

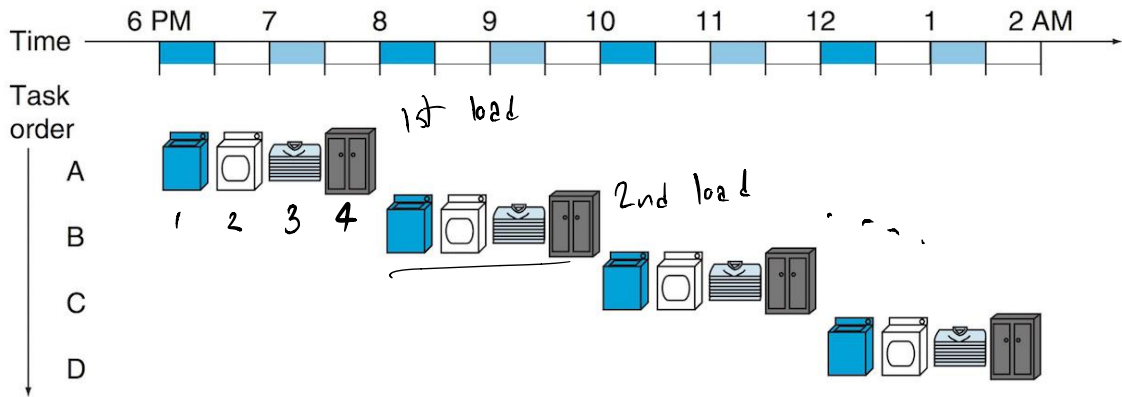
# Pipelining

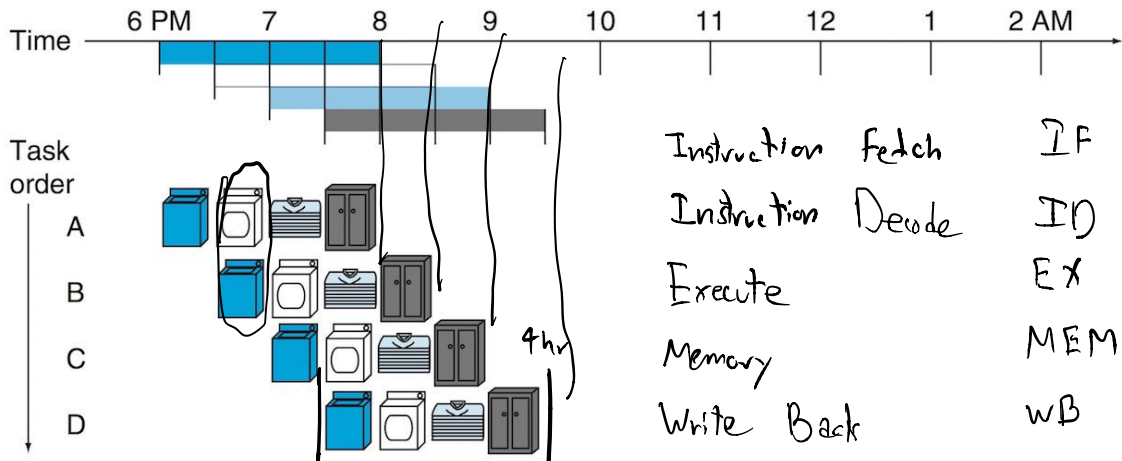
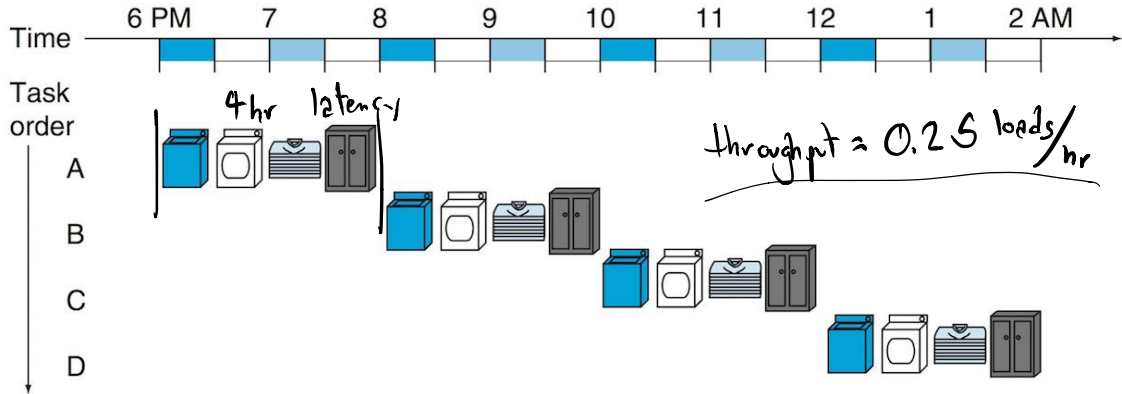
CSC 211 – December 1, 2020

# Datapath Lab Q&A

No questions!

# What is Pipelining?





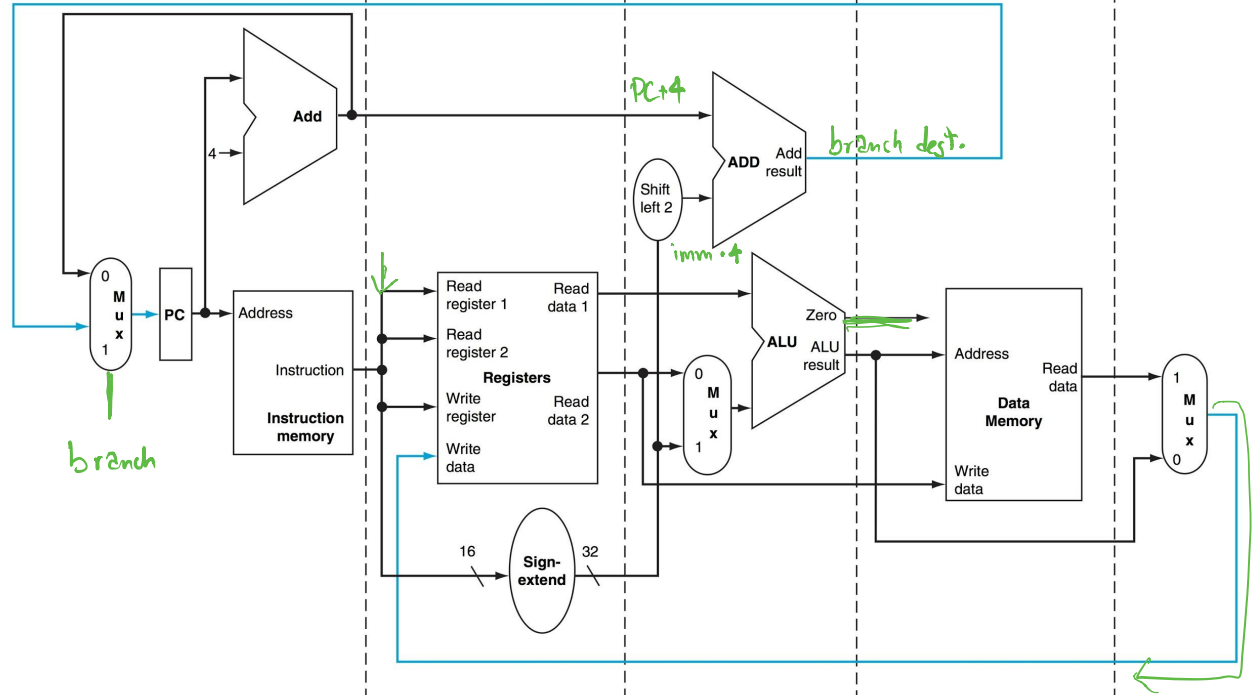
IF: Instruction fetch

ID: Instruction decode/  
register file read

EX: Execute/  
address calculation

MEM: Memory access

WB: Write back





# Comparing Performance

total time from start to end of the instruction's execution.

Compute maximum instruction latency and for single-cycle and pipelined datapaths using these parameters:

$$\text{IF} = 150\text{ps}$$

$$\text{ID} = 100\text{ps}$$

$$\text{EX} = 100\text{ps}$$

$$\text{MEM} = 200\text{ps}$$

$$\text{WB} = 100\text{ps}$$

$$\text{Pipeline registers} = 20\text{ps}$$

pipelining only

Latency

Throughput

Single-Cycle

$$650\text{ps}$$

$$\frac{1000\text{ps}}{650\text{ps}/\text{inst.}} = \sim 1.5 \text{ inst./ns}$$

$$= 1.5 \text{ billion inst./sec.}$$

Pipelined

$$\underline{5 \cdot 220\text{ps} = 1100\text{ps}}$$

$$\frac{1000\text{ps}}{220\text{ps}/\text{inst.}} = \sim 4.5 \text{ inst./ns}$$

$$= 4.5 \text{ billion inst./sec.}$$

## Little's Law

$$\underline{L} = \lambda \cdot \underline{W}$$

$L$  = length of queue  
(or # of tasks running concurrently)

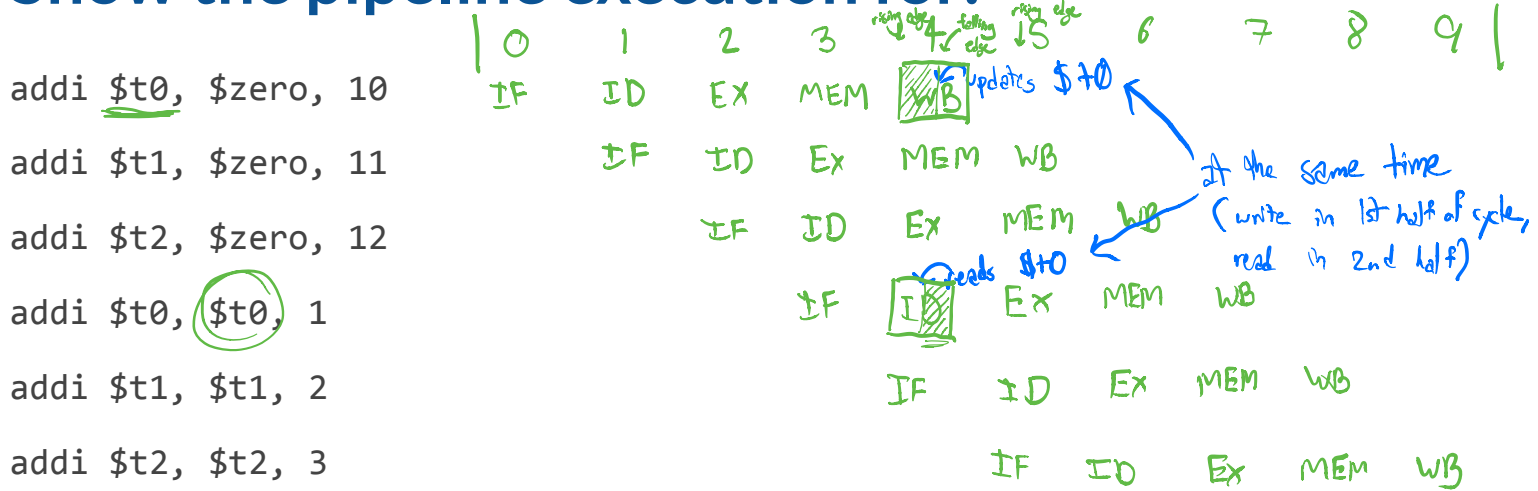
$\lambda$  = completion rate  
(inst. throughput)

$W$  = wait time  
(inst. latency)

$$\lambda = \frac{L}{W} \Rightarrow \lambda = \frac{5 \text{ inst.}}{1100 \text{ ps}} \rightarrow \frac{1 \text{ inst.}}{220 \text{ ps}}$$

# Pipeline Practice

# Show the pipeline execution for:



time  $\longrightarrow$

add \$t0, \$t0, \$zero  $\Rightarrow$  nop

# Show the pipeline execution for:

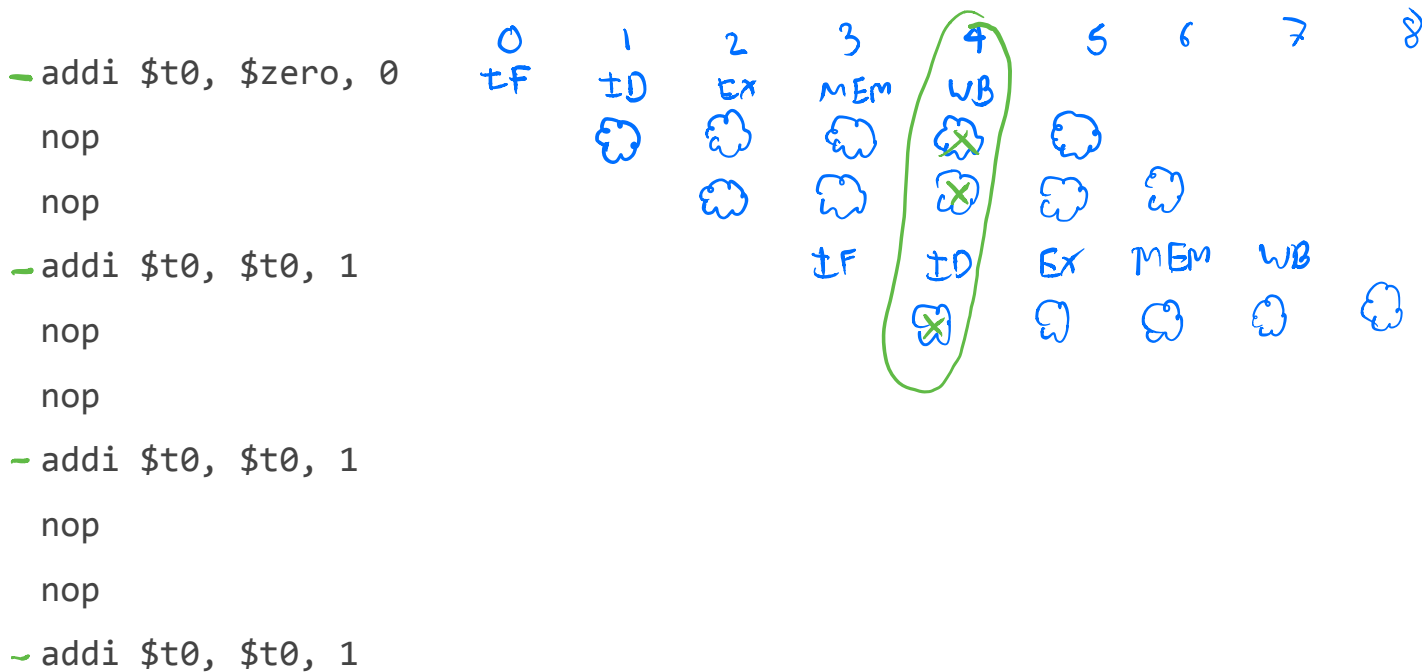
Assume \$t0 is initially 10



What is the value of \$t0 at the end?  $\$t0 = 12$

How can we fix this instruction sequence?

# Here's the fixed version



Static Scheduling Continued: reorder instructions

## Show the pipeline execution for:

	0	1	2	3	4	5	6	7	8
lw \$s0, 0(\$t0)	IF	ID	EX	MEM	WB				
addi \$t0, \$t0, 4		IF	ID	EX	MEM	WB			
lw \$s1, 0(\$t1) ←			IF	ID	EX	MEM	WB		
addi \$t1, \$t1, 4				IF	ID	EX	MEM	WB	
add \$s2, \$s0, \$s1 ←					IF	ID	EX	MEM	WB

Where will this sequence go wrong?  
(where does it differ from a single-cycle datapath?)

The last instruction uses  
\$s1 before it has been written

How can we fix this inst. sequence w/ static scheduling?

Swap the order of  
the 2nd and 3rd insts

# Show the pipeline execution for:

A addi \$t0, \$zero, 0 -

B add \$t0, \$t0, \$s0

C add \$t0, \$t0, \$s1

D addi \$t1, \$zero, 1 -

E add \$t1, \$t1, \$s2

F add \$t1, \$t1, \$s3

A

D

nop

B

E

nop

C

F

1. Where does this inst. sequence go wrong on a pipelined datapath?

Instructions B, C, E and F read values before they are ready.

2. How can we fix it w/ static scheduling?

See schedule above

# Show the pipeline execution for:

```
addi $s0, $zero, 0
```

```
j somewhere
```

```
addi $s0, $s0, 1
```

```
addi $s1, $s1, 1
```

```
...
```

somewhere:

```
addi $s1, $zero, 1
```

# Controlling a Pipelined Datapath

