

CS/CoE 1541 – Exam 1 (Spring 2019). Name: _____

Question 1 (15 points): In this problem, consider the execution of the following code segment on a 5-stage pipeline with forwarding/stalling hardware and branches resolved in ID. Assume that the loop I2-I5 in the following code executes 1000 times (branches are not delayed branches).

```

I1: sub  $3, $0, $0
I2: lw   $1, 0($2)
I3: add  $3, $3, $1
I4: addi $2, $2, 4
I5: bneq $5, $2, I2    /*branch to I2 if contents of $5 and $2 are not equal*/
I6: sw   $3, 0($2)
    
```

(a) Complete the following trace for the execution of the above code (up to cycle 15) assuming no branch prediction (that is always predicting that a branch is not taken).

	IF	ID	EX	MEM	WB
Cycle 1					
Cycle 2					
Cycle 3					
Cycle 4					
Cycle 5					
Cycle 6					
Cycle 7					
Cycle 8					
Cycle 9					
Cycle 10					
Cycle 11					
Cycle 12					
Cycle 13					
Cycle 14					

(b) How many cycles would it take to complete the execution of the 1000 iterations of the loop (4000 instructions) – still assuming no branch prediction?

(c) What would be the CPI while executing the loop (with no branch prediction)?

(d) What would be the CPI while executing the loop is we assume a perfect branch predictor?

Question 4 (15 points): Consider a pipelined CPU with **6-stages** (IF, ID, EX1, EX2, MEM, WB) with branch conditions and target addresses being resolved in the EX2 stage. The delays for the six stages are 1.7, 1.9, 1.8, 2.0, 1.7 and 2.0 ns respectively. Assume that, on average, 30% of the instructions executed on this architecture are load/store instructions and 20% are branch instructions with 30% of the executed branches taken. The CPI of the architecture is 3.0 if the stalls due to control hazards are ignored.

(a) What is the CPI for the architecture when stalls due to control hazards are included?

- $CPI_a =$

(b) Repeat part (a) if the architecture is modified such that the logic for calculating the branch target and address is moved from EX2 to ID. This modification, however, will increase the delay of the ID stage from 1.9 to 2.1 ns and reduce the delay of the EX2 stage from 2.0 to 1.8 ns.

- $CPI_b =$

(c) Repeat part (a) if branches are resolved in EX2 but hardware for branch prediction is added to IF resulting in an 90% prediction accuracy. That is, after a branch, the correct instruction is fetched 90% of the time and the wrong instruction 10% of the time. The branch predictor will increase the delay of the IF stage from 1.7 to 1.9 nanoseconds.

- $CPI_c =$

(d) What is the cycle time for each of the above three architectures?

- $T_a =$ the cycle time for the original architecture =

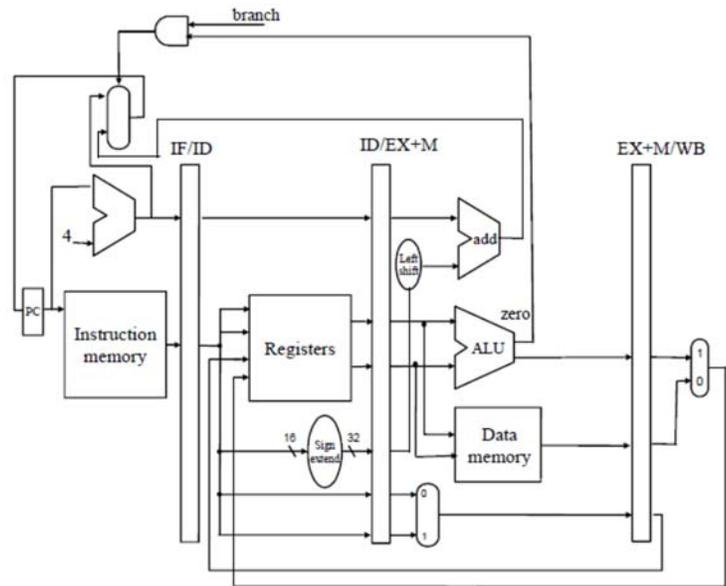
- $T_b =$ the cycle time when branch resolution is moved to ID as described in (b) =

- $T_c =$ the cycle time when a branch predictor is added as described in (c) =

(e) How would you determine which of the three architectures is the most efficient?

Question 5 (15 points): In this question, you will explore changing the MIPS ISA so that *lw/sw* instructions do not use an immediate constant. That is, the memory address is found in the register without the capability of adding a constant (this is called register indirect addressing). With this modification, it is possible to have a 4-stage pipeline by merging the EX and MEM stages into one stage (call it EX+M), as shown in the figure below.

The advantage of the 4-stage pipeline is that forwarding completely eliminates data hazards. Its disadvantage is that the **number of executed instructions in a program increases by 10%** because each instruction of the form “*lw \$r1, I(\$r2)*” has to be replaced by two instructions (“*addi \$r2, \$r2, I*” followed by “*lw \$r1, \$r2*”) and each instruction of the form “*sw \$r1, I(\$r2)*” has to be replaced by two instructions (“*addi \$r2, \$r2, I*” followed by “*sw \$r1, \$r2*”).



- (a) Assume that it takes 90ps to read from the register file, 90ps to write into the register file and 200ps to fetch or store into memory (instruction or data). Assume also that the Adder and ALU delays are 150ps and 240ps, respectively and ignore all other delays. What is the minimum cycle time for this 4-stage pipeline?

- (b) To compare the performance of the 4-stage pipeline with the original 5-stage pipeline, we assume that both have the same cycle time, use forwarding and resolve branches in the EX stage.
- a. Given that the CPI for the 5-stage pipeline is 2.05 and that the probability of a load-use data hazard in the 5-stage pipeline is 5% (the probability that an instruction is a *lw* instruction that loads into a register R and that this *lw* instruction is immediately followed by an instruction which reads from R). Compute the CPI of the 4-stage pipeline.

$$CPI_{5\text{-stage}} = 2.05$$

$$CPI_{4\text{-stage}} =$$

- b. Which pipeline would be more efficient? – should quantitatively support your answer

Question 6 (15 points): Consider the scheduling of the following loop on a superscalar architecture with two pipelines, one for ALU/branch instructions and the other for lw/sw instructions.

- L1: lw \$t0, 1000(\$s4)
- L2: addi \$s1, \$s1, 1
- L3: lw \$t1, 1004(\$s4)
- L4: addi \$s5, \$t1, 4
- L5: add \$t0, \$t0, \$s1
- L6: sw \$s1, 5000(\$s4)
- L7: beq \$s4, \$s6, L

a) Show the scheduling of one iteration of the loop assuming that the architecture has forwarding and stalling hardware but that the compiler **does not** reorder the instructions:

ALU/branch pipeline	Load/store pipeline	
		Cycle 1
		Cycle 2
		Cycle 3
		...
		...
		...
		...
		...

b) It is possible to improve the efficiency of the code on the superscalar (without affecting the correctness of the code) by reordering the instructions. Specifically, by moving instruction L___ to the position between L___ and L___.

c) Show the pipeline schedule after you reorder the instructions as specified in part b.

ALU/branch pipeline	Load/store pipeline	
		Cycle 1
		Cycle 2
		Cycle 3
		...
		...
		...
		...
		...

The following slide (from lecture 1) shows the machine instruction formats for the three types of instructions:

Rtype rd, rs, rt

I-type rt, rs (I)

J-type L



Overview of MIPS

- only three instruction formats

•

R	op	rs	rt	rd	shamt	funct
I	op	rs	rt	16 bit constant/offset		
J	op	26 bit relative address				

- Arithmetic and logic instructions use the **R-type** format
 - Memory and branch instructions (*lw, sw, beq, bne*) use the **I-type** format
 - Jump instructions use the **J-type** format
- Refer to section A.10 for complete specifications of MIPS instructions