

Computer System Architecture

6.823 Quiz #2

April 9th, 2021

Name: _____

90 Minutes
16 Pages

Notes:

- Not all questions are equally hard. Look over the whole quiz and budget your time carefully.
- Please state any assumptions you make, and show your work.
- Please write your answers by hand, on paper or a tablet.
- Please email all 16 pages of questions with your answers, including this cover page. Alternatively, you may email scans (or photographs) of separate sheets of paper. Emails should be sent to 6823-staff@csail.mit.edu
- Please ensure your name is written on every page you turn in.
- Do not discuss a quiz's contents with students who have not yet taken the quiz.
- Please sign the following statement before starting the quiz. If you are emailing separate sheets of paper, copy the statement onto the first page and sign it.

I certify that I will start and finish the quiz on time, and that
I will not give or receive unauthorized help on this quiz.

Sign here: _____

Part A	_____	30 Points
Part B	_____	40 Points
Part C	_____	30 Points

TOTAL _____ **100 Points**

Part A: Branch Prediction (30 points)

Ben Bitdiddle is trying to design a branch predictor for his processor. The C code is shown below:

```
for (int i = 0; i < 1000000; i++) {
    if (i % 2 == 0) { // branch B1
        // do something A
        ...
    }
    if (i % 5 == 0) { // branch B2
        // do something B
        ...
    }
}
```

The corresponding assembly code is shown below:

```

                                ADDI R1, R0, 0
                                LOOP: MODI R2, R1, 2
0xc44:                          BNE  R2, R0, M2          // branch B1
                                (do something A)
                                ...
                                M2:  MODI R2, R1, 5
0xc84:                          BNE  R2, R0, END        // branch B2
                                (do something B)
                                ...
                                END: ADDI R1, R1, 1
0xcc0:                          BNE  R1, 1000000, LOOP    // branch LP
```

The MODI (modulo-immediate) instruction is defined as follows:

```
MODI rd, rs, imm: rd <- rs mod imm
```

Note that whenever the condition of the if-statement is met, we do **not** take the branch BNE.

Question 1 (3 points)

In steady state, how often are each of the three branches taken in the above code?

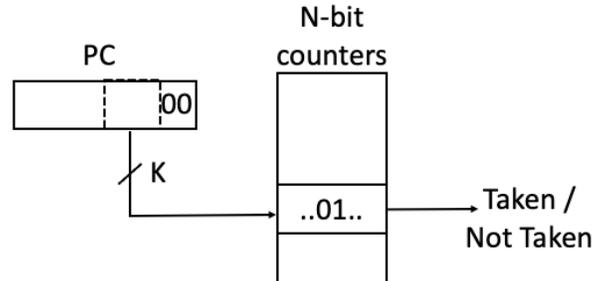
Question 2 (3 points)

Ben's first attempt is to design a static branch predictor. The static predictor predicts not taken for all forward branches, and taken for all backwards branches.

For each branch in Ben's code, what is the accuracy of this static predictor?

Question 3 (9 points)

Ben tries to improve his prediction accuracy by designing a bimodal predictor.



The bimodal predictor consists of a table of N -bit counters. The predictor uses the K least significant bits of the PC in a word (not byte) address to index into the N -bit counters. The most significant bit of the counter is used to make a prediction for the branch (1 for Taken, 0 for Not Taken). If a branch is found to be taken, the counter is incremented by one (up to a maximum value of $2^N - 1$), and decremented by one (down to a minimum value of 0) otherwise. **Assume that all bits of the N -bit counters are initialized to zeros in the beginning.**

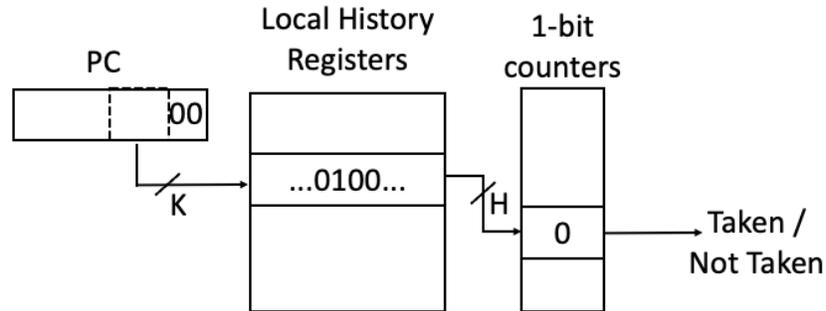
a) How many bits are needed so that the three branches (B1, B2, and LP) are mapped to distinct entries in the table of N -bit counters?

b) What is the minimum value of N required to achieve as good of an accuracy in steady state **for all branches** as the static predictor from Question 2?

c) Given your value of N , what is the accuracy of Ben's bimodal predictor for the three branches in steady state?

Question 4 (5 points)

Ben now considers a 2-level local history predictor design.

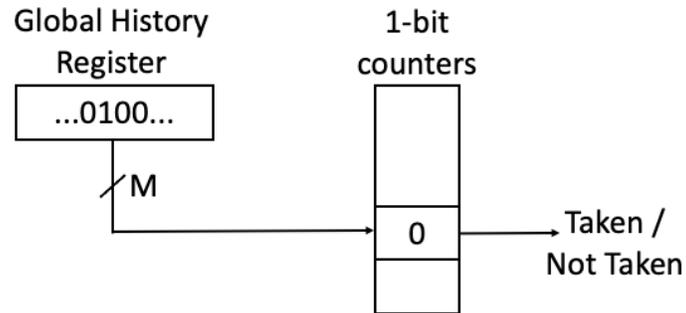


The **K** least significant bits of the PC in a word are now used to index into a table of local history registers. Each local history register stores **H** bits, and is used to keep track of the history of the PC. The H bits are used in turn to index into a table of 1-bit counters indicating the prediction (Taken if 1, Not taken if 0).

What is the minimum value of H such that all branches are predicted perfectly in steady state? Assume that the value of K is large enough that different branches are mapped to distinct local history registers.

Question 5 (5 points)

Ben now designs a global history predictor as follows.



The predictor consists of a single global history register of M bits that stores the outcomes of all branches in the program. Every time the processor encounters a branch, it shifts in a 1 to the history register from the right if the branch is taken, and a 0 if the branch is not taken. The global history register is used to index into a table of 1-bit counters indicating the prediction (Taken if 1, Not taken if 0).

Can this predictor achieve perfect prediction for all branches of Ben's code in steady state? Briefly explain your reasoning.

Question 6 (5 points)

Ben is given the following piece of C code:

```
for (int i = 0; i < N; i++) {  
    for (int j = 0; j < M; j++) { // branch LP  
        ... // Do something  
    }  
}
```

Assume that $M > 2$. Ben wants to choose a branch predictor that maximizes the prediction accuracy for the inner loop (LP) of this code (assume that the outer loop is completely ignored by the branch predictor). Ben is given two choices: a bimodal predictor with 1-bit counters from Question 3, and a local history predictor from Question 4 with M -bit local histories. For what values of N will the local history predictor outperform the bimodal predictor? Assume that all structures for both predictors have their entries initialized to all zeros at the beginning of the outer loop.

Hint: First try to evaluate the two predictors with $N = 1$. What are their prediction accuracies?

- e) Assume M3 data becomes available, I14 is issued, and the branch is found to be predicted incorrectly and should have not been taken. The global history register is shifted to the right by one bit to recover the previous history.

- f) Assume Instruction I17 is an `addi R5, R4, 1` instruction. I17 enters the Register Rename and Allocation stage. It updates the rename table entry of R5 to A5.

- g) Assume A3 data becomes available, and instruction I16 encodes an offset of 0 (not shown in the handout). Instruction I16 is issued and finds no matching loads in the load buffer.

- h) Assume Instruction I19 is a branch instruction. When I19 enters the decode stage, the branch predictor consults entry 7 and predicts the branch to be taken.

- i) Assume M1 data becomes available. I10 is issued, executes, and writes back its result to the `data` field of A2 and sets the corresponding `pd` bit.

- j) Assume all instructions up to I12 commit. I13 commits, and overwrites register R1 in the register file with the value from its `data` field.

Question 2 (5 points)

Answer the following yes/no questions about our out-of-order machine with split reservation stations.

- a) Do writes to the register file happen in program order?
- b) If a store finds a matching younger load in the load buffer, should the load and all following instructions be killed?
- c) Is the rename table always updated after an instruction commits?
- d) Are the 2-bit counters in the branch predictor updated lazily?
- e) Suppose our programs have 50% arithmetic and 50% memory instructions. Should both reservation stations be sized equally for maximum utilization?

Question 3 (5 points)

In our current implementation of reservation stations, at most one instruction is committed per cycle. Is there a way to provide higher commit throughput (commit more than one instruction/cycle) without making each of the reservation stations check more than one entry per cycle for commit? If so, briefly describe such a mechanism. If not, state your reasoning.

Part C: Advanced Memory Operations (30 points)

For this part, please refer to the accompanying store sets Handout for a detailed explanation of how store sets work.

Consider the following C code:

```
int A[100];

for (int i = 0; i < 100; i++) {
    if (i % 2 == 0) {
        int j = long_latency_op1(); // Some long latency operation
        A[j] = x; // Store X
    }
    else {
        int j = long_latency_op2(); // Some long latency operation
        A[j] = y; // Store Y
    }
    int z = A[i]; // Load Z
    large_work(); // Large amount of work
}
```

Note that Store X, Store Y, and Load Z are the only instructions that access memory in this piece of code.

Ben runs the C code on his out-of-order processor which implements store sets to predict memory dependences as described in the store sets handout. Assume the following when evaluating how the above code executes on Ben's processor:

- The value j is equal to i at every iteration of the loop. In addition, assume that the store address $A[j]$ is available much later than the load address $A[i]$.
- Because the stores depend on long-latency operations, the load is always ready to be issued before the store.
- All branches are predicted perfectly.
- The store sets of all loads are invalidated before the start of this code segment.
- The program performs a large amount of work at the end of the loop such that stores and loads from different iterations cannot be in flight at the same time.

Question 1 (3 points)

Out of the 100 iterations, how many times will the processor stall a load due to predicting a correct memory dependence via store sets?

Now Ben runs the following C code on his machine:

```
int A[100];

for (int i = 0; i < 100; i++) {
    int x = foo();
    A[i] = x; // Store X
    ...
    int z = A[99 - i]; // Load Y
    ...
}
```

Store X and Load Y are the only instructions that access memory in this piece of code. Again, assume perfect branch prediction. Also assume that the processor has a large enough reorder buffer such that loads and stores from different iterations can be in flight simultaneously, and stores and loads are issued at roughly the same rate.

Question 2 (4 points)

Will our store set implementation predict a memory dependence between Store X and Load Y at some point? If so, what proportion of those predictions are false dependences (i.e., dependences that did not actually exist)? You don't need to give the exact answer, a rough proportion (within 5-10%) is fine. Justify your answer briefly.

Question 3 (5 points)

Alyssa P. Hacker notes that Ben's store set implementation is always building store sets, but never forgetting useless ones. A simple solution would be to periodically clear all store sets, but Ben thinks this is too costly. Derive a simple mechanism that can remove stores from store sets that keep predicting a false dependence. *Hint: Think about how we design branch predictors.*

Associative searches on store buffers are very expensive in terms of cycle time and power. Recall that when a load issues, it must perform an associative lookup on the store buffer to find any stores that it must forward data from.

With accurate memory dependence prediction, we can eliminate the store buffer completely. If we can detect that a load is dependent on an earlier store during decode, it can simply read the store's data input physical register instead of waiting for the store to write to value to its store buffer entry. This technique is called **Speculative Memory Bypassing (SMB)**.

For example,

```

SW   P3, 0(P4) ; Store the value to memory
...
LW   P6, 0(P5) ; Load the value from memory

```

Here, the stores and loads are denoted with the actual physical registers they read and write during execution (i.e., they are already renamed). Notice that if we predict that $\text{Reg}[P4] == \text{Reg}[P5]$, the LW can directly read the register P3 and simply move the value to register P6. We call these *bypassing loads*. Loads that are predicted to not depend on a previous store (i.e., *non-bypassing loads*) can simply execute as before and read the data from the cache.

Since the physical register file now replaces the store buffer as a forwarding intermediary, we can eliminate the store buffer completely. Stores are now performed in order, when it is the store's turn to commit, updating the cache directly without an intermediary buffer.

We must also take care of two situations where we mis-predict the memory dependence: (1) A non-bypassing load that should have bypassed from a previous store, and (2) A bypassing load that should have read from the cache or bypassed from a different store. To preserve correct execution, every load is re-executed in program order when it is about to commit, always reading from the cache. This is called a *verification load*. If the original (speculative) load and the verification load produce different values, the load and all later instructions are discarded.

Ben modifies his processor to perform SMB and verification loads, allowing him to remove the store buffer entirely, and uses store sets to predict memory dependences.

Question 4 (4 points)

With Ben's modifications, do we still need any associative searches into the load buffer? Justify your reasoning.

Question 5 (4 points)

If the majority of the program's loads are predicted as non-bypassing, are there any significant performance penalties associated with performing verification loads?

Question 6 (5 points)

Alyssa notes that the processor does not need to perform verification loads for all non-bypassing loads. If we can keep some information about previously committed stores, we can "filter" certain verification loads. Briefly describe the condition under which a verification load can be avoided.

Question 7 (5 points)

Explain a scenario where the processor with SMB (and no store buffer) would perform *better* than the original design (with load and store buffers, and store sets for memory dependence prediction). Then, explain a scenario where it would perform *worse*. State clearly any assumptions you make, and briefly justify your reasoning for each scenario.