

CS433 Homework 2 (Chapter 3)

Assigned on 9/19/2017
Due in class on 10/5/2017

Instructions:

1. Please write your name and NetID clearly on the first page.
2. Refer to the course fact sheet for policies on collaboration.
3. Due **IN CLASS** on 10/5/2017.

Problem 1 [40 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) IS and WB take one cycle each.
- 3) One Instruction is issued per cycle.
- 4) Memory accesses use the integer functional unit to perform effective address calculation during the EX stage.
- 5) For stores, memory is accessed during the EX stage (Tomasulo's algorithm without speculation) or commit stage (Tomasulo's algorithm with speculation).
- 6) All loads access memory during the EX stage.
- 7) Loads and Stores stay in EX for one cycle.
- 8) Functional units are not pipelined.
- 9) 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 .
- 10) An instruction waiting for data on the CDB can move to its EX stage in the cycle after the CDB broadcast.
- 11) Only one instruction can write to the CDB in one clock cycle.
- 12) Branches and stores do not need the CDB.
- 13) 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.

- 14) 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).
- 15) Assume that the BNEQZ occupies the integer functional unit for its computation and spends one cycle in EX.

Part A [10 points]

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; registers, functional units, 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 two entries are given.

Instruction	IS	EX	WB	Reason for stalls
L.D F0, 0(R1)	1	2	3	
ADD.D F2, F0, F4	2	4-8	9	RAW on F0 (from 1)
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? If so, find an example from the code and explain. Otherwise, explain why.

Part C [2 points]

Would adding another multiplier functional unit be more advantageous than pipelining the multiplier for the above code segment? Explain your answer.

Part D [18 points]

For this part, assume *hardware speculation* and *dual-issue* added to the Tomasulo pipeline you used for the 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 issue 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 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. 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)
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 the part D, which of the following optimizations will cause a performance improvement of at least one cycle per loop iteration: (1) triple issue, (2) three instruction commits per cycle, or (3) reorder buffer of size 14? Explain why.

Part F [2 Points]

For the code in the part D, assume a floating point division 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 [10 points]

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

```
ADD.D R1, R0, R0

L1:  ADD.D 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 each branch. For instance, the branch 1 will be executed 12 times, and those 12 times should be recorded in the table for the branch 1. Similarly, the 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? Fill the following tables to record the prediction and action of each branch.

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

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

Part B [6 Points]

Now assume that 2-bit saturation counters are used. When the processor starts to execute the above code, both counters contain value 0. What is the number of correct predictions? Fill the following tables to record the prediction and action of each branch.

Step	Counter Values	Branch 1 Prediction	Actual Branch 1 Outcome
1	00	N	T
2			T
3			
4			
5			
6			
7			
8			
9			
10			
11			
12			

Step	Counter Values	Branch 2 Prediction	Actual Branch 2 Outcome
1	00	N	T
2			
3			
4			

Problem 3 [10 Points]

Suppose we have a deeply pipelined processor, for which we implement a branch target buffer (BTB) for conditional branches only. Assume the followings:

- (1) 15% of the total instructions are conditional branches.
- (2) the BTB hit rate is 90%.
- (3) the BTB miss penalty is always 3 cycles.
- (4) even if hit on the BTB, if the prediction was incorrect, there is a 4-cycle penalty.
- (5) the branch prediction accuracy is 90%.
- (6) the base CPI without branch stall is 1.

Part (A) [5 points]

How much faster is this processor compared to a processor that does not have a branch predictor and has a fixed two-cycle branch penalty?

Part (B) [5 points]

For this part, let's assume that the BTB is extended to store one target instruction and used to convert an unconditional branch into the target instruction at the IF stage (i.e., branch folding). Also, assume that 10% of the total instructions are unconditional branches and an unconditional branch originally takes 1 cycle to execute without the BTB. Calculate the CPI with the extended BTB.

Problem 4 (*for graduate students only*) [10 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 as described 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 dependence 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 (IPC) of this processor?

Part (B) [3 points]

Consider the system in the part A except that now every 48th instruction is a load that misses in the cache and the miss latency is 500 cycles. What is the average IPC of this processor?

Part (C) [5 points]

Consider the system in the 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).