

Hazard3

Table of Contents

1. Introduction	1
2. Instruction Cycle Counts	2
2.1. RV32I	2
2.2. M Extension	3
2.3. C Extension	4
2.4. Privileged Instructions (including Zicsr)	4
3. CSRs	5
3.1. Standard CSRs	5
3.1.1. mvendorid	5
3.1.2. marchid	5
3.1.3. mimpid	5
3.1.4. mstatus	5
3.1.5. misa	5
3.2. Custom CSRs	5
3.2.1. midcr	5
3.2.2. meie0	6
3.2.3. meip0	6
3.2.4. mlei	7
3.2.5. Maybe-adds	7
4. Debug	8
4.1. Implementation-defined behaviour	8
4.2. UART DTM	9

Chapter 1. Introduction

Hazard3 is a 3-stage RISC-V processor, providing the following architectural support:

- **RV32I**: 32-bit base instruction set
- **M** extension: integer multiply/divide/modulo
- **C** extension: compressed instructions
- **Zicsr** extension: CSR access
- M-mode privileged instructions **ECALL**, **EBREAK**, **MRET**
- The machine-mode (M-mode) privilege state, and standard M-mode CSRs

The following are planned for future implementation:

- Support for **WFI** instruction
- Debug support
- **A** extension: atomic memory access
 - **LR/SC** fully supported
 - AMONone PMA on all of memory (AMOs are decoded but unconditionally trigger access fault without attempting memory access)
- Some nonstandard M-mode CSRs for interrupt control etc

Chapter 2. Instruction Cycle Counts

All timings are given assuming perfect bus behaviour (no stalls). Stalling of the **I** bus can delay execution indefinitely, as can stalling of the **D** bus during a load or store.

2.1. RV32I

Instruction	Cycles	Note
Integer Register-register		
<code>add rd, rs1, rs2</code>	1	
<code>sub rd, rs1, rs2</code>	1	
<code>slt rd, rs1, rs2</code>	1	
<code>sltu rd, rs1, rs2</code>	1	
<code>and rd, rs1, rs2</code>	1	
<code>or rd, rs1, rs2</code>	1	
<code>xor rd, rs1, rs2</code>	1	
<code>sll rd, rs1, rs2</code>	1	
<code>srl rd, rs1, rs2</code>	1	
<code>sra rd, rs1, rs2</code>	1	
Integer Register-immediate		
<code>addi rd, rs1, imm</code>	1	<code>nop</code> is a pseudo-op for <code>addi x0, x0, 0</code>
<code>slti rd, rs1, imm</code>	1	
<code>sltiu rd, rs1, imm</code>	1	
<code>andi rd, rs1, imm</code>	1	
<code>ori rd, rs1, imm</code>	1	
<code>xori rd, rs1, imm</code>	1	
<code>slli rd, rs1, imm</code>	1	
<code>srli rd, rs1, imm</code>	1	
<code>srai rd, rs1, imm</code>	1	
Large Immediate		
<code>lui rd, imm</code>	1	
<code>auipc rd, imm</code>	1	
Control Transfer		
<code>jal rd, label</code>	2 ^[1]	
<code>jalr rd, rs1, imm</code>	2 ^[1]	
<code>beq rs1, rs2, label</code>	1 or 2 ^[1]	1 if nontaken, 2 if taken.

Instruction	Cycles	Note
bne rs1, rs2, label	1 or 2 ^[1]	1 if nontaken, 2 if taken.
blt rs1, rs2, label	1 or 2 ^[1]	1 if nontaken, 2 if taken.
bge rs1, rs2, label	1 or 2 ^[1]	1 if nontaken, 2 if taken.
bltu rs1, rs2, label	1 or 2 ^[1]	1 if nontaken, 2 if taken.
bgeu rs1, rs2, label	1 or 2 ^[1]	1 if nontaken, 2 if taken.
Load and Store		
lw rd, imm(rs1)	1 or 2	1 if next instruction is independent, 2 if dependent. ^[2]
lh rd, imm(rs1)	1 or 2	1 if next instruction is independent, 2 if dependent. ^[2]
lhu rd, imm(rs1)	1 or 2	1 if next instruction is independent, 2 if dependent. ^[2]
lb rd, imm(rs1)	1 or 2	1 if next instruction is independent, 2 if dependent. ^[2]
lbu rd, imm(rs1)	1 or 2	1 if next instruction is independent, 2 if dependent. ^[2]
sw rs2, imm(rs1)	1	
sh rs2, imm(rs1)	1	
sb rs2, imm(rs1)	1	

2.2. M Extension

Timings assume the core is configured with `MULDIV_UNROLL = 2` and `MUL_FAST = 1`. I.e. the sequential multiply/divide circuit processes two bits per cycle, and a separate dedicated multiplier is present for the `mul` instruction.

Instruction	Cycles	Note
32 × 32 → 32 Multiply		
mul rd, rs1, rs2	1 or 2	1 if next instruction is independent, 2 if dependent.
32 × 32 → 64 Multiply, Upper Half		
mulh rd, rs1, rs2	18 to 20	Depending on sign correction
mulhsu rd, rs1, rs2	18 to 20	Depending on sign correction
mulhu rd, rs1, rs2	18	
Divide and Remainder		
div	18 or 19	Depending on sign correction
divu	18	
rem	18 or 19	Depending on sign correction
remu	18	

2.3. C Extension

All C extension 16-bit instructions on Hazard3 are aliases of base RV32I instructions. They perform identically to their 32-bit counterparts.

A consequence of the C extension is that 32-bit instructions can be non-naturally-aligned. This has no penalty during sequential execution, but branching to a 32-bit instruction that is not 32-bit-aligned carries a 1 cycle penalty, because the instruction fetch is cracked into two naturally-aligned bus accesses.

2.4. Privileged Instructions (including Zicsr)

Instruction	Cycles	Note
CSR Access		
<code>csrrw rd, csr, rs1</code>	1	
<code>csrrc rd, csr, rs1</code>	1	
<code>csrrs rd, csr, rs1</code>	1	
<code>csrrwi rd, csr, imm</code>	1	
<code>csrrci rd, csr, imm</code>	1	
<code>csrrsi rd, csr, imm</code>	1	
Trap Request		
<code>ecall</code>	3	Time given is for jumping to <code>mtvec</code>
<code>ebreak</code>	3	Time given is for jumping to <code>mtvec</code>

[1] A branch to a 32-bit instruction which is not 32-bit-aligned requires one additional cycle, because two naturally-aligned bus cycles are required to fetch the target instruction.

[2] If an instruction uses load data (from stage 3) in stage 2, a 1-cycle bubble is inserted after the load. Load-data to store-data dependency does not experience this, because the store data is used in stage 3. However, load-data to store-address (or e.g. load-to-add) does qualify.

Chapter 3. CSRs

The RISC-V privileged specification affords flexibility as to which CSRs are implemented, and how they behave. This section documents the concrete behaviour of Hazard3's standard and nonstandard M-mode CSRs, as implemented.

3.1. Standard CSRs

3.1.1. mvendorid

Address: `0xf11`

Read-only, constant. Value is configured when the processor is instantiated. Should contain either all-zeroes, or a valid JEDEC JEP106 vendor ID.

3.1.2. marchid

Address: `0xf12`

Read-only, constant. Architecture identifier for Hazard3, value can be altered when the processor is instantiated. Default is currently all zeroes as unregistered.

3.1.3. mimpid

Address: `0xf12`

Read-only, constant. Value is configured when the processor is instantiated. Should contain either all-zeroes, or some number specifying a version of Hazard3 (e.g. git hash).

3.1.4. mstatus

blah blah

3.1.5. misa

Read-only, constant. Value depends on which ISA extensions Hazard5 is configured with.

3.2. Custom CSRs

These are all allocated in the space `0xbc0` through `0xbff` which is available for custom read/write M-mode CSRs, and `0xfc0` through `0xfff` which is available for custom read-only M-mode CSRs.

3.2.1. midcr

Address: `0xbc0`

Implementation-defined control register. Miscellaneous nonstandard controls.

Bits	Name	Description
31:1	-	RES0
0	<code>eivect</code>	Modified external interrupt vectoring. If 0, use standard behaviour: all external interrupts set interrupt <code>mcause</code> of 11 and vector to <code>mtvec + 0x2c</code> . If 1, external interrupts use distinct interrupt <code>mcause</code> numbers 16 upward, and distinct vectors <code>mtvec + (irq + 16) * 4</code> . Resets to 0. Has no effect when <code>mtvec[0]</code> is 0.

3.2.2. `meie0`

Address: `0xbe0`

External interrupt enable register 0. Contains a read-write bit for each external interrupt request IRQ0 through IRQ31. A 1 bit indicates that interrupt is currently enabled.

Addresses `0xbe1` through `0xbe3` are reserved for further `meie` registers, supporting up to 128 external interrupts.

An external interrupt is taken when all of the following are true:

- The interrupt is currently asserted in `meip0`
- The matching interrupt enable bit is set in `meie0`
- The standard M-mode interrupt enable `mstatus.mie` is set
- The standard M-mode global external interrupt enable `mie.meie` is set

`meie0` resets to **all-ones**, for compatibility with software which is only aware of `mstatus` and `mie`. Because `mstatus.mie` and `mie.meie` are both initially clear, the core will not take interrupts straight out of reset, but it is strongly recommended to configure `meie0` before setting the global interrupt enable, to avoid interrupts from unexpected sources.

3.2.3. `meip0`

Address: `0xfe0`

External IRQ pending register 0. Contains a read-only bit for each external interrupt request IRQ0 through IRQ31. A 1 bit indicates that interrupt is currently asserted. IRQs are assumed to be level-sensitive, and the relevant `meip0` bit is cleared by servicing the requestor so that it deasserts its interrupt request.

Addresses `0xfe1` through `0xfe3` are reserved for further `meip` registers, supporting up to 128 external interrupts.

When any bit is set in both `meip0` and `meie0`, the standard external interrupt pending bit `mip.meip` is also set. In other words, `meip0` is filtered by `meie0` to generate the standard `mip.meip` flag. So, an external interrupt is taken when *all* of the following are true:

- An interrupt is currently asserted in `meip0`

- The matching interrupt enable bit is set in `meie0`
- The standard M-mode interrupt enable `mstatus.mie` is set
- The standard M-mode global external interrupt enable `mie.meie` is set

In this case, the processor jumps to either:

- `mtvec` directly, if vectoring is disabled (`mtvec[0]` is 0)
- `mtvec + 0x2c`, if vectoring is enabled (`mtvec[0]` is 1) and modified external IRQ vectoring is disabled (`midcr.eivect` is 0)
- `mtvect + (mlei + 16) * 4`, if vectoring is enabled (`mtvec[0]` is 1) and modified external IRQ vectoring is enabled (`midcr.eivect` is 1).
 - `mlei` is a read-only CSR containing the lowest-numbered pending-and-enabled external interrupt.

3.2.4. mlei

Address: `0xfe4`

Lowest external interrupt. Contains the index of the lowest-numbered external interrupt which is both asserted in `meip0` and enabled in `meie0`. Can be used for faster software vectoring when modified external interrupt vectoring (`midcr.eivect = 1`) is not in use.

Bits	Name	Description
31:5	-	RES0
4:0	-	Index of the lowest-numbered active external interrupt. A LSB-first priority encode of <code>meip0</code> & <code>meie0</code> . Zero when no external interrupts are both pending and enabled.

3.2.5. Maybe-adds

An option to clear a bit in `meie0` when that interrupt is taken, and set it when an `mret` has a matching `mcause` for that interrupt. Makes preemption support easier.

Chapter 4. Debug

Currently the plan is for Hazard3, with its associated debug module (DM), to support the following:

- Run/halt/reset control as required
- Abstract GPR access as required
- Program buffer: 2 words plus `impebreak`
- Automatic program buffer execution triggered by abstract GPR access (`abstractauto`)
- Some minimum useful trigger unit — likely just breakpoints, no watchpoints

The core itself will implement the following, enabling the DM to provide a compliant debug interface:

- Debug mode CSRs
- Ability to enter debug mode with correct update of `dpc` etc
 - Synchronously via exception, `ebreak` or trigger match
 - Asynchronously via external halt request
- Ability to exit debug mode to M mode
- Address query/match interface for external trigger unit
- Ability to inject words into the instruction prefetch queue when the processor is halted
- Ability to suppress exception entry when executing instructions in debug mode, and provide an external signal to indicate the exception took place
- A read/write data bus which allows the DM to intercept core CSR accesses

The DM implements abstract GPR access by injecting a dummy CSR access instruction, and manipulating the CSR port to get data in/out of the core. A `csrr` is used to write to a core register, and a `csrw` to read from a core register. By injecting a `csrrw`, the DM can *swap* a GPR with one of its own internal registers, though this is not exposed through the abstract GPR access command.

The debugger implements memory and CSR access using the Program Buffer, which uses the same instruction injection interface used by the DM to implement abstract GPR access. The `abstractauto` feature allows the DM to execute the program buffer automatically following every abstract GPR access, which can be used for e.g. autoincrementing read/write memory bursts.

4.1. Implementation-defined behaviour

This is not an exhaustive list (yet).

DM feature support:

- Abstract CSR and memory access are not implemented
- The Program Buffer is implemented, size 2 words, `impebreak` = 1.
- A single data register (`data0`) is implemented as a per-hart CSR accessible by the DM

- `abstractauto` is supported on the program buffer registers and the data register
- Multiple hart selection (`hase1` = 1) is not supported

Core behaviour:

- All control transfer instructions are illegal in debug mode (depend on value of PC)
- `auipc` is illegal in debug mode (depends on value of PC)
- The `dret` instruction is not supported (a special purpose DM-to-core signal is used to signal resume)
- Entering and exiting debug mode does not clear an atomic load reservation; the host may explicitly clear a reservation using a dummy `sc` instruction via the program buffer.
- The `dscratch` CSRs are not implemented
- `data0` is implemented as a scratch CSR mapped at `0x7b2` (the location of `dscratch0`), readable and writable by the debugger.
- `dcsr.stepie` is hardwired to 0 (no interrupts during single stepping)
- `dcsr.stopcount` and `dcsr.stoptime` are hardwired to 1 (no counter/timer increment in debug mode)
- `dcsr.mprven` is hardwired to 0
- `dcsr.prv` is hardwired to 3 (M-mode)

4.2. UART DTM

Hazard3 defines a minimal UART Debug Transport Module, which allows the Debug Module to be accessed via a standard 8n1 asynchronous serial port. The UART DTM is always accessed by the host using a two-wire serial interface (TXD RXD) running at 1 Mbaud. The interface between the DTM and DM is an AMBA 3 APB port with a 32-bit data bus and 8-bit address bus.

This is a quick hack, and not suitable for production systems:

- Debug hardware should not expect a frequency reference for a UART to be present
- The UART DTM does not implement any flow control or error detection/correction

The host may send the following commands:

Command	To DTM	From DTM
<code>0x00</code> NOP	-	-
<code>0x01</code> Read ID	-	4-byte ID, same format as JTAG-DTM ID (JEP106-compatible)
<code>0x02</code> Read DMI	1 address byte	4 data bytes
<code>0x03</code> Write DMI	1 address byte, 4 data bytes	data bytes echoed back
<code>0xa5</code> Disconnect	-	-

Initially after power-on the DTM is in the Dormant state, and will ignore any commands. The host sends the magic sequence "SUP?" (0x53, 0x55, 0x50, 0x3f) to wake the DTM, and then issues a Read ID command to check the link is up. The DTM can be returned to the Dormant state at any time using the 0xa5 Disconnect command.

So that the host can queue up batches of commands in its transmit buffer, without overrunning the DTM's transmit bandwidth, it's recommended to pad each command with NOPs so that it is strictly larger than the response. For example, a Read ID should be followed by four NOPs, and a Read DMI should be followed by 3 NOPs.

To recover command framing, write 6 NOP commands (the length of the longest commands). This will be interpreted as between 1 and 6 NOPs depending on the DTM's state.

This interface assumes the DMI data transfer takes very little time compared with the UART access (typically less than one baud period). When the host-to-DTM bandwidth is kept greater than the DTM-to-host bandwidth, thanks to appropriate NOP padding, the host can queue up batches of commands in its transmit buffer, and this should never overrun the DTM's response channel. So, the 1 Mbaud 8n1 UART link provides 67 kB/s of half-duplex data bandwidth between host and DM, which is enough to get your system off the ground.