    NullifyCurrentInstruction()
endif
    
```

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
- Unimplemented Operation

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.



Format: BC1FL offset (cc = 0 implied) **MIPS II**
 BC1FL cc, offset **MIPS IV**

Purpose: To test an FP condition code and do a PC-relative conditional branch; execute the delay slot only if the branch is taken.

Description: if (cc = 0) then branch_likely

An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.

If the FP condition code bit *cc* is false (0), branch to the effective target address after the instruction in the delay slot is executed. If the branch is not taken, the instruction in the delay slot is not executed.

An FP condition code is set by the FP compare instruction, *C.cond.fmt*.

The MIPS I architecture defines a single floating-point condition code, implemented as the coprocessor 1 condition signal (Cp1Cond) and the C bit in the FP *Control and Status* register. MIPS I, II, and III architectures must have the *cc* field set to 0, which is implied by the first format in the *Format* section.

The MIPS IV architecture adds seven more condition code bits to the original condition code 0. FP compare and conditional branch instructions specify the condition code bit to set or test. Both assembler formats are valid for MIPS IV.

Restrictions:

MIPS II, III: There must be at least one instruction between the compare instruction that sets a condition code and the branch instruction that tests it. Hardware does not detect a violation of this restriction.

MIPS IV: None.

Operation:

MIPS II, and III define a single condition code; MIPS IV adds 7 more condition codes. This operation specification is for the general "Branch On Condition" operation with the *tf* (true/false) and *nd* (nullify delay slot) fields as variables. The individual instructions BC1F, BC1FL, BC1T, and BC1TL have specific values for *tf* and *nd*.

MIPS II and MIPS III:

```

I-1: condition ← COC[1] = tf
I: target_offset ← (offset15)GPRLEN-(16+2) || offset || 02
I+1: if condition then
    PC ← PC + target
else if nd then
    NullifyCurrentInstruction()
endif
    
```

MIPS IV:

```


I: condition ← FCC[cc] = tf
    target_offset ← (offset15)GPRLEN-(16+2) || offset || 02
I+1: if condition then
    PC ← PC + target
else if nd then
    NullifyCurrentInstruction()
endif
    
```

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point Unimplemented Operation

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.



Format: BC1T offset (cc = 0 implied) **MIPS I**
 BC1T cc, offset **MIPS IV**

Purpose: To test an FP condition code and do a PC-relative conditional branch.

Description: if (cc = 1) then branch

An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.

If the FP condition code bit *cc* is true (1), branch to the effective target address after the instruction in the delay slot is executed

An FP condition code is set by the FP compare instruction, *C.cond.fmt*.

The MIPS I architecture defines a single floating-point condition code, implemented as the coprocessor 1 condition signal (Cp1Cond) and the C bit in the FP *Control and Status* register. MIPS I, II, and III architectures must have the *cc* field set to 0, which is implied by the first format in the *Format* section.

The MIPS IV architecture adds seven more condition code bits to the original condition code 0. FP compare and conditional branch instructions specify the condition code bit to set or test. Both assembler formats are valid for MIPS IV.

Restrictions:

MIPS I, II, III: There must be at least one instruction between the compare instruction that sets a condition code and the branch instruction that tests it. Hardware does not detect a violation of this restriction.

MIPS IV: None

Operation:

MIPS I, II, and III define a single condition code; MIPS IV adds 7 more condition codes. This operation specification is for the general "Branch On Condition" operation with the *tf* (true/false) and *nd* (nullify delay slot) fields as variables. The individual instructions BC1F, BC1FL, BC1T, and BC1TL have specific values for *tf* and *nd*.

MIPS I

```

I-1: condition ← COC[1] = tf
I: target ← (offset15)GPREN-(16+2) || offset || 02
I+1: if condition then
    PC ← PC + target
endif
    
```

MIPS II and MIPS III:

```

I-1: condition ← COC[1] = tf
I: target ← (offset15)GPREN-(16+2) || offset || 02
I+1: if condition then
    PC ← PC + target
else if nd then
    NullifyCurrentInstruction()
endif
    
```

MIPS IV:

```


I: condition ← FCC[cc] = tf
    target ← (offset15)GPREN-(16+2) || offset || 02
I+1: if condition then
    PC ← PC + target
else if nd then
    NullifyCurrentInstruction()
endif
    
```

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
- Unimplemented Operation

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.



Format: BC1TL offset (cc = 0 implied) **MIPS II**
 BC1TL cc, offset **MIPS IV**

Purpose: To test an FP condition code and do a PC-relative conditional branch; execute the delay slot only if the branch is taken.

Description: if (cc = 1) then branch_likely

An 18-bit signed offset (the 16-bit *offset* field shifted left 2 bits) is added to the address of the instruction following the branch (**not** the branch itself), in the branch delay slot, to form a PC-relative effective target address.

If the FP condition code bit *cc* is true (1), branch to the effective target address after the instruction in the delay slot is executed. If the branch is not taken, the instruction in the delay slot is not executed.

An FP condition code is set by the FP compare instruction, *C.cond.fmt*.

The MIPS I architecture defines a single floating-point condition code, implemented as the coprocessor 1 condition signal (Cp1Cond) and the C bit in the FP *Control and Status* register. MIPS I, II, and III architectures must have the *cc* field set to 0, which is implied by the first format in the *Format* section.

The MIPS IV architecture adds seven more condition code bits to the original condition code 0. FP compare and conditional branch instructions specify the condition code bit to set or test. Both assembler formats are valid for MIPS IV.

Restrictions:

MIPS II, III: There must be at least one instruction between the compare instruction that sets a condition code and the branch instruction that tests it. Hardware does not detect a violation of this restriction.

MIPS IV: None.

Operation:

MIPS II, and III define a single condition code; MIPS IV adds 7 more condition codes. This operation specification is for the general "Branch On Condition" operation with the *tf* (true/false) and *nd* (nullify delay slot) fields as variables. The individual instructions BC1F, BC1FL, BC1T, and BC1TL have specific values for *tf* and *nd*.

MIPS II and MIPS III:

```

I-1: condition ← COC[1] = tf
I: target ← (offset15)GPRLEN-(16+2) || offset || 02
I+1: if condition then
    PC ← PC + target
    else if nd then
        NullifyCurrentInstruction()
    endif
    
```

MIPS IV:

```

I: condition ← FCC[cc] = tf
    target ← (offset15)GPRLEN-(16+2) || offset || 02
I+1: if condition then
    PC ← PC + target
    else if nd then
        NullifyCurrentInstruction()
    endif
    
```

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
- Unimplemented Operation

Programming Notes:

With the 18-bit signed instruction offset, the conditional branch range is ± 128 KBytes. Use jump (J) or jump register (JR) instructions to branch to more distant addresses.

31	26	25	21	20	16	15	11	10	8	7	6	5	4	3	0	
COP1 0 1 0 0 0 1						fmt	ft	fs	cc	0 0 0	FC 1 1	cond				
						6	5	5	5	3	2	2				

Format: C.cond.S fs, ft (cc = 0 implied) MIPS I
 C.cond.D fs, ft (cc = 0 implied)
 C.cond.S cc, fs, ft MIPS IV
 C.cond.D cc, fs, ft

Purpose: To compare FP values and record the Boolean result in a condition code.

Description: $cc \leftarrow fs \text{ compare_cond } ft$

The value in FPR *fs* is compared to the value in FPR *ft*; the values are in format *fmt*. The comparison is exact and neither overflows nor underflows. If the comparison specified by *cond*_{2,1} is true for the operand values, then the result is true, otherwise it is false. If no exception is taken, the result is written into condition code *cc*; true is 1 and false is 0.

If *cond*₃ is set and at least one of the values is a NaN, an Invalid Operation condition is raised; the result depends on the FP exception model currently active.

- ◆ *Precise exception model:* The Invalid Operation flag is set in the FCSR. If the Invalid Operation enable bit is set in the FCSR, no result is written and an Invalid Operation exception is taken immediately. Otherwise, the Boolean result is written into condition code *cc*.
- ◆ *Imprecise exception model (R8000 normal mode):* The Boolean result is written into condition code *cc*. No FCSR flag is set. If the Invalid Operation enable bit is set in the FCSR, an Invalid Operation exception is taken, imprecisely, at some future time.

There are four mutually exclusive ordering relations for comparing floating-point values; one relation is always true and the others are false. The familiar relations are *greater than*, *less than*, and *equal*. In addition, the IEEE floating-point standard defines the relation *unordered* which is true when at least one operand value is NaN; NaN compares unordered with everything, including itself. Comparisons ignore the sign of zero, so +0 equals -0.

The comparison condition is a logical predicate, or equation, of the ordering relations such as "less than or equal", "equal", "not less than", or "unordered or equal". Compare distinguishes sixteen comparison predicates. The Boolean result of the instruction is obtained by substituting the Boolean value of each ordering relation for the two FP values into equation. If the *equal* relation is true, for example, then all four example predicates above would yield a true result. If the *unordered* relation is true then only the final predicate, "unordered or equal" would yield a true result.

Logical negation of a compare result allows eight distinct comparisons to test for sixteen predicates as shown in Table 5.24. Each mnemonic tests for both a predicate and its logical negation. For each mnemonic, compare tests the truth of the first predicate. When the first predicate is true, the result is true as shown in the "if predicate is true" column (note that the False predicate is never true and False/True do not follow the normal pattern). When the first predicate is true, the second predicate must be false, and vice versa. The truth of the second predicate is the logical negation of the instruction result. After a compare instruction, test for the truth of the first predicate with the Branch on FP True (BC1T) instruction and the truth of the second with Branch on FP False (BC1F).

Instr	Comparison Predicate				Comparison CC Result		Instr		
cond Mnemonic	name of predicate and logically negated predicate (abbreviation)	relation values				If predicate is true	Inv Op excp if Q NaN	cond field	
		>	<	=	?			3	2.0
F	False [this predicate is always False, True (T) it never has a True result]	F	F	F	F	F	No	0	0
UN	Unordered Ordered (OR)	F	F	F	T	T			1
EQ	Equal Not Equal (NEQ)	F	F	T	F	T			2
UEQ	Unordered or Equal Ordered or Greater than or Less than (OGL)	F	F	T	T	T			3
OLT	Ordered or Less Than Unordered or Greater than or Equal (UGE)	F	T	F	F	T			4
ULT	Unordered or Less Than Ordered or Greater than or Equal (OGE)	F	T	F	T	T			5
OLE	Ordered or Less than or Equal Unordered or Greater Than (UGT)	F	T	T	F	T			6
ULE	Unordered or Less than or Equal Ordered or Greater Than (OGT)	F	T	T	T	T			7

key: "?" = *unordered*, ">" = *greater than*, "<" = *less than*, "=" is *equal*, "T" = True, "F" = False

Table 5.24 FPU Comparisons Without Special Operand Exceptions

There is another set of eight compare operations, distinguished by a *cond*₃ value of 1, testing the same sixteen conditions. For these additional comparisons, if at least one of the operands is a NaN, including Quiet NaN, then an Invalid Operation condition is raised. If the Invalid Operation condition is enabled in the FCSR, then an Invalid Operation exception occurs.

Instr	Comparison Predicate name of predicate and logically negated predicate (abbreviation)	relation values				Comparison CC Result		Instr										
		>	<	=	?	If predicate is true	Inv Op excp if Q NaN	cond field										
								3	2..0									
SF	Signaling False [this predicate always False]	F	F	F	F	F	Yes	1	0									
	Signaling True (ST)	T	T	T	T	T												
NGLE	Not Greater than or Less than or Equal	F	F	F	T	T				Yes	1	0						
	Greater than or Less than or Equal (GLE)	T	T	T	F	F												
SEQ	Signaling Equal	F	F	T	F	T							Yes	1	0			
	Signaling Not Equal (SNE)	T	T	F	T	F												
NGL	Not Greater than or Less than	F	F	T	T	T										Yes	1	0
	Greater than or Less than (GL)	T	T	F	F	F												
LT	Less than	F	T	F	F	T	Yes	1	0									
	Not Less Than (NLT)	T	F	T	T	F												
NGE	Not Greater than or Equal	F	T	F	T	T				Yes	1	0						
	Greater than or Equal (GE)	T	F	T	F	F												
LE	Less than or Equal	F	T	T	F	T							Yes	1	0			
	Not Less than or Equal (NLE)	T	F	F	T	F												
NGT	Not Greater than	F	T	T	T	T										Yes	1	0
	Greater than (GT)	T	F	F	F	F												

key: "?" = unordered, ">" = greater than, "<" = less than, "=" is equal, "T" = True, "F" = False

Table 5.25 FPU Comparisons With Special Operand Exceptions for QNaNs

The instruction encoding is an extension made in the MIPS IV architecture. In previous architecture levels the *cc* field for this instruction must be 0.

The MIPS I architecture defines a single floating-point condition code, implemented as the coprocessor 1 condition signal (Cp1Cond) and the C bit in the FP *Control and Status* register. MIPS I, II, and III architectures must have the *cc* field set to 0, which is implied by the first format in the *Format* section.

The MIPS IV architecture adds seven more condition code bits to the original condition code 0. FP compare and conditional branch instructions specify the condition code bit to set or test. Both assembler formats are valid for MIPS IV.

Restrictions:

The fields *fs* and *ft* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.

The operands must be values in format *fmt*. If they are not, the result is undefined and the value of the operand FPRs becomes undefined.

MIPS I, II, III: There must be at least one instruction between the compare instruction that sets a condition code and the branch instruction that tests it. Hardware does not detect a violation of this restriction.

Operation:

```

if NaN(Value FPR(fs, fmt)) or NaN(ValueFPR(ft, fmt)) then
    less ← false
    equal ← false
    unordered ← true
    if t then
        SignalException(InvalidOperation)
    endif
else
    less ← ValueFPR(fs, fmt) < ValueFPR(ft, fmt)
    equal ← ValueFPR(fs, fmt) = ValueFPR(ft, fmt)
    unordered ← false
endif
condition ← (cond2 and less) or (cond1 and equal) or (cond0 and unordered)
FCC[cc] ← condition
if cc = 0 then
    COC[1] ← condition
endif
    
```


Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Unimplemented Operation
 - Invalid Operation

Programming Notes:

FP computational instructions, including compare, that receive an operand value of Signaling NaN, will raise the Invalid Operation condition. The comparisons that raise the Invalid Operation condition for Quiet NaNs in addition to SNaNs, permit a simpler programming model if NaNs are errors. Using these compares, programs do not need explicit code to check for QNaNs causing the *unordered* relation. Instead, they take an exception and allow the exception handling system to deal with the error when it occurs. For example, consider a comparison in which we want to know if two numbers are equal, but for which unordered would be an error.

```
# comparisons using explicit tests for QNaN
c.eq.d  $f2,$f4  # check for equal
nop
bc1t    L2      # it is equal
c.un.d  $f2,$f4  # it is not equal, but might be unordered
bc1t    ERROR#  unordered goes off to an error handler
# not-equal-case code here
...
# equal-case code here
L2:
# -----
# comparison using comparisons that signal QNaN
c.seq.d $f2,$f4  # check for equal
nop
bc1t    L2      # it is equal
nop
# it is not unordered here...
# not-equal-case code here
...
#equal-case code here
L2:
```



Format: CEIL.L.S fd, fs **MIPS III**
 CEIL.L.D fd, fs

Purpose: To convert an FP value to 64-bit fixed-point, rounding up.

Description: fd ← convert_and_round(fs)

The value in FPR *fs* in format *fmt*, is converted to a value in 64-bit long fixed-point format rounding toward $+\infty$ (rounding mode 2). The result is placed in FPR *fd*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{63} to $2^{63}-1$, the result cannot be represented correctly and an IEEE Invalid Operation condition exists. The result depends on the FP exception model currently active.

- ◆ *Precise exception model:* The Invalid Operation flag is set in the FCSR. If the Invalid Operation enable bit is set in the FCSR, no result is written to *fd* and an Invalid Operation exception is taken immediately. Otherwise, the default result, $2^{63}-1$, is written to *fd*.
- ◆ *Imprecise exception model (R8000 normal mode):* The default result, $2^{63}-1$, is written to *fd*. No FCSR flag is set. If the Invalid Operation enable bit is set in the FCSR, an Invalid Operation exception is taken, imprecisely, at some future time.

Restrictions:

The fields *fs* and *fd* must specify valid FPRs; *fs* for type *fmt* and *fd* for long fixed-point. If they are not valid, the result is undefined.


The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR(*fd*, L, ConvertFmt(ValueFPR(*fs*, *fmt*), *fmt*, L))

Exceptions:

- | | |
|----------------------|-------------------------|
| Coprocessor Unusable | |
| Reserved Instruction | |
| Floating-Point | |
| Invalid Operation | Unimplemented Operation |
| Inexact | Overflow |



Format: CEIL.W.S *fd, fs* **MIPS II**
 CEIL.W.D *fd, fs*

Purpose: To convert an FP value to 32-bit fixed-point, rounding up.

Description: $fd \leftarrow \text{convert_and_round}(fs)$

The value in FPR *fs* in format *fmt*, is converted to a value in 32-bit word fixed-point format rounding toward $+\infty$ (rounding mode 2). The result is placed in FPR *fd*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{31} to $2^{31}-1$, the result cannot be represented correctly and an IEEE Invalid Operation condition exists. The result depends on the FP exception model currently active.

- ◆ *Precise exception model: The Invalid Operation flag is set in the FCSR. If the Invalid Operation enable bit is set in the FCSR, no result is written to *fd* and an Invalid Operation exception is taken immediately. Otherwise, the default result, $2^{31}-1$, is written to *fd*.*
- ◆ *Imprecise exception model (R8000 normal mode): The default result, $2^{31}-1$, is written to *fd*. No FCSR flag is set. If the Invalid Operation enable bit is set in the FCSR, an Invalid Operation exception is taken, imprecisely, at some future time.*

Restrictions:

The fields *fs* and *fd* must specify valid FPRs; *fs* for type *fmt* and *fd* for word fixed-point. If they are not valid, the result is undefined.


The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR(*fd*, *W*, ConvertFmt(ValueFPR(*fs*, *fmt*), *fmt*, *W*))

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Invalid Operation
 - Unimplemented Operation
 - Inexact
 - Overflow



Format: CFC1 rt, fs **MIPS I**

Purpose: To copy a word from an FPU control register to a GPR.

Description: $rt \leftarrow FP_Control[fs]$

Copy the 32-bit word from FP (coprocessor 1) control register *fs* into GPR *rt*, sign-extending it if the GPR is 64 bits.

Restrictions:

There are only a couple control registers defined for the floating-point unit. The result is not defined if *fs* specifies a register that does not exist.

For MIPS I, MIPS II, and MIPS III, the contents of GPR *rt* are undefined for the instruction immediately following CFC1.

Operation: MIPS I - III


I: temp $\leftarrow FCR[fs]$
 I+1: GPR[rt] $\leftarrow sign_extend(temp)$

Operation: MIPS IV

temp $\leftarrow FCR[fs]$
 GPR[rt] $\leftarrow sign_extend(temp)$

Exceptions:

Coprocessor Unusable



Format: CTC1 rt, fs MIPS I

Purpose: To copy a word from a GPR to an FPU control register.

Description: FP_Control[fs] ← rt

Copy the low word from GPR *rt* into FP (coprocessor 1) control register *fs*.

Writing to control register 31, the *Floating-Point Control and Status Register* or FCSR, causes the appropriate exception if any cause bit and its corresponding enable bit are both set. The register will be written before the exception occurs.

Restrictions:

There are only a couple control registers defined for the floating-point unit. The result is not defined if *fs* specifies a register that does not exist.

For MIPS I, MIPS II, and MIPS III, the contents of floating-point control register *fs* are undefined for the instruction immediately following CTC1.

Operation: MIPS I - III


- I: temp ← GPR[rt]_{31..0}
- I+1: FCR[fs] ← temp
- COC[1] ← FCR[31]₂₃

Operation: MIPS IV

- temp ← GPR[rt]_{31..0}
- FCR[fs] ← temp
- COC[1] ← FCR[31]₂₃

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Unimplemented Operation
 - Invalid Operation
 - Division-by-zero
 - Inexact
 - Overflow
 - Underflow



Format: CVT.D.S *fd, fs* MIPS I
 CVT.D.W *fd, fs*
 CVT.D.L *fd, fs* MIPS III

Purpose: To convert an FP or fixed-point value to double FP.

Description: $fd \leftarrow \text{convert_and_round}(fs)$

The value in FPR *fs* in format *fmt* is converted to a value in double floating-point format rounded according to the current rounding mode in FCSR. The result is placed in FPR *fd*.

If *fmt* is S or W, then the operation is always exact.

Restrictions:

The fields *fs* and *fd* must specify valid FPRs; *fs* for type *fmt* and *fd* for double floating-point. If they are not valid, the result is undefined.

The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR (*fd*, D, ConvertFmt(ValueFPR(*fs*, *fmt*), *fmt*, D))

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Invalid Operation
 - Unimplemented Operation
 - Inexact
 - Overflow
 - Underflow

31	26	25	21	20	16	15	11	10	6	5	0
COP1 010001	fmt		0 00000		fs		fd		CVT.L 100101		
6	5		5		5		5		6		

Format: CVT.L.S *fd, fs* **MIPS III**
 CVT.L.D *fd, fs*

Purpose: To convert an FP value to a 64-bit fixed-point.

Description: $fd \leftarrow \text{convert_and_round}(fs)$

Convert the value in format *fmt* in FPR *fs* to long fixed-point format, round according to the current rounding mode in FCSR, and place the result in FPR *fd*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{63} to $2^{63}-1$, the result cannot be represented correctly and an IEEE Invalid Operation condition exists. The result depends on the FP exception model currently active:

- ◆ *Precise exception model: The Invalid Operation flag is set in the FCSR. If the Invalid Operation enable bit is set in the FCSR, no result is written to *fd* and an Invalid Operation exception is taken immediately. Otherwise, the default result, $2^{63}-1$, is written to *fd*.*
- ◆ *Imprecise exception model (R8000 normal mode): The default result, $2^{63}-1$, is written to *fd*. No FCSR flag is set. If the Invalid Operation enable bit is set in the FCSR, an Invalid Operation exception is taken, imprecisely, at some future time.*

Restrictions:

The fields *fs* and *fd* must specify valid FPRs; *fs* for type *fmt* and *fd* for long fixed-point. If they are not valid, the result is undefined.


The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR (*fd*, L, ConvertFmt(ValueFPR(*fs*, *fmt*), *fmt*, L))

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Invalid Operation
 - Unimplemented Operation
 - Inexact
 - Overflow



Format: CVT.S.D fd, fs MIPS I
 CVT.S.W fd, fs
 CVT.S.L fd, fs MIPS III

Purpose: To convert an FP or fixed-point value to single FP.

Description: $fd \leftarrow \text{convert_and_round}(fs)$

The value in FPR *fs* in format *fmt* is converted to a value in single floating-point format rounded according to the current rounding mode in FCSR. The result is placed in FPR *fd*.

Restrictions:

The fields *fs* and *fd* must specify valid FPRs; *fs* for type *fmt* and *fd* for single floating-point. If they are not valid, the result is undefined.


The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR(*fd*, *S*, ConvertFmt(ValueFPR(*fs*, *fmt*), *fmt*, *S*))

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Invalid Operation
 - Unimplemented Operation
 - Inexact
 - Overflow
 - Underflow



Format: CVT.W.S *fd, fs* **MIPS I**
 CVT.W.D *fd, fs*

Purpose: To convert an FP value to 32-bit fixed-point.

Description: $fd \leftarrow \text{convert_and_round}(fs)$

The value in FPR *fs* in format *fmt* is converted to a value in 32-bit word fixed-point format rounded according to the current rounding mode in FCSR. The result is placed in FPR *fd*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{31} to $2^{31}-1$, the result cannot be represented correctly and an IEEE Invalid Operation condition exists. The result depends on the FP exception model currently active.

- ◆ *Precise exception model: The Invalid Operation flag is set in the FCSR. If the Invalid Operation enable bit is set in the FCSR, no result is written to *fd* and an Invalid Operation exception is taken immediately. Otherwise, the default result, $2^{31}-1$, is written to *fd*.*
- ◆ *Imprecise exception model (R8000 normal mode): The default result, $2^{31}-1$, is written to *fd*. No FCSR flag is set. If the Invalid Operation enable bit is set in the FCSR, an Invalid Operation exception is taken, imprecisely, at some future time.*

Restrictions:

The fields *fs* and *fd* must specify valid FPRs; *fs* for type *fmt* and *fd* for word fixed-point. If they are not valid, the result is undefined.


The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR(*fd*, W, ConvertFmt(ValueFPR(*fs*, *fmt*), *fmt*, W))

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Invalid Operation
 - Unimplemented Operation
 - Inexact
 - Overflow



Format: DIV.S fd, fs, ft MIPS I
 DIV.D fd, fs, ft

Purpose: To divide FP values.

Description: $fd \leftarrow fs / ft$

The value in FPR *fs* is divided by the value in FPR *ft*. The result is calculated to infinite precision, rounded according to the current rounding mode in FCSR, and placed into FPR *fd*. The operands and result are values in format *fmt*.

Restrictions:

The fields *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.


The operands must be values in format *fmt*. If they are not, the result is undefined and the value of the operand FPRs becomes undefined.

Operation:

StoreFPR (fd, fmt, ValueFPR(fs, fmt) / ValueFPR(ft, fmt))

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Inexact Unimplemented Operation
 - Division-by-zero Invalid Operation
 - Overflow Underflow



Format: DMFC1 rt, fs **MIPS III**

Purpose: To copy a doubleword from an FPR to a GPR.

Description: $rt \leftarrow fs$

The doubleword contents of FPR *fs* are placed into GPR *rt*.

If the coprocessor 1 general registers are 32-bits wide (a native 32-bit processor or 32-bit register emulation mode in a 64-bit processor), FPR *fs* is held in an even/odd register pair. The low word is taken from the even register *fs* and the high word is from *fs+1*.

Restrictions:

If *fs* does not specify an FPR that can contain a doubleword, the result is undefined.

For MIPS III, the contents of GPR *rt* are undefined for the instruction immediately following DMFC1.

Operation: MIPS I - III

```

I:  if SizeFGR() = 64 then           /* 64-bit wide FGRs */
      data ← FGR[fs]
    elseif fs0 = 0 then               /* valid specifier, 32-bit wide FGRs */
      data ← FGR[fs+1] || FGR[fs]
    else                               /* undefined for odd 32-bit FGRs */
      UndefinedResult()
    endif
I+1: GPR[rt] ← data
  
```


Operation: MIPS IV

```

if SizeFGR() = 64 then           /* 64-bit wide FGRs */
  data ← FGR[fs]
elseif fs0 = 0 then             /* valid specifier, 32-bit wide FGRs */
  data ← FGR[fs+1] || FGR[fs]
else                             /* undefined for odd 32-bit FGRs */
  UndefinedResult()
endif
GPR[rt] ← data
  
```

Exceptions:

- Reserved Instruction
- Coprocessor Unusable



Format: DMTC1 rt, fs MIPS III

Purpose: To copy a doubleword from a GPR to an FPR.

Description: fs ← rt

The doubleword contents of GPR *rt* are placed into FPR *fs*.

If coprocessor 1 general registers are 32-bits wide (a native 32-bit processor or 32-bit register emulation mode in a 64-bit processor), FPR *fs* is held in an even/odd register pair. The low word is placed in the even register *fs* and the high word is placed in *fs+1*.

Restrictions:

If *fs* does not specify an FPR that can contain a doubleword, the result is undefined.

For MIPS III, the contents of FPR *fs* are undefined for the instruction immediately following DMTC1.

Operation: MIPS I - III

```

I: data ← GPR[rt]
I+1: if SizeFGR() = 64 then /* 64-bit wide FGRs */
    FGR[fs] ← data
elseif fs0 = 0 then /* valid specifier, 32-bit wide FGRs */
    FGR[fs+1] ← data63..32
    FGR[fs] ← data31..0
else /* undefined result for odd 32-bit FGRs */
    UndefinedResult()
endif
    
```


Operation: MIPS IV

```

data ← GPR[rt]
if SizeFGR() = 64 then /* 64-bit wide FGRs */
    FGR[fs] ← data
elseif fs0 = 0 then /* valid specifier, 32-bit wide FGRs */
    FGR[fs+1] ← data63..32
    FGR[fs] ← data31..0
else /* undefined result for odd 32-bit FGRs */
    UndefinedResult()
endif
    
```

Exceptions:

- Reserved Instruction
- Coprocessor Unusable



Format: FLOOR.L.S fd, fs MIPS III
 FLOOR.L.D fd, fs

Purpose: To convert an FP value to 64-bit fixed-point, rounding down.

Description: $fd \leftarrow \text{convert_and_round}(fs)$

The value in FPR *fs* in format *fmt*, is converted to a value in 64-bit long fixed-point format rounding toward $-\infty$ (rounding mode 3). The result is placed in FPR *fd*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{63} to $2^{63}-1$, the result cannot be represented correctly and an IEEE Invalid Operation condition exists. The result depends on the FP exception model currently active.

- ◆ *Precise exception model: The Invalid Operation flag is set in the FCSR. If the Invalid Operation enable bit is set in the FCSR, no result is written to fd and an Invalid Operation exception is taken immediately. Otherwise, the default result, $2^{63}-1$, is written to fd.*
- ◆ *Imprecise exception model (R8000 normal mode): The default result, $2^{63}-1$, is written to fd. No FCSR flag is set. If the Invalid Operation enable bit is set in the FCSR, an Invalid Operation exception is taken, imprecisely, at some future time.*

Restrictions:

The fields *fs* and *fd* must specify valid FPRs; *fs* for type *fmt* and *fd* for long fixed-point. If they are not valid, the result is undefined.


The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR(fd, L, ConvertFmt(ValueFPR(fs, fmt), fmt, L))

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Invalid Operation Unimplemented Operation
 - Inexact Overflow



Format: FLOOR.W.S fd, fs MIPS II
 FLOOR.W.D fd, fs

Purpose: To convert an FP value to 32-bit fixed-point, rounding down.

Description: $fd \leftarrow \text{convert_and_round}(fs)$

The value in FPR *fs* in format *fmt*, is converted to a value in 32-bit word fixed-point format rounding toward $-\infty$ (rounding mode 3). The result is placed in FPR *fd*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{31} to $2^{31}-1$, the result cannot be represented correctly and an IEEE Invalid Operation condition exists. The result depends on the FP exception model currently active.

- ◆ *Precise exception model:* The Invalid Operation flag is set in the FCSR. If the Invalid Operation enable bit is set in the FCSR, no result is written to *fd* and an Invalid Operation exception is taken immediately. Otherwise, the default result, $2^{31}-1$, is written to *fd*.
- ◆ *Imprecise exception model (R8000 normal mode):* The default result, $2^{31}-1$, is written to *fd*. No FCSR flag is set. If the Invalid Operation enable bit is set in the FCSR, an Invalid Operation exception is taken, imprecisely, at some future time.

Restrictions:

The fields *fs* and *fd* must specify valid FPRs; *fs* for type *fmt* and *fd* for word fixed-point. If they are not valid, the result is undefined.


The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR(*fd*, W, ConvertFmt(ValueFPR(*fs*, *fmt*), *fmt*, W))

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Invalid Operation Unimplemented Operation
 - Inexact Overflow



Format: LDC1 ft, offset(base)

MIPS II

Purpose: To load a doubleword from memory to an FPR.

Description: $ft \leftarrow \text{memory}[\text{base}+\text{offset}]$

The contents of the 64-bit doubleword at the memory location specified by the aligned effective address are fetched and placed in FPR *ft*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

If coprocessor 1 general registers are 32-bits wide (a native 32-bit processor or 32-bit register emulation mode in a 64-bit processor), FPR *ft* is held in an even/odd register pair. The low word is placed in the even register *ft* and the high word is placed in *ft+1*.

Restrictions:

If *ft* does not specify an FPR that can contain a doubleword, the result is undefined.

An Address Error exception occurs if $\text{EffectiveAddress}_{2..0} \neq 0$ (not doubleword-aligned).

MIPS IV: The low-order 3 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.


Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr2..0 ≠ 03 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
data ← LoadMemory(uncached, DOUBLEWORD, pAddr, vAddr, DATA)
if SizeFGR() = 64 then                /* 64-bit wide FGRs */
    FGR[ft] ← data
elseif ft0 = 0 then                  /* valid specifier, 32-bit wide FGRs */
    FGR[ft+1] ← data63..32
    FGR[ft] ← data31..0
else                                  /* undefined result for odd 32-bit FGRs */
    UndefinedResult()
endif
    
```

Exceptions:

- Coprocessor unusable
- Reserved Instruction
- TLB Refill, TLB Invalid
- Address Error



Format: LDXC1 fd, index(base) **MIPS IV**

Purpose: To load a doubleword from memory to an FPR (GPR+GPR addressing).

Description: $fd \leftarrow \text{memory}[\text{base} + \text{index}]$

The contents of the 64-bit doubleword at the memory location specified by the aligned effective address are fetched and placed in FPR *fd*. The contents of GPR *index* and GPR *base* are added to form the effective address.

If coprocessor 1 general registers are 32-bits wide (a native 32-bit processor or 32-bit register emulation mode in a 64-bit processor), FPR *fd* is held in an even/odd register pair. The low word is placed in the even register *fd* and the high word is placed in *fd+1*.

Restrictions:

If *fd* does not specify an FPR that can contain a doubleword, the result is undefined.

The Region bits of the effective address must be supplied by the contents of *base*. If $\text{EffectiveAddress}_{63..62} \neq \text{base}_{63..62}$, the result is undefined.

An Address Error exception occurs if $\text{EffectiveAddress}_{2..0} \neq 0$ (not doubleword-aligned).

MIPS IV: The low-order 3 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.


Operation:

```

vAddr ← GPR[base] + GPR[index]
if vAddr2..0 ≠ 03 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
mem ← LoadMemory(unchched, DOUBLEWORD, pAddr, vAddr, DATA)
if SizeFGR() = 64 then /* 64-bit wide FGRs */
    FGR[fd] ← data
elseif fd0 = 0 then /* valid specifier, 32-bit wide FGRs */
    FGR[fd+1] ← data63..32
    FGR[fd] ← data31..0
else /* undefined result for odd 32-bit FGRs */
    UndefinedResult()
endif
    
```

Exceptions:

- TLB Refill, TLB Invalid
- Address Error
- Reserved Instruction
- Coprocessor Unusable



Format: LWC1 ft, offset(base) MIPS I

Purpose: To load a word from memory to an FPR.

Description: $ft \leftarrow \text{memory}[\text{base}+\text{offset}]$

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched and placed into the low word of coprocessor 1 general register *ft*. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

If coprocessor 1 general registers are 64-bits wide, bits 63..32 of register *ft* become undefined.

Restrictions:

An Address Error exception occurs if $\text{EffectiveAddress}_{1..0} \neq 0$ (not word-aligned).

MIPS IV: The low-order 2 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation: 32-bit Processors

```

I: /* "mem" is aligned 64-bits from memory. Pick out correct bytes. */
   vAddr ← sign_extend(offset) + GPR[base]
   if vAddr1..0 ≠ 02 then SignalException(AddressError) endif
   (pAddr, uncached) ← AddressTranslation(vAddr, DATA, LOAD)
   mem ← LoadMemory(uncached, WORD, pAddr, vAddr, DATA)
I+1: FGR[ft] ← mem
    
```


Operation: 64-bit Processors

```

/* "mem" is aligned 64-bits from memory. Pick out correct bytes. */
vAddr ← sign_extend(offset) + GPR[base]
if vAddr1..0 ≠ 02 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation(vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
mem ← LoadMemory(uncached, WORD, pAddr, vAddr, DATA)
bytesel ← vAddr2..0 xor (BigEndianCPU || 02)
if SizeFGR() = 64 then /* 64-bit wide FGRs */
   FGR[ft] ← undefined32 || mem31+8*bytesel..8*bytesel
else /* 32-bit wide FGRs */
   FGR[ft] ← mem31+8*bytesel..8*bytesel
endif
    
```

Exceptions:

- Coprocessor unusable
- Reserved Instruction
- TLB Refill, TLB Invalid
- Address Error



Format: LWXC1 fd, index(base) **MIPS IV**

Purpose: To load a word from memory to an FPR (GPR+GPR addressing).

Description: $fd \leftarrow \text{memory}[\text{base}+\text{index}]$

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched and placed into the low word of coprocessor 1 general register *fd*. The contents of GPR *index* and GPR *base* are added to form the effective address.

If coprocessor 1 general registers are 64-bits wide, bits 63..32 of register *fd* become undefined.

Restrictions:

The Region bits of the effective address must be supplied by the contents of *base*. If $\text{EffectiveAddress}_{63..62} \neq \text{base}_{63..62}$, the result is undefined.

An Address Error exception occurs if $\text{EffectiveAddress}_{1..0} \neq 0$ (not word-aligned).

MIPS IV: The low-order 2 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation:


```

vAddr ← GPR[base] + GPR[index]
if vAddr1..0 ≠ 02 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
/* "mem" is aligned 64-bits from memory. Pick out correct bytes. */
mem ← LoadMemory(uncached, WORD, pAddr, vAddr, DATA)
bytesel ← vAddr2..0 xor (BigEndianCPU || 02)
if SizeFGR() = 64 then /* 64-bit wide FGRs */
    FGR[fd] ← undefined32 || mem31+8*bytesel..8*bytesel
else /* 32-bit wide FGRs */
    FGR[fd] ← mem31+8*bytesel..8*bytesel
endif

```

Exceptions:

- TLB Refill, TLB Invalid
- Address Error
- Reserved Instruction
- Coprocessor Unusable



Format: MADD.S *fd, fr, fs, ft* MIPS IV
MADD.D *fd, fr, fs, ft*

Purpose: To perform a combined multiply-then-add of FP values.

Description: $fd \leftarrow (fs \times ft) + fr$

The value in FPR *fs* is multiplied by the value in FPR *ft* to produce a product. The value in FPR *fr* is added to the product. The result sum is calculated to infinite precision, rounded according to the current rounding mode in FCSR, and placed into FPR *fd*. The operands and result are values in format *fmt*.

The accuracy of the result depends which of two alternative arithmetic models is used by the implementation for the computation. The numeric models are explained in the section **Arithmetic Instructions** towards the beginning of this Chapter.

Restrictions:

The fields *fr*, *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.


The operands must be values in format *fmt*. If they are not, the result is undefined and the value of the operand FPRs becomes undefined.

Operation:

- $vfr \leftarrow \text{ValueFPR}(fr, fmt)$
- $vfs \leftarrow \text{ValueFPR}(fs, fmt)$
- $vft \leftarrow \text{ValueFPR}(ft, fmt)$
- StoreFPR(*fd*, *fmt*, $vfr + vfs * vft$)

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Inexact Unimplemented Operation
 - Invalid Operation Overflow
 - Underflow



Format: MFC1 rt, fs MIPS I

Purpose: To copy a word from an FPU (CP1) general register to a GPR.

Description: $rt \leftarrow fs$

The low word from FPR *fs* is placed into the low word of GPR *rt*. If GPR *rt* is 64 bits wide, then the value is sign extended.

Restrictions:

For MIPS I, MIPS II, and MIPS III the contents of GPR *rt* are undefined for the instruction immediately following MFC1.

Operation: MIPS I - III


I: word \leftarrow FGR[fs]_{31..0}
I+1: GPR[rt] \leftarrow sign_extend(word)

Operation: MIPS IV

word \leftarrow FGR[fs]_{31..0}
 GPR[rt] \leftarrow sign_extend(word)

Exceptions:

Coprocessor Unusable



Format: MOV.S *fd*, *fs* MIPS I
 MOV.D *fd*, *fs*

Purpose: To move an FP value between FPRs.

Description: $fd \leftarrow fs$

The value in FPR *fs* is placed into FPR *fd*. The source and destination are values in format *fmt*.

The move is non-arithmetic; it causes no IEEE 754 exceptions.

Restrictions:

The fields *fs* and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.


The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR(*fd*, *fmt*, ValueFPR(*fs*, *fmt*))

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Unimplemented Operation



Format: MOVF rd, rs, cc MIPS IV

Purpose: To test an FP condition code then conditionally move a GPR.

Description: if (cc = 0) then rd ← rs

If the floating-point condition code specified by *cc* is zero, then the contents of **GPR** *rs* are placed into **GPR** *rd*.

Restrictions:

None


Operation:

```

active ← FCC[cc] = tf
if active then
    GPR[rd] ← GPR[rs]
endif
    
```

Exceptions:

Reserved Instruction
Coprocessor Unusable



Format: MOVF.S fd, fs, cc MIPS IV
 MOVF.D fd, fs, cc

Purpose: To test an FP condition code then conditionally move an FP value.

Description: if (cc = 0) then fd ← fs

If the floating-point condition code specified by *cc* is zero, then the value in FPR *fs* is placed into FPR *fd*. The source and destination are values in format *fmt*.

If the condition code is not zero, then FPR *fs* is not copied and FPR *fd* contains its previous value in format *fmt*. If *fd* did not contain a value either in format *fmt* or previously unused data from a load or move-to operation that could be interpreted in format *fmt*, then the value of *fd* becomes undefined.

The move is non-arithmetic; it causes no IEEE 754 exceptions.

Restrictions:

The fields *fs* and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.

The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.


Operation:

```

if FCC[cc] = tf then
    StoreFPR(fd, fmt, ValueFPR(fs, fmt))
else
    StoreFPR(fd, fmt, ValueFPR(fd, fmt))
endif
    
```

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Unimplemented operation



Format: MOVN.S fd, fs, rt MIPS IV
 MOVN.D fd, fs, rt

Purpose: To test a GPR then conditionally move an FP value.

Description: if (rt ≠ 0) then fd ← fs

If the value in GPR *rt* is not equal to zero then the value in FPR *fs* is placed in FPR *fd*. The source and destination are values in format *fmt*.

If GPR *rt* contains zero, then FPR *fs* is not copied and FPR *fd* contains its previous value in format *fmt*. If *fd* did not contain a value either in format *fmt* or previously unused data from a load or move-to operation that could be interpreted in format *fmt*, then the value of *fd* becomes undefined.

The move is non-arithmetic; it causes no IEEE 754 exceptions.

Restrictions:

The fields *fs* and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.

The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.


Operation:

```

if GPR[rt] ≠ 0 then
    StoreFPR(fd, fmt, ValueFPR(fs, fmt))
else
    StoreFPR(fd, fmt, ValueFPR(fd, fmt))
endif
    
```

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point Unimplemented operation



Format: MOVT rd, rs, cc MIPS IV

Purpose: To test an FP condition code then conditionally move a GPR.

Description: if (cc = 1) then rd ← rs

If the floating-point condition code specified by *cc* is one then the contents of **GPR** *rs* are placed into **GPR** *rd*.

Restrictions:


None

Operation:

```
if FCC[cc] = tf then
    GPR[rd] ← GPR[rs]
endif
```

Exceptions:

Reserved Instruction
Coprocesor Unusable



Format: MOVT.S *fd, fs, cc* MIPS IV
 MOVT.D *fd, fs, cc*

Purpose: To test an FP condition code then conditionally move an FP value.

Description: if (*cc* = 1) then *fd* ← *fs*

If the floating-point condition code specified by *cc* is one then the value in FPR *fs* is placed into FPR *fd*. The source and destination are values in format *fmt*.

If the condition code is not one, then FPR *fs* is not copied and FPR *fd* contains its previous value in format *fmt*. If *fd* did not contain a value either in format *fmt* or previously unused data from a load or move-to operation that could be interpreted in format *fmt*, then the value of *fd* becomes undefined.

The move is non-arithmetic; it causes no IEEE 754 exceptions.

Restrictions:

The fields *fs* and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.

The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.


Operation:

```

if FCC[cc] = tf then
  StoreFPR(fd, fmt, ValueFPR(fs, fmt))
else
  StoreFPR(fd, fmt, ValueFPR(fd, fmt))
endif
  
```

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point Unimplemented operation



Format: MOVZ.S *fd, fs, rt* MIPS IV
 MOVZ.D *fd, fs, rt*

Purpose: To test a GPR then conditionally move an FP value.

Description: if (*rt* = 0) then *fd* ← *fs*

If the value in GPR *rt* is equal to zero then the value in FPR *fs* is placed in FPR *fd*. The source and destination are values in format *fmt*.

If GPR *rt* is not zero, then FPR *fs* is not copied and FPR *fd* contains its previous value in format *fmt*. If *fd* did not contain a value either in format *fmt* or previously unused data from a load or move-to operation that could be interpreted in format *fmt*, then the value of *fd* becomes undefined.

The move is non-arithmetic; it causes no IEEE 754 exceptions.

Restrictions:

The fields *fs* and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.

The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.


Operation:

```

if GPR[rt] = 0 then
    StoreFPR(fd, fmt, ValueFPR(fs, fmt))
else
    StoreFPR(fd, fmt, ValueFPR(fd, fmt))
endif
    
```

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Unimplemented operation



Format: MSUB.S fd, fr, fs, ft **MIPS IV**
 MSUB.D fd, fr, fs, ft

Purpose: To perform a combined multiply-then-subtract of FP values.

Description: $fd \leftarrow (fs \times ft) - fr$

The value in FPR *fs* is multiplied by the value in FPR *ft* to produce an intermediate product. The value in FPR *fr* is subtracted from the product. The subtraction result is calculated to infinite precision, rounded according to the current rounding mode in FCSR, and placed into FPR *fd*. The operands and result are values in format *fmt*.

The accuracy of the result depends which of two alternative arithmetic models is used by the implementation for the computation. The numeric models are explained in the section **Arithmetic Instructions** earlier on in this Chapter.

Restrictions:

The fields *fr*, *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.


The operands must be values in format *fmt*. If they are not, the result is undefined and the value of the operand FPRs becomes undefined.

Operation:

$vfr \leftarrow \text{ValueFPR}(fr, \text{fmt})$
 $vfs \leftarrow \text{ValueFPR}(fs, \text{fmt})$
 $vft \leftarrow \text{ValueFPR}(ft, \text{fmt})$
 StoreFPR(*fd*, *fmt*, (*vfs* * *vft*) - *vfr*)

Exceptions:

- | | |
|----------------------|-------------------------|
| Reserved Instruction | |
| Coprocessor Unusable | |
| Floating-Point | |
| Inexact | Unimplemented Operation |
| Invalid Operation | Overflow |
| Underflow | |



Format: MTC1 rt, fs **MIPS I**

Purpose: To copy a word from a GPR to an FPU (CP1) general register.

Description: fs ← rt

The low word in GPR *rt* is placed into the low word of floating-point (coprocessor 1) general register *fs*. If coprocessor 1 general registers are 64-bits wide, bits 63..32 of register *fs* become undefined.

Restrictions:

For MIPS I, MIPS II, and MIPS III the value of FPR *fs* is undefined for the instruction immediately following MTC1.

Operation: MIPS I - III

```

I: data ← GPR[rt]31..0
I+1: if SizeFGR() = 64 then /* 64-bit wide FGRs */
    FGR[fs] ← undefined32 || data
else /* 32-bit wide FGRs */
    FGR[fs] ← data
endif
    
```


Operation: MIPS IV

```

data ← GPR[rt]31..0
if SizeFGR() = 64 then /* 64-bit wide FGRs */
    FGR[fs] ← undefined32 || data
else /* 32-bit wide FGRs */
    FGR[fs] ← data
endif
    
```

Exceptions:

Coprocessor Unusable



Format: MUL.S fd, fs, ft MIPS I
 MUL.D fd, fs, ft

Purpose: To multiply FP values.

Description: $fd \leftarrow fs \times ft$

The value in FPR *fs* is multiplied by the value in FPR *ft*. The result is calculated to infinite precision, rounded according to the current rounding mode in FCSR, and placed into FPR *fd*. The operands and result are values in format *fmt*.

Restrictions:

The fields *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.


The operands must be values in format *fmt*. If they are not, the result is undefined and the value of the operand FPRs becomes undefined.

Operation:

StoreFPR (fd, fmt, ValueFPR(fs, fmt) * ValueFPR(ft, fmt))

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Inexact Unimplemented Operation
 - Invalid Operation Overflow
 - Underflow



Format: NEG.S *fd*, *fs* MIPS I
 NEG.D *fd*, *fs*

Purpose: To negate an FP value.

Description: $fd \leftarrow -(fs)$

The value in FPR *fs* is negated and placed into FPR *fd*. The value is negated by changing the sign bit value. The operand and result are values in format *fmt*.

This operation is arithmetic; a NaN operand signals invalid operation.

Restrictions:

The fields *fs* and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.


The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR(*fd*, *fmt*, Negate(ValueFPR(*fs*, *fmt*)))

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Unimplemented Operation
 - Invalid Operation



Format: NMADD.S fd, fr, fs, ft MIPS IV
 NMADD.D fd, fr, fs, ft

Purpose: To negate a combined multiply-then-add of FP values.

Description: $fd \leftarrow -((fs \times ft) + fr)$

The value in FPR *fs* is multiplied by the value in FPR *ft* to produce an intermediate product. The value in FPR *fr* is added to the product. The result sum is calculated to infinite precision, rounded according to the current rounding mode in FCSR, negated by changing the sign bit, and placed into FPR *fd*. The operands and result are values in format *fmt*.

The accuracy of the result depends which of two alternative arithmetic models is used by the implementation for the computation.

Restrictions:

The fields *fr*, *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.


The operands must be values in format *fmt*. If they are not, the result is undefined and the value of the operand FPRs becomes undefined.

Operation:

$vfr \leftarrow \text{ValueFPR}(fr, fmt)$
 $vfs \leftarrow \text{ValueFPR}(fs, fmt)$
 $vft \leftarrow \text{ValueFPR}(ft, fmt)$
 StoreFPR(*fd*, *fmt*, $-(vfr + vfs * vft)$)

Exceptions:

Coprocessor Unusable	
Reserved Instruction	
Floating-Point	
Inexact	Unimplemented Operation
Invalid Operation	Overflow
Underflow	



Format: NMSUB.S *fd, fr, fs, ft* **MIPS IV**
 NMSUB.D *fd, fr, fs, ft*

Purpose: To negate a combined multiply-then-subtract of FP values.

Description: $fd \leftarrow - ((fs \times ft) - fr)$

The value in FPR *fs* is multiplied by the value in FPR *ft* to produce an intermediate product. The value in FPR *fr* is subtracted from the product. The result is calculated to infinite precision, rounded according to the current rounding mode in FCSR, negated by changing the sign bit, and placed into FPR *fd*. The operands and result are values in format *fmt*.

The accuracy of the result depends which of two alternative arithmetic models is used by the implementation for the computation. The numeric models are explained in the section **Arithmetic Instructions** earlier on in this Chapter.

Restrictions:

The fields *fr*, *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.


The operands must be values in format *fmt*. If they are not, the result is undefined and the value of the operand FPRs becomes undefined.

Operation:

- $vfr \leftarrow \text{ValueFPR}(fr, fmt)$
- $vfs \leftarrow \text{ValueFPR}(fs, fmt)$
- $vft \leftarrow \text{ValueFPR}(ft, fmt)$
- $\text{StoreFPR}(fd, fmt, -((vfs * vft) - vfr))$

Exceptions:

- Reserved Instruction
- Coprocessor Unusable
- Floating-Point
 - Inexact Unimplemented Operation
 - Invalid Operation Overflow
 - Underflow



Format: PREFX hint, index(base) **MIPS IV**

Purpose: To prefetch locations from memory (GPR+GPR addressing).

Description: prefetch_memory[base+index]

PREFX adds the contents of GPR *index* 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. The *hint* field supplies information about the way that the data is expected to be used.

PREFX is an advisory instruction. It may change the performance of the program. For all *hint* values, it neither changes architecturally-visible state nor alters the meaning of the program. An implementation may do nothing when executing a PREFX instruction.

If MIPS IV instructions are supported and enabled and Coprocessor 1 is enabled (allowing access to CP1X), PREFX does not cause addressing-related exceptions. If it raises an exception condition, the exception condition is ignored. If an addressing-related exception condition 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.

PREFX will never generate a memory operation for a location with an uncached memory access type.

If PREFX results in a memory operation, the memory access type used for the operation is determined by the memory access type of the effective address, just as it would be if the memory operation had been caused by a load or store to the effective address.

PREFX enables the processor to take some action, typically prefetching the data into cache, to improve program performance. The action taken for a specific PREFX instruction is both system and context dependent. Any action, including doing nothing, is permitted that does not change architecturally-visible state or alter the meaning of a program. It is expected that implementations will either do nothing or take an action that will increase the performance of the program.

For a cached location, the expected, and useful, action is for the processor to prefetch a block of data that includes the effective address. The size of the block, and the level of the memory hierarchy it is fetched into are implementation specific.

The *hint* field supplies information about the way the data is expected to be used. No *hint* value causes an action that modifies architecturally-visible state. A processor may use a *hint* value to improve the effectiveness of the prefetch action. The defined *hint* values and the recommended prefetch action are shown in the table below. The *hint* table may be extended in future implementations.

Table 0-1 Values of Hint Field for Prefetch Instruction

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	store	Data is expected to be stored or modified. Fetch data as if for a store.
2-3		Not yet defined.
4	load_streamed	Data is expected to be loaded (not modified) but not reused extensively; it will “stream” through cache. Fetch data as if for a load and place it in the cache so that it will not displace data prefetched as “retained”.
5	store_streamed	Data is expected to be stored or modified but not reused extensively; it will “stream” through cache. Fetch data as if for a store and place it in the cache so that it will not displace data prefetched as “retained”.
6	load_retained	Data is expected to be loaded (not modified) and reused extensively; it should be “retained” in the cache. Fetch data as if for a load and place it in the cache so that it will not be displaced by data prefetched as “streamed”.
7	store_retained	Data is expected to be stored or modified and reused extensively; it should be “retained” in the cache. Fetch data as if for a store and place it in the cache so that will not be displaced by data prefetched as “streamed”.
8-31		Not yet defined.

Restrictions:

The Region bits of the effective address must be supplied by the contents of *base*. If $\text{EffectiveAddress}_{63..62} \neq \text{base}_{63..62}$, the result of the instruction is undefined.

Operation:

$vAddr \leftarrow \text{GPR}[\text{base}] + \text{GPR}[\text{index}]$
 $(pAddr, \text{uncached}) \leftarrow \text{AddressTranslation}(vAddr, \text{DATA}, \text{LOAD})$
 Prefetch(uncached, pAddr, vAddr, DATA, hint)

Exceptions:

- Reserved Instruction
- Coprocessor Unusable


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 is determined.

Implementation Notes:

It is recommended that a reserved *hint* field value either cause a default prefetch action that is expected to be useful for most cases of data use, such as the “load” *hint*, or cause the instruction to be treated as a NOP.



Format: RECIP.S fd, fs **MIPS IV**
 RECIP.D fd, fs

Purpose: To approximate the reciprocal of an FP value (quickly).

Description: $fd \leftarrow 1.0 / fs$

The reciprocal of the value in FPR *fs* is approximated and placed into FPR *fd*.
 The operand and result are values in format *fmt*.

The numeric accuracy of this operation is implementation dependent; it does not meet the accuracy specified by the IEEE 754 Floating-Point standard. The computed result differs from the both the exact result and the IEEE-mandated representation of the exact result by no more than one unit in the least-significant place (ulp).

It is implementation dependent whether the result is affected by the current rounding mode in FCSR.

Restrictions:

The fields *fs* and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.


The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR(*fd*, *fmt*, 1.0 / valueFPR(*fs*, *fmt*))

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Inexact Unimplemented Operation
 - Division-by-zero Invalid Operation
 - Overflow Underflow



Format: ROUND.L.S fd, fs **MIPS III**
 ROUND.L.D fd, fs

Purpose: To convert an FP value to 64-bit fixed-point, rounding to nearest.

Description: $fd \leftarrow \text{convert_and_round}(fs)$

The value in FPR *fs* in format *fmt*, is converted to a value in 64-bit long fixed-point format rounding to nearest/even (rounding mode 0). The result is placed in FPR *fd*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{63} to $2^{63}-1$, the result cannot be represented correctly and an IEEE Invalid Operation condition exists. The result depends on the FP exception model currently active.

- ◆ *Precise exception model: The Invalid Operation flag is set in the FCSR. If the Invalid Operation enable bit is set in the FCSR, no result is written to fd and an Invalid Operation exception is taken immediately. Otherwise, the default result, $2^{63}-1$, is written to fd.*
- ◆ *Imprecise exception model (R8000 normal mode): The default result, $2^{63}-1$, is written to fd. No FCSR flag is set. If the Invalid Operation enable bit is set in the FCSR, an Invalid Operation exception is taken, imprecisely, at some future time.*

Restrictions:

The fields *fs* and *fd* must specify valid FPRs; *fs* for type *fmt* and *fd* for long fixed-point. If they are not valid, the result is undefined.


The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR(fd, L, ConvertFmt(ValueFPR(fs, fmt), fmt, L))

Exceptions:

- | | |
|----------------------|-------------------------|
| Coprocessor Unusable | |
| Reserved Instruction | |
| Floating-Point | |
| Inexact | Unimplemented Operation |
| Overflow | Invalid Operation |



Format: ROUND.W.S fd, fs MIPS II
 ROUND.W.D fd, fs

Purpose: To convert an FP value to 32-bit fixed-point, rounding to nearest.

Description: $fd \leftarrow \text{convert_and_round}(fs)$

The value in FPR *fs* in format *fmt*, is converted to a value in 32-bit word fixed-point format rounding to nearest/even (rounding mode 0). The result is placed in FPR *fd*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{31} to $2^{31}-1$, the result cannot be represented correctly and an IEEE Invalid Operation condition exists. The result depends on the FP exception model currently active.

- ◆ *Precise exception model:* The Invalid Operation flag is set in the FCSR. If the Invalid Operation enable bit is set in the FCSR, no result is written to *fd* and an Invalid Operation exception is taken immediately. Otherwise, the default result, $2^{31}-1$, is written to *fd*.
- ◆ *Imprecise exception model (R8000 normal mode):* The default result, $2^{31}-1$, is written to *fd*. No FCSR flag is set. If the Invalid Operation enable bit is set in the FCSR, an Invalid Operation exception is taken, imprecisely, at some future time.

Restrictions:

The fields *fs* and *fd* must specify valid FPRs; *fs* for type *fmt* and *fd* for word fixed-point. If they are not valid, the result is undefined.


The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR(*fd*, W, ConvertFmt(ValueFPR(*fs*, *fmt*), *fmt*, W))

Exceptions:

- | | |
|----------------------|-------------------------|
| Coprocessor Unusable | |
| Reserved Instruction | |
| Floating-Point | |
| Inexact | Unimplemented Operation |
| Invalid Operation | Overflow |



Format: RSQRT.S fd, fs **MIPS IV**
 RSQRT.D fd, fs

Purpose: To approximate the reciprocal of the square root of an FP value (quickly).

Description: $fd \leftarrow 1.0 / \text{sqrt}(fs)$

The reciprocal of the positive square root of the value in FPR *fs* is approximated and placed into FPR *fd*. The operand and result are values in format *fmt*.

The numeric accuracy of this operation is implementation dependent; it does not meet the accuracy specified by the IEEE 754 Floating-Point standard. The computed result differs from the both the exact result and the IEEE-mandated representation of the exact result by no more than two units in the least-significant place (ulp).

It is implementation dependent whether the result is affected by the current rounding mode in FCSR.

Restrictions:

The fields *fs* and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.


The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR(fd, fmt, 1.0 / SquareRoot(valueFPR(fs, fmt)))

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Inexact Unimplemented Operation
 - Division-by-zero Invalid Operation
 - Overflow Underflow



Format: SDC1 ft, offset(base) MIPS II

Purpose: To store a doubleword from an FPR to memory.

Description: $\text{memory}[\text{base}+\text{offset}] \leftarrow \text{ft}$

The 64-bit doubleword in FPR *ft* is stored in memory at the location specified by the aligned effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

If coprocessor 1 general registers are 32-bits wide (a native 32-bit processor or 32-bit register emulation mode in a 64-bit processor), FPR *ft* is held in an even/odd register pair. The low word is taken from the even register *ft* and the high word is from *ft+1*.

Restrictions:

If *ft* does not specify an FPR that can contain a doubleword, the result is undefined.

An Address Error exception occurs if $\text{EffectiveAddress}_{2..0} \neq 0$ (not doubleword-aligned).

MIPS IV: The low-order 3 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr2..0 ≠ 03 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation(vAddr, DATA, STORE)
if SizeFGR() = 64 then /* 64-bit wide FGRs */
    data ← FGR[ft]
elseif ft0 = 0 then /* valid specifier, 32-bit wide FGRs */
    data ← FGR[ft+1] || FGR[ft]
else /* undefined for odd 32-bit FGRs */
    UndefinedResult()
endif
StoreMemory(uncached, DOUBLEWORD, data, pAddr, vAddr, DATA)

```

Exceptions:

- Coprocessor unusable
- Reserved Instruction
- TLB Refill, TLB Invalid
- TLB Modified
- Address Error



Format: SDXC1 fs, index(base) MIPS IV

Purpose: To store a doubleword from an FPR to memory (GPR+GPR addressing).

Description: memory[base+index] ← fs

The 64-bit doubleword in FPR *fs* is stored in memory at the location specified by the aligned effective address. The contents of GPR *index* and GPR *base* are added to form the effective address.

If coprocessor 1 general registers are 32-bits wide (a native 32-bit processor or 32-bit register emulation mode in a 64-bit processor), FPR *fs* is held in an even/odd register pair. The low word is taken from the even register *fs* and the high word is from *fs*+1.

Restrictions:

If *fs* does not specify an FPR that can contain a doubleword, the result is undefined.

The Region bits of the effective address must be supplied by the contents of *base*. If $\text{EffectiveAddress}_{63..62} \neq \text{base}_{63..62}$, the result is undefined.

An Address Error exception occurs if $\text{EffectiveAddress}_{2..0} \neq 0$ (not doubleword-aligned).

MIPS IV: The low-order 3 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.


Operation:

```

vAddr ← GPR[base] + GPR[index]
if vAddr2..0 ≠ 03 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation(vAddr, DATA, STORE)
if SizeFGR() = 64 then /* 64-bit wide FGRs */
    data ← FGR[fs]
elseif fs0 = 0 then /* valid specifier, 32-bit wide FGRs */
    data ← FGR[fs+1] || FGR[fs]
else /* undefined for odd 32-bit FGRs */
    UndefinedResult()
endif
StoreMemory(uncached, DOUBLEWORD, data, pAddr, vAddr, DATA)
    
```

Exceptions:

- TLB Refill, TLB Invalid
- TLB Modified
- Address Error
- Reserved Instruction
- Coprocessor Unusable



Format: SQRT.S *fd, fs* MIPS II
 SQRT.D *fd, fs*

Purpose: To compute the square root of an FP value.

Description: $fd \leftarrow \text{SQRT}(fs)$

The square root of the value in FPR *fs* is calculated to infinite precision, rounded according to the current rounding mode in FCSR, and placed into FPR *fd*. The operand and result are values in format *fmt*.

If the value in FPR *fs* corresponds to -0 , the result will be -0 .

Restrictions:

If the value in FPR *fs* is less than 0, an Invalid Operation condition is raised.

The fields *fs* and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.


The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR(*fd*, *fmt*, SquareRoot(ValueFPR(*fs*, *fmt*)))

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Unimplemented Operation
 - Invalid Operation
 - Inexact



Format: SUB.S fd, fs, ft **MIPS I**
 SUB.D fd, fs, ft

Purpose: To subtract FP values.

Description: $fd \leftarrow fs - ft$

The value in FPR *ft* is subtracted from the value in FPR *fs*. The result is calculated to infinite precision, rounded according to the current rounding mode in FCSR, and placed into FPR *fd*. The operands and result are values in format *fmt*.

Restrictions:

The fields *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is undefined.


The operands must be values in format *fmt*. If they are not, the result is undefined and the value of the operand FPRs becomes undefined.

Operation:

StoreFPR (fd, fmt, ValueFPR(fs, fmt) – ValueFPR(ft, fmt))

Exceptions:

- Coprocessor Unusable
- Reserved Instruction
- Floating-Point
 - Inexact Unimplemented Operation
 - Invalid Operation Overflow
 - Underflow



Format: SWC1 ft, offset(base) **MIPS I**

Purpose: To store a word from an FPR to memory.

Description: memory[base+offset] ← ft

The low 32-bit word from FPR *ft* is stored in memory at the location specified by the aligned effective address. The 16-bit signed *offset* is added to the contents of GPR *base* to form the effective address.

Restrictions:

An Address Error exception occurs if EffectiveAddress_{1..0} ≠ 0 (not word-aligned).

MIPS IV: The low-order 2 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.

Operation: 32-bit Processors


```
vAddr ← sign_extend(offset) + GPR[base]
if vAddr1..0 ≠ 02 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, STORE)
data ← FGR[ft]
StoreMemory (uncached, WORD, data, pAddr, vAddr, DATA)
```

Operation: 64-bit Processors

```
vAddr ← sign_extend(offset) + GPR[base]
if vAddr1..0 ≠ 02 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..3 // (pAddr2..0 xor (ReverseEndian || 02))
bytesel ← vAddr2..0 xor (BigEndianCPU || 02)
/* the bytes of the word are moved into the correct byte lanes */
if SizeFGR() = 64 then /* 64-bit wide FGRs */
    data ← 032-8*bytesel || FGR[ft]31..0 || 08*bytesel /* top or bottom wd of 64-bit data */
else /* 32-bit wide FGRs */
    data ← 032-8*bytesel || FGR[ft] || 08*bytesel /* top or bottom wd of 64-bit data */
endif
StoreMemory (uncached, WORD, data, pAddr, vAddr, DATA)
```

Exceptions:

- Coprocessor unusable
- Reserved Instruction
- TLB Refill, TLB Invalid
- TLB Modified
- Address Error



Format: SWXC1 fs, index(base) **MIPS IV**

Purpose: To store a word from an FPR to memory (GPR+GPR addressing).

Description: memory[base+index] ← fs

The low 32-bit word from FPR *fs* is stored in memory at the location specified by the aligned effective address. The contents of GPR *index* and GPR *base* are added to form the effective address.

Restrictions:

The Region bits of the effective address must be supplied by the contents of *base*. If $\text{EffectiveAddress}_{63..62} \neq \text{base}_{63..62}$, the result is undefined.

An Address Error exception occurs if $\text{EffectiveAddress}_{1..0} \neq 0$ (not word-aligned).

MIPS IV: The low-order 2 bits of the *offset* field must be zero. If they are not, the result of the instruction is undefined.


Operation:

```

vAddr ← GPR[base] + GPR[index]
if vAddr1..0 ≠ 02 then SignalException(AddressError) endif
(pAddr, uncached) ← AddressTranslation(vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..3 || (pAddr2..0 xor (ReverseEndian || 02))
bytesel ← vAddr2..0 xor (BigEndianCPU || 02)
/* the bytes of the word are moved into the correct byte lanes */
if SizeFGR() = 64 then /* 64-bit wide FGRs */
    data ← 032-8*bytesel || FGR[fs]31..0 || 08*bytesel /* top or bottom wd of 64-bit data */
else /* 32-bit wide FGRs */
    data ← 032-8*bytesel || FGR[fs] || 08*bytesel /* top or bottom wd of 64-bit data */
endif
StoreMemory(uncached, WORD, data, pAddr, vAddr, DATA)
    
```

Exceptions:

- TLB Refill, TLB Invalid
- TLB Modified
- Address Error
- Reserved Instruction
- Coprocessor Unusable



Format: TRUNC.L.S fd, fs **MIPS III**
 TRUNC.L.D fd, fs

Purpose: To convert an FP value to 64-bit fixed-point, rounding toward zero.

Description: fd ← convert_and_round(fs)

The value in FPR *fs* in format *fmt*, is converted to a value in 64-bit long fixed-point format rounding toward zero (rounding mode 1). The result is placed in FPR *fd*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{63} to $2^{63}-1$, the result cannot be represented correctly and an IEEE Invalid Operation condition exists. The result depends on the FP exception model currently active.

- ◆ *Precise exception model: The Invalid Operation flag is set in the FCSR. If the Invalid Operation enable bit is set in the FCSR, no result is written to fd and an Invalid Operation exception is taken immediately. Otherwise, the default result, $2^{63}-1$, is written to fd.*
- ◆ *Imprecise exception model (R8000 normal mode): The default result, $2^{63}-1$, is written to fd. No FCSR flag is set. If the Invalid Operation enable bit is set in the FCSR, an Invalid Operation exception is taken, imprecisely, at some future time.*

Restrictions:

The fields *fs* and *fd* must specify valid FPRs; *fs* for type *fmt* and *fd* for long fixed-point. If they are not valid, the result is undefined.


The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

StoreFPR(fd, L, ConvertFmt(ValueFPR(fs, fmt), fmt, L))

Exceptions:

- | | |
|----------------------|-------------------------|
| Coprocessor Unusable | |
| Reserved Instruction | |
| Floating-Point | |
| Inexact | Unimplemented Operation |
| Invalid Operation | Overflow |



Format: TRUNC.W.S fd, fs MIPS II
 TRUNC.W.D fd, fs

Purpose: To convert an FP value to 32-bit fixed-point, rounding toward zero.

Description: $fd \leftarrow \text{convert_and_round}(fs)$

The value in FPR *fs* in format *fmt*, is converted to a value in 32-bit word fixed-point format using rounding toward zero (rounding mode 1)). The result is placed in FPR *fd*.

When the source value is Infinity, NaN, or rounds to an integer outside the range -2^{31} to $2^{31}-1$, the result cannot be represented correctly and an IEEE Invalid Operation condition exists. The result depends on the FP exception model currently active.

- ◆ *Precise exception model:* The Invalid Operation flag is set in the FCSR. If the Invalid Operation enable bit is set in the FCSR, no result is written to *fd* and an Invalid Operation exception is taken immediately. Otherwise, the default result, $2^{31}-1$, is written to *fd*.
- ◆ *Imprecise exception model (R8000 normal mode):* The default result, $2^{31}-1$, is written to *fd*. No FCSR flag is set. If the Invalid Operation enable bit is set in the FCSR, an Invalid Operation exception is taken, imprecisely, at some future time.

Restrictions:

The fields *fs* and *fd* must specify valid FPRs; *fs* for type *fmt* and *fd* for word fixed-point. If they are not valid, the result is undefined.

The operand must be a value in format *fmt*. If it is not, the result is undefined and the value of the operand FPR becomes undefined.

Operation:

$\text{StoreFPR}(fd, W, \text{ConvertFmt}(\text{ValueFPR}(fs, \text{fmt}), \text{fmt}, W))$

Exceptions:

- | | |
|----------------------|-------------------------|
| Coprocessor Unusable | |
| Reserved Instruction | |
| Floating-Point | |
| Inexact | Invalid Operation |
| Overflow | Unimplemented Operation |




FPU Instructions Encoding


Notes

An FPU instruction is a single 32-bit aligned word. The distinct FP instruction layouts are shown below. Variable information is in lower-case labels, such as "offset". Upper-case labels and any numbers indicate constant data. A table follows all the layouts that explains the fields used in them. Note that the same field may have different names in different instruction layout pictures. The field name is mnemonic to the function of that field in the instruction layout. The opcode tables and the instruction decode discussion use the canonical field names: opcode, fmt, nd, tf, and function. The other fields are not used for instruction decode.


Immediate: load/store using register + offset addressing.




Register: 2-register and 3-register formatted arithmetic operations.




Register Immediate: data transfer -- CPU / FPU register.




Condition code, Immediate: conditional branches on FPU cc using PC + offset.




Register to Condition Code: formatted FP compare.




Condition Code, Register FP: FPU register move-conditional on FP cc.




Register-4: 4-register formatted arithmetic operations.




Register Index: Load/store using register + register addressing.



Register Index hint: Prefetch using register + register addressing.



Condition Code, Register Integer: CPU register move-conditional on FP cc.



BC	Branch Conditional instruction subcode (op=COP1)
base	CPU register: base address for address calculations
COP1	Coprocessor 1 primary opcode value in op field.
COP1X	Coprocessor 1 eXtended primary opcode value in op field.
cc	condition code specifier. For architecture levels prior to MIPS IV it must be zero.
fd	FPU register: destination (arithmetic, loads, move-to) or source (stores, move-from)
fmt	destination and/or operand type ("format") specifier
fr	FPU register: source
fs	FPU register: source
ft	FPU register: source (for stores, arithmetic) or destination (for loads)
function	function field specifying a function within a particular op operation code.
function: op4 + fmt3	op4 is a 3-bit function field specifying which 4-register arithmetic operation for COP1X, fmt3 is a 3-bit field specifying the format of the operands and destination. The combinations are shown as several distinct instructions in the opcode tables.
hint	hint field made available to cache controller for prefetch operation
index	CPU register, holds index address component for address calculations
MOVC	Value in function field for conditional move. There is one value for the instruction with op=COP1, another for the instruction with op=SPECIAL.
nd	nullify delay. If set, branch is Likely and delay slot instruction is not executed. This must be zero for MIPS I.
offset	signed offset field used in address calculations
op	primary operation code (COP1, COP1X, LWC1, SWC1, LDC1, SDC1, SPECIAL)
PREFX	Value in function field for prefetch instruction for op=COP1X
rd	CPU register: destination
rs	CPU register: source
rt	CPU register: source / destination
SPECIAL	SPECIAL primary opcode value in op field.
sub	Operation subcode field for COP1 register immediate mode instructions.
tf	true/false. The condition from FP compare is tested for equality with tf bit.

FPU (CP1) Instruction Opcode Bit Encoding

This section describes the encoding of the Floating-Point Unit (FPU) instructions for the four levels of the MIPS architecture, MIPS I through MIPS IV. Each architecture level includes the instructions in the previous level;¹ MIPS IV includes all instructions in MIPS I, MIPS II, and MIPS III. This section presents eight different views of the instruction encoding.

- ◆ *Separate encoding tables for each architecture level.*
- ◆ *A MIPS IV encoding table showing the architecture level at which each opcode was originally defined and subsequently modified (if modified).*
- ◆ *Separate encoding tables for each architecture revision showing the changes made during that revision.*

Instruction Decode

Instruction field names are printed in **bold** in this section.

¹ An exception to this rule is that the reserved, but never implemented, Coprocessor 3 instructions were removed or changed to another use starting in MIPS III.

The primary **opcode** field is decoded first. The **opcode** values LWC1, SWC1, LDC1, and SDC1 fully specify FPU load and store instructions. The **opcode** values *COP1*, *COP1X*, and *SPECIAL* specify instruction classes. Instructions within a class are further specified by values in other fields.

COP1 Instruction Class

The **opcode**=*COP1* instruction class encodes most of the FPU instructions. The class is further decoded by examining the **fmt** field. The **fmt** values fully specify the CPU ↔ FPU register move instructions and specify the *S*, *D*, *W*, *L*, and *BC* instruction classes.

The **opcode**=*COP1* + **fmt**=*BC* instruction class encodes the conditional branch instructions. The class is further decoded, and the instructions fully specified, by examining the **nd** and **tf** fields.

The **opcode**=*COP1* + **fmt**=(*S*, *D*, *W*, or *L*) instruction classes encode instructions that operate on formatted (typed) operands. Each of these instruction classes is further decoded by examining the **function** field. With one exception the **function** values fully specify instructions. The exception is the *MOVCF* instruction class.

The **opcode**=*COP1* + **fmt**=(*S* or *D*) + **function**=*MOVCF* instruction class encodes the *MOVT.fmt* and *MOVF.fmt* conditional move instructions (to move FP values based on FP condition codes). The class is further decoded, and the instructions fully specified, by examining the **tf** field.

COP1X Instruction Class

The **opcode**=*COP1X* instruction class encodes the indexed load/store instructions, the indexed prefetch, and the multiply accumulate instructions. The class is further decoded, and the instructions fully specified, by examining the **function** field.

SPECIAL Instruction Class

The **opcode**=*SPECIAL* instruction class is further decoded by examining the **function** field. The only **function** value that applies to FPU instruction encoding is the *MOVCI* instruction class. The remainder of the **function** values encode CPU instructions.


The **opcode**=*SPECIAL* + **function**=*MOVCI* instruction class encodes the *MOVT* and *MOVF* conditional move instructions (to move CPU registers based on FP condition codes). The class is further decoded, and the instructions fully specified, by examining the **tf** field.

Instruction Subsets of MIPS III and MIPS IV Processors


MIPS III processors, such as the RC4000, RC4200, RC4300, RC4400, and RC4600, have a processor mode in which only the MIPS II instructions are valid. The MIPS II encoding table describes the MIPS II-only mode.

MIPS IV processors, such as the R8000 and R10000, have processor modes in which only the MIPS II or MIPS III instructions are valid. The MIPS II encoding table describes the MIPS II-only mode. The MIPS III encoding table describes the MIPS III-only mode.

Instructions encoded by the **opcode** field.




opcode		bits 28..26							
bits		0	1	2	3	4	5	6	7
31..29		000	001	010	011	100	101	110	111
0	000	c							
1	001								
2	010	COP1 δ							
3	011								
4	100								
5	101								
6	110	LWC1							
7	111	SWC1							




fmt		bits 23..21							
bits		0	1	2	3	4	5	6	7
25..24		000	001	010	011	100	101	110	111
0	00	MFC1	*	CFC1	*	MTC1	*	CTC1	*
1	01	BC δ	*	*	*	*	*	*	*
2	10	S δ	D δ	*	*	W δ	*	*	*
3	11	*	*	*	*	*	*	*	*

Instructions encoded by the **tf** field when opcode=COP1 and fmt=BC.




tf		bit 16	
t	f	0	1
		BC1F	BC1T


Instructions encoded by the **function** field when opcode=*COP1* and fmt = *S*, *D*, or *W*



function	bits 2..0								
bits	0	1	2	3	4	5	6	7	
5..3	000	001	010	011	100	101	110	111	
0	000	ADD	SUB	MUL	DIV	*	ABS	MOV	NEG
1	001	*	*	*	*	*	*	*	
2	010	*	*	*	*	*	*	*	
3	011	*	*	*	*	*	*	*	
4	100	*	CVT.D	*	*	CVT.W	*	*	
5	101	*	*	*	*	*	*	*	
6	110	C.F α	C.UN α	C.EQ α	C.UEQ α	C.OLT α	C.ULT α	C.OLE α	C.ULE α
7	111	C.SF α	C.NGLE α	C.SEQ α	C.NGL α	C.LT α	C.NGE α	C.LE α	C.NGT α




function	bits 2..0								
bits	0	1	2	3	4	5	6	7	
5..3	000	001	010	011	100	101	110	111	
0	000	ADD	SUB	MUL	DIV	*	ABS	MOV	NEG
1	001	*	*	*	*	*	*	*	
2	010	*	*	*	*	*	*	*	
3	011	*	*	*	*	*	*	*	
4	100	CVT.S	*	*	*	CVT.W	*	*	
5	101	*	*	*	*	*	*	*	
6	110	C.F α	C.UN α	C.EQ α	C.UEQ α	C.OLT α	C.ULT α	C.OLE α	C.ULE α
7	111	C.SF α	C.NGLE α	C.SEQ α	C.NGL α	C.LT α	C.NGE α	C.LE α	C.NGT α




function	bits 2..0							
bits	0	1	2	3	4	5	6	7
5..3	000	001	010	011	100	101	110	111
0	000	*	*	*	*	*	*	*
1	001	*	*	*	*	*	*	*
2	010	*	*	*	*	*	*	*
3	011	*	*	*	*	*	*	*
4	100	CVT.S	CVT.D	*	*	*	*	*
5	101	*	*	*	*	*	*	*
6	110	*	*	*	*	*	*	*
7	111	*	*	*	*	*	*	*

Instructions encoded by the **opcode** field.




opcode		bits 28..26							
bits		0	1	2	3	4	5	6	7
31..29		000	001	010	011	100	101	110	111
0	000	c							
1	001								
2	010	COP1 δ							
3	011								
4	100								
5	101								
6	110	LWC1						LDC1	
7	111	SWC1						SDC1	




fmt		bits 23..21							
bits		0	1	2	3	4	5	6	7
25..24		000	001	010	011	100	101	110	111
0	00	MFC1	*	CFC1	*	MTC1	*	CTC1	*
1	01	BC δ	*	*	*	*	*	*	*
2	10	S δ	D δ	*	*	W δ	*	*	*
3	11	*	*	*	*	*	*	*	*

Instructions encoded by the **nd** and **tf** fields when opcode=COP1 and fmt=BC.




nd		tf	
bit 17		0	1
0		BC1F	BC1T
1		BC1FL	BC1TL

Instructions encoded by the **function** field when opcode=*COP1* and fmt = *S*, *D*, or *W*




function		bits 2..0							
bits	5..3	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000	ADD	SUB	MUL	DIV	SQRT	ABS	MOV	NEG
1	001	*	*	*	*	ROUND. W	TRUNC.W	CEIL.W	FLOOR.W
2	010	*	*	*	*	*	*	*	*
3	011	*	*	*	*	*	*	*	*
4	100	*	CVT.D	*	*	CVT.W	*	*	*
5	101	*	*	*	*	*	*	*	*
6	110	C.F α	C.UN α	C.EQ α	C.UEQ α	C.OLT α	C.ULT α	C.OLE α	C.ULE α
7	111	C.SF α	C.NGLE α	C.SEQ α	C.NGL α	C.LT α	C.NGE α	C.LE α	C.NGT α



function		bits 2..0							
bits	5..3	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000	ADD	SUB	MUL	DIV	SQRT	ABS	MOV	NEG
1	001	*	*	*	*	ROUND. W	TRUNC.W	CEIL.W	FLOOR.W
2	010	*	*	*	*	*	*	*	*
3	011	*	*	*	*	*	*	*	*
4	100	CVT.S	*	*	*	CVT.W	*	*	*
5	101	*	*	*	*	*	*	*	*
6	110	C.F α	C.UN α	C.EQ α	C.UEQ α	C.OLT α	C.ULT α	C.OLE α	C.ULE α
7	111	C.SF α	C.NGLE α	C.SEQ α	C.NGL α	C.LT α	C.NGE α	C.LE α	C.NGT α


encoding when
fmt = W



function bits 2..0


bits	0	1	2	3	4	5	6	7
5..3	000	001	010	011	100	101	110	111
0	000	*	*	*	*	*	*	*
1	001	*	*	*	*	*	*	*
2	010	*	*	*	*	*	*	*
3	011	*	*	*	*	*	*	*
4	100	CVT.S	CVT.D	*	*	*	*	*
5	101	*	*	*	*	*	*	*
6	110	*	*	*	*	*	*	*
7	111	*	*	*	*	*	*	*

Instructions encoded by the opcode field.



opcode bits 28..26


bits	0	1	2	3	4	5	6	7
31..29	000	001	010	011	100	101	110	111
0	000	COP1 δ						
1	001							
2	010							
3	011							
4	100	c						
5	101	LWC1						
6	110							
7	111	SWC1						
		LDC1						
		SDC1						




fmt bits 23..21

bits	0	1	2	3	4	5	6	7
25..24	000	001	010	011	100	101	110	111
0	00	MFC1	DMFC1	CFC1	*	MTC1	DMTC1	CTC1
1	01	BC δ	*	*	*	*	*	*
2	10	S δ	D δ	*	*	W δ	L δ	*
3	11	*	*	*	*	*	*	*


Instructions encoded by the **nd** and **tf** fields when opcode=*COP1* and fmt=*BC*.



Instructions encoded by the **function** field when opcode=*COP1* and fmt = *S, D, W, or L*




function		bits 2..0							
bits	5..3	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000	ADD	SUB	MUL	DIV	SQRT	ABS	MOV	NEG
1	001	ROUND.L	TRUNC.L	CEIL.L	FLOOR.L	ROUND.W	TRUNC.W	CEIL.W	FLOOR.W
2	010	*	*	*	*	*	*	*	*
3	011	*	*	*	*	*	*	*	*
4	100	*	CVT.D	*	*	CVT.W	CVT.L	*	*
5	101	*	*	*	*	*	*	*	*
6	110	C.F α	C.UN α	C.EQ α	C.UEQ α	C.OLT α	C.ULT α	C.OLE α	C.ULE α
7	111	C.SF α	C.NGLE α	C.SEQ α	C.NGL α	C.LT α	C.NGE α	C.LE α	C.NGT α



function		bits 2..0							
bits	5..3	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000	ADD	SUB	MUL	DIV	SQRT	ABS	MOV	NEG
1	001	ROUND.L	TRUNC.L	CEIL.L	FLOOR.L	ROUND.W	TRUNC.W	CEIL.W	FLOOR.W
2	010	*	*	*	*	*	*	*	*
3	011	*	*	*	*	*	*	*	*
4	100	CVT.S	*	*	*	CVT.W	CVT.L	*	*
5	101	*	*	*	*	*	*	*	*
6	110	C.F α	C.UN α	C.EQ α	C.UEQ α	C.OLT α	C.ULT α	C.OLE α	C.ULE α
7	111	C.SF α	C.NGLE α	C.SEQ α	C.NGL α	C.LT α	C.NGE α	C.LE α	C.NGT α


encoding when
fmt = W or L



function bits 2..0


bits	0	1	2	3	4	5	6	7
5..3	000	001	010	011	100	101	110	111
0	000	*	*	*	*	*	*	*
1	001	*	*	*	*	*	*	*
2	010	*	*	*	*	*	*	*
3	011	*	*	*	*	*	*	*
4	100	CVT.S	CVT.D	*	*	*	*	*
5	101	*	*	*	*	*	*	*
6	110	*	*	*	*	*	*	*
7	111	*	*	*	*	*	*	*

Instructions encoded by the opcode field.



opcode bits 28..26


bits	0	1	2	3	4	5	6	7
31..29	000	001	010	011	100	101	110	111
0	000	SPECIAL δ, β						
1	001							
2	010	COP1 δ		COP1X δ, λ			c	
3	011							
4	100							
5	101							
6	110	LWC1					LDC1	
7	111	SWC1					SDC1	




fmt bits 23..21

bits	0	1	2	3	4	5	6	7
25..24	000	001	010	011	100	101	110	111
0	00	MFC1	DMFC1	CFC1	*	MTC1	DMTC1	CTC1
1	01	BC δ	*	*	*	*	*	*
2	10	S δ	D δ	*	*	W δ	L δ	*
3	11	*	*	*	*	*	*	*


Instructions encoded by the **nd** and **tf** fields when opcode=*COP1* and fmt=*BC*.




Instructions encoded by the **function** field when opcode=*COP1* and fmt = *S, D, W, or L*



function		bits 2..0							
bits	5..3	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000	ADD	SUB	MUL	DIV	SQRT	ABS	MOV	NEG
1	001	ROUND.L	TRUNC.L	CEIL.L	FLOOR.L	ROUND.W	TRUNC.W	CEIL.W	FLOOR.W
2	010	*	MOVCF δ	MOVZ	MOVN	*	RECIP	RSQRT	
3	011	*	*	*	*	*	*	*	*
4	100	*	CVT.D	*	*	CVT.W	CVT.L	*	*
5	101	*	*	*	*	*	*	*	*
6	110	C.F α	C.UN α	C.EQ α	C.UEQ α	C.OLT α	C.ULT α	C.OLE α	C.ULE α
7	111	C.SF α	C.NGLE α	C.SEQ α	C.NGL α	C.LT α	C.NGE α	C.LE α	C.NGT α




function		bits 2..0							
bits	5..3	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000	ADD	SUB	MUL	DIV	SQRT	ABS	MOV	NEG
1	001	ROUND.L	TRUNC.L	CEIL.L	FLOOR.L	ROUND.W	TRUNC.W	CEIL.W	FLOOR.W
2	010	*	MOVCF δ	MOVZ	MOVN	*	RECIP	RSQRT	
3	011	*	*	*	*	*	*	*	*
4	100	CVT.S	*	*	*	CVT.W	CVT.L	*	*
5	101	*	*	*	*	*	*	*	*
6	110	C.F α	C.UN α	C.EQ α	C.UEQ α	C.OLT α	C.ULT α	C.OLE α	C.ULE α
7	111	C.SF α	C.NGLE α	C.SEQ α	C.NGL α	C.LT α	C.NGE α	C.LE α	C.NGT α





function	bits 2..0							
bits	0	1	2	3	4	5	6	7
5..3	000	001	010	011	100	101	110	111
0	000	*	*	*	*	*	*	*
1	001	*	*	*	*	*	*	*
2	010	*	*	*	*	*	*	*
3	011	*	*	*	*	*	*	*
4	100	CVT.S	CVT.D	*	*	*	*	*
5	101	*	*	*	*	*	*	*
6	110	*	*	*	*	*	*	*
7	111	*	*	*	*	*	*	*

Instructions encoded by the **function** field when opcode=COP1X.



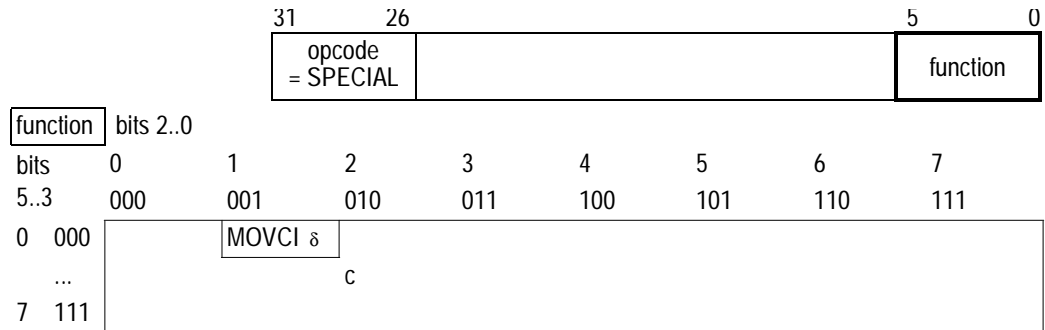
function	bits 2..0							
bits	0	1	2	3	4	5	6	7
5..3	000	001	010	011	100	101	110	111
0	000	LWXC1	LDXC1	*	*	*	*	*
1	001	SWXC1	SDXC1	*	*	*	*	PREFX
2	010	*	*	*	*	*	*	*
3	011	*	*	*	*	*	*	*
4	100	MADD.S	MADD.D	*	*	*	*	*
5	101	MSUB.S	MSUB.D	*	*	*	*	*
6	110	NMADD.S	NMADD.D	*	*	*	*	*
7	111	NMSUB.S	NMSUB.D	*	*	*	*	*

Instructions encoded by the **tf** field when opcode=COP1, fmt = S or D, and function=MOVCF.




These are the MOVf.fmt and MOVt.fmt instructions. They should not be confused with MOVf and MOVt.

Instruction class encoded by the **function** field when opcode=*SPECIAL*.

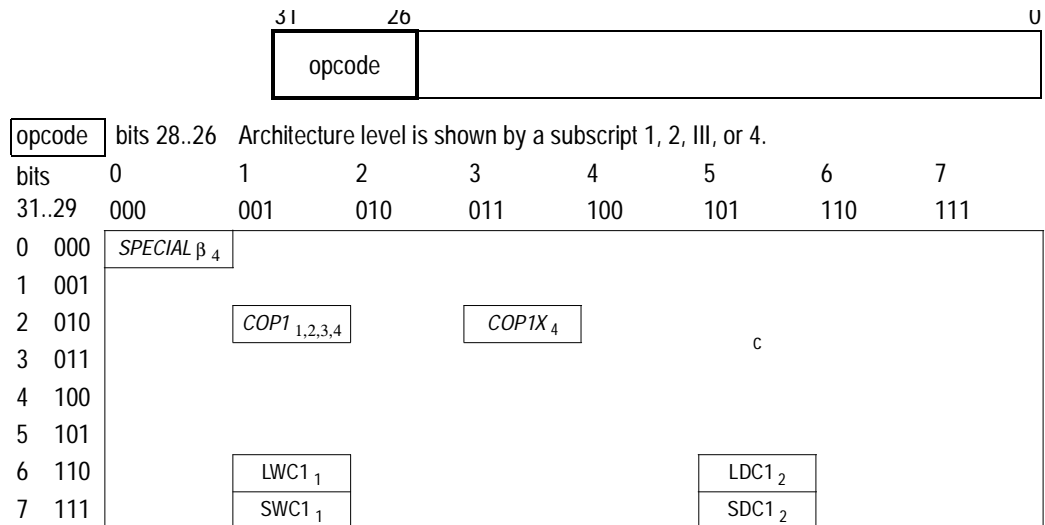



Instructions encoded by the **tf** field when opcode = *SPECIAL* and function=*MOVCI*.



The architecture level in which each MIPS IV encoding was defined is indicated by a subscript 1, 2, 3, or 4 (for architecture level I, II, III, or IV). If an instruction or instruction class was later extended, the extending level is indicated after the defining level.

Instructions encoded by the **opcode** field.






fmt bits 23..21 Architecture level is shown by a subscript 1, 2, 3, or 4.

bits	0	1	2	3	4	5	6	7
25..24	000	001	010	011	100	101	110	111
0 00	MFC1 ₁	DMFC1 ₃	CFC1 ₁	* ₁	MTC1 ₁	DMTC1 ₃	CTC1 ₁	* ₁
1 01	BC _{1,2,4}	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁
2 10	S _{1,2,3,4}	D _{1,2,3,4}	* ₁	* ₁	W _{1,2,3,4}	L _{3,4}	* ₁	* ₁
3 11	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁


Instructions encoded by the **nd** and **tf** fields when opcode=COP1 and fmt=BC.



tf bit 16 Architecture level is shown by a subscript 1, 2, 3, or 4.


n	0	1
d	BC1F _{1,4}	BC1T _{1,4}
bit 17	BC1FL _{2,4}	BC1TL _{2,4}

Instructions encoded by the **function** field when opcode=COP1 and fmt = S, D, W, or L




function bits 2..0 Architecture level is shown by a subscript 1, 2, 3, or 4.

bits	0	1	2	3	4	5	6	7
5..3	000	001	010	011	100	101	110	111
0 000	ADD ₁	SUB ₁	MUL ₁	DIV ₁	SQRT ₂	ABS ₁	MOV ₁	NEG ₁
1 001	ROUND.L ₃	TRUNC.L ₃	CEIL.L ₃	FLOOR.L ₃	ROUND.W ₂	TRUNC.W ₂	CEIL.W ₂	FLOOR.W ₂
2 010	* ₁	MOVCF ₄	MOVZ ₄	MOVN ₄	* ₁	RECIP ₄	RSQRT ₄	* ₁
3 011	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁
4 100	* ₁	CVT.D _{1,3}	* ₁	* ₁	CVT.W ₁	CVT.L ₃	* ₁	* ₁
5 101	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁
6 110	C.F _{1,4}	C.UN _{1,4}	C.EQ _{1,4}	C.UEQ _{1,4}	C.OLT _{1,4}	C.ULT _{1,4}	C.OLE _{1,4}	C.ULE _{1,4}
7 111	C.SF _{1,4}	C.NGLE _{1,4}	C.SEQ _{1,4}	C.NGL _{1,4}	C.LT _{1,4}	C.NGE _{1,4}	C.LE _{1,4}	C.NGT _{1,4}



function bits 2..0 Architecture level is shown by a subscript 1, 2, 3, or 4.

bits	0	1	2	3	4	5	6	7	
5..3	000	001	010	011	100	101	110	111	
0	000	ADD ₁	SUB ₁	MUL ₁	DIV ₁	SQRT ₂	ABS ₁	MOV ₁	NEG ₁
1	001	ROUND.L ₃	TRUNC.L ₃	CEIL.L ₃	FLOOR.L ₃	ROUND.W ₂	TRUNC.W ₂	CEIL.W ₂	FLOOR.W ₂
2	010	* ₁	MOVCF ₄	MOVZ ₄	MOVN ₄	* ₁	RECIP ₄	RSQRT ₄	* ₁
3	011	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁
4	100	CVT.S _{1,3}	* ₁	* ₁	* ₁	CVT.W ₁	CVT.L ₃	* ₁	* ₁
5	101	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁
6	110	C.F _{1,4}	C.UN _{1,4}	C.EQ _{1,4}	C.UEQ _{1,4}	C.OLT _{1,4}	C.ULT _{1,4}	C.OLE _{1,4}	C.ULE _{1,4}
7	111	C.SF _{1,4}	C.NGLE _{1,4}	C.SEQ _{1,4}	C.NGL _{1,4}	C.LT _{1,4}	C.NGE _{1,4}	C.LE _{1,4}	C.NGT _{1,4}



function bits 2..0 Architecture level is shown by a subscript 1, 2, 3, or 4.

bits	0	1	2	3	4	5	6	7
5..3	000	001	010	011	100	101	110	111
0	000	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁
1	001	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁
2	010	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁
3	011	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁
4	100	CVT.S _{1,3}	CVT.D _{1,3}	* ₁	* ₁	* ₁	* ₁	* ₁
5	101	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁
6	110	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁
7	111	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁	* ₁

Instructions encoded by the **function** field when opcode=*COP1X*.

		31	26					5	0
		opcode = COP1X						function	
function	bits 2..0	Architecture level is shown by a subscript 1, 2, 3, or 4.							
bits	0	1	2	3	4	5	6	7	
5..3	000	001	010	011	100	101	110	111	
0	000	LWXC1 ₄	LDXC1 ₄	* ₄	* ₄	* ₄	* ₄	* ₄	
1	001	SWXC1 ₄	SDXC1 ₄	* ₄	* ₄	* ₄	* ₄	PREFX ₄	
2	010	* ₄	* ₄	* ₄	* ₄	* ₄	* ₄	* ₄	
3	011	* ₄	* ₄	* ₄	* ₄	* ₄	* ₄	* ₄	
4	100	MADD.S ₄	MADD.D ₄	* ₄	* ₄	* ₄	* ₄	* ₄	
5	101	MSUB.S ₄	MSUB.D ₄	* ₄	* ₄	* ₄	* ₄	* ₄	
6	110	NMADD.S ₄	NMADD.D ₄	* ₄	* ₄	* ₄	* ₄	* ₄	
7	111	NMSUB.S ₄	NMSUB.D ₄	* ₄	* ₄	* ₄	* ₄	* ₄	

Instructions encoded by the **tf** field when opcode=*COP1*, fmt = *S* or *D*, and function=*MOVCF*.

		31	26 25		21	16			5	0
		opcode = COP1			fmt = S, D	t f				function = MOVCF
t	bit 16	0		1		These are the MOVF.fmt and MOVT.fmt instructions. They should not be confused with MOVF and MOVT.				
f		MOVCF (fmt) ₄		MOVT (fmt) ₄						

Instruction class encoded by the **function** field when opcode=*SPECIAL*.


		31	26					5	0
		opcode = SPECIAL						function	
function	bits 2..0	Architecture level is shown by a subscript 1, 2, 3, or 4.							
bits	0	1	2	3	4	5	6	7	
5..3	000	001	010	011	100	101	110	111	
0	000	MOVCI ₄							
...		c							
7	111								

Instructions encoded by the **tf** field when opcode = *SPECIAL* and function=*MOVCI*.

		31	26					16			5	0
		opcode = SPECIAL						t f				function = MOVCI
t	bit 16	0		1		These are the MOVF and MOVT instructions. They should not be confused with MOVF.fmt and MOVT.fmt.						
f		MOVF ₄		MOVT ₄								

An instruction encoding is shown if the instruction is added or extended in this architecture revision. An instruction class, like COP1, is shown if the instruction class is added in this architecture revision.


Instructions encoded by the **opcode** field.



opcode bits 28..26

bits	0	1	2	3	4	5	6	7
31..29	000	001	010	011	100	101	110	111

0	000							
1	001							
2	010							
3	011							
4	100							
5	101							
6	110						LDC1	
7	111						SDC1	




fmt bits 23..21

bits	0	1	2	3	4	5	6	7
25..24	000	001	010	011	100	101	110	111

0	00							
1	01							
2	10							
3	11							


Instructions encoded by the **nd** and **tf** fields when opcode=COP1 and fmt=BC.




tf bit 16

nd	0	1	
bit 17	1	BC1FL	BC1TL


Instructions encoded by the **function** field when opcode=*COP1* and fmt = *S*, *D*, or *W*



	function	bits 2..0							
	bits	0	1	2	3	4	5	6	7
	5..3	000	001	010	011	100	101	110	111
0	000					SQRT			
1	001					ROUND. W	TRUNC.W	CEIL.W	FLOOR.W
2	010								
3	011								
4	100								
5	101								
6	110								
7	111								




	function	bits 2..0							
	bits	0	1	2	3	4	5	6	7
	5..3	000	001	010	011	100	101	110	111
0	000					SQRT			
1	001					ROUND. W	TRUNC.W	CEIL.W	FLOOR.W
2	010								
3	011								
4	100								
5	101								
6	110								
7	111								



function		bits 2..0							
bits	0	1	2	3	4	5	6	7	U
5..3	000	001	010	011	100	101	110	111	
0	000								
1	001								
2	010								
3	011								
4	100								
5	101								
6	110								
7	111								


An instruction encoding is shown if the instruction is added or extended in this architecture revision. An instruction class, like COP1, is shown if the instruction class is added in this architecture revision.

Instructions encoded by the **opcode** field.




opcode		bits 28..26							
bits	0	1	2	3	4	5	6	7	U
31..29	000	001	010	011	100	101	110	111	
0	000								
1	001								
2	010								
3	011								
4	100								
5	101								
6	110								
7	111								

Instructions encoded by the **fmt** field when opcode=COP1.




fmt		bits 23..21							
bits	0	1	2	3	4	5	6	7	U
25..24	000	001	010	011	100	101	110	111	
0	00		DMFC1				DMTC1		
1	01								
2	10						L 8		
3	11								


Instructions encoded by the **nd** and **tf** fields when opcode=*COP1* and fmt=*BC*.



Instructions encoded by the **function** field when opcode=*COP1* and fmt = *S, D, or L*.




function		bits 2..0							
bits	5..3	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000								
1	001	ROUND.L	TRUNC.L	CEIL.L	FLOOR.L				
2	010								
3	011								
4	100						CVT.L		
5	101								
6	110								
7	111								



function		bits 2..0							
bits	5..3	0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	000								
1	001	ROUND.L	TRUNC.L	CEIL.L	FLOOR.L				
2	010								
3	011								
4	100						CVT.L		
5	101								
6	110								
7	111								

encoding when
fmt = L




function bits 2..0

bits	0	1	2	3	4	5	6	7
5..3	000	001	010	011	100	101	110	111
0	000	*	*	*	*	*	*	*
1	001	*	*	*	*	*	*	*
2	010	*	*	*	*	*	*	*
3	011	*	*	*	*	*	*	*
4	100	CVT.S	CVT.D	*	*	*	*	*
5	101	*	*	*	*	*	*	*
6	110	*	*	*	*	*	*	*
7	111	*	*	*	*	*	*	*

An instruction encoding is shown if the instruction is added or extended in this architecture revision. An instruction class, like COP1X, is shown if the instruction class is added in this architecture revision.


Instructions encoded by the opcode field.



opcode bits 28..26

bits	0	1	2	3	4	5	6	7
31..29	000	001	010	011	100	101	110	111
0	000	COP1X δ						
1	001							
2	010							
3	011							
4	100							
5	101							
6	110							
7	111							

Instructions encoded by the fmt field when opcode=COP1.



fmt bits 23..21

bits	0	1	2	3	4	5	6	7
25..24	000	001	010	011	100	101	110	111
0	00							
1	01							
2	10							
3	11							

Instructions encoded by the **nd** and **tf** fields when opcode=*COP1* and fmt=*BC*.






Table 6-15 (cont.)


Instructions encoded by the **function** field when opcode=*COP1* and fmt = *S, D, W, or L*.



function		bits 2..0							
bits	0	1	2	3	4	5	6	7	
5..3	000	001	010	011	100	101	110	111	
0	000								
1	001								
2	010	MOVCF δ	MOVZ	MOVN		RECIP	RSQRT		
3	011								
4	100								
5	101								
6	110	C.F	C.UN	C.EQ	C.UEQ	C.OLT	C.ULT	C.OLE	C.ULE
7	111	C.SF	C.NGLE	C.SEQ	C.NGL	C.LT	C.NGE	C.LE	C.NGT




function		bits 2..0							
bits	0	1	2	3	4	5	6	7	
5..3	000	001	010	011	100	101	110	111	
0	000								
1	001								
2	010	MOVCF δ	MOVZ	MOVN		RECIP	RSQRT		
3	011								
4	100								
5	101								
6	110	C.F	C.UN	C.EQ	C.UEQ	C.OLT	C.ULT	C.OLE	C.ULE
7	111	C.SF	C.NGLE	C.SEQ	C.NGL	C.LT	C.NGE	C.LE	C.NGT



function	bits 2..0							
bits	0	1	2	3	4	5	6	7
5..3	000	001	010	011	100	101	110	111
0	000							
1	001							
2	010							
3	011							
4	100							
5	101							
6	110							
7	111							



Table 6-15 (cont.) FPU Instruction Encoding Changes - MIPS IV Revision.

Instructions encoded by the **function** field when opcode=COP1X.




function	bits 2..0							
bits	0	1	2	3	4	5	6	7
5..3	000	001	010	011	100	101	110	111
0	000	LWXC1	LDXC1	*	*	*	*	*
1	001	SWXC1	SDXC1	*	*	*	*	PREFX
2	010	*	*	*	*	*	*	*
3	011	*	*	*	*	*	*	*
4	100	MADD.S	MADD.D	*	*	*	*	*
5	101	MSUB.S	MSUB.D	*	*	*	*	*
6	110	NMADD.S	NMADD.D	*	*	*	*	*
7	111	NMSUB.S	NMSUB.D	*	*	*	*	*

Instructions encoded by the **tf** field when opcode=COP1, fmt = S or D, and function=MOVCF.




These are the MOVf.fmt and MOVt.fmt instructions. They should not be confused with MOVf and MOVt.

Instruction class encoded by the **function** field when opcode=*SPECIAL*.



Instructions encoded by the **tf** field when opcode = *SPECIAL* and function=*MOVCI*.



Key to all FPU (CP1) instruction encoding tables:

- * This opcode is reserved for future use. An attempt to execute it causes either a Reserved Instruction exception or a Floating Point Unimplemented Operation Exception. The choice of exception is implementation specific.
- α The table shows 16 compare instructions with values named *C*.condition where "condition" is a comparison condition such as "EQ". These encoding values are all documented in the instruction description titled "C.cond.fmt".
- β The *SPECIAL* instruction class was defined in MIPS I for CPU instructions. An FPU instruction was first added to the instruction class in MIPS IV.
- δ (also *italic* opcode name) This opcode indicates an instruction class. The instruction word must be further decoded by examining additional tables that show values for another instruction field.
- λ The *COP1X* opcode in MIPS IV was the *COP3* opcode in MIPS I and II and a reserved instruction in MIPS III.
- χ These opcodes are not FPU operations. For further information on them, look in the CPU Instruction Encoding information Chapter 3.
- (*fmt*) This opcode is a conditional move of formatted FP registers - either MOVf.D, MOVf.S, MOVt.D, or MOVt.S. It should not be confused with the similarly-named MOVf or MOVt instruction that moves CPU registers.



Index

Notes

Numerics

64-bit RISController Family Primary Cache Indexing 2-25

A

About This Manual i

ABS.fmt 5-2

Access Functions for Floating-Point Registers 1-19

AccessLength Specifications for Loads/Stores 1-18

ADD 2-2

Add Immediate Unsigned Word 2-4

Add Immediate Word 2-3

Add Unsigned Word 2-5

Add Word 2-2

ADD.fmt 5-3

ADDI 2-3

ADDIU 2-4

ADDU 2-5

ALU Instructions With an Immediate Operand 1-5

AND 2-6

And 2-6

And Immediate 2-7

ANDI 2-7

Architecture Level in Which CPU Instructions are Defined or Extended 3-8

Arithmetic Instructions 4-2

Arithmetic Logic Unit 1-4

Assembler Format 1-13

Atomic Update CPU Load/Store Instructions 1-4

Atomic Update Loads and Stores 1-4

B

BC1F 5-4

BC1FL 5-6

BC1T 5-8

BC1TL 5-10

BEQ 2-8

BEQL 2-9

BGEZ 2-10

BGEZAL 2-11

BGEZALL 2-12

BGEZL 2-13

BGTZ 2-14

BGTZL 2-15

BLEZ 2-16

BLEZL 2-17

BLTZ 2-18, 2-21
 BLTZAL 2-19
 BLTZALL 2-20
 BNE 2-22
 BNEL 2-23
 Branch on Equal 2-8
 Branch on Equal Likely 2-9
 Branch on FP False 5-4
 Branch on FP False Likely 5-6
 Branch on FP True 5-8
 Branch on FP True Likely 5-10
 Branch on Greater Than or Equal to Zero 2-10
 Branch on Greater Than or Equal to Zero and Link 2-11
 Branch on Greater Than or Equal to Zero and Link Likely 2-12
 Branch on Greater Than or Equal to Zero Likely 2-13
 Branch on Greater Than Zero 2-14
 Branch on Greater Than Zero Likely 2-15
 Branch on Less Than or Equal to Zero 2-16
 Branch on Less Than or Equal to Zero Likely 2-17
 Branch on Less Than Zero 2-18
 Branch on Less Than Zero And Link 2-19
 Branch on Less Than Zero And Link Likely 2-20
 Branch on Less Than Zero Likely 2-21
 Branch on Not Equal 2-22
 Branch on Not Equal Likely 2-23
 BREAK 2-24
 Breakpoint 2-24
 Bytes Loaded by LDL Instruction 2-68, 2-78
 Bytes Loaded by LDR Instruction 2-70, 2-80
 Bytes Loaded by LWL Instruction 2-91
 Bytes Loaded by LWR Instruction 2-94
 Bytes Stored by SDL Instruction 2-127
 Bytes Stored by SDR Instruction 2-129
 Bytes Stored by SWL Instruction 2-147
 Bytes Stored by SWR Instruction 2-150
C
 C.cond.fmt 5-12
 CACHE 2-25
 Cache 2-25
 Cache Coherence Algorithms and Access Types 1-11
 CEIL.L.fmt 5-17
 CEIL.W.fmt 5-18
 CFC1 2-28, 5-19
 CLO 2-29
 CLZ 2-30
 Computational Instructions 1-4
 Conditional Branch Instructions 4-5
 Conditional Move Instructions 1-8

Conversion Instructions 4-3
 COP0 2-31
 Coprocessor 0 - COP0 3-2
 Coprocessor 1 - COP1, COP1X, MOVCI, and CP1 load/store 3-3
 Coprocessor 2 - COP2 and CP2 load/store 3-3
 Coprocessor 3 - COP3 and CP3 load/store 3-3
 Coprocessor Definition and Use in the MIPS Architecture 1-9
 Coprocessor General Register Access Functions 1-16
 Coprocessor Instructions 1-9
 Coprocessor Load and Store Instructions 1-4, 1-10
 Coprocessor Load/Store Instructions 1-4
 Coprocessor Operation 2-31
 Coprocessor Operation Instructions 1-10
 Coprocessor Operations 1-10
 Count Leading Ones 2-29
 Count Leading Zeros 2-30
 CPU Conditional Move Instructions 1-9
 CPU Conditional Move on FPU True/False Instructions 4-5
 CPU Functional Instruction Groups 1-1
 CPU Instruction Encoding 3-1, 3-2
 immediate 3-1
 jump 3-1
 register 3-1
 CPU Instruction Encoding - MIPS I Architecture 3-4
 CPU Instruction Encoding - MIPS II Architecture 3-5
 CPU Instruction Encoding - MIPS III Architecture 3-6
 CPU Instruction Encoding - MIPS IV Architecture 3-7
 CPU Instruction Encoding Changes - MIPS II Revision 3-9
 CPU Instruction Encoding Changes - MIPS III Revision 3-10
 CPU Instruction Encoding Changes - MIPS IV Revision 3-11
 CPU Instruction Formats 3-1
 CPU Instructions Basics 1-1
 CPU Loads and Stores 1-3
 CTC1 2-32, 5-20
 CVT.D.fmt 5-21
 CVT.L.fmt 5-22
 CVT.S.fmt 5-23
 CVT.W.fmt 5-24
D
 DADD 2-33
 DADDI 2-34
 DADDIU 2-35
 DADDU 2-36
 Data Transfer Instructions 4-1
 DDIV 2-37
 DDIVU 2-38
 Delayed Loads 1-3
 Description of an Instruction 4-6

DIV 2-39
 DIV.fmt 5-25
 Divide Unsigned Word 2-40
 Divide Word 2-39
 DIVU 2-40
 DMFC1 2-43, 5-26
 DMFCO 2-41
 DMTC0 2-42
 DMTC1 2-44, 5-27
 DMULT 2-45
 DMULTU 2-46
 Doubleword Add 2-33
 Doubleword Add Immediate 2-34
 Doubleword Add Immediate Unsigned 2-35
 Doubleword Add Unsigned 2-36
 Doubleword Divide 2-37
 Doubleword Divide Unsigned 2-38
 Doubleword Move From Floating-Point 2-43, 5-26
 Doubleword Move From System Control Coprocessor 2-41
 Doubleword Move To Floating-Point 2-44, 5-27
 Doubleword Move to System Control Coprocessor 2-42
 Doubleword Multiply 2-45
 Doubleword Multiply Unsigned 2-46
 Doubleword Shift Left Logical 2-47
 Doubleword Shift Left Logical Plus 32 2-48
 Doubleword Shift Left Logical Variable 2-49
 Doubleword Shift Right Arithmetic 2-50
 Doubleword Shift Right Arithmetic Plus 32 2-51
 Doubleword Shift Right Arithmetic Variable 2-52
 Doubleword Shift Right Logical 2-53
 Doubleword Shift Right Logical Plus 32 2-54
 Doubleword Shift Right Logical Variable 2-55
 Doubleword Subtract 2-56
 Doubleword Subtract Unsigned 2-57
 DSLL 2-47
 DSLL32 2-48
 DSLLV 2-49
 DSRA 2-50
 DSRA32 2-51
 DSRAV 2-52
 DSRL 2-53
 DSRL32 2-54
 DSRLV 2-55
 DSUB 2-56
 DSUBU 2-57
E
 ERET 2-58
 Error Exception Trap 2-58

Exception Instructions 1-8
 Exclusive OR 2-174
 Exclusive OR Immediate 2-175
F
 Floating-Point Absolute Value 5-2
 Floating-Point Add 5-3
 Floating-Point Ceiling Convert to Long Fixed-Point 5-17
 Floating-Point Ceiling Convert to Word Fixed-Point 5-18
 Floating-Point Compare 5-12
 Floating-Point Convert to Double Floating-Point 5-21
 Floating-Point Convert to Long Fixed-Point 5-22
 Floating-Point Convert to Single Floating-Point 5-23
 Floating-Point Convert to Word Fixed-Point 5-24
 Floating-Point Divide 5-25
 Floating-Point Floor Convert to Long Fixed-Point 5-28
 Floating-Point Floor Convert to Word Fixed-Point 5-29
 Floating-Point Move 5-36
 Floating-Point Move Conditional on FP False 5-38
 Floating-Point Move Conditional on FP True 5-41
 Floating-Point Move Conditional on Not Zero 5-39
 Floating-Point Move Conditional on Zero 5-42
 Floating-Point Multiply 5-45
 Floating-Point Multiply Add 5-34
 Floating-Point Multiply Subtract 5-43
 Floating-Point Negate 5-46
 Floating-Point Negative Multiply Add 5-47
 Floating-Point Negative Multiply Subtract 5-48
 Floating-Point Round to Long Fixed-Point 5-52
 Floating-Point Round to Word Fixed-Point 5-53
 Floating-Point Subtract 5-58
 Floating-Point Truncate to Word Fixed-Point 5-62
 Floating-Pt Truncate to Long Fixed-Pt 5-61
 FLOOR.L.fmt 5-28
 FLOOR.W.fmt 5-29
 Formatted Operand Value Move Instructions 4-4
 FPU Approximate Arithmetic Operations 4-3
 FPU Comparisons With Special Operand Exceptions for QNaNs 5-14
 FPU Comparisons Without Special Operand Exceptions 5-13
 FPU Conditional Branch Instructions 4-5
 FPU Conditional Move on True/False Instructions 4-4
 FPU Conditional Move on Zero/Nonzero Instruction 4-4
 FPU Conversion Operations Using a Directed Rounding Mode 4-4
 FPU Formatted Operand Move Instructions 4-4
 FPU IEEE Arithmetic Operations 4-3
 FPU Instruction Set Details 4-1
 FPU Instructions 4-1
 FPU Instructions Basics 4-1
 FPU Loads and Stores Using Register + Offset Address Mode 4-2

FPU Loads and Stores Using Register + Register Address Mode 4-2
 FPU Move To/From Instructions 4-2
 FPU Multiply-Accumulate Arithmetic Operations 4-3
 FPU Operand Format Field (fmt, fmt3) Decoding 4-5

I

Implementation-Specific Access Types 1-11
 Individual CPU Instruction Descriptions 1-21
 Individual FPU Instruction Descriptions 4-7
 Instruction Decode 3-2
 Instruction Description 1-13
 Instruction Descriptions 1-12
 Instruction Encoding Picture 1-12
 Instruction Exceptions 1-14
 Instruction Mnemonic and Name 1-12
 Instruction Operation 1-13
 Instruction Purpose 1-13
 Instruction Restrictions 1-13
 Instruction Subsets of MIPS III and MIPS IV Processors 3-2

J

J 2-59
 JAL 2-60
 JALR 2-61, 2-71
 JR 2-62, 2-72
 Jump 2-59
 Jump and Branch Instructions 1-6
 Jump And Link 2-60
 Jump And Link Register 2-61, 2-71
 Jump Instructions Jumping Within a 256 Megabyte Region 1-7
 Jump Instructions to Absolute Address 1-7
 Jump Register 2-62, 2-72

L

LB 2-63, 2-73
 LBU 2-64
 LD 2-65, 2-75
 LDC1 2-66, 2-76, 5-30
 LDL 2-67, 2-77
 LDR 2-69, 2-79
 LDXC1 5-31
 LH 2-81
 LHU 2-82
 LL 2-83
 LLD 2-85
 Load and Store Instructions 1-2
 Load and Store Memory Functions 1-17
 Load Byte 2-63
 Load Byte Unsigned 2-64, 2-74
 Load Doubleword 2-65, 2-75
 Load Doubleword Indexed to Floating-Point 5-31

Load Doubleword Left 2-67, 2-77
 Load Doubleword Right 2-69, 2-79
 Load Doubleword to Coprocessor 2-66, 2-76
 Load Doubleword to Floating-Point 5-30
 Load Halfword 2-81
 Load Halfword Unsigned 2-82
 Load Linked Doubleword 2-85
 Load Linked Word 2-83
 Load Upper Immediate 2-86
 Load Word 2-87
 Load Word Indexed to Floating-Point 5-33
 Load Word Left 2-90
 Load Word Right 2-93
 Load Word To Coprocessor 2-88
 Load Word to Floating-Point 5-32
 Load Word Unsigned 2-96
 Load/Store Operations Using Register + Offset Addressing Mode 1-2
 Load/Store Operations Using Register + Register Addressing Mode 1-2
 LUI 2-86
 LW 2-87
 LWC1 2-88, 5-32
 LWL 2-90
 LWR 2-93
 LWU 2-96
 LWXC1 5-33
M
 MAD 2-97
 MADD.fmt 5-34
 MADU 2-98
 Memory Access Types 1-10
 cached coherent 1-10
 cached noncoherent 1-10
 uncached 1-10
 MFC1 5-35
 MFCz 2-99
 MFHI 2-100
 MFLO 2-101
 MIPS Architecture Extensions 1-1
 Miscellaneous Functions 1-20
 Miscellaneous Instructions 1-8, 4-5
 Mixing References with Different Access Types 1-10
 MOV.fmt 5-36
 Move Conditional on FP False 5-37
 Move Conditional on FP True 5-40
 Move Conditional on Not Zero 2-102
 Move Conditional on Zero 2-103
 Move Control Word from Floating-Point 2-28, 5-19
 Move Control Word to Floating-Point 5-20

Move Control word to Floating-Point 2-32
 Move From Coprocessor 2-99
 Move From HI Register 2-100
 Move From LO Register 2-101
 Move To Coprocessor 2-106
 Move To HI Register 2-107
 Move To LO Register 2-108
 Move Word From Floating-Point 5-35
 Move Word to Floating-Point 5-44
 MOVF 5-37
 MOVF.fmt 5-38
 MOVN 2-102
 MOVN.fmt 5-39
 MOVT 5-40
 MOVT.fmt 5-41
 MOVZ 2-103
 MOVZ.fmt 5-42
 MSUB 2-104, 2-105
 MSUB.fmt 5-43
 MTC1 5-44
 MTCz 2-106
 MTHI 2-107
 MTLO 2-108
 MUL 2-109
 MUL.fmt 5-45
 MULT 2-110
 Multiply 2-109
 Multiply Accumulate 2-97
 Multiply and Divide Instructions 1-6
 Multiply Subtract 2-104
 Multiply Subtract Unsigned 2-105
 Multiply Unsigned Word 2-111
 Multiply Word 2-110
 Multiply/Add Unsigned 2-98
 Multiply/Divide Instructions 1-6
 MULTU 2-111
N
 NEG.fmt 5-46
 NMADD.fmt 5-47
 NMSUB.fmt 5-48
 Non-CPU Instructions 3-2
 NOR 2-112
 Normal CPU Load/Store Instructions 1-3
 Not Or 2-112
O
 Operand ALU Instructions 1-5
 Operation Notation Conventions and Functions 4-7
 Operation Section Notation and Functions 1-14

OR 2-113
 Or 2-113
 Or Immediate 2-114
 ORI 2-114
P
 PC-Relative Conditional Branch Instructions, Comparing 2 Registers 1-7
 PC-Relative Conditional Branch Instructions, Comparing Against Zero 1-7
 PFU Load/Store Instructions Using Register + Register Addressing 1-4
 PREF 2-115
 Prefetch 2-115
 Prefetch Indexed 5-49
 Prefetch Instructions 1-9
 Prefetch Using Register + Offset Address Mode 1-9
 Prefetch Using Register + Register Address Mode 1-9
 PREFIX 5-49
 Probe TLB for Matching Entry 2-168
 Programming and Implementation Notes 1-14
 Pseudocode Functions 1-16
 Pseudocode Language 1-14
 Pseudocode Symbols 1-14
R
 RECIP.fmt 5-51
 Reciprocal Approximation 5-51
 Reciprocal Square Root Approximation 5-54
 REGIMM Instruction Class 3-2
 Restore From Exception 2-117
 RFE 2-117
 ROUND.L.fmt 5-52
 ROUND.W.fmt 5-53
 RSQRT.fmt 5-54
S
 SC 2-119
 SCD 2-122
 SD 2-124
 SDC1 2-125, 5-55
 SDL 2-126
 SDR 2-128
 SDXC1 5-56
 Serialization Instructions 1-8
 Set On Less Than 2-133
 Set on Less Than Immediate 2-134
 Set on Less Than Immediate Unsigned 2-135
 Set on Less Than Unsigned 2-136
 SH 2-130
 Shift Instructions 1-5
 Shift Word Left Logical 2-131
 Shift Word Left Logical Variable 2-132
 Shift Word Right Arithmetic 2-137

Shift Word Right Arithmetic Variable 2-138
 Shift Word Right Logical 2-139
 Shift Word Right Logical Variable 2-140
 SLL 2-131
 SLLV 2-132
 SLT 2-133
 SLTI 2-134
 SLTIU 2-135
 SLTU 2-136
 SPECIAL Instruction Class 3-2
 SQRTfmt 5-57
 SRA 2-137
 SRAV 2-138
 SRL 2-139
 SRLV 2-140
 Store Conditional Doubleword 2-122
 Store Conditional Word 2-119
 Store Doubleword 2-124
 Store Doubleword From Coprocessor 2-125
 Store Doubleword from Floating-Point 5-55
 Store Doubleword Indexed from Floating-Point 5-56
 Store Doubleword Left 2-126
 Store Doubleword Right 2-128
 Store Halfword 2-130
 Store Word 2-143
 Store Word From Coprocessor 2-144
 Store Word Indexed from Floating-Point 5-60
 Store Word Left 2-146
 Store Word Right 2-149
 SUB 2-141
 SUB.fmt 5-58
 Subtract Unsigned Word 2-142
 Subtract Word 2-141
 SUBU 2-142
 Summary of Manual Contents i
 SW 2-143
 SWC1 2-144
 SWL 2-146
 SWR 2-149
 SWXC1 5-60
 Symbols in Instruction Operation Statements 1-14
 SYNC 2-152
 Synchronize Shared Memory 2-152
 SYSCALL 2-155
 System Call 2-155
 System Call and Breakpoint Instructions 1-8
T
 TEQ 2-156

TEQI 2-157
 TGE 2-158
 TGEI 2-159
 TGEIU 2-160
 TGEU 2-161
 TLBP 2-168
 TLBWR 2-171
 TLT 2-162
 TLTl 2-163
 TLTlU 2-164
 TLTU 2-165
 TNE 2-166
 TNEI 2-167
 Trap if Equal 2-156
 Trap if Equal Immediate 2-157
 Trap if Greater or Equal 2-158
 Trap if Greater or Equal Immediate 2-159
 Trap If Greater Or Equal Immediate Unsigned 2-160
 Trap If Greater or Equal Unsigned 2-161
 Trap if Less Than 2-162
 Trap if Less Than Immediate 2-163
 Trap if Less Than Immediate Unsigned 2-164
 Trap if Less Than Unsigned 2-165
 Trap if Not Equal 2-166
 Trap if Not Equal Immediate 2-167
 Trap-on-Condition Instructions, Comparing an Immediate 1-8
 Trap-on-Condition Instructions, Comparing Two Registers 1-8
 TRUNC.L.fmt 5-61
 TRUNC.W.fmt 5-62

U

Unaligned CPU Load/Store Instructions 1-3
 Unaligned Doubleword Load using LDL and LDR 2-67, 2-77
 Unaligned Doubleword Load using LDR and LDL 2-69, 2-79
 Unaligned Doubleword Store with SDL and SDR 2-126
 Unaligned Doubleword Store with SDR and SDL 2-128
 Unaligned Word Load using LWL and LWR 2-90
 Unaligned Word Load using LWR and LWL 2-93
 Unaligned Word Store using SWL and SWR 2-146
 Unaligned Word Store using SWR and SWL 2-149

V

Valid Formats for FPU Operations 4-6
 Valid Operands for FP Instructions 4-5
 Values of Hint Field for Prefetch Instruction in RC32364 2-115
 Values of Hint Field for Prefetch Instruction in RC5000 2-116

W

WAIT 2-173
 Wait 2-173
 Write Random TLB Entry 2-171

X
XOR 2-174
XORI 2-175

IMPORTANT NOTICE AND DISCLAIMER

RENESAS ELECTRONICS CORPORATION AND ITS SUBSIDIARIES (“RENESAS”) PROVIDES TECHNICAL SPECIFICATIONS AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES “AS IS” AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT OF THIRD-PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for developers who are designing with Renesas products. You are solely responsible for (1) selecting the appropriate products for your application, (2) designing, validating, and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, or other requirements. These resources are subject to change without notice. Renesas grants you permission to use these resources only to develop an application that uses Renesas products. Other reproduction or use of these resources is strictly prohibited. No license is granted to any other Renesas intellectual property or to any third-party intellectual property. Renesas disclaims responsibility for, and you will fully indemnify Renesas and its representatives against, any claims, damages, costs, losses, or liabilities arising from your use of these resources. Renesas' products are provided only subject to Renesas' Terms and Conditions of Sale or other applicable terms agreed to in writing. No use of any Renesas resources expands or otherwise alters any applicable warranties or warranty disclaimers for these products.

(Disclaimer Rev.1.01)

Corporate Headquarters

TOYOSU FORESIA, 3-2-24 Toyosu,
Koto-ku, Tokyo 135-0061, Japan
www.renesas.com

Trademarks

Renesas and the Renesas logo are trademarks of Renesas Electronics Corporation. All trademarks and registered trademarks are the property of their respective owners.

Contact Information

For further information on a product, technology, the most up-to-date version of a document, or your nearest sales office, please visit www.renesas.com/contact-us/.