



Nios II Processor Reference Guide

Updated for Intel® Quartus® Prime Design Suite: **20.1**



Subscribe

Send Feedback

NII-PRG | 2020.04.24

Latest document on the web: [PDF](#) | [HTML](#)



Contents

- 1. Introduction..... 8**
 - 1.1. Nios II Processor System Basics..... 8
 - 1.2. Getting Started with the Nios II Processor.....9
 - 1.3. Customizing Nios II Processor Designs.....10
 - 1.4. Configurable Soft Processor Core Concepts.....11
 - 1.4.1. Configurable Soft Processor Core..... 11
 - 1.4.2. Flexible Peripheral Set and Address Map.....11
 - 1.4.3. Automated System Generation..... 12
 - 1.5. Intel FPGA IP Evaluation Mode.....13
 - 1.6. Introduction Revision History.....13

- 2. Processor Architecture..... 14**
 - 2.1. Processor Implementation.....15
 - 2.2. Register File..... 16
 - 2.3. Arithmetic Logic Unit..... 17
 - 2.3.1. Unimplemented Instructions.....17
 - 2.3.2. Custom Instructions..... 17
 - 2.4. Reset and Debug Signals..... 18
 - 2.5. Exception and Interrupt Controllers..... 18
 - 2.5.1. Exception Controller..... 19
 - 2.5.2. EIC Interface.....19
 - 2.5.3. Internal Interrupt Controller..... 20
 - 2.6. Memory and I/O Organization.....20
 - 2.6.1. Instruction and Data Buses..... 21
 - 2.6.2. Cache Memory.....23
 - 2.6.3. Tightly-Coupled Memory..... 25
 - 2.6.4. Address Map..... 26
 - 2.6.5. Memory Management Unit.....26
 - 2.6.6. Memory Protection Unit..... 27
 - 2.7. JTAG Debug Module..... 27
 - 2.7.1. JTAG Target Connection..... 28
 - 2.7.2. Download and Execute Software..... 28
 - 2.7.3. Software Breakpoints..... 29
 - 2.7.4. Hardware Breakpoints..... 29
 - 2.7.5. Hardware Triggers.....29
 - 2.7.6. Trace Capture.....30
 - 2.8. Processor Architecture Revision History..... 31


- 3. Programming Model..... 33**
 - 3.1. Operating Modes.....33
 - 3.1.1. Supervisor Mode..... 33
 - 3.1.2. User Mode..... 34
 - 3.2. Memory Management Unit..... 34
 - 3.2.1. Recommended Usage..... 34
 - 3.2.2. Memory Management..... 35
 - 3.2.3. Address Space and Memory Partitions.....36
 - 3.2.4. TLB Organization..... 38
 - 3.2.5. TLB Lookups.....39



3.3. Memory Protection Unit.....	39
3.3.1. Memory Regions.....	40
3.3.2. Overlapping Regions.....	41
3.3.3. Enabling the MPU.....	42
3.4. Registers.....	42
3.4.1. General-Purpose Registers.....	42
3.4.2. Control Registers.....	43
3.4.3. Shadow Register Sets.....	62
3.5. Working with the MPU.....	65
3.5.1. MPU Region Read and Write Operations.....	65
3.5.2. MPU Initialization.....	66
3.5.3. Debugger Access.....	67
3.6. Working with ECC.....	67
3.6.1. Enabling ECC.....	67
3.6.2. Handling ECC Errors.....	67
3.6.3. Injecting ECC Errors.....	68
3.7. Exception Processing.....	71
3.7.1. Terminology.....	71
3.7.2. Exception Overview.....	72
3.7.3. Exception Latency.....	74
3.7.4. Reset Exceptions.....	74
3.7.5. Break Exceptions.....	75
3.7.6. Interrupt Exceptions.....	77
3.7.7. Instruction-Related Exceptions.....	80
3.7.8. Other Exceptions.....	85
3.7.9. Exception Processing Flow.....	85
3.7.10. Determining the Cause of Interrupt and Instruction-Related Exceptions.....	90
3.7.11. Handling Nested Exceptions.....	91
3.7.12. Handling Nonmaskable Interrupts.....	93
3.7.13. Masking and Disabling Exceptions.....	93
3.8. Memory and Peripheral Access.....	95
3.8.1. Cache Memory.....	96
3.9. Instruction Set Categories.....	97
3.9.1. Data Transfer Instructions.....	97
3.9.2. Arithmetic and Logical Instructions.....	98
3.9.3. Move Instructions.....	98
3.9.4. Comparison Instructions.....	99
3.9.5. Shift and Rotate Instructions.....	99
3.9.6. Program Control Instructions.....	100
3.9.7. Other Control Instructions.....	101
3.9.8. Custom Instructions.....	101
3.9.9. No-Operation Instruction.....	102
3.9.10. Potential Unimplemented Instructions.....	102
3.10. Programming Model Revision History.....	102
4. Instantiating the Nios II Processor.....	103
4.1. Main Nios II Tab.....	103
4.2. Vectors Tab.....	104
4.2.1. Reset Vector.....	104
4.2.2. Exception Vector.....	105
4.2.3. Fast TLB Miss Exception Vector.....	105



- 4.3. Caches and Memory Interfaces Tab.....106
 - 4.3.1. Instruction Cache.....107
 - 4.3.2. Flash Accelerator.....108
 - 4.3.3. Data Cache.....108
 - 4.3.4. Tightly-coupled Memories.....108
 - 4.3.5. Peripheral Region.....109
- 4.4. Arithmetic Instructions Tab.....109
 - 4.4.1. Arithmetic Instructions.....109
 - 4.4.2. Arithmetic Implementation.....110
- 4.5. MMU and MPU Settings Tab.....110
 - 4.5.1. MMU.....111
 - 4.5.2. MPU.....112
- 4.6. JTAG Debug Tab.....112
- 4.7. Advanced Features Tab.....114
 - 4.7.1. ECC.....114
 - 4.7.2. Interrupt Controller Interfaces.....115
 - 4.7.3. Shadow Register Sets.....115
 - 4.7.4. Reset Signals.....115
 - 4.7.5. CPU ID Control Register Value.....116
 - 4.7.6. Generate Trace File.....116
 - 4.7.7. Exception Checking.....116
 - 4.7.8. Branch Prediction.....117
 - 4.7.9. RAM Memory Protection.....117
- 4.8. The Quartus Prime IP File.....117
- 4.9. Instantiating the Nios II Processor Revision History.....117
- 5. Nios II Core Implementation Details.....118**
 - 5.1. Device Family Support.....118
 - 5.2. Nios II/f Core.....119
 - 5.2.1. Overview.....119
 - 5.2.2. Arithmetic Logic Unit.....120
 - 5.2.3. Memory Access.....121
 - 5.2.4. Tightly-Coupled Memory.....124
 - 5.2.5. Memory Management Unit.....125
 - 5.2.6. Memory Protection Unit.....125
 - 5.2.7. Execution Pipeline.....126
 - 5.2.8. Instruction Performance.....127
 - 5.2.9. Exception Handling.....128
 - 5.2.10. ECC.....130
 - 5.2.11. JTAG Debug Module.....131
 - 5.3. Nios II/s Core.....132
 - 5.3.1. Overview.....132
 - 5.3.2. Arithmetic Logic Unit.....132
 - 5.3.3. Memory Access.....134
 - 5.3.4. Tightly-Coupled Memory.....135
 - 5.3.5. Execution Pipeline.....135
 - 5.3.6. Instruction Performance.....136
 - 5.3.7. Exception Handling.....137
 - 5.3.8. JTAG Debug Module.....137
 - 5.4. Nios II/e Core.....137
 - 5.4.1. Overview.....138



5.4.2. Arithmetic Logic Unit.....	138
5.4.3. Memory Access.....	138
5.4.4. Instruction Execution Stages.....	138
5.4.5. Instruction Performance.....	139
5.4.6. Exception Handling.....	139
5.4.7. JTAG Debug Module.....	139
5.5. Nios II Core Implementation Details Revision History.....	140
6. Nios II Processor Versions.....	141
6.1. Nios II Versions Revision History.....	141
6.2. Architecture Revisions.....	141
6.3. Core Revisions.....	142
6.3.1. Nios II/f Core.....	142
6.3.2. Nios II/s Core.....	142
6.3.3. Nios II/e Core.....	142
6.4. JTAG Debug Module Revisions.....	142
6.5. Nios II Processor Versions Revision History.....	142
7. Application Binary Interface.....	143
7.1. Data Types.....	143
7.2. Memory Alignment.....	143
7.3. Register Usage.....	144
7.4. Stacks.....	145
7.4.1. Frame Pointer Elimination.....	146
7.4.2. Call Saved Registers.....	146
7.4.3. Further Examples of Stacks.....	146
7.4.4. Function Prologues.....	148
7.5. Arguments and Return Values.....	149
7.5.1. Arguments.....	150
7.5.2. Return Values.....	150
7.6. DWARF-2 Definition.....	151
7.7. Object Files.....	151
7.8. Relocation.....	151
7.9. ABI for Linux Systems.....	154
7.9.1. Linux Toolchain Relocation Information.....	154
7.9.2. Linux Function Calls.....	157
7.9.3. Linux Operating System Call Interface.....	158
7.9.4. Linux Process Initialization.....	159
7.9.5. Linux Position-Independent Code.....	159
7.9.6. Linux Program Loading and Dynamic Linking.....	161
7.9.7. Linux Conventions.....	164
7.10. Development Environment.....	164
7.11. Application Binary Interface Revision History.....	165
8. Instruction Set Reference.....	166
8.1. Word Formats.....	166
8.1.1. I-Type.....	166
8.1.2. R-Type.....	166
8.1.3. J-Type.....	167
8.2. Instruction Opcodes.....	167
8.3. Assembler Pseudo-Instructions.....	169
8.4. Assembler Macros.....	169



8.5. Instruction Set Reference..... 170

- 8.5.1. add 171
- 8.5.2. addi 172
- 8.5.3. and 173
- 8.5.4. andhi 174
- 8.5.5. andi 174
- 8.5.6. beq 175
- 8.5.7. bge 175
- 8.5.8. bgeu..... 176
- 8.5.9. bgt..... 176
- 8.5.10. bgtu 177
- 8.5.11. ble 177
- 8.5.12. bleu 177
- 8.5.13. blt 178
- 8.5.14. bltu..... 178
- 8.5.15. bne 179
- 8.5.16. br 179
- 8.5.17. break 180
- 8.5.18. bret 181
- 8.5.19. call 181
- 8.5.20. callr 182
- 8.5.21. cmpeq 182
- 8.5.22. cmpeqi 183
- 8.5.23. cmpge 183
- 8.5.24. cmpgei 184
- 8.5.25. cmpgeu 185
- 8.5.26. cmpgeui 185
- 8.5.27. cmpgt 186
- 8.5.28. cmpgti 186
- 8.5.29. cmpgtu 186
- 8.5.30. cmpgtui 187
- 8.5.31. cmple..... 187
- 8.5.32. cmplei 187
- 8.5.33. cmpleu 188
- 8.5.34. cmpleui 188
- 8.5.35. cmplt 188
- 8.5.36. cmplti 189
- 8.5.37. cmpltu 190
- 8.5.38. cmpltui 190
- 8.5.39. cmpne 191
- 8.5.40. cmpnei 191
- 8.5.41. custom 192
- 8.5.42. div 193
- 8.5.43. divu 194
- 8.5.44. eret 194
- 8.5.45. flushd 195
- 8.5.46. flushda 196
- 8.5.47. flushi 198
- 8.5.48. flushp 198
- 8.5.49. initd 199
- 8.5.50. initda 200



8.5.51. initi	201
8.5.52. jmp	202
8.5.53. jmpj	202
8.5.54. ldb / ldbio	203
8.5.55. ldbu / ldbuio	204
8.5.56. ldh / ldhio	205
8.5.57. ldhu / ldhuio	206
8.5.58. ldw / ldwio	207
8.5.59. mov	208
8.5.60. movhi	208
8.5.61. movi	209
8.5.62. movia	209
8.5.63. movui	209
8.5.64. mul	210
8.5.65. muli	211
8.5.66. mulxss	211
8.5.67. mulxsu	212
8.5.68. mulxuu	213
8.5.69. nextpc	214
8.5.70. nop	214
8.5.71. nor	214
8.5.72. or	215
8.5.73. orhi	215
8.5.74. ori	216
8.5.75. rdctl	216
8.5.76. rdprs	217
8.5.77. ret	217
8.5.78. rol	218
8.5.79. roli	218
8.5.80. ror	219
8.5.81. sll	219
8.5.82. slli	220
8.5.83. sra	220
8.5.84. srai	221
8.5.85. srl	221
8.5.86. srli	222
8.5.87. stb / stbio l.....	222
8.5.88. sth / sthio	223
8.5.89. stw / stwio	224
8.5.90. sub.....	225
8.5.91. subi	226
8.5.92. sync	226
8.5.93. trap	227
8.5.94. wrctl	228
8.5.95. wrprs	228
8.5.96. xor	229
8.5.97. xorhi	229
8.5.98. xori	230
8.6. Instruction Set Reference Revision History.....	230

1. Introduction

This handbook describes the Nios[®] II processor from a high-level conceptual description to the low-level details of implementation. The chapters in this handbook describe the Nios II processor architecture, the programming model, and the instruction set. The Nios II processor is only available in the Intel[®] Quartus[®] Prime 14.1 release and above.

We have ended development of new Nios II Classic processor features with the Intel Quartus Prime 14.0 release. New features are implemented only in the Nios II processor core. Although the Classic processor remains supported, we recommend that you use the Nios II core for future designs.

This handbook assumes you have a basic familiarity with embedded processor concepts. You do not need to be familiar with any specific Intel FPGA technology or with Intel FPGA development tools. This handbook limits discussion of hardware implementation details of the processor system. The Nios II processors are designed for Intel FPGA devices, and so this handbook does describe some FPGA implementation concepts. Your familiarity with FPGA technology provides a deeper understanding of the engineering trade-offs related to the design and implementation of the Nios II processor.

This chapter introduces the Intel FPGA Nios II embedded processor family and describes the similarities and differences between the Nios II processor and traditional embedded processors.

Related Information

[Nios II Processor webpage](#)

1.1. Nios II Processor System Basics

The Nios II processor is a general-purpose RISC processor core with the following features:

- Full 32-bit instruction set, data path, and address space
- 32 general-purpose registers
- Optional shadow register sets
- 32 interrupt sources
- External interrupt controller interface for more interrupt sources
- Single-instruction 32 × 32 multiply and divide producing a 32-bit result
- Dedicated instructions for computing 64-bit and 128-bit products of multiplication
- Optional floating-point instructions for single-precision floating-point operations
- Single-instruction barrel shifter



- Access to a variety of on-chip peripherals, and interfaces to off-chip memories and peripherals
- Hardware-assisted debug module enabling processor start, stop, step, and trace under control of the Nios II software development tools
- Optional memory management unit (MMU) to support operating systems that require MMUs
- Optional memory protection unit (MPU)
- Software development environment based on the GNU C/C++ tool chain and the Nios II Software Build Tools (SBT) for Eclipse
- Integration with Intel FPGA's Signal Tap II* Embedded Logic Analyzer, enabling real-time analysis of instructions and data along with other signals in the FPGA design
- Instruction set architecture (ISA) compatible across all Nios II processor systems
- Performance up to 250 DMIPS
- Error correcting code (ECC) support for all Nios II processor internal RAM blocks

A Nios II processor system is equivalent to a microcontroller or "computer on a chip" that includes a processor and a combination of peripherals and memory on a single chip. A Nios II processor system consists of a Nios II processor core, a set of on-chip peripherals, on-chip memory, and interfaces to off-chip memory, all implemented on a single Intel FPGA device. Like a microcontroller family, all Nios II processor systems use a consistent instruction set and programming model.

1.2. Getting Started with the Nios II Processor

The easiest way to start designing effectively is to use an Intel development kit that includes a ready-made development board and the Nios II Embedded Design Suite (EDS) containing all the software development tools necessary to write Nios II software.

The Nios II EDS includes the following two closely-related software development tool flows:

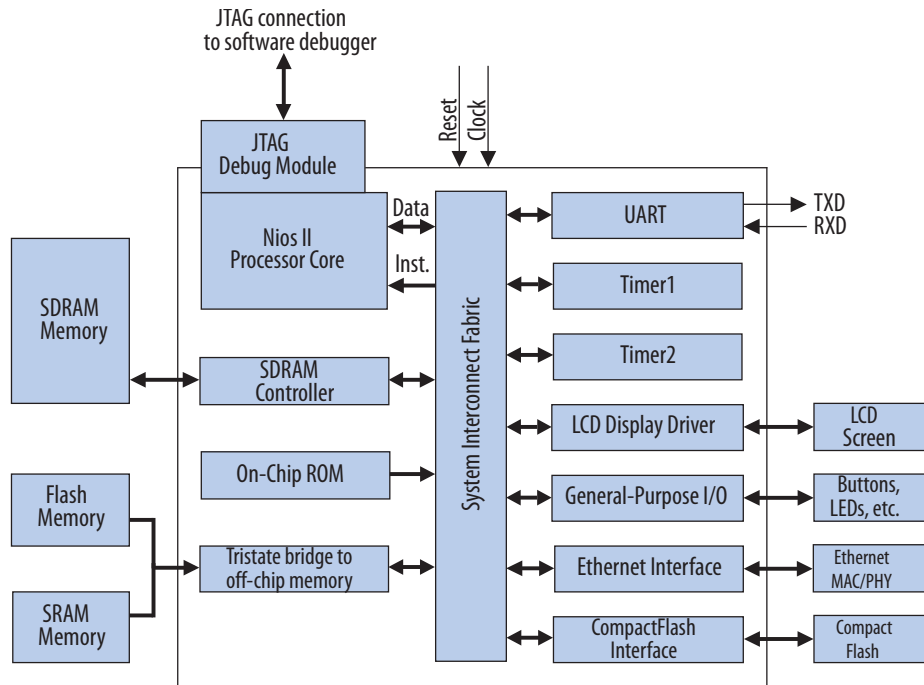
- The Nios II SBT
- The Nios II SBT for Eclipse

Both tools flows are based on the GNU C/C++ compiler. The Nios II SBT for Eclipse™ provides a familiar and established environment for software development. Using the Nios II SBT for Eclipse, you can immediately begin developing and simulating Nios II software applications.

The Nios II SBT also provides a command line interface.

Using the Nios II hardware reference designs included in an Intel development kit, you can prototype an application running on a board before building a custom hardware platform.

Figure 1. Example of a Nios II Processor System



If the prototype system adequately meets design requirements using an Intel-provided reference design, you can copy the reference design and use it without modification in the final hardware platform. Otherwise, you can customize the Nios II processor system until it meets cost or performance requirements.

Related Information

[All Development Kits - Intel FPGA webpage](#)

1.3. Customizing Nios II Processor Designs

In practice, most FPGA designs implement some extra logic in addition to the processor system. Intel FPGAs provide flexibility to add features and enhance performance of the Nios II processor system. You can also eliminate unnecessary processor features and peripherals to fit the design in a smaller, lower-cost device.



Because the pins and logic resources in Intel FPGA devices are programmable, many customizations are possible:

- You can rearrange the pins on the chip to simplify the board design. For example, you can move address and data pins for external SDRAM memory to any side of the chip to shorten board traces.
- You can use extra pins and logic resources on the chip for functions unrelated to the processor. Extra resources can provide a few extra gates and registers as glue logic for the board design; or extra resources can implement entire systems. For example, a Nios II processor system consumes only 5% of a large Intel FPGA, leaving the rest of the chip's resources available to implement other functions.
- You can use extra pins and logic on the chip to implement additional peripherals for the Nios II processor system. Intel FPGA offers a library of peripherals that easily connect to Nios II processor systems.

1.4. Configurable Soft Processor Core Concepts

This section introduces Nios II concepts that are unique or different from other discrete microcontrollers. The concepts described in this section provide a foundation for understanding the rest of the features discussed in this handbook.

1.4.1. Configurable Soft Processor Core

The Nios II processor is a configurable soft IP core, as opposed to a fixed, off-the-shelf microcontroller. You can add or remove features on a system-by-system basis to meet performance or price goals. Soft means the processor core is not fixed in silicon and can be targeted to any Intel FPGA family.

You are not required to create a new Nios II processor configuration for every new design. Intel FPGA provides ready-made Nios II system designs that you can use as is. If these designs meet your system requirements, there is no need to configure the design further. In addition, you can use the Nios II instruction set simulator to begin writing and debugging Nios II applications before the final hardware configuration is determined.

1.4.2. Flexible Peripheral Set and Address Map

A flexible peripheral set is one of the most notable differences between Nios II processor systems and fixed microcontrollers. Because the Nios II processor is implemented in programmable logic, you can easily build made-to-order Nios II processor systems with the exact peripheral set required for the target applications.

Intel FPGA provides software constructs to access memory and peripherals generically, independently of address location. Therefore, the flexible peripheral set and address map does not affect application developers.

There are two broad classes of peripherals: standard peripherals and custom peripherals.

1.4.2.1. Standard Peripherals

Intel FPGA provides a set of peripherals commonly used in microcontrollers, such as timers, serial communication interfaces, general-purpose I/O, SDRAM controllers, and other memory interfaces. The list of available peripherals continues to increase as Intel FPGA and third-party vendors release new peripherals.



Related Information

[Embedded Peripherals IP User Guide](#)

1.4.2.2. Custom Components

You can also create custom components and integrate them in Nios II processor systems. For performance-critical systems that spend most CPU cycles executing a specific section of code, it is a common technique to create a custom peripheral that implements the same function in hardware.

This approach offers a double performance benefit:

- Hardware implementation is faster than software.
- Processor is free to perform other functions in parallel while the custom peripheral operates on data.

Related Information

[Creating Platform Designer Components](#)

1.4.2.3. Custom Instructions

Like custom peripherals, custom instructions allow you to increase system performance by augmenting the processor with custom hardware. You can achieve significant performance improvements, often on the order of 10 to 100 times, by implementing performance-critical operations in hardware using custom instruction logic.

The custom logic is integrated into the Nios II processor's arithmetic logic unit (ALU). Similar to native Nios II instructions, custom instruction logic can take values from up to two source registers and optionally write back a result to a destination register.

Because the processor is implemented on reprogrammable Intel FPGAs, software and hardware engineers can work together to iteratively optimize the hardware and test the results of software running on hardware.

From the software perspective, custom instructions appear as machine-generated assembly macros or C functions, so programmers do not need to understand assembly language to use custom instructions.

1.4.3. Automated System Generation

Intel FPGA's Platform Designer system integration tools fully automate the process of configuring processor features and generating a hardware design that you program in an Intel FPGA device. The Platform Designer graphical user interface (GUI) enables you to configure Nios II processor systems with any number of peripherals and memory interfaces. You can create entire processor systems without performing any schematic or HDL design entry. Platform Designer can also import HDL design files, providing an easy mechanism to integrate custom logic in a Nios II processor system.

After system generation, you can download the design onto a board, and debug software executing on the board. To the software developer, the processor architecture of the design is set. Software development proceeds in the same manner as for traditional, nonconfigurable processors.



1.5. Intel FPGA IP Evaluation Mode

You can evaluate the Nios II processor without a license. With Intel FPGA's free Intel FPGA IP Evaluation Mode feature, you can perform the following actions:

- Simulate the behavior of a Nios II processor within your system.
- Verify the functionality of your design, as well as evaluate its size and speed quickly and easily.
- Generate time-limited device programming files for designs that include Nios II processors.
- Program a device and verify your design in hardware.

You only need to purchase a license for the Nios II processor when you are completely satisfied with its functionality and performance, and want to take your design to production.


Related Information

[AN 320: Using Intel FPGA IP Evaluation Mode](#)

For more information about Intel FPGA IP Evaluation Mode.

1.6. Introduction Revision History

Document Version	Changes
2019.10.17	Removed sections: <i>Installing Windows* Subsystem for Linux* (WSL) on Windows</i> and placed them in the Nios II Software Developer Handbook .
2019.07.01	Added sections: <i>Installing Windows Subsystem for Linux (WSL) on Windows</i> .
2019.04.30	Added section: <i>Installing Eclipse IDE into Nios II EDS</i> .
2018.04.18	<ul style="list-style-type: none">• Implemented editorial enhancements.• Changed the document part number to NII-PRG.• Corrected incorrect callouts of <i>ldio/stio</i> to <i>ldwio/ stwio</i> throughout the document.
2016.10.28	Maintenance release.
2015.04.02	Initial release



2. Processor Architecture

This chapter describes the hardware structure of the Nios II processor, including a discussion of all the functional units of the Nios II architecture and the fundamentals of the Nios II processor hardware implementation.

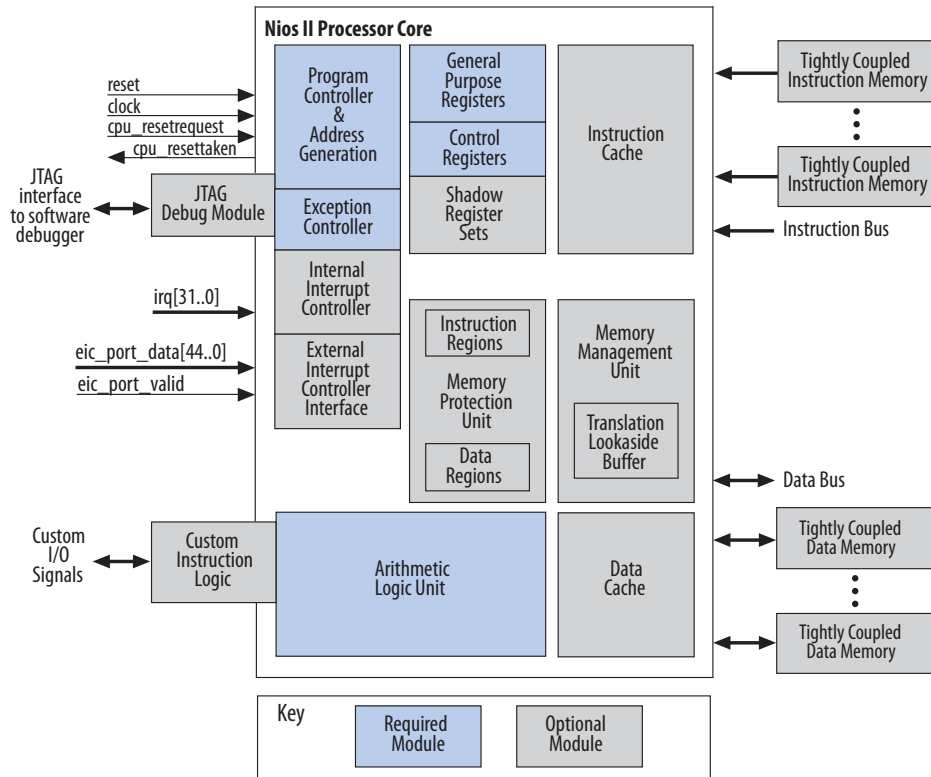
The Nios II architecture describes an instruction set architecture (ISA). The ISA in turn necessitates a set of functional units that implement the instructions. A Nios II processor core is a hardware design that implements the Nios II instruction set and supports the functional units described in this document. The processor core does not include peripherals or the connection logic to the outside world. It includes only the circuits required to implement the Nios II architecture.

The Nios II architecture defines the following functional units:

- Register file
- Arithmetic logic unit (ALU)
- Interface to custom instruction logic
- Exception controller
- Internal or external interrupt controller
- Instruction bus
- Data bus
- Memory management unit (MMU)
- Memory protection unit (MPU)
- Instruction and data cache memories
- Tightly-coupled memory interfaces for instructions and data
- JTAG debug module



Figure 2. Nios II Processor Core Block Diagram



2.1. Processor Implementation

The functional units of the Nios II architecture form the foundation for the Nios II instruction set. However, this does not indicate that any unit is implemented in hardware. The Nios II architecture describes an instruction set, not a particular hardware implementation. A functional unit can be implemented in hardware, emulated in software, or omitted entirely.

A Nios II implementation is a set of design choices embodied by a particular Nios II processor core. All implementations support the instruction set defined in the *Instruction Set Reference* chapter.

Each implementation achieves specific objectives, such as smaller core size or higher performance. This flexibility allows the Nios II architecture to adapt to different target applications.

Implementation variables generally fit one of three trade-off patterns: more or less of a feature; inclusion or exclusion of a feature; hardware implementation or software emulation of a feature. An example of each trade-off follows:

- More or less of a feature—For example, to fine-tune performance, you can increase or decrease the amount of instruction cache memory. A larger cache increases execution speed of large programs, while a smaller cache conserves on-chip memory resources.
- Inclusion or exclusion of a feature—For example, to reduce cost, you can choose to omit the JTAG debug module. This decision conserves on-chip logic and memory resources, but it eliminates the ability to use a software debugger to debug applications.
- Hardware implementation or software emulation—For example, in control applications that rarely perform complex arithmetic, you can choose for the division instruction to be emulated in software. Removing the divide hardware conserves on-chip resources but increases the execution time of division operations.

For information about which Nios II cores supports what features, refer to the *Nios II Core Implementation Details* chapter of the *Nios II Processor Reference Handbook*.

For complete details about user-selectable parameters for the Nios II processor, refer to the *Instantiating the Nios II Processor* chapter of the *Nios II Processor Reference Handbook*.

Related Information

- [Nios II Core Implementation Details](#) on page 118
- [Instruction Set Reference](#) on page 166

2.2. Register File

The Nios II architecture supports a flat register file, consisting of thirty-two 32-bit general-purpose integer registers, and up to thirty-two 32-bit control registers. The architecture supports supervisor and user modes that allow system code to protect the control registers from errant applications.

The Nios II processor can optionally have one or more shadow register sets. A shadow register set is a complete set of Nios II general-purpose registers. When shadow register sets are implemented, the `CRS` field of the `status` register indicates which register set is currently in use. An instruction access to a general-purpose register uses whichever register set is active.

A typical use of shadow register sets is to accelerate context switching. When shadow register sets are implemented, the Nios II processor has two special instructions, `rdprs` and `wrprs`, for moving data between register sets. Shadow register sets are typically manipulated by an operating system kernel, and are transparent to application code. A Nios II processor can have up to 63 shadow register sets.

The Nios II architecture allows for the future addition of floating-point registers.

For details about shadow register set implementation and usage, refer to “Registers” and “Exception Processing” in the *Programming Model* chapter of the *Nios II Processor Reference Handbook*.



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

Related Information

- [Programming Model](#) on page 33
- [Instruction Set Reference](#) on page 166

2.3. Arithmetic Logic Unit

The Nios II ALU operates on data stored in general-purpose registers. ALU operations take one or two inputs from registers, and store a result back in a register. The ALU supports the data operations described in the table below. To implement any other operation, software computes the result by performing a combination of the fundamental operations.

Table 1. Operations Supported by the Nios II ALU

Category	Details
Arithmetic	The ALU supports addition, subtraction, multiplication, and division on signed and unsigned operands.
Relational	The ALU supports the equal, not-equal, greater-than-or-equal, and less-than relational operations (<code>==</code> , <code>!=</code> , <code>>=</code> , <code><</code>) on signed and unsigned operands.
Logical	The ALU supports AND, OR, NOR, and XOR logical operations.
Shift and Rotate	The ALU supports shift and rotate operations, and can shift/rotate data by 0 to 31 bit positions per instruction. The ALU supports arithmetic shift right and logical shift right/left. The ALU supports rotate left/right.

2.3.1. Unimplemented Instructions

Some Nios II processor core implementations do not provide hardware to support the entire Nios II instruction set. In such a core, instructions without hardware support are known as unimplemented instructions.

The processor generates an exception whenever it issues an unimplemented instruction so your exception handler can call a routine that emulates the operation in software. Unimplemented instructions do not affect the programmer's view of the processor.

For a list of potential unimplemented instructions, refer to the *Programming Model* chapter of the *Nios II Processor Reference Handbook*.

Related Information

[Programming Model](#) on page 33

2.3.2. Custom Instructions

The Nios II architecture supports user-defined custom instructions. The Nios II ALU connects directly to custom instruction logic, enabling you to implement operations in hardware that are accessed and used exactly like native instructions.

Refer to "Custom Instruction Tab" in the *Instantiating the Nios II Processor* chapter of the *Nios II Processor Reference Handbook* for additional information.

Related Information

Nios II Custom Instruction User Guide

2.4. Reset and Debug Signals

The table below describes the reset and debug signals that the Nios II processor core supports.

Table 2. Nios II Processor Debug and Reset Signals

Signal Name	Type	Purpose
reset	Reset	This is a global hardware reset signal that forces the processor core to reset immediately.
cpu_resetrequest	Reset	<p>This is an optional, local reset signal that causes the processor to reset without affecting other components in the Nios II system. The processor finishes executing any instructions in the pipeline, and then enters the reset state. This process can take several clock cycles, so be sure to continue asserting the <code>cpu_resetrequest</code> signal until the processor core asserts a <code>cpu_resettaken</code> signal.</p> <p>The processor core asserts a <code>cpu_resettaken</code> signal for 1 cycle when the reset is complete and then periodically if <code>cpu_resetrequest</code> remains asserted. The processor remains in the reset state for as long as <code>cpu_resetrequest</code> is asserted. While the processor is in the reset state, it periodically reads from the reset address. It discards the result of the read, and remains in the reset state.</p> <p>The processor does not respond to <code>cpu_resetrequest</code> when the processor is under the control of the JTAG debug module, that is, when the processor is paused. The processor responds to the <code>cpu_resetrequest</code> signal if the signal is asserted when the JTAG debug module relinquishes control, both momentarily during each single step as well as when you resume execution.</p>
debug_reset_request	Reset	This reset output signal appears when the JTAG Debug module is enabled. This signal is triggered by the JTAG debugger or <code>nios2-download -r</code> command. This signal must be connected to the <code>reset</code> input signal of the Nios II processor which allows the JTAG debugger to reset the processor. This signal can be connected to the reset input signal of other components when needed.
debugreq	Debug	This is an optional signal that temporarily suspends the processor for debugging purposes. When you assert the signal, the processor pauses in the same manner as when a breakpoint is encountered, transfers execution to the routine located at the break address, and asserts a <code>debugack</code> signal. Asserting the <code>debugreq</code> signal when the processor is already paused has no effect.
reset_req	Reset	This optional signal prevents the memory corruption by performing a reset handshake before the processor resets.

For more information about adding reset signals and debug signals to the Nios II processor, refer to **Advanced Features Tab** and **JTAG Debug Module Tab** in the *Instantiating the Nios II Processor* chapter respectively.

2.5. Exception and Interrupt Controllers

The Nios II processor includes hardware for handling exceptions, including hardware interrupts. It also includes an optional external interrupt controller (EIC) interface. The EIC interface enables you to speed up interrupt handling in a complex system by adding a custom interrupt controller.



2.5.1. Exception Controller

The Nios II architecture provides a simple, nonvectored exception controller to handle all exception types. Each exception, including internal hardware interrupts, causes the processor to transfer execution to an exception address. An exception handler at this address determines the cause of the exception and dispatches an appropriate exception routine.

Exception addresses are specified with the Platform Designer Nios II Processor parameter editor.

All exceptions are precise. Precise means that the processor has completed execution of all instructions preceding the faulting instruction and not started execution of instructions following the faulting instruction. Precise exceptions allow the processor to resume program execution once the exception handler clears the exception.

2.5.2. EIC Interface

An EIC provides high performance hardware interrupts to reduce your program's interrupt latency. An EIC is typically used in conjunction with shadow register sets and when you need more than the 32 interrupts provided by the Nios II internal interrupt controller.

The Nios II processor connects to an EIC through the EIC interface. When an EIC is present, the internal interrupt controller is not implemented; Platform Designer connects interrupts to the EIC.

The EIC selects among active interrupts and presents one interrupt to the Nios II processor, with interrupt handler address and register set selection information. The interrupt selection algorithm is specific to the EIC implementation, and is typically based on interrupt priorities. The Nios II processor does not depend on any specific interrupt prioritization scheme in the EIC.

For every external interrupt, the EIC presents an interrupt level. The Nios II processor uses the interrupt level in determining when to service the interrupt.

Any external interrupt can be configured as an NMI. NMIs are not masked by the `status.PIE` bit, and have no interrupt level.

An EIC can be software-configurable.

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

For a typical example of an EIC, refer to the *Vectored Interrupt Controller* chapter in the Embedded Peripherals IP User Guide.

For details about EIC usage, refer to "Exception Processing" in the *Programming Model* chapter of the *Nios II Processor Reference Handbook*.

Related Information

- [Embedded Peripherals IP User Guide](#)
For a typical example of an EIC, refer to the *Vectored Interrupt Controller* chapter in the *Embedded Peripherals IP User Guide*.

- [Programming Model](#) on page 33

2.5.3. Internal Interrupt Controller

The Nios II architecture supports 32 internal hardware interrupts. The processor core has 32 level-sensitive interrupt request (IRQ) inputs, `irq0` through `irq31`, providing a unique input for each interrupt source. IRQ priority is determined by software. The architecture supports nested interrupts.

Your software can enable and disable any interrupt source individually through the `ienable` control register, which contains an interrupt-enable bit for each of the IRQ inputs. Software can enable and disable interrupts globally using the PIE bit of the `status` control register. A hardware interrupt is generated if and only if all of the following conditions are true:

- The PIE bit of the `status` register is 1
- An interrupt-request input, `irq<n>`, is asserted
- The corresponding bit `n` of the `ienable` register is 1

The interrupt vector custom instruction is less efficient than using the EIC interface with the Intel FPGA vectored interrupt controller component, and thus is deprecated in Platform Designer. Intel recommends using the EIC interface.

2.6. Memory and I/O Organization

This section explains hardware implementation details of the Nios II memory and I/O organization. The discussion covers both general concepts true of all Nios II processor systems, as well as features that might change from system to system.

The flexible nature of the Nios II memory and I/O organization are the most notable difference between Nios II processor systems and traditional microcontrollers. Because Nios II processor systems are configurable, the memories and peripherals vary from system to system. As a result, the memory and I/O organization varies from system to system.

A Nios II core uses one or more of the following to provide memory and I/O access:

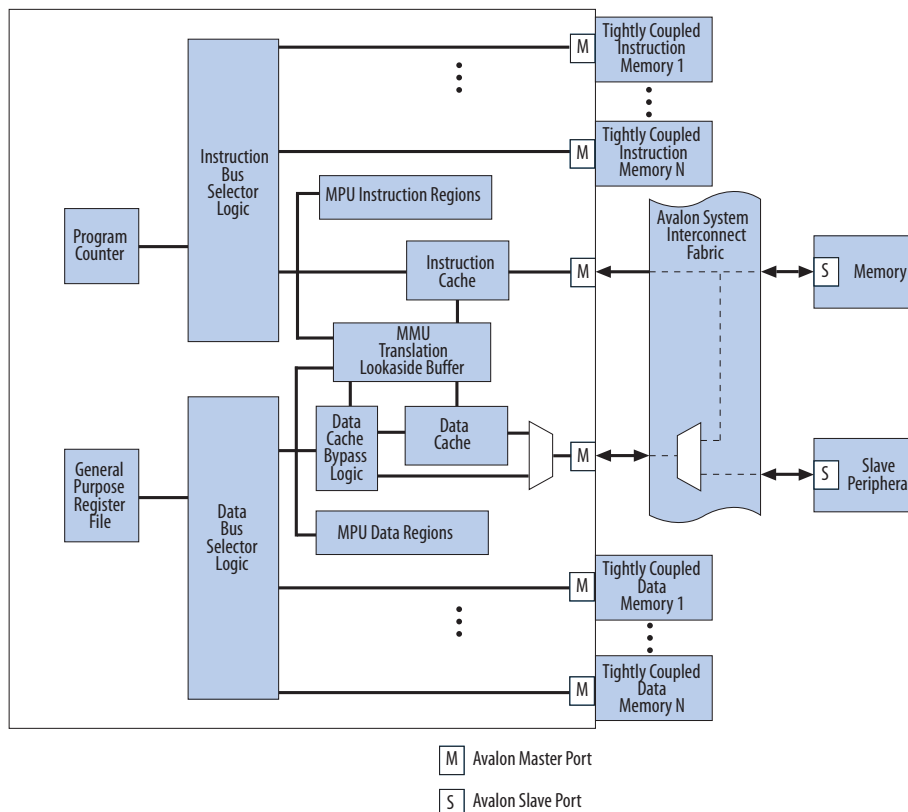
- Instruction master port—An Avalon[®] Memory-Mapped (Avalon-MM) master port that connects to instruction memory via system interconnect fabric
- Instruction cache—Fast cache memory internal to the Nios II core
- Data master port—An Avalon-MM master port that connects to data memory and peripherals via system interconnect fabric
- Data cache—Fast cache memory internal to the Nios II core
- Tightly-coupled instruction or data memory port—Interface to fast on-chip memory outside the Nios II core

The Nios II architecture handles the hardware details for the programmer, so programmers can develop Nios II applications without specific knowledge of the hardware implementation.

For details that affect programming issues, refer to the *Programming Model* chapter of the *Nios II Processor Reference Handbook*.



Figure 3. Nios II Memory and I/O Organization



Related Information

[Programming Model](#) on page 33

2.6.1. Instruction and Data Buses

The Nios II architecture supports separate instruction and data buses, classifying it as a Harvard architecture. Both the instruction and data buses are implemented as Avalon-MM master ports that adhere to the Avalon-MM interface specification. The data master port connects to both memory and peripheral components, while the instruction master port connects only to memory components.

Note: The Nios II instruction and data masters have a combined address map. The memory model is arranged so that instructions and data are in the same address space.

Related Information

[Avalon Interface Specifications](#)

Refer to the Avalon Interface Specifications for details of the Avalon-MM interface.

2.6.1.1. Memory and Peripheral Access

The Nios II architecture provides memory-mapped I/O access. Both data memory and peripherals are mapped into the address space of the data master port. The Nios II architecture uses little-endian byte ordering. Words and halfwords are stored in memory with the more-significant bytes at higher addresses.

The Nios II architecture does not specify anything about the existence of memory and peripherals; the quantity, type, and connection of memory and peripherals are system-dependent. Typically, Nios II processor systems contain a mix of fast on-chip memory and slower off-chip memory. Peripherals typically reside on-chip, although interfaces to off-chip peripherals also exist.

2.6.1.2. Instruction Master Port

The Nios II instruction bus is implemented as a 32-bit Avalon-MM master port. The instruction master port performs a single function: it fetches instructions to be executed by the processor. The instruction master port does not perform any write operations.

The instruction master port is a pipelined Avalon-MM master port. Support for pipelined Avalon-MM transfers minimizes the impact of synchronous memory with pipeline latency and increases the overall f_{MAX} of the system. The instruction master port can issue successive read requests before data has returned from prior requests. The Nios II processor can prefetch sequential instructions and perform branch prediction to keep the instruction pipe as active as possible.

The instruction master port always retrieves 32 bits of data. The instruction master port relies on dynamic bus-sizing logic contained in the system interconnect fabric. By virtue of dynamic bus sizing, every instruction fetch returns a full instruction word, regardless of the width of the target memory. Consequently, programs do not need to be aware of the widths of memory in the Nios II processor system.

The Nios II architecture supports on-chip cache memory for improving average instruction fetch performance when accessing slower memory. Refer to the "Cache Memory" section of this chapter for details.

The Nios II architecture supports tightly-coupled memory, which provides guaranteed low-latency access to on-chip memory. Refer to the "Tightly-Coupled Memory" section of this chapter for details.


Related Information

- [Cache Memory](#) on page 23
- [Tightly-Coupled Memory](#) on page 25

2.6.1.3. Data Master Port

The Nios II data bus is implemented as a 32-bit Avalon-MM master port. The data master port performs two functions:

- Read data from memory or a peripheral when the processor executes a load instruction
- Write data to memory or a peripheral when the processor executes a store instruction



Byte-enable signals on the master port specify which of the four byte-lane(s) to write during store operations. Load and store operations can complete in a single clock cycle when the data master port is connected to zero-wait-state memory.

Note: Nios II only supports a fixed 32-byte linesize for data cache.

The Nios II architecture supports on-chip cache memory for improving average data transfer performance when accessing slower memory. Refer to the "Cache Memory" section of this chapter for details.

The Nios II architecture supports tightly-coupled memory, which provides guaranteed low-latency access to on-chip memory. Refer to "Tightly-Coupled Memory" section of this chapter for details.

Related Information

- [Cache Memory](#) on page 23
- [Tightly-Coupled Memory](#) on page 25

2.6.1.4. Shared Memory for Instructions and Data

Usually the instruction and data master ports share a single memory that contains both instructions and data. While the processor core has separate instruction and data buses, the overall Nios II processor system might present a single, shared instruction/data bus to the outside world. The outside view of the Nios II processor system depends on the memory and peripherals in the system and the structure of the system interconnect fabric.

The data and instruction master ports never cause a gridlock condition in which one port starves the other. For highest performance, assign the data master port higher arbitration priority on any memory that is shared by both instruction and data master ports.

2.6.2. Cache Memory

The Nios II architecture supports cache memories on both the instruction master port (instruction cache) and the data master port (data cache). Cache memory resides on-chip as an integral part of the Nios II processor core. The cache memories can improve the average memory access time for Nios II processor systems that use slow off-chip memory such as SDRAM for program and data storage.

The instruction and data caches are enabled perpetually at run-time, but methods are provided for software to bypass the data cache so that peripheral accesses do not return cached data. Cache management and cache coherency are handled by software. The Nios II instruction set provides instructions for cache management.

2.6.2.1. Configurable Cache Memory Options

The cache memories are optional. The need for higher memory performance (and by association, the need for cache memory) is application dependent. Many applications require the smallest possible processor core, and can trade-off performance for size.

A Nios II processor core might include one, both, or neither of the cache memories. Furthermore, for cores that provide data as well as instruction cache, the sizes of the cache memories are user-configurable. The inclusion of cache memory does not affect the functionality of programs, but it does affect the speed at which the processor fetches instructions and reads/writes data.

2.6.2.2. Effective Use of Cache Memory

The effectiveness of cache memory to improve performance is based on the following premises:

- Regular memory is located off-chip, and access time is long compared to on-chip memory
- The largest, performance-critical instruction loop is smaller than the instruction cache
- The largest block of performance-critical data is smaller than the data cache

Optimal cache configuration is application specific, although you can make decisions that are effective across a range of applications. For example, if a Nios II processor system includes only fast, on-chip memory (i.e., it never accesses slow, off-chip memory), an instruction or data cache is unlikely to offer any performance gain. As another example, if the critical loop of a program is 2 KB, but the size of the instruction cache is 1 KB, an instruction cache does not improve execution speed. In fact, an instruction cache may degrade performance in this situation.

If an application always requires certain data or sections of code to be located in cache memory for performance reasons, the tightly-coupled memory feature might provide a more appropriate solution. Refer to the "Tightly-Coupled Memory" section for details.

2.6.2.3. Cache Bypass Methods

The Nios II architecture provides the following methods for bypassing the data cache:

- I/O load and store instructions
- Bit-31 cache bypass
- Peripheral Region

Note: By default the Bit-31 cache bypass is turned on.

2.6.2.3.1. I/O Load and Store Instructions Method

The load and store I/O instructions such as `ldwio` and `stwio` bypass the data cache and force an Avalon-MM data transfer to a specified address.

2.6.2.3.2. The Bit-31 Cache Bypass Method

The bit-31 cache bypass method on the data master port uses bit 31 of the address as a tag that indicates whether the processor should transfer data to/from cache, or bypass it. This is a convenience for software, which might need to cache certain addresses and bypass others. Software can pass addresses as parameters between functions, without having to specify any further information about whether the addressed data is cached or not.

To determine which cores implement which cache bypass methods, refer to the *Nios II Core Implementation Details* chapter of the *Nios II Processor Reference Handbook*.



Related Information

[Nios II Core Implementation Details](#) on page 118

2.6.2.3.3. Peripheral Region

Nios II cores optionally support a new peripheral region mechanism to indicate cacheability. The peripheral region cacheability mechanism allows a user at Platform Designer generation time to specify a region of address space that is treated as non-cacheable. The peripheral region is any integer power of 2 bytes from a minimum of 4096 bytes up to a maximum of 2 GBytes and must be located at a base address aligned to the size of the peripheral region. The peripheral region is available as long as an MMU is not present.

2.6.3. Tightly-Coupled Memory

Tightly-coupled memory provides guaranteed low-latency memory access for performance-critical applications. Compared to cache memory, tightly-coupled memory provides the following benefits:

- Performance similar to cache memory
- Software can guarantee that performance-critical code or data is located in tightly-coupled memory
- No real-time caching overhead, such as loading, invalidating, or flushing memory

Physically, a tightly-coupled memory port is a separate master port on the Nios II processor core, similar to the instruction or data master port. A Nios II core can have zero, one, or multiple tightly-coupled memories. The Nios II architecture supports tightly-coupled memory for both instruction and data access. Each tightly-coupled memory port connects directly to exactly one memory with guaranteed low, fixed latency. The memory is external to the Nios II core and is located on chip.

2.6.3.1. Accessing Tightly-Coupled Memory

Tightly-coupled memories occupy normal address space, the same as other memory devices connected via system interconnect fabric. The address ranges for tightly-coupled memories (if any) are determined at system generation time.

Software accesses tightly-coupled memory using regular load and store instructions. From the software's perspective, there is no difference accessing tightly-coupled memory compared to other memory.

Note: The tightly-coupled master requires a fixed memory latency of 1 cycle. Hence, the transaction with a slave in a different clock domain may not be successful since the transfer would take more than 1 cycle.

2.6.3.2. Effective Use of Tightly-Coupled Memory

A system can use tightly-coupled memory to achieve maximum performance for accessing a specific section of code or data. For example, interrupt-intensive applications can place exception handler code into a tightly-coupled memory to minimize interrupt latency. Similarly, compute-intensive digital signal processing (DSP) applications can place data buffers into tightly-coupled memory for the fastest possible data access.

If the application's memory requirements are small enough to fit entirely on chip, it is possible to use tightly-coupled memory exclusively for code and data. Larger applications must selectively choose what to include in tightly-coupled memory to maximize the cost-performance trade-off.

Related Information

[Using Tightly Coupled Memory with the Nios II Processor Tutorial](#)

For additional tightly-coupled memory guidelines.

2.6.4. Address Map

The address map for memories and peripherals in a Nios II processor system is design dependent. You specify the address map in Platform Designer.

There are three addresses that are part of the processor and deserve special mention:

- Reset address
- Exception address
- Break handler address


Programmers access memories and peripherals by using macros and drivers. Therefore, the flexible address map does not affect application developers.

2.6.5. Memory Management Unit

The optional Nios II MMU provides the following features and functionality:

- Virtual to physical address mapping
- Memory protection
- 32-bit virtual and physical addresses, mapping a 4-GB virtual address space into as much as 4 GB of physical memory
- 4-KB page and frame size
- Low 512 MB of physical address space available for direct access
- Hardware translation lookaside buffers (TLBs), accelerating address translation
 - Separate TLBs for instruction and data accesses
 - Read, write, and execute permissions controlled per page
 - Default caching behavior controlled per page
 - TLBs acting as n-way set-associative caches for software page tables
 - TLB sizes and associativities configurable in the Nios II Processor parameter editor
- Format of page tables (or equivalent data structures) determined by system software
- Replacement policy for TLB entries determined by system software
- Write policy for TLB entries determined by system software

For more information about the MMU implementation, refer to the *Programming Model* chapter of the *Nios II Processor Reference Handbook*.



You can optionally include the MMU when you instantiate the Nios II processor in your Nios II hardware system. When present, the MMU is always enabled, and the data and instruction caches are virtually-indexed, physically-tagged caches. Several parameters are available, allowing you to optimize the MMU for your system needs.

For complete details about user-selectable parameters for the Nios II MMU, refer to the *Instantiating the Nios II Processor* chapter of the *Nios II Processor Reference Handbook*.

Note: The Nios II MMU is optional and mutually exclusive from the Nios II MPU. Nios II systems can include either an MMU or MPU, but cannot include both an MMU and MPU on the same Nios II processor core.

Related Information

[Programming Model](#) on page 33

2.6.6. Memory Protection Unit

The optional Nios II MPU provides the following features and functionality:

- Memory protection
- Up to 32 instruction regions and 32 data regions
- Variable instruction and data region sizes
- Amount of region memory defined by size or upper address limit
- Read and write access permissions for data regions
- Execute access permissions for instruction regions
- Overlapping regions

For more information about the MPU implementation, refer to the *Programming Model* chapter of the *Nios II Processor Reference Handbook*.

You can optionally include the MPU when you instantiate the Nios II processor in your Nios II hardware system. When present, the MPU is always enabled. Several parameters are available, allowing you to optimize the MPU for your system needs.

For complete details about user-selectable parameters for the Nios II MPU, refer to the *Instantiating the Nios II Processor* chapter of the *Nios II Processor Reference Handbook*.

Note: The Nios II MPU is optional and mutually exclusive from the Nios II MMU. Nios II systems can include either an MPU or MMU, but cannot include both an MPU and MMU on the same Nios II processor core.

Related Information

[Programming Model](#) on page 33

2.7. JTAG Debug Module

The Nios II architecture supports a JTAG debug module that provides on-chip emulation features to control the processor remotely from a host PC. PC-based software debugging tools communicate with the JTAG debug module and provide facilities, such as the following features:

- Downloading programs to memory
- Starting and stopping execution
- Setting breakpoints and watchpoints
- Analyzing registers and memory
- Collecting real-time execution trace data

Note: The Nios II MMU does not support the JTAG debug module trace.

The debug module connects to the JTAG circuitry in an Intel FPGA. External debugging probes can then access the processor via the standard JTAG interface on the FPGA. On the processor side, the debug module connects to signals inside the processor core. The debug module has nonmaskable control over the processor, and does not require a software stub linked into the application under test. All system resources visible to the processor in supervisor mode are available to the debug module. For trace data collection, the debug module stores trace data in memory either on-chip or in the debug probe.

The debug module gains control of the processor either by asserting a hardware break signal, or by writing a break instruction into program memory to be executed. In both cases, the processor transfers execution to the routine located at the break address. The break address is specified with the Nios II Processor parameter editor in Platform Designer.

Soft processor cores such as the Nios II processor offer unique debug capabilities beyond the features of traditional, fixed processors. The soft nature of the Nios II processor allows you to debug a system in development using a full-featured debug core, and later remove the debug features to conserve logic resources. For the release version of a product, the JTAG debug module functionality can be reduced, or removed altogether.

The following sections describe the capabilities of the Nios II JTAG debug module hardware. The usage of all hardware features is dependent on host software, such as the Nios II Software Build Tools for Eclipse, which manages the connection to the target processor and controls the debug process.

2.7.1. JTAG Target Connection

The JTAG target connection provides the ability to connect to the processor through the standard JTAG pins on the Intel FPGA. This provides basic capabilities to start and stop the processor, and examine and edit registers and memory. The JTAG target connection is the minimum requirement for the Nios II flash programmer.

Note: While the processor has no minimum clock frequency requirements, Intel recommends that your design's system clock frequency be at least four times the JTAG clock frequency to ensure that the on-chip instrumentation (OCI) core functions properly.

2.7.2. Download and Execute Software

Downloading software refers to the ability to download executable code and data to the processor's memory via the JTAG connection. After downloading software to memory, the JTAG debug module can then exit debug mode and transfer execution to the start of executable code.



2.7.3. Software Breakpoints

Software breakpoints allow you to set a breakpoint on instructions residing in RAM. The software breakpoint mechanism writes a break instruction into executable code stored in RAM. When the processor executes the break instruction, control is transferred to the JTAG debug module.

2.7.4. Hardware Breakpoints

Hardware breakpoints allow you to set a breakpoint on instructions residing in nonvolatile memory, such as flash memory. The hardware breakpoint mechanism continuously monitors the processor's current instruction address. If the instruction address matches the hardware breakpoint address, the JTAG debug module takes control of the processor.

Hardware breakpoints are implemented using the JTAG debug module's hardware trigger feature.

2.7.5. Hardware Triggers

Hardware triggers activate a debug action based on conditions on the instruction or data bus during real-time program execution. Triggers can do more than halt processor execution. For example, a trigger can be used to enable trace data collection during real-time processor execution.

Hardware trigger conditions are based on either the instruction or data bus. Trigger conditions on the same bus can be logically ANDed, enabling the JTAG debug module to trigger, for example, only on write cycles to a specific address.

Table 3. Trigger Conditions

Condition	Bus	Description
Specific address	Data, Instruction	Trigger when the bus accesses a specific address.
Specific data value	Data	Trigger when a specific data value appears on the bus.
Read cycle	Data	Trigger on a read bus cycle.
Write cycle	Data	Trigger on a write bus cycle.
Armed	Data, Instruction	Trigger only after an armed trigger event. Refer to the Armed Triggers section.
Range	Data	Trigger on a range of address values, data values, or both. Refer to the Triggering on Ranges of Values section.

When a trigger condition occurs during processor execution, the JTAG debug module triggers an action, such as halting execution, or starting trace capture. The table below lists the trigger actions supported by the Nios II JTAG debug module.

Table 4. Trigger Actions

Action	Description
Break	Halt execution and transfer control to the JTAG debug module.
External trigger	Assert a trigger signal output. This trigger output can be used, for example, to trigger an external logic analyzer.
Trace on	Turn on trace collection.

continued...



Action	Description
Trace off	Turn off trace collection.
Trace sample	Store one sample of the bus to trace buffer.
Arm	Enable an armed trigger.

Note: For the *Trace sample* trigger action, only conditions on the data bus can trigger this action.

2.7.5.1. Armed Triggers

The JTAG debug module provides a two-level trigger capability, called armed triggers. Armed triggers enable the JTAG debug module to trigger on event B, only after event A. In this example, event A causes a trigger action that enables the trigger for event B.

2.7.5.2. Triggering on Ranges of Values

The JTAG debug module can trigger on ranges of data or address values on the data bus. This mechanism uses two hardware triggers together to create a trigger condition that activates on a range of values within a specified range.

2.7.6. Trace Capture

Trace capture refers to ability to record the instruction-by-instruction execution of the processor as it executes code in real-time. The JTAG debug module offers the following trace features:

- Capture execution trace (instruction bus cycles).
- Capture data trace (data bus cycles).
- For each data bus cycle, capture address, data, or both.
- Start and stop capturing trace in real time, based on triggers.
- Manually start and stop trace under host control.
- Optionally stop capturing trace when trace buffer is full, leaving the processor executing.
- Store trace data in on-chip memory buffer in the JTAG debug module. (This memory is accessible only through the JTAG connection.)
- Store trace data to larger buffers in an off-chip debug probe.

Certain trace features require additional licensing or debug tools from third-party debug providers. For example, an on-chip trace buffer is a standard feature of the Nios II processor, but using an off-chip trace buffer requires additional debug software and hardware provided by Imagination Technologies™, LLC or Lauterbach GmbH.

Related Information

Lauterbach.com

For more information, refer to the Lauterbach GmbH website.



2.7.6.1. Execution vs. Data Trace

The JTAG debug module supports tracing the instruction bus (execution trace), the data bus (data trace), or both simultaneously. Execution trace records only the addresses of the instructions executed, enabling you to analyze where in memory (that is, in which functions) code executed. Data trace records the data associated with each load and store operation on the data bus.

The JTAG debug module can filter the data bus trace in real time to capture the following:

- Load addresses only
- Store addresses only
- Both load and store addresses
- Load data only
- Load address and data
- Store address and data
- Address and data for both loads and stores
- Single sample of the data bus upon trigger event

2.7.6.2. Trace Frames

A frame is a unit of memory allocated for collecting trace data. However, a frame is not an absolute measure of the trace depth.

To keep pace with the processor executing in real time, execution trace is optimized to store only selected addresses, such as branches, calls, traps, and interrupts. From these addresses, host-side debug software can later reconstruct an exact instruction-by-instruction execution trace. Furthermore, execution trace data is stored in a compressed format, such that one frame represents more than one instruction. As a result of these optimizations, the actual start and stop points for trace collection during execution might vary slightly from the user-specified start and stop points.

Data trace stores 100% of requested loads and stores to the trace buffer in real time. When storing to the trace buffer, data trace frames have lower priority than execution trace frames. Therefore, while data frames are always stored in chronological order, execution and data trace are not guaranteed to be exactly synchronized with each other.


2.8. Processor Architecture Revision History

Document Version	Changes
2019.05.20	<i>Reset and Debug Signals</i> : Added definition for the <code>debug_reset_request</code> signal.
2019.04.30	Maintenance release
2018.04.18	<ul style="list-style-type: none"> • Implemented editorial enhancements. • Removed the following sections: <ul style="list-style-type: none"> – Introduction to Nios II Floating Point Custom Instructions – Nios II Floating Point Hardware 2 Component – Nios II Floating Point Hardware (FPH1) Component

continued...



Document Version	Changes
	Instead, refer to the <i>Nios II Custom Instruction User Guide</i> .
2016.10.28	Maintenance release
2015.04.02	Initial release



3. Programming Model

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 these concepts and focuses on how these concepts are specifically implemented in the Nios II processor. Where possible, this chapter uses industry-standard terminology.

Note: 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 ignore the sections covering operating modes, the MMU, the MPU, or the control registers exclusively used by the MMU and MPU.

Related Information

[Nios II Software Developer's Handbook](#)

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

3.1. 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.

3.1.1. 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.

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

3.1.2. 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 section for more information about the MMU memory partitions. The MMU provides operating systems access permissions on a per-page basis. Refer to Virtual Addressing 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 for more information about MPU regions.

Related Information

- [Address Space and Memory Partitions](#) on page 36
- [Memory Regions](#) on page 40
- [Virtual Addressing](#) on page 35

3.2. 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.

3.2.1. 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 Intel FPGA[®] 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.

Note: The Intel FPGA 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 section.

Related Information

[Memory Protection Unit](#) on page 39

3.2.2. 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

3.2.2.1. 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-KB pages and physical memory into 4-KB 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.

There is a 20 bit virtual page number (VPN) and a 12 bit page offset.

Table 5. MMU Virtual Address Fields

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
Virtual Page Number															
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.

3.2.2.2. 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.

3.2.3. Address Space and Memory Partitions

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

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

3.2.3.1. Virtual Memory Address Space

The 4-GB virtual memory space is divided into partitions. The upper 2 GB of memory is reserved for the operating system and the lower 2 GB is reserved for user processes.

Table 6. Virtual Memory Partitions

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

Note: All partitions except the user partition in the "Virtual Memory Partition" table are supervisor-only partitions.

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

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

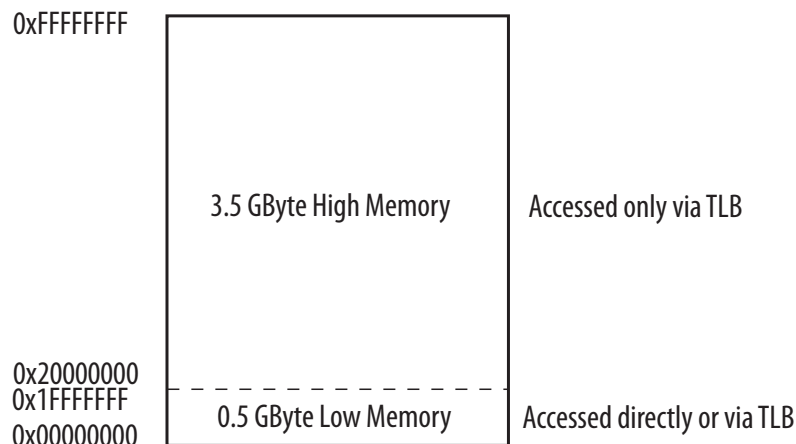


I/O and kernel partitions bypass the TLB. The kernel MMU and user partitions use the TLB. If all software runs in the kernel partition, the MMU is effectively disabled.

3.2.3.2. Physical Memory Address Space

The 4-GB physical memory is divided into low memory and high memory. The lowest ½ GB of physical address space is low memory. The upper 3½ GB of physical address space is high memory.

Figure 4. 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.

Note: 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. The Nios II Processor parameter editor in Platform Designer prevents you from choosing an address outside of low physical memory.

3.2.3.3. 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 described in the above table, Virtual Memory Partitions. 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.

3.2.4. 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.

Note:

You can configure the number of TLB entries and the number of ways (set associativity) of the TLB with the Nios II Processor parameter editor in Platform Designer. By default, the TLB is a 16-way cache. The default number of entries depends on the target device, as follows:

- Cyclone III[®], Stratix III[®], Stratix IV—256 entries, requiring one M9K RAM
For more information, refer to the *Instantiating the Nios II Processor* 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 information about instruction and data caches.

The tag portion of a TLB entry contains information used when matching a virtual address to a TLB entry.

Table 7. 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 <code>tlbmisc</code> control register, effectively extending the virtual address. The field size is configurable in the Nios_II Processor parameter editor, 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 8. 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.
<i>continued...</i>	



Field Name	Description
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.

Note: 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.

Related Information

- [Programming Model](#) on page 33
- [Nios II Core Implementation Details](#) on page 118

3.2.5. TLB Lookups

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

Example 1. 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
    
```

Example 2. 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
    
```

Refer to “Instruction-Related Exceptions” for information about TLB exceptions.

Related Information

[Instruction-Related Exceptions](#) on page 80

3.3. 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 the Memory Management Unit section.

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.

MPU protects user application. Therefore for interrupt service, the system must have access to the regions that may cause potential violation because MPU generates exception post access and does not prevent access to the memory region in hardware.

Related Information

[Memory Management Unit](#) on page 34

3.3.1. 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)

3.3.1.1. 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-KB region must have a base address that is a multiple of 4 KB. If the base address is not properly aligned, the behavior is undefined.

3.3.1.2. Region Type

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

3.3.1.3. 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.



3.3.1.4. Region Size or Upper Address Limit

A Platform Designer 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 256 bytes but can be configured 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.

3.3.1.5. 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.

3.3.1.6. 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 `ldwio` or `stwio` instructions. The bit-31 cache and Peripheral Region features are available when the MMU is not present.

Refer to the Cache Memory section for more information on cache bypass and Peripheral Region.

Related Information

[Cache Memory](#) on page 96

3.3.2. 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.

3.3.3. 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 the Working with the MPU section for more information.

Related Information

[Working with the MPU](#) on page 65

3.4. 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 register type.

3.4.1. General-Purpose Registers

The Nios II architecture provides thirty-two 32-bit general-purpose registers, `r0` through `r31`. 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`.

Table 9. The Nios II General-Purpose Registers

Register	Name	Function	Register	Name	Function
r0	zero	0x00000000	r16		Callee-saved register
r1	at	Assembler temporary	r17		Callee-saved register
r2		Return value	r18		Callee-saved register
r3		Return value	r19		Callee-saved register
r4		Register arguments	r20		Callee-saved register
r5		Register arguments	r21		Callee-saved register
r6		Register arguments	r22		Callee-saved register
r7		Register arguments	r23		Callee-saved register
r8		Caller-saved register	r24	et	Exception temporary
r9		Caller-saved register	r25	bt	Breakpoint temporary

continued...



Register	Name	Function	Register	Name	Function
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	sstatus	Status register
r15		Caller-saved register	r31	ra	Return address

Note: 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.

Note: 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 Status Register section.

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

Related Information

[Application Binary Interface](#) on page 143

3.4.2. 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.

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

The Nios II architecture supports up to 32 control registers. All non-reserved control registers have names recognized by the assembler.

Table 10. Control Register Names and Bits

Register	Name	Register Contents
0	status	Refer to The status Register on page 45
1	estatus	Refer to The estatus Register on page 47
2	bstatus	Refer to The bstatus Register
3	ienable	Internal interrupt-enable bits The ienable Register Available only when the external interrupt controller interface is not present. Otherwise reserved.
4	ipending	Pending internal interrupt bits The ipending Register Available only when the external interrupt controller interface is not present. Otherwise reserved.
<i>continued...</i>		



Register	Name	Register Contents
5	cpuid	Unique processor identifier
6	Reserved	Reserved
7	exception	Refer to The exception Register
8	pteaddr	Refer to The pteaddr Register Available only when the MMU is present. Otherwise reserved.
9	tlbacc	Refer to The tlbacc Register Available only when the MMU is present. Otherwise reserved.
10	tlbmisc	Refer to The tlbmisc Register Available only when the MMU is present. Otherwise reserved.
11	eccinj	Refer to The eccinj Register Available only when ECC is present.
12	badaddr	Refer to The badaddr Register
13	config	Refer to The config Register on page 55 Available only when the MPU or ECC is present. Otherwise reserved.
14	mpubase	Refer to The mpubase Register Available only when the MPU is present. Otherwise reserved.
15	mpuacc	Refer to The mpuacc Register for MASK variations table. Available only when the MPU is present. Otherwise reserved.
16-31	Reserved	Reserved

The following sections describe the non-reserved 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.

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

The Nios II architecture supports up to 32 control registers. All non-reserved control registers have names recognized by the assembler.

Table 11. Control Register Names and Bits

Register	Name	Register Contents
0	status	Refer to The status Register on page 45
1	estatus	Refer to The estatus Register on page 47
2	bstatus	Refer to The bstatus Register
3	ienable	Internal interrupt-enable bits The ienable Register
<i>continued...</i>		



Register	Name	Register Contents
		Available only when the external interrupt controller interface is not present. Otherwise reserved.
4	ipending	Pending internal interrupt bits The ipending Register Available only when the external interrupt controller interface is not present. Otherwise reserved.
5	cpuid	Unique processor identifier
6	Reserved	Reserved
7	exception	Refer to The exception Register
8	pteaddr	Refer to The pteaddr Register Available only when the MMU is present. Otherwise reserved.
9	tlbacc	Refer to The tlbacc Register Available only when the MMU is present. Otherwise reserved.
10	tlbmisc	Refer to The tlbmisc Register Available only when the MMU is present. Otherwise reserved.
11	eccinj	Refer to The eccinj Register Available only when ECC is present.
12	badaddr	Refer to The badaddr Register
13	config	Refer to The config Register on page 55 Available when the MPU or ECC is present. Otherwise reserved.
14	mpubase	Refer to The mpubase Register Available only when the MPU is present. Otherwise reserved.
15	mpuacc	Refer to The mpuacc Register for MASK variations table. Available only when the MPU is present. Otherwise reserved.
16–31	Reserved	Reserved

The following sections describe the non-reserved control registers.

3.4.2.1. 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 12. [status Control Register Fields](#)

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

Table 13. status Control Register Field Descriptions

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 ⁽¹⁾
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 ⁽⁴⁾
PRS	<p>PRS is the previous register set field. The processor copies the CRS field to the PRS field upon one of the following events:</p> <ul style="list-style-type: none"> In a processor with no MMU, on any exception In a processor with an MMU, on one of the following: <ul style="list-style-type: none"> Break exception Nonbreak exception when <code>status.EH</code> is zero <p>The processor copies CRS to PRS immediately after copying the status register to <code>estatus</code>, <code>bstatus</code> or <code>sstatus</code>.</p> <p>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.</p> <p>1 Ensure that system software writes only valid register set numbers to the PRS field. Processor behavior is undefined with an unimplemented register set number.</p>	Read/Write	0	Shadow register sets only ⁽⁴⁾
CRS	<p>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.</p> <p>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.</p>	Read ⁽²⁾	0	Shadow register sets only ⁽⁴⁾
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 ⁽⁴⁾
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 ⁽⁴⁾
EH ⁽³⁾	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	Read/Write	0	MMU or ECC only ⁽⁴⁾

continued...

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

(2) The CRS field is read-only. For information about manually changing register sets, refer to the External Interrupt Controller Interface section.



Bit	Description	Access	Reset	Available
	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.			
U ⁽³⁾	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 ⁽⁴⁾
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

Related Information

External Interrupt Controller Interface on page 77

3.4.2.2. The estatus Register

The `estatus` register holds a saved copy of the `status` register during nonbreak exception processing.

Table 14. estatus Control Register Fields

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

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

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`.

Note: 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 section.

⁽³⁾ The state where both EH and U are one is illegal and causes undefined results.

⁽⁴⁾ When this field is unimplemented, the field value always reads as 0, and the processor behaves accordingly.

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 the Exception Processing section for more information.

Related Information

- [Exception Processing](#) on page 71
- [The `sstatus` Register](#) on page 63

3.4.2.3. The `bstatus` Register

The `bstatus` register holds a saved copy of the `status` register during break exception processing.

Table 15. `bstatus` Control Register Fields

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

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

The Status Control Register Field Description table 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 the Processing a Break section for more information.

Related Information

[Processing a Break](#) on page 76

3.4.2.4. 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 `irqn` interrupt is enabled; a bit value of zero means that the corresponding interrupt is disabled. Refer to the Exception Processing section for more information.

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

Related Information

[Exception Processing](#) on page 71



3.4.2.5. The `ipending` Register

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

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

3.4.2.6. The `cpuid` Register

The `cpuid` register holds a constant value that you define in the Nios II Processor parameter editor to uniquely identify each processor in a multiprocessor system. In Platform Designer, unique values must be assigned manually. Writing to the `cpuid` register has no effect.

3.4.2.7. The exception Register

Nios II/f Processor provides information useful to system software for exception processing in the exception and `badaddr` registers when an exception occurs.

Note: The exception register is not available for Nios II/e core, therefore you cannot add a MMU or MPU to the processor configuration. For more information, refer to the *Nios II Core Implementation Details* chapter of this document.

For information about controlling the extra exception information option, refer to the *Instantiating the Nios II Processor* chapter of this document.

Table 16. exception Control Register Fields

Bit Fields																
ECCFTL	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	
Reserved																
16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Reserved										Cause			Rsvd			

Table 17. exception Control Register Field Descriptions

Field	Description	Access	Reset	Available
ECCFTL	The Nios II processor writes to ECCFTL when it detects a potentially fatal ECC error. When <code>ECCFTL = 1</code> , the Nios II processor detects an ECC register file error. When <code>ECCFTL = 0</code> , another ECC exception occurred.	Read	0	Only with ECC
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 the Nios II Exceptions (In Decreasing Priority Order table lists the CAUSE field value for each exception. CAUSE is not written on a break or an external interrupt.	Read	0	Only with Nios II/f

Related Information

[Programming Model](#) on page 33

3.4.2.8. 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 18. pteaddr Control Register Fields

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
PTBASE															
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
VPN														Rsvd	

Table 19. 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.

3.4.2.9. 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 can 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 20. tlbacc Control Register Fields

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

Issuing a `wrcctl` instruction to the `tlbacc` register writes the `tlbacc` register with the specified value. If `tlbmisc.WE = 1`, the `wrcctl` 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 21. 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, <code>tlbmisc.PID</code> is included in the TLB lookup. When G = 1, <code>tlbmisc.PID</code> 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.

3.4.2.10. The tlbmisc Register

The `tlbmisc` register contains the remaining TLB-related fields and is only available in systems with an MMU.

Table 22. tlbmisc Control Register Fields

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

Table 23. tlbmisc Control Register Field Descriptions

Field	Description	Access	Reset	Available
EE	If this field is a 1, a software-triggered ECC error (1, 2, or 3 bit error) occurred because software initiated a TLB read operation. Only set this field to 1 if <code>CONFIG.ECCEN</code> is 1.	Read/Write	0	Only with MMU and EEC
WAY	The <code>WAY</code> field controls the mapping from the VPN to a particular TLB entry. This field size is variable. Unused upper bits must be written as zero.	Read/Write	0	Only with MMU
RD	<code>RD</code> is the read flag. Setting <code>RD</code> to one triggers a TLB read operation.	Write	0	Only with MMU
WE	<code>WE</code> is the TLB write enable flag. When <code>WE = 1</code> , a write to <code>tlbacc</code> writes through to a TLB entry.	Read/Write	0	Only with MMU
PID	<code>PID</code> is the process identifier field. This field size is variable. Unused upper bits must be written as zero.	Read/Write	0	Only with MMU
DBL	<code>DBL</code> is the double TLB miss exception flag.	Read	0	Only with MMU
BAD	<code>BAD</code> is the bad virtual address exception flag.	Read	0	Only with MMU
PERM	<code>PERM</code> is the TLB permission violation exception flag.	Read	0	Only with MMU
D	<code>D</code> is the data access exception flag. When <code>D = 1</code> , the exception is a data access exception. When <code>D = 0</code> , the exception is an instruction access exception.	Read	0	Only with MMU

For `DBL`, `BAD`, and `PERM` fields you can also use `exception.CAUSE` to determine these exceptions.

The following sections provide more information about the `tlbmisc` fields.

3.4.2.10.1. 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`



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.

3.4.2.10.2. 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`.

3.4.2.10.3. 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.

Note: The number of ways (sets) is configurable in Platform Designer at generation time, up to a maximum of 16.

3.4.2.10.4. 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.

Note: 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 `wrc1l` instruction, the `PID` value read from the TLB has priority over the value written by the `wrc1l` instruction.

The size of the `PID` field is configured in Platform Designer 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.

3.4.2.10.5. 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.

3.4.2.10.6. 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 Nios II Exceptions (In Decreasing Priority Order) table in the "Exception Overview" section for more information on these exceptions.

Related Information

[Exception Overview](#) on page 72

3.4.2.10.7. 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.

3.4.2.10.8. 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

3.4.2.11. The badaddr Register

Nios II/f processor provides information useful to system software for exception processing in the exception and `badaddr` registers when an exception occurs.

Note: The exception register is not available for Nios II/e core, therefore you cannot add a MMU or MPU to the processor configuration. For more information, refer to the *Nios II Core Implementation Details* chapter of this document.

When an exception occurs in Nios II/f processor, the `badaddr` register contains the byte instruction or data address associated with certain exceptions at the time the exception occurred. The Nios II Exceptions Table lists which exceptions write the `badaddr` register along with the value written.



Table 24. badaddr Control Register Fields

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

Table 25. 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 the Nios II Exceptions Table lists which exceptions write the BADDR field.	Read	0	Only with Nios II/f

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.

Related Information

- [Exception Overview](#) on page 72
- [Programming Model](#) on page 33

3.4.2.12. 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 26. config Control Register Fields

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
Reserved															
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Reserved												ECCE XE	ECCE N	ANI	PE

Table 27. config Control Register Field Descriptions

Field	Description	Access	Reset	Available
ANI	ANI is the automatic nested interrupt mode bit. If ANI is set to zero, the processor clears <code>status.PIE</code> on each interrupt, disabling fast nested interrupts. If ANI is set to one, the processor keeps <code>status.PIE</code> 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
ECCEXE	ECCEXE is the ECC error exception enable bit. When ECCEXE = 1, the Nios II processor generates ECC error exceptions.	Read/Write	0	Only with ECC
ECCEN	ECCEN is the ECC enable bit. When ECCEN = 0, the Nios II processor ignores all ECC errors. When ECCEN = 1, the Nios II processor recovers all recoverable ECC errors.	Read/Write	0	Only with ECC
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

3.4.2.13. 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 28. mpubase Control Register Fields

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

Table 29. 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

⁽⁵⁾ This field size is variable. Unused upper bits must be written as zero.

⁽⁶⁾ This field size is variable. Unused upper bits and unused lower bits must be written as zero.



The `BASE` field specifies the base address of an MPU region. The 24-bit `BASE` field corresponds to bits 8 through 31 of the base address, making the base address always a multiple of 256 bytes. If the minimum region size (set in Platform Designer at generation time) is larger than 256 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 two least-significant bits of the `BASE` field (bits 8 through 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 the MPU Region Read and Write Operations section for more information on MPU region read and write operations.

Related Information

[MPU Region Read and Write Operations](#) on page 65

3.4.2.14. 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 can 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.

A Platform Designer generation-time option controls whether the `mpuacc` register contains a `MASK` or `LIMIT` field.

Table 30. `mpuacc` Control Register Fields for `MASK` Variation

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
MASK[n-1:p] ⁽⁷⁾															
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
MASK[n-1:p] ⁽⁷⁾								0	MT			PERM		RD	WR

Table 31. `mpuacc` Control Register Fields for `LIMIT` Variation

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
LIMIT[n:p] ⁽⁷⁾															
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
LIMIT[n:p] ⁽⁷⁾								0	MT			PERM		RD	WR

Table 32. mpuacc Control Register Field Descriptions

Field	Description	Access	Reset	Available
MASK	MASK specifies the size of the region.	Read/Write	0	Only with MPU
LIMIT	LIMIT specifies the upper address limit of the region.	Read/Write	0	Only with MPU
MT	(MT) Memory Type: <ul style="list-style-type: none"> 0 = peripheral (non-cacheable, non-write bufferable) 1 = normal (cacheable, write bufferable) 2 = device (non-cacheable, write bufferable) 3 = reserved 	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
WR	WR is the write region flag. When WR = 1, wrctl instructions to the mpuacc register perform a write operation.	Write	0	Only with MPU

The MASK and LIMIT fields are mutually exclusive. Refer to [Table 30](#) on page 57 and [Table 31](#) on page 57.

Related Information

- [The LIMIT Field](#) on page 59
- [The MASK Field](#) on page 58

3.4.2.14.1. 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.

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

MASK Region Size Encodings Table lists the MASK field encodings for all possible region sizes in a full 31-bit byte address space.

Table 33. MASK Region Size Encodings

MASK Encoding	Region Size
0xFFFFFFFF	256 bytes
0xFFFFFFFFE	512 bytes
<i>continued...</i>	

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



MASK Encoding	Region Size
0xFFFFFC	1 KB
0xFFFFF8	2 KB
0xFFFFF0	4 KB
0xFFFFE0	8 KB
0xFFFFC0	16 KB
0xFFFF80	32 KB
0xFFFF00	64 KB
0xFFFE00	128 KB
0xFFFC00	256 KB
0xFF8000	512 KB
0xFF0000	1 MB
0xFFE000	2 MB
0xFFC000	4 MB
0xFF8000	8 MB
0xFF0000	16 MB
0xFE0000	32 MB
0xFC0000	64 MB
0xF80000	128 MB
0xF00000	256 MB
0xE00000	512 MB
0xC00000	1 GB
0x800000	2 GB
0x000000	4 GB

The MASK field contains the following value, where `region_size` is in bytes:

$$\text{MASK} = 0xFFFFFFFF \ll \log_2(\text{region_size} \gg 8)$$

3.4.2.14.2. 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.

3.4.2.14.3. The MT Flag

The MT flag determines the default memory type of an MPU data region. . The MT flag only applies to data regions. For instruction regions, the MT bit must be written with 0 for instruction regions 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 `ldwio` or `stwio` instruction. The encoding of the MT field is setup to be backwards-compatible with the Nios II Classic core MPU where bit 5 of MPUACC contains the cacheable bit (0 = non-cacheable, 1 = cacheable) and bit 6 is zero.

Instruction	Bit 31 Bypass Data Cache	MT Flag Value	Data Access
ldwio/stwio	don't care	don't care	non-cacheable
ldw/stw	1	don't care	non-cacheable
ldw/stw	0	0 or 2	non-cacheable
ldw/stw	0	1	cacheable

Note: The bit 31 cache bypass and peripheral region features are supported when the MPU is present. Refer to the Cache memory section for more information on cache bypass.

3.4.2.14.4. The PERM Field

The PERM field specifies the allowed access permissions.

Table 34. Instruction Region Permission Values

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

Table 35. 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

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



3.4.2.14.5. The RD Flag

Setting the RD flag signifies that an MPU region read operation should be performed when a `wrc1` instruction is issued to the `mpuacc` register. Refer to the MPU Region Read and Write Operations section for more information. The RD flag always returns 0 when read by a `rdc1` instruction.

Related Information

[MPU Region Read and Write Operations](#) on page 65

3.4.2.14.6. The WR Flag

Setting the WR flag signifies that an MPU region write operation should be performed when a `wrc1` instruction is issued to the `mpuacc` register. Refer to the MPU Region Read and Write Operations section for more information. The WR flag always returns 0 when read by a `rdc1` instruction.

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

Related Information

[MPU Region Read and Write Operations](#) on page 65

3.4.2.14.7. The eccinj Register

The `eccinj` register injects 1 and 2 bit errors to the Nios II processor's internal RAM blocks that support ECC. Injecting errors allows the software to test the ECC error exception handling code. The error(s) are injected in the data bits, not the parity bits. The `eccinj` register is only available when ECC is present.

Table 36. eccinj Control Register Fields

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
Reserved										DC WB		DTCM 3		DTCM 2	
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
DTCM 1		DTCM 0		TLB		DC DAT		DC TAG		ICDAT		ICTAG		RF	

Software writes 0x1 to inject a 1 bit ECC error or 0x2 to inject a 2-bit ECC error to the RAM field. Hardware sets the value of the inject field to 0x0 after the error injection has occurred.

Table 37. eccinj Control Register Field Descriptions

Field	Description	Access	Reset	Available
RF	Inject an ECC error in the register file's RAM.	Read/Write	0	Only with ECC
ICTAG	Inject an ECC error in the instruction cache Tag RAM.	Read/Write	0	Only with ECC

continued...

Field	Description	Access	Reset	Available
ICDAT	Inject an ECC error in the instruction cache data RAM.	Read/Write	0	Only with ECC
DCTAG	Inject ECC error in data cache tag RAM.	Read/Write	0	
DCDAT	Inject an ECC error in the data cache data RAM. Injection occurs on next store instruction that writes the data cache or the next line fill.	Read/Write	0	
TLB	Inject an ECC error in the MMU TLB RAM. Errors are injected in the tag portion of the VPN field.	Read/Write	0	Only with ECC
DTCM0	Inject ECC error in DTCM0. Injection occurs on next store instruction that writes this DTCM.	Read/Write	0	
DTCM1	Inject ECC error in DTCM1. Injection occurs on next store instruction that writes this DTCM.	Read/Write	0	
DTCM2	Inject ECC error in DTCM2. Injection occurs on next store instruction that writes this DTCM.	Read/Write	0	
DTCM3	Inject ECC error in DTCM3. Injection occurs on next store instruction that writes this DTCM.	Read/Write	0	
DC WB	Inject ECC error in data cache victim line buffer RAM. Injection occurs on the first word written into the victim buffer RAM when a dirty line is being written back.	Read/Write	0	

Refer to “Working with ECC” for more information about when errors are injected.

Related Information

[Working with ECC](#) on page 67

3.4.3. 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`.

Note: 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 the Exception Processing section.



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*.

Related Information

- [Instruction Set Reference](#) on page 166
- [Exception Processing](#) on page 71

3.4.3.1. 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.

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 38. `sstatus` Control Register Fields

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

Table 39. `sstatus` Control Register Field Descriptions

Bit	Description	Access	Reset	Available
SRS ⁽⁸⁾	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	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	Undefined	⁽⁹⁾
NMI	NMI is the nonmaskable interrupt mode bit. The processor sets NMI to 1 when it takes a nonmaskable interrupt.	Read/Write	Undefined	⁽⁹⁾
PRS	⁽⁹⁾	Read/Write	Undefined	⁽⁹⁾
CRS	⁽⁹⁾	Read/Write	Undefined	⁽⁹⁾

continued...

⁽⁸⁾ If the EIC interface and shadow register sets are not present, SRS always reads as 0, and the processor behaves accordingly.

Bit	Description	Access	Reset	Available
IL	(9)	Read/Write	Undefined	(9)
IH	(9)	Read/Write	Undefined	(9)
EH	(9)	Read/Write	Undefined	(9)
U	(9)	Read/Write	Undefined	(9)
PIE	(9)	Read/Write	Undefined	(9)

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".

For details about `status.EH`, refer to the Nios II Processor Status After Taking Exceptions Table.

Related Information

- [The status Register](#) on page 45
- [Exceptions and Processor Status](#) on page 88
- [Requested Register Set](#) on page 78
- [The status Register](#) on page 45

3.4.3.1.1. 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.

(9) Refer to the status Control Register Field Descriptions Table



3.4.3.1.2. 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.

3.4.3.2. 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.

3.5. Working with the MPU

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

3.5.1. 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.MT`, 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 `wrc1` instruction to the `mpubase` register with the `mpubase.INDEX` and `mpubase.D` fields set to identify the MPU region.
- Execute a `wrc1` 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 `rdc1` instruction to the `mpubase` register to read the loaded the `mpubase` register value.
- Execute a `rdc1` 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.MT`, and `mpuacc.PERM` values for the MPU region.

Note: Values for the `mpubase` register are not actually retrieved until the `wrc1` 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.MT`, 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.

3.5.2. MPU Initialization

Your system software must provide a data structure that contains the region information described in the "Memory Regions" section of this chapter 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, Intel FPGA 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 the "MPU Region Read and Write Operations" section of this chapter. 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.


Note: 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 the "MPU Region Read and Write Operations" section of this chapter. Repeat this dual `wrctl` instruction sequence until all desired regions are defined.

Related Information

- [MPU Region Read and Write Operations](#) on page 65
- [Memory Regions](#) on page 40



3.5.3. 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.

3.6. Working with ECC

3.6.1. Enabling ECC

The ECC is disabled on system reset. Before enabling the ECC, initialize the Nios II RAM blocks to avoid spurious ECC errors.

The Nios II processor executes the `INITI` instruction on each cache line, which initializes the instruction cache RAM. The RAM does not require special initialization because any detected ECC errors are ignored if the line is invalid; the line is invalid after `INITI` instructions initialize the tag RAM.

Nios II processor instructions that write to every register (except register 0) initialize the register file RAM blocks. If shadow register sets are present, this step is performed for all registers in the shadow register set using the `WRPRS` instruction.

Nios II processor instructions that write every TLB RAM location initialize the MMU TLB RAM. This RAM does not require special initialization.

3.6.1.1. Disabling ECC

Disable ECC in software by writing 0 to `CONFIG.ECCEN`. Software can re-enable ECC without reinitializing the ECC-protected RAMs because the ECC parity bits are written to the RAM blocks even if ECC is disabled.

3.6.2. Handling ECC Errors

ECC error exceptions occur when unrecoverable ECC errors are detected. The software's ability to recover from the ECC error depends on the nature of the error.

Typically, software can recover from an unrecoverable MMU TLB ECC error (2 bit error) because the TLB is a software-managed cache of the operating system page tables stored in the main memory (e.g., SDRAM). Software can invalidate the TLB entry, return to the instruction that took the ECC error exception, and execute the TLB's mishandled code to load a TLB entry from the page tables.

In general, software cannot recover from a register file ECC error (2 bit error) because the correct value of a register is not known. If the exception handler reads a register that has a 2 bit ECC error associated with it, another ECC error occurs and an exception handler loop can occur.

Exception handler loops occur when an ECC error exception occurs in the exception handler before it is ready to handle nested exceptions. To minimize the occurrence or exception handler loops, locate the ECC error exception handler code in normal cacheable memory, ensure that all data accesses are to non-cacheable memory, and minimize register reading.

The ECC error signals (`ecc_event_bus`) provide the `EEH` signal for external logic to detect a possible exception handler loop and reset the Nios II processor.

3.6.3. Injecting ECC Errors

This section describes the code sequence for injecting ECC errors for each ECC-protected RAM, assuming the ECC is enabled and interrupts are disabled for the duration of the code sequence.

3.6.3.1. Instruction Cache Tag RAM

1. Ensure all code up to the JMP instruction is in the same instruction cache line or is located in an ITCM.
2. Use a `FLUSHI` instruction to flush an instruction cache line other than the line containing the executing code.
3. Use a `FLUSHP` instruction to flush the pipeline.
4. Use a `WRCTL` instruction to set `ECCINJ.ICTAG` to `INJS` or `INJD`. This setting causes an ECC error to occur on the start of the next line fill.
5. Use a `JMP` instruction to jump to an instruction address in the flushed line.
6. The ECC error is injected when writing the tag RAM at the start of the line fill.
7. Use a `RDCTL` instruction to ensure that the value of `ECCINJ.ICTAG` is `NOINJ`.
8. The ECC error triggers after the target of the `JMP` instruction.

3.6.3.2. Instruction Cache Data RAM

1. Ensure all code up to the JMP instruction is in the same instruction cache line or is located in an ITCM.
2. Use a `FLUSHI` instruction to flush an instruction cache line other than the line containing the executing code.
3. Use a `FLUSHP` instruction to flush the pipeline.
4. Use a `WRCTL` instruction to set `ECCINJ.ICDAT` to `INJS` or `INJD`. This setting causes an ECC error to occur on the start of the next line fill.
5. Use a `JMP` instruction to jump to an instruction address in the flushed line.
6. The ECC error is injected when writing the tag RAM at the start of the line fill.
7. Use a `RDCTL` instruction to ensure that the value of `ECCINJ.ICDAT` is `NOINJ`.
8. Execute the target of the `JMP` instruction twice (first to inject the ECC error and second to be triggered by it).



3.6.3.3. ITCMs

Software running on the Nios II cannot directly inject an ECC error in an ITCM because the Nios II only writes ITCMs when correcting ECC errors. To inject an ECC in an ITCM, the TCM RAM must also be connected to a DTCM master. The Nios II provided DTCM error injection mechanism (i.e. `ECCINJ` register) is used to inject an error in the TCM RAM as follows:

1. Use a `WRCTL` instruction to set `ECCINJ` so that it injects ECC errors in the DTCM connected to the ITCM.
2. Use a `STW` instruction to write the DTCM.
3. Use a `RDCTL` instruction to ensure the value of the `ECCINJ` field written by the `WRCTL` is `NOINJ`.
4. Use a `JMP` instruction to jump to an instruction address in the ITCM.
5. The ECC error should be triggered on the target of the `JMP` instruction.

3.6.3.4. Register File RAM Blocks

1. Use a `WRCTL` instruction to set `ECCINJ.RF` to `INJS` or `INJD` (as desired).
2. Execute any instruction that writes any register except `R0`.
3. Use a `RDCTL` instruction to ensure that the value of `ECCINJ.RF` is `NOINJ`.
4. Use an instruction to read the desired register from `rA` such as `OR rd, r0, rx` where `rx` is the register written in the previous step. This action triggers the ECC error.
5. Use an instruction to read the desired register from `rB` such as `OR rd, rx, r0` where `rx` is the register written in the previous step.

3.6.3.5. Data Cache Tag RAM

1. Use a `LOAD` instruction from a data address to get the line in the cache. The line should be clean.
2. Use a `WRCTL` instruction to set `ECCINJ.DCTAG` to `INJS` or `INJD`.
3. Use a `STORE` instruction from a data address mapped to that line. The `STORE` instruction should hit in the data cache and write the tag RAM to set the dirty bit.
4. The ECC error is injected when the tag RAM is written.
5. Use a `RDCTL` instruction to ensure the value of `ECCINJ.DCTAG` is `NOINJ`. Before the `RDCTL`, use a `FLUSHP` instruction to avoid the RAW hazard on `ECCINJ`.
6. Do another `LOAD` or `STORE` instruction to the same line.
7. The ECC error should be triggered on this second `LOAD/STORE` instruction.

3.6.3.6. Data Cache Data RAM (Clean Line)

1. Use a `FLUSHDA` instruction to ensure the line isn't in the data cache.
2. Use a `LOAD` instruction to load a clean data cache line.
3. Use a `WRCTL` instruction to set `ECCINJ.DCDAT` field to `INJS` or `INJD`.

4. Use a `LOAD` instruction to an address in the data cache line to inject the error.
5. Use a `RDCTL` instruction to ensure the values of the field written by the `WRCTL` to `ECCINJ` is `NOINJ`. Before the `RDCTL`, use a `FLUSHP` instruction to avoid the RAW hazard on `ECCINJ`.
6. Use a `LOAD` instruction from the same address.
7. The ECC error should be triggered on the `LOAD` instruction.

3.6.3.7. Data Cache Data RAM (Dirty Line)

1. Use a `LOAD` instruction to load a data cache line.
2. Use a `WRCTL` instruction to set `ECCINJ.DCDAT` field to `INJS` or `INJD` (as desired).
3. Use a `STORE` instruction to an address in the data cache line.
4. Use a `RDCTL` instruction to ensure the values of the field written by the `WRCTL` to `ECCINJ` is `NOINJ`. Before the `RDCTL`, use a `FLUSHP` instruction to avoid the RAW hazard on `ECCINJ`.
5. Either use a `LOAD` instruction from the same address or trigger a writeback of the dirty line (e.g. `FLUSHDA` instruction)
6. The ECC error should be triggered on the `LOAD` instruction unless it is only detected during the writeback of a dirty line. In the writeback of a dirty line case, the ECC error is triggered an undefined number of instructions later.

3.6.3.8. Data Cache Victim Line Buffer RAM

1. Use a `LOAD` instruction to load a data cache line.
2. Use a `WRCTL` instruction to set `ECCINJ.DCWB` field to `INJS` or `INJD` (as desired).
3. Use a `STORE` instruction to an address in the data cache line.
4. Use a `RDCTL` instruction to ensure the values of the field written by the `WRCTL` to `ECCINJ` is `NOINJ`. Before the `RDCTL`, use a `FLUSHP` instruction to avoid the RAW hazard on `ECCINJ`.
5. Either use a `LOAD` instruction from the same address or trigger a writeback of the dirty line (e.g. `FLUSHDA` instruction)
6. The ECC error should be triggered on the `LOAD` instruction unless it is only detected during the writeback of a dirty line. In the writeback of a dirty line case, the ECC error is triggered an undefined number of instructions later.

3.6.3.9. DTCMs

1. Use a `WRCTL` instruction to set the `ECCINJ.DTCM` field to `INJS` or `INJD` for the desired DTCM.
2. Use a `STW` instruction to write an address in the DTCM.
3. Use a `RDCTL` instruction to ensure the value of the field written by the `WRCTL` to `ECCINJ` is `NOINJ`.
4. Use a `LOAD` instruction from the same address in the DTCM.
5. The ECC error should be triggered on the `LOAD` instruction.



3.6.3.10. MMU TLB RAM

1. Use a WRCTL instruction to set ECCINJ.TLB to INJS or INJD.
2. Use a WRCTL instruction to write a TLB entry. The ECC error is injected at this time and any associated uTLB entry can be flushed.
3. Use a RDCTL instruction to ensure the value of ECCINJ.TLB is NOINJ.
4. Perform an instruction/data access to cause the hardware to read the TLB entry (copied into uTLB) and the ECC decoder should detect the ECC error at this time. Alternatively, initiate a software read of the TLB (by writing TLBMISC.RD to 1).
5. If a software read was initiated, the TLBMISC.EE field should be set to 1 on any instruction after the WRCTL that triggered the software read.
6. If a hardware read was initiated, the ECC error should be triggered on the first instruction after the hardware read.

3.7. 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.

3.7.1. Terminology

Intel FPGA Nios II Gen 2 and 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 Intel-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 non-overhead 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.
- *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.

3.7.2. 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 you specify in the Nios II processor IP core setup parameters.
- *Break exception*—Occurs when the JTAG debug module requests control. Control is transferred to the break address you specify in the Nios II processor IP core setup parameters.
- *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 the Nios II Exceptions Table. Control is transferred to the exception address you specify in the Nios II processor IP core setup parameters.

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 40. Nios II Exceptions (In Decreasing Priority Order)

Exception	Type	Available	Cause	Address	Vector
Reset	Reset	Always	0		Reset
Hardware break	Break	Always	—		Break
Processor-only reset request	Reset	Always	1		Reset
ECC Data Cache Writeback Error	Instruction-related	ECC and data cache	22		General exception
Internal interrupt	Interrupt	Internal interrupt controller	2	ea-4 ⁽¹¹⁾	General exception
External nonmaskable interrupt	Interrupt	External interrupt controller interface	—	ea-4 ⁽¹¹⁾	Requested handler address ⁽¹²⁾
<i>continued...</i>					



Exception	Type	Available	Cause	Address	Vector
External maskable interrupt	Interrupt	External interrupt controller interface	2	ea-4 ⁽¹¹⁾	Requested handler address ⁽¹²⁾
ECC TLB error (instruction)	Instruction-related	MMU and ECC	18	ea-4 ⁽¹¹⁾	General exception
Supervisor-only instruction address ⁽¹⁰⁾	Instruction-related	MMU	9	ea-4 ⁽¹¹⁾	General exception
Fast TLB miss (instruction) ⁽¹⁰⁾	Instruction-related	MMU	12	pteaddr.VPN, ea-4 ⁽¹¹⁾	Fast TLB Miss exception
Double TLB miss (instruction) ⁽¹⁰⁾	Instruction-related	MMU	12	pteaddr.VPN, ea-4 ⁽¹¹⁾	General exception
TLB permission violation (execute) ⁽¹⁰⁾	Instruction-related	MMU	13	pteaddr.VPN, ea-4 ⁽¹¹⁾	General exception
ECC register file error	Instruction-related	ECC	20	ea-4 ⁽¹¹⁾	General exception
MPU region violation (instruction) ⁽¹⁰⁾	Instruction-related	MPU	16	ea-4 ⁽¹¹⁾	General exception
Bus Instruction Fetch Error		M Core	23	ea-4 ⁽¹¹⁾	General exception
ECC Fetch Error (instruction fetch)		ECC and ITCM	19	ea-4 ⁽¹¹⁾	General exception
ECC Register File Error		ECC	20	ea-4 ⁽¹¹⁾	General exception
Supervisor-only instruction	Instruction-related	MMU or MPU	10	ea-4 ⁽¹¹⁾	General exception
Trap instruction	Instruction-related	Always	3	ea-4 ⁽¹¹⁾	General exception
Illegal instruction	Instruction-related	Illegal instruction detection on, MMU, or MPU	5	ea-4 ⁽¹¹⁾	General exception
Unimplemented instruction	Instruction-related	Always	4	ea-4 ⁽¹¹⁾	General exception
Break instruction	Instruction-related	Always	—	ba-4 ⁽¹¹⁾	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
ECC TLB error (data)	Instruction-related	MMU and ECC	18	badaddr (data address)	General exception
Division error	Instruction-related	Division error detection on	8	ea-4 ⁽¹¹⁾	General exception

continued...

Exception	Type	Available	Cause	Address	Vector
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
MPU region violation (data)	Instruction-related	MPU	17	badaddr (data address)	General exception
Bus Data Region Violation		M core	24	badaddr (data address)	General exception
ECC Data Error		ECC and (data cache OR DTCM)	21	badaddr (data address)	General exception

Related Information

- [Requested Handler Address](#) on page 78
- [General-Purpose Registers](#) on page 42

3.7.3. 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.

3.7.3.1. 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.

3.7.4. Reset Exceptions

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

-
- (10) It is possible for any instruction fetch to cause this exception.
 - (11) Refer to the Nios II General-Purpose Registers Table for descriptions of the `ea` and `ba` registers.
 - (12) For a description of the requested handler address, refer to the Requested Handler Address section of this chapter.



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.

Note: 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.

Note: 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
- Nios II C-to-hardware (C2H) acceleration compiler logic.

For more information refer to the Nios II Custom Instruction User Guide for reset conditions.

Related Information

[Nios II Custom Instruction User Guide](#)

3.7.5. 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.

3.7.5.1. 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.

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

1. Writes the address of the instruction following the break to the `ba` register (`r30`) in the normal register set.
2. Clears `status.U` to zero, forcing the processor into supervisor mode, when the system contains an MMU or MPU.
3. Sets `status.EH` to one, indicating the processor is handling an exception, when the system contains an MMU.
4. Copies `status.CRS` to `status.PRS` and then sets `status.CRS` to 0.
5. Transfers execution to the break handler, stored at the break vector specified in the Nios II Processor parameter editor.


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

3.7.5.2. Understanding 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.

3.7.5.3. 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.



3.7.6. 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

3.7.6.1. 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.

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

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

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.

For a typical EIC implementation, refer to the *Vectored Interrupt Controller* chapter in the *Embedded Peripherals IP User Guide*.

Related Information

- [Embedded Peripherals IP User Guide](#)
- [Requested NMI Mode](#) on page 78
- [Requested Register Set](#) on page 78
- [Requested Interrupt Level](#) on page 78

3.7.6.1.1. 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.

3.7.6.1.2. 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.

3.7.6.1.3. 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.

3.7.6.1.4. 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.



3.7.6.1.5. 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 the Requested Register Set section of this chapter.

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

Related Information

[Requested Register Set](#) on page 78

3.7.6.2. 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:

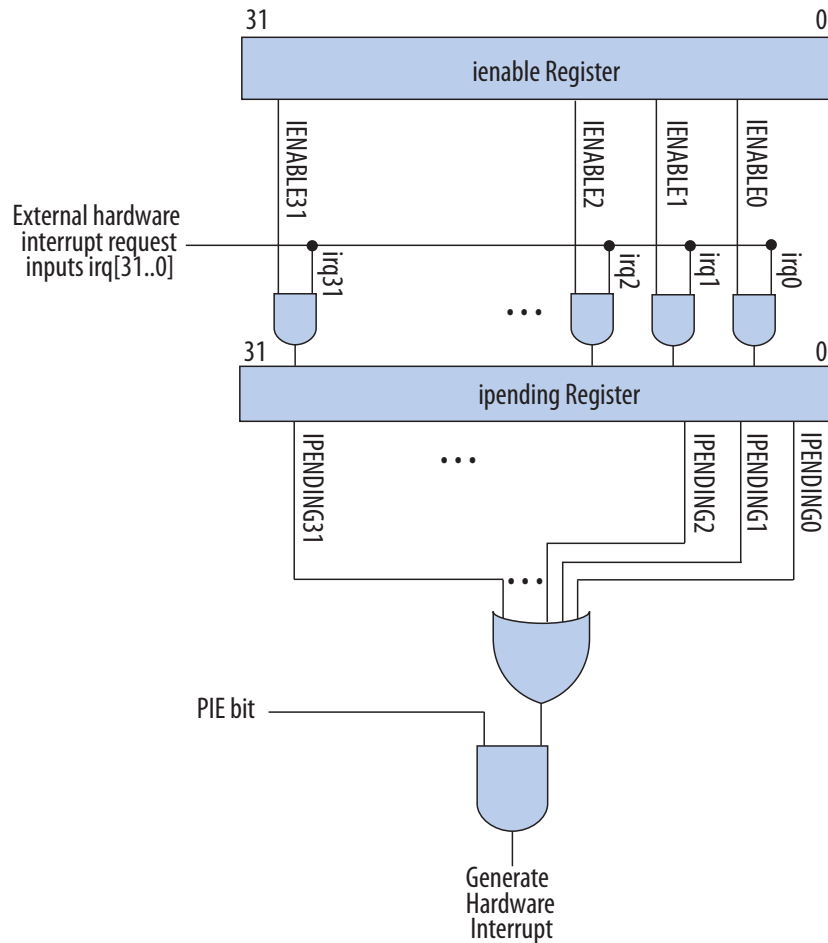
- The `PIE` bit of the `status` control register is one.
- An interrupt-request input, `irqn`, is asserted.
- The corresponding bit `n` of the `ienable` control register is one.

Upon hardware interrupt, the processor clears the `PIE` bit to zero, disabling further interrupts, and performs the other steps outlined in the "Exception Processing Flow" section of this chapter.

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.

Note: 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 5. Relationship Between ienable, ipending, PIE and Hardware Interrupts



Related Information


Exception Processing Flow on page 85

3.7.7. Instruction-Related Exceptions

Instruction-related exceptions occur during execution of Nios II instructions. When they occur, the processor perform the steps outlined in the "Exception Processing Flow" section of this chapter.

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

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

Related Information

[Exception Processing Flow](#) on page 85

3.7.7.1. Trap Instruction

When a program issues the `trap` instruction, the processor 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.

3.7.7.2. Break Instruction

The break instruction is treated as a break exception. For more information, refer to the "Break Exceptions" section of this chapter.

Related Information

[Break Exceptions](#) on page 75

3.7.7.3. 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.

For more information, refer to the "Potential Unimplemented Instructions" section of this chapter.

Related Information

[Potential Unimplemented Instructions](#) on page 102

3.7.7.4. 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. Illegal instruction checking is always on regardless of MMU or MPU settings.

For information about controlling this option, refer to the *Instantiating the Nios II Processor* 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.

Note:

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 information about each specific Nios II core.

Related Information

- [Instruction Set Reference](#) on page 166
- [Programming Model](#) on page 33
- [Nios II Core Implementation Details](#) on page 118

3.7.7.5. 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 the "Operating Modes" section of this chapter for more information.

Related Information

[Operating Modes](#) on page 33

3.7.7.6. 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 the Virtual Memory Partitions Table.

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

Related Information

[Virtual Memory Address Space](#) on page 36



3.7.7.7. 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.

3.7.7.8. 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.

For information about controlling this option, refer to the *Instantiating the Nios II Processor* 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.

Related Information

[Programming Model](#) on page 33

3.7.7.9. 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.

For information about controlling this option, refer to the *Instantiating the Nios II Processor* 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.

Related Information

[Programming Model](#) on page 33

3.7.7.10. Division Error

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

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.

Related Information

[Programming Model](#) on page 33

3.7.7.11. 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 with the Nios II Processor parameter editor in Platform Designer, 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.

3.7.7.12. 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.



3.7.7.13. 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.

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

3.7.7.14. 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.

3.7.8. 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.

3.7.9. Exception Processing Flow

Except for the break exception (refer to the Processing a Break section of this chapter), this section describes how the processor responds to exceptions, including interrupts and instruction-related exceptions.

Related Information

- [Exception Handling](#)
For details about 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 a Break](#) on page 76

3.7.9.1. 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 with the Nios II Processor parameter editor in Platform Designer, is called the exception vector in the Nios II Processor parameter editor. 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.

Note: 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 with the Nios II Processor parameter editor in Platform Designer, is called the fast TLB miss exception vector in the Nios II Processor parameter editor.

3.7.9.2. 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 41. 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 ⁽¹³⁾	Yes	Yes	Yes

The Nios II processor supports fast nested interrupts with shadow register sets, as described in the "Shadow Register Set" section of this chapter.

Keeping `status.PIE` set allows higher level interrupts to be taken immediately, without requiring the interrupt handler to set `status.PIE` to 1.

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.

Related Information

- [Shadow Register Sets](#) on page 79
- [Nios II Core Implementation Details](#) on page 118

3.7.9.3. 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 as well as before re-enabling `PIE`. The ISR also saves `estatus` and `ea (r29)` before re-enabling `PIE`.

⁽¹³⁾ 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.

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" for more information.

Related Information

[Handling Nested Exceptions](#) on page 91

3.7.9.4. Exceptions and Processor Status

The Nios II Processor Status After Taking Exception Table 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.

Table 42. Nios II Processor Status After Taking Exception

Processor Status Register or Field	System Status Before Taking Exception						
	External Interrupt Asserted (14)				Internal Interrupt Asserted or Noninterrupt Exception		
	status.EH==1 (29)		status.EH==0		status.EH==1	status.EH==0	
	RRS==0 (30)	RRS!=0	RRS==0	RRS!=0		TLB Miss (31)	No TLB Miss
TLB Permission Violation (31)					No TLB Permission Violation		
<code>pteaddr.VPN</code> (15)	No change				VPN (16)		No change
<code>status.PRS</code> (30)	No change		<code>status.CRS</code> (30) (32)		No change		
<code>pc</code>	RHA				General exception vector (17)	Fast TLB exception vector (18)	General exception vector(30)
<code>sstatus</code> (19)(33)	No change			<code>status</code> (32) (20)	No change		
<code>estatus</code> (33)	No change		<code>status</code> (32)	No change			<code>status</code> (32)
<code>ea</code>	No change		return address (21)		No change	return address	
<code>tlbmisc.D</code> (29)	No change				(22)		

continued...

- (14) If the Nios II processor does not have an EIC interface, external interrupts do not occur.
- (15) If the Nios II processor does not have an MMU, this register is not implemented.
- (16) The VPN of the address triggering the exception
- (17) Invokes the general exception handler
- (18) Invokes the fast TLB miss exception handler
- (19) If the Nios II processor does not have shadow register sets, this register is not implemented.
- (20) `sstatus.SRS` is set to 1 if `RRS` is not equal to `status.CRS`.
- (21) The address following the instruction being executed when the exception occurs
- (22) Set to 1 on a data access exception, set to 0 otherwise



Processor Status Register or Field	System Status Before Taking Exception							
	External Interrupt Asserted (14)				Internal Interrupt Asserted or Noninterrupt Exception			
	status.EH==1 (29)		status.EH==0		status.EH==1	status.EH==0		
	RRS==0 (30)	RRS!=0	RRS==0	RRS!=0		TLB Miss (31)	No TLB Miss	
TLB Permission Violation (31)					No TLB Permission Violation			
tlbmisc.DBL (29)	No change				(23)			
tlbmisc.PERM(29)	No change				(24)			
tlbmisc.BAD (29)	No change				(25)			
status.PIE	No change				0 (26)			
status.EH(29)	No change				1 (27)			
status.IH (35)	1				No change			
status.NMI (35)	RNMI				No change			
status.IL (35)	RIL				No change			
status.RSIE (30)(35)	0				No change			
status.CRS (30)	RRS				No change			
status.U (29)	0 (28)							

- (14) If the Nios II processor does not have an EIC interface, external interrupts do not occur.
- (23) Set to 1 on a double TLB miss, set to 0 otherwise
- (24) Set to 1 on a TLB permission violation, set to 0 otherwise
- (25) Set to 1 on a bad virtual address exception, set to 0 otherwise
- (26) Disables exceptions and nonmaskable interrupts
- (27) If the MMU is implemented, indicates that the processor is handling an exception.
- (28) Puts the processor in supervisor mode.
- (29) 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.
- (30) 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.
- (31) If the Nios II processor does not have an MMU, TLB-related exceptions do not occur.
- (32) The pre-exception value
- (33) Saves the processor's pre-exception status
- (34) If the MMU is implemented, indicates that the processor is handling an exception.
- (35) If the Nios II processor does not have an EIC interface, this field is not implemented.

3.7.10. 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.

3.7.10.1. Nios II/f Exception Processing

The `CAUSE` field of the `exception` register contains a code for the highest-priority exception occurring at the time. The `BADDR` field of the `badaddr` register contains the byte instruction address or data address for certain exceptions.

Refer to the Nios II Exceptions table for more information in the Exception Overview section.

Note: 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.

3.7.10.2. Nios II/e Exception Processing

Example 3. Determining Exception Cause for Nios II/e Exception Processing

```

/* 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)
}
}

```



3.7.11. 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 keep 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 3–44.

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 the Instruction-Related Exceptions section of this chapter.

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.

Related Information

- [Exception Processing Flow](#) on page 85
- [Instruction-Related Exceptions](#) on page 80
- [Processor Architecture](#) on page 14

3.7.11.1. 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.

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

3.7.11.2. 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 certain conditions. With the following conditions satisfied, ISRs need not save and restore register contents on entry and exit:

- Automatic nested interrupts are enabled.
- Each interrupt is assigned to a dedicated shadow register set.
- All interrupts with the same RIL are assigned to dedicated shadow register sets.
- Multiple interrupts with different RILs can be assigned to a single shadow register set. However, with multiple register sets, you must not allow the RILs assigned to one shadow register set to overlap the RILs assigned to another register set.

The following tables demonstrate the validity of register set assignments when preemption within a register set is enabled.

Table 43. Example of Illegal RIL Assignment

RIL	Register Set 1	Register Set 2
1	IRQ0	
2	IRQ1	
3		IRQ2
4	IRQ3	
5		IRQ4
6		IRQ5
7		IRQ6


Table 44. Example of Legal RIL Assignment

RIL	Register Set 1	Register Set 2
1	IRQ0	
2	IRQ1	
3	IRQ3	
4		IRQ2
5		IRQ4
6		IRQ5
7		IRQ6

Note: 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.



Note: 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.

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 the Nested "Exceptions with Internal Interrupt Controller" section of this chapter.

Related Information

[Nested Exceptions with the Internal Interrupt Controller](#) on page 91

3.7.12. Handling Nonmaskable Interrupts

Writing an NMI handler involves the same basic techniques as writing any other interrupt handler. However, nonmaskable interrupts always preempt maskable interrupts, and cannot be preempted. This knowledge 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, when each NMI is assigned to a dedicated shadow register set. Therefore, NMIs can be handled transparently.

Note: 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.

Note: 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.

3.7.13. 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.

3.7.13.1. 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.

3.7.13.2. Masking Interrupts with an External Interrupt Controller

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

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.

3.7.13.3. 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 `irqn` interrupt is enabled; a bit value of zero means that the corresponding interrupt is disabled.

Refer to the "Exception Processing" section of this chapter for more information.

An ISR can adjust `ienable` so that IRQs of equal or lower priority are disabled. Refer to the "Handling Nested Exceptions" section of this chapter for more information.

Related Information

- [Handling Nested Exceptions](#) on page 91
- [Exception Processing](#) on page 71

3.7.13.4. 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 the "Nested Exceptions with an External Interrupt Controller" section of this chapter.
- 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.

Note: 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`.

Note: The `eret` instruction 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.

Note: 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.

Related Information

[Nested Exceptions with the Internal Interrupt Controller](#) on page 91

3.7.13.4.1. 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.

3.8. Memory and Peripheral Access

Nios II addresses are 32 bits, allowing access up to a 4-gigabyte address space. The MMU supports the full 32-bit physical address. Bit 31 bypass is optional, you can access full 32-bit addressing without the MMU.

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 uses little-endian byte ordering. For data wider than 8 bits stored in memory, the more-significant bits are located in higher addresses.

The Nios II architecture supports register and immediate addressing.

Related Information

[Nios II Core Implementation Details](#) on page 118

3.8.1. 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 the optimal bit-31 cache bypass. With bit-31 cache bypass, the address space of processor cores is 2 GB, 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 the [Address Space and Memory Partitions](#) section, and the [TLB Organization](#) section of this chapter.
- 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" section of this chapter for more information.

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

Related Information

- [Cache and Tightly Coupled Memory](#)
- [Reset Exceptions](#) on page 74
- [TLB Organization](#) on page 38
- [Address Space and Memory Partitions](#) on page 36
- [Instruction Set Reference](#) on page 166
- [Processor Architecture](#) on page 14



- [Nios II Core Implementation Details](#) on page 118

3.8.1.1. 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 KB), 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 KB 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 KB, 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-KB 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.

3.9. Instruction Set Categories

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

3.9.1. 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 as well as write buffering to improve memory bandwidth. The architecture provides instructions for both cached and uncached accesses.

Table 45. Wide Data Transfer Instructions

Instruction	Description
ldw stw	The <code>ldw</code> and <code>stw</code> 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 <code>ldwio</code> and <code>stwio</code> .
ldwio stwio	<code>ldwio</code> and <code>stwio</code> instructions load and store 32-bit data words from/to peripherals without caching and buffering. Access cycles for <code>ldwio</code> and <code>stwio</code> instructions are guaranteed to occur in instruction order and are never suppressed.

Table 46. 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 in the following table cell.
ldbio ldbuio stbio ldhio ldhuio sthio	These operations load/store byte and half-word data from/to peripherals without caching or buffering.

3.9.2. Arithmetic and Logical Instructions

Logical instructions support `and`, `or`, `xor`, and `nor` operations. Arithmetic instructions support addition, subtraction, multiplication, and division operations.

Table 47. Arithmetic and Logical Instructions

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 <code>and</code> , <code>or</code> , and <code>xor</code> 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 <code>and</code> , <code>or</code> , and <code>xor</code> , 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.
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 <code>add</code> , <code>sub</code> , and <code>mul</code> 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 <code>mul</code> .
mulxsu	This instruction is used in computing a 128-bit result of a 64x64 signed multiplication.

3.9.3. Move Instructions

These instructions provide move operations to copy the value of a register or an immediate value to another register.



Table 48. Move Instructions

Instruction	Description
mov movhi movl movui movia	mov copies the value of one register to another register. movl moves a 16-bit signed immediate value to a register, and sign-extends the value to 32 bits. movui and movhi move a 16-bit immediate 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.

3.9.4. 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.

Table 49. Comparison Instructions

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

3.9.5. 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.

Table 50. Shift and Rotate Instructions

Instruction	Description
rol ror roli	The <code>rol</code> and <code>roli</code> instructions provide left bit-rotation. <code>roli</code> uses an immediate value to specify the number of bits to rotate. The <code>ror</code> instructions provides right bit-rotation. There is no immediate version of <code>ror</code> , because <code>roli</code> can be used to implement the equivalent operation.
sll slli sra srl srai srli	These shift instructions implement the <code><<</code> and <code>>></code> operators of the C programming language. The <code>sll</code> , <code>slli</code> , <code>srl</code> , <code>srli</code> instructions provide left and right logical bit-shifting operations, inserting zeros. The <code>sra</code> and <code>srai</code> instructions provide arithmetic right bit-shifting, duplicating the sign bit in the most significant bit. <code>slli</code> , <code>srli</code> and <code>srai</code> use an immediate value to specify the number of bits to shift.

3.9.6. Program Control Instructions

The Nios II architecture supports the unconditional jump, branch, and call instructions. These instructions do not have delay slots.

Table 51. Unconditional Jump and Call Instructions

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 <code>ra</code> .
callr	This instruction calls a subroutine at the absolute address contained in a register, and stores the return address in register <code>ra</code> . This instruction serves the roll of dereferencing a C function pointer.
ret	The <code>ret</code> instruction is used to return from subroutines called by <code>call</code> or <code>callr</code> . <code>ret</code> loads and executes the instruction specified by the address in register <code>ra</code> .
jmp	The <code>jmp</code> instruction jumps to an absolute address contained in a register. <code>jmp</code> is used to implement switch statements of the C programming language.
jmp <i>i</i>	The <code>jmp<i>i</i></code> instruction jumps to an absolute address using an immediate value to determine the absolute address.
br	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. The conditional branches support the following 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 52. 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 the "Comparison Instructions" section of this chapter for a description of the relational operations implemented.



Related Information

[Comparison Instructions](#) on page 99

3.9.7. Other Control Instructions

Table 53. Other Control Instructions

Instruction	Description
trap eret	The <code>trap</code> and <code>eret</code> instructions generate and return from exceptions. These instructions are similar to the <code>call/ret</code> pair, but are used for exceptions. <code>trap</code> saves the <code>status</code> register in the <code>estatus</code> register, saves the return address in the <code>ea</code> register, and then transfers execution to the general exception handler. <code>eret</code> returns from exception processing by restoring <code>status</code> from <code>estatus</code> , and executing the instruction specified by the address in <code>ea</code> .
break bret	The <code>break</code> and <code>bret</code> instructions generate and return from breaks. <code>break</code> and <code>bret</code> 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 <code>status</code> register. The value is read from or stored to a general-purpose register.
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. <code>wrprs</code> can set <code>r0</code> to 0 in a shadow register set. System software must use <code>wrprs</code> to initialize <code>r0</code> to 0 in each shadow register set before using that register set.

3.9.8. Custom Instructions

The `custom` instruction provides low-level access to custom instruction logic. The inclusion of custom instructions is specified with the Nios II Processor parameter editor in Platform Designer, and the function implemented by custom instruction logic is design dependent.

For more information, refer to the “Custom Instructions” section of the *Processor Architecture* chapter of the *Nios II Processor Reference Handbook*

For continued more information refer to the *Nios II Custom Instruction User Guide*.

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

Related Information

- [Nios II Custom Instruction User Guide](#)
- [Processor Architecture](#) on page 14

3.9.9. No-Operation Instruction

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

3.9.10. 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 the "Unimplemented Instruction" section of this chapter for more information regarding unimplemented instruction processing.

Related Information

[Unimplemented Instruction](#) on page 81

3.10. Programming Model Revision History

Document Version	Changes
2020.04.24	Added clarification about the MT flag and cache access in section: <i>The MT Flag</i> .
2019.12.20	<ul style="list-style-type: none"> • Corrections to <code>BASE</code> field description in <i>The mpubase Register</i> • Correction to <i>The LIMIT Field</i>
2019.04.30	Maintenance release
2018.04.18	<ul style="list-style-type: none"> • Implemented editorial enhancements. • Corrected information about <code>mpubase</code>, and <code>mpuacc</code> register. • Corrected the Table: <i>MASK Region Size Encodings</i>.
2016.10.28	Removed extra exception information option.
2015.04.02	Initial release

4. Instantiating the Nios II Processor

This chapter describes the Nios II Processor parameter editor in Platform Designer. The Nios II Processor parameter editor allows you to specify the processor features for a particular Nios II hardware system. This chapter covers the features of the Nios II processor that you can configure with the Nios II Processor parameter editor; it is not a user guide for creating complete Nios II processor systems.

To get started designing custom Nios II systems, refer to the Nios II Hardware Development Tutorial.

Development kits for Intel FPGA devices, available on the All Development Kits page of the Intel FPGA website, also provide ready-made hardware design examples that demonstrate different configurations of the Nios II processor.

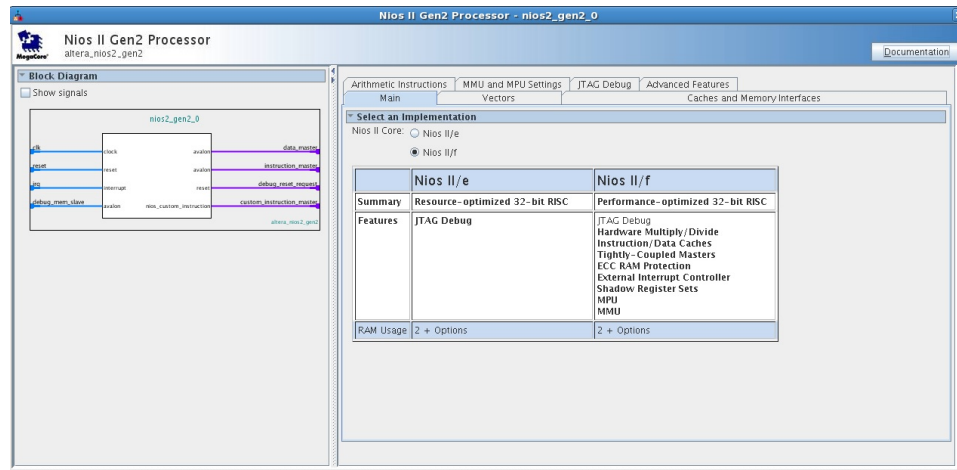
Related Information

[Nios II Hardware Development Tutorial](#)

4.1. Main Nios II Tab

The main purpose of the **Main** tab is to select the processor core. The core you select on this tab affects other options available on this and other tabs.

Figure 6. Nios II Platform Designer Main Tab



Intel FPGA offers the following Nios II cores:

- **Nios II/f**—The Nios II/f **fast** core is designed for fast performance. As a result, this core presents the most configuration options allowing you to fine tune the processor for performance.
- **Nios II/e**—The Nios II/e **economy** core is designed to achieve the smallest possible core size. As a result, this core has a limited feature set, and many settings are not available when the Nios II/e core is selected.

The **Main** tab displays a selector guide table that lists the basic properties of each core.

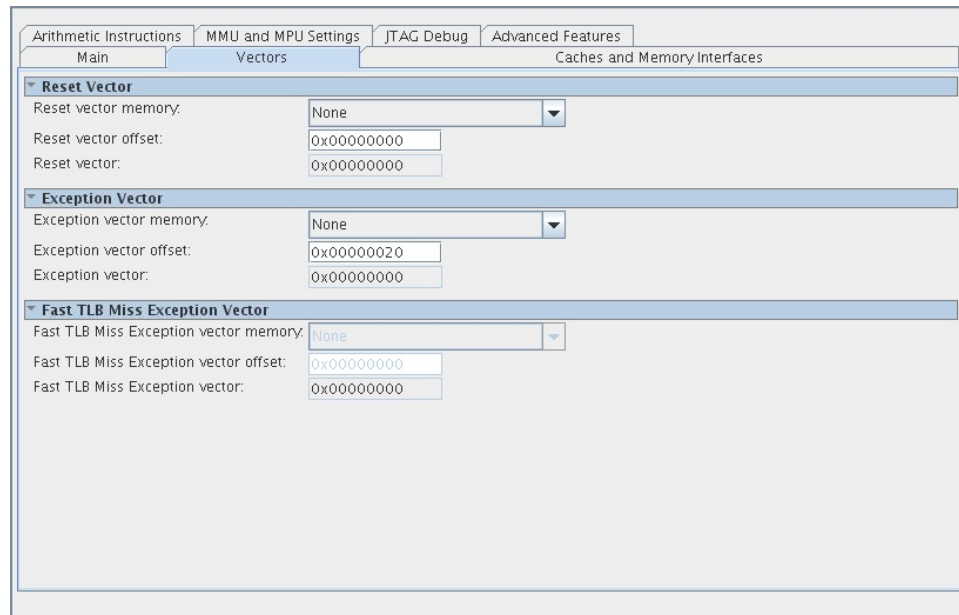
For implementation information about each core, refer to the *Nios II Core Implementation Details* chapter of the *Nios II Processor Reference Handbook*.

Related Information

[Nios II Core Implementation Details](#) on page 118

4.2. Vectors Tab

Figure 7. Nios II Platform Designer Vector Tab



4.2.1. Reset Vector

Parameters in this section select the memory module where the reset code (boot loader) resides, and the location of the reset vector (reset address). The reset vector cannot be configured until your system memory components are in place.

The **Reset vector memory** list, which includes all memory modules mastered by the Nios II processor, selects the reset vector memory module. In a typical system, select a nonvolatile memory module for the reset code.



Reset vector offset specifies the location of the reset vector relative to the memory module's base address. Platform Designer calculates the physical address of the reset vector when you modify the memory module, the offset, or the memory module's base address. In Platform Designer, **Reset vector** displays the read-only, calculated address. The address is always a physical address, even when an MMU is present.

Note: Platform Designer provides an **Absolute** option, which allows you to specify an absolute address in **Reset vector offset**. Use an absolute address when the memory storing the reset handler is located outside of the processor system and subsystems of the processor system.

For information about reset exceptions, refer to the *Programming Model* chapter of the *Nios II Processor Reference Handbook*.

Related Information

[Programming Model](#) on page 33

4.2.2. Exception Vector

Parameters in this section select the memory module where the general exception vector (exception address) resides, and the location of the general exception vector. The general exception vector cannot be configured until your system memory components are in place.

The **Exception vector memory** list, which includes all memory modules mastered by the Nios II processor, selects the exception vector memory module. In a typical system, select a low-latency memory module for the exception code.

Exception vector offset specifies the location of the exception vector relative to the memory module's base address. Platform Designer calculates the physical address of the exception vector when you modify the memory module, the offset, or the memory module's base address. In Platform Designer, **Exception vector** displays the read-only, calculated address.. The address is always a physical address, even when an MMU is present.

Note: Platform Designer provides an **Absolute** option, which allows you to specify an absolute address in **Exception vector offset**. Use an absolute address when the memory storing the exception handler is located outside of the processor system and subsystems of the processor system.

For information about exceptions, refer to the *Programming Model* chapter of the *Nios II Processor Reference Handbook*.

Related Information

[Programming Model](#) on page 33

4.2.3. Fast TLB Miss Exception Vector

The fast TLB miss exception vector is a special exception vector used exclusively by the MMU to handle TLB miss exceptions. Parameters in this section select the memory module where the fast TLB miss exception vector (exception address) resides, and the location of the fast TLB miss exception vector. The fast TLB miss exception vector cannot be configured until your system memory components are in place.



The Fast TLB Miss Exception vector memory list, which includes all memory modules mastered by the Nios II processor, selects the exception vector memory module. In a typical system, select a low-latency memory module for the exception code.

Note: Platform Designer provides an Absolute option, which allows you to specify an absolute address in Fast TLB Miss Exception vector offset. Use an absolute address when the memory storing the exception handler is located outside of the processor system and subsystems of the processor system.

Fast TLB Miss Exception vector offset specifies the location of the exception vector relative to the memory module's base address. Platform Designer calculates the physical address of the exception vector when you modify the memory module, the offset, or the memory module's base address. In Platform Designer, Fast TLB Miss Exception vector displays the readonly, calculated address. The address is always a physical address, even when an MMU is present.

Note: The Nios II MMU is optional and mutually exclusive from the Nios II MPU. Nios II systems can include either an MMU or MPU, but cannot include both an MMU and MPU in the same design.

For information about the Nios II MMU, refer to the *Programming Model* chapter of the Nios II Processor Reference Handbook.

To function correctly with the MMU, 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 parameter editor.

Related Information

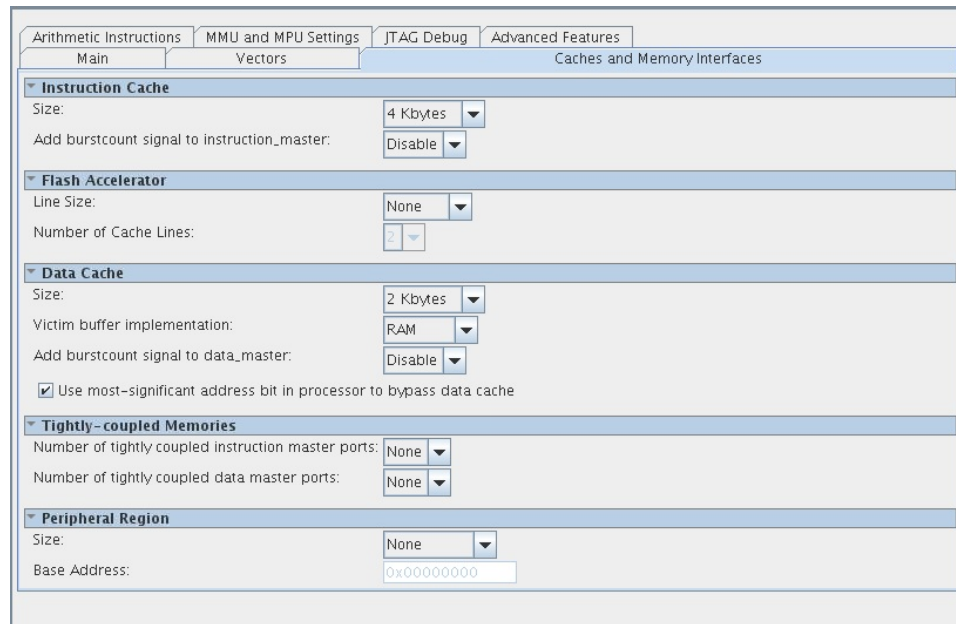
[Programming Model](#) on page 33

4.3. Caches and Memory Interfaces Tab

The **Caches and Memory Interfaces** tab allows you to configure the cache and tightly-coupled memory usage for the instruction and data master ports.



Figure 8. Nios II Platform Designer Caches and Memory Interfaces Tab



4.3.1. Instruction Cache

The Instruction cache parameters provide the following options for the Nios II/f core:

- **Size**—Specifies the size of the instruction cache. Valid sizes are from 512 bytes to 64 KBytes, or None.

Choosing None disables the instruction cache. The Avalon-MM instruction master port from the Nios II processor will still be available. In this case, you must include a tightly-coupled instruction memory.

- **Add burstcount signal to instruction_master** —The Nios II processor can fill its instruction cache lines using burst transfers. Usually you enable bursts on the processor's instruction master when instructions are stored in DRAM, and disable bursts when instructions are stored in SRAM.

Bursting to DRAM typically improves memory bandwidth, but might consume additional FPGA resources. Be aware that when bursts are enabled, accesses to slaves might go through additional hardware (called burst adapters) which might decrease your f_{MAX} .

When the Nios II processor transfers execution to the first word of a cache line, the processor fills the line by executing a sequence of word transfers that have ascending addresses, such as 0, 4, 8, 12, 16, 20, 24, 28.

However, when the Nios II processor transfers execution to an instruction that is not the first word of a cache line, the processor fetches the required (or "critical") instruction first, and then fills the rest of the cache line. The addresses of a burst increase until the last word of the cache line is filled, and then continue with the first word of the cache line. For example, with a 32-byte cache line, transferring control to address 8 results in a burst with the following address sequence: 8, 12, 16, 20, 24, 28, 0, 4.

4.3.2. Flash Accelerator

The Flash accelerator is a small fully-associative cache for real-time applications. Use the Flash accelerator when executing directly from on-chip memories such as flash.

You can choose either 64-bits or 128-bits for your line size option. After choosing a line size you are able to choose either 2 or 4 cache lines for your system. After instantiating the Flash Accelerator the `flash_instruction_master` signal becomes available.

Note: It is recommended to use the Flash Accelerator with MAX 10 on-chip flash.

4.3.3. Data Cache

The **Data Cache** parameters provide the following options for the Nios II/f core:

- **Size**—Specifies the size of the data cache. Valid sizes are from 512 bytes to 64 KBytes, or None. Depending on the value specified for Data cache, the following options are available:
- **Victim buffer implementation**—Specifies whether to use RAM or registers. The data cache victim buffer temporarily holds a dirty cache line while the data is written back to external memory.
- **Add burstcount signal to data_master**—The Nios II processor can fill its data cache lines using burst transfers. Usually you enable bursts on the processor's data bus when processor data is stored in DRAM, and disable bursts when processor data is stored in SRAM.

Bursting to DRAM typically improves memory bandwidth but might consume additional FPGA resources. Be aware that when bursts are enabled, accesses to slaves might go through additional hardware (called burst adapters) which might decrease your fMAX.

The burst length is always 8 for a 32-byte line size. Data cache bursts are always aligned on the cache line boundary. For example, with a 32-byte Nios II data cache line, a cache miss to the address 8 results in a burst with the following address sequence: 0, 4, 8, 12, 16, 20, 24 and 28.

- **Use most-significant address bit in processor to bypass data cache**—This option is enabled by default, the data master interfaces only support up to 31-bit byte address. Disable this option to support full 32-bit byte address

Although the Nios II processor can operate entirely out of tightly-coupled memory without the need for Avalon-MM instruction or data masters, software debug is not possible when either the Avalon-MM instruction or data master is omitted.

Note: By default this feature is turned on for backwards compatibility with the Nios II Classic core.

4.3.4. Tightly-coupled Memories

Number of tightly coupled instruction master port(s) (Include tightly coupled instruction master port(s))—Specifies one to four tightly-coupled instruction master ports for the Nios II processor. In Platform Designer, select the number from the **Number of tightly coupled instruction master port(s)** list. Tightly coupled



memory ports appear on the connection panel of the Nios II processor on the Platform Designer System Contents tab. You must connect each port to exactly one memory component in the system.

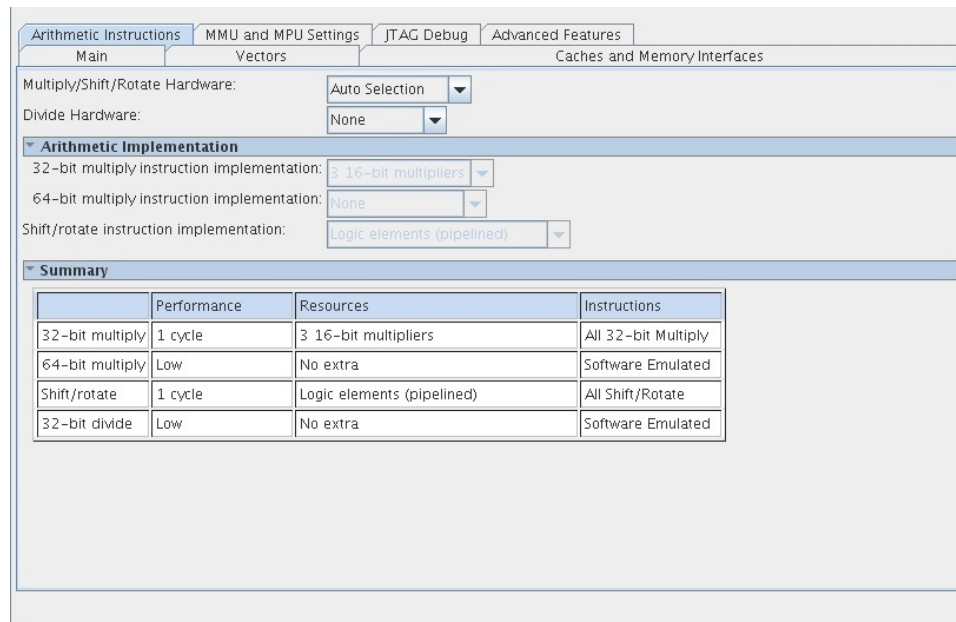
4.3.5. Peripheral Region

The Peripheral Region section in the Caches and Memory Interfaces tab has a maximum size of 2 Gbytes. You can set the base address once the size has been selected. All addresses in the peripheral region produce uncacheable data accesses.

4.4. Arithmetic Instructions Tab

Nios II/f cores offer hardware multiply and divide options. You can choose the best option to balance embedded multiplier usage, logic element (LE) usage, and performance.

Figure 9. Nios II Platform Designer Arithmetic Instructions Tab



4.4.1. Arithmetic Instructions

- **Multiply/Shift/Rotate Hardware:**—You have the option of either choosing auto or manually selected hardware. It is recommended to choose auto, which selects the hardware according to the device family in your current Platform Designer project.
- **Divide Hardware:**—Platform Designer allows you to choose SRT Radix-2 as your divide hardware option if selected.

Table 54. Divide Hardware: SRT Radix-2

Hardware	Performance	Resources	Instructions
32-bit divide	35 cycles	Logic elements	All Divide

4.4.2. Arithmetic Implementation

This section is only available if you choose to manually select your multiply/shift/rotate hardware.

- 32-bit multiply instruction implementation**—You have three options for the 32-bit multiply instruction. Choosing the 1 32-bit multiplier option allocates all resources to the 32-bit multiplier making it the resource for the 64-bit multiply instruction and shift/rotate instruction.

Table 55. 32-bit multiply instruction implementation options

Performance	Resources	Instruction
11 cycles	Logic elements	All 32-bit Multiply
1 cycle	3 16-bit multipliers	All 32-bit Multiply
1 cycle	1 32-bit multipliers	All 32-bit Multiply

- 64-bit multiply instruction implementation**— This option can only be used if the 32-bit multiply instruction selection is set to the 3 16-bit multipliers option. The Nios II only supports up to a 32 x 32 bit multiplication. The 64-bit option is achieved by using the 32-bit multiplier along with the multiply extended instructions (*mulxss*, *mulxsu*, *mulxuu*), which can be found in the *Instruction Set Reference* chapter of this manual.

Table 56. 64-bit multiply instruction implementation options

Performance	Resources	Instruction
2 cycles	1 16-bit multiplier	All 64-bit Multiply

- Shift/rotate instruction implementation**—Platform Designer gives you the option of either choosing non-pipelined or pipelined.

Table 57. Shift/rotate instruction implementation options

Performance	Resources	Instructions
2 - 11 cycles	Logic elements (non-pipelined)	All Shift/Rotate
1 cycle	Logic elements (pipelined)	All Shift/Rotate

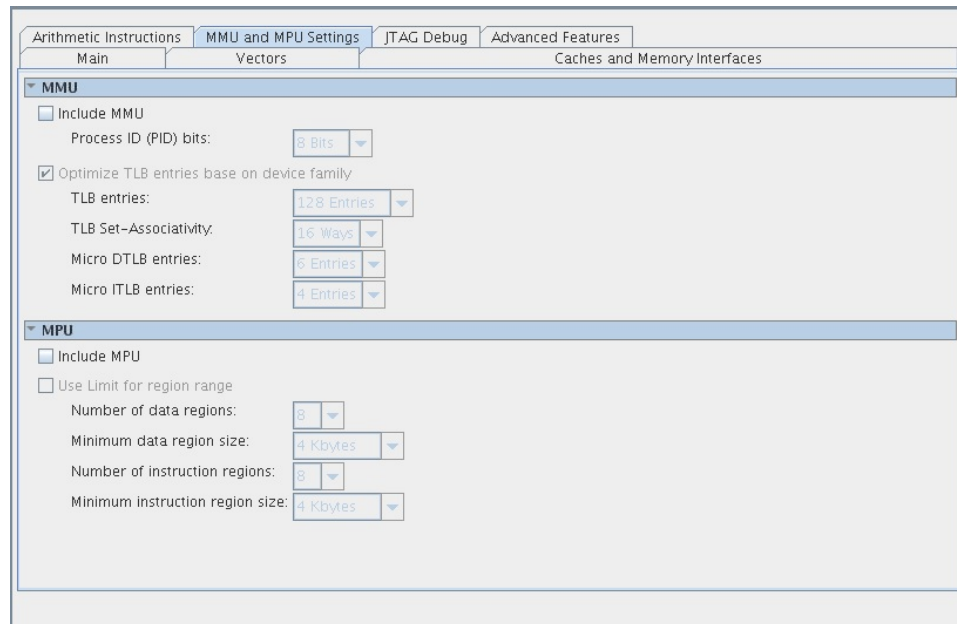
Note: Highly recommend to chose auto selection. Platform Designer makes the selections according to the device family previously selected.

4.5. MMU and MPU Settings Tab

The **MMU and MPU Settings** tab presents settings for configuring the MMU and MPU on the Nios II processor. You can select the features appropriate for your target application.



Figure 10. Nios II Platform Designer MMU and MPU Settings Tab



Note: This tab is only available for the Nios II/f core.

4.5.1. MMU

When **Include MMU** on the **MMU and MPU Settings** tab is on, the MMU settings on the **MMU and MPU Settings** tab provide the following options for the MMU in the Nios II/f core. Typically, you should not need to change any of these settings from their default values.

- **Process ID (PID) bits**—Specifies the number of bits to use to represent the process identifier.
- **Optimize number of TLB entries based on device family**—When on, specifies the optimal number of TLB entries to allocate based on the device family of the target hardware and disables **TLB entries**.
- **TLB entries**—Specifies the number of entries in the translation lookaside buffer (TLB).
- **TLB Set-Associativity**—Specifies the number of set-associativity ways in the TLB.
- **Micro DTLB entries**—Specifies the number of entries in the micro data TLB.
- **Micro ITLB entries**—Specifies the number of entries in the micro instruction TLB.

For information about the MMU, refer to the *Programming Model* chapter of the *Nios II Processor Reference Handbook*.

For specifics on the Nios II/f core, refer to the *Nios II Core Implementation Details* chapter of the *Nios II Processor Reference Handbook*.

Related Information

- [Programming Model](#) on page 33

- [Nios II Core Implementation Details](#) on page 118

4.5.2. MPU

When **Include MPU** on the **MMU and MPU Settings** tab is on, the MPU settings on the **MMU and MPU Settings** tab provide the following options for the MPU in the Nios II/f core.

- **Use Limit for region range**—Controls whether the amount of memory in the region is defined by size or by upper address limit. When on, the amount of memory is based on the given upper address limit. When off, the amount of memory is based on the given size.
- **Number of data regions**—Specifies the number of data regions to allocate. Allowed values range from 2 to 32.
- **Minimum data region size**—Specifies the minimum data region size. Allowed values range from 256 bytes to 1 MB and must be a power of two.
- **Number of instruction regions**—Specifies the number of instruction regions to allocate. Allowed values range from 2 to 32.
- **Minimum instruction region size**—Specifies the minimum instruction region size. Allowed values range from 256 bytes to 1 MB and must be a power of two.

Note:

The maximum region size is the size of the Nios II instruction and data addresses automatically determined when the Nios II system is generated in Platform Designer. Maximum region size is based on the address range of slaves connected to the Nios II instruction and data masters.

For information about the MPU, refer to the *Programming Model* chapter of the *Nios II Processor Reference Handbook*.

For specifics on the Nios II/f core, refer to the *Nios II Core Implementation Details* chapter of the *Nios II Processor Reference Handbook*.

Related Information

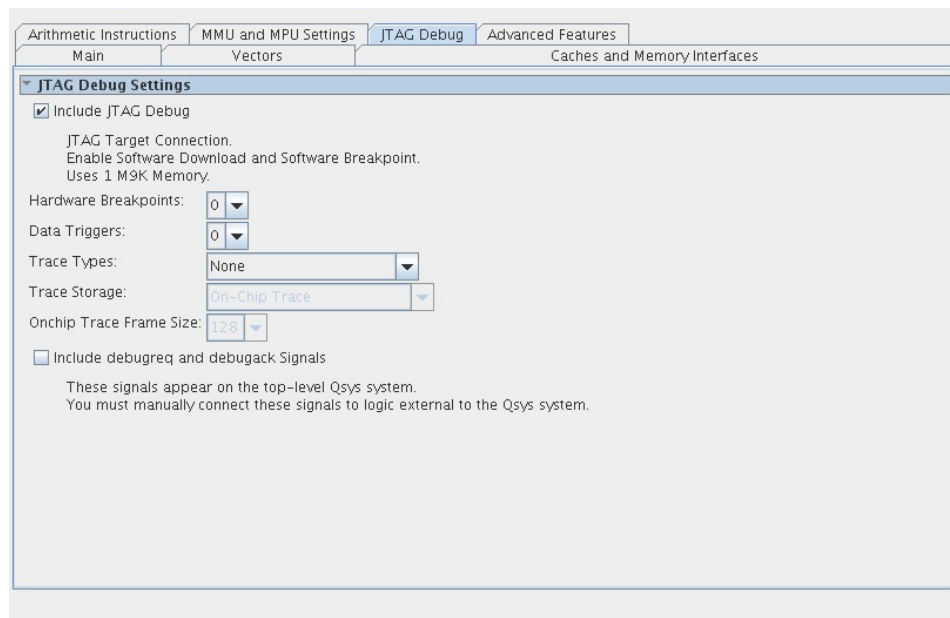
- [Programming Model](#) on page 33
- [Nios II Core Implementation Details](#) on page 118

4.6. JTAG Debug Tab

The **JTAG Debug** tab presents settings for configuring the JTAG debug module on the Nios II processor. You can select the debug features appropriate for your target application.



Figure 11. Nios II Platform Designer JTAB Debug Tab



Soft processor cores such as the Nios II processor offer unique debug capabilities beyond the features of traditional fixed processors. The soft nature of the Nios II processor allows you to debug a system in development using a full-featured debug core, and later remove the debug features to conserve logic resources. For the release version of a product, you might choose to reduce the JTAG debug module functionality, or remove it altogether.

Table 58. Debug Configuration Features

Feature	Description
JTAG Target Connection	Connects to the processor through the standard JTAG pins on the Intel FPGA. This connection provides the basic capabilities to start and stop the processor, and examine/edit registers and memory.
Download Software	Downloads executable code to the processor's memory via the JTAG connection.
Software Breakpoints	Sets a breakpoint on instructions residing in RAM.
Hardware Breakpoints	Sets a breakpoint on instructions residing in nonvolatile memory, such as flash memory.
Data Triggers	Triggers based on address value, data value, or read or write cycle. You can use a trigger to halt the processor on specific events or conditions, or to activate other events, such as starting execution trace, or sending a trigger signal to an external logic analyzer. Two data triggers can be combined to form a trigger that activates on a range of data or addresses.
Instruction Trace	Captures the sequence of instructions executing on the processor in real time.
Data Trace	Captures the addresses and data associated with read and write operations executed by the processor in real time.
On-Chip Trace	Stores trace data in on-chip memory.
Off-Chip Trace	Stores trace data in an external debug probe. Off-chip trace instantiates a PLL inside the Nios II core. Off-chip trace requires a debug probe from Imagination Technologies or Lauterbach GmbH.

The **Include debugreq and debugack signals** debug signals setting provides the following functionality. When on, the Nios II processor includes debug request and acknowledge signals. These signals let another device temporarily suspend the Nios II processor for debug purposes. The signals are exported to the top level of your Platform Designer system.

For more information about the debug signals, refer to the Processor Architecture chapter of the Nios II Processor Reference Handbook.

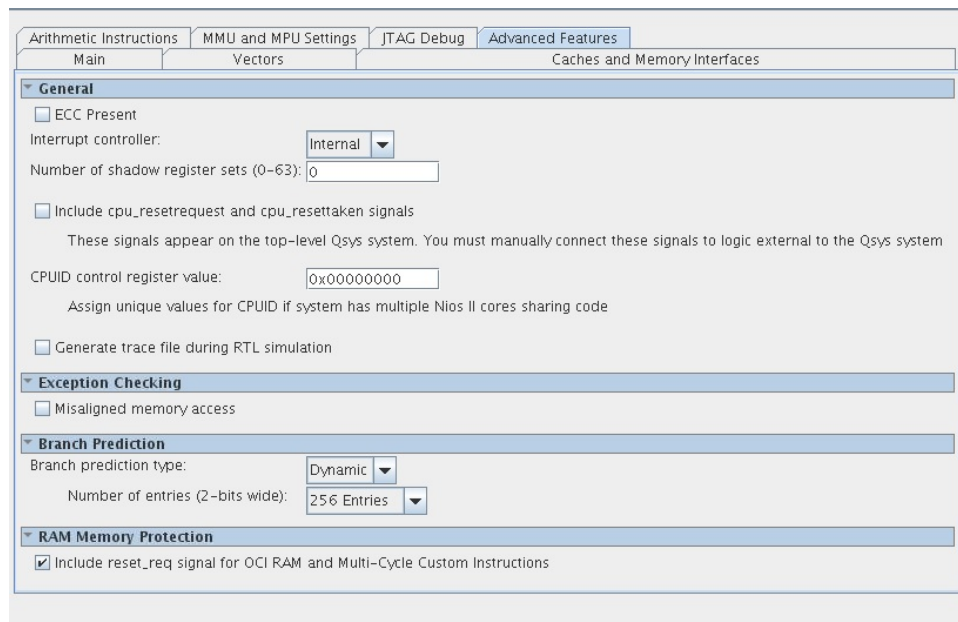
You can set the onchip trace buffer size to sizes from 128 to 64K trace frames, using OCI Onchip Trace. Larger buffer sizes consume more on-chip M4K RAM blocks. Every M4K RAM block can store up to 128 trace frames.

Related Information

[Processor Architecture](#) on page 14

4.7. Advanced Features Tab


Figure 12. Nios II Platform Designer Advanced Features Tab



4.7.1. ECC

ECC is only available for the Nios II/f core and provides ECC support for Nios II internal RAM blocks, such as instruction cache, MMU TLB, and register file. The SECDED ECC algorithm is based on Hamming codes, which detect 1 or 2 bit errors and corrects 1 bit errors. If the Nios II processor does not attempt to correct any errors and only detects them, the ECC algorithm can detect 3 bit errors.

Refer to "ECC" section in the Nios II Core Implementation Details chapter for more information about ECC support in the Nios II/f core.



Related Information

[ECC](#) on page 130

4.7.2. Interrupt Controller Interfaces

The Interrupt controller setting determines which of the following configurations is implemented:

- Internal interrupt controller
- External interrupt controller (EIC) interface

The EIC interface is available only on the Nios II/f core.

Note: When the EIC interface and shadow register sets are implemented on the Nios II core, you must ensure that your software is built with the latest Nios II Embedded Design Suite (EDS) version. Earlier versions have an implementation of the `eret` instruction that is incompatible with shadow register sets.

For details about the EIC interface, refer to “Exception Processing” in the Programming Model chapter of the Nios II Processor Reference Handbook.

Related Information

[Exception Processing](#) on page 71

4.7.3. Shadow Register Sets

The Number of shadow register sets setting determines whether the Nios II core implements shadow register sets. The Nios II core can be configured with up to 63 shadow register sets.

Shadow register sets are available only on the Nios II/f core.

Note: When the EIC interface and shadow register sets are implemented on the Nios II core, you must ensure that your software is built with the Nios II EDS version 9.0 or higher.

For details about shadow register sets, refer to “Registers” in the Programming Model chapter of the Nios II Processor Reference Handbook.

Related Information

[Registers](#) on page 42

4.7.4. Reset Signals

The Include **cpu_resetrequest** and **cpu_resettaken** signals reset signals setting provides the following functionality. When on, the Nios II processor includes processor-only reset request signals. These signals let another device individually reset the Nios II processor without resetting the entire system. The signals are exported to the top level of your system.

Note: You must manually connect these signals to logic external to your Platform Designer system.

For more information on the reset signals, refer to the Processor Architecture chapter of the Nios II Processor Reference Handbook.

Related Information

[Processor Architecture](#) on page 14

4.7.5. CPU ID Control Register Value

In Platform Designer, the automatically-assigned CPUID control register value is always 0x00000000, so Intel recommends always assigning the value manually.

To assign the value yourself, turn on Assign cpuid control register value manually and type a 32-bit value (in hexadecimal or decimal format) in the cpuid control register value box.

4.7.6. Generate Trace File

Through this selection, Platform Designer creates a trace file called "`system_name`"_"`cpu_name`".tr. Use the `nios2-trace` command to display it.

4.7.7. Exception Checking

The Exception Checking settings provide the following options:

Misaligned memory access—Misaligned memory access detection is only available for the Nios II/f core. When Misaligned memory access is on, the processor checks for misaligned memory accesses.

Note:

When your system contains an MMU or MPU, the processor automatically generates misaligned memory access exceptions. Therefore, the **Misaligned memory access** check box is always disabled when **Include MMU** or **Include MPU** on the **Core Nios II** tab are on.

There are two misaligned memory address exceptions:

- **Misaligned data address**—Data addresses of load and store instructions are checked for misalignment. A data address is considered misaligned if the byte address is not a multiple of the data width of the load or store instruction (4 bytes for word, 2 bytes for half-word). Byte load and store instructions are always aligned so never generate a misaligned data address exception.
- **Misaligned destination address**—Destination instruction addresses of `br`, `callr`, `jmp`, `ret`, `eret`, and `bret` instructions are checked for misalignment. A destination instruction address is considered misaligned if the target byte address of the instruction is not a multiple of four.

Your exception handler can use this code to quickly determine the proper action to take, rather than have to determine the cause of an exception through instruction decoding. Additionally, some exceptions also store the instruction or data address associated with the exception in the `badaddr` register.

For further descriptions of exceptions, exception handling, and control registers, refer to the Programming Model chapter of the Nios II Processor Reference Handbook.

Related Information

[Programming Model](#) on page 33



4.7.8. Branch Prediction

This section in the tab allows you to choose either dynamic or static for your branch prediction type. As for the number of entries (2-bits wide) you can choose to have 256, 4096, or up to 8192 entries.

4.7.9. RAM Memory Protection

Through the RAM Memory Protection section you can include the `reset_req` signal for OCI RAM and Multi-Cycle Custom Instructions.

4.8. The Quartus Prime IP File

The Quartus[®] Prime IP file (`.qip`) is a file generated by the MegaWizard[™] Plug-In Manager, that contains information about a generated IP core. You are prompted to add this `.qip` file to the current project at the time of Quartus Prime file generation. In most cases, the `.qip` file contains all of the necessary assignments and information required to process the core or system in the Quartus Prime compiler. Generally, a single `.qip` file is generated for each IP core and for each Platform Designer system. However, some complex components generate a separate `.qip` file, so the system `.qip` file references the component `.qip` file.

4.9. Instantiating the Nios II Processor Revision History

Document Version	Changes
2019.04.30	Maintenance release
2018.04.18	Implemented editorial enhancements.
2016.10.28	Removed extra exception information option.
2015.04.02	Initial release

5. Nios II Core Implementation Details

This document describes all of the Nios II processor core implementations available at the time of publishing. This document describes only implementation-specific features of each processor core. All cores support the Nios II instruction set architecture.

For more information regarding the Nios II instruction set architecture, refer to the *Instruction Set Reference* chapter of the *Nios II Processor Reference Handbook*.

For common core information and details on a specific core, refer to the *Nios II Performance Benchmarks*.

Related Information

- [Instruction Set Reference](#) on page 166
- [Nios II Performance Benchmarks](#)

5.1. Device Family Support


All Nios II cores provide the same support for target Intel FPGA device families.

Table 59. Device Family Support

Device Family	Support
Max 10	Final
Arria® GX	Final
Arria II GX	Final
Arria II GZ	Final
Arria V	Final
Intel Arria 10	Final
Cyclone® II	Final
Cyclone III	Final
Cyclone III LS	Final
Cyclone IV GX	Final
Cyclone IV E	Final
Cyclone V	Final
Intel Cyclone 10 LP	Final
Stratix® II	Final
Stratix II GX	Final
Stratix III	Final
<i>continued...</i>	

Intel Corporation. All rights reserved. Agilix, Altera, Arria, Cyclone, Enpirion, Intel, the Intel logo, MAX, Nios, Quartus and Stratix words and logos are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Intel warrants performance of its FPGA and semiconductor products to current specifications in accordance with Intel's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Intel assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Intel. Intel customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

*Other names and brands may be claimed as the property of others.



Device Family	Support
Stratix IV E	Final
Stratix IV GT	Final
Stratix IV GX	Final
Stratix V	Final
Intel Stratix 10	Final

Preliminary support—The core is verified with preliminary timing models for this device family. The core meets all functional requirements, but might still be undergoing timing analysis for the device family. It can be used in production designs with caution.

Final support—The core is verified with final timing models for this device family. The core meets all functional and timing requirements for the device family and can be used in production designs.

5.2. Nios II/f Core

The Nios II/f fast core is designed for high execution performance. Performance is gained at the expense of core size. Intel FPGA designed the Nios II/f core with the following design goals in mind:

- Maximize the instructions-per-cycle execution efficiency
- Optimize interrupt latency
- Maximize f_{MAX} performance of the processor core

The resulting core is optimal for performance-critical applications, as well as for applications with large amounts of code and data, such as systems running a full-featured operating system.

5.2.1. Overview

The Nios II/f core:

- Has separate optional instruction and data caches
- Provides optional MMU to support operating systems that require an MMU
- Provides optional MPU to support operating systems and runtime environments that desire memory protection but do not need virtual memory management
- Can access up to 4 GB of external address space when bit31 is not enabled
- Supports optional external interrupt controller (EIC) interface to provide customizable interrupt prioritization
- Supports optional shadow register sets to improve interrupt latency
- Supports optional tightly-coupled memory for instructions and data
- Employs a 6-stage pipeline to achieve maximum DMIPS/MHz
- Performs dynamic or static branch prediction
- Provides optional hardware multiply, divide, and shift options to improve arithmetic performance



- Supports the addition of custom instructions
- Optional ECC support for internal RAM blocks (data cache, data cache victim buffer RAM, instruction and data tightly-coupled memories, instruction cache, MMU TLB, and register file)
- Supports the JTAG debug module
- Supports optional JTAG debug module enhancements, including hardware breakpoints and real-time trace

The following sections discuss the noteworthy details of the Nios II/f core implementation. This document does not discuss low-level design issues or implementation details that do not affect Nios II hardware or software designers.

5.2.2. Arithmetic Logic Unit

The Nios II/f core provides several arithmetic logic unit (ALU) options to improve the performance of multiply, divide, and shift operations.

5.2.2.1. Multiply and Divide Performance

The Nios II/f core provides the following hardware multiplier options:

- DSP Block—Includes DSP block multipliers available on the target device. This option is available only on Intel FPGAs that have a hardware multiplier that supports 32-bit multiplication.
- Embedded Multipliers—Includes dedicated embedded multipliers available on the target device. This option is available only on Intel FPGAs that have embedded multipliers.
- Logic Elements—Includes hardware multipliers built from logic element (LE) resources.
- None—Does not include multiply hardware. In this case, multiply operations are emulated in software.

The Nios II/f core also provides a hardware divide option that includes LE-based divide circuitry in the ALU.

Including an ALU option improves the performance of one or more arithmetic instructions.

Note: The performance of the embedded multipliers differ, depending on the target FPGA family.

Table 60. Hardware Multiply and Divide Details for the Nios II/f Core

ALU Option	Hardware Details	Cycles per Instruction	Result Latency Cycles	Supported Instructions
No hardware multiply or divide	Multiply and divide instructions generate an exception	–	–	None
Logic elements	ALU includes 32 x 4-bit multiplier	11	+2	mul, muli
<i>continued...</i>				



ALU Option	Hardware Details	Cycles per Instruction	Result Latency Cycles	Supported Instructions
32-bit multiplier	ALU includes 32 x 32-bit multiplier	1	+2	mul, muli, mulxss, mulxsu, mulxuu
16-bit multiplier	ALU includes 3 16 x 16-bit multiplier	1	+2	mul, muli
16-bit multiplier	ALU includes 4 16 x 16-bit multiplier	2	+2	mul, muli, mulxss, mulxsu, mulxuu
Hardware divide	ALU includes SRT Radix-2 divide circuit	35	+2	div, divu

The cycles per instruction value determines the maximum rate at which the ALU can dispatch instructions and produce each result. The latency value determines when the result becomes available. If there is no data dependency between the results and operands for back-to-back instructions, then the latency does not affect throughput. However, if an instruction depends on the result of an earlier instruction, then the processor stalls through any result latency cycles until the result is ready.

In the following code example, a multiply operation (with 1 instruction cycle and 2 result latency cycles) is followed immediately by an add operation that uses the result of the multiply. On the Nios II/f core, the `addi` instruction, like most ALU instructions, executes in a single cycle. However, in this code example, execution of the `addi` instruction is delayed by two additional cycles until the multiply operation completes.

```
mul r1, r2, r3      ; r1 = r2 * r3
addi r1, r1, 100   ; r1 = r1 + 100 (Depends on result of mul)
```

In contrast, the following code does not stall the processor.

```
mul r1, r2, r3      ; r1 = r2 * r3
or r5, r5, r6       ; No dependency on previous results
or r7, r7, r8       ; No dependency on previous results
addi r1, r1, 100    ; r1 = r1 + 100 (Depends on result of mul)
```

5.2.2.2. Shift and Rotate Performance

The performance of shift operations depends on the hardware multiply option. When a hardware multiplier is present, the ALU achieves shift and rotate operations in three or four clock cycles. Otherwise, the ALU includes dedicated shift circuitry that achieves one-bit-per-cycle shift and rotate performance.

Related Information

[Instruction Performance](#) on page 127

5.2.3. Memory Access

The Nios II/f core provides optional instruction and data caches. The cache size for each is user-definable, between 512 bytes and 64 KB.



The memory address width in the Nios II/f core depends on whether the optional MMU is present. Without an MMU, the Nios II/f core supports the bit-31 cache bypass method for accessing I/O on the data master port. Therefore addresses are 31 bits wide, reserving bit 31 for the cache bypass function. With an MMU, cache bypass is a function of the memory partition and the contents of the translation lookaside buffer (TLB). Therefore bit-31 cache bypass is disabled, and 32 address bits are available to address memory.

5.2.3.1. Instruction and Data Master Ports

The instruction master port is a pipelined Avalon Memory-Mapped (Avalon-MM) master port. The core also includes a data cache with a fixed 32-byte line size, making the data master port a pipelined Avalon-MM master port.

The instruction and data master ports on the Nios II/f core are optional. A master port can be excluded, as long as the core includes at least one tightly-coupled memory to take the place of the missing master port.

Note: Although the Nios II processor can operate entirely out of tightly-coupled memory without the need for Avalon-MM instruction or data masters, software debug is not possible when either the Avalon-MM instruction or data master is omitted.

Support for pipelined Avalon-MM transfers minimizes the impact of synchronous memory with pipeline latency. The pipelined instruction and data master ports can issue successive read requests before prior requests complete.

5.2.3.2. Instruction and Data Caches

This section first describes the similar characteristics of the instruction and data cache memories, and then describes the differences.

Both the instruction and data cache addresses are divided into fields based on whether or not an MMU is present in your system.

Table 61. Cache Byte Address Fields

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
tag													line		
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
line											offset				

Table 62. Cache Virtual Byte Address Fields

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



Table 63. Cache Physical Byte Address Fields

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

5.2.3.2.1. Instruction Cache

The instruction cache memory has the following characteristics:

- Direct-mapped cache implementation.
- 32 bytes (8 words) per cache line.
- The instruction master port reads an entire cache line at a time from memory, and issues one read per clock cycle.
- Critical word first.
- Virtually-indexed, physically-tagged, when MMU present.

The size of the tag field depends on the size of the cache memory and the physical address size. The size of the line field depends only on the size of the cache memory. The offset field is always five bits (i.e., a 32-byte line). The maximum instruction byte address size is 31 bits in systems without an MMU present. In systems with an MMU, the maximum instruction byte address size is 32 bits and the tag field always includes all the bits of the physical frame number (PFN).

The instruction cache is optional. However, excluding instruction cache from the Nios II/f core requires that the core include at least one tightly-coupled instruction memory.

5.2.3.2.2. Data Cache

- Direct-mapped cache implementation
- Line size of 32-bytes
- The data master port reads an entire cache line at a time from memory, and issues one read per clock cycle.
- Write-back
- Write-allocate (i.e., on a store instruction, a cache miss allocates the line for that address)
- Virtually-indexed, physically-tagged, when MMU present

The size of the tag field depends on the size of the cache memory and the physical address size. The size of the line field depends only on the size of the cache memory. The size of the offset field depends on the line size. Line sizes of 32 bytes have offset widths of 5-bits. The maximum data byte address size is 31 bits in systems without an MMU present. In systems with an MMU, the maximum data byte address size is 32 bits and the tag field always includes all the bits of the PFN.

The data cache is optional. If the data cache is excluded from the core, the data master port can also be excluded.

The Nios II instruction set provides several different instructions to clear the data cache. There are two important questions to answer when determining the instruction to use. Do you need to consider the tag field when looking for a cache match? Do you need to write dirty cache lines back to memory before clearing? Below the table lists the most appropriate instruction to use for each case.

Table 64. Data Cache Clearing Instructions

Instruction	Ignore Tag Field	Consider Tag Field
Write Dirty Lines	flushd	flushda
Do Not Write Dirty Lines	initd	initda

For more information regarding the Nios II instruction set, refer to the *Instruction Set Reference* chapter of the *Nios II Processor Reference Handbook*.

The Nios II/f core implements all the data cache bypass methods.

For information regarding the data cache bypass methods, refer to the *Processor Architecture* chapter of the *Nios II Processor Reference Handbook*

Mixing cached and uncached accesses to the same cache line can result in invalid data reads. For example, the following sequence of events causes cache incoherency.

1. The Nios II core writes data to cache, creating a dirty data cache line.
2. The Nios II core reads data from the same address, but bypasses the cache.

Note:

Avoid mixing cached and uncached accesses to the same cache line, regardless whether you are reading from or writing to the cache line. If it is necessary to mix cached and uncached data accesses, flush the corresponding line of the data cache after completing the cached accesses and before performing the uncached accesses.

Related Information

- [Instruction Set Reference](#) on page 166
- [Processor Architecture](#) on page 14

5.2.3.2.3. Bursting

When the data cache is enabled, you can enable bursting on the data master port. Consult the documentation for memory devices connected to the data master port to determine whether bursting can improve performance.

5.2.4. Tightly-Coupled Memory

The Nios II/f core provides optional tightly-coupled memory interfaces for both instructions and data. A Nios II/f core can use up to four each of instruction and data tightly-coupled memories. When a tightly-coupled memory interface is enabled, the Nios II core includes an additional memory interface master port. Each tightly-coupled memory interface must connect directly to exactly one memory slave port.

When tightly-coupled memory is present, the Nios II core decodes addresses internally to determine if requested instructions or data reside in tightly-coupled memory. If the address resides in tightly-coupled memory, the Nios II core fetches the



instruction or data through the tightly-coupled memory interface. Software accesses tightly-coupled memory with the usual load and store instructions, such as `ldw` or `ldwio`.

Accessing tightly-coupled memory bypasses cache memory. The processor core functions as if cache were not present for the address span of the tightly-coupled memory. Instructions for managing cache, such as `initd` and `flushd`, do not affect the tightly-coupled memory, even if the instruction specifies an address in tightly-coupled memory.

When the MMU is present, tightly-coupled memories are always mapped into the kernel partition and can only be accessed in supervisor mode.

5.2.5. Memory Management Unit

The Nios II/f core provides options to improve the performance of the Nios II MMU.

For information about the MMU architecture, refer to the *Programming Model* chapter of the *Nios II Processor Reference Handbook*.

Related Information

[Programming Model](#) on page 33

5.2.5.1. Micro Translation Lookaside Buffers

The translation lookaside buffer (TLB) consists of one main TLB stored in on-chip RAM and two separate micro TLBs (μ TLB) for instructions (μ ITLB) and data (μ DTLB) stored in LE-based registers.

The TLBs have a configurable number of entries and are fully associative. The default configuration has 6 μ DTLB entries and 4 μ ITLB entries. The hardware chooses the least-recently used μ TLB entry when loading a new entry.

The μ TLBs are not visible to software. They act as an inclusive cache of the main TLB. The processor firsts look for a hit in the μ TLB. If it misses, it then looks for a hit in the main TLB. If the main TLB misses, the processor takes an exception. If the main TLB hits, the TLB entry is copied into the μ TLB for future accesses.

The hardware automatically flushes the μ TLB on each TLB write operation and on a `wrc1` to the `tlbmisc` register in case the process identifier (PID) has changed.

5.2.6. Memory Protection Unit

The Nios II/f core provides options to improve the performance of the Nios II MPU.

For information about the MPU architecture, refer to the *Programming Model* chapter of the *Nios II Processor Reference Handbook*.

Related Information

[Programming Model](#) on page 33

5.2.7. Execution Pipeline

This section provides an overview of the pipeline behavior for the benefit of performance-critical applications. Designers can use this information to minimize unnecessary processor stalling. Most application programmers never need to analyze the performance of individual instructions.

The Nios II/f core employs a 6-stage pipeline.

Table 65. Implementation Pipeline Stages for Nios II/f Core

Stage Letter	Stage Name
F	Fetch
D	Decode
E	Execute
M	Memory
A	Align
W	Writeback

Up to one instruction is dispatched or retired per cycle. Instructions are dispatched and retired in order. Dynamic branch prediction is implemented using a 2-bit branch history table. The pipeline stalls for the following conditions:

- Multi-cycle instructions
- Avalon-MM instruction master port read accesses
- Avalon-MM data master port read/write accesses
- Data dependencies on long latency instructions (for example: load, multiply, shift).

5.2.7.1. Pipeline Stalls

The pipeline is set up so that if a stage stalls, no new values enter that stage or any earlier stages. No “catching up” of pipeline stages is allowed, even if a pipeline stage is empty.

Only the A-stage and D-stage are allowed to create stalls.

The A-stage stall occurs if any of the following conditions occurs:

- An A-stage memory instruction is waiting for Avalon-MM data master requests to complete. Typically this happens when a load or store misses in the data cache, or a `flushd` instruction needs to write back a dirty line.
- An A-stage shift/rotate instruction is still performing its operation. This only occurs with the multicycle shift circuitry (i.e., when the hardware multiplier is not available).
- An A-stage divide instruction is still performing its operation. This only occurs when the optional divide circuitry is available.
- An A-stage multicycle custom instruction is asserting its stall signal. This only occurs if the design includes multicycle custom instructions.



The D-stage stall occurs if an instruction is trying to use the result of a late result instruction too early and no M-stage pipeline flush is active. The late result instructions are loads, shifts, rotates, `rdctl`, multiplies (if hardware multiply is supported), divides (if hardware divide is supported), and multicycle custom instructions (if present).

5.2.7.2. Branch Prediction

The Nios II/f core performs dynamic and static branch predictions to minimize the cycle penalty associated with taken branches.

5.2.8. Instruction Performance

All instructions take one or more cycles to execute. Some instructions have other penalties associated with their execution. Late result instructions have two cycles placed between them and an instruction that uses their result. Instructions that flush the pipeline cause up to three instructions after them to be cancelled. This creates a three-cycle penalty and an execution time of four cycles. Instructions that require Avalon-MM transfers are stalled until any required Avalon-MM transfers (up to one write and one read) are completed.

Table 66. Instruction Execution Performance for Nios II/f Core 4byte/line data cache

Instruction	Cycles	Penalties
Normal ALU instructions (e.g., <code>add</code> , <code>cmplt</code>)	1	
Combinatorial custom instructions	1	
Multicycle custom instructions	> 1	Late result
Branch (correctly predicted, taken)	2	
Branch (correctly predicted, not taken)	1	
Branch (mispredicted)	4	Pipeline flush
<code>trap</code> , <code>break</code> , <code>eret</code> , <code>bret</code> , <code>flushp</code> , <code>wrctl</code> , <code>wrprs</code> ; illegal and unimplemented instructions	4 or 5	Pipeline flush
<code>call</code> , <code>jmp</code> , <code>rdprs</code>	2	
<code>jmp</code> , <code>ret</code> , <code>callr</code>	3	
<code>rdctl</code>	1	Late result
load (without Avalon-MM transfer)	1	Late result
load (with Avalon-MM transfer)	> 1	Late result
store (without Avalon-MM transfer)	1	
store (with Avalon-MM transfer)	> 1	
<code>flushd</code> , <code>flushda</code> (without Avalon-MM transfer)	2	
<code>flushd</code> , <code>flushda</code> (with Avalon-MM transfer)	> 2	
<code>initd</code> , <code>initda</code>	2	
<code>flushi</code> , <code>initi</code>	4	
Multiply		Late result

continued...



Instruction	Cycles	Penalties
Divide		Late result
Shift/rotate (with hardware multiply using embedded multipliers)	1	Late result
Shift/rotate (with hardware multiply using LE-based multipliers)	2	Late result
Shift/rotate (without hardware multiply present)	1 to 32	Late result
All other instructions	1	

For Multiply and Divide, the number of cycles depends on the hardware multiply or divide option. Refer to "Arithmetic Logic Unit" and "Instruction and Data Caches" s for details.

In the default Nios II/f configuration, instructions `trap`, `break`, `eret`, `bret`, `flushp`, `wrcctl`, `wrprs` require four clock cycles. If any of the following options are present, they require five clock cycles:

- MMU
- MPU
- Division exception
- Misaligned load/store address exception
- EIC port
- Shadow register sets


Related Information

- [Data Cache](#) on page 123
- [Instruction and Data Caches](#) on page 122
- [Arithmetic Logic Unit](#) on page 120

5.2.9. Exception Handling

The Nios II/f core supports the following exception types:

- Hardware interrupts
- Software trap
- Illegal instruction
- Unimplemented instruction
- Supervisor-only instruction (MMU or MPU only)
- Supervisor-only instruction address (MMU or MPU only)
- Supervisor-only data address (MMU or MPU only)
- Misaligned data address
- Misaligned destination address
- Division error
- Error-correcting code (ECC)
- Fast translation lookaside buffer (TLB) miss (MMU only)



- Double TLB miss (MMU only)
- TLB permission violation (MMU only)
- MPU region violation (MPU only)

5.2.9.1. External Interrupt Controller Interface

The EIC interface enables you to speed up interrupt handling in a complex system by adding a custom interrupt controller.

The EIC interface is an Avalon-ST sink with the following input signals:

- `eic_port_valid`
- `eic_port_data`

Signals are rising-edge triggered, and synchronized with the Nios II clock input.

The EIC interface presents the following signals to the Nios II processor through the `eic_port_data` signal:

- Requested handler address (RHA)—The 32-bit address of the interrupt handler associated with the requested interrupt.
- Requested register set (RRS)—The six-bit number of the register set associated with the requested interrupt.
- Requested interrupt level (RIL)—The six-bit interrupt level. If RIL is 0, no interrupt is requested.
- Requested nonmaskable interrupt (RNMI) flag—A one-bit flag indicating whether the interrupt is to be treated as nonmaskable.

Table 67. `eic_port_data` Signal

Bit Fields														
44	...													
RHA														
...	13	12	11	10	9	8	7	6	5	4	3	2	1	0
RHA		RRS						RNMI	RIL					

Following Avalon-ST protocol requirements, the EIC interface samples `eic_port_data` only when `eic_port_valid` is asserted (high). When `eic_port_valid` is not asserted, the processor latches the previous values of RHA, RRS, RIL and RNMI. To present new values on `eic_port_data`, the EIC must transmit a new packet, asserting `eic_port_valid`. An EIC can transmit a new packet once per clock cycle.

For an example of an EIC implementation, refer to the *Vectored Interrupt Controller* chapter in the Embedded Peripherals IP User Guide.

Related Information

[Embedded Peripherals IP User Guide](#)



5.2.10. ECC

The Nios II/f core has the option to add ECC support for the following Nios II internal RAM blocks.


- Instruction cache
 - ECC errors (1, 2, or 3 bits) that occur in the instruction cache are recoverable; the Nios II processor flushes the cache line and reads from external memory instead of correcting the ECC error.
- Register file
 - 1 bit ECC errors are recoverable
 - 2 bit ECC errors are not recoverable and generate ECC exceptions
- MMU TLB
 - 1 bit ECC errors triggered by hardware reads are recoverable
 - 2 bit ECC errors triggered by hardware reads are not recoverable and generate ECC exception.
 - 1 or 2 bit ECC errors triggered by software reads to the `TLBMISC` register do not trigger an exception, instead, `TLBMISC.EE` is set to 1. Software must read this field and invalidate/overwrite the TLB entry.
- Data Cache
 - tag RAM—The `ECCINJ.DCTAG` field is used to inject ECC errors into the tag RAM.
 - data RAM—The `ECCINJ.DCDAT` field is used to inject ECC errors into the data RAM
- Tightly-Coupled Memories (TCMs)—Nios II includes the ECC encoder/decoder logic for each TCM and the TCM master port data width is increased to allow the Nios II to read and write the ECC parity bits. The TCM must be a RAM and must store the ECC parity bits along with the data bits.
 - Instruction Tightly-Coupled Memories (ITCM)—Nios II supports up to 4 ITCMs
 - Data Tightly-Coupled Memories (DTCM)—Nios II supports up to 4 DTCMs

The ECC interface is an Avalon-ST source with the output signal `ecc_event_bus`. This interface allows external logic to monitor ECC errors in the Nios II processor.

Table 68. ECC Error Signals

Bit	Field	Description	Effect on Software	Available
0	EEH	ECC error exception while in exception handler mode (i.e., <code>STATUS.EH = 1</code>).	Likely fatal	Always
1	RF_RE	Recoverable (1 bit) ECC error in register file RAM	None	Always
2	RF_UE	Unrecoverable (2 bit) ECC error in register file RAM	Likely fatal	Always
3	ICTAG_RE	Recoverable (1, 2, or 3 bit) ECC error in instruction cache tag RAM	None	Instruction cache present
4	ICDAT_RE	Recoverable (1, 2, or 3 bit) ECC error in instruction cache data RAM.	None	Instruction cache present

continued...



Bit	Field	Description	Effect on Software	Available
5	ITCM0_RE	Recoverable (1-bit) ECC error in ITCM0	None	ITCM0 present
6	ITCM0_UE	Unrecoverable (2-bit) ECC error in ITCM0	Possibly fatal	ITCM0 present
7	ITCM1_RE	Recoverable (1-bit) ECC error in ITCM1	None	ITCM1 present
8	ITCM1_UE	Unrecoverable (2-bit) ECC error in ITCM1	Likely fatal	ITCM1 present
9	ITCM2_RE	Recoverable (1-bit) ECC error in ITCM2	None	ITCM2 present
10	ITCM2_UE	Unrecoverable (2-bit) ECC error in ITCM2	Likely fatal	ITCM2 present
11	ITCM3_RE	Recoverable (1-bit) ECC error in ITCM3	None	ITCM3 present
12	ITCM3_UE	Unrecoverable (2-bit) ECC error in ITCM3	Likely fatal	ITCM3 present
13	DCTAG_RE	Recoverable (1-bit) ECC error in data cache tag RAM	None	Data cache present
14	DCTAG_UE	Unrecoverable (2-bit) ECC error in data cache tag RAM	Likely fatal	Data cache present
15	DCDAT_RE	Recoverable (1-bit with dirty line, 2-bit or 3-bit with clean line) ECC error in data cache data RAM. Excludes recoverable errors found during writeback of a dirty line.	None	Data cache present
16	DCDAT_UE	Unrecoverable (2-bit with dirty line) ECC error in data cache data RAM. Excludes unrecoverable errors found during writeback of a dirty line.	Likely fatal	Data cache present
17	DCWB_RE	Recoverable (1-bit) ECC error in data cache data RAM or victim line buffer RAM during writeback of a dirty line.	None	Data cache present
18	DCWB_UE	Unrecoverable (2-bit) ECC error in data cache data RAM or victim line buffer RAM during writeback of a dirty line.	Likely fatal	Data cache present
19	TLB_RE	Recoverable (1 bit) ECC error in TLB RAM (hardware read of TLB)	None	MMU present
20	TLB_UE	Unrecoverable (2 bit) ECC error in TLB RAM (hardware read of TLB)	Possibly fatal	MMU present
21	TLB_SW	Software-triggered (1, 2, or 3 bit) ECC error in software read of TLB	Possibly fatal	MMU present
22	DTCM0_RE	Recoverable (1-bit) ECC error in DTCM0	None	DTCM0 present
23	DTCM0_UE	Unrecoverable (2-bit) ECC error in DTCM0	Likely fatal	DTCM0 present
24	DTCM1_RE	Recoverable (1-bit) ECC error in DTCM1	None	DTCM1 present
25	DTCM1_UE	Unrecoverable (2-bit) ECC error in DTCM1	Likely fatal	DTCM1 present
26	DTCM2_RE	Recoverable (1-bit) ECC error in DTCM2	None	DTCM2 present
27	DTCM2_UE	Unrecoverable (2-bit) ECC error in DTCM2	Likely fatal	DTCM2 present
28	DTCM3_RE	Recoverable (1-bit) ECC error in DTCM3	None	DTCM3 present
29	DTCM3_UE	Unrecoverable (2-bit) ECC error in DTCM3	Likely fatal	DTCM3 present

5.2.11. JTAG Debug Module

The Nios II/f core supports the JTAG debug module to provide a JTAG interface to software debugging tools. The Nios II/f core supports an optional enhanced interface that allows real-time trace data to be routed out of the processor and stored in an external debug probe.

Note: The Nios II MMU does not support the JTAG debug module trace.

5.3. Nios II/s Core

The Nios II/s standard core is designed for small core size. On-chip logic and memory resources are conserved at the expense of execution performance. The Nios II/s core uses approximately 20% less logic than the Nios II/f core, but execution performance also drops by roughly 40%. Intel FPGA designed the Nios II/s core with the following design goals in mind:

- Do not cripple performance for the sake of size.
- Remove hardware features that have the highest ratio of resource usage to performance impact.

The resulting core is optimal for cost-sensitive, medium-performance applications. This includes applications with large amounts of code as well as data, such as systems running an operating system in which performance is not the highest priority.

5.3.1. Overview


The Nios II/s core:

- Has an instruction cache, but no data cache
- Can access up to 2 GB of external address space
- Supports optional tightly-coupled memory for instructions
- Employs a 5-stage pipeline
- Performs static branch prediction
- Provides hardware multiply, divide, and shift options to improve arithmetic performance
- Supports the addition of custom instructions
- Supports the JTAG debug module
- Supports optional JTAG debug module enhancements, including hardware breakpoints and real-time trace

The following sections discuss the noteworthy details of the Nios II/s core implementation. This document does not discuss low-level design issues or implementation details that do not affect Nios II hardware or software designers.

5.3.2. Arithmetic Logic Unit

The Nios II/s core provides several ALU options to improve the performance of multiply, divide, and shift operations.



5.3.2.1. Multiply and Divide Performance

The Nios II/s core provides the following hardware multiplier options:

- **DSP Block**—Includes DSP block multipliers available on the target device. This option is available only on Intel FPGAs that have DSP Blocks.
- **Embedded Multipliers**—Includes dedicated embedded multipliers available on the target device. This option is available only on Intel FPGAs that have embedded multipliers.
- **Logic Elements**—Includes hardware multipliers built from logic element (LE) resources.
- **None**—Does not include multiply hardware. In this case, multiply operations are emulated in software.

The Nios II/s core also provides a hardware divide option that includes LE-based divide circuitry in the ALU.

Including an ALU option improves the performance of one or more arithmetic instructions.

Note: The performance of the embedded multipliers differ, depending on the target FPGA family.

Table 69. Hardware Multiply and Divide Details for the Nios II/s Core

ALU Option	Hardware Details	Cycles per instruction	Supported Instructions
No hardware multiply or divide	Multiply and divide instructions generate an exception	–	None
LE-based multiplier	ALU includes 32 x 4-bit multiplier	11	mul, muli
Embedded multiplier on Stratix III families	ALU includes 32 x 32-bit multiplier	3	mul, muli, mulxss, mulxsu, mulxuu
Embedded multiplier on Cyclone III families	ALU includes 32 x 16-bit multiplier	5	mul, muli
Hardware divide	ALU includes multicycle divide circuit	4 – 66	div, divu

5.3.2.2. Shift and Rotate Performance

The performance of shift operations depends on the hardware multiply option. When a hardware multiplier is present, the ALU achieves shift and rotate operations in three or four clock cycles. Otherwise, the ALU includes dedicated shift circuitry that achieves one-bit-per-cycle shift and rotate performance.

Refer to the "Instruction Execution Performance for Nios II/s Core" table in the "Instruction Performance" section for details.

Related Information

[Instruction Performance](#) on page 136



5.3.3. Memory Access

The Nios II/s core provides instruction cache, but no data cache. The instruction cache size is user-definable, between 512 bytes and 64 KB. The Nios II/s core can address up to 2 GB of external memory. The Nios II architecture reserves the most-significant bit of data addresses for the bit-31 cache bypass method. In the Nios II/s core, bit 31 is always zero.

For information regarding data cache bypass methods, refer to the *Processor Architecture* chapter of the *Nios II Processor Reference Handbook*.

Related Information

[Processor Architecture](#) on page 14

5.3.3.1. Instruction and Data Master Ports

The instruction master port is a pipelined Avalon Memory-Mapped (Avalon-MM) master port. The core also includes a data cache with a fixed 32-byte line size, making the data master port a pipelined Avalon-MM master port.

The instruction and data master ports on the Nios II/f core are optional. A master port can be excluded, as long as the core includes at least one tightly-coupled memory to take the place of the missing master port.

Note: Although the Nios II processor can operate entirely out of tightly-coupled memory without the need for Avalon-MM instruction or data masters, software debug is not possible when either the Avalon-MM instruction or data master is omitted.

Support for pipelined Avalon-MM transfers minimizes the impact of synchronous memory with pipeline latency. The pipelined instruction and data master ports can issue successive read requests before prior requests complete.


5.3.3.2. Instruction Cache

The instruction cache for the Nios II/s core is nearly identical to the instruction cache in the Nios II/f core. The instruction cache memory has the following characteristics:

- Direct-mapped cache implementation
- The instruction master port reads an entire cache line at a time from memory, and issues one read per clock cycle.
- Critical word first

Table 70. Instruction Byte Address Fields

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
tag													line		
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
line											offset				



The size of the tag field depends on the size of the cache memory and the physical address size. The size of the line field depends only on the size of the cache memory. The offset field is always five bits (i.e., a 32-byte line). The maximum instruction byte address size is 31 bits.

The instruction cache is optional. However, excluding instruction cache from the Nios II/s core requires that the core include at least one tightly-coupled instruction memory.

5.3.4. Tightly-Coupled Memory

The Nios II/s core provides optional tightly-coupled memory interfaces for instructions. A Nios II/s core can use up to four tightly-coupled instruction memories. When a tightly-coupled memory interface is enabled, the Nios II core includes an additional memory interface master port. Each tightly-coupled memory interface must connect directly to exactly one memory slave port.

When tightly-coupled memory is present, the Nios II core decodes addresses internally to determine if requested instructions reside in tightly-coupled memory. If the address resides in tightly-coupled memory, the Nios II core fetches the instruction through the tightly-coupled memory interface. Software does not require awareness of whether code resides in tightly-coupled memory or not.

Accessing tightly-coupled memory bypasses cache memory. The processor core functions as if cache were not present for the address span of the tightly-coupled memory. Instructions for managing cache, such as `init_i` and `flush_i`, do not affect the tightly-coupled memory, even if the instruction specifies an address in tightly-coupled memory.

5.3.5. Execution Pipeline

This section provides an overview of the pipeline behavior for the benefit of performance-critical applications. Designers can use this information to minimize unnecessary processor stalling. Most application programmers never need to analyze the performance of individual instructions.

The Nios II/s core employs a 5-stage pipeline.

Table 71. Implementation Pipeline Stages for Nios II/s Core

Stage Letter	Stage Name
F	Fetch
D	Decode
E	Execute
M	Memory
W	Writeback



Up to one instruction is dispatched or retired per cycle. Instructions are dispatched and retired in-order. Static branch prediction is implemented using the branch offset direction; a negative offset (backward branch) is predicted as taken, and a positive offset (forward branch) is predicted as not taken. The pipeline stalls for the following conditions:

- Multi-cycle instructions (e.g., shift/rotate without hardware multiply)
- Avalon-MM instruction master port read accesses
- Avalon-MM data master port read/write accesses
- Data dependencies on long latency instructions (for example: load, multiply, shift operations)

5.3.5.1. Pipeline Stalls

The pipeline is set up so that if a stage stalls, no new values enter that stage or any earlier stages. No “catching up” of pipeline stages is allowed, even if a pipeline stage is empty.

Only the M-stage is allowed to create stalls.

The M-stage stall occurs if any of the following conditions occurs:

- An M-stage load/store instruction is waiting for Avalon-MM data master transfer to complete.
- An M-stage shift/rotate instruction is still performing its operation when using the multicycle shift circuitry (i.e., when the hardware multiplier is not available).
- An M-stage shift/rotate/multiply instruction is still performing its operation when using the hardware multiplier (which takes three cycles).
- An M-stage multicycle custom instruction is asserting its stall signal. This only occurs if the design includes multicycle custom instructions.

5.3.5.2. Branch Prediction

The Nios II/s core performs static branch prediction to minimize the cycle penalty associated with taken branches.

5.3.6. Instruction Performance

All instructions take one or more cycles to execute. Some instructions have other penalties associated with their execution. Instructions that flush the pipeline cause up to three instructions after them to be cancelled. This creates a three-cycle penalty and an execution time of four cycles. Instructions that require an Avalon-MM transfer are stalled until the transfer completes.

Table 72. Instruction Execution Performance for Nios II/s Core

Instruction	Cycles	Penalties
Normal ALU instructions (e.g., add, cmp1t)	1	
Combinatorial custom instructions	1	
Multicycle custom instructions	> 1	
<i>continued...</i>		



Instruction	Cycles	Penalties
Branch (correctly predicted taken)	2	
Branch (correctly predicted not taken)	1	
Branch (mispredicted)	4	Pipeline flush
trap, break, eret, bret, flushp, wrctl, unimplemented	4	Pipeline flush
jmp, jmp_i, ret, call, callr	4	Pipeline flush
rdctl	1	
load, store	> 1	
flushi, initi	4	
Multiply		
Divide		
Shift/rotate (with hardware multiply using embedded multipliers)	3	
Shift/rotate (with hardware multiply using LE-based multipliers)	4	
Shift/rotate (without hardware multiply present)	1 to 32	
All other instructions	1	

5.3.7. Exception Handling

The Nios II/s core supports the following exception types:

- Internal hardware interrupt
- Software trap
- Illegal instruction
- Unimplemented instruction

5.3.8. JTAG Debug Module

The Nios II/s core supports the JTAG debug module to provide a JTAG interface to software debugging tools. The Nios II/s core supports an optional enhanced interface that allows real-time trace data to be routed out of the processor and stored in an external debug probe.

5.4. Nios II/e Core

The Nios II/e economy core is designed to achieve the smallest possible core size. Intel FPGA designed the Nios II/e core with a singular design goal: reduce resource utilization any way possible, while still maintaining compatibility with the Nios II instruction set architecture. Hardware resources are conserved at the expense of execution performance.

The resulting core is optimal for cost-sensitive applications as well as applications that require simple control logic.

5.4.1. Overview

The Nios II/e core:

- Executes at most one instruction per six clock cycles
- Full 32-bit addressing
- Can access up to 4 GB of external address space
- Supports the addition of custom instructions
- Supports the JTAG debug module
- Does not provide hardware support for potential unimplemented instructions
- Has no instruction cache or data cache
- Does not perform branch prediction

The following sections discuss the noteworthy details of the Nios II/e core implementation. This document does not discuss low-level design issues, or implementation details that do not affect Nios II hardware or software designers.

5.4.2. Arithmetic Logic Unit

The Nios II/e core does not provide hardware support for any of the potential unimplemented instructions. All unimplemented instructions are emulated in software.

The Nios II/e core employs dedicated shift circuitry to perform shift and rotate operations. The dedicated shift circuitry achieves one-bit-per-cycle shift and rotate operations.

5.4.3. Memory Access

The Nios II/e core does not provide instruction cache or data cache. All memory and peripheral accesses generate an Avalon-MM transfer. The Nios II/e core can address up to 4 GB of external memory, full 32-bit addressing.

For information regarding data cache bypass methods, refer to the Processor Architecture chapter of the *Nios II Processor Reference Handbook*.

The Nios II/e core does not provide instruction cache or data cache. All memory and peripheral accesses generate an Avalon-MM transfer.

For information regarding data cache bypass methods, refer to the Processor Architecture chapter of the *Nios II Processor Reference Handbook*.

Related Information

[Processor Architecture](#) on page 14

5.4.4. Instruction Execution Stages

This section provides an overview of the pipeline behavior as a means of estimating assembly execution time. Most application programmers never need to analyze the performance of individual instructions.



5.4.5. Instruction Performance

The Nios II/e core dispatches a single instruction at a time, and the processor waits for an instruction to complete before fetching and dispatching the next instruction. Because each instruction completes before the next instruction is dispatched, branch prediction is not necessary. This greatly simplifies the consideration of processor stalls. Maximum performance is one instruction per six clock cycles. To achieve six cycles, the Avalon-MM instruction master port must fetch an instruction in one clock cycle. A stall on the Avalon-MM instruction master port directly extends the execution time of the instruction.

Table 73. Instruction Execution Performance for Nios II/e Core

Instruction	Cycles
Normal ALU instructions (e.g., add, cmplt)	6
All branch, jmp, jmp _l , ret, call, call _r	6
trap, break, eret, bret, flushp, wrctl, rdctl, unimplemented	6
All load word	6 + Duration of Avalon-MM read transfer
All load halfword	9 + Duration of Avalon-MM read transfer
All load byte	10 + Duration of Avalon-MM read transfer
All store	6 + Duration of Avalon-MM write transfer
All shift, all rotate	7 to 38
All other instructions	6
Combinatorial custom instructions	6
Multicycle custom instructions	6


5.4.6. Exception Handling

The Nios II/e core supports the following exception types:

- Internal hardware interrupt
- Software trap
- Illegal instruction
- Unimplemented instruction

5.4.7. JTAG Debug Module

The Nios II/e core supports the JTAG debug module to provide a JTAG interface to software debugging tools. The JTAG debug module on the Nios II/e core does not support hardware breakpoints or trace.



5.5. Nios II Core Implementation Details Revision History

Document Version	Changes
2019.04.30	Maintenance release
2018.04.18	Implemented editorial enhancements.
2017.05.08	Added link to <i>Nios II Performance Benchmarks</i> .
2016.10.28	Maintenance release.
2015.04.02	Initial release

6. Nios II Processor Versions

Each release of the Nios II Embedded Design Suite (EDS) introduces improvements to the Nios II processor, the software development tools, or both. This chapter catalogs the history of revisions to the Nios II processor; it does not track revisions to development tools, such as the Nios II Software Build Tools (SBT).

Improvements to the Nios II processor might affect:

- Features of the Nios II architecture—An example of an architecture revision is adding instructions to support floating-point arithmetic.
- Implementation of a specific Nios II core—An example of a core revision is increasing the maximum possible size of the data cache memory for the Nios II/f core.
- Features of the JTAG debug module—An example of a JTAG debug module revision is adding an additional trigger input to the JTAG debug module, allowing it to halt processor execution on a new type of trigger event.

Intel FPGA implements Nios II revisions such that code written for an existing Nios II core also works on future revisions of the same core.

6.1. Nios II Versions Revision History

The number for any version of the Nios II processor is determined by the version of the Nios II EDS.

Table 74. Document Revision History

Document Version	Changes
2019.04.30	Maintenance release
2015.01.01	Initial release of the Nios II processor.

6.2. Architecture Revisions

Architecture revisions augment the fundamental capabilities of the Nios II architecture, and affect all Nios II cores. A change in the architecture mandates a revision to all Nios II cores to accommodate the new architectural enhancement. For example, when Intel FPGA adds a new instruction to the instruction set, Intel FPGA consequently must update all Nios II cores to recognize the new instruction.

Table 75. Nios II Architecture Revisions

Version	Release Date	Notes
14.0	January 2015	Initial release of the Nios II processor architecture.

6.3. Core Revisions

Core revisions introduce changes to an existing Nios II core. Core revisions most commonly fix identified bugs, or add support for an architecture revision. Not every Nios II core is revised with every release of the Nios II architecture.

6.3.1. Nios II/f Core

Table 76. Nios II/f Core Revisions

Version	Release Date	Notes
14.0	January 2015	Initial release of the Nios IIf core.

6.3.2. Nios II/s Core

Table 77. Nios II/s Core Revisions

Version	Release Date	Notes
14.0	January 2015	Initial release of the Nios II s core.

6.3.3. Nios II/e Core

Table 78. Nios II/e Core Revisions

Version	Release Date	Notes
14.0	January 2015	Initial release of the Nios II e core.

6.4. JTAG Debug Module Revisions

JTAG debug module revisions augment the debug capabilities of the Nios II processor, or fix bugs isolated within the JTAG debug module logic.

Table 79. JTAG Debug Module Revisions

Version	Release Date	Notes
14.0	January 2015	Initial release of the JTAG debug module.

6.5. Nios II Processor Versions Revision History

Document Version	Changes
2019.04.30	Maintenance release
2018.04.18	Implemented editorial enhancements.
2016.10.28	Maintenance release.
2015.04.02	Initial release

7. Application Binary Interface

This chapter describes the Application Binary Interface (ABI) for the Nios[®] II processor. The ABI describes:

- How data is arranged in memory
- Behavior and structure of the stack
- Function calling conventions

7.1. Data Types

Table 80. Representation of Data C/C++ Types

Type	Size (Bytes)	Representation
char, signed char	1	two's complement (ASCII)
unsigned char	1	binary (ASCII)
short, signed short	2	two's complement
unsigned short	2	binary
int, signed int	4	two's complement
unsigned int	4	binary
long, signed long	4	two's complement
unsigned long	4	binary
float	4	IEEE
double	8	IEEE
pointer	4	binary
long long	8	two's complement
unsigned long long	8	binary

7.2. Memory Alignment

Contents in memory are aligned as follows:

- A function must be aligned to a minimum of 32-bit boundary.
- The minimum alignment of a data element is its natural size. A data element larger than 32 bits need only be aligned to a 32-bit boundary.
- Structures, unions, and strings must be aligned to a minimum of 32 bits.
- Bit fields inside structures are always 32-bit aligned.

7.3. Register Usage

The ABI adds additional usage conventions to the Nios II register file defined in the *Programming Model* chapter of the *Nios II Processor Reference Handbook*.

Table 81. Nios II ABI Register Usage

Register	Name	Used by Compiler	Callee Saved ⁽³⁶⁾	Normal Usage
r0	zero	v		0x00000000
r1	at			Assembler temporary
r2		v		Return value (least-significant 32 bits)
r3		v		Return value (most-significant 32 bits)
r4		v		Register arguments (first 32 bits)
r5		v		Register arguments (second 32 bits)
r6		v		Register arguments (third 32 bits)
r7		v		Register arguments (fourth 32 bits)
r8		v		Caller-saved general-purpose registers
r9		v		
r10		v		
r11		v		
r12		v		
r13		v		
r14		v		
r15		v		
r16		v	v	Callee-saved general-purpose registers
r17		v	v	
r18		v	v	
r19		v	v	
r20		v	v	
r21		v	v	
r22		v	⁽³⁷⁾	
r23		v	⁽³⁸⁾	

continued...

⁽³⁶⁾ A function can use one of these registers if it saves it first. The function must restore the register's original value before exiting.

⁽³⁷⁾ In the GNU Linux operating system, r22 points to the global offset table (GOT). Otherwise, it is available as a callee-saved general-purpose register.



Register	Name	Used by Compiler	Callee Saved ⁽³⁶⁾	Normal Usage
r24	et			Exception temporary
r25	bt			Break temporary
r26	gp	v		Global pointer
r27	sp	v		Stack pointer
r28	fp	v	⁽³⁹⁾	Frame pointer
r29	ea			Exception return address
r30	ba			<ul style="list-style-type: none"> • Normal register set: Break return address • Shadow register sets: SSTATUS register
r31	ra	v		Return address

The endianness of values greater than 8 bits is little endian. The upper 8 bits of a value are stored at the higher byte address.

Related Information

- [Frame Pointer Elimination](#) on page 146
- [Programming Model](#) on page 33

7.4. Stacks

The stack grows downward (i.e. towards lower addresses). The stack pointer points to the last used slot. The frame pointer points to the saved frame pointer near the top of the stack frame.

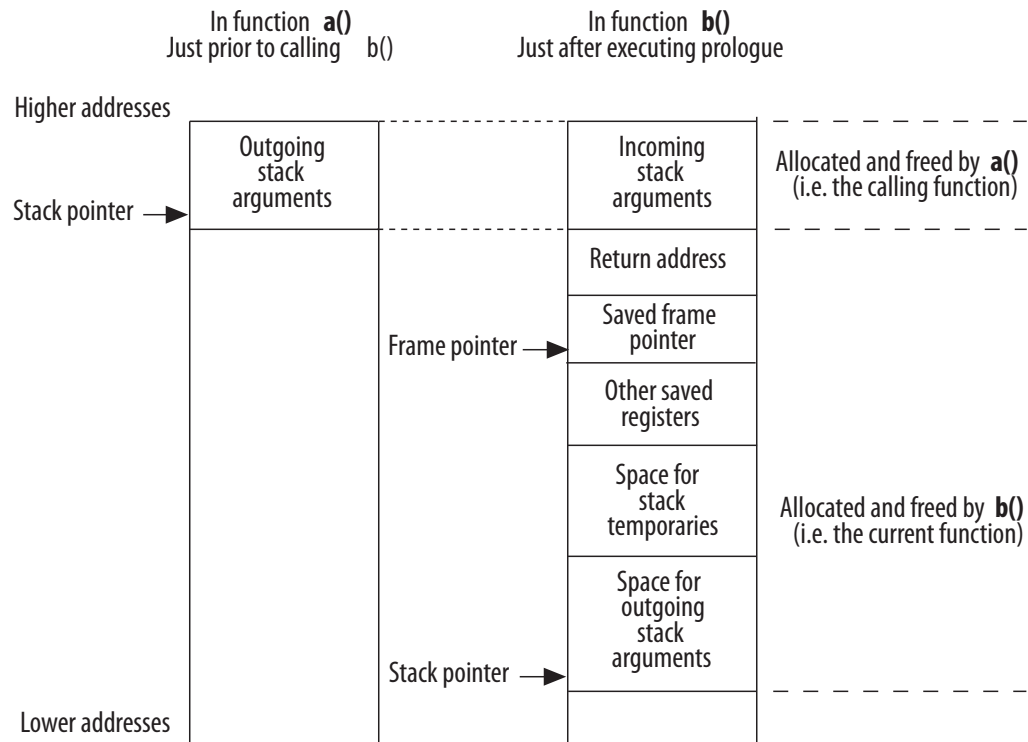
The figure below shows an example of the structure of a current frame. In this case, function a () calls function b (), and the stack is shown before the call and after the prologue in the called function has completed.

⁽³⁶⁾ A function can use one of these registers if it saves it first. The function must restore the register's original value before exiting.

⁽³⁸⁾ In the GNU Linux operating system, r23 is used as the thread pointer. Otherwise, it is available as a callee-saved general-purpose register.

⁽³⁹⁾ If the frame pointer is not used, the register is available as a callee-saved temporary register. Refer to "Frame Pointer Elimination" .

Figure 13. Stack Pointer, Frame Pointer and the Current Frame



Each section of the current frame is aligned to a 32-bit boundary. The ABI requires the stack pointer be 32-bit aligned at all times.

7.4.1. Frame Pointer Elimination

The frame pointer is provided for debugger support. If you are not using a debugger, you can optimize your code by eliminating the frame pointer, using the `-fomit-frame-pointer` compiler option. When the frame pointer is eliminated, register `fp` is available as a temporary register.

7.4.2. Call Saved Registers

The compiler is responsible for generating code to save registers that need to be saved on entry to a function, and to restore the registers on exit. If there are any such registers, they are saved on the stack, from high to low addresses, in the following order: `ra`, `fp`, `sp`, `gp`, `r25`, `r24`, `r23`, `r22`, `r21`, `r20`, `r19`, `r18`, `r17`, `r16`, `r15`, `r14`, `r13`, `r12`, `r11`, `r10`, `r9`, `r8`, `r7`, `r6`, `r5`, `r4`, `r3`, and `r2`. Stack space is not allocated for registers that are not saved.

7.4.3. Further Examples of Stacks

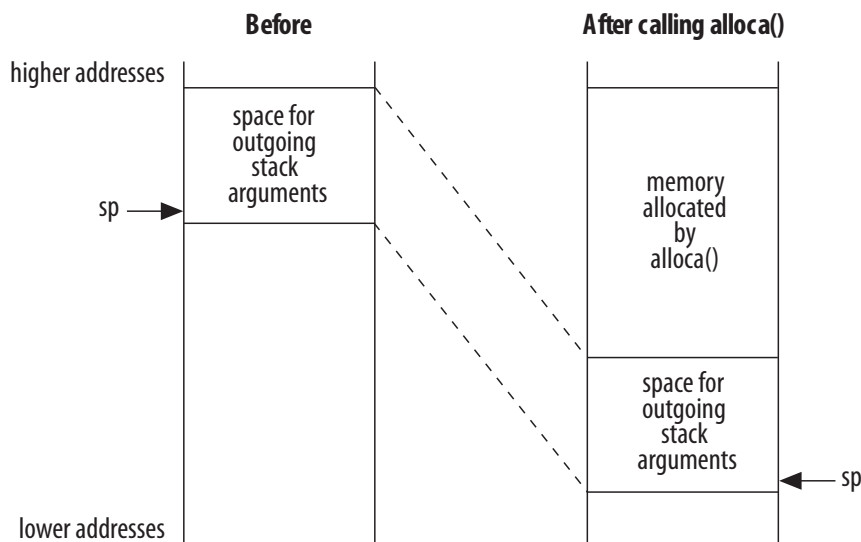
There are a number of special cases for stack layout, which are described in this section.

7.4.3.1. Stack Frame for a Function With `alloca()`

The Nios II stack frame implementation provides support for the `alloca()` function, defined in the Berkeley Software Distribution (BSD) extension to C, and implemented by the gcc compiler. The space allocated by `alloca()` replaces the outgoing arguments and the outgoing arguments get new space allocated at the bottom of the frame.

Note: The Nios II C/C++ compiler maintains a frame pointer for any function that calls `alloca()`, even if `-fomit-frame-pointer` is specified.

Figure 14. Stack Frame after Calling `alloca()`

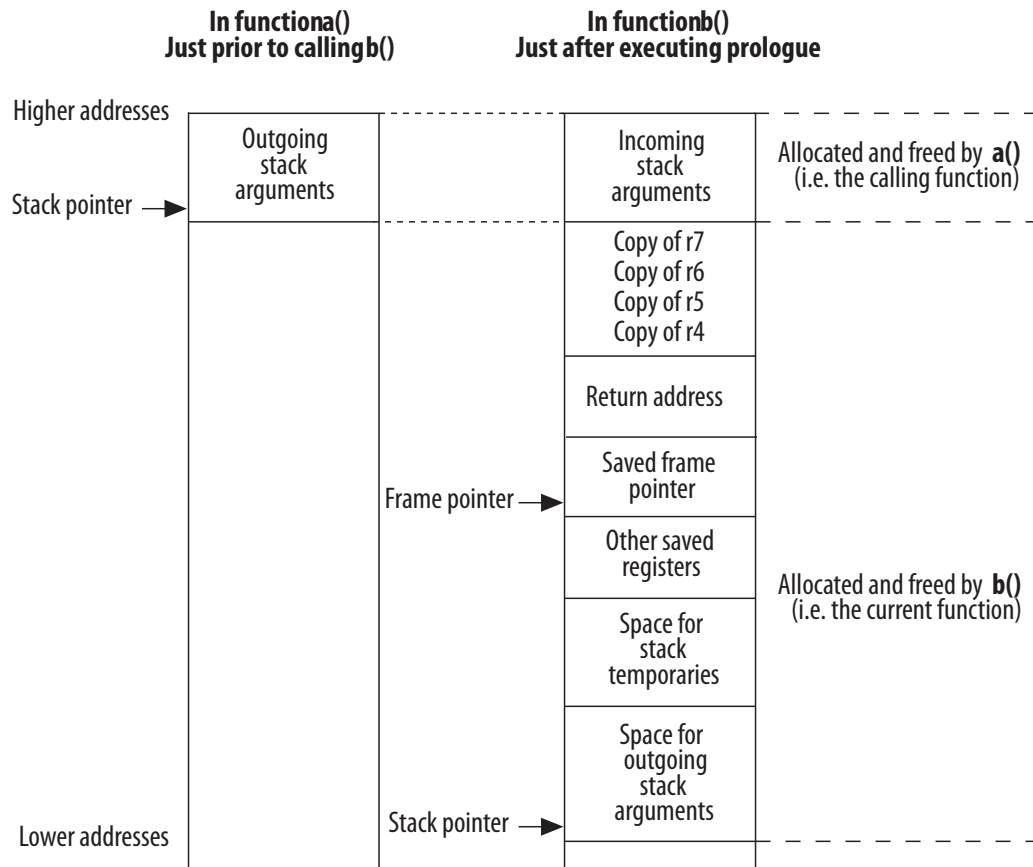


7.4.3.2. Stack Frame for a Function with Variable Arguments

Functions that take variable arguments (`varargs`) still have their first 16 bytes of arguments arriving in registers `r4` through `r7`, just like other functions.

In order for `varargs` to work, functions that take variable arguments allocate 16 extra bytes of storage on the stack. They copy to the stack the first 16 bytes of their arguments from registers `r4` through `r7` as shown below.

Figure 15. Stack Frame Using Variable Arguments



7.4.3.3. Stack Frame for a Function with Structures Passed By Value

Functions that take `struct` value arguments still have their first 16 bytes of arguments arriving in registers `r4` through `r7`, just like other functions.

If part of a structure is passed using registers, the function might need to copy the register contents back to the stack. This operation is similar to that required in the variable arguments case as shown in the figure above, *Stack Frame Using Variable Arguments*.

Related Information

[Stack Frame for a Function with Variable Arguments](#) on page 147

7.4.4. Function Prologues

The Nios II C/C++ compiler generates function prologues that allocate the stack frame of a function for storage of stack temporaries and outgoing arguments. In addition, each prologue is responsible for saving the state of the calling function. This entails saving certain registers on the stack. These registers, the callee-saved registers, are listed in Nios II ABI Register Usage Table in the Register Usage section. A function prologue is required to save a callee-saved register only if the function uses the register.



Given the function prologue algorithm, when doing a back trace, a debugger can disassemble instructions and reconstruct the processor state of the calling function.

Note: An even better way to find out what the prologue has done is to use information stored in the DWARF-2 debugging fields of the executable and linkable format (**.elf**) file.

The instructions found in a Nios II function prologue perform the following tasks:

- Adjust the stack pointer (to allocate the frame)
- Store registers to the frame
- Set the frame pointer to the location of the saved frame pointer

Example 4. A function prologue

```
/* Adjust the stack pointer */
addi    sp, sp, -16    /* make a 16-byte frame */

/* Store registers to the frame */
stw     ra, 12(sp)    /* store the return address */
stw     fp, 8(sp)     /* store the frame pointer*/
stw     r16, 4(sp)    /* store callee-saved register */
stw     r17, 0(sp)    /* store callee-saved register */

/* Set the new frame pointer */
addi    fp, sp, 8
```

Related Information

[Register Usage](#) on page 144

7.4.4.1. Prologue Variations

The following variations can occur in a prologue:

- If the function's frame size is greater than 32,767 bytes, extra temporary registers are used in the calculation of the new stack pointer as well as for the offsets of where to store callee-saved registers. The extra registers are needed because of the maximum size of immediate values allowed by the Nios II processor.
- If the frame pointer is not in use, the final instruction, recalculating the frame pointer, is not generated.
- If variable arguments are used, extra instructions store the argument registers on the stack.
- If the compiler designates the function as a leaf function, the return address is not saved.
- If optimizations are on, especially instruction scheduling, the order of the instructions might change and become interlaced with instructions located after the prologue.

7.5. Arguments and Return Values

This section discusses the details of passing arguments to functions and returning values from functions.

7.5.1. Arguments

The first 16 bytes to a function are passed in registers `r4` through `r7`. The arguments are passed as if a structure containing the types of the arguments were constructed, and the first 16 bytes of the structure are located in `r4` through `r7`.

A simple example:

```
int function (int a, int b);
```

The equivalent structure representing the arguments is:

```
struct { int a; int b; };
```

The first 16 bytes of the `struct` are assigned to `r4` through `r7`. Therefore `r4` is assigned the value of `a` and `r5` the value of `b`.

The first 16 bytes to a function taking variable arguments are passed the same way as a function not taking variable arguments. The called function must clean up the stack as necessary to support the variable arguments.

Refer to [Stack Frame for a Function with Variable Arguments](#)

Related Information

[Stack Frame for a Function with Variable Arguments](#) on page 147

7.5.2. Return Values

Return values of types up to 8 bytes are returned in `r2` and `r3`. For return values greater than 8 bytes, the caller must allocate memory for the result and must pass the address of the result memory as a hidden zero argument.

The hidden zero argument is best explained through an example.

Example 5. Returned struct

```
/* b() computes a structure-type result and returns it */
STRUCT b(int i, int j)
{
    ...
    return result;
}
void a(...)
{
    ...
    value = b(i, j);
}
```

In the example above, if the result type is no larger than 8 bytes, `b()` returns its result in `r2` and `r3`.

If the return type is larger than 8 bytes, the Nios II C/C++ compiler treats this program as if `a()` had passed a pointer to `b()`. The example below shows how the Nios II C/C++ compiler sees the code in the Returned Struct example above.



Example 6. Returned struct is Larger than 8 Bytes

```
void b(STRUCT *p_result, int i, int j)
{
    ...
    *p_result = result;
}

void a(...)
{
    STRUCT value;
    ...
    b(&value, i, j);
}
```

7.6. DWARF-2 Definition

Registers r0 through r31 are assigned numbers 0 through 31 in all DWARF-2 debugging sections.

7.7. Object Files

Table 82. Nios II-Specific ELF Header Values

Member	Value
e_ident[EI_CLASS]	ELFCLASS32
e_ident[EI_DATA]	ELFDATA2LSB
e_machine	EM_ALTERA_NIOS2 == 113

7.8. Relocation

In a Nios II object file, each relocatable address reference possesses a relocation type. The relocation type specifies how to calculate the relocated address. The bit mask specifies where the address is found in the instruction.


Table 83. Nios II Relocation Calculation

Name	Value	Overflow check (40)	Relocated Address R	Bit Mask M	Bit Shift B
R_NIOS2_NONE	0	n/a	None	n/a	n/a
R_NIOS2_S16	1	Yes	S + A	0x003FFFC0	6
R_NIOS2_U16	2	Yes	S + A	0x003FFFC0	6
R_NIOS2_PCREL16	3	Yes	((S + A) - 4) - PC	0x003FFFC0	6
R_NIOS2_CALL26 ⁽⁴¹⁾	4	Yes	(S + A) >> 2	0xFFFFFC0	6

continued...

(40) For relocation types where no overflow check is performed, the relocated address is truncated to fit the instruction.


(41) Linker is permitted to clobber register AT in the course of resolving overflows



Name	Value	Overflow check (40)	Relocated Address R	Bit Mask M	Bit Shift B
R_NIOS2_CALL26_NOAT	41	No	$(S + A) \gg 2$	0xFFFFF0C0	6
R_NIOS2_IMM5	5	Yes	$(S + A) \& 0x1F$	0x000007C0	6
R_NIOS2_CACHE_OPX	6	Yes	$(S + A) \& 0x1F$	0x07C00000	22
R_NIOS2_IMM6	7	Yes	$(S + A) \& 0x3F$	0x00000FC0	6
R_NIOS2_IMM8	8	Yes	$(S + A) \& 0xFF$	0x00003FC0	6
R_NIOS2_HI16	9	No	$((S + A) \gg 16) \& 0xFFFF$	0x003FFFC0	6
R_NIOS2_LO16	10	No	$(S + A) \& 0xFFFF$	0x003FFFC0	6
R_NIOS2_HIADJ16	11	No	Adj(S+A)	0x003FFFC0	6
R_NIOS2_BFD_RELOC_32	12	No	S + A	0xFFFFFFFF	0
R_NIOS2_BFD_RELOC_16	13	Yes	$(S + A) \& 0xFFFF$	0x0000FFFF	0
R_NIOS2_BFD_RELOC_8	14	Yes	$(S + A) \& 0xFF$	0x000000FF	0
R_NIOS2_GPREL	15	No	$(S + A - GP) \& 0xFFFF$	0x003FFFC0	6
R_NIOS2_GNU_VTINHERIT	16	n/a	None	n/a	n/a
R_NIOS2_GNU_VTENTRY	17	n/a	None	n/a	n/a
R_NIOS2_UJMP	18	No	$((S + A) \gg 16) \& 0xFFFF,$ $(S + A + 4) \& 0xFFFF$	0x003FFFC0	6
R_NIOS2_CJMP	19	No	$((S + A) \gg 16) \& 0xFFFF,$ $(S + A + 4) \& 0xFFFF$	0x003FFFC0	6
R_NIOS2_CALLR	20	No	$((S + A) \gg 16) \& 0xFFFF,$ $(S + A + 4) \& 0xFFFF$	0x003FFFC0	6
R_NIOS2_ALIGN	21	n/a	None	n/a	n/a
R_NIOS2_GOT16	22 ⁽⁴²⁾	Yes	G	0x003FFFC0	6
R_NIOS2_CALL16	23 ⁽⁴²⁾	Yes	G	0x003FFFC0	6
R_NIOS2_GOTOFF_LO	24 ⁽⁴²⁾	No	$(S + A - GOT) \& 0xFFFF$	0x003FFFC0	6
R_NIOS2_GOTOFF_HA	25 ⁽⁴²⁾	No	Adj (S + A - GOT)	0x003FFFC0	6
R_NIOS2_PCREL_LO	26 ⁽⁴²⁾	No	$(S + A - PC) \& 0xFFFF$	0x003FFFC0	6
R_NIOS2_PCREL_HA	27 ⁽⁴²⁾	No	Adj (S + A - PC)	0x003FFFC0	6
R_NIOS2_TLS_GD16	28 ⁽⁴²⁾	Yes	Refer to Thread-Local Storage section	0x003FFFC0	6
R_NIOS2_TLS_LDM16	29 ⁽⁴²⁾	Yes	Refer to Thread-Local Storage section	0x003FFFC0	6

continued...

⁽⁴⁰⁾ For relocation types where no overflow check is performed, the relocated address is truncated to fit the instruction.



Name	Value	Overflow check (40)	Relocated Address R	Bit Mask M	Bit Shift B
R_NIOS2_TLS_LDO16	30 ⁽⁴²⁾	Yes	Refer to Thread-Local Storage section	0x003FFFC0	6
R_NIOS2_TLS_IE16	31 ⁽⁴²⁾	Yes	Refer to Thread-Local Storage section	0x003FFFC0	6
R_NIOS2_TLS_LE16	32 ⁽⁴²⁾	Yes	Refer to Thread-Local Storage section	0x003FFFC0	6
R_NIOS2_TLS_DTPMOD	33 ⁽⁴²⁾	No	Refer to Thread-Local Storage section	0xFFFFFFFF	0
R_NIOS2_TLS_DTPREL	34 ⁽⁴²⁾	No	Refer to Thread-Local Storage section	0xFFFFFFFF	0
R_NIOS2_TLS_TPREL	35 ⁽⁴²⁾	No	Refer to Thread-Local Storage section	0xFFFFFFFF	0
R_NIOS2_COPY	36 ⁽⁴²⁾	No	Refer to Copy Relocation section.	n/a	n/a
R_NIOS2_GLOB_DAT	37 ⁽⁴²⁾	No	S	0xFFFFFFFF	0
R_NIOS2_JUMP_SLOT	38 ⁽⁴²⁾	No	Refer to Jump Slot Relocation section.	0xFFFFFFFF	0
R_NIOS2_RELATIVE	39 ⁽⁴²⁾	No	BA+A	0xFFFFFFFF	0
R_NIOS2_GOTOFF	40 ⁽⁴²⁾	No	S+A	0xFFFFFFFF	0
R_NIOS2_GOT_LO	42 ⁽⁴²⁾	No	G & 0xFFFF	0x003FFFC0	6
R_NIOS2_GOT_HA	43 ⁽⁴²⁾	No	Adj(G)	0x003FFFC0	6
R_NIOS2_CALL_LO	44 ⁽⁴²⁾	No	G & 0xFFFF	0x003FFFC0	6
R_NIOS2_CALL_HA	45 ⁽⁴²⁾	No	Adj(G)	0x003FFFC0	6

Expressions in the table above use the following conventions:

- S: Symbol address
- A: Addend
- PC: Program counter
- GP: Global pointer
- Adj(X): $((X \gg 16) \& 0xFFFF) + ((X \gg 15) \& 0x1) \& 0xFFFF$
- BA: The base address at which a shared library is loaded
- GOT: The value of the Global Offset Table (GOT) pointer (Linux only)
- G: The offset into the GOT for the GOT slot for symbol S (Linux only)

⁽⁴⁰⁾ For relocation types where no overflow check is performed, the relocated address is truncated to fit the instruction.

⁽⁴²⁾ Relocation support is provided for Linux systems.



With the information in the table above, any Nios II instruction can be relocated by manipulating it as an unsigned 32-bit integer, as follows:

$$Xr = ((R \ll B) \& M | (X \& \sim M));$$

where:

- R is the relocated address, calculated in the above table
- B is the bit shift
- M is the bit mask
- X is the original instruction
- Xr is the relocated instruction

Related Information

- [Jump Slot Relocation](#) on page 156
- [Copy Relocation](#) on page 155
- [Thread-Local Storage](#) on page 156

7.9. ABI for Linux Systems

This section describes details specific to Linux systems beyond the Linux-specific information in Nios II ABI Register Usage Table and the Nios II Relocation Calculation Table.

Related Information

- [Relocation](#) on page 151
- [Register Usage](#) on page 144

7.9.1. Linux Toolchain Relocation Information

Dynamic relocations can appear in the runtime relocation sections of executables and shared objects, but never appear in object files (with the exception of R_NIOS2_TLS_DTPREL, which is used for debug information). No other relocations are dynamic.

Table 84. Dynamic Relocations

R_NIOS2_TLS_DTPMOD
R_NIOS2_TLS_DTPREL
R_NIOS2_TLS_TPREL
R_NIOS2_COPY
R_NIOS2_GLOB_DAT
R_NIOS2_JUMP_SLOT
R_NIOS2_RELATIVE



A global offset table (GOT) entry referenced using R_NIOS2_GOT16, R_NIOS2_GOT_LO as well as R_NIOS2_GOT_HA must be resolved at load time. A GOT entry referenced only using R_NIOS2_CALL16, R_NIOS2_CALL_LO as well as R_NIOS2_CALL_HA can initially refer to a procedure linkage table (PLT) entry and then be resolved lazily.

Because the TP-relative relocations are 16-bit relocations, no dynamic object using local dynamic or local executable thread-local storage (TLS) can have more than 64 KB of TLS data. New relocations might be added to support this in the future.

Several new assembler operators are defined to generate the Linux-specific relocations, as listed in the table below.

Table 85. Relocation and Operator

Relocation	Operator
R_NIOS2_GOT16	%got
R_NIOS2_CALL16	%call
R_NIOS2_GOTOFF_LO	%gotoff_hiadj
R_NIOS2_GOTOFF_HA	%gotoff_lo
R_NIOS2_PCREL_LO	%hiadj
R_NIOS2_PCREL_HA	%lo
R_NIOS2_TLS_GD16	%tls_gd
R_NIOS2_TLS_LDM16	%tls_ldm
R_NIOS2_TLS_LDO16	%tls_ldo
R_NIOS2_TLS_IE16	%tls_ie
R_NIOS2_TLS_LE16	%tls_le
R_NIOS2_TLS_DTPREL	%tls_ldo
R_NIOS2_GOTOFF	%gotoff
R_NIOS2_GOT_LO	%got_lo
R_NIOS2_GOT_HA	%got_hiadj
R_NIOS2_CALL_LO	%call_lo
R_NIOS2_CALL_HA	%call_hiadj

The %hiadj and %lo operators generate PC-relative or non-PC-relative relocations, depending whether the expression being relocated is PC-relative. For instance, %hiadj(_gp_got - .) generates R_NIOS2_PCREL_HA. %tls_ldo generates R_NIOS2_TLS_LDO16 when used as an immediate operand, and R_NIOS2_TLS_DTPREL when used with the .word directive.

7.9.1.1. Copy Relocation

The R_NIOS2_COPY relocation is used to mark variables allocated in the executable that are defined in a shared library. The variable's initial value is copied from the shared library to the relocated location.

7.9.1.2. Jump Slot Relocation

Jump slot relocations are used for the PLT.

For information about the PLT, refer to "Procedure Linkage Table" section.

Related Information

- [Procedure Linkage Table](#) on page 162
- [Procedure Linkage Table](#) on page 162

7.9.1.3. Thread-Local Storage

The Nios II processor uses the Variant I model for thread-local storage.

The end of the thread control block (TCB) is located 0x7000 bytes before the thread pointer. The TCB is eight bytes long. The first word is the dynamic thread pointer (DTV) pointer and the second word is reserved. Each module's dynamic thread pointer is biased by 0x8000 (when retrieved using `__tls_get_addr`). The thread library can store additional private information before the TCB.

In the GNU Linux toolchain, the GOT pointer (`_gp_got`) is always kept in `r22`, and the thread pointer is always kept in `r23`.

In the following examples, any registers can be used, except that the argument to `__tls_get_addr` is always passed in `r4` and its return value is always returned in `r2`. Calls to `__tls_get_addr` must use the normal position-independent code (PIC) calling convention in PIC code; these sequences are for example only, and the compiler might generate different sequences. No linker relaxations are defined.

Example 7. General Dynamic Model

```
addi r4, r22, %tls_gd(x)           # R_NIOS2_TLS_GD16 x
call __tls_get_addr                # R_NIOS2_CALL26 __tls_get_addr
# Address of x in r2
```

In the general dynamic model, a two-word GOT slot is allocated for `x`, as shown in "GOT Slot for General Dynamic Model" example.

Example 8. GOT Slot for General Dynamic Model

```
GOT[n]                R_NIOS2_TLS_DTPMOD x
GOT[n+1]              R_NIOS2_TLS_DTPREL x
```

Example 9. Local Dynamic Model

```
addi r4, r22, %tls_ldm(x)         # R_NIOS2_TLS_LDM16 x
call __tls_get_addr                # R_NIOS2_CALL26 __tls_get_addr
addi r5, r2, %tls_ldo(x)          # R_NIOS2_TLS_LDO16 x
# Address of x in r5
ldw r6, %tls_ldo(x2)(r2)         # R_NIOS2_TLS_LDO16 x2
# Value of x2 in r6
```

One 2-word GOT slot is allocated for all `R_NIOS2_TLS_LDM16` operations in the linked object. Any thread-local symbol in this object can be used, as shown in "GOT Slot with Thread-Local Storage" example.



Example 10. GOT Slot with Thread-Local Storage

```
GOT[n]          R_NIOS2_TLS_DTPMOD x
GOT[n+1]        0
```

Example 11. Initial Exec Model

```
ldw    r4, %tls_ie(x)(r22)    # R_NIOS2_TLS_IE16 x
add    r4, r23, r4
# Address of x in r4
```

A single GOT slot is allocated to hold the offset of x from the thread pointer, as shown in "GOT Slot for Initial Exec Model" example.

Example 12. GOT Slot for Initial Exec Model

```
GOT[n]          R_NIOS2_TLS_TPREL x
```

Example 13. Local Exec Model

```
addi   r4, r23, %tls_le(x)    # R_NIOS2_TLS_LE16 x
# Address of x in r4
```

There is no GOT slot associated with the local exec model.

Debug information uses the GNU extension DW_OP_GNU_push_tls_address.

Example 14. Debug Information

```
.byte 0x03          # DW_OP_addr
.word %tls_ldo(x)   # R_NIOS2_TLS_DTPREL x
.byte 0xe0          # DW_OP_GNU_push_tls_address
```

7.9.2. Linux Function Calls

Register `r23` is reserved for the thread pointer on GNU Linux systems. It is initialized by the C library and it may be used directly for TLS access, but not modified. On non-Linux systems `r23` is a general-purpose, callee-saved register.

The global pointer, `r26` or `gp`, is globally fixed. It is initialized in startup code and always valid on entry to a function. This method does not allow for multiple `gp` values, so `gp`-relative data references are only possible in the main application (that is, from position dependent code). `gp` is only used for small data access, not GOT access, because code compiled as PIC may be used from shared libraries. The linker may take advantage of `gp` for shorter PLT sequences when the addresses are in range. The compiler needs an option to disable use of `gp`; the option is necessary for applications with excessive amounts of small data. For comparison, XUL (Mozilla display engine, 16 MB code, 2 MB data) has only 27 KB of small data and the limit is 64 KB. This option is separate from `-G 0`, because `-G 0` creates ABI incompatibility. A file compiled with `-G 0` puts global `int` variables into `.data` but files compiled with `-G 8` expect such `int` variables to be in `.sdata`.



PIC code which needs a GOT pointer needs to initialize the pointer locally using `nextpc`; the GOT pointer is not passed during function calls. This approach is compatible with both static relocatable binaries and System V style shared objects. A separate ABI is needed for shared objects with independently relocatable text and data.

Stack alignment is 32-bit. The frame pointer points at the top of the stack when it is in use, to simplify backtracing. Insert `alloca` between the local variables and the outgoing arguments. The stack pointer points to the bottom of the outgoing argument area.

A large `struct` return value is handled by passing a pointer in the first argument register (not the disjoint return value register).

7.9.3. Linux Operating System Call Interface

Table 86. Signals for Unhandled Instruction-Related Exceptions

Exception	Signal
Supervisor-only instruction address	SIGSEGV
TLB permission violation (execute)	SIGSEGV
Supervisor-only instruction	SIGILL
Unimplemented instruction	SIGILL
Illegal instruction	SIGILL
Break instruction	SIGTRAP
Supervisor-only data address	SIGSEGV
Misaligned data address	SIGBUS
Misaligned destination address	SIGBUS
Division error	SIGFPE
TLB Permission Violation (read)	SIGSEGV
TLB Permission Violation (write)	SIGSEGV

There are no floating-point exceptions. The optional floating point unit (FPU) does not support exceptions and any process wanting exact IEEE conformance needs to use a soft-float library (possibly accelerated by use of the attached FPU).

The `break` instruction in a user process might generate a `SIGTRAP` signal for that process, but is not required to. Userspace programs should not use the `break` instruction and userspace debuggers should not insert one. If no hardware debugger is connected, the OS should assure that the `break` instruction does not cause the system to stop responding.

For information about userspace debugging, refer to "Userspace Breakpoints".

The page size is 4 KB. Virtual addresses in user mode are all below 2 GB due to the MMU design. The NULL page is not mapped.

Related Information

[Userspace Breakpoints](#) on page 164



7.9.4. Linux Process Initialization

The stack pointer, `sp`, points to the argument count on the stack.

Table 87. Stack Initial State at User Process Start

Purpose	Start Address	Length
Unspecified	High addresses	
Referenced strings		Varies
Unspecified		
Null auxiliary vector entry		4 bytes
Auxiliary vector entries		8 bytes each
NULL terminator for <code>envp</code>		4 bytes
Environment pointers	$sp + 8 + 4 \times argc$	4 bytes each
NULL terminator for <code>argv</code>	$sp + 4 + 4 \times argc$	4 bytes
Argument pointers	$sp + 4$	4 bytes each
Argument count	<code>sp</code>	4 bytes
Unspecified	Low addresses	

If the application should register a destructor function with `atexit`, the pointer is placed in `r4`. Otherwise `r4` is zero.

The contents of all other registers are unspecified. User code should set `fp` to zero to mark the end of the frame chain.

The auxiliary vector is a series of pairs of 32-bit tag and 32-bit value, terminated by an `AT_NULL` tag.

7.9.5. Linux Position-Independent Code

Every position-independent code (PIC) function which uses global data or global functions must load the value of the GOT pointer into a register. Any available register may be used. If a caller-saved register is used the function must save and restore it around calls. If a callee-saved register is used it must be saved and restored around the current function. Examples in this document use `r22` for the GOT pointer.

The GOT pointer is loaded using a PC-relative offset to the `_gp_got` symbol, as shown below.

Example 15. Loading the GOT Pointer

```
nextpc r22
l:
    orhi r1, %hiadj(_gp_got - 1b) # R_NIOS2_PCREL_HA _gp_got
    addi r1, r1, %lo(_gp_got - 1b) # R_NIOS2_PCREL_LO _gp_got - 4
    add r22, r22, r1
    # GOT pointer in r22
```

Data may be accessed by loading its location from the GOT. A single word GOT entry is generated for each referenced symbol.

Example 16. Small GOT Model Entry for Global Symbols

```
addi   r3, r22, %got(x)      # R_NIOS2_GOT16
GOT[n]                                R_NIOS2_GLOB_DAT x
```

Example 17. Large GOT Model Entry for Global Symbols

```
movhi  r3,      %got_hiadj(x) # R_NIOS2_GOT_HA
addi   r3, r3, %got_lo(x)     # R_NIOS2_GOT_LO
add    r3, r3, r22
GOT[n]                                R_NIOS2_GLOB_DAT x
```

For local symbols, the symbolic reference to *x* is replaced by a relative relocation against symbol zero, with the link time address of *x* as an addend, as shown in the example below.

Example 18. Local Symbols for small GOT Model

```
addi   r3, r22, %got(x)      # R_NIOS2_GOT16
GOT[n]                                R_NIOS2_RELATIVE +x
```

Example 19. Local Symbols for large GOT Model

```
movhi  r3,      %got_hiadj(x) # R_NIOS2_GOT_HA
addi   r3, r3, %got_lo(x)     # R_NIOS2_GOT_LO
add    r3, r3, r22
GOT[n]                                R_NIOS2_RELATIVE +x
```

The `call` and `jmp` instructions are not available in position-independent code. Instead, all calls are made through the GOT. Function addresses may be loaded with `%call`, which allows lazy binding. To initialize a function pointer, load the address of the function with `%got` instead. If no input object requires the address of the function its GOT entry is placed in the PLT GOT for lazy binding, as shown in the example below.

For information about the PLT, refer to the "Procedure Linkage Table" section.

Example 20. Small GOT Model entry in PLT GOT

```
ldw    r3, %call(fun)(r22)   # R_NIOS2_CALL16 fun
callr  r3
PLTGOT[n]                                R_NIOS_JUMP_SLOT fun
```

Example 21. Large GOT Model entry in PLT GOT

```
movhi  r3,      %call_hiadj(x) # R_NIOS2_CALL_HA
addi   r3, r3, %call_lo(x)     # R_NIOS2_CALL_LO
add    r3, r3, r22
ldw    r3, 0(r3)
```




```
callr r3
PLTGOT[n]                                R_NIOS_JUMP_SLOT fun
```

When a function or variable resides in the current shared object at compile time, it can be accessed via a PC-relative or GOT-relative offset, as shown below.

Example 22. Accessing Function or Variable in Current Shared Object

```
orhi    r3, %gotoff_hiadj(x)             # R_NIOS2_GOTOFF_HA x
addi    r3, r3, %gotoff_lo(x)           # R_NIOS2_GOTOFF_LO x
add     r3, r22, r3
# Address of x in r3
```

Multiway branches such as switch statements can be implemented with a table of GOT-relative offsets, as shown below.

Example 23. Switch Statement Implemented with Table

```
# Scaled table offset in r4
orhi    r3, %gotoff_hiadj(Ltable)       # R_NIOS2_GOTOFF_HA Ltable
addi    r3, r3, %gotoff_lo(Ltable)     # R_NIOS2_GOTOFF_LO Ltable
add     r3, r22, r3                    # r3 == &Ltable
add     r3, r3, r4
ldw     r4, 0(r3)                       # r3 == Ltable[index]
add     r4, r4, r22                    # Convert offset into
destination
  jmp     r4
...
Ltable:
  .word %gotoff(Label1)
  .word %gotoff(Label2)
  .word %gotoff(Label3)
```

Related Information

[Procedure Linkage Table](#) on page 162

7.9.6. Linux Program Loading and Dynamic Linking

7.9.6.1. Global Offset Table

Because shared libraries are position-independent, they can not contain absolute addresses for symbols. Instead, addresses are loaded from the GOT.

The first word of the GOT is filled in by the link editor with the unrelocated address of the `__DYNAMIC`, which is at the start of the dynamic section. The second and third words are reserved for the dynamic linker.

For information about the dynamic linker, refer to the "Procedure Linkage Table" section.

The linker-defined symbol `__GLOBAL_OFFSET_TABLE__` points to the reserved entries at the beginning of the GOT. The linker-defined symbol `__gp_got` points to the base address used for GOT-relative relocations. The value of `__gp_got` might vary between object files if the linker creates multiple GOT sections.

Related Information

Procedure Linkage Table on page 162

7.9.6.2. Function Addresses

Function addresses use the same `SHN_UNDEF` and `st_value` convention for PLT entries as in other architectures, such as `x86_64`.

7.9.6.3. Procedure Linkage Table

Function calls in a position-dependent executable may use the `call` and `jmp` instructions, which address the contents of a 256-MB segment. They may also use the `%lo`, `%hi`, and `%hiadj` operators to take the address of a function. If the function is in another shared object, the link editor creates a callable stub in the executable called a PLT entry. The PLT entry loads the address of the called function from the PLT GOT (a region at the start of the GOT) and transfers control to it.

The PLT GOT entry needs a relocation referring to the final symbol, of type `R_NIOS2_JUMP_SLOT`. The dynamic linker may immediately resolve it, or may leave it unmodified for lazy binding. The link editor fills in an initial value pointing to the lazy binding stubs at the start of the PLT section.

Each PLT entry appears as shown in the example below.

Example 24. PLT Entry

```
.PLTn:
  orhi   r15, r0, %hiadj(plt_got_slot_address)
  ldw    r15, %lo(plt_got_slot_address)(r15)
  jmp    r15
```

The example below shows the PLT entry when the PLT GOT is close enough to the small data area for a relative jump.

Example 25. PLT Entry Near Small Data Area

```
.PLTn:
  ldw    r15, %gprel(plt_got_slot_address)(gp)
  jmp    r15
```

Example 26. Initial PLT Entry

```
res_0:
  br .PLTresolve
  ...
.PLTresolve:
  orhi   r14, r0, %hiadj(res_0)
  addi   r14, r14, %lo(res_0)
  sub    r15, r15, r14
  orhi   r13, %hiadj(_GLOBAL_OFFSET_TABLE_)
  ldw    r14, %lo(_GLOBAL_OFFSET_TABLE_+4)(r13)
  ldw    r13, %lo(_GLOBAL_OFFSET_TABLE_+8)(r13)
  jmp    r13
```

In front of the initial PLT entry, a series of branches start of the initial entry (the `nextpc` instruction). There is one branch for each PLT entry, labelled `res_0` through `res_N`. The last several branches may be replaced by `nop` instructions to improve



performance. The link editor arranges for the Nth PLT entry to point to the Nth branch; `res_N - res_0` is four times the index into the `.rela.plt` section for the corresponding `R_JUMP_SLOT` relocation.

The dynamic linker initializes `GOT[1]` to a unique identifier for each library and `GOT[2]` to the address of the runtime resolver routine. In order for the two loads in `.PLTresolve` to share the same `%hiadj, _GLOBAL_OFFSET_TABLE_` must be aligned to a 16-byte boundary.

The runtime resolver receives the original function arguments in `r4` through `r7`, the shared library identifier from `GOT[1]` in `r14`, and the relocation index times four in `r15`. The resolver updates the corresponding PLT GOT entry so that the PLT entry transfers control directly to the target in the future, and then transfers control to the target.

In shared objects, the `call` and `jmp` instructions can not be used because the library load address is not known at link time. Calls to functions outside the current shared object must pass through the GOT. The program loads function addresses using `%call`, and the link editor may arrange for such entries to be lazily bound. Because PLT entries are only used for lazy binding, shared object PLTs are smaller, as shown below.

Example 27. Shared Object PLT

```
.PLTn:
    orhi    r15, r0, %hiadj(index * 4)
    addi   r15, r15, %lo(index * 4)
    br     .PLTresolve
```

Example 28. Initial PLT Entry

```
.PLTresolve:
    nextpc  r14
    orhi   r13, r0, %hiadj(_GLOBAL_OFFSET_TABLE_)
    add   r13, r13, r14
    ldw  r14, %lo(_GLOBAL_OFFSET_TABLE_+4)(r13)
    ldw  r13, %lo(_GLOBAL_OFFSET_TABLE_+8)(r13)
    jmp  r13
```

If the initial PLT entry is out of range, the resolver can be inline, because it is only one instruction longer than a long branch, as shown below.

Example 29. Initial PLT Entry Out of Range

```
.PLTn:
    orhi    r15, r0, %hiadj(index * 4)
    addi   r15, r15, %lo(index * 4)
    nextpc  r14
    orhi   r13, r0, %hiadj(_GLOBAL_OFFSET_TABLE_)
    add   r13, r13, r14
    ldw  r14, %lo(_GLOBAL_OFFSET_TABLE_+4)(r13)
    ldw  r13, %lo(_GLOBAL_OFFSET_TABLE_+8)(r13)
    jmp  r13
```

7.9.6.4. Linux Program Interpreter

The program interpreter is `/lib/ld.so.1`.



7.9.6.5. Linux Initialization and Termination Functions

The implementation is responsible for calling `DT_INIT()`, `DT_INIT_ARRAY()`, `DT_PREINIT_ARRAY()`, `DT_FINI()`, and `DT_FINI_ARRAY()`.

7.9.7. Linux Conventions

7.9.7.1. System Calls

The Linux system call interface relies on the `trap` instruction with immediate argument zero. The system call number is passed in register `r2`. The arguments are passed in `r4`, `r5`, `r6`, `r7`, `r8`, and `r9` as necessary. The return value is written in `r2` on success, or a positive error number is written to `r2` on failure. A flag indicating successful completion, to distinguish error values from valid results, is written to `r7`; 0 indicates `syscall` success and 1 indicates `r2` contains a positive `errno` value.

7.9.7.2. Userspace Breakpoints

Userspace breakpoints are accomplished using the `trap` instruction with immediate operand 31 (all ones). The OS must distinguish this instruction from a `trap 0` system call and generate a `trap` signal.

7.9.7.3. Atomic Operations

The Nios II architecture does not have atomic operations (such as load linked and store conditional). Atomic operations are emulated using a kernel system call via the `trap` instruction. The toolchain provides intrinsic functions which perform the system call. Applications must use those functions rather than the system call directly. Atomic operations may be added in a future processor extension.

7.9.7.4. Processor Requirements

Linux requires that a hardware multiplier be present. The full 64-bit multiplier (`mulx` instructions) is not required.

7.10. Development Environment

The following object macros are defined:

- `NIOS2`
- `__NIOS2`
- `__NIOS2__`
- `nios2`
- `__nios2`
- `__nios2__`
- `nios2_little_endian`
- `__nios2_little_endian`
- `__nios2_little_endian__`



The object macro `__nios2_arch__` is predefined to 1 when you compile a program for Nios II R1 ISA and is predefined to 2 when you compile for Nios II R2 ISA.

7.11. Application Binary Interface Revision History

Document Version	Changes
2019.04.30	Maintenance release
2018.04.18	<ul style="list-style-type: none">Implemented editorial enhancements.Updated the information about object macros in <i>Development Environment</i>.
2016.10.28	Maintenance release.
2015.04.02	Initial release

8. Instruction Set Reference

This section introduces the Nios II instruction word format and provides a detailed reference of the Nios II instruction set.

8.1. Word Formats

There are three types of Nios II instruction word format: I-type, R-type, and J-type.

8.1.1. I-Type

The defining characteristic of the I-type instruction word format is that it contains an immediate value embedded within the instruction word. I-type instructions words contain:

- A 6-bit opcode field OP
- Two 5-bit register fields A and B
- A 16-bit immediate data field IMM16

In most cases, fields A and IMM16 specify the source operands, and field B specifies the destination register. IMM16 is considered signed except for logical operations and unsigned comparisons.

I-type instructions include arithmetic and logical operations such as `addi` and `andi`; branch operations; load and store operations; and cache management operations.


Table 88. I-Type Instruction Format

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										OP					

8.1.2. R-Type

The defining characteristic of the R-type instruction word format is that all arguments and results are specified as registers. R-type instructions contain:

- A 6-bit opcode field OP
- Three 5-bit register fields A, B, and C
- An 11-bit opcode-extension field OPX



In most cases, fields A and B specify the source operands, and field C specifies the destination register.

Some R-Type instructions embed a small immediate value in the five low-order bits of OPX. Unused bits in OPX are always 0.

R-type instructions include arithmetic and logical operations such as `add` and `nor`; comparison operations such as `cmpeq` and `cmplt`; the `custom` instruction; and other operations that need only register operands.

Table 89. R-Type Instruction Format

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					OPX
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
OPX										OP					

8.1.3. J-Type

J-type instructions contain:

- A 6-bit opcode field
- A 26-bit immediate data field

J-type instructions, such as `call` and `jmp_i`, transfer execution anywhere within a 256-MB range.

Table 90. J-Type Instruction Format

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
IMM26															
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM26										OP					

8.2. Instruction Opcodes

The OP field in the Nios II instruction word specifies the major class of an opcode as listed in the two tables below. Most values of OP are encodings for I-type instructions. One encoding, OP = 0x00, is the J-type instruction `call`. Another encoding, OP = 0x3a, is used for all R-type instructions, in which case, the OPX field differentiates the instructions. All undefined encodings of OP and OPX are reserved.

Table 91. OP Encodings

OP	Instruction	OP	Instruction	OP	Instruction	OP	Instruction
0x00	<code>call</code>	0x10	<code>cmplt_i</code>	0x20	<code>cmpeq_i</code>	0x30	<code>cmplt_tui</code>
0x01	<code>jmp_i</code>	0x11		0x21		0x31	
0x02		0x12		0x22		0x32	<code>custom</code>

continued...



OP	Instruction	OP	Instruction	OP	Instruction	OP	Instruction
0x03	ldbu	0x13	initda	0x23	ldbuio	0x33	initd
0x04	addi	0x14	ori	0x24	muli	0x34	orhi
0x05	stb	0x15	stw	0x25	stbio	0x35	stwio
0x06	br	0x16	blt	0x26	beq	0x36	bltu
0x07	ldb	0x17	ldw	0x27	ldbio	0x37	ldwio
0x08	cmpgei	0x18	cmpnei	0x28	cmpgeui	0x38	rdprs
0x09		0x19		0x29		0x39	
0x0A		0x1A		0x2A		0x3A	R-type
0x0B	ldhu	0x1B	flushda	0x2B	ldhuio	0x3B	flushd
0x0C	andi	0x1C	xori	0x2C	andhi	0x3C	xorhi
0x0D	sth	0x1D		0x2D	sthio	0x3D	
0x0E	bge	0x1E	bne	0x2E	bgeu	0x3E	
0x0F	ldh	0x1F		0x2F	ldhio	0x3F	

Table 92. OPX Encodings for R-Type Instructions

OPX	Instruction	OPX	Instruction	OPX	Instruction	OPX	Instruction
0x00		0x10	cmplt	0x20	cmpeq	0x30	cmpltu
0x01	eret	0x11		0x21		0x31	add
0x02	roli	0x12	slli	0x22		0x32	
0x03	rol	0x13	sll	0x23		0x33	
0x04	flushp	0x14	wrprs	0x24	divu	0x34	break
0x05	ret	0x15		0x25	div	0x35	
0x06	nor	0x16	or	0x26	rdctl	0x36	sync
0x07	mulxuu	0x17	mulxsu	0x27	mul	0x37	
0x08	cmpge	0x18	cmpne	0x28	cmpgeu	0x38	
0x09	bret	0x19		0x29	initi	0x39	sub
0x0A		0x1A	srl	0x2A		0x3A	srai
0x0B	ror	0x1B	srl	0x2B		0x3B	sra
0x0C	flushi	0x1C	nextpc	0x2C		0x3C	
0x0D	jmp	0x1D	callr	0x2D	trap	0x3D	
0x0E	and	0x1E	xor	0x2E	wrctl	0x3E	
0x0F		0x1F	mulxss	0x2F		0x3F	



8.3. Assembler Pseudo-Instructions

Pseudo-instructions are used in assembly source code like regular assembly instructions. Each pseudo-instruction is implemented at the machine level using an equivalent instruction. The `movia` pseudo-instruction is the only exception, being implemented with two instructions. Most pseudo-instructions do not appear in disassembly views of machine code.

Table 93. Assembler Pseudo-Instructions

Pseudo-Instruction	Equivalent Instruction
<code>bgt rB, rA, label</code>	<code>blt rA, rB, label</code>
<code>bgtu rB, rA, label</code>	<code>bltu rA, rB, label</code>
<code>ble rB, rA, label</code>	<code>bge rA, rB, label</code>
<code>bleu rB, rA, label</code>	<code>bgeu rA, rB, label</code>
<code>cmpgt rC, rB, rA</code>	<code>cmplt rC, rA, rB</code>
<code>cmpgti rB, rA, IMMED</code>	<code>cmpgei rB, rA, (IMMED+1)</code>
<code>cmpgtu rC, rB, rA</code>	<code>cmpltu rC, rA, rB</code>
<code>cmpgtui rB, rA, IMMED</code>	<code>cmpgeui rB, rA, (IMMED+1)</code>
<code>cmple rC, rB, rA</code>	<code>cmpge rC, rA, rB</code>
<code>cmplei rB, rA, IMMED</code>	<code>cmplti rB, rA, (IMMED+1)</code>
<code>cmpleu rC, rB, rA</code>	<code>cmpgeu rC, rA, rB</code>
<code>cmpleui rB, rA, IMMED</code>	<code>cmpltui rB, rA, (IMMED+1)</code>
<code>mov rC, rA</code>	<code>add rC, rA, r0</code>
<code>movhi rB, IMMED</code>	<code>orhi rB, r0, IMMED</code>
<code>movi rB, IMMED</code>	<code>addi, rB, r0, IMMED</code>
<code>movia rB, label</code>	<code>orhi rB, r0, %hiadj(label)</code> <code>addi, rB, r0, %lo(label)</code>
<code>movui rB, IMMED</code>	<code>ori rB, r0, IMMED</code>
<code>nop</code>	<code>add r0, r0, r0</code>
<code>subi rB, rA, IMMED</code>	<code>addi rB, rA, (-IMMED)</code>

Refer to the *Application Binary Interface* chapter of the *Nios II Processor Reference Handbook* for more information about global pointers.

Related Information

[Application Binary Interface](#) on page 143

8.4. Assembler Macros

The Nios II assembler provides macros to extract halfwords from labels and from 32-bit immediate values. These macros return 16-bit signed values or 16-bit unsigned values depending on where they are used. When used with an instruction that requires

a 16-bit signed immediate value, these macros return a value ranging from -32768 to 32767. When used with an instruction that requires a 16-bit unsigned immediate value, these macros return a value ranging from 0 to 65535.

Table 94. Assembler Macros

Macro	Description	Operation
<code>%lo(immed32)</code>	Extract bits [15..0] of immed32	<code>immed32 & 0xFFFF</code>
<code>%hi(immed32)</code>	Extract bits [31..16] of immed32	<code>(immed32 >> 16) & 0xFFFF</code>
<code>%hiadj(immed32)</code>	Extract bits [31..16] and adds bit 15 of immed32	<code>((immed32 >> 16) & 0xFFFF) + ((immed32 >> 15) & 0x1)</code>
<code>%gprel(immed32)</code>	Replace the immed32 address with an offset from the global pointer	<code>immed32 - _gp</code>

Refer to the *Application Binary Interface* chapter of the *Nios II Processor Reference Handbook* for more information about global pointers.

Related Information

[Application Binary Interface](#) on page 143

8.5. Instruction Set Reference

The following pages list all Nios II instruction mnemonics in alphabetical order.

Table 95. Notation Conventions

Notation	Meaning
<code>X – Y</code>	X is written with Y
<code>PC – X</code>	The program counter (PC) is written with address X; the instruction at X is the next instruction to execute
<code>PC</code>	The address of the assembly instruction in question
<code>rA, rB, rC</code>	One of the 32-bit general-purpose registers
<code>prs.rA</code>	General-purpose register rA in the previous register set
<code>IMMn</code>	An n-bit immediate value, embedded in the instruction word
<code>IMMED</code>	An immediate value
<code>X_n</code>	The nth bit of X, where n = 0 is the LSB
<code>X_{n..m}</code>	Consecutive bits n through m of X
<code>0xNNMM</code>	Hexadecimal notation
<code>X : Y</code>	Bitwise concatenation For example, <code>(0x12 : 0x34) = 0x1234</code>
<code>σ(X)</code>	The value of X after being sign-extended to a full register-sized signed integer
<code>X >> n</code>	The value X after being right-shifted n bit positions
<code>X << n</code>	The value X after being left-shifted n bit positions
<code>X & Y</code>	Bitwise logical AND
<code>X Y</code>	Bitwise logical OR

continued...



Notation	Meaning
$X \wedge Y$	Bitwise logical XOR
$\sim X$	Bitwise logical NOT (one's complement)
Mem8[X]	The byte located in data memory at byte address X
Mem16[X]	The halfword located in data memory at byte address X
Mem32[X]	The word located in data memory at byte address X
label	An address label specified in the assembly file
(signed) rX	The value of rX treated as a signed number
(unsigned) rX	The value of rX treated as an unsigned number

Note: All register operations apply to the current register set, except as noted.

The following exceptions are not listed for each instruction because they can occur on any instruction fetch:

- Supervisor-only instruction address
- Fast TLB miss (instruction)
- Double TLB miss (instruction)
- TLB permission violation (execute)
- MPU region violation (instruction)

For information about these and all Nios II exceptions, refer to the *Programming Model* chapter of the *Nios II Processor Reference Handbook*.

Related Information

[Programming Model](#) on page 33

8.5.1. add

Instruction	add
Operation	$rC \leftarrow rA + rB$
Assembler Syntax	add rC, rA, rB
Example	add r6, r7, r8
Description	Calculates the sum of rA and rB. Stores the result in rC. Used for both signed and unsigned addition.
Usage	Carry Detection (unsigned operands): Following an add operation, a carry out of the MSB can be detected by checking whether the unsigned sum is less than one of the unsigned operands. The carry bit can be written to a register, or a conditional branch can be taken based on the carry condition. The following code shows both cases: <pre>add rC, rA, rB cmpltu rD, rC, rA add rC, rA, rB bltu rC, rA, label # The original add operation # rD is written with the carry bit</pre>

continued...



```
# The original add operation
# Branch if carry generated
Overflow Detection (signed operands):
An overflow is detected when two positives are added and
the sum is negative, or when two negatives are added and
the sum is positive. The overflow condition can control a
conditional branch, as shown in the following code:
add rC, rA, rB
xor rD, rC, rA
xor rE, rC, rB
and rD, rD, rE
blt rD, r0,label
# The original add operation
# Compare signs of sum and rA
# Compare signs of sum and rB
# Combine comparisons
# Branch if overflow occurred
```

Exceptions

None

Instruction Type

R

Instruction Fields

A = Register index of operand rA
 B = Register index of operand rB
 C = Register index of operand rC

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x31
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x31					0					0x3A					

8.5.2. addi

Instruction

addi

Operation

$rB \leftarrow rA + \sigma(\text{IMM16})$

Assembler Syntax

addi rB, rA, IMM16

Example

addi r6, r7, -100

Description

Sign-extends the 16-bit immediate value and adds it to the value of rA. Stores the sum in rB.

Usage

Carry Detection (unsigned operands):

Following an addi operation, a carry out of the MSB can be detected by checking whether the unsigned sum is less than one of the unsigned operands. The carry bit can be written to a register, or a conditional branch can be taken based on the carry condition. The following code shows both cases:

```
addi rB, rA, IMM16
cmpltu rD, rB, rA
addi rB, rA, IMM16
bltu rB, rA, label
# The original add operation
# rD is written with the carry bit
```

continued...



```
# The original add operation
# Branch if carry generated
Overflow Detection (signed operands):
An overflow is detected when two positives are added and
the sum is negative, or when two negatives are added and
the sum is positive. The overflow condition can control a
conditional branch, as shown in the following code:
addi rB, rA, IMM16
xor rC, rB, rA
xorhi rD, rB, IMM16
and rC, rC, rD
blt rC, r0,label
# The original add operation
# Compare signs of sum and rA
# Compare signs of sum and IMM16
# Combine comparisons
# Branch if overflow occurred
```

Exceptions

None

Instruction Type

I

Instruction Fields

A = Register index of operand rA
 B = Register index of operand rB
 IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x04					

8.5.3. and

Instruction

bitwise logical and

Operation

$rC \leftarrow rA \& rB$

Assembler Syntax

and rC, rA, rB

Example

and r6, r7, r8

Description

Calculates the bitwise logical AND of rA and rB and stores the result in rC.

Exceptions

None

Instruction Type

R

Instruction Fields

A = Register index of operand rA
 B = Register index of operand rB
 C = Register index of operand rC



Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x0e
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x0e					0					0x3A					

8.5.4. andhi

Instruction	bitwise logical and immediate into high halfword
Operation	$rB \leftarrow rA \& (IMM16 : 0x0000)$
Assembler Syntax	<code>andhi rB, rA, IMM16</code>
Example	<code>andhi r6, r7, 100</code>
Description	Calculates the bitwise logical AND of rA and (IMM16 : 0x0000) and stores the result in rB.
Exceptions	None
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit unsigned immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x2c					

8.5.5. andi

Instruction	bitwise logical and immediate
Operation	$rB \leftarrow rA \& (0x0000 : IMM16)$
Assembler Syntax	<code>andi rB, rA, IMM16</code>
Example	<code>andi r6, r7, 100</code>
Description	Calculates the bitwise logical AND of rA and (0x0000 : IMM16) and stores the result in rB.
Exceptions	None
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit unsigned immediate value



Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x0c					

8.5.6. beq

Instruction	branch if equal
Operation	if (rA == rB) then PC ← PC + 4 + σ(IMM16) else PC ← PC + 4
Assembler Syntax	beq rA, rB, label
Example	beq r6, r7, label
Description	If rA == rB, then beq transfers program control to the instruction at label. In the instruction encoding, the offset given by IMM16 is treated as a signed number of bytes relative to the instruction immediately following beq. The two least-significant bits of IMM16 are always zero, because instruction addresses must be word-aligned.
Exceptions	Misaligned destination address
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x26					

8.5.7. bge

Instruction	branch if greater than or equal signed
Operation	if ((signed) rA >= (signed) rB) then PC ← PC + 4 + σ(IMM16) else PC ← PC + 4
Assembler Syntax	bge rA, rB, label
Example	bge r6, r7, top_of_loop
Description	If (signed) rA >= (signed) rB, then bge transfers program control to the instruction at label. In the instruction encoding, the offset given by IMM16 is treated as a signed number of bytes relative to the instruction immediately

continued...



Exceptions	Misaligned destination address
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A						B						IMM16			
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x0e					

8.5.8. bgeu

Instruction	branch if greater than or equal unsigned
Operation	if ((unsigned) rA >= (unsigned) rB) then PC ← PC + 4 + σ(IMM16) else PC ← PC + 4
Assembler Syntax	bgeu rA, rB, label
Example	bgeu r6, r7, top_of_loop
Description	If (unsigned) rA >= (unsigned) rB, then bgeu transfers program control to the instruction at label. In the instruction encoding, the offset given by IMM16 is treated as a signed number of bytes relative to the instruction immediately following bgeu. The two least-significant bits of IMM16 are always zero, because instruction addresses must be word-aligned.
Exceptions	Misaligned destination address
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A						B						IMM16			
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x2e					

8.5.9. bgt

Instruction	branch if greater than signed
Operation	if ((signed) rA > (signed) rB)

continued...



Assembler Syntax	then PC ← label else PC ← PC + 4
Example	<code>bgt rA, rB, label</code>
Description	<code>bgt r6, r7, top_of_loop</code>
Pseudo-instruction	If (signed) rA > (signed) rB, then <code>bgt</code> transfers program control to the instruction at label. <code>bgt</code> is implemented with the <code>blt</code> instruction by swapping the register operands.

8.5.10. bgtu

Instruction	branch if greater than unsigned
Operation	if ((unsigned) rA > (unsigned) rB) then PC ← label else PC ← PC + 4
Assembler Syntax	<code>bgtu rA, rB, label</code>
Example	<code>bgtu r6, r7, top_of_loop</code>
Description	If (unsigned) rA > (unsigned) rB, then <code>bgtu</code> transfers program control to the instruction at label.
Pseudo-instruction	<code>bgtu</code> is implemented with the <code>bltu</code> instruction by swapping the register operands.

8.5.11. ble

Instruction	branch if less than or equal signed
Operation	if ((signed) rA ≤ (signed) rB) then PC ← label else PC ← PC + 4
Assembler Syntax	<code>ble rA, rB, label</code>
Example	<code>ble r6, r7, top_of_loop</code>
Description	If (signed) rA ≤ (signed) rB, then <code>ble</code> transfers program control to the instruction at label.
Pseudo-instruction	<code>ble</code> is implemented with the <code>bge</code> instruction by swapping the register operands.

8.5.12. bleu

Instruction	branch if less than or equal to unsigned
Operation	if ((unsigned) rA ≤ (unsigned) rB) then PC ← label else PC ← PC + 4
Assembler Syntax	<code>bleu rA, rB, label</code>

continued...



Example	<code>bleu r6, r7, top_of_loop</code>
Description	If (unsigned) $rA \leq$ (unsigned) rB , then <code>bleu</code> transfers program counter to the instruction at label.
Pseudo-instruction	<code>bleu</code> is implemented with the <code>bgeu</code> instruction by swapping the register operands.

8.5.13. blt

Instruction	branch if less than signed
Operation	if ((signed) $rA <$ (signed) rB) then $PC \leftarrow PC + 4 + \sigma(\text{IMM16})$ else $PC \leftarrow PC + 4$
Assembler Syntax	<code>blt rA, rB, label</code>
Example	<code>blt r6, r7, top_of_loop</code>
Description	If (signed) $rA <$ (signed) rB , then <code>blt</code> transfers program control to the instruction at label. In the instruction encoding, the offset given by IMM16 is treated as a signed number of bytes relative to the instruction immediately following <code>blt</code> . The two least-significant bits of IMM16 are always zero, because instruction addresses must be word-aligned.
Exceptions	Misaligned destination address
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x16					

8.5.14. bltu

Instruction	branch if less than unsigned
Operation	if ((unsigned) $rA <$ (unsigned) rB) then $PC \leftarrow PC + 4 + \sigma(\text{IMM16})$ else $PC \leftarrow PC + 4$
Assembler Syntax	<code>bltu rA, rB, label</code>
Example	<code>bltu r6, r7, top_of_loop</code>
Description	If (unsigned) $rA <$ (unsigned) rB , then <code>bltu</code> transfers program control to the instruction at label. In the instruction encoding, the offset given by IMM16 is treated as a signed number of bytes relative to the instruction immediately following <code>bltu</code> . The two least-significant bits of IMM16 are always zero, because instruction addresses must be word-aligned.

continued...



Exceptions	Misaligned destination address
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x36					

8.5.15. bne

Instruction	branch if not equal
Operation	if (rA != rB) then PC ← PC + 4 + σ(IMM16) else PC ← PC + 4
Assembler Syntax	bne rA, rB, label
Example	bne r6, r7, top_of_loop
Description	If rA != rB, then <code>bne</code> transfers program control to the instruction at label. In the instruction encoding, the offset given by IMM16 is treated as a signed number of bytes relative to the instruction immediately following <code>bne</code> . The two least-significant bits of IMM16 are always zero, because instruction addresses must be word-aligned.
Exceptions	Misaligned destination address
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x1e					

8.5.16. br

Instruction	unconditional branch
Operation	PC ← PC + 4 + σ(IMM16)
Assembler Syntax	br label
Example	br top_of_loop

continued...



Description	Transfers program control to the instruction at label. In the instruction encoding, the offset given by IMM16 is treated as a signed number of bytes relative to the instruction immediately following <code>br</code> . The two least-significant bits of IMM16 are always zero, because instruction addresses must be word-aligned.
Exceptions	Misaligned destination address
Instruction Type	I
Instruction Fields	IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
0					0					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x06					

8.5.17. break

Instruction	<code>break</code>
Operation	<code>bstatus ← status</code> <code>PIE ← 0</code> <code>U ← 0</code> <code>ba ← PC + 4</code> <code>PC ← break handler address</code>
Assembler Syntax	<code>break</code> <code>break imm5</code>
Example	<code>break</code>
Description	<p>Breaks program execution and transfers control to the debugger break-processing routine. Saves the address of the next instruction in register <code>ba</code> and saves the contents of the <code>status</code> register in <code>bstatus</code>. Disables interrupts, then transfers execution to the break handler.</p> <p>The 5-bit immediate field <code>imm5</code> is ignored by the processor, but it can be used by the debugger.</p> <p><code>break</code> with no argument is the same as <code>break 0</code>.</p>
Usage	<p><code>break</code> is used by debuggers exclusively. Only debuggers should place <code>break</code> in a user program, operating system, or exception handler. The address of the break handler is specified with the Nios_II Processor parameter editor in Platform Designer.</p> <p>Some debuggers support <code>break</code> and <code>break 0</code> instructions in source code. These debuggers treat the <code>break</code> instruction as a normal breakpoint.</p>
Exceptions	Break
Instruction Type	R
Instruction Fields	IMM5 = Type of breakpoint



Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
0					0					0x1e					0x34
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x34					IMM5					0x3a					

8.5.18. bret

Instruction	breakpoint return
Operation	status ← bstatus PC ← ba
Assembler Syntax	bret
Example	bret
Description	Copies the value of bstatus to the status register, then transfers execution to the address in ba.
Usage	bret is used by debuggers exclusively and should not appear in user programs, operating systems, or exception handlers.
Exceptions	Misaligned destination address Supervisor-only instruction
Instruction Type	R
Instruction Fields	None

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
0x1e					0					0x1e					0x09
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x09					0					0x3a					

8.5.19. call

Instruction	call subroutine
Operation	ra ← PC + 4 PC ← (PC _{31..28} : IMM26 × 4)
Assembler Syntax	call label
Example	call write_char
Description	Saves the address of the next instruction in register ra, and transfers execution to the instruction at address (PC _{31..28} : IMM26 × 4).
Usage	call can transfer execution anywhere within the 256-MB range determined by PC _{31..28} . The Nios II GNU linker does not automatically handle cases in which the address is out of this range.

continued...



Exceptions	None
Instruction Type	J
Instruction Fields	IMM26 = 26-bit unsigned immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
IMM26															
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM26										0					

8.5.20. callr

Instruction	call subroutine in register
Operation	ra ← PC + 4 PC ← rA
Assembler Syntax	callr rA
Example	callr r6
Description	Saves the address of the next instruction in the return address register, and transfers execution to the address contained in register rA.
Usage	callr is used to dereference C-language function pointers.
Exceptions	Misaligned destination address
Instruction Type	R
Instruction Fields	A = Register index of operand rA

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A				0				0x1f				0x1d			
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x1d				0				0x3a							

8.5.21. cmpeq

Instruction	compare equal
Operation	if (rA == rB) then rC ← 1 else rC ← 0
Assembler Syntax	cmpeq rC, rA, rB
Example	cmpeq r6, r7, r8
Description	If rA == rB, then stores 1 to rC; otherwise, stores 0 to rC.
Usage	cmpeq performs the == operation of the C programming language. Also, cmpeq can be used to implement the C logical negation operator "!".

continued...



Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x20
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x20					0					0x3a					

8.5.22. cmpeqi

Instruction	compare equal immediate
Operation	if (rA == σ (IMM16)) then rB ← 1 else rB ← 0
Assembler Syntax	cmpeqi rB, rA, IMM16
Example	cmpeqi r6, r7, 100
Description	Sign-extends the 16-bit immediate value IMM16 to 32 bits and compares it to the value of rA. If rA == σ (IMM16), cmpeqi stores 1 to rB; otherwise stores 0 to rB.
Usage	cmpeqi performs the == operation of the C programming language.
Exceptions	None
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x20					

8.5.23. cmpge

Instruction	compare greater than or equal signed
Operation	if ((signed) rA >= (signed) rB) then rC ← 1

continued...




Assembler Syntax	else rC ← 0
Example	cmpge rC, rA, rB
Description	cmpge r6, r7, r8
Usage	If rA ≥ rB, then stores 1 to rC; otherwise stores 0 to rC.
Exceptions	cmpge performs the signed ≥ operation of the C programming language.
Instruction Type	None
Instruction Fields	R
	A = Register index of operand rA
	B = Register index of operand rB
	C = Register index of operand rC

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x08
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x08					0					0x3a					

8.5.24. cmpgei

Instruction	compare greater than or equal signed immediate
Operation	if ((signed) rA ≥ (signed) σ(IMM16)) then rB ← 1 else rB ← 0
Assembler Syntax	cmpgei rB, rA, IMM16
Example	cmpgei r6, r7, 100
Description	Sign-extends the 16-bit immediate value IMM16 to 32 bits and compares it to the value of rA. If rA ≥ σ(IMM16), then cmpgei stores 1 to rB; otherwise stores 0 to rB.
Usage	cmpgei performs the signed ≥ operation of the C programming language.
Exceptions	None
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x08					



8.5.25. cmpgeu

Instruction	compare greater than or equal unsigned
Operation	if ((unsigned) rA >= (unsigned) rB) then rC ← 1 else rC ← 0
Assembler Syntax	cmpgeu rC, rA, rB
Example	cmpgeu r6, r7, r8
Description	If rA >= rB, then stores 1 to rC; otherwise stores 0 to rC.
Usage	cmpgeu performs the unsigned >= operation of the C programming language.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x28
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x28					0					0x3a					

8.5.26. cmpgeui

Instruction	compare greater than or equal unsigned immediate
Operation	if ((unsigned) rA >= (unsigned) (0x0000 : IMM16)) then rB ← 1 else rB ← 0
Assembler Syntax	cmpgeui rB, rA, IMM16
Example	cmpgeui r6, r7, 100
Description	Zero-extends the 16-bit immediate value IMM16 to 32 bits and compares it to the value of rA. If rA >= (0x0000 : IMM16), then cmpgeui stores 1 to rB; otherwise stores 0 to rB.
Usage	cmpgeui performs the unsigned >= operation of the C programming language.
Exceptions	None
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit unsigned immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x28					

8.5.27. cmpgt

Instruction	compare greater than signed
Operation	if ((signed) rA > (signed) rB) then rC ← 1 else rC ← 0
Assembler Syntax	cmpgt rC, rA, rB
Example	cmpgt r6, r7, r8
Description	If rA > rB, then stores 1 to rC; otherwise stores 0 to rC.
Usage	cmpgt performs the signed > operation of the C programming language.
Pseudo-instruction	cmpgt is implemented with the <code>cmplt</code> instruction by swapping its rA and rB operands.

8.5.28. cmpgti

Instruction	compare greater than signed immediate
Operation	if ((signed) rA > (signed) IMMED) then rB ← 1 else rB ← 0
Assembler Syntax	cmpgti rB, rA, IMMED
Example	cmpgti r6, r7, 100
Description	Sign-extends the immediate value IMMED to 32 bits and compares it to the value of rA. If rA > σ(IMMED), then <code>cmpgti</code> stores 1 to rB; otherwise stores 0 to rB.
Usage	<code>cmpgti</code> performs the signed > operation of the C programming language. The maximum allowed value of IMMED is 32766. The minimum allowed value is -32769.
Pseudo-instruction	<code>cmpgti</code> is implemented using a <code>cmpgei</code> instruction with an IMM16 immediate value of IMMED + 1.

8.5.29. cmpgtu

Instruction	compare greater than unsigned
Operation	if ((unsigned) rA > (unsigned) rB) then rC ← 1 else rC ← 0
Assembler Syntax	cmpgtu rC, rA, rB

continued...



Example	<code>cmpgtu r6, r7, r8</code>
Description	If $rA > rB$, then stores 1 to rC ; otherwise stores 0 to rC .
Usage	<code>cmpgtu</code> performs the unsigned $>$ operation of the C programming language.
Pseudo-instruction	<code>cmpgtu</code> is implemented with the <code>cmpltu</code> instruction by swapping its rA and rB operands.

8.5.30. `cmpgtui`

Instruction	compare greater than unsigned immediate
Operation	if $((\text{unsigned}) rA > (\text{unsigned}) \text{IMMED})$ then $rB \leftarrow 1$ else $rB \leftarrow 0$
Assembler Syntax	<code>cmpgtui rB, rA, IMMED</code>
Example	<code>cmpgtui r6, r7, 100</code>
Description	Zero-extends the immediate value <code>IMMED</code> to 32 bits and compares it to the value of rA . If $rA > \text{IMMED}$, then <code>cmpgtui</code> stores 1 to rB ; otherwise stores 0 to rB .
Usage	<code>cmpgtui</code> performs the unsigned $>$ operation of the C programming language. The maximum allowed value of <code>IMMED</code> is 65534. The minimum allowed value is 0.
Pseudo-instruction	<code>cmpgtui</code> is implemented using a <code>cmpgeui</code> instruction with an <code>IMM16</code> immediate value of <code>IMMED + 1</code> .

8.5.31. `cmple`

Instruction	compare less than or equal signed
Operation	if $((\text{signed}) rA \leq (\text{signed}) rB)$ then $rC \leftarrow 1$ else $rC \leftarrow 0$
Assembler Syntax	<code>cmple rC, rA, rB</code>
Example	<code>cmple r6, r7, r8</code>
Description	If $rA \leq rB$, then stores 1 to rC ; otherwise stores 0 to rC .
Usage	<code>cmple</code> performs the signed \leq operation of the C programming language.
Pseudo-instruction	<code>cmple</code> is implemented with the <code>cmple</code> instruction by swapping its rA and rB operands.

8.5.32. `cmplei`

Instruction	compare less than or equal signed immediate
Operation	if $((\text{signed}) rA < (\text{signed}) \text{IMMED})$ then $rB \leftarrow 1$ else $rB \leftarrow 0$
Assembler Syntax	<code>cmplei rB, rA, IMMED</code>

continued...



Example	<code>cmplei r6, r7, 100</code>
Description	Sign-extends the immediate value IMMED to 32 bits and compares it to the value of rA. If $rA \leq \sigma(\text{IMMED})$, then <code>cmplei</code> stores 1 to rB; otherwise stores 0 to rB.
Usage	<code>cmplei</code> performs the signed \leq operation of the C programming language. The maximum allowed value of IMMED is 32766. The minimum allowed value is -32769.
Pseudo-instruction	<code>cmplei</code> is implemented using a <code>cmplti</code> instruction with an IMM16 immediate value of IMMED + 1.

8.5.33. cmpleu

Instruction	compare less than or equal unsigned
Operation	if $((\text{unsigned}) rA < (\text{unsigned}) rB)$ then $rC \leftarrow 1$ else $rC \leftarrow 0$
Assembler Syntax	<code>cmpleu rC, rA, rB</code>
Example	<code>cmpleu r6, r7, r8</code>
Description	If $rA \leq rB$, then stores 1 to rC; otherwise stores 0 to rC.
Usage	<code>cmpleu</code> performs the unsigned \leq operation of the C programming language.
Pseudo-instruction	<code>cmpleu</code> is implemented with the <code>cmpgeu</code> instruction by swapping its rA and rB operands.


8.5.34. cmpleui

Instruction	compare less than or equal unsigned immediate
Operation	if $((\text{unsigned}) rA \leq (\text{unsigned}) \text{IMMED})$ then $rB \leftarrow 1$ else $rB \leftarrow 0$
Assembler Syntax	<code>cmpleui rB, rA, IMMED</code>
Example	<code>cmpleui r6, r7, 100</code>
Description	Zero-extends the immediate value IMMED to 32 bits and compares it to the value of rA. If $rA \leq \text{IMMED}$, then <code>cmpleui</code> stores 1 to rB; otherwise stores 0 to rB.
Usage	<code>cmpleui</code> performs the unsigned \leq operation of the C programming language. The maximum allowed value of IMMED is 65534. The minimum allowed value is 0.
Pseudo-instruction	<code>cmpleui</code> is implemented using a <code>cmpltui</code> instruction with an IMM16 immediate value of IMMED + 1.

8.5.35. cmplt

Instruction	compare less than signed
Operation	if $((\text{signed}) rA < (\text{signed}) rB)$ then $rC \leftarrow 1$ else $rC \leftarrow 0$

continued...



Assembler Syntax	<code>cmplt rC, rA, rB</code>
Example	<code>cmplt r6, r7, r8</code>
Description	If $rA < rB$, then stores 1 to rC; otherwise stores 0 to rC.
Usage	<code>cmplt</code> performs the signed $<$ operation of the C programming language.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x10
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x10					0					0x3a					

8.5.36. `cmplti`

Instruction	compare less than signed immediate
Operation	if ((signed) rA < (signed) σ (IMM16)) then rB ← 1 else rB ← 0
Assembler Syntax	<code>cmplti rB, rA, IMM16</code>
Example	<code>cmplti r6, r7, 100</code>
Description	Sign-extends the 16-bit immediate value IMM16 to 32 bits and compares it to the value of rA. If $rA < \sigma$ (IMM16), then <code>cmplti</code> stores 1 to rB; otherwise stores 0 to rB.
Usage	<code>cmplti</code> performs the signed $<$ operation of the C programming language.
Exceptions	None
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x10					

8.5.37. `cmpltu`

Instruction	compare less than unsigned
Operation	if ((unsigned) rA < (unsigned) rB) then rC ← 1 else rC ← 0
Assembler Syntax	<code>cmpltu rC, rA, rB</code>
Example	<code>cmpltu r6, r7, r8</code>
Description	If rA < rB, then stores 1 to rC; otherwise stores 0 to rC.
Usage	<code>cmpltu</code> performs the unsigned < operation of the C programming language.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x30
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x30					0					0x3a					

8.5.38. `cmpltui`

Instruction	compare less than unsigned immediate
Operation	if ((unsigned) rA < (unsigned) (0x0000 : IMM16)) then rB ← 1 else rB ← 0
Assembler Syntax	<code>cmpltui rB, rA, IMM16</code>
Example	<code>cmpltui r6, r7, 100</code>
Description	Zero-extends the 16-bit immediate value IMM16 to 32 bits and compares it to the value of rA. If rA < (0x0000 : IMM16), then <code>cmpltui</code> stores 1 to rB; otherwise stores 0 to rB.
Usage	<code>cmpltui</code> performs the unsigned < operation of the C programming language.
Exceptions	None
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit unsigned immediate value



Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x30					

8.5.39. cmpne

Instruction	compare not equal
Operation	if (rA != rB) then rC ← 1 else rC ← 0
Assembler Syntax	cmpne rC, rA, rB
Example	cmpne r6, r7, r8
Description	If rA != rB, then stores 1 to rC; otherwise stores 0 to rC.
Usage	cmpne performs the != operation of the C programming language.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC

Bit Fields																
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	
A					B					C						0x18
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
0x18					0					0x3a						

8.5.40. cmpnei

Instruction	compare not equal immediate
Operation	if (rA != σ (IMM16)) then rB ← 1 else rB ← 0
Assembler Syntax	cmpnei rB, rA, IMM16
Example	cmpnei r6, r7, 100
Description	Sign-extends the 16-bit immediate value IMM16 to 32 bits and compares it to the value of rA. If rA != σ (IMM16), then cmpnei stores 1 to rB; otherwise stores 0 to rB.
Usage	cmpnei performs the != operation of the C programming language.

continued...



Exceptions	None
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x18					

8.5.41. custom

Instruction	custom instruction
Operation	if writerc == 1 then rC ← f _N (rA, rB, A, B, C) else Ø ← f _N (rA, rB, A, B, C)
Assembler Syntax	custom N, xC, xA, xB Where xA means either general purpose register rA, or custom register cA.
Example	custom 0, c6, r7, r8
Description	The custom opcode provides access to up to 256 custom instructions allowed by the Nios II architecture. The function implemented by a custom instruction is user-defined and is specified with the Nios_II Processor parameter editor in Platform Designer. The 8-bit immediate N field specifies which custom instruction to use. Custom instructions can use up to two parameters, xA and xB, and can optionally write the result to a register xC.
Usage	To access a custom register inside the custom instruction logic, clear the bit readra, readrb, or writerc that corresponds to the register field. In assembler syntax, the notation cN refers to register N in the custom register file and causes the assembler to clear the c bit of the opcode. For example, custom 0, c3, r5, r0 performs custom instruction 0, operating on general-purpose registers r5 and r0, and stores the result in custom register 3.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand A B = Register index of operand B C = Register index of operand C readra = 1 if instruction uses rA, 0 otherwise readrb = 1 if instruction uses rB, 0 otherwise writerc = 1 if instruction provides result for rC, 0 otherwise N = 8-bit number that selects instruction



Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					readr a
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
readr b	writer c	N								0x32					

8.5.42. div

Instruction	divide
Operation	$rC \leftarrow rA \div rB$
Assembler Syntax	<code>div rC, rA, rB</code>
Example	<code>div r6, r7, r8</code>
Description	<p>Treating rA and rB as signed integers, this instruction divides rA by rB and then stores the integer portion of the resulting quotient to rC. After attempted division by zero, the value of rC is undefined. There is no divide-by-zero exception. After dividing -2147483648 by -1, the value of rC is undefined (the number +2147483648 is not representable in 32 bits). There is no overflow exception.</p> <p>Nios II processors that do not implement the <code>div</code> instruction cause an unimplemented instruction exception.</p>
Usage	<p>Remainder of Division:</p> <p>If the result of the division is defined, then the remainder can be computed in rD using the following instruction sequence:</p> <pre>div rC, rA, rB mul rD, rC, rB sub rD, rA, rD # The original div operation # rD = remainder</pre>
Exceptions	<p>Division error</p> <p>Unimplemented instruction</p>
Instruction Type	R
Instruction Fields	<p>A = Register index of operand rA</p> <p>B = Register index of operand rB</p> <p>C = Register index of operand rC</p>

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x18
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x25					0					0x3a					

8.5.43. divu

Instruction	divide unsigned
Operation	$rC \leftarrow rA \div rB$
Assembler Syntax	<code>divu rC, rA, rB</code>
Example	<code>divu r6, r7, r8</code>
Description	<p>Treating rA and rB as unsigned integers, this instruction divides rA by rB and then stores the integer portion of the resulting quotient to rC. After attempted division by zero, the value of rC is undefined. There is no divide-by-zero exception.</p> <p>Nios II processors that do not implement the <code>divu</code> instruction cause an unimplemented instruction exception.</p>
Usage	<p>Remainder of Division:</p> <p>If the result of the division is defined, then the remainder can be computed in rD using the following instruction sequence:</p> <pre>divu rC, rA, rB mul rD, rC, rB sub rD, rA, rD # The original divu operation # rD = remainder</pre>
Exceptions	<p>Division error Unimplemented instruction</p>
Instruction Type	R
Instruction Fields	<p>A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC</p>

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x24
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x24					0					0x3a					

8.5.44. eret

Instruction	exception return
Operation	$status \leftarrow estatus$ $PC \leftarrow ea$
Assembler Syntax	<code>eret</code>
Example	<code>eret</code>
Description	Copies the value of <code>estatus</code> into the <code>status</code> register, and transfers execution to the address in <code>ea</code> .

continued...



Usage	Use <code>eret</code> to return from traps, external interrupts, and other exception handling routines. Note that before returning from hardware interrupt exceptions, the exception handler must adjust the <code>ea</code> register.
Exceptions	Misaligned destination address Supervisor-only instruction
Instruction Type	R
Instruction Fields	None

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
0x1d					0x1e					C					0x01
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x01					0					0x3a					

8.5.45. `flushd`

Instruction	<code>flushd</code> flush data cache line
Operation	Flushes the data cache line associated with address <code>rA + σ(IMM16)</code> .
Assembler Syntax	<code>flushd IMM16(rA)</code>
Example	<code>flushd -100(r6)</code>
Description	<p>If the Nios II processor implements a direct mapped data cache, <code>flushd</code> writes the data cache line that is mapped to the specified address back to memory if the line is dirty, and then clears the data cache line. Unlike <code>flushda</code>, <code>flushd</code> writes the dirty data back to memory even when the addressed data is not currently in the cache. This process comprises the following steps:</p> <ul style="list-style-type: none"> • Compute the effective address specified by the sum of <code>rA</code> and the signed 16-bit immediate value. • Identify the data cache line associated with the computed effective address. Each data cache effective address comprises a <code>tag</code> field and a <code>line</code> field. When identifying the data cache line, <code>flushd</code> ignores the <code>tag</code> field and only uses the <code>line</code> field to select the data cache line to clear. • Skip comparing the cache line tag with the effective address to determine if the addressed data is currently cached. Because <code>flushd</code> ignores the cache line tag, <code>flushd</code> flushes the cache line regardless of whether the specified data location is currently cached. • If the data cache line is dirty, write the line back to memory. A cache line is dirty when one or more words of the cache line have been modified by the processor, but are not yet written to memory. • Clear the valid bit for the line. <p>If the Nios II processor core does not have a data cache, the <code>flushd</code> instruction performs no operation.</p>
Usage	Use <code>flushd</code> to write dirty lines back to memory even if the addressed memory location is not in the cache, and then flush the cache line. By contrast, refer to “ <code>flushda</code> flush data

continued...



cache address”, “initd initialize data cache line”, and “initda initialize data cache address” for other cache-clearing options.

For more information on data cache, refer to the *Cache and Tightly Coupled Memory* chapter of the *Nios II Software Developer’s Handbook*.

Exceptions

None

Instruction Type

I

Instruction Fields

A = Register index of operand rA

IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					0					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x3b					


Related Information

- [Cache and Tightly-Coupled Memory](#)
- [flushda](#) on page 196
- [initda](#) on page 200
- [initd](#) on page 199

8.5.46. flushda

Instruction	flush data cache address
Operation	Flushes the data cache line currently caching address rA + σ(IMM16)
Assembler Syntax	flushda IMM16(rA)
Example	flushda -100(r6)
Description	If the Nios II processor implements a direct mapped data cache, flushda writes the data cache line that is mapped to the specified address back to memory if the line is dirty, and then clears the data cache line. Unlike flushd,

continued...



`flushda` writes the dirty data back to memory only when the addressed data is currently in the cache. This process comprises the following steps:

- Compute the effective address specified by the sum of `rA` and the signed 16-bit immediate value.
- Identify the data cache line associated with the computed effective address. Each data cache effective address comprises a `tag` field and a `line` field. When identifying the line, `flushda` uses both the `tag` field and the `line` field.
- Compare the cache line tag with the effective address to determine if the addressed data is currently cached. If the `tag` fields do not match, the effective address is not currently cached, so the instruction does nothing.
- If the data cache line is dirty and the `tag` fields match, write the dirty cache line back to memory. A cache line is dirty when one or more words of the cache line have been modified by the processor, but are not yet written to memory.
- Clear the valid bit for the line.
If the Nios II processor core does not have a data cache, the `flushda` instruction performs no operation.

Usage

Use `flushda` to write dirty lines back to memory only if the addressed memory location is currently in the cache, and then flush the cache line. By contrast, refer to “`flushd` flush data cache line”, “`initd` initialize data cache line”, and “`initda` initialize data cache address” for other cache-clearing options.

For more information on the Nios II data cache, refer to the *Cache and Tightly Coupled Memory* chapter of the *Nios II Software Developer’s Handbook*.

Exceptions

Supervisor-only data address
Fast TLB miss (data)
Double TLB miss (data)
MPU region violation (data)

Instruction Type

I


Instruction Fields

`A` = Register index of operand `rA`
`IMM16` = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					0					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x1b					

Related Information

- [Cache and Tightly-Coupled Memory](#)
- [initda](#) on page 200
- [initd](#) on page 199
- [flushd](#) on page 195



8.5.47. flushi

Instruction	flush instruction cache line
Operation	Flushes the instruction cache line associated with address rA.
Assembler Syntax	<code>flushi rA</code>
Example	<code>flushi r6</code>
Description	<p>Ignoring the tag, <code>flushi</code> identifies the instruction cache line associated with the byte address in rA, and invalidates that line.</p> <p>If the Nios II processor core does not have an instruction cache, the <code>flushi</code> instruction performs no operation.</p> <p>For more information about the data cache, refer to the <i>Cache and Tightly Coupled Memory</i> chapter of the <i>Nios II Software Developer's Handbook</i>.</p>
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					0					0					0x0c
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x0c					0					0x3a					

Related Information

[Cache and Tightly-Coupled Memory](#)

8.5.48. flushp

Instruction	flush pipeline
Operation	Flushes the processor pipeline of any prefetched instructions.
Assembler Syntax	<code>flushp</code>
Example	<code>flushp</code>
Description	Ensures that any instructions prefetched after the <code>flushp</code> instruction are removed from the pipeline.
Usage	Use <code>flushp</code> before transferring control to newly updated instruction memory.
Exceptions	None
Instruction Type	R
Instruction Fields	None



Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					0					0					0x04
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x04					0					0x3a					

8.5.49. `initd`

Instruction	initialize data cache line
Operation	Initializes the data cache line associated with address $rA + \sigma(\text{IMM16})$.
Assembler Syntax	<code>initd IMM16(rA)</code>
Example	<code>initd 0(r6)</code>
Description	<p>If the Nios II processor implements a direct mapped data cache, <code>initd</code> clears the data cache line without checking for (or writing) a dirty data cache line that is mapped to the specified address back to memory. Unlike <code>initda</code>, <code>initd</code> clears the cache line regardless of whether the addressed data is currently cached. This process comprises the following steps:</p> <ul style="list-style-type: none"> • Compute the effective address specified by the sum of <code>rA</code> and the signed 16-bit immediate value. • Identify the data cache line associated with the computed effective address. Each data cache effective address comprises a <code>tag</code> field and a <code>line</code> field. When identifying the line, <code>initd</code> ignores the <code>tag</code> field and only uses the <code>line</code> field to select the data cache line to clear. • Skip comparing the cache line tag with the effective address to determine if the addressed data is currently cached. Because <code>initd</code> ignores the cache line tag, <code>initd</code> flushes the cache line regardless of whether the specified data location is currently cached. • Skip checking if the data cache line is dirty. Because <code>initd</code> skips the dirty cache line check, data that has been modified by the processor, but not yet written to memory is lost. • Clear the valid bit for the line. <p>If the Nios II processor core does not have a data cache, the <code>initd</code> instruction performs no operation.</p>
Usage	<p>Use <code>initd</code> after processor reset and before accessing data memory to initialize the processor's data cache. Use <code>initd</code> with caution because it does not write back dirty data. By contrast, refer to "flushd flush data cache line", "flushda flush data cache address", and "initda initialize data cache address" for other cache-clearing options. Intel recommends using <code>initd</code> only when the processor comes out of reset.</p> <p>For more information on data cache, refer to the <i>Cache and Tightly Coupled Memory</i> chapter of the <i>Nios II Software Developer's Handbook</i>.</p>
Exceptions	Supervisor-only instruction
Instruction Type	I
Instruction Fields	<p>A = Register index of operand <code>rA</code> IMM16 = 16-bit signed immediate value</p>

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					0					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x33					

Related Information

- [Cache and Tightly-Coupled Memory](#)
- [flushda](#) on page 196
- [initda](#) on page 200
- [flushd](#) on page 195

8.5.50. initda

Instruction	initialize data cache address
Operation	Initializes the data cache line currently caching address $rA + \sigma(\text{IMM16})$
Assembler Syntax	<code>initda IMM16(rA)</code>
Example	<code>initda -100(r6)</code>
Description	<p>If the Nios II processor implements a direct mapped data cache, <code>initda</code> clears the data cache line without checking for (or writing) a dirty data cache line that is mapped to the specified address back to memory. Unlike <code>initd</code>, <code>initda</code> clears the cache line only when the addressed data is currently cached. This process comprises the following steps:</p> <ul style="list-style-type: none"> • Compute the effective address specified by the sum of <code>rA</code> and the signed 16-bit immediate value. • Identify the data cache line associated with the computed effective address. Each data cache effective address comprises a <code>tag</code> field and a <code>line</code> field. When identifying the line, <code>initda</code> uses both the <code>tag</code> field and the <code>line</code> field. • Compare the cache line tag with the effective address to determine if the addressed data is currently cached. If the <code>tag</code> fields do not match, the effective address is not currently cached, so the instruction does nothing. • Skip checking if the data cache line is dirty. Because <code>initd</code> skips the dirty cache line check, data that has been modified by the processor, but not yet written to memory is lost. • Clear the valid bit for the line. <p>If the Nios II processor core does not have a data cache, the <code>initda</code> instruction performs no operation.</p>
Usage	Use <code>initda</code> to skip writing dirty lines back to memory and to flush the cache line only if the addressed memory location is currently in the cache. By contrast, refer to “ <code>flushd</code> flush data cache line”, “ <code>flushda</code> flush data cache address”, and “ <code>initd</code> initialize data cache line” on page 8–55 for other cache-clearing options. Use <code>initda</code> with caution because it does not write back dirty data.

continued...



For more information on the Nios II data cache, refer to the *Cache and Tightly Coupled Memory* chapter of the *Nios II Software Developer's Handbook*.

Exceptions	Supervisor-only data address Fast TLB miss (data) Double TLB miss (data) MPU region violation (data) Unimplemented instruction
Instruction Type	I
Instruction Fields	A = Register index of operand rA IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					0					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x13					

Related Information

- [Cache and Tightly-Coupled Memory](#)
- [flushda](#) on page 196
- [initd](#) on page 199
- [flushd](#) on page 195

8.5.51. `initd`

Instruction	<code>initd</code> initialize instruction cache line
Operation	Initializes the instruction cache line associated with address rA.
Assembler Syntax	<code>initd rA</code>
Example	<code>initd r6</code>
Description	Ignoring the tag, <code>initd</code> identifies the instruction cache line associated with the byte address in <code>ra</code> , and <code>initd</code> invalidates that line. If the Nios II processor core does not have an instruction cache, the <code>initd</code> instruction performs no operation.
Usage	This instruction is used to initialize the processor's instruction cache. Immediately after processor reset, use <code>initd</code> to invalidate each line of the instruction cache. For more information on instruction cache, refer to the <i>Cache and Tightly Coupled Memory</i> chapter of the <i>Nios II Software Developer's Handbook</i> .
Exceptions	Supervisor-only instruction
Instruction Type	R
Instruction Fields	A = Register index of operand rA

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					0					0					0x29
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x29					0					0x3a					

Related Information

Cache and Tightly-Coupled Memory

8.5.52. jmp

Instruction	computed jump
Operation	PC ← rA
Assembler Syntax	jmp rA
Example	jmp r12
Description	Transfers execution to the address contained in register rA.
Usage	It is illegal to jump to the address contained in register r31. To return from subroutines called by call or callr, use ret instead of jmp.
Exceptions	Misaligned destination address
Instruction Type	R
Instruction Fields	A = Register index of operand rA

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					0					0					0x0d
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x0d					0					0x3a					

8.5.53. jmp_i

Instruction	jump immediate
Operation	PC ← (PC _{31..28} : IMM26 × 4)
Assembler Syntax	jmp_i label
Example	jmp_i write_char
Description	Transfers execution to the instruction at address (PC _{31..28} : IMM26 × 4).
Usage	jmp_i is a low-overhead local jump. jmp_i can transfer execution anywhere within the 256-MB range determined by PC _{31..28} . The Nios II GNU linker does not automatically handle cases in which the address is out of this range.
Exceptions	None
Instruction Type	J
Instruction Fields	IMM26 = 26-bit unsigned immediate value



Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
IMM26															
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM26										0x01					

8.5.54. ldb / ldbio

Instruction	load byte from memory or I/O peripheral
Operation	$rB \leftarrow \sigma(\text{Mem8}[rA + \sigma(\text{IMM16})])$
Assembler Syntax	<pre>ldb rB, byte_offset(rA) ldbio rB, byte_offset(rA)</pre>
Example	<pre>ldb r6, 100(r5)</pre>
Description	Computes the effective byte address specified by the sum of rA and the instruction's signed 16-bit immediate value. Loads register rB with the desired memory byte, sign extending the 8-bit value to 32 bits. In Nios II processor cores with a data cache, this instruction may retrieve the desired data from the cache instead of from memory.
Usage	Use the <code>ldbio</code> instruction for peripheral I/O. In processors with a data cache, <code>ldbio</code> bypasses the cache and is guaranteed to generate an Avalon-MM data transfer. In processors without a data cache, <code>ldbio</code> acts like <code>ldb</code> . For more information on data cache, refer to the <i>Cache and Tightly Coupled Memory</i> chapter of the <i>Nios II Software Developer's Handbook</i> .
Exceptions	Supervisor-only data address Misaligned data address TLB permission violation (read) Fast TLB miss (data) Double TLB miss (data) MPU region violation (data)
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Table 96. ldb

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x07					

Table 97. ldbio

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x27					

Related Information

Cache and Tightly-Coupled Memory

8.5.55. ldbu / ldbuio

Instruction	load unsigned byte from memory or I/O peripheral
Operation	$rB \leftarrow 0x000000 : Mem8[rA + \sigma(IMM16)]$
Assembler Syntax	ldbu rB, byte_offset(rA) ldbuio rB, byte_offset(rA)
Example	ldbu r6, 100(r5)
Description	Computes the effective byte address specified by the sum of rA and the instruction's signed 16-bit immediate value. Loads register rB with the desired memory byte, zero extending the 8-bit value to 32 bits.
Usage	In processors with a data cache, this instruction may retrieve the desired data from the cache instead of from memory. Use the ldbuio instruction for peripheral I/O. In processors with a data cache, ldbuio bypasses the cache and is guaranteed to generate an Avalon-MM data transfer. In processors without a data cache, ldbuio acts like ldbu. For more information on data cache, refer to the <i>Cache and Tightly Coupled Memory</i> chapter of the <i>Nios II Software Developer's Handbook</i> .
Exceptions	Supervisor-only data address Misaligned data address TLB permission violation (read) Fast TLB miss (data) Double TLB miss (data) MPU region violation (data)
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Table 98. ldbu

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x03					



Table 99. ldbuio

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x23					

Related Information

Cache and Tightly-Coupled Memory

8.5.56. ldh / ldhio

Instruction	load halfword from memory or I/O peripheral
Operation	$rB \leftarrow \sigma(\text{Mem16}[rA + \sigma(\text{IMM16})])$
Assembler Syntax	ldh rB, byte_offset(rA) ldhio rB, byte_offset(rA)
Example	ldh r6, 100(r5)
Description	Computes the effective byte address specified by the sum of rA and the instruction's signed 16-bit immediate value. Loads register rB with the memory halfword located at the effective byte address, sign extending the 16-bit value to 32 bits. The effective byte address must be halfword aligned. If the byte address is not a multiple of 2, the operation is undefined.
Usage	In processors with a data cache, this instruction may retrieve the desired data from the cache instead of from memory. Use the ldhio instruction for peripheral I/O. In processors with a data cache, ldhio bypasses the cache and is guaranteed to generate an Avalon-MM data transfer. In processors without a data cache, ldhio acts like ldh. For more information on data cache, refer to the <i>Cache and Tightly Coupled Memory</i> chapter of the <i>Nios II Software Developer's Handbook</i> .
Exceptions	Supervisor-only data address Misaligned data address TLB permission violation (read) Fast TLB miss (data) Double TLB miss (data) MPU region violation (data)
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Table 100. Idh

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x0f					

Table 101. Idhio

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x2f					

Related Information

Cache and Tightly-Coupled Memory

8.5.57. Idhu / Idhuio

Instruction	load unsigned halfword from memory or I/O peripheral
Operation	$rB \leftarrow 0x0000 : \text{Mem16}[rA + \sigma(\text{IMM16})]$
Assembler Syntax	<code>ldhu rB, byte_offset(rA)</code> <code>ldhuio rB, byte_offset(rA)</code>
Example	<code>ldhu r6, 100(r5)</code>
Description	Computes the effective byte address specified by the sum of rA and the instruction's signed 16-bit immediate value. Loads register rB with the memory halfword located at the effective byte address, zero extending the 16-bit value to 32 bits. The effective byte address must be halfword aligned. If the byte address is not a multiple of 2, the operation is undefined.
Usage	In processors with a data cache, this instruction may retrieve the desired data from the cache instead of from memory. Use the <code>ldhuio</code> instruction for peripheral I/O. In processors with a data cache, <code>ldhuio</code> bypasses the cache and is guaranteed to generate an Avalon-MM data transfer. In processors without a data cache, <code>ldhuio</code> acts like <code>ldhu</code> . For more information on data cache, refer to the <i>Cache and Tightly Coupled Memory</i> chapter of the <i>Nios II Software Developer's Handbook</i> .
Exceptions	Supervisor-only data address Misaligned data address TLB permission violation (read) Fast TLB miss (data) Double TLB miss (data) MPU region violation (data)
Instruction Type	I
Instruction Fields	A = Register index of operand rA

continued...



B = Register index of operand rB

IMM16 = 16-bit signed immediate value

Table 102. Idhu

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x0b					

Table 103. Idhuio

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x2b					

Related Information

Cache and Tightly-Coupled Memory

8.5.58. ldw / ldwio

Instruction	load 32-bit word from memory or I/O peripheral
Operation	$rB \leftarrow \text{Mem32}[rA + \sigma(\text{IMM16})]$
Assembler Syntax	<code>ldw rB, byte_offset(rA)</code> <code>ldwio rB, byte_offset(rA)</code>
Example	<code>ldw r6, 100(r5)</code>
Description	Computes the effective byte address specified by the sum of rA and the instruction's signed 16-bit immediate value. Loads register rB with the memory word located at the effective byte address. The effective byte address must be word aligned. If the byte address is not a multiple of 4, the operation is undefined.
Usage	In processors with a data cache, this instruction may retrieve the desired data from the cache instead of from memory. Use the <code>ldwio</code> instruction for peripheral I/O. In processors with a data cache, <code>ldwio</code> bypasses the cache and memory. Use the <code>ldwio</code> instruction for peripheral I/O. In processors with a data cache, <code>ldwio</code> bypasses the cache and is guaranteed to generate an Avalon-MM data transfer. In processors without a data cache, <code>ldwio</code> acts like <code>ldw</code> . For more information on data cache, refer to the <i>Cache and Tightly Coupled Memory</i> chapter of the <i>Nios II Software Developer's Handbook</i> .
Exceptions	Supervisor-only data address Misaligned data address TLB permission violation (read) Fast TLB miss (data) Double TLB miss (data)

continued...



Instruction Type	MPU region violation (data)
Instruction Fields	I A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Table 104. Idw

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x17					

Table 105. Idwio

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x37					

Related Information

Cache and Tightly-Coupled Memory

8.5.59. mov

Instruction	move register to register
Operation	rC ← rA
Assembler Syntax	mov rC, rA
Example	mov r6, r7
Description	Moves the contents of rA to rC.
Pseudo-instruction	mov is implemented as add rC, rA, r0.

8.5.60. movhi

Instruction	move immediate into high halfword
Operation	rB ← (IMMED : 0x0000)
Assembler Syntax	movhi rB, IMMED
Example	movhi r6, 0x8000
Description	Writes the immediate value IMMED into the high halfword of rB, and clears the lower halfword of rB to 0x0000.
Usage	The maximum allowed value of IMMED is 65535. The minimum allowed value is 0. To load a 32-bit constant into a register, first load the upper 16 bits using a movhi pseudo-

continued...



Pseudo-instruction

instruction. The `%hi()` macro can be used to extract the upper 16 bits of a constant or a label. Then, load the lower 16 bits with an `ori` instruction. The `%lo()` macro can be used to extract the lower 16 bits of a constant or label as shown in the following code:

```
movhi rB, %hi(value)
ori rB, rB, %lo(value)
```

An alternative method to load a 32-bit constant into a register uses the `%hiadj()` macro and the `addi` instruction as shown in the following code:

```
movhi rB, %hiadj(value)
addi rB, rB, %lo(value)
```

`movhi` is implemented as `orhi rB, r0, IMMED`.

8.5.61. movi

Instruction

move signed immediate into word

Operation

$rB \leftarrow \sigma(\text{IMMED})$

Assembler Syntax

```
movi rB, IMMED
```

Example

```
movi r6, -30
```

Description

Sign-extends the immediate value IMMED to 32 bits and writes it to rB.

Usage

The maximum allowed value of IMMED is 32767. The minimum allowed value is -32768. To load a 32-bit constant into a register, refer to the `movhi` instruction.

Pseudo-instruction

`movi` is implemented as `addi rB, r0, IMMED`.

8.5.62. movia

Instruction

move immediate address into word

Operation

$rB \leftarrow \text{label}$

Assembler Syntax

```
movia rB, label
```

Example

```
movia r6, function_address
```

Description

Writes the address of label to rB.

Pseudo-instruction

`movia` is implemented as:

```
orhi rB, r0, %hiadj(label)
addi rB, rB, %lo(label)
```

8.5.63. movui

Instruction

move unsigned immediate into word

Operation

$rB \leftarrow (0x0000 : \text{IMMED})$

Assembler Syntax

```
movui rB, IMMED
```

Example

```
movui r6, 100
```

continued...



Description	Zero-extends the immediate value IMMED to 32 bits and writes it to rB.
Usage	The maximum allowed value of IMMED is 65535. The minimum allowed value is 0. To load a 32-bit constant into a register, refer to the <code>movhi</code> instruction.
Pseudo-instruction	<code>movui</code> is implemented as <code>ori rB, r0, IMMED</code> .

8.5.64. mul

Instruction	multiply
Operation	$rC \leftarrow (rA \times rB)_{31..0}$
Assembler Syntax	<code>mul rC, rA, rB</code>
Example	<code>mul r6, r7, r8</code>
Description	<p>Multiplies rA times rB and stores the 32 low-order bits of the product to rC. The result is the same whether the operands are treated as signed or unsigned integers.</p> <p>Nios II processors that do not implement the <code>mul</code> instruction cause an unimplemented instruction exception.</p>
Usage	<p>Carry Detection (unsigned operands):</p> <p>Before or after the multiply operation, the carry out of the MSB of rC can be detected using the following instruction sequence:</p> <pre>mul rC, rA, rB mulxuu rD, rA, rB cmpne rD, rD, r0 # The mul operation (optional) # rD is nonzero if carry occurred # rD is 1 if carry occurred, 0 if not</pre> <p>The <code>mulxuu</code> instruction writes a nonzero value into rD if the multiplication of unsigned numbers generates a carry (unsigned overflow). If a 0/1 result is desired, follow the <code>mulxuu</code> with the <code>cmpne</code> instruction.</p> <p>Overflow Detection (signed operands):</p> <p>After the multiply operation, overflow can be detected using the following instruction sequence:</p> <pre>mul rC, rA, rB cmplt rD, rC, r0 mulxss rE, rA, rB add rD, rD, rE cmpne rD, rD, r0 # The original mul operation # rD is nonzero if overflow # rD is 1 if overflow, 0 if not</pre> <p>The <code>cmplt-mulxss-add</code> instruction sequence writes a nonzero value into rD if the product in rC cannot be represented in 32 bits (signed overflow). If a 0/1 result is desired, follow the instruction sequence with the <code>cmpne</code> instruction.</p>
Exceptions	Unimplemented instruction
Instruction Type	R
Instruction Fields	<p>A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC</p>



Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x27
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x27					0					0x3a					

8.5.65. muli

Instruction	multiply immediate
Operation	$rB \leftarrow (rA \times \sigma(\text{IMM16}))_{31..0}$
Assembler Syntax	<code>muli rB, rA, IMM16</code>
Example	<code>muli r6, r7, -100</code>
Description	<p>Sign-extends the 16-bit immediate value IMM16 to 32 bits and multiplies it by the value of rA. Stores the 32 low-order bits of the product to rB. The result is independent of whether rA is treated as a signed or unsigned number.</p> <p>Nios II processors that do not implement the <code>muli</code> instruction cause an unimplemented instruction exception.</p> <p>Carry Detection and Overflow Detection: For a discussion of carry and overflow detection, refer to the <code>mul</code> instruction.</p>
Exceptions	Unimplemented instruction
Instruction Type	I
Instruction Fields	<p>A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value</p>

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x24					

8.5.66. mulxss

Instruction	multiply extended signed/signed
Operation	$rC \leftarrow ((\text{signed}) rA) \times ((\text{signed}) rB)_{63..32}$
Assembler Syntax	<code>mulxss rC, rA, rB</code>
Example	<code>mulxss r6, r7, r8</code>
Description	<p>Treating rA and rB as signed integers, <code>mulxss</code> multiplies rA times rB, and stores the 32 high-order bits of the product to rC.</p> <p>Nios II processors that do not implement the <code>mulxss</code> instruction cause an unimplemented instruction exception.</p>

continued...



Usage	Use <code>mulxss</code> and <code>mul</code> to compute the full 64-bit product of two 32-bit signed integers. Furthermore, <code>mulxss</code> can be used as part of the calculation of a 128-bit product of two 64-bit signed integers. Given two 64-bit integers, each contained in a pair of 32-bit registers, (<code>S1 : U1</code>) and (<code>S2 : U2</code>), their 128-bit product is $(U1 \times U2) + ((S1 \times U2) \ll 32) + ((U1 \times S2) \ll 32) + ((S1 \times S2) \ll 64)$. The <code>mulxss</code> and <code>mul</code> instructions are used to calculate the 64-bit product $S1 \times S2$.
Exceptions	Unimplemented instruction
Instruction Type	R
Instruction Fields	A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x1f
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x1f					0					0x3a					

8.5.67. `mulxsu`

Instruction	<code>multiply extended signed/unsigned</code>
Operation	$rC \leftarrow ((\text{signed } rA) \times ((\text{unsigned } rB)))_{63..32}$
Assembler Syntax	<code>mulxsu rC, rA, rB</code>
Example	<code>mulxsu r6, r7, r8</code>
Description	Treating <code>rA</code> as a signed integer and <code>rB</code> as an unsigned integer, <code>mulxsu</code> multiplies <code>rA</code> times <code>rB</code> , and stores the 32 high-order bits of the product to <code>rC</code> . Nios II processors that do not implement the <code>mulxsu</code> instruction cause an unimplemented instruction exception.
Usage	<code>mulxsu</code> can be used as part of the calculation of a 128-bit product of two 64-bit signed integers. Given two 64-bit integers, each contained in a pair of 32-bit registers, (<code>S1 : U1</code>) and (<code>S2 : U2</code>), their 128-bit product is: $(U1 \times U2) + ((S1 \times U2) \ll 32) + ((U1 \times S2) \ll 32) + ((S1 \times S2) \ll 64)$. The <code>mulxsu</code> and <code>mul</code> instructions are used to calculate the two 64-bit products $S1 \times U2$ and $U1 \times S2$.
Exceptions	Unimplemented instruction
Instruction Type	R
Instruction Fields	A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC




Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x17
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x17					0					0x3a					

8.5.68. mulxuu

Instruction	multiply extended unsigned/unsigned
Operation	$rC \leftarrow ((\text{unsigned}) rA) \times ((\text{unsigned}) rB))_{63..32}$
Assembler Syntax	<code>mulxuu rC, rA, rB</code>
Example	<code>mulxuu r6, r7, r8</code>
Description	<p>Treating <code>rA</code> and <code>rB</code> as unsigned integers, <code>mulxuu</code> multiplies <code>rA</code> times <code>rB</code> and stores the 32 high-order bits of the product to <code>rC</code>.</p> <p>Nios II processors that do not implement the <code>mulxuu</code> instruction cause an unimplemented instruction exception.</p>
Usage	<p>Use <code>mulxuu</code> and <code>mul</code> to compute the 64-bit product of two 32-bit unsigned integers. Furthermore, <code>mulxuu</code> can be used as part of the calculation of a 128-bit product of two 64-bit signed integers. Given two 64-bit signed integers, each contained in a pair of 32-bit registers, (<code>S1 : U1</code>) and (<code>S2 : U2</code>), their 128-bit product is $(U1 \times U2) + ((S1 \times U2) \ll 32) + ((U1 \times S2) \ll 32) + ((S1 \times S2) \ll 64)$. The <code>mulxuu</code> and <code>mul</code> instructions are used to calculate the 64-bit product $U1 \times U2$.</p> <p><code>mulxuu</code> also can be used as part of the calculation of a 128-bit product of two 64-bit unsigned integers. Given two 64-bit unsigned integers, each contained in a pair of 32-bit registers, (<code>T1 : U1</code>) and (<code>T2 : U2</code>), their 128-bit product is $(U1 \times U2) + ((U1 \times T2) \ll 32) + ((T1 \times U2) \ll 32) + ((T1 \times T2) \ll 64)$. The <code>mulxuu</code> and <code>mul</code> instructions are used to calculate the four 64-bit products $U1 \times U2$, $U1 \times T2$, $T1 \times U2$, and $T1 \times T2$.</p>
Exceptions	Unimplemented instruction
Instruction Type	R
Instruction Fields	<p>A = Register index of operand <code>rA</code></p> <p>B = Register index of operand <code>rB</code></p> <p>C = Register index of operand <code>rC</code></p>

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x07
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x07					0					0x3a					



8.5.69. nextpc

Instruction	get address of following instruction
Operation	$rC \leftarrow PC + 4$
Assembler Syntax	<code>nextpc rC</code>
Example	<code>nextpc r6</code>
Description	Stores the address of the next instruction to register rC.
Usage	A relocatable code fragment can use <code>nextpc</code> to calculate the address of its data segment. <code>nextpc</code> is the only way to access the PC directly.
Exceptions	None
Instruction Type	R
Instruction Fields	C = Register index of operand rC

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
0					0					C					0x1c
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x1c					0					0x3a					

8.5.70. nop

Instruction	no operation
Operation	None
Assembler Syntax	<code>nop</code>
Example	<code>nop</code>
Description	<code>nop</code> does nothing.
Pseudo-instruction	<code>nop</code> is implemented as <code>add r0, r0, r0</code> .

8.5.71. nor

Instruction	bitwise logical nor
Operation	$rC \leftarrow \sim(rA \mid rB)$
Assembler Syntax	<code>nor rC, rA, rB</code>
Example	<code>nor r6, r7, r8</code>
Description	Calculates the bitwise logical NOR of rA and rB and stores the result in rC.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC



3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
A				B				C				0x06				0				0x3a										

8.5.72. or

Instruction	bitwise logical or
Operation	$rC \leftarrow rA rB$
Assembler Syntax	or rC, rA, rB
Example	or r6, r7, r8
Description	Calculates the bitwise logical OR of rA and rB and stores the result in rC.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A				B				C				0x16			
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x16				0				0x3a							

8.5.73. orhi

Instruction	bitwise logical or immediate into high halfword
Operation	$rB \leftarrow rA (IMM16 : 0x0000)$
Assembler Syntax	orhi rB, rA, IMM16
Example	orhi r6, r7, 100
Description	Calculates the bitwise logical OR of rA and (IMM16 : 0x0000) and stores the result in rB.
Exceptions	None
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A				B				IMM16							
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16								0x34							



8.5.74. ori

Instruction	bitwise logical or immediate
Operation	$rB \leftarrow rA \mid (0x0000 : IMM16)$
Assembler Syntax	<code>ori rB, rA, IMM16</code>
Example	<code>ori r6, r7, 100</code>
Description	Calculates the bitwise logical OR of rA and (0x0000 : IMM16) and stores the result in rB.
Exceptions	None
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit unsigned immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x14					

8.5.75. rdctl

Instruction	read from control register
Operation	$rC \leftarrow ctIN$
Assembler Syntax	<code>rdctl rC, ctIN</code>
Example	<code>rdctl r3, ct131</code>
Description	Reads the value contained in control register ctIN and writes it to register rC.
Exceptions	Supervisor-only instruction
Instruction Type	R
Instruction Fields	C = Register index of operand rC N = Control register index of operand ctIN

Bit Fields																
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	
0					0					C						0x26
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
0x26					N					0x3a						



8.5.76. rdprs

Instruction	read from previous register set
Operation	$rB \leftarrow prs.rA + \sigma(IMM16)$
Assembler Syntax	<code>rdprs rB, rA, IMM16</code>
Example	<code>rdprs r6, r7, 0</code>
Description	Sign-extends the 16-bit immediate value IMM16 to 32 bits, and adds it to the value of rA from the previous register set. Places the result in rB in the current register set.
Usage	The previous register set is specified by status.PRS. By default, status.PRS indicates the register set in use before an exception, such as an external interrupt, caused a register set change. To read from an arbitrary register set, software can insert the desired register set number in status.PRS prior to executing rdprs. If shadow register sets are not implemented on the Nios II core, rdprs is an illegal instruction.
Exceptions	Supervisor-only instruction Illegal instruction
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A						B						IMM16			
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x38					

8.5.77. ret

Instruction	return from subroutine
Operation	$PC \leftarrow ra$
Assembler Syntax	<code>ret</code>
Example	<code>ret</code>
Description	Transfers execution to the address in ra.
Usage	Any subroutine called by call or callr must use ret to return.
Exceptions	Misaligned destination address
Instruction Type	R
Instruction Fields	None

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
0x1f					0					0					0x05
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x05					0					0x3a					

8.5.78. rol

Instruction	rotate left
Operation	$rC \leftarrow rA$ rotated left $rB_{4..0}$ bit positions
Assembler Syntax	<code>rol rC, rA, rB</code>
Example	<code>rol r6, r7, r8</code>
Description	Rotates rA left by the number of bits specified in $rB_{4..0}$ and stores the result in rC . The bits that shift out of the register rotate into the least-significant bit positions. Bits 31–5 of rB are ignored.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x03
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x03					0					0x3a					

8.5.79. roli

Instruction	rotate left immediate
Operation	$rC \leftarrow rA$ rotated left IMM5 bit positions
Assembler Syntax	<code>roli rC, rA, IMM5</code>
Example	<code>roli r6, r7, 3</code>
Description	Rotates rA left by the number of bits specified in IMM5 and stores the result in rC . The bits that shift out of the register rotate into the least-significant bit positions.
Usage	In addition to the rotate-left operation, <code>roli</code> can be used to implement a rotate-right operation. Rotating left by $(32 - IMM5)$ bits is the equivalent of rotating right by IMM5 bits.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA C = Register index of operand rC IMM5 = 5-bit unsigned immediate value



Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					0					C					0x02
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x02					IMM5					0x3a					

8.5.80. ror

Instruction	rotate right
Operation	$rC \leftarrow rA$ rotated right $rB_{4..0}$ bit positions
Assembler Syntax	<code>ror rC, rA, rB</code>
Example	<code>ror r6, r7, r8</code>
Description	Rotates rA right by the number of bits specified in $rB_{4..0}$ and stores the result in rC. The bits that shift out of the register rotate into the most-significant bit positions. Bits 31– 5 of rB are ignored.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x0b
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x0b					0					0x3a					

8.5.81. sll

Instruction	shift left logical
Operation	$rC \leftarrow rA \ll (rB_{4..0})$
Assembler Syntax	<code>sll rC, rA, rB</code>
Example	<code>sll r6, r7, r8</code>
Description	Shifts rA left by the number of bits specified in $rB_{4..0}$ (inserting zeroes), and then stores the result in rC. <code>sll</code> performs the <code><<</code> operation of the C programming language.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x13
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x13					0					0x3a					

8.5.82. slli

Instruction	shift left logical immediate
Operation	$rC \leftarrow rA \ll IMM5$
Assembler Syntax	<code>slli rC, rA, IMM5</code>
Example	<code>slli r6, r7, 3</code>
Description	Shifts rA left by the number of bits specified in IMM5 (inserting zeroes), and then stores the result in rC.
Usage	slli performs the << operation of the C programming language.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA C = Register index of operand rC IMM5 = 5-bit unsigned immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					0					C					0x12
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x12					IMM5					0x3a					

8.5.83. sra

Instruction	shift right arithmetic
Operation	$rC \leftarrow (\text{signed}) rA \gg ((\text{unsigned}) rB_{4..0})$
Assembler Syntax	<code>sra rC, rA, rB</code>
Example	<code>sra r6, r7, r8</code>
Description	Shifts rA right by the number of bits specified in rB _{4..0} (duplicating the sign bit), and then stores the result in rC. Bits 31–5 are ignored.
Usage	sra performs the signed >> operation of the C programming language.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC



Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x3b
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x3b					0					0x3a					

8.5.84. srai

Instruction	shift right arithmetic immediate
Operation	$rC \leftarrow (\text{signed}) rA \gg ((\text{unsigned}) IMM5)$
Assembler Syntax	<code>srai rC, rA, IMM5</code>
Example	<code>srai r6, r7, 3</code>
Description	Shifts <i>rA</i> right by the number of bits specified in <i>IMM5</i> (duplicating the sign bit), and then stores the result in <i>rC</i> .
Usage	<code>srai</code> performs the signed \gg operation of the C programming language.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand <i>rA</i> C = Register index of operand <i>rC</i> IMM5 = 5-bit unsigned immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					0					C					0x3a
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x3a					IMM5					0x3a					

8.5.85. srl

Instruction	shift right logical
Operation	$rC \leftarrow (\text{unsigned}) rA \gg ((\text{unsigned}) rB_{4..0})$
Assembler Syntax	<code>srl rC, rA, rB</code>
Example	<code>srl r6, r7, r8</code>
Description	Shifts <i>rA</i> right by the number of bits specified in <i>rB_{4..0}</i> (inserting zeroes), and then stores the result in <i>rC</i> . Bits 31–5 are ignored.
Usage	<code>srl</code> performs the unsigned \gg operation of the C programming language.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand <i>rA</i> B = Register index of operand <i>rB</i> C = Register index of operand <i>rC</i>

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x1b
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x1b					0					0x3a					

8.5.86. srl

Instruction	shift right logical immediate
Operation	$rC \leftarrow (\text{unsigned}) rA \gg ((\text{unsigned}) IMM5)$
Assembler Syntax	<code>srl rC, rA, IMM5</code>
Example	<code>srl r6, r7, 3</code>
Description	Shifts rA right by the number of bits specified in IMM5 (inserting zeroes), and then stores the result in rC.
Usage	srl performs the unsigned >> operation of the C programming language.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA C = Register index of operand rC IMM5 = 5-bit unsigned immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x1a
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x1a					IMM5					0x3a					

8.5.87. stb / stbio

Instruction	store byte to memory or I/O periphera
Operation	$Mem8[rA + \sigma(IMM16)] \leftarrow rB_{7..0}$
Assembler Syntax	<code>stb rB, byte_offset(rA)</code> <code>stbio rB, byte_offset(rA)</code>
Example	<code>stb r6, 100(r5)</code>
Description	Computes the effective byte address specified by the sum of rA and the instruction's signed 16-bit immediate value. Stores the low byte of rB to the memory byte specified by the effective address.
Usage	In processors with a data cache, this instruction may not generate an Avalon-MM bus cycle to noncache data memory immediately. Use the stbio instruction for peripheral I/O. In processors with a data cache, stbio bypasses the cache and is guaranteed to generate an Avalon-MM data transfer. In processors without a data cache, stbio acts like stb.

continued...



Exceptions	Supervisor-only data address Misaligned data address TLB permission violation (write) Fast TLB miss (data) Double TLB miss (data) MPU region violation (data)
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Table 106. *stb*

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x05					

Table 107. *stbio*

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x25					

8.5.88. *sth* / *sthio*

Instruction	store halfword to memory or I/O peripheral
Operation	$\text{Mem16}[\text{rA} + \sigma(\text{IMM16})] \leftarrow \text{rB}_{15..0}$
Assembler Syntax	<i>sth</i> rB, byte_offset(rA) <i>sthio</i> rB, byte_offset(rA)
Example	<i>sth</i> r6, 100(r5)
Description	Computes the effective byte address specified by the sum of rA and the instruction's signed 16-bit immediate value. Stores the low halfword of rB to the memory location specified by the effective byte address. The effective byte address must be halfword aligned. If the byte address is not a multiple of 2, the operation is undefined.
Usage	In processors with a data cache, this instruction may not generate an Avalon-MM data transfer immediately. Use the <i>sthio</i> instruction for peripheral I/O. In processors with a data cache, <i>sthio</i> bypasses the cache and is guaranteed to generate an Avalon-MM data transfer. In processors without a data cache, <i>sthio</i> acts like <i>sth</i> .
Exceptions	Supervisor-only data address Misaligned data address TLB permission violation (write) Fast TLB miss (data)

continued...



Instruction Type	Double TLB miss (data) MPU region violation (data)
Instruction Fields	I A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Table 108. sth

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x0d					

Table 109. sthio

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x2d					

8.5.89. stw / stwio

Instruction	store word to memory or I/O peripheral
Operation	Mem32[rA + σ (IMM16)] ← rB
Assembler Syntax	stw rB, byte_offset(rA) stwio rB, byte_offset(rA)
Example	stw r6, 100(r5)
Description	Computes the effective byte address specified by the sum of rA and the instruction's signed 16-bit immediate value. Stores rB to the memory location specified by the effective byte address. The effective byte address must be word aligned. If the byte address is not a multiple of 4, the operation is undefined.
Usage	In processors with a data cache, this instruction may not generate an Avalon-MM data transfer immediately. Use the stwio instruction for peripheral I/O. In processors with a data cache, stwio bypasses the cache and is guaranteed to generate an Avalon-MM bus cycle. In processors without a data cache, stwio acts like stw.
Exceptions	Supervisor-only data address Misaligned data address TLB permission violation (write) Fast TLB miss (data) Double TLB miss (data)

continued...



Instruction Type	MPU region violation (data)
Instruction Fields	I A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit signed immediate value

Table 110. stw

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x15					

Table 111. stwio

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x35					

8.5.90. sub

Instruction	subtract
Operation	rC ← rA – rB
Assembler Syntax	sub rC, rA, rB
Example	sub r6, r7, r8
Description	Subtract rB from rA and store the result in rC.
Usage	<p>Carry Detection (unsigned operands):</p> <p>The carry bit indicates an unsigned overflow. Before or after a sub operation, a carry out of the MSB can be detected by checking whether the first operand is less than the second operand. The carry bit can be written to a register, or a conditional branch can be taken based on the carry condition. Both cases are shown in the following code:</p> <pre>sub rC, rA, rB cmpltu rD, rA, rB sub rC, rA, rB bltu rA, rB, label # The original sub operation (optional) # rD is written with the carry bit # The original sub operation (optional) # Branch if carry generated</pre> <p>Overflow Detection (signed operands):</p> <p>Detect overflow of signed subtraction by comparing the sign of the difference that is written to rC with the signs of the operands. If rA and rB have different signs, and the sign of</p>

continued...



rC is different than the sign of rA, an overflow occurred. The overflow condition can control a conditional branch, as shown in the following code:

```
sub rC, rA, rB
xor rD, rA, rB
xor rE, rA, rC
and rD, rD, rE
blt rD, r0, label
# The original sub operation
# Compare signs of rA and rB
# Compare signs of rA and rC
# Combine comparisons
# Branch if overflow occurred
```

Exceptions

None

Instruction Type

R

Instruction Fields

A = Register index of operand rA

B = Register index of operand rB

C = Register index of operand rC

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x39
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x39					0					0x3a					

8.5.91. subi

Instruction

subtract immediate

Operation

$rB \leftarrow rA - \sigma(\text{IMMED})$

Assembler Syntax

`subi rB, rA, IMMED`

Example

`subi r8, r8, 4`

Description

Sign-extends the immediate value IMMED to 32 bits, subtracts it from the value of rA and then stores the result in rB.

Usage

The maximum allowed value of IMMED is 32768. The minimum allowed value is -32767.

Pseudo-instruction

`subi` is implemented as `addi rB, rA, -IMMED`

8.5.92. sync

Instruction

memory synchronization

Operation

None

Assembler Syntax

`sync`

Example

`sync`

continued...



Description	Forces all pending memory accesses to complete before allowing execution of subsequent instructions. In processor cores that support in-order memory accesses only, this instruction performs no operation.
Exceptions	None
Instruction Type	R
Instruction Fields	None

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
0					0					0					0x36
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x36					0					0x3a					

8.5.93. trap

Instruction	trap
Operation	$estatus \leftarrow status$ $PIE \leftarrow 0$ $U \leftarrow 0$ $ea \leftarrow PC + 4$ $PC \leftarrow \text{exception handler address}$
Assembler Syntax	<pre>trap trap imm5</pre>
Example	trap
Description	<p>Saves the address of the next instruction in register <i>ea</i>, saves the contents of the <i>status</i> register in <i>estatus</i>, disables interrupts, and transfers execution to the exception handler. The address of the exception handler is specified with the Nios_II Processor parameter editor in Platform Designer.</p> <p>The 5-bit immediate field <i>imm5</i> is ignored by the processor, but it can be used by the debugger.</p> <p>trap with no argument is the same as trap 0.</p>
Usage	To return from the exception handler, execute an <i>eret</i> instruction.
Exceptions	Trap
Instruction Type	R
Instruction Fields	IMM5 = Type of breakpoint

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
0					0					0x1d					0x2d
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x2d					IMM5					0x3a					



8.5.94. wrctl

Instruction	write to control register
Operation	$ctlN \leftarrow rA$
Assembler Syntax	<code>wrctl ctlN, rA</code>
Example	<code>wrctl ctl6, r3</code>
Description	Writes the value contained in register rA to the control register ctlN.
Exceptions	Supervisor-only instruction
Instruction Type	R
Instruction Fields	A = Register index of operand rA N = Control register index of operand ctlN

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					0					0					0x2e
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x2e					N					0x3a					

8.5.95. wrprs

Instruction	write to previous register set
Operation	$prs.rC \leftarrow rA$
Assembler Syntax	<code>wrprs rC, rA</code>
Example	<code>wrprs r6, r7</code>
Description	Copies the value of rA in the current register set to rC in the previous register set. This instruction can set r0 to 0 in a shadow register set.
Usage	<p>The previous register set is specified by <code>status.PRS</code>. By default, <code>status.PRS</code> indicates the register set in use before an exception, such as an external interrupt, caused a register set change.</p> <p>To write to an arbitrary register set, software can insert the desired register set number in <code>status.PRS</code> prior to executing <code>wrprs</code>.</p> <p>System software must use <code>wrprs</code> to initialize r0 to 0 in each shadow register set before using that register set.</p> <p>If shadow register sets are not implemented on the Nios II core, <code>wrprs</code> is an illegal instruction.</p>
Exceptions	Supervisor-only instruction Illegal instruction
Instruction Type	R
Instruction Fields	A = Register index of operand rA C = Register index of operand rC



Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					0					C					0x14
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x14					0					0x3a					

8.5.96. xor

Instruction	bitwise logical exclusive or
Operation	$rC \leftarrow rA \wedge rB$
Assembler Syntax	<code>xor rC, rA, rB</code>
Example	<code>xor r6, r7, r8</code>
Description	Calculates the bitwise logical exclusive-or of rA and rB and stores the result in rC.
Exceptions	None
Instruction Type	R
Instruction Fields	A = Register index of operand rA B = Register index of operand rB C = Register index of operand rC

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					C					0x1e
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0x1e					0					0x3a					

8.5.97. xorhi

Instruction	bitwise logical exclusive or immediate into high halfword
Operation	$rB \leftarrow rA \wedge (\text{IMM16} : 0x0000)$
Assembler Syntax	<code>xorhi rB, rA, IMM16</code>
Example	<code>xorhi r6, r7, 100</code>
Description	Calculates the bitwise logical exclusive XOR of rA and (IMM16 : 0x0000) and stores the result in rB.
Exceptions	None
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit unsigned immediate value



Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x3c					

8.5.98. xori

Instruction	bitwise logical exclusive or immediate
Operation	$rB \leftarrow rA \wedge (0x0000 : IMM16)$
Assembler Syntax	<code>xori rB, rA, IMM16</code>
Example	<code>xori r6, r7, 100</code>
Description	Calculates the bitwise logical exclusive OR of rA and (0x0000 : IMM16) and stores the result in rB.
Exceptions	None
Instruction Type	I
Instruction Fields	A = Register index of operand rA B = Register index of operand rB IMM16 = 16-bit unsigned immediate value

Bit Fields															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
A					B					IMM16					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IMM16										0x1c					

8.6. Instruction Set Reference Revision History

Document Version	Changes
2020.04.24	Corrected the <i>Assembler Pseudo-Instructions</i> .
2019.08.21	<ul style="list-style-type: none"> Corrected Operation for: <ul style="list-style-type: none"> cmpeqi custom ldw / ldwio Corrected Instruction Type for: <ul style="list-style-type: none"> cmpge cmpgei Corrected Bit Fields for: <ul style="list-style-type: none"> custom
2018.04.18	Implemented editorial enhancements.
2016.10.28	Maintenance release.
2015.04.02	Initial release