

6.823
Computer System Architecture
Lab 3

Assigned Apr. 8, 2014

Worth 10% of Course Grade

Due Apr. 30, 2014

<http://csg.csail.mit.edu/6.823/>

Warning: This lab is competitive and open-ended. Do not wait until the last minute to attempt to complete the lab.

Summary

Branches create instruction hazards in a pipelined processor. When an instruction is loaded the processor must decide which instruction to fetch next. Most of the time, the subsequent instruction in x86 is located in the bytes immediately following the current instruction address. Thus just fetching the next bytes for execution is a reasonable strategy. For branch instructions the next instruction may lie elsewhere in memory. In this case loading the next instruction will likely be incorrect, and the pipeline will stall as the processor fetches the correct next instruction.

Incorrectly fetching the next instruction following a branch instruction can be extremely detrimental to performance, since branch instructions account for 1/6 the total number of instructions. The mechanism that architects use to deal with instruction hazards is branch prediction. At the high level, branch prediction is somewhat simple to describe a branch predictor interface: the branch predictor accepts an address from the processor and returns a branch direction. The actual address is predicted by another structure such as a BTB. Notice that a trivial branch predictor implementation always returns a 'not taken' result, i.e., step to the next sequential instruction. In this lab, you will research predictors and implement a branch predictor model.

Setting up

First, set up your environment for Pin. You'll need to do this each time you log in to work on the labs.

```
% add 6.823
% source /mit/6.823/Spring14/setup.csh (or setup.sh)
```

To obtain the materials for lab 3, use the following commands, assuming that you start in your individual repository (cd \$USER) from the previous lab:

```
% svn export $LAB3ROOT
% svn add lab3handout
% svn commit -m "Lab 3 Initial Check-in"
```

In the lab3handout directory that just got created, you should find a make file, some sample source code, and a test script. Type the following at the command prompt:

```
% cd lab3handout
% make
```

Make will build the bpredictor Pin tool. Attempting to run the Pin tool will result in a failure message, stating that no Pin tool has been implemented. The caches Pin tool can be invoked from the command line in the following manner:

```
% pin -t bpredictor - [target_executable]
```

We have provided a test perl script lab3test.pl. The perl script will invoke Pin using the bpredictor Pintool on multiple SPEC binaries. To invoke the perl script, type:

```
% ./lab3test.pl
```

Although we provide a script that will test your Pin tool, the script will not verify your Pin tool. Since the lab is open-ended in the first place, your results may differ substantially from your classmates. Further, we reserve the right to run test cases not included in the released test script. You are encouraged to compare your results with your classmates.

Lab Task

The purpose of this lab is to explore branch predictor architectures. To implement you will extend the BranchPredictor class to implement your own branch predictor. Logically, this branch predictor will hang off of the processor core and accept branch address queries from the processor. Given a branch address, the branch predictor will respond with a bit denoting whether the branch is taken or not taken. The Branch Predictor class consists of the following functions:

makePrediction: This function will accept the address of a branch instruction. The function will return true if the branch is predicted to be taken and false otherwise.

updatePrediction: This function is a feedback from the processor core. It will take parameters the branch address, your original prediction, the actual direction of the branch. You will use this information to update your internal branch predictor state as you see fit.

In the pin tool that we have implemented, when we encounter a branch instruction, we will call makePrediction to get your prediction. We will then determine the correct prediction and call your updatePrediction function.

Notice that your branch predictor will only need to predict whether the branch is taken or not. You will not have to present a potential branch target. Your branch predictor is allowed to use no more than 33 Kilobits of persistent storage. Using more than 33 Kilobits will disqualify your branch predictor from the competition portion of the assignment. Needless to say, this would be extremely detrimental to your grade. You

may these bits in any manner you see fit, and you may use any amount of computation that your branch predictor requires. At the end of the execution, we will dump statistics about the performance of your branch predictor, which we will track in our code.

As a side note, the sub-field of computer architecture related to branch prediction is rather rich and contains many important ideas. Although implementing the branch predictors proposed in class will likely obtain half of the correctness credit, if you hope to beat the TA's reference branch predictor, you will probably have to do some research on your own.

Although your solution will not be graded on its performance in terms of wall clock time, you should note that your Teaching Assistants are impatient people. The TA solution runs the sample testbench in less than one hour on the class machines (with nominal load). For grading purposes, we will allow your pin tool to run for an order of magnitude more time than ours requires (around ten hours). After ten hours, we will kill your submission and assign a grade based on progress to that point. Do not write horrendously inefficient code: it makes kittens sad.

When you have completed the lab to your satisfaction, submit your changes to the svn repository. The deadline for submission is 23:59:59 EDT 30 April 2014. We'll grade whatever code you have checked in by the deadline. **No Late Submissions will be accepted!** Seriously.

Lab Questions

Your response to the lab questions should be typed in lab3questions.pdf (or lab3questions.doc) in the lab3handout directory. The course material necessary to answer some of the questions will be covered after the distribution of this handout, so if some of the material looks foreign, it probably just hasn't been covered yet. Some questions may require coding, and as such should not be put off until the last minute.

1. **It's absolutely pitiful to code like that.** Your TAs have written the lab handout in a slightly suboptimal way, in terms of simulation speed. Figure out the inefficiency and fix it. Explain your change. Does it make a big difference in speed?
2. **Hardware? Don't talk about hardware.** Explain the operation of your implementation. Detail the general algorithm that you used and explain why you chose it. Why does the algorithm that you finally chose outperform the other algorithms you tested and those discussed in class? Describe how you allocated the storage bits of your branch predictor. Don't forget to cite your sources.
3. **Are you kidding me? Hardware?** We gave you unlimited computation to develop your branch predictor. Could your implementation be built into a processor? Why or why not. Discuss which elements of your algorithm would cause problems if implemented in silicon.
4. **I'm just hoping we can implement a software model.** The branch predictor you implemented is likely to be missing several of the elements of a real branch

predictor. For example, in a real machine several predictions might be made before an update occurs. Why might that happen? Explain the consequences of this on your predictor in terms of impact on prediction and hardware changes that might be added to address that impact. Don't neglect the impact on misprediction recovery.

When you have answered these questions to your satisfaction, put them in a file called lab3questions.pdf (or lab3questions.doc) in your lab3handout directory, then run the following to add them and commit them.

```
% svn add lab3questions.pdf
% svn commit -m "Lab 3 Questions Check-In"
```

As with the lab code, we'll grade whatever you have checked in by the deadline.

Lab Grading

10%: Submission compiles

60% - Competitive grade based on the following break-down based on aggregate performance on the public and private benchmarks

20% - below 80% accuracy

30% - Any implementation achieving greater than 80% accuracy

40% - Any implementation achieving 90% accuracy on the benchmarks

50% - Any implementation achieving 95% accuracy on the benchmarks

60% - Any implementation achieving 98% accuracy on the benchmarks

30% - Quality of lab question responses.

Advice on Mine Sweeping

There may be bugs in either our code or infrastructure. If you notice any 'interesting' or 'unexpected' behavior it could be a problem in the code or infrastructure that we provided. Report these bugs immediately to the TA, preferably in an email with the subject 6.823 Bug Report. This will help to ensure prompt fixing of any issues that may arise.

Guides for the perplexed

<http://www.pintool.org/> - Pin home page

<http://tig.csail.mit.edu/twiki/bin/view/TIG/UsingSubversionAtCSAIL> – an SVN tutorial