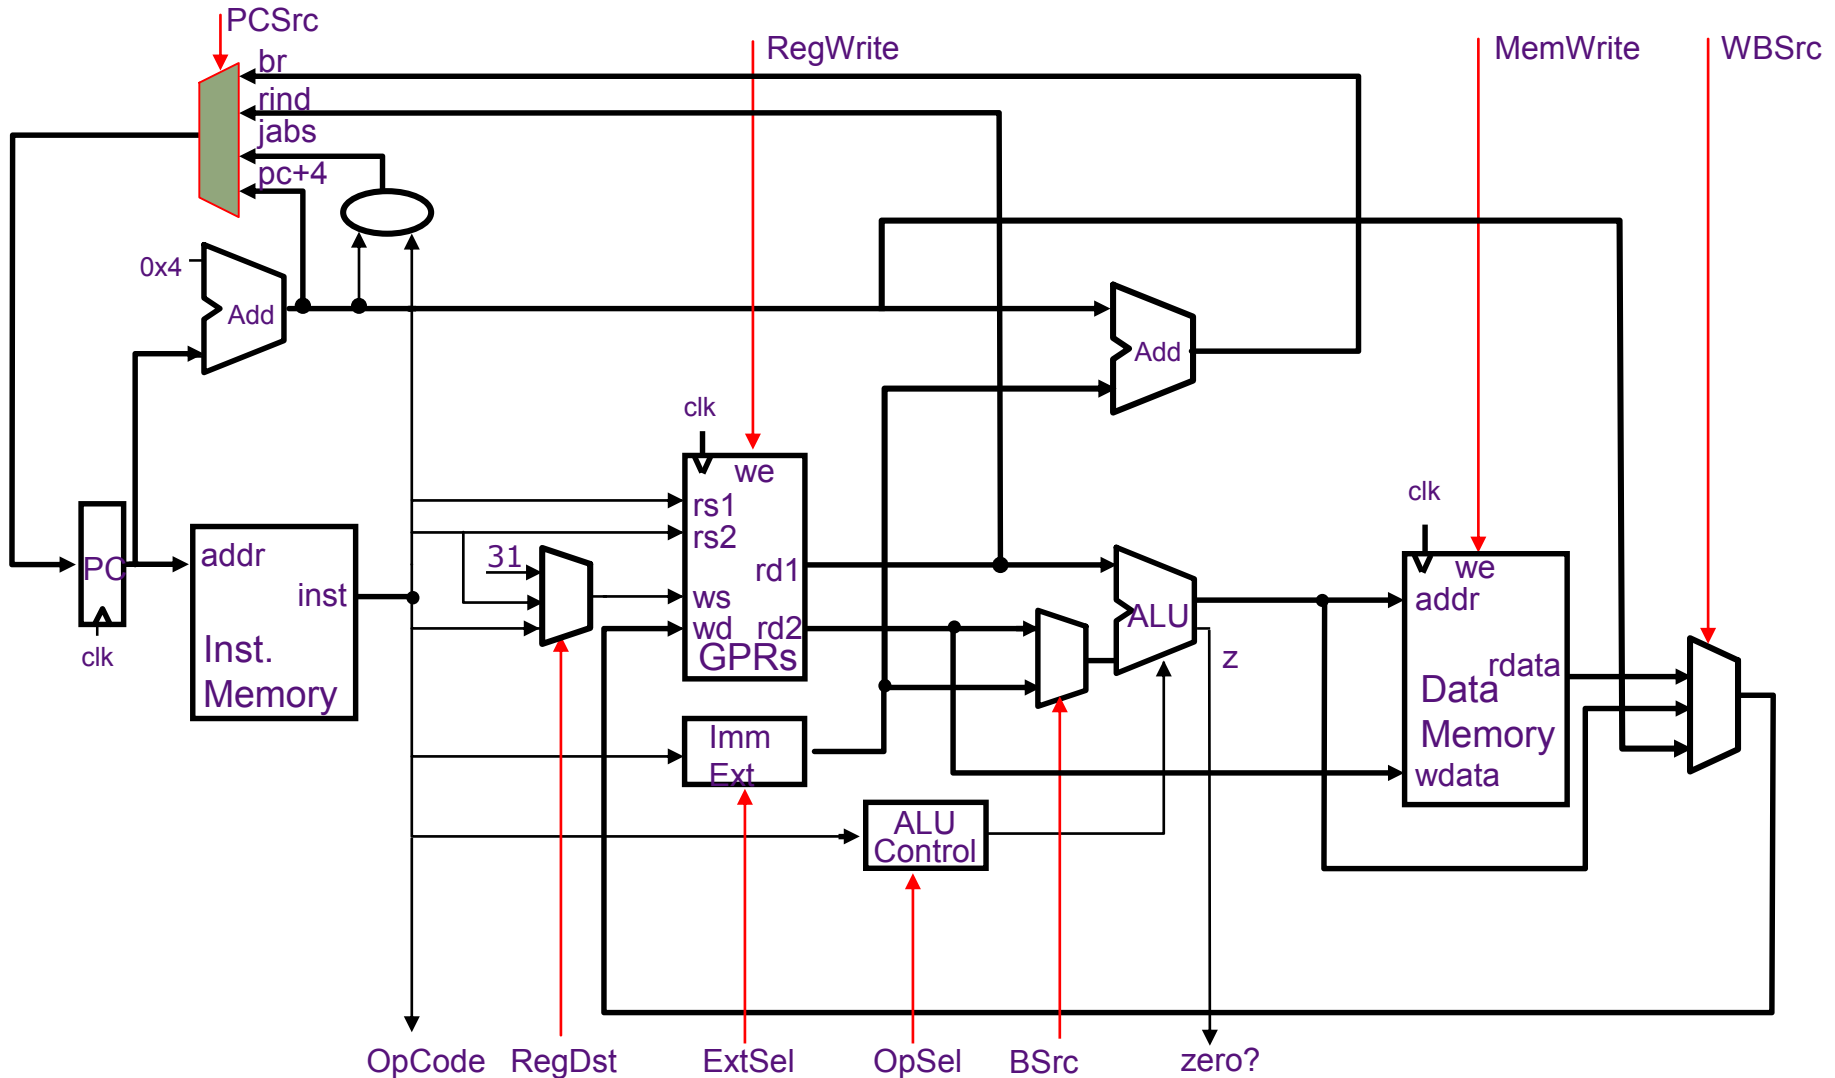


# Instruction Pipelining and Hazards

*Daniel Sanchez*


Computer Science and Artificial Intelligence Laboratory  
M.I.T.

# Reminder: Harvard-Style Single-Cycle Datapath for MIPS

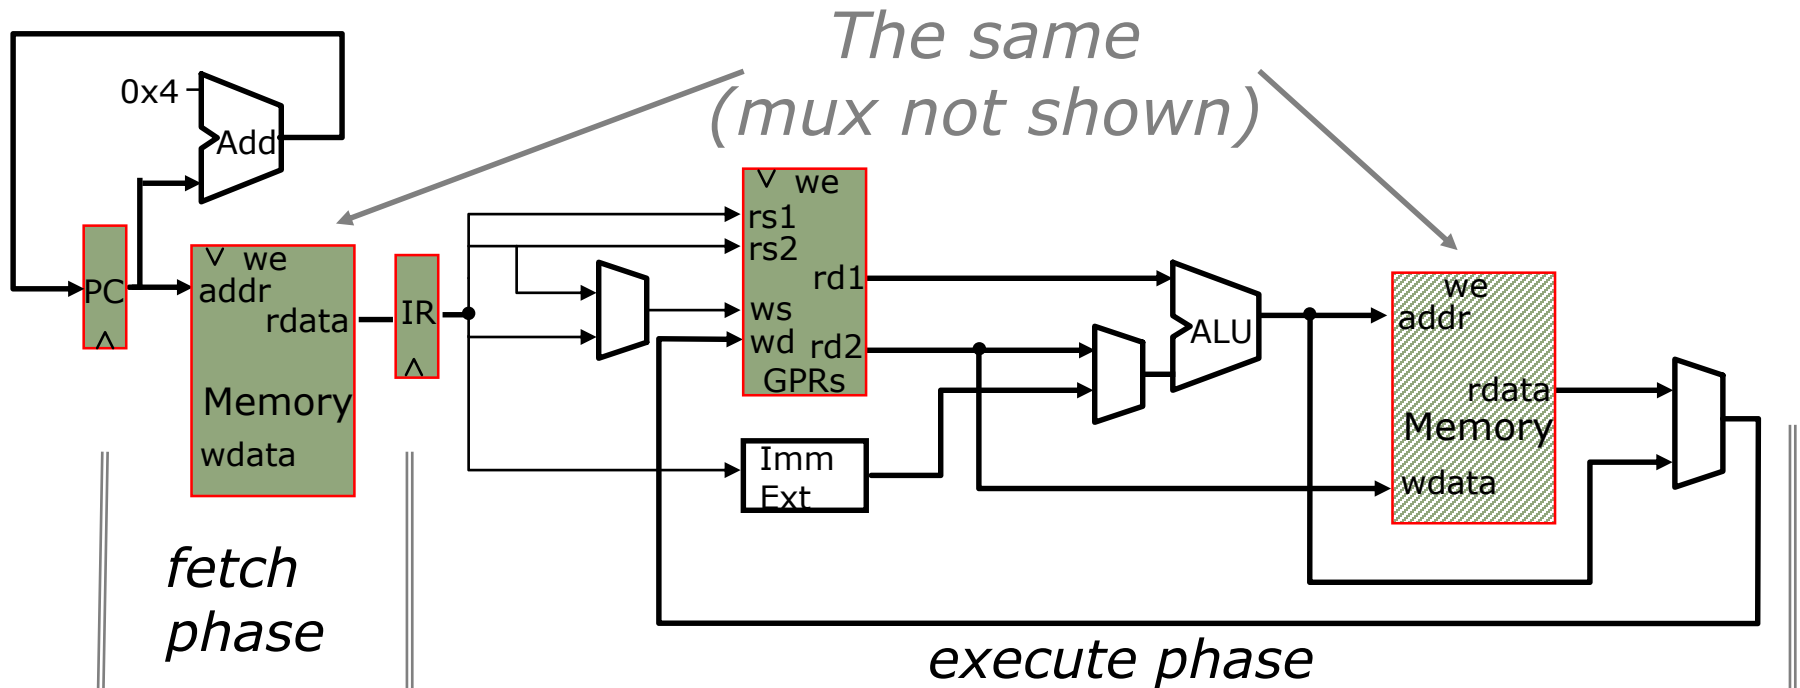


# Reminder: Princeton Microarchitecture

## Datapath & Control for 2 cycles-per-instruction




# Princeton Microarchitecture (redrawn)



Only one of the phases is active in any cycle  
⇒ a lot of datapath not used at any given time

# Princeton Microarchitecture

## *Overlapped execution*




*Can we overlap instruction fetch and execute?*

*Which action should be prioritized?*

*What do we do with Fetch?*

# Stalling the instruction fetch

## Princeton Microarchitecture




When stall condition is indicated

- *don't fetch a new instruction and don't change the PC*
- *insert a nop in the IR*
- *set the Memory Address mux to ALU (not shown)*

*What if IR contains a jump or branch instruction?*

# Need to stall on branches

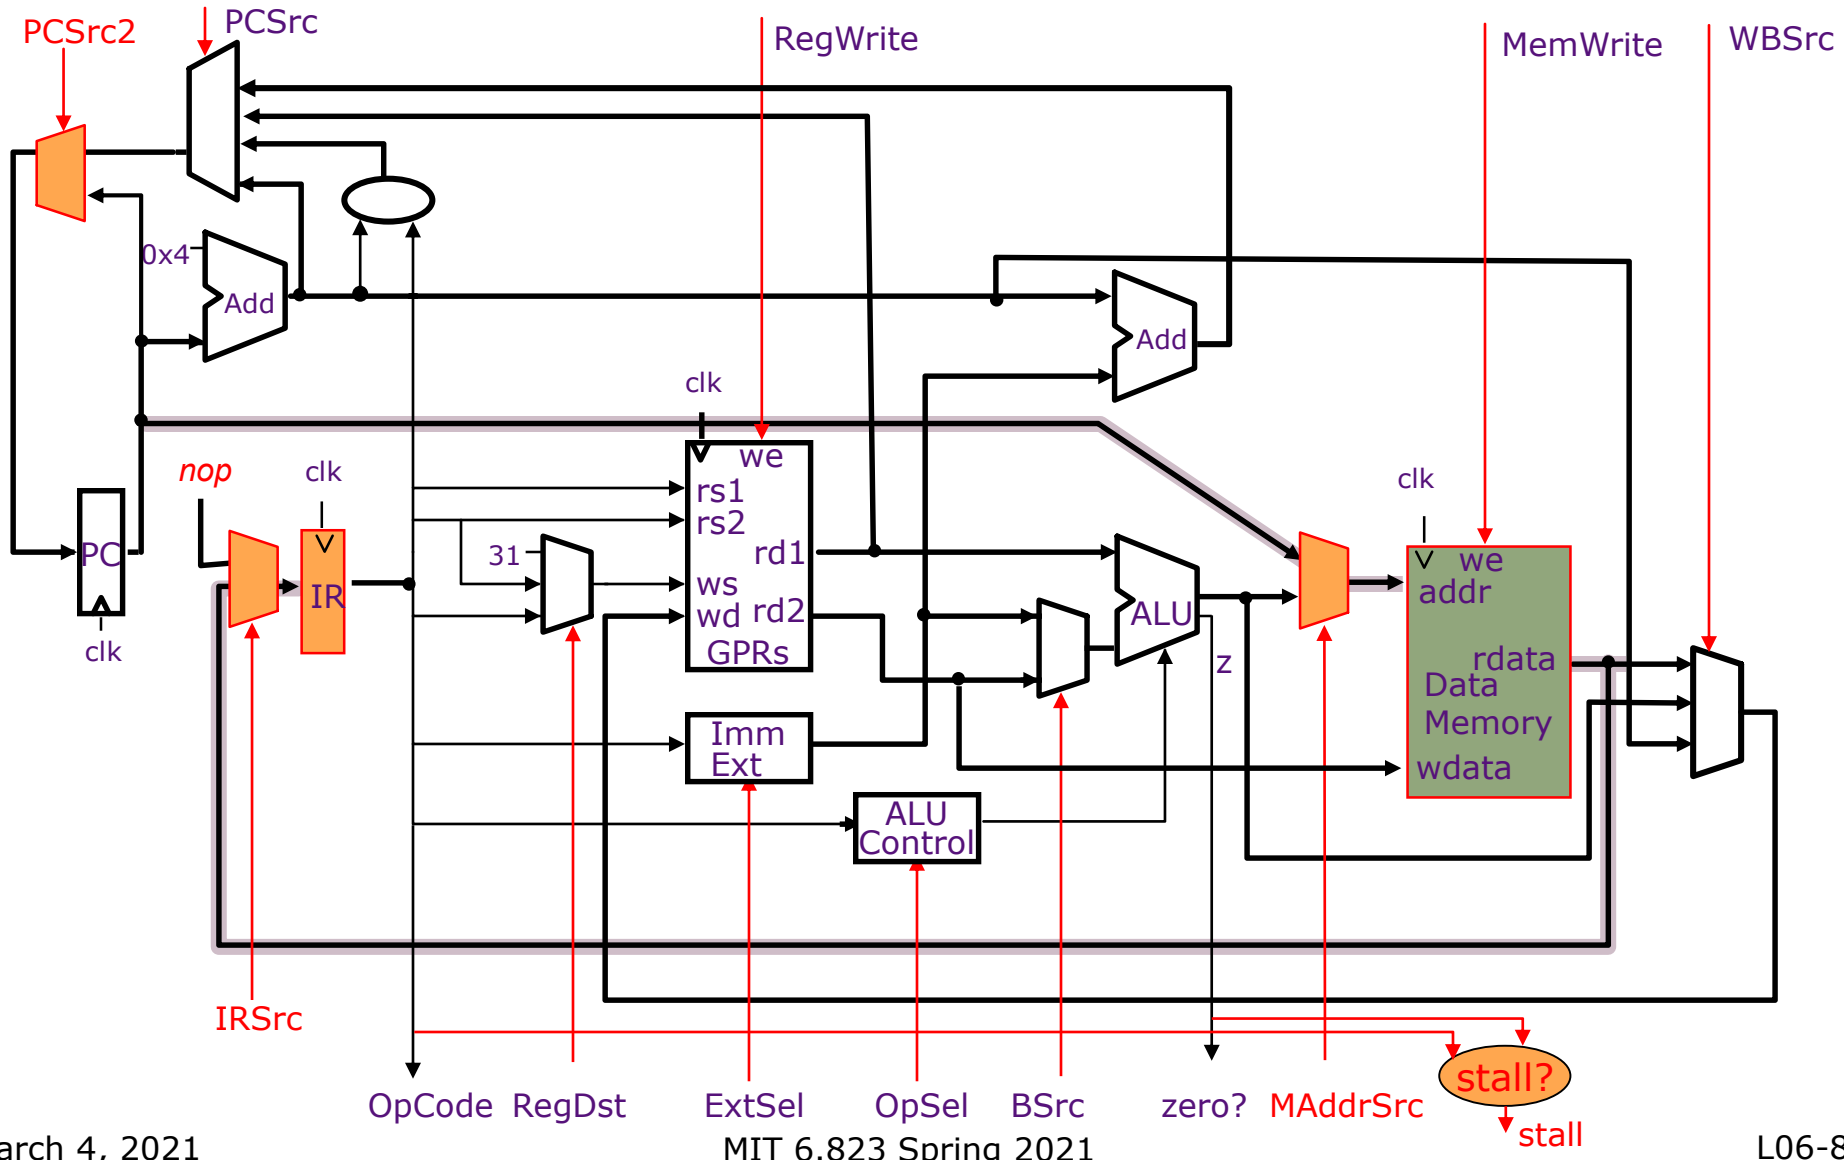
## Princeton Microarchitecture



When IR contains a jump or taken branch

- *no "structural conflict" for the memory*
- *but we do not have the correct PC value in the PC*
- *memory cannot be used – Address Mux setting is irrelevant*
- *insert a nop in the IR*
- *insert the nextPC (branch-target) address in the PC*

# Pipelined Princeton Microarchitecture





# Pipelined Princeton: Control Table

Opcode	Stall	Ext Sel	B Src	Op Sel	Mem W	Reg W	WB Src	Reg Dst	PC Src1	PC Src2	IR Src	MAddr Src
ALU	no	*	Reg	Func	no	yes	ALU	rd	pc+4	npc	mem	pc
ALUi	no	sE <sub>16</sub>	Imm	Op	no	yes	ALU	rt	pc+4	npc	mem	pc
ALUiu	no	uE <sub>16</sub>	Imm	Op	no	yes	ALU	rt	pc+4	npc	mem	pc
LW	yes	sE <sub>16</sub>	Imm	+	no	yes	Mem	rt	pc+4	pc	nop	ALU
SW	yes	sE <sub>16</sub>	Imm	+	yes	no	*	*	pc+4	pc	nop	ALU
BEQZ <sub>Z=1</sub>	yes	sE <sub>16</sub>	*	0?	no	no	*	*	br	npc	nop	*
BEQZ <sub>Z=0</sub>	no	sE <sub>16</sub>	*	0?	no	no	*	*	pc+4	npc	mem	pc
J	yes	*	*	*	no	no	*	*	jabs	npc	nop	*
JAL	yes	*	*	*	no	yes	PC	R31	jabs	npc	nop	*
JR	yes	*	*	*	no	no	*	*	rind	npc	nop	*
JALR	yes	*	*	*	no	yes	PC	R31	rind	npc	nop	*
NOP	no	*	*	*	no	no	*	*	pc+4	npc	mem	pc

BSrc = Reg / Imm ; WBSrc = ALU / Mem / PC; IRSrc = nop/mem; MAddrSrc = pc/ALU

RegDst = rt / rd / R31; PCSrc1 = pc+4 / br / rind / jabs; PCSrc2 = pc/nPC

stall & IRSrc columns are identical

# Pipelined Princeton Architecture

---


*Clock:*  $t_{\text{C-Princeton}} > t_{\text{RF}} + t_{\text{ALU}} + t_{\text{M}} + t_{\text{WB}}$

*CPI:*  $(1 - f) + 2f$  cycles per instruction  
where  $f$  is the fraction of  
instructions that cause a stall

*What is a likely value of  $f$ ?*

# An Ideal Pipeline

---




- All objects go through the same stages
- No sharing of resources between any two stages
- Propagation delay through all pipeline stages is equal
- The scheduling of an object entering the pipeline is not affected by the objects in other stages

*These conditions generally hold for industrial assembly lines.*

*But what about an instruction pipeline?*

# Pipelined Datapath



Clock period can be reduced by dividing the execution of an instruction into multiple cycles

$$t_C > \max \{t_{IM}, t_{RF}, t_{ALU}, t_{DM}, t_{RW}\} \quad (= t_{DM} \text{ probably})$$

*However, CPI will increase unless instructions are pipelined*

# How to divide datapath into stages


---

Suppose memory is significantly slower than other stages. For example, suppose

$$\begin{aligned}t_{IM} &= 10 \text{ units} \\t_{DM} &= 10 \text{ units} \\t_{ALU} &= 5 \text{ units} \\t_{RF} &= 1 \text{ unit} \\t_{RW} &= 1 \text{ unit}\end{aligned}$$

Since the slowest stage determines the clock, it may be possible to combine some stages without any loss of performance

# Alternative Pipelining



$$t_C > \max \{t_{IM}, t_{RF}+t_{ALU}, t_{DM}+t_{RW}\} = t_{DM} + t_{RW}$$

Write-back stage takes much less time than other stages.  
 Suppose we combined it with the memory phase

# Maximum Speedup by Pipelining


---

Assumptions	Unpipelined $t_c$	Pipelined Speedup $t_c$
1. $t_{IM} = t_{DM} = 10,$ $t_{ALU} = 5,$ $t_{RF} = t_{RW} = 1$ 4-stage pipeline		
2. $t_{IM} = t_{DM} = t_{ALU} = t_{RF} = t_{RW} = 5$ 4-stage pipeline		
3. $t_{IM} = t_{DM} = t_{ALU} = t_{RF} = t_{RW} = 5$ 5-stage pipeline		

*What seems to be the message here?*

# 5-Stage Pipelined Execution


## Instruction Flow Diagram






# 5-Stage Pipelined Execution

## Resource Usage Diagram



# Pipelined Execution

## *ALU Instructions*



*Not quite correct!*



# How instructions can interact with each other in a pipeline

---

- An instruction in the pipeline may need a resource being used by another instruction in the pipeline  
→ *structural hazard*
- An instruction may depend on a value produced by an earlier instruction
  - Dependence may be for a data calculation  
→ *data hazard*
  - Dependence may be for calculating the next PC  
→ *control hazard (branches, interrupts)*



# Resolving Data Hazards

---

Strategy 1: *Wait for the result to be available by freezing earlier pipeline stages* → *stall*

Strategy 2: *Route data as soon as possible after it is calculated to the earlier pipeline stage* → *bypass*

Strategy 3: *Speculate on the dependence*

*Two cases:*

*Guessed correctly* → *do nothing*

*Guessed incorrectly* → *kill and restart*

# Resolving Data Hazards (1)

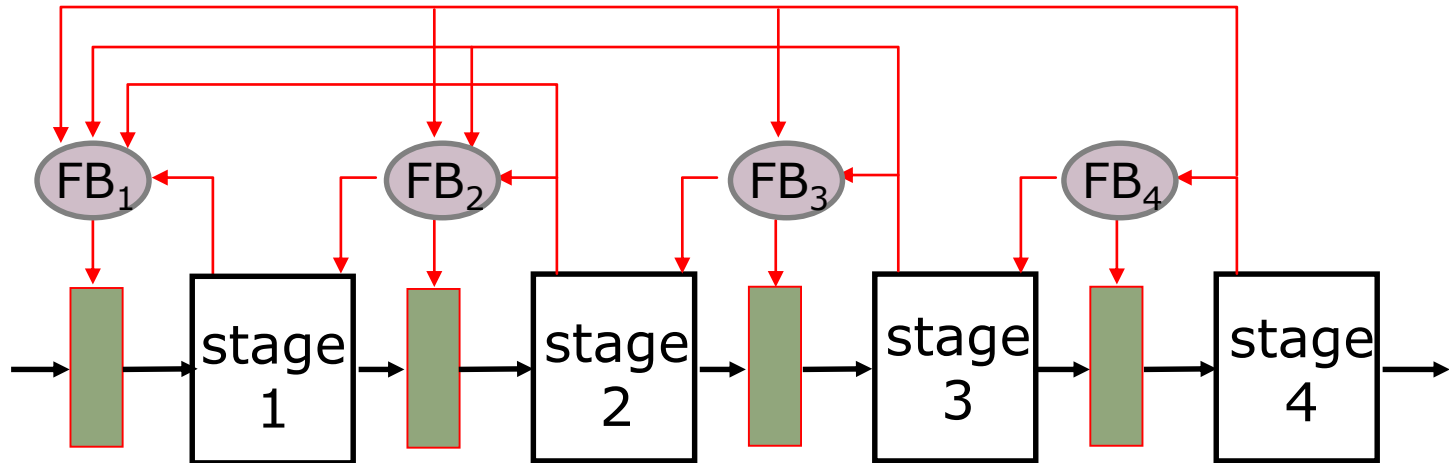
---

*Strategy 1:*

*Wait for the result to be available by freezing earlier pipeline stages → **stall***

# Feedback to Resolve Hazards

---

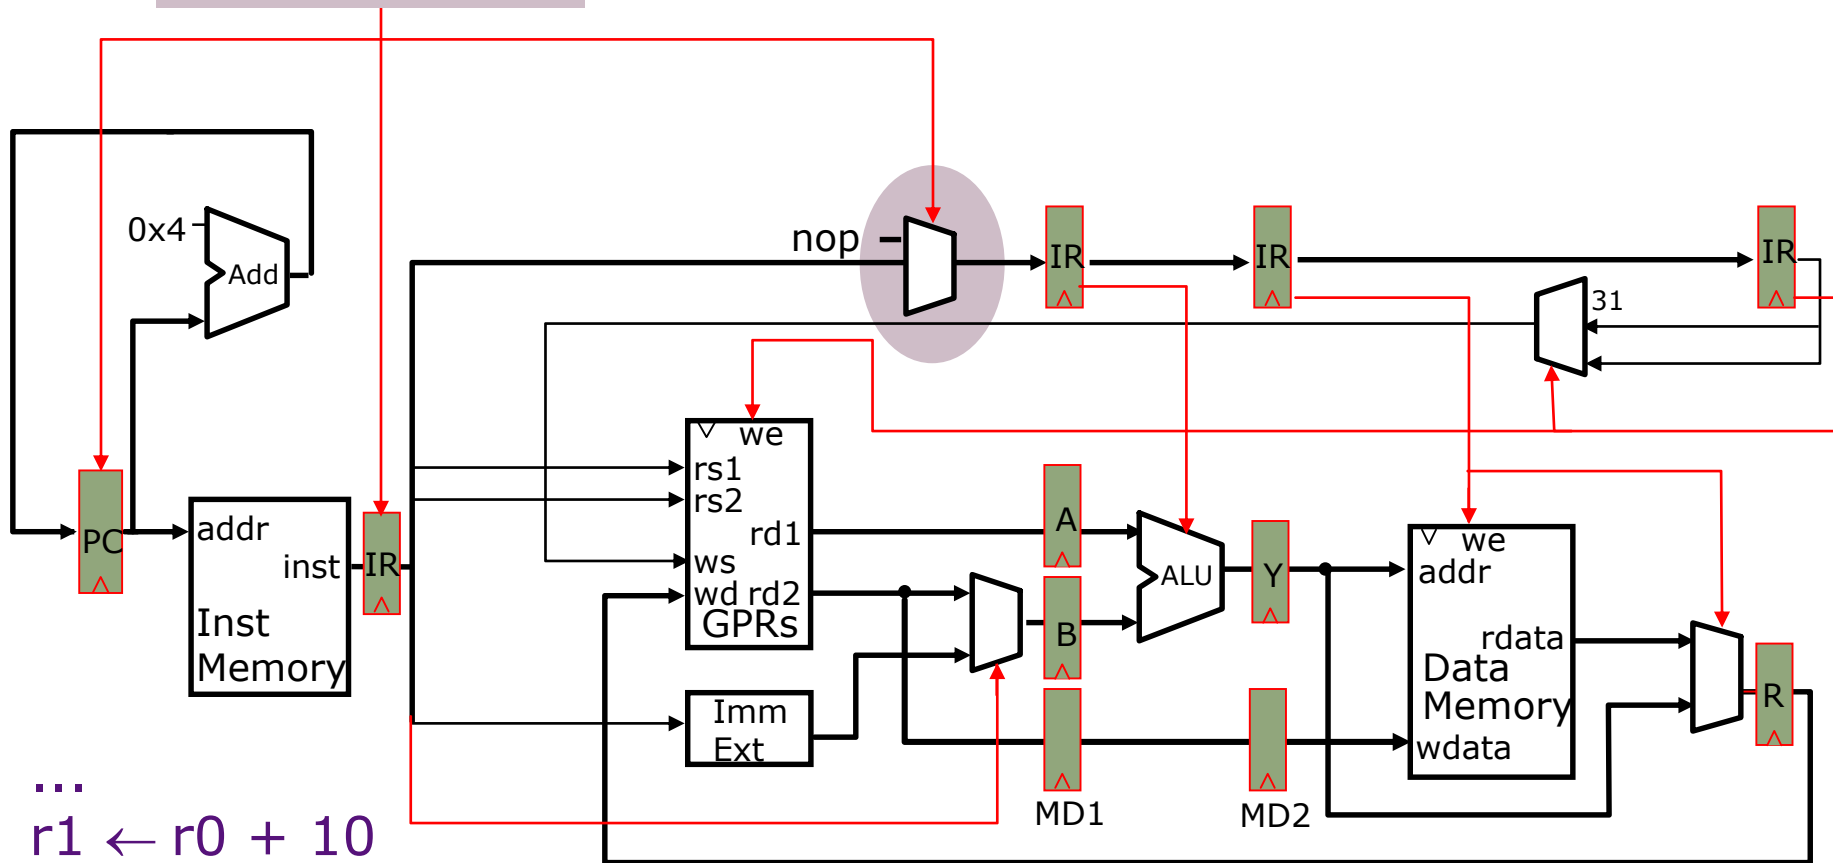


- Later stages provide dependence information to earlier stages which can *stall (or kill) instructions*
- Controlling a pipeline in this manner works provided *the instruction at stage  $i+1$  can complete without any interference from instructions in stages 1 to  $i$*  (otherwise deadlocks may occur)





# Resolving Data Hazards by Stalling

*Stall Condition*

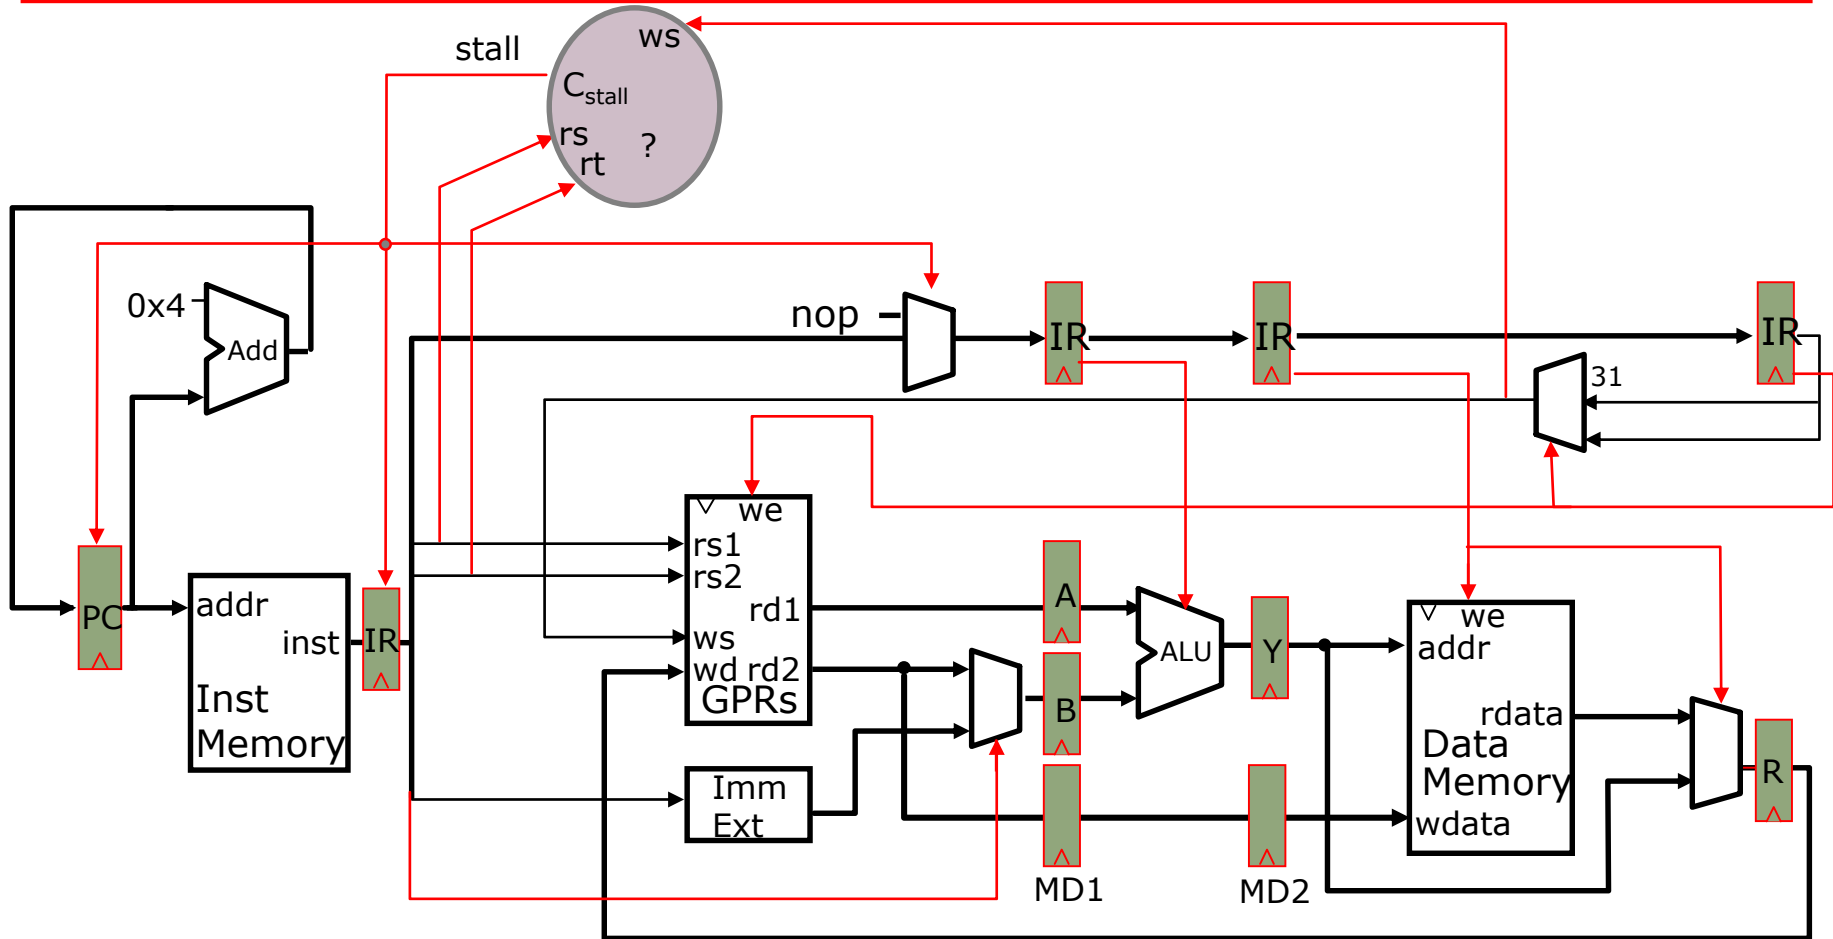


# Stalled Stages and Pipeline Bubbles



*nop* ⇒ *pipeline bubble*

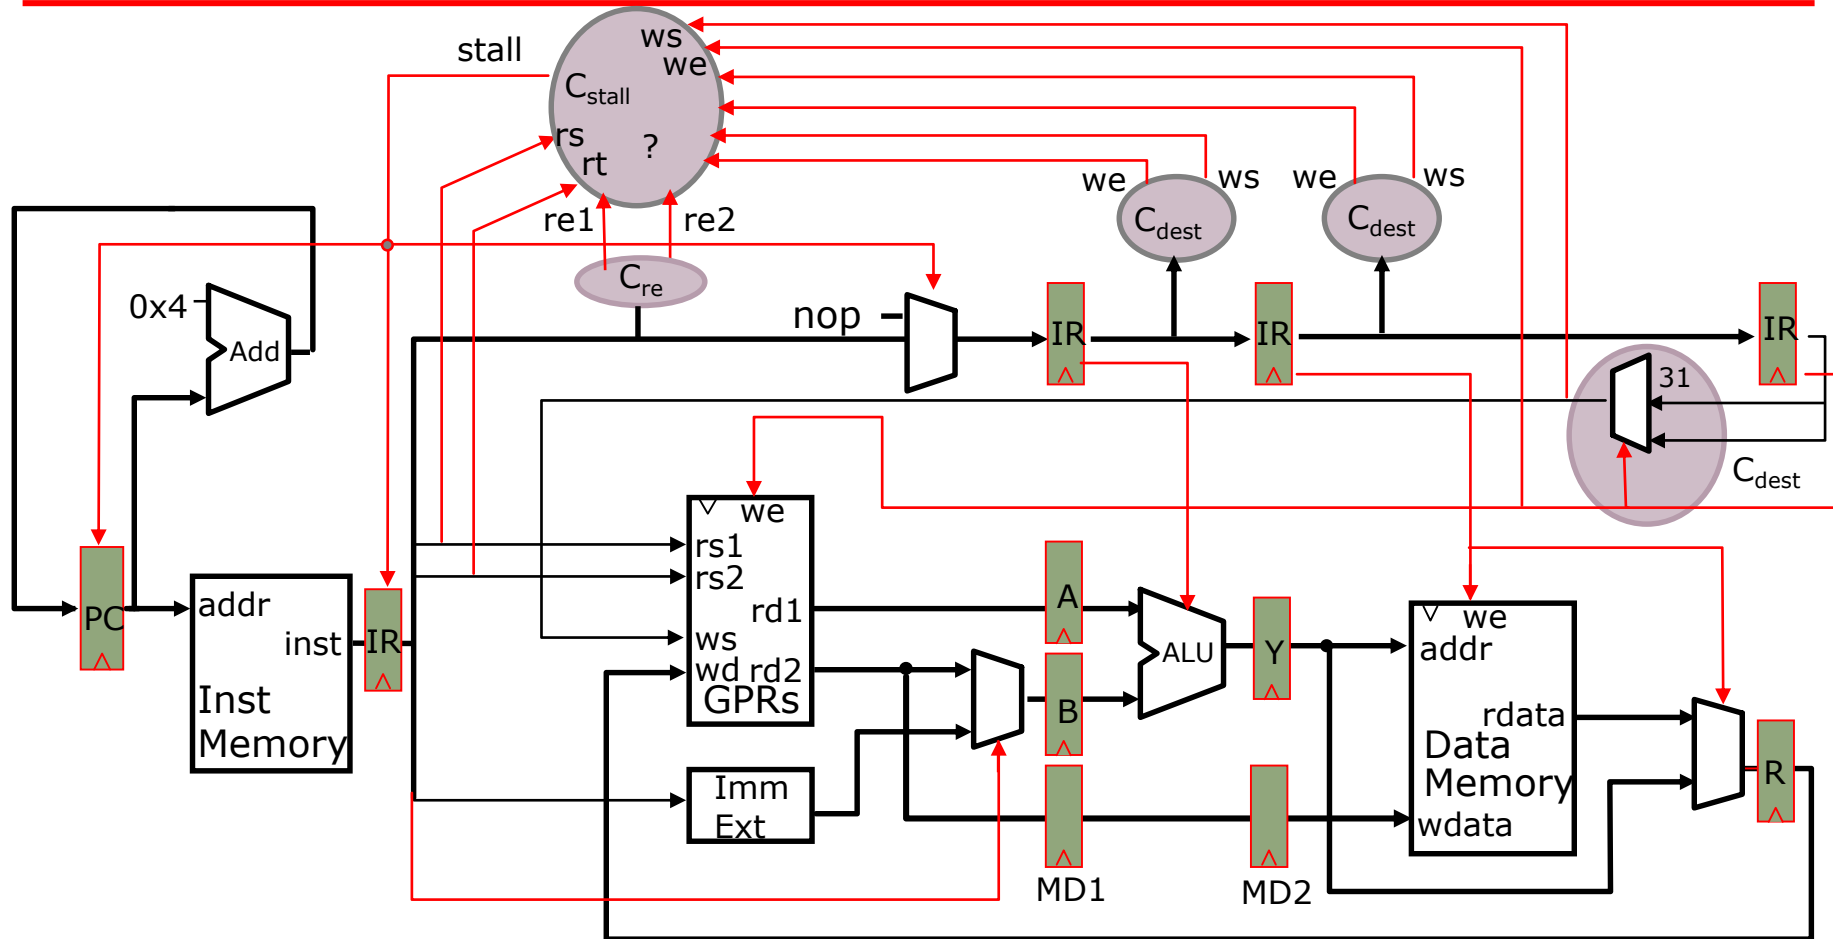
# Stall Control Logic



Compare the *source registers* of the instruction in the decode stage with the *destination register* of the *uncommitted instructions*.

# Stall Control Logic

*ignoring jumps & branches*



Should we always stall if the rs field matches some rd?

# Source & Destination Registers

*R-type:*

op	rs	rt	rd		func
----	----	----	----	--	------

*I-type:*

op	rs	rt	immediate16
----	----	----	-------------

*J-type:*

op	immediate26
----	-------------

*source(s) destination*

ALU	$rd \leftarrow (rs) \text{ func } (rt)$	<i>rs, rt</i>	<i>rd</i>
ALUi	$rt \leftarrow (rs) \text{ op } \text{imm}$	<i>rs</i>	<i>rt</i>
LW	$rt \leftarrow M [(rs) + \text{imm}]$	<i>rs</i>	<i>rt</i>
SW	$M [(rs) + \text{imm}] \leftarrow (rt)$	<i>rs, rt</i>	
BZ	<i>cond (rs)</i>		
	<i>true:</i> $PC \leftarrow (PC) + \text{imm}$	<i>rs</i>	
	<i>false:</i> $PC \leftarrow (PC) + 4$	<i>rs</i>	
J	$PC \leftarrow (PC) + \text{imm}$		
JAL	$r31 \leftarrow (PC), PC \leftarrow (PC) + \text{imm}$		<i>31</i>
JR	$PC \leftarrow (rs)$	<i>rs</i>	
JALR	$r31 \leftarrow (PC), PC \leftarrow (rs)$	<i>rs</i>	<i>31</i>

# Deriving the Stall Signal

$C_{dest}$

$ws = \text{Case opcode}$

ALU  $\Rightarrow rd$   
ALUi, LW  $\Rightarrow rt$   
JAL, JALR  $\Rightarrow R31$

$we = \text{Case opcode}$

ALU, ALUi, LW  $\Rightarrow (ws \neq 0)$   
JAL, JALR  $\Rightarrow on$   
...  $\Rightarrow off$

$C_{re}$

$re1 = \text{Case opcode}$

ALU, ALUi,  
LW, SW, BZ,  
JR, JALR  $\Rightarrow on$   
J, JAL  $\Rightarrow off$

$re2 = \text{Case opcode}$

ALU, SW  $\Rightarrow on$   
...  $\Rightarrow off$


$C_{stall}$

$$\begin{aligned} \text{stall} = & ((rs_D == ws_E) \cdot we_E + \\ & (rs_D == ws_M) \cdot we_M + \\ & (rs_D == ws_W) \cdot we_W) \cdot re1_D + \\ & ((rt_D == ws_E) \cdot we_E + \\ & (rt_D == ws_M) \cdot we_M + \\ & (rt_D == ws_W) \cdot we_W) \cdot re2_D \end{aligned}$$

*This is not  
the full story!*

# Hazards due to Loads & Stores

*Stall Condition*



...  
 $M[(r1)+7] \leftarrow (r2)$   
 $r4 \leftarrow M[(r3)+5]$

*Is there any possible data hazard in this instruction sequence?*

# Load & Store Hazards

---

```
...  
M[(r1)+7] ← (r2)  
r4 ← M[(r3)+5]  
...
```

$(r1)+7 = (r3)+5 \Rightarrow$  *data hazard*

However, the hazard is avoided because *our memory system completes writes in one cycle!*

Load/Store hazards are sometimes resolved in the pipeline and sometimes in the memory system itself.

*More on this later in the course.*



Next lecture:  
Control Hazards,  
Bypassing,  
and Speculation