

# Computer Organization

## Final Exam

1. The following figure is the complete data path for the pipelined MIPS processor with necessary supports for forwarding and hazarding detection. Assume the pipeline is **optimized for branches that are not taken** and that we moved the branch execution to the ID stage. Your task is to determine the correct values of the signals A, B, C, and D under the following conditions.

(1) (5%) At the clock cycle 5 when executing the following program. The initial conditions are: \$1=0, \$2=5, \$3=4, \$4=7, \$5=3, \$6=8, \$7=3, \$8=0, \$9=9, Mem[20]=5.

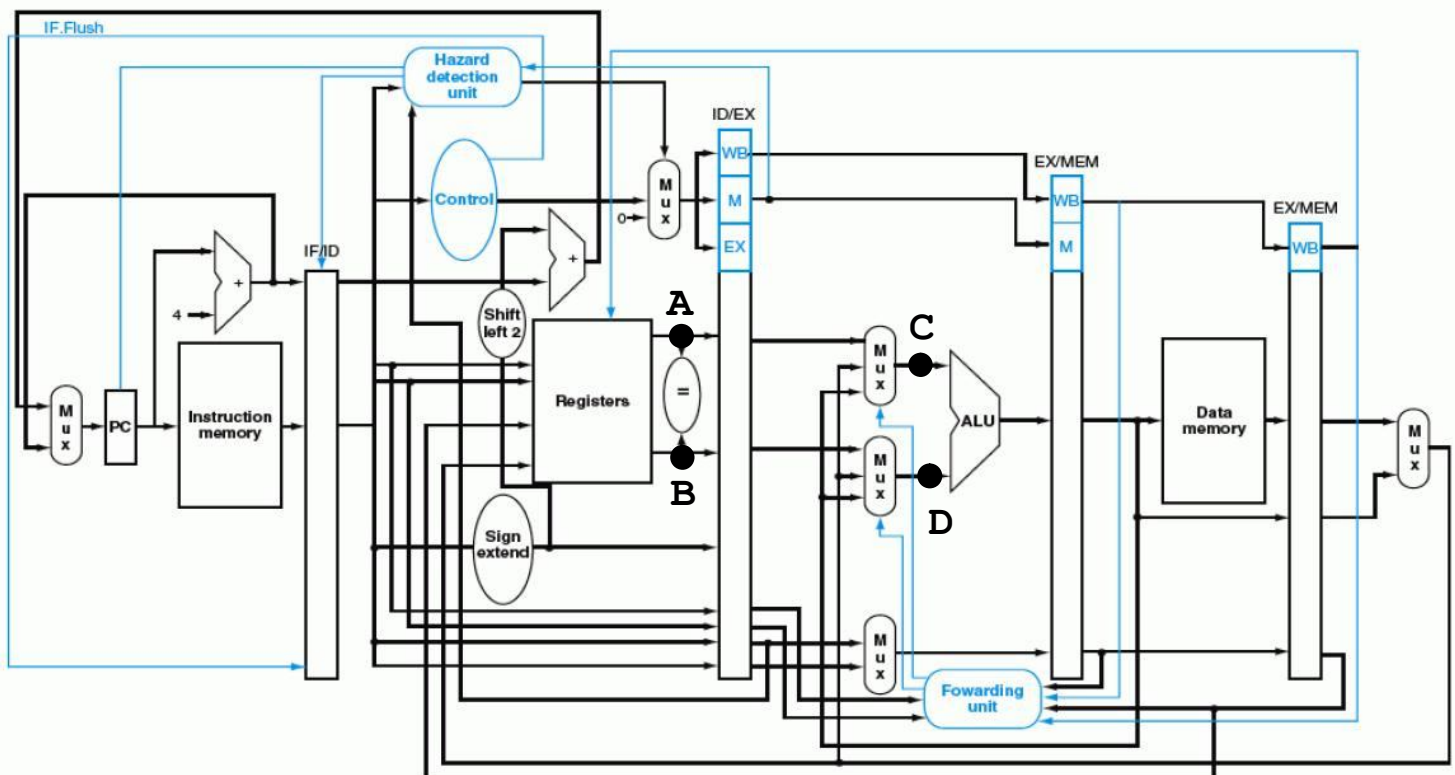
```

lw    $10, 20($1)
sub   $11, $2, $3
and   $12, $4, $5
or    $13, $6, $7
add   $14, $8, $9
    
```

(2) (5%) At the clock cycle 3 when executing the following program. The initial conditions are: \$1=3, \$2=5, \$3=3, \$4=0, \$5=3, \$6=8, \$7=0, \$8=5, Mem[50]=13.

```

36   sub   $10, $4, $8
40   beq   $1, $3, 7    #PC-relative branch to 40 + 4 + 7*4 = 72
44   and   $12, $2, $5
48   or    $13, $2, $6
52   add   $14, $4, $2
56   slt   $15, $6, $7
.....
72   lw    $4, 50($7)
    
```



2. The individual stages of the data path have the following latencies:

IF	ID	EX	MEM	WB
200 ps	150 ps	120 ps	190 ps	140 ps

(1) (4%) What is the clock cycle time in a pipelined and non-pipelined processor?

(2) (6%) If we can split one stage of the pipelined data path into two new stage, each with half the latency of the original stage, which stage would you split and what is the new clock cycle time of the processor?

3. The following problems refer to the following sequence of instructions:

<b>lw</b>	<b>\$5, -16(\$5)</b>
<b>sw</b>	<b>\$5, -16(\$5)</b>
<b>add</b>	<b>\$5, \$5, \$5</b>

(1) (5%) Indicate dependences and their type.

(2) (5%) Assume there is no forwarding in this pipelined processor. Indicate hazards and add "nop" instructions to eliminate them.

4. (10%) Assume the following MIPS code is executed on a pipelined processor with a five stage pipeline, full forwarding, and a predict-taken branch predictor:

	add	\$1, \$5, \$3
Label1:	sw	\$1, 0(\$2)
	add	\$2, \$2, \$3
	beq	\$2, \$4, Label1 ; Not taken
	add	\$5, \$5, \$1
	sw	\$1, 0(\$2)

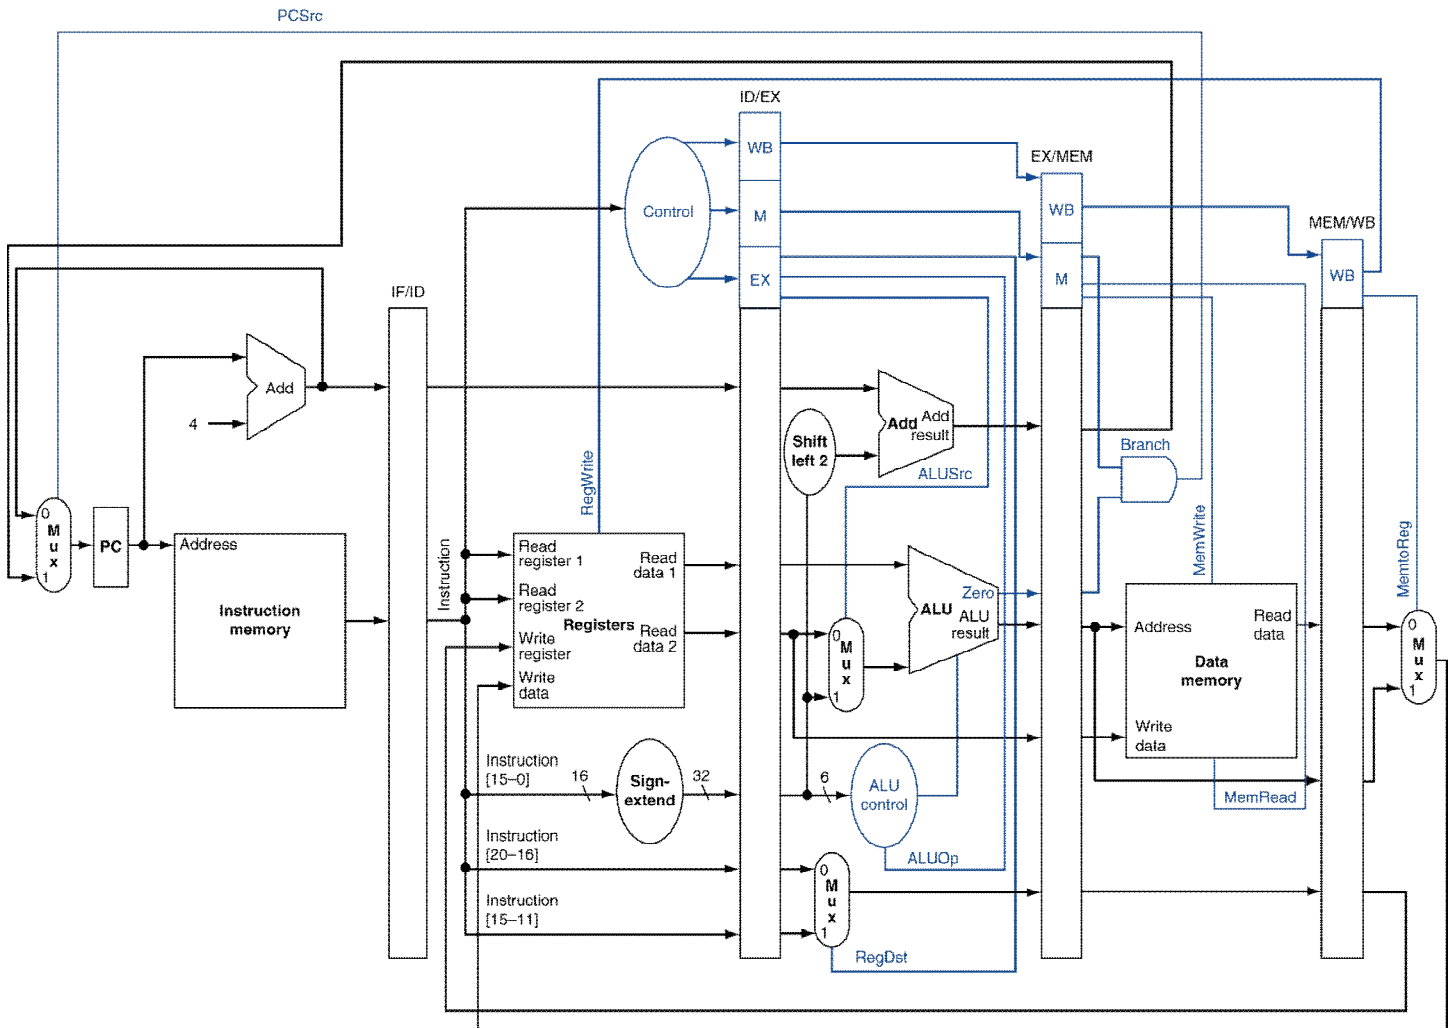
Draw the pipeline execution diagram for this code, assuming there are no delay slots and that branches execute in the EX stage.

5. Please refer to the execution of the instruction "slt \$2, \$1, \$3" in the pipelined data path from the following figure (in the next page).

(1) (7%) For each stage of the pipeline, what are the values of control signals asserted by this instruction in that pipeline stage?

EX	MEM	WB
ALUSrc = _____	Branch = _____	MemtoReg = _____
ALUOp = 10	MemWrite = _____	RegWrite = _____
RegDst = _____	MemRead = _____	

(2) (3%) What is the value of the PCSrc signal for this instruction?



6. For a direct-mapped cache design with 32-bit address, the following bits of the address are used to access the cache.

Tag	Index	Offset
31-12	11-5	4-0

- (1) (2%) What is the cache line size (in words)?
- (2) (2%) How many entries does the cache have?
- (3) (6%) For the following “**byte-addressed**” cache references, how many blocks are replaced?

Address	0	4	16	132	232	160	1024	30	140	3100	180	2180
---------	---	---	----	-----	-----	-----	------	----	-----	------	-----	------

7. (10%) There are three small caches, each consisting of four one-word blocks. One cache is fully associative, a second is two-way set associative, and the third is direct-mapped. Find the numbers of misses for each organization given the following sequence of block addresses: 0, 8, 0, 6, 8. Assume that the least recently used replacement policy for the set-associative cache.

8. Below is a list of 32-bit memory address reference, given as “**byte addresses**”.

Address	24	856	700	856	24	336	260	696	256	420	340	860
---------	----	-----	-----	-----	----	-----	-----	-----	-----	-----	-----	-----

- (1) (6%) Given a direct-mapped cache with 16 one-word blocks. List if each reference is a hit or a miss,

assuming the cache is initially empty.

- (2) (8%) There are three direct-mapped cache designs possible, all with a total of eight words of data: C1 has one-word blocks, C2 has two-word blocks, and C3 has four-word blocks. In terms of miss rate, which cache design is the best? If the miss stall time is 25 cycles, and C1 has an access time of 2 cycles, C2 takes 3 cycles, and C3 takes 5 cycles, which is the best cache design?

9. Media applications that play audio or video files are part of a class of workloads called “streaming” workloads; i.e., they bring in large amounts of data but do not reuse much of it. Consider a video streaming workload that accesses a 512 KB working set sequentially with the following address stream:

Address	0	4	8	12	16	20	24	28	32...
---------	---	---	---	----	----	----	----	----	-------

- (1) (6%) Assume a 64 KB direct-mapped cache with a 32-byte line. What is the miss rate for the address stream above? How is this miss rate sensitive to the size of the cache or the working set? How would you categorize the misses this workload is experiencing, based on the 3C model.

- (2) (6%) Cache block size (B) can affect both miss rate and miss latency. Assuming the following miss rate table, assuming a 1-CPI machine with an average of 1.35 references (both instruction and data) per instruction. Given the following miss rates for various block sizes, what is the optimal block size for a **miss latency of  $20 \times B$  cycles**?

8 Bytes	16 Bytes	32 Bytes	64 Bytes	128 Bytes
8%	3%	1.8%	1.5%	2%

10. (5%) Consider a virtual memory system with the following properties:

- A. 40-bit virtual byte addresses
- B. 16-KB pages
- C. 36-bit physical byte addresses

What is the total size of the page table for each process on this machine? Assume that the **valid, protection, dirty, and use bits** take a total of 4 bits and that all the virtual pages are in use.