CMPUT 429 / CMPE 382 Mid-Term

Section 1. Short answer

Use 1 or at most 2 sentences to answer each of the following questions.

1. Define bandwidth.. Where in a CPU is this important? (5 marks)

2. Define latency. Does pipelining reduce latency? Explain. (5 marks)

3. How does an exception affect a pipeline? (5 marks)

Section 2. Short calculation / coding

4. The following sequence of diagrams depicts an instruction passing through the DLX pipeline. Which of the following instructions is it, and why? Be careful: There may be more than one answer. (15 marks)

- LW (load), SW (store word), ADD, SUBI (subtract immediate), BEQZ, SLTI (set if less than immediate), J (jump)

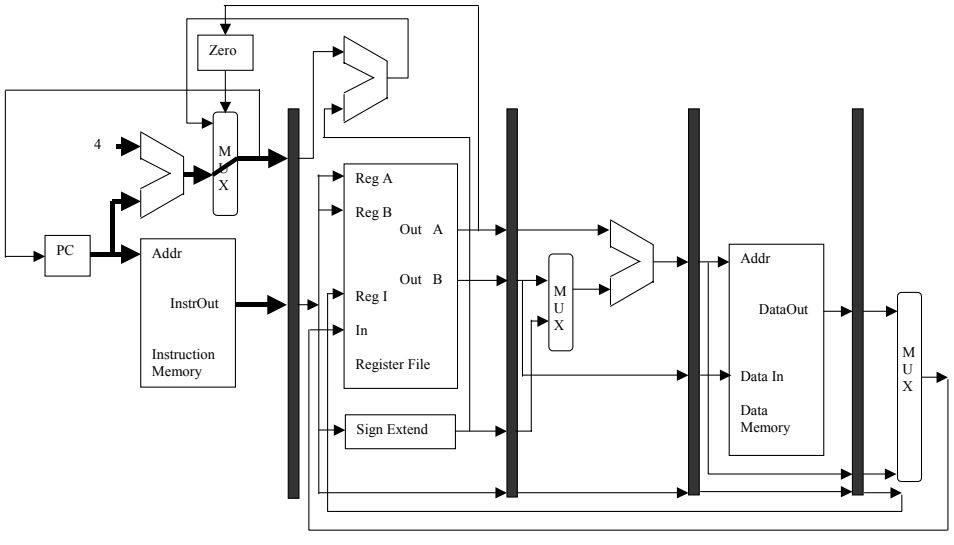

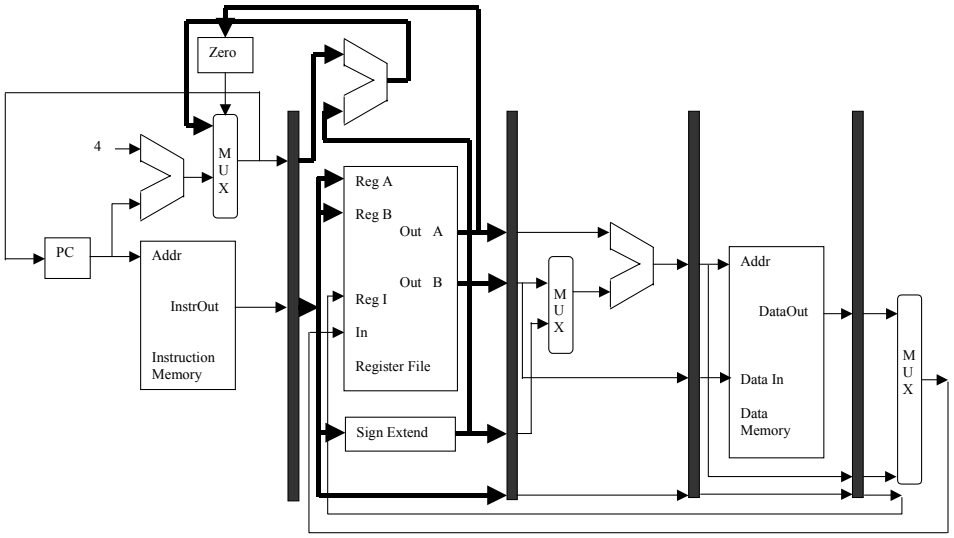
Zero
4
M
U
X
PC
Addr
InstrOut
Instruction
Memory
Reg A
Reg B
Out A
Out B
Reg I
In
Register File
Sign Extend
M
U
X
Addr
DataOut
Data In
Data
Memory
M
U
X

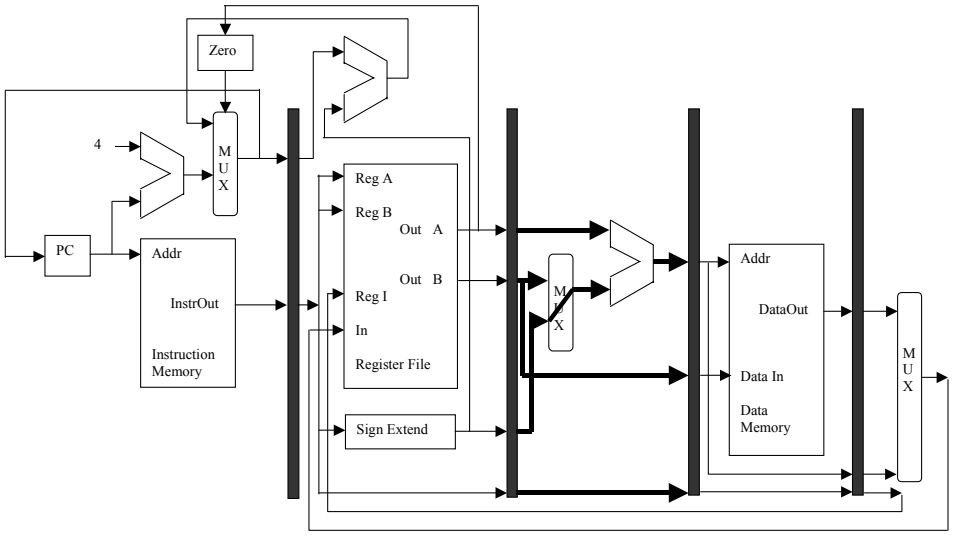
Zero
4
M
U
X
PC
Addr
InstrOut
Instruction
Memory
Reg A
Reg B
Out A
Out B
Reg I
In
Register File
Sign Extend
M
U
X
Addr
DataOut
Data In
Data
Memory
M
U
X

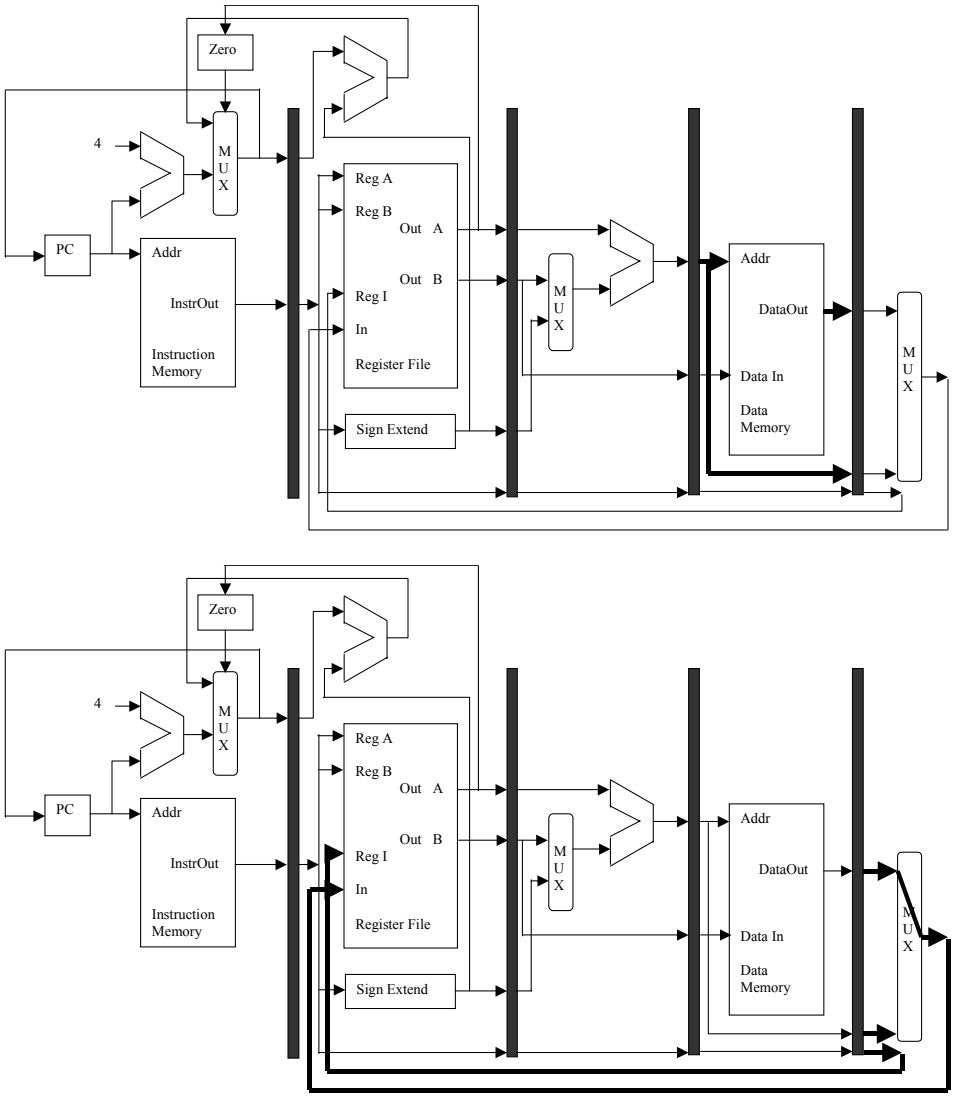

5. List all the dependencies (RAW, WAR, WAW) in the following fragment, and indicate whether they are loop-carried or not. (20 marks)

```
for (j = 1; j < 100; j-- ) {
        a[j] = b[j] + b[j+1];          /* S1 */
        b[j] = a[j-1] – c[j+1];        /* S2 */
}
```

Section 3. Design problem (10 marks per part)

You have an existing design, thoughtfully called Machine A, with a 5 stage pipeline running at 500 MHz. You are considering a new machine - Machine B – which will have a 7 stage pipeline running at 550 MHz. Both machines will have the same instruction set architecture but more effort can be expended on the compiler for machine A simply because it has been in production longer. As a result, there will be fewer stalls on Machine A. Code for both A and B has the same frequency of branches, at 20%. They have the same frequency of loads and stores: 30% loads and 10% stores. The branch delay slot on machine A can be filled 85% of the time, while on B it is only filled 80% of the time. The load delay slot on both machines can be filled 90% of the time.

a) Assuming that the pipe stages have been evenly balanced in Machine A, what is the best clock rate you would expect for Machine B, neglecting pipelining overhead ?

b) Assume that, neglecting the effect of stalls, all instructions have a CPI of 1.0. Including the effects of stalls for branches and loads, what is the effective CPI of Machine A? Machine B?

c) What is the MIPS rating for Machine A? Machine B? Is this a useful value to use when comparing these machines? Why or why not?

d) What is the comparative running time for a program on these two machines? Should you invest the effort required to build Machine B? Why or why not?

e) How fast could Machine A be made, in MIPS, just by improving the compiler's ability to fill branch delay slots?