

[1] What are the two main ways to define performance?

[1] The Operating System needs to know what two things in order to deal with an interrupt?

[1] 20 years ago, dynamic power consumption was more than an order of magnitude larger than leakage power. **Approximately** what percentage of power consumption is due to leakage now?

[1] Predicting the direction of a branch is not enough. What else is necessary?

[1] Advances in technology allow us to make both transistors and wires smaller and smaller. If everything else is held constant, what happens to the power density as things shrink?

[2] When we talk about the number of operands in an instruction (a 1-operand or a 2-operand instruction, for example), what do we mean?

[2] Artificial intelligence techniques can provide more accurate branch prediction than existing approaches - why are these techniques not used?

[2] Over the years, clock rates have grown by a factor of 1000 while power consumed only increased by a factor of 30. How was this accomplished? Why is this technique no longer effective?

[2] Do benchmark programs remain valid indefinitely? Why or why not?

[2] The multicycle MIPS design uses a single memory, but the pipelined MIPS design has two. Why is that?

[3] It is difficult for the internal processing elements on a CMOS chip to cross the chip boundary and communicate with things that are on other chips. Explain why that is.

[4] You are in charge of designing a new embedded machine. You would like to support 4096 different operations, 64 registers, and use a 2-operand instruction format.

a) Draw what the instruction would look like.

b) Assume you discover you must use a fixed 20 bit instruction size. Will this instruction format work? If not, explain why and give two different ways to fix the problem.

[7] An important program spends 50% of its time doing Integer arithmetic operations, 30% of its time doing loads and stores, and 20% of its time doing floating point arithmetic. By redesigning the hardware you can make the Floating Point unit 90% faster (take 10% as long), the load/store unit 70% faster (take 30% as long), or the integer unit 40% faster (take 60% as long). Which should you do, and why?

[4] In your first design of a 5-stage pipeline (F,D,E,M,W), F takes 26 time units, D takes 27, E takes 50, M takes 52, and W takes 51.

a) What will the clock cycle time be for this pipeline?

b) Is it a balanced pipeline? If not, explain what you could do to fix it. What would the cycle time be now?

[8] Here is a code sequence.

load R5, 0(R13)

or R1, R5, R4

add R1, R4, R2

sub R2, R4, R6

store R2, 5(R6)

Indicate all dependencies (draw lines/arrows between them, and write beside each line/arrow which hazard is involved).

[9] You are given a machine with an 8 stage pipeline, which writes to the register file during the D1 cycle and reads during D2. This machine uses neither a branch delay slot nor a load delay slot. These are the stages:

F D1 D2 E1 E2 M1 M2 W

a. Assuming this machine has a branch predictor and the branch condition is calculated by the end of the D1 stage, how big is the branch penalty (measured in cycles) when the prediction is incorrect? What if the branch condition is calculated by the end of E1?

b. Assume on a load the value is available at the end of M2. How many stall cycles stalls would this machine need on a load if it has forwarding logic and you are forwarding to E1? What if you were forwarding to E2?

c) What type of data hazard does the above pipeline need to worry about?

d) If the above pipeline were modified to support out of order completion, what new data hazard would be introduced?

e) If in addition to completing out of order, instructions were allowed to issue out of order, what new data hazard would be introduced?

[15] Here is a code sequence.

load R5, 8(R9)

load R2, 4(R6)

add R3, R1, R2

sub R1, R7, R6

load R4, 0(R1)

and R4, R9, R10

Assuming a standard 5-stage pipeline that does not support hazard detection and does no forwarding,

- a)** Indicate all dependencies (draw lines/arrows between them, and write beside each line/arrow which hazard is involved).
- b)** Insert as many No Operations (NOPS) as required in order to ensure this code runs correctly. (Remember, writes to the register file occur on the first half of the cycle, and reads occur during the second half).
- c)** Circle the NOPs that can be removed if forwarding and hazard detection logic is implemented.
- d)** Assuming the pipeline has been modified so that it does support hazard detection and forwarding, schedule the code to remove as many stalls as possible. How many NOPS are left?
- e)** If R9 contains the value 0x12, what address in memory is accessed when executing the first load instruction?

In this question, we are going to wire up a 12-bit processor. The machine is word-addressable, where a word is 12 bits. immediates are sign extended, Offset is not. The machine has 3 different instruction formats: R, I, and J. Memory takes a single cycle to return a value.

R-type:	<i>(Arithmetic and logical: $rd = rs1 \text{ OP } rs2$)</i>			
Opcode	rd	rs1	rs2	funct
11-8	7-6	5-4	3-2	1-0
I-type:	<i>(Arithmetic and logical: $rd = rs1 \text{ OP } Immediate$)</i>			
	<i>(Load: $rd = mem[rs1+Immediate]$)</i>			
	<i>(Store: $mem[rs1+Immediate]=rd$)</i>			
	<i>(Branch: $if (rd \text{ OP } rs1)=COND, PC=PC+Immediate$)</i>			
Opcode	rd	rs1	Immediate	
11-8	7-6	5-4	3-0	
J-type:	<i>(Jump: $PC = Offset$)</i>			
Opcode	Offset			
11-8	7-0			

The ALU can perform 7 functions, written this way: OP [ALU2 ALU1 ALU0]
 Decrement X [111], Increment X [110], Add [011], SUB [100], AND [101], OR [001], NOT [010]

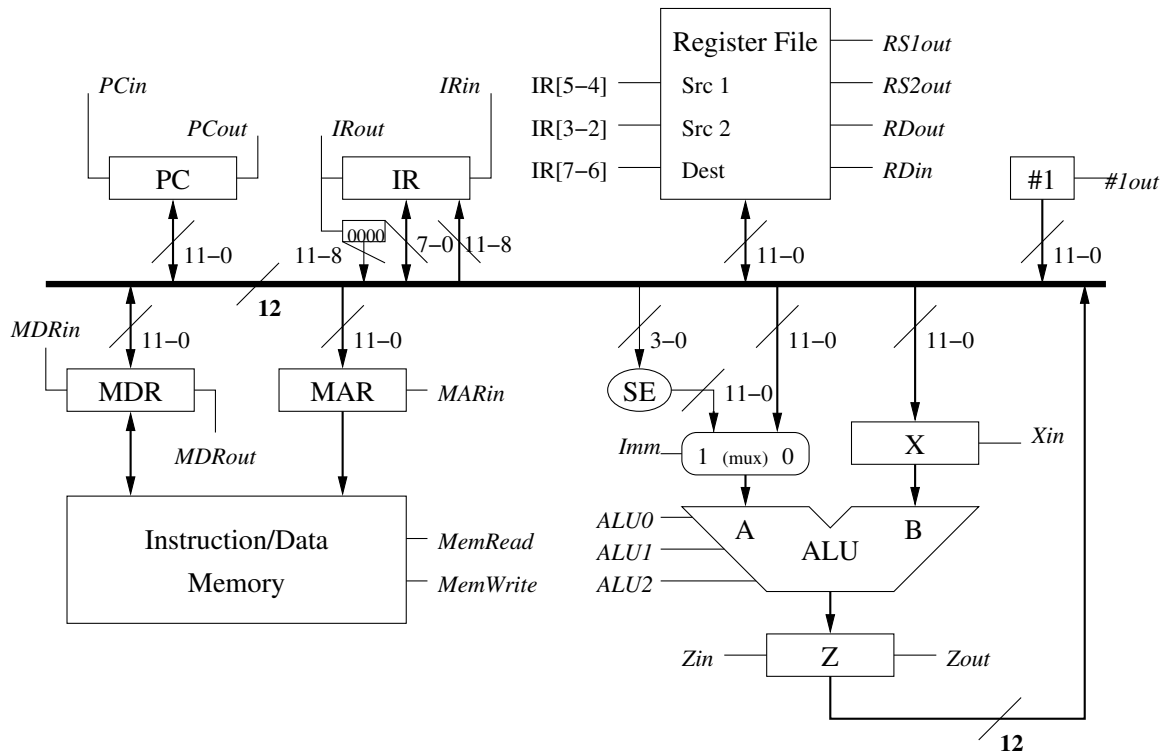
Here are some of the instructions that have been defined:

Name	Opcode(Funct)	Name	Opcode(Funct)	Name	Opcode(Funct)
NOP	0000(00)	lw	0001(xx)	sw	0011(xx)
NOT	1000(00)	BEQZ	0100(xx)	j	0101(xx)
AND	1000(01)	AND Imm	1001(xx)	ADD	1100(01)
OR	1000(10)	OR Imm	1010(xx)	ADD Imm	1101(xx)
XOR	1000(11)	XOR Imm	1011(xx)	SUB	1100(01)

Here are the 21 control signals.

PCin	PCout	IRin	IRout	MDRin	MDRout	MARin
Zin	Zout	RDin	RDout	MemRead	MmWrite	Imm
RS1out	RS2out	#1out	ALU0	ALU1	ALU2	Xin

Here is a diagram of the machine.



[11] Fill in the microcode steps necessary to perform an instruction fetch (incrementing the PC is considered part of the instruction fetch process).

Step	PCin	IRin	MARin	MDRin	MemRead	MemWrite	Imm	SE	Xin	Zin	ALU0	ALU1	ALU2	#1out	ALU2	ALU1	ALU0	MemRead	MemWrite	Imm
0																				
1																				
2																				
3																				
4																				
5																				


[8] Now that you have done the instruction fetch, fill in the microcode steps necessary to perform the following instruction: **AND R1,R2,-1** Assume memory can respond in a single cycle.

S t e p	P C i n n	I R i n n	M A R i n n	M D R i n n	R D i n n	X i n n	Z i n n	I R o u t	P C o u t	M D R o u t	Z o u t	R D o u t	R S 1 o u t	R S 2 o u t	# 1 o u t	A L U 2	A L U 1	A L U 0	M r e a d	M w r i t e	I m m
0																					
1																					
2																					
3																					
4																					

[2] Assuming you have done the instruction fetch, fill in the microcode steps necessary to perform the following instruction: **J 0x23**

S t e p	P C i n n	I R i n n	M A R i n n	M D R i n n	R D i n n	X i n n	Z i n n	I R o u t	P C o u t	M D R o u t	Z o u t	R D o u t	R S 1 o u t	R S 2 o u t	# 1 o u t	A L U 2	A L U 1	A L U 0	M r e a d	M w r i t e	I m m
0																					
1																					
2																					
3																					

Here is the state diagram for a "random" machine:



[3] Assuming there are 3 state variables (Y2-Y0), that State0 = $\overline{Y2} * \overline{Y1} * \overline{Y0}$ (000) and State7 = $Y2 * Y1 * Y0$ (111), write down the exact boolean equation for the **RegWrite** signal.

[3] Assuming the same situation as in the previous question, write down the exact boolean equation for **NextState2**.

[2] If you were using a minimized microcode configuration for this machine, circle on the diagram where the dispatch rom points are.

[6] In class, we talked about the cycle by cycle steps that occur on different interrupts. For example, here is what happens if there is an illegal operand interrupt generated by instruction i (note this machine has a 6 stage pipeline):

	1	2	3	4	5	6	7	8	9	
i	IF	ID	EX	M1	M2	WB	<- Interrupt detected			
i+1		IF	ID	EX	M1	M2	WB	<- Instruction Squashed		
i+2			IF	ID	EX	M1	M2	WB	<- Trap Handler fetched	
i+3				IF	ID	EX	M1	M2	WB	
i+4					IF	ID	EX	M1	M2	WB

Fill out the following table if instruction i+1 experiences a fault in the M1 stage (page fault, for example) and we are assuming **imprecise** interrupts:

	1	2	3	4	5	6	7	8	9	10	
i	IF	ID	EX	M1	M2	WB					
i+1		IF	ID	EX	M1	M2	WB				
i+2			IF	ID	EX	M1	M2	WB			
i+3				IF	ID	EX	M1	M2	WB		
i+4					IF	ID	EX	M1	M2	WB	
i+5						IF	ID	EX	M1	M2	WB

Assuming we are supporting **precise** interrupts, what happens in this case?

	1	2	3	4	5	6	7	8	9	10			
i	IF	ID	EX	M1	M2	WB							
i+1		IF	ID	EX	M1	M2	WB	<- M1 stage has Page Fault					
i+2			IF	ID	EX	M1	M2	WB	<- Illegal Instruction detected				
i+3				IF	ID	EX	M1	M2	WB	<- IF stage has page fault			
i+4					IF	ID	EX	M1	M2	WB			
i+5						IF	ID	EX	M1	M2	WB		
i+6							IF	ID	EX	M1	M2	WB	
i+7								IF	ID	EX	M1	M2	WB