

UCR

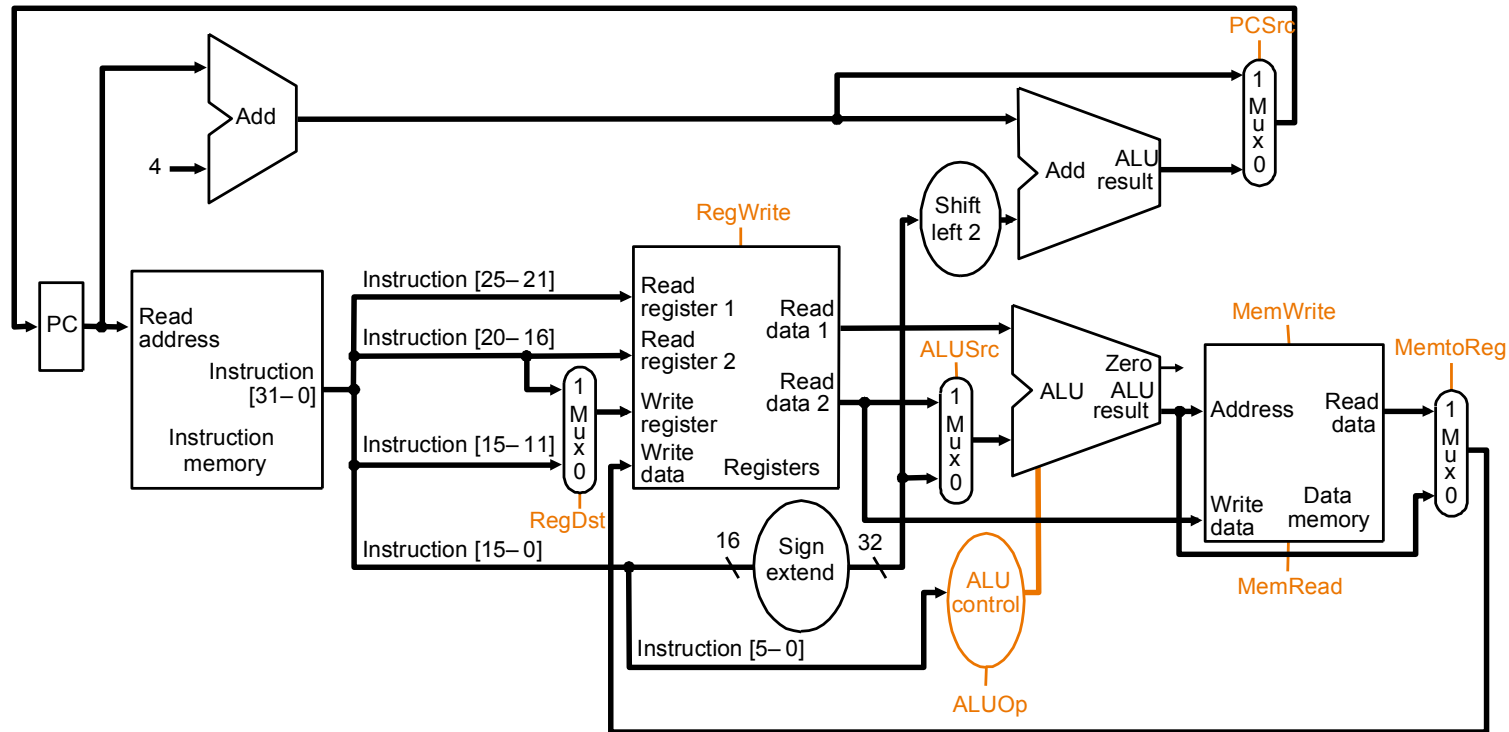
CS161 – Design and Architecture of Computer Systems

Multi-Cycle CPU Design

UNIVERSITY OF CALIFORNIA, RIVERSIDE

Single Cycle Implementation

- Calculate cycle time assuming negligible delays except:
 - memory (2ns), ALU and adders (2ns), register file access (1ns)



Single Cycle – Steps of each instruction



Inst. Type	Functional Units Used				
R-type	Instruction fetch	Register read	ALU	Register write	
Load	Instruction fetch	Register read	ALU	Memory access	Register write
Store	Instruction fetch	Register read	ALU	Memory access	
Branch	Instruction fetch	Register read	ALU		
Jump	Instruction fetch				

Single Cycle – How long is the cycle?

Inst. Type	Inst. Mem.	Reg. File (read)	ALU (s)	Data Mem.	Reg. File (write)	Total	Inst. %
R-type	2	1	2	0	1	6 ns	44
Load	2	1	2	2	1	8 ns	24
Store	2	1	2	2	0	7 ns	12
Branch	2	1	2	0	0	5 ns	18
Jump	2	0	0	0	0	2 ns	2

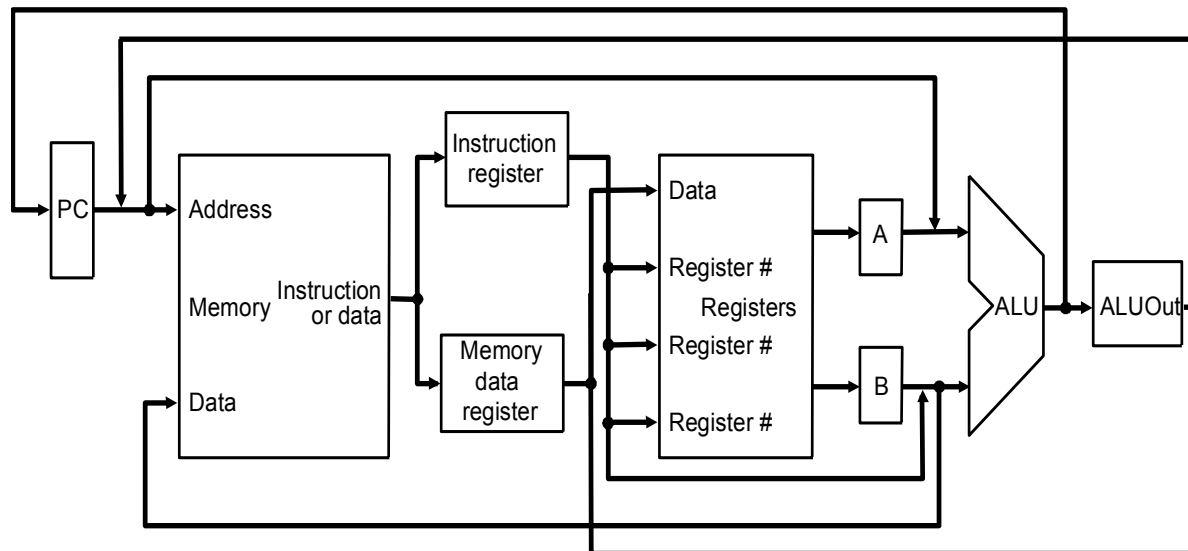
The cycle time must accommodate the longest operation: *lw*.
Cycle time = 8 ns but the CPI = 1.

If we can accommodate variable number of cycles for each instruction and a cycle time of 1ns.

$$\text{CPI} = 6 \cdot 44\% + 8 \cdot 24\% + 7 \cdot 12\% + 5 \cdot 18\% + 2 \cdot 2\% = 6.3$$

Where we are headed

- Single Cycle Problems:
 - what if we had a more complicated instruction like floating point?
 - waste of area
- One Solution:
 - use a “smaller” cycle time
 - have different instructions take different numbers of cycles
 - a “multicycle” datapath:

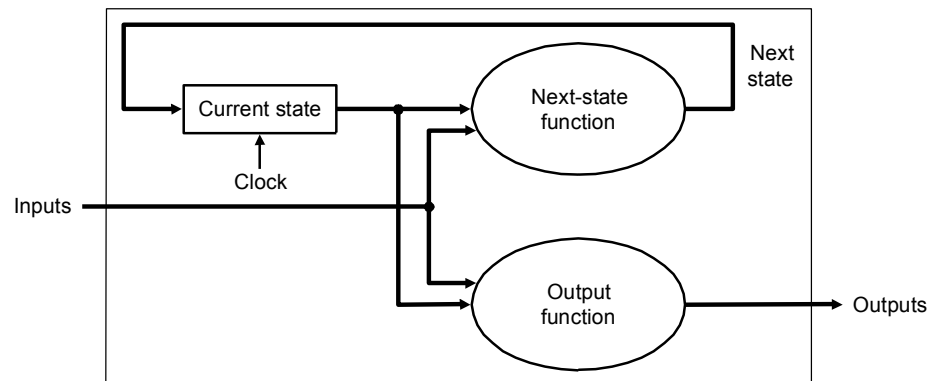


Multicycle Approach

- ▶ We will be reusing functional units
 - ▶ ALU used to compute address and to increment PC
 - ▶ Memory used for instruction and data
- ▶ Our control signals will not be determined solely by instruction
 - ▶ e.g., what should the ALU do for a “subtract” instruction?
- ▶ We’ll use a finite state machine for control

Review: finite state machines

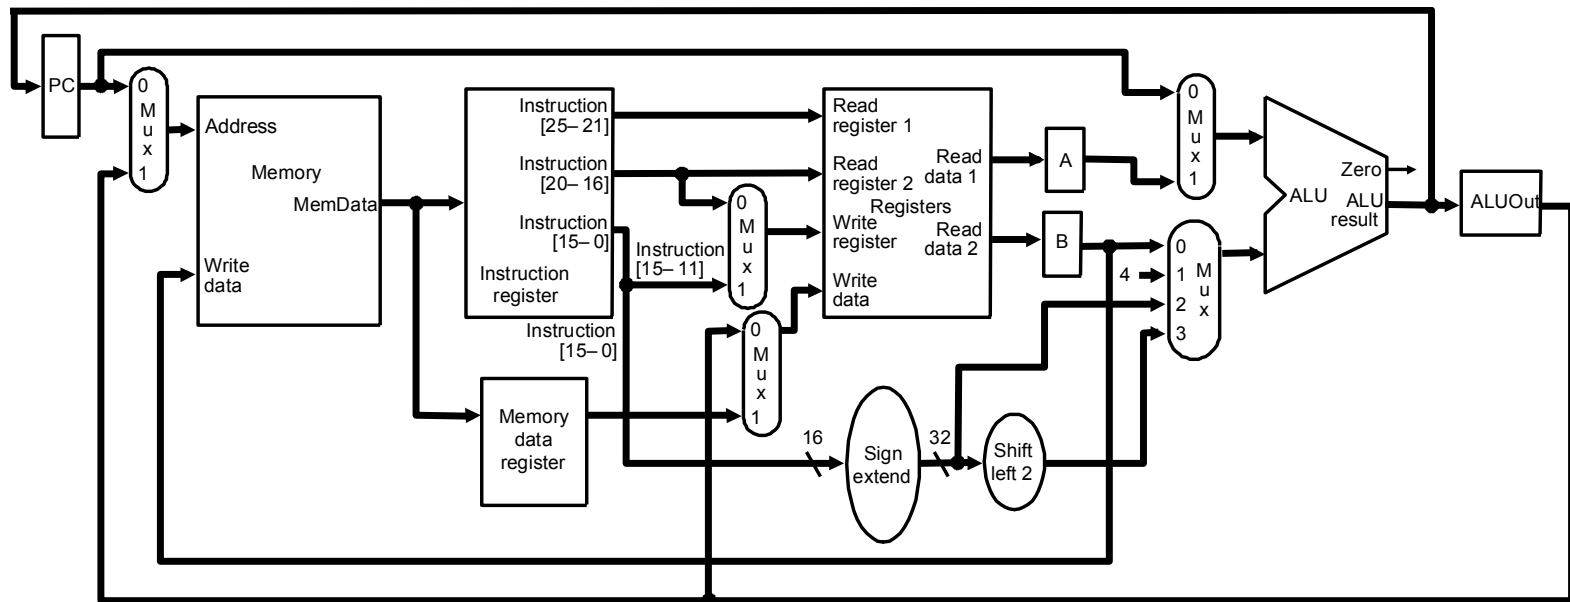
- ▶ Finite state machines:
 - ▶ a set of states and
 - ▶ next state function (determined by current state and the input)
 - ▶ output function (determined by current state and possibly input)



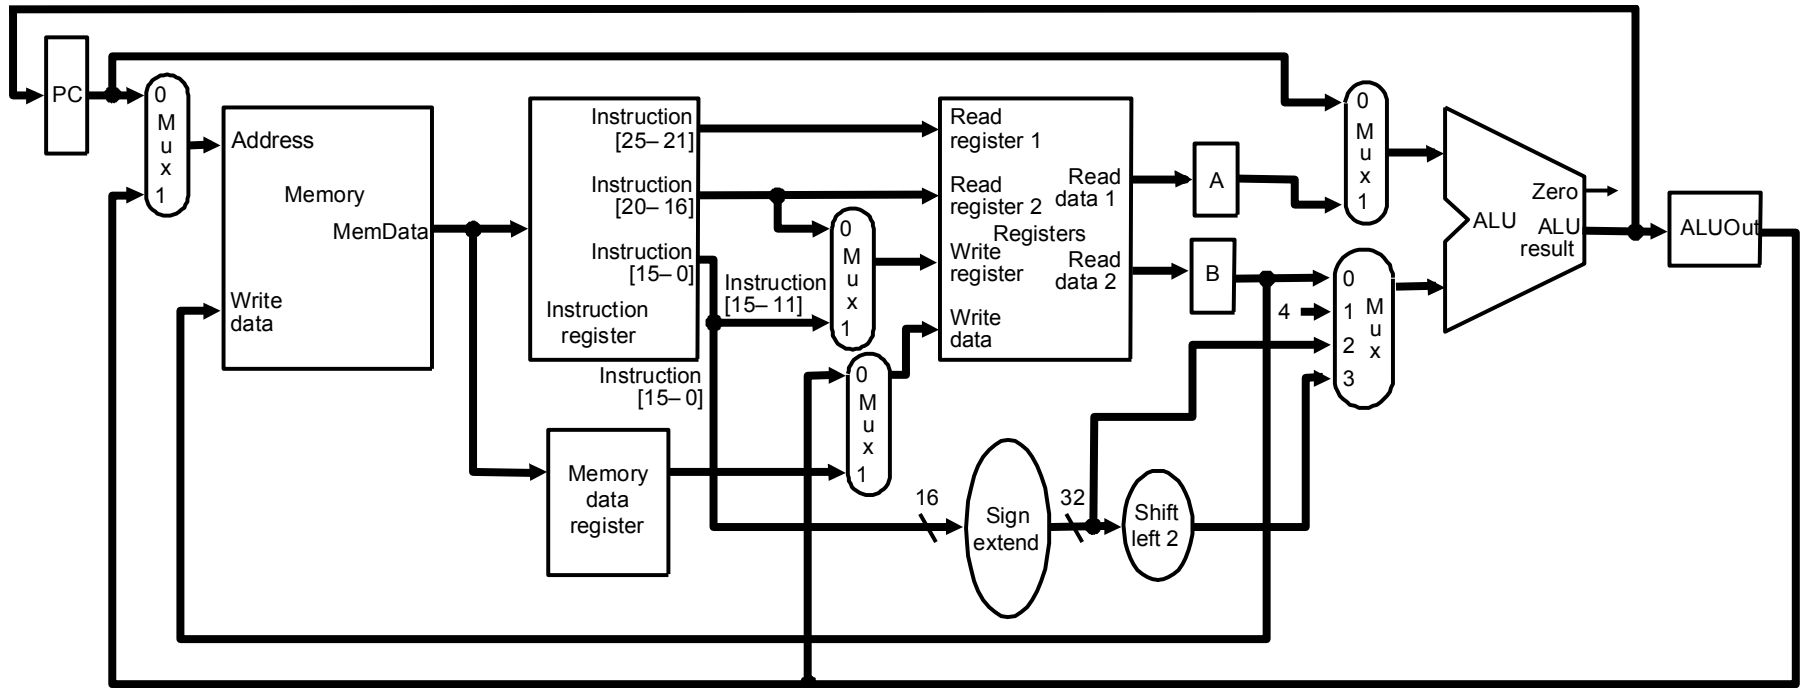
- ▶ We'll use a Moore machine (output based only on current state)

Multicycle Approach

- Break up the instructions into steps, each step takes a cycle
 - balance the amount of work to be done
 - restrict each cycle to use only one major functional unit
- At the end of a cycle
 - store values for use in later cycles (easiest thing to do)
 - introduce additional “internal” registers



Muticycle Datapath



Five Execution Steps

- Instruction Fetch (F)
- Instruction Decode and Register Fetch (D)
- Execution, Memory Address Computation, or Branch Completion (EX)
- Memory Access or R-type instruction completion (M)
- Write-back step (W)

INSTRUCTIONS TAKE FROM 3 - 5 CYCLES!

Step 1: Instruction Fetch


- ❑ Use PC to get instruction and put it in the Instruction Register.
- ❑ Increment the PC by 4 and put the result back in the PC.
- ❑ Can be described concisely using RTL "Register-Transfer Language"

```
IR = Memory[PC];
```

```
PC = PC + 4;
```

- ❑ *Can we figure out the values of the control signals?*

Datapath of Multicycle Implementation



Step 2: Instruction Decode and Register Fetch



- Read registers rs and rt in case we need them
- Compute the branch address in case the instruction is a branch
- RTL:

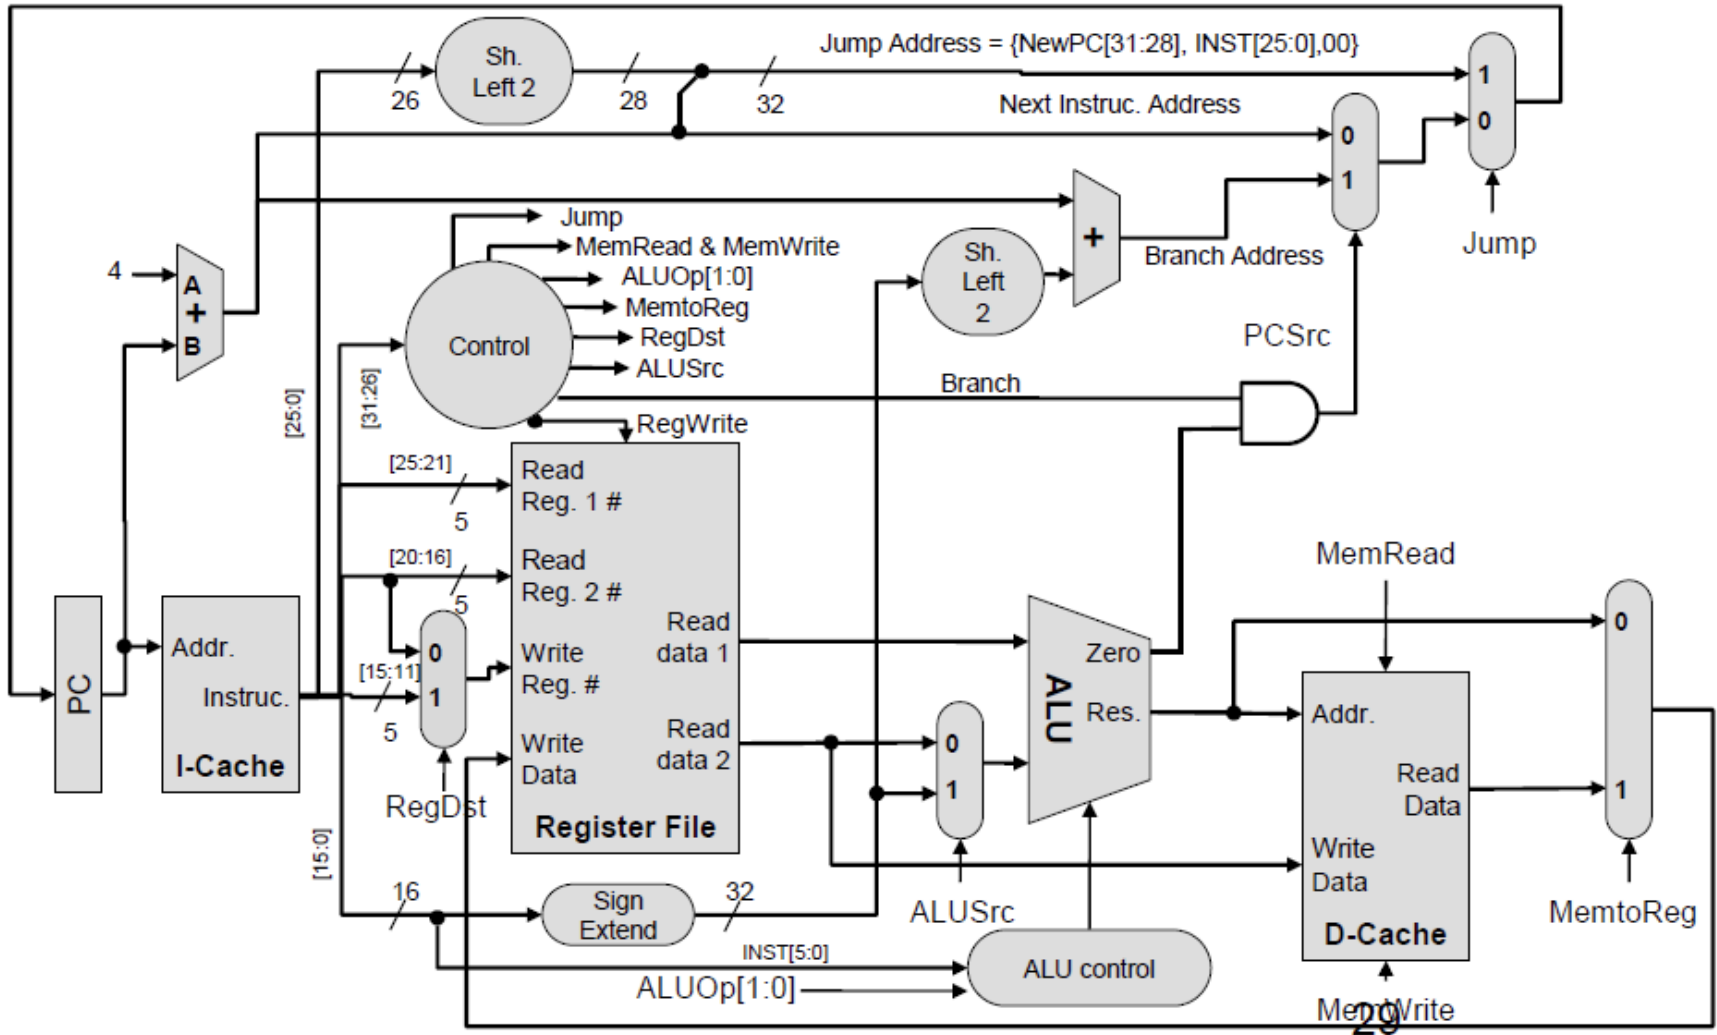
```
A = Reg[IR[25-21]];
```

```
B = Reg[IR[20-16]];
```

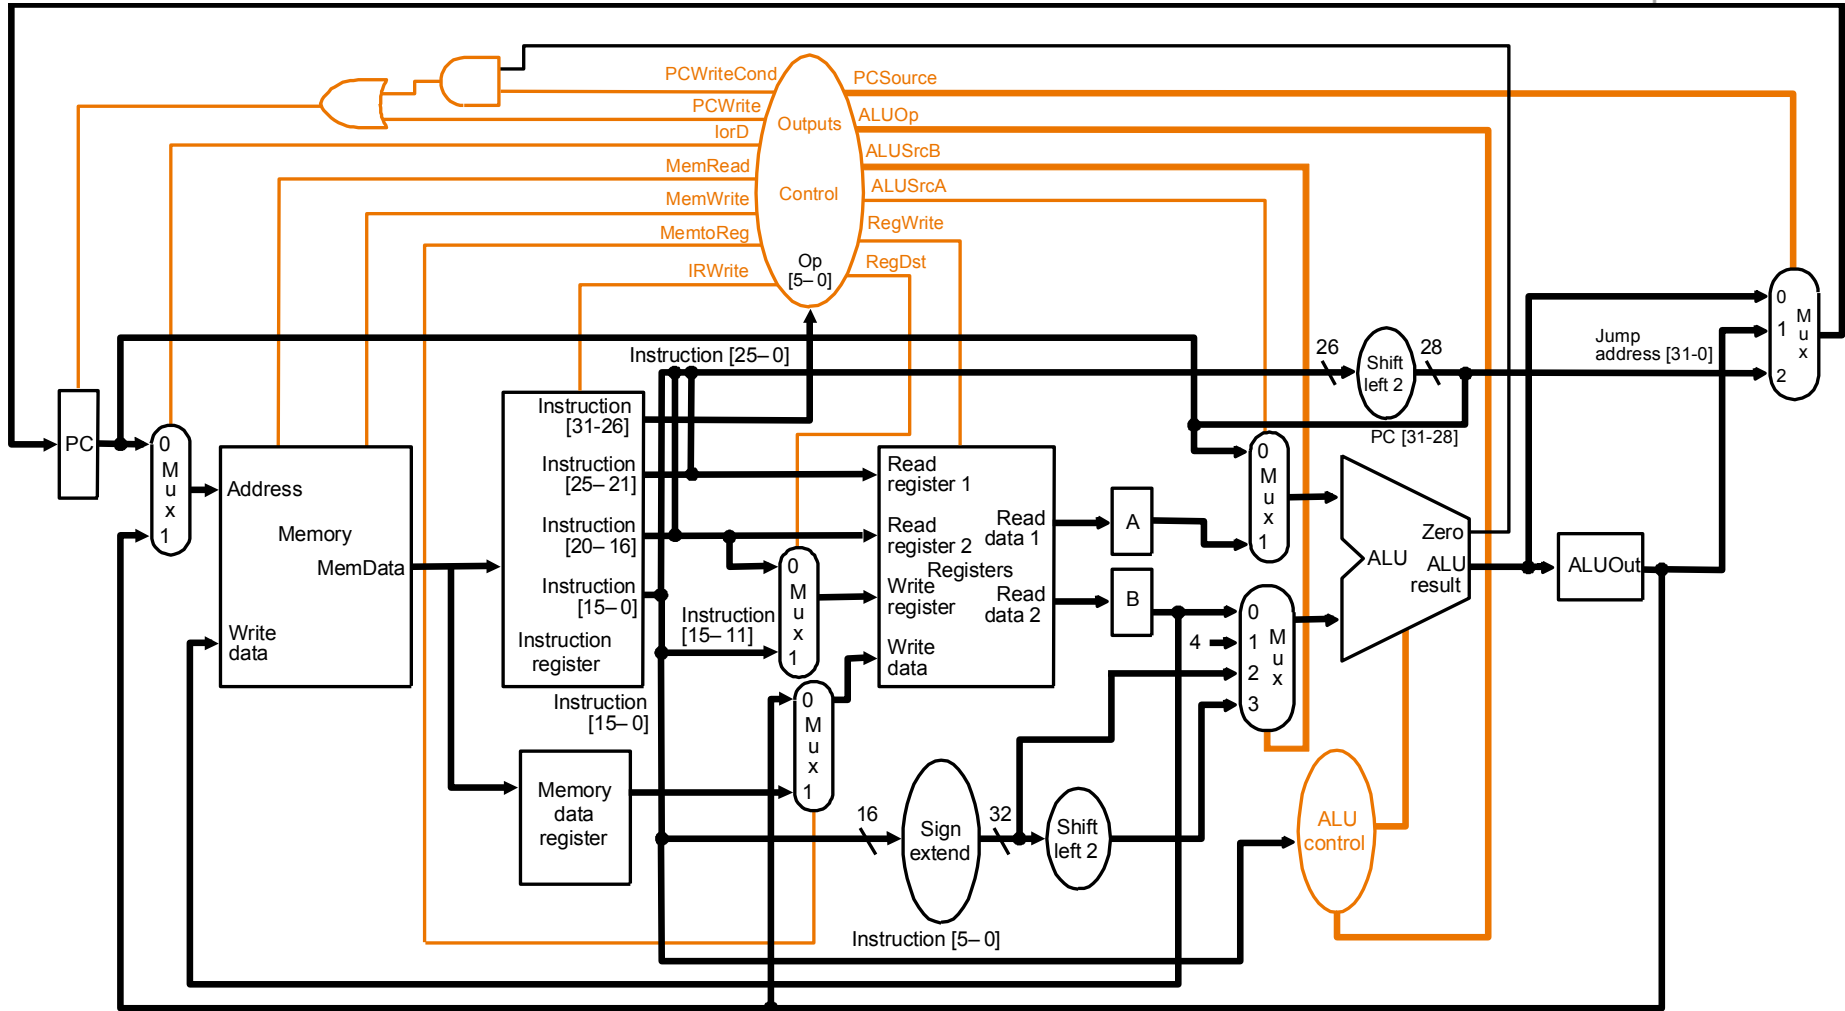
```
ALUOut = PC + (sign-extend(IR[15-0]) << 2);
```

- We aren't setting any control lines based on the instruction type
(we are busy "decoding" it in our control logic)

Single Cycle Implementation



Datapath of Multicycle Implementation



Step 3 (instruction dependent)



- › ALU is performing one of three functions, based on instruction type

- › Memory Reference:

```
ALUOut = A + sign-extend(IR[15-0]);
```


- › R-type:

```
ALUOut = A op B;
```

- › Branch:

```
if (A==B) PC = ALUOut;
```


Datapath of Multicycle Implementation



Step 4 (R-type or memory-access)

- › Loads and stores access memory


```
MDR = Memory[ALUOut];
    or
Memory[ALUOut] = B;
```

- › R-type instructions finish

```
Reg[IR[15-11]] = ALUOut;
```

The write actually takes place at the end of the cycle on the edge

Datapath of Multicycle Implementation



Step 5 Write-back step



□ $\text{Reg}[\text{IR}[20-16]] = \text{MDR};$

What about all the other instructions?

Summary:



Step name	Action for R-type instructions	Action for memory-reference instructions	Action for branches	Action for jumps
Instruction fetch	$IR = \text{Memory}[PC]$ $PC = PC + 4$			
Instruction decode/register fetch	$A = \text{Reg} [IR[25-21]]$ $B = \text{Reg} [IR[20-16]]$ $ALUOut = PC + (\text{sign-extend} (IR[15-0]) \ll 2)$			
Execution, address computation, branch/ jump completion	$ALUOut = A \text{ op } B$	$ALUOut = A + \text{sign-extend} (IR[15-0])$	if $(A == B)$ then $PC = ALUOut$	$PC = PC [31-28] \parallel (IR[25-0] \ll 2)$
Memory access or R-type completion	$\text{Reg} [IR[15-11]] = ALUOut$	Load: $MDR = \text{Memory}[ALUOut]$ or Store: $\text{Memory} [ALUOut] = B$		
Memory read completion		Load: $\text{Reg}[IR[20-16]] = MDR$		

Simple Questions

- How many cycles will it take to execute this code?

```
lw $t2, 0($t3)
lw $t3, 4($t3)
beq $t2, $t3, Label ← #assume not
add $t5, $t2, $t3
sw $t5, 8($t3)
```

Label: ...

- What is going on during the 8th cycle of execution?
- In what cycle does the actual addition of $\$t2$ and $\$t3$ takes place?

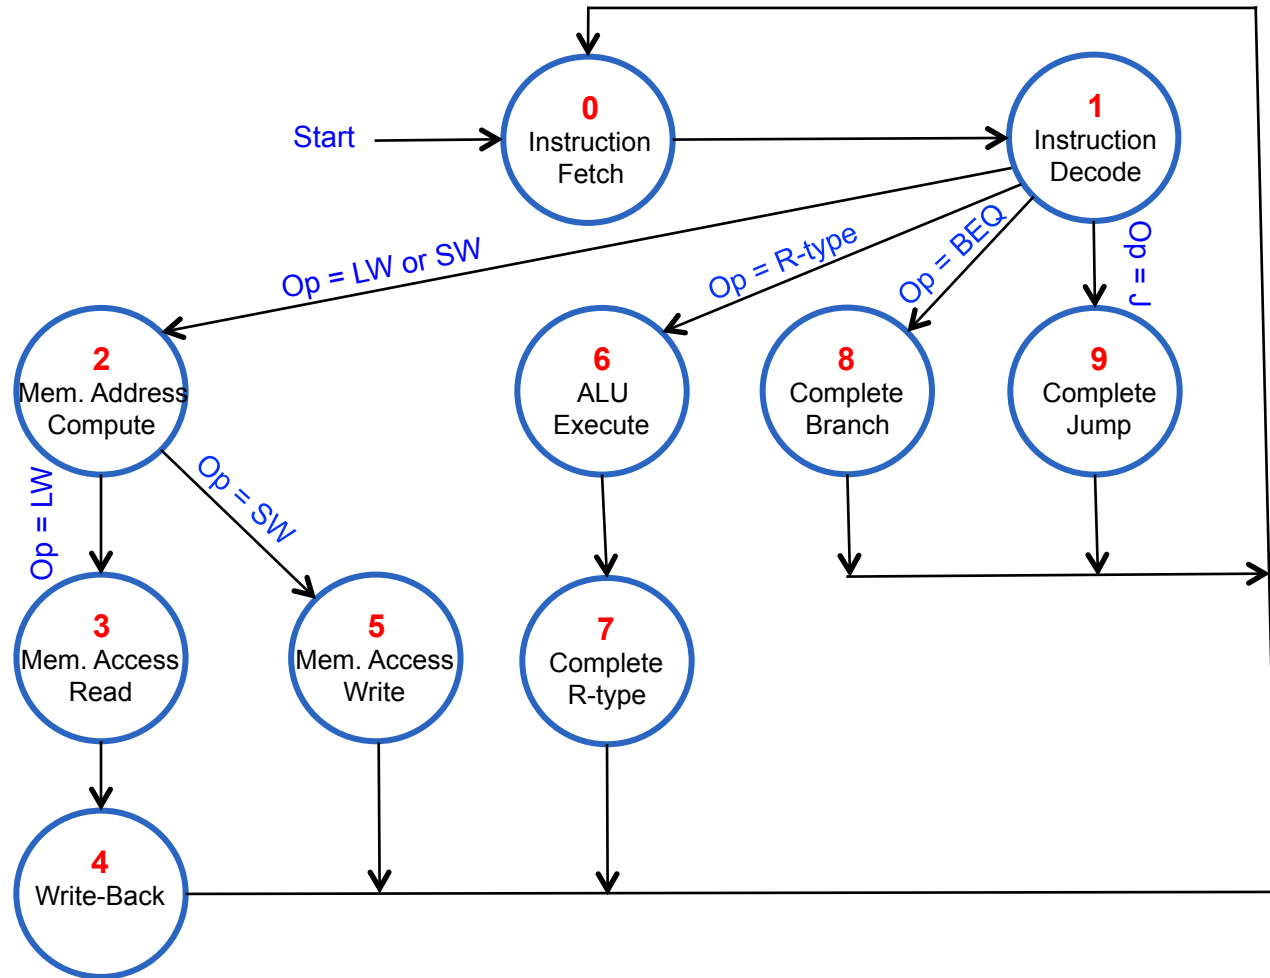


Implementing the Control

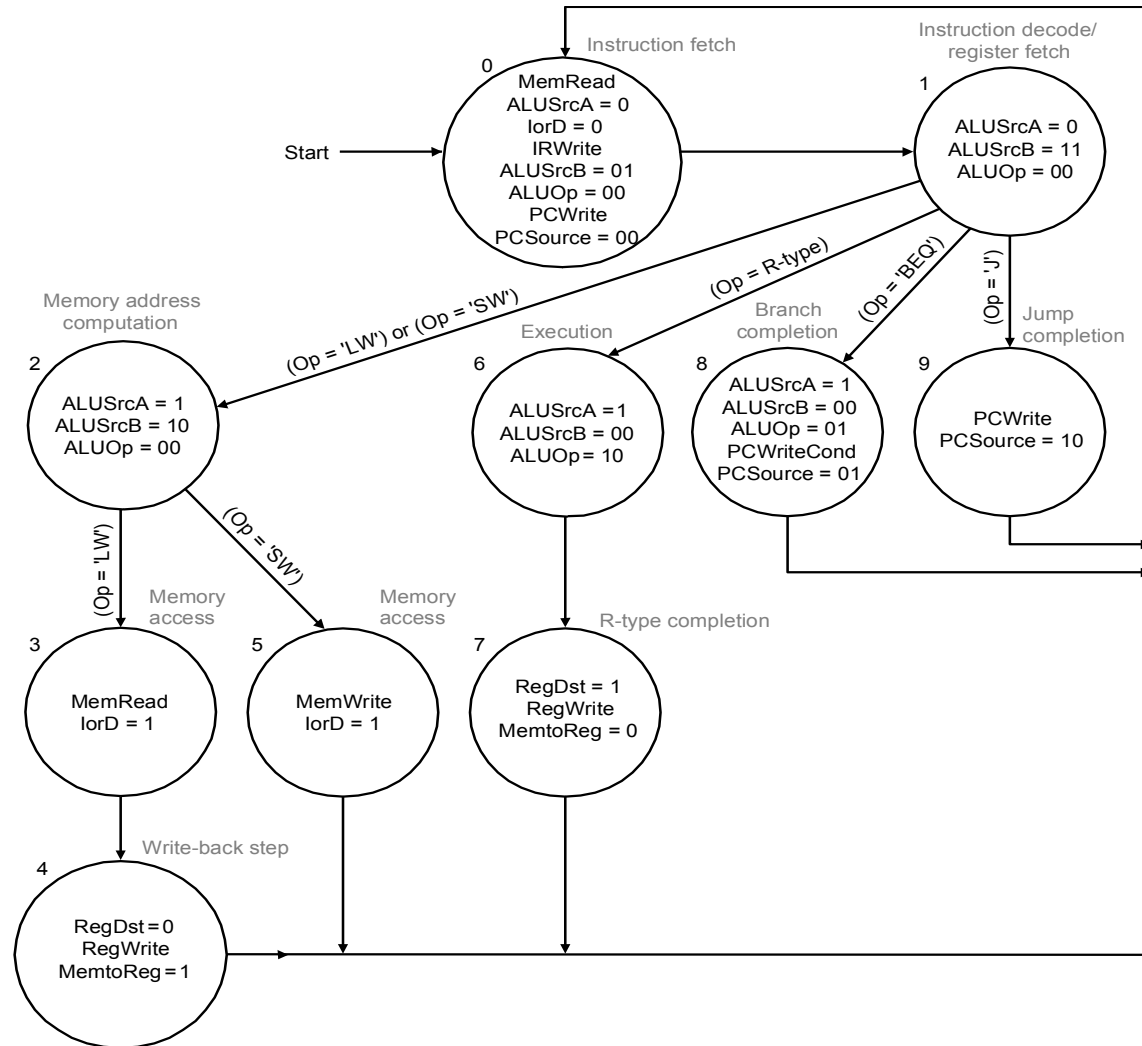


- ▶ Value of control signals is dependent upon:
 - ▶ what instruction is being executed
 - ▶ which step is being performed
- ▶ Use the information we've accumulated to specify a finite state machine
 - ▶ specify the finite state machine graphically, or
 - ▶ use microprogramming
- ▶ Implementation can be derived from specification

Graphical Specification of FSM



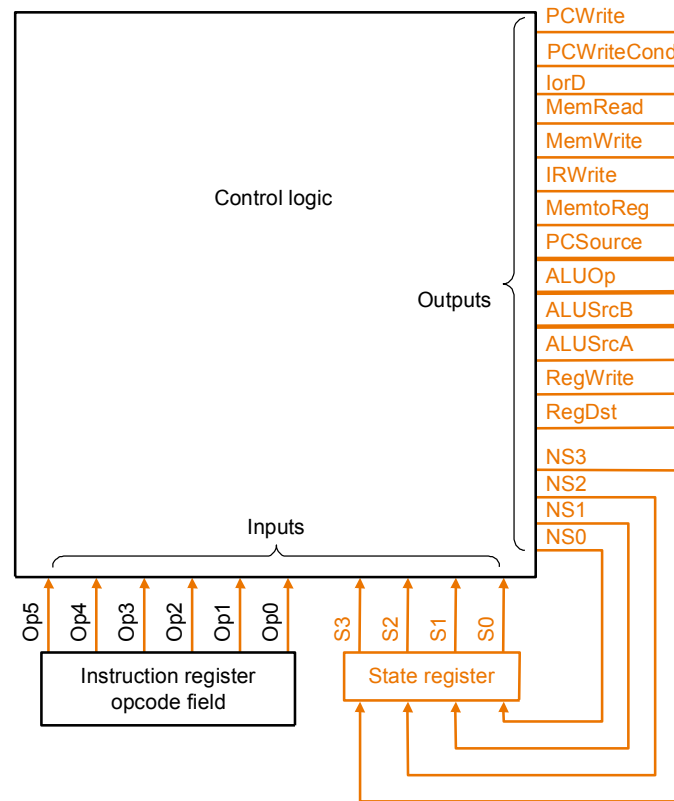
Detailed FSM




Finite State Machine for Control



Implementation:



Inputs



R-format

lw

sw

beq

Outputs

RegDst

ALUSrc

MemtoReg

RegWrite


MemRead

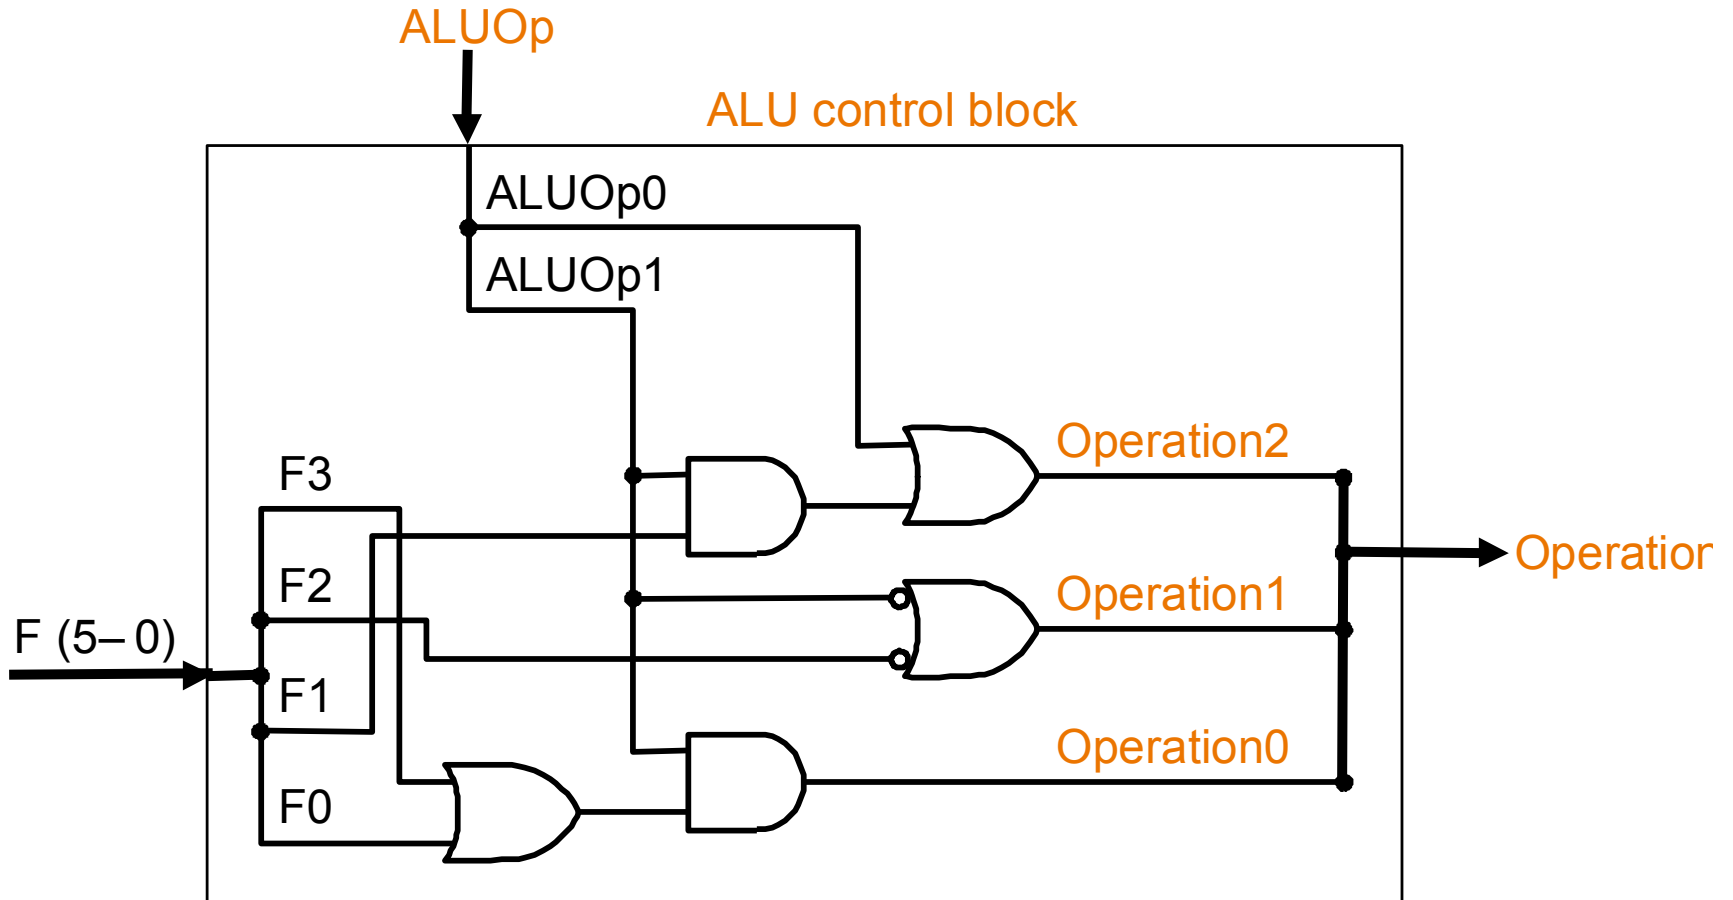
MemWrite

Branch

ALUOp1


ALUOp0






PLA Implementation

- Programmable Logic Array
- If I picked a horizontal or vertical line could you explain it?



ROM Implementation

- ❑ ROM = "Read Only Memory"
 - values of memory locations are fixed ahead of time
- ❑ A ROM can be used to implement a truth table
 - if the address is m -bits, we can address 2^m entries in the ROM.
 - our outputs are the bits of data that the address points to.



0	0	0	0	0	1	1
0	0	1	1	1	0	0
0	1	0	1	1	0	0
0	1	1	1	0	0	0
1	0	0	0	0	0	0
1	0	1	0	0	0	1
1	1	0	0	1	1	0
1	1	1	0	1	1	1

m is the "height", and n is the "width"

ROM Implementation



- How many inputs are there?
 - 6 bits for opcode, 4 bits for state = 10 address lines
(i.e., $2^{10} = 1024$ different addresses)
- How many outputs are there?
 - 16 datapath-control outputs, 4 state bits = 20 outputs
- ROM is $2^{10} \times 20 = 20\text{K}$ bits
- Rather wasteful, since for lots of the entries, the outputs are the same
 - i.e., opcode is often ignored

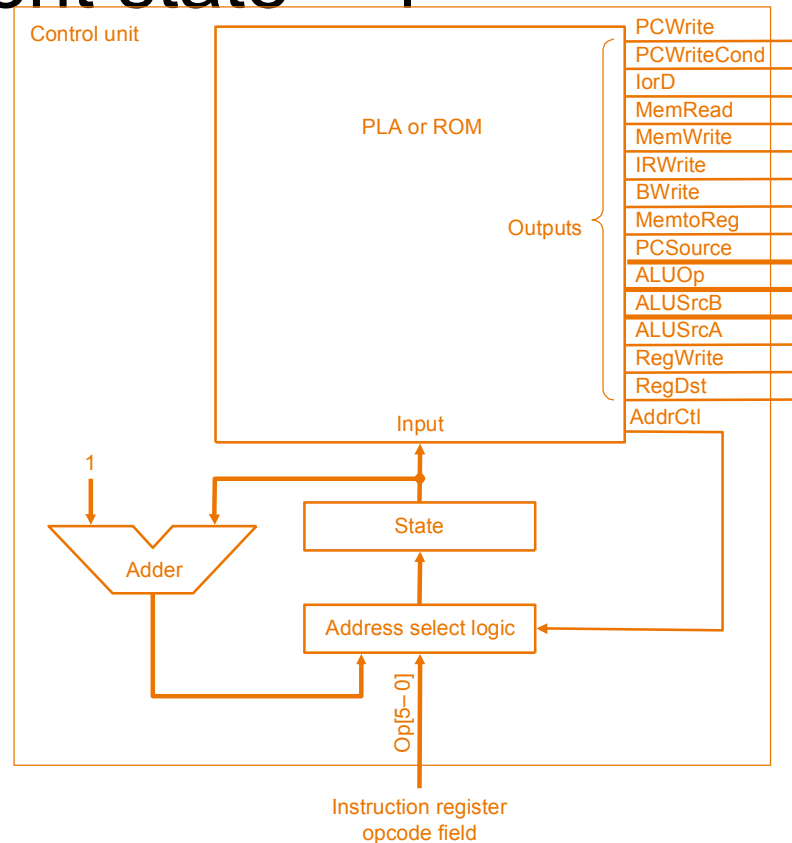
ROM vs PLA



- › Break up the table into two parts
 - 4 state bits tell you the 16 outputs, $2^4 \times 16$ bits of ROM
 - 10 bits tell you the 4 next state bits, $2^{10} \times 4$ bits of ROM
 - Total: 4.3K bits of ROM
- › PLA is much smaller
 - can share product terms
 - only need entries that produce an active output
 - can take into account don't cares
- › Size is $(\#inputs \times \#product\text{-terms}) + (\#outputs \times \#product\text{-terms})$
For this example = $(10 \times 17) + (20 \times 17) = 460$ PLA cells
- › PLA cells usually about the size of a ROM cell (slightly bigger)

Another Implementation Style

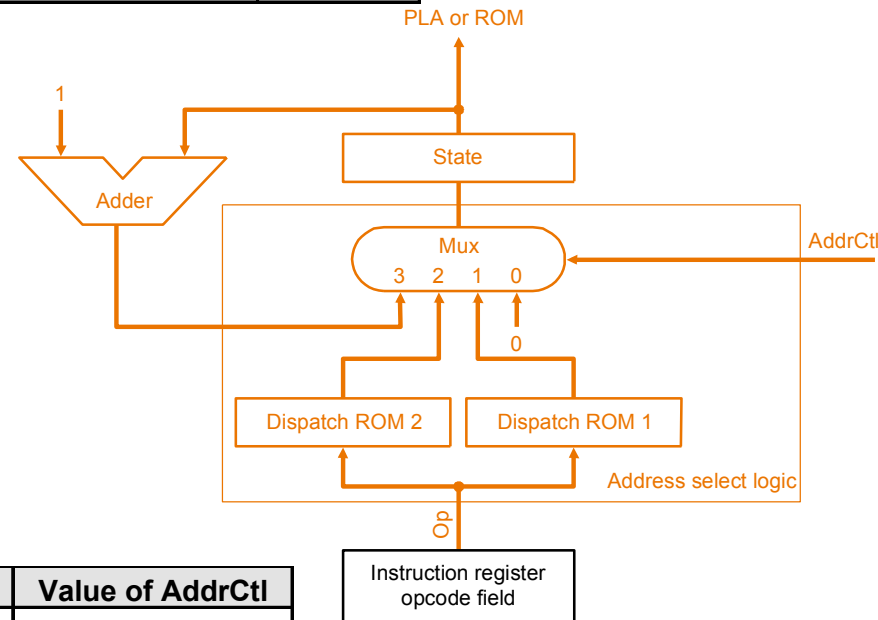
- ❑ Complex instructions: the "next state" is often current state + 1



Details

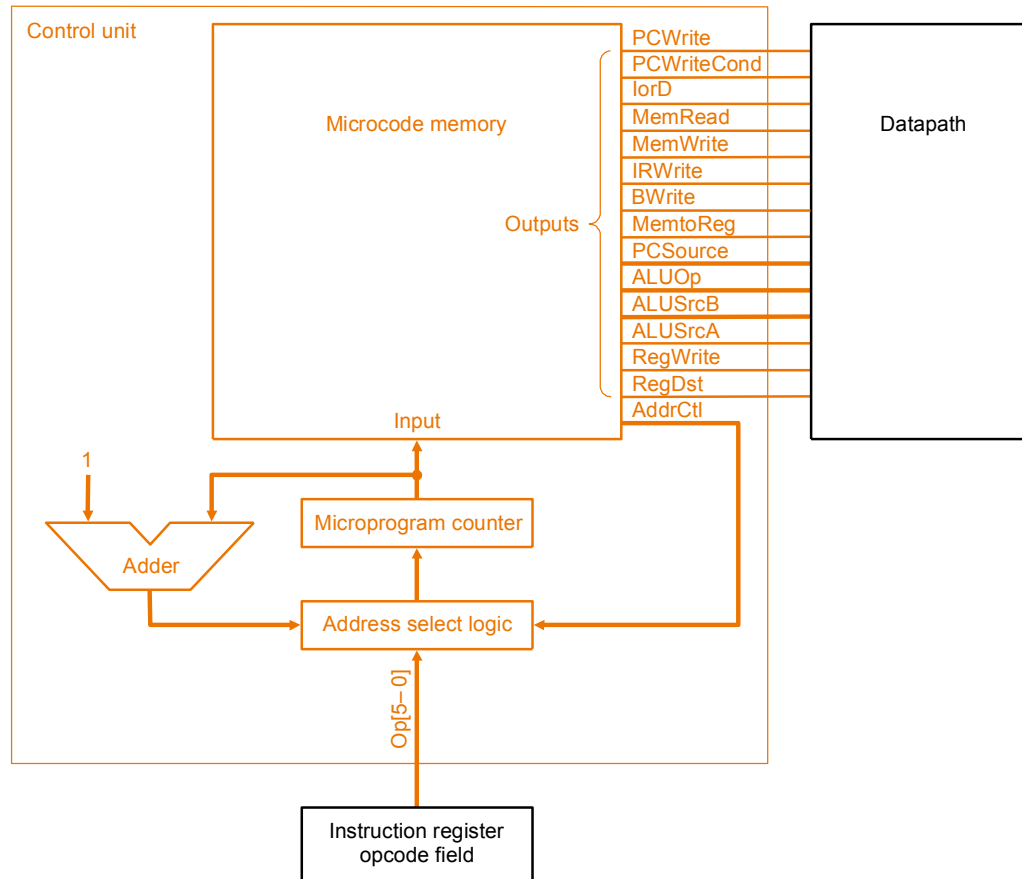
Dispatch ROM 1		
Op	Opcode name	Value
000000	R-format	0110
000010	jmp	1001
000100	beq	1000
100011	lw	0010
101011	sw	0010

Dispatch ROM 2		
Op	Opcode name	Value
100011	lw	0011
101011	sw	0101



State number	Address-control action	Value of AddrCtl
0	Use incremented state	3
1	Use dispatch ROM 1	1
2	Use dispatch ROM 2	2
3	Use incremented state	3
4	Replace state number by 0	0
5	Replace state number by 0	0
6	Use incremented state	3
7	Replace state number by 0	0
8	Replace state number by 0	0
9	Replace state number by 0	0

Microprogramming



Microprogramming



- › How all CISC ISAs were built.
- › μ -store: RAM-memory with N bits/word
 - › N = number of control lines in CPU.
 - › Execution of each instruction starts IF and ID.
 - › Instruction Opcode points to a location in μ -store
 - › Output consecutive control words until DONE.
- › μ -store is updatable
- › μ -controller is = a mini-CPU running the main CPU
- › M-programming worked when CPU = multiple boards
 - › μ -controller on one board: much faster

Microprogramming



- ❑ A specification methodology
 - appropriate if hundreds of opcodes, modes, cycles, etc.
 - signals specified symbolically using microinstructions

Label	ALU control	SRC1	SRC2	Register control	Memory	PCWrite control	Sequencing
Fetch	Add	PC	4		Read PC	ALU	Seq
	Add	PC	Extshft	Read			Dispatch 1
Mem1	Add	A	Extend				Dispatch 2
LW2					Read ALU		Seq
				Write MDR			Fetch
SW2					Write ALU		Fetch
Rformat1	Func code	A	B				Seq
				Write ALU			Fetch
BEQ1	Subt	A	B			ALUOut-cond	Fetch
JUMP1						Jump address	Fetch

- ❑ *Will two implementations of the same architecture have the same microcode?*
- ❑ *What would a microassembler do?*

Microinstruction format



Field name	Value	Signals active	Comment
ALU control	Add	ALUOp = 00	Cause the ALU to add.
	Subt	ALUOp = 01	Cause the ALU to subtract; this implements the compare for branches.
	Func code	ALUOp = 10	Use the instruction's function code to determine ALU control.
SRC1	PC	ALUSrcA = 0	Use the PC as the first ALU input.
	A	ALUSrcA = 1	Register A is the first ALU input.
SRC2	B	ALUSrcB = 00	Register B is the second ALU input.
	4	ALUSrcB = 01	Use 4 as the second ALU input.
	Extend	ALUSrcB = 10	Use output of the sign extension unit as the second ALU input.
	Extshft	ALUSrcB = 11	Use the output of the shift-by-two unit as the second ALU input.
Register control	Read		Read two registers using the rs and rt fields of the IR as the register numbers and putting the data into registers A and B.
	Write ALU	RegWrite, RegDst = 1, MemtoReg = 0	Write a register using the rd field of the IR as the register number and the contents of the ALUOut as the data.
	Write MDR	RegWrite, RegDst = 0, MemtoReg = 1	Write a register using the rt field of the IR as the register number and the contents of the MDR as the data.
Memory	Read PC	MemRead, lorD = 0	Read memory using the PC as address; write result into IR (and the MDR).
	Read ALU	MemRead, lorD = 1	Read memory using the ALUOut as address; write result into MDR.
	Write ALU	MemWrite, lorD = 1	Write memory using the ALUOut as address, contents of B as the data.
PC write control	ALU	PCSource = 00 PCWrite	Write the output of the ALU into the PC.
	ALUOut-cond	PCSource = 01, PCWriteCond	If the Zero output of the ALU is active, write the PC with the contents of the register ALUOut.
	jump address	PCSource = 10, PCWrite	Write the PC with the jump address from the instruction.
Sequencing	Seq	AddrCtl = 11	Choose the next microinstruction sequentially.
	Fetch	AddrCtl = 00	Go to the first microinstruction to begin a new instruction.
	Dispatch 1	AddrCtl = 01	Dispatch using the ROM 1.
	Dispatch 2	AddrCtl = 10	Dispatch using the ROM 2.

Maximally vs. Minimally Encoded



- › No encoding:
 - › 1 bit for each datapath operation
 - › faster, requires more memory (logic)
 - › used for Vax 780 — an astonishing 400K of memory!
- › Lots of encoding:
 - › send the microinstructions through logic to get control signals
 - › uses less memory, slower
- › Historical context of CISC:
 - › Too much logic to put on a single chip with everything else
 - › Use a ROM (or even RAM) to hold the microcode
 - › It's easy to add new instructions

Microcode: Trade-offs

- ❑ Distinction between specification and implementation is sometimes blurred
- ❑ Specification Advantages:
 - Easy to design and write
 - Design architecture and microcode in parallel
- ❑ Implementation (off-chip ROM) Advantages
 - Easy to change since values are in memory
 - Can emulate other architectures
 - Can make use of internal registers
- ❑ Implementation Disadvantages, SLOWER now that:
 - Control is implemented on same chip as processor
 - ROM is no longer faster than RAM
 - No need to go back and make changes

The Big Picture

