

Data and Control Hazards

Prof. Hakim Weatherspoon

CS 3410, Spring 2015

Computer Science

Cornell University

See P&H Chapter: 4.6-4.8

Announcements

Prelim next week

Tuesday at 7:30.

Go to location based on netid

[a-g]* → MRS146: Morrison Hall 146

[h-l]* → RRB125: Riley-Robb Hall 125

[m-n]* → RRB105: Riley-Robb Hall 105

[o-s]* → MVRG71: M Van Rensselaer Hall G71

[t-z]* → MVRG73: M Van Rensselaer Hall G73

Prelim reviews

Yesterday, Tue, Feb 24 @ 7:30pm in Olin 255

SATURDAY, Feb 28 @ 7:30pm in Upson B17

Prelim conflicts

Contact Deniz Altinbuken <deniz@cs.cornell.edu>

Announcements

Prelim1:

- Time: We will start at 7:30pm sharp, so come early
- Location: on previous slide
- Closed Book
 - Cannot use electronic device or outside material
- Practice prelims are online in CMS

Material covered everything up to end of this week

- Everything up to and including data hazards
- Appendix B (logic, gates, FSMs, memory, ALUs)
- Chapter 4 (pipelined [and non] MIPS processor with hazards)
- Chapters 2 (Numbers / Arithmetic, simple MIPS instructions)
- Chapter 1 (Performance)
- HW1, Lab0, Lab1, Lab2, C-Lab0, C-Lab1

Goals for Today

Data Hazards

- Data dependencies
- Problem, detection, and solutions
 - (delaying, stalling, forwarding, bypass, etc)
- Hazard detection unit
- Forwarding unit

Control Hazards

- What is the next instruction to execute if a branch is taken? Not taken?
- How to resolve control hazards
- Optimizations

Hazards

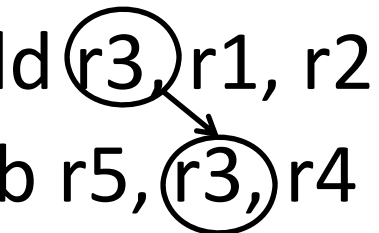
3 kinds

- Structural hazards
 - Multiple instructions want to use same unit
- Data hazards
 - Results of instruction needed before
- Control hazards
 - Don't know which side of branch to take

How to handle data hazards

- What to do if data hazard detected?

i.e. add (r3), r1, r2
sub r5, (r3), r4



- Options
 - Nothing
 - Change the ISA to match implementation
 - Stall
 - Pause current and subsequent instructions till safe
 - Slow down the pipeline (add bubbles to pipeline)
 - Forward/bypass
 - Forward data value to where it is needed

Data Hazards

Data Hazards

- register file reads occur in stage 2 (ID)
- register file writes occur in stage 5 (WB)
- next instructions may read values about to be written

Data Hazards

Stall

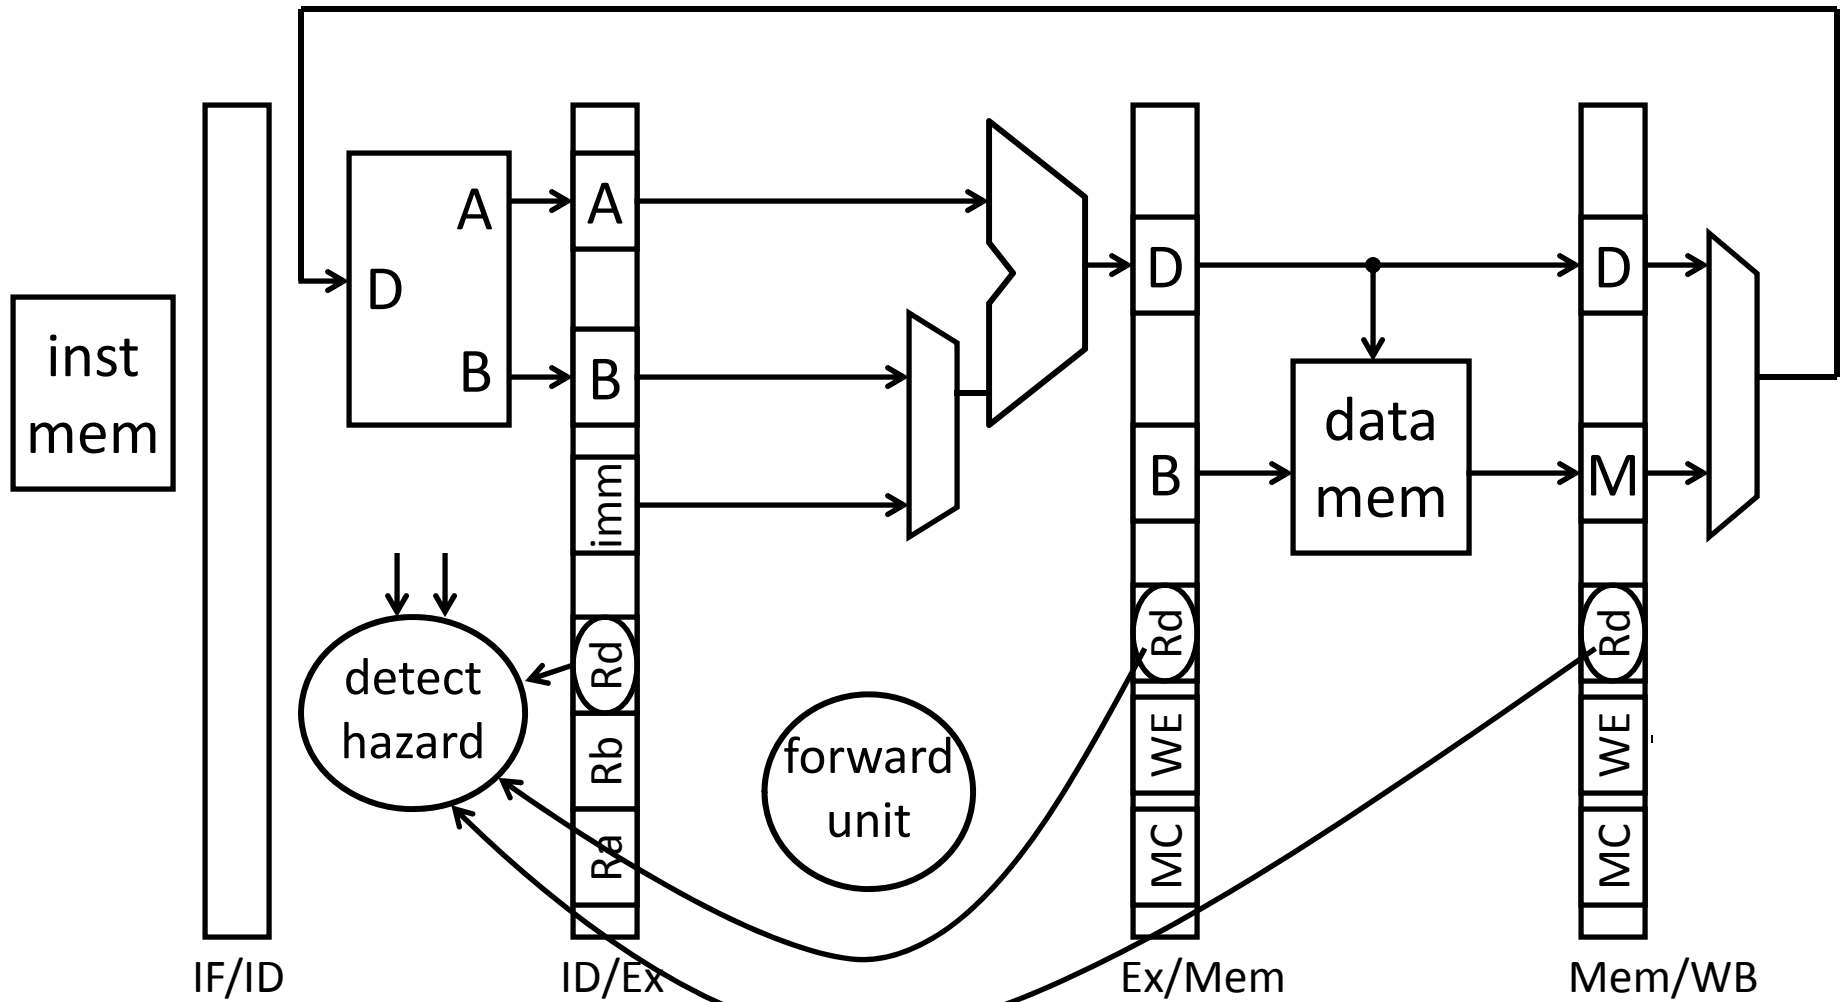
- Pause current and all subsequent instructions

Forward/Bypass

- Try to steal correct value from elsewhere in pipeline
- Otherwise, fall back to stalling or require a delay slot

Tradeoffs?

Data Hazards



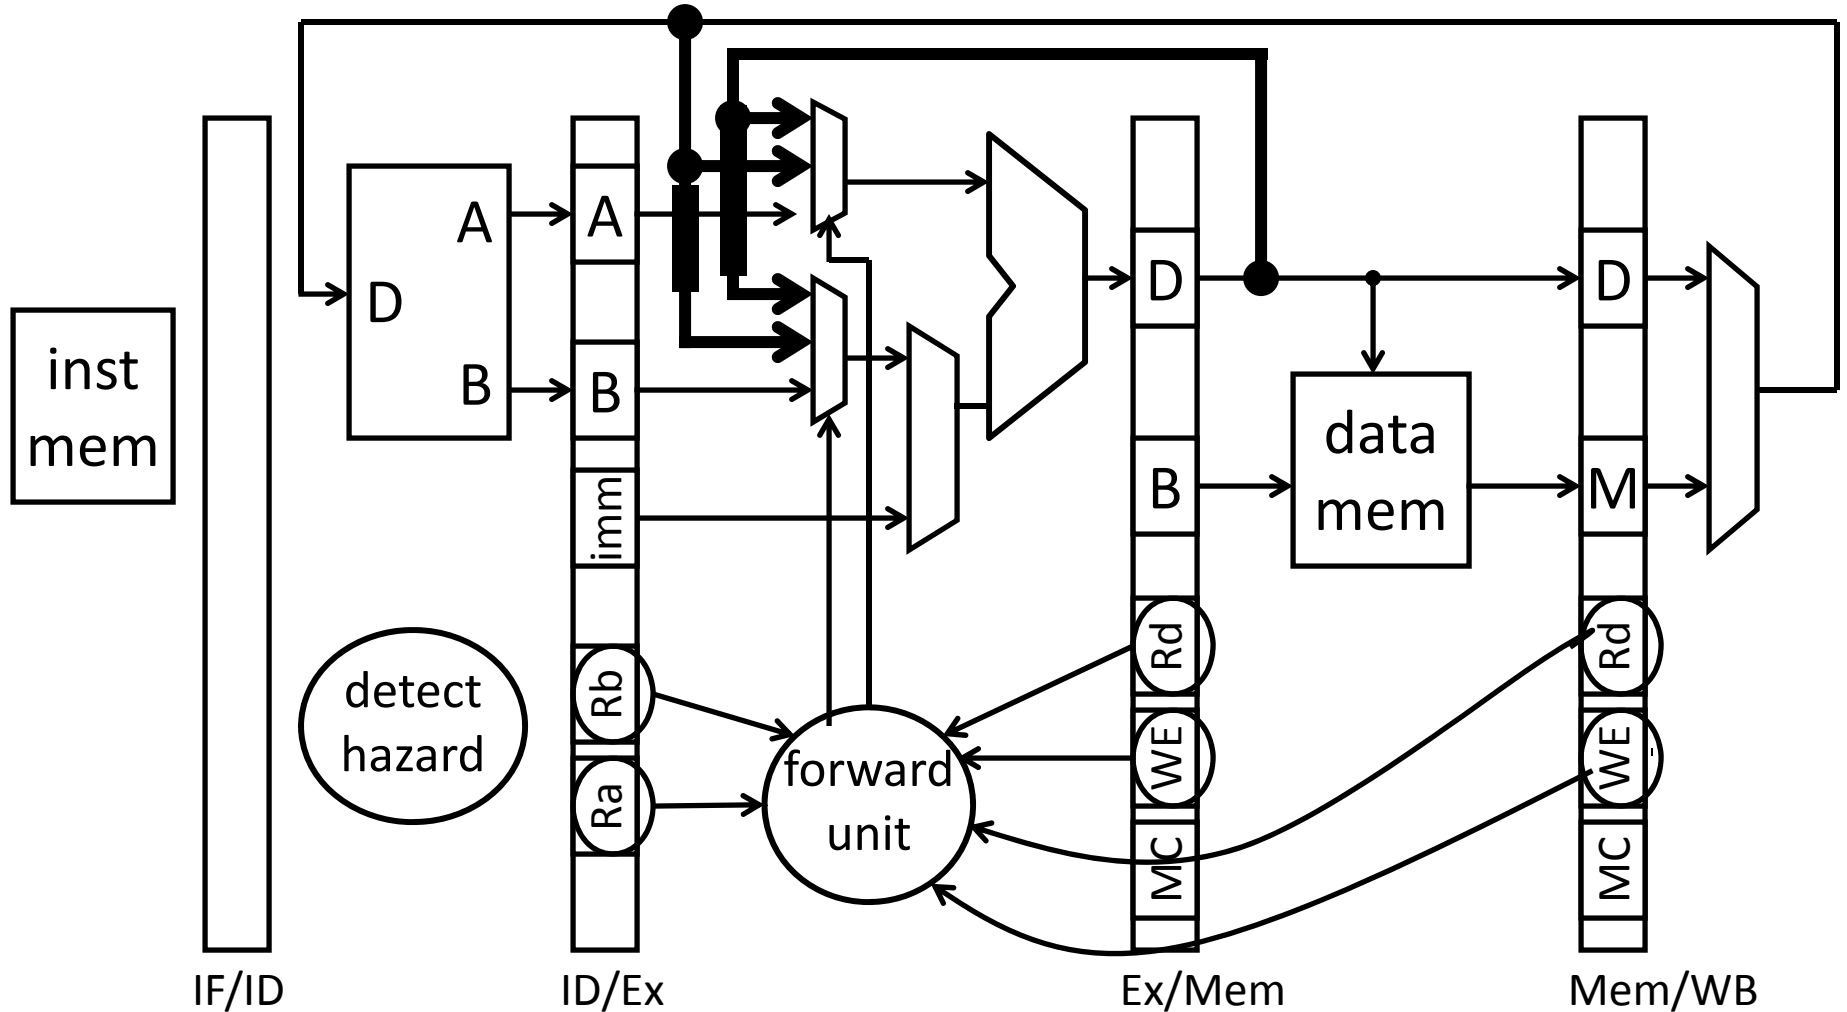
$$\text{stall} = \text{If}(\text{IF/ID.Ra} \neq 0 \ \&\&$$

$$(\text{IF/ID.Ra} == \text{ID/Ex.Rd}$$

$$\text{IF/ID.Ra} == \text{Ex/M.Rd}$$

$$\text{IF/ID.Ra} == \text{M/W.Rd}))$$

Data Hazards



Three types of forwarding/bypass

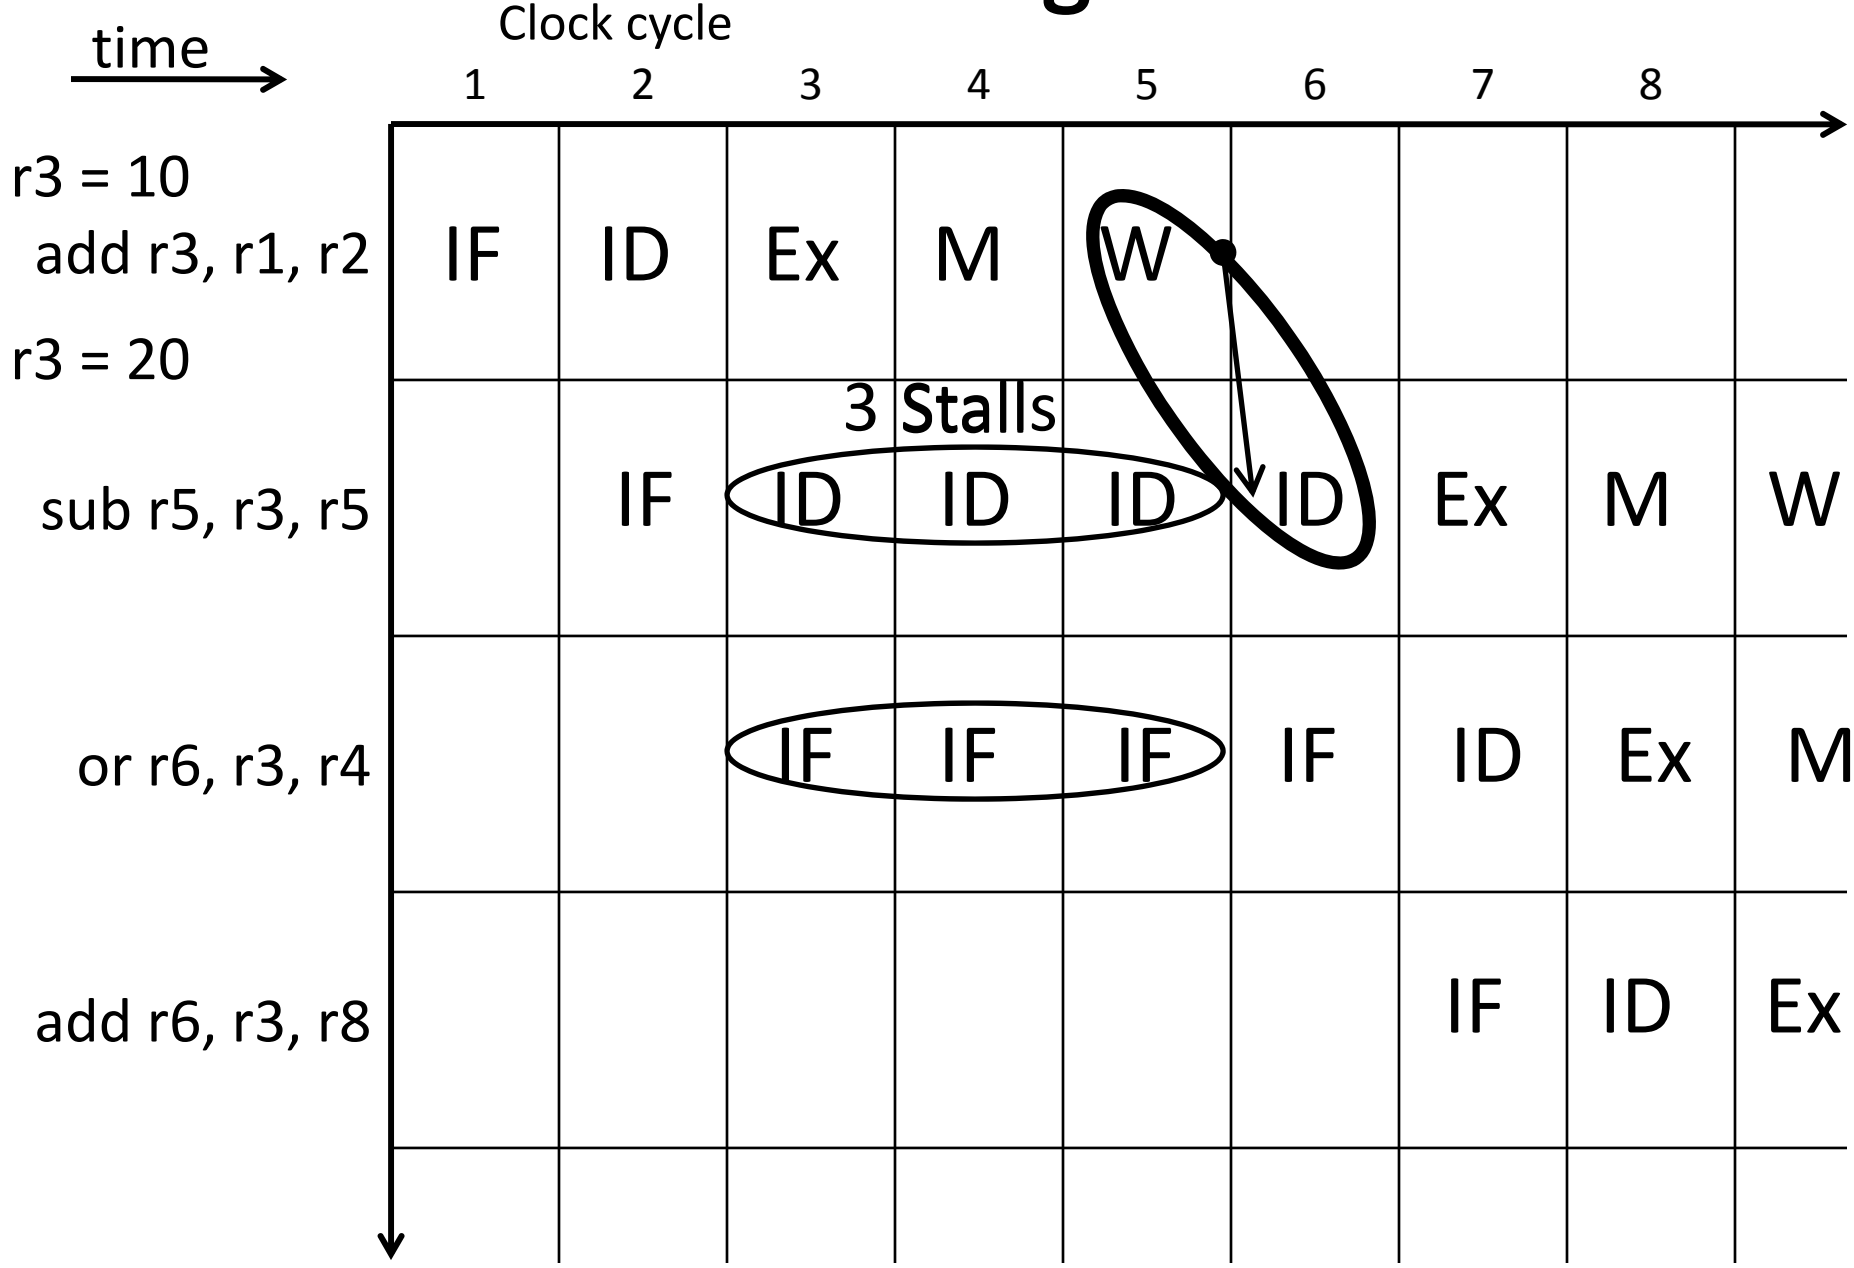
- Forwarding from Ex/Mem registers to Ex stage ($M \rightarrow Ex$)
- Forwarding from Mem/WB register to Ex stage ($W \rightarrow Ex$)
- RegisterFile Bypass

Stalling

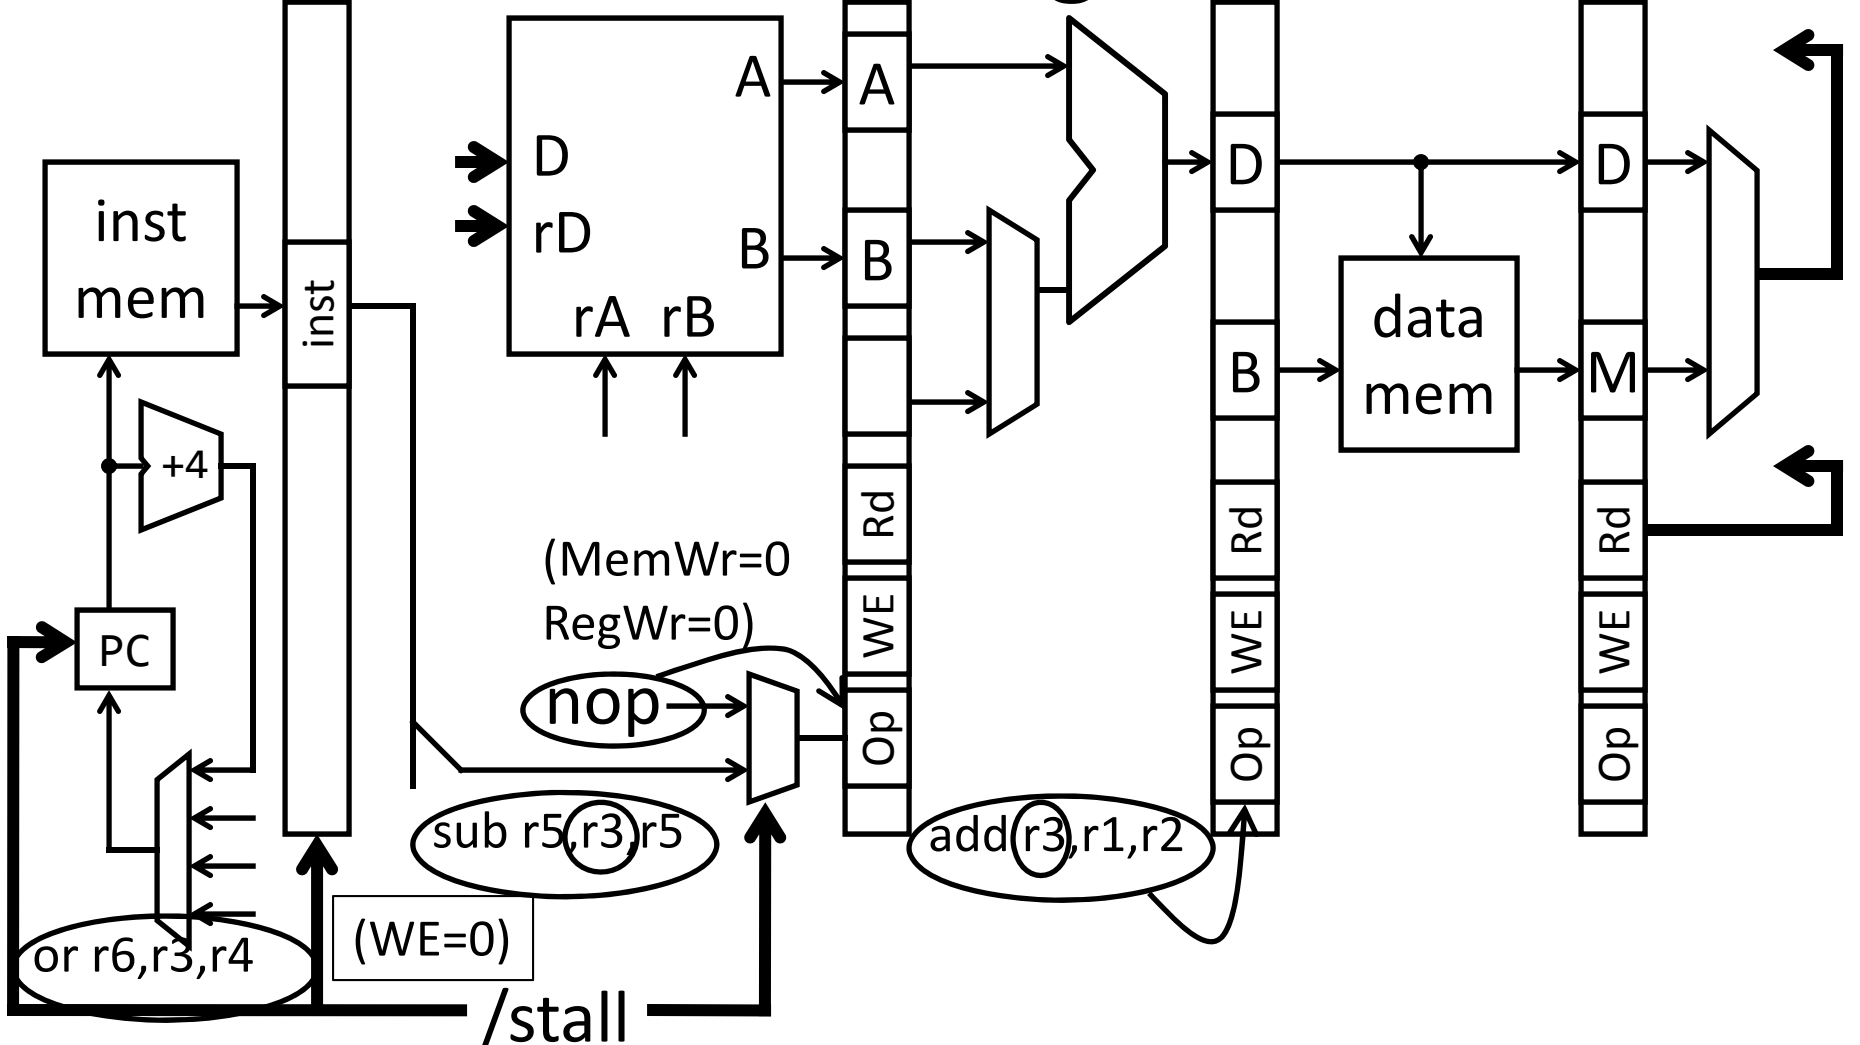
Pause current and all subsequent instructions

“slow down the pipeline”

Stalling

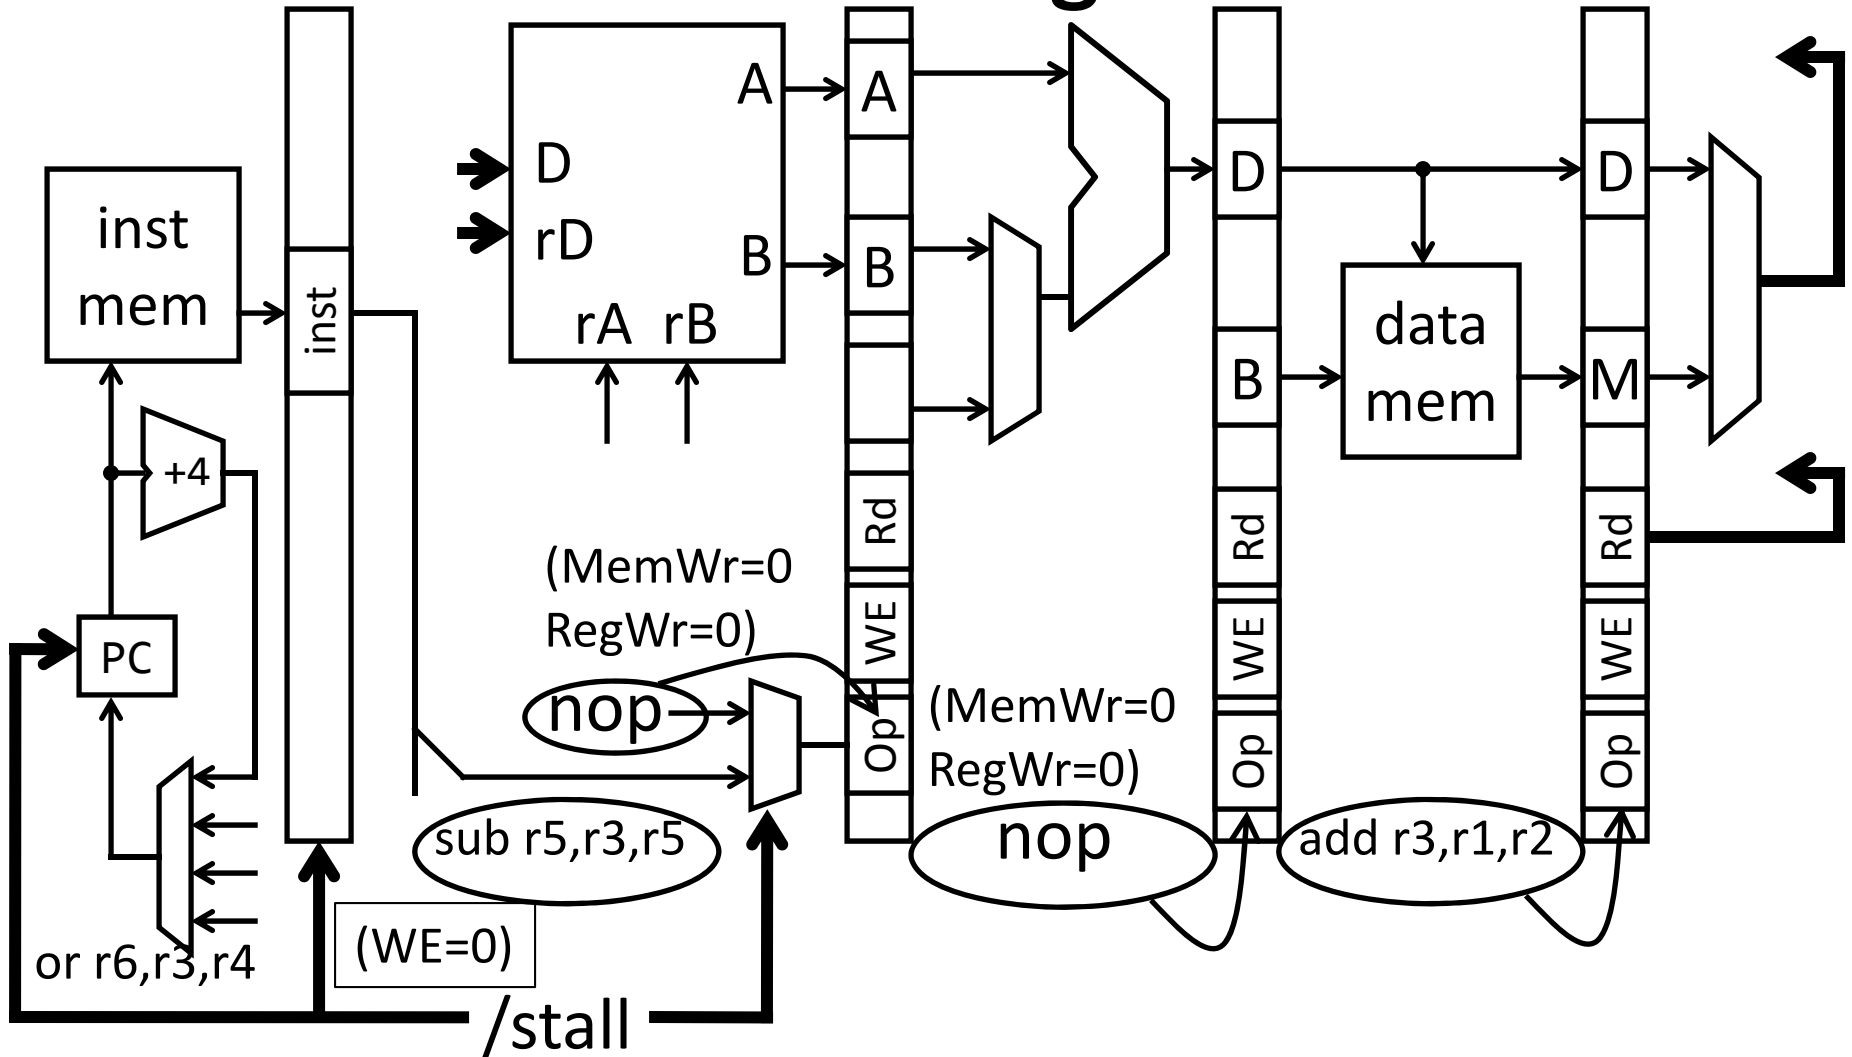


Stalling



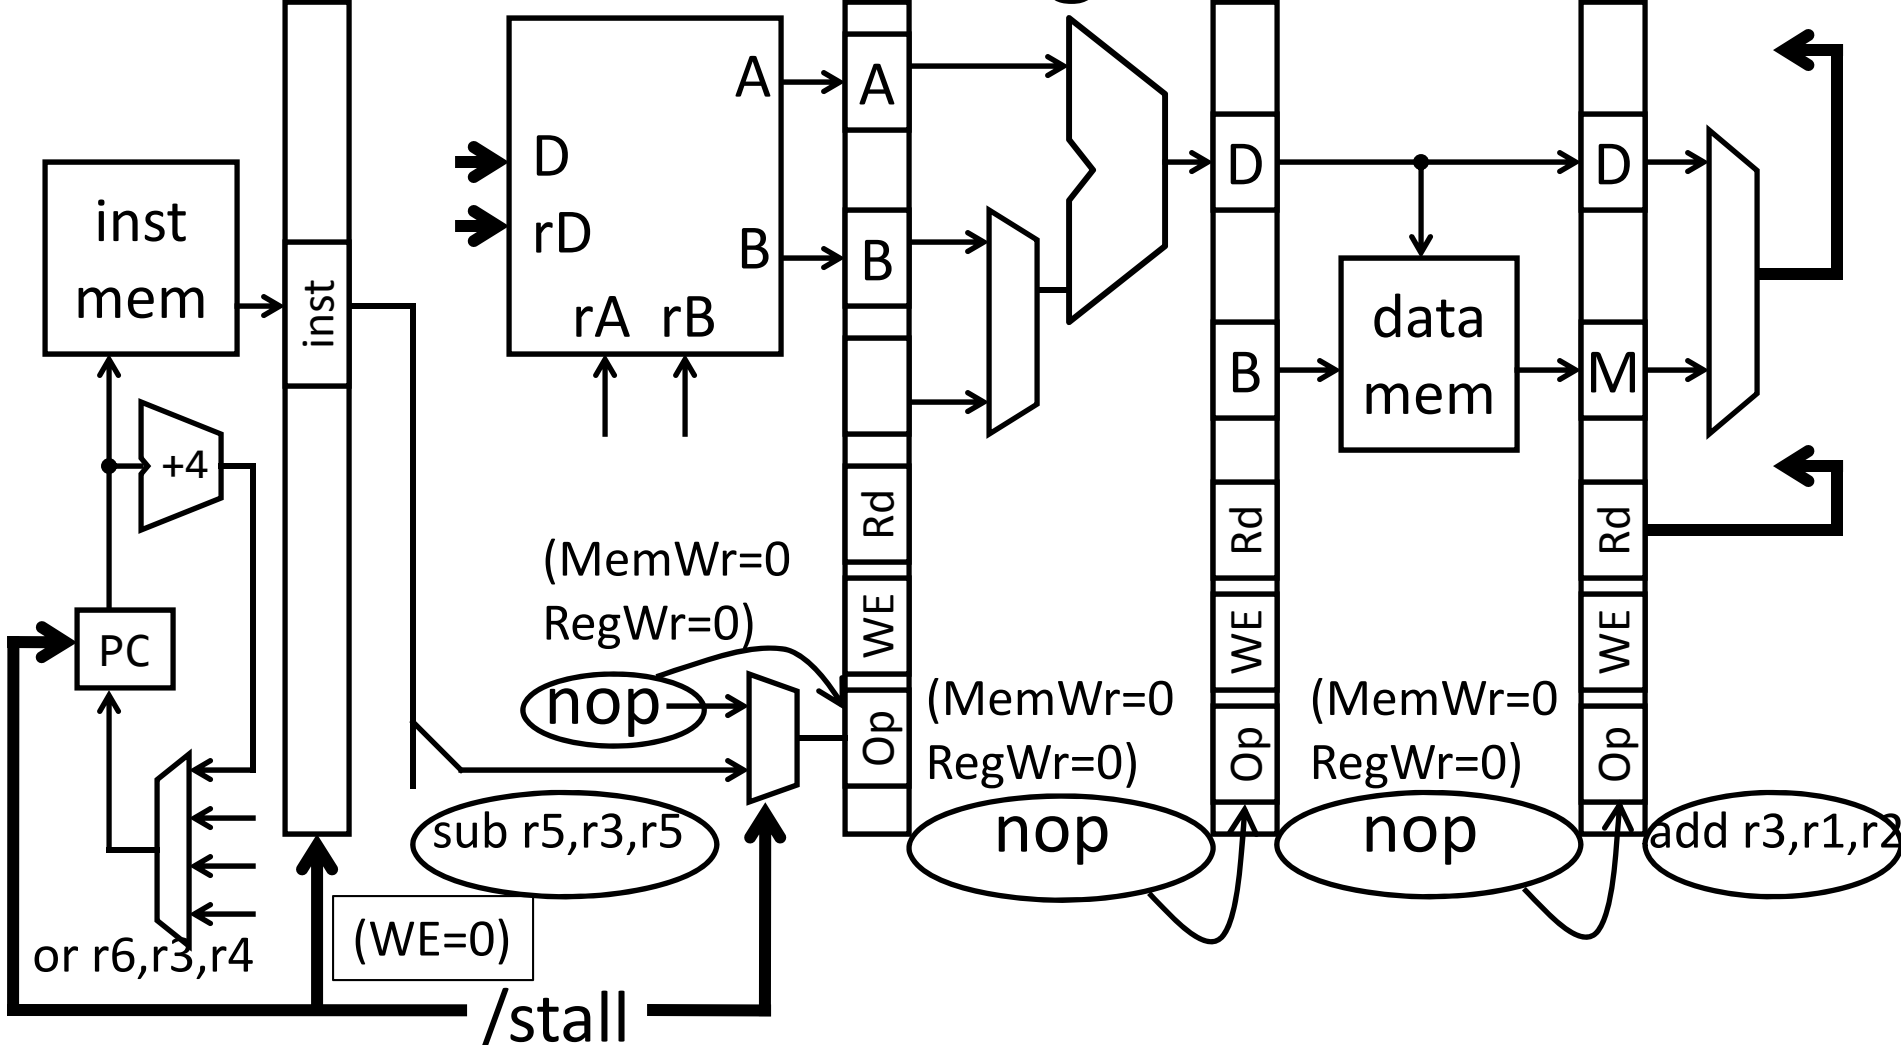
`NOP = If(IF/ID.rA ≠ 0 &&`
`(IF/ID.rA==ID/Ex.Rd`
`IF/ID.rA==Ex/M.Rd`
`IF/ID.rA==M/W.Rd))`

Stalling



$$\text{NOP} = \text{If}(\text{IF}/\text{ID}.\text{rA} \neq 0 \ \&\& \\
 (\text{IF}/\text{ID}.\text{rA} == \text{ID}/\text{Ex}.\text{Rd} \\
 \text{IF}/\text{ID}.\text{rA} == \text{Ex}/\text{M}.\text{Rd} \\
 \text{IF}/\text{ID}.\text{rA} == \text{M}/\text{W}.\text{Rd}))$$

Stalling



$NOP = \text{If}(\text{IF}/\text{ID}.rA \neq 0 \ \&\&$
 $\quad (\text{IF}/\text{ID}.rA == \text{ID}/\text{Ex}.Rd$
 $\quad \text{IF}/\text{ID}.rA == \text{Ex}/\text{M}.Rd$
 $\quad \text{IF}/\text{ID}.rA == \text{M}/\text{W}.Rd))$

Stalling

How to stall an instruction in ID stage

- prevent IF/ID pipeline register update
 - stalls the ID stage instruction
- convert ID stage instr into nop for later stages
 - innocuous “bubble” passes through pipeline
- prevent PC update
 - stalls the next (IF stage) instruction

Forwarding

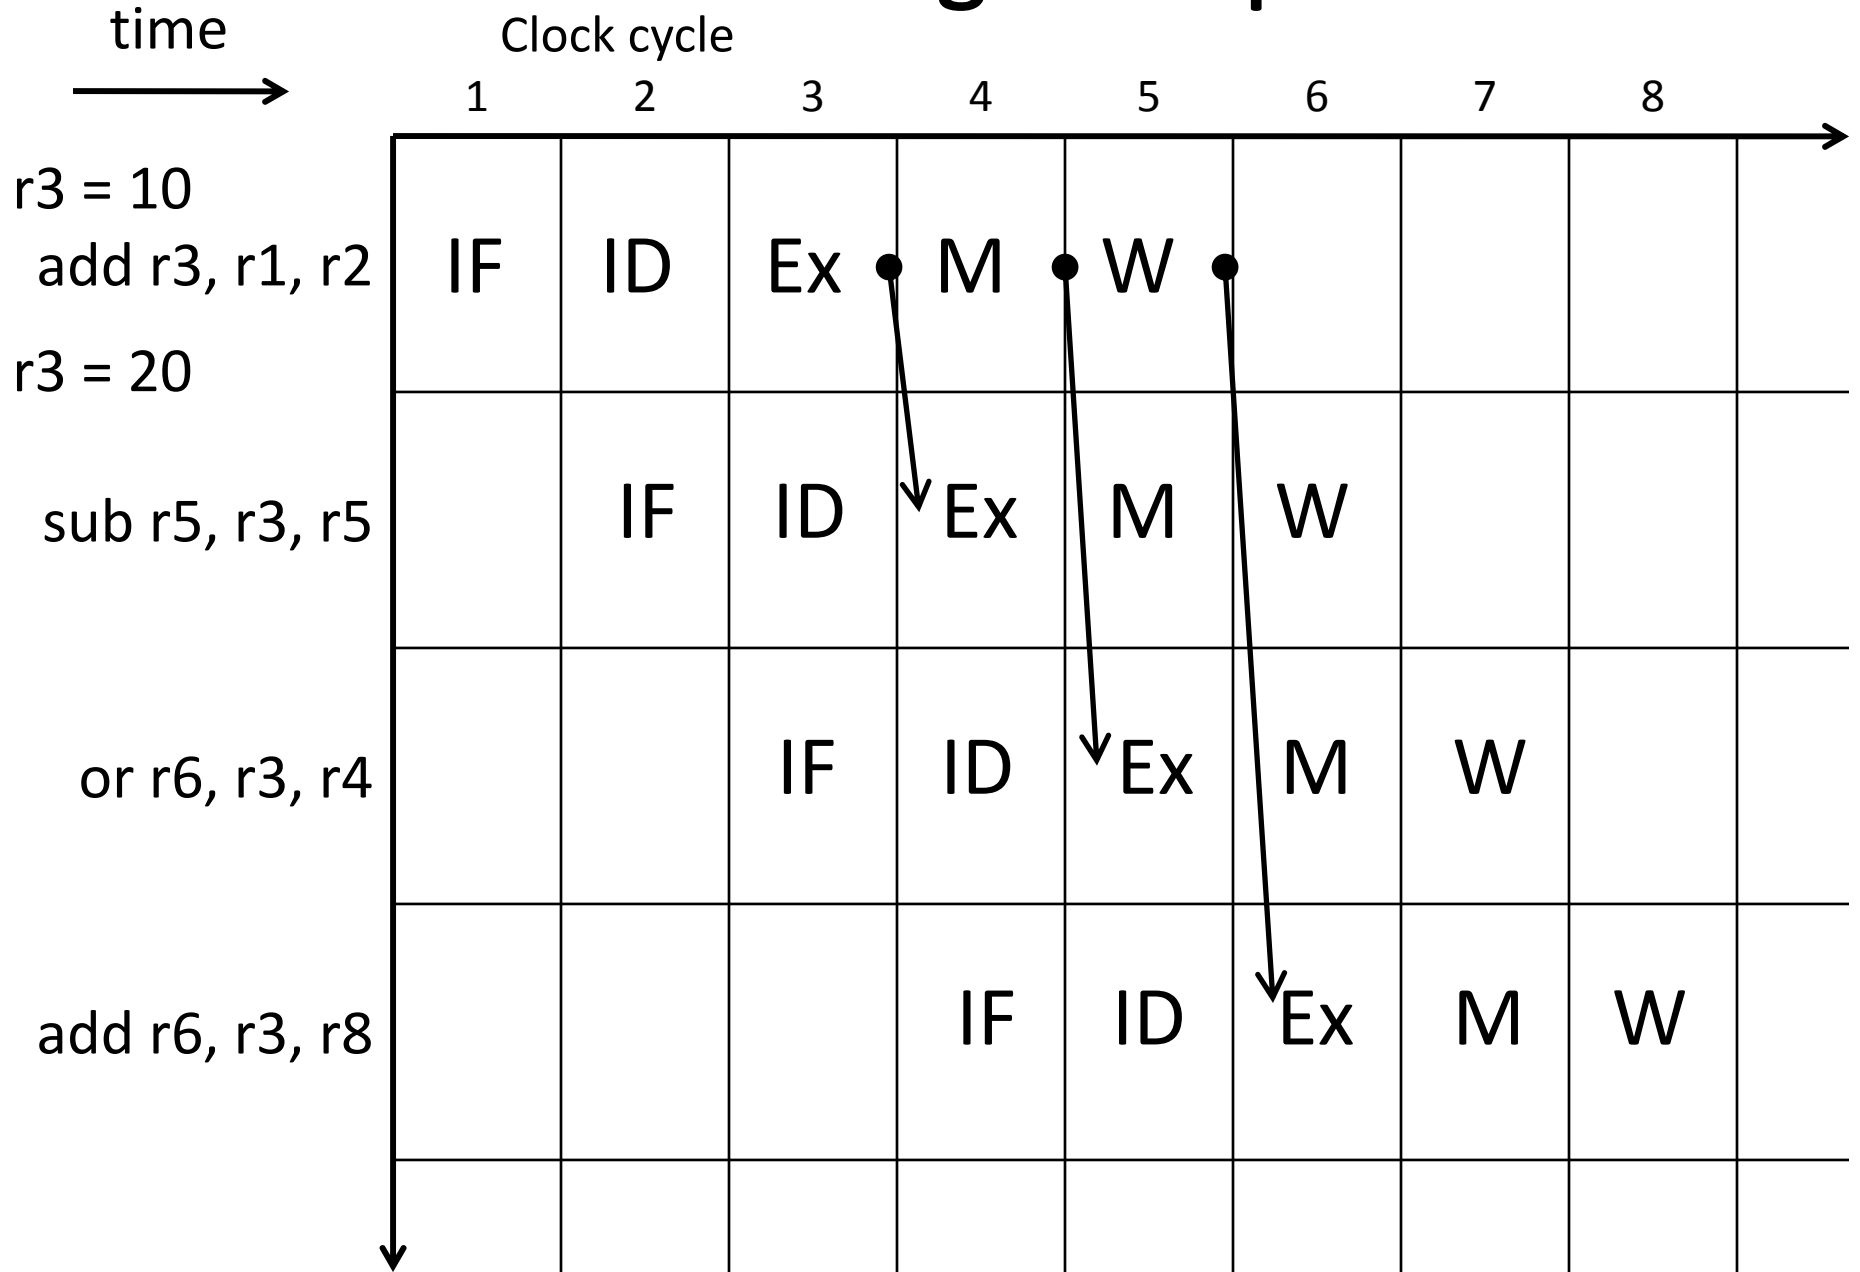
Alternative to stalling, we can forward

Forwarding bypasses some pipelined stages forwarding a result to a dependent instruction operand (register)

Three types of forwarding/bypass

- Forwarding from Ex/Mem registers to Ex stage (M→Ex)
- Forwarding from Mem/WB register to Ex stage (W→Ex)
- RegisterFile Bypass

Forwarding Example



Forwarding

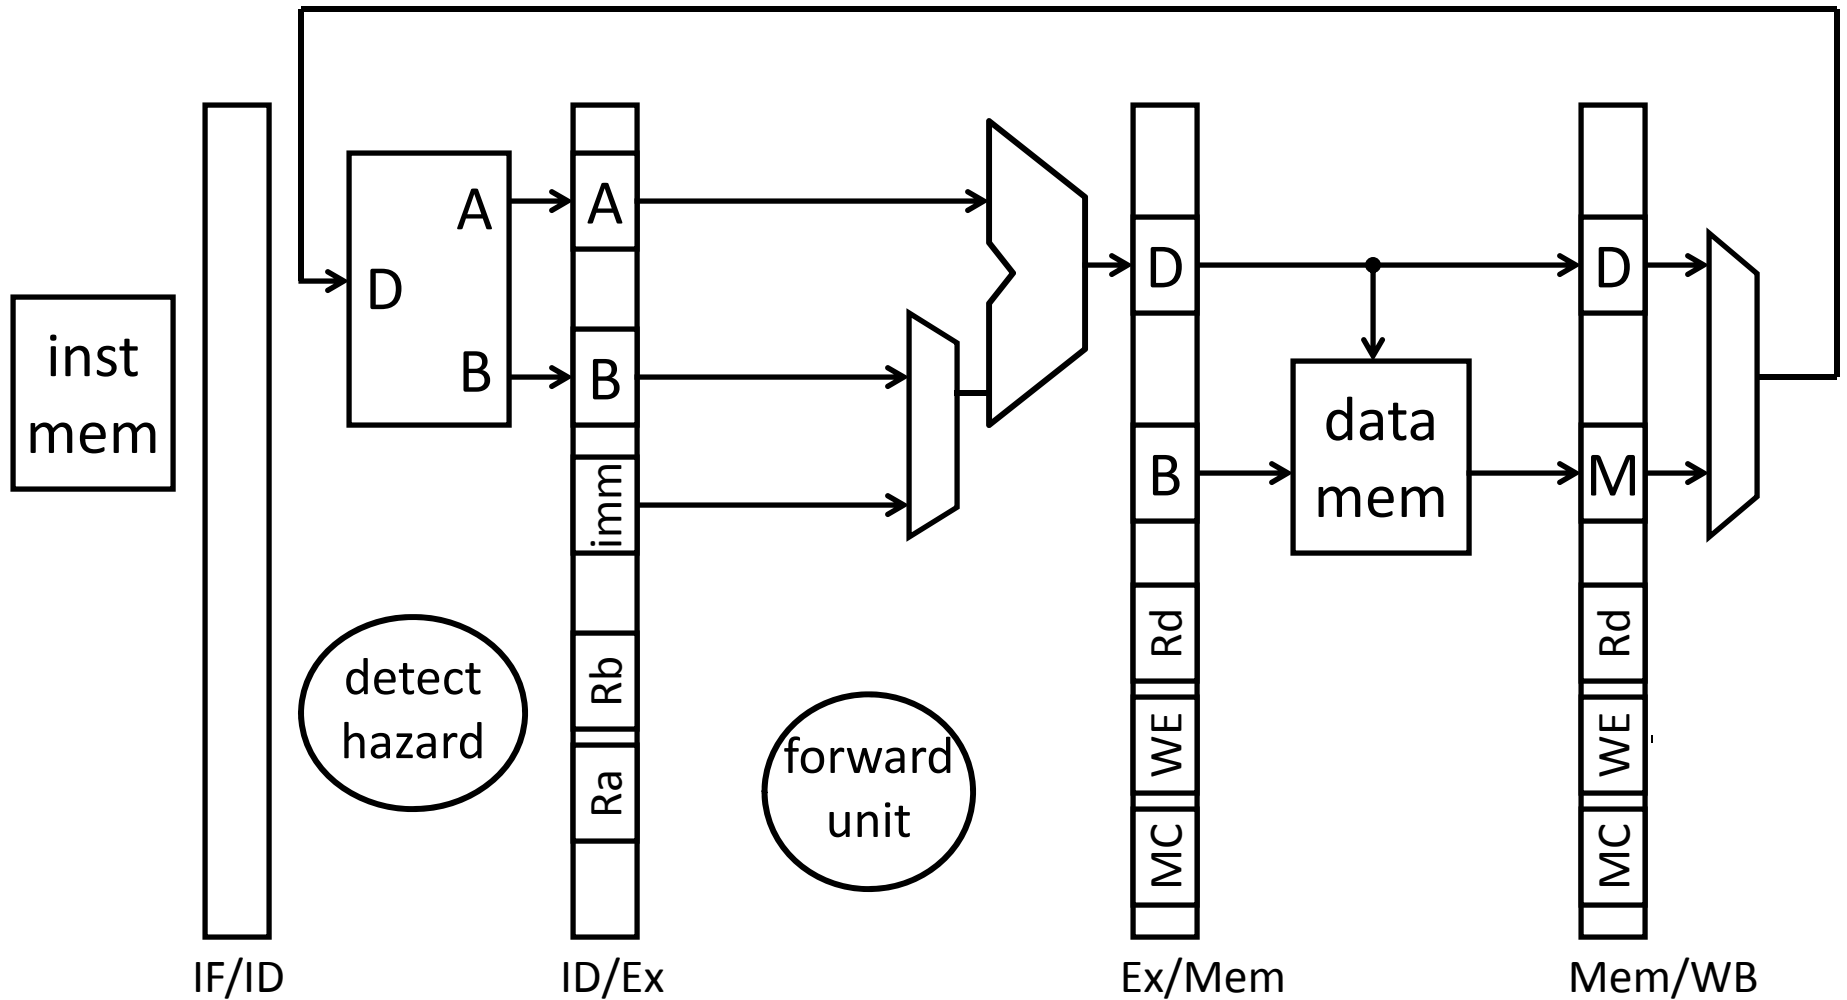
Alternative to stalling, we can forward

Forwarding bypasses some pipelined stages forwarding a result to a dependent instruction operand (register)

Three types of forwarding/bypass

- Forwarding from Ex/Mem registers to Ex stage (M→Ex)
- Forwarding from Mem/WB register to Ex stage (W→Ex)
- RegisterFile Bypass

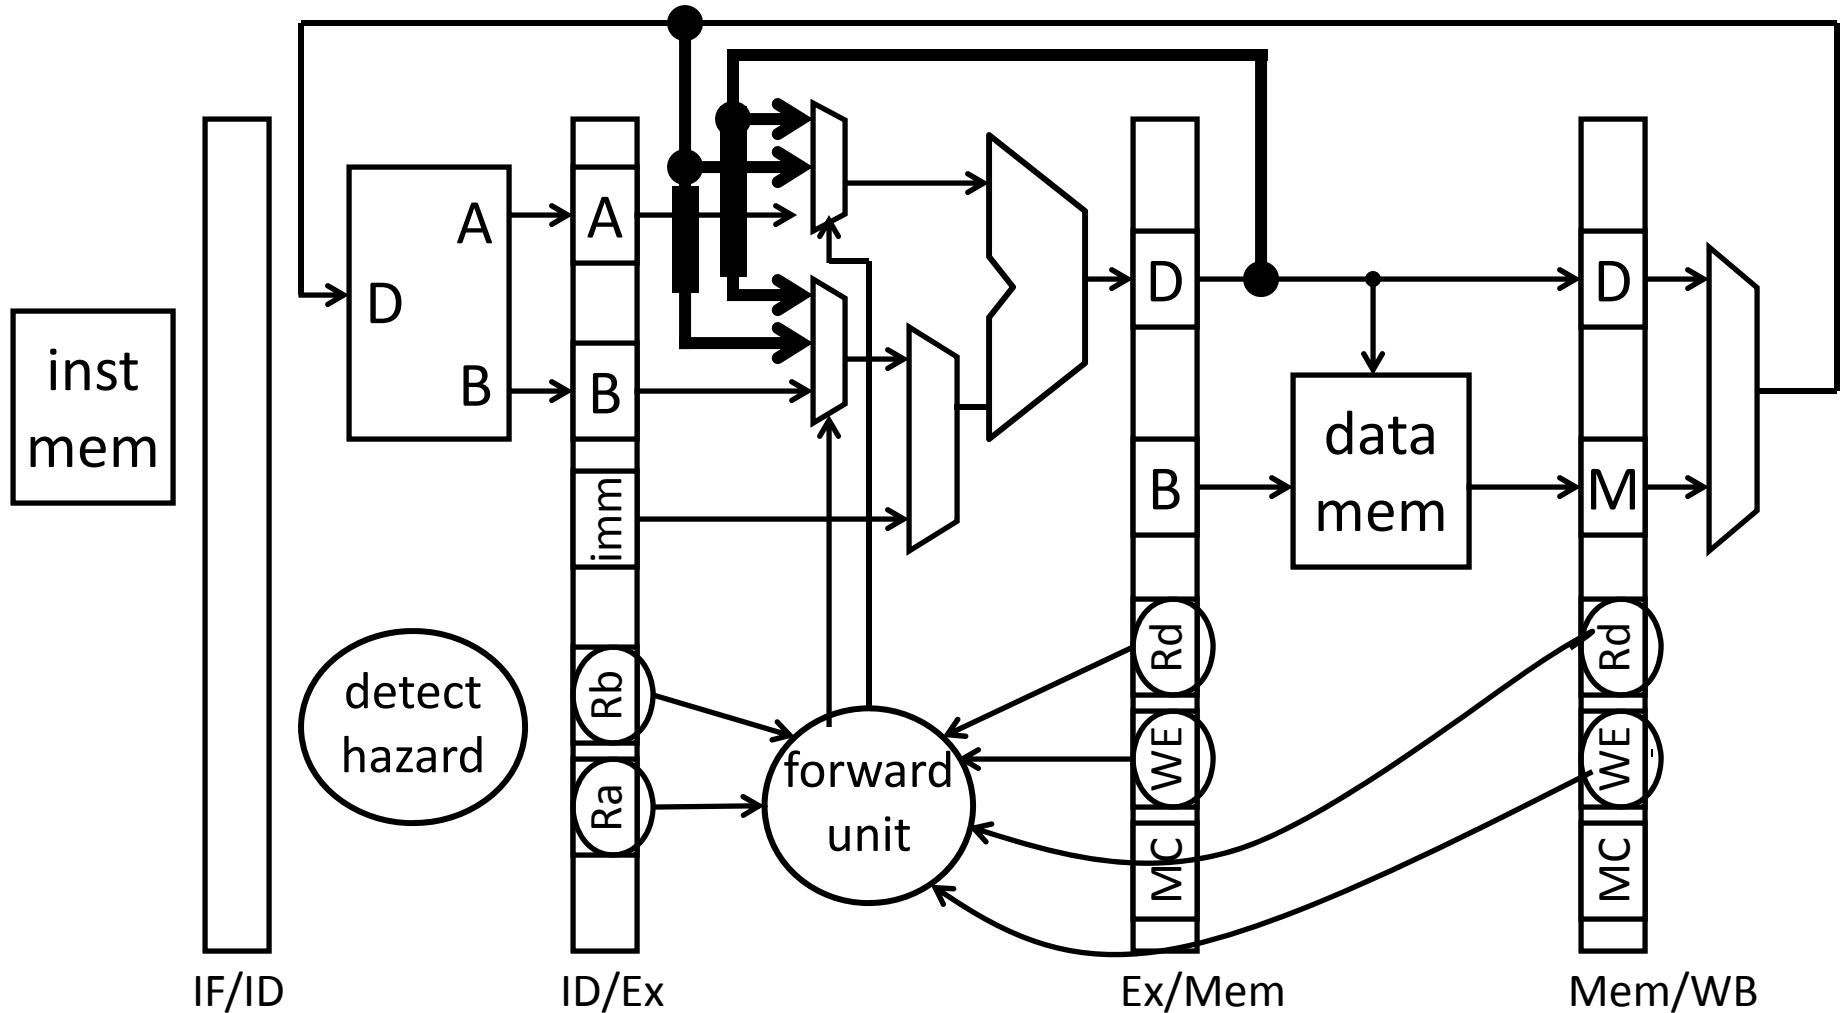
Forwarding Datapath



Three types of forwarding/bypass

- Forwarding from Ex/Mem registers to Ex stage (M→Ex)
- Forwarding from Mem/WB register to Ex stage (W → Ex)
- RegisterFile Bypass

Forwarding Datapath



Three types of forwarding/bypass

- Forwarding from Ex/Mem registers to Ex stage (M→Ex)
- Forwarding from Mem/WB register to Ex stage (W → Ex)
- RegisterFile Bypass

Forwarding Datapath 1

Ex/MEM to EX Bypass

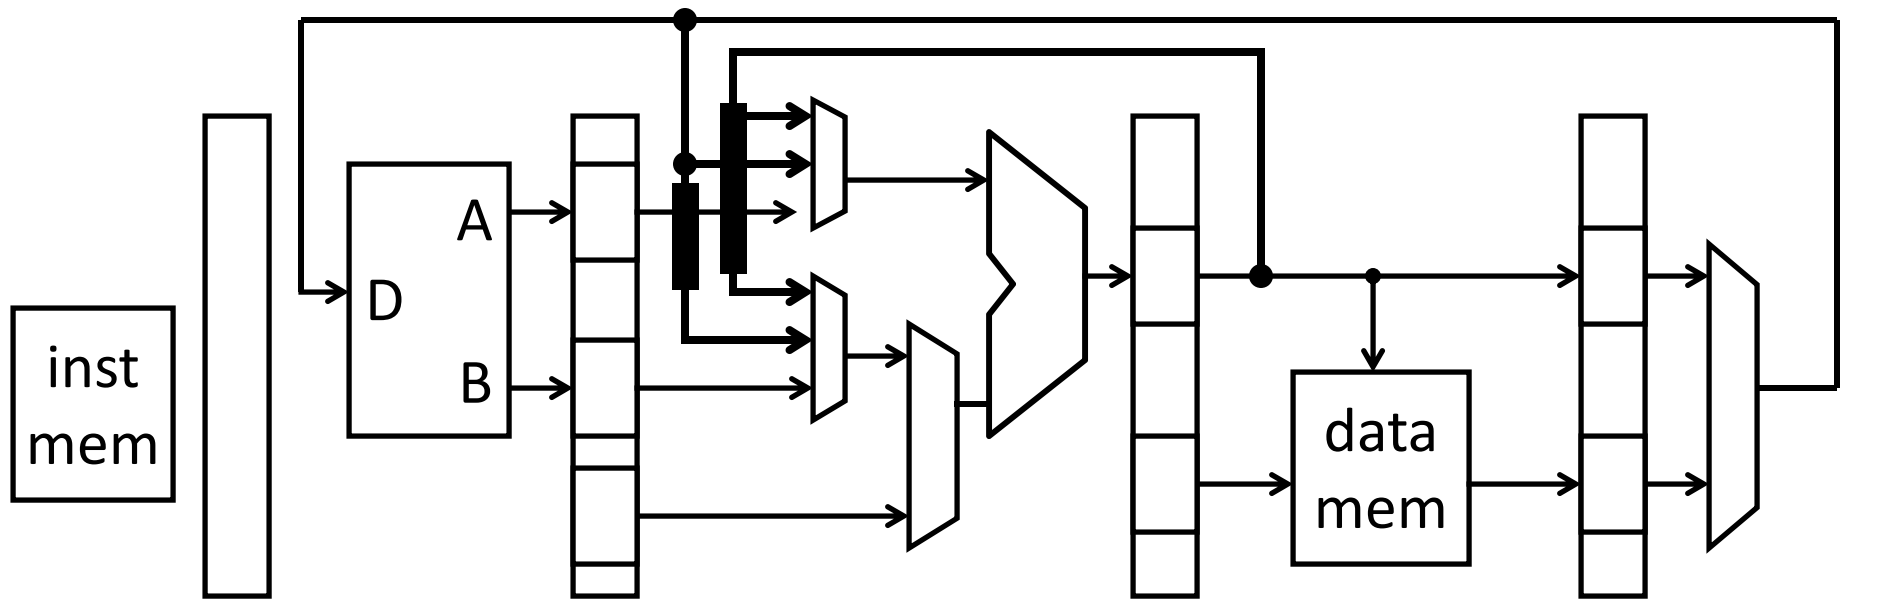
- EX needs ALU result that is still in MEM stage
- Resolve:

Add a bypass from EX/MEM.D to start of EX

How to detect? Logic in Ex Stage:

$$\text{forward} = (\text{Ex/M.WE} \ \&\& \ \text{EX/M.Rd} \neq 0 \ \&\& \\ \text{ID/Ex.Ra} == \text{Ex/M.Rd}) \\ || \text{ (same for Rb)}$$

Forwarding Datapath 2



add r3, r1, r2	IF	ID	Ex	M	W			
sub r5, r3, r1		IF	ID	Ex	M	W		
or r6, r3, r4			IF	ID	Ex	M	W	

Arrows in the table indicate forwarding: from the M stage of 'add r3, r1, r2' to the Ex stage of 'sub r5, r3, r1', and from the W stage of 'add r3, r1, r2' to the Ex stage of 'or r6, r3, r4'.

Forwarding Datapath 2

Mem/WB to EX Bypass

- EX needs value being written by WB
- Resolve:

Add bypass from WB final value to start of EX

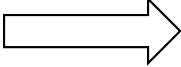
How to detect? Logic in Ex Stage:

$$\text{forward} = (\text{M/WB.WE} \ \&\& \ \text{M/WB.Rd} \ \neq \ 0 \ \&\& \\ \text{ID/Ex.Ra} \ == \ \text{M/WB.Rd} \ \&\& \\ \ || \ \text{(same for Rb)})$$

Is this it?

Not quite!

Forwarding Datapath 2

add r3, r1, r2		add r3, r1, r2
sub r5, r3, r5		sub r3, r3, r5
or r6, r3, r4		or r6, r3, r4
add r6, r3, r8		add r6, r3, r8

How to detect? Logic in Ex Stage. Forward from M/WB reg if:
M/WB (WE on, Rd != 0) and (M/WB.Rd == ID/Ex.Ra)
also NOT(Ex/M.Rd == ID/Ex.Ra) and (WE, Rd != 0))

Rb same as Ra

Check pg. 311

Register File Bypass

Register File Bypass

- Reading a value that is currently being written

Detect:

$((Ra == MEM/WB.Rd) \text{ or } (Rb == MEM/WB.Rd))$
and (WB is writing a register)

Resolve:

Add a bypass around register file (WB to ID)

Better: just negate register file clock

- writes happen at end of first half of each clock cycle
- reads happen during second half of each clock cycle

Are we done yet?

add r3, r1, r2

lw(r4), 20(r8)

or r6, r3, r4

add r6, r3, r8

Memory “Load-Use Data Hazard”

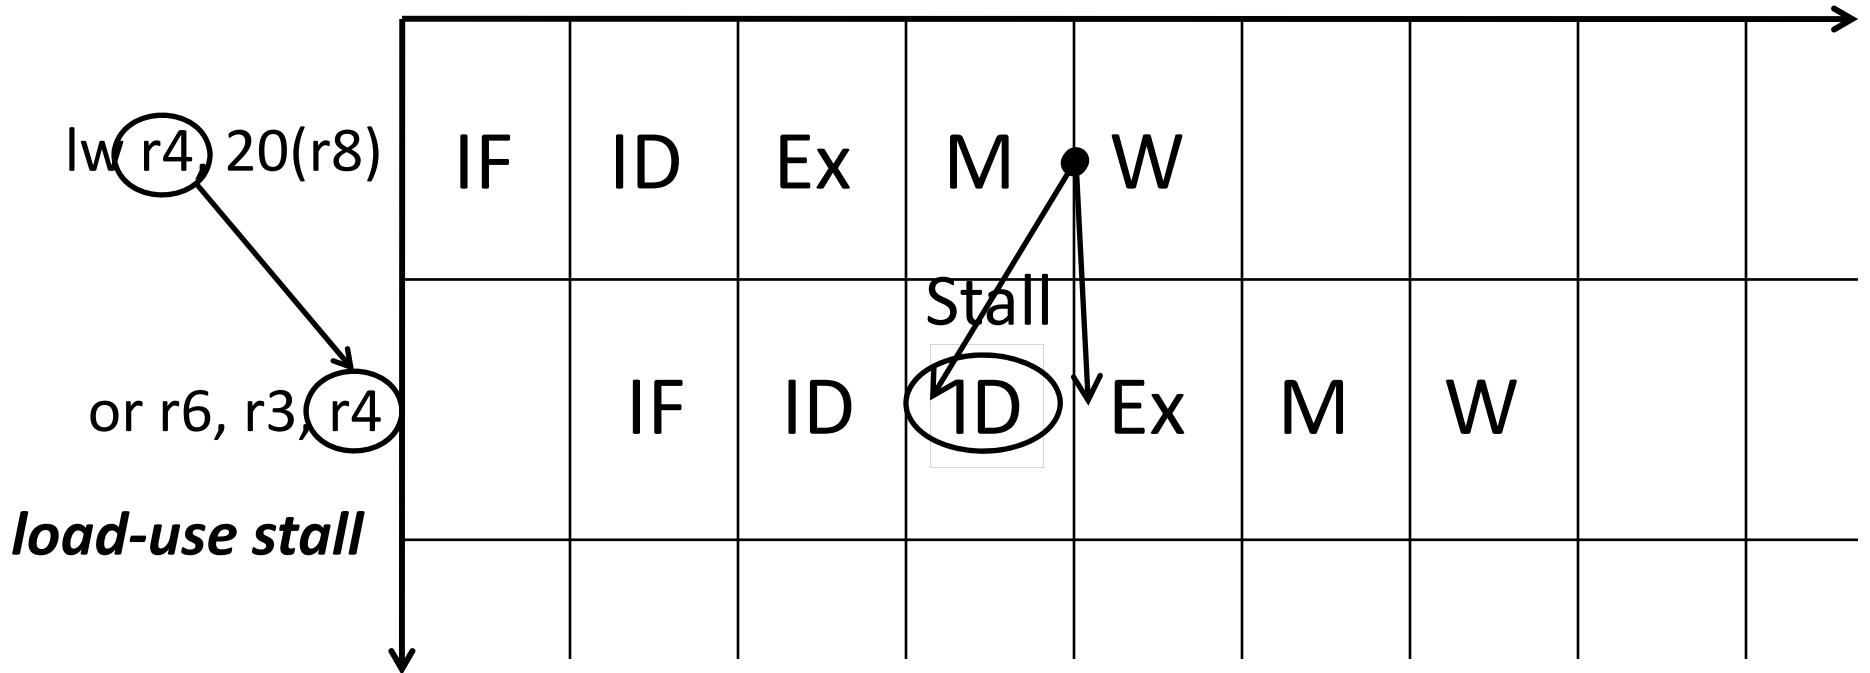
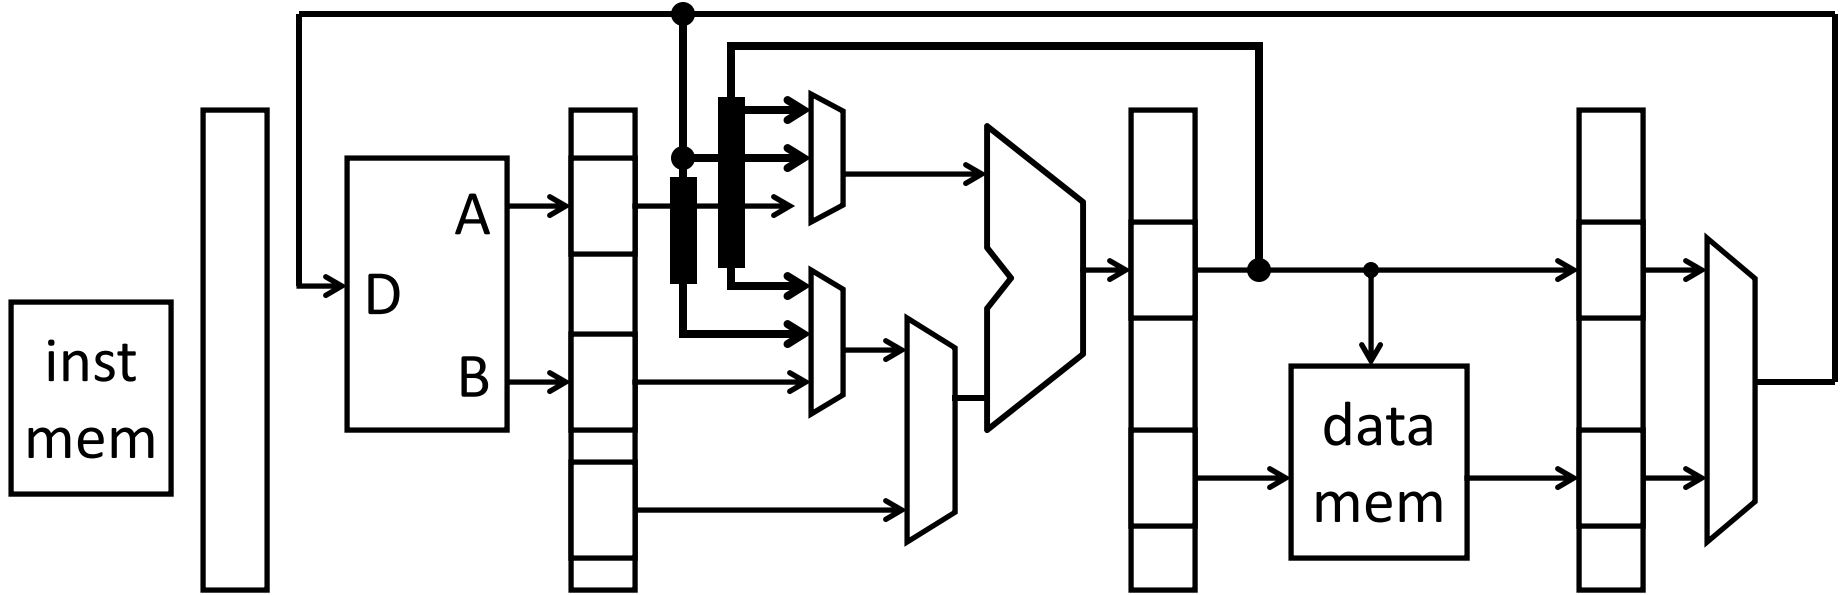
Memory Load Data Hazard

What happens if data dependency after a load word instruction?

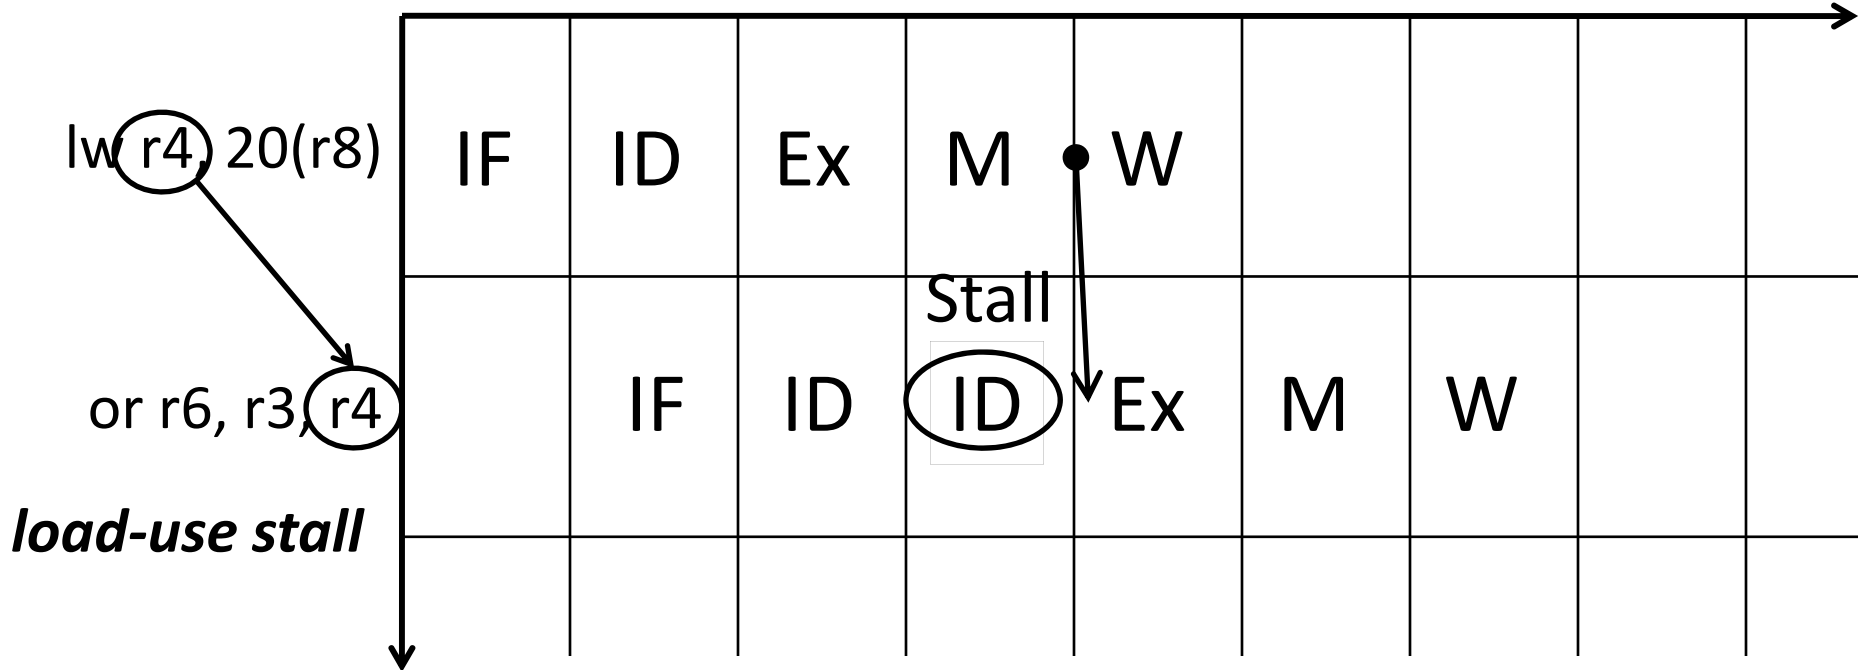
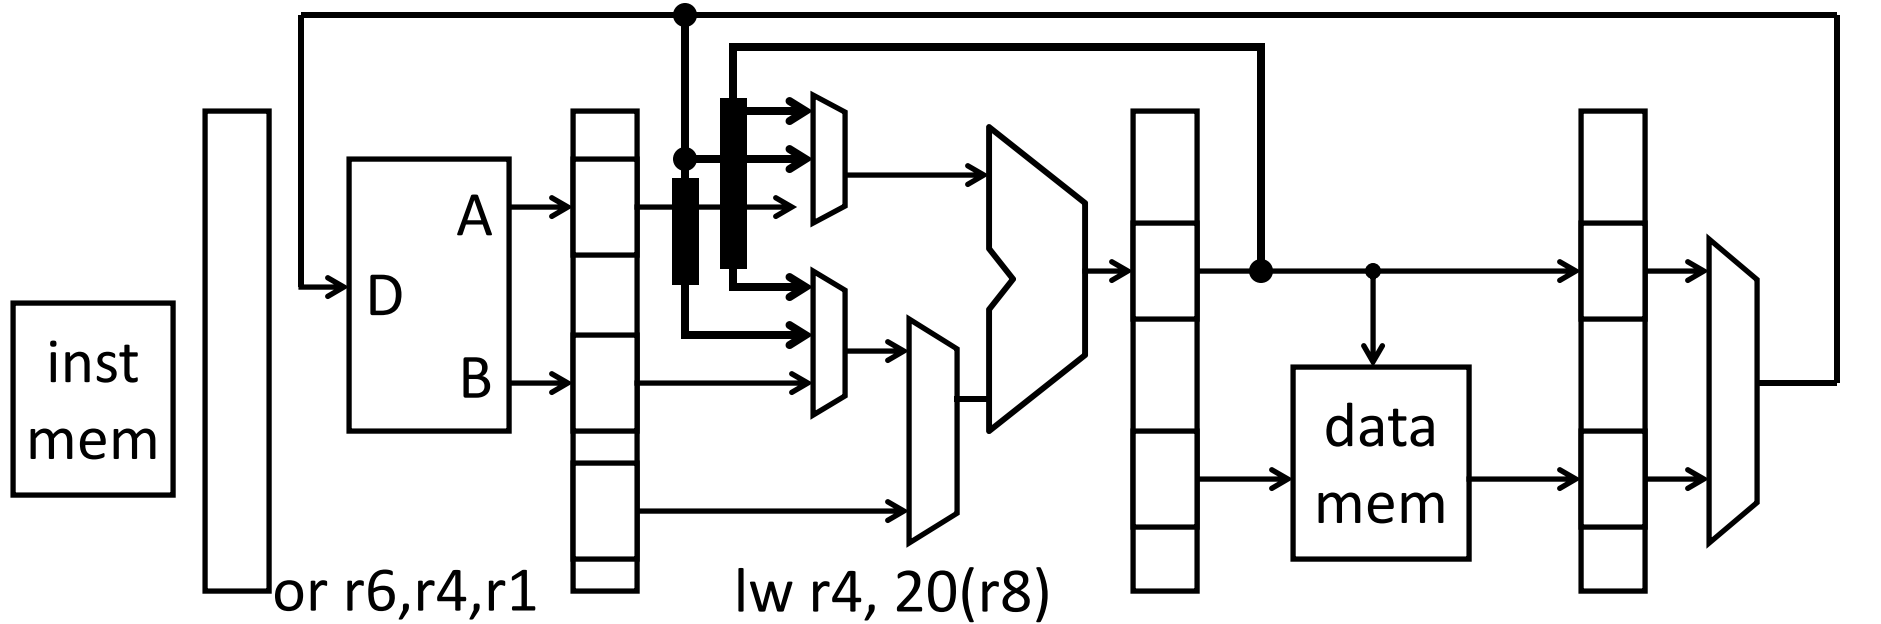
Memory Load Data Hazard

- Value not available until after the M stage
- So: next instruction can't proceed if hazard detected

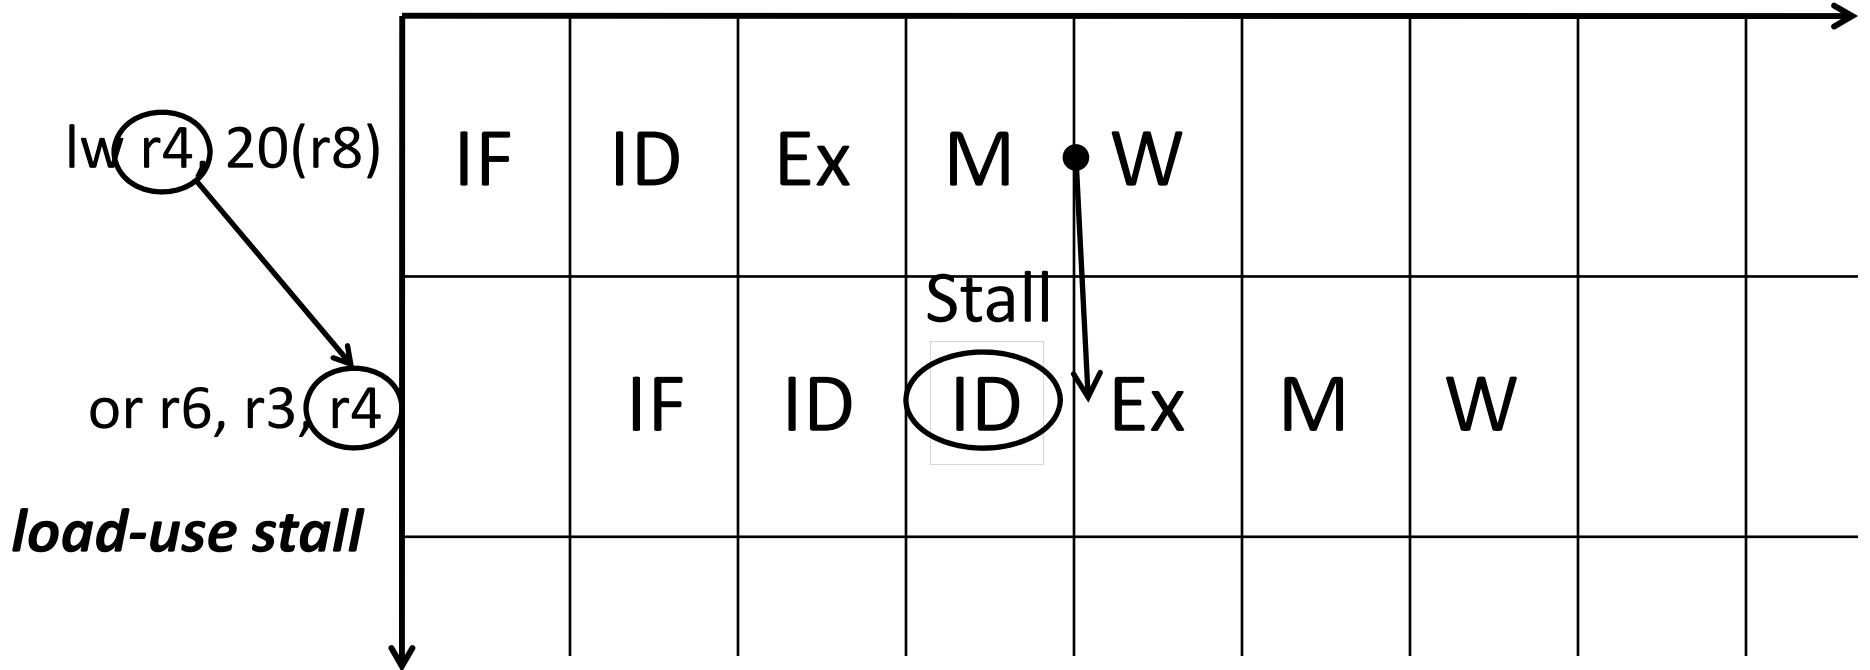
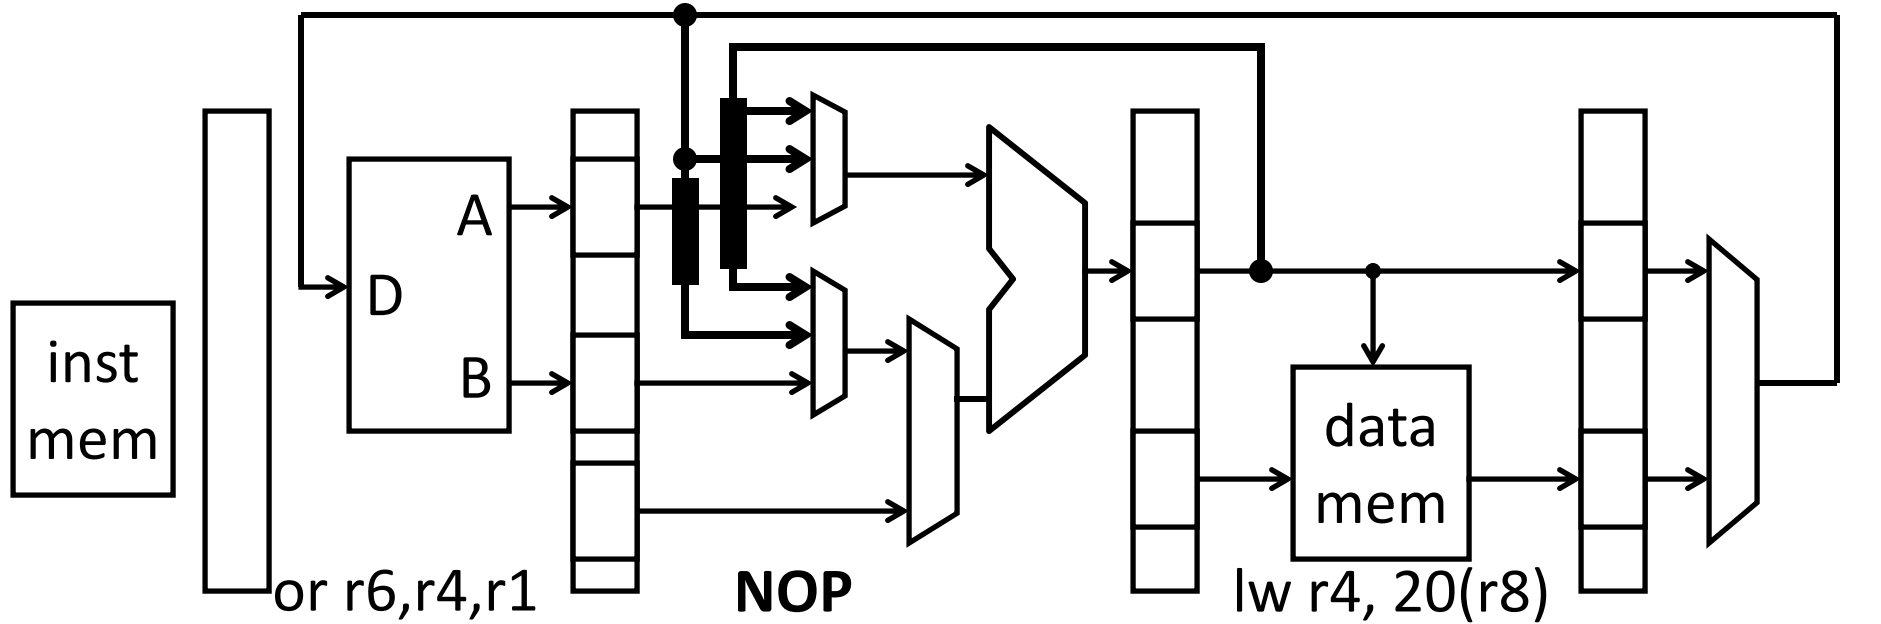
Memory Load Data Hazard



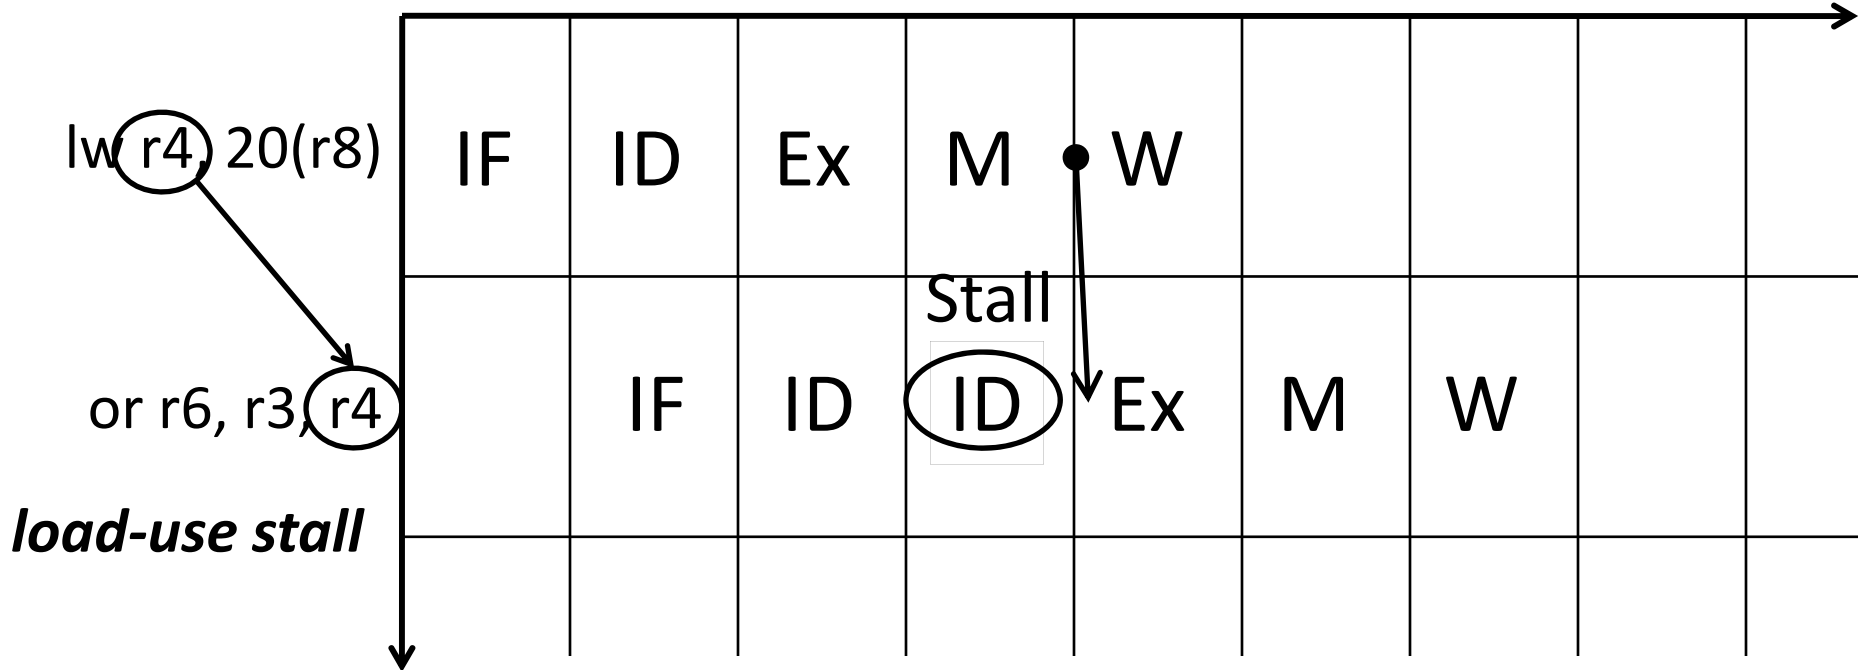
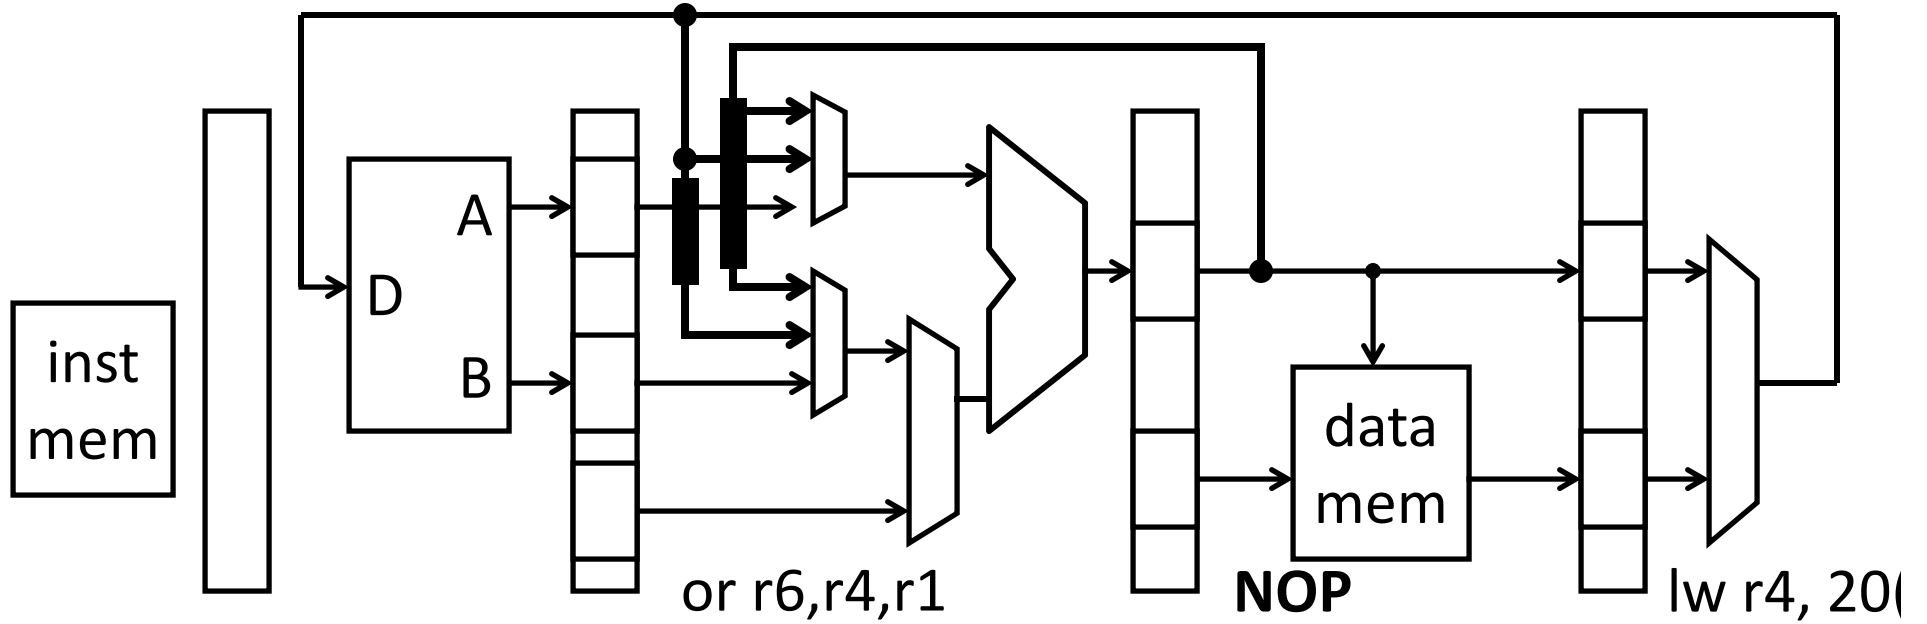
Memory Load Data Hazard



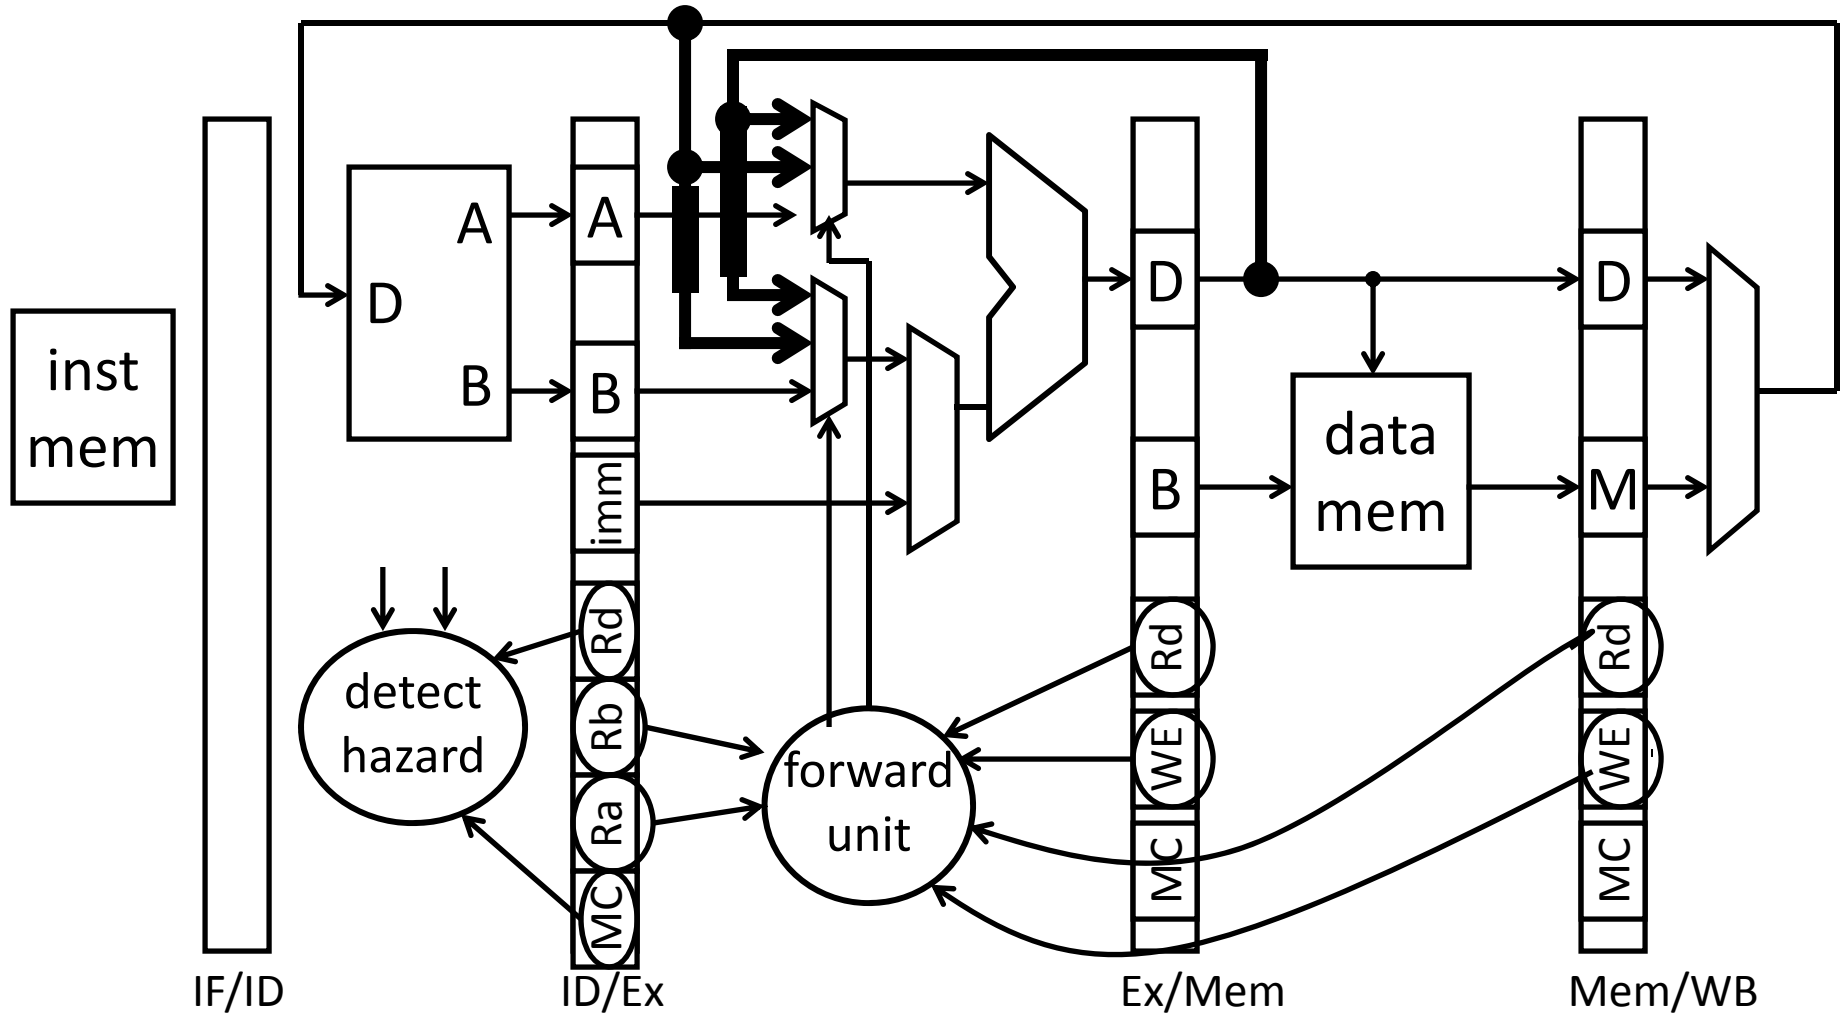
Memory Load Data Hazard



Memory Load Data Hazard



Memory Load Data Hazard



Stall =

If(ID/Ex.MemRead &&
IF/ID.Ra == ID/Ex.Rd

Memory Load Data Hazard

Load Data Hazard

- Value not available until WB stage
- So: next instruction can't proceed if hazard detected

Resolution:

- MIPS 2000/3000: one delay slot
 - ISA says results of loads are not available until one cycle later
 - Assembler inserts nop, or reorders to fill delay slot
- MIPS 4000 onwards: stall
 - But really, programmer/compiler reorders to avoid stalling in the load delay slot

For stall, how to detect? Logic in ID Stage

- Stall = ID/Ex.MemRead &&
(IF/ID.Ra == ID/Ex.Rd || IF/ID.Rb == ID/Ex.Rd)

Quiz

Find all hazards, and say how they are resolved:

add r3, r1, r2

nand r5, r3, r4

add r2, r6, r3

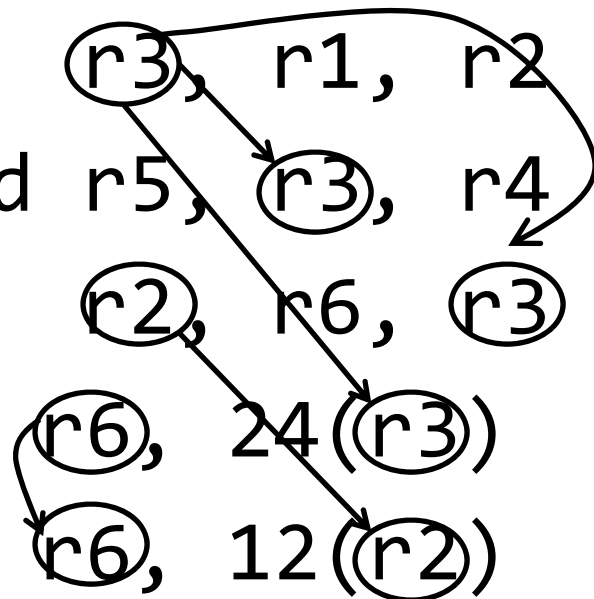
lw r6, 24(r3)

sw r6, 12(r2)

Quiz

Find all hazards, and say how they are resolved:

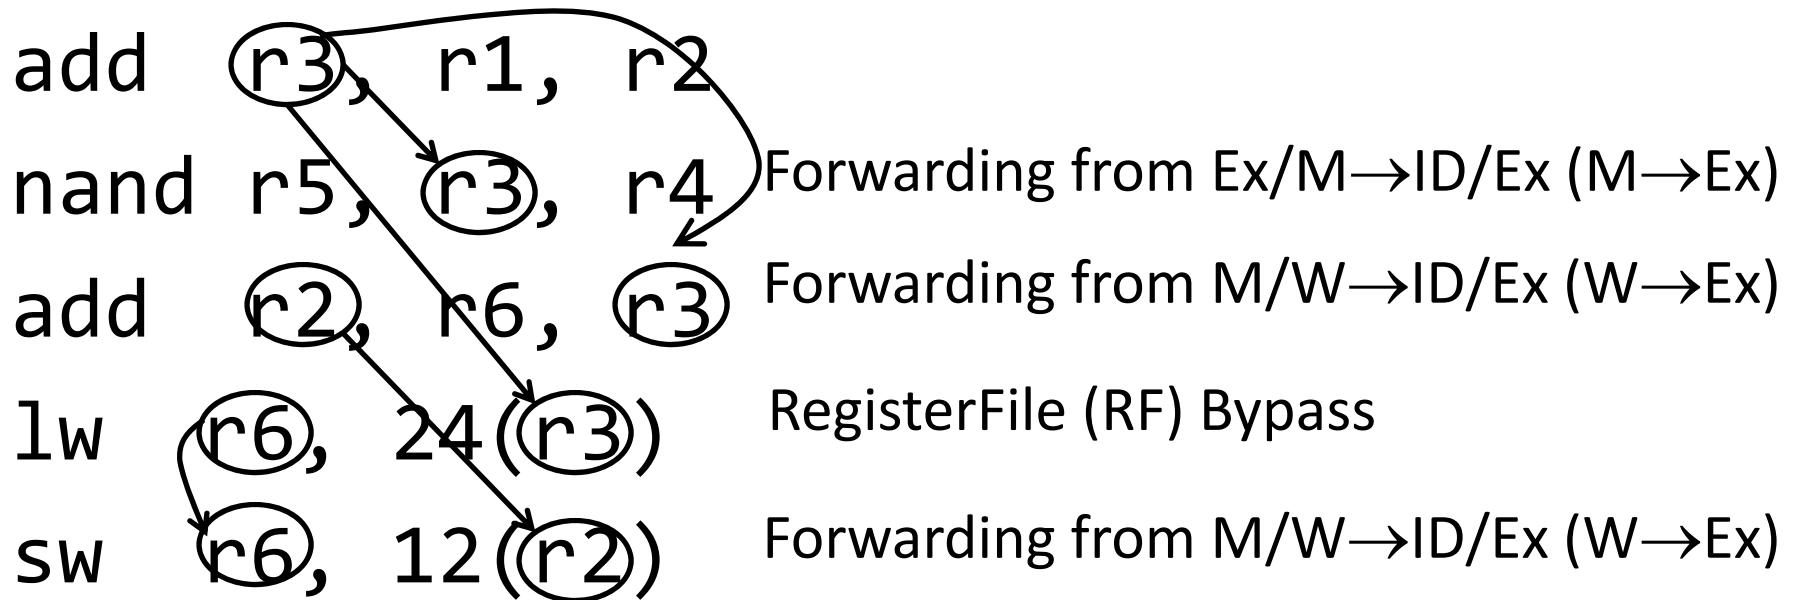
add r3, r1, r2
nand r5, r3, r4
add r2, r6, r3
lw r6, 24(r3)
sw r6, 12(r2)



5 Hazards

Quiz

Find all hazards, and say how they are resolved:



Stall

+ Forwarding from M/W→ID/Ex (W→Ex)

5 Hazards

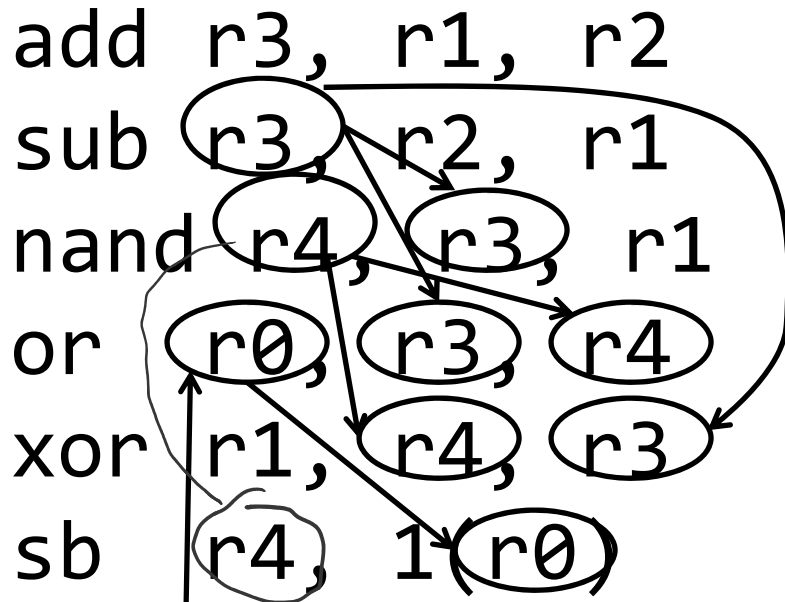
Quiz

Find all hazards, and say how they are resolved:

```
add r3, r1, r2
sub r3, r2, r1
nand r4, r3, r1
or  r0, r3, r4
xor r1, r4, r3
sb  r4, 1(r0)
```


Quiz 2

Find all hazards, and say how they are resolved:



Hours and hours of debugging!

Data Hazard Recap

Delay Slot(s)

- Modify ISA to match implementation

Stall

- Pause current and all subsequent instructions

Forward/Bypass

- Try to steal correct value from elsewhere in pipeline
- Otherwise, fall back to stalling or require a delay slot

Why are we learning about this?

Logic and gates

Numbers & arithmetic

States & FSMs

Memory

A simple CPU

Performance

Pipelining

Hazards: Data and Control

Why are we learning about this?

Logic and gates

Numbers & arithmetic

States & FSMs

Memory

A simple CPU

Performance

Pipelining

Hazards: Data and Control

Computer Organization is fundamental to CS

Control Hazards

What about branches?

A control hazard occurs if there is a control instruction (e.g. BEQ) and the program counter (PC) following the control instruction is not known until the control instruction computes if the branch should be taken

e.g.

0x10: beq r1, r2, L

0x14: add r3, r0, r3

0x18: sub r5, r4, r6

0x1C: L: or r3, r2, r4

Control Hazards

Control Hazards

- instructions are fetched in stage 1 (IF)
- branch and jump decisions occur in stage 3 (EX)
- i.e. next PC is not known until **2 cycles after** branch/jump

What happens to instr following a branch, if branch **not** taken?

A) Stall

B) Forward/Bypass

C) Zap/Flush

D) All the above

E) None of the above

e.g.

0x10: beq r1, r2, L

0x14: add r3, r0, r3

0x18: sub r5, r4, r6

0x1C: L: or r3, r2, r4

Control Hazards

Control Hazards

- instructions are fetched in stage 1 (IF)
- branch and jump decisions occur in stage 3 (EX)
- i.e. next PC is not known until **2 cycles after** branch/jump

What happens to instr following a branch, if branch **taken**?

A) Stall

B) Forward/Bypass

C) Zap/Flush

D) All the above

E) None of the above

e.g.

0x10: beq r1, r2, L

0x14: add r3, r0, r3

0x18: sub r5, r4, r6

0x1C: L: or r3, r2, r4

Control Hazards

Control Hazards

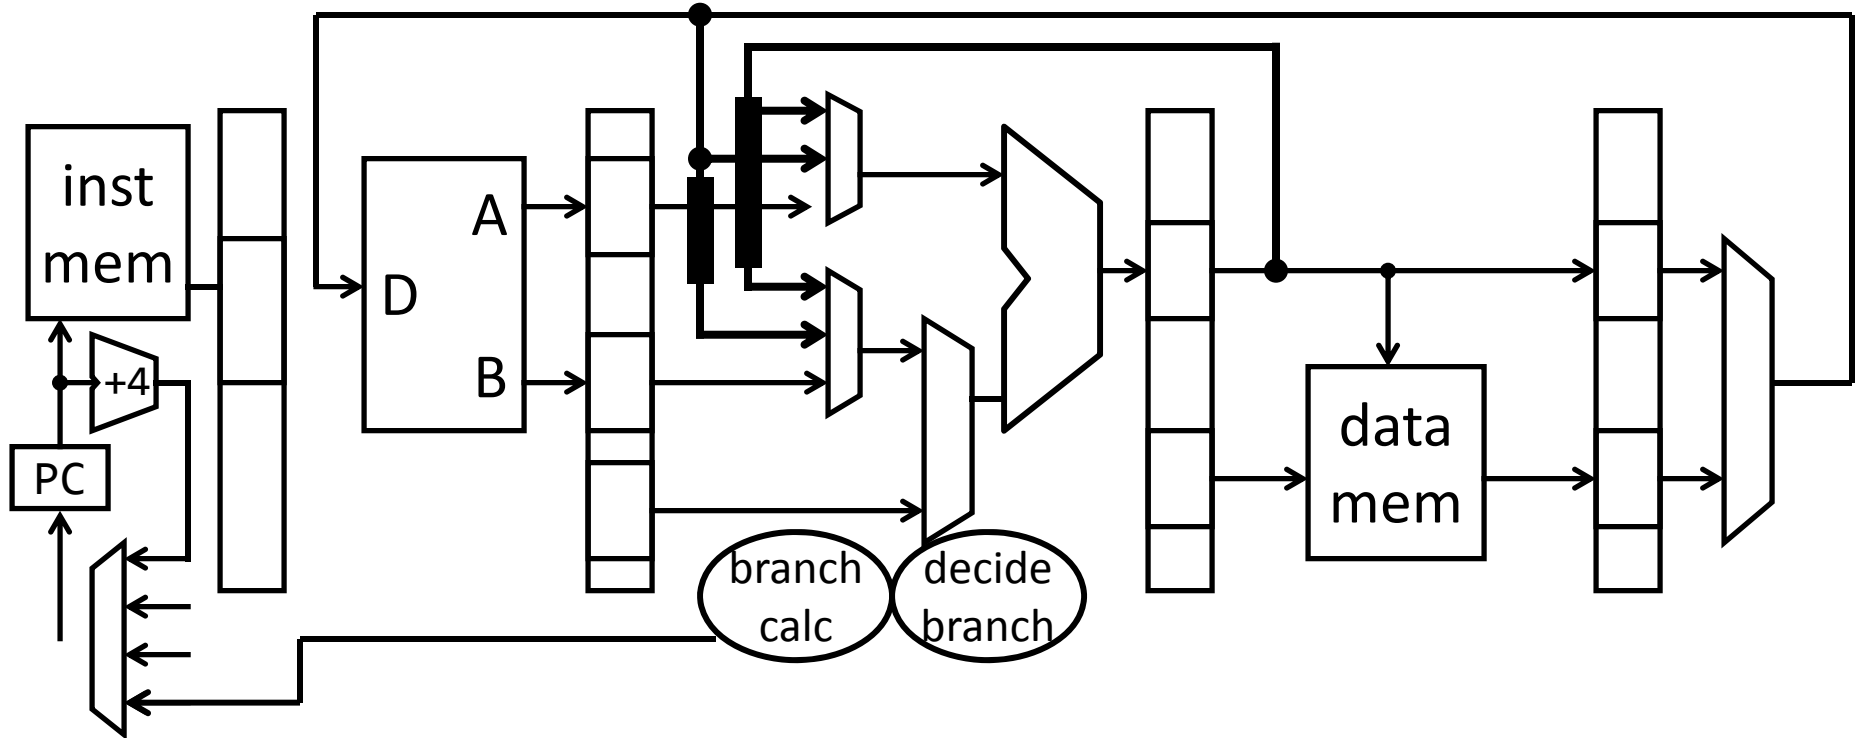
- instructions are fetched in stage 1 (IF)
- branch and jump decisions occur in stage 3 (EX)
- i.e. next PC is not known until **2 cycles after** branch/jump

What happens to instr following a branch, if branch **taken**?

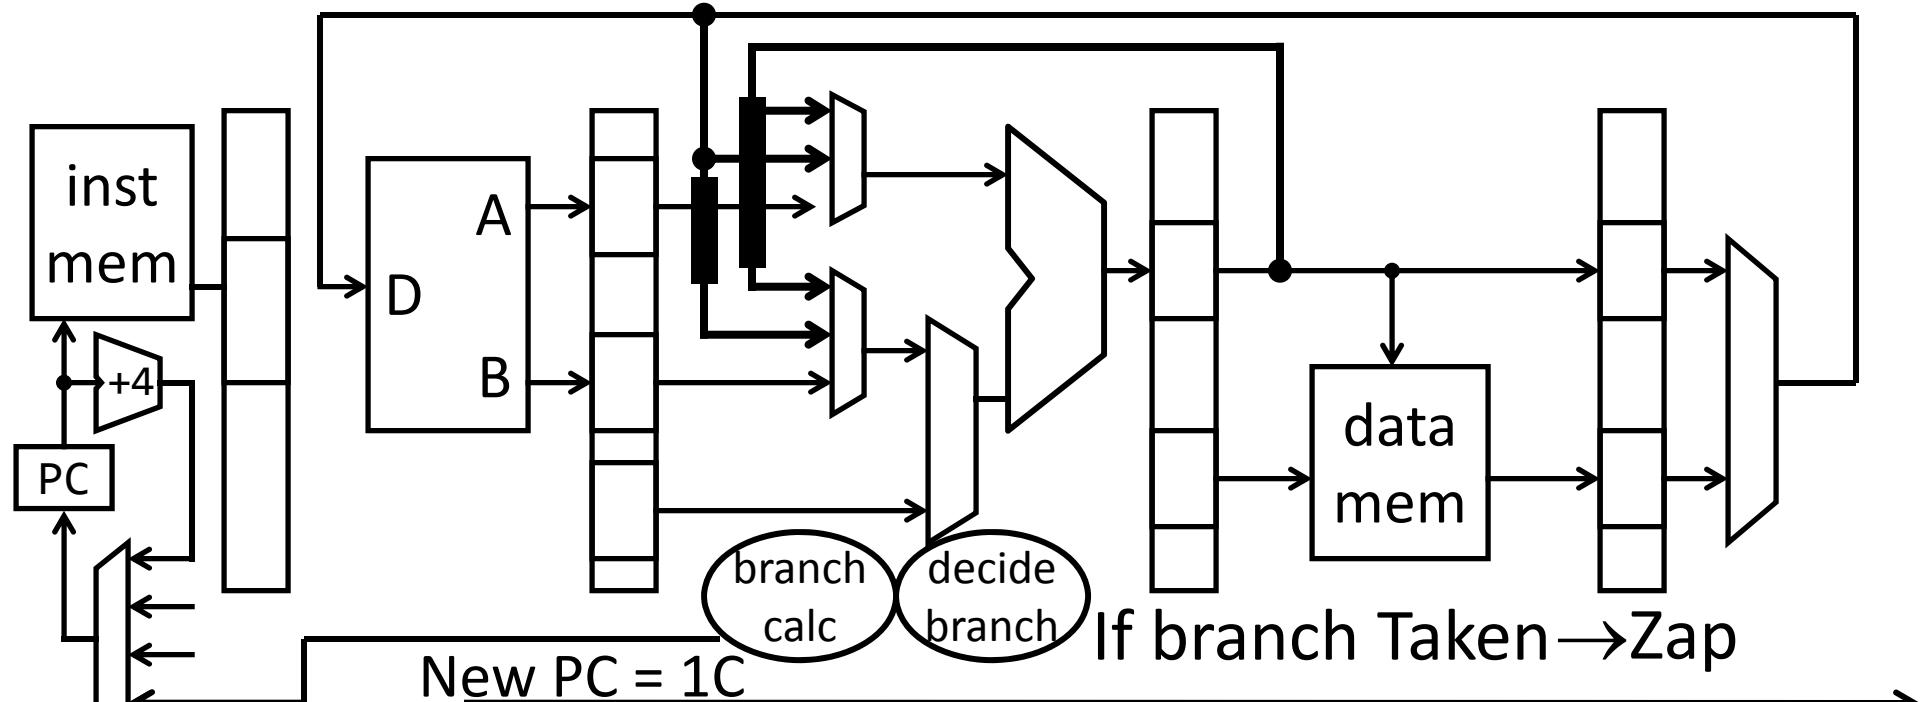
Stall (+ Zap/Flush)

- prevent PC update
- clear IF/ID pipeline register
 - instruction just fetched might be wrong one, so convert to nop
- allow branch to continue into EX stage

Control Hazards

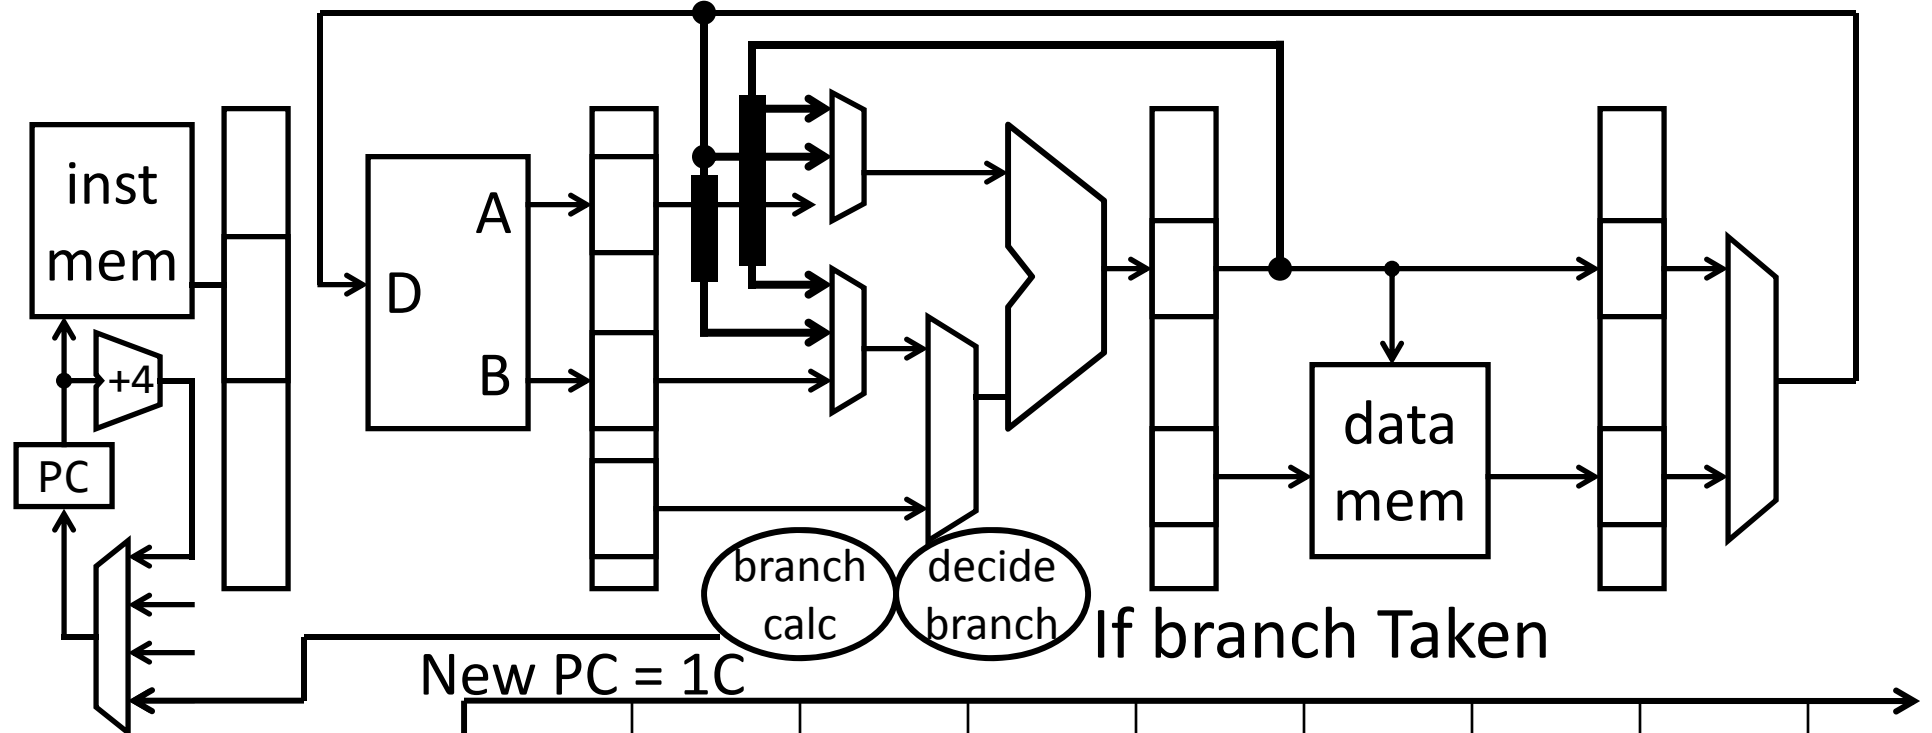


Control Hazards



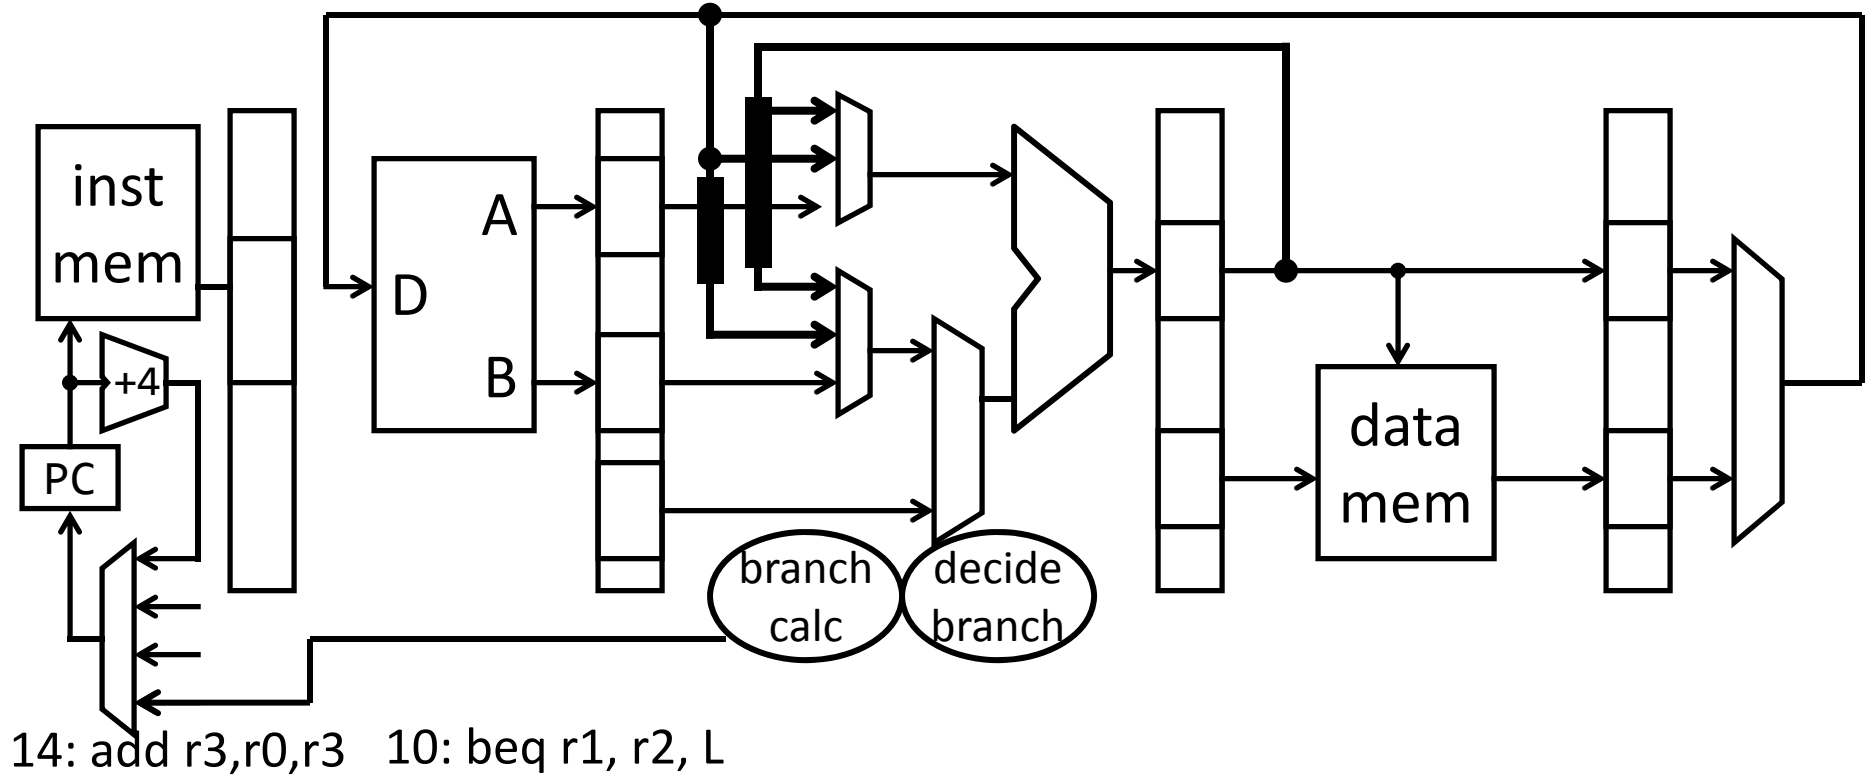
10: beq r1, r2, (L)	IF	ID	Ex	M	W				
14: add r3, r0, r3		IF	ID	NOP	NOP	NOP			
18: sub r5, r4, r6			IF	NOP	NOP	NOP	NOP		
1C: (L) or r3, r2, r4				IF	ID	Ex	M	W	

Control Hazards

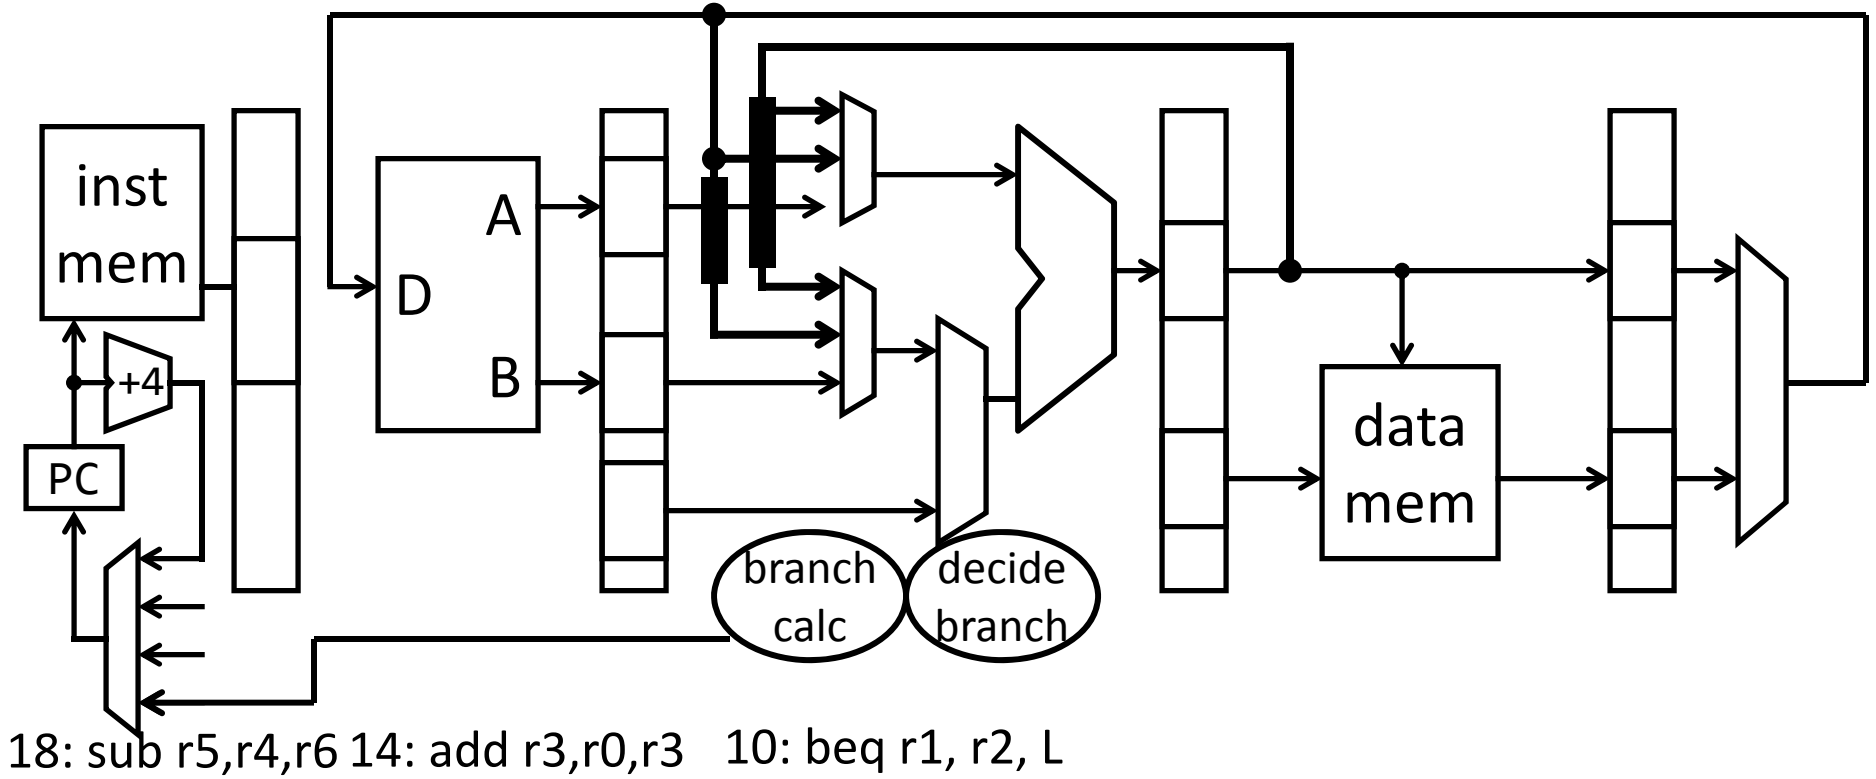


10:	beq r1, r2, L	IF	ID	Ex	M	W			
14	add r3, r0, r3		IF	ID	NOP	NOP	NOP		
18	sub r5, r4, r6			IF	NOP	NOP	NOP	NOP	
1C:	L: or r3, r2, r4				IF	ID	Ex	M	W

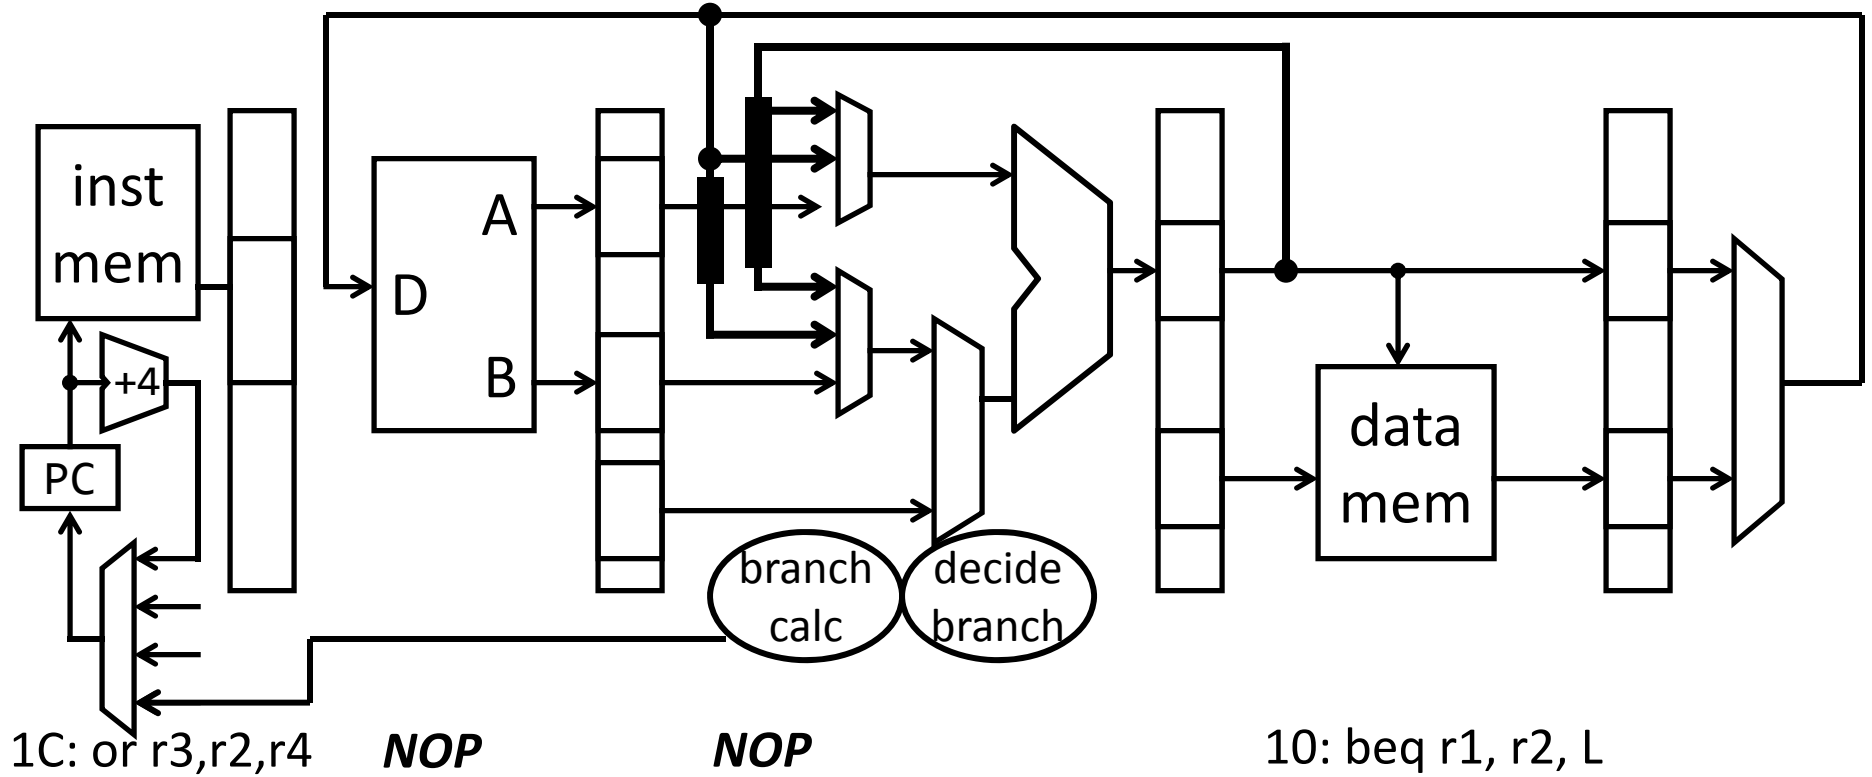
Control Hazards



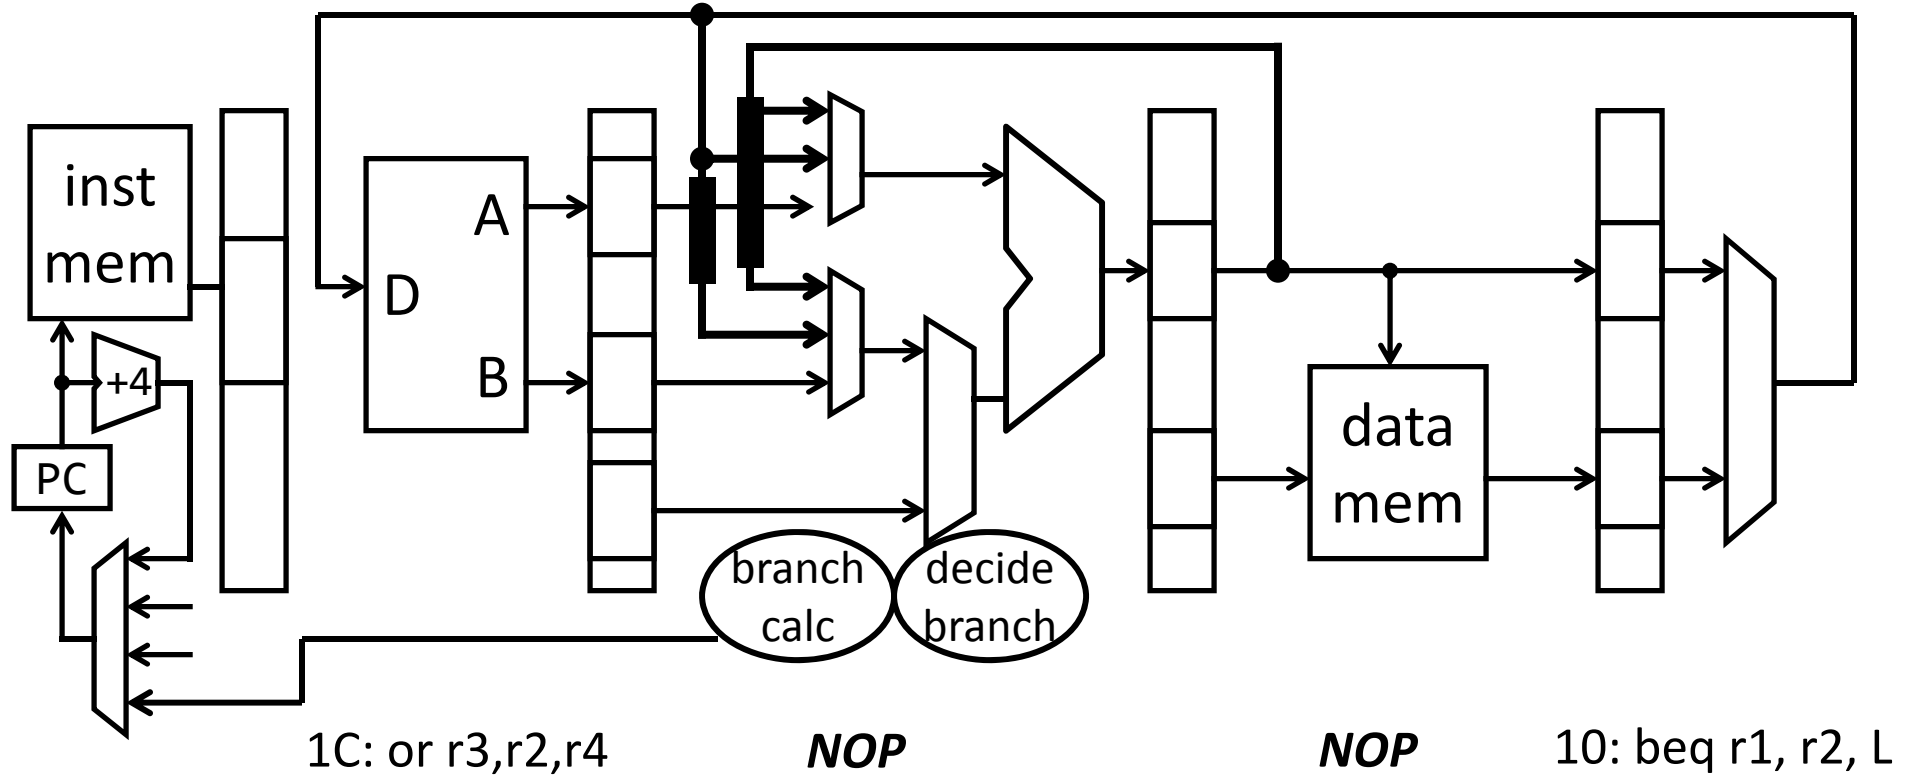
Control Hazards



Control Hazards



Control Hazards



Takeaway

Control hazards occur because the PC following a control instruction is not known until control instruction computes if branch should be taken or not.

If branch taken, then need to zap/flush instructions.

There is a performance penalty for branches:

Need to stall, then may need to zap (flush) subsequent instructions that have already been fetched.

Next Goal

Can we reduce the cost of a control hazard?

Reduce the cost of control hazard?

Can we forward/bypass values for branches?

- We can move branch calc from EX to ID
- will require new bypasses into ID stage; or can just zap the second instruction

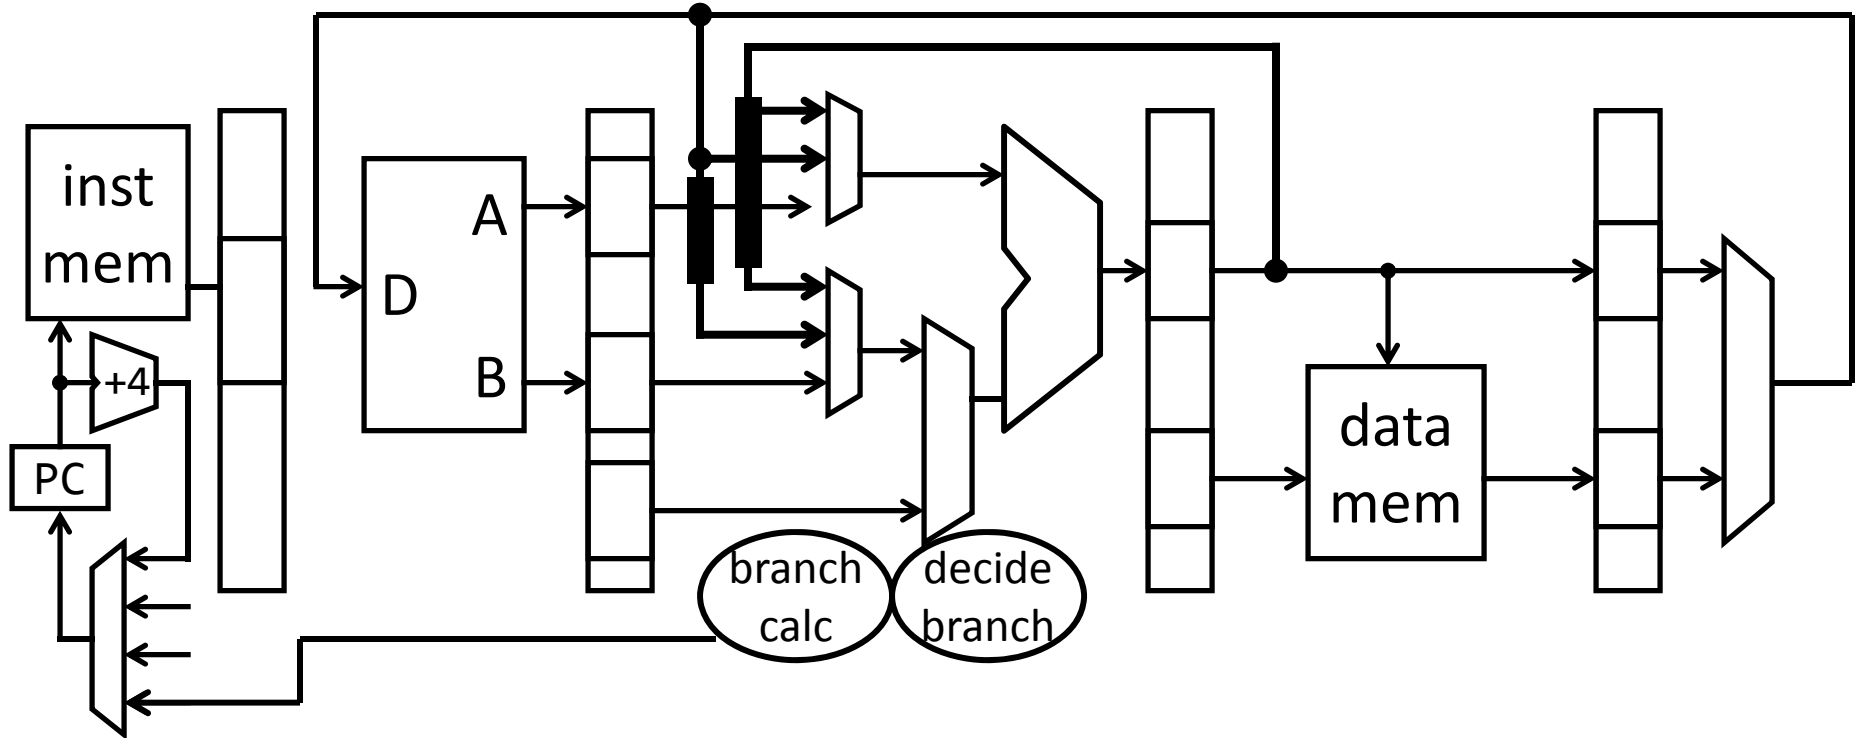
What happens to instructions following a branch, if branch taken?

- Still need to zap/flush instructions

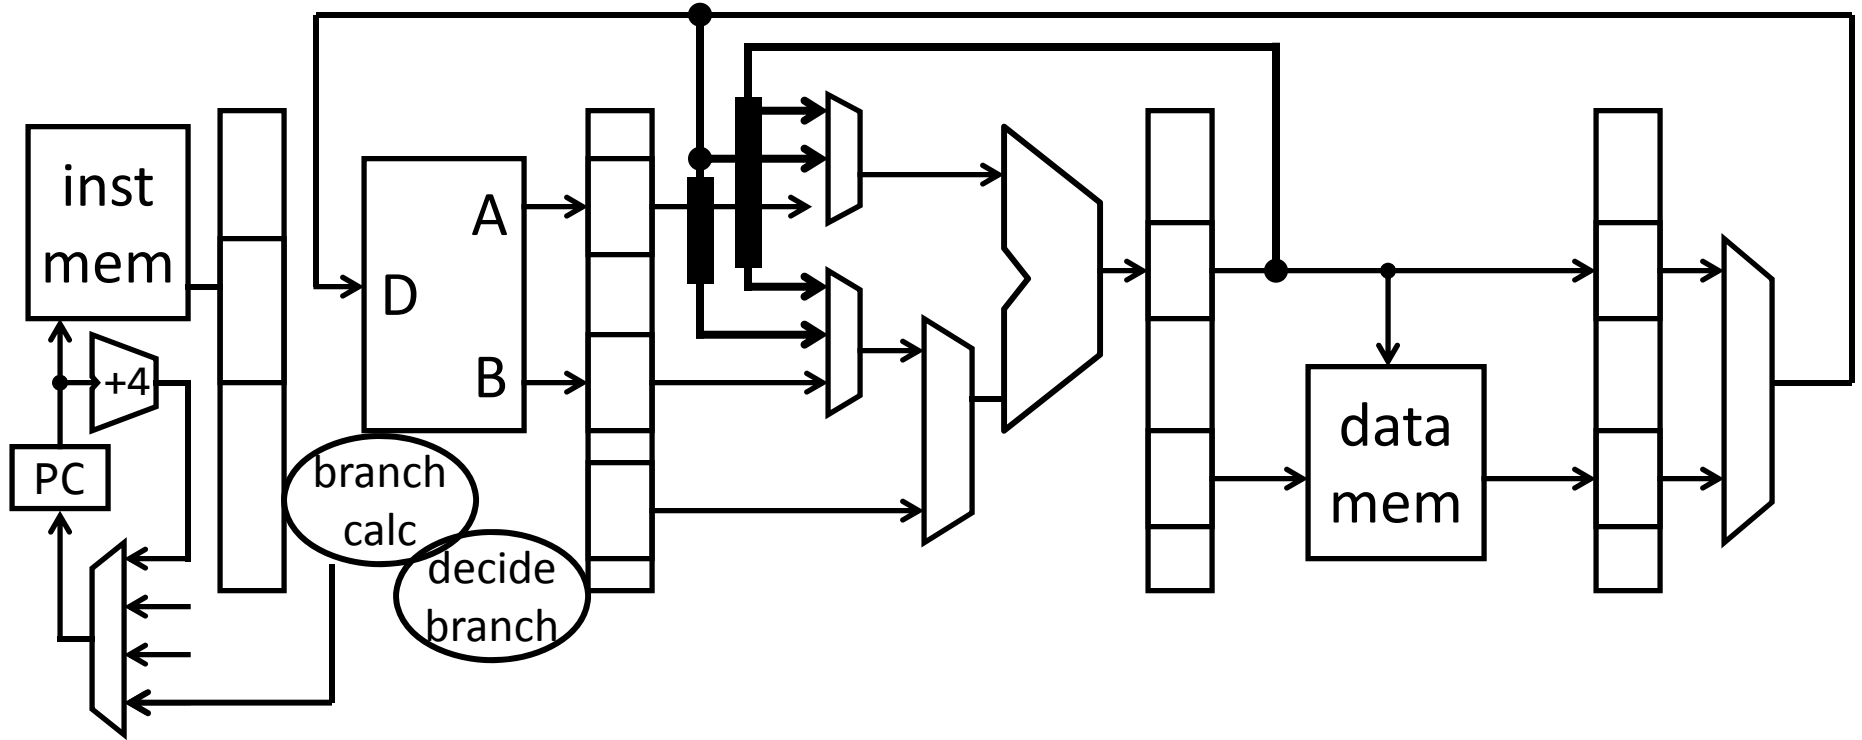
Is there still a performance penalty for branches

- Yes, need to stall, then may need to zap (flush) subsequent instructions that have already been fetched

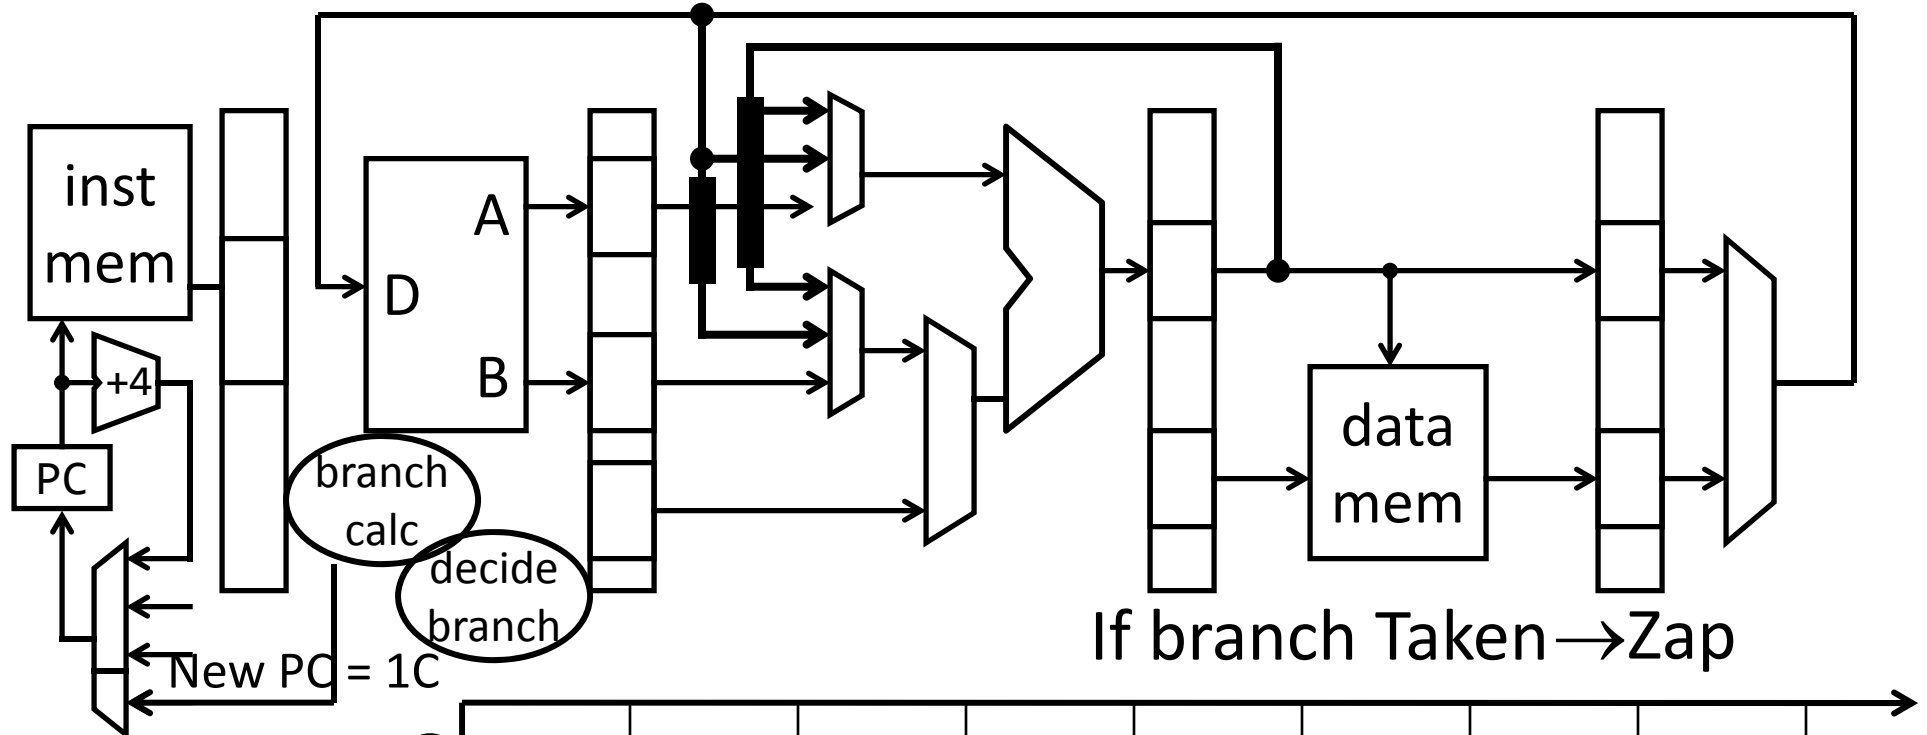
Control Hazards



Control Hazards



Control Hazards



10: beq r1, r2, (L)

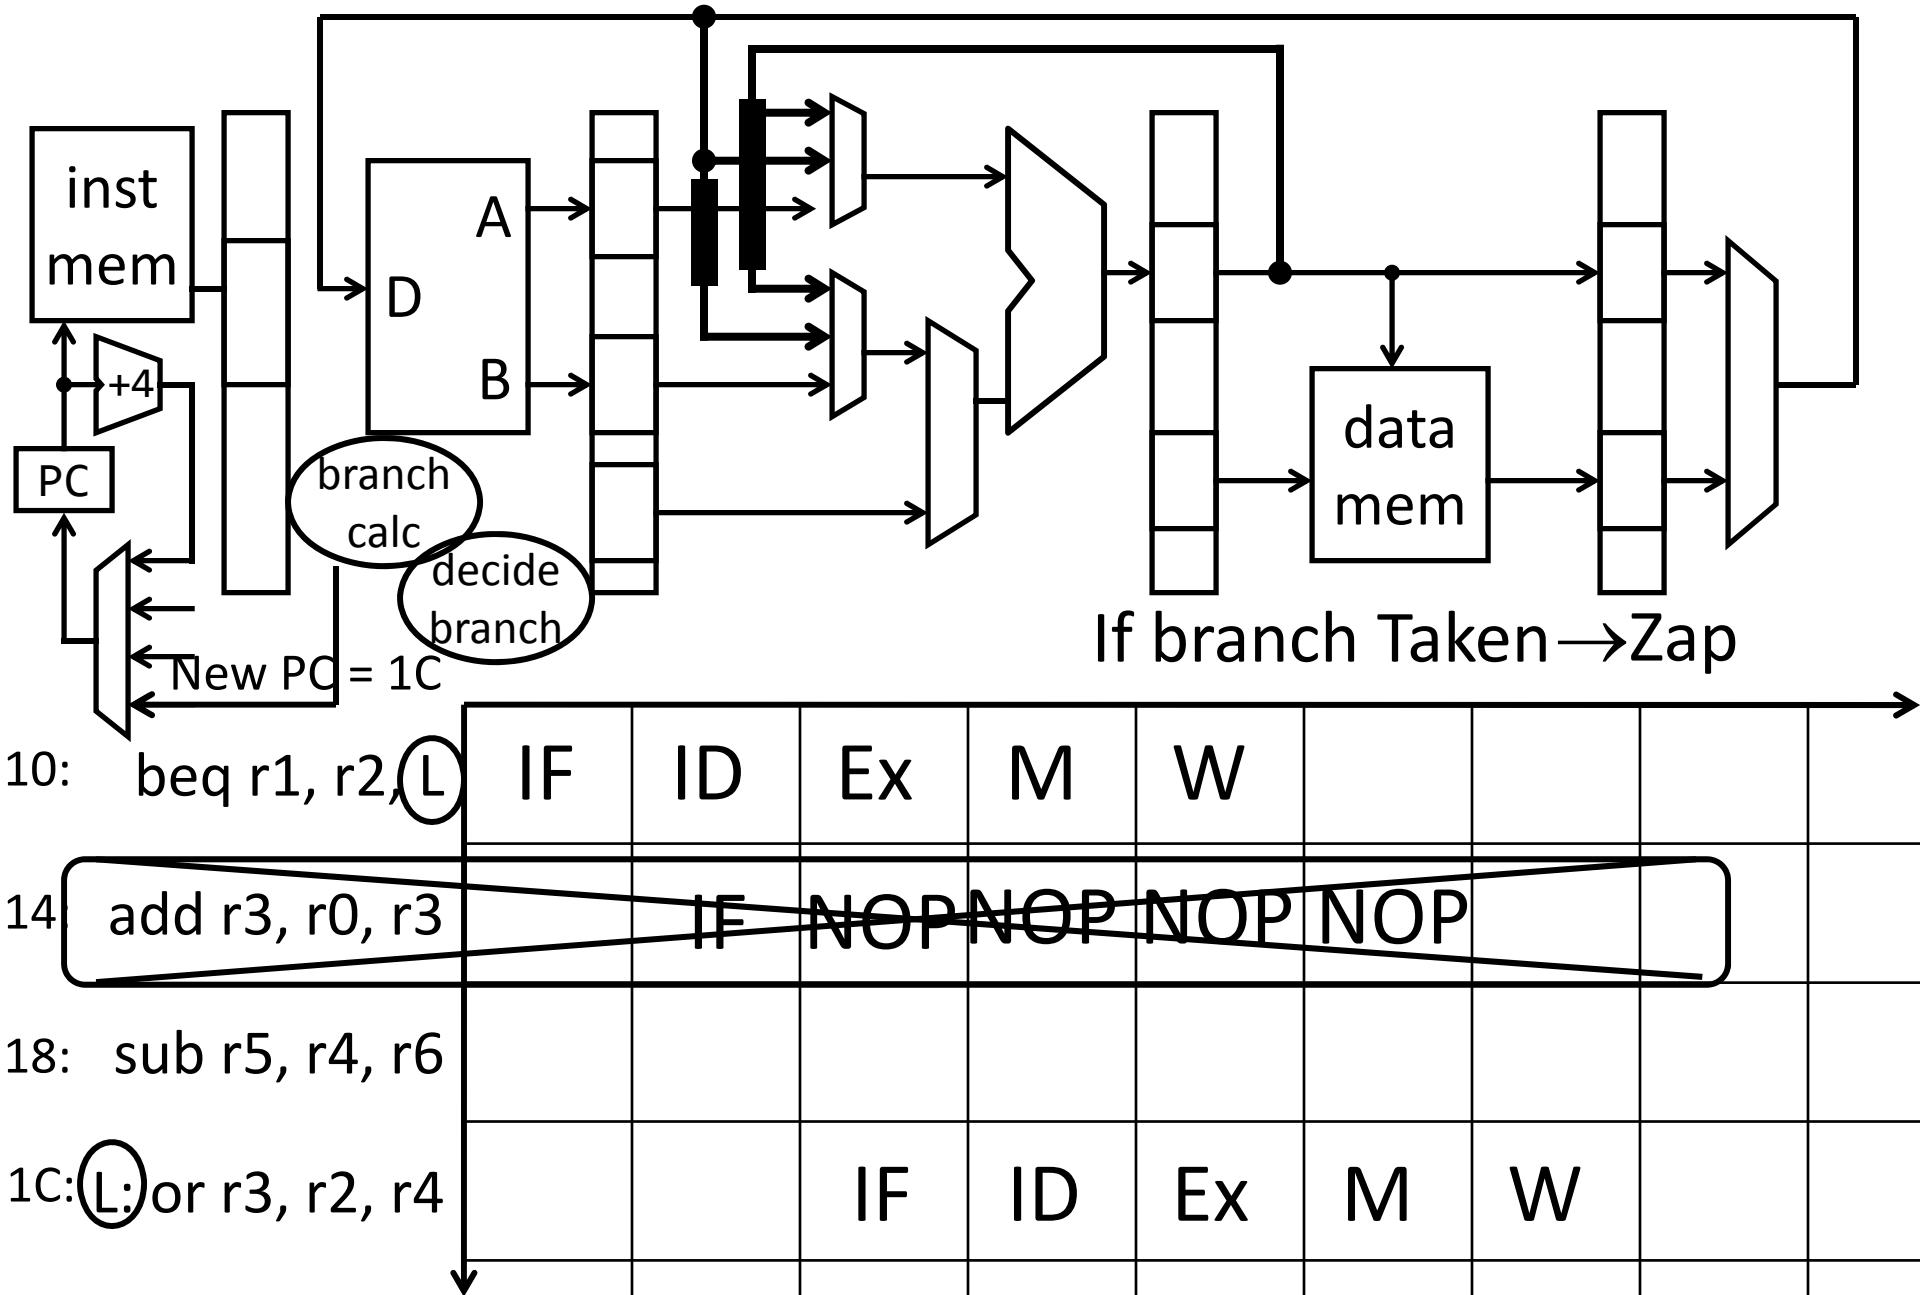
14: add r3, r0, r3

18: sub r5, r4, r6

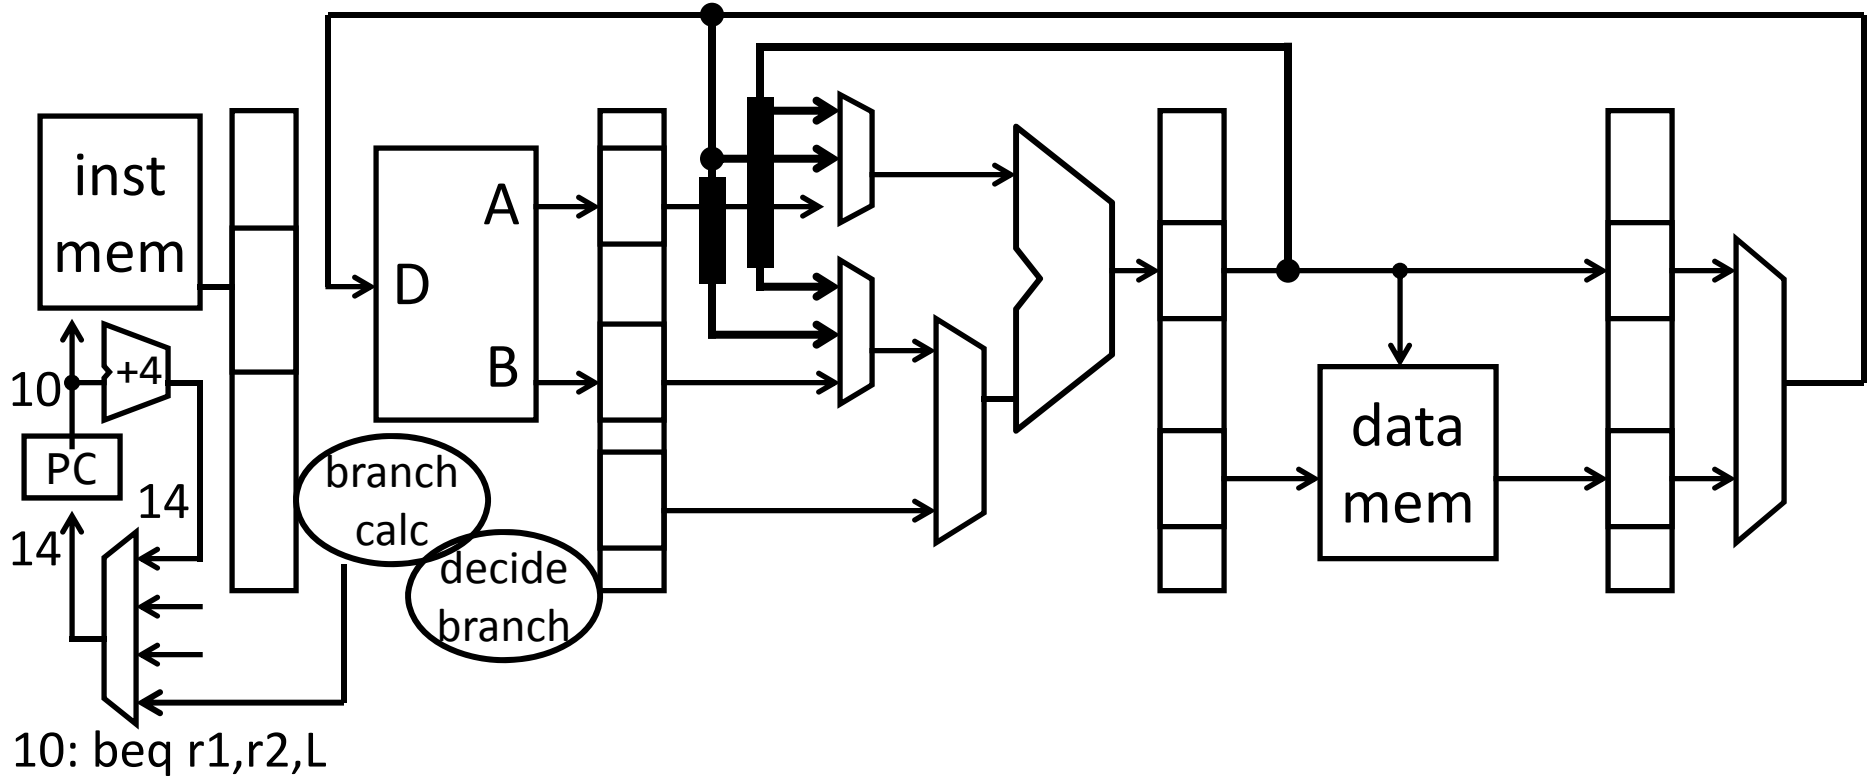
1C: (L) or r3, r2, r4

IF	ID	Ex	M	W				
	IF	NOP	NOP	NOP	NOP			
		IF	ID	Ex	M	W		

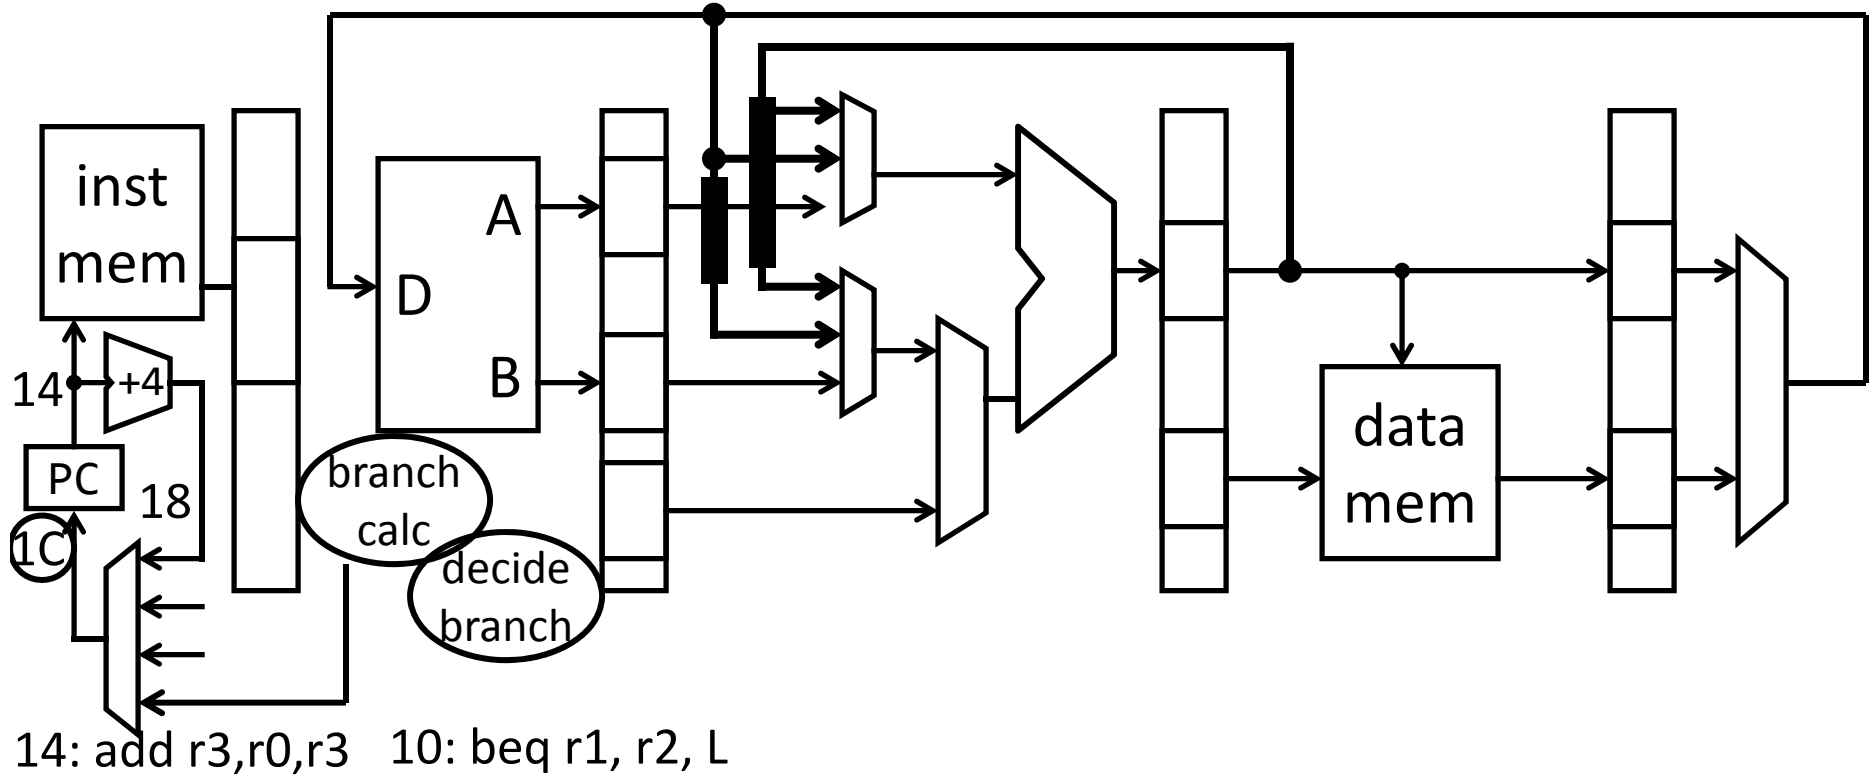
Control Hazards



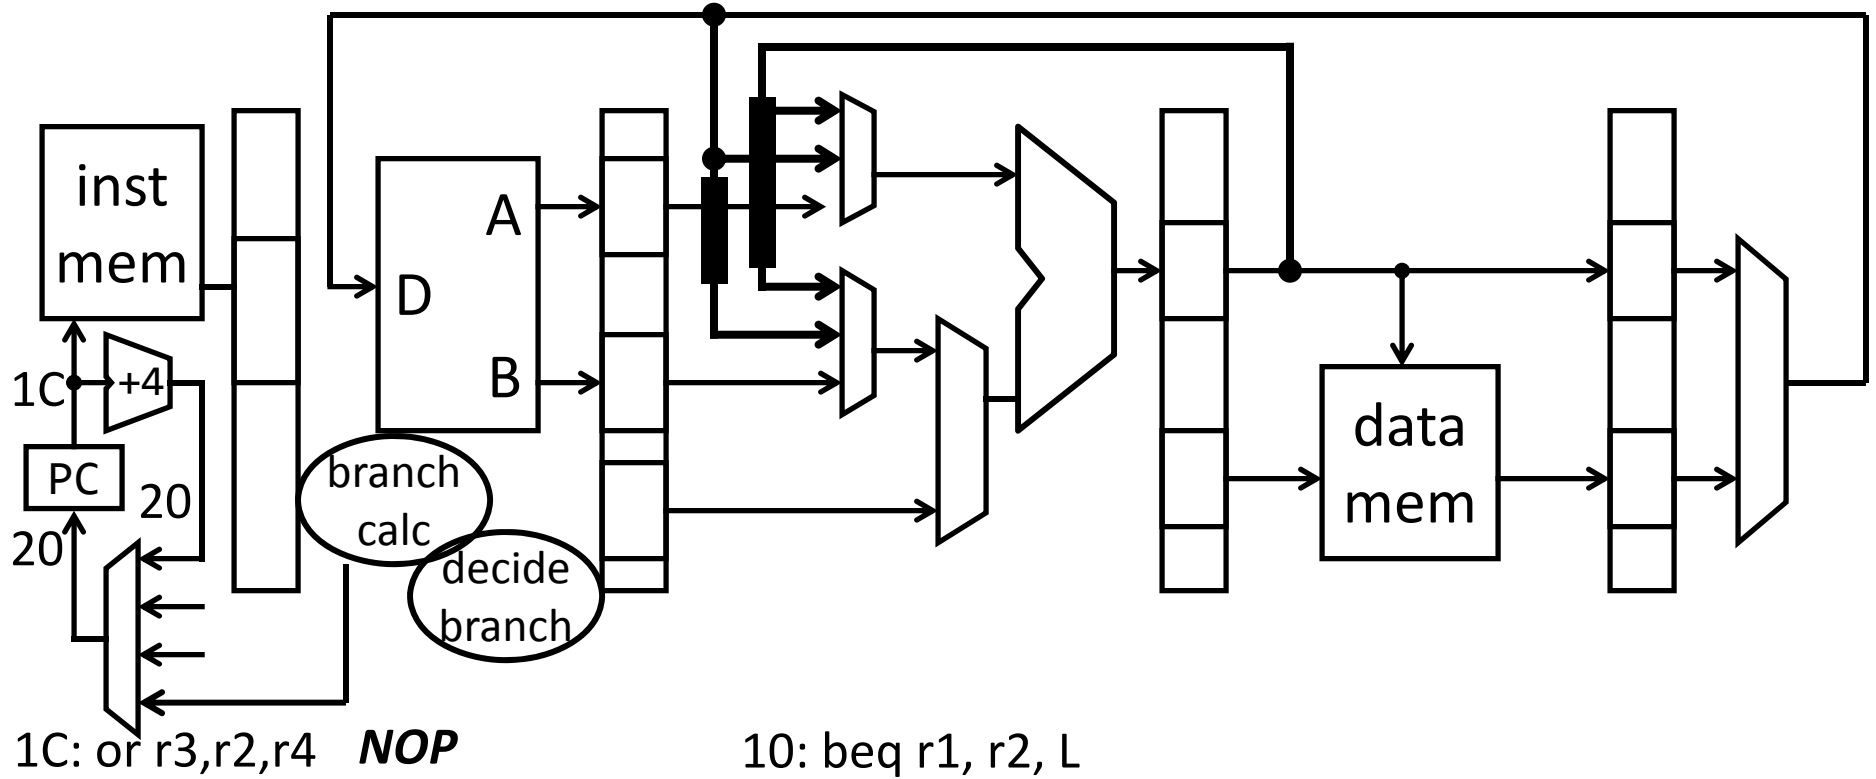
Control Hazards



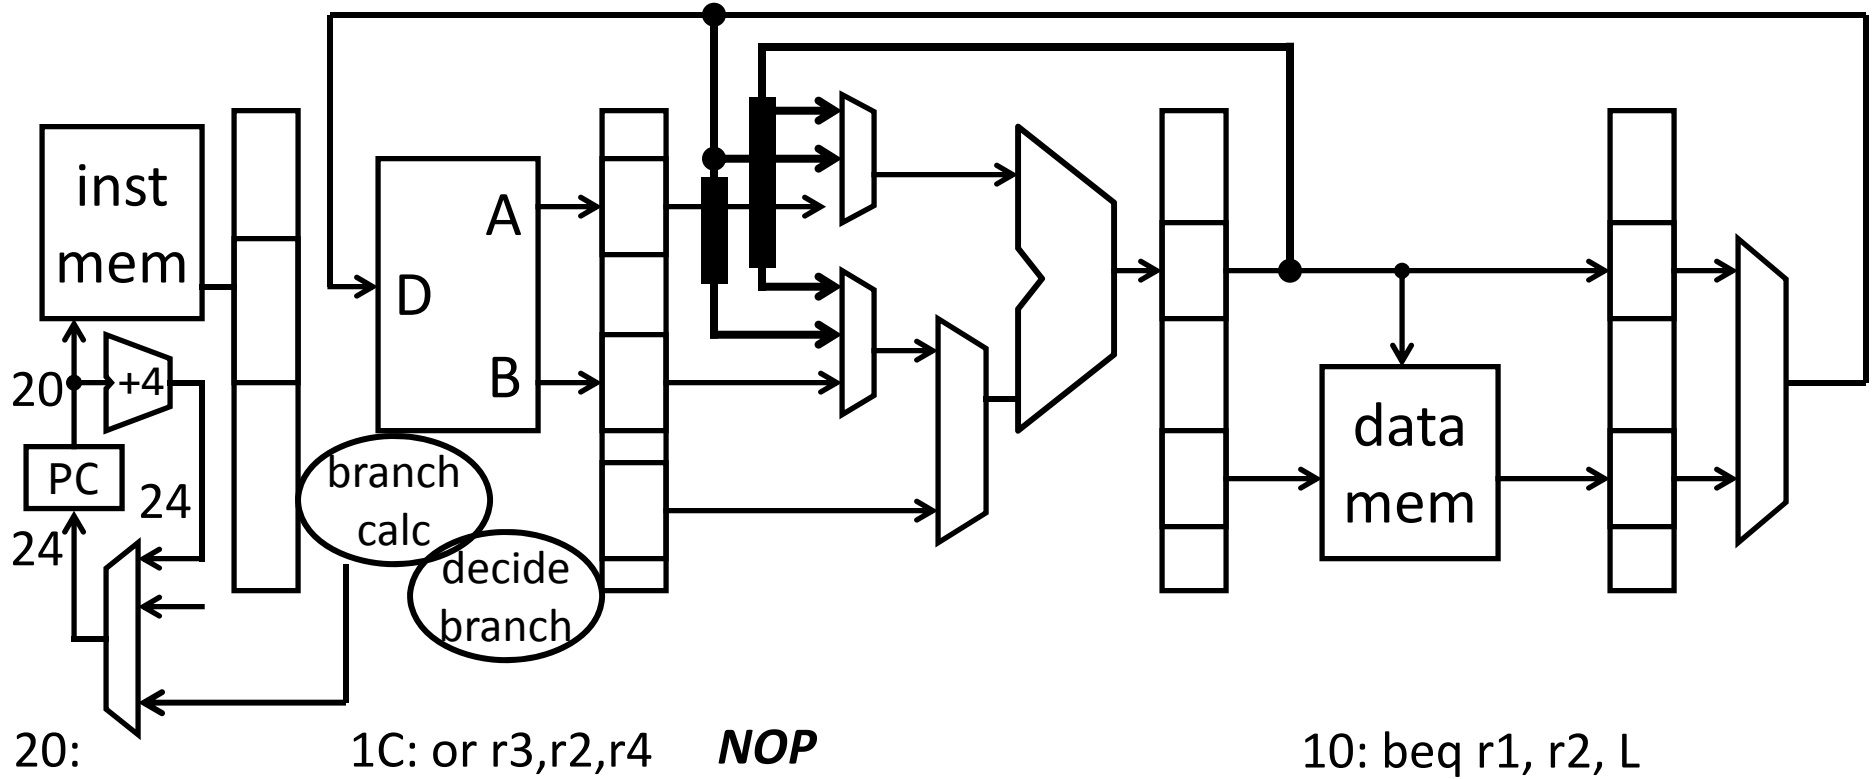
Control Hazards



Control Hazards



Control Hazards



Control Hazards

Control Hazards

- instructions are fetched in stage 1 (IF)
- branch and jump decisions occur in stage 3 (EX)
i.e. next PC is not known until **2 cycles after** branch/jump
- Can optimize and move branch and jump decision to stage 2 (ID)
i.e. next PC is not known until **1 cycles after** branch/jump

Stall (+ Zap)

- prevent PC update
- clear IF/ID pipeline register
 - instruction just fetched might be wrong one, so convert to nop
- allow branch to continue into EX stage

Takeaway

Control hazards occur because the PC following a control instruction is not known until control instruction computes if branch should be taken or not

If branch taken, then need to zap/flush instructions. There still a performance penalty for branches: Need to stall, then may need to zap (flush) subsequent instructions that have already been fetched

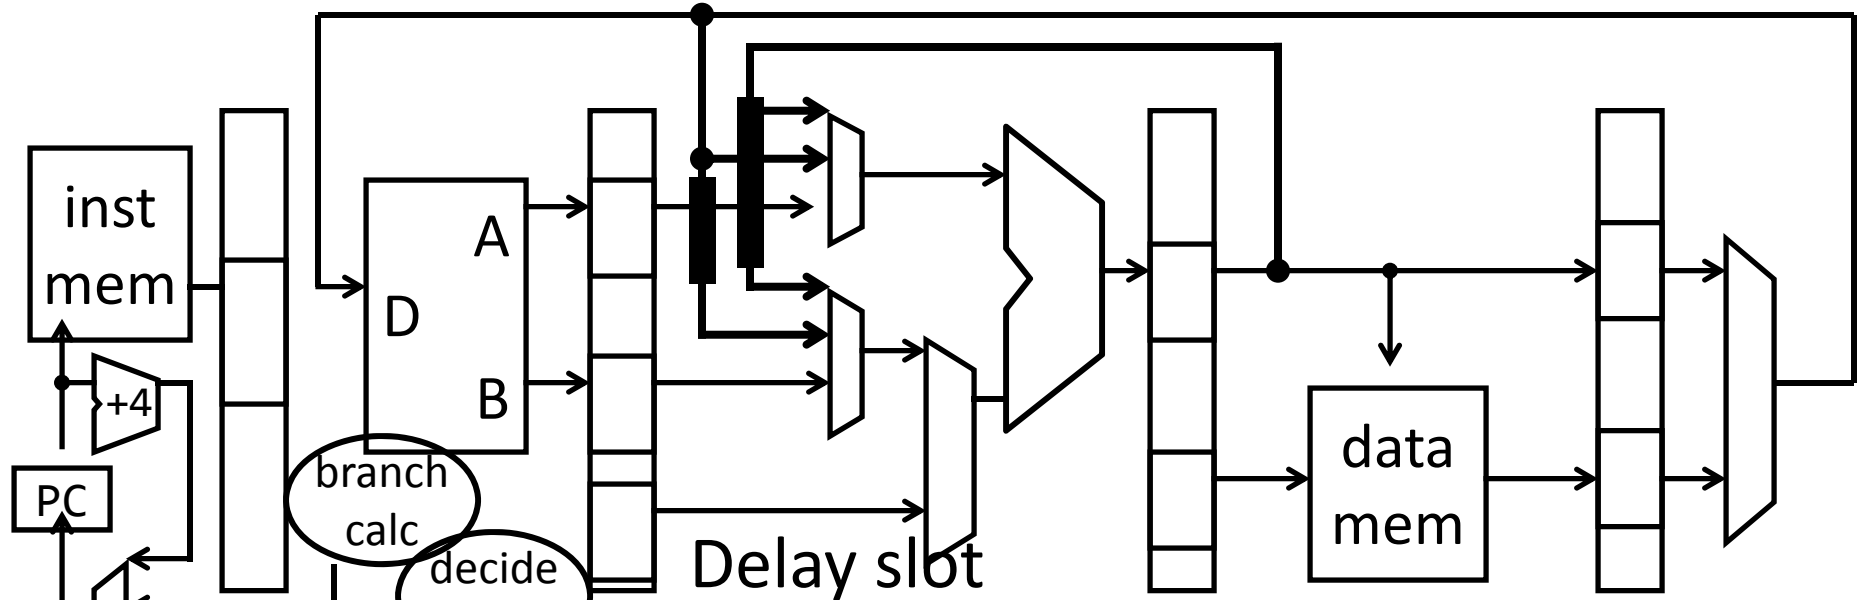
We can reduce cost of a control hazard by moving branch decision and calculation from Ex stage to ID stage. This reduces the cost from flushing two instructions to only flushing one.

Reduce cost of Control Hazards More

Delay Slot

- ISA says N instructions after branch/jump *always* executed
 - MIPS has 1 branch delay slot
 - i.e. Whether branch taken or not, instruction following branch is *always* executed

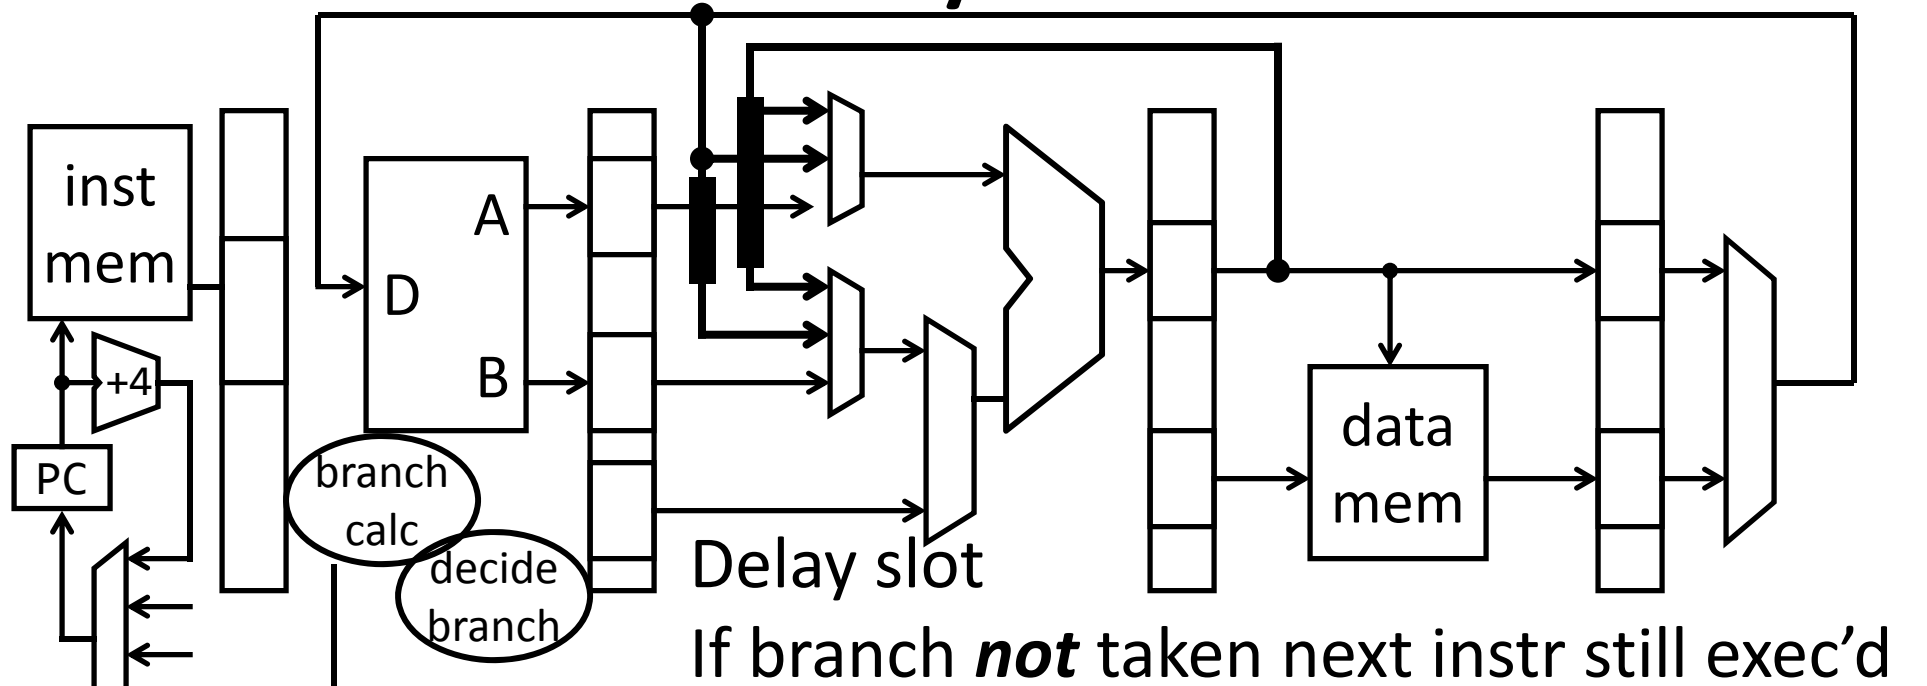
Delay Slot



If branch ***taken*** next instr still exec'd

10: beq r1, r2, L	IF	ID	Ex	M	W				
14: add r3, r0, r3		IF	ID	Ex	M	W			
18: sub r5, r4, r6									
1C: L: or r3, r2, r4			IF	ID	Ex	M	W		

Delay Slot



10: beq r1, r2, L	IF	ID	Ex	M	W				
14: add r3, r0, r3		IF	ID	Ex	M	W			
18: sub r5, r4, r6			IF	ID	Ex	M	W		
1C: L: or r3, r2, r4				IF	ID	Ex	M	W	

Control Hazards

Control Hazards

- instructions are fetched in stage 1 (IF)
- branch and jump decisions occur in stage 3 (EX)
i.e. next PC is not known until 2 cycles after branch/jump
- Can optimize and move branch and jump decision to stage 2 (ID)
i.e. next PC is not known until **1 cycles after** branch/jump

Stall (+ Zap)

- prevent PC update
- clear IF/ID pipeline register
 - instruction just fetched might be wrong one, so convert to nop
- allow branch to continue into EX stage

Delay Slot

- ISA says N instructions after branch/jump always executed
 - MIPS has 1 branch delay slot

Takeaway

Control hazards occur because the PC following a control instruction is not known until control instruction computes if branch should be taken or not. If branch taken, then need to zap/flush instructions. There still a performance penalty for branches: Need to stall, then may need to zap (flush) subsequent instructions that have already been fetched.

We can reduce cost of a control hazard by moving branch decision and calculation from Ex stage to ID stage. This reduces the cost from flushing two instructions to only flushing one.

Delay Slots can potentially increase performance due to control hazards by putting a useful instruction in the delay slot since the instruction in the delay slot will ***always*** be executed. Requires software (compiler) to make use of delay slot. Put nop in delay slot if not able to put useful instruction in delay slot.

Reduce cost of Ctrl Haz even further?

Speculative Execution

- “*Guess*” direction of the branch
 - Allow instructions to move through pipeline
 - Zap them later if wrong guess
- Useful for long pipelines

Speculative Execution: Loops

Pipeline so far

- “Guess” (predict) that the branch will not be taken

We can do better!

- Make prediction based on last branch
- Predict “take branch” if last branch “taken”
- Or Predict “do not take branch” if last branch “not taken”

- Need one bit to keep track of last branch

Speculative Execution: Loops

What is accuracy of branch predictor?

Wrong twice per loop!

Once on loop enter and exit

We can do better with 2 bits

While (r3 ≠ 0) {... r3--;}

Top: BEQZ r3, End



J Top

End:

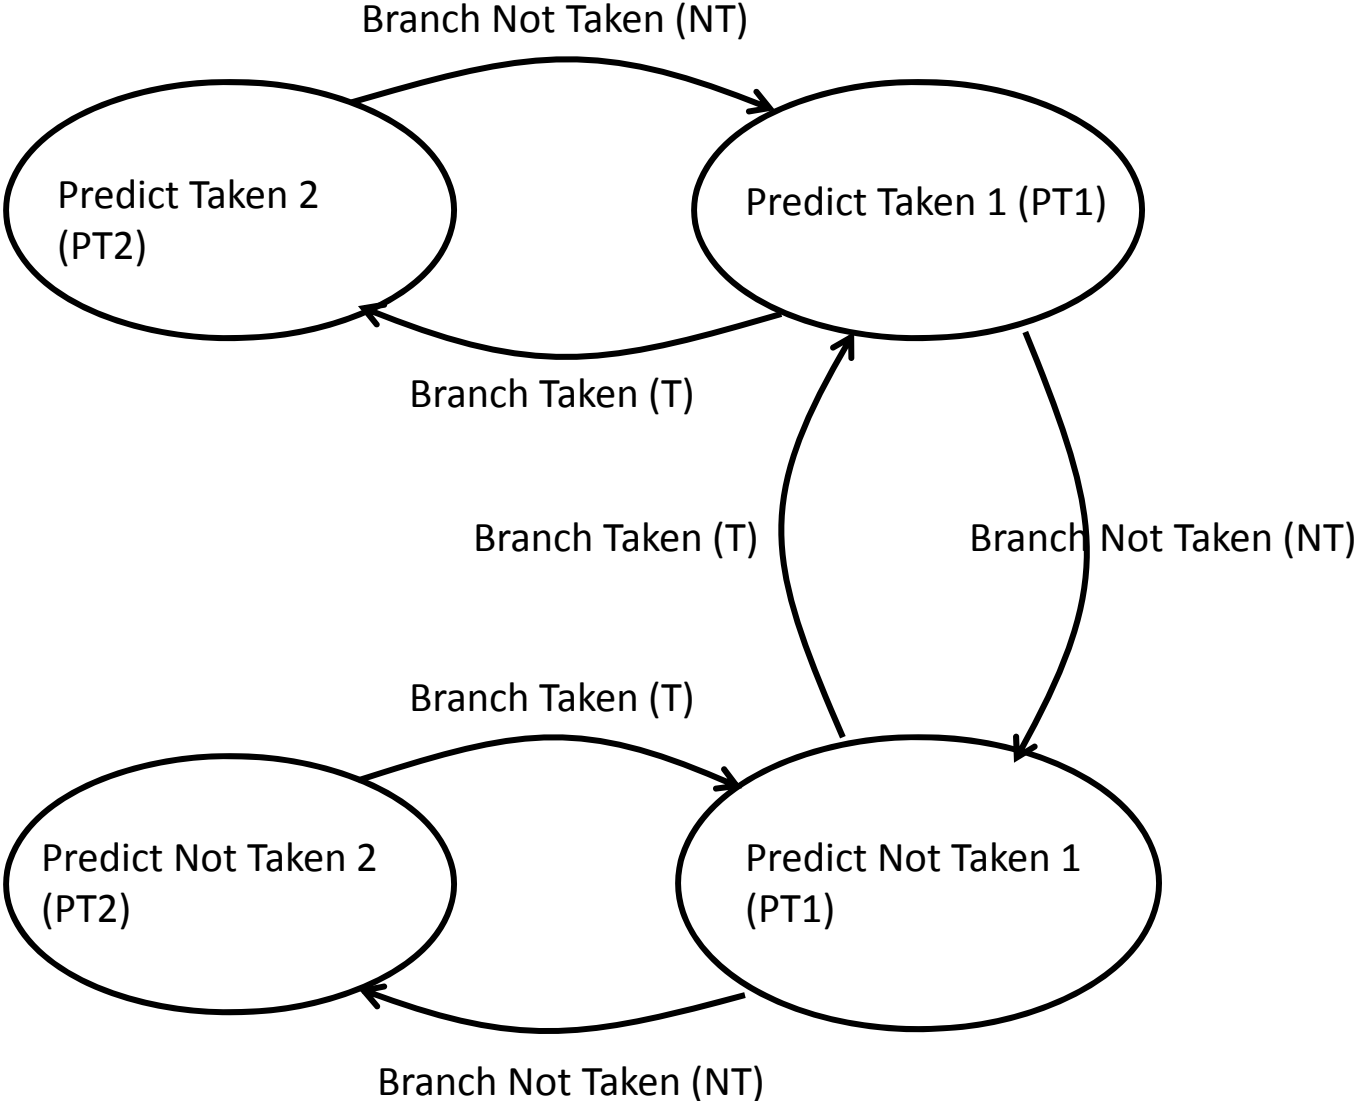
While (r3 ≠ 0) {... r3--;}

Top2: BEQZ r3, End2

J Top

End2:

Speculative Execution: Branch Execution



Summary

Control hazards

- Is branch taken or not?
- Performance penalty: stall and flush

Reduce cost of control hazards

- Move branch decision from Ex to ID
 - 2 nops to 1 nop
- Delay slot
 - Compiler puts useful work in delay slot. ISA level.
- Branch prediction
 - Correct. Great!
 - Wrong. Flush pipeline. Performance penalty

Hazards Summary

Data hazards

Control hazards

Structural hazards

- resource contention
- so far: impossible because of ISA and pipeline design

Hazards Summary

Data hazards

- register file reads occur in stage 2 (IF)
- register file writes occur in stage 5 (WB)
- next instructions may read values soon to be written

Control hazards

- branch instruction may change the PC in stage 3 (EX)
- next instructions have already started executing

Structural hazards

- resource contention
- so far: impossible because of ISA and pipeline design