



# LatticeMico32 Processor Reference Manual

Lattice Semiconductor Corporation  
5555 NE Moore Court  
Hillsboro, OR 97124  
(503) 268-8000

July 2009

---

---

## Copyright

Copyright © 2008 Lattice Semiconductor Corporation.

This document may not, in whole or part, be copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine-readable form without prior written consent from Lattice Semiconductor Corporation.

## Trademarks

Lattice Semiconductor Corporation, L Lattice Semiconductor Corporation (logo), L (stylized), L (design), Lattice (design), LSC, E<sup>2</sup>CMOS, Extreme Performance, FlashBAK, flexiFlash, flexiMAC, flexiPCS, FreedomChip, GAL, GDX, Generic Array Logic, HDL Explorer, IPexpress, ISP, ispATE, ispClock, ispDOWNLOAD, ispGAL, ispGDS, ispGDX, ispGDXV, ispGDX2, ispGENERATOR, ispJTAG, ispLEVER, ispLeverCORE, ispLSI, ispMACH, ispPAC, ispTRACY, ispTURBO, ispVIRTUAL MACHINE, ispVM, ispXP, ispXPGA, ispXPLD, LatticeEC, LatticeECP, LatticeECP-DSP, LatticeECP2, LatticeECP2M, LatticeMico8, LatticeMico32, LatticeSC, LatticeSCM, LatticeXP, LatticeXP2, MACH, MachXO, MACO, ORCA, PAC, PAC-Designer, PAL, Performance Analyst, PURESPEED, Reveal, Silicon Forest, Speedlocked, Speed Locking, SuperBIG, SuperCOOL, SuperFAST, SuperWIDE, sysCLOCK, sysCONFIG, sysDSP, sysHSI, sysI/O, sysMEM, The Simple Machine for Complex Design, TransFR, UltraMOS, and specific product designations are either registered trademarks or trademarks of Lattice Semiconductor Corporation or its subsidiaries in the United States and/or other countries. ISP, Bringing the Best Together, and More of the Best are service marks of Lattice Semiconductor Corporation.

HyperTransport is a licensed trademark of the HyperTransport Technology Consortium in the U.S. and other jurisdictions.

Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

## Disclaimers

NO WARRANTIES: THE INFORMATION PROVIDED IN THIS DOCUMENT IS "AS IS" WITHOUT ANY EXPRESS OR IMPLIED WARRANTY OF ANY KIND INCLUDING WARRANTIES OF ACCURACY, COMPLETENESS, MERCHANTABILITY, NONINFRINGEMENT OF INTELLECTUAL PROPERTY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL LATTICE SEMICONDUCTOR CORPORATION (LSC) OR ITS SUPPLIERS BE LIABLE FOR ANY DAMAGES WHATSOEVER (WHETHER DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL, INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION, OR LOSS OF INFORMATION) ARISING OUT OF THE USE OF OR INABILITY TO USE THE INFORMATION PROVIDED IN THIS DOCUMENT, EVEN IF LSC HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME JURISDICTIONS PROHIBIT THE EXCLUSION OR LIMITATION OF CERTAIN LIABILITY, SOME OF THE ABOVE LIMITATIONS MAY NOT APPLY TO YOU.

LSC may make changes to these materials, specifications, or information, or to the products described herein, at any time without notice. LSC makes no commitment to update this documentation. LSC reserves the right to discontinue any product or service without notice and assumes no obligation

---

to correct any errors contained herein or to advise any user of this document of any correction if such be made. LSC recommends its customers obtain the latest version of the relevant information to establish, before ordering, that the information being relied upon is current.

## Type Conventions Used in This Document

Convention	Meaning or Use
<b>Bold</b>	Items in the user interface that you select or click. Text that you type into the user interface.
<i>&lt;Italic&gt;</i>	Variables in commands, code syntax, and path names.
<b>Ctrl+L</b>	Press the two keys at the same time.
<i>Courier</i>	Code examples. Messages, reports, and prompts from the software.
...	Omitted material in a line of code.
.	Omitted lines in code and report examples.
[ ]	Optional items in syntax descriptions. In bus specifications, the brackets are required.
( )	Grouped items in syntax descriptions.
{ }	Repeatable items in syntax descriptions.
	A choice between items in syntax descriptions.

---

# Contents

<b>Chapter 1</b>	<b>LatticeMico32 Processor and Systems</b>	<b>1</b>
<b>Chapter 2</b>	<b>Programmer's Model</b>	<b>5</b>
	Pipeline Architecture	5
	Data Types	6
	Register Architecture	7
	General-Purpose Registers	7
	Control and Status Registers	9
	Memory Architecture	14
	Address Space	14
	Endianness	14
	Address Alignment	15
	Stack Layout	15
	Caches	16
	Inline Memories	18
	Exceptions	20
	Exception Processing	21
	Exception Handlers	22
	Nested Exceptions	26
	Remapping the Exception Table	26
	Reset Summary	26
	Using Breakpoints	27
	Using Watchpoints	27
	Debug Architecture	28
	DC – Debug Control	29
	DEBA – Debug Exception Base Address	29
	JTX – JTAG UART Transmit Register	29
	JRX – JTAG UART Receive Register	30
	BPn – Breakpoint	30
	WPn – Watchpoint	31
	Instruction Set Categories	31

	Arithmetic	31
	Logic	31
	Comparison	32
	Shift	32
	Data Transfer	32
	Program Flow Control	32
<b>Chapter 3</b>	<b>Configuring the LatticeMico32 Processor</b>	<b>35</b>
	Configuration Options	35
	EBR Use	38
<b>Chapter 4</b>	<b>WISHBONE Interconnect Architecture</b>	<b>39</b>
	Introduction to WISHBONE Interconnect	39
	WISHBONE Registered Feedback Mode	40
	CTI_IO()	40
	BTE_IO()	41
	Component Signals	42
	Master Port and Signal Descriptions	43
	Slave Port and Signal Descriptions	44
	Arbitration Schemes	46
	Shared-Bus Arbitration	46
	Slave-Side Arbitration	46
<b>Chapter 5</b>	<b>Instruction Set</b>	<b>49</b>
	Instruction Formats	49
	Opcode Look-Up Table	50
	Pseudo-Instructions	51
	Instruction Descriptions	51
	add	52
	addi	52
	and	53
	andhi	53
	andi	54
	b	54
	be	55
	bg	55
	bge	56
	bgeu	56
	bgu	57
	bi	57
	bne	58
	break	58
	bret	59
	call	59
	calli	60
	cmpe	60
	cmpei	61
	cmpg	61
	cmpgi	62
	cmpge	62
	cmpgei	63

cmpgeu	63
cmpgeui	64
cmpgu	64
cmpgui	65
cmpne	65
cmpnei	66
divu	66
eret	67
lb	67
lbu	68
lh	68
lhu	69
lw	69
modu	70
mul	71
muli	71
mv	72
mvhi	72
nor	73
nori	73
not	74
or	74
ori	75
orhi	75
rcsr	76
ret	76
sb	76
scall	77
sextb	77
sexth	78
sh	79
sl	79
sli	80
sr	80
sri	81
sru	81
srui	82
sub	82
sw	83
wcsr	83
xnor	84
xnori	84
xor	85
xori	85

<b>Index</b>	<b>87</b>
--------------	-----------



# LatticeMico32 Processor and Systems

As systems become more complex, there are a growing number of L\_2 and L\_3 protocols that continue to burden a local host processor. These tend to incrementally add processing requirements to the local processor, starving other critical functions of processor machine cycles. To alleviate the local host processor's processing requirements, embedded processors are being utilized to support the main processor in a distributed processing architecture. These embedded processors offer localized control, OA&M functionality, and statistics gathering and processing features, thereby saving the host processor many unnecessary clock cycles, which can be used for higher-level functions.

A soft processor provides added flexibility in the implementation of your design. Functionality that can be implemented in software rather than hardware allows much greater freedom in terms of the types of changes that can be made. With software-based processing, it is possible for the hardware logic to remain stable and functional upgrades can be made through software modification. Additionally, it is much quicker and simpler to implement functionality in software than it is to design it in hardware, leading to a reduced time to market.

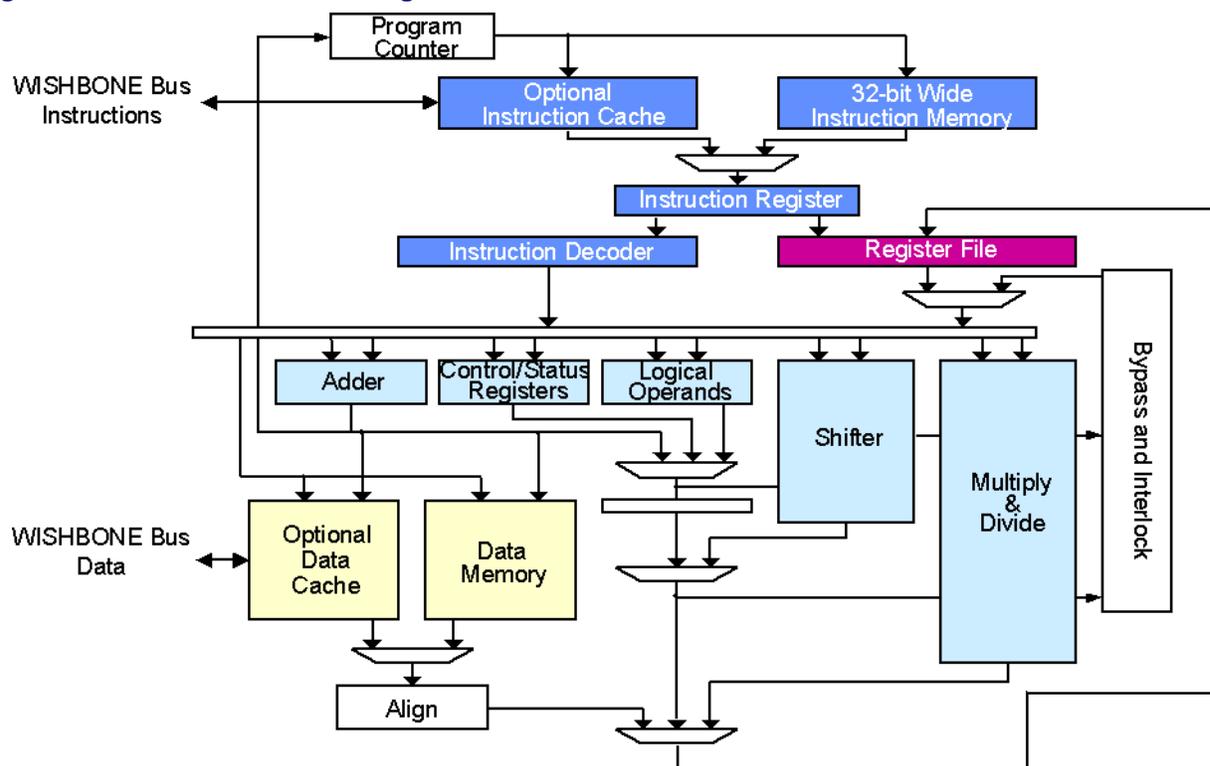
The LatticeMico32™ is a configurable 32-bit soft processor core for Lattice Field Programmable Gate Array (FPGA) devices. By combining a 32-bit wide instruction set with 32 general-purpose registers, the LatticeMico32 provides the performance and flexibility suitable for a wide variety of markets, including communications, consumer, computer, medical, industrial, and automotive. With separate instruction and data buses, this Harvard architecture processor allows for single-cycle instruction execution as the instruction and data memories can be accessed simultaneously. Additionally, the LatticeMico32 uses a Reduced Instruction Set Computer (RISC) architecture, thereby providing a simpler instruction set and faster performance. As a result, the processor core consumes minimal device resources, while maintaining the

performance required for a broad application set. Some of the key features of this 32-bit processor include:

- ◆ RISC architecture
- ◆ 32-bit data path
- ◆ 32-bit instructions
- ◆ 32 general-purpose registers
- ◆ Up to 32 external interrupts
- ◆ Optional instruction cache
- ◆ Optional data cache
- ◆ Dual WISHBONE memory interfaces (instruction and data)

Figure 1 shows a block diagram of the LatticeMico32 processor core.

**Figure 1: LatticeMico32 Block Diagram**



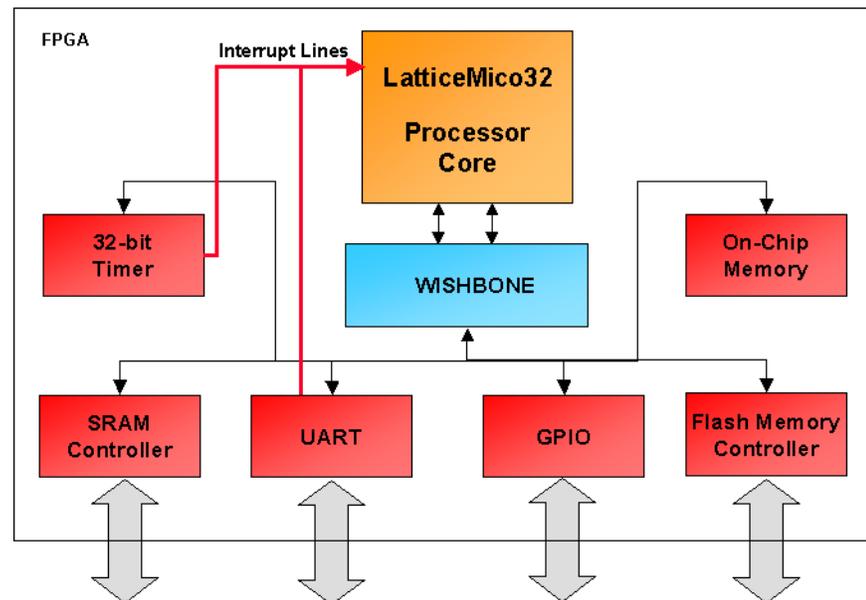
To accelerate the development of processor systems, several optional peripheral components are available with the LatticeMico32 processor. Specifically, these components are connected to the processor through a WISHBONE bus interface, a royalty-free, public-domain specification. By using this open source bus interface, you can incorporate your own

WISHBONE components into your embedded designs. The components include:

- ◆ Memory controllers
  - ◆ Asynchronous SRAM
  - ◆ Double data rate (DDR)
  - ◆ On-chip
- ◆ Input/output (I/O) ports
  - ◆ 32-bit timer
  - ◆ Direct memory access (DMA) controller
  - ◆ General-purpose I/O (GPIO)
  - ◆ I<sup>2</sup>C master controller
  - ◆ Serial peripheral interface (SPI)
  - ◆ Universal asynchronous receiver transmitter (UART)

Figure 2 shows a complete embedded system using the LatticeMico32 processor along with several components.

**Figure 2: LatticeMico32 Processor Embedded System**



This manual describes the architecture of the LatticeMico32 processor. It includes information on configuration options, pipeline architecture, register architecture, memory architecture, debug architecture, and the instruction set. It is intended to be used as a reference when you design processors for use on supported Lattice Semiconductor field programmable gate arrays (FPGAs).



## Programmer's Model

This chapter describes the pipeline architecture of the LatticeMico32 processor.

---

### Pipeline Architecture

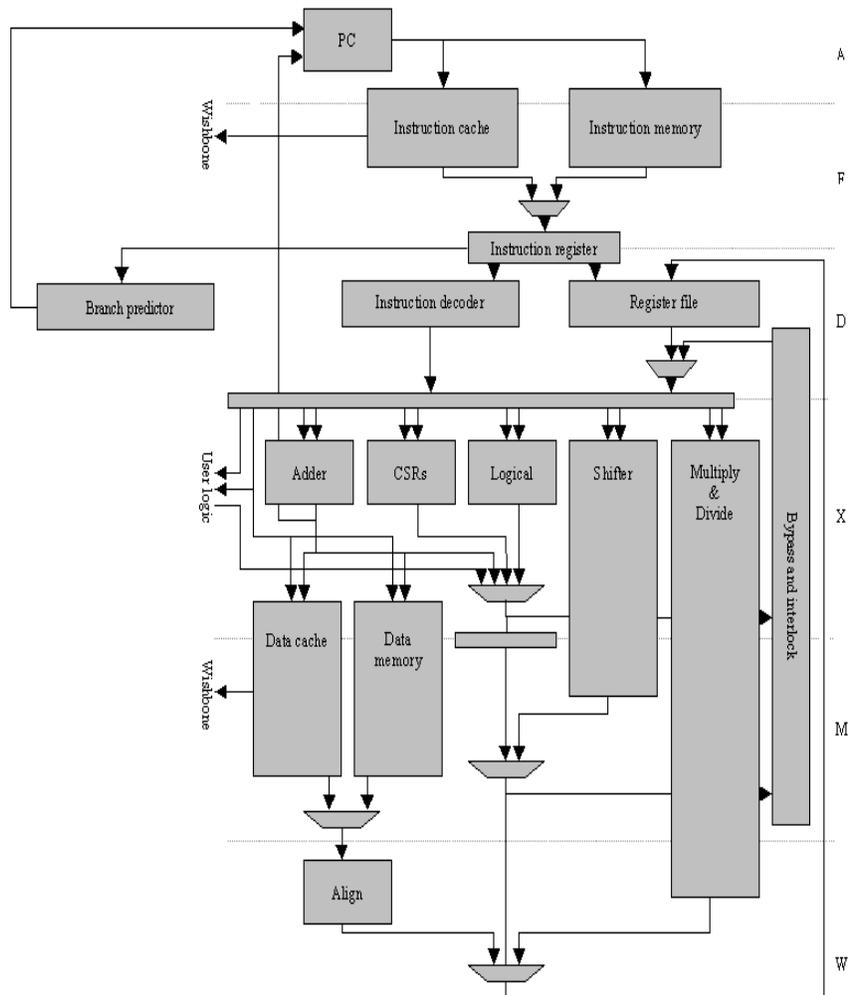
---

The LatticeMico32 processor uses a 32-bit, 6-stage pipeline, as shown in Figure 3 on page 6. It is fully bypassed and interlocked. The bypass logic is responsible for forwarding results back through the pipeline, allowing most instructions to be effectively executed in a single cycle. The interlock is responsible for detecting read-after-write hazards and stalling the pipeline until the hazard has been resolved. This avoids the need to insert nop directives between dependent instructions, keeping code size to a minimum, as well as simplifying assembler-level programming.

The six pipeline stages are:

- ◆ Address – The address of the instruction to execute is calculated and sent to the instruction cache.
- ◆ Fetch – The instruction is read from memory.
- ◆ Decode – The instruction is decoded, and operands are either fetched from the register file or bypassed from the pipeline. PC-relative branches are predicted by a static branch predictor.
- ◆ Execute – The operation specified by the instruction is performed. For simple instructions such as addition or a logical operation, execution finishes in this stage, and the result is made available for bypassing.
- ◆ Memory – For more complicated instructions such as loads, stores, multiplies, or shifts, a second execution stage is required.
- ◆ Writeback – Results produced by the instructions are written back to the register file

**Figure 3: LatticeMico32 Pipeline**



## Data Types

The LatticeMico32 processor supports the data types listed in Table 1.

**Table 1: Data Types**

Type	Range	Bits	Encoding	C Compiler Type
Unsigned byte	$[0, 2^8-1]$	8	Binary	Unsigned character
Signed byte	$[-2^7, 2^7-1]$	8	Two's complement	Character
Unsigned half-word	$[0, 2^{16}-1]$	16	Binary	Unsigned short
Signed half-word	$[-2^{15}, 2^{15}-1]$	16	Two's complement	Short

**Table 1: Data Types**

Type	Range	Bits	Encoding	C Compiler Type
Unsigned word	$[0, 2^{32}-1]$	32	Binary	Unsigned int/ unsigned long
Signed word	$[-2^{31}, 2^{31}-1]$	32	Two's complement	Int/long

In addition to the above, the extended data types in Table 2 can be emulated through a compiler.

**Table 2: Extended Data Types**

Data Type	Range	Bits	Encoding	C Compiler Type
Unsigned double-word	$[0, 2^{64}-1]$	64	Binary	Unsigned long long
Signed double-word	$[-2^{63}, 2^{63}-1]$	64	Two's complement	Long long
Single-precision real	$[1.1754e-38, 3.4028e+38]$	32	IEEE 754	Float
Double-precision real	$[2.2250e-308, 1.7976e+308]$	64	IEEE 754	Double

## Register Architecture

This section describes the register architecture of the LatticeMico32 processor.

### General-Purpose Registers

The LatticeMico32 features the following 32-bit registers:

- ◆ By convention, register 0 (r0) must always hold the value 0, and this is required for correct operation by both the LatticeMico32 assembler and the C compiler. On power-up, the value of 0 in r0 is not hardwired, so you must initialize it to load r0 with the 0 value.
- ◆ Registers 1 through 28 are truly general purpose and can be used as the source or destination register for any instruction. After reset, the values in all of these registers are undefined.
- ◆ Register 29 (ra) is used by the call instruction to save the return address but is otherwise general purpose.
- ◆ Register 30 (ea) is used to save the value of the Program Counter (PC) when an exception occurs, so it should not be used by user-level programs.
- ◆ Register 31 (ba) saves the value of the Program Counter (PC) when a breakpoint or watchpoint exception occurs, so it should not be used by user-level programs.

After reset, the values in all of the above 32-bit registers are undefined. To ensure that register 0 contains 0, the first instruction executed after reset should be `xor r0, r0, r0`.

Table 3 lists the general-purpose registers and specifies their use by the C compiler. In this table, the callee is the function called by the caller function.

**Table 3: General-Purpose Registers**

Register Name	Function	Saver
r0	Holds the value zero	
r1	General-purpose/argument 0/return value 0	Caller
r2	General-purpose/argument 1/return value 1	Caller
r3	General-purpose/argument 2	Caller
r4	General-purpose/argument 3	Caller
r5	General-purpose/argument 4	Caller
r6	General-purpose/argument 5	Caller
r7	General-purpose/argument 6	Caller
r8	General-purpose/argument 7	Caller
r9	General-purpose	Caller
r10	General-purpose	Caller
r11	General-purpose	Callee
r12	General-purpose	Callee
r13	General-purpose	Callee
r14	General-purpose	Callee
r15	General-purpose	Callee
r16	General-purpose	Callee
r17	General-purpose	Callee
r18	General-purpose	Callee
r19	General-purpose	Callee
r20	General-purpose	Callee
r21	General-purpose	Callee
r22	General-purpose	Callee
r23	General-purpose	Callee
r24	General-purpose	Callee
r25	General-purpose	Callee
r26/gp	General-purpose/global pointer	Callee

**Table 3: General-Purpose Registers (Continued)**

Register Name	Function	Saver
r27/fp	General-purpose/frame pointer	Callee
r28/sp	Stack pointer	Callee
r29/ra	General-purpose/return address	Caller
r30/ea	Exception address	
r31/ba	Breakpoint address	

## Control and Status Registers

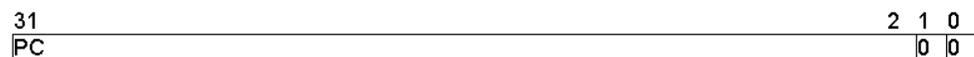
Table 4 shows all of the names of the control and status registers (CSR), whether the register can be read from or written to, and the index used when accessing the register. Some of the registers are optional, depending on the configuration of the processor (see “Configuring the LatticeMico32 Processor” on page 35). All signal levels are active high.

**Table 4: Control and Status Registers**

Name	Access	Index	Optional	Description
PC			No	Program counter
IE	R/W	0x0	Yes	Interrupt enable
EID	R	—	No	LM32 revision number. EID is not accessible via Assembly, since it does not have an index.
IM	R/W	0x1	Yes	Interrupt mask
IP	R	0x2	Yes	Interrupt pending
ICC	W	0x3	Yes	Instruction cache control
DCC	W	0x4	Yes	Data cache control
CC	R	0x5	Yes	Cycle counter
CFG	R	0x6	No	Configuration
EBA	R/W	0x7	No	Exception base address

### PC – Program Counter

The PC CSR is a 32-bit register that contains the address of the instruction currently being executed. Because all instructions are four bytes wide, the two least significant bits of the PC are always zero. After reset, the value of the PC CSR is h00000000.

**Figure 4: Format of the PC CSR**

## EID – Exception ID

The EID CSR is a 3-bit value that indicates what has caused the program execution to stop. The values and corresponding events are the following:

- 0 – Reset
- 1 – Breakpoint Exception
- 2 – Instruction Bus Error
- 3 – Watchpoint Exception
- 4 – Data Bust Error
- 5 – Divide-by-zero
- 6 – Interrupt
- 7 – System Call

## IE – Interrupt Enable

The IE CSR contains a single-bit flag, IE, that determines whether interrupts are enabled. This flag has priority over the IM CSR. In addition, there are two bits, BIE and EIE, that are used to save the value of the IE field when either a breakpoint or other exception occurs. Each interrupt is associated with a mask bit (IE bit) indexed with each interrupt. After reset, the value of the IE CSR is h00000000.

**Figure 5: Format of the IE CSR**



**Table 5: Fields of the IE CSR**

Field	Values	Description
IE	0 – Interrupts disabled 1 – Interrupts enabled	Determines whether interrupts are enabled.
EIE	0 – Interrupts disabled 1 – Interrupts enabled	Holds a copy of the IE field when an exception occurs.
BIE	0 – Interrupts disabled 1 – Interrupts enabled	Holds a copy of the IE field when a breakpoint occurs.

### IM – Interrupt Mask

The IM CSR contains an enable bit for each of the 32 interrupts. Bit 0 corresponds to interrupt 0. In order for an interrupt to be raised, both an enable bit in this register and the IE flag in the IE CSR must be set to 1. After reset, the value of the IM CSR is h00000000.

### IP – Interrupt Pending

The IP CSR contains a pending bit for each of the 32 interrupts. A pending bit is set when the corresponding interrupt request line is asserted low. Bit 0 corresponds to interrupt 0. Bits in the IP CSR can be cleared by writing a 1 with the wcsr instruction. Writing a 0 has no effect. After reset, the value of the IP CSR is h00000000.

### ICC – Instruction Cache Control

The ICC CSR provides a control bit that, when written with any value, causes the contents of the entire instruction cache to be invalidated.

**Figure 6: Format of the ICC CSR**



Field	Values	Description
I	Any – Invalidate instruction cache	When written, the contents of the instruction cache are invalidated.

### DCC – Data Cache Control

The DCC CSR provides a control bit that, when written with any value, causes the contents of the entire data cache to be invalidated.

**Figure 7: Format of the DCC CSR**



**Table 6: Fields of the DCC CSR**

Field	Values	Description
I	Any – Invalidate data cache	When written, the contents of the data cache are invalidated.

### CC – Cycle Counter

The CC CSR is an optional 32-bit register that is incremented on each clock cycle. It can be used to profile ghost code sequences.

### CFG – Configuration

The CFG CSR details the configuration of a particular instance of a LatticeMico32 processor.

**Figure 8: Format of the CFG CSR**



**Table 7: Fields of the CFG CSR**

Field	Values	Description
M	0 – Multiply is not implemented 1 – Multiply is implemented	Indicates whether a hardware multiplier is implemented.
D	0 – Divide is not implemented 1 – Divide is implemented	Indicates whether a hardware divider is implemented.
S	0 – Barrel shift is not implemented 1 – Barrel shift is implemented	Indicates whether a hardware barrel-shifter is implemented.
U		Reserved.
X	0 – Sign extend is not implemented 1 – Sign extend is implemented	Indicates whether the sign-extension instructions are implemented.
CC	0 – Cycle counter is not implemented 1 – Cycle counter is implemented	Indicates whether the CC CSR is implemented.
IC	0 – Instruction cache is not implemented 1 – Instruction cache is implemented	Indicates whether an instruction cache is implemented.



# Memory Architecture

This section describes the memory architecture of the LatticeMico32 processor.

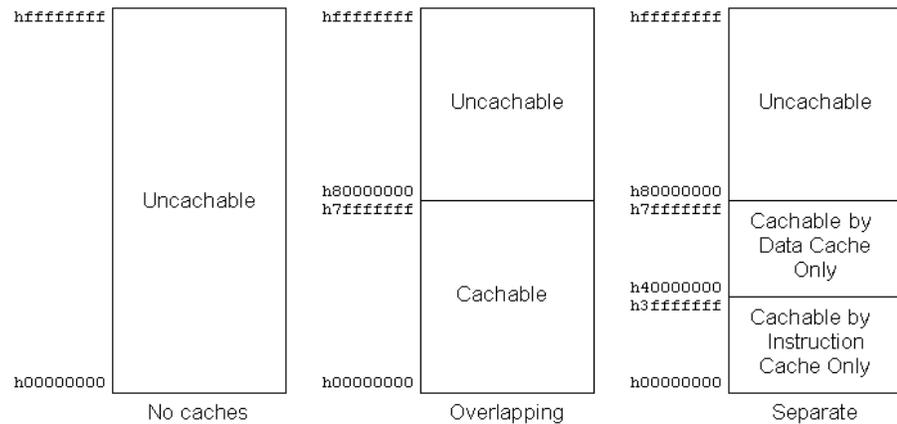
## Address Space

The LatticeMico32 processor has a flat 32-bit, byte-addressable address space. For LatticeMico32 processors with caches, the portion of the address space that is cacheable can be configured separately for both the instruction and data cache. This allows for the size of the cache tag RAMs to be optimized to be as small as is required (the fewer the number of cacheable addresses, the smaller the tag RAMs will be).

If an instruction cache is used, attempts to fetch instructions from outside of the range of cacheable addresses result in undefined behavior, so only one cached region is supported. Portions of the memory image are not cached, so if a miss occurs, it will not be fetched.

Figure 10 illustrates some possible configurations. Typically, the parts of the address space that are cacheable are used for storing code or program data, with I/O components being mapped into uncacheable addresses.

**Figure 10: Cacheable Addresses**



## Endianness

The LatticeMico32 processor is big-endian, which means that multi-byte objects, such as half-words and words, are stored with the most significant byte at the lowest address.

## Address Alignment

All memory accesses must be aligned to the size of the access, as shown in Table 8. No check is performed for unaligned access. All unaligned accesses result in undefined behavior.

**Table 8: Memory Access Alignment Requirements**

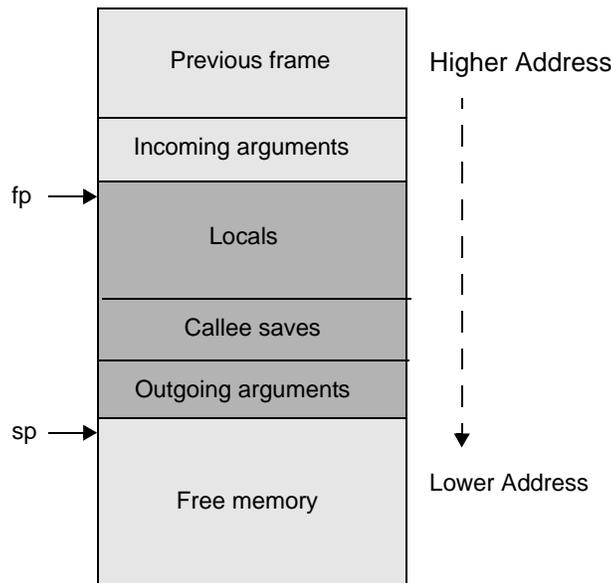
Access Size	Address Requirements
Byte	None
Half-word	Address must be half-word aligned (bit 0 must be 0)
Word	Address must be word aligned (bits 1 and 0 must be 0)

## Stack Layout

Figure 11 shows the conventional layout of a stack frame. The stack grows toward lower memory as data is pushed onto it. The stack pointer (sp) points to the first unused location, and the frame pointer (fp) points at the first location used in the active frame. In many cases, a compiler may be able to eliminate the frame pointer, because data can often be accessed by using a negative displacement from the stack pointer, freeing up the frame pointer for use as a general-purpose register.

As illustrated in Table 3 on page 8, the first eight function arguments are passed in registers. Any remaining arguments are passed on the stack, as illustrated in Figure 11.

**Figure 11: Stack Layout**



## Caches

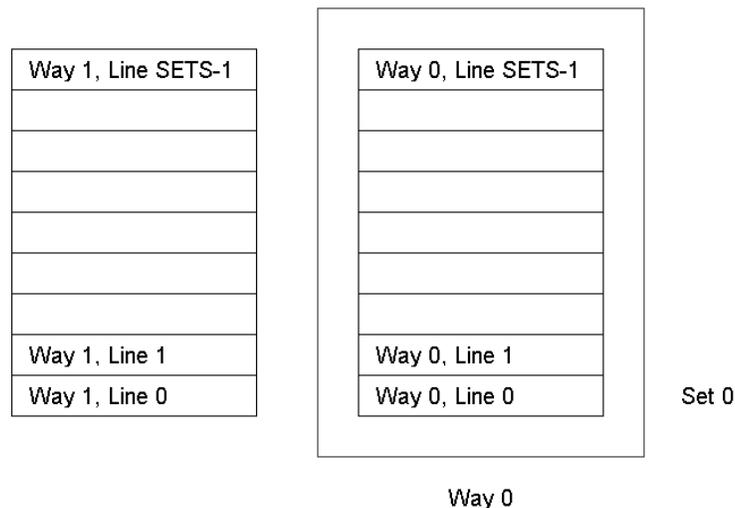
A cache is a fast memory (single-cycle access) that stores a copy of a limited subset of the data held in main memory, which may take the CPU several cycles to access. A cache helps improve overall performance by exploiting the fact that the same data is typically accessed several times in a short interval. By storing a local copy of the data in the processor's cache, the multiple cycles required to access the data can be reduced to just a single cycle for all subsequent accesses once the data is loaded into the cache.

### Cache Architecture

When a cache accesses a data item, it is also likely to access data at adjacent addresses (such as with arrays or structures) by loading data into the cache in lines. A line can consist of 4, 8, or 16 adjacent bytes, and is specified by the `BYTES_PER_LINE` option.

A one-way associative (direct-mapped) cache consists of an array of cache lines known as a "way." To allow the cache to operate at a high frequency, data from main memory can only be stored in a specific cache line. A two-way associative cache consists of a two-dimensional array of cache lines. It requires slightly more logic to implement but allows data from main memory to be stored in one of two places in the cache. It helps performance by reducing cache conflicts that occur when a program is accessing multiple data items that would map to the same cache line in a one-way associative cache. The number of lines in each way is specified by the `ICACHE_SETS` and `DCACHE_SETS` options. The ways are assigned in a round-robin fashion. Each time a cache miss occurs the way number is switched.

**Figure 12: Cache Organization**



The LatticeMico32 caches are write-through, which means that whenever a store instruction writes to an address that is cached, the data is written to both the cache and main memory. A read-miss allocation policy means that a cache line is only fetched from memory for a load instruction. If a cache miss

occurs for a store instruction, the data is written directly to memory without the cache being updated.

The LatticeMico32 processor supports a range of cache configurations, as detailed in Table 9.

**Table 9: Cache Configurations**

Attribute	Values
Size	0 kB, 1 kB, 2 kB, 4 kB, 8 kB, 16 kB, 32 kB
Sets	128, 256, 512, 1024
Associativity	1, 2
Bytes-per-line	4, 8, 16
Write policy	Write-through
Update policy	Read miss only

The LatticeMico32 caches are initialized automatically by embedded logic, so they do not require a program to initialize or enable them.

### Invalidating the Caches

The contents of the instruction cache can be invalidated by writing to the ICC CSR. It is recommended that you follow the write to the ICC CSR with four nops, as follows:

```
wcsr ICC, r0
nop
nop
nop
nop
```

The contents of the data cache can similarly be invalidated by writing to the DCC CSR as follows:

```
wcsr DCC, r0
```

It is recommended that you avoid placing a load or store instruction immediately before or after the wcsr instruction.

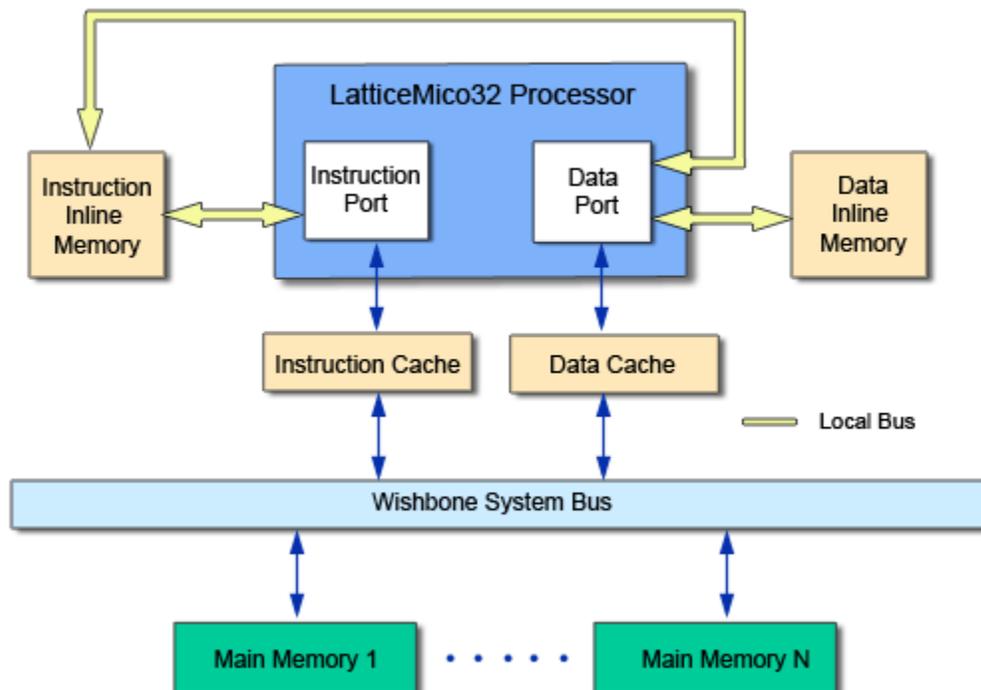
The LatticeMico32 caches are not kept consistent with respect to each other. This means that if a store instruction writes to an area of memory that is currently cached by the instruction cache, the instruction cache will not be automatically updated to reflect the store. It is your responsibility to invalidate the instruction cache after the write has taken place, if necessary.

Similarly, the caches do not snoop bus activity to monitor for writes by peripherals (by DMA for example) to addresses that are cached. It is again your responsibility to ensure that the cache is invalidated before reading memory that may have been written by a peripheral.

## Inline Memories

The LatticeMico32 processor enables you to optionally connect to on-chip memory, through instruction and data ports, by using a local bus rather than the Wishbone interface. Memory connected to the CPU in such a manner is referred to as inline memory. Figure 13 shows a functional block diagram of the LatticeMico32 processor with inline memories.

**Figure 13: LatticeMico32 Inline Memories**



There are two types of inline memories:

- ◆ **Instruction Inline Memory** – This memory component is connected to the Instruction Port of the LatticeMico32 CPU and is used to hold only program memory of any software application.
- ◆ **Data Inline Memory** – This memory component is connected to the Data Port of the LatticeMico32 CPU and is used to hold read-only or read/write data of any software application.

### Note

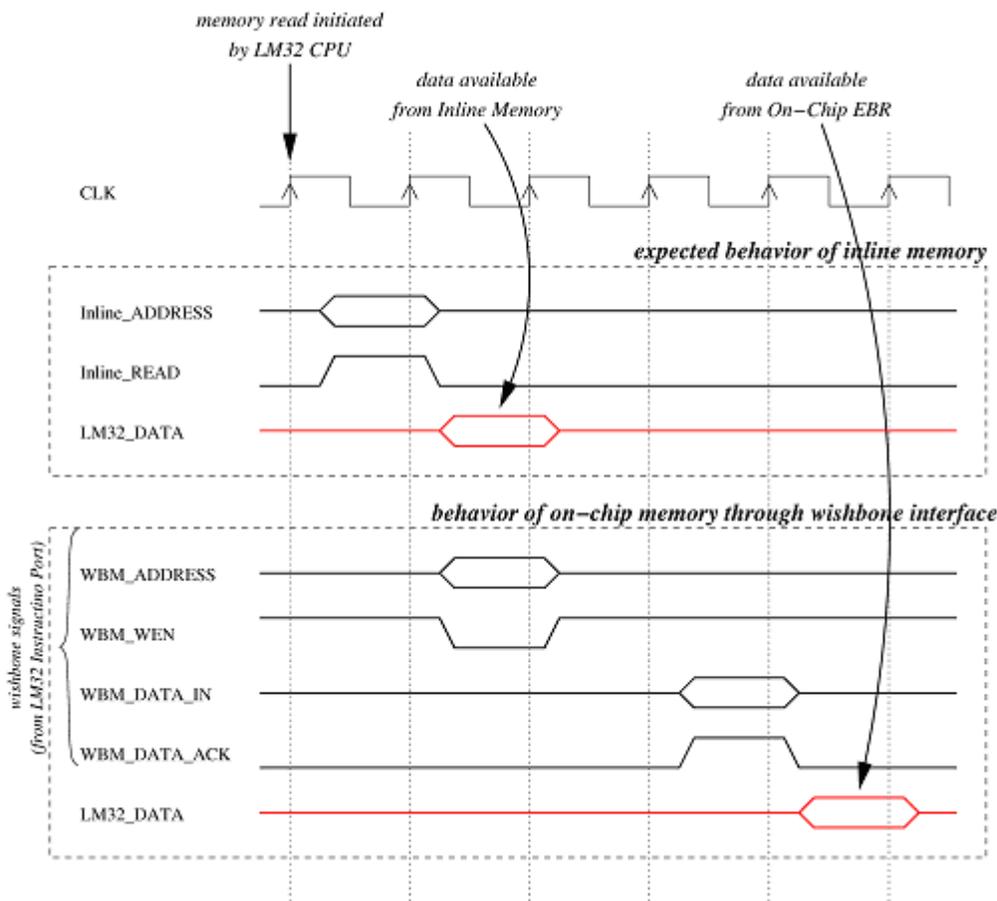
The Instruction Inline Memory is also connected to the Data Port of the LatticeMico32 CPU in order to facilitate loading of the memory image of the software application through the command line `lm32-elf-gdb` or through the C/C++ SPE Debugger.

While it is possible to create a LatticeMico32 platform that contains inline memories as the sole memory components, inline memories can co-exist in a platform with other Wishbone-based memory components. Inline memories act as types of main memories, but with the difference that the contents of these memories are not cached.

### Performance Advantage Over Wishbone-based Memory without Caches

The direct connection between CPU and EBR-based inline memory has the advantage of providing a single-cycle read/write access to the CPU. Figure 14 shows cycle-level analysis of potential performance benefits of inline memory when compared to on-chip memory (EBR) that is connected to the CPU through the Wishbone interface.

Figure 14: Cycle-level Analysis



This diagram compares the number of cycles it takes to service read access from the LatticeMico32 CPU by the inline memory versus the Wishbone-based on-chip EBR. The read access initiated to inline memory will be completed in the next cycle, whereas a read access initiated to EBR will take four cycles. A similar behavior can be seen for writes initiated by the

LatticeMico32 CPU. This shows that deploying program code or data to inline memory can provide at least a 3x speedup over Wishbone-based memories.

### Performance Advantage Over Wishbone-based Memory with Caches

It is common to configure the LatticeMico32 CPU with Instruction and Data caches to reduce the performance impact of accessing Wishbone-based memories, since they theoretically provide a single-cycle access. In practice, however, you will encounter situations in which a single-cycle cache access is not possible. In these situations, inline memory affords a performance advantage. Such situations include the following scenarios:

- ◆ Any cache access (read or write) that results in a miss will initiate an access to memory components on the Wishbone Interface. As a result, the cache access will not take multiple cycles to complete.
- ◆ The data cache in LatticeMico32 is write-through, meaning that any write to the data cache from LatticeMico32 will immediately result in access to memory components on the Wishbone interface. This means that all data cache writes are multicycle accesses.

---

## Exceptions

---

Exceptions are events either inside or outside of the processor that cause a change in the normal flow of program execution. The LatticeMico32 processor can raise eight types of exceptions, as shown in Table 10. One clock must be held in active high until an acknowledge (ACK\_I) is accepted. The remaining are internal and no ACK\_I is generated. The exceptions are listed in a decreasing order of priority, so if multiple exceptions occur simultaneously, the exception with the highest priority is raised.

**Table 10: Exceptions**

Exception	ID	Condition
Reset	0	Raised when the processor's reset pin is asserted.
Breakpoint	1	Raised when either a break instruction is executed or when a hardware breakpoint is triggered.
InstructionBusError	2	Raised when an instruction fetch fails, typically due to the requested address being invalid.
Watchpoint	3	Raised when a hardware watchpoint is triggered.
DataBusError	4	Raised when a data access fails, typically because either the requested address is invalid or the type of access is not allowed.
DivideByZero	5	Raised when an attempt is made to divide by zero.

**Table 10: Exceptions**

Exception	ID	Condition
Interrupt	6	Raised when one of the processor's interrupt pins is asserted, providing that the corresponding field in the interrupt mask (IM) CSR is set and the global interrupt enable flag, IE.IE, is set.
SystemCall	7	Raised when an scall instruction is executed.

## Exception Processing

Exceptions occur in the execute pipeline stage. If there is an instruction in the memory pipeline stage, that instruction is first allowed to finish. All instructions from the Execute stage back are then killed and do not cause any user-transparent state changes. For example, no flags are set.

DataBusError exceptions are imprecise, so the exception may not necessarily be raised on the load or store instruction that caused the bus error (it will be raised on some later instruction). This imprecise behavior allows for high-performance implementations where the pipeline need not be stalled until a store instruction commits to memory. It is acceptable because bus errors are non-recoverable.

The sequence of operations performed by the processor after an exception depends upon the type of the highest priority exception that has occurred.

### Non-Debug Exceptions

When a non-debug related exception occurs, the address of the instruction in the execute stage of the processor's pipeline is saved in the ea register. The global interrupt enable held in the IE.IE register is copied to the IE.EIE register and then cleared, disabling interrupts. Finally, the PC is set to the address of the exception handler, which is derived from the address contained in the EBA, if the DC.RE bit is not set, or the DEBA, if the DC.RE bit is set.

```
ea = PC
IE.EIE = IE.IE
IE.IE = 0
PC = (DC.RE ? DEBA:EBA) + (ID * 32)
```

### Debug Exceptions

When a debug-related exception occurs, such as a breakpoint or watchpoint exception, the address of the instruction in the execute stage of the processor's pipeline is saved in the ba register. The global interrupt enable held in the IE.IE register is copied to the IE.BIE register and then cleared, disabling interrupts. Finally, the PC is set to the address of the exception handler, which is derived from the address contained in the DEBA.

```
ba = PC
IE.BIE = IE.IE
IE.IE = 0
PC = DEBA + (ID * 32)
```

## Nesting Exceptions

Because different registers are used to save a state when a debug-related exception occurs (ba and IE.BIE instead of ea and IE.EIE), limited nesting of exceptions is possible, allowing the interrupt handler code to be debugged. Any further nesting of exceptions requires software support.

## Interrupts

The LatticeMico32 processor supports up to 32 maskable, active-low, level-sensitive interrupts. Each interrupt line has a corresponding mask bit in the IM CSR. The mask enable is active high. A global interrupt enable flag is implemented in the IE CSR. Software can query the status of the interrupts and acknowledge them through the IP CSR.

If more interrupt sources or more sophisticated interrupt detection methods are required, external interrupt controllers can be cascaded onto the processor's interrupt pins to provide the needed functionality.

## Exception Handlers

When an exception occurs, the CPU branches to an address that is an offset from either the EBA CSR or the DEBA CSR. The offset is calculated by multiplying the exception ID by 32. Exception IDs are shown in Figure 10 on page 20. Since all LatticeMico32 instructions are four bytes long, this means each exception handler can be eight instructions long. If further instructions are required, the handler can call a subroutine.

Whether the EBA or DEBA is used as the base address depends upon the type of the exception that occurred and whether DC.RE is set. Having two different base addresses for the exception table allows a debug monitor to exist in a different memory from the main program code. For example, the debug monitor may exist in an on-chip ROM, whereas the main program code may be in a DDR or SRAM. The DC.RE flag allows either interrupts to run at full speed when debugging or for the debugger to take complete control and handle all exceptions.

When an exception occurs, the only state that is automatically saved by the CPU is the PC, which is saved in either ea or ba, and the interrupt enable flag, IE.IE, which is saved in either IE.EIE or IE.BIE. It is the responsibility of the exception handler to save and restore any other registers that it uses, if it returns to the previously executing code. This will likely be the case for interrupt exception handlers but may not be so for handlers that typically result in the termination of a program, such as divide by zero or bus error handlers.

The following piece of code shows how the exception handlers can be implemented. The nops are required to ensure that the next exception handler is aligned at the correct address.

To ensure that this code is at the correct address, it is common practice to place it in its own section. Place the following assembler directive at the start of the code:

```
.section    .boot, "ax", @progbits
```

**Figure 15: Exception Handler Example**

```
/* Exception handlers */

_reset_handler:
    xor    r0, r0, r0
    bi    _crt0
    nop
    nop
    nop
    nop
    nop
    nop
    nop

_breakpoint_handler:
    sw    (sp+0), ra
    calli save_all
    mvi   r1, SIGTRAP
    calli raise
    bi    restore_all_and_bret
    nop
    nop
    nop

_instruction_bus_error_handler:
    sw    (sp+0), ra
    calli save_all
    mvi   r1, SIGSEGV
    calli raise
    bi    restore_all_and_eret
    nop
    nop
    nop

_watchpoint_handler:
    sw    (sp+0), ra
    calli save_all
    mvi   r1, SIGTRAP
    calli raise
    bi    restore_all_and_bret
    nop
    nop
    nop

_data_bus_error_handler:
    sw    (sp+0), ra
    calli save_all
    mvi   r1, SIGSEGV
    calli raise
    bi    restore_all_and_eret
    nop
    nop
    nop
```

**Figure 15: Exception Handler Example (Continued)**

```

_divide_by_zero_handler:
    sw      (sp+0), ra
    calli   save_all
    mvi     r1, SIGFPE
    calli   raise
    bi      restore_all_and_eret
    nop
    nop
    nop

_interrupt_handler:
    sw      (sp+0), ra
    calli   save_all
    mvi     r1, SIGINT
    calli   raise
    bi      restore_all_and_eret
    nop
    nop
    nop

_system_call_handler:
    sw      (sp+0), ra
    calli   save_all
    mv      r1, sp
    calli   handle_scall
    bi      restore_all_and_eret
    nop
    nop
    nop

_save_all:
    addi    sp, sp, -56
    /* Save all caller save registers onto the stack */
    sw      (sp+4), r1
    sw      (sp+8), r2
    sw      (sp+12), r3
    sw      (sp+16), r4
    sw      (sp+20), r5
    sw      (sp+24), r6
    sw      (sp+28), r7
    sw      (sp+32), r8
    sw      (sp+36), r9
    sw      (sp+40), r10
    sw      (sp+48), ea
    sw      (sp+52), ba
    /* ra needs to be moved from initial stack location */
    lw      r1, (sp+56)
    sw      (sp+44), r1
    ret

```

**Figure 15: Exception Handler Example (Continued)**

```

/* Restore all registers and return from exception */
_restore_all_and_eret:
    lw    r1, (sp+4)
    lw    r2, (sp+8)
    lw    r3, (sp+12)
    lw    r4, (sp+16)
    lw    r5, (sp+20)
    lw    r6, (sp+24)
    lw    r7, (sp+28)
    lw    r8, (sp+32)
    lw    r9, (sp+36)
    lw    r10, (sp+40)
    lw    ra, (sp+44)
    lw    ea, (sp+48)
    lw    ba, (sp+52)
    addi  sp, sp, 56
    eret

/* Restore all registers and return from breakpoint */
_restore_all_and_bret:
    lw    r1, (sp+4)
    lw    r2, (sp+8)
    lw    r3, (sp+12)
    lw    r4, (sp+16)
    lw    r5, (sp+20)
    lw    r6, (sp+24)
    lw    r7, (sp+28)
    lw    r8, (sp+32)
    lw    r9, (sp+36)
    lw    r10, (sp+40)
    lw    ra, (sp+44)
    lw    ea, (sp+48)
    lw    ba, (sp+52)
    addi  sp, sp, 56
    bret

```

Then in the linker script, place the code at the reset value of EBA or DEBA, as shown in Figure 16.

**Figure 16: Placing Exception Handler in Memory**

```

MEMORY
{
    ram : ORIGIN = 0x00000000, LENGTH = 0x00100000
}

SECTIONS
{
    .boot : { *(.boot) } > ram
}

```

## Nested Exceptions

If desirable, it is possible to support nested exceptions with a small amount of software code. There are two situations where this may be needed:

- ◆ To allow system calls to be callable from exception handlers
- ◆ To prioritize interrupts

To enable nested exceptions, an exception handler must save all the state that is modified when an exception occurs, including the ea and ba registers, as well as the IE CSR. These registers can simply be saved on the stack. When returning from the exception handler, these registers must, obviously, be restored from the values saved on the stack.

To support nested prioritized interrupts, an exception handler should save the registers just outlined, then save the IM CSR, and then mask all lower-priority interrupts. IE.IE can then be set to re-enable interrupts. When the interrupt handler has finished, IE.IE should be cleared before all the saved registers, including IM, are restored.

## Remapping the Exception Table

In order to increase performance, the exception table can be remapped at run time by writing a new value to EBA. It would be used in a system in which the power-up value of EBA points to a slow, non-volatile memory, such as a FLASH memory, but the code is executed from a faster, non-volatile RAM, such as DDR or SRAM.

## Reset Summary

During reset, the following occurs:

- ◆ All CSRs are set to their reset values as listed in "Control and Status Registers" on page 9.
- ◆ Interrupts are disabled.
- ◆ All hardware breakpoints and watchpoints are disabled.
- ◆ If implemented, the contents of the caches are invalidated.
- ◆ A reset exception is raised, which causes the PC to be set to the value in the EBA CSR, where program execution starts.

The register file is not reset, so it is the responsibility of the reset exception handler to set register 0 to 0. This should be achieved by executing the following sequence: `xor r0, r0, r0`.

## Using Breakpoints

The LatticeMico32 architecture supports both software and hardware breakpoints. Software breakpoints should be used for setting breakpoints in code that resides in volatile memory, such as DDR or SRAM, while hardware breakpoints should be used for setting breakpoints in code that resides in non-volatile memory, such as FLASH or ROM.

A software breakpoint is simply a break instruction. In order to set a breakpoint, it is simply a case of replacing the instruction at the desired address with the break instruction. When the break instruction is executed, a breakpoint exception is raised, and the `ba` register contains the address of the break instruction that was executed. It is then up to the exception handler to either restore the instruction that was overwritten and continue execution, or to take some other action, depending upon why the breakpoint was set.

It is typically either not possible or very slow to write a break instruction to non-volatile RAM. For processors with breakpoints greater than 0, it is possible to set a hardware breakpoint by writing the address of the instruction on which the breakpoint should be set to one of the `BP $n$`  CSRs. The processor then constantly compares the values in these `BP $n$`  CSRs with the address of the instruction being executed. If a match occurs, and the breakpoint is enabled (by the LSB being set to 1), a breakpoint exception will be raised. As with software breakpoints, the address of the instruction that caused the breakpoint is saved in the `ba` register. If the breakpoint exception handler wishes to resume program execution, it must clear the enable bit in the relevant `BP` CSR; otherwise, the breakpoint exception is raised as soon as execution resumes.

## Using Watchpoints

The LatticeMico32 architecture supports hardware watchpoints. Watchpoints are a mechanism by which a program can watch out for specific memory accesses. For example, a program can set up a watchpoint that will cause a watchpoint exception to be raised every time the address 0 is accessed (something that is useful for tracking down null pointer errors in C programs).

To set up a watchpoint, the memory address that is being watched must be written to one of the `WP $n$`  CSRs. The watchpoint then needs to be enabled by writing the corresponding `C` field in the `DC` CSR. This field takes one of the four values that indicate the following:

- ◆ The watchpoint is disabled.
- ◆ The watchpoint exception is only raised on read accesses.
- ◆ The watchpoint exception is only raised on write accesses.
- ◆ The watchpoint exception is raised on either read or write accesses.

## Debug Architecture

This section describes the debug architecture of the LatticeMico32 processor.

The LatticeMico32 debug architecture provides:

- ◆ Software breakpoints
- ◆ Hardware breakpoints
- ◆ Hardware watchpoints
- ◆ Single-step capability
- ◆ Ability to remap exception handlers when debugging is enabled
- ◆ Hardware support for debugging interrupt handlers

Table 11 shows the debug control and status registers.

**Table 11: Debug Control and Status Registers**

Name	Access	Index	Description
DC	W	0x8	Debug control
DEBA	R/W	0x9	Debug exception base address
JTX	R/W	0xe	JTAG UART transmit
JRX	R/W	0xf	JTAG UART receive
BP0	W	0x10	Breakpoint address 0
BP1	W	0x11	Breakpoint address 1
BP2	W	0x12	Breakpoint address 2
BP3	W	0x13	Breakpoint address 3
WP0	W	0x18	Watchpoint address 0
WP1	W	0x19	Watchpoint address 1
WP2	W	0x1a	Watchpoint address 2
WP3	W	0x1b	Watchpoint address 3



**Table 13: Fields of the JTX CSR**

Field	Values	Description
TXD		Transmits data
F	0 – Empty 1 – Full	Indicates whether the transmit data register is full

## JRX – JTAG UART Receive Register

The JRX CSR can be used for receiving data through a JTAG interface. This CSR is only implemented if JTAG\_UART\_ENABLED equals TRUE.

**Figure 20: Format of the JRX CSR****Table 14: Fields of the JTX CSR**

Field	Values	Description
RXD		Receives data.
F	0 – Empty 1 – Full	Indicates whether the receive data register is full.

## BP<sub>n</sub> – Breakpoint

The BP<sub>n</sub> CSRs hold an instruction breakpoint address and a control bit that determines whether the breakpoint is enabled. Because instructions are always word-aligned, only the 30 most significant bits of the breakpoint address are needed. After reset, the value of the BP<sub>n</sub> CSRs is h00000000.

These CSRs are only implemented if DEBUG\_ENABLED equals TRUE.

**Figure 21: Format of the BP<sub>n</sub> CSRs****Table 15: BP<sub>n</sub> CSR Fields**

Field	Value	Description
E	b0 – Breakpoint is disabled b1 – Breakpoint is enabled	Breakpoint enable
A		Breakpoint address (Bits 31:2)

## WP<sub>n</sub> – Watchpoint

The WP<sub>n</sub> CSRs hold data watchpoint addresses. After reset, the value of the WP<sub>n</sub> CSRs is h00000000. These CSRs are only implemented if DEBUG\_ENABLED equals TRUE.

---

## Instruction Set Categories

---

LatticeMico32 supports a variety of instructions for arithmetic, logic, data comparison, data movement, and program control. Not all instructions are available in all configurations of the processor. Support for some types of instructions can be eliminated to reduce the amount of FPGA resources used. See “Configuring the LatticeMico32 Processor” on page 35.

Instructions ending with the letter “i” use an immediate value instead of a register. Instructions ending with “hi” use a 16-bit immediate and the high 16 bits from a register. Instructions ending with the letter “u” treat the data as unsigned integers.

For descriptions of individual instructions, see “Instruction Set” on page 49.

### Arithmetic

The instruction set includes the standard 32-bit integer arithmetic operations. Support for the multiply and divide instructions is optional.

- ◆ Add: add, addi
- ◆ Subtract: sub
- ◆ Multiply: mul, muli
- ◆ Divide and modulus: divu, modu

There are also instructions to sign-extend byte and half-word data to word size. Support for these instructions is optional.

- ◆ Sign-extend: sextb, sixth

### Logic

The instruction set includes the standard 32-bit bitwise logic operations. Most of the logic instructions also have 16-bit immediate or high 16-bit versions.

- ◆ AND: and, andi, andhi
- ◆ OR: or, ori, orhi
- ◆ Exclusive-OR: xor, xori
- ◆ Complement: not
- ◆ NOR: nor, nori
- ◆ Exclusive-NOR: xnor, xnori

## Comparison

The instruction set has basic comparison instructions with versions for register-to-register and register-to-16-bit-immediate and signed and unsigned comparisons. The instructions return 1 if true and 0 if false.

- ◆ Equal: cmpe, cmpei
- ◆ Not equal: cmpne, cmpnei
- ◆ Greater: cmpg, cmpgi, cmpgu, cmpgui
- ◆ Greater or equal: cmpge, cmpgei, cmpgeu, cmpgeui

## Shift

The instruction set supports left and right shifting of data in general-purpose registers. The number of bits to shift can be given through a register or a 5-bit immediate. The right shift instruction has signed and unsigned versions (also known as arithmetic and logical shifting). Support for shift instructions is optional.

- ◆ Left shift: sl, sli
- ◆ Right shift: sr, sri, sru, srui

## Data Transfer

Data transfer includes instructions that move data of byte, half-word, and word sizes between memory and registers. Memory addresses are relative and given as the sum of a general-purpose register and a signed 16-bit immediate, for example, (r2+32).

- ◆ Load register from memory: lb, lbu, lh, lhu, lw  
Byte and half-word values are either sign-extended or zero-extended to fill the register.
- ◆ Store register to memory: sb, sh, sw  
Byte and half-word values are taken from the lowest order part of the register.

There are also instructions for moving data from one register to another, including general-purpose and control and status registers.

- ◆ Move between general-purpose registers: mv
- ◆ Move immediate to high 16 bits of register: mvhi
- ◆ Read and write control and status register: rcsr, wcsr

## Program Flow Control

Program flow control instructions include branches, function and exception calls, and returns. The conditional branches and the immediate versions of the unconditional branch and call instructions establish the next instruction's address by adding a signed immediate to the PC register. Since the immediate is signed, the jump can be to a lower or higher address.

- ◆ Unconditional branch: b, bi

- ◆ Branch if equal: be
- ◆ Branch if not equal: bne
- ◆ Branch if greater: bg, bgu
- ◆ Branch if greater or equal: bge, bgeu
- ◆ Function call and return: call, calli, ret
- ◆ System call: scall
- ◆ Return from exception: eret
- ◆ Software breakpoint and return: break, bret



## Configuring the LatticeMico32 Processor

This chapter describes possible configuration options that you can use for the LatticeMico32 processor. You are expected to use the Lattice Mico System Builder (MSB) tool to configure the LatticeMico32 processor. Use the processor's configuration GUI, located in the MSB, to specify the Verilog parameters of the processor's RTL. For more information on the processor's configuration GUI, refer to LatticeMico32 online Help.

---

### Configuration Options

---

Table 16 describes the Verilog parameters for the LatticeMico32 processor.

**Table 16: Verilog Configuration Options**

Parameter Name	Values	Default	Description
MC_MULTIPLY_ENABLED	TRUE, FALSE	FALSE	Enables LUT-based multicycle multiplier. mul, multi instructions are implemented. Multiply instructions take 32 cycles to complete.
PL_MULTIPLY_ENABLED	TRUE, FALSE	TRUE	Enables pipelined multiplier (uses DSP blocks if available). mul, multi instructions are implemented. Multiply instructions take 3 cycles to complete.
DIVIDE_ENABLED	TRUE, FALSE	FALSE	Determines whether the divide and modulus instructions (divu, modu) are implemented.

**Table 16: Verilog Configuration Options (Continued)**

Parameter Name	Values	Default	Description
MC_BARREL_SHIFT_ENABLED	TRUE, FALSE	FALSE	Enables LUT-based multicycle barrel shifter. Enables shift instructions (sr, sri, sru, srui, sl, sli). Each shift instruction can take up to 32 cycles. If both SIGN_EXTEND_ENABLED and PL_BARREL_SHIFT_ENABLED are FALSE, this option must be set to TRUE.
PL_BARREL_SHIFT_ENABLED	TRUE, FALSE	TRUE	Enables pipelined barrel shifter. Enables shift instructions (sr, sri, sru, srui, sl, sli). Shift instructions take 3 cycles to complete. If both MC_BARREL_SHIFT_ENABLED and SIGN_EXTEND_ENABLED are FALSE, this option must be set to TRUE.
SIGN_EXTEND_ENABLED	TRUE, FALSE	FALSE	Determines whether the sign-extension instructions (sextb, sexth) are implemented. If both MC_BARREL_SHIFT_ENABLED and PL_BARREL_SHIFT_ENABLED are FALSE, this option must be set to TRUE.
DEBUG_ENABLED	TRUE, FALSE	TRUE	Determines whether software-based debugging support is implemented (that is, a ROM monitor is required to debug).
HW_DEBUG_ENABLED	TRUE, FALSE	TRUE	Determines whether hardware-based debugging support is implemented (that is, a ROM monitor is not required to debug). If this option is set to TRUE, DEBUG_ENABLED and JTAG_ENABLED must also be set to TRUE.
ROM_DEBUG_ENABLED	TRUE, FALSE	FALSE	Determines whether support for debugging ROM-based programs is implemented. If this option is set to TRUE, DEBUG_ENABLED must also be set to TRUE.
BREAKPOINTS	0-4		Specifies the number of breakpoint CSRs. If this option is set to a non-zero value, ROM_DEBUG_ENABLED must be set to TRUE.
WATCHPOINTS	0-4		Specifies the number of watchpoint CSRs. If this option is set to a non-zero value, SW_DEBUG_ENABLED must be set to TRUE.
JTAG_ENABLED	TRUE, FALSE	TRUE	Determines whether a JTAG interface is implemented.
JTAG_UART_ENABLED	TRUE, FALSE	TRUE	Determines whether a JTAG UART is implemented. If this option is set to TRUE, JTAG_ENABLED must be set to TRUE.
CYCLE_COUNTER_ENABLED	TRUE, FALSE	FALSE	Determines whether a cycle counter is implemented.
ICACHE_ENABLED	TRUE, FALSE	TRUE	Determines whether an instruction cache is implemented.

**Table 16: Verilog Configuration Options (Continued)**

Parameter Name	Values	Default	Description
ICACHE_BASE_ADDRESS	Any address aligned to the size of the cacheable region.	0	Specifies the base address of region cacheable by instruction cache.
ICACHE_LIMIT	Any integer multiple of the capacity of the cache added to the base address of the cacheable region	0x7FFFFFFF	Specifies the upper limit of region cacheable by instruction cache.
ICACHE_SETS	128, 256, 512, 1024	512	Specifies the number of sets in the instruction cache.
ICACHE_ASSOCIATIVITY	1, 2	1	Specifies the associativity of instruction cache.
ICACHE_BYTES_PER_LINE	4, 8, 16	4	Specifies the number of bytes per instruction cache line.
DCACHE_ENABLED	TRUE, FALSE	TRUE	Determines whether a data cache is implemented.
DCACHE_BASE_ADDRESS	Any address aligned to the size of the cacheable region	0	Specifies the base address of region cacheable by data cache.
DCACHE_LIMIT	Any integer multiple of the capacity of the cache added to the base address of the cacheable region	0xFFFFFFFF	Specifies the upper limit of region cacheable by data cache.
DCACHE_SETS	128, 256, 512, 1024	512	Specifies the number of sets in the data cache.
DCACHE_ASSOCIATIVITY	1, 2	1	Specifies the associativity of the data cache.
DCACHE_BYTES_PER_LINE	4, 8, 16	4	Specifies the number of bytes per data cache line.
INTERRUPTS	0–32	32	Specifies the number of external interrupts.
EBA_RESET	Any 256-byte aligned address	0	Specifies the reset value of the EBA CSR.

**Table 16: Verilog Configuration Options (Continued)**

Parameter Name	Values	Default	Description
DEBA_RESET	Any 256-byte aligned address	0	Specifies the reset value of the DEB_CSR.
EBR_POSEDGE_REGISTER_FILE	TRUE, FALSE	FALSE	Use EBR to implement register file instead of distributed RAM (LUTs).

## EBR Use

The following details of embedded block RAM (EBR) use with different configurations are based on the LatticeECP family of FPGAs.

- ◆ Software-based debugging (DEBUG\_ENABLED) requires two EBRs.
- ◆ The instruction and data caches (ICACHE\_ENABLED and DCACHE\_ENABLED, respectively) require EBR based on the size of the cache:

$$\text{cache size} = \text{sets} \times \text{bytes per cache line} \times \text{associativity}$$

$$\text{number of EBR} = \text{cache size}/1024 + 1$$

For example, the default LatticeMico32 processor in the MSB has software-based debugging, an instruction cache, and a data cache. Both caches have 512 sets, 16 bytes per cache line, and an associativity of 1.

For each cache:

$$\text{cache size} = 512 \times 16 \times 1 = 8192$$

$$\text{number of EBR} = 8192/1024 + 1 = 9$$

Total number of EBRs required:

Software-based debugging	2
Instruction cache	9
Data cache	9
	20

## WISHBONE Interconnect Architecture

This chapter describes the standard WISHBONE interconnect architecture that is employed by LatticeMico32 System. It focuses on the items that you must be aware of to begin designing and programming the functions of your system interconnects.

---

### Introduction to WISHBONE Interconnect

---

LatticeMico32 System uses a standard WISHBONE interconnect architecture to connect the processor to its on-chip component resources, such as the LatticeMico32 UART and the LatticeMico32 SPI.

The WISHBONE interconnect works as a general-purpose interface, defining the standard data exchanges between the processor module and its components. The interconnect does not interfere with the regulation of the processor or component application-specific functions. Like microcomputer buses, the WISHBONE bus is flexible enough to be tailored to a specific application, robust enough to provide a number of bus cycles and data path widths to solve various system issues, and universal enough to allow a number of suppliers to create design products for it, making it more cost-effective.

For more information on the WISHBONE System-on-Chip (SoC) Interconnection Architecture for Portable IP Cores, as it is formally known, refer to the OPENCORES.ORG Web site at [www.opencores.org/projects.cgi/web/wishbone](http://www.opencores.org/projects.cgi/web/wishbone). The subject matter is very detailed and goes beyond the scope of this manual.

## WISHBONE Registered Feedback Mode

This section describes the WISHBONE Registered Feedback mode. To implement an advanced synchronous cycle termination scheme, Registered Feedback mode bus cycles use the Cycle Type Identifier, CTI\_O( ) and CTI\_I( ), address tags. Both master and slave interfaces support CTI\_O( ) and CTI\_I( ) for improved bandwidth. The type of burst information is provided by the Burst Type Extension, BTE\_O( ) and BTE\_I( ) address tags.

All WISHBONE Registered Feedback-compatible cores must support WISHBONE Classic bus cycles.

Design new IP cores to support WISHBONE Registered Feedback bus cycles to ensure maximum throughput in all systems.

### CTI\_IO( )

The cycle-type identifier CTI\_IO( ) address tag provides additional information about the current cycle. The master sends this information to the slave. The slave can use this information to prepare the response for the next cycle.

**Table 17: Cycle Type Identifiers**

CTI_O(2:0)	Description
000	Classic cycle
001	Constant address burst cycle
010	Incrementing burst cycle
011	Reserved
100	Reserved
101	Reserved
110	Reserved
111	End of burst

Observe the following allowances and rules:

- ◆ Master and slave interfaces may be designed to support the CTI\_I( ) and CTI\_O( ) signals. Also, master and slave interfaces may be designed to support a limited number of burst types.
- ◆ Master and slave interfaces that do support the CTI\_I( ) and CTI\_O( ) signals must at least support the Classic cycle CTI\_IO( )=000 and the End-of-Cycle CTI\_IO( )=111.
- ◆ Master and slave interfaces that are designed to support a limited number of burst types must complete the unsupported cycles as though they were WISHBONE Classic cycle, that is, CTI\_IO( )=000.
- ◆ For description languages that allow default values for input ports (like VHDL), CTI\_I( ) may be assigned a default value of 000.

- ◆ In addition to the WISHBONE Classic rules for generating cycle termination signals ACK\_O, RTY\_O, and ERR\_O, a SLAVE may assert a termination cycle without checking the STB\_I signal.
- ◆ ACK\_O, RTY\_O, and ERR\_O may be asserted while STB\_O is negated.
- ◆ A cycle terminates when the cycle termination signal, STB\_I and STB\_O are asserted. Even if ACK\_O/ACK\_I is asserted, the other signals are only valid when STB\_O/STB\_I is also asserted.

To avoid the inherent wait state in synchronous termination schemes, the slave must generate the response as soon as possible, that is, the next cycle. It can use the CTI\_I( ) signals to determine the response for the next cycle, but if it cannot determine the state of STB\_I for the next cycle, it must generate the response independent of STB\_I.

## BTE\_IO( )

The burst-type extension BTE\_IO( ) address tag provides additional information about the current burst. The master sends this information to the slave. This information is only relevant for incrementing bursts. In the future, other burst types may use these signals. See Table 18 for BTE\_IO(1:0) signal incrementing and decrementing bursts.

**Table 18: Burst Type Extension Signal Bursts**

BTE_IO(1:0)	Description
00	Linear burst
01	4-beat wrap burst
10	8-beat wrap burst
11	16-beat wrap burst

Observe the following allowances and rules:

- ◆ Master and slave interfaces that support incrementing burst cycles must support the BTE\_O( ) and BTE\_I( ) signals.
- ◆ Master and slave interfaces may be designed to support a limited number of burst extensions.
- ◆ Master and slave interfaces that are designed to support a limited number of burst extensions must complete the unsupported cycles as though they were WISHBONE Classic cycle, that is, CTI\_IO( )= 000.

## Component Signals

In Mico System Builder (MSB), you define which components are in the platform and what needs to communicate with what. When the platform generator is run in MSB, it uses this information to build the WISHBONE-based interconnect of the platform. This generated interconnect is a set of Verilog wires connecting the various processor and component ports. To do this, the components must implement certain ports and follow a specific port-naming convention.

Table 19 defines the suffixes that must be used on the names of a component's ports. The suffixes of the ports of a master port are different than those of a slave port. The generated interconnect creates signals with names that end with the same suffix as the component port to which the signal is attached. Table 19 also notes which signals are mandatory and which are optional to support the basic WISHBONE bus cycle.

The prefixes used in the port and signal naming are not described in this section.

The port and signal descriptions that follow refer to the port or signal that ends with the string in the title.

**Table 19: List of Component Port and Signal Name Suffixes**

Master Ports			Slave Ports		
Name	Width	Optional (O)/Mandatory (M)	Name	Width	Optional (O)/Mandatory (M)
_ADR_O	32 bits	M	_ADR_I	32 bits	M
_DAT_O	32 bits	M	_DAT_I	32 bits	M
_DAT_I	32 bits	M	_DAT_O	32 bits	M
_SEL_O	4 bits	M	_SEL_I	4 bits	M
_WE_O	1 bit	M	_WE_I	1 bit	M
_ACK_I	1 bit	M	_ACK_O	1 bit	M
_ERR_I	1 bit	O	_ERR_O	1 bit	O
_RTY_I	1 bit	O	_RTY_O	1 bit	O
_CTI_O	3 bits	O	_CTI_I	3 bits	O
_BTE_O	2 bits	O	_BTE_I	2 bits	O
_LOCK_O	1 bit	O	_LOCK_I	1 bit	O
_CYC_O	1 bit	M	_CYC_I	1 bit	M
_STB_O	1 bit	M	_STB_I	1 bit	M

## Master Port and Signal Descriptions

This section describes the master ports and signals listed in Table 19.

### ADR\_O [31:2]

The address output array ADR\_O( ) is used to pass a binary address. ADR\_O( ) actually has a full 32 bits. But, because all addressing is on DWORD (4-byte) boundaries, the lowest two bits are always zero.

### DAT\_O [31:0]

The data output array DAT\_O( ) is used to store a binary value for output.

### DAT\_I [31:0]

The data input array DAT\_I( ) is used to store a binary value for input.

### SEL\_O [3:0]

The Select Output array SEL\_O( ) indicates where valid data is expected on the DAT\_I( ) signal array during READ cycles and where it is placed on the DAT\_O( ) signal array during WRITE cycles. The array boundaries are determined by the granularity of a port.

### WE\_O

The write enable output WE\_O indicates whether the current local bus cycle is a READ or WRITE cycle. The signal is negated during READ cycles and is asserted during WRITE cycles.

### ACK\_I

This signal is called the acknowledge input ACK\_I. When asserted, the signal indicates the normal termination of a bus cycle by the slave. Also see the ERR\_I and RTY\_I signal descriptions.

### ERR\_I

The Error Input ERR\_I indicates an abnormal cycle termination by the slave. The source of the error and the response generated by the master depends on the master functionality. Also see the ACK\_I and RTY\_I signal descriptions.

### RTY\_I

The Retry Input RTY\_I indicates that the interface is not ready to accept or send data, so the cycle should be retried. The core functionality defines when and how the cycle is retried. Also see the ERR\_I and RTY\_I signal descriptions.

### **CTI\_O [2:0]**

For descriptions of the cycle-type identifier CTI\_O( ), see “CTI\_IO( )” on page 40.

### **BTE\_O [1:0]**

For descriptions of the burst-type extension BTE\_O( ), see “BTE\_IO( )” on page 41.

### **LOCK\_O**

The lock output LOCK\_O, when asserted, indicates that the current bus cycle cannot be interrupted. Lock is asserted to request complete ownership of the bus. After the transfer starts, the INTERCON does not grant the bus to any other master until the current master negates LOCK\_O or CYC\_O.

### **CYC\_O**

The cycle output CYC\_O, when asserted, indicates that a valid bus cycle is in progress. The signal is asserted for the duration of all bus cycles. For example, during a BLOCK transfer cycle there can be multiple data transfers. The CYC\_O signal is asserted during the first data transfer and remains asserted until the last data transfer. The CYC\_O signal is useful for interfaces with multi-port interfaces, such as dual-port memories. In these cases, the CYC\_O signal requests the use of a common bus from an arbiter.

### **STB\_O**

The strobe output STB\_O indicates a valid data transfer cycle. It is used to qualify various other signals on the interface, such as SEL\_O( ). The slave asserts either the ACK\_I, ERR\_I, or RTY\_I signals in response to every assertion of the STB\_O signal.

## **Slave Port and Signal Descriptions**

This section describes the slave ports and signals listed in the Table 19.

### **ADR\_I [31:2]**

The address input array ADR\_I( ) is used to pass a binary address. ADR\_I( ) actually has a full 32 bits. But, because all addressing is on DWORD (4-byte) boundaries, the lowest two bits are always zero.

### **DAT\_I [31:0]**

The data input array DAT\_I( ) is used to store a binary value for input.

### **DAT\_O [31:0]**

The data output array DAT\_O( ) is used to store a binary value for output.

### **SEL\_I [3:0]**

The select input array SEL\_I( ) indicates where valid data is placed on the DAT\_I( ) signal array during WRITE cycles and where it should be present on the DAT\_O( ) signal array during READ cycles. The array boundaries are determined by the granularity of a port.

## **WE\_I**

The write enable Input WE\_I indicates whether the current local bus cycle is a READ or WRITE cycle. The signal is negated during READ cycles and is asserted during WRITE cycles.

## **ACK\_O**

The acknowledge output ACK\_O, when asserted, indicates the termination of a normal bus cycle by the slave. Also see the ERR\_O and RTY\_O signal descriptions.

## **ERR\_O**

The error output ERR\_O indicates an abnormal cycle termination by the slave. The source of the error and the response generated by the master depends on the master functionality. Also see the ACK\_O and RTY\_O signal descriptions.

## **RTY\_O**

The retry output RTY\_O indicates that the slave interface is not ready to accept or send data, so the cycle should be retried. The core functionality defines when and how the cycle is retried. Also see the ERR\_O and RTY\_O signal descriptions.

## **CTI\_I**

For descriptions of the cycle-type identifier CTI\_I( ), see “CTI\_IO( )” on page 40.

## **BTE\_I [1:0]**

For descriptions of the burst-type extension BTE\_i( ), see “BTE\_IO( )” on page 41.

## **LOCK\_I**

The lock input LOCK\_I, when asserted, indicates that the current bus cycle is uninterruptible. A slave that receives the LOCK LOCK\_I signal is accessed by a single master only until either LOCK\_I or CYC\_I is negated.

## **CYC\_I [2:0]**

The Cycle Input CYC\_I, when asserted, indicates that a valid bus cycle is in progress. The signal is asserted for the duration of all bus cycles. For example, during a BLOCK transfer cycle there can be multiple data transfers. The CYC\_I signal is asserted during the first data transfer and remains asserted until the last data transfer.

## **STB\_I**

The strobe input STB\_I, when asserted, indicates a valid data transfer cycle. A slave responds to other WISHBONE signals only when this STB\_I is asserted, except for the RST\_I signal, to which it should always respond. The slave asserts either the ACK\_O, ERR\_O, or RTY\_O signals in response to every assertion of the STB\_I signal.

## Arbitration Schemes

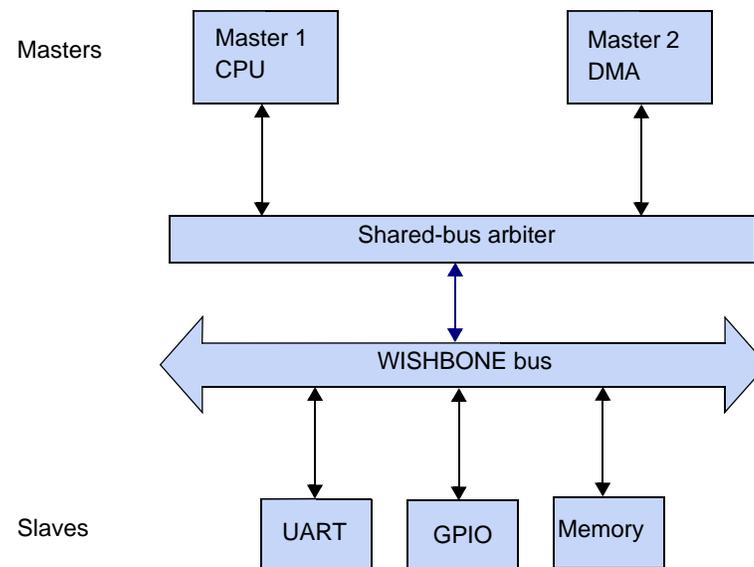
MSB supports the following arbitration schemes for platform generation:

- ◆ Shared-bus arbitration schemes
- ◆ Slave-side fixed arbitration schemes
- ◆ Slave-side round-robin arbitration schemes

### Shared-Bus Arbitration

The shared-bus arbitration scheme is shown in Figure 22.

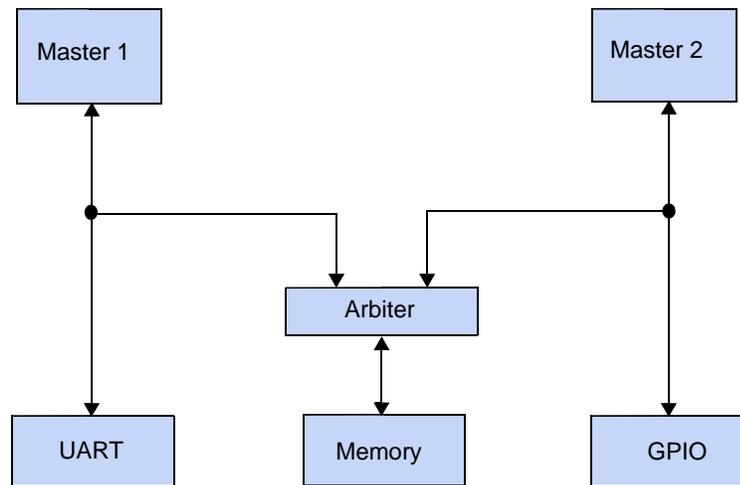
**Figure 22: Bus Architecture with Shared-Bus Arbitration**



In the shared-bus arbitration scheme, one or more bus masters and bus slaves connect to a shared bus. A single arbiter controls the bus, that is, the path between masters and slaves. Each bus master requests control of the bus from the arbiter, and the arbiter grants access to a single master at a time. Once a master has control of the bus, it performs transfers with a bus slave. If multiple masters attempt to access the bus at the same time, the arbiter allocates the bus resources to a single master according to fixed arbitration rules, forcing all other masters to wait.

### Slave-Side Arbitration

Slave-side arbitration is shown in Figure 23.

**Figure 23: Bus Architecture with Slave-Side Arbitration**

In slave-side arbitration, each multi-master slave has its own arbiter. A master port never waits to access a slave port, unless a different master port attempts to access the same slave port at the same time. As a result, multiple master ports active at the same time simultaneously transfer data with independent slave ports.

In the slave-side arbitration scheme, arbitration is only required when two or more masters contend for the same slave port. This scheme is called slave-side arbitration because it is implemented when two or more masters connect to a single slave.

### Slave-Side Fixed Arbitration

In the slave-side fixed arbitration scheme, when two or more masters request control of the bus for the same slave simultaneously, the master with the highest priority gains access to the bus. At every slave transfer, only requesting masters are included in the arbitration. The master with the highest priority is granted access to the bus.

### Slave-Side Round-Robin Arbitration

In the slave-side round-robin arbitration scheme, when multiple masters contend for access to a slave port, the arbiter grants access to the bus in round-robin order. At every slave transfer, only requesting masters are included in the round-robin arbitration.



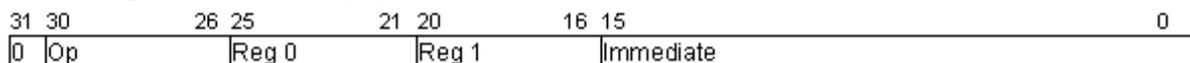
## Instruction Set

This chapter includes descriptions of all of the instruction opcodes of the LatticeMico32 processor.

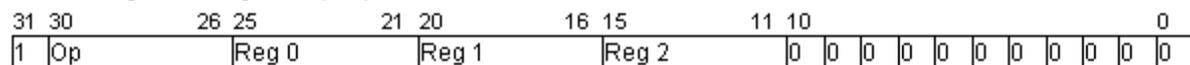
### Instruction Formats

All LatticeMico32 instructions are 32 bits wide. They are in four basic formats, as shown in Figure 24 through Figure 27.

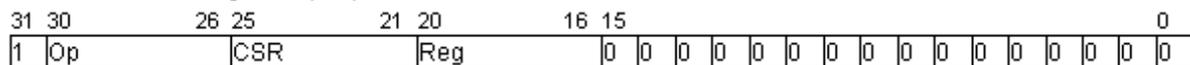
**Figure 24: Register Immediate (RI) Format**



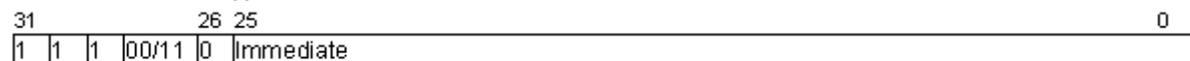
**Figure 25: Register Register (RR) Format**



**Figure 26: Control Register (CR) Format**



**Figure 27: Immediate (I) Format**



## Opcode Look-Up Table

Opcode	Decimal	Hexadecimal	Mnemonic	Opcode	Decimal	Hexadecimal	Mnemonic
000000	00	00	sru	100000	32	20	sru
000001	01	01	nori	100001	33	21	nor
000010	02	02	muli	100010	34	22	mul
000011	03	03	sh	100011	35	23	divu
000100	04	04	lb	100100	36	24	rcsr
000101	05	05	sri	100101	37	25	sr
000110	06	06	xori	100110	38	26	xor
000111	07	07	lh	100111	39	27	div
001000	08	08	andi	101000	40	28	and
001001	09	09	xnori	101001	41	29	xnor
001010	10	0A	lw	101010	42	2A	reserved
001011	11	0B	lhu	101011	43	2B	raise
001100	12	0C	sb	101100	44	2C	sextb
001101	13	0D	addi	101101	45	2D	add
001110	14	0E	ori	101110	46	2E	or
001111	15	0F	sli	101111	47	2F	sl
010000	16	10	lbu	110000	48	30	b
010001	17	11	be	110001	49	31	modu
010010	18	12	bg	110010	50	32	sub
010011	19	13	bge	110011	51	33	reserved
010100	20	14	bgeu	110100	52	34	wcsr
010101	21	15	bgu	110101	53	35	mod
010110	22	16	sw	110110	54	36	call
010111	23	17	bne	110111	55	37	sexth
011000	24	18	andhi	111000	56	38	bi
011001	25	19	cmpei	111001	57	39	cmpe
011010	26	1A	cmpgi	111010	58	3A	cmpg
011011	27	1B	cmpgei	111011	59	3B	cmpge
011100	28	1C	cmpgeui	111100	60	3C	cmpgeu
011101	29	1D	cmpgui	111101	61	3D	cmpgu
011110	30	1E	orhi	111110	62	3E	calli
011111	31	1F	cmpnei	111111	63	3F	cmpne

## Pseudo-Instructions

To aid the semantics of assembler programs, the LatticeMico32 assembler implements a variety of pseudo-instructions. Table 20 lists these instructions and to what actual instructions they are mapped. Disassemblers show the actual implementation.

**Table 20: Pseudo-Instructions**

Mnemonic	Implementation	Description
ret	b ra	Returns from function call.
mv rX, rY	or rX, rY, r0	Moves value in rY to rX.
mvhi rX, imm16	orhi rX, r0, imm16	Moves the 16-bit, left-shifted immediate into rX.
not rX, rY	xnor rX, rY, r0	Is the bitwise complement of the value in rY and stores the result in rX.
mvi	addi rd, r0, imm16	Adds 16-bit immediate to r0 and stores the result in rd.  <b>Note:</b> GCC compiler tool chain expects r0 contents to be zero.
nop	addi r0, r0, 0	Adds 0 to r0 and saves it to r0, resulting in no operation (nop).

## Instruction Descriptions

Some of the following tables include these parameters:

- ◆ Syntax – Describes the assembly language syntax for the instruction.
- ◆ Issue – The “issue” cycles mean the number of cycles that the microprocessor takes to place this instruction in the pipeline. For example, if the issue is 1 cycle, the next instruction will be introduced into the pipeline the very next cycle. If the issue is 4, the next instruction will be introduced three cycles later. The branches and calls are issue 4 cycles, which means that the pipeline stalls for the next three cycles.
- ◆ Semantics – Describes how the instruction creates a result from the inputs and where it puts the result. The Semantics feature refers to terms used in the assembly language syntax for the instruction.

The Semantics feature also uses the following terms:

- ◆ gpr – Refers to a general-purpose register.
- ◆ PC – Refers to a program counter.
- ◆ csr – Refers to a control and status register.
- ◆ IE.BIE – Refers to the BIE bit of the IE (interrupt enable) register.
- ◆ IE.IE – Refers to the IE bit of the IE (interrupt enable) register.
- ◆ IE.EIE – Refers to the EIE bit of the IE (interrupt enable) register.



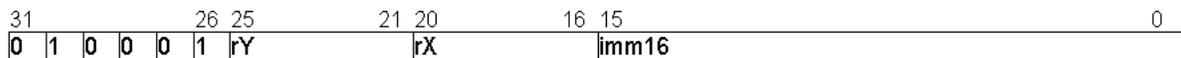




**Table 26: b Instruction Features**

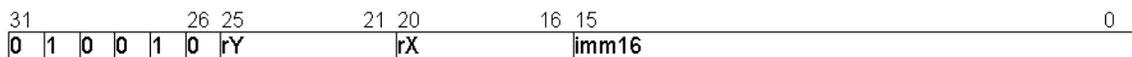
Feature	Description
Syntax	b rX
Example	b r3
Semantics	PC = gpr[rX]
Issue	4 cycles
See Also	bi, branch with immediate

## be

**Figure 34: be Instruction****Table 27: be Instruction Features**

Feature	Description
Operation	Branch if equal
Description	Compares the value in rX with the value in rY, branching to the address given by the sum of the PC and the sign-extended immediate if the values are equal.
Syntax	be rX, rY, imm16
Example	be r4, r2, label
Semantics	if (gpr[rX] == gpr[rY]) PC = PC + sign_extend(imm16 << 2)
Issue	1 cycle (not taken), 4 cycles (taken)
See Also	bne, branch if not equal

## bg

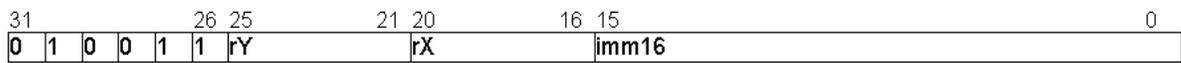
**Figure 35: bg Instruction****Table 28: bg Instruction Features**

Feature	Description
Operation	Branch if greater
Description	Compares the value in rX with the value in rY, branching to the address given by the sum of the PC and the sign-extended immediate if the value in rX is greater than the value in rY. The values in rX and rY are treated as signed integers.

**Table 28: bg Instruction Features**

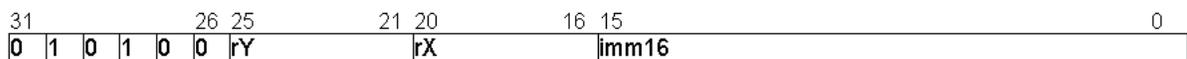
Feature	Description
Syntax	<code>bg rX, rY, imm16</code>
Example	<code>bg r4, r2, label</code>
Semantics	<code>if (gpr[rX] &gt; gpr[rY]) PC = PC + sign_extend(imm16 &lt;&lt; 2)</code>
Issue	1 cycle (not taken), 4 cycles (taken)
See Also	<code>bgu</code> , branch if greater, unsigned

## bge

**Figure 36: bge Instruction****Table 29: bge Instruction Features**

Feature	Description
Operation	Branch if greater or equal
Description	Compares the value in rX with the value in rY, branching to the address given by the sum of the PC and the sign-extended immediate if the value in rX is greater or equal to the value in rY. The values in rX and rY are treated as signed integers.
Syntax	<code>bge rX, rY, imm16</code>
Example	<code>bge r4, r2, label</code>
Semantics	<code>if (gpr[rX] &gt;= gpr[rY]) PC = PC + sign_extend(imm16 &lt;&lt; 2)</code>
Issue	1 cycle (not taken), 4 cycles (taken)
See Also	<code>bgeu</code> , branch if greater or equal, unsigned

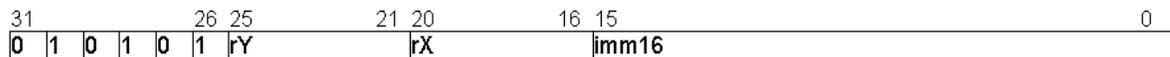
## bgeu

**Figure 37: bgeu Instruction**

**Table 30: bgeu Instruction Features**

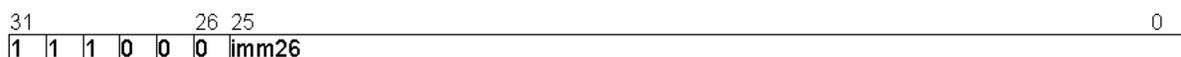
Feature	Description
Operation	Branch if greater or equal, unsigned
Description	Compares the value in rX with the value in rY, branching to the address given by the sum of the PC and the sign-extended immediate if the value in rX is greater or equal to the value in rY. The values in rX and rY are treated as unsigned integers.
Syntax	bgeu rX, rY, imm16
Example	bgeu r4, r2, label
Semantics	if (gpr[rX] >= gpr[rY]) PC = PC + sign_extend(imm16 << 2)
Issue	1 cycle (not taken), 4 cycles (taken)
See Also	bge, branch if greater or equal, signed

## bgu

**Figure 38: bgu Instruction****Table 31: bgu Instruction Features**

Feature	Description
Operation	Branch if greater, unsigned
Description	Compares the value in rX with the value in rY, branching to the address given by the sum of the PC and the sign-extended immediate if the value in rX is greater than the value in rY. The values in rX and rY are treated as unsigned integers.
Syntax	bgu rX, rY, imm16
Example	bgu r4, r2, label
Semantics	if (gpr[rX] > gpr[rY]) PC = PC + sign_extend(imm16 << 2)
Issue	1 cycle (not taken), 4 cycles (taken)
See Also	bg, branch if greater, signed

## bi

**Figure 39: bi Instruction**

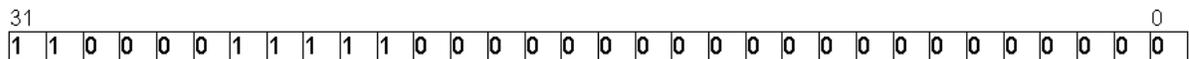


**Table 34: break Instruction Features**

Feature	Description
Operation	Software breakpoint
Description	Raises a breakpoint exception.
Syntax	<code>break</code>
Example	<code>break</code>
Semantics	$gpr[ba] = PC$ $IE.BIE = IE.IE$ $IE.IE = 0$ $PC = DEBA + ID * 32$
Issue	4 cycles
See Also	<code>bret</code> , return from breakpoint

## bret

**Figure 42: bret Instruction**

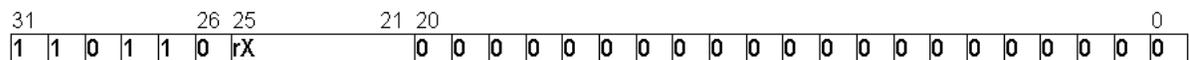


**Table 35: bret Instruction Features**

Feature	Description
Operation	Return from breakpoint
Description	Unconditional branch to the address in the breakpoint address register (ba), updating interrupt enable with value saved in breakpoint interrupt enable register.
Syntax	<code>bret</code>
Example	<code>bret</code>
Semantics	$PC = gpr[ba]$ $IE.IE = IE.BIE$
Issue	4 cycles
See Also	<code>break</code> , breakpoint

## call

**Figure 43: call Instruction**

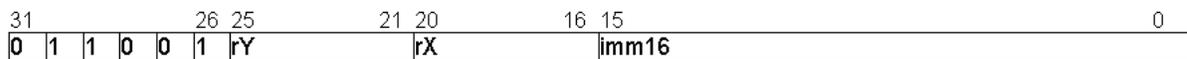




**Table 38: cmpe Instruction Features**

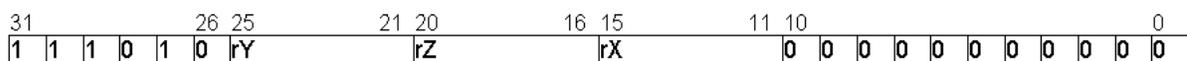
Feature	Description
Operation	Compare equal
Description	Compares the value in rY with the value in rZ, storing 1 in rX if they are equal, otherwise 0.
Syntax	<code>cmpe rX, rY, rZ</code>
Example	<code>cmpe r14, r15, r17</code>
Semantics	<code>gpr[rX] = gpr[rY] == gpr[rZ]</code>
Result	2 cycles
Issue	1 cycle
See Also	<code>cmpei</code> , compare equal with immediate

## cmpei

**Figure 46: cmpei Instruction****Table 39: cmpei Instruction Features**

Feature	Description
Operation	Compare equal
Description	Compares the value in rY with the sign-extended immediate, storing 1 in rX if they are equal, 0 otherwise.
Syntax	<code>cmpei rX, rY, imm16</code>
Example	<code>cmpei r4, r2, 0x5555</code>
Semantics	<code>gpr[rX] = gpr[rY] == sign_extend(imm16)</code>
Result	2 cycles
Issue	1 cycle
See Also	<code>cmpe</code> , compare equal between registers

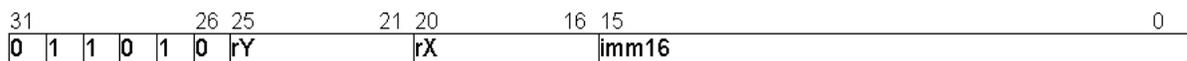
## cmpg

**Figure 47: cmpg Instruction**

**Table 40: cmpg Instruction Features**

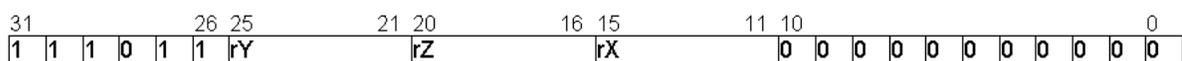
Feature	Description
Operation	Compare greater
Description	Compares the value in rY with the value in rZ, storing 1 in rX if the value in rY is greater than the value in rZ, 0 otherwise. Both operands are treated as signed integers.
Syntax	<code>cmpg rX, rY, rZ</code>
Example	<code>cmpg r14, r15, r17</code>
Semantics	<code>gpr[rX] = gpr[rY] &gt; gpr[rZ]</code>
Result	2 cycles
Issue	1 cycle
See Also	<code>cmpgi</code> , compare greater with immediate; <code>cmpgu</code> , compare greater, unsigned; <code>cmpgui</code> , compare greater with immediate, unsigned

## cmpgi

**Figure 48: cmpgi Instruction****Table 41: cmpgi Instruction Features**

Feature	Description
Operation	Compare greater
Description	Compares the value in rY with the sign-extended immediate, storing 1 in rX if the value in rY is greater than the immediate, 0 otherwise. Both operands are treated as signed integers.
Syntax	<code>cmpgi rX, rY, imm16</code>
Example	<code>cmpgi r4, r2, 0x5555</code>
Semantics	<code>gpr[rX] = gpr[rY] &gt; sign_extend(imm16)</code>
Result	2 cycles
Issue	1 cycle
See Also	<code>cmpg</code> , compare greater between registers; <code>cmpgu</code> , compare greater, unsigned; <code>cmpgui</code> , compare greater with immediate, unsigned

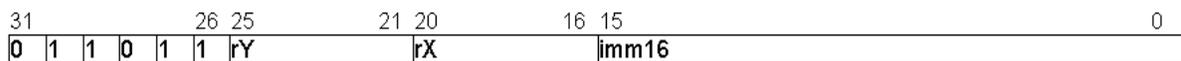
## cmpge

**Figure 49: cmpge Instruction**

**Table 42: cmpge Instruction Features**

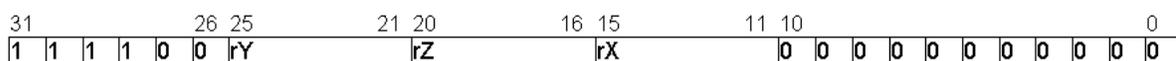
Feature	Description
Operation	Compare greater or equal
Description	Compares the value in rY with the value in rZ, storing 1 in rX if the value in rY is greater or equal to the value in rZ, 0 otherwise. Both operands are treated as signed integers.
Syntax	<code>cmpge rX, rY, rZ</code>
Example	<code>cmpge r14, r15, r17</code>
Semantics	<code>gpr[rX] = gpr[rY] &gt;= gpr[rZ]</code>
Result	2 cycles
Issue	1 cycle
See Also	<code>cmpgei</code> , compare with immediate; <code>cmpgeu</code> , compare, unsigned; <code>cmpgeui</code> , compare with immediate, unsigned

## cmpgei

**Figure 50: cmpgei Instruction****Table 43: cmpgei Instruction Features**

Feature	Description
Operation	Compare greater or equal
Description	Compares the value in rY with the sign-extended immediate, storing 1 in rX if the value in rY is greater or equal to the immediate, 0 otherwise. Both operands are treated as signed integers.
Syntax	<code>cmpgei rX, rY, imm16</code>
Example	<code>cmpgei r4, r2, 0x5555</code>
Semantics	<code>gpr[rX] = gpr[rY] &gt;= sign_extend(imm16)</code>
Result	2 cycles
Issue	1 cycle
See Also	<code>cmpge</code> , compare between registers; <code>cmpgeu</code> , compare, unsigned; <code>cmpgeui</code> , compare with immediate, unsigned

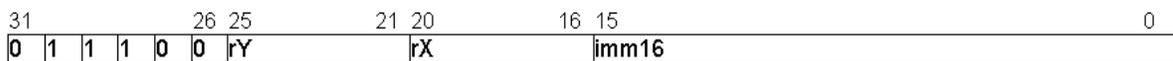
## cmpgeu

**Figure 51: cmpgeu Instruction**

**Table 44: cmpgeu Instruction Features**

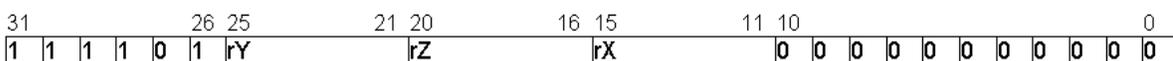
Feature	Description
Operation	Compare greater or equal
Description	Compares the value in rY with the value in rZ, storing 1 in rX if the value in rY is greater or equal to the value in rZ, 0 otherwise. Both operands are treated as unsigned integers.
Syntax	<code>cmpgeu rX, rY, rZ</code>
Example	<code>cmpgeu r14, r15, r17</code>
Semantics	<code>gpr[rX] = gpr[rY] &gt;= gpr[rZ]</code>
Result	2 cycles
Issue	1 cycle
See Also	<code>cmpge</code> , compare between registers; <code>cmpgei</code> , compare with immediate; <code>cmpgeui</code> , compare with immediate, unsigned

## cmpgeui

**Figure 52: cmpgeui Instruction****Table 45: cmpgeui Instruction Features**

Feature	Description
Operation	Compare greater or equal
Description	Compares the value in rY with the zero-extended immediate, storing 1 in rX if the value in rY is greater or equal to the immediate, 0 otherwise. Both operands are treated as unsigned integers.
Syntax	<code>cmpgeui rX, rY, imm16</code>
Example	<code>cmpgeui r4, r2, 0x5555</code>
Semantics	<code>gpr[rX] = gpr[rY] &gt;= zero_extend(imm16)</code>
Result	2 cycles
Issue	1 cycle
See Also	<code>cmpge</code> , compare between registers; <code>cmpgei</code> , compare with immediate; <code>cmpgeu</code> , compare, unsigned

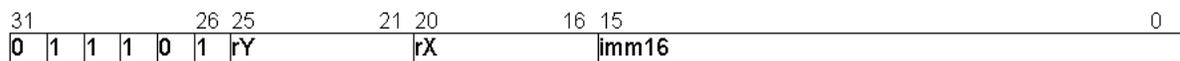
## cmpgu

**Figure 53: cmpgu Instruction**

**Table 46: cmpgu Instruction Features**

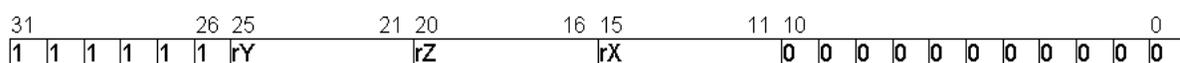
Feature	Description
Operation	Compare greater unsigned
Description	Compares the value in rY with the value in rZ, storing 1 in rX if the value in rY is greater than the value in rZ, 0 otherwise. Both operands are treated as unsigned integers.
Syntax	<code>cmpgu rX, rY, rZ</code>
Example	<code>cmpgu r14, r15, r17</code>
Semantics	<code>gpr[rX] = gpr[rY] &gt; gpr[rZ]</code>
Result	2 cycles
Issue	1 cycle
See Also	<code>cmpg</code> , compare greater, signed; <code>cmpgi</code> , compare greater with immediate; <code>cmpgui</code> , compare greater with immediate, unsigned

## cmpgui

**Figure 54: cmpgui Instruction****Table 47: cmpgui Instruction Features**

Feature	Description
Operation	Compare greater unsigned
Description	Compares the value in rY with the zero-extended immediate, storing 1 in rX if the value in rY is greater than the immediate, 0 otherwise. Both operands are treated as unsigned integers.
Syntax	<code>cmpgui rX, rY, imm16</code>
Example	<code>cmpgui r4, r2, 0x5555</code>
Semantics	<code>gpr[rX] = gpr[rY] &gt; zero_extend(imm16)</code>
Result	2 cycles
Issue	1 cycle
See Also	<code>cmpg</code> , compare greater, signed; <code>cmpgi</code> , compare greater with immediate; <code>cmpgu</code> , compare greater, unsigned

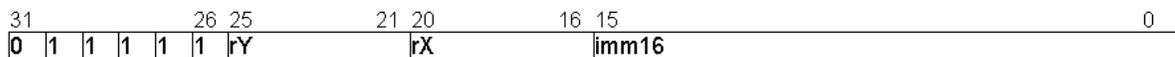
## cmpne

**Figure 55: cmpne Instruction**

**Table 48: cmpne Instruction Features**

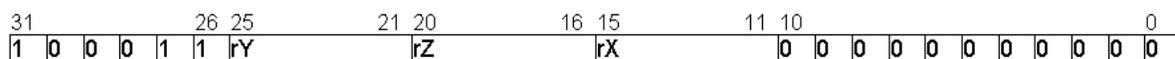
Feature	Description
Operation	Compare not equal
Description	Compares the value in rY with the value in rZ, storing 1 in rX if they are not equal, 0 otherwise.
Syntax	<code>cmpne rX, rY, rZ</code>
Example	<code>cmpne r14, r15, r17</code>
Semantics	<code>gpr[rX] = gpr[rY] != gpr[rZ]</code>
Result	2 cycles
Issue	1 cycle
See Also	<code>cmpnei</code> , compare not equal with immediate

## cmpnei

**Figure 56: cmpnei Instruction****Table 49: cmpnei Instruction Features**

Feature	Description
Operation	Compare not equal
Description	Compares the value in rY with the sign-extended immediate, storing 1 in rX if they are not equal, 0 otherwise.
Syntax	<code>cmpnei rX, rY, imm16</code>
Example	<code>cmpnei r4, r2, 0x5555</code>
Semantics	<code>gpr[rX] = gpr[rY] != sign_extend(imm16)</code>
Result	2 cycles
Issue	1 cycle
See Also	<code>cmpne</code> , compare not equal between registers

## divu

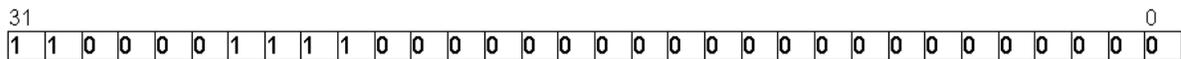
**Figure 57: divu Instruction**

**Table 50: divu Instruction Features**

Feature	Description
Operation	Unsigned iinteger division
Description	Divides the value in rY by the value in rZ, storing the quotient in rX. Both operands are treated as unsigned integers. Available only if the processor was configured with the DIVIDE_ENABLED option.
Syntax	<code>divu rX, rY, rZ</code>
Example	<code>divu r14, r15, r17</code>
Semantics	$gpr[rX] = gpr[rY] / gpr[rZ]$
Result	34 cycles
Issue	34 cycles
See Also	<code>modu</code> , <code>modulus</code>

## eret

**Figure 58: eret Instruction**

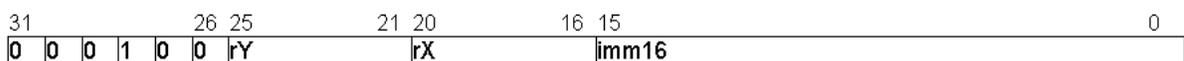


**Table 51: eret Instruction Features**

Feature	Description
Operation	Return from exception
Description	Unconditionally branches to the address in the exception address register (ea), updating interrupt enable with value saved in exception interrupt enable register.
Syntax	<code>eret</code>
Example	<code>eret</code>
Semantics	$PC = gpr[ea]$ $IE.IE = IE.EIE$
Result	
Issue	3 cycles
See Also	<code>scall</code> , <code>system call</code>

## lb

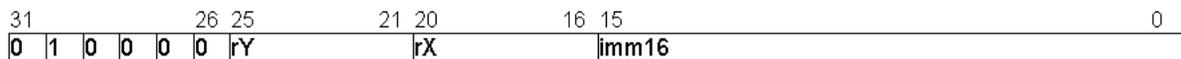
**Figure 59: lb Instruction**



**Table 52: lb Instruction Features**

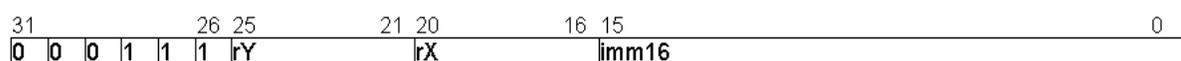
Feature	Description
Operation	Load byte from memory
Description	Loads a byte from memory at the address specified by the sum of the value in rY added to the sign-extended immediate, storing the sign-extended result into rX.
Syntax	<code>lb rX, (rY+imm16)</code>
Example	<code>lb r4, (r2+5)</code>
Semantics	<code>address = gpr[rY] + sign_extend(imm16)</code> <code>gpr[rX] = sign_extend(memory[address])</code>
Result	3 cycles
Issue	1 cycle
See Also	lbu, load byte, unsigned; lh, load half-word, signed; lhu, load half-word, unsigned; lw, load word

## lbu

**Figure 60: lbu Instruction****Table 53: lbu Instruction Features**

Feature	Description
Operation	Load unsigned byte from memory
Description	Loads a byte from memory at the address specified by the sum of the value in rY added to the sign-extended immediate, storing the zero-extended result into rX.
Syntax	<code>lbu rX, (rY+imm16)</code>
Example	<code>lbu r4, (r2+5)</code>
Semantics	<code>address = gpr[rY] + sign_extend(imm16)</code> <code>gpr[rX] = zero_extend(memory[address])</code>
Result	3 cycles
Issue	1 cycle
See Also	lb, load byte, signed; lh, load half-word, signed; lhu, load half-word, unsigned; lw, load word

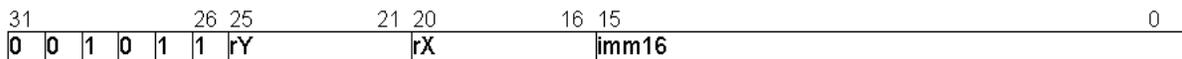
## lh

**Figure 61: lh Instruction**

**Table 54: lh Instruction Features**

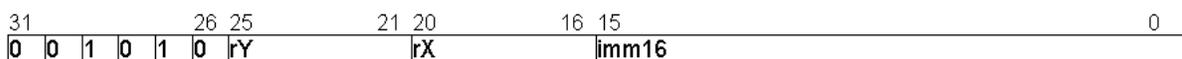
Feature	Description
Operation	Load half-word from memory
Description	Loads a half-word from memory at the address specified by the sum of the value in rY added to the sign-extended immediate, storing the sign-extended result into rX.
Syntax	lh rX, (rY+imm16)
Example	lh r4, (r2+6)
Semantics	<pre>address = gpr[rY] + sign_extend(imm16) gpr[rX] = sign_extend((memory[address] &lt;&lt; 8)                       (memory[address+1]))</pre>
Result	3 cycles
Issue	1 cycle
See Also	lb, load byte, signed; lbu, load byte, unsigned; lhu, load half-word, unsigned; lw, load word

## lhu

**Figure 62: lhu Instruction****Table 55: lhu Instruction Features**

Feature	Description
Operation	Load unsigned half-word from memory
Description	Loads a half-word from memory at the address specified by the sum of the value in rY added to the sign-extended immediate, storing the zero-extended result into rX.
Syntax	lhu rX, (rY+imm16)
Example	lhu r4, (r2+6)
Semantics	<pre>address = gpr[rY] + sign_extend(imm16) gpr[rX] = zero_extend((memory[address] &lt;&lt; 8)                       (memory[address+1]))</pre>
Result	3 cycles
Issue	1 cycle
See Also	lb, load byte, signed; lbu, load byte, unsigned; lh, load half-word, signed; lw, load word

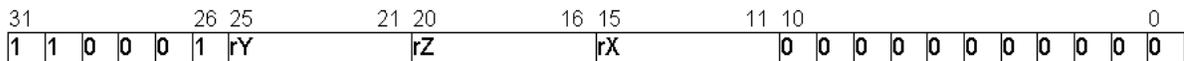
## lw

**Figure 63: lw Instruction**

**Table 56: lw Instruction Features**

Feature	Description
Operation	Load word from memory
Description	Loads a word from memory at address specified by the sum of the value in rY added to the sign-extended immediate, storing the result in rX.
Syntax	<code>lw rX, (rY+imm16)</code>
Example	<code>lw r4, (r2+8)</code>
Semantics	<pre> address = gpr[rY] + sign_extend(imm16) gpr[rX] = (memory[address] &lt;&lt; 24)             (memory[address+1] &lt;&lt; 16)             (memory[address+2] &lt;&lt; 8)             (memory[address+3]) </pre>
Result	3 cycles
Issue	1 cycle
See Also	lb, load byte, signed; lbu, load byte, unsigned; lh, load half-word, signed; lhu, load half-word, unsigned

## modu

**Figure 64: modu Instruction****Table 57: modu Instruction Features**

Feature	Description
Operation	Unsigned integer modulus
Description	Divides the value in rY by the value in rZ, storing the remainder in rX. Both operands are treated as unsigned integers.  Available only if the processor was configured with the DIVIDE_ENABLED option.
Syntax	<code>modu rX, rY, rZ</code>
Example	<code>modu r14, r15, r17</code>
Semantics	<code>gpr[rX] = gpr[rY] % gpr[rZ]</code>
Result	34 cycles
Issue	34 cycles
See Also	divu, divide

## mul

Figure 65: mul Instruction

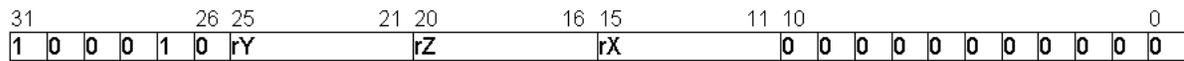


Table 58: mul Instruction Features

Feature	Description
Operation	Integer multiply
Description	Multiplies the value in rY by the value in rZ, storing the low 32 bits of the product in rX.  Available only if the processor was configured with either the MC_MULTIPLY_ENABLED or PL_MULTIPLY_ENABLED option.
Syntax	mul rX, rY, rZ
Example	mul r14, r15, r17
Semantics	$\text{gpr}[rX] = \text{gpr}[rY] * \text{gpr}[rZ]$
Result	3 cycles
Issue	1 cycle
See Also	muli, multiply with immediate

## muli

Figure 66: muli Instruction

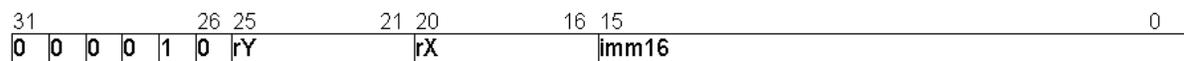


Table 59: muli Instruction Features

Feature	Description
Operation	Integer multiply
Description	Multiplies the value in rY by the sign-extended immediate, storing the low 32 bits of the product in rX.  Available only if the processor was configured with either the MC_MULTIPLY_ENABLED or PL_MULTIPLY_ENABLED option.
Syntax	muli rX, rY, imm16
Example	muli r4, r2, 0x5555
Semantics	$\text{gpr}[rX] = \text{gpr}[rY] * \text{sign\_extend}(\text{imm16})$
Result	3 cycles
Issue	1 cycle
See Also	mul, multiply between registers

## mv

Feature	Description
Operation	Move
Description	Moves the value in rY to rX. This is a pseudo-instruction implemented with: <code>or rX, rY, r0</code> .
Syntax	<code>mv rX, rY</code>
Example	<code>mv r4, r2</code>
Semantics	<code>gpr[rX] = gpr[rY]   gpr[r0]</code>
Result	1 cycle
Issue	1 cycle
See Also	<code>mvhi</code> , move immediate into high 16 bits

## mvhi

Feature	Description
Operation	Move high 16 bits
Description	Moves the 16-bit, left-shifted immediate into rX. This is a pseudo-instruction implemented with: <code>orhi rX, r0, imm16</code> .
Syntax	<code>mvhi rX, imm16</code>
Example	<code>mvhi r4, 0x5555</code>
Semantics	<code>gpr[rX] = gpr[r0]   (imm16 &lt;&lt; 16)</code>
Result	1 cycle
Issue	1 cycle
See Also	<code>mv</code> , move between registers

## nor

Figure 67: nor Instruction

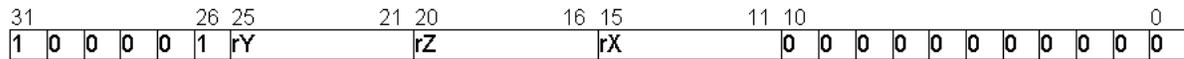


Table 60: nor Instruction Features

Feature	Description
Operation	Bitwise logical NOR
Description	Bitwise NOR of the value in rY with the value in rZ, storing the result in rX.
Syntax	<code>nor rX, rY, rZ</code>
Example	<code>nor r14, r15, r17</code>
Semantics	$\text{gpr}[rX] = \sim(\text{gpr}[rY] \mid \text{gpr}[rZ])$
Result	1 cycle
Issue	1 cycle
See Also	nori, NOR with immediate

## nori

Figure 68: nori Instruction

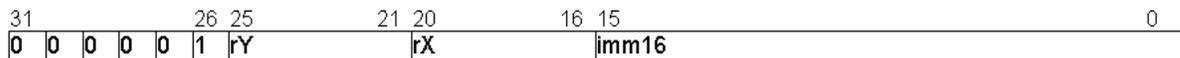


Table 61: nori Instruction Features

Feature	Description
Operation	Bitwise logical NOR
Description	Bitwise NOR of the value in rY with the zero-extended immediate, storing the result in rX.
Syntax	<code>nori rX, rY, imm16</code>
Example	<code>nori r4, r2, 0x5555</code>
Semantics	$\text{gpr}[rX] = \sim(\text{gpr}[rY] \mid \text{zero\_extend}(\text{imm16}))$
Result	1 cycle
Issue	1 cycle
See Also	nor, NOR between registers

## not

Feature	Description
Operation	Bitwise complement
Description	Bitwise complement of the value in rY, storing the result in rX. This is a pseudo-instruction implemented with: <code>xnor rX, rY, r0</code> .
Syntax	<code>not rX, rY</code>
Example	<code>not r4, r2</code>
Semantics	$\text{gpr}[rX] = \sim(\text{gpr}[rY] \wedge \text{gpr}[r0])$
Result	1 cycle
Issue	1 cycle

## or

Figure 69: or Instruction

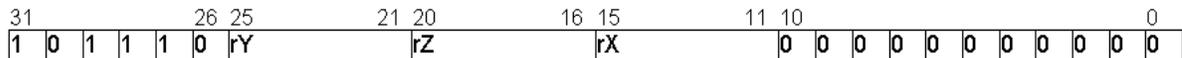


Table 62: or Instruction Features

Feature	Description
Operation	Bitwise logical OR
Description	Bitwise OR of the value in rY with the value in rZ, storing the result in rX.
Syntax	<code>or rX, rY, rZ</code>
Example	<code>or r14, r15, r17</code>
Semantics	$\text{gpr}[rX] = \text{gpr}[rY] \mid \text{gpr}[rZ]$
Result	1 cycle
Issue	1 cycle
See Also	<code>ori</code> , OR with immediate; <code>orhi</code> , OR with high 16 bits

## ori

Figure 70: ori Instruction

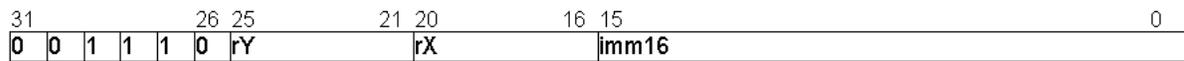


Table 63: ori Instruction Features

Feature	Description
Operation	Bitwise logical OR
Description	Bitwise OR of the value in rY with the zero-extended immediate, storing the result in rX.
Syntax	<code>ori rX, rY, imm16</code>
Example	<code>ori r4, r2, 0x5555</code>
Semantics	<code>gpr[rX] = gpr[rY]   zero_extend(imm16)</code>
Result	1 cycle
Issue	1 cycle
See Also	or, OR between registers; orhi, OR with high 16 bits

## orhi

Figure 71: orhi Instruction

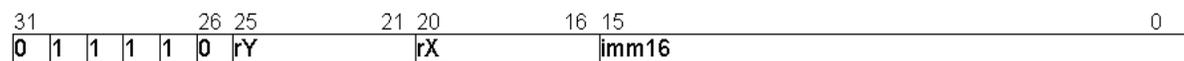


Table 64: orhi Instruction Features

Feature	Description
Operation	Bitwise logical OR (high 16-bits)
Description	Bitwise OR of the value in rY with the 16-bit, left-shifted immediate, storing the result in rX.
Syntax	<code>orhi rX, rY, imm16</code>
Example	<code>orhi r4, r2, 0x5555</code>
Semantics	<code>gpr[rX] = gpr[rY]   (imm16 &lt;&lt; 16)</code>
Result	1 cycle
Issue	1 cycle
See Also	or, OR between registers; ori, OR with immediate

## rcsr

Figure 72: rcsr Instruction

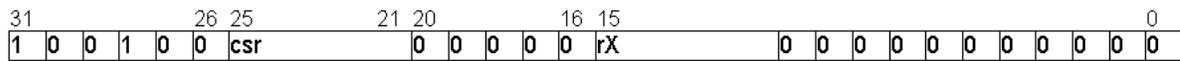


Table 65: rcsr Instruction Features

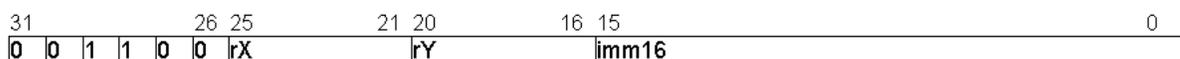
Feature	Description
Operation	Read control and status register
Description	Reads the value of the specified control and status register and stores it in rX.
Syntax	<code>rcsr rX, csr</code>
Example	<code>rcsr r15, IM</code>
Semantics	<code>gpr[rX] = csr</code>
Result	1 cycle
Issue	1 cycle
See Also	wcsr, write control and status register

## ret

Feature	Description
Operation	Return from function call
Description	Unconditional branch to address in ra. This is a pseudo-instruction implemented with: <code>b ra</code> .
Syntax	<code>ret</code>
Example	<code>ret</code>
Semantics	<code>PC = gpr[ra]</code>
Result	
Issue	4 cycles
See Also	call, function call from register; calli, function call with immediate

## sb

Figure 73: sb Instruction

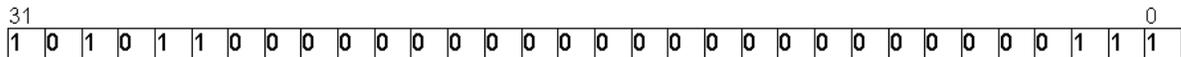


**Table 66: sb Instruction Features**

Feature	Description
Operation	Store byte to memory
Description	Stores the lower byte in rY into memory at the address specified by the sum of the value in rX added to the sign-extended immediate.
Syntax	<code>sb(rX+imm16), rY</code>
Example	<code>sb(r2+8), r4</code>
Semantics	<code>address = gpr[rX] + sign_extend(imm16)</code> <code>memory[address] = gpr[rY] &amp; 0xff</code>
Result	
Issue	1 cycle
See Also	sh, store half-word; sw, store word

## scall

**Figure 74: scall Instruction**

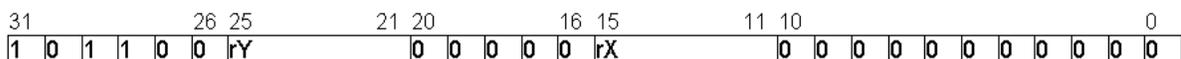


**Table 67: scall Instruction Features**

Feature	Description
Operation	System call
Description	Raises a system call exception.
Syntax	<code>scall</code>
Example	<code>scall</code>
Semantics	<code>gpr[ea] = PC</code> <code>IE.EIE = IE.IE</code> <code>IE.IE = 0</code> <code>PC = (DC.RE ? DEBA : EBA) + ID * 32</code>
Result	
Issue	4 cycles
See Also	eret, return from exception

## sxtb

**Figure 75: sxtb Instruction**





## sh

Figure 77: sh Instruction

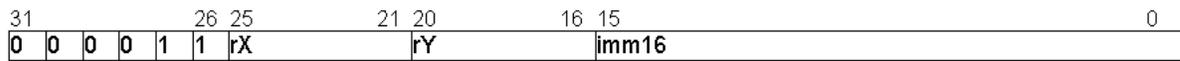


Table 70: sh Instruction Features

Feature	Description
Operation	Store half-word to memory
Description	Stores the lower half-word in rY into memory at the address specified by the sum of the value in rX added to the sign-extended immediate.
Syntax	sh (rX+imm16), rY
Example	sh (r2+8), r4
Semantics	$address = gpr[rX] + sign\_extend(imm16)$ $memory[address] = gpr[rY] \& 0xff$ $memory[address+1] = (gpr[rY] \gg 8) \& 0xff$
Result	
Issue	1 cycle
See Also	sb, store byte; sw, store word

## sl

Figure 78: sl Instruction

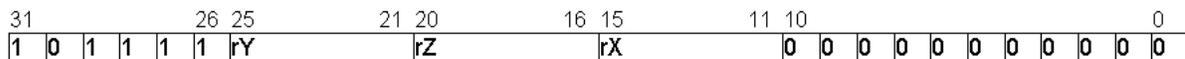


Table 71: sl Instruction Features

Feature	Description
Operation	Shift left
Description	Shifts the value in rY left by the number of bits specified by the value in rZ, storing the result in rX.  Available only if the processor was configured with either the MC_BARREL_SHIFT_ENABLED or PL_BARREL_SHIFT_ENABLED option.
Syntax	sl rX, rY, rZ
Example	sl r14, r15, r17
Semantics	$gpr[rX] = gpr[rY] \ll (gpr[rZ] \& 0x1f)$
Result	2 cycles
Issue	1 cycle
See Also	sli, shift left with immediate

## sli

Figure 79: sli Instruction

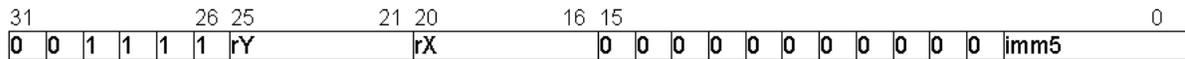


Table 72: sli Instruction Features

Feature	Description
Operation	Shift left
Description	Shifts the value in rY left by the number of bits specified by the immediate, storing the result in rX.  Available only if the processor was configured with either the MC_BARREL_SHIFT_ENABLED or PL_BARREL_SHIFT_ENABLED option.
Syntax	sli rX, rY, imm5
Example	sli r4, r2, 17
Semantics	$\text{gpr}[\text{rX}] = \text{gpr}[\text{rY}] \ll \text{imm5}$
Result	2 cycles
Issue	1 cycle
See Also	sl, shift left from register

## sr

Figure 80: sr Instruction

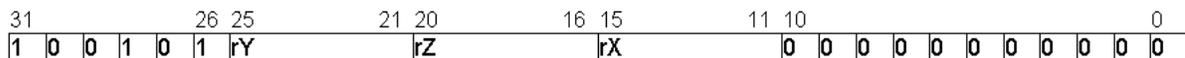


Table 73: sr Instruction Features

Feature	Description
Operation	Shift right (arithmetic)
Description	Shifts the signed value in rY right by the number of bits specified by the value in rZ, storing the result in rX.  Available only if the processor was configured with either the MC_BARREL_SHIFT_ENABLED or PL_BARREL_SHIFT_ENABLED option.
Syntax	sr rX, rY, rZ
Example	sr r14, r15, r17
Semantics	$\text{gpr}[\text{rX}] = \text{gpr}[\text{rY}] \gg (\text{gpr}[\text{rZ}] \& 0x1f)$
Result	2 cycles
Issue	1 cycle
See Also	sri, shift right with immediate; sru, shift right, unsigned; srui, shift right with immediate, unsigned

## sri

Figure 81: sri Instruction

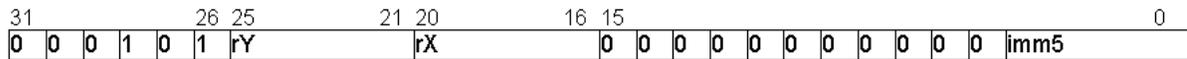


Table 74: sri Instruction Features

Feature	Description
Operation	Shift right (arithmetic)
Description	Shifts the signed value in rY right by the number of bits specified by the immediate, storing the result in rX.  Available only if the processor was configured with either the MC_BARREL_SHIFT_ENABLED or PL_BARREL_SHIFT_ENABLED option.
Syntax	sri rX, rY, imm5
Example	sri r4, r2, 12
Semantics	$\text{gpr}[rX] = \text{gpr}[rY] \gg \text{imm5}$
Result	2 cycles
Issue	1 cycle
See Also	sr, shift right from register; sru, shift right, unsigned; srui, shift right with immediate, unsigned

## sru

Figure 82: sru Instruction

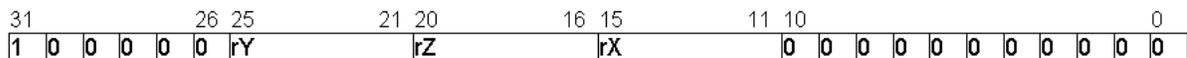


Table 75: sru Instruction Features

Feature	Description
Operation	Shift right, unsigned (logical)
Description	Shifts the unsigned value in rY right by the number of bits specified by the value in rZ, storing the result in rX.  Available only if the processor was configured with either the MC_BARREL_SHIFT_ENABLED or PL_BARREL_SHIFT_ENABLED option.
Syntax	sru rX, rY, rZ
Example	sru r14, r15, r17
Semantics	$\text{gpr}[rX] = \text{gpr}[rY] \gg (\text{gpr}[rZ] \& 0x1f)$
Result	2 cycles
Issue	1 cycle
See Also	sr, shift right from register; sri, shift right with immediate; srui, shift right with immediate, unsigned

## srui

Figure 83: srui Instruction

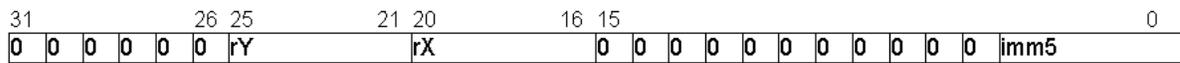


Table 76: srui Instruction Features

Feature	Description
Operation	Shifts right, unsigned (logical)
Description	Shifts the unsigned value in rY right by the number of bits specified by the immediate, storing the result in rX.  Available only if the processor was configured with either the MC_BARREL_SHIFT_ENABLED or PL_BARREL_SHIFT_ENABLED option.
Syntax	srui rX, rY, imm5
Example	srui r4, r2, 5
Semantics	$gpr[rX] = gpr[rY] \gg imm5$
Result	2 cycles
Issue	1 cycle
See Also	sr, shift right from register; sri, shift right with immediate; sru, shift right, unsigned

## sub

Figure 84: sub Instruction

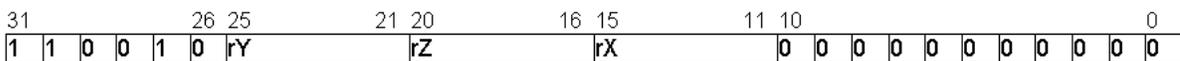


Table 77: sub Instruction Features

Feature	Description
Operation	Integer subtraction
Description	Subtracts the value in rZ from the value in rY, storing the result in rX.
Syntax	sub rX, rY, rZ
Example	sub r14, r15, r17
Semantics	$gpr[rX] = gpr[rY] - gpr[rZ]$
Result	1 cycle
Issue	1 cycle
See Also	addi, add with signed immediate



## xnor

Figure 87: xnor Instruction

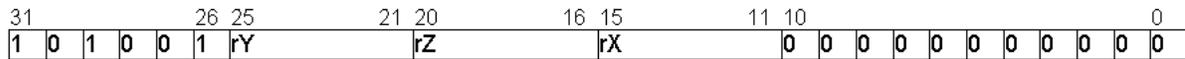


Table 80: xnor Instruction Features

Feature	Description
Operation	Bitwise logical exclusive-NOR
Description	Bitwise exclusive-NOR of the value in rY with the value in rZ, storing the result in rX.
Syntax	<code>xnor rX, rY, rZ</code>
Example	<code>xnor r14, r15, r17</code>
Semantics	$\text{gpr}[rX] = \sim(\text{gpr}[rY] \wedge \text{gpr}[rZ])$
Result	1 cycle
Issue	1 cycle
See Also	xnori, XNOR with immediate

## xnori

Figure 88: xnori Instruction

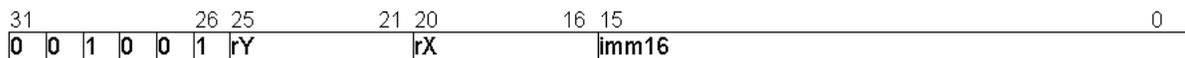


Table 81: xnori Instruction Features

Feature	Description
Operation	Bitwise logical exclusive-NOR
Description	Bitwise exclusive-NOR of the value in rY with the zero-extended immediate, storing the result in rX.
Syntax	<code>xnori rX, rY, imm16</code>
Example	<code>xnori r4, r2, 0x5555</code>
Semantics	$\text{gpr}[rX] = \sim(\text{gpr}[rY] \wedge \text{zero\_extend}(\text{imm16}))$
Result	1 cycle
Issue	1 cycle
See Also	xnor, XNOR between registers

## XOR

Figure 89: xor Instruction

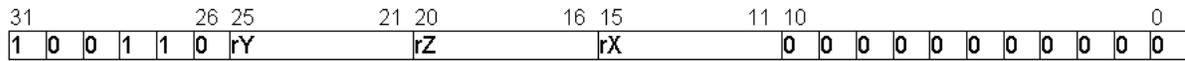


Table 82: xor Instruction Features

Feature	Description
Operation	Bitwise logical exclusive-OR
Description	Bitwise exclusive-OR of the value in rY with the value in rZ, storing the result in rX.
Syntax	<code>xor rX, rY, rZ</code>
Example	<code>xor r14, r15, r17</code>
Semantics	$\text{gpr}[rX] = \text{gpr}[rY] \wedge \text{gpr}[rZ]$
Result	1 cycle
Issue	1 cycle
See Also	xori, XOR with immediate

## xori

Figure 90: xori Instruction

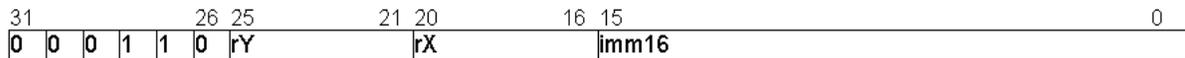


Table 83: xori Instruction Features

Feature	Description
Operation	Bitwise logical exclusive-OR
Description	Bitwise exclusive-OR of the value in rY with the zero-extended immediate, storing the result in rX.
Syntax	<code>xori rX, rY, imm16</code>
Example	<code>xori r4, r2, 0x5555</code>
Semantics	$\text{gpr}[rX] = \text{gpr}[rY] \wedge \text{zero\_extend}(\text{imm16})$
Result	1 cycle
Issue	1 cycle
See Also	xori, XOR between registers



# Index

## A

A field **30**  
ACK **20**  
ACK\_I **43**  
ACK\_O **45**  
Acknowledge Input **43**  
Acknowledge Output **45**  
add instruction **52**  
addi instruction **52**  
address alignment **15**  
Address Input array **44**  
Address Output array **43**  
Address pipeline stage **5**  
address space **14**  
ADR\_I() **44**  
ADR\_O() **43**  
and instruction **53**  
andhi instruction **53**  
andi instruction **54**  
arbitration schemes **46**  
arithmetic instructions **31**  
associativity in caches **16**

## B

b instruction **54**  
ba register **7**  
be instruction **55**  
bg instruction **55**  
bge instruction **56**  
bgeu instruction **56**  
bgu instruction **57**  
bi instruction **57**  
BIE field **10**  
big-endian **14**  
bne instruction **58**

BP field **12**  
BP registers **27, 28, 30**  
break instruction **58**  
breakpoint address register **7**  
breakpoint exceptions **20, 21, 23, 27**  
breakpoint registers **27, 28, 30**  
bret instruction **59**  
BTE\_I() **41**  
BTE\_O() **41**  
burst type extension **41**  
bypass in pipeline **5**

## C

cache configurations **17, 38**  
cacheable addresses **14**  
caches **16, 38**  
call instruction **59**  
calli instruction **60**  
CC field **12**  
CC register **9, 12**  
CFG register **9, 12**  
cmpe instruction **60**  
cmpei instruction **61**  
cmpg instruction **61**  
cmpge instruction **62**  
cmpgei instruction **63**  
cmpgeu instruction **63**  
cmpgeui instruction **64**  
cmpgi instruction **62**  
cmpgu instruction **64**  
cmpgui instruction **65**  
cmpne instruction **65**  
cmpnei instruction **66**  
comparison instructions **32**  
component signals

- introduction to **42**
- master signals **43**
- slave signals **44**
- configuration options **35**
- configuration register **9, 12**
- control and status registers
  - configuration **9, 12**
  - cycle counter **9, 12**
  - data cache control **9, 12**
  - exception base address **9, 13, 26**
  - instruction cache control **9, 11**
  - interrupt enable **9, 10**
  - interrupt mask **9, 11**
  - interrupt pending **9, 11**
  - introduction **9**
  - program counter **9**
- CR format for instructions **49**
- CTI\_I() **40**
- CTI\_O() **40**
- CYC\_I **45**
- CYC\_O **44**
- cycle counter **9, 12**
- Cycle Input **45**
- Cycle Output **44**
- cycle type identifier **40**
- CYCLE\_COUNTER\_ENABLED **36**

## D

- D field **12**
- DAT\_I() **43, 44**
- DAT\_O() **43, 44**
- data cache control **9, 12**
- Data Input array **43, 44**
- Data Output array **43, 44**
- data transfer instructions **32**
- data types **6**
- DataBusError **20, 21, 23**
- DC field **12**
- DC register **28, 29**
- DCACHE\_ASSOCIATIVITY **37**
- DCACHE\_BASE\_ADDRESS **37**
- DCACHE\_BYTES\_PER\_LINE **37**
- DCACHE\_ENABLED **37**
- DCACHE\_LIMIT **37**
- DCACHE\_SETS **37**
- DCC register **9, 12**
- DEBA register **28, 29**
- debug **28, 38**
- debug control **28, 29**
- debug control and status registers **28**
- debug exception base aAddress **28**
- debug exception base address **29**
- debug exceptions **21**
- debug monitor **22**
- DEBUG\_ENABLED **36, 38**
- decode pipeline stage **5**
- DIVIDE\_ENABLED **35**
- DivideByZero **20, 24**

- divu instruction **66**

## E

- E field **30**
- ea register **7**
- EBA register **9, 13, 26**
- EBR **38**
- EIE field **10**
- embedded block RAM **38**
- endianness **14**
- eret instruction **67**
- ERR\_I **43**
- ERR\_O **45**
- Error Input **43**
- Error Output **45**
- exception address register **7**
- exception base address **9, 13, 26**
- exceptions
  - breakpoints **27**
  - debug **21**
  - handlers **22**
  - interrupts **22, 26**
  - introduction to **20**
  - nested **26**
  - non-debug **21**
  - processing **21**
  - reset **26**
  - watchpoints **27**
- execute pipeline stage **5**
- extended data types **7**

## F

- F field in JTAG UART Receive Register **30**
- F field in JTAG UART Transmit Register **29**
- fetch pipeline stage **5**
- fixed slave-side arbitration scheme **47**
- fp register **9, 15**
- frame pointer **9, 15**

## G

- G field **12**
- general-purpose registers **7**
- global pointer **8**
- gp register **8**

## H

- H field **12**
- handlers, exceptions **22**

## I

- I bit in data cache control **12**
- I bit in instruction cache control **11**
- I format for instructions **49**
- IC field **12**
- ICACHE\_ASSOCIATIVITY **37**
- ICACHE\_BASE\_ADDRESS **37**
- ICACHE\_BYTES\_PER\_LINE **37**
- ICACHE\_ENABLED **36**

- ICACHE\_LIMIT **37**
  - ICACHE\_SETS **37**
  - ICC register **9, 11**
  - IE field **10**
  - IE register **9, 10**
  - IM register **9, 11**
  - initializing caches **17**
  - inline memories **18**
  - instruction cache control **9, 11**
  - instruction set
    - categories **31**
    - descriptions **51**
      - add **52**
      - addi **52**
      - and **53**
      - andhi **53**
      - andi **54**
      - b **54**
      - be **55**
      - bg **55**
      - bge **56**
      - bgeu **56**
      - bgu **57**
      - bi **57**
      - bne **58**
      - break **58**
      - bret **59**
      - call **59**
      - calli **60**
      - cmpe **60**
      - cmpei **61**
      - cmpg **61**
      - cmpge **62**
      - cmpgei **63**
      - cmpgeu **63**
      - cmpgeui **64**
      - cmpgi **62**
      - cmpgu **64**
      - cmpgui **65**
      - cmpne **65**
      - cmpnei **66**
      - divu **66**
      - eret **67**
      - lb **67**
      - lbu **68**
      - lhu **68, 69**
      - lw **69**
      - modu **70**
      - mul **71**
      - muli **71**
      - mv **72**
      - mvh **72**
      - nor **73**
      - nori **73**
      - not **74**
      - or **74**
      - orhi **75**
      - ori **75**
      - rcsr **76**
      - ret **76**
      - sb **76**
      - scall **77**
      - sxtb **77**
      - sxth **78**
      - sh **79**
      - sl **79**
      - sli **80**
      - sr **80**
      - sri **81**
      - sru **81**
      - sru\_i **82**
      - sub **82**
      - sw **83**
      - wcsr **83**
      - xnor **84**
      - xnori **84**
      - xor **85**
      - xori **85**
    - formats **49**
    - opcodes **50**
    - pseudo-instructions **51**
  - InstructionBusError **20, 23**
  - INT field **12**
  - interconnect architecture *see* WISHBONE
    - interconnect
  - interlock in pipeline **5**
  - Interrupt **21, 22, 24, 26**
  - interrupt enable **10**
  - interrupt mask **9, 11**
  - interrupt pending **9, 11**
  - interrupt renable **9**
  - invalidating caches **17**
  - IP register **9, 11**
- J**
- J field **12**
  - JRX register **28, 30**
  - JTAG UART Receive Register **28, 30**
  - JTAG UART Transmit Register **29**
  - JTAG UART transmit register **28**
  - JTX register **28, 29**
- L**
- lb instruction **67**
  - lbu instruction **68**
  - lh instruction **68**
  - lhu instruction **69**
  - lines in caches **16**
  - Lock Input **45**
  - Lock Output **44**
  - LOCK\_I **45**
  - LOCK\_O **44**
  - logic instructions **31**
  - lw instruction **69**
- M**
- M field **12**

- master signals **43**
  - memory architecture
    - address alignment **15**
    - address space **14**
    - cacheable addresses **14**
    - endianness **14**
    - exceptions **20**
    - stack layout **15**
  - memory pipeline stage **5**
  - model, programmer's **5**
  - modu instruction **70**
  - monitor, debug **22**
  - mul instruction **71**
  - muli instruction **71**
  - mv instruction **72**
  - mvh instruction **72**
- N**
- nested exceptions **26**
  - non-debug exceptions **21**
  - nor instruction **73**
  - nori instruction **73**
  - not instruction **74**
- O**
- opcodes **50**
  - OPENCORES.ORG **39**
  - or instruction **74**
  - orhi instruction **75**
  - ori instruction **75**
- P**
- PC register **9**
  - pipeline **5, 21**
  - processing exceptions **21**
  - program counter **9**
  - program flow control instructions **32**
  - programmer's model **5**
  - pseudo-instructions **51**
- R**
- R field **12**
  - r0 register **7**
  - ra register **7**
  - rcsr instruction **76**
  - read-miss **16**
  - registered feedback mode **40**
  - registers
    - control and status **9**
    - debug control and status **28**
    - general-purpose **7**
  - reset **20, 23, 26**
  - resources **38**
  - ret instruction **76**
  - Retry Input **43**
  - Retry Output **45**
  - return address register **7**
  - REV field **12**
  - RI format for instructions **49**
  - round-robin slave-side arbitration scheme **47**
  - RR format for instructions **49**
  - RTY\_I **43**
  - RTY\_O **45**
  - RXD field **30**
- S**
- S field **12**
  - sb instruction **76**
  - scall instruction **77**
  - SEL\_I() **44**
  - SEL\_O() **43**
  - Select Input array **44**
  - Select Output array **43**
  - sextb instruction **77**
  - sexth instruction **78**
  - sh instruction **79**
  - shared-bus arbitration scheme **46**
  - shift instructions **32**
  - SIGN\_EXTEND\_ENABLED **36**
  - signals, component **42**
  - sl instruction **79**
  - slave signals **44**
  - slave-side arbitration scheme **46**
    - fixed **47**
    - round-robin **47**
  - sli instruction **80**
  - SoC Interconnection Architecture for Portable IP Cores *see* WISHBONE interconnect
  - sp register **9, 15**
  - sr instruction **80**
  - sri instruction **81**
  - sru instruction **81**
  - srui instruction **82**
  - stack layout **15**
  - stack pointer **9, 15**
  - stages in pipeline **5**
  - STB\_I **45**
  - STB\_O **44**
  - Strobe Input **45**
  - Strobe Output **44**
  - sub instruction **82**
  - sw instruction **83**
  - SystemCall **21, 24**
  - System-on-Chip Interconnection Architecture for Portable IP Cores *see* WISHBONE interconnect
- T**
- TXD field **29**
- W**
- watchpoint exceptions **20, 21, 23, 27**
  - Watchpoint registers **27, 31**
  - watchpoint registers **28**
  - ways in caches **16**
  - wcsr instruction **83**

WE\_I **45**  
WE\_O **43**  
WISHBONE interconnect **39**  
    component signals **42**  
    introduction to **39**  
    master signals **43**  
    registered feedback mode **40**  
    slave signals **44**  
WP field **12**  
WP registers **27, 28, 31**  
Write Enable Input **45**  
Write Enable Output **43**  
writeback pipeline stage **5**

**X**

X field **12**  
xnor instruction **84**  
xnori instruction **84**  
xor instruction **85**  
xori instruction **85**