




# Chapter 4

## The Processor

# Recall .... ISA ?



- Instruction Format or Encoding
  - how is it decoded?
- Location of operands and result
  - where other than memory?
  - how many explicit operands?
  - how are memory operands located?
  - which can or cannot be in memory?
- Data type and size
- Operations
  - what are supported?
- Successor instruction
  - what instruction is executed next?


# Introduction

- CPU performance factors
  - Instruction count
    - Determined by ISA and compiler
  - CPI and Cycle time
    - Determined by CPU hardware
- We will examine two MIPS implementations
  - A simplified version
  - A more realistic pipelined version
- Simple subset, shows most aspects
  - Memory reference: l w, s w
  - Arithmetic/logical: add, sub, and, or, sl t
  - Control transfer: beq, j

# How to Design a Processor?

1. Analyze instruction set (datapath requirements)
  - The meaning of each instruction is given by the *register transfers*
  - Datapath must include storage element
  - Datapath must support each register transfer
2. Select set of datapath components and establish clocking methodology
3. Assemble datapath meeting the requirements
4. Analyze implementation of each instruction to determine setting of control points effecting register transfer
5. Assemble the control logic

# Datapath vs Control




- Datapath: Storage, FU, interconnect sufficient to perform the desired functions
  - Inputs are Control Points
  - Outputs are signals
- Controller: State machine to orchestrate operation on the data path
  - Based on desired function and signals


# Step 1: Analyze Instruction Set

- All MIPS instructions are 32 bits long with 3 formats:


- R-type:



- I-type:



- J-type:



6 bits

26 bits


- The different fields are:

- op: operation of the instruction
- rs, rt, rd: source and destination register
- shamt: shift amount
- funct: selects variant of the “op” field
- address / immediate
- target address: target address of jump

# Our Example: A MIPS Subset


- R-Type:

- add rd, rs, rt
- sub rd, rs, rt
- and rd, rs, rt
- or rd, rs, rt
- slt rd, rs, rt



- Load/Store:

- lw rt,rs,imm16
- sw rt,rs,imm16



- Imm operand:


- addi rt,rs,imm16

- Branch:

- beq rs,rt,imm16

- Jump:

- j target



# Instruction Execution

- PC → instruction memory, fetch instruction
- Register numbers → register file, read registers
- Depending on instruction class
  - Use ALU to calculate
    - Arithmetic result
    - Memory address for load/store
    - Branch target address
  - Access data memory for load/store
  - PC ← target address or PC + 4



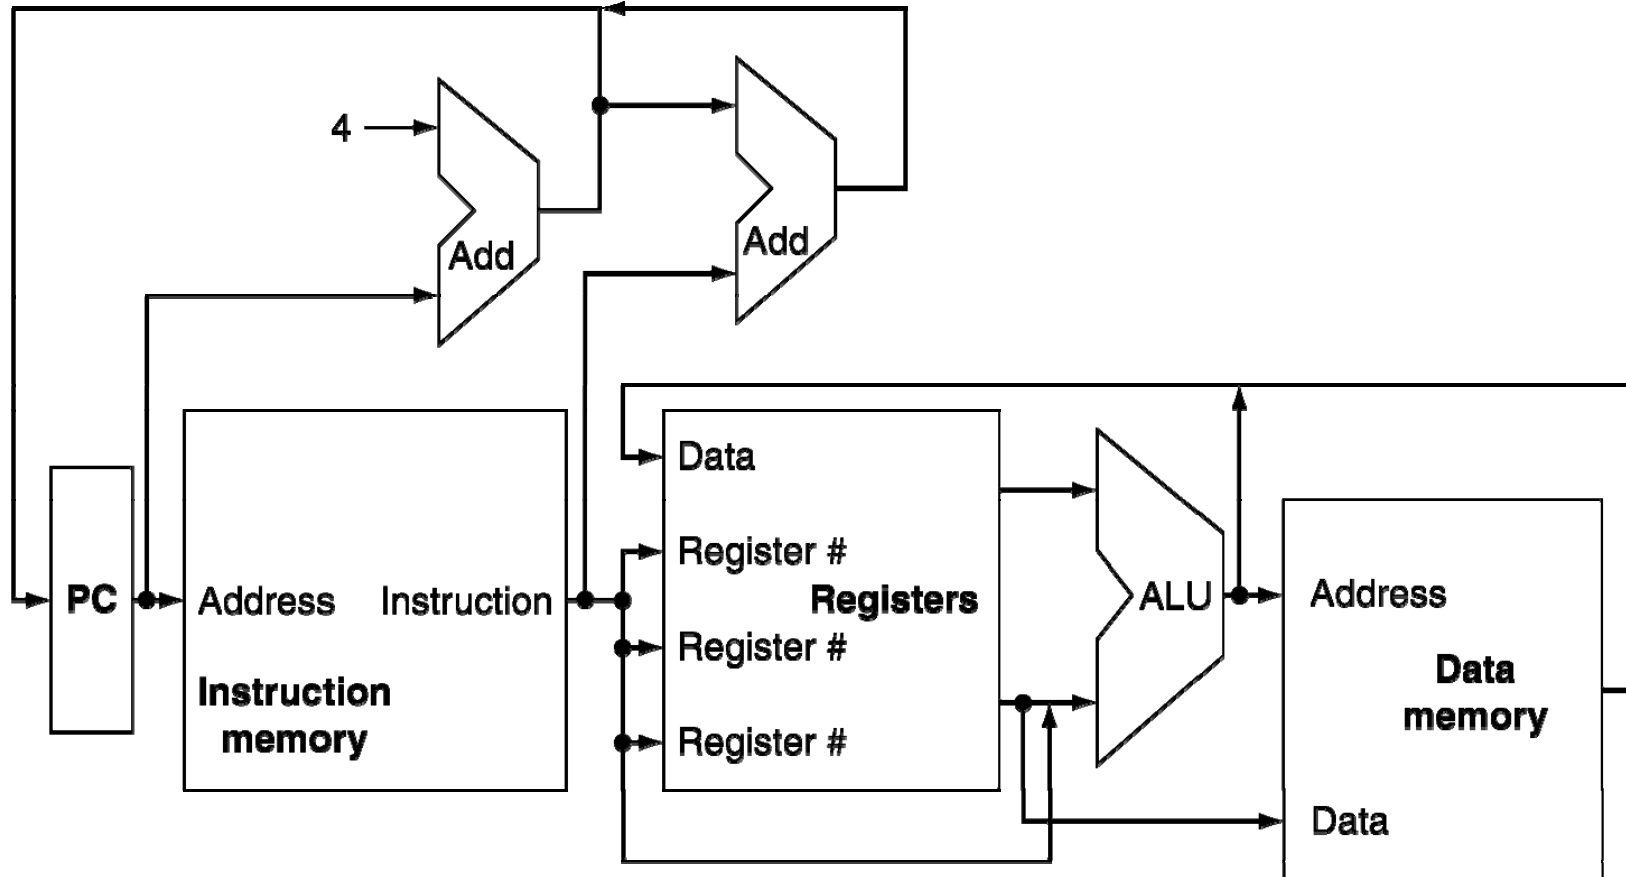
# Logical Register Transfers

- RTL gives the meaning of the instructions
- All start by fetching the instruction, read registers, then use ALU => simplicity and regularity help

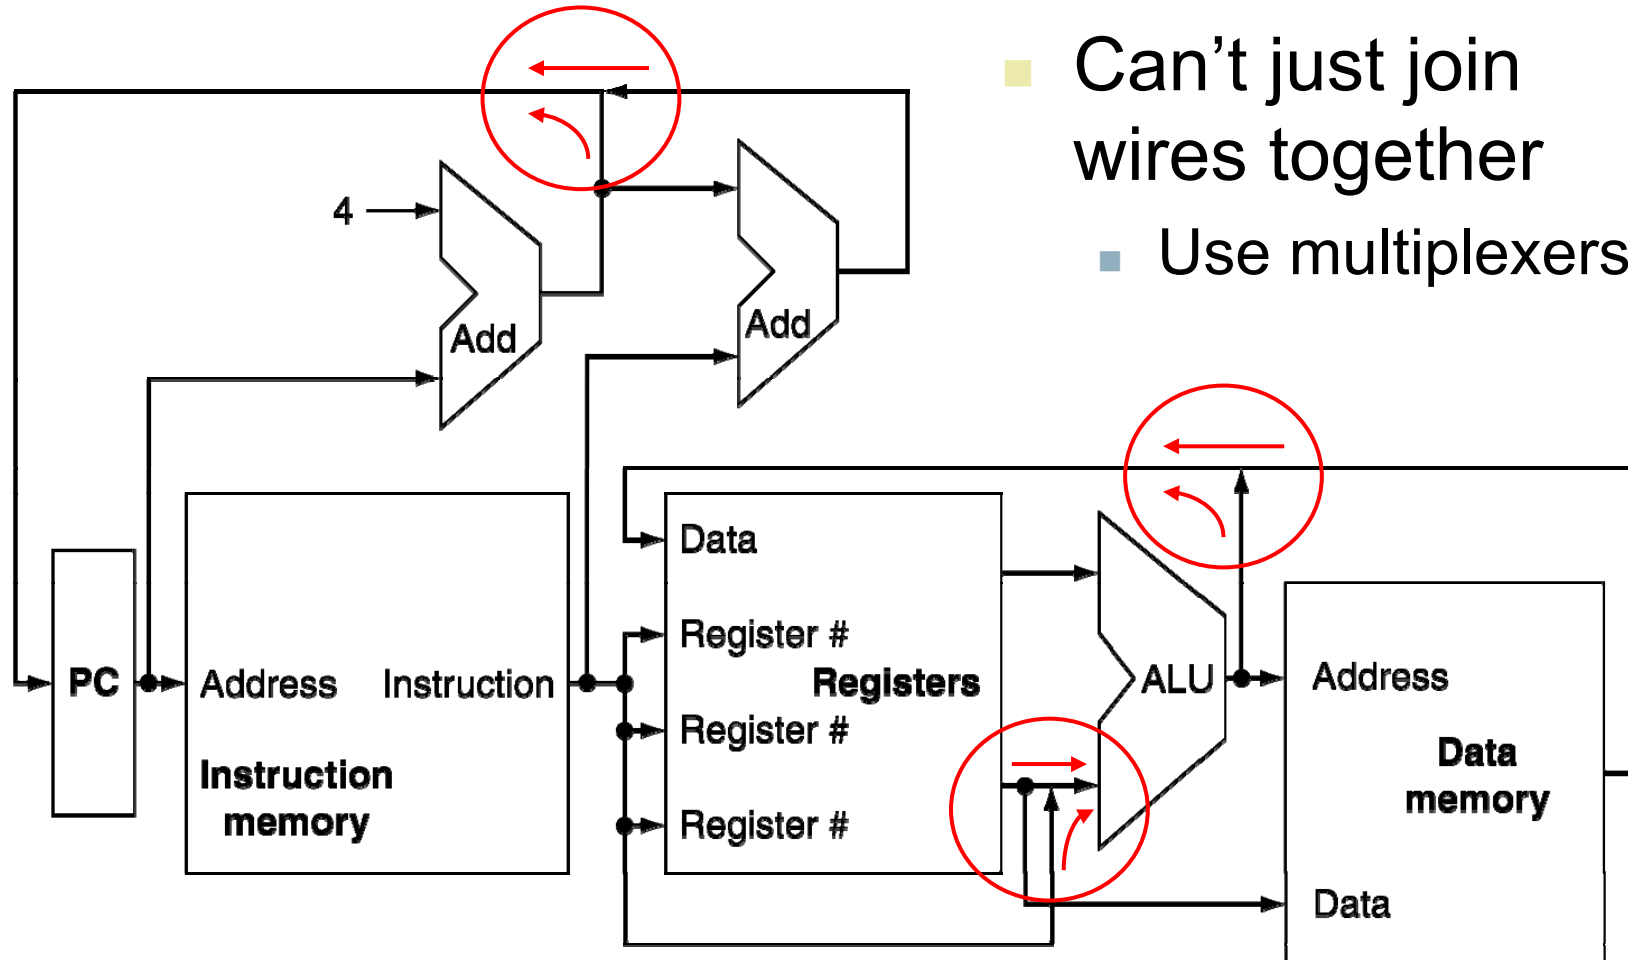
MEM[ PC ] = op | rs | rt | rd | shamt | funct  
 or = op | rs | rt | Imm16  
 or = op | Imm26 (added at the end)

Inst	Register transfers
ADD	$R[rd] \leftarrow R[rs] + R[rt]; \quad PC \leftarrow PC + 4$
SUB	$R[rd] \leftarrow R[rs] - R[rt]; \quad PC \leftarrow PC + 4$
LOAD	$R[rt] \leftarrow \text{MEM}[ R[rs] + \text{sign\_ext}(\text{Imm16}) ]; \quad PC \leftarrow PC + 4$
STORE	$\text{MEM}[ R[rs] + \text{sign\_ext}(\text{Imm16}) ] \leftarrow R[rt]; \quad PC \leftarrow PC + 4$
ADDI	$R[rt] \leftarrow R[rs] + \text{sign\_ext}(\text{Imm16}); \quad PC \leftarrow PC + 4$
BEQ	if $(R[rs] == R[rt])$ then $PC \leftarrow PC + 4 + \text{sign\_ext}(\text{Imm16})$    00 else $PC \leftarrow PC + 4$

# MIPS Datapath (simplified)

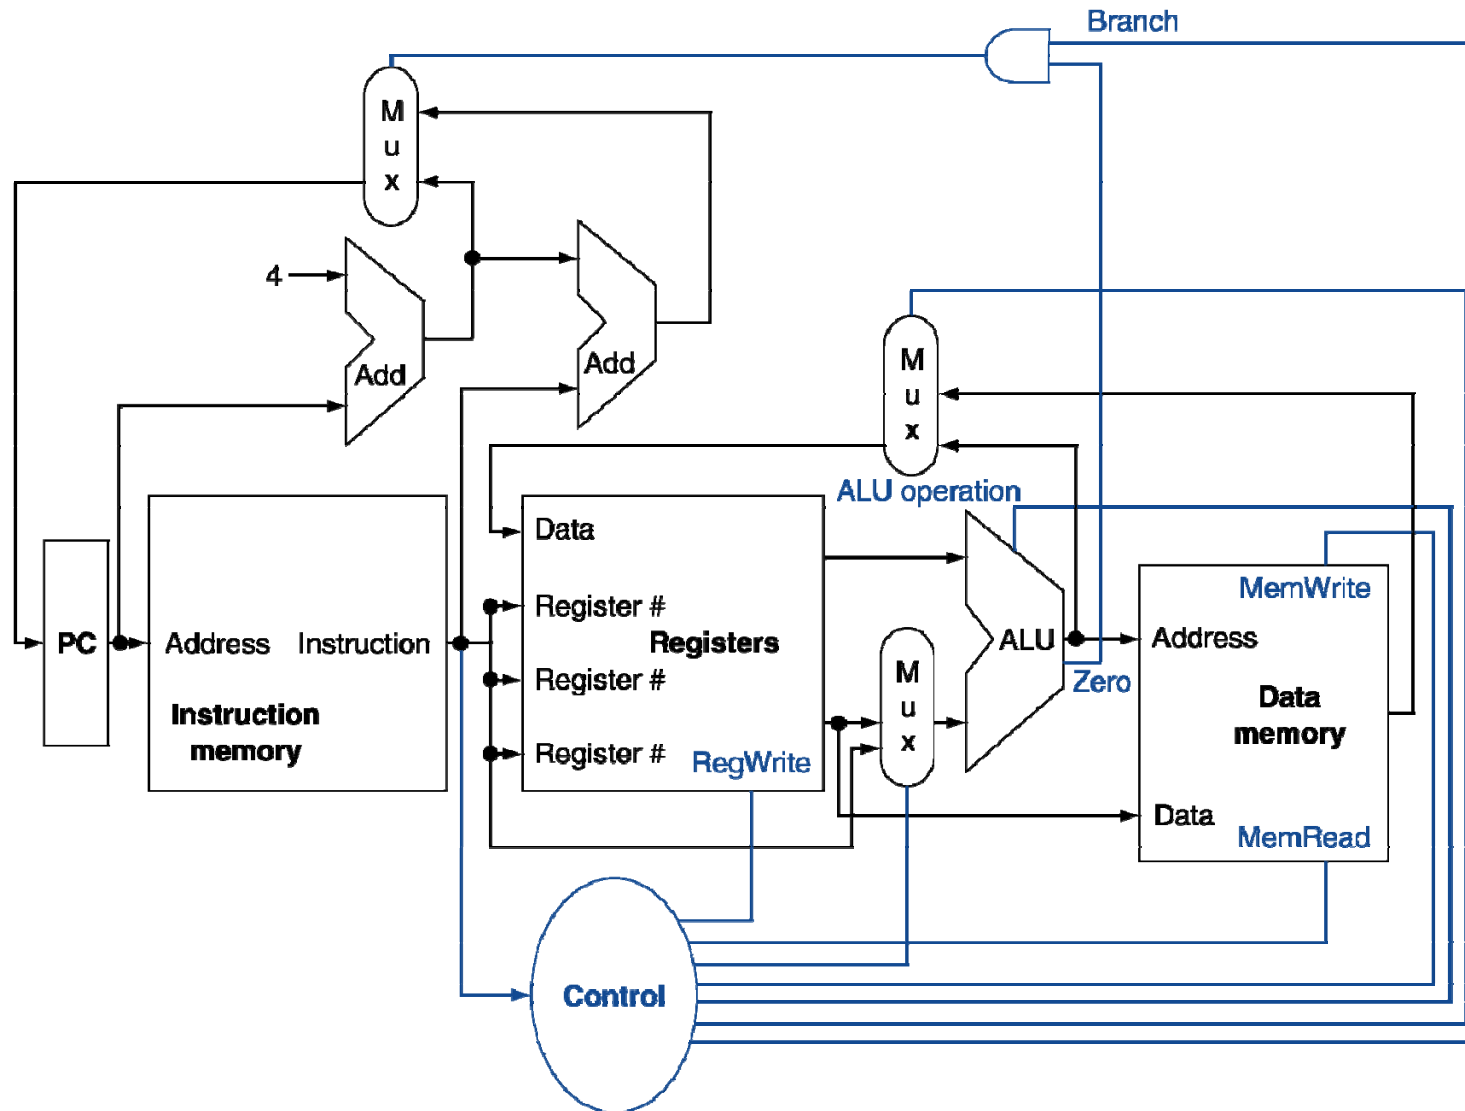


# Multiplexers




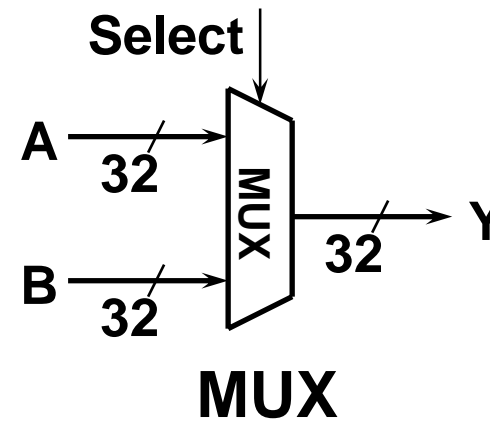

- Can't just join wires together
  - Use multiplexers

# Control



# Step 2a: Datapath Components

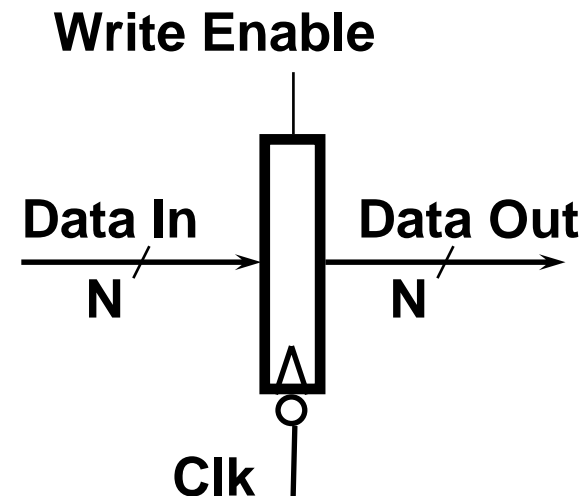
- Basic building blocks of combinational logic elements :



# Step 2b: Datapath Components

## Storage elements:

- Register:
  - Similar to the D Flip Flop, except
    - N-bit input and output
    - Write Enable input
      - negated (0): Data Out will not change
      - asserted (1): Data Out will become Data In




# Logic Design Basics

- Information encoded in binary
  - Low voltage = 0, High voltage = 1
  - One wire per bit
  - Multi-bit data encoded on multi-wire buses
- Combinational element
  - Operate on data
  - Output is a function of input
- State (sequential) elements
  - Store information

# Combinational Elements


- AND-gate

- $Y = A \& B$




- Adder

- $Y = A + B$




- Multiplexer

- $Y = S ? I1 : I0$



- Arithmetic/Logic Unit



- $Y = F(A, B)$







# Sequential Elements

- Register: stores data in a circuit
  - Uses a clock signal to determine when to update the stored value
  - Edge-triggered: update when Clk changes from 0 to 1





# Sequential Elements

- Register with write control
  - Only updates on clock edge when write control input is 1
  - Used when stored value is required later



# Clocking Methodology

- Combinational logic transforms data during clock cycles
  - Between clock edges
  - Input from state elements, output to state element
  - Longest delay determines clock period

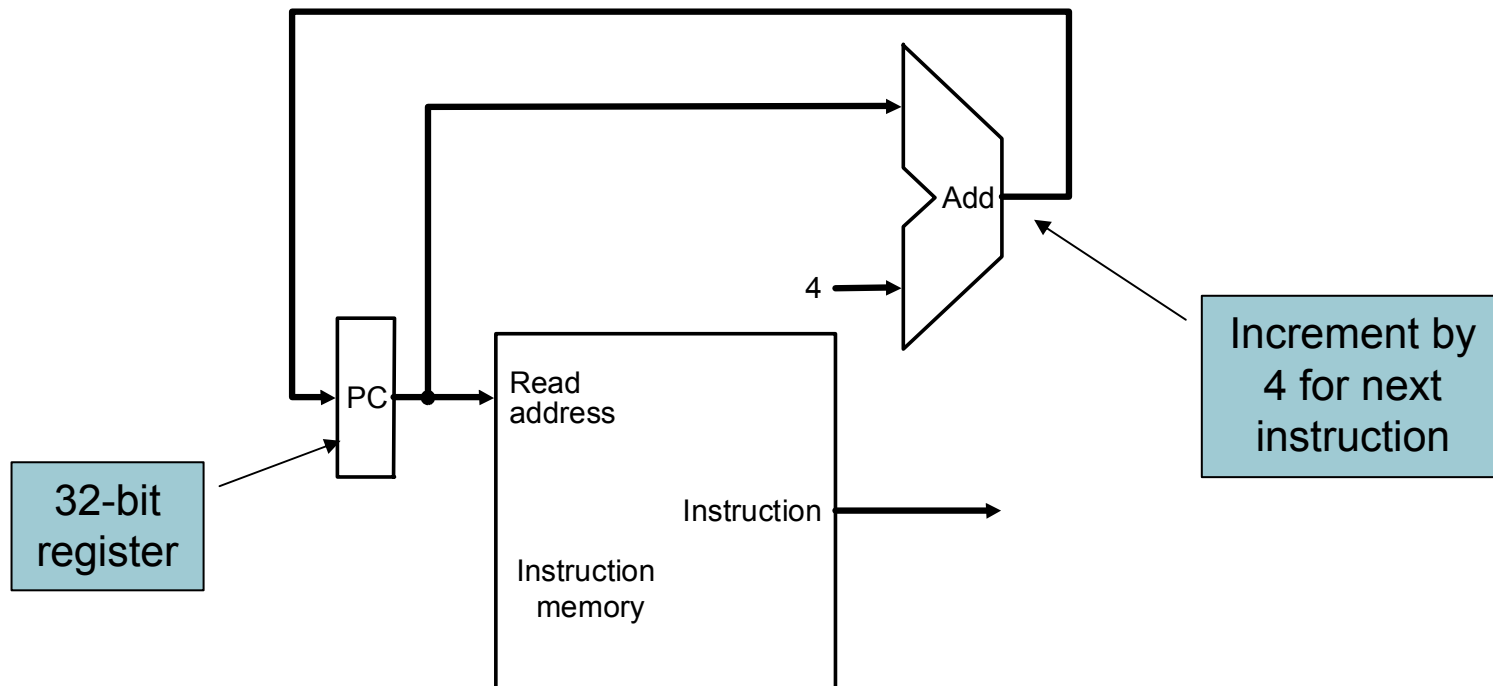


# Building a Datapath

- Datapath
  - Elements that process data and addresses in the CPU
    - Registers, ALUs, mux's, memories, ...
- We will build a MIPS datapath incrementally
  - **Refining the overview design**
  - **Step 3a: Datapath Assembly**


# Step 3a: Datapath Assembly

- **Instruction fetch unit:** common operations
  - Fetch the instruction:  $\text{mem}[\text{PC}]$
  - Update the program counter:
    - Sequential code:  $\text{PC} \leftarrow \text{PC} + 4$
    - Branch and Jump:  $\text{PC} \leftarrow \text{"Something else"}$



# Step 3b: Add and Subtract

- $R[rd] \leftarrow R[rs] \text{ op } R[rt]$       Ex: `add rd, rs, rt`
  - Ra, Rb, Rw come from inst.'s rs, rt, and rd fields
  - ALU and RegWrite: control logic after decode



Two read ports and one write port

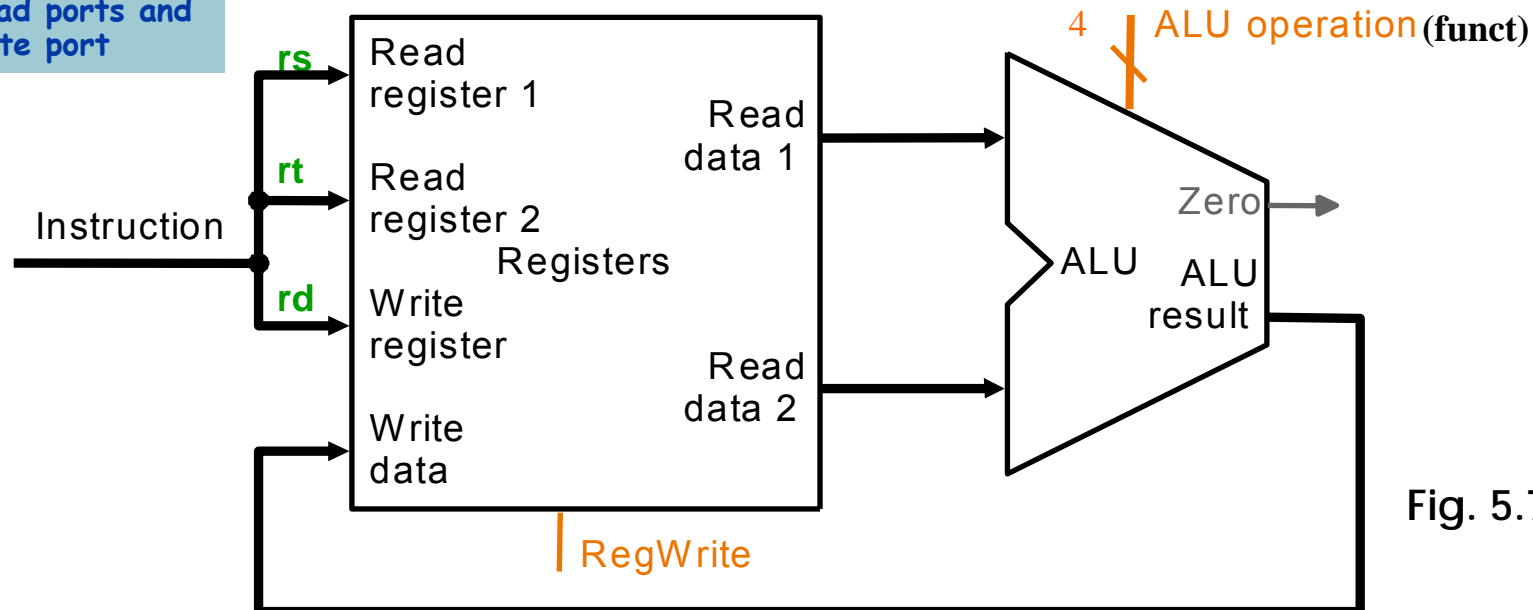
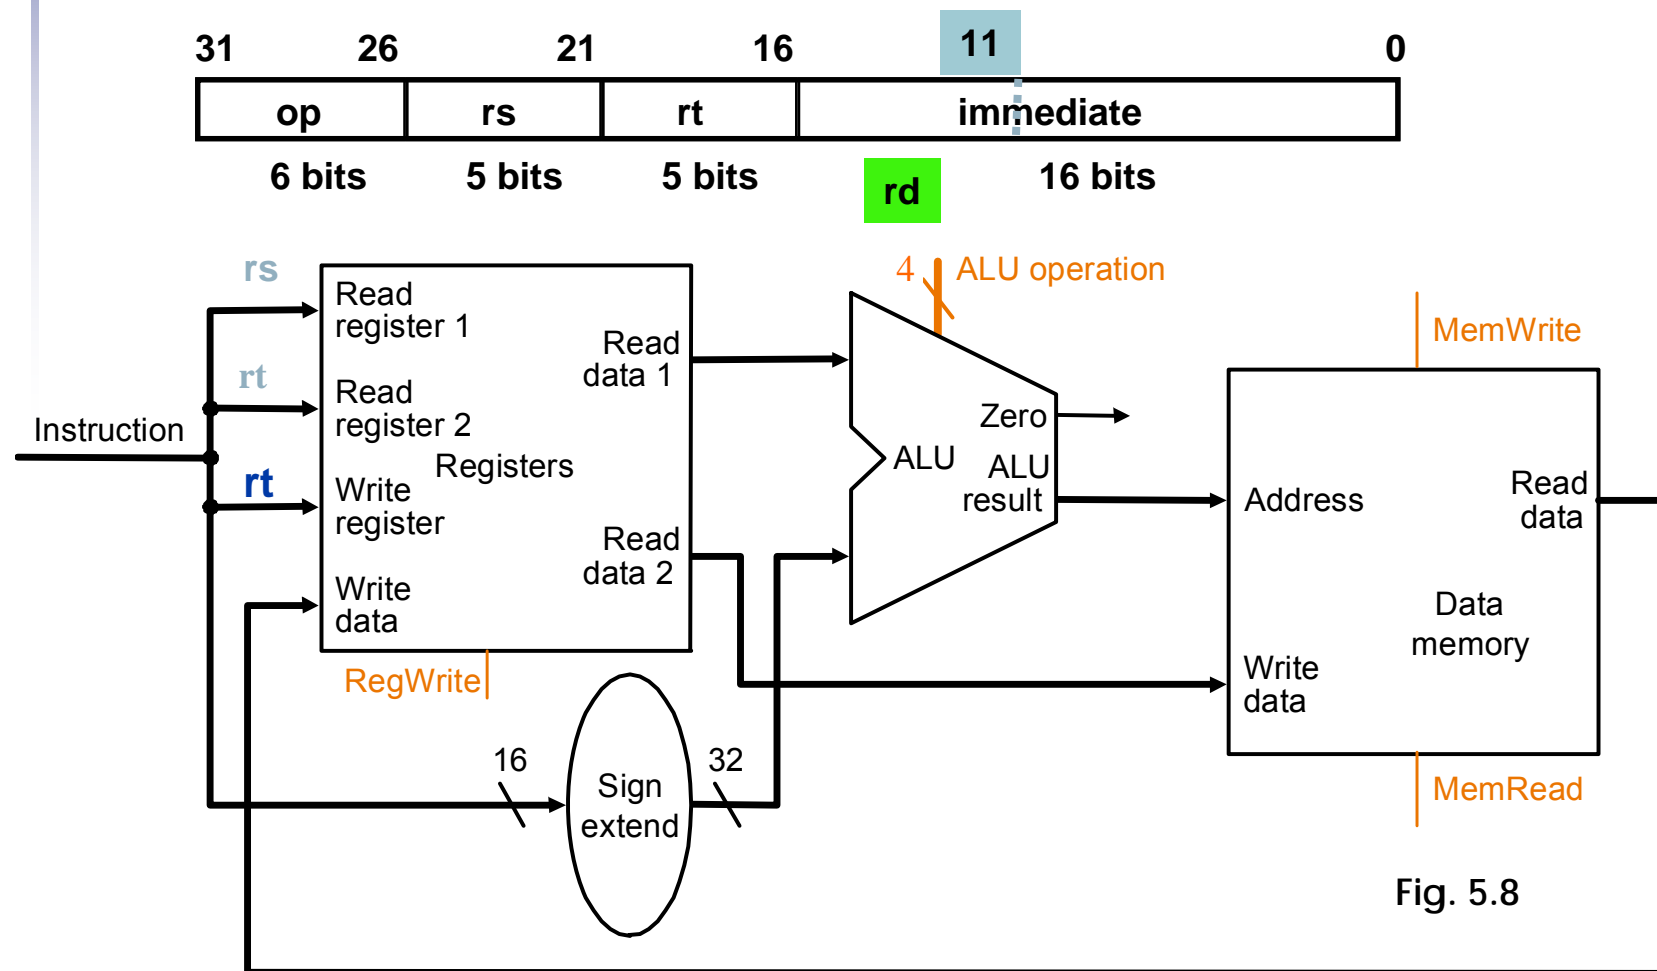


Fig. 5.7

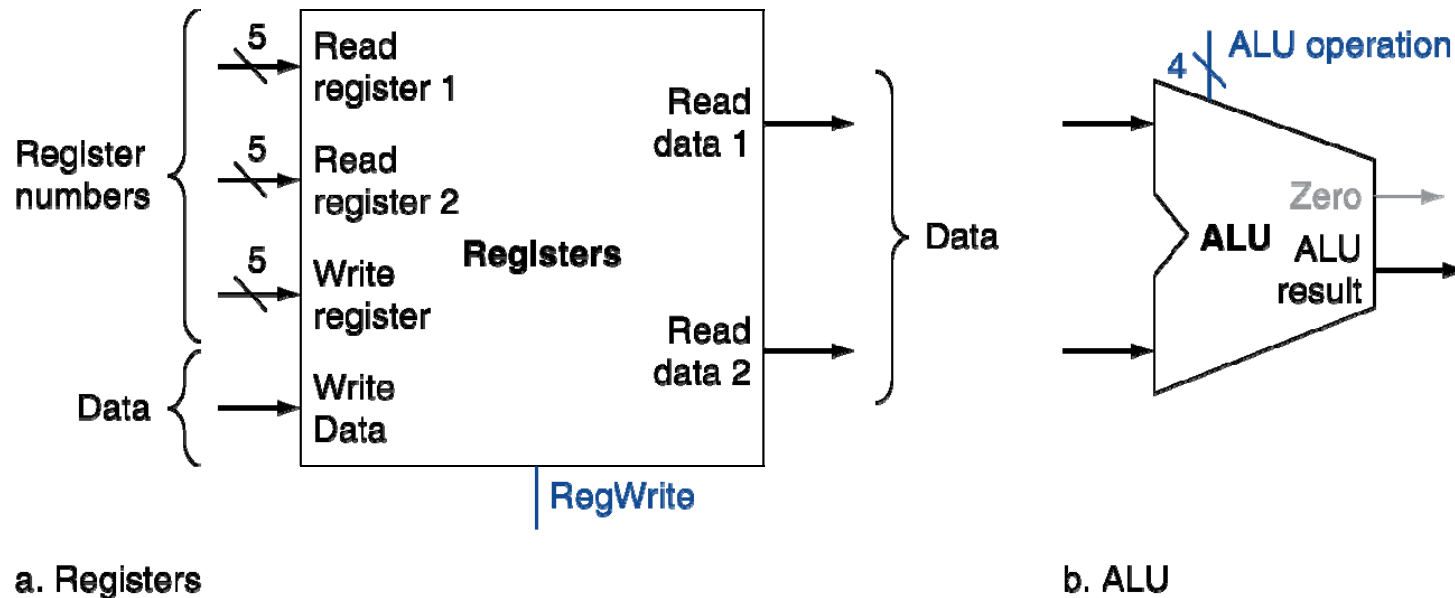
# Step 3c: Store/Load Operations

- $R[rt] \leftarrow \text{Mem}[R[rs] + \text{SignExt}[\text{imm16}]]$  Ex: `lw rt,rs,imm16`



# R-Format Instructions

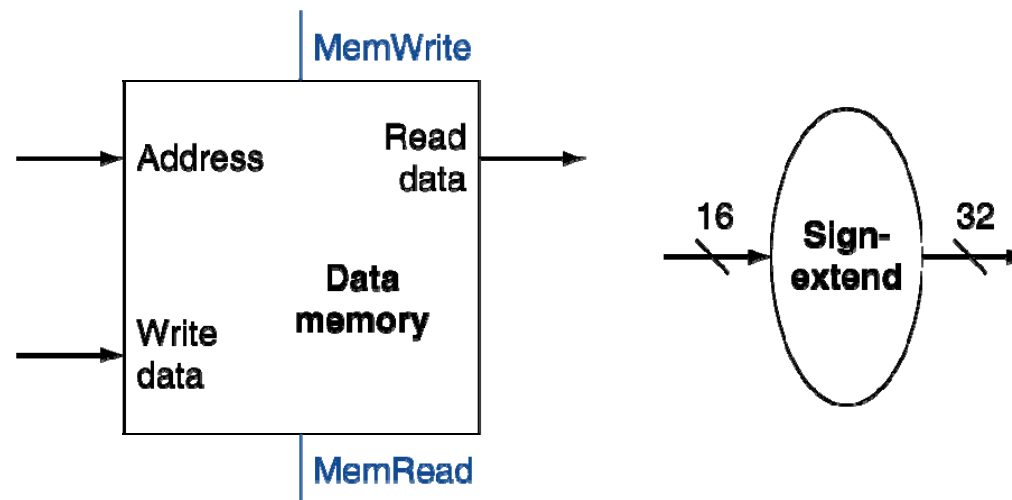
- Read two register operands
- Perform arithmetic/logical operation
- Write register result





# Load/Store Instructions

- Read register operands
- Calculate address using 16-bit offset
  - Use ALU, but sign-extend offset
- Load: Read memory and update register
- Store: Write register value to memory




a. Data memory unit

b. Sign extension unit

# Step 3d: Branch Operations

- `beq rs, rt, imm16`



`mem[PC]`

Fetch inst. from memory

`Equal <- R[rs] == R[rt]`

Calculate branch condition

`if (COND == 0)`

Calculate next inst. address

`PC <- PC + 4 + ( SignExt(imm16) x 4 )`

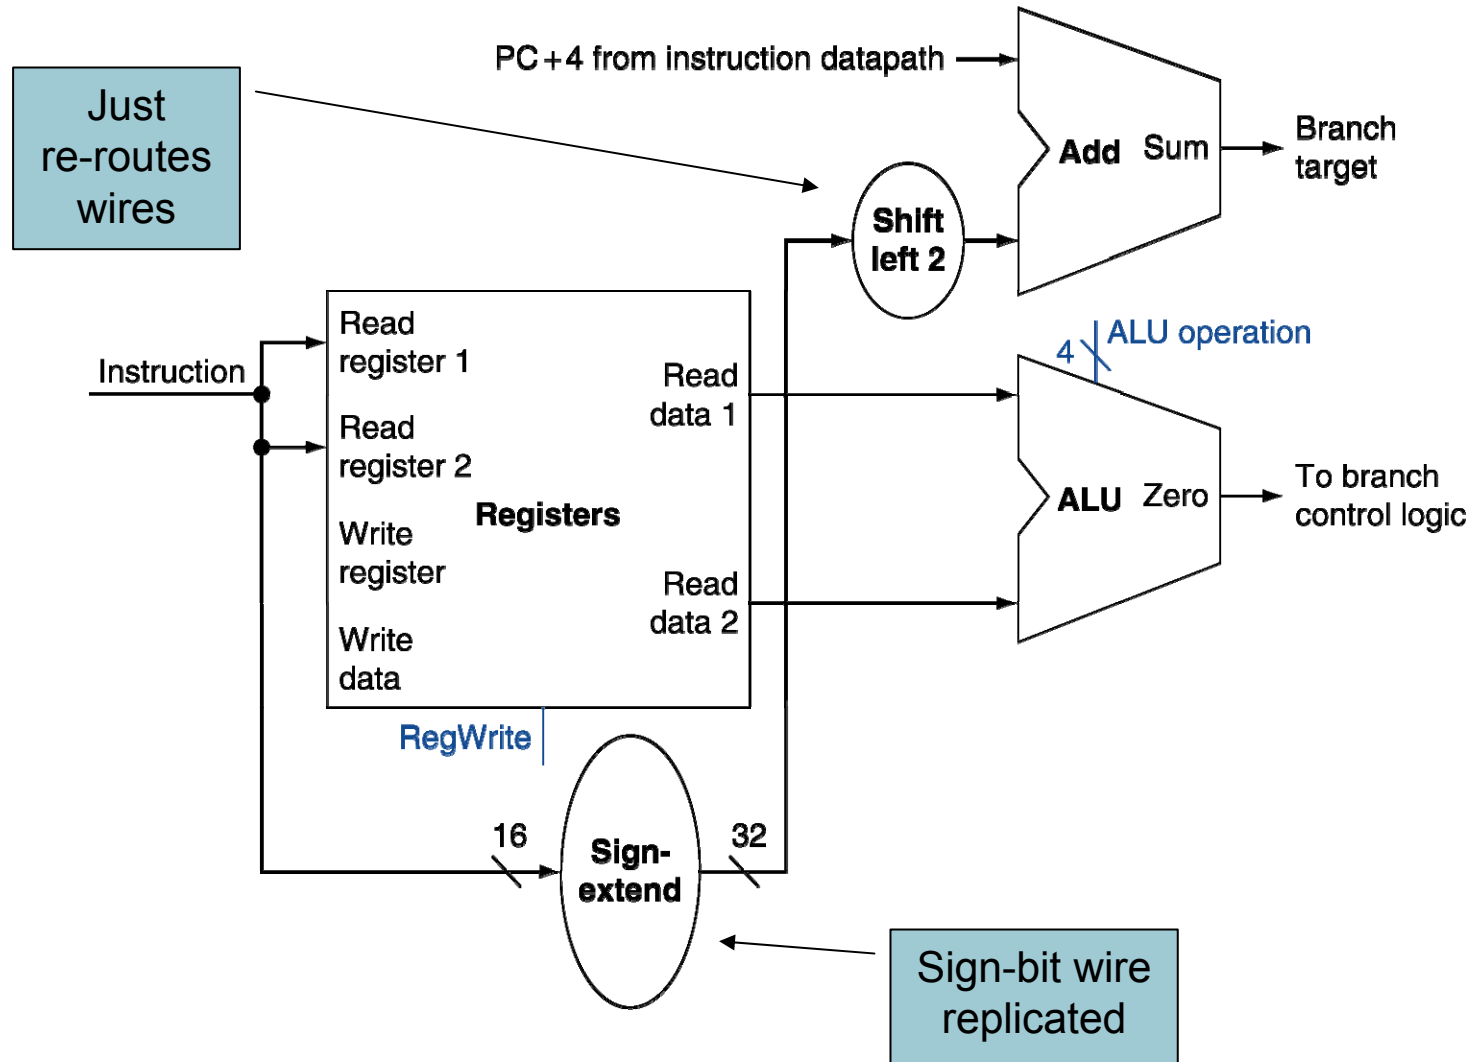
`else`

`PC <- PC + 4`

# Branch Instructions

- Read register operands
- Compare operands
  - Use ALU, subtract and check Zero output
- Calculate target address
  - Sign-extend displacement
  - Shift left 2 places (word displacement)
  - Add to PC + 4
    - Already calculated by instruction fetch

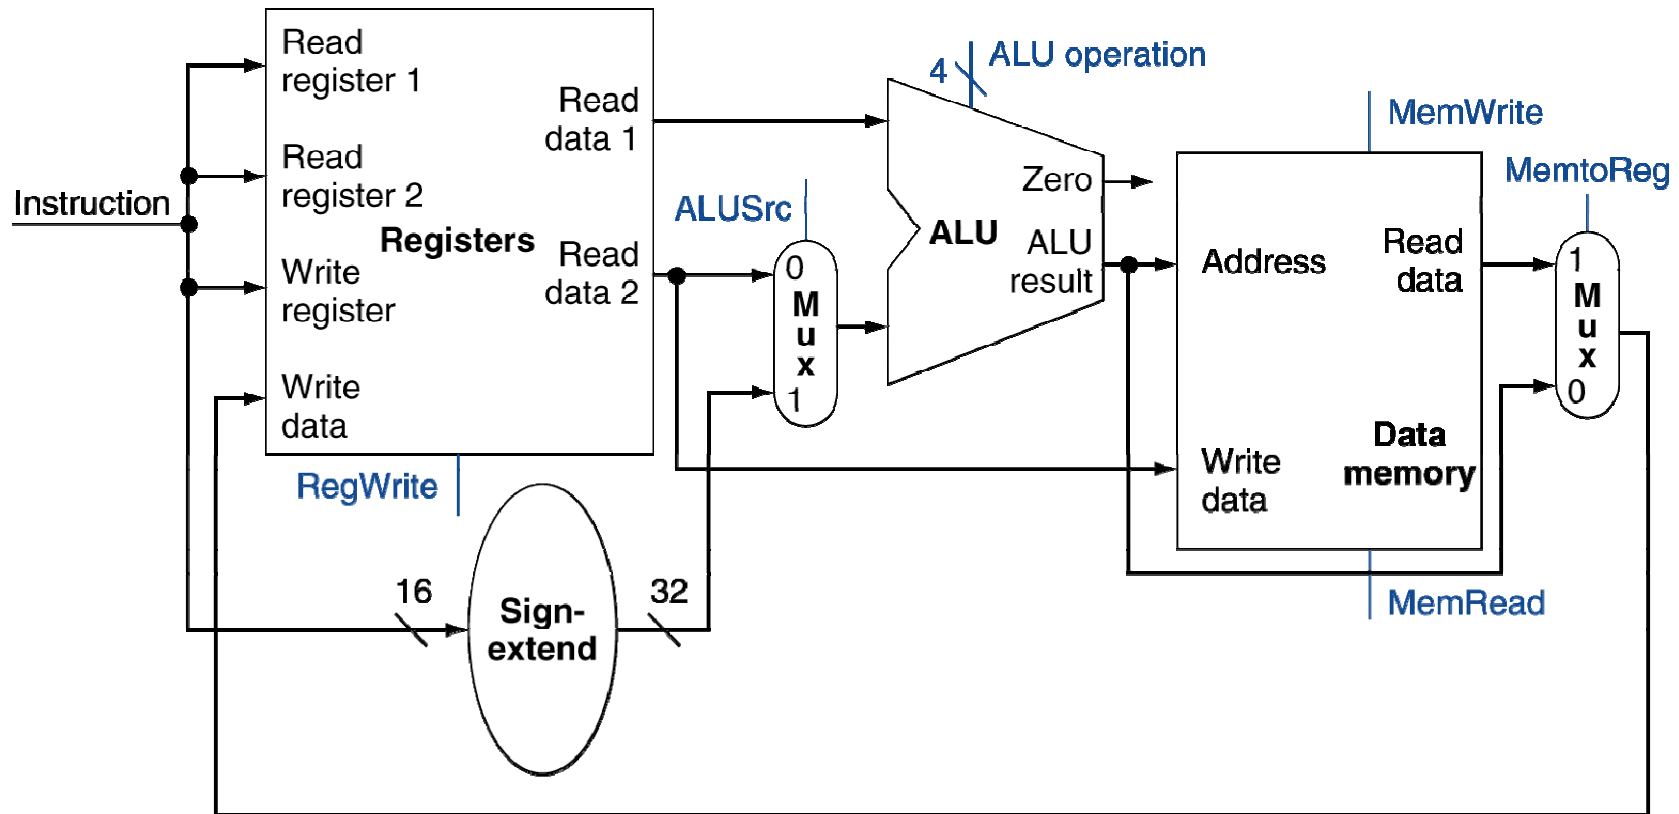
# Branch Instructions



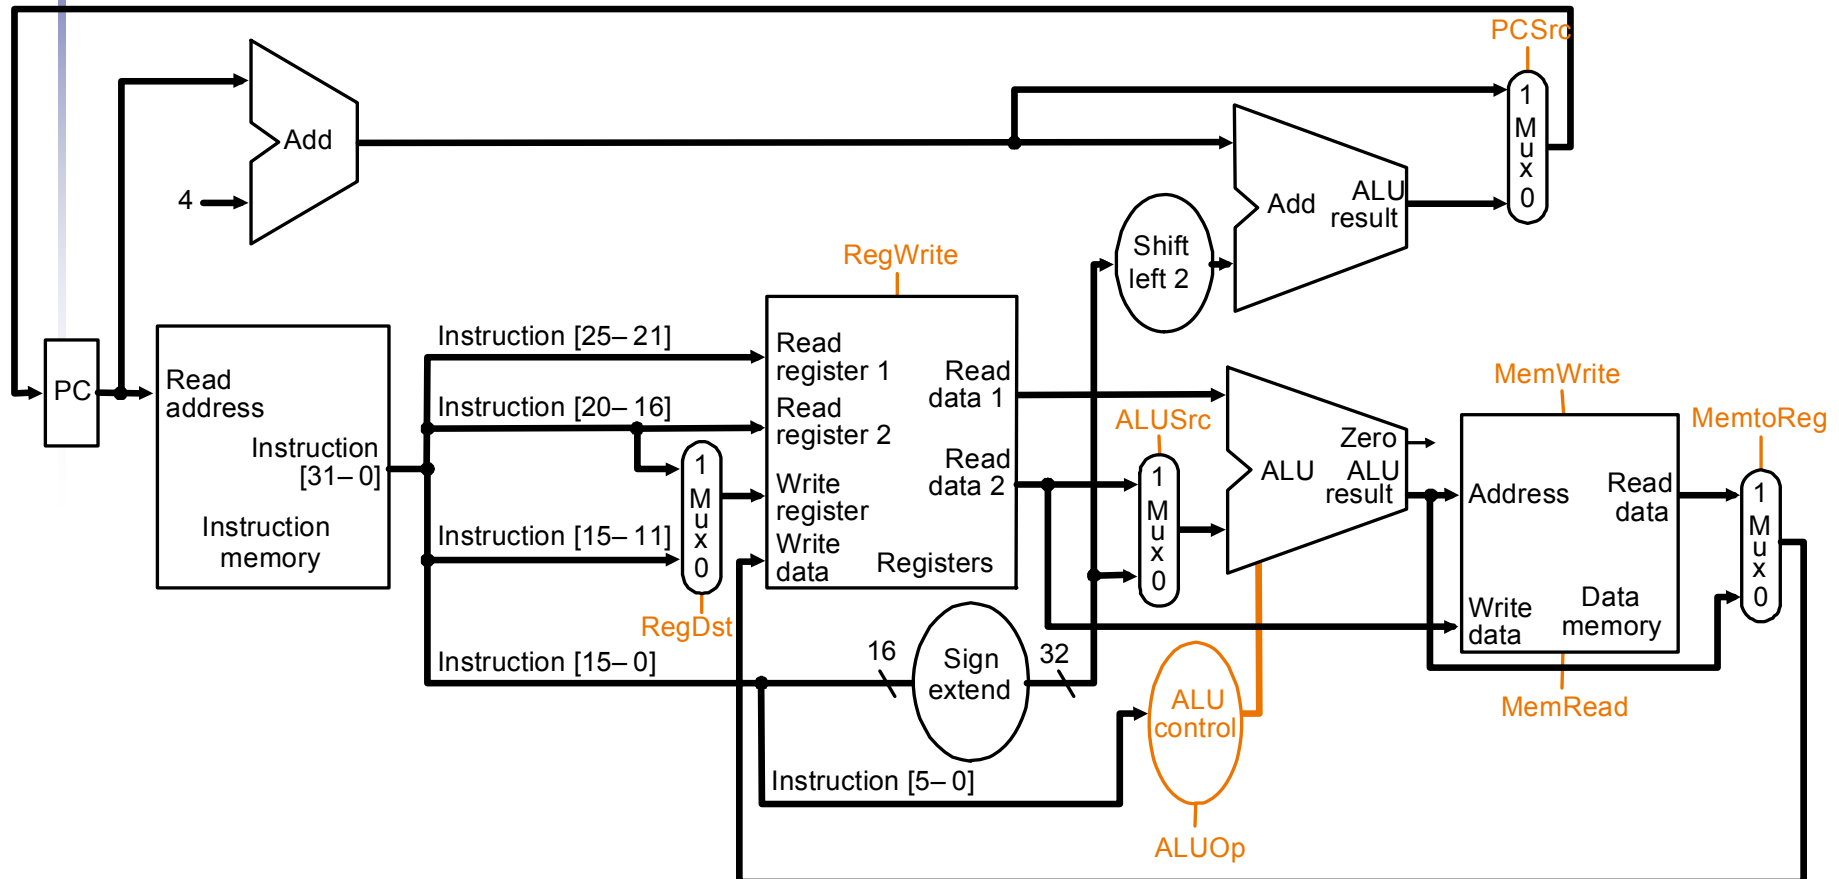
# Composing the Elements

- First-cut data path does an instruction in one clock cycle
  - Each datapath element can only do one function at a time
  - Hence, we need separate instruction and data memories
- Use **multiplexers** where alternate data sources are used for different instructions

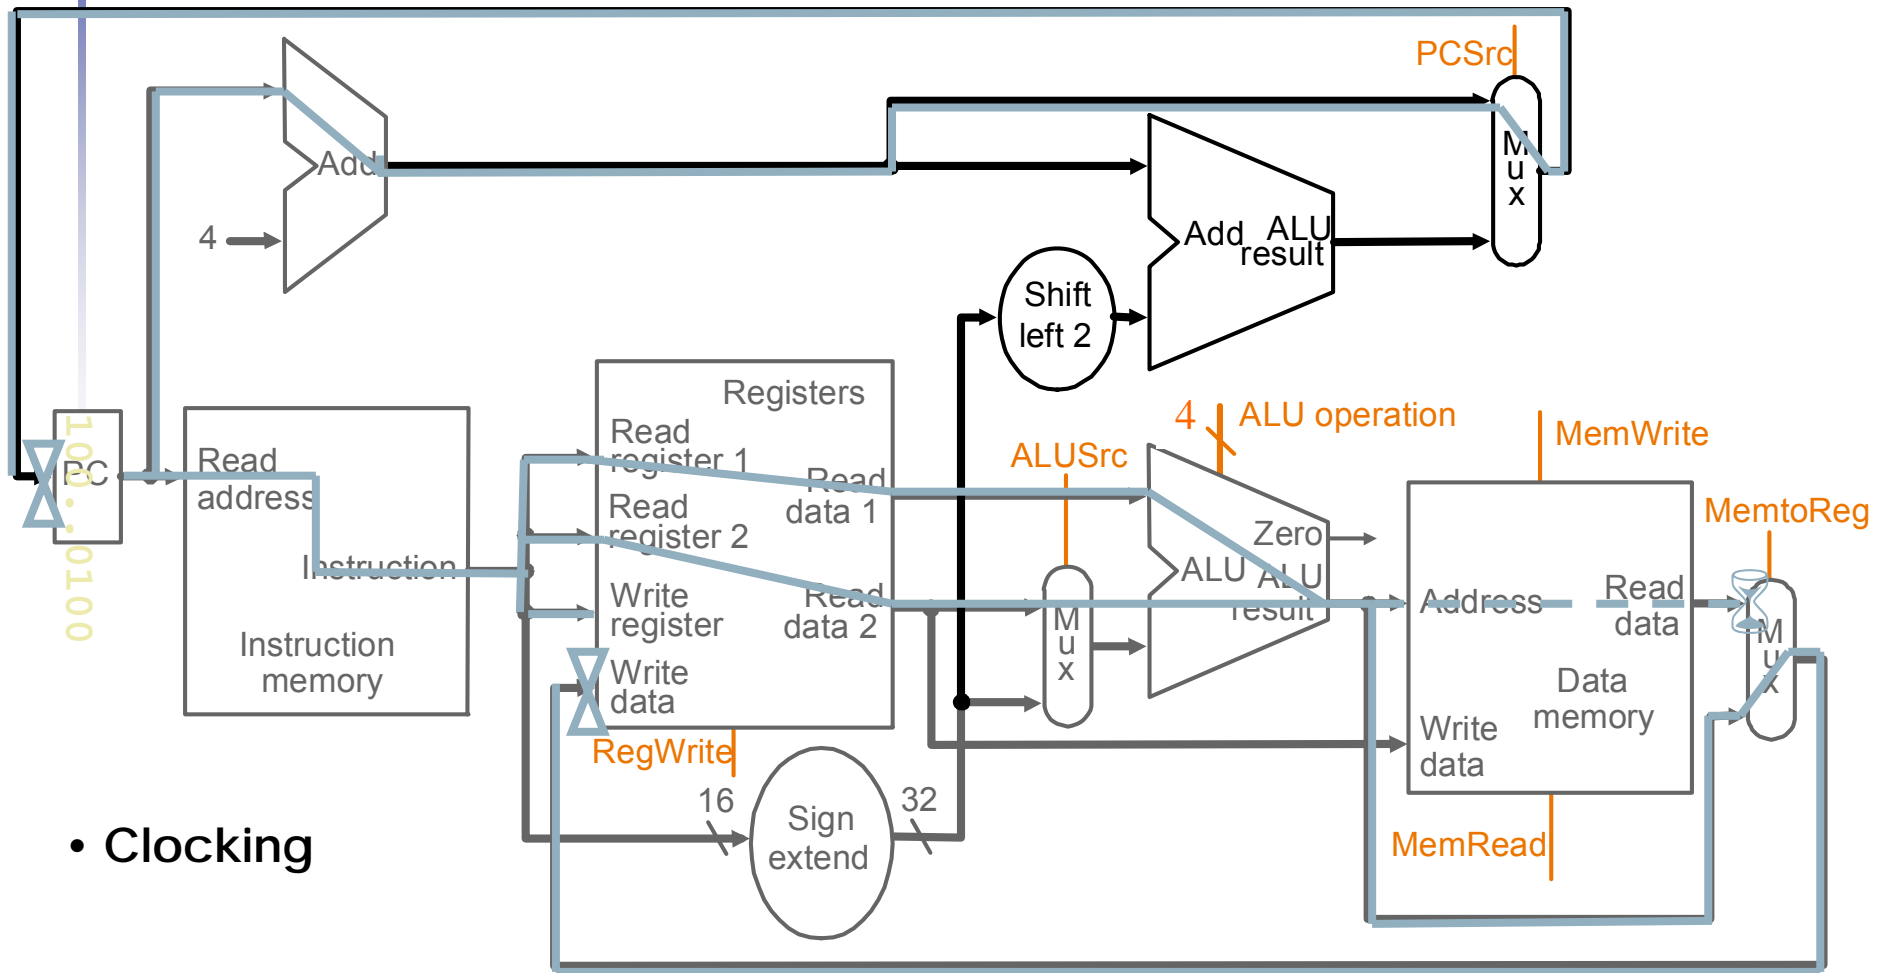
# R-Type/Load/Store Datapath



# A Single Cycle Datapath



# Data Flow during add



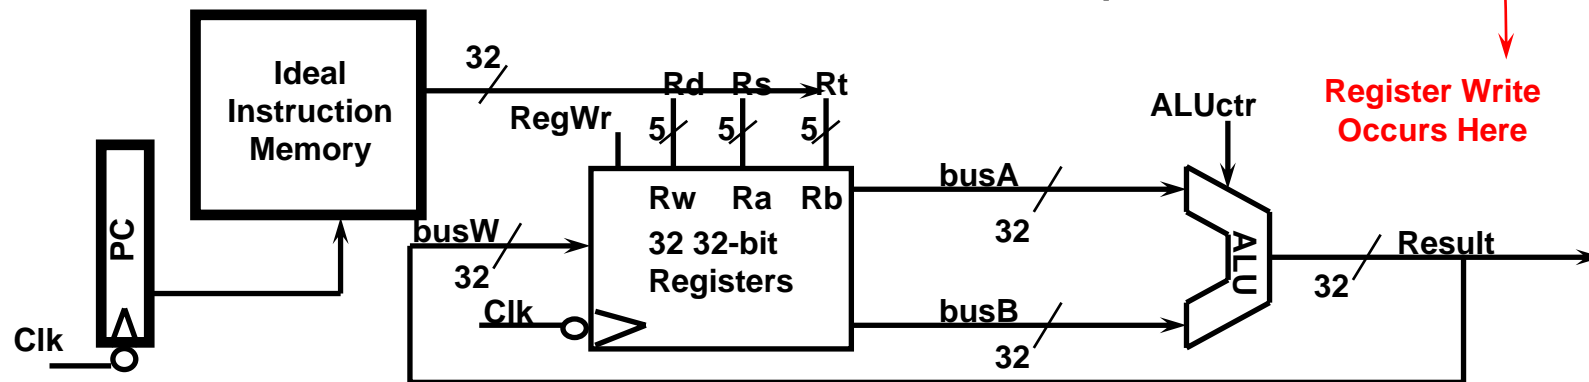
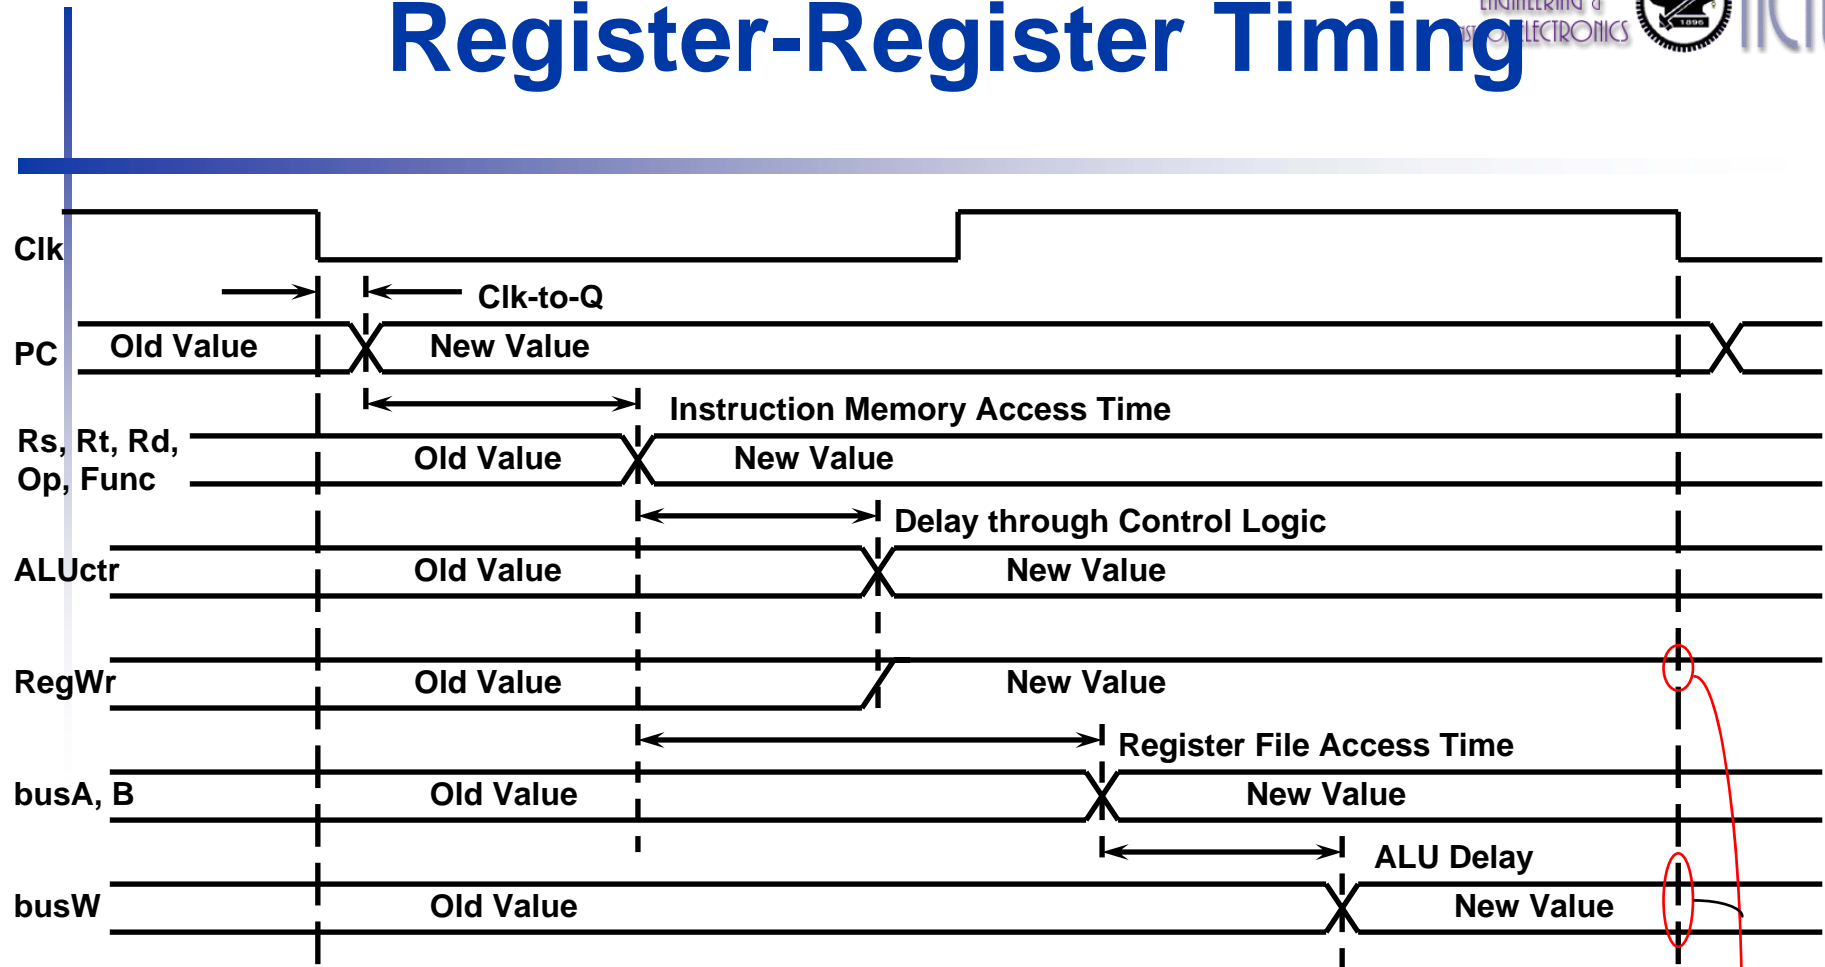
• Clocking



# Clocking Methodology

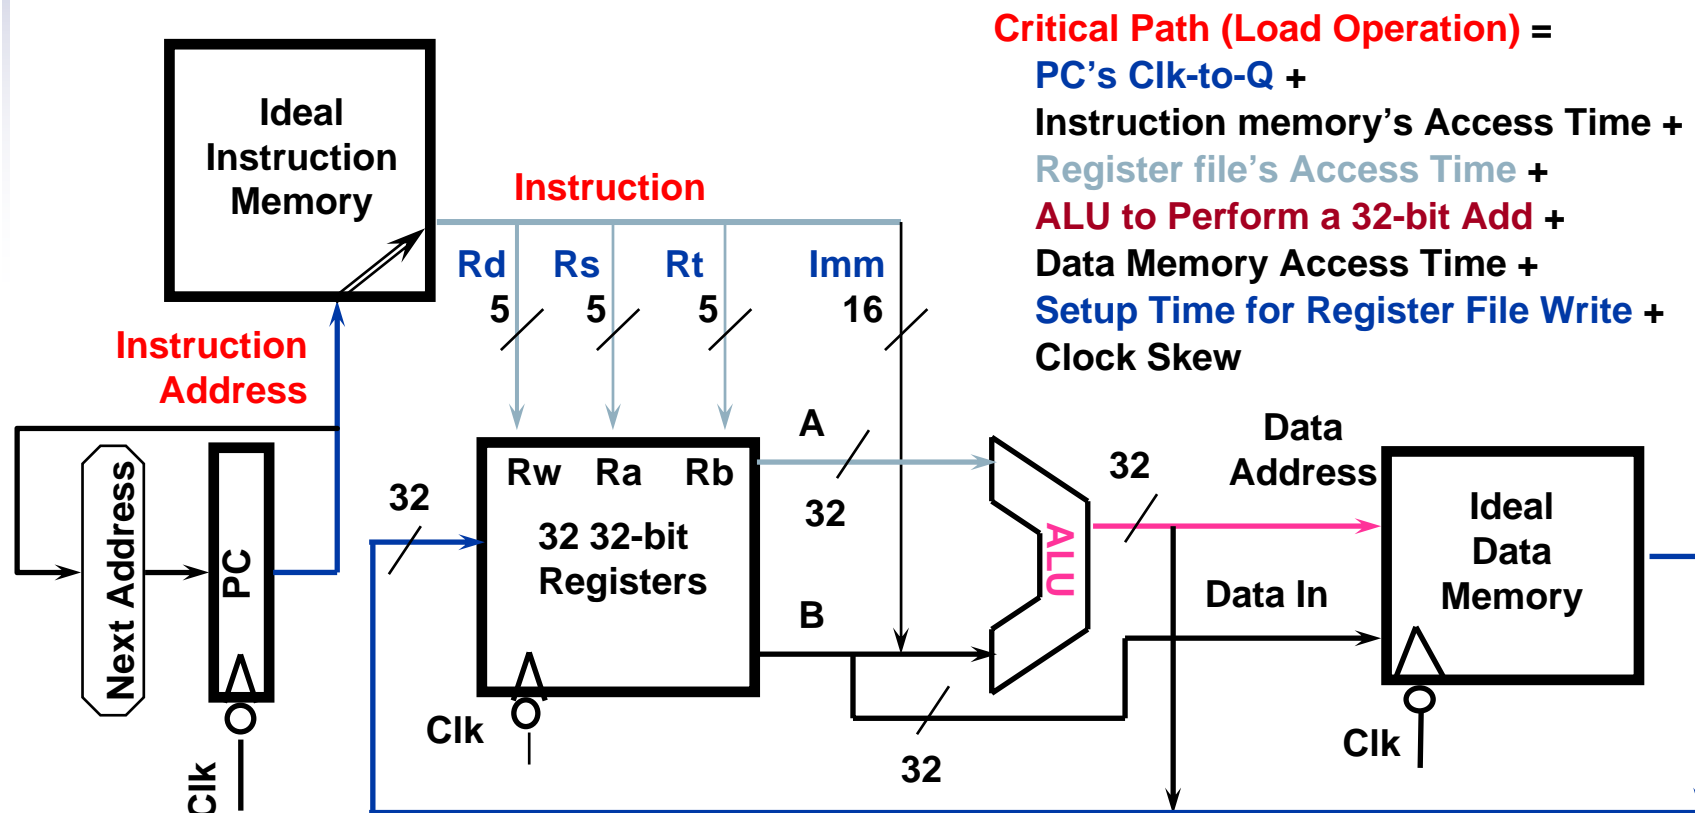
- Define when signals are read and written
- Assume edge-triggered (**synchronous design**):
  - Values in storage (state) elements updated only on a clock edge  
=> clock edge should arrive only after input signals stable
  - Any combinational circuit must have inputs from and outputs to storage elements
  - **Clock cycle**: time for signals to propagate from one storage element, through combinational circuit, to reach the second storage element
  - A register can be read, its value propagated through some combinational circuit, new value is written back to the same register, all in same cycle => no feedback within a single cycle

# Register-Register Timing

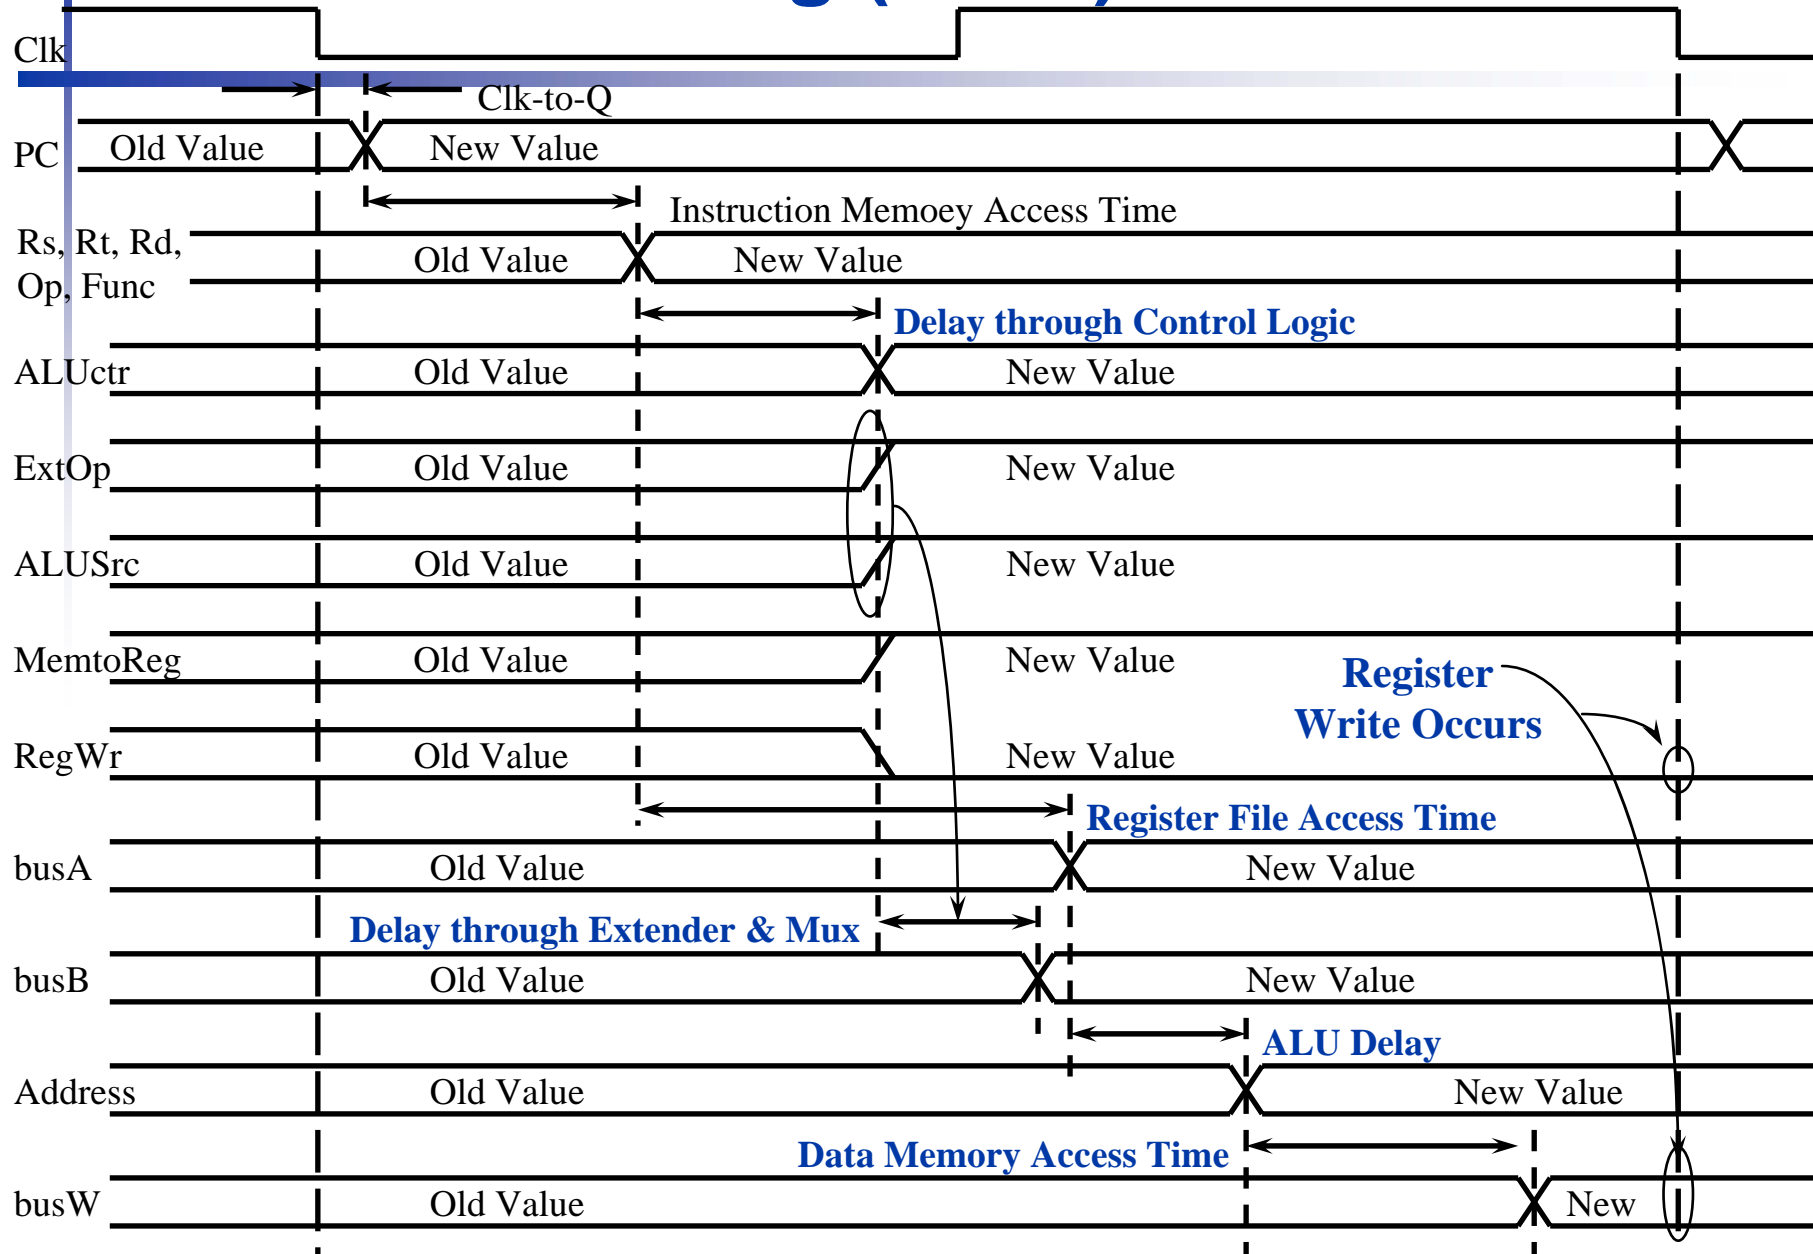


# The Critical Path

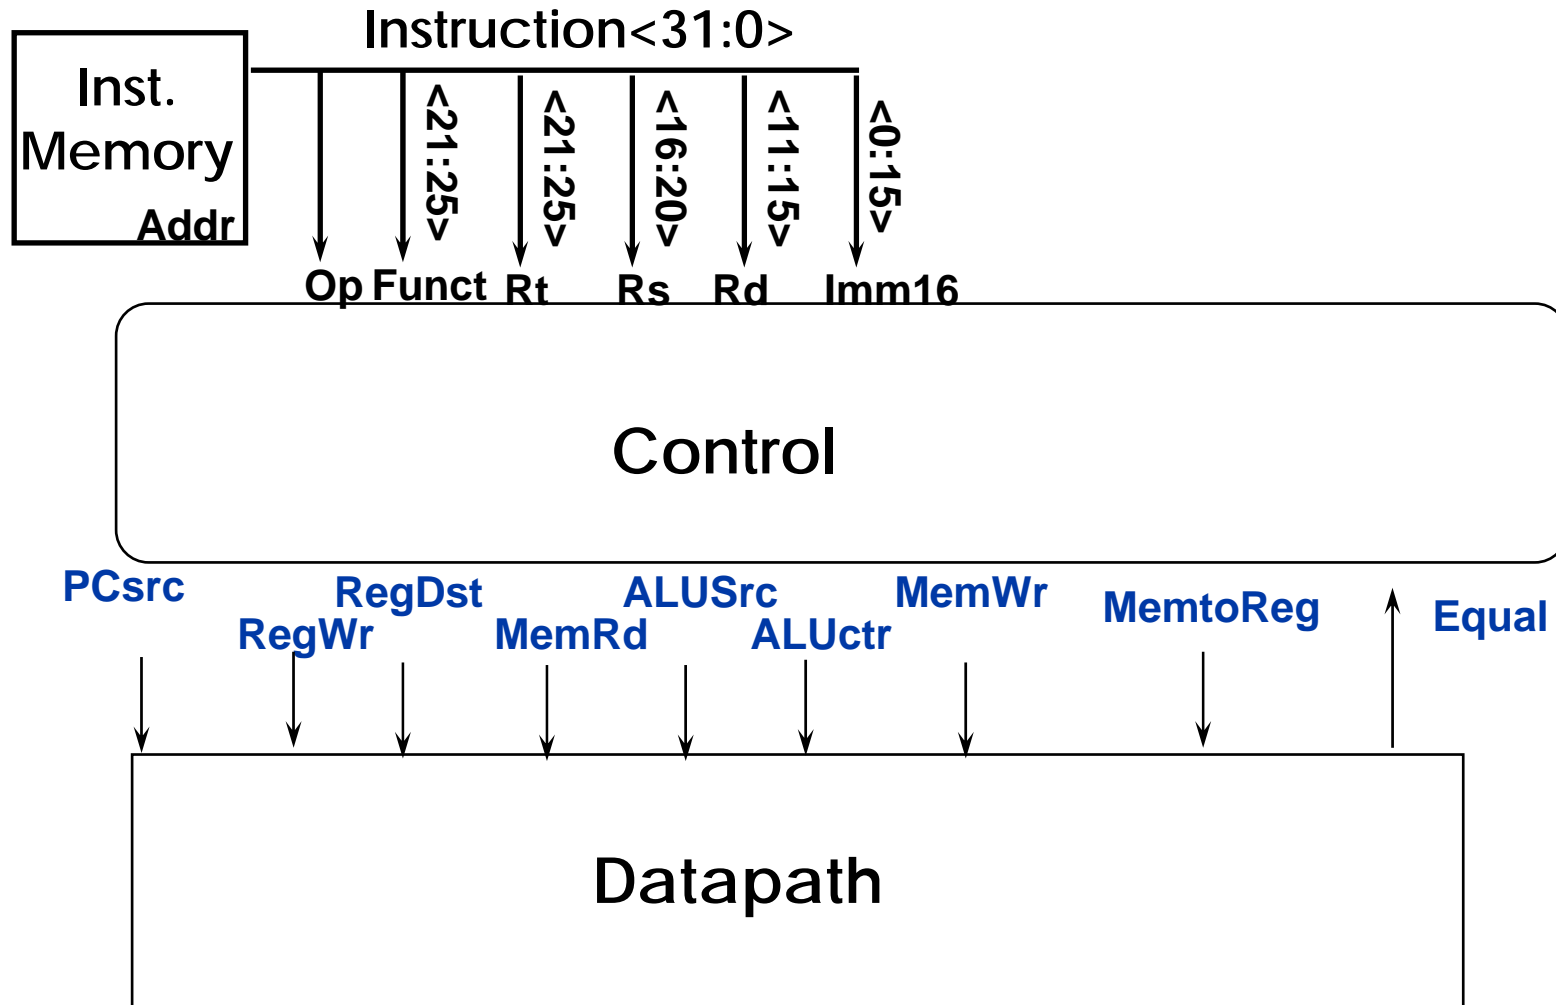
- Register file and ideal memory:
  - During read, behave as combinational logic:
    - Address valid => Output valid after access time



# Worst Case Timing (Load)



# Step 4: Control Points and Signals



# Control ?

- To select the operations to perform
  - ALU, read/write, etc
- To control the flow of data
  - Multiplexor inputs

# ALU Control

- ALU used for
  - Load/Store: F = add
  - Branch: F = subtract
  - R-type: F depends on funct field

ALU control	Function
0000	AND
0001	OR
0010	add
0110	subtract
0111	set-on-less-than
1100	NOR

# ALU Control

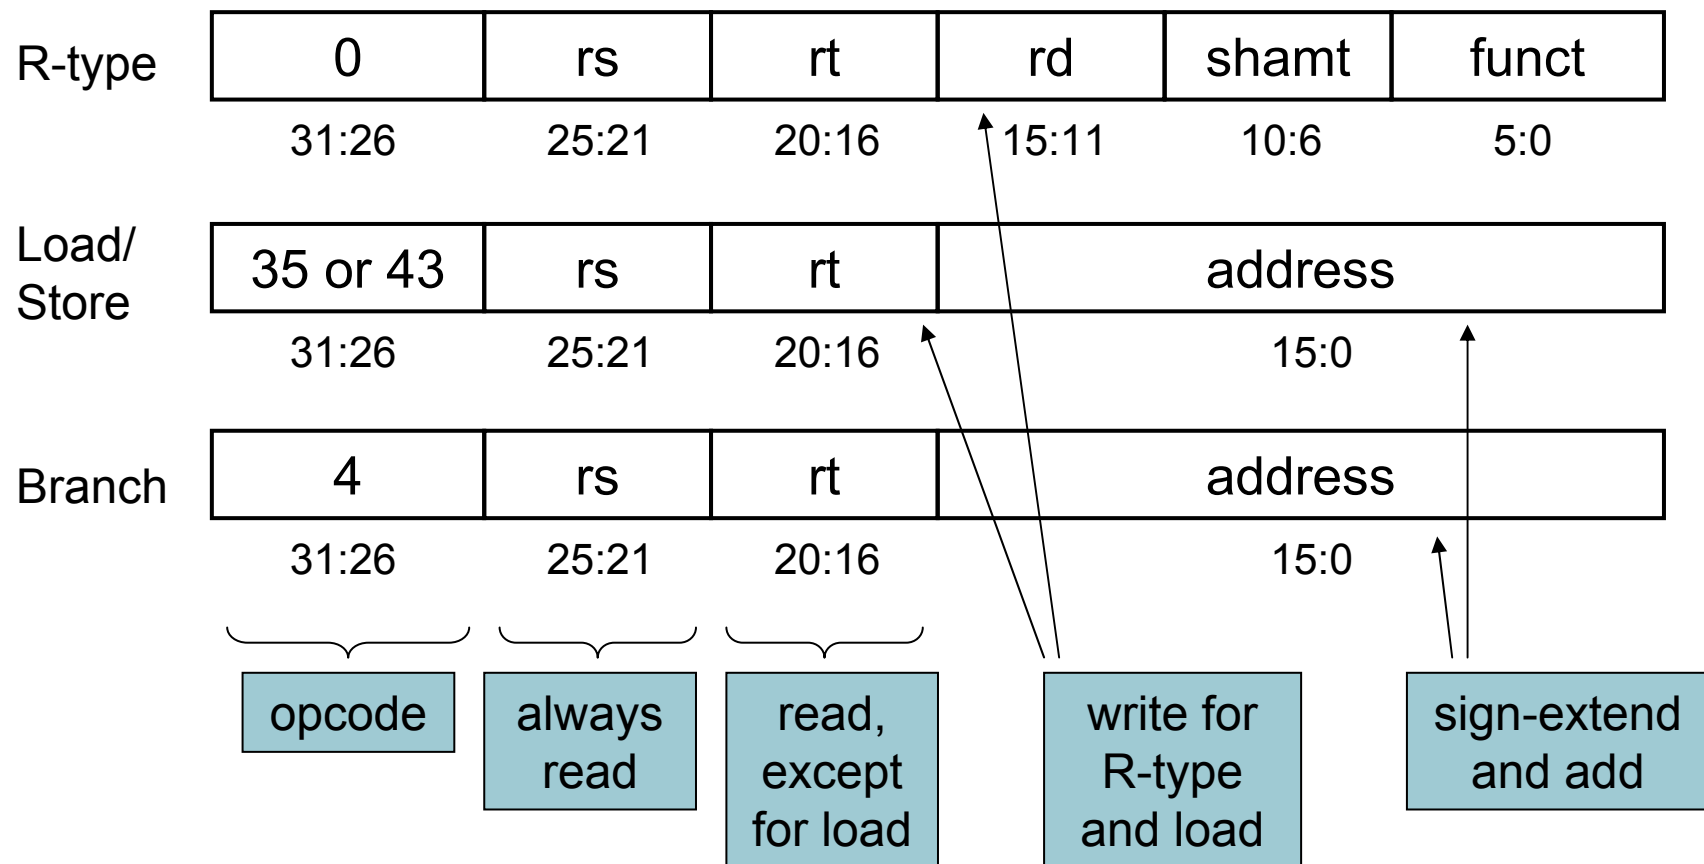
- Assume 2-bit ALUOp derived from opcode
  - Combinational logic derives ALU control

opcode	ALUOp	Operation	funct	ALU function	ALU control
lw	00	load word	XXXXXX	add	0010
sw	00	store word	XXXXXX	add	0010
beq	01	branch equal	XXXXXX	subtract	0110
R-type	10	add	100000	add	0010
		subtract	100010	subtract	0110
		AND	100100	AND	0000
		OR	100101	OR	0001
		set-on-less-than	101010	set-on-less-than	0111



# The Main Control Unit

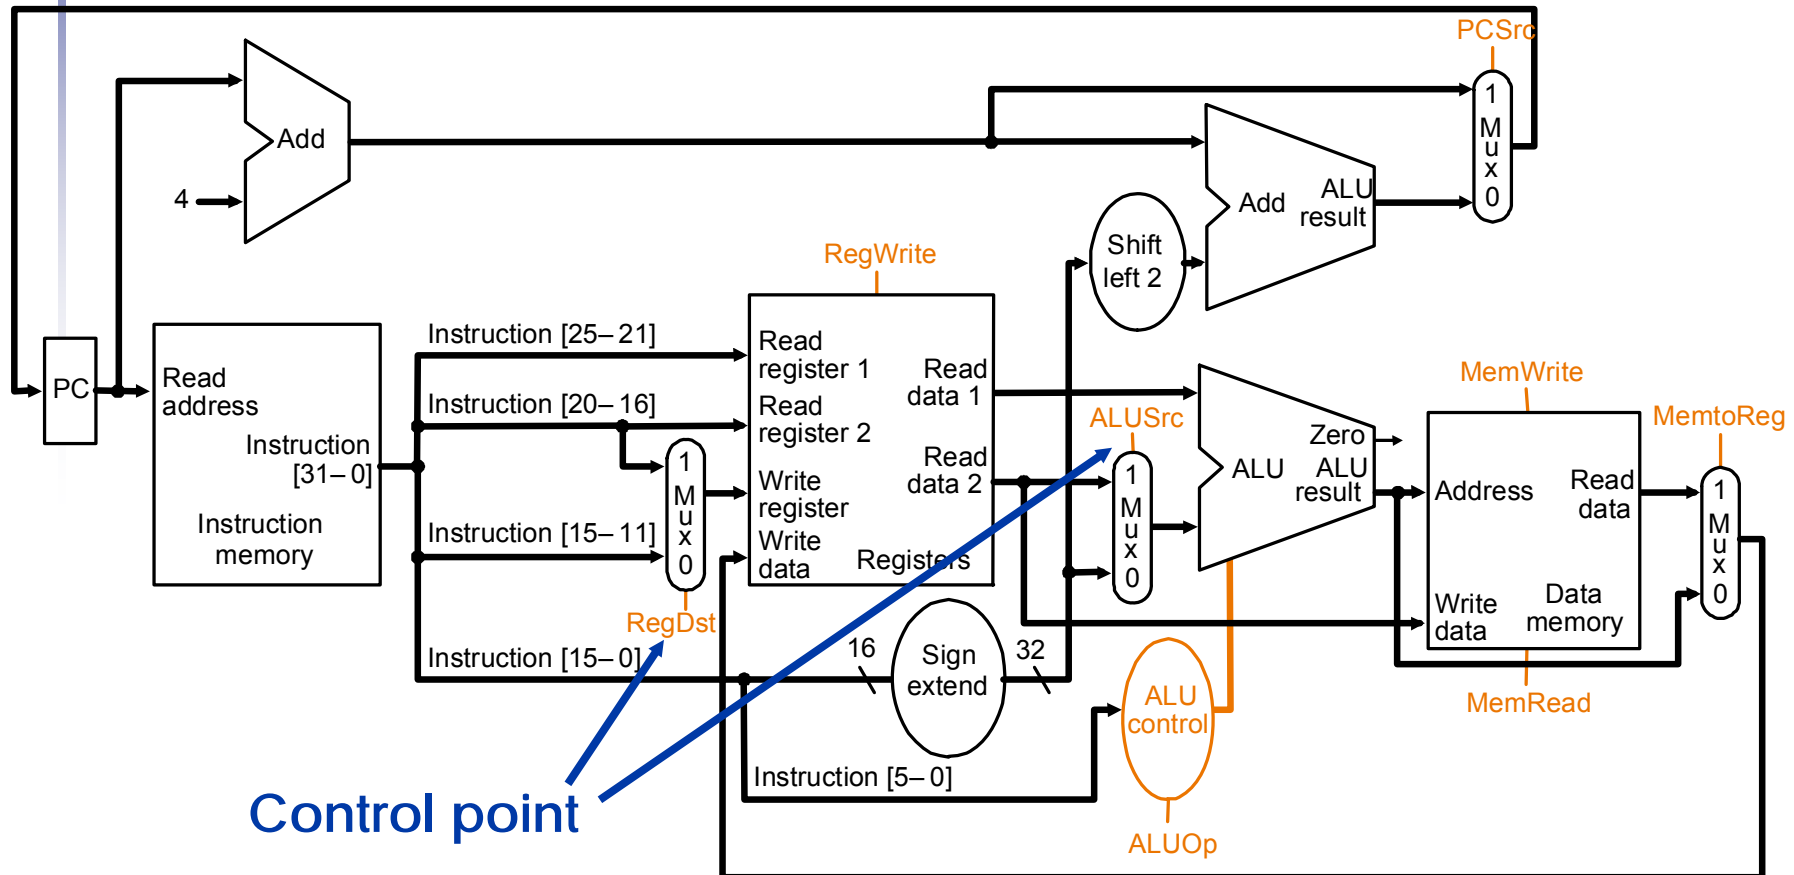
- Control signals derived from instruction



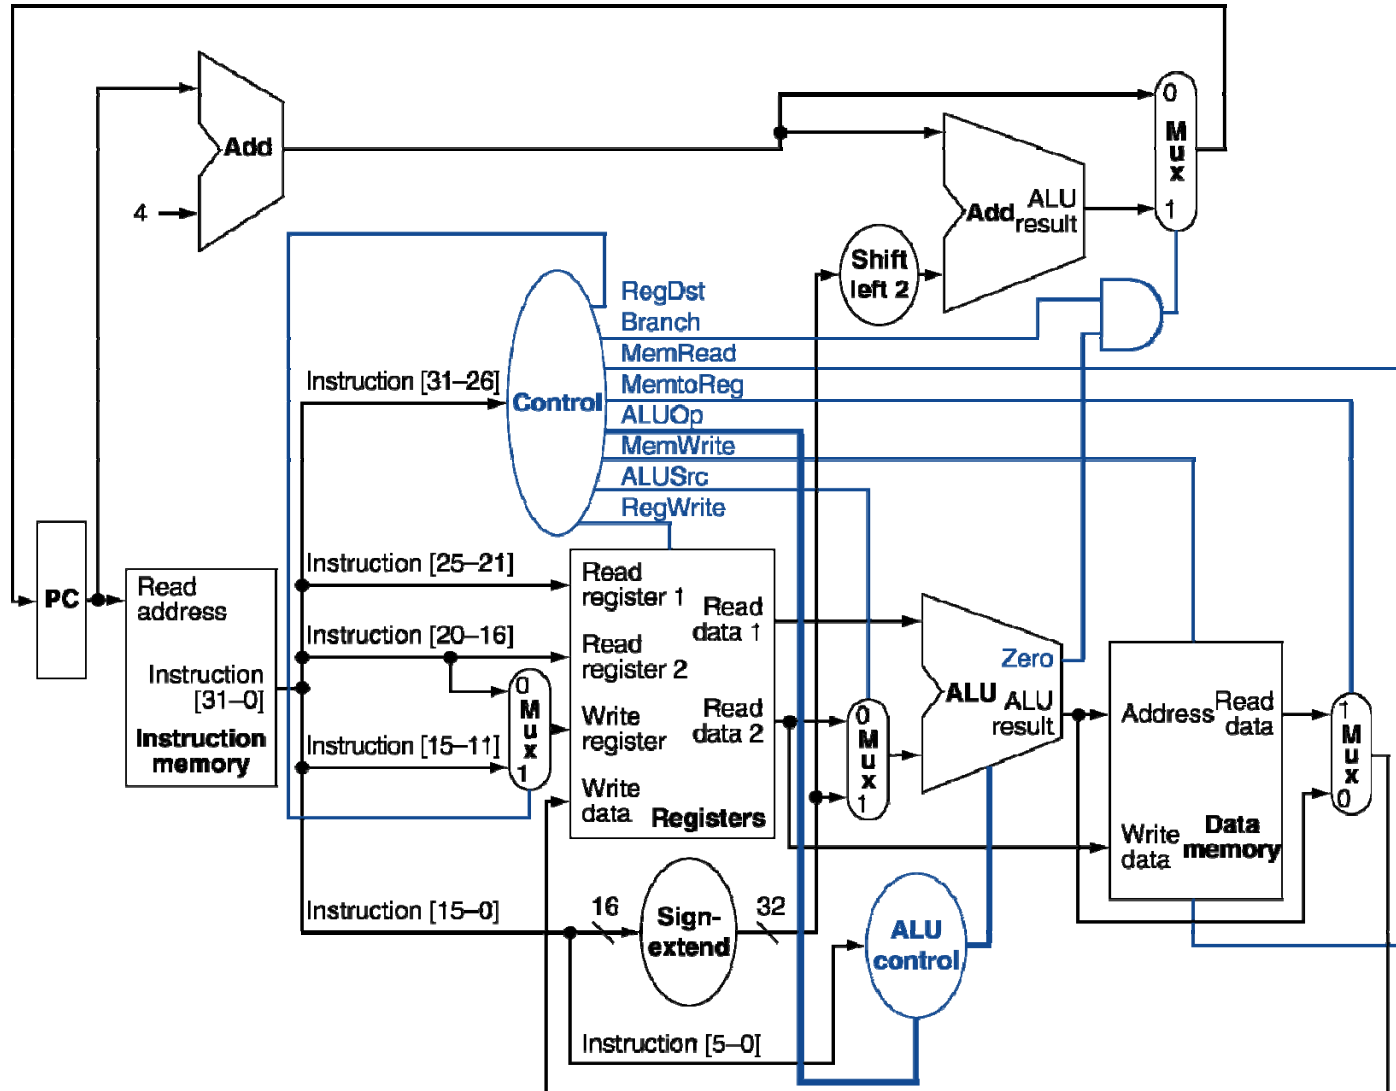
# Designing Main Control

- Some observations:
    - opcode (Op[5-0]) is always in bits 31-26
    - two registers to be read are always in rs (bits 25-21) and rt (bits 20-16) (for R-type, beq, sw)
    - base register for lw and sw is always in rs (25-21)
    - 16-bit offset for beq, lw, sw is always in 15-0
    - destination register is in one of two positions:
      - lw: in bits 20-16 (rt)
      - R-type: in bits 15-11 (rd)
- => need a multiplex to select the address for written register

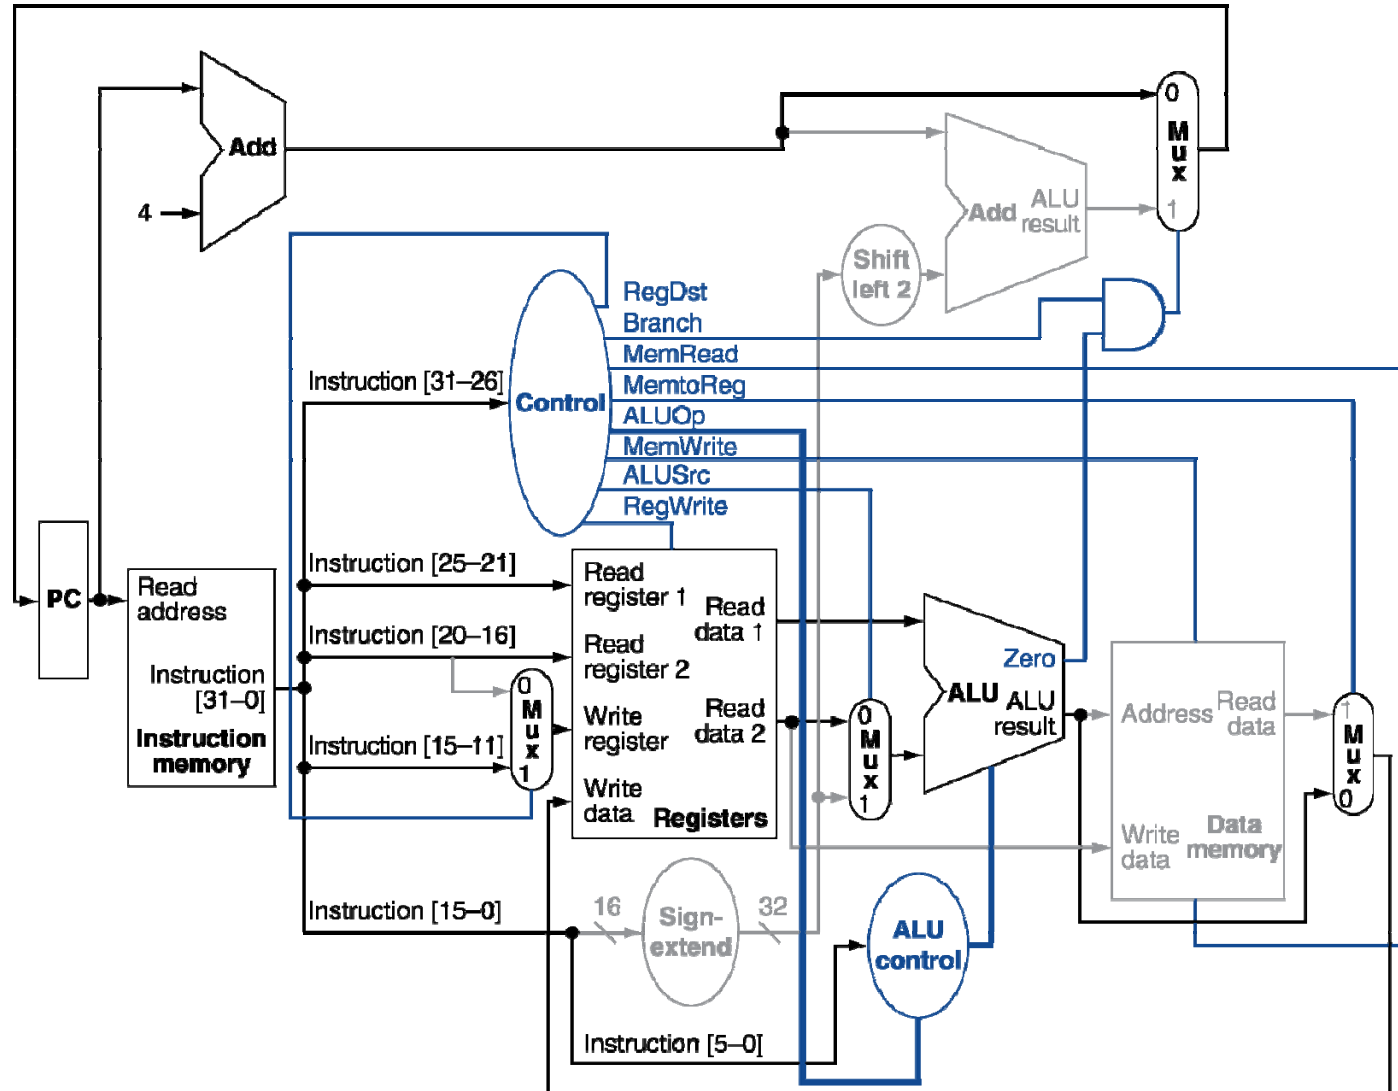
# Datapath with Mux and Control



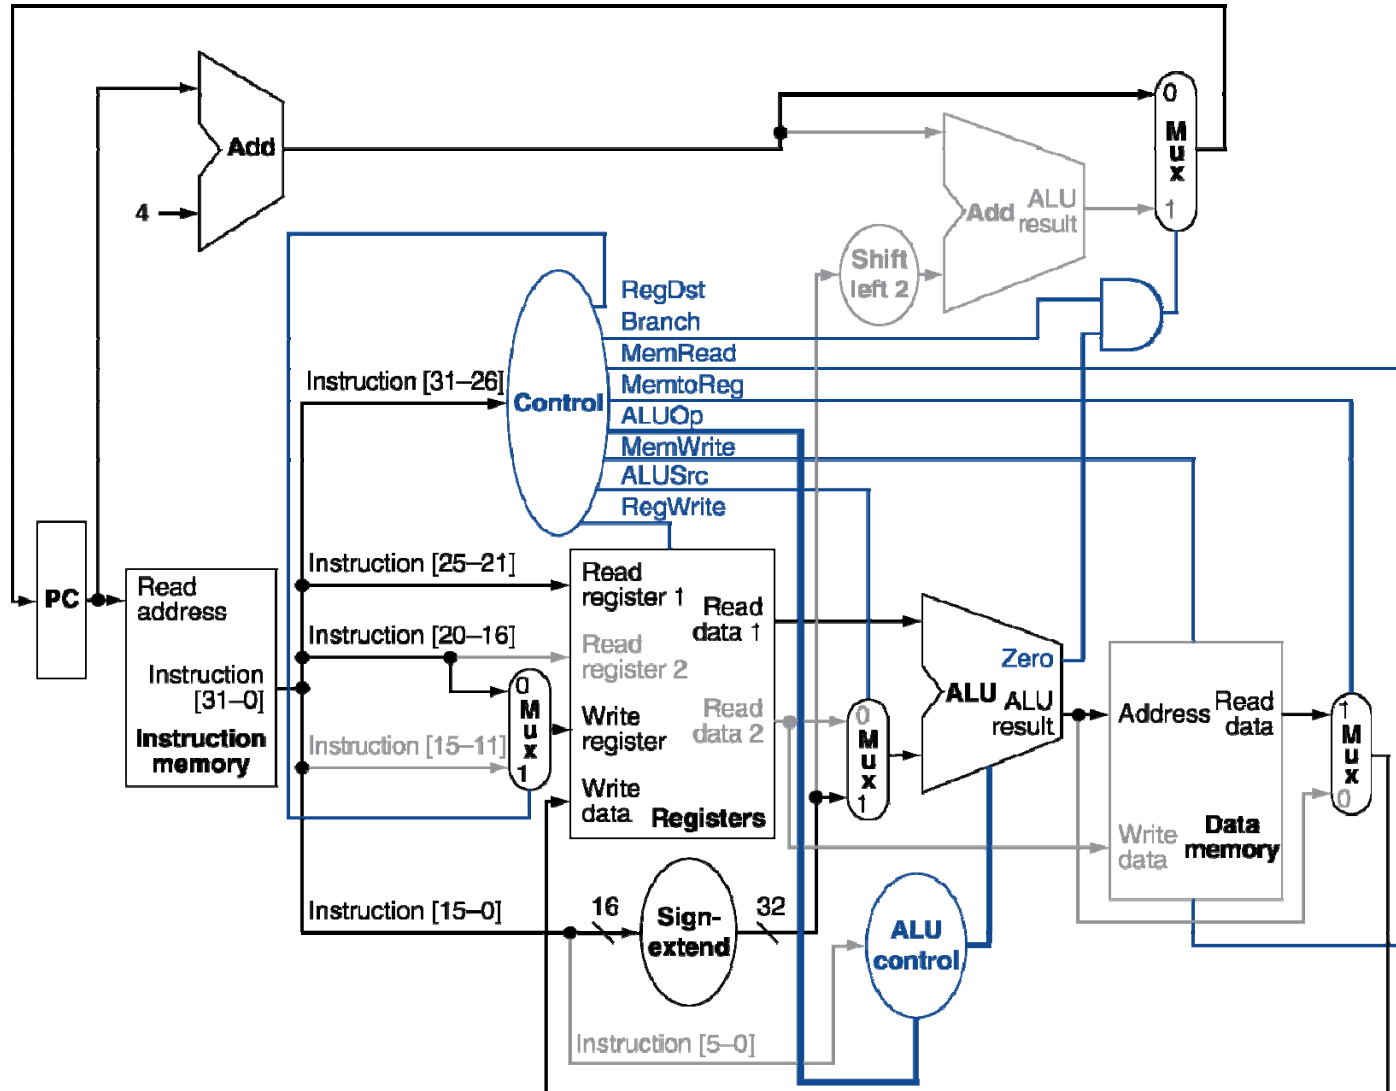
# Datapath With Control



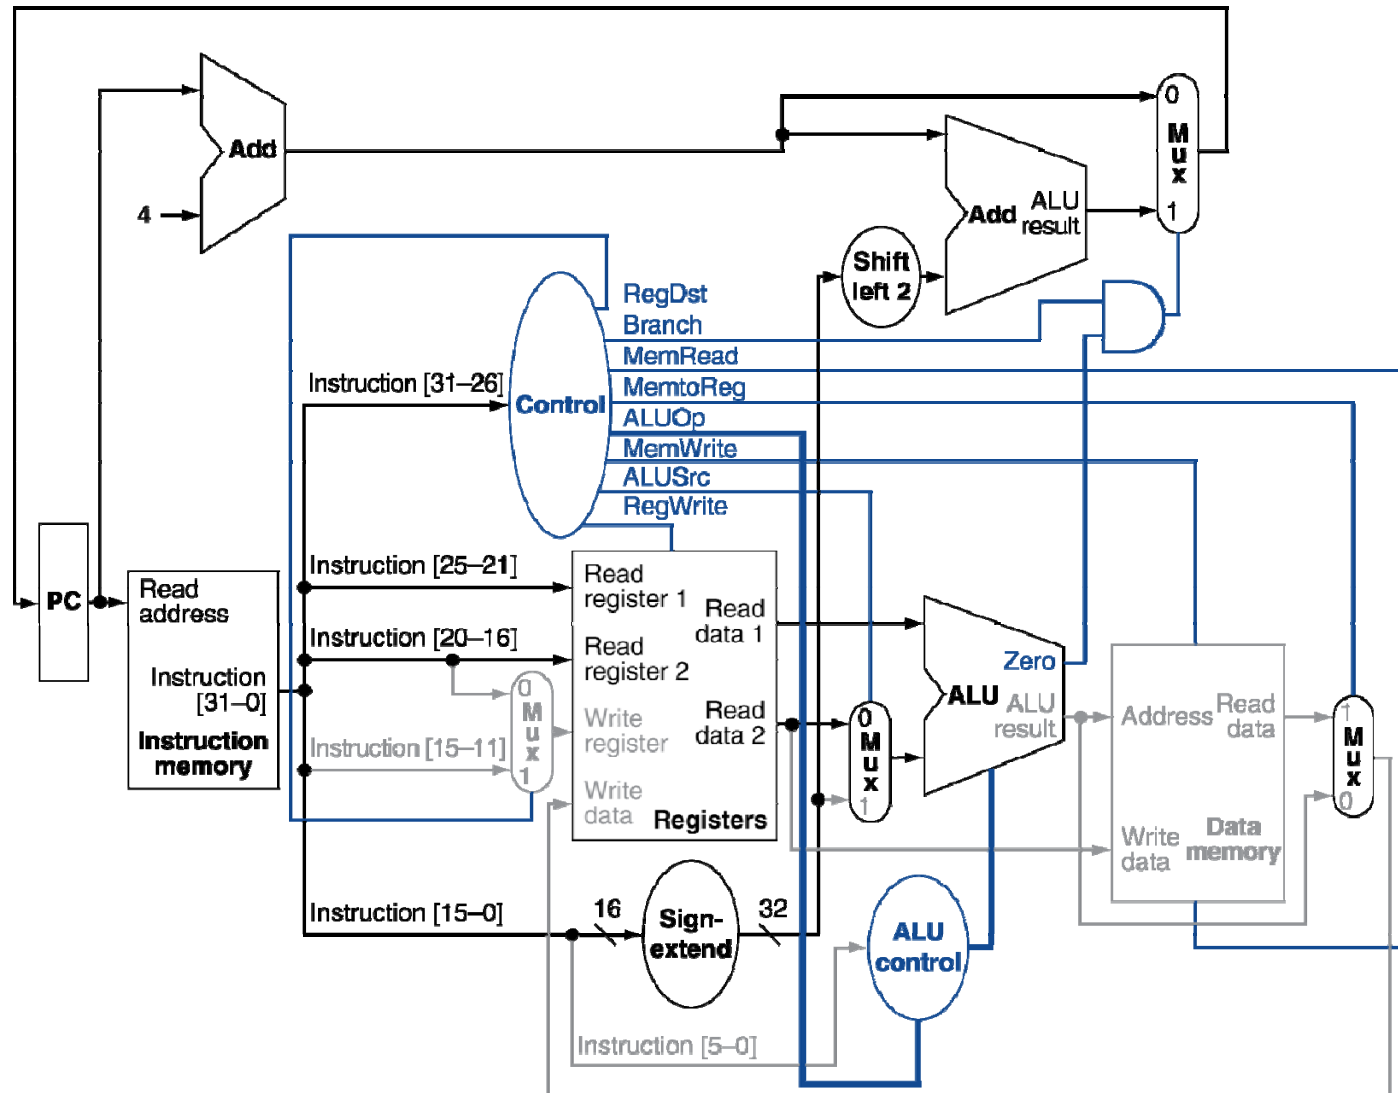
# R-Type Instruction




# Load Instruction



# Branch-on-Equal Instruction



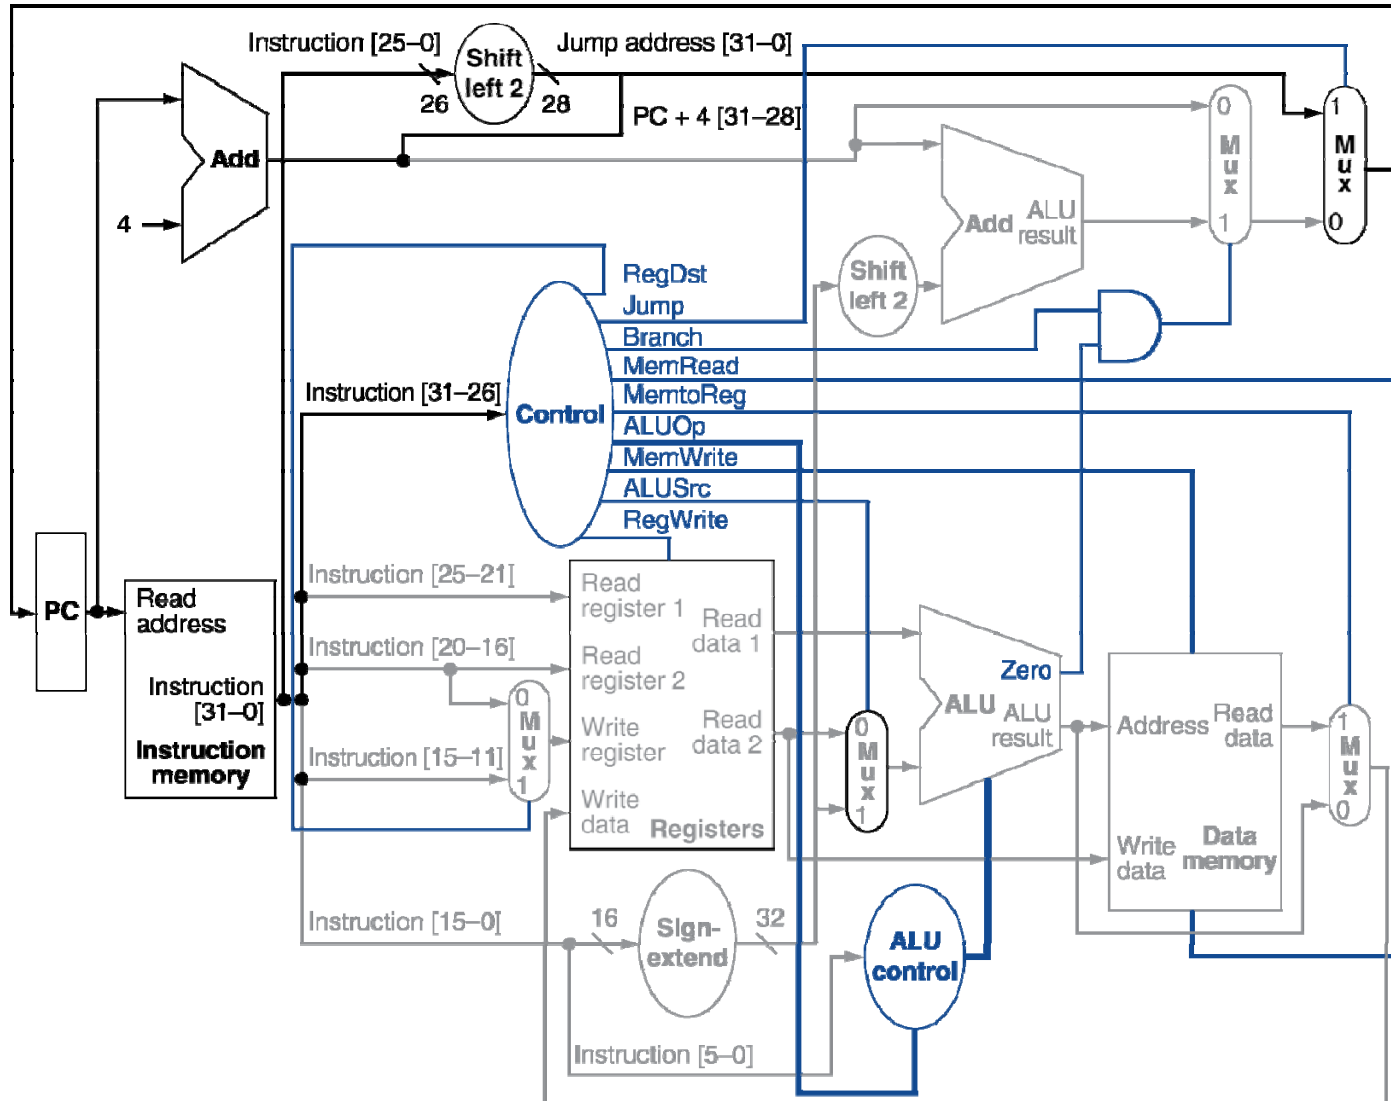
# Implementing Jumps



- **Jump uses word address**
- Update PC with concatenation of
  - Top 4 bits of old PC
  - 26-bit jump address
  - 00
- Need an extra control signal decoded from opcode



# Datapath With Jumps Added

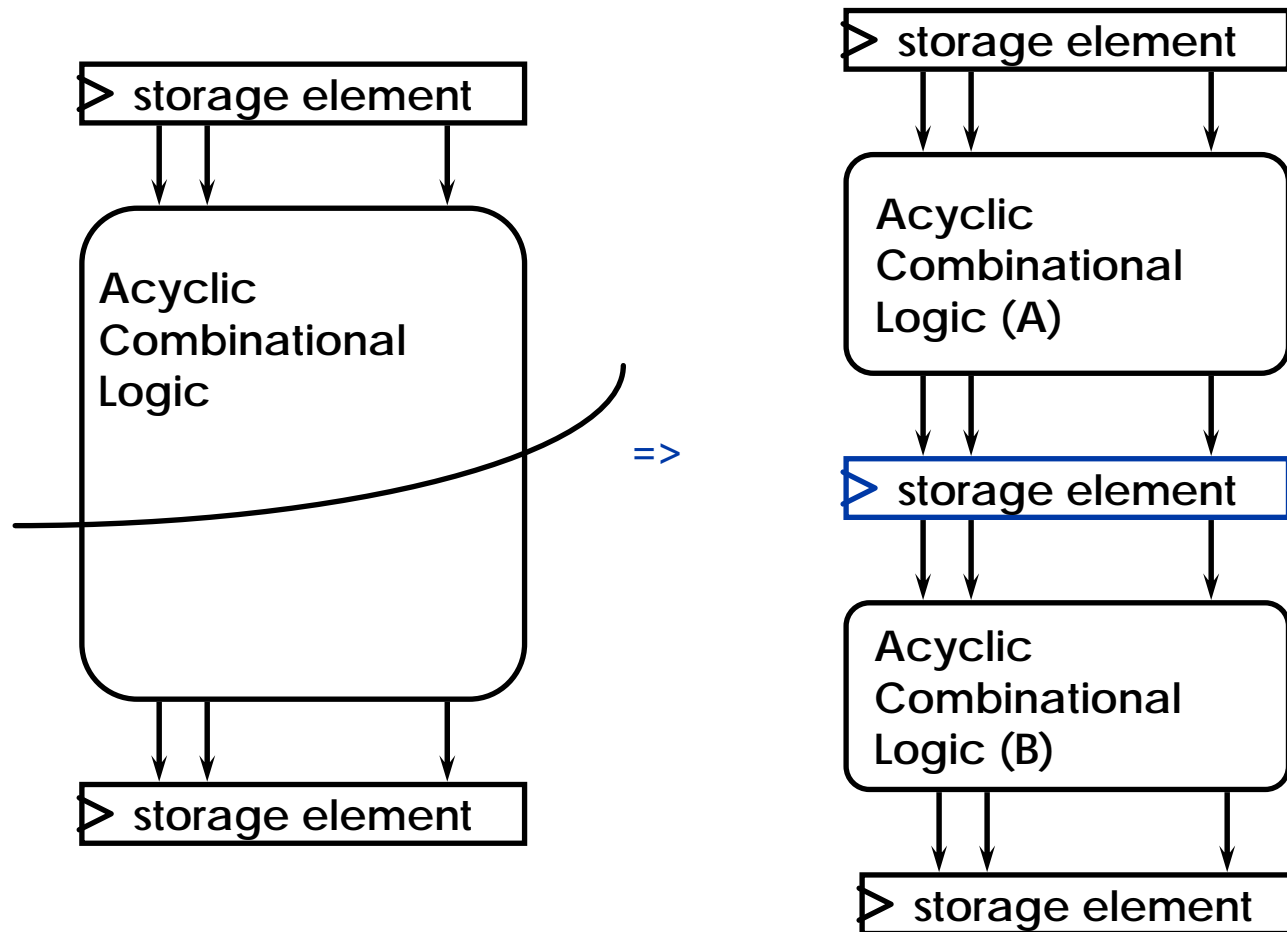


# Performance Issues

- Longest delay determines clock period
  - Critical path: load instruction
  - Instruction memory → register file → ALU → data memory → register file
- Not feasible to vary period for different instructions
- Violates design principle
  - Making the common case fast
- **We will improve performance by pipelining**

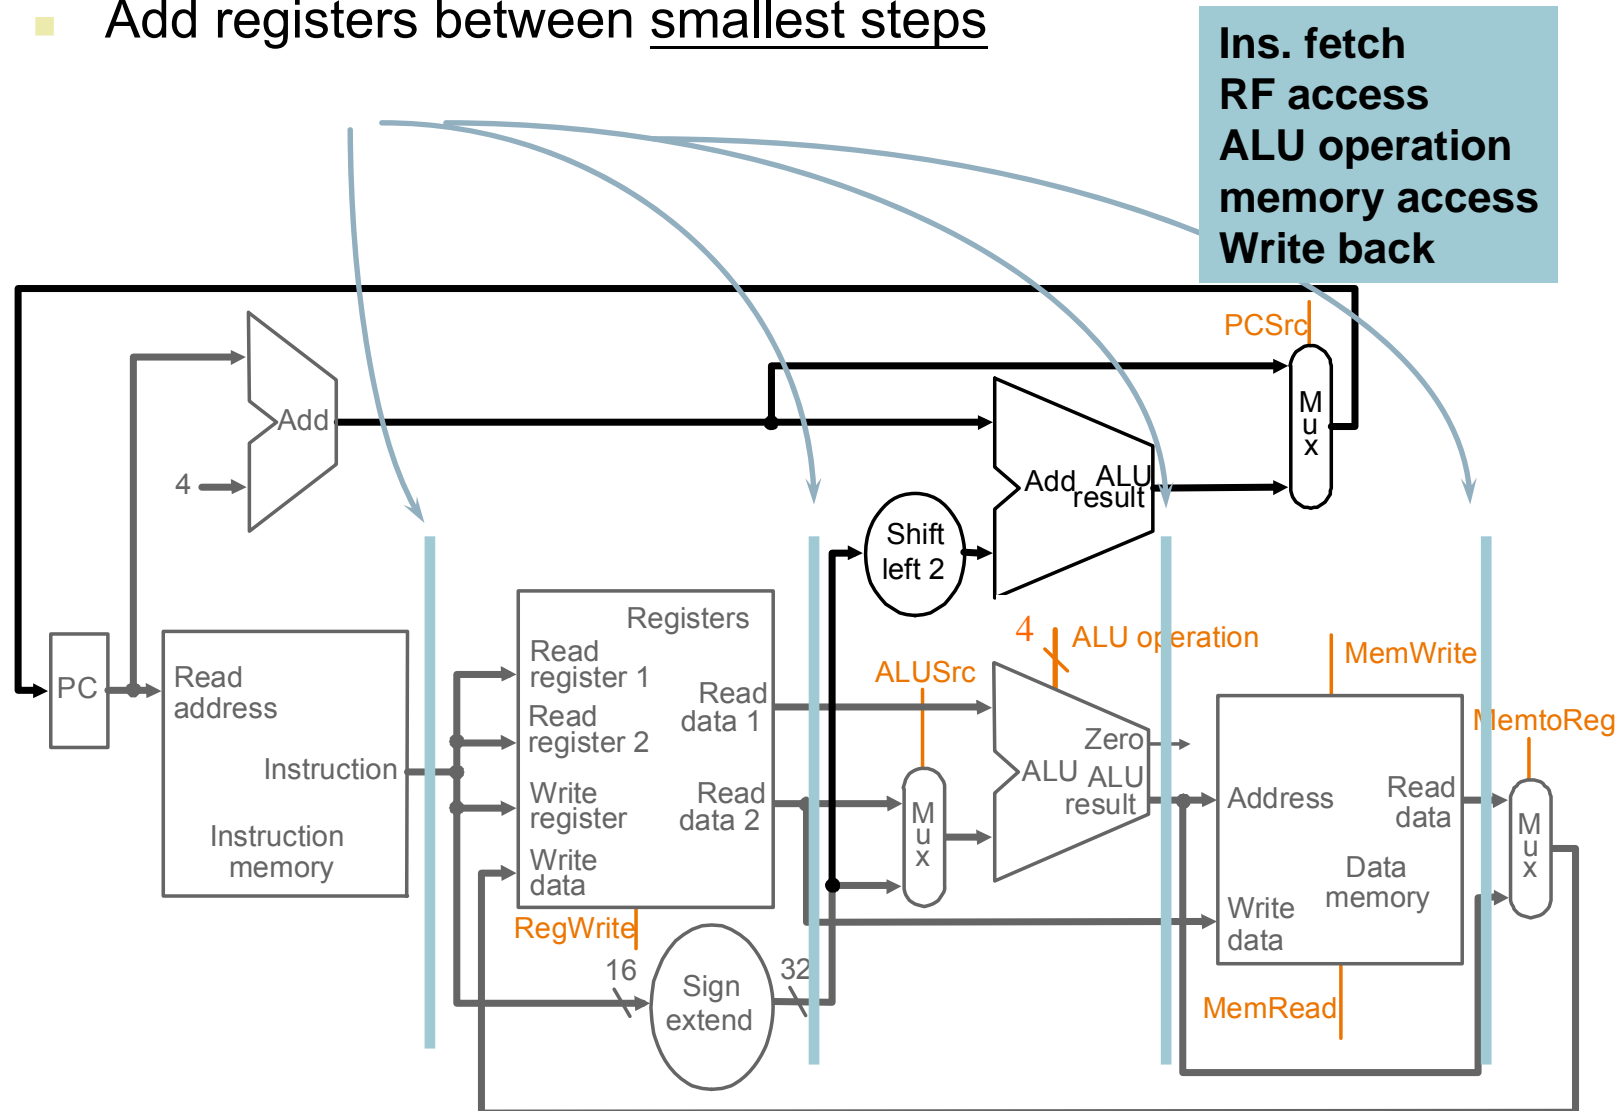
# Multistage Implementation

- Critical path reduction





# Partition Single-Cycle Datapath

- Add registers between smallest steps



# Pipelining Analogy

- Pipelined laundry: overlapping execution
  - Parallelism improves performance



# 5-Stage MIPS Pipeline

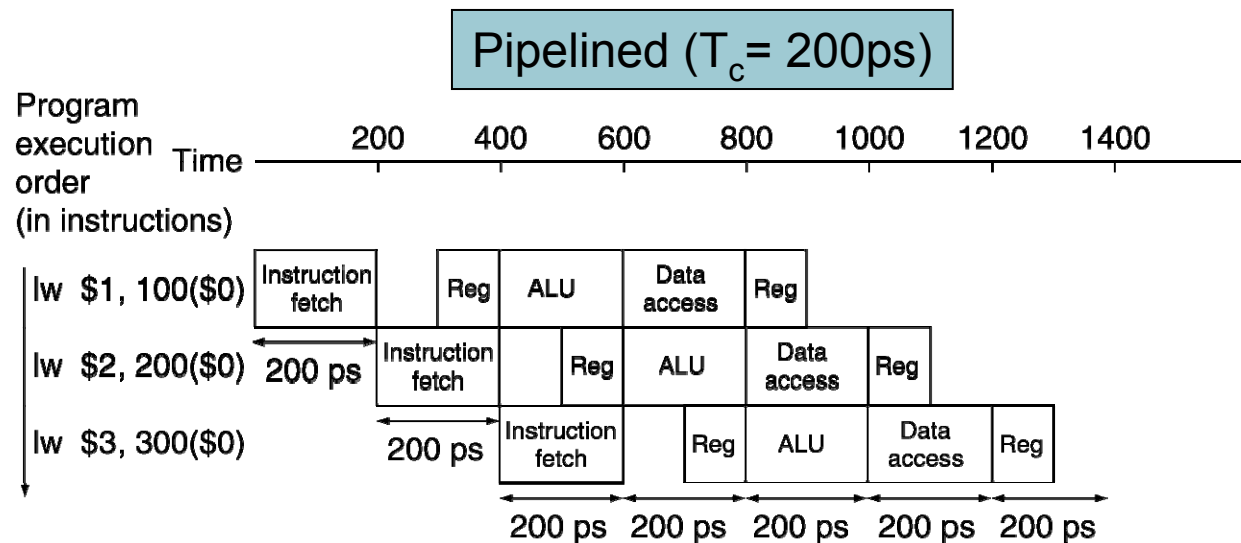
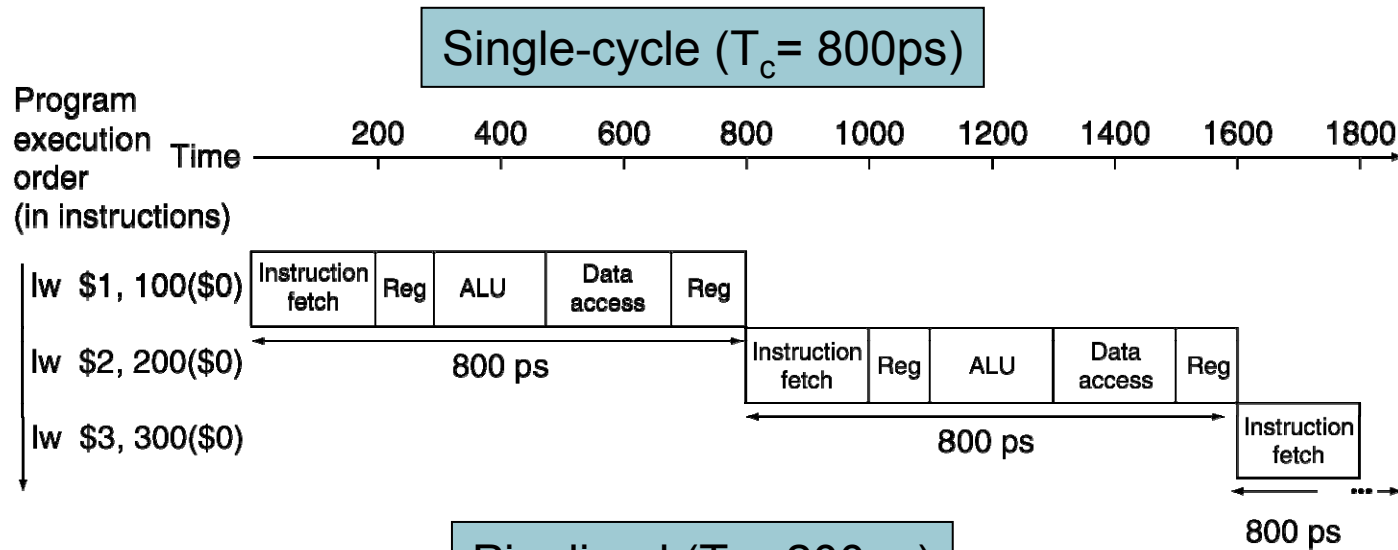
- Five stages, one step per stage
  1. **IF:** Instruction fetch from memory
  2. **ID:** Instruction decode & register read
  3. **EX:** Execute operation or calculate address
  4. **MEM:** Access memory operand
  5. **WB:** Write result back to register

# Pipeline Performance

- Assume time for stages is
  - 100ps for register read or write
  - 200ps for other stages
- Compare pipelined datapath with single-cycle datapath

Instr	Instr fetch	Register read	ALU op	Memory access	Register write	Total time
lw	200ps	100 ps	200ps	200ps	100 ps	800ps
sw	200ps	100 ps	200ps	200ps		700ps
R-format	200ps	100 ps	200ps		100 ps	600ps
beq	200ps	100 ps	200ps			500ps

# Pipeline Performance



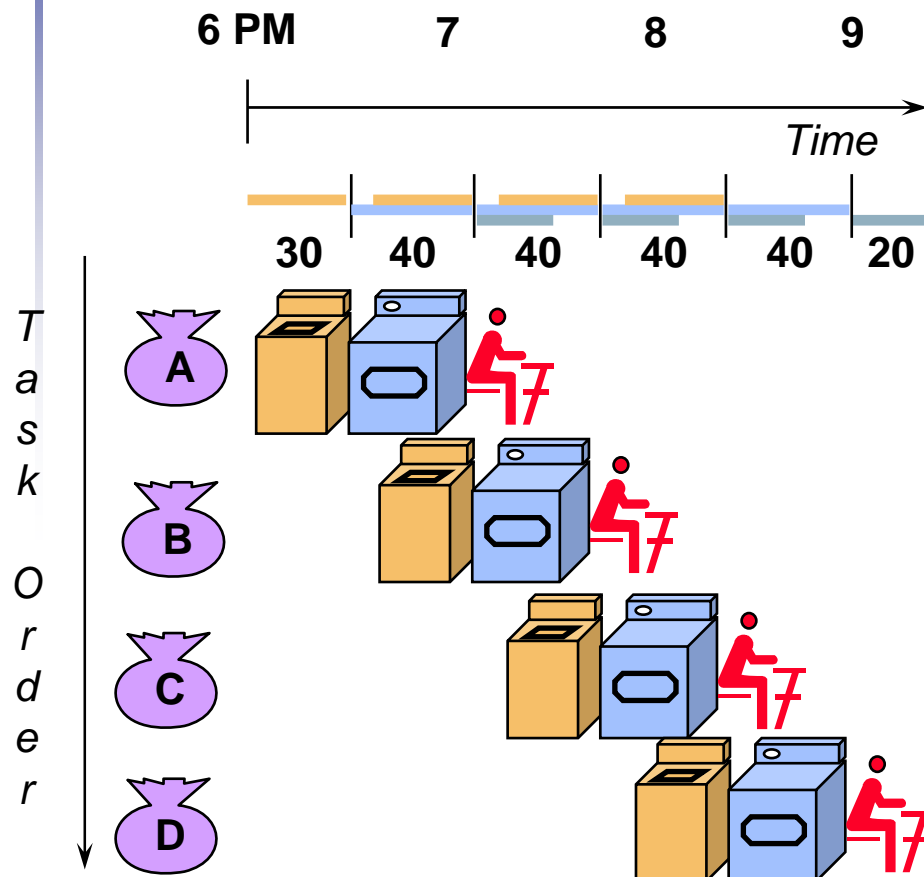


# Pipeline Speedup

- If all stages are balanced
  - i.e., all take the same time
  - Time between instructions<sub>pipelined</sub>  

$$= \frac{\text{Time between instructions}_{\text{nonpipelined}}}{\text{Number of stages}}$$
- If not balanced, speedup is less
- Speedup due to increased throughput
  - Latency (time for each instruction) does not decrease

# Pipelining Lessons



- Doesn't help **latency** of single task, but **throughput** of entire
- Pipeline rate limited by **slowest** stage
- **Multiple** tasks working at same time using different resources
- Potential speedup = **Number pipe stages**
- Unbalanced stage length; time to **"fill"** & **"drain"** the pipeline reduce speedup
- **Stall for dependences**



# Designing a Pipelined Processor

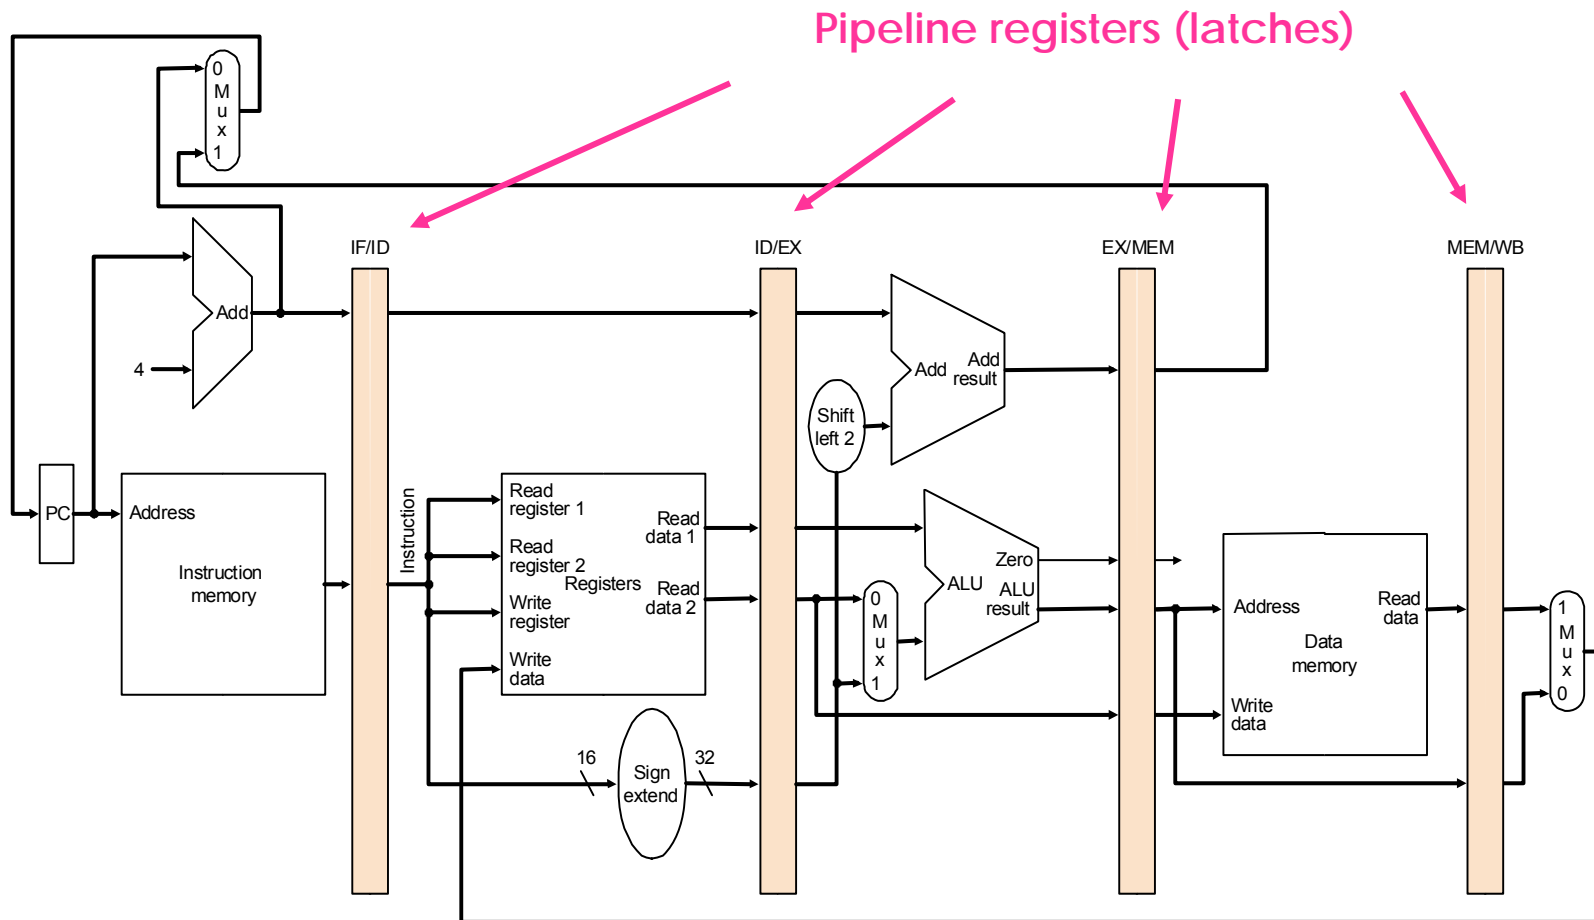
- Examine the datapath and control diagram
  - Starting with single cycle datapath
- Partition datapath into stages:
  - IF (instruction fetch), ID (instruction decode and register file read), EX (execution or address calculation), MEM (data memory access), WB (write back)
- Associate resources with stages
- Ensure that flows do not conflict, or figure out how to resolve
- Assert control in appropriate stage

# MIPS ISA Designed for Pipelining

- All instructions are 32-bits
  - Easier to fetch and decode in one cycle
  - c.f. x86: 1- to 17-byte instructions
- Few and regular instruction formats
  - Can decode and read registers in one step
- Load/store addressing
  - Can calculate address in 3<sup>rd</sup> stage, access memory in 4<sup>th</sup> stage
- Alignment of memory operands
  - Memory access takes only one cycle

# Pipelined Datapath


Use registers between stages to carry data and control



# MIPS ISA Micro-operations

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

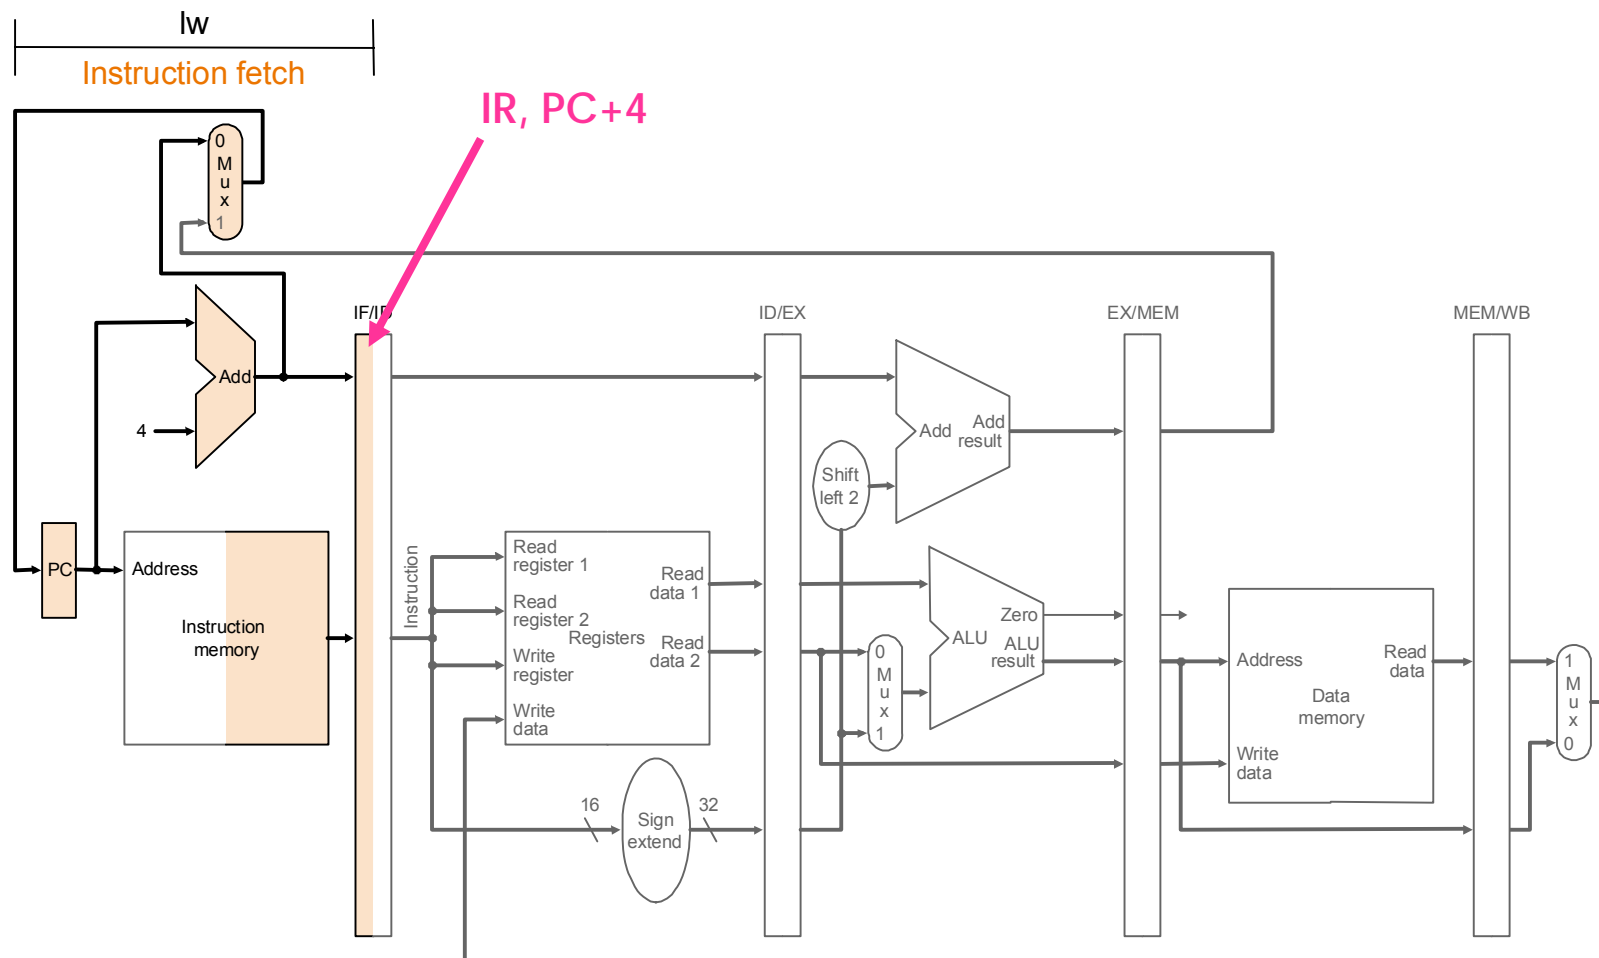
# Consider Load Instruction



- IF: Instruction Fetch
  - Fetch the instruction from the Instruction Memory
- ID: Instruction Decode
  - Registers fetch and instruction decode
- EX: Calculate the memory address
- MEM: Read the data from the Data Memory
- WB: Write the data back to the register file

# IF Stage of lw

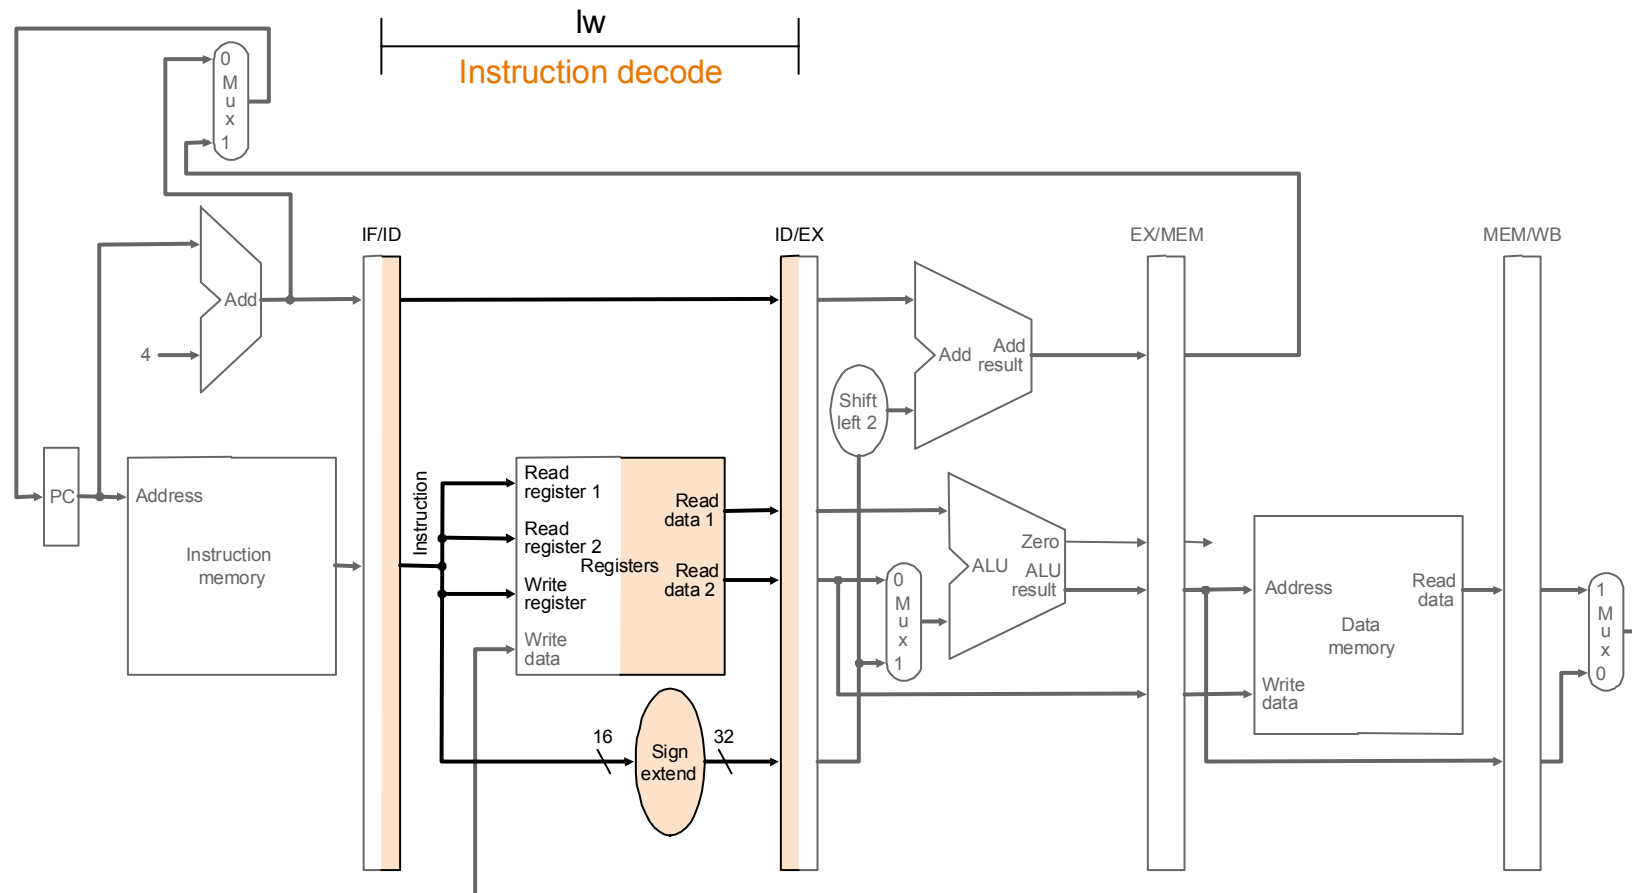
- Ex: `lw rt,rs,imm16`





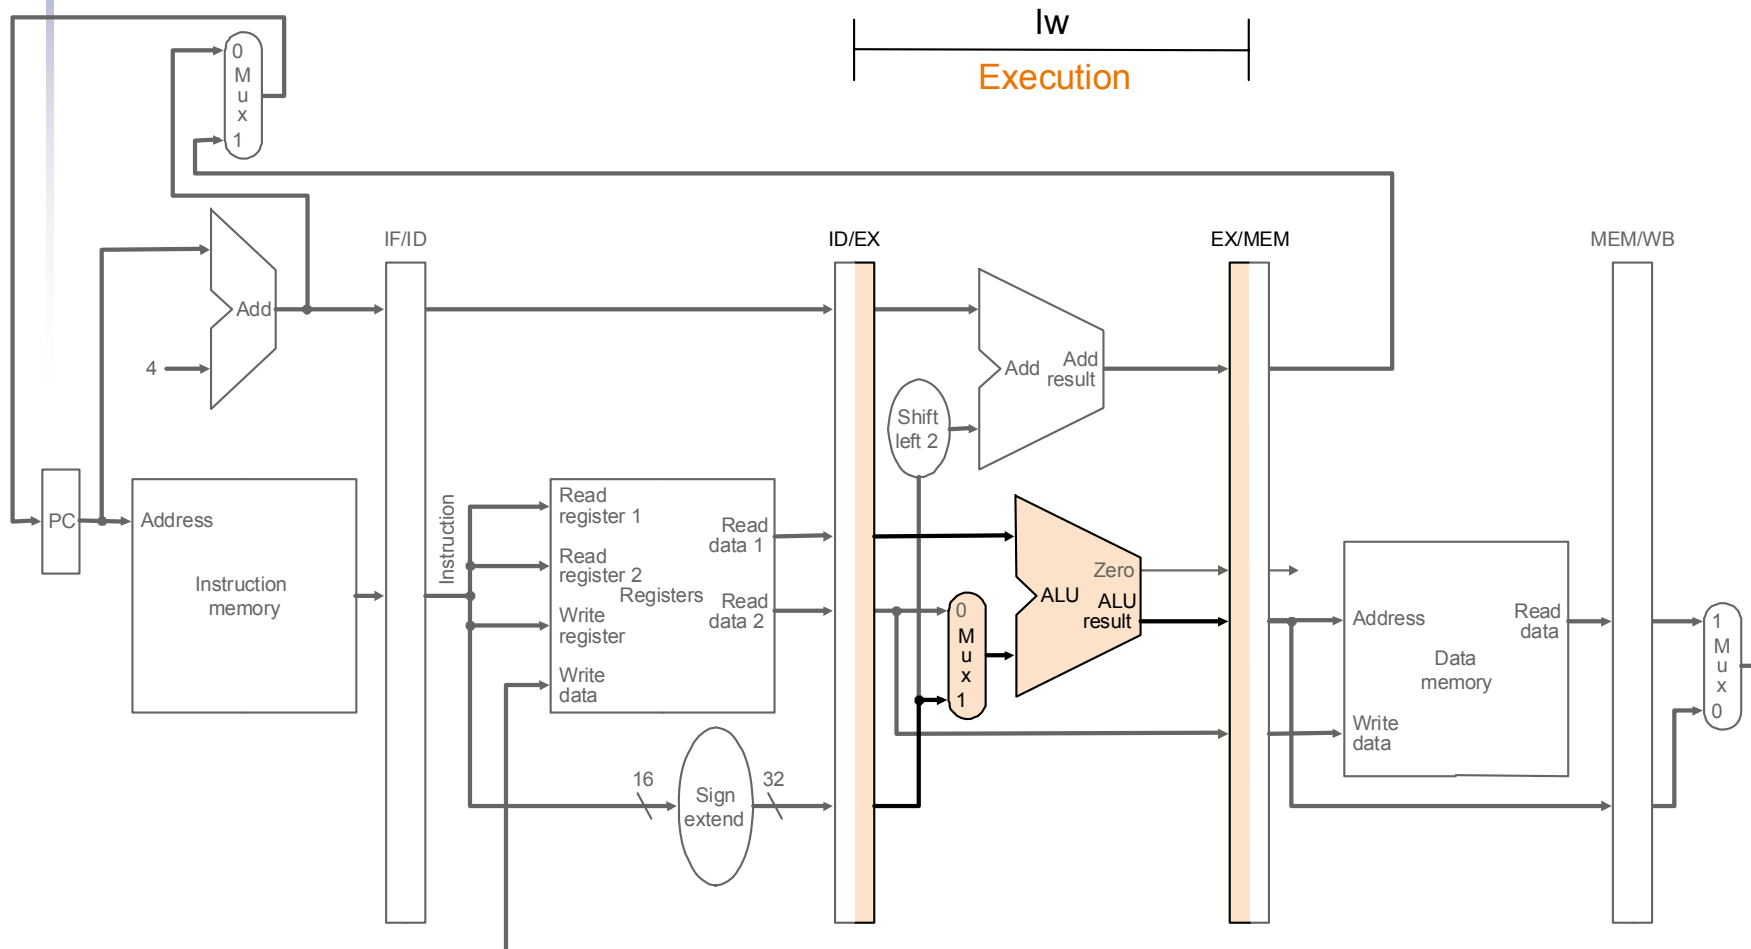
# ID Stage of lw

- Ex: `lw rt,rs,imm16`
- `A = Reg[IR[25-21]];`



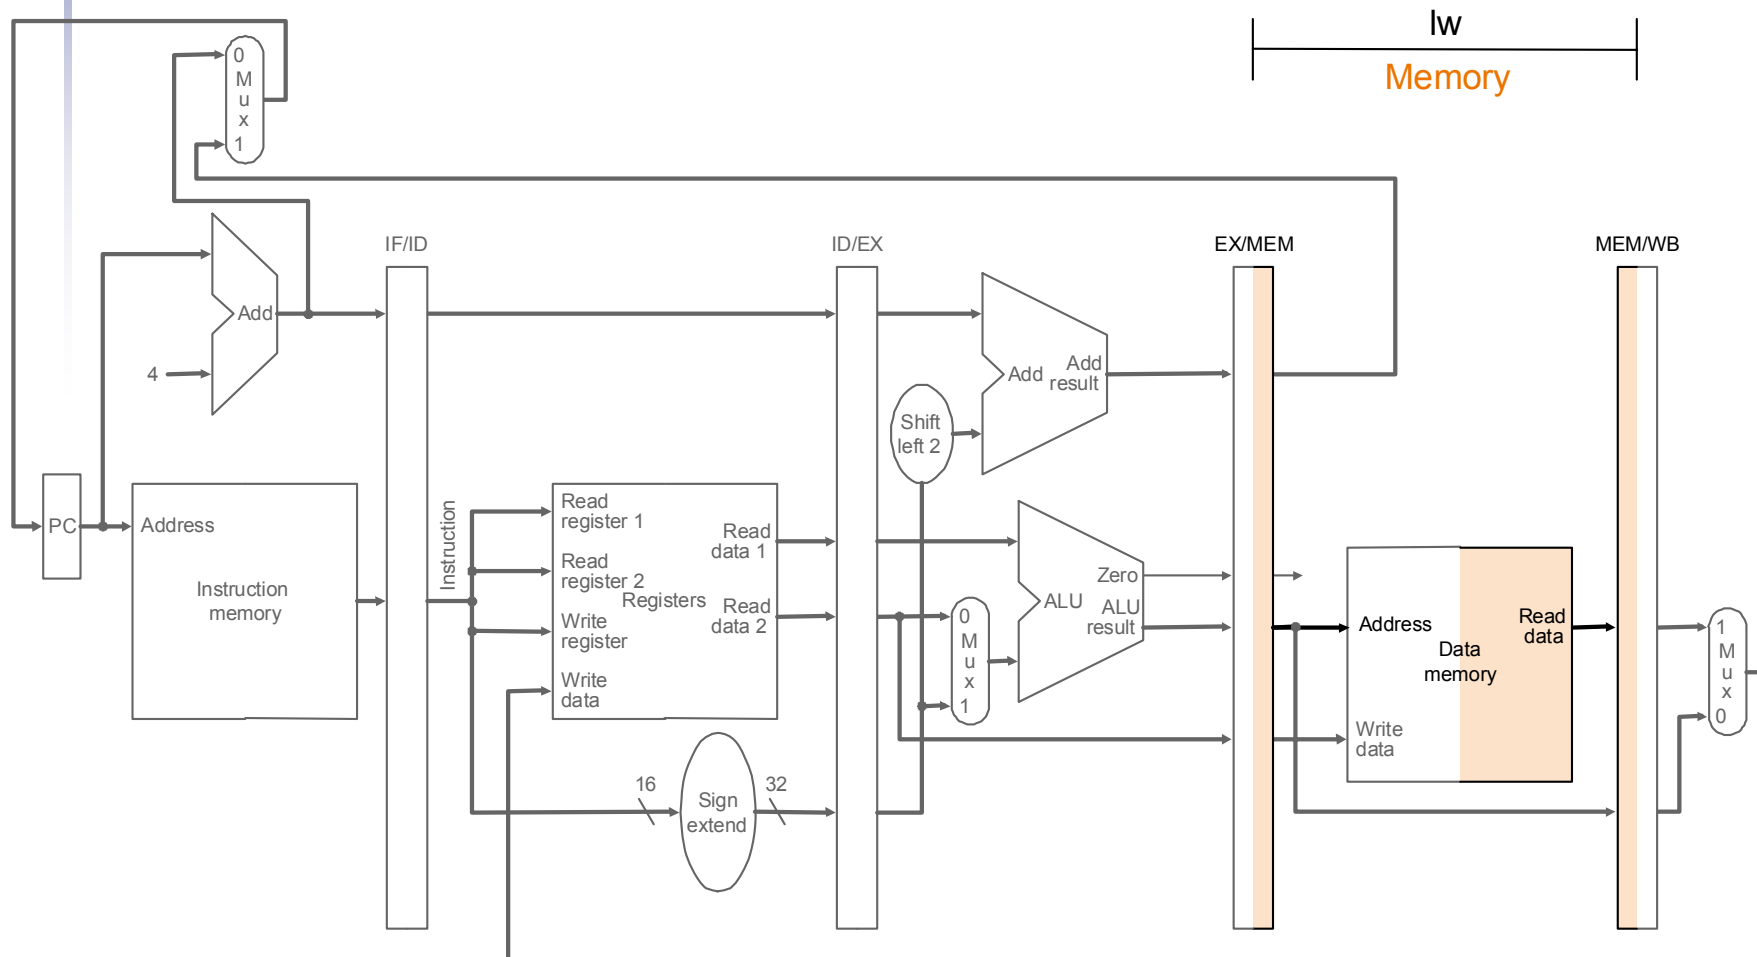
# EX Stage of lw

- Ex: `lw rt,rs,imm16`
- $ALU_{out} = A + \text{sign-ext}(IR[15-0])$



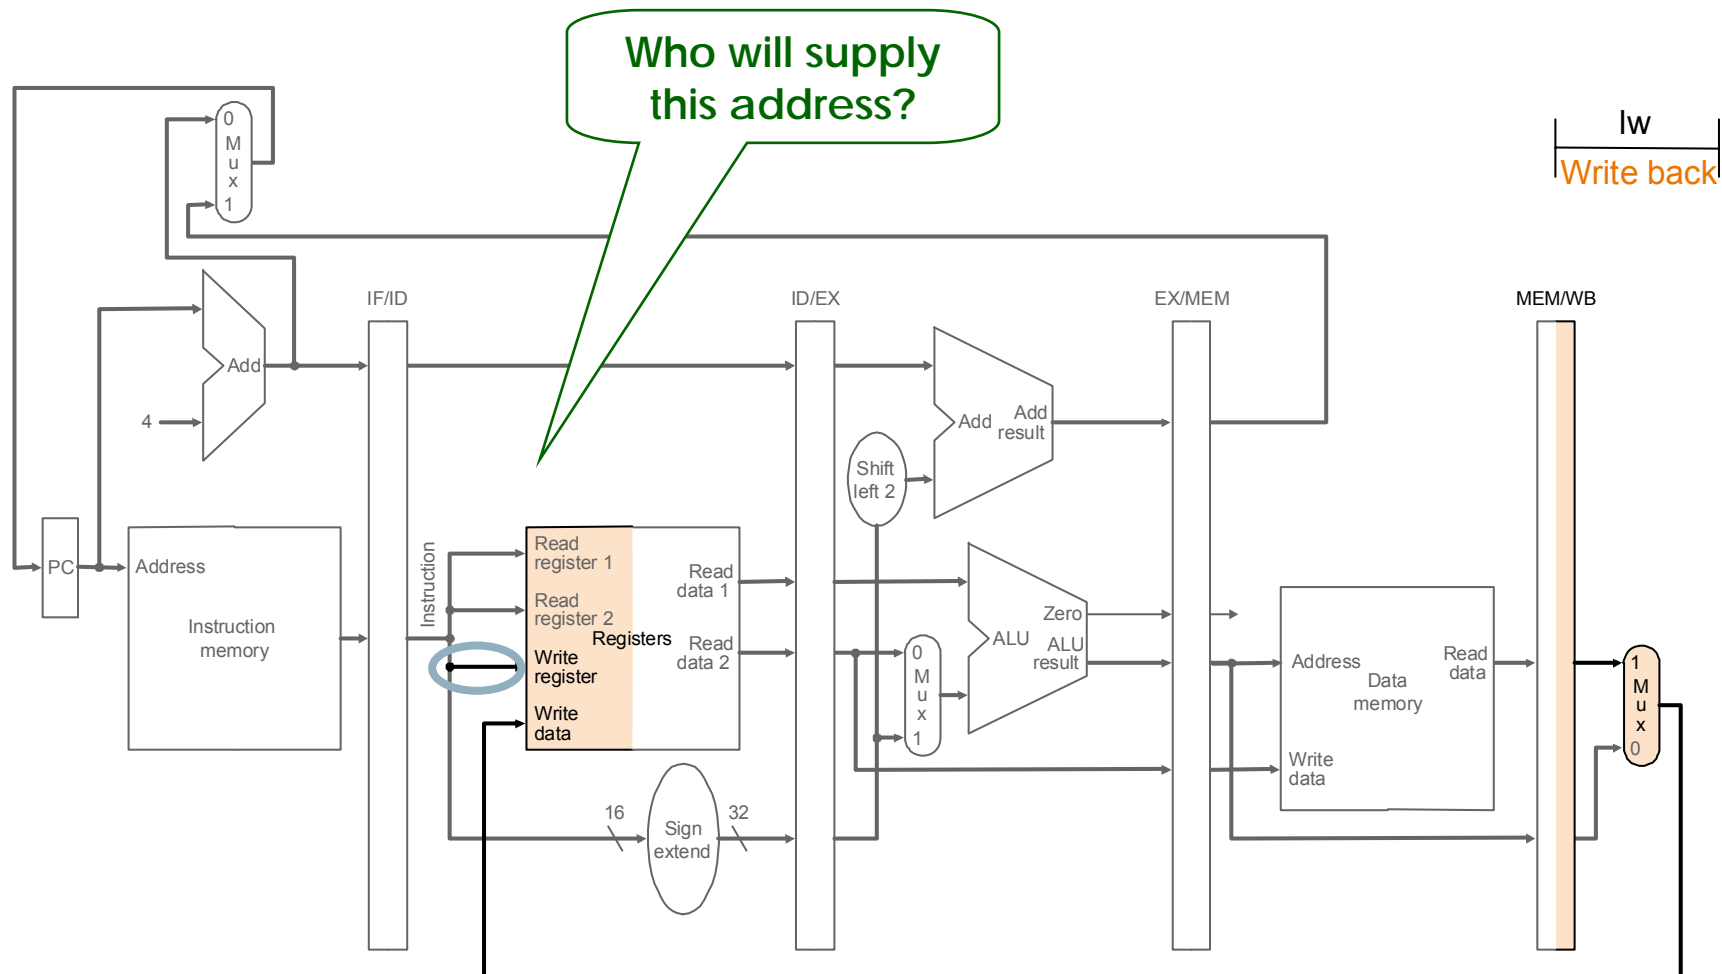
# MEM State of lw

- Ex: `lw rt,rs,imm16`
- **MDR = mem[ALUout]**

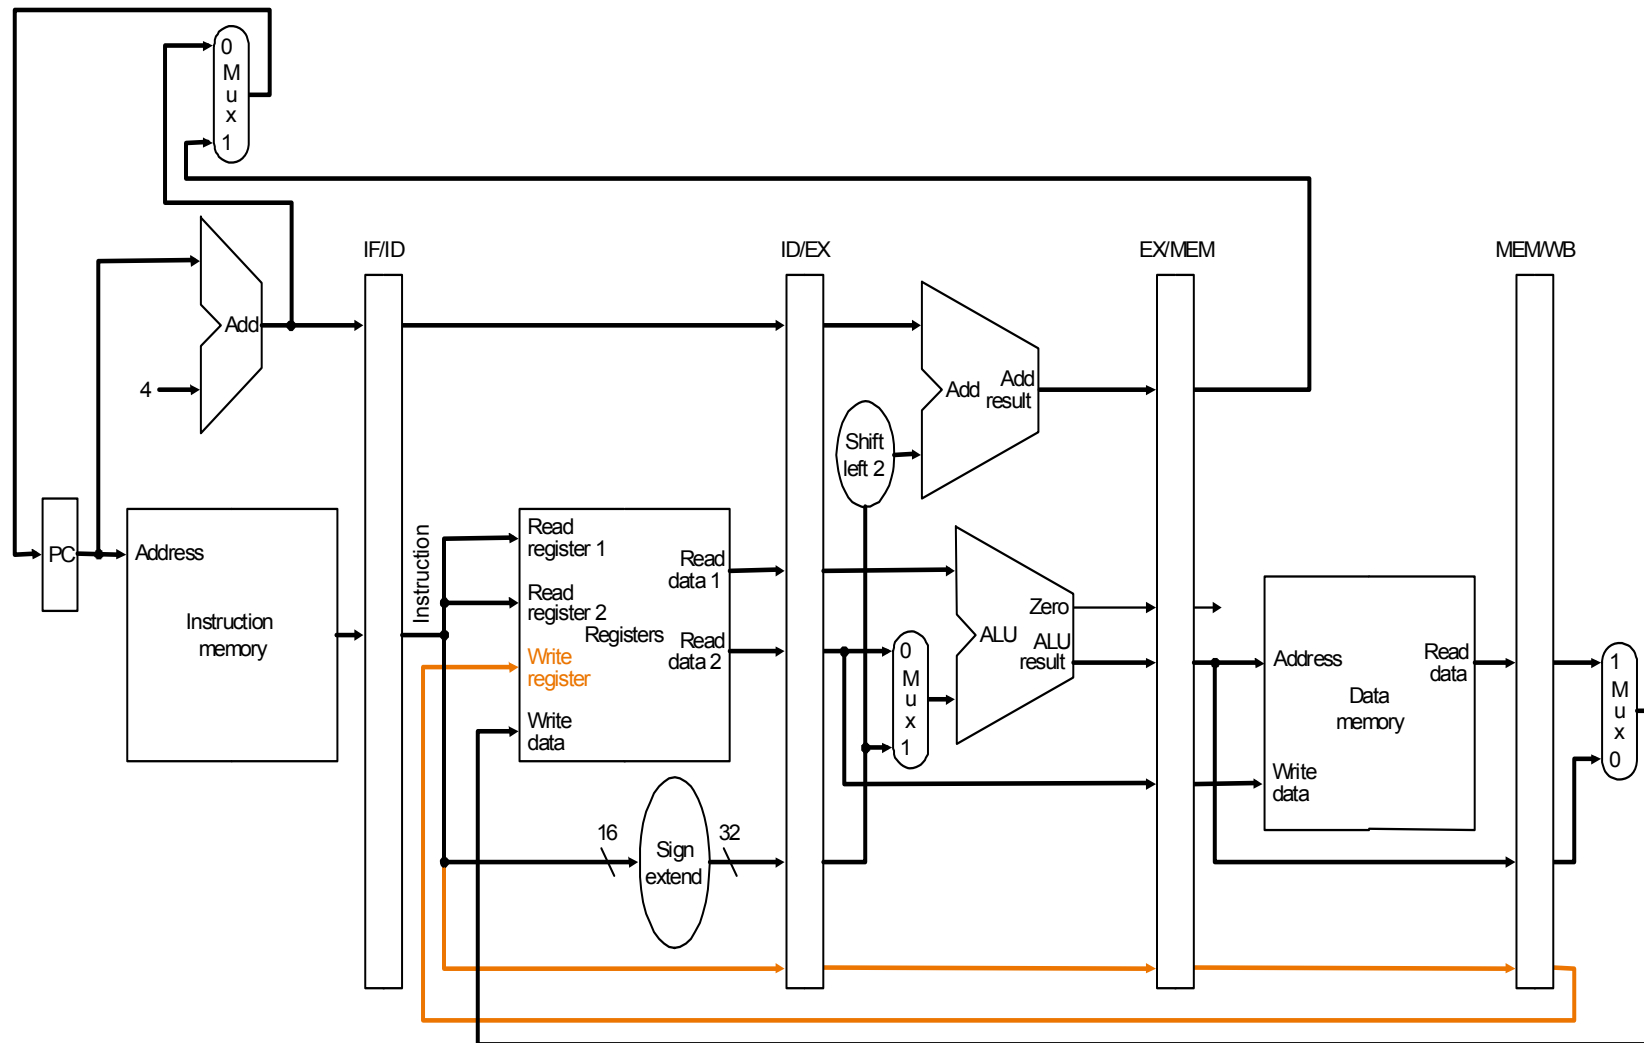


# WB Stage of lw

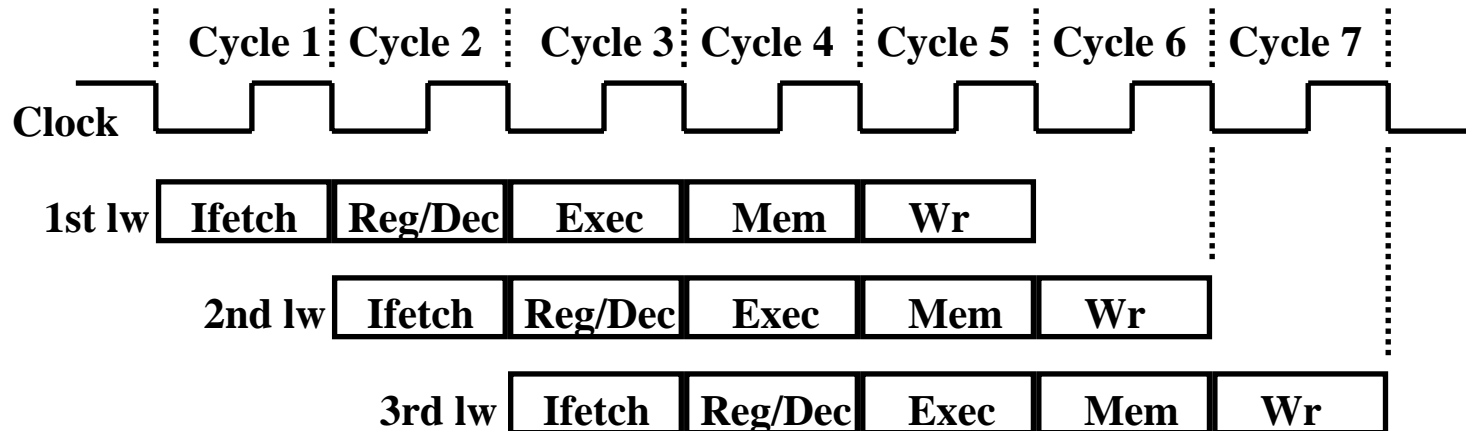
- Ex: `lw rt,rs,imm16`
- $\text{Reg}[\text{IR}[20-16]] = \text{MDR}$



# Pipelined Datapath




# Pipelining 1w Instructions



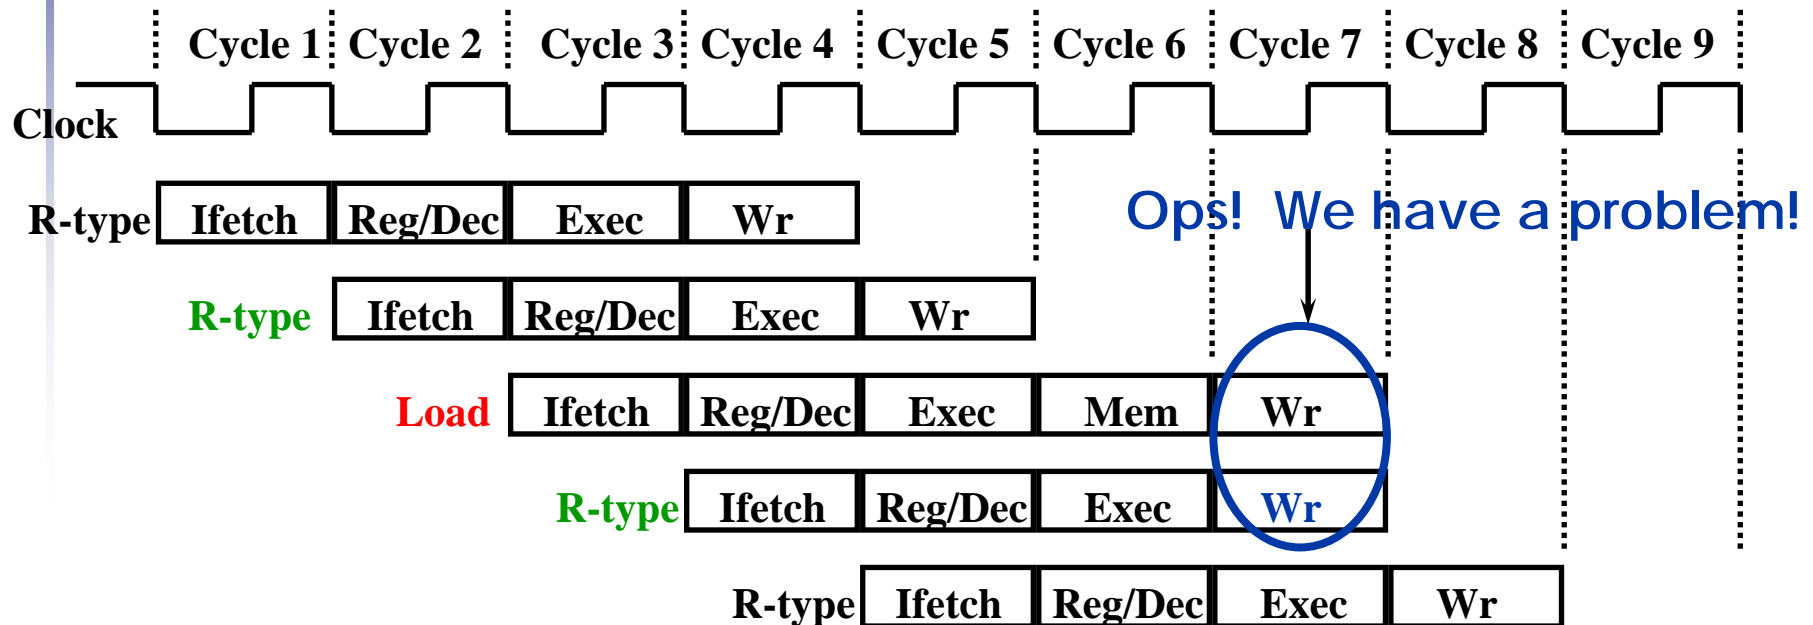
- 5 functional units in the pipeline datapath are:
  - **Instruction Memory** for the Ifetch stage
  - Register File's **Read ports** (busA and busB) for the Reg/Dec stage
  - **ALU** for the Exec stage
  - **Data Memory** for the MEM stage
  - Register File's **Write port** (busW) for the WB stage

# The Four Stages of R-type Instruction



- IF: fetch the instruction from the Instruction Memory
- ID: registers fetch and instruction decode
- EX: ALU operates on the two register operands
- WB: write ALU output back to the register file

# Hazard Problem



- We have a *structural hazard*:
  - Two instructions try to write to the register file at the same time!
  - Only one write port



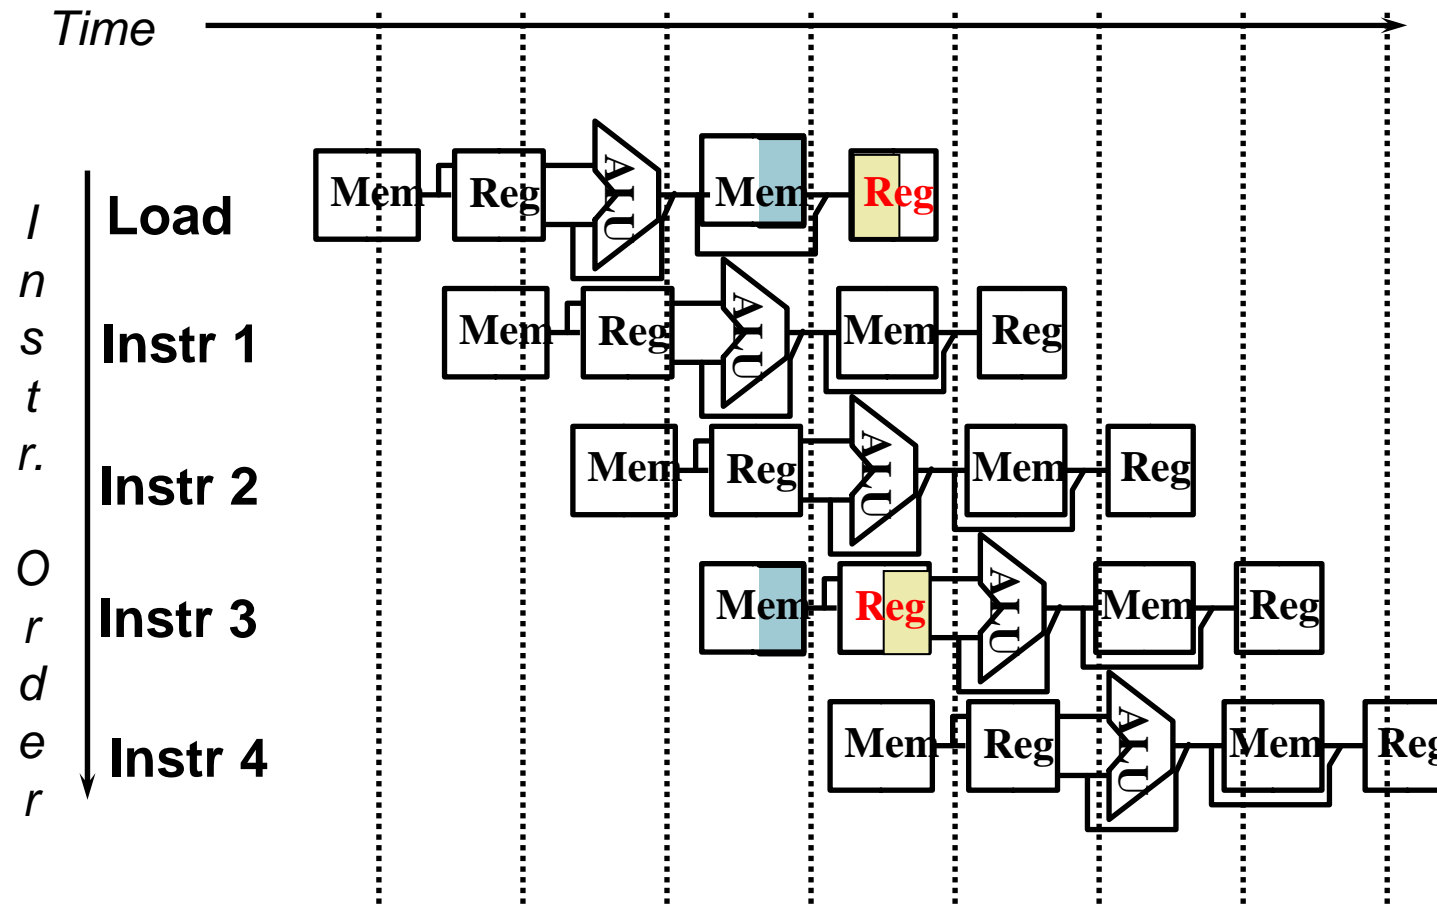
# Pipeline Hazards

- Situations that prevent starting the next instruction in the next cycle
- Structure hazard
  - A required resource is busy
- Data hazard
  - Need to wait for previous instruction to complete its data read/write
- Control hazard
  - Deciding on control action depends on previous instruction
- *Several ways to solve: forwarding, adding pipeline bubble, making instructions same length*

# Structure Hazards

- Conflict for use of a resource
- In MIPS pipeline with a single memory
  - Load/store requires data access
  - Instruction fetch would have to *stall* for that cycle
    - Would cause a pipeline “bubble”
- Hence, pipelined datapaths require separate instruction/data memories
  - Or separate instruction/data caches

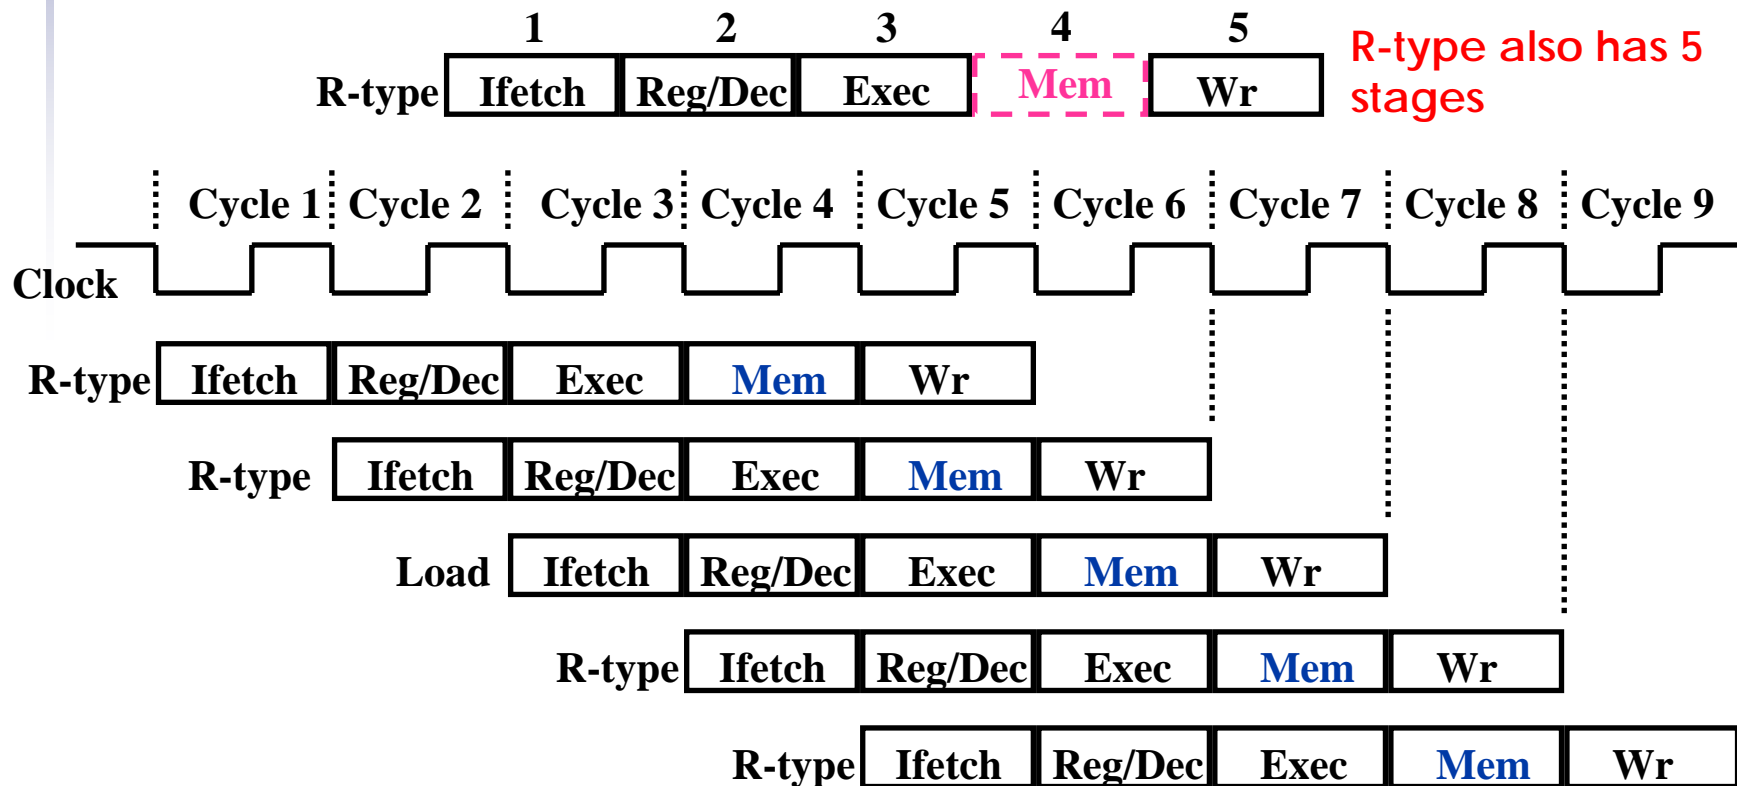
# Structural Hazard Solution: Seperate I/D Memory




1. Use 2 memory: data memory and instruction memory
2. First half cycle for write and the second half cycle for read

# Structural Hazard Solution: Delay R-type's Write

- Delay R-type's register write by one cycle:
  - R-type also use Reg File's write port at Stage 5
  - MEM is a NOP stage: nothing is being done.



# The Four Stages of sw




- IF: fetch the instruction from the Instruction Memory
- ID: registers fetch and instruction decode
- EX: calculate the memory address
- MEM: write the data into the Data Memory

Add an extra stage:

- WB: NOP

# The Three Stages of beq



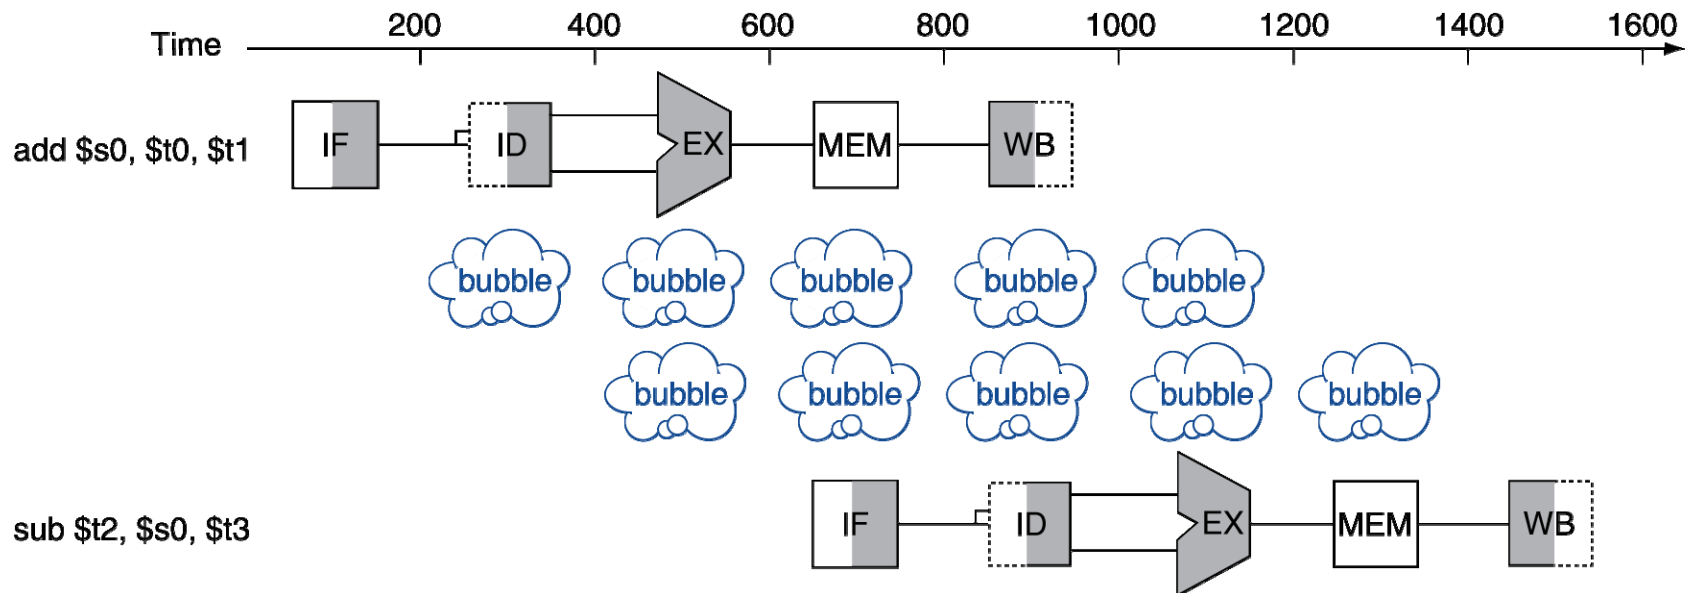
- IF: fetch the instruction from the Instruction Memory
- ID: registers fetch and instruction decode
- EX:
  - compares the two register operand
  - select correct branch target address
  - latch into PC

Add two extra stages:

- MEM: NOP
- WB: NOP

# Data Hazards

- An instruction depends on completion of data access by a previous instruction
  - add **\$s0**, \$t0, \$t1
  - sub \$t2, **\$s0**, \$t3



# Types of Data Hazards

Three types: (inst. i1 followed by inst. i2)

- **RAW (read after write):**

i2 tries to read operand before i1 writes it

- **WAR (write after read):**

i2 tries to write operand before i1 reads it

- Gets wrong operand, e.g., autoincrement addr.
- Can't happen in MIPS 5-stage pipeline because:
  - All instructions take 5 stages, and reads are always in stage 2, and writes are always in stage 5

- **WAW (write after write):**

i2 tries to write operand before i1 writes it

- Leaves wrong result ( i1's not i2's); occur only in pipelines that write in more than one stage
- Can't happen in MIPS 5-stage pipeline because:
  - All instructions take 5 stages, and writes are always in stage 5

- **RAR?**

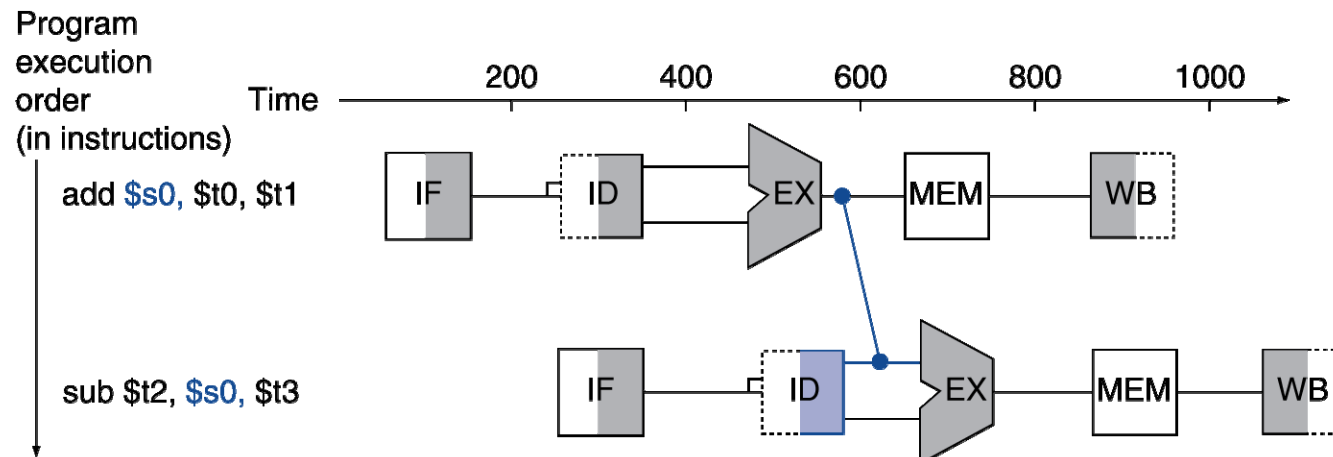


# Handling Data Hazards

- Use simple, fixed designs
  - Eliminate WAR by always fetching operands early (ID) in pipeline
  - Eliminate WAW by doing all write backs in order (last stage, static)
  - These features have a lot to do with ISA design
- **Internal forwarding** in register file:
  - Write in first half of clock and read in second half
  - Read delivers what is written, resolve hazard between sub and add
- Detect and resolve remaining ones
  - Compiler inserts NOP
  - Forward
  - Stall

# Forwarding (aka Bypassing)

- Use result when it is computed
  - Don't wait for it to be stored in a register
  - Requires extra connections in the datapath



# Example

- Consider the following code sequence

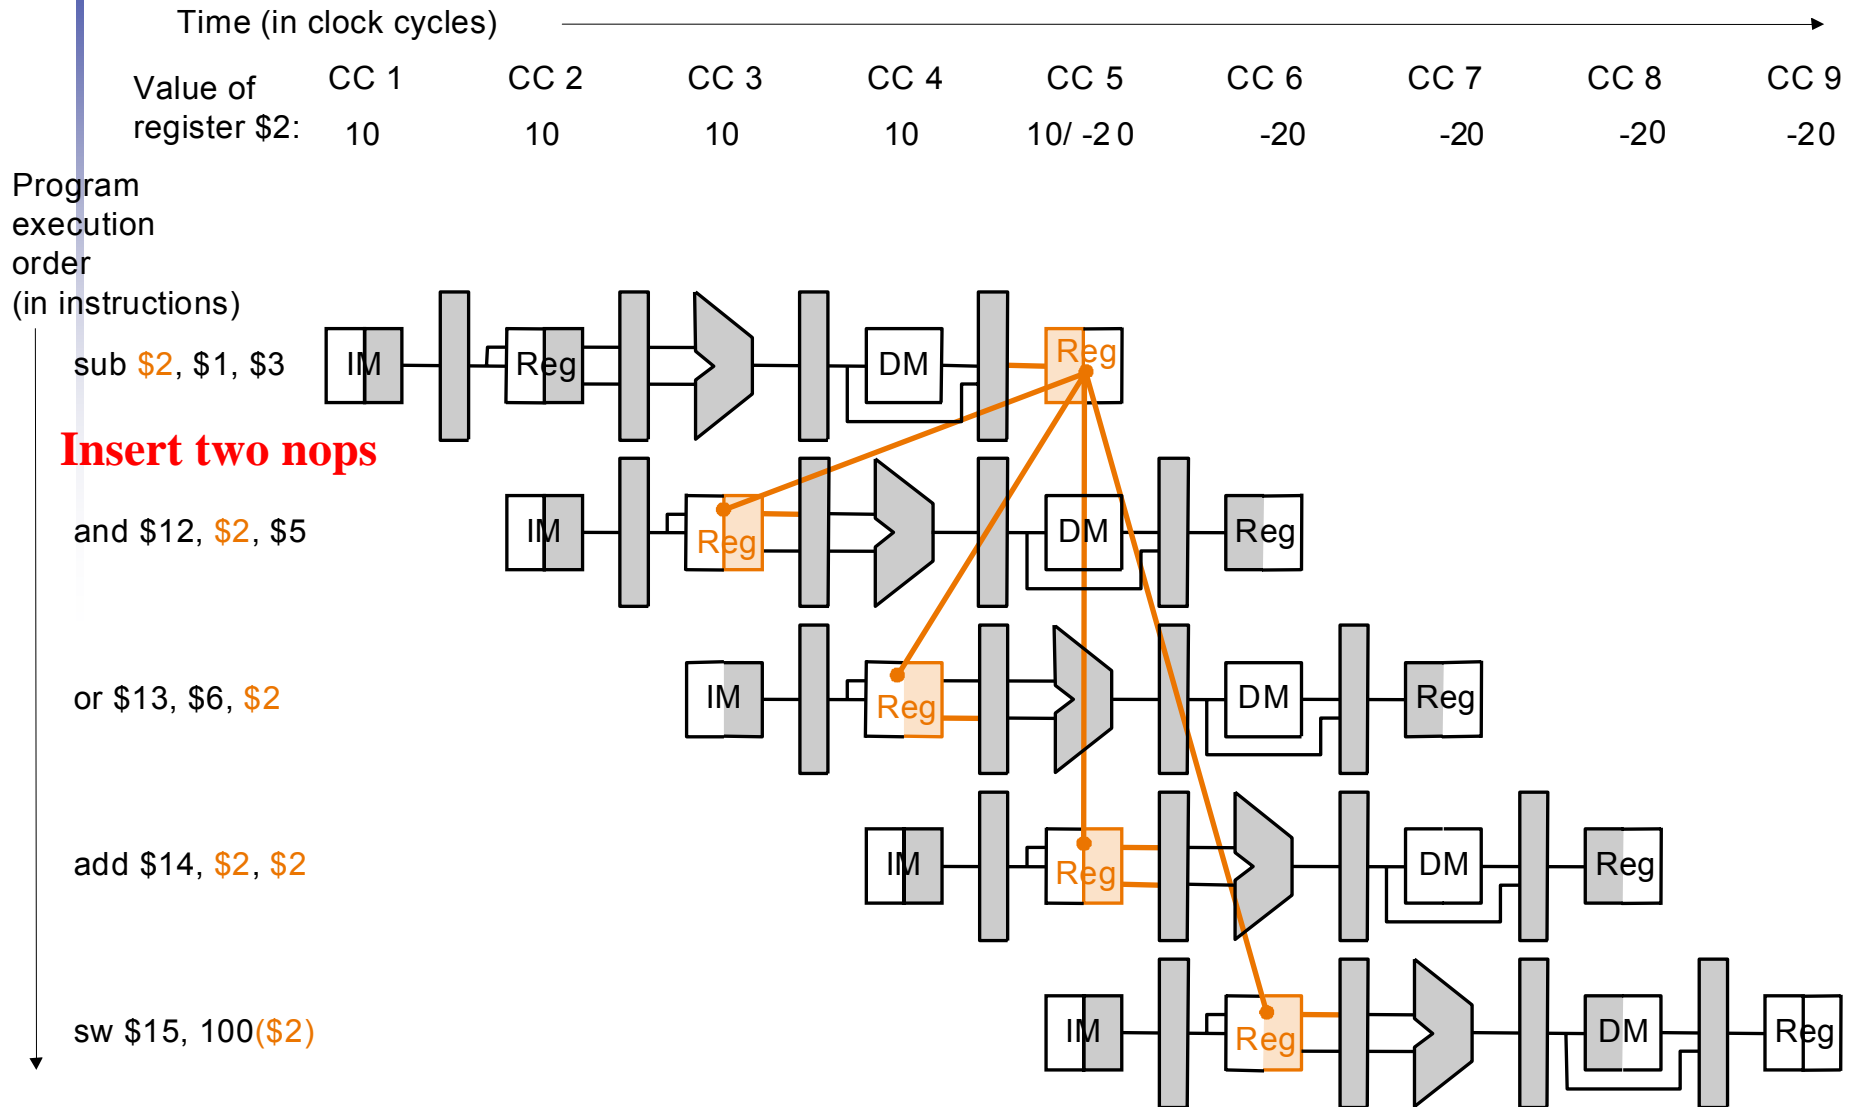


```

sub $2, $1, $3
and $12, $2, $5
or $13, $6, $2
add $14, $2, $2
sw $15, 100($2)

```

# Data Hazards Solution: Inserting NOPs by Software



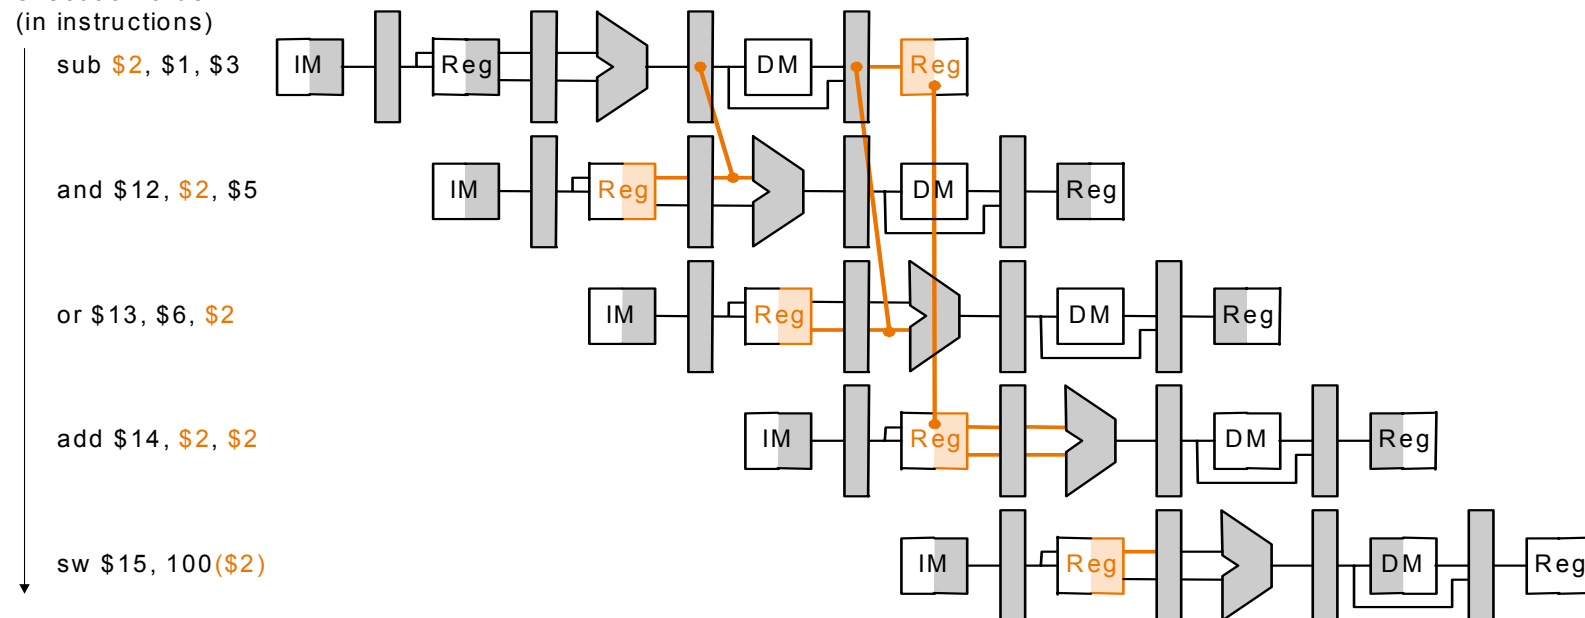
# Data Hazards Solution: Internal Forwarding Logic

- Use temporary results, e.g., those in pipeline registers, don't wait for them to be written

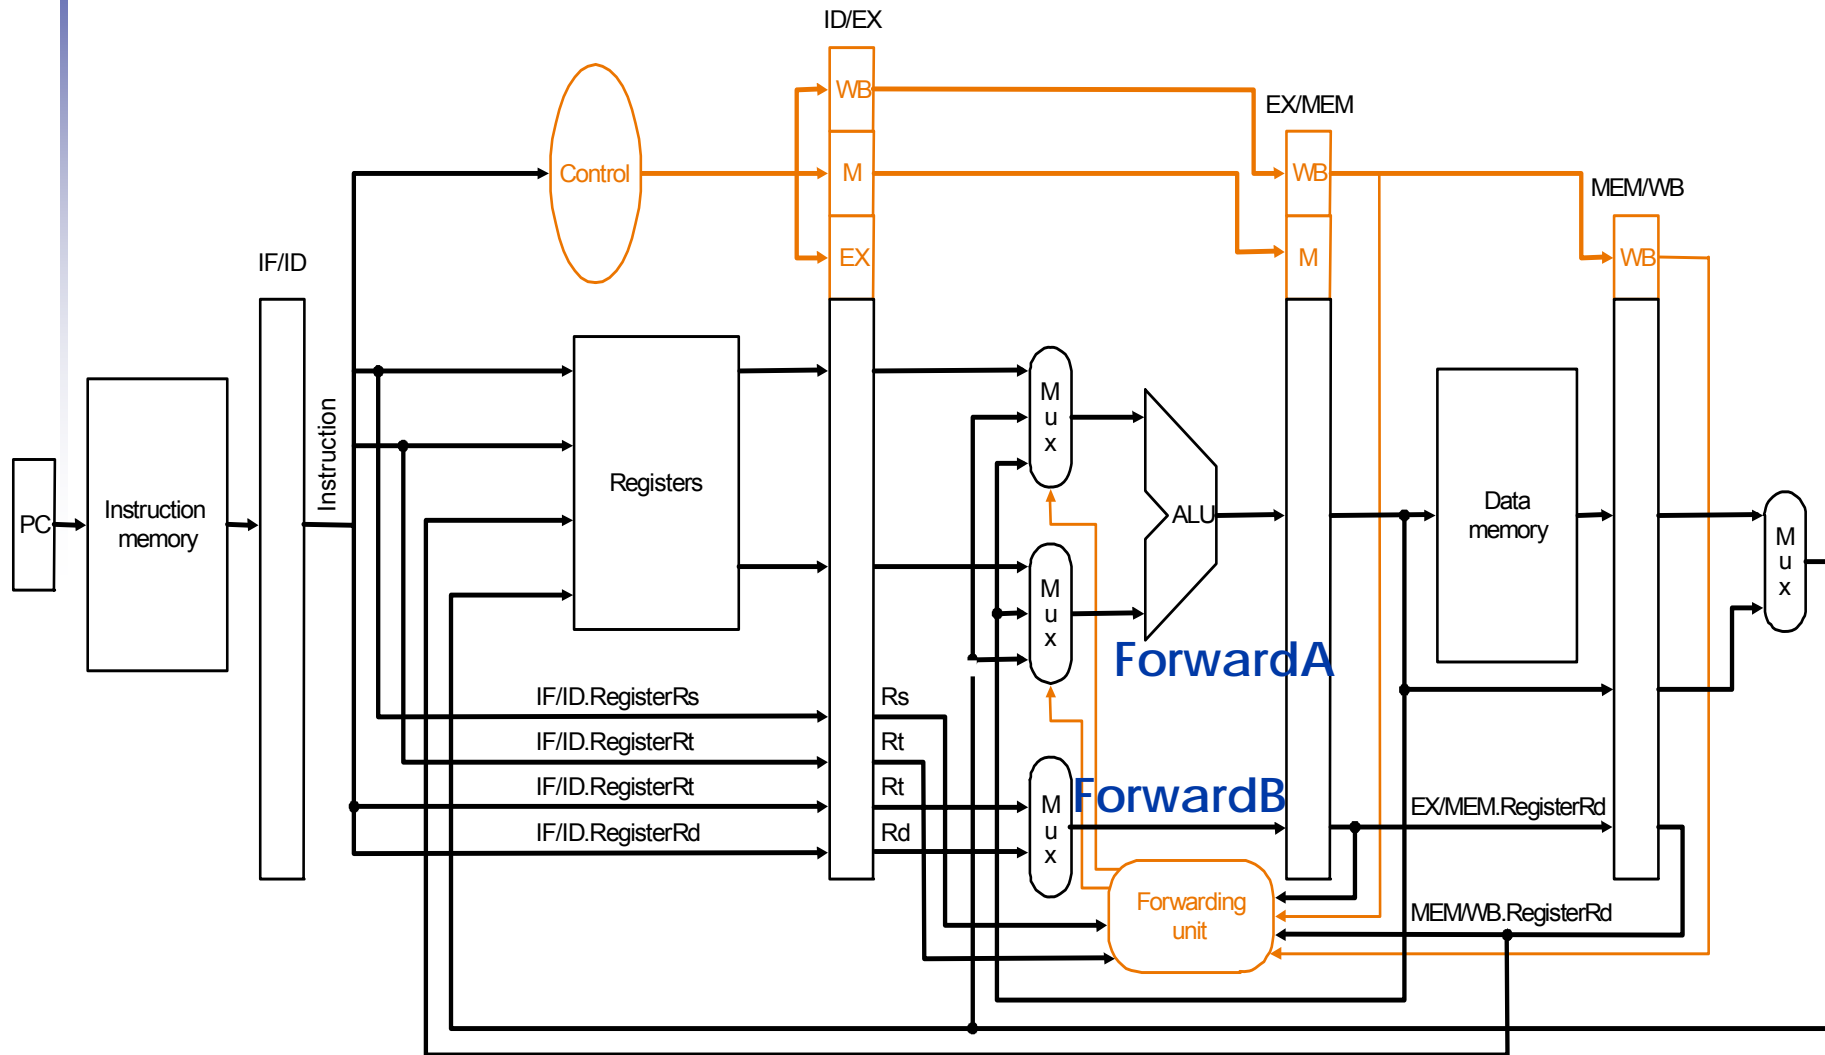
Time (in clock cycles) →

	CC 1	CC 2	CC 3	CC 4	CC 5	CC 6	CC 7	CC 8	CC 9
Value of register \$2 :	10	10	10	10	10/-20	-20	-20	-20	-20
Value of EX/MEM :	X	X	X	-20	X	X	X	X	X
Value of MEM/WB :	X	X	X	X	-20	X	X	X	X

Program  
execution order  
(in instructions)

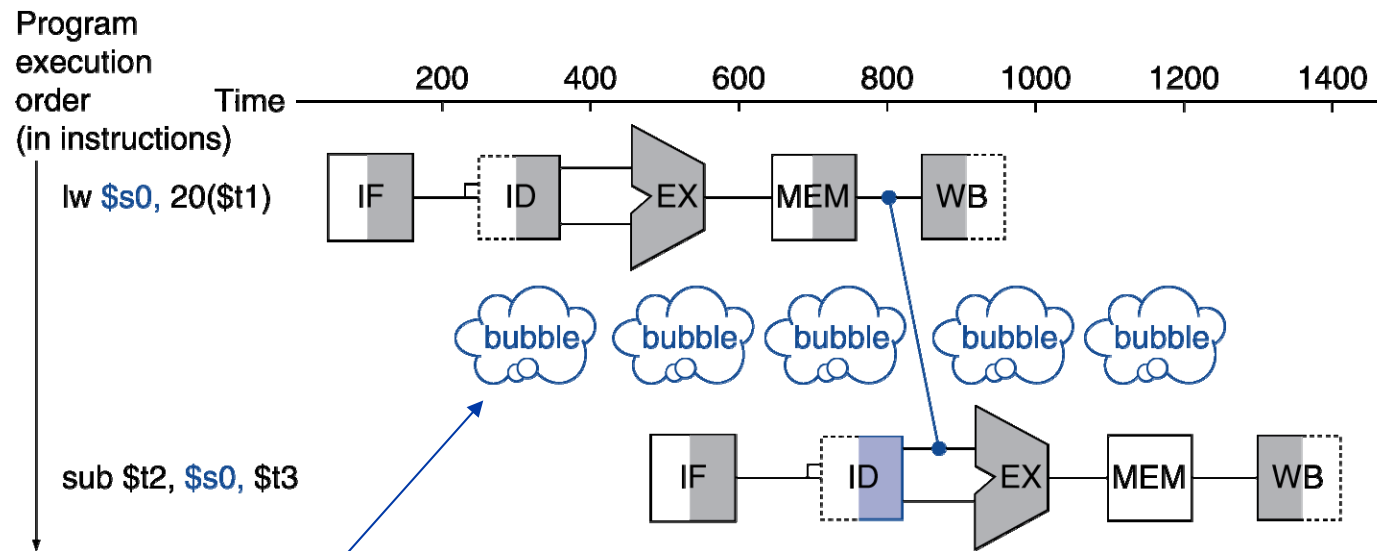


# Pipeline with Forwarding



# Load-Use Data Hazard

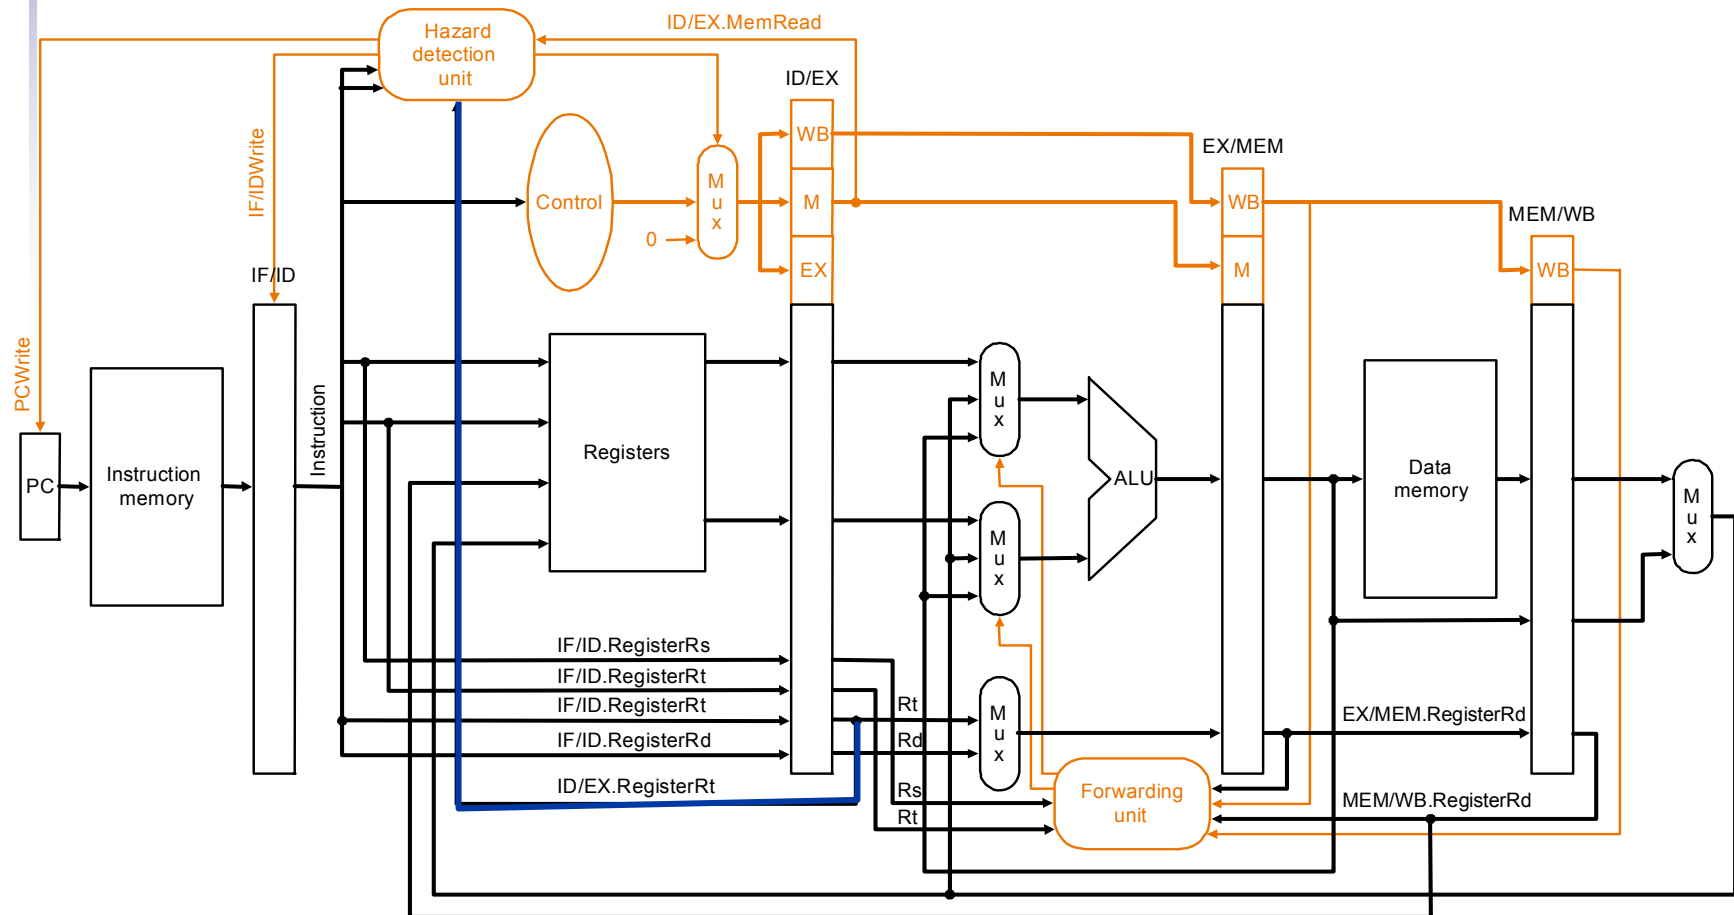
- **Can't always avoid stalls by forwarding**
  - If value not computed when needed
  - Can't forward backward in time!



Software Check or Hardware Handling

# Pipeline with Stalling Unit

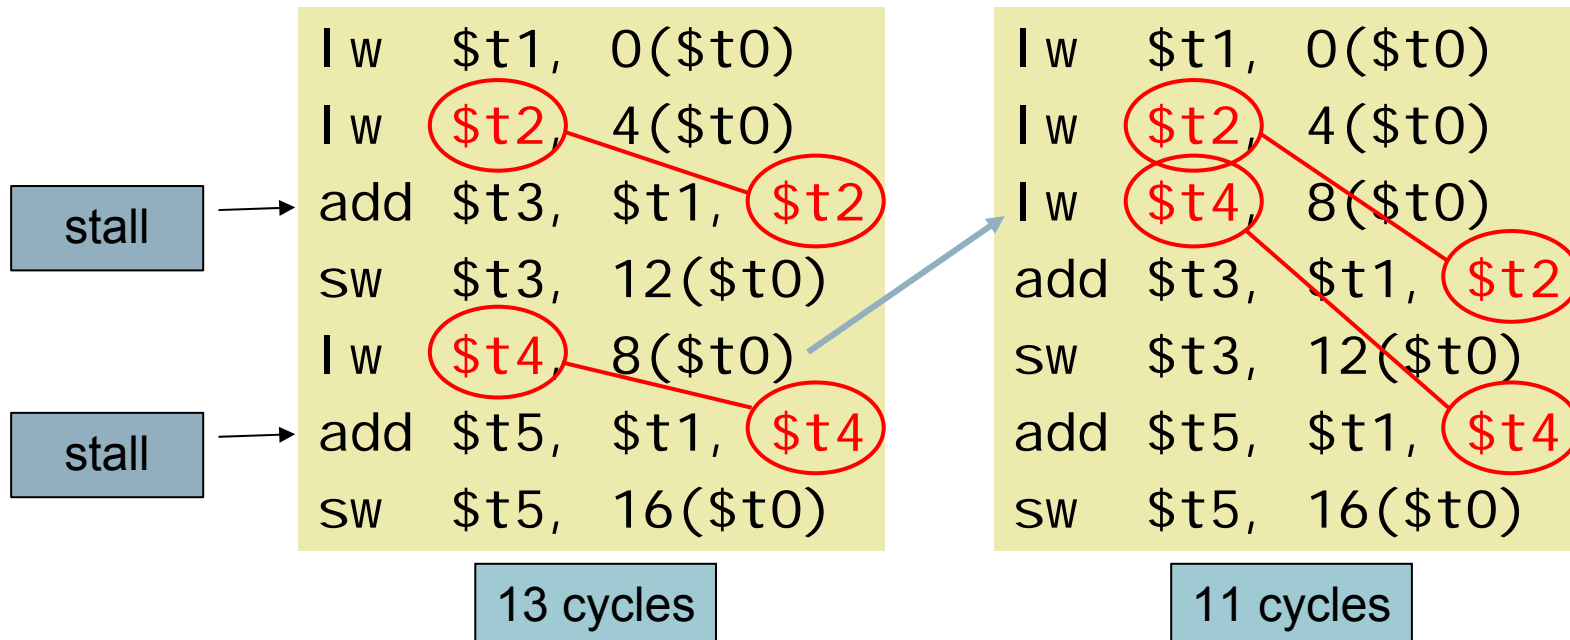
- Forwarding controls ALU inputs, hazard detection controls PC, IF/ID, control signals





# Code Scheduling to Avoid Stalls

- Reorder code to avoid use of load result in the next instruction
- C code for  $A = B + E; C = B + F;$

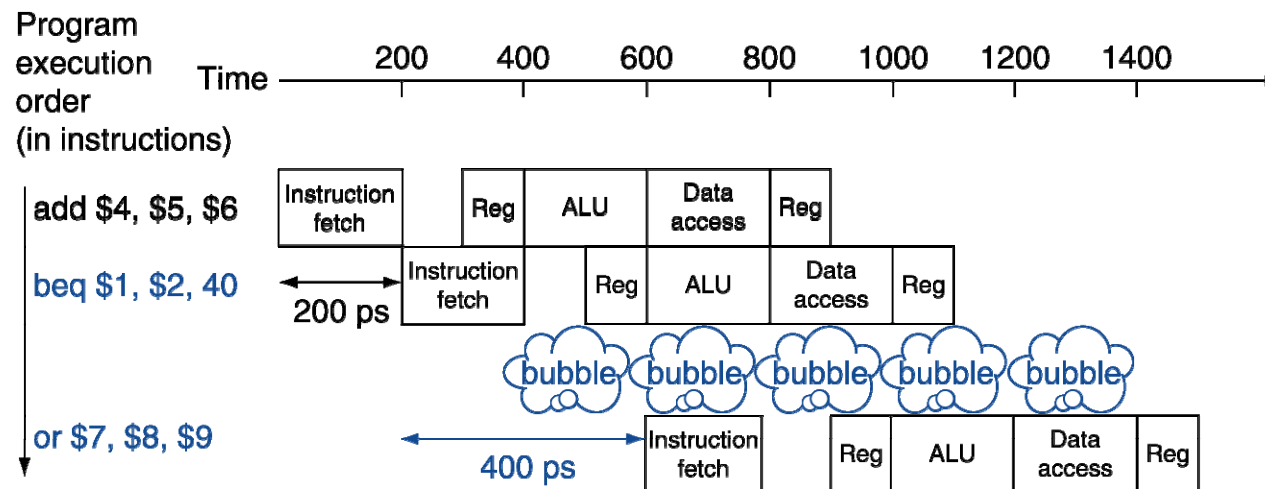


# Control Hazards

- Branch determines flow of control
  - Fetching next instruction depends on branch outcome
  - Pipeline can't always fetch correct instruction
    - Still working on ID stage of branch
- In MIPS pipeline
  - Need to compare registers and compute target early in the pipeline
  - Add hardware to do it in ID stage

# Stall on Branch

- Wait until branch outcome determined before fetching next instruction



# Handling Branch Hazard

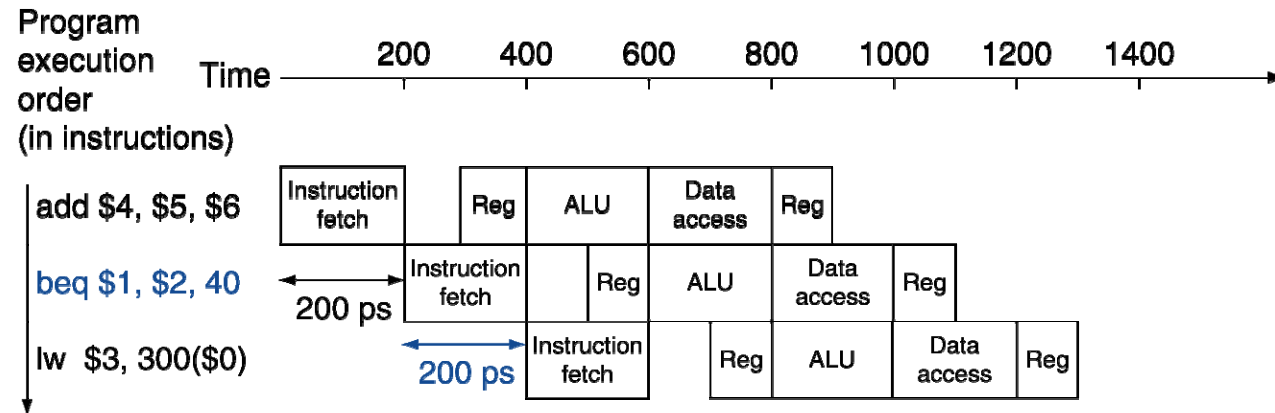
- Reduce delay of taken branch by moving branch execution earlier in the pipeline
  - Move up branch address calculation to ID
  - Check branch equality at ID (using XOR) by comparing the two registers read during ID
  - Branch decision made at ID => one instruction to flush
  - Add a control signal, IF.Flush, to zero instruction field of IF/ID => making the instruction an NOP
- (Static) Predict branch always not taken
  - Need to add hardware for flushing inst. if wrong
  - Branch decision made at MEM => need to flush instruction in IF/ID, ID/EX by changing control values to 0
- Dynamic branch prediction
- Compiler rescheduling, delay branch

# Branch Prediction

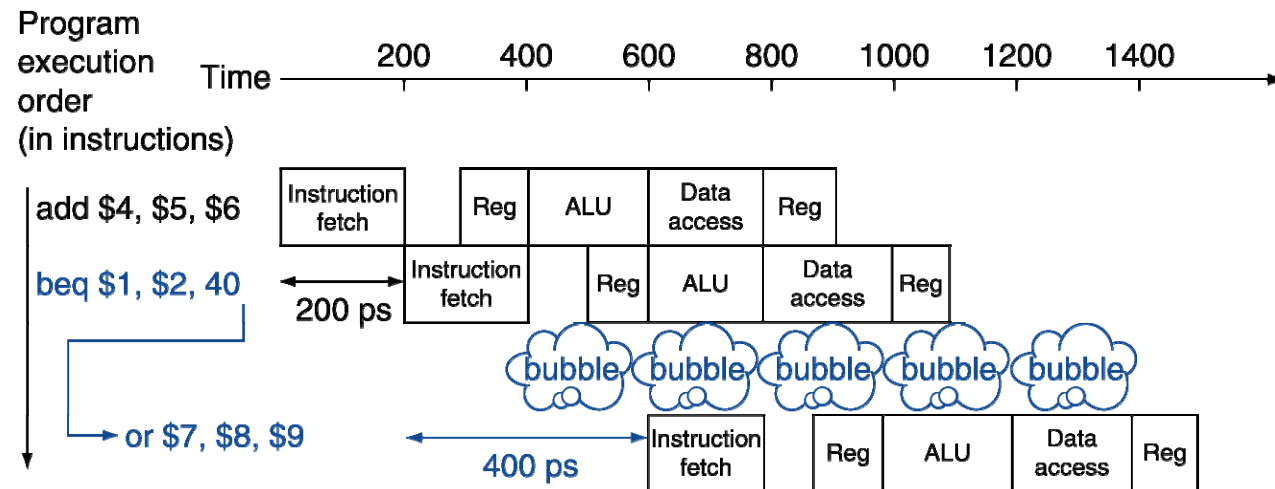
- Longer pipelines can't readily determine branch outcome early
  - Stall penalty becomes unacceptable
- Predict outcome of branch
  - Only stall if prediction is wrong
- In MIPS pipeline
  - Can predict branches not taken
  - Fetch instruction after branch, with no delay

# MIPS with Predict Not Taken

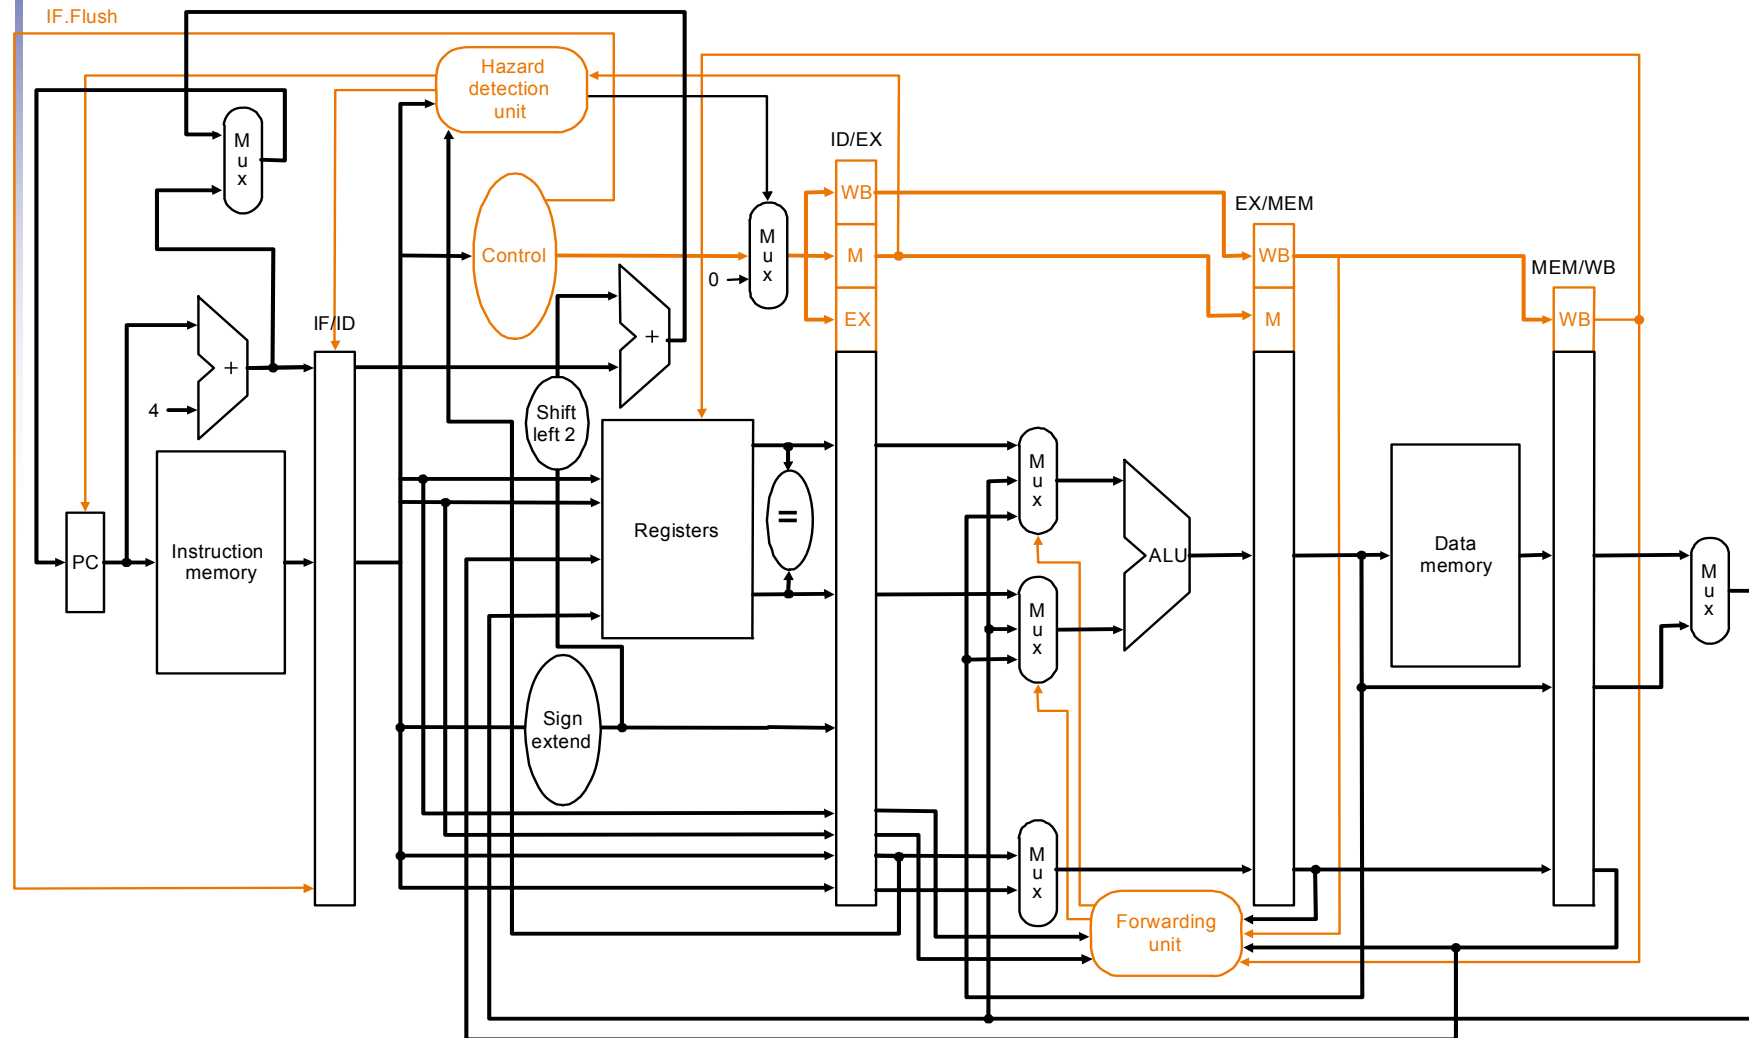
Prediction  
correct



Prediction  
incorrect



# Pipeline with Flushing



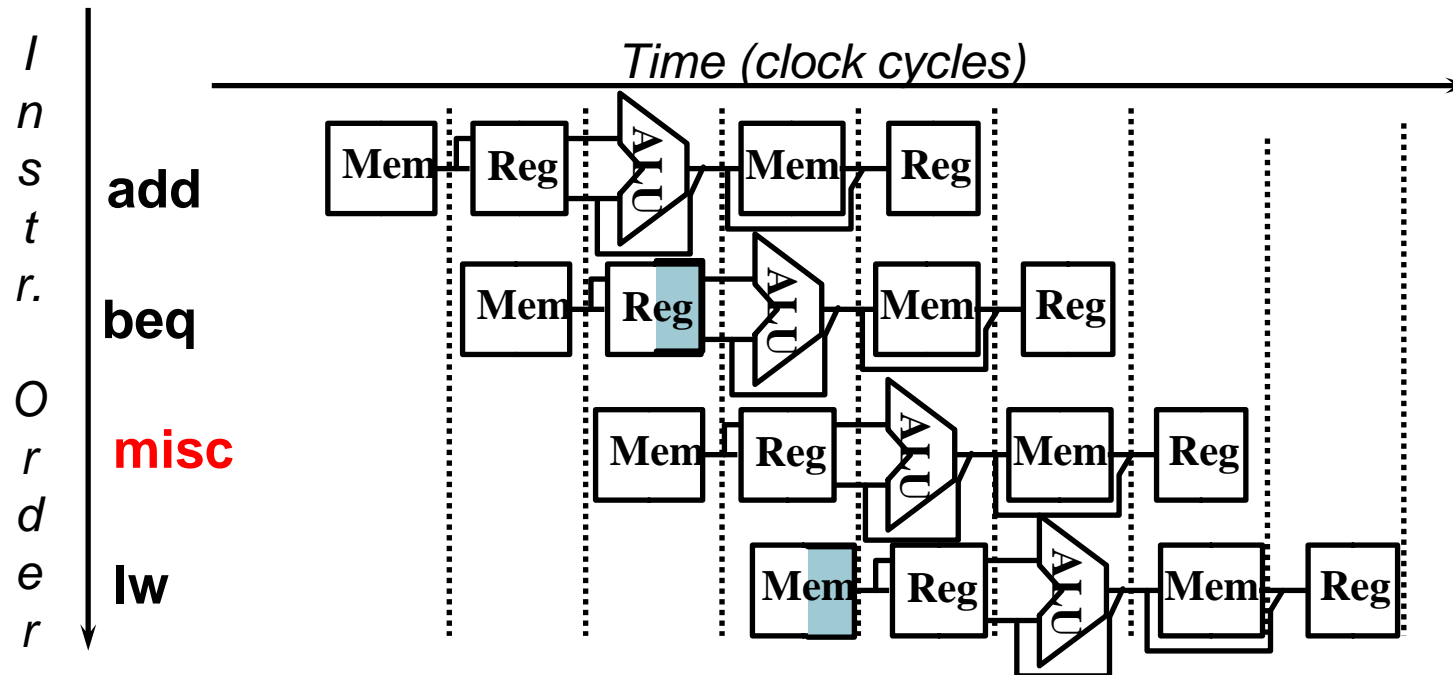
# More-Realistic Branch Prediction

- Static branch prediction
  - Based on typical branch behavior
  - Example: loop and if-statement branches
    - Predict backward branches taken
    - Predict forward branches not taken
- Dynamic branch prediction
  - Hardware measures actual branch behavior
    - e.g., record recent history of each branch
  - Assume future behavior will continue the trend
    - When wrong, stall while re-fetching, and update history



# Delayed Branch

- Predict-not-taken + branch decision at ID  
 => the following instruction is always executed  
 => branches take effect 1 cycle later



- 0 clock cycle penalty per branch instruction if can find instruction to put in slot ( $\cong 50\%$  of time)

# Pipeline Summary

## The BIG Picture

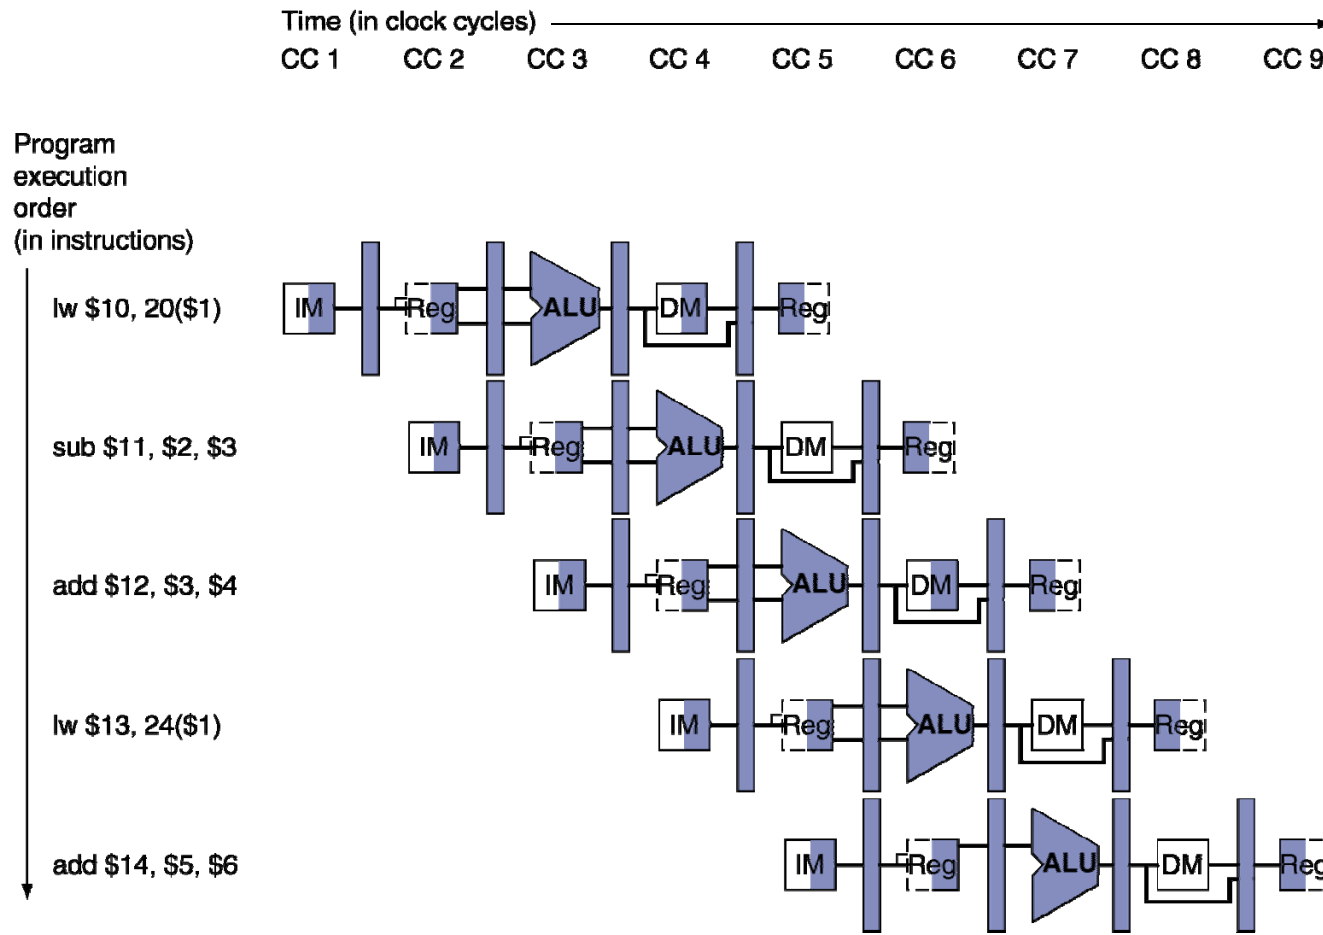
- Pipelining improves performance by increasing instruction throughput
  - Executes multiple instructions in parallel
  - Each instruction has the same latency
- Subject to hazards
  - Structure, data, control
- Instruction set design affects complexity of pipeline implementation

# Pipeline Operation

- Cycle-by-cycle flow of instructions through the pipelined datapath
  - “Single-clock-cycle” pipeline diagram
    - Shows pipeline usage in a single cycle
    - Highlight resources used
  - c.f. “multi-clock-cycle” diagram
    - Graph of operation over time

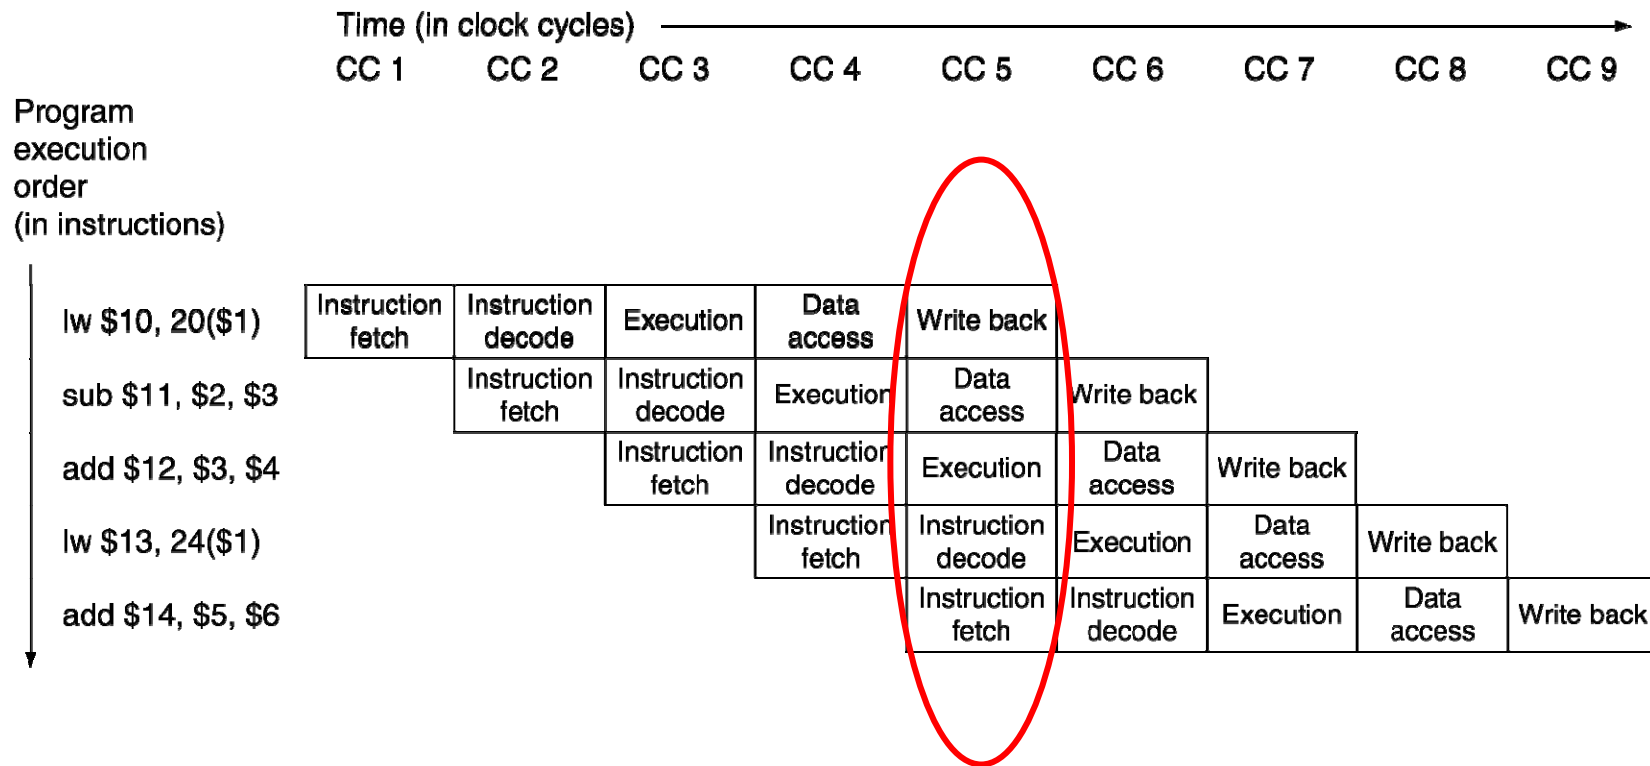
# Multi-Cycle Pipeline Diagram

- Form showing resource usage



# Multi-Cycle Pipeline Diagram

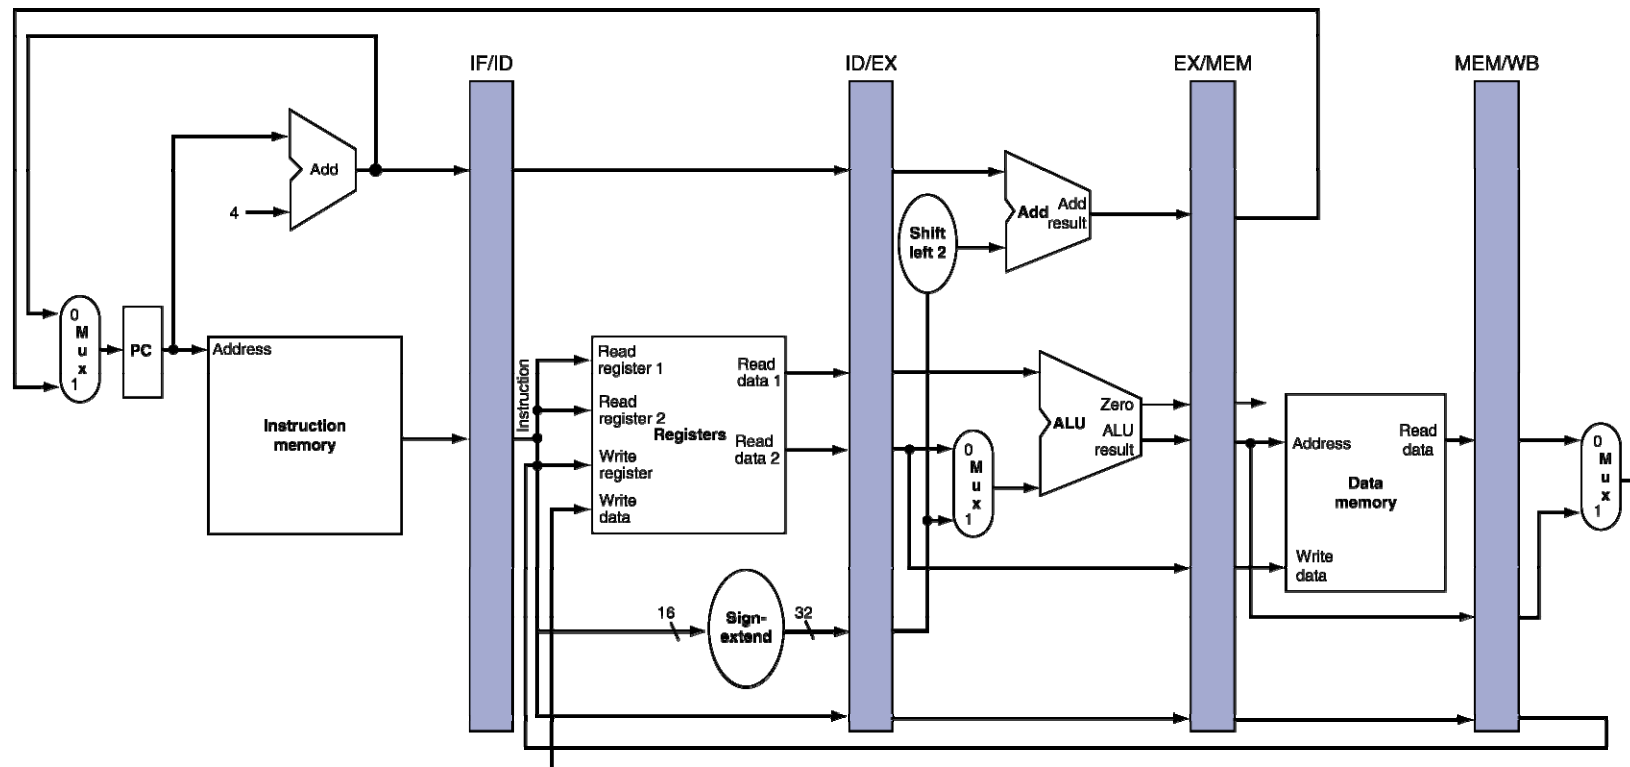
## ■ Traditional form



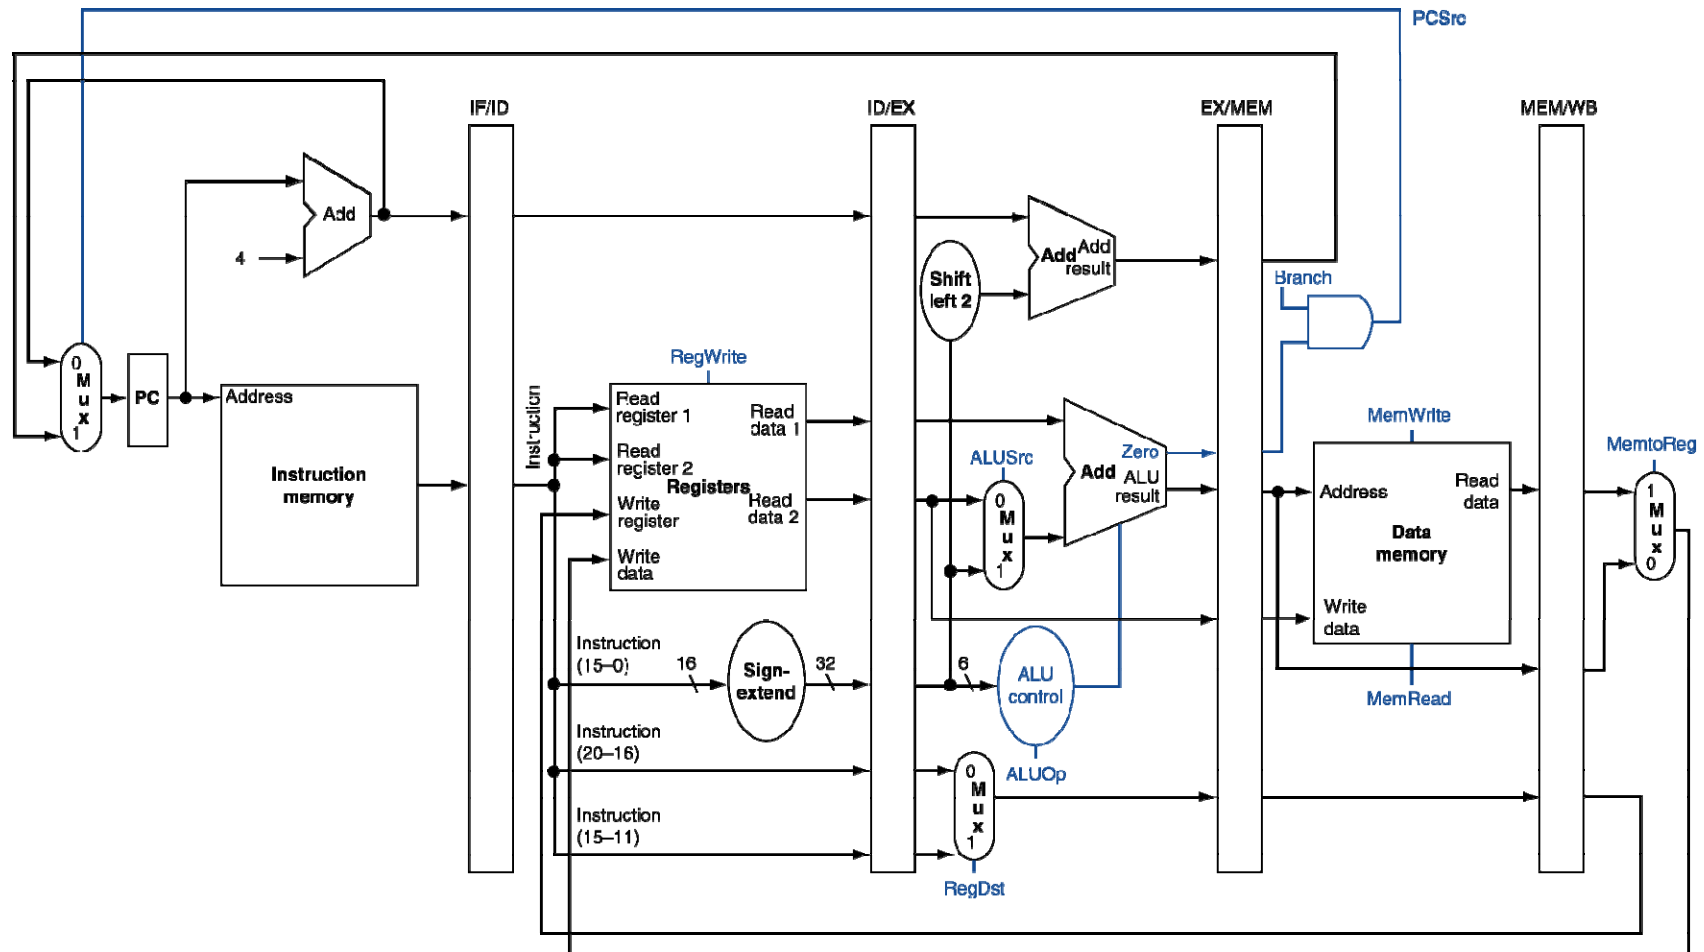
# Single-Cycle Pipeline Diagram

- State of pipeline in a given cycle: cc5

add \$14, \$5, \$6	lw \$13, 24 (\$1)	add \$12, \$3, \$4	sub \$11, \$2, \$3	lw \$10, 20(\$1)
Instruction fetch	Instruction decode	Execution	Memory	Write-back

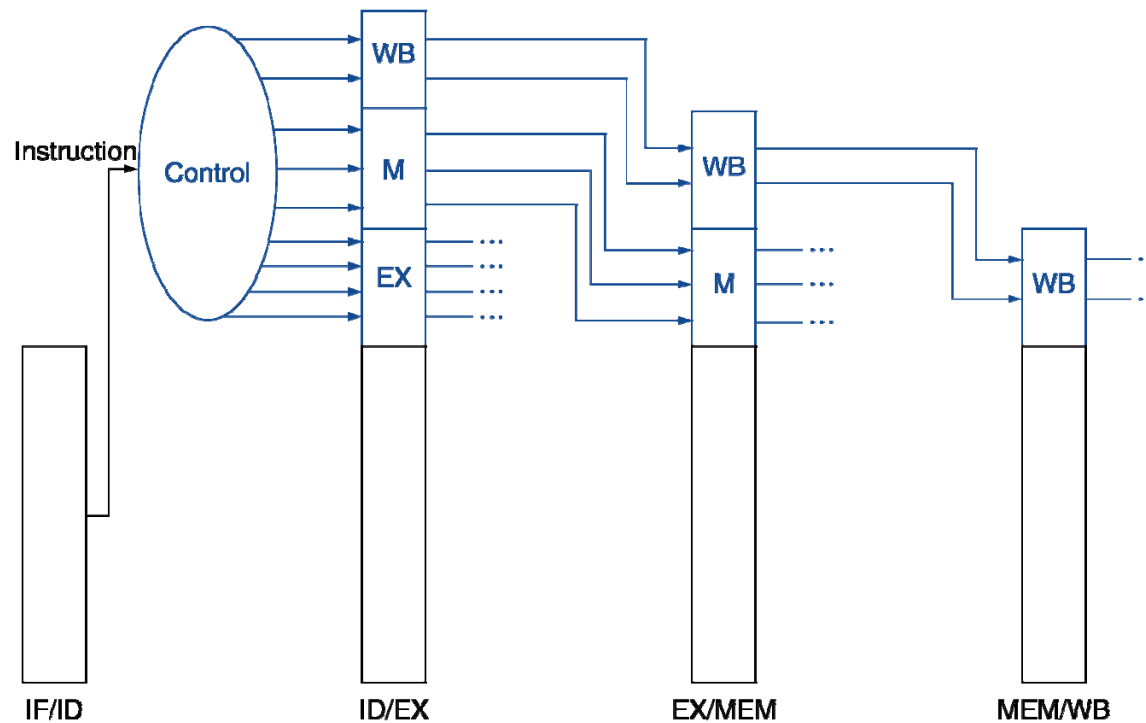


# Pipelined Control (Simplified)



# Data Stationary Pipelined Control

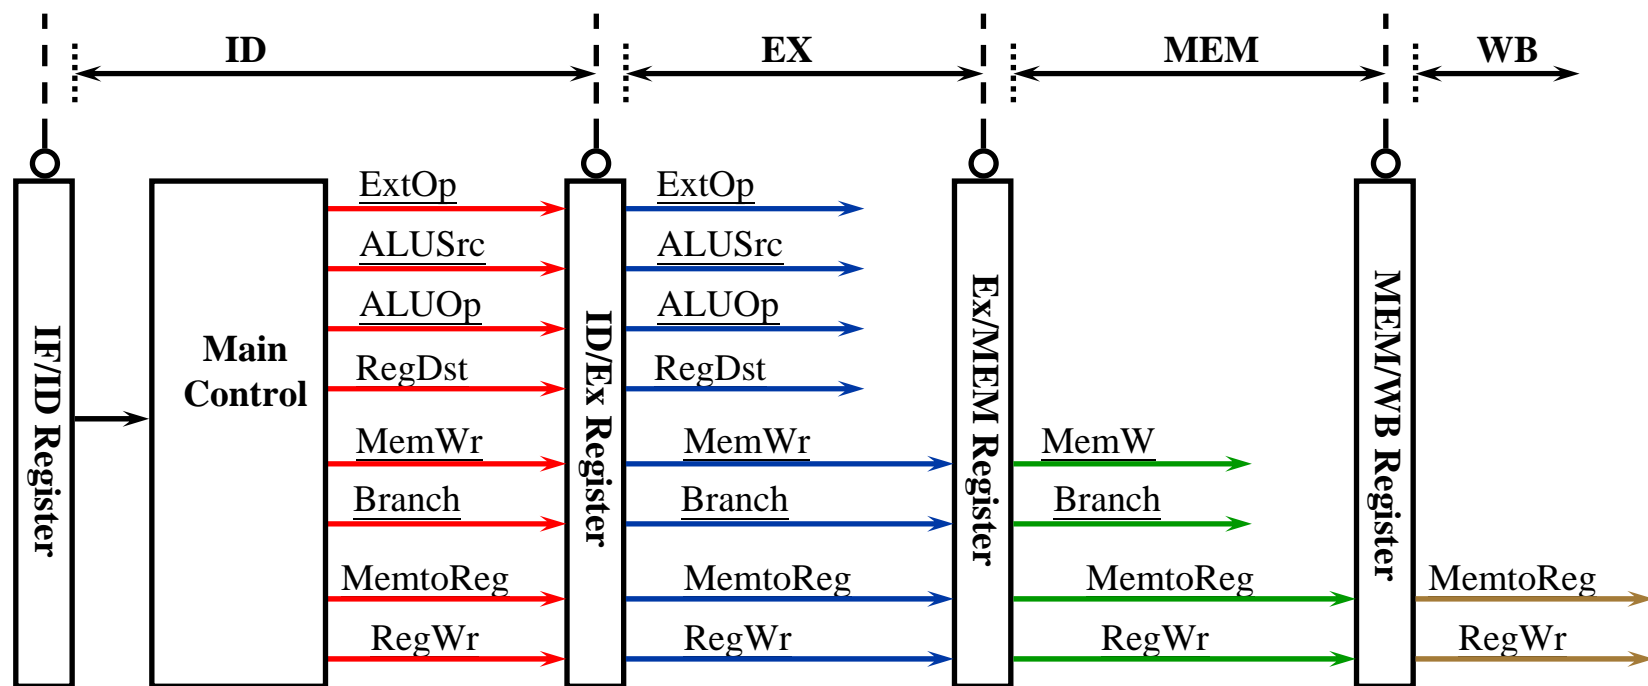
- Control signals derived from instruction
  - Main control generates control signals during ID
  - Pass control signals along just like the data



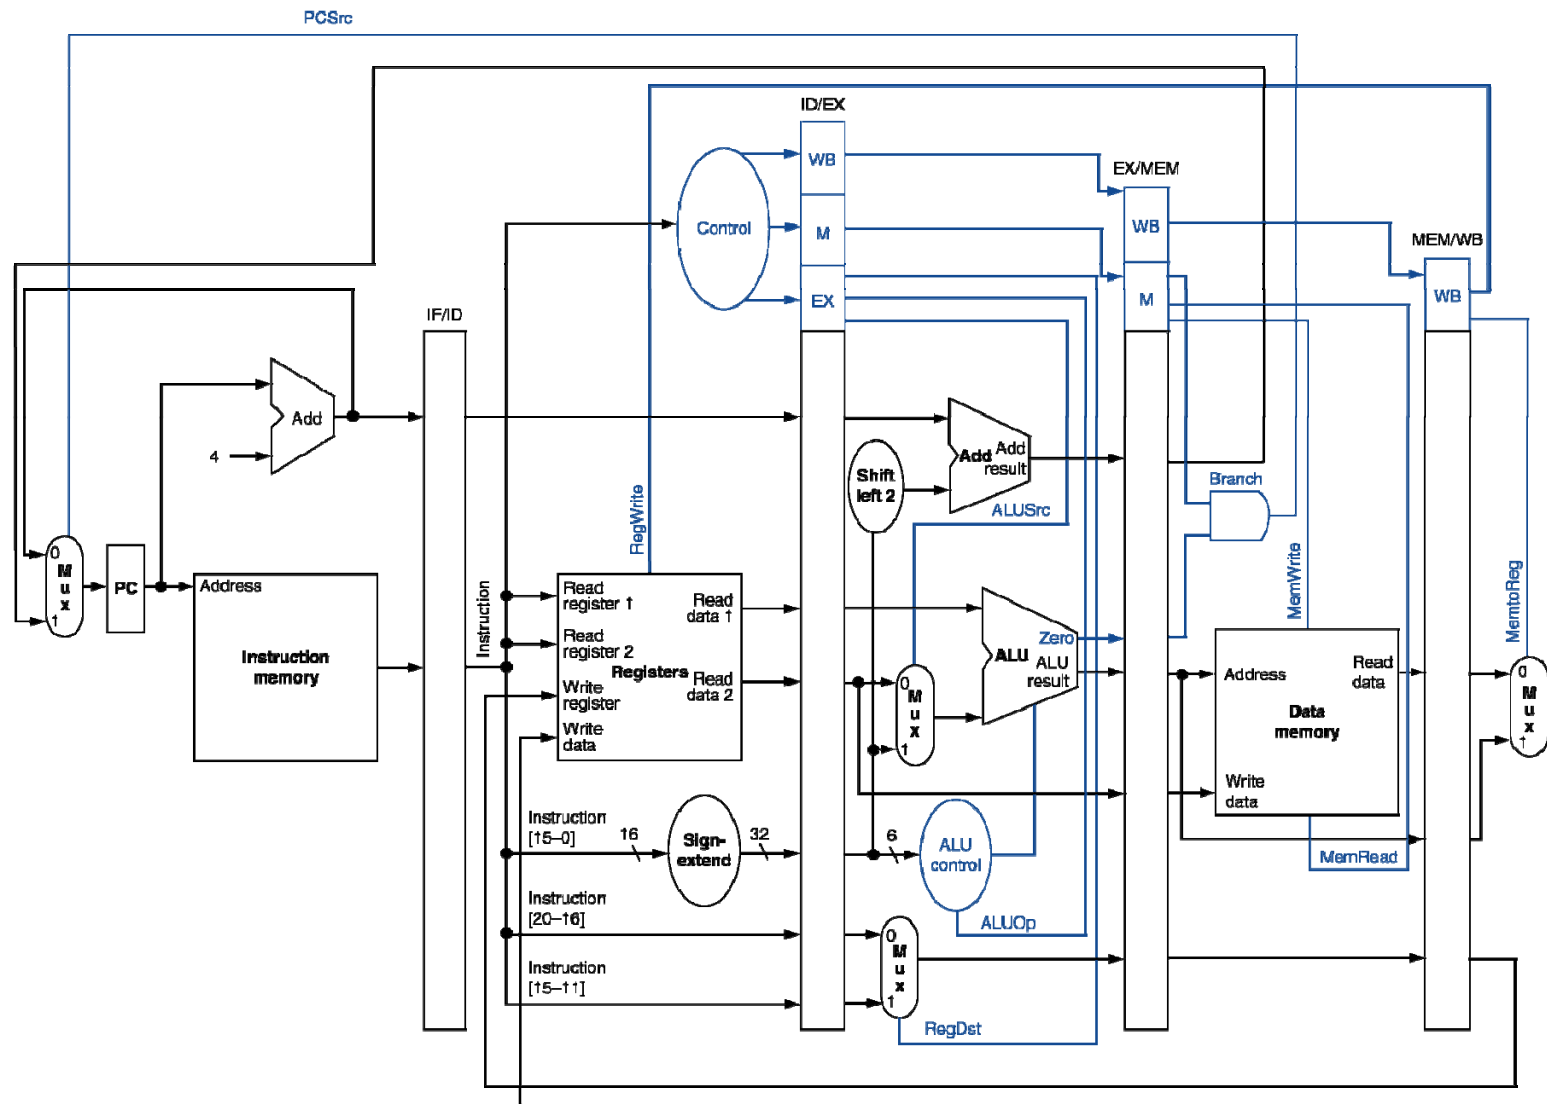


# Data Stationary Control

- Signals for EX (ExtOp, ALUSrc, ...) are used 1 cycle later
- Signals for MEM (MemWr, Branch) are used 2 cycles later
- Signals for WB (MemtoReg, MemWr) are used 3 cycles later



# Pipelined Control



# Data Hazards in ALU Instructions

- Consider this sequence:

sub \$2, \$1, \$3

and \$12, \$2, \$5

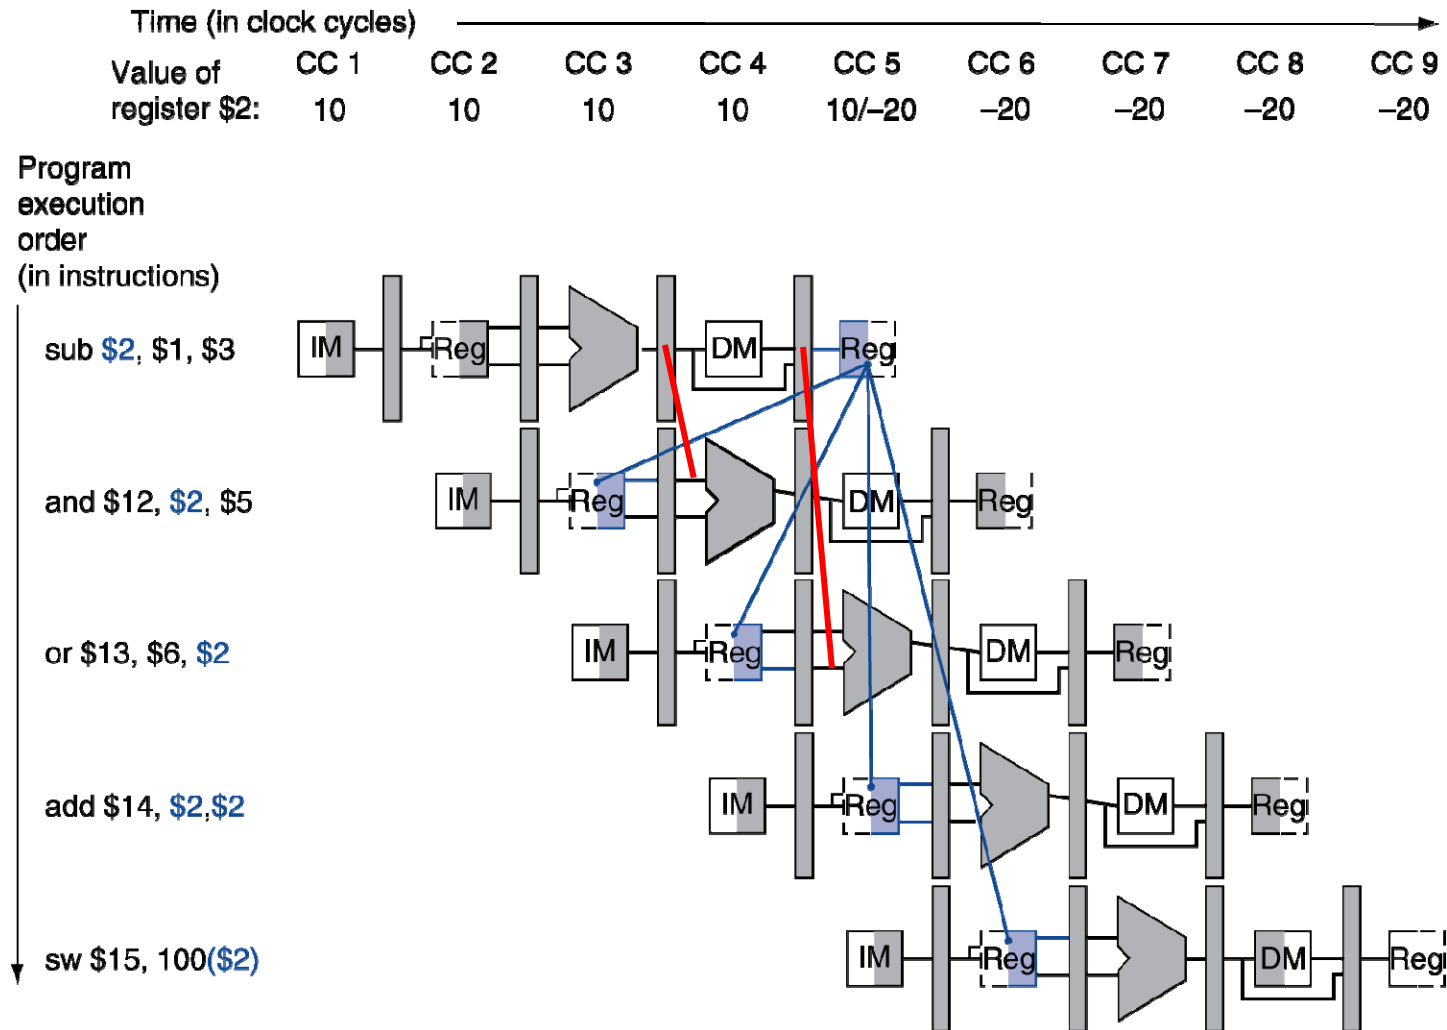
or \$13, \$6, \$2

add \$14, \$2, \$2

sw \$15, 100(\$2)

- We can resolve hazards with forwarding
  - How do we detect when to forward?

# Dependencies & Forwarding





# Detecting the Need to Forward

- Pass register numbers along pipeline
  - e.g., ID/EX.RegisterRs = register number for Rs sitting in ID/EX pipeline register
- ALU operand register numbers in EX stage are given by
  - ID/EX.RegisterRs, ID/EX.RegisterRt
- Data hazards when
 

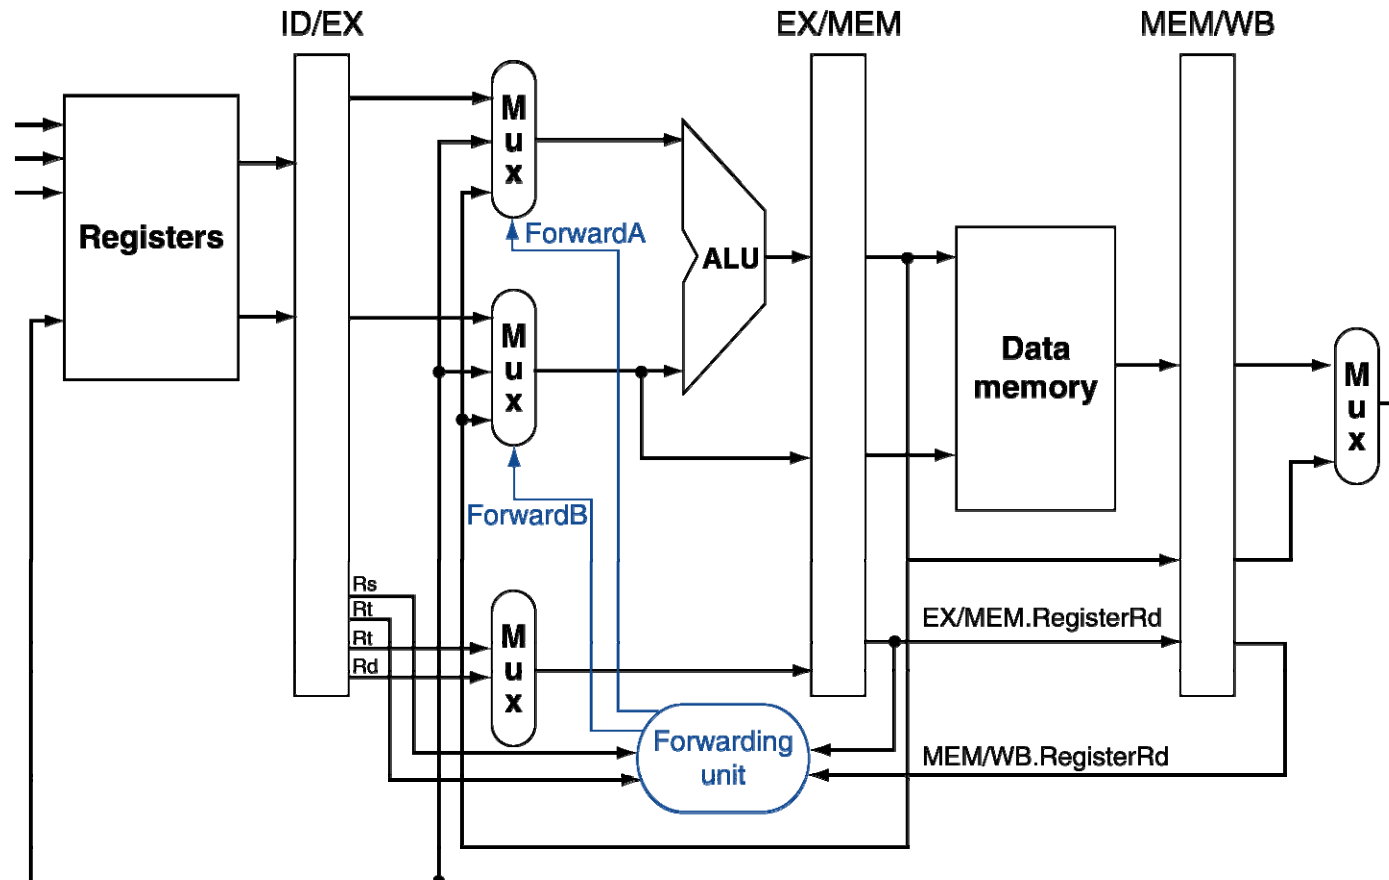
1a. EX/MEM.RegisterRd = ID/EX.RegisterRs 1b. EX/MEM.RegisterRd = ID/EX.RegisterRt	}	Fwd from EX/MEM pipeline reg
2a. MEM/WB.RegisterRd = ID/EX.RegisterRs 2b. MEM/WB.RegisterRd = ID/EX.RegisterRt	}	Fwd from MEM/WB pipeline reg



# Detecting the Need to Forward

- But only if forwarding instruction will write to a register!
  - EX/MEM.RegWrite, MEM/WB.RegWrite
- And only if Rd for that instruction is not \$zero
  - EX/MEM.RegisterRd  $\neq$  0,  
MEM/WB.RegisterRd  $\neq$  0

# Forwarding Paths



b. With forwarding

# Forwarding Conditions

- EX hazard
  - if (EX/MEM.RegWrite and (EX/MEM.RegisterRd  $\neq$  0)  
and (EX/MEM.RegisterRd = ID/EX.RegisterRs))  
ForwardA = 10
  - if (EX/MEM.RegWrite and (EX/MEM.RegisterRd  $\neq$  0)  
and (EX/MEM.RegisterRd = ID/EX.RegisterRt))  
ForwardB = 10
- MEM hazard
  - if (MEM/WB.RegWrite and (MEM/WB.RegisterRd  $\neq$  0)  
and (MEM/WB.RegisterRd = ID/EX.RegisterRs))  
ForwardA = 01
  - if (MEM/WB.RegWrite and (MEM/WB.RegisterRd  $\neq$  0)  
and (MEM/WB.RegisterRd = ID/EX.RegisterRt))  
ForwardB = 01



# Double Data Hazard

- Consider the sequence:

```
add $1, $1, $2
```

```
add $1, $1, $3
```

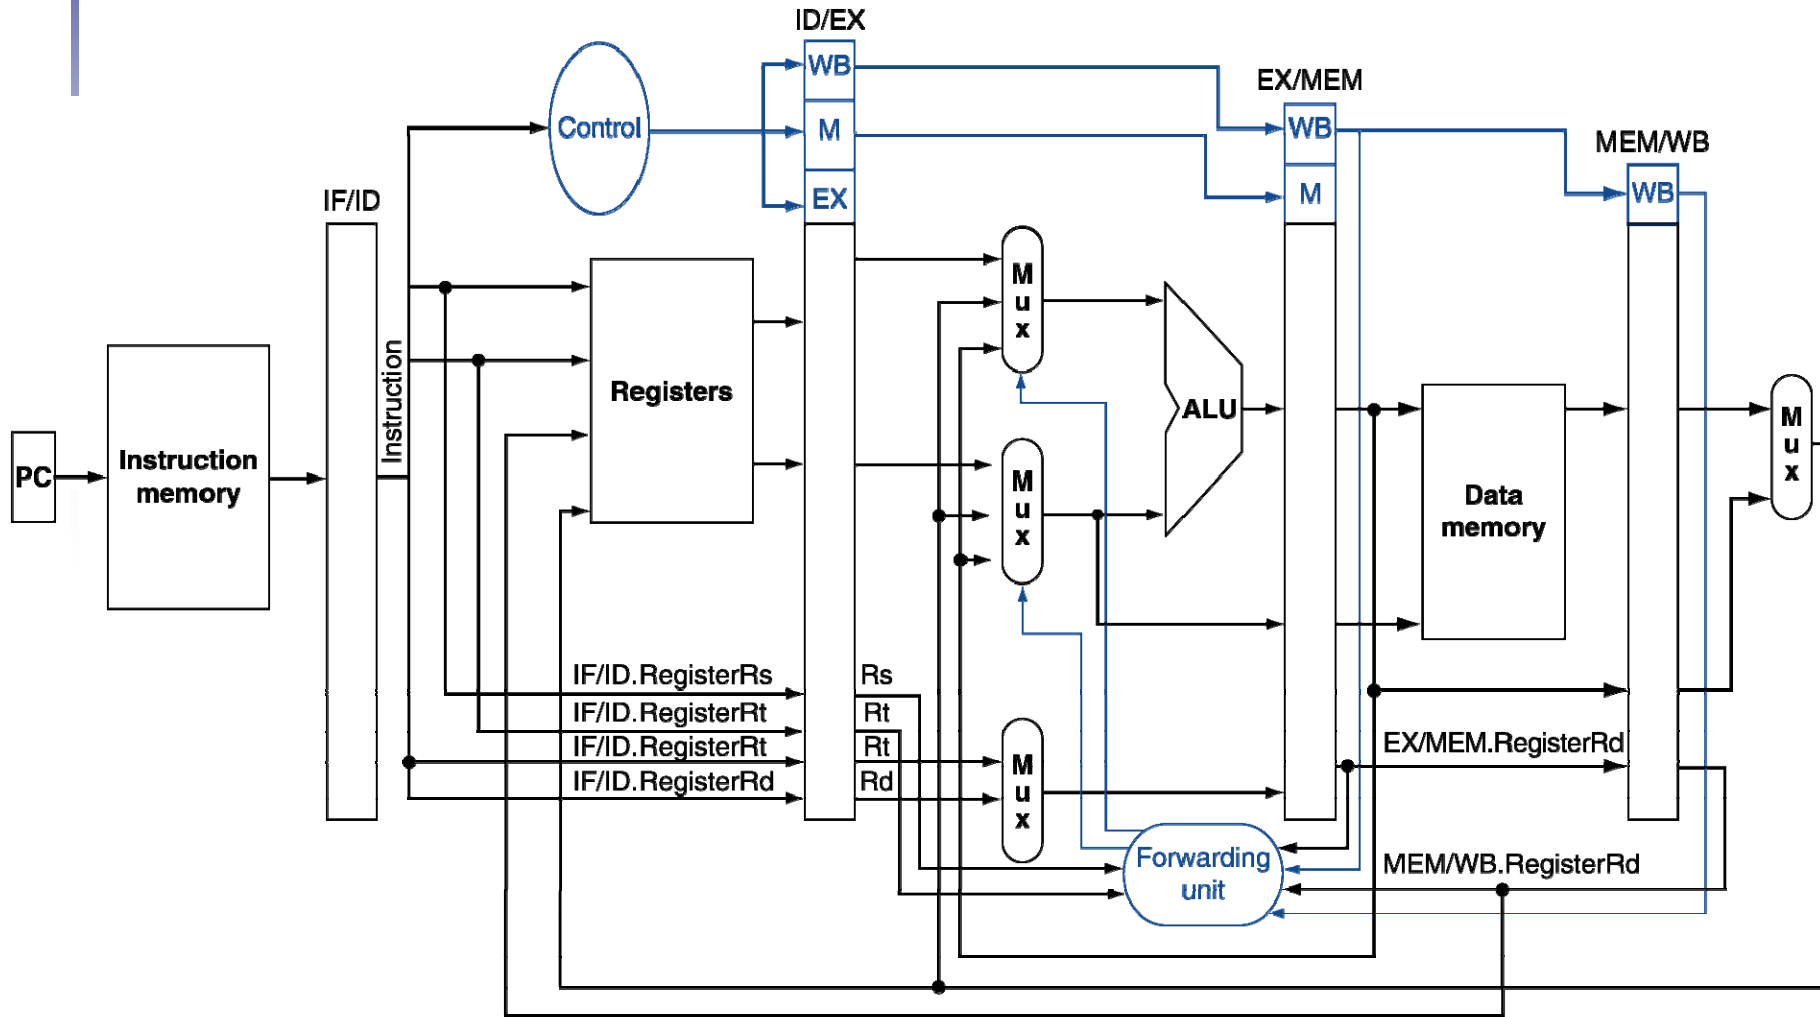
```
add $1, $1, $4
```

- Both hazards occur
  - Want to use the most recent
- Revise MEM hazard condition
  - Only fwd if EX hazard condition isn't true

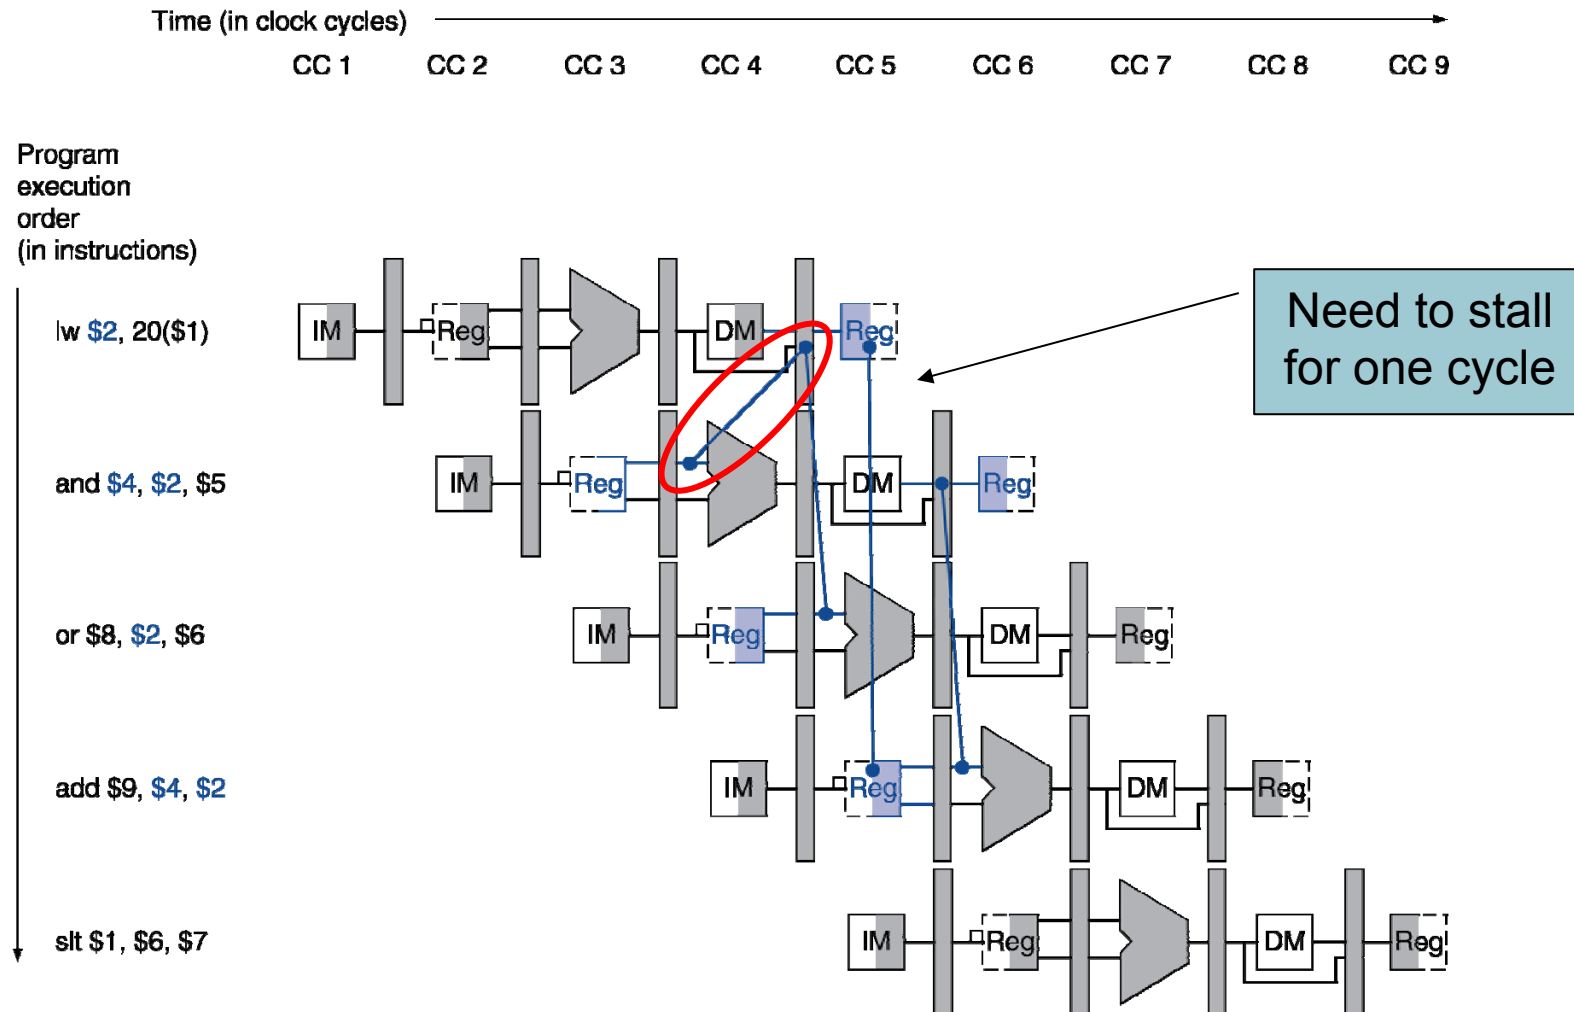
# Revised Forwarding Condition

- MEM hazard
  - if (MEM/WB.RegWrite and (MEM/WB.RegisterRd  $\neq$  0)  
 and not (EX/MEM.RegWrite and (EX/MEM.RegisterRd  $\neq$  0)  
 and (EX/MEM.RegisterRd = ID/EX.RegisterRs))  
 and (MEM/WB.RegisterRd = ID/EX.RegisterRs))  
 ForwardA = 01
  - if (MEM/WB.RegWrite and (MEM/WB.RegisterRd  $\neq$  0)  
 and not (EX/MEM.RegWrite and (EX/MEM.RegisterRd  $\neq$  0)  
 and (EX/MEM.RegisterRd = ID/EX.RegisterRt))  
 and (MEM/WB.RegisterRd = ID/EX.RegisterRt))  
 ForwardB = 01

# Datapath with Forwarding



# Load-Use Data Hazard



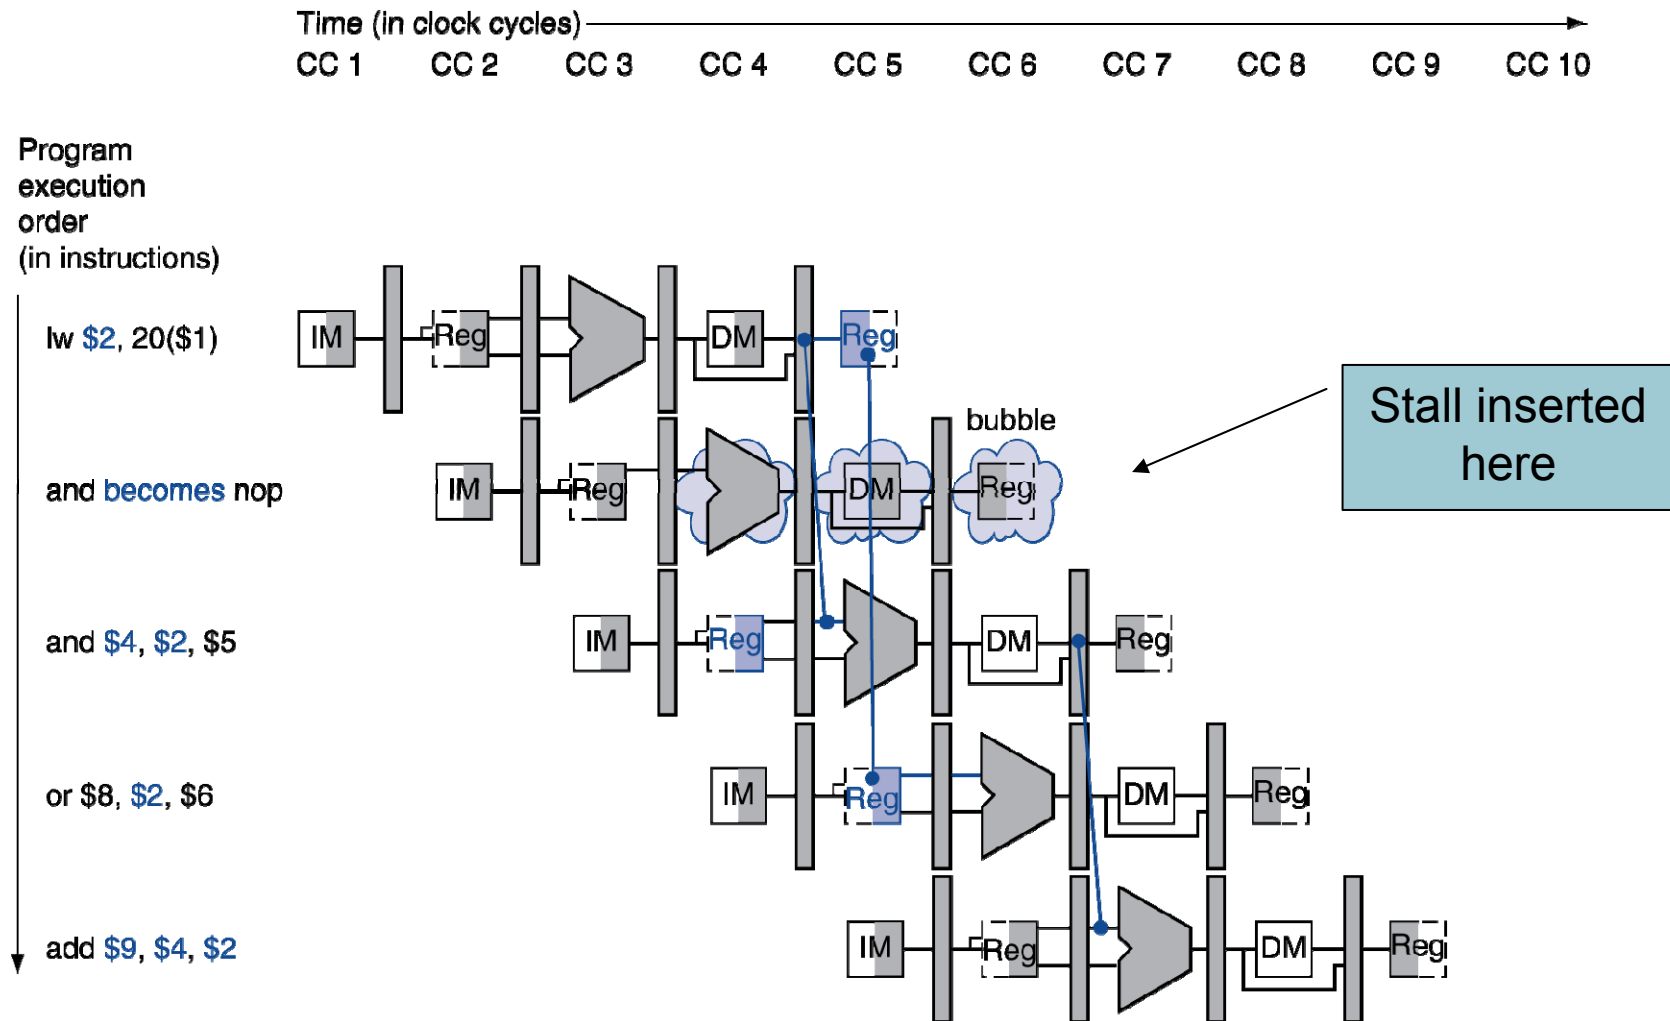
# Load-Use Hazard Detection

- Check when using instruction is decoded in ID stage
- ALU operand register numbers in ID stage are given by
  - IF/ID.RegisterRs, IF/ID.RegisterRt
- Load-use hazard when
  - ID/EX.MemRead and
    - ((ID/EX.RegisterRt = IF/ID.RegisterRs) or (ID/EX.RegisterRt = IF/ID.RegisterRt))
- If detected, stall and insert bubble

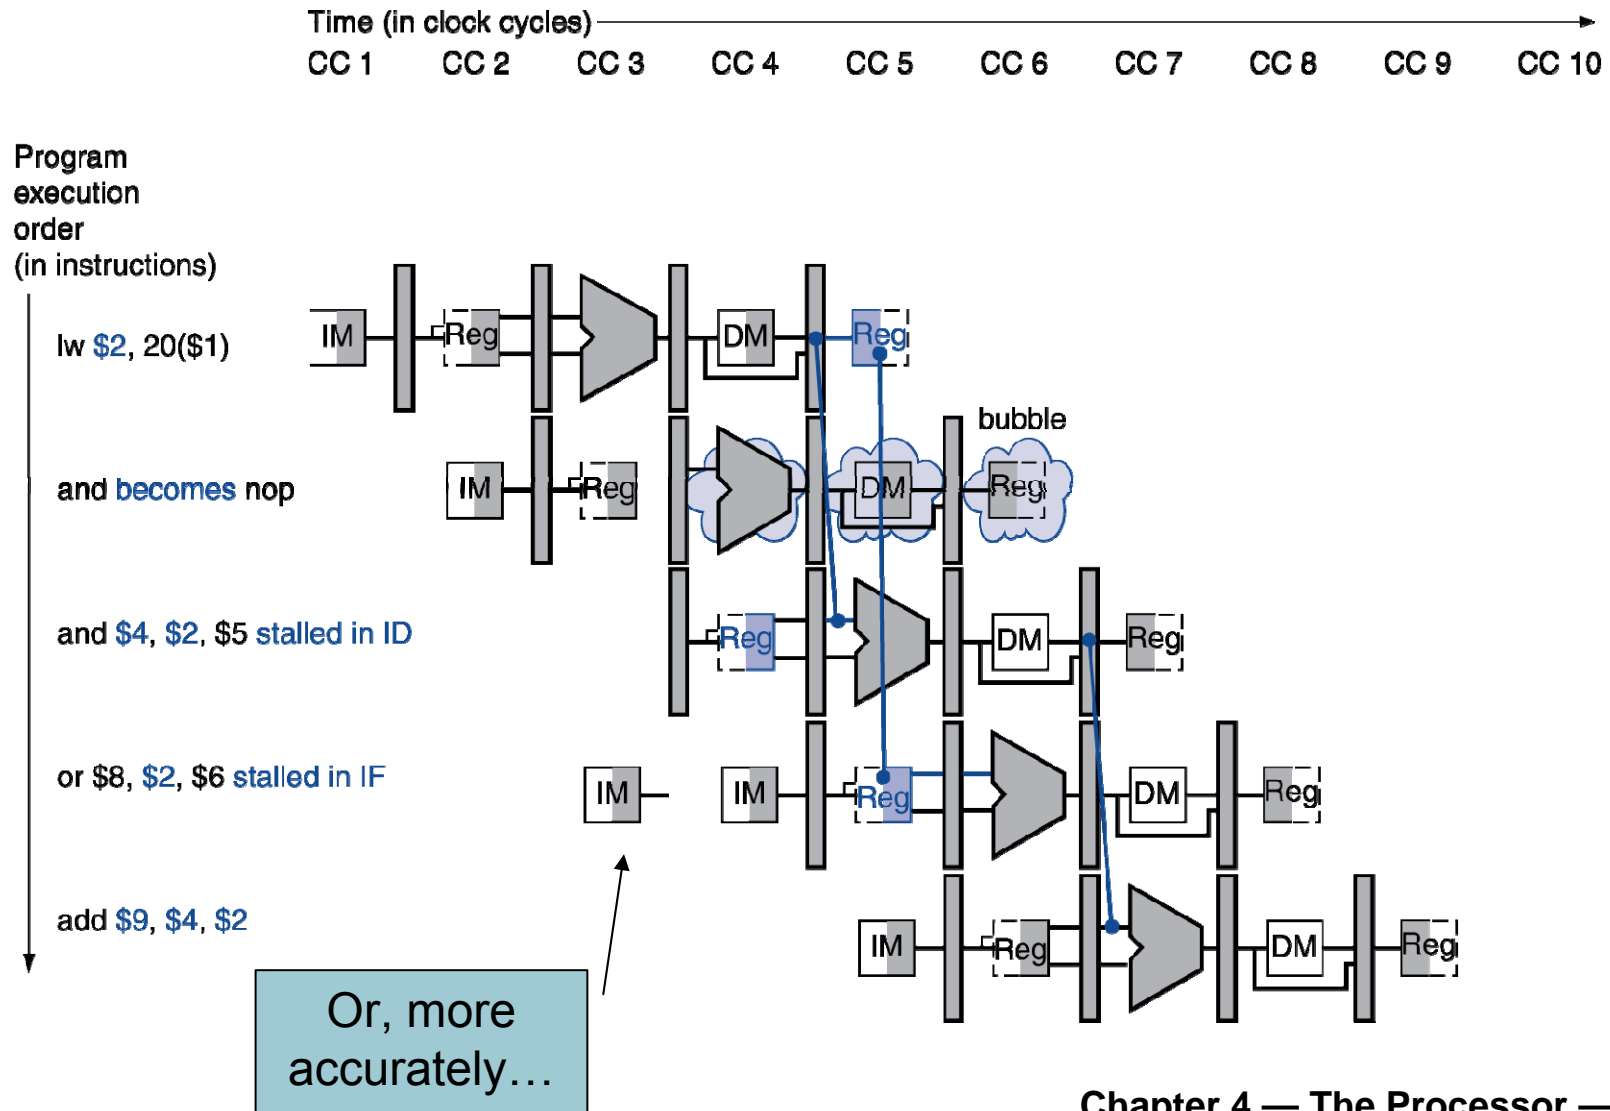
# How to Stall the Pipeline

- Force control values in ID/EX register to 0
  - EX, MEM and WB do nop (no-operation)
- Prevent update of PC and IF/ID register
  - Using instruction is decoded again
  - Following instruction is fetched again
  - 1-cycle stall allows MEM to read data for I w
    - Can subsequently forward to EX stage

# Stall/Bubble in the Pipeline

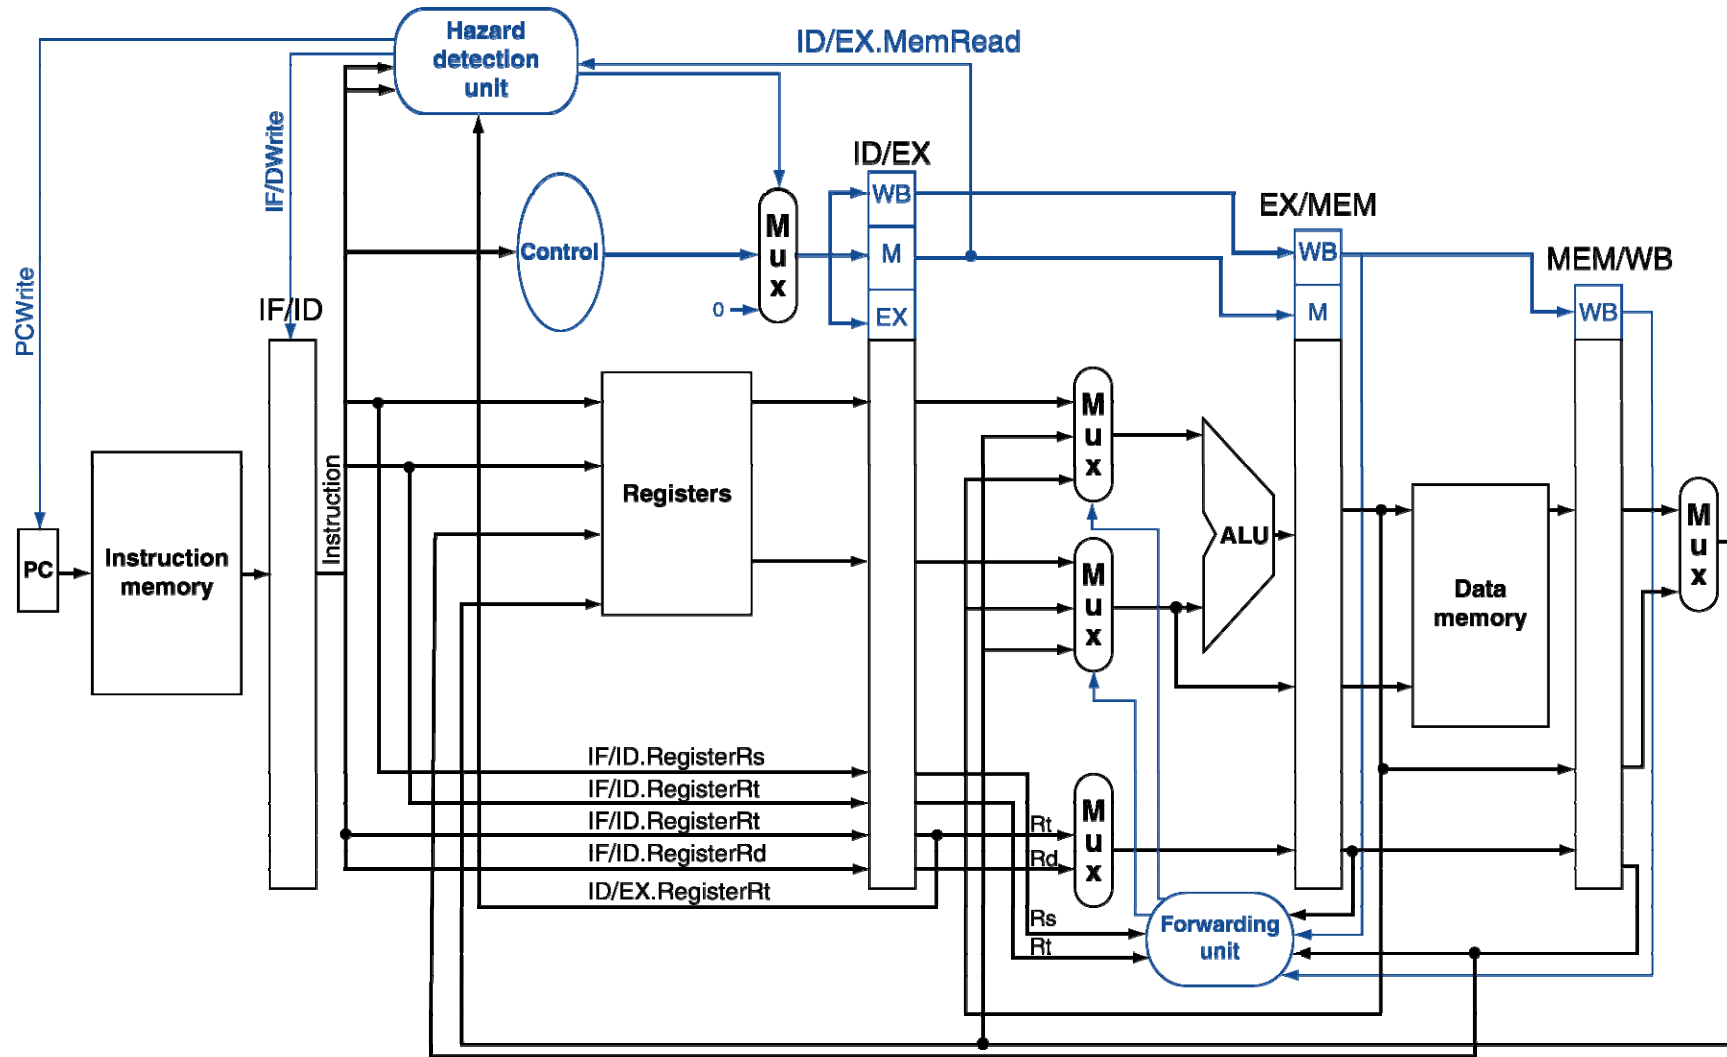


# Stall/Bubble in the Pipeline





# Datapath with Hazard Detection



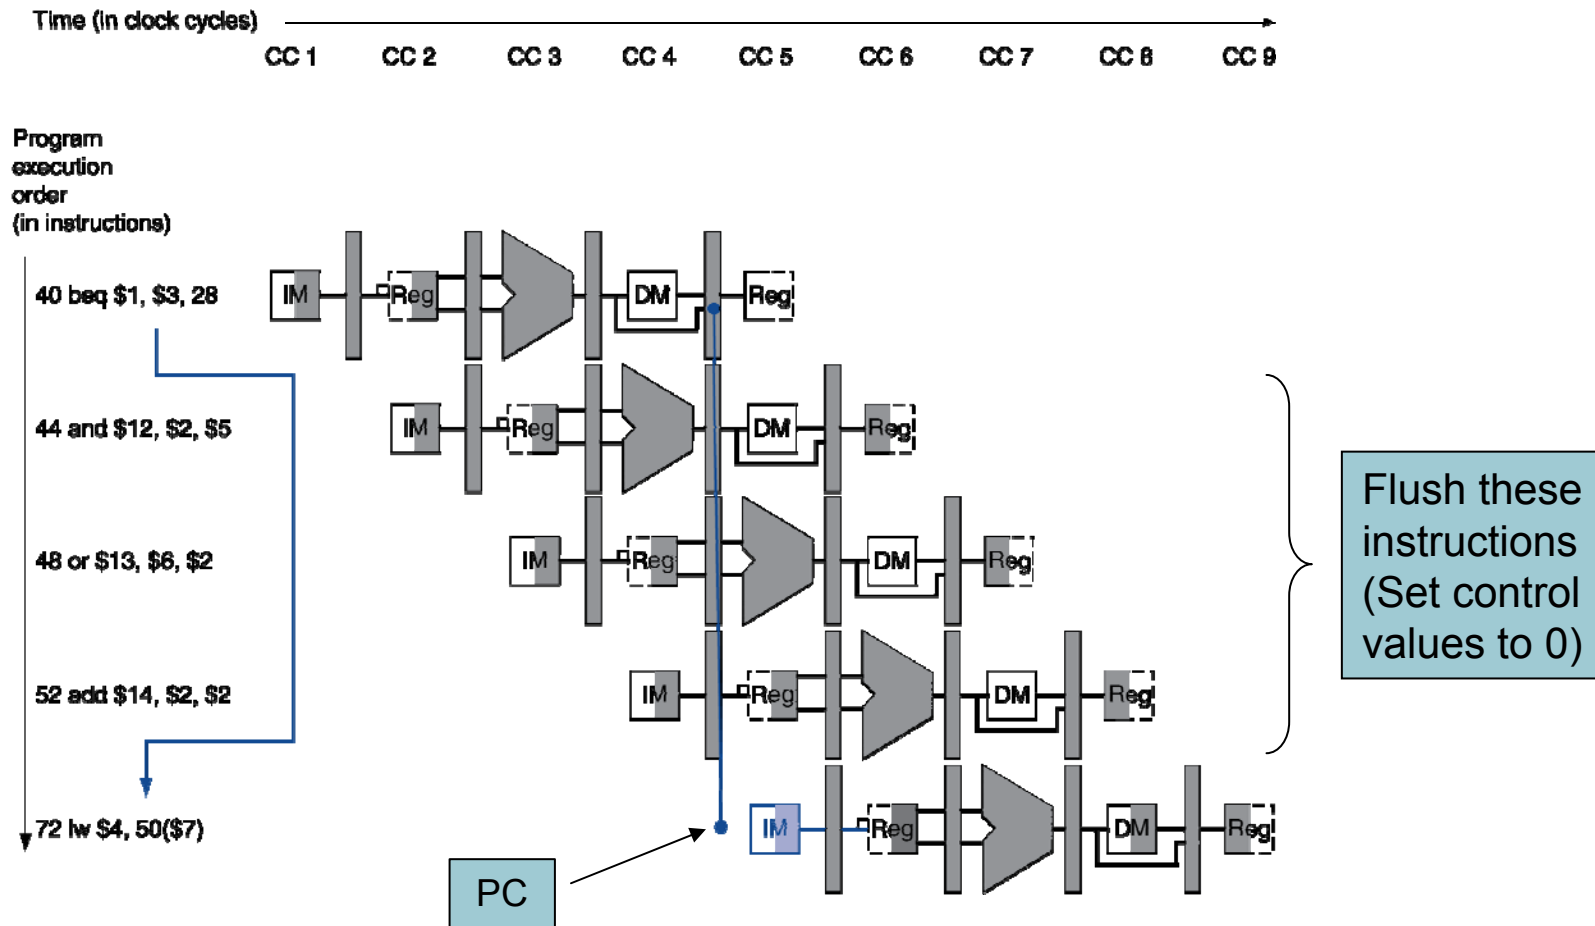
# Stalls and Performance

## The BIG Picture

- Stalls reduce performance
  - But are required to get correct results
- Compiler can arrange code to avoid hazards and stalls
  - Requires knowledge of the pipeline structure

# Branch Hazards

- If branch outcome determined in MEM



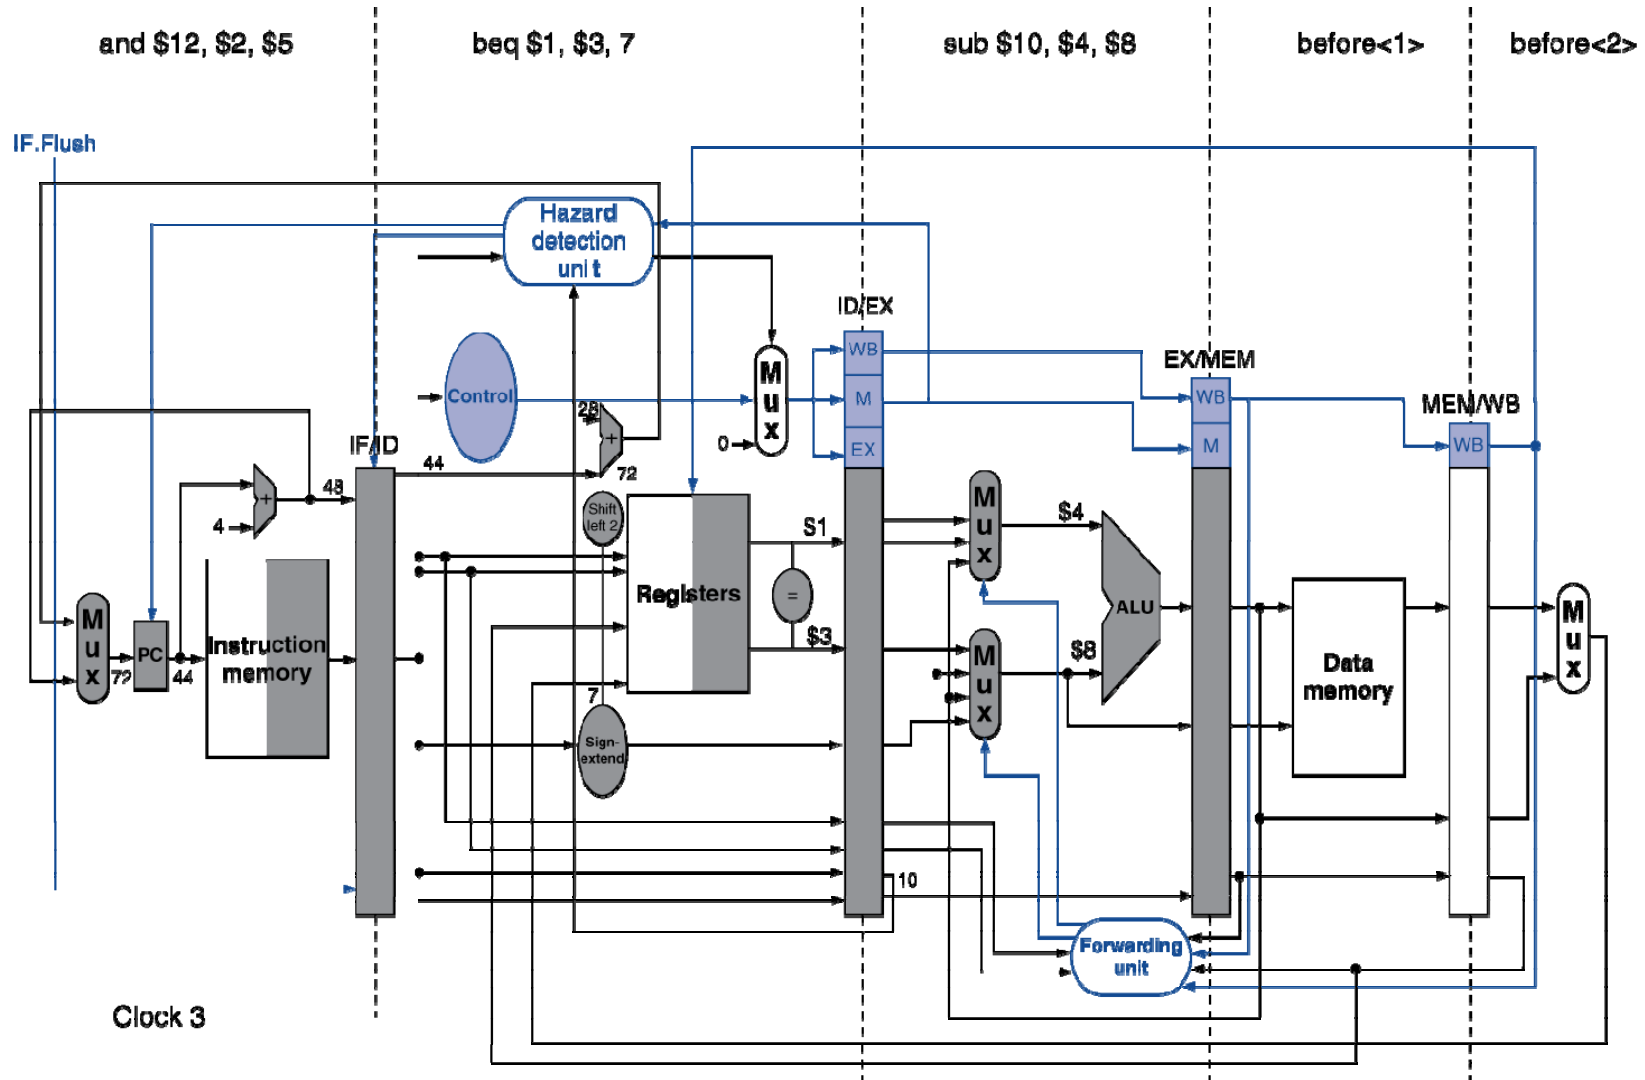


# Reducing Branch Delay

- Move hardware to determine outcome to ID stage
  - Target address adder
  - Register comparator
- Example: branch taken

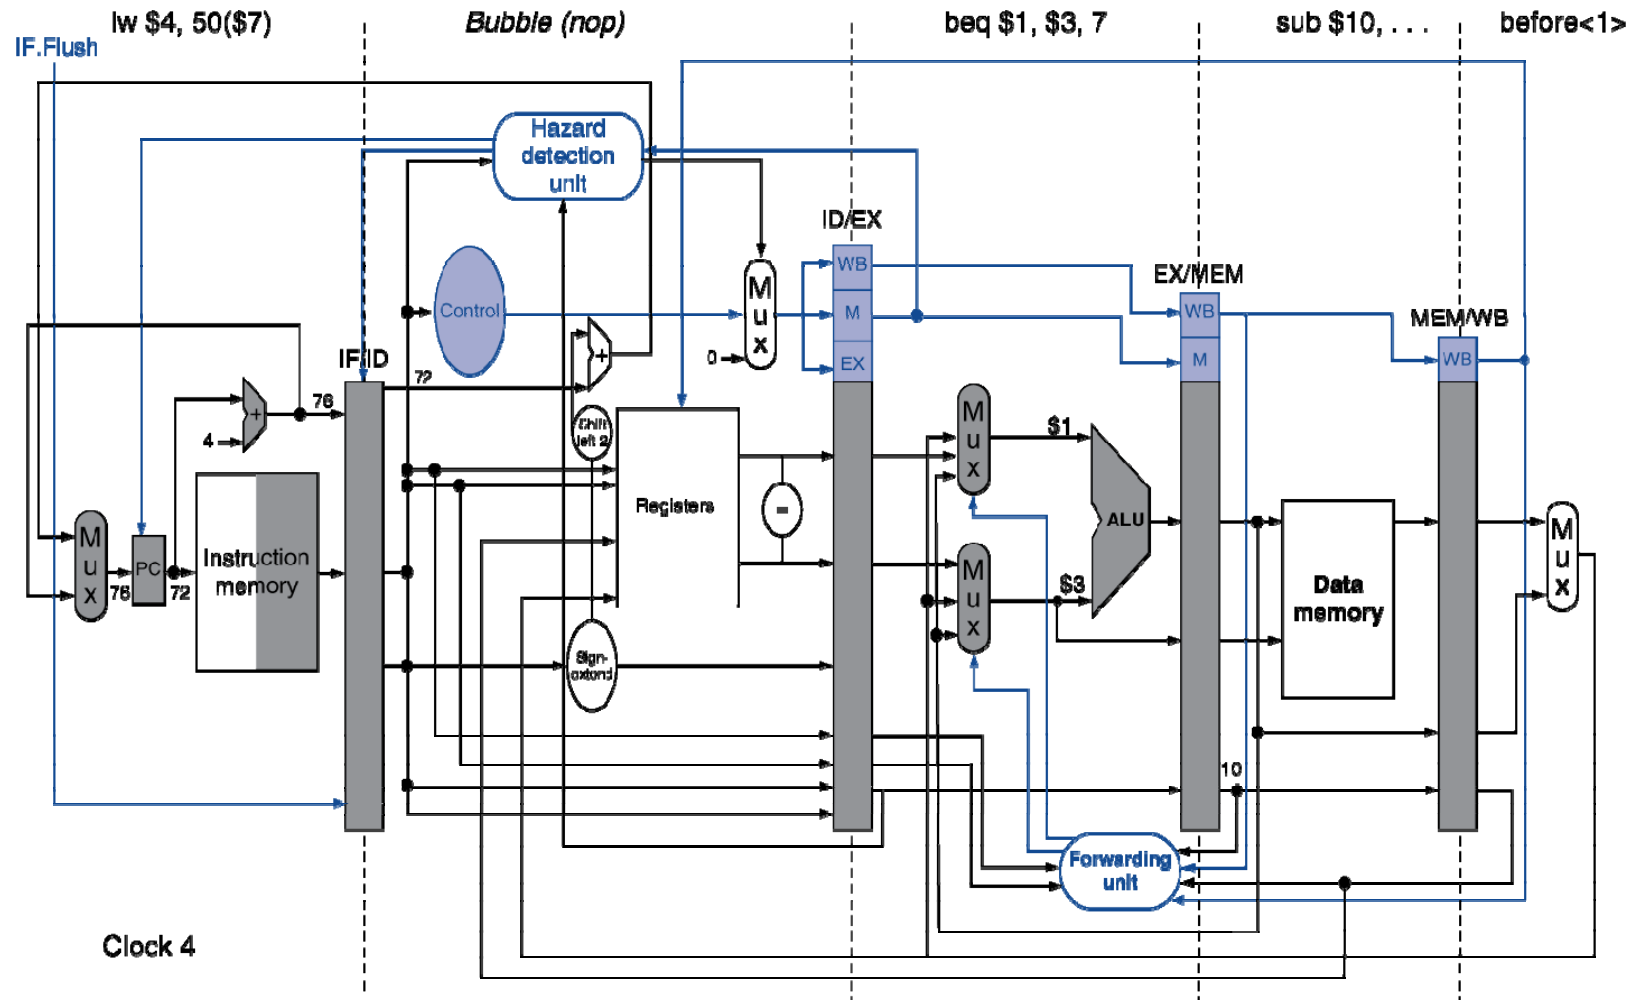
```
36:  sub  $10, $4, $8
40:  beq  $1,  $3,  7
44:  and  $12, $2, $5
48:  or   $13, $2, $6
52:  add  $14, $4, $2
56:  slt  $15, $6, $7
    ...
72:  lw   $4,  50($7)
```

# Example: Branch Taken



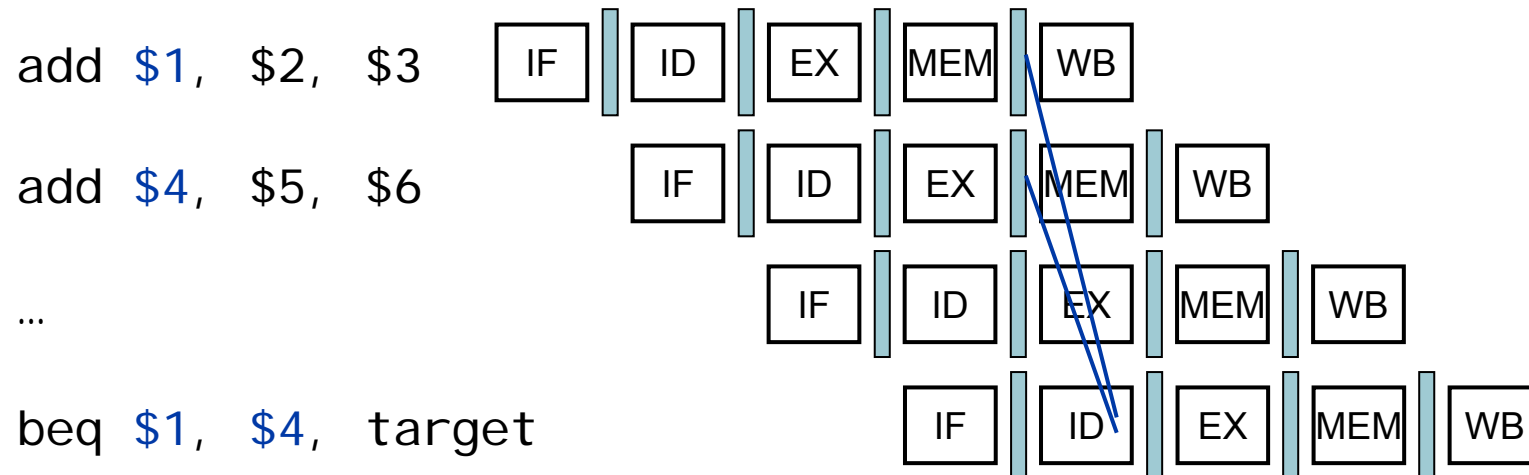
Clock 3

# Example: Branch Taken



# Data Hazards for Branches

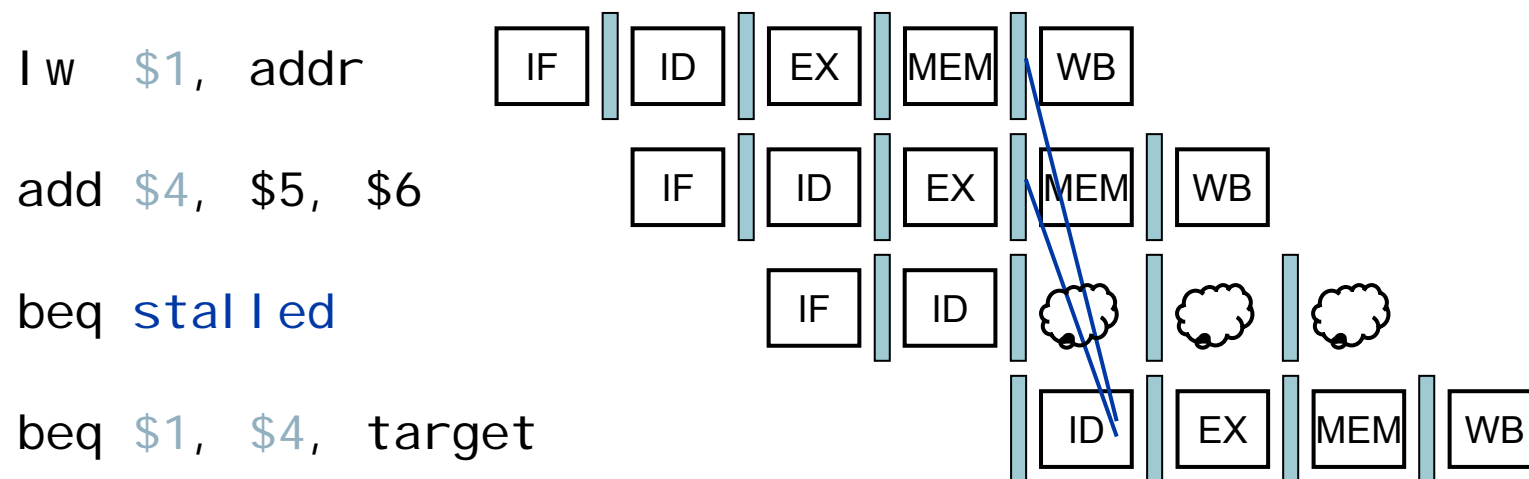
- If a comparison register is a destination of 2<sup>nd</sup> or 3<sup>rd</sup> preceding ALU instruction



- Can resolve using forwarding

# Data Hazards for Branches


- If a comparison register is a destination of preceding ALU instruction or 2nd preceding load instruction
  - Need 1 stall cycle





# Data Hazards for Branches

- If a comparison register is a destination of immediately preceding load instruction
  - Need 2 stall cycles





# Dynamic Branch Prediction

- In deeper and superscalar pipelines, branch penalty is more significant
- Use dynamic prediction
  - Branch prediction buffer (aka branch history table)
  - Indexed by recent branch instruction addresses
  - Stores outcome (taken/not taken)
  - To execute a branch
    - Check table, expect the same outcome
    - Start fetching from fall-through or target
    - If wrong, flush pipeline and flip prediction

# 1-Bit Predictor: Shortcoming

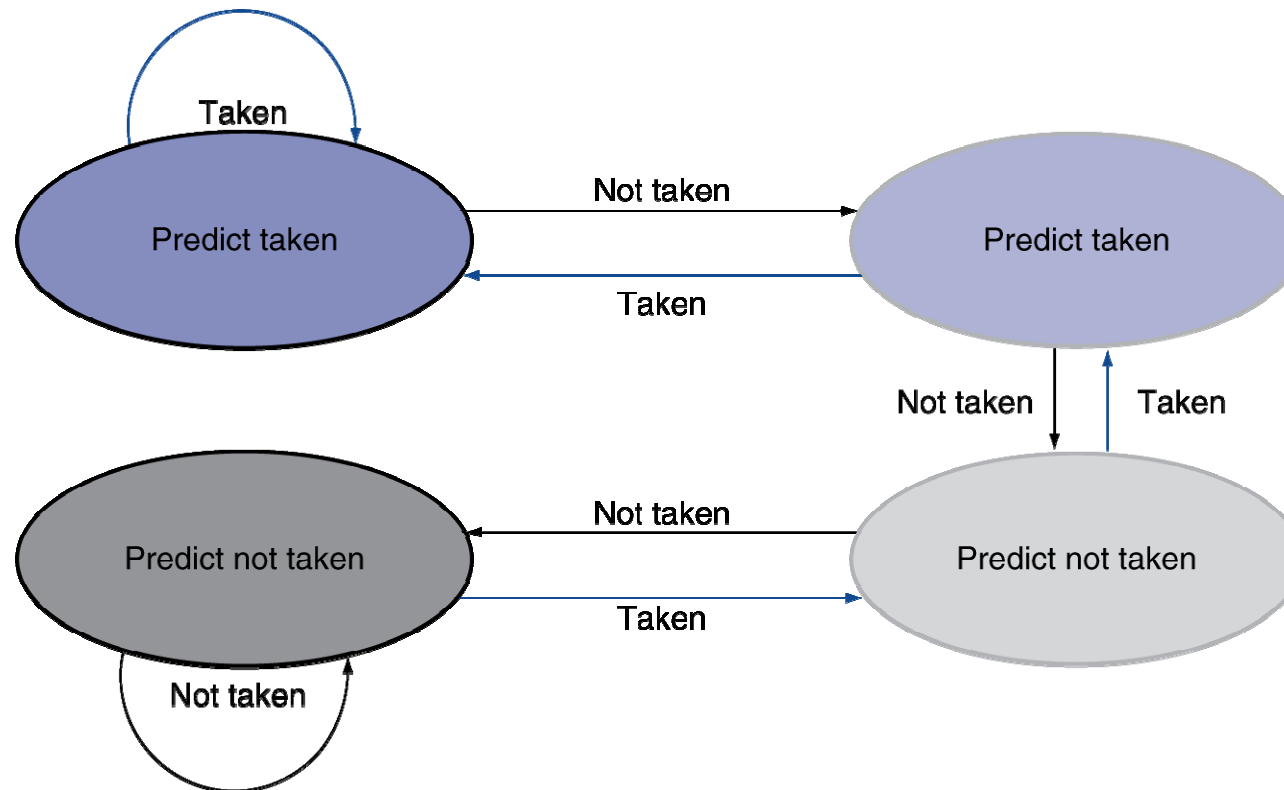
- Inner loop branches mispredicted twice!



- Mispredict as taken on last iteration of inner loop
- Then mispredict as not taken on first iteration of inner loop next time around

# 2-Bit Predictor

- Only change prediction on two successive mispredictions





# Calculating the Branch Target Address

- Even with predictor, still need to calculate the target address
  - 1-cycle penalty for a taken branch
- Branch target buffer
  - Cache of target addresses
  - Indexed by PC when instruction fetched
    - If hit and instruction is branch predicted taken, can fetch target immediately

# Exceptions and Interrupts

- “Unexpected” events requiring change in flow of control
  - Different ISAs use the terms differently
- Exception
  - Arises within the CPU
    - e.g., undefined opcode, overflow, syscall, ...
- Interrupt
  - From an external I/O controller
- Dealing with them without sacrificing performance is hard

# Handling Exceptions

- In MIPS, exceptions managed by a System Control Coprocessor (CP0)
- Save PC of offending (or interrupted) instruction
  - In MIPS: Exception Program Counter (EPC)
- Save indication of the problem
  - In MIPS: Cause register
  - We'll assume 1-bit
    - 0 for undefined opcode, 1 for overflow
- Jump to handler at 8000 00180



# An Alternate Mechanism

- Vectored Interrupts
  - Handler address determined by the cause
- Example:
  - Undefined opcode:    C000 0000
  - Overflow:                C000 0020
  - ....:                    C000 0040
- Instructions either
  - Deal with the interrupt, or
  - Jump to real handler



# Handler Actions

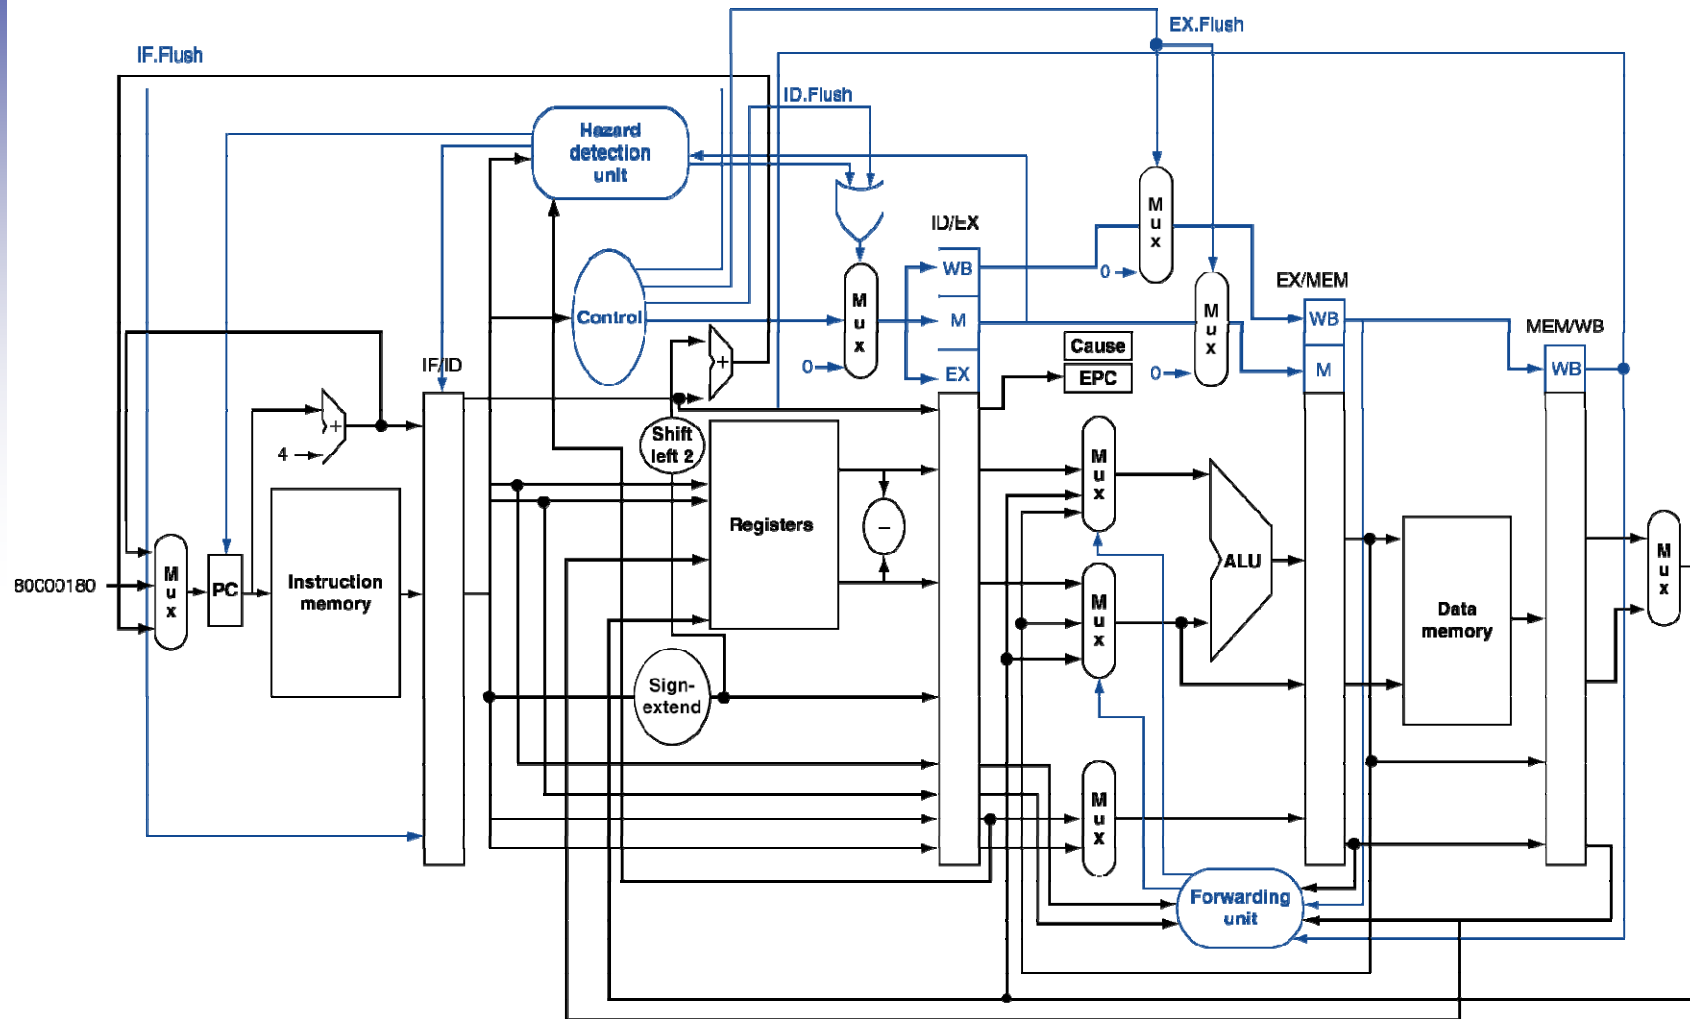
- Read cause, and transfer to relevant handler
- Determine action required
- If restartable
  - Take corrective action
  - use EPC to return to program
- Otherwise
  - Terminate program
  - Report error using EPC, cause, ...



# Exceptions in a Pipeline

- Another form of control hazard
- Consider overflow on add in EX stage
  - add \$1, \$2, \$1
  - Prevent \$1 from being clobbered
  - Complete previous instructions
  - Flush add and subsequent instructions
  - Set Cause and EPC register values
  - Transfer control to handler
- Similar to mispredicted branch
  - Use much of the same hardware

# Pipeline with Exceptions



# Exception Properties

- Restartable exceptions
  - Pipeline can flush the instruction
  - Handler executes, then returns to the instruction
    - Refetched and executed from scratch
- PC saved in EPC register
  - Identifies causing instruction
  - Actually PC + 4 is saved
    - Handler must adjust

# Exception Example

- Exception on `add` in

```

40      sub    $11, $2, $4
44      and    $12, $2, $5
48      or     $13, $2, $6
4C      add    $1,  $2, $1
50      slt   $15, $6, $7
54      lw    $16, 50($7)

```

...

- Handler

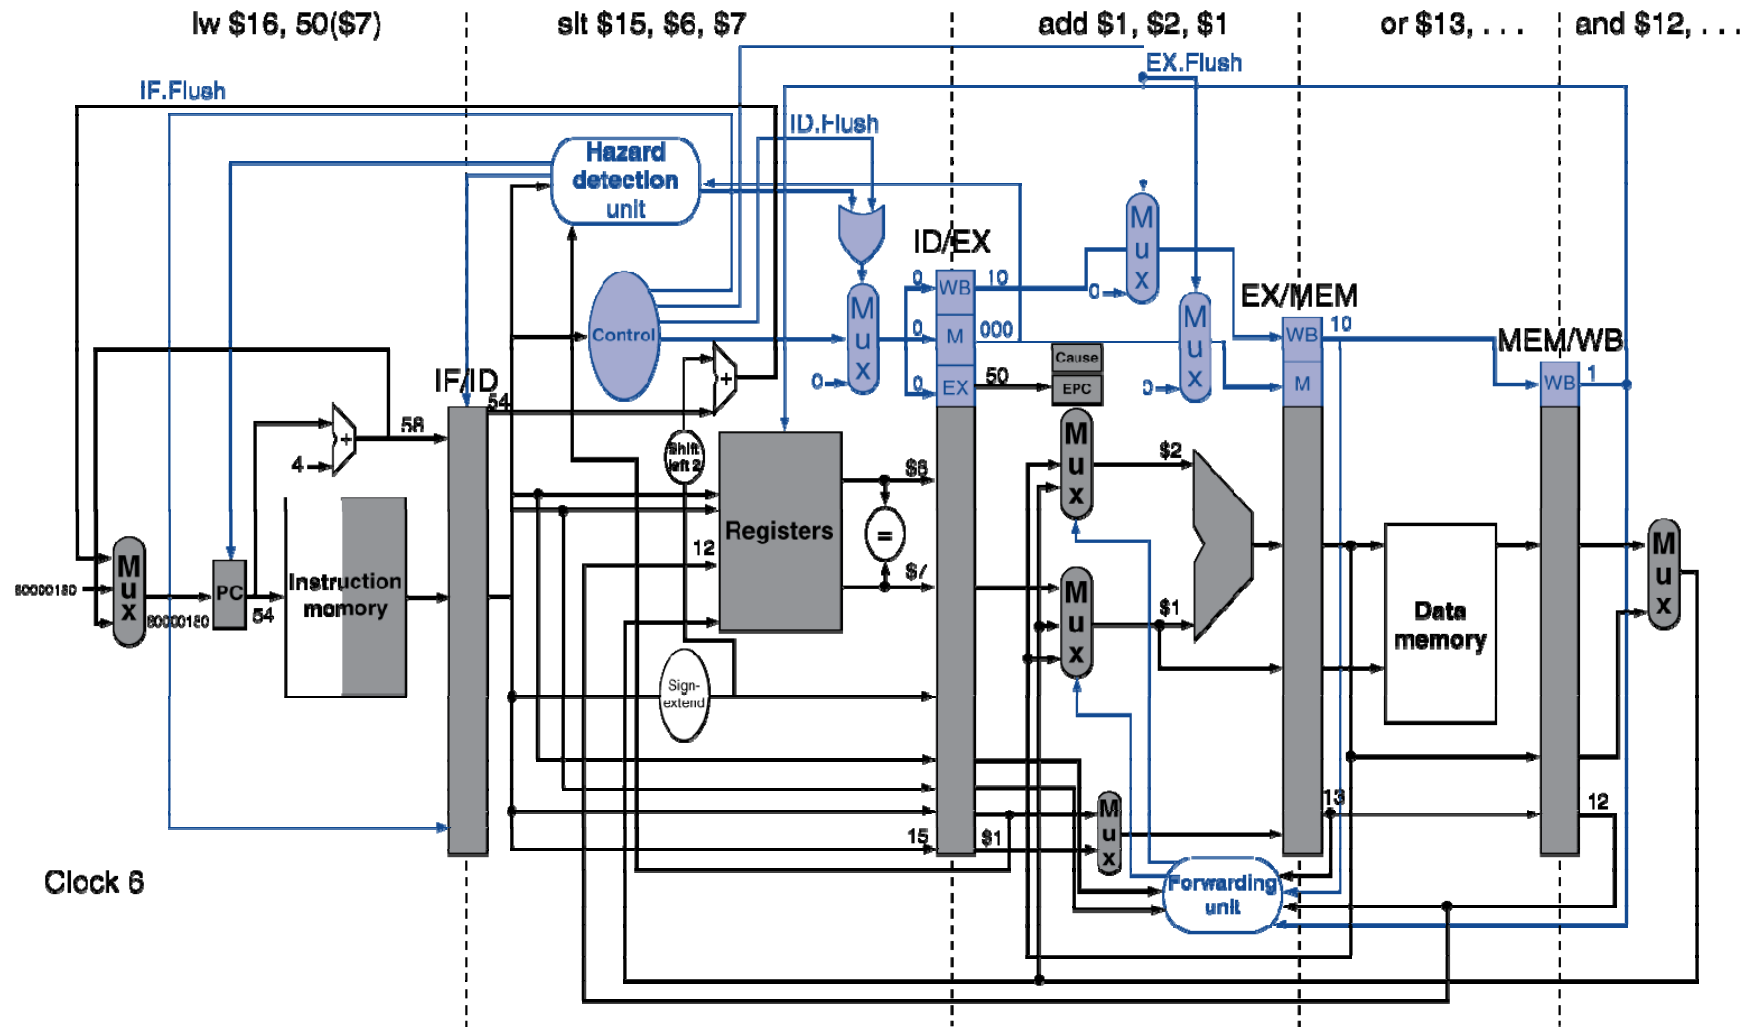
```

80000180    sw    $25, 1000($0)
80000184    sw    $26, 1004($0)

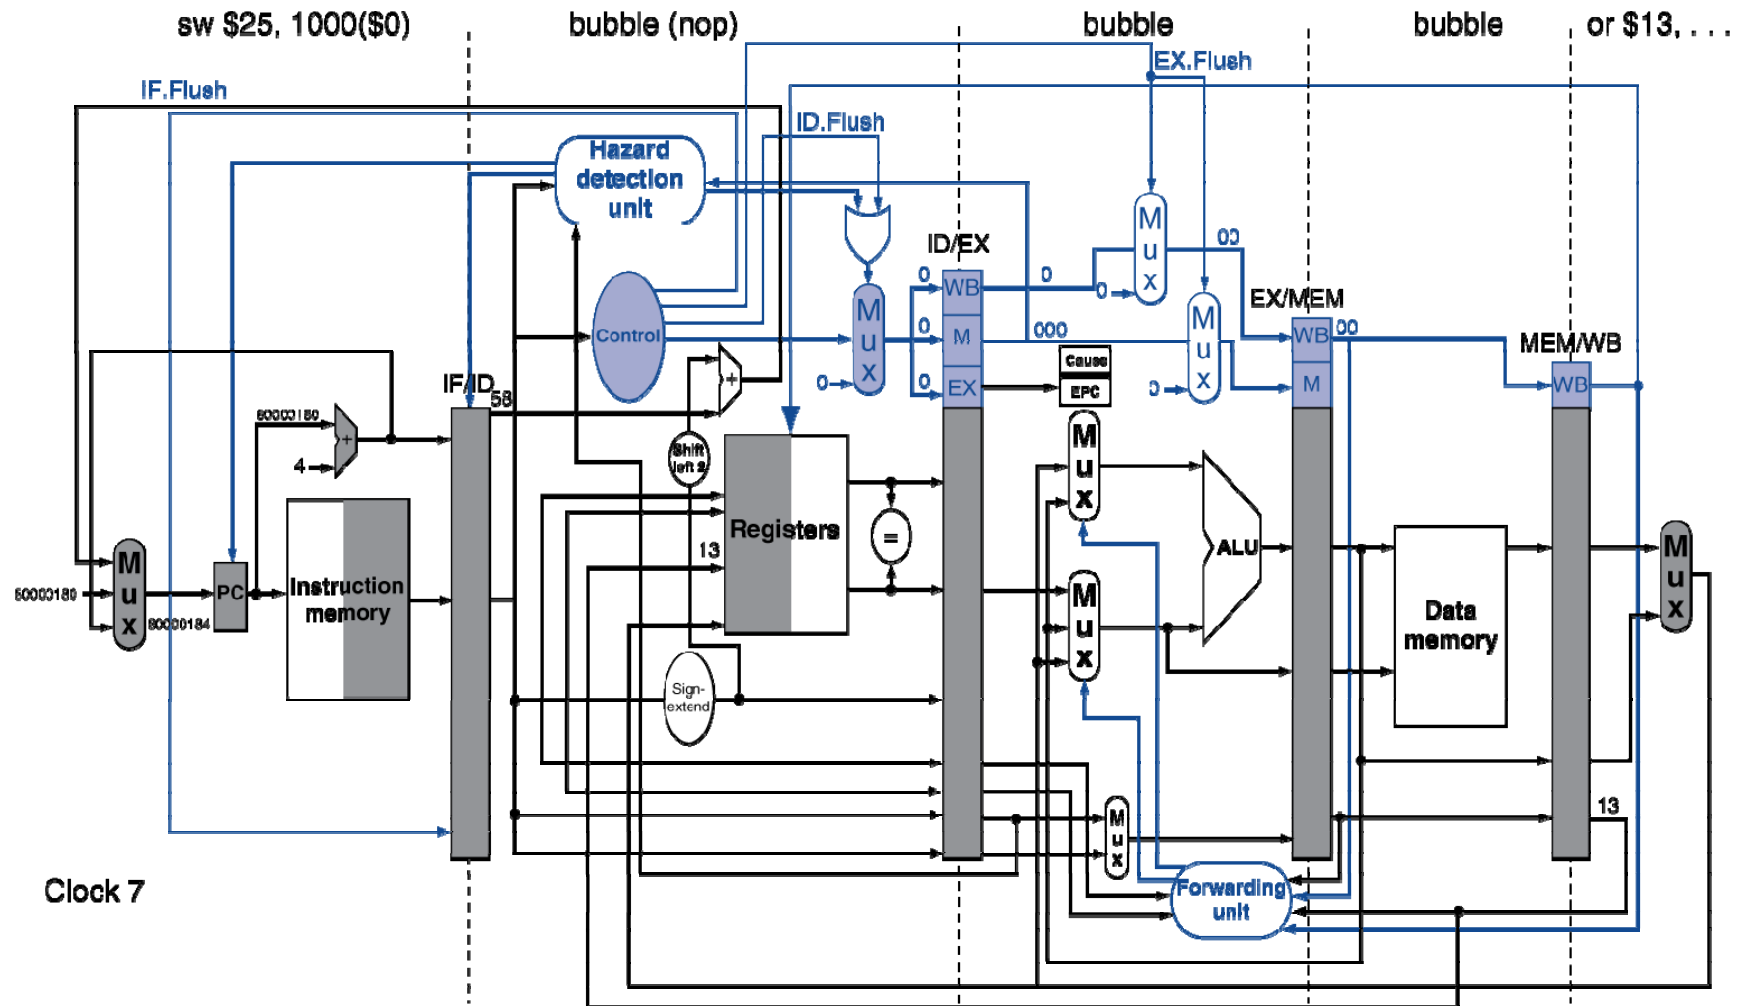
```

...

# Exception Example



# Exception Example



# Multiple Exceptions

- Pipelining overlaps multiple instructions
  - Could have multiple exceptions at once
- Simple approach: deal with exception from earliest instruction
  - Flush subsequent instructions
  - “Precise” exceptions
- In complex pipelines
  - Multiple instructions issued per cycle
  - Out-of-order completion
  - Maintaining precise exceptions is difficult!



# Imprecise Exceptions

- Just stop pipeline and save state
  - Including exception cause(s)
- Let the handler work out
  - Which instruction(s) had exceptions
  - Which to complete or flush
    - May require “manual” completion
- Simplifies hardware, but more complex handler software
- Not feasible for complex multiple-issue out-of-order pipelines

# Instruction-Level Parallelism (ILP)

- Pipelining: executing multiple instructions in parallel
- To increase ILP
  - Deeper pipeline
    - Less work per stage  $\Rightarrow$  shorter clock cycle
  - Multiple issue
    - Replicate pipeline stages  $\Rightarrow$  multiple pipelines
    - Start multiple instructions per clock cycle
    - $CPI < 1$ , so use Instructions Per Cycle (IPC)
    - E.g., 4GHz 4-way multiple-issue
      - 16 BIPS, peak  $CPI = 0.25$ , peak  $IPC = 4$
    - But dependencies reduce this in practice

# Multiple Issue

- Static multiple issue
  - Compiler groups instructions to be issued together
  - Packages them into “issue slots”
  - Compiler detects and avoids hazards
- Dynamic multiple issue
  - CPU examines instruction stream and chooses instructions to issue each cycle
  - Compiler can help by reordering instructions
  - CPU resolves hazards using advanced techniques at runtime

# Speculation

- “Guess” what to do with an instruction
  - Start operation as soon as possible
  - Check whether guess was right
    - If so, complete the operation
    - If not, roll-back and do the right thing
- Common to static and dynamic multiple issue
- Examples
  - Speculate on branch outcome
    - Roll back if path taken is different
  - Speculate on load
    - Roll back if location is updated

# Compiler/Hardware Speculation

- Compiler can reorder instructions
  - e.g., move load before branch
  - Can include “fix-up” instructions to recover from incorrect guess
- Hardware can look ahead for instructions to execute
  - Buffer results until it determines they are actually needed
  - Flush buffers on incorrect speculation



# Speculation and Exceptions

- What if exception occurs on a speculatively executed instruction?
  - e.g., speculative load before null-pointer check
- Static speculation
  - Can add ISA support for deferring exceptions
- Dynamic speculation
  - Can buffer exceptions until instruction completion (which may not occur)

# Static Multiple Issue

- Compiler groups instructions into “issue packets”
  - Group of instructions that can be issued on a single cycle
  - Determined by pipeline resources required
- Think of an issue packet as a very long instruction
  - Specifies multiple concurrent operations
  - ⇒ Very Long Instruction Word (VLIW)

# Scheduling Static Multiple Issue

- Compiler must remove some/all hazards
  - Reorder instructions into issue packets
  - No dependencies with a packet
  - Possibly some dependencies between packets
    - Varies between ISAs; compiler must know!
  - Pad with nop if necessary



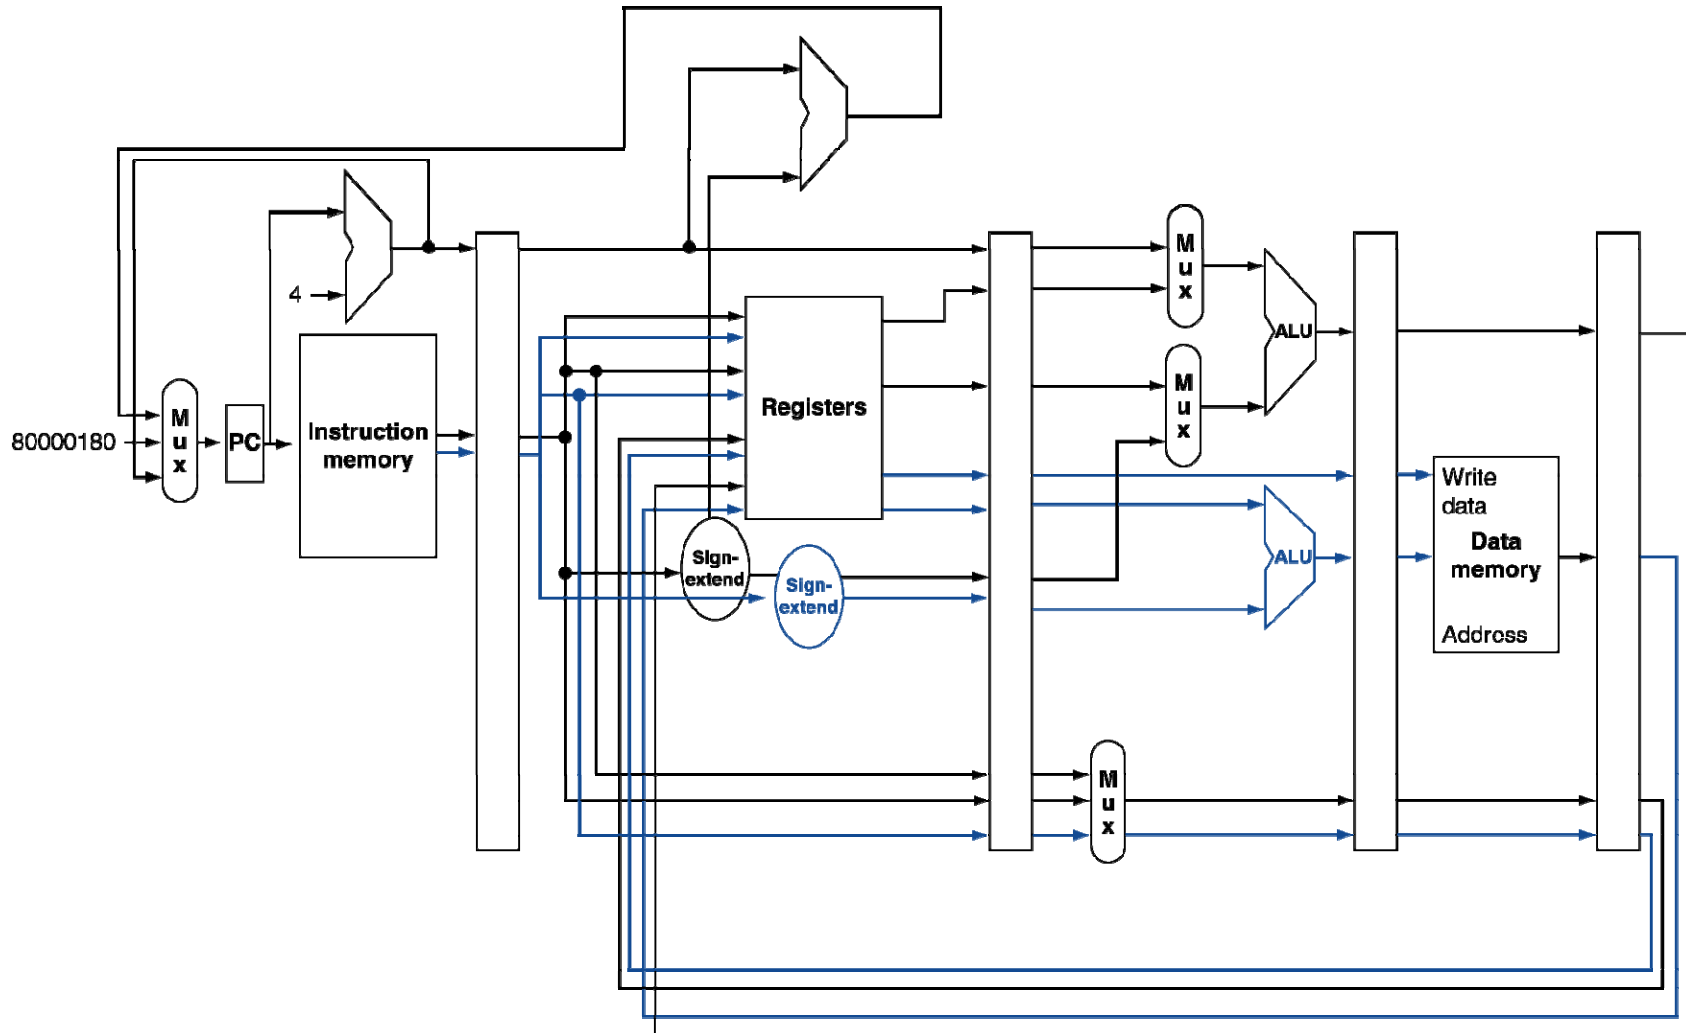


# MIPS with Static Dual Issue

- Two-issue packets
  - One ALU/branch instruction
  - One load/store instruction
  - 64-bit aligned
    - ALU/branch, then load/store
    - Pad an unused instruction with nop

Address	Instruction type	Pipeline Stages						
		IF	ID	EX	MEM	WB		
n	ALU/branch	IF	ID	EX	MEM	WB		
n + 4	Load/store	IF	ID	EX	MEM	WB		
n + 8	ALU/branch		IF	ID	EX	MEM	WB	
n + 12	Load/store		IF	ID	EX	MEM	WB	
n + 16	ALU/branch			IF	ID	EX	MEM	WB
n + 20	Load/store			IF	ID	EX	MEM	WB

# MIPS with Static Dual Issue



# Hazards in the Dual-Issue MIPS

- More instructions executing in parallel
- EX data hazard
  - Forwarding avoided stalls with single-issue
  - Now can't use ALU result in load/store in same packet
    - `add $t0, $s0, $s1`
    - `load $s2, 0($t0)`
    - Split into two packets, effectively a stall
- Load-use hazard
  - Still one cycle use latency, but now two instructions
- More aggressive scheduling required

# Scheduling Example

- Schedule this for dual-issue MIPS

```

Loop: lw    $t0, 0($s1)      # $t0=array element
      addu  $t0, $t0, $s2    # add scalar in $s2
      sw    $t0, 0($s1)     # store result
      addi  $s1, $s1, -4    # decrement pointer
      bne   $s1, $zero, Loop # branch $s1!=0
  
```

	ALU/branch	Load/store	cycle
Loop:	nop	lw \$t0, 0(\$s1)	1
	addi \$s1, \$s1, -4	nop	2
	addu \$t0, \$t0, \$s2	nop	3
	bne \$s1, \$zero, Loop	sw \$t0, 4(\$s1)	4

- $IPC = 5/4 = 1.25$  (c.f. peak  $IPC = 2$ )

# Loop Unrolling

- Replicate loop body to expose more parallelism
  - Reduces loop-control overhead
- Use different registers per replication
  - Called “register renaming”
  - Avoid loop-carried “anti-dependencies”
    - Store followed by a load of the same register
    - Aka “name dependence”
      - Reuse of a register name

# Loop Unrolling Example

	ALU/branch	Load/store	cycle
Loop:	addi \$s1, \$s1, -16	lw \$t0, 0(\$s1)	1
	nop	lw \$t1, 12(\$s1)	2
	addu \$t0, \$t0, \$s2	lw \$t2, 8(\$s1)	3
	addu \$t1, \$t1, \$s2	lw \$t3, 4(\$s1)	4
	addu \$t2, \$t2, \$s2	sw \$t0, 16(\$s1)	5
	addu \$t3, \$t4, \$s2	sw \$t1, 12(\$s1)	6
	nop	sw \$t2, 8(\$s1)	7
	bne \$s1, \$zero, Loop	sw \$t3, 4(\$s1)	8

- $IPC = 14/8 = 1.75$ 
  - Closer to 2, but at cost of registers and code size

# Dynamic Multiple Issue

- “Superscalar” processors
- CPU decides whether to issue 0, 1, 2, ... each cycle
  - Avoiding structural and data hazards
- Avoids the need for compiler scheduling
  - Though it may still help
  - Code semantics ensured by the CPU

# Dynamic Pipeline Scheduling

- Allow the CPU to execute instructions out of order to avoid stalls
  - But commit result to registers in order

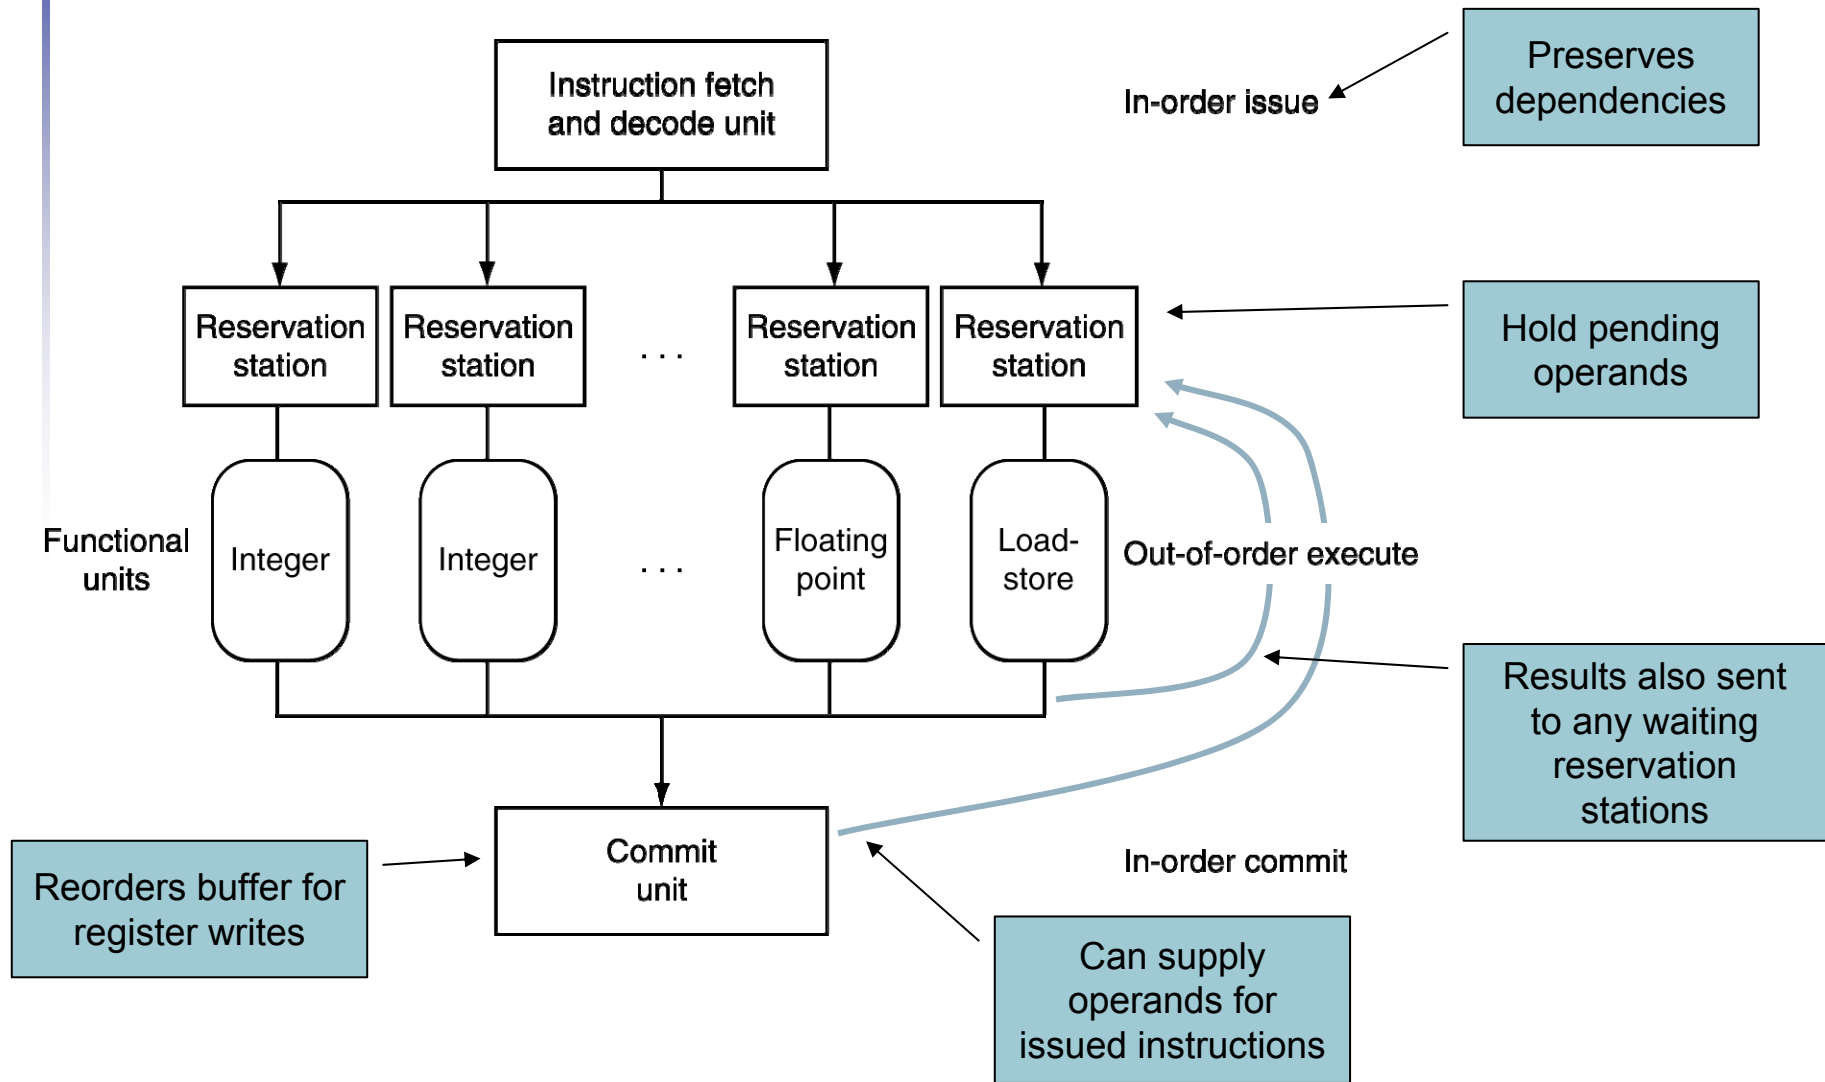
- Example

```
lw      $t0, 20($s2)
addu   $t1, $t0, $t2
sub     $s4, $s4, $t3
slti   $t5, $s4, 20
```

- Can start sub while addu is waiting for lw



# Dynamically Scheduled CPU



# Register Renaming

- Reservation stations and reorder buffer effectively provide register renaming
- On instruction issue to reservation station
  - If operand is available in register file or reorder buffer
    - Copied to reservation station
    - No longer required in the register; can be overwritten
  - If operand is not yet available
    - It will be provided to the reservation station by a function unit
    - Register update may not be required

# Speculation

- Predict branch and continue issuing
  - Don't commit until branch outcome determined
- Load speculation
  - Avoid load and cache miss delay
    - Predict the effective address
    - Predict loaded value
    - Load before completing outstanding stores
    - Bypass stored values to load unit
  - Don't commit load until speculation cleared

# Why Do Dynamic Scheduling?

- Why not just let the compiler schedule code?
- Not all stalls are predicable
  - e.g., cache misses
- Can't always schedule around branches
  - Branch outcome is dynamically determined
- Different implementations of an ISA have different latencies and hazards

# Does Multiple Issue Work?

## The BIG Picture

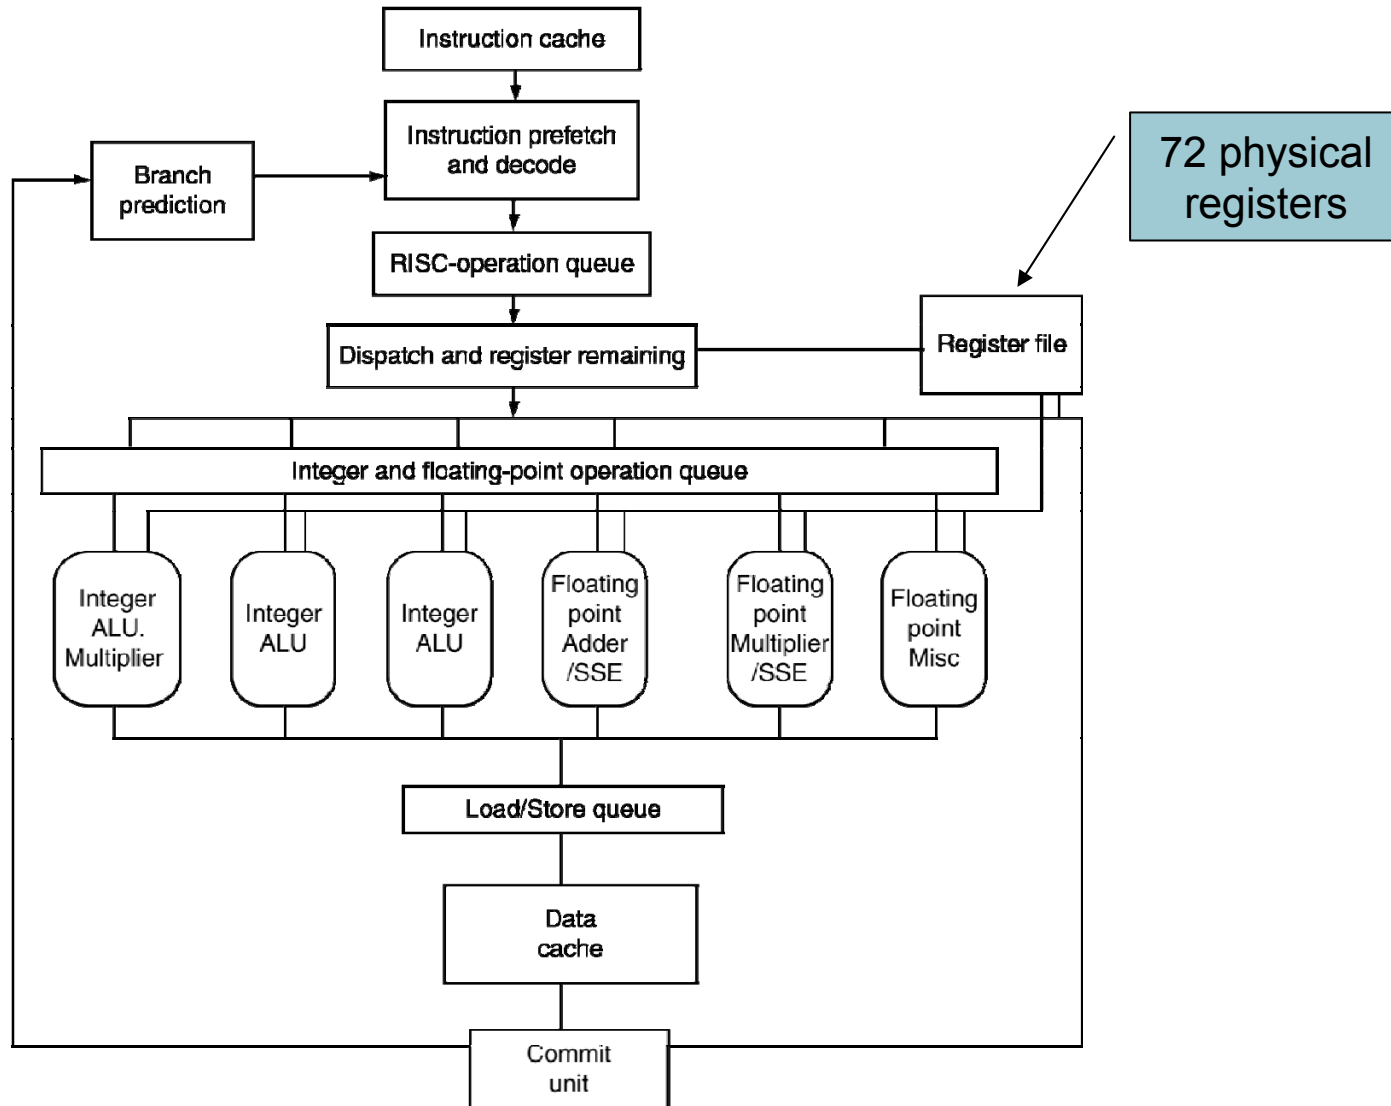
- Yes, but not as much as we'd like
- Programs have real dependencies that limit ILP
- Some dependencies are hard to eliminate
  - e.g., pointer aliasing
- Some parallelism is hard to expose
  - Limited window size during instruction issue
- Memory delays and limited bandwidth
  - Hard to keep pipelines full
- Speculation can help if done well

# Power Efficiency

- Complexity of dynamic scheduling and speculations requires power
- Multiple simpler cores may be better

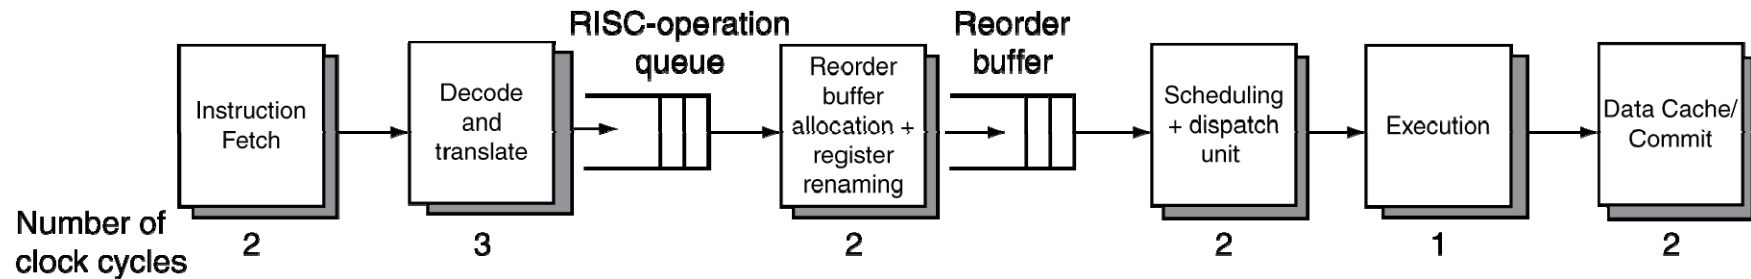
Microprocessor	Year	Clock Rate	Pipeline Stages	Issue width	Out-of-order/ Speculation	Cores	Power
i486	1989	25MHz	5	1	No	1	5W
Pentium	1993	66MHz	5	2	No	1	10W
Pentium Pro	1997	200MHz	10	3	Yes	1	29W
P4 Willamette	2001	2000MHz	22	3	Yes	1	75W
P4 Prescott	2004	3600MHz	31	3	Yes	1	103W
Core	2006	2930MHz	14	4	Yes	2	75W
UltraSparc III	2003	1950MHz	14	4	No	1	90W
UltraSparc T1	2005	1200MHz	6	1	No	8	70W

# The Opteron X4 Microarchitecture



# The Opteron X4 Pipeline Flow

- For integer operations



- FP is 5 stages longer
- Up to 106 RISC-ops in progress
- Bottlenecks
  - Complex instructions with long dependencies
  - Branch mispredictions
  - Memory access delays



# Fallacies

- Pipelining is easy (!)
  - The basic idea is easy
  - The devil is in the details
    - e.g., detecting data hazards
- Pipelining is independent of technology
  - So why haven't we always done pipelining?
  - More transistors make more advanced techniques feasible
  - Pipeline-related ISA design needs to take account of technology trends
    - e.g., predicated instructions

# Pitfalls

- Poor ISA design can make pipelining harder
  - e.g., complex instruction sets (VAX, IA-32)
    - Significant overhead to make pipelining work
    - IA-32 micro-op approach
  - e.g., complex addressing modes
    - Register update side effects, memory indirection
  - e.g., delayed branches
    - Advanced pipelines have long delay slots

# Concluding Remarks

- ISA influences design of datapath and control
- Datapath and control influence design of ISA
- Pipelining improves instruction throughput using parallelism
  - More instructions completed per second
  - Latency for each instruction not reduced
- Hazards: structural, data, control
- Multiple issue and dynamic scheduling (ILP)
  - Dependencies limit achievable parallelism
  - Complexity leads to the power wall