

CS 433: Computer Architecture – Fall 2022

Homework 3

Total Points: Undergraduates (69 points), Graduates (79 points)

Undergraduate students should solve only the first 4 problems.

Graduate students should solve all problems.

Due Date: September 29, 2022 at 10:00 pm CT

(See course information slides for more details)

Directions:

- **All students must write and sign the following statement at the end of their homework submission. "I have read the honor code for this class in the course information handout and have done this homework in conformance with that code. I understand fully the penalty for violating the honor code policies for this class." No credit will be given for a submission that does not contain this signed statement.**
- **On top of the first page of your homework solution, please write your name and NETID, your partner's name and NETID, and whether you are an undergrad or grad student.**
- **Please show all the work that you used to arrive at your answer. Answers without justification will not receive credit. Errors in numerical calculations will not be penalized. Cascading errors will usually be penalized only once.**
- **See course information slides for more details.**

Problem 1 [39 points]

This problem concerns Tomasulo's algorithm. Consider the following architecture specification:

Functional Unit Type	Cycles in Ex	Number of Functional Units
Integer	1	1
FP Adder	5	1
FP Multiplier	8	1
FP Divider	15	1

- (1) Assume that you have unlimited reservation stations.
- (2) Memory accesses use the integer functional unit to perform effective address calculation during the EX stage. For stores, memory is accessed during the EX stage (Tomasulo's algorithm without speculation) or commit stage (Tomasulo's algorithm with speculation). All loads access memory during the EX stage. Loads and Stores stay in EX for 1 cycle.
- (3) Functional units are *not* pipelined.
- (4) If an instruction moves to its WB stage in cycle x , then an instruction that is waiting on the same functional unit (due to a structural hazard) can start executing in cycle x .
- (5) An instruction waiting for data on the CDB can move to its EX stage in the cycle after the CDB broadcast.
- (6) Only one instruction can write to the CDB in one clock cycle. Branches and stores do not need the CDB.
- (7) Whenever there is a conflict for a functional unit or the CDB, assume that the oldest (by program order) of the conflicting instructions gets access, while others are stalled.
- (8) Assume that the result from the integer functional unit is also broadcast on the CDB and forwarded to dependent instructions through the CDB (just like any floating point instruction).
- (9) Branches do not have a delay slot.
- (10) Assume that the BNEQZ occupies the integer functional unit for its computation and spends one cycle in EX.

Part A [9 points]

For this part, assume a single-issue machine. Fill in the table below with the cycle number where each instruction occupies the given stage. If a stall occurs, describe the reason for the stall. The reason should include the type of hazard; the register, functional unit, etc. that caused the dependence; and the instruction on which the given instruction is dependent (use the IS stage cycle number to identify the instruction). The first entry has been filled for you.

Instruction	IS	EX	WB	Reasons for stall
L.D F0, 0(R1)	1	2	3	
ADD.D F2, F0, F4	2			
MUL.D F4, F2, F6				
ADD.D F6, F8, F10				
DADDI R1, R1, #8				
L.D F1, 0(R2)				
MUL.D F1, F1, F8				
ADD.D F6, F3, F5				
DADDI R2, R2, #8				

Part B [2 points]

Would pipelining any of the functional units reduce the total execution time of the above code segment? Explain your answer.

Part C [2 points]

Would adding another multiplier functional unit be more advantageous than pipelining the multiplier for the above code segment? Assume that pipelining the multiplier results in 8 stages with 1 cycle per stage. Explain your answer.

Part D [18 points]

For this part, assume *hardware speculation* and *dual-issue* added to the Tomasulo pipeline you used for part A. That is, assume that an instruction can issue even before the branch has completed (or started) its execution (as with perfect branch and target prediction). However, assume that an instruction after a branch cannot issue in the same cycle as the branch; the earliest it can issue is in the cycle immediately after the branch (to give time to access the branch history table and/or buffer). Any other pair of instructions can be issued in the same cycle. Assume that a store calculates its target address in EX and performs its memory access during the Commit stage. Recall that stores and branches (BNEZ) do not write back.

Additionally, assume that your reorder buffer has *12 entries* (at the beginning of execution the ROB is empty). Furthermore, *two instructions can commit each cycle*.

Fill in the cycle numbers in each pipeline stage for each instruction in the first two iterations of the loop represented below, assuming the branch is always taken. If a stall occurs, describe the reason for the stall. The reason should include the type of hazard; the register, functional unit, etc. that caused the dependence; and the instruction on which the given instruction is dependent (use the IS stage cycle number to identify the instruction). Additionally, note any stalls due to commit ordering. The entries for the first two instructions of the first iteration are filled in for you. CM stands for the commit stage.

Instruction	IS	EX	WB	CM	Reason for Stalls
Iteration 1					
LP: L.D F0, 0(R1)	1	2	3	4	
ADD.D F0, F0, F6	1	4-8	9	10	RAW on F0 (from 1.1)
DIV.D F2, F2, F0					
L.D F0, 8(R1)					
DIV.D F4, F0, F8					
S.D F4, 16(R1)					
DADDI R1, R1, #-24					
BNEZ R1, LP					
Iteration 2					
LP: L.D F0, 0(R1)					
ADD.D F0, F0, F6					
DIV.D F2, F2, F0					
L.D F0, 8(R1)					
DIV.D F4, F0, F8					
S.D F4, 16(R1)					
DADDI R1, R1, #-24					
BNEZ R1, LP					

Part E [6 Points]

For the code in Part D, which of the following optimizations will improve the total execution time, in terms of number of cycles: triple issue, three instructions committed per cycle, or increasing the reorder buffer size to 14 entries? Explain why (consider each of these as independent optimizations).

Part F [2 Points]

For the code in Part D, assume a floating point divide by 0 incurs an exception. In which clock cycle will the system invoke a jump to the interrupt service routine if F8 used in the fifth instruction has the value 0? Assume that the exception is identified as soon as EX begins and the instruction with the exception gives up the execution unit in the same cycle (i.e., it is available for another instruction in the next cycle). Explain your answer.

Problem 2: Dynamic Branch Prediction [14 points]

Consider the following MIPS code. The register R0 is always 0.

```
DADDI    R1, R0, R0
```

```
L1: DADDI    R2, R0, R0
```

```
L2: DADDI    R2, R2, #1
```

```
    DSUBI    R3, R2, #3
```

```
    BNEQZ   R3, L2          -- Branch 1
```

```
    DADDI    R1, R1, #1
```

```
    DSUBI    R4, R1, #4
```

```
    BNEQZ   R4, L1          -- Branch 2
```

Each table below refers to only one branch. For instance, branch 1 will be executed 12 times. Those 12 times should be recorded in the table for branch 1. Similarly, branch 2 is executed 4 times.

Part A [4 points]

Assume that 1-bit branch predictors are used. When the processor starts to execute the above code, both predictors contain value N (Not taken). What is the number of correct predictions? Use the following tables to record the prediction and action of each branch. Several entries are filled in for you.

Branch 1:

Step	Branch 1 Prediction	Actual Branch 1 Action
1	N	T
2		T
3		
4		
5		
6		
7		
8		
9		
10		
11		
12		

Branch 2:

Step	Branch 2 Prediction	Actual Branch 2 Action
1	N	T
2		
3		
4		

Part B [4 Points]

Now assume that 2-bit saturation counters are used. When the processor starts to execute the above code, both counters contain value 00. What is the number of correct predictions? Use the following tables to record the prediction and action of each branch. You have to follow the 2-bit saturation counters taught in class for branch prediction.

Branch 1:

Step	Counter Value	Branch 1 Prediction	Actual Branch 1 Action
1	0 0	N	T
2			T
3			
4			
5			
6			
7			
8			
9			
10			
11			
12			

Branch 2:

Step	Counter Value	Branch 2 Prediction	Actual Branch 2 Action
1	0 0	N	T
2			
3			
4			

Part C [6 points]

Now assume that 2 level *global* correlating predictors of the form (2, 1) are used. (Note that global here means that the history used captures the history of all previous branches. It does not mean that there is only one set of prediction bits for all branches.) When the processor starts to execute the above code, the outcome of the previous two branches is not taken (N). Also assume that the initial state of predictors of all branches is not taken (N). What is the number of correct predictions? Use the following table to record your steps. Record the "New State" of predictors in the form W/X/Y/Z where,

W - state corresponds to the case where the last branch and the branch before the last are both TAKEN.

X - state corresponds to the case where the last branch is TAKEN and the branch before the last is NOT TAKEN.

Y - state corresponds to the case where the last branch is NOT TAKEN and the branch before the last is TAKEN.

Z - state corresponds to the case where the last branch and the branch before the last are both NOT TAKEN.

Note: The state of the predictor at position W/X/Y/Z can be *N* for predict-not-taken or *T* for predict-taken.

Branch 1:

Step	Branch 1 Prediction	Actual Branch 1 Action	New State
1	N	T	N/N/N/T
2		T	
3			
4			
5			
6			
7			
8			
9			
10			
11			
12			

Branch 2:

Step	Branch 2 Prediction	Actual Branch 2 Action	New State
1	N	T	N/N/T/N
2			
3			
4			

Problem 3 [8 Points]

Part A [6 points]

Suppose we have a deeply pipelined processor, for which we implement a branch-target buffer for conditional branches only. Assume that the misprediction penalty is always four cycles and buffer miss penalty is always three cycles. Assume a 90% hit rate, 90% accuracy, and 15% branch frequency. How much faster is the processor with the branch-target buffer versus a processor that has a fixed two-cycle branch penalty? Assume that the base CPI without branch stalls is one cycle.

Part B [2 points]

Now consider a branch-target buffer design that distinguishes conditional and unconditional branches, storing the target address for a conditional branch and the target instruction for an unconditional branch. What is the penalty or benefit in clock cycles when an unconditional branch is found in the buffer? Explain.

Problem 4 [8 points]

This problem concerns the implications of the reorder buffer size on performance. Consider a processor implementing Tomasulo's algorithm with reservation stations and the reorder buffer scheme described in detail in the lecture notes. Assume infinite processor resources unless stated otherwise; e.g., infinite execution units and infinite reservation stations. Assume a perfect branch predictor and assume there are no data dependences in the instruction stream we are considering. Assume the maximum instruction fetch rate is 12 instructions per cycle. (The other stages in the pipeline have no constraints; e.g., the processor can decode an unbounded number of instructions per cycle.)

Part A [2 points]

Suppose all instructions take one cycle to execute and the processor has an infinite reorder buffer. What is the average instructions-per-cycle rate or IPC for this processor?

Part B [2 points]

Consider the system in part (a) except that now every 48th instruction is a load that misses in the cache such that completing the load takes 500 cycles. What is the average instructions-per-cycle or IPC for this processor?

Part C [4 points]

Consider the system in part (b) except that now the reorder buffer size is 48 entries. What is the average IPC for this processor? If the IPC is less than 12, then what is the smallest reorder buffer size for which the IPC will be 12 again (assume the reorder buffer size can only be a multiple of 12).

NOTE: ONLY GRADUATE STUDENTS SHOULD SOLVE THE NEXT PROBLEM.

Problem 5 - GRADUATE STUDENTS PROBLEM [10 points]

In class, we studied the use of a reorder buffer to maintain precise interrupts. With this mechanism, an instruction does not update the register file with its newly computed value until it is committed. Instead, the new value is stored in the reorder buffer. This requires later instructions to possibly read source operand values from the reorder buffer.

An alternative is to use a history buffer. This mechanism updates the register file as soon as the instruction computes a new value, but it stores the previous value of the register in the history buffer. On an interrupt, appropriate old values are restored.

Consider using the above history buffer idea to maintain precise interrupts with the standard Tomasulo algorithm design (with reservation stations) as covered in class (no reorder buffer). Explain the modifications to the Tomasulo pipeline for this purpose.

Hint: As with the reorder buffer, we need to split the Write stage into Complete and Commit.

Separate your answer into the following parts. You need only give the conceptual changes from the basic Tomasulo algorithm (e.g., at the level discussed in the lecture notes for the reorder buffer).

Part A [2 points] Explain how the fields of the history buffer would be different from the reorder buffer.

Part B [1 point] Describe the changes to the Issue stage.

Part C [1 point] Describe the changes to the Execute stage.

Part D [2 points] Describe the changes to the Complete stage.

Part E [2.5 points] Describe the changes to the Commit stage (for all instructions, including for stores and branches).

Part F [1.5 points] How is an interrupt handled?