

CS/EE 3710 — Computer Architecture Lab

Checkpoint #2 — Datapath Infrastructure

Overview

In order to complete the datapath for your *insert-name-here* machine, the register file and ALU that you designed in checkpoint 1 needs a little support. In particular, the extra registers that were hinted at in checkpoint 1 need to be fleshed out and the initial memory interface needs to be completed. In order to have a complete datapath, you need to make the program counter complete, the MAR (memory address register) and MDR (memory data register) need to be specified, the memory access process needs to be figured out (at least the initial version that uses block RAM on the FPGA), and other registers like the instruction register and immediate register need to be instantiated, as well as sign extenders, etc. Once this is done, the remaining tasks are the instruction decoding, the control state machine, and (the biggie) figuring out what support you need for I/O for your application.

Program Counter

The program counter is a dedicated special register in the machine that holds the address of the next instruction to execute. It needs to be capable of being updated in every way the PC needs to be updated. For your machine, this means that the PC needs to be incremented by one word (the normal case), added to a (sign-extended) displacement (for branches) or loaded from a register (for jumps). Your datapath needs to be able to perform all of these operations. If your PC is already set up to feed into your ALU, then you could do the branch displacement calculation by setting some input muxes so that the PC and the immediate go to the ALU and the ALU function is set to add. You might load from a register by setting the ALU such that the appropriate register source makes it through the ALU without modification. This value can then be latched into the PC. For the increment case, you could either put the PC through one side of the ALU, and select a constant 1 for the other argument (put a constant value on one of the input muxes or something similar to that approach), or you could build your PC as a loadable counter. If you use the counter approach you can load the counter for the update and displacement functions, and count the counter for the increment-the-pc function. The choice is yours. The advantage of the counter is that you may not have to use the complete datapath for each PC increment, the advantage of the increment-through-the-alu approach is that every pc-update function goes through the same process, but with different mux settings. Remember that when you use the ALU to update the PC, you should *not* update the condition codes! Finally, remember that for a JAL instruction, the PC needs to have a path into the register file so that PC+1 (i.e. the address of the next instruction following the JAL) can be stored in the link register. Your datapath must allow this operation.

Another issue with the PC has to do with signed and unsigned arithmetic. Recall that the signed arithmetic is all done with two's complement numbers. This means that the range of numbers in a 16-bit word is -32,768 to 32,767. On the other hand, if you use those 16 bits to encode an unsigned number, you can represent 0 to 65,534 (64k). Since addresses are usually considered unsigned numbers, we need to consider what it means to have an unsigned PC that is operated on by a two's complement ALU, especially in the face of signed offsets that might require subtraction! Consider what happens if you try to use two's complement signed arithmetic to take a large unsigned number (large enough so the high-order-bit is 1) and add a negative signed number to it to try to subtract something (so the high-order bit is also one here since it is a negative signed number). Does it work? Are there constraints on when it works? I recommend trying this out on some smaller numbers to get a feel for what's going on. Trying this on 4- or 5-bit numbers is a good way to test things out to make sure you understand what's happening.

Note that our PC is addressing 16-bit words and not bytes! I think the best way to think about it is that with 16 bits of PC you can directly address 64k locations. Each of those locations is a 16-bit word. If you really want to think about it as bytes then the address space (without playing any tricks with segment

registers or things like that to increase the address space) is 128k bytes, and the PC 15:0 is the word-portion of that 17-bit byte-address-space. But that's probably the wrong way to think about it. 64k words in the address space with each word being 16 bits is easier I think.

Instruction Register

This is a 16-bit register that holds the current instruction that you are executing. From this register you can decode all the information needed to execute the instruction. This information will consist of register addresses (both sources and destinations), function code information for ALU and shifter, mux settings for the various muxes in the control paths, and other information about which instruction is being executed for so the control state machine knows what to do. The main issue with the instruction register is where it gets its data from. See the MDR discussion for more details...

Sign Extension

Various instructions in our machine make use of sign-extended immediate values. Recall from reading the 3710-ISA handout that immediate values in arithmetic operations are sign-extended from the 8-bits that are in the instruction encoding. Logical immediate operations are zero-extended instead of sign-extended. Check the 3710-ISA handout for details. Sign extension can be easily done using Verilog concatenations as inputs to various circuit components like muxes.

Memory and Memory-Mapped I/O System

Load and Store instructions on our version of the machine also point to word locations just like the PC (load and store addresses are word-addresses). There's no way to load a single byte on our machine. Thus, the Load and Store instructions also use word addresses. In this machine I/O is memory-mapped which means that access to I/O devices is by loading and storing to special locations in the memory space. Each of those locations will be a 16-bit location because all loads and stores deal with 16-bit data in our machine.

The memory map for one example of our baseline machine is shown in the slides on the web page (last slide in the CR-16 intro slides). In that case the memory was separated into four equal sized chunks. These chunks are either 16k words if you're thinking of 16-bit words as the basic unit, or 32k bytes if you're thinking of bytes as the basic unit (I think words are easier in this case). Looking at the top two bits of the address can tell you which quadrant of memory you're in. If you're in the top quadrant of memory (word address C000 to FFFF) then you're accessing I/O space in this example instead of code/data space. I/O devices would be mapped into specific locations, or ranges of locations, in that space. Remember, the larger you can make the range of addresses that correspond to an I/O device, the fewer bits you need to check to see if you're accessing that device.

Let's be a little more specific with a couple of required I/O devices: the LEDs and the switches on the FPGA board. The LEDs could be mapped to memory as the space defined by C0xx. This means that whenever you do a STOR to an address in the range of C000 to C0FF that value (or the low half of the 16-bit value since there are only 8 LEDs) will get written to the LEDs instead of to the memory. This means a couple things from your point of view. First you need to have an LED register that is written only when addresses in that range are being written to, and the outputs of that register should be connected (through the UCF file) to the LEDs on the Spartan3e board.

Because the LEDs are a write-only I/O device, you could re-use that range of addresses for the switches. Or you could define a different memory range for reading the value from the switches. I think it probably makes more sense to use the same range. Either way the switches are also mapped into the I/O space. When you read from, say, C0xx, you should get the value on the slider switches into the register. I would probably map the sliders to the lowest nibble of the register, and the four pushbuttons around the rotary switch to the next nibble. That way you can read, under program control, the value of 8 different switches and buttons on the Spartan3e board.

Whatever you decide to do with those bits, you will need to build an address decoder for your memory system. In its simplest form it looks like lab2 and needs to make sure that memory loads and stores in the

range of 0000 to BFFF go to the memory that you're using for code/data, and loads and stores outside that range do something else. In particular, anything in the range C000 to C0FF will go to the LED register and come from the switches. Address decoding at its simplest looks at the address on the address bus and uses combinational logic to control the enable signals of memory elements based on which address is being accessed.

Block RAM interface

The most convenient memory to use for your processor is the Block RAM that lives on the Spartan3e chip. We have a Spartan3e500 FPGA which has 20 Block RAM blocks on the chip. Each block is 18k bits. The RAM is configurable in a large number of ways (see the Mini-MIPS slides on the class web site, slide 23 or so). The most convenient for our 16 bit processor is probably the 1k x 16 version. You can actually make each block 1k x 18 if you have any need for those extra two bits. See the slides for details about the read-first/write-first issues, and remember that the block RAMs are clocked on both reads and writes.

You can instantiate Block RAM in your project in (at least) two ways:

1. Inference with behavioral Verilog: This is the technique we used in `exmem.v` in lab2 with mips. The `exmem.v` code was written according to a specific Verilog template that ISE understands. Using that template will result in Block RAM being used on the chip. You can get that Verilog template from the ISE tools. Go to Edit→LanguageTemplates in the ISE WEBpack tool. Then select Verilog→Synthesis Options→Coding Examples→RAM→Block RAM→Single Port to see all the different Verilog code templates. You'll see that there are Verilog behavioral templates for all the different types of Block RAM configurations (single/dual port, no-change/read-first/write-first, examples with enables). You can initialize the behavioral Block RAM using the `$readmemh` or `$readmemb` command as documented in the Verilog code template.
2. Direct instantiation of a Block RAM: In this technique you directly instantiate the Block RAM module in your Verilog code by name. To see examples of these start with the Edit→LanguageTemplates in the ISE WEBpack tool. Then open Verilog→Device Primitive Instantiation→FPGA→RAM / ROM→Block Ram→VirtexII/II-PRO Spartan-3/3E. From there you can see different examples for dual port and single port varieties. In the single-port category, for example, you can see an example of instantiating a `1kx16+2 RAMB16_S18` component in Verilog. You'll see that the mechanism for initializing the directly instantiated Block RAM is different than for the behavioral Block RAM. For direct instantiation you use the `.INIT` statements to specify the initial contents of the RAM.

As a side note, there are also nice Verilog templates for making smaller RAMs with "Distributed RAM" which can make nice register files. Distributed RAM is made using the Look Up Tables (LUTs) inside the logic cells on the Xilinx part.

Back to the main point, you need to pay attention to the total amount of Block RAM available on the Spartan part. Our part has 20 Block RAMs, each one is 18k bits. So, if you configure each of them as 1k x 16, then you can get a max of 20 of those 1k x 16 blocks. That doesn't fill up the available memory space of your processor. That's 20k addresses and your 16 bit address space allows 64k addresses. You have two alternatives: live within 20k words of code/data space, or use other, external, types of memory. There are good and bad points for each of these options.

Living in 20k words with VGA: This is probably fine as long as you don't plan on large, colorful VGA frame buffers. It will be hard to write applications that actually consume more than 20k of code/data space. But, if you use a VGA frame buffer mapped into your processor's RAM space, that can chew up space in a hurry. Consider 640x480 (standard VGA, low res by today's standards) is 307,200 pixels. If you use the VGA connector on the Spartan3e board, you get only one bit each for R G and