

Bluespec for a Pipelined SMIPSV2 Processor

6.375 Laboratory 3
March 12, 2006

The third laboratory assignment is to implement a pipelined SMIPSV2 in Bluespec SystemVerilog. As with Lab One, your deliverables for this lab are (a) your working Bluespec code checked into CVS and (b) written answers to the critical thinking questions. The lab is due at the start of class on Monday, March 20. Again you may submit your question answers on paper in class, or electronically in CVS (plain text or PDF only). As always, you may discuss your design with others in the class, but you must turn in your own work.

The goal of this lab is to familiarize yourself with the Bluespec language and compiler, with particular emphasis on scheduling Guarded Atomic Actions. For this lab we have decided to impose the following restriction: do not use Wires, RWires, PulseWires, or any kind of primitive Wire to complete these assignments. We feel that it is important for you to gain insight into what the Bluespec scheduler is doing and why. These combinational primitives can muddy the picture and make it difficult to reason about your design. Instead use the Bypass FIFO and Ephemeral History Register modules we provide to abstract combinational behavior. After completing this lab you will have a much better understanding of when Wires are and are not appropriate, and will be better equipped to use these constructs in your final project.

The SMIPSV2 processor you design in this lab will implement the same ISA as the previous lab, but will be a different microarchitecture. The major difference between this design and the previous lab is the presence of Instruction and Data caches. Even on a cache hit, data will not be returned until the following cycle.

This means that the processor must be split into a 4-stage pipeline, as shown in Figure 1. This has several ramifications, including: (a) you must correctly detect dependencies and stall the pipeline when they occur; (b) when a branch is mis-predicted you should not only kill the instructions which were fetched, but also correctly ignore all responses from memory requests these instructions have made; (c) FIFOs must be sized appropriately and have good scheduling properties for your design to achieve high throughput.

Before proceeding you should do *Tutorial 8: GAA-to-RTL Synthesis using the Bluespec Compiler*. This will familiarize you with the Bluespec compiler and simulating a simple SMIPSV2 microarchitecture generated from BSV source code. Additionally you should understand all the concepts presented in lecture *Bluespec-6, Modelling Processors*. The processor in that example, although much simpler than the one presented here, will serve as a good starting point for your design. As with Lab One you should start with the harness at `/mit/6.375/lab-harnesses/lab3-harness.gz`.

1 Processor Overview

The design presented in Figure 1 is mostly straightforward: the pcGen stage generates the next instruction address using the PC register, and sends it to the `instReqQ` FIFO. It also writes the current PC + 4 to the `pcQ` for use in branch instructions. The response from the ICACHE is placed

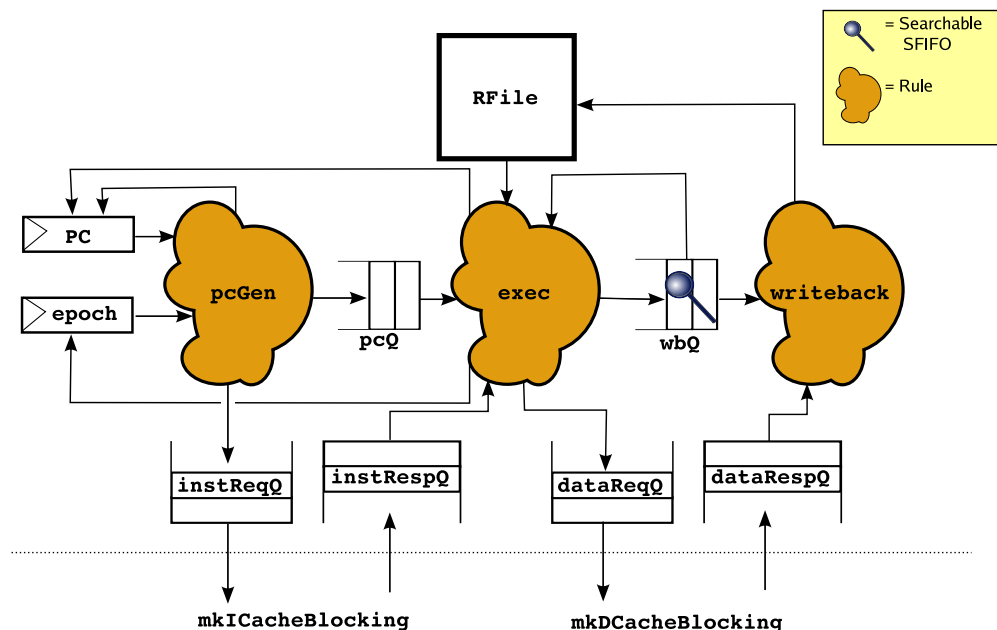


Figure 1: Pipeline Rules for SMIPSV2 Processor

into the `instRespQ` FIFO. Then the Execute stage takes the response out of the FIFO, reads the register file, and performs the operation. The result of the execution is placed into the writeback queue `wbQ` (the details of which we will explain in Section 2). If the instruction was a load or a store the memory request is placed into the `dataReqQ` queue. The Writeback stage is responsible for taking the final result from the `wbQ` and the `dataRespQ` and actually updating the register file.


For your design you will probably want one rule for each stage, so a `pcGen`, `execute`, and `writeback` rule. (Plus a rule for dealing with the epoch as we will explain in Section 3.) The interaction of these rules will be key, so carefully craft their predicates and actions. When is it safe to execute the next instruction? What actions should the writeback rule perform? Careful thought about the *intent* of each rule will aid you during implementation.

2 Stalling the Pipeline

One significant part of this assignment is generating the stall signal for the execute stage. When is it safe to read the register file and execute the next instruction? Essentially, this means detecting Read-After-Write (RAW) hazards. So if our design is executing the following program:

```
A) addiu r6, r10, 1
B) xor   r4, r5, r6
```

The processor will eventually reach the following state:



Instruction B cannot be executed because it reads `r6`, which is being written by instruction A. This means we need to stall until the writeback of A is complete. We will implement this using `SFIFO`, a searchable FIFO. `SFIFO` is just like a normal Bluespec FIFO except it has the following interface:

```
interface SFIFO#(type any_T, search_T);
  //Standard FIFO methods
  method Action enq(any_T data);
  method Action deq();
  method any_T first();
  method Action clear();
  //New SFIFO methods
  method Bool find(search_T searchVal);
  method Bool find2(search_T searchVal);
endinterface
```

Note that `SFIFO` has two extra methods, `find()` and `find2()`. These methods take a searchtype parameter, and return a boolean. Specifically, they return `True` if the given parameter is present in the FIFO, and `False` otherwise. `find()` and `find2()` have no implicit condition — they always ready — they will simply return `False` if the FIFO is empty.

Why does `SFIFO` include two methods `find()` and `find2()`? This so you can search it twice for instructions that have two operands. For instance, in the above example the execute rule can check if `r5` is in the `writebackQ` using `find()`, and `r6` using `find2()`. However if the instruction had been `xori` instead of `xor`, it just would have used `find()` because `xori` just has one register argument.

What type should you store in the writeback queue and how should you search it? First let us consider what the result of the execute stage should be. Well, if the instruction was an ALU op, then the result should be the destination register and the data to put into it. If it was a Load, then we need the destination register to put the response from memory into. If it was a store then we need to record this fact so we can receive the acknowledgement from memory. Finally let's treat the `cp0_fromhost` register specially.

```

typedef union tagged
{
    struct {Data data, Rindx dest} WB_ALU;
    Data WB_Host;
    Rindx WB_Load;
    void WB_Store;
}
WBResult deriving (Eq, Bits);

```

Now you must define what it means to search the fifo by writing a function. This function should take an `Rindx` (the thing we're looking for) and a `WBResult` (the thing we're searching). The function should return true if the search value is "found" in the FIFO.

```

function Bool findf(Rindx searchval, WBResult val);
    //You write this
endfunction

```

When you instantiate the SFIFO, you should pass in the appropriate types and find function. What does it mean to pass in a function to a hardware module in Bluespec? Essentially it means that when the compiler instantiates the module it will do so with the combinational logic you provide. Think of the SFIFO as a black box — a black box with a hole in it. The function you provide fills that hole.

Thus the types of the writeback queue is as follows:

```

//Searchable for stall signal
SFIFO#(WBResult, Rindx) wbQ <- mkSFIFO(findf);

```

All in all, the best way to encapsulate the stall signal is probably by writing a function called `stallfunc()`. `stallfunc()` takes an instruction and an SFIFO and returns False if can be executed, and True if it must stall.

So your design will probably look something like this:

```

function Bool stallfunc(Instr inst, SFIFO#(WBResult, Rindx) f);
... //You write this ...
endfunction
module mkProc (IProc);
... //State elements ...
rule execute (instRespQ.first() matches tagged LoadResp .ld
              &&& unpack(ld.data) matches .inst
              &&& !stallfunc(inst, wbQ));
...
case (inst) matches //Execute the instruction
tagged LW .it:
...
endrule

```

NOTE: This rule predicate is not quite complete, as we will learn in the next section.

3 Dealing with Branches

You may have noticed that branches are resolved in the Execute stage. Why is this a problem? Because if the branch has been taken (or, with a branch predictor, if the branch has been mispredicted) then the `pcGen` stage has made a memory request for an instruction which we must ignore. (In a design with a non-blocking cache it may even have made more than one.) There are many ways to handle this, but the simplest way is to use an *epoch*.

An epoch is a conceptual grouping of all instructions in-between branch mispredictions. We can track this by having the *pcGen* rule send the epoch as the tag for all load requests:

```

rule pcGen ...
...
instReqQ.enq( Load{ addr:pc, tag:epoch } );
...
endrule

```

Note that this is okay because our memory system is in-order, so the tag is essentially unused. If the memory system was allowed to respond out-of-order then we would have to actually create an appropriate tag to differentiate responses. In this case we could devote some bits of the tag field to the epoch, and some to the differentiation — but we do not need to worry about this for this lab.

When a mispredict occurs we clear all queues which are holding instructions issued after the branch instruction, and increment the epoch. Then we have a separate rule that discards responses from the wrong epoch.

```

rule discard ( instRespQ.first() matches tagged LoadResp .ld
               &&& ld.tag != epoch);
  traceTiny("stage","D");
  instRespQ.deq();
endrule

```

Now when we execute we must check that we can execute the instruction, and that it is from the correct epoch:

```

rule execute (instRespQ.first() matches tagged LoadResp .ld
              &&& ld.tag == epoch
              &&& unpack(ld.data) matches .inst
              &&& !stallfunc(inst, wbQ));
  ...
  Bool branchTaken = False;
  Addr newPC;
  ...
  if (branchTaken)
  begin
    //Clear appropriate FIFOs here
    epoch <= epoch + 1;
    pc <= newPC;
    ...
  end

```

4 Working with BSVCLib

In Figure 1 we represent all FIFOs in the design with the same picture, as if they were the same hardware structure. In reality, in order to achieve good throughput you will need to (a) appropriately size all FIFOs and (b) ensure that they have the correct scheduling properties to ensure maximum concurrency between rules.

To this end we are providing you with a library of predefined FIFOs with various properties. In BSVCLib, located at `/mit/6.375/install/bsvclib`, we've provided a library of useful Bluespec modules, including SFIFO and a Bypass FIFO that you should use in your design. To use SFIFO import the SFIFO package. To use the Bypass FIFO, import BFIFO.

The properties of various FIFOs are given in Figure 2. When choosing a FIFO remember to consider both its size, and its schedule. What case do you expect to be the most common? How does the memory latency affect things? Is extra area and an extra cycle of latency worth an improvement in throughput. Be sure to run the benchmark suite and examine which rules fire to check how your change impacts throughput.

Sometimes there may be places where you wish no FIFO existed at all, IE a wire. The problem with such a combinational structure is that you must be able to guarantee that your rules will always fire when values are on the wire — no communication should be dropped under any circumstance. Rather than making you reason in such a way for this lab, we are providing you with a safer abstraction: `mkBFIFO1()`, a Bypass FIFO. This FIFO behaves like a wire (`enq()` before `deq()`) as

FIFO Variant	Package	Size	Sched	Comment
<code>mkFIFO()</code>	FIFO	2	<code>deq < enq</code>	Default. Use this to get your design working. <code>deq</code> and <code>enq</code> may happen simultaneously when contains 1 element
<code>mkFIFO1()</code>	FIFO	1	<code>deq ME enq</code>	<code>deq</code> if full, <code>enq</code> if empty. Mutually exclusive.
<code>mkSizedFIFO(<i>n</i>)</code>	FIFO	<i>n</i>	<code>deq < enq</code>	<code>deq</code> and <code>enq</code> may happen simultaneously when neither full nor empty
<code>mkLFIFO()</code>	FIFO	1	<code>deq < enq</code>	<code>deq</code> and <code>enq</code> may happen simultaneously when full
<code>mkBFIFO()</code>	BFIFO	1	<code>enq < deq</code>	If <code>enq</code> and <code>deq</code> happen when empty, value is bypassed
<code>mkSizedBFIFO(<i>n</i>)</code>	BFIFO	<i>n</i>	<code>enq < deq</code>	A larger buffer for when no <code>deq</code> happens
<code>mkSFIFO()</code>	SFIFO	1	<code>deq < find < enq</code>	Uses SFIFO interface. Properties are like <code>mkFIFO</code>
<code>mkSizedSFIFO(<i>n</i>)</code>	SFIFO	<i>n</i>	<code>deq < find < enq</code>	Uses SFIFO interface. Properties are like <code>mkSizedFIFO</code>

Figure 2: Properties of various FIFO modules. For all FIFOs: `first < (enq, deq) < clear`.

long as both occur. Otherwise the value is buffered in the FIFO, so the `deq()` can occur later. You should ensure that your design is functionally correct with normal FIFOs before you attempt to introduce Bypass FIFOs.

5 Lab Hints and Tips

This section contains several hints and tips which should help you in completing the lab assignment.

Tip 1: Correctness First, Performance Second

Design for correctness first, then worry about performance. Initially you should implement the design using the standard `mkFIFO` and `mkSFIFO` queues. You can implement naive stalling logic by having your `findf` function simply return `True`. This makes the execute stage stall any time the writeback queue contains any value. Pay special attention to the writeback queue `wbqQ` and the `dataReqQ`. When should the execute rule write into which queue? What assumptions will the writeback rule make about their contents?

After all the queues are in place implement the epoch, making the fetch send the appropriate tag to the memory system. Then change the execute rule to use the writeback queues instead of updating the register file itself. Implement the writeback rule, taking into consideration the different situations: Load Responses, Store Responses, etc.

At this point your design should be able to pass `make run-asm-tests`. Even such a naive design should show an improvement in IPC on the benchmarks over `smips-4mcycle`.

Then, once your system is complete, begin to improve performance iteratively. First implement the real stall signal using `SFIFO`. Then examine the schedule and see which rules are blocking each other. What can you do about that? Is this a place where a different FIFO could affect things? After making a change (particularly a change which involves adding `mkBFIFO1` make sure to rerun `make run-asm-tests` to ensure correctness.

Tip 2: Check the Scheduler Output

As you begin to refine your design, pay close attention to the output of the Bluespec compiler. Sometimes rule conflicts can point to a bug in the design. For example, early in my implementation the scheduler deduced that both the execute and writeback rules were using the Register File write method, as it turned out, this was because I had a bug where I forgot to change one instruction in the execute case statement, so it was still updating the register file directly. After viewing the schedule I realized that something must be wrong.

This step will become even more important as you begin to change your FIFOs to improve throughput. Long FIFOs will tend to “decouple” rules so that they become more independent. Shorter FIFOs will do the opposite, as the rules will interact through the state elements directly. Make sure you understand all of this before you begin using the Bypass FIFO!

Tip 3: Observe rule CAN_FIRE and WILL_FIRE signals

As you refine, a good strategy is to run a single benchmark and observe the waveform. For each rule the compiler outputs a CAN_FIRE signal and a WILL_FIRE signal. While the scheduler output can give you a sense of static priorities, use these signals to observe the dynamic interactions of these rules.

A related tip is to observe the signals for FIFO methods. Each FIFO will have wires corresponding to `enq`, `deq`, and `first`, as well as `RDY_enq`, `EN_deq`, and so on. Observing when enqueues and dequeues are occurring can be extremely helpful, particularly if your design is deadlocking for some reason.

Tip 4: Think very carefully about sizing FIFOs

Although the scheduling properties of the FIFOs are important, so too is their length. For instance, consider the `pcQ` FIFO. At the beginning of time it starts out empty, then the `pcGen` rule enqueues into it, and makes a memory request. Even if the memory request comes back the following cycle because of a cache hit, the response still has to go into the `instRespQ`. Therefore it seems that making both the `instRespQ` and the `pcQ` Bypass FIFOs could result in the execute rule reading the PC for an incorrect instruction.

Although it is possible to find a good configuration using an experimental approach, reasoning about the system using high-level knowledge can point you towards the optimum configuration. Always check the effect of your changes on the generated schedule.

6 Critical Thinking Questions

As with Lab One, the primary deliverable for this lab assignment is your working Bluespec source code for the pipelined SMIPSV2 processor. In addition, you should prepare written answers to the following questions and turn them in at the beginning of class on the due date. Question 3 involves adding a branch predictor, as you did in Lab One. Again, this may require significant design work. Again, do not worry if you are unable to finish the branch predictor question. As long as you make a reasonable effort you will not be penalized.

Question 1: Design Choices and Synthesis

Discuss and motivate any design choices you made. Is there any way in which your implementation differs from the diagram in Figure 1? What FIFOs did you end up using, and why is this a good configuration? What is the relationship between your Execute and Writeback rules? Are they Conflict-Free, Sequentially Composable, or Conflicting? Why has the scheduler deduced this?

Push your design through the synthesis flow from Lab Two. For now do not try to optimize these numbers, just report them. Report on area, clock period, and Instructions per Cycle (IPC) for the various benchmarks. Also calculate Instructions per Second (IPS, IPC / Clock Period). Are there any surprises in the results?

Question 2: Refining using Ephemeral History Registers

In the BSVCLib (at `mit/6.375/install/bsvclib/`) there's an package called `EHR2`. In this package you'll find `mkEHR2_Reg` an implementation of an Ephemeral History Register. The interface is similar to a normal register except with "subscripts" added:

```
interface EHR2_Reg#(type valType);
  method valType read_0();
  method valType read_1();
  method Action write_0(valType val);
  method Action write_1(valType val);
endinterface
```

Recall that the schedule for the EHR is as follows:

$$\text{read}_0 < \text{write}_0 < \text{read}_1 < \text{write}_1$$

This package also contains `mkEHR2_RFile`, a register file which has the Ephemeral History structure. It has the following interface:

```
interface EHR2_RFile#(type addrType, type valType);
  method Action wr_0(addrType rindx, valType data);
  method Action wr_1(addrType rindx, valType data);
  method valType rd1_0(addrType rindx);
  method valType rd1_1(addrType rindx);
  method valType rd2_0(addrType rindx);
  method valType rd2_1(addrType rindx);
endinterface
```

The schedule for this register file is as follows;

$$\text{rd1}_0, \text{rd2}_0 < \text{wr}_0 < \text{rd1}_1, \text{rd2}_1 < \text{wr}_1$$

Finally, the package provides `mkEHR2_FIFO` and `mkEHR2_SFIFO`. These have the following interface:

```
interface EHR2_FIFO#(type valType);
  method Action enq_0(valType d);
  method Action enq_1(valType d);
  method Action deq_0();
  method Action deq_1();
  method valType first_0();
  method valType first_1();
  method Action clear_0();
  method Action clear_1();
endinterface
```

These have the following schedule:

```
(first_0, find_0, find2_0) < deq_0 < enq_0 < clear_0 <
(first_1, find_1, find2_1) < deq_1 < enq_1 < clear_1
```

Note that there is no bypassing version of these FIFOs. With EHRs this same bypassing effect can be achieved by using `enq_0` with `first_1` and `deq_1`.

Using these constructs, create a version of your design which has the following scheduling property:

```
writeback < execute < pcGen
```

To do this you will need to mechanically transform your design as described in class. First change all the state elements to the appropriate EHR type. Then label all method calls in `writeback` with subscript 0, and proceed from there.

What throughput does this version of your design achieve? Synthesize this version of the design as you did for Question 1. How have the results changed and what was the cause of the biggest change? In your opinion is the difference worth it? How does this method compare to experimentally sizing your FIFOs as in Question 1?

Question 3: Adding a Branch Predictor

As with lab one, create a version of your design which adds a branch predictor and a branch history table. Similarly with Lab One, you should attempt this question after all other questions and tasks are finished. The purpose of this task is to think about designing modules in Bluespec from scratch, and then using them in a design.

This will encompass:

- Creating a new interface for the branch predictor.
- Creating a module which provides that interface.

- Changing your design to use this new module properly.

What was the effect does this new module have on your design's schedule? On the overall throughput of your processor? Synthesize this new design as you did in Question 1 and report on the results. In your opinion, is the increase in throughput worth the penalty to area?

Question 4: Synthesizing Bluespec by Hand

Ben Bitdiddle is writing a 32-bit barrel shifter in Bluespec. The module can take any 32-bit number and shift it right or left by any amount. To minimize area Ben decides to implement the design using a circular shifter and a counter.

Unfortunately, Ben accidentally forgets to write predicates for his rules:

```
typedef enum {Left, Right} Direction;
module mkBarrelShifter (Shifter);
  Reg#(Bit#(32)) r <- mkReg(0);
  Reg#(Bit#(5)) cnt <- mkReg(0);
  Reg#(Direction) dir <- mkRegU();
  rule shiftLeft (True);
    r <= r << 1;
    cnt <= cnt - 1;
  endrule
  rule shiftRight (True);
    r <= r >> 1;
    cnt <= cnt - 1;
  endrule
  method Action shift(Direction d, Bit#(32) data, Bit#(5) amt) if (cnt == 0);
    dir <= d;
    r <= data;
    cnt <= amt;
  endmethod
  method Bit#(32) result() if (cnt == 0);
    return r;
  endmethod
endmodule
```

What schedule will the compiler deduce for Ben's design? (If you must make a choice at some point make it arbitrarily.) What hardware will the compiler generate? Diagram the resulting datapath, clearly labelling which part corresponds to the scheduling logic.

What predicates should Ben have provided? How do they change the schedule? Redraw the correct datapath which will result, again highlighting the scheduling logic.

Note that we understand that you could do this problem by typing the code into the compiler and seeing what happens. You could even base your diagram on the resulting Verilog! Unfortunately, the compiler will not be available to help you during the quiz... hint hint....