




Pipelining (part 1)


Human pipeline: laundry




Human pipeline: laundry



Waste (1)




Waste (1)



3

Waste (2)




4

Latency — Time for One



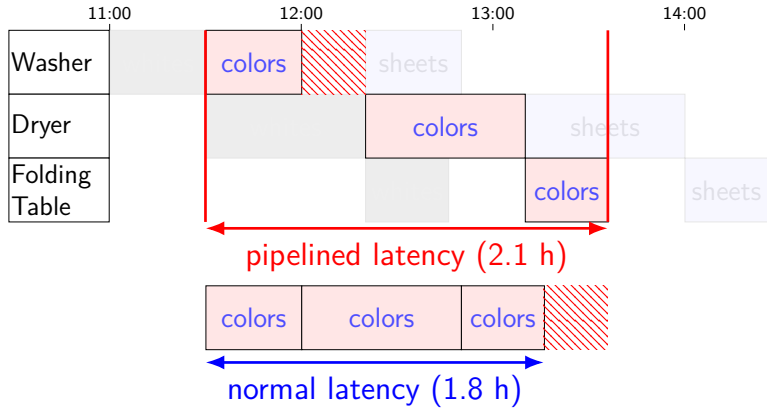
5

Latency — Time for One



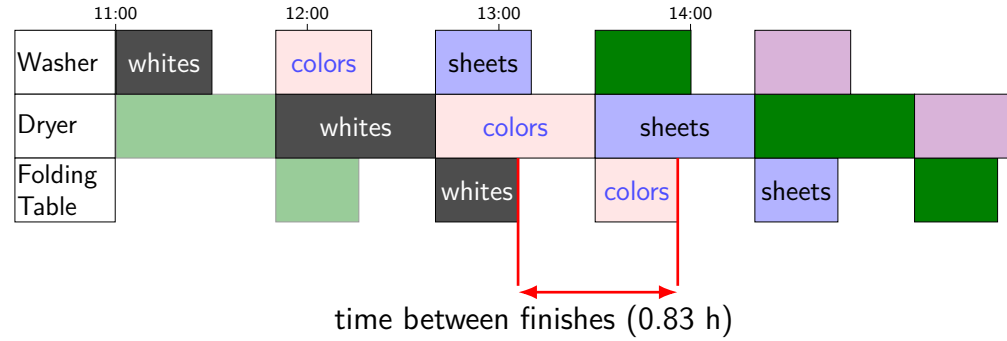
5

Latency — Time for One



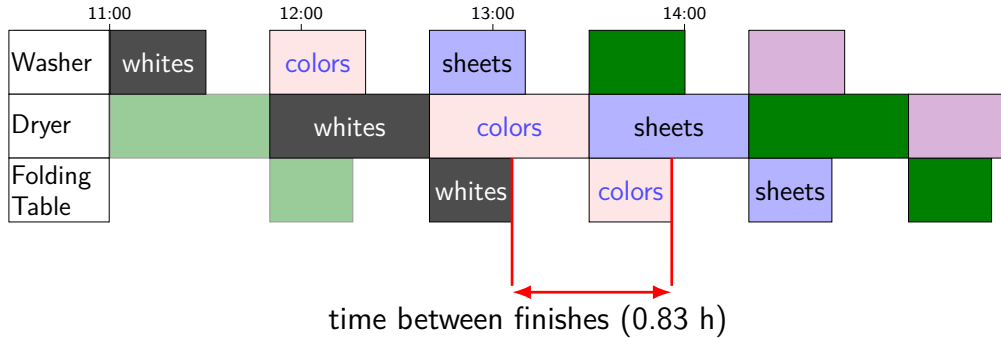
5

Throughput — Rate of Many



6


Throughput — Rate of Many



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

6


Throughput — Rate of Many





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

6


times three circuit




times three circuit




times three circuit




100 ps latency \implies 10 results/ns throughput



times three and repeat




times three and repeat



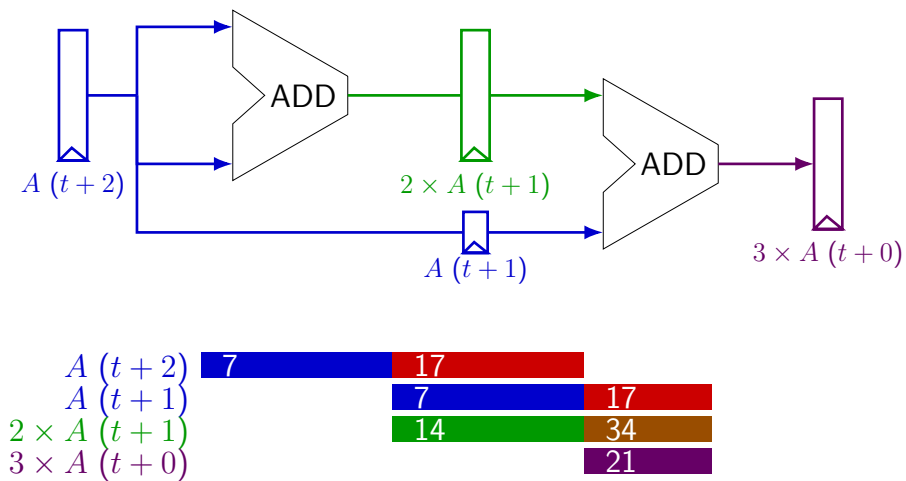
8

pipelined times three



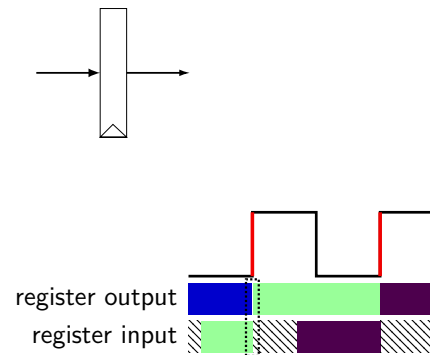
9

pipelined times three




9

register tolerances



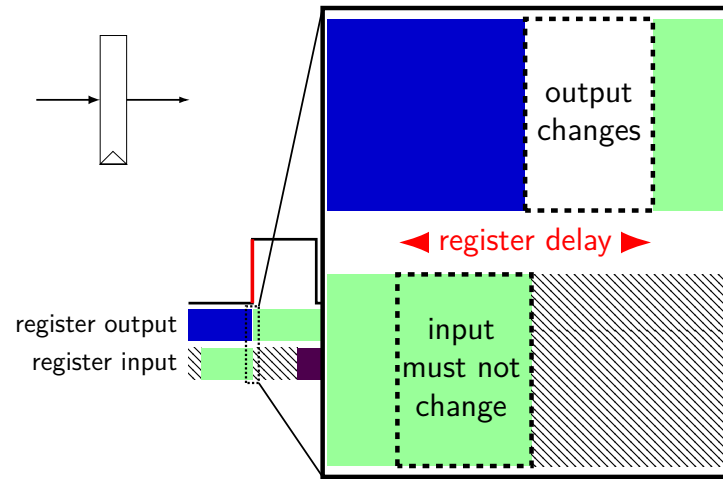
10

register tolerances



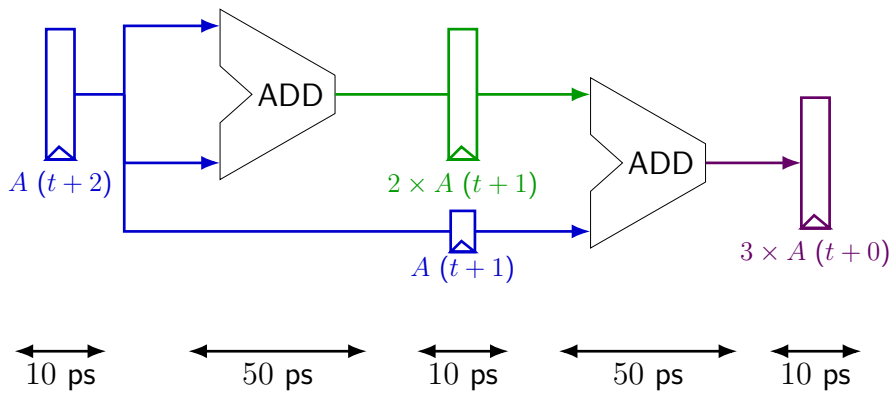
10

register tolerances



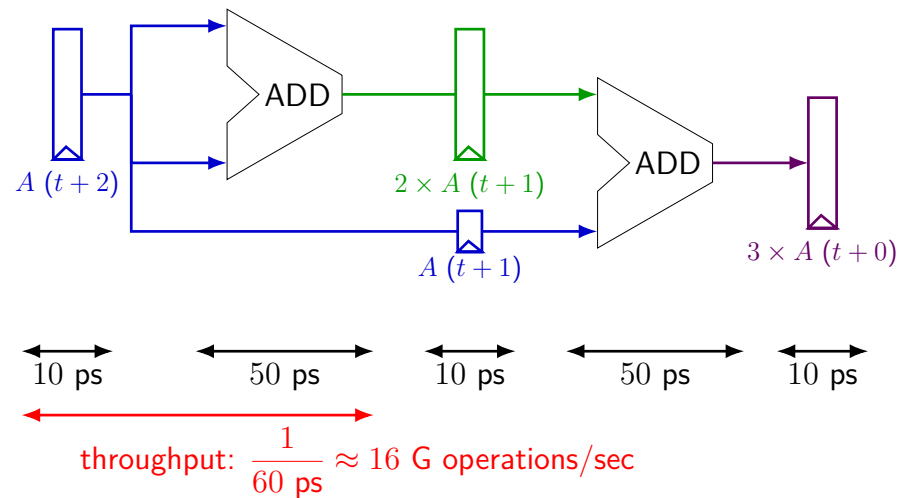
10

times three pipeline timing



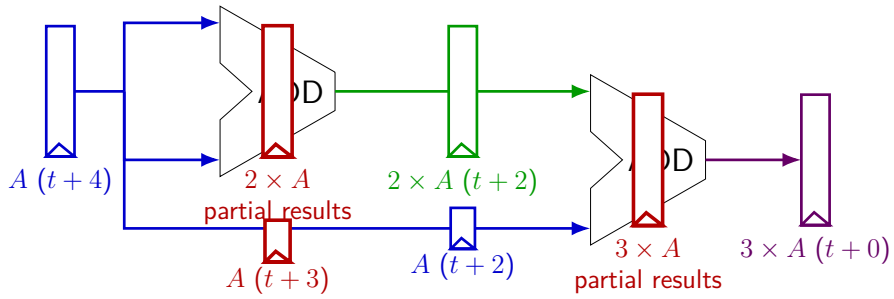
11

times three pipeline timing




11

deeper pipeline




12

deeper pipeline



12

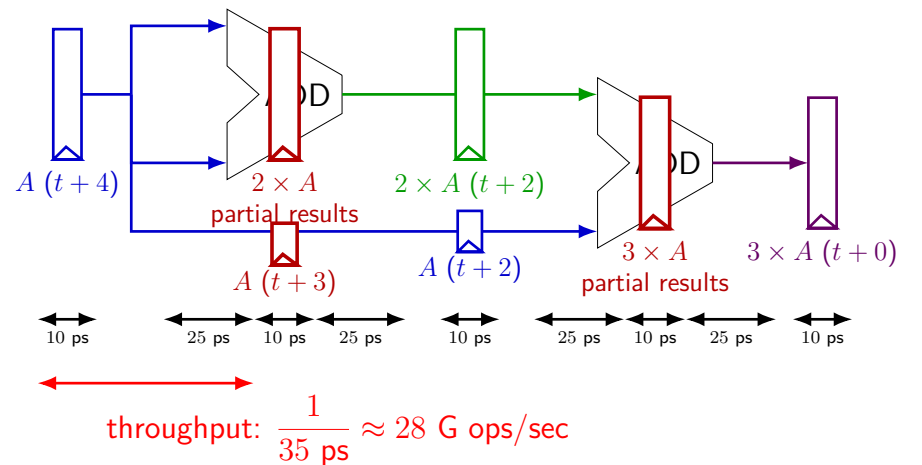
deeper pipeline



exercise: throughput now?


12

deeper pipeline



12

deeper pipeline

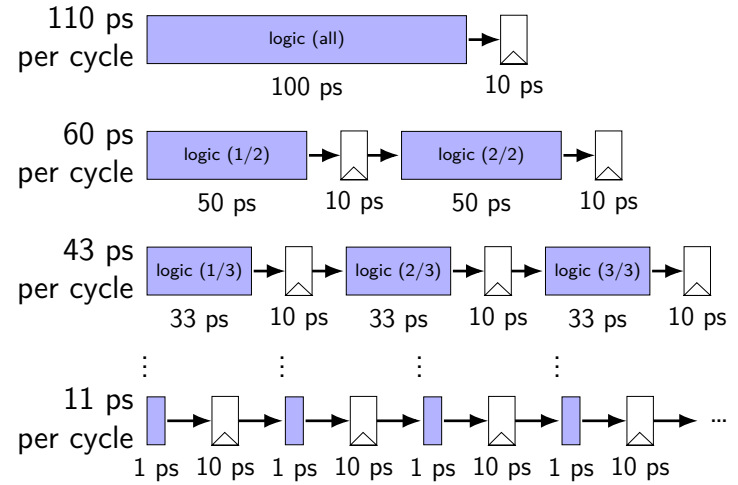


Problem: How much faster can we get?

Problem: Can we even do this?

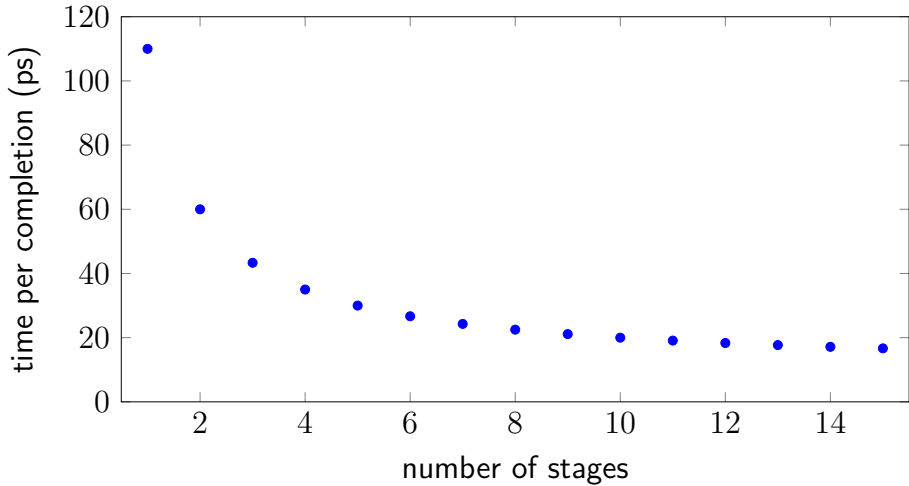
13

diminishing returns: register delays




14

diminishing returns: register delays



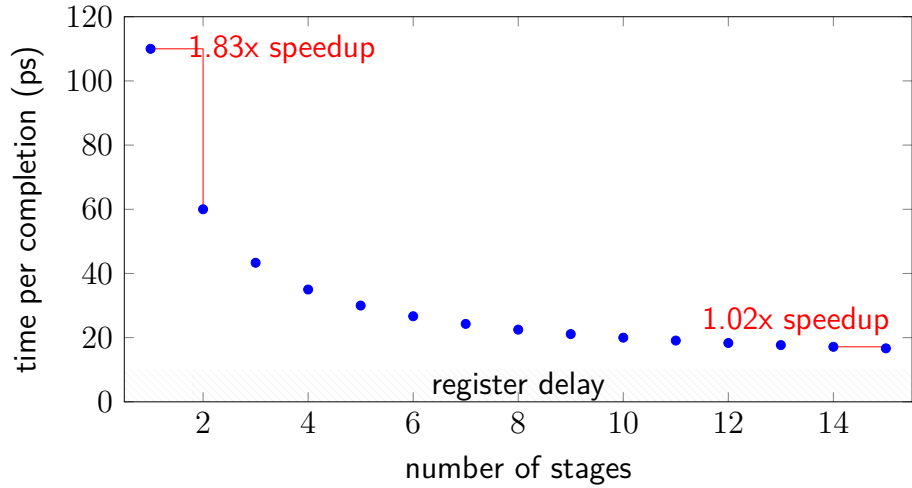
15

diminishing returns: register delays




15

diminishing returns: register delays




15

diminishing returns: register delays




16

diminishing returns: register delays



16

deeper pipeline




Problem: How much faster can we get?

Problem: **Can we even do this?**

17


deeper pipeline



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

18

deeper pipeline

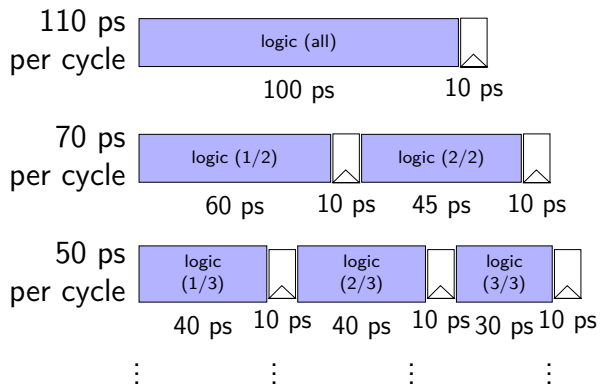


18

diminishing returns: uneven split

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

Probably not...

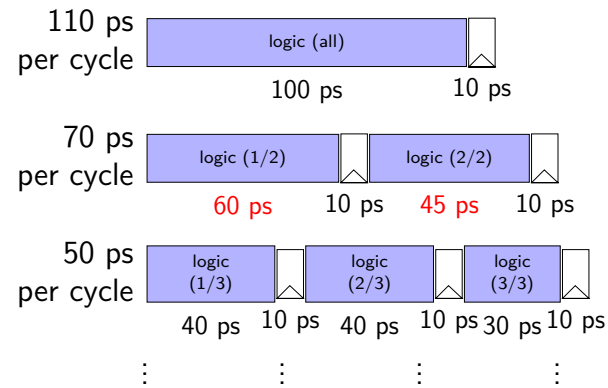


19

diminishing returns: uneven split

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

Probably not...

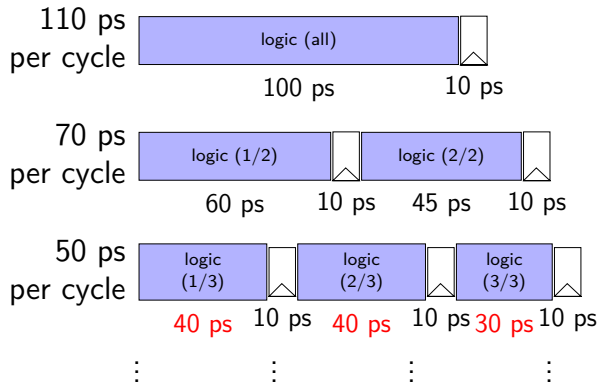


19

diminishing returns: uneven split

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

Probably not...



19

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

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

writes happen
at end of cycle

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

reads — "magic"
like combinatorial logic
as values available

20

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

21

textbook stages

conceptual order only pipeline stages

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

Decode: read register file

Execute: arithmetic (ALU)

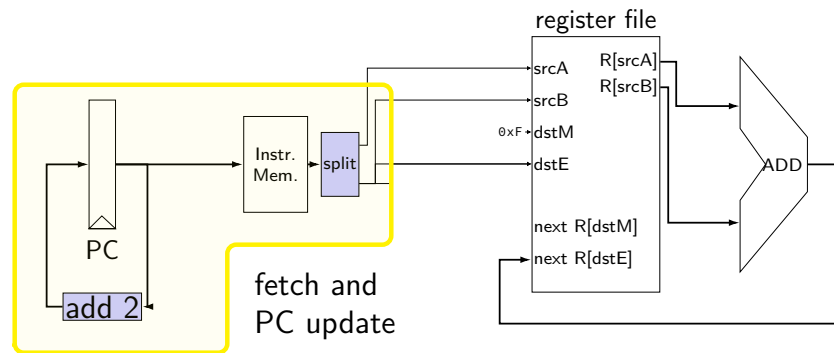
Memory: read/write data memory

Writeback: write register file

5 stages
one instruction in each
compute next to start immediately

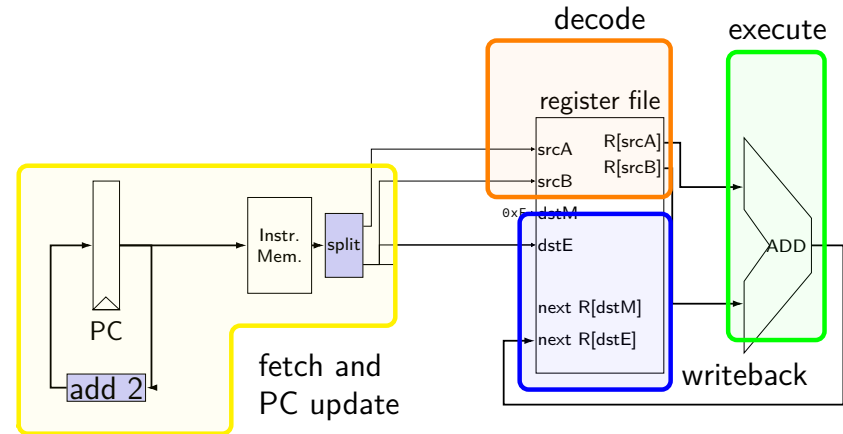
21

addq CPU



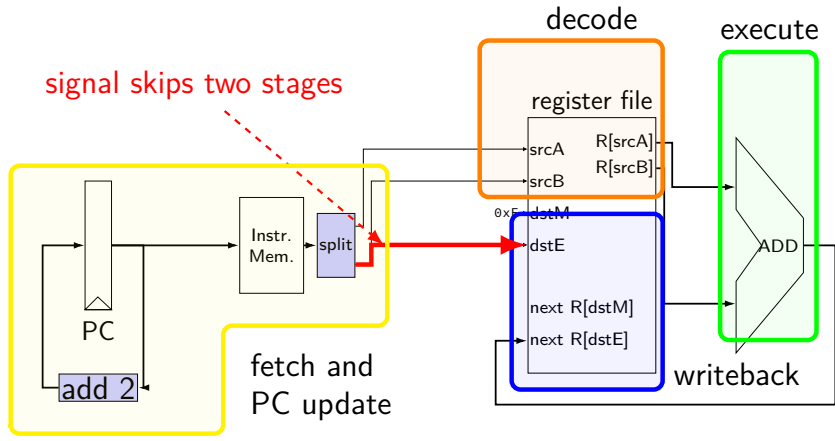
22

addq CPU

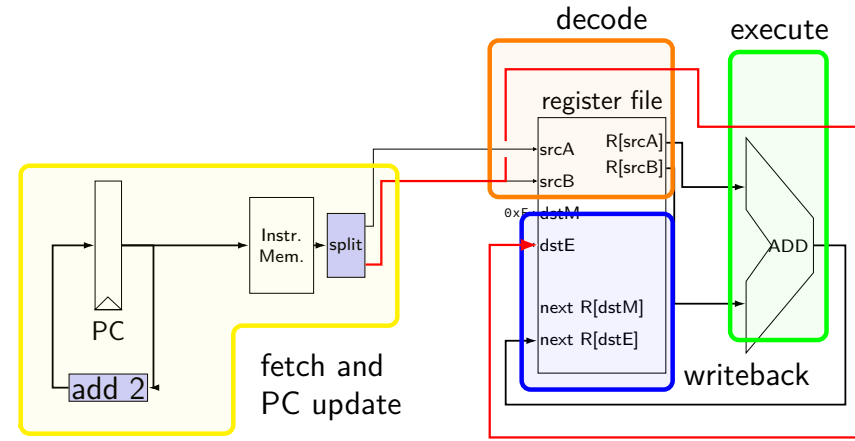


22

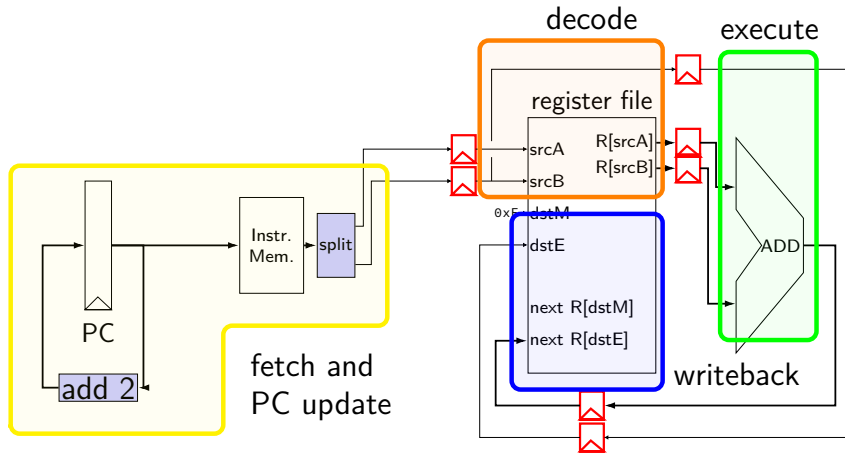
addq CPU



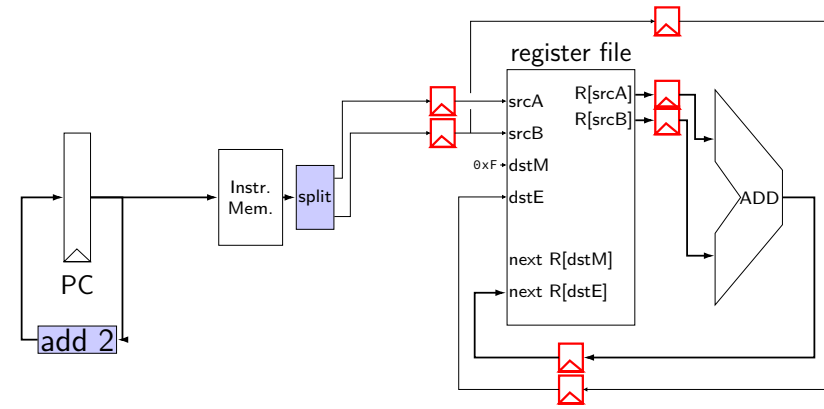
addq CPU



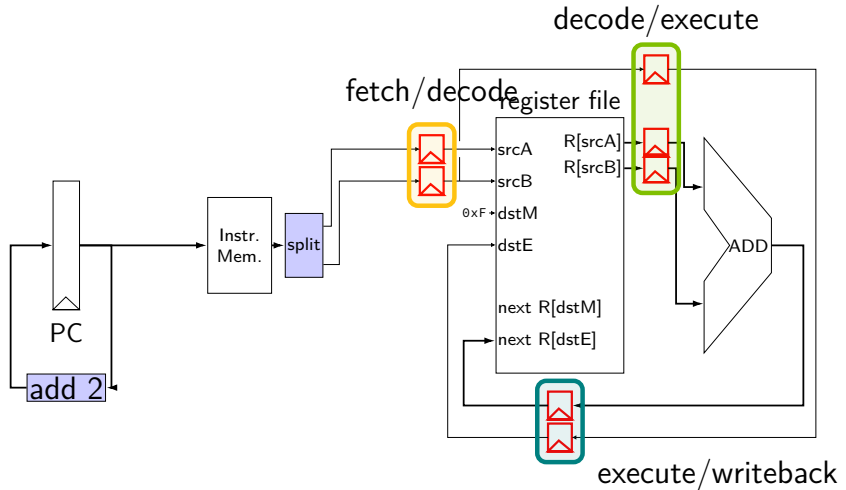
pipelined addq processor



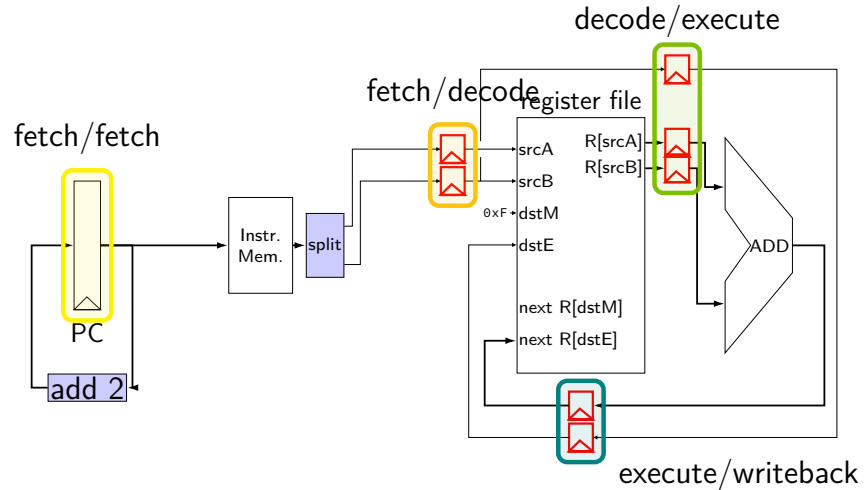
pipelined addq processor



pipelined addq processor

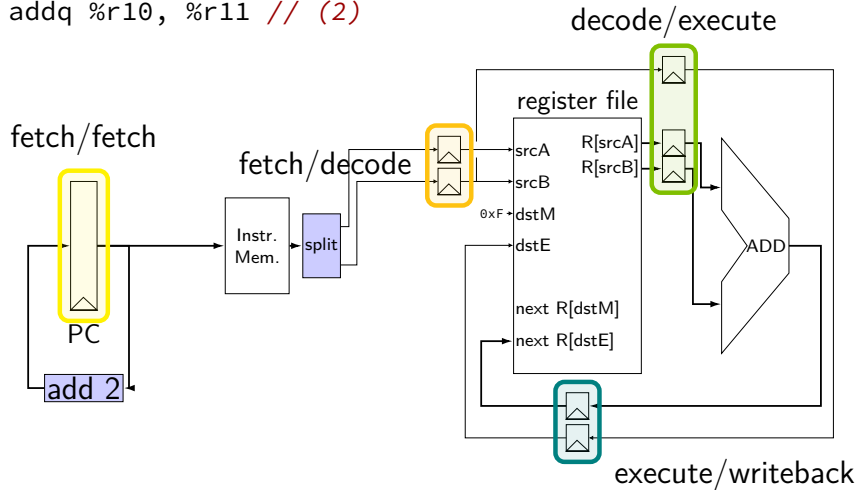


pipelined addq processor



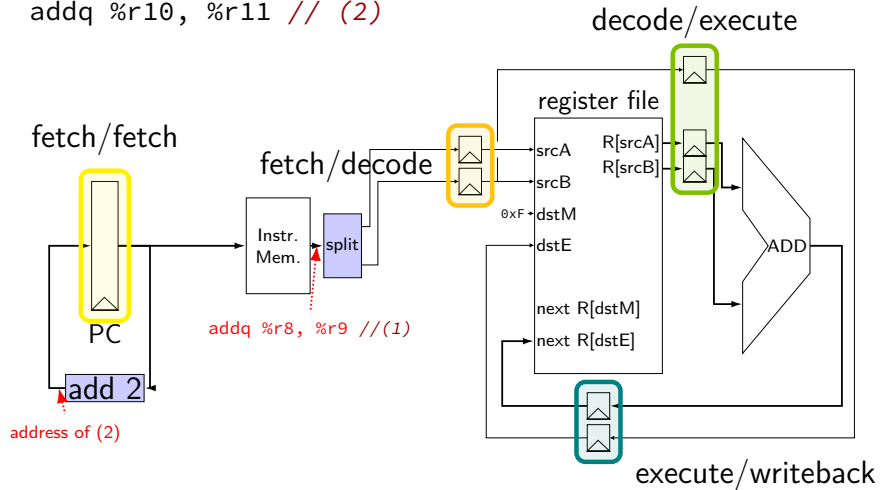
addq execution

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



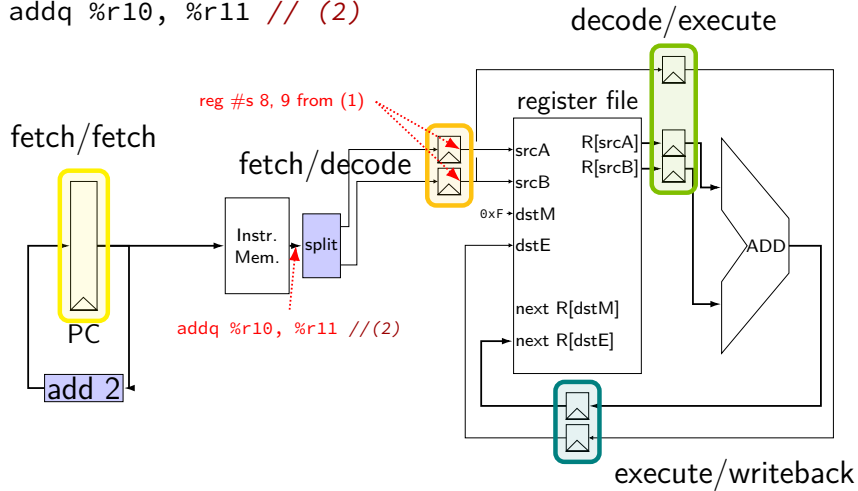
addq execution

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



addq execution

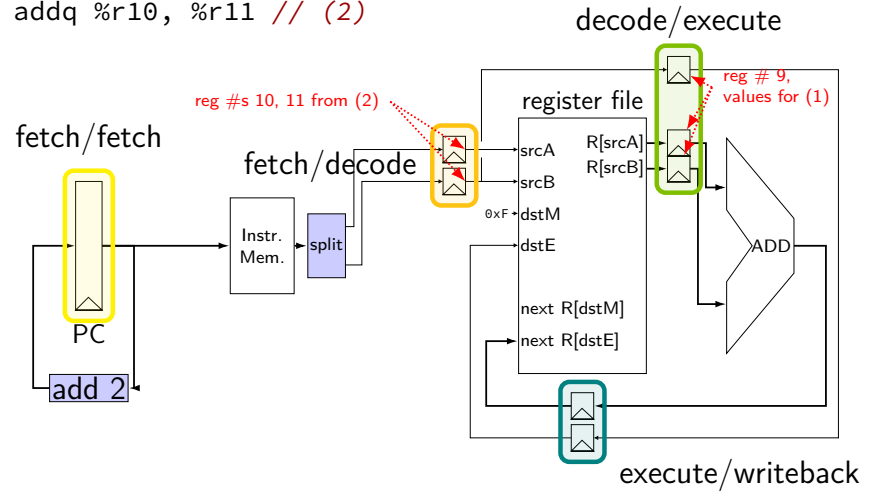
addq %r8, %r9 // (1)
 addq %r10, %r11 // (2)



24

addq execution

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




24

addq processor timing

// initially %r8 = 800,
 // %r9 = 900, etc.

addq %r8, %r9
 addq %r10, %r11
 addq %r12, %r13
 addq %r9, %r8




	fetch	fetch/decode	decode/execute			execute/writeback		
cycle	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

25

addq processor timing

// initially %r8 = 800,
 // %r9 = 900, etc.

addq %r8, %r9
 addq %r10, %r11
 addq %r12, %r13
 addq %r9, %r8




	fetch	fetch/decode	decode/execute			execute/writeback		
cycle	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

25

addq processor timing

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




	fetch	fetch/decode		decode/execute			execute/writeback	
cycle	PC	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstM]	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

25

addq processor timing

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




	fetch	fetch/decode		decode/execute			execute/writeback	
cycle	PC	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstM]	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

25

addq processor timing

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




	fetch	fetch/decode		decode/execute			execute/writeback	
cycle	PC	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstM]	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

25

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)

26