

Department of Electrical and Computer Engineering  
The University of Texas at Austin

EE 460N Spring 2023  
Instructor: Yale N. Patt  
TAs: Michael Chen, Ali Mansoorshahi  
Exam 1  
March 1, 2023

Name: \_\_\_\_\_

Problem 1 (25 points): \_\_\_\_\_

Problem 2 (15 points): \_\_\_\_\_

Problem 3 (30 points): \_\_\_\_\_

Problem 4 (30 points): \_\_\_\_\_

Total (100 points): \_\_\_\_\_

Note: Please be sure that your answers to all questions (and all supporting work that is required) are contained in the space provided.

Note: Please be sure your name is recorded on each sheet of the exam.

Please read the following sentence, and if you agree, sign where requested:  
I have not given nor received any unauthorized help on this exam.

Signature: \_\_\_\_\_

**GOOD LUCK!**

Name: \_\_\_\_\_

**Problem 1 (25 points):** Answer the following questions.

Note: For each of the five answers below, if you leave the box **empty**, you will receive one point of the five.

**Part a (5 points):** A program is running on the LC-3b when an exception occurs during the execution of the STW instruction shown below.

```
ADD, R1, R2, R3
STW R1, A
BRz NEXT
```

If there was no exception, the BRz instruction would use the Z bit set by the ADD instruction since STW does not set condition codes.

However, because an exception occurs, an exception service routine is executed, then the STW instruction, and then the BRz instruction. When the exception service routine finishes execution the condition codes contain the values set by the last instruction in the service routine that sets condition codes.

How do we make sure that BRz uses the Z bit set by the ADD and not the Z bit set by some instruction in the service routine?

Save the CC on the stack when switching context

**Part b (5 points):** In class I told you the role of the architect includes predicting the rate of improvement of technology. What is the problem with being too conservative in the architect's prediction?

the product will underperform

**Part c (5 points):** A microarchitect notes that for a specified cycle time, a certain function will need two clock cycles to complete, but if the cycle time is increased by 10%, the function will only need one clock cycle. Is it ever a good idea to increase the cycle time by 10% in order to perform that function in one clock cycle? If yes, when?

if the benefit of that instruction taking 1 cycle is more than everything being 10% slower.

Name: \_\_\_\_\_

**Part d (5 points):** Which of the following are not part of the ISA: PC, condition codes, MAR, memory addressability, PSR, pipeline, data types

MAR, Pipeline

**Part e (5 points):** A Load/Store ISA has a restriction on what a single instruction can do. What is that restriction?

it can only operate on registers or load/store from memory. not both

Name: \_\_\_\_\_

**Problem 2 (15 points):** Memory location x4180 contains a branch instruction (call it A) that is always taken when the Branch History Register contains 110101. Memory location x4B00 contains a branch instruction (call it B) that is always not taken when the Branch History Register contains 110101.

If we use my original two level predictor, what is the address of the 2-bit counter in the Pattern History Table that gets updated by the two branch instructions?

x35

Unfortunately, Branches A and B interfere with each other, producing many mispredictions. Gshare can solve the problem by having the two branches use different 2-bit counters when the Branch History Register is 110101. This can be done by XORing the history with bits 14, 12, 11, 9, 8, and 7 of the PC to obtain the address of the 2-bit counter to be used for predicting the direction of branches A and B. (Note: Bit 14 is XORed with the left-most bit in the Branch History Register.)

What is the address of the counter that A uses with Gshare?

x16


What is the address of the counter that B uses with Gshare?

x1B

Name: \_\_\_\_\_

**Problem 3 (30 points):** Our customer wants us to add a SEARCH instruction to the LC-3b which will search a region of memory for a specified value.

The instruction format is shown below. We will use the unused opcode 1010.

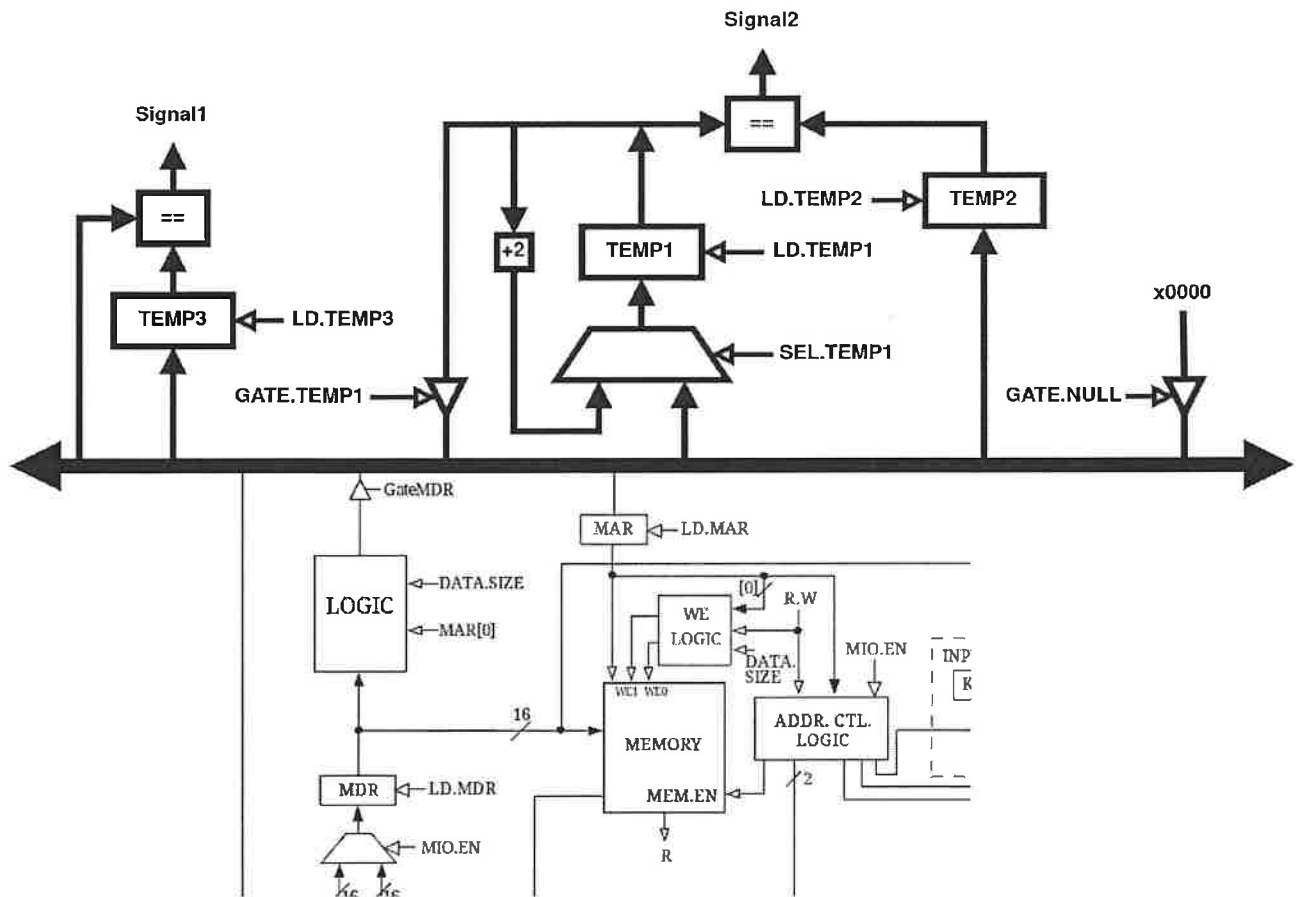


BR (Base Register) contains the starting address of the memory region to be searched. imm6 is a positive 2's complement integer containing the number of words in the region. SR (Source Register) contains the value being searched.

SEARCH examines the locations of the region, starting with the address in BR. If SEARCH finds the value, it loads the address of that location into BR. If SEARCH does not find the value in the region, it loads x0000 into BR. Assume there can be only one instance of the value in the region being searched.

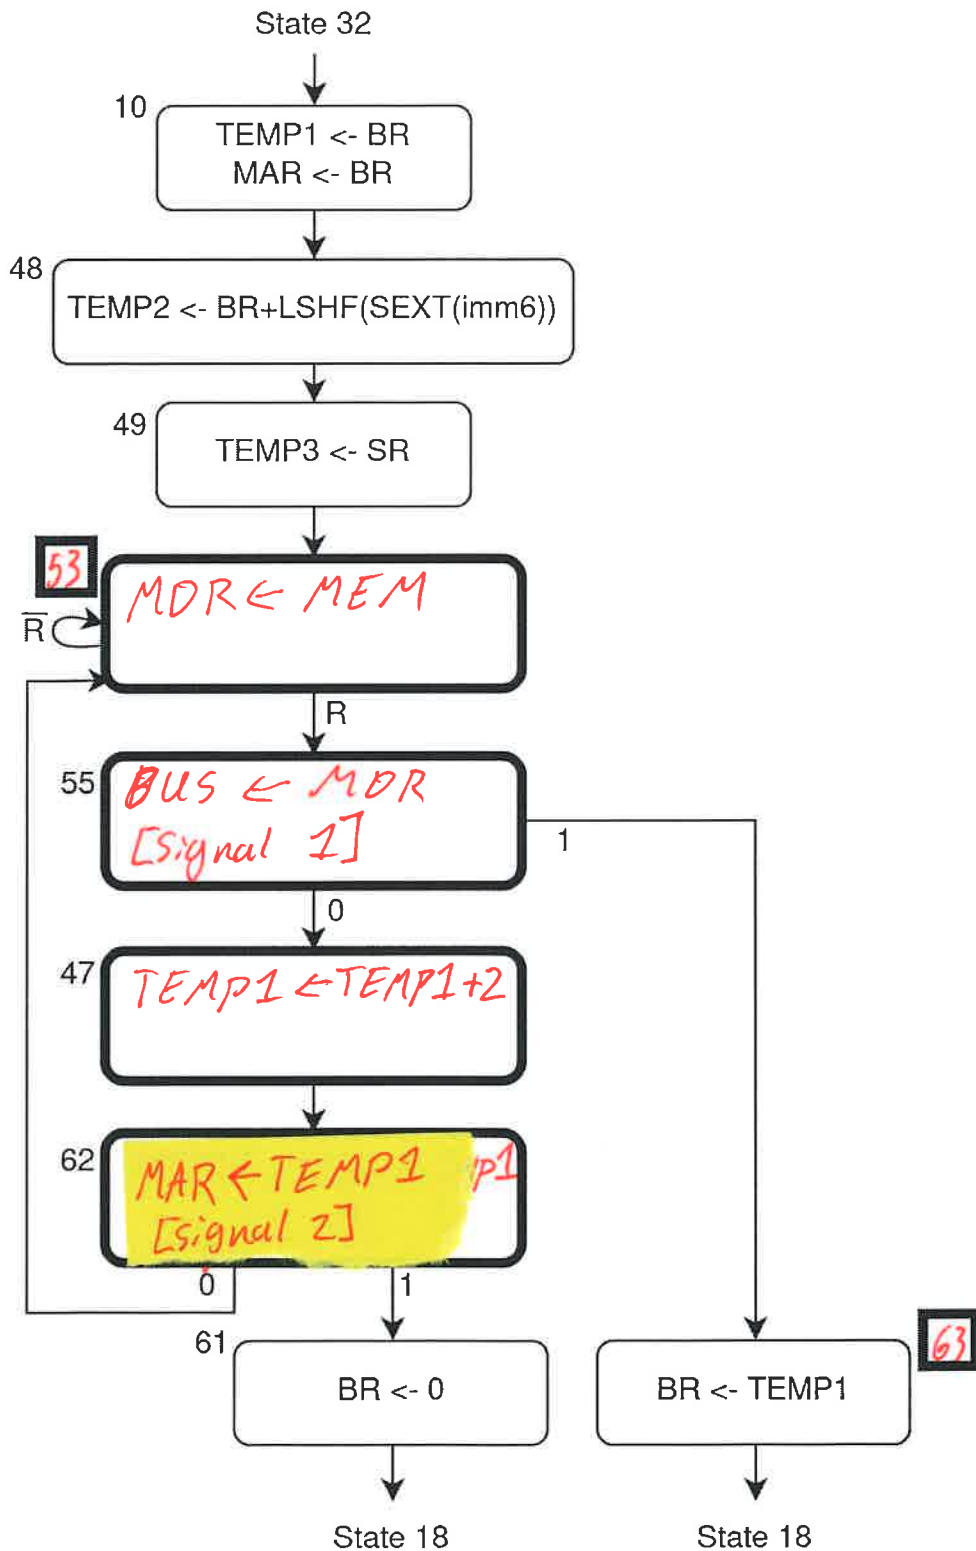
For example, if BR is 1, imm6 is 3, SR is 4, R1 = 0x4000, and R4 = 0xBEEF, if  $M[x4004] = 0xBEEF$ , SEARCH will load x4004 into R1. If no memory location in the region contains 0xBEEF, SEARCH will load x0000 into R1.

Your job: Augment the LC-3b to implement SEARCH. This will require the additions to the datapath shown in bold-face below.



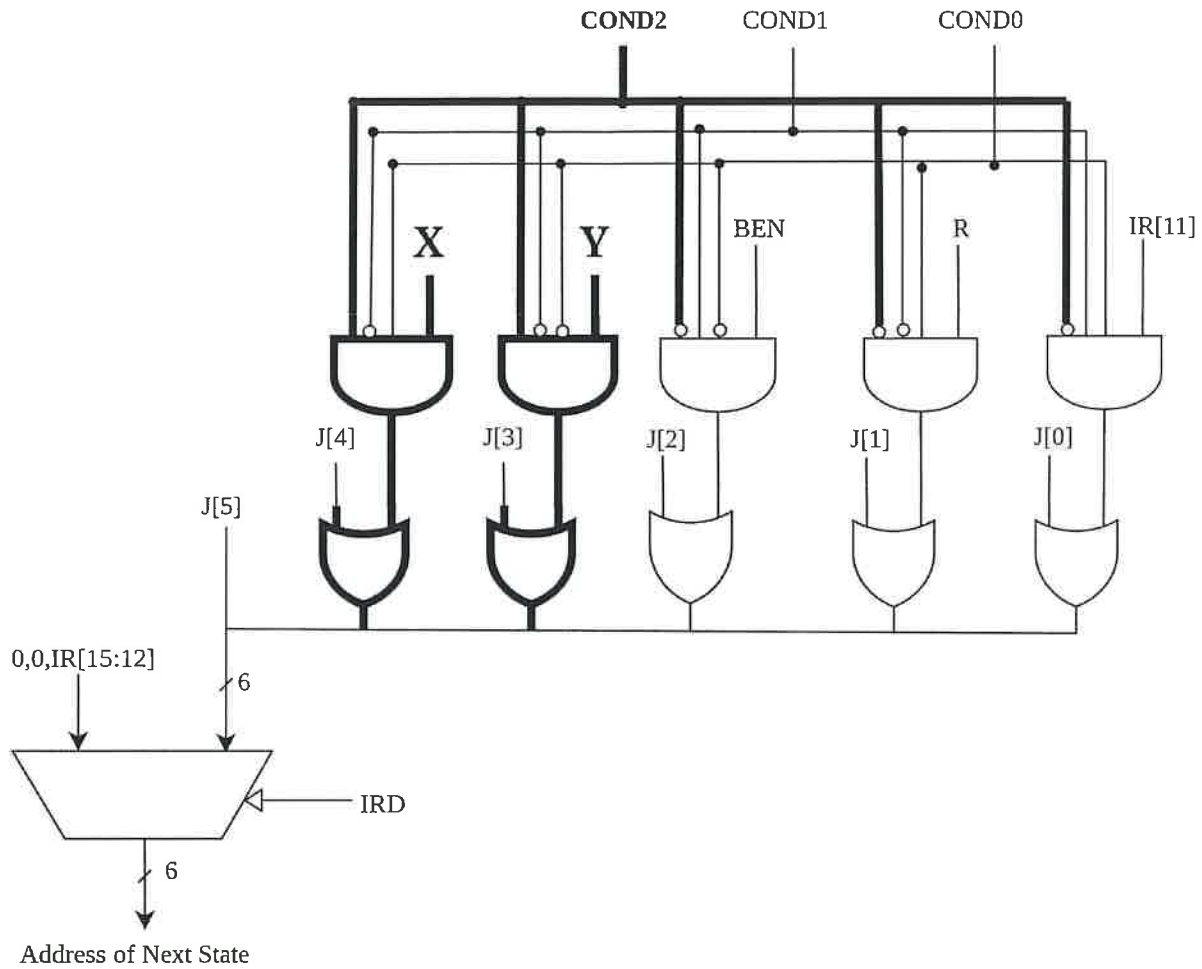
Name: \_\_\_\_\_

**Part a (10 points) The State Machine:** Complete the state machine shown below, i.e. fill in the missing information in the four states and the two state numbers shown in boldface.



Name: \_\_\_\_\_

**Part b (4 points):** The modified micro-sequencer is shown below. The additions are shown in bold.



What is X? Be precise.

*Signal 1*

What is Y? Be precise.

*Signal 2*

Name: \_\_\_\_\_

**Part c (7 points):** In the table below, each row represents a state, each column represents a control signal, and each entry represents the value of that control signal for that state. The entry "x" means don't care. The values of the control signals for states 48, 53, 55, and an unidentified state are shown. Your job is to determine what control signal is specified by each column, and the state number of the unidentified state shown in boldface.

The possible control signals to choose from are: COND, J, GATE.MARMUX, GATE.MDR, GATE.PC, MIO.EN, LD.MAR, and GATE.ALU.

State	<i>COND</i>	<i>J</i>	<i>Gate.MARMUX</i>	<i>MIO.EN</i>	<i>Gate.MDR</i>
48	0	49	1	0	0
53	1	53	0	1	0
55	5	47	0	0	1
<b>62</b>	4	53	0	0	0

**Part d (9 points) Analysis:** We could have performed the function of the new SEARCH instruction in software, rather than add the SEARCH instruction to the ISA. Which choice provides the better performance? Explain in 20 words or fewer.

*instruction due to only fetching 1 time.*

Without further modification to the data path, could we have combined states 10, 48, and 49 into a single state? Yes/No. Explain.

*NO. ALL USE BUS*



Name: \_\_\_\_\_

**Problem 4 (30 points):** A microarchitecture capable of out-of-order execution is tasked with executing a program consisting of six instructions.

The microarchitecture has the following characteristics:

- The instruction cycle consists of five stages: FETCH, DECODE, REGISTER RENAME, EXECUTE, WRITE-BACK.
- ADD requires two clock cycles to EXECUTE, MUL takes three, all else take one each.
- There is one of each functional unit.
- Reservation Stations are allocated and filled in at the end of REGISTER RENAME.
- Instructions with no dependencies can begin EXECUTE directly after REGISTER RENAME.
- Data forwarding is in use. Results are broadcast in the last stage of EXECUTE, and dependent instructions can begin EXECUTE in the next cycle.
- Reservation Stations are deallocated at the end of WRITEBACK, and can be allocated for another instruction in the next clock cycle.

As we showed in class last week, a program can be represented as a data flow graph, with each instruction represented as a node, the sources of that instruction shown as input arcs to the node, the result of each instruction shown as an output arc, and all dependencies shown as arcs from the node producing a result to all the nodes that need that result as a source.

For example, the following data flow graph represents the two instruction sequence

MUL R1,R2,R3  
ADD R4,R3,R1


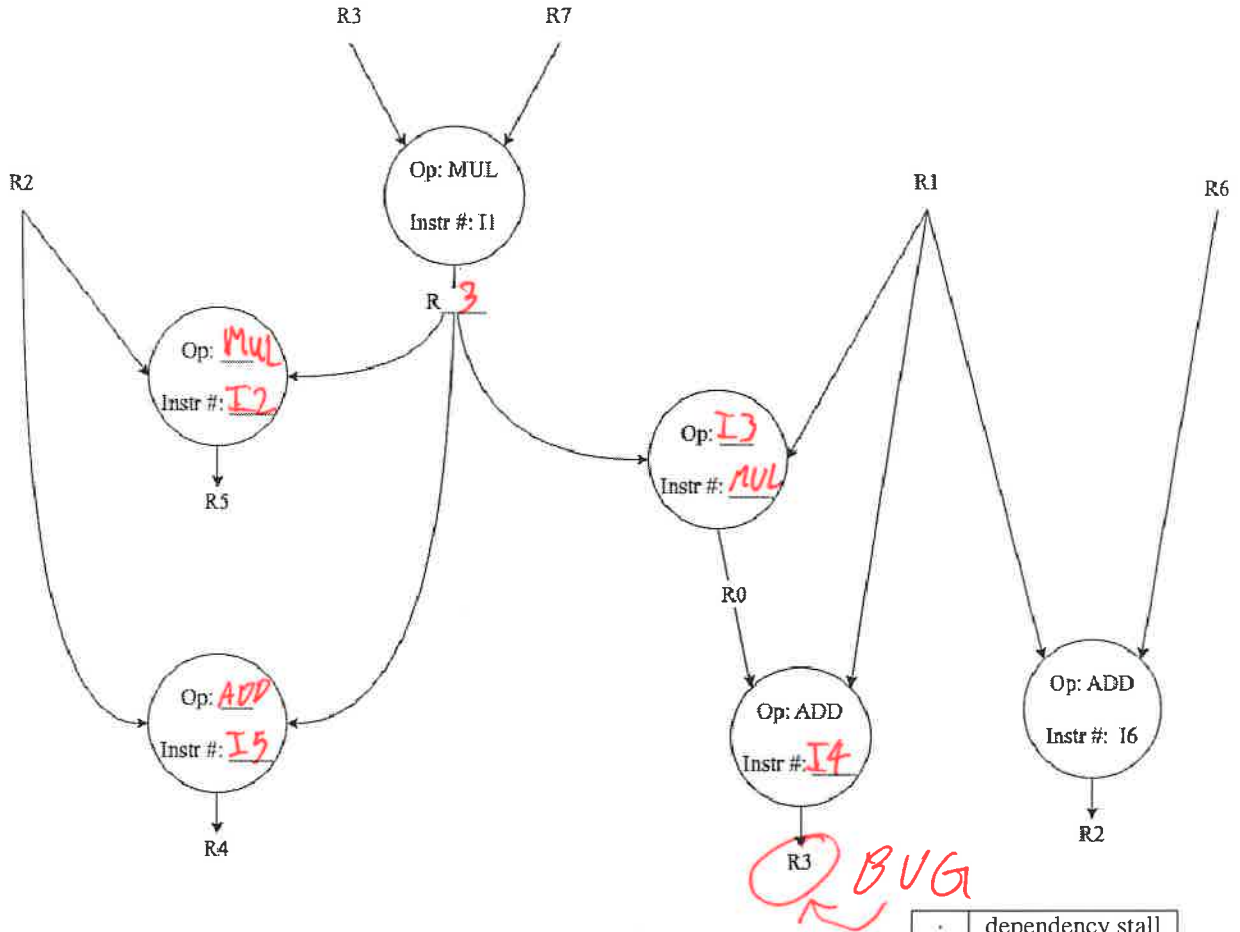


Figure 1: Dataflow Example

Name: \_\_\_\_\_

The data flow graph of the six instruction program is shown below.



The timing diagram below specifies what each instruction is doing during each clock cycle from the time it is fetched until WRITEBACK. All FETCH and WRITEBACK stages have been filled in. Use the symbols to the right to fill in the timing diagram:

.	dependency stall
×	structural stall
D	DECODE
A	ADD Execute
M	MUL Execute

Instruction	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
I1	F	D	RR	M	M	M	W													
I2		F	D	RR	.	.	M	M	M	W										
I3			F	D	RR	.	.	X	X	M	M	M	W							
I4				F	D	<del>RR</del>	<del>RR</del>	RR	.	.	.	.	A	A	W					
I5					F	<del>RR</del>	<del>RR</del>	D	<del>RR</del>	X	RR	A	A	W						
I6						F	F	F	<del>RR</del>	X	D	X	X	RR	A	A	W			

R3 α  
R5 β  
R0 π  
R3 α  
R4 β  
R2 τ

Name: \_\_\_\_\_

The state of the Register Alias Table is shown at three points in time, before cycle 1, at the end of cycle 7, and at the end of cycle 14.

	V	Tag	Val
R0	1	-	0
R1	1	-	1
R2	1	-	2
R3	1	-	3
R4	1	-	4
R5	1	-	5
R6	1	-	6
R7	1	-	7

Initial

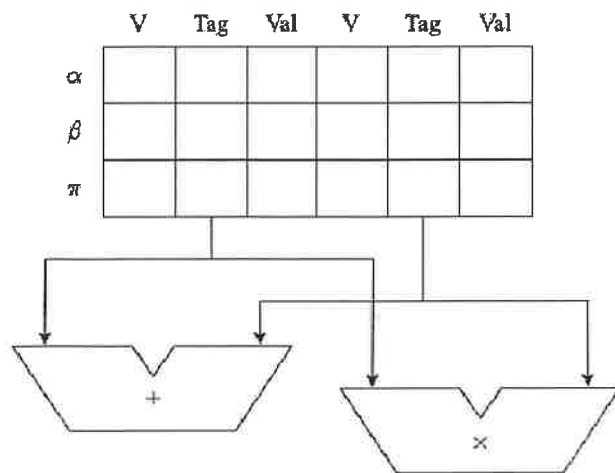
	V	Tag	Val
R0	0	$\tau$	0
R1	1	-	1
R2	1	-	2
R3	1	$\alpha$	21
R4	1	-	4
R5	$\theta$	$\beta$	5
R6	1	-	6
R7	1	-	7

End of cycle 7



	V	Tag	Val
R0	1	$\tau$	21
R1	1	-	1
R2	0	$\tau$	2
R3	0	$\alpha$	21
R4	1	$\beta$	23
R5	1	$\beta$	42
R6	1	-	6
R7	1	-	7

End of cycle 14

We also know the reservation stations can be one of the following options:



Option 1: Centralized



Option 2: Distributed

**Part a (6 points):** Fill in the missing entries in the dataflow graph above by identifying all nodes (opcode and instruction number) and labeling the architectural registers (R0 to R7) that specify the sources and destinations of all instructions.

**Part b (18 points):** Fill in the missing entries in Register Alias Tables and the Timing Diagrams.

Name: \_\_\_\_\_

**Part c (6 points):** Are the reservation stations centralized or distributed? Are the functional units pipelined or not pipelined. Answer the questions in the boxes provided. One or two word answers are sufficient.

Is it Centralized or Distributed

*Centralized*

ADD, Pipelined or not pipelined

*Pipelined*

MUL, Pipelined or not pipelined

*NOT*