



Pipeline Hazards

Arvind

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

*Based on the material prepared by
Arvind and Krste Asanovic*

Technology Assumptions

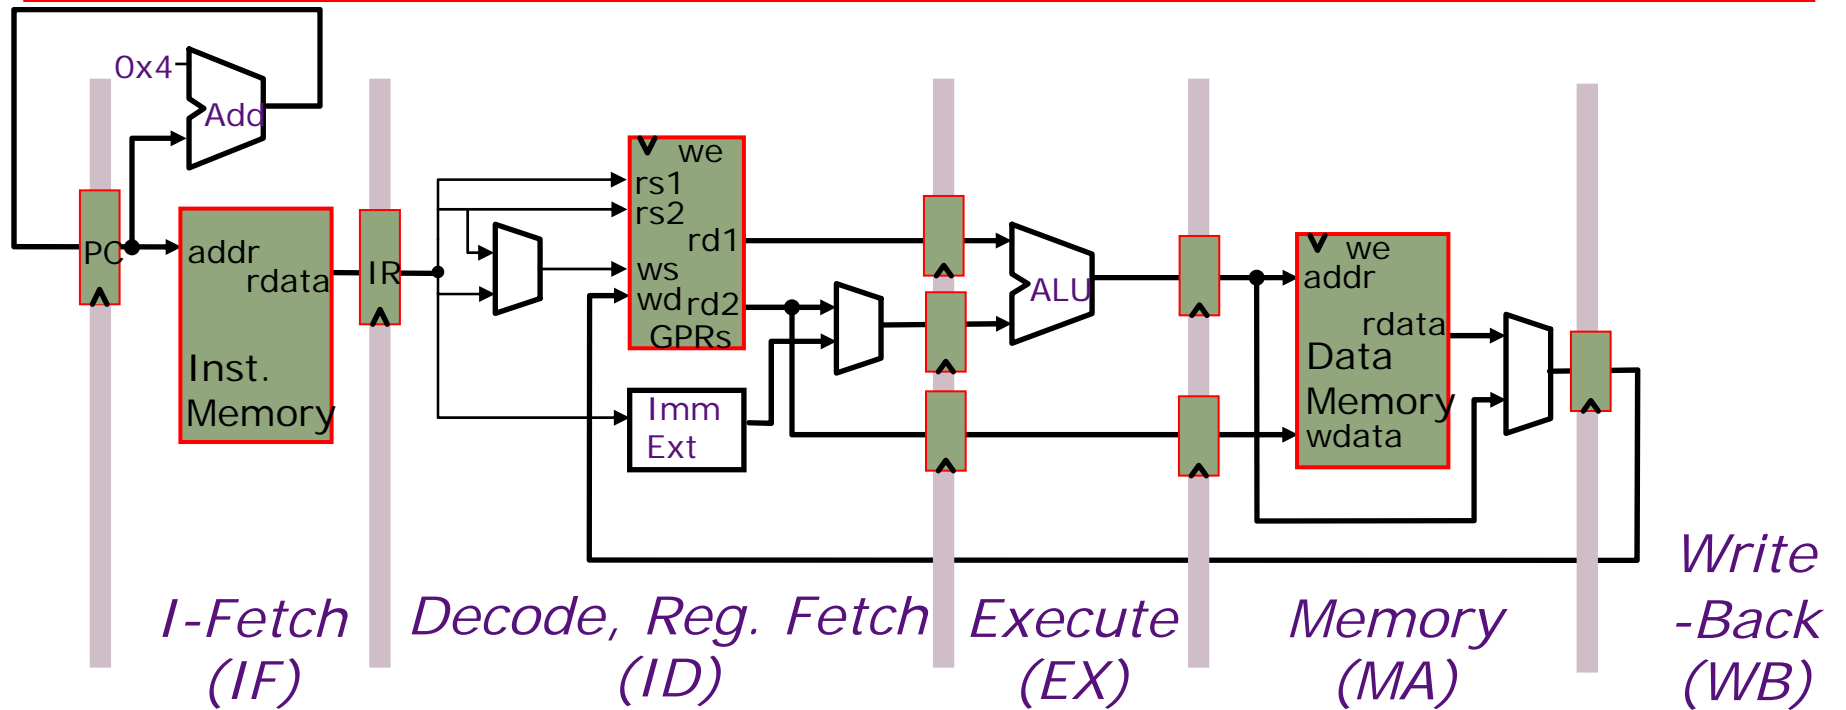
- A small amount of very fast memory (caches) backed up by a large, slower memory
- Fast ALU (at least for integers)
- Multiported Register files (slower!)

It makes the following timing assumption valid

$$t_{IM} \approx t_{RF} \approx t_{ALU} \approx t_{DM} \approx t_{RW}$$

A 5-stage pipelined Harvard architecture will be the focus of our detailed design

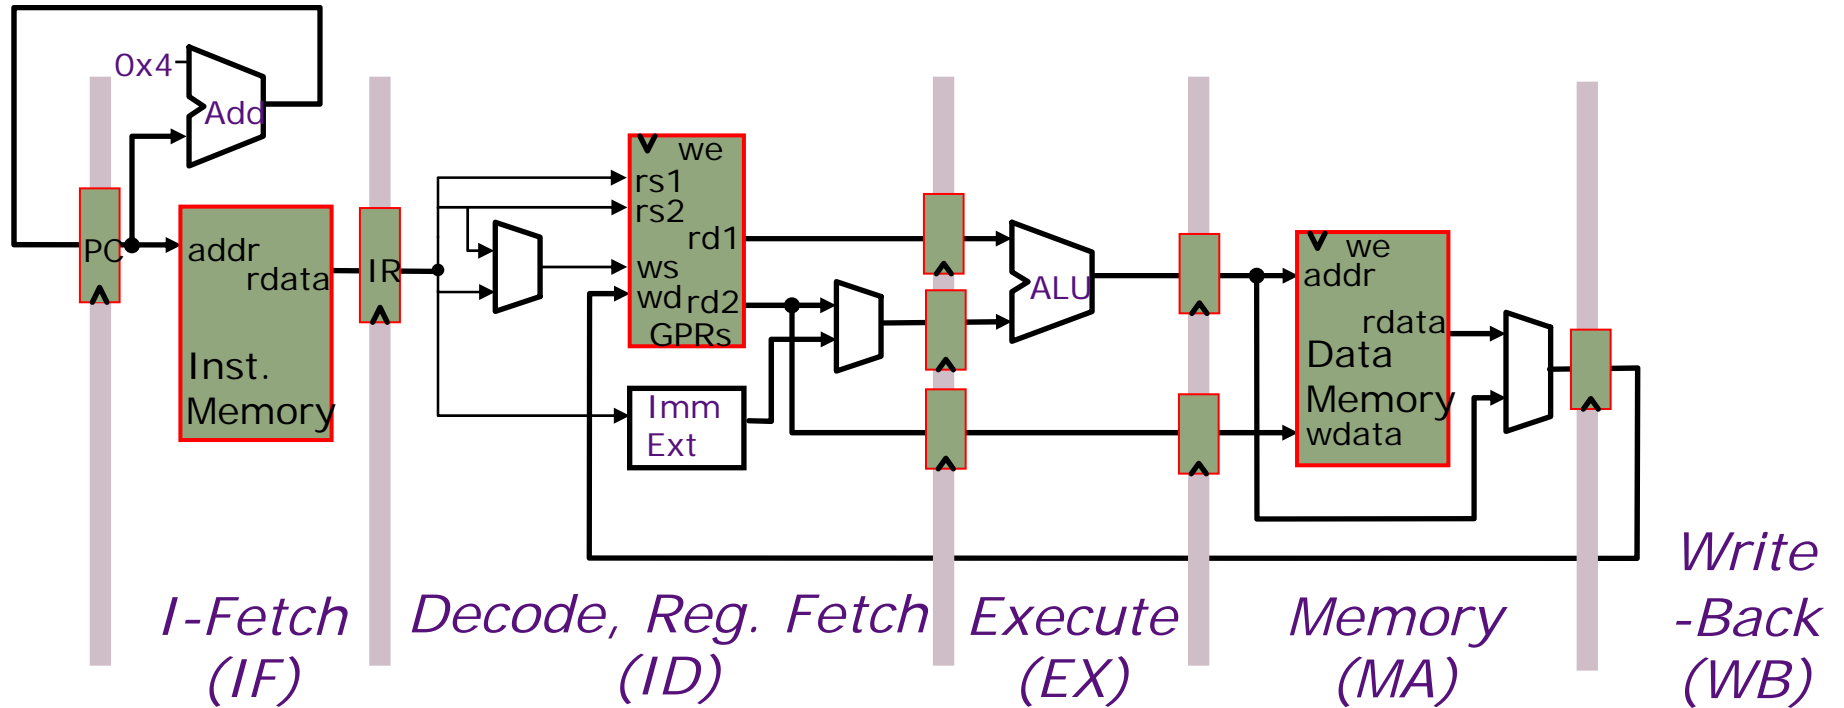
5-Stage Pipelined Execution



<i>time</i>	t0	t1	t2	t3	t4	t5	t6	t7
instruction1	IF ₁	ID ₁	EX ₁	MA ₁	WB ₁				
instruction2		IF ₂	ID ₂	EX ₂	MA ₂	WB ₂			
instruction3			IF ₃	ID ₃	EX ₃	MA ₃	WB ₃		
instruction4				IF ₄	ID ₄	EX ₄	MA ₄	WB ₄	
instruction5					IF ₅	ID ₅	EX ₅	MA ₅	WB ₅


5-Stage Pipelined Execution

Resource Usage Diagram



	<i>time</i>	t0	t1	t2	t3	t4	t5	t6	t7	...
Resources	<i>IF</i>	₁	₂	₃	₄	₅				
	<i>ID</i>		₁	₂	₃	₄	₅			
	<i>EX</i>			₁	₂	₃	₄	₅		
	<i>MA</i>				₁	₂	₃	₄	₅	
	<i>WB</i>					₁	₂	₃	₄	₅


Pipelined Execution: ALU Instructions



Not quite correct!

We need an Instruction Reg (IR) for each stage


IR's and Control points



Are control points connected properly?

- ALU instructions
- Load/Store instructions
- Write back


Pipelined MIPS Datapath *without jumps*



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 produce data that is needed by a later instruction
 - *data hazard*
- In the extreme case, an instruction may determine the next instruction to be executed
 - *control hazard (branches, interrupts,...)*

Data Hazards



...
 $r1 \leftarrow r0 + 10$
 $r4 \leftarrow r1 + 17$
...

r1 is stale. Oops!

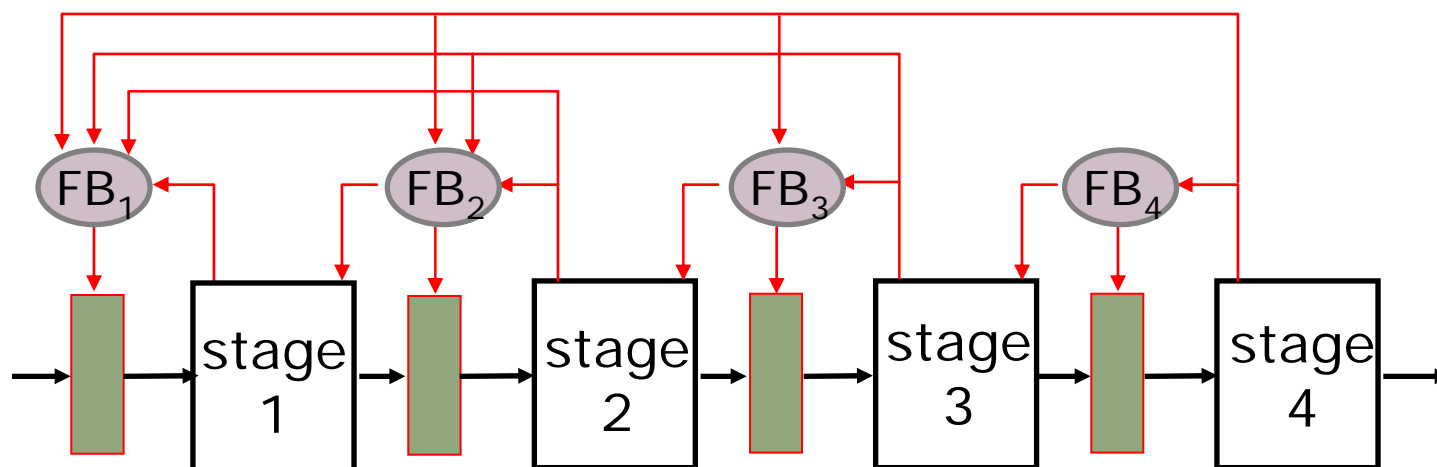
Resolving Data Hazards

Freeze earlier pipeline stages until the data becomes available \Rightarrow *interlocks*

If data is available somewhere in the datapath provide a *bypass* to get it to the right stage

Speculate about the hazard resolution and *kill* the instruction later if the speculation is wrong.

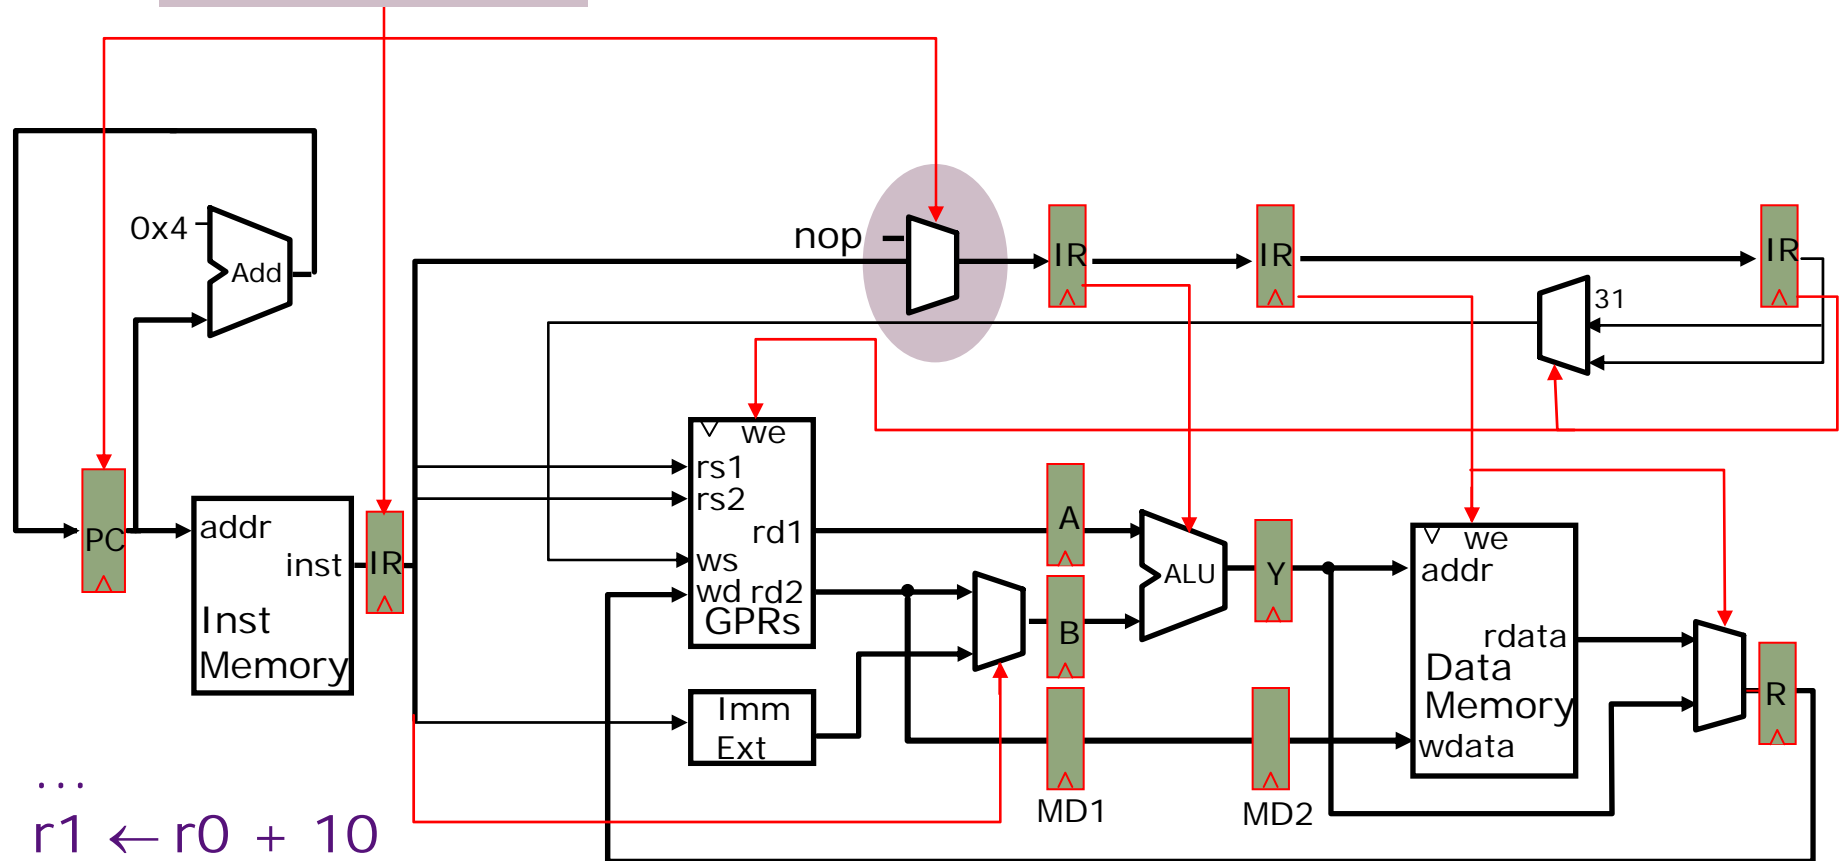
Feedback to Resolve Hazards



- Detect a hazard and provide feedback to previous stages to *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)



Interlocks to resolve Data Hazards

Stall Condition



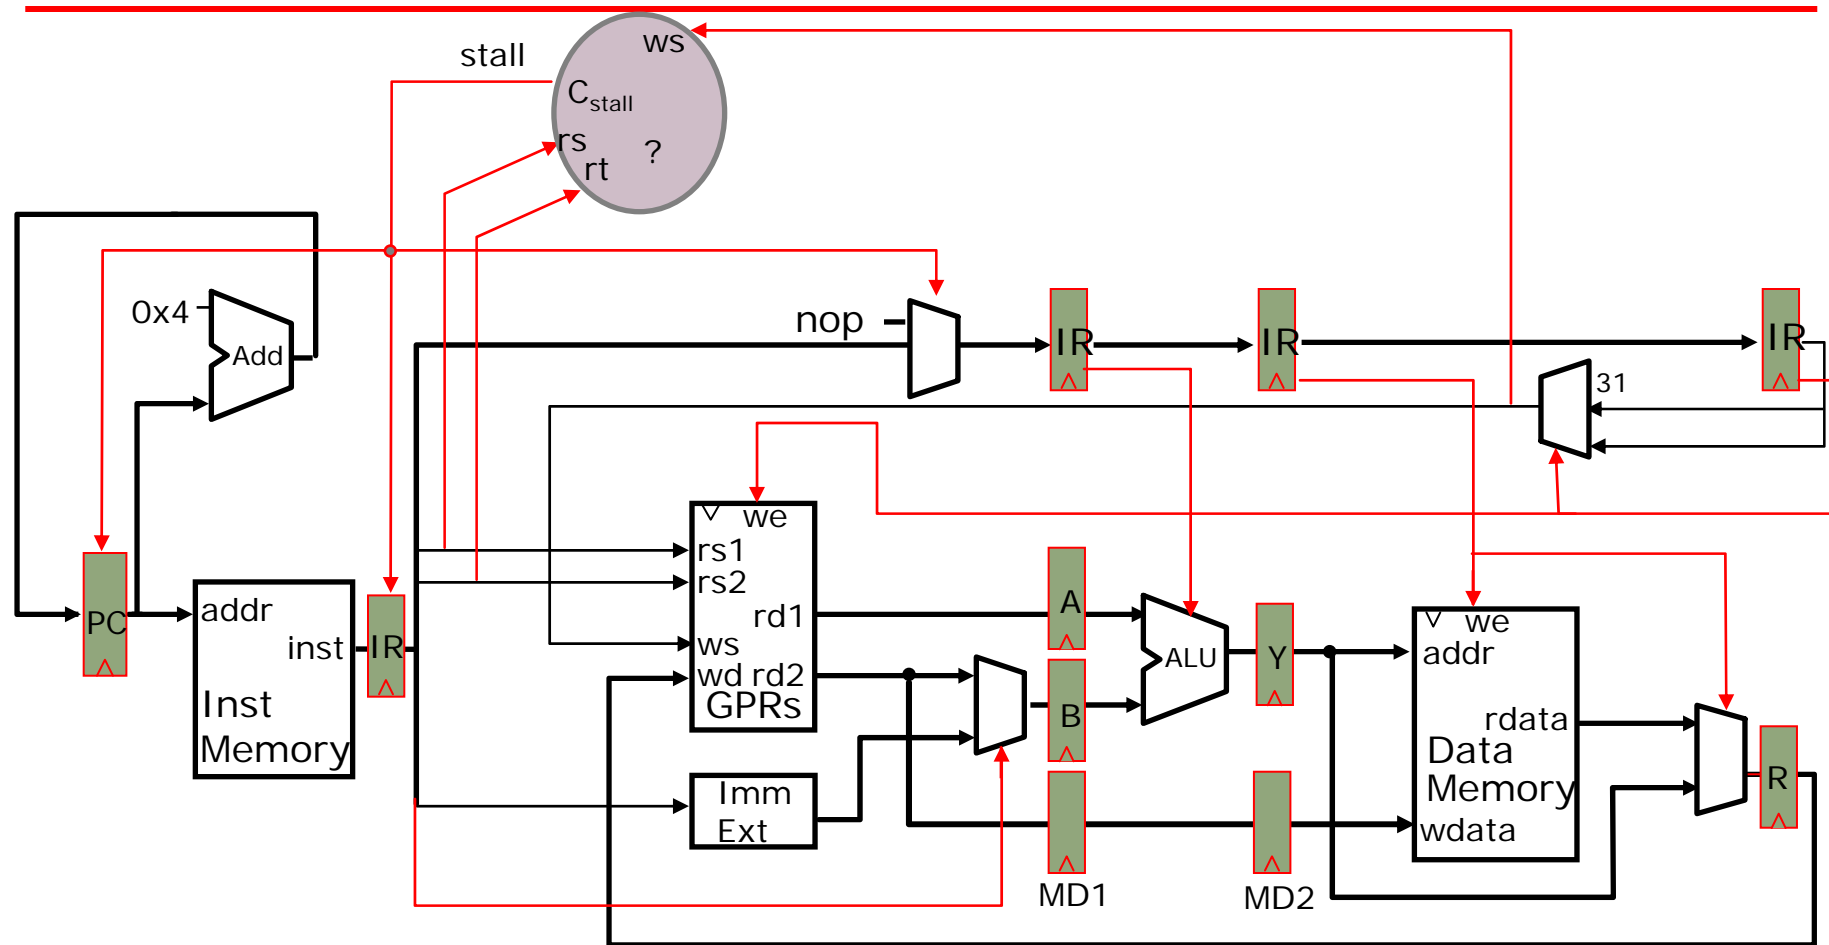
...
 $r1 \leftarrow r0 + 10$
 $r4 \leftarrow r1 + 17$
 ...

Stalled Stages and Pipeline Bubbles



nop ⇒ *pipeline bubble*

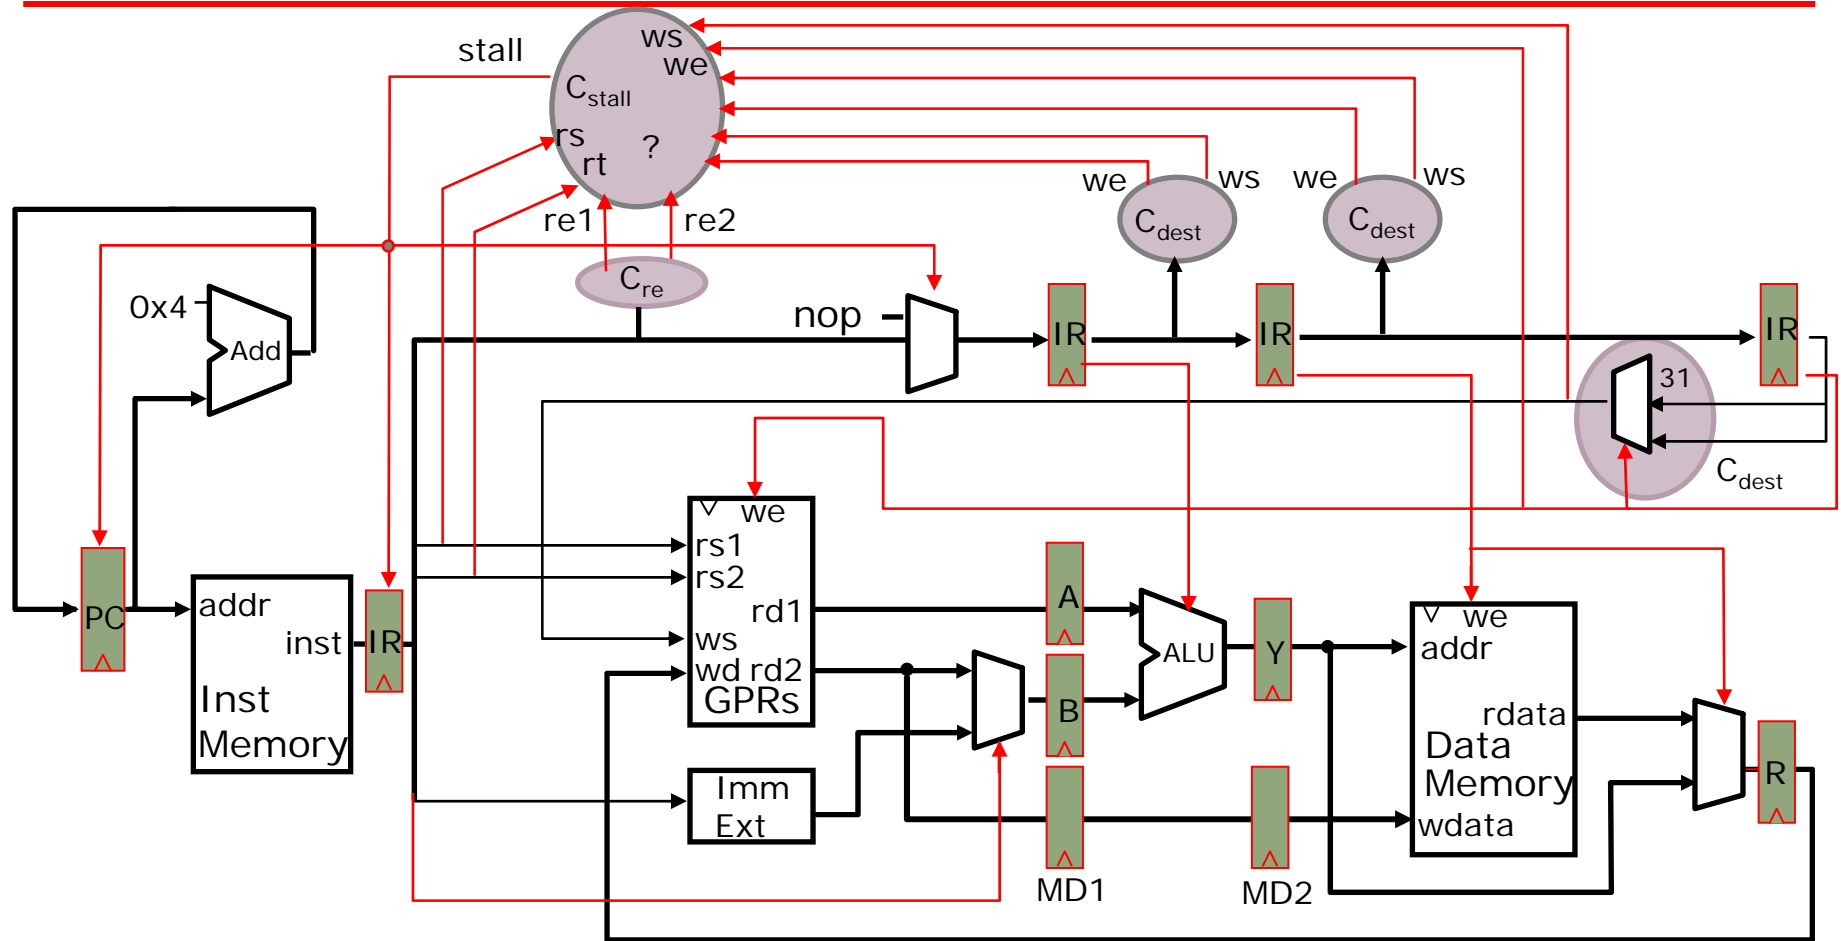
Interlock Control Logic



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

Interlocks Control Logic

ignoring jumps & branches



Should we always stall if the rs field matches some rd?
 not every instruction writes a register \Rightarrow we
 not every instruction reads a register \Rightarrow re

Source & Destination Registers

R-type:

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

I-type:

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

J-type:

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

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

Deriving the Stall Signal

C_{dest}

ws = *Case opcode*
 ALU \Rightarrow rd
 ALUi, LW \Rightarrow rt
 JAL, JALR \Rightarrow R31

we = *Case opcode*
 ALU, ALUi, LW \Rightarrow (ws \neq 0)
 JAL, JALR \Rightarrow on
 ... \Rightarrow off

C_{re}

re1 = *Case opcode*
 ALU, ALUi,
 LW, SW, BZ,
 JR, JALR \Rightarrow on
 J, JAL \Rightarrow off

re2 = *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

What if
 $(r1)+7 = (r3)+5$?



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

*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 !*

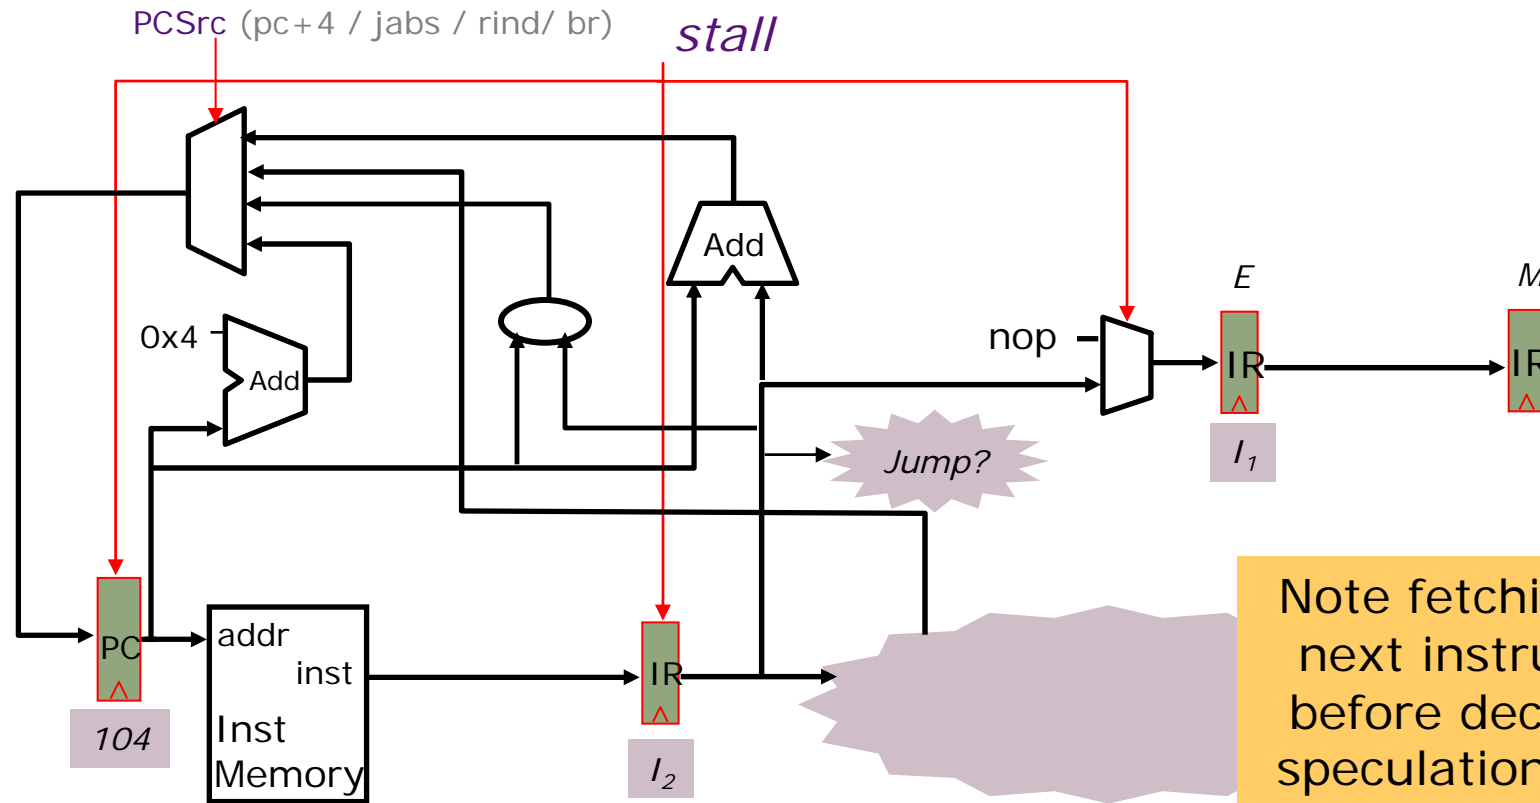
Load/Store hazards, even when they do exist, are often resolved in the memory system itself.

More on this later in the course.



Five-minute break to stretch your legs

Complications due to Jumps

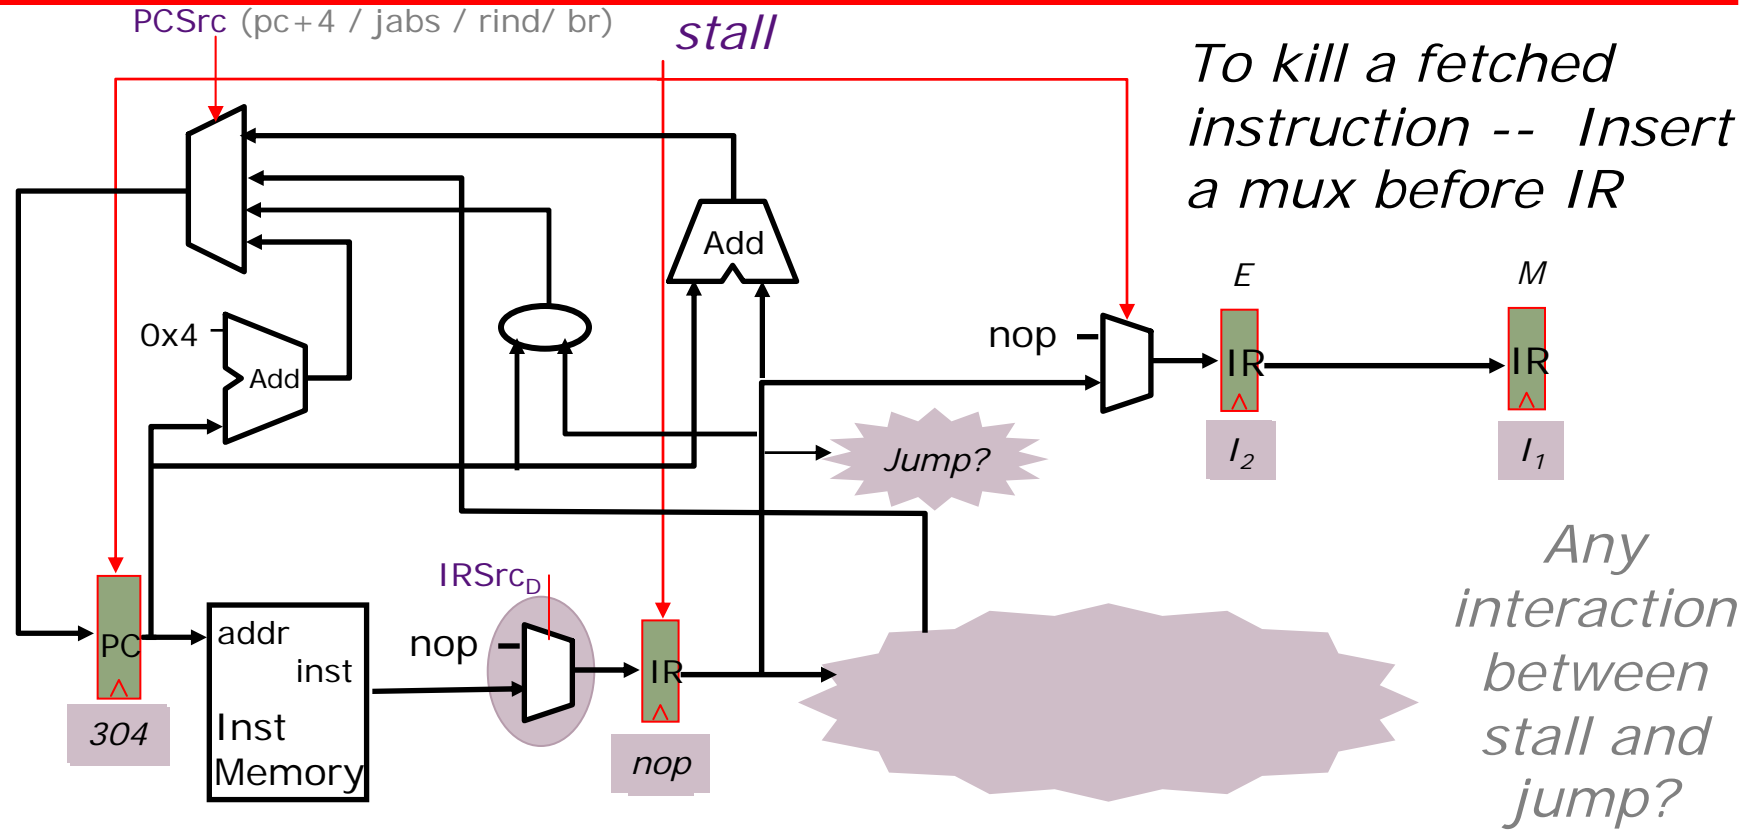


I_1	096	ADD	
I_2	100	J	200
I_3	104	ADD	<i>kill</i>
I_4	304	ADD	

A jump instruction kills (not stalls) the following instruction

How?



Pipelining Jumps



I ₁	096	ADD	
I ₂	100	J	200
I ₃	104	ADD	<i>kill</i>
I ₄	304	ADD	

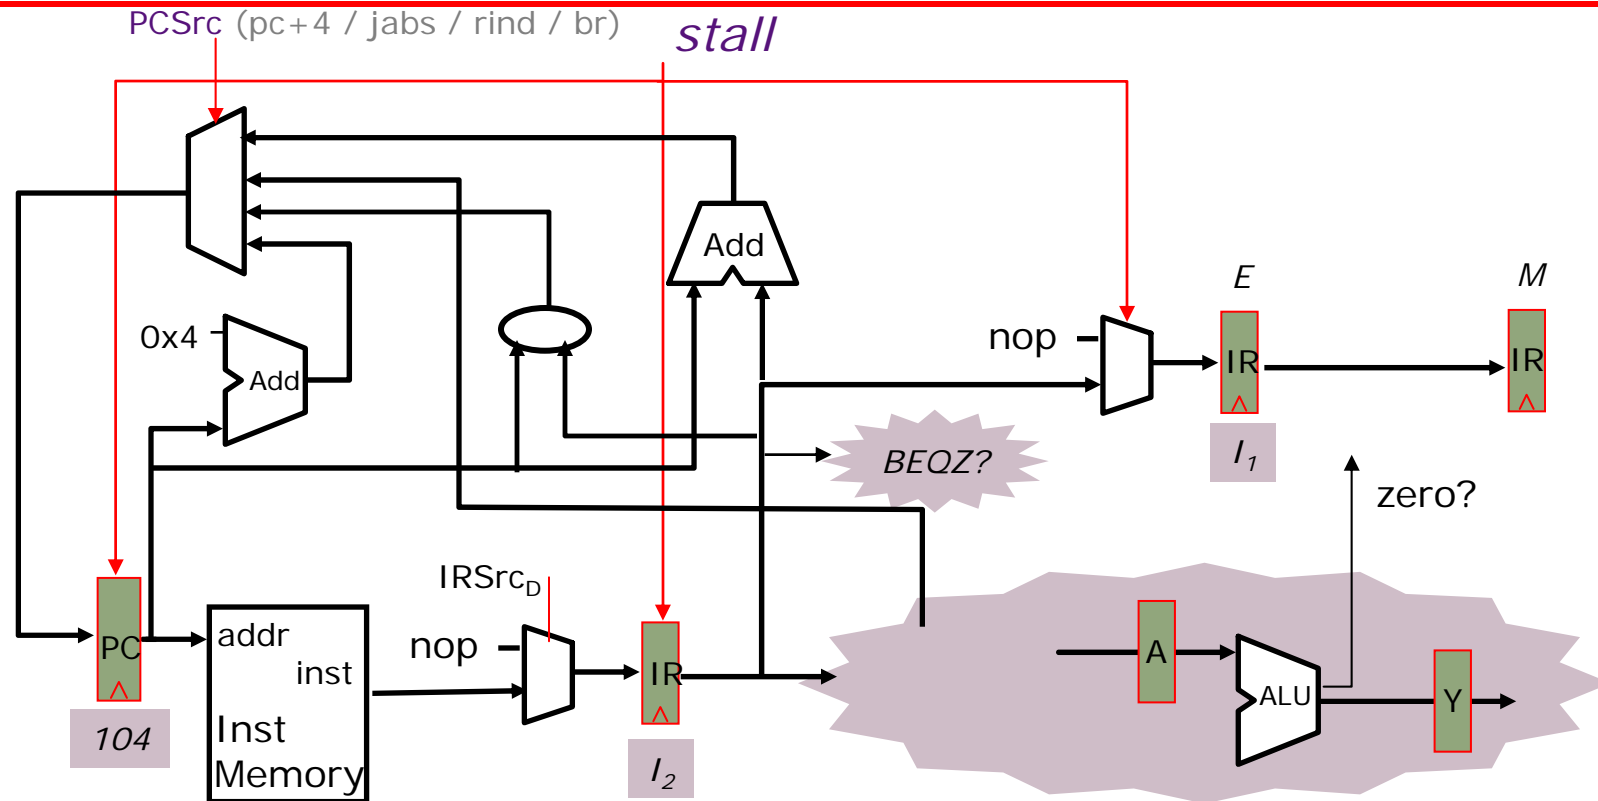
IRSrc_D = Case opcode_D
 J, JAL ⇒ nop
 ... ⇒ IM

Jump Pipeline Diagrams



nop ⇒ *pipeline bubble*

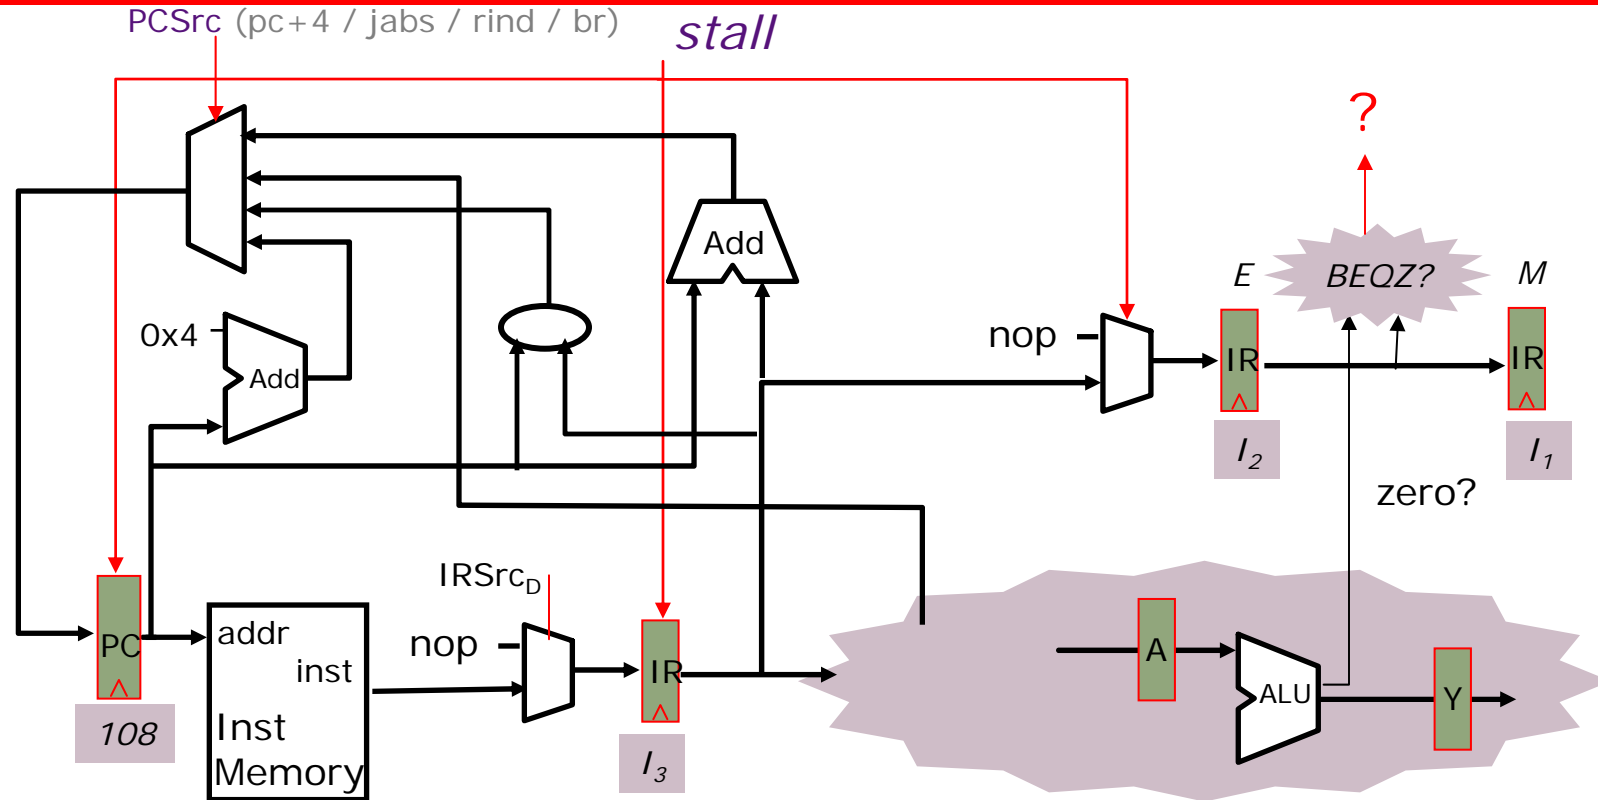
Pipelining Conditional Branches



I_1	096	ADD
I_2	100	BEQZ r1 200
I_3	104	ADD
I_4	304	ADD

Branch condition is not known until the execute stage
what action should be taken in the decode stage ?

Pipelining Conditional Branches




If the branch is taken

- kill the two following instructions
- the instruction at the decode stage is not valid

⇒ *stall signal is not valid*

I ₁	096	ADD
I ₂	100	BEQZ r1 200
I ₃	104	ADD
I ₄	304	ADD

Pipelining Conditional Branches



If the branch is taken

- kill the two following instructions
- the instruction at the decode stage is not valid

⇒ *stall signal is not valid*

I_1	096	ADD
I_2	100	BEQZ r1 200
I_3	104	ADD
I_4	304	ADD

New Stall Signal

$$\text{stall} = (((rs_D = ws_E).we_E + (rs_D = ws_M).we_M + (rs_D = ws_W).we_W).re1_D \\ + ((rt_D = ws_E).we_E + (rt_D = ws_M).we_M + (rt_D = ws_W).we_W).re2_D \\) . !((opcode_E = BEQZ).z + (opcode_E = BNEZ).!z)$$

Don't stall if the branch is taken. Why?

Instruction at the decode stage is invalid

Control Equations for PC and IR Muxes

$$\text{PCSrc} = \text{Case opcode}_E$$

BEQZ.z, BNEZ.!z	\Rightarrow br
...	\Rightarrow
	Case opcode _D
J, JAL	\Rightarrow jabs
JR, JALR	\Rightarrow rind
...	\Rightarrow pc+4

Give priority to the older instruction, i.e., execute stage instruction over decode stage instruction



$$\text{IRSrc}_D = \text{Case opcode}_E$$

BEQZ.z, BNEZ.!z	\Rightarrow nop
...	\Rightarrow
	Case opcode _D
J, JAL, JR, JALR	\Rightarrow nop
...	\Rightarrow IM

$$\text{IRSrc}_E = \text{Case opcode}_E$$

BEQZ.z, BNEZ.!z	\Rightarrow nop
...	\Rightarrow stall.nop + !stall.IR _D

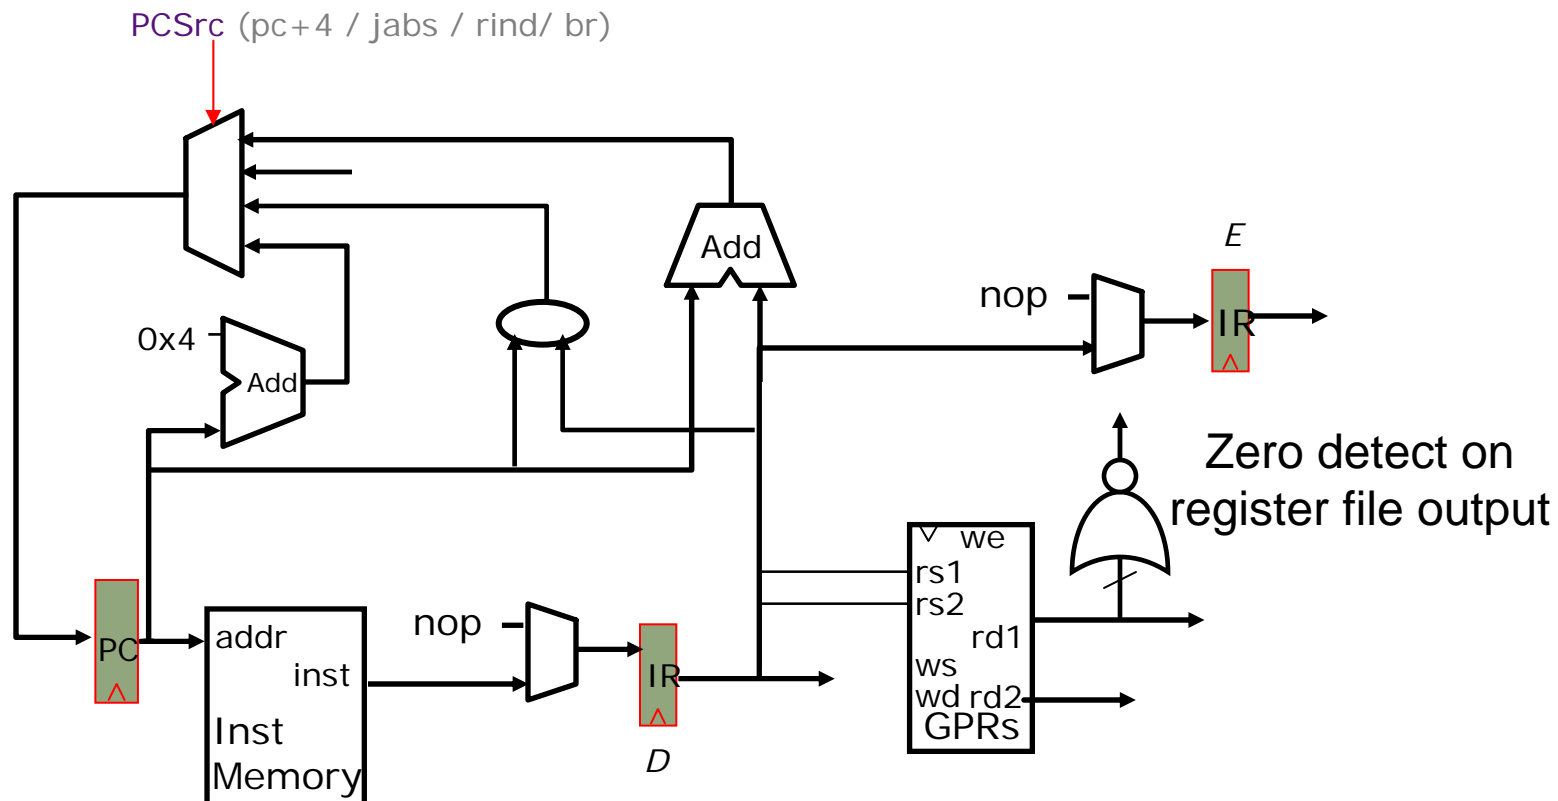
Branch Pipeline Diagrams (resolved in execute stage)



nop ⇒ *pipeline bubble*

Reducing Branch Penalty (resolve in decode stage)

- One pipeline bubble can be removed if an extra comparator is used in the Decode stage



Pipeline diagram now same as for jumps

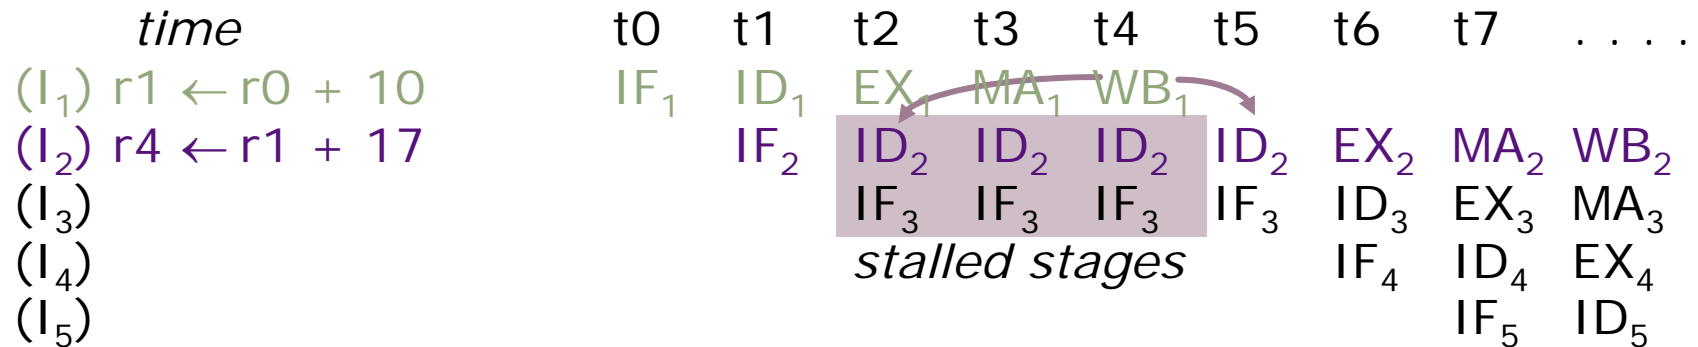
Branch Delay Slots (expose control hazard to software)

- Change the ISA semantics so that the instruction that follows a jump or branch is always executed
 - gives compiler the flexibility to put in a useful instruction where normally a pipeline bubble would have resulted.

I ₁	096	ADD	
I ₂	100	BEQZ r1 200	<i>Delay slot instruction</i>
I ₃	104	ADD	← <i>executed regardless of</i>
I ₄	304	ADD	<i>branch outcome</i>


- Other techniques include branch prediction, which can dramatically reduce the branch penalty... *to come later*

Bypassing

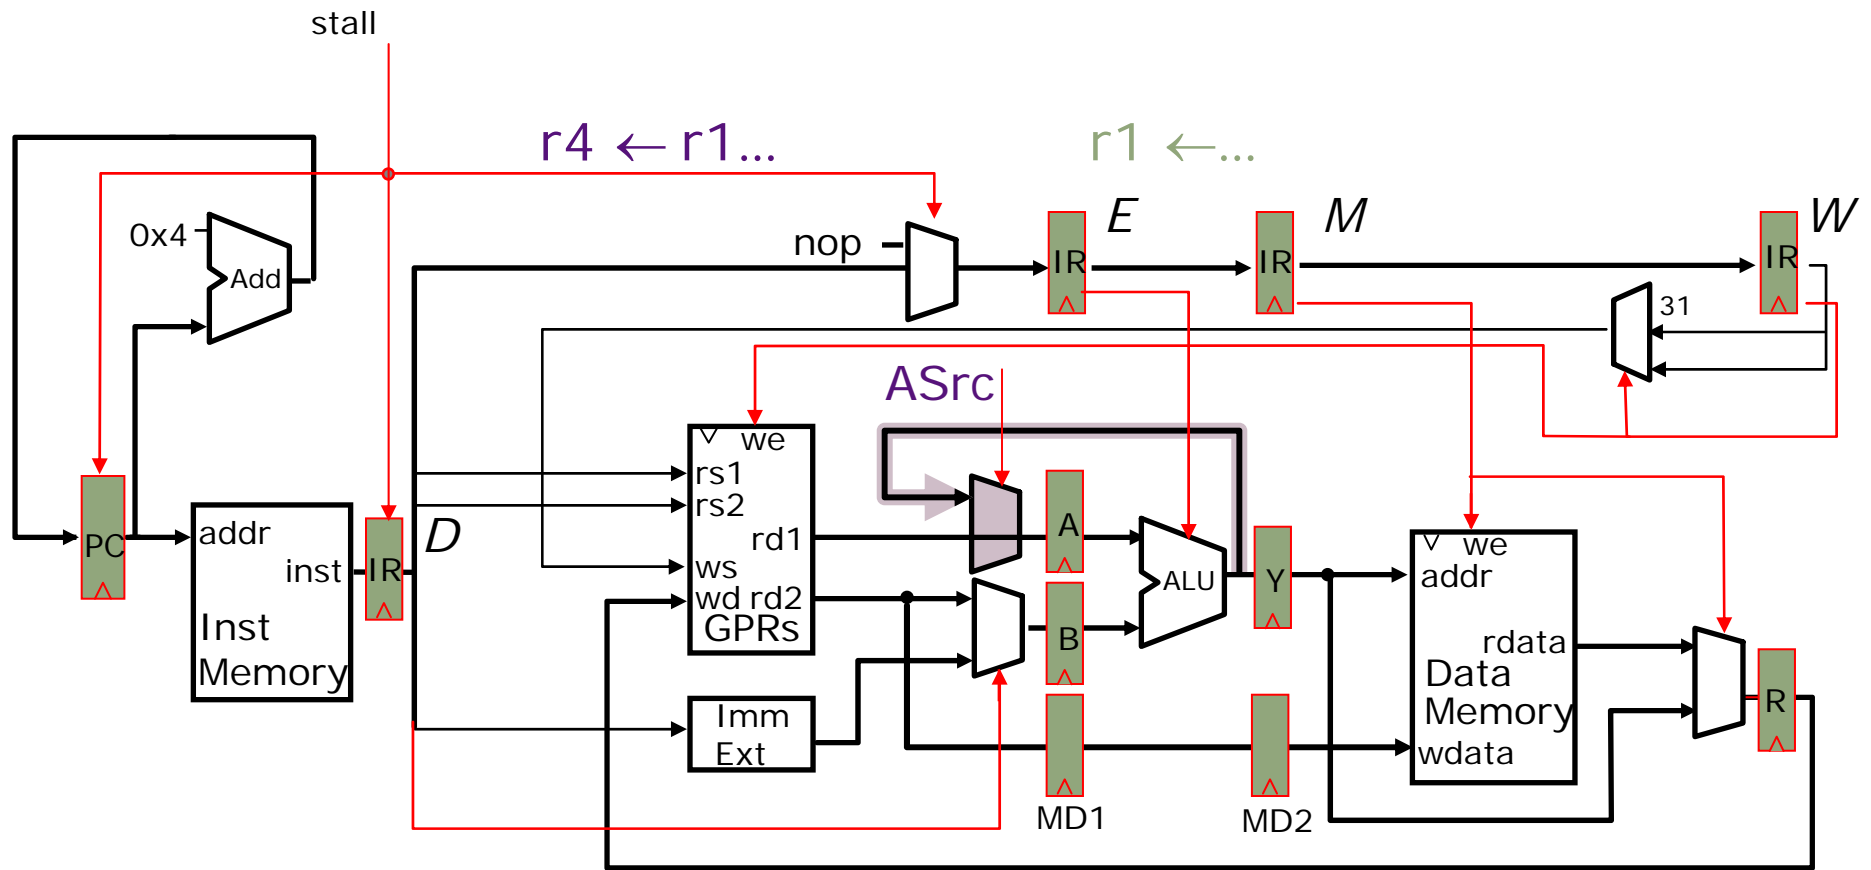


Each *stall or kill* introduces a bubble in the pipeline
 $\Rightarrow CPI > 1$

A new datapath, i.e., a *bypass*, can get the data from the output of the ALU to its input



Adding a Bypass



When does this bypass help?

...
 $(I_1) \quad r1 \leftarrow r0 + 10$
 $(I_2) \quad r4 \leftarrow r1 + 17$
yes

$r1 \leftarrow M[r0 + 10]$
 $r4 \leftarrow r1 + 17$
no

JAL 500
 $r4 \leftarrow r31 + 17$
no

The Bypass Signal

Deriving it from the Stall Signal

$$\text{stall} = (\overline{(\text{rs}_D = \text{ws}_E)} \cdot \text{we}_E + (\text{rs}_D = \text{ws}_M) \cdot \text{we}_M + (\text{rs}_D = \text{ws}_W) \cdot \text{we}_W) \cdot \text{re1}_D \\ + ((\text{rt}_D = \text{ws}_E) \cdot \text{we}_E + (\text{rt}_D = \text{ws}_M) \cdot \text{we}_M + (\text{rt}_D = \text{ws}_W) \cdot \text{we}_W) \cdot \text{re2}_D$$

ws = Case opcode
 ALU ⇒ rd
 ALUi, LW ⇒ rt
 JAL, JALR ⇒ R31

we = Case opcode
 ALU, ALUi, LW ⇒ (ws ≠ 0)
 JAL, JALR ⇒ on
 ... ⇒ off

$$\text{ASrc} = (\text{rs}_D = \text{ws}_E) \cdot \text{we}_E \cdot \text{re1}_D$$

Is this correct?

No because only ALU and ALUi instructions can benefit from this bypass

Split we_E into two components: we-bypass, we-stall

Bypass and Stall Signals

Split we_E into two components: we-bypass, we-stall

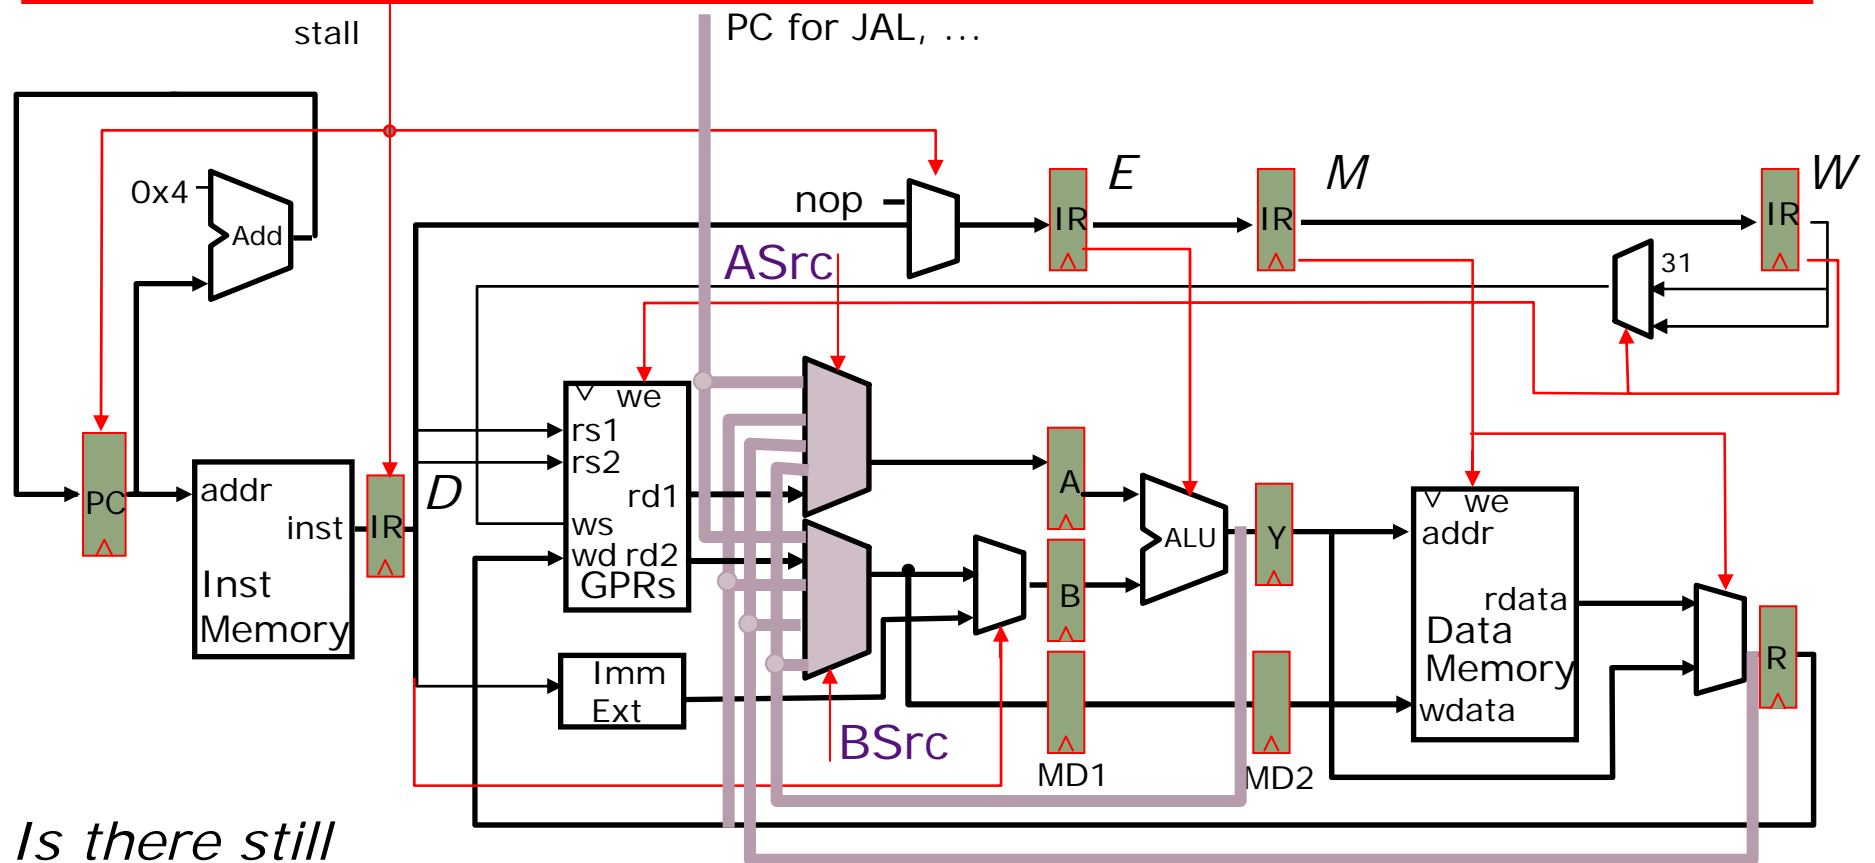
$we_bypass_E = \text{Case opcode}_E$ ALU, ALUi $\Rightarrow (ws \neq 0)$... $\Rightarrow \text{off}$
--

$we_stall_E = \text{Case opcode}_E$ LW $\Rightarrow (ws \neq 0)$ JAL, JALR $\Rightarrow \text{on}$... $\Rightarrow \text{off}$

$ASrc = (rs_D = ws_E) \cdot we_bypass_E \cdot re1_D$

$stall = ((rs_D = ws_E) \cdot we_stall_E +$ $(rs_D = ws_M) \cdot we_M + (rs_D = ws_W) \cdot we_W) \cdot re1_D$ $+ ((rt_D = ws_E) \cdot we_E + (rt_D = ws_M) \cdot we_M + (rt_D = ws_W) \cdot we_W) \cdot re2_D$
--

Fully Bypassed Datapath



*Is there still
a need for the
stall signal ?*

$$\text{stall} = (rs_D = ws_E) \cdot (\text{opcode}_E = LW_E) \cdot (ws_E \neq 0) \cdot re1_D \\ + (rt_D = ws_E) \cdot (\text{opcode}_E = LW_E) \cdot (ws_E \neq 0) \cdot re2_D$$

Why an Instruction may not be dispatched every cycle (CPI > 1)

- Full bypassing may be too expensive to implement
 - typically all frequently used paths are provided
 - some infrequently used bypass paths may increase cycle time and counteract the benefit of reducing CPI
- Loads have two cycle latency
 - Instruction after load cannot use load result
 - MIPS-I ISA defined *load delay slots*, a software-visible pipeline hazard (compiler schedules independent instruction or inserts NOP to avoid hazard). Removed in MIPS-II.
- Conditional branches may cause bubbles
 - kill following instruction(s) if no delay slots

Machines with software-visible delay slots may execute significant number of NOP instructions inserted by the compiler.



Thank you !