

Poll Q: What affects throughput?

Peak throughput depends on...

	Single Cycle	Multi-Cycle	Pipeline
A	Longest Instruction	Cycle Time	Average Instruction
B	Longest Instruction	Cycle Time	Longest Instruction
C	Longest Instruction	Average Instruction	Cycle Time
D	Average Instruction	Longest Instruction	Cycle Time
E	<i>None of the above</i>		

Poll Q: What affects throughput?

Peak throughput depends on...

	Single Cycle	Multi-Cycle	Pipeline
C	Longest Instruction	Average Instruction	Cycle Time

Throughput is useful work over time – one measure: insts / sec

$$ET = \text{Inst} * \frac{CPI * CT}{\text{inst}}$$

Handwritten annotations: 'S' above Inst, 'inst' above CPI, and an arrow pointing to the fraction with 'S' above and 'inst' below.

→ Single Cycle: $ET = \text{Inst} * 1 * \text{BIG}$


Multi Cycle: $ET = \text{Inst} * [3..5] * CT$

→ Pipeline: $ET = \text{Inst} * 5 * CT$

Handwritten annotations: 'BIG' circled in the first equation; '[3..5]' circled in the second; '5' written below the third, with an arrow pointing to 'CT'.

Pipelining in Modern CPUs

- CPU Datapath
- Arithmetic Units — *pipelined multiplier*
- System Buses
- Software (at multiple levels)
- etc...

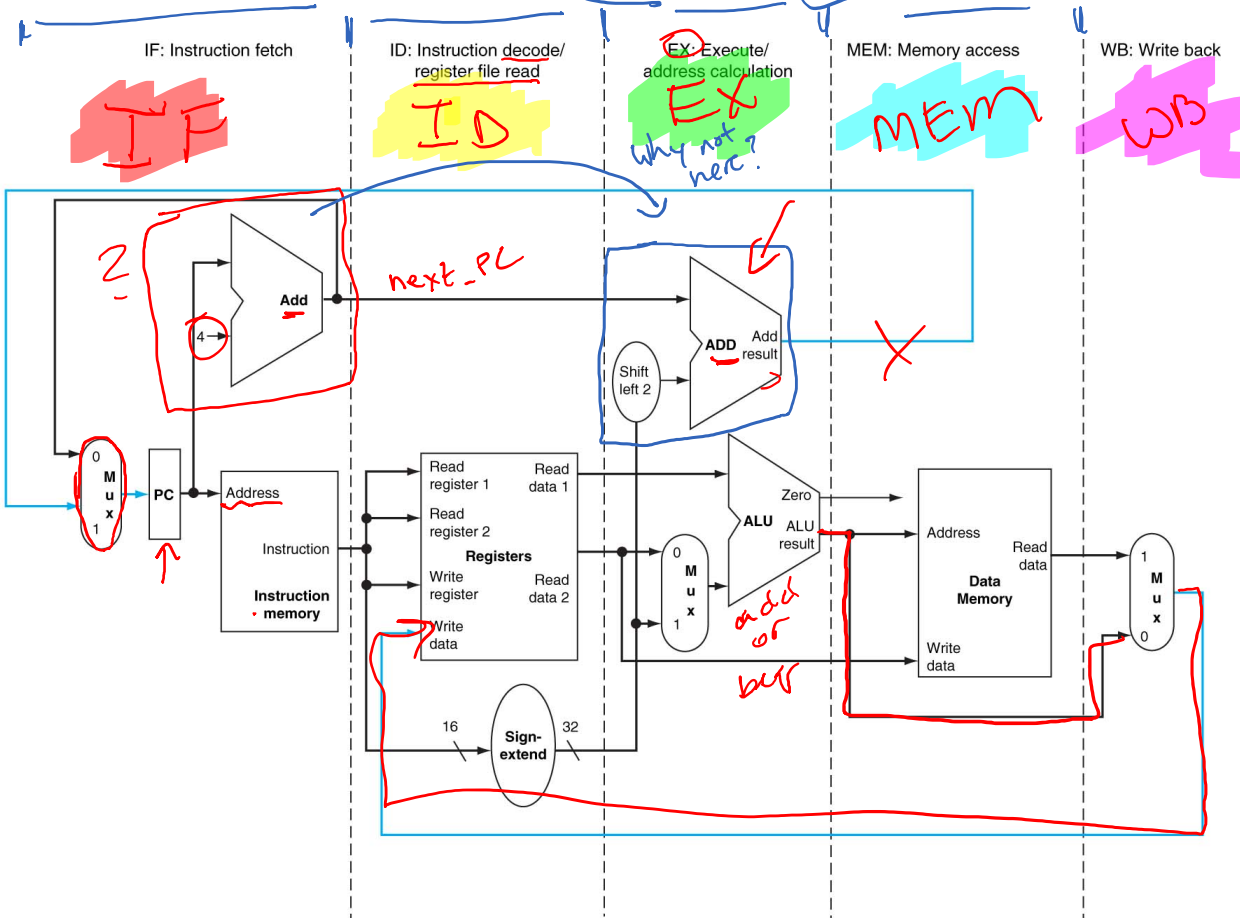
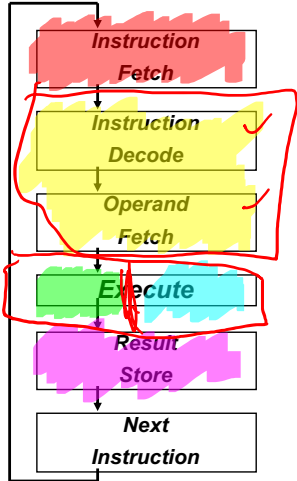


A Pipelined Datapath

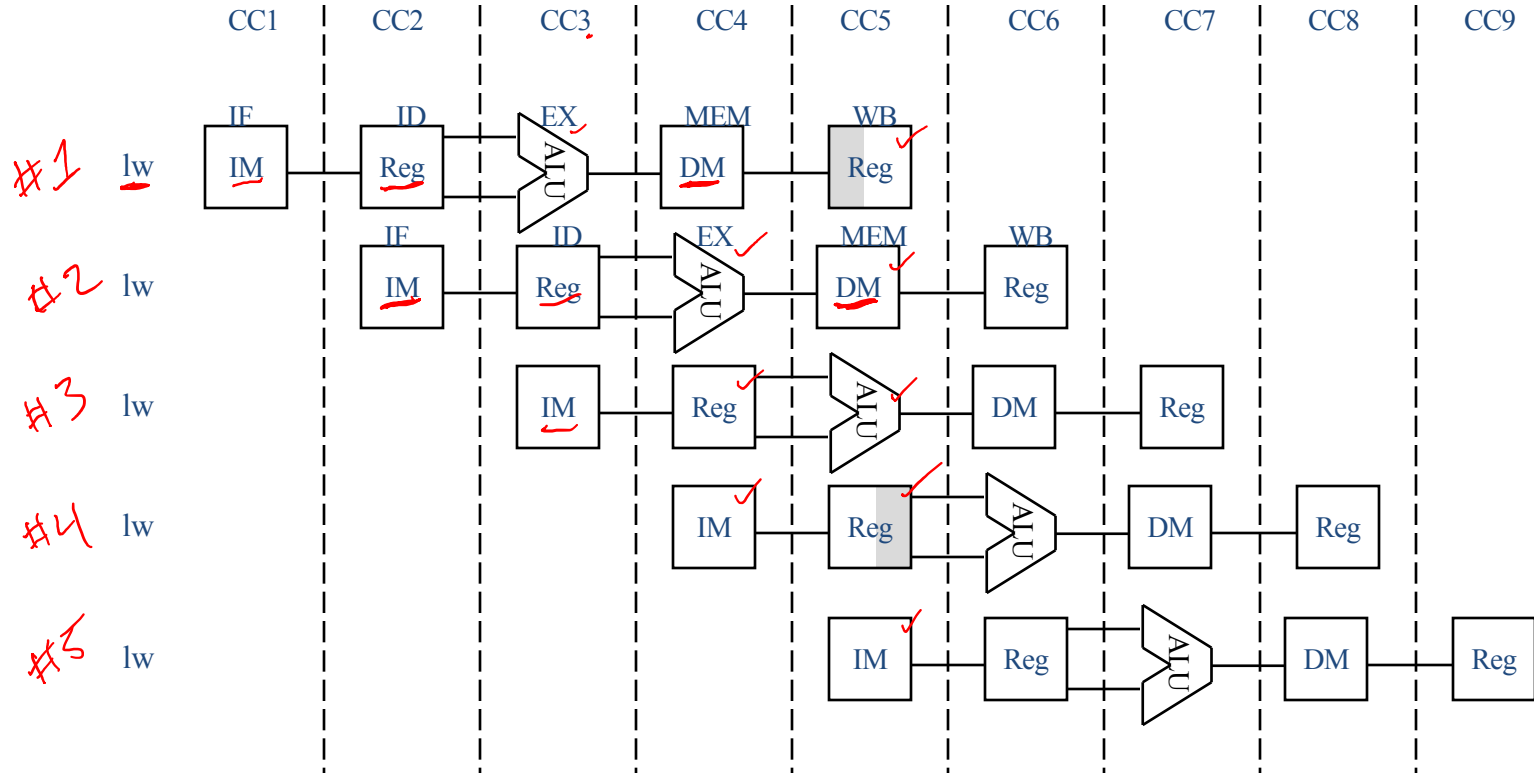
IF	Instruction fetch
ID	Instruction decode and register fetch
EX	Execution and effective address calculation
MEM	Memory access
WB	Write back

Pipelined Datapath (roughly)

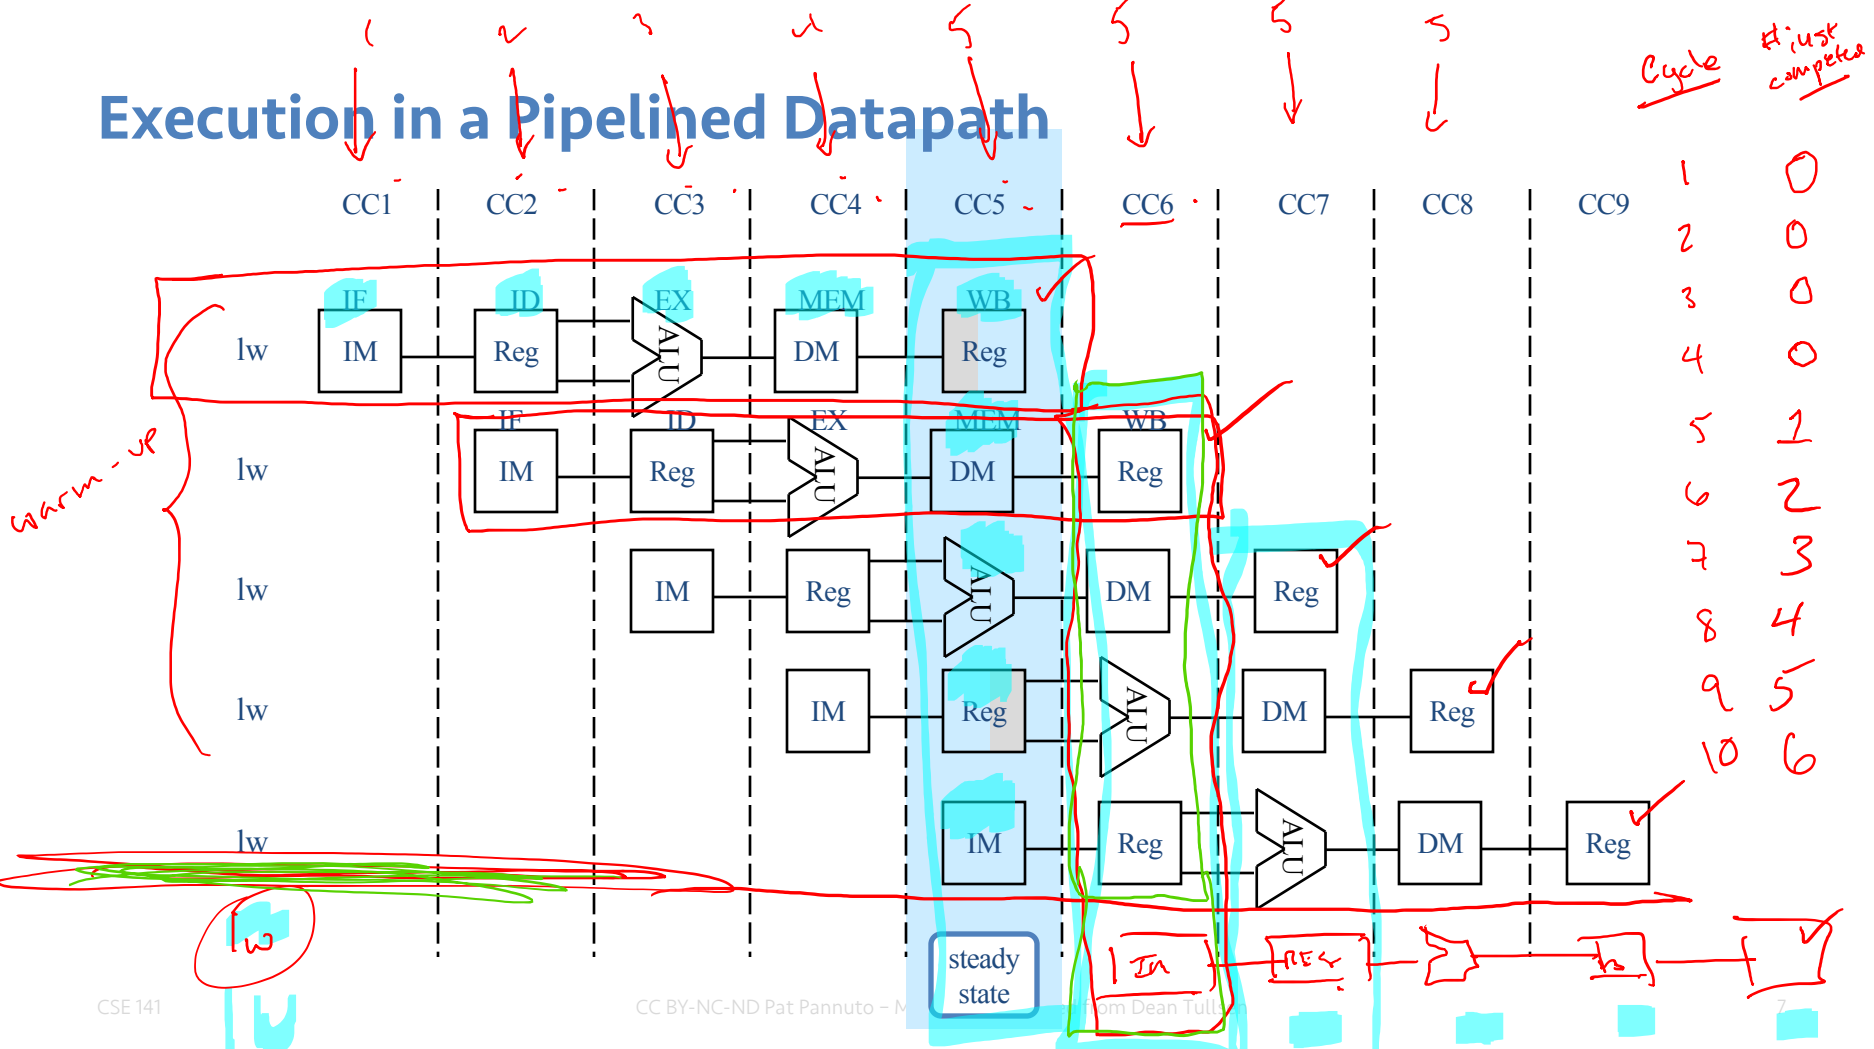
Conceptual model




Execution in a Pipelined Datapath



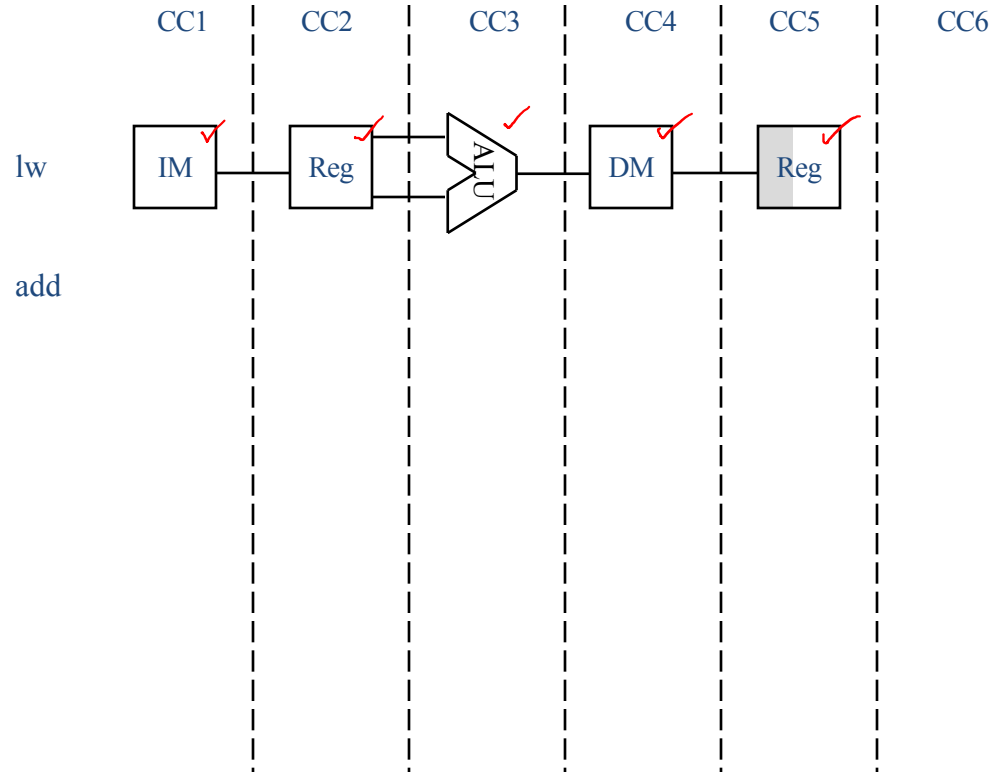
Execution in a Pipelined Datapath



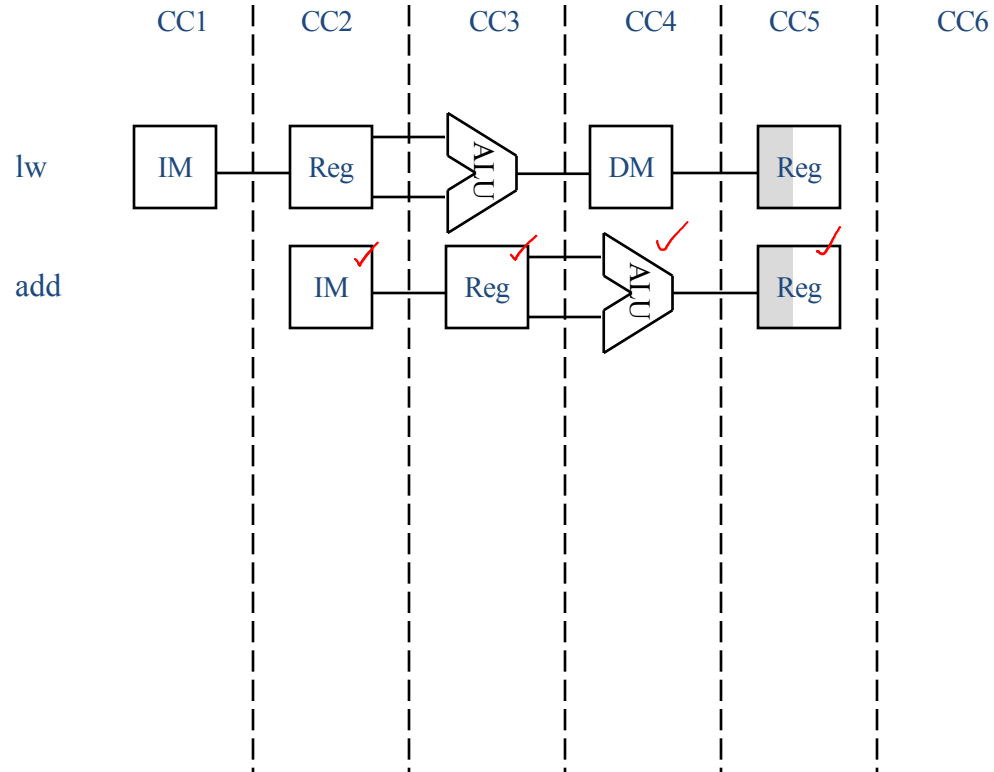
Mixed Instructions in the Pipeline



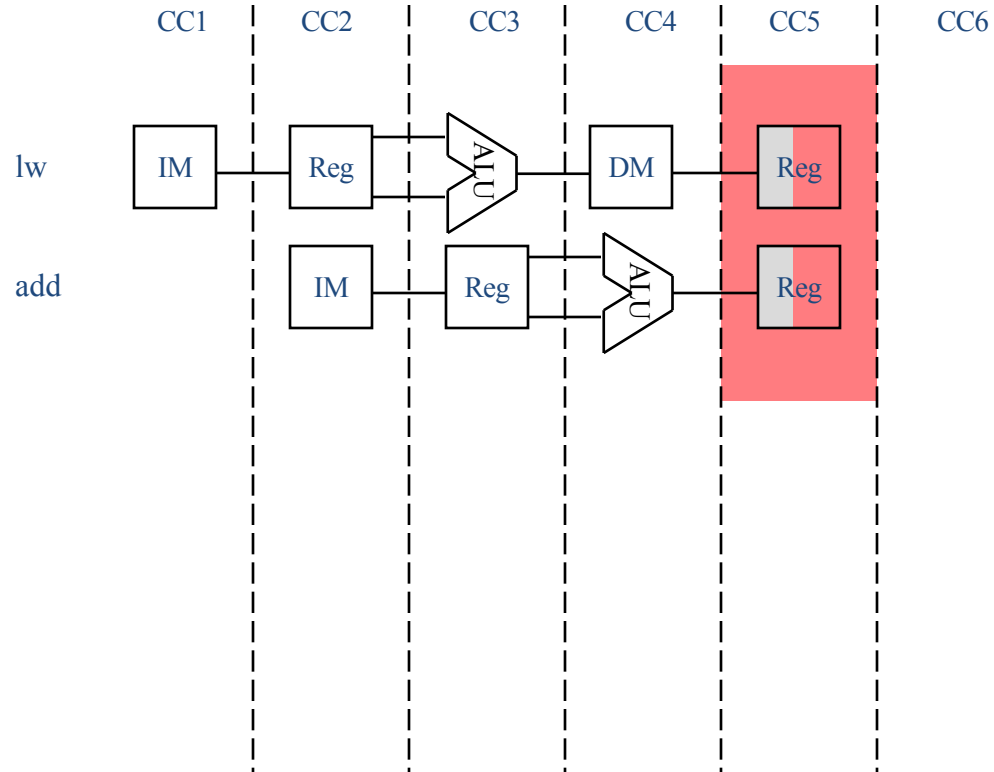
Mixed Instructions in the Pipeline



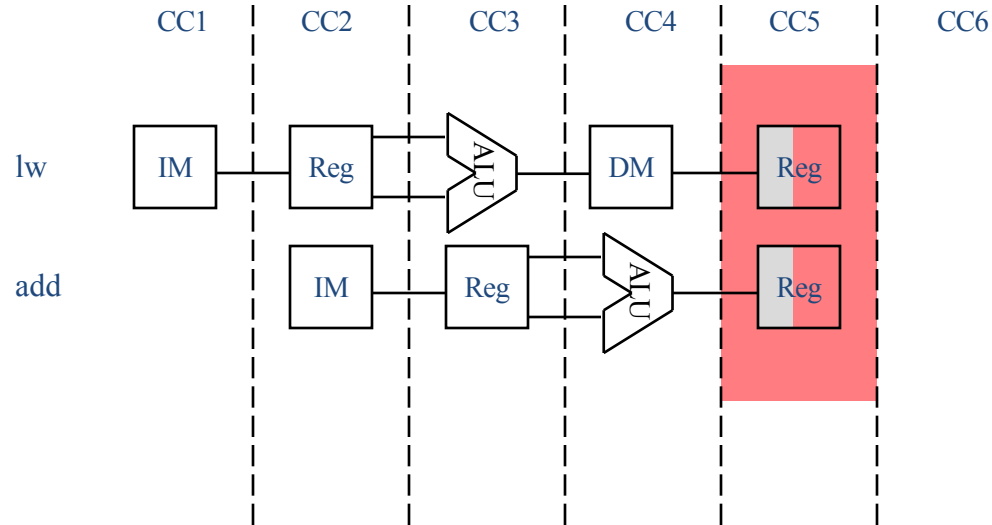
Mixed Instructions in the Pipeline



Mixed Instructions in the Pipeline



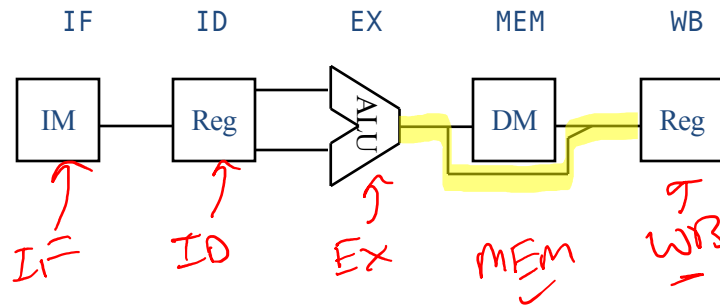
Mixed Instructions in the Pipeline




This is called a **structural hazard** – too many instructions want to use the same resource.
In our pipeline, we can make this hazard disappear (next slide).
In more complex pipelines, structural hazards are again possible.

Pipeline Principles

- All instructions that share a pipeline should have the same *stages* in the same *order*.
 - therefore, *add* does nothing during Mem stage
 - sw does nothing during WB stage
- All intermediate values must be latched each cycle.



Pipeline stages



Cycle	# inst
1	0
2	0
3	0
4	0
5	1
6	2
7	3

- What is the performance implication of making every instruction go through all 5 stages? (e.g., instead of 4 for add, 3 for beq, etc.)

(Choose **BEST** answer)

<input checked="" type="checkbox"/> A	Decreases peak throughput by 20%
<input checked="" type="checkbox"/> B	Increases program latency by 20%
<input checked="" type="checkbox"/> C	No significant impact on peak throughput or program latency
<input checked="" type="checkbox"/> D	Depends on how many R type instructions, beq, etc.
<input type="checkbox"/> E	None of the above

Handwritten annotations: 25% next to A, 20% next to B, 50% next to C, and 50% next to D. A red circle highlights option C.

Pipelined Datapath

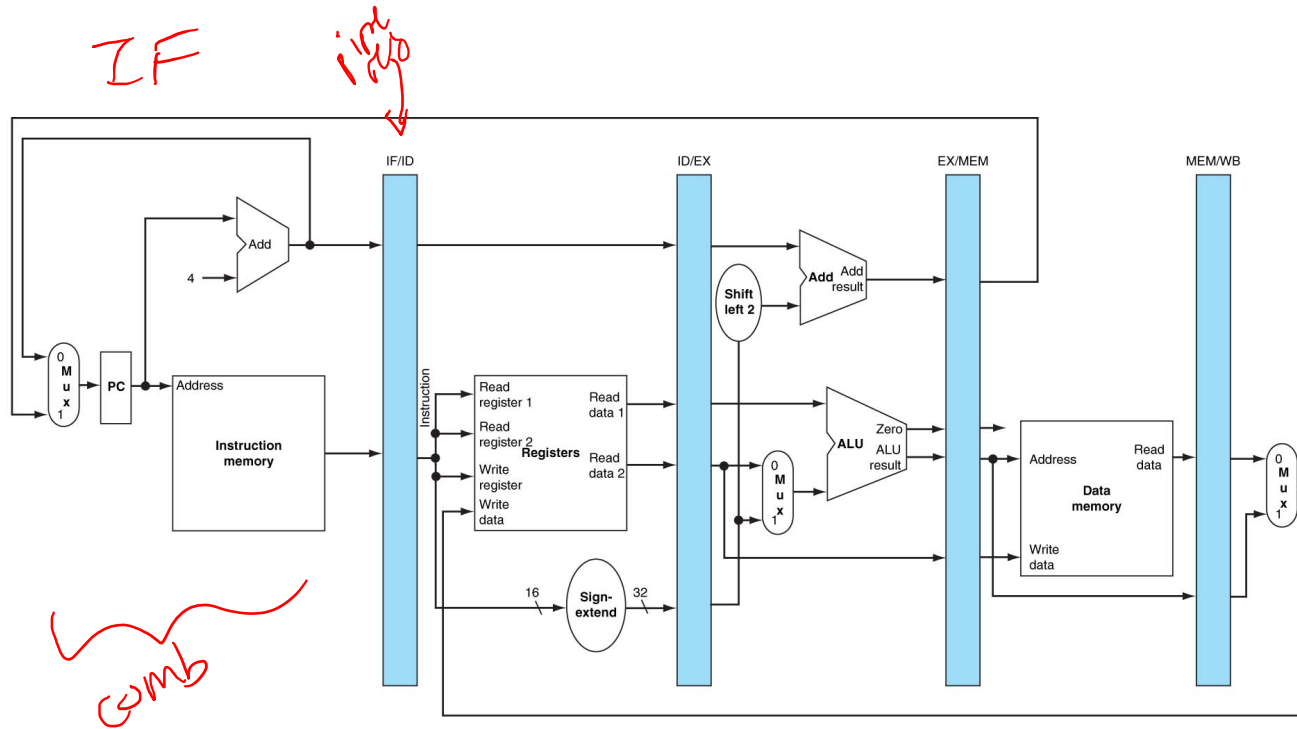
Instruction Fetch

Instruction Decode/
Register Fetch

Execute/
Address Calculation

Memory Access

Write Back



Pipelined Datapath

Instruction Fetch

Instruction Decode/
Register Fetch

Execute/
Address Calculation

Memory Access

Write Back

