

EECS 370 Winter 2009

Homework 5

Assigned: Tuesday, March 17, 2009
Due: Tuesday, March 31, 2009 (in class)

Name: _____ Uniqname: _____

Instructions:

1. Please write your name and uniqname in the spaces provided above. *Attach this cover sheet* to your completed solutions to the problems listed and turn them in at class on the due date. Submissions without a completed, attached cover sheet cannot be graded.
2. Your answers should be neat, clear, and concise. Computer-written work is recommended (but not required). Show all your work, and state any special or non-obvious assumptions you make.
3. You may discuss your solution methods or your answers with other students, but the solutions you submit must be your own.

Scores:

Problem #	Points
1	/4
2	/8
3	/10
4	/10
5	/8
Total	/40

Problem 1 (4 points)

Suppose we are given a “magical” register file that you can read from and write to in negligible time. For everything else, assume your performance conforms to Figure 6.2 on page 373 of your textbook.

a) Compared to using the old register file (100 ps reads and writes) in Figure 6.2, will the magical register file affect the performance of a pipelined architecture? If yes, by how much? If not, why?

b) What if, instead of a magical register file, you use a new register file that now takes 25% more time?

Problem 2 (8 points)

You have analyzed a program and discovered that 60% of branches are taken and 40% are not taken. You have also discovered that 35% of all branches branch forward, while the remaining 65% branch backward (this statistic holds for both taken and not taken branches). Evaluate the accuracy (i.e. % correct predictions on average) of the following branch prediction schemes when running your program. Show your work.

a) Random Prediction (i.e. 50/50)

b) All branches predicted taken

c) All branches prediction not taken

d) Branches predicted taken if offset negative, not taken otherwise.

Problem 3 (10 points)

Consider the standard 5 stage LC-2K9 pipeline as presented in lectures. Assume all data dependencies can be handled by forwarding. Assume a speculated “never taken” policy and that branches are resolved after the EX cycle at stage 4.

a) What is the average CPI (clocks per instruction) for a program in which 20% of all instructions are branches (this is roughly true for most programs) and branches are predicted correctly 70% of the time?

b) You decide to implement a more accurate branch predictor in the standard LC2K7 pipeline. Simulations suggest that this branch predictor will be accurate 80% of the time. What would be the average CPI of your new pipeline?

c) Now suppose that you were motivated by the above results and decided to evaluate an even more accurate branch predictor. The design under evaluation has an average branch prediction accuracy of 95.5%. However, this branch predictor is so large and complex that we must increase our fetch stage (IF) leading to a clock cycle increase of 15%. Is implementing this branch predictor in our design a good idea (compared to the design of part (b))? Why?

Problem 4 (10 points)

Consider a normal 5-stage LC-2K8 pipeline as discussed in class. Branches are resolved in the MEM stage and branches are predicted as **not taken**. 30% of load instructions incur a one cycle stall due to data hazards that cannot be resolved via forwarding. Assume the instruction memory has a 2% miss rate and the data memory has a 10% miss rate for both reads and writes. For an instruction or data memory read, a miss causes a 15 cycle stall. For a data memory write, a miss causes a 10 cycle stall. Analysis has shown the following dynamic instruction breakdown for program X: 10% not taken branches, 7% taken branches, 20% loads, 23% stores, 40% R-type instructions (add/nand).

(a) What is the CPI of this machine when running program X? (4 points)

(b) Now assume the pipeline was stretched out to 10 stages (numbered 1 to 10, with instruction fetch beginning in stage 1, and instructions finally completing in stage 10). Branches are resolved in stage 7. The clock frequency is increased by 50% than that of the original processor with the 5 stage pipeline. For an instruction or data memory read, a miss now causes a 35 cycle stall. 30% of load instructions still incur a one cycle stall due to data hazards that cannot be resolved via forwarding. What is the CPI for this new pipeline design? (3 points)

(c) Which of these machines performs better, the original 5 stage pipeline or the new 10 stage pipeline? How many times faster is it than the other one? (3 points)

Problem 5 (8 points)

Consider a typical 5-stage pipelined processor (Fetch, Decode, Execute, Memory, and Write back stages) where:

- Branches are predicted not-taken and resolved in Memory stage.

- Forwarding paths are available. When forwarding cannot resolve the problem, the instruction stalls in Execute.
- All instructions take 1 cycle to execute other than the “multiply” instruction which takes 2 cycles (and thus causes a stall in the pipeline).
- Address calculations are done in the Execute stage.

For a given program, 20% of instructions are loads, 25% are adds, 10% are multiply, 20% are branches, 25% are stores; and 60% of branches are not taken.

Also, for a load instruction, *each operand* of the following instruction has a 40% chance of being dependent on the result of the load instruction. You may assume the probabilities are independent of each other.

What would be the CPI for this program? Show your work for full credit.

Hint: You may have to use the probability formula

$$\begin{aligned}
 P(A \cup B) &= P(A) + P(B) - P(A \cap B) \\
 &= P(A) + P(B) \quad \text{if A and B are mutually exclusive}
 \end{aligned}$$