

CSE 560 – Practice Problem Set 4 Solution

1. Use the following RISC code fragment:

- (1) loop: load r1, 0(r2)
- (2) addi r1 ← r1, #1
- (3) store 0(r2), r1
- (4) addi r2 ← r2, #4
- (5) sub r4 ← r3, r2
- (6) bnez r4, loop
- (7) next instruction
- (8) next instruction

You may assume that the initial value of r3 is r2 + 396.

(a) Show the timing of this instruction sequence (i.e., draw a pipeline diagram) for our 5-stage pipeline without any forwarding or bypassing hardware but assuming a register read and a write in the same clock cycle “forwards” through the register file. Assume that the branch is handled by flushing the pipeline (use the notation f in the pipeline diagram). If all memory references hit in the cache (i.e., take one cycle), how many cycles does this loop take to execute?

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
(1)	F	D	X	M	W													
(2)		F	D	d*	d*	X	M	W										
(3)			F	p*	p*	D	d*	d*	X	M	W							
(4)						F	p*	p*	D	X	M	W						
(5)									F	D	d*	d*	X	M	W			
(6)										F	p*	p*	D	d*	d*	X	M	W
(7)													F	p*	p*	D	f*	
(8)																F	f*	
(1)																	F	D

17 – 1 cycles = 16 cycles to execute

(b) Show the timing of this instruction sequence for our 5-stage pipeline with normal forwarding and bypassing hardware. Assume that the branch is handled by predicting it as not taken. If all memory references hit in the cache, how many cycles does this loop take to execute?

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
(1)	F	D	X	M	W													
(2)		F	D	d*	X	M	W											
(3)			F	p*	D	X	M	W										
(4)					F	D	X	M	W									
(5)						F	D	X	M	W								
(6)							F	D	X	M	W							
(7)								F	D	f*								
(8)									F	f*								
(1)									F	D								

10 – 1 cycles = 9 cycles to execute

2. Your job is to build the latest and greatest handheld device, and your company has purchased the rights to use a family of processor designs (notice the similarity to ARM's business model, think of your company as a customer of ARM). The immediate task is to choose which micro-architecture (within the family) that best meets the needs of your new device. The two choices are B and L (for big.LITTLE, ARM systems that support two different types of cores in a single chip, a point that is irrelevant to the current problem).
 - Since both designs under consideration (B and L) are ARM processors, the compiler and the ISA can be considered to be the same.
 - Both designs use the same 5-stage pipeline with stages F, D, X, M, W
 - The mix of instruction types is as follows:
 - i. 20% conditional branch instructions
 - ii. 20% load instructions
 - iii. 10% store instructions
 - iv. 50% data manipulation (ALU) instructions (of which, 20% of these are floating point instructions)
 - In design B, floating point instructions are fully pipelined and therefore have a negligible number of pipeline stalls (i.e, you can ignore them).
 - In design L, the floating point execution hardware is not pipelined, resulting in an average stall of 3 clock cycles for each floating point instruction.
 - In design B, a hybrid branch predictor is present in the fetch (F) pipeline stage that predicts correctly 90% of the time. Mispredictions trigger a 2-clock cycle flush.
 - In design L, a simpler branch predictor is present in the decode (D) pipeline stage that predicts correctly 80% of the time. Mispredictions trigger a 2-clock cycle flush. Don't forget that correct predictions will have a 1-clock cycle bubble for this design.
 - Both designs can be clocked at 2 GHz.

- (a) Write an analytical expression for CPI_B and CPI_L , including identifiable components for each instruction type.

Variable	Meaning
CPI_{branch}	CPI for conditional branch instructions
$CPI_{loadstore}$	CPI for load and store instructions
CPI_{data}	CPI for data instructions

$$CPI_{branch,B} = 1 + (1-0.9) \times 2 = 1.2 \quad \text{10\% of the time the branch is mispredicted}$$

$$CPI_{branch,L} = 1 + (0.8) \times 1 + (1-0.8) \times 2 = 2.2 \quad \text{20\% of the time the branch is mispredicted}$$

$$CPI_{loadstore,B} = CPI_{loadstore,L} = 1$$

$$CPI_{data,B} = 1$$

$$CPI_{data,L} = 1 + (0.2) \times 3 = 1.6$$

$$CPI_B = (0.2) \times CPI_{branch,B} + (0.2+0.1) \times CPI_{loadstore,B} + (0.5) \times CPI_{data,B}$$

$$CPI_L = (0.2) \times CPI_{branch,L} + (0.2+0.1) \times CPI_{loadstore,L} + (0.5) \times CPI_{data,L}$$

- (b) Substitute numbers into the above expressions, giving numerical values for CPI_B and CPI_L .

$$CPI_B = .2(1.2) + 0.3(1) + 0.5(1) = 1.04$$

$$CPI_L = .2(2.2) + 0.3(1) + 0.5(1.6) = 1.54$$

- (c) For a benchmark program with 1 billion (10^{10}) instructions, what is the time to completion for each design?

$$T_B = IC \times CPI_B \times 0.5ns = (10^{10}) \times (1.04) \times (0.5 \times 10^{-9}) = 5.2 \text{ seconds}$$

$$T_L = IC \times CPI_L \times 0.5ns = (10^{10}) \times (1.54) \times (0.5 \times 10^{-9}) = 7.7 \text{ seconds}$$

- (d) Which processor is faster, and what % faster is it?

$$\text{Speedup} = T_L / T_B = 7.7 / 5.2 = 1.5, \text{ design B is 1.5 times faster than design L}$$

However, the question really asked for % faster,

$$\% \text{faster} = 100(T_L / T_B - 1) = 50\%, \text{ design B is 50\% faster than design L}$$