

18-447 Lecture 8: Data Hazard and Resolution

James C. Hoe

Department of ECE

Carnegie Mellon University

Housekeeping

- Your goal today
 - detect and resolve data hazards in in-order instruction pipelines
- Notices
 - Lab 2, **status check next week, due week of 2/24**
 - HW 2, **due 2/19 before class**
- Readings
 - P&H Ch 4

Instruction Pipeline Reality

- **Not identical tasks**
 - coalescing instruction types into one “multi-function” pipe
 - external fragmentation (some idle stages)
- **Not uniform suboperations**
 - group or sub-divide steps into stages to minimize variance
 - internal fragmentation (some too-fast stages)
- **Not independent tasks**
 - dependency detection and resolution
 - next lecture(s)

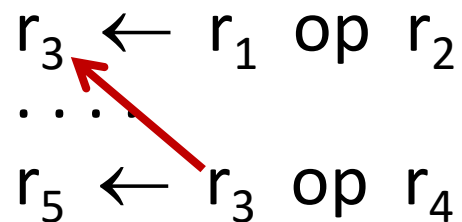
Review



Even more messy if not RISC

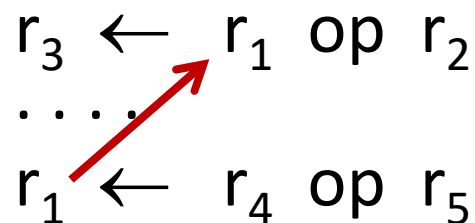
Data Dependence

Data dependence

$$\begin{array}{l} r_3 \leftarrow r_1 \text{ op } r_2 \\ \dots \\ r_5 \leftarrow r_3 \text{ op } r_4 \end{array}$$


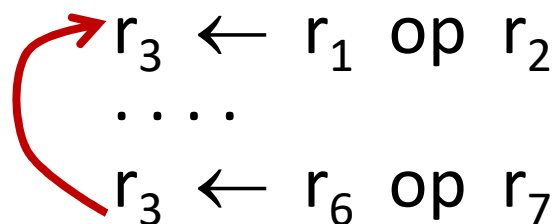
Read-after-Write (RAW)

Anti-dependence

$$\begin{array}{l} r_3 \leftarrow r_1 \text{ op } r_2 \\ \dots \\ r_1 \leftarrow r_4 \text{ op } r_5 \end{array}$$


Write-after-Read (WAR)

Output-dependence

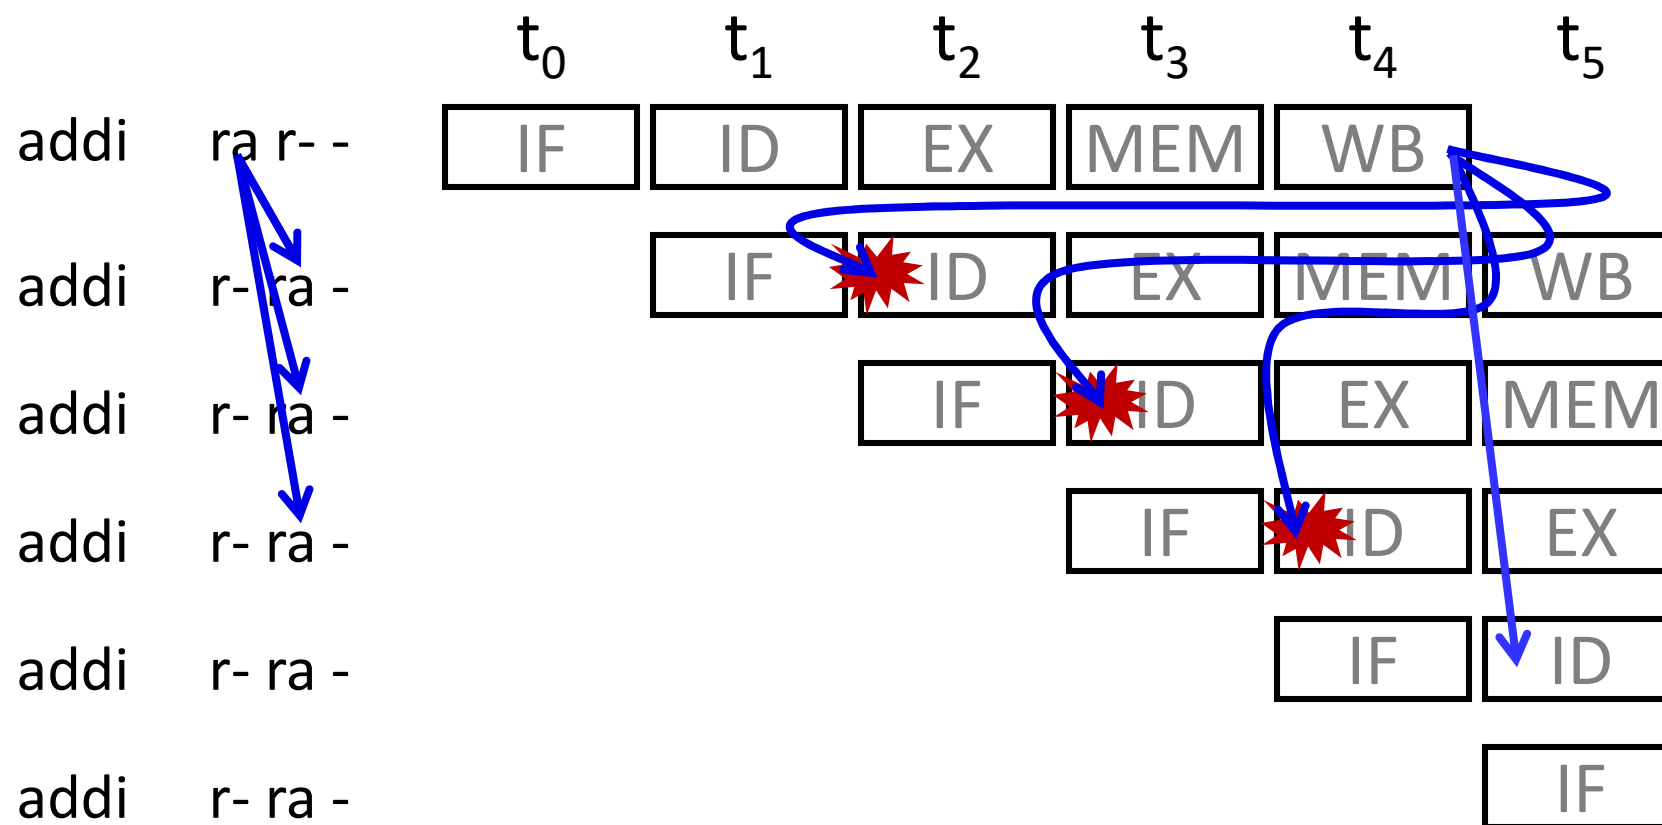
$$\begin{array}{l} r_3 \leftarrow r_1 \text{ op } r_2 \\ \dots \\ r_3 \leftarrow r_6 \text{ op } r_7 \end{array}$$


Write-after-Write (WAW)

Don't forget memory instructions

Review

RAW Dependency and Hazard

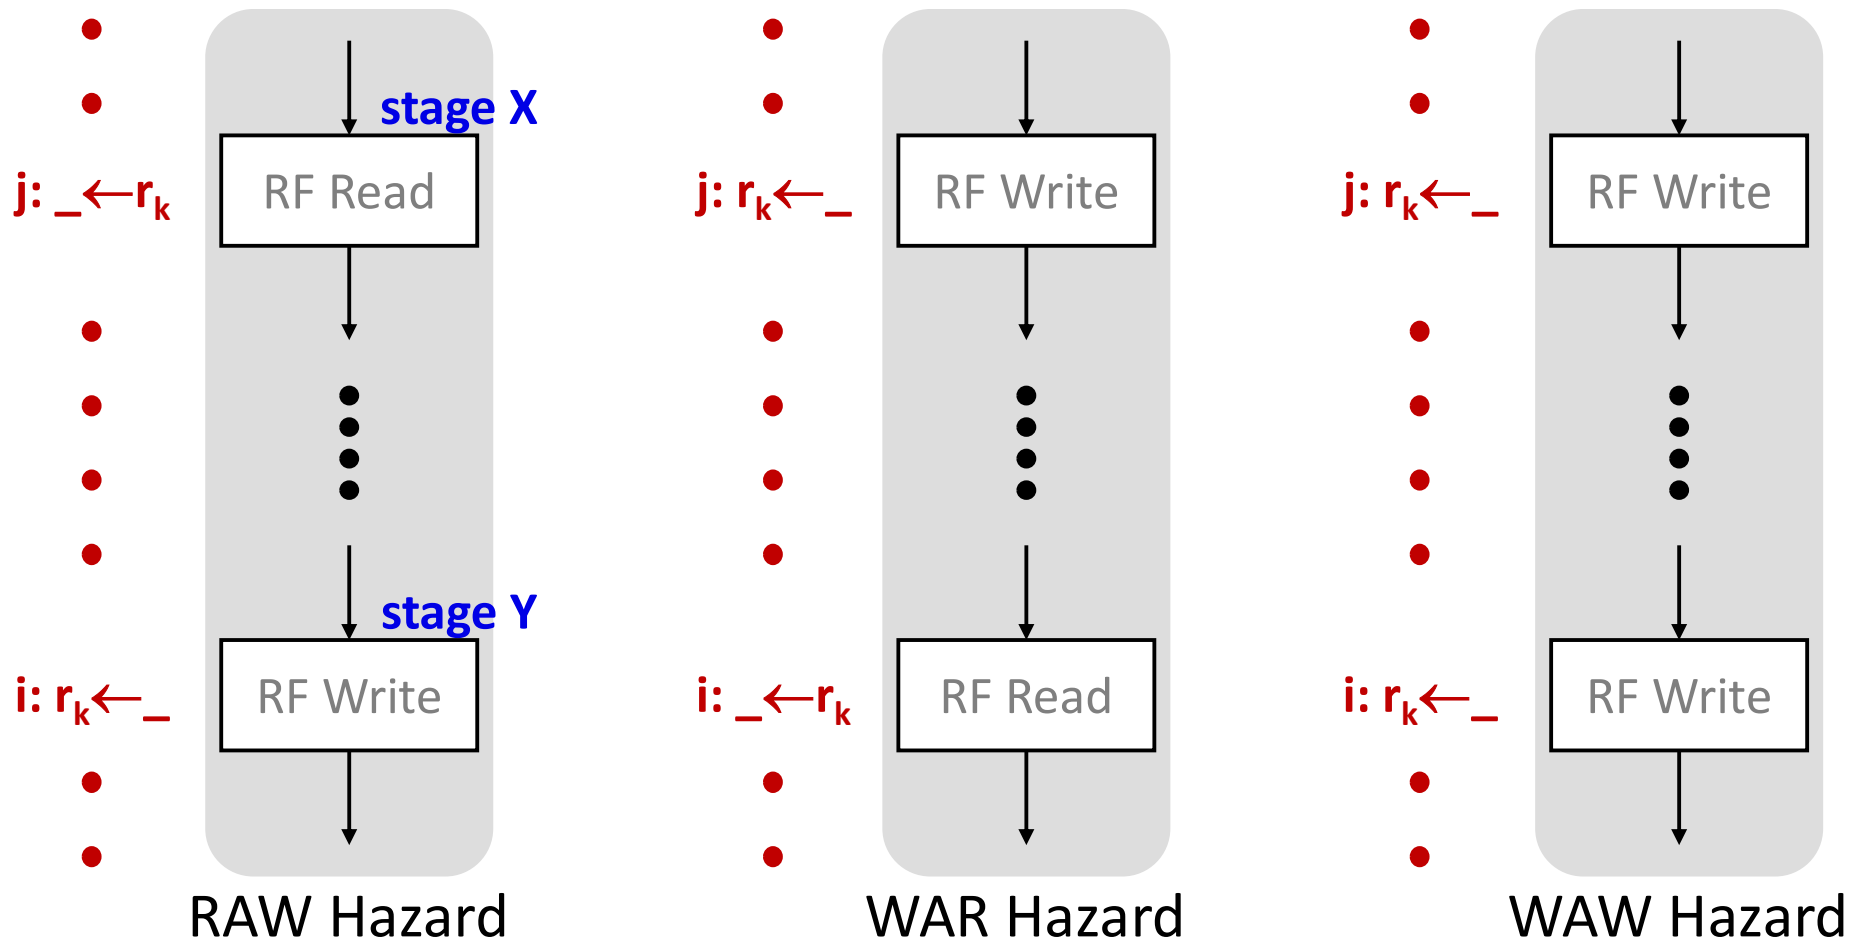


Register Data Hazard Analysis

	R/I-Type	LW	SW	Bxx	Jal	Jalr
IF						
ID	read RF	read RF	read RF	read RF		read RF
EX						
MEM						
WB	write RF	write RF			write RF	write RF

- For a given pipeline, when is there a register data hazard between 2 dependent instructions?
 - dependence type: RAW, WAR, WAW?
 - instruction types involved?
 - distance between the two instructions?

Hazard in In-order Pipeline



$\text{dist}_{\text{dependence}}(i, j) \leq \text{dist}_{\text{hazard}}(X, Y) \Rightarrow \text{Hazard!!}$

$\text{dist}_{\text{dependence}}(i, j) > \text{dist}_{\text{hazard}}(X, Y) \Rightarrow \text{Safe}$

RAW Hazard Analysis Example

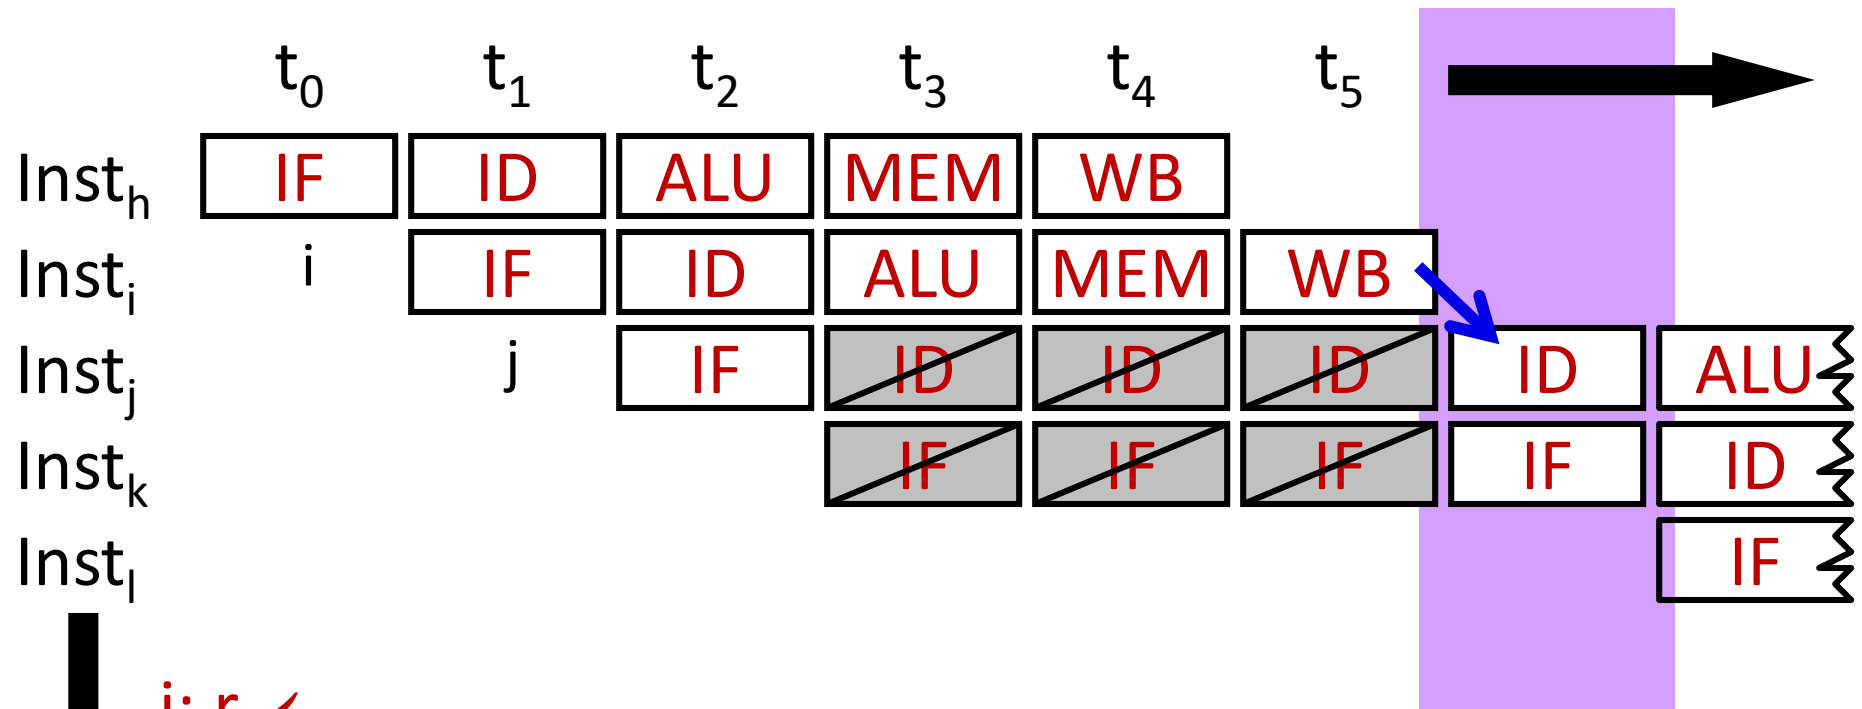
	R/I-Type	LW	SW	Bxx	Jal	Jalr
IF						
ID	read RF	read RF	read RF	read RF		read RF
EX						
MEM						
WB	write RF	write RF			write RF	write RF

- Older I_A and younger I_B have RAW hazard iff
 - I_B (R/I, LW, SW, Bxx or JALR) reads a register written by I_A (R/I, LW, or JAL/R)
 - $\text{dist}(I_A, I_B) \leq \text{dist}(\text{ID}, \text{WB}) = 3$

What about WAW and WAR hazard?

What about memory data hazard?

Pipeline Stall: universal hazard resolution



$i: r_x \leftarrow _$
 bubble
 bubble
 bubble
 $j: _ \leftarrow r_x$

$$\text{dist}(i,j)=4$$

Stall==make younger instruction
wait until hazard passes

1. stop all up-stream stages
2. drain all down-stream stages

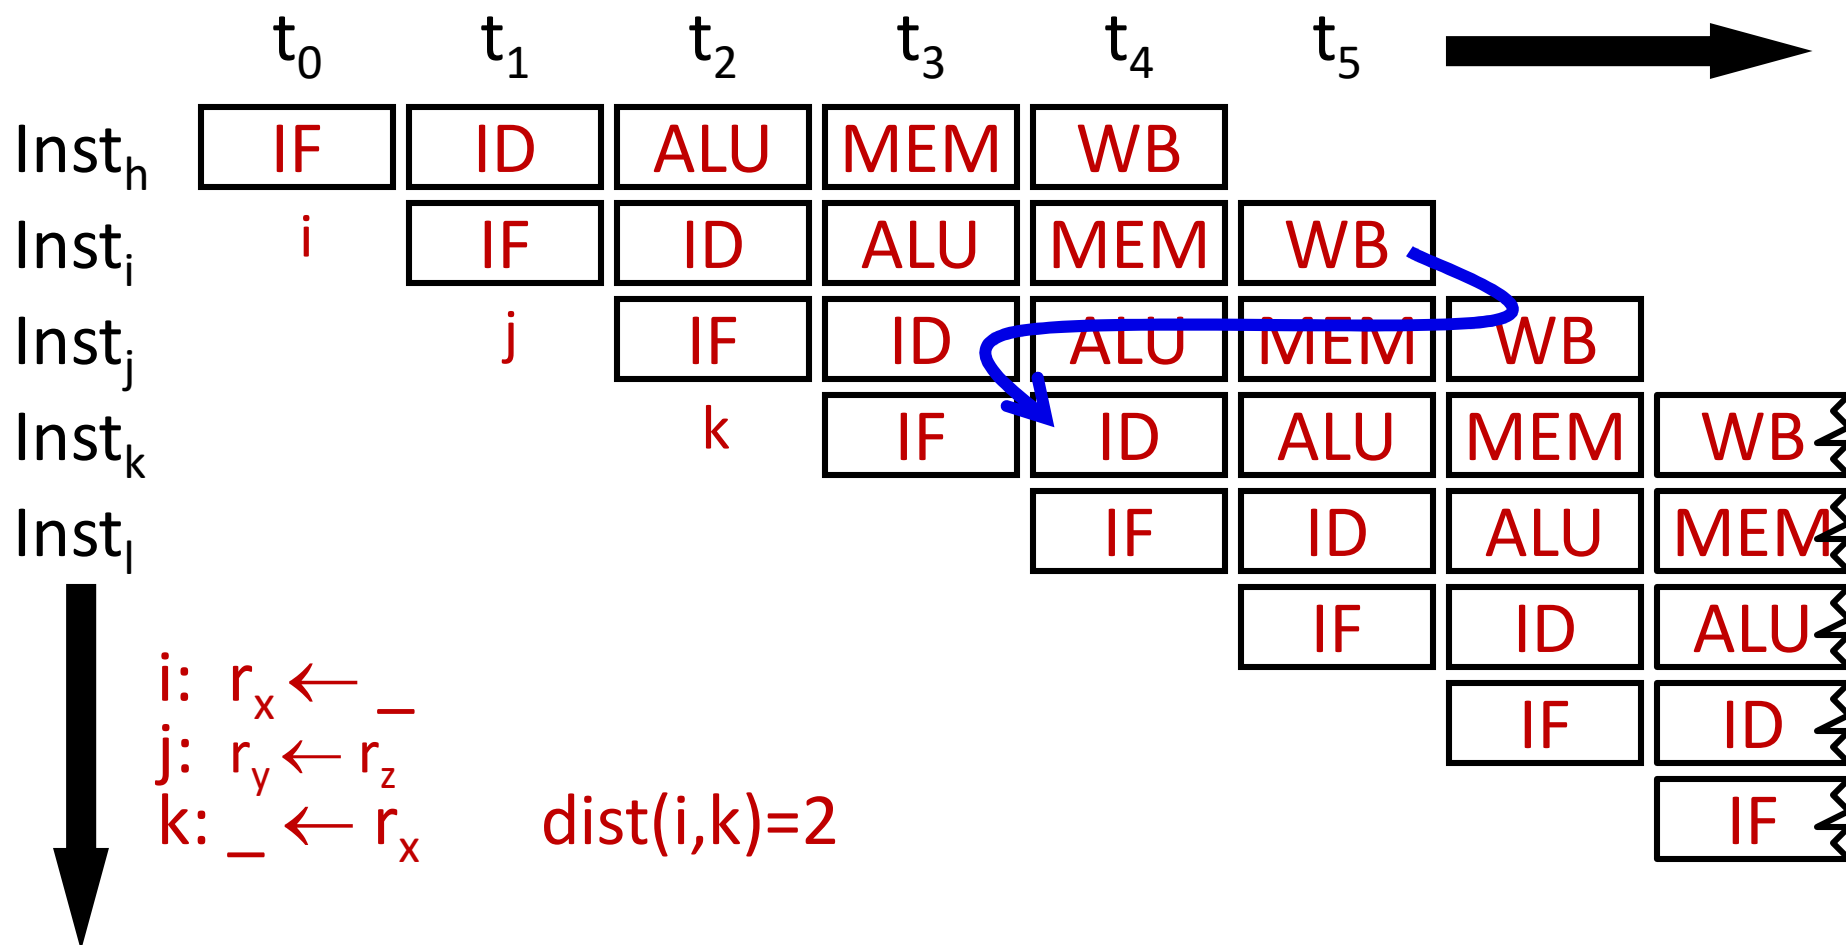
Pipeline Stall

	t ₀	t ₁	t ₂	t ₃	t ₄	t ₅	t ₆	t ₇	t ₈	t ₉	t ₁₀
IF	i	j	k	k	k	k	l				
ID	h	i	j	j	j	j	k	l			
EX		h	i	bub	bub	bub	j	k	l		
MEM			h	i	bub	bub	bub	j	k	l	
WB				h	i	bub	bub	bub	j	k	l

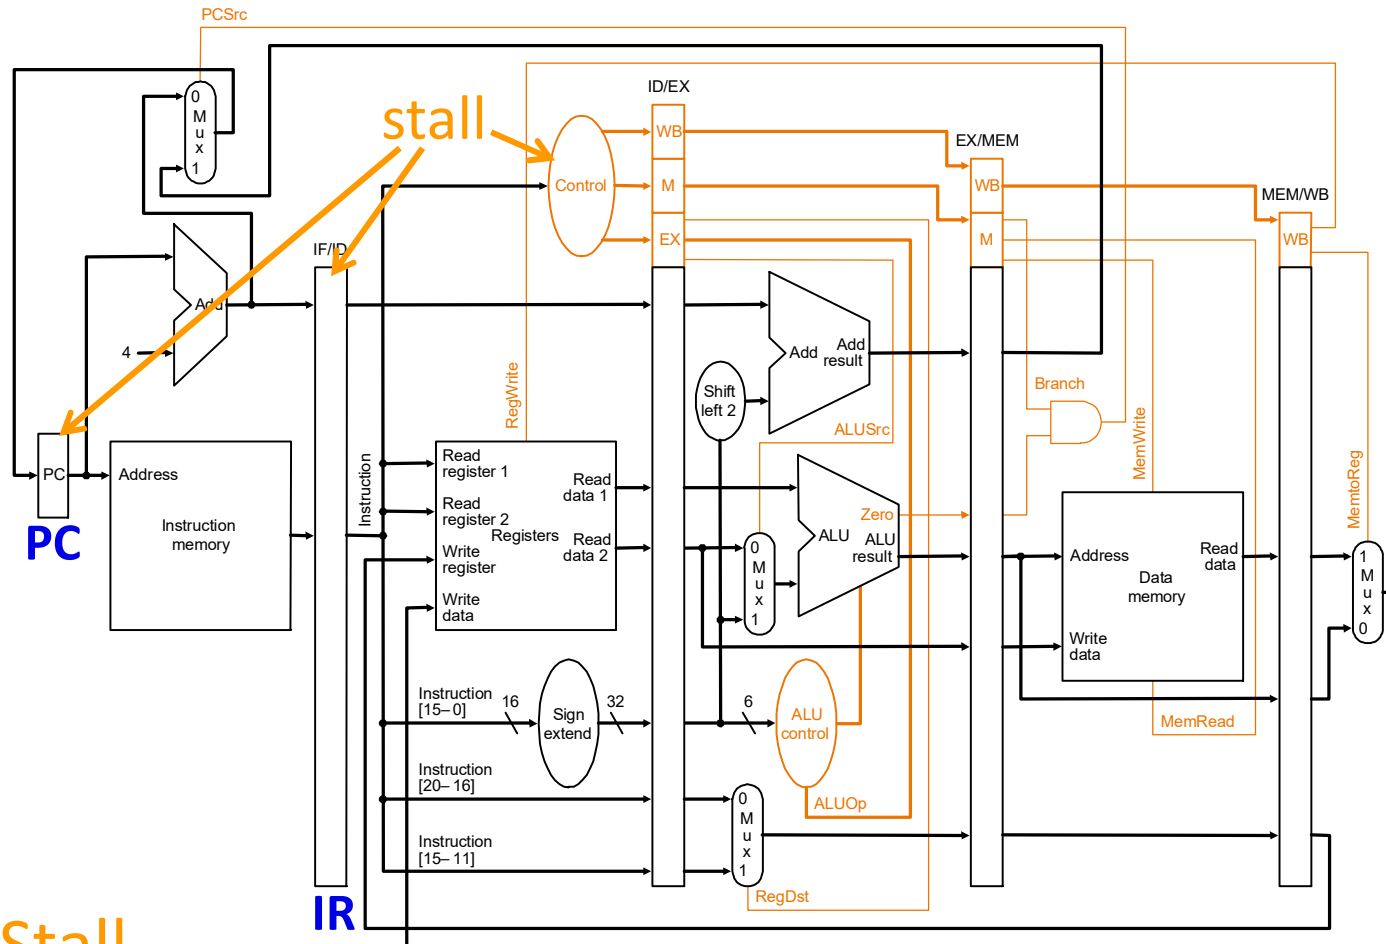
i: rx ← _

j: _ ← rx

Pop Quiz: What happens in this case?



Stall



- Stall

- disable **PC** and **IR** latching
- set $\text{RegWrite}_{ID} = 0$ and $\text{MemWrite}_{ID} = 0$

Stall Condition

- Older I_A and younger I_B have RAW hazard iff
 - I_B (R/I, LW, SW, Bxx or JALR) reads a register written by I_A (R/I, LW, or JAL/R)
 - $\text{dist}(I_A, I_B) \leq \text{dist}(\text{ID}, \text{WB}) = 3$
- More plainly, before I_B in ID reads a register, I_B needs to check if any I_A in EX, MEM or WB is going to update it (if so, value in RF is “stale”)

Watch out for x0!!

Stall Condition

- Helper functions
 - $use_rs1(I)$ returns true if I uses $rs1$ && $rs1 \neq x0$
- Stall IF and ID when
 - $(rs1_{ID} == rd_{EX}) \ \&\& \ use_rs1(IR_{ID}) \ \&\& \ RegWrite_{EX}$ or
 - $(rs1_{ID} == rd_{MEM}) \ \&\& \ use_rs1(IR_{ID}) \ \&\& \ RegWrite_{MEM}$ or
 - $(rs1_{ID} == rd_{WB}) \ \&\& \ use_rs1(IR_{ID}) \ \&\& \ RegWrite_{WB}$ or
 - $(rs2_{ID} == rd_{EX}) \ \&\& \ use_rs2(IR_{ID}) \ \&\& \ RegWrite_{EX}$ or
 - $(rs2_{ID} == rd_{MEM}) \ \&\& \ use_rs2(IR_{ID}) \ \&\& \ RegWrite_{MEM}$ or
 - $(rs2_{ID} == rd_{WB}) \ \&\& \ use_rs2(IR_{ID}) \ \&\& \ RegWrite_{WB}$

It is crucial that EX, MEM and WB continue to advance during stall

Impact of Stall on Performance

- Each stall cycle corresponds to 1 lost ALU cycle
- A program with **N** instructions and **S** stall cycles:

$$\text{average IPC} = \frac{N}{N+S}$$

- **S** depends on
 - frequency of hazard-causing dependencies
 - distance between hazard-causing instruction pairs
 - distance between hazard-causing dependencies

(suppose i_1, i_2 and i_3 all depend on i_0 , once i_1 's hazard is resolved by stalling, i_2 and i_3 do not stall)

Sample Assembly [P&H]


for (j=i-1; j>=0 && v[j] > v[j+1]; j-=1) { }

	addi	\$s1, \$s0, -1	3 stalls
for2tst:	slti	\$t0, \$s1, 0	3 stalls
	bne	\$t0, \$zero, exit2	
	sll	\$t1, \$s1, 2	3 stalls
	add	\$t2, \$a0, \$t1	3 stalls
	lw	\$t3, 0(\$t2)	
	lw	\$t4, 4(\$t2)	3 stalls
	slt	\$t0, \$t4, \$t3	3 stalls
	beq	\$t0, \$zero, exit2	
		
	addi	\$s1, \$s1, -1	
	j	for2tst	

exit2:

Data Forwarding (or Register Bypassing)

- What does “**ADD rx ry rz**” mean? Get inputs from **RF[ry]** and **RF[rz]** and put result in **RF[rx]**?
- But, **RF** is just a part of an abstraction
 - a way to connect dataflow between instructions
 - “inputs to **ADD** are resulting values of the last instructions to assign to **RF[ry]** and **RF[rz]**”
 - **RF** doesn't have to exist as a literal object
- If only dataflow matters, don't wait for WB . . .

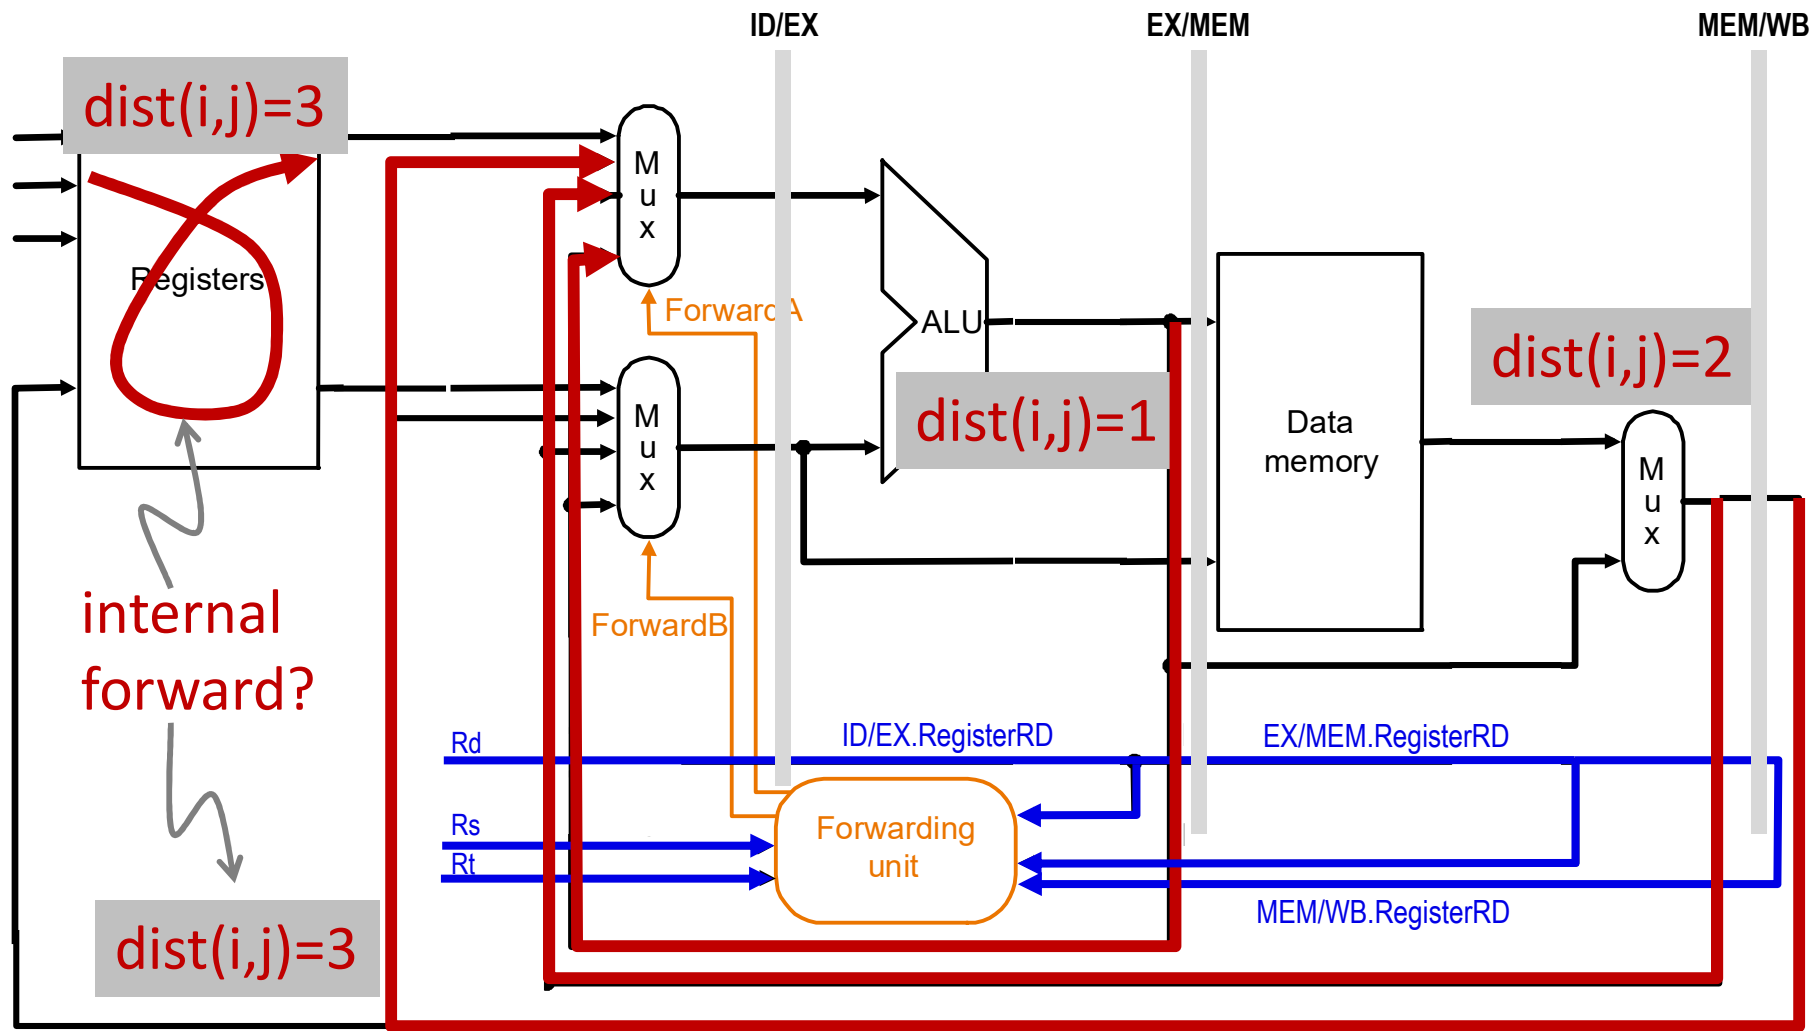


Resolving RAW Hazard by Forwarding

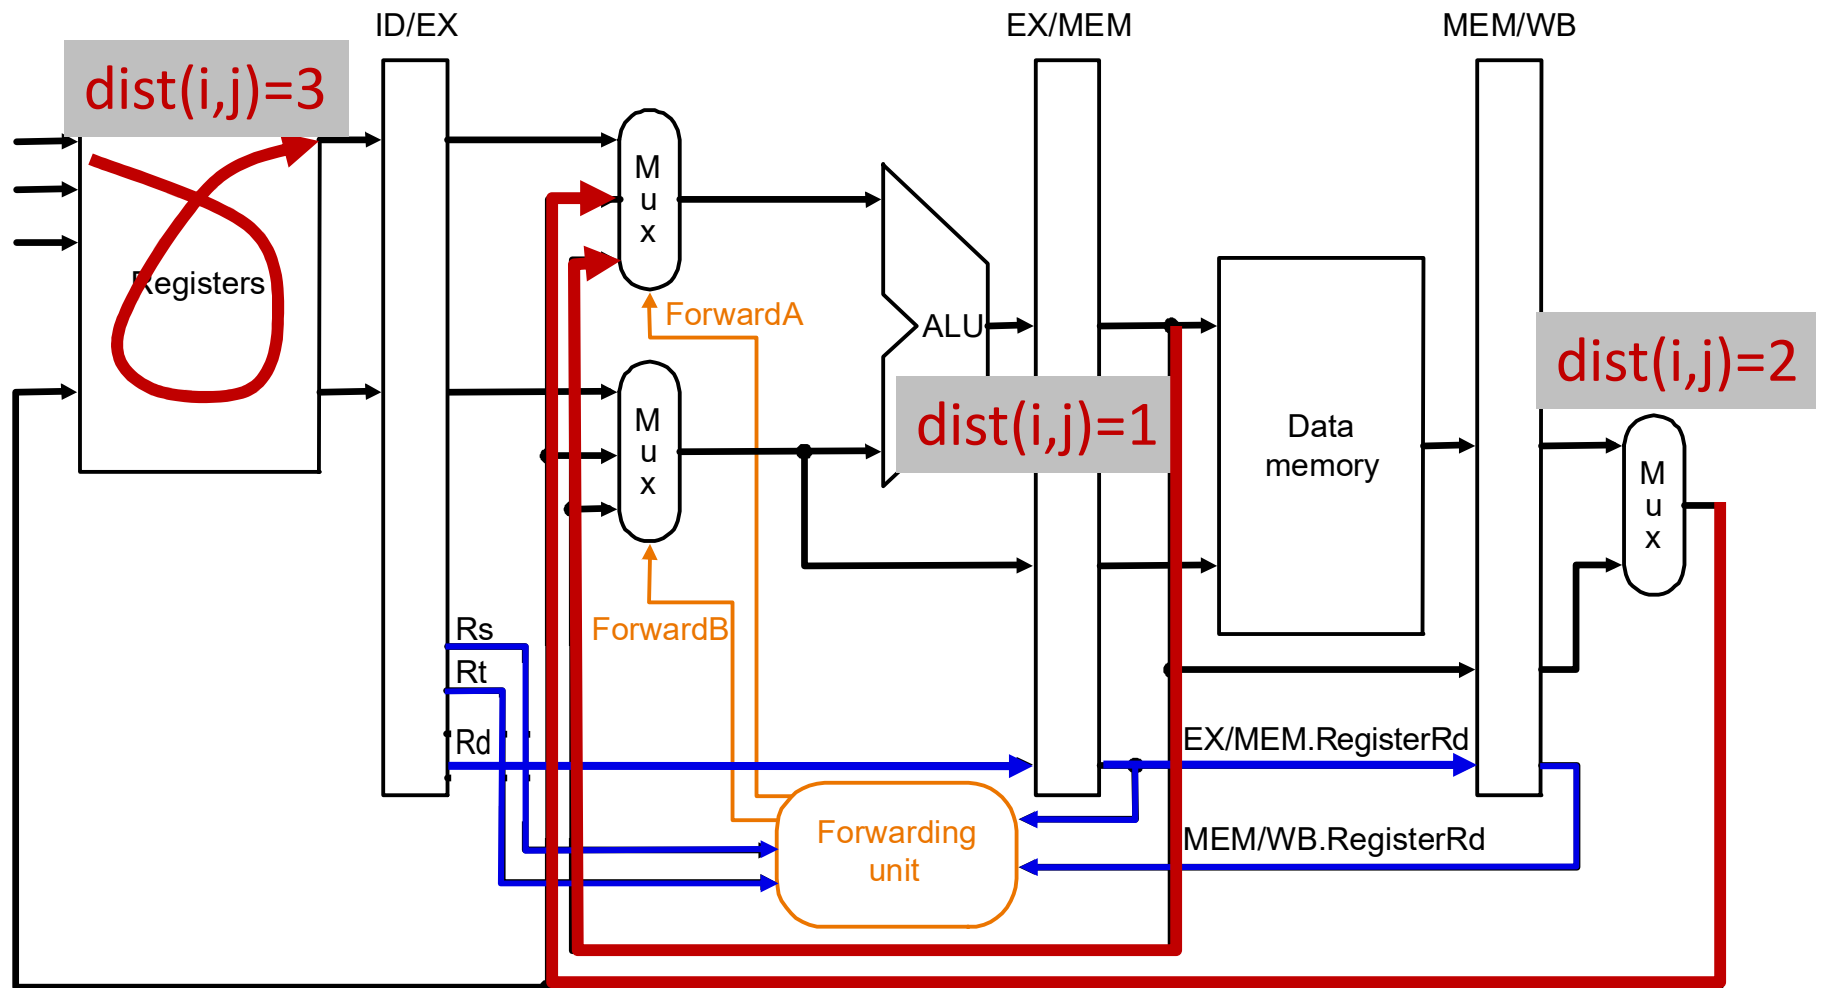
A hazard exits

- Older I_A and younger I_B have RAW hazard iff
 - I_B (R/I, LW, SW, Bxx or JALR) reads a register written by I_A (R/I, LW, or JAL/R)
 - $\text{dist}(I_A, I_B) \leq \text{dist}(\text{ID}, \text{WB}) = 3$
- More plainly, before I_B in ID reads a register, I_B needs to check if any I_A in EX, MEM or WB is going to update it (if so, value in RF is “stale”)
- Before: I_B need to stall for RF to update
- Now: I_B need to stall for I_A to produce result
 - retrieve I_A result from datapath when ready
 - must retrieve from **youngest** if multiple hazards

Forwarding Paths (v1)



Forwarding Paths (v2)



better if EX is the fastest stage

Forwarding Logic (for v1)

```

if ( $rs1_{ID} \neq 0$ ) && ( $rs1_{ID} == rd_{EX}$ ) &&  $RegWrite_{EX}$  then
    forward writeback value from EX           // dist=1
else if ( $rs1_{ID} \neq 0$ ) && ( $rs1_{ID} == rd_{MEM}$ ) &&  $RegWrite_{MEM}$  then
    forward writeback value from MEM          // dist=2
else if ( $rs1_{ID} \neq 0$ ) && ( $rs1_{ID} == rd_{WB}$ ) &&  $RegWrite_{WB}$  then
    forward writeback value from WB           // dist=3
else
    use  $A_{ID}$                                  // dist > 3

```

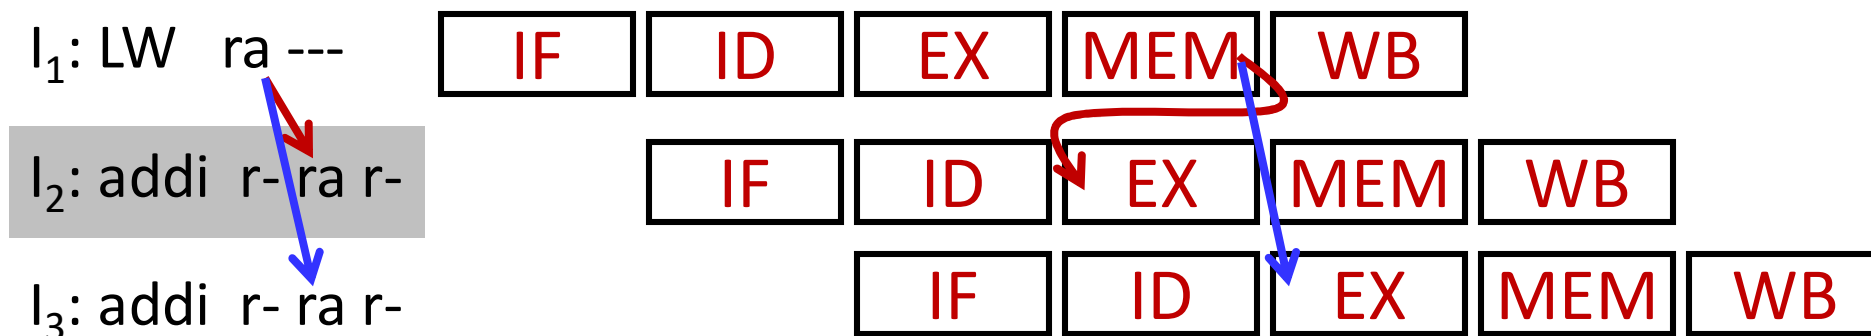
Must check in right order
 Why doesn't *use_rs1()* appear?
 Isn't it bad to forward from LW in EX?

Data Hazard Analysis (with Forwarding)

	R/I-Type	LW	SW	Bxx	Jal	Jalr
IF						
ID						
EX	use produce	use	use	use	produce	use produce
MEM		produce	(use)			
WB						

- Even with forwarding, RAW dependence on immediate preceding LW results in hazard
- $\text{Stall} = \{ [(rs1_{ID} == rd_{EX}) \ \&\& \ use_rs1(IR_{ID})] \ || \ i.e., \ op_{EX}=Lx$
 $[(rs2_{ID} == rd_{EX}) \ \&\& \ use_rs2(IR_{ID})] \} \ \&\& \ MemRead_{EX}$

Historical: MIPS Load “Delay Slot”



- R2000 defined LW with arch. latency of 1 inst
 - invalid for I_2 (in LW's delay slot) to ask for LW's result
 - any dependence on LW at least distance 2
- Delay slot vs dynamic stalling
 - fill with an independent instruction (no difference)
 - if not, fill with a NOP (no difference)
- Can't lose on 5-stage . . . good idea?

Hint: 1. non-atomic instruction; 2. μ arch influence

Sample Assembly [P&H]

for (j=i-1; j>=0 && v[j] > v[j+1]; j-=1) { }

```

    addi    $s1, $s0, -1
for2tst:   slti    $t0, $s1, 0
           bne    $t0, $zero, exit2
           sll    $t1, $s1, 2
           add    $t2, $a0, $t1
           lw     $t3, 0($t2)
           lw     $t4, 4($t2)
           slt    $t0, $t4, $t3
           beq    $t0, $zero, exit2
           .....
           addi   $s1, $s1, -1
           j     for2tst

```

1 stall or
1 nop (MIPS)


exit2:

Why not very deep pipelines?

- With only 5 stages, still plenty of combinational logic between registers
- “Superpipelining” \Rightarrow increase pipelining such that even intrinsic operations (e.g. ALU, RF access, memory access) require multiple stages
- What’s the problem?

Inst₀: r1 \leftarrow r2 + r3

Inst₁: r4 \leftarrow r1 + 2



Terminology

- Dependency
 - ordering requirement between instructions
- Pipeline Hazard:
 - (potential) violation of dependencies
- Hazard Resolution:
 - static \Rightarrow schedule instructions at compile time to avoid hazards
 - dynamic \Rightarrow detect hazard and adjust pipeline operation
Stall, Flush or Forward
- Pipeline Interlock (i.e., stall)

Dependencies and Pipelining


(architecture vs. microarchitecture)

Sequential and atomic
instruction semantics

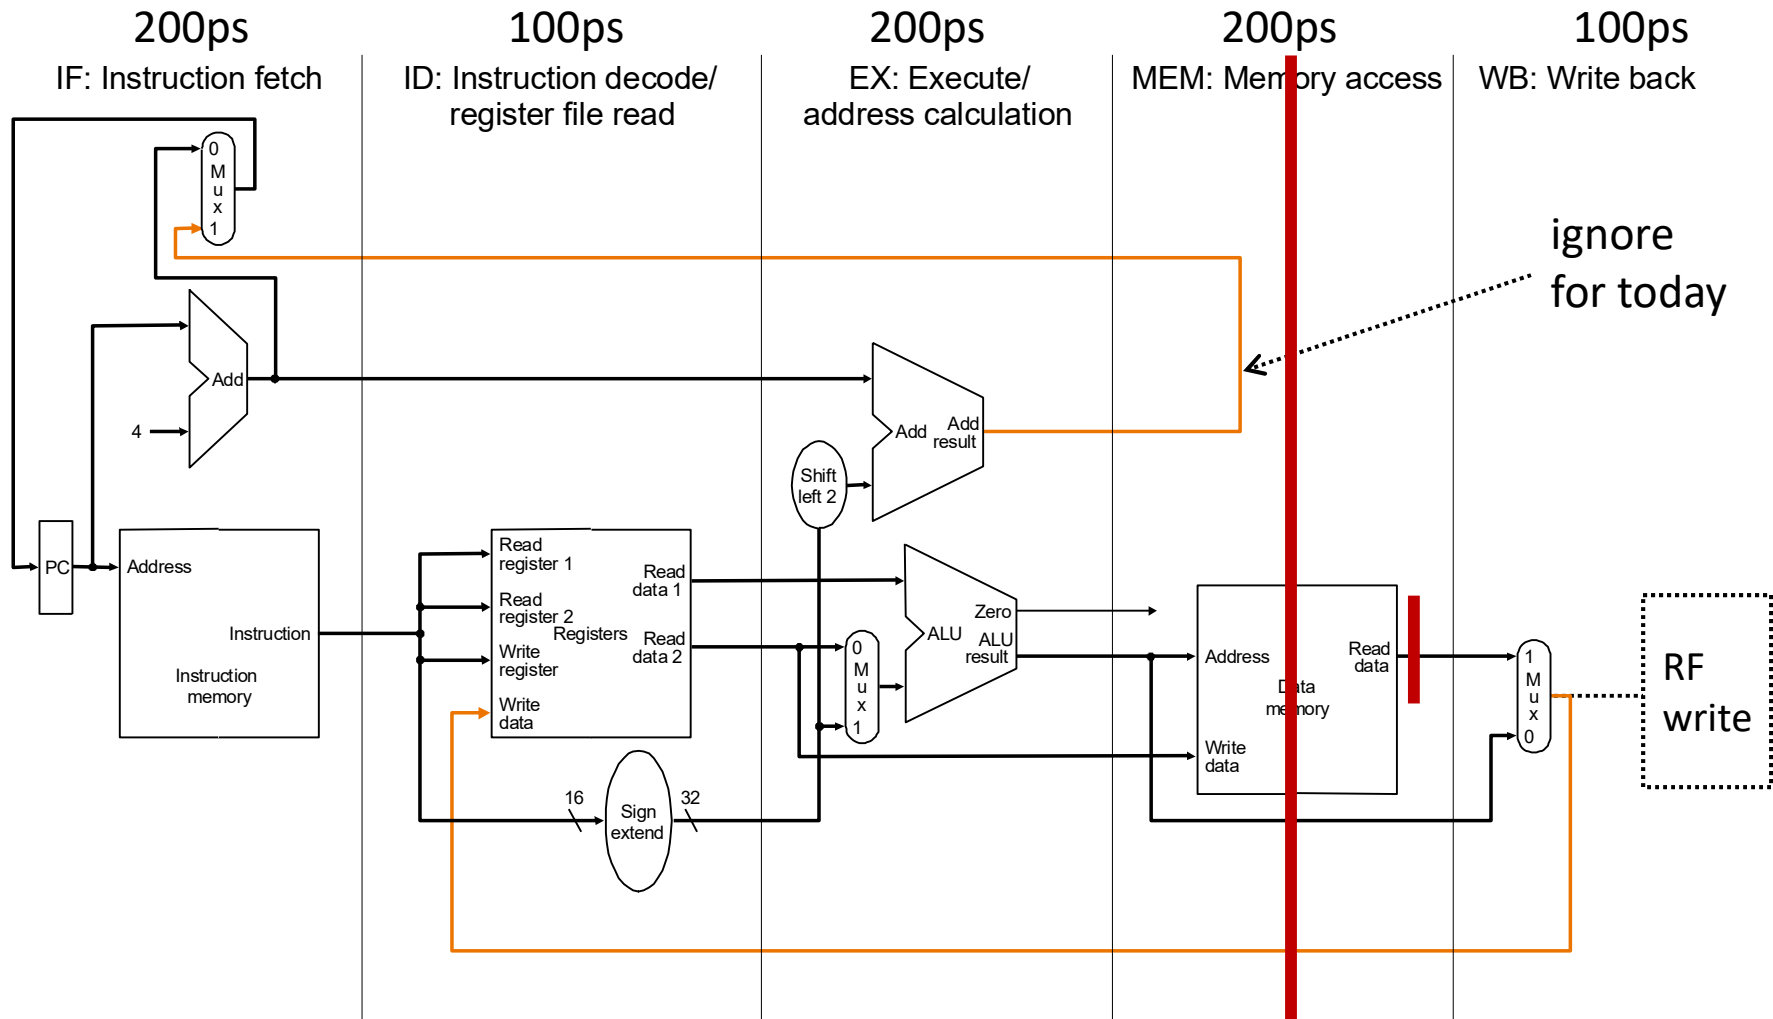


Defines what is correct;
doesn't say do it this way

True dependence between two
instructions may only require
ordering of certain sub-operations



Lab 2 with 6 Stages



Lab 2: Data Hazard Analysis with Stalling

	R/I-Type	LW	SW	Bxx	Jal	Jalr
IF						
ID	read RF	read RF	read RF	read RF		read RF
EX						
MEM1						
MEM2						
WB	write RF	write RF			write RF	write RF

Lab 2: Data Hazard Analysis with Forwarding

	R/I-Type	LW	SW	Bxx	Jal	Jalr
IF						
ID						
EX	use produce	use	use	use	produce	use produce
MEM1			(use)			
MEM2						
WB		produce				