



ZIP CPU SPECIFICATION

Dan Gisselquist, Ph.D.
dgisselq (at) opencores.org

November 17, 2015

Copyright (C) 2015, Gisselquist Technology, LLC

This project is free software (firmware): you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program. If not, see <http://www.gnu.org/licenses/> for a copy.

Revision History

Rev.	Date	Author	Description
0.6	11/17/2015	Gisselquist	Added graphics to illustrate pipeline discussion.
0.5	9/29/2015	Gisselquist	Added pipelined memory access discussion.
0.4	9/19/2015	Gisselquist	Added DMA controller, improved stall information, and self-assessment info.
0.3	8/22/2015	Gisselquist	First completed draft
0.2	8/19/2015	Gisselquist	Still Draft, more complete
0.1	8/17/2015	Gisselquist	Incomplete First Draft

Contents

	Page
1	Introduction 1
1.1	Characteristics of a SwiC 1
1.2	Lessons Learned 3
2	CPU Architecture 8
2.1	Simplified Bus 8
2.2	Register Set 8
2.3	Conditional Instructions 10
2.4	Traditional Interrupt Handling 11
2.5	Operand B 11
2.6	Address Modes 12
2.7	Move Operands 12
2.8	Multiply Operations 12
2.9	Floating Point 13
2.10	Native Instructions 13
2.11	Derived Instructions 13
2.12	Pipeline Stages 13
2.13	Pipeline Stalls 19
3	Peripherals 26
3.1	Interrupt Controller 26
3.2	Counter 27
3.3	Timer 27
3.4	Watchdog Timer 27
3.5	Bus Watchdog 28
3.6	Jiffies 28
3.7	Direct Memory Access Controller 28
4	Operation 30
4.1	System High 30
4.2	Traditional Interrupt Handling 31
4.3	Example: Idle Task 34
4.4	Example: Memory Copy 35
4.5	Example: Context Switch 35
5	Registers 41
5.1	Peripheral Registers 42
5.2	Debug Port Registers 45
6	Wishbone Datasheets 47
7	Clocks 49
8	I/O Ports 50
9	Initial Assessment 52
9.1	The Good 52
9.2	The Not so Good 53
9.3	The Next Generation 54

Figures

Figure		Page
1.1	Zip CPU internal pipeline architecture	2
1.2	An Ideal Pipeline: One instruction per clock cycle	4
1.3	Instructions wait for each other	5
1.4	Instructions proceed independently	5
1.5	A typical branch delay slot approach	5
1.6	The branch delay slot breaks with a slow memory	6
1.7	How the CPU halts when sleeping	6
1.8	Instructions can stack up behind a stalled instruction	7
2.1	Zip CPU Register File	9
2.2	A conditional branch generates 5 stall cycles	19
2.3	An expedited delay costs only 2 stall cycles	20
2.4	A (not taken) conditional branch followed by a memory operation	21
2.5	Pipeline handling of a load instruction	22
2.6	Pipeline handling of a store instruction	23
2.7	Pipeline handling of a store followed by a load instruction	24
3.1	Zip System Peripherals	27

Tables

Table		Page
2.1	Condition Code Register Bit Assignment	9
2.2	Condition Code / Status Register Bits	10
2.3	Conditions for conditional operand execution	11
2.4	An example of a double conditional	11
2.5	Bit allocation for Operand B	12
2.6	Zip CPU Instruction Set	14
2.7	Derived Instructions	15
2.8	Derived Instructions, continued	16
2.9	Derived Instructions, continued	17
2.10	Derived Instructions, continued	18
4.1	Executing an idle from supervisor mode	31
4.2	Traditional Interrupt handling	32
4.3	Example Saving Minimal User Context	33
4.4	Example Restoring Minimal User Context	34
4.5	Example Idle Loop	34
4.6	Example Memory Copy code in C	35
4.7	Example Memory Copy code in Zip Assembly	35
4.8	Checking for whether the user issued a TRAP instruction	36
4.9	Example Storing User Task Context	37
4.10	Example Watchdog Reset	37
4.11	Example checking for active interrupts	38
4.12	Example Restoring User Task Context	40
5.1	Zip System Internal/Peripheral Registers	41
5.2	Zip System Debug Registers	41
5.3	Interrupt Controller Register Bits	42
5.4	Timer Register Bits	43
5.5	Jiffies Register Bits	43
5.6	Counter Register Bits	43
5.7	DMA Control Register Bits	44
5.8	Debug Control Register Bits	45
5.9	Debug Register Addresses	46
6.1	Wishbone Datasheet for the Debug Interface	47
6.2	Wishbone Datasheet for the CPU as Master	48
7.1	List of Clocks	49
8.1	CPU Master Wishbone I/O Ports	50
8.2	CPU Debug Wishbone I/O Ports	51
8.3	I/O Ports	51

Preface

Many people have asked me why I am building the Zip CPU. ARM processors are good and effective. Xilinx makes and markets Microblaze, Altera Nios, and both have better toolsets than the Zip CPU will ever have. OpenRISC is also available, RISC-V may be replacing it. Why build a new processor?

The easiest, most obvious answer is the simple one: Because I can.

There's more to it, though. There's a lot that I would like to do with a processor, and I want to be able to do it in a vendor independent fashion. First, I would like to be able to place this processor inside an FPGA. Without paying royalties, ARM is out of the question. I would then like to be able to generate Verilog code, both for the processor and the system it sits within, that can run equivalently on both Xilinx and Altera chips, and that can be easily ported from one manufacturer's chipsets to another. Even more, before purchasing a chip or a board, I would like to know that my soft core works. I would like to build a test bench to test components with, and Verilator is my chosen test bench. This forces me to use all Verilog, and it prevents me from using any proprietary cores. For this reason, Microblaze and Nios are out of the question.

Why not OpenRISC? That's a hard question. The OpenRISC team has done some wonderful work on an amazing processor, and I'll have to admit that I am envious of what they've accomplished. I would like to port binutils to the Zip CPU, as I would like to port GCC and GDB. They are way ahead of me. The OpenRISC processor, however, is complex and hefty at about 4,500 LUTs. It has a lot of features of modern CPUs within it that ... well, let's just say it's not the little guy on the block. The Zip CPU is lighter weight, costing only about 2,300 LUTs with no peripherals, and 3,200 LUTs with some very basic peripherals.

My final reason is that I'm building the Zip CPU as a learning experience. The Zip CPU has allowed me to learn a lot about how CPUs work on a very micro level. For the first time, I am beginning to understand many of the Computer Architecture lessons from years ago.

To summarize: Because I can, because it is open source, because it is light weight, and as an exercise in learning.

Dan Gisselquist, Ph.D.

1.

Introduction

The original goal of the Zip CPU was to be a very simple CPU. You might think of it as a poor man's alternative to the OpenRISC architecture. For this reason, all instructions have been designed to be as simple as possible, and are all designed to be executed in one instruction cycle per instruction, barring pipeline stalls. Indeed, even the bus has been simplified to a constant 32-bit width, with no option for more or less. This has resulted in the choice to drop push and pop instructions, pre-increment and post-decrement addressing modes, and more.

For those who like buzz words, the Zip CPU is:

- A 32-bit CPU: All registers are 32-bits, addresses are 32-bits, instructions are 32-bits wide, etc.
- A RISC CPU. There is no microcode for executing instructions. All instructions are designed to be completed in one clock cycle.
- A Load/Store architecture. (Only load and store instructions can access memory.)
- Wishbone compliant. All peripherals are accessed just like memory across this bus.
- A Von-Neumann architecture. (The instructions and data share a common bus.)
- A pipelined architecture, having stages for **Prefetch**, **Decode**, **Read-Operand**, the **ALU/Memory** unit, and **Write-back**. See Fig. 1.1 for a diagram of this structure.
- Completely open source, licensed under the GPL.¹

The Zip CPU also has one very unique feature: the ability to do pipelined loads and stores. This allows the CPU to access on-chip memory at one access per clock, minus a stall for the initial access.

1.1 Characteristics of a SwiC

Here, we shall define a soft core internal to an FPGA as a "System within a Chip," or a SwiC. SwiCs have some very unique properties internal to them that have influenced the design of the Zip CPU. Among these are the bus, memory, and available peripherals.

Most other approaches to soft core CPU's employ a Harvard architecture. This allows these other CPU's to have two separate bus structures: one for the program fetch, and the other for thememory. The Zip CPU is fairly unique in its approach because it uses a von Neumann architecture. This was

¹Should you need a copy of the Zip CPU licensed under other terms, please contact me.

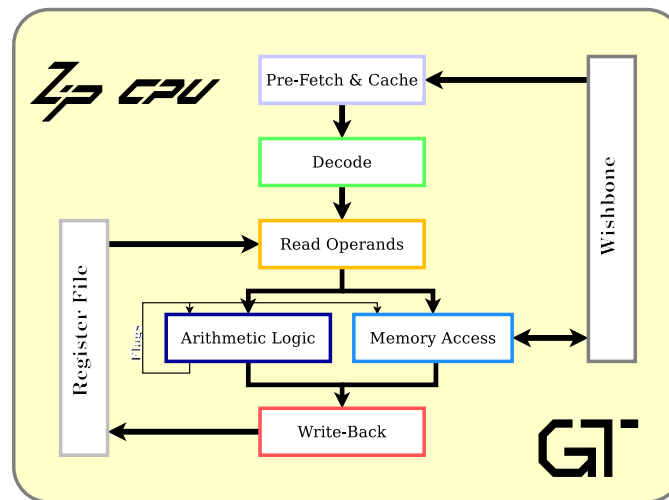


Figure 1.1: Zip CPU internal pipeline architecture

done for simplicity. By using a von Neumann architecture, only one bus needs to be implemented within any FPGA. This helps to minimize real-estate, while maintaining a high clock speed. The disadvantage is that it can severely degrade the overall instructions per clock count.

Soft core's within an FPGA have an additional characteristic regarding memory access: it is slow. Memory on chip may be accessed at a single cycle per access, but small FPGA's have a limited amount of memory on chip. Going off chip, however, is expensive. Two examples will prove this point. On the XuLA2 board, Flash can be accessed at 128 cycles per 32-bit word, or 64 cycles per subsequent word in a pipelined architecture. Likewise, the SDRAM chip on the XuLA2 board allows 6 cycle access for a write, 10 cycles per read, and 2 cycles for any subsequent pipelined access read or write. Either way you look at it, this memory access will be slow and this doesn't account for any logic delays should the bus implementation logic get complicated.

As may be noticed from the above discussion about memory speed, a second characteristic of memory is that all memory accesses may be pipelined, and that pipelined memory access is faster than non-pipelined access. Therefore, a SwiC soft core should support pipelined operations, but it should also allow a higher priority subsystem to get access to the bus (no starvation).

As a further characteristic of SwiC memory options, on-chip cache's are expensive. If you want to have a minimum of logic, cache logic may not be the highest on the priority list.

In sum, memory is slow. While one processor on one FPGA may be able to fill its pipeline, the same processor on another FPGA may struggle to get more than one instruction at a time into the pipeline. Any SwiC must be able to deal with both cases: fast and slow memories.

A final characteristic of SwiC's within FPGA's is the peripherals. Specifically, FPGA's are highly reconfigurable. Soft peripherals can easily be created on chip to support the SwiC if necessary. As an example, a simple 30-bit peripheral could easily support reversing 30-bit numbers: a read from the peripheral returns it's bit-reversed address. This is cheap within an FPGA, but expensive in instructions.

Indeed, anything that must be done fast within an FPGA is likely to already be done elsewhere in the fabric. This leaves the CPU with the role of handling sequential tasks that need a lot of state.

This means that the SwiC needs to live within a very unique environment, separate and different from the traditional SoC. That isn't to say that a SwiC cannot be turned into a SoC, just that this SwiC has not been designed for that purpose.

1.2 Lessons Learned

Now, however, that I've worked on the Zip CPU for a while, it is not nearly as simple as I originally hoped. Worse, I've had to adjust to create capabilities that I was never expecting to need. These include:

- **External Debug:** Once placed upon an FPGA, some external means is still necessary to debug this CPU. That means that there needs to be an external register that can control the CPU: reset it, halt it, step it, and tell whether it is running or not. My chosen interface includes a second register similar to this control register. This second register allows the external controller or debugger to examine registers internal to the CPU.
- **Internal Debug:** Being able to run a debugger from within a user process requires an ability to step a user process from within a debugger. It also requires a break instruction that can be substituted for any other instruction, and substituted back. The break is actually difficult: the break instruction cannot be allowed to execute. That way, upon a break, the debugger should be able to jump back into the user process to step the instruction that would've been at the break point initially, and then to replace the break after passing it.

Incidentally, this break messes with the prefetch cache and the pipeline: if you change an instruction partially through the pipeline, the whole pipeline needs to be cleansed. Likewise if you change an instruction in memory, you need to make sure the cache is reloaded with the new instruction.

- **Prefetch Cache:** My original implementation had a very simple prefetch stage. Any time the PC changed the prefetch would go and fetch the new instruction. While this was perhaps this simplest approach, it cost roughly five clocks for every instruction. This was deemed unacceptable, as I wanted a CPU that could execute instructions in one cycle. I therefore have a prefetch cache that issues pipelined wishbone accesses to memory and then pushes instructions at the CPU. Sadly, this accounts for about 20% of the logic in the entire CPU, or 15% of the logic in the entire system.
- **Operating System:** In order to support an operating system, interrupts and so forth, the CPU needs to support supervisor and user modes, as well as a means of switching between them. For example, the user needs a means of executing a system call. This is the purpose of the **'trap'** instruction. This instruction needs to place the CPU into supervisor mode (here equivalent to disabling interrupts), as well as handing it a parameter such as identifying which O/S function was called.

My initial approach to building a trap instruction was to create an external peripheral which, when written to, would generate an interrupt and could return the last value written to it. In practice, this approach didn't work at all: the CPU executed two instructions while waiting




Figure 1.2: An Ideal Pipeline: One instruction per clock cycle

for the trap interrupt to take place. Since then, I’ve decided to keep the rest of the CC register for that purpose so that a write to the CC register, with the GIE bit cleared, could be used to execute a trap. This has other problems, though, primarily in the limitation of the uses of the CC register. In particular, the CC register is the best place to put CPU state information and to “announce” special CPU features (floating point, etc). So the trap instruction still switches to interrupt mode, but the CC register is not nearly as useful for telling the supervisor mode processor what trap is being executed.

Modern timesharing systems also depend upon a **Timer** interrupt to handle task swapping. For the Zip CPU, this interrupt is handled external to the CPU as part of the CPU System, found in `zipsystem.v`. The timer module itself is found in `ziptimer.v`.

- **Pipeline Stalls:** My original plan was to not support pipeline stalls at all, but rather to require the compiler to properly schedule all instructions so that stalls would never be necessary. After trying to build such an architecture, I gave up, having learned some things:

First, an ideal pipeline might look something like Fig. 1.2. Notice that, in this figure, all the pipeline stages are complete and full. Every instruction takes one clock and there are no delays. However, as the discussion above pointed out, the memory associated with a SwiC may not allow single clock access. It may be instead that you can only read every two clocks. In that case, what shall the pipeline look like? Should it look like Fig. 1.3, where instructions are held back until the pipeline is full, or should it look like Fig. 1.4, where each instruction is allowed to move through the pipeline independently? For better or worse, the Zip CPU allows instructions to move through the pipeline independently.

One approach to avoiding stalls is to use a branch delay slot, such as is shown in Fig. 1.5. In this figure, instructions BR (a branch), BD (a branch delay instruction), are followed by instructions after the branch: IA, IB, etc. Since it takes a processor a clock cycle to execute a branch, the delay slot allows the processor to do something useful in that branch. The problem the Zip CPU has with this approach is, what happens when the pipeline looks like Fig. 1.6? In this case, the branch delay slot never gets filled in the first place, and so the pipeline squashes it before it gets executed. If not that, then what happens when handling interrupts or debug




Figure 1.3: Instructions wait for each other




Figure 1.4: Instructions proceed independently




Figure 1.5: A typical branch delay slot approach




Figure 1.8: Instructions can stack up behind a stalled instruction

are valid and the next step is stalled. This allows the pipeline to fill any time a later stage stalls, as illustrated in Fig. 1.8.

This approach is also different from other pipeline approaches. Instead of keeping the entire pipeline filled, each stage is treated independently. Therefore, individual stages may move forward as long as the subsequent stage is available, regardless of whether the stage behind it is filled.

- **Verilog Modules:** When examining how other processors worked here on open cores, many of them had one separate module per pipeline stage. While this appeared to me to be a fascinating and commendable idea, my own implementation didn't work out quite so nicely.

As an example, the decode module produces a *lot* of control wires and registers. Creating a module out of this, with only the simplest of logic within it, seemed to be more a lesson in passing wires around, rather than encapsulating logic.

Another example was the register writeback section. I would love this section to be a module in its own right, and many have made them such. However, other modules depend upon writeback results other than just what's placed in the register (i.e., the control wires). For these reasons, I didn't manage to fit this section into its own module.

The result is that the majority of the CPU code can be found in the `zipcpu.v` file.

With that introduction out of the way, let's move on to the instruction set.

2.

CPU Architecture

The Zip CPU supports a set of two operand instructions, where the second operand (always a register) is the result. The only exception is the store instruction, where the first operand (always a register) is the source of the data to be stored.

2.1 Simplified Bus

The bus architecture of the Zip CPU is that of a simplified WISHBONE bus. It has been simplified in this fashion: all operations are 32-bit operations. The bus is neither little endian nor big endian. For this reason, all words are 32-bits. All instructions are also 32-bits wide. Everything has been built around the 32-bit word.

2.2 Register Set

The Zip CPU supports two sets of sixteen 32-bit registers, a supervisor and a user set as shown in Fig. 2.1. The supervisor set is used in interrupt mode when interrupts are disabled, whereas the user set is used otherwise. Of this register set, the Program Counter (PC) is register 15, whereas the status register (SR) or condition code register (CC) is register 14. By convention, the stack pointer will be register 13 and noted as (SP)—although there is nothing special about this register other than this convention. The CPU can access both register sets via move instructions from the supervisor state, whereas the user state can only access the user registers.

The status register is special, and bears further mention. As shown in Fig. 2.1, the lower 11 bits of the status register form a set of CPU state and condition codes. Writes to other bits of this register are preserved.

Of the condition codes, the bottom four bits are the current flags: Zero (Z), Carry (C), Negative (N), and Overflow (V).

The next bit is a clock enable (0 to enable) or sleep bit (1 to put the CPU to sleep). Setting this bit will cause the CPU to wait for an interrupt (if interrupts are enabled), or to completely halt (if interrupts are disabled).

The sixth bit is a global interrupt enable bit (GIE). When this sixth bit is a '1' interrupts will be enabled, else disabled. When interrupts are disabled, the CPU will be in supervisor mode, otherwise it is in user mode. Thus, to execute a context switch, one only need enable or disable interrupts. (When an interrupt line goes high, interrupts will automatically be disabled, as the CPU goes and deals with its context switch.) Special logic has been added to keep the user mode from setting the




Figure 2.1: Zip CPU Register File

Bit #	Access	Description
31...11	R/W	Reserved for future uses
10	R	(Reserved for) Bus-Error Flag
9	R	Trap, or user interrupt, Flag. Cleared on return to userspace.
8	R	Illegal Instruction Flag
7	R/W	Break-Enable
6	R/W	Step
5	R/W	Global Interrupt Enable (GIE)
4	R/W	Sleep. When GIE is also set, the CPU waits for an interrupt.
3	R/W	Overflow
2	R/W	Negative. The sign bit was set as a result of the last ALU instruction.
1	R/W	Carry
0	R/W	Zero. The last ALU operation produced a zero.

Table 2.1: Condition Code Register Bit Assignment

Bit	Meaning
9	Soft trap, set on a trap from user mode, cleared when returning to user mode
8	Illegal instruction error flag
7	Halt on break, to support an external debugger
6	Step, single step the CPU in user mode
5	GIE, or Global Interrupt Enable
4	Sleep
3	V, or overflow bit.
2	N, or negative bit.
1	C, or carry bit.
0	Z, or zero bit.

Table 2.2: Condition Code / Status Register Bits

sleep register and clearing the GIE register at the same time, with clearing the GIE register taking precedence.

The seventh bit is a step bit. This bit can be set from supervisor mode only. After setting this bit, should the supervisor mode process switch to user mode, it would then accomplish one instruction in user mode before returning to supervisor mode. Then, upon return to supervisor mode, this bit will be automatically cleared. This bit has no effect on the CPU while in supervisor mode.

This functionality was added to enable a userspace debugger functionality on a user process, working through supervisor mode of course.

The eighth bit is a break enable bit. This controls whether a break instruction in user mode will halt the processor for an external debugger (break enabled), or whether the break instruction will simply send the CPU into interrupt mode. Encountering a break in supervisor mode will halt the CPU independent of the break enable bit. This bit can only be set within supervisor mode.

This functionality was added to enable an external debugger to set and manage breakpoints.

The ninth bit is an illegal instruction bit. When the CPU tries to execute either a non-existent instruction, or an instruction from an address that produces a bus error, the CPU will (if implemented) switch to supervisor mode while setting this bit. The bit will automatically be cleared upon any return to user mode.

The tenth bit is a trap bit. It is set whenever the user requests a soft interrupt, and cleared on any return to userspace command. This allows the supervisor, in supervisor mode, to determine whether it got to supervisor mode from a trap or from an external interrupt or both.

These status register bits are summarized in Tbl. 2.2.

2.3 Conditional Instructions

Most, although not quite all, instructions may be conditionally executed. From the four condition code flags, eight conditions are defined. These are shown in Tbl. 2.3. There is no condition code for less than or equal, not C or not V. Sorry, I ran out of space in 3-bits. Conditioning on a non-supported condition is still possible, but it will take an extra instruction and a pipeline stall. (Ex: *Stall*); `TST $4, CC`; `STO.NZ R0, (R1)`) As an alternative, the condition may often be reversed, recov-

Code	Mnemonic	Condition
3'h0	None	Always execute the instruction
3'h1	.Z	Only execute when 'Z' is set
3'h2	.NE	Only execute when 'Z' is not set
3'h3	.GE	Greater than or equal ('N' not set, 'Z' irrelevant)
3'h4	.GT	Greater than ('N' not set, 'Z' not set)
3'h5	.LT	Less than ('N' set)
3'h6	.C	Carry set
3'h7	.V	Overflow set

Table 2.3: Conditions for conditional operand execution

```

CMP 1,R0
;Condition codes are now set based upon R0-1
CMP.Z 2,R1
;If R0 ≠ 1, conditions are unchanged.
;If R0 = 1, conditions are set based upon R1-2.
;Now do something based upon the conjunction of both conditions.
;While we use the example of a STO, it could be any instruction.
STO.Z R0,(R2)

```

Table 2.4: An example of a double conditional

ering those extra two clocks. Thus instead of `CMP Rx,Ry; BNW label` you can issue a `CMP Ry,Rx; BV label`.

Conditionally executed ALU instructions will not further adjust the condition codes, with the exception of `CMP` and `TST` instructions. Conditional `CMP` or `TST` instructions will adjust conditions whenever their conditionals are true. In this way, multiple conditions may be evaluated without branches. For example, to do something if `R0` is one and `R1` is two, one might try code such as Tbl. 2.4.

2.4 Traditional Interrupt Handling

2.5 Operand B

Many instruction forms have a 21-bit source “Operand B” associated with them. This Operand B is either equal to a register plus a signed immediate offset, or an immediate offset by itself. This value is encoded as shown in Tbl. 2.5.

Sixteen and twenty bit immediate values don’t make sense for all instructions. For example, what is the point of a 20-bit immediate when executing a 16-bit multiply? Likewise, why have a 16-bit immediate when adding to a logical or arithmetic shift? In these cases, the extra bits are reserved for future instruction possibilities.

Bit 20	19 ... 16	15 ... 0
1'b0	20-bit Signed Immediate value	
1'b1	4-bit Register	16-bit Signed immediate offset

Table 2.5: Bit allocation for Operand B

2.6 Address Modes

The Zip CPU supports two addressing modes: register plus immediate, and immediate address. Addresses are therefore encoded in the same fashion as Operand B's, shown above.

A lot of long hard thought was put into whether to allow pre/post increment and decrement addressing modes. Finding no way to use these operators without taking two or more clocks per instruction,¹ these addressing modes have been removed from the realm of possibilities. This means that the Zip CPU has no native way of executing push, pop, return, or jump to subroutine operations. Each of these instructions can be emulated with a set of instructions from the existing set.

2.7 Move Operands

The previous set of operands would be perfect and complete, save only that the CPU needs access to non-supervisory registers while in supervisory mode. Therefore, the MOV instruction is special and offers access to these registers ... when in supervisory mode. To keep the compiler simple, the extra bits are ignored in non-supervisory mode (as though they didn't exist), rather than being mapped to new instructions or additional capabilities. The bits indicating which register set each register lies within are the A-Usr and B-Usr bits. When set to a one, these refer to a user mode register. When set to a zero, these refer to a register in the current mode, whether user or supervisor. Further, because a load immediate instruction exists, there is no move capability between an immediate and a register: all moves come from either a register or a register plus an offset.

This actually leads to a bit of a problem: since the MOV instruction encodes which register set each register is coming from or moving to, how shall a compiler or assembler know how to compile a MOV instruction without knowing the mode of the CPU at the time? For this reason, the compiler will assume all MOV registers are supervisor registers, and display them as normal. Anything with the user bit set will be treated as a user register. The CPU will quietly ignore the supervisor bits while in user mode, and anything marked as a user register will always be valid.

2.8 Multiply Operations

The Zip CPU supports two Multiply operations, a 16x16 bit signed multiply (MPYS) and a 16x16 bit unsigned multiply (MPYU). In both cases, the operand is a register plus a 16-bit immediate, subject to the rule that the register cannot be the PC or CC registers. The PC register field has been stolen to create a multiply by immediate instruction. The CC register field is reserved.

¹The two clocks figure comes from the design of the register set, allowing only one write per clock. That write is either from the memory unit or the ALU, but never both.

2.9 Floating Point

The Zip CPU does not (yet) support floating point operations. However, the instruction set reserves two possibilities for future floating point operations.

The first floating point operation hole in the instruction set involves setting a proposed (but non-existent) floating point bit in the CC register. The next instruction would then simply interpret its operands as floating point instructions. Not all instructions, however, have floating point equivalents. Further, the immediate fields do not apply in floating point mode, and must be set to zero. Not all instructions make sense as floating point operations. Therefore, only the CMP, SUB, ADD, and MPY instructions may be issued as floating point instructions. Other instructions allow the examining of the floating point bit in the CC register. In all cases, the floating point bit is cleared one instruction after it is set.

The other possibility for floating point operations involves exploiting the hole in the instruction set that the NOOP and BREAK instructions reside within. These two instructions use 24-bits of address space, when only a single bit is necessary. A simple adjustment to this space could create instructions with 4-bit register addresses for each register, a 3-bit field for conditional execution, and a 2-bit field for which operation. In this fashion, such a floating point capability would only fill 13-bits of the 24-bit field, still leaving lots of room for expansion.

In both cases, the Zip CPU would support 32-bit single precision floats only, since other choices would complicate the pipeline.

The current architecture does not support a floating point not-implemented interrupt. Any soft floating point emulation must be done deliberately.

2.10 Native Instructions

The instruction set for the Zip CPU is summarized in Tbl. 2.6.

As you can see, there's lots of room for instruction set expansion. The NOOP and BREAK instructions are the only instructions within one particular 24-bit hole. The rest of this space is reserved for future enhancements.

2.11 Derived Instructions

The Zip CPU supports many other common instructions, but not all of them are single cycle instructions. The derived instruction tables, Tbls. 2.7, 2.8, 2.9 and 2.10, help to capture some of how these other instructions may be implemented on the Zip CPU. Many of these instructions will have assembly equivalents, such as the branch instructions, to facilitate working with the CPU.

2.12 Pipeline Stages

As mentioned in the introduction, and highlighted in Fig. 1.1, the Zip CPU supports a five stage pipeline.

1. **Prefetch:** Reads instruction from memory and into a cache, if so configured. This stage is actually pipelined itself, and so it will stall if the PC ever changes. Stalls are also created here if the instruction isn't in the prefetch cache.

Op Code	31...24		23...16	15...8	7...0	Sets CC?	
CMP (Sub)	4'h0	D. Reg	Cond.	Operand B			Yes
TST (And)	4'h1	D. Reg	Cond.	Operand B			Yes
MOV	4'h2	D. Reg	Cond.	A-Usr	B-Reg	B-Usr	15'bit signed offset
LODI	4'h3	R. Reg	24'bit Signed Immediate				
NOOP	4'h4	4'he	24'h00				
BREAK	4'h4	4'he	24'h01				
<i>Reserved</i>	4'h4	4'he	24'bits, but not 0 or 1.				
LODIHI	4'h4	4'hf	Cond.	1'b1	R. Reg	16-bit Immediate	
LODILO	4'h4	4'hf	Cond.	1'b0	R. Reg	16-bit Immediate	
16-b MPYU	4'h4	R. Reg	Cond.	1'b0	Reg	16-bit Offset	
16-b MPYU(I)	4'h4	R. Reg	Cond.	1'b0	4'hf	16-bit Offset	
16-b MPYS	4'h4	R. Reg	Cond.	1'b1	Reg	16-bit Offset	
16-b MPYS(I)	4'h4	R. Reg	Cond.	1'b1	4'hf	16-bit Offset	
ROL	4'h5	R. Reg	Cond.	Operand B, truncated to low order 5 bits			
LOD	4'h6	R. Reg	Cond.	Operand B address			
STO	4'h7	D. Reg	Cond.	Operand B address			
SUB	4'h8	R. Reg	Cond.	Operand B			Yes
AND	4'h9	R. Reg	Cond.	Operand B			Yes
ADD	4'ha	R. Reg	Cond.	Operand B			Yes
OR	4'hb	R. Reg	Cond.	Operand B			Yes
XOR	4'hc	R. Reg	Cond.	Operand B			Yes
LSL/ASL	4'hd	R. Reg	Cond.	Operand B, imm. truncated to 6 bits			Yes
ASR	4'he	R. Reg	Cond.	Operand B, imm. truncated to 6 bits			Yes
LSR	4'hf	R. Reg	Cond.	Operand B, imm. truncated to 6 bits			Yes

Table 2.6: Zip CPU Instruction Set

Mapped	Actual	Notes
ABS Rx	TST -1,Rx	Absolute value, depends upon derived NEG.
ADD Ra,Rx ADDC Rb,Ry	NEG.LT Rx Add Ra,Rx ADD.C \$1,Ry Add Rb,Ry	Add with carry
BRA.Cond +/--\$Addr	MOV.cond \$Addr+PC,PC	Branch or jump on condition. Works for 15-bit signed address offsets.
BRA.Cond +/--\$Addr	LDI \$Addr,Rx ADD.cond Rx,PC	Branch/jump on condition. Works for 23 bit address offsets, but costs a register, an extra instruction, and sets the flags.
BNC PC+\$Addr	Test \$Carry,CC MOV.Z PC+\$Addr,PC	Example of a branch on an unsupported condition, in this case a branch on not carry
BUSY	MOV \$-1(PC),PC	Execute an infinite loop
CLRF.NZ Rx	XOR.NZ Rx,Rx	Clear Rx, and flags, if the Z-bit is not set
CLR Rx	LDI \$0,Rx	Clears Rx, leaves flags untouched. This instruction cannot be conditional.
EXCH.W Rx	ROL \$16,Rx	Exchanges the top and bottom 16-bit words of Rx
HALT	Or \$SLEEP,CC	This only works when issued in interrupt/supervisor mode. In user mode this is simply a wait until interrupt instruction.
INT	LDI \$0,CC	
IRET	OR \$GIE,CC	Also known as an RTU instruction (Return to Userspace)
JMP R6+\$Addr	MOV \$Addr(R6),PC	
JSR PC+\$Addr	SUB \$1,SP MOV \$3+PC,R0 STO R0,1(SP) MOV \$Addr+PC,PC ADD \$1,SP	Jump to Subroutine. Note the required cleanup instruction after returning. This could easily be turned into a three instruction operand, removing the preliminary stack instruction before and the cleanup after, by adjusting how any stack frame was built for this routine to include space at the top of the stack for the PC. Note also that jumping to a subroutine costs a copy register, R0 in this case.
JSR PC+\$Addr	MOV \$3+PC,R12 MOV \$addr+PC,PC	This is the high speed version of a subroutine call, necessitating a register to hold the last PC address. In its favor, this method doesn't suffer the mandatory memory access of the other approach.
LDI.l \$val,Rx	LDIHI (\$val>>16)&0xffff, Rx LDILO (\$val&0xffff),Rx	Sadly, there's not enough instruction space to load a complete immediate value into any register. Therefore, fully loading any register takes two cycles. The LDIHI (load immediate high) and LDILO (load immediate low) instructions have been created to facilitate this.

Table 2.7: Derived Instructions

Mapped	Actual	Notes
LOD.b \$addr, Rx	LDI \$addr, Ra LDI \$addr, Rb LSR \$2, Ra AND \$3, Rb LOD (Ra), Rx LSL \$3, Rb SUB \$32, Rb ROL Rb, Rx AND \$0ffh, Rx	This CPU is designed for 32'bit word length instructions. Byte addressing is not supported by the CPU or the bus, so it therefore takes more work to do. Note also that in this example, \$Addr is a byte-wise address, where all other addresses in this document are 32-bit wordlength addresses. For this reason, we needed to drop the bottom two bits. This also limits the address space of character accesses using this method from 16 MB down to 4MB.
LSL \$1, Rx LSLC \$1, Ry	LSL \$1, Ry LSL \$1, Rx OR.C \$1, Ry	Logical shift left with carry. Note that the instruction order is now backwards, to keep the conditions valid. That is, LSL sets the carry flag, so if we did this the other way with Rx before Ry, then the condition flag wouldn't have been right for an OR correction at the end.
LSR \$1, Rx LSRC \$1, Ry	CLR Rz LSR \$1, Ry LDIHI.C \$8000h, Rz LSR \$1, Rx OR Rz, Rx	Logical shift right with carry
NEG Rx	XOR \$-1, Rx ADD \$1, Rx	
NEG.C Rx	MOV.C \$-1+Rx, Rx XOR.C \$-1, Rx	
NOOP	NOOP	While there are many operations that do nothing, such as MOV Rx,Rx, or OR \$0,Rx, these operations have consequences in that they might stall the bus if Rx isn't ready yet. For this reason, we have a dedicated NOOP instruction.
NOT Rx	XOR \$-1, Rx	
POP Rx	LOD \$1(SP), Rx ADD \$1, SP	Note that for interrupt purposes, one can never depend upon the value at (SP). Hence you read from it, then increment it, lest having incremented it first something then comes along and writes to that value before you can read the result.

Table 2.8: Derived Instructions, continued

PUSH Rx	SUB \$1,SP STO Rx,\$1(SP)	Note that for pipelined operation, it helps to coalesce all the SUB's into one command, and place the STO's right after each other.
PUSH Rx-Ry	SUB \$n,SP STO Rx,\$n(SP) ... STO Ry,\$1(SP)	Multiple pushes at once only need the single subtract from the stack pointer. This derived instruction is analogous to a similar one on the Motorola 68k architecture, although the Zip Assembler does not support this instruction (yet). This instruction also supports pipelined memory access.
RESET	STO \$1,\$watchdog(R12) NOOP NOOP	This depends upon the peripheral base address being in R12. Another opportunity might be to jump to the reset address from within supervisor mode.
RET	LOD \$1(SP),PC	Note that this depends upon the calling context to clean up the stack, as outlined for the JSR instruction.
RET	MOV R12,PC	This is the high(er) speed version, that doesn't touch the stack. As such, it doesn't suffer a stall on memory read/write to the stack.
STEP Rr,Rt	LSR \$1,Rr XOR.C Rt,Rr	Step a Galois implementation of a Linear Feedback Shift Register, Rr, using taps Rt
STO.b Rx,\$addr	LDI \$addr,Ra LDI \$addr,Rb LSR \$2,Ra AND \$3,Rb SUB \$32,Rb LOD (Ra),Ry AND \$0ffh,Rx AND ~\$0ffh,Ry ROL Rb,Rx OR Rx,Ry STO Ry,(Ra)	This CPU and it's bus are <i>not</i> optimized for byte-wise operations. Note that in this example, \$addr is a byte-wise address, whereas in all of our other examples it is a 32-bit word address. This also limits the address space of character accesses from 16 MB down to 4MB.F Further, this instruction implies a byte ordering, such as big or little endian.
SWAP Rx,Ry	XOR Ry,Rx XOR Rx,Ry XOR Ry,Rx	While no extra registers are needed, this example does take 3-clocks.
TRAP #X	LDI \$x,RO AND ~\$GIE,CC	This works because whenever a user lowers the \$GIE flag, it sets a TRAP bit within the CC register. Therefore, upon entering the supervisor state, the CPU only need check this bit to know that it got there via a TRAP. The trap could be made conditional by making the LDI and the AND conditional. In that case, the assembler would quietly turn the LDI instruction into an LDILO and LDIHI pair, but the effect would be the same.

Table 2.9: Derived Instructions, continued

TST Rx	TST \$-1,Rx	Set the condition codes based upon Rx. Could also do a CMP \$0,Rx, ADD \$0,Rx, SUB \$0,Rx, etc, AND \$-1,Rx, etc. The TST and CMP approaches won't stall future pipeline stages looking for the value of Rx.
WAIT	Or \$GIE \$SLEEP,CC	Wait until the next interrupt, then jump to supervisor/interrupt mode.

Table 2.10: Derived Instructions, continued

The Zip CPU supports one of two prefetch methods, depending upon a flag set at build time within the `zipcpu.v` file. The simplest is a non-cached implementation of a prefetch. This implementation is fairly small, and ideal for users of the Zip CPU who need the extra space on the FPGA fabric. However, because this non-cached version has no cache, the maximum number of instructions per clock is limited to about one per five.

The second prefetch module is a pipelined prefetch with a cache. This module tries to keep the instruction address within a window of valid instruction addresses. While effective, it is not a traditional cache implementation. One unique feature of this cache implementation, however, is that it can be cleared in a single clock. A disappointing feature, though, was that it needs an extra internal pipeline stage to be implemented.

2. **Decode:** Decodes an instruction into op code, register(s) to read, and immediate offset. This stage also determines whether the flags will be set or whether the result will be written back.
3. **Read Operands:** Read registers and apply any immediate values to them. There is no means of detecting or flagging arithmetic overflow or carry when adding the immediate to the operand. This stage will stall if any source operand is pending.
4. Split into two tracks: An **ALU** which will accomplish a simple instruction, and the **MemOps** stage which handles LOD (load) and STO (store) instructions.
 - Loads will stall the entire pipeline until complete.
 - Condition codes are available upon completion of the ALU stage
 - Issuing an instruction to the memory unit while the memory unit is busy will stall the entire pipeline. If the bus deadlocks, only a reset will release the CPU. (Watchdog timer, anyone?)
 - The Zip CPU currently has no means of reading and acting on any error conditions on the bus.
5. **Write-Back:** Conditionally write back the result to the register set, applying the condition. This routine is bi-entrant: either the memory or the simple instruction may request a register write.

The Zip CPU does not support out of order execution. Therefore, if the memory unit stalls, every other instruction stalls. Memory stores, however, can take place concurrently with ALU operations, although memory reads (loads) cannot.




Figure 2.2: A conditional branch generates 5 stall cycles

2.13 Pipeline Stalls

The processing pipeline can and will stall for a variety of reasons. Some of these are obvious, some less so. These reasons are listed below:

- When the prefetch cache is exhausted
This reason should be obvious. If the prefetch cache doesn't have the instruction in memory, the entire pipeline must stall until enough of the prefetch cache is loaded to support the next instruction.
- While waiting for the pipeline to load following any taken branch, jump, return from interrupt or switch to interrupt context (5 stall cycles)

Fig. 2.2 illustrates the situation for a conditional branch. In this case, the branch instruction, BC, is nominally followed by instructions I0 and so forth. However, since the branch is taken, the next instruction must be IA. Therefore, the pipeline needs to be cleared and reloaded. Given that there are five stages to the pipeline, that accounts for four of the five stalls. The last stall cycle is lost in the pipelined prefetch stage which needs at least one clock with a valid PC before it can produce a new output. **Note: When I did this myself, I counted six stall cycles, for a total of seven cycles for this instruction. Is five really the right answer?**

The Zip CPU handles MOV \$X(PC), PC, ADD \$X, PC, and LDI \$X, PC instructions specially, however. These instructions, when not conditioned on the flags, can execute with only 2 stall cycles, such as is shown in Fig. 2.3.² In this example, BR is a branch always taken, I1 is the instruction following the branch in memory, while IA is the first instruction at the branch address. (CLR denotes a clear-pipeline operation, and does not represent any instruction.)

- When reading from a prior register while also adding an immediate offset

²Note that this behavior is slated to be improved upon in subsequent releases. With a better prefetch, it should be possible to drop this down to one or zero stall cycles.




Figure 2.3: An expedited delay costs only 2 stall cycles

1. OPCODE ?,RA
2. (*stall*)
3. OPCODE I+RA, RB

Since the addition of the immediate register within OpB decoding gets applied during the read operand stage so that it can be nicely settled before the ALU, any instruction that will write back an operand must be separated from the opcode that will read and apply an immediate offset by one instruction. The good news is that this stall can easily be mitigated by proper scheduling. That is, any instruction that does not add an immediate to RA may be scheduled into the stall slot.

- When any (conditional) write to either the CC or PC Register is followed by a memory operation
 1. OPCODE RA,PC *Ex: a branch opcode*
 2. (*stall, even if jump not taken*)
 3. LOD \$X(RA),RB

A timing diagram of this pipeline situation is shown in Fig. 2.4, for a conditional branch, BC, a memory operation, Mem (which must be a load here), and ALU instructions I1 and so forth. Since branches take place in the writeback stage, the Zip CPU will stall the pipeline for one clock anytime there may be a possible jump—forcing the memory operation to stay in the operand decode stage. This prevents an instruction from executing a memory access after the jump but before the jump is recognized.

This stall may be mitigated by shuffling the operations immediately following a potential branch so that an ALU operation follows the branch instead of a memory operation.

- When reading from the CC register after setting the flags
 1. ALUOP RA,RB *Ex: a compare opcode*

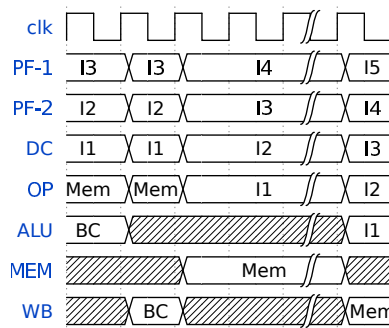


Figure 2.4: A (not taken) conditional branch followed by a memory operation

2. *(stall)*
3. `TST sys.ccv,CC`
4. `BZ somewhere`

The reason for this stall is simply performance: many of the flags are determined via combinatorial logic *during* the writeback cycle. Trying to then place these into the input for one of the operands for an ALU instruction during the same cycle created a time delay loop that would no longer execute in a single 100 MHz clock cycle. (The time delay of the multiply within the ALU wasn't helping either ...).

This stall may be eliminated via proper scheduling, by placing an instruction that does not set flags in between the ALU operation and the instruction that references the CC register. For example, `MOV $addr+PC,uPC` followed by an `RTU (OR $GIE,CC)` instruction will not incur this stall, whereas an `OR $BREAKEN,CC` followed by an `OR $STEP,CC` will incur the stall, while a `LDI $BREAKEN|$STEP,CC` will not since it doesn't read the condition codes.

- When waiting for a memory read operation to complete
 1. `LOD address,RA`
 2. *(multiple stalls, bus dependent, 4 clocks best)*
 3. `OPCODE I+RA,RB`

Remember, the Zip CPU does not support out of order execution. Therefore, anytime the memory unit becomes busy both the memory unit and the ALU must stall until the memory unit is cleared. This is illustrated in Fig. 2.5, since it is especially true of a load instruction, which must still write its operand back to the register file. Note that there is an extra stall at the end of the memory cycle, so that the memory unit will be idle for one clock before an instruction will be accepted into the ALU. Store instructions are different, as shown in Fig. 2.6, since they can be busy with the bus without impacting later write back pipeline stages. Hence, only loads stall the pipeline.

This, of course, also assumes that the memory being accessed is a single cycle memory and that there are no stalls to get to the memory. Slower memories, such as the Quad SPI flash, will

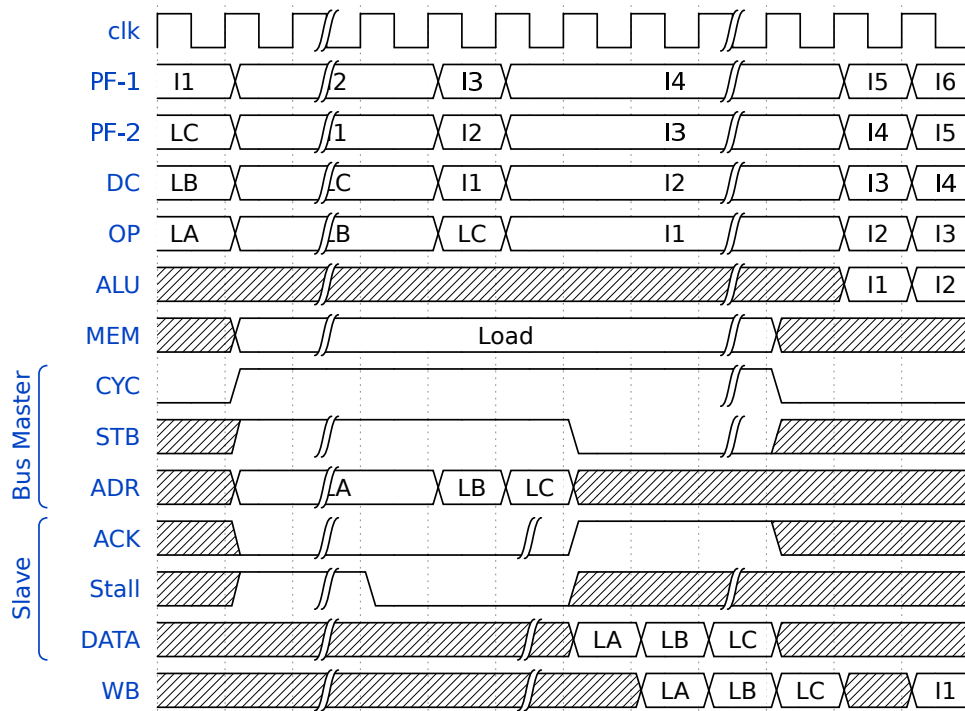


Figure 2.5: Pipeline handling of a load instruction

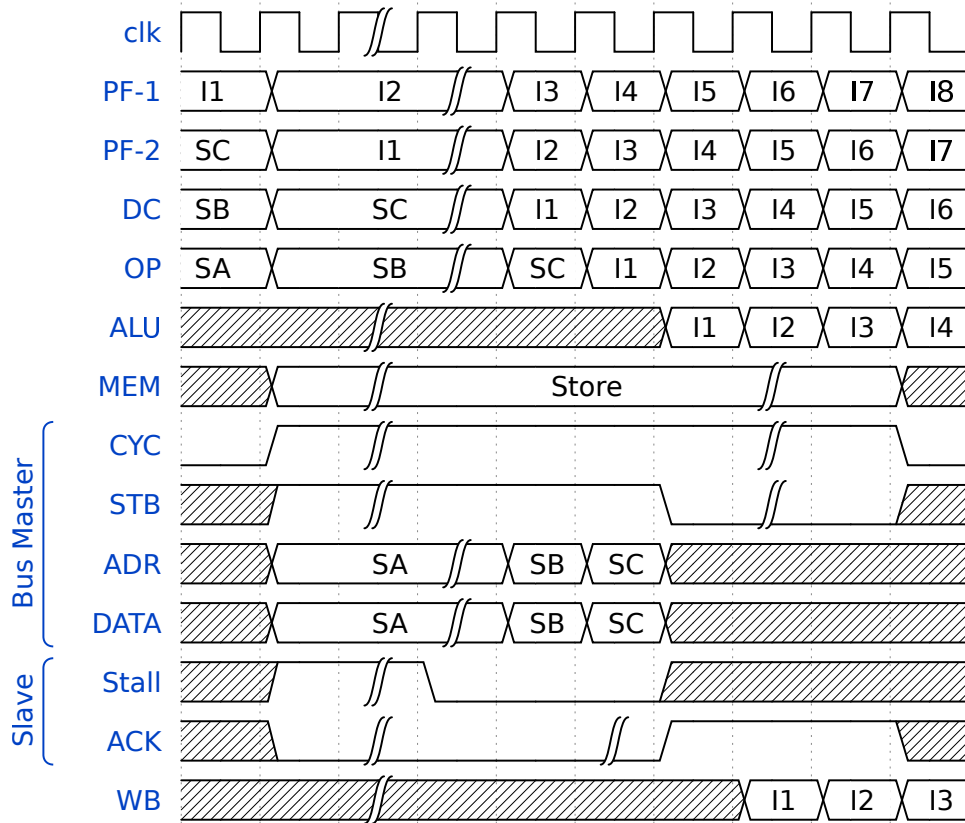


Figure 2.6: Pipeline handling of a store instruction

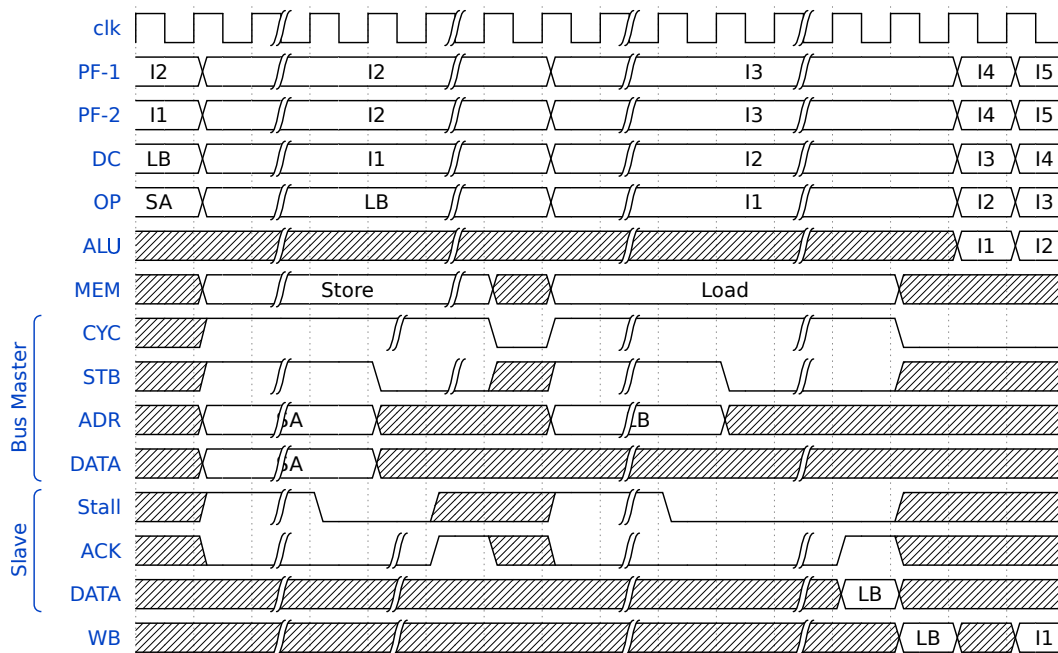


Figure 2.7: Pipeline handling of a store followed by a load instruction

take longer—perhaps even as long as forty clocks. During this time the CPU and the external bus will be busy, and unable to do anything else. Likewise, if it takes a couple of clock cycles for the bus to be free, as shown in both Figs. 2.5 and 2.6, there will be stalls.

- Memory operation followed by a memory operation
 1. `STO address,RA`
 2. *(multiple stalls, bus dependent, 4 clocks best)*
 3. `LOD address,RB`
 4. *(multiple stalls, bus dependent, 4 clocks best)*

In this case, the LOD instruction cannot start until the STO is finished, as illustrated by Fig. 2.7. With proper scheduling, it is possible to do something in the ALU while the memory unit is busy with the STO instruction, but otherwise this pipeline will stall while waiting for it to complete before the load instruction can start.

The Zip CPU does have the capability of supporting pipelined memory access, but only under the following conditions: all accesses within the pipeline must all be reads or all be writes, all must use the same register for their address, and there can be no stalls or other instructions between pipelined memory access instructions. Further, the offset to memory must be increasing by one address each instruction. These conditions work well for saving or storing registers

to the stack. Indeed, if you noticed, both Fig. 2.5 and Fig. 2.6 illustrated pipelined memory accesses.

- When waiting for a conditional memory read operation to complete
 1. LOD.Z address,RA
 2. *(multiple stalls, bus dependent, 7 clocks best)*
 3. OPCODE I+RA,RB

In this case, the Zip CPU doesn't warn the prefetch cache to get off the bus two cycles before using the bus, so there's a potential for an extra three cycle cost due to bus contention between the prefetch and the CPU.

This is true for both the LOD and the STO instructions, with the exception that the STO instruction will continue in parallel with any ALU instructions that follow it.

3.

Peripherals

While the previous chapter describes a CPU in isolation, the Zip System includes a minimum set of peripherals as well. These peripherals are shown in Fig. 3.1 and described here. They are designed to make the Zip CPU more useful in an Embedded Operating System environment.

3.1 Interrupt Controller

Perhaps the most important peripheral within the Zip System is the interrupt controller. While the Zip CPU itself can only handle one interrupt, and has only the one interrupt state: disabled or enabled, the interrupt controller can make things more interesting.

The Zip System interrupt controller module supports up to 15 interrupts, all controlled from one register. Bit 31 of the interrupt controller controls overall whether interrupts are enabled (1'b1) or disabled (1'b0). Bits 16–30 control whether individual interrupts are enabled (1'b1) or disabled (1'b0). Bit 15 is an indicator showing whether or not any interrupt is active, and bits 0–15 indicate whether or not an individual interrupt is active.

The interrupt controller has been designed so that bits can be controlled individually without having any knowledge of the rest of the controller setting. To enable an interrupt, write to the register with the high order global enable bit set and the respective interrupt enable bit set. No other bits will be affected. To disable an interrupt, write to the register with the high order global enable bit cleared and the respective interrupt enable bit set. To clear an interrupt, write a '1' to that interrupt's status pin. Zero's written to the register have no affect, save that a zero written to the master enable will disable all interrupts.

As an example, suppose you wished to enable interrupt #4. You would then write to the register a 0x80100010 to enable interrupt #4 and to clear any past active state. When you later wish to disable this interrupt, you would write a 0x00100010 to the register. As before, this both disables the interrupt and clears the active indicator. This also has the side effect of disabling all interrupts, so a second write of 0x80000000 may be necessary to re-enable any other interrupts.

The Zip System currently hosts two interrupt controllers, a primary and a secondary. The primary interrupt controller has one interrupt line (perhaps more if you configure it for more) which may come from an external interrupt controller, and one interrupt line from the secondary controller. Other primary interrupts include the system timers, the jiffies interrupt, and the manual cache interrupt. The secondary interrupt controller maintains an interrupt state for all of the processor accounting counters.

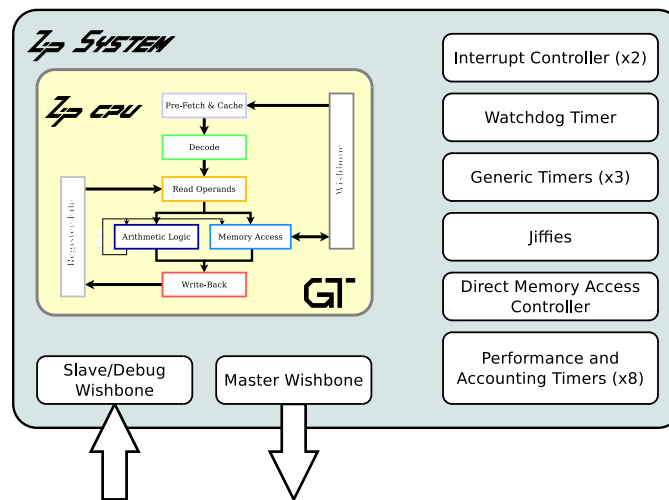


Figure 3.1: Zip System Peripherals

3.2 Counter

The Zip Counter is a very simple counter: it just counts. It cannot be halted. When it rolls over, it issues an interrupt. Writing a value to the counter just sets the current value, and it starts counting again from that value.

Eight counters are implemented in the Zip System for process accounting. This may change in the future, as nothing as yet uses these counters.

3.3 Timer

The Zip Timer is also very simple: it simply counts down to zero. When it transitions from a one to a zero it creates an interrupt.

Writing any non-zero value to the timer starts the timer. If the high order bit is set when writing to the timer, the timer becomes an interval timer and reloads its last start time on any interrupt. Hence, to mark seconds, one might set the timer to 100 million (the number of clocks per second), and set the high bit. Ever after, the timer will interrupt the CPU once per second (assuming a 100 MHz clock). This reload capability also limits the maximum timer value to $2^{31} - 1$ (about 21 seconds using a 100 MHz clock), rather than $2^{32} - 1$.

3.4 Watchdog Timer

The watchdog timer is no different from any of the other timers, save for one critical difference: the interrupt line from the watchdog timer is tied to the reset line of the CPU. Hence writing a '1' to the watchdog timer will always reset the CPU. To stop the Watchdog timer, write a '0' to it. To start it, write any other number to it—as with the other timers.

While the watchdog timer supports interval mode, it doesn't make as much sense as it did with the other timers.

3.5 Bus Watchdog

There is an additional watchdog timer on the Wishbone bus. This timer, however, is hardware configured and not software configured. The timer is reset at the beginning of any bus transaction, and only counts clocks during such bus transactions. If the bus transaction takes longer than the number of counts the timer allots, it will raise a bus error flag to terminate the transaction. This is useful in the case of any peripherals that are misbehaving. If the bus watchdog terminates a bus transaction, the CPU may then read from its port to find out which memory location created the problem.

Aside from its unusual configuration, the bus watchdog is just another implementation of the fundamental timer described above.

3.6 Jiffies

This peripheral is motivated by the Linux use of 'jiffies' whereby a process can request to be put to sleep until a certain number of 'jiffies' have elapsed. Using this interface, the CPU can read the number of 'jiffies' from the peripheral (it only has the one location in address space), add the sleep length to it, and write the result back to the peripheral. The zipjiffies peripheral will record the value written to it only if it is nearer the current counter value than the last current waiting interrupt time. If no other interrupts are waiting, and this time is in the future, it will be enabled. (There is currently no way to disable a jiffie interrupt once set, other than to disable the interrupt line in the interrupt controller.) The processor may then place this sleep request into a list among other sleep requests. Once the timer expires, it would write the next Jiffy request to the peripheral and wake up the process whose timer had expired.

Indeed, the Jiffies register is nothing more than a glorified counter with an interrupt. Unlike the other counters, the Jiffies register cannot be set. Writes to the jiffies register create an interrupt time. When the Jiffies register later equals the value written to it, an interrupt will be asserted and the register then continues counting as though no interrupt had taken place.

The purpose of this register is to support alarm times within a CPU. To set an alarm for a particular process N clocks in advance, read the current Jiffies value, and N , and write it back to the Jiffies register. The O/S must also keep track of values written to the Jiffies register. Thus, when an 'alarm' trips, it should be removed from the list of alarms, the list should be sorted, and the next alarm in terms of Jiffies should be written to the register.

3.7 Direct Memory Access Controller

The Direct Memory Access (DMA) controller can be used to either move memory from one location to another, to read from a peripheral into memory, or to write from a peripheral into memory all without CPU intervention. Further, since the DMA controller can issue (and does issue) pipeline wishbone accesses, any DMA memory move will by nature be faster than a corresponding program accomplishing the same move. To put this to numbers, it may take a program 18 clocks per word

transferred, whereas this DMA controller can move one word in two clocks—provided it has bus access. (The CPU gets priority over the bus.)

When copying memory from one location to another, the DMA controller will copy in units of a given transfer length—up to 1024 words at a time. It will read that transfer length into its internal buffer, and then write to the destination address from that buffer. If the CPU interrupts a DMA transfer, it will release the bus, let the CPU complete whatever it needs to do, and then restart its transfer by writing the contents of its internal buffer and then re-entering its read cycle again.

When coupled with a peripheral, the DMA controller can be configured to start a memory copy on an interrupt line going high. Further, the controller can be configured to issue reads from (or to) the same address instead of incrementing the address at each clock. The DMA completes once the total number of items specified (not the transfer length) have been transferred.

In each case, once the transfer is complete and the DMA unit returns to idle, the DMA will issue an interrupt.

4.

Operation

The Zip CPU, and even the Zip System, is not a System on a Chip (SoC). It needs to be connected to its operational environment in order to be used. Specifically, some per system adjustments need to be made:

1. The Zip System depends upon an external 32-bit Wishbone bus. This must exist, and must be connected to the Zip CPU for it to work.
2. The Zip System needs to be told of its `RESET_ADDRESS`. This is the program counter of the first instruction following a reset.
3. If you want the Zip System to start up on its own, you will need to set the `START_HALTED` parameter to zero. Otherwise, if you wish to manually start the CPU, that is if upon reset you want the CPU start start in its halted, reset state, then set this parameter to one.
4. The third parameter to set is the number of interrupts you will be providing from external to the CPU. This can be anything from one to nine, but it cannot be zero. (Wire this line to a 1'b0 if you do not wish to support any external interrupts.)
5. Finally, you need to place into some wishbone accessible address, whether RAM or (more likely) ROM, the initial instructions for the CPU.

If you have enabled your CPU to start automatically, then upon power up the CPU will immediately start executing your instructions.

This is, however, not how I have used the Zip CPU. I have instead used the Zip CPU in a more controlled environment. For me, the CPU starts in a halted state, and waits to be told to start. Further, the `RESET` address is a location in RAM. After bringing up the board I am using, and further the bus that is on it, the RAM memory is then loaded externally with the program I wish the Zip System to run. Once the RAM is loaded, I release the CPU. The CPU then runs until its halt condition, at which point its task is complete.

Eventually, I intend to place an operating system onto the ZipSystem, I'm just not there yet.

The rest of this chapter examines some common programming models, and how they might be applied to the Zip System, and then finish with a couple of examples.

4.1 System High

The easiest and simplest way to run the Zip CPU is in the system high mode. In this mode, the CPU runs your program in supervisor mode from reboot to power down, and is never interrupted.

```
supervisor_idle:
    ; While not strictly required, the following move helps to
    ; ensure that the prefetch doesn't try to fetch an instruction
    ; outside of the CPU's address space when it switches to user
    ; mode.
    MOV supervisor_idle_continue,uPC
    ; Put the processor into user mode and to sleep in the same
    ; instruction.
    OR $SLEEP|$GIE,CC
supervisor_idle_continue:
    ; Now, if we haven't done this inline, we need to return
    ; to whatever function called us.
    RETN
```

Table 4.1: Executing an idle from supervisor mode

You will need to poll the interrupt controller to determine when any external condition has become active. This mode is useful, and can handle many microcontroller tasks.

Even better, in system high mode, all of the user registers are available to the system high program as variables. Accessing these registers can be done in a single clock cycle, which would move them to the active register set or move them back. While this may seem like a load or store instruction, none of these register accesses will suffer from memory delays.

The one thing that cannot be done in supervisor mode is a wait for interrupt instruction. This, however, is easily rectified by jumping to a user task within the supervisors memory space, such as Tbl. 4.1.

4.2 Traditional Interrupt Handling

Although the Zip CPU does not have a traditional interrupt architecture, it is possible to create the more traditional interrupt approach via software. In this mode, the programmable interrupt controller is used together with the supervisor state to create the illusion of more traditional interrupt handling.

To set this up, upon reboot the supervisor task:

1. Creates a (single) user context, a user stack, and sets the user program counter to the entry of the user task
2. Creates a task table of ISR entries
3. Enables the master interrupt enable via the interrupt controller, albeit without enabling any of the fifteen potential underlying interrupts.
4. Switches to user mode, as the first part of the while loop in Tbl. 4.2.

```

while(true) {
    rtu();
    if (trap) { // Here, we allow users to install ISRs, or
                // whatever else they may wish to do in supervisor mode.
    } else {
        volatile int *pic = PIC_ADDRESS;

        // Save the user context before running any ISRs. This could easily be
        // implemented as an inline assembly routine or macro
        SAVE_PARTIAL_CONTEXT;
        // At this point, we know an interrupt has taken place: Ask the programmable
        // interrupt controller (PIC) which interrupts are enabled and which are active.
        int picv = *pic;
        // Turn off all active interrupts
        // Globally disable interrupt generation in the process
        int active = (picv >> 16) & picv & 0x07fff;
        *pic = (active<<16);
        // We build a mask of interrupts to re-enable in picv.
        picv = 0;
        for(int i=0,msk=1; i<15; i++, msk<<=1) {
            if ((active & msk)&&(isr_table[i])) {
                mov(isr_table[i],uPC);
                // Acknowledge this particular interrupt. While we could acknowledge all
                // interrupts at once, by acknowledging only those with ISR's we allow
                // the user process to use peripherals manually, and to manually check
                // whether or no those other interrupts had occurred.
                *pic = msk;
                rtu();
                // The ISR will only exit on a trap in the Zip architecture. There is
                // no RETI instruction. Since the PIC holds all interrupts disabled,
                // there is no need to check for further interrupts.
                //
                // The tricky part is that, because of how the PIC is built, the ISR cannot
                // re-enable its own interrupt without re-enabling all interrupts. Hence, we
                // look at R0 upon ISR completion to know if an interrupt needs to be
                // re-enabled.
                mov(uR0,tmp);
                picv |= (tmp & 0x7fff) << 16;
            }
        }
        RESTORE_PARTIAL_CONTEXT;
        // Re-activate all (requested) interrupts
        *pic = picv | 0x80000000;
    }
}

```

Table 4.2: Traditional Interrupt handling

```
SAVE_PARTIAL_CONTEXT:
; We save R0, CC, and PC only
MOV  -3(uSP),R3
MOV  uR0,R0
MOV  uCC,R1
MOV  uPC,R2
STO  R0,1(R3) ; Exploit memory pipelining:
STO  R1,2(R3) ; All instructions write to stack
STO  R2,3(R3) ; All offsets increment by one
MOV  R3,uSP ; Return the updated stack pointer
```

Table 4.3: Example Saving Minimal User Context

We can work through the interrupt handling process by examining Tbl. 4.2. First, remember, the CPU is always running either the user or the supervisor context. Once the supervisor switches to user mode, control does not return until either an interrupt or a trap has taken place. (Okay, there's also the possibility of a bus error, or an illegal instruction such as an unimplemented floating point instruction—but for now we'll just focus on the trap instruction.) Therefore, if the trap bit isn't set, then we know an interrupt has taken place.

To process an interrupt, we steal the user's stack: the PC and CC registers are saved on the stack, as outlined in Tbl. 4.3. This is much cheaper than the full context swap of a preemptive multitasking kernel, but it also depends upon the ISR saving any state it uses. Further, if multiple ISR's get called at once, this loses its optimality property very quickly.

As Sec. 3.1 discusses, the top of the PIC register stores which interrupts are enabled, and the bottom stores which have tripped. (Interrupts may trip without being enabled, they just will not generate an interrupt to the CPU.) Our first step is to query the register to find out our interrupt state, and then to disable any interrupts that have tripped. To do that, we write a one to the enable half of the register while also clearing the top bit (master interrupt enable). This has the consequence of disabling any and all further interrupts, not just the ones that have tripped. Hence, upon completion, we re-enable the master interrupt bit again. Finally, we keep track of which interrupts have tripped.

Using the bit mask of interrupts that have tripped, we walk through all fifteen possible interrupts. If there is an ISR installed, we acknowledge and reset the interrupt within the PIC, and then call the ISR. The ISR, however, cannot re-enable its interrupt without re-enabling the master interrupt bit. Thus, to keep things simple, when the ISR is finished it places its interrupt mask back into R0, or clears R0. This tells the supervisor mode process which interrupts to re-enable. Any other registers that the ISR uses must be saved and restored. (This is only truly optimal if only a single ISR is called.) As a final instruction, the ISR clears the GIE bit executing a user trap. (Remember, the Zip CPU has no RETI instruction to restore the stack and return to userland. It needs to go through the supervisor mode to get there.)

Then, once all interrupts are handled, the user context is restored in a fashion similar to Tbl. 4.4. Again, this is short and sweet simply because any other registers that needed saving were saved within the ISR.


```

RESTORE_PARTIAL_CONTEXT:
; We restore R0, CC, and PC only
MOV uSP,R3 ; Return the updated stack pointer
LOD R0,1(R3),R0 ; Exploit memory pipelining:
LOD R1,2(R3),R1 ; All instructions write to stack
LOD R2,3(R3),R2 ; All offsets increment by one
MOV R0,uR0
MOV R1,uCC
MOV R2,uPC
MOV 3(R3),uSP

```

Table 4.4: Example Restoring Minimal User Context

```

idle_task:
; Wait for the next interrupt, then switch to supervisor task
WAIT
; When we come back, it's because the supervisor wishes to
; wait for an interrupt again, so go back to the top.
BRA idle_task

```

Table 4.5: Example Idle Loop

There you have it: the Zip CPU, with its non-traditional interrupt architecture, can still process interrupts in a very traditional fashion.

4.3 Example: Idle Task

One task every operating system needs is the idle task, the task that takes place when nothing else can run. On the Zip CPU, this task is quite simple, and it is shown in assemble in Tbl. 4.5. When this task runs, the CPU will fill up all of the pipeline stages up the ALU. The `WAIT` instruction, upon leaving the ALU, places the CPU into a sleep state where nothing more moves. Sure, there may be some more settling, the pipe cache continue to read until full, other instructions may issue until the pipeline fills, but then everything will stall. Then, once an interrupt takes place, control passes to the supervisor task to handle the interrupt. When control passes back to this task, it will be on the next instruction. Since that next instruction sends us back to the top of the task, the idle task thus does nothing but wait for an interrupt.

This should be the lowest priority task, the task that runs when nothing else can. It will help lower the FPGA power usage overall—at least its dynamic power usage.

```

void memcpy(void *dest, void *src, int len) {
    for(int i=0; i<len; i++)
        *dest++ = *src++;
}

```

Table 4.6: Example Memory Copy code in C

```

memcpy:
    ; R0 = *dest, R1 = *src, R2 = LEN
    ; The following will operate in 17 clocks per word minus one clock
    CMP 0,R2
    LOD.Z -1(SP),PC ; A conditional return
    ; (One stall on potentially writing to PC)
    LOD (R1),R3
    ; (4 stalls, cannot be scheduled away)
    STO R3,(R2) ; (4 schedulable stalls, has no impact now)
    ADD 1,R1
    SUB 1,R2
    BNZ loop
    ; (5 stalls, if branch taken, to clear and refill the pipeline)
    RET

```

Table 4.7: Example Memory Copy code in Zip Assembly

4.4 Example: Memory Copy

One common operation is that of a memory move or copy. Consider the C code shown in Tbl. 4.6. This same code can be translated in Zip Assembly as shown in Tbl. 4.7. This example points out several things associated with the Zip CPU. First, a straightforward implementation of a for loop is not the fastest loop structure. For this reason, we have placed the test to continue at the end. Second, all pointers are void pointers to arbitrary 32-bit data types. The Zip CPU does not have explicit support for smaller or larger data types, and so this memory copy cannot be applied at a byte level. Third, we've optimized the conditional jump to a return instruction into a conditional return instruction.

4.5 Example: Context Switch

Fundamental to any multiprocessing system is the ability to switch from one task to the next. In the ZipSystem, this is accomplished in one of a couple ways. The first step is that an interrupt happens. Anytime an interrupt happens, the CPU needs to execute the following tasks in supervisor mode:

1. Check for a trap instruction. That is, if the user task requested a trap, we may not wish to adjust the context, check interrupts, or call the scheduler. Tbl. 4.8 shows the rudiments of

```

return_to_user:
    ; The instruction before the context switch processing must
    ; be the RTU instruction that enacted user mode in the first
    ; place. We show it here just for reference.
    RTU

trap_check:
    MOV uCC,R0
    TST $TRAP,R0
    BNZ swap_out
    ; Do something here to execute the trap
    ; Don't need to call the scheduler, so we can just return
    BRA return_to_user

```

Table 4.8: Checking for whether the user issued a TRAP instruction

this code, while showing nothing of how the actual trap would be implemented.

You may also wish to note that the instruction before the first instruction in our context swap *must be* a return to userspace instruction. Remember, the supervisor process is re-entered where it left off. This is different from many other processors that enter interrupt mode at some vector or other. In this case, we always enter supervisor mode right where we last left.¹

2. Capture user counters. If the operating system is keeping track of system usage via the accounting counters, those counters need to be copied and accumulated into some master counter at this point.
3. Preserve the old context. This involves pushing all the user registers onto the user stack and then copying the resulting stack address into the tasks task structure, as shown in Tbl. 4.9. For the sake of discussion, we assume the supervisor maintains a pointer to the current task's structure in supervisor register R12, and that `stack` is an offset to the beginning of this structure indicating where the stack pointer is to be kept within it.
For those who are still interested, the full code for this context save can be found as an assembler macro within the assembler include file, `sys.i`.
4. Reset the watchdog timer. If you are using the watchdog timer, it should be reset on a context swap, to know that things are still working. Example code for this is shown in Tbl. 4.10.
5. Interrupt handling. An interrupt handler within the Zip System is nothing more than a task. At context swap time, the supervisor needs to disable all of the interrupts that have tripped, and then enable all of the tasks that would deal with each of these interrupts. These can be user tasks, run at higher priority than any other user tasks. Either way, they will need to re-enable their own interrupt themselves, if the interrupt is still relevant.

An example of this master interrupt handling is shown in Tbl. 4.11.

¹The one exception to this rule is upon reset where supervisor mode is entered at a pre-programmed wishbone memory address.

```
swap_out:
    MOV -15(uSP),R5
    STO R5,stack(R12)
    MOV uR0,R0
    MOV uR1,R1
    MOV uR2,R2
    MOV uR3,R3
    MOV uR4,R4
    STO R0,1(R5) ; Exploit memory pipelining:
    STO R1,2(R5) ; All instructions write to stack
    STO R2,3(R5) ; All offsets increment by one
    STO R3,4(R5) ; Longest pipeline is 5 cycles.
    STO R4,5(R5)
    ... ; Need to repeat for all user registers
    MOV uR10,R0
    MOV uR11,R1
    MOV uR12,R2
    MOV uCC,R3
    MOV uPC,R4
    STO R0,11(R5)
    STO R1,12(R5)
    STO R2,13(R5)
    STO R3,14(R5)
    STO R4,15(R5)
    ; We can skip storing the stack, uSP, since it'll be stored
    ; elsewhere (in the task structure)
```

Table 4.9: Example Storing User Task Context

```
'define WATCHDOG_ADDRESS 32'hc000_0002
'define WATCHDOG_TICKS 32'd1_000_000 ; = 10 ms
    LDI WATCHDOG_ADDRESS,R0
    LDI WATCHDOG_TICKS,R1
    STO R1,(R0)
```

Table 4.10: Example Watchdog Reset

pre_handler:

```
LDI PIC_ADDRESS,R0
; Start by grabbing the interrupt state from the interrupt
; controller. We'll store this into the register R7 so that
; we can keep and preserve this information for the scheduler
; to use later.
LOD (R0),R1
MOV R1,R7
; As a next step, we need to acknowledge and disable all active
; interrupts. We'll start by calculating all of our active
; interrupts.
AND 0x07fff,R1
; Put the active interrupts into the upper half of R1
ROL 16,R1
LDILO 0x0ffff,R1
AND R7,R1
; Acknowledge and disable active interrupts
; This also disables all interrupts from the controller, so
; we'll need to re-enable interrupts in general shortly
STO R1,(R0)
; We leave our active interrupt mask in R7 so the scheduler can
; release any tasks that depended upon them.
```

Table 4.11: Example checking for active interrupts

6. Calling the scheduler. This needs to be done to pick the next task to switch to. It may be an interrupt handler, or it may be a normal user task. From a priority standpoint, it would make sense that the interrupt handlers all have a higher priority than the user tasks, and that once they have been called the user tasks may then be called again. If no task is ready to run, run the idle task to wait for an interrupt.

This suggests a minimum of four task priorities:

- (a) Interrupt handlers, executed with their interrupts disabled
- (b) Device drivers, executed with interrupts re-enabled
- (c) User tasks
- (d) The idle task, executed when nothing else is able to execute

For our purposes here, we'll just assume that a pointer to the current task is maintained in R12, that a JSR `scheduler` is called, and that the next current task is likewise placed into R12.

7. Restore the new tasks context. Given that the scheduler has returned a task that can be run at this time, the stack pointer needs to be pulled out of the tasks task structure, placed into the user register, and then the rest of the user registers need to be popped back off of the stack to run this task. An example of this is shown in Tbl. 4.12, assuming as before that the task pointer is found in supervisor register R12. As with storing the user context, the full code associated with restoring the user context can be found in the assembler include file, `sys.i`.
8. Clear the userspace accounting registers. In order to keep track of per process system usage, these registers need to be cleared before reactivating the userspace process. That way, upon the next interrupt, we'll know how many clocks the userspace program has encountered, and how many instructions it was able to issue in those many clocks.
9. Jump back to the instruction just before saving the last tasks context, because that location in memory contains the return from interrupt command that we are going to need to execute, in order to guarantee that we return back here again.

```
swap_in:
    LOD stack(R12),R5
    MOV 15(R1),uSP
    ; Be sure to exploit the memory pipelining capability
    LOD 1(R5),R0
    LOD 2(R5),R1
    LOD 3(R5),R2
    LOD 4(R5),R3
    LOD 5(R5),R4
    MOV R0,uR0
    MOV R1,uR1
    MOV R2,uR2
    MOV R3,uR3
    MOV R4,uR4
    ... ; Need to repeat for all user registers
    LOD 11(R5),R0
    LOD 12(R5),R1
    LOD 13(R5),R2
    LOD 14(R5),R3
    LOD 15(R5),R4
    MOV R0,uR10
    MOV R1,uR11
    MOV R2,uR12
    MOV R3,uCC
    MOV R4,uPC
    BRA return_to_user
```

Table 4.12: Example Restoring User Task Context

5.

Registers

The ZipSystem registers fall into two categories, ZipSystem internal registers accessed via the ZipCPU shown in Tbl. 5.1, and the two debug registers shown in Tbl. 5.2.

Name	Address	Width	Access	Description
PIC	0xc0000000	32	R/W	Primary Interrupt Controller
WDT	0xc0000001	32	R/W	Watchdog Timer
	0xc0000002	32	R/W	<i>(Reserved for future use)</i>
CTRIC	0xc0000003	32	R/W	Secondary Interrupt Controller
TMRA	0xc0000004	32	R/W	Timer A
TMRB	0xc0000005	32	R/W	Timer B
TMRC	0xc0000006	32	R/W	Timer C
JIFF	0xc0000007	32	R/W	Jiffies
MTASK	0xc0000008	32	R/W	Master Task Clock Counter
MMSTL	0xc0000009	32	R/W	Master Stall Counter
MPSTL	0xc000000a	32	R/W	Master Pre-Fetch Stall Counter
MICNT	0xc000000b	32	R/W	Master Instruction Counter
UTASK	0xc000000c	32	R/W	User Task Clock Counter
UMSTL	0xc000000d	32	R/W	User Stall Counter
UPSTL	0xc000000e	32	R/W	User Pre-Fetch Stall Counter
UICNT	0xc000000f	32	R/W	User Instruction Counter
DMACTRL	0xc0000010	32	R/W	DMA Control Register
DMALEN	0xc0000011	32	R/W	DMA total transfer length
DMA_SRC	0xc0000012	32	R/W	DMA source address
DMA_DST	0xc0000013	32	R/W	DMA destination address

Table 5.1: Zip System Internal/Peripheral Registers

Name	Address	Width	Access	Description
ZIPCTRL	0	32	R/W	Debug Control Register
ZIPDATA	1	32	R/W	Debug Data Register

Table 5.2: Zip System Debug Registers

Bit #	Access	Description
31	R/W	Master Interrupt Enable
30...16	R/W	Interrupt Enables, write '1' to change
15	R	Current Master Interrupt State
15...0	R/W	Input Interrupt states, write '1' to clear

Table 5.3: Interrupt Controller Register Bits

5.1 Peripheral Registers

The peripheral registers, listed in Tbl. 5.1, are shown in the CPU's address space. These may be accessed by the CPU at these addresses, and when so accessed will respond as described in Chapt. 3. These registers will be discussed briefly again here.

The Zip CPU Interrupt controller has four different types of bits, as shown in Tbl. 5.3. The high order bit, or bit-31, is the master interrupt enable bit. When this bit is set, then any time an interrupt occurs the CPU will be interrupted and will switch to supervisor mode, etc.

Bits 30 ... 16 are interrupt enable bits. Should the interrupt line go ghile while enabled, an interrupt will be generated. To set an interrupt enable bit, one needs to write the master interrupt enable while writing a '1' to this the bit. To clear, one need only write a '0' to the master interrupt enable, while leaving this line high.

Bits 15 ... 0 are the current state of the interrupt vector. Interrupt lines trip when they go high, and remain tripped until they are acknowledged. If the interrupt goes high for longer than one pulse, it may be high when a clear is requested. If so, the interrupt will not clear. The line must go low again before the status bit can be cleared.

As an example, consider the following scenario where the Zip CPU supports four interrupts, 3...0.

1. The Supervisor will first, while in the interrupts disabled mode, write a 32'h800f000f to the controller. The supervisor may then switch to the user state with interrupts enabled.
2. When an interrupt occurs, the supervisor will switch to the interrupt state. It will then cycle through the interrupt bits to learn which interrupt handler to call.
3. If the interrupt handler expects more interrupts, it will clear its current interrupt when it is done handling the interrupt in question. To do this, it will write a '1' to the low order interrupt mask, such as writing a 32'h80000001.
4. If the interrupt handler does not expect any more interrupts, it will instead clear the interrupt from the controller by writing a 32'h00010001 to the controller.
5. Once all interrupts have been handled, the supervisor will write a 32'h80000000 to the interrupt register to re-enable interrupt generation.
6. The supervisor should also check the user trap bit, and possible soft interrupt bits here, but this action has nothing to do with the interrupt control register.

Bit #	Access	Description
31	R/W	Auto-Reload
30...0	R/W	Current timer value

Table 5.4: Timer Register Bits

Bit #	Access	Description
31...0	R	Current jiffy value
31...0	W	Value/time of next interrupt

Table 5.5: Jiffies Register Bits

- The supervisor will then leave interrupt mode, possibly adjusting whichever task is running, by executing a return from interrupt command.

Leaving the interrupt controller, we show the timer registers bit definitions in Tbl. 5.4. As you may recall, the timer just counts down to zero and then trips an interrupt. Writing to the current timer value sets that value, and reading from it returns that value. Writing to the current timer value while also setting the auto-reload bit will send the timer into an auto-reload mode. In this mode, upon setting its interrupt bit for one cycle, the timer will also reset itself back to the value of the timer that was written to it when the auto-reload option was written to it. To clear and stop the timer, just simply write a '32'h00' to this register.

The Jiffies register is somewhat similar in that the register always changes. In this case, the register counts up, whereas the timer always counted down. Reads from this register, as shown in Tbl. 5.5, always return the time value contained in the register. Writes greater than the current Jiffy value, that is where the new value minus the old value is greater than zero while ignoring truncation, will set a new Jiffy interrupt time. At that time, the Jiffy vector will clear, and another interrupt time may either be written to it, or it will just continue counting without activating any more interrupts.

The Zip CPU also supports several counter peripherals, mostly in the way of process accounting. This peripherals have a single register associated with them, shown in Tbl. 5.6. Writes to this register set the new counter value. Reads read the current counter value.

The current design operation of these counters is that of performance counting. Two sets of four registers are available for keeping track of performance. The first is a task counter. This just counts clock ticks. The second counter is a prefetch stall counter, then an master stall counter. These allow the CPU to be evaluated as to how efficient it is. The fourth and final counter is an instruction counter, which counts how many instructions the CPU has issued.

Bit #	Access	Description
31...0	R/W	Current counter value

Table 5.6: Counter Register Bits

Bit #	Access	Description
31	R	DMA Active
30	R	Wishbone error, transaction aborted. This bit is cleared the next time this register is written to.
29	R/W	Set to '1' to prevent the controller from incrementing the source address, '0' for normal memory copy.
28	R/W	Set to '1' to prevent the controller from incrementing the destination address, '0' for normal memory copy.
27 ... 16	W	The DMA Key. Write a 12'hfed to these bits to start the activate any DMA transfer.
27	R	Always reads '0', to force the deliberate writing of the key.
26 ... 16	R	Indicates the number of items in the transfer buffer that have yet to be written.
15	R/W	Set to '1' to trigger on an interrupt, or '0' to start immediately upon receiving a valid key.
14 ... 10	R/W	Select among one of 32 possible interrupt lines.
9 ... 0	R/W	Intermediate transfer length minus one. Thus, to transfer one item at a time set this value to 0. To transfer 1024 at a time, set it to 1024.

Table 5.7: DMA Control Register Bits

It is envisioned that these counters will be used as follows: First, every time a master counter rolls over, the supervisor (Operating System) will record the fact. Second, whenever activating a user task, the Operating System will set the four user counters to zero. When the user task has completed, the Operating System will read the timers back off, to determine how much of the CPU the process had consumed.

The final peripheral to discuss is the DMA controller. This controller has four registers. Of these four, the length, source and destination address registers should need no further explanation. They are full 32-bit registers specifying the entire transfer length, the starting address to read from, and the starting address to write to. The registers can be written to when the DMA is idle, and read at any time. The control register, however, will need some more explanation.

The bit allocation of the control register is shown in Tbl. 5.7. This control register has been designed so that the common case of memory access need only set the key and the transfer length. Hence, writing a 32'h0fed03ff to the control register will start any memory transfer. On the other hand, if you wished to read from a serial port (constant address) and put the result into a buffer every time a word was available, you might wish to write 32'h2fed8000—this assumes, of course, that you have a serial port wired to the zero bit of this interrupt control. (The DMA controller does not use the interrupt controller, and cannot clear interrupts.) As a third example, if you wished to write to an external FIFO anytime it was less than half full (had fewer than 512 items), and interrupt line 2 indicated this condition, you might wish to issue a 32'h1fed8dff to this port.

Bit #	Access	Description
31...14	R	Reserved
13	R	CPU GIE setting
12	R	CPU is sleeping
11	W	Command clear PF cache
10	R/W	Command HALT, Set to '1' to halt the CPU
9	R	Stall Status, '1' if CPU is busy
8	R/W	Step Command, set to '1' to step the CPU, also sets the halt bit
7	R	Interrupt Request
6	R/W	Command RESET
5...0	R/W	Debug Register Address

Table 5.8: Debug Control Register Bits

5.2 Debug Port Registers

Accessing the Zip System via the debug port isn't as straight forward as accessing the system via the wishbone bus. The debug port itself has been reduced to two addresses, as outlined earlier in Tbl. 5.2. Access to the Zip System begins with the Debug Control register, shown in Tbl. 5.8.

The first step in debugging access is to determine whether or not the CPU is halted, and to halt it if not. To do this, first write a '1' to the Command HALT bit. This will halt the CPU and place it into debug mode. Once the CPU is halted, the stall status bit will drop to zero. Thus, if bit 10 is high and bit 9 low, the debug port is open to examine the internal state of the CPU.

At this point, the external debugger may examine internal state information from within the CPU. To do this, first write again to the command register a value (with command halt still high) containing the address of an internal register of interest in the bottom 6 bits. Internal registers that may be accessed this way are listed in Tbl. 5.9. Primarily, these "registers" include access to the entire CPU register set, as well as the internal peripherals. To read one of these registers once the address is set, simply issue a read from the data port. To write one of these registers or peripheral ports, simply write to the data port after setting the proper address.

In this manner, all of the CPU's internal state may be read and adjusted.

As an example of how to use this, consider what would happen in the case of an external break point. If and when the CPU hits a break point that causes it to halt, the Command HALT bit will activate on its own, the CPU will then raise an external interrupt line and wait for a debugger to examine its state. After examining the state, the debugger will need to remove the breakpoint by writing a different instruction into memory and by writing to the command register while holding the clear cache, command halt, and step CPU bits high, (32'hd00). The debugger may then replace the breakpoint now that the CPU has gone beyond it, and clear the cache again (32'h500).

To leave this debug mode, simply write a '32'h0' value to the command register.

Name	Address	Width	Access	Description
sR0	0	32	R/W	Supervisor Register R0
sR1	0	32	R/W	Supervisor Register R1
sSP	13	32	R/W	Supervisor Stack Pointer
sCC	14	32	R/W	Supervisor Condition Code Register
sPC	15	32	R/W	Supervisor Program Counter
uR0	16	32	R/W	User Register R0
uR1	17	32	R/W	User Register R1
uSP	29	32	R/W	User Stack Pointer
uCC	30	32	R/W	User Condition Code Register
uPC	31	32	R/W	User Program Counter
PIC	32	32	R/W	Primary Interrupt Controller
WDT	33	32	R/W	Watchdog Timer
CTRIC	35	32	R/W	Secondary Interrupt Controller
TMRA	36	32	R/W	Timer A
TMRB	37	32	R/W	Timer B
TMRC	38	32	R/W	Timer C
JIFF	39	32	R/W	Jiffies peripheral
MTASK	40	32	R/W	Master task clock counter
MMSTL	41	32	R/W	Master memory stall counter
MPSTL	42	32	R/W	Master Pre-Fetch Stall counter
MICNT	43	32	R/W	Master instruction counter
UTASK	44	32	R/W	User task clock counter
UMSTL	45	32	R/W	User memory stall counter
UPSTL	46	32	R/W	User Pre-Fetch Stall counter
UICNT	47	32	R/W	User instruction counter
DMACMD	48	32	R/W	DMA command and status register
DMALEN	49	32	R/W	DMA transfer length
DMARD	50	32	R/W	DMA read address
DMAWR	51	32	R/W	DMA write address

Table 5.9: Debug Register Addresses

6.

Wishbone Datasheets

The Zip System supports two wishbone ports, a slave debug port and a master port for the system itself. These are shown in Tbl. 6.1 and Tbl. 6.2 respectively. I do not recommend that you connect

Description	Specification																				
Revision level of wishbone	WB B4 spec																				
Type of interface	Slave, Read/Write, single words only																				
Address Width	1-bit																				
Port size	32-bit																				
Port granularity	32-bit																				
Maximum Operand Size	32-bit																				
Data transfer ordering	(Irrelevant)																				
Clock constraints	Works at 100 MHz on a Basys-3 board																				
Signal Names	<table border="1"> <thead> <tr> <th>Signal Name</th> <th>Wishbone Equivalent</th> </tr> </thead> <tbody> <tr> <td>i_clk</td> <td>CLK_I</td> </tr> <tr> <td>i_dbg_cyc</td> <td>CYC_I</td> </tr> <tr> <td>i_dbg_stb</td> <td>STB_I</td> </tr> <tr> <td>i_dbg_we</td> <td>WE_I</td> </tr> <tr> <td>i_dbg_addr</td> <td>ADR_I</td> </tr> <tr> <td>i_dbg_data</td> <td>DAT_I</td> </tr> <tr> <td>o_dbg_ack</td> <td>ACK_O</td> </tr> <tr> <td>o_dbg_stall</td> <td>STALL_O</td> </tr> <tr> <td>o_dbg_data</td> <td>DAT_O</td> </tr> </tbody> </table>	Signal Name	Wishbone Equivalent	i_clk	CLK_I	i_dbg_cyc	CYC_I	i_dbg_stb	STB_I	i_dbg_we	WE_I	i_dbg_addr	ADR_I	i_dbg_data	DAT_I	o_dbg_ack	ACK_O	o_dbg_stall	STALL_O	o_dbg_data	DAT_O
	Signal Name	Wishbone Equivalent																			
	i_clk	CLK_I																			
	i_dbg_cyc	CYC_I																			
	i_dbg_stb	STB_I																			
	i_dbg_we	WE_I																			
	i_dbg_addr	ADR_I																			
	i_dbg_data	DAT_I																			
	o_dbg_ack	ACK_O																			
	o_dbg_stall	STALL_O																			
o_dbg_data	DAT_O																				

Table 6.1: Wishbone Datasheet for the Debug Interface

these together through the interconnect. Rather, the debug port of the CPU should be accessible regardless of the state of the master bus.

You may wish to notice that neither the `ERR` nor the `RETRY` wires have been implemented. What this means is that the CPU is currently unable to detect a bus error condition, and so may stall indefinitely (hang) should it choose to access a value not on the bus, or a peripheral that is not yet properly configured.

Description	Specification																				
Revision level of wishbone	WB B4 spec																				
Type of interface	Master, Read/Write, single cycle or pipelined																				
Address Width	32-bit bits																				
Port size	32-bit																				
Port granularity	32-bit																				
Maximum Operand Size	32-bit																				
Data transfer ordering	(Irrelevant)																				
Clock constraints	Works at 100 MHz on a Basys-3 board																				
Signal Names	<table border="1"> <thead> <tr> <th>Signal Name</th> <th>Wishbone Equivalent</th> </tr> </thead> <tbody> <tr> <td>i_clk</td> <td>CLK_0</td> </tr> <tr> <td>o_wb_cyc</td> <td>CYC_0</td> </tr> <tr> <td>o_wb_stb</td> <td>STB_0</td> </tr> <tr> <td>o_wb_we</td> <td>WE_0</td> </tr> <tr> <td>o_wb_addr</td> <td>ADR_0</td> </tr> <tr> <td>o_wb_data</td> <td>DAT_0</td> </tr> <tr> <td>i_wb_ack</td> <td>ACK_I</td> </tr> <tr> <td>i_wb_stall</td> <td>STALL_I</td> </tr> <tr> <td>i_wb_data</td> <td>DAT_I</td> </tr> </tbody> </table>	Signal Name	Wishbone Equivalent	i_clk	CLK_0	o_wb_cyc	CYC_0	o_wb_stb	STB_0	o_wb_we	WE_0	o_wb_addr	ADR_0	o_wb_data	DAT_0	i_wb_ack	ACK_I	i_wb_stall	STALL_I	i_wb_data	DAT_I
	Signal Name	Wishbone Equivalent																			
	i_clk	CLK_0																			
	o_wb_cyc	CYC_0																			
	o_wb_stb	STB_0																			
	o_wb_we	WE_0																			
	o_wb_addr	ADR_0																			
	o_wb_data	DAT_0																			
	i_wb_ack	ACK_I																			
i_wb_stall	STALL_I																				
i_wb_data	DAT_I																				

Table 6.2: Wishbone Datasheet for the CPU as Master

7.

Clocks

This core is based upon the Basys-3 development board sold by Digilent. The Basys-3 development board contains one external 100 MHz clock, which is sufficient to run the Zip CPU core. I hesitate

Name	Source	Rates (MHz)		Description
		Max	Min	
i_clk	External	100 MHz	100 MHz	System clock.

Table 7.1: List of Clocks

to suggest that the core can run faster than 100 MHz, since I have had struggled with various timing violations to keep it at 100 MHz. So, for now, I will only state that it can run at 100 MHz.

8.

I/O Ports

The I/O ports to the Zip CPU may be grouped into three categories. The first is that of the master wishbone used by the CPU, then the slave wishbone used to command the CPU via a debugger, and then the rest. The first two of these were already discussed in the wishbone chapter. They are listed here for completeness in Tbl. 8.1 and 8.2 respectively.

There are only four other lines to the CPU: the external clock, external reset, incoming external interrupt line(s), and the outgoing debug interrupt line. These are shown in Tbl. 8.3. The clock line was discussed briefly in Chapt. 7. We typically run it at 100 MHz. The reset line is an active high reset. When asserted, the CPU will start running again from its reset address in memory. Further, depending upon how the CPU is configured and specifically on the `START_HALTED` parameter, it may or may not start running automatically. The `i_ext_int` line is for an external interrupt. This line may be as wide as 6 external interrupts, depending upon the setting of the `EXTERNAL_INTERRUPTS` line. As currently configured, the ZipSystem only supports one such interrupt line by default. For us, this line is the output of another interrupt controller, but that's a board specific setup detail. Finally, the Zip System produces one external interrupt whenever the CPU halts to wait for the debugger.

Port	Width	Direction	Description
<code>o_wb_cyc</code>	1	Output	Indicates an active Wishbone cycle
<code>o_wb_stb</code>	1	Output	WB Strobe signal
<code>o_wb_we</code>	1	Output	Write enable
<code>o_wb_addr</code>	32	Output	Bus address
<code>o_wb_data</code>	32	Output	Data on WB write
<code>i_wb_ack</code>	1	Input	Slave has completed a R/W cycle
<code>i_wb_stall</code>	1	Input	WB bus slave not ready
<code>i_wb_data</code>	32	Input	Incoming bus data

Table 8.1: CPU Master Wishbone I/O Ports

Port	Width	Direction	Description
i_wb_cyc	1	Input	Indicates an active Wishbone cycle
i_wb_stb	1	Input	WB Strobe signal
i_wb_we	1	Input	Write enable
i_wb_addr	1	Input	Bus address, command or data port
i_wb_data	32	Input	Data on WB write
o_wb_ack	1	Output	Slave has completed a R/W cycle
o_wb_stall	1	Output	WB bus slave not ready
o_wb_data	32	Output	Incoming bus data

Table 8.2: CPU Debug Wishbone I/O Ports

Port	Width	Direction	Description
i_clk	1	Input	The master CPU clock
i_rst	1	Input	Active high reset line
i_ext_int	1...6	Input	Incoming external interrupts
o_ext_int	1	Output	CPU Halted interrupt

Table 8.3: I/O Ports

9.

Initial Assessment

Having now worked with the Zip CPU for a while, it is worth offering an honest assessment of how well it works and how well it was designed. At the end of this assessment, I will propose some changes that may take place in a later version of this Zip CPU to make it better.

9.1 The Good

- The Zip CPU is light weight and fully featured as it exists today. For anyone who wishes to build a general purpose CPU and then to experiment with building and adding particular features, the Zip CPU makes a good starting point—it is fairly simple. Modifications should be simple enough.
- The Zip CPU was designed to be an implementable soft core that could be placed within an FPGA, controlling actions internal to the FPGA. It fits this role rather nicely. It does not fit the role of a system on a chip very well, but then it was never intended to be a system on a chip but rather a system within a chip.
- The extremely simplified instruction set of the Zip CPU was a good choice. Although it does not have many of the commonly used instructions, PUSH, POP, JSR, and RET among them, the simplified instruction set has demonstrated an amazing versatility. I will contend therefore and for anyone who will listen, that this instruction set offers a full and complete capability for whatever a user might wish to do with two exceptions: bitwise character access and accelerated floating-point support.
- This simplified instruction set is easy to decode.
- The simplified bus transactions (32-bit words only) were also very easy to implement.
- The pipelined load/store approach is novel, and can be used to greatly increase the speed of the processor.
- The novel approach of having a single interrupt vector, which just brings the CPU back to the instruction it left off at within the last interrupt context doesn't appear to have been that much of a problem. If most modern systems handle interrupt vectoring in software anyway, why maintain hardware support for it?
- My goal of a high rate of instructions per clock may not be the proper measure. For example, if instructions are being read from a SPI flash device, such as is common among FPGA

implementations, these same instructions may suffer stalls of between 64 and 128 cycles per instruction just to read the instruction from the flash. Executing the instruction in a single clock cycle is no longer the appropriate measure. At the same time, it should be possible to use the DMA peripheral to copy instructions from the FLASH to a temporary memory location, after which they may be executed at a single instruction cycle per access again.

9.2 The Not so Good

- While one of the stated goals was to use a small amount of logic, 3k LUTs isn't that impressively small. Indeed, it's really much too expensive when compared against other 8 and 16-bit CPUs that have less than 1k LUTs.

Still, . . . it's not bad, it's just not astonishingly good.

- The fact that the instruction width equals the bus width means that the instruction fetch cycle will always be interfering with any load or store memory operation, with the only exception being if the instruction is already in the cache.

This could be fixed in one of three ways: the instruction set architecture could be modified to handle Very Long Instruction Words (VLIW) so that each 32-bit word would encode two or more instructions, the instruction fetch bus width could be increased from 32-bits to 64-bits or more, or the instruction bus could be separated from the data bus. Any and all of these approaches would increase the overall LUT count.

- The (non-existent) floating point unit was an after-thought, isn't even built as a potential option, and most likely won't support the full IEEE standard set of FPU instructions—even for single point precision. This (non-existent) capability would benefit the most from an out-of-order execution capability, which the Zip CPU does not have.

Still, sharing FPU registers with the main register set was a good idea and worth preserving, as it simplifies context swapping.

Perhaps this really isn't a problem, but rather a feature. By not implementing FPU instructions, the Zip CPU maintains a lower LUT count than it would have if it did implement these instructions.

- The CPU has no character support. This is both good and bad. Realistically, the CPU works just fine without it. Characters can be supported as subsets of 32-bit words without any problem. Practically, though, it will make compiling non-Zip CPU code difficult—especially anything that assumes `sizeof(int)=4*sizeof(char)`, or that tries to create unions with characters and integers and then attempts to reference the address of the characters within that union.
- The Zip CPU does not support a data cache. One can still be built externally, but this is a limitation of the CPU proper as built. Further, under the theory of the Zip CPU design (that of an embedded soft-core processor within an FPGA, where any “address” may reference either memory or a peripheral that may have side-effects), any data cache would need to be based upon an initial knowledge of whether or not it is supporting memory (cachable) or peripherals. This knowledge must exist somewhere, and that somewhere is currently (and by design) external to the CPU.

This may also be written off as a “feature” of the Zip CPU, since the addition of a data cache can greatly increase the LUT count of a soft core.

The Zip CPU compensates for this via its pipelined load and store instructions.

- Many other instruction sets offer three operand instructions, whereas the Zip CPU only offers two operand instructions. This means that it takes the Zip CPU more instructions to do many of the same operations. The good part of this is that it gives the Zip CPU a greater amount of flexibility in its immediate operand mode, although that increased flexibility isn’t necessarily as valuable as one might like.
- The Zip CPU does not currently detect and trap on either illegal instructions or bus errors. Attempts to access non-existent memory quietly return erroneous results, rather than halting the process (user mode) or halting or resetting the CPU (supervisor mode).
- The Zip CPU doesn’t support out of order execution. I suppose it could be modified to do so, but then it would no longer be the “simple” and low LUT count CPU it was designed to be. The two primary results are that 1) loads may unnecessarily stall the CPU, even if other things could be done while waiting for the load to complete, 2) bus errors on stores will never be caught at the point of the error, and 3) branch prediction becomes more difficult.
- Although switching to an interrupt context in the Zip CPU design doesn’t require a tremendous swapping of registers, in reality it still does—since any task swap still requires saving and restoring all 16 user registers. That’s a lot of memory movement just to service an interrupt.
- The Zip CPU is by no means generic: it will never handle addresses larger than 32-bits (16GB) without a complete and total redesign. This may limit its utility as a generic CPU in the future, although as an embedded CPU within an FPGA this isn’t really much of a limit or restriction.
- While the Zip CPU has its own assembler, it has no linker and does not (yet) support a compiler. The standard C library is an even longer shot. My dream of having binutils and gcc support has not been realized and at this rate may not be realized. (I’ve been intimidated by the challenge everytime I’ve looked through those codes.)

9.3 The Next Generation

This section could also be labeled as my “To do” list.

Given the feedback listed above, perhaps its time to consider what changes could be made to improve the Zip CPU in the future. I offer the following as proposals:

- **Remove the low LUT goal.** It wasn’t really achieved, and the proposals below will only increase the amount of logic the Zip CPU requires. While I expect that the Zip CPU will always be somewhat of a light weight, it will never be the smallest kid on the block.

I’m actually struggling with this idea. The whole goal of the Zip CPU was to be light weight. Wouldn’t it make more sense to create and maintain options whereby it would remain lightweight? For example, if the process accounting registers are anything but light weight, why keep them? Why not instead make some compile flags that just turn them off, keeping the CPU lightweight? The same holds for the prefetch cache.

- The ‘.V’ condition was never used in any code other than my test code. Suggest changing it to a ‘.LE’ condition, which seems to be more useful.
- **Consider a more traditional Instruction Cache.** The current pipelined instruction cache just reads a window of memory into its cache. If the CPU leaves that window, the entire cache is invalidated. A more traditional cache, however, might allow common subroutines to stay within the cache without invalidating the entire cache structure.
- **Very Long Instruction Word (VLIW).** The goal here would be to create a new instruction set whereby two instructions would be encoded in each 32-bit word. While this may speed up CPU operation, it would necessitate an instruction redesign.