

CSE471: Computer Design & Organization

Assignment 2

Due: Tuesday, April 18

The purpose of this assignment is to supplement your knowledge of the design of different dynamic branch prediction schemes with intuition about their relative performance.

You are expected to work in teams of 2 people, hopefully with a different partner than the last assignment.

1. For this assignment your applications should be `cc1` and `compress` with the same input as in homework 1.
2. Simulate the same microarchitecture as in the first assignment, but vary the branch prediction strategy between:
 - 2-bit dynamic branch prediction
 - correlating branch prediction, using the default number of history bits and pattern history table size

This means you will have separate simulations for the two branch prediction techniques.

SimpleScalar determines the branch target address for both branch prediction schemes with a branch target buffer.

3. Generate and analyze your results. In particular, address the following issues:
 - Branch frequency.
 - How often are branches executed in your programs?
 - If there is a difference between the two programs, do you have an hypothesis to explain it?
 - How does this compare with the “average” frequency of every 4-6 instructions?
 - If there is a difference from the average, do you have an hypothesis to explain it?
 - Branch characterization.
 - For conditional branches, record the frequencies of each element of the Cartesian product: (*forward*, *backward*) x (*taken*, *not taken*). You can express your answer in a table like this one:

	taken	not taken
forward	#	#
backward	#	#

Do you need SimpleScalar to generate additional stats to do this? If so, write the code to do so.

- Discuss whether or not your results support the use of the traditional static scheme of backward branches taken and forward branches not taken.
- Prediction accuracy.
 - Record the number of correct predictions for all branch prediction schemes, including what the static scheme would have achieved. (You do not have to run a simulation, just use the branch characterization results.)
 - Which prediction scheme has the most correct predictions for each program? Why? If you need to run additional simulations to justify your analysis and conclusion, do so.
 - This question is only applicable if each of your programs performed best with a different branch prediction scheme. If you had to pick only one prediction scheme that would be used for both programs, which would it be and why?
- 4. Is there a difference between the number of instructions fetched and the number of instructions committed? Why or why not?
- 5. Is there a limit to the number of useful history bits, or is more history always better? Perform simulations to determine what the situation really is. What do you think explains your data?

Keep in mind that when doing a sensitivity analysis of this kind, you should keep all factors constant except the one you are varying.

- 6. Write up your experiments, the results and an analysis of the results in a report, as outlined in the report handout and illustrated in the sample reports. In the results section, devote a different subsection to each of the issues listed above in items 3 through 5. Use tables and graphs to illustrate the results.
- 7. In addition to the experimental portion of the assignment, answer the following question. For a (3,2) correlated prediction scheme which has a global history register and one pattern history table, show the contents of the history register and the PHT after the following sequence of branch executions have taken place: 1 not-taken branch, 2 taken branches, 2 not-taken branches, 2 taken branches, and 2 not-taken branches.
 - The prediction bits in the pattern history table use 2-bit prediction, implemented with saturating counters. The counters are incremented when a branch is taken and decremented when a branch is not taken. Each counter is initialized to the weakly taken state, as illustrated on slide 5 in the lecture notes. Assume this state has the value 10.
 - The global history register is initialized to zero.
 - The addresses of the ten branch instructions are: ffff0000, ffff0001, ffff0010, ffff0000, ffff0000, ffff0000, ffff0011, ffff0010, ffff0000.

Extra credit only, meaning you don't have to do this if you prefer not to:

After getting a sense of how prediction correctness varies with changes in the size of the hardware data structures, the question that microarchitects usually face, is:

“If I have a certain amount of chip area that I can devote to branch prediction, how should I divide it among the various branch data structures for correlating branch prediction, the global history registers, the pattern history table, and the branch target buffer?”

For extra credit, you can try to answer that question for a bit budget of 2K bits. To focus your thinking, you could consider issues, such as the following:

- Given 2K bits, how many of those bits should you devote to the history registers and how many should be used for the pattern history table?
- Should there be one global history register that is used by all branches, or should there be several, each of which is used by fewer branches?
- If several, how many?
- How many entries should there be in the PHT?

Be sure that all your hardware structures have powers of two entries, so that they can be accessed easily by the PC or branch history bits. This will probably mean that you won't use up your entire bit budget.

Why do you think your design of choice got the performance it did, relative to others you tried?