



MIPS® Architecture For Programmers Volume II-A: The MIPS32® Instruction Set

Document Number: MD00086

Revision 3.00

March 25, 2010

**MIPS Technologies, Inc.
955 East Arques Avenue
Sunnyvale, CA 94085-4521**

Copyright © 2001-2003,2005,2008-2010 MIPS Technologies Inc. All rights reserved.

Copyright © 2001-2003,2005,2008-2010 MIPS Technologies, Inc. All rights reserved.

Unpublished rights (if any) reserved under the copyright laws of the United States of America and other countries.

This document contains information that is proprietary to MIPS Technologies, Inc. ("MIPS Technologies"). Any copying, reproducing, modifying or use of this information (in whole or in part) that is not expressly permitted in writing by MIPS Technologies or an authorized third party is strictly prohibited. At a minimum, this information is protected under unfair competition and copyright laws. Violations thereof may result in criminal penalties and fines.

Any document provided in source format (i.e., in a modifiable form such as in FrameMaker or Microsoft Word format) is subject to use and distribution restrictions that are independent of and supplemental to any and all confidentiality restrictions. UNDER NO CIRCUMSTANCES MAY A DOCUMENT PROVIDED IN SOURCE FORMAT BE DISTRIBUTED TO A THIRD PARTY IN SOURCE FORMAT WITHOUT THE EXPRESS WRITTEN PERMISSION OF MIPS TECHNOLOGIES, INC.

MIPS Technologies reserves the right to change the information contained in this document to improve function, design or otherwise. MIPS Technologies does not assume any liability arising out of the application or use of this information, or of any error or omission in such information. Any warranties, whether express, statutory, implied or otherwise, including but not limited to the implied warranties of merchantability or fitness for a particular purpose, are excluded. Except as expressly provided in any written license agreement from MIPS Technologies or an authorized third party, the furnishing of this document does not give recipient any license to any intellectual property rights, including any patent rights, that cover the information in this document.

The information contained in this document shall not be exported, reexported, transferred, or released, directly or indirectly, in violation of the law of any country or international law, regulation, treaty, Executive Order, statute, amendments or supplements thereto. Should a conflict arise regarding the export, reexport, transfer, or release of the information contained in this document, the laws of the United States of America shall be the governing law.

The information contained in this document constitutes one or more of the following: commercial computer software, commercial computer software documentation or other commercial items. If the user of this information, or any related documentation of any kind, including related technical data or manuals, is an agency, department, or other entity of the United States government ("Government"), the use, duplication, reproduction, release, modification, disclosure, or transfer of this information, or any related documentation of any kind, is restricted in accordance with Federal Acquisition Regulation 12.212 for civilian agencies and Defense Federal Acquisition Regulation Supplement 227.7202 for military agencies. The use of this information by the Government is further restricted in accordance with the terms of the license agreement(s) and/or applicable contract terms and conditions covering this information from MIPS Technologies or an authorized third party.

MIPS, MIPS I, MIPS II, MIPS III, MIPS IV, MIPS V, MIPS-3D, MIPS16, MIPS16e, MIPS32, MIPS64, MIPS-Based, MIPSsim, MIPSpro, MIPS Technologies logo, MIPS-VERIFIED, MIPS-VERIFIED logo, 4K, 4Kc, 4Km, 4Kp, 4KE, 4KEc, 4KEm, 4KEp, 4KS, 4KSc, 4KSd, M4K, M14K, 5K, 5Kc, 5Kf, 24K, 24Kc, 24Kf, 24KE, 24KEc, 24KEf, 34K, 34Kc, 34Kf, 74K, 74Kc, 74Kf, 1004K, 1004Kc, 1004Kf, R3000, R4000, R5000, ASMACRO, Atlas, "At the core of the user experience.", BusBridge, Bus Navigator, CLAM, CorExtend, CoreFPGA, CoreLV, EC, FPGA View, FS2, FS2 FIRST SILICON SOLUTIONS logo, FS2 NAVIGATOR, HyperDebug, HyperJTAG, JALGO, Logic Navigator, Malta, MDMX, MED, MGB, microMIPS, OCI, PDtrace, the Pipeline, Pro Series, SEAD, SEAD-2, SmartMIPS, SOC-it, System Navigator, and YAMON are trademarks or registered trademarks of MIPS Technologies, Inc. in the United States and other countries.

All other trademarks referred to herein are the property of their respective owners.

Template: nB1.03, Built with tags: 2B ARCH FPU_PS FPU_PSandARCH MIPS32

MIPS® Architecture For Programmers Volume II-A: The MIPS32® Instruction Set, Revision 3.00

Copyright © 2001-2003,2005,2008-2010 MIPS Technologies Inc. All rights reserved.

Contents

Chapter 1: About This Book	11
1.1: Typographical Conventions	11
1.1.1: Italic Text	11
1.1.2: Bold Text	12
1.1.3: Courier Text	12
1.2: UNPREDICTABLE and UNDEFINED	12
1.2.1: UNPREDICTABLE	12
1.2.2: UNDEFINED	13
1.2.3: UNSTABLE	13
1.3: Special Symbols in Pseudocode Notation	13
1.4: For More Information	16
Chapter 2: Guide to the Instruction Set	17
2.1: Understanding the Instruction Fields	17
2.1.1: Instruction Fields	19
2.1.2: Instruction Descriptive Name and Mnemonic	19
2.1.3: Format Field	19
2.1.4: Purpose Field	20
2.1.5: Description Field	20
2.1.6: Restrictions Field	20
2.1.7: Operation Field	21
2.1.8: Exceptions Field	21
2.1.9: Programming Notes and Implementation Notes Fields	22
2.2: Operation Section Notation and Functions	22
2.2.1: Instruction Execution Ordering	22
2.2.2: Pseudocode Functions	22
2.3: Op and Function Subfield Notation	31
2.4: FPU Instructions	31
Chapter 3: The MIPS32® Instruction Set	33
3.1: Compliance and Subsetting	33
3.2: Alphabetical List of Instructions	34
ABS.fmt	44
ADD	45
ADD.fmt	46
ADDI	47
ADDIU	48
ADDU	49
ALNV.PS	50
AND	52
ANDI	53
B	54
BAL	55
BC1F	56
BC1FL	58
BC1T	60
BC1TL	62

BC2F	64
BC2FL	65
BC2T	66
BC2TL	67
BEQ	68
BEQL	69
BGEZ	70
BGEZAL	71
BGEZALL	72
BGEZL	74
BGTZ	75
BGTZL	76
BLEZ	77
BLEZL	78
BLTZ	79
BLTZAL	80
BLTZALL	81
BLTZL	83
BNE	84
BNEL	85
BREAK	86
C.cond.fmt	87
CACHE	91
CEIL.L.fmt	97
CEIL.W.fmt	98
CFC1	99
CFC2	100
CLO	101
CLZ	102
COP2	103
CTC1	104
CTC2	106
CVT.D.fmt	107
CVT.L.fmt	108
CVT.PS.S	109
CVT.S.fmt	110
CVT.S.PL	111
CVT.S.PU	112
CVT.W.fmt	113
DERET	114
DI	115
DIV	116
DIV.fmt	118
DIVU	119
EHB	120
EI	121
ERET	122
EXT	123
FLOOR.L.fmt	125
FLOOR.W.fmt	126
INS	127
J	129
JAL	130

JALR.....	131
JALR.HB.....	133
JALX.....	137
JR.....	138
JR.HB.....	140
LB.....	142
LBU.....	143
LDC1.....	144
LDC2.....	145
LDXC1.....	146
LH.....	147
LHU.....	148
LL.....	149
LUI.....	150
LUXC1.....	151
LW.....	152
LWC1.....	153
LWC2.....	154
LWL.....	155
LWR.....	157
LWXC1.....	159
MADD.....	160
MADD.fmt.....	161
MADDDU.....	162
MFC0.....	163
MFC1.....	164
MFC2.....	165
MFHC1.....	166
MFHC2.....	167
MFHI.....	168
MFLO.....	169
MOV.fmt.....	170
MOVF.....	171
MOVF.fmt.....	172
MOVN.....	173
MOVN.fmt.....	174
MOVT.....	175
MOVT.fmt.....	176
MOVZ.....	177
MOVZ.fmt.....	178
MSUB.....	179
MSUB.fmt.....	180
MSUBU.....	181
MTC0.....	182
MTC1.....	183
MTC2.....	184
MTHC1.....	185
MTHC2.....	186
MTHI.....	187
MTLO.....	188
MUL.....	189
MUL.fmt.....	190
MULT.....	191

MULTU	192
NEG.fmt.....	193
NMADD.fmt	194
NMSUB.fmt.....	195
NOP	196
NOR.....	197
OR	198
ORI	199
PAUSE	201
PLL.PS	203
PLU.PS.....	204
PREF	205
PREFX.....	208
PUL.PS	209
PUU.PS	210
RDHWR.....	211
RDPGPR	213
RECIP.fmt.....	214
ROTR	215
ROTRV	216
ROUND.L.fmt	217
ROUND.W.fmt	218
RSQRT.fmt	219
SB	220
SC.....	221
SDBBP	223
SDC1	224
SDC2	225
SDXC1.....	226
SEB	227
SEH	228
SH.....	229
SLL	230
SLLV	231
SLT	232
SLTI	233
SLTIU	234
SLTU	235
SQRT.fmt.....	236
SRA	237
SRAV	238
SRL.....	239
SRLV	240
SSNOP	241
SUB	242
SUB.fmt	243
SUBU.....	244
SUXC1.....	245
SW	246
SWC1	247
SWC2	248
SWL.....	249
SWR	251

SWXC1	253
SYNC	254
SYNCL	259
SYSCALL	261
TEQ	262
TEQL	263
TGE	264
TGEI	265
TGEIU	266
TGEU	267
TLBP	268
TLBR	269
TLBWI	271
TLBWR	273
TLT	274
TLTI	275
TLTIU	276
TLTU	277
TNE	278
TNEI	279
TRUNC.L.fmt	280
TRUNC.W.fmt	281
WAIT	282
WRPGPR	283
WSBH	284
XOR	285
XORI	286
Appendix A: Instruction Bit Encodings	287
A.1: Instruction Encodings and Instruction Classes	287
A.2: Instruction Bit Encoding Tables	287
A.3: Floating Point Unit Instruction Format Encodings	295
Appendix B: Revision History	297

Figures

Figure 2.1: Example of Instruction Description	18
Figure 2.2: Example of Instruction Fields	19
Figure 2.3: Example of Instruction Descriptive Name and Mnemonic	19
Figure 2.4: Example of Instruction Format	19
Figure 2.5: Example of Instruction Purpose	20
Figure 2.6: Example of Instruction Description	20
Figure 2.7: Example of Instruction Restrictions	21
Figure 2.8: Example of Instruction Operation	21
Figure 2.9: Example of Instruction Exception	21
Figure 2.10: Example of Instruction Programming Notes	22
Figure 2.11: COP_LW Pseudocode Function	23
Figure 2.12: COP_LD Pseudocode Function	23
Figure 2.13: COP_SW Pseudocode Function	23
Figure 2.14: COP_SD Pseudocode Function	24
Figure 2.15: CoprocessorOperation Pseudocode Function	24
Figure 2.16: AddressTranslation Pseudocode Function	24
Figure 2.17: LoadMemory Pseudocode Function	25
Figure 2.18: StoreMemory Pseudocode Function	25
Figure 2.19: Prefetch Pseudocode Function	26
Figure 2.20: SyncOperation Pseudocode Function	27
Figure 2.21: ValueFPR Pseudocode Function	27
Figure 2.22: StoreFPR Pseudocode Function	28
Figure 2.23: CheckFPEException Pseudocode Function	29
Figure 2.24: FPConditionCode Pseudocode Function	29
Figure 2.25: SetFPConditionCode Pseudocode Function	29
Figure 2.26: SignalException Pseudocode Function	30
Figure 2.27: SignalDebugBreakpointException Pseudocode Function	30
Figure 2.28: SignalDebugModeBreakpointException Pseudocode Function	30
Figure 2.29: NullifyCurrentInstruction PseudoCode Function	31
Figure 2.30: JumpDelaySlot Pseudocode Function	31
Figure 2.31: PolyMult Pseudocode Function	31
Figure 3.1: Example of an ALNV.PS Operation	50
Figure 3.2: Usage of Address Fields to Select Index and Way	91
Figure 3.3: Operation of the EXT Instruction	123
Figure 3.4: Operation of the INS Instruction	127
Figure 3.5: Unaligned Word Load Using LWL and LWR	155
Figure 3.6: Bytes Loaded by LWL Instruction	156
Figure 3.7: Unaligned Word Load Using LWL and LWR	157
Figure 3.8: Bytes Loaded by LWR Instruction	158
Figure 3.9: Unaligned Word Store Using SWL and SWR	249
Figure 3.10: Bytes Stored by an SWL Instruction	250
Figure 3.11: Unaligned Word Store Using SWR and SWL	251
Figure 3.12: Bytes Stored by SWR Instruction	252
Figure A.1: Sample Bit Encoding Table	288

Tables

Table 1.1: Symbols Used in Instruction Operation Statements.....	13
Table 2.1: AccessLength Specifications for Loads/Stores	26
Table 3.1: CPU Arithmetic Instructions	34
Table 3.2: CPU Branch and Jump Instructions	35
Table 3.3: CPU Instruction Control Instructions	36
Table 3.4: CPU Load, Store, and Memory Control Instructions	36
Table 3.5: CPU Logical Instructions.....	36
Table 3.6: CPU Insert/Extract Instructions	37
Table 3.7: CPU Move Instructions	37
Table 3.8: CPU Shift Instructions	37
Table 3.9: CPU Trap Instructions.....	38
Table 3.10: Obsolete CPU Branch Instructions	38
Table 3.11: FPU Arithmetic Instructions.....	39
Table 3.12: FPU Branch Instructions	39
Table 3.13: FPU Compare Instructions.....	39
Table 3.14: FPU Convert Instructions	40
Table 3.15: FPU Load, Store, and Memory Control Instructions	40
Table 3.16: FPU Move Instructions.....	41
Table 3.17: Obsolete FPU Branch Instructions.....	41
Table 3.18: Coprocessor Branch Instructions	41
Table 3.19: Coprocessor Execute Instructions	42
Table 3.20: Coprocessor Load and Store Instructions.....	42
Table 3.21: Coprocessor Move Instructions.....	42
Table 3.22: Obsolete Coprocessor Branch Instructions.....	42
Table 3.23: Privileged Instructions	42
Table 3.24: EJTAG Instructions	43
Table 3.25: FPU Comparisons Without Special Operand Exceptions	88
Table 3.26: FPU Comparisons With Special Operand Exceptions for QNaNs	89
Table 3.27: Usage of Effective Address.....	91
Table 3.28: Encoding of Bits[17:16] of CACHE Instruction.....	92
Table 3.29: Encoding of Bits [20:18] of the CACHE Instruction	93
Table 3.30: Values of <i>hint</i> Field for PREF Instruction	205
Table 3.31: RDHWR Register Numbers	211
Table 3.32: Encodings of the Bits[10:6] of the SYNC instruction; the SType Field.....	256
Table A.1: Symbols Used in the Instruction Encoding Tables	288
Table A.2: MIPS32 Encoding of the Opcode Field	289
Table A.3: MIPS32 SPECIAL Opcode Encoding of Function Field.....	290
Table A.4: MIPS32 <i>REGIMM</i> Encoding of <i>rt</i> Field	290
Table A.5: MIPS32 SPECIAL2 Encoding of Function Field	290
Table A.6: MIPS32 <i>SPECIAL3</i> Encoding of Function Field for Release 2 of the Architecture.....	291
Table A.7: MIPS32 <i>MOVCI</i> Encoding of <i>tf</i> Bit	291
Table A.8: MIPS32 <i>SRL</i> Encoding of Shift/Rotate	291
Table A.9: MIPS32 <i>SRLV</i> Encoding of Shift/Rotate.....	291
Table A.10: MIPS32 <i>BSHFL</i> Encoding of <i>sa</i> Field.....	292
Table A.11: MIPS32 <i>COP0</i> Encoding of <i>rs</i> Field	292
Table A.12: MIPS32 <i>COP0</i> Encoding of Function Field When <i>rs=CO</i>	292
Table A.13: MIPS32 <i>COP1</i> Encoding of <i>rs</i> Field	293

Table A.14: MIPS32 <i>COP1</i> Encoding of Function Field When <i>rs=S</i>	293
Table A.15: MIPS32 <i>COP1</i> Encoding of Function Field When <i>rs=D</i>	293
Table A.16: MIPS32 <i>COP1</i> Encoding of Function Field When <i>rs=W</i> or <i>L</i>	294
Table A.17: MIPS64 <i>COP1</i> Encoding of Function Field When <i>rs=PS</i>	294
Table A.18: MIPS32 <i>COP1</i> Encoding of <i>tf</i> Bit When <i>rs=S, D, or PS</i> , Function= <i>MOVCF</i>	294
Table A.19: MIPS32 <i>COP2</i> Encoding of <i>rs</i> Field	295
Table A.20: MIPS64 <i>COP1X</i> Encoding of Function Field	295
Table A.21: Floating Point Unit Instruction Format Encodings.....	295

About This Book

The MIPS® Architecture For Programmers Volume II-A: The MIPS32® Instruction Set comes as part of a multi-volume set.

- Volume I-A describes conventions used throughout the document set, and provides an introduction to the MIPS32® Architecture
- Volume I-B describes conventions used throughout the document set, and provides an introduction to the microMIPS32™ Architecture
- Volume II-A provides detailed descriptions of each instruction in the MIPS32® instruction set
- Volume II-B provides detailed descriptions of each instruction in the microMIPS32™ instruction set
- Volume III describes the MIPS32® and microMIPS32™ Privileged Resource Architecture which defines and governs the behavior of the privileged resources included in a MIPS® processor implementation
- Volume IV-a describes the MIPS16e™ Application-Specific Extension to the MIPS32® Architecture. Beginning with Release 3 of the Architecture, microMIPS is the preferred solution for smaller code size.
- Volume IV-b describes the MDMX™ Application-Specific Extension to the MIPS64® Architecture and microMIPS64™. It is not applicable to the MIPS32® document set nor the microMIPS32™ document set
- Volume IV-c describes the MIPS-3D® Application-Specific Extension to the MIPS® Architecture
- Volume IV-d describes the SmartMIPS® Application-Specific Extension to the MIPS32® Architecture and the microMIPS32™ Architecture
- Volume IV-e describes the MIPS® DSP Application-Specific Extension to the MIPS® Architecture
- Volume IV-f describes the MIPS® MT Application-Specific Extension to the MIPS® Architecture
- Volume IV-h describes the MIPS® MCU Application-Specific Extension to the MIPS® Architecture

1.1 Typographical Conventions

This section describes the use of *italic*, **bold** and `courier` fonts in this book.

1.1.1 Italic Text

- is used for *emphasis*

About This Book

- is used for *bits, fields, registers*, that are important from a software perspective (for instance, address bits used by software, and programmable fields and registers), and various *floating point instruction formats*, such as *S, D,* and *PS*
- is used for the memory access types, such as *cached* and *uncached*

1.1.2 Bold Text

- represents a term that is being **defined**
- is used for **bits** and **fields** that are important from a hardware perspective (for instance, **register** bits, which are not programmable but accessible only to hardware)
- is used for ranges of numbers; the range is indicated by an ellipsis. For instance, **5..1** indicates numbers 5 through 1
- is used to emphasize **UNPREDICTABLE** and **UNDEFINED** behavior, as defined below.

1.1.3 Courier Text

`Courier` fixed-width font is used for text that is displayed on the screen, and for examples of code and instruction pseudocode.

1.2 UNPREDICTABLE and UNDEFINED

The terms **UNPREDICTABLE** and **UNDEFINED** are used throughout this book to describe the behavior of the processor in certain cases. **UNDEFINED** behavior or operations can occur only as the result of executing instructions in a privileged mode (i.e., in Kernel Mode or Debug Mode, or with the CPO usable bit set in the Status register). Unprivileged software can never cause **UNDEFINED** behavior or operations. Conversely, both privileged and unprivileged software can cause **UNPREDICTABLE** results or operations.

1.2.1 UNPREDICTABLE

UNPREDICTABLE results may vary from processor implementation to implementation, instruction to instruction, or as a function of time on the same implementation or instruction. Software can never depend on results that are **UNPREDICTABLE**. **UNPREDICTABLE** operations may cause a result to be generated or not. If a result is generated, it is **UNPREDICTABLE**. **UNPREDICTABLE** operations may cause arbitrary exceptions.

UNPREDICTABLE results or operations have several implementation restrictions:

- Implementations of operations generating **UNPREDICTABLE** results must not depend on any data source (memory or internal state) which is inaccessible in the current processor mode
- **UNPREDICTABLE** operations must not read, write, or modify the contents of memory or internal state which is inaccessible in the current processor mode. For example, **UNPREDICTABLE** operations executed in user mode must not access memory or internal state that is only accessible in Kernel Mode or Debug Mode or in another process
- **UNPREDICTABLE** operations must not halt or hang the processor

1.2.2 UNDEFINED

UNDEFINED operations or behavior may vary from processor implementation to implementation, instruction to instruction, or as a function of time on the same implementation or instruction. **UNDEFINED** operations or behavior may vary from nothing to creating an environment in which execution can no longer continue. **UNDEFINED** operations or behavior may cause data loss.

UNDEFINED operations or behavior has one implementation restriction:

- **UNDEFINED** operations or behavior must not cause the processor to hang (that is, enter a state from which there is no exit other than powering down the processor). The assertion of any of the reset signals must restore the processor to an operational state

1.2.3 UNSTABLE

UNSTABLE results or values may vary as a function of time on the same implementation or instruction. Unlike **UNPREDICTABLE** values, software may depend on the fact that a sampling of an **UNSTABLE** value results in a legal transient value that was correct at some point in time prior to the sampling.

UNSTABLE values have one implementation restriction:

- Implementations of operations generating **UNSTABLE** results must not depend on any data source (memory or internal state) which is inaccessible in the current processor mode

1.3 Special Symbols in Pseudocode Notation

In this book, algorithmic descriptions of an operation are described as pseudocode in a high-level language notation resembling Pascal. Special symbols used in the pseudocode notation are listed in [Table 1.1](#).

Table 1.1 Symbols Used in Instruction Operation Statements

Symbol	Meaning
\leftarrow	Assignment
$=, \neq$	Tests for equality and inequality
\parallel	Bit string concatenation
x^y	A y -bit string formed by y copies of the single-bit value x
$b\#n$	A constant value n in base b . For instance $10\#100$ represents the decimal value 100, $2\#100$ represents the binary value 100 (decimal 4), and $16\#100$ represents the hexadecimal value 100 (decimal 256). If the "b#" prefix is omitted, the default base is 10.
$0bn$	A constant value n in base 2. For instance $0b100$ represents the binary value 100 (decimal 4).
$0xn$	A constant value n in base 16. For instance $0x100$ represents the hexadecimal value 100 (decimal 256).
$x_{y..z}$	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
$*, \times$	2's complement or floating point multiplication (both used for either)
div	2's complement integer division

Table 1.1 Symbols Used in Instruction Operation Statements (Continued)

Symbol	Meaning
mod	2's complement modulo
/	Floating point division
<	2's complement less-than comparison
>	2's complement greater-than comparison
≤	2's complement less-than or equal comparison
≥	2's complement greater-than or equal comparison
nor	Bitwise logical NOR
xor	Bitwise logical XOR
and	Bitwise logical AND
or	Bitwise logical OR
GPRLen	The length in bits (32 or 64) of the CPU general-purpose registers
<i>GPR[x]</i>	CPU general-purpose register <i>x</i> . The content of <i>GPR[0]</i> is always zero. In Release 2 of the Architecture, <i>GPR[x]</i> is a short-hand notation for <i>SGPR[SRSCtl_{CSS}, x]</i> .
<i>SGPR[s,x]</i>	In Release 2 of the Architecture and subsequent releases, multiple copies of the CPU general-purpose registers may be implemented. <i>SGPR[s,x]</i> refers to GPR set <i>s</i> , register <i>x</i> .
<i>FPR[x]</i>	Floating Point operand register <i>x</i>
<i>FCC[CC]</i>	Floating Point condition code <i>CC</i> . <i>FCC[0]</i> has the same value as <i>COC[1]</i> .
<i>FPR[x]</i>	Floating Point (Coprocessor unit 1), general register <i>x</i>
<i>CPR[z,x,s]</i>	Coprocessor unit <i>z</i> , general register <i>x</i> , select <i>s</i>
CP2CPR[x]	Coprocessor unit 2, general register <i>x</i>
<i>CCR[z,x]</i>	Coprocessor unit <i>z</i> , control register <i>x</i>
CP2CCR[x]	Coprocessor unit 2, control register <i>x</i>
<i>COC[z]</i>	Coprocessor unit <i>z</i> condition signal
<i>Xlat[x]</i>	Translation of the MIPS16e GPR number <i>x</i> into the corresponding 32-bit GPR number
BigEndianMem	Endian mode as configured at chip reset (0 → Little-Endian, 1 → Big-Endian). Specifies the endianness of the memory interface (see LoadMemory and StoreMemory pseudocode function descriptions), and the endianness of Kernel and Supervisor mode execution.
BigEndianCPU	The endianness for load and store instructions (0 → Little-Endian, 1 → Big-Endian). In User mode, this endianness may be switched by setting the <i>RE</i> bit in the <i>Status</i> register. Thus, BigEndianCPU may be computed as (BigEndianMem XOR ReverseEndian).
ReverseEndian	Signal to reverse the endianness of load and store instructions. This feature is available in User mode only, and is implemented by setting the <i>RE</i> bit of the <i>Status</i> register. Thus, ReverseEndian may be computed as (SR _{RE} and User mode).
<i>LLbit</i>	Bit of virtual state used to specify operation for instructions that provide atomic read-modify-write. <i>LLbit</i> is set when a linked load occurs and is tested 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.

Table 1.1 Symbols Used in Instruction Operation Statements (Continued)

Symbol	Meaning						
I , I+n , I-n :	<p>This occurs as a prefix to <i>Operation</i> description lines and functions as a label. It indicates the instruction time during which the pseudocode appears to “execute.” 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 — that is, during the instruction time of another instruction. When this happens, the instruction operation is written in sections labeled 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 has the portion of the instruction operation description that writes the result register in a section labeled 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 labeled I for the following instruction. Within one pseudocode sequence, the effects of the statements take place in order. However, between sequences of statements for different instructions that occur “at the same time,” there is no defined order. Programs must not depend on a particular order of evaluation between such sections.</p>						
PC	<p>The <i>Program Counter</i> 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 <i>PC</i> during an instruction time. If no value is assigned to <i>PC</i> during an instruction time by any pseudocode statement, it is automatically incremented by either 2 (in the case of a 16-bit MIPS16e instruction) or 4 before the next instruction time. A taken branch assigns the target address to the <i>PC</i> during the instruction time of the instruction in the branch delay slot.</p> <p>In the MIPS Architecture, the <i>PC</i> value is only visible indirectly, such as when the processor stores the restart address into a GPR on a jump-and-link or branch-and-link instruction, or into a Coprocessor 0 register on an exception. The <i>PC</i> value contains a full 32-bit address all of which are significant during a memory reference.</p>						
ISA Mode	<p>In processors that implement the MIPS16e Application Specific Extension or the microMIPS base architectures, the <i>ISA Mode</i> is a single-bit register that determines in which mode the processor is executing, as follows:</p> <table border="1" data-bbox="602 1121 1271 1241"> <thead> <tr> <th>Encoding</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>The processor is executing 32-bit MIPS instructions</td> </tr> <tr> <td>1</td> <td>The processor is executing MIIPS16e instructions</td> </tr> </tbody> </table> <p>In the MIPS Architecture, the <i>ISA Mode</i> value is only visible indirectly, such as when the processor stores a combined value of the upper bits of <i>PC</i> and the <i>ISA Mode</i> into a GPR on a jump-and-link or branch-and-link instruction, or into a Coprocessor 0 register on an exception.</p>	Encoding	Meaning	0	The processor is executing 32-bit MIPS instructions	1	The processor is executing MIIPS16e instructions
Encoding	Meaning						
0	The processor is executing 32-bit MIPS instructions						
1	The processor is executing MIIPS16e instructions						
PABITS	<p>The number of physical address bits implemented is represented by the symbol PABITS. As such, if 36 physical address bits were implemented, the size of the physical address space would be $2^{\text{PABITS}} = 2^{36}$ bytes.</p>						
FP32RegistersMode	<p>Indicates whether the FPU has 32-bit or 64-bit floating point registers (FPRs). In MIPS32 Release 1, the FPU has 32 32-bit FPRs in which 64-bit data types are stored in even-odd pairs of FPRs. In MIPS64, (and optionally in MIPS32 Release2 and MIPSr3) the FPU has 32 64-bit FPRs in which 64-bit data types are stored in any FPR.</p> <p>In MIPS32 Release 1 implementations, FP32RegistersMode is always a 0. MIPS64 implementations have a compatibility mode in which the processor references the FPRs as if it were a MIPS32 implementation. In such a case FP32RegistersMode is computed from the FR bit in the <i>Status</i> register. If this bit is a 0, the processor operates as if it had 32 32-bit FPRs. If this bit is a 1, the processor operates with 32 64-bit FPRs. The value of FP32RegistersMode is computed from the FR bit in the <i>Status</i> register.</p>						
InstructionInBranchDelaySlot	<p>Indicates whether the instruction at the Program Counter address was executed in the delay slot of a branch or jump. This condition reflects the <i>dynamic</i> state of the instruction, not the <i>static</i> state. That is, the value is false if a branch or jump occurs to an instruction whose <i>PC</i> immediately follows a branch or jump, but which is not executed in the delay slot of a branch or jump.</p>						

Table 1.1 Symbols Used in Instruction Operation Statements (Continued)

Symbol	Meaning
SignalException(exception, argument)	Causes an exception to be signaled, using the exception parameter as the type of exception and the argument parameter as an exception-specific argument). Control does not return from this pseudocode function—the exception is signaled at the point of the call.

1.4 For More Information

Various MIPS RISC processor manuals and additional information about MIPS products can be found at the MIPS URL: <http://www.mips.com>

For comments or questions on the MIPS32® Architecture or this document, send Email to support@mips.com.

Guide to the Instruction Set


This chapter provides a detailed guide to understanding the instruction descriptions, which are listed in alphabetical order in the tables at the beginning of the next chapter.

2.1 Understanding the Instruction Fields

Figure 2.1 shows an example instruction. Following the figure are descriptions of the fields listed below:

- “Instruction Fields” on page 19
- “Instruction Descriptive Name and Mnemonic” on page 19
- “Format Field” on page 19
- “Purpose Field” on page 20
- “Description Field” on page 20
- “Restrictions Field” on page 20
- “Operation Field” on page 21
- “Exceptions Field” on page 21
- “Programming Notes and Implementation Notes Fields” on page 22

Figure 2.1 Example of Instruction Description




2.1.1 Instruction Fields

Fields encoding the instruction word are shown in register form at the top of the instruction description. The following rules are followed:

- The values of constant fields and the *opcode* names are listed in uppercase (SPECIAL and ADD in Figure 2.2). Constant values in a field are shown in binary below the symbolic or hexadecimal value.
- All variable fields are listed with the lowercase names used in the instruction description (*rs*, *rt*, and *rd* in Figure 2.2).
- Fields that contain zeros but are not named are unused fields that are required to be zero (bits 10:6 in Figure 2.2). If such fields are set to non-zero values, the operation of the processor is UNPREDICTABLE.

Figure 2.2 Example of Instruction Fields



2.1.2 Instruction Descriptive Name and Mnemonic

The instruction descriptive name and mnemonic are printed as page headings for each instruction, as shown in Figure 2.3.

Figure 2.3 Example of Instruction Descriptive Name and Mnemonic



2.1.3 Format Field

The assembler formats for the instruction and the architecture level at which the instruction was originally defined are given in the *Format* field. 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 their order of extension (for an example, see C.cond.fmt). 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.

Figure 2.4 Example of Instruction Format



The assembler format is shown with literal parts of the assembler instruction printed in uppercase characters. The variable parts, the operands, are shown as the lowercase names of the appropriate fields. The architectural level at which the instruction was first defined, for example “MIPS32” is shown at the right side of the page.

There can be more than one assembler format for each architecture level. 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 lists both ADD.S and ADD.D.

The assembler format lines sometimes include parenthetical comments to help explain variations in the formats (once again, see [C.cond.fmt](#)). These comments are not a part of the assembler format.

2.1.4 Purpose Field

The *Purpose* field gives a short description of the use of the instruction.

Figure 2.5 Example of Instruction Purpose

Purpose: Add Word

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

2.1.5 Description Field

If a one-line symbolic description of the instruction is feasible, it appears 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.

Figure 2.6 Example of Instruction Description

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

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, the destination register is not modified and an Integer Overflow exception occurs.
- If the addition does not overflow, the 32-bit result is placed into GPR *rd*.

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 /Status* register.

2.1.6 Restrictions Field

The *Restrictions* field documents any possible restrictions that may affect the instruction. Most restrictions fall into one of the following six categories:

- Valid values for instruction fields (for example, see floating point [ADD.fmt](#))
- ALIGNMENT requirements for memory addresses (for example, see [LW](#))
- Valid values of operands (for example, see [ALNV.PS](#))

- Valid operand formats (for example, see floating point [ADD.fmt](#))
- Order of instructions necessary to guarantee correct execution. These ordering constraints avoid pipeline hazards for which some processors do not have hardware interlocks (for example, see [MUL](#)).
- Valid memory access types (for example, see [LL/SC](#))

Figure 2.7 Example of Instruction Restrictions**Restrictions:**

None

2.1.7 Operation Field

The *Operation* field describes the operation of the instruction as pseudocode in a high-level language notation resembling Pascal. 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 are omitted for legibility.

Figure 2.8 Example of Instruction Operation**Operation:**

```
temp ← (GPR[rs]31 || GPR[rs]31..0) + (GPR[rt]31 || GPR[rt]31..0)
if temp32 ≠ temp31 then
    SignalException(IntegerOverflow)
else
    GPR[rd] ← temp
endif
```

See 2.2 “[Operation Section Notation and Functions](#)” on page 22 for more information on the formal notation used here.

2.1.8 Exceptions Field

The *Exceptions* field lists the exceptions that can be caused by *Operation* of the instruction. It omits exceptions that can be caused by the instruction fetch, for instance, TLB Refill, and also omits exceptions that can be caused by asynchronous external events such as an Interrupt. Although a 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 dependent upon the implementation.

Figure 2.9 Example of Instruction Exception**Exceptions:**

Integer Overflow

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

2.1.9 Programming Notes and Implementation Notes Fields

The *Notes* 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.

Figure 2.10 Example of Instruction Programming Notes

Programming Notes:

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

2.2 Operation Section Notation and Functions

In an instruction description, the *Operation* section uses a high-level language notation to describe the operation performed by each instruction. Special symbols used in the pseudocode are described in the previous chapter. Specific pseudocode functions are described below.

This section presents information about the following topics:

- “[Instruction Execution Ordering](#)” on page 22
- “[Pseudocode Functions](#)” on page 22

2.2.1 Instruction Execution Ordering

Each of the high-level language statements in the *Operations* section are executed sequentially (except as constrained by conditional and loop constructs).

2.2.2 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. These functions are defined in this section, and include the following:

- “[Coprocessor General Register Access Functions](#)” on page 22
- “[Memory Operation Functions](#)” on page 24
- “[Floating Point Functions](#)” on page 27
- “[Miscellaneous Functions](#)” on page 30

2.2.2.1 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 functions described in this section.

COP_LW

The COP_LW function defines 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 memword in coprocessor general register *rt*.

Figure 2.11 COP_LW Pseudocode Function

```
COP_LW (z, rt, memword)
  z: The coprocessor unit number
  rt: Coprocessor general register specifier
  memword: A 32-bit word value supplied to the coprocessor

  /* Coprocessor-dependent action */

endfunction COP_LW
```

COP_LD

The COP_LD function defines 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 memdouble in coprocessor general register *rt*.

Figure 2.12 COP_LD Pseudocode Function

```
COP_LD (z, rt, memdouble)
  z: The coprocessor unit number
  rt: Coprocessor general register specifier
  memdouble: 64-bit doubleword value supplied to the coprocessor.

  /* Coprocessor-dependent action */

endfunction COP_LD
```

COP_SW

The COP_SW function 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 *rt*.

Figure 2.13 COP_SW Pseudocode Function

```
dataword ← COP_SW (z, rt)
  z: The coprocessor unit number
  rt: Coprocessor general register specifier
  dataword: 32-bit word value

  /* Coprocessor-dependent action */

endfunction COP_SW
```

COP_SD

The COP_SD function 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 low-order doubleword in coprocessor general register *rt*.

Figure 2.14 COP_SD Pseudocode Function

```

datadouble ← COP_SD (z, rt)
  z: The coprocessor unit number
  rt: Coprocessor general register specifier
  datadouble: 64-bit doubleword value

  /* Coprocessor-dependent action */

endfunction COP_SD

```

CoprocessorOperation

The CoprocessorOperation function performs the specified Coprocessor operation.

Figure 2.15 CoprocessorOperation Pseudocode Function

```

CoprocessorOperation (z, cop_fun)

  /* z:          Coprocessor unit number */
  /* cop_fun:    Coprocessor function from function field of instruction */

  /* Transmit the cop_fun value to coprocessor z */

endfunction CoprocessorOperation

```

2.2.2.2 Memory Operation Functions

Regardless of byte ordering (big- or little-endian), the address of a halfword, word, or doubleword is the smallest byte address of the bytes that form 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* pseudocode for load and store operations, the following functions summarize the handling of virtual addresses and the access of 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 2.1](#). The bytes within the addressed unit of memory (word for 32-bit processors or doubleword for 64-bit processors) that are used can be determined directly from the *AccessLength* and the two or three low-order bits of the address.

AddressTranslation

The AddressTranslation function translates a virtual address to a physical address and its cacheability and coherency attribute, describing the mechanism used to resolve the memory reference.

Given the virtual address *vAddr*, and whether the reference is to Instructions or Data (*IorD*), find the corresponding physical address (*pAddr*) and the cacheability and coherency attribute (*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 or fixed mapping MMU determines 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.

Figure 2.16 AddressTranslation Pseudocode Function

```

(pAddr, CCA) ← AddressTranslation (vAddr, IorD, LorS)

  /* pAddr: physical address */
  /* CCA:   Cacheability&Coherency Attribute, the method used to access caches*/

```



```

/*      and memory and resolve the reference */

/* vAddr: virtual address */
/* IorD: Indicates whether access is for INSTRUCTION or DATA */
/* LorS: Indicates whether access is for LOAD or STORE */

/* See the address translation description for the appropriate MMU */
/* type in Volume III of this book for the exact translation mechanism */

endfunction AddressTranslation

```

LoadMemory

The LoadMemory function loads a value from memory.

This action uses cache and main memory as specified in both the Cacheability and Coherency Attribute (*CCA*) and the access (*IorD*) to find the contents of *AccessLength* memory bytes, starting at physical location *pAddr*. The data is returned in a fixed-width naturally aligned memory element (*MemElem*). The low-order 2 (or 3) bits of the address and the *AccessLength* indicate which of the bytes within *MemElem* need to be passed to the processor. If the memory access type of the reference is *uncached*, only the referenced bytes are read from memory and marked as valid within the memory element. If the access type is *cached* but 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, this block is the entire memory element.

Figure 2.17 LoadMemory Pseudocode Function

```

MemElem ← LoadMemory (CCA, AccessLength, pAddr, vAddr, IorD)

/* 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:     Cacheability&CoherencyAttribute=method used to access caches */
/*          and memory and resolve the reference */

/* AccessLength: Length, in bytes, of access */
/* pAddr:     physical address */
/* vAddr:     virtual address */
/* IorD:     Indicates whether access is for Instructions or Data */

endfunction LoadMemory

```

StoreMemory

The StoreMemory function stores 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 Cacheability and Coherency Attribute (*CCA*). The *MemElem* contains the data for an aligned, fixed-width memory element (a word for 32-bit processors, a doubleword for 64-bit processors), though only the bytes that are actually stored to memory need be valid. The low-order two (or three) bits of *pAddr* and the *AccessLength* field indicate which of the bytes within the *MemElem* data should be stored; only these bytes in memory will actually be changed.

Figure 2.18 StoreMemory Pseudocode Function

```

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

```

```

/* CCA:      Cacheability&Coherency Attribute, 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, either 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 */

endfunction StoreMemory

```

Prefetch

The Prefetch function prefetches 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.

Figure 2.19 Prefetch Pseudocode Function

```

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

/* CCA:      Cacheability&Coherency Attribute, 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 */

endfunction Prefetch

```

Table 2.1 lists the data access lengths and their labels for loads and stores.

Table 2.1 AccessLength Specifications for Loads/Stores

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)

SyncOperation

The SyncOperation function orders loads and stores to synchronize shared memory.

This action makes the effects of the synchronizable loads and stores indicated by *stype* occur in the same order for all processors.

Figure 2.20 SyncOperation Pseudocode Function

```
SyncOperation(stype)

    /* stype: Type of load/store ordering to perform. */

    /* Perform implementation-dependent operation to complete the */
    /* required synchronization operation */

endfunction SyncOperation
```

2.2.2.3 Floating Point Functions

The pseudocode shown in below specifies how the unformatted contents loaded or moved to CPI 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

The ValueFPR function returns a formatted value from the floating point registers.

Figure 2.21 ValueFPR Pseudocode Function

```
value ← ValueFPR(fpr, fmt)

    /* value: The formattted value from the FPR */

    /* fpr:   The FPR number */
    /* fmt:   The format of the data, one of: */
    /*       S, D, W, L, PS, */
    /*       OB, QH, */
    /*       UNINTERPRETED_WORD, */
    /*       UNINTERPRETED_DOUBLEWORD */
    /* The UNINTERPRETED values are used to indicate that the datatype */
    /* is not known as, for example, in SWC1 and SDC1 */

case fmt of
    S, W, UNINTERPRETED_WORD:
        valueFPR ← FPR[fpr]

    D, UNINTERPRETED_DOUBLEWORD:
        if (FP32RegistersMode = 0)
            if (fpr0 ≠ 0) then
                valueFPR ← UNPREDICTABLE
            else
                valueFPR ← FPR[fpr+1]31..0 || FPR[fpr]31..0
            endif
        else
            valueFPR ← FPR[fpr]
        endif

    L, PS:
        if (FP32RegistersMode = 0) then
            valueFPR ← UNPREDICTABLE
```

```

        else
            valueFPR ← FPR[fpr]
        endif

    DEFAULT:
        valueFPR ← UNPREDICTABLE

    endcase
endfunction ValueFPR

```

The pseudocode shown below specifies the way 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 receives a value from the StoreFPR(), it is not valid to interpret the value with ValueFPR() in a different format.

StoreFPR

Figure 2.22 StoreFPR Pseudocode Function

```

StoreFPR (fpr, fmt, value)

/* fpr:   The FPR number */
/* fmt:   The format of the data, one of: */
/*        S, D, W, L, PS, */
/*        OB, QH, */
/*        UNINTERPRETED_WORD, */
/*        UNINTERPRETED_DOUBLEWORD */
/* value: The formatted value to be stored into the FPR */

/* The UNINTERPRETED values are used to indicate that the datatype */
/* is not known as, for example, in LWC1 and LDC1 */

case fmt of
    S, W, UNINTERPRETED_WORD:
        FPR[fpr] ← value

    D, UNINTERPRETED_DOUBLEWORD:
        if (FP32RegistersMode = 0)
            if (fpr0 ≠ 0) then
                UNPREDICTABLE
            else
                FPR[fpr]   ← UNPREDICTABLE32 || value31..0
                FPR[fpr+1] ← UNPREDICTABLE32 || value63..32
            endif
        else
            FPR[fpr] ← value
        endif

    L, PS:
        if (FP32RegistersMode = 0) then
            UNPREDICTABLE
        else
            FPR[fpr] ← value
        endif

endcase

```

```
endfunction StoreFPR
```

The pseudocode shown below checks for an enabled floating point exception and conditionally signals the exception.

CheckFPEException

Figure 2.23 CheckFPEException Pseudocode Function

```
CheckFPEException()

/* A floating point exception is signaled if the E bit of the Cause field is a 1 */
/* (Unimplemented Operations have no enable) or if any bit in the Cause field */
/* and the corresponding bit in the Enable field are both 1 */

    if ( (FCSR17 = 1) or
        ((FCSR16..12 and FCSR11..7) ≠ 0) ) then
        SignalException(FloatingPointException)
    endif

endfunction CheckFPEException
```

FPCConditionCode

The FPCConditionCode function returns the value of a specific floating point condition code.

Figure 2.24 FPCConditionCode Pseudocode Function

```
tf ← FPCConditionCode(cc)

/* tf: The value of the specified condition code */

/* cc: The Condition code number in the range 0..7 */

if cc = 0 then
    FPCConditionCode ← FCSR23
else
    FPCConditionCode ← FCSR24+cc
endif

endfunction FPCConditionCode
```

SetFPCConditionCode

The SetFPCConditionCode function writes a new value to a specific floating point condition code.

Figure 2.25 SetFPCConditionCode Pseudocode Function

```
SetFPCConditionCode(cc, tf)
    if cc = 0 then
        FCSR ← FCSR31..24 || tf || FCSR22..0
    else
        FCSR ← FCSR31..25+cc || tf || FCSR23+cc..0
    endif

endfunction SetFPCConditionCode
```

2.2.2.4 Miscellaneous Functions

This section lists miscellaneous functions not covered in previous sections.

SignalException

The `SignalException` function signals an exception condition.

This action results in an exception that aborts the instruction. The instruction operation pseudocode never sees a return from this function call.

Figure 2.26 SignalException Pseudocode Function

```
SignalException(Exception, argument)

/* Exception:   The exception condition that exists. */
/* argument:   A exception-dependent argument, if any */

endfunction SignalException
```

SignalDebugBreakpointException

The `SignalDebugBreakpointException` function signals a condition that causes entry into Debug Mode from non-Debug Mode.

This action results in an exception that aborts the instruction. The instruction operation pseudocode never sees a return from this function call.

Figure 2.27 SignalDebugBreakpointException Pseudocode Function

```
SignalDebugBreakpointException()

endfunction SignalDebugBreakpointException
```

SignalDebugModeBreakpointException

The `SignalDebugModeBreakpointException` function signals a condition that causes entry into Debug Mode from Debug Mode (i.e., an exception generated while already running in Debug Mode).

This action results in an exception that aborts the instruction. The instruction operation pseudocode never sees a return from this function call.

Figure 2.28 SignalDebugModeBreakpointException Pseudocode Function

```
SignalDebugModeBreakpointException()

endfunction SignalDebugModeBreakpointException
```

NullifyCurrentInstruction

The `NullifyCurrentInstruction` function nullifies the current instruction.

The instruction is aborted, inhibiting not only the functional effect of the instruction, but also inhibiting all exceptions detected during fetch, decode, or execution of the instruction in question. For branch-likely instructions, nullification kills the instruction in the delay slot of the branch likely instruction.

Figure 2.29 NullifyCurrentInstruction PseudoCode Function

```
NullifyCurrentInstruction()

endfunction NullifyCurrentInstruction
```

JumpDelaySlot

The JumpDelaySlot function is used in the pseudocode for the PC-relative instructions in the MIPS16e ASE. The function returns TRUE if the instruction at *vAddr* is executed in a jump delay slot. A jump delay slot always immediately follows a JR, JAL, JALR, or JALX instruction.

Figure 2.30 JumpDelaySlot Pseudocode Function

```
JumpDelaySlot(vAddr)

/* vAddr:Virtual address */

endfunction JumpDelaySlot
```

PolyMult

The PolyMult function multiplies two binary polynomial coefficients.

Figure 2.31 PolyMult Pseudocode Function

```
PolyMult(x, y)
  temp ← 0
  for i in 0 .. 31
    if  $x_i = 1$  then
      temp ← temp xor ( $y_{(31-i)..0} || 0^i$ )
    endif
  endfor

  PolyMult ← temp

endfunction PolyMult
```

2.3 Op and Function Subfield Notation

In some instructions, the instruction subfields *op* and *function* can have constant 5- or 6-bit values. When reference is made to these instructions, uppercase mnemonics are used. For instance, in the floating point ADD instruction, *op*=COPI and *function*=ADD. In other cases, a single field has both fixed and variable subfields, so the name contains both upper- and lowercase characters.

2.4 FPU Instructions

In the detailed description of each FPU instruction, all variable subfields in an instruction format (such as *fs*, *ft*, *immediate*, and so on) are shown in lowercase. The instruction name (such as ADD, SUB, and so on) is shown in uppercase.

For the sake of clarity, an alias is sometimes used for a variable subfield in the formats of specific instructions. For example, *rs=base* in the format for load and store instructions. Such an alias is always lowercase since it refers to a variable subfield.

Guide to the Instruction Set

Bit encodings for mnemonics are given in Volume I, in the chapters describing the CPU, FPU, MDMX, and MIPS16e instructions.

See “[Op and Function Subfield Notation](#)” on page 31 for a description of the *op* and *function* subfields.

The MIPS32® Instruction Set

3.1 Compliance and Subsetting

To be compliant with the MIPS32 Architecture, designs must implement a set of required features, as described in this document set. To allow flexibility in implementations, the MIPS32 Architecture does provide subsetting rules. An implementation that follows these rules is compliant with the MIPS32 Architecture as long as it adheres strictly to the rules, and fully implements the remaining instructions. Supersetting of the MIPS32 Architecture is only allowed by adding functions to the *SPECIAL2* major opcode, by adding control for co-processors via the *COP2*, *LWC2*, *SWC2*, *LDC2*, and/or *SDC2*, or via the addition of approved Application Specific Extensions.

Note: The use of COP3 as a customizable coprocessor has been removed in the Release 2 of the MIPS32 architecture. The use of the COP3 is now reserved for the future extension of the architecture. Implementations using Release 1 of the MIPS32 architecture are strongly discouraged from using the COP3 opcode for a user-available coprocessor as doing so will limit the potential for an upgrade path to a 64-bit floating point unit.

The instruction set subsetting rules are as follows:

- All CPU instructions must be implemented - no subsetting is allowed (unless described in this list).
- The FPU and related support instructions, including the MOVF and MOVT CPU instructions, may be omitted. Software may determine if an FPU is implemented by checking the state of the FP bit in the *Config1* CP0 register. If the FPU is implemented, it must include S, D, and W formats, operate instructions, and all supporting instructions. Software may determine which FPU data types are implemented by checking the appropriate bit in the *FIR* CP1 register. The following allowable FPU subsets are compliant with the MIPS32 architecture:
 - No FPU
 - FPU with S, D, and W formats and all supporting instructions
- Coprocessor 2 is optional and may be omitted. Software may determine if Coprocessor 2 is implemented by checking the state of the C2 bit in the *Config1* CP0 register. If Coprocessor 2 is implemented, the Coprocessor 2 interface instructions (BC2, CFC2, COP2, CTC2, LDC2, LWC2, MFC2, MTC2, SDC2, and SWC2) may be omitted on an instruction-by-instruction basis.
- Supervisor Mode is optional. If Supervisor Mode is not implemented, bit 3 of the *Status* register must be ignored on write and read as zero.
- The standard TLB-based memory management unit may be replaced with a simpler MMU (e.g., a Fixed Mapping MMU). If this is done, the rest of the interface to the Privileged Resource Architecture must be preserved. If a TLB-based memory management unit is implemented, it must be the standard TLB-based MMU as described in the Privileged Resource Architecture chapter. Software may determine the type of the MMU by checking the MT field in the *Config* CP0 register.

The MIPS32® Instruction Set

- The Privileged Resource Architecture includes several implementation options and may be subsetted in accordance with those options.
- Instruction, CP0 Register, and CP1 Control Register fields that are marked “Reserved” or shown as “0” in the description of that field are reserved for future use by the architecture and are not available to implementations. Implementations may only use those fields that are explicitly reserved for implementation dependent use.
- Supported ASEs are optional and may be subsetted out. In most cases, software may determine if a supported ASE is implemented by checking the appropriate bit in the *Config1* or *Config3* CP0 register. If they are implemented, they must implement the entire ISA applicable to the component, or implement subsets that are approved by the ASE specifications.
- EJTAG is optional and may be subsetted out. If it is implemented, it must implement only those subsets that are approved by the EJTAG specification.
- The JALX instruction is only implemented when there are other instruction sets available on the device (microMIPS or MIPS16e).
- If any instruction is subsetted out based on the rules above, an attempt to execute that instruction must cause the appropriate exception (typically Reserved Instruction or Coprocessor Unusable).

3.2 Alphabetical List of Instructions

Table 3.1 through Table 3.24 provide a list of instructions grouped by category. Individual instruction descriptions follow the tables, arranged in alphabetical order.

Table 3.1 CPU Arithmetic Instructions

Mnemonic	Instruction
ADD	Add Word
ADDI	Add Immediate Word
ADDIU	Add Immediate Unsigned Word
ADDU	Add Unsigned Word
CLO	Count Leading Ones in Word
CLZ	Count Leading Zeros in Word
DIV	Divide Word
DIVU	Divide Unsigned Word
MADD	Multiply and Add Word to Hi, Lo
MADDU	Multiply and Add Unsigned Word to Hi, Lo
MSUB	Multiply and Subtract Word to Hi, Lo
MSUBU	Multiply and Subtract Unsigned Word to Hi, Lo
MUL	Multiply Word to GPR
MULT	Multiply Word
MULTU	Multiply Unsigned Word

Table 3.1 CPU Arithmetic Instructions (Continued)

Mnemonic	Instruction	
SEB	Sign-Extend Byte	Release 2 & subsequent
SEH	Sign-Extend Halfword	Release 2 & subsequent
SLT	Set on Less Than	
SLTI	Set on Less Than Immediate	
SLTIU	Set on Less Than Immediate Unsigned	
SLTU	Set on Less Than Unsigned	
SUB	Subtract Word	
SUBU	Subtract Unsigned Word	

Table 3.2 CPU Branch and Jump Instructions

Mnemonic	Instruction	
B	Unconditional Branch	
BAL	Branch and Link	
BEQ	Branch on Equal	
BGEZ	Branch on Greater Than or Equal to Zero	
BGEZAL	Branch on Greater Than or Equal to Zero and Link	
BGTZ	Branch on Greater Than Zero	
BLEZ	Branch on Less Than or Equal to Zero	
BLTZ	Branch on Less Than Zero	
BLTZAL	Branch on Less Than Zero and Link	
BNE	Branch on Not Equal	
J	Jump	
JAL	Jump and Link	
JALR	Jump and Link Register	
JALR.HB	Jump and Link Register with Hazard Barrier	Release 2 & subsequent
JALX	Jump and Link Exchange	microMIPS or MIPS16e also implemented
JR	Jump Register	
JR.HB	Jump Register with Hazard Barrier	Release 2 & subsequent

Table 3.3 CPU Instruction Control Instructions

Mnemonic	Instruction	
EHB	Execution Hazard Barrier	Release 2 & subsequent
NOP	No Operation	
PAUSE	Wait for LLBit to Clear	Release 2.1 & subsequent
SSNOP	Superscalar No Operation	

Table 3.4 CPU Load, Store, and Memory Control Instructions

Mnemonic	Instruction	
LB	Load Byte	
LBU	Load Byte Unsigned	
LH	Load Halfword	
LHU	Load Halfword Unsigned	
LL	Load Linked Word	
LW	Load Word	
LWL	Load Word Left	
LWR	Load Word Right	
PREF	Prefetch	
SB	Store Byte	
SC	Store Conditional Word	
SH	Store Halfword	
SW	Store Word	
SWL	Store Word Left	
SWR	Store Word Right	
SYNC	Synchronize Shared Memory	
SYNCI	Synchronize Caches to Make Instruction Writes Effective	Release 2 & subsequent

Table 3.5 CPU Logical Instructions

Mnemonic	Instruction
AND	And
ANDI	And Immediate

Table 3.5 CPU Logical Instructions (Continued)

Mnemonic	Instruction
LUI	Load Upper Immediate
NOR	Not Or
OR	Or
ORI	Or Immediate
XOR	Exclusive Or
XORI	Exclusive Or Immediate

Table 3.6 CPU Insert/Extract Instructions

Mnemonic	Instruction	
EXT	Extract Bit Field	Release 2 & subsequent
INS	Insert Bit Field	Release 2 & subsequent
WSBH	Word Swap Bytes Within Halfwords	Release 2 & subsequent

Table 3.7 CPU Move Instructions

Mnemonic	Instruction	
MFHI	Move From HI Register	
MFLO	Move From LO Register	
MOVF	Move Conditional on Floating Point False	
MOVN	Move Conditional on Not Zero	
MOVT	Move Conditional on Floating Point True	
MOVZ	Move Conditional on Zero	
MTHI	Move To HI Register	
MTLO	Move To LO Register	
RDHWR	Read Hardware Register	Release 2 & subsequent

Table 3.8 CPU Shift Instructions

Mnemonic	Instruction	
ROTR	Rotate Word Right	Release 2 & subsequent
ROTRV	Rotate Word Right Variable	Release 2 & subsequent
SLL	Shift Word Left Logical	

Table 3.8 CPU Shift Instructions (Continued)

Mnemonic	Instruction
SLLV	Shift Word Left Logical Variable
SRA	Shift Word Right Arithmetic
SRAV	Shift Word Right Arithmetic Variable
SRL	Shift Word Right Logical
SRLV	Shift Word Right Logical Variable

Table 3.9 CPU Trap Instructions

Mnemonic	Instruction
BREAK	Breakpoint
SYSCALL	System Call
TEQ	Trap if Equal
TEQI	Trap if Equal Immediate
TGE	Trap if Greater or Equal
TGEI	Trap if Greater of Equal Immediate
TGEIU	Trap if Greater or Equal Immediate Unsigned
TGEU	Trap if Greater or Equal Unsigned
TLT	Trap if Less Than
TLTI	Trap if Less Than Immediate
TLTIU	Trap if Less Than Immediate Unsigned
TLTU	Trap if Less Than Unsigned
TNE	Trap if Not Equal
TNEI	Trap if Not Equal Immediate

Table 3.10 Obsolete¹ CPU Branch Instructions

Mnemonic	Instruction
BEQL	Branch on Equal Likely
BGEZALL	Branch on Greater Than or Equal to Zero and Link Likely
BGEZL	Branch on Greater Than or Equal to Zero Likely
BGTZL	Branch on Greater Than Zero Likely
BLEZL	Branch on Less Than or Equal to Zero Likely

Table 3.10 Obsolete¹ CPU Branch Instructions (Continued)

Mnemonic	Instruction
BLTZALL	Branch on Less Than Zero and Link Likely
BLTZL	Branch on Less Than Zero Likely
BNEL	Branch on Not Equal Likely

1. Software is strongly encouraged to avoid use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS32 architecture.

Table 3.11 FPU Arithmetic Instructions

Mnemonic	Instruction
ABS.fmt	Floating Point Absolute Value
ADD.fmt	Floating Point Add
DIV.fmt	Floating Point Divide
MADD.fmt	Floating Point Multiply Add
MSUB.fmt	Floating Point Multiply Subtract
MUL.fmt	Floating Point Multiply
NEG.fmt	Floating Point Negate
NMADD.fmt	Floating Point Negative Multiply Add
NMSUB.fmt	Floating Point Negative Multiply Subtract
RECIP.fmt	Reciprocal Approximation
RSQRT.fmt	Reciprocal Square Root Approximation
SQRT.fmt	Floating Point Square Root
SUB.fmt	Floating Point Subtract

Table 3.12 FPU Branch Instructions

Mnemonic	Instruction
BC1F	Branch on FP False
BC1T	Branch on FP True

Table 3.13 FPU Compare Instructions

Mnemonic	Instruction
C.cond.fmt	Floating Point Compare

Table 3.14 FPU Convert Instructions

Mnemonic	Instruction	
ALNV.PS	Floating Point Align Variable	64-bit FPU Only
CEIL.L.fmt	Floating Point Ceiling Convert to Long Fixed Point	64-bit FPU Only
CEIL.W.fmt	Floating Point Ceiling Convert to Word Fixed Point	
CVT.D.fmt	Floating Point Convert to Double Floating Point	
CVT.L.fmt	Floating Point Convert to Long Fixed Point	64-bit FPU Only
CVT.PS.S	Floating Point Convert Pair to Paired Single	64-bit FPU Only
CVT.S.PL	Floating Point Convert Pair Lower to Single Floating Point	64-bit FPU Only
CVT.S.PU	Floating Point Convert Pair Upper to Single Floating Point	64-bit FPU Only
CVT.S.fmt	Floating Point Convert to Single Floating Point	
CVT.W.fmt	Floating Point Convert to Word Fixed Point	
FLOOR.L.fmt	Floating Point Floor Convert to Long Fixed Point	64-bit FPU Only
FLOOR.W.fmt	Floating Point Floor Convert to Word Fixed Point	
PLL.PS	Pair Lower Lower	64-bit FPU Only
PLU.PS	Pair Lower Upper	64-bit FPU Only
PUL.PS	Pair Upper Lower	64-bit FPU Only
PUU.PS	Pair Upper Upper	64-bit FPU Only
ROUND.L.fmt	Floating Point Round to Long Fixed Point	64-bit FPU Only
ROUND.W.fmt	Floating Point Round to Word Fixed Point	
TRUNC.L.fmt	Floating Point Truncate to Long Fixed Point	64-bit FPU Only
TRUNC.W.fmt	Floating Point Truncate to Word Fixed Point	

Table 3.15 FPU Load, Store, and Memory Control Instructions

Mnemonic	Instruction	
LDC1	Load Doubleword to Floating Point	
LDXC1	Load Doubleword Indexed to Floating Point	64-bit FPU Only
LUXC1	Load Doubleword Indexed Unaligned to Floating Point	64-bit FPU Only
LWC1	Load Word to Floating Point	
LWXC1	Load Word Indexed to Floating Point	64-bit FPU Only
PREFX	Prefetch Indexed	
SDC1	Store Doubleword from Floating Point	

Table 3.15 FPU Load, Store, and Memory Control Instructions (Continued)

Mnemonic	Instruction	
SDXC1	Store Doubleword Indexed from Floating Point	64-bit FPU Only
SUXC1	Store Doubleword Indexed Unaligned from Floating Point	64-bit FPU Only
SWC1	Store Word from Floating Point	
SWXC1	Store Word Indexed from Floating Point	64-bit FPU Only

Table 3.16 FPU Move Instructions

Mnemonic	Instruction	
CFC1	Move Control Word from Floating Point	
CTC1	Move Control Word to Floating Point	
MFC1	Move Word from Floating Point	
MFHC1	Move Word from High Half of Floating Point Register	Release 2 & subsequent
MOV.fmt	Floating Point Move	
MOV.fmt	Floating Point Move Conditional on Floating Point False	
MOVN.fmt	Floating Point Move Conditional on Not Zero	
MOVT.fmt	Floating Point Move Conditional on Floating Point True	
MOVZ.fmt	Floating Point Move Conditional on Zero	
MTC1	Move Word to Floating Point	
MTHC1	Move Word to High Half of Floating Point Register	Release 2 & subsequent

Table 3.17 Obsolete¹ FPU Branch Instructions

Mnemonic	Instruction
BC1FL	Branch on FP False Likely
BC1TL	Branch on FP True Likely

1. Software is strongly encouraged to avoid use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS32 architecture.

Table 3.18 Coprocessor Branch Instructions

Mnemonic	Instruction
BC2F	Branch on COP2 False
BC2T	Branch on COP2 True

Table 3.19 Coprocessor Execute Instructions

Mnemonic	Instruction
COP2	Coprocessor Operation to Coprocessor 2

Table 3.20 Coprocessor Load and Store Instructions

Mnemonic	Instruction
LDC2	Load Doubleword to Coprocessor 2
LWC2	Load Word to Coprocessor 2
SDC2	Store Doubleword from Coprocessor 2
SWC2	Store Word from Coprocessor 2

Table 3.21 Coprocessor Move Instructions

Mnemonic	Instruction	
CFC2	Move Control Word from Coprocessor 2	
CTC2	Move Control Word to Coprocessor 2	
MFC2	Move Word from Coprocessor 2	
MFHC2	Move Word from High Half of Coprocessor 2 Register	Release 2 & subsequent
MTC2	Move Word to Coprocessor 2	
MTHC2	Move Word to High Half of Coprocessor 2 Register	Release 2 & subsequent

Table 3.22 Obsolete¹ Coprocessor Branch Instructions

Mnemonic	Instruction
BC2FL	Branch on COP2 False Likely
BC2TL	Branch on COP2 True Likely

1. Software is strongly encouraged to avoid use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS32 architecture.

Table 3.23 Privileged Instructions


Mnemonic	Instruction
CACHE	Perform Cache Operation

Table 3.23 Privileged Instructions (Continued)

Mnemonic	Instruction	
DI	Disable Interrupts	Release 2 & subsequent
EI	Enable Interrupts	Release 2 & subsequent
ERET	Exception Return	
MFC0	Move from Coprocessor 0	
MTC0	Move to Coprocessor 0	
RDPGPR	Read GPR from Previous Shadow Set	Release 2 & subsequent
TLBP	Probe TLB for Matching Entry	
TLBR	Read Indexed TLB Entry	
TLBWI	Write Indexed TLB Entry	
TLBWR	Write Random TLB Entry	
WAIT	Enter Standby Mode	
WRPGPR	Write GPR to Previous Shadow Set	Release 2 & subsequent

Table 3.24 EJTAG Instructions

Mnemonic	Instruction
DERET	Debug Exception Return
SDBBP	Software Debug Breakpoint



Format: ABS.fmt
 ABS.S fd, fs
 ABS.D fd, fs
 ABS.PS fd, fs

MIPS32
MIPS32
MIPS64, MIPS32 Release 2

Purpose: Floating Point Absolute Value

Description: $FPR[fd] \leftarrow \text{abs}(FPR[fs])$

The absolute value of the value in FPR *fs* is placed in FPR *fd*. The operand and result are values in format *fmt*. ABS.PS takes the absolute value of the two values in FPR *fs* independently, and ORs together any generated exceptions.

Cause bits are ORed into the *Flag* bits if no exception is taken.

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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of ABS.PS is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:


StoreFPR(fd, fmt, AbsoluteValue(ValueFPR(fs, fmt)))

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation, Invalid Operation



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

MIPS32

Purpose: Add Word

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

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

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, the destination register is not modified and an Integer Overflow exception occurs.
- If the addition does not overflow, the 32-bit result is placed into GPR *rd*.

Restrictions:

None

Operation:


```
temp ← (GPR[rs]31 || GPR[rs]31..0) + (GPR[rt]31 || GPR[rt]31..0)
if temp32 ≠ temp31 then
    SignalException(IntegerOverflow)
else
    GPR[rd] ← temp
endif
```

Exceptions:

Integer Overflow

Programming Notes:

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



Format: ADD.fmt
 ADD.S fd, fs, ft
 ADD.D fd, fs, ft
 ADD.PS fd, fs, ft

MIPS32
MIPS32
MIPS64, MIPS32 Release 2

Purpose: Floating Point Add

To add floating point values

Description: $FPR[fd] \leftarrow FPR[fs] + FPR[ft]$

The value in FPR *ft* is added to the value in FPR *fs*. The result is calculated to infinite precision, rounded by using to the current rounding mode in *FCSR*, and placed into FPR *fd*. The operands and result are values in format *fmt*. ADD.PS adds the upper and lower halves of FPR *fs* and FPR *ft* independently, and ORs together any generated exceptions.

Cause bits are ORed into the *Flag* bits if no exception is taken.

Restrictions:

The fields *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *fmt*. If they are not valid, the result is **UNPREDICTABLE**.

The operands must be values in format *fmt*; if they are not, the result is **UNPREDICTABLE** and the value of the operand FPRs becomes **UNPREDICTABLE**.

The result of ADD.PS is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:


StoreFPR (fd, fmt, ValueFPR(fs, fmt) +_{fmt} ValueFPR(ft, fmt))

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation, Invalid Operation, Inexact, Overflow, Underflow



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

MIPS32

Purpose: Add Immediate Word

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

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

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, the destination register is not modified and an Integer Overflow exception occurs.
- If the addition does not overflow, the 32-bit result is placed into GPR *rt*.

Restrictions:

None

Operation:


```
temp ← (GPR[rs]31 || GPR[rs]31..0) + sign_extend(immediate)
if temp32 ≠ temp31 then
    SignalException(IntegerOverflow)
else
    GPR[rt] ← temp
endif
```

Exceptions:

Integer Overflow

Programming Notes:

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



Format: ADDIU *rt*, *rs*, *immediate*

MIPS32

Purpose: Add Immediate Unsigned Word

To add a constant to a 32-bit integer

Description: $GPR[rt] \leftarrow GPR[rs] + immediate$

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:

None

Operation:


```
temp ← GPR[rs] + sign_extend(immediate)
GPR[rt] ← 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. This instruction is appropriate for unsigned arithmetic, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as C language arithmetic.



Format: ADDU *rd*, *rs*, *rt*

MIPS32

Purpose: Add Unsigned Word

To add 32-bit integers

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

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:

None

Operation:

```
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. This instruction is appropriate for unsigned arithmetic, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as C language arithmetic.

Exceptions:

Coprocesor Unusable, Reserved Instruction

Programming Notes:

ALNV.PS is designed to be used with LUXC1 to load 8 bytes of data from any 4-byte boundary. For example:

```

/* Copy T2 bytes (a multiple of 16) of data T0 to T1, T0 unaligned, T1 aligned.
   Reads one dw beyond the end of T0. */
LUXC1    F0, 0(T0) /* set up by reading 1st src dw */
LI       T3, 0     /* index into src and dst arrays */
ADDIU   T4, T0, 8 /* base for odd dw loads */
ADDIU   T5, T1, -8/* base for odd dw stores */
LOOP:
LUXC1    F1, T3(T4)
ALNV.PS  F2, F0, F1, T0/* switch F0, F1 for little-endian */
SDC1     F2, T3(T1)
ADDIU   T3, T3, 16
LUXC1    F0, T3(T0)
ALNV.PS  F2, F1, F0, T0/* switch F1, F0 for little-endian */
BNE     T3, T2, LOOP
SDC1     F2, T3(T5)
DONE:

```

ALNV.PS is also useful with SUXC1 to store paired-single results in a vector loop to a possibly misaligned address:

```


/* T1[i] = T0[i] + F8, T0 aligned, T1 unaligned. */
   CVT.PS.S F8, F8, F8/* make addend paired-single */

/* Loop header computes 1st pair into F0, stores high half if T1 */
/* misaligned */

LOOP:
LDC1     F2, T3(T4)/* get T0[i+2]/T0[i+3] */
ADD.PS   F1, F2, F8/* compute T1[i+2]/T1[i+3] */
ALNV.PS  F3, F0, F1, T1/* align to dst memory */
SUXC1    F3, T3(T1)/* store to T1[i+0]/T1[i+1] */
ADDIU   T3, 16     /* i = i + 4 */
LDC1     F2, T3(T0)/* get T0[i+0]/T0[i+1] */
ADD.PS   F0, F2, F8/* compute T1[i+0]/T1[i+1] */
ALNV.PS  F3, F1, F0, T1/* align to dst memory */
BNE     T3, T2, LOOP
SUXC1    F3, T3(T5)/* store to T1[i+2]/T1[i+3] */

/* Loop trailer stores all or half of F0, depending on T1 alignment */

```



Format: AND *rd*, *rs*, *rt*

MIPS32

Purpose: And

To do a bitwise logical AND

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

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`

MIPS32

Purpose: And Immediate

To do a bitwise logical AND with a constant

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

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:

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

Exceptions:

None



Format: B offset

Assembly Idiom

Purpose: Unconditional Branch

To do an unconditional branch

Description: branch

B offset is the assembly idiom used to denote an unconditional branch. The actual instruction is interpreted by the hardware as BEQ r0, r0, offset.

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.

Restrictions:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:


I: target_offset \leftarrow sign_extend(offset || 0²)
I+1: PC \leftarrow PC + target_offset

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 addresses outside this range.



Format: BAL offset

Assembly Idiom

Purpose: Branch and Link

To do an unconditional PC-relative procedure call

Description: `procedure_call`

BAL offset is the assembly idiom used to denote an unconditional branch. The actual instruction is interpreted by the hardware as BGEZAL r0, offset.

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

An 18-bit signed offset (the 16-bit *offset* field shifted left 2bits) 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.

Restrictions:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

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 **UNPREDICTABLE**. 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: $\text{target_offset} \leftarrow \text{sign_extend}(\text{offset} \parallel 0^2)$
 $\text{GPR}[31] \leftarrow \text{PC} + 8$
I+1: $\text{PC} \leftarrow \text{PC} + \text{target_offset}$

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 addresses outside this range.



Format: BC1F offset (cc = 0 implied)
BC1F cc, offset

MIPS32
MIPS32

Purpose: Branch on FP False

To test an FP condition code and do a PC-relative conditional branch

Description: if FPConditionCode(cc) = 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 FP condition code bit *cc* is false (0), the program branches 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.

Restrictions:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

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*.

```

I:    condition ← FPConditionCode(cc) = 0
        target_offset ← (offset15)GPRLLEN-(16+2) || offset || 02
I+1:  if condition then
        PC ← PC + target_offset
        endif

```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

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 addresses outside this range


Historical Information:

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/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 and MIPS32 architectures add 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 and MIPS32.

In the MIPS I, II, and III architectures there must be at least one instruction between the compare instruction that sets

the condition code and the branch instruction that tests it. Hardware does not detect a violation of this restriction.



Format: BC1FL offset (cc = 0 implied)
BC1FL cc, offset

MIPS32
MIPS32

Purpose: Branch on FP False Likely

To test an FP condition code and make a PC-relative conditional branch; execute the instruction in the delay slot only if the branch is taken.

Description: if FPConditionCode(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), the program branches 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.

Restrictions:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

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*.

```

I:    condition ← FPConditionCode(cc) = 0
        target_offset ← (offset15)GPRLEN-(16+2) || offset || 02
I+1:  if condition then
        PC ← PC + target_offset
        else
        NullifyCurrentInstruction()
        endif

```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

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 addresses outside this range.

Software is strongly encouraged to avoid the use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS Architecture.

Some implementations always predict the branch will be taken, so there is a significant penalty if the branch is not taken. Software should only use this instruction when there is a very high probability (98% or more) that the branch will be taken. If the branch is not likely to be taken or if the probability of a taken branch is unknown, software is


encouraged to use the BC1F instruction instead.

Historical Information:

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/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 and MIPS32 architectures add 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 and MIPS32.

In the MIPS II and III architectures, 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.



Format: BC1T offset (cc = 0 implied)
BC1T cc, offset

MIPS32
MIPS32

Purpose: Branch on FP True

To test an FP condition code and do a PC-relative conditional branch

Description: if FPConditionCode(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), the program branches 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.

Restrictions:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

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*.

```

I:    condition ← FPConditionCode(cc) = 1
        target_offset ← (offset15)GPRLLEN-(16+2) || offset || 02
I+1:  if condition then
        PC ← PC + target_offset
        endif

```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

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 addresses outside this range.


Historical Information:

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/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 and MIPS32 architectures add 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 and MIPS32.

In the MIPS I, II, and III architectures there must be at least one instruction between the compare instruction that sets

the condition code and the branch instruction that tests it. Hardware does not detect a violation of this restriction.



Format: BC1TL offset (cc = 0 implied)
BC1TL cc, offset

MIPS32
MIPS32

Purpose: Branch on FP True Likely

To test an FP condition code and do a PC-relative conditional branch; execute the instruction in the delay slot only if the branch is taken.

Description: if FPConditionCode(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), the program branches 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.

Restrictions:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

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*.

```

I:      condition ← FPConditionCode(cc) = 1
          target_offset ← (offset15)GPRLEN-(16+2) || offset || 02
I+1:   if condition then
           PC ← PC + target_offset
         else
           NullifyCurrentInstruction()
         endif

```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

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 addresses outside this range.

Software is strongly encouraged to avoid the use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS Architecture.

Some implementations always predict the branch will be taken, so there is a significant penalty if the branch is not taken. Software should only use this instruction when there is a very high probability (98% or more) that the branch


will be taken. If the branch is not likely to be taken or if the probability of a taken branch is unknown, software is encouraged to use the BC1T instruction instead.

Historical Information:

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/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 and MIPS32 architectures add 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 and MIPS32.

In the MIPS II and III architectures, 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.



Format: BC2F offset (cc = 0 implied)
BC2F cc, offset

MIPS32
MIPS32

Purpose: Branch on COP2 False

To test a COP2 condition code and do a PC-relative conditional branch

Description: if COP2Condition(cc) = 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 COP2 condition specified by *cc* is false (0), the program branches to the effective target address after the instruction in the delay slot is executed.

Restrictions:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

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 BC2F, BC2FL, BC2T, and BC2TL have specific values for *tf* and *nd*.

```

I:    condition ← COP2Condition(cc) = 0
        target_offset ← (offset15)GPRLEN-(16+2) || offset || 02
I+1:  if condition then
        PC ← PC + target_offset
        endif


```

Exceptions:

Coprocessor Unusable, 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 addresses outside this range.



Format: BC2FL offset (cc = 0 implied)
BC2FL cc, offset

MIPS32
MIPS32

Purpose: Branch on COP2 False Likely

To test a COP2 condition code and make a PC-relative conditional branch; execute the instruction in the delay slot only if the branch is taken.

Description: if COP2Condition(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 COP2 condition specified by *cc* is false (0), the program branches 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:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

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 BC2F, BC2FL, BC2T, and BC2TL have specific values for *tf* and *nd*.

```

I:    condition ← COP2Condition(cc) = 0
        target_offset ← (offset15)GPRLLEN-(16+2) || offset || 02
I+1:  if condition then
        PC ← PC + target_offset
        else
        NullifyCurrentInstruction()
        endif

```

Exceptions:


Coprocessor Unusable, 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 addresses outside this range.

Software is strongly encouraged to avoid the use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS Architecture.

Some implementations always predict the branch will be taken, so there is a significant penalty if the branch is not taken. Software should only use this instruction when there is a very high probability (98% or more) that the branch will be taken. If the branch is not likely to be taken or if the probability of a taken branch is unknown, software is encouraged to use the BC2F instruction instead.



Format: BC2T offset (cc = 0 implied)
BC2T cc, offset

MIPS32
MIPS32

Purpose: Branch on COP2 True

To test a COP2 condition code and do a PC-relative conditional branch

Description: if COP2Condition(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 COP2 condition specified by *cc* is true (1), the program branches to the effective target address after the instruction in the delay slot is executed.

Restrictions:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

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 BC2F, BC2FL, BC2T, and BC2TL have specific values for *tf* and *nd*.

```

I:    condition ← COP2Condition(cc) = 1
        target_offset ← (offset15)GPRLEN-(16+2) || offset || 02
I+1:  if condition then
        PC ← PC + target_offset
        endif


```

Exceptions:

Coprocessor Unusable, 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 addresses outside this range.



Format: BC2TL offset (cc = 0 implied)
BC2TL cc, offset

MIPS32
MIPS32

Purpose: Branch on COP2 True Likely

To test a COP2 condition code and do a PC-relative conditional branch; execute the instruction in the delay slot only if the branch is taken.

Description: if COP2Condition(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 COP2 condition specified by *cc* is true (1), the program branches 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:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

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 BC2F, BC2FL, BC2T, and BC2TL have specific values for *tf* and *nd*.

```

I:    condition ← COP2Condition(cc) = 1
        target_offset ← (offset15)GPRLLEN-(16+2) || offset || 02
I+1: if condition then
        PC ← PC + target_offset
        else
        NullifyCurrentInstruction()
        endif

```

Exceptions:


Coprocessor Unusable, 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 addresses outside this range.

Software is strongly encouraged to avoid the use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS Architecture.

Some implementations always predict the branch will be taken, so there is a significant penalty if the branch is not taken. Software should only use this instruction when there is a very high probability (98% or more) that the branch will be taken. If the branch is not likely to be taken or if the probability of a taken branch is unknown, software is encouraged to use the BC2T instruction instead.



Format: BEQ *rs*, *rt*, *offset*

MIPS32

Purpose: Branch on Equal

To compare GPRs then do a PC-relative conditional branch

Description: if $GPR[rs] = GPR[rt]$ 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:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I:    target_offset ← sign_extend(offset || 02)
        condition ← (GPR[rs] = GPR[rt])
I+1:  if condition then
        PC ← PC + target_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 addresses outside this range.

BEQ *r0*, *r0* *offset*, expressed as B *offset*, is the assembly idiom used to denote an unconditional branch.



Format: BEQL *rs*, *rt*, *offset*

MIPS32

Purpose: Branch on Equal Likely

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

Description: if $GPR[rs] = GPR[rt]$ 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:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I:    target_offset ← sign_extend(offset || 02)
        condition ← (GPR[rs] = GPR[rt])
I+1:  if condition then
        PC ← PC + target_offset
        else
        NullifyCurrentInstruction()
        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 addresses outside this range.

Software is strongly encouraged to avoid the use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS Architecture.

Some implementations always predict the branch will be taken, so there is a significant penalty if the branch is not taken. Software should only use this instruction when there is a very high probability (98% or more) that the branch will be taken. If the branch is not likely to be taken or if the probability of a taken branch is unknown, software is encouraged to use the BEQ instruction instead.

Historical Information:

In the MIPS I architecture, this instruction signaled a Reserved Instruction Exception.



Format: BGEZ *rs*, *offset*

MIPS32

Purpose: Branch on Greater Than or Equal to Zero

To test a GPR then do a PC-relative conditional branch

Description: if $GPR[rs] \geq 0$ 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:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I:    target_offset ← sign_extend(offset || 02)
        condition ← GPR[rs] ≥ 0GPRLEN
I+1:  if condition then
            PC ← PC + target_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 addresses outside this range.



Format: BGEZAL *rs*, *offset*

MIPS32

Purpose: Branch on Greater Than or Equal to Zero and Link

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

Description: if $GPR[rs] \geq 0$ 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 continues 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:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

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

Operation:

```

I:    target_offset ← sign_extend(offset || 02)
        condition ← GPR[rs] ≥ 0GPRLEN
        GPR[31] ← PC + 8
I+1:  if condition then
        PC ← PC + target_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 addresses outside this range.

BGEZAL *r0*, *offset*, expressed as BAL *offset*, is the assembly idiom used to denote a PC-relative branch and link. BAL is used in a manner similar to JAL, but provides PC-relative addressing and a more limited target PC range.



Format: BGEZALL *rs*, *offset*

MIPS32

Purpose: Branch on Greater Than or Equal to Zero and Link Likely

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

Description: if $GPR[rs] \geq 0$ 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 continues 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 reexecuted. The result of executing such an instruction is **UNPREDICTABLE**. This restriction permits an exception handler to resume execution by reexecuting the branch when an exception occurs in the branch delay slot.

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I:    target_offset ← sign_extend(offset || 02)
        condition ← GPR[rs] ≥ 0GPRLEN
        GPR[31] ← PC + 8
I+1:  if condition then
        PC ← PC + target_offset
        else
        NullifyCurrentInstruction()
        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 addresses outside this range.

Software is strongly encouraged to avoid the use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS Architecture.

Some implementations always predict the branch will be taken, so there is a significant penalty if the branch is not taken. Software should only use this instruction when there is a very high probability (98% or more) that the branch will be taken. If the branch is not likely to be taken or if the probability of a taken branch is unknown, software is encouraged to use the BGEZAL instruction instead.

Historical Information:

In the MIPS I architecture, this instruction signaled a Reserved Instruction Exception.



Format: BGEZL *rs*, *offset*

MIPS32

Purpose: Branch on Greater Than or Equal to Zero Likely

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

Description: if $GPR[rs] \geq 0$ 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:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I:    target_offset ← sign_extend(offset || 02)
        condition ← GPR[rs] ≥ 0GPRLEN
I+1:  if condition then
            PC ← PC + target_offset
        else
            NullifyCurrentInstruction()
        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 addresses outside this range.

Software is strongly encouraged to avoid the use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS Architecture.

Some implementations always predict the branch will be taken, so there is a significant penalty if the branch is not taken. Software should only use this instruction when there is a very high probability (98% or more) that the branch will be taken. If the branch is not likely to be taken or if the probability of a taken branch is unknown, software is encouraged to use the BGEZ instruction instead.

Historical Information:

In the MIPS I architecture, this instruction signaled a Reserved Instruction Exception.



Format: BGTZ *rs*, *offset*

MIPS32

Purpose: Branch on Greater Than Zero

To test a GPR then do a PC-relative conditional branch

Description: if $GPR[rs] > 0$ 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:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I:    target_offset ← sign_extend(offset || 02)
        condition ← GPR[rs] > 0GPRLEN
I+1:  if condition then
            PC ← PC + target_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 addresses outside this range.



Format: BGTZL *rs*, *offset*

MIPS32

Purpose: Branch on Greater Than Zero Likely

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

Description: if $GPR[rs] > 0$ 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:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I:    target_offset ← sign_extend(offset || 02)
        condition ← GPR[rs] > 0GPRLEN
I+1:  if condition then
            PC ← PC + target_offset
        else
            NullifyCurrentInstruction()
        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 addresses outside this range.

Software is strongly encouraged to avoid the use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS Architecture.

Some implementations always predict the branch will be taken, so there is a significant penalty if the branch is not taken. Software should only use this instruction when there is a very high probability (98% or more) that the branch will be taken. If the branch is not likely to be taken or if the probability of a taken branch is unknown, software is encouraged to use the BGTZ instruction instead.

Historical Information:

In the MIPS I architecture, this instruction signaled a Reserved Instruction Exception.



Format: BLEZ *rs*, *offset*

MIPS32

Purpose: Branch on Less Than or Equal to Zero

To test a GPR then do a PC-relative conditional branch

Description: if $GPR[rs] \leq 0$ 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:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I:    target_offset ← sign_extend(offset || 02)
        condition ← GPR[rs] ≤ 0GPRLEN
I+1:  if condition then
            PC ← PC + target_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 addresses outside this range.



Format: BLEZL *rs*, *offset*

MIPS32

Purpose: Branch on Less Than or Equal to Zero Likely

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

Description: if $GPR[rs] \leq 0$ 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:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I:    target_offset ← sign_extend(offset || 02)
        condition ← GPR[rs] ≤ 0GPRLEN
I+1:  if condition then
            PC ← PC + target_offset
        else
            NullifyCurrentInstruction()
        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 addresses outside this range.

Software is strongly encouraged to avoid the use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS Architecture.

Some implementations always predict the branch will be taken, so there is a significant penalty if the branch is not taken. Software should only use this instruction when there is a very high probability (98% or more) that the branch will be taken. If the branch is not likely to be taken or if the probability of a taken branch is unknown, software is encouraged to use the BLEZ instruction instead.

Historical Information:

In the MIPS I architecture, this instruction signaled a Reserved Instruction Exception.



Format: BLTZ *rs*, *offset*

MIPS32

Purpose: Branch on Less Than Zero

To test a GPR then do a PC-relative conditional branch

Description: if $GPR[rs] < 0$ 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:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I:   target_offset ← sign_extend(offset || 02)
        condition ← GPR[rs] < 0GPRLEN
I+1: if condition then
        PC ← PC + target_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 addresses outside this range.



Format: BLTZAL *rs*, *offset*

MIPS32

Purpose: Branch on Less Than Zero and Link

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

Description: if GPR[*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 continues 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 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 reexecuted. The result of executing such an instruction is **UNPREDICTABLE**. This restriction permits an exception handler to resume execution by reexecuting the branch when an exception occurs in the branch delay slot.

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I:    target_offset ← sign_extend(offset || 02)
        condition ← GPR[rs] < 0GPRLEN
        GPR[31] ← PC + 8
I+1:  if condition then
        PC ← PC + target_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 addresses outside this range.



Format: BLTZALL *rs*, *offset*

MIPS32

Purpose: Branch on Less Than Zero and Link Likely

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

Description: if $GPR[rs] < 0$ 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 continues 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 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 reexecuted. The result of executing such an instruction is **UNPREDICTABLE**. This restriction permits an exception handler to resume execution by reexecuting the branch when an exception occurs in the branch delay slot.

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I:    target_offset ← sign_extend(offset || 02)
        condition ← GPR[rs] < 0GPRLEN
        GPR[31] ← PC + 8
I+1:  if condition then
            PC ← PC + target_offset
        else
            NullifyCurrentInstruction()
        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 addresses outside this range.

Software is strongly encouraged to avoid the use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS Architecture.

Some implementations always predict the branch will be taken, so there is a significant penalty if the branch is not taken. Software should only use this instruction when there is a very high probability (98% or more) that the branch will be taken. If the branch is not likely to be taken or if the probability of a taken branch is unknown, software is encouraged to use the BLTZAL instruction instead.

Historical Information:

In the MIPS I architecture, this instruction signaled a Reserved Instruction Exception.



Format: BLTZL *rs*, *offset*

MIPS32

Purpose: Branch on Less Than Zero Likely

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

Description: if $GPR[rs] < 0$ 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:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I:    target_offset ← sign_extend(offset || 02)
        condition ← GPR[rs] < 0GPRLEN
I+1:  if condition then
        PC ← PC + target_offset
        else
        NullifyCurrentInstruction()
        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 addresses outside this range.

Software is strongly encouraged to avoid the use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS Architecture.

Some implementations always predict the branch will be taken, so there is a significant penalty if the branch is not taken. Software should only use this instruction when there is a very high probability (98% or more) that the branch will be taken. If the branch is not likely to be taken or if the probability of a taken branch is unknown, software is encouraged to use the BLTZ instruction instead.

Historical Information:

In the MIPS I architecture, this instruction signaled a Reserved Instruction Exception.



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

MIPS32

Purpose: Branch on Not Equal

To compare GPRs then do a PC-relative conditional branch

Description: if $GPR[rs] \neq GPR[rt]$ 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:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I:    target_offset ← sign_extend(offset || 02)
        condition ← (GPR[rs] ≠ GPR[rt])
I+1:  if condition then
            PC ← PC + target_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 addresses outside this range.



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

MIPS32

Purpose: Branch on Not Equal Likely

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

Description: if $GPR[rs] \neq GPR[rt]$ 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:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I:    target_offset ← sign_extend(offset || 02)
        condition ← (GPR[rs] ≠ GPR[rt])
I+1:  if condition then
            PC ← PC + target_offset
        else
            NullifyCurrentInstruction()
        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 addresses outside this range.

Software is strongly encouraged to avoid the use of the Branch Likely instructions, as they will be removed from a future revision of the MIPS Architecture.

Some implementations always predict the branch will be taken, so there is a significant penalty if the branch is not taken. Software should only use this instruction when there is a very high probability (98% or more) that the branch will be taken. If the branch is not likely to be taken or if the probability of a taken branch is unknown, software is encouraged to use the BNE instruction instead.

Historical Information:

In the MIPS I architecture, this instruction signaled a Reserved Instruction Exception.



Format: BREAK

MIPS32

Purpose: Breakpoint

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



Format: C.cond.fmt
 C.cond.S fs, ft (cc = 0 implied) MIPS32
 C.cond.D fs, ft (cc = 0 implied) MIPS32
 C.cond.PS fs, ft(cc = 0 implied) MIPS64, MIPS32 Release 2
 C.cond.S cc, fs, ft MIPS32
 C.cond.D cc, fs, ft MIPS32
 C.cond.PS cc, fs, ft MIPS64, MIPS32 Release 2

Purpose: Floating Point Compare

To compare FP values and record the Boolean result in a condition code

Description: $FPUConditionCode(cc) \leftarrow FPR[fs] \text{ compare_cond } FPR[ft]$

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, the result is true; otherwise, the result is false. If no exception is taken, the result is written into condition code *CC*; true is 1 and false is 0.

c.cond.PS compares the upper and lower halves of FPR *fs* and FPR *ft* independently and writes the results into condition codes *CC + 1* and *CC* respectively. The *CC* number must be even. If the number is not even the operation of the instruction is **UNPREDICTABLE**.

If one of the values is an SNaN, or *cond₃* is set and at least one of the values is a QNaN, an Invalid Operation condition is raised and 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*.

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 among the 16 comparison predicates. The Boolean result of the instruction is obtained by substituting the Boolean value of each ordering relation for the two FP values in the equation. If the *equal* relation is true, for example, then all four example predicates above yield a true result. If the *unordered* relation is true then only the final predicate, *unordered or equal*, yields a true result.

Logical negation of a compare result allows eight distinct comparisons to test for the 16 predicates as shown in Table 3.25. 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, and the second predicate must be false, and vice versa. (Note that the False predicate is never true and False/True do not follow the normal pattern.)

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 can be made with the Branch on FP True (BC1T) instruction and the truth of the second can be made with Branch on FP False (BC1F).

Table 3.26 shows another set of eight compare operations, distinguished by a *cond₃* value of 1 and testing the same 16 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*, an Invalid Operation

exception occurs.

Table 3.25 FPU Comparisons Without Special Operand Exceptions

Instruction	Comparison Predicate				Comparison CC Result		Instruction		
Cond Mnemonic	Name of Predicate and Logically Negated Predicate (Abbreviation)	Relation Values				If Predicate Is True	Inv Op Excp. if QNaN?	Condition Field	
		>	<	=	?			3	2..0
F	False [this predicate is always False]	F	F	F	F	F	No	0	0
	True (T)	T	T	T	T				
UN	Unordered	F	F	F	T	T	No	0	1
	Ordered (OR)	T	T	T	F	F			2
EQ	Equal	F	F	T	F	T	No	0	3
	Not Equal (NEQ)	T	T	F	T	F			4
UEQ	Unordered or Equal	F	F	T	T	T	No	0	5
	Ordered or Greater Than or Less Than (OGL)	T	T	F	F	F			6
OLT	Ordered or Less Than	F	T	F	F	T	No	0	7
	Unordered or Greater Than or Equal (UGE)	T	F	T	T	F			
ULT	Unordered or Less Than	F	T	F	T	T	No	0	
	Ordered or Greater Than or Equal (OGE)	T	F	T	F	F			
OLE	Ordered or Less Than or Equal	F	T	T	F	T	No	0	
	Unordered or Greater Than (UGT)	T	F	F	T	F			
ULE	Unordered or Less Than or Equal	F	T	T	T	T	No	0	
	Ordered or Greater Than (OGT)	T	F	F	F	F			

Key: ? = unordered, > = greater than, < = less than, = is equal, T = True, F = False

Table 3.26 FPU Comparisons With Special Operand Exceptions for QNaNs

Instruction Cond Mnemonic	Comparison Predicate Name of Predicate and Logically Negated Predicate (Abbreviation)	Relation Values				Comparison CC Result If Predicate Is True	Inv Op Excp If QNaN?	Instruction	
		>	<	=	?			Condition 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				
NGLE	Not Greater Than or Less Than or Equal	F	F	F	T	T			1
	Greater Than or Less Than or Equal (GLE)	T	T	T	F				
SEQ	Signaling Equal	F	F	T	F	T			2
	Signaling Not Equal (SNE)	T	T	F	T				
NGL	Not Greater Than or Less Than	F	F	T	T	T			3
	Greater Than or Less Than (GL)	T	T	F	F				
LT	Less Than	F	T	F	F	T			4
	Not Less Than (NLT)	T	F	T	T				
NGE	Not Greater Than or Equal	F	T	F	T	T			5
	Greater Than or Equal (GE)	T	F	T	F				
LE	Less Than or Equal	F	T	T	F	T			6
	Not Less Than or Equal (NLE)	T	F	F	T				
NGT	Not Greater Than	F	T	T	T	T			7
	Greater Than (GT)	T	F	F	F				

Key: ? = unordered, > = greater than, < = less than, = is equal, T = True, F = False

Restrictions:

The fields *fs* and *ft* must specify FPRs valid for operands of type *fmt*; if they are not valid, the result is **UNPREDICTABLE**.

The operands must be values in format *fmt*; if they are not, the result is **UNPREDICTABLE** and the value of the operand FPRs becomes **UNPREDICTABLE**.

The result of C.cond.PS is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode, or if the condition code number is odd.

Operation:

```

if SNaN(ValueFPR(fs, fmt)) or SNaN(ValueFPR(ft, fmt)) or
   QNaN(ValueFPR(fs, fmt)) or QNaN(ValueFPR(ft, fmt)) then
  less ← false
  equal ← false
  unordered ← true
  if (SNaN(ValueFPR(fs,fmt)) or SNaN(ValueFPR(ft,fmt))) or
     (cond3 and (QNaN(ValueFPR(fs,fmt)) or QNaN(ValueFPR(ft,fmt)))) then
    SignalException(InvalidOperation)
  endif
endif
else
  less ← ValueFPR(fs, fmt) <fmt ValueFPR(ft, fmt)
  equal ← ValueFPR(fs, fmt) =fmt ValueFPR(ft, fmt)
  unordered ← false
endif

```

```

condition ← (cond2 and less) or (cond1 and equal)
           or (cond0 and unordered)
SetFPConditionCode(cc, condition)

```

For c.cond.PS, the pseudo code above is repeated for both halves of the operand registers, treating each half as an independent single-precision values. Exceptions on the two halves are logically ORed and reported together. The results of the lower half comparison are written to condition code CC; the results of the upper half comparison are written to condition code CC+1.

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation, Invalid Operation


Programming Notes:

FP computational instructions, including compare, that receive an operand value of Signaling NaN raise the Invalid Operation condition. 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

```



Format: CACHE op, offset(base)

MIPS32

Purpose: Perform Cache Operation

To perform the cache operation specified by op.


Description:

The 16-bit offset is sign-extended and added to the contents of the base register to form an effective address. The effective address is used in one of the following ways based on the operation to be performed and the type of cache as described in the following table.

Table 3.27 Usage of Effective Address

Operation Requires an	Type of Cache	Usage of Effective Address
Address	Virtual	The effective address is used to address the cache. An address translation may or may not be performed on the effective address (with the possibility that a TLB Refill or TLB Invalid exception might occur)
Address	Physical	The effective address is translated by the MMU to a physical address. The physical address is then used to address the cache
Index	N/A	<p>The effective address is translated by the MMU to a physical address. It is implementation dependent whether the effective address or the translated physical address is used to index the cache. As such, a kseg0 address should always be used for cache operations that require an index. See the Programming Notes section below.</p> <p>Assuming that the total cache size in bytes is CS, the associativity is A, and the number of bytes per tag is BPT, the following calculations give the fields of the address which specify the way and the index:</p> $\begin{aligned} \text{OffsetBit} &\leftarrow \text{Log}_2(\text{BPT}) \\ \text{IndexBit} &\leftarrow \text{Log}_2(\text{CS} / \text{A}) \\ \text{WayBit} &\leftarrow \text{IndexBit} + \text{Ceiling}(\text{Log}_2(\text{A})) \\ \text{Way} &\leftarrow \text{Addr}_{\text{WayBit}-1.. \text{IndexBit}} \\ \text{Index} &\leftarrow \text{Addr}_{\text{IndexBit}-1.. \text{OffsetBit}} \end{aligned}$ <p>For a direct-mapped cache, the Way calculation is ignored and the Index value fully specifies the cache tag. This is shown symbolically in the figure below.</p>

Figure 3.2 Usage of Address Fields to Select Index and Way



A TLB Refill and TLB Invalid (both with cause code equal TLBL) exception can occur on any operation. For index operations (where the address is used to index the cache but need not match the cache tag) software should use unmapped addresses to avoid TLB exceptions. This instruction never causes TLB Modified exceptions nor TLB Refill exceptions with a cause code of TLBS.

The effective address may be an arbitrarily-aligned by address. The CACHE instruction never causes an Address Error Exception due to a non-aligned address.

A Cache Error exception may occur as a by-product of some operations performed by this instruction. For example, if a Writeback operation detects a cache or bus error during the processing of the operation, that error is reported via a Cache Error exception. Similarly, a Bus Error Exception may occur if a bus operation invoked by this instruction is terminated in an error. However, cache error exceptions must not be triggered by an Index Load Tag or Index Store tag operation, as these operations are used for initialization and diagnostic purposes.

An Address Error Exception (with cause code equal AdEL) may occur if the effective address references a portion of the kernel address space which would normally result in such an exception. It is implementation dependent whether such an exception does occur.

It is implementation dependent whether a data watch is triggered by a cache instruction whose address matches the Watch register address match conditions.

The CACHE instruction and the memory transactions which are sourced by the CACHE instruction, such as cache refill or cache writeback, obey the ordering and completion rules of the SYNC instruction.

Bits [17:16] of the instruction specify the cache on which to perform the operation, as follows:

Table 3.28 Encoding of Bits[17:16] of CACHE Instruction

Code	Name	Cache
0b00	I	Primary Instruction
0b01	D	Primary Data or Unified Primary
0b10	T	Tertiary
0b11	S	Secondary

Bits [20:18] of the instruction specify the operation to perform. To provide software with a consistent base of cache operations, certain encodings must be supported on all processors. The remaining encodings are recommended

For implementations which implement multiple level of caches and where the hardware maintains the smaller cache as a proper subset of a larger cache (every address which is resident in the smaller cache is also resident in the larger cache; also known as the inclusion property), it is recommended that the CACHE instructions which operate on the larger, outer-level cache; should first operate on the smaller, inner-level cache. For example, a Hit_Writeback_Invalidate operation targeting the Secondary cache, should first operate on the primary data cache first. If the CACHE instruction implementation does not follow this policy then any software which flushes the caches must mimic this behavior. That is, the software sequences must first operate on the inner cache then operate on the outer cache. The software must place a SYNC instruction after the CACHE instruction whenever there are possible writebacks from the inner cache to ensure that the writeback data is resident in the outer cache before operating on the outer cache. If neither the CACHE instruction implementation nor the software cache flush sequence follow this policy, then the inclusion property of the caches can be broken, which might be a condition that the cache management hardware can not properly deal with.

For implementations which implement multiple level of caches without the inclusion property, the use of a SYNC instruction after the CACHE instruction is still needed whenever writeback data has to be resident in the next level of memory hierarchy.

For multiprocessor implementations that maintain coherent caches, some of the Hit type of CACHE instruction operations may optionally affect all coherent caches within the implementation. If the effective address uses a coherent Cache Coherency Attribute (CCA), then the operation is *globalized*, meaning it is broadcast to all of the coherent caches within the system. If the effective address does not use one of the coherent CCAs, there is no broadcast of the operation. If multiple levels of caches are to be affected by one CACHE instruction, all of the affected cache levels must be processed in the same manner - either all affected cache levels use the globalized behavior or all affected cache levels use the non-globalized behavior.

Table 3.29 Encoding of Bits [20:18] of the CACHE Instruction

Code	Caches	Name	Effective Address Operand Type	Operation	Compliance Implemented
0b000	I	Index Invalidate	Index	Set the state of the cache block at the specified index to invalid. This required encoding may be used by software to invalidate the entire instruction cache by stepping through all valid indices.	Required
	D	Index Writeback Invalidate / Index Invalidate	Index	For a write-back cache: If the state of the cache block at the specified index is valid and dirty, write the block back to the memory address specified by the cache tag. After that operation is completed, set the state of the cache block to invalid. If the block is valid but not dirty, set the state of the block to invalid.	Required
	S, T	Index Writeback Invalidate / Index Invalidate	Index	For a write-through cache: Set the state of the cache block at the specified index to invalid. This required encoding may be used by software to invalidate the entire data cache by stepping through all valid indices. Note that Index Store Tag should be used to initialize the cache at power up.	Required if S, T cache is implemented
0b001	All	Index Load Tag	Index	Read the tag for the cache block at the specified index into the <i>TagLo</i> and <i>TagHi</i> Coprocessor 0 registers. If the <i>DataLo</i> and <i>DataHi</i> registers are implemented, also read the data corresponding to the byte index into the <i>DataLo</i> and <i>DataHi</i> registers. This operation must not cause a Cache Error Exception. The granularity and alignment of the data read into the <i>DataLo</i> and <i>DataHi</i> registers is implementation-dependent, but is typically the result of an aligned access to the cache, ignoring the appropriate low-order bits of the byte index.	Recommended

Table 3.29 Encoding of Bits [20:18] of the CACHE Instruction (Continued)

Code	Caches	Name	Effective Address Operand Type	Operation	Compliance Implemented
0b010	All	Index Store Tag	Index	Write the tag for the cache block at the specified index from the <i>TagLo</i> and <i>TagHi</i> Coprocessor 0 registers. This operation must not cause a Cache Error Exception. This required encoding may be used by software to initialize the entire instruction or data caches by stepping through all valid indices. Doing so requires that the <i>TagLo</i> and <i>TagHi</i> registers associated with the cache be initialized first.	Required
0b011	All	Implementation Dependent	Unspecified	Available for implementation-dependent operation.	Optional
0b100	I, D	Hit Invalidate	Address	If the cache block contains the specified address, set the state of the cache block to invalid. This required encoding may be used by software to invalidate a range of addresses from the instruction cache by stepping through the address range by the line size of the cache.	Required (Instruction Cache Encoding Only), Recommended otherwise
	S, T	Hit Invalidate	Address	In multiprocessor implementations with coherent caches, the operation may optionally be broadcast to all coherent caches within the system.	Optional, if <i>Hit_Invalidate_D</i> is implemented, the S and T variants are recommended.
0b101	I	Fill	Address	Fill the cache from the specified address.	Recommended
	D	Hit Writeback Invalidate / Hit Invalidate	Address	For a write-back cache: If the cache block contains the specified address and it is valid and dirty, write the contents back to memory. After that operation is completed, set the state of the cache block to invalid. If the block is valid but not dirty, set the state of the block to invalid. For a write-through cache: If the cache block contains the specified address, set the state of the cache block to invalid. This required encoding may be used by software to invalidate a range of addresses from the data cache by stepping through the address range by the line size of the cache.	Required
	S, T	Hit Writeback Invalidate / Hit Invalidate	Address	In multiprocessor implementations with coherent caches, the operation may optionally be broadcast to all coherent caches within the system.	Required if S, T cache is implemented

Table 3.29 Encoding of Bits [20:18] of the CACHE Instruction (Continued)

Code	Caches	Name	Effective Address Operand Type	Operation	Compliance Implemented
0b110	D	Hit Writeback	Address	<p>If the cache block contains the specified address and it is valid and dirty, write the contents back to memory. After the operation is completed, leave the state of the line valid, but clear the dirty state. For a write-through cache, this operation may be treated as a nop.</p> <p>In multiprocessor implementations with coherent caches, the operation may optionally be broadcast to all coherent caches within the system.</p>	Recommended
	S, T	Hit Writeback	Address		Optional, if Hit_Writeback_D is implemented, the S and T variants are recommended.
0b111	I, D	Fetch and Lock	Address	<p>If the cache does not contain the specified address, fill it from memory, performing a writeback if required, and set the state to valid and locked. If the cache already contains the specified address, set the state to locked. In set-associative or fully-associative caches, the way selected on a fill from memory is implementation dependent.</p> <p>The lock state may be cleared by executing an Index Invalidate, Index Writeback Invalidate, Hit Invalidate, or Hit Writeback Invalidate operation to the locked line, or via an Index Store Tag operation to the line that clears the lock bit. Note that clearing the lock state via Index Store Tag is dependent on the implementation-dependent cache tag and cache line organization, and that Index and Index Writeback Invalidate operations are dependent on cache line organization. Only Hit and Hit Writeback Invalidate operations are generally portable across implementations.</p> <p>It is implementation dependent whether a locked line is displaced as the result of an external invalidate or intervention that hits on the locked line. Software must not depend on the locked line remaining in the cache if an external invalidate or intervention would invalidate the line if it were not locked.</p> <p>It is implementation dependent whether a Fetch and Lock operation affects more than one line. For example, more than one line around the referenced address may be fetched and locked. It is recommended that only the single line containing the referenced address be affected.</p>	Recommended

Restrictions:

The operation of this instruction is **UNDEFINED** for any operation/cache combination that is not implemented.

The operation of this instruction is **UNDEFINED** if the operation requires an address, and that address is uncacheable.

The operation of the instruction is **UNPREDICTABLE** if the cache line that contains the CACHE instruction is the target of an invalidate or a writeback invalidate.

If this instruction is used to lock all ways of a cache at a specific cache index, the behavior of that cache to subsequent cache misses to that cache index is **UNDEFINED**.

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

Any use of this instruction that can cause cacheline writebacks should be followed by a subsequent SYNC instruction to avoid hazards where the writeback data is not yet visible at the next level of the memory hierarchy.

Operation:

```
vAddr ← GPR[base] + sign_extend(offset)
(pAddr, uncached) ← AddressTranslation(vAddr, DataReadReference)
CacheOp(op, vAddr, pAddr)
```

Exceptions:

TLB Refill Exception.

TLB Invalid Exception

Coprocessor Unusable Exception

Address Error Exception

Cache Error Exception

Bus Error Exception

Programming Notes:

For cache operations that require an index, it is implementation dependent whether the effective address or the translated physical address is used as the cache index. Therefore, the index value should always be converted to a kseg0 address by ORing the index with 0x80000000 before being used by the cache instruction. For example, the following code sequence performs a data cache Index Store Tag operation using the index passed in GPR a0:

```
li    a1, 0x80000000    /* Base of kseg0 segment */
or    a0, a0, a1        /* Convert index to kseg0 address */
cache DCIndexStTag, 0(a1) /* Perform the index store tag operation */
```


31	26 25	21 20	16 15	11 10	6 5	0
COP1 010001	fmt	0 00000	fs	fd	CEIL.L 001010	
6	5	5	5	5	6	

Format: CEIL.L.fmt
 CEIL.L.S fd, fs MIPS64, MIPS32 Release 2 CEIL.L.D fd, fs MIPS64, MIPS32 Release 2

Purpose: Fixed Point Ceiling Convert to Long Fixed Point

To convert an FP value to 64-bit fixed point, rounding up

Description: $FPR[fd] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in 64-bit long fixed point format and 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, an IEEE Invalid Operation condition exists, and 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*.

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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of this instruction is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:


$\text{StoreFPR}(fd, L, \text{ConvertFmt}(\text{ValueFPR}(fs, \text{fmt}), \text{fmt}, L))$

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact, Overflow



Format: CEIL.W.fmt
 CEIL.W.S fd, fs
 CEIL.W.D fd, fs

MIPS32
MIPS32

Purpose: Floating Point Ceiling Convert to Word Fixed Point

To convert an FP value to 32-bit fixed point, rounding up

Description: $FPR[fd] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in 32-bit word fixed point format and 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, an IEEE Invalid Operation condition exists, and 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*.

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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Operation:


`StoreFPR(fd, W, ConvertFmt(ValueFPR(fs, fmt), fmt, W))`

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact, Overflow



Format: CFC1 *rt*, *fs*

MIPS32

Purpose: Move Control Word From Floating Point

To copy a word from an FPU control register to a GPR

Description: $GPR[rt] \leftarrow FP_Control[fs]$

Copy the 32-bit word from FP (coprocessor 1) control register *fs* into GPR *rt*.

Restrictions:

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

Operation:

```

if fs = 0 then
    temp ← FIR
elseif fs = 25 then
    temp ← 024 || FCSR31..25 || FCSR23
elseif fs = 26 then
    temp ← 014 || FCSR17..12 || 05 || FCSR6..2 || 02
elseif fs = 28 then
    temp ← 020 || FCSR11..7 || 04 || FCSR24 || FCSR1..0
elseif fs = 31 then
    temp ← FCSR
else
    temp ← UNPREDICTABLE
endif
GPR[rt] ← temp

```


Exceptions:

Coprocessor Unusable, Reserved Instruction

Historical Information:

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

MIPS V and MIPS32 introduced the three control registers that access portions of FCSR. These registers were not available in MIPS I, II, III, or IV.



Format: CFC2 rt, Impl

MIPS32

The syntax shown above is an example using CFC1 as a model. The specific syntax is implementation dependent.

Purpose: Move Control Word From Coprocessor 2

To copy a word from a Coprocessor 2 control register to a GPR

Description: $GPR[rt] \leftarrow CP2CCR[Impl]$

Copy the 32-bit word from the Coprocessor 2 control register denoted by the *Impl* field. The interpretation of the *Impl* field is left entirely to the Coprocessor 2 implementation and is not specified by the architecture.

Restrictions:


The result is **UNPREDICTABLE** if *Impl* specifies a register that does not exist.

Operation:

```
temp ← CP2CCR[Impl]
GPR[rt] ← temp
```

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: CLO rd, rs

MIPS32

Purpose: Count Leading Ones in Word

To count the number of leading ones in a word

Description: $GPR[rd] \leftarrow \text{count_leading_ones } GPR[rs]$

Bits 31..0 of GPR *rs* are scanned from most significant to least significant bit. The number of leading ones is counted and the result is written to GPR *rd*. If all of bits **31..0** were set in GPR *rs*, the result written to GPR *rd* is 32.

Restrictions:


To be compliant with the MIPS32 and MIPS64 Architecture, software must place the same GPR number in both the *rt* and *rd* fields of the instruction. The operation of the instruction is **UNPREDICTABLE** if the *rt* and *rd* fields of the instruction contain different values.

Operation:

```
temp ← 32
for i in 31 .. 0
  if GPR[rs]i = 0 then
    temp ← 31 - i
    break
  endif
endfor
GPR[rd] ← temp
```

Exceptions:

None



Format: CLZ *rd*, *rs*

MIPS32

Purpose: Count Leading Zeros in Word

Count the number of leading zeros in a word

Description: $GPR[rd] \leftarrow \text{count_leading_zeros } GPR[rs]$

Bits **31..0** of GPR *rs* are scanned from most significant to least significant bit. The number of leading zeros is counted and the result is written to GPR *rd*. If no bits were set in GPR *rs*, the result written to GPR *rd* is 32.

Restrictions:


To be compliant with the MIPS32 and MIPS64 Architecture, software must place the same GPR number in both the *rt* and *rd* fields of the instruction. The operation of the instruction is **UNPREDICTABLE** if the *rt* and *rd* fields of the instruction contain different values.

Operation:

```
temp ← 32
for i in 31 .. 0
  if GPR[rs]i = 1 then
    temp ← 31 - i
    break
  endif
endfor
GPR[rd] ← temp
```

Exceptions:

None



Format: COP2 func

MIPS32

Purpose: Coprocessor Operation to Coprocessor 2

To perform an operation to Coprocessor 2

Description: `CoprocessorOperation(2, cofun)`

An implementation-dependent operation is performed to Coprocessor 2, with the *cofun* value passed as an argument. The operation may specify and reference internal coprocessor registers, and may change the state of the coprocessor conditions, but does not modify state within the processor. Details of coprocessor operation and internal state are described in the documentation for each Coprocessor 2 implementation.


Restrictions:

Operation:

`CoprocessorOperation(2, cofun)`

Exceptions:

Coprocessor Unusable
Reserved Instruction



Format: CTC1 rt, fs

MIPS32

Purpose: Move Control Word to Floating Point

To copy a word from a GPR to an FPU control register

Description: FP_Control[fs] ← GPR[rt]

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

Writing to the floating point *Control/Status* register, the *FCSR*, causes the appropriate exception if any *Cause* bit and its corresponding *Enable* bit are both set. The register is written before the exception occurs. Writing to *FEXR* to set a cause bit whose enable bit is already set, or writing to *FENR* to set an enable bit whose cause bit is already set causes the appropriate exception. The register is written before the exception occurs and the *EPC* register contains the address of the CTC1 instruction.

Restrictions:

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

Operation:

```

temp ← GPR[rt]31..0
if fs = 25 then /* FCCR */
  if temp31..8 ≠ 024 then
    UNPREDICTABLE
  else
    FCSR ← temp7..1 || FCSR24 || temp0 || FCSR22..0
  endif
elseif fs = 26 then /* FEXR */
  if temp31..18 ≠ 0 or temp11..7 ≠ 0 or temp2..0 ≠ 0 then
    UNPREDICTABLE
  else
    FCSR ← FCSR31..18 || temp17..12 || FCSR11..7 ||
      temp6..2 || FCSR1..0
  endif
elseif fs = 28 then /* FENR */
  if temp31..12 ≠ 0 or temp6..3 ≠ 0 then
    UNPREDICTABLE
  else
    FCSR ← FCSR31..25 || temp2 || FCSR23..12 || temp11..7
      || FCSR6..2 || temp1..0
  endif
elseif fs = 31 then /* FCSR */
  if (FCSRImpl field is not implemented) and(temp22..18 ≠ 0) then
    UNPREDICTABLE
  elseif (FCSRImpl field is implemented) and temp20..18 ≠ 0 then
    UNPREDICTABLE
  else
    FCSR ← temp
  endif
else

```



```
    UNPREDICTABLE  
endif
```

CheckFPEException() **Exceptions:**

Coprocessor Unusable, Reserved Instruction


Floating Point Exceptions:

Unimplemented Operation, Invalid Operation, Division-by-zero, Inexact, Overflow, Underflow

Historical Information:

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

MIPS V and MIPS32 introduced the three control registers that access portions of FCSR. These registers were not available in MIPS I, II, III, or IV.



Format: CTC2 *rt*, *Impl*

MIPS32

The syntax shown above is an example using CTC1 as a model. The specific syntax is implementation dependent.

Purpose: Move Control Word to Coprocessor 2

To copy a word from a GPR to a Coprocessor 2 control register

Description: CP2CCR[*Impl*] ← GPR[*rt*]

Copy the low word from GPR *rt* into the Coprocessor 2 control register denoted by the *Impl* field. The interpretation of the *Impl* field is left entirely to the Coprocessor 2 implementation and is not specified by the architecture.

Restrictions:


The result is **UNPREDICTABLE** if *rd* specifies a register that does not exist.

Operation:

```
temp ← GPR[rt]
CP2CCR[Impl] ← temp
```

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: CVT.D.fmt
 CVT.D.S fd, fs
 CVT.D.W fd, fs
 CVT.D.L fd, fs

MIPS32
MIPS32
MIPS64, MIPS32 Release 2

Purpose: Floating Point Convert to Double Floating Point

To convert an FP or fixed point value to double FP

Description: $FPR[fd] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in double floating point format and 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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

For CVT.D.L, the result of this instruction is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:


StoreFPR (fd, D, ConvertFmt(ValueFPR(fs, fmt), fmt, D))

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact



Format: CVT.L.fmt
 CVT.L.S fd, fs
 CVT.L.D fd, fs

MIPS64, MIPS32 Release 2
MIPS64, MIPS32 Release 2

Purpose: Floating Point Convert to Long Fixed Point

To convert an FP value to a 64-bit fixed point

Description: $FPR[fd] \leftarrow \text{convert_and_round}(FPR[fs])$

Convert the value in format *fmt* in FPR *fs* to long fixed point format and round 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^{63} to $2^{63}-1$, the result cannot be represented correctly, an IEEE Invalid Operation condition exists, and 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*.

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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of this instruction is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:

StoreFPR (fd, L, ConvertFmt(ValueFPR(fs, fmt), fmt, L))

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact, Overflow

31	26 25	21 20	16 15	11 10	6 5	0
COP1 010001	fmt 10000	ft	fs	fd	CVT.PS 100110	
6	5	5	5	5	6	

Format: CVT.PS.S *fd*, *fs*, *ft*


MIPS64, MIPS32 Release 2

Purpose: Floating Point Convert Pair to Paired Single

To convert two FP values to a paired single value

Description: $FPR[fd] \leftarrow FPR[fs]_{31..0} || FPR[ft]_{31..0}$

The single-precision values in FPR *fs* and *ft* are written into FPR *fd* as a paired-single value. The value in FPR *fs* is written into the upper half, and the value in FPR *ft* is written into the lower half.



CVT.PS.S is similar to PLL.PS, except that it expects operands of format *S* instead of *PS*.

The move is non-arithmetic; it causes no IEEE 754 exceptions.

Restrictions:

The fields *fs* and *ft* must specify FPRs valid for operands of type *S*; if they are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *S*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of this instruction is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:


$StoreFPR(fd, S, ValueFPR(fs, S) || ValueFPR(ft, S))$

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation



Format: CVT.S.fmt
 CVT.S.D fd, fs **MIPS32**
 CVT.S.W fd, fs **MIPS32**
 CVT.S.L fd, fs **MIPS64, MIPS32 Release 2**

Purpose: Floating Point Convert to Single Floating Point

To convert an FP or fixed point value to single FP

Description: $FPR[fd] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in single floating point format and 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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

For CVT.S.L, the result of this instruction is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:


StoreFPR(fd, S, ConvertFmt(ValueFPR(fs, fmt), fmt, S))

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact, Overflow, Underflow



Format: CVT.S.PL *fd*, *fs*

MIPS64, MIPS32 Release 2

Purpose:

Floating Point Convert Pair Lower to Single Floating Point

To convert one half of a paired single FP value to single FP

Description: $FPR[fd] \leftarrow FPR[fs]_{31..0}$

The lower paired single value in FPR *fs*, in format *PS*, is converted to a value in single floating point format. The result is placed in FPR *fd*. This instruction can be used to isolate the lower half of a paired single value.

The operation is non-arithmetic; it causes no IEEE 754 exceptions.

Restrictions:

The fields *fs* and *fd* must specify valid FPRs—*fs* for type *PS* and *fd* for single floating point. If they are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *PS*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of CVT.S.PL is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.


Operation:

`StoreFPR (fd, S, ConvertFmt(ValueFPR(fs, PS), PL, S))`

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:



Format: CVT.S.PU *fd*, *fs*

MIPS64, MIPS32 Release 2

Purpose: Floating Point Convert Pair Upper to Single Floating Point

To convert one half of a paired single FP value to single FP

Description: $FPR[fd] \leftarrow FPR[fs]_{63..32}$

The upper paired single value in FPR *fs*, in format *PS*, is converted to a value in single floating point format. The result is placed in FPR *fd*. This instruction can be used to isolate the upper half of a paired single value.

The operation is non-arithmetic; it causes no IEEE 754 exceptions.

Restrictions:

The fields *fs* and *fd* must specify valid FPRs—*fs* for type *PS* and *fd* for single floating point. If they are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *PS*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of CVT.S.PU is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.


Operation:

`StoreFPR (fd, S, ConvertFmt(ValueFPR(fs, PS), PU, S))`

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:



Format: CVT.W.fmt
 CVT.W.S fd, fs
 CVT.W.D fd, fs

MIPS32
MIPS32

Purpose: Floating Point Convert to Word Fixed Point

To convert an FP value to 32-bit fixed point

Description: $FPR[fd] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in 32-bit word fixed point format and 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, an IEEE Invalid Operation condition exists, and 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*.

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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Operation:


`StoreFPR(fd, W, ConvertFmt(ValueFPR(fs, fmt), fmt, W))`

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact, Overflow



Format: DERET

EJTAG

Purpose: Debug Exception Return

To Return from a debug exception.

Description:

DERET clears execution and instruction hazards, returns from Debug Mode and resumes non-debug execution at the instruction whose address is contained in the *DEPC* register. DERET does not execute the next instruction (i.e. it has no delay slot).

Restrictions:

A DERET placed between an LL and SC instruction does not cause the SC to fail.

If the *DEPC* register with the return address for the DERET was modified by an MTC0 or a DMTC0 instruction, a CP0 hazard exists that must be removed via software insertion of the appropriate number of SSNOP instructions (for implementations of Release 1 of the Architecture) or by an EHB, or other execution hazard clearing instruction (for implementations of Release 2 of the Architecture).

DERET implements a software barrier that resolves all execution and instruction hazards created by Coprocessor 0 state changes (for Release 2 implementations, refer to the SYNCI instruction for additional information on resolving instruction hazards created by writing the instruction stream). The effects of this barrier are seen starting with the instruction fetch and decode of the instruction at the PC to which the DERET returns.

This instruction is legal only if the processor is executing in Debug Mode. The operation of the processor is **UNDEFINED** if a DERET is executed in the delay slot of a branch or jump instruction.

Operation:

```


DebugDM ← 0
DebugIEXI ← 0
if IsMIPS16Implemented() | (Config3ISA > 0) then
    PC ← DEPC31..1 || 0
    ISAMode ← DEPC0
else
    PC ← DEPC
endif
ClearHazards()

```

Exceptions:

Coprocessor Unusable Exception

Reserved Instruction Exception



Format: DI
DI rt

MIPS32 Release 2
MIPS32 Release 2

Purpose: Disable Interrupts

To return the previous value of the *Status* register and disable interrupts. If DI is specified without an argument, GPR r0 is implied, which discards the previous value of the Status register.

Description: $GPR[rt] \leftarrow Status; Status_{IE} \leftarrow 0$

The current value of the *Status* register is loaded into general register *rt*. The Interrupt Enable (IE) bit in the *Status* register is then cleared.

Restrictions:

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction Exception.

Operation:

This operation specification is for the general interrupt enable/disable operation, with the *sc* field as a variable. The individual instructions DI and EI have a specific value for the *sc* field.

```
data ← Status
GPR[rt] ← data
StatusIE ← 0
```


Exceptions:

Coprocessor Unusable
Reserved Instruction (Release 1 implementations)

Programming Notes:

The effects of this instruction are identical to those accomplished by the sequence of reading *Status* into a GPR, clearing the IE bit, and writing the result back to *Status*. Unlike the multiple instruction sequence, however, the DI instruction can not be aborted in the middle by an interrupt or exception.

This instruction creates an execution hazard between the change to the Status register and the point where the change to the interrupt enable takes effect. This hazard is cleared by the EHB, JALR.HB, JR.HB, or ERET instructions. Software must not assume that a fixed latency will clear the execution hazard.



Format: DIV *rs*, *rt*

MIPS32

Purpose: Divide Word

To divide a 32-bit signed integers

Description: $(HI, LO) \leftarrow GPR[rs] / GPR[rt]$

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:

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

Operation:

```

q ← GPR[rs]31..0 div GPR[rt]31..0
LO ← q
r ← GPR[rs]31..0 mod GPR[rt]31..0
HI ← r

```

Exceptions:

None

Programming Notes:

No arithmetic exception occurs under any circumstances. If divide-by-zero or overflow conditions are 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 within 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 a zero is detected.


By default, most compilers for the MIPS architecture will emit additional instructions to check for the divide-by-zero and overflow cases when this instruction is used. In many compilers, the assembler mnemonic “DIV r0, rs, rt” can be used to prevent these additional test instructions to be emitted.

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 interlocks 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.

Historical Perspective:

In MIPS 1 through MIPS III, if either of the two instructions preceding the divide is an MFHI or MFLO, the result of the MFHI or MFLO is **UNPREDICTABLE**. Reads of the *HI* or *LO* special register must be separated from subse-

quent instructions that write to them by two or more instructions. This restriction was removed in MIPS IV and MIPS32 and all subsequent levels of the architecture.



Format: DIV.fmt
 DIV.S fd, fs, ft
 DIV.D fd, fs, ft

MIPS32
MIPS32

Purpose: Floating Point Divide

To divide FP values

Description: $FPR[fd] \leftarrow FPR[fs] / FPR[ft]$

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 **UNPREDICTABLE**.

The operands must be values in format *fmt*; if they are not, the result is **UNPREDICTABLE** and the value of the operand FPRs becomes **UNPREDICTABLE**.

Operation:


StoreFPR (fd, fmt, ValueFPR(fs, fmt) / ValueFPR(ft, fmt))

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Invalid Operation, Unimplemented Operation, Division-by-zero, Overflow, Underflow



Format: DIVU *rs*, *rt*

MIPS32

Purpose: Divide Unsigned Word

To divide a 32-bit unsigned integers

Description: $(HI, LO) \leftarrow GPR[rs] / GPR[rt]$

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:

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

Operation:

```

q ← (0 || GPR[rs]31..0) div (0 || GPR[rt]31..0)
r ← (0 || GPR[rs]31..0) mod (0 || GPR[rt]31..0)
LO ← sign_extend(q31..0)
HI ← sign_extend(r31..0)

```

Exceptions:

None

Programming Notes:

See “Programming Notes” for the [DIV](#) instruction.

Historical Perspective:

In MIPS 1 through MIPS III, if either of the two instructions preceding the divide is an MFHI or MFLO, the result of the MFHI or MFLO is UNPREDICTABLE. Reads of the *HI* or *LO* special register must be separated from subsequent instructions that write to them by two or more instructions. This restriction was removed in MIPS IV and MIPS32 and all subsequent levels of the architecture.

31	26 25	21 20	16 15	11 10	6 5	0
SPECIAL 000000	0 00000	0 00000	0 00000	3 00011	SLL 000000	
6	5	5	5	5	6	

Format: EHB

MIPS32 Release 2

Purpose: Execution Hazard Barrier

To stop instruction execution until all execution hazards have been cleared.

Description:

EHB is the assembly idiom used to denote execution hazard barrier. The actual instruction is interpreted by the hardware as SLL r0, r0, 3.

This instruction alters the instruction issue behavior on a pipelined processor by stopping execution until all execution hazards have been cleared. Other than those that might be created as a consequence of setting Status_{CU0}, there are no execution hazards visible to an unprivileged program running in User Mode. All execution hazards created by previous instructions are cleared for instructions executed immediately following the EHB, even if the EHB is executed in the delay slot of a branch or jump. The EHB instruction does not clear instruction hazards—such hazards are cleared by the JALR.HB, JR.HB, and ERET instructions.

Restrictions:

None

Operation:


```
ClearExecutionHazards()
```

Exceptions:

None

Programming Notes:

In MIPS32 Release 2 implementations, this instruction resolves all execution hazards. On a superscalar processor, EHB alters the instruction issue behavior in a manner identical to SSNOP. For backward compatibility with Release 1 implementations, the last of a sequence of SSNOPs can be replaced by an EHB. In Release 1 implementations, the EHB will be treated as an SSNOP, thereby preserving the semantics of the sequence. In Release 2 implementations, replacing the final SSNOP with an EHB should have no performance effect because a properly sized sequence of SSNOPs will have already cleared the hazard. As EHB becomes the standard in MIPS implementations, the previous SSNOPs can be removed, leaving only the EHB.



Format: EI
EI rt

MIPS32 Release 2
MIPS32 Release 2

Purpose: Enable Interrupts

To return the previous value of the *Status* register and enable interrupts. If EI is specified without an argument, GPR r0 is implied, which discards the previous value of the Status register.

Description: GPR[rt] ← Status; Status_{IE} ← 1

The current value of the *Status* register is loaded into general register *rt*. The Interrupt Enable (IE) bit in the *Status* register is then set.

Restrictions:

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction Exception.

Operation:

This operation specification is for the general interrupt enable/disable operation, with the *sc* field as a variable. The individual instructions DI and EI have a specific value for the *sc* field.

```
data ← Status
GPR[rt] ← data
StatusIE ← 1
```

Exceptions:

Coprocessor Unusable
Reserved Instruction (Release 1 implementations)

Programming Notes:

The effects of this instruction are identical to those accomplished by the sequence of reading *Status* into a GPR, setting the IE bit, and writing the result back to *Status*. Unlike the multiple instruction sequence, however, the EI instruction can not be aborted in the middle by an interrupt or exception.

This instruction creates an execution hazard between the change to the Status register and the point where the change to the interrupt enable takes effect. This hazard is cleared by the EHB, JALR.HB, JR.HB, or ERET instructions. Software must not assume that a fixed latency will clear the execution hazard.

31	26	25	24	6	5	0	
COP0 010000	CO 1	0 000 0000 0000 0000 0000				ERET 011000	
6	1	19				6	

Format: ERET

MIPS32

Purpose: Exception Return

To return from interrupt, exception, or error trap.

Description:

ERET clears execution and instruction hazards, conditionally restores $SRSCtl_{CSS}$ from $SRSCtl_{IIS\Sigma}$ in a Release 2 implementation, and returns to the interrupted instruction at the completion of interrupt, exception, or error processing. ERET does not execute the next instruction (i.e., it has no delay slot).

Restrictions:

The operation of the processor is **UNDEFINED** if an ERET is executed in the delay slot of a branch or jump instruction.

An ERET placed between an LL and SC instruction will always cause the SC to fail.

ERET implements a software barrier that resolves all execution and instruction hazards created by Coprocessor 0 state changes (for Release 2 implementations, refer to the SYNCI instruction for additional information on resolving instruction hazards created by writing the instruction stream). The effects of this barrier are seen starting with the instruction fetch and decode of the instruction at the PC to which the ERET returns.

In a Release 2 implementation, ERET does not restore $SRSCtl_{CSS}$ from $SRSCtl_{PSS}$ if $Status_{BEV} = 1$, or if $Status_{ERL} = 1$ because any exception that sets $Status_{ERL}$ to 1 (Reset, Soft Reset, NMI, or cache error) does not save $SRSCtl_{CSS}$ in $SRSCtl_{PSS}$. If software sets $Status_{ERL}$ to 1, it must be aware of the operation of an ERET that may be subsequently executed.

Operation:


```

if StatusERL = 1 then
    temp ← ErrorEPC
    StatusERL ← 0
else
    temp ← EPC
    StatusEXL ← 0
    if (ArchitectureRevision ≥ 2) and (SRSCtlHSS > 0) and (StatusBEV = 0) then
        SRSCtlCSS ← SRSCtlPSS
    endif
endif
if IsMIPS16Implemented() | (Config3ISA > 0) then
    PC ← temp31..1 || 0
    ISAMode ← temp0
else
    PC ← temp
endif
LLbit ← 0
ClearHazards()

```

Exceptions:

Coprocessor Unusable Exception



Format: EXT *rt*, *rs*, *pos*, *size*

MIPS32 Release 2

Purpose: Extract Bit Field

To extract a bit field from GPR *rs* and store it right-justified into GPR *rt*.

Description: $GPR[rt] \leftarrow \text{ExtractField}(GPR[rs], \text{msbd}, \text{lsb})$

The bit field starting at bit *pos* and extending for *size* bits is extracted from GPR *rs* and stored zero-extended and right-justified in GPR *rt*. The assembly language arguments *pos* and *size* are converted by the assembler to the instruction fields *msbd* (the most significant bit of the destination field in GPR *rt*), in instruction bits **15..11**, and *lsb* (least significant bit of the source field in GPR *rs*), in instruction bits **10..6**, as follows:


```
msbd ← size-1
lsb ← pos
```

The values of *pos* and *size* must satisfy all of the following relations:

```
0 ≤ pos < 32
0 < size ≤ 32
0 < pos+size ≤ 32
```

Figure 3-9 shows the symbolic operation of the instruction.

Figure 3.3 Operation of the EXT Instruction



Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction Exception.

The operation is **UNPREDICTABLE** if $\text{lsb}+\text{msbd} > 31$.


Operation:

```
if (lsb + msbd) > 31) then
    UNPREDICTABLE
endif
temp ← 032-(msbd+1) || GPR[rs]msbd+lsb..lsb
```

$\text{GPR}[\text{rt}] \leftarrow \text{temp}$

Exceptions:

Reserved Instruction



Format: FLOOR.L.fmt
 FLOOR.L.S fd, fs
 FLOOR.L.D fd, fs

MIPS64, MIPS32 Release 2
MIPS64, MIPS32 Release 2

Purpose: Floating Point Floor Convert to Long Fixed Point

To convert an FP value to 64-bit fixed point, rounding down

Description: $FPR[fd] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in 64-bit long fixed point format and rounded 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, an IEEE Invalid Operation condition exists, and 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*.

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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of this instruction is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:


$\text{StoreFPR}(fd, L, \text{ConvertFmt}(\text{ValueFPR}(fs, fmt), fmt, L))$

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact, Overflow



Format: FLOOR.W.fmt

FLOOR.W.S fd, fs

FLOOR.W.D fd, fs

MIPS32

MIPS32

Purpose: Floating Point Floor Convert to Word Fixed Point

To convert an FP value to 32-bit fixed point, rounding down

Description: $FPR[fd] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in 32-bit word fixed point format and rounded 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, an IEEE Invalid Operation condition exists, and 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*.

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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Operation:


`StoreFPR(fd, W, ConvertFmt(ValueFPR(fs, fmt), fmt, W))`

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Invalid Operation, Unimplemented Operation, Inexact, Overflow



Format: `INS rt, rs, pos, size`

MIPS32 Release 2

Purpose: Insert Bit Field

To merge a right-justified bit field from GPR *rs* into a specified field in GPR *rt*.

Description: $GPR[rt] \leftarrow InsertField(GPR[rt], GPR[rs], msb, lsb)$

The right-most *size* bits from GPR *rs* are merged into the value from GPR *rt* starting at bit position *pos*. The result is placed back in GPR *rt*. The assembly language arguments *pos* and *size* are converted by the assembler to the instruction fields *msb* (the most significant bit of the field), in instruction bits 15..11, and *lsb* (least significant bit of the field), in instruction bits 10..6, as follows:

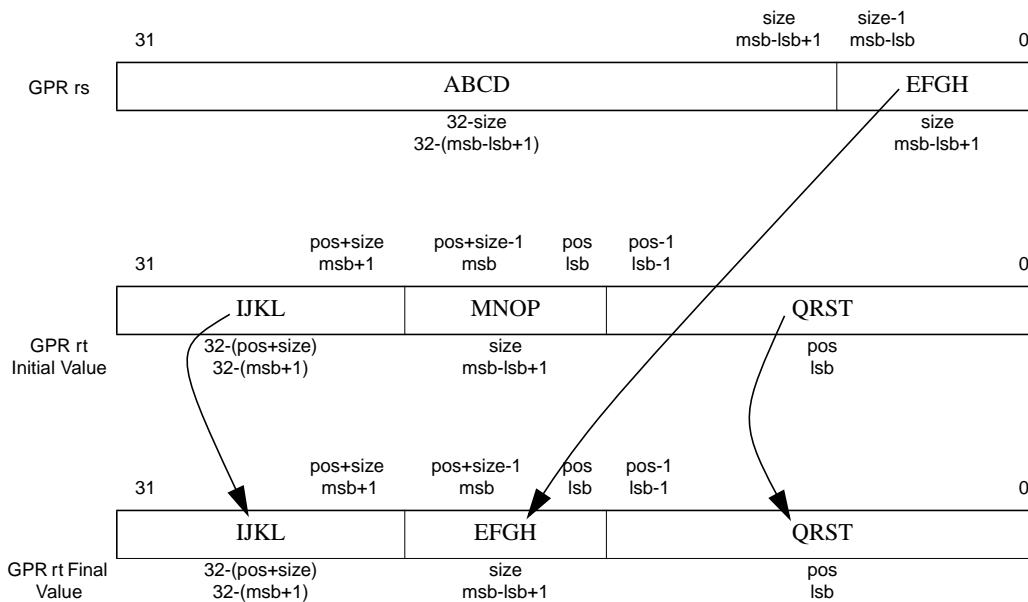
$msb \leftarrow pos+size-1$
 $lsb \leftarrow pos$

The values of *pos* and *size* must satisfy all of the following relations:

$0 \leq pos < 32$
 $0 < size \leq 32$
 $0 < pos+size \leq 32$

Figure 3-10 shows the symbolic operation of the instruction.

Figure 3.4 Operation of the INS Instruction



Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction Exception.


The operation is **UNPREDICTABLE** if $lsb > msb$.

Operation:

```
if lsb > msb) then
    UNPREDICTABLE
endif
GPR[rt] ← GPR[rt]31..msb+1 || GPR[rs]msb-lsb..0 || GPR[rt]lsb-1..0
```

Exceptions:

Reserved Instruction



Format: J target

MIPS32

Purpose: Jump

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 2bits. 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 that follows the jump, in the branch delay slot, before executing the jump itself.

Restrictions:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

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


Exceptions:

None

Programming Notes:

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

This definition creates the following boundary case: When the jump instruction is in the last word of a 256MB region, it can branch only to the following 256MB region containing the branch delay slot.



Format: JAL target

MIPS32

Purpose: Jump and Link

To execute a procedure call within the current 256MB-aligned region

Description:

Place the return address link in GPR 31. The return link is the address of the second instruction following the branch, at which location execution continues after a procedure call.

This is a PC-region branch (not PC-relative); the effective target address is in the “current” 256MB-aligned region. The low 28 bits of the target address is the *instr_index* field shifted left 2bits. 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 that follows the jump, in the branch delay slot, before executing the jump itself.

Restrictions:

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

I: GPR[31] ← PC + 8
I+1: PC ← PC_{GPRLEN-1..28} || instr_index || 0²


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 256MB region aligned on a 256MB boundary. It allows a branch from anywhere in the region to anywhere in the region, an action not allowed by a signed relative offset.

This definition creates the following boundary case: When the branch instruction is in the last word of a 256MB region, it can branch only to the following 256MB region containing the branch delay slot.



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

MIPS32
MIPS32

Purpose: Jump and Link Register

To execute a procedure call to an instruction address in a register

Description: GPR[rd] ← return_addr, PC ← GPR[rs]

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

For processors that do not implement the MIPS16e ASE nor microMIPS32/64 ISA:

- Jump to the effective target address in GPR *rs*. Execute the instruction that follows the jump, in the branch delay slot, before executing the jump itself. Bit 0 of the target address is always zero so that no Address Exceptions occur when bit 0 of the source register is one.

For processors that do implement the MIPS16e ASE or microMIPS32/64 ISA:

- Jump to the effective target address in GPR *rs*. Execute the instruction that follows the jump, in the branch delay slot, before executing the jump itself. Set the *ISA Mode* bit to the value in GPR *rs* bit 0. Bit 0 of the target address is always zero so that no Address Exceptions occur when bit 0 of the source register is one.

In release 1 of the architecture, the only defined hint field value is 0, which sets default handling of JALR. In Release 2 of the architecture, bit 10 of the hint field is used to encode a hazard barrier. See the [JALR.HB](#) instruction description for additional information.

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 **UNPREDICTABLE**. This restriction permits an exception handler to resume execution by re-executing the branch when an exception occurs in the branch delay slot.

If only one instruction set is implemented, then the effective target address must obey the alignment rules of the instruction set. If multiple instruction sets are implemented, the effective target address must obey the alignment rules of the intended instruction set of the target address as specified by the bit 0 or GPR *rs*.

For processors that do not implement the microMIPS32/64 ISA, the effective target address in GPR *rs* must be naturally-aligned. For processors that do not implement the MIPS16e ASE nor microMIPS32/64 ISA, if either of the two least-significant bits are not zero, an Address Error exception occurs when the branch target is subsequently fetched as an instruction.

For processors that do implement the MIPS16e ASE or microMIPS32/64 ISA, if target ISAMode bit is 0 (GPR *rs* bit 0) is zero and bit 1 is one, an Address Error exception occurs when the jump target is subsequently fetched as an instruction.

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

I: temp ← GPR[rs]
GPR[rd] ← PC + 8


```
I+1:if Config1CA = 0 then
    PC ← temp
else
    PC ← tempGPRLEN-1..1 || 0
    ISAMode ← temp0
endif
```

Exceptions:

None

Programming Notes:

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



Format: JALR.HB rs (rd = 31 implied)
JALR.HB rd, rs

MIPS32 Release 2
MIPS32 Release 2

Purpose: Jump and Link Register with Hazard Barrier

To execute a procedure call to an instruction address in a register and clear all execution and instruction hazards

Description: GPR[rd] ← return_addr, PC ← GPR[rs], clear execution and instruction hazards

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

For processors that do not implement the MIPS16 ASE nor microMIPS32/64 ISA:

- Jump to the effective target address in GPR *rs*. Execute the instruction that follows the jump, in the branch delay slot, before executing the jump itself. Bit 0 of the target address is always zero so that no Address Exceptions occur when bit 0 of the source register is one.

For processors that do implement the MIPS16 ASE or microMIPS32/64 ISA:

- Jump to the effective target address in GPR *rs*. Execute the instruction that follows the jump, in the branch delay slot, before executing the jump itself. Set the *ISA Mode* bit to the value in GPR *rs* bit 0. Bit 0 of the target address is always zero so that no Address Exceptions occur when bit 0 of the source register is one.

JALR.HB implements a software barrier that resolves all execution and instruction hazards created by Coprocessor 0 state changes (for Release 2 implementations, refer to the SYNCI instruction for additional information on resolving instruction hazards created by writing the instruction stream). The effects of this barrier are seen starting with the instruction fetch and decode of the instruction at the PC to which the JALR.HB instruction jumps. An equivalent barrier is also implemented by the ERET instruction, but that instruction is only available if access to Coprocessor 0 is enabled, whereas JALR.HB is legal in all operating modes.

This instruction clears both execution and instruction hazards. Refer to the [EHB](#) instruction description for the method of clearing execution hazards alone.

JALR.HB uses bit 10 of the instruction (the upper bit of the hint field) to denote the hazard barrier operation.

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 **UNPREDICTABLE**. This restriction permits an exception handler to resume execution by re-executing the branch when an exception occurs in the branch delay slot.

If only one instruction set is implemented, then the effective target address must obey the alignment rules of the instruction set. If multiple instruction sets are implemented, the effective target address must obey the alignment rules of the intended instruction set of the target address as specified by the bit 0 or GPR *rs*.

For processors that do not implement the microMIPS32/64 ISA, the effective target address in GPR *rs* must be naturally-aligned. For processors that do not implement the MIPS16 ASE nor microMIPS32/64 ISA, if either of the two least-significant bits are not zero, an Address Error exception occurs when the branch target is subsequently fetched as an instruction.

For processors that do implement the MIPS16 ASE or microMIPS32/64 ISA, if bit 0 is zero and bit 1 is one, an Address Error exception occurs when the jump target is subsequently fetched as an instruction.

After modifying an instruction stream mapping or writing to the instruction stream, execution of those instructions has **UNPREDICTABLE** behavior until the instruction hazard has been cleared with JALR.HB, JR.HB, ERET, or DERET. Further, the operation is **UNPREDICTABLE** if the mapping of the current instruction stream is modified.

JALR.HB does not clear hazards created by any instruction that is executed in the delay slot of the JALR.HB. Only hazards created by instructions executed before the JALR.HB are cleared by the JALR.HB.

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I: temp ← GPR[rs]
      GPR[rd] ← PC + 8
I+1: if Config1_CA = 0 then
      PC ← temp
      else
      PC ← temp_GPREN-1..1 || 0
      ISAMode ← temp_0
      endif
      ClearHazards()

```

Exceptions:

None

Programming Notes:

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


This instruction implements the final step in clearing execution and instruction hazards before execution continues. A hazard is created when a Coprocessor 0 or TLB write affects execution or the mapping of the instruction stream, or after a write to the instruction stream. When such a situation exists, software must explicitly indicate to hardware that the hazard should be cleared. Execution hazards alone can be cleared with the EHB instruction. Instruction hazards can only be cleared with a JR.HB, JALR.HB, or ERET instruction. These instructions cause hardware to clear the hazard before the instruction at the target of the jump is fetched. Note that because these instructions are encoded as jumps, the process of clearing an instruction hazard can often be included as part of a call (JALR) or return (JR) sequence, by simply replacing the original instructions with the HB equivalent.

Example: Clearing hazards due to an ASID change

```

/*
 * Code used to modify ASID and call a routine with the new
 * mapping established.
 *
 * a0 = New ASID to establish
 * a1 = Address of the routine to call
 */
mfc0   v0, C0_EntryHi      /* Read current ASID */
li     v1, ~M_EntryHiASID /* Get negative mask for field */
and    v0, v0, v1         /* Clear out current ASID value */
or     v0, v0, a0         /* OR in new ASID value */
mtc0   v0, C0_EntryHi     /* Rewrite EntryHi with new ASID */
jalr.hb a1                 /* Call routine, clearing the hazard */
nop

```

Format: JALX target

MIPS32 with (microMIPS32 or MIPS16e)

Purpose: Jump and Link Exchange

To execute a procedure call within the current 256 MB-aligned region and change the *ISA Mode* from MIPS32 to microMIPS32 or MIPS16e.

Description:

Place the return address link in GPR 31. The return link is the address of the second instruction following the branch, at which location execution continues after a procedure call. The value stored in GPR 31 bit 0 reflects the current value of the *ISA Mode* bit.

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, toggling the *ISA Mode* bit. Execute the instruction that follows the jump, in the branch delay slot, before executing the jump itself.

Restrictions:

This instruction only supports 32-bit aligned branch target addresses.

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

If the microMIPS base architecture is not implemented and the MIPS16e ASE is not implemented, a Reserved Instruction Exception is initiated.

Operation:

I: GPR[31] ← PC + 8
I+1: PC ← PC_{GPRLEN-1..28} || instr_index || 0²
 ISAMode ← (not ISAMode)


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 from anywhere in the region to anywhere in the region, an action not allowed by a signed relative offset.

This definition creates the following boundary case: When the branch instruction is in the last word of a 256 MB region, it can branch only to the following 256 MB region containing the branch delay slot.



Format: JR *rs*

MIPS32

Purpose: Jump Register

To execute a branch to an instruction address in a register

Description: $PC \leftarrow GPR[rs]$

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

For processors that implement the MIPS16e ASE or microMIPS32/64 ISA, set the *ISA Mode* bit to the value in GPR *rs* bit 0. Bit 0 of the target address is always zero so that no Address Exceptions occur when bit 0 of the source register is one

Restrictions:

If only one instruction set is implemented, then the effective target address must obey the alignment rules of the instruction set. If multiple instruction sets are implemented, the effective target address must obey the alignment rules of the intended instruction set of the target address as specified by the bit 0 or GPR *rs*.

For processors that do not implement the microMIPS ISA, the effective target address in GPR *rs* must be naturally-aligned. For processors that do not implement the MIPS16e ASE or microMIPS ISA, if either of the two least-significant bits are not zero, an Address Error exception occurs when the branch target is subsequently fetched as an instruction.

For processors that do implement the MIPS16e ASE or microMIPS ISA, if bit 0 is zero and bit 1 is one, an Address Error exception occurs when the jump target is subsequently fetched as an instruction.

In release 1 of the architecture, the only defined hint field value is 0, which sets default handling of JR. In Release 2 of the architecture, bit 10 of the hint field is used to encode an instruction hazard barrier. See the [JR.HB](#) instruction description for additional information.

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

```

I: temp ← GPR[rs]
I+1: if Config1CA = 0 then
    PC ← temp
  else
    PC ← tempGPRLEN-1..1 || 0
    ISAMode ← temp0
  endif


```

Exceptions:

None

Programming Notes:

Software should use the value 31 for the *rs* field of the instruction word on return from a JAL, JALR, or BGEZAL, and should use a value other than 31 for remaining uses of JR.



Format: JR.HB rs

MIPS32 Release 2

Purpose: Jump Register with Hazard Barrier

To execute a branch to an instruction address in a register and clear all execution and instruction hazards.

Description: $PC \leftarrow GPR[rs]$, clear execution and instruction hazards

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

JR.HB implements a software barrier that resolves all execution and instruction hazards created by Coprocessor 0 state changes (for Release 2 implementations, refer to the SYNCI instruction for additional information on resolving instruction hazards created by writing the instruction stream). The effects of this barrier are seen starting with the instruction fetch and decode of the instruction at the PC to which the JR.HB instruction jumps. An equivalent barrier is also implemented by the ERET instruction, but that instruction is only available if access to Coprocessor 0 is enabled, whereas JR.HB is legal in all operating modes.

This instruction clears both execution and instruction hazards. Refer to the [EHB](#) instruction description for the method of clearing execution hazards alone.

JR.HB uses bit 10 of the instruction (the upper bit of the hint field) to denote the hazard barrier operation.

For processors that implement the MIPS16e ASE or microMIPS32/64 ISA, set the *ISA Mode* bit to the value in GPR *rs* bit 0. Bit 0 of the target address is always zero so that no Address Exceptions occur when bit 0 of the source register is one.

Restrictions:

If only one instruction set is implemented, then the effective target address must obey the alignment rules of the instruction set. If multiple instruction sets are implemented, the effective target address must obey the alignment rules of the intended instruction set of the target address as specified by the bit 0 or GPR *rs*.

For processors that do not implement the microMIPS ISA, the effective target address in GPR *rs* must be naturally-aligned. For processors that do not implement the MIPS16 ASE or microMIPS ISA, if either of the two least-significant bits are not zero, an Address Error exception occurs when the branch target is subsequently fetched as an instruction.

For processors that do implement the MIPS16 ASE or microMIPS ISA, if bit 0 is zero and bit 1 is one, an Address Error exception occurs when the jump target is subsequently fetched as an instruction.

After modifying an instruction stream mapping or writing to the instruction stream, execution of those instructions has **UNPREDICTABLE** behavior until the hazard has been cleared with JALR.HB, JR.HB, ERET, or DERET. Further, the operation is **UNPREDICTABLE** if the mapping of the current instruction stream is modified.

JR.HB does not clear hazards created by any instruction that is executed in the delay slot of the JR.HB. Only hazards created by instructions executed before the JR.HB are cleared by the JR.HB.

Processor operation is **UNPREDICTABLE** if a branch, jump, ERET, DERET, or WAIT instruction is placed in the delay slot of a branch or jump.

Operation:

I: temp \leftarrow GPR[rs]

```

I+1:if Config1CA = 0 then
    PC ← temp
else
    PC ← tempGPRLEN-1..1 || 0
    ISAMode ← temp0
endif
ClearHazards()

```

Exceptions:

None

Programming Notes:

This instruction implements the final step in clearing execution and instruction hazards before execution continues. A hazard is created when a Coprocessor 0 or TLB write affects execution or the mapping of the instruction stream, or after a write to the instruction stream. When such a situation exists, software must explicitly indicate to hardware that the hazard should be cleared. Execution hazards alone can be cleared with the EHB instruction. Instruction hazards can only be cleared with a JR.HB, JALR.HB, or ERET instruction. These instructions cause hardware to clear the hazard before the instruction at the target of the jump is fetched. Note that because these instructions are encoded as jumps, the process of clearing an instruction hazard can often be included as part of a call (JALR) or return (JR) sequence, by simply replacing the original instructions with the HB equivalent.

Example: Clearing hazards due to an ASID change

```

/*
 * Routine called to modify ASID and return with the new
 * mapping established.
 *
 * a0 = New ASID to establish
 */
mfc0   v0, C0_EntryHi      /* Read current ASID */
li     v1, ~M_EntryHiASID /* Get negative mask for field */
and    v0, v0, v1          /* Clear out current ASID value */
or     v0, v0, a0          /* OR in new ASID value */
mtc0   v0, C0_EntryHi      /* Rewrite EntryHi with new ASID */
jr.hb  ra                  /* Return, clearing the hazard */
nop

```

Example: Making a write to the instruction stream visible

```

/*
 * Routine called after new instructions are written to
 * make them visible and return with the hazards cleared.
 */
{Synchronize the caches - see the SYNCI and CACHE instructions}
sync                                /* Force memory synchronization */
jr.hb  ra                            /* Return, clearing the hazard */
nop


```

Example: Clearing instruction hazards in-line

```

la     AT, 10f
jr.hb  AT                            /* Jump to next instruction, clearing */
nop                                       /* hazards */
10:

```



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

MIPS32

Purpose: Load Byte

To load a byte from memory as a signed value

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

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:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..2 || (pAddr1..0 xor ReverseEndian2)
memword ← LoadMemory (CCA, BYTE, pAddr, vAddr, DATA)
byte ← vAddr1..0 xor BigEndianCPU2
GPR[rt] ← sign_extend(memword7+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Watch



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

MIPS32

Purpose: Load Byte Unsigned

To load a byte from memory as an unsigned value

Description: $GPR[rt] \leftarrow memory[GPR[base] + offset]$

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:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
pAddr ← pAddrPSIZE-1..2 || (pAddr1..0 xor ReverseEndian2)
memword ← LoadMemory (CCA, BYTE, pAddr, vAddr, DATA)
byte ← vAddr1..0 xor BigEndianCPU2
GPR[rt] ← zero_extend(memword7+8*byte..8*byte)

```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Watch



Format: LDC1 ft, offset(base)

MIPS32

Purpose: Load Doubleword to Floating Point

To load a doubleword from memory to an FPR

Description: $FPR[ft] \leftarrow \text{memory}[GPR[base] + \text{offset}]$

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.

Restrictions:

An Address Error exception occurs if $\text{EffectiveAddress}_{2..0} \neq 0$ (not doubleword-aligned).

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr2..0 ≠ 03 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
paddr ← paddr xor ((BigEndianCPU xor ReverseEndian) || 02)
memlsw ← LoadMemory(CCA, WORD, pAddr, vAddr, DATA)
paddr ← paddr xor 0b100
memmsw ← LoadMemory(CCA, WORD, pAddr, vAddr+4, DATA)
memdoubleword ← memmsw || memlsw
StoreFPR(ft, UNINTERPRETED_DOUBLEWORD, memdoubleword)

```

Exceptions:

Coprocessor Unusable, Reserved Instruction, TLB Refill, TLB Invalid, Address Error, Watch



Format: LDC2 *rt*, *offset*(*base*)

MIPS32

Purpose: Load Doubleword to Coprocessor 2

To load a doubleword from memory to a Coprocessor 2 register

Description: $CPR[2,rt,0] \leftarrow \text{memory}[GPR[base] + \text{offset}]$

The contents of the 64-bit doubleword at the memory location specified by the aligned effective address are fetched and placed in Coprocessor 2 register *rt*. 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 $\text{EffectiveAddress}_{2..0} \neq 0$ (not doubleword-aligned).

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr2..0 ≠ 03 then SignalException(AddressError) endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)
paddr ← paddr xor ((BigEndianCPU xor ReverseEndian) || 02)
memlsw ← LoadMemory(CCA, WORD, pAddr, vAddr, DATA)
paddr ← paddr xor 0b100
memmsw ← LoadMemory(CCA, WORD, pAddr, vAddr+4, DATA)
←memlsw
←memmsw

```

Exceptions:

Coprocessor Unusable, Reserved Instruction, TLB Refill, TLB Invalid, Address Error, Watch



Format: LDXC1 *fd*, *index*(*base*)

MIPS64
MIPS32 Release 2

Purpose: Load Doubleword Indexed to Floating Point

To load a doubleword from memory to an FPR (GPR+GPR addressing)

Description: $FPR[fd] \leftarrow \text{memory}[GPR[base] + GPR[index]]$

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.

Restrictions:

An Address Error exception occurs if $\text{EffectiveAddress}_{2..0} \neq 0$ (not doubleword-aligned).

Operation:


```

vAddr ← GPR[base] + GPR[index]
if vAddr2..0 ≠ 03 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
paddr ← paddr xor ((BigEndianCPU xor ReverseEndian) || 02)
memlsw ← LoadMemory(CCA, WORD, pAddr, vAddr, DATA)
paddr ← paddr xor 0b100
memmsw ← LoadMemory(CCA, WORD, pAddr, vAddr+4, DATA)
memdoubleword ← memmsw || memlsw
StoreFPR(fd, UNINTERPRETED_DOUBLEWORD, memdoubleword)

```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Reserved Instruction, Coprocessor Unusable, Watch



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

MIPS32

Purpose: Load Halfword Unsigned

To load a halfword from memory as an unsigned value

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

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.

Operation:


```

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

```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Watch



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

MIPS32

Purpose: Load Linked Word

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

Description: $GPR[rt] \leftarrow memory[GPR[base] + offset]$

The LL and SC instructions provide the primitives to implement atomic read-modify-write (RMW) operations for synchronizable memory locations.

The contents of the 32-bit word at the memory location specified by the aligned effective address are fetched and written into GPR *rt*. The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address.

This begins a RMW sequence on the current processor. There can be only one active RMW sequence per processor. When an LL is executed it starts an 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.

Executing LL on one processor does not cause an action that, by itself, causes 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 synchronizable by all processors and I/O devices sharing the location; if it is not, the result is **UNPREDICTABLE**. Which storage is synchronizable is a function of both CPU and system implementations. See the documentation of the SC instruction for the formal definition.

The effective address must be naturally-aligned. If either of the 2 least-significant bits of the effective address is non-zero, an Address Error exception occurs.

Operation:

```


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

```

Exceptions:

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

Programming Notes:



Format: LUI *rt*, *immediate*

MIPS32

Purpose: Load Upper Immediate

To load a constant into the upper half of a word

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

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

Restrictions:


None

Operation:

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

Exceptions:

None



Format: LUXC1 fd, index(base)

MIPS64
MIPS32 Release 2

Purpose: Load Doubleword Indexed Unaligned to Floating Point

To load a doubleword from memory to an FPR (GPR+GPR addressing), ignoring alignment

Description: $FPR[fd] \leftarrow memory[(GPR[base] + GPR[index])_{PSIZE-1..3}]$

The contents of the 64-bit doubleword at the memory location specified by the effective address are fetched and placed into the low word of FPR *fd*. The contents of GPR *index* and GPR *base* are added to form the effective address. The effective address is doubleword-aligned; EffectiveAddress_{2..0} are ignored.

Restrictions:


The result of this instruction is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:

```
vAddr ← (GPR[base]+GPR[index])63..3 || 03
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
pAddr ← pAddr xor ((BigEndianCPU xor ReverseEndian) || 02)
memlsw ← LoadMemory(CCA, WORD, pAddr, vAddr, DATA)
pAddr ← pAddr xor 0b100
memmsw ← LoadMemory(CCA, WORD, pAddr, vAddr+4, DATA)
memdoubleword ← memmsw || memlsw
StoreFPR(ft, UNINTERPRETED_DOUBLEWORD, memdoubleword)
```

Exceptions:

Coprocessor Unusable, Reserved Instruction, TLB Refill, TLB Invalid, Watch



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

MIPS32

Purpose: Load Word

To load a word from memory as a signed value

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

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 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Operation:


```

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

```

Exceptions:

TLB Refill, TLB Invalid, Bus Error, Address Error, Watch



Format: LWC1 *ft*, *offset*(*base*)

MIPS32

Purpose: Load Word to Floating Point

To load a word from memory to an FPR

Description: $FPR[ft] \leftarrow \text{memory}[GPR[base] + \text{offset}]$

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 FPR *ft*. If FPRs are 64 bits wide, bits 63..32 of FPR *fs* become **UNDEFINED**. 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 $\text{EffectiveAddress}_{1..0} \neq 0$ (not word-aligned).

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)


memword ← LoadMemory(CCA, WORD, pAddr, vAddr, DATA)

StoreFPR(ft, UNINTERPRETED_WORD,
         memword)

```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Reserved Instruction, Coprocessor Unusable, Watch



Format: LWC2 *rt*, *offset*(*base*)

MIPS32

Purpose: Load Word to Coprocessor 2

To load a word from memory to a COP2 register

Description: $CPR[2,rt,0] \leftarrow \text{memory}[GPR[base] + offset]$

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 *COP2* (Coprocessor 2) general register *rt*. 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 $\text{EffectiveAddress}_{1..0} \neq 0$ (not word-aligned).

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr12..0 ≠ 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, LOAD)


memword ← LoadMemory(CCA, DOUBLEWORD, pAddr, vAddr, DATA)

CPR[2,rt,0] ← memword

```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Reserved Instruction, Coprocessor Unusable, Watch



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

MIPS32

Purpose: Load Word Left

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


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

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 4 consecutive bytes forming a word (*W*) in memory starting at an arbitrary byte boundary.

The most-significant 1 to 4 bytes of *W* is in the aligned word containing the *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.

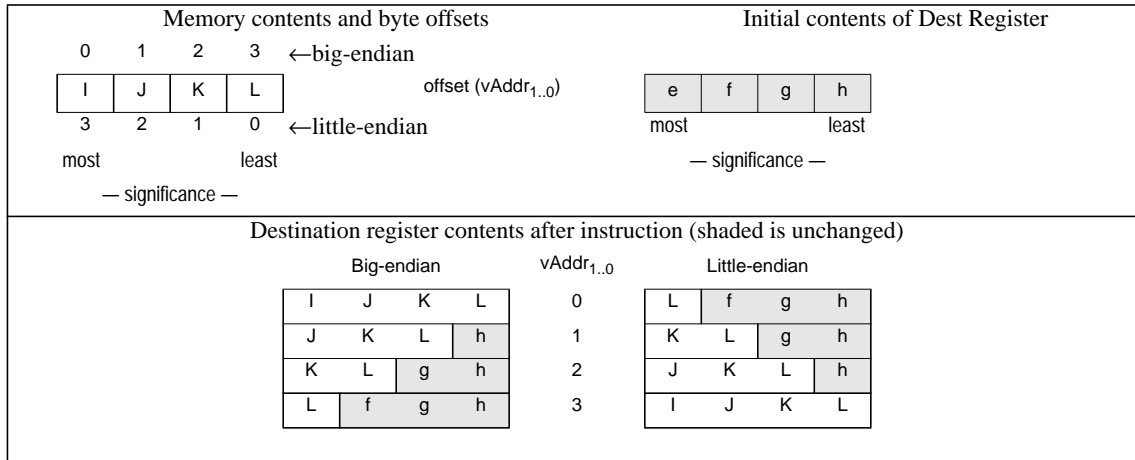
The figure below illustrates this operation using big-endian byte ordering for 32-bit and 64-bit registers. The 4 consecutive bytes in 2..5 form an unaligned word starting at location 2. A part of *W*, 2 bytes, is in the aligned word containing the most-significant byte at 2. First, LWL loads these 2 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

Figure 3.5 Unaligned Word Load Using LWL and LWR



The bytes loaded from memory to the destination register depend on both the offset of the effective address within an aligned word, that is, the low 2 bits of the address ($vAddr_{1,0}$), and the current byte-ordering mode of the processor (big- or little-endian). The figure below shows the bytes loaded for every combination of offset and byte ordering.

Figure 3.6 Bytes Loaded by LWL Instruction

**Restrictions:**

None

Operation:

```

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

```

Exceptions:

None


TLB Refill, TLB Invalid, Bus Error, Address Error, Watch

Programming Notes:

The architecture provides no direct support for treating unaligned words as unsigned values, that is, zeroing bits 63..32 of the destination register when bit 31 is loaded.

Historical Information:

In the MIPS I architecture, the LWL and LWR instructions were exceptions to the load-delay scheduling restriction. A LWL or LWR instruction which was immediately followed by another LWL or LWR instruction, and used the same destination register would correctly merge the 1 to 4 loaded bytes with the data loaded by the previous instruction. All such restrictions were removed from the architecture in MIPS II.



Format: LWR rt, offset(base)

MIPS32

Purpose: Load Word Right

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

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

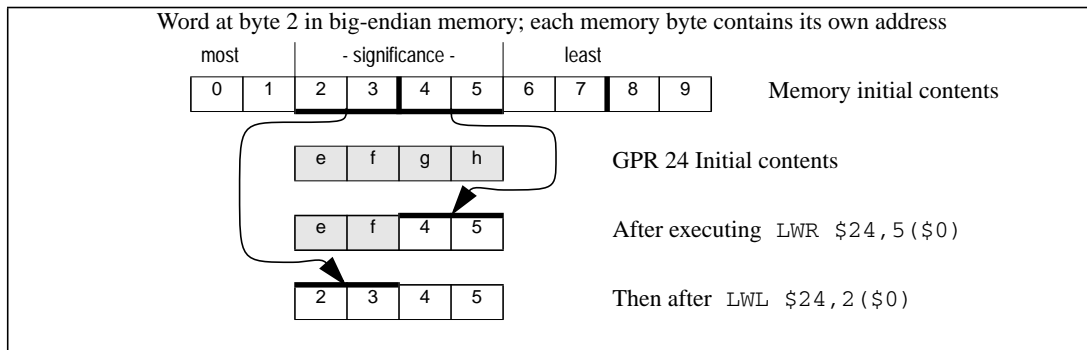
The 16-bit signed *offset* is added to the contents of GPR *base* to form an effective address (*EffAddr*). *EffAddr* is the address of the least-significant of 4 consecutive bytes forming a word (*W*) in memory starting at an arbitrary byte boundary.

A part of *W*, the least-significant 1 to 4 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.

Executing both LWR and LWL, in either order, delivers a sign-extended word value in the destination register.


The figure below illustrates this operation using big-endian byte ordering for 32-bit and 64-bit registers. The 4 consecutive bytes in 2..5 form an unaligned word starting at location 2. A part of *W*, 2 bytes, is in the aligned word containing the least-significant byte at 5. First, LWR loads these 2 bytes into the right part of the destination register. Next, the complementary LWL loads the remainder of the unaligned word.

Figure 3.7 Unaligned Word Load Using LWL and LWR



The bytes loaded from memory to the destination register depend on both the offset of the effective address within an aligned word, that is, the low 2 bits of the address ($vAddr_{1,0}$), and the current byte-ordering mode of the processor (big- or little-endian). The figure below shows the bytes loaded for every combination of offset and byte ordering.

Figure 3.8 Bytes Loaded by LWR Instruction

**Restrictions:**

None

Operation:

```

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

```

Exceptions:


TLB Refill, TLB Invalid, Bus Error, Address Error, Watch

Programming Notes:

The architecture provides no direct support for treating unaligned words as unsigned values, that is, zeroing bits 63..32 of the destination register when bit 31 is loaded.

Historical Information:

In the MIPS I architecture, the LWL and LWR instructions were exceptions to the load-delay scheduling restriction. A LWL or LWR instruction which was immediately followed by another LWL or LWR instruction, and used the same destination register would correctly merge the 1 to 4 loaded bytes with the data loaded by the previous instruction. All such restrictions were removed from the architecture in MIPS II.



Format: LWXC1 fd, index(base)

MIPS64
MIPS32 Release 2

Purpose: Load Word Indexed to Floating Point

To load a word from memory to an FPR (GPR+GPR addressing)

Description: $FPR[fd] \leftarrow \text{memory}[GPR[base] + GPR[index]]$

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 FPR *fd*. If FPRs are 64 bits wide, bits 63..32 of FPR *fs* become **UNDEFINED**. The contents of GPR *index* and GPR *base* are added to form the effective address.

Restrictions:

An Address Error exception occurs if $\text{EffectiveAddress}_{1..0} \neq 0$ (not word-aligned).

Operation:

```

vAddr ← GPR[base] + GPR[index]
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)


memword ← LoadMemory(CCA, WORD, pAddr, vAddr, DATA)

StoreFPR(fd, UNINTERPRETED_WORD,
         memword)

```

Exceptions:

TLB Refill, TLB Invalid, Address Error, Reserved Instruction, Coprocessor Unusable, Watch



Format: MADD *rs*, *rt*

MIPS32

Purpose: Multiply and Add Word to Hi,Lo

To multiply two words and add the result to Hi, Lo

Description: $(HI, LO) \leftarrow (HI, LO) + (GPR[rs] \times GPR[rt])$

The 32-bit word value in GPR *rs* is multiplied by the 32-bit word value in GPR *rt*, treating both operands as signed values, to produce a 64-bit result. The product is added to the 64-bit concatenated values of *HI* and *LO*. The most significant 32 bits of the result are written into *HI* and the least significant 32 bits are written into *LO*. No arithmetic exception occurs under any circumstances.

Restrictions:

None

This instruction does not provide the capability of writing directly to a target GPR.

Operation:


```
temp ← (HI || LO) + (GPR[rs] × GPR[rt])
HI ← temp63..32
LO ← temp31..0
```

Exceptions:

None

Programming Notes:

Where the size of the operands are known, software should place the shorter operand in GPR *rt*. This may reduce the latency of the instruction on those processors which implement data-dependent instruction latencies.



Format: MADDU *rs*, *rt*

MIPS32

Purpose: Multiply and Add Unsigned Word to Hi,Lo

To multiply two unsigned words and add the result to *HI*, *LO*.

Description: $(HI, LO) \leftarrow (HI, LO) + (GPR[rs] \times GPR[rt])$

The 32-bit word value in GPR *rs* is multiplied by the 32-bit word value in GPR *rt*, treating both operands as unsigned values, to produce a 64-bit result. The product is added to the 64-bit concatenated values of *HI* and *LO*. The most significant 32 bits of the result are written into *HI* and the least significant 32 bits are written into *LO*. No arithmetic exception occurs under any circumstances.

Restrictions:

None

This instruction does not provide the capability of writing directly to a target GPR.

Operation:


```
temp ← (HI || LO) + (GPR[rs] × GPR[rt])
HI ← temp63..32
LO ← temp31..0
```

Exceptions:

None

Programming Notes:

Where the size of the operands are known, software should place the shorter operand in GPR *rt*. This may reduce the latency of the instruction on those processors which implement data-dependent instruction latencies.



Format: MFC0 rt, rd
MFC0 rt, rd, sel

MIPS32
MIPS32

Purpose: Move from Coprocessor 0

To move the contents of a coprocessor 0 register to a general register.

Description: $GPR[rt] \leftarrow CPR[0,rd,sel]$

The contents of the coprocessor 0 register specified by the combination of *rd* and *sel* are loaded into general register *rt*. Note that not all coprocessor 0 registers support the *sel* field. In those instances, the *sel* field must be zero.

Restrictions:

The results are **UNDEFINED** if coprocessor 0 does not contain a register as specified by *rd* and *sel*.


Operation:

```
data ← CPR[0,rd,sel]
GPR[rt] ← data
```

Exceptions:

Coprocessor Unusable

Reserved Instruction



Format: MFC1 *rt*, *fs*

MIPS32

Purpose: Move Word From Floating Point

To copy a word from an FPU (CP1) general register to a GPR

Description: $GPR[rt] \leftarrow FPR[fs]$

The contents of FPR *fs* are loaded into general register *rt*.

Restrictions:

Operation:


```
data ← ValueFPR(fs, UNINTERPRETED_WORD)
GPR[rt] ← data
```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Historical Information:

For MIPS I, MIPS II, and MIPS III the contents of GPR *rt* are **UNPREDICTABLE** for the instruction immediately following MFC1.



Format: MFC2 rt, Impl
MFC2, rt, Impl, sel

MIPS32
MIPS32

The syntax shown above is an example using MFC1 as a model. The specific syntax is implementation dependent.

Purpose: Move Word From Coprocessor 2

To copy a word from a COP2 general register to a GPR

Description: $GPR[rt] \leftarrow CP2CPR[Impl]$

The contents of the coprocessor 2 register denoted by the *Impl* field are and placed into general register *rt*. The interpretation of the *Impl* field is left entirely to the Coprocessor 2 implementation and is not specified by the architecture.

Restrictions:


The results are **UNPREDICTABLE** if *Impl* specifies a coprocessor 2 register that does not exist.

Operation:

$data \leftarrow CP2CPR[Impl]$
 $GPR[rt] \leftarrow data$

Exceptions:

Coprocessor Unusable



Format: MFHC1 *rt*, *fs*

MIPS32 Release 2

Purpose: Move Word From High Half of Floating Point Register

To copy a word from the high half of an FPU (CP1) general register to a GPR

Description: $GPR[rt] \leftarrow FPR[fs]_{63..32}$

The contents of the high word of FPR *fs* are loaded into general register *rt*. This instruction is primarily intended to support 64-bit floating point units on a 32-bit CPU, but the semantics of the instruction are defined for all cases.

Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction Exception.

The results are **UNPREDICTABLE** if $Status_{FR} = 0$ and *fs* is odd.


Operation:

$data \leftarrow Value_{FPR}(fs, UNINTERPRETED_DOUBLEWORD)_{63..32}$
 $GPR[rt] \leftarrow data$

Exceptions:

Coprocessor Unusable

Reserved Instruction



Format: MFHC2 *rt*, *Impl*
MFHC2, *rt*, *rd*, *sel*

MIPS32 Release 2
MIPS32 Release 2

The syntax shown above is an example using MFHC1 as a model. The specific syntax is implementation dependent.

Purpose: Move Word From High Half of Coprocessor 2 Register

To copy a word from the high half of a COP2 general register to a GPR

Description: $GPR[rt] \leftarrow CP2CPR[Impl]_{63..32}$

The contents of the high word of the coprocessor 2 register denoted by the *Impl* field are placed into GPR *rt*. The interpretation of the *Impl* field is left entirely to the Coprocessor 2 implementation and is not specified by the architecture.

Restrictions:

The results are **UNPREDICTABLE** if *Impl* specifies a coprocessor 2 register that does not exist, or if that register is not 64 bits wide.

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction Exception.


Operation:

$data \leftarrow CP2CPR[Impl]_{63..32}$
 $GPR[rt] \leftarrow data$

Exceptions:

Coprocessor Unusable

Reserved Instruction



Format: MFHI rd

MIPS32

Purpose: Move From HI Register

To copy the special purpose *HI* register to a GPR

Description: $GPR[rd] \leftarrow HI$

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

Restrictions:

None

Operation:


$GPR[rd] \leftarrow HI$

Exceptions:

None

Historical Information:

In the MIPS I, II, and III architectures, the two instructions which follow the MFHI must not modify the *HI* register. If this restriction is violated, the result of the MFHI is **UNPREDICTABLE**. This restriction was removed in MIPS IV and MIPS32, and all subsequent levels of the architecture.



Format: MFLO rd

MIPS32

Purpose: Move From LO Register

To copy the special purpose *LO* register to a GPR

Description: $GPR[rd] \leftarrow LO$

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

Restrictions:

None

Operation:


$GPR[rd] \leftarrow LO$

Exceptions:

None

Historical Information:

In the MIPS I, II, and III architectures, the two instructions which follow the MFLO must not modify the *HI* register. If this restriction is violated, the result of the MFLO is **UNPREDICTABLE**. This restriction was removed in MIPS IV and MIPS32, and all subsequent levels of the architecture.



Format: MOV.fmt
 MOV.S fd, fs
 MOV.D fd, fs
 MOV.PS fd, fs

MIPS32
MIPS32
MIPS64, MIPS32 Release 2

Purpose: Floating Point Move

To move an FP value between FPRs

Description: $FPR[fd] \leftarrow FPR[fs]$

The value in FPR *fs* is placed into FPR *fd*. The source and destination are values in format *fmt*. In paired-single format, both the halves of the pair are copied to *fd*.

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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of MOV.PS is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:


StoreFPR(fd, fmt, ValueFPR(fs, fmt))

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation



Format: MOVF rd, rs, cc

MIPS32

Purpose: Move Conditional on Floating Point False

To test an FP condition code then conditionally move a GPR

Description: if FPConditionCode(cc) = 0 then GPR[rd] ← GPR[rs]

If the floating point condition code specified by *CC* is zero, then the contents of GPR *rs* are placed into GPR *rd*.


Restrictions:

Operation:

```
if FPConditionCode(cc) = 0 then
    GPR[rd] ← GPR[rs]
endif
```

Exceptions:

Reserved Instruction, Coprocessor Unusable



Format: MOVF.fmt
 MOVF.S fd, fs, cc
 MOVF.D fd, fs, cc
 MOVF.PS fd, fs, cc

MIPS32
MIPS32
MIPS64
MIPS32 Release 2

Purpose: Floating Point Move Conditional on Floating Point False

To test an FP condition code then conditionally move an FP value

Description: if FPConditionCode(cc) = 0 then FPR[fd] ← FPR[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* retains 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 **UNPREDICTABLE**.

MOV.FPS conditionally merges the lower half of FPR *fs* into the lower half of FPR *fd* if condition code *CC* is zero, and independently merges the upper half of FPR *fs* into the upper half of FPR *fd* if condition code *CC*+1 is zero. The *CC* field must be even; if it is odd, the result of this operation is **UNPREDICTABLE**.

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 **UNPREDICTABLE**. The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of MOV.FPS is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:


```
if FPConditionCode(cc) = 0 then
  StoreFPR(fd, fmt, ValueFPR(fs, fmt))
else
  StoreFPR(fd, fmt, ValueFPR(fd, fmt))
```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation



Format: MOVN rd, rs, rt

MIPS32

Purpose: Move Conditional on Not Zero

To conditionally move a GPR after testing a GPR value

Description: if GPR[rt] \neq 0 then GPR[rd] \leftarrow GPR[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]  $\neq$  0 then
    GPR[rd]  $\leftarrow$  GPR[rs]
endif
```

Exceptions:

None

Programming Notes:

The non-zero value tested might be the *condition true* result from the SLT, SLTI, SLTU, and SLTIU comparison instructions or a boolean value read from memory.



Format: MOVN.fmt
 MOVN.S fd, fs, rt
 MOVN.D fd, fs, rt
 MOVN.PS fd, fs, rt

MIPS32
MIPS32
MIPS64, MIPS32 Release 2

Purpose: Floating Point Move Conditional on Not Zero

To test a GPR then conditionally move an FP value

Description: if GPR[rt] \neq 0 then FPR[fd] \leftarrow FPR[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 **UNPREDICTABLE**.

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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of MOVN.PS is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:


```
if GPR[rt]  $\neq$  0 then
  StoreFPR(fd, fmt, ValueFPR(fs, fmt))
else
  StoreFPR(fd, fmt, ValueFPR(fd, fmt))
endif
```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation



Format: `MOVT rd, rs, cc`

MIPS32

Purpose: Move Conditional on Floating Point True

To test an FP condition code then conditionally move a GPR

Description: `if FPConditionCode(cc) = 1 then GPR[rd] ← GPR[rs]`

If the floating point condition code specified by *CC* is one, then the contents of GPR *rs* are placed into GPR *rd*.


Restrictions:

Operation:

```
if FPConditionCode(cc) = 1 then
  GPR[rd] ← GPR[rs]
endif
```

Exceptions:

Reserved Instruction, Coprocessor Unusable



Format: MOVT.fmt
 MOVT.S fd, fs, cc
 MOVT.D fd, fs, cc
 MOVT.PS fd, fs, cc

MIPS32
MIPS32
MIPS64, MIPS32 Release 2

Purpose: Floating Point Move Conditional on Floating Point True

To test an FP condition code then conditionally move an FP value

Description: if FPConditionCode(cc) = 1 then FPR[fd] ← FPR[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.

MOVT.PS conditionally merges the lower half of FPR *fs* into the lower half of FPR *fd* if condition code *CC* is one, and independently merges the upper half of FPR *fs* into the upper half of FPR *fd* if condition code *CC*+1 is one. The *CC* field should be even; if it is odd, the result of this operation is **UNPREDICTABLE**.

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 **UNPREDICTABLE**. The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of MOVT.PS is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:


```
if FPConditionCode(cc) = 0 then
  StoreFPR(fd, fmt, ValueFPR(fs, fmt))
else
  StoreFPR(fd, fmt, ValueFPR(fd, fmt))
endif
```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation



Format: MOVZ rd, rs, rt

MIPS32

Purpose: Move Conditional on Zero

To conditionally move a GPR after testing a GPR value

Description: if GPR[rt] = 0 then GPR[rd] ← GPR[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:

None

Programming Notes:

The zero value tested might be the *condition false* result from the SLT, SLTI, SLTU, and SLTIU comparison instructions or a boolean value read from memory.



Format: MOVZ.fmt
 MOVZ.S fd, fs, rt
 MOVZ.D fd, fs, rt
 MOVZ.PS fd, fs, rt

MIPS32
MIPS32
MIPS64, MIPS32 Release 2

Purpose: Floating Point Move Conditional on Zero

To test a GPR then conditionally move an FP value

Description: if GPR[rt] = 0 then FPR[fd] ← FPR[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 **UNPREDICTABLE**.

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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of MOVZ.PS is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

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 Exceptions:

Unimplemented Operation



Format: MTC0 rt, rd
MTC0 rt, rd, sel

MIPS32
MIPS32

Purpose: Move to Coprocessor 0

To move the contents of a general register to a coprocessor 0 register.

Description: $CPR[0, rd, sel] \leftarrow GPR[rt]$

The contents of general register *rt* are loaded into the coprocessor 0 register specified by the combination of *rd* and *sel*. Not all coprocessor 0 registers support the *sel* field. In those instances, the *sel* field must be set to zero.

Restrictions:

The results are **UNDEFINED** if coprocessor 0 does not contain a register as specified by *rd* and *sel*.


Operation:

$data \leftarrow GPR[rt]$
 $CPR[0, rd, sel] \leftarrow data$

Exceptions:

Coprocessor Unusable

Reserved Instruction



Format: MTC1 *rt*, *fs*

MIPS32

Purpose: Move Word to Floating Point

To copy a word from a GPR to an FPU (CP1) general register

Description: $FPR[fs] \leftarrow GPR[rt]$

The low word in GPR *rt* is placed into the low word of FPR *fs*.

Restrictions:

Operation:


```
data ← GPR[rt]31..0
StoreFPR(fs, UNINTERPRETED_WORD, data)
```

Exceptions:

Coprocessor Unusable

Historical Information:

For MIPS I, MIPS II, and MIPS III the value of FPR *fs* is **UNPREDICTABLE** for the instruction immediately following MTC1.



Format: MTC2 rt, Impl
MTC2 rt, Impl, sel

MIPS32
MIPS32

The syntax shown above is an example using MTC1 as a model. The specific syntax is implementation dependent.

Purpose: Move Word to Coprocessor 2

To copy a word from a GPR to a COP2 general register

Description: $CP2CPR[Impl] \leftarrow GPR[rt]$

The low word in GPR *rt* is placed into the low word of coprocessor 2 general register denoted by the *Impl* field. The interpretation of the *Impl* field is left entirely to the Coprocessor 2 implementation and is not specified by the architecture.

Restrictions:

The results are **UNPREDICTABLE** if *Impl* specifies a coprocessor 2 register that does not exist.


Operation:

$data \leftarrow GPR[rt]$
 $CP2CPR[Impl] \leftarrow data$

Exceptions:

Coprocessor Unusable

Reserved Instruction



Format: MTHC1 *rt*, *fs*

MIPS32 Release 2

Purpose: Move Word to High Half of Floating Point Register

To copy a word from a GPR to the high half of an FPU (CP1) general register

Description: $FPR[fs]_{63..32} \leftarrow GPR[rt]$

The word in GPR *rt* is placed into the high word of FPR *fs*. This instruction is primarily intended to support 64-bit floating point units on a 32-bit CPU, but the semantics of the instruction are defined for all cases.

Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction Exception.

The results are **UNPREDICTABLE** if $Status_{FR} = 0$ and *fs* is odd.

Operation:

```
newdata ← GPR[rt]
olddata ← ValueFPR(fs, UNINTERPRETED_DOUBLEWORD)_{31..0}
StoreFPR(fs, UNINTERPRETED_DOUBLEWORD, newdata || olddata)
```


Exceptions:

Coprocessor Unusable

Reserved Instruction

Programming Notes

When paired with MTC1 to write a value to a 64-bit FPR, the MTC1 must be executed first, followed by the MTHC1. This is because of the semantic definition of MTC1, which is not aware that software will be using an MTHC1 instruction to complete the operation, and sets the upper half of the 64-bit FPR to an **UNPREDICTABLE** value.



Format: MTHC2 *rt*, *Impl*
MTHC2 *rt*, *Impl*, *sel*

MIPS32 Release 2
MIPS32 Release 2

The syntax shown above is an example using MTHC1 as a model. The specific syntax is implementation dependent.

Purpose: Move Word to High Half of Coprocessor 2 Register

To copy a word from a GPR to the high half of a COP2 general register

Description: $CP2CPR[Impl]_{63..32} \leftarrow GPR[rt]$

The word in GPR *rt* is placed into the high word of coprocessor 2 general register denoted by the *Impl* field. The interpretation of the *Impl* field is left entirely to the Coprocessor 2 implementation and is not specified by the architecture.

Restrictions:

The results are **UNPREDICTABLE** if *Impl* specifies a coprocessor 2 register that does not exist, or if that register is not 64 bits wide.

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction Exception.

Operation:

```
data ← GPR[rt]
CP2CPR[Impl] ← data || CPR[2,rd,sel]_{31..0}
```


Exceptions:

Coprocessor Unusable

Reserved Instruction

Programming Notes

When paired with MTC2 to write a value to a 64-bit CPR, the MTC2 must be executed first, followed by the MTHC2. This is because of the semantic definition of MTC2, which is not aware that software will be using an MTHC2 instruction to complete the operation, and sets the upper half of the 64-bit CPR to an **UNPREDICTABLE** value.



Format: MTHI rs

MIPS32

Purpose: Move to HI Register

To copy a GPR to the special purpose *HI* register

Description: $HI \leftarrow GPR[rs]$

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

Restrictions:

A computed result written to the *HI/LO* pair by DIV, DIVU, MULT, or MULTU must be read by MFHI or MFLO before a new result can be written into either *HI* or *LO*.

If an MTHI instruction is executed following one of these arithmetic instructions, but before an MFLO or MFHI instruction, the contents of *LO* are **UNPREDICTABLE**. The following example shows this illegal situation:

```
MULT  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 UNPREDICTABLE value
```

Operation:


$HI \leftarrow GPR[rs]$

Exceptions:

None

Historical Information:

In MIPS I-III, if either of the two preceding instructions is MFHI, the result of that MFHI is **UNPREDICTABLE**. Reads of the *HI* or *LO* special register must be separated from any subsequent instructions that write to them by two or more instructions. In MIPS IV and later, including MIPS32 and MIPS64, this restriction does not exist.



Format: MTLO rs

MIPS32

Purpose: Move to LO Register

To copy a GPR to the special purpose *LO* register

Description: $LO \leftarrow GPR[rs]$

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

Restrictions:

A computed result written to the *HI/LO* pair by DIV, DIVU, MULT, or MULTU must be read by MFHI or MFLO before a new result can be written into either *HI* or *LO*.

If an MTLO instruction is executed following one of these arithmetic instructions, but before an MFLO or MFHI instruction, the contents of *HI* are **UNPREDICTABLE**. The following example shows this illegal situation:

```
MULT  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 UNPREDICTABLE value
```

Operation:


$LO \leftarrow GPR[rs]$

Exceptions:

None

Historical Information:

In MIPS I-III, if either of the two preceding instructions is MFHI, the result of that MFHI is **UNPREDICTABLE**. Reads of the *HI* or *LO* special register must be separated from any subsequent instructions that write to them by two or more instructions. In MIPS IV and later, including MIPS32 and MIPS64, this restriction does not exist.



Format: MUL rd, rs, rt

MIPS32

Purpose: Multiply Word to GPR

To multiply two words and write the result to a GPR.

Description: $GPR[rd] \leftarrow GPR[rs] \times GPR[rt]$

The 32-bit word value in GPR *rs* is multiplied by the 32-bit value in GPR *rt*, treating both operands as signed values, to produce a 64-bit result. The least significant 32 bits of the product are written to GPR *rd*. The contents of *HI* and *LO* are **UNPREDICTABLE** after the operation. No arithmetic exception occurs under any circumstances.

Restrictions:

Note that this instruction does not provide the capability of writing the result to the *HI* and *LO* registers.

Operation:

```
temp ← GPR[rs] × GPR[rt]
GPR[rd] ← temp31..0
HI ← UNPREDICTABLE
LO ← UNPREDICTABLE
```

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 GPR *rd* before the results are written interlocks 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.

Where the size of the operands are known, software should place the shorter operand in GPR *rt*. This may reduce the latency of the instruction on those processors which implement data-dependent instruction latencies.



Format: MUL.fmt
 MUL.S fd, fs, ft
 MUL.D fd, fs, ft
 MUL.PS fd, fs, ft

MIPS32
MIPS32
MIPS64
MIPS32 Release 2

Purpose: Floating Point Multiply

To multiply FP values

Description: $FPR[fd] \leftarrow FPR[fs] \times FPR[ft]$

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*. MUL.PS multiplies the upper and lower halves of FPR *fs* and FPR *ft* independently, and ORs together any generated exceptional conditions.

Restrictions:

The fields *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *fmt*; if they are not valid, the result is **UNPREDICTABLE**.

The operands must be values in format *fmt*; if they are not, the result is **UNPREDICTABLE** and the value of the operand FPRs becomes **UNPREDICTABLE**.

The result of MUL.PS is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:


$$\text{StoreFPR}(fd, fmt, \text{ValueFPR}(fs, fmt) \times_{fmt} \text{ValueFPR}(ft, fmt))$$

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Unimplemented Operation, Invalid Operation, Overflow, Underflow



Format: MULT *rs*, *rt*

MIPS32

Purpose: Multiply Word

To multiply 32-bit signed integers

Description: (HI, LO) \leftarrow GPR[*rs*] \times GPR[*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:

None

Operation:

$$\begin{aligned} \text{prod} &\leftarrow \text{GPR}[\text{rs}]_{31..0} \times \text{GPR}[\text{rt}]_{31..0} \\ \text{LO} &\leftarrow \text{prod}_{31..0} \\ \text{HI} &\leftarrow \text{prod}_{63..32} \end{aligned}$$

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 interlocks 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.

Where the size of the operands are known, software should place the shorter operand in GPR *rt*. This may reduce the latency of the instruction on those processors which implement data-dependent instruction latencies.



Format: MULTU *rs*, *rt*

MIPS32

Purpose: Multiply Unsigned Word

To multiply 32-bit unsigned integers

Description: $(HI, LO) \leftarrow GPR[rs] \times GPR[rt]$

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:

None

Operation:

$$\begin{aligned} \text{prod} &\leftarrow (0 \parallel GPR[rs]_{31..0}) \times (0 \parallel GPR[rt]_{31..0}) \\ LO &\leftarrow \text{prod}_{31..0} \\ HI &\leftarrow \text{prod}_{63..32} \end{aligned}$$

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 interlocks 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.

Where the size of the operands are known, software should place the shorter operand in GPR *rt*. This may reduce the latency of the instruction on those processors which implement data-dependent instruction latencies.



Format: NEG.fmt
 NEG.S fd, fs
 NEG.D fd, fs
 NEG.PS fd, fs

MIPS32
MIPS32
MIPS64, MIPS32 Release 2

Purpose: Floating Point Negate

To negate an FP value

Description: $FPR[fd] \leftarrow -FPR[fs]$

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*. NEG.PS negates the upper and lower halves of FPR *fs* independently, and ORs together any generated exceptional conditions.

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 **UNPREDICTABLE**. The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of NEG.PS is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:


StoreFPR(fd, fmt, Negate(ValueFPR(fs, fmt)))

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Unimplemented Operation, Invalid Operation



Format: NMADD.fmt
 NMADD.S fd, fr, fs, ft
 NMADD.D fd, fr, fs, ft
 NMADD.PS fd, fr, fs, ft

MIPS64, MIPS32 Release 2
MIPS64, MIPS32 Release 2
MIPS64, MIPS32 Release 2

Purpose: Floating Point Negative Multiply Add

To negate a combined multiply-then-add of FP values

Description: $FPR[fd] \leftarrow - ((FPR[fs] \times FPR[ft]) + FPR[fr])$

The value in FPR *fs* is multiplied by the value in FPR *ft* to produce an intermediate product. The intermediate product is rounded according to the current rounding mode in *FCSR*. 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 results and flags are as if separate floating-point multiply and add and negate instructions were executed.

NMADD.PS applies the operation to the upper and lower halves of FPR *fr*, FPR *fs*, and FPR *ft* independently, and ORs together any generated exceptional conditions.

Cause bits are ORed into the *Flag* bits if no exception is taken.

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 **UNPREDICTABLE**.

The operands must be values in format *fmt*; if they are not, the result is **UNPREDICTABLE** and the value of the operand FPRs becomes **UNPREDICTABLE**.

The result of NMADD.PS is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:


```
vfr ← ValueFPR(fr, fmt)
vfs ← ValueFPR(fs, fmt)
vft ← ValueFPR(ft, fmt)
StoreFPR(fd, fmt, -(vfr +fmt (vfs ×fmt vft)))
```

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Unimplemented Operation, Invalid Operation, Overflow, Underflow



Format: NMSUB.fmt
 NMSUB.S fd, fr, fs, ft
 NMSUB.D fd, fr, fs, ft
 NMSUB.PS fd, fr, fs, ft

MIPS64, MIPS32 Release 2
MIPS64, MIPS32 Release 2
MIPS64, MIPS32 Release 2

Purpose: Floating Point Negative Multiply Subtract

To negate a combined multiply-then-subtract of FP values

Description: $FPR[fd] \leftarrow - ((FPR[fs] \times FPR[ft]) - FPR[fr])$

The value in FPR *fs* is multiplied by the value in FPR *ft* to produce an intermediate product. The intermediate product is rounded according to the current rounding mode in *FCSR*. 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 results and flags are as if separate floating-point multiply and subtract and negate instructions were executed.

NMSUB.PS applies the operation to the upper and lower halves of FPR *fr*, FPR *fs*, and FPR *ft* independently, and ORs together any generated exceptional conditions.

Cause bits are ORed into the *Flag* bits if no exception is taken.

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 **UNPREDICTABLE**.

The operands must be values in format *fmt*; if they are not, the result is **UNPREDICTABLE** and the value of the operand FPRs becomes **UNPREDICTABLE**.

The result of NMSUB.PS is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:

$vfr \leftarrow \text{ValueFPR}(fr, fmt)$
 $vfs \leftarrow \text{ValueFPR}(fs, fmt)$
 $vft \leftarrow \text{ValueFPR}(ft, fmt)$
 $\text{StoreFPR}(fd, fmt, -((vfs \times_{fmt} vft) -_{fmt} vfr))$

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Unimplemented Operation, Invalid Operation, Overflow, Underflow

31	26 25	21 20	16 15	11 10	6 5	0
SPECIAL 000000	0 00000	0 00000	0 00000	0 00000	0 00000	SLL 000000
6	5	5	5	5	5	6

Format: NOP

Assembly Idiom

Purpose: No Operation

To perform no operation.

Description:

NOP is the assembly idiom used to denote no operation. The actual instruction is interpreted by the hardware as SLL r0, r0, 0.

Restrictions:

None

Operation:


None

Exceptions:

None

Programming Notes:

The zero instruction word, which represents SLL, r0, r0, 0, is the preferred NOP for software to use to fill branch and jump delay slots and to pad out alignment sequences.



Format: NOR *rd*, *rs*, *rt*

MIPS32

Purpose: Not Or

To do a bitwise logical NOT OR

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

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*

MIPS32

Purpose: Or

To do a bitwise logical OR

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

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*

MIPS32

Purpose: Or Immediate

To do a bitwise logical OR with a constant

Description: $GPR[rt] \leftarrow GPR[rs] \text{ or } immediate$

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:

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

Exceptions:

None

31	26 25 24	21 20	16 15	11 10	6 5	0
SPECIAL 000000	0 00000	0 00000	0 00000	5 00101	SLL 000000	
6	5	5	5	5	6	

Format: PAUSE

MIPS32 Release 2/MT ASE

Purpose: Wait for the LLBit to clear

Description:

Locks implemented using the LL/SC instructions are a common method of synchronization between threads of control. A typical lock implementation does a load-linked instruction and checks the value returned to determine whether the software lock is set. If it is, the code branches back to retry the load-linked instruction, thereby implementing an active busy-wait sequence. The PAUSE instruction is intended to be placed into the busy-wait sequence to block the instruction stream until such time as the load-linked instruction has a chance to succeed in obtaining the software lock.

The precise behavior of the PAUSE instruction is implementation-dependent, but it usually involves descheduling the instruction stream until the LLBit is zero. In a single-threaded processor, this may be implemented as a short-term WAIT operation which resumes at the next instruction when the LLBit is zero or on some other external event such as an interrupt. On a multi-threaded processor, this may be implemented as a short term YIELD operation which resumes at the next instruction when the LLBit is zero. In either case, it is assumed that the instruction stream which gives up the software lock does so via a write to the lock variable, which causes the processor to clear the LLBit as seen by this thread of execution.

The encoding of the instruction is such that it is backward compatible with all previous implementations of the architecture. The PAUSE instruction can therefore be placed into existing lock sequences and treated as a NOP by the processor, even if the processor does not implement the PAUSE instruction.

Restrictions:

The operation of the processor is **UNPREDICTABLE** if a PAUSE instruction is placed in the delay slot of a branch or a jump.

Operation:

```

if LLBit ≠ 0 then
    EPC ← PC + 4                /* Resume at the following instruction */
    DescheduleInstructionStream()
endif

```

Exceptions:

None

Programming Notes:

The PAUSE instruction is intended to be inserted into the instruction stream after an LL instruction has set the LLBit and found the software lock set. The program may wait forever if a PAUSE instruction is executed and there is no possibility that the LLBit will ever be cleared.

An example use of the PAUSE instruction is included in the following example:

```


acquire_lock:

```

```
ll    t0, 0(a0)           /* Read software lock, set hardware lock */
bnez  t0, acquire_lock_retry: /* Branch if software lock is taken */
addiu t0, t0, 1           /* Set the software lock */
sc    t0, 0(a0)           /* Try to store the software lock */
bnez  t0, 10f             /* Branch if lock acquired successfully */
sync
acquire_lock_retry:
pause                               /* Wait for LLBIT to clear before retry */
b     acquire_lock             /* and retry the operation */
nop
10:

    Critical region code

release_lock:
sync
sw    zero, 0(a0)          /* Release software lock, clearing LLBIT */
                                /* for any PAUSEd waiters */
```



Format: PLL.PS *fd*, *fs*, *ft*

MIPS64, MIPS32 Release 2

Purpose: Pair Lower Lower

To merge a pair of paired single values with realignment

Description: $FPR[fd] \leftarrow \text{lower}(FPR[fs]) \ || \ \text{lower}(FPR[ft])$

A new paired-single value is formed by catenating the lower single of FPR *fs* (bits **31..0**) and the lower single of FPR *ft* (bits **31..0**).

The move is non-arithmetic; it causes no IEEE 754 exceptions.

Restrictions:

The fields *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *PS*. If they are not valid, the result is **UNPREDICTABLE**.


The result of this instruction is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:

$\text{StoreFPR}(fd, PS, \text{ValueFPR}(fs, PS)_{31..0} \ || \ \text{ValueFPR}(ft, PS)_{31..0})$

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: PLU.PS *fd*, *fs*, *ft*

MIPS64, MIPS32 Release 2

Purpose: Pair Lower Upper

To merge a pair of paired single values with realignment

Description: $FPR[fd] \leftarrow \text{lower}(FPR[fs]) \ || \ \text{upper}(FPR[ft])$

A new paired-single value is formed by catenating the lower single of FPR *fs* (bits **31..0**) and the upper single of FPR *ft* (bits **63..32**).

The move is non-arithmetic; it causes no IEEE 754 exceptions.

Restrictions:

The fields *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *PS*. If they are not valid, the result is **UNPREDICTABLE**.


The result of this instruction is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:

$\text{StoreFPR}(fd, PS, \text{ValueFPR}(fs, PS)_{31..0} \ || \ \text{ValueFPR}(ft, PS)_{63..32})$

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: `PREF hint,offset(base)`

MIPS32

Purpose: Prefetch

To move data between memory and cache.

Description: `prefetch_memory(GPR[base] + offset)`

PREF adds the 16-bit signed *offset* to the contents of GPR *base* to form an effective byte address. The *hint* field supplies information about the way that the data is expected to be used.

PREF enables the processor to take some action, typically causing data to be moved to or from the 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 as long as it does not change architecturally visible state or alter the meaning of a program. Implementations are expected either to do nothing, or to take an action that increases the performance of the program. The PrepareForStore function is unique in that it may modify the architecturally visible state.

PREF does not cause addressing-related exceptions, including TLB exceptions. If the address specified would cause an addressing exception, the exception condition is ignored and no data movement occurs. However even if no data is moved, some action that is not architecturally visible, such as writeback of a dirty cache line, can take place.

It is implementation dependent whether a Bus Error or Cache Error exception is reported if such an error is detected as a byproduct of the action taken by the PREF instruction.

PREF neither generates a memory operation nor modifies the state of a cache line for a location with an *uncached* memory access type, whether this type is specified by the address segment (e.g., *kseg1*), the programmed cacheability and coherency attribute of a segment (e.g., the use of the K0, KU, or K23 fields in the *Config* register), or the per-page cacheability and coherency attribute provided by the TLB.

If PREF results in a memory operation, the memory access type and cacheability&coherency attribute used for the operation are determined by the memory access type and cacheability&coherency attribute 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.

For a cached location, the expected and useful action for the processor is 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.

In coherent multiprocessor implementations, if the effective address uses a coherent Cacheability and Coherency Attribute (CCA), then the instruction causes a coherent memory transaction to occur. This means a prefetch issued on one processor can cause data to be evicted from the cache in another processor.

The PREF instruction and the memory transactions which are sourced by the PREF instruction, such as cache refill or cache writeback, obey the ordering and completion rules of the SYNC instruction.

Table 3.30 Values of *hint* Field for PREF Instruction

Value	Name	Data Use and Desired Prefetch Action
0	load	Use: Prefetched data is expected to be read (not modified). Action: Fetch data as if for a load.

Table 3.30 Values of *hint* Field for PREF Instruction

1	store	Use: Prefetched data is expected to be stored or modified. Action: Fetch data as if for a store.
2-3	Reserved	Reserved for future use - not available to implementations.
4	load_streamed	Use: Prefetched data is expected to be read (not modified) but not reused extensively; it “streams” through cache. Action: Fetch data as if for a load and place it in the cache so that it does not displace data prefetched as “retained.”
5	store_streamed	Use: Prefetched data is expected to be stored or modified but not reused extensively; it “streams” through cache. Action: Fetch data as if for a store and place it in the cache so that it does not displace data prefetched as “retained.”
6	load_retained	Use: Prefetched data is expected to be read (not modified) and reused extensively; it should be “retained” in the cache. Action: Fetch data as if for a load and place it in the cache so that it is not displaced by data prefetched as “streamed.”
7	store_retained	Use: Prefetched data is expected to be stored or modified and reused extensively; it should be “retained” in the cache. Action: Fetch data as if for a store and place it in the cache so that it is not displaced by data prefetched as “streamed.”
8-20	Reserved	Reserved for future use - not available to implementations.
21-24	Implementation Dependent	Unassigned by the Architecture - available for implementation-dependent use.
25	writeback_invalidate (also known as “nudge”)	Use: Data is no longer expected to be used. Action: For a writeback cache, schedule a writeback of any dirty data. At the completion of the writeback, mark the state of any cache lines written back as invalid. If the cache line is not dirty, it is implementation dependent whether the state of the cache line is marked invalid or left unchanged. If the cache line is locked, no action is taken.
26-29	Implementation Dependent	Unassigned by the Architecture - available for implementation-dependent use.

Table 3.30 Values of *hint* Field for PREF Instruction

30	PrepareForStore	<p>Use: Prepare the cache for writing an entire line, without the overhead involved in filling the line from memory.</p> <p>Action: If the reference hits in the cache, no action is taken. If the reference misses in the cache, a line is selected for replacement, any valid and dirty victim is written back to memory, the entire line is filled with zero data, and the state of the line is marked as valid and dirty.</p> <p>Programming Note: Because the cache line is filled with zero data on a cache miss, software must not assume that this action, in and of itself, can be used as a fast bzero-type function.</p>
31	Implementation Dependent	Unassigned by the Architecture - available for implementation-dependent use.

Restrictions:

None

Operation:

```
vAddr ← GPR[base] + sign_extend(offset)
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
Prefetch(CCA, pAddr, vAddr, DATA, hint)
```

Exceptions:

Bus Error, Cache Error

Prefetch does not take any TLB-related or address-related exceptions under any circumstances.

Programming Notes:


Prefetch cannot move data to or 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. A prefetch may be used using an address pointer before the validity of the pointer is determined without worrying about an addressing exception.

It is implementation dependent whether a Bus Error or Cache Error exception is reported if such an error is detected as a byproduct of the action taken by the PREF instruction. Typically, this only occurs in systems which have high-reliability requirements.

Prefetch operations have no effect on cache lines that were previously locked with the CACHE instruction.

Hint field encodings whose function is described as “streamed” or “retained” convey usage intent from software to hardware. Software should not assume that hardware will always prefetch data in an optimal way. If data is to be truly retained, software should use the Cache instruction to lock data into the cache.



Format: `PREFX hint, index(base)`

MIPS64
MIPS32 Release 2

Purpose: Prefetch Indexed

To move data between memory and cache.

Description: `prefetch_memory[GPR[base] + GPR[index]]`

PREFX adds the contents of GPR *index* to the contents of GPR *base* to form an effective byte address. The *hint* field supplies information about the way the data is expected to be used.

The only functional difference between the PREF and PREFX instructions is the addressing mode implemented by the two. Refer to the [PREF](#) instruction for all other details, including the encoding of the *hint* field.

Restrictions:

Operation:

```
vAddr ← GPR[base] + GPR[index]
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, LOAD)
Prefetch(CCA, pAddr, vAddr, DATA, hint)
```


Exceptions:

Coprocessor Unusable, Reserved Instruction, Bus Error, Cache Error

Programming Notes:

The PREFX instruction is only available on processors that implement floating point and should never be generated by compilers in situations other than those in which the corresponding load and store indexed floating point instructions are generated.

Also refer to the corresponding section in the [PREF](#) instruction description.



Format: PUL.PS *fd*, *fs*, *ft*

MIPS64, MIPS32 Release 2

Purpose: Pair Upper Lower

To merge a pair of paired single values with realignment

Description: $FPR[fd] \leftarrow upper(FPR[fs]) \ || \ lower(FPR[ft])$

A new paired-single value is formed by concatenating the upper single of FPR *fs* (bits **63..32**) and the lower single of FPR *ft* (bits **31..0**).

The move is non-arithmetic; it causes no IEEE 754 exceptions.

Restrictions:

The fields *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *PS*. If they are not valid, the result is **UNPREDICTABLE**.


The result of this instruction is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:

$StoreFPR(fd, PS, ValueFPR(fs, PS)_{63..32} \ || \ ValueFPR(ft, PS)_{31..0})$

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: PUU.PS *fd*, *fs*, *ft*

MIPS64, MIPS32 Release 2

Purpose: Pair Upper Upper

To merge a pair of paired single values with realignment

Description: $FPR[fd] \leftarrow upper(FPR[fs]) \ || \ upper(FPR[ft])$

A new paired-single value is formed by concatenating the upper single of FPR *fs* (bits **63..32**) and the upper single of FPR *ft* (bits **63..32**).

The move is non-arithmetic; it causes no IEEE 754 exceptions.

Restrictions:

The fields *fs*, *ft*, and *fd* must specify FPRs valid for operands of type *PS*. If they are not valid, the result is **UNPREDICTABLE**.


The result of this instruction is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:

$StoreFPR(fd, PS, ValueFPR(fs, PS)_{63..32} \ || \ ValueFPR(ft, PS)_{63..32})$

Exceptions:

Coprocessor Unusable, Reserved Instruction



Format: RDHWR *rt*,*rd*

MIPS32 Release 2

Purpose: Read Hardware Register

To move the contents of a hardware register to a general purpose register (GPR) if that operation is enabled by privileged software.

Description: GPR[*rt*] ← HWR[*rd*]

If access is allowed to the specified hardware register, the contents of the register specified by *rd* is loaded into general register *rt*. Access control for each register is selected by the bits in the coprocessor 0 *HWREna* register.

The available hardware registers, and the encoding of the *rd* field for each, are shown in Table 3.31.

Table 3.31 RDHWR Register Numbers

Register Number (<i>rd</i> Value)	Mnemonic	Description										
0	CPUNum	Number of the CPU on which the program is currently running. This register provides read access to the coprocessor 0 <i>EBase</i> _{CPUNum} field.										
1	SYNCL_Step	Address step size to be used with the SYNCL instruction, or zero if no caches need be synchronized. See that instruction's description for the use of this value.										
2	CC	High-resolution cycle counter. This register provides read access to the coprocessor 0 <i>Count</i> Register.										
3	CCRes	Resolution of the CC register. This value denotes the number of cycles between update of the register. For example: <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>CCRes Value</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">1</td> <td>CC register increments every CPU cycle</td> </tr> <tr> <td style="text-align: center;">2</td> <td>CC register increments every second CPU cycle</td> </tr> <tr> <td style="text-align: center;">3</td> <td>CC register increments every third CPU cycle</td> </tr> <tr> <td colspan="2" style="text-align: center;">etc.</td> </tr> </tbody> </table>	CCRes Value	Meaning	1	CC register increments every CPU cycle	2	CC register increments every second CPU cycle	3	CC register increments every third CPU cycle	etc.	
CCRes Value	Meaning											
1	CC register increments every CPU cycle											
2	CC register increments every second CPU cycle											
3	CC register increments every third CPU cycle											
etc.												
4-28		These registers numbers are reserved for future architecture use. Access results in a Reserved Instruction Exception.										
29	ULR	User Local Register. This register provides read access to the coprocessor 0 <i>UserLocal</i> register, if it is implemented. In some operating environments, the <i>UserLocal</i> register is a pointer to a thread-specific storage block.										
30-31		These register numbers are reserved for implementation-dependent use. If they are not implemented, access results in a Reserved Instruction Exception.										

Restrictions:

In implementations of Release 1 of the Architecture, this instruction resulted in a Reserved Instruction Exception.


Access to the specified hardware register is enabled if Coprocessor 0 is enabled, or if the corresponding bit is set in the *HWREna* register. If access is not allowed or the register is not implemented, a Reserved Instruction Exception is signaled.

Operation:

```
case rd
  0: temp ← EBaseCPUNum
  1: temp ← SYNCI_StepSize()
  2: temp ← Count
  3: temp ← CountResolution()
  29: temp ← UserLocal
  30: temp ← Implementation-Dependent-Value
  31: temp ← Implementation-Dependent-Value
  otherwise: SignalException(ReservedInstruction)
endcase
GPR[rt] ← temp
```

Exceptions:

Reserved Instruction



Format: RDPGPR rd, rt

MIPS32 Release 2

Purpose: Read GPR from Previous Shadow Set

To move the contents of a GPR from the previous shadow set to a current GPR.

Description: $GPR[rd] \leftarrow SGPR[SRSCtl_{PSS}, rt]$

The contents of the shadow GPR register specified by $SRSCtl_{PSS}$ (signifying the previous shadow set number) and rt (specifying the register number within that set) is moved to the current GPR rd .

Restrictions:

In implementations prior to Release 2 of the Architecture, this instruction resulted in a Reserved Instruction Exception.


Operation:

$GPR[rd] \leftarrow SGPR[SRSCtl_{PSS}, rt]$

Exceptions:

Coprocessor Unusable

Reserved Instruction



Format: RECIP.fmt
 RECIP.S fd, fs
 RECIP.D fd, fs

MIPS64, MIPS32 Release 2
MIPS64, MIPS32 Release 2

Purpose: Reciprocal Approximation

To approximate the reciprocal of an FP value (quickly)

Description: $FPR[fd] \leftarrow 1.0 / FPR[fs]$

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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Operation:


StoreFPR(fd, fmt, 1.0 / valueFPR(fs, fmt))

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Division-by-zero, Unimplemented Op, Invalid Op, Overflow, Underflow



Format: ROTR *rd*, *rt*, *sa*

SmartMIPS Crypto, MIPS32 Release 2

Purpose: Rotate Word Right

To execute a logical right-rotate of a word by a fixed number of bits

Description: $GPR[rd] \leftarrow GPR[rt] \leftrightarrow(\text{right}) sa$

The contents of the low-order 32-bit word of GPR *rt* are rotated right; the word result is placed in GPR *rd*. The bit-rotate amount is specified by *sa*.

Restrictions:

Operation:


```

if ((ArchitectureRevision() < 2) and (Config3SM = 0)) then
    UNPREDICTABLE
endif
s ← sa
temp ← GPR[rt]s-1..0 || GPR[rt]31..s
GPR[rd] ← temp

```

Exceptions:

Reserved Instruction



Format: ROTRV *rd*, *rt*, *rs*

SmartMIPS Crypto, MIPS32 Release 2

Purpose: Rotate Word Right Variable

To execute a logical right-rotate of a word by a variable number of bits

Description: $GPR[rd] \leftarrow GPR[rt] \leftrightarrow(\text{right}) GPR[rs]$

The contents of the low-order 32-bit word of GPR *rt* are rotated right; the word result is placed in GPR *rd*. The bit-rotate amount is specified by the low-order 5 bits of GPR *rs*.

Restrictions:

Operation:


```

if ((ArchitectureRevision() < 2) and (Config3SM = 0)) then
  UNPREDICTABLE
endif
s ← GPR[rs]4..0
temp ← GPR[rt]s-1..0 || GPR[rt]31..s
GPR[rd] ← temp

```

Exceptions:

Reserved Instruction



Format: ROUND.L.fmt
 ROUND.L.S fd, fs
 ROUND.L.D fd, fs

MIPS64, MIPS32 Release 2
MIPS64, MIPS32 Release 2

Purpose: Floating Point Round to Long Fixed Point

To convert an FP value to 64-bit fixed point, rounding to nearest

Description: $FPR[fd] \leftarrow \text{convert_and_round}(FPR[fs])$

The value in FPR *fs*, in format *fmt*, is converted to a value in 64-bit long fixed point format and rounded 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. In this case 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*.

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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

The result of this instruction is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:


`StoreFPR(fd, L, ConvertFmt(ValueFPR(fs, fmt), fmt, L))`

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Unimplemented Operation, Invalid Operation, Overflow



Format: ROUND.W.fmt

ROUND.W.S fd, fs

ROUND.W.D fd, fs

MIPS32

MIPS32

Purpose: Floating Point Round to Word Fixed Point

To convert an FP value to 32-bit fixed point, rounding to nearest

Description: $FPR[fd] \leftarrow \text{convert_and_round}(FPR[fs])$

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. In this case 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*.

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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Operation:


`StoreFPR(fd, W, ConvertFmt(ValueFPR(fs, fmt), fmt, W))`

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Unimplemented Operation, Invalid Operation, Overflow



Format: RSQRT.fmt
 RSQRT.S fd, fs
 RSQRT.D fd, fs

MIPS64, MIPS32 Release 2
MIPS64, MIPS32 Release 2

Purpose: Reciprocal Square Root Approximation

To approximate the reciprocal of the square root of an FP value (quickly)

Description: $FPR[fd] \leftarrow 1.0 / \text{sqrt}(FPR[fs])$

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 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).

The effect of the current *FCSR* rounding mode on the result is implementation dependent.

Restrictions:

The fields *fs* and *fd* must specify FPRs valid for operands of type *fmt*; if they are not valid, the result is **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Operation:


`StoreFPR(fd, fmt, 1.0 / SquareRoot(valueFPR(fs, fmt)))`

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Division-by-zero, Unimplemented Operation, Invalid Operation, Overflow, Underflow



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

MIPS32

Purpose: Store Byte

To store a byte to memory

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

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:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddr_PSIZE-1..2 || (pAddr_1..0 xor ReverseEndian2)
bytesel ← vAddr_1..0 xor BigEndianCPU2
dataword ← GPR[rt]_31-8*bytesel..0 || 08*bytesel
StoreMemory (CCA, BYTE, dataword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Bus Error, Address Error, Watch



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

MIPS32

Purpose: Store Conditional Word

To store a word to memory to complete an atomic read-modify-write

Description: if *atomic_update* then $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{GPR}[\text{rt}]$, $\text{GPR}[\text{rt}] \leftarrow 1$
else $\text{GPR}[\text{rt}] \leftarrow 0$

The LL and SC instructions provide primitives to implement atomic read-modify-write (RMW) operations for synchronizable memory locations.

The 32-bit word in GPR *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 an effective address.

The SC completes the RMW sequence begun by the preceding LL instruction executed on the processor. To complete the RMW sequence atomically, the following occur:

- The 32-bit word of GPR *rt* is stored into memory at the location specified by the aligned effective address.
- A 1, indicating success, is written into GPR *rt*.

Otherwise, memory is not modified and a 0, indicating failure, is written into GPR *rt*.

If either of the following events occurs between the execution of LL and SC, the SC fails:

- A coherent store is completed by another processor or coherent I/O module into the block of synchronizable physical memory containing the word. The size and alignment of the block is implementation dependent, but it is at least one word and at most the minimum page size.
- An ERET instruction is executed.

If either 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 not predictable. Portable programs should not cause one of these events.

- A memory access instruction (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 is **UNPREDICTABLE**:

- Execution of SC must have been preceded by execution of an LL instruction.
- An RMW sequence executed without intervening events that would cause the SC to fail must use the same address in the LL and SC. The address is the same if the virtual address, physical address, and cacheability & coherency attribute are identical.

Atomic RMW is provided only for synchronizable memory locations. A synchronizable memory location is one that is associated with the state and logic necessary to implement the LL/SC semantics. Whether a memory location is synchronizable depends on the processor and system configurations, and on the memory access type used for the

location:

- **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, and they may not be mixed.
- **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*.
- **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.

Restrictions:

The addressed location must have a memory access type of *cached noncoherent* or *cached coherent*; if it does not, the result is **UNPREDICTABLE**.

The effective address must be naturally-aligned. If either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Operation:

```
vAddr ← sign_extend(offset) + GPR[base]
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
dataword ← GPR[rt]
if LLbit then
    StoreMemory (CCA, WORD, dataword, pAddr, vAddr, DATA)
endif
GPR[rt] ← 031 || LLbit
```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch


Programming Notes:

LL and SC are used to atomically update memory locations, as shown 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                   # branch-delay slot
```

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, and 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.



Format: SDBBP code

EJTAG

Purpose: Software Debug Breakpoint

To cause a debug breakpoint exception

Description:

This instruction causes a debug exception, passing control to the debug exception handler. If the processor is executing in Debug Mode when the SDBBP instruction is executed, the exception is a Debug Mode Exception, which sets the `Debug_DExcCode` field to the value 0x9 (Bp). The code field can be used for passing information to the debug exception handler, and is retrieved by the debug exception handler only by loading the contents of the memory word containing the instruction, using the DEPC register. The CODE field is not used in any way by the hardware.

Restrictions:

Operation:


```

If DebugDM = 0 then
    SignalDebugBreakpointException()
else
    SignalDebugModeBreakpointException()
endif

```

Exceptions:

Debug Breakpoint Exception
 Debug Mode Breakpoint Exception



Format: SDC1 ft, offset(base)

MIPS32

Purpose: Store Doubleword from Floating Point

To store a doubleword from an FPR to memory

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{FPR}[\text{ft}]$

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.

Restrictions:

An Address Error exception occurs if $\text{EffectiveAddress}_{2..0} \neq 0$ (not doubleword-aligned).

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr2..0 ≠ 03 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)
datadoubleword ← ValueFPR(ft, UNINTERPRETED_DOUBLEWORD)
paddr ← paddr xor ((BigEndianCPU xor ReverseEndian) || 02)
StoreMemory(CCA, WORD, datadoubleword31..0, pAddr, vAddr, DATA)
paddr ← paddr xor 0b100
StoreMemory(CCA, WORD, datadoubleword63..32, pAddr, vAddr+4, DATA)

```

Exceptions:

Coprocessor Unusable, Reserved Instruction, TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch



Format: SDC2 *rt*, *offset*(*base*)

MIPS32

Purpose: Store Doubleword from Coprocessor 2

To store a doubleword from a Coprocessor 2 register to memory

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{CPR}[2, \text{rt}, 0]$

The 64-bit doubleword in Coprocessor 2 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:

An Address Error exception occurs if $\text{EffectiveAddress}_{2..0} \neq 0$ (not doubleword-aligned).

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr2..0 ≠ 03 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)
lsw ← CPR[2,rt,0]
msw ← CPR[2,rt+1,0]
paddr ← paddr xor ((BigEndianCPU xor ReverseEndian) || 02)
StoreMemory(CCA, WORD, lsw, pAddr, vAddr, DATA)
paddr ← paddr xor 0b100
StoreMemory(CCA, WORD, msw, pAddr, vAddr+4, DATA)

```

Exceptions:

Coprocessor Unusable, Reserved Instruction, TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch



Format: SDXC1 fs, index(base)

MIPS64
MIPS32 Release 2

Purpose: Store Doubleword Indexed from Floating Point

To store a doubleword from an FPR to memory (GPR+GPR addressing)

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{GPR}[\text{index}]] \leftarrow \text{FPR}[\text{fs}]$

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.

Restrictions:

An Address Error exception occurs if $\text{EffectiveAddress}_{2..0} \neq 0$ (not doubleword-aligned).

Operation:


```

vAddr ← GPR[base] + GPR[index]
if vAddr2..0 ≠ 03 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)
datadoubleword ← ValueFPR(fs, UNINTERPRETED_DOUBLEWORD)
paddr ← paddr xor ((BigEndianCPU xor ReverseEndian) || 02)
StoreMemory(CCA, WORD, datadoubleword31..0, pAddr, vAddr, DATA)
paddr ← paddr xor 0b100
StoreMemory(CCA, WORD, datadoubleword63..32, pAddr, vAddr+4, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Coprocessor Unusable, Address Error, Reserved Instruction, Watch.



Format: SEB rd, rt

MIPS32 Release 2

Purpose: Sign-Extend Byte

To sign-extend the least significant byte of GPR *rt* and store the value into GPR *rd*.

Description: $GPR[rd] \leftarrow \text{SignExtend}(GPR[rt]_{7..0})$

The least significant byte from GPR *rt* is sign-extended and stored in GPR *rd*.

Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction Exception.

Operation:

$GPR[rd] \leftarrow \text{sign_extend}(GPR[rt]_{7..0})$


Exceptions:

Reserved Instruction

Programming Notes:

For symmetry with the SEB and SEH instructions, one would expect that there would be ZEB and ZEH instructions that zero-extend the source operand. Similarly, one would expect that the SEW and ZEW instructions would exist to sign- or zero-extend a word to a doubleword. These instructions do not exist because there are functionally-equivalent instructions already in the instruction set. The following table shows the instructions providing the equivalent functions.

Expected Instruction	Function	Equivalent Instruction
ZEB rx, ry	Zero-Extend Byte	ANDI rx, ry, 0xFF
ZEH rx, ry	Zero-Extend Halfword	ANDI rx, ry, 0xFFFF



Format: SEH rd, rt

MIPS32 Release 2

Purpose: Sign-Extend Halfword

To sign-extend the least significant halfword of GPR *rt* and store the value into GPR *rd*.

Description: $GPR[rd] \leftarrow \text{SignExtend}(GPR[rt]_{15..0})$

The least significant halfword from GPR *rt* is sign-extended and stored in GPR *rd*.

Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction Exception.

Operation:

$GPR[rd] \leftarrow \text{sign_extend}(GPR[rt]_{15..0})$

Exceptions:

Reserved Instruction

Programming Notes:


The SEH instruction can be used to convert two contiguous halfwords to sign-extended word values in three instructions. For example:

```
lw    t0, 0(a1)           /* Read two contiguous halfwords */
seh   t1, t0              /* t1 = lower halfword sign-extended to word */
sra   t0, t0, 16         /* t0 = upper halfword sign-extended to word */
```

Zero-extended halfwords can be created by changing the SEH and SRA instructions to ANDI and SRL instructions, respectively.

For symmetry with the SEB and SEH instructions, one would expect that there would be ZEB and ZEH instructions that zero-extend the source operand. Similarly, one would expect that the SEW and ZEW instructions would exist to sign- or zero-extend a word to a doubleword. These instructions do not exist because there are functionally-equivalent instructions already in the instruction set. The following table shows the instructions providing the equivalent functions.

Expected Instruction	Function	Equivalent Instruction
ZEB rx, ry	Zero-Extend Byte	ANDI rx, ry, 0xFF
ZEH rx, ry	Zero-Extend Halfword	ANDI rx, ry, 0xFFFF



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

MIPS32

Purpose: Store Halfword

To store a halfword to memory

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

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.

Operation:

```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr0 ≠ 0 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..2 || (pAddr1..0 xor (ReverseEndian || 0))
bytesel ← vAddr1..0 xor (BigEndianCPU || 0)
dataword ← GPR[rt]31-8*bytesel..0 || 08*bytesel
StoreMemory (CCA, HALFWORD, dataword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch

31	26 25	21 20	16 15	11 10	6 5	0
SPECIAL 000000	0 00000	rt	rd	sa	SLL 000000	
6	5	5	5	5	6	

Format: SLL rd, rt, sa

MIPS32

Purpose: Shift Word Left Logical

To left-shift a word by a fixed number of bits

Description: $GPR[rd] \leftarrow GPR[rt] \ll sa$

The contents of the low-order 32-bit word of GPR *rt* are shifted left, inserting zeros into the emptied bits; the word result is placed in GPR *rd*. The bit-shift amount is specified by *sa*.

Restrictions:

None

Operation:

```
s ← sa
temp ← GPR[rt](31-s)..0 || 0s
GPR[rd] ← temp
```


Exceptions:

None

Programming Notes:

SLL r0, r0, 0, expressed as NOP, is the assembly idiom used to denote no operation.

SLL r0, r0, 1, expressed as SSNOP, is the assembly idiom used to denote no operation that causes an issue break on superscalar processors.



Format: SLLV rd, rt, rs

MIPS32

Purpose: Shift Word Left Logical Variable

To left-shift a word by a variable number of bits

Description: $GPR[rd] \leftarrow GPR[rt] \ll rs$

The contents of the low-order 32-bit word of GPR *rt* are shifted left, inserting zeros into the emptied bits; the result word is placed in GPR *rd*. The bit-shift amount is specified by the low-order 5 bits of GPR *rs*.

Restrictions:

None

Operation:

```

s ← GPR[rs]4..0
temp ← GPR[rt](31-s)..0 || 0s
GPR[rd] ← temp


```

Exceptions:

None

Programming Notes:

None



Format: SLT rd, rs, rt

MIPS32

Purpose: Set on Less Than

To record the result of a less-than comparison

Description: $GPR[rd] \leftarrow (GPR[rs] < GPR[rt])$

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, it is 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*

MIPS32

Purpose: Set on Less Than Immediate

To record the result of a less-than comparison with a constant

Description: $GPR[rt] \leftarrow (GPR[rs] < immediate)$

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, it is 0 (false).

The arithmetic comparison does not cause an Integer Overflow exception.

Restrictions:

None

Operation:


```

if GPR[rs] < sign_extend(immediate) then
    GPR[rt] ← 0GPREN-1 || 1
else
    GPR[rt] ← 0GPREN
endif

```

Exceptions:

None



Format: SLTIU *rt*, *rs*, *immediate*

MIPS32

Purpose: Set on Less Than Immediate Unsigned

To record the result of an unsigned less-than comparison with a constant

Description: $GPR[rt] \leftarrow (GPR[rs] < immediate)$

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, it is 0 (false).

Because the 16-bit *immediate* is sign-extended before comparison, the instruction can 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[rt] ← 0GPRLEN-1 || 1
else
    GPR[rt] ← 0GPRLEN
endif

```

Exceptions:

None



Format: SLTU rd, rs, rt

MIPS32

Purpose: Set on Less Than Unsigned

To record the result of an unsigned less-than comparison

Description: $GPR[rd] \leftarrow (GPR[rs] < GPR[rt])$

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, it is 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

MIPS32

Purpose: Shift Word Right Arithmetic

To execute an arithmetic right-shift of a word by a fixed number of bits

Description: $GPR[rd] \leftarrow GPR[rt] \gg sa$ (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 amount is specified by *sa*.

Restrictions:


None

Operation:

```
s ← sa
temp ← (GPR[rt]31)s || GPR[rt]31..s
GPR[rd] ← temp
```

Exceptions:

None



Format: SRAV *rd*, *rt*, *rs*

MIPS32

Purpose: Shift Word Right Arithmetic Variable

To execute an arithmetic right-shift of a word by a variable number of bits

Description: $GPR[rd] \leftarrow GPR[rt] \gg rs$ (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 amount is specified by the low-order 5 bits of GPR *rs*.

Restrictions:

None

Operation:


```

s ← GPR[rs]4..0
temp ← (GPR[rt]31)s || GPR[rt]31..s
GPR[rd] ← temp

```

Exceptions:

None



Format: SRL rd, rt, sa

MIPS32

Purpose: Shift Word Right Logical

To execute a logical right-shift of a word by a fixed number of bits

Description: $GPR[rd] \leftarrow GPR[rt] \gg sa$ (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 amount is specified by *sa*.

Restrictions:


None

Operation:

```
s ← sa
temp ← 0s || GPR[rt]31..s
GPR[rd] ← temp
```

Exceptions:

None



Format: SRLV rd, rt, rs

MIPS32

Purpose: Shift Word Right Logical Variable

To execute a logical right-shift of a word by a variable number of bits

Description: $GPR[rd] \leftarrow GPR[rt] \gg GPR[rs]$ (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 amount is specified by the low-order 5 bits of GPR *rs*.

Restrictions:

None

Operation:

```
s ← GPR[rs]4..0
temp ← 0s || GPR[rt]31..s
GPR[rd] ← temp
```

Exceptions:

None

31	26 25	21 20	16 15	11 10	6 5	0
SPECIAL 000000	0 00000	0 00000	0 00000	1 00001	SLL 000000	
6	5	5	5	5	6	

Format: SSNOP

MIPS32

Purpose: Superscalar No Operation

Break superscalar issue on a superscalar processor.

Description:

SSNOP is the assembly idiom used to denote superscalar no operation. The actual instruction is interpreted by the hardware as SLL r0, r0, 1.

This instruction alters the instruction issue behavior on a superscalar processor by forcing the SSNOP instruction to single-issue. The processor must then end the current instruction issue between the instruction previous to the SSNOP and the SSNOP. The SSNOP then issues alone in the next issue slot.

On a single-issue processor, this instruction is a NOP that takes an issue slot.

Restrictions:

None

Operation:

None

Exceptions:

None

Programming Notes:


SSNOP is intended for use primarily to allow the programmer control over CP0 hazards by converting instructions into cycles in a superscalar processor. For example, to insert at least two cycles between an MTC0 and an ERET, one would use the following sequence:

```

mtc0   x,y
ssnop
ssnop
eret

```

Based on the normal issues rules of the processor, the MTC0 issues in cycle T. Because the SSNOP instructions must issue alone, they may issue no earlier than cycle T+1 and cycle T+2, respectively. Finally, the ERET issues no earlier than cycle T+3. Note that although the instruction after an SSNOP may issue no earlier than the cycle after the SSNOP is issued, that instruction may issue later. This is because other implementation-dependent issue rules may apply that prevent an issue in the next cycle. Processors should not introduce any unnecessary delay in issuing SSNOP instructions.



Format: SUB *rd*, *rs*, *rt*

MIPS32

Purpose: Subtract Word

To subtract 32-bit integers. If overflow occurs, then trap

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

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:

None

Operation:


```
temp ← (GPR[rs]31 || GPR[rs]31..0) - (GPR[rt]31 || GPR[rt]31..0)
if temp32 ≠ temp31 then
    SignalException(IntegerOverflow)
else
    GPR[rd] ← temp31..0
endif
```

Exceptions:

Integer Overflow

Programming Notes:

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



Format: SUBU *rd*, *rs*, *rt*

MIPS32

Purpose: Subtract Unsigned Word

To subtract 32-bit integers

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

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 and placed into GPR *rd*.

No integer overflow exception occurs under any circumstances.

Restrictions:

None

Operation:


```
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 unsigned arithmetic, such as address arithmetic, or integer arithmetic environments that ignore overflow, such as C language arithmetic.



Format: SUXC1 *fs*, *index*(*base*)

MIPS64, MIPS32 Release 2

Purpose: Store Doubleword Indexed Unaligned from Floating Point

To store a doubleword from an FPR to memory (GPR+GPR addressing) ignoring alignment

Description: $\text{memory}[(\text{GPR}[\text{base}] + \text{GPR}[\text{index}])_{\text{PSIZE}-1..3}] \leftarrow \text{FPR}[\text{fs}]$

The contents of the 64-bit doubleword in FPR *fs* is stored at the memory location specified by the effective address. The contents of GPR *index* and GPR *base* are added to form the effective address. The effective address is doubleword-aligned; EffectiveAddress_{2..0} are ignored.

Restrictions:


The result of this instruction is **UNPREDICTABLE** if the processor is executing in 16 FP registers mode.

Operation:

```
vAddr ← (GPR[base]+GPR[index])63..3 || 03
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)
datadoubleword ← ValueFPR(fs, UNINTERPRETED_DOUBLEWORD)
paddr ← paddr xor ((BigEndianCPU xor ReverseEndian) || 02)
StoreMemory(CCA, WORD, datadoubleword31..0, pAddr, vAddr, DATA)
paddr ← paddr xor 0b100
StoreMemory(CCA, WORD, datadoubleword63..32, pAddr, vAddr+4, DATA)
```

Exceptions:

Coprocessor Unusable, Reserved Instruction, TLB Refill, TLB Invalid, TLB Modified, Watch



Format: `SW rt, offset(base)`

MIPS32

Purpose: Store Word

To store a word to memory

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

The least-significant 32-bit word of 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 either of the 2 least-significant bits of the address is non-zero, an Address Error exception occurs.

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
dataword ← GPR[rt]
StoreMemory (CCA, WORD, dataword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch



SWC1 ft, offset(base)

MIPS32

Purpose: Store Word from Floating Point

To store a word from an FPR to memory

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{FPR}[\text{ft}]$

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 $\text{EffectiveAddress}_{1..0} \neq 0$ (not word-aligned).

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
if vAddr1..0 ≠ 03 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)
dataword ← ValueFPR(ft, UNINTERPRETED_WORD)
StoreMemory(CCA, WORD, dataword, pAddr, vAddr, DATA)

```

Exceptions:

Coprocessor Unusable, Reserved Instruction, TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch



Format: SWC2 rt, offset(base)

MIPS32

Purpose: Store Word from Coprocessor 2

To store a word from a COP2 register to memory

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{offset}] \leftarrow \text{CPR}[2, \text{rt}, 0]$

The low 32-bit word from COP2 (Coprocessor 2) 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:

An Address Error exception occurs if $\text{EffectiveAddress}_{1..0} \neq 0$ (not word-aligned).

Operation:

```

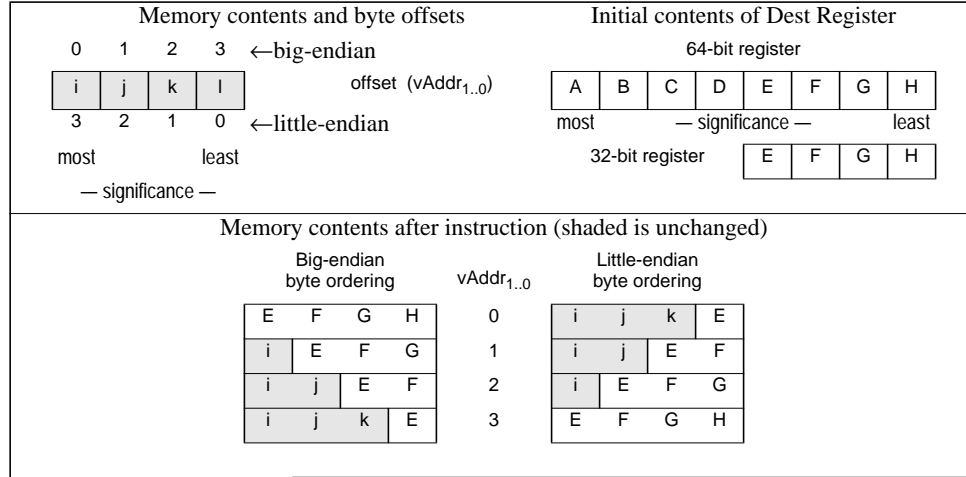
vAddr ← sign_extend(offset) + GPR[base]
if vAddr1..0 ≠ 02 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)
dataword ← CPR[2,rt,0]
StoreMemory(CCA, WORD, dataword, pAddr, vAddr, DATA)

```

Exceptions:

Coprocessor Unusable, Reserved Instruction, TLB Refill, TLB Invalid, TLB Modified, Address Error, Watch

Figure 3.10 Bytes Stored by an SWL Instruction

**Restrictions:**

None

Operation:

```


vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..2 || (pAddr1..0 xor ReverseEndian2)
If BigEndianMem = 0 then
    pAddr ← pAddrPSIZE-1..2 || 02
endif
byte ← vAddr1..0 xor BigEndianCPU2
dataword ← 024-8*byte || GPR[rt]31..24-8*byte
StoreMemory(CCA, byte, dataword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Bus Error, Address Error , Watch

Figure 3.12 Bytes Stored by SWR Instruction

**Restrictions:**

None

Operation:


```

vAddr ← sign_extend(offset) + GPR[base]
(pAddr, CCA) ← AddressTranslation (vAddr, DATA, STORE)
pAddr ← pAddrPSIZE-1..2 || (pAddr1..0 xor ReverseEndian2)
If BigEndianMem = 0 then
    pAddr ← pAddrPSIZE-1..2 || 02
endif
byte ← vAddr1..0 xor BigEndianCPU2
dataword ← GPR[rt]31-8*byte || 08*byte
StoreMemory(CCA, WORD-byte, dataword, pAddr, vAddr, DATA)

```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Bus Error, Address Error, Watch



Format: SWXC1 fs, index(base)

MIPS64
MIPS32 Release 2

Purpose: Store Word Indexed from Floating Point

To store a word from an FPR to memory (GPR+GPR addressing)

Description: $\text{memory}[\text{GPR}[\text{base}] + \text{GPR}[\text{index}]] \leftarrow \text{FPR}[\text{fs}]$

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:


An Address Error exception occurs if $\text{EffectiveAddress}_{1..0} \neq 0$ (not word-aligned).

Operation:

```
vAddr ← GPR[base] + GPR[index]
if vAddr1..0 ≠ 03 then
    SignalException(AddressError)
endif
(pAddr, CCA) ← AddressTranslation(vAddr, DATA, STORE)
dataword ← ValueFPR(fs, UNINTERPRETED_WORD)
StoreMemory(CCA, WORD, dataword, pAddr, vAddr, DATA)
```

Exceptions:

TLB Refill, TLB Invalid, TLB Modified, Address Error, Reserved Instruction, Coprocessor Unusable, Watch



Format: SYNC (stype = 0 implied)
 SYNC stype

MIPS32
MIPS32

Purpose: To order loads and stores for shared memory.

Description:

These types of ordering guarantees are available through the SYNC instruction:

- Completion Barriers
- Ordering Barriers

Simple Description for Completion Barrier:

- The barrier affects only *uncached* and *cached coherent* loads and stores.
- The specified memory instructions (loads or stores or both) that occur before the SYNC instruction must be completed before the specified memory instructions 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.

Detailed Description for Completion Barrier:

- Every synchronizable specified memory instruction (loads or stores or both) that occurs in the instruction stream before the SYNC instruction must be already globally performed before any synchronizable specified memory instructions that occur after the SYNC are allowed to be performed, with respect to any other processor or coherent I/O module.
- The barrier does not guarantee the order in which instruction fetches are performed.
- A stype value of zero will always be defined such that it performs the most complete set of synchronization operations that are defined. This means stype zero always does a completion barrier that affects both loads and stores preceding the SYNC instruction and both loads and stores that are subsequent to the SYNC instruction. Non-zero values of stype may be defined by the architecture or specific implementations to perform synchronization behaviors that are less complete than that of stype zero. If an implementation does not use one of these non-zero values to define a different synchronization behavior, then that non-zero value of stype must act the same as stype zero completion barrier. This allows software written for an implementation with a lighter-weight barrier to work on another implementation which only implements the stype zero completion barrier.
- A completion barrier is required, potentially in conjunction with SSNOP (in Release 1 of the Architecture) or EHB (in Release 2 of the Architecture), to guarantee that memory reference results are visible across operating mode changes. For example, a completion barrier is required on some implementations on entry to and exit from Debug Mode to guarantee that memory effects are handled correctly.

SYNC behavior when the stype field is zero:

- A completion barrier that affects preceding loads and stores and subsequent loads and stores.

Simple Description for Ordering Barrier:

- The barrier affects only *uncached* and *cached coherent* loads and stores.
- The specified memory instructions (loads or stores or both) that occur before the SYNC instruction must always be ordered before the specified memory instructions after the SYNC.
- Memory instructions which are ordered before other memory instructions are processed by the load/store datapath first before the other memory instructions.

Detailed Description for Ordering Barrier:

- Every synchronizable specified memory instruction (loads or stores or both) that occurs in the instruction stream before the SYNC instruction must reach a stage in the load/store datapath after which no instruction re-ordering is possible before any synchronizable specified memory instruction which occurs after the SYNC instruction in the instruction stream reaches the same stage in the load/store datapath.
- If any memory instruction before the SYNC instruction in program order, generates a memory request to the external memory and any memory instruction after the SYNC instruction in program order also generates a memory request to external memory, the memory request belonging to the older instruction must be globally performed before the time the memory request belonging to the younger instruction is globally performed.
- The barrier does not guarantee the order in which instruction fetches are performed.

As compared to the completion barrier, the ordering barrier is a lighter-weight operation as it does not require the specified instructions before the SYNC to be already completed. Instead it only requires that those specified instructions which are subsequent to the SYNC in the instruction stream are never re-ordered for processing ahead of the specified instructions which are before the SYNC in the instruction stream. This potentially reduces how many cycles the barrier instruction must stall before it completes.

The Acquire and Release barrier types are used to minimize the memory orderings that must be maintained and still have software synchronization work.

Implementations that do not use any of the non-zero values of stype to define different barriers, such as ordering barriers, must make those stype values act the same as stype zero.

For the purposes of this description, the CACHE, PREF and PREFX instructions are treated as loads and stores. That is, these instructions and the memory transactions sourced by these instructions obey the ordering and completion rules of the SYNC instruction.

Table 3.32 lists the available completion barrier and ordering barriers behaviors that can be specified using the stype field..

Table 3.32 Encodings of the Bits[10:6] of the SYNC instruction; the STYPE Field

Code	Name	Older instructions which must reach the load/store ordering point before the SYNC instruction completes.	Younger instructions which must reach the load/store ordering point only after the SYNC instruction completes.	Older instructions which must be globally performed when the SYNC instruction completes	Compliance
0x0	SYNC or SYNC 0	Loads, Stores	Loads, Stores	Loads, Stores	Required
0x4	SYNC_WMB or SYNC 4	Stores	Stores		Optional
0x10	SYNC_MB or SYNC 16	Loads, Stores	Loads, Stores		Optional
0x11	SYNC_ACQUIRE or SYNC 17	Loads	Loads, Stores		Optional
0x12	SYNC_RELEASE or SYNC 18	Loads, Stores	Stores		Optional
0x13	SYNC_RMB or SYNC 19	Loads	Loads		Optional
0x1-0x3, 0x5-0xF					Implementation-Specific and Vendor Specific Sync Types
0x14 - 0x1F	RESERVED				Reserved for MIPS Technologies for future extension of the architecture.

Terms:

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*.

Load/Store Datapath: The portion of the processor which handles the load/store data requests coming from the processor pipeline and processes those requests within the cache and memory system hierarchy.

Restrictions:

The effect of SYNC on the global order of loads and stores for memory access types other than *uncached* and *cached coherent* is **UNPREDICTABLE**.

Operation:

```
SyncOperation(stype)
```

Exceptions:

None

Programming Notes:

A processor executing load and store instructions observes the order in which loads and stores using the same memory access type occur in the instruction stream; this is known as *program order*.

A *parallel program* has multiple instruction streams that can execute simultaneously 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 store—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 neither is it 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 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 effect of all loads and stores in one group is seen by all processors before the effect of any load or store in the subsequent 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 implementation of 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 generally does not operate on a system that is not strongly ordered. However, a program that does use SYNC works on both types of systems. (System-specific documentation describes the actions needed to reliably

share data in parallel programs for that system.)

The behavior of a load or store using one memory access type is **UNPREDICTABLE** 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 appear to all processors; it does 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.

```
# 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 read
   LW  R1, DATA      # Read (valid) shared DATA value
```

The code fragments above shows 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.

Software written to use a SYNC instruction with a non-zero stype value, expecting one type of barrier behavior, should only be run on hardware that actually implements the expected barrier behavior for that non-zero stype value or on hardware which implements a superset of the behavior expected by the software for that stype value. If the hardware does not perform the barrier behavior expected by the software, the system may fail.



Format: SYNCI offset (base)

MIPS32 Release 2

Purpose: Synchronize Caches to Make Instruction Writes Effective

To synchronize all caches to make instruction writes effective.

Description:

This instruction is used after a new instruction stream is written to make the new instructions effective relative to an instruction fetch, when used in conjunction with the SYNC and JALR.HB, JR.HB, or ERET instructions, as described below. Unlike the CACHE instruction, the SYNCI instruction is available in all operating modes in an implementation of Release 2 of the architecture.

The 16-bit offset is sign-extended and added to the contents of the base register to form an effective address. The effective address is used to address the cache line in all caches which may need to be synchronized with the write of the new instructions. The operation occurs only on the cache line which may contain the effective address. One SYNCI instruction is required for every cache line that was written. See the Programming Notes below.

A TLB Refill and TLB Invalid (both with cause code equal TLBL) exception can occur as a byproduct of this instruction. This instruction never causes TLB Modified exceptions nor TLB Refill exceptions with a cause code of TLBS.

A Cache Error exception may occur as a byproduct of this instruction. For example, if a writeback operation detects a cache or bus error during the processing of the operation, that error is reported via a Cache Error exception. Similarly, a Bus Error Exception may occur if a bus operation invoked by this instruction is terminated in an error.

An Address Error Exception (with cause code equal AdEL) may occur if the effective address references a portion of the kernel address space which would normally result in such an exception. It is implementation dependent whether such an exception does occur.

It is implementation dependent whether a data watch is triggered by a SYNCI instruction whose address matches the Watch register address match conditions. In multiprocessor implementations where instruction caches are not coherently maintained by hardware, the SYNCI instruction may optionally affect all coherent icaches within the system. If the effective address uses a coherent Cacheability and Coherency Attribute (CCA), then the operation may be *globalized*, meaning it is broadcast to all of the coherent instruction caches within the system. If the effective address does not use one of the coherent CCAs, there is no broadcast of the SYNCI operation. If multiple levels of caches are to be affected by one SYNCI instruction, all of the affected cache levels must be processed in the same manner - either all affected cache levels use the globalized behavior or all affected cache levels use the non-globalized behavior.

In multiprocessor implementations where instruction caches are coherently maintained by hardware, the SYNCI instruction should behave as a NOP instruction.

Restrictions:

The operation of the processor is **UNPREDICTABLE** if the effective address references any instruction cache line that contains instructions to be executed between the SYNCI and the subsequent JALR.HB, JR.HB, or ERET instruction required to clear the instruction hazard.

The SYNCI instruction has no effect on cache lines that were previously locked with the CACHE instruction. If correct software operation depends on the state of a locked line, the CACHE instruction must be used to synchronize the caches.

The SYNCI instruction acts on the current processor at a minimum. It is implementation specific whether it affects the caches on other processors in a multi-processor system, except as required to perform the operation on the current

processor (as might be the case if multiple processors share an L2 or L3 cache).

Full visibility of the new instruction stream requires execution of a subsequent SYNC instruction, followed by a JALR.HB, JR.HB, DERET, or ERET instruction. The operation of the processor is **UNPREDICTABLE** if this sequence is not followed.

Operation:

```
vaddr ← GPR[base] + sign_extend(offset)
SynchronizeCacheLines(vaddr)      /* Operate on all caches */
```

Exceptions:

Reserved Instruction Exception (Release 1 implementations only)

TLB Refill Exception

TLB Invalid Exception

Address Error Exception

Cache Error Exception

Bus Error Exception


Programming Notes:

When the instruction stream is written, the SYNCI instruction should be used in conjunction with other instructions to make the newly-written instructions effective. The following example shows a routine which can be called after the new instruction stream is written to make those changes effective. Note that the SYNCI instruction could be replaced with the corresponding sequence of CACHE instructions (when access to Coprocessor 0 is available), and that the JR.HB instruction could be replaced with JALR.HB, ERET, or DERET instructions, as appropriate. A SYNC instruction is required between the final SYNCI instruction in the loop and the instruction that clears instruction hazards.

```
/*
 * This routine makes changes to the instruction stream effective to the
 * hardware. It should be called after the instruction stream is written.
 * On return, the new instructions are effective.
 *
 * Inputs:
 *   a0 = Start address of new instruction stream
 *   a1 = Size, in bytes, of new instruction stream
 */

    beq    a1, zero, 20f      /* If size==0, */
    nop                                /* branch around */
    addu   a1, a0, a1        /* Calculate end address + 1 */
    rdhwr  v0, HW_SYNCI_Step /* Get step size for SYNCI from new */
                                /* Release 2 instruction */

    beq    v0, zero, 20f     /* If no caches require synchronization, */
    nop                                /* branch around */
10: synci 0(a0)             /* Synchronize all caches around address */
    addu   a0, a0, v0        /* Add step size in delay slot */
    sltu   v1, a0, a1        /* Compare current with end address */
    bne    v1, zero, 10b     /* Branch if more to do */
    nop                                /* branch around */
    sync                                /* Clear memory hazards */
20: jr.hb ra                /* Return, clearing instruction hazards */
    nop
```



Format: SYSCALL

MIPS32

Purpose: System Call

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*

MIPS32

Purpose: Trap if Equal

To compare GPRs and do a conditional trap

Description: if GPR[*rs*] = GPR[*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:

Trap



Format: TEQI *rs*, *immediate*

MIPS32

Purpose: Trap if Equal Immediate

To compare a GPR to a constant and do a conditional trap

Description: if GPR[*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:

Trap



Format: TGE *rs*, *rt*

MIPS32

Purpose: Trap if Greater or Equal

To compare GPRs and do a conditional trap

Description: if $GPR[rs] \geq GPR[rt]$ 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:

Trap



Format: TGEI rs, immediate

MIPS32

Purpose: Trap if Greater or Equal Immediate

To compare a GPR to a constant and do a conditional trap

Description: if $GPR[rs] \geq immediate$ 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:

Trap



Format: TGEIU rs, immediate

MIPS32

Purpose: Trap if Greater or Equal Immediate Unsigned

To compare a GPR to a constant and do a conditional trap

Description: if $GPR[rs] \geq immediate$ 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 can 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:

Trap



Format: TGEU *rs*, *rt*

MIPS32

Purpose: Trap if Greater or Equal Unsigned

To compare GPRs and do a conditional trap

Description: if $GPR[rs] \geq GPR[rt]$ 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:

Trap

$$TLB[i]_{C0} \ || \ TLB[i]_{D0} \ || \ TLB[i]_{V0} \ || \ TLB[i]_G$$
Exceptions:

Coprocesor Unusable

Machine Check

Exceptions:

Coprocesor Unusable

Machine Check

31	26	25	24	6	5	0
COP0 010000	CO 1	0 000 0000 0000 0000 0000			TLBWR 000110	
6	1	19			6	

Format: TLBWR

MIPS32

Purpose: Write Random TLB Entry

To write a TLB entry indexed by the *Random* register.

Description:

The TLB entry pointed to by the *Random* register is written from the contents of the *EntryHi*, *EntryLo0*, *EntryLo1*, and *PageMask* registers. It is implementation dependent whether multiple TLB matches are detected on a TLBWR. In such an instance, a Machine Check Exception is signaled. In Release 2 of the Architecture, multiple TLB matches may only be reported on a TLB write. The information written to the TLB entry may be different from that in the *EntryHi*, *EntryLo0*, and *EntryLo1* registers, in that:

- The value written to the VPN2 field of the TLB entry may have those bits set to zero corresponding to the one bits in the Mask field of the *PageMask* register (the least significant bit of VPN2 corresponds to the least significant bit of the Mask field). It is implementation dependent whether these bits are preserved or zeroed during a TLB write.
- The value written to the PFN0 and PFN1 fields of the TLB entry may have those bits set to zero corresponding to the one bits in the Mask field of *PageMask* register (the least significant bit of PFN corresponds to the least significant bit of the Mask field). It is implementation dependent whether these bits are preserved or zeroed during a TLB write.
- The single G bit in the TLB entry is set from the logical AND of the G bits in the *EntryLo0* and *EntryLo1* registers.

Restrictions:

If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

Operation:

```


i ← Random
TLB[i]Mask ← PageMaskMask
TLB[i]VPN2 ← EntryHiVPN2 and not PageMaskMask # Implementation dependent
TLB[i]ASID ← EntryHiASID
TLB[i]G ← EntryLo1G and EntryLo0G
TLB[i]PFN1 ← EntryLo1PFN and not PageMaskMask # Implementation dependent
TLB[i]C1 ← EntryLo1C
TLB[i]D1 ← EntryLo1D
TLB[i]V1 ← EntryLo1V
TLB[i]PFN0 ← EntryLo0PFN and not PageMaskMask # Implementation dependent
TLB[i]C0 ← EntryLo0C
TLB[i]D0 ← EntryLo0D
TLB[i]V0 ← EntryLo0V

```

Exceptions:

Coprocessor Unusable

Machine Check



Format: TLT *rs*, *rt*

MIPS32

Purpose: Trap if Less Than

To compare GPRs and do a conditional trap

Description: if GPR[*rs*] < GPR[*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:

Trap



Format: TLTI rs, immediate

MIPS32

Purpose: Trap if Less Than Immediate

To compare a GPR to a constant and do a conditional trap

Description: if GPR[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:

Trap



Format: TLTIU rs, immediate

MIPS32

Purpose: Trap if Less Than Immediate Unsigned

To compare a GPR to a constant and do a conditional trap

Description: if GPR[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 can 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:

Trap



Format: TLTU *rs*, *rt*

MIPS32

Purpose: Trap if Less Than Unsigned

To compare GPRs and do a conditional trap

Description: if GPR[*rs*] < GPR[*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:

Trap



Format: TNE *rs*, *rt*

MIPS32

Purpose: Trap if Not Equal

To compare GPRs and do a conditional trap

Description: if GPR[*rs*] \neq GPR[*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]  $\neq$  GPR[rt] then
    SignalException(Trap)
endif
```

Exceptions:

Trap



Format: TNEI rs, immediate

MIPS32

Purpose: Trap if Not Equal Immediate

To compare a GPR to a constant and do a conditional trap

Description: if GPR[rs] \neq 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]  $\neq$  sign_extend(immediate) then
    SignalException(Trap)
endif
```

Exceptions:

Trap



Format: TRUNC.W.fmt
 TRUNC.W.S fd, fs **MIPS32**
 TRUNC.W.D fd, fs **MIPS32**

Purpose: Floating Point Truncate to Word Fixed Point

To convert an FP value to 32-bit fixed point, rounding toward zero

Description: $FPR[fd] \leftarrow \text{convert_and_round}(FPR[fs])$

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. In this case 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*.

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 **UNPREDICTABLE**.

The operand must be a value in format *fmt*; if it is not, the result is **UNPREDICTABLE** and the value of the operand FPR becomes **UNPREDICTABLE**.

Operation:


`StoreFPR(fd, W, ConvertFmt(ValueFPR(fs, fmt), fmt, W))`

Exceptions:

Coprocessor Unusable, Reserved Instruction

Floating Point Exceptions:

Inexact, Invalid Operation, Overflow, Unimplemented Operation



Format: WAIT

MIPS32

Purpose: Enter Standby Mode

Wait for Event

Description:

The WAIT instruction performs an implementation-dependent operation, usually involving a lower power mode. Software may use bits 24:6 of the instruction to communicate additional information to the processor, and the processor may use this information as control for the lower power mode. A value of zero for bits 24:6 is the default and must be valid in all implementations.

The WAIT instruction is typically implemented by stalling the pipeline at the completion of the instruction and entering a lower power mode. The pipeline is restarted when an external event, such as an interrupt or external request occurs, and execution continues with the instruction following the WAIT instruction. It is implementation-dependent whether the pipeline restarts when a non-enabled interrupt is requested. In this case, software must poll for the cause of the restart. The assertion of any reset or NMI must restart the pipeline and the corresponding exception must be taken.

If the pipeline restarts as the result of an enabled interrupt, that interrupt is taken between the WAIT instruction and the following instruction (EPC for the interrupt points at the instruction following the WAIT instruction).

Restrictions:

The operation of the processor is **UNDEFINED** if a WAIT instruction is placed in the delay slot of a branch or a jump.


If access to Coprocessor 0 is not enabled, a Coprocessor Unusable Exception is signaled.

Operation:

```
I: Enter implementation dependent lower power mode
I+1: /* Potential interrupt taken here */
```

Exceptions:

Coprocessor Unusable Exception



Format: WRPGPR rd, rt

MIPS32 Release 2

Purpose: Write to GPR in Previous Shadow Set

To move the contents of a current GPR to a GPR in the previous shadow set.

Description: $SGPR[SRSCtl_{pSS}, rd] \leftarrow GPR[rt]$

The contents of the current GPR *rt* is moved to the shadow GPR register specified by *SRSCtl_{pSS}* (signifying the previous shadow set number) and *rd* (specifying the register number within that set).

Restrictions:

In implementations prior to Release 2 of the Architecture, this instruction resulted in a Reserved Instruction Exception.


Operation:

$SGPR[SRSCtl_{pSS}, rd] \leftarrow GPR[rt]$

Exceptions:

Coprocessor Unusable

Reserved Instruction



Format: WSBH rd, rt

MIPS32 Release 2

Purpose: Word Swap Bytes Within Halfwords

To swap the bytes within each halfword of GPR *rt* and store the value into GPR *rd*.

Description: $GPR[rd] \leftarrow \text{SwapBytesWithinHalfwords}(GPR[rt])$

Within each halfword of GPR *rt* the bytes are swapped, and stored in GPR *rd*.

Restrictions:

In implementations prior to Release 2 of the architecture, this instruction resulted in a Reserved Instruction Exception.

Operation:

$$GPR[rd] \leftarrow GPR[r]_{23..16} || GPR[r]_{31..24} || GPR[r]_{7..0} || GPR[r]_{15..8}$$

Exceptions:

Reserved Instruction

Programming Notes:


The WSBH instruction can be used to convert halfword and word data of one endianness to another endianness. The endianness of a word value can be converted using the following sequence:

```
lw    t0, 0(a1)           /* Read word value */
wsbh  t0, t0              /* Convert endiannes of the halfwords */
rotr  t0, t0, 16         /* Swap the halfwords within the words */
```

Combined with SEH and SRA, two contiguous halfwords can be loaded from memory, have their endianness converted, and be sign-extended into two word values in four instructions. For example:

```
lw    t0, 0(a1)           /* Read two contiguous halfwords */
wsbh  t0, t0              /* Convert endiannes of the halfwords */
seh   t1, t0              /* t1 = lower halfword sign-extended to word */
sra   t0, t0, 16         /* t0 = upper halfword sign-extended to word */
```

Zero-extended words can be created by changing the SEH and SRA instructions to ANDI and SRL instructions, respectively.



Format: XOR *rd*, *rs*, *rt*

MIPS32

Purpose: Exclusive OR

To do a bitwise logical Exclusive OR

Description: $GPR[rd] \leftarrow GPR[rs] \text{ XOR } GPR[rt]$

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*

MIPS32

Purpose: Exclusive OR Immediate

To do a bitwise logical Exclusive OR with a constant

Description: $GPR[rt] \leftarrow GPR[rs] \text{ XOR } immediate$

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:

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

Exceptions:

None

Instruction Bit Encodings

A.1 Instruction Encodings and Instruction Classes

Instruction encodings are presented in this section; field names are printed here and throughout the book in *italics*.

When encoding an instruction, the primary *opcode* field is encoded first. Most *opcode* values completely specify an instruction that has an *immediate* value or offset.

Opcode values that do not specify an instruction instead specify an instruction class. Instructions within a class are further specified by values in other fields. For instance, *opcode* REGIMM specifies the *immediate* instruction class, which includes conditional branch and trap *immediate* instructions.


A.2 Instruction Bit Encoding Tables

This section provides various bit encoding tables for the instructions of the MIPS32® ISA.

Figure A.1 shows a sample encoding table and the instruction *opcode* field this table encodes. Bits 31..29 of the *opcode* field are listed in the leftmost columns of the table. Bits 28..26 of the *opcode* field are listed along the topmost rows of the table. Both decimal and binary values are given, with the first three bits designating the row, and the last three bits designating the column.

An instruction's encoding is found at the intersection of a row (bits 31..29) and column (bits 28..26) value. For instance, the *opcode* value for the instruction labelled EX1 is 33 (decimal, row and column), or 011011 (binary). Similarly, the *opcode* value for EX2 is 64 (decimal), or 110100 (binary).

Figure A.1 Sample Bit Encoding Table



Tables A.2 through A.20 describe the encoding used for the MIPS32 ISA. Table A.1 describes the meaning of the symbols used in the tables.

Table A.1 Symbols Used in the Instruction Encoding Tables

Symbol	Meaning
*	Operation or field codes marked with this symbol are reserved for future use. Executing such an instruction must cause a Reserved Instruction Exception.
δ	(Also <i>italic</i> field name.) Operation or field codes marked with this symbol denotes a field class. The instruction word must be further decoded by examining additional tables that show values for another instruction field.
β	Operation or field codes marked with this symbol represent a valid encoding for a higher-order MIPS ISA level or a new revision of the Architecture. Executing such an instruction must cause a Reserved Instruction Exception.
∇	Operation or field codes marked with this symbol represent instructions which were only legal if 64-bit operations were enabled on implementations of Release 1 of the Architecture. In Release 2 of the architecture, operation or field codes marked with this symbol represent instructions which are legal if 64-bit floating point operations are enabled. In other cases, executing such an instruction must cause a Reserved Instruction Exception (non-coprocessor encodings or coprocessor instruction encodings for a coprocessor to which access is allowed) or a Coprocessor Unusable Exception (coprocessor instruction encodings for a coprocessor to which access is not allowed).

Table A.1 Symbols Used in the Instruction Encoding Tables (Continued)

Symbol	Meaning
θ	Operation or field codes marked with this symbol are available to licensed MIPS partners. To avoid multiple conflicting instruction definitions, MIPS Technologies will assist the partner in selecting appropriate encodings if requested by the partner. The partner is not required to consult with MIPS Technologies when one of these encodings is used. If no instruction is encoded with this value, executing such an instruction must cause a Reserved Instruction Exception (<i>SPECIAL2</i> encodings or coprocessor instruction encodings for a coprocessor to which access is allowed) or a Coprocessor Unusable Exception (coprocessor instruction encodings for a coprocessor to which access is not allowed).
σ	Field codes marked with this symbol represent an EJTAG support instruction and implementation of this encoding is optional for each implementation. If the encoding is not implemented, executing such an instruction must cause a Reserved Instruction Exception. If the encoding is implemented, it must match the instruction encoding as shown in the table.
ε	Operation or field codes marked with this symbol are reserved for MIPS Application Specific Extensions. If the ASE is not implemented, executing such an instruction must cause a Reserved Instruction Exception.
ϕ	Operation or field codes marked with this symbol are obsolete and will be removed from a future revision of the MIPS32 ISA. Software should avoid using these operation or field codes.
\oplus	Operation or field codes marked with this symbol are valid for Release 2 implementations of the architecture. Executing such an instruction in a Release 1 implementation must cause a Reserved Instruction Exception.

Table A.2 MIPS32 Encoding of the Opcode Field

opcode		bits 28..26							
		0	1	2	3	4	5	6	7
bits 31..29		000	001	010	011	100	101	110	111
0	000	<i>SPECIAL</i> δ	<i>REGIMM</i> δ	J	JAL	BEQ	BNE	BLEZ	BGTZ
1	001	ADDI	ADDIU	SLTI	SLTIU	ANDI	ORI	XORI	LUI
2	010	<i>COP0</i> δ	<i>COP1</i> δ	<i>COP2</i> $\theta\delta$	<i>COP1X</i> ¹ δ	BEQL ϕ	BNEL ϕ	BLEZL ϕ	BGTZL ϕ
3	011	β	β	β	β	<i>SPECIAL2</i> δ	JALX ε	ε	<i>SPECIAL3</i> ² $\delta\oplus$
4	100	LB	LH	LWL	LW	LBU	LHU	LWR	β
5	101	SB	SH	SWL	SW	β	β	SWR	CACHE
6	110	LL	LWC1	LWC2 θ	PREF	β	LDC1	LDC2 θ	β
7	111	SC	SWC1	SWC2 θ	*	β	SDC1	SDC2 θ	β

1. In Release 1 of the Architecture, the *COP1X* opcode was called *COP3*, and was available as another user-available coprocessor. In Release 2 of the Architecture, a full 64-bit floating point unit is available with 32-bit CPUs, and the *COP1X* opcode is reserved for that purpose on all Release 2 CPUs. 32-bit implementations of Release 1 of the architecture are strongly discouraged from using this opcode for a user-available coprocessor as doing so will limit the potential for an upgrade path to a 64-bit floating point unit.
2. Release 2 of the Architecture added the *SPECIAL3* opcode. Implementations of Release 1 of the Architecture signaled a Reserved Instruction Exception for this opcode.

Table A.3 MIPS32 SPECIAL Opcode Encoding of Function Field

function		bits 2..0							
		0	1	2	3	4	5	6	7
bits 5..3		000	001	010	011	100	101	110	111
0	000	SLL ¹	MOVCI δ	SRL δ	SRA	SLLV	*	SRLV δ	SRAV
1	001	JR ²	JALR ²	MOVZ	MOVN	SYSCALL	BREAK	*	SYNC
2	010	MFHI	MTHI	MFLO	MTLO	β	*	β	β
3	011	MULT	MULTU	DIV	DIVU	β	β	β	β
4	100	ADD	ADDU	SUB	SUBU	AND	OR	XOR	NOR
5	101	*	*	SLT	SLTU	β	β	β	β
6	110	TGE	TGEU	TLT	TLTU	TEQ	*	TNE	*
7	111	β	*	β	β	β	*	β	β

1. Specific encodings of the *rt*, *rd*, and *sa* fields are used to distinguish among the SLL, NOP, SSNOP, EHB and PAUSE functions.
2. Specific encodings of the *hint* field are used to distinguish JR from JR.HB and JALR from JALR.HB

Table A.4 MIPS32 REGIMM Encoding of *rt* Field

rt		bits 18..16							
		0	1	2	3	4	5	6	7
bits 20..19		000	001	010	011	100	101	110	111
0	00	BLTZ	BGEZ	BLTZL ϕ	BGEZL ϕ	*	*	*	*
1	01	TGEI	TGEIU	TLTI	TLTIU	TEQI	*	TNEI	*
2	10	BLTZAL	BGEZAL	BLTZALL ϕ	BGEZALL ϕ	*	*	*	*
3	11	*	*	*	*	*	*	*	SYNCP \oplus

Table A.5 MIPS32 SPECIAL2 Encoding of Function Field

function		bits 2..0							
		0	1	2	3	4	5	6	7
bits 5..3		000	001	010	011	100	101	110	111
0	000	MADD	MADDU	MUL	θ	MSUB	MSUBU	θ	θ
1	001	θ	θ	θ	θ	θ	θ	θ	θ
2	010	θ	θ	θ	θ	θ	θ	θ	θ
3	011	θ	θ	θ	θ	θ	θ	θ	θ
4	100	CLZ	CLO	θ	θ	β	β	θ	θ
5	101	θ	θ	θ	θ	θ	θ	θ	θ
6	110	θ	θ	θ	θ	θ	θ	θ	θ
7	111	θ	θ	θ	θ	θ	θ	θ	SDBBP σ

Table A.6 MIPS32 SPECIAL3¹ Encoding of Function Field for Release 2 of the Architecture

function		bits 2..0							
		0	1	2	3	4	5	6	7
bits 5..3		000	001	010	011	100	101	110	111
0	000	EXT \oplus	β	β	β	INS \oplus	β	β	β
1	001	*	*	*	*	*	*	*	*
2	010	*	*	*	*	*	*	*	*
3	011	*	*	*	*	*	*	*	*
4	100	BSHFL $\oplus \delta$	*	*	*	β	*	*	*
5	101	*	*	*	*	*	*	*	*
6	110	*	*	*	*	*	*	*	*
7	111	*	*	*	RDHWR \oplus	*	*	*	*

1. Release 2 of the Architecture added the SPECIAL3 opcode. Implementations of Release 1 of the Architecture signaled a Reserved Instruction Exception for this opcode and all function field values shown above.

Table A.7 MIPS32 MOVCI Encoding of *tf* Bit

tf	bit 16	
	0	1
	MOVF	MOV \mathcal{T}

Table A.8 MIPS32¹ SRL Encoding of Shift/Rotate

R	bit 21	
	0	1
	SRL	ROTR

1. Release 2 of the Architecture added the ROTR instruction. Implementations of Release 1 of the Architecture ignored bit 21 and treated the instruction as an SRL.

Table A.9 MIPS32¹ SRLV Encoding of Shift/Rotate

R	bit 6	
	0	1
	SRLV	ROTRV

1. Release 2 of the Architecture added the ROTRV instruction. Implementations of Release 1 of the Architecture ignored bit 6 and treated the instruction as an SRLV.

Table A.10 MIPS32 BSHFL Encoding of sa Field¹

sa		bits 8..6							
		0	1	2	3	4	5	6	7
bits 10..9		000	001	010	011	100	101	110	111
0	00			WSBH					
1	01								
2	10	SEB							
3	11	SEH							

1. The sa field is sparsely decoded to identify the final instructions. Entries in this table with no mnemonic are reserved for future use by MIPS Technologies and may or may not cause a Reserved Instruction exception.

Table A.11 MIPS32 COP0 Encoding of rs Field

rs		bits 23..21							
		0	1	2	3	4	5	6	7
bits 25..24		000	001	010	011	100	101	110	111
0	00	MFC0	β	*	*	MTC0	β	*	*
1	01	*	*	RDPGPR \oplus	MFMC0 ¹ $\delta\oplus$	*	*	WRPGPR \oplus	*
2	10	C0 δ							
3	11								

1. Release 2 of the Architecture added the MFMC0 function, which is further decoded as the DI (bit 5 = 0) and EI (bit 5 = 1) instructions.

Table A.12 MIPS32 COP0 Encoding of Function Field When rs=CO

function		bits 2..0							
		0	1	2	3	4	5	6	7
bits 5..3		000	001	010	011	100	101	110	111
0	000	*	TLBR	TLBWI	*	*	*	TLBWR	*
1	001	TLBP	*	*	*	*	*	*	*
2	010	*	*	*	*	*	*	*	*
3	011	ERET	*	*	*	*	*	*	DERET σ
4	100	WAIT	*	*	*	*	*	*	*
5	101	*	*	*	*	*	*	*	*
6	110	*	*	*	*	*	*	*	*
7	111	*	*	*	*	*	*	*	*

Table A.13 MIPS32 COP1 Encoding of *rs* Field

rs		bits 23..21							
		0	1	2	3	4	5	6	7
bits 25..24		000	001	010	011	100	101	110	111
0	00	MFC1	β	CFC1	MFHC1 \oplus	MTC1	β	CTC1	MTHC1 \oplus
1	01	$BC1 \delta$	$BC1ANY2 \delta \epsilon \nabla$	$BC1ANY4 \delta \epsilon \nabla$	*	*	*	*	*
2	10	S δ	D δ	*	*	W δ	L δ	PS δ	*
3	11	*	*	*	*	*	*	*	*

Table A.14 MIPS32 COP1 Encoding of Function Field When *rs*=S

function		bits 2..0							
		0	1	2	3	4	5	6	7
bits 5..3		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	*	*	*	*	RECIP2 $\epsilon \nabla$	RECIP1 $\epsilon \nabla$	RSQRT1 $\epsilon \nabla$	RSQRT2 $\epsilon \nabla$
4	100	*	CVT.D	*	*	CVT.W	CVT.L ∇	CVT.PS ∇	*
5	101	*	*	*	*	*	*	*	*
6	110	C.F CABS.F $\epsilon \nabla$	C.UN CABS.UN $\epsilon \nabla$	C.EQ CABS.EQ $\epsilon \nabla$	C.UEQ CABS.UEQ $\epsilon \nabla$	C.OLT CABS.OLT $\epsilon \nabla$	C.ULT CABS.ULT $\epsilon \nabla$	C.OLE CABS.OLE $\epsilon \nabla$	C.ULE CABS.ULE $\epsilon \nabla$
7	111	C.SF CABS.SF $\epsilon \nabla$	C.NGLE CABS.NGLE $\epsilon \nabla$	C.SEQ CABS.SEQ $\epsilon \nabla$	C.NGL CABS.NGL $\epsilon \nabla$	C.LT CABS.LT $\epsilon \nabla$	C.NGE CABS.NGE $\epsilon \nabla$	C.LE CABS.LE $\epsilon \nabla$	C.NGT CABS.NGT $\epsilon \nabla$

Table A.15 MIPS32 COP1 Encoding of Function Field When *rs*=D

function		bits 2..0							
		0	1	2	3	4	5	6	7
bits 5..3		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	*	*	*	*	RECIP2 $\epsilon \nabla$	RECIP1 $\epsilon \nabla$	RSQRT1 $\epsilon \nabla$	RSQRT2 $\epsilon \nabla$
4	100	CVT.S	*	*	*	CVT.W	CVT.L ∇	*	*
5	101	*	*	*	*	*	*	*	*
6	110	C.F CABS.F $\epsilon \nabla$	C.UN CABS.UN $\epsilon \nabla$	C.EQ CABS.EQ $\epsilon \nabla$	C.UEQ CABS.UEQ $\epsilon \nabla$	C.OLT CABS.OLT $\epsilon \nabla$	C.ULT CABS.ULT $\epsilon \nabla$	C.OLE CABS.OLE $\epsilon \nabla$	C.ULE CABS.ULE $\epsilon \nabla$
7	111	C.SF CABS.SF $\epsilon \nabla$	C.NGLE CABS.NGLE $\epsilon \nabla$	C.SEQ CABS.SEQ $\epsilon \nabla$	C.NGL CABS.NGL $\epsilon \nabla$	C.LT CABS.LT $\epsilon \nabla$	C.NGE CABS.NGE $\epsilon \nabla$	C.LE CABS.LE $\epsilon \nabla$	C.NGT CABS.NGT $\epsilon \nabla$

Table A.16 MIPS32 COP1 Encoding of Function Field When $rs=W$ or L ¹

function		bits 2..0							
		0	1	2	3	4	5	6	7
bits 5..3		000	001	010	011	100	101	110	111
0	000	*	*	*	*	*	*	*	*
1	001	*	*	*	*	*	*	*	*
2	010	*	*	*	*	*	*	*	*
3	011	*	*	*	*	*	*	*	*
4	100	CVT.S	CVT.D	*	*	*	*	CVT.PS.PW ϵ ∇	*
5	101	*	*	*	*	*	*	*	*
6	110	*	*	*	*	*	*	*	*
7	111	*	*	*	*	*	*	*	*

1. Format type L is legal only if 64-bit floating point operations are enabled.

Table A.17 MIPS64 COP1 Encoding of Function Field When $rs=PS$ ¹

function		bits 2..0							
		0	1	2	3	4	5	6	7
bits 5..3		000	001	010	011	100	101	110	111
0	000	ADD ∇	SUB ∇	MUL ∇	*	*	ABS ∇	MOV ∇	NEG ∇
1	001	*	*	*	*	*	*	*	*
2	010	*	MOVCF δ ∇	MOVZ ∇	MOVN ∇	*	*	*	*
3	011	ADDR ϵ ∇	*	MULR ϵ ∇	*	RECIP2 ϵ ∇	RECIP1 ϵ ∇	RSQRT1 ϵ ∇	RSQRT2 ϵ ∇
4	100	CVT.S.PU ∇	*	*	*	CVT.PW.PS ϵ ∇	*	*	*
5	101	CVT.S.PL ∇	*	*	*	PLL.PS ∇	PLU.PS ∇	PUL.PS ∇	PUU.PS ∇
6	110	C.F ∇ CABS.F ϵ ∇	C.UN ∇ CABS.UN ϵ ∇	C.EQ ∇ CABS.EQ ϵ ∇	C.UEQ ∇ CABS.UEQ ϵ ∇	C.OLT ∇ CABS.OLT ϵ ∇	C.ULT ∇ CABS.ULT ϵ ∇	C.OLE ∇ CABS.OLE ϵ ∇	C.ULE ∇ CABS.ULE ϵ ∇
7	111	C.SF ∇ CABS.SF ϵ ∇	C.NGLE ∇ CABS.NGLE ϵ ∇	C.SEQ ∇ CABS.SEQ ϵ ∇	C.NGL ∇ CABS.NGL ϵ ∇	C.LT ∇ CABS.LT ϵ ∇	C.NGE ∇ CABS.NGE ϵ ∇	C.LE ∇ CABS.LE ϵ ∇	C.NGT ∇ CABS.NGT ϵ ∇

1. Format type PS is legal only if 64-bit floating point operations are enabled.

Table A.18 MIPS32 COP1 Encoding of tf Bit When $rs=S, D,$ or $PS,$ Function= $MOVCF$

tf	bit 16	
	0	1
	MOV.F.fmt	MOV.T.fmt

Table A.19 MIPS32 COP2 Encoding of *rs* Field

rs		<i>bits 23..21</i>							
		0	1	2	3	4	5	6	7
<i>bits 25..24</i>		000	001	010	011	100	101	110	111
0	00	MFC2 θ	β	CFC2 θ	MFHC2 $\theta\oplus$	MTC2 θ	β	CTC2 θ	MTHC2 $\theta\oplus$
1	01	BC2 θ	*	*	*	*	*	*	*
2	10	C2 $\theta\delta$							
3	11								

Table A.20 MIPS64 COP1X Encoding of Function Field¹

function		<i>bits 2..0</i>							
		0	1	2	3	4	5	6	7
<i>bits 5..3</i>		000	001	010	011	100	101	110	111
0	000	LWXC1 ∇	LDXC1 ∇	*	*	*	LUXC1 ∇	*	*
1	001	SWXC1 ∇	SDXC1 ∇	*	*	*	SUXC1 ∇	*	PREFX ∇
2	010	*	*	*	*	*	*	*	*
3	011	*	*	*	*	*	*	ALNV.PS ∇	*
4	100	MADD.S ∇	MADD.D ∇	*	*	*	*	MADD.PS ∇	*
5	101	MSUB.S ∇	MSUB.D ∇	*	*	*	*	MSUB.PS ∇	*
6	110	NMADD.S ∇	NMADD.D ∇	*	*	*	*	NMADD.PS ∇	*
7	111	NMSUB.S ∇	NMSUB.D ∇	*	*	*	*	NMSUB.PS ∇	*

1. COP1X instructions are legal only if 64-bit floating point operations are enabled.

A.3 Floating Point Unit Instruction Format Encodings

Instruction format encodings for the floating point unit are presented in this section. This information is a tabular presentation of the encodings described in tables Table A.13 and Table A.20 above.

Table A.21 Floating Point Unit Instruction Format Encodings

<i>fmt</i> field (bits 25..21 of COP1 opcode)		<i>fmt3</i> field (bits 2..0 of COP1X opcode)		Mnemonic	Name	Bit Width	Data Type
Decimal	Hex	Decimal	Hex				
0..15	00..0F	—	—	Used to encode Coprocessor 1 interface instructions (MFC1, CTC1, etc.). Not used for format encoding.			
16	10	0	0	S	Single	32	Floating Point
17	11	1	1	D	Double	64	Floating Point
18..19	12..13	2..3	2..3	Reserved for future use by the architecture.			
20	14	4	4	W	Word	32	Fixed Point

Table A.21 Floating Point Unit Instruction Format Encodings

<i>fmt</i> field (bits 25..21 of COP1 opcode)		<i>fmt3</i> field (bits 2..0 of COP1X opcode)		Mnemonic	Name	Bit Width	Data Type
Decimal	Hex	Decimal	Hex				
21	15	5	5	L	Long	64	Fixed Point
22	16	6	6	PS	Paired Sin- gle	2 × 32	Floating Point
23	17	7	7	Reserved for future use by the architecture.			
24..31	18..1F	—	—	Reserved for future use by the architecture. Not available for <i>fmt3</i> encoding.			

Revision History

In the left hand page margins of this document you may find vertical change bars to note the location of significant changes to this document since its last release. Significant changes are defined as those which you should take note of as you use the MIPS IP. Changes to correct grammar, spelling errors or similar may or may not be noted with change bars. Change bars will be removed for changes which are more than one revision old.

Please note: Limitations on the authoring tools make it difficult to place change bars on changes to figures. Change bars on figure titles are used to denote a potential change in the figure itself.

Revision	Date	Description
0.90	November 1, 2000	Internal review copy of reorganized and updated architecture documentation.
0.91	November 15, 2000	Internal review copy of reorganized and updated architecture documentation.
0.92	December 15, 2000	Changes in this revision: <ul style="list-style-type: none"> • Correct sign in description of MSUBU. • Update JR and JALR instructions to reflect the changes required by MIPS16.
0.95	March 12, 2001	Update for second external review release
1.00	August 29, 2002	Update based on all review feedback: <ul style="list-style-type: none"> • Add missing optional select field syntax in mtc0/mfc0 instruction descriptions. • Correct the PREF instruction description to acknowledge that the Prepare-ForStore function does, in fact, modify architectural state. • To provide additional flexibility for Coprocessor 2 implementations, extend the <i>sel</i> field for DMFC0, DMTC0, MFC0, and MTC0 to be 8 bits. • Update the PREF instruction to note that it may not update the state of a locked cache line. • Remove obviously incorrect documentation in DIV and DIVU with regard to putting smaller numbers in register <i>rt</i>. • Fix the description for MFC2 to reflect data movement from the coprocessor 2 register to the GPR, rather than the other way around. • Correct the pseudo code for LDC1, LDC2, SDC1, and SDC2 for a MIPS32 implementation to show the required word swapping. • Indicate that the operation of the CACHE instruction is UNPREDICTABLE if the cache line containing the instruction is the target of an invalidate or writeback invalidate. • Indicate that an Index Load Tag or Index Store Tag operation of the CACHE instruction must not cause a cache error exception. • Make the entire right half of the MFC2, MTC2, CFC2, CTC2, DMFC2, and DMTC2 instructions implementation dependent, thereby acknowledging that these fields can be used in any way by a Coprocessor 2 implementation. • Clean up the definitions of LL, SC, LLD, and SCD. • Add a warning that software should not use non-zero values of the <i>stpe</i> field of the SYNC instruction. • Update the compatibility and subsetting rules to capture the current requirements.

Revision History

Revision	Date	Description
1.90	September 1, 2002	<p>Merge the MIPS Architecture Release 2 changes in for the first release of a Release 2 processor. Changes in this revision include:</p> <ul style="list-style-type: none">• All new Release 2 instructions have been included: DI, EHB, EI, EXT, INS, JALR.HB, JR.HB, MFHC1, MFHC2, MTHC1, MTHC2, RDHWR, RDPGPR, ROTR, ROTRV, SEB, SEH, SYNCI, WRPGPR, WSBH.• The following instruction definitions changed to reflect Release 2 of the Architecture: DERET, ERET, JAL, JALR, JR, SRL, SRLV• With support for 64-bit FPUs on 32-bit CPUs in Release 2, all floating point instructions that were previously implemented by MIPS64 processors have been modified to reflect support on either MIPS32 or MIPS64 processors in Release 2.• All pseudo-code functions have been updated, and the Are64bitFPOperationsEnabled function was added.• Update the instruction encoding tables for Release 2.
2.00	June 9, 2003	<p>Continue with updates to merge Release 2 changes into the document. Changes in this revision include:</p> <ul style="list-style-type: none">• Correct the target GPR (from rd to rt) in the SLTI and SLTIU instructions. This appears to be a day-one bug.• Correct CPR number, and missing data movement in the pseudocode for the MTC0 instruction.• Add note to indicate that the CACHE instruction does not take Address Error Exceptions due to mis-aligned effective addresses.• Update SRL, ROTR, SRLV, ROTRV, DSRL, DROTR, DSRLV, DROTRV, DSRL32, and DROTR32 instructions to reflect a 1-bit, rather than a 4-bit decode of shift vs. rotate function.• Add programming note to the PrepareForStore PREF hint to indicate that it can not be used alone to create a bzero-like operation.• Add note to the PREF and PREFX instruction indicating that they may cause Bus Error and Cache Error exceptions, although this is typically limited to systems with high-reliability requirements.• Update the SYNCI instruction to indicate that it should not modify the state of a locked cache line.• Establish specific rules for when multiple TLB matches can be reported (on writes only). This makes software handling easier.
2.50	July 1, 2005	<p>Changes in this revision:</p> <ul style="list-style-type: none">• Correct figure label in LWR instruction (it was incorrectly specified as LWL).• Update all files to FrameMaker 7.1.• Include support for implementation-dependent hardware registers via RDHWR.• Indicate that it is implementation-dependent whether prefetch instructions cause EJTAG data breakpoint exceptions on an address match, and suggest that the preferred implementation is not to cause an exception.• Correct the MIPS32 pseudocode for the LDC1, LDXC1, LUXC1, SDC1, SDXC1, and SUXC1 instructions to reflect the Release 2 ability to have a 64-bit FPU on a 32-bit CPU. The correction simplifies the code by using the ValueFPR and StoreFPR functions, which correctly implement the Release 2 access to the FPRs.• Add an explicit recommendation that all cache operations that require an index be done by converting the index to a kseg0 address before performing the cache operation.• Expand on restrictions on the PREF instruction in cases where the effective address has an uncached coherency attribute.•

Revision	Date	Description
2.60	June 25, 2008	Changes in this revision: <ul style="list-style-type: none"> • Applied the new B0.01 template. • Update RDHWR description with the UserLocal register. • added PAUSE instruction • Ordering SYNCs • CMP behavior of CACHE, PREF*, SYNCI • CVT.S.PL, CVT.S.PU are non-arithmetic (no exceptions) • *MADD.fmt & *MSUB.fmt are non-fused. • various typos fixed
2.61	July 10, 2008	<ul style="list-style-type: none"> • Revision History file was incorrectly copied from Volume III. • Removed index conditional text from PAUSE instruction description. • SYNC instruction - added additional format "SYNC stype"
2.62	January 2, 2009	<ul style="list-style-type: none"> • LWC1, LWXC1 - added statement that upper word in 64bit registers are UNDEFINED. • CVT.S.PL and CVT.S.PU descriptions were still incorrectly listing IEEE exceptions. • Typo in CFC1 Description. • CCRes is accessed through \$3 for RDHWR, not \$4.
3.00	March 25, 2010	<ul style="list-style-type: none"> • JALX instruction description added. • Sub-setting rules updated for JALX. •

Revision History