CS433: Computer Architecture – Fall 2022
Homework 2
Total Points: 41 points
All students should solve all problems
Due Date: September 20, 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 [5 points]

Consider two different machines. The first has a single cycle datapath (i.e., a single stage, non-pipelined machine) with a cycle time of 15 ns. The second is a pipelined machine with 5 pipeline stages and a cycle time of 3ns.

Part (A) [1 point]

What is the speedup of the pipelined machine versus the single cycle machine assuming there are no stalls?

Part (B) [2 points]

What is the speedup of the pipelined machine versus the single cycle machine if the pipeline stalls 1 cycle for 25% of the instructions?

Part (C) [2 points]

Now consider a 4 stage pipeline machine with a cycle time of 3.1 ns. Again, assuming no stalls, is this implementation faster or slower than the original 5 stage pipeline? Explain your answer.

Problem 2 [4 points]

Consider two different 5-stage pipeline machines (IF ID EX MEM WB). The first machine resolves branches in the ID stage, uses one branch delay slot, and can fill 80% of the delay slots with useful instructions. The second machine resolves branches in the EX stage (i.e., it determines whether the branch is taken and the target address of a taken branch in the EX stage) and uses a predict-not-taken scheme. Assume that the cycle times of the machines are identical. Given that 35% of the instructions are branches, 25% of branches are taken, and that stalls are due to branches alone, which machine is faster? To get any credit, you must justify your answer.

Problem 3 [10 points]

Consider the following loop.

loop:

1.   *ADDI R2, R2, #1*
2.   *LD R4, 0(R3)*
3.   *LD R5, 4(R3)*
4.   *ADD R6, R4, R5*
5.   *MUL R4, R6, R7*
6.   *SUBI R3, R3, #8*
7.   *BNEZ R2, loop*
8.   *ADD R11, R12, R13*

Part (A) [4 points]

Identify all data dependencies (potential data hazards) in the given code snippet. Assume the loop takes exactly one iteration to complete. Specify if the data dependence is RAW, WAW or WAR.**Part (B) [2 points]**

Assume a 5-stage pipeline (IF ID EX MEM WB) without any forwarding or bypassing hardware, but with support for a register read and write in the same cycle. Also assume that branches are resolved in the ID stage and handled by stalling the pipeline. All stages take 1 cycle. Again, the loop takes one iteration to complete. Which dependencies from part (a) cause stalls? How many cycles does the loop take to execute?

Part (C) [2 points]

Assume that the pipeline now supports full forwarding and bypassing. Furthermore, branches are handled as predicted-not-taken. As before, the loop takes one iteration to complete. Which dependencies from part (a) still cause stalls and why? How many cycles does the loop take to execute now?

Part (D) [2 points]

If the pipeline from part (c) instead uses a branch delay slot, how would you schedule the instructions in the loop to minimize stalls? For this part, assume the *loop takes multiple iterations to complete*. Explain your answer.

Problem 4 [8 points]

High-performance processors often have deep pipeliness. Imagine that you have a 10-stage pipeline in which every stage of the 5-stage pipeline discussed in class has been split in two with stages labeled IF1, IF2, ID1, ID2, EX1, etc. Data is forwarded from the end of a pair of stages to the beginning of the two stages where they are needed. For example, the output of the second execute stage is forwarded as the input of the first execute stage, still causing a 1-cycle delay.

Part (A) [6 points]

For the above 10-stage pipeline, show the timing of the instruction sequence below for the first 15 cycles. Assume full forwarding and bypassing hardware as described above. Assume that branches are (magically) predicted as taken and that there is magic hardware that ensures that branches that are actually taken do not incur any stalls Also assume that the loop executes for more than one iteration.

```
1   Loop: ld    x1, 0(x2)      // load x1 from address 0+x2
2         addi x1, x1, #1      // x1 = x1 + 1
3         sd    x1, 0(x2)      // store x1 at address 0+x2
4         addi x2, x2, #4      // x2 = x2 + 4
5         sub  x4, x3, x2      // x4 = x3 - x2
6         benz       x4, Loop  // branch to Loop if x4 != 0
```

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	IF1	IF2	ID1	ID2	EX1	EX2	M1	M2	W1	W2					...
2															...
3															...
4															...
5															...
6															...
7															...

Part (B) [2 points]

Assume that in the 5-stage pipeline, the longest stage requires 0.8 ns, and the pipeline register delay is 0.1 ns. What is the clock cycle time of the 5-stage pipeline? If the 10-stage pipeline splits all stages in half, what is the cycle time of the 10-stage machine? Assume the pipeline register delay does not change.

Problem 5 [14 points]

For this problem, we will explore a pipeline for a register-memory architecture. The architecture has two instruction formats: a register-register format and a register-memory format. In the register-memory format, one of the operands for an ALU instruction could come from memory.

There is a single memory-addressing mode (offset + base register). The only non-branch register-memory instructions available have the format:

Op Rdest, Rsrc1, Rsrc2

or

Op Rdest, Rsrc1, MEM

where Op is one of the following: Add, Subtract, And, Or, Load (in which case Rsrc1 is ignored), or Store. Rsrc1, Rsrc2, and Rdest are registers. MEM is a (base register, offset) pair.

Branches compare two registers and, depending on the outcome of the comparison, move to a target address. The target address can be specified as a PC-relative offset or in a register (with no offset). Assume that the pipeline structure of the machine is as follows:

IF RF ALU1 MEM ALU2 WB

The first ALU stage is used for effective address calculation for memory references and branches. The second ALU stage is used for operations and branch comparison. RF is both decode and register-fetch stage. Assume that when a register read and a register write of the same register occur in the same cycle, the write data is forwarded.

Part (A) [4 points]

Find the number of adders, counting any adder or incrementor, needed to minimize the number of structural hazards. Justify why you need this number of adders.

Part (B) [4 points]

Find the number of register read and write ports and memory read and write ports needed to minimize the number of structural hazards. Justify why you need this number of ports for the register file and memory.

Part (C) [3 points]

Will data forwarding from the ALU2 stage to any of ALU1, MEM, or ALU2 stages reduce or avoid stalls? Explain your answer for each stage.

Part (D) [3 points]

Will data forwarding from the MEM stage to any of ALU1, MEM, or ALU2 stages reduce or avoid stalls? Explain your answer for each stage.