

18

Question 1: Single-cycle CPU implementation (20 points)

You are asked to implement a hypothetical instruction $sw+$ in the single-cycle pipeline. $sw+$ is a "store word, with post increment" that is found in some real architectures (e.g., IA-64). It is encoded as an I-type instruction and performs the following operations:

$M[R[rs]] = R[rt]$	Field	op	rs	rt	imm
$R[rs] = R[rs] + imm$	Bits	31-26	25-21	20-16	15-0

Part (a)

The single-cycle datapath from lecture appears below. Show what changes are needed to support $sw+$ instruction. You should only add wires and muxes to the datapath; do not modify the main functional units themselves (the memory, register file and ALU). Try to keep your diagram neat! (10 points)

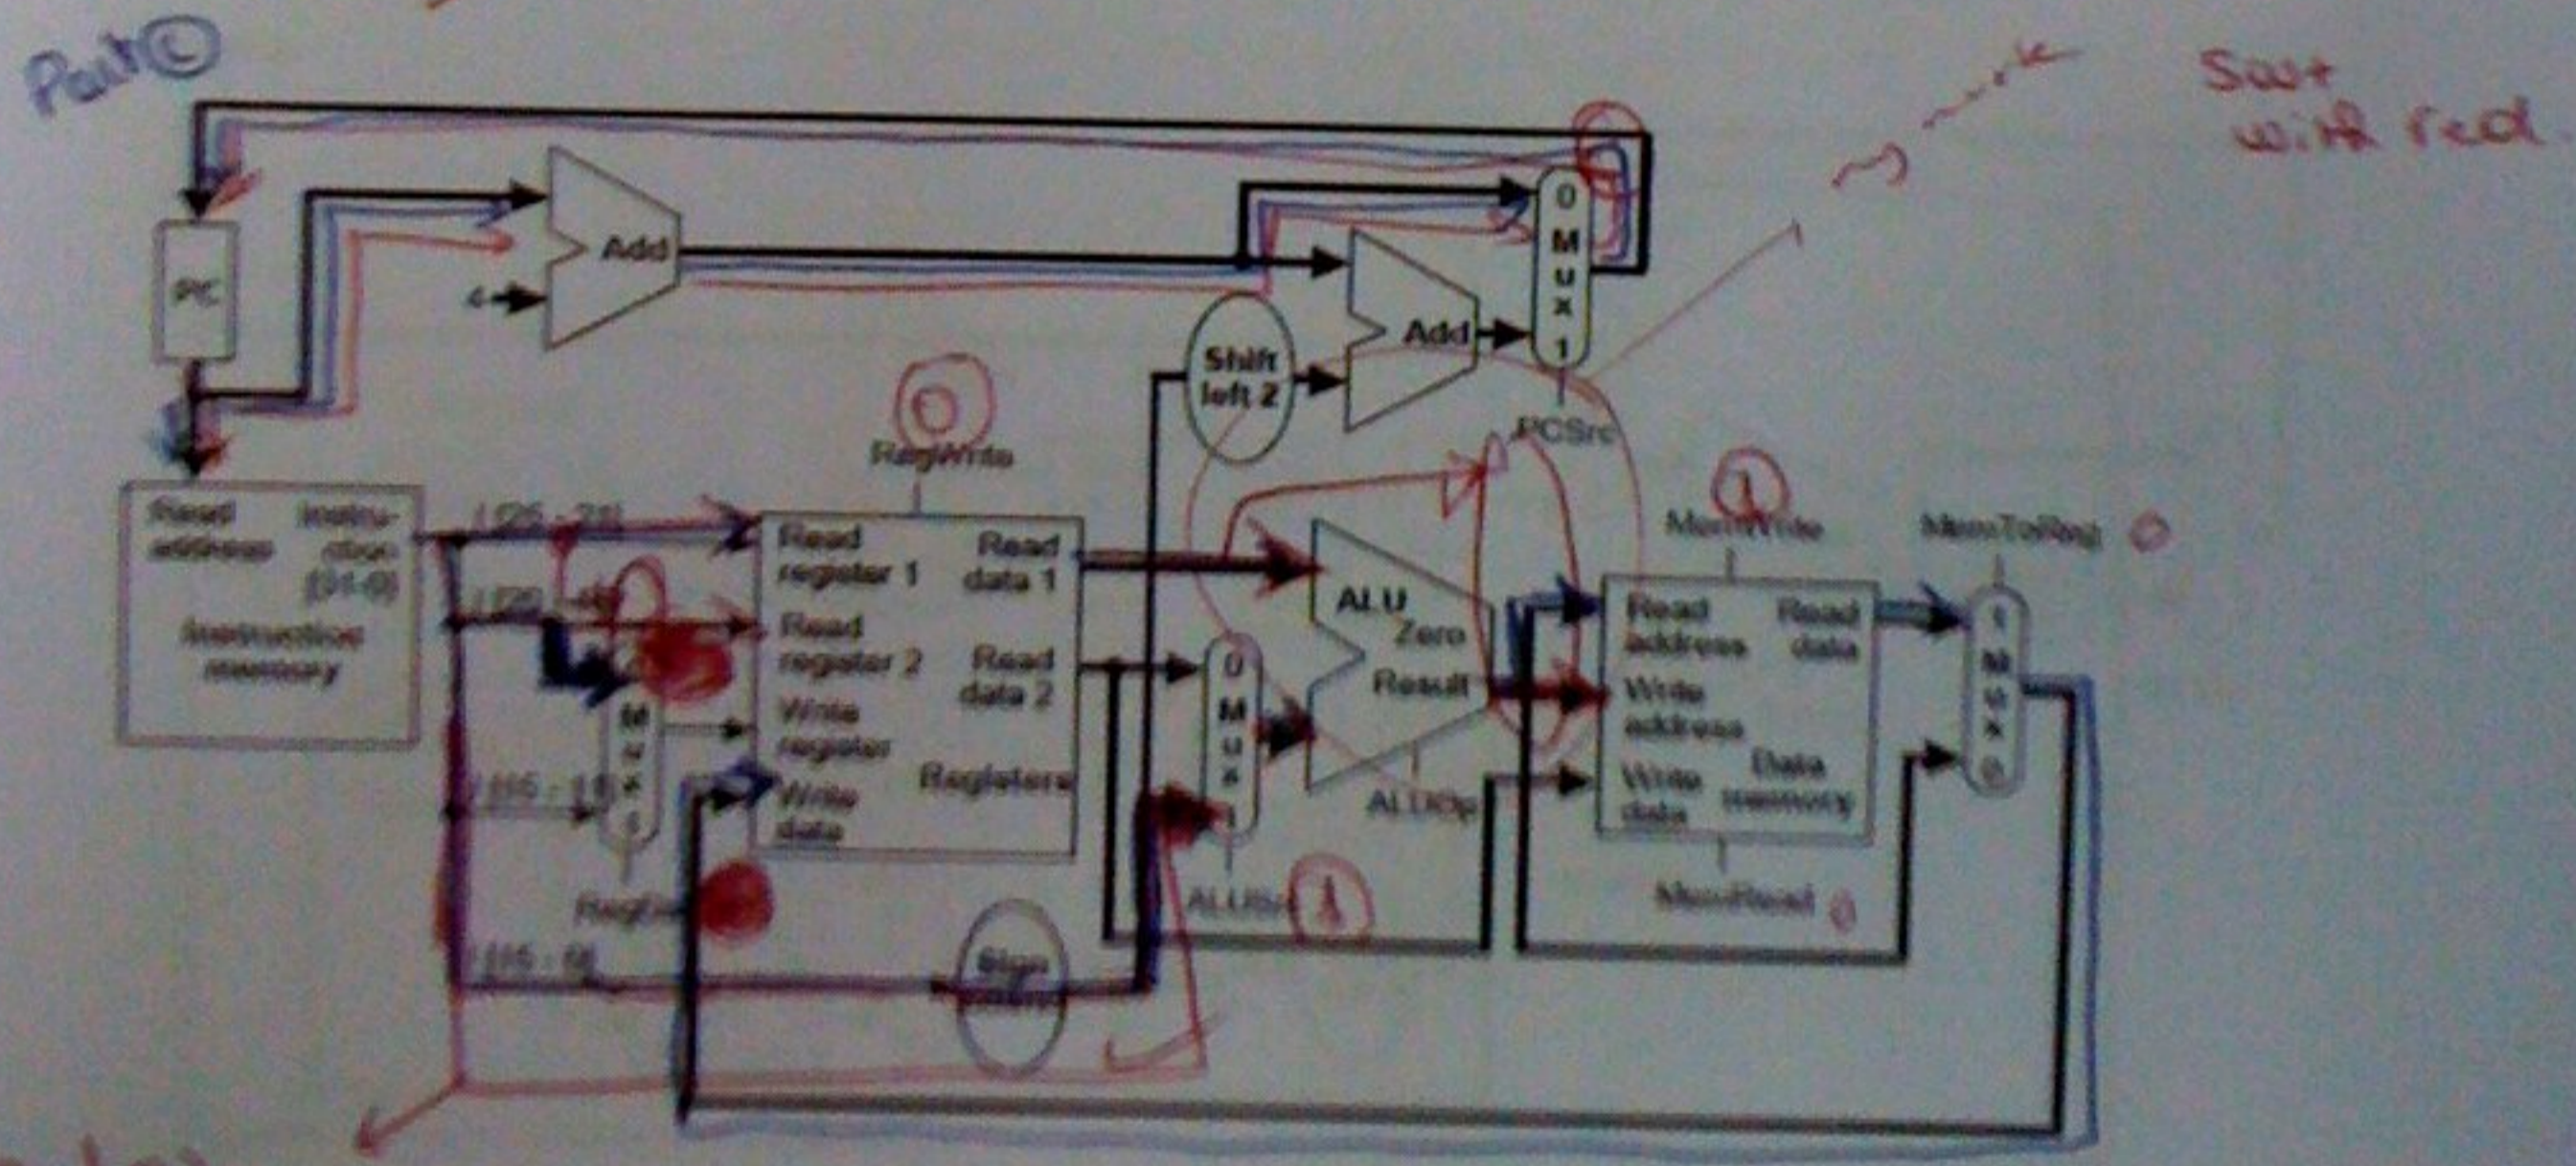
Note: While we're primarily concerned about correctness, full points will only be rewarded to solutions that do not lengthen the clock cycle. Assume that the ALU, Memory, and Register file all take 2ns, and everything else is instantaneous.

Part (b)

On the diagram below, write (next to the signal's name) values of all control signals required for the $sw+$ instruction. (5 points)

Part (c)

On the diagram below, clearly mark all wires that are active during the execution of $sw+$ instruction. (5 points) *marked in blue*



Part (a)
 The $sw+$ will add a wire from the instruction memory to the multiplexer before the ALU and this wire will transfer the immediate values.

Question 2: Multi-cycle CPU implementation and its performance (20 points)

Assume that the ALU can perform the max2 operation (*i.e.*, return the greater of 2 inputs):

$$\text{alu_result} = (\text{A_input} > \text{B_input}) ? \text{A_input} : \text{B_input};$$

ALUOp for this instruction is MAX2.

Given this improved ALU, implement the **max4** instruction that writes into register *rd* the largest value of 4 registers:

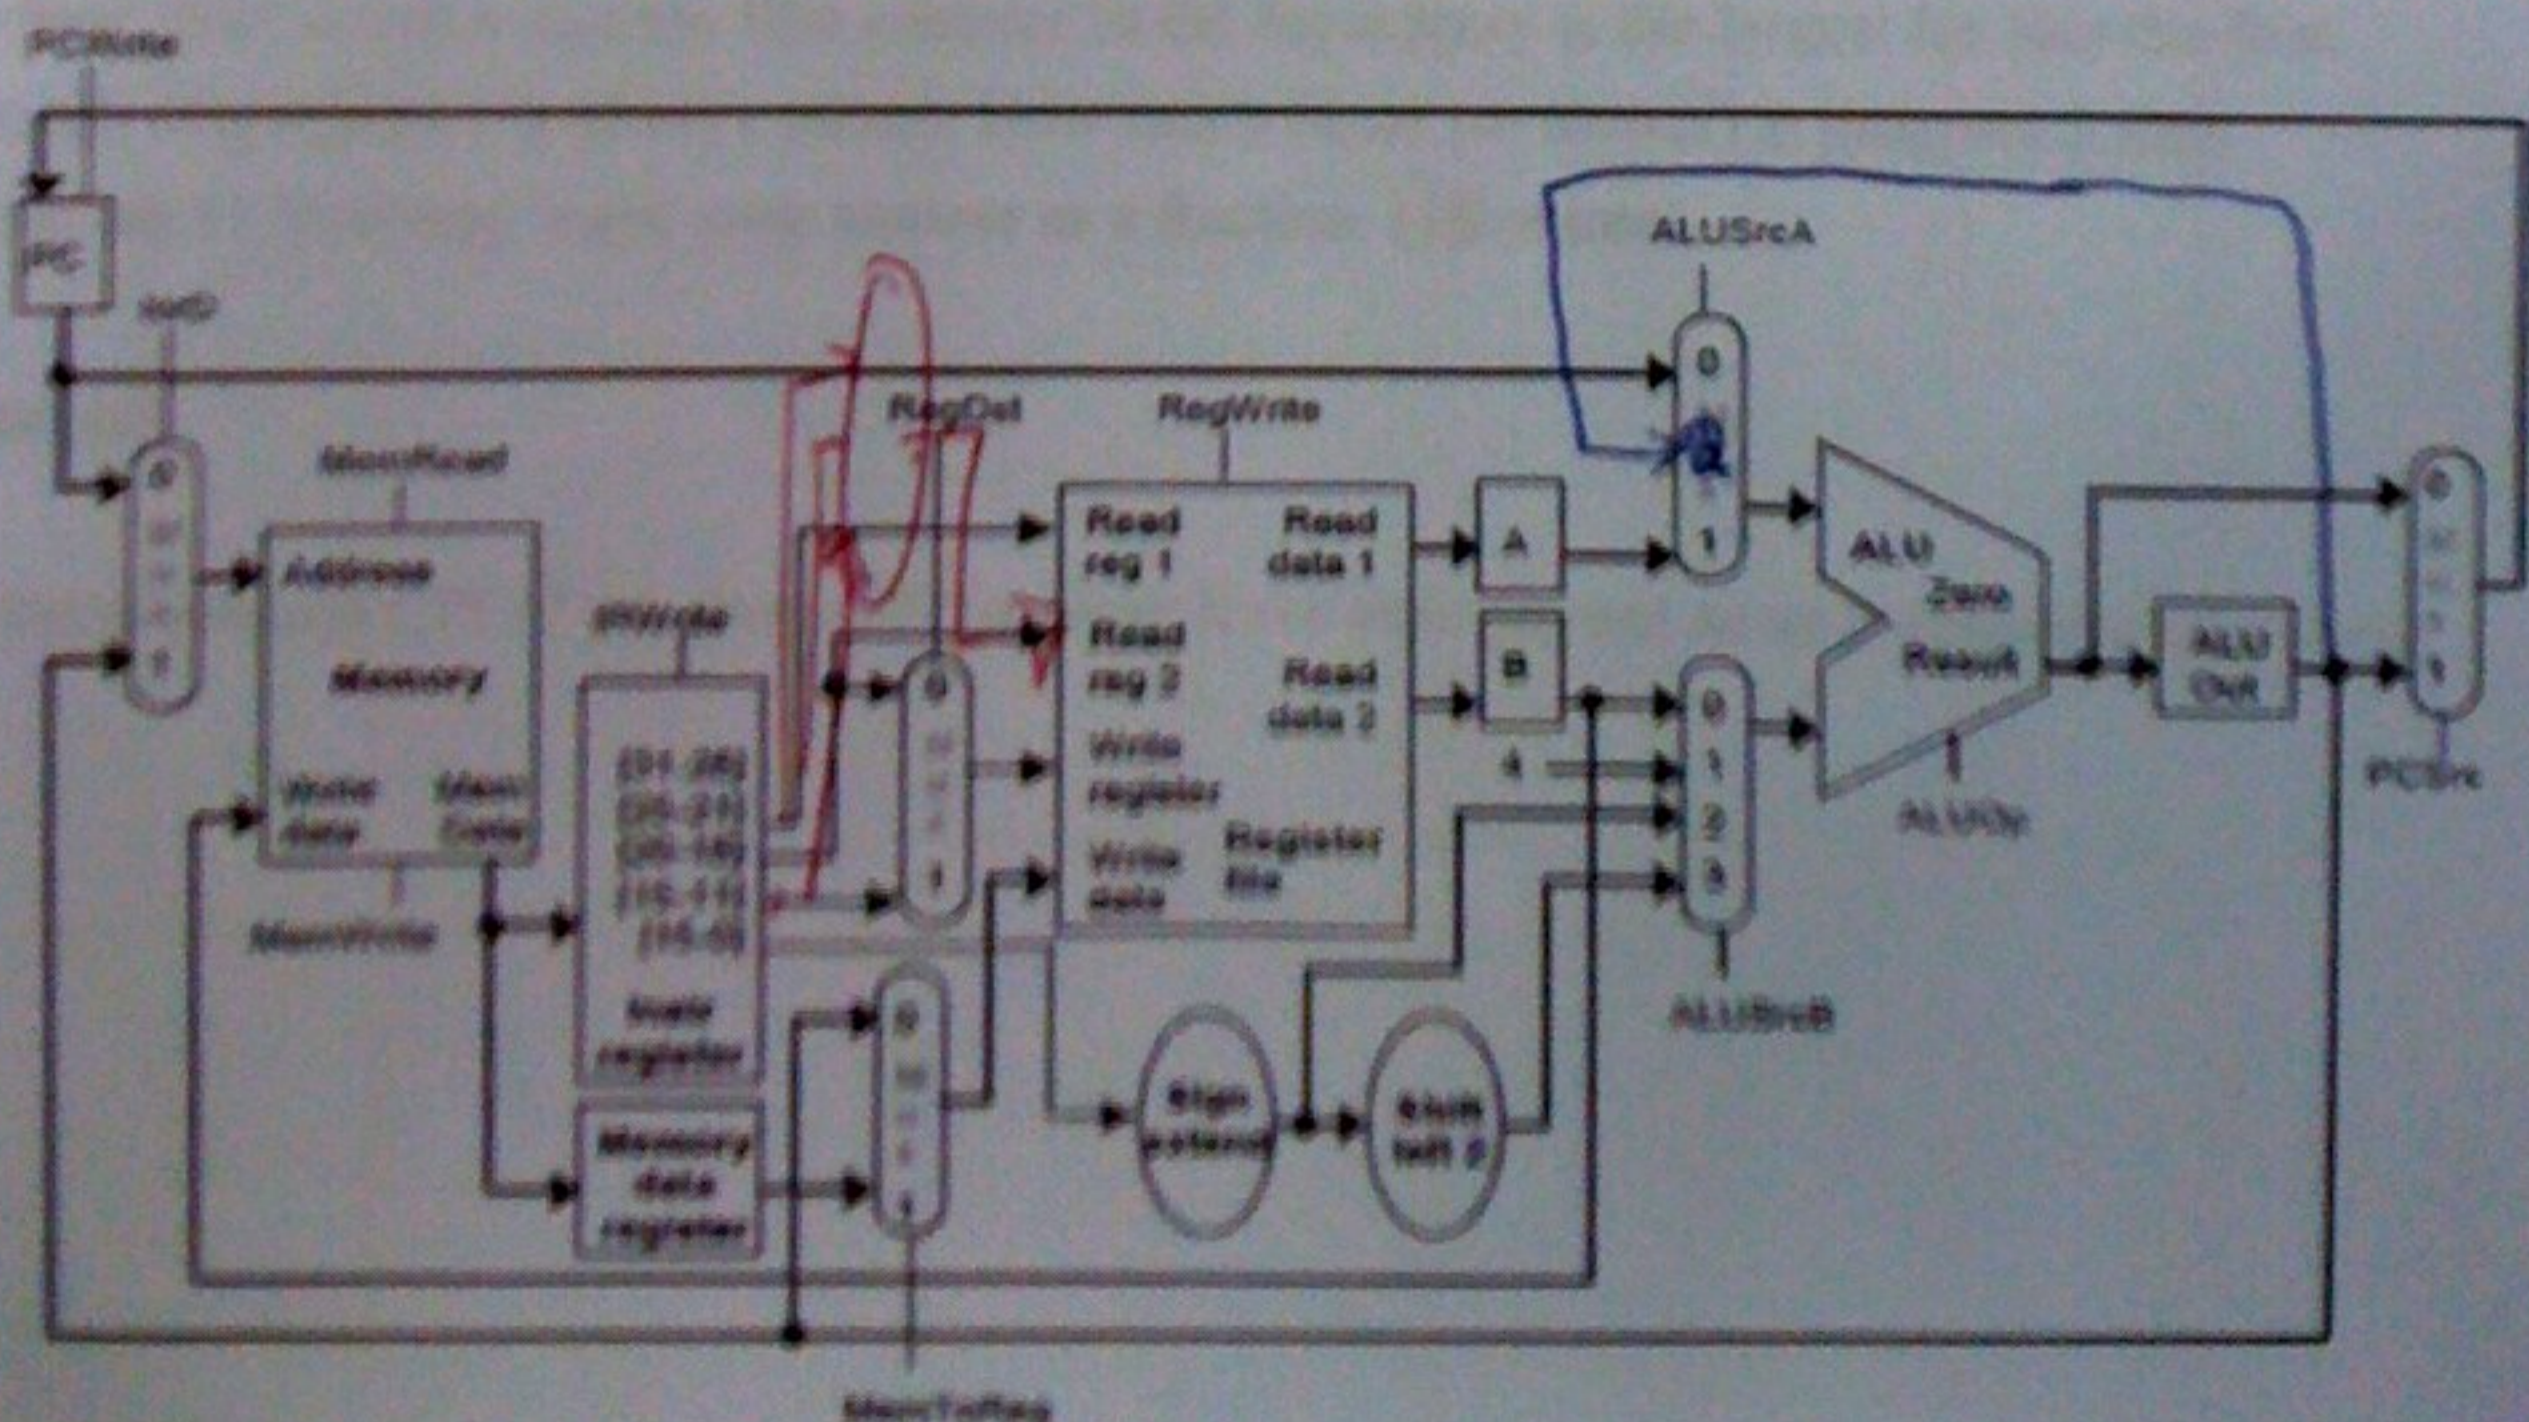
$$\text{max4 } rs, rt, rd, rm \# rd = \max(rs, rt, rd, rm)$$

Note that register *rd* is both an *input* and an *output*. Instruction **max4** has the following format:

Field	op	rs	rt	rd	rm	func
Bits	31-26	25-21	20-16	15-11	10-6	5-0

Part (a)

The multicycle datapath from lecture appears below. Show what changes are needed to support **max4**. You should only add wires and multiplexers to the datapath; do not modify the main functional units themselves (the memory, register file, and ALU). Try to keep your diagram neat! (10 points)



Note: While we're primarily concerned about correctness, full points will only be rewarded to solutions that use a minimal number of cycles and do not lengthen the clock cycle. Assume that the ALU, Memory and Registerfile all take 2ns, and everything else is instantaneous.

Part (b)

The **max4** instruction can be used in place of three branch instructions, reducing the number of instructions that need to be executed. Below are two functionally equivalent programs; the second of which uses the **max4** instruction:

Program 1	Program 2
5 lw v0, 0(a0)	lw t0, 4(a0) 5
4 add a1, a0, 12	lw t1, 8(a0) 5
label: add a0, a0, 4 4	lw t2, 12(a0) 5
5 lw t1, 0(a0)	max4 v0, t0, t1, t2 5
4 slt t2, t1, v0	jr \$ra 3
3 bne t2, \$zero, skip	lw v0, 0(a0) 5
move v0, t1	
skip: bne a0, a1, label 3	
3 jr \$ra	

Assume the datapath and control that you implemented in parts (a) and (b); assume **slt** and **move** can be considered as R-type instructions and **jr** and **bne** take the same amount of time as **beq**. Also assume that the first element in the input array is the largest (so that the first branch of program 1 is always taken). How much faster (or slower) is program 2 than program 1? You may leave your answer as a fraction. (10 points)

Part (c)

What is average CPI of program 1? You may leave your answer as an expression. (5 points)

Note: While we're primarily concerned about correctness, full points will only be rewarded to solutions that use a minimal number of cycles and do not lengthen the clock cycle. Assume that the ALU, Memory and Registerfile all take 2ns, and everything else is instantaneous.

Part (b)

The `max4` instruction can be used in place of three branch instructions, reducing the number of instructions that need to be executed. Below are two functionally equivalent programs; the second of which uses the `max4` instruction:

Program 1

```

5 lw v0, 0(a0)
4 add a1, a0, 12
label: add a0, a0, 4 4
5 lw t1, 0(a0)
4 slt t2, t1, v0
3 bne t2, $zero, skip
move v0, t1
skip: bne a0, a1, label 3
3 jr $ra
  
```

Program 2

```

lw t0, 4(a0) 5
lw t1, 8(a0) 5
lw t2, 12(a0) 5
max4 v0, t0, t1, t2 3
jr $ra 3
lw v0, 0(a0) 5
  
```

Assume the datapath and control that you implemented in parts (a) and (b); assume `slt` and `move` can be considered as R-type instructions and `jr` and `bne` take the same amount of time as `beq`. Also assume that the first element in the input array is the largest (so that the first branch of program 1 is always taken). How much faster (or slower) is program 2 than program 1? You may leave your answer as a fraction. (10 points)

Part (c)

What is average CPI of program 1? You may leave your answer as an expression. (5 points)

Question 3: More performance (20 points)

Part (a)

Assume the following delays for the main functional units:

Functional Unit	Time delay
Memory	5 ns
ALU	4 ns
Register File	3 ns

15

lw = Memory 5
 Register 3
 ALU 4
 Memory 5
 Register 3

Given the following instructions: lw, sw, add, beq, calculate:

- minimum time to perform each instruction
- time required on a single-cycle datapath (from q. 1)
- time required on multi-cycle datapath (from q. 2).

Write your answers in the table below. State any assumptions. (10/15 points)

Instruction	Time to perform the instruction:		
	minimum time	in single-cycle	in multi-cycle
lw	20 ns	20 ns	5 x 5 = 25 ns
sw	17 ns	20 ns	5 x 4 = 20 ns
add	15 ns	20 ns	5 x 4 = 20 ns
beq	12 ns	20 ns	5 x 3 = 15 ns

↑
 we will take the data speed


Part (b)

How much faster than a 1GHz single-cycle MIPS processor would a 3.0GHz Pentium4 x86 processor be if it achieved a CPI of 0.5 on a workload? (5 points)


Question 4: Pipelining and forwarding (20 points)

Part (a)

Show or list all of the dependencies in this program. For each dependency, indicate which instructions *and* register are involved. (8 points)



10



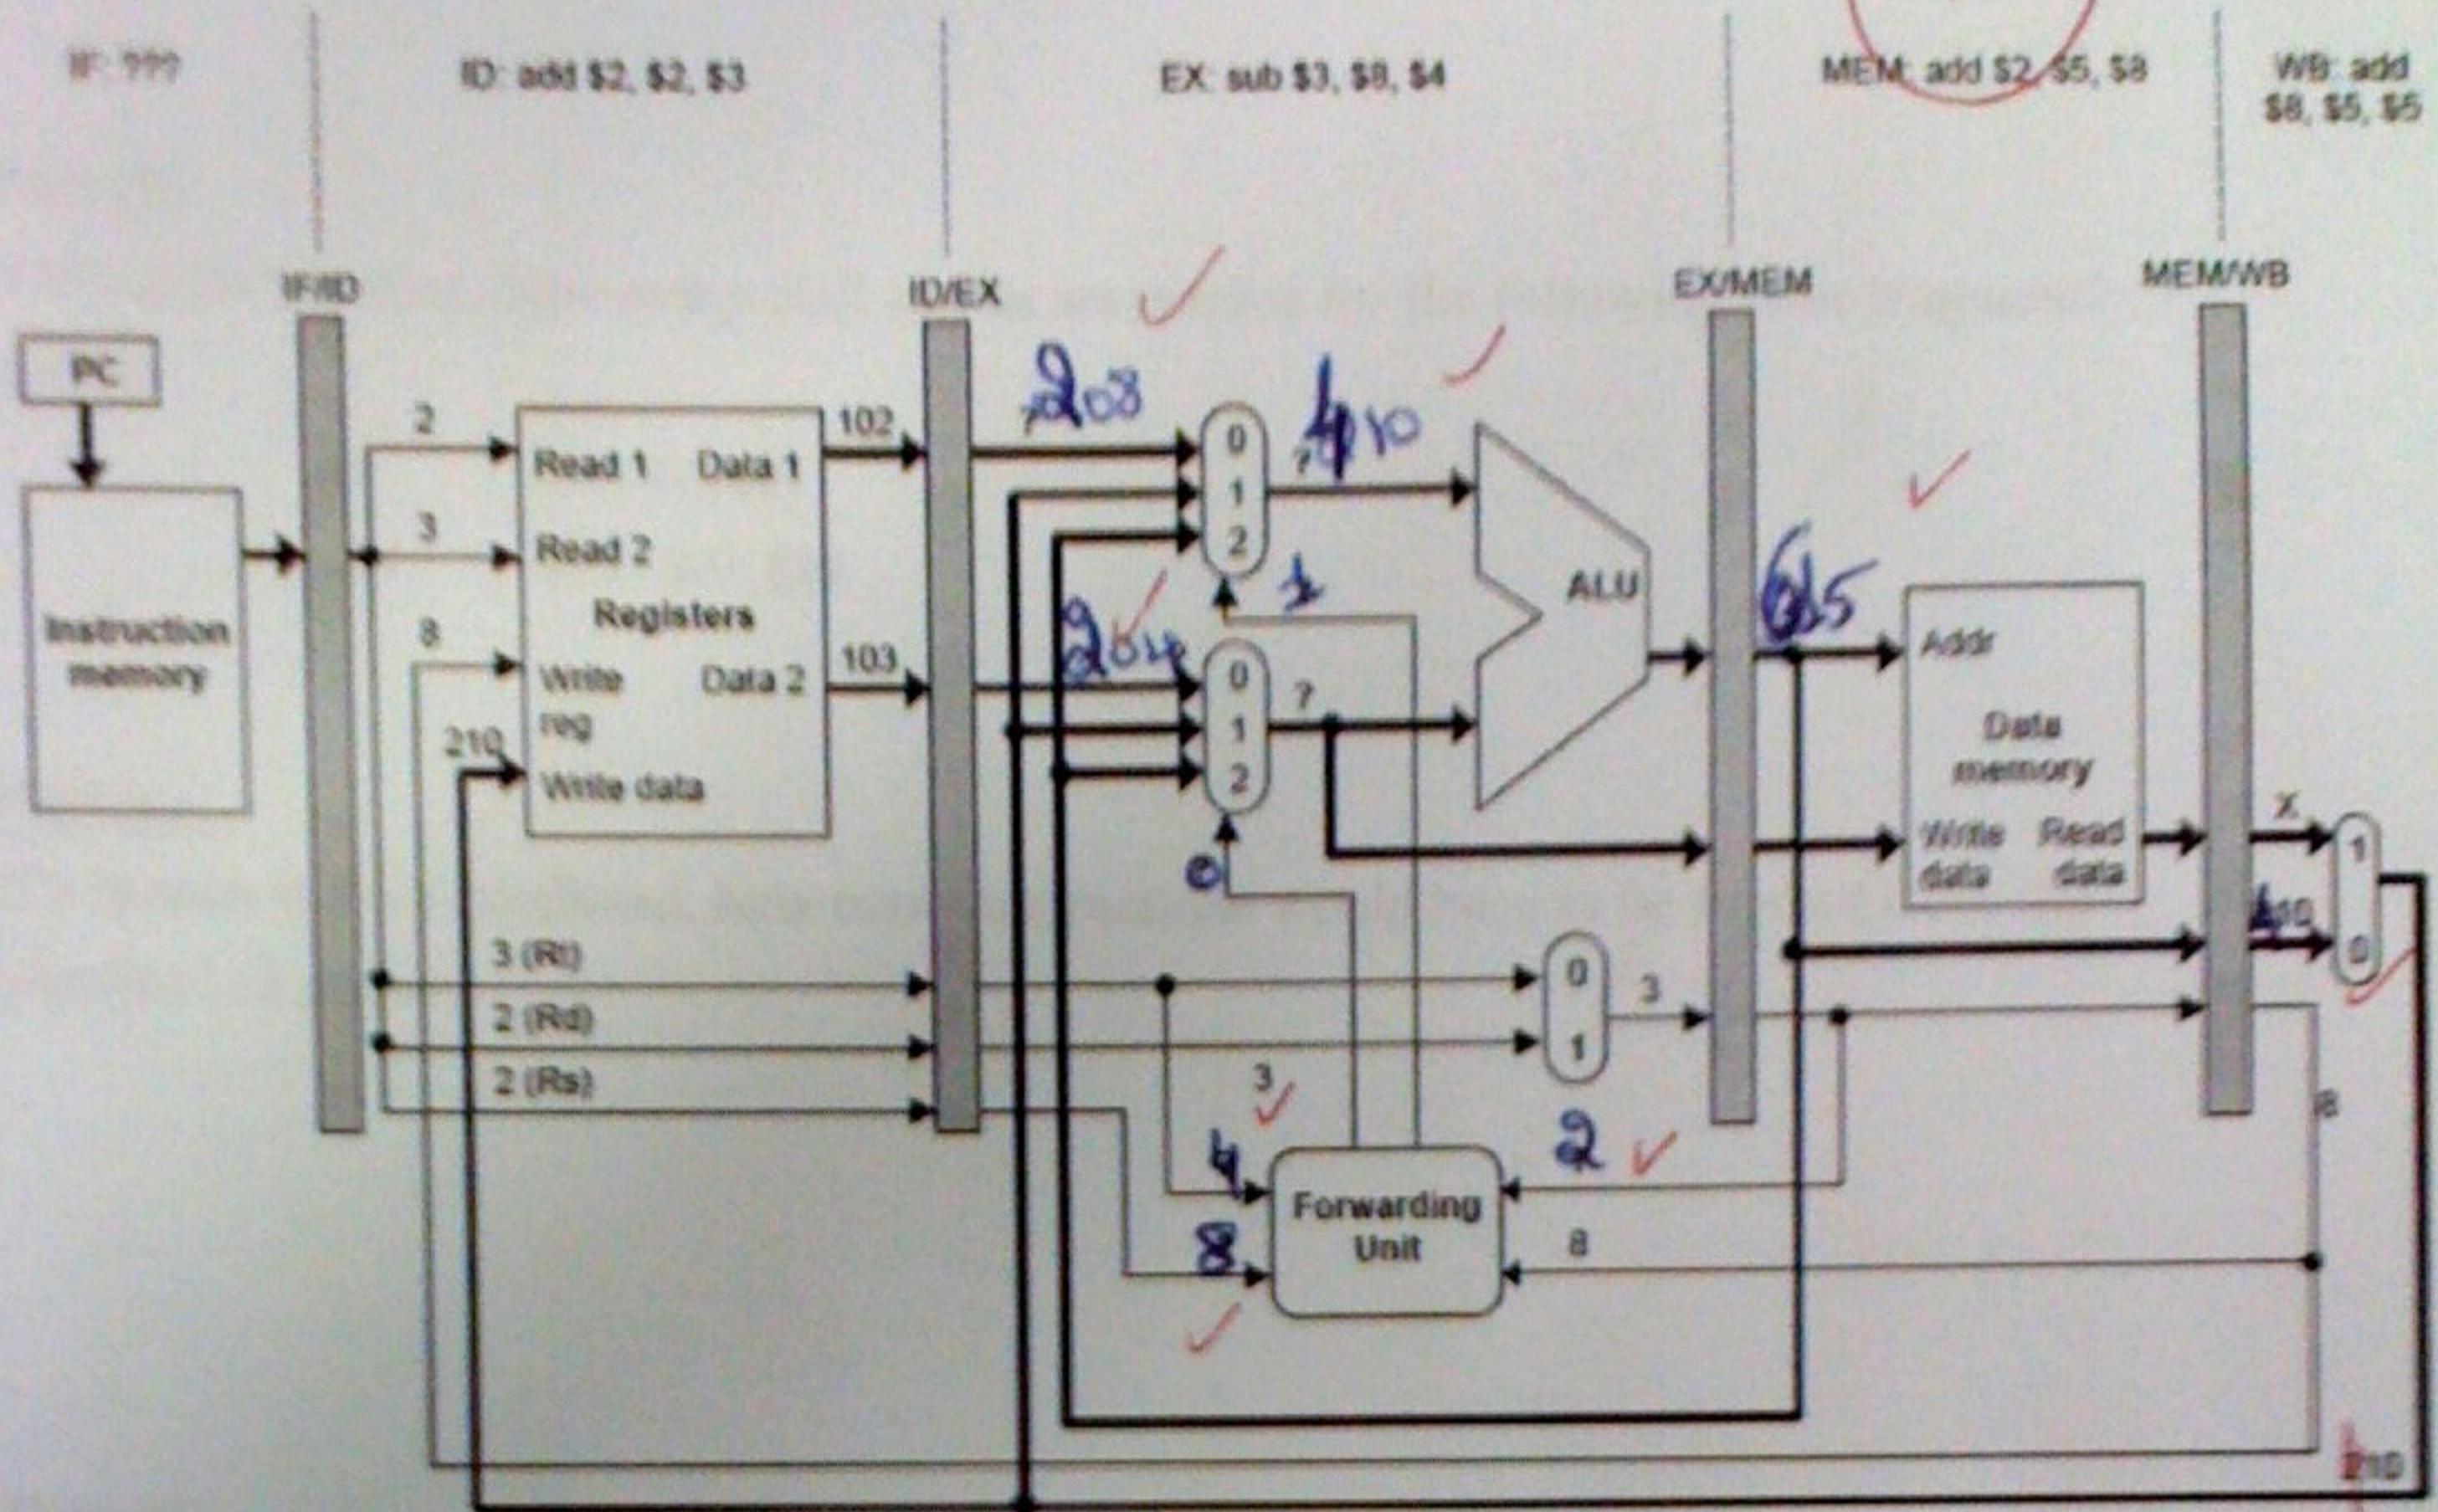
Part (b)

The pipelined datapath on the next page shows the fifth cycle of executing this program, including values for several of the stages. Fill in the ten remaining values, marked with a ? symbol, in the EX and MEM stages. (20 points; 2 points each)

N.B.:

- Write your answers directly on the diagram, but write clearly.
- Show decimal values.
- Assume that registers initially contain their number plus 200: \$2 contains 202, \$8 contains 208, etc.
- Write "X" for any numbers that cannot be determined.

10



Question 5: Pipelining performance (20 points)

IFG	FET	ROT	EXP	REN	WLD	REG	EXE	DET	WRB
1	2	3	4	5	6	7	8	9	10

One CPU manufacturer has proposed the 10-stage pipeline above for a 500MHz (2ns clock cycle) machine. Here are the correspondences between this and the MIPS pipeline:

- Instructions are fetched in the FET stage.
- Register reading is performed in the REG stage.
- ALU operations and memory accesses are both done in the EXE stage.
- Branches are resolved in the DET stage.
- WRB is the writeback stage.

Part (a)

How much time is required to execute one million instructions on this processor, assuming there are no dependencies or branches in the code? (5 points)

Part (b)

Without forwarding, how many stall cycles are needed for the following code fragment? (5 points)

```
lw    St0, 0($a0)
add   Sv1, St0, St0
```

Part (c)

If a branch is miss-predicted, how many instructions would have to be flushed from the pipeline? (5 points)

Part (d)

Assume that a program executes one million instructions. Of these, 15% are load instructions which stall, and 10% of the instructions are branches. The CPU predicts branches correctly 75% of the time. How much time will it take to execute this program? (5 points)

Performance

1. Formula for computing the CPU time of a program P running on a machine X:

$$\text{CPU time}_{XP} = \text{Number of instructions executed}_P \times \text{CPI}_{XP} \times \text{Clock cycle time}_X$$

2. CPI is the average number of clock cycles per instruction:

$$\text{CPI} = \text{Number of cycles needed} / \text{Number of instructions executed}$$

3. Speedup is a metric for relative performance of 2 executions:

$$\begin{aligned} \text{Speedup} &= \text{Performance after improvement} / \text{Performance before improvement} \\ &= \text{Execution time before improvement} / \text{Execution time after improvement} \end{aligned}$$