

ECE 252 / CPS 220

Advanced Computer Architecture I

Lecture 7

Pipelining – Part 2

Benjamin Lee
Electrical and Computer Engineering
Duke University

www.duke.edu/~bcl15
www.duke.edu/~bcl15/class/class_ece252fall11.html



ECE252 Administrivia

29 September – Homework #2 Due

- Use blackboard forum for questions
- Attend office hours with questions
- Email for separate meetings

4 October – Class Discussion

Roughly one reading per class. Do not wait until the day before!

1. Srinivasan et al. “Optimizing pipelines for power and performance”
2. Mahlke et al. “A comparison of full and partial predicated execution support for ILP processors”
3. Palacharla et al. “Complexity-effective superscalar processors”
4. Yeh et al. “Two-level adaptive training branch prediction”



Data Hazards and Scheduling

Try producing faster code for

- $A = B + C$; $D = E - F$;
- Assume A, B, C, D, E, and F are in memory
- Assume pipelined processor

Slow Code

LW	Rb, b
LW	Rc, c
ADD	Ra, Rb, Rc
SW	a, Ra
LW	Re, e
LW	Rf, f
SUB	Rd, Re, Rf
SW	d, Rd

Fast Code

LW	Rb, b
LW	Rc, c
LW	Re, e
ADD	Ra, Rb, Rc
LW	Rf, f
SW	a, Ra
SUB	Rd, Re, Rf
SW	d, Rd



Compiler Scheduling

Reduce stalls by moving instructions

- Basic pipeline scheduling eliminates back-to-back load-use pairs
- What are the limitations of scheduling?

Scheduling Scope

- Requires an independent instruction to place between load-use pairs
- Little scope for scheduling, 1-add, 3-ld/st

Slow Code

```
ld      r2, 4(sp)
ld      r3, 8(sp)
add     r3, r2, r1
st      r1, 0(sp)
```

Fast Code

```
ld      r2, 4(sp)
ld      r3, 8(sp)
add     r3, r2, r1
st      r1, 0(sp)
```



Compiler Scheduling

Number of registers

- Registers hold “live” values
- Example code contains 7 different values, including sp
- Before: max 3 values live → 3 registers sufficient
- After: max 4 values live → 3 registers insufficient
- Original code re-uses r1 and r2, re-scheduling causes WAR violations

<u>Before</u>			<u>After</u>	
ld	r2, 4 (sp)		ld	r2, 4 (sp)
ld	r1 , 8 (sp)		ld	r1, 8 (sp)
add	r1 , r2, r1 #stall		ld	r2 , 16 (sp)
st	r1, 0 (sp)		add	r1, r2 , r1 #WAR
ld	r2, 16 (sp)		ld	r1 , 20 (sp)
ld	r1 , 20 (sp)		st	r1 , 0 (sp) #WAR
sub	r2, r1, r1 #stall		sub	r2, r1, r1
st	r1, 12 (sp)		st	r1, 12 (sp)



Compiler Scheduling

Alias Analysis

- Determine whether load/stores reference same memory locations
- Determines if loads/stores can be re-ordered
- Previous Examples: easy, all loads/stores use the same base register (sp)
- New Example: Can compiler tell that r8 = sp?

<u>Before</u>				<u>After</u>	
ld	r2, 4 (sp)			ld	r2, 4 (sp)
ld	r3, 8 (sp)			ld	r3, 8 (sp)
add	r3, r2, r1	#stall		ld	r5, 0 (r8)
st	r1, 0 (sp)			add	r3, r2, r1
ld	r5, 0 (r8)			ld	r6, 4 (r8)
ld	r6, 4 (r8)			st	r1, 0 (sp)
sub	r5, r6, r4	#stall		sub	r5, r6, r4
st	r4, 8 (r8)			st	r4, 8 (r8)



Control Hazards

Arises when calculating next program counter (PC)

- Pipeline stalls if required values not yet available

Jumps

- Jumps (immediate) requires opcode, offset, current PC
- Jump (register) requires opcode, register value

Conditional Branches

- Requires opcode, current PC, register (for condition), offset

Sequential Successor Instructions

- Opcode, current PC



Program Counter Calculations

Identify change in control flow during decode

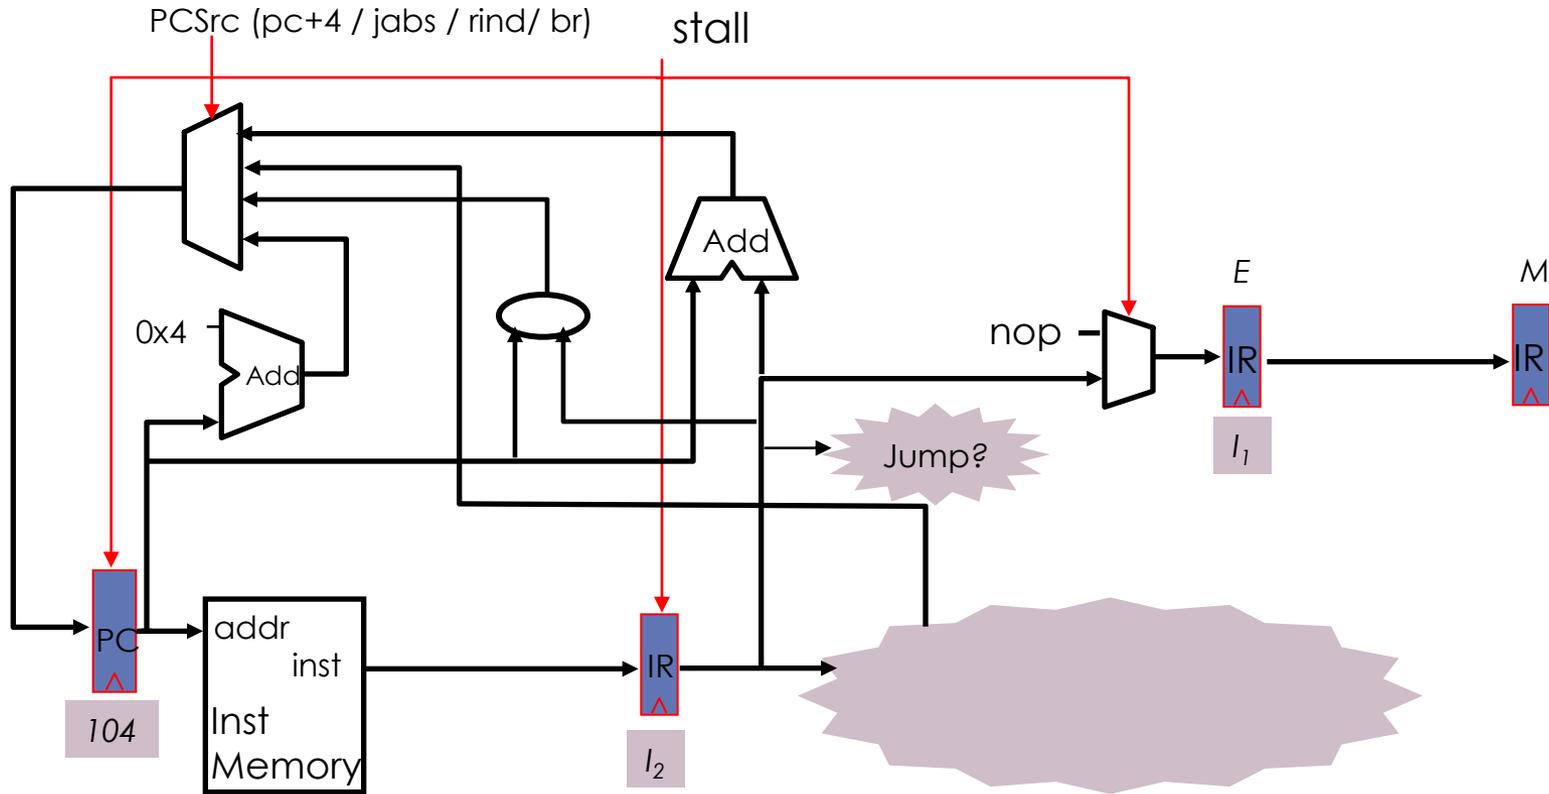
	Time									
	t0	t1	t2	t3	t4	t5	t6	t7	
(I ₁) r1 ← (r0) + 10	IF ₁	ID ₁	EX ₁	MA ₁	WB ₁					
(I ₂) r3 ← (r2) + 17		IF ₂	IF ₂	ID ₂	EX ₂	MA ₂	WB ₂			
(I ₃)			IF ₃	IF ₃	ID ₃	EX ₃	MA ₃	WB ₃		
(I ₄)					IF ₄	IF ₄	ID ₄	EX ₄	MA ₄	WB ₄

Resource Usage

	time									
	t0	t1	t2	t3	t4	t5	t6	t7	
IF	I ₁	nop	I ₂	nop	I ₃	nop	I ₄			
ID		I ₁	nop	I ₂	nop	I ₃	nop	I ₄		
EX			I ₁	nop	I ₂	nop	I ₃	nop	I ₄	
MA				I ₁	nop	I ₂	nop	I ₃	nop	I ₄
WB					I ₁	nop	I ₂	nop	I ₃	nop



Speculate $PC \leftarrow PC + 4$

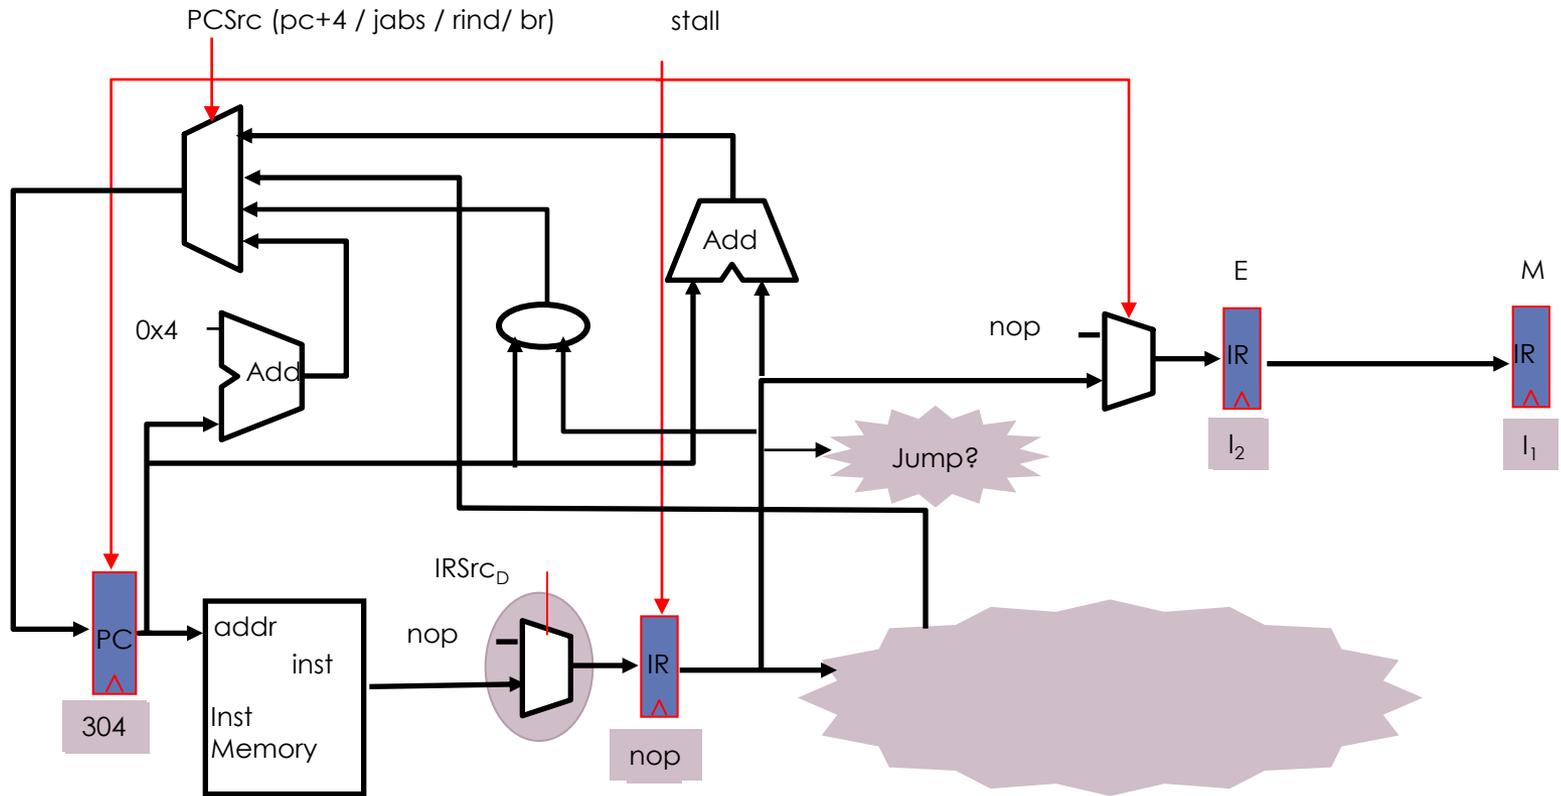


	addr	instr	
l_1	096	ADD	
l_2	100	J 304	
l_3	104	ADD	kill
l_4	304	ADD	

A jump instruction kills (not stalls) the following instruction. How?



Pipelining Jumps



	addr	instr	
I_1	096	ADD	
I_2	100	J 304	
I_3	104	ADD	kill
I_4	304	ADD	

To kill a fetched instruction, add mux before IR to insert "nops"

J, JAL: $IR \leftarrow nop$
 otherwise: $IR \leftarrow inst$



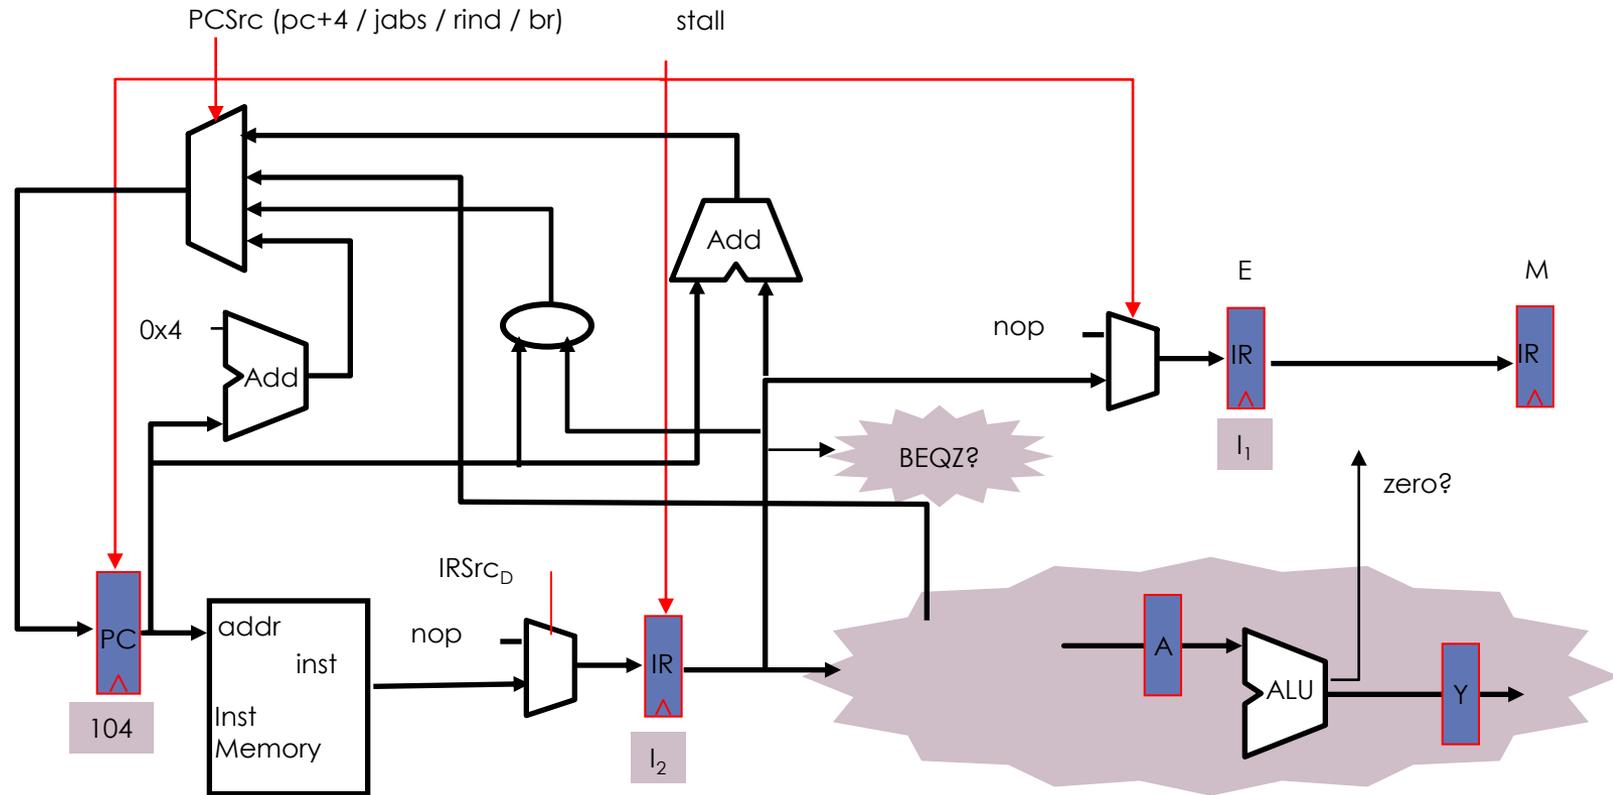
Pipelining Jumps

	time								
	t0	t1	t2	t3	t4	t5	t6	t7
(I ₁) 096: ADD	IF ₁	ID ₁	EX ₁	MA ₁	WB ₁				
(I ₂) 100: J 304		IF ₂	ID ₂	EX ₂	MA ₂	WB ₂			
(I ₃) 104: ADD			IF ₃	nop	nop	nop	nop		
(I ₄) 304: ADD			IF ₄	ID ₄	EX ₄	MA ₄	WB ₄		

	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	nop	I ₄	I ₅			
EX			I ₁	I ₂	nop	I ₄	I ₅		
MA				I ₁	I ₂	nop	I ₄	I ₅	
WB					I ₁	I ₂	nop	I ₄	I ₅



Pipelining Conditional Branches

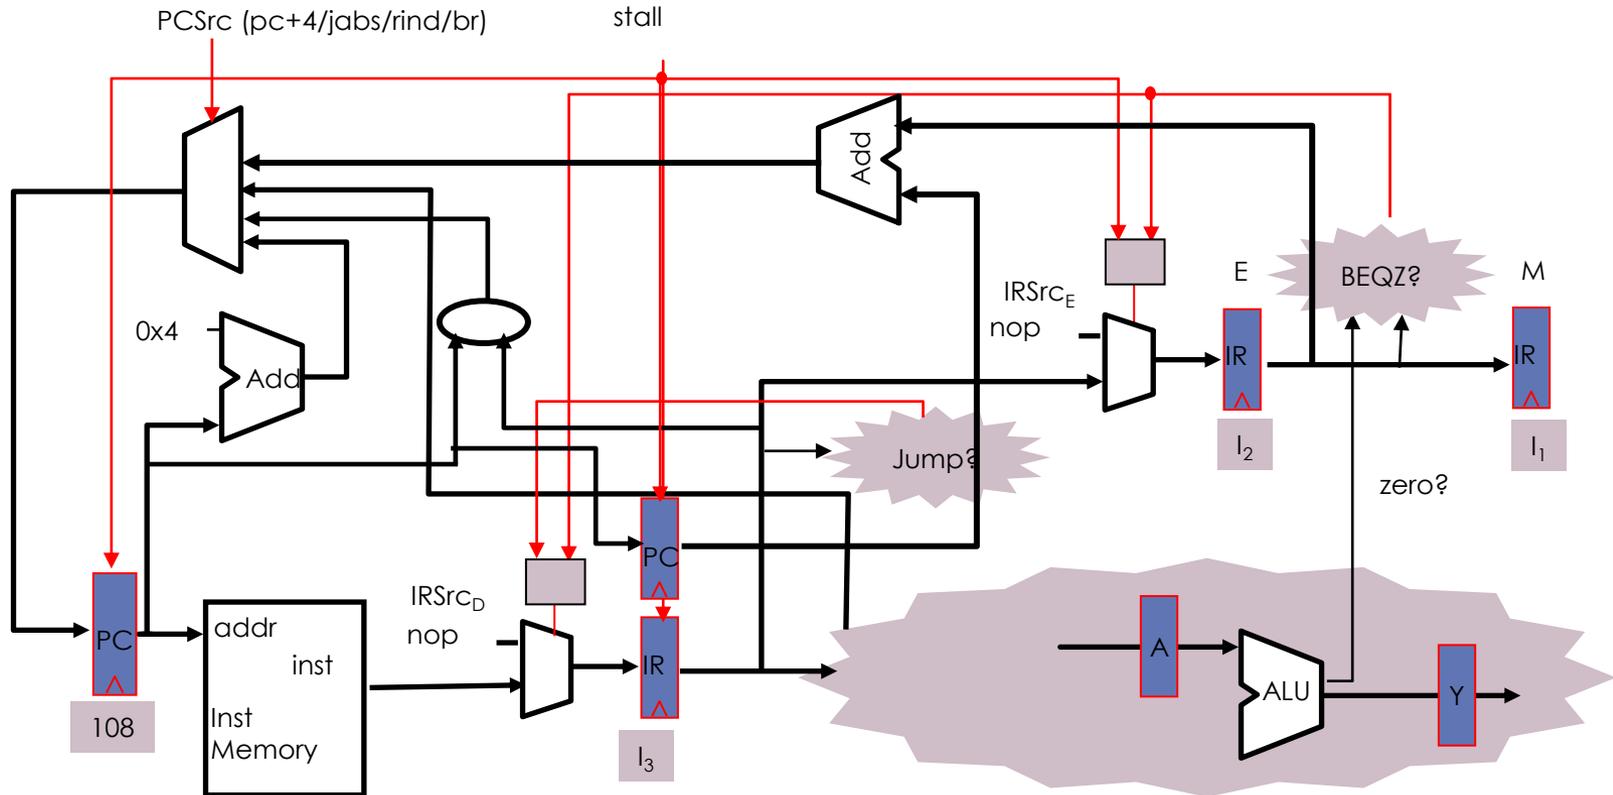


	<u>addr</u>	<u>instr</u>
l_1	096	ADD
l_2	100	BEQZ r1 200
l_3	104	ADD
l_4	304	ADD

Branch condition computed in execute stage. What should be done in decode stage?



Pipelining Conditional Branches



	addr	instr
l_1	096	ADD
l_2	100	BEQZ r1 200
l_3	104	ADD
l_4	304	ADD

If branch is taken, kill two following instructions. And because instruction in decode stage is invalid, update stall signal.



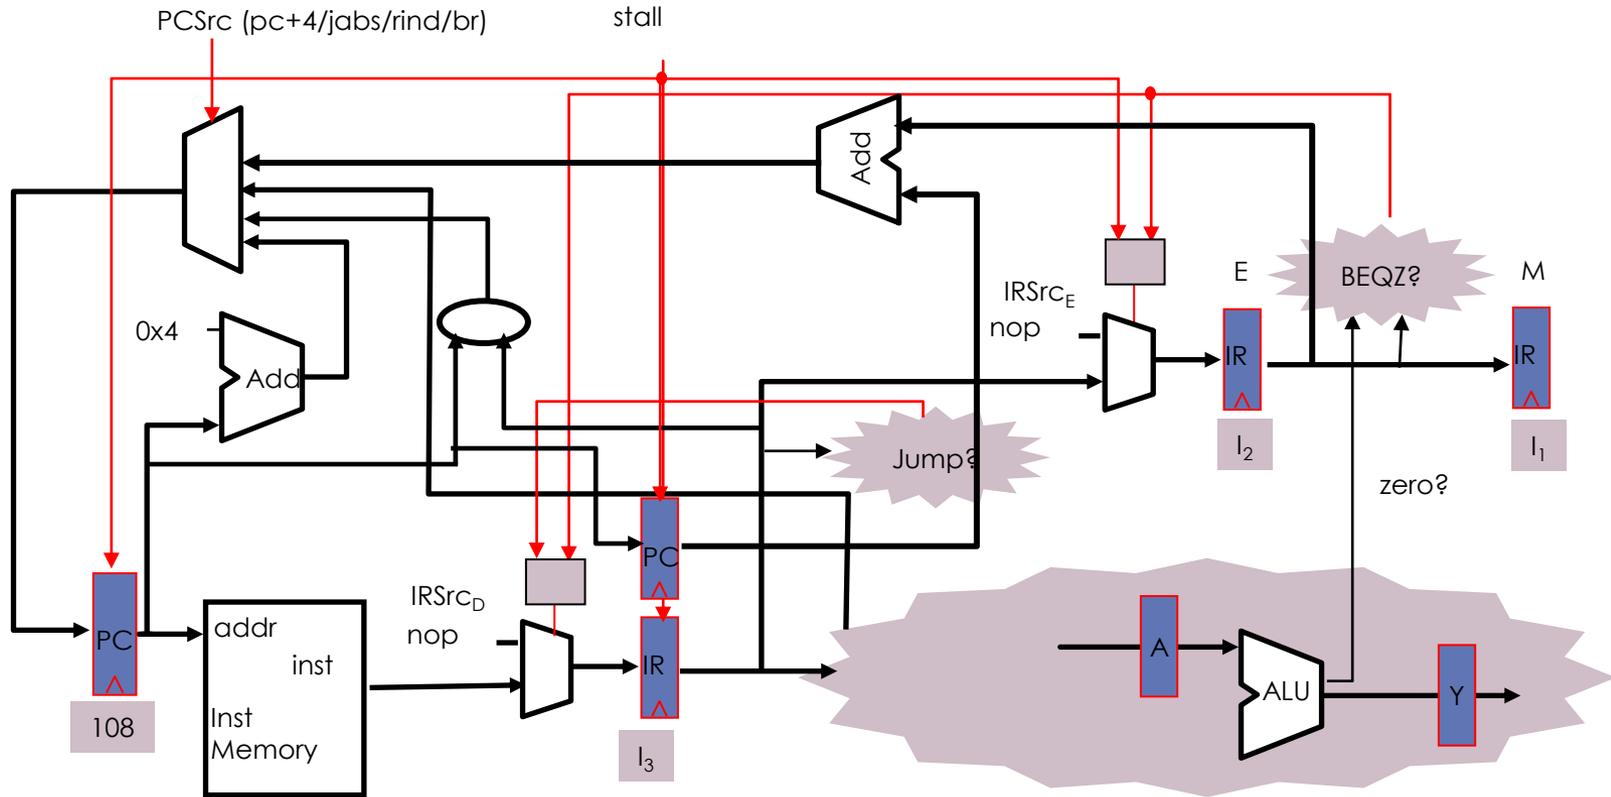
Update Stall Signal

```
Stall ← <<original stall signal>>  
      & !(      (opcodeE == BEQZ) & zero?      # branch condition true  
              + (opcodeE == BNEZ) & !zero?    # branch condition true  
      )
```

Do not stall if branch is taken. Why?
Instruction in the decode stage is invalid.
Kill instruction instead.



Pipelining Conditional Branches



	<u>addr</u>	<u>instr</u>
l_1	096	ADD
l_2	100	BEQZ r1 200
l_3	104	ADD
l_4	304	ADD

If branch is taken, kill two following instructions. And because instruction in decode stage is invalid, update stall signal.



Derive PCSrc Signal

Derive mux control signal for PCSrc.

if((opcodeE == BEQZ & z) + (opcodeE == BNEZ & !z)), PCSrc \leftarrow br
else if ((opcodeD == J) + (opcodeD == JAL)), PCSrc \leftarrow jabs
else if ((opcodeD == JR) + (opcodeD == JALR)), PCSrc \leftarrow rind
otherwise, PCSrc \leftarrow PC + 4

Derive mux control signal for IRSrcD.

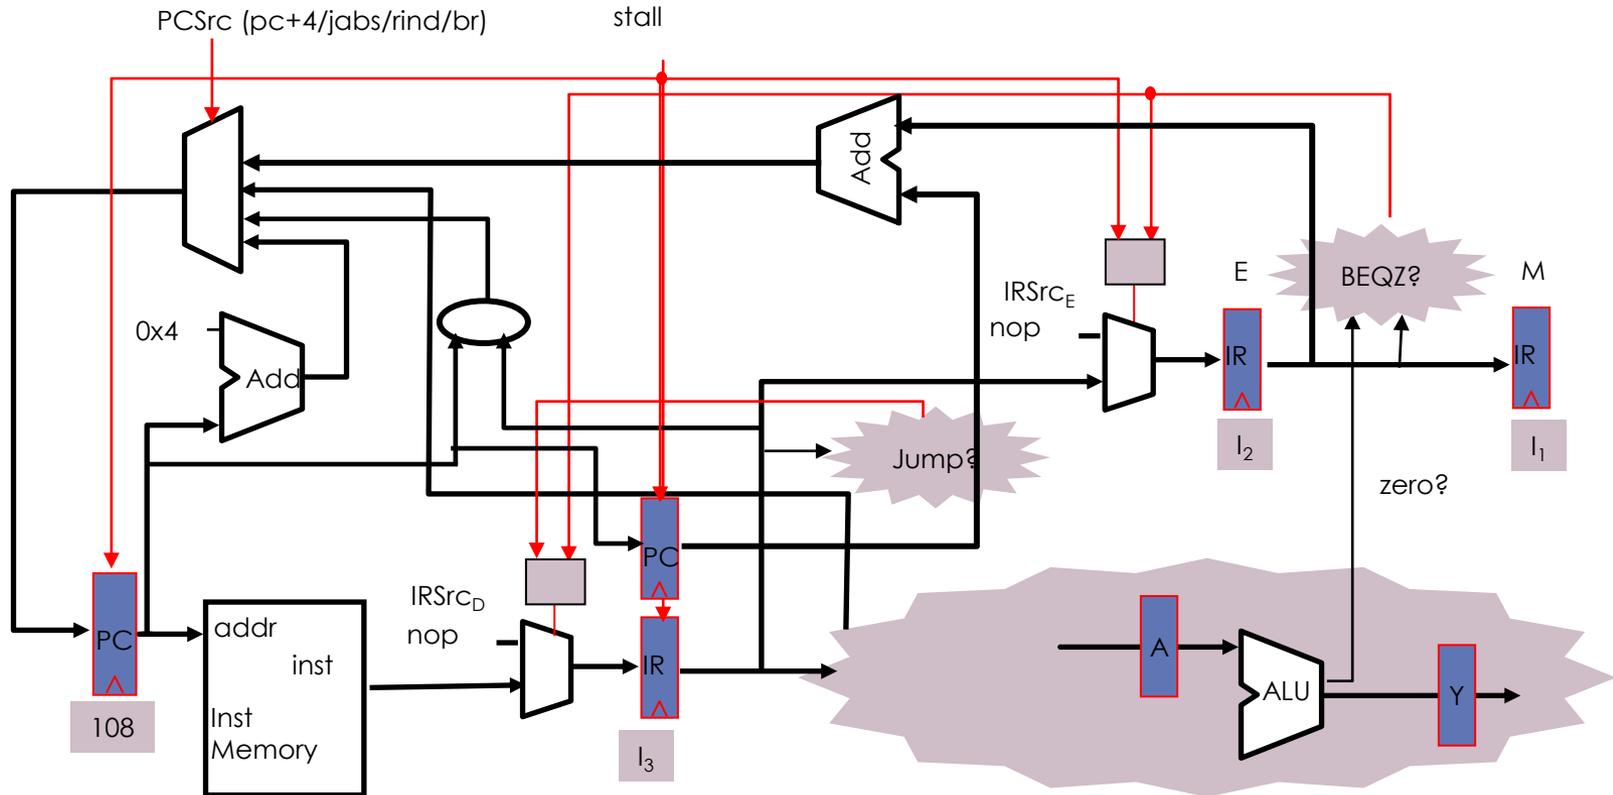
if((opcodeE == BEQZ & z) + (opcodeE == BNEZ & !z)), ICSrcD \leftarrow nop
else if((opcodeD==J) + (opcodeD==JAL) +
 (opcodeD==JAR) + (opcodeD==JALR)), ICSrcD \leftarrow nop
otherwise, ICSrcD \leftarrow Instr

Derive mux control signal for IRSrcE.

if((opcodeE == BEQZ & z) + (opcodeE == BNEZ & !z)), ICSrcE \leftarrow nop
otherwise, IRSrcE \leftarrow (stall & nop) + (!stall & IRD)



Pipelining Conditional Branches



	<u>addr</u>	<u>instr</u>
l_1	096	ADD
l_2	100	BEQZ r1 200
l_3	104	ADD
l_4	304	ADD

If branch is taken, kill two following instructions. And because instruction in decode stage is invalid, update stall signal.



Derive IRSrcD Signal

Derive mux control signal for PCSrc.

if((opcodeE == BEQZ & z) + (opcodeE == BNEZ & !z)), PCSrc \leftarrow br
else if ((opcodeD == J) + (opcodeD == JAL)), PCSrc \leftarrow jabs
else if ((opcodeD == JR) + (opcodeD == JALR)), PCSrc \leftarrow rind
otherwise, PCSrc \leftarrow PC + 4

Derive mux control signal for IRSrcD.

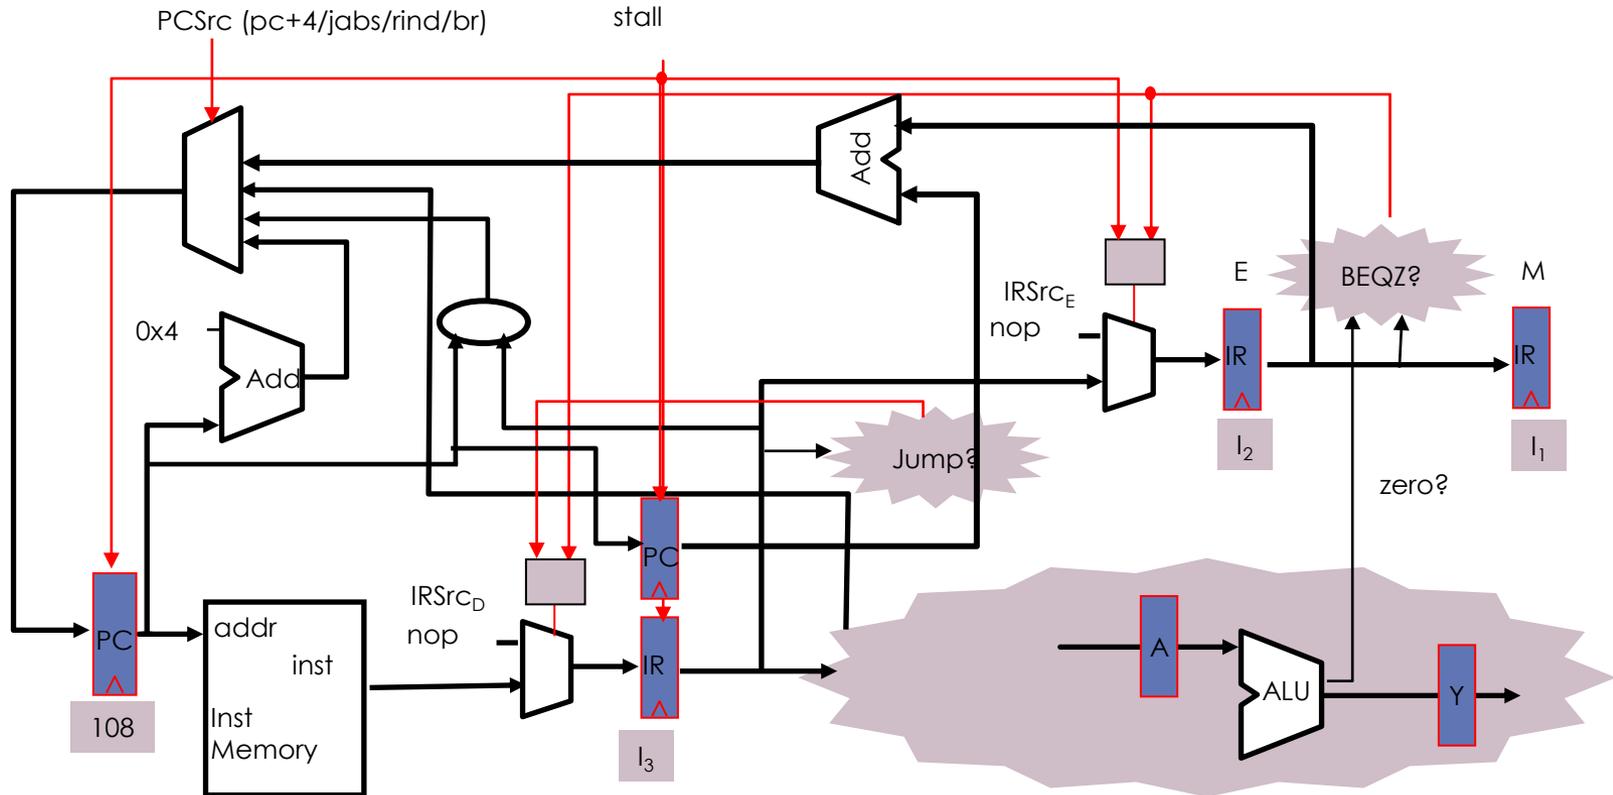
**if((opcodeE == BEQZ & z) + (opcodeE == BNEZ & !z)), ICSrcD \leftarrow nop
else if((opcodeD==J) + (opcodeD==JAL) +
 (opcodeD==JAR) + (opcodeD==JALR)), ICSrcD \leftarrow nop
otherwise, ICSrcD \leftarrow Instr**

Derive mux control signal for IRSrcE.

if((opcodeE == BEQZ & z) + (opcodeE == BNEZ & !z)), ICSrcE \leftarrow nop
otherwise, IRSrcE \leftarrow (stall & nop) + (!stall & IRD)



Pipelining Conditional Branches



	<u>addr</u>	<u>instr</u>
l_1	096	ADD
l_2	100	BEQZ r1 200
l_3	104	ADD
l_4	304	ADD

If branch is taken, kill two following instructions. And because instruction in decode stage is invalid, update stall signal.



Derive IRSrcE Signal

Derive mux control signal for PCSrc.

if((opcodeE == BEQZ & z) + (opcodeE == BNEZ & !z)), PCSrc \leftarrow br
else if ((opcodeD == J) + (opcodeD == JAL)), PCSrc \leftarrow jabs
else if ((opcodeD == JR) + (opcodeD == JALR)), PCSrc \leftarrow rind
otherwise, PCSrc \leftarrow PC + 4

Derive mux control signal for IRSrcD.

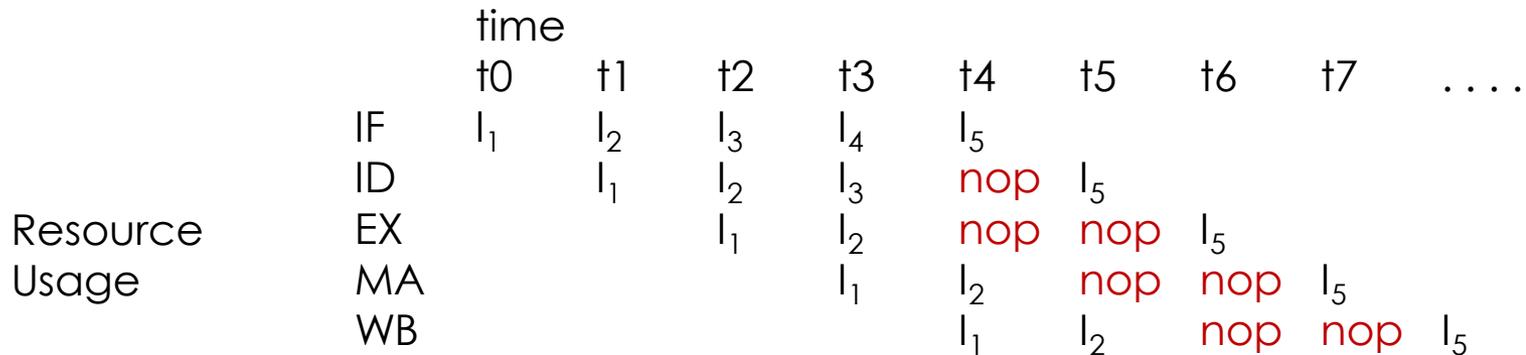
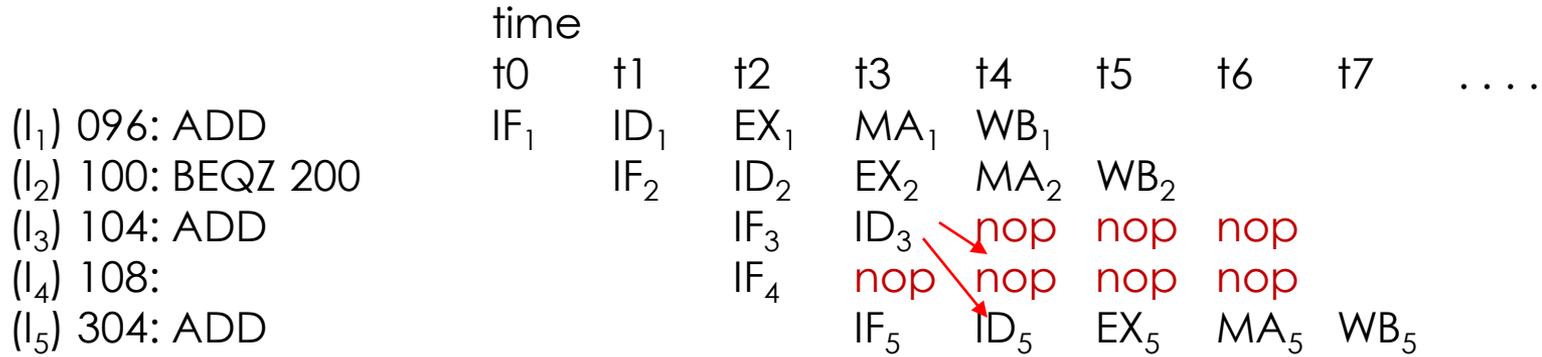
if((opcodeE == BEQZ & z) + (opcodeE == BNEZ & !z)), IRSrcD \leftarrow nop
else if((opcodeD==J) + (opcodeD==JAL) +
 (opcodeD==JAR) + (opcodeD==JALR)), IRSrcD \leftarrow nop
otherwise, IRSrcD \leftarrow Instr

Derive mux control signal for IRSrcE.

**if((opcodeE == BEQZ & z) + (opcodeE == BNEZ & !z)), IRSrcE \leftarrow nop
otherwise, IRSrcE \leftarrow (stall & nop) + (!stall & IRD)**



Pipelining Branches





Resolving Branch Conditions

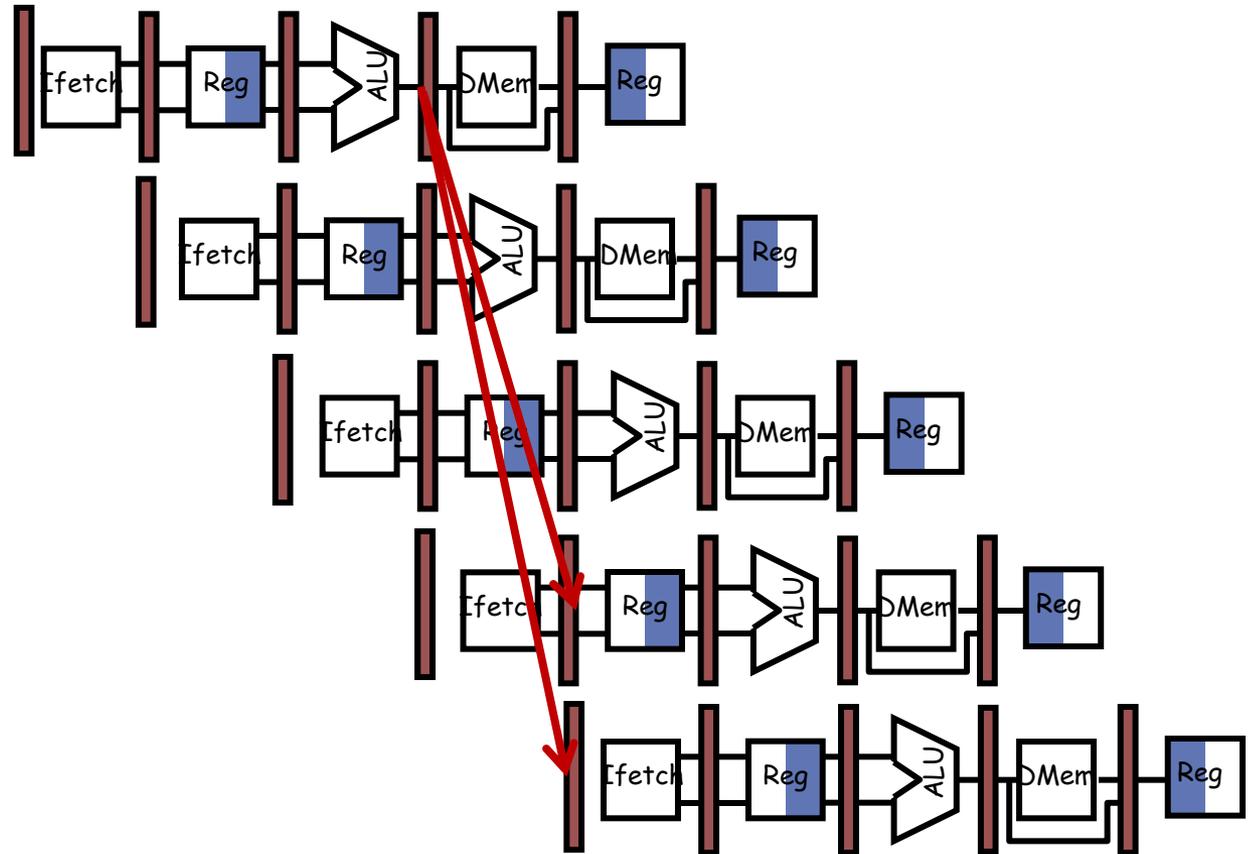
beq r1, r3, 36

and r2, r3, r5

or r6, r1, r7

add r8, r1, r9

xor r10, r1, r11



Three stage stall. What about the 3 instructions in between? Attempt to reduce. Stall/Kill signals in decode stage, reduces stalls from 3 to 2 stages.



Solution 1: Resolve Earlier

Large performance impact

- Suppose CPI = 1, 30% branch
- If branch stalls for 3 cycles, new CPI is 1.9

Solution – Branch Computation

- Determine whether branch is taken or not earlier in pipeline (e.g., beq)
- Compute target branch address earlier (e.g., PC addition)

Solution – MIPS

- MIPS branch tests if a register is equal to zero (e.g., beq)
- Move zero test to ID/RF stage
- Introduce adder to calculate new PC in ID/RF stage
- With early branch resolution and kill/stall signals in decode, require 1-cycle per branch, not 3-cycles

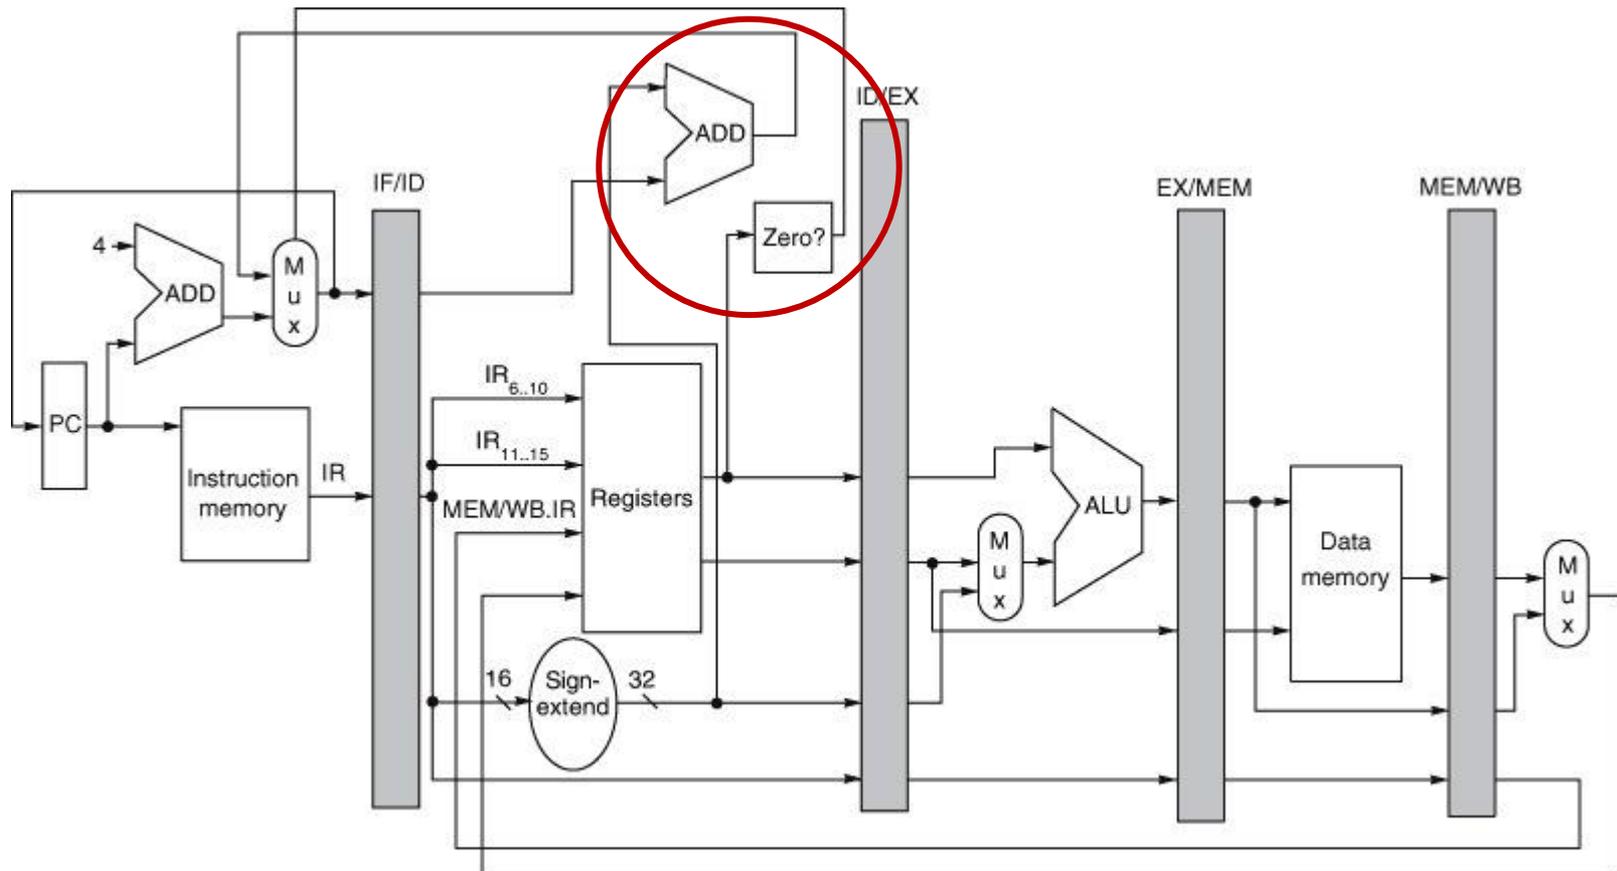


Pipelined MIPS Datapath

Figure A.24, Page A-38

Add sufficient logic in decode stage to generate “zero?” signal

Branch is resolved in ID stage instead of EX stage, eliminating one stall cycle



© 2007 Elsevier, Inc. All rights reserved.



Solution 2: Predict Condition

Stall until branch direction is clear

Predict Branch Not Taken

- Execute successor instructions in sequence
- “Squash” instructions in pipeline if branch taken
- Advantage: 47% of MIPS branches not taken
- Advantage: PC+4 already calculated for instruction fetch

Predict Branch Taken

- Advantage: 53% of MIPS branches taken
- Disadvantage: Target address not yet calculated, 1-cycle penalty

More sophisticated branch prediction later...



Solution 3: Change ISA Semantics

Delayed Branch

- Change ISA semantics so that instruction following jump/branch always executed.
- Define branch to take place after a following instruction
- Gives compiler flexibility to schedule useful instructions into a branch-induced stall

- Branch delay of length n

Branch instruction

Sequential successor 1

Sequential successor 2

...

Sequential successor n

Branch target if taken

- MIPS uses $n=1$ delay slot to calculate branch outcome, target address



Scheduling Delay Slots

Figure A.24, Page A-38



© 2007 Elsevier, Inc. All rights reserved.

(a) Fills delay slot and reduces instruction count, (b) DSUB needs copying and increases instruction count, (c) OR executes if branch fails so issue speculatively



Delayed Branches

Compiler Effectiveness (n=1 branch delay slot)

- Fill about 60% of branch delay slots
- About 80% of instructions executed are useful computation

Disadvantages of Delayed Branches

- As pipelines deepen, branch delay grows and requires more slots
- Less popular than dynamic approaches (e.g., branch prediction)



Pipelining in Practice

Why is $IPC < 1$?

Full forwarding may be too expensive to implement

- Implement only frequently used forwarding paths
- Implementing infrequently used forwarding paths might impact length of pipeline stage, increase clock period, and reduce IPC forwarding gains

Multi-cycle Instructions (e.g., loads)

- Instruction following a multi-cycle instruction cannot use its results
- MIPS-I defined load-delay slots, a software-visible pipeline hazard.
- Rely on compiler to schedule useful instructions, nops

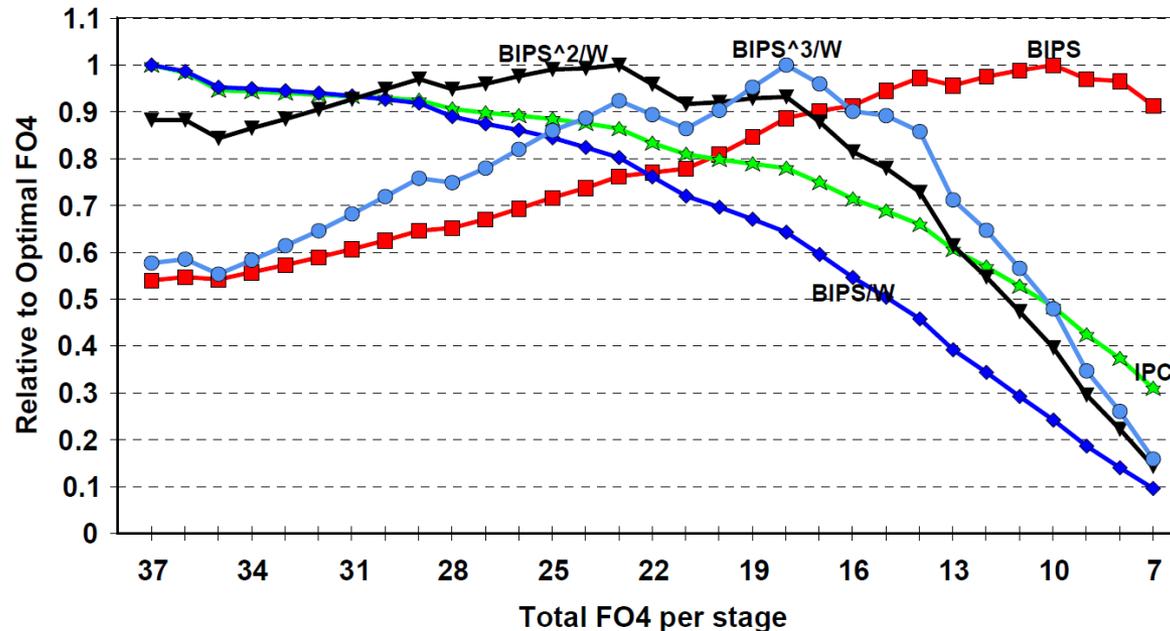
Conditional Branches

- Without delay slots, kill following instructions
- With delay slots, rely on compiler to schedule useful instructions, nops



Optimal Pipeline Depth

- Srinivasan et al. "Optimizing pipelines for power and performance," 2002.
- Performance (BIPS) versus Power (W)
- FO4 is measure of delay: Delay of inverter that is driven by inverter 4x smaller and that is driving inverter 4x larger.
- Quantify amount of logic per pipeline stage in FO4 delays
- (shorter delays \rightarrow deeper pipelines)

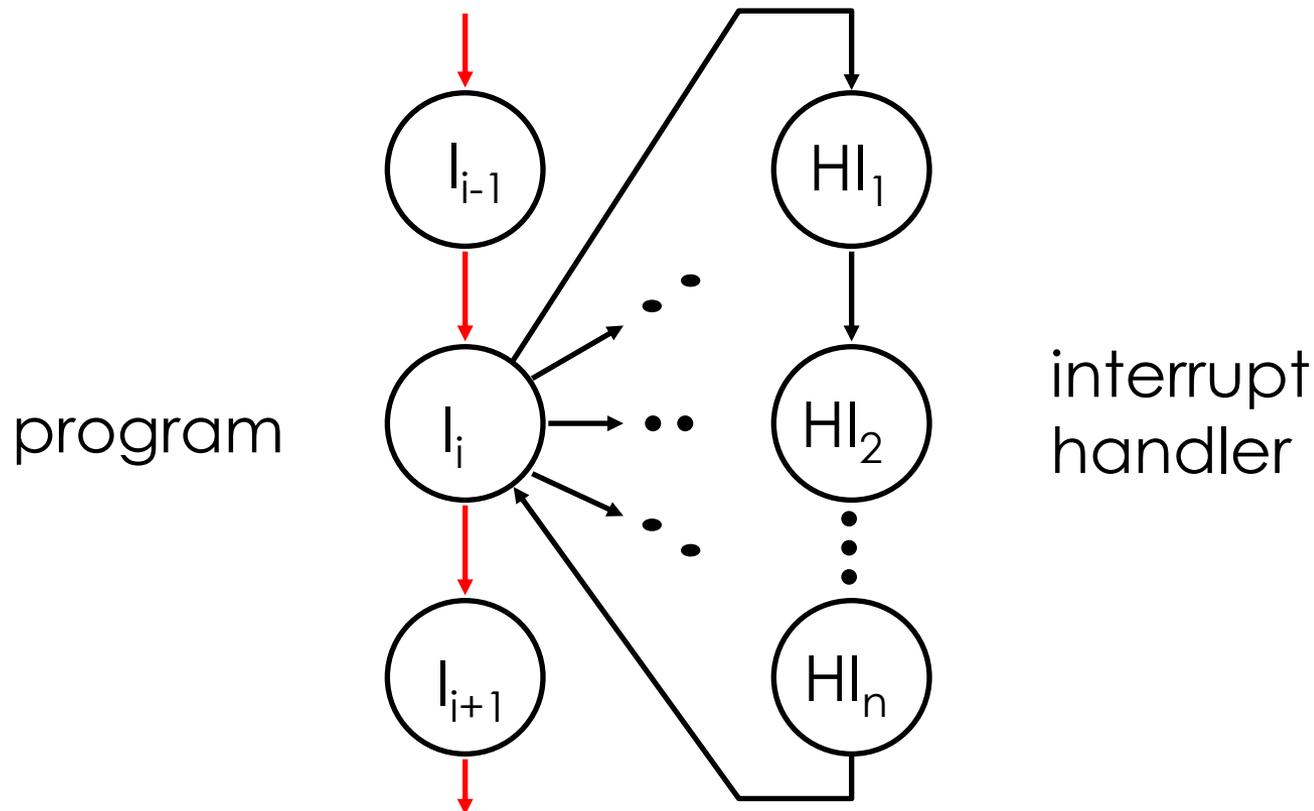




Interrupts and Exceptions

Interrupts alter normal control flow

- Event that needs to be processed by another (system) program
- Event is considered unexpected or rare from program's perspective





Causes of Interrupts

Interrupt

- An event that requests the attention of the processor

Asynchronous Interrupt – External Event

- Input/output device service-request
- Timer expiration
- Power disruptions, hardware failure

Synchronous Interrupt – Internal Event

- Undefined opcode, privileged instruction
- Arithmetic overflow, FPU exception
- Misaligned memory access
- Virtual memory exceptions – page faults, TLB misses, protection violations
- Traps – system calls, jumps into kernel
- Also known as exceptions



Asynchronous Interrupts

Service Request

- An I/O device requests attention by asserting one of the prioritized interrupt request lines

Invoking Interrupt Handler

- Processor decides to process interrupt
- Stops current program at instruction j , completing all instructions up to $j-1$. Defines a precise interrupt.
- Saves PC of instruction j in a special register (e.g., EPC)
- Disables interrupts and transfers control to designated interrupt handler running in kernel mode.



Interrupt Handler

PC Processing

- Save EPC before re-enabling interrupts, thereby allowing nested interrupts
- Need an instruction to move EPC into general-purpose registers
- Need a way to mask further interrupts until EPC saved

Status Register

- Read status register to determine cause of interrupt
- Executes handler code

Exiting Interrupt Handler

- Use special indirect jump instruction RFE (return-from-exception)
- Enables interrupts
- Restores processor to user mode
- Restores hardware status and control state



Synchronous Interrupts

Exceptions

- A synchronous interrupt (exception) is caused by a particular instruction

Instruction Re-start

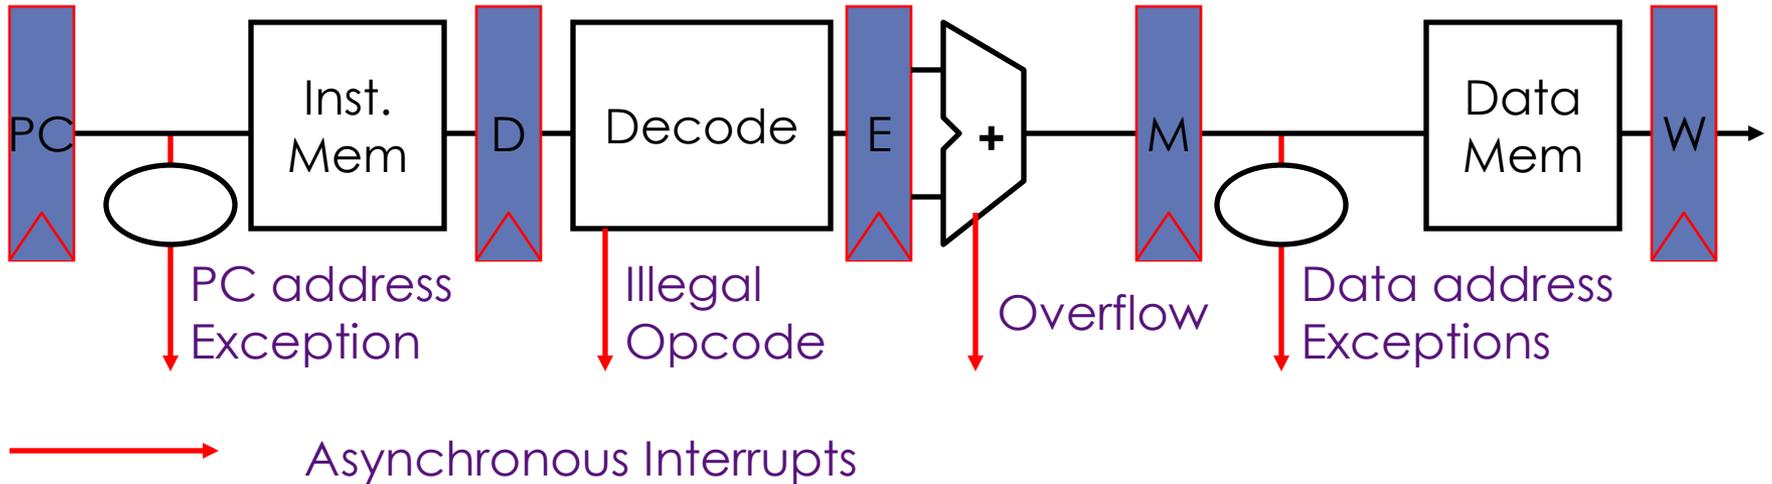
- Generally, instruction cannot be completed
- Instruction needs re-start after exception has been handled
- Processor must undo the effect of partially executed instructions

System Calls

- If the interrupt arises from a system calls (traps), trapping instruction considered complete
- System calls require a special jump instruction and changing into privileged kernel mode



Pipelining and Interrupt Handling

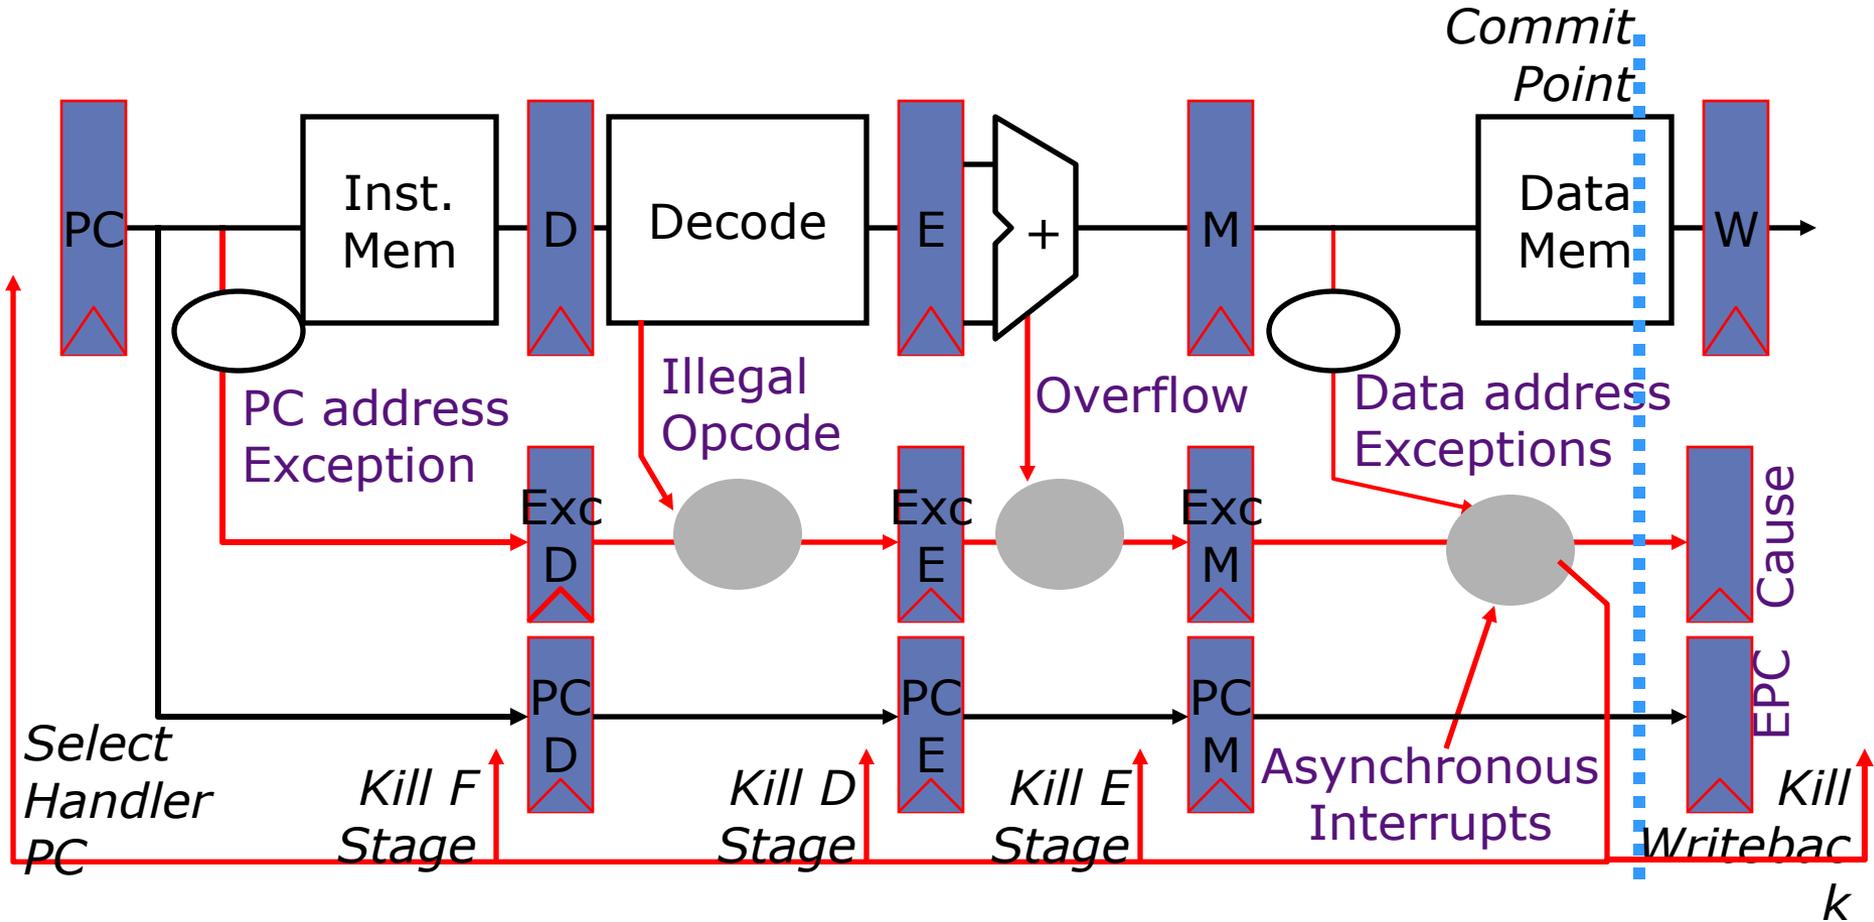


Synchronous: How does the processor handle multiple, simultaneous exceptions in different pipeline stages?

Asynchronous: How does the processor handle external interrupts?



Pipelining and Interrupt Handling





Pipelining and Interrupt Handling

Propagate exception flags through pipeline until commit point

Internal Interrupts

- An instruction might generate multiple exception flags
- For a given instruction, exceptions in earlier pipe stages over-ride those in later pipe stages, thereby prioritizing exceptions earlier in time

Inject external interrupts at commit point

- External interrupts over-ride internal interrupts

Check exception flags at commit point

- If exception flagged, update cause and EPC register
- Kill instructions in all pipeline stages
- Inject handler PC into fetch stage



Speculating about Exceptions

Predict

- Exceptions are rare. Predicting that no exceptions occurred is accurate.

Check Prediction

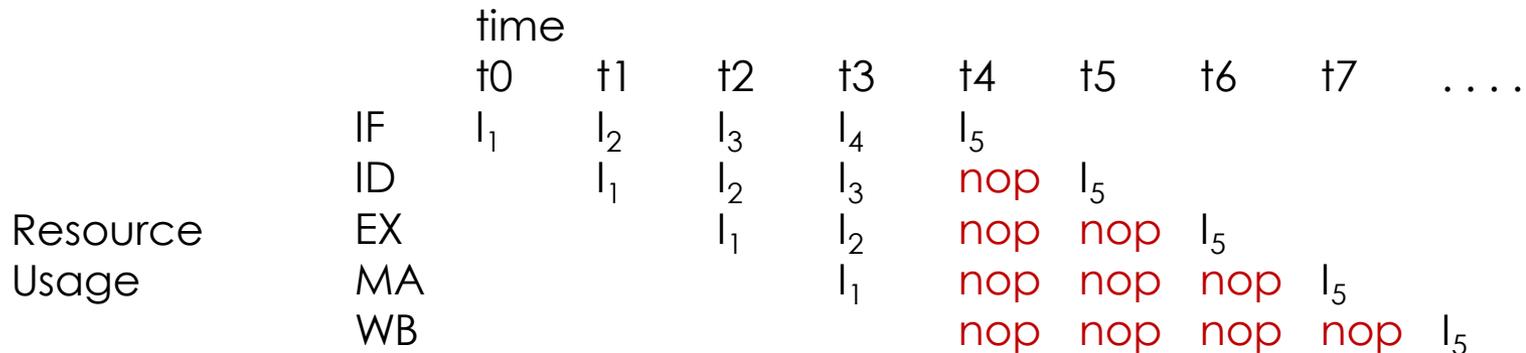
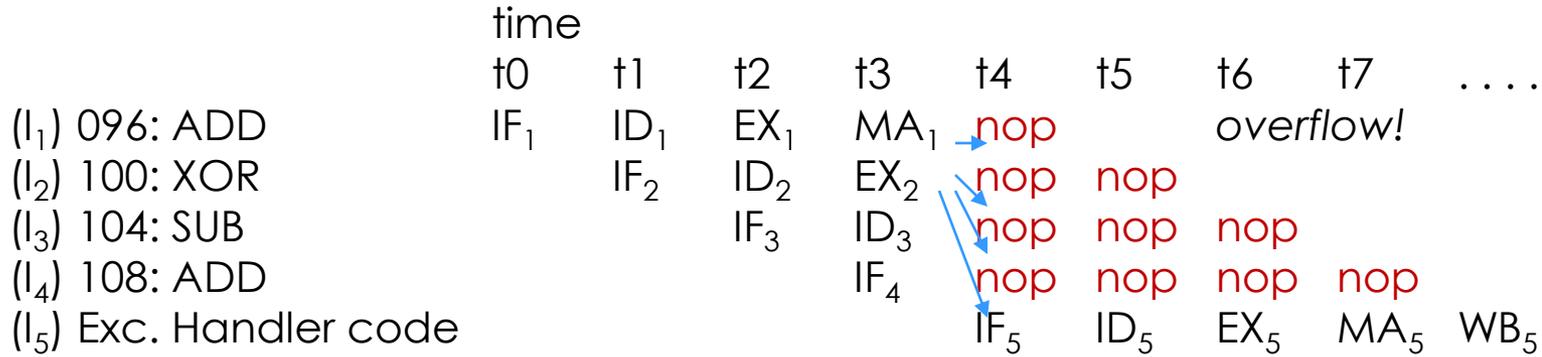
- Exceptions detected at end of pipeline (commit point). Invoke special hardware for various exception types

Recovery Mechanism

- Architectural state modified at end of pipeline (commit point).
- Discard partially executed instructions after an exception
- Launch exception handler after flushing pipeline



Pipelining and Exceptions





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)
- Alvin Lebeck (Duke)
- David Patterson (UCB)
- Daniel Sorin (Duke)