EEL308 Minor Exam 2 Semester 1302 **Full marks: 20** **21 March 2014**

(Permitted references: Textbook (Patterson, Hennessy) and lecture notes of this course only.)

Absolutely no Internet access, use of Wikipedia, etc. This and/or any other malpractice will earn you an 'F' grade in the entire course. No explanations will be entertained.

Please begin your answer to each question on a new page.

Good luck! ☺

1. Consider the following piece of MIPS assembly code:

```
loop:   lw    $1, 0($2)
        add   $3, $1, $2
        sw    $3, 0($2)
        addi  $2, $2, 4
        bne   $2, $4, loop
```

Let's suppose we have a single-cycle MIPS implementation with the stages for instruction fetch, decode, execution, memory access and register write taking 200 ps, 100 ps, 200 ps, 200 ps and 100 ps respectively.

(a) Using this implementation calculate the time taken for one iteration of the above loop.

(b) Now, let's take the five-stage MIPS pipelined implementation (with the complete control path; and branch decision in the ID stage). What will be the total time taken to complete one iteration of the above loop with the same sequence of instructions?

(c) Let's say the content of $4 is large enough that we can statically unroll this loop 10 times, i.e., 10 iterations, without any branch misprediction. However, we are not allowed to change the order of the first three instructions within any iteration. Compare the single-cycle implementation with the pipelined implementation in this scenario (mention the times taken), and determine the pipeline speedup, if any. (Note that we are still executing everything in-order.)

(1+2+4=7 marks)

2. Let's suppose we have an architecture where the branch misprediction penalty is 2 clock cycles. In a local branch predictor scheme using 4-bit local history and indexed by the last 8-bits of the branch address, the misprediction frequency is 10%. In the benchmark programs used to evaluate this architecture, the frequency of branches is 40%. The average CPI using this scheme is 1.5. After detailed analysis, we come up with a (2,2) branch predictor whose prediction buffer has the same size as the earlier case, but the misprediction frequency is now reduced to 2%.

(a) Find the CPI under the new branch prediction scheme.

(b) Find the number of entries in the new branch prediction buffer.

(3+2=5 marks)

3. What kinds of stalls are tackled by (a) register renaming (be specific) and (b) hardware speculation? (Hint: one phrase answer) **(1 mark)**

4. In the following figure, there are erroneous control values in one stage. Identify and correct them. **(2 marks)**

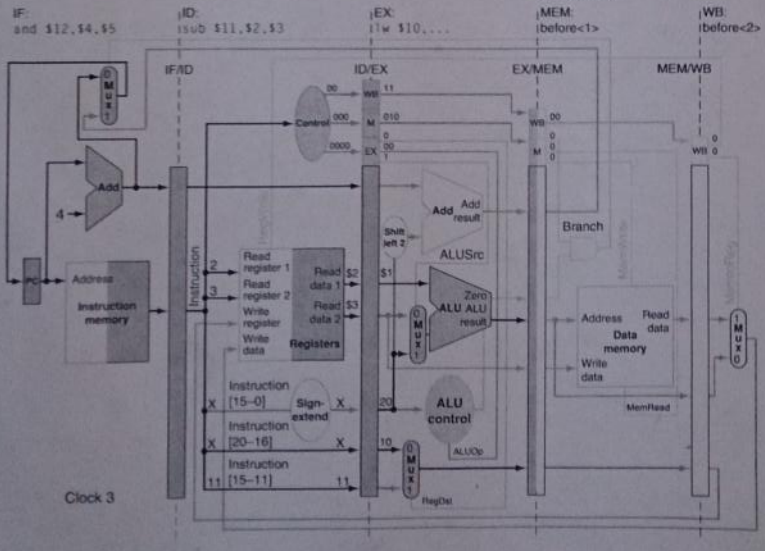

5. Consider a dynamic scheduling scheme using a scoreboard (with a store buffer) for the following piece of MIPS assembly code:

```
MUL.D    F0, F2, F4
ADD.D    F6, F0, F8
S.D      F6, 8(R1)
ADD.D    F8, F10, F4
MUL.D    F6, F10, F8
```

The **execution stage** for MUL.D takes 20 cycles, that for ADD.D takes 2 cycles, and that for S.D 1 cycle. The first instruction is issued in cycle 1. The instruction issue width is 1 instruction per cycle. There are two FP multipliers, one FP adder, one FP division unit and one integer unit. Also, each functional unit directly supplies its output to a waiting instruction. Note that issue stalls for destination register contention. In the following table, enter the cycle numbers during which each instruction is in one of the four stages.

(5 marks)

		Instruction status			
Instruction		Issue	Read operands	Execution complete	Write result
MUL.D	F0, F2, F4	1			
ADD.D	F6, F0, F8				
S.D	F6, 8(R1)				
ADD.D	F8, F10, F4				
MUL.D	F12, F10, F8				

MAL 180 2 =
MAL 358 1 =
MUL =
2D E