

Architecture of Computer Systems


Lecture 4 - Pipelining



Last time in Lecture 3

- Microcoding became less attractive as gap between RAM and ROM speeds reduced
- Complex instruction sets difficult to pipeline, so difficult to increase performance as gate count grew
- Load-Store RISC ISAs designed for efficient pipelined implementations
 - Very similar to vertical microcode
 - Inspired by earlier Cray machines (more on these later)
- Iron Law explains architecture design space
 - Trade instructions/program, cycles/instruction, and time/cycle

An Ideal Pipeline



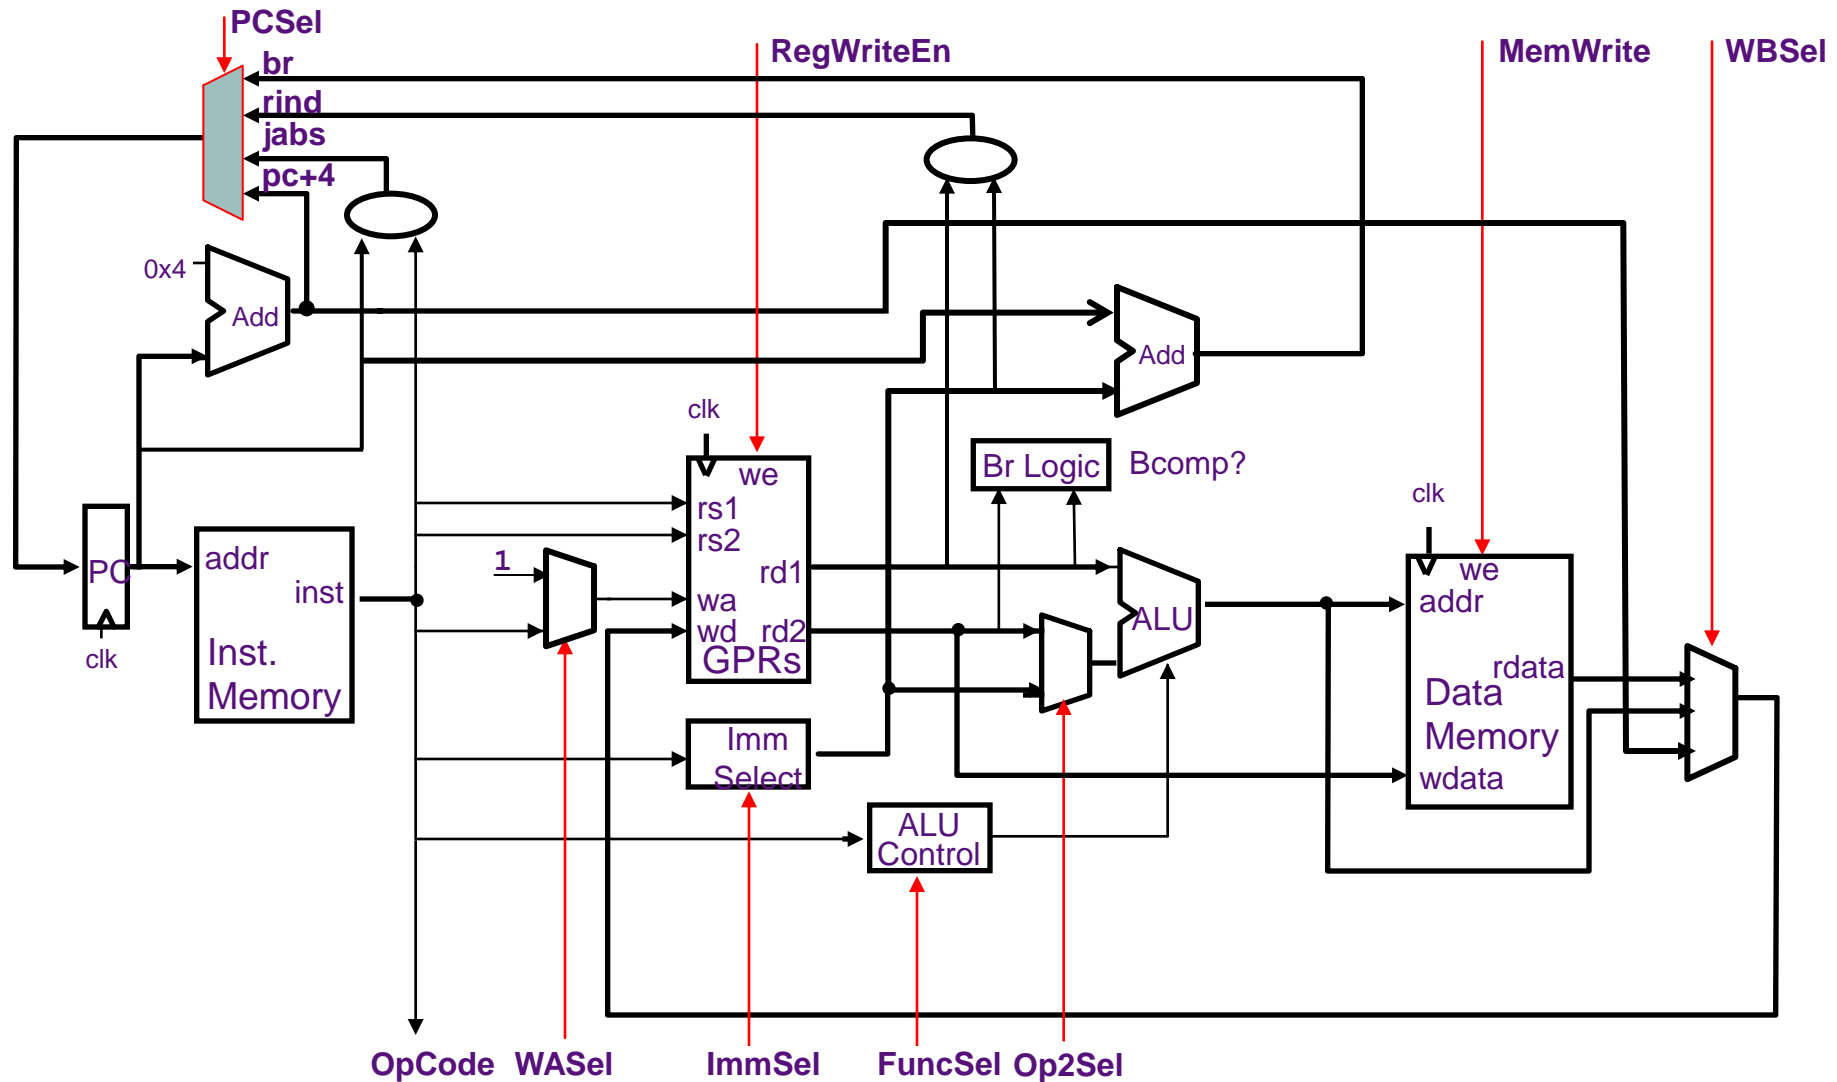
- All objects go through the same stages
- No sharing of resources between any two stages
- Propagation delay through all pipeline stages is equal
- The scheduling of an object entering the pipeline is not affected by the objects in other stages

These conditions generally hold for industrial assembly lines, but instructions depend on each other!

Pipelined RISC-V

- To pipeline RISC-V:
- First build RISC-V without pipelining with $CPI=1$
- Next, add pipeline registers to reduce cycle time while maintaining $CPI=1$

Lecture 3: Unpipelined Datapath for RISC-V



Lecture 3: Hardwired Control

Table

OpCode	ImmSel	Op2Sel	FuncSel	MemWr	RFWen	WBSel	WASel	PCSel
ALU	*	Reg	Func	no	yes	ALU	rd	pc+4
ALUi	IType ₁₂	Imm	Op	no	yes	ALU	rd	pc+4
LW	IType ₁₂	Imm	+	no	yes	Mem	rd	pc+4
SW	BsType ₁₂	Imm	+	yes	no	*	*	pc+4
BEQ _{true}	BrType ₁₂	*	*	no	no	*	*	br
BEQ _{false}	BrType ₁₂	*	*	no	no	*	*	pc+4
J	*	*	*	no	no	*	*	jabs
JAL	*	*	*	no	yes	PC	X1	jabs
JALR	*	*	*	no	yes	PC	rd	rind

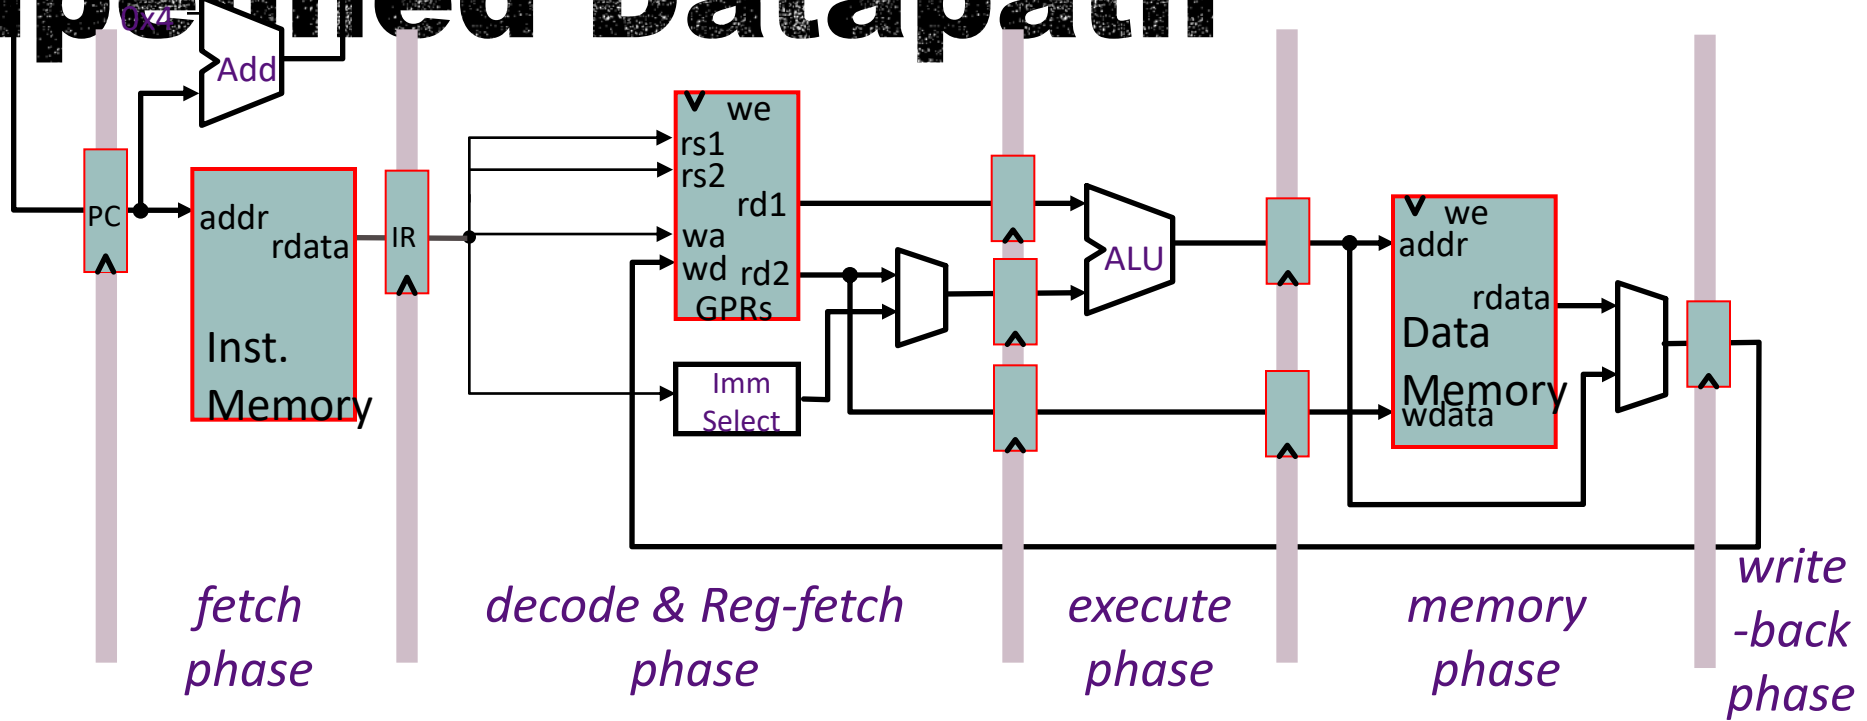
Op2Sel = Reg / Imm

WASel = rd / X1

WBSel = ALU / Mem / PC

PCSel = pc+4 / br / rind / jabs

Pipelined Datapath



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

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

However, CPI will increase unless instructions are pipelined

“Iron Law” of Processor Performance


$$\text{Performance} = \frac{\text{Instructions}}{\text{Cycles}} \times \frac{\text{Cycles}}{\text{Time}}$$

$\frac{\text{Instructions}}{\text{Program}} * \text{Instruction} * \text{Cycle}$


- Instructions per program depends on source code, compiler technology, and ISA
- Cycles per instructions (CPI) depends on ISA and μ architecture
- Time per cycle depends upon the μ architecture and base technology

	Microarchitecture	CPI	cycle time
Lecture 2	Microcoded	>1	short
Lecture 3	Single-cycle unpipelined	1	long
Lecture 4	Pipelined	1	short


CPI Examples



Unpipelined machine



Pipelined machine



Technology Assumptions

A small amount of very fast memory (caches)
backed up by a large, slower memory

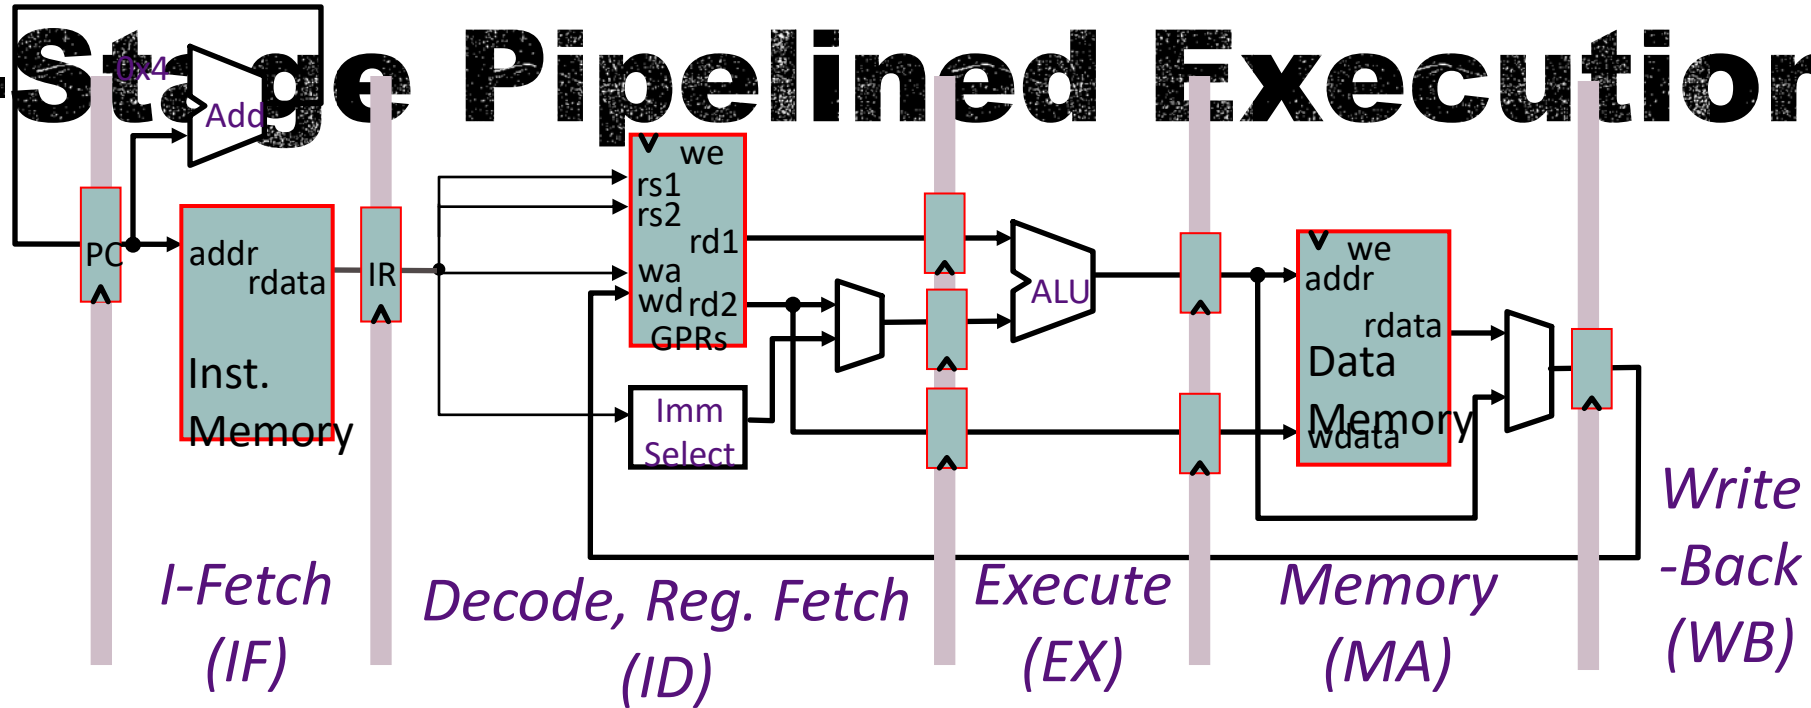
- Fast ALU (at least for integers)
- Multiported Register files (slower!)

Thus, the following timing assumption is reasonable

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

A 5-stage pipeline will be focus of our detailed design
- *some commercial designs have over 30 pipeline stages to do an integer add!*

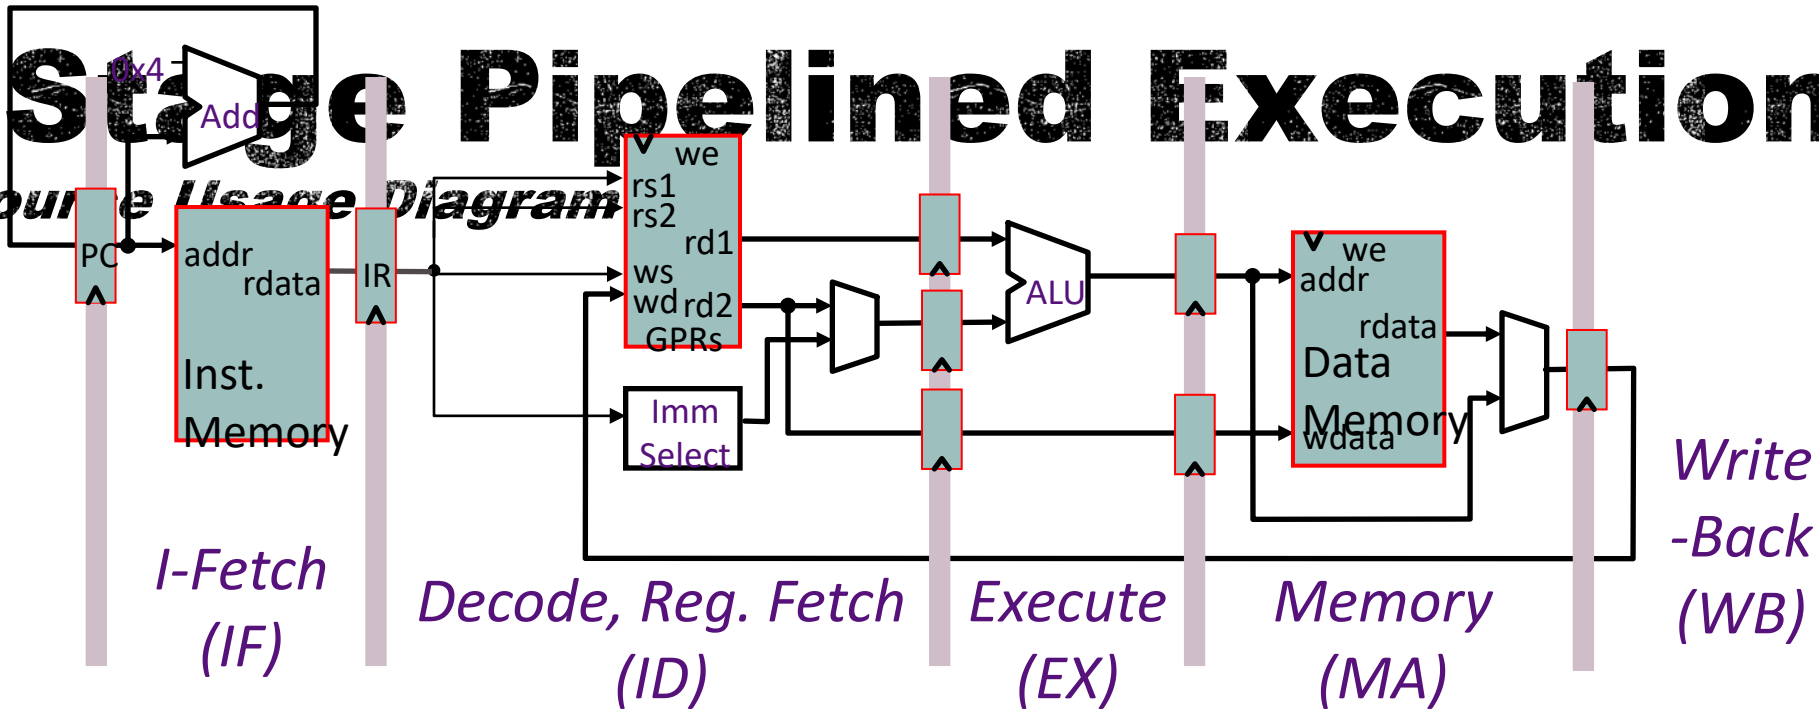
5-Stage Pipelined Execution



time	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

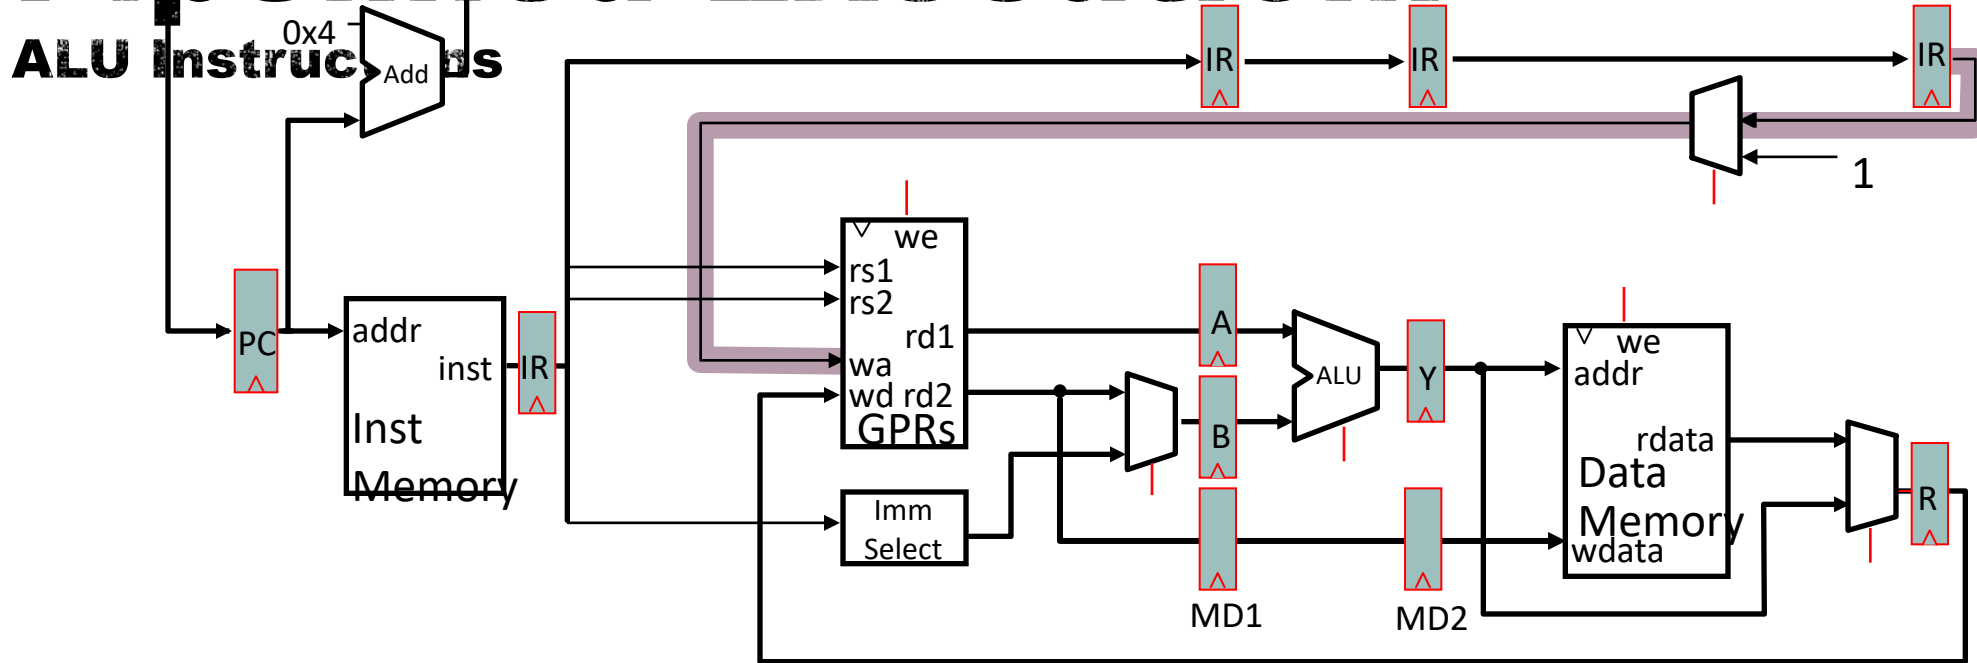
5-Stage Pipelined Execution

Resource Usage Diagram



time	t0	t1	t2	t3	t4	t5	t6	t7
IF	I_1	I_2	I_3	I_4	I_5				
ID		I_1	I_2	I_3	I_4	I_5			
EX			I_1	I_2	I_3	I_4	I_5		
MA				I_1	I_2	I_3	I_4	I_5	
WB					I_1	I_2	I_3	I_4	I_5

Pipelined Execution:

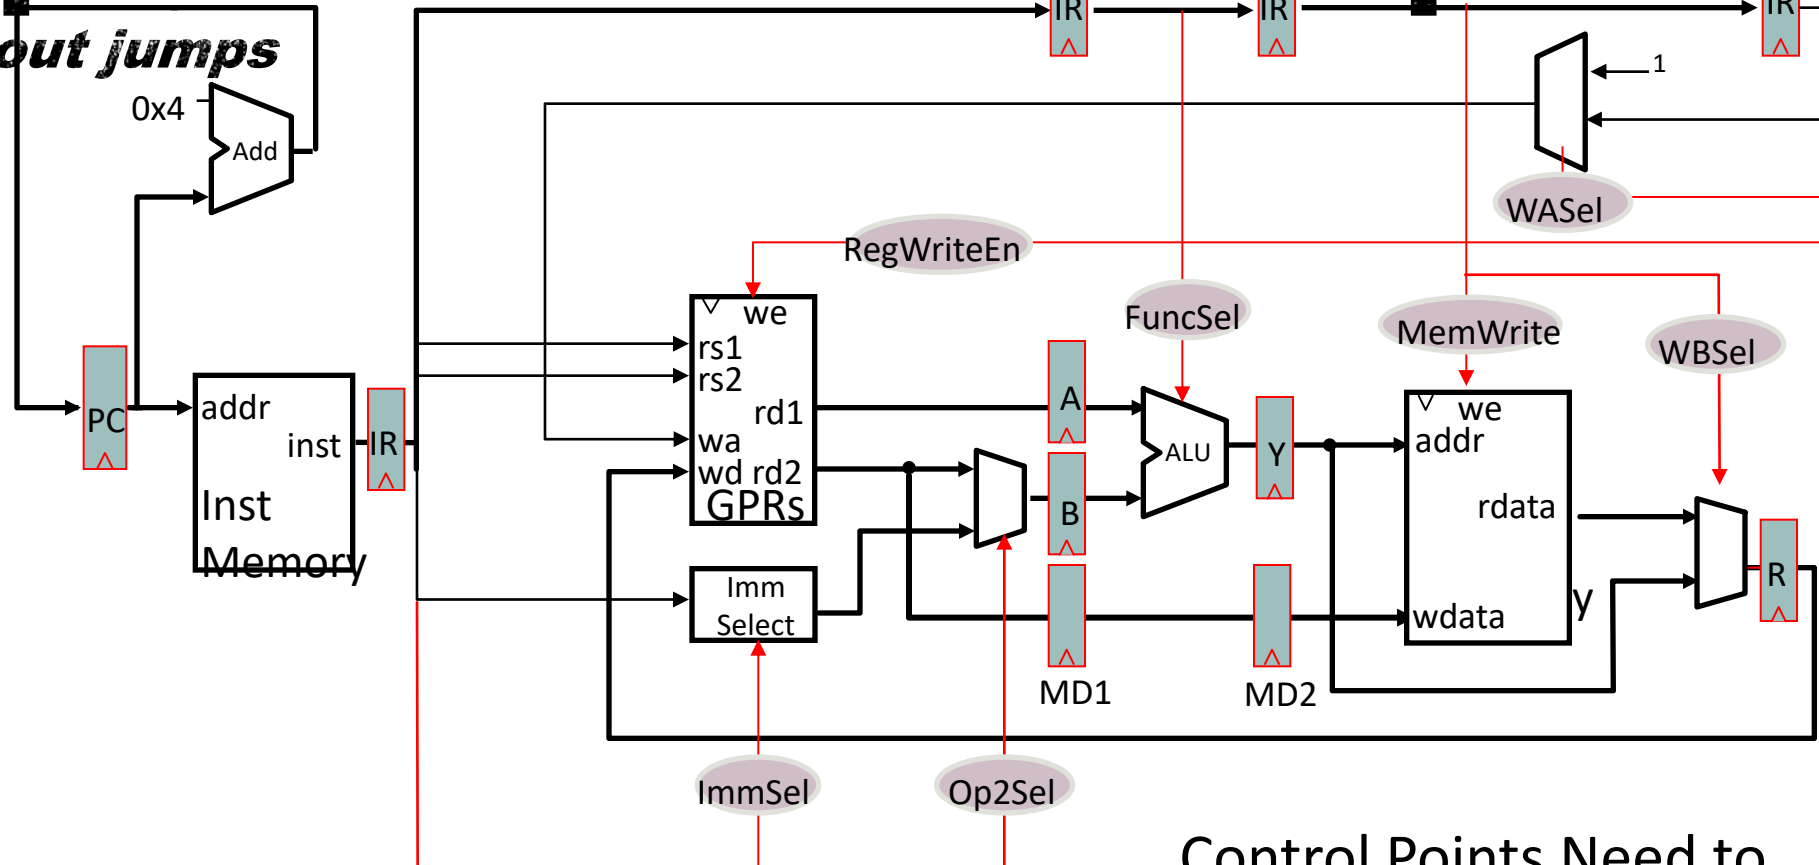


Not quite correct!

We need an Instruction Reg (IR) for each stage

Pipelined RISC-V Datapath

without jumps



Control Points Need to Be Connected

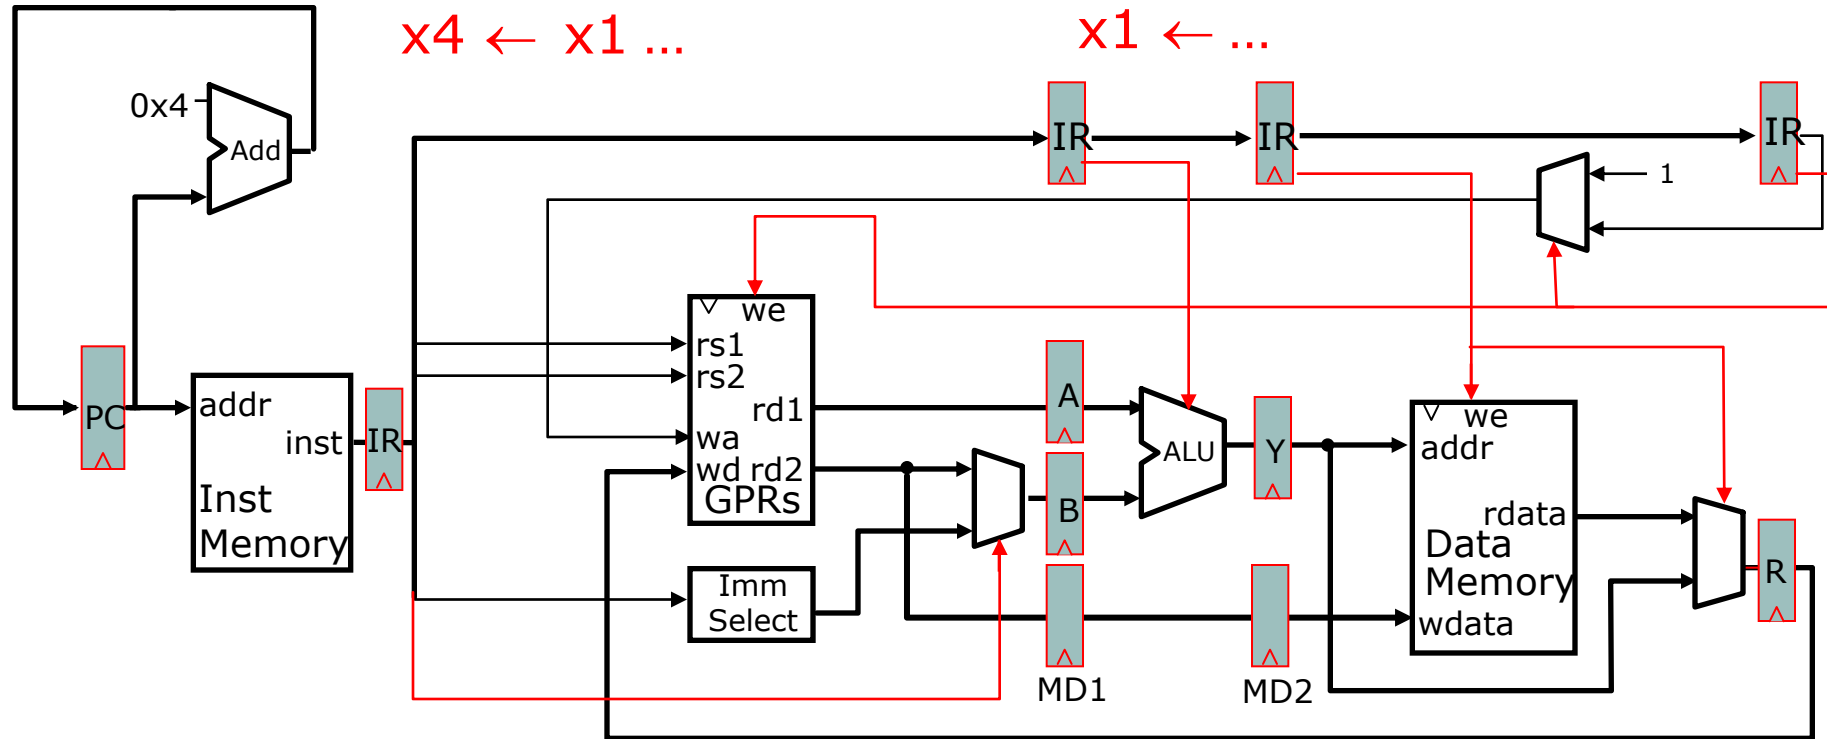
Instructions interact with each other in pipeline

- An instruction in the pipeline may need a resource being used by another instruction in the pipeline → *structural hazard*
- An instruction may depend on something produced by an earlier instruction
 - Dependence may be for a data value
→ *data hazard*
 - Dependence may be for the next instruction's address
→ *control hazard (branches, exceptions)*

Resolving Structural Hazards

- Structural hazard occurs when two instructions need same hardware resource at same time
 - Can resolve in hardware by stalling newer instruction till older instruction finished with resource
- A structural hazard can always be avoided by adding more hardware to design
 - E.g., if two instructions both need a port to memory at same time, could avoid hazard by adding second port to memory
- Our 5-stage pipeline has no structural hazards by design
 - Thanks to RISC-V ISA, which was designed for pipelining

Data Hazards



$x4 \leftarrow x1 \dots$

$x1 \leftarrow \dots$

...
 $x1 \leftarrow x0 + 10$
 $x4 \leftarrow x1 + 17$
 ...

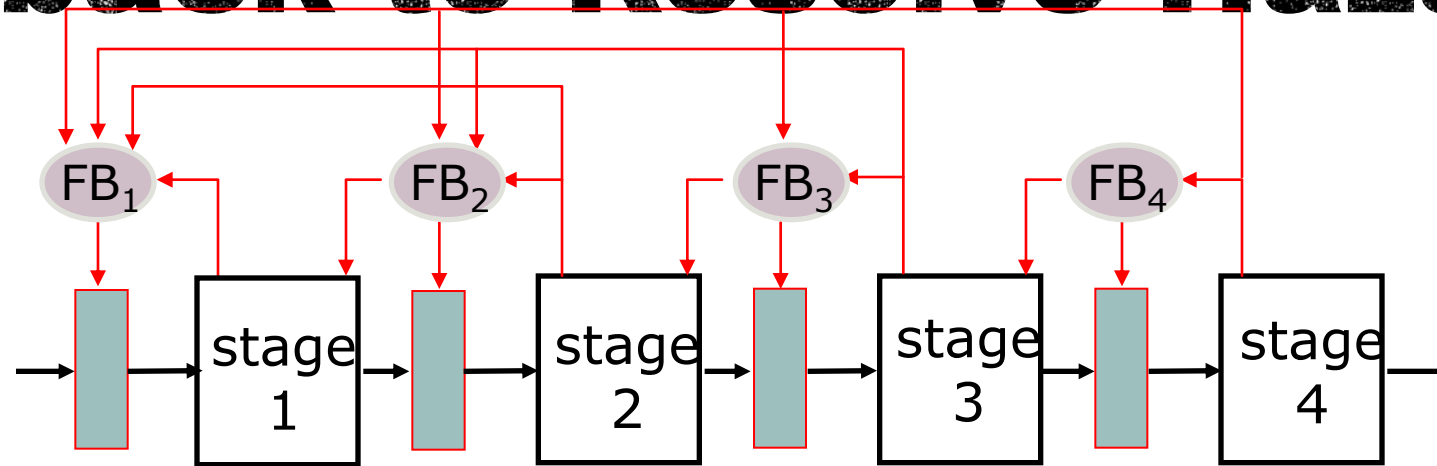
$x1$ is stale. Oops!

Resolving Data Hazards (1)

Strategy 1:

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

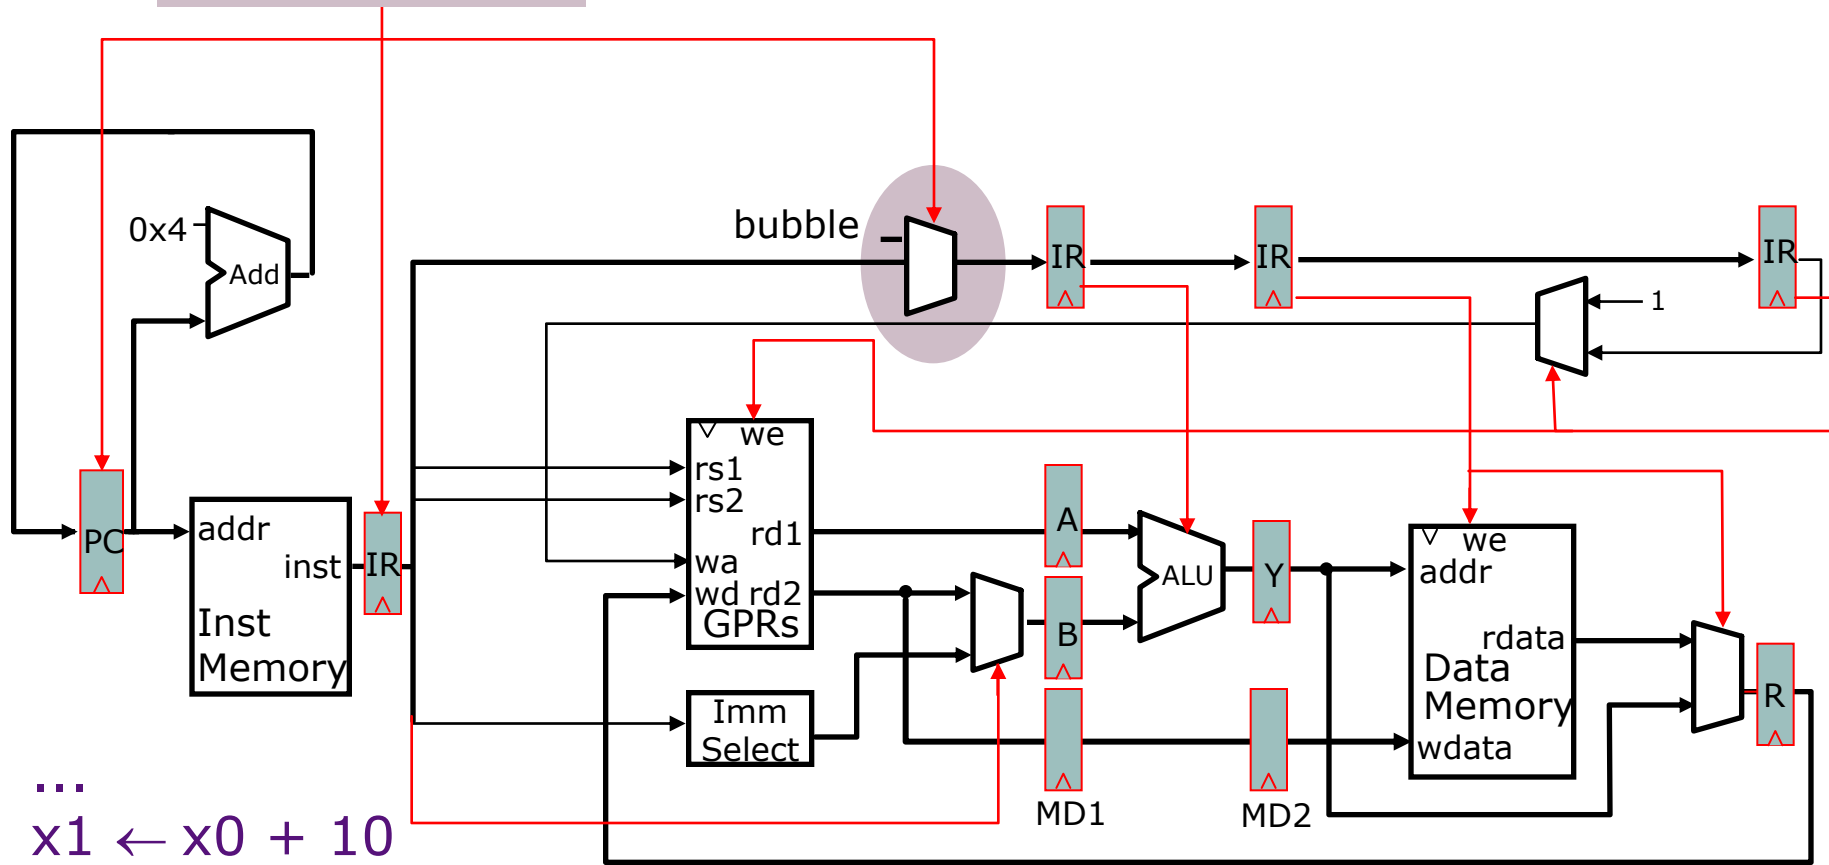
Feedback to Resolve Hazards



- 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)



Interlocks to resolve Data Hazards

Stall Condition

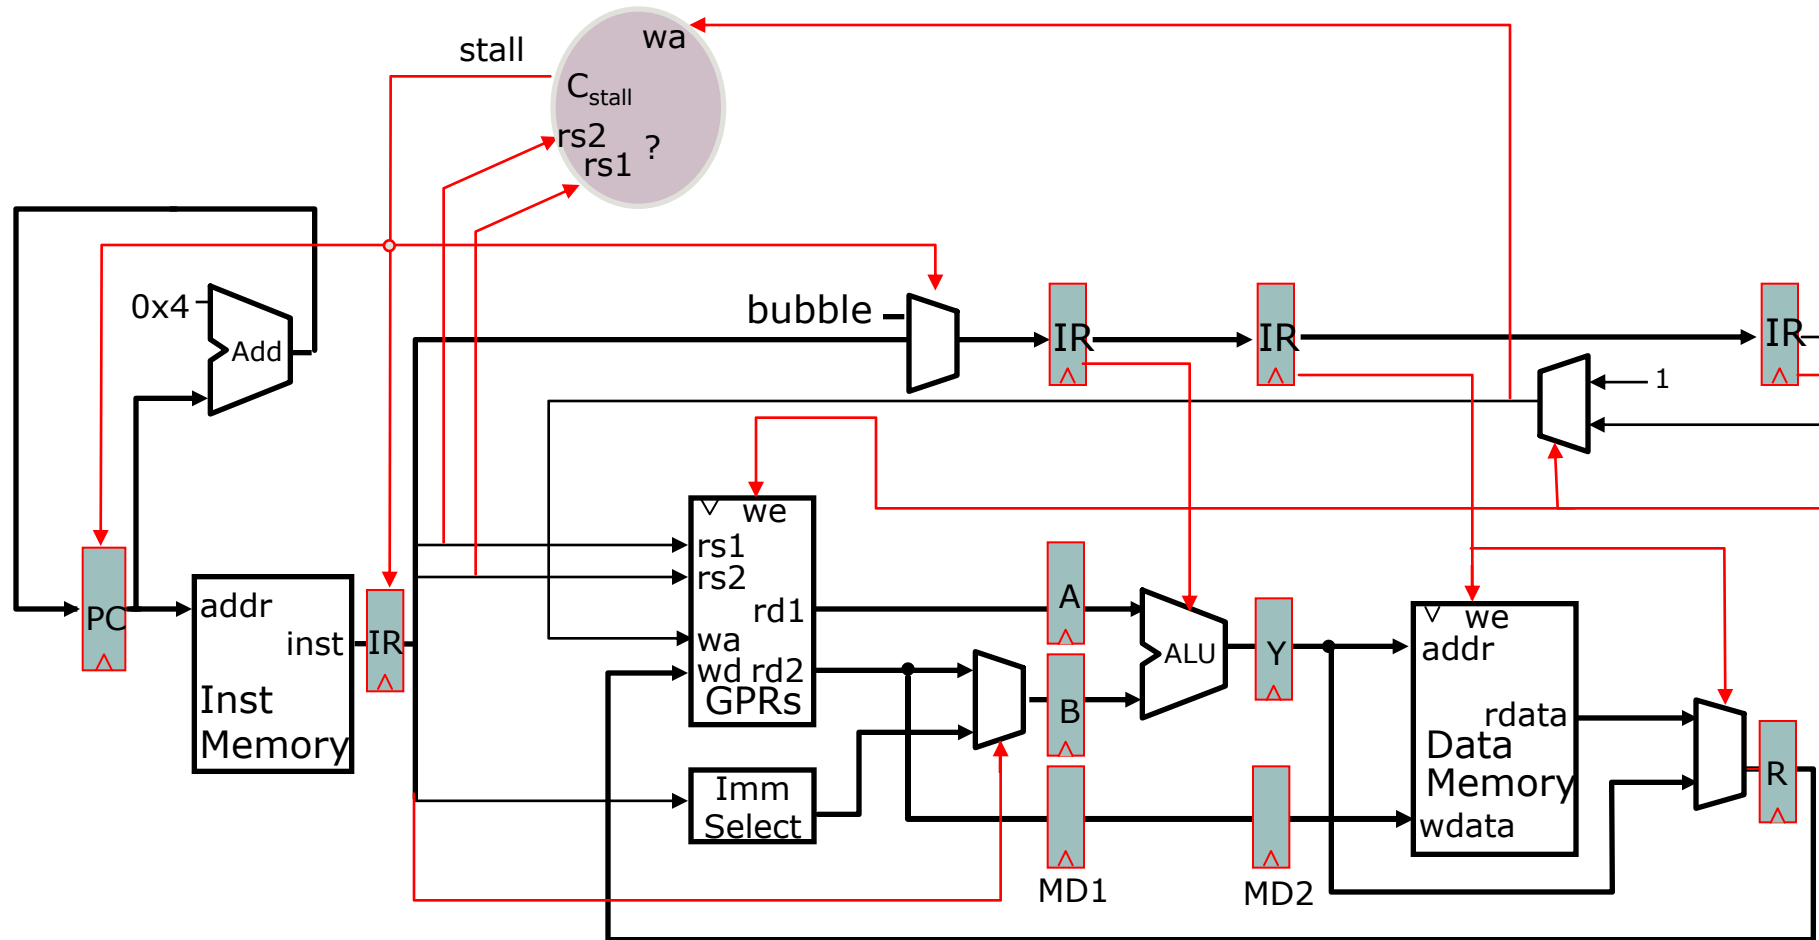


...
 $x1 \leftarrow x0 + 10$
 $x4 \leftarrow x1 + 17$
...

Stalled Stages and Pipeline Bubbles



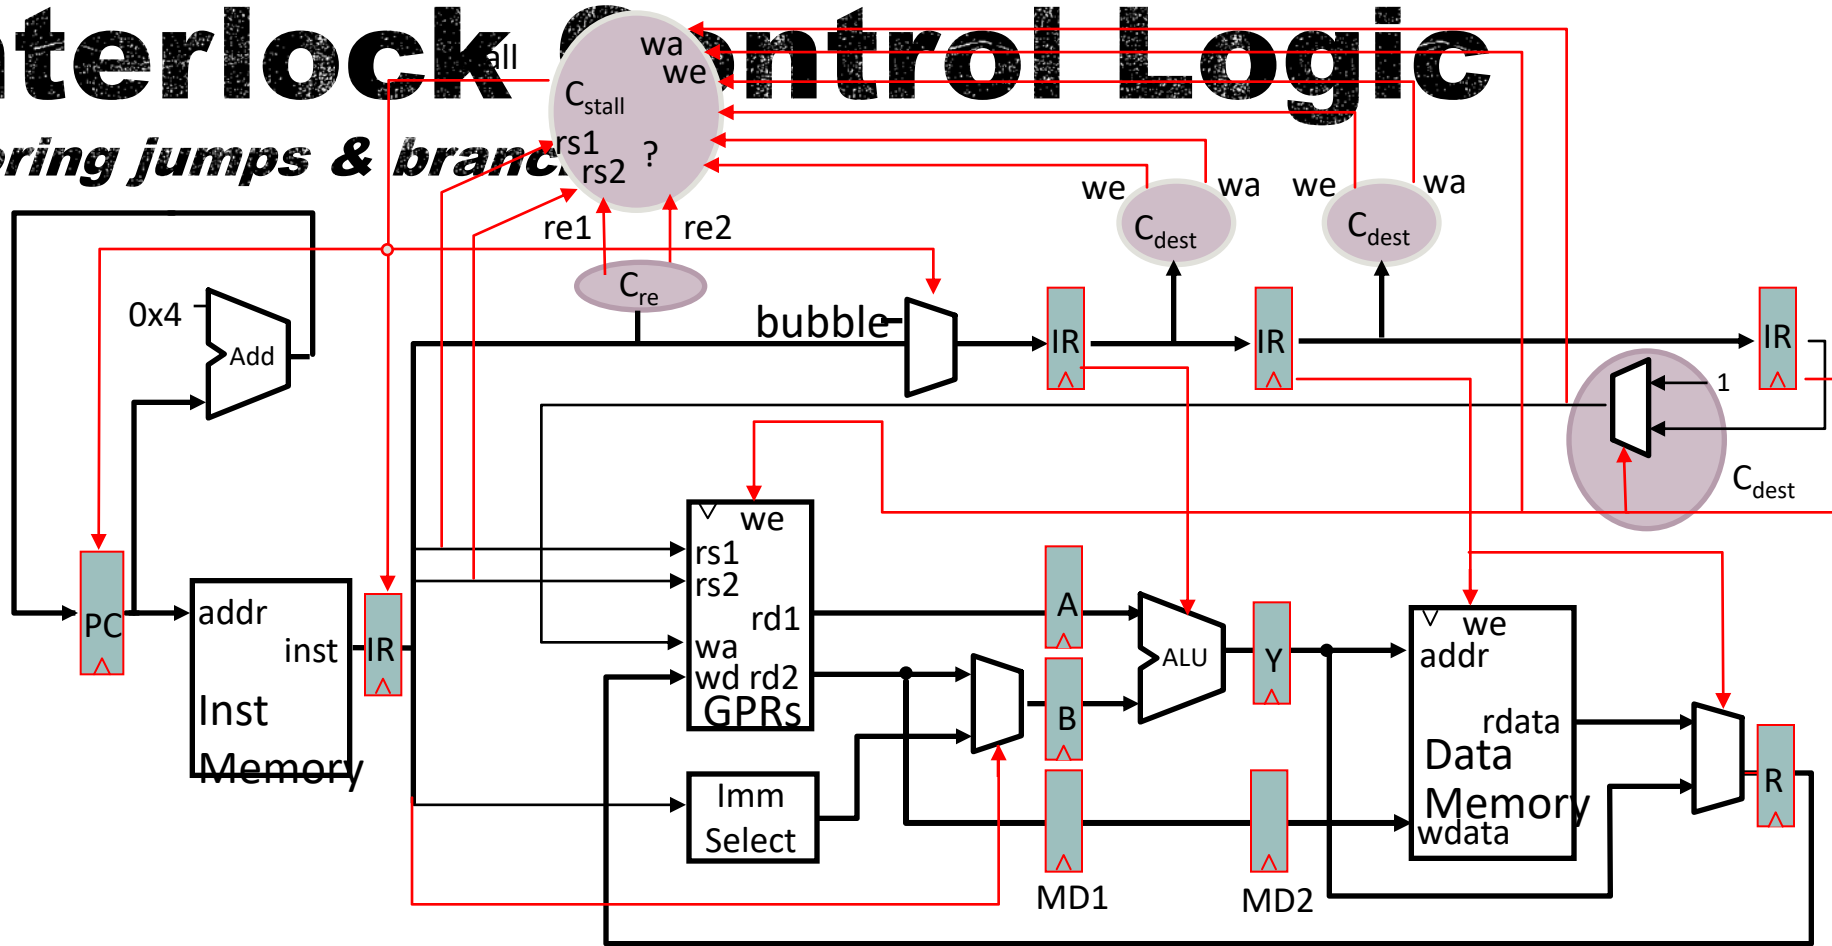
Interlock Control Logic



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

Interlock Control Logic

ignoring jumps & branches



Should we always stall if an rs field matches some rd?
 not every instruction writes a register wa
 not every instruction reads a register ra

Source Register

rd	rs1	rs2	func10	opcode
rd	rs1	Imm[11:0]	func3	opcode
Imm[11:7]	rs1	rs2	Imm[6:0]	func3
Jump offset[24:0]				opcode

ALU
ALUI/LW/JALR
SW/Bcond

	<i>source(s)</i>		<i>destination</i>	
ALU	rd	rs1, rs2	rs1, rs2	rd
ALUI	rd	rs1	rd	
LW	rd	rs1	rd	
SW		rs1, rs2	-	
Bcond		rs1, rs2	-	
	<i>true:</i> PC + imm			
	<i>false:</i> PC + 4			
J		-	-	
JAL	x1	-	x1	
JALR	rd	rs1	rd	

Deriving the Stall Signal

C_{dest}

$ws = \text{Case opcode}$

JAL $??X1$

else $??rd$

$we = \text{Case opcode}$

ALU, ALUi, LW, JALR $?(ws \neq 0)$

JAL $??on$

... $??off$

C_{re}

$re1 = \text{Case opcode}$

ALU, ALUi,

LW, SW, Bcond,

JALR $??on$

J, JAL $??off$

$re2 = \text{Case opcode}$

ALU, SW, Bcond,

... $??off$

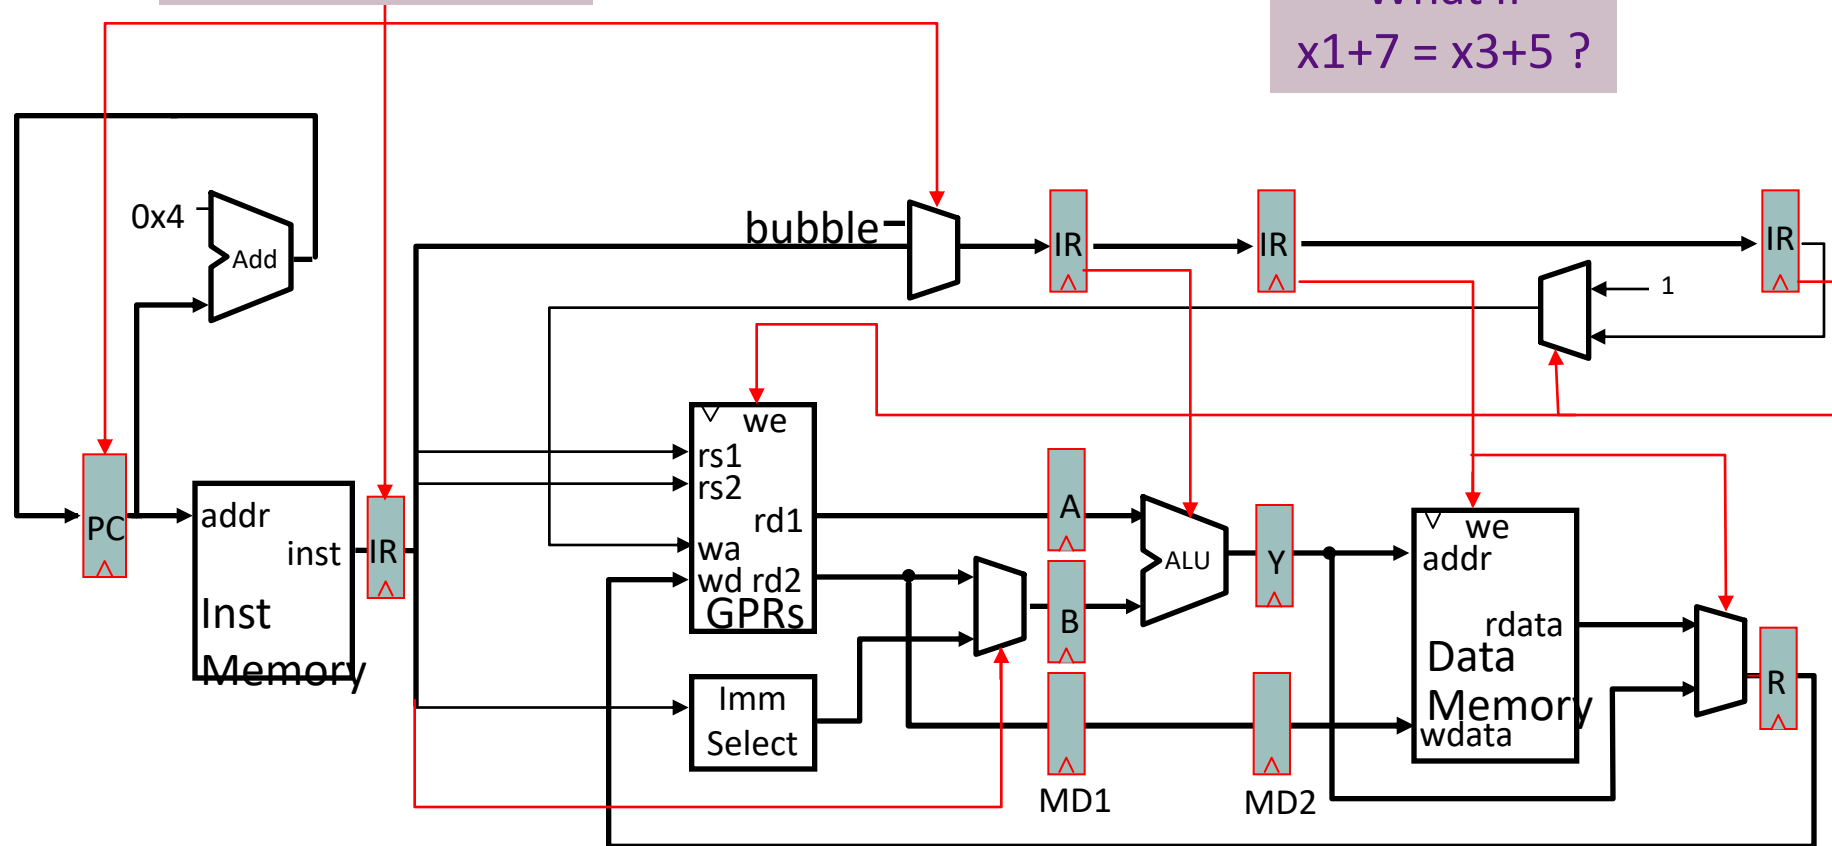
$$C_{stall} \quad stall = ((rs1_D = ws_E).we_E + (rs1_D = ws_M).we_M + (rs1_D = ws_W).we_W) \cdot re1_D + ((rs2_D = ws_E).we_E + (rs2_D = ws_M).we_M + (rs2_D = ws_W).we_W) \cdot re2_D$$

This is not the full story!

Hazards due to Loads & Stores

Stall Condition

What if $x1+7 = x3+5$?



...
 $M[x1+7] \neq x2$
 $x4 \neq M[x3+5]$
 ...

Is there any possible data hazard in this instruction sequence?

Load & Store Hazards

$M[x1+7] \text{ ? } x2$
 $x4 \text{ ? } M[x3+5]$
...

$x1+7 = x3+5 \text{ ? } \textit{data hazard}$

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.

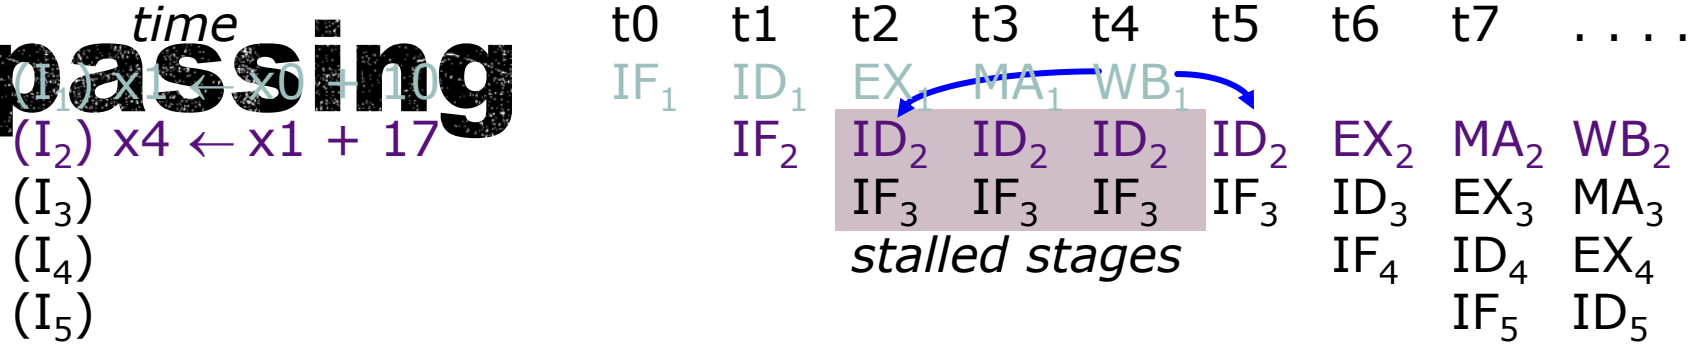
More on this later in the course.

Resolving Data Hazards (2)

Strategy 2:

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


Bypassing



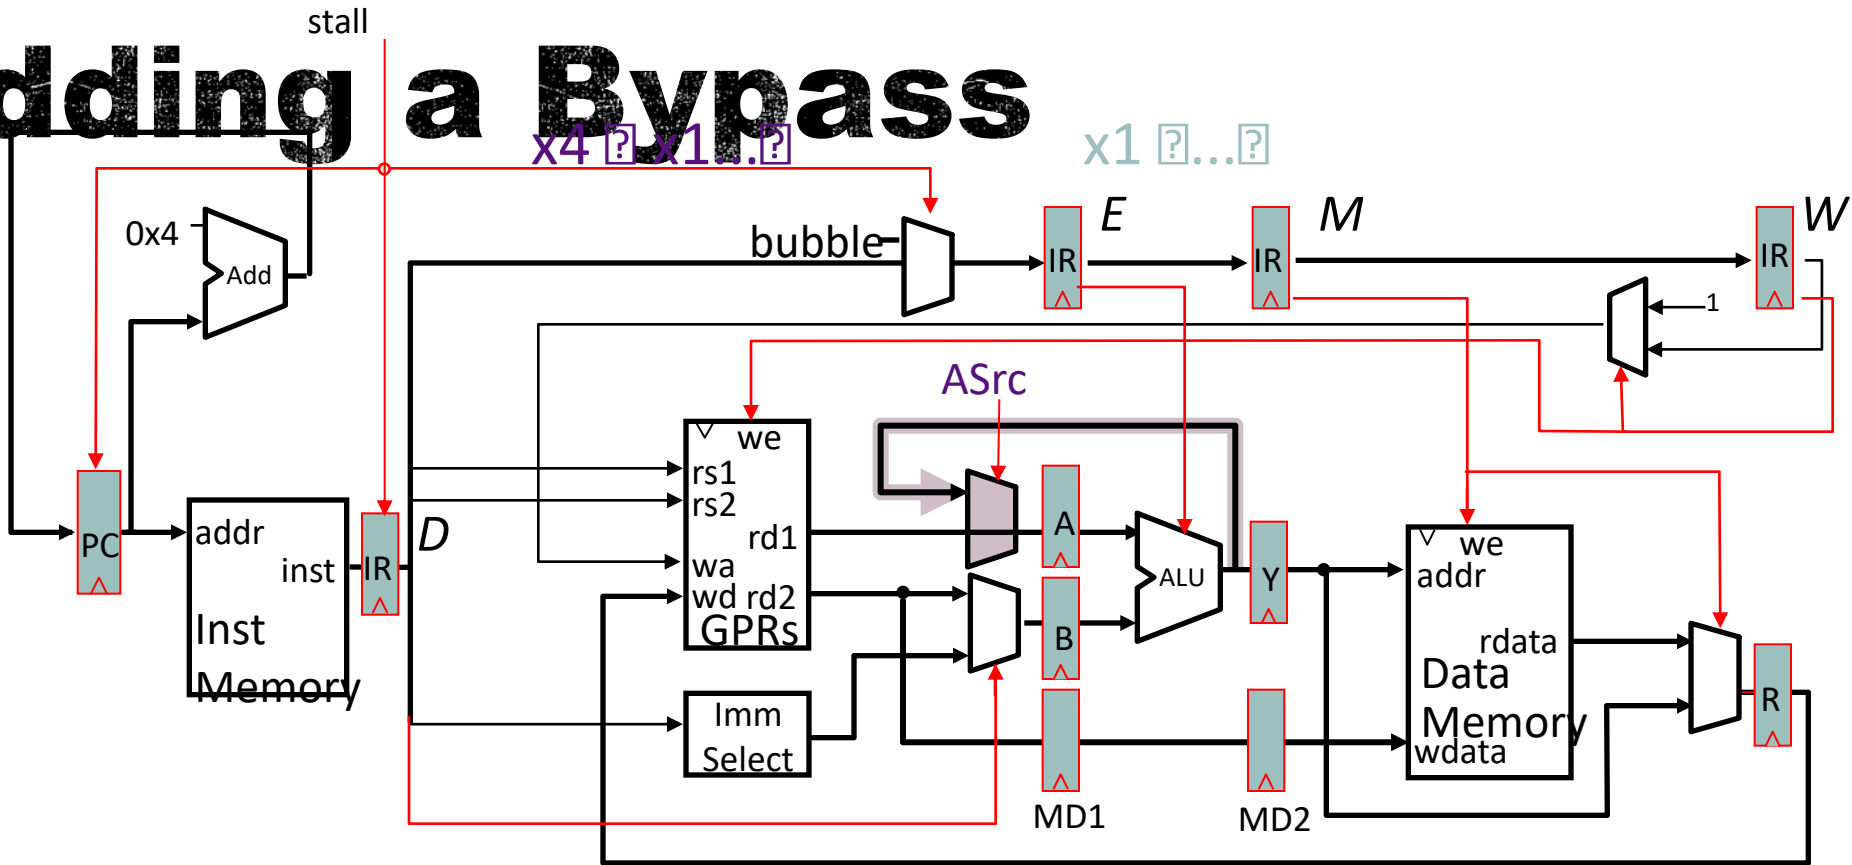
Each *stall or kill* introduces a bubble in the pipeline

$$CPI > 1$$

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?

...
 $(I_1) \ x1 \ \boxed{??} \ x0 + 10$
 $(I_2) \ x4 \ \boxed{??} \ x1 + 17$
yes

$x1 \ \boxed{??} \ M[x0 + 10]$
 $x4 \ \boxed{??} \ x1 + 17$
no

JAL 500
 $x4 \ \boxed{??} \ x1 + 17$
no

The Bypass Signal

Deriving it from the Stall Signal

$$\text{stall} = ((rs1_D = ws_E).we_E + (rs1_D = ws_M).we_M + (rs1_D = ws_W).we_W).re1_D \\ + ((rs2_D = ws_E).we_E + (rs2_D = ws_M).we_M + (rs2_D = ws_W).we_W).re2_D$$

ws = Case opcode

JAL $??X1$

else $??rd$

we = Case opcode

ALU, ALUi, LW, JALR $?(ws \neq 0)$

JAL $??on$

... $??off$

$$ASrc = (rs1_D = ws_E).we_E.re1_D$$

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

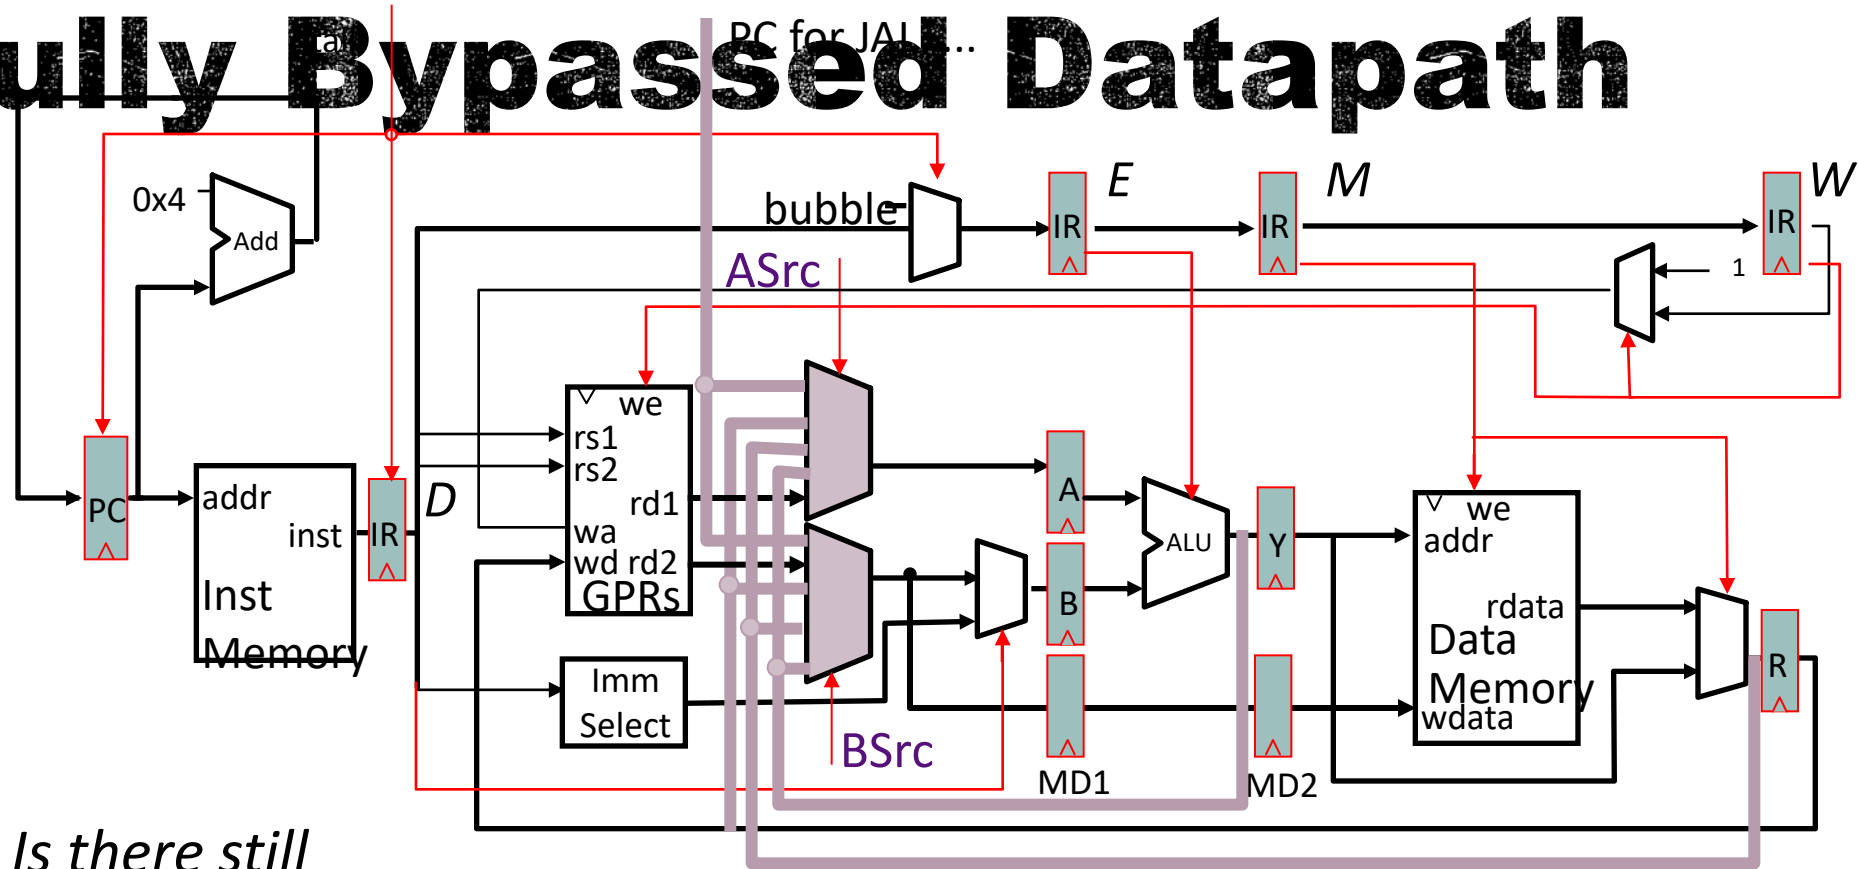
$we\text{-}bypass_E = \text{Case opcode}_E$
 ALU, ALUi ? (ws ? 0)
 ... ??off

$we\text{-}stall_E = \text{Case opcode}_E$
 LW, JALR ? (ws ? 0)
 JAL ??on
 ... ??off

$ASrc = (rs1_D = ws_E).we\text{-}bypass_E . re1_D$

$stall = ((rs1_D = ws_E).we\text{-}stall_E +$
 $(rs1_D = ws_M).we_M + (rs1_D = ws_W).we_W). re1_D$
 $+ ((rs2_D = ws_E).we_E + (rs2_D = ws_M).we_M + (rs2_D = ws_W).we_W). re2_D$

Fully Bypassed Datapath




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


$$\text{stall} = (\text{rs1}_D = \text{ws}_E) \cdot (\text{opcode}_E = \text{LW}_E) \cdot (\text{ws}_E \neq 0) \cdot \text{re1}_D \\ + (\text{rs2}_D = \text{ws}_E) \cdot (\text{opcode}_E = \text{LW}_E) \cdot (\text{ws}_E \neq 0) \cdot \text{re2}_D$$

Pipeline CPI Examples


Measure from when first instruction finishes to when last instruction in sequence finishes.



3 instructions finish in 3 cycles
 $CPI = 3/3 = 1$



3 instructions finish in 4 cycles
 $CPI = 4/3 = 1.33$



3 instructions finish in 5 cycles
 $CPI = 5/3 = 1.67$

Resolving Data Hazards (3)

Strategy 3: Speculate on the dependence!

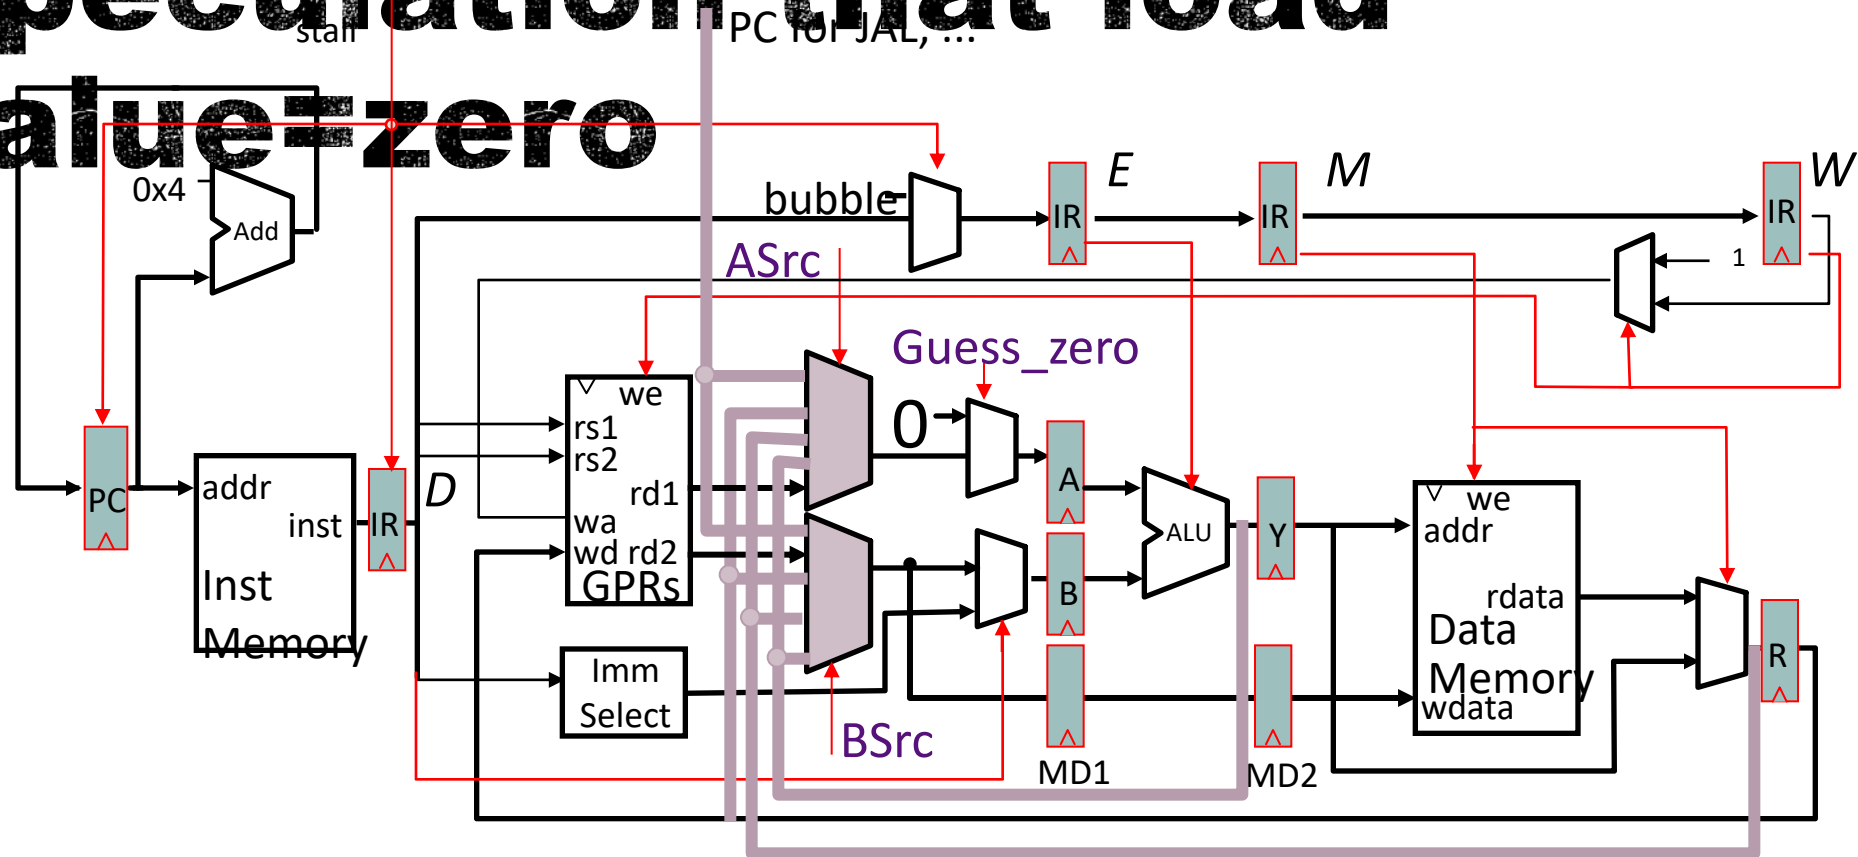
Two cases:

Gessed correctly → do nothing

Gessed incorrectly → kill and restart

.... We'll later see examples of this approach in more complex processors.

Speculation that load value=zero



$$\text{Guess_zero} = (\text{rs1}_D = \text{ws}_E) \cdot (\text{opcode}_E = \text{LW}_E) \cdot (\text{ws}_E \neq 0) \cdot \text{re1}_D$$

Also need to add circuitry to remember that this was a guess and flush pipeline if load not zero!


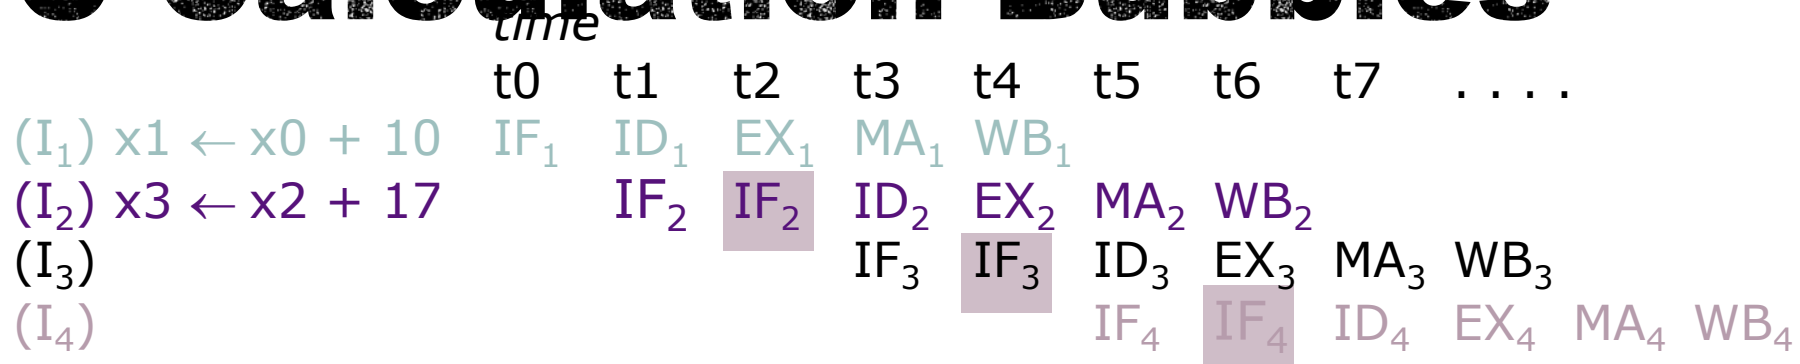
Not worth doing in practice – why?

Control Hazards

What do we need to calculate next PC?

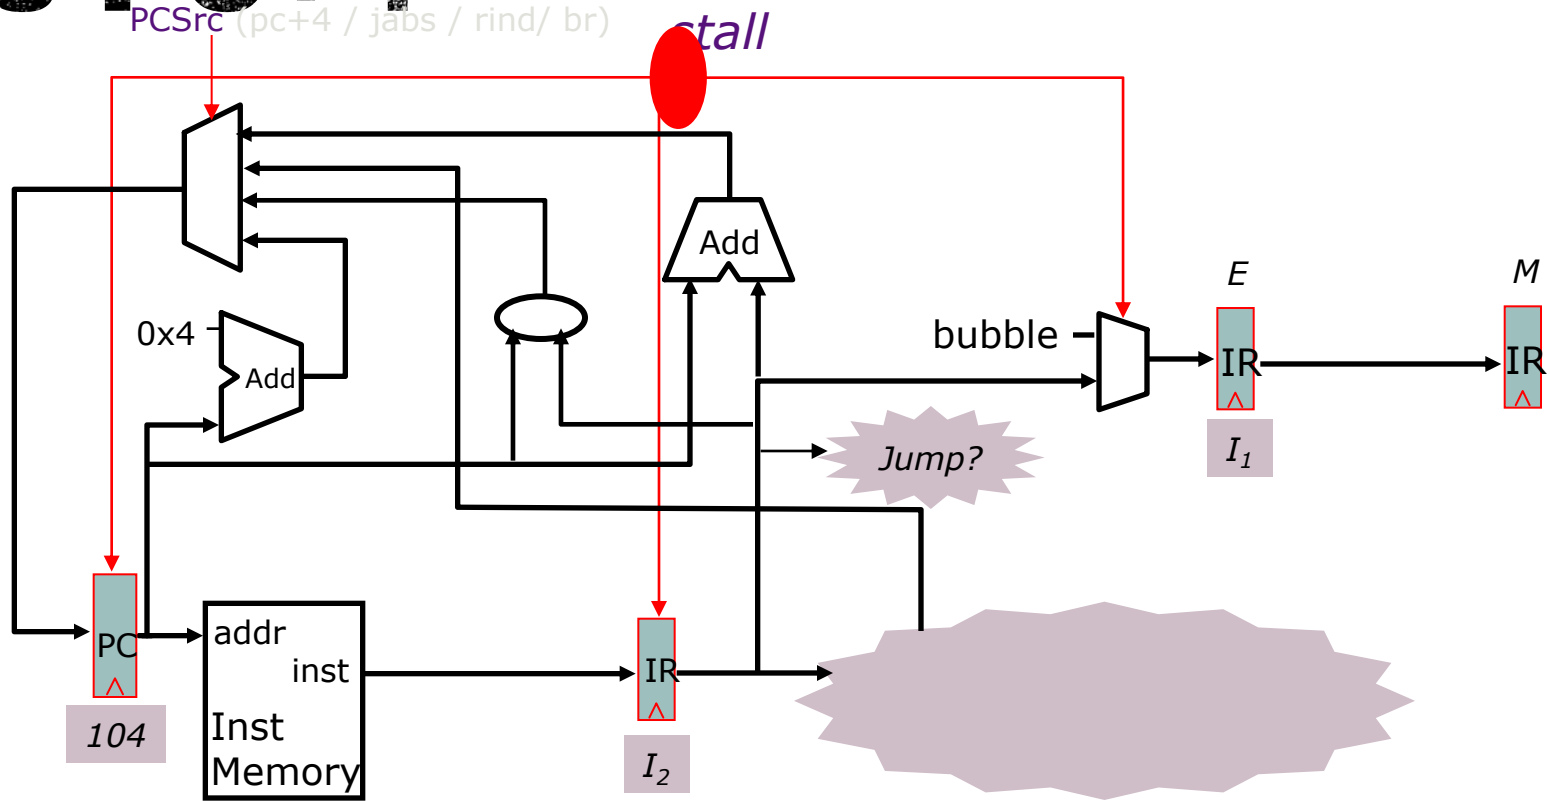
- For Jumps
 - Opcode, PC and offset
- For Jump Register
 - Opcode, Register value, and PC
- For Conditional Branches
 - Opcode, Register (for condition), PC and offset
- For all other instructions
 - Opcode and PC (and have to know it's not one of above)

PC Calculation Bubbles



- ⇒ *pipeline bubble*

Speculate next address is PC+4



I_1 096ADD

I_2 100J 304

I_3 104ADD

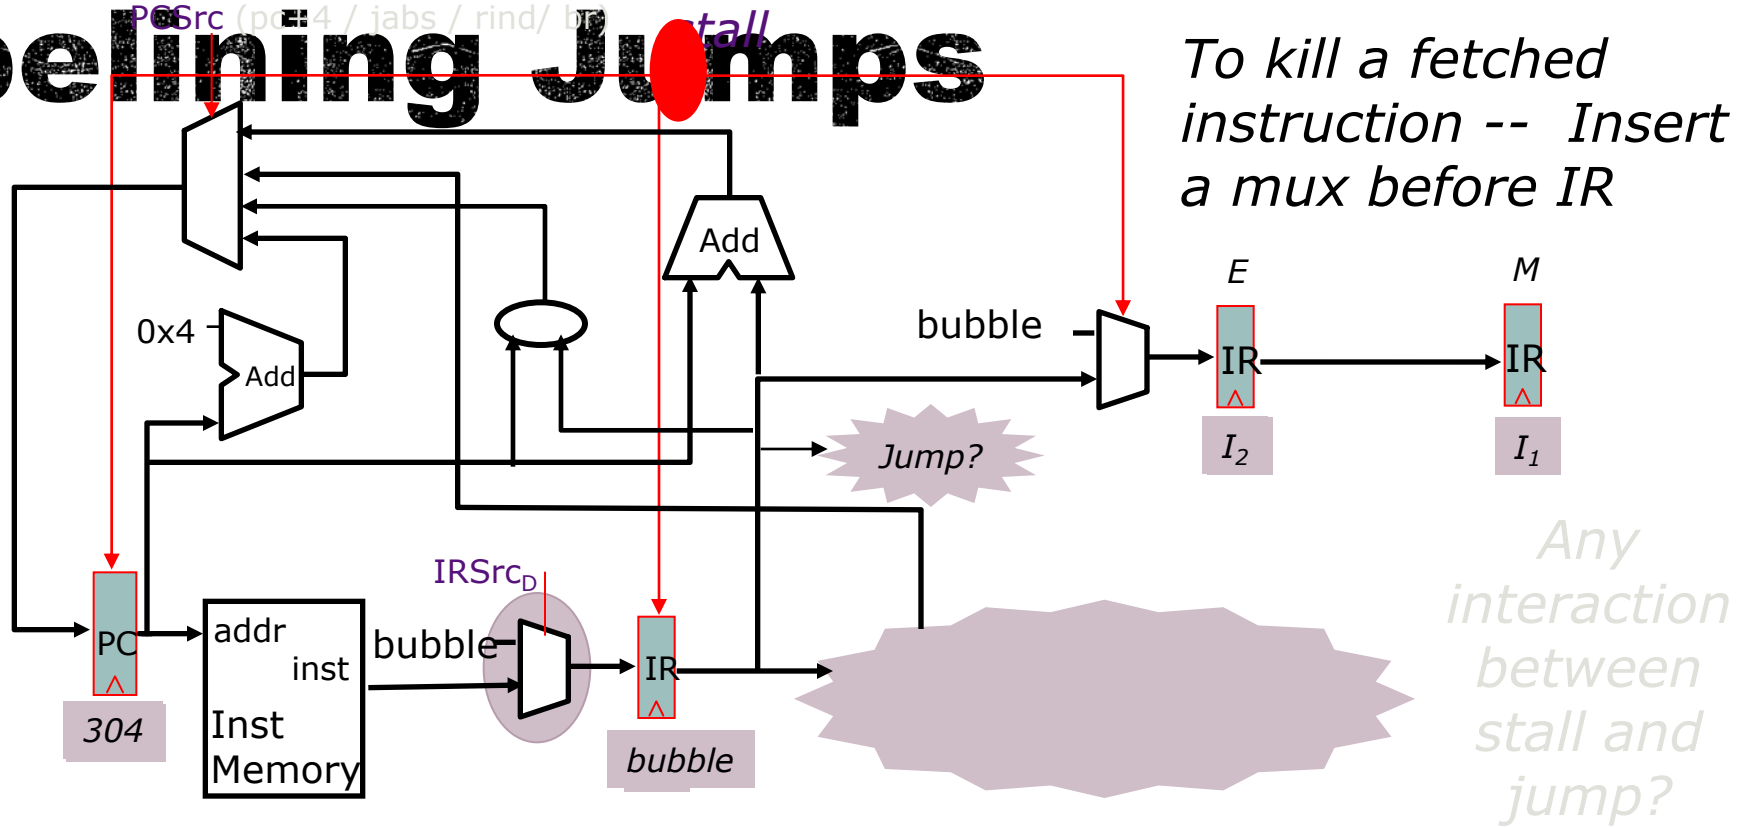
I_4 304ADD

kill

A jump instruction kills (not stalls) the following instruction

How?

Pipelining Jumps



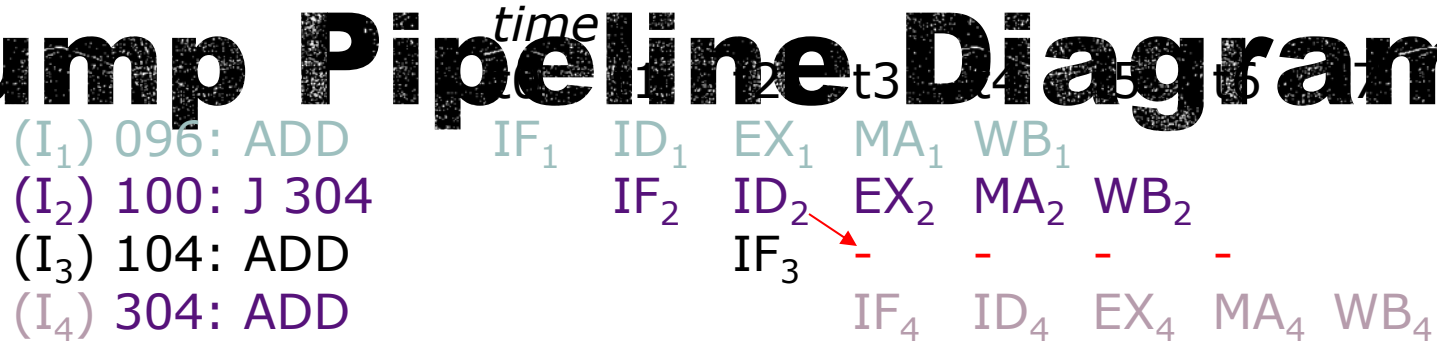
To kill a fetched instruction -- Insert a mux before IR

Any interaction between stall and jump?

- I₁ 096 ADD
- I₂ 100 J 304
- I₃ 104 ADD — kill
- I₄ 304 ADD

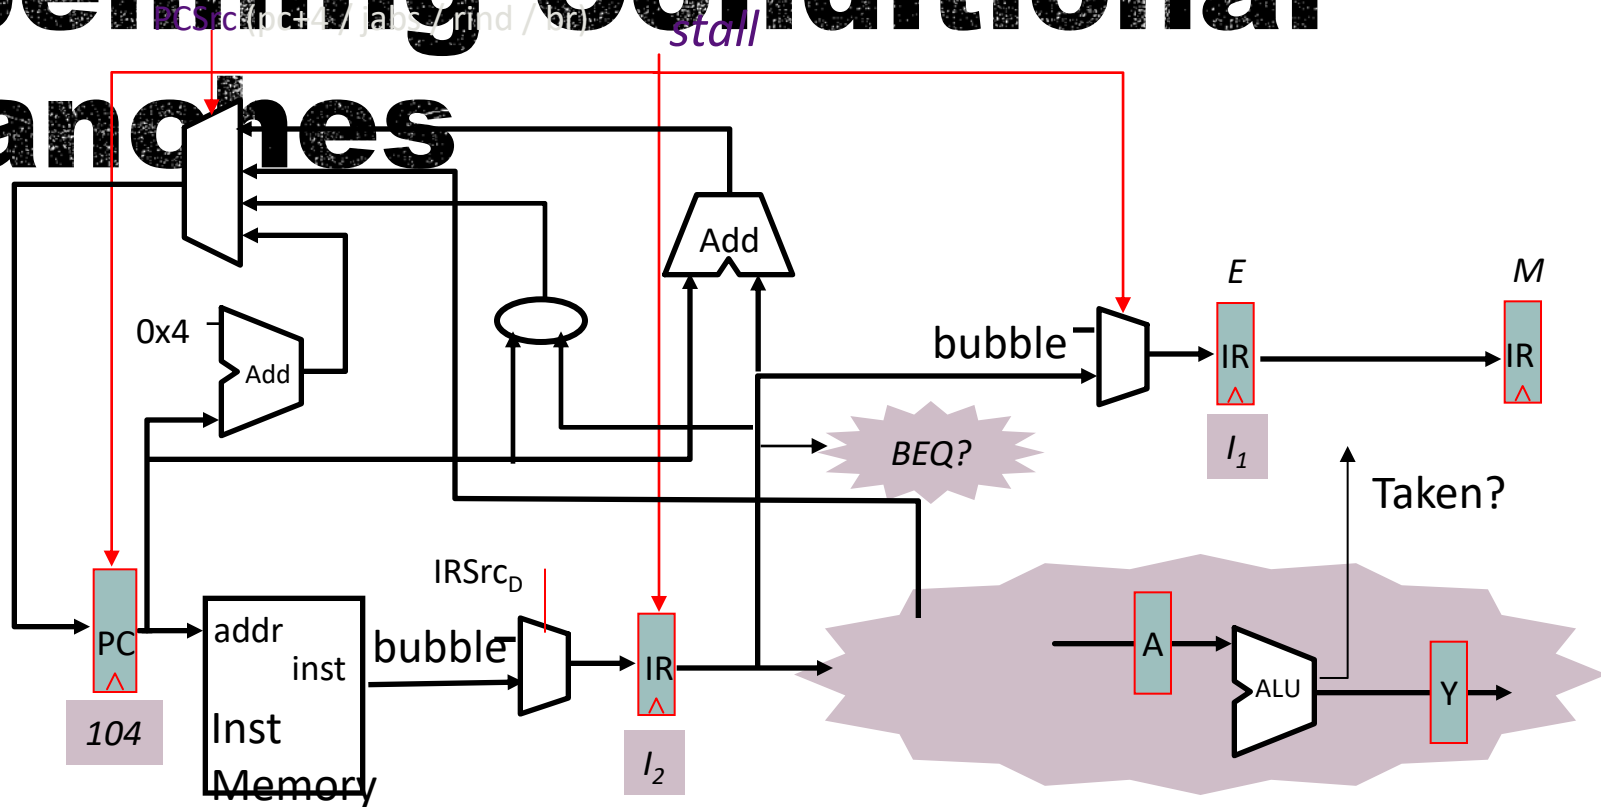
IRSrc_D = Case opcode_D
 J, JAL ⇒ bubble
 ... ⇒ IM

Jump Pipeline Diagrams.



- ⇒ *pipeline bubble*

Pipelining Conditional Branches


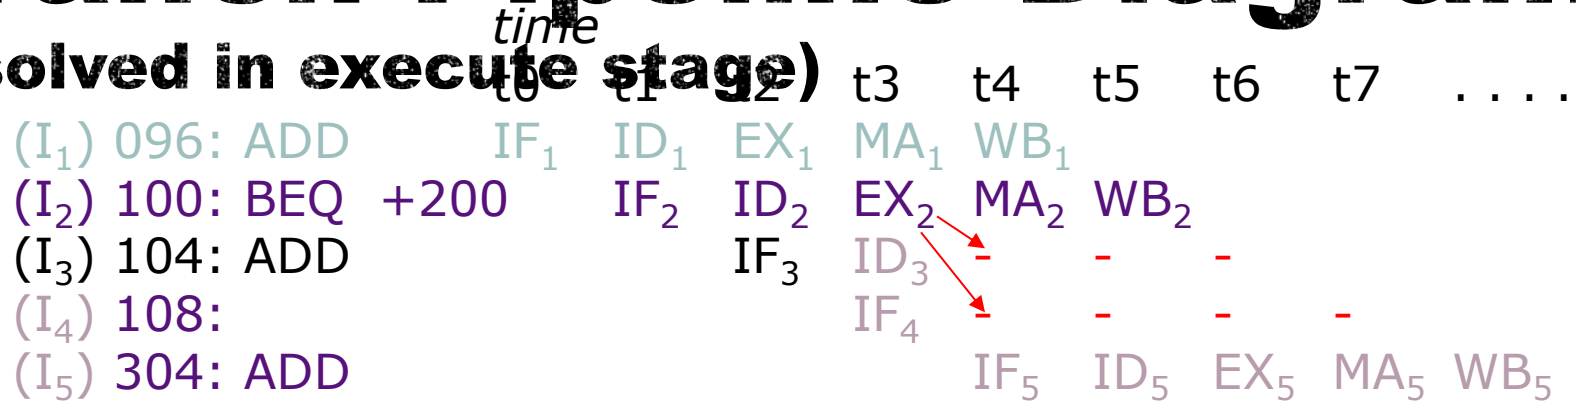


- I_1 096 ADD
- I_2 100 BEQ x1,x2 +200
- I_3 104 ADD
- I_4 304 ADD

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

Branch Pipeline Diagrams

(resolved in execute stage)



- ⇒ pipeline bubble

Acknowledgements

- These slides contain material developed and copyright by:
 - Arvind (MIT)
 - Krste Asanovic (MIT/UCB)
 - Joel Emer (Intel/MIT)
 - James Hoe (CMU)
 - John Kubiatowicz (UCB)
 - David Patterson (UCB)
- MIT material derived from course 6.823
- UCB material derived from course CS252