

18-447

# Computer Architecture

## Lecture 5: Single-Cycle Microarchitectures

Prof. Onur Mutlu

Carnegie Mellon University

Spring 2014, 1/24/2014

# Assignments

---

- Lab 1 due today
- Lab 2 out (start early)
- HW1 due next week
- HW0 issues
  - Make sure your forms are correctly filled in and readable
  - Extended deadline to resubmit: Sunday night (January 26)

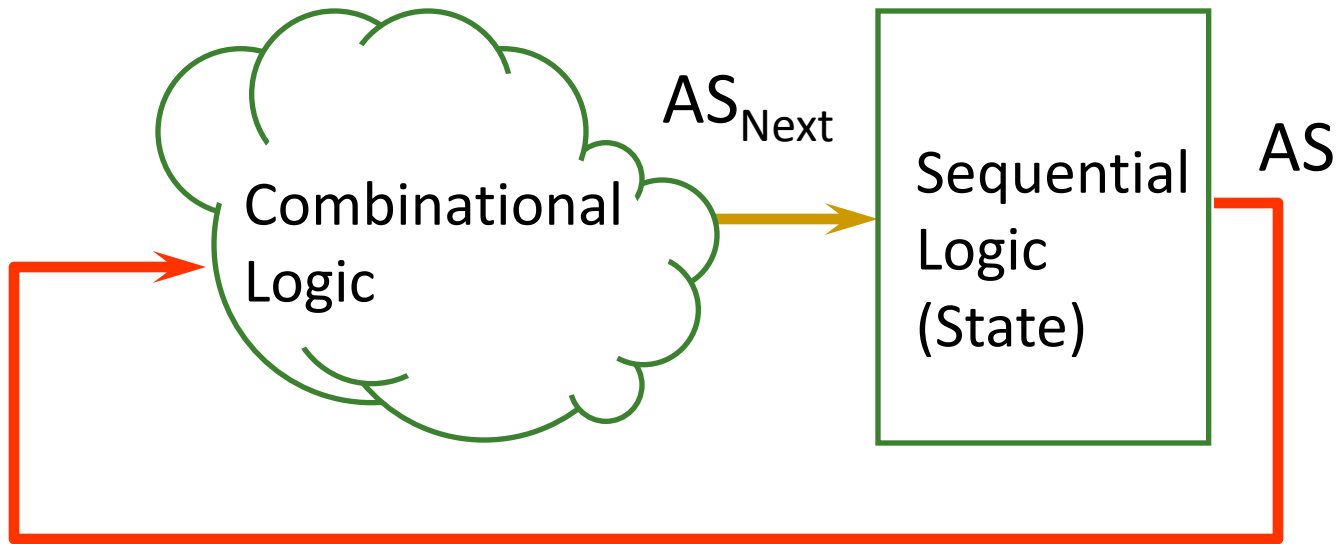
# A Single-Cycle Microarchitecture

## *A Closer Look*

# Remember...

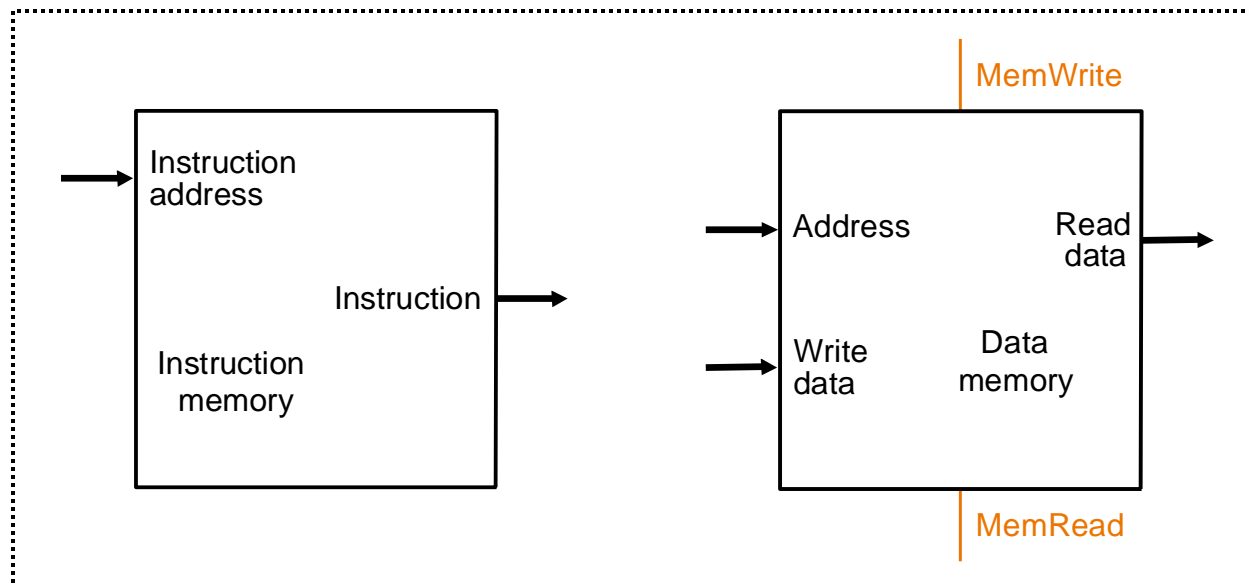
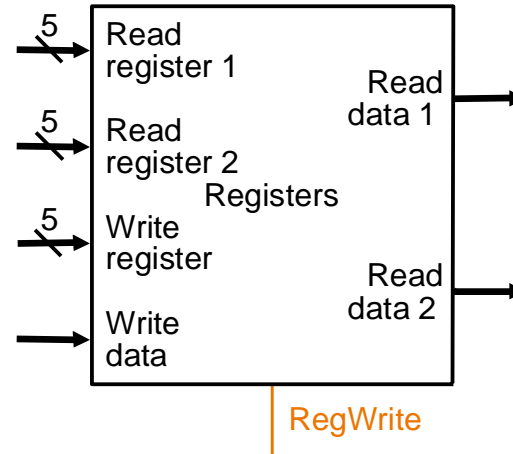

---

- Single-cycle machine



# Let's Start with the State Elements

- Data and control inputs



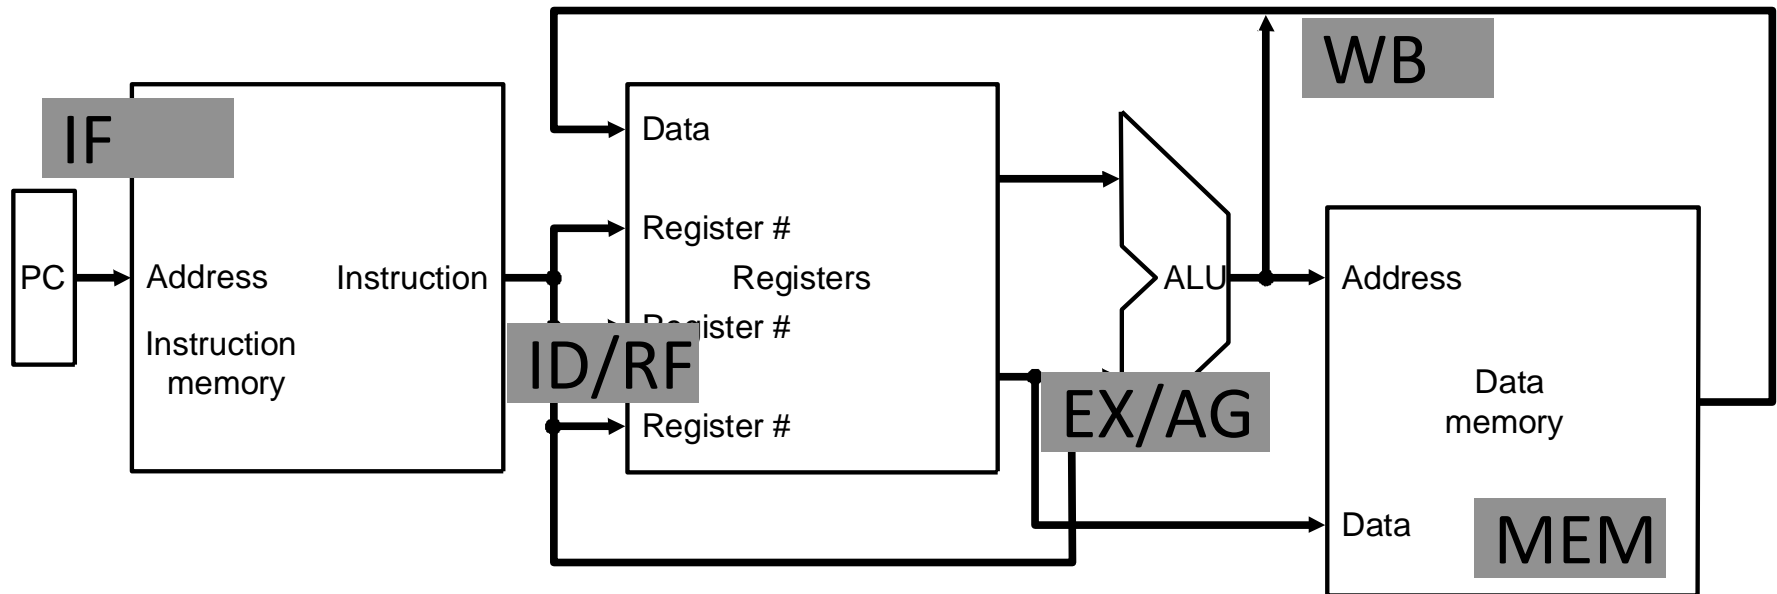
# For Now, We Will Assume

---

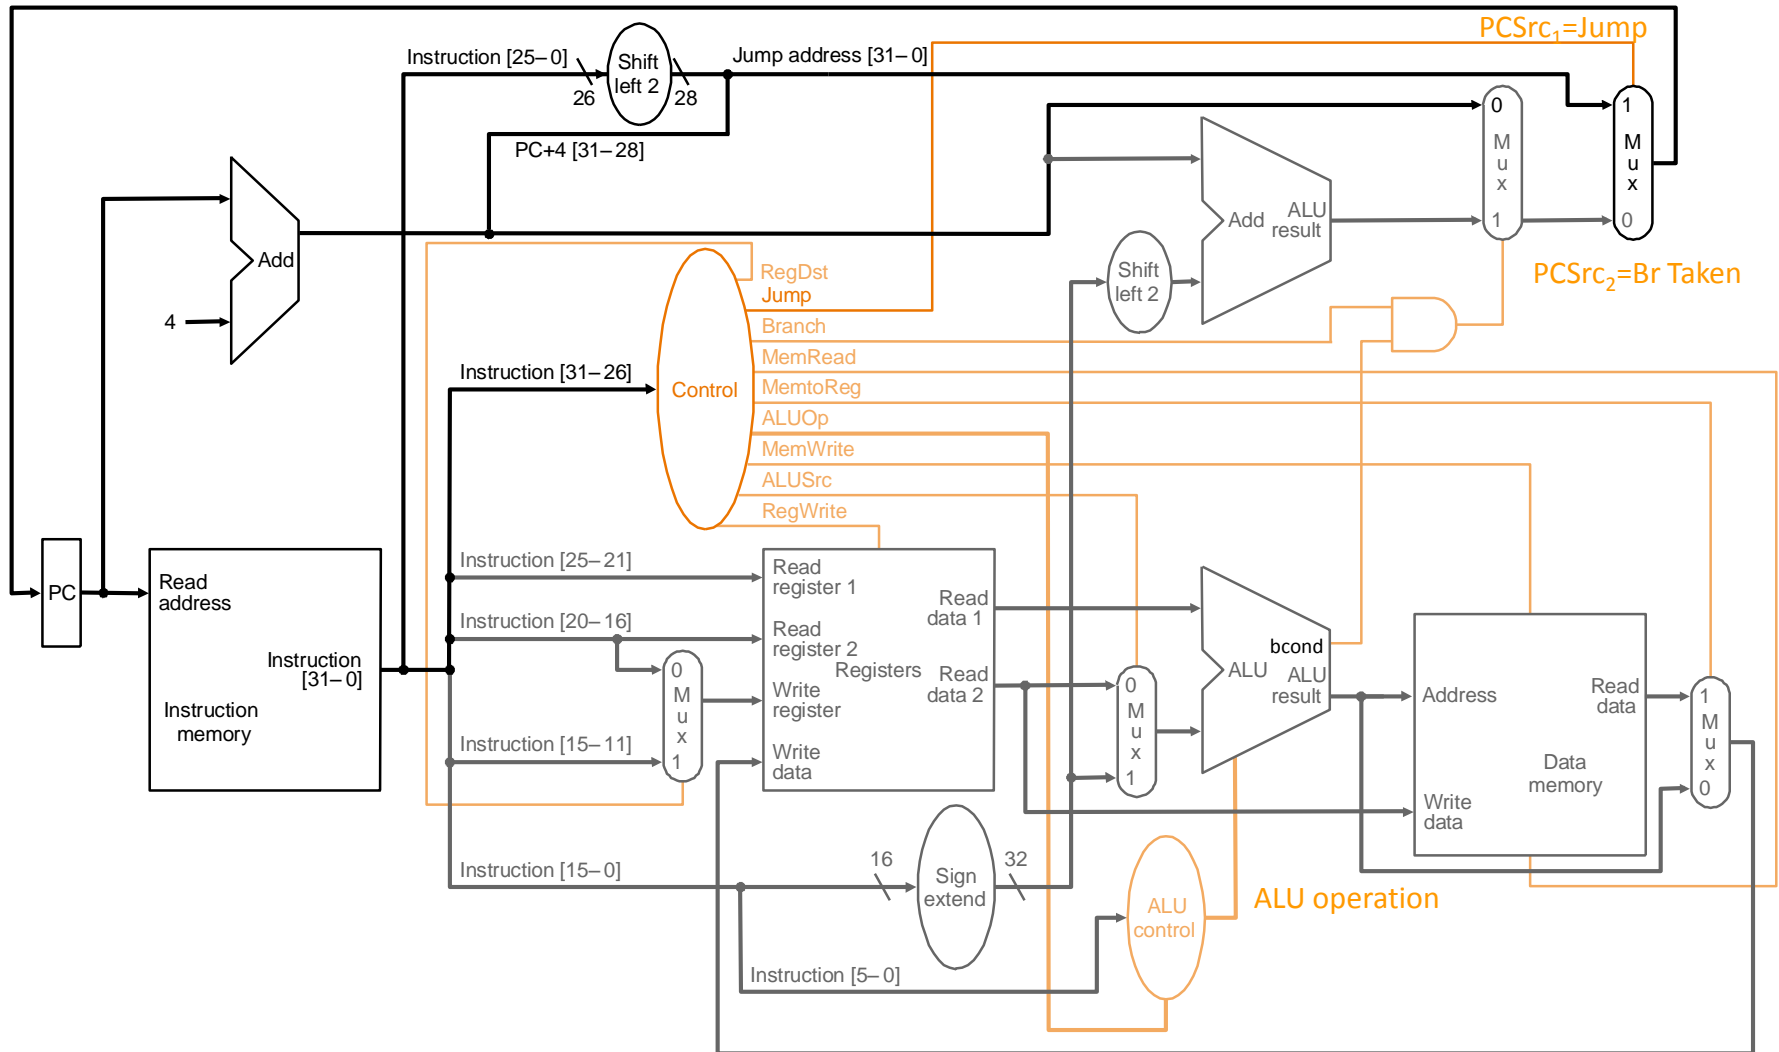
- “Magic” memory and register file
- Combinational read
  - output of the read data port is a combinational function of the register file contents and the corresponding read select port
- Synchronous write
  - the selected register is updated on the positive edge clock transition when write enable is asserted
    - Cannot affect read output in between clock edges
- Single-cycle, synchronous memory
  - Contrast this with memory that tells when the data is ready
  - i.e., Ready bit: indicating the read or write is done

# Instruction Processing

- 5 generic steps (P&H)
  - ❑ Instruction fetch (IF)
  - ❑ Instruction decode and register operand fetch (ID/RF)
  - ❑ Execute/Evaluate memory address (EX/AG)
  - ❑ Memory operand fetch (MEM)
  - ❑ Store/writeback result (WB)



# What Is To Come: The Full MIPS Datapath



\*\*Based on original figure from [P&H CO&D, COPYRIGHT 2004 Elsevier. ALL RIGHTS RESERVED.]



# Single-Cycle Datapath for *Arithmetic and Logical Instructions*


# R-Type ALU Instructions

---

- Assembly (e.g., register-register signed addition)

ADD rd<sub>reg</sub> rs<sub>reg</sub> rt<sub>reg</sub>

- Machine encoding



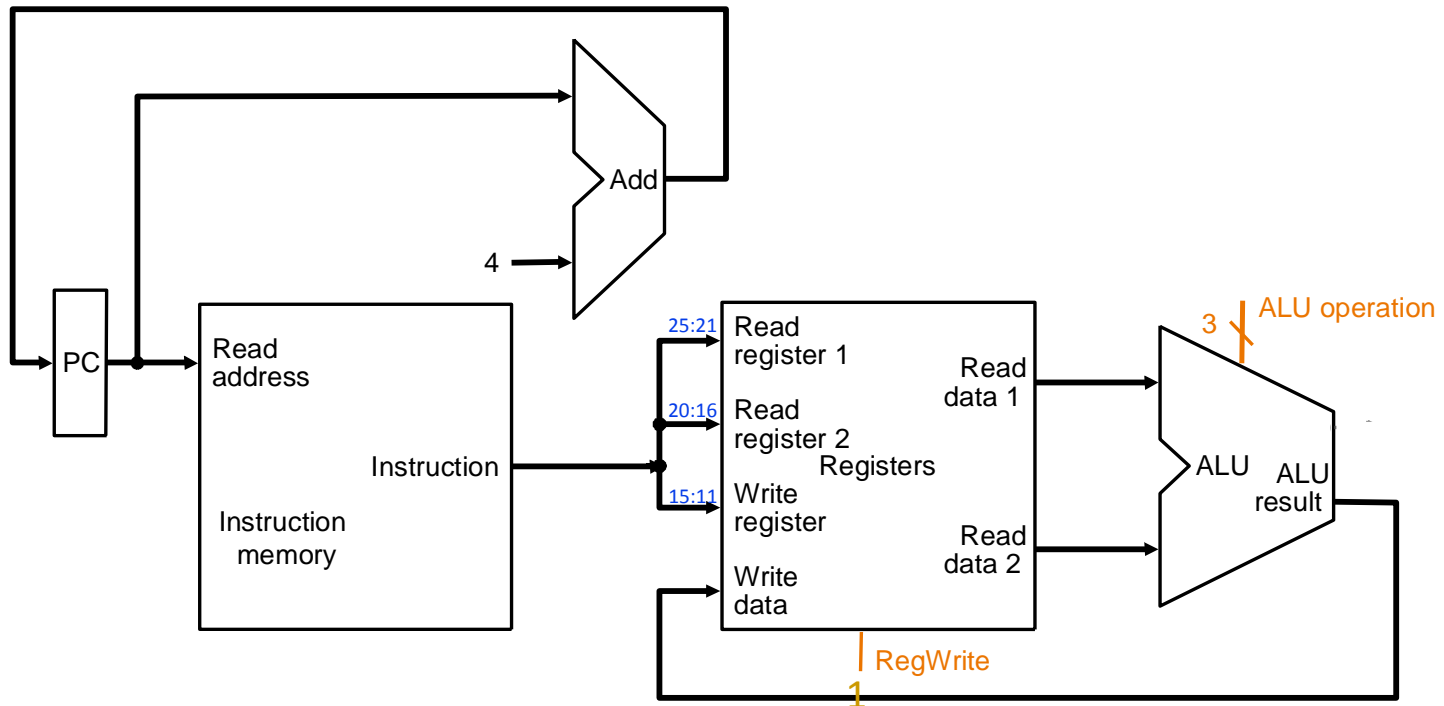
- Semantics

if MEM[PC] == ADD rd rs rt

GPR[rd] ← GPR[rs] + GPR[rt]

PC ← PC + 4

# ALU Datapath



Combinational  
state update logic

if MEM[PC] == ADD rd rs rt  
 $GPR[rd] \leftarrow GPR[rs] + GPR[rt]$   
 $PC \leftarrow PC + 4$


# I-Type ALU Instructions

---

- Assembly (e.g., register-immediate signed additions)

ADDI  $rt_{reg}$   $rs_{reg}$   $immediate_{16}$

- Machine encoding



I-type

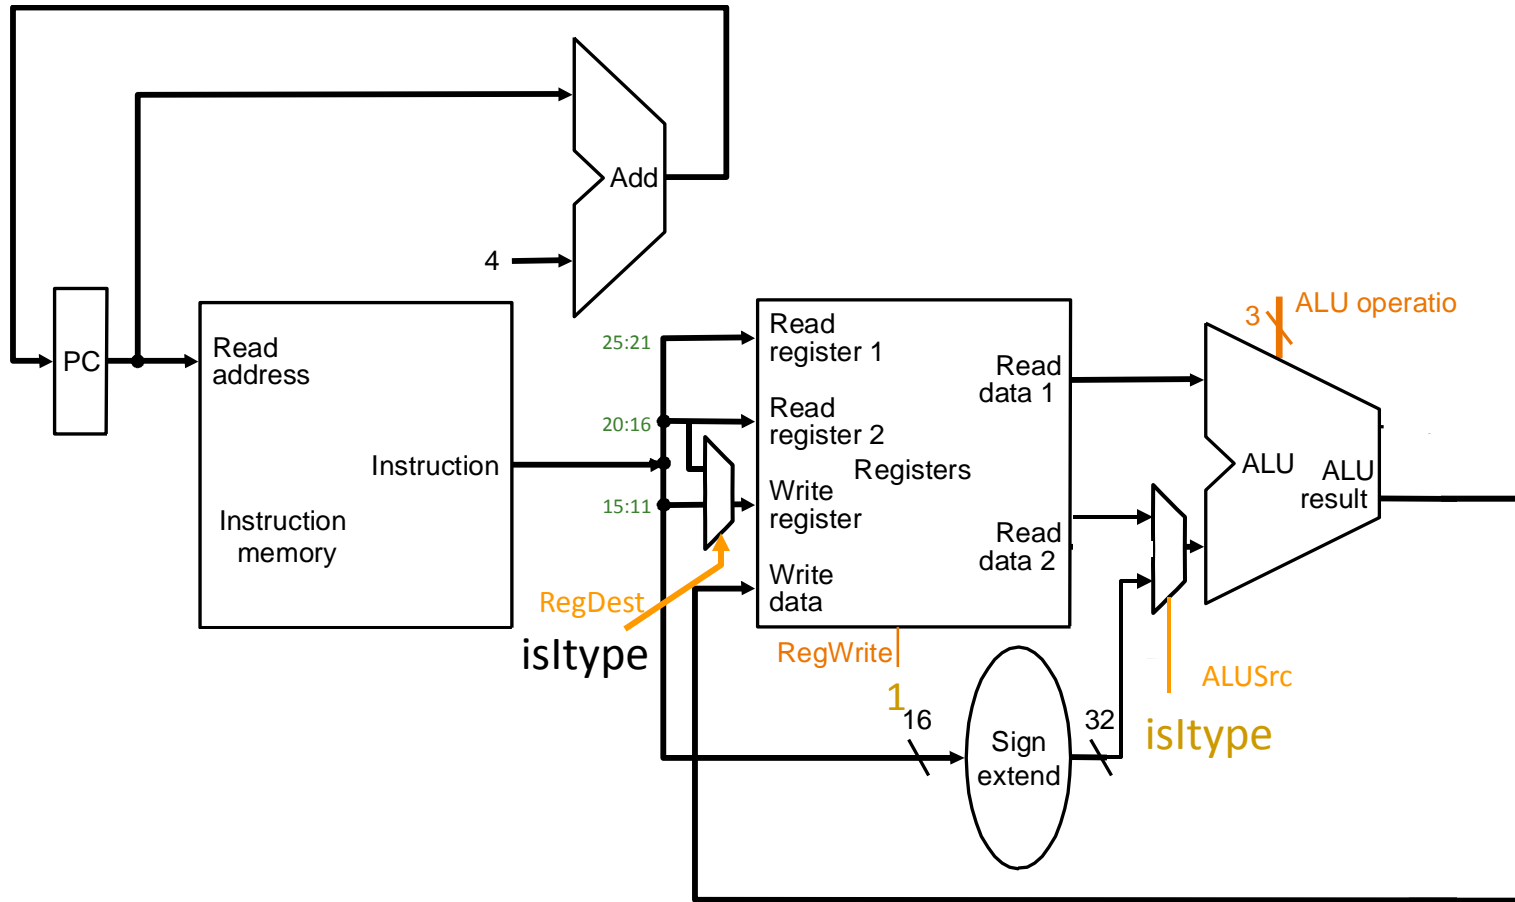
- Semantics

if MEM[PC] == ADDI  $rt$   $rs$   $immediate$

$GPR[rt] \leftarrow GPR[rs] + \text{sign-extend}(immediate)$

$PC \leftarrow PC + 4$

# Datapath for R and I-Type ALU Insts.



if MEM[PC] == ADDI rt rs immediate  
 $GPR[rt] \leftarrow GPR[rs] + \text{sign-extend}(\text{immediate})$   
 $PC \leftarrow PC + 4$

Combinational  
 state update logic

# Single-Cycle Datapath for *Data Movement Instructions*


# Load Instructions

---

- Assembly (e.g., load 4-byte word)

LW  $rt_{reg}$   $offset_{16}$  ( $base_{reg}$ )

- Machine encoding



I-type

- Semantics

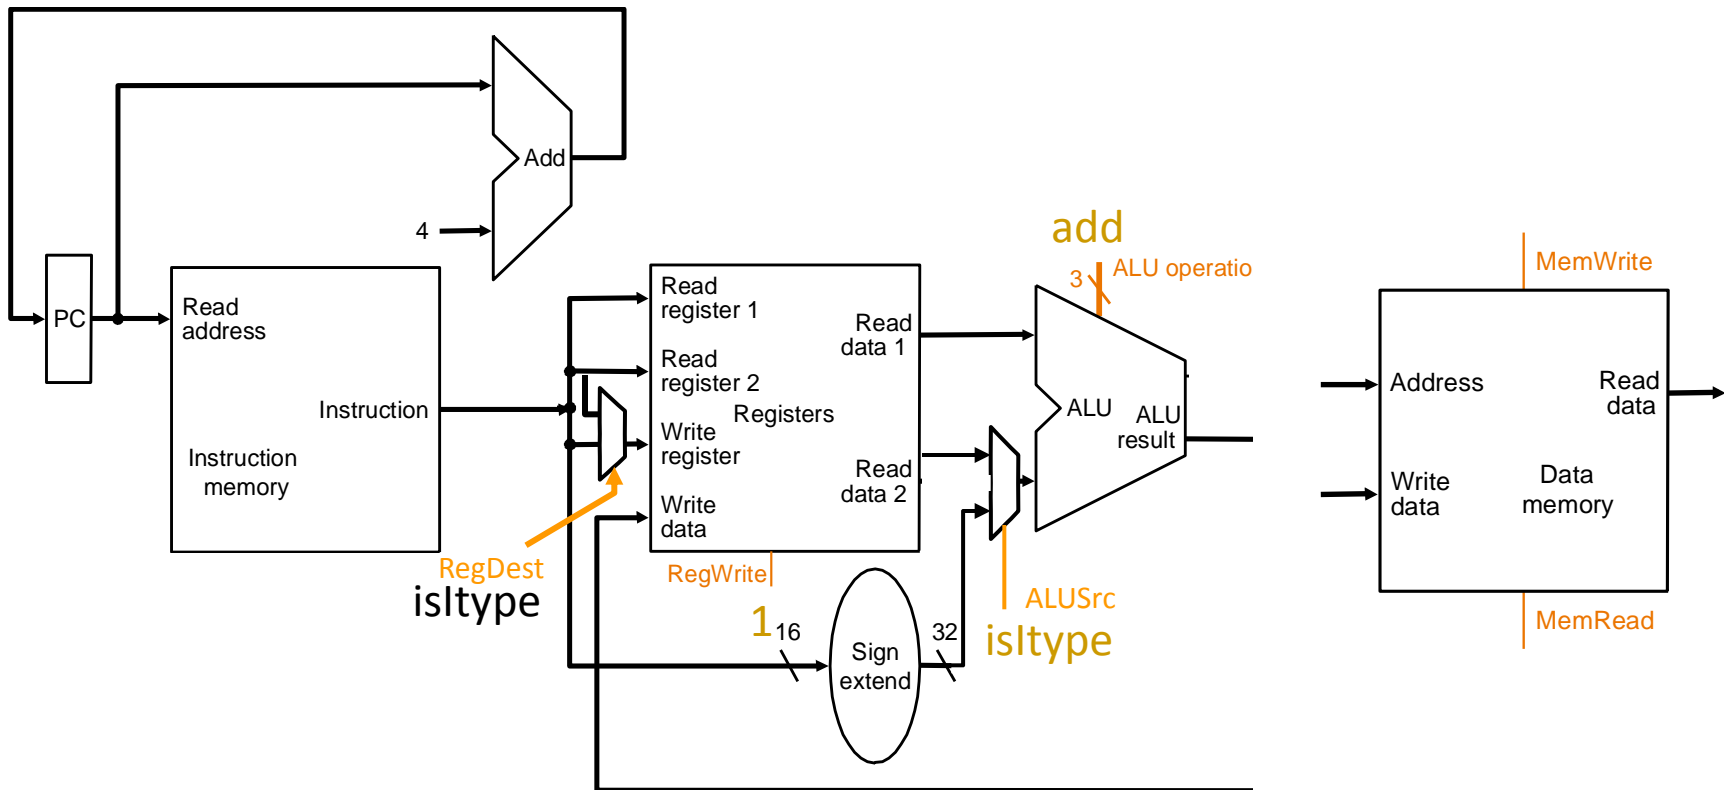
if  $MEM[PC] == LW\ rt\ offset_{16}\ (base)$

EA = sign-extend(offset) + GPR[base]

GPR[rt] ← MEM[ translate(EA) ]

PC ← PC + 4

# LW Datapath



if  $MEM[PC] == LW \text{ rt offset}_{16} \text{ (base)}$   
 $EA = \text{sign-extend}(\text{offset}) + GPR[\text{base}]$   
 $GPR[\text{rt}] \leftarrow MEM[\text{translate}(EA)]$   
 $PC \leftarrow PC + 4$



Combinational  
state update logic 16




# Store Instructions

---

- Assembly (e.g., store 4-byte word)

$SW\ rt_{reg}\ offset_{16}\ (base_{reg})$

- Machine encoding



I-type

- Semantics

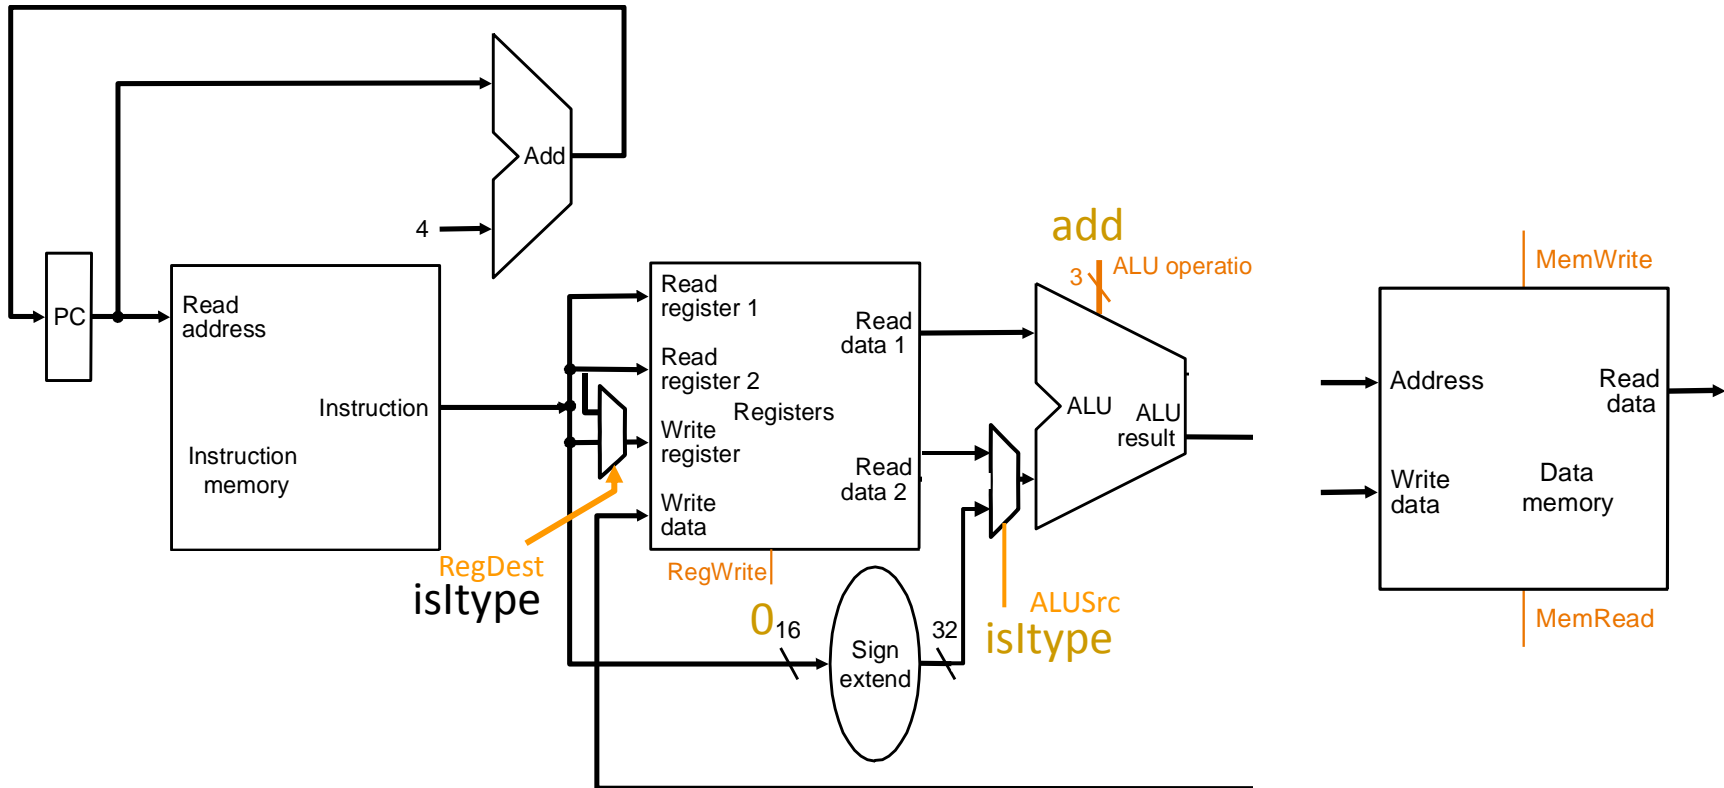
if  $MEM[PC] == SW\ rt\ offset_{16}\ (base)$

$EA = \text{sign-extend}(\text{offset}) + GPR[\text{base}]$

$MEM[\text{translate}(EA)] \leftarrow GPR[rt]$

$PC \leftarrow PC + 4$

# SW Datapath



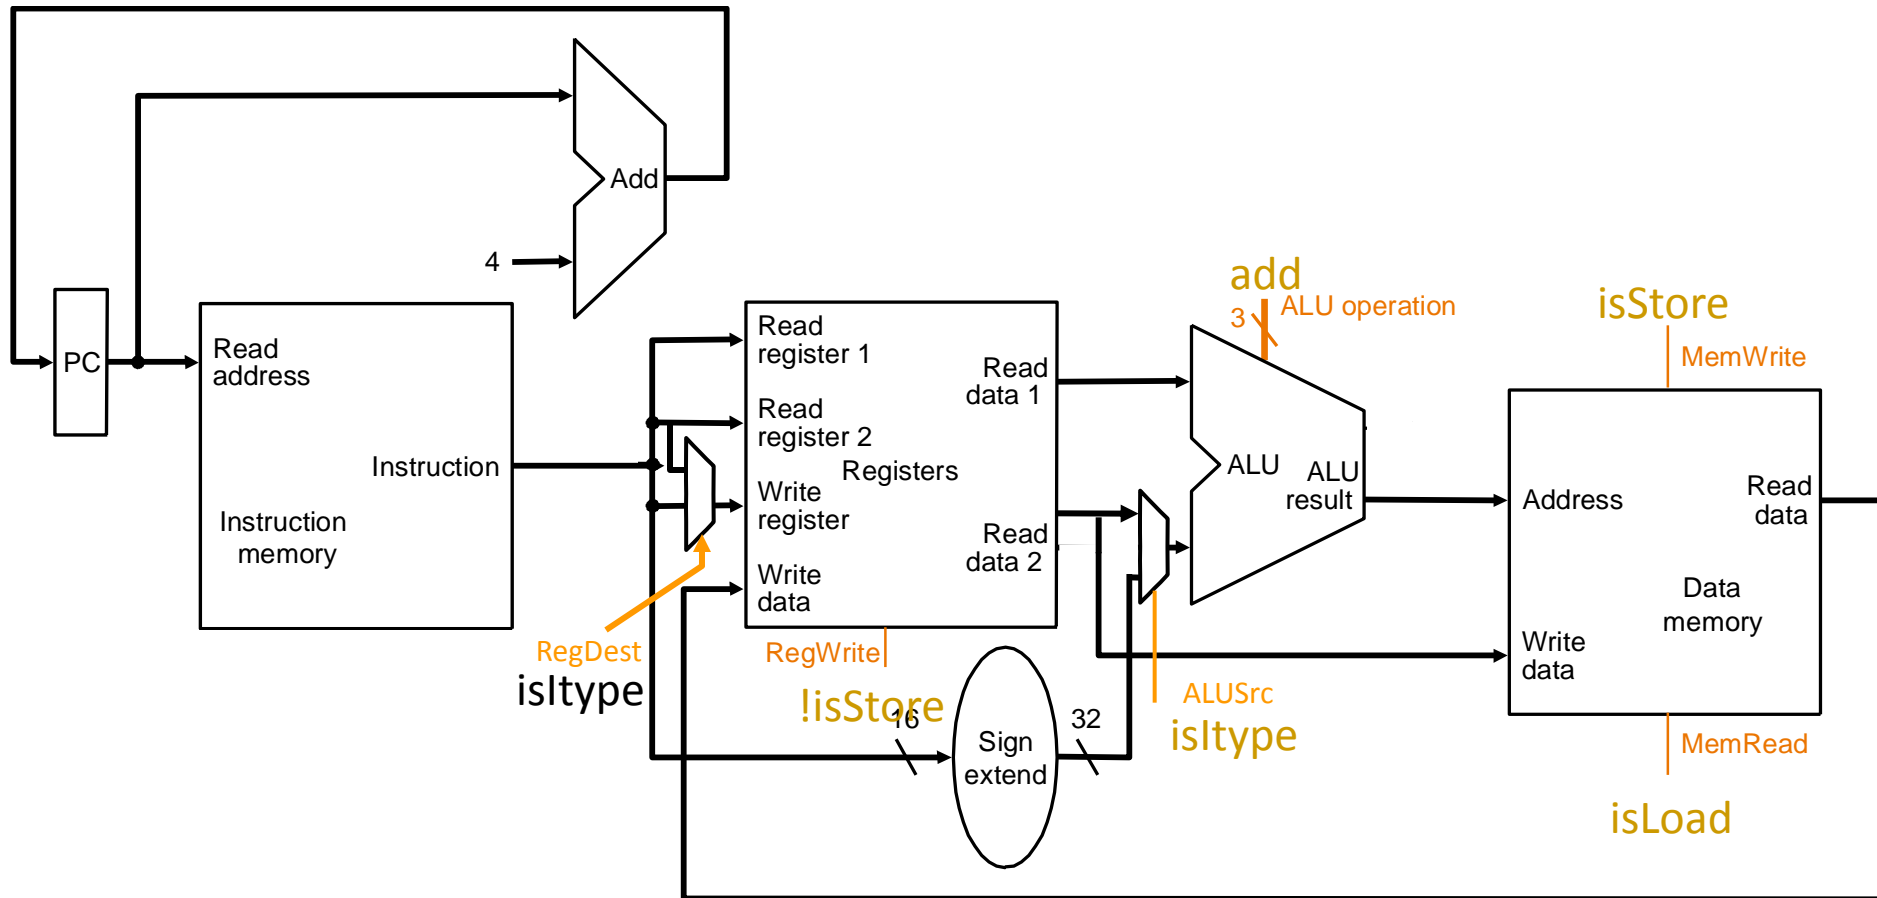
if  $MEM[PC] == SW\ rt\ offset_{16}\ (base)$   
 $EA = \text{sign-extend}(\text{offset}) + GPR[\text{base}]$   
 $MEM[\text{translate}(EA)] \leftarrow GPR[\text{rt}]$   
 $PC \leftarrow PC + 4$

IF	ID	EX	MEM	WB
----	----	----	-----	----

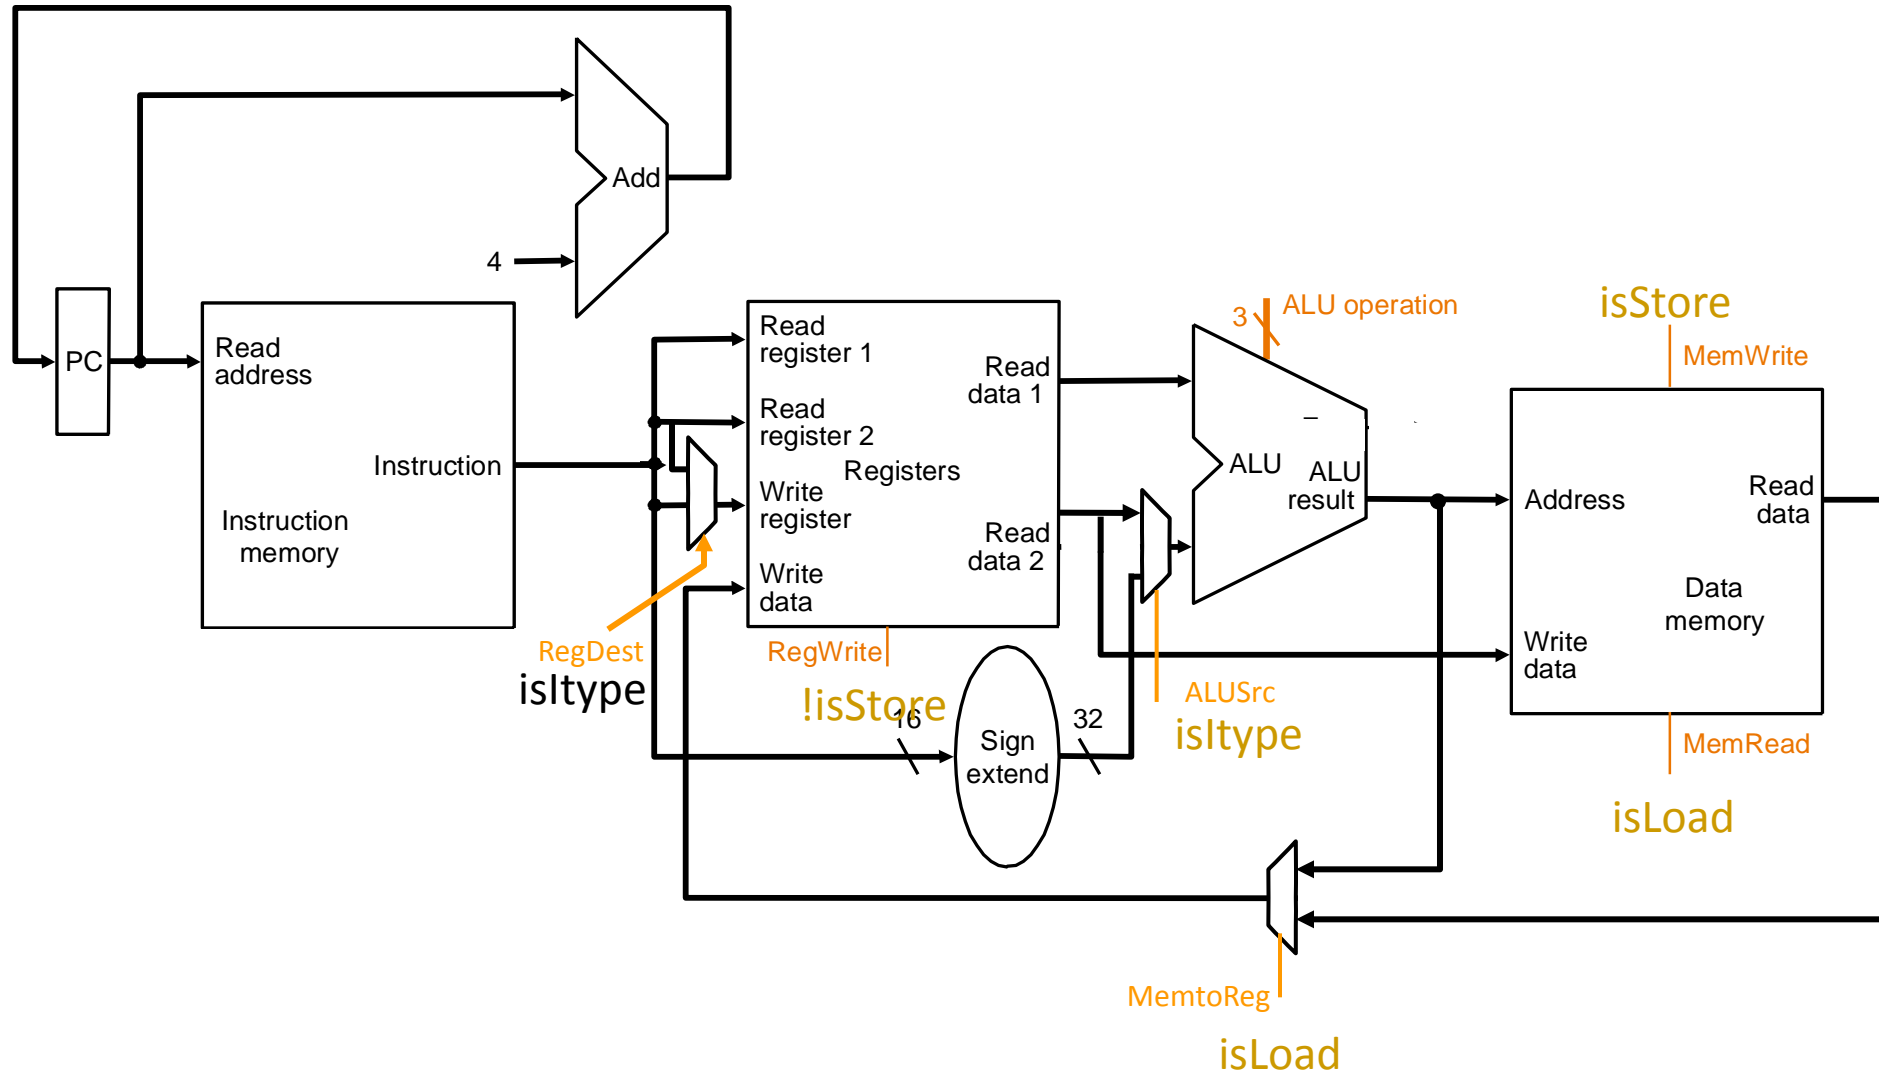


Combinational  
state update logic 18

# Load-Store Datapath



# Datapath for Non-Control-Flow Insts.



# Single-Cycle Datapath for *Control Flow Instructions*


# Unconditional Jump Instructions

---

- Assembly

J immediate<sub>26</sub>

- Machine encoding



J-type


- Semantics

if MEM[PC] == J immediate<sub>26</sub>

target = { PC[31:28], immediate<sub>26</sub>, 2' b00 }

PC ← target

# Unconditional Jump Datapath



\*\*Based on original figure from [P&H CO&D, COPYRIGHT 2004 Elsevier. ALL RIGHTS RESERVED.]

```

if MEM[PC]==J immediate26
    PC = { PC[31:28], immediate26, 2' b00 }
    
```

What about JR, JAL, JALR?


# Conditional Branch Instructions

---

- Assembly (e.g., branch if equal)

BEQ  $rs_{reg}$   $rt_{reg}$   $immediate_{16}$

- Machine encoding



I-type

- Semantics (assuming no branch delay slot)

if  $MEM[PC] == BEQ\ rs\ rt\ immediate_{16}$

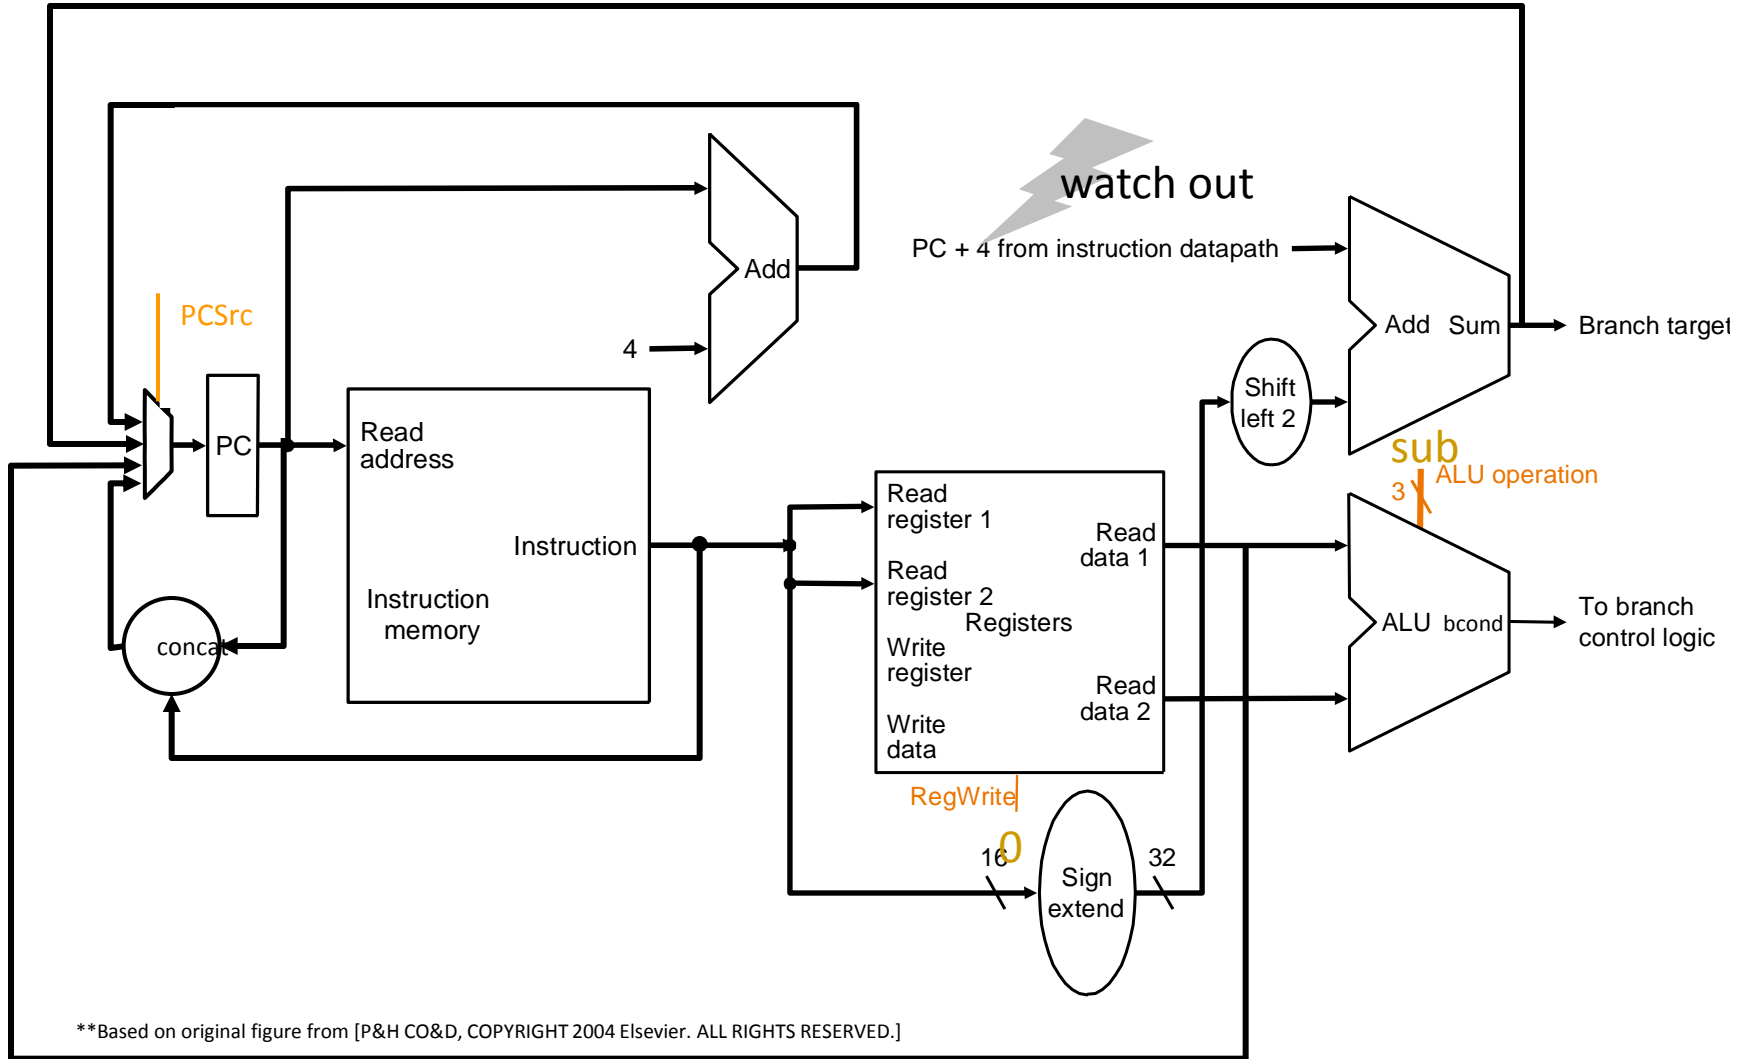
target =  $PC + 4 + \text{sign-extend}(immediate) \times 4$

if  $GPR[rs] == GPR[rt]$  then  $PC \leftarrow \text{target}$

else  $PC \leftarrow PC + 4$

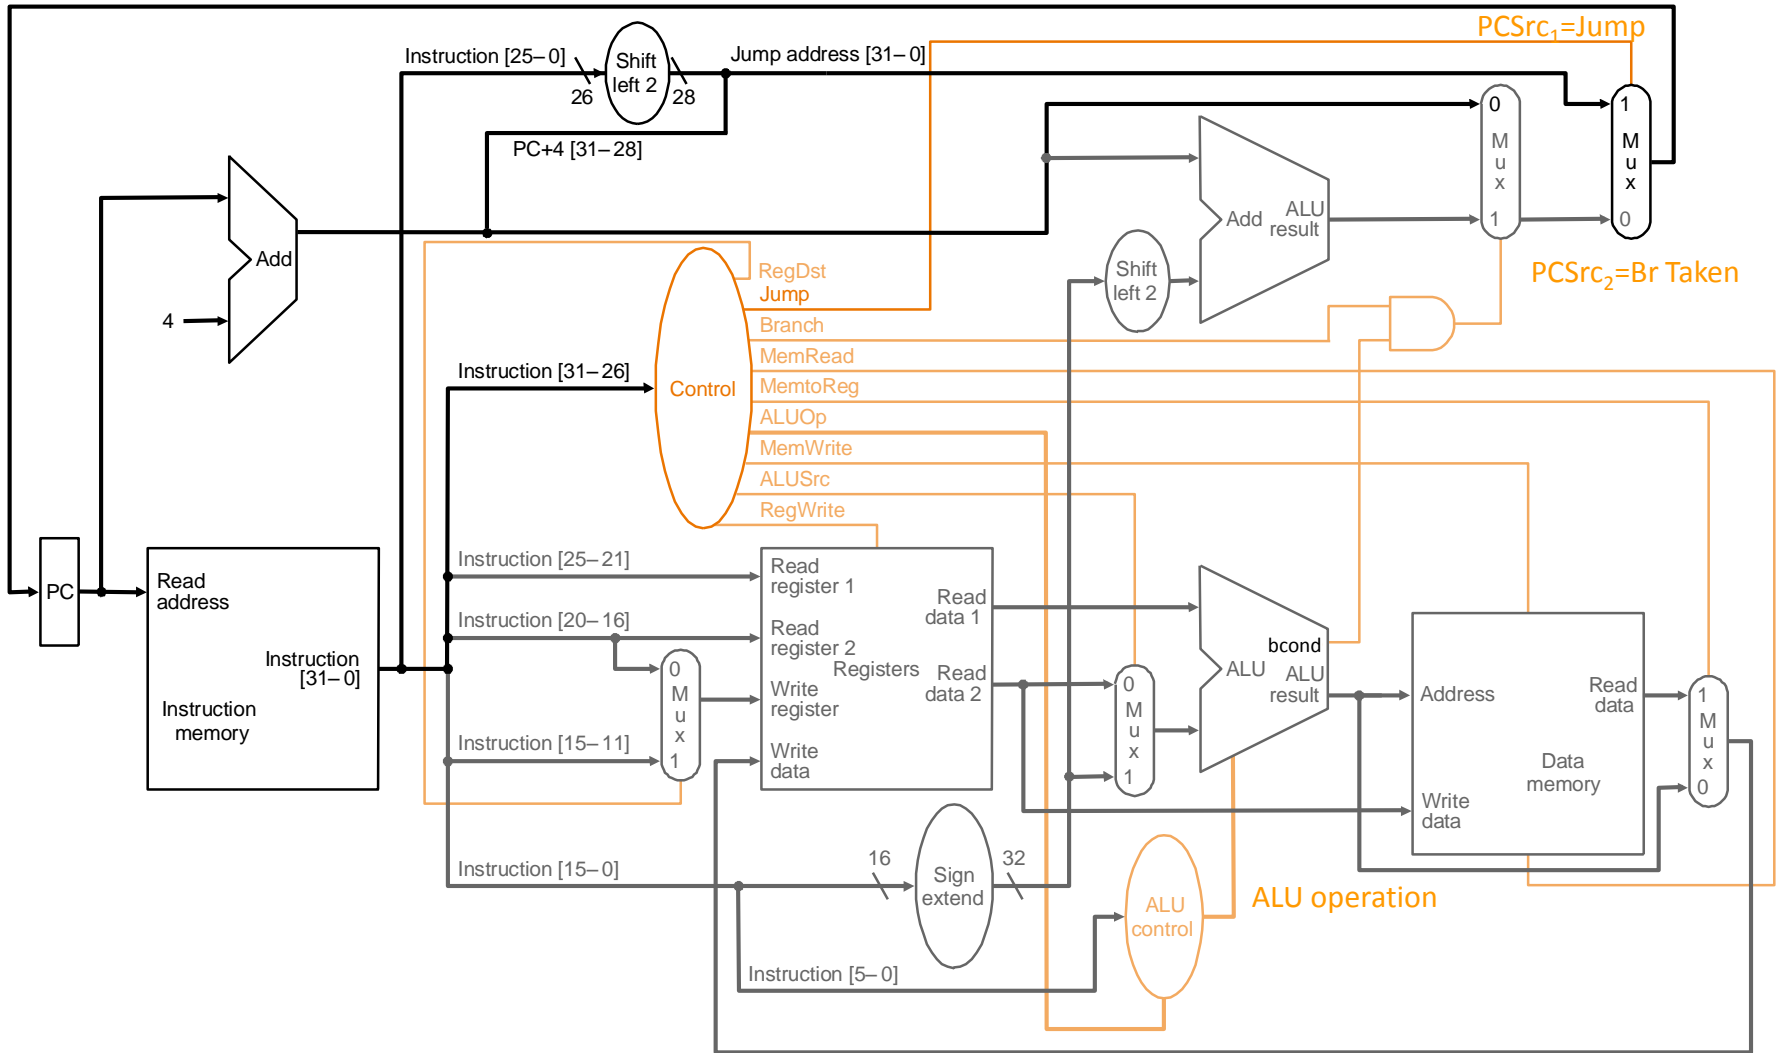


# Conditional Branch Datapath (For You to Fix)



How to uphold the delayed branch semantics?

# Putting It All Together

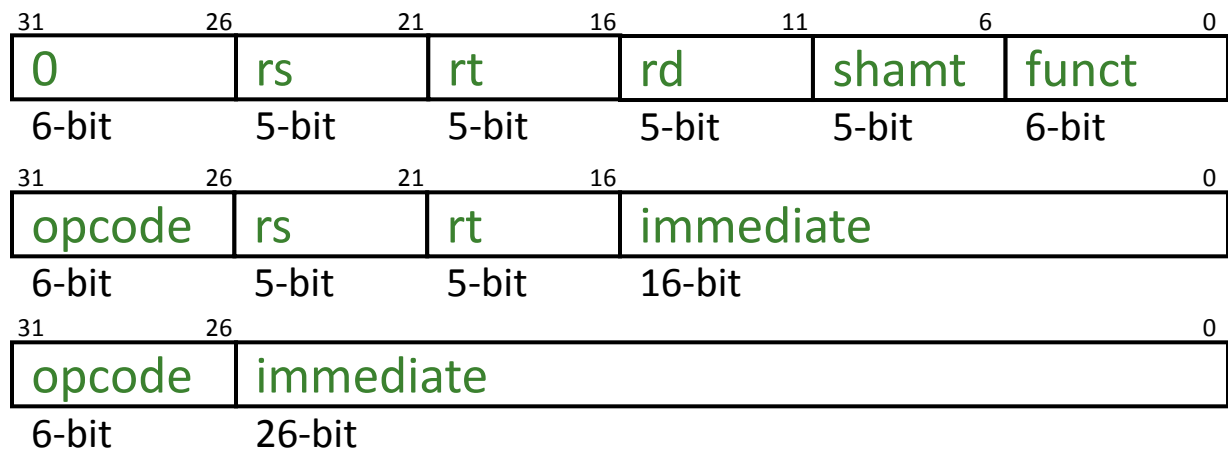


\*\*Based on original figure from [P&H CO&D, COPYRIGHT 2004 Elsevier. ALL RIGHTS RESERVED.]

# Single-Cycle Control Logic

# Single-Cycle Hardwired Control

- As combinational function of  $\text{Inst} = \text{MEM}[\text{PC}]$



R-type

I-type

J-type

- Consider
  - All R-type and I-type ALU instructions
  - LW and SW
  - BEQ, BNE, BLEZ, BGTZ
  - J, JR, JAL, JALR

# Single-Bit Control Signals

---

	When De-asserted	When asserted	Equation
RegDest	GPR write select according to <code>rt</code> , i.e., <code>inst[20:16]</code>	GPR write select according to <code>rd</code> , i.e., <code>inst[15:11]</code>	<code>opcode==0</code>
ALUSrc	2 <sup>nd</sup> ALU input from 2 <sup>nd</sup> GPR read port	2 <sup>nd</sup> ALU input from sign-extended 16-bit immediate	<code>(opcode!=0) &amp;&amp;</code> <code>(opcode!=BEQ) &amp;&amp;</code> <code>(opcode!=BNE)</code>
MemtoReg	Steer ALU result to GPR write port	steer memory load to GPR wr. port	<code>opcode==LW</code>
RegWrite	GPR write disabled	GPR write enabled	<code>(opcode!=SW) &amp;&amp;</code> <code>(opcode!=Bxx) &amp;&amp;</code> <code>(opcode!=J) &amp;&amp;</code> <code>(opcode!=JR))</code>

# Single-Bit Control Signals

---

	When De-asserted	When asserted	Equation
MemRead	Memory read disabled	Memory read port return load value	<code>opcode==LW</code>
MemWrite	Memory write disabled	Memory write enabled	<code>opcode==SW</code>
PCSrc <sub>1</sub>	According to PCSrc <sub>2</sub>	next PC is based on 26-bit immediate jump target	<code>(opcode==J)    (opcode==JAL)</code>
PCSrc <sub>2</sub>	next PC = PC + 4	next PC is based on 16-bit immediate branch target	<code>(opcode==Bxx) &amp;&amp; "bcond is satisfied"</code>

# ALU Control

---

- case opcode

- '0' ⇒ select operation according to **funct**

- 'ALUi' ⇒ selection operation according to **opcode**

- 'LW' ⇒ select addition

- 'SW' ⇒ select addition

- 'Bxx' ⇒ select bcond generation function

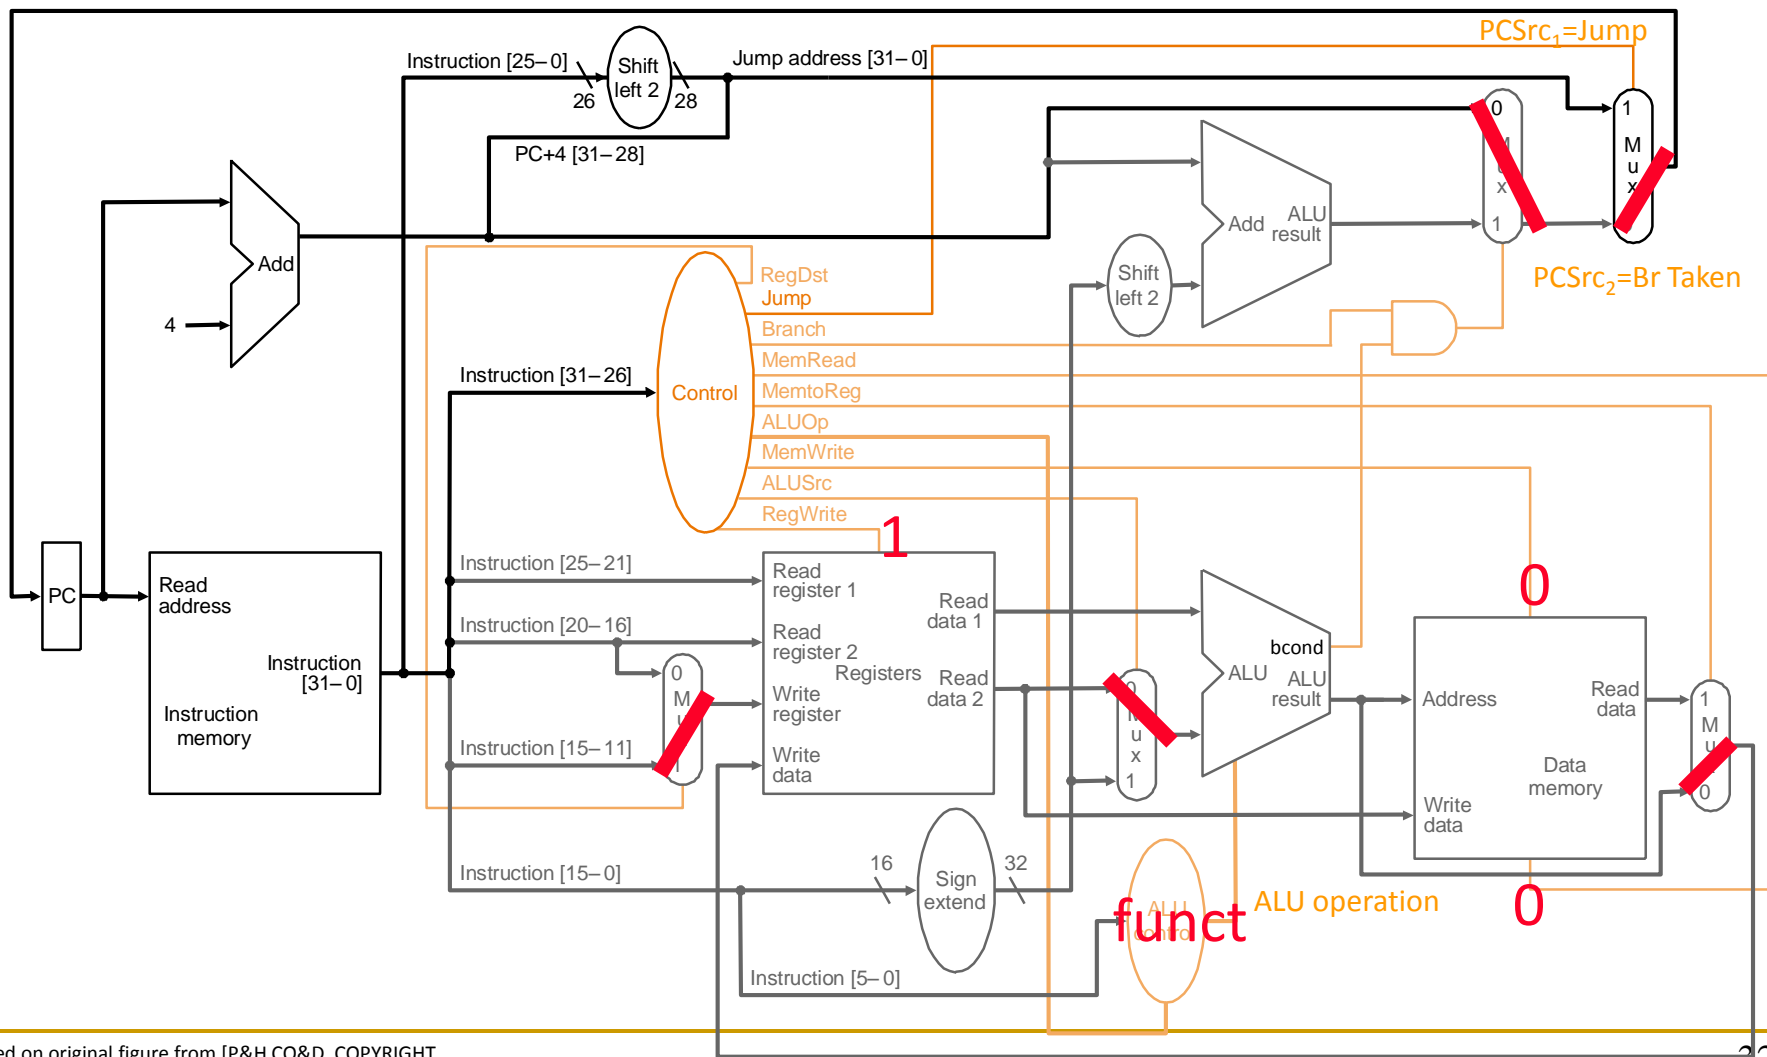
- \_\_\_ ⇒ don't care

- Example ALU operations

- ADD, SUB, AND, OR, XOR, NOR, etc.

- bcond on equal, not equal, LE zero, GT zero, etc.

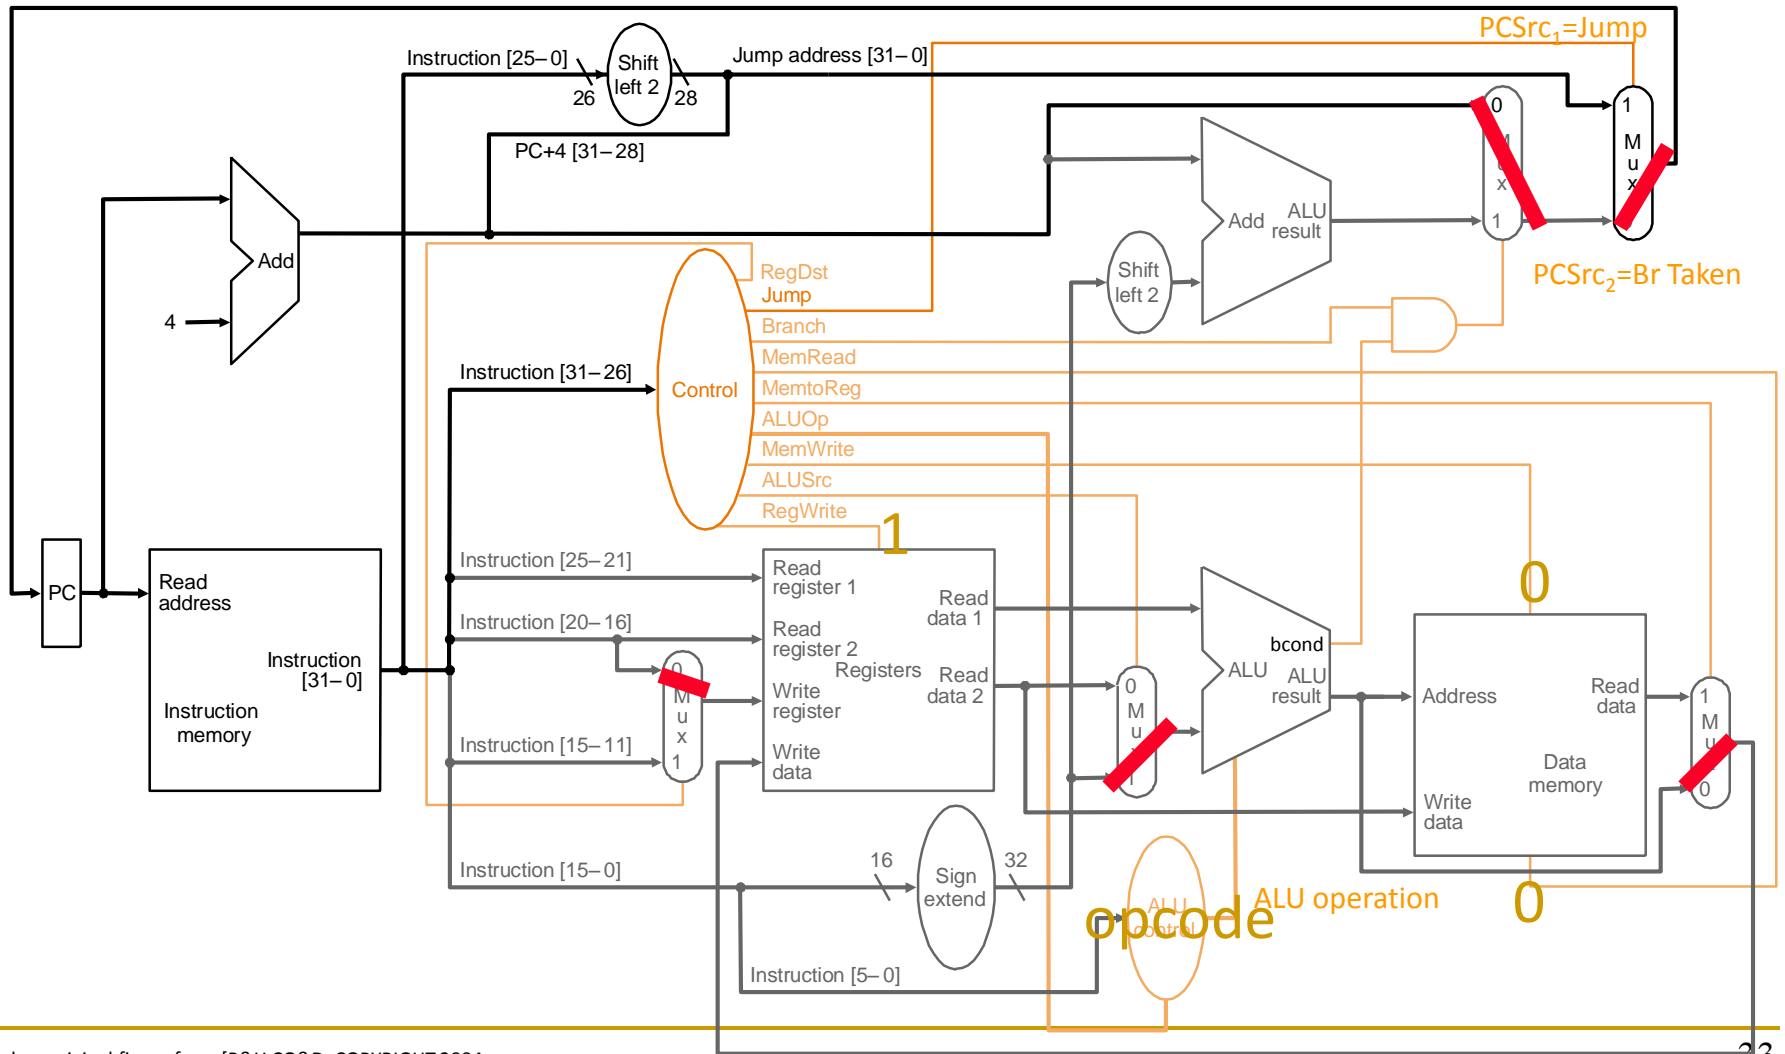
# R-Type ALU



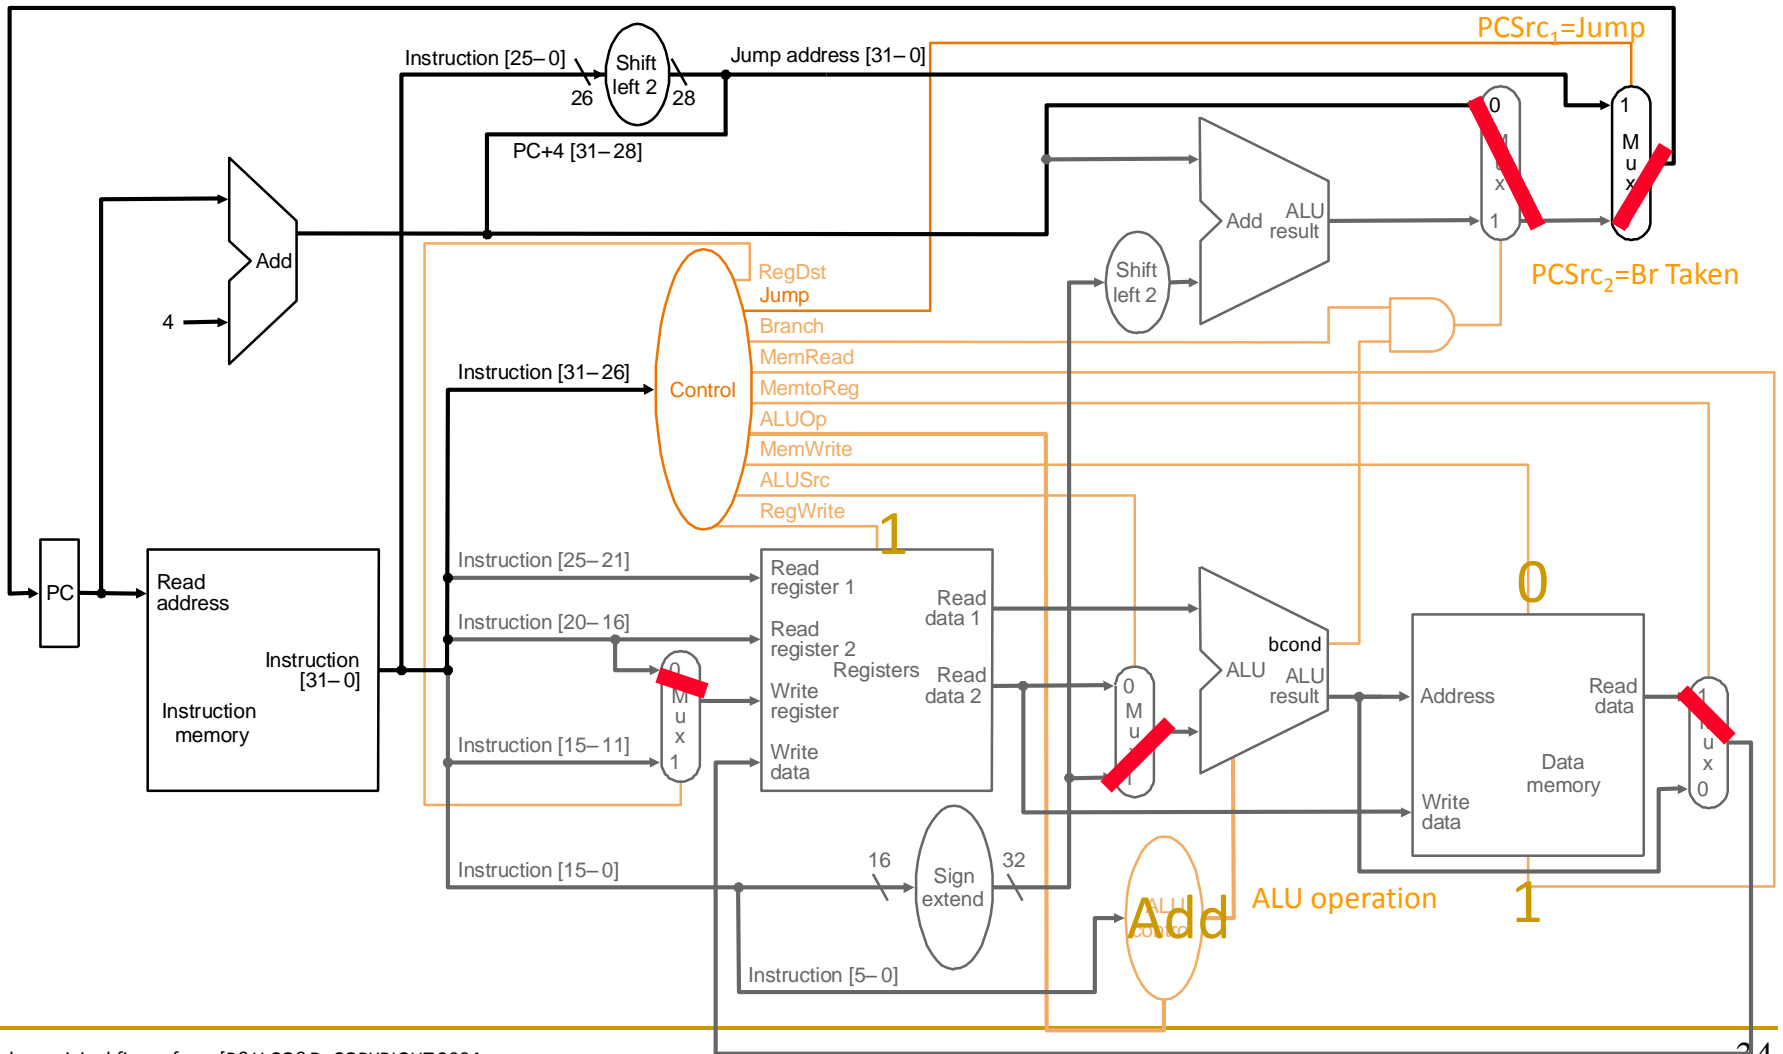
\*\*Based on original figure from [P&H CO&D, COPYRIGHT 2004 Elsevier. ALL RIGHTS RESERVED.]

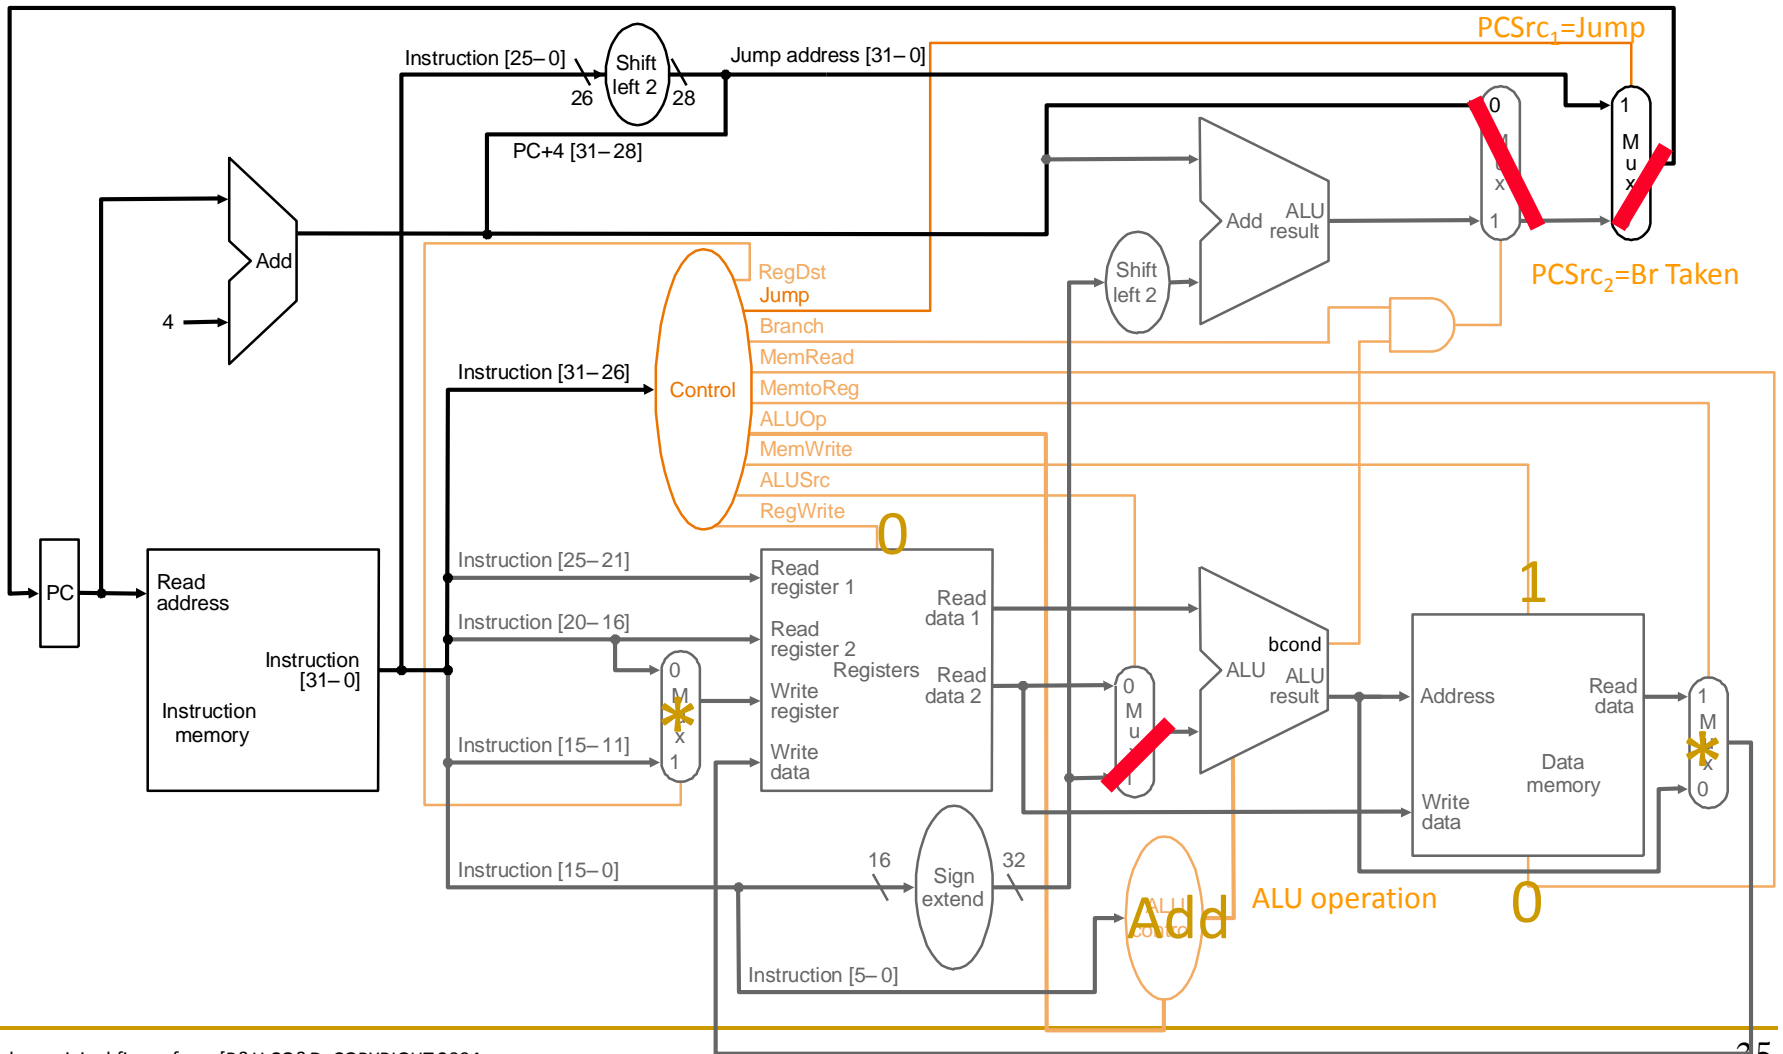


# I-Type ALU



\*\*Based on original figure from [P&H CO&D, COPYRIGHT 2004 Elsevier. ALL RIGHTS RESERVED.]

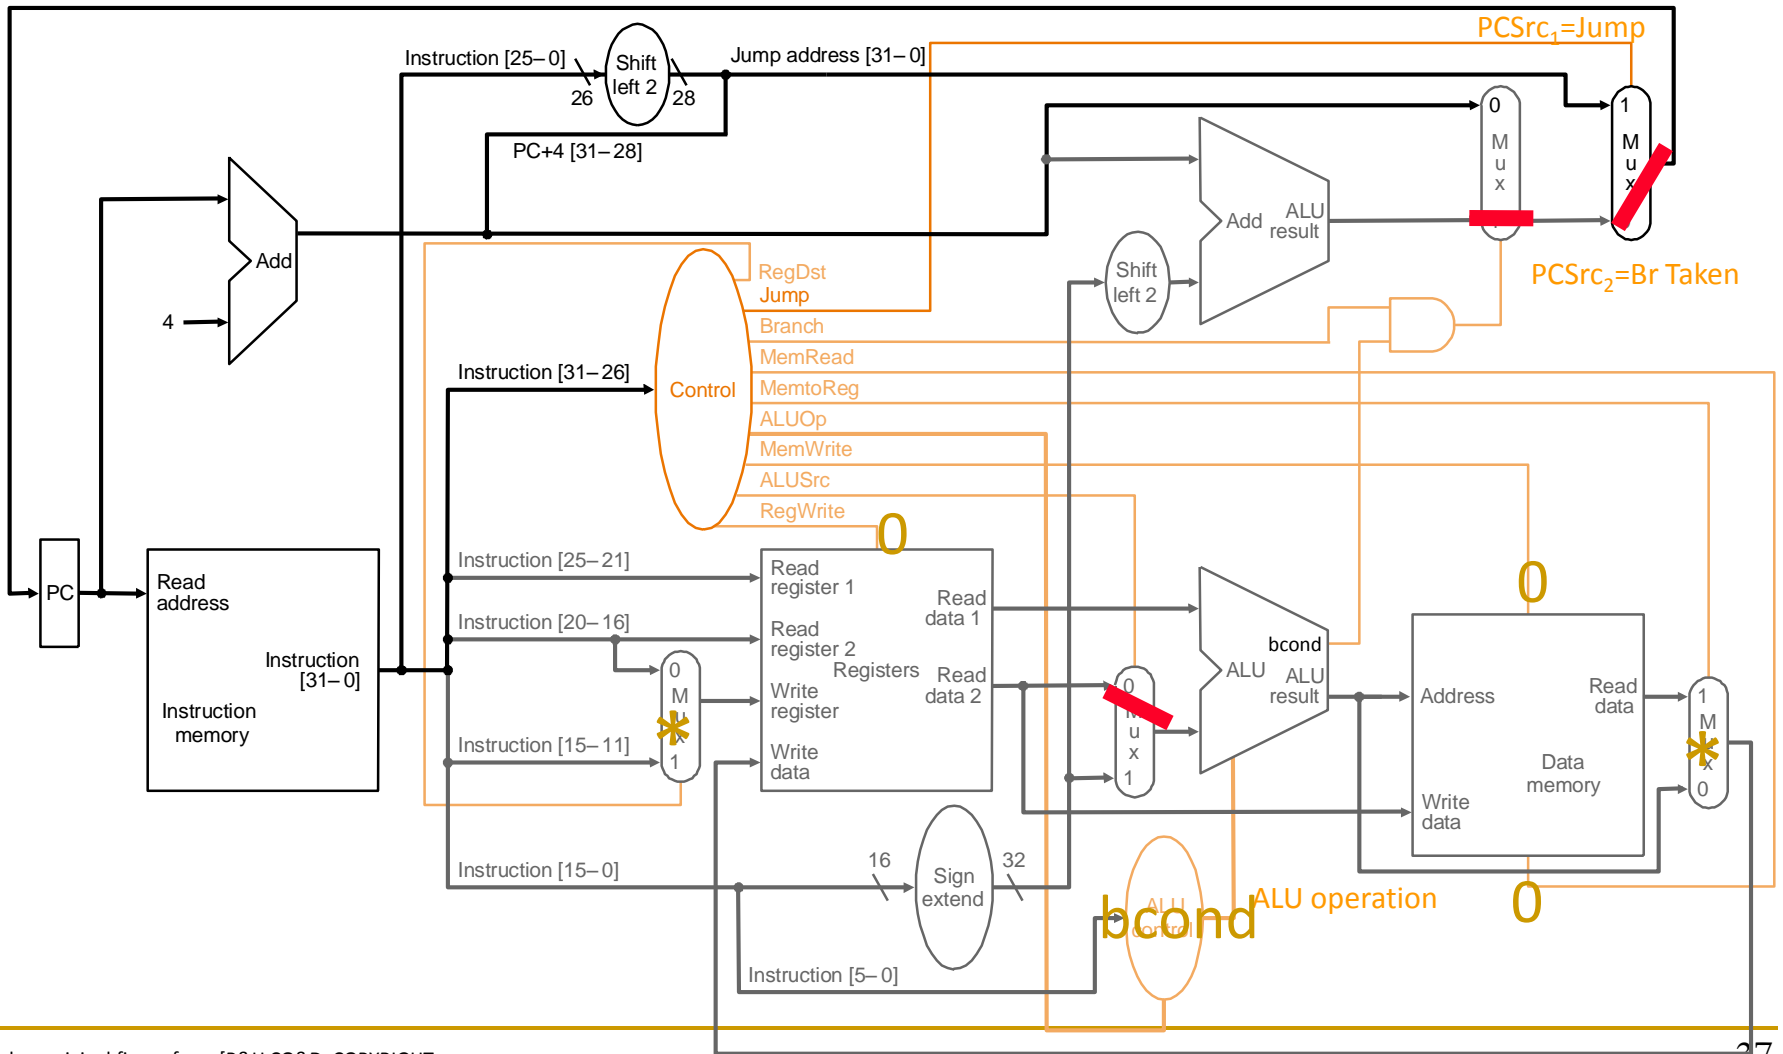




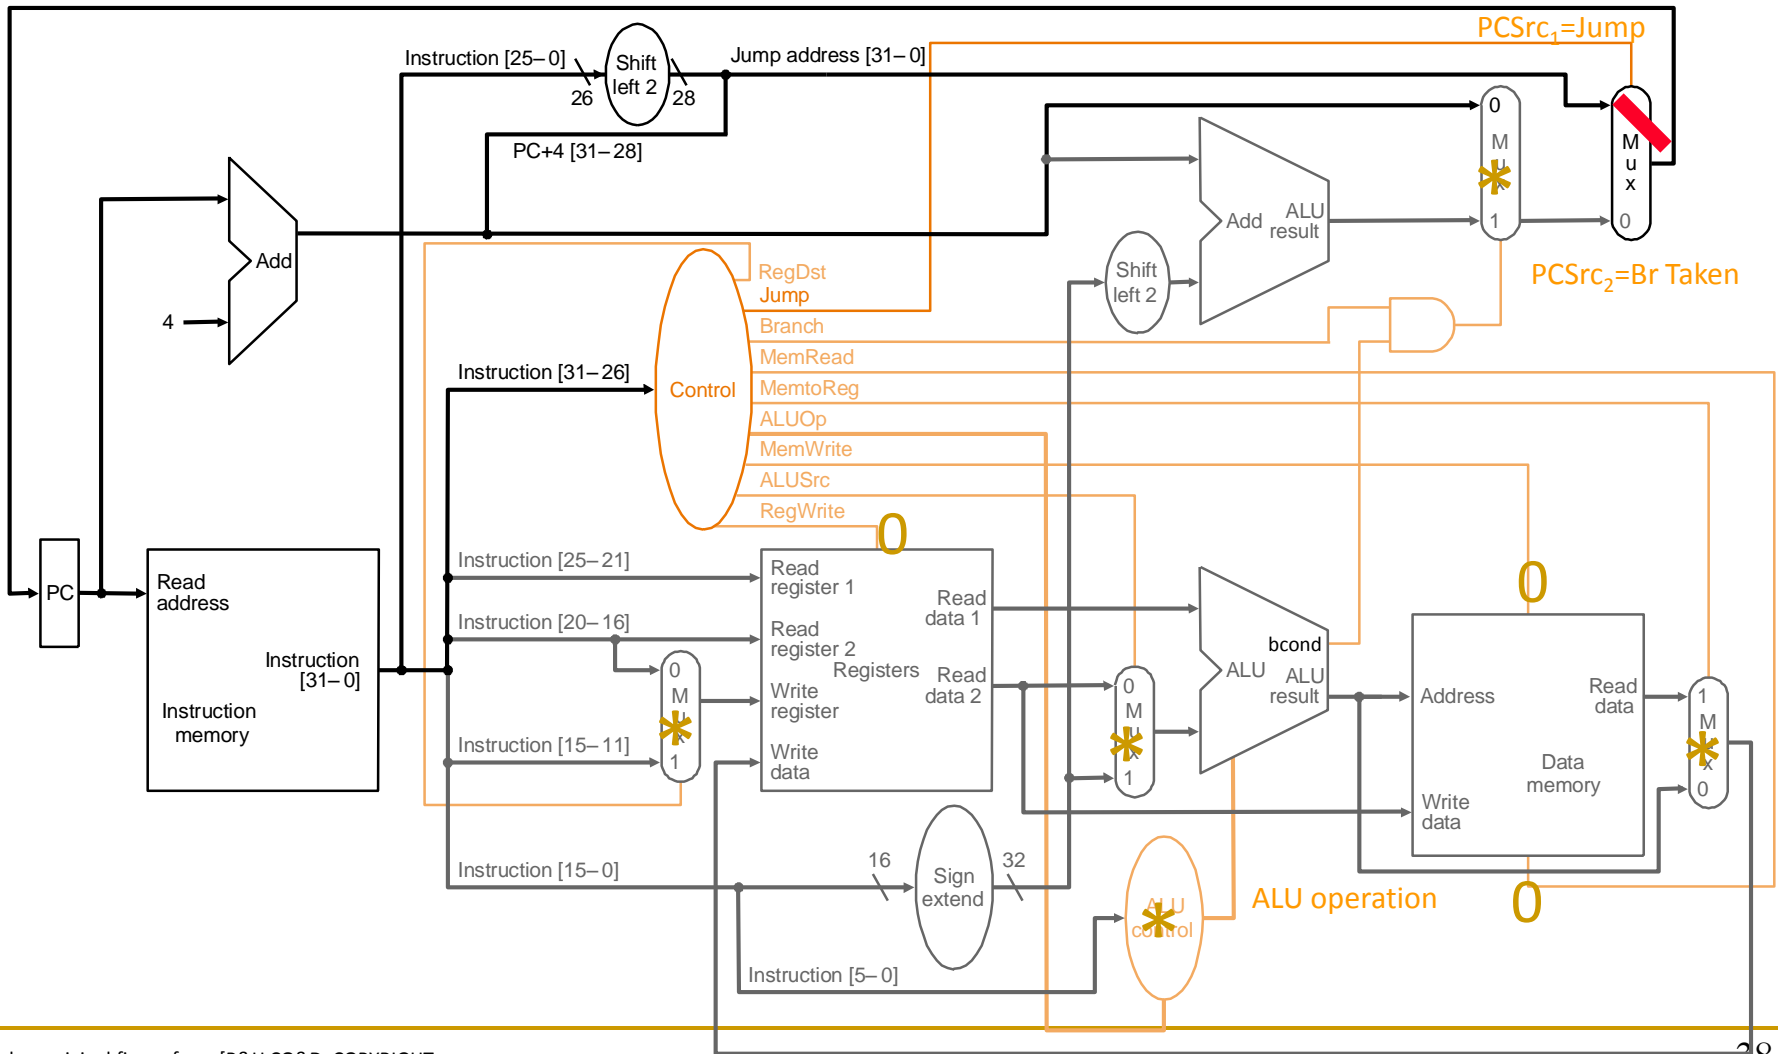


# Branch Taken

Some control signals are dependent on the processing of data



# Jump



# What is in That Control Box?

---

- Combinational Logic → **Hardwired Control**
  - Idea: Control signals generated combinatorially based on instruction
  - Necessary in a single-cycle microarchitecture...
- Sequential Logic → **Sequential/Microprogrammed Control**
  - Idea: A memory structure contains the control signals associated with an instruction
  - Control Store

# Evaluating the Single-Cycle Microarchitecture



# A Single-Cycle Microarchitecture

---

- Is *this* a good idea/design?
- When is this a good design?
- When is this a bad design?
- How can we design a better microarchitecture?

# A Single-Cycle Microarchitecture: Analysis

---

- Every instruction takes 1 cycle to execute
  - CPI (Cycles per instruction) is strictly 1
- How long each instruction takes is determined by how long the slowest instruction takes to execute
  - Even though many instructions do not need that long to execute
- Clock cycle time of the microarchitecture is determined by how long it takes to complete the slowest instruction
  - Critical path of the design is determined by the processing time of the slowest instruction

# What is the Slowest Instruction to Process?

---

- Let's go back to the basics
- All six phases of the instruction processing cycle take a *single machine clock cycle* to complete
  - Fetch
    - 1. Instruction fetch (IF)
  - Decode
    - 2. Instruction decode and
  - Evaluate Address
    - register operand fetch (ID/RF)
  - Fetch Operands
    - 3. Execute/Evaluate memory address (EX/AG)
  - Execute
    - 4. Memory operand fetch (MEM)
  - Store Result
    - 5. Store/writeback result (WB)
- Do each of the above phases take the same time (latency) for all instructions?

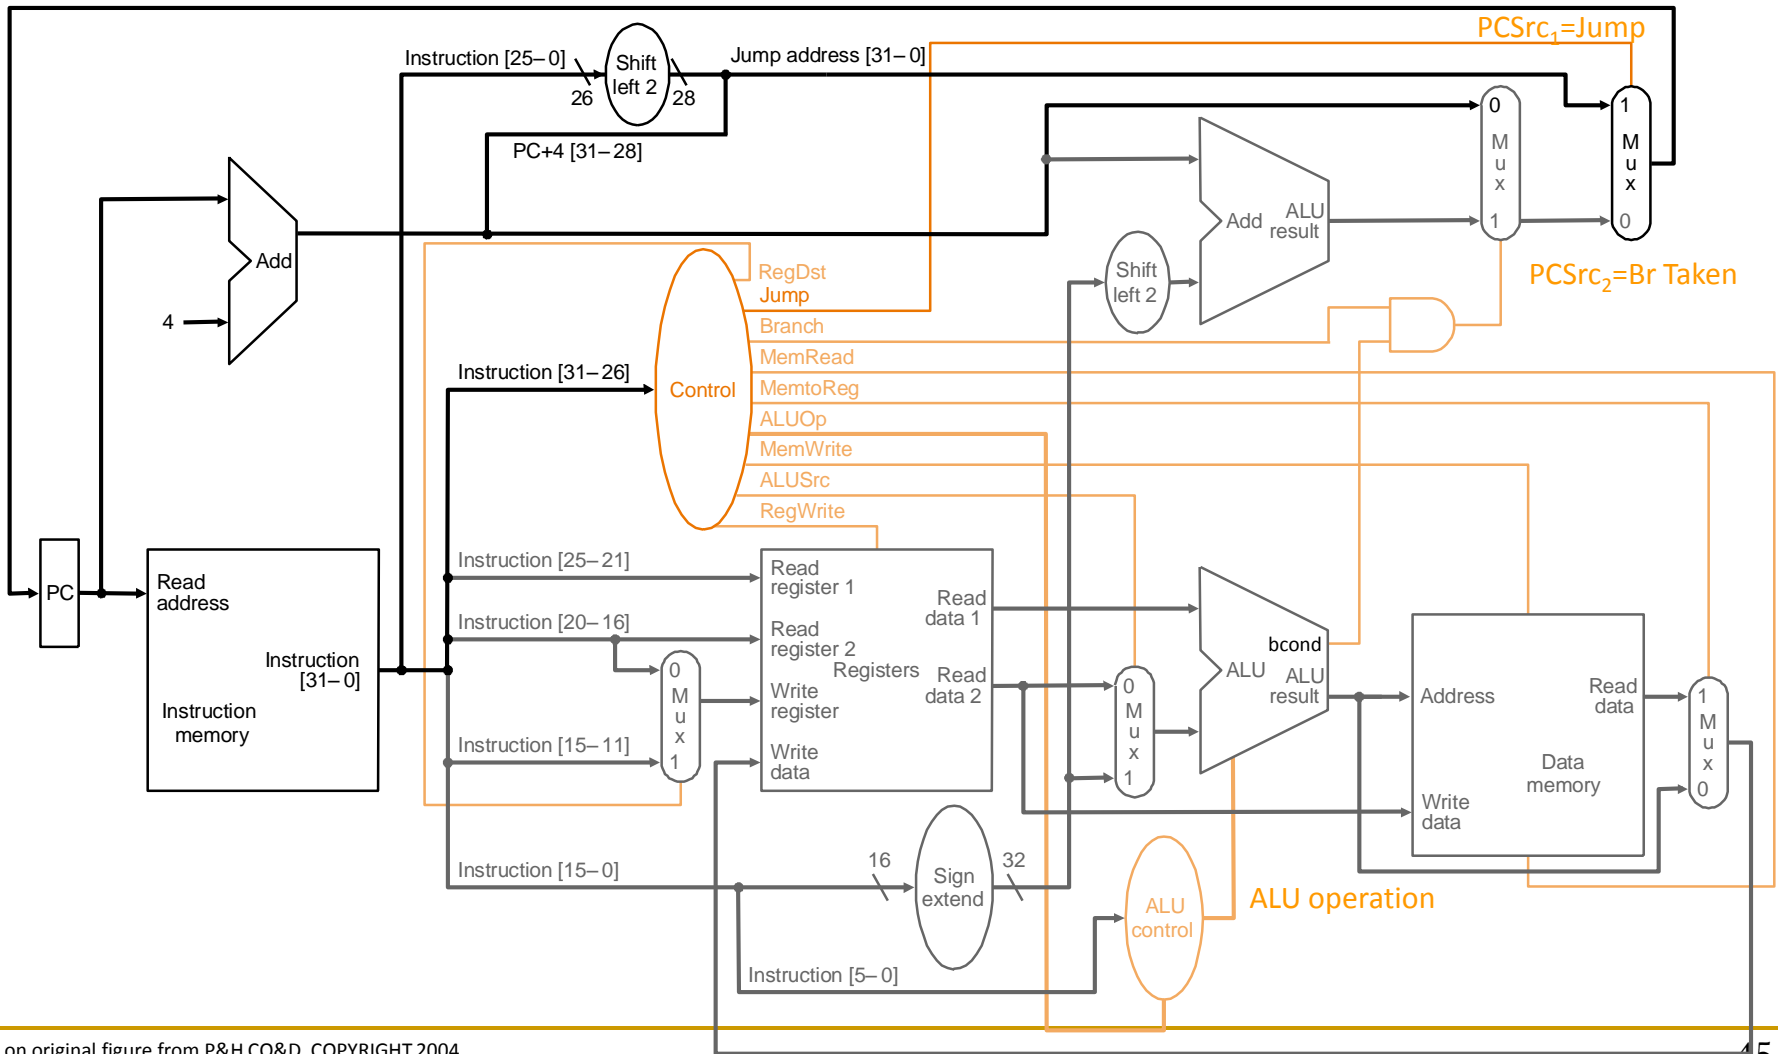
# Single-Cycle Datapath Analysis

---

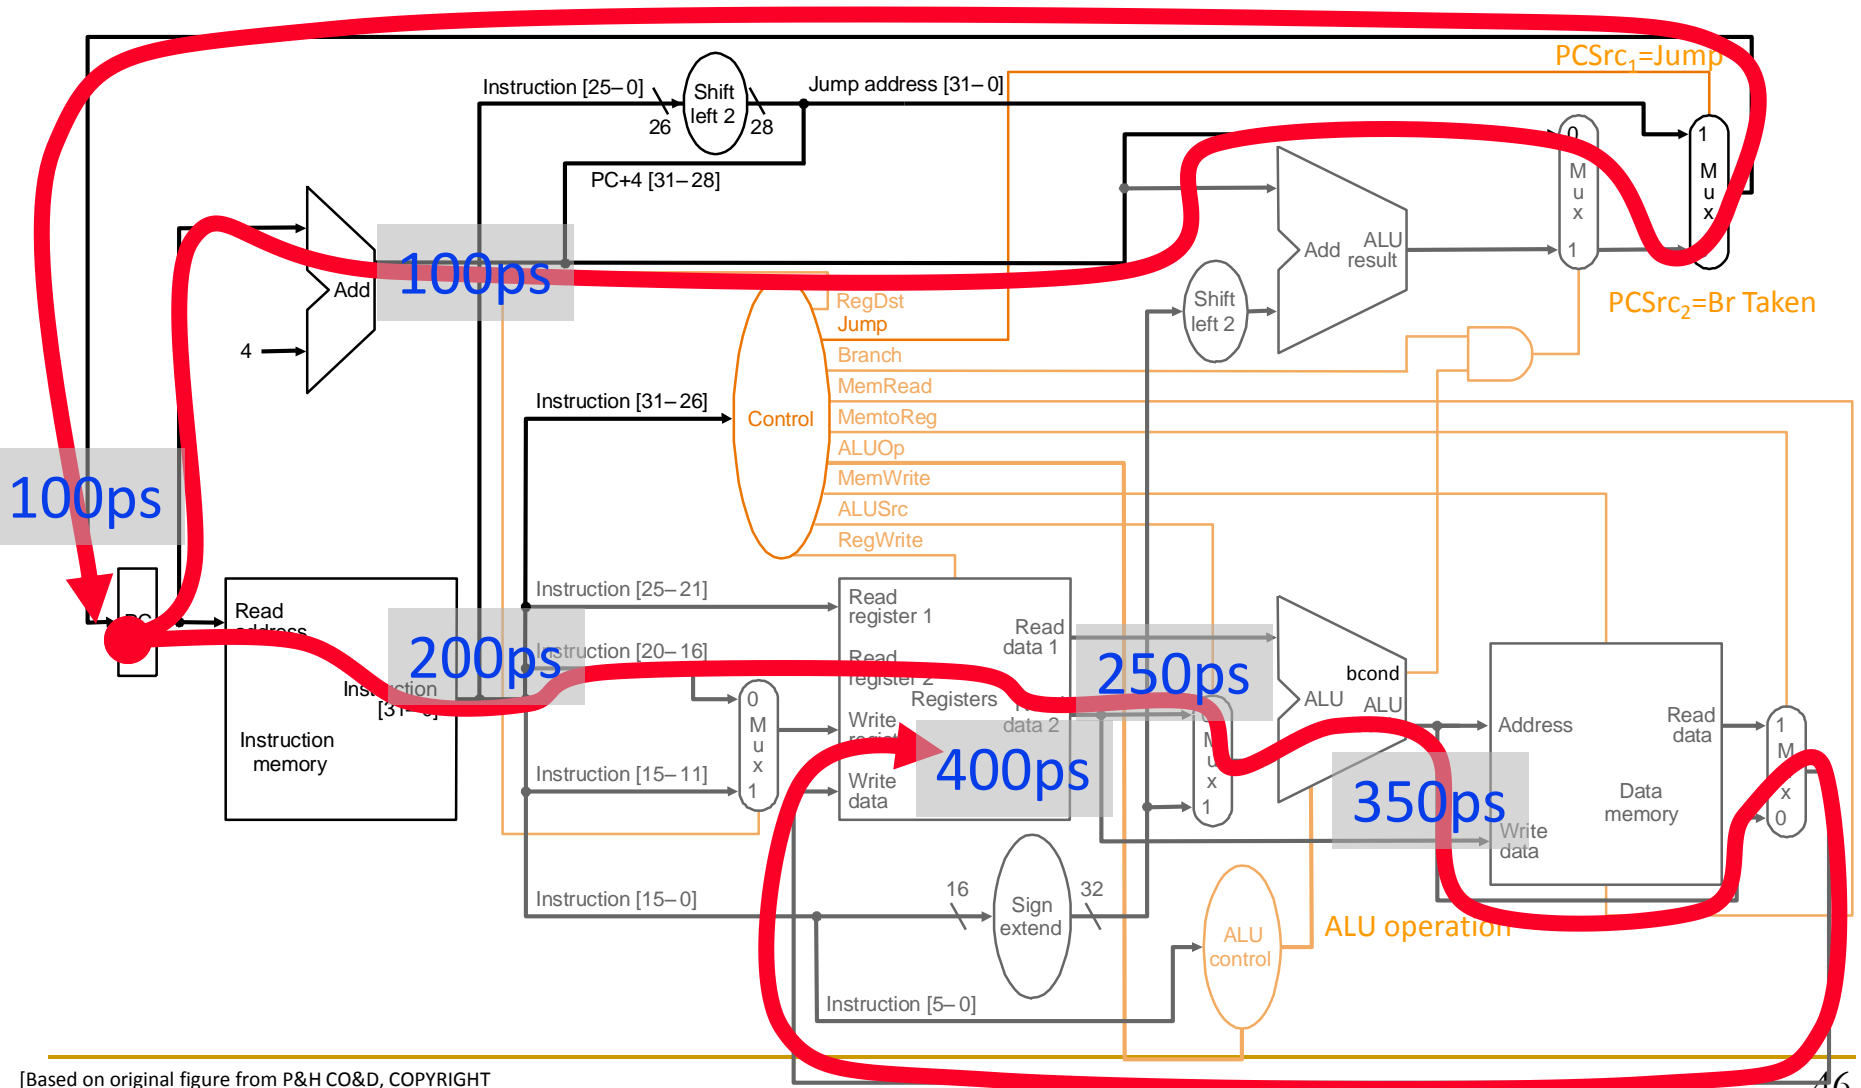
- Assume
  - ❑ memory units (read or write): 200 ps
  - ❑ ALU and adders: 100 ps
  - ❑ register file (read or write): 50 ps
  - ❑ other combinational logic: 0 ps


steps	IF	ID	EX	MEM	WB	Delay
resources	mem	RF	ALU	mem	RF	
R-type	200	50	100		50	400
I-type	200	50	100		50	400
LW	200	50	100	200	50	600
SW	200	50	100	200		550
Branch	200	50	100			350
Jump	200					200

# Let's Find the Critical Path

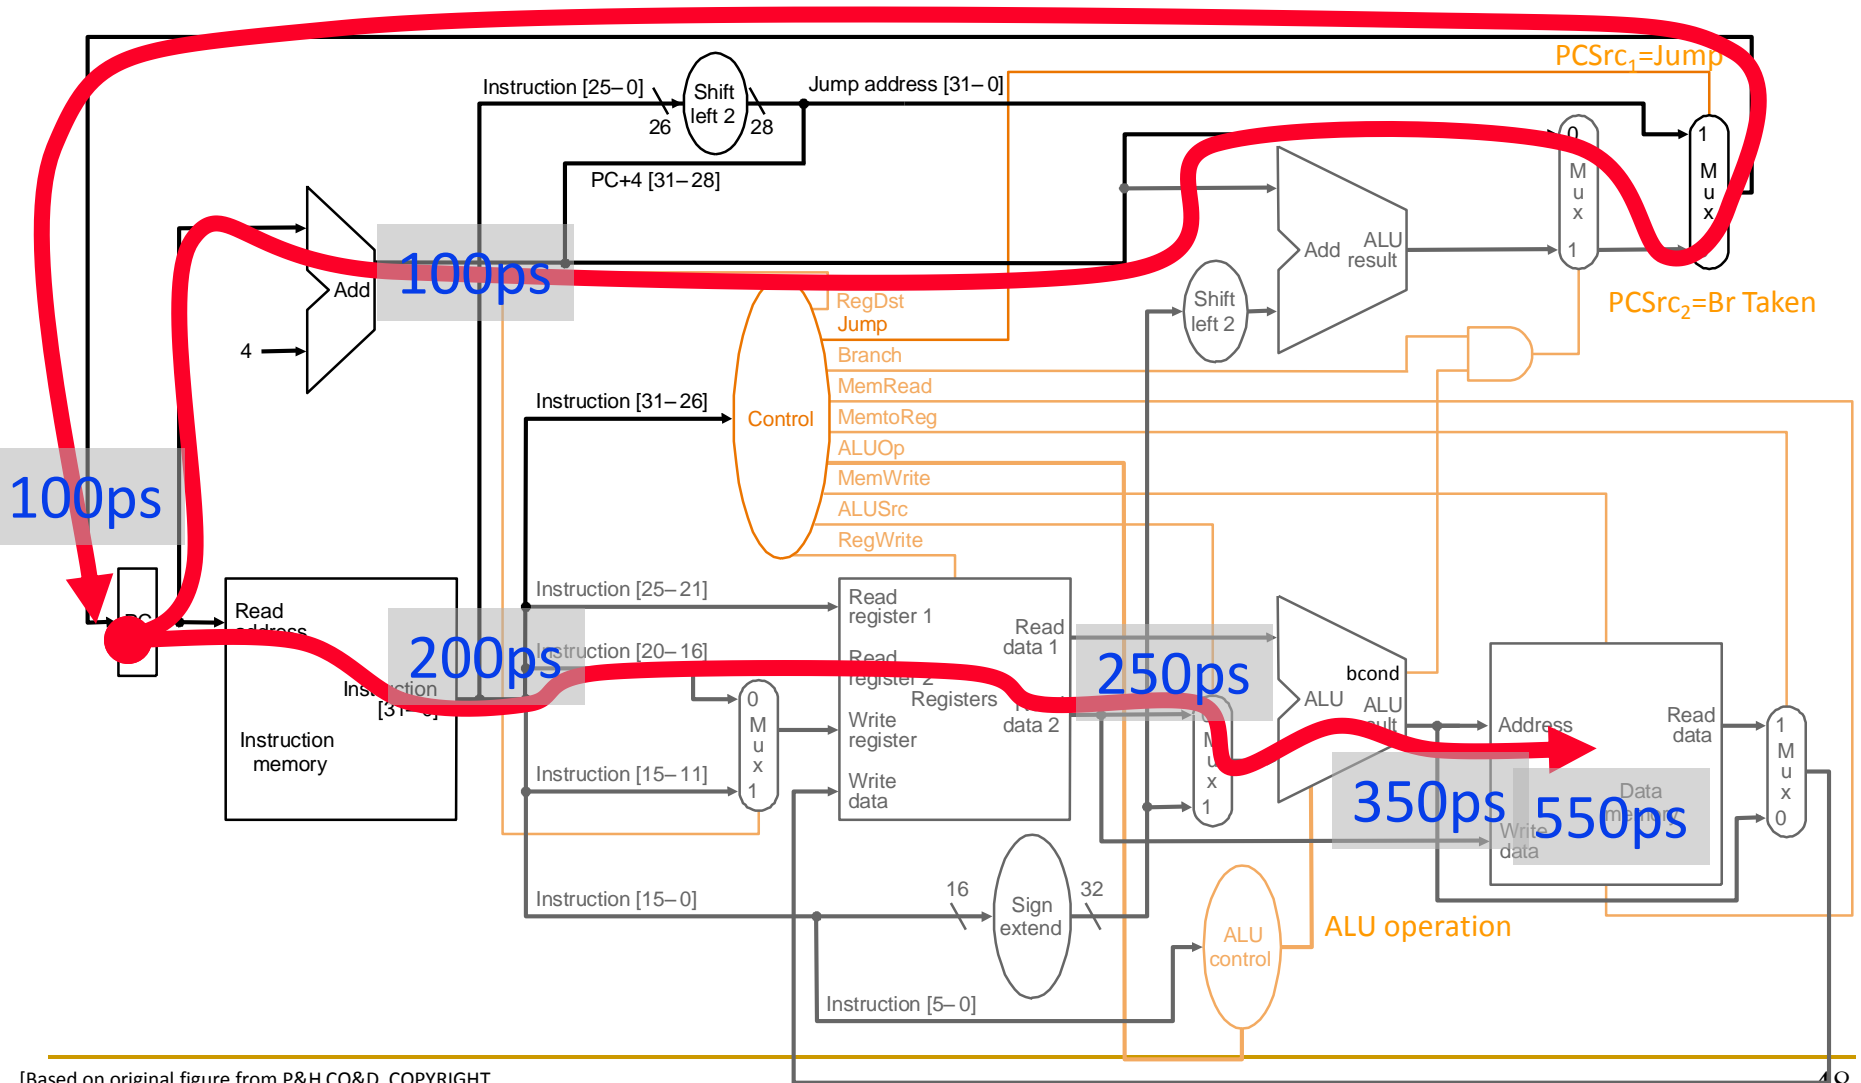


# R-Type and I-Type ALU



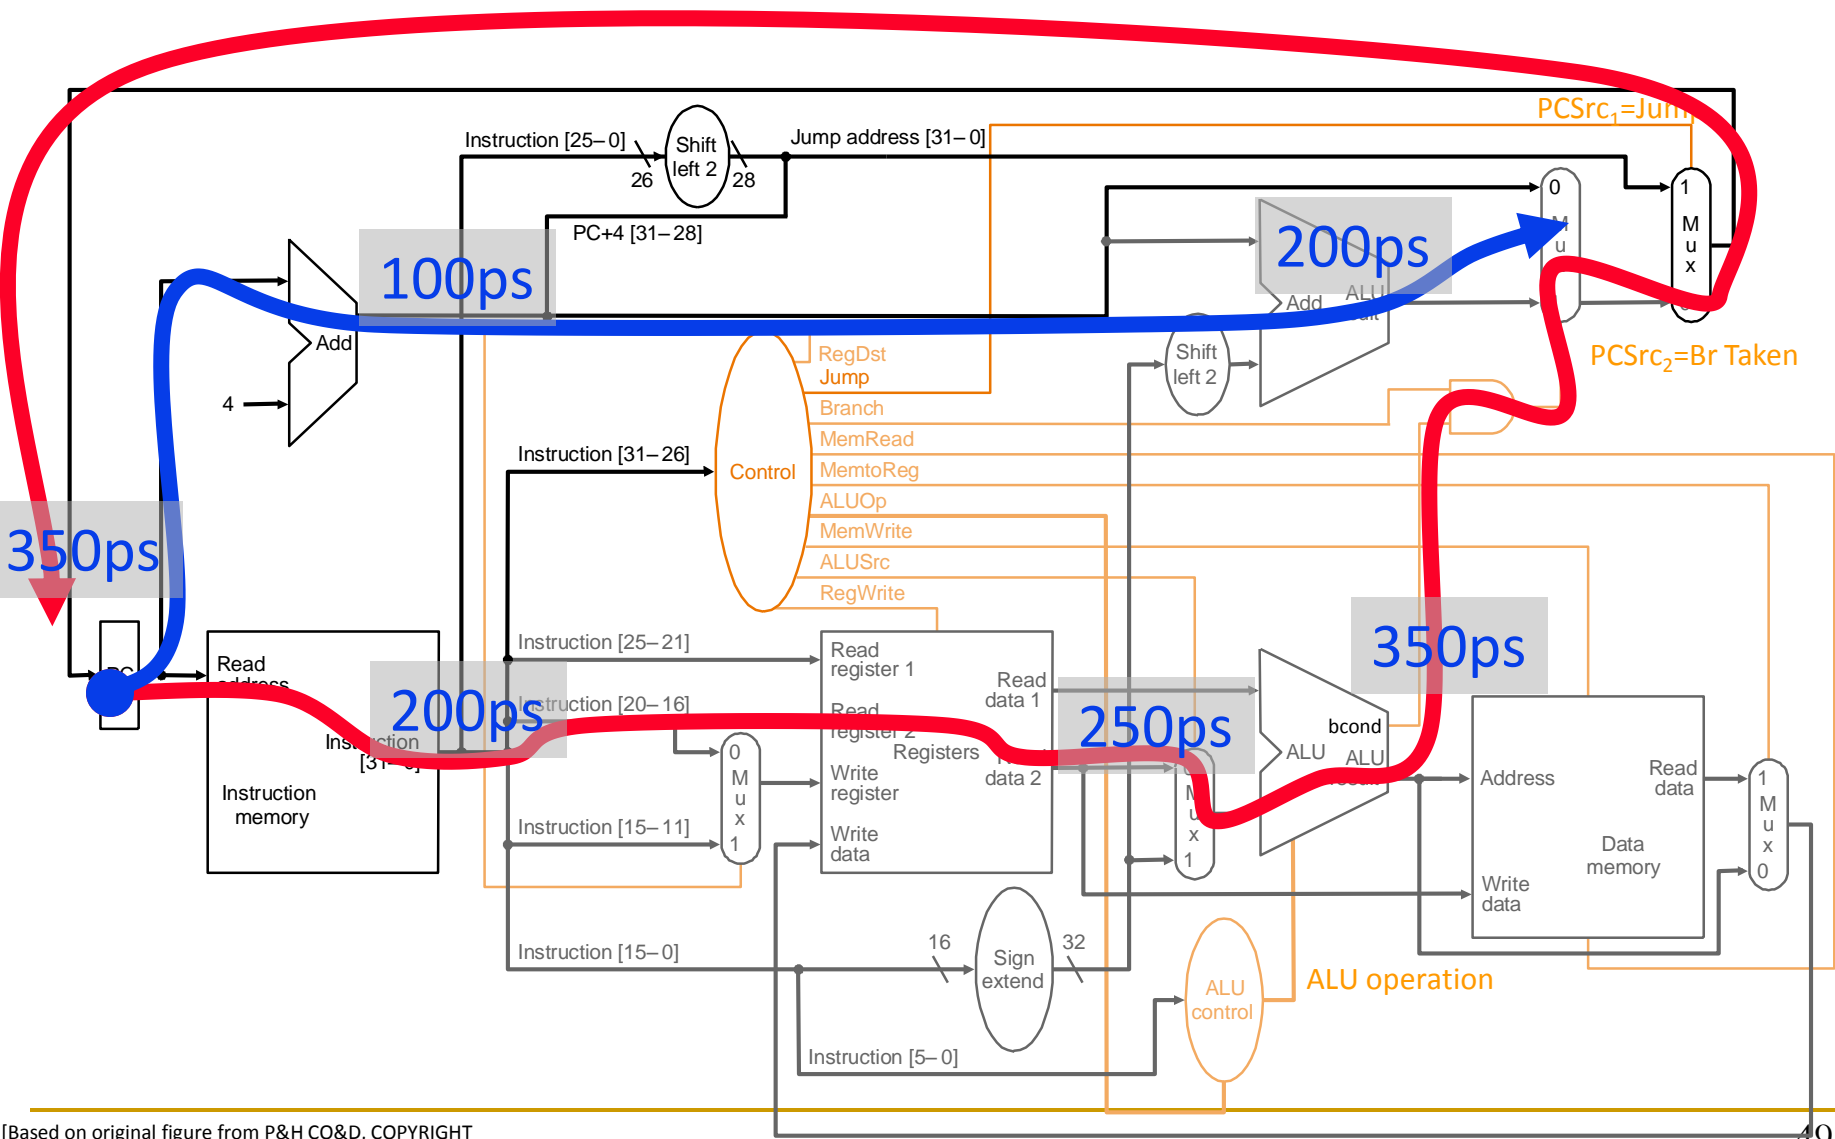


[Based on original figure from P&H CO&D, COPYRIGHT 2004 Elsevier. ALL RIGHTS RESERVED.]



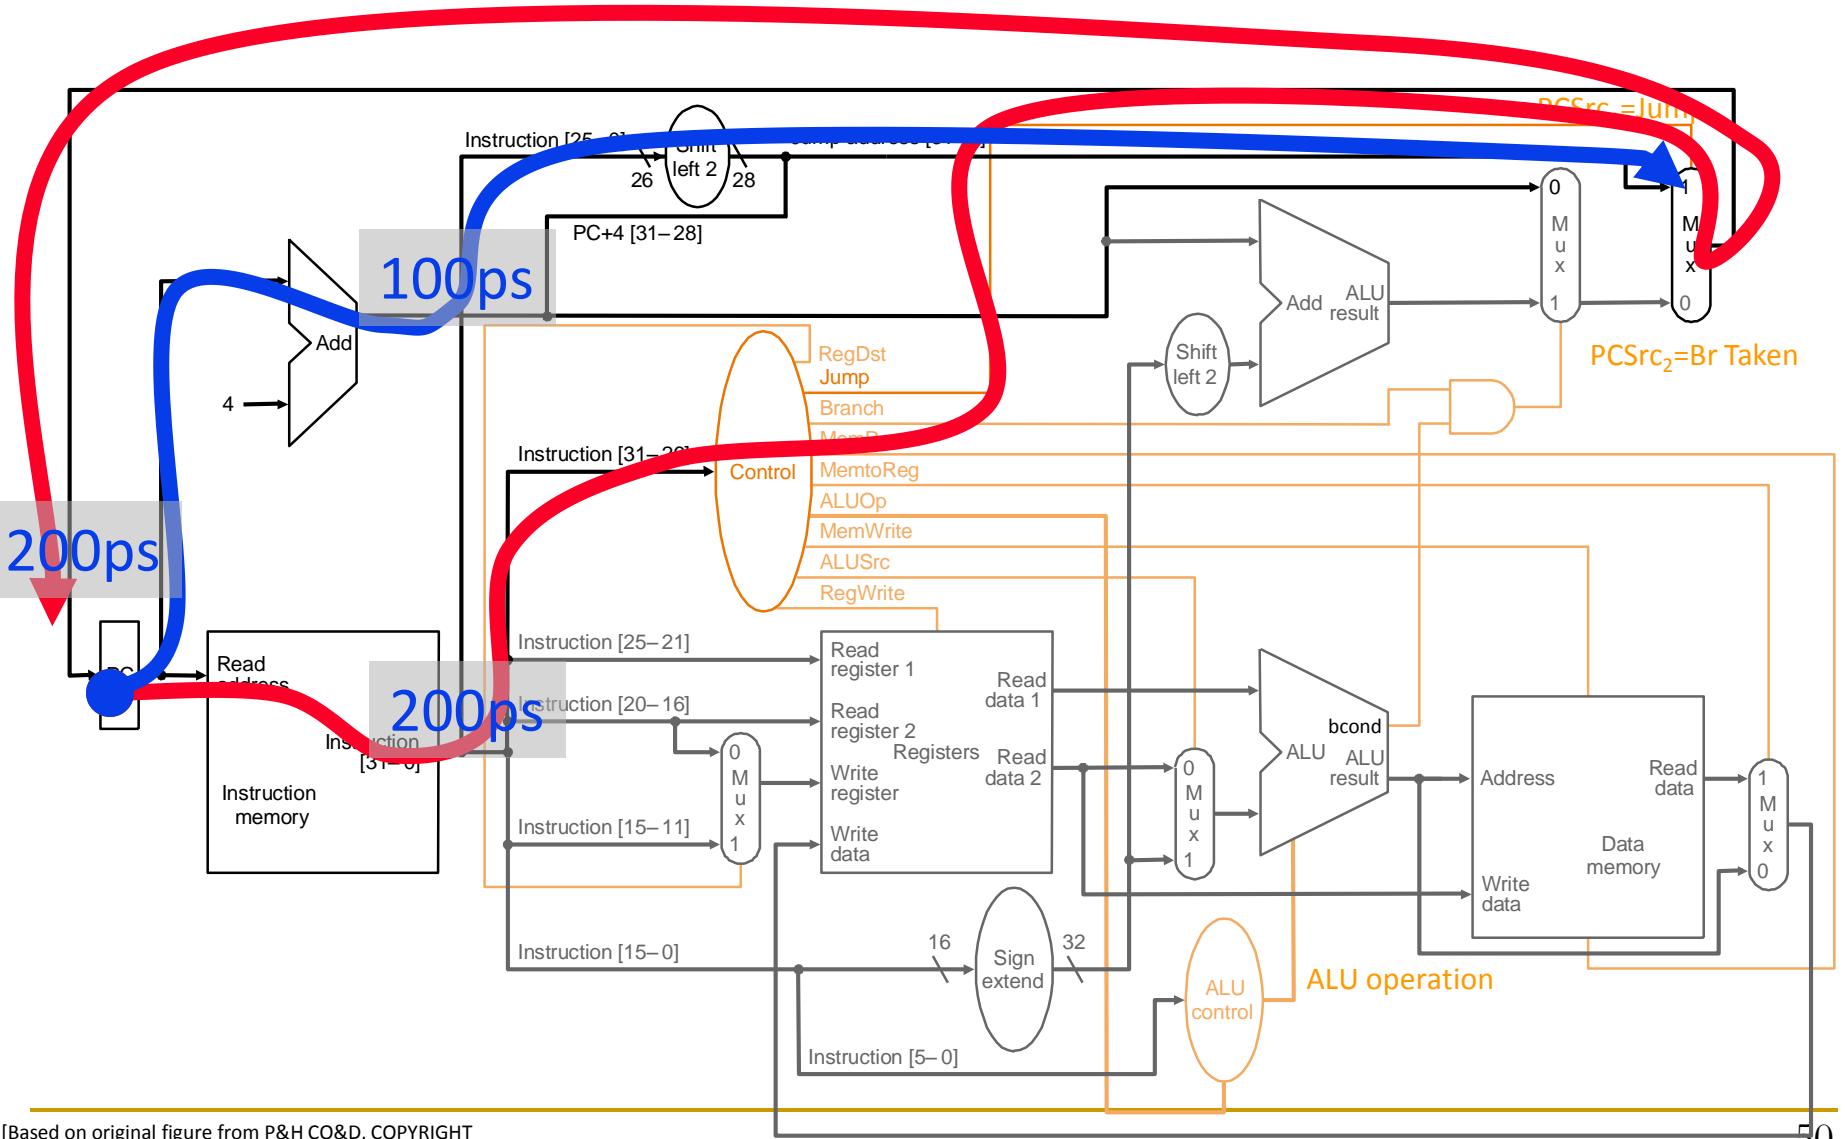


# Branch Taken



[Based on original figure from P&H CO&D, COPYRIGHT 2004 Elsevier. ALL RIGHTS RESERVED.]

# Jump



[Based on original figure from P&H CO&D, COPYRIGHT 2004 Elsevier. ALL RIGHTS RESERVED.]

# What About Control Logic?

---

- How does that affect the critical path?
- Food for thought for you:
  - Can control logic be on the critical path?
  - A note on CDC 5600: control store access too long...

# What is the Slowest Instruction to Process?

---

- Memory is not magic
- What if memory *sometimes* takes 100ms to access?
- Does it make sense to have a simple register to register add or jump to take {100ms+all else to do a memory operation}?
- And, what if you need to access memory more than once to process an instruction?
  - Which instructions need this?
  - Do you provide multiple ports to memory?

# Single Cycle uArch: Complexity

---

- Contrived
  - All instructions run as slow as the slowest instruction
- Inefficient
  - All instructions run as slow as the slowest instruction
  - Must provide worst-case combinational resources in parallel as required by any instruction
  - Need to replicate a resource if it is needed more than once by an instruction during different parts of the instruction processing cycle
- Not necessarily the simplest way to implement an ISA
  - Single-cycle implementation of REP MOVS, INDEX, POLY?
- Not easy to optimize/improve performance
  - Optimizing the common case does not work (e.g. common instructions)
  - Need to optimize the worst case all the time

# Microarchitecture Design Principles

---

- **Critical path design**
  - Find the maximum combinational logic delay and decrease it
- **Bread and butter (common case) design**
  - Spend time and resources on where it matters
    - i.e., improve what the machine is really designed to do
  - Common case vs. uncommon case
- **Balanced design**
  - Balance instruction/data flow through hardware components
  - Balance the hardware needed to accomplish the work
- *How does a single-cycle microarchitecture fare in light of these principles?*

# Multi-Cycle Microarchitectures

# Multi-Cycle Microarchitectures

---

- Goal: Let each instruction take (close to) only as much time it really needs
- Idea
  - Determine clock cycle time independently of instruction processing time
  - Each instruction takes as many clock cycles as it needs to take
    - Multiple state transitions per instruction
    - The states followed by each instruction is different



# Remember: The “Process instruction” Step


---

- ISA specifies abstractly what  $A'$  should be, given an instruction and  $A$ 
  - It defines an abstract finite state machine where
    - State = programmer-visible state
    - Next-state logic = instruction execution specification
  - From ISA point of view, there are no “intermediate states” between  $A$  and  $A'$  during instruction execution
    - One state transition per instruction
- Microarchitecture implements how  $A$  is transformed to  $A'$ 
  - There are many choices in implementation
  - We can have programmer-invisible state to optimize the speed of instruction execution: multiple state transitions per instruction
    - Choice 1:  $AS \rightarrow AS'$  (transform  $A$  to  $A'$  in a single clock cycle)
    - Choice 2:  $AS \rightarrow AS+MS1 \rightarrow AS+MS2 \rightarrow AS+MS3 \rightarrow AS'$  (take multiple clock cycles to transform  $AS$  to  $AS'$ )


# Multi-Cycle Microarchitecture

---


AS = Architectural (programmer visible) state  
at the beginning of an instruction




Step 1: Process part of instruction in one clock cycle



Step 2: Process part of instruction in the next clock cycle



...



AS' = Architectural (programmer visible) state  
at the end of a clock cycle

# Benefits of Multi-Cycle Design

---

## ■ Critical path design

- Can keep reducing the critical path independently of the worst-case processing time of any instruction

## ■ Bread and butter (common case) design

- Can optimize the number of states it takes to execute “important” instructions that make up much of the execution time

## ■ Balanced design

- No need to provide more capability or resources than really needed
  - An instruction that needs resource X multiple times does not require multiple X's to be implemented
  - Leads to more efficient hardware: Can reuse hardware components needed multiple times for an instruction

# Remember: Performance Analysis

---

- Execution time of an instruction
    - $\{\text{CPI}\} \times \{\text{clock cycle time}\}$
  - Execution time of a program
    - Sum over all instructions [ $\{\text{CPI}\} \times \{\text{clock cycle time}\}$ ]
    - $\{\#\text{ of instructions}\} \times \{\text{Average CPI}\} \times \{\text{clock cycle time}\}$
  - Single cycle microarchitecture performance
    - $\text{CPI} = 1$
    - Clock cycle time = long
  - Multi-cycle microarchitecture performance
    - $\text{CPI} = \text{different for each instruction}$ 
      - Average CPI  $\rightarrow$  hopefully small
    - Clock cycle time = short
- Now, we have two degrees of freedom to optimize independently**