Introduction

This chapter describes the Nios® II programming model, covering processor features at the assembly language level. Fully understanding the contents of this chapter requires prior knowledge of computer architecture, operating systems, virtual memory and memory management, software processes and process management, exception handling, and instruction sets. This chapter assumes you have a detailed understanding of the aforementioned concepts and focuses on how these concepts are specifically implemented in the Nios II processor. Where possible, this chapter uses industry-standard terminology.

This chapter discusses the following topics from the system programmer's perspective:

- Operating modes, page 3–2—Defines the relationships between executable code and memory.
- Memory management unit (MMU), page 3–3—Describes virtual memory support for full-featured operating systems.
- Memory protection unit (MPU), page 3–8—Describes memory protection without virtual memory management.
- Registers, page 3–10—Describes the Nios II register sets.
- Working With the MPU, page 3–28—Provides an overview of MPU initialization and operation.
- Exception processing, page 3–30—Describes how the Nios II processor responds to exceptions.
- Memory and Peripheral Access, page 3–51—Describes Nios II addressing.
- Instruction set categories, page 3–53—Introduces the Nios II instruction set.

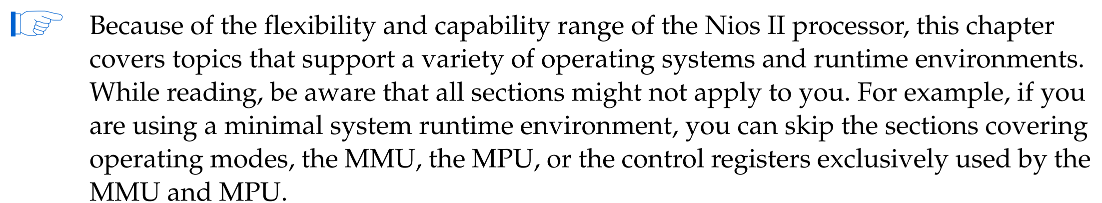

Because of the flexibility and capability range of the Nios II processor, this chapter covers topics that support a variety of operating systems and runtime environments. While reading, be aware that all sections might not apply to you. For example, if you are using a minimal system runtime environment, you can skip the sections covering operating modes, the MMU, the MPU, or the control registers exclusively used by the MMU and MPU.

High-level software development tools are not discussed here. Refer to the *Nios II Software Developer's Handbook* for information about developing software.

Operating Modes

Operating modes control how the processor operates, manages system memory, and accesses peripherals. The Nios II architecture supports these operating modes:

- Supervisor mode
- User mode

The following sections define the modes, their relationship to your system software and application code, and their relationship to the Nios II MMU and Nios II MPU. Refer to "Memory Management Unit" on page 3–3 for more information about the Nios II MMU. Refer to "Memory Protection Unit" on page 3–8 for more information about the Nios II MPU.

Supervisor Mode

Supervisor mode allows unrestricted operation of the processor. All code has access to all processor instructions and resources. The processor may perform any operation the Nios II architecture provides. Any instruction may be executed, any I/O operation may be initiated, and any area of memory may be accessed.

Operating systems and other system software run in supervisor mode. In systems with an MMU, application code runs in user mode, and the operating system, running in supervisor mode, controls the application's access to memory and peripherals. In systems with an MPU, your system software controls the mode in which your application code runs. In Nios II systems without an MMU or MPU, all application and system code runs in supervisor mode.

Code that needs direct access to and control of the processor runs in supervisor mode. For example, the processor enters supervisor mode whenever a processor exception (including processor reset or break) occurs. Software debugging tools also use supervisor mode to implement features such as breakpoints and watchpoints.

For systems without an MMU or MPU, all code runs in supervisor mode.

User Mode

User mode is available only when the Nios II processor in your hardware design includes an MMU or MPU. User mode exists solely to support operating systems. Operating systems (that make use of the processor's user mode) run your application code in user mode. The user mode capabilities of the processor are a subset of the supervisor mode capabilities. Only a subset of the instruction set is available in user mode.

The operating system determines which memory addresses are accessible to user mode applications. Attempts by user mode applications to access memory locations without user access enabled are not permitted and cause an exception. Code running in user mode uses system calls to make requests to the operating system to perform I/O operations, manage memory, and access other system functionality in the supervisor memory.

The Nios II MMU statically divides the 32-bit virtual address space into user and supervisor partitions. Refer to "Address Space and Memory Partitions" on page 3–4 for more information about the MMU memory partitions. The MMU provides operating systems access permissions on a per-page basis. Refer to "Virtual Addressing" on page 3–3 for more information about MMU pages.

The Nios II MPU supervisor and user memory divisions are determined by the operating system or runtime environment. The MPU provides user access permissions on a region basis. Refer to "Memory Regions" on page 3–8 for more information about MPU regions.

Memory Management Unit

The Nios II processor provides an MMU to support full-featured operating systems. Operating systems that require virtual memory rely on an MMU to manage the virtual memory. When present, the MMU manages memory accesses including translation of virtual addresses to physical addresses, memory protection, cache control, and software process memory allocation.

Recommended Usage

Including the Nios II MMU in your Nios II hardware system is optional. The MMU is only useful with an operating system that takes advantage of it.

Many Nios II systems have simpler requirements where minimal system software or a small-footprint operating system (such as the Altera® hardware abstraction library (HAL) or a third party real-time operating system) is sufficient. Such software is unlikely to function correctly in a hardware system with an MMU-based Nios II processor. Do not include an MMU in your Nios II system unless your operating system requires it.

The Altera HAL and HAL-based real-time operating systems do not support the MMU.

If your system needs memory protection, but not virtual memory management, refer to "Memory Protection Unit" on page 3–8.

Memory Management

Memory management comprises two key functions:

- Virtual addressing—Mapping a virtual memory space into a physical memory space
- Memory protection—Allowing access only to certain memory under certain conditions

Virtual Addressing

A virtual address is the address that software uses. A physical address is the address which the hardware outputs on the address lines of the Avalon® bus. The Nios II MMU divides virtual memory into 4 KByte pages and physical memory into 4 KByte frames.

The MMU contains a hardware translation lookaside buffer (TLB). The operating system is responsible for creating and maintaining a page table (or equivalent data structures) in memory. The hardware TLB acts as a software managed cache for the page table. The MMU does not perform any operations on the page table, such as hardware table walks. Therefore the operating system is free to implement its page table in any appropriate manner.

Table 3–1 shows how the Nios II MMU divides up the virtual address. There is a 20 bit virtual page number (VPN) and a 12 bit page offset.

Table 3–1. MMU Virtual Address Fields

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Virtual Page Number																				Page Offset											

As input, the TLB takes a VPN plus a process identifier (to guarantee uniqueness). As output, the TLB provides the corresponding physical frame number (PFN).

Distinct processes can use the same virtual address space. The process identifier, concatenated with the virtual address, distinguishes identical virtual addresses in separate processes. To determine the physical address, the Nios II MMU translates a VPN to a PFN and then concatenates the PFN with the page offset. The bits in the page offset are not translated.

Memory Protection

The Nios II MMU maintains read, write, and execute permissions for each page. The TLB provides the permission information when translating a VPN. The operating system can control whether or not each process is allowed to read data from, write data to, or execute instructions on each particular page. The MMU also controls whether accesses to each data page are cacheable or uncacheable by default.

Whenever an instruction attempts to access a page that either has no TLB mapping, or lacks the appropriate permissions, the MMU generates an exception. The Nios II processor's precise exceptions enable the system software to update the TLB, and then re-execute the instruction if desired.

Address Space and Memory Partitions

The MMU provides a 4 GByte virtual address space, and is capable of addressing up to 4 GBytes of physical memory.

The amount of actual physical memory, determined by the configuration of your hardware system, might be less than the available 4 GBytes of physical address space.

Virtual Memory Address Space

The 4 GByte virtual memory space is divided into partitions. The upper 2 GBytes of memory is reserved for the operating system and the lower 2 GBytes is reserved for user processes. Table 3–2 names and describes the partitions.

© November 2009 Altera Corporation

Table 3–2. Virtual Memory Partitions

Partition	Virtual Address Range	Used By	Memory Access	User Mode Access	Default Data Cacheability
I/O *(1)*	`0xE0000000–0xFFFFFFFF`	Operating system	Bypasses TLB	No	Disabled
Kernel *(1)*	`0xC0000000–0xDFFFFFFF`	Operating system	Bypasses TLB	No	Enabled
Kernel MMU *(1)*	`0x80000000–0xBFFFFFFF`	Operating system	Uses TLB	No	Set by TLB
User	`0x00000000–0x7FFFFFFF`	User processes	Uses TLB	Set by TLB	Set by TLB

Note to Table 3–2:

(1) Supervisor-only partition

Each partition has a specific size, purpose, and relationship to the TLB:

- The 512 MByte I/O partition provides access to peripherals.
- The 512 MByte kernel partition provides space for the operating system kernel.
- The 1 GByte kernel MMU partition is used by the TLB miss handler and kernel processes.
- The 2 GByte user partition is used by application processes.

I/O and kernel partitions bypass the TLB. The kernel MMU and user partitions use the TLB. If all software runs in the kernel partition, the MMU is effectively disabled.

Physical Memory Address Space

The 4 GByte physical memory is divided into low memory and high memory. The lowest 0.5 GBytes of physical address space is low memory. The upper 3.5 GBytes of physical address space is high memory. Figure 3–1 shows how physical memory is divided.

Figure 3–1. Division of Physical Memory

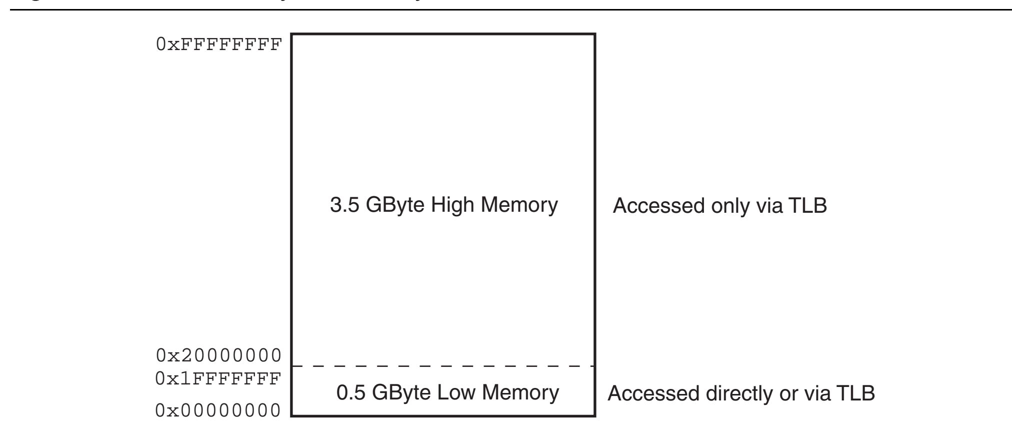

High physical memory can only be accessed through the TLB. Any physical address in low memory (29-bits or less) can be accessed through the TLB or by bypassing the TLB. When bypassing the TLB, a 29-bit physical address is computed by clearing the top three bits of the 32-bit virtual address.

To function correctly, the base physical address of all exception vectors (reset, general exception, break, and fast TLB miss) must point to low physical memory so that hardware can correctly map their virtual addresses into the kernel partition. This restriction is enforced by the Nios II Processor MegaWizard™ interface in SOPC Builder.

Data Cacheability

Each partition has a rule that determines the default data cacheability property of each memory access. When data cacheability is enabled on a partition of the address space, a data access to that partition can be cached, if a data cache is present in the system. When data cacheability is disabled, all access to that partition goes directly to the Avalon switch fabric. Bit 31 is not used to specify data cacheability, as it is in Nios II cores without MMUs. Virtual memory partitions that bypass the TLB have a default data cacheability property, as shown in Table 3–2. For partitions that are mapped through the TLB, data cacheability is controlled by the TLB on a per-page basis.

Non-I/O load and store instructions use the default data cacheability property. I/O load and store instructions are always noncacheable, so they ignore the default data cacheability property.

TLB Organization

A TLB functions as a cache for the operating system's page table. In Nios II processors with an MMU, one main TLB is shared by instruction and data accesses. The TLB is stored in on-chip RAM and handles translations for instruction fetches and instructions that perform data accesses.

The TLB is organized as an *n*-way set-associative cache. The software specifies the way (set) when loading a new entry.

You can configure the number of TLB entries and the number of ways (set associativity) of the TLB in SOPC Builder at system generation time. By default, the TLB is a 16-way cache. The default number of entries depends on the target device, as follows:

- Cyclone®, Cyclone II, Stratix®, Stratix II, Stratix II GX—128 entries, requiring one M4K RAM
- Cyclone III, Stratix III, Stratix IV—256 entries, requiring one M9K RAM

For further detail, refer to the *Instantiating the Nios II Processor in SOPC Builder* chapter of the *Nios II Processor Reference Handbook*.

The operating system software is responsible for guaranteeing that multiple TLB entries do not map the same virtual address. The hardware behavior is undefined when multiple entries map the same virtual address.

Each TLB entry consists of a tag and data portion. This is analogous to the tag and data portion of instruction and data caches.

Refer to the *Nios II Core Implementation Details* chapter of the *Nios II Processor Reference Handbook* for details on instruction and data caches.

The tag portion of a TLB entry contains information used when matching a virtual address to a TLB entry. Table 3–3 describes the tag portion of a TLB entry.

Table 3–3. TLB Tag Portion Contents

Field Name	Description
`VPN`	`VPN` is the virtual page number field. This field is compared with the top 20 bits of the virtual address.
`PID`	`PID` is the process identifier field. This field is compared with the value of the current process identifier stored in the `tlbmisc` control register, effectively extending the virtual address. The field size is configurable at system generation time, and can be between 8 and 14 bits.
`G`	`G` is the global flag. When `G` = 1, the `PID` is ignored in the TLB lookup.

The TLB data portion determines how to translate a matching virtual address to a physical address. Table 3–4 describes the data portion of a TLB entry.

Table 3–4. TLB Data Portion Contents

Field Name	Description
`PFN`	`PFN` is the physical frame number field. This field specifies the upper bits of the physical address. The size of this field depends on the range of physical addresses present in the system. The maximum size is 20 bits.
`C`	`C` is the cacheable flag. Determines the default data cacheability of a page. Can be overridden for data accesses using I/O load and store family of Nios II instructions.
`R`	`R` is the readable flag. Allows load instructions to read a page.
`W`	`W` is the writable flag. Allows store instructions to write a page.
`X`	`X` is the executable flag. Allows instruction fetches from a page.

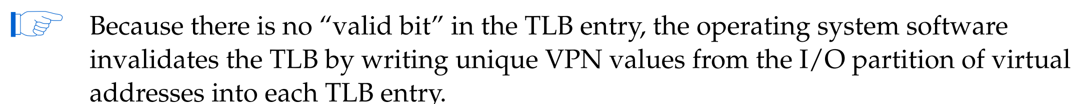
Because there is no “valid bit” in the TLB entry, the operating system software invalidates the TLB by writing unique VPN values from the I/O partition of virtual addresses into each TLB entry.

TLB Lookups

A TLB lookup attempts to convert a virtual address (VADDR) to a physical address (PADDR).

The TLB lookup algorithm for instruction fetches is shown in Figure 3–2. The TLB lookup algorithm for data accesses is shown in Figure 3–3.

Refer to “Instruction-Related Exceptions” on page 3–38 for details on TLB exceptions.

Figure 3–2. TLB Lookup Algorithm for Instruction Fetches

```
if (VPN match && (G == 1 || PID match))
   if (X == 1)
      PADDR = concat(PFN, VADDR[11:0])
   else
      take TLB permission violation exception
else
   if (EH bit of status register == 1)
      take double TLB miss exception
   else
      take fast TLB miss exception
```

Figure 3–3. TLB Lookup Algorithm for Data Access Operations

```
if (VPN match && (G == 1 || PID match))
   if ((load && R == 1) || (store && W == 1) || flushda)
      PADDR = concatenate(PFN, VADDR[11:0])
   else
      take TLB permission violation exception
else
   if (EH bit of status register == 1)
      take double TLB miss exception
   else
      take fast TLB miss exception
```

Memory Protection Unit

The Nios II processor provides an MPU for operating systems and runtime environments that desire memory protection but do not require virtual memory management. For information about memory protection with virtual memory management, refer to “Memory Management Unit” on page 3–3.

When present and enabled, the MPU monitors all Nios II instruction fetches and data memory accesses to protect against errant software execution. The MPU is a hardware facility that system software uses to define memory regions and their associated access permissions. The MPU triggers an exception if software attempts to access a memory region in violation of its permissions, allowing you to intervene and handle the exception as appropriate. The precise exception effectively prevents the illegal access to memory.

The MPU extends the Nios II processor to support user mode and supervisor mode. Typically, system software runs in supervisor mode and end-user applications run in user mode, although all software can run in supervisor mode if desired. System software defines which MPU regions belong to supervisor mode and which belong to user mode.

Memory Regions

The MPU contains up to 32 instruction regions and 32 data regions. Each region is defined by the following attributes:

- Base address
- Region type
- Region index

- Region size or upper address limit
- Access permissions
- Default cacheability (data regions only)

Base Address

The base address specifies the lowest address of the region. The base address is aligned on a region-sized boundary. For example, a 4 Kbyte region must have a base address that is a multiple of 4 Kbytes. If the base address is not properly aligned, the behavior is undefined.

Region Type

Each region is identified as either an instruction region or a data region.

Region Index

Each region has an index ranging from zero to the number of regions of its region type minus one. Index zero has the highest priority.

Region Size or Upper Address Limit

An SOPC Builder generation-time option controls whether the amount of memory in the region is defined by size or upper address limit. The size is an integer power of two bytes. The limit is the highest address of the region plus one. The minimum supported region size is 64 bytes but can be configured at system generation time for larger minimum sizes to save logic resources. The maximum supported region size equals the Nios II address space (a function of the address ranges of slaves connected to the Nios II masters). Any access outside of the Nios II address space is considered not to match any region and triggers an MPU region violation exception.

When regions are defined by size, the size is encoded as a binary mask to facilitate the following MPU region address range matching:

```
(address & region_mask) == region_base_address
```

When regions are defined by limit, the limit is encoded as an unsigned integer to facilitate the following MPU region address range matching:

```
(address >= region_base) && (address < region_limit)
```

The region limit uses a less-than instead of a less-than-or-equal-to comparison because less-than provides a more efficient implementation. The limit is one bit larger than the address so that full address range may be included in a range. Defining the region by limit results in slower and larger address range match logic than defining by size but allows finer granularity in region sizes.

Access Permissions

The access permissions consist of execute permissions for instruction regions and read/write permissions for data regions. Any instruction that performs a memory access that violates the access permissions triggers an exception. Additionally, any instruction that performs a memory access that does not match any region triggers an exception.

Default Cacheability

The default cacheability specifies whether normal load and store instructions access the data cache or bypass the data cache. The default cacheability is only present for data regions. You can override the default cacheability by using the `ldio` or `stio` instructions. The bit 31 cache bypass feature is available when the MPU is present. Refer to "Cache Memory" on page 3–52 for more information on cache bypass.

Overlapping Regions

The memory addresses of regions can overlap. Overlapping regions have several uses including placing markers or small holes inside of a larger region. For example, the stack and heap may be located in the same region, growing from opposite ends of the address range. To detect stack/heap overflows, you can define a small region between the stack and heap with no access permissions and assign it a higher priority than the larger region. Any access attempts to the hole region trigger an exception informing system software about the stack/heap overflow.

If regions overlap so that a particular access matches more than one region, the region with the highest priority (lowest index) determines the access permissions and default cacheability.

Enabling the MPU

The MPU is disabled on system reset. System software enables and disables the MPU by writing to a control register. Before enabling the MPU, you must create at least one instruction and one data region, otherwise unexpected results can occur. Refer to "Working with the MPU" on page 3–28 for more information.

Registers

The Nios II register set includes general-purpose registers and control registers. In addition, the Nios II/f core can optionally have shadow register sets. This section discusses each type of registers.

General-Purpose Registers

The Nios II architecture provides thirty-two 32-bit general-purpose registers, `r0` through `r31`, as shown in Table 3–5. Some registers have names recognized by the assembler. For example, the `zero` register (`r0`) always returns the value zero, and writing to `zero` has no effect. The `ra` register (`r31`) holds the return address used by procedure calls and is implicitly accessed by the `call`, `callr` and `ret` instructions. C and C++ compilers use a common procedure-call convention, assigning specific meaning to registers `r1` through `r23` and `r26` through `r28`.

 © November 2009 Altera Corporation

Table 3–5. The Nios II General Purpose Registers

Register	Name	Function	Register	Name	Function
r0	zero	0x00000000	r16		
r1	at	Assembler temporary	r17		
r2		Return value	r18		
r3		Return value	r19		
r4		Register arguments	r20		
r5		Register arguments	r21		
r6		Register arguments	r22		
r7		Register arguments	r23		
r8		Caller-saved register	r24	et	Exception temporary
r9		Caller-saved register	r25	bt	Breakpoint temporary *(1)*
r10		Caller-saved register	r26	gp	Global pointer
r11		Caller-saved register	r27	sp	Stack pointer
r12		Caller-saved register	r28	fp	Frame pointer
r13		Caller-saved register	r29	ea	Exception return address
r14		Caller-saved register	r30	ba	Breakpoint return address *(2)*
r15		Caller-saved register	r31	ra	Return address

Notes to Table 3–5:

(1) r25 is used exclusively by the JTAG debug module. It is used as the breakpoint temporary (bt) register in the normal register set. In shadow register sets, r25 is reserved.

(2) r30 is used as the breakpoint return address (ba) in the normal register set, and as the shadow register set status (sstatus) in each shadow register set. For details about sstatus, refer to "The sstatus Register" on page 3–26.

For more information, refer to the *Application Binary Interface* chapter of the *Nios II Processor Reference Handbook*.

Control Registers

Control registers report the status and change the behavior of the processor. Control registers are accessed differently than the general-purpose registers. The special instructions rdctl and wrctl provide the only means to read and write to the control registers and are only available in supervisor mode.

When writing to control registers, all undefined bits must be written as zero.

The Nios II architecture supports up to 32 control registers. Table 3–6 shows details of the defined control registers. All nonreserved control registers have names recognized by the assembler.

Table 3–6. Control Register Names and Bits (Part 1 of 2)

Register	Name	Register Contents
0	status	Refer to Table 3–7 on page 3–12
1	estatus	Refer to Table 3–9 on page 3–14
2	bstatus	Refer to Table 3–10 on page 3–15

Table 3–6. Control Register Names and Bits (Part 2 of 2)

Register	Name	Register Contents
3	`ienable`	Internal interrupt-enable bits *(3)*
4	`ipending`	Pending internal interrupt bits *(3)*
5	`cpuid`	Unique processor identifier
6	Reserved	Reserved
7	`exception`	Refer to Table 3–11 on page 3–16
8	`pteaddr` *(1)*	Refer to Table 3–13 on page 3–16
9	`tlbacc` *(1)*	Refer to Table 3–15 on page 3–17
10	`tlbmisc` *(1)*	Refer to Table 3–17 on page 3–18
11	Reserved	Reserved
12	`badaddr`	Refer to Table 3–19 on page 3–21
13	`config` *(2)*	Refer to Table 3–21 on page 3–21
14	`mpubase` *(2)*	Refer to Table 3–23 on page 3–22
15	`mpuacc` *(2)*	Refer to Table 3–25 on page 3–23
16–31	Reserved	Reserved

Notes to Table 3–6:

(1) Available only when the MMU is present. Otherwise reserved.

(2) Available only when the MPU is present. Otherwise reserved.

(3) Available only when the external interrupt controller interface is not present. Otherwise reserved.

The following sections describe the nonreserved control registers.

The status Register

The value in the `status` register determines the state of the Nios II processor. All `status` bits are set to predefined values at processor reset. Some bits are exclusively used by and available only to certain features of the processor, such as the MMU, MPU or external interrupt controller (EIC) interface. Table 3–7 shows the layout of the `status` register.

Table 3–7. status Control Register Fields

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Reserved								`RSIE`	`NMI`	`PRS`						`CRS`						`IL`						`IH`	`EH`	`U`	`PIE`

Table 3–8 gives details of the fields defined in the `status` register.

Table 3–8. status Control Register Field Descriptions (Part 1 of 2)

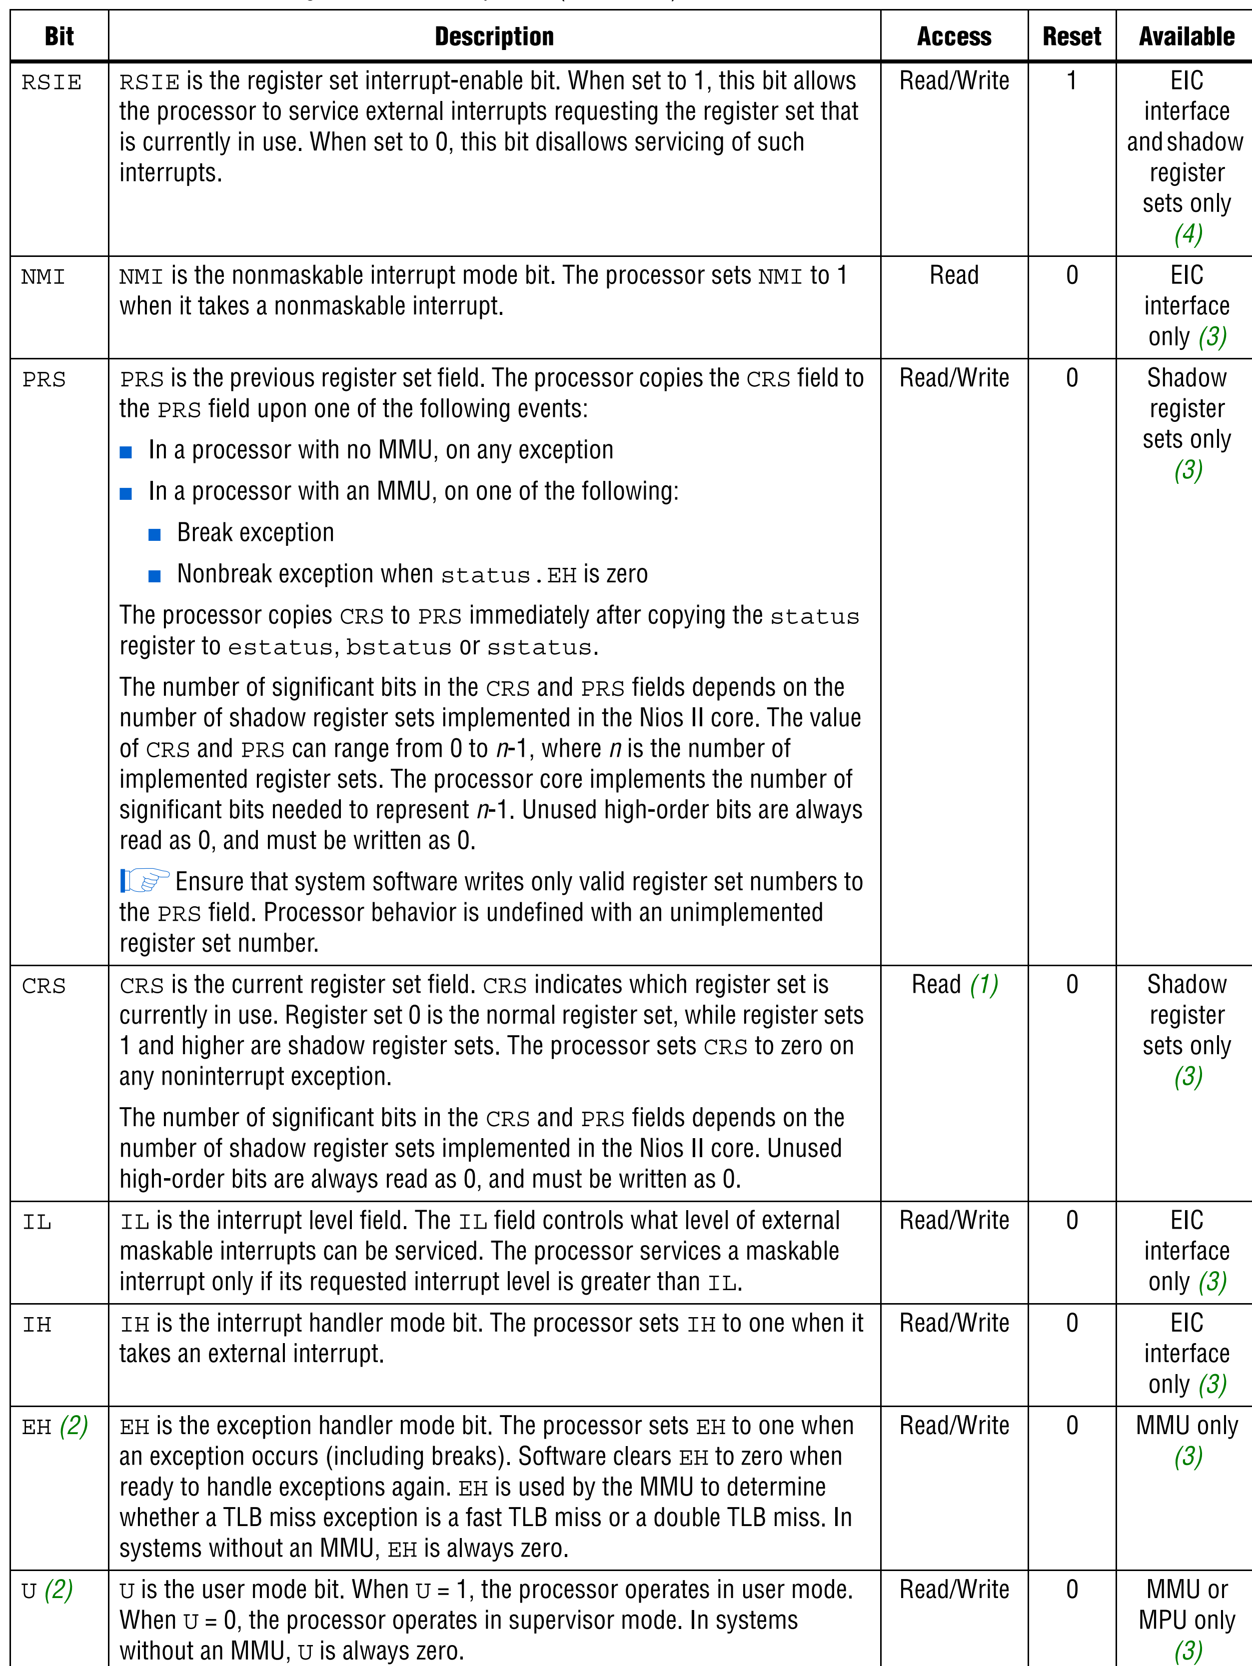

Bit	Description	Access	Reset	Available
`RSIE`	`RSIE` is the register set interrupt-enable bit. When set to 1, this bit allows the processor to service external interrupts requesting the register set that is currently in use. When set to 0, this bit disallows servicing of such interrupts.	Read/Write	1	EIC interface and shadow register sets only *(4)*
`NMI`	`NMI` is the nonmaskable interrupt mode bit. The processor sets `NMI` to 1 when it takes a nonmaskable interrupt.	Read	0	EIC interface only *(3)*

`PRS`	`PRS` is the previous register set field. The processor copies the `CRS` field to the `PRS` field upon one of the following events:
	■ In a processor with no MMU, on any exception
	■ In a processor with an MMU, on one of the following:
	  ■ Break exception
	  ■ Nonbreak exception when `status.EH` is zero
	The processor copies `CRS` to `PRS` immediately after copying the `status` register to `estatus`, `bstatus` or `sstatus`.
	The number of significant bits in the `CRS` and `PRS` fields depends on the number of shadow register sets implemented in the Nios II core. The value of `CRS` and `PRS` can range from 0 to *n*-1, where *n* is the number of implemented register sets. The processor core implements the number of significant bits needed to represent *n*-1. Unused high-order bits are always read as 0, and must be written as 0.
	Ensure that system software writes only valid register set numbers to the `PRS` field. Processor behavior is undefined with an unimplemented register set number.
		Read/Write	0	Shadow register sets only *(3)*

`CRS`	`CRS` is the current register set field. `CRS` indicates which register set is currently in use. Register set 0 is the normal register set, while register sets 1 and higher are shadow register sets. The processor sets `CRS` to zero on any noninterrupt exception. The number of significant bits in the `CRS` and `PRS` fields depends on the number of shadow register sets implemented in the Nios II core. Unused high-order bits are always read as 0, and must be written as 0.	Read *(1)*	0	Shadow register sets only *(3)*
`IL`	`IL` is the interrupt level field. The `IL` field controls what level of external maskable interrupts can be serviced. The processor services a maskable interrupt only if its requested interrupt level is greater than `IL`.	Read/Write	0	EIC interface only *(3)*
`IH`	`IH` is the interrupt handler mode bit. The processor sets `IH` to one when it takes an external interrupt.	Read/Write	0	EIC interface only *(3)*
`EH` *(2)*	`EH` is the exception handler mode bit. The processor sets `EH` to one when an exception occurs (including breaks). Software clears `EH` to zero when ready to handle exceptions again. `EH` is used by the MMU to determine whether a TLB miss exception is a fast TLB miss or a double TLB miss. In systems without an MMU, `EH` is always zero.	Read/Write	0	MMU only *(3)*
`U` *(2)*	`U` is the user mode bit. When `U` = 1, the processor operates in user mode. When `U` = 0, the processor operates in supervisor mode. In systems without an MMU, `U` is always zero.	Read/Write	0	MMU or MPU only *(3)*

Table 3–8. status Control Register Field Descriptions (Part 2 of 2)

Bit	Description	Access	Reset	Available
PIE	PIE is the processor interrupt-enable bit. When PIE = 0, internal and maskable external interrupts and noninterrupt exceptions are ignored. When PIE = 1, internal and maskable external interrupts can be taken, depending on the status of the interrupt controller. Noninterrupt exceptions are unaffected by PIE.	Read/Write	0	Always

Notes to Table 3–8:

(1) The CRS field is read-only. For information about manually changing register sets, refer to "External Interrupt Controller Interface" on page 3–35.

(2) The state where both EH and U are one is illegal and causes undefined results.

(3) When this field is unimplemented, the field value always reads as 0, and the processor behaves accordingly.

(4) When this field is unimplemented, the field value always reads as 1, and the processor behaves accordingly.

The estatus Register

The estatus register holds a saved copy of the status register during nonbreak exception processing. Table 3–9 shows the layout of the estatus register.

Table 3–9. estatus Control Register Fields

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Reserved								RSIE	NMI	PRS						CRS						IL						IH	EH	U	PIE

All fields in the estatus register have read/write access. All fields reset to 0.

Table 3–8 describes the details of the fields defined in the estatus register.

When the Nios II processor takes an interrupt, if status.eh is zero (that is, the MMU is in nonexception mode), the processor copies the contents of the status register to estatus.

If shadow register sets are implemented, and the interrupt requests a shadow register set, the Nios II processor copies status to sstatus, not to estatus.

For details about the sstatus register, refer to "The sstatus Register" on page 3–26.

The exception handler can examine estatus to determine the pre-exception status of the processor. When returning from an exception, the eret instruction restores the pre-exception value of status. The instruction restores the pre-exception value by copying either estatus or sstatus back to status, depending on the value of status.CRS.

Refer to "Exception Processing" on page 3–30 for more information.

The bstatus Register

The bstatus register holds a saved copy of the status register during break exception processing. Table 3–10 shows the layout of the bstatus register.

Table 3–10. bstatus Control Register Fields

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Reserved								RSIE	NMI	PRS						CRS						IL						IH	EH	U	PIE

All fields in the `bstatus` register have read/write access. All fields reset to 0.

Table 3–8 describes the details of the fields defined in the `bstatus` register.

When a break occurs, the value of the `status` register is copied into `bstatus`. Using `bstatus`, the debugger can restore the `status` register to the value prior to the break. The `bret` instruction causes the processor to copy bstatus back to status. Refer to "Processing a Break" on page 3–34 for more information.

The ienable Register

The `ienable` register controls the handling of internal hardware interrupts. Each bit of the `ienable` register corresponds to one of the interrupt inputs, `irq0` through `irq31`. A value of one in bit *n* means that the corresponding `irq`*n* interrupt is enabled; a bit value of zero means that the corresponding interrupt is disabled. Refer to "Exception Processing" on page 3–30 for more information.

When the internal interrupt controller is not implemented, the value of the `ienable` register is always 0.

The ipending Register

The value of the `ipending` register indicates the value of the interrupt signals driven into the processor. A value of one in bit *n* means that the corresponding `irq`*n* input is asserted. Writing a value to the `ipending` register has no effect.

The `ipending` register is present only when the internal interrupt controller is implemented.

The cpuid Register

The `cpuid` register holds a constant value that uniquely identifies each processor in a multiprocessor system. The `cpuid` value is determined at system generation time and is guaranteed to be unique for each processor in the system. Writing to the `cpuid` register has no effect.

The exception Register

When the extra exception information option is enabled, the Nios II processor provides information useful to system software for exception processing in the `exception` and `badaddr` registers when an exception occurs. When your system contains an MMU or MPU, the extra exception information is always enabled. When no MMU or MPU is present, the Nios II Megawizard interface gives you the option to have the processor provide the extra exception information.

To see how to control the extra exception information option, refer to the *Instantiating the Nios II Processor in SOPC Builder* chapter of the *Nios II Processor Reference Handbook.*

Table 3–11 shows the layout of the `exception` register.

Table 3–11. exception Control Register Field Descriptions

Field	Description	Access	Reset	Available
`CAUSE`	`CAUSE` is written by the Nios II processor when certain exceptions occur. `CAUSE` contains a code for the highest-priority exception occurring at the time. The Cause column in Table 3–33 on page 3–31 shows the `CAUSE` field value for each exception. `CAUSE` is not written on a break or an external interrupt.	Read	0	Only with extra exception information

Table 3–12 gives details of the fields defined in the `exception` register.

Table 3–12. exception Control Register Fields

<table>
<tr><th>31</th><th>30</th><th>29</th><th>28</th><th>27</th><th>26</th><th>25</th><th>24</th><th>23</th><th>22</th><th>21</th><th>20</th><th>19</th><th>18</th><th>17</th><th>16</th><th>15</th><th>14</th><th>13</th><th>12</th><th>11</th><th>10</th><th>9</th><th>8</th><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td colspan="25">Reserved</td><td colspan="5">CAUSE</td><td colspan="2">Rsvd</td></tr>
</table>

The pteaddr Register

The `pteaddr` register contains the virtual address of the operating system's page table and is only available in systems with an MMU. The `pteaddr` register layout accelerates fast TLB miss exception handling. Table 3–13 shows the layout of the `pteaddr` register.

Table 3–13. pteaddr Control Register Fields

<table>
<tr><th>31</th><th>30</th><th>29</th><th>28</th><th>27</th><th>26</th><th>25</th><th>24</th><th>23</th><th>22</th><th>21</th><th>20</th><th>19</th><th>18</th><th>17</th><th>16</th><th>15</th><th>14</th><th>13</th><th>12</th><th>11</th><th>10</th><th>9</th><th>8</th><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td colspan="10">PTBASE</td><td colspan="20">VPN</td><td colspan="2">Rsvd</td></tr>
</table>

Table 3–14 gives details of the fields defined in the `pteaddr` register.

Table 3–14. pteaddr Control Register Field Descriptions

Field	Description	Access	Reset	Available
`PTBASE`	`PTBASE` is the base virtual address of the page table.	Read/Write	0	Only with MMU
`VPN`	`VPN` is the virtual page number. `VPN` can be set by both hardware and software.	Read/Write	0	Only with MMU

Software writes to the `PTBASE` field when switching processes. Hardware never writes to the `PTBASE` field.

Software writes to the `VPN` field when writing a TLB entry. Hardware writes to the `VPN` field on a fast TLB miss exception, a TLB permission violation exception, or on a TLB read operation. The `VPN` field is not written on any exceptions taken when an exception is already active, that is, when `status.EH` is already one.

The tlbacc Register

The `tlbacc` register is used to access TLB entries and is only available in systems with an MMU. The `tlbacc` register holds values that software will write into a TLB entry or has previously read from a TLB entry. The `tlbacc` register provides access to only a portion of a complete TLB entry. `pteaddr.VPN` and `tlbmisc.PID` hold the remaining TLB entry fields.

Table 3–15 shows the layout of the `tlbacc` register.

Table 3–15. tlbacc Control Register Fields

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IG							C	R	W	X	G	PFN																			

Table 3–16 gives details of the fields defined in the `tlbacc` register.

Issuing a `wrctl` instruction to the `tlbacc` register writes the `tlbacc` register with the specified value. If `tlbmisc.WE` = 1, the `wrctl` instruction also initiates a TLB write operation, which writes a TLB entry. The TLB entry written is specified by the line portion of `pteaddr.VPN` and the `tlbmisc.WAY` field. The value written is specified by the value written into `tlbacc` along with the values of `pteaddr.VPN` and `tlbmisc.PID`. A TLB write operation also increments `tlbmisc.WAY`, allowing software to quickly modify TLB entries.

Issuing a `rdctl` instruction to the `tlbacc` register returns the value of the `tlbacc` register. The `tlbacc` register is written by hardware when software triggers a TLB read operation (that is, when `wrctl` sets `tlbmisc.RD` to one).

Table 3–16. tlbacc Control Register Field Descriptions

Field	Description	Access	Reset	Available
`IG`	`IG` is ignored by hardware and available to hold operating system specific information. Read as zero but can be written as nonzero.	Read/Write	0	Only with MMU
`C`	`C` is the data cacheable flag. When `C` = 0, data accesses are uncacheable. When `C` = 1, data accesses are cacheable.	Read/Write	0	Only with MMU
`R`	`R` is the readable flag. When `R` = 0, load instructions are not allowed to access memory. When `R` = 1, load instructions are allowed to access memory.	Read/Write	0	Only with MMU
`W`	`W` is the writable flag. When `W` = 0, store instructions are not allowed to access memory. When `W` = 1, store instructions are allowed to access memory.	Read/Write	0	Only with MMU
`X`	`X` is the executable flag. When `X` = 0, instructions are not allowed to execute. When `X` = 1, instructions are allowed to execute.	Read/Write	0	Only with MMU
`G`	`G` is the global flag. When `G` = 0, `tlbmisc.PID` is included in the TLB lookup. When `G` = 1, `tlbmisc.PID` is ignored and only the virtual page number is used in the TLB lookup.	Read/Write	0	Only with MMU
`PFN`	`PFN` is the physical frame number field. All unused upper bits must be zero.	Read/Write	0	Only with MMU

The `tlbacc` register format is the recommended format for entries in the operating system page table. The `IG` bits are ignored by the hardware on `wrctl` to `tlbacc` and read back as zero on `rdctl` from `tlbacc`. The operating system can use the `IG` bits to hold operating system specific information without having to clear these bits to zero on a TLB write operation.

The tlbmisc Register

The `tlbmisc` register contains the remaining TLB-related fields and is only available in systems with an MMU. Table 3–17 shows the layout of the `tlbmisc` register.

Table 3–17. tlbmisc Control Register Fields

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Reserved								`WAY` *(1)*				`RD`	`WE`	`PID` *(1)*														`DBL`	`BAD`	`PERM`	`D`

Notes to Table 3–17:

(1) This field size is variable. Unused upper bits must be written as zero.

Table 3–18 gives details of the fields defined in the `tlbmisc` register.

Table 3–18. tlbmisc Control Register Field Descriptions

Field	Description	Access	Reset	Available
`WAY`	The `WAY` field controls the mapping from the VPN to a particular TLB entry.	Read/Write	0	Only with MMU
`RD`	`RD` is the read flag. Setting `RD` to one triggers a TLB read operation.	Write	0	Only with MMU
`WE`	`WE` is the TLB write enable flag. When `WE` = 1, a write to `tlbacc` writes through to a TLB entry.	Read/Write	0	Only with MMU
`PID`	`PID` is the process identifier field.	Read/Write	0	Only with MMU
`DBL` *(1)*	`DBL` is the double TLB miss exception flag.	Read	0	Only with MMU
`BAD` *(1)*	`BAD` is the bad virtual address exception flag.	Read	0	Only with MMU
`PERM` *(1)*	`PERM` is the TLB permission violation exception flag.	Read	0	Only with MMU
`D`	`D` is the data access exception flag. When `D` = 1, the exception is a data access exception. When `D` = 0, the exception is an instruction access exception.	Read	0	Only with MMU

Notes to Table 3–18:

(1) You can also use `exception.CAUSE` to determine these exceptions.

The following sections provide further details of the `tlbmisc` fields.

The RD Flag

System software triggers a TLB read operation by setting `tlbmisc.RD` (with a `wrctl` instruction). A TLB read operation loads the following register fields with the contents of a TLB entry:

- The tag portion of `pteaddr.VPN`
- `tlbmisc.PID`
- The `tlbacc` register

The TLB entry to be read is specified by the following values:

- the line portion of `pteaddr.VPN`
- `tlbmisc.WAY`

© November 2009 Altera Corporation

When system software changes the fields that specify the TLB entry, there is no immediate effect on `pteaddr.VPN`, `tlbmisc.PID`, or the `tlbacc` register. The registers retain their previous values until the next TLB read operation is initiated. For example, when the operating system sets `pteaddr.VPN` to a new value, the contents of `tlbacc` continues to reflect the previous TLB entry. `tlbacc` does not contain the new TLB entry until after an explicit TLB read.

The WE Flag

When `WE` = 1, a write to `tlbacc` writes the `tlbacc` register and a TLB entry. When `WE` = 0, a write to `tlbacc` only writes the `tlbacc` register.

Hardware sets the `WE flag` to one on a TLB permission violation exception, and on a TLB miss exception when `status.EH` = 0. When a TLB write operation writes the `tlbacc` register, the write operation also writes to a TLB entry when `WE` = 1.

The WAY Field

The `WAY` field controls the mapping from the VPN to a particular TLB entry. `WAY` specifies the set to be written to in the TLB. The MMU increments `WAY` when system software performs a TLB write operation. Unused upper bits in `WAY` must be written as zero.

The number of ways (sets) is configurable in SOPC Builder at generation time, up to a maximum of 16.

The PID Field

`PID` is a unique identifier for the current process that effectively extends the virtual address. The process identifier can be less than 14 bits. Any unused upper bits must be zero.

`tlbmisc.PID` contains the `PID` field from a TLB tag. The operating system must set the `PID` field when switching processes, and before each TLB write operation.

Use of the process identifier is optional. To implement memory management without process identifiers, clear `tlbmisc.PID` to zero. Without a process identifier, all processes share the same virtual address space.

The MMU sets `tlbmisc.PID` on a TLB read operation. When the software triggers a TLB read, by setting `tlbmisc.RD` to one with the `wrctl` instruction, the `PID` value read from the TLB has priority over the value written by the `wrctl` instruction.

The size of the `PID` field is configured in SOPC Builder at system generation, and can be from 8 to 14 bits. If system software defines a process identifier smaller than the `PID` field, unused upper bits must be written as zero.

The DBL Flag

During a general exception, the processor sets `DBL` to one when a double TLB miss condition exists. Otherwise, the processor clears `DBL` to zero.

The `DBL` flag indicates whether the most recent exception is a double TLB miss condition. When a general exception occurs, the MMU sets `DBL` to one if a double TLB miss is detected, and clears `DBL` to zero otherwise.

The BAD Flag

During a general exception, the processor sets BAD to one when a bad virtual address condition exists, and clears BAD to zero otherwise. The following exceptions set the BAD flag to one:

- Supervisor-only instruction address
- Supervisor-only data address
- Misaligned data address
- Misaligned destination address

Refer to Table 3–33 on page 3–31 for more information on these exceptions.

The PERM Flag

During a general exception, the processor sets PERM to one for a TLB permission violation exception, and clears PERM to zero otherwise.

The D Flag

The D flag indicates whether the exception is an instruction access exception or a data access exception. During a general exception, the processor sets D to one when the exception is related to a data access, and clears D to zero for all other nonbreak exceptions.

The following exceptions set the D flag to one:

- Fast TLB miss (data)
- Double TLB miss (data)
- TLB permission violation (read or write)
- Misaligned data address
- Supervisor-only data address

The badaddr Register

When the extra exception information option is enabled, the Nios II processor provides information useful to system software for exception processing in the exception and badaddr registers when an exception occurs. When your system contains an MMU or MPU, the extra exception information is always enabled. When no MMU or MPU is present, the Nios II Megawizard interface gives you the option to have the processor provide the extra exception information.

To see how to control the extra exception information option, refer to the *Instantiating the Nios II Processor in SOPC Builder* chapter of the *Nios II Processor Reference Handbook*.

When the option for extra exception information is enabled and a processor exception occurs, the badaddr register contains the byte instruction or data address associated with certain exceptions at the time the exception occurred. Table 3–33 on page 3–31 shows which exceptions write the badaddr register along with the value written. Table 3–19 shows the layout of the badaddr register.

Table 3–19. badaddr Control Register Fields

<table>
<tr><th>31</th><th>30</th><th>29</th><th>28</th><th>27</th><th>26</th><th>25</th><th>24</th><th>23</th><th>22</th><th>21</th><th>20</th><th>19</th><th>18</th><th>17</th><th>16</th><th>15</th><th>14</th><th>13</th><th>12</th><th>11</th><th>10</th><th>9</th><th>8</th><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td colspan="32">BADDR</td></tr>
</table>

Table 3–20 gives details of the fields defined in the badaddr register.

Table 3–20. badaddr Control Register Field Descriptions

Field	Description	Access	Reset	Available
BADDR	BADDR contains the byte instruction address or data address associated with an exception when certain exceptions occur. The Address column of Table 3–33 on page 3–31 shows which exceptions write the BADDR field.	Read	0	Only with extra exception information

The BADDR field allows up to a 32-bit instruction address or data address. If an MMU or MPU is present, the BADDR field is 32 bits because MMU and MPU instruction and data addresses are always full 32-bit values. When an MMU is present, the BADDR field contains the virtual address.

If there is no MMU or MPU and the Nios II address space is less than 32 bits, unused high-order bits are written and read as zero. If there is no MMU, bit 31 of a data address (used to bypass the data cache) is always zero in the BADDR field.

The config Register

The config register configures Nios II runtime behaviors that do not need to be preserved during exception processing (in contrast to the information in the status register). Table 3–21 shows the layout of the config register.

Table 3–21. config Control Register Fields

<table>
<tr><th>31</th><th>30</th><th>29</th><th>28</th><th>27</th><th>26</th><th>25</th><th>24</th><th>23</th><th>22</th><th>21</th><th>20</th><th>19</th><th>18</th><th>17</th><th>16</th><th>15</th><th>14</th><th>13</th><th>12</th><th>11</th><th>10</th><th>9</th><th>8</th><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td colspan="30">Reserved</td><td>ANI</td><td>PE</td></tr>
</table>

Table 3–22 gives details of the fields defined in the config register

Table 3–22. config Control Register Field Descriptions

Field	Description	Access	Reset	Available
PE	PE is the memory protection enable bit. When PE =1, the MPU is enabled. When PE = 0, the MPU is disabled. In systems without an MPU, PE is always zero.	Read/Write	0	Only with MPU
ANI	ANI is the automatic nested interrupt mode bit. If ANI is set to zero, the processor clears status.PIE on each interrupt, disabling fast nested interrupts. If ANI is set to one, the processor leaves status.PIE set to one at the time of an interrupt, enabling fast nested interrupts. If the EIC interface and shadow register sets are not implemented in the Nios II core, ANI always reads as zero, disabling fast nested interrupts.	Read/Write	0	Only with the EIC interface and shadow register sets

The mpubase Register

The `mpubase` register works in conjunction with the `mpuacc` register to set and retrieve MPU region information and is only available in systems with an MPU. Table 3–23 shows the layout of the `mpubase` register.

Table 3–23. mpubase Control Register Fields

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	BASE *(2)*																									INDEX *(1)*					D

Notes to Table 3–23:

(1) This field size is variable. Unused upper bits must be written as zero.

(2) This field size is variable. Unused upper bits and unused lower bits must be written as zero.

Table 3–24 gives details of the fields defined in the `mpubase` register

Table 3–24. mpubase Control Register Field Descriptions

Field	Description	Access	Reset	Available
`BASE`	`BASE` is the base memory address of the region identified by the `INDEX` and `D` fields.	Read/Write	0	Only with MPU
`INDEX`	`INDEX` is the region index number.	Read/Write	0	Only with MPU
`D`	`D` is the region access bit. When `D` =1, `INDEX` refers to a data region. When `D` = 0, `INDEX` refers to an instruction region.	Read/Write	0	Only with MPU

The `BASE` field specifies the base address of an MPU region. The 25-bit `BASE` field corresponds to bits 6 through 30 of the base address, making the base address always a multiple of 64 bytes. If the minimum region size set in SOPC Builder at generation time is larger than 64 bytes, unused low-order bits of the `BASE` field must be written as zero and are read as zero. For example, if the minimum region size is 1024 bytes, the four least-significant bits of the `BASE` field (bits 6 though 9 of the `mpubase` register) must be zero. Similarly, if the Nios II address space is less than 31 bits, unused high-order bits must also be written as zero and are read as zero.

The `INDEX` and `D` fields specify the region information to access when an MPU region read or write operation is performed. The `D` field specifies whether the region is a data region or an instruction region. The `INDEX` field specifies which of the 32 data or instruction regions to access. If there are fewer than 32 instruction or 32 data regions, unused high-order bits must be written as zero and are read as zero.

Refer to "MPU Region Read and Write Operations" on page 3–28 for more information on MPU region read and write operations.

The mpuacc Register

The `mpuacc` register works in conjunction with the `mpubase` register to set and retrieve MPU region information and is only available in systems with an MPU. The `mpuacc` register consists of attributes that will be set or have been retrieved which define the MPU region. The `mpuacc` register only holds a portion of the attributes that define an MPU region. The remaining portion of the MPU region definition is held by the `BASE` field of the `mpubase` register.

An SOPC Builder generation-time option controls whether the `mpuacc` register contains a `MASK` or `LIMIT` field. Table 3–25 shows the layout of the `mpuacc` register with the `MASK` field. Table 3–26 shows the layout of the `mpuacc` register with the `LIMIT` field.

Table 3–25. mpuacc Control Register Fields for MASK Variation

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	`MASK` *(1)*																									C	`PERM`			`RD`	`WR`

Note to Table 3–25:

(1) This field size is variable. Unused upper bits and unused lower bits must be written as zero.

Table 3–26. mpuacc Control Register Fields for LIMIT Variation

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
`LIMIT` *(1)*																										C	`PERM`			`RD`	`WR`

Note to Table 3–26:

(1) This field size is variable. Unused upper bits and unused lower bits must be written as zero.

Table 3–27 gives details of the fields defined in the `mpuacc` register.

Table 3–27. mpuacc Control Register Field Descriptions

Field	Description	Access	Reset	Available
`MASK` *(1)*	`MASK` specifies the size of the region.	Read/Write	0	Only with MPU
`LIMIT` *(1)*	`LIMIT` specifies the upper address limit of the region.	Read/Write	0	Only with MPU
`C`	`C` is the data cacheable flag. `C` only applies to MPU data regions and determines the default cacheability of a data region. When `C` = 0, the data region is uncacheable. When `C` = 1, the data region is cacheable.	Read/Write	0	Only with MPU
`PERM`	`PERM` specifies the access permissions for the region.	Read/Write	0	Only with MPU
`RD`	`RD` is the read region flag. When `RD` = 1, `wrctl` instructions to the `mpuacc` register perform a read operation.	Write	0	Only with MPU
`WE`	`WR` is the write region flag. When `WR` = 1, `wrctl` instructions to the `mpuacc` register perform a write operation.	Write	0	Only with MPU

Note to Table 3–27:

(1) The `MASK` and `LIMIT` fields are mutually exclusive. Refer to Table 3–25 and Table 3–26.

The following sections provide further details of the `mpuacc` fields.

The MASK Field

When the amount of memory reserved for a region is defined by size, the `MASK` field specifies the size of the memory region. The `MASK` field is the same number of bits as the `BASE` field of the `mpubase` register.

Unused high-order or low-order bits must be written as zero and are read as zero.

Table 3–28 shows the MASK field encodings for all possible region sizes in a full 31-bit byte address space.

Table 3–28. MASK Region Size Encodings

MASK Encoding	Region Size
0x1FFFFFF	64 bytes
0x1FFFFFE	128 bytes
0x1FFFFFC	256 bytes
0x1FFFFF8	512 bytes
0x1FFFFF0	1 Kbyte
0x1FFFFE0	2 Kbytes
0x1FFFFC0	4 Kbytes
0x1FFFF80	8 Kbytes
0x1FFFF00	16 Kbytes
0x1FFFE00	32 Kbytes
0x1FFFC00	64 Kbytes
0x1FFF800	128 Kbytes
0x1FFF000	256 Kbytes
0x1FFE000	512 Kbytes
0x1FFC000	1 Mbyte
0x1FF8000	2 Mbytes
0x1FF0000	4 Mbytes
0x1FE0000	8 Mbytes
0x1FC0000	16 Mbytes
0x1F80000	32 Mbytes
0x1F00000	64 Mbytes
0x1E00000	128 Mbytes
0x1C00000	256 Mbytes
0x1800000	512 Mbytes
0x1000000	1 Gbyte
0x0000000	2 Gbytes

Bit 31 of the mpuacc register is not used by the MASK field. Because memory region size is already a power of two, one less bit is needed. The MASK field contains the following value, where region_size is in bytes:

```
MASK = 0x1FFFFFF << log2(region_size >> 6)
```

The LIMIT Field

When the amount of memory reserved for a region is defined by an upper address limit, the `LIMIT` field specifies the upper address of the memory region plus one. For example, to achieve a memory range for byte addresses `0x4000` to `0x4fff` with a 256 byte minimum region size, the `BASE` field of the `mpubase` register is set to `0x40` (`0x4000 >> 8`) and the `LIMIT` field is set to `0x50` (`0x5000 >> 8`). Because the `LIMIT` field is one more bit than the number of bits of the `BASE` field of the `mpubase` register, bit 31 of the `mpuacc` register is available to the `LIMIT` field.

The C Flag

The `C` flag determines the default data cacheability of an MPU region. The `C` flag only applies to data regions. For instruction regions, the `C` bit must be written with 0 and is always read as 0.

When data cacheability is enabled on a data region, a data access to that region can be cached, if a data cache is present in the system. You can override the default cacheability and force an address to noncacheable with an `ldio` or `stio` instruction.

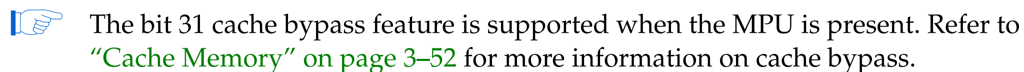

The bit 31 cache bypass feature is supported when the MPU is present. Refer to "Cache Memory" on page 3–52 for more information on cache bypass.

The PERM Field

The `PERM` field specifies the allowed access permissions. Table 3–29 shows possible values of the `PERM` field for instruction regions and Table 3–30 shows possible values of the `PERM` field for data regions.

Table 3–29. Instruction Region Permission Values

Value	Supervisor Permissions	User Permissions
0	None	None
1	Execute	None
2	Execute	Execute

Table 3–30. Data Region Permission Values

Value	Supervisor Permissions	User Permissions
0	None	None
1	Read	None
2	Read	Read
4	Read/Write	None
5	Read/Write	Read
6	Read/Write	Read/Write

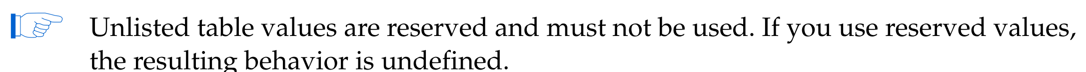

Unlisted table values are reserved and must not be used. If you use reserved values, the resulting behavior is undefined.

The RD Flag

Setting the `RD` flag signifies that an MPU region read operation should be performed when a `wrctl` instruction is issued to the `mpuacc` register. Refer to "MPU Region Read and Write Operations" on page 3–28 for more information. The `RD` flag always returns 0 when read by a `rdctl` instruction.

The WR Flag

Setting the `WR` flag signifies that an MPU region write operation should be performed when a `wrctl` instruction is issued to the `mpuacc` register. Refer to "MPU Region Read and Write Operations" on page 3–28 for more information. The `WR` flag always returns 0 when read by a `rdctl` instruction.

Setting both the `RD` and `WR` flags to one results in undefined behavior.

Shadow Register Sets

The Nios II processor can optionally have one or more shadow register sets. A shadow register set is a complete alternate set of Nios II general-purpose registers, which can be used to maintain a separate runtime context for an interrupt service routine (ISR).

When shadow register sets are implemented, `status.CRS` indicates the register set currently in use. A Nios II core can have up to 63 shadow register sets. If n is the configured number of shadow register sets, the shadow register sets are numbered from 1 to n. Register set 0 is the normal register set.

A shadow register set behaves precisely the same as the normal register set. The register set currently in use can only be determined by examining `status.CRS`.

When shadow register sets and the EIC interface are implemented on the Nios II core, you must ensure that your software is built with the Nios II EDS version 9.0 or later. Earlier versions have an implementation of the `eret` instruction that is incompatible with shadow register sets.

Shadow register sets are typically used in conjunction with the EIC interface. This combination can substantially reduce interrupt latency.

For details of EIC interface usage, refer to "Exception Processing" on page 3–30.

System software can read from and write to any shadow register set by setting `status.PRS` and using the `rdprs` and `wrprs` instructions.

For details of the `rdprs` and `wrprs` instructions, refer to the *Instruction Set Reference* chapter of the *Nios II Processor Reference Handbook.*

The sstatus Register

The value in the `sstatus` register preserves the state of the Nios II processor during external interrupt handling. The value of `sstatus` is undefined at processor reset. Some bits are exclusively used by and available only to certain features of the processor. Table 3–31 shows the layout of the `sstatus` register.

The `sstatus` register is physically stored in general-purpose register `r30` in each shadow register set. The normal register set does not have an `sstatus` register, but each shadow register set has a separate `sstatus` register.

Table 3–31. sstatus Control Register Fields

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
SRS	Reserved							RSIE	NMI	PRS						CRS						IL						IH	EH	U	PIE

Table 3–32 gives details of the fields defined in the `sstatus` register.

Table 3–32. sstatus Control Register Field Descriptions

Bit	Description	Access	Reset	Available
SRS	SRS is the switched register set bit. The processor sets SRS to 1 when an external interrupt occurs, if the interrupt required the processor to switch to a different register set.	Read/Write	Undefined	EIC interface and shadow register sets only
RSIE	*(1)*	Read/Write	Undefined	*(1)*
NMI	*(1)*	Read/Write	Undefined	*(1)*
PRS	*(1)*	Read/Write	Undefined	*(1)*
CRS	*(1)*	Read/Write	Undefined	*(1)*
IL	*(1)*	Read/Write	Undefined	*(1)*
IH	*(1)*	Read/Write	Undefined	*(1)*
EH	*(1)*	Read/Write	Undefined	*(1)*
U	*(1)*	Read/Write	Undefined	*(1)*
PIE	*(1)*	Read/Write	Undefined	*(1)*

Note to Table 3–32:

(1) Refer to Table 3–8 on page 3–13.

(2) If the EIC interface and shadow register sets are not present SRS always reads as 0, and the processor behaves accordingly.

The `sstatus` register is present in the Nios II core if both the EIC interface and shadow register sets are implemented. There is one copy of `sstatus` for each shadow register set.

When the Nios II processor takes an interrupt, if a shadow register set is requested (RRS = 0) and the MMU is not in exception handler mode (`status.EH` = 0), the processor copies `status` to `sstatus`.

For details about RRS, refer to "Requested Register Set" on page 3–36. For details about `status.EH`, refer to Table 3–35 on page 3–46.

Changing Register Sets

Modifying `status.CRS` immediately switches the Nios II processor to another register set. However, software cannot write to `status.CRS` directly. To modify `status.CRS`, insert the desired value into the saved copy of the `status` register, and then execute the `eret` instruction, as follows:

- If the processor is currently running in the normal register set, insert the new register set number in `estatus.CRS`, and execute `eret`.

- If the processor is currently running in a shadow register set, insert the new register set number in `sstatus.CRS`, and execute `eret`.

Before executing `eret` to change the register set, system software must set individual external interrupt masks correctly to ensure that registers in the shadow register set cannot be corrupted. If an interrupt is assigned to the register set, system software must ensure that one of the following conditions is true:

- The ISR is written to preserve register contents.
- The individual interrupt is disabled. The method for disabling an individual external interrupt is specific to the EIC implementation.

Stacks and Shadow Register Sets

Depending on system requirements, the system software can create a dedicated stack for each register set, or share a stack among several register sets. If a stack is shared, the system software must copy the stack pointer each time the register set changes. Use the `rdprs` instruction to copy the stack register between the current register set and another register set.

Initialization with Shadow Register Sets

At initialization, system software must carry out the following tasks to ensure correct software functioning with shadow register sets:

- After the `gp` register is initialized in the normal register set, copy it to all shadow register sets, to ensure that all code can correctly address the small data sections.
- Copy the `zero` register from the normal register set to all shadow register sets, using the `wrprs` instruction.

Working with the MPU

This section provides a basic overview of MPU initialization and the MPU region read and write operations.

MPU Region Read and Write Operations

MPU region read and write operations are operations that access MPU region attributes through the `mpubase` and `mpuacc` control registers. The `mpubase.BASE`, `mpuacc.MASK`, `mpuacc.LIMIT`, `mpuacc.C`, and `mpuacc.PERM` fields comprise the MPU region attributes.

MPU region read operations retrieve the current values for the attributes of a region. Each MPU region read operation consists of the following actions:

- Execute a `wrctl` instruction to the `mpubase` register with the `mpubase.INDEX` and `mpubase.D` fields set to identify the MPU region.
- Execute a `wrctl` instruction to the `mpuacc` register with the `mpuacc.RD` field set to one and the `mpuacc.WR` field cleared to zero. This action loads the `mpubase` and `mpuacc` register values.

- Execute a `rdctl` instruction to the `mpubase` register to read the loaded the `mpubase` register value.
- Execute a `rdctl` instruction to the `mpuacc` register to read the loaded the `mpuacc` register value.

The MPU region read operation retrieves `mpubase.BASE`, `mpuacc.MASK` or `mpuacc.LIMIT`, `mpuacc.C`, and `mpuacc.PERM` values for the MPU region.

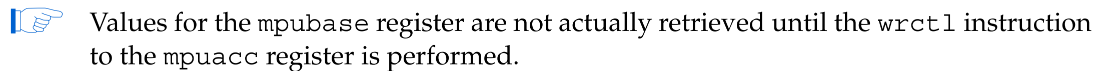

Values for the `mpubase` register are not actually retrieved until the `wrctl` instruction to the `mpuacc` register is performed.

MPU region write operations set new values for the attributes of a region. Each MPU region write operation consists of the following actions:

- Execute a `wrctl` instruction to the `mpubase` register with the `mpubase.INDEX` and `mpubase.D` fields set to identify the MPU region.
- Execute a `wrctl` instruction to the `mpuacc` register with the `mpuacc.WR` field set to one and the `mpuacc.RD` field cleared to zero.

The MPU region write operation sets the values for `mpubase.BASE`, `mpuacc.MASK` or `mpuacc.LIMIT`, `mpuacc.C`, and `mpuacc.PERM` as the new attributes for the MPU region.

Normally, a `wrctl` instruction flushes the pipeline to guarantee that any side effects of writing control registers take effect immediately after the `wrctl` instruction completes execution. However, `wrctl` instructions to the `mpubase` and `mpuacc` control registers do not automatically flush the pipeline. Instead, system software is responsible for flushing the pipeline as needed (either by using a `flushp` instruction or a `wrctl` instruction to a register that does flush the pipeline). Because a context switch typically requires reprogramming the MPU regions for the new thread, flushing the pipeline on each `wrctl` instruction would create unnecessary overhead.

MPU Initialization

Your system software must provide a data structure that contains the region information described in "Memory Regions" on page 3–8 for each active thread. The data structure ideally contains two 32-bit values that correspond to the `mpubase` and `mpuacc` register formats.

The MPU is disabled on system reset. Before enabling the MPU, Altera recommends initializing all MPU regions. Enable desired instruction and data regions by writing each region's attributes to the `mpubase` and `mpuacc` registers as described in "MPU Region Read and Write Operations" on page 3–28. You must also disable unused regions. When using region size, clear `mpuacc.MASK` to zero. When using limit, set the `mpubase.BASE` to a nonzero value and clear `mpuacc.LIMIT` to zero.

You must enable at least one instruction and one data region, otherwise unpredictable behavior might occur.

To perform a context switch, use a `wrctl` to write a zero to the `PE` field of the `config` register to disable the MPU, define all MPU regions from the new thread's data structure, and then use another `wrctl` to write a one to `config.PE` to enable the MPU.

Define each region using the pair of `wrctl` instructions described in "MPU Region Read and Write Operations" on page 3–28. Repeat this dual `wrctl` instruction sequence until all desired regions are defined.

Debugger Access

The debugger can access all MPU-related control registers using the normal `wrctl` and `rdctl` instructions. During debugging, the Nios II ignores the MPU, effectively temporarily disabling it.

Exception Processing

Exception processing is the act of responding to an exception, and then returning, if possible, to the pre-exception execution state.

All Nios II exceptions are precise. Precise exceptions enable the system software to re-execute the instruction, if desired, after handling the exception.

Terminology

Altera Nios II documentation uses the following terminology to discuss exception processing:

- Exception—a transfer of control away from a program's normal flow of execution, caused by an event, either internal or external to the processor, which requires immediate attention.
- Interrupt—an exception caused by an explicit request signal from an external device; also: hardware interrupt.
- Interrupt controller—hardware that interfaces the processor to interrupt request signals from external devices.
- Internal interrupt controller—the nonvectored interrupt controller that is integral to the Nios II processor. The internal interrupt controller is available in all revisions of the Nios II processor.
- Vectored interrupt controller (VIC)—an Altera-provided external interrupt controller.
- Exception (interrupt) latency—The time elapsed between the event that causes the exception (assertion of an interrupt request) and the execution of the first instruction at the handler address.
- Exception (interrupt) response time—The time elapsed between the event that causes the exception (assertion of an interrupt request) and the execution of nonoverhead exception code, that is, specific to the exception type (device).
- Global interrupts—All maskable exceptions on the Nios II processor, including internal interrupts and maskable external interrupts, but not including nonmaskable interrupts.
- Worst-case latency—The value of the exception (interrupt) latency, assuming the maximum disabled time or maximum masked time, and assuming that the exception (interrupt) occurs at the beginning of the masked/disabled time.

 © November 2009 Altera Corporation

- Maximum disabled time—The maximum amount of continuous time that the system spends with maskable interrupts disabled.
- Maximum masked time—The maximum amount of continuous time that the system spends with a single interrupt masked.
- Shadow register set—a complete alternate set of Nios II general-purpose registers, which can be used to maintain a separate runtime context for an ISR.

Exception Overview

Each of the Nios II exceptions falls into one of the following categories:

- Reset exception—Occurs when the Nios II processor is reset. Control is transferred to the reset address configured when the Nios II processor is instantiated.
- Break exception—Occurs when the JTAG debug module requests control. Control is transferred to the break address configured when the Nios II processor is instantiated.
- Interrupt exception—Occurs when a peripheral device signals a condition requiring service
- Instruction-related exception—Occurs when any of several internal conditions occurs, as detailed in Table 3–33 on page 3–31.

Table 3–33 shows all possible Nios II exceptions in order of highest to lowest priority. The following table columns specify information for the exceptions:

- Exception—Gives the name of the exception.
- Type—Specifies the exception type.
- Available—Specifies when support for that exception is present.
- Cause—Specifies the value of the `CAUSE` field of the `exception` register, for exceptions that write the `exception.CAUSE` field.
- Address—Specifies the instruction or data address associated with the exception.
- Vector—Specifies which exception vector address the processor passes control to when the exception occurs.

Table 3–33. Nios II Exceptions (In Decreasing Priority Order) (Part 1 of 3)

Exception	Type	Available	Cause	Address	Vector
Reset	Reset	Always	0		Reset
Hardware Break	Break	Always	—		Break
Processor-only Reset Request	Reset	Always	1		Reset
Internal Interrupt	Interrupt	Internal interrupt controller	2	`ea-4` *(2)*	General exception
External nonmaskable interrupt	Interrupt	External interrupt controller interface	—	`ea-4` *(2)*	Requested handler address *(3)*

Table 3–33. Nios II Exceptions (In Decreasing Priority Order) (Part 2 of 3)

Exception	Type	Available	Cause	Address	Vector
External maskable interrupt	Interrupt	External interrupt controller interface	2	`ea-4` *(2)*	Requested handler address *(3)*
Supervisor-only Instruction Address *(1)*	Instruction-related	MMU	9	`ea-4` *(2)*	General exception
Fast TLB Miss (instruction) *(1)*	Instruction-related	MMU	12	`pteaddr.VPN`, `ea-4` *(2)*	Fast TLB Miss exception
Double TLB Miss (instruction) *(1)*	Instruction-related	MMU	12	`pteaddr.VPN`, `ea-4` *(2)*	General exception
TLB Permission Violation (execute) *(1)*	Instruction-related	MMU	13	`pteaddr.VPN`, `ea-4` *(2)*	General exception
MPU Region Violation (instruction) *(1)*	Instruction-related	MPU	16	`ea-4` *(2)*	General exception
Supervisor-only Instruction	Instruction-related	MMU or MPU	10	`ea-4` *(2)*	General exception
Trap Instruction	Instruction-related	Always	3	`ea-4` *(2)*	General exception
Illegal Instruction	Instruction-related	Illegal instruction detection on, MMU, or MPU	5	`ea-4` *(2)*	General exception
Unimplemented Instruction	Instruction-related	Always	4	`ea-4` *(2)*	General exception
Break Instruction	Instruction-related	Always	—	`ba-4` *(2)*	Break
Supervisor-only Data Address	Instruction-related	MMU	11	`badaddr` (data address)	General exception
Misaligned Data Address	Instruction-related	Illegal memory access detection on, MMU, or MPU	6	`badaddr` (data address)	General exception
Misaligned Destination Address	Instruction-related	Illegal memory access detection on, MMU, or MPU	7	`badaddr` (destination address)	General exception
Division Error	Instruction-related	Division error detection on	8	`ea-4` *(2)*	General exception
Fast TLB Miss (data)	Instruction-related	MMU	12	`pteaddr.VPN`, `badaddr` (data address)	Fast TLB Miss exception
Double TLB Miss (data)	Instruction-related	MMU	12	`pteaddr.VPN`, `badaddr` (data address)	General exception
TLB Permission Violation (read)	Instruction-related	MMU	14	`pteaddr.VPN`, `badaddr` (data address)	General exception
TLB Permission Violation (write)	Instruction-related	MMU	15	`pteaddr.VPN`, `badaddr` (data address)	General exception

Table 3–33. Nios II Exceptions (In Decreasing Priority Order) (Part 3 of 3)

Exception	Type	Available	Cause	Address	Vector
MPU Region Violation (data)	Instruction-related	MPU	17	`badaddr` (data address)	General exception

Notes to Table 3–33:

(1) It is possible for any instruction fetch to cause this exception.

(2) Refer to Table 3–6 on page 3–11 for descriptions of the `ea` and `ba` registers.

(3) For a description of the requested handler address, refer to "Requested Handler Address" on page 3–36.

Exception Latency

Exception latency specifies how quickly the system can respond to an exception. Exception latency depends on the type of exception, the software and hardware configuration, and the processor state.

Interrupt Latency

The interrupt controller can mask individual interrupts. Each interrupt can have a different maximum masked time. The worst-case interrupt latency for interrupt *i* is determined by that interrupt's maximum masked time, or by the maximum disabled time, whichever is greater.

Reset Exceptions

When a processor reset signal is asserted, the Nios II processor performs the following steps:

1. Sets `status.RSIE` to 1, and clears all other fields of the `status` register.
2. Invalidates the instruction cache line associated with the reset vector.
3. Begins executing the reset handler, located at the reset vector.

All noninterrupt exception handlers must run in the normal register set.

Clearing the `status.PIE` field disables maskable interrupts. If the MMU or MPU is present, clearing the `status.U` field forces the processor into supervisor mode.

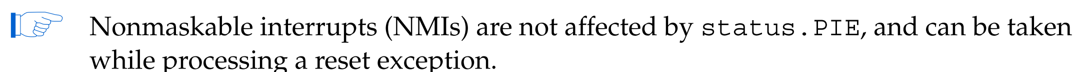
Nonmaskable interrupts (NMIs) are not affected by `status.PIE`, and can be taken while processing a reset exception.

Invalidating the reset cache line guarantees that instruction fetches for reset code comes from uncached memory.

Aside from the instruction cache line associated with the reset vector, the contents of the cache memories are indeterminate after reset. To ensure cache coherency after reset, the reset handler located at the reset vector must immediately initialize the instruction cache. Next, either the reset handler or a subsequent routine should proceed to initialize the data cache.

The reset state is undefined for all other system components, including but not limited to:

- General-purpose registers, except for `zero` (`r0`) in the normal register set, which is permanently zero.
- Control registers, except for `status`. `status.RSIE` is reset to 1, and the remaining fields are reset to 0.
- Instruction and data memory.
- Cache memory, except for the instruction cache line associated with the reset vector.
- Peripherals. Refer to the appropriate peripheral data sheet or specification for reset conditions.
- Custom instruction logic. Refer to the *Nios II Custom Instruction User Guide* for reset conditions.
- Nios II C-to-hardware (C2H) acceleration compiler logic.

Break Exceptions

A break is a transfer of control away from a program's normal flow of execution for the purpose of debugging. Software debugging tools can take control of the Nios II processor via the JTAG debug module.

Break processing is the means by which software debugging tools implement debug and diagnostic features, such as breakpoints and watchpoints. Break processing is a type of exception processing, but the break mechanism is independent from general exception processing. A break can occur during exception processing, enabling debug tools to debug exception handlers.

The processor enters the break processing state under either of the following conditions:

- The processor executes the `break` instruction. This is often referred to as a software break.
- The JTAG debug module asserts a hardware break.

Processing a Break

A break causes the processor to take the following steps:

1. Stores the contents of the `status` register to `bstatus`.
2. Clears `status.PIE` to zero, disabling maskable interrupts.

Nonmaskable interrupts (NMIs) are not affected by `status.PIE`, and can be taken while processing a break exception.

3. Writes the address of the instruction following the break to the `ba` register (`r30`) in the normal register set.
4. Clears `status.U` to zero, forcing the processor into supervisor mode, when the system contains an MMU or MPU.

5. Sets `status.EH` to one, indicating the processor is handling an exception, when the system contains an MMU.
6. Copies `status.CRS` to `status.PRS` and then sets `status.CRS` to 0.
7. Transfers execution to the break handler, stored at the break vector specified at system generation time.

All noninterrupt exception handlers, including the break handler, must run in the normal register set.

Register Usage

The `bstatus` control register and general-purpose registers `bt` (`r25`) and `ba` (`r30`) in the normal register set are reserved for debugging. Code is not prevented from writing to these registers, but debug code might overwrite the values. The break handler can use `bt` (`r25`) to help save additional registers.

Returning From a Break

After processing a break, the break handler releases control of the processor by executing a `bret` instruction. The `bret` instruction restores `status` by copying the contents of `bstatus` and returns program execution to the address in the `ba` register (`r30`) in the normal register set. Aside from `bt` and `ba`, all registers are guaranteed to be returned to their pre-break state after returning from the break handler.

Interrupt Exceptions

A peripheral device can request an interrupt by asserting an interrupt request (IRQ) signal. IRQs interface to the Nios II processor through an interrupt controller. You can configure the Nios II processor with either of the following interrupt controller options:

- The external interrupt controller interface
- The internal interrupt controller

External Interrupt Controller Interface

The Nios II EIC interface enables you to connect the Nios II processor to an external interrupt controller component. The EIC can monitor and prioritize IRQ signals, and determine which interrupt to present to the Nios II processor. An EIC can be software-configurable.

The Nios II processor does not depend on any particular implementation of an EIC. The degree of EIC configurability, and EIC configuration methods, are implementation-specific. This section discusses the EIC interface, and general features of EICs. For usage details, refer to the documentation for the specific EIC in your system.

For a typical EIC implementation, refer to the *Vectored Interrupt Controller* chapter in *Volume 5: Embedded Peripherals* of the *Quartus II Handbook*.

When an IRQ is asserted, the EIC presents the following information to the Nios II processor:

- The requested handler address (RHA)—Refer to "Requested Handler Address"
- The requested interrupt level (RIL)—Refer to "Requested Interrupt Level"
- The requested register set (RRS)—Refer to "Requested Register Set"
- Requested nonmaskable interrupt (RNMI) mode—Refer to "Requested NMI Mode"

The Nios II processor EIC interface connects to a single EIC, but an EIC can support a daisy-chained configuration. In a daisy-chained configuration, multiple EICs can monitor and prioritize interrupts. The EIC directly connected to the processor presents the processor with the highest-priority interrupt from all EICs in the daisy chain.

An EIC component can support an arbitrary level of daisy-chaining, potentially allowing the Nios II processor to handle an arbitrary number of prioritized interrupts.

Requested Handler Address

The RHA specifies the address of the handler associated with the interrupt. The availability of an RHA for each interrupt allows the Nios II processor to jump directly to the interrupt handler, reducing interrupt latency.

The RHA for each interrupt is typically software-configurable. The method for specifying the RHA is dependent on the specific EIC implementation.

If the Nios II processor is implemented with an MMU, the processor treats handler addresses as virtual addresses.

Requested Interrupt Level

The Nios II processor uses the RIL to decide when to take a maskable interrupt. The interrupt is taken only when the RIL is greater than `status.IL`.

The RIL is ignored for nonmaskable interrupts.

Requested Register Set

If shadow register sets are implemented on the Nios II core, the EIC specifies a register set when it asserts an interrupt request. When it takes the interrupt, the Nios II processor switches to the requested register set. When an interrupt has a dedicated register set, the interrupt handler avoids the overhead of saving registers.

The method of assigning register sets to interrupts depends on the specific EIC implementation. Register set assignments can be software-configurable.

Multiple interrupts can be configured to share a register set. In this case, the interrupt handlers must be written so as to avoid register corruption. For example, one of the following conditions must be true:

- The interrupts cannot pre-empt one another. For example, all interrupts are at the same level.
- Registers are saved in software. For example, each interrupt handler saves its own registers on entry, and restores them on exit.

Typically, the Nios II processor is configured so that when it takes an interrupt, other interrupts in the same register set are disabled. If interrupt preemption within a register set is desired, the interrupt handler can re-enable interrupts in its register set.

By default, the Nios II processor disables maskable interrupts when it takes an interrupt request. To enable nested interrupts, system software or the ISR itself must re-enable interrupts after the interrupt is taken.

Alternatively, to take full advantage of nested interrupts with shadow register sets, system software can set the `config.ANI` flag. When `config.ANI` = 1, the Nios II processor leaves maskable interrupts enabled after it takes an interrupt.

Requested NMI Mode

Any interrupt can be nonmaskable, depending on the configuration of the EIC. An NMI typically signals a critical system event requiring immediate handling, to ensure either system stability or real-time performance.

`status.IL` and RIL are ignored for nonmaskable interrupts.

Shadow Register Sets

Although shadow register sets can be implemented independently of the EIC interface, typically the two features are used together. Combining shadow register sets with an appropriate EIC, you can minimize or eliminate the context switch overhead for critical interrupts.

For the best interrupt performance, assign a dedicated register set to each of the most time-critical interrupts. Less-critical interrupts can share register sets, provided the ISRs are protected from register corruption as noted in “Requested Register Set”.

The method for mapping interrupts to register sets is specific to the particular EIC implementation.

Internal Interrupt Controller

When the internal interrupt controller is implemented, a peripheral device can request a hardware interrupt by asserting one of the Nios II processor’s 32 interrupt-request inputs, `irq0` through `irq31`. A hardware interrupt is generated if and only if all three of these conditions are true:

- The `PIE` bit of the `status` control register is one.
- An interrupt-request input, `irq`*n*, is asserted.
- The corresponding bit *n* of the `ienable` control register is one.

Upon hardware interrupt, the processor clears the `PIE` bit to zero, disabling further interrupts, and performs the other steps outlined in “Exception Processing Flow” on page 3–43.

The value of the `ipending` control register shows which interrupt requests (IRQ) are pending. By peripheral design, an IRQ bit is guaranteed to remain asserted until the processor explicitly responds to the peripheral. Figure 3–4 shows the relationship between `ipending`, `ienable`, `PIE`, and the generation of an interrupt.

Although shadow register sets can be implemented in any Nios II/f processor, the internal interrupt controller does not have features to take advantage of it as external interrupt controllers do.

Figure 3–4. Relationship Between ienable, ipending, PIE and Hardware Interrupts

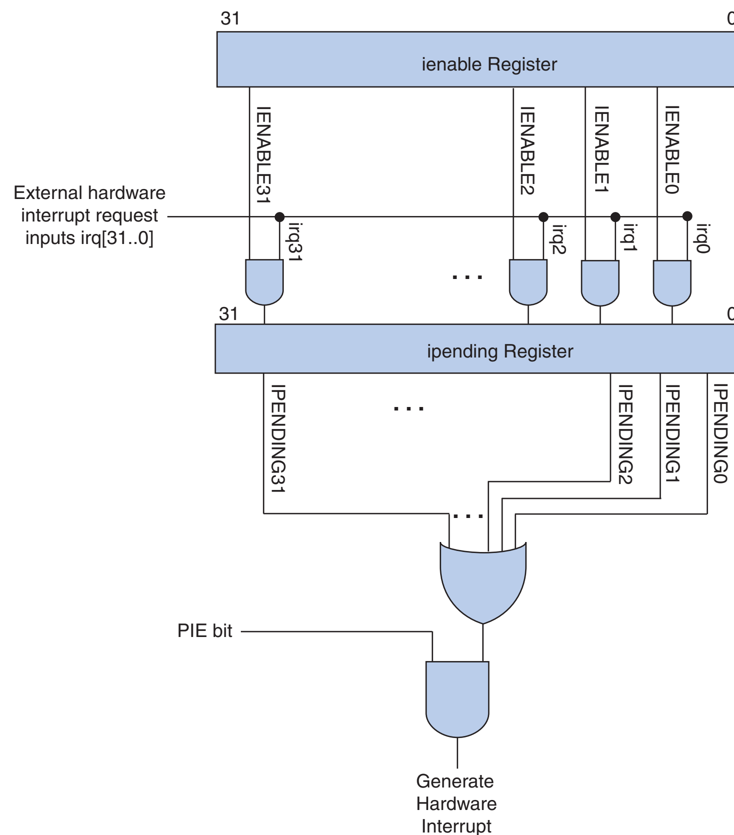

Instruction-Related Exceptions

Instruction-related exceptions occur during execution of Nios II instructions and perform the steps outlined in “Exception Processing Flow” on page 3–43.

The Nios II processor generates the following instruction-related exceptions.

- Trap instruction
- Break instruction
- Unimplemented instruction
- Illegal instruction
- Supervisor-only instruction
- Supervisor-only instruction address
- Supervisor-only data address
- Misaligned data address
- Misaligned destination address
- Division error
- Fast TLB miss
- Double TLB miss

- TLB permission violation
- MPU region violation

All noninterrupt exception handlers must run in the normal register set.

Trap Instruction

When a program issues the `trap` instruction, it generates a software trap exception. A program typically issues a software trap when the program requires servicing by the operating system. The general exception handler for the operating system determines the reason for the trap and responds appropriately.

Break Instruction

The break instruction is treated as a break exception. Refer to "Break Exceptions" on page 3–34 for details.

Unimplemented Instruction

When the processor issues a valid instruction that is not implemented in hardware, an unimplemented instruction exception is generated. The general exception handler determines which instruction generated the exception. If the instruction is not implemented in hardware, control is passed to an exception routine that might choose to emulate the instruction in software. Refer to "Potential Unimplemented Instructions" on page 3–58 for more information.

Illegal Instruction

Illegal instructions are instructions with an undefined opcode or opcode-extension field. The Nios II processor can check for illegal instructions and generate an exception when an illegal instruction is encountered. When your system contains an MMU or MPU, illegal instruction checking is always on. When no MMU or MPU is present, you have the option to have the processor check for illegal instructions.

To see how to control this option, refer to the *Instantiating the Nios II Processor in SOPC Builder* chapter of the *Nios II Processor Reference Handbook.*

When the processor issues an instruction with an undefined opcode or opcode-extension field, and illegal instruction exception checking is turned on, an illegal instruction exception is generated.

Refer to the OP Encodings and OPX Encodings for R-Type Instructions tables in the *Instruction Set Reference* chapter of the *Nios II Processor Reference Handbook* to see the unused opcodes and opcode extensions.

All undefined opcodes are reserved. The processor does occasionally use some undefined encodings internally. Executing one of these undefined opcodes does not trigger an illegal instruction exception. Refer to the *Nios II Core Implementation Details* chapter of the *Nios II Processor Reference Handbook* for details on each specific Nios II core.

Supervisor-only Instruction

When your system contains an MMU or MPU and the processor is in user mode (`status.U = 1`), executing a supervisor-only instruction results in a supervisor-only instruction exception. The supervisor-only instructions are `initd`, `initi`, `eret`, `bret`, `rdctl`, and `wrctl`.

This exception is implemented only in Nios II processors configured to use supervisor mode and user mode. Refer to "Operating Modes" on page 3–2 for more information.

Supervisor-only Instruction Address

When your system contains an MMU and the processor is in user mode (`status.U = 1`), attempts to access a supervisor-only instruction address result in a supervisor-only instruction address exception. Any instruction fetch can cause this exception. For definitions of supervisor-only address ranges, refer to Table 3–2 on page 3–5.

This exception is implemented only in Nios II processors that include the MMU.

Supervisor-only Data Address

When your system contains an MMU and the processor is in user mode (`status.U = 1`), any attempt to access a supervisor-only data address results in a supervisor-only data address exception. Instructions that can cause a supervisor-only data address exception are all loads, all stores, and `flushda`.

This exception is implemented only in Nios II processors that include the MMU.

Misaligned Data Address

The Nios II processor can check for misaligned data addresses of load and store instructions and generate an exception when a misaligned data address is encountered. When your system contains an MMU or MPU, misaligned data address checking is always on. When no MMU or MPU is present, you have the option to have the processor check for misaligned data addresses.

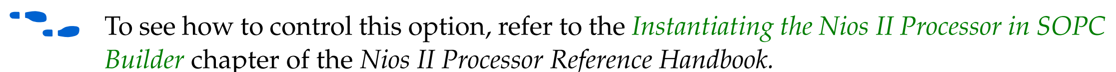
To see how to control this option, refer to the *Instantiating the Nios II Processor in SOPC Builder* chapter of the *Nios II Processor Reference Handbook.*

A data address is considered misaligned if the byte address is not a multiple of the width of the load or store instruction data width (four bytes for word, two bytes for half-word). Byte load and store instructions are always aligned so never take a misaligned address exception.

Misaligned Destination Address

The Nios II processor can check for misaligned destination addresses of the `callr`, `jmp`, `ret`, `eret`, `bret`, and all branch instructions and generate an exception when a misaligned destination address is encountered. When your system contains an MMU or MPU, misaligned destination address checking is always on. When no MMU or MPU is present, you have the option to have the processor check for misaligned destination addresses.

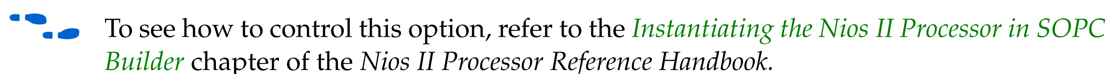
To see how to control this option, refer to the *Instantiating the Nios II Processor in SOPC Builder* chapter of the *Nios II Processor Reference Handbook.*

A destination address is considered misaligned if the target byte address of the instruction is not a multiple of four.

Division Error

The Nios II processor can check for division errors and generate an exception when a division error is encountered.

To see how to control this option, refer to the *Instantiating the Nios II Processor in SOPC Builder* chapter of the *Nios II Processor Reference Handbook.*

The division error exception detects divide instructions that produce a quotient that can't be represented. The two cases are divide by zero and a signed division that divides the largest negative number -2147483648 (0x80000000) by -1 (0xffffffff). Division error detection is only available if divide instructions are supported by hardware.

Fast TLB Miss

Fast TLB miss exceptions are implemented only in Nios II processors that include the MMU. The MMU has a special exception vector (fast TLB miss), specified in SOPC Builder at system generation time, specifically to handle TLB miss exceptions quickly. Whenever the processor cannot find a TLB entry matching the VPN (optionally extended by a process identifier), the result is a TLB miss exception. At the time of the exception, the processor first checks `status.EH`. When `status.EH` = 0, no other exception is already in process, so the processor considers the TLB miss a fast TLB miss, sets `status.EH` to one, and transfers control to the fast TLB miss exception handler (rather than to the general exception handler).

There are two kinds of fast TLB miss exceptions:

- Fast TLB miss (instruction)—Any instruction fetch can cause this exception.
- Fast TLB miss (data)—Load, store, `initda`, and `flushda` instructions can cause this exception.

The fast TLB miss exception handler can inspect the `tlbmisc.D` field to determine which kind of fast TLB miss exception occurred.

Double TLB Miss

Double TLB miss exceptions are implemented only in Nios II processors that include the MMU. When a TLB miss exception occurs while software is currently processing an exception (that is, when `status.EH` = 1), a double TLB miss exception is generated. Specifically, whenever the processor cannot find a TLB entry matching the VPN (optionally extended by a process identifier) and `status.EH` = `1`, the result is a double TLB miss exception. The most common scenario is that a double TLB miss exception occurs during processing of a fast TLB miss exception. The processor preserves register values from the original exception and transfers control to the general exception handler which processes the newly-generated exception.

There are two kinds of double TLB miss exceptions:

- Double TLB miss (instruction)—Any instruction fetch can cause this exception.
- Double TLB miss (data)—Load, store, `initda`, and `flushda` instructions can cause this exception.

The general exception handler can inspect either the `exception.CAUSE` or `tlbmisc.D` field to determine which kind of double TLB miss exception occurred.

TLB Permission Violation

TLB permission violation exceptions are implemented only in Nios II processors that include the MMU. When a TLB entry is found matching the VPN (optionally extended by a process identifier), but the permissions do not allow the access to complete, a TLB permission violation exception is generated.

There are three kinds of TLB permission violation exceptions:

- TLB permission violation (execute)—Any instruction fetch can cause this exception.
- TLB permission violation (read)—Any load instruction can cause this exception.
- TLB permission violation (write)—Any store instruction can cause this exception.

The general exception handler can inspect the `exception.CAUSE` field to determine which permissions were violated.

The data cache management instructions (`initd`, `initda`, `flushd`, and `flushda`) ignore the TLB permissions and do not generate TLB permission violation exceptions.

MPU Region Violation

MPU region violation exceptions are implemented only in Nios II processors that include the MPU. An MPU region violation exception is generated under any of the following conditions:

- An instruction fetch or data address matched a region but the permissions for that region did not allow the action to complete.
- An instruction fetch or data address did not match any region.

The general exception handler reads the MPU region attributes to determine if the address did not match any region or which permissions were violated.

There are two kinds of MPU region violation exceptions:

- MPU region violation (instruction)—Any instruction fetch can cause this exception.
- MPU region violation (data)—Load, store, `initda`, and `flushda` instructions can cause this exception.

The general exception handler can inspect the `exception.CAUSE` field to determine which kind of MPU region violation exception occurred.

Other Exceptions

The preceding sections describe all of the exception types defined by the Nios II architecture at the time of publishing. However, some processor implementations might generate exceptions that do not fall into the categories listed in the preceding sections. Therefore, a robust exception handler must provide a safe response (such as issuing a warning) in the event that it cannot identify the cause of an exception.

Exception Processing Flow

Except for the break exception (refer to "Processing a Break" on page 3–34), this section describes how the processor responds to exceptions, including interrupts and instruction-related exceptions.

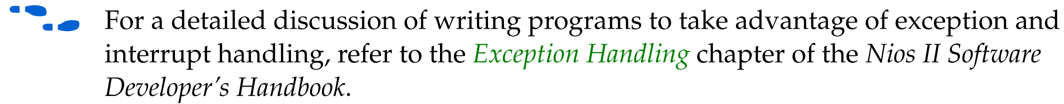

For a detailed discussion of writing programs to take advantage of exception and interrupt handling, refer to the *Exception Handling* chapter of the *Nios II Software Developer's Handbook*.

Processing General Exceptions

The general exception handler is a routine that determines the cause of each exception (including the double TLB miss exception), and then dispatches an exception routine to respond to the exception. The address of the general exception handler, specified in SOPC Builder at system generation time, is called the exception vector in the Nios II Megawizard interface. At run time this address is fixed, and software cannot modify it. Programmers do not directly access exception vectors, and can write programs without awareness of the address.

If the EIC interface is present, the general exception handler processes only noninterrupt exceptions.

The fast TLB miss exception handler only handles the fast TLB miss exception. It is built for speed to process TLB misses quickly. The fast TLB miss exception handler address, specified in SOPC Builder at system generation time, is called the fast TLB miss exception vector in the Nios II Megawizard interface.

Exception Flow with the EIC Interface

If the EIC interface is present, interrupt processing differs markedly from noninterrupt exception processing. The EIC interface provides the following information to the Nios II processor for each interrupt request:

- RHA—The requested handler address for the interrupt handler assigned to the requested interrupt.
- RRS—The requested register set to be used when the interrupt handler executes. If shadow register sets are not implemented, RRS must always be 0.
- RIL—The requested interrupt level specifies the priority of the interrupt.
- RNMI—The requested NMI flag specifies whether to treat the interrupt as nonmaskable.

For further information about the RHA, RRS, RIL and RNMI, refer to "The Nios II/f Core" in the *Nios II Core Implementation Details* chapter of the *Nios II Processor Reference Handbook*.

When the EIC interface presents an interrupt to the Nios II processor, the processor uses several criteria, as follows, to determine whether to take the interrupt:

- Nonmaskable interrupts—The processor takes any NMI as long as it is not processing a previous NMI.
- Maskable interrupts—The processor takes a maskable interrupt if maskable interrupts are enabled, and if the requested interrupt level is higher than that of the interrupt currently being processed (if any). However, if shadow register sets are implemented, the processor takes the interrupt only if the interrupt requests a register set different from the current register set, or if the register set interrupt enable flag (`status.RSIE`) is set.

Table 3–34 summarizes the conditions under which the Nios II processor takes an external interrupt.

Table 3–34. Conditions Required to Take External Interrupt

RNMI == 1		RNMI == 0					
status.NMI == 0	status.NMI == 1	status.PIE == 0	status.PIE == 1				
			RIL <= status.IL	RIL > status.IL			
				Processor Has Shadow Register Sets			No Shadow Register Sets
				RRS == status.CRS		RRS != status.CRS	
				status.RSIE == 0	status.RSIE == 1		
Yes	No	No	No	No *(1)*	Yes	Yes	Yes

Note to Table 3–34:

(1) Nested interrupts using the same register set are allowed only if system software has explicitly permitted them by setting `status.RSIE`. This restriction ensures that such interrupts are taken only if the handler is coded to save the register context.

The Nios II processor supports fast nested interrupts with shadow register sets, as described in “Shadow Register Sets” on page 3–26. When shadow register sets are implemented, the `config.ANI` field is set to 0 at reset.

Software must set `config.ANI` to 1 to enable fast nested interrupts. If `config.ANI` is set to 1 when a maskable external interrupt occurs, `status.PIE` not cleared. Leaving `status.PIE` set allows higher level interrupts to be taken immediate, without requiring the interrupt handler to set `status.PIE` to 1.

System software can disable fast nested interrupts by setting `config.ANI` to 0. In this state, the processor disables maskable interrupts when taking an exception, just as it does without shadow register sets. An individual interrupt handler can re-enable interrupts by setting `status.PIE` to 1, if desired.

Exception Flow with the Internal Interrupt Controller

A general exception handler determines which of the pending interrupts has the highest priority, and then transfers control to the appropriate ISR. The ISR stops the interrupt from being visible (either by clearing it at the source or masking it using `ienable`) before returning and/or before re-enabling `PIE`. The ISR also saves `estatus` and `ea` (`r29`) before re-enabling `PIE`.

Interrupts can be re-enabled by writing one to the `PIE` bit, thereby allowing the current ISR to be interrupted. Typically, the exception routine adjusts `ienable` so that IRQs of equal or lower priority are disabled before re-enabling interrupts. Refer to "Handling Nested Exceptions" on page 3–48 for more information.

Exceptions and Processor Status

Table 3–35 lists all changes to the Nios II processor state as a result of nonbreak exception processing actions performed by hardware. For systems with an MMU, `status.EH` indicates whether or not exception processing is already in progress. When `status.EH` = 1, exception processing is already in progress and the states of the exception registers are preserved to retain the original exception states.

Determining the Cause of Interrupt and Instruction-Related Exceptions

The general exception handler must determine the cause of each exception and then transfer control to an appropriate exception routine.

With Extra Exception Information

When you have included the extra exception information in your Nios II system, the `CAUSE` field of the `exception` register (refer to "The exception Register" on page 3–15) contains a code for the highest-priority exception occurring at the time and the `BADDR` field of the `badaddr` register (refer to "The badaddr Register" on page 3–20) contains the byte instruction address or data address for certain exceptions. Refer to Table 3–33 on page 3–31 for more information.

External interrupts do not set `exception.CAUSE`.

To determine the cause of an exception, simply read the cause of the exception from `exception.CAUSE` and then transfer control to the appropriate exception routine.

Extra exception information is always enabled in Nios II systems containing an MMU or MPU.

Without Extra Exception Information

When you have not included the extra exception information in your Nios II system, your exception handler must determine the cause of exception itself. For this reason, Altera recommends always enabling the extra exception information.

When the extra exception information is not available, use the sequence in Example 3–1 on page 3–47 to determine the cause of an exception.

Table 3–35. Nios II Processor Status After Taking Exception

Processor Status Register or Field	System Status Before Taking Exception							
	External Interrupt Asserted *(1)*				Internal Interrupt Asserted or Noninterrupt Exception			
	status.EH==1 *(2)*		status.EH==0		status.EH==1	status.EH==0		
						TLB Miss *(4)*	No TLB Miss	
	RRS==0 *(3)*	RRS!=0	RRS==0	RRS!=0			TLB Permission Violation *(4)*	No TLB Permission Violation
`pteaddr.VPN` *(5)*	No change					VPN *(6)*		No change
`status.PRS` *(3)*	No change		`status.CRS` *(3) (7)*		No change			
`pc`	RHA				General exception vector *(8)*	Fast TLB exception vector *(9)*	General exception vector *(3)*	
`sstatus` *(10) (11)*	No change			`status` *(7) (12)*	No change			
`estatus` *(11)*	No change		`status` *(7)*	No change			`status` *(7)*	
`ea`	No change		return address *(13)*		No change	return address		
`tlbmisc.D` *(2)*	No change					*(14)*		
`tlbmisc.DBL` *(2)*	No change					*(15)*		
`tlbmisc.PERM` *(2)*	No change					*(16)*		
`tlbmisc.BAD` *(2)*	No change					*(17)*		
`status.PIE`	`config.ANI` *(18)*				0 *(19)*			
`status.EH` *(2)*	No change				1 *(20)*			
`status.IH` *(21)*	1				No change			
`status.NMI` *(21)*	RNMI				No change			
`status.IL` *(21)*	RIL				No change			
`status.RSIE` *(3) (21)*	0				No change			
`status.CRS` *(3)*	RRS				No change			
`status.U` *(2)*	0 *(22)*							

Notes to Table 3–35: (Part 1 of 2)

(1) If the Nios II processor does not have an EIC interface, external interrupts do not occur.

(2) If the Nios II processor does not have an MMU, this field is not implemented. Its value is always 0, and the processor behaves accordingly.

(3) If the Nios II processor does not have shadow register sets, this field is not implemented. Its value is always 0, and the processor behaves accordingly.

(4) If the Nios II processor does not have an MMU, TLB-related exceptions do not occur.

(5) If the Nios II processor does not have an MMU, this register is not implemented.

(6) The VPN of the address triggering the exception

(7) The pre-exception value

© November 2009 Altera Corporation

Notes to Table 3–35: (Part 2 of 2)

(8) Invokes the general exception handler
(9) Invokes the fast TLB miss exception handler
(10) If the Nios II processor does not have shadow register sets, this register is not implemented.
(11) Saves the processor's pre-exception status
(12) `sstatus.SRS` is set to 1 if RRS is not equal to `status.CRS`.
(13) The address following the instruction being executed when the exception occurs
(14) Set to 1 on a data access exception, set to 0 otherwise
(15) Set to 1 on a double TLB miss, set to 0 otherwise
(16) Set to 1 on a TLB permission violation, set to 0 otherwise
(17) Set to 1 on a bad virtual address exception, set to 0 otherwise
(18) Disables exceptions and nonmaskable interrupts, unless automatic nested interrupts are explicitly enabled by `config.ANI`
(19) Disables exceptions and nonmaskable interrupts
(20) If the MMU is implemented, indicates that the processor is handling an exception.
(21) If the Nios II processor does not have an EIC interface, this field is not implemented.
(22) Puts the processor in supervisor mode.

Example 3–1. Determining Exception Cause Without Extra Exception Information

```
/* With an internal interrupt controller, check for interrupt
   exceptions. With an external interrupt controller, ipending is
   always 0, and this check can be omitted. */
if (estatus.PIE == 1 and ipending != 0) {
    handle interrupt

/* Decode exception from instruction */
/* Note: Because the exception register is included with the MMU and */
/* MPU, you never need to determine MMU or MPU exceptions by decoding */
} else {
    decode instruction at $ea-4
    if (instruction is trap)
        handle trap exception
    else if (instruction is load or store)
        handle misaligned data address exception
    else if (instruction is branch, bret, callr, eret, jmp, or ret)
        handle misaligned destination address exception
    else if (instruction is unimplemented)
        handle unimplemented instruction exception
    else if (instruction is illegal)
        handle illegal instruction exception
    else if (instruction is divide) {
        if (denominator == 0)
            handle division error exception
        else if (instruction is signed divide and numerator == 0x80000000
                                          and denominator == 0xffffffff)
            handle division error exception
        }
    }

    /* Not any known exception */
    } else {
        handle unknown exception (If internal interrupt controller
            is implemented, could be spurious interrupt)
    }
}
```

Handling Nested Exceptions

The Nios II processor supports several types of nested exceptions, depending on which optional features are implemented. Nested exceptions can occur under the following circumstances:

- An exception handler enables maskable interrupts
- An EIC is present, and an NMI occurs
- An EIC is present, and the processor is configured to leave maskable interrupts enabled when taking an interrupt
- An exception handler triggers an instruction-related exception

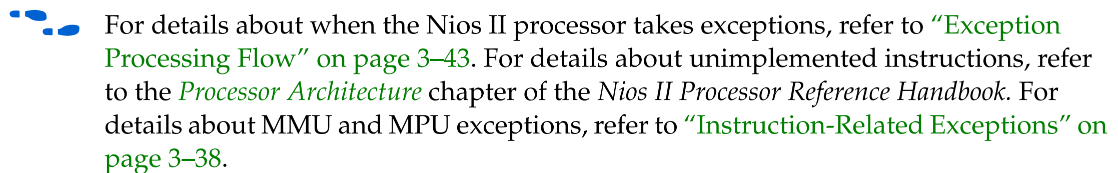

For details about when the Nios II processor takes exceptions, refer to “Exception Processing Flow” on page 3–43. For details about unimplemented instructions, refer to the *Processor Architecture* chapter of the *Nios II Processor Reference Handbook.* For details about MMU and MPU exceptions, refer to “Instruction-Related Exceptions” on page 3–38.

A system can be designed to eliminate the possibility of nested exceptions. However, if nested exceptions are possible, the exception handlers must be carefully written to prevent each handler from corrupting the context in which a pre-empted handler runs.

If an exception handler issues a `trap` instruction, an optional instruction, or an instruction which could generate an MMU or MPU exception, it must save and restore the contents of the `estatus` and `ea` registers.

Nested Exceptions with the Internal Interrupt Controller

You can enable nested exceptions in each exception handler on a case-by-case basis. If you want to allow a given exception handler to be pre-empted, set `status.PIE` to 1 near the beginning of the handler. Enabling maskable interrupts early in the handler minimizes the worst-case latency of any nested exceptions.

Ensure that all pre-empting handlers preserve the register contents.

Nested Exceptions with an External Interrupt Controller

With an EIC, handling of nested interrupts is more sophisticated than with the internal interrupt controller. Handling of noninterrupt exceptions, however, is the same.

When individual external interrupts have dedicated shadow register sets, the Nios II processor supports fast interrupt handling with no overhead for saving register contents. To take full advantage of fast interrupt handling, system software must set up the following conditions:

- Each interrupt is assigned to a dedicated shadow register set
- All interrupts with the same RIL are assigned to dedicated shadow register sets. That is, multiple interrupts with different RILs must not be assigned to the same shadow register set.
- Automatic nested interrupts are enabled (`config.ANI` is set to 1).

© November 2009 Altera Corporation

With these conditions satisfied, ISRs need not save and restore register contents on entry and exit.

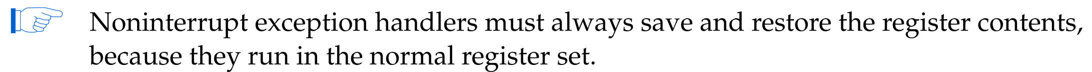
Noninterrupt exception handlers must always save and restore the register contents, because they run in the normal register set.

Multiple interrupts can share a register set, with some loss of performance. There are two techniques for sharing register sets:

- Set `status.RSIE` to 0. When an ISR is running in a given register set, the processor does not take any maskable interrupt assigned to the same register set. Such interrupts must wait for the running ISR to complete, regardless of their interrupt level.

  This technique can result in a priority inversion.

- Ensure that each ISR saves and restores registers on entry and exit, and set `status.RSIE` to 1 after registers are saved. When an ISR is running in a given register set, the processor takes an interrupt in the same register set if it has a higher interrupt level.

System software can globally disable fast nested interrupts by setting `config.ANI` to 0. In this state, the Nios II processor disables interrupts when taking a maskable interrupt (nonmaskable interrupts always disable maskable interrupts). Individual ISRs can re-enable nested interrupts by setting `status.PIE` to 1, as described in “Nested Exceptions with the Internal Interrupt Controller” on page 3–48.

Handing Nonmaskable Interrupts

Writing an NMI handler involves the same basic techniques as writing any other interrupt handler. However, nonmaskable interrupts always pre-empt maskable interrupts, and cannot be pre-empted. This can simplify handler design in some ways, but it means that an NMI handler can have a significant impact on overall interrupt latency. For the best system performance, perform the absolute minimum of processing in your NMI handlers, and defer less-critical processing to maskable interrupt handlers or foreground software.

NMIs leave intact the processor state associated with maskable interrupts and other exceptions, as well as normal, nonexception processing, provided each NMI is assigned to a dedicated shadow register set. Therefore NMIs can be handled transparently.

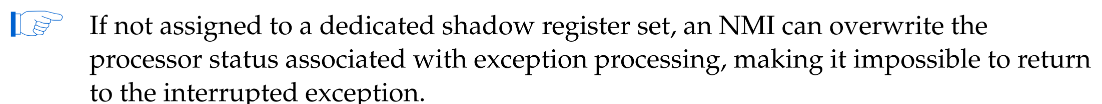
If not assigned to a dedicated shadow register set, an NMI can overwrite the processor status associated with exception processing, making it impossible to return to the interrupted exception.

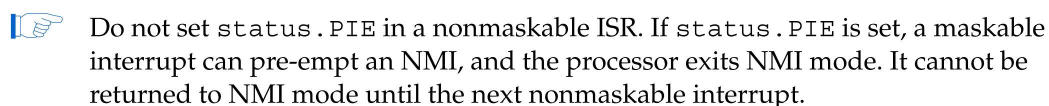
Do not set `status.PIE` in a nonmaskable ISR. If `status.PIE` is set, a maskable interrupt can pre-empt an NMI, and the processor exits NMI mode. It cannot be returned to NMI mode until the next nonmaskable interrupt.

Returning From Interrupt and Instruction-Related Exceptions

The `eret` instruction is used to resume execution at the pre-exception address.

You must ensure that when an exception handler modifies registers, they are restored when it returns. This can be taken care of in either of the following ways:

- In the case of ISRs, if the EIC interface and shadow register sets are implemented, and the ISR has a dedicated register set, no software action is required. The Nios II processor returns to the previous register set when it executes `eret`, which restores the register contents. For details, refer to "Nested Exceptions with an External Interrupt Controller".
- In the case of noninterrupt exceptions, for ISRs in a system with the internal interrupt controller, and for ISRs without a dedicated shadow register set, the exception handler must save registers on entry and restore them on exit. Saving the register contents on the stack is a typical, re-entrant implementation.

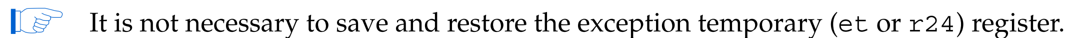

It is not necessary to save and restore the exception temporary (`et` or `r24`) register.

When executing the `eret` instruction, the processor performs the following tasks:

1. Restores the previous contents of `status` as follows:
   - If `status.CRS` is 0, copies `estatus` to `status`
   - If `status.CRS` is nonzero, copies `sstatus` to `status`
2. Transfers program execution to the address in the `ea` register (`r29`) in the register set specified by the original value of `status.CRS`.

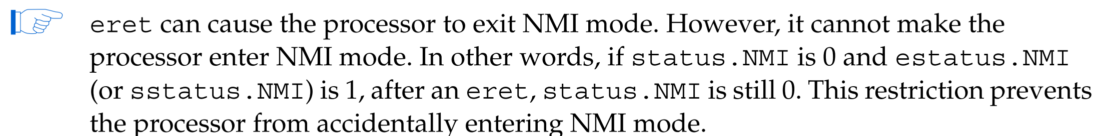

`eret` can cause the processor to exit NMI mode. However, it cannot make the processor enter NMI mode. In other words, if `status.NMI` is 0 and `estatus.NMI` (or `sstatus.NMI`) is 1, after an `eret`, `status.NMI` is still 0. This restriction prevents the processor from accidentally entering NMI mode.

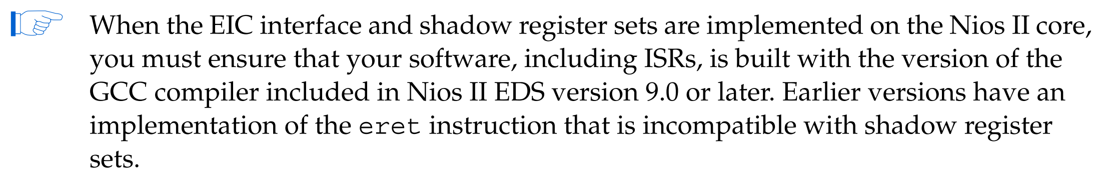

When the EIC interface and shadow register sets are implemented on the Nios II core, you must ensure that your software, including ISRs, is built with the version of the GCC compiler included in Nios II EDS version 9.0 or later. Earlier versions have an implementation of the `eret` instruction that is incompatible with shadow register sets.

Return Address Considerations

The return address requires some consideration when returning from exception processing routines. After an exception occurs, `ea` contains the address of the instruction following the point where the exception occurred.

When returning from instruction-related exceptions, execution must resume from the instruction following the instruction where the exception occurred. Therefore, `ea` contains the correct return address.

On the other hand, hardware interrupt exceptions must resume execution from the interrupted instruction itself. In this case, the exception handler must subtract 4 from `ea` to point to the interrupted instruction.

Masking and Disabling Exceptions

The Nios II processor provides several methods for temporarily turning off some or all exceptions from software. The available methods depend on the hardware configuration. This section discusses all potentially available methods.

Disabling Maskable Interrupts

Software can disable and enable maskable interrupts with the `status.PIE` bit. When `PIE` = 0, maskable interrupts are ignored. When `PIE` = 1, internal and maskable external interrupts can be taken, depending on the status of the interrupt controller.

Masking Interrupts with an External Interrupt Controller

Masking Individual Interrupts

Typical EIC implementations allow system software to mask individual interrupts. The method of masking individual interrupts is implementation-specific.

Interrupt Levels

The `status.IL` field controls what level of external maskable interrupts can be serviced. The processor services a maskable interrupt only if its requested interrupt level is greater than `status.IL`.

An ISR can make run-time adjustments to interrupt nesting by manipulating `status.IL`. For example, if an ISR is running at level 5, to temporarily allow pre-emption by another level 5 interrupt, it can set `status.IL` to 4.

To enable all external interrupts, set `status.IL` to 0. To disable all external interrupts, set `status.IL` to 63.

Masking Interrupts with the Internal Interrupt Controller

The `ienable` register controls the handling of internal hardware interrupts. Each bit of the `ienable` register corresponds to one of the interrupt inputs, `irq0` through `irq31`. A value of one in bit *n* means that the corresponding `irq`*n* interrupt is enabled; a bit value of zero means that the corresponding interrupt is disabled. Refer to "Exception Processing" on page 3–30 for more information.

An ISR can adjust `ienable` so that IRQs of equal or lower priority are disabled. Refer to "Handling Nested Exceptions" on page 3–48 for more information.

Memory and Peripheral Access

Nios II addresses are 32 bits, allowing access up to a 4 gigabyte address space. Nios II core implementations without MMUs restrict addresses to 31 bits or fewer. The MMU supports the full 32-bit physical address.

For details, refer to the *Nios II Core Implementation Details* chapter of the *Nios II Processor Reference Handbook*.

Peripherals, data memory, and program memory are mapped into the same address space. The locations of memory and peripherals within the address space are determined at system generation time. Reading or writing to an address that does not map to a memory or peripheral produces an undefined result.

The processor's data bus is 32 bits wide. Instructions are available to read and write byte, half-word (16-bit), or word (32-bit) data.

The Nios II architecture is little endian. For data wider than 8 bits stored in memory, the more-significant bits are located in higher addresses.

The Nios II architecture supports register+immediate addressing.

Cache Memory

The Nios II architecture and instruction set accommodate the presence of data cache and instruction cache memories. Cache management is implemented in software by using cache management instructions. Instructions are provided to initialize the cache, flush the caches whenever necessary, and to bypass the data cache to properly access memory-mapped peripherals.

The Nios II architecture provides the following mechanisms to bypass the cache:

- When no MMU is present, bit 31 of the address is reserved for bit-31 cache bypass. With bit-31 cache bypass, the address space of processor cores is 2 GBytes, and the high bit of the address controls the caching of data memory accesses.
- When the MMU is present, cacheability is controlled by the MMU, and bit 31 functions as a normal address bit. For details, refer to "Address Space and Memory Partitions" on page 3–4, and "TLB Organization" on page 3–6.
- Cache bypass instructions, such as `ldwio` and `stwio`.

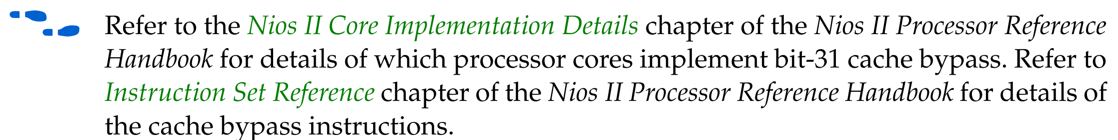

Refer to the *Nios II Core Implementation Details* chapter of the *Nios II Processor Reference Handbook* for details of which processor cores implement bit-31 cache bypass. Refer to *Instruction Set Reference* chapter of the *Nios II Processor Reference Handbook* for details of the cache bypass instructions.

Code written for a processor core with cache memory behaves correctly on a processor core without cache memory. The reverse is not true. If it is necessary for a program to work properly on multiple Nios II processor core implementations, the program must behave as if the instruction and data caches exist. In systems without cache memory, the cache management instructions perform no operation, and their effects are benign.

For a complete discussion of cache management, refer to the *Cache and Tightly Coupled Memory* chapter of the *Nios II Software Developer's Handbook*.

Some consideration is necessary to ensure cache coherency after processor reset. Refer to "Reset Exceptions" on page 3–33 for more information.

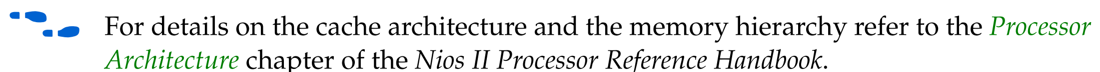

For details on the cache architecture and the memory hierarchy refer to the *Processor Architecture* chapter of the *Nios II Processor Reference Handbook*.

 © November 2009 Altera Corporation

Virtual Address Aliasing

A virtual address alias occurs when two virtual addresses map to the same physical address. When an MMU and caches are present and the caches are larger than a page (4 KBytes), the operating system must prevent illegal virtual address aliases. Because the caches are virtually-indexed and physically-tagged, a portion of the virtual address is used to select the cache line. If the cache is 4 KBytes or less in size, the portion of the virtual address used to select the cache line fits with bits 11:0 of the virtual address which have the same value as bits 11:0 of the physical address (they are untranslated bits of the page offset). However, if the cache is larger than 4 KBytes, bits beyond the page offset (bits 12 and up) are used to select the cache line and these bits are allowed to be different than the corresponding physical address.

For example, in a 64 KByte direct-mapped cache with a 16-byte line, bits 15:4 are used to select the line. Assume that virtual address `0x1000` is mapped to physical address `0xF000` and virtual address `0x2000` is also mapped to physical address `0xF000`. This is an illegal virtual address alias because accesses to virtual address `0x1000` use line 0x1 and accesses to virtual address `0x2000` use line 0x2 even though they map to the same physical address. This results in two copies of the same physical address in the cache. With an *n*-byte direct-mapped cache, there could be *n*/4096 copies of the same physical address in the cache if illegal virtual address aliases are not prevented. The bits of the virtual address that are used to select the line and are translated bits (bits 12 and up) are known as the color of the address. An operating system avoids illegal virtual address aliases by ensuring that if multiple virtual addresses map the same physical address, the virtual addresses have the same color. Note though, the color of the virtual addresses does not need to be the same as the color as the physical address because the cache tag contains all the bits of the PFN.

Instruction Set Categories

This section introduces the Nios II instructions categorized by type of operation performed.

Data Transfer Instructions

The Nios II architecture is a load-store architecture. Load and store instructions handle all data movement between registers, memory, and peripherals. Memories and peripherals share a common address space. Some Nios II processor cores use memory caching and/or write buffering to improve memory bandwidth. The architecture provides instructions for both cached and uncached accesses.

Table 3–36 describes the wide (32-bit) load and store instructions.

Table 3–36. Wide Data Transfer Instructions

Instruction	Description
`ldw` `stw`	The `ldw` and `stw` instructions load and store 32-bit data words from/to memory. The effective address is the sum of a register's contents and a signed immediate value contained in the instruction. Memory transfers can be cached or buffered to improve program performance. This caching and buffering might cause memory cycles to occur out of order, and caching might suppress some cycles entirely. Data transfers for I/O peripherals should use `ldwio` and `stwio`.
`ldwio` `stwio`	`ldwio` and `stwio` instructions load and store 32-bit data words from/to peripherals without caching and buffering. Access cycles for `ldwio` and `stwio` instructions are guaranteed to occur in instruction order and are never suppressed.

The data transfer instructions in Table 3–37 support byte and half-word transfers.

Table 3–37. Narrow Data Transfer Instructions

Instruction	Description

`ldb`
`ldbu`
`stb`
`ldh`
`ldhu`
`sth`
	`ldb`, `ldbu`, `ldh` and `ldhu` load a byte or half-word from memory to a register. `ldb` and `ldh` sign-extend the value to 32 bits, and `ldbu` and `ldhu` zero-extend the value to 32 bits.
	`stb` and `sth` store byte and half-word values, respectively.
	Memory accesses can be cached or buffered to improve performance. To transfer data to I/O peripherals, use the "io" versions of the instructions, described below.

`ldbio`
`ldbuio`
`stbio`
`ldhio`
`ldhuio`
`sthio`
	These operations load/store byte and half-word data from/to peripherals without caching or buffering.

Arithmetic and Logical Instructions

Logical instructions support `and`, `or`, `xor`, and `nor` operations. Arithmetic instructions support addition, subtraction, multiplication, and division operations. Refer to Table 3–38.

Table 3–38. Arithmetic and Logical Instructions (Part 1 of 2)

Instruction	Description

`and`
`or`
`xor`
`nor`
	These are the standard 32-bit logical operations. These operations take two register values and combine them bit-wise to form a result for a third register.

`andi`
`ori`
`xori`
	These operations are immediate versions of the `and`, `or`, and `xor` instructions. The 16-bit immediate value is zero-extended to 32 bits, and then combined with a register value to form the result.

`andhi`
`orhi`
`xorhi`
	In these versions of `and`, `or`, and `xor`, the 16-bit immediate value is shifted logically left by 16 bits to form a 32-bit operand. Zeroes are shifted in from the right.

Table 3–38. Arithmetic and Logical Instructions (Part 2 of 2)

Instruction	Description

`add`
`sub`
`mul`
`div`
`divu`
	These are the standard 32-bit arithmetic operations. These operations take two registers as input and store the result in a third register.

`addi`
`subi`
`muli`
	These instructions are immediate versions of the `add`, `sub`, and `mul` instructions. The instruction word includes a 16-bit signed value.

`mulxss` `mulxuu`	These instructions provide access to the upper 32 bits of a 32x32 multiplication operation. Choose the appropriate instruction depending on whether the operands should be treated as signed or unsigned values. It is not necessary to precede these instructions with a `mul`.
`mulxsu`	This instruction is used in computing a 128-bit result of a 64x64 signed multiplication.

Move Instructions

These instructions provide move operations to copy the value of a register or an immediate value to another register. Refer to Table 3–39.

Table 3–39. Move Instructions

Instruction	Description

`mov`
`movhi`
`movi`
`movui`
`movia`
	`mov` copies the value of one register to another register. `movi` moves a 16-bit signed immediate value to a register, and sign-extends the value to 32 bits. `movui` and `movhi` move an immediate 16-bit value into the lower or upper 16-bits of a register, inserting zeros in the remaining bit positions. Use `movia` to load a register with an address.

Comparison Instructions

The Nios II architecture supports a number of comparison instructions. All of these compare two registers or a register and an immediate value, and write either one (if true) or zero to the result register. These instructions perform all the equality and relational operators of the C programming language. Refer to Table 3–40.

Table 3–40. Comparison Instructions (Part 1 of 2)

Instruction	Description
`cmpeq`	==
`cmpne`	!=
`cmpge`	signed >=
`cmpgeu`	unsigned >=
`cmpgt`	signed >
`cmpgtu`	unsigned >
`cmple`	unsigned <=
`cmpleu`	unsigned <=
`cmplt`	signed <

Table 3–40. Comparison Instructions (Part 2 of 2)

Instruction	Description
`cmpltu`	unsigned <

`cmpeqi`
`cmpnei`
`cmpgei`
`cmpgeui`
`cmpgti`
`cmpgtui`
`cmplei`
`cmpleui`
`cmplti`
`cmpltui`
	These instructions are immediate versions of the comparison operations. They compare the value of a register and a 16-bit immediate value. Signed operations sign-extend the immediate value to 32-bits. Unsigned operations fill the upper bits with zero.

Shift and Rotate Instructions

The following instructions provide shift and rotate operations. The number of bits to rotate or shift can be specified in a register or an immediate value. Refer to Table 3–41.

Table 3–41. Shift and Rotate Instructions

Instructio n	Description

`rol`
`ror`
`roli`
	The `rol` and `roli` instructions provide left bit-rotation. `roli` uses an immediate value to specify the number of bits to rotate. The `ror` instructions provides right bit-rotation.
	There is no immediate version of `ror`, because `roli` can be used to implement the equivalent operation.

`sll`
`slli`
`sra`
`srl`
`srai`
`srli`
	These shift instructions implement the `<<` and `>>` operators of the C programming language. The `sll`, `slli`, `srl`, `srli` instructions provide left and right logical bit-shifting operations, inserting zeros. The `sra` and `srai` instructions provide arithmetic right bit-shifting, duplicating the sign bit in the most significant bit. `slli`, `srli` and `srai` use an immediate value to specify the number of bits to shift.

Program Control Instructions

The Nios II architecture supports the unconditional jump and call instructions listed in Table 3–42. These instructions do not have delay slots.

Table 3–42. Unconditional Jump and Call Instructions (Part 1 of 2)

Instruction	Description
`call`	This instruction calls a subroutine using an immediate value as the subroutine's absolute address, and stores the return address in register `ra`.
`callr`	This instruction calls a subroutine at the absolute address contained in a register, and stores the return address in register `ra`. This instruction serves the roll of dereferencing a C function pointer.
`ret`	The `ret` instruction is used to return from subroutines called by `call` or `callr`. `ret` loads and executes the instruction specified by the address in register `ra`.
`jmp`	The `jmp` instruction jumps to an absolute address contained in a register. `jmp` is used to implement switch statements of the C programming language.

 © November 2009 Altera Corporation

Table 3–42. Unconditional Jump and Call Instructions (Part 2 of 2)

Instruction	Description
`jmpi`	The `jmpi` instruction jumps to an absolute address using an immediate value to determine the absolute address.
`or`	This instruction branches relative to the current instruction. A signed immediate value gives the offset of the next instruction to execute.

The conditional-branch instructions compare register values directly, and branch if the expression is true. Refer to Table 3–43. The conditional branches support the equality and relational comparisons of the C programming language:

- == and !=
- < and <= (signed and unsigned)
- > and >= (signed and unsigned)

The conditional-branch instructions do not have delay slots.

Table 3–43. Conditional-Branch Instructions

Instruction	Description

`bge`
`bgeu`
`bgt`
`bgtu`
`ble`
`bleu`
`blt`
`bltu`
`beq`
`bne`
	These instructions provide relative branches that compare two register values and branch if the expression is true. Refer to "Comparison Instructions" on page 3–55 for a description of the relational operations implemented.

Other Control Instructions

Table 3–44 shows other control instructions.

Table 3–44. Other Control Instructions (Part 1 of 2)

Instruction	Description
`trap` `eret`	The `trap` and `eret` instructions generate and return from exceptions. These instructions are similar to the `call`/`ret` pair, but are used for exceptions. `trap` saves the `status` register in the `estatus` register, saves the return address in the `ea` register, and then transfers execution to the general exception handler. `eret` returns from exception processing by restoring `status` from `estatus`, and executing the instruction specified by the address in `ea`.
`break` `bret`	The `break` and `bret` instructions generate and return from breaks. `break` and `bret` are used exclusively by software debugging tools. Programmers never use these instructions in application code.
`rdctl` `wrctl`	These instructions read and write control registers, such as the `status` register. The value is read from or stored to a general-purpose register.

Table 3–44. Other Control Instructions (Part 2 of 2)

Instruction	Description
flushd flushda flushi initd initda initi	These instructions are used to manage the data and instruction cache memories.
flushp	This instruction flushes all prefetched instructions from the pipeline. This is necessary before jumping to recently-modified instruction memory.
sync	This instruction ensures that all previously-issued operations have completed before allowing execution of subsequent load and store operations.
rdprs wrprs	These instructions read and write a general-purpose registers between the current register set and another register set. wrprs can set r0 to 0 in a shadow register set. System software must use wrprs to initialize r0 to 0 in each shadow register set before using that register set.

Custom Instructions

The custom instruction provides low-level access to custom instruction logic. The inclusion of custom instructions is specified in SOPC Builder at system generation time, and the function implemented by custom instruction logic is design dependent.

For further details, refer to the "Custom Instructions" section of the *Processor Architecture* chapter of the *Nios II Processor Reference Handbook* and the *Nios II Custom Instruction User Guide*.

Machine-generated C functions and assembly macros provide access to custom instructions, and hide implementation details from the user. Therefore, most software developers never use the custom assembly instruction directly.

No-Operation Instruction

The Nios II assembler provides a no-operation instruction, nop.

Potential Unimplemented Instructions

Some Nios II processor cores do not support all instructions in hardware. In this case, the processor generates an exception after issuing an unimplemented instruction. Only the following instructions can generate an unimplemented instruction exception:

- mul
- muli
- mulxss
- mulxsu
- mulxuu
- div

- `divu`
- `initda`

All other instructions are guaranteed not to generate an unimplemented instruction exception.

An exception routine must exercise caution if it uses these instructions, because they could generate another exception before the previous exception is properly handled. Refer to "Unimplemented Instruction" on page 3–39 for more information regarding unimplemented instruction processing.

Referenced Documents

This chapter references the following documents:

- *Nios II Software Developer's Handbook*
- *Processor Architecture* chapter of the *Nios II Processor Reference Handbook*
- *Application Binary Interface* chapter of the *Nios II Processor Reference Handbook*
- *Instruction Set Reference* chapter of the *Nios II Processor Reference Handbook*
- *Instantiating the Nios II Processor in SOPC Builder* chapter of the *Nios II Processor Reference Handbook*
- *Nios II Core Implementation Details* chapter of the *Nios II Processor Reference Handbook*
- *Exception Handling* chapter of the *Nios II Software Developer's Handbook*
- *Cache and Tightly Coupled Memory* chapter of the *Nios II Software Developer's Handbook*
- *Vectored Interrupt Controller* chapter in *Volume 5: Embedded Peripherals* of the *Quartus II Handbook*
- *Nios II Custom Instruction User Guide*

Document Revision History

Table 3–45 shows the revision history for this document.

Table 3–45. Document Revision History (Part 1 of 2)

Date & Document Version	Changes Made	Summary of Changes
November 2009 v9.1.0	■ Added external interrupt controller interface information. ■ Added shadow register set information.	Added shadow register sets and external interrupt controller support
March 2009 v9.0.0	Maintenance release.	—
November 2008 v8.1.0	Maintenance release.	—
May 2008 v8.0.0	Added text to describe the MMU, MPU, and advanced exceptions.	Added MMU, MPU, and advanced exceptions.

Table 3–45. Document Revision History (Part 2 of 2)

Date & Document Version	Changes Made	Summary of Changes
October 2007 v7.2.0	■ Reworked text to refer to break and reset as exceptions. ■ Grouped exceptions, break, reset, and interrupts all under Exception Processing. ■ Added table showing all Nios II exceptions (by priority). ■ Removed "ctl" references to control registers. ■ Added `jmpi` instruction to tables.	—
May 2007 v7.1.0	■ Added table of contents to Introduction section. ■ Added Referenced Documents section.	—
March 2007 v7.0.0	Maintenance release.	—
November 2006 v6.1.0	Maintenance release.	—
May 2006 v6.0.0	Maintenance release.	—
October 2005 v5.1.0	Maintenance release.	—
May 2005 v5.0.0	Maintenance release.	—
September 2004 v1.1	■ Added details for new control register `ctl5`. ■ Updated details of debug and break processing to reflect new behavior of the `break` instruction.	—
May 2004 v1.0	Initial release.	—

© November 2009 Altera Corporation