

CS 152 Laboratory Exercise 1 (Version A)

Professor: Krste Asanović
TA: Yunsup Lee
Department of Electrical Engineering & Computer Science
University of California, Berkeley

January 29, 2013

1 Introduction and goals

The goal of this laboratory assignment is to familiarize you with the `Chisel` simulation environment while also allowing you to conduct some simple experiments. By modifying an existing instruction tracer module, you will collect instruction mix statistics and make some architectural recommendations based on the results.

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 or three. 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.

1.1 Graded Items

You will turn in a hard copy of your results to the professor or TA. Please label each section of the results clearly. The directed items need to be turned in for evaluation. You only need to turn in *one* of the problems found in the Open-Ended Portion.

1. (Directed) Problem 4.4: recorded instruction mixes for each benchmark and answers
2. (Directed) Problem 4.5: thought problem answers
3. (Directed) Problem 4.6: CPI analysis answers
4. (Directed) Problem 4.7: design problem answers
5. (Open-ended) Problem 5.1: source code and recorded ratio
6. (Open-ended) Problem 5.2: data and the modified section of `Chisel` source code
7. (Open-ended) Problem 5.3: instruction definition, test code, worksheet, modified section of `Chisel` source code
8. (Open-ended) Problem 5.4 and Problem 5.5: design proposals and supporting data
9. (Directed) Problem 6: Feedback on this lab

Lab reports must be in *readable* English and not raw dumps of log-files. It is *highly* recommended that your lab report be typed. Charts, tables, and figures - when appropriate - are great ways to succinctly summarize your data.

2 The RISC-V Instruction Set Architecture

The processors in this lab that you will be studying implement the RISC-V ISA, recently developed at UC Berkeley for use in education and research.

An entire tool-chain is provided. The `riscv-gcc`, `riscv-g++` cross-compilers build new binaries from RISC-V assembly, C, and C++ source codes. The `riscv-objdump` tool can disassemble existing RISC-V binaries to show the exact sequence of instructions being executed.

The ISA simulator `riscv-isa-sim` can execute RISC-V binaries. The ISA simulator serves as the golden reference for the ISA. It is not cycle-accurate, but it executes very quickly. The front-end server, `fesvr`, loads a RISC-V binary and connects to either the ISA simulator or a `Chisel`-created simulator.

The RISC-V ISA manual can be found in the “resources” section of the CS 152 website or directly at <http://www.eecs.berkeley.edu/Pubs/TechRpts/2011/EECS-2011-62.html>. For Lab 1, all processors implement the 32-bit variant, known as RV32.

3 Chisel

`Chisel` is a new *hardware construction language* developed at UC Berkeley for the rapid design and development of hardware. `Chisel` raises the level of abstraction by allowing designers to utilize concepts such as object orientation, functional programming, parameterized types, and type inference. Unlike HDL languages such as Verilog which were designed first to be *simulation* languages, `Chisel` is designed *specifically* for constructing actual hardware.

`Chisel` can generate both low-level Verilog code, for mapping designs to FPGA and ASICs, and high-speed C++-based cycle-accurate software simulators.

`Chisel`, an acronym for Constructing Hardware In a Scala Embedded Language, is a domain-specific language embedded inside of Scala. `Chisel` code describing a processor is actually a legal Scala program whose execution outputs either Verilog code or C++ code. Figure 1 shows the design flow of `Chisel`. For more details, please visit <http://chisel.eecs.berkeley.edu>.

3.1 Chisel in This Lab

Provided with this lab are four different processors: a 1-stage pipeline, a 2-stage pipeline, a 5-stage pipeline, and a micro-coded pipeline. All are implemented in `Chisel`.

In this lab, you will compile the provided `Chisel` processors into C++ software simulators, and use the C++ simulators to quickly run cycle-accurate experiments regarding instruction mixes and pipeline hazards. A tutorial on the `Chisel` language can be found at <http://chisel.eecs.berkeley.edu>. Students will not be required to write `Chisel` code as part of this lab, beyond changing and adding parameters as directed.

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.




Figure 1: Chisel Design Flow. Only the C++ Simulator path will be exercised for this lab.

4 Directed Portion (2/7 of lab grade)

4.1 Terminology and conventions

Host machine refers to the machine on which you run the simulators. For this lab, the *Host machine* will be an Instructional machine (`inst$`). *Target machine* refers to the simulated machine. In this case, the provided RISC-V processors are your target machines.

4.2 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. We will provide you with an instructional computing account for this purpose.

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`. (see <http://inst.eecs.berkeley.edu/cgi-bin/clients.cgi?choice=servers> for more information about available machines).

First, download the lab materials¹:

```
inst$ cd ~
inst$ cp -R ~/cs152/sp13/lab1 .

inst$ cd lab1
inst$ export LAB1ROOT=$PWD
```

We will refer to `~/lab1` as `#{LAB1ROOT}` in the rest of the handout to denote the location of the Lab 1 directory.

The directory structure is shown below:

- `#{LAB1ROOT}/`
 - `Makefile`
 - `src/ Chisel source code for each processor.`
 - * `common/ Common source code shared between all processors.`
 - * `rv32_1stage/ Source code for the RISC-V 1-stage processor`
 - * `rv32_2stage/ ...`
 - * `rv32_5stage/ ...`
 - * `rv32_ucose/ ...`
 - `emulator/`
 - * `common/ Common emulation infrastructure shared between all processors.`
 - * `rv32_1stage/ C++ simulation tools and output files.`
 - * `rv32_2stage/ ...`
 - * `rv32_5stage/ ...`
 - * `rv32_ucose/ ...`
 - `test/ Local source code for benchmarks and tests.`
 - * `riscv-bmarks/ Local benchmarks written in C.`
 - * `riscv-tests32/ Local tests written in assembly.`
 - `install/ Compiled assembly tests and benchmarks.`
 - `chisel/ Chisel voodoo. You can safely ignore this directory.`
 - `sbt/ Scala voodoo. You can safely ignore this directory.`
 - `uncore/ Chisel source code for the uncore. You can safely ignore this directory.`

Of particular note is that the source code for the `Chisel` processors can be found in `#{LAB1ROOT}/src/`. While you do not have to understand the code to do this assignment, it may be interesting to see the entire workings of a processor. While it is not recommended that you modify any of the processors while collecting data for them in the directed lab portion (except as directed), feel free in your own time (or perhaps as part of the open-ended portion) to change and tweak the processors as you see fit.

¹The capital “R” in “cp -R” is critical, as the -R option maintains the symbolic links found in the `#{LAB1ROOT}` directory.




Figure 2: The Testing Environment. The front-end server (`fesvr`) loads the RISC-V binary from the `Host` file system, starts the `Target` system simulator, and sends the RISC-V binary code to the `Target` simulator to populate the simulated processor’s instruction memory with the program. Once the `fesvr` finishes sending the test code, the `fesvr` sends the “Start” command and the `Target` processor begins execution.

4.3 First Steps: Building the 1-Stage Processor

In this lab, four different processors are provided: a 1-stage processor, a 2-stage processor, a 5-stage processor, and a micro-coded processor. The 5-stage processor implements both a fully-bypassed pipeline and a no-bypassing/fully interlocked pipeline.

The following command will set up your bash environment, giving you access to the entire CS152 lab tool-chain. Run it before each session:²

```
inst$ source ~cs152/sp13/cs152.bashrc
```

Navigate to the `chisel/` directory and execute the following command:

```
inst$ cd ${LAB1ROOT}/
inst$ make run-emulator
```

If this is your first time running `sbt`, this command may take a while.³ The command `make run-emulator` does the following:

- runs `sbt`, the Scala Built Tool, selects the `rv32_1stage` project, and runs the `Chisel` code which generates a C++ cycle-accurate description of the processor. The generated C++ code can be found in `${LAB1ROOT}/emulator/rv32_1stage/generated-src/`.
- compiles the generated C++ code into a binary called `emulator`.
- calls the RISC-V front-end server (called `fesvr`), which opens a socket to your C++ binary `emulator`, and sends it a RISC-V binary for the target processor to execute (See Figure 2). All of the RISC-V tests and benchmarks will be executed when calling “`make run-emulator`”.⁴

A `PASSED` should be generated by each program. If you see any `FAILED`, verify you are running on a recommended instructional machine. Otherwise, contact your TA.

²Or better yet, add this command to your bash profile.

³If you get a `java.lang.OutOfMemoryError` exception, run `make run-emulator` again.

⁴Which tests and benchmarks are executed can be found in the `${LAB1ROOT}/emulator/common/Makefile.include`.

Building Other Processors

By default, `make run-emulator` builds the 1-stage processor (`rv32_1stage`). To build the other processors:

```
inst$ cd ${LAB1ROOT}/
inst$ export MK_TARGET_PROC=rv32_2stage
inst$ make run-emulator
```

The bash variable `${MK_TARGET_PROC}` tells the Makefile which processor you want to build. The valid options are:

```
inst$ export MK_TARGET_PROC=rv32_1stage
inst$ export MK_TARGET_PROC=rv32_2stage
inst$ export MK_TARGET_PROC=rv32_5stage
inst$ export MK_TARGET_PROC=rv32_ucose
```

4.4 Instruction Mix Tracing Using the 1-Stage Processor

For this section of the lab you will track the instruction mixes of several RISC-V benchmark programs provided to you.

```
inst$ cd ${LAB1ROOT}/
inst$ export MK_TARGET_PROC=rv32_1stage
inst$ make run-emulator
inst$ cd ${LAB1ROOT}/emulator/rv32_1stage/
inst$ ls
inst$ vim vvadd.riscv.out
```

We have provided a set of benchmarks for you to gather results from: `median`, `multiply`, `qsort`, `towers`, and `vvadd`. Using your editor of choice, look at the output files generated from `make run-emulator`.⁵ The processor state is written to the output file on every cycle. At the end of the file, statistics from the “Trace” object can be found:

⁵To speed up parsing data out of all of the benchmark output files, type “`grep \/# *.riscv.out`” to dump all trace information to stdout.

```

#----- Tracer Data -----
#
#      CPI   : 1.00
#      IPC   : 1.00
#      cycles: 3029
#
#      Bubbles      : 0.000 %
#      Nop instr    : 0.000 %
#      Arith instr  : 59.756 %
#      Ld/St instr  : 30.175 %
#      branch instr: 9.937 %
#      misc instr   : 0.132 %
#-----

```

A few things to note: software compiler-generated NOPs *do* count towards the instruction count but machine-inserted “bubbles”⁶ *do not*, and “Tracer” only counts data when the processor’s “Stats Enable” register is set (control register “cr10”). Also, the denominator used for calculating the percentages is “cycles.”

Note how the mix of different types of instructions vary between benchmarks. Record the mix you observed for each benchmark (remember: don’t provide raw dumps!). Which benchmark has the highest arithmetic intensity? Which benchmark seems most likely to be memory bound? Which benchmark seems most likely to be dependent on branch predictor performance?⁷

4.5 CPI analysis problem

Consider the results gathered from the RV32 1-stage processor. Suppose you were to design a new machine such that the average CPI of loads and stores is 2 cycles, integer arithmetic instructions take 1 cycle, and other instructions take 1.5 cycles on average. What is the overall CPI of the machine for each benchmark?

What is the relative performance for each benchmark if loads/stores are sped up to have an average CPI of 1 cycle? Is this still a worthwhile modification if it means that the cycle time is increased 30%? Is it worthwhile for all benchmarks, or only some? Explain.

4.6 CPI Analysis Using the 5-Stage Processor

For this section we will analyze the effects of branching and bypassing in a 5-stage processor.⁸

The 5-stage processor provided in this lab has been parameterized to support both full-bypassing (but must still stall for load-use hazards) and fully-interlocked. The fully-interlocked variant provides *no* bypassing, and instead must stall (interlock) the `instruction fetch` and `decode` stages until all hazards have been resolved.

⁶A “bubble” is inserted, for example, when the 2-stage processor takes a branch and must kill the Instruction Fetch stage.

⁷If you would like to see the disassembly of any benchmark, you can visit `~cs152/install/riscv-bmarks/`, and view the `*.riscv.dump` files. You can also use `riscv-objdump` to create your own disassembly files.

⁸The 2-stage processor will not be explicitly used in this lab, but it is provided to show how pipelining in a very simple processor is implemented. Likewise, the micro-coded processor is also not explicitly used in this lab.

First, we must set the pipeline to “Full-Bypassing”. Navigate to the `Chisel` source code:

```
inst$ cd ${LAB1ROOT}/src/rv32_5stage
inst$ vim consts.scala
```

The file `consts.scala` provides all constants and machine parameters for the processor. Change the parameter on line 19 to `val USE_FULL_BYPASSING=true;`. You can see how this parameter changes the pipeline by looking at the data path in `dpath.scala` (lines 178-204) and the control path in `cpath.scala` (lines 178-195). The data path holds the bypass muxes used when full bypassing is activated. The control path holds the `stall` logic, which must account for more situations when no bypassing is supported.

After turning “full bypassing” on, compile and run the processor as follows:

```
inst$ cd ${LAB1ROOT}/
inst$ export MK_TARGET_PROC=rv32_5stage
inst$ make run-emulator
inst$ cd ${LAB1ROOT}/emulator/rv32_5stage/
inst$ ls
inst$ vim vvadd.riscv.out
```

Record the CPI value for all benchmarks. Is it what you expected?

Now turn “full bypassing” off in `consts.scala`, and re-run the results (make sure it recompiled your `Chisel` code).

Record the new CPI values for all benchmarks. How does full bypassing versus full interlocking perform? If adding full bypassing hurt the cycle time of the processor by 25%, would it be worth it? Argue your case. Be quantitative.

4.7 Design problem

Imagine that you are being asked by your employer to evaluate a potential modification to the design of a 5-stage RISC-V. The proposed modification is that the Execute/Address Calculation stage and the Memory Access stage be merged into a single pipeline stage. In this combined stage, the ALU and Memory will operate in parallel. Data access instructions will use memory while leaving the ALU idle, and arithmetic instructions will use the ALU while leaving memory idle. These changes are beneficial in terms of area and power efficiency. Think to yourself why this is the case, and if you are still unsure, ask about it in Section or OH.

In RISC-V, the effective address of a load or store is calculated by summing the contents of one register (`rs1`) with an immediate value (`imm`).

The problem with the new design is that there is now no way to perform any address calculation in the middle of a load or store instruction, since loads and stores do not get to access the ALU. Proponents of the new design advocate changing the ISA to allow only one addressing mode: register direct addressing. Only one source register is used, and the value it contains is the memory address to be accessed. No offset can be specified.

In RISC-V, the only way to perform register direct addressing register-immediate address calculation with `imm = 0`.

With the proposed design, any load or store instruction which uses register-immediate addressing with $imm \neq 0$ will take two instructions. First, the register and immediate values must be summed with an add instruction, and then this calculated address can be loaded from or stored to in the next instruction. Load and store instructions which currently use an offset of zero will not require extra instructions on the new design.

Your job is to determine the percentage increase in the total number of instructions that would have to be executed under the new design. This will require a more detailed analysis of the different types of loads and stores executed by our benchmark codes.

In order to track more specific statistics about the instructions being executed, you will need to modify the “Tracer” class found in the source files `tracer.cpp` and `tracer.h` (located in the `/${LAB1ROOT}/emulator/common/` directory).

Modify “Tracer” to detect the percentage of instructions that are loads and stores with non-zero offsets.

Follow the steps laid out in the `tracer.h` file to accomplish this task. There is existing code provided in “Tracer” which you can follow to implement your modifications.

Use the provided RISC-V ISA specification (found in “extra course materials” on the CS 152 webpage) to determine which bits of the instruction correspond to which fields.

After modifying `tracer.h` and `tracer.cpp`, you can re-compile and re-run your data with “make run” in the `/${LAB1ROOT}/emulator/rv32_5stage/` directory.⁹

What percentages of the instruction mix do the various types of load and store instructions make up? Evaluate the new design in terms of the percentage increase in the number of instructions that will have to be executed. Which design would you advise your employer to adopt? (Justify your position.)

5 Open-ended Portion (5/7 of lab grade)

Pick *one* of the following questions. The open-ended portion is worth a large fraction of the grade of the lab, and the grade depends on how complex and interesting a project you complete, so spend the appropriate amount of time and energy on it. Also, have fun with it!

5.1 Mix Manufacturing

The goal of this section is to investigate how effectively (or ineffectively) the compiler will handle complicated C code created by you.

Using no more than 15 lines of C code, attempt to produce RISC-V assembly code with the maximum ratio of branch to non-branch instructions when run on the 5-stage processor (fully bypassed). In other words, try to produce as many branch instructions as possible. Your C code can contain as many poor coding practices as you like, but limit yourself to one statement per line and do not cheat by calling functions or executing any code not contained within the 15 line block. Your code must terminate. You can use code that creates jumps, but jump instructions do not count; only conditional branches count.

⁹You do not need to run “make run-emulator” from the `/${LAB1ROOT}/` directory if you do not need to recompile the Chisel code for your target processor. Instead you can stay in the emulator `/${LAB1ROOT}/emulator/rv32_5stage/` directory, and only rebuild the C++ simulator.

Finally, run your code on the 5-stage processor (fully bypassed). What is the resulting CPI? As you added more branches, did the CPI get higher or lower? Explain why the CPI went in the direction it did.

Write your code into the file `${LAB1ROOT}/test/riscv-bmarks/mix_manufacturing/mix_manufacturing_main.c`. Modify it to add your custom code. In your write-up, summarize some of the ideas you tried.¹⁰

To test for correctness you can just compile and run it as follows:

```
inst$ cd ${LAB1ROOT}/test/riscv-bmarks/  
inst$ make; make run
```

This invokes the RISC-V ISA simulator, which quickly tests the correctness of the code. You may also view the disassembly by opening the file `mix_manufacturing.rv32u.dump`.

However, to get a cycle-accurate trace of the code to determine the effect your program has on CPI, you will have to run the code on the RV32 5-stage processor:

```
inst$ cd ${LAB1ROOT}/test/riscv-bmarks/  
inst$ make  
inst$ make install  
inst$ cd ${LAB1ROOT}/  
inst$ export MK_TARGET_PROC=rv32_5stage  
inst$ make run-emulator
```

Report the ratio of branches to non-branches that you achieved in your code. You will submit this mix report, the achieved CPI of the 5-stage processor, your lines of C code, and the disassembly of your C code.

5.2 Bypass-path Analysis

As an engineer working for a new start-up processor design company, you find yourself 3% over budget area-wise on your company's latest 5-stage processor (your company makes very small processors, and every bit of area counts!). However, if you remove one bypass path you can make the budget and ship on time!

Using the `Chisel` source code found in `${LAB1ROOT}/src/rv32_5stage/`, analyze the impact on CPI when different bypass paths are removed from the design. The files `dpath.scala` and `cpath.scala` hold the relevant code for modifying the bypass paths and stall logic. Make sure that your modified pipeline passes all of the assembly tests!

Show your data to support the bypass path you feel can be removed with the least impact on CPI. Also, include in an appendix snippets of your modified `Chisel` code.

Feel free to email your TA or attend his office hours if you need help understanding `Chisel`, the processor, or anything else regarding this problem.

¹⁰Note: the provided processors only support load word and store word; therefore, you should avoid `char` variables which synthesize load byte unsigned and store byte instructions.

5.3 Define and Implement Your Favorite Complex Instruction

In the problem set, we have asked you to implement two complex instructions (ADDm and STRCPY) in the micro-coded processor. Imagine you are adding a new instruction to the RISC-V ISA. Propose a new complex instruction (other than MOVN/MOVZ) that needs an EZ/NZ uBr in its implementation, and has at least one memory operand. Write an assembly test for the proposed instruction. Implement the instruction in the micro-coded processor, and test your new instruction.

To define an instruction, you first need to come up with an encoding. Consult the RISC-V ISA document to first come up with an instruction format (see section 2.3 of the RISC-V technical report), and then spot an opcode space that is empty (look at Table 2 of the RISC-V technical report). You will have to add your instruction definition to `src/common/instructions.scala` (please search for FIXME in the file). Let's take a look at the definition for MOVN:

```
val MOVN = Bits("b?????_?????_?????_0000001101_1110111",32)
```

Note that the ? character is used for bit locations that can change (e.g., register specifiers). Bit locations that are fixed have the actual bits written down. _ characters are ignored. The name of the variable is used as a label for the dispatcher in the microcode.

Once you come up with an instruction encoding, you'll have to write an assembly test to test your instruction. To help you out, we wrote an assembly test for the MOVN instruction. Please take a look at `test/riscv-tests32/riscv-vu_movn.S`. Since the assembler doesn't know about the instruction, we manually write down the instruction with a `.word` assembly construct. We also write some assembly code to load values to registers and memory. Finally, the code checks whether it computes the right value. We have provided you with an empty assembly file that you can use at `test/riscv-tests32/riscv-vu_yourinst.S` (please search for FIXME in the file). Compile your assembly test:

```
inst$ cd ${LAB1ROOT}/test/riscv-tests32/  
inst$ make
```

Next, work out the implementation in a worksheet you have used in the problem set (worksheet 2.A or 2.B). Once you have figured out all the states and control signals, start adding your microcode to `src/rv32_ucose/microcode.scala` (please search for FIXME in the file). You will be able to see the microcode instructions for all RISC-V instructions implemented in the micro-coded machine. Again, to provide you an example, we have implemented the MOVN instruction in `microcode.scala`. Once you are done, build the processor and run the assembly test:

```
inst$ cd ${LAB1ROOT}/  
inst$ export MK_TARGET_PROC=rv32_ucose  
inst$ make run-emulator
```

Look at the log output at `emulator/rv32_ucose/riscv-vu_yourinst.out` to observe the machine state. See if the micro-coded processor has executed your instruction correctly. Change your implementation if necessary.

Feel free to email your TA or attend his office hours if you need help understanding Chisel, the processor, or anything else regarding this problem.

5.4 Processor Design

Propose a processor modification of your own to a 5-stage pipeline. Justify your design modification's overhead, cost, or motivation by explaining which instructions are affected by the changes you propose and in what way. You may have to draw a diagram explaining your proposed changes for clarity, and you will very likely have to modify the “Trace” object to track specific types of instructions not previously traced. A further tactic might be to show that while some instructions are impacted negatively, these instructions are not a significant portion of certain benchmarks. Feel free to be creative. Try to quantitatively make your case, but you do *not* need to implement your proposed processor design.

5.5 ISA Design

Imagine you are modifying the RISC-V ISA to allow for two sizes of instructions (let's pretend that it can't already, and that you are creating a competing proposal). With this modification, you will be able reduce the total code size of programs compiled for your machine by allowing some instructions to be encoded in a more compact form. Pick an instruction format (they are listed in the RISC-V handout) and make a proposal for why it should be the one which is allowed to be reduced in size. What is your proposal for reducing the size of this format, and how long is this smaller encoding? By how much does including the variable length instruction reduce code size for the benchmarks we looked at earlier? You can use the full RISC-V manual (<http://www.eecs.berkeley.edu/Pubs/TechRpts/2011/EECS-2011-62.html>) if you want to find additional instruction types which use your proposed format. Be quantitative, and use the provided testing infrastructure to do experiments and make informed decisions. For example, if you propose shorter immediates, run traces on the provided benchmarks to show *exactly* how many instructions you can shorten (i.e., how many instructions can utilize the smaller immediate). Consult with your TA if you need further information on probing the state of the emulator.

5.6 Your Own Idea

We are also open to your own ideas. Particularly enterprising individuals can even modify the provided `Chisel` processors as part of a study of your own design. However, you must first consult with the Professor and/or TA to ensure your idea is of sufficient merit and of manageable complexity.

6 The Third Portion: Feedback

This is a relatively brand new lab, and as such, we would like your feedback! Please fill out the survey form at <http://tinyurl.com/cs152-sp13-lab1-survey>. The link will also be posted on piazza.

How many hours did the directed portion take you? How many hours did you spend on the open-ended portion? Was this lab boring? Did you learn anything? Is there anything you would change? Feel free to write as little or as much as you want (a point will be taken off only if left completely empty).

7 Acknowledgments

Many people have contributed to versions of this lab over the years. This lab was originally developed for CS152 at UC Berkeley by Christopher Celio, and heavily inspired by the previous set of CS 152 labs (which targeted the Simics emulators) written by Henry Cook. This lab was made possible through the work of Jonathan Bachrach, who lead the development of **Chisel**, and through the work of Andrew Waterman, Yunsup Lee, David Patterson, and Krste Asanović who developed the RISC-V ISA.

8 Appendix: Processor Diagrams




Figure 3: The Bus-Based RISC-V Implementation.




Figure 4: The RV32 1-Stage Processor.

Figure 5: The RV32 2-Stage Processor.

16




Figure 6: The RV32 5-Stage Processor.

