

Data Hazards & Dynamic Instruction Scheduling (II)

Hung-Wei Tseng

Recap: Three pipeline hazards

- Structural hazards — resource conflicts cannot support simultaneous execution of instructions in the pipeline
- Control hazards — the PC can be changed by an instruction in the pipeline
- Data hazards — an instruction depending on a the result that's not yet generated or propagated when the instruction needs that

Recap: addressing hazards

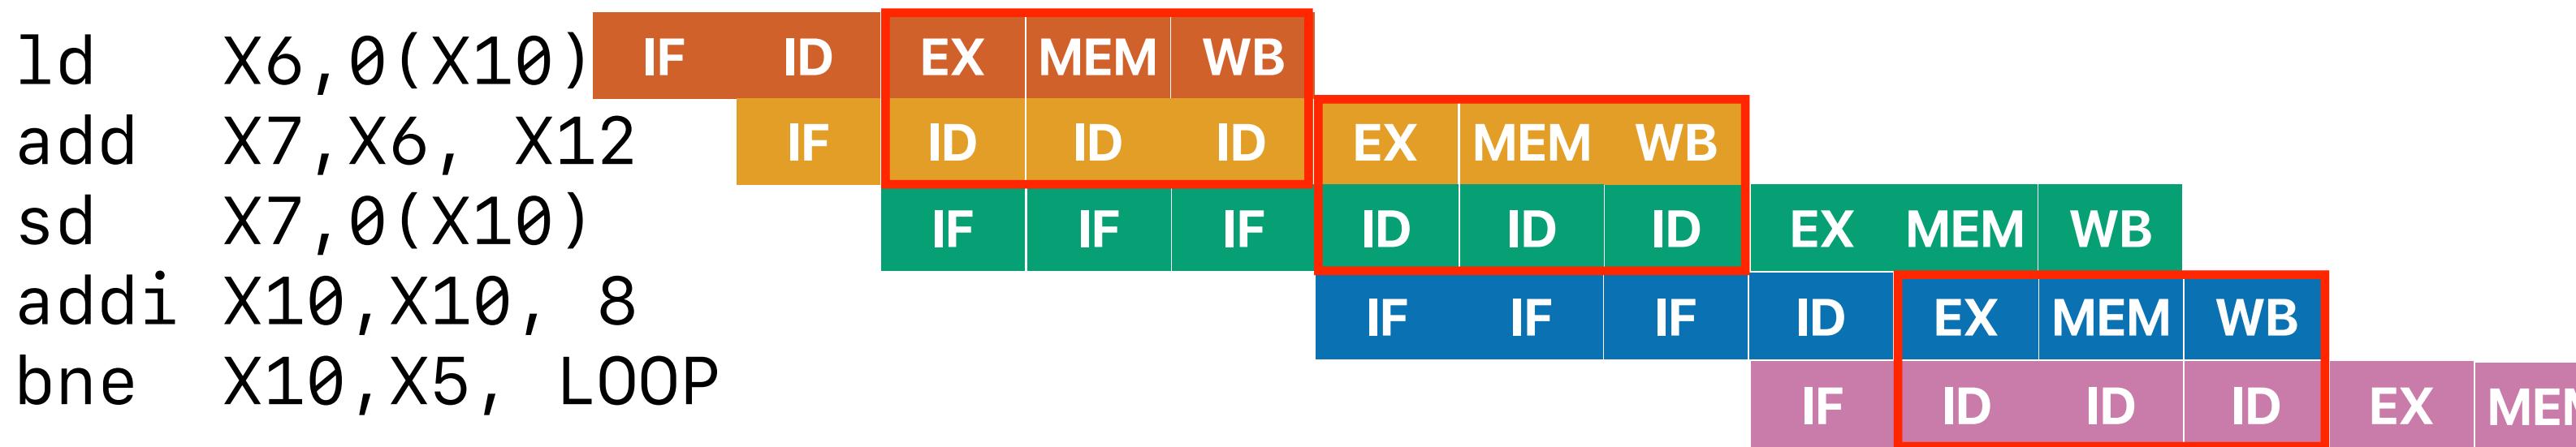
- Structural hazards
 - Stall
 - Modify hardware design
- Control hazards
 - Stall
 - Static prediction
 - Dynamic prediction

Data hazards

- An instruction currently in the pipeline cannot receive the “logically” correct value for execution
- Data dependencies
 - The output of an instruction is the input of a later instruction
 - May result in data hazard if the later instruction that consumes the result is still in the pipeline

How many of data hazards?

- How many pairs of instructions in the following RISC-V instructions will result in data hazards/stalls in a basic 5-stage RISC-V pipeline?



- A. 1
- B. 2
- C. 3
- D. 4
- E. 5

Team scores



7

12.5

9

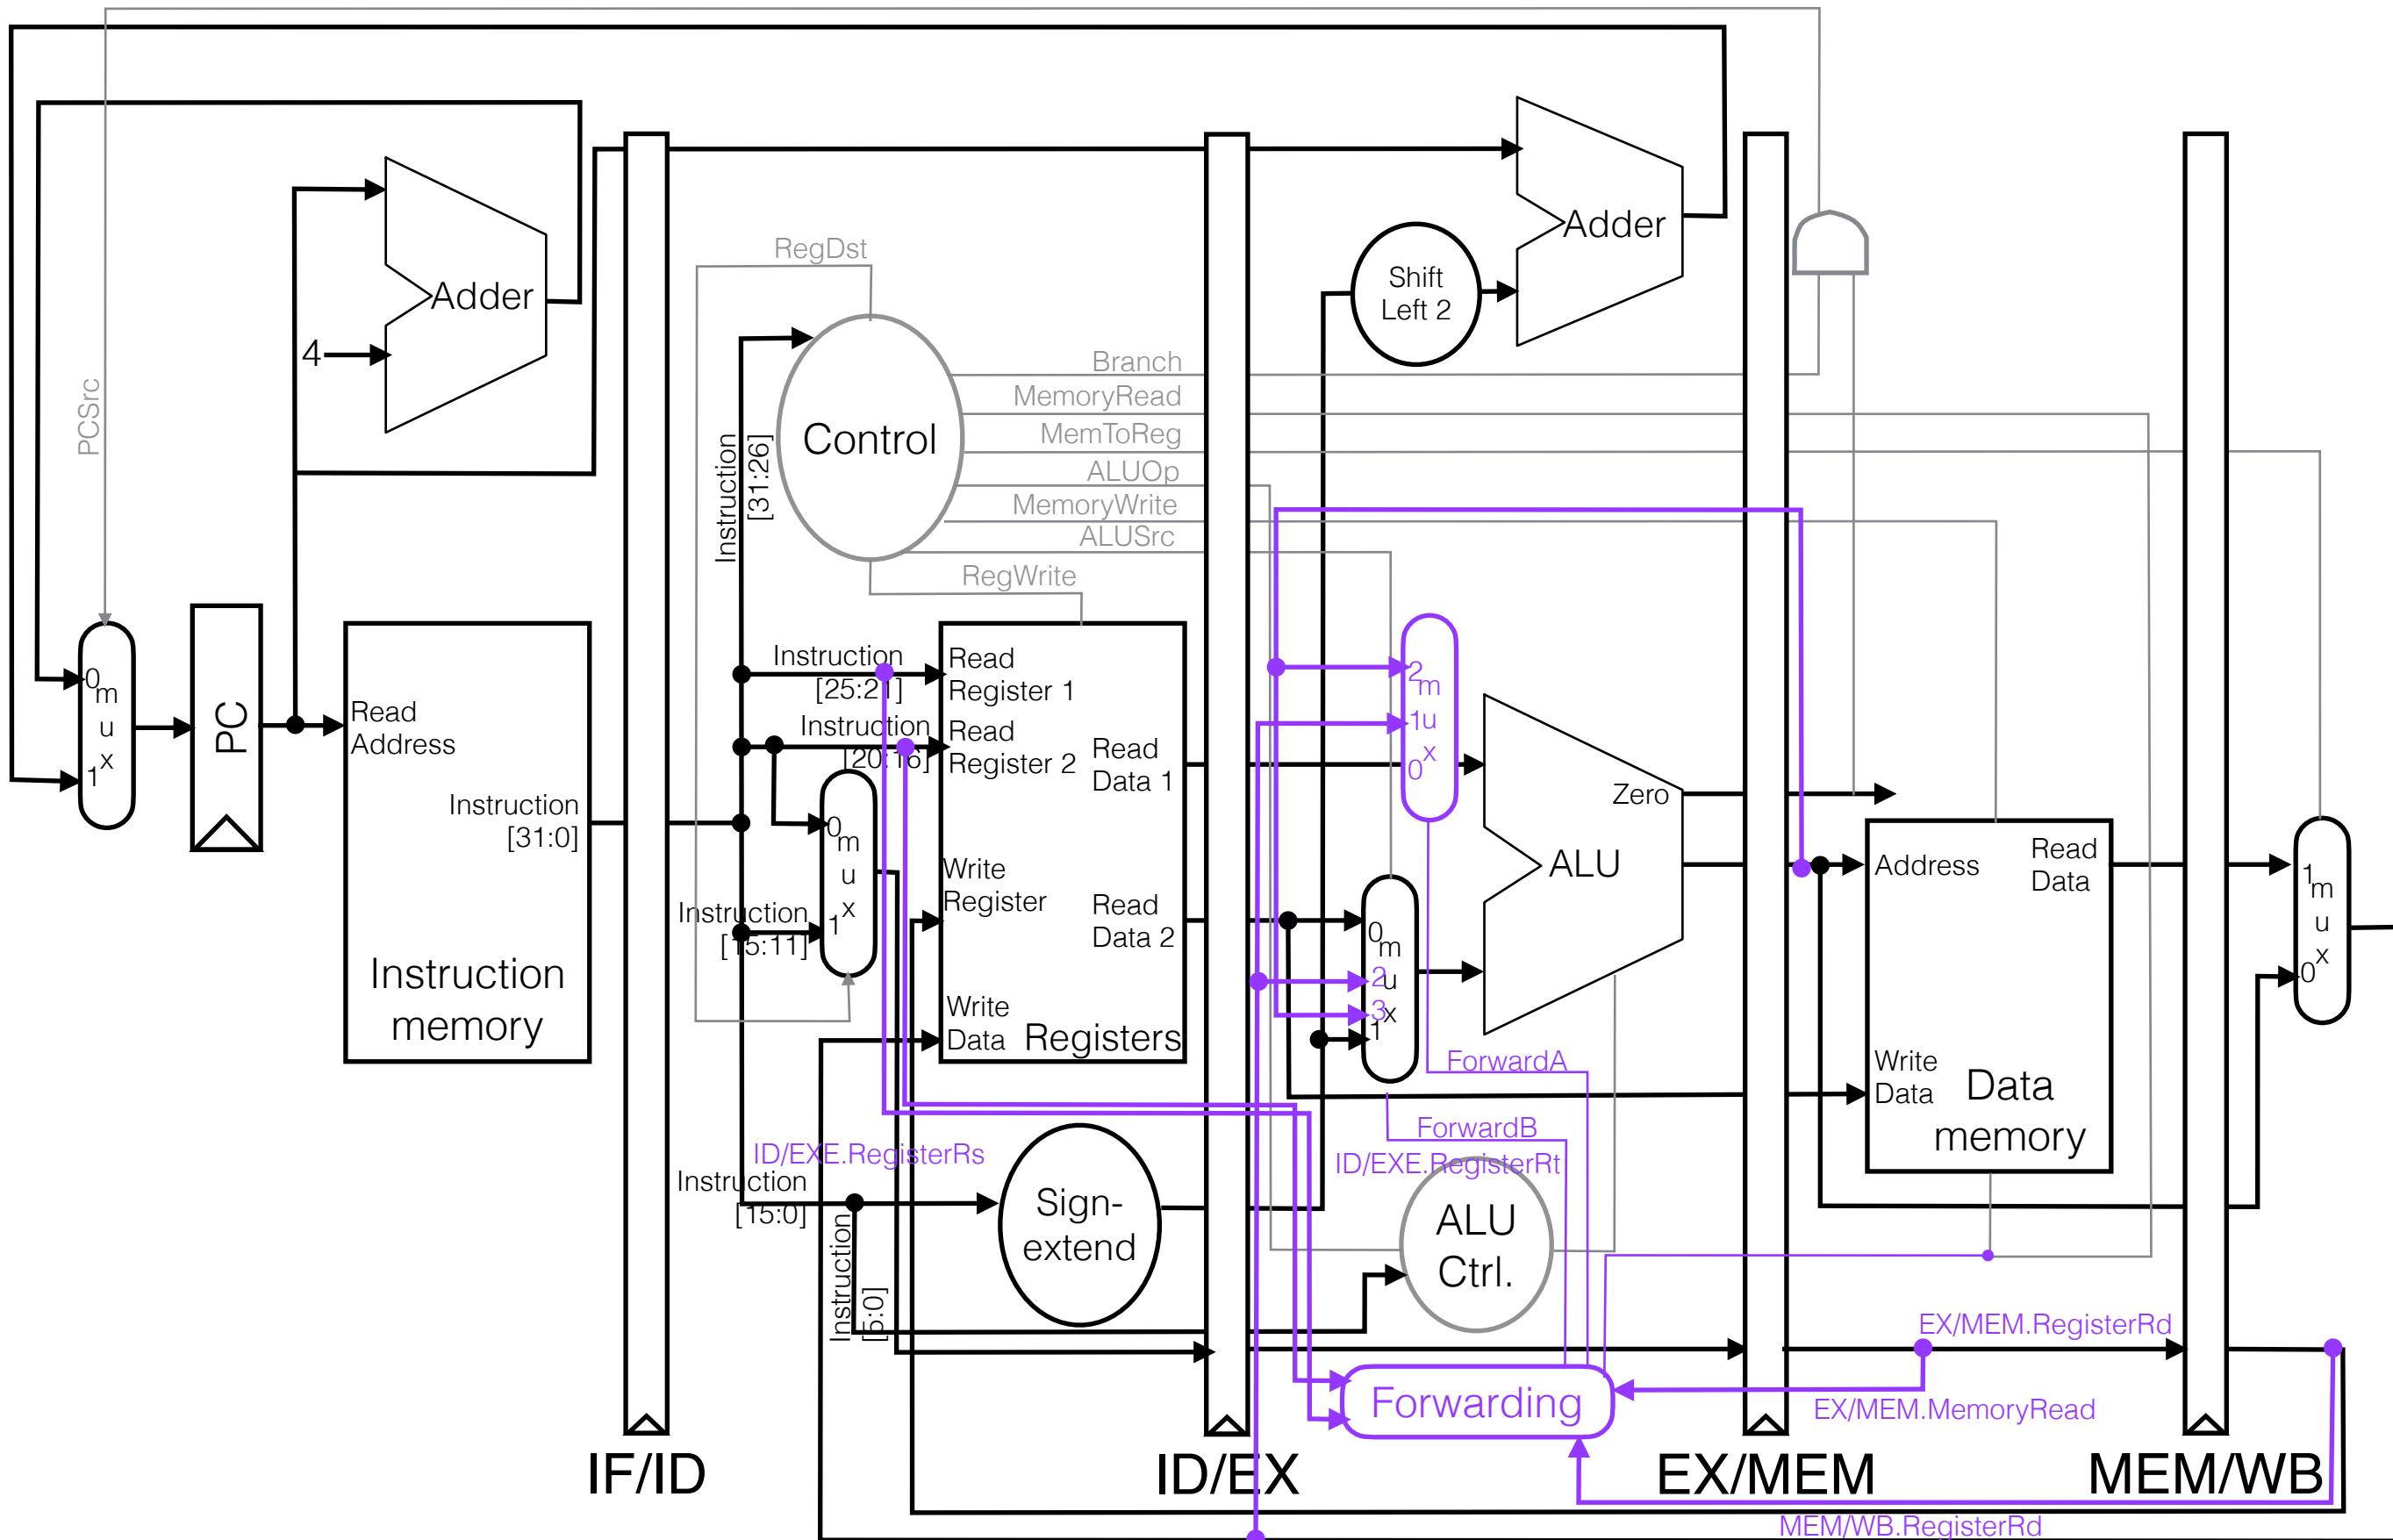
7

Outline

- Data forwarding
- Tomasulo's algorithm

Solution 2: Data forwarding

- Add logics/wires to forward the desired values to the demanding instructions
- In our five stage pipeline — if the instruction entering the EXE stage consumes a result from a previous instruction that is entering MEM stage or WB stage
 - A source of the instruction entering EXE stage is the destination of an instruction entering MEM/WB stage
 - The previous instruction must be an instruction that updates register file



How many of data hazards w/ Data Forwarding?

- How many pairs of instructions in the following RISC-V instructions will result in data hazards/stalls in a basic 5-stage RISC-V pipeline with "full" data forwarding?

```
ld    X6, 0(X10)
add   X7, X6, X12
sd    X7, 0(X10)
addi  X10, X10, 8
bne   X10, X5,  LOOP
```

- A. 0
- B. 1
- C. 2
- D. 3
- E. 4

How many of data hazards w/ Data Forwarding?

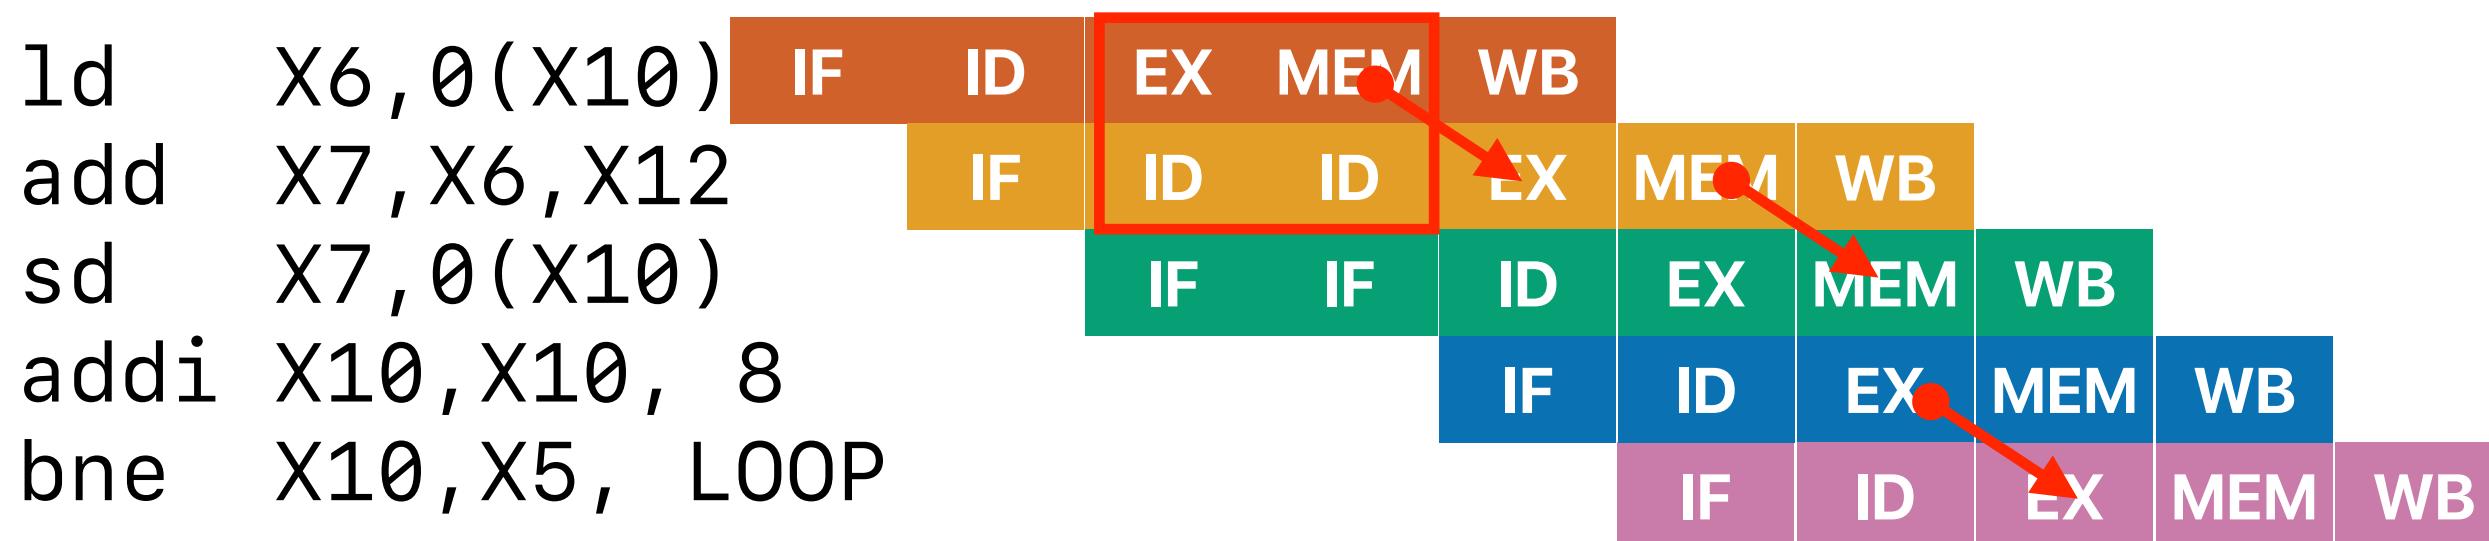
- How many pairs of instructions in the following RISC-V instructions will result in data hazards/stalls in a basic 5-stage RISC-V pipeline with "full" data forwarding?

```
ld    X6, 0(X10)
add   X7, X6, X12
sd    X7, 0(X10)
addi  X10, X10, 8
bne   X10, X5,  LOOP
```

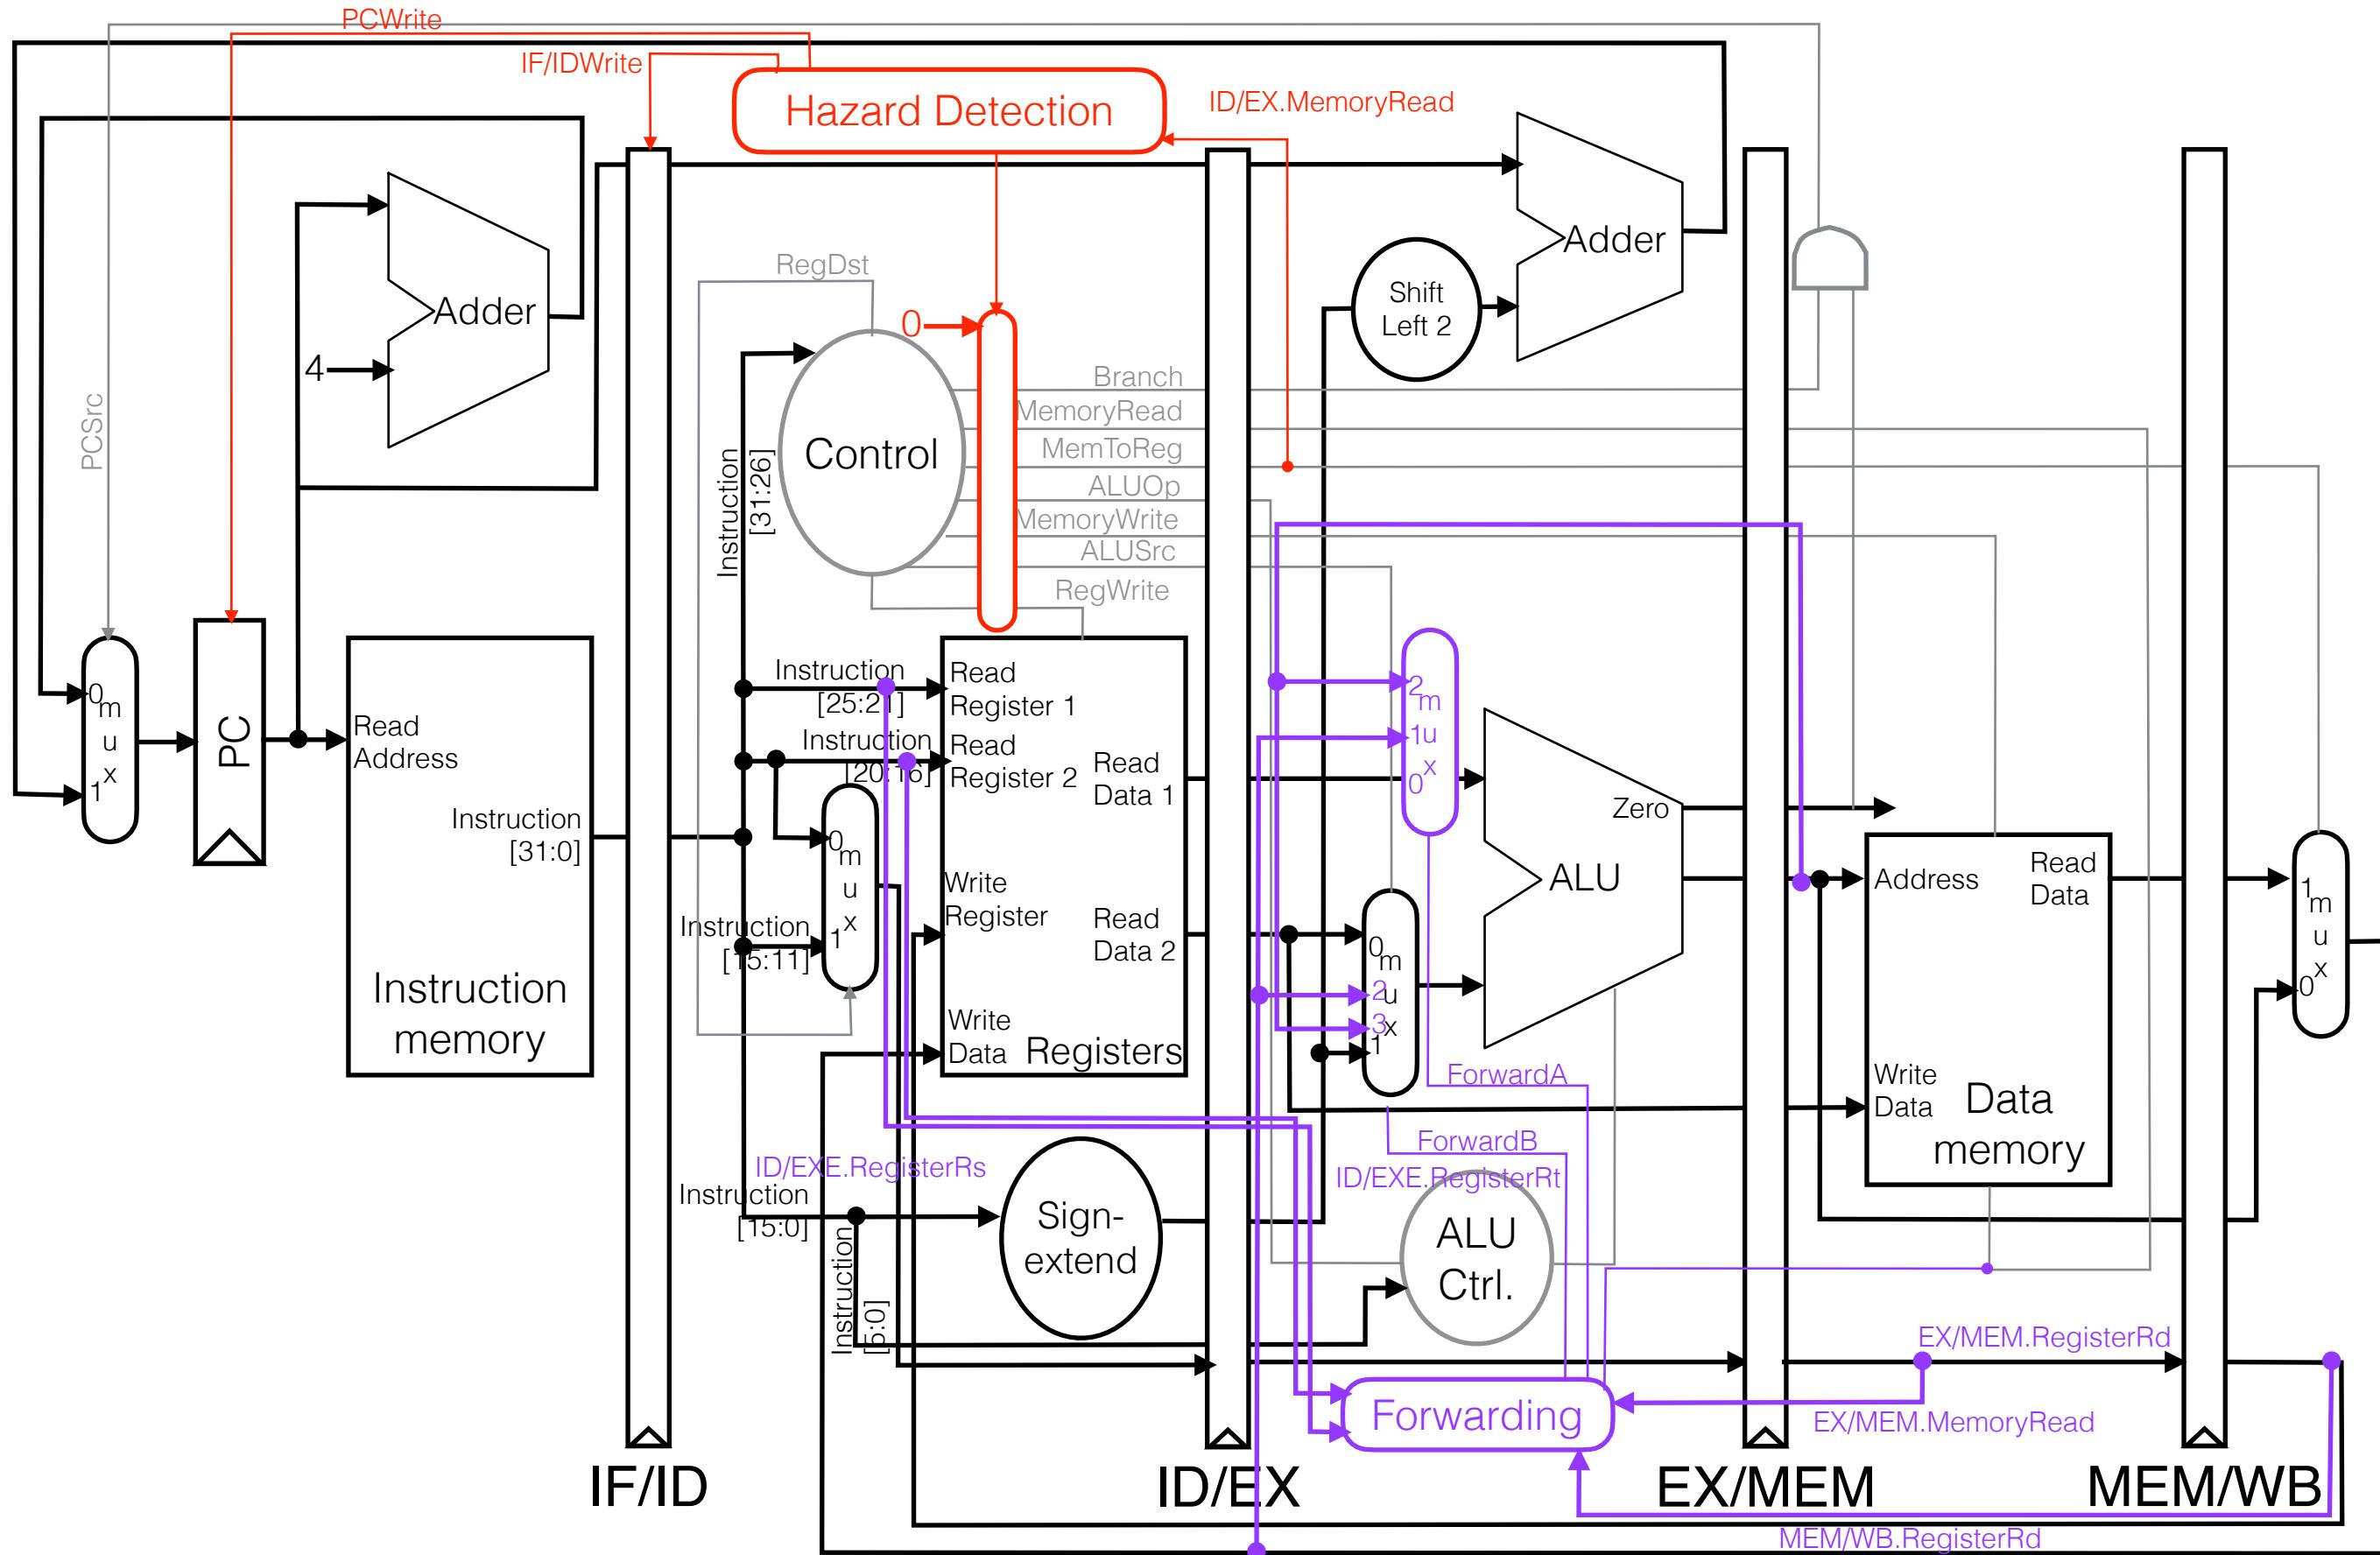
- A. 0
- B. 1
- C. 2
- D. 3
- E. 4

How many of data hazards w/ Data Forwarding?

- How many pairs of instructions in the following RISC-V instructions will result in data hazards/stalls in a basic 5-stage RISC-V pipeline with "full" data forwarding?

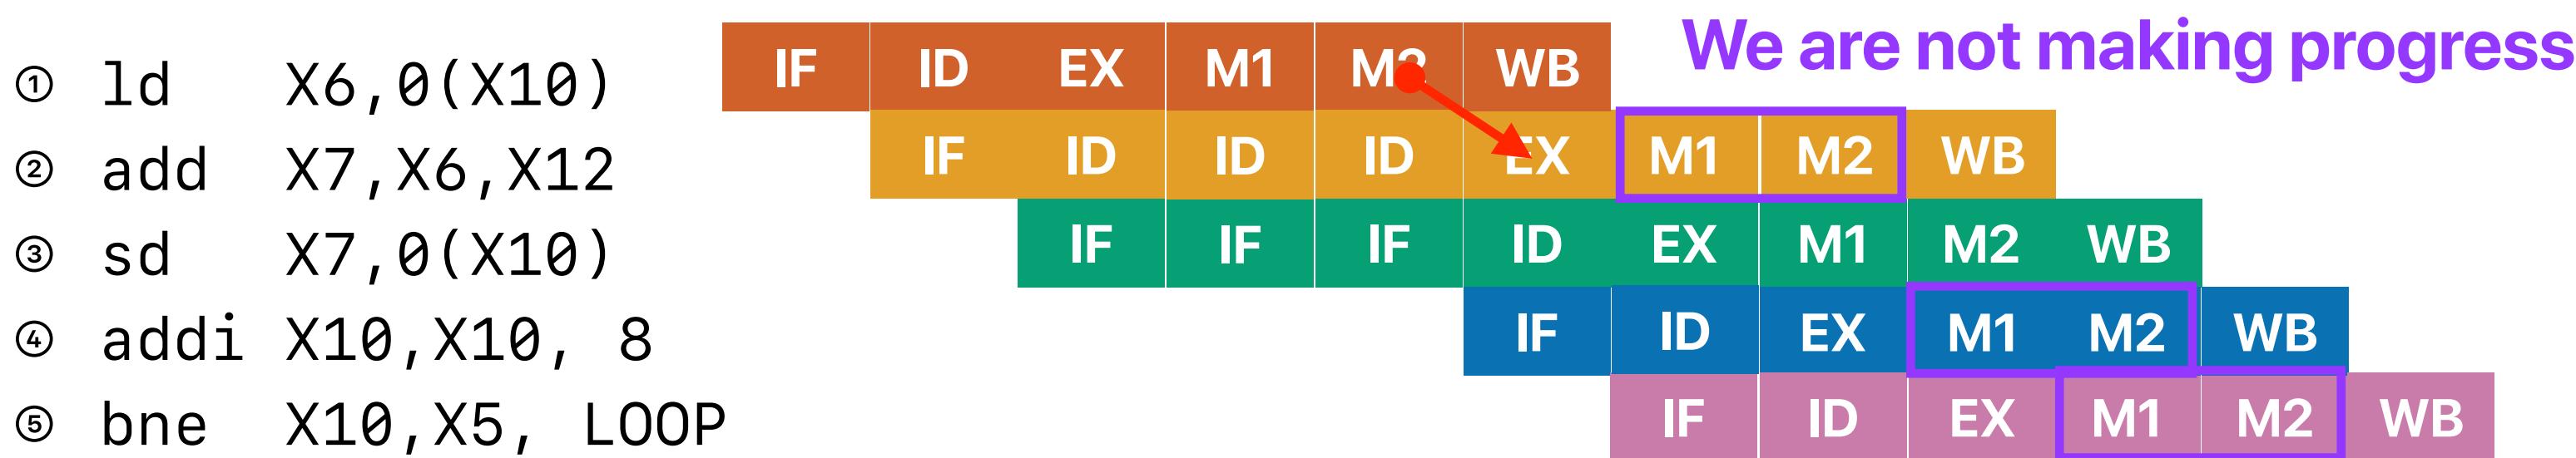


- A. 0
- B. 1**
- C. 2
- D. 3
- E. 4



Problems with data forwarding

- What if our pipeline gets deeper? — Considering a newly designed pipeline where memory stage is split into 2 stages and the memory access finishes at the 2nd memory stage. By reordering which pair of the following instruction stream can we eliminate all stalls without affecting the correctness of the code?



The effect of code optimization

- By reordering which pair of the following instruction stream can we eliminate all stalls without affecting the correctness of the code?
 - ① ld X6, 0(X10)
 - ② add X7, X6, X12
 - ③ sd X7, 0(X10)
 - ④ addi X10, X10, 8
 - ⑤ bne X10, X5, LOOP
- A. (1) & (2)
- B. (2) & (3)
- C. (3) & (4)
- D. (4) & (5)
- E. None of the pairs can be reordered

The effect of code optimization

- By reordering which pair of the following instruction stream can we eliminate all stalls without affecting the correctness of the code?
 - ① ld X6, 0(X10)
 - ② add X7, X6, X12
 - ③ sd X7, 0(X10)
 - ④ addi X10, X10, 8
 - ⑤ bne X10, X5, LOOP
 - A. (1) & (2)
 - B. (2) & (3)
 - C. (3) & (4)
 - D. (4) & (5)
 - E. None of the pairs can be reordered

The effect of code optimization

- By reordering which pair of the following instruction stream can we eliminate all stalls without affecting the correctness of the code?
 - ① ld X6, 0(X10)
 - ② add X7, X6, X12
 - ③ sd X7, 0(X10)
 - ④ addi X10, X10, 8
 - ⑤ bne X10, X5, LOOP
 - A. (1) & (2)
 - B. (2) & (3)
 - C. (3) & (4)
 - D. (4) & (5)
 - E. None of the pairs can be reordered

If we can predict the future ...

- Consider the following dynamic instructions:

- ① ld X6, 0(X10)
- ② add X7, X6, X12
- ③ sd X7, 0(X10)
- ④ addi X10, X10, 8
- ⑤ bne X10, X5, LOOP
- ⑥ ld X6, 0(X10)
- ⑦ add X7, X6, X12
- ⑧ sd X7, 0(X10)
- ⑨ addi X10, X10, 8
- ⑩ bne X10, X5, LOOP

Which of the following pair can we reorder without affecting the correctness if the **branch prediction is perfect**?

- A. (2) and (4)
- B. (3) and (5)
- C. (5) and (6)
- D. (6) and (9)
- E. (9) and (10)

If we can predict the future ...

- Consider the following dynamic instructions:

- ① ld X6, 0(X10)
- ② add X7, X6, X12
- ③ sd X7, 0(X10)
- ④ addi X10, X10, 8
- ⑤ bne X10, X5, LOOP
- ⑥ ld X6, 0(X10)
- ⑦ add X7, X6, X12
- ⑧ sd X7, 0(X10)
- ⑨ addi X10, X10, 8
- ⑩ bne X10, X5, LOOP

Which of the following pair can we reorder without affecting the correctness if the **branch prediction is perfect**?

- A. (2) and (4)
- B. (3) and (5)
- C. (5) and (6)
- D. (6) and (9)
- E. (9) and (10)

If we can predict the future ...

- Consider the following dynamic instructions:

- ① ld X6, 0(X10)
- ② add X7, X6, X12
- ③ sd X7, 0(X10)
- ④ addi X10, X10, 8
- ⑤ bne X10, X5, LOOP
- ⑥ ld X6, 0(X10)
- ⑦ add X7, X6, X12
- ⑧ sd X7, 0(X10)
- ⑨ addi X10, X10, 8
- ⑩ bne X10, X5, LOOP

Can we use “branch prediction” to predict the future and reorder instructions across the branch?

Which of the following pair can we reorder without affecting the correctness if the **branch prediction is perfect**?

- A. (2) and (4)
- B. (3) and (5)
- C. (5) and (6)
- D. (6) and (9)
- E. (9) and (10)

Dynamic instruction scheduling/ Out-of-order (OoO) execution

Tips of drawing a pipeline diagram

- Each instruction has to go through all 5 pipeline stages: IF, ID, EXE, MEM, WB in order — only valid if it's single-issue, RISC-V 5-stage pipeline
- An instruction can enter the next pipeline stage in the next cycle if
 - No other instruction is occupying the next stage
 - This instruction has completed its own work in the current stage
 - The next stage has all its inputs ready
- Fetch a new instruction only if
 - We know the next PC to fetch
 - We can predict the next PC
 - Flush an instruction if the branch resolution says it's mis-predicted.

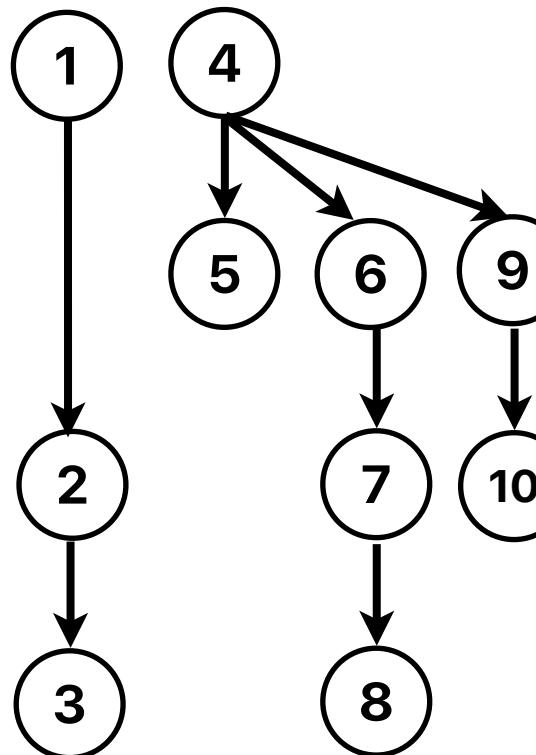
What do you need to execution an instruction?

- Whenever the instruction is decoded — put decoded instruction somewhere
- Whenever the inputs are ready — **all data dependencies are resolved**
- Whenever the target functional unit is available

Scheduling instructions: based on data dependencies

- Draw the data dependency graph, put an arrow if an instruction depends on the other.

- ① ld X6, 0(X10)
- ② add X7, X6, X12
- ③ sd X7, 0(X10)
- ④ addi X10, X10, 8
- ⑤ bne X10, X5, LOOP
- ⑥ ld X6, 0(X10)
- ⑦ add X7, X6, X12
- ⑧ sd X7, 0(X10)
- ⑨ addi X10, X10, 8
- ⑩ bne X10, X5, LOOP



- **In theory**, instructions without dependencies can be executed in parallel or out-of-order
- Instructions with dependencies can never be reordered

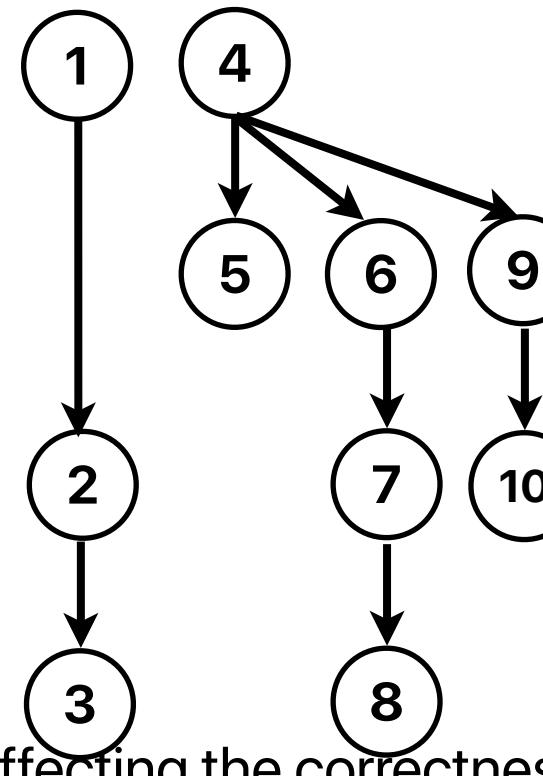
If we can predict the future ...

- Consider the following dynamic instructions:

- ① ld X6, 0(X10)
- ② add X7, X6, X12
- ③ sd X7, 0(X10)
- ④ addi X10, X10, 8
- ⑤ bne X10, X5, LOOP
- ⑥ ld X6, 0(X10)
- ⑦ add X7, X6, X12
- ⑧ sd X7, 0(X10)
- ⑨ addi X10, X10, 8
- ⑩ bne X10, X5, LOOP

Which of the following pair can we reorder without affecting the correctness if the **branch prediction is perfect**?

- A. (2) and (4)
- B. (3) and (5)
- C. (5) and (6)
- D. (6) and (9)
- E. (9) and (10)

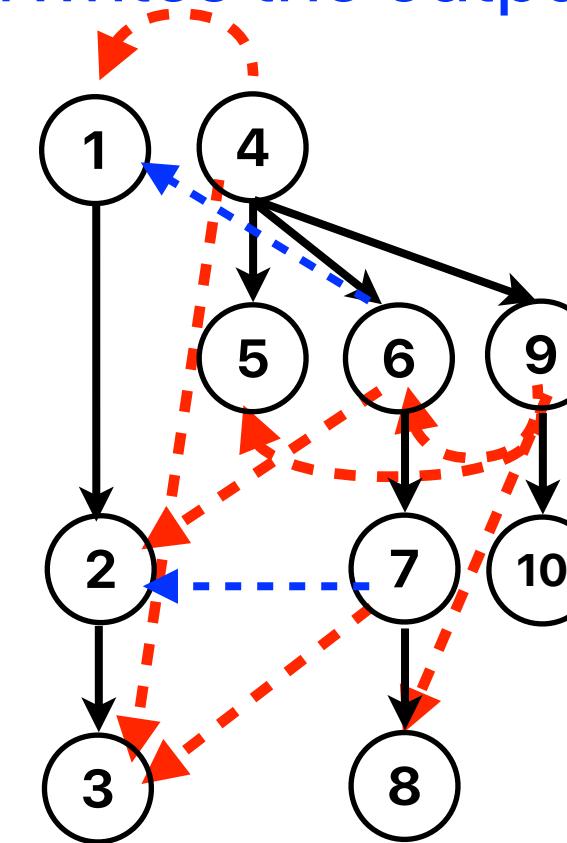


We still can only reorder (5) and (6)
even though (2) & (4) are not
depending on each other!

False dependencies

- We are still limited by **false dependencies**
- They are not “true” dependencies because they don’t have an arrow in data dependency graph
 - WAR (Write After Read): a later instruction overwrites the source of an earlier one
 - 4 and 1 4 and 3, 6 and 2, 7 and 3, 9 and 5, 9 and 6, 9 and 8
 - WAW (Write After Write): a later instruction overwrites the output of an earlier one
 - 6 and 1, 7 and 2

①	ld	X6, 0(X10)
②	add	X7, X6, X12
③	sd	X7, 0(X10)
④	addi	X10, X10, 8
⑤	bne	X10, X5, LOOP
⑥	ld	X6, 0(X10)
⑦	add	X7, X6, X12
⑧	sd	X7, 0(X10)
⑨	addi	X10, X10, 8
⑩	bne	X10, X5, LOOP



False dependencies

- Consider the following dynamic instructions

- ① ld X12, 0(X20)
- ② add X12, X10, X12
- ③ sub X18, X12, X10
- ④ ld X12, 8(X20)
- ⑤ add X14, X18, X12
- ⑥ add X18, X14, X14
- ⑦ sd X14, 16(X20)
- ⑧ addi X20, X20, 8

which of the following pair is not a “false dependency”

- A. (1) and (4)
- B. (1) and (8)
- C. (5) and (7)
- D. (4) and (8)
- E. (7) and (8)

False dependencies

- Consider the following dynamic instructions

- ① ld X12, 0(X20)
- ② add X12, X10, X12
- ③ sub X18, X12, X10
- ④ ld X12, 8(X20)
- ⑤ add X14, X18, X12
- ⑥ add X18, X14, X14
- ⑦ sd X14, 16(X20)
- ⑧ addi X20, X20, 8

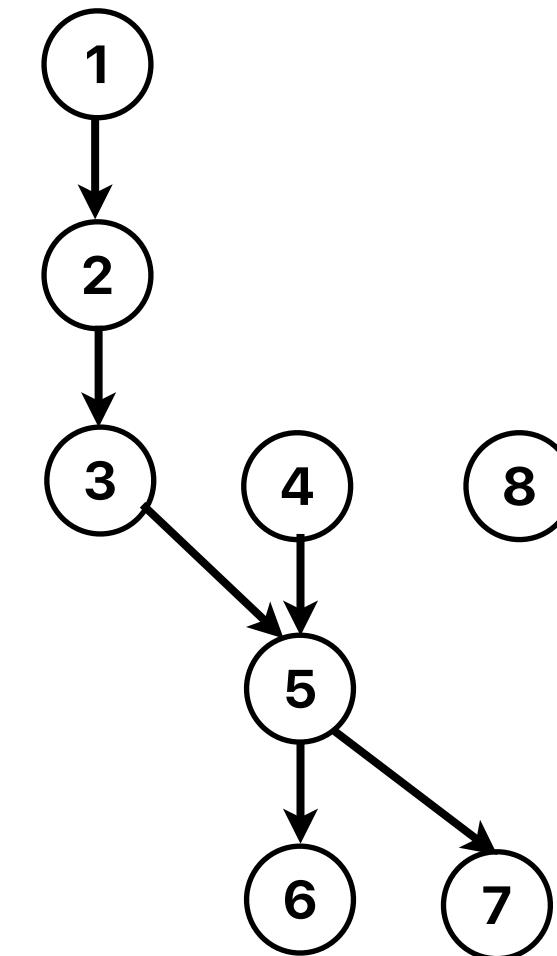
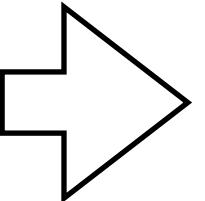
which of the following pair is not a “false dependency”

- A. (1) and (4)
- B. (1) and (8)
- C. (5) and (7)
- D. (4) and (8)
- E. (7) and (8)

False dependencies

- Consider the following dynamic instructions

- ① ld X12, 0(X20)
- ② add X12, X10, X12
- ③ sub X18, X12, X10
- ④ ld X12, 8(X20)
- ⑤ add X14, X18, X12
- ⑥ add X18, X14, X14
- ⑦ sd X14, 16(X20)
- ⑧ addi X20, X20, 8



which of the following pair is not a “false dependency”

- | | |
|----------------|------------------------------|
| A. (1) and (4) | WAW |
| B. (1) and (8) | WAR |
| C. (5) and (7) | True dependency (RAW) |
| D. (4) and (8) | WAR |
| E. (7) and (8) | WAR |

Out-of-order execution

- Any sequence of instructions has set of RAW, WAW, and WAR hazards that constrain its execution.
- Can we design a processor that extracts as much parallelism as possible, while still respecting these dependences?



Tomasulo's Algorithm

IBM 360 Model 91

- The most powerful commercialized machine in 1968

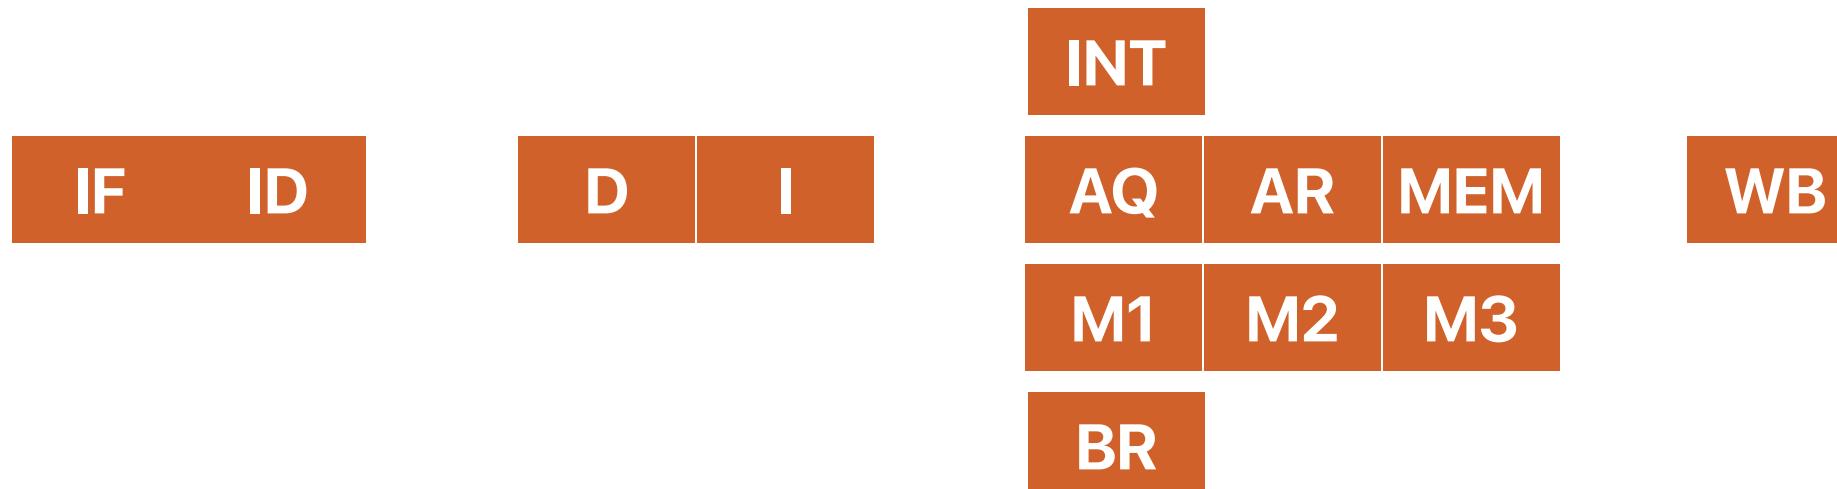


CDC 6600

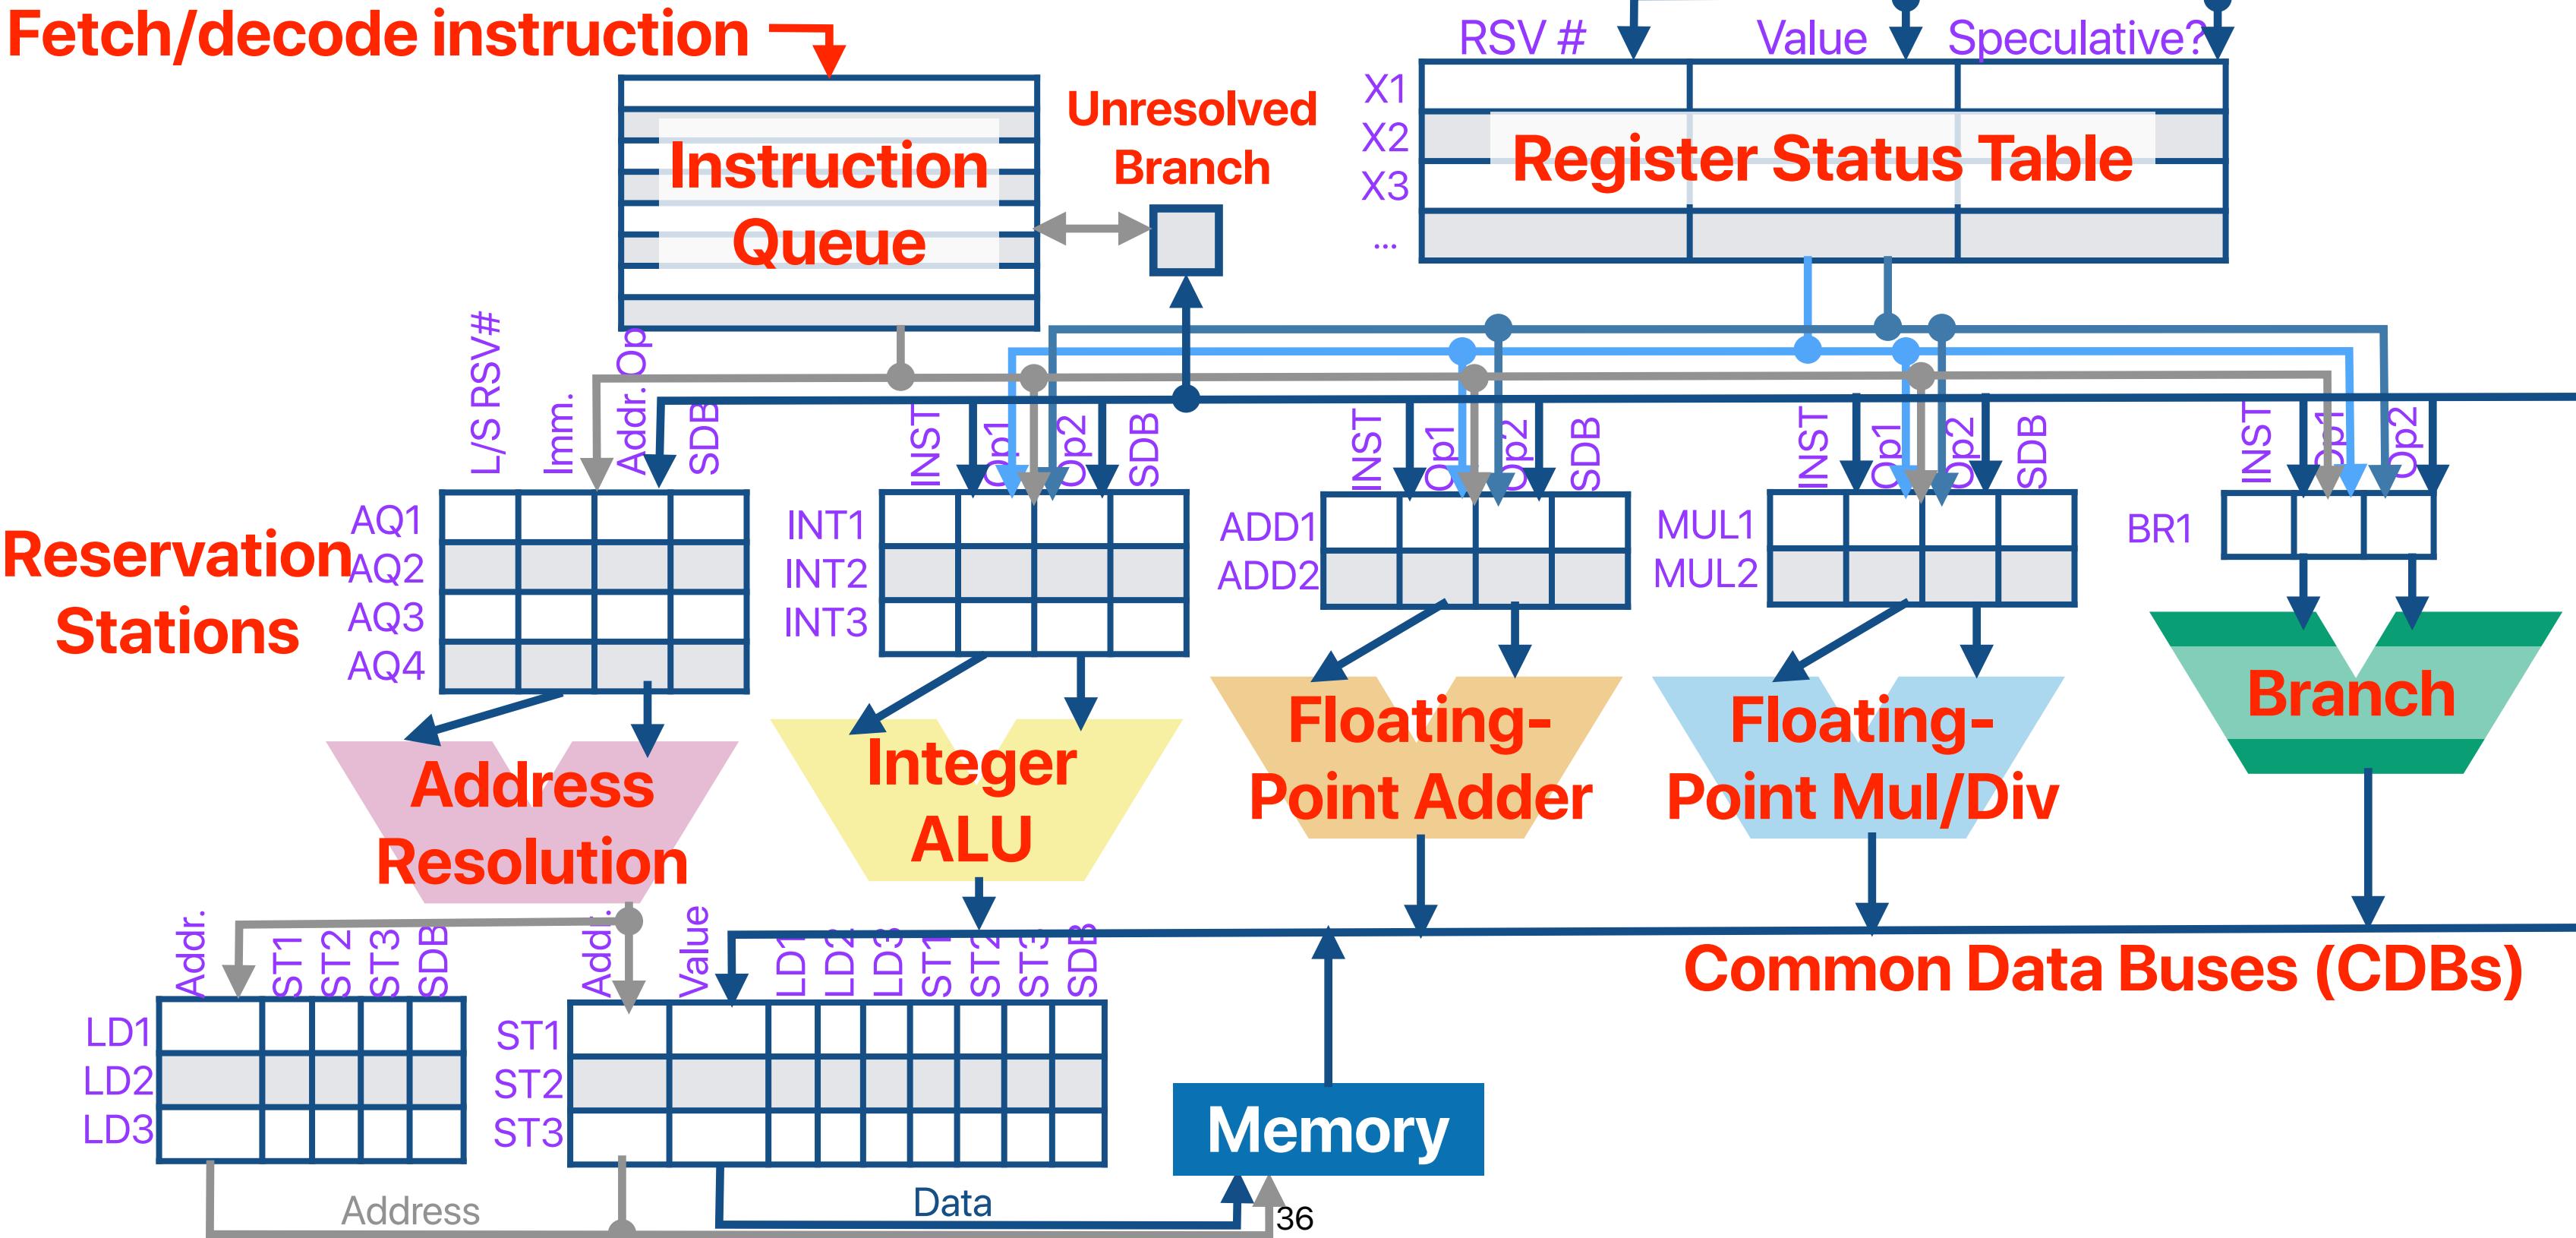
- The first machine with out-of-order execution.
- With performance of up to 3 "M" FLOPS
- It uses score-boarding to achieve so.
 - Instruction storage added to each functional execution unit
 - Instructions issue to FU when no structural hazards, begin execution when dependences satisfied. Thus, instructions issued to different FUs can execute out of order.
 - "scoreboard" tracks RAW, WAR, WAW hazards, tells each instruction when to proceed.
 - No forwarding
 - No register renaming

Pipeline in Tomasulo

- Dispatch (D) — allocate a “reservation station” for a decoded instruction
- Issue (I) — collect pending values/branch outcome from common data bus
- Execute (INT, AQ/AQ/MEM, M1/M2/M3, BR) — send the instruction to its corresponding pipeline if no structural hazards
- Write Back (WB) — broadcast the result through CDB



Overview of a processor supporting Tomasulo's algorithm



Tomasulo in motion

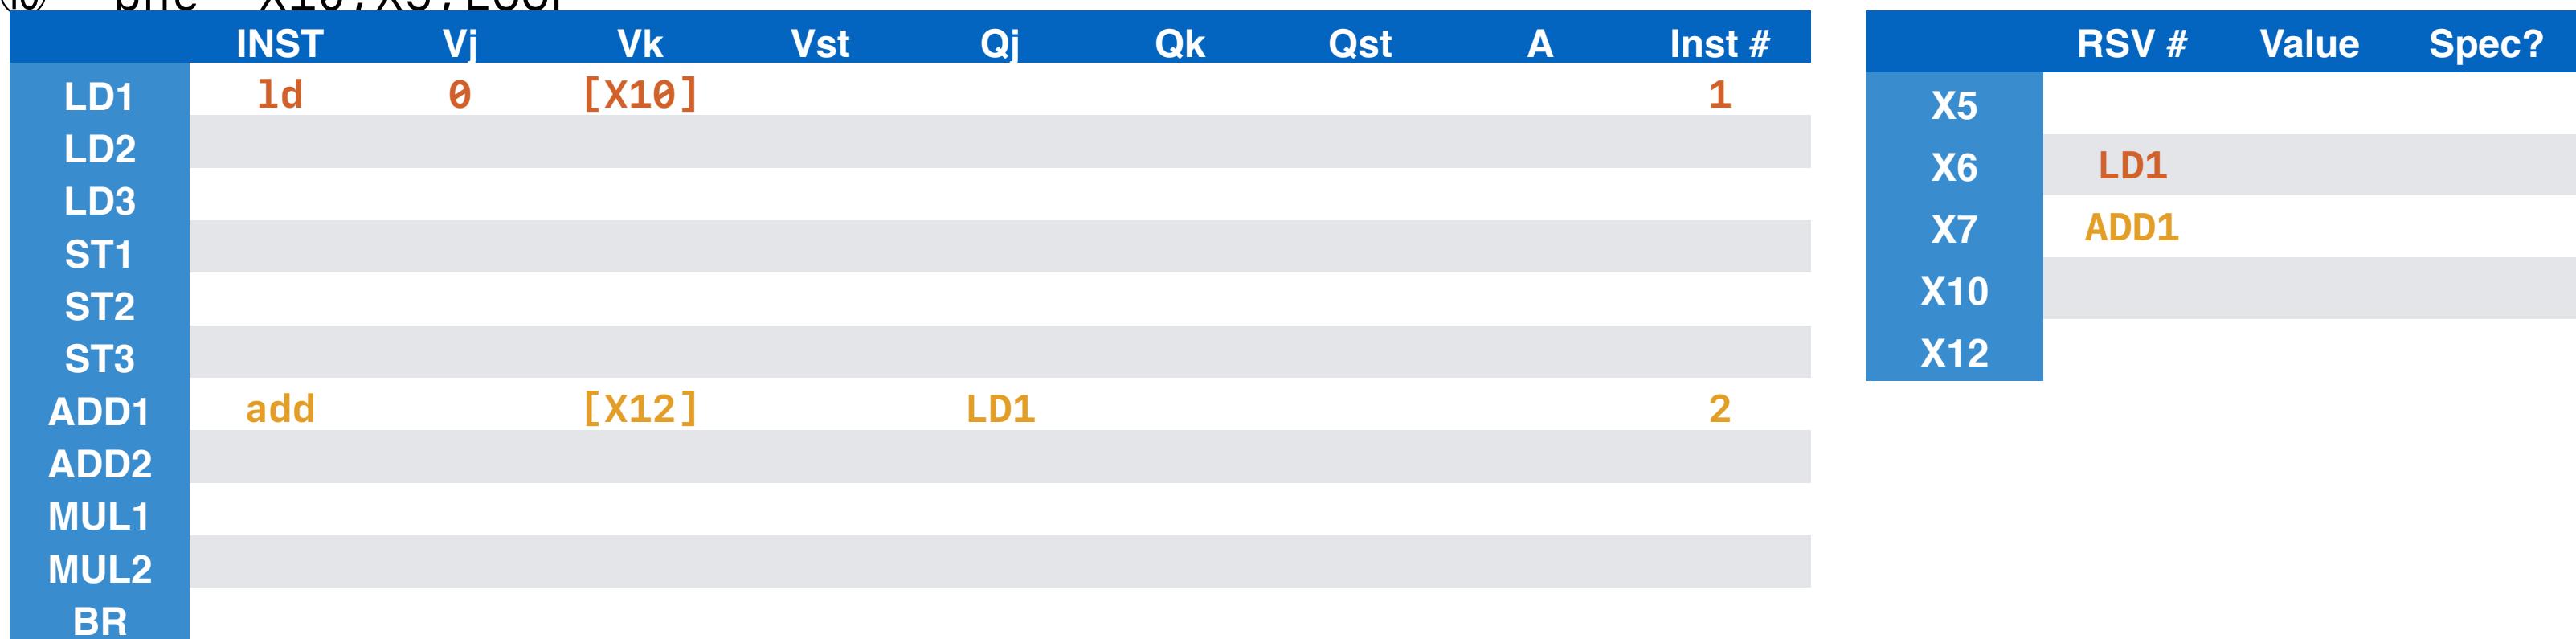
- ① ld X6, 0(X10)
 - ② add X7, X6, X12
 - ③ sd X7, 0(X10)
 - ④ addi X10, X10, 8
 - ⑤ bne X10, X5, LOOP
 - ⑥ ld X6, 0(X10)
 - ⑦ add X7, X6, X12
 - ⑧ sd X7, 0(X10)
 - ⑨ addi X10, X10, 8
 - ⑩ bne X10, X5, LOOP

Tomasulo in motion

- ① ld X6, 0(X10)
- ② add X7, X6, X12
- ③ sd X7, 0(X10)
- ④ addi X10, X10, 8
- ⑤ bne X10, X5, LOOP
- ⑥ ld X6, 0(X10)
- ⑦ add X7, X6, X12
- ⑧ sd X7, 0(X10)
- ⑨ addi X10, X10, 8
- ⑩ bne X10, X5, LOOP

Tomasulo in motion

- ① ld X6, 0(X10)
 - ② add X7, X6, X12
 - ③ sd X7, 0(X10)
 - ④ addi X10, X10, 8
 - ⑤ bne X10, X5, LOOP
 - ⑥ ld X6, 0(X10)
 - ⑦ add X7, X6, X12
 - ⑧ sd X7, 0(X10)
 - ⑨ addi X10, X10, 8
 - ⑩ bne X10, X5, LOOP



Tomasulo in motion

①	ld	X6, 0(X10)	D	AQ	AR
②	add	X7, X6, X12		D	I
③	sd	X7, 0(X10)			D

- ① ld X6, 0(X10)
 - ② add X7, X6, X12
 - ③ sd X7, 0(X10)
 - ④ addi X10, X10, 8
 - ⑤ bne X10, X5, LOOP
 - ⑥ ld X6, 0(X10)
 - ⑦ add X7, X6, X12
 - ⑧ sd X7, 0(X10)
 - ⑨ addi X10, X10, 8
 - ⑩ bne X10, X5, LOOP

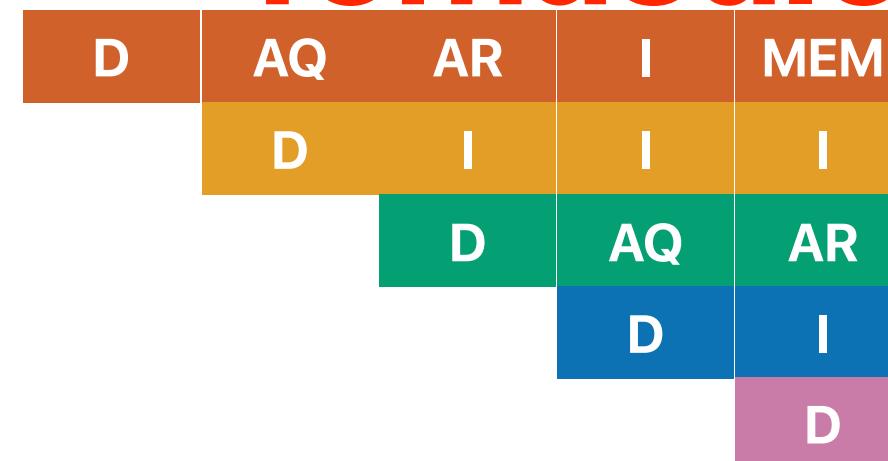
Tomasulo in motion

			D	AQ	AR	I
②	add	X7,X6,X12		D	I	I
③	sd	X7,0(X10)		D		AQ
④	addi	X10,X10,8				D

- ① ld X6, 0(X10)
 - ② add X7, X6, X12
 - ③ sd X7, 0(X10)
 - ④ addi X10, X10, 8
 - ⑤ bne X10, X5, LOOP
 - ⑥ ld X6, 0(X10)
 - ⑦ add X7, X6, X12
 - ⑧ sd X7, 0(X10)
 - ⑨ addi X10, X10, 8
 - ⑩ bne X10, X5, LOOP

Tomasulo in motion

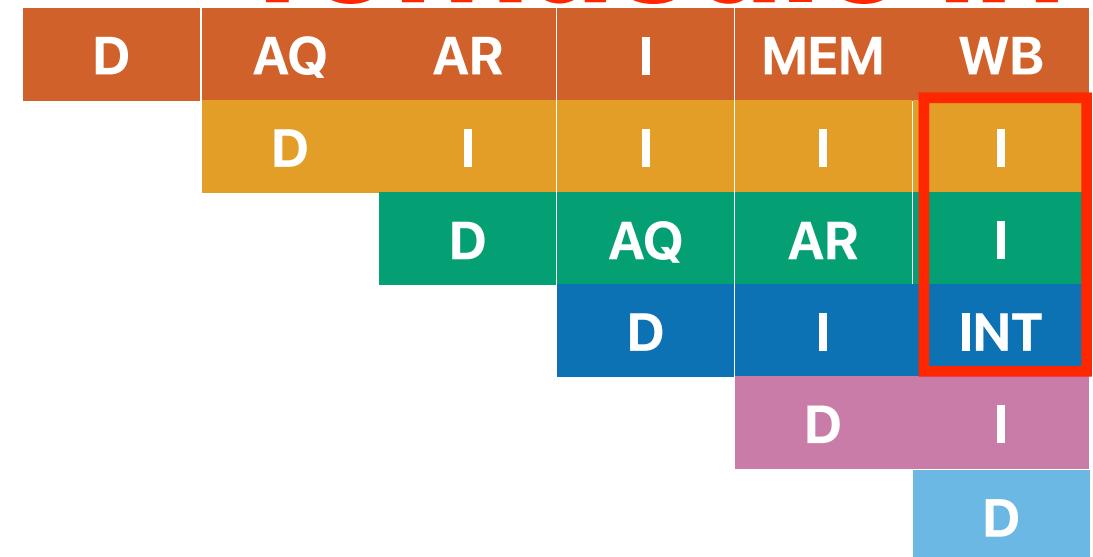
- ① ld X6, 0(X10)
- ② add X7, X6, X12
- ③ sd X7, 0(X10)
- ④ addi X10, X10, 8
- ⑤ bne X10, X5, LOOP
- ⑥ ld X6, 0(X10)
- ⑦ add X7, X6, X12
- ⑧ sd X7, 0(X10)
- ⑨ addi X10, X10, 8
- ⑩ bne X10, X5, LOOP



	INST	Vj	Vk	Vst	Qj	Qk	Qst	A	Inst #	RSV #	Value	Spec?
LD1	ld	0	[X10]						1	X5		
LD2										X6	LD1	
LD3										X7	ADD1	
ST1	sd	0	[X10]				ADD1		3	X10	ADD2	
ST2										X12		
ST3												
ADD1	add		[X12]		LD1				2			
ADD2	addi	8	[X10]						4			
MUL1												
MUL2												
BR	bne		[X5]		ADD2				5			

Tomasulo in motion

- ① ld X6, 0(X10)
- ② add X7, X6, X12
- ③ sd X7, 0(X10)
- ④ addi X10, X10, 8
- ⑤ bne X10, X5, LOOP
- ⑥ ld X6, 0(X10)
- ⑦ add X7, X6, X12
- ⑧ sd X7, 0(X10)
- ⑨ addi X10, X10, 8
- ⑩ bne X10, X5, LOOP



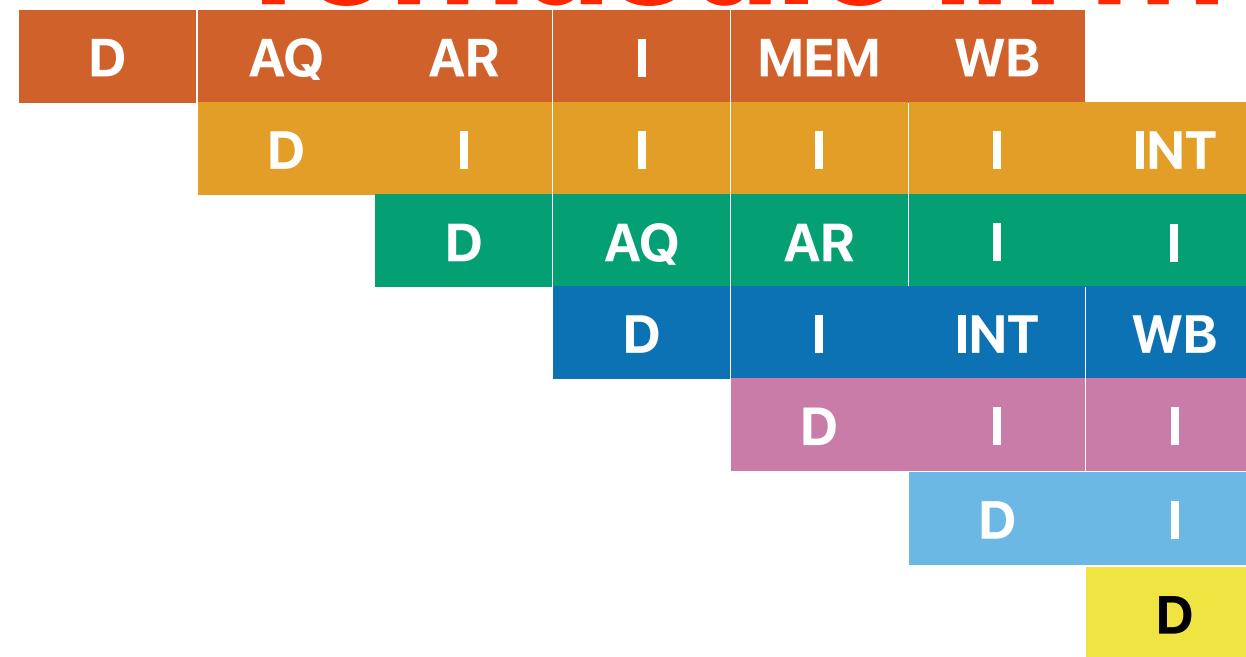
addi is now ahead of add!

	INST	Vj	Vk	Vst	Qj	Qk	Qst	A	Inst #
LD1	ld	0	[X10]						1
LD2	ld	0			ADD2				6
LD3									
ST1	sd	0	[X10]				ADD1		3
ST2									
ST3									
ADD1	add	LD1	[X12]						2
ADD2	addi	8	[X10]						4
MUL1									
MUL2									
BR	bne		[X5]		ADD2				5

RSV #	Value	Spec?
X5		
X6	LD2	LD1 1
X7	ADD1	
X10	ADD2	
X12		

Tomasulo in motion

- ① ld X6, 0(X10)
- ② add X7, X6, X12
- ③ sd X7, 0(X10)
- ④ addi X10, X10, 8
- ⑤ bne X10, X5, LOOP
- ⑥ ld X6, 0(X10)
- ⑦ add X7, X6, X12
- ⑧ sd X7, 0(X10)
- ⑨ addi X10, X10, 8
- ⑩ bne X10, X5, LOOP

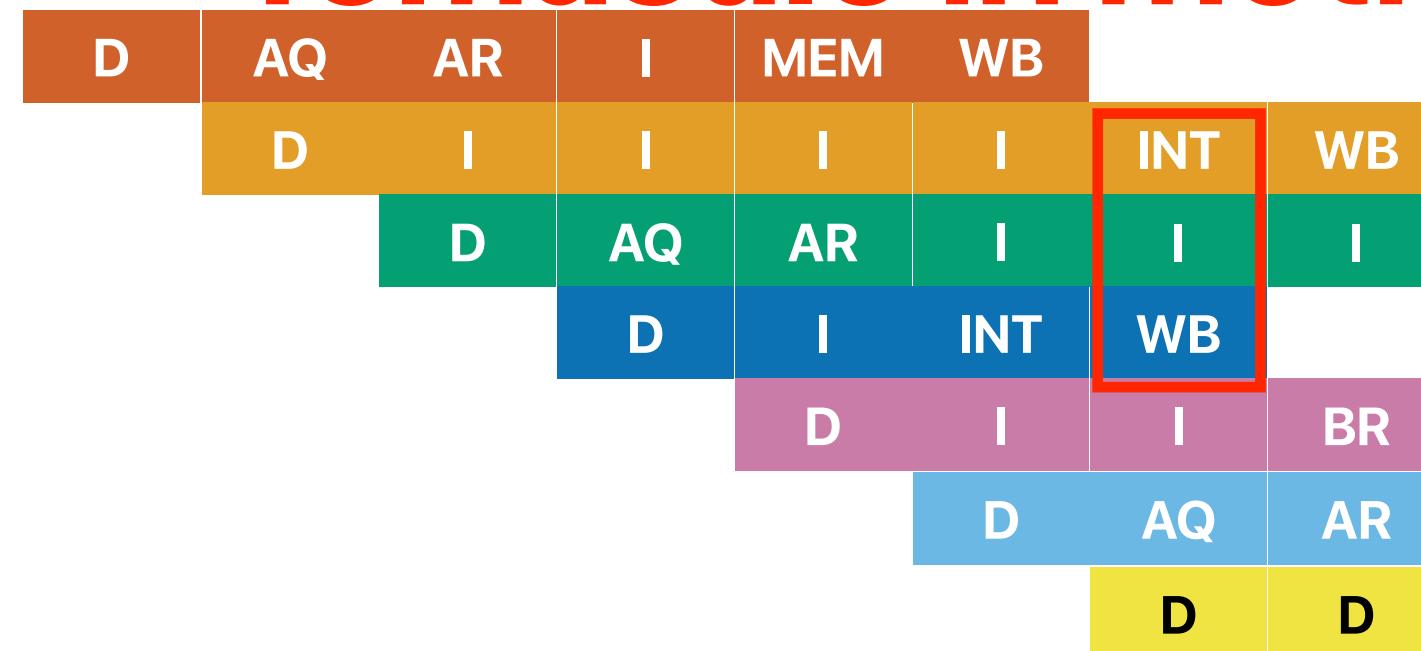


	INST	Vj	Vk	Vst	Qj	Qk	Qst	A	Inst #
LD1	ld	0	[X10]					1	
LD2	ld	0	ADD2					6	
LD3									
ST1	sd	0	[X10]				ADD1	3	
ST2									
ST3									
ADD1	add	LD1	[X12]					2	
ADD2	addi	8	[X10]					4	
MUL1									
MUL2									
BR	bne	ADD2	[X5]						5

RSV #	Value	Spec?
X5		
X6	LD2	1
X7	ADD1	
X10	ADD2	
X12		

Tomasulo in motion

- ① ld X6, 0(X10)
- ② add X7, X6, X12
- ③ sd X7, 0(X10)
- ④ addi X10, X10, 8
- ⑤ bne X10, X5, LOOP
- ⑥ ld X6, 0(X10)
- ⑦ add X7, X6, X12
- ⑧ sd X7, 0(X10)
- ⑨ addi X10, X10, 8
- ⑩ bne X10, X5, LOOP



addi is finished
ahead of **add** and **sd**!

	INST	Vj	Vk	Vst	Qj	Qk	Qst	A	Inst #
LD1	ld	0	[X10]					1	
LD2	ld	0	ADD2					6	
LD3									
ST1	sd	0	[X10]	ADD1				3	
ST2									
ST3									
ADD1	add	LD1	[X12]					2	
ADD2	add		[X12]		[LD2]			7	
MUL1									
MUL2									
BR	bne	ADD2	[X5]						

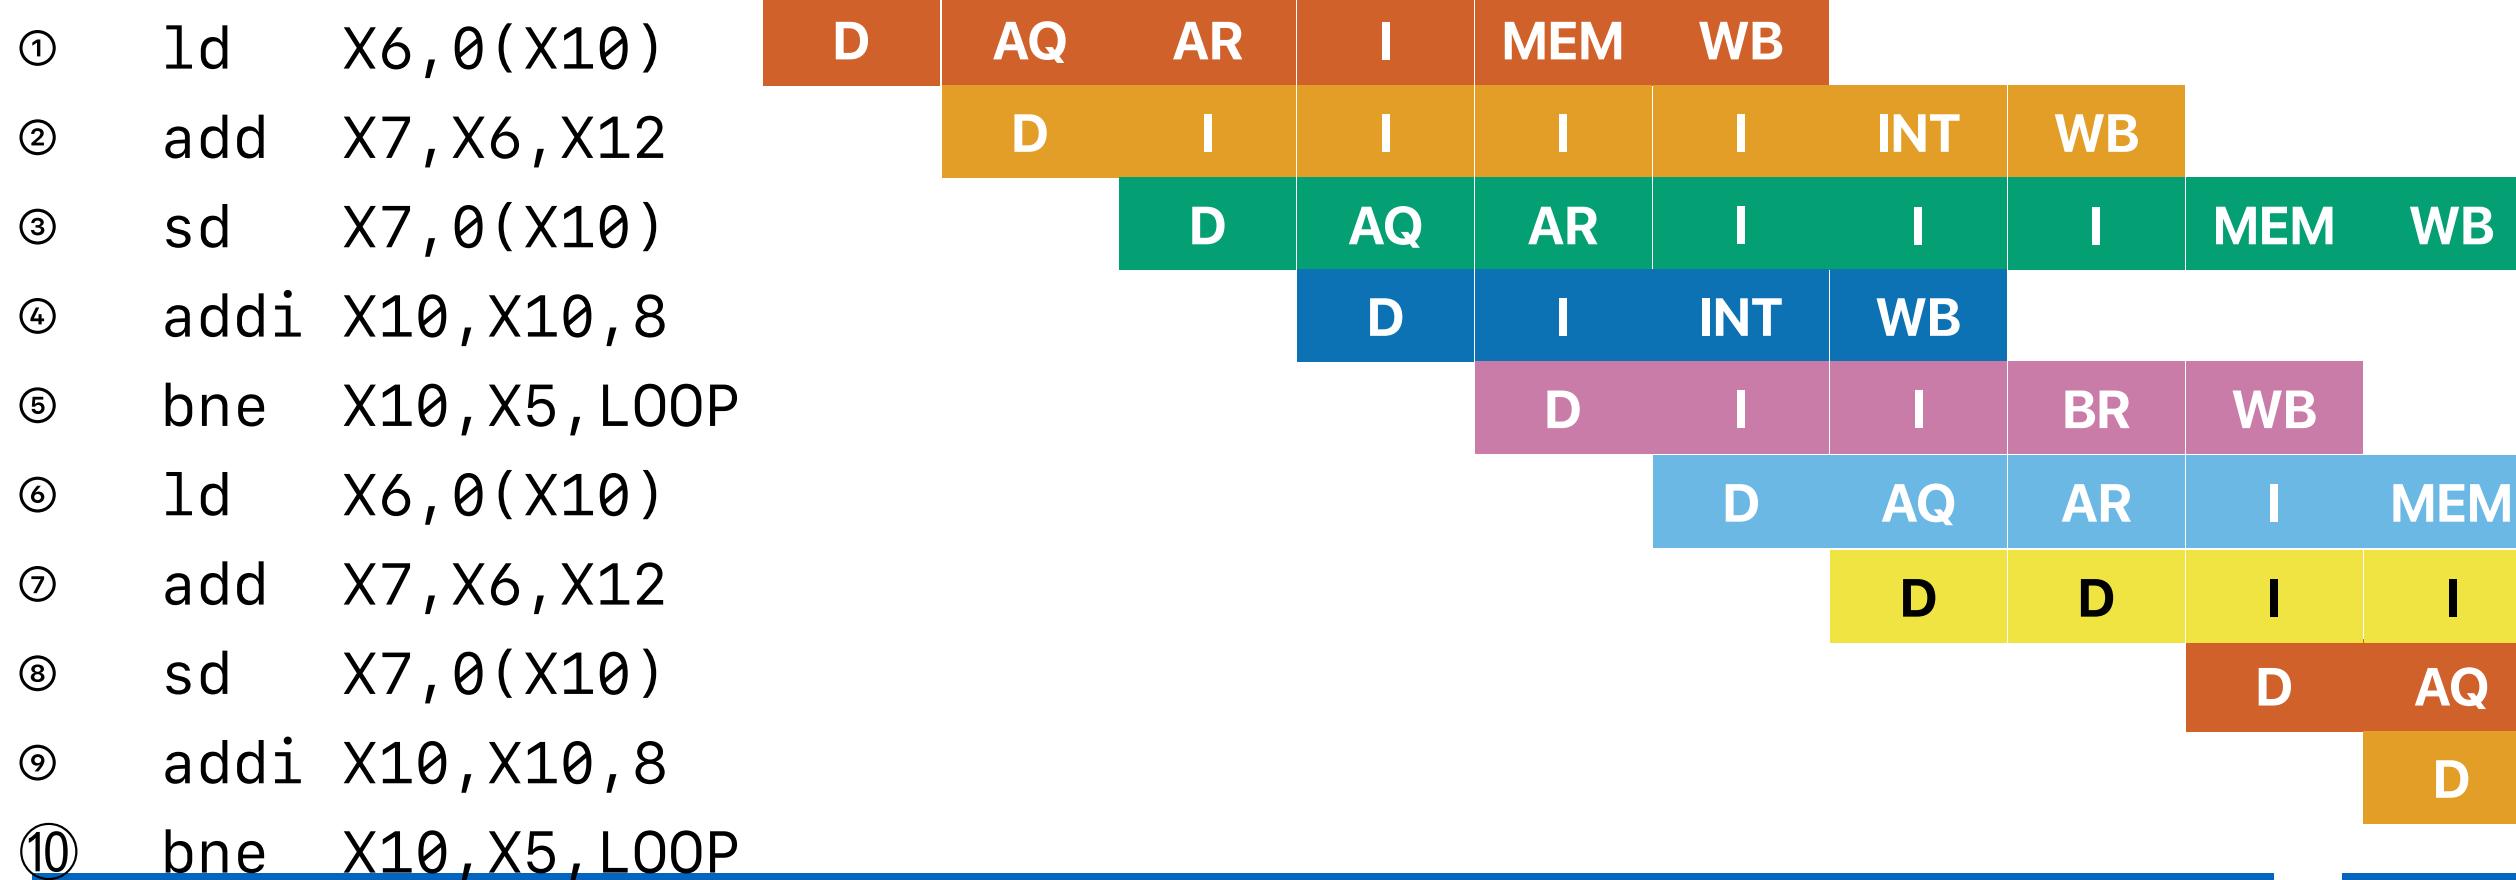
	RSV #	Value	Spec?
X5			
X6	LD2		1
X7	ADD2	ADD1	1
X10			
X12			

Tomasulo in motion

			D	AQ	AR	I	MEM	WB		
①	ld	X6, 0(X10)	D			I				
②	add	X7, X6, X12		D	I	I	I	INT	WB	
③	sd	X7, 0(X10)		D	AQ	AR	I	I	I	MEM
④	addi	X10, X10, 8			D	I	INT	WB		
⑤	bne	X10, X5, LOOP				D	I	I	BR	WB
⑥	ld	X6, 0(X10)				D	AQ	AR	I	
⑦	add	X7, X6, X12				D	D	D	I	
⑧	sd	X7, 0(X10)							D	
⑨	addi	X10, X10, 8								
⑩	bne	X10, X5, LOOP								

	INST	Vj	Vk	Vst	Qj	Qk	Qst	A	Inst #	RSV #	Value	Spec?
LD1	ld	0	[X10]						1	X5		
LD2	ld	0	ADD2						6	X6	LD2	
LD3										X7	ADD2	
ST1	sd	0	[X10]	ADD1					3	X10		ADD2
ST2	sd	0	[X10]				ADD2		8	X12		
ST3												
ADD1	add	LD1	[X12]						2			
ADD2	add		[X12]		[LD2]				7			
MUL1												
MUL2												
BR	bne	ADD2	[X5]						5			

Tomasulo in motion



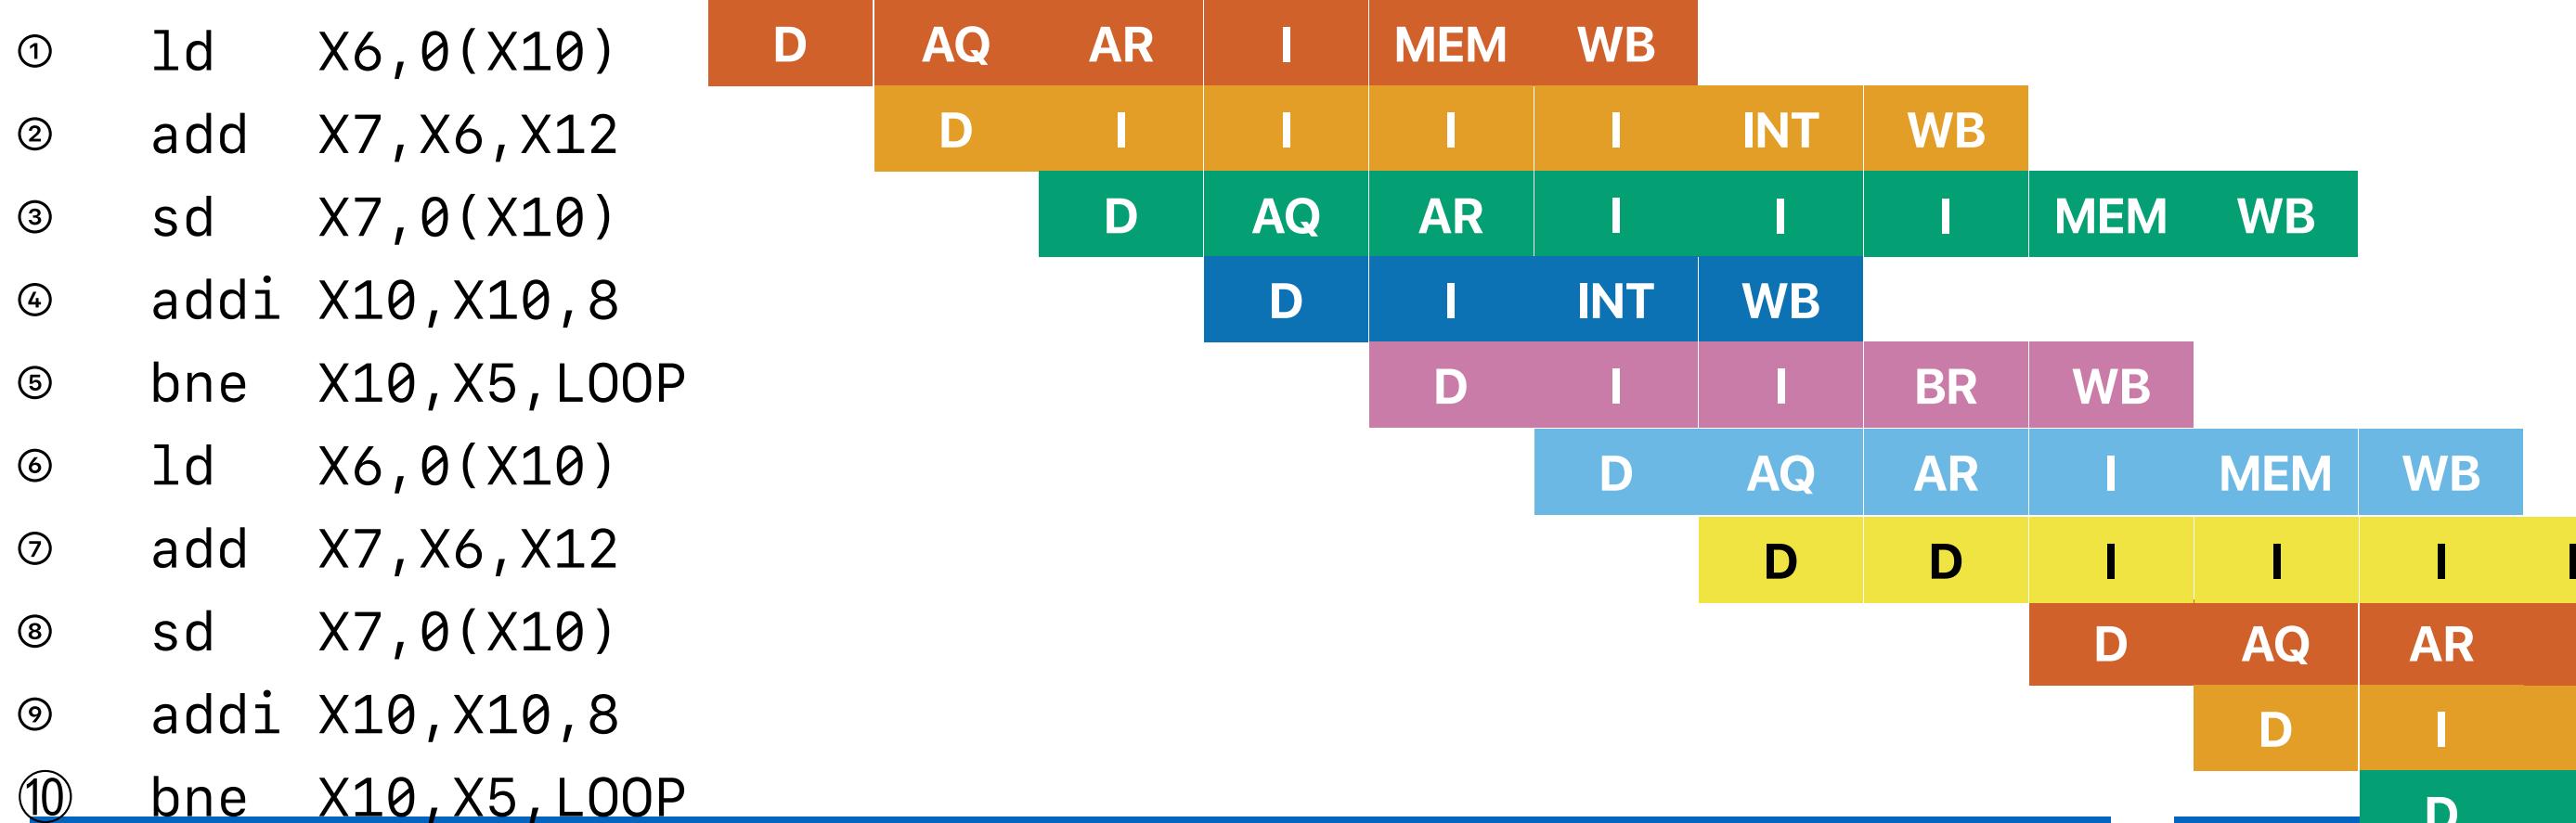
	INST	Vj	Vk	Vst	Qj	Qk	Qst	A	Inst #	RSV #	Value	Spec?
LD1	ld	0	[X10]						1	X5		
LD2	ld	0	ADD2						6	X6	LD2	
LD3										X7	ADD2	
ST1	sd	0	[X10]	ADD1					3	X10	ADD1	
ST2	sd	0	[X10]				ADD2		8	X12		
ST3												
ADD1	add	8	ADD2						9			
ADD2	add		[X12]		[LD2]				7			
MUL1												
MUL2												
BR	bne	ADD2	[X5]						5			

Tomasulo in motion

			D	AQ	AR	I	MEM	WB				
①	ld	X6, 0(X10)	D	AQ	AR	I	MEM	WB				
②	add	X7, X6, X12	D	I	I	I	I	INT	WB			
③	sd	X7, 0(X10)	D	AQ	AR	I	I	I	I	MEM	WB	
④	addi	X10, X10, 8	D	I	INT	WB						
⑤	bne	X10, X5, LOOP	D	I	I	BR	WB					
⑥	ld	X6, 0(X10)	D	AQ	AR	I	MEM	WB				
⑦	add	X7, X6, X12	D	D	I	I	I	I				
⑧	sd	X7, 0(X10)	D	AQ	AR							
⑨	addi	X10, X10, 8	D	I								
⑩	bne	X10, X5, LOOP	D									

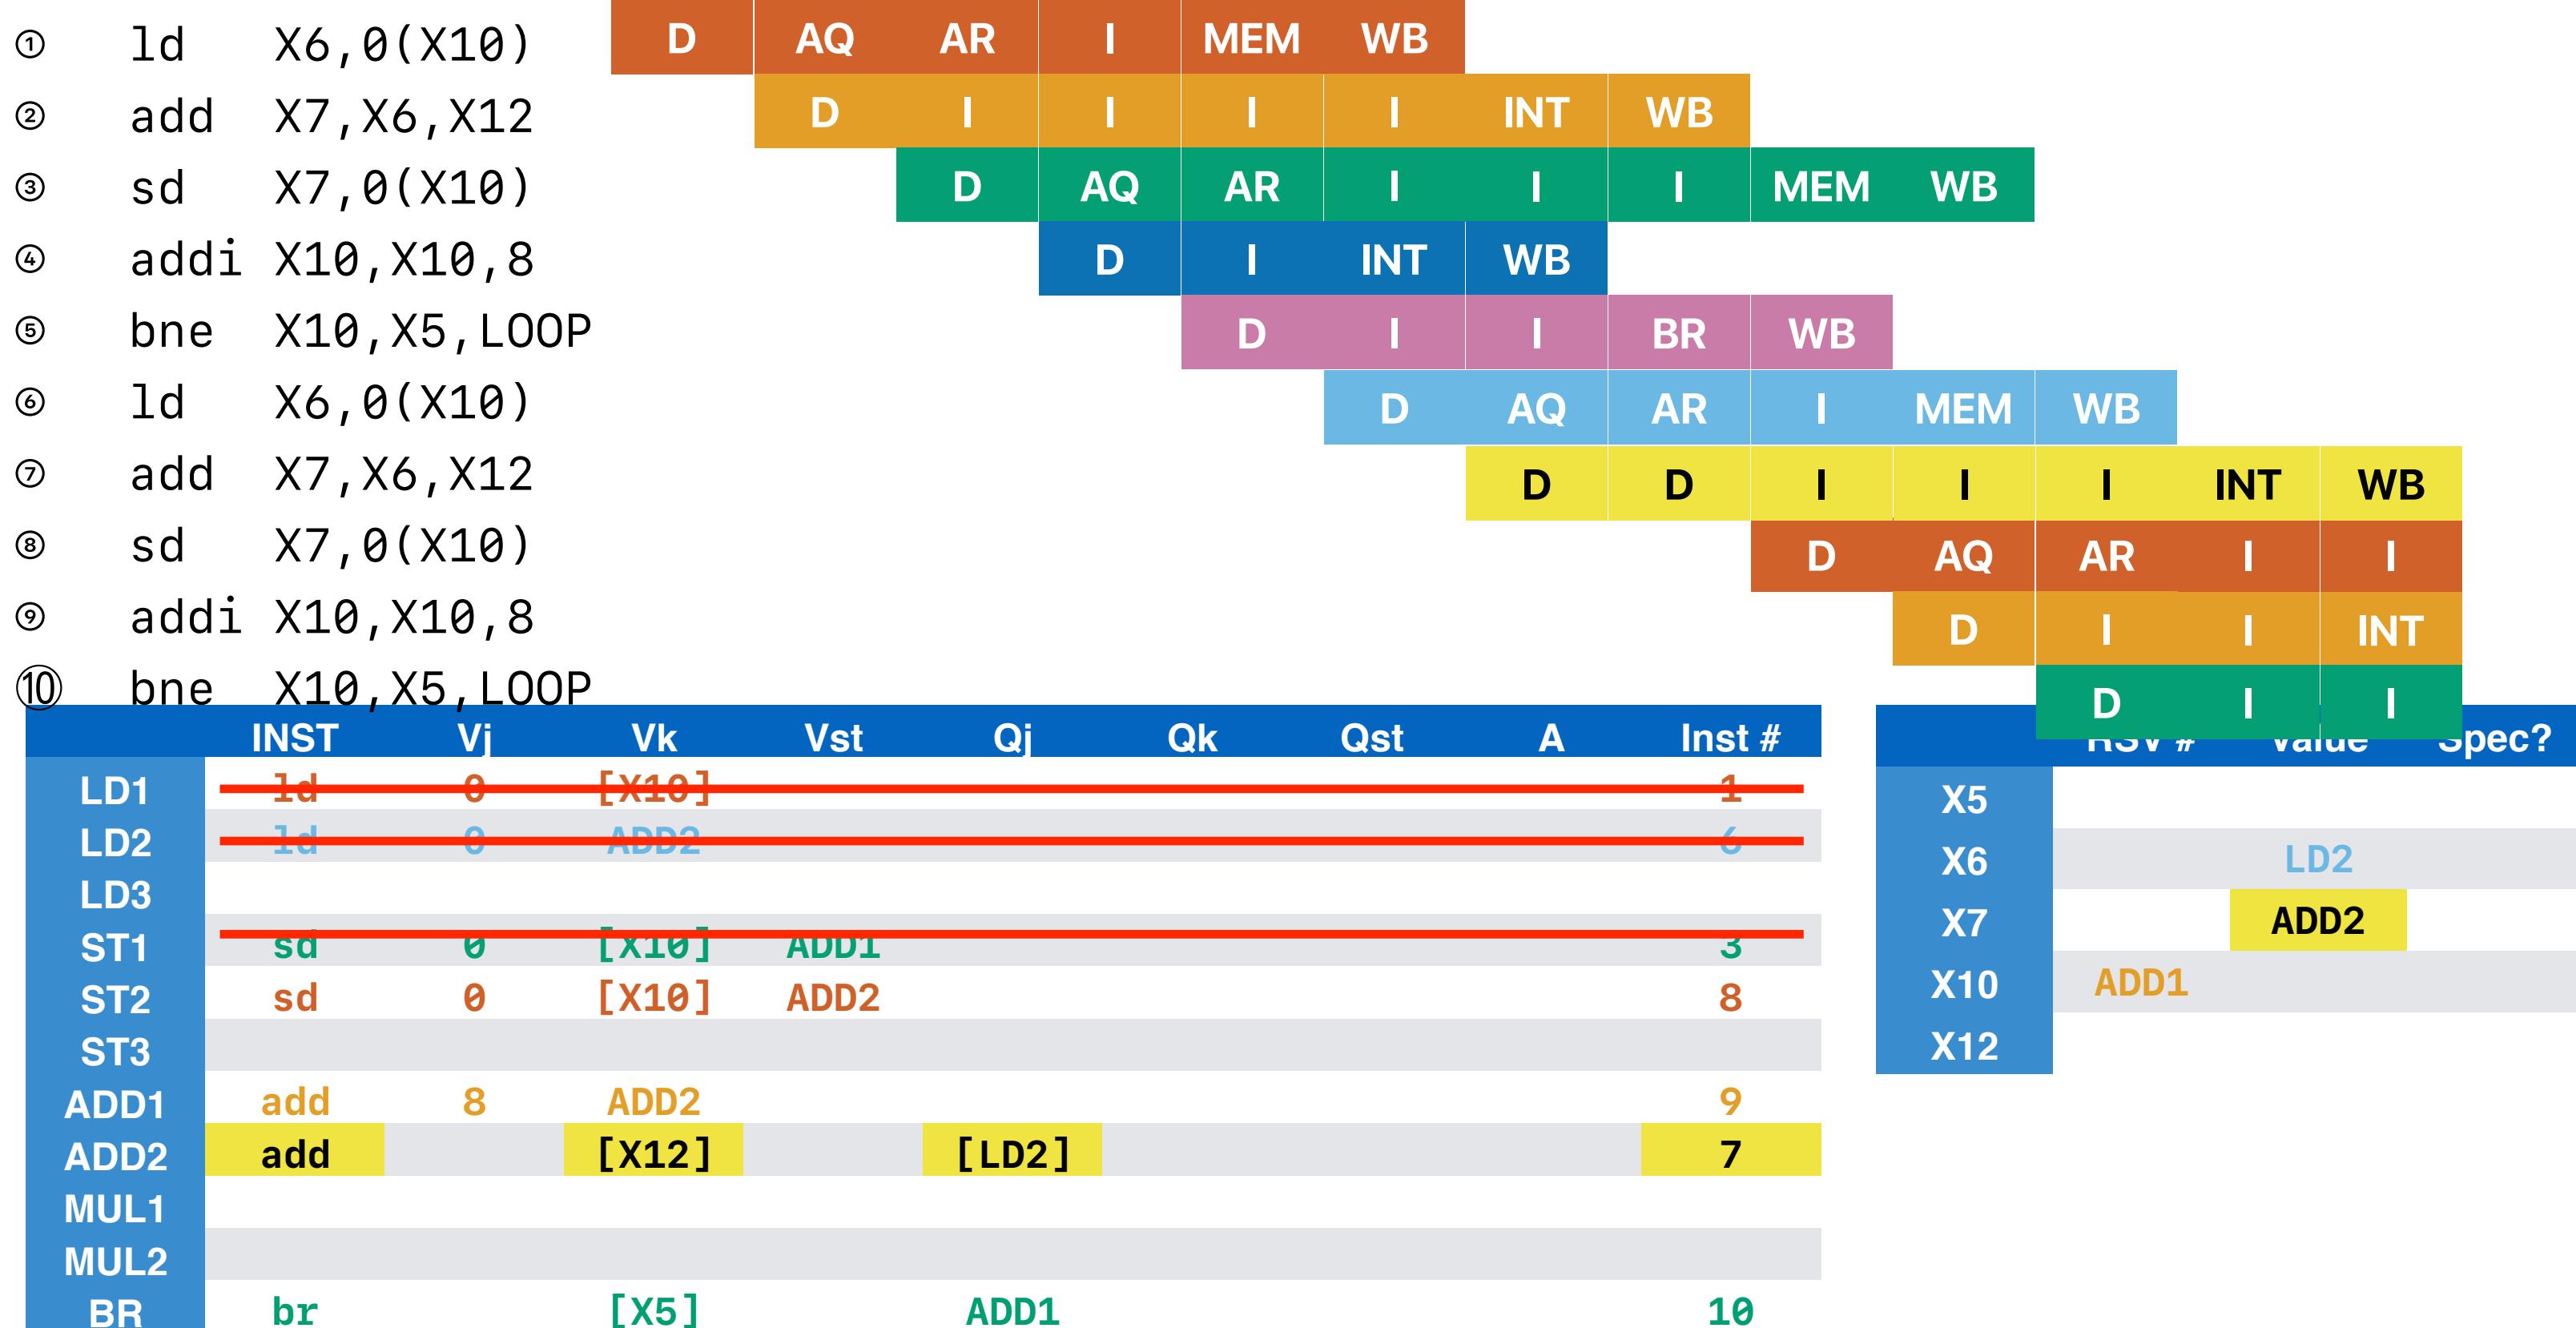
	INST	Vj	Vk	Vst	Qj	Qk	Qst	A	Inst #	Resv #	Value	Spec?
LD1	ld	0	[X10]						1		X5	
LD2	ld	0	ADD2						6		X6	LD2
LD3											X7	ADD2
ST1	sd	0	[X10]	ADD1					3		X10	ADD1
ST2	sd	0	[X10]				ADD2		8		X12	
ST3												
ADD1	add	8	ADD2						9			
ADD2	add		[X12]		[LD2]				7			
MUL1												
MUL2												
BR	br		[X5]		ADD1				10			

Tomasulo in motion

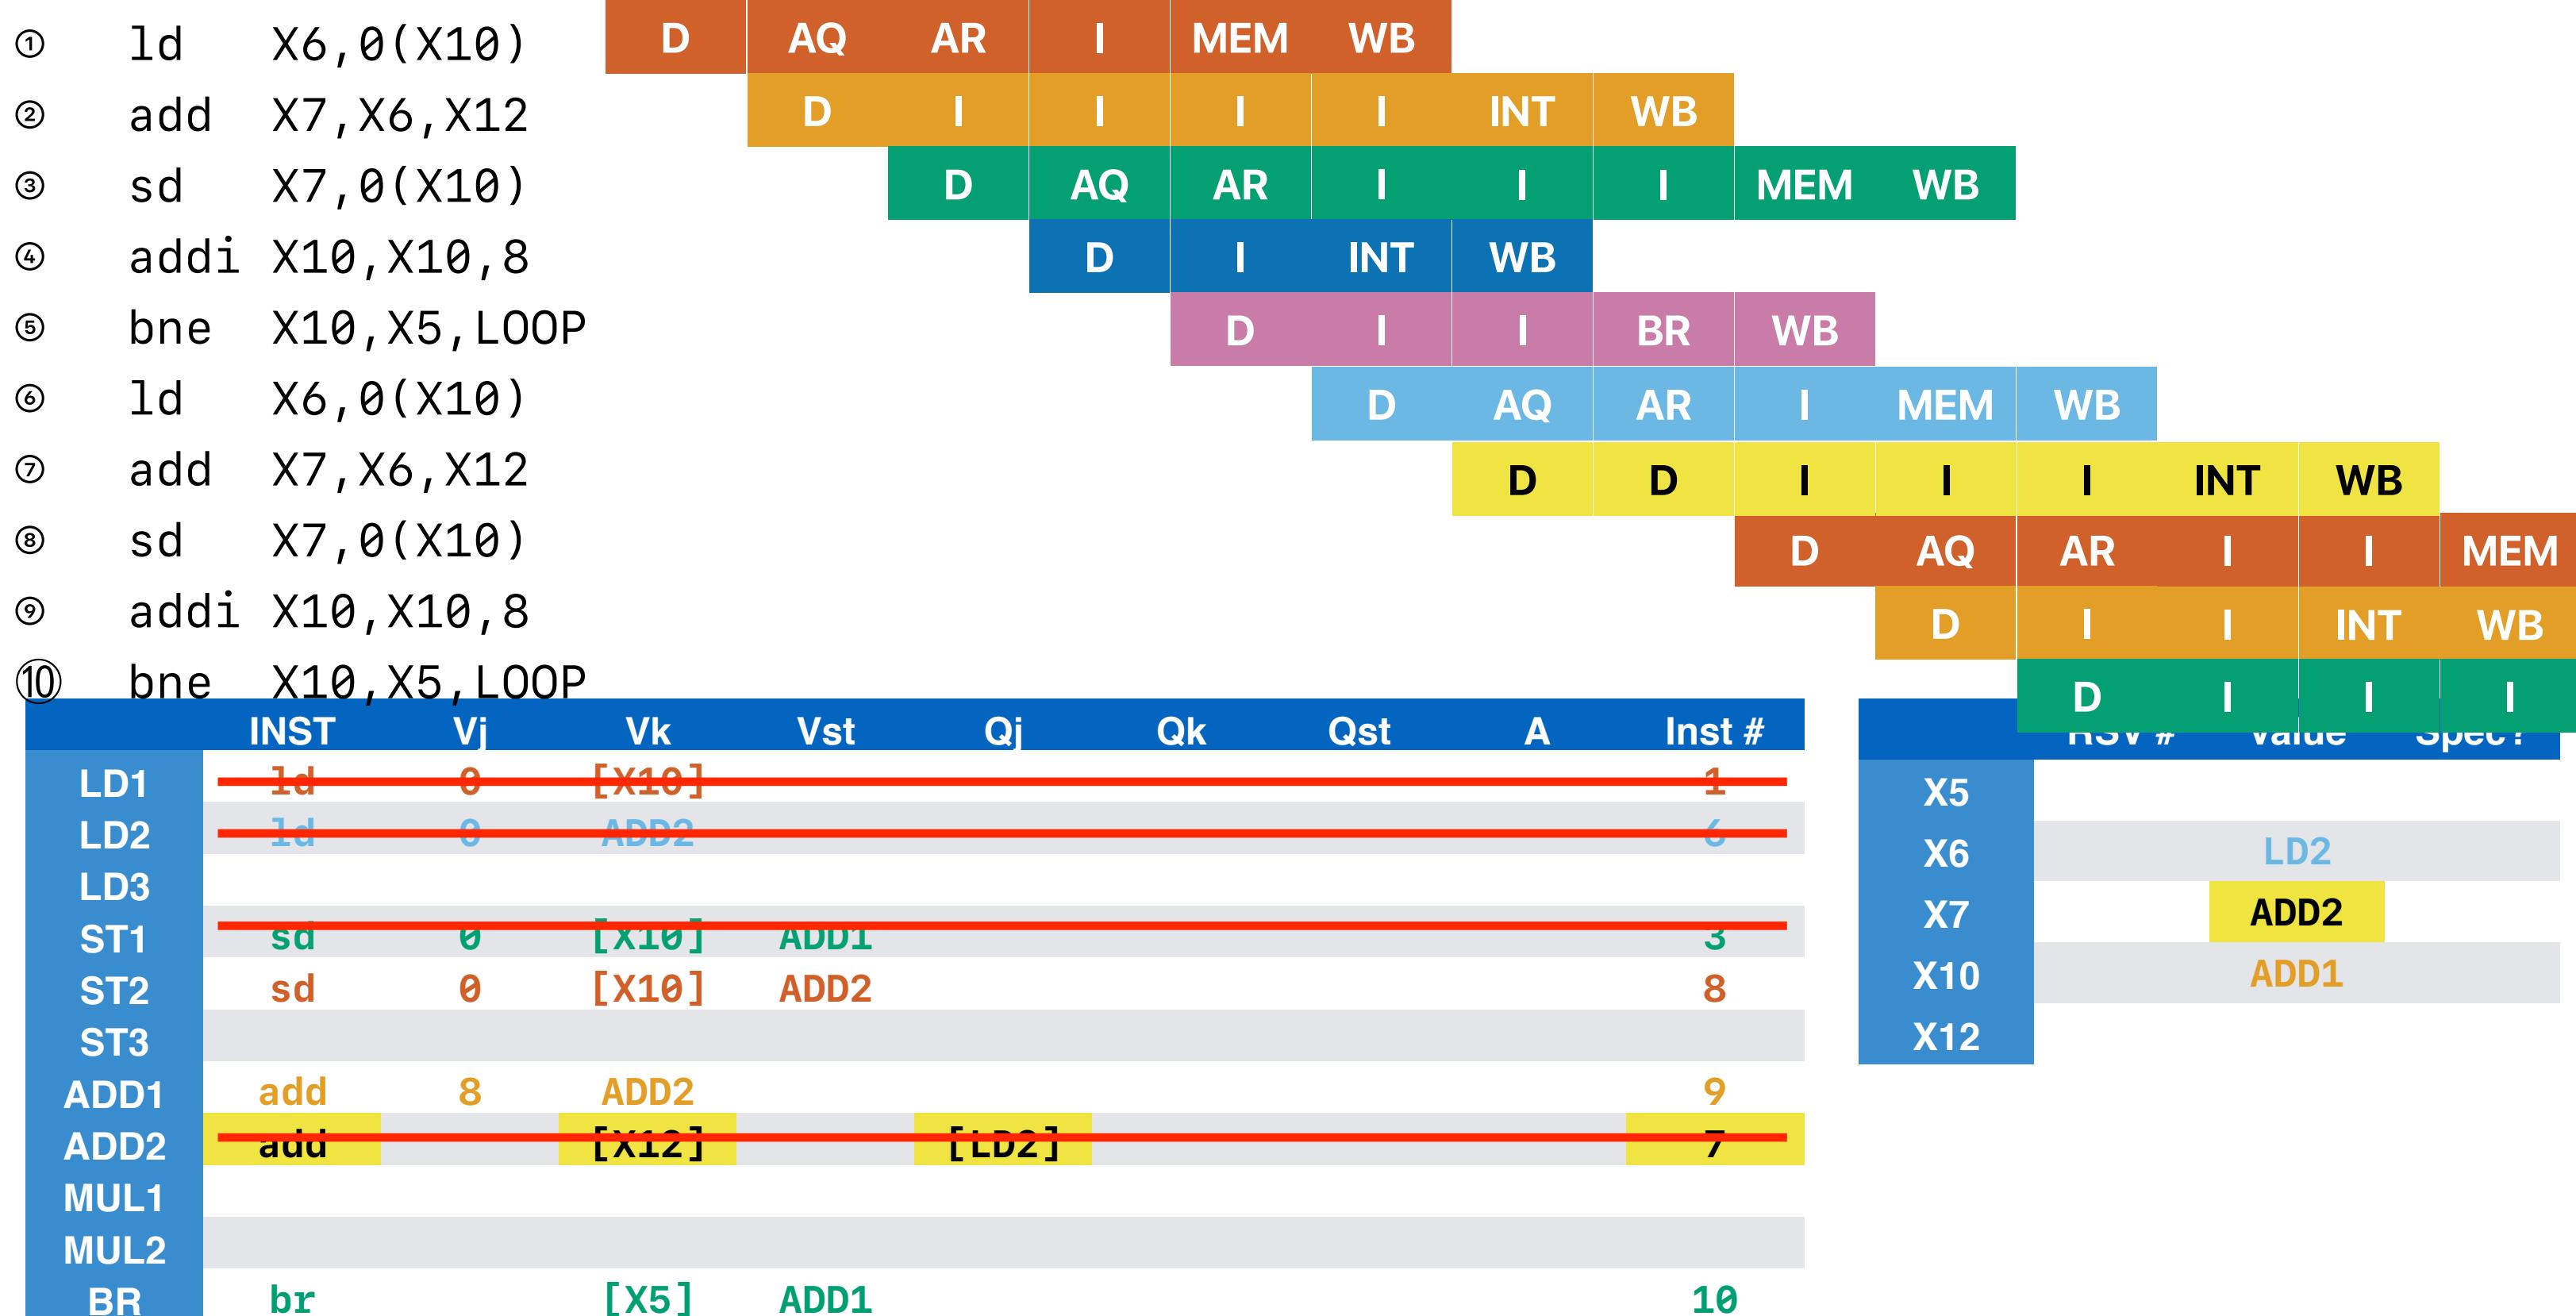


	INST	Vj	Vk	Vst	Qj	Qk	Qst	A	Inst #	Resv	Value	Spec?
LD1	ld	0	[X10]						1		X5	
LD2	ld	0	ADD2						6		X6	LD2
LD3											X7	ADD2
ST1	sd	0	[X10]	ADD1					3		X10	ADD1
ST2	sd	0	[X10]				ADD2		8		X12	
ST3												
ADD1	add	8	ADD2						9			
ADD2	add		[X12]		[LD2]				7			
MUL1												
MUL2												
BR	br		[X5]		ADD1				10			

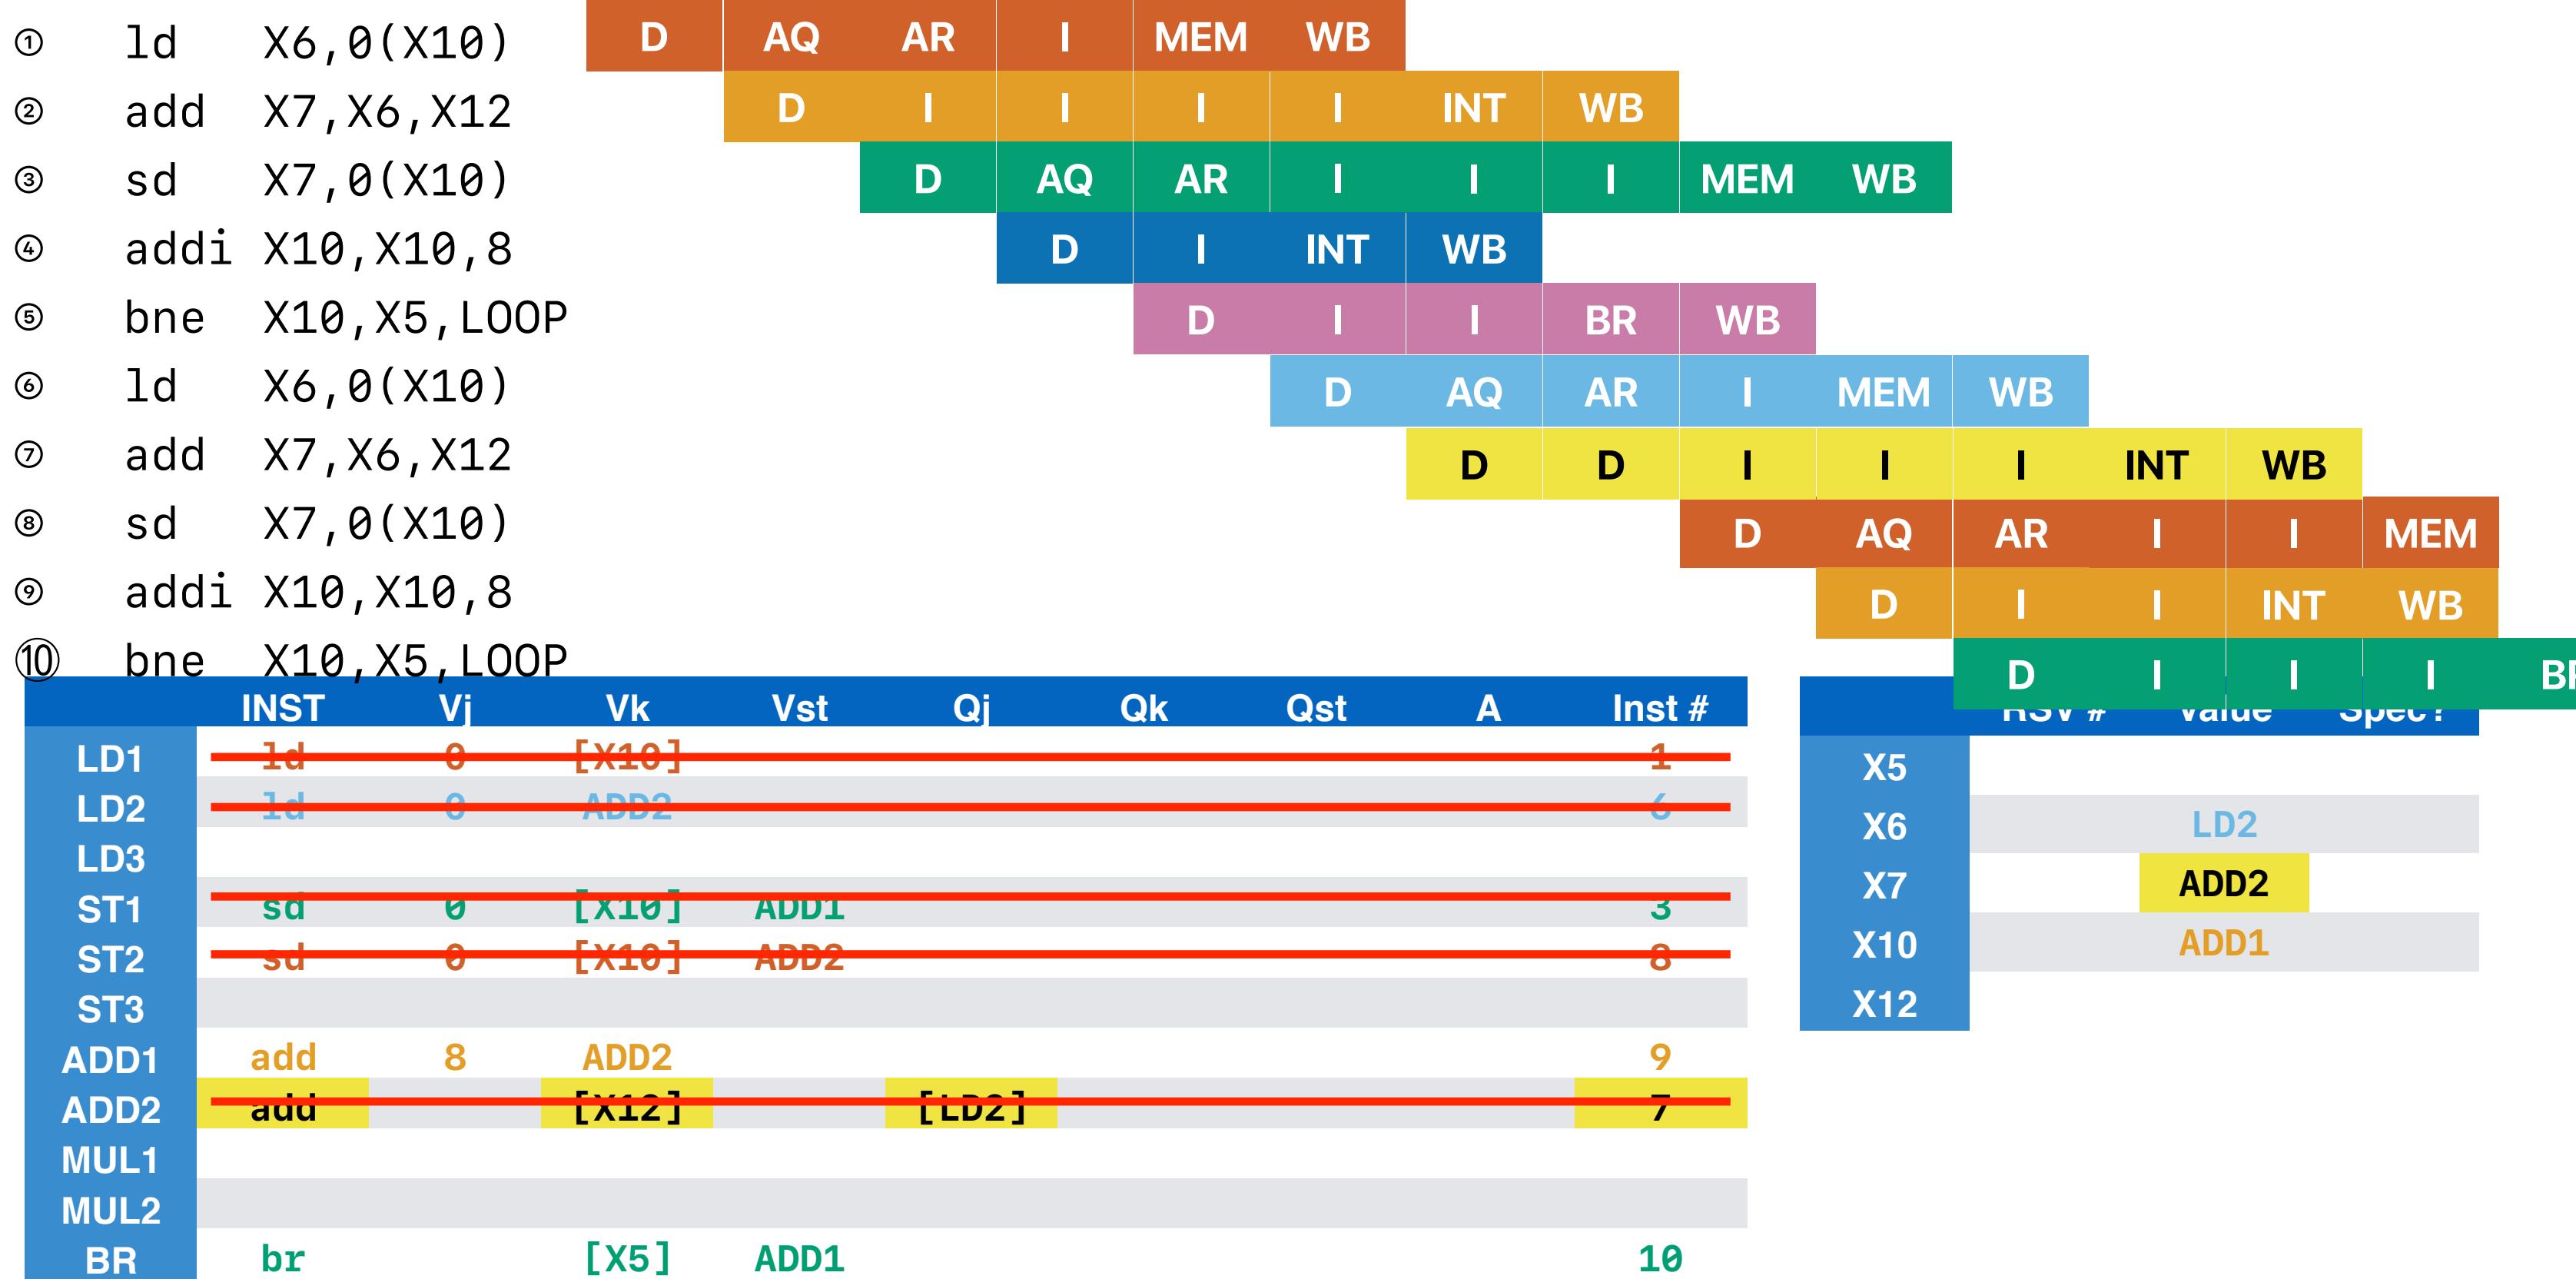
Tomasulo in motion



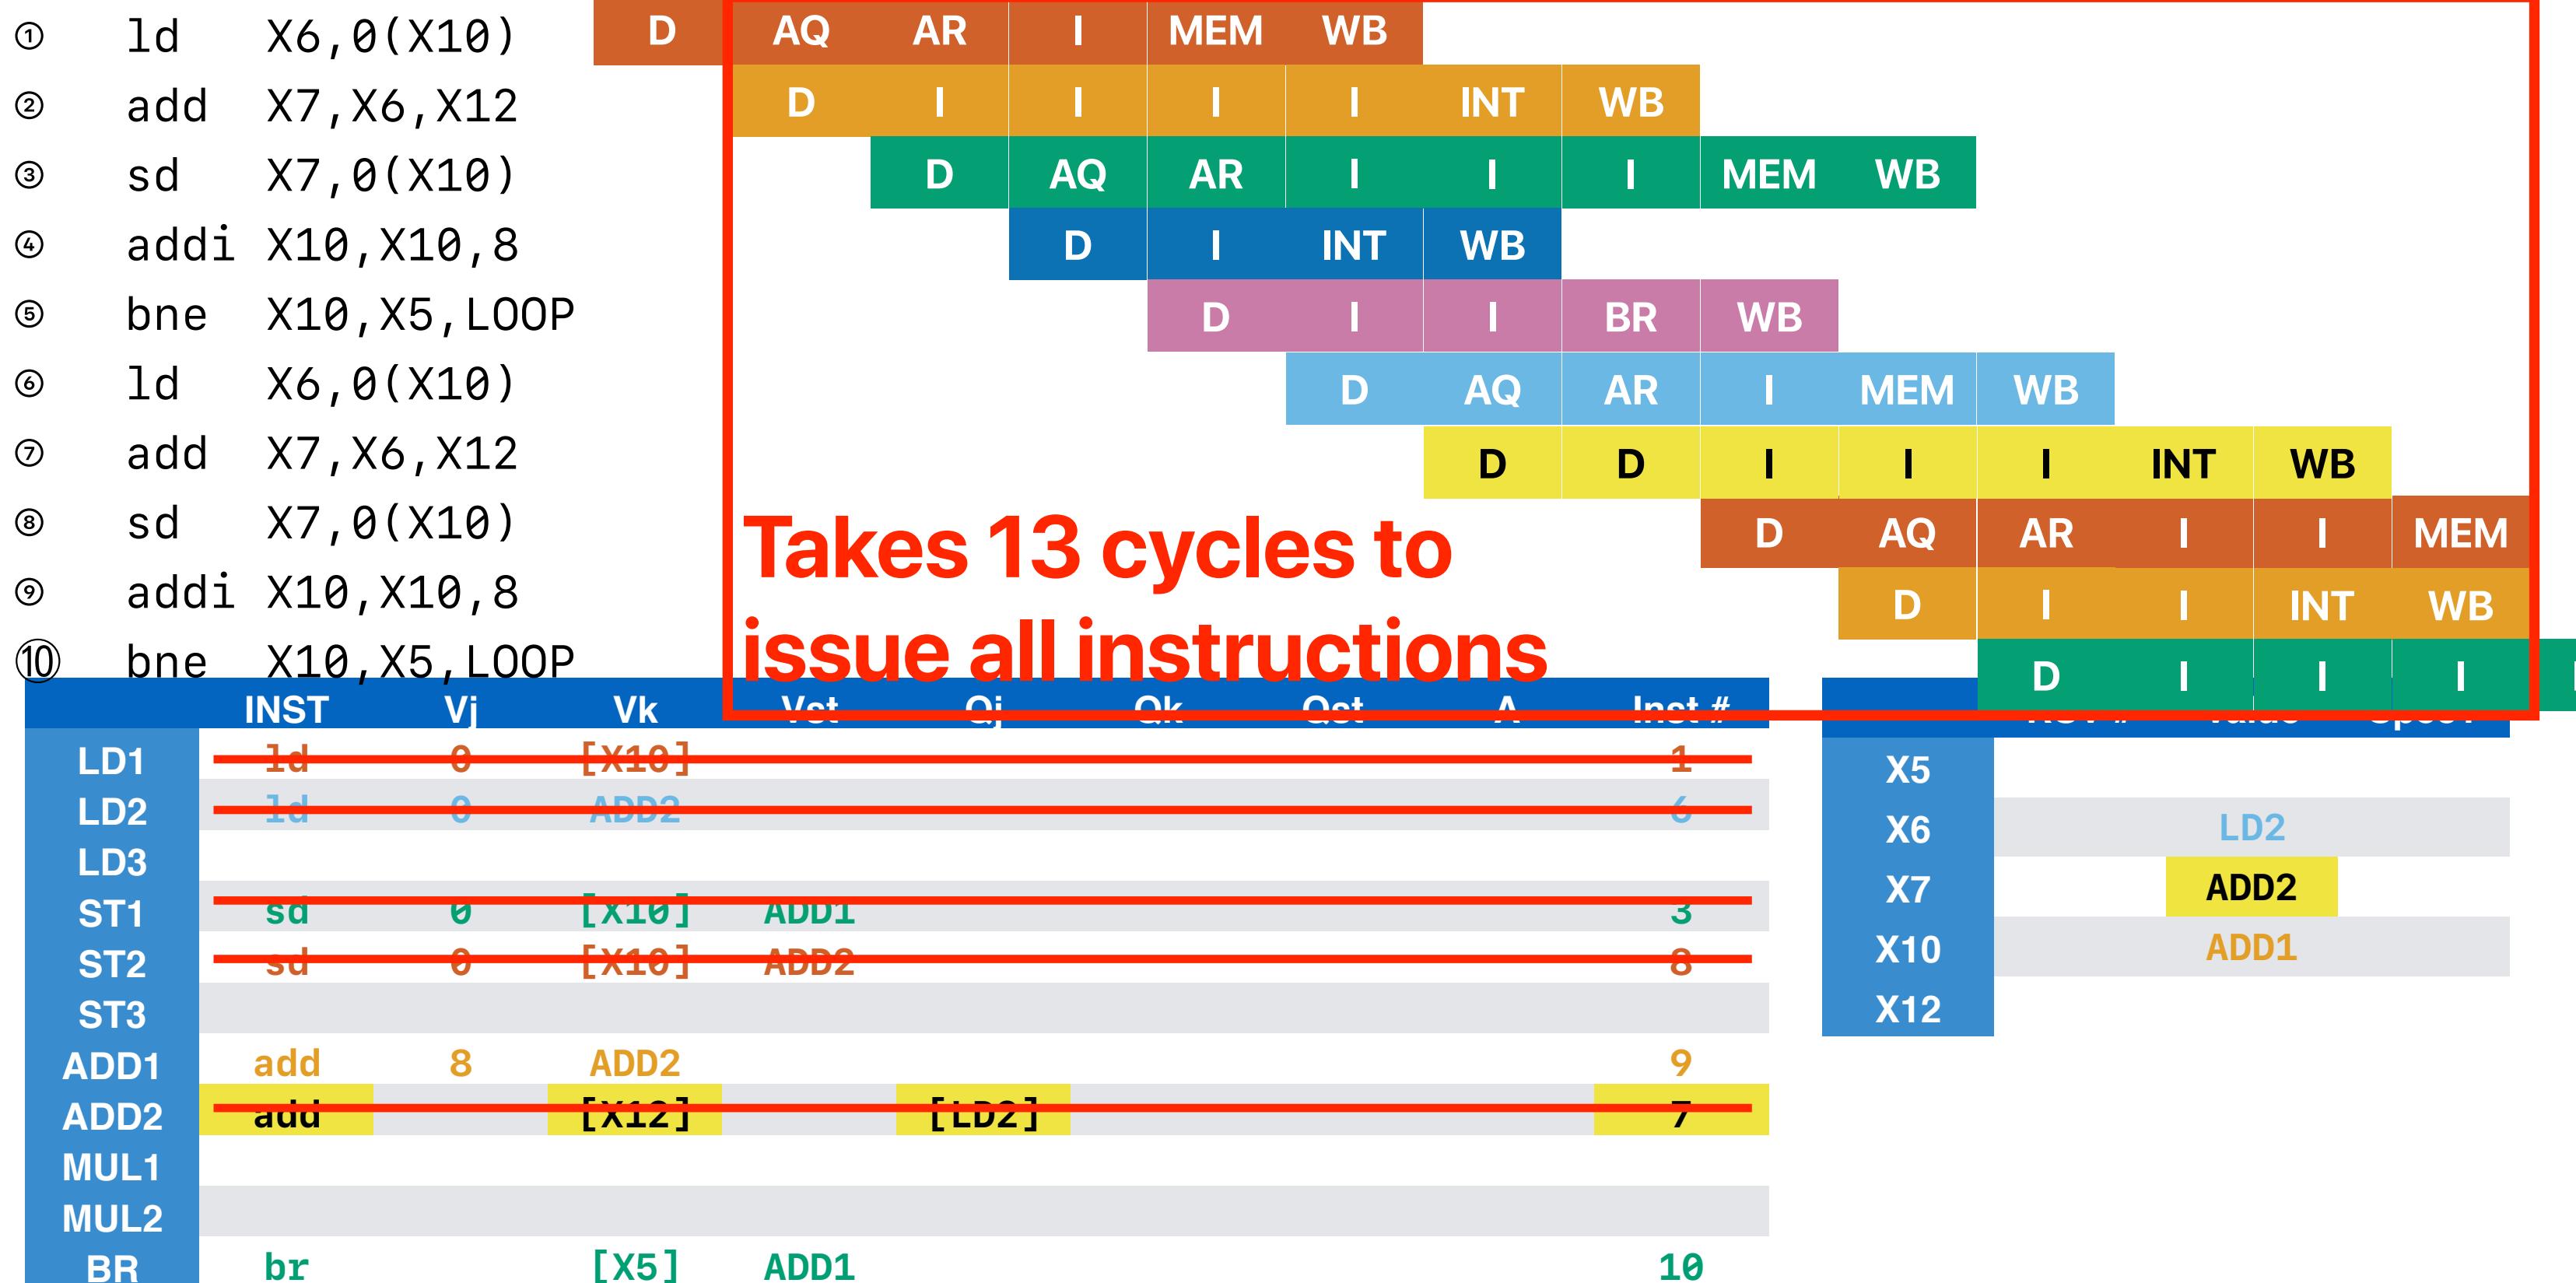
Tomasulo in motion



Tomasulo in motion



Tomasulo in motion



Announcement

- Project is up — check the website
- Assignment #4 is up — start EARLY!!!
- Office Hours on Zoom (the office hour link, not the lecture one)
 - Hung-Wei/Prof. Usagi: M 8p-9p, W 2p-3p
 - Quan Fan: F 1p-3p
- Regarding projected grades
 - Based on your “weighted total” column in iLearn — we only have 50% offered so far, that’s why the max is only 48 now
 - Our final grading is based on “relative ranking” and scale may change

Computer Science & Engineering

203

つづく

