C152 Laboratory Exercise 3

Instructor: John Lazzaro
TA: Eric Love
Author: Christopher Celio
Department of Electrical Engineering & Computer Science
University of California, Berkeley

Apr 9, 2014

1   Introduction and goals

The goal of this laboratory assignment is to allow you to conduct a variety of experiments in the `Chisel` simulation environment.

You will be provided a complete implementation of a speculative out–of–order processor. Students will run experiments on it, analyze the design, and make recommendations for future development. You can also choose to improve the design as part of the open-ended portion.

The lab has two sections, a directed portion and an open–ended portion. Everyone will do the directed portion the same way, and grades will be assigned based on correctness. The open–ended portion will allow you to pursue more creative investigations, and your grade will be based on the effort made to complete the task or the arguments you provide in support of your ideas.

Students are encouraged to discuss solutions to the lab assignments with other students, but must run through the directed portion of the lab by themselves and turn in their own lab report. For the open-ended portion of each lab, students can work individually or in groups of two. Any open-ended lab assignment completed as a group should be written up and handed in separately. Students are free to take part in different groups for different lab assignments.

You are only required to do one of the open–ended assignments. These assignments are in general starting points or suggestions. Alternatively, you can propose and complete your own open–ended project as long as it is sufficiently rigorous. If you feel uncertain about the rigor of a proposal, feel free to consult the TA or the professor.

1.1   `Chisel` & The Berkeley Out–of–Order Machine

The `Chisel` infrastructure is very similar to Lab 1, with the addition of a new processor, the RISC-V Berkeley Out–of–Order Machine, or "BOOM". BOOM is heavily inspired by the MIPS R10k and the Alpha 21264 out–of–order processors[1, 3]. Like the R10k and the 21264, BOOM is a unified physical register file design (also known as "explicit register renaming"). BOOM is a superscalar processor; the width of its various stages is parameterizable.

The Chip

For this lab, you will be given an entire functioning processor, implemented in `Chisel`. The `Chisel` source code describes an entire "chip" with an interface to the outside world via a DRAM memory link. On-chip is an out–of–order core, which is where the focus of this lab will be. The core, in this case the BOOM processor, is directly connected to an instruction cache (16kB) and a non-blocking data cache (32kB). Any miss in either cache will require a trip to DRAM[2] (located "off-chip").

The BOOM Pipeline

NPC = next pc calculation
I\$ = instruction cache access
BTB = branch target buffer
BHT = branch history table
BP = branch predict
Dc/Rn = Decode / Rename
IS/RF = Issue / RF Read
EX = Execute
D\$ = data cache access
Fmt = formatting LD data
WB = RF writeback
Com = ROB Commit

Figure 1: The Berkeley Out of Order Machine Processor. Shown is a single-issue version.

Conceptually, BOOM is broken up into 10 stages: *Fetch, Decode, Register Rename, Dispatch, Issue, Register Read, Execute, Memory, Writeback,* and *Commit.* However, many of those stages are combined in the current implementation, yielding *six* stages: *Fetch, Decode/Rename/Dispatch, Issue/RegisterRead, Execute, Memory,* and *Writeback* (*Commit* occurs asynchronously).

Fetch Instructions are *fetched* from the Instruction Memory and placed into a four-entry deep FIFO, known as the *fetch buffer.*[1]

Decode *Decode* pulls instructions out of the *fetch buffer* and generates the appropriate "micro-op" to place into the pipeline.[2]

---

[1]While the fetch buffer is four-entries deep, it can instantly read out the first instruction on the front of the FIFO. Put another way, instructions don't need to spend four cycles moving their way through the *fetch buffer* if there are no instructions in front of them. In fact, BOOM's fetch buffer is a "flow-through queue" meaning if the queue is empty, the Decode stage can immediately operate on the instructions returning from the instruction cache without placing them in the queue first.

[2]Because RISC-V is a RISC ISA, nearly all instructions generate only a single micro-op, with the exception of store instructions, which generate a "store address generation" micro-op and a "store data generation" micro-op.

Rename The ISA, or "logical", register specifiers are then *renamed* into "physical" register specifiers.

Dispatch The instruction is then *dispatched*, or written, into the *Issue Window*.

Issue Instructions sitting in the *Issue Window* wait until all of their operands are ready, and are then *issued*. This is the beginning of the out–of–order piece of the pipeline.

RF Read Issued instructions first *read* their operands from the unified physical register file...

Execute and then enter the *Execute* stage where the integer ALU resides. Issued memory operations perform their address calculations in the *Execute* stage, and then the address is sent to the data cache (if it is a load) in which the data is accessed during the *Memory* stage. The calculated addresses are also written into the Load/Store Unit at the end of the *Execute* stage.

Memory The Load/Store Unit consists of three queues: a Load Address Queue (LAQ), a Store Address Queue (SAQ), and a Store Data Queue (SDQ) (see Figure 5). Loads are optimistically fired to memory when their address is added to the LAQ during the *Execute* stage. In parallel, the incoming load compares its address against the SAQ to find if there are any store addresses that the load depends on. If the store data is present, the load receives the data from the store (*store data forwarding*) and the memory request is killed. If the data is not present, the load is put to sleep. Loads that are put to sleep are reissued to memory at a later time. The current policy in BOOM is to retry the load at the head of the LAQ .[3]

Stores are fired to memory at *commit*, when both its address and its data are present.

Writeback ALU operations and load operations are *written* back to the physical register file.[4]

Commit The Reorder Buffer, or ROB, tracks the status of each instruction in the pipeline. When the head of the ROB is not-busy, it *commits* the instruction. For stores, the ROB signals to the store at the head of the Store Queue that it can now write its data to memory. For loads, the ROB signals the Load/Store Unit to verify that the load did not fail a memory ordering dependence (i.e., a load issued before a store it depended on committed). If the load did fail, the entire pipeline must be killed and restarted. Exceptions are also taken at this point, which requires slowly unwinding the ROB to return the rename map tables to their proper state.

BOOM supports full branch speculation and branch prediction. Each instruction, no matter where it is in the pipeline, is accompanied by a branch mask that marks which branches the instruction is "speculated under". A mispredicted branch requires killing all instructions that depended on that branch. When a branch instructions passes through *Rename*, copies of the

---

[3]A higher performance processor would allow loads to track *why* they were put to sleep, and to wake and retry loads once the issue has been resolved.

[4]Even the single-issue version of BOOM provides ALU operations and memory operations each their own write port, meaning the register file is a two-read, two-write register file (two different destinations can be written simultaneously). Also notice that in a *unified physical register file design*, speculative instructions are being written back.

Register Rename Table and the *Free List* are made. On a mispredict, the saved processor state is restored.

The "front-end" contains a Branch History Table, composed of simple n-bit history counters indexed by the PC. The BHT is read in parallel with instruction cache access. As an instruction is returned from the cache and inserted into the *fetch buffer*, the instruction is quickly checked to see if it is a branch. If the instruction is a branch, the prediction is used to redirect the *front-end* on a *TAKE BRANCH* prediction. BOOM also has a Branch Target Buffer (BTB) that directs the *next pc*.

BOOM implements the RISC-V variant RV64IMS[5]. RV64IMS is the 64-bit variant which provides the basic integer instructions, multiply/divide instructions, and the supervisor-level ISA. BOOM does *not* provide hardware support for floating point.

See Figure 5 for a more detailed diagram of the pipeline. Additional information on BOOM can be found in the appendices and the CS152 BOOM Section notes; in particular, the issue window, the load/store unit, and the execution pipeline are covered in greater detail.

---

[5]Atomic memory operations and floating point instructions are the main categories of instructions that BOOM does not implement.

1.2   Graded Items

Submit your lab report as a file `report.pdf` in the top level of the `lab3` directory in your private repo on github. Some of the open-ended questions also request source code - please also submit this to your git repo. Please label each section of the results clearly. The following items need to be turned in for evaluation:

1. Problem 2.2: CPI, branch predictor statistics, and answers

2. Problem 2.3: CPI statistics and answers

3. Problem 2.4: Issue Window statistics and answers

4. Problem 2.5: Issue Window statistics, instrumentation code, and answers

5. Problem 3.1, /3.2 modifications and evaluations (open-ended portion)

2   Directed Portion

The questions in the directed portion of the lab use `Chisel`. A tutorial (and other documentation) on the `Chisel` language can be found at (`http://chisel.eecs.berkeley.edu`).[6] Although students will not be required to write `Chisel` code as part of this lab, students will need to write instrumentation code in C++ code which probes the state of a `Chisel` processor.

WARNING: `Chisel` is an ongoing project at Berkeley and continues to undergo rapid development. Any documentation on `Chisel` may be out of date, especially regarding syntax. Feel free to consult with your TA with any questions you may have, and report any bugs you encounter. Likewise, BOOM will pass all tests and benchmarks for the default parameters, however, changing parameters or adding new branch predictors will create new instruction interleavings which may expose bugs in the processor itself.

2.1   Setting Up Your `Chisel` Workspace

To complete this lab you will log in to an instructional server, which is where you will use `Chisel` and the RISC-V tool-chain.

The tools for this lab were set up to run on any of the twelve instructional Linux servers `t7400-1.eecs`, `t7400-2.eecs`, ..., `t7400-12.eecs`.

---

[6]Chisel documentation can also be found within the lab itself. Look under `$Lab3Root/chisel/doc/` for more information.

First, download the lab materials into your private repo directory:[7,8]

```
inst$ cp -R ~cs152/sp14/lab3 ./lab3

inst$ cd ./lab3
inst$ export LAB3ROOT=$PWD
```

You should make sure to git add this directory structure and commit it to your repo now, before you generate lots of extra stuff while compiling the emulator

We will refer to `./lab3` as `${LAB3ROOT}` in the rest of the handout to denote the location of the Lab 3 directory.

The directory structure is shown below:

- `${LAB3ROOT}/`
    - doc/ Useful documentation and related materials.
    - csrc/ C++ simulation source code. The out-of-order tracer lives here, as does emulator.cc
    - bsrc/ Chisel source code for BOOM.
        * **consts.scala** BOOM configuration file. Parameters such as machine width, ROB size, bypassing-enable, etc.
    - emulator/Compiled emulator binaries and object files.
        * **Makefile** The Makefile you'll need to run experiments using the benchmarks
        * generated-src/ C++ code generated by Chisel for the emulator
            · **Top.h** Declarations of all Chisel signals in C++
        * **output/** results of make run and make run-bmarks-test (CPI etc.)
    - riscv-test/ Source code for benchmarks and tests.
        * benchmarks/ Benchmarks written in C.
            · lsu_forwarding/ A benchmark you can write for open-ended problem 3.2.
            · lsu_failures/ A benchmark you can write for open-ended problem 3.2.
        * riscv-tests/ Tests written in assembly.
    - install/ Install directory for tests and benchmarks that are visible to BOOM.
    - chisel/ The Chisel source code.
    - src/ Chisel code for interfacing BOOM with the surrounding simulator infrastructure
    - uncore/ Chisel source code to the uncore (i.e., outside of the "tile").
    - riscv-pk/ Repository of the RISC-V proxy kernel. Serves as the OS for BOOM.
    - dramsim2/ A DRAM simulator that the Chisel emulator hooks into.
    - sbt/ Chisel/Scala voodoo.
    - Makefile The high-level Makefile.

---

[7]The capital "R" in "cp -R" is critical, as the -R option maintains the symbolic links used.

[8]The actual name of the Lab3 directory might have letters appended to it to denote different *versions*. Newer versions will be necessary as bugs are ironed out.

The most interesting items have been bolded:**emulator/Makefile** to build and test the processor, the `Chisel` source code found in `bsrc`, and the output files found in `emulator/output/`.

The following command will set up your bash environment, giving you access to the entire CS152 lab tool-chain. Run it before each session:[9] [10]

```
inst$ source ~cs152/sp14/.bashrc
```

To compile the `Chisel` source code for BOOM, compile the resulting C++ simulator, and run all tests and benchmarks, run the following Bash script:

```
inst$ cd ${LAB3ROOT}/emulator
inst$ make run
```

To "clean" everything, simply run the same script with an additional parameter:

```
inst$ make clean
```

The entire build and test process should take around ten to fifteen minutes on the t7400 machines.[11] The first time you `make run`, you may have to wait a little longer while the Scala Build Tool downloads the `Chisel` compiler.

2.2   Gathering the CPI and Branch Prediction Accuracy of BOOM

For this problem, collect and report the **CPI** and **branch predictor accuracy** for the benchmarks *bubble, dhrystone, median, multiply, qsort, towers,* and *vvadd.* You will do this twice for BOOM: with and without branch prediction turned on. First, turn off branch prediction as follows:

```
inst$ vim ${LAB3ROOT}/bsrc/consts.scala
```

Change the lines `USE_BRANCH_PREDICTOR` and `ENABLE_BTB` to be set to "false". Then compile the resulting simulator and run it through the benchmarks as follows:

```
inst$ cd ${LAB3ROOT}/emulator
inst$ make run-bmarks-test
inst$ make stats
```

The `make run` and `make run-bmarks-test` commands compile `Chisel` code into C++ code, and then compile theat C++ code into a cycle-accurate simulator. Finally, they call the RISC-V front-end server to start the simulator and run a suite of tests and/or benchmarks, respectively, on the target processor. The `make stats` command greps for * in the generated *.riscv.out files in `output/` to display the tracer's statistics. Try running `make run-bmarks-test` now.

---

[9]Or better yet, add this command to your bash profile.

[10]If you see errors about "htif_pthread.h", then you probably have an improperly set environment.

[11]The generated C++ source code is ∼5MB in size, so some patience is required while it compiles.

Then do it again, but with branch prediction turned on.[12]

The default parameters for BOOM are summarized in Table 1. While some of these parameters (instruction window, ROB, LD/ST unit) are on the small side, the machine is generally well fed because it only fetches and dispatches one instruction at a time, and the pipeline is not very long.[13]

Table 1: The BOOM Parameters for Problem 2.2.

	Default
Register File	64 physical registers
ROB	32 entries
Inst Window	12 entries
LD Queue	8 entries
ST Queue	8 entries
Max Branches	8 branches
Branch Prediction	128 two-bit counters
Issue loads ASAP	on
ALU Bypassing	**off**
BTB	**off**

Table 2: CPI for the in-order 5-stage pipeline and the out-of-order "6-stage" pipeline. Fill in the rest of the table.

	bubble	dhry	median	multiply	qsort	towers	vvadd
5-stage (interlocking)	1.90	n/a	2.42	2.00	2.24	1.56	2.75
5-stage (bypassing)	1.37	n/a	2.19	1.68	1.47	1.37	2.31
BOOM (PC+4)							
BOOM (BHT)							

Table 3: Branch prediction accuracy for *predict PC+4* and a simple 2-bit BHT prediction scheme. Fill in the rest of the table.

	bubble	dhry	median	multiply	qsort	towers	vvadd
BOOM (PC+4)							
BOOM (BHT)							

Compare your collected results with the in-order, 5-stage processor. Notice that BOOM is a

---

[12]The default branch predictor provided with BOOM is a branch history table made up of 128 two-bit counters, indexed by PC.

[13]Also, by keeping many of BOOM's data structures small, it keeps compile time fast(er) and allows us to easily visualize the entire state on the machine when viewing the debug versions of the *.out files generated by simulation.

6-stage processor (with **no** bypassing enabled for this problem), so it can be most closely compared to the in-order, 5-stage with no bypassing (i.e., interlocked). Explain the results you gathered. Are they what you expected? Was out-of-order issue an improvement on the CPI for these benchmarks? Was using a BHT always a win for BOOM? Why or why not? (Don't forget to include the accuracy numbers of the branch predictor!).[14]

Additional Notes: Jump-and-Link-Register (JALR) is included in the branch accuracy statistics, while Jump-and-Link (JAL) is *not*. Jump-and-Link (JAL) is handled perfectly by the front-end[15], while JALR is always predicted as *not taken*.[16] The CPI is calculated at the *Commit* stage. Finally, the branch predictor accuracy is calculated based on the signals in the *Execute* stage, which means that the reported accuracy is also including *misspeculated* instructions.[17]

2.3 Bottlenecks to performance

Building an out–of–order processor is hard. Building an out–of–order processor that is well balanced and high performance is *really hard*. Any one piece of of the processor can bottleneck the machine and lead to poor performance.

For this problem you will set the parameters of the machine to a low-featured "worst-case" baseline (see Table 4).

Table 4: BOOM Parameters: worst-case baseline versus "default" for the rest of the lab questions.

	Worst-case	Default
Register File	33 physical registers	64 physical registers
Branch Prediction	off	128 two-bit counters
Branch Target Buffer	off	4 entries
Issue loads ASAP	off	on
ALU Bypassing	off	**on**

Begin by setting BOOM to the values in the "worst-case" column from Table 4. All of the necessary parameters can be found in `bsrc/consts.scala`.[18]

Run the benchmarks (`make run-bmarks-test; make stats`) to collect the data for the first row in Table 5. The performance should be dreadful. And when you see an error, don't panic.[19]

---

[14]Hint: when a branch is mispredicted for BOOM, what is the branch penalty?

[15]Once the JAL instruction returns from the instruction cache, the front-end can quickly decode it, calculate the target address, and redirect the PC. The JAL continues towards the backend where it is treated as a NOP, its execution already handled by the front-end!

[16]This is a sign that yours truly needs to build a *Return Address Stack* to properly predict JALR instructions.

[17]The branch predictor itself is updated in the *Commit* stage.

[18]The exact name of the variables, in order, are "PHYS_REG_COUNT", "USE_BRANCH_PREDICTOR", "EN-ABLE_BTB", "ENABLE_SPECULATE_LOADS", and "ENABLE_ALU_BYPASSING".

[19]You will probably see `Dhrystone` throw an error. Upon inspecting `emulator/output/dhrystone.riscv.out`, you should see that `Dhrystone` timed out. This is normal (it just ran that slow!). The CPI measurements are still valid and you can move on with the rest of the experiment.

Table 5: CPI for the in-order 5-stage pipeline and the out-of-order "6-stage" pipeline. Gradually turn on additional features as you move down the table. Fill in the rest of the table.

	bubble	dhry	median	multiply	qsort	towers	vvadd
5-stage (interlocking)	1.90	n/a	2.42	2.00	2.24	1.56	2.75
5-stage (bypassing)	1.37	n/a	2.19	1.68	1.47	1.37	2.31
BOOM (worst case baseline)							
BOOM (64 regs)							
BOOM (BHT & BTB)							
BOOM (fast loads)							
BOOM (bypassing)							

Now we will slowly add back the features we took away. For the 2nd row, return the physical register count to 64 registers (from 33), and rerun the benchmarks (thought problem: why is 33 registers the smallest allowed amount?)[20]

For the 3rd row, add back branch prediction (both the BHT and the BTB). Then for the 4th row add back load speculation (without load speculation, the loads will execute in order). And for the last row, enable ALU bypassing (the last row in the table should have all of the "default" values set).

Collecting this data is pretty straight-forward but admittedly time consuming (~5-10 minutes per row in the table), so *do* walk away from the computer, go outside, get coffee, or watch *Arrested Development* while your computer hums away. The idea here is to get a feel for the performance numbers when certain features are missing. But not to worry, the lab picks up very quickly in the next section!

2.4   Chisel: Analyzing the Issue Window, Part I

So far, you have been running BOOM in single-issue mode; all stages of the pipeline handle only a single instruction at a time. However, it is more than possible to implement an out-of-order processor that allows different stages to handle different amounts of instructions at a time (for example, committing two instructions at a time for a single-issue machine makes a considerable amount of sense. Why?[21]).

In fact, for this problem, your TA is wondering "Just how much performance is being left on the table by only allowing one instruction to be issued out of the *Issue Window* at a time?"

Your job is to quantify this, and answer your TA's question.

You will solve this question by writing C++ code in the "Out–of–Order Tracer" object that probes the state of BOOM every cycle. The OOOTracer object is found in `${LAB3ROOT}/csrc/oootracer.cpp/.h`. It is the same Tracer object you saw in Lab 1, with a few modifications. For this question, you will add any counters you need in the appropriate locations.[22] The main piece

---

[20]Answer: the ISA has 32 registers, and you need one additional register to act as a temporary once you have allocated all of the ISA registers.

[21]Answer: because waiting on hazards to resolve can back the machine up, potentially making the ROB commit the bottleneck.

[22]Grep for "Step".

of your logic will go into the `Tracer_t::monitor_issue_window()` function. Read the instructions provided in `$LAB3ROOT/csrc/oootracer.cpp` for additional information. See Appendix A for details on how the *Issue Window* works.

To answer this question, **count the number of cycles in which at least two issue slots are requesting to be issued**. Make sure you are only counting cycles in which "Tracer.paused" is not asserted. Some sample code is provided in `Tracer_t::monitor_issue_window()` to show how to detect that issue slot #0 is "valid".

Also, make sure to edit `consts.scala` and set `FETCH_WIDTH` equal to 2. The reason for doing this is that the limited fetch bandwidth of a one-wide machine will make the apparent benefit of dual-issue artificially small.[23]

For all benchmarks, report how many cycles contain two instructions requesting to be issued. Do you think it would be beneficial to issue up to two instructions every cycle out of the issue window?

2.5 Chisel: Analyzing the Issue Window, Part II

Issuing two instructions simultaneously could be very expensive: it would require adding two more read ports and a *third* write port to the register file to handle the worst case of two ALU operations being issued and writing back in the same cycle that a load from memory comes back.

Instead, your TA proposes to issue two instructions simultaneously *if and only if* one instruction is an ALU operation and the second instruction is a memory operation. This will require adding a second ALU to perform address calculations, and an additional read port to read out the base address or store data required for load and store micro-ops.

To answer this question, augment your previous C++ probing code by checking the micro-op code, or "uopc" tag, on each issue slot (See Figure 3): **count the number of cycles in which at least one ALU micro-op and one memory micro-op requests to be issued**. The values of each "uopc" can be found in `bsrc/consts.scala` (roughly lines 192-64).

Consider any non-Load and non-Store to be an ALU operation, for the purposes of this question.

Report your results for the benchmarks, and submit your code via Github. Having collected data for Sections 2.4 and 2.5, what is your final recommendation on supporting multiple issue in BOOM? Is single-issue out of the *Issue Window* good enough, or would ALU/Mem dual-issue or even full dual-issue be worth the added costs?

2.6 Chisel: Analyzing the Issue Window, Part III

Now that you've analyzed the theoretical benefit of wider issue, you can see how much actual performance the current BOOM implementation manages to gain from going two-wide. Open up `consts.scala` once more and set `ISSUE_WIDTH` equal to two. Now run `make run` and look at the new CPIs. How much did performance actually improve?

---

[23]Imagine, a program containing only `addi` instructions with no dependencies. Each instruction will be removed from the issue window as soon as it is put there, and so there will never be more than one requesting instruction at any time). However, if instructions were fetched in pairs, it would become obvious that two instructions could be issued every cycle.

3 Open-ended Portion

All open-ended questions should use the following parameters, as shown in Table 6 (unless otherwise specified).

Table 6: The Default BOOM Parameters for the Open-ended Questions.

Default	
Fetch Width	1 wide
Issue Width	1 wide
Register File	64 physical registers
ROB	32 entries
Inst Window	12 entries
LD Queue	8 entries
ST Queue	8 entries
Max Branches	8 branches
Branch Prediction	128 two-bit counters
Issue loads ASAP	on
ALU Bypassing	**on**

3.1 Designing Issue-select Policies

You are a new employee at Processors-R-Us charged with analyzing and improving the performance of your company's latest processor. Under heavy pressure to make the looming tape-out deadline (contracts with your customers are pretty strict), the lead architect and your boss, Chris, decided to cut corners on the issue window and copied the MIPS R10K's simple issue-select scheme.[24] However, you believe you can come up with a better design in time for the impending tape-out (and also make the lawyers happy).

Issue-select is a critical part of an out-of-order processor. After decode and rename, micro-ops sit in the issue window waiting for their operands to become ready. Once ready, micro-ops then *request* to be issued. However, many more micro-ops may be requesting than can be issued. The **role of the issue-select policy is to choose which of the possible requesting micro-ops should be issued**. Your job for this question will be to design, implement, and experiment with different issue-select policies in `Chisel`.

Currently, BOOM mimics the MIPS R10K's policy - *micro-ops are selected based on their static position in the issue window and not based on age or other metrics.* Considering that Micro-ops are inserted into the issue window wherever an empty slot exists, one can imagine that this policy is very suboptimal![25]

An intern has recently added to BOOM a new "two-level issue-select" scheme, and he suggests you can use this to implement a better *issue-select* policy. Each issue slot now provides *two* request bits - a *request* signal that denotes the micro-op is requesting to be issued, and a *request_hp*

---

[24]This hypothetical is in no way auto-biographical.

[25]Chris should have first taken a quick look at the MIPS R12K, the successor to the R10K. If he had, he would have found out that one of the few things changed was the move to an aged-based issue-select policy!
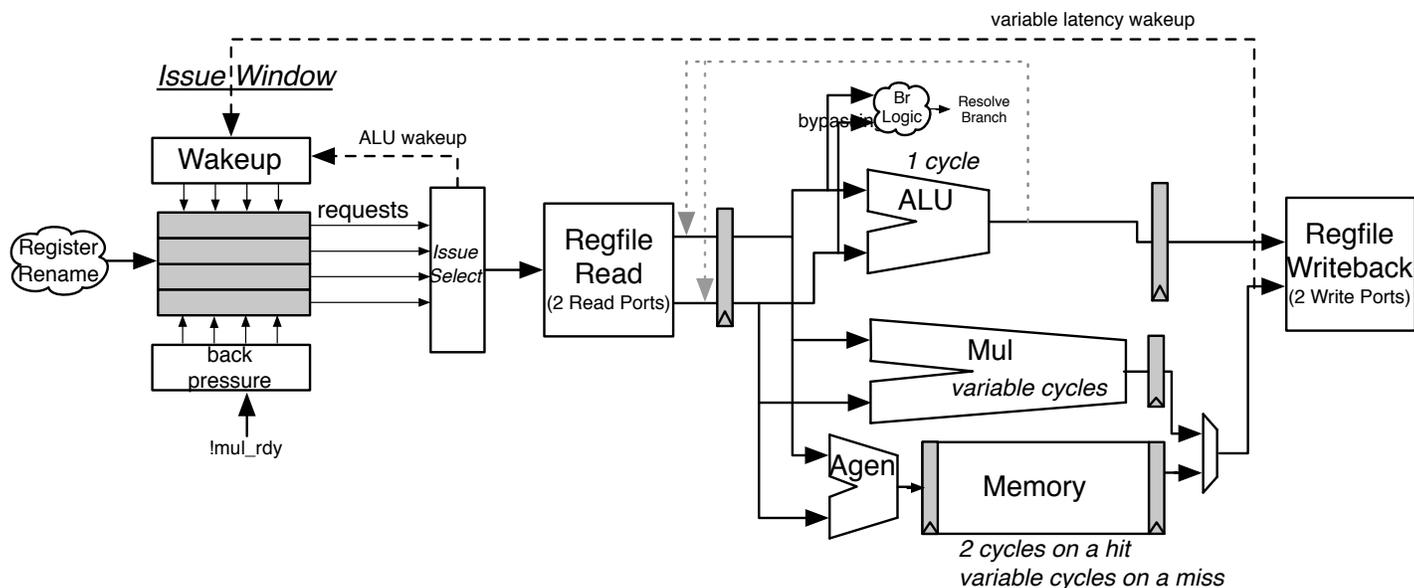
Figure 2: The issue logic and execution pipeline for a single-issue pipeline. For Question 3.1, you will modify the logic in the issue window slot.

signal to denote that the request is "high priority". The issue-select logic first selects amongst all "high priority" requestors before then filling any remaining issue ports from the "low priority" requestors.[26]

Your job is to modify the logic in the issue window slot that generates the *high_priority* signal. In doing so, you can change the policy of which micro-ops will get picked first by the issue-select logic!

There are a lot of possible schemes for choosing *high_priority*. The intern suggests that branches should be given *high_priority*. His theory is that mispredicted branches are expensive, so it is best to resolve branches as soon as possible to prevent too much misspeculated work from being executed. Another idea of his is to give priority to load instructions, as they are longer latency instructions that may require a trip to DRAM to resolve. Another co-worker suggests trying psuedo-age priority. For that, she explains, each micro-op counts how many cycles it's been sitting in the issue window; once it's been sitting for "too long" (a tunable parameter), it then asserts *high_priority*. "Randomly assigning *high_priority* may also be interesting!" exclaims another co-worker.

To implement giving branches high priority, open the `issue.scala` code in the BOOM source directory. Search for "High Priority Request" (around lines 150- this is the code for the issue slot). Set the *high_priority* variable to check if the slot's micro-op is a branch or jump:

```
high priority := slotUop.is_br_or_jmp
```

---

[26]A very common approach amongst modern, high-performance out-of-order processors is to use a "collapsing" queue. New instructions are inserted at the bottom of the queue, and the instructions are issued in priority starting from the top of the queue. As instructions are issued, the queue shifts all the instructions up to fill in the gaps. In this manner, the issue-select is biased to picking the oldest ready instructions first. As one can imagine, a collapsing queue is both very power hungry and very tough on the critical path of the processor. For BOOM a much simpler two-level scheme was chosen, which does not require the complications of moving micro-ops around.

Take a look in `dpath.scala` for the definition of the `MicroOp` bundle (around lines 120-170). There you can see all of the information stored in a micro-op bundle that may be useful for this problem (e.g., *is_load* tells you if the micro-op is a load instruction).

For this project you should explore a number of ideas (try and be creative). Some of your ideas (such as age-based schemes) will require a parameter sweep to find the optimal tuning parameters. Provide data on which schemes worked and which schemes did not, and explain your theories on why each scheme performed the way it did. And finally, which scheme do you propose to your team as what should be used for the next tape-out?

To help shed insight into the effects of issue-select logic, you should run your designs on two different versions of BOOM:

Table 7: New BOOM parameters for use with the issue-select policy question.

	Single Issue	Dual Issue
Issue Width	1	2
Fetch Width	2	2

In both cases, BOOM will be fetching/decoding/rename/dispatching two instructions every cycle. But for the single-issue version, one can imagine which micro-op you choose to issue will be even more important because of the mismatch between fetch bandwidth and issue bandwidth (lots of micro-ops should be getting backed up in the issue window!). See Appendix E for some starting information on writing `Chisel` code.

Caveats: this is a new feature running on real RTL running on very small benchmarks - we cannot promise that the results you get will be interesting! Also, the initial state of the hardware is random and as such, rerunning the benchmarks can provide different performance results each time. You should probably average multiple runs together and have error bars on your graphs to properly establish your designs are providing a statistically significant change in performance! And it's okay if your designs do not work out (we care more about your effort and insight).

3.2 Writing torture benchmarks: create code that exercises different features in the LSU.

The goal of this open-ended assignment is to purposefully design a set of benchmarks which stress different parts of BOOM. This problem is broken down into two parts:

- Write two benchmarks to stress the Load/Store Unit
- Write a benchmark(s) to introspect a parameter within BOOM

3.2.1 Part 1: Load/Store Unit Micro-benchmarks

You may have noticed that many of the benchmarks do not use all of the (very complicated) features in the Load/Store Unit. For example, few benchmarks perform any store data forwarding. For this part, you will implement two (small) benchmarks, each attempting to exercise a different characteristic.

- Maximize store data forwarding
- Maximize memory ordering failures

As a reminder, "store data forwarding" is when a load is able to use the data waiting in the store data queue (SDQ) before the store has committed (there is a `store->load` dependence in the program). A memory ordering failure is when a load that depends on a store (a `store->load` dependence) is issued to memory before the store has been issued to memory - the load has received the wrong data.

There is no line limit for the code used in this problem. Each benchmark must run for at least twenty thousand cycles (as provided by the SetStats() printout).

Two skeleton benchmarks are provided for you in `riscv-tests/benchmarks/lsu_forwarding/` and `riscv-tests/benchmarks/lsu_failures/`. To build and test them under the RISC-V ISA simulator:

```
inst$ cd ${LAB3ROOT}/riscv-tests/benchmarks/
inst$ make run
```

Once you are satisfied with your code and would like to run it on BOOM, type:

```
inst$ cd ${LAB3ROOT}/riscv-tests/benchmarks/
inst$ make; make run-riscv
```

... to compile your benchmarks and run them on the ISA simulator.[27]
Finally, you can run BOOM as usual:

```
inst$ cd ${LAB3ROOT}/emulator
inst$ make run-bmarks-test
```

---

[27]See Appendix F for more information on how to add a new benchmark (aside from `lsu_forwarding`, `lsu_failiures`, or `param_introspection`) to the benchmark suite and run it on BOOM.

Be creative! When you are finished, submit your code via Github. In your report, discuss some of the ideas you considered, and describe how your final benchmarks work.

Finally, it is possible that you may uncover bugs in BOOM through your stress testing: if you do, consider your benchmarking efforts a success! (save a copy of any offending code and let your TA know about any bugs you find).

How to use the benchmark templates We've provided a basic template for your benchmark codes in this section as C source files in `riscv-tests/benchmarks/lsu_forwarding/lsu_forwarding.c` and in analogous locations for the other two parts of this problem. If you open up this file, you'll notice that it contains a `main()` function that calls `setStats()` before and after calling a function called `benchmark()`:

```
//---------------------------------------------------------------------
// the main computations
//
// Call functions, run highly nested loops, write inline assembly... go crazy.
// Just be sure to look at the objdumps to make sure the compiler is not
// undoing your work.

__attribute__ ((noinline)) void benchmark()
{
   // maximize store-data forwarding
}
```

The comment instructs you to check that the compiler hasn't undone your work. This means you should examine the automatically-generated disassembly file, in this case `lsu_forwarding.riscv.dump` [28], and verify that the assembly code is what you'd expected.

If your code involves new functions, you might want to make sure the compiler isn't inlining them and optimizing away the function calls. This property is assured for the `benchmark()` function via the addition of the `__attribute__ ((noinline))` directive, and you can use this for your functions too. Assuming it's not optimized away, you can find your function in the disassembly by grepping for "`function-name>:`".

You should be able to use `printf`s to debug your code. It is highly recommended that you debug your benchmark using `make run-riscv` in the `benchmarks` directory before you try running on BOOM, since the former is a much faster method of testing.

3.2.2 Part 2: Parameter Introspection

Now the *real* challenge! Pick a non-binary parameter in BOOM's design and try to discover its value via a benchmark you design and implement yourself!

The basic strategy is as follows. Step 1) implement a micro-benchmark that stresses a certain parameter of the machine and measure the machine's performance. Step 2) go into `bsrc/consts.scala` to change the parameter you are studying, and rerun your benchmark. Step 3) Repeat to gather

---

[28]You can always use the RISC-V disassembly tool, `riscv-objdump -D lsu_forwarding.riscv`, to generate this output manually.

more results. Step 4) Build a model to describe how performance is affected by modifying your parameter.

Your model should be good enough that the TA can take your model and benchmark, run it on a machine and discover the value of the parameter in question without knowing its value a priori (even better if the TA can change other parameters of the machine so your model is not simply a lookup table).

Here are a set of parameters to choose from:[29]

- ROB size
- Number of physical registers
- Maximum number of branches
- Number of issue slots
- Number of entries in the load and store queues
- Number of entries in the fetch buffer
- Number of entries in the BHT
- Data cache associativity

Submit your code, describe how it works, and what ideas you explored. Also submit your data and your model showing how well it works on BOOM.

Naturally, this is a challenging task. The goal of this project is to make you think very carefully about out-of-order micro-architecture and write code to defeat the processor. There may not necessarily be a "clean" answer here.

Warning: not all parameters are created equally. Some will be harder challenges than others, and we cannot guarantee that all parameters will be doable. But with a dose of cleverness, you might be surprised what you can discover! (especially when you can white-box test your ideas).

---

[29]You may not use cache size(number of sets) as a parameter, as that is too easy.

4   Acknowledgments

This lab was originally developed for CS152 at UC Berkeley by Christopher Celio, and partially inspired by the previous set of CS152 labs written by Henry Cook.

References

[1] R.E. Kessler. The Alpha 21264 Microprocessor. *IEEE Micro*, 19(2):24–36, 1999.

[2] P. Rosenfeld, E. Cooper-Balis, and B. Jacob. Dramsim2: A cycle accurate memory system simulator. *Computer Architecture Letters*, 10(1):16 –19, jan.-june 2011.

[3] K.C. Yeager. The MIPS R10000 Superscalar Microprocessor. *IEEE Micro*, 16(2):28–41, 1996.

A   Appendix: The Issue Window

Figure 3 shows a single issue slot from the *Issue Window*.[30]

Instructions (actually they are "micro-ops" by this stage) are *dispatched* into the *Issue Window*. From here, they wait for all of their operands to be ready ("p" stands for *presence* bit, which marks when an operand is *present* in the register file).

Once ready, the *issue slot* will assert its "request" signal, and wait to be *issued*. Currently, BOOM uses a fixed priority encoding to give the lower ID entries priority. Question 3.1 covers a two-level issue select scheme in which uops asserting a *high_priority* request are selected first, before choosing from the remaining, lower priority requestors.
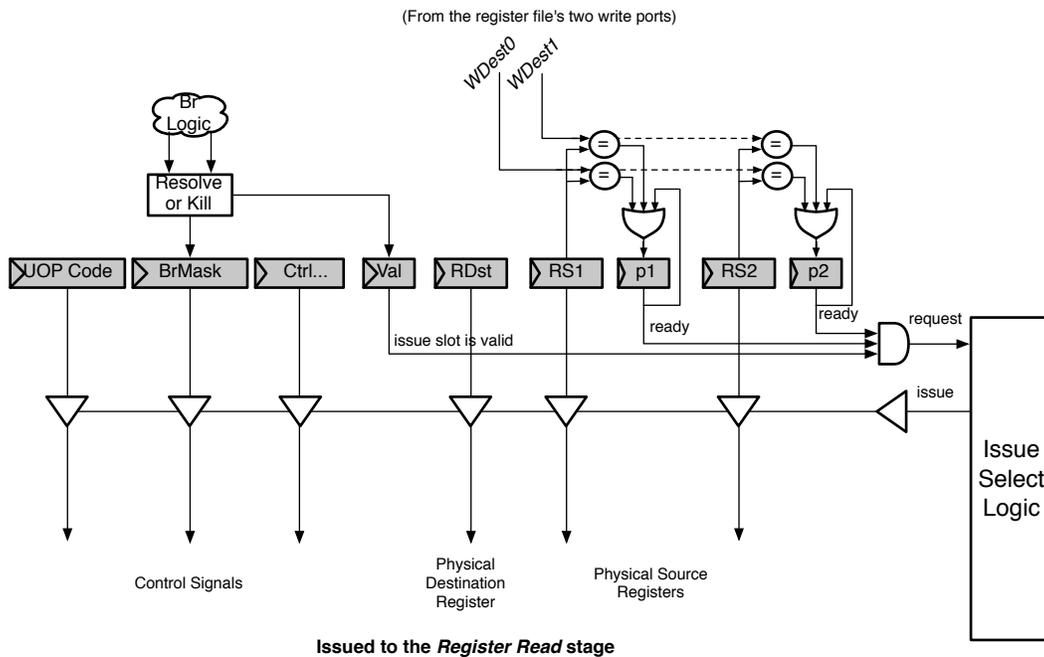
Figure 3: A single issue slot from the Issue Window.

---

[30]Conceptually, a bus is shown for implementing the driving of the signals sent to the *Register Read* Stage. In reality, for now anyways, BOOM actually uses muxes.

B   Appendix: The BOOM Source Code

The BOOM source code can be found in `{LAB3ROOT}/bsrc`.

Some of the code structure is shown below:

- `bsrc/`
  - consts.scala All constants and adjustable parameters.
  - tile.scala The tile, instantiates memory and the core.
  - core.scala The top-level of the processor core component.
  - issue.scala The issue-window and associated issue-select logic.
  - icache.scala Instruction cache.
  - dpath.scala Main chunk of the BOOM datapath and control code.
  - brpredictor.scala Branch predictor. Uses a table of n-bit history counters.
  - prefetcher.scala Data prefetcher.
  - decode.scala Decode table.
  - rename.scala Register renaming logic, map tables, and free-list logic.
  - rob.scala Re-order Buffer.
  - lsu.scala Load/Store Unit.
  - dcachewrapper.scala Instantiates the DC and translates into OoO-speak.
  - functional_unit.scala Encapsulates the functional units.
  - execute.scala Logic that ties functional units together into execution units.
  - util.scala Some helpful utility functions.

C   Appendix: How to Read Chisel Signals in the C++ Test-Harness Code

In this lab, we will be exercising the C++ tool-flow of `Chisel` (`Chisel` can also emit a Verilog version of a design). Often, whether for debugging purposes or for instrumentation, we will often want to probe the state of a Chisel design from the C++ test-harness.

As an example, let's probe the "micro-op opcode" signal ("uop code", or "uopc") that is stored in the *issue slot* of the *Issue Window* (see Figure 3). If we look through the `Chisel` code of BOOM, we see that the `IntegerIssueSlot` component is what describes the *issue slot*, and that it contains the variable `slotUop`.[31] The `slotUop` is a `Chisel` "Bundle", or group of signals (think "struct" in C). We can see the definition of the "MicroOp" bundle in `dpath.scala`, near lines 50-100. One of the field variables within the bundle is `uopc`: this is the "micro-op opcode". Thus, the "uopc" within the *issue slot* would be "slotUop.uopc" in `Chisel`, or `slotUop_uopc` once it has been generated into `C++` code. The `slotUop` signal is a `Reg` type, or *register*, and is written to on the positive-edge of the clock signal when the *issue slot*'s *write-enable* signal is asserted.

Each `IntegerIssueSlot` component is instantiated inside the `DatPath` component, which itself is instantiated inside the `Core` component, which in turn is instantiated inside the `BoomTile` component, which in turn is instantiated inside the `Top` component (phew!). When `Chisel` generates the resulting C++ code, the signal `slotUop_uopc` contains its entire parentage in its name-mangled C++ name.

---

[31]The code for the issue slot can be found in `bsrc/dpath.scala`, near lines 200-300.

C.1 Finding the C++ Variable

The best way to find the C++ variable name for `slotUop_uopc` is to look through the gener-
ated C++ code in `${LAB3ROOT}/emulator/generated-src/Top.h`, which holds *all* Chisel signals.
Grepping for `slotUop_uopc` we find the variables:

```
dat_t<9> Top_BoomTile_core_dpath_issue_unit_IntegerIssueSlot_6__slotUop_uopc;
dat_t<9> Top_BoomTile_core_dpath_issue_unit_IntegerIssueSlot_6__slotUop_uopc_shadow;

dat_t<9> Top_BoomTile_core_dpath_issue_unit_IntegerIssueSlot_5__slotUop_uopc
Top_BoomTile_core_dpath_issue_unit_IntegerIssueSlot_5__slotUop_uopc_shadow;
etc....
```

First, since there are 8 issue slots in BOOM by default, we will find 4 chunks of "slotUop_uopc"
signals. `Chisel` will automatically add 1,2,3... to the component's name when it finds multiple
instantiations of it. Second, we see the full path name to `slotUop_uopc`: the top-level module is
"Top", followed by "BoomTile", "core", "d" (for datapath), and finally "IntegerIssueSlot."

Third, we see an additional version of the `slotUop_uopc` variable: a `_shadow` version. You can
safely ignore these variables.[32]

C.2 Reading out the value from the C++ Variable

Although we have now found the variable we are interested in
(`Top_BoomTile_core_dpath_issue_unit_IntegerIssueSlot_0__slotUop_uopc`,
`Top_BoomTile_core_dpath_issue_unit_IntegerIssueSlot_1__slotUop_uopc`, etc.), we can see
that it is of type `dat_t<9>`. This is a special templated class type that encapsulates all `Chisel`
variables. In this case, it is describing an 9-bit wide value. The problem is we may occasionally
want to describe variables of over 128 bits in our `Chisel` design, but natively C and C++ can only
handle double the size of the native host machine's register. Thus, `Chisel` uses its own data-type
class which maps to an array of `uint64_t` variables under the hood.

The important thing to know is that we can use the function `.lo_word()` to pull out the lowest
64-bits from a `dat_t<>` variable.

```
Top_t *tile = new Top_t(); // instantiate our Chisel design
uint64_t slot0_uopc = tile->Top_BoomTile_core_dpath_issue_unit_IntegerIssueSlot_0_slotUop_uopc
uint64_t slot1_uopc = tile->Top_BoomTile_core_dpath_issue_unit_IntegerIssueSlot_1__slotUop_uop
etc...
```

D   Appendix: The Load/Store Unit

The Load/Store Unit is responsible for deciding when to fire memory operations to the memory
system. There are three queues: the Load Address Queue (LAQ), the Store Address Queue (SAQ),
and the Store Data Queue (SDQ). Load instructions generate a "uopLD" micro-op. When issued,
"uopLD" calculates the load address and places its result in the LAQ. Store instructions generate
two micro-ops, "uopSTA" (Store Address Generation) and "uopSTD" (Store Data Generation).

---

[32]They exist because `slot_uopc` is a *register*. On `clock_lo` the *shadow* version is updated, and on `clock_hi` the
regular version is updated to the *shadow* version's value.

The STA micro-op calculates the store address and places its result in the SAQ queue. The STD micro-op moves the store data from the register file to the SDQ. Each of these micro-ops will issue out of the *Issue Window* as soon their operands are ready.

D.1 Store Instructions

Entries in the Store Queue[33] are allocated in the *Decode* stage (the appropriate bit in the `stq_entry_val` vector is set). A "valid" bit denotes when an entry in the SAQ or SDQ holds a valid address or data (`saq_val` and `sdq_val` respectively). Store instructions are fired to the memory system at *Commit*; the ROB notifies the Store Queue when it can fire the next store. By design, stores are fired to the memory in program order.

D.2 Load Instructions

Entries in the Load Queue (LAQ) are allocated in the *Decode* stage (`laq_entry_val`). In *Decode*, each load entry is also given a *store mask* (`laq_st_mask`), which marks which stores in the Store Queue the given load depends on. When a store is fired to memory and leaves the Store Queue, the appropriate bit in the *store mask* is cleared.

Once a load address has been computed and placed in the LAQ, the corresponding *valid* bit is set (`laq_val`).

Loads are optimistically fired to memory on arrival to the LSU (getting loads fired early is a huge benefit of out–of–order pipelines). Simultaneously, the load instruction compares its address with all of the store addresses that it depends on. If there is a match, the memory request is killed. If the corresponding store data is present, then the store data is *forwarded* to the load and the load marks itself as having *succeeded*. If the store data is not present, then the load goes to *sleep*. Loads that have been put to sleep are woken up later. Once the sleeping load is at the head of the LAQ, it is retried.[34]

D.3 Memory Ordering Failures

The Load/Store Unit has to be careful regarding `store->load` dependences. For the best performance, loads need to be fired to memory as soon as possible.

```
sw x1 -> 0(x2)
ld x3 <- 0(x4)
```

However, if `x2` and `x4` reference the same memory address, then the load in our example *depends* on the earlier store. If the load issues to memory before the store has been issued, the load will read the wrong value from memory, and a *memory ordering failure* has occurred. On an ordering failure, the pipeline must be flushed and the rename map tables reset. This is an incredibly expensive operation.

To discover ordering failures, when a store address is computed and enters the SAQ, it checks the entire LAQ for any address matches. If there is a match, the store checks to see if the load

---

[33]When I refer to the *Store Queue*, I really mean both the SAQ and SDQ.

[34]Higher performance processors will track *why* a load was put to sleep and wake it up once the blocking cause has been alleviated.

has *executed*, and if it got its data from memory or if the data was forwarded from an older store. In either case, a memory ordering failure has occurred. The ordering failure information is sent to the ROB and the load is marked as having thrown an exception (a memory ordering failure is treated as an exception, as it requires clearing the pipeline, resetting the map tables, and retrying the load).

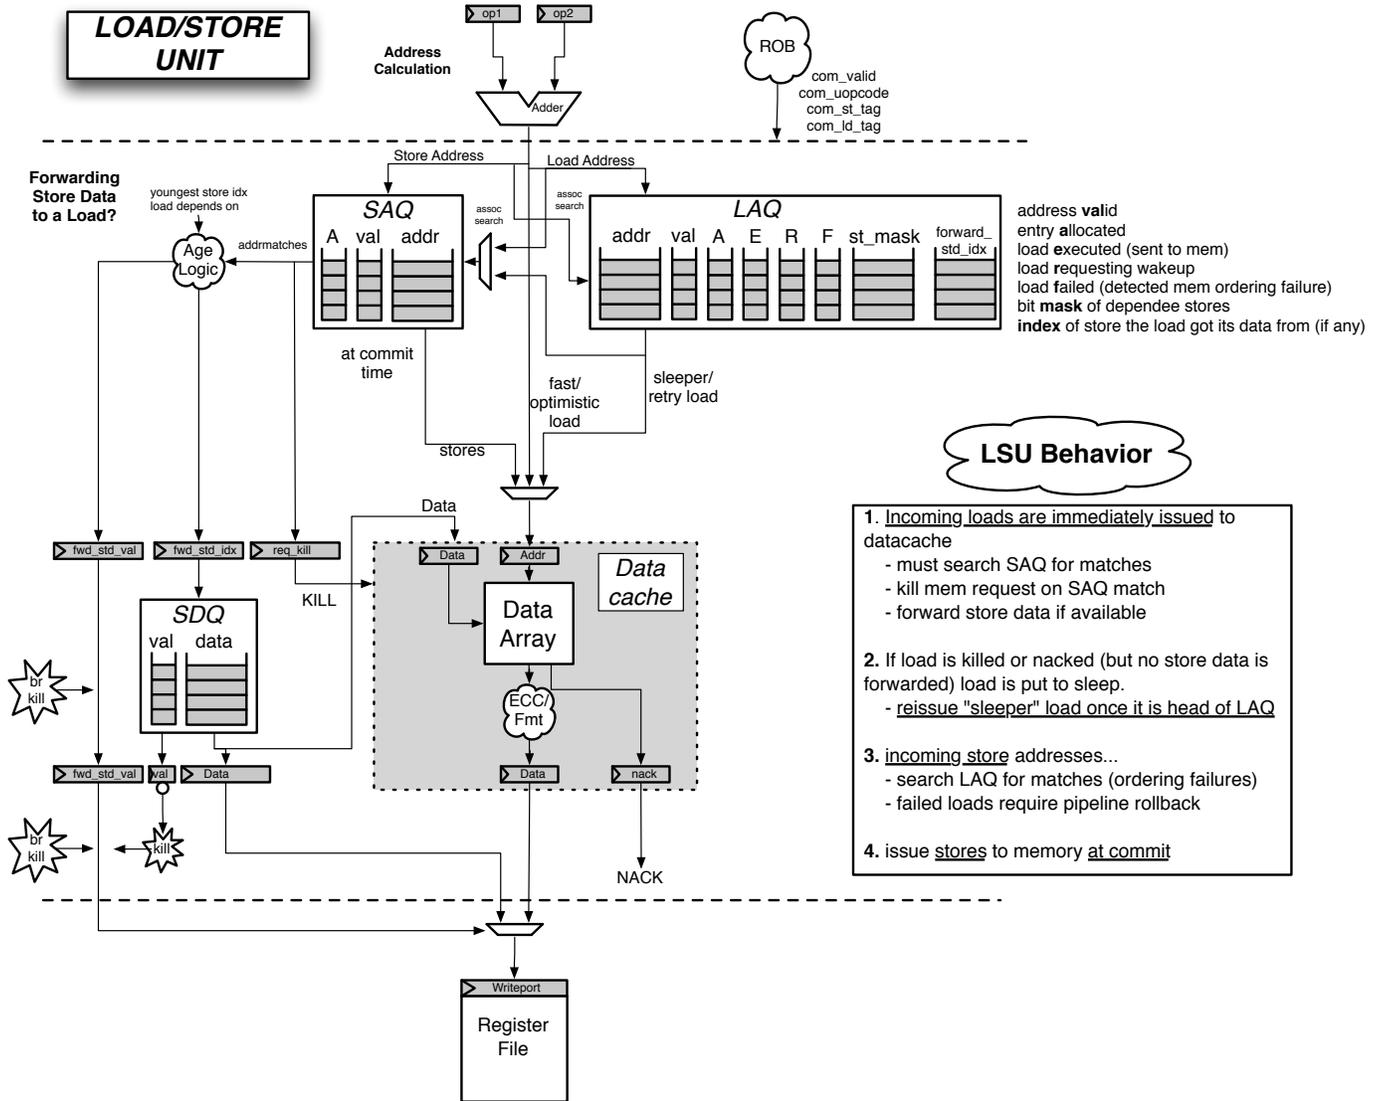See Figure 4 for more information about the Load/Store Unit.

LOAD/STORE UNIT

op1   op2

Address Calculation

Adder

ROB

com_valid
com_uopcode
com_st_tag
com_ld_tag

Store Address   Load Address

Forwarding Store Data to a Load?

youngest store idx
load depends on

Age Logic   addrmatches

SAQ
A   val   addr

assoc search

assoc search

LAQ
addr   val   A   E   R   F   st_mask   forward_std_idx

address **val**id
entry **a**llocated
load **e**xecuted (sent to mem)
load **r**equesting wakeup
load **f**ailed (detected mem ordering failure)
bit **mask** of dependee stores
index of store the load got its data from (if any)

at commit time

fast/ optimistic load

sleeper/ retry load

stores

LSU Behavior

Data

fwd_std_val   fwd_std_idx   req_kill

KILL

SDQ
val   data

br kill

fwd_std_val   val   Data

br kill   kill

Data   Addr

Data cache

Data Array

ECC/ Fmt

Data   nack

NACK

1. Incoming loads are immediately issued to datacache
  - must search SAQ for matches
  - kill mem request on SAQ match
  - forward store data if available

2. If load is killed or nacked (but no store data is forwarded) load is put to sleep.
  - reissue "sleeper" load once it is head of LAQ

3. incoming store addresses...
  - search LAQ for matches (ordering failures)
  - failed loads require pipeline rollback

4. issue stores to memory at commit

Writeport

Register File

Figure 4: The Load/Store Unit.

23

E   Appendix: Some Notes on Writing `Chisel` Code

Some questions in this lab require writing `Chisel` code. Here is a short write-up on getting started. See `https://chisel.eecs.berkeley.edu` for more information.

Chisel has three important signal types. One is `Bool()`, which unsurprisingly holds a single boolean value. Use this for true/false control signals and state. The `UInt(width=n)` type holds an unsigned integer that is `n` bits wide. Raw bits that don't have any meaning as an integer should be stored with the `Bits(width = n)`. New signals can be declared as follows:

```
val my_bits = Bits(width = 4)
my_bits := Bits(0) // hard-wire my_bits to zero
val msb_is_hi = my_bits(3) === Bits(1)
```

In the above example, the `my_bits` is declared as a wire of width 4, and then hard-wired to 0. The `Chisel` operator `:=` *wires* one signal to another, whereas the single `=` is a *Scala* operator that assigns a value to a Scala variable. So the above code declares a new Scala value, `my_bits`, that contains a new instance of a `Chisel` wire of width four.

The `msb_hi` signal is a `Bool()` that is true when the MSB of `my_bits` is hi (so, in this case, it's always false).

The operator `signal(i, j)` extracts the `i`th through `j`th bits of `signal`. `my_bits(3, 0)` would extract *all* bits of `my_bits`, for example. You can use the triple-equals comparison operator `===` to check if two signals are equal to each other.

To add new registers, you can use the `Chisel Reg()` keyword:

```
val counter = Reg(init=UInt(0, width=4))
counter := counter + UInt(1)
```

This code declares a register called `counter` that stores 4-bit unsigned integer values, and then says that the register is to be incremented by one on every clock cycle. Because the register is only four bits wide, it will overflow to zero in 16 cycles. You can control when new values are assigned to a register using a `when` statement:

```
when(slotUop.is_ret) {
  counter := counter + UInt(1)
}
```

This only increments the counter when the uop in its slot is a function return instruction.

This should be enough `Chisel` advice to get you through this problem. You'll likely find it helpful to glance at some other `Chisel` code while you work, particularly in issue.scala and dpath.scala. Note how the fields of the `MicroOp` bundle are declared at the beginning of dpath.scala. You can find more documentation on `Chisel` at `https://chisel.eecs.berkeley.edu/`.

F   Appendix: Adding a New Benchmark

For some of these questions, you will either need to create new benchmarks, or add existing benchmarks to the build system.

F.1 Creating a new benchmark

The source code for all benchmarks can be found in `test/riscv-bmarks/`. Each benchmark is given its own directory. To create a new benchmark, it is easiest to copy an existing benchmark directory.

```
inst$ cp -r vvadd my_new_bench
inst$ cd my_new_bench; ls
bmark.mk dataset1.h vvadd_gendata.pl vvadd_main.c
```

First, open `bmark.mk`. You will want to find and replace all instances of "vvadd" with "my_new_bench". If your benchmark requires multiple C files, add them to `vvadd_c_src`. Likewise, any assembly files can be added to `vvadd_riscv_src`. The build system should be able to take care of actually building and linking your different source files.

The vvadd benchmark has a Perl script `vvadd_gendata.pl` to generate a random input set of arrays, which are stored in `dataset1.h`. This removes the processor from having to generate and test its own input vectors. You can safely ignore these files (or repurpose them for your own use).

The `vvadd_main.c` file holds the main code for vvadd. Rename this file to fit the file name declared in `bmark.mk` (probably `my_new_bench_main.c`). Now you can add your own code in here. You can delete pretty much everything except for the vital `setStats(int)` function. The `setStats(int)` turns on and off the statistic tracking used by `oootracer.*`. *Note: setStats() should only be called twice in your benchmark, as only the last set of data will be stored by oootracer.*. Note also that the value returned from `main()` will indicate to the proxy-kernel whether or not the test was a success (return 0 for success or an error code of your choice on failure).

F.2 Adding a benchmark to the build system

Once you are happy with your new benchmark, you need to modify two `Makefile`s. First, open `riscv-test/benchmarks/Makefile`, and find the `bmarks` variable. Add "my_new_bench" to the listing. You can now build your benchmark and test it on the ISA simulator.

```
inst$ make; make run;
```

Once you are satisfied, you must "install" the benchmark to `install/riscv-bmarks`. This is where BOOM looks for benchmarks to run.

```
inst$ make install
```

One final `Makefile` modification is required. Open `${LAB3ROOT}/Makefrag` and find the variable `global_bmarks`. Add your benchmark to this variable as well but make sure to add the extension `.riscv`. Now running `make run-bmarks-test` in `emulator/` will run your benchmark on BOOM as well!
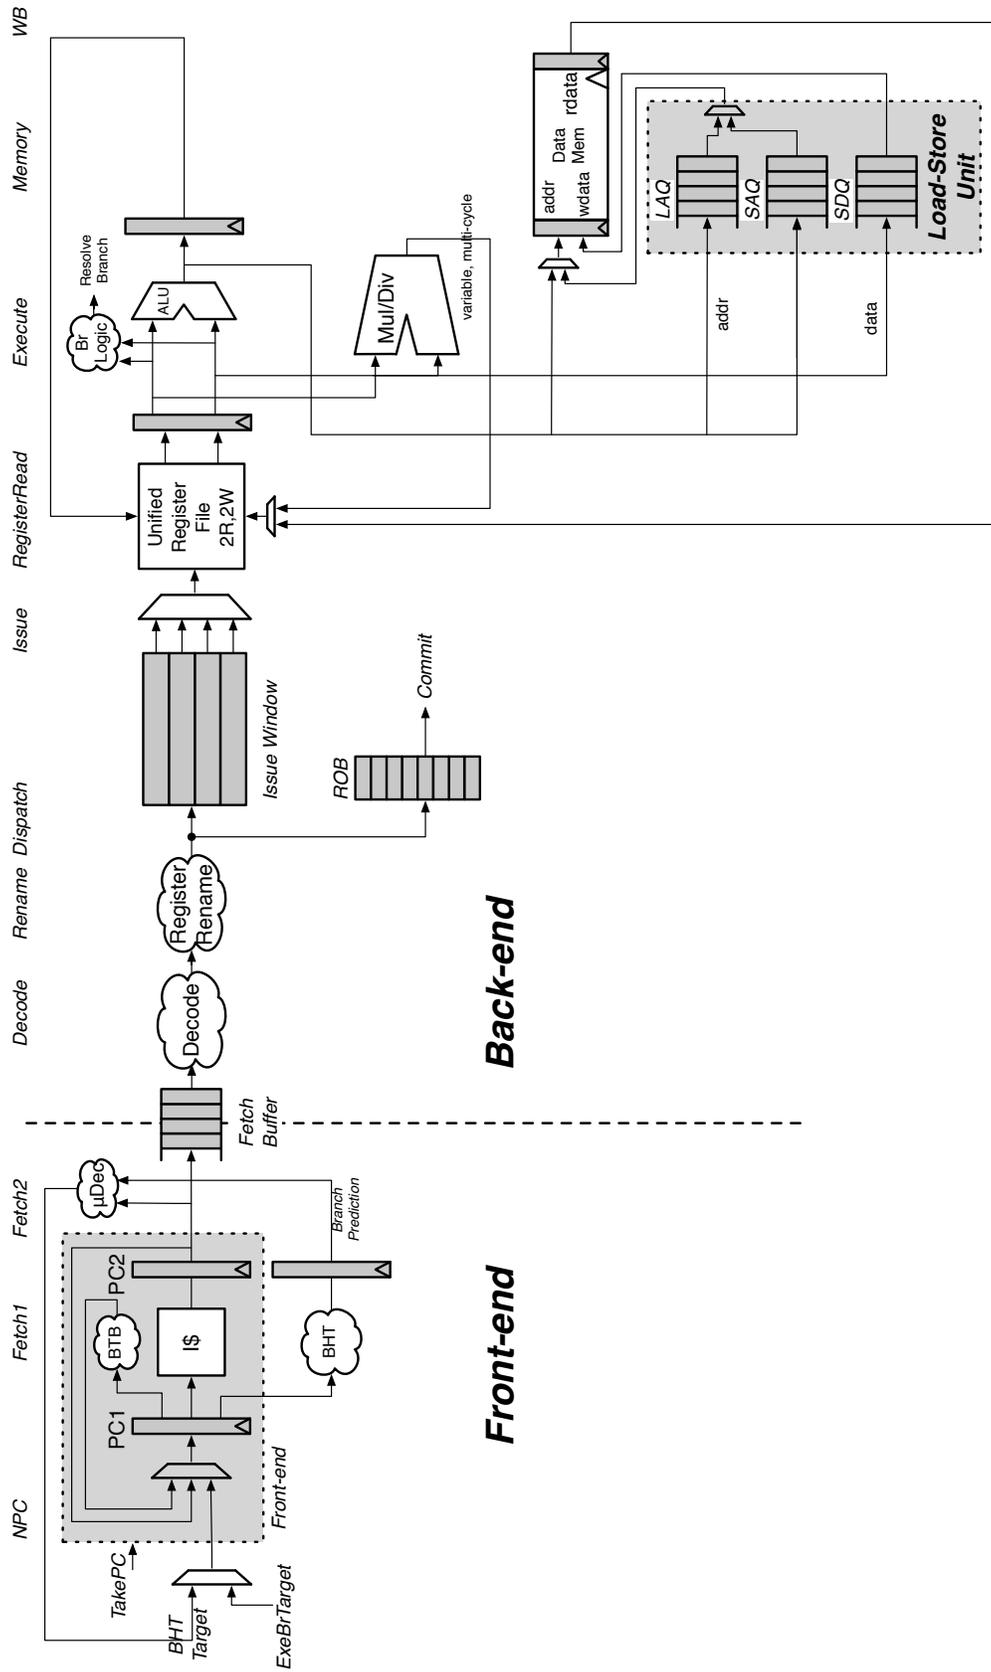
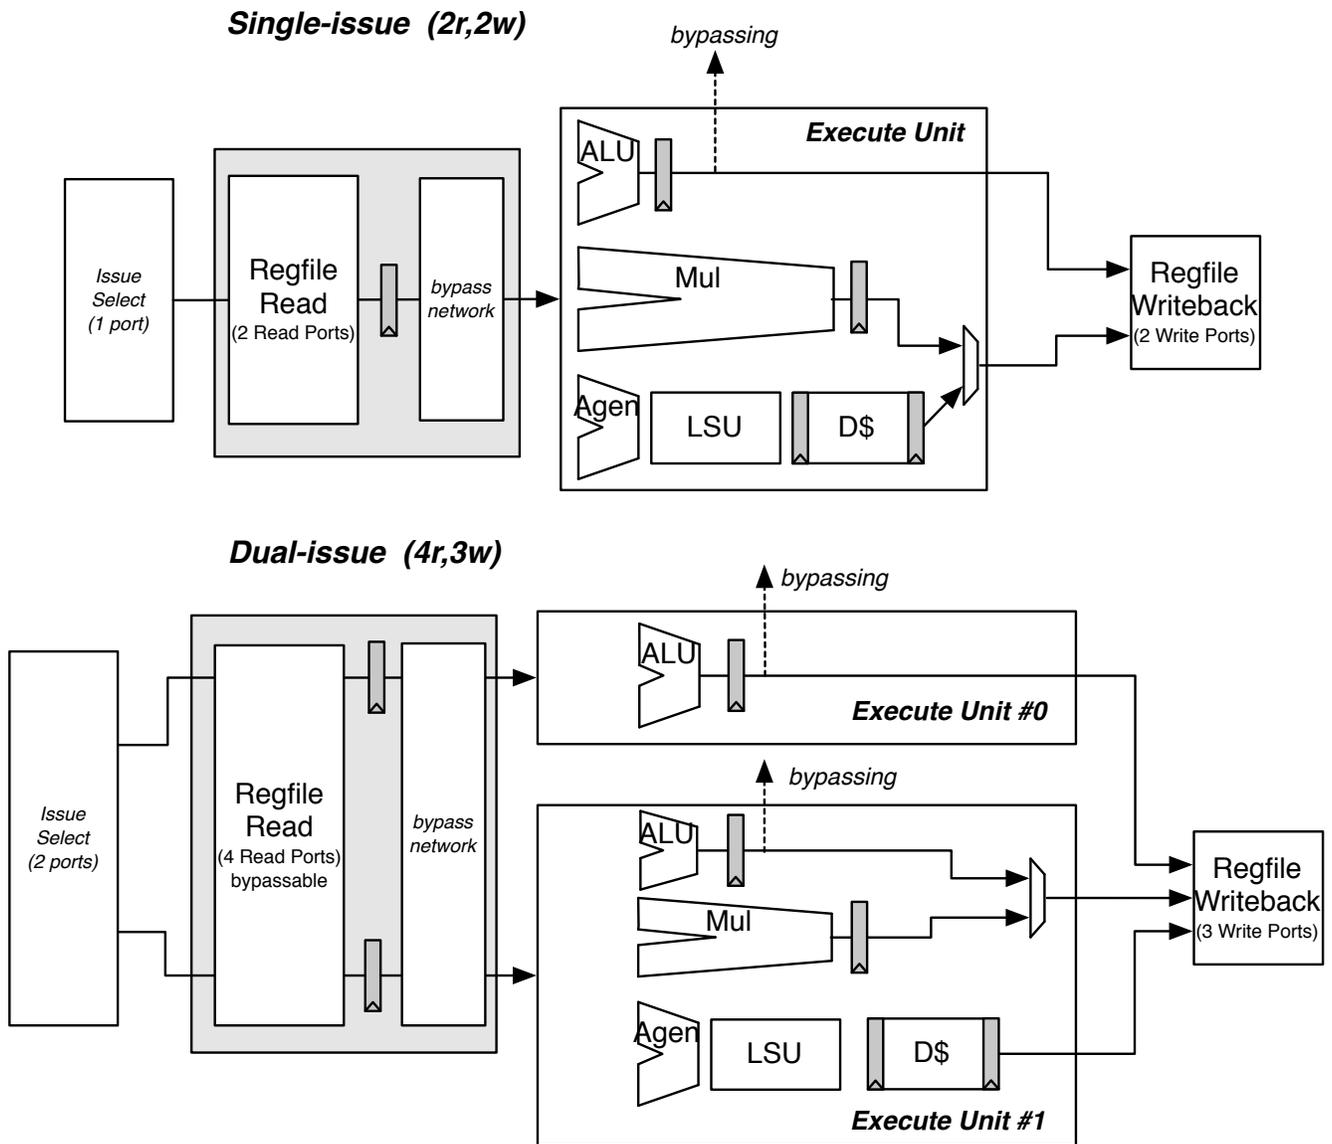Figure 5: A more detailed diagram of BOOM.

Figure 6: The issue logic and execution pipeline for single-issue and dual-issue pipelines. The issue logic can only issue one micro-op to each execution unit, but within each execution unit are multiple functional units. For example, the dual-issue pipeline can handle two ALU micro-ops per cycle, or one ALU micro-op and one memory micro-op per cycle, but not two memory micro-ops. Also notice that only the ALU functional units allow for the bypassing of back-to-back instructions.