

Pipelining

1

Changelog

Changes made in this version not seen in first lecture:

- 5 October 2017: put addq timing slide before critical path slides
- 5 October 2017: slide 25: arrows to fetch/decode registers point to outputs, not inputs
- 5 October 2017: slide 27: sum with no pipelining was 550 ps, not 500 ps
- 5 October 2017: slide 33: e_dstE and W_dstE labels were swapped
- 5 October 2017: slide 34: rA should have been D_rA, e_valA should have been d_valA

1

logistics

exam graded — scores on gradebook

keys on Collab

view exam, submit regrade requests via TPEGS


note: two questions dropped **outside of TPEGS**

median: 83.5; 25th percentile: 77.3; 75th percentile: 90.7

HCL homework due next Wednesday

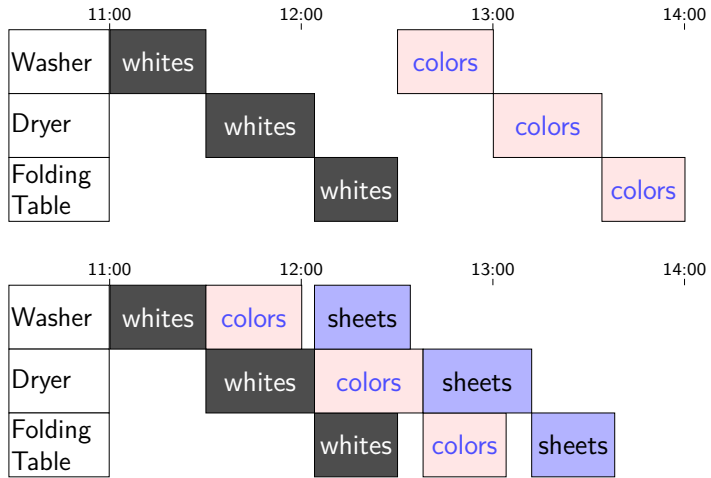
2

Human pipeline: laundry




3

Human pipeline: laundry




3

Waste (1)




4

Waste (1)




4

Waste (2)




5

Latency — Time for One



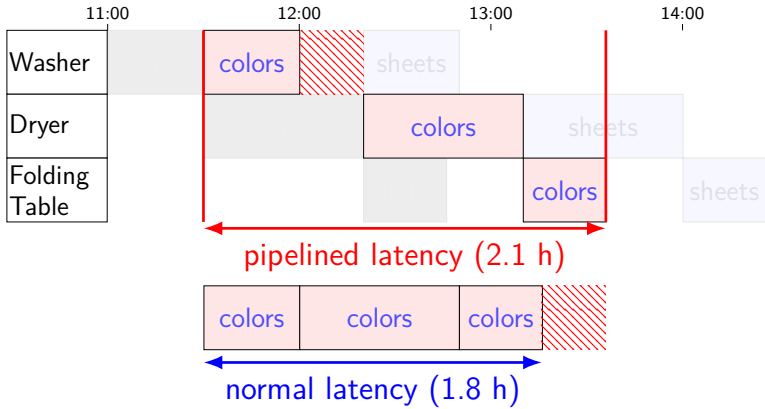
6

Latency — Time for One



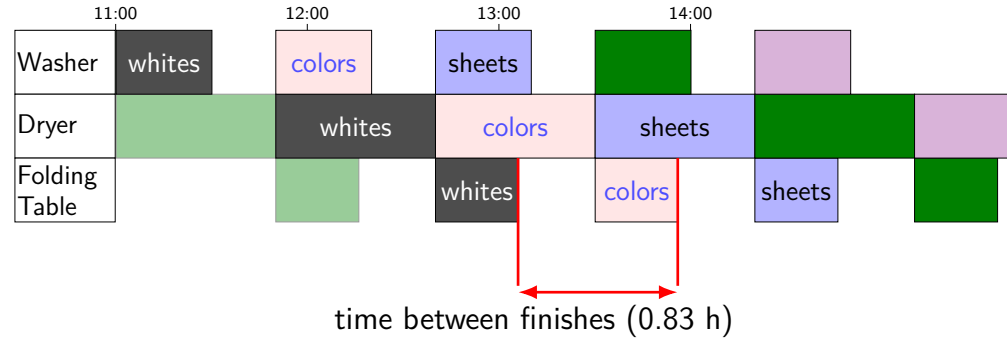
6

Latency — Time for One




6

Throughput — Rate of Many



7


Throughput — Rate of Many



$$\frac{1 \text{ load}}{0.83\text{h}} = 1.2 \text{ loads/h}$$

7


Throughput — Rate of Many



$$\frac{1 \text{ load}}{0.83\text{h}} = 1.2 \text{ loads/h}$$



7

times three circuit




8

times three circuit




8


times three circuit





100 ps latency \Rightarrow 10 results/ns throughput




times three and repeat




times three and repeat



pipelined times three




pipelined times three



$A(t+2)$	7	17	
$A(t+1)$		7	17
$2 \times A(t+1)$		14	34
$3 \times A(t+0)$			21


10

register tolerances



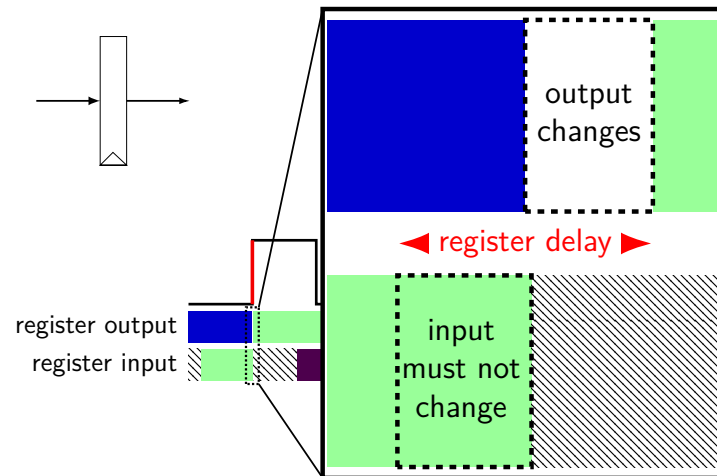
11

register tolerances



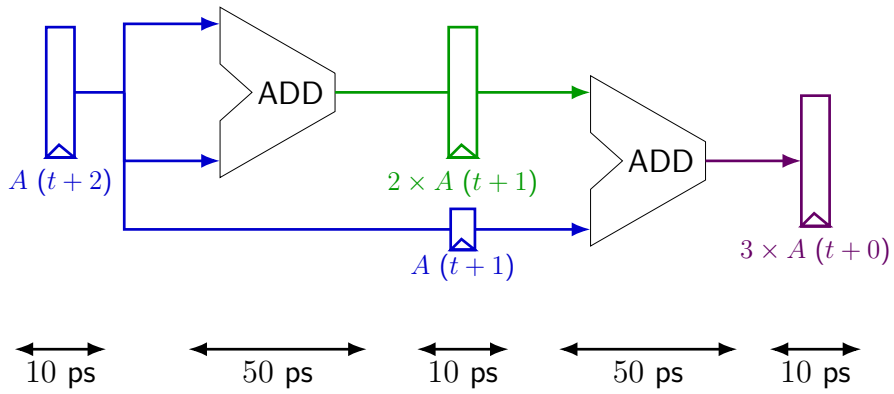
11

register tolerances



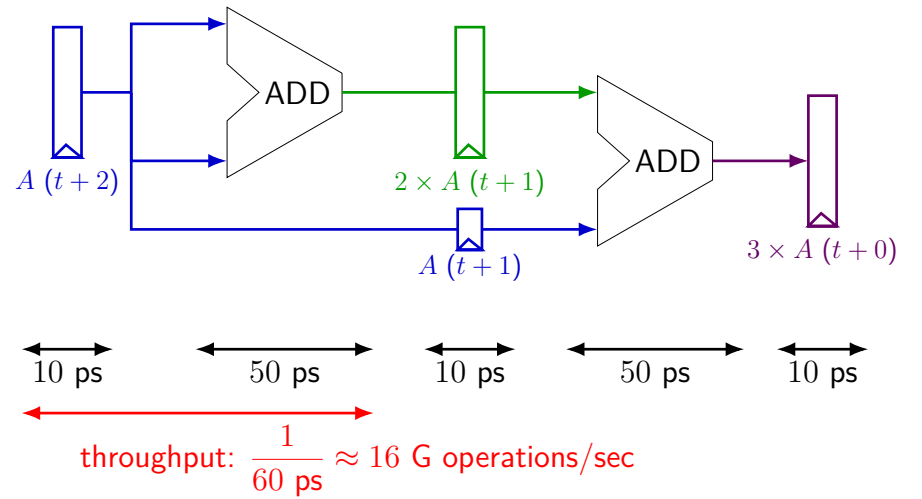
11

times three pipeline timing



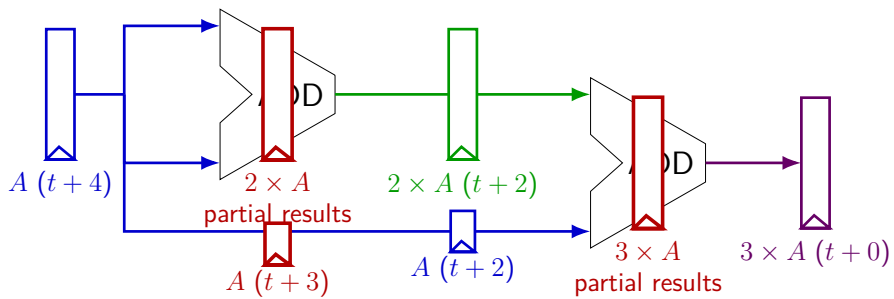
12

times three pipeline timing




12

deeper pipeline




13

deeper pipeline



13


deeper pipeline



exercise: throughput now?

13


deeper pipeline



$$\text{throughput: } \frac{1}{35 \text{ ps}} \approx 28 \text{ G ops/sec}$$

13

deeper pipeline

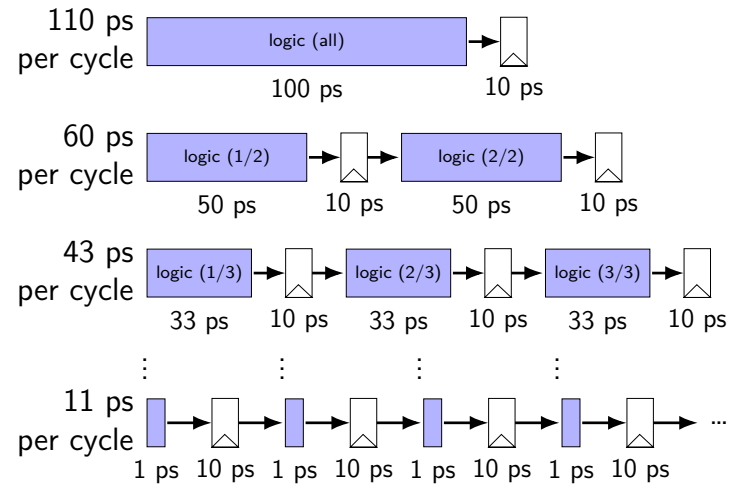


Problem: How much faster can we get?

Problem: Can we even do this?


14

diminishing returns: register delays




15


diminishing returns: register delays



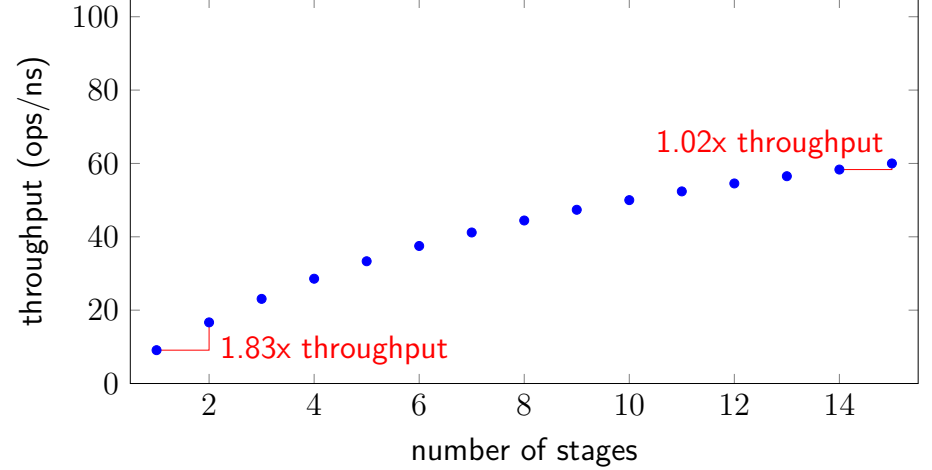
diminishing returns: register delays




diminishing returns: register delays



diminishing returns: register delays




diminishing returns: register delays



17

deeper pipeline




Problem: How much faster can we get?

Problem: Can we even do this?

18


deeper pipeline



exercise: throughput now? (didn't split second add evenly)

19

deeper pipeline

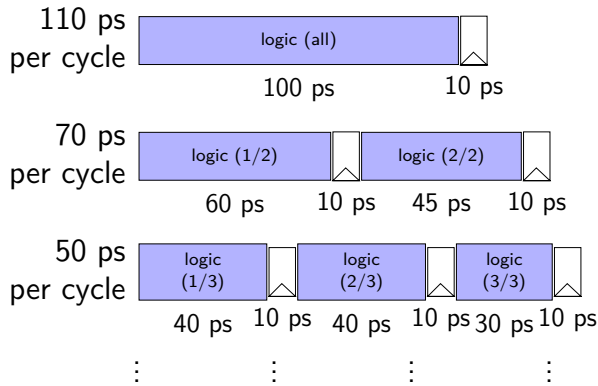


19

diminishing returns: uneven split

Can we split up some logic (e.g. adder) arbitrarily?

Probably not...

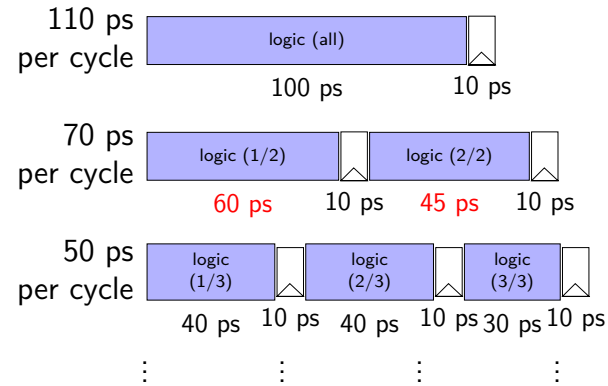


20

diminishing returns: uneven split

Can we split up some logic (e.g. adder) arbitrarily?

Probably not...

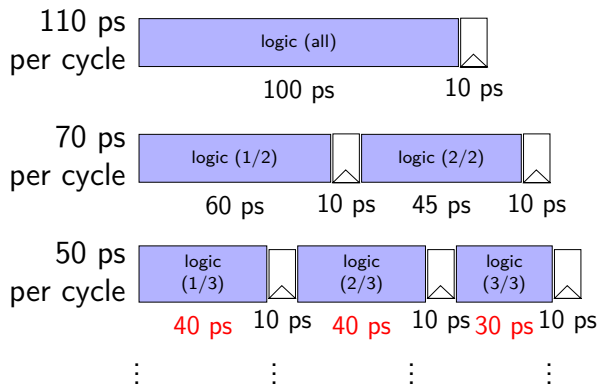


20

diminishing returns: uneven split

Can we split up some logic (e.g. adder) arbitrarily?

Probably not...



20

textbook SEQ 'stages'

conceptual order only

Fetch: read instruction memory

Decode: read register file

Execute: arithmetic (ALU)

Memory: read/write data memory

Writeback: write register file

PC Update: write PC register

21

textbook SEQ 'stages'

conceptual order only

Fetch: read instruction memory

Decode: read register file

Execute: arithmetic (ALU)

Memory: read/write data memory

Writeback: write register file

PC Update: write PC register

writes happen
at end of cycle

21

textbook SEQ 'stages'

conceptual order only

Fetch: read instruction memory

Decode: read register file

Execute: arithmetic (ALU)

Memory: read/write data memory

Writeback: write register file

PC Update: write PC register

reads — "magic"
like combinatorial logic
as values available

21

textbook stages

~~conceptual order only~~ pipeline stages

Fetch/PC Update: read instruction memory;
compute next PC

Decode: read register file

Execute: arithmetic (ALU)

Memory: read/write data memory

Writeback: write register file

22

textbook stages

~~conceptual order only~~ pipeline stages

Fetch/PC Update: read instruction memory;
compute next PC

Decode: read register file

Execute: arithmetic (ALU)

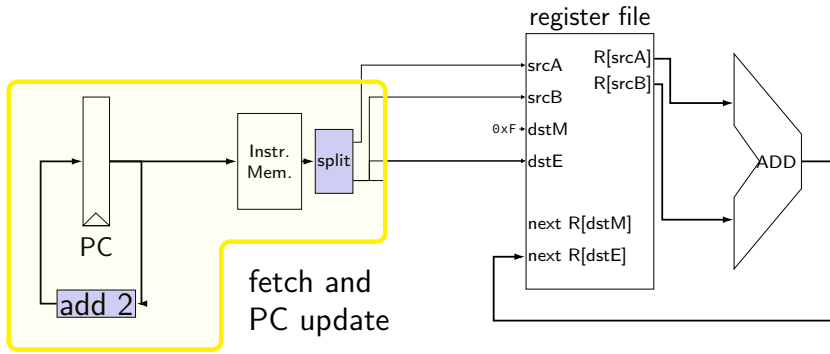
Memory: read/write data memory

Writeback: write register file

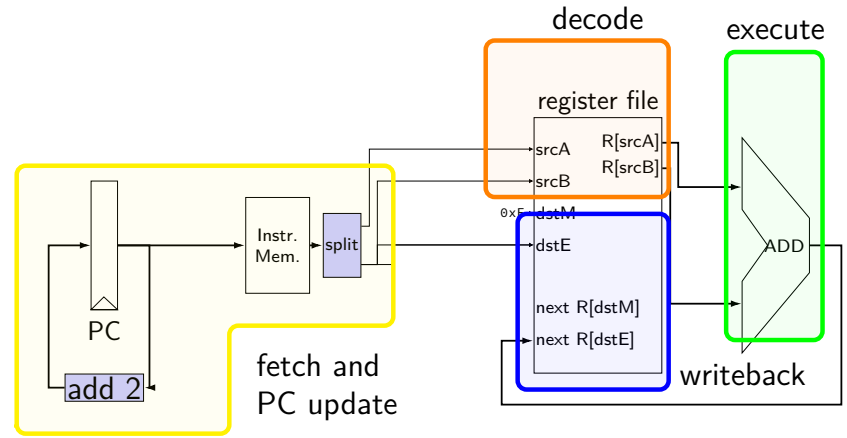
5 stages
one instruction in each
compute next to start immediately

22

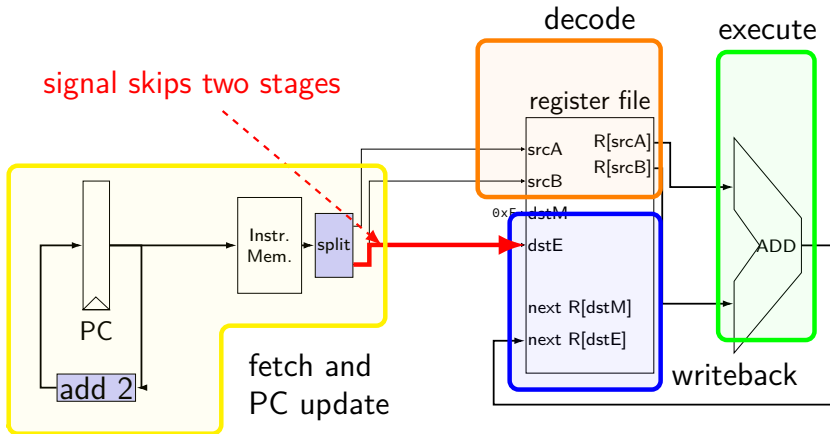
addq CPU



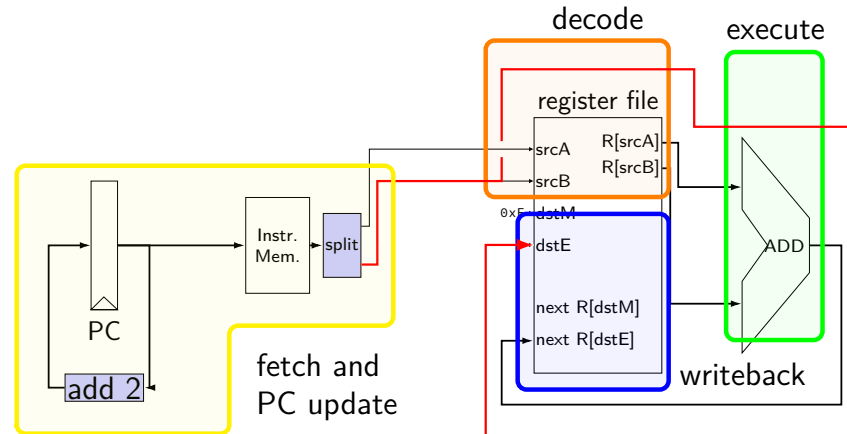
addq CPU



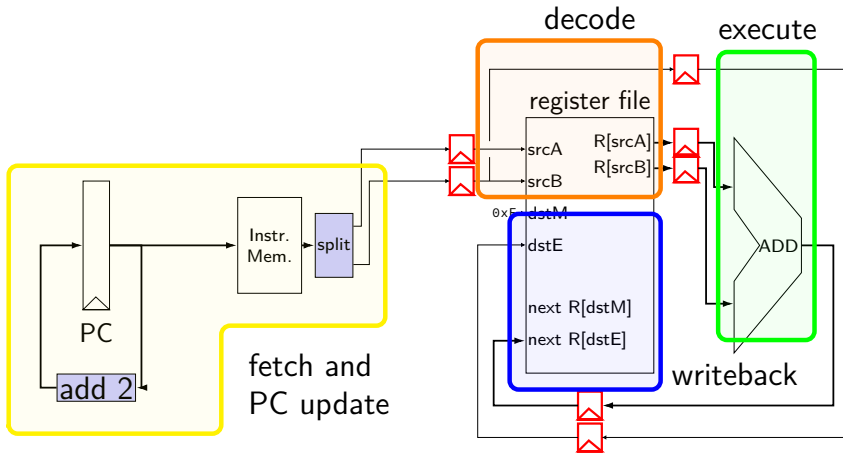
addq CPU



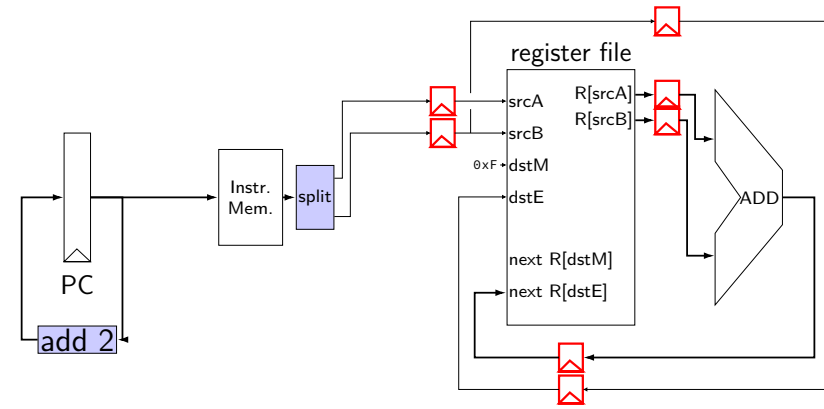
addq CPU



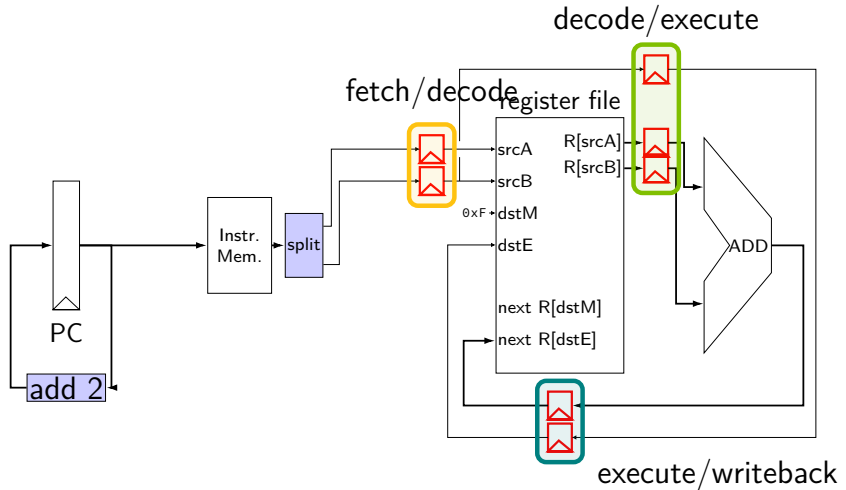
pipelined addq processor



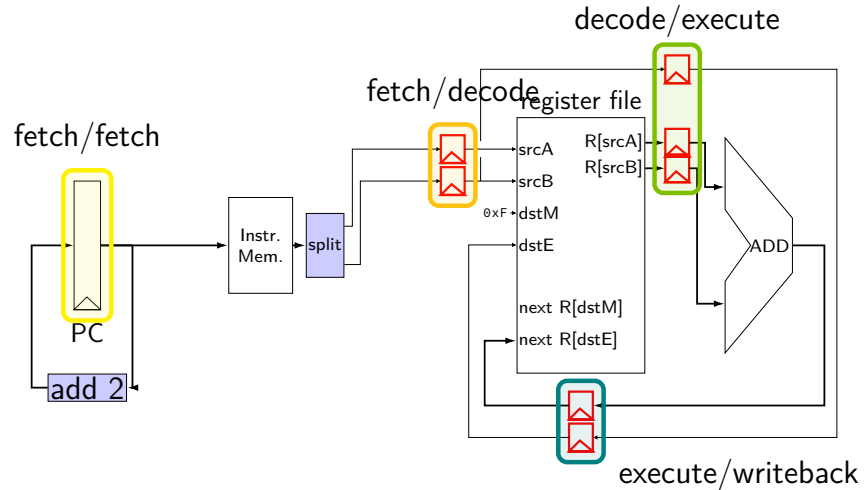
pipelined addq processor



pipelined addq processor

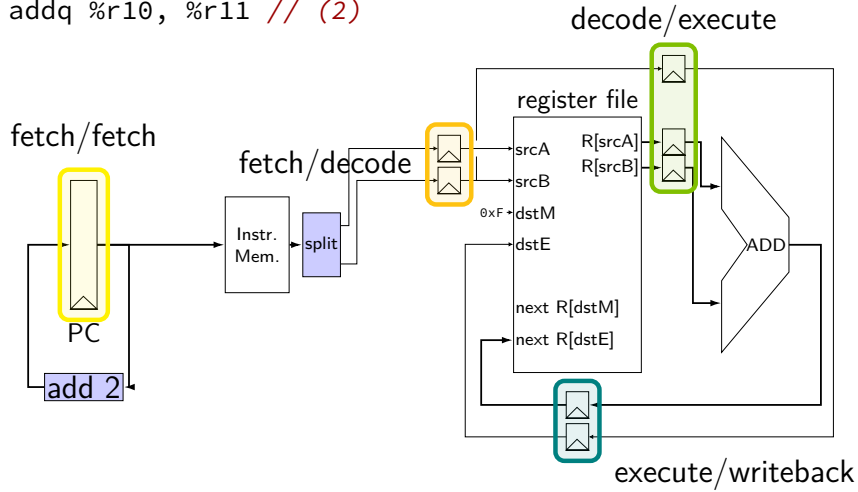


pipelined addq processor



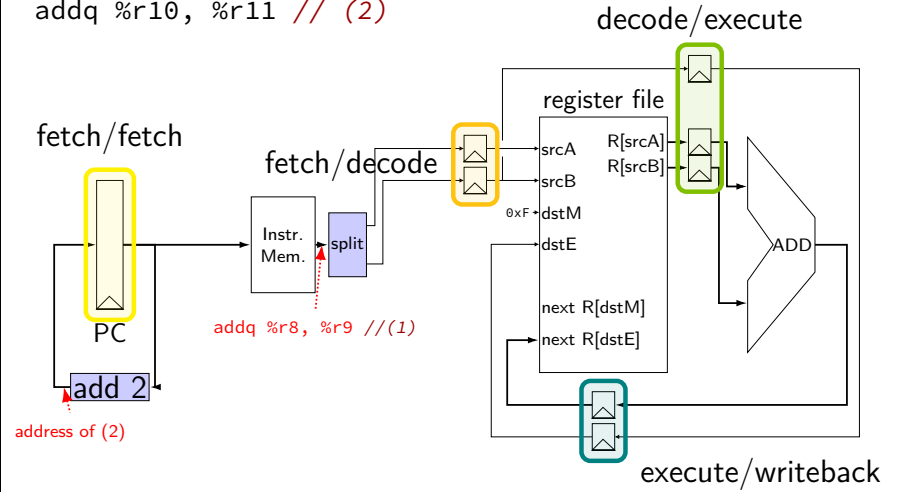
addq execution

```
addq %r8, %r9 // (1)  
addq %r10, %r11 // (2)
```



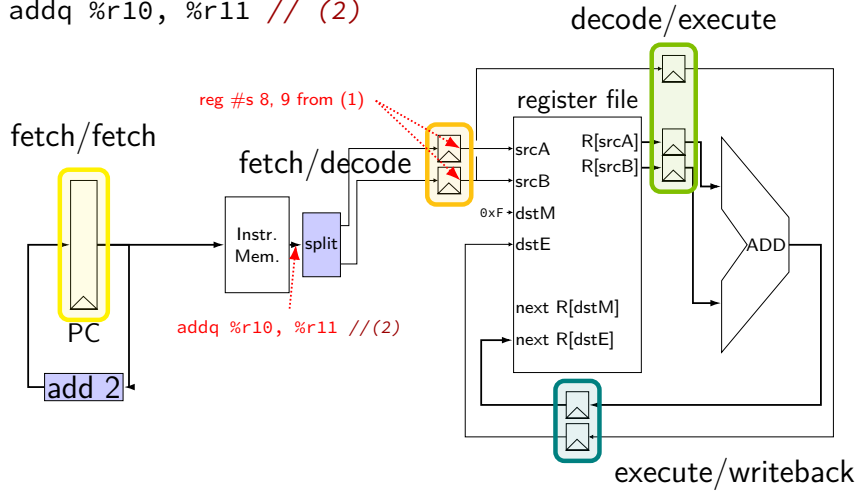
addq execution

```
addq %r8, %r9 // (1)  
addq %r10, %r11 // (2)
```



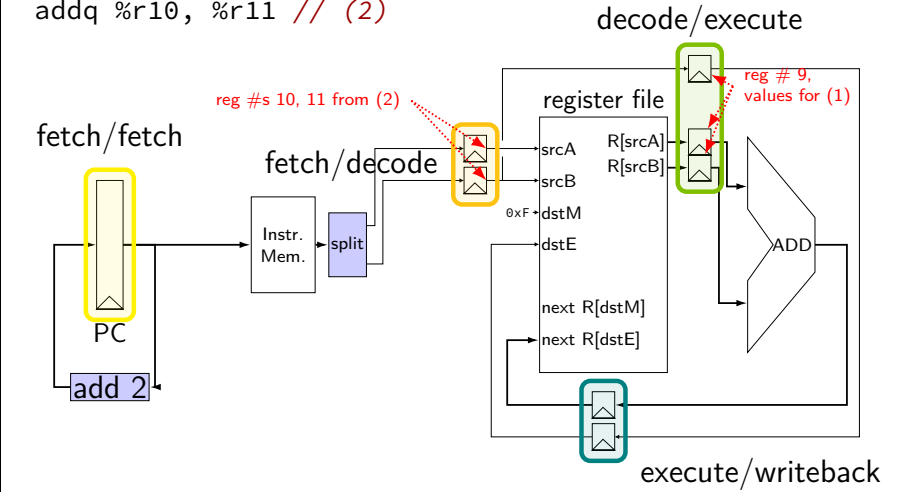
addq execution

```
addq %r8, %r9 // (1)  
addq %r10, %r11 // (2)
```




addq execution

```
addq %r8, %r9 // (1)  
addq %r10, %r11 // (2)
```



addq processor timing


```
// initially %r8 = 800,
//           %r9 = 900, etc.
addq %r8, %r9
addq %r10, %r11
addq %r12, %r13
addq %r9, %r8
```



cycle	fetch	fetch/decode		decode/execute			execute/writeback	
	PC	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstE]	dstE
0	0x0							
1	0x2	8	9					
2	0x4	10	11	800	900	9		
3	0x6	12	13	1000	1100	11	1700	9
4		9	8	1200	1300	13	2100	11
5				1700	800	8	2500	13
6							2500	8

addq processor timing


```
// initially %r8 = 800,
//           %r9 = 900, etc.
addq %r8, %r9
addq %r10, %r11
addq %r12, %r13
addq %r9, %r8
```



cycle	fetch	fetch/decode		decode/execute			execute/writeback	
	PC	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstE]	dstE
0	0x0							
1	0x2	8	9					
2	0x4	10	11	800	900	9		
3	0x6	12	13	1000	1100	11	1700	9
4		9	8	1200	1300	13	2100	11
5				1700	800	8	2500	13
6							2500	8

addq processor timing


```
// initially %r8 = 800,
//           %r9 = 900, etc.
addq %r8, %r9
addq %r10, %r11
addq %r12, %r13
addq %r9, %r8
```



cycle	fetch	fetch/decode		decode/execute			execute/writeback	
	PC	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstE]	dstE
0	0x0							
1	0x2	8	9					
2	0x4	10	11	800	900	9		
3	0x6	12	13	1000	1100	11	1700	9
4		9	8	1200	1300	13	2100	11
5				1700	800	8	2500	13
6							2500	8

addq processor timing


```
// initially %r8 = 800,
//           %r9 = 900, etc.
addq %r8, %r9
addq %r10, %r11
addq %r12, %r13
addq %r9, %r8
```



cycle	fetch	fetch/decode		decode/execute			execute/writeback	
	PC	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstE]	dstE
0	0x0							
1	0x2	8	9					
2	0x4	10	11	800	900	9		
3	0x6	12	13	1000	1100	11	1700	9
4		9	8	1200	1300	13	2100	11
5				1700	800	8	2500	13
6							2500	8

addq processor timing

```
// initially %r8 = 800,
//           %r9 = 900, etc.
addq %r8, %r9
addq %r10, %r11
addq %r12, %r13
addq %r9, %r8
```




cycle	fetch	fetch/decode		decode/execute			execute/writeback	
	PC	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstM]	next R[dstE]
0	0x0							
1	0x2	8	9					
2	0x4	10	11	800	900	9		
3	0x6	12	13	1000	1100	11	1700	9
4		9	8	1200	1300	13	2100	11
5				1700	800	8	2500	13
6							2500	8

26

addq processor performance

example delays:

path	time
add 2	80 ps
instruction memory	200 ps
register file read	125 ps
add	100 ps
register file write	125 ps



no pipelining: 1 instruction per 550 ps

add up everything but add 2 (critical (slowest) path)

pipelining: 1 instruction per 200 ps + pipeline register delays

slowest path through stage + pipeline register delays

latency: 800 ps + pipeline register delays (4 cycles)

27

critical path

every path from state output to state input needs enough time

output — may change on rising edge of clock

input — must be stable sufficiently before rising edge of clock

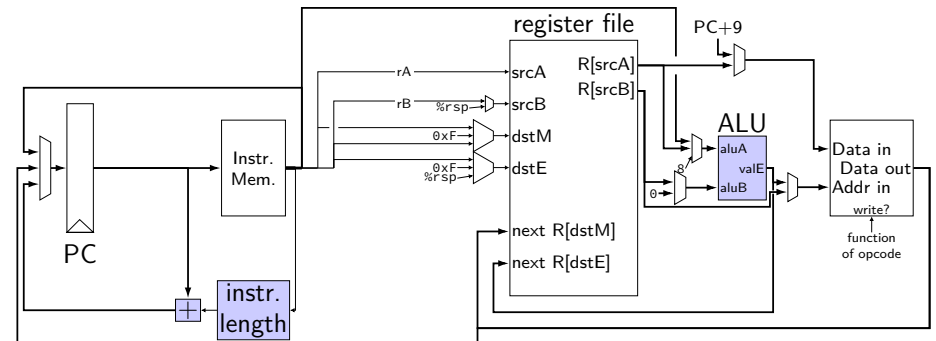
critical path: **slowest** of all these paths — determines cycle time

times three: slowest part of ALU ended up mattering

matters with or without pipelining

28

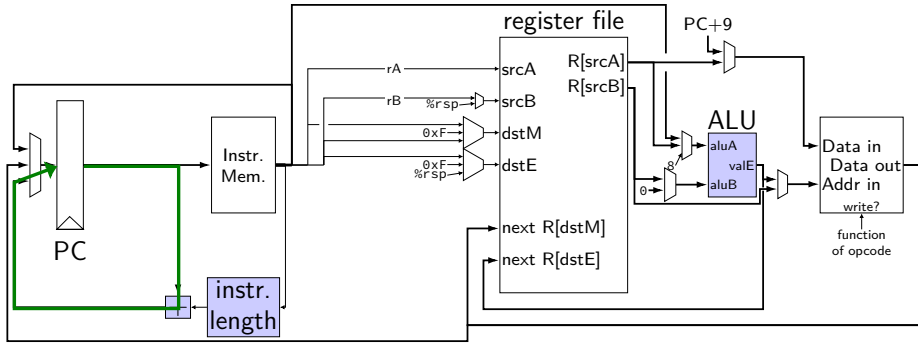
SEQ paths



29

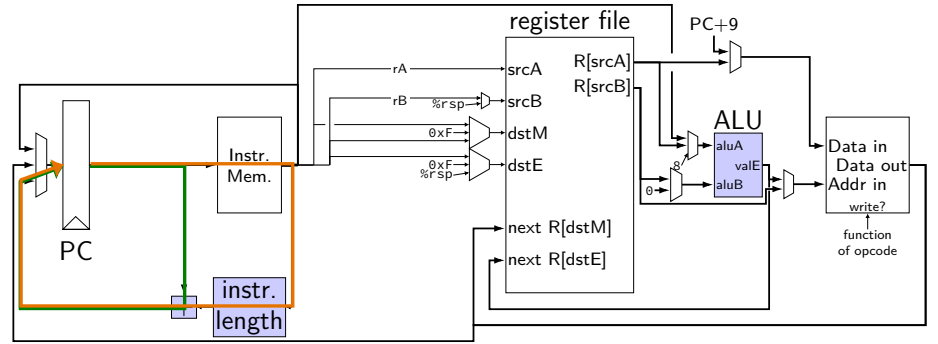
SEQ paths

path 1: 25 picoseconds



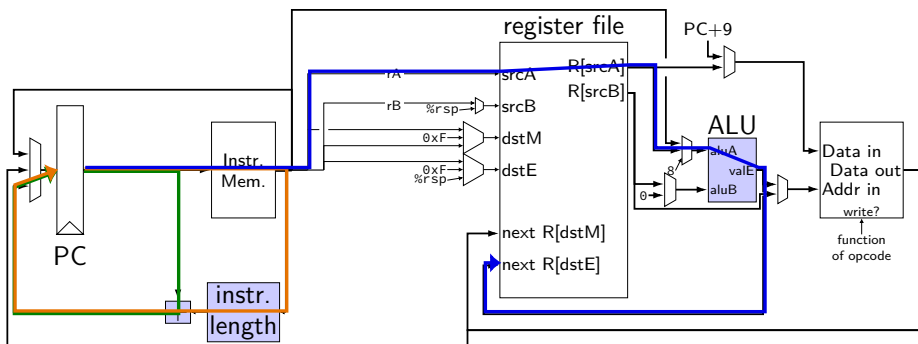
SEQ paths

path 1: 25 picoseconds path 2: 50 picoseconds



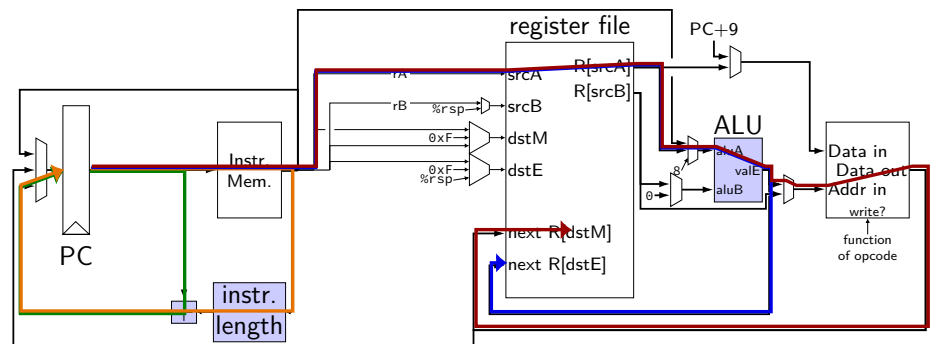
SEQ paths

path 1: 25 picoseconds path 2: 50 picoseconds
path 3: 400 picoseconds



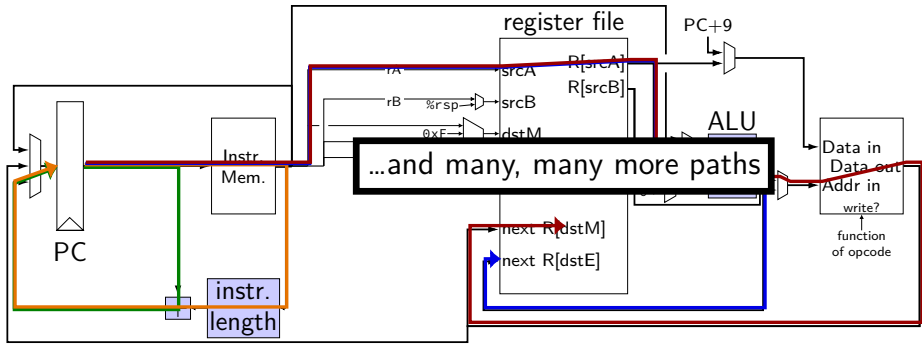
SEQ paths

path 1: 25 picoseconds path 2: 50 picoseconds
path 3: 400 picoseconds path 4: 900 picoseconds
... ..




SEQ paths

path 1: 25 picoseconds path 2: 50 picoseconds
path 3: 400 picoseconds path 4: 900 picoseconds
... ...



29

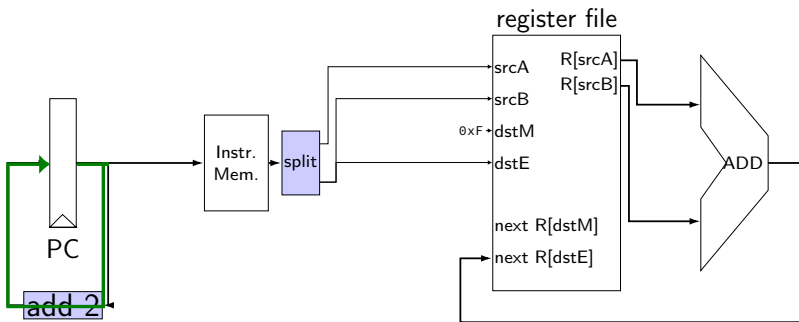
sequential addq paths



30

sequential addq paths

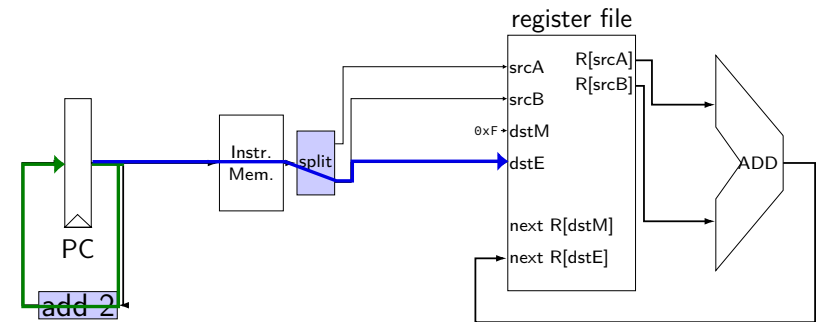
path 1: 25 picoseconds



30

sequential addq paths

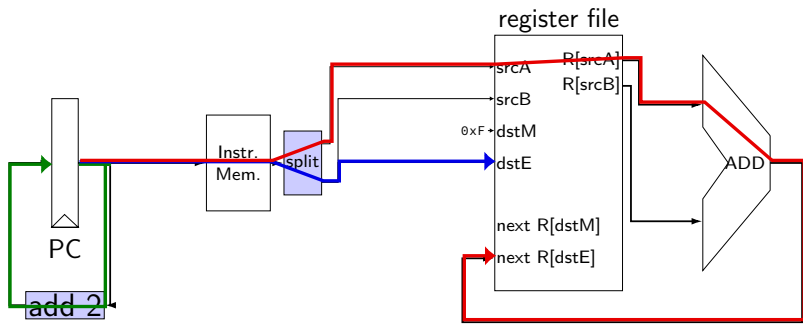
path 1: 25 picoseconds
path 2: 375 picoseconds



30

sequential addq paths


path 1: 25 picoseconds
 path 2: 375 picoseconds
 path 3: 500 picoseconds



30

sequential addq paths


path 1: 25 picoseconds
 path 2: 375 picoseconds
 path 3: 500 picoseconds
 path 4: 500 picoseconds



30

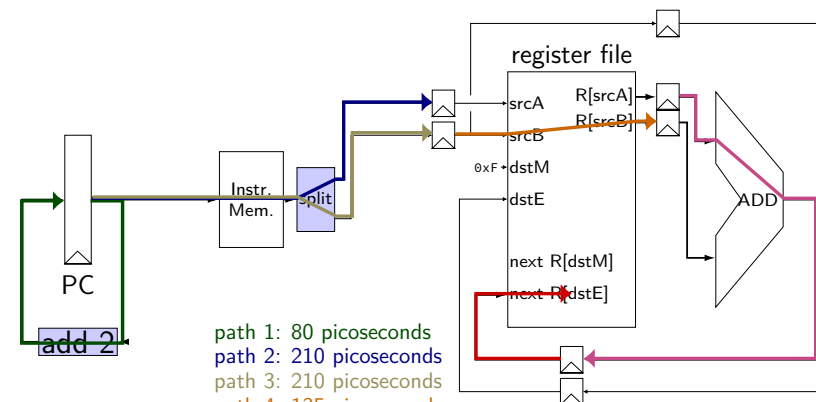
sequential addq paths

path 1: 25 picoseconds
 path 2: 375 picoseconds
 path 3: 500 picoseconds
 path 4: 500 picoseconds
 overall cycle time: 500 picoseconds (longest path)



30

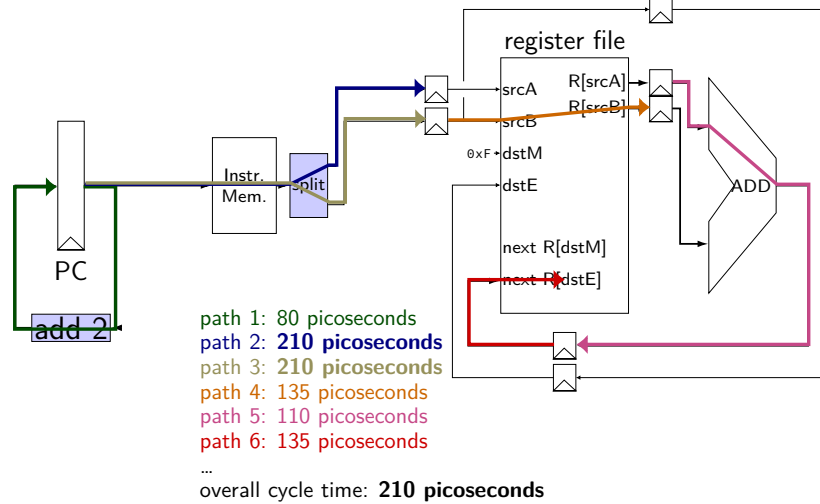
pipelined addq paths



path 1: 80 picoseconds
 path 2: 210 picoseconds
 path 3: 210 picoseconds
 path 4: 135 picoseconds
 path 5: 110 picoseconds
 path 6: 135 picoseconds
 ...
 overall cycle time: 210 picoseconds

31

pipelined addq paths



31


exercise

path	time
add 2	50 ps
instruction memory	200 ps
register file read	125 ps
add	100 ps
register file write	125 ps

pipeline register delay: 10ps

how will throughput improve if we double the speed of the instruction memory?

- A. 2.00x
- B. 1.70x to 1.99x
- C. 1.60x to 1.69x
- D. 1.50x to 1.59x
- E. less than 1.50x



32

exercise


path	time
add 2	50 ps
instruction memory	200 ps
register file read	125 ps
add	100 ps
register file write	125 ps

pipeline register delay: 10ps

how will throughput improve if we double the speed of the instruction memory?

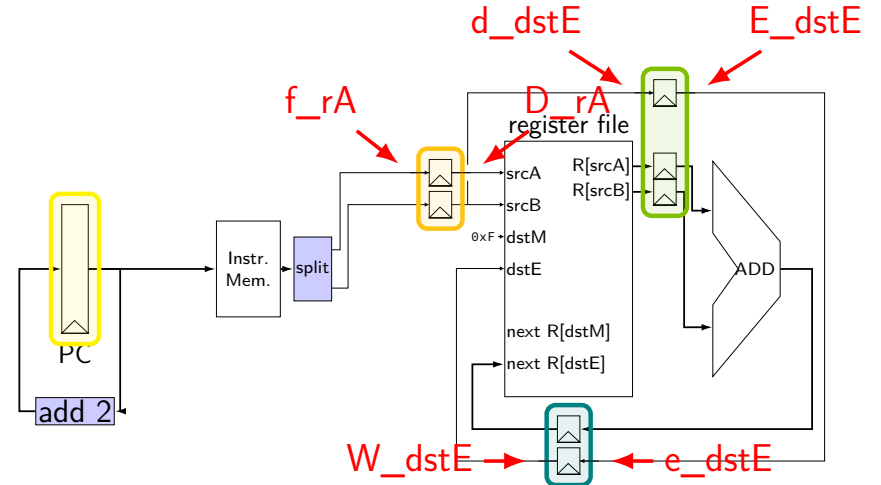
- A. 2.00x
- B. 1.70x to 1.99x
- C. 1.60x to 1.69x
- D. 1.50x to 1.59x
- E. less than 1.50x

$$\frac{1}{135} \div \frac{1}{210} = 1.56x \text{ — D}$$



32

pipeline register naming convention



33

pipeline register naming convention

f — fetch sends values here

D — decode receives values here

d — decode sends values here

...

34

addq HCL

```

...
/* f: from fetch */
f_rA = i10bytes[12..16];
f_rB = i10bytes[12..16];

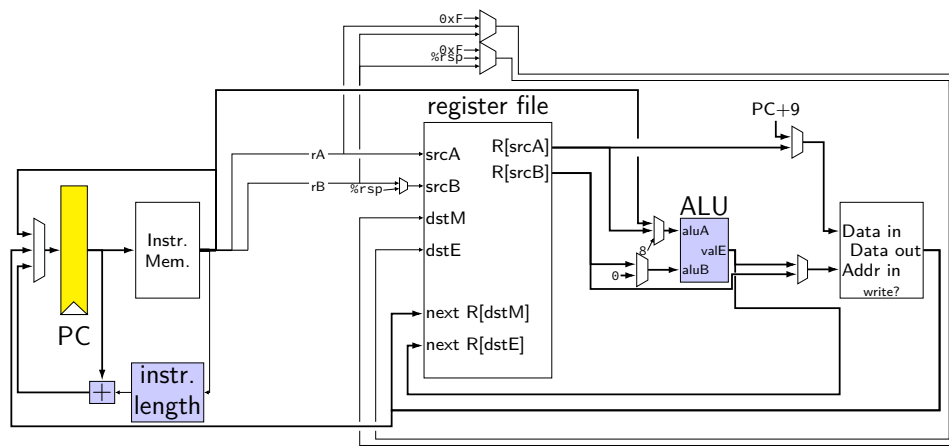
/* fetch to decode */
/* f_rA -> D_rA, etc. */
register fd {
    rA : 4 = REG_NONE;
    rB : 4 = REG_NONE;
}

/* D: to decode
   d: from decode */
d_dstE = D_rB;
/* use register file: */
reg_srcA = D_rA;
d_valA = reg_outputA;
...

/* decode to execute */
register dE {
    dstE : 4 = REG_NONE;
    valA : 64 = 0;
    valB : 64 = 0;
}
    
```

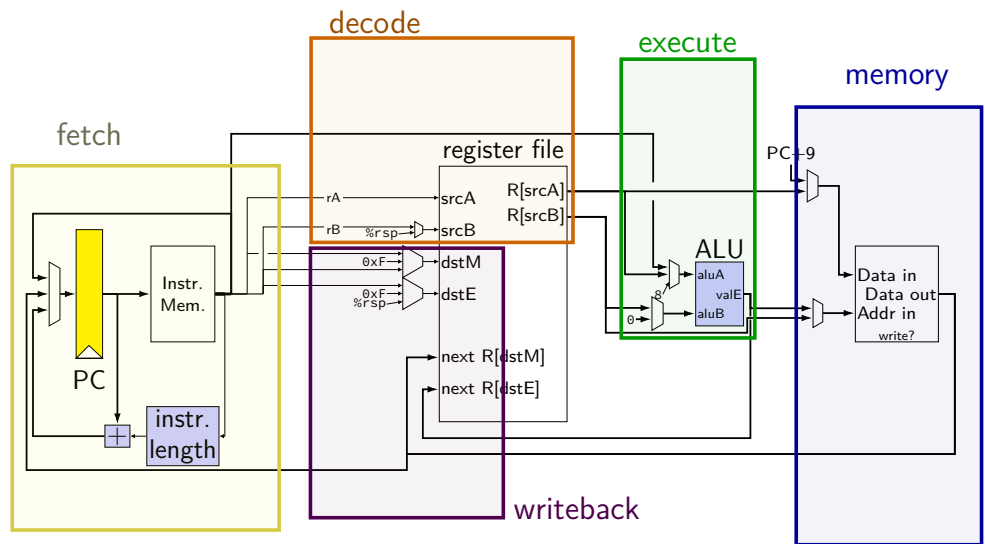
35

SEQ without stages



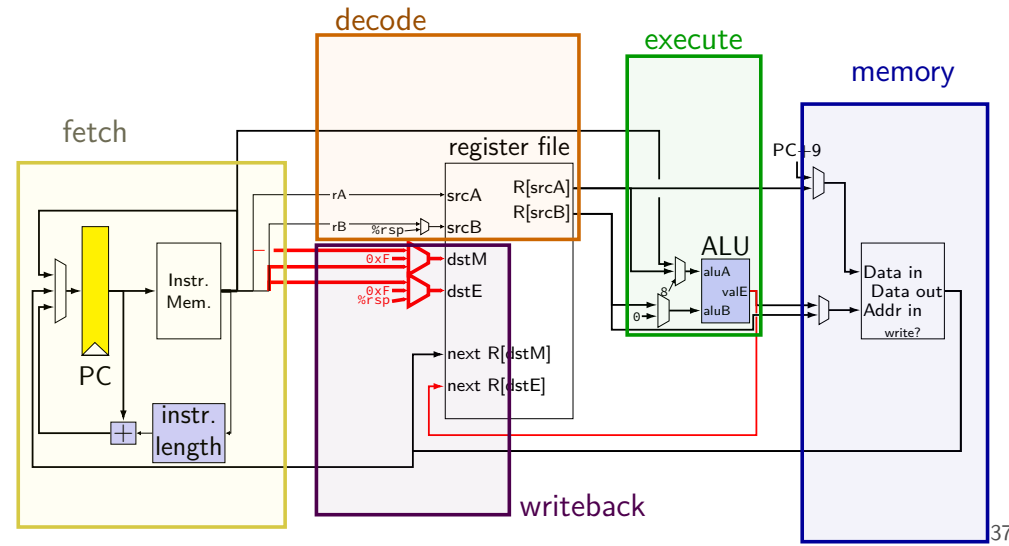
36

SEQ with stages

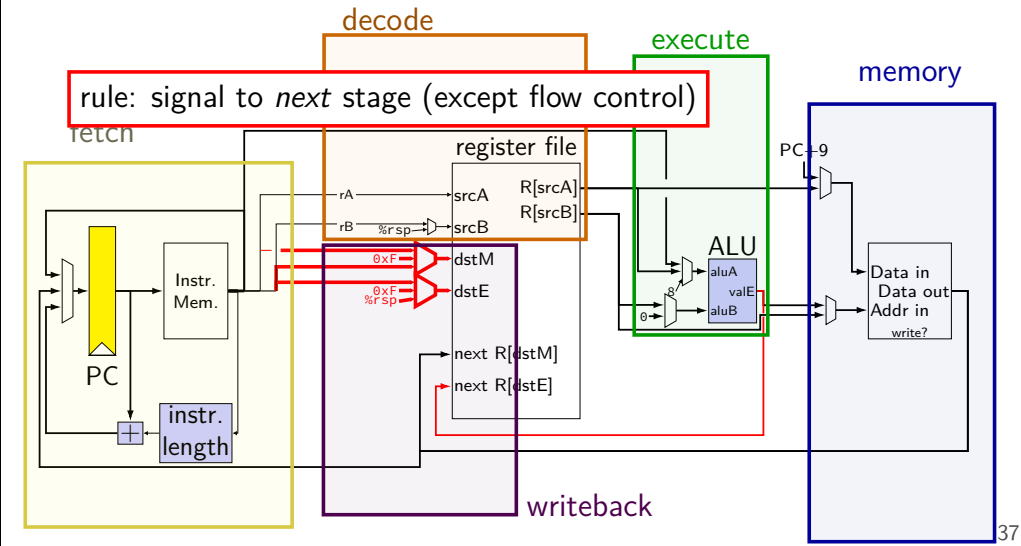


37

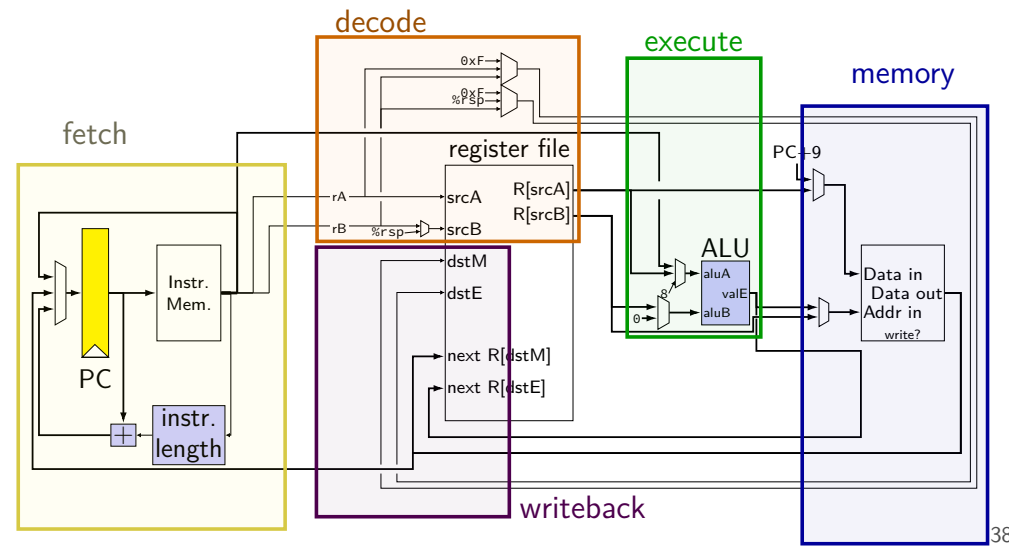
SEQ with stages



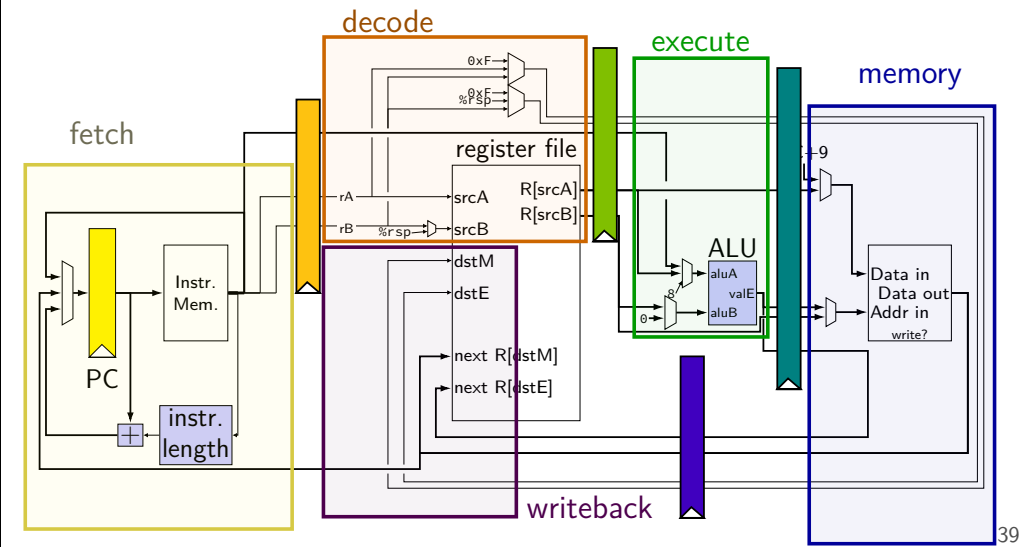
SEQ with stages



SEQ with stages (actually sequential)



adding pipeline registers



adding pipeline registers

