

last time

pipelining idea

- laundry analogy

- divide work into *stages*

- first: instruction A does stage 1

- then: instruction A does stage 2 while instruction B does stage 1

- ...

latency (start to finish) versus throughput (rate of finishing)

pipelining with circuits

- send values from one stage to next only

- add registers between each pair of stages

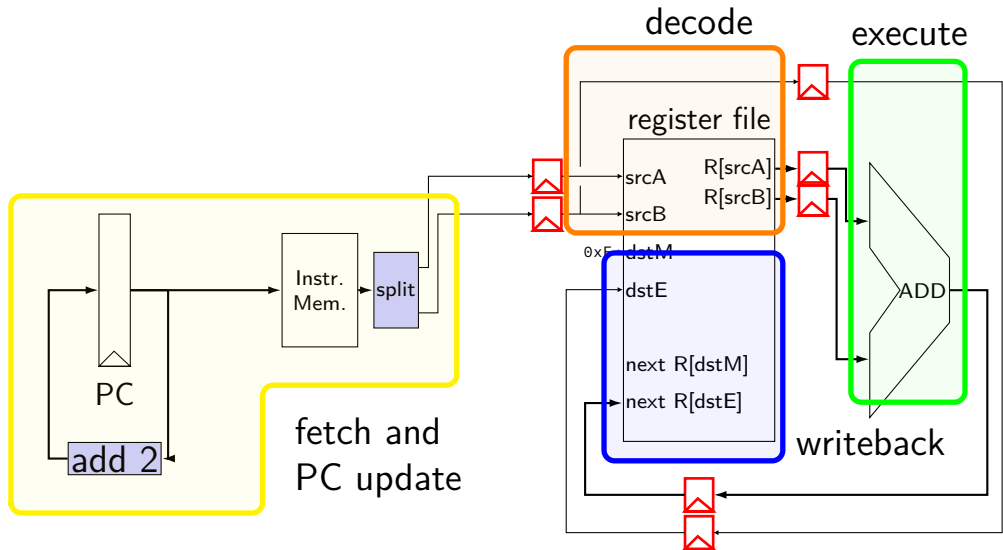
diminishing returns in pipelining

- only as fast as slowest stage

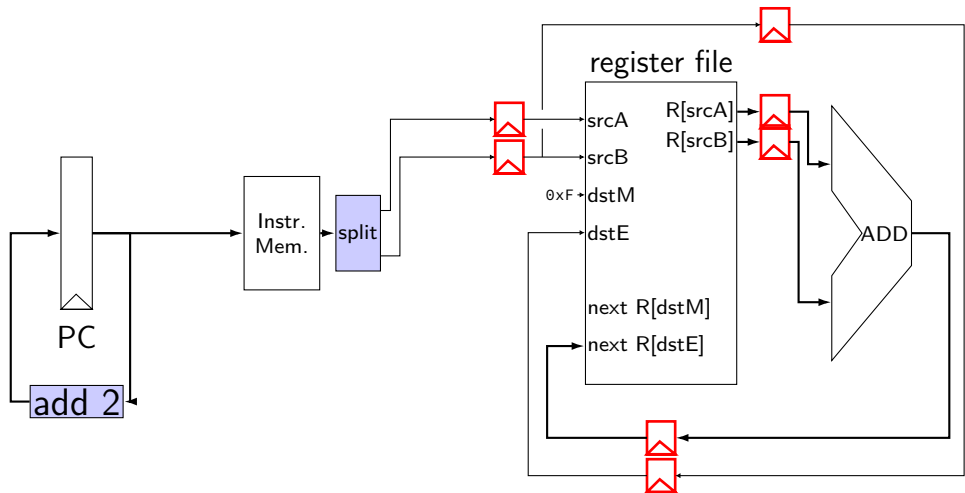
- cannot divide work evenly in practice

- registers take time

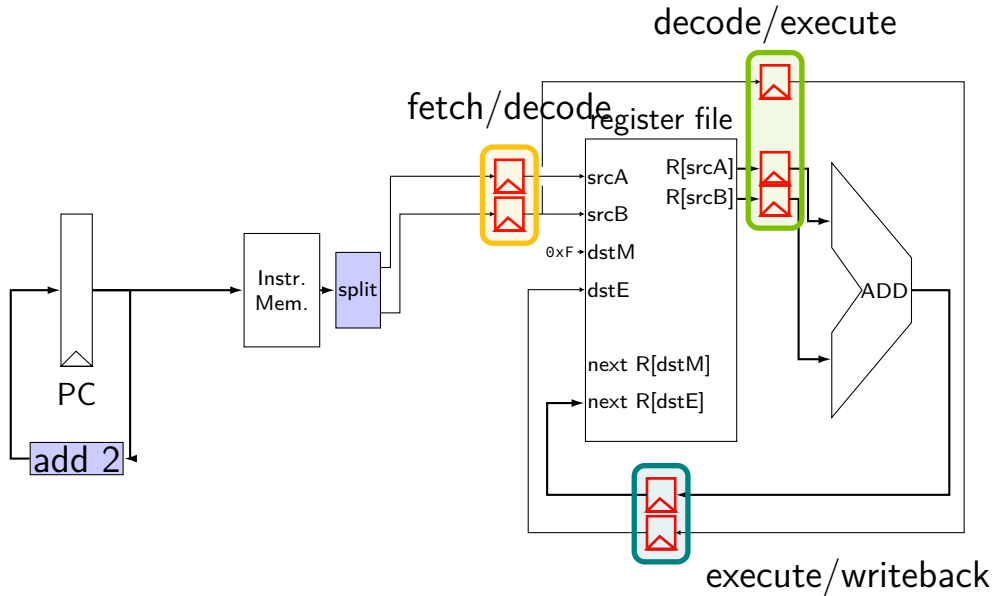
pipelined addq processor



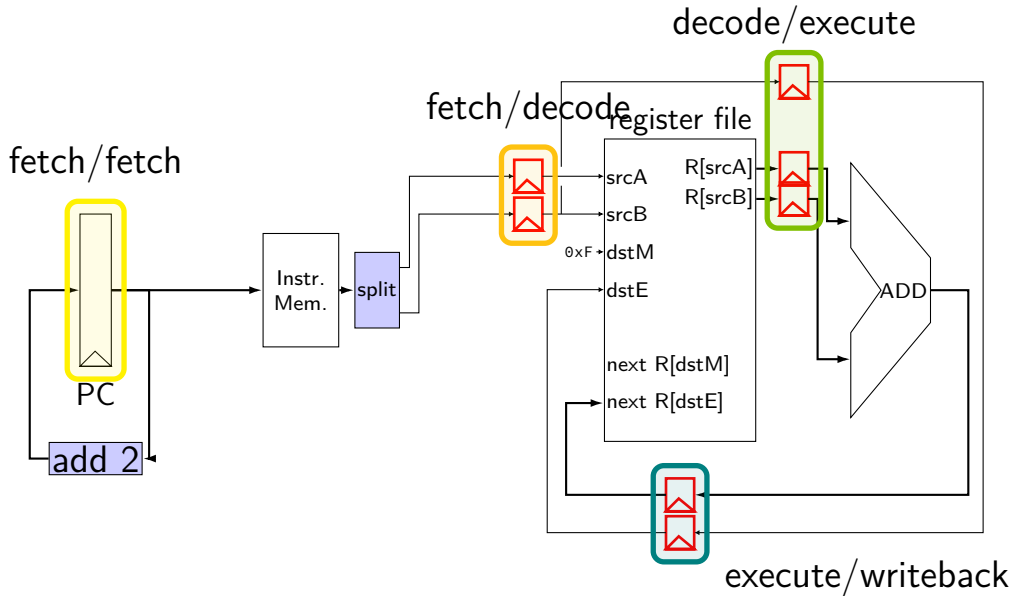
pipelined addq processor



pipelined addq processor

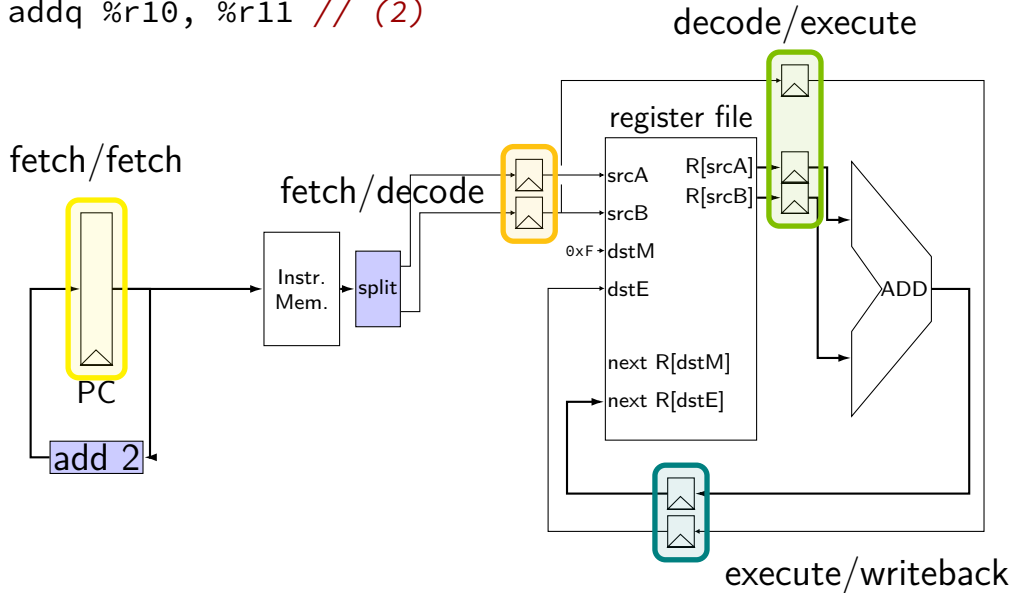


pipelined addq processor



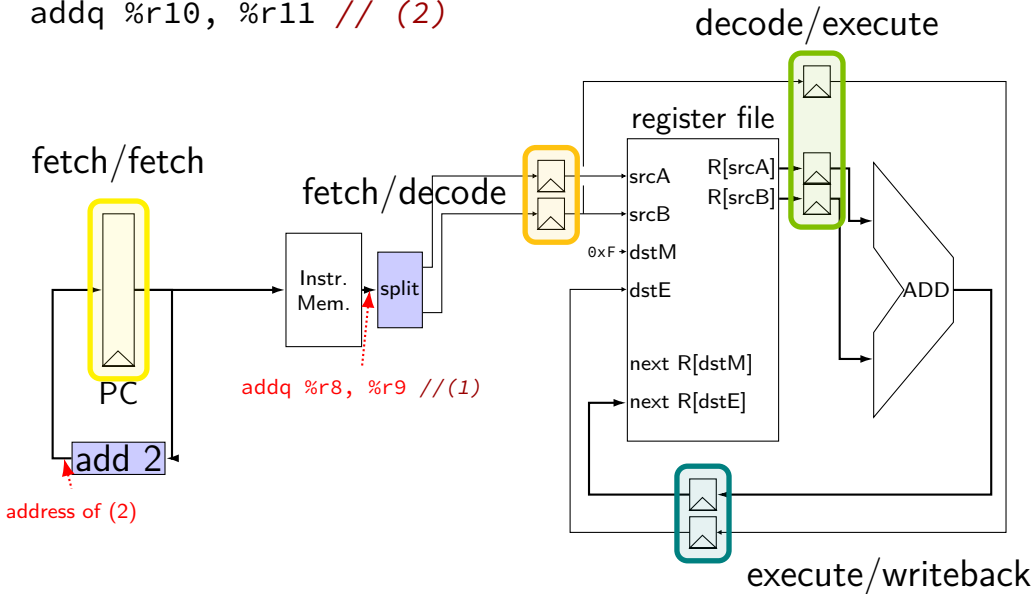
addq execution

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



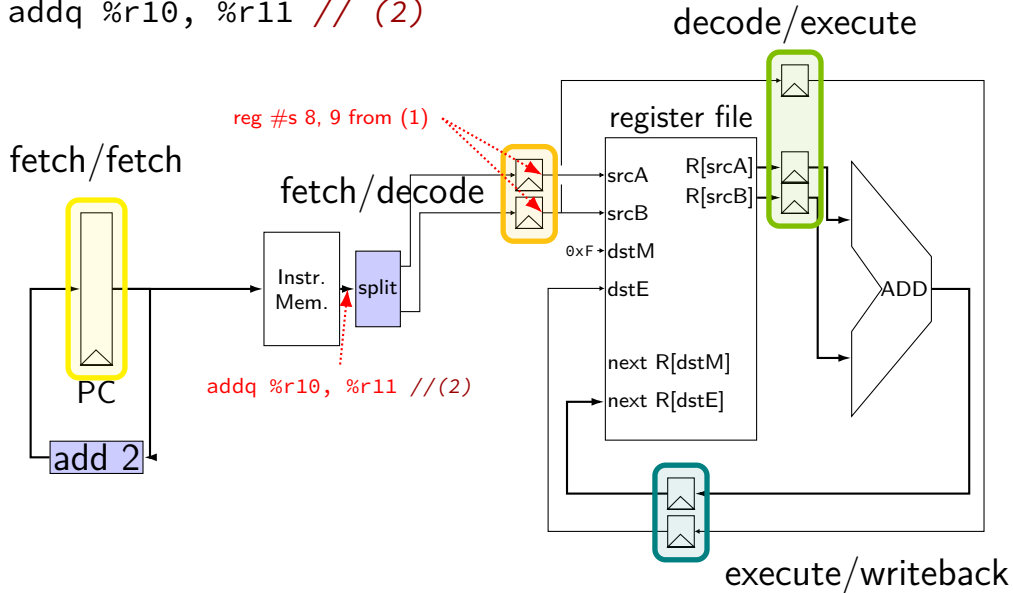
addq execution

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
```



	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

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

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

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

critical path

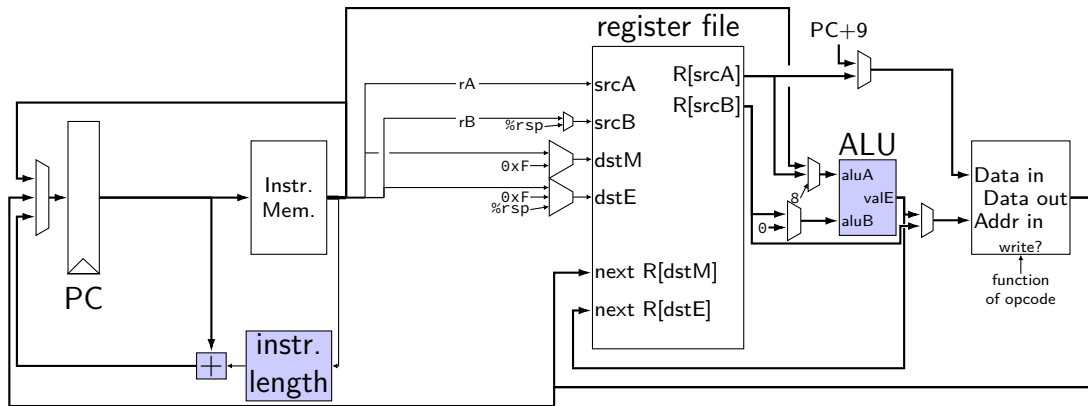
every path from state output to state input needs enough time
output — may change on rising edge of clock
input — must be stable sufficiently before rising edge of clock

critical path: **slowest** of all these paths — determines cycle time
times three: slowest stage ended up mattering

have to choose *one* clock cycle length
can't vary clock depending on what instruction is running

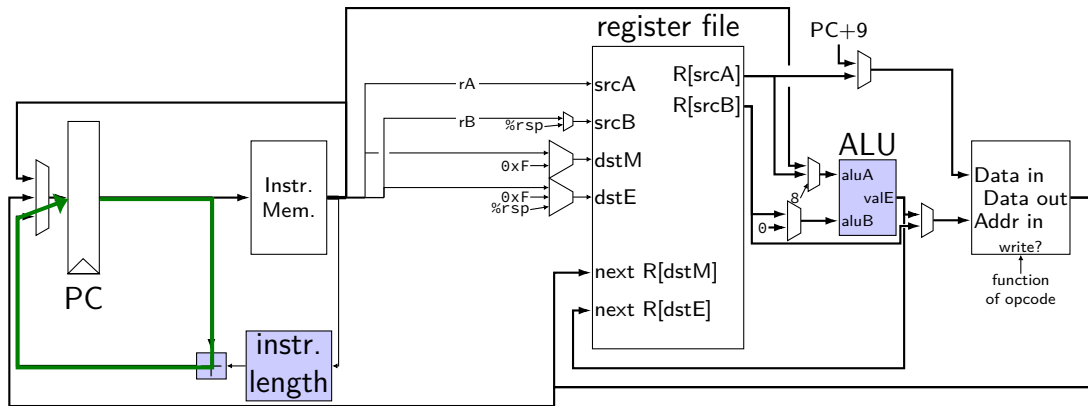
matters with or without pipelining

SEQ paths



SEQ paths

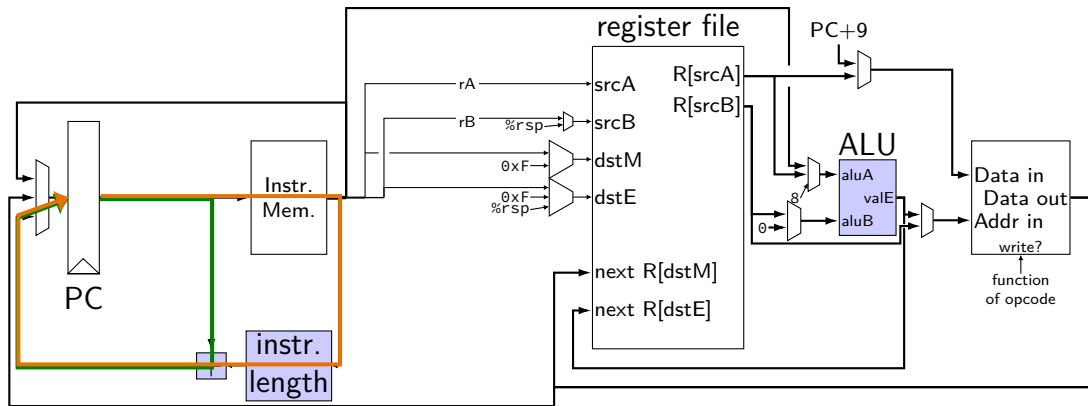
path 1: 25 picoseconds



SEQ paths

path 1: 25 picoseconds

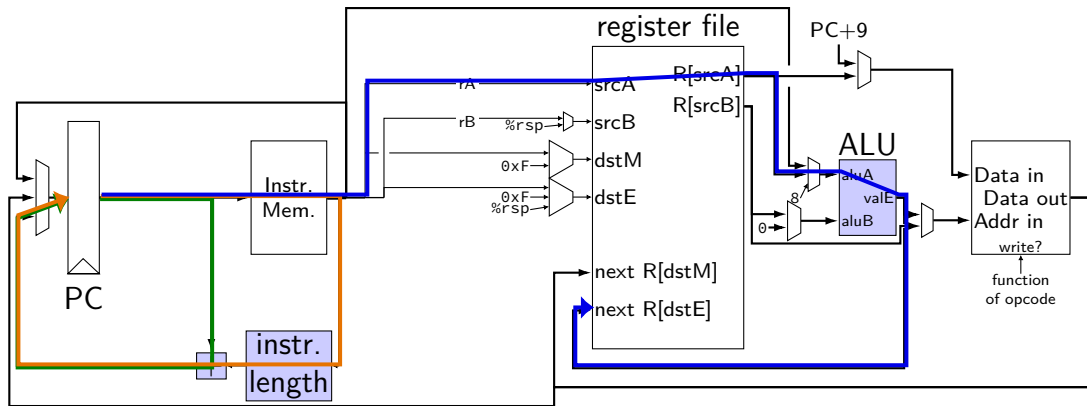
path 2: 50 picoseconds



SEQ paths

path 1: 25 picoseconds
path 3: 400 picoseconds

path 2: 50 picoseconds



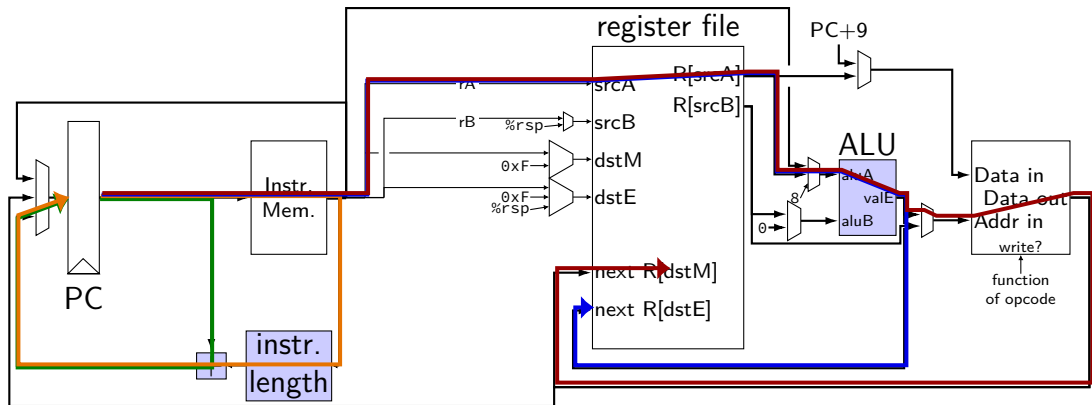
SEQ paths

path 1: 25 picoseconds
path 3: 400 picoseconds

path 2: 50 picoseconds
path 4: 900 picoseconds

...

...



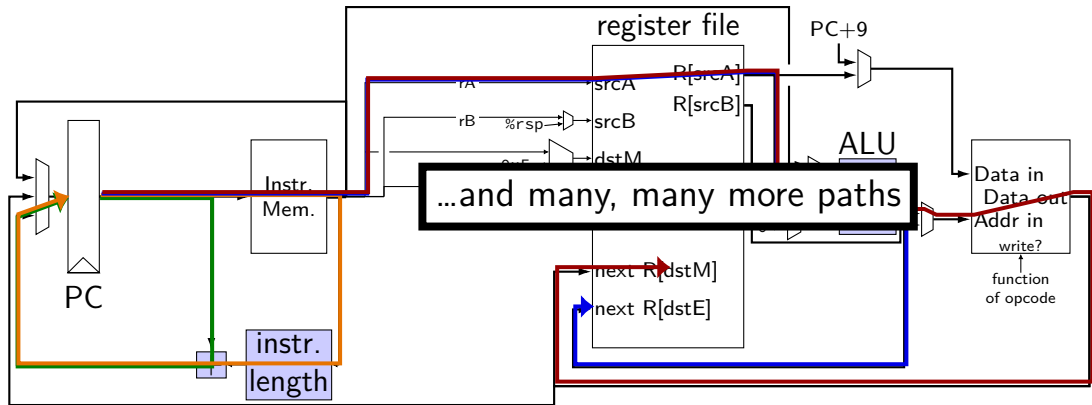
SEQ paths

path 1: 25 picoseconds
path 3: 400 picoseconds

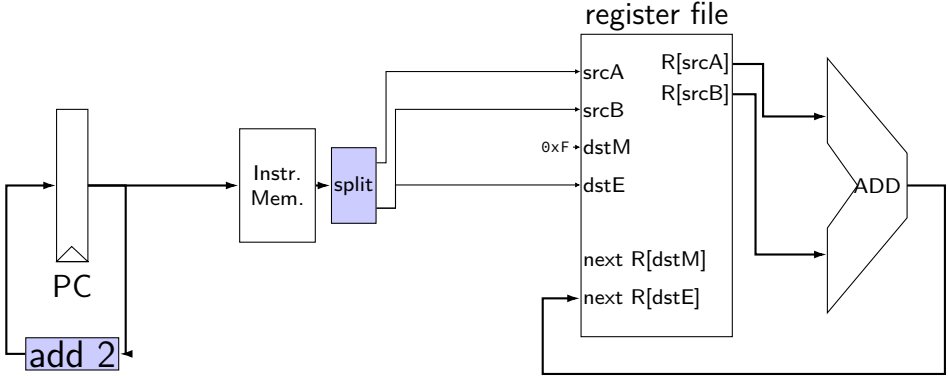
path 2: 50 picoseconds
path 4: 900 picoseconds

...

...

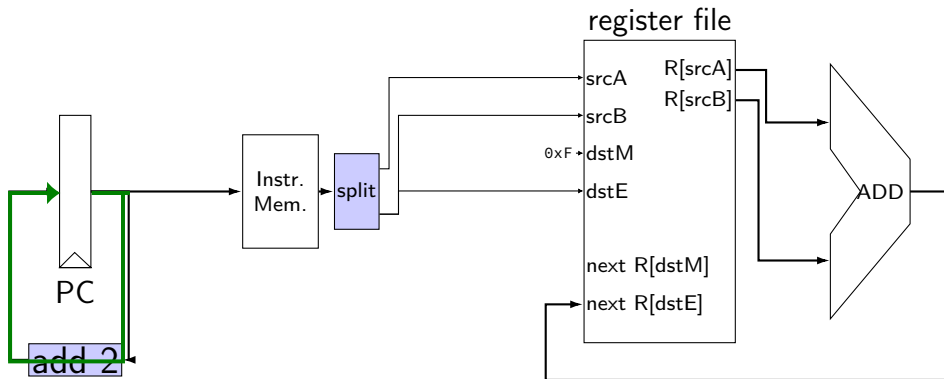


sequential addq paths



sequential addq paths

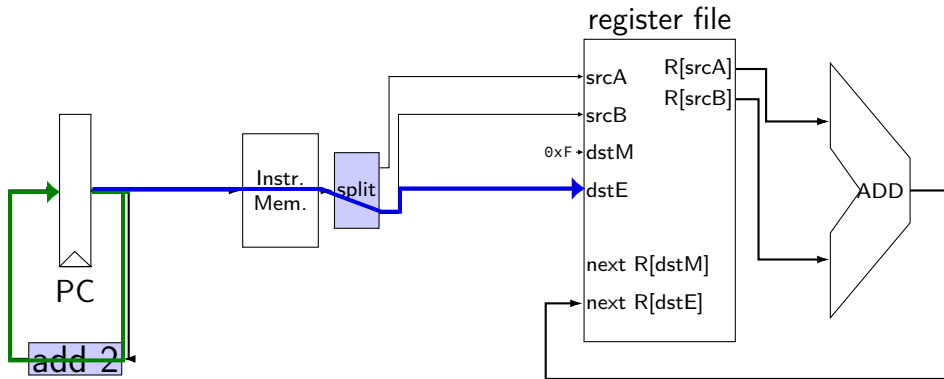
path 1: 25 picoseconds



sequential addq paths

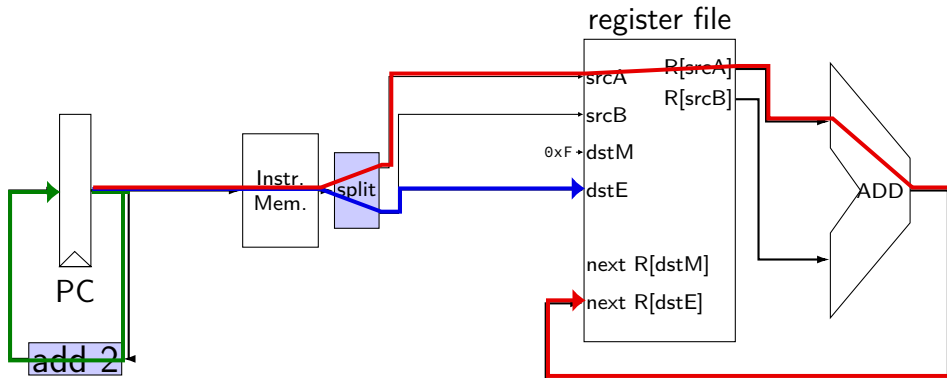
path 1: 25 picoseconds

path 2: 375 picoseconds



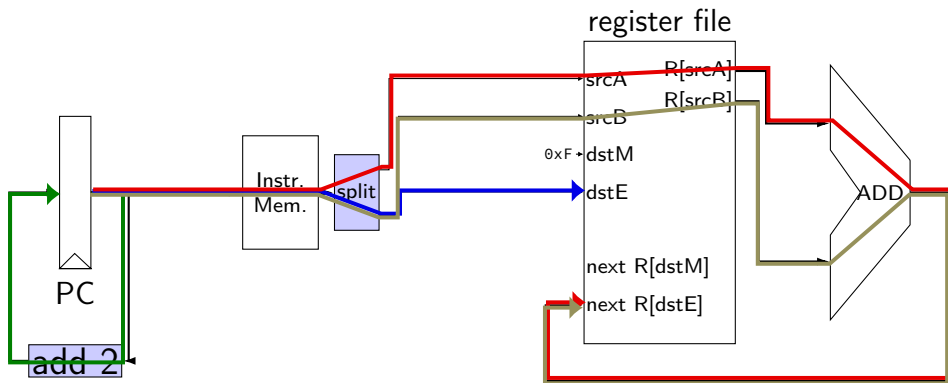
sequential addq paths

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



sequential addq paths

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



sequential addq paths

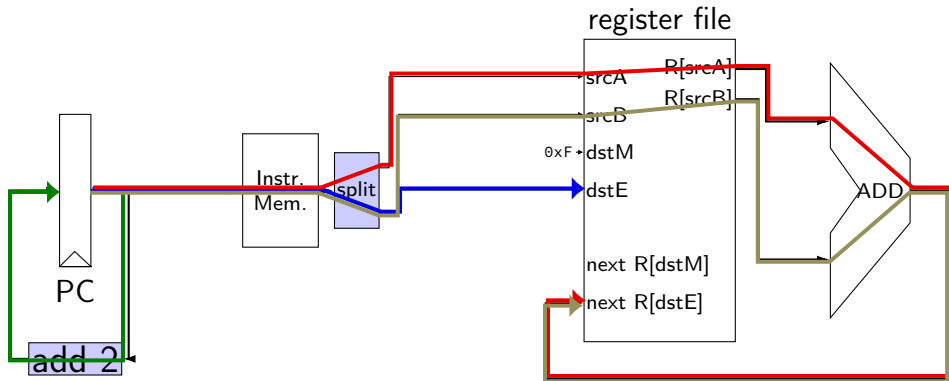
path 1: 25 picoseconds

path 2: 375 picoseconds

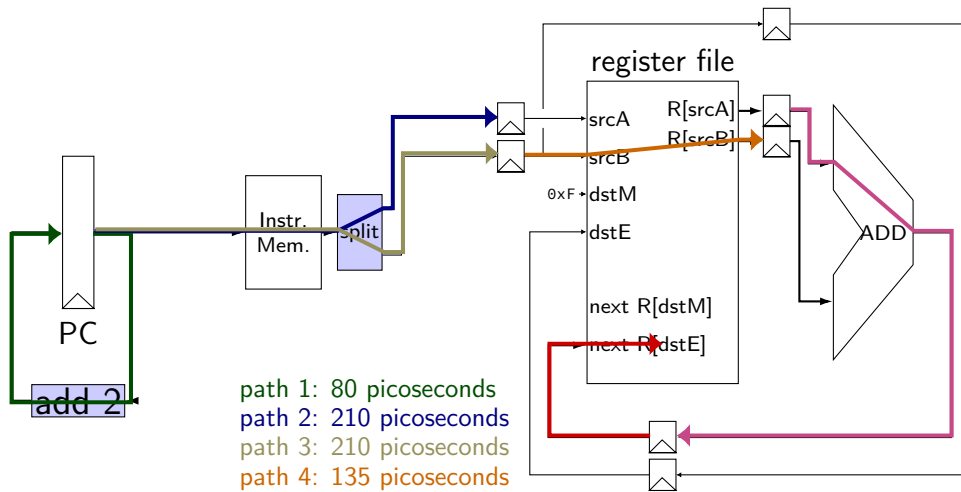
path 3: **500 picoseconds**

path 4: **500 picoseconds**

overall cycle time: **500 picoseconds** (longest path)



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

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**

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

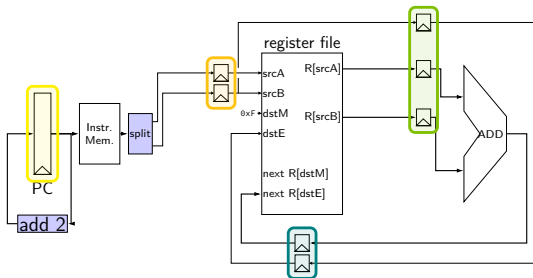
slowest path through stage + pipeline register delays

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

addq processor timing exercise 1

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



exercise 1: when instruction 1 stores its result in %rbx, what did instruction 3 *just complete*?

```
addq %rax, %rbx /* 1 */  
addq %rcx, %rdx /* 2 */  
addq %r8, %r9  /* 3 */
```


- A.** fetch **B.** decode (register read) **C.** execute
D. writeback **E.** nothing (it's not running)

addq processor performance exercise 2

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

now, cycle time = 200 ps + register delays per cycle



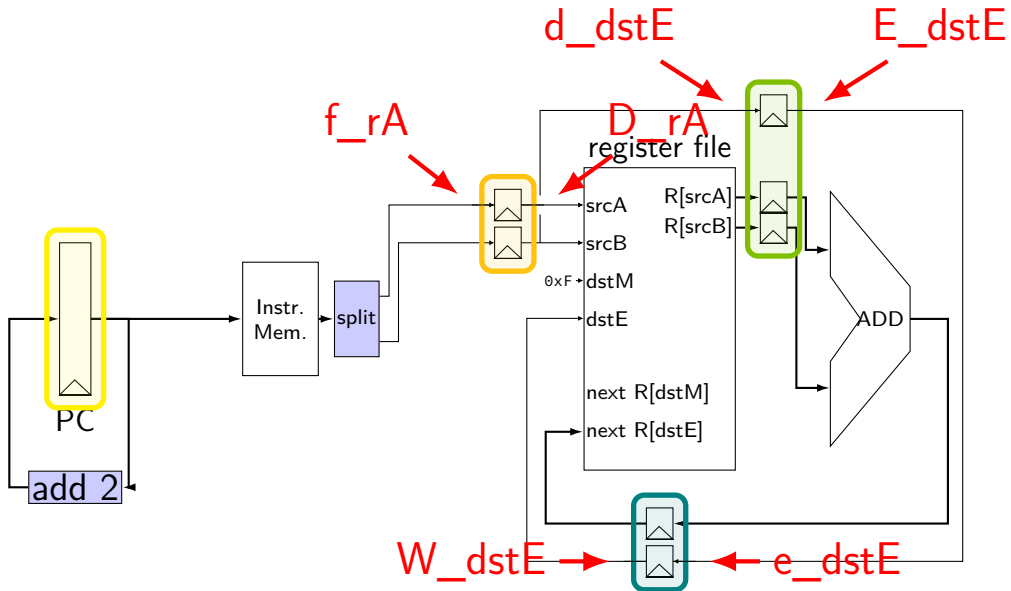
exercise 2:

suppose we combine add and register file read together — new cycle time?

- A.** 200 ps + reg delay
- C.** 250 ps + reg delay
- E.** 325 ps + reg delay

- B.** 225 ps + reg delay
- D.** 280 ps + reg delay
- E.** something else

pipeline register naming convention



pipeline register naming convention

f — fetch sends values here

D — decode receives values here

d — decode sends values here

...

addq HCL

```
...
/* f: from fetch */
f_rA = i10bytes[12..16];
f_rB = i10bytes[8..12];

/* 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;
}
```

```
...
```

addq fetch/decode

unpipelined

```
/* Fetch+PC Update*/  
pc = P_pc;  
p_pc = pc + 2;  
rA = i10bytes[12..16];  
rB = i10bytes[8..12];  
/* Decode */  
reg_srcA = rA;  
reg_srcB = rB;  
reg_dstE = rB;  
valA = reg_outputA;  
valB = reg_outputB;
```

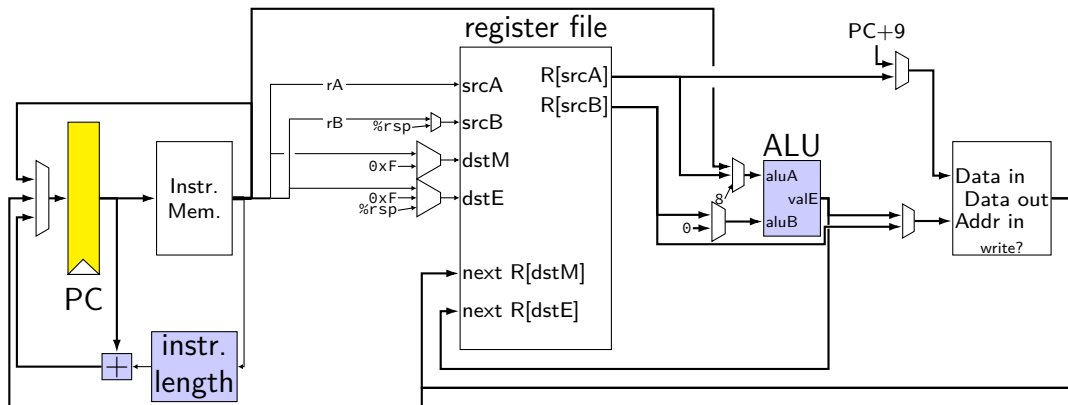
pipelined

```
/* Fetch+PC Update*/  
pc = P_pc;  
p_pc = pc + 2;  
f_rA = i10bytes[12..16];  
f_rB = i10bytes[8..12];  
/* Decode */  
reg_srcA = D_rA;  
reg_srcB = D_rB;  
d_dstE = D_rB;  
d_valA = reg_outputA;  
d_valB = reg_outputB;
```

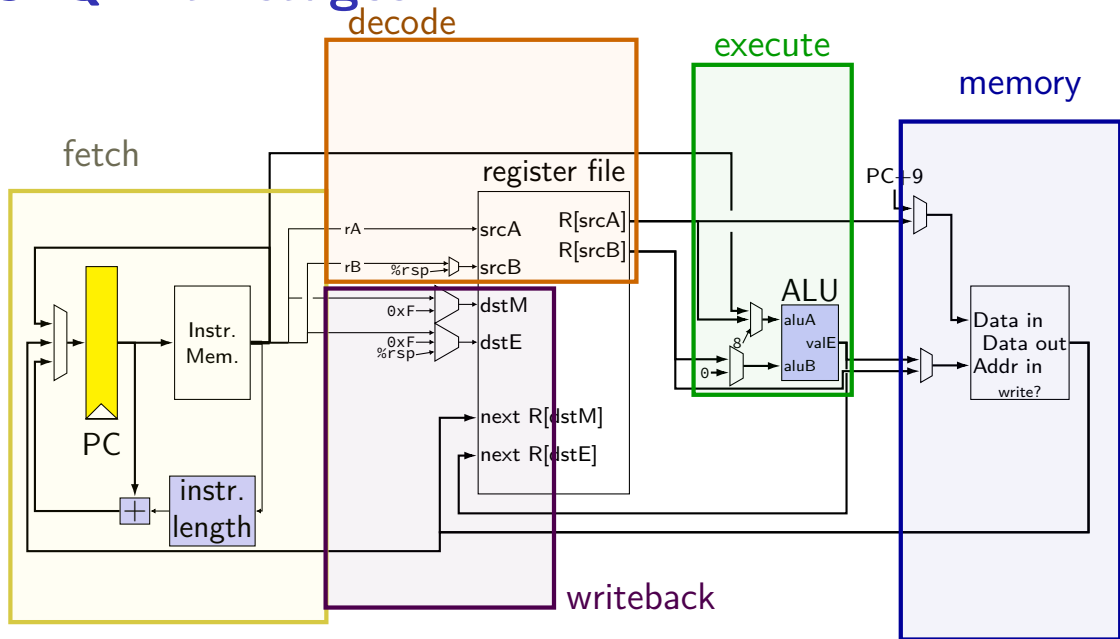
addq pipeline registers

```
register pP {
    pc : 64 = 0;
};
/* Fetch+PC Update*/
register fD {
    rA : 4 = REG_NONE; rB : 4 = REG_NONE;
};
/* Decode */
register dE {
    valA : 64 = 0; valB : 64 = 0; dstE : 4 = REG_NONE;
}
/* Execute */
register eW {
    valE : 64 = 0; dstE : 4 = REG_NONE;
}
/* Writeback */
```

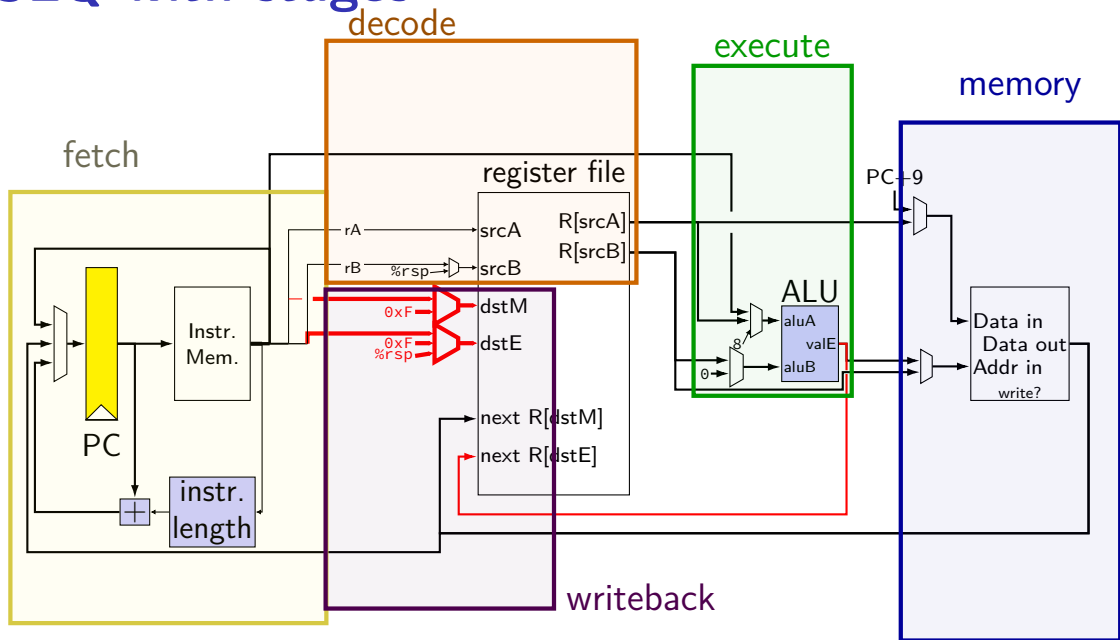
SEQ without stages



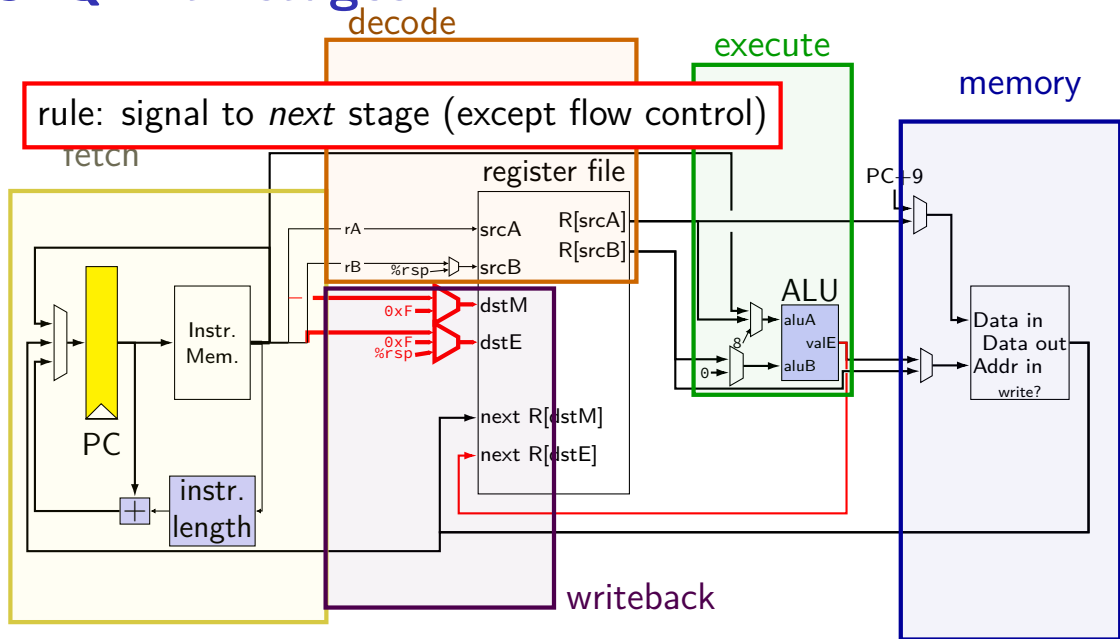
SEQ with stages



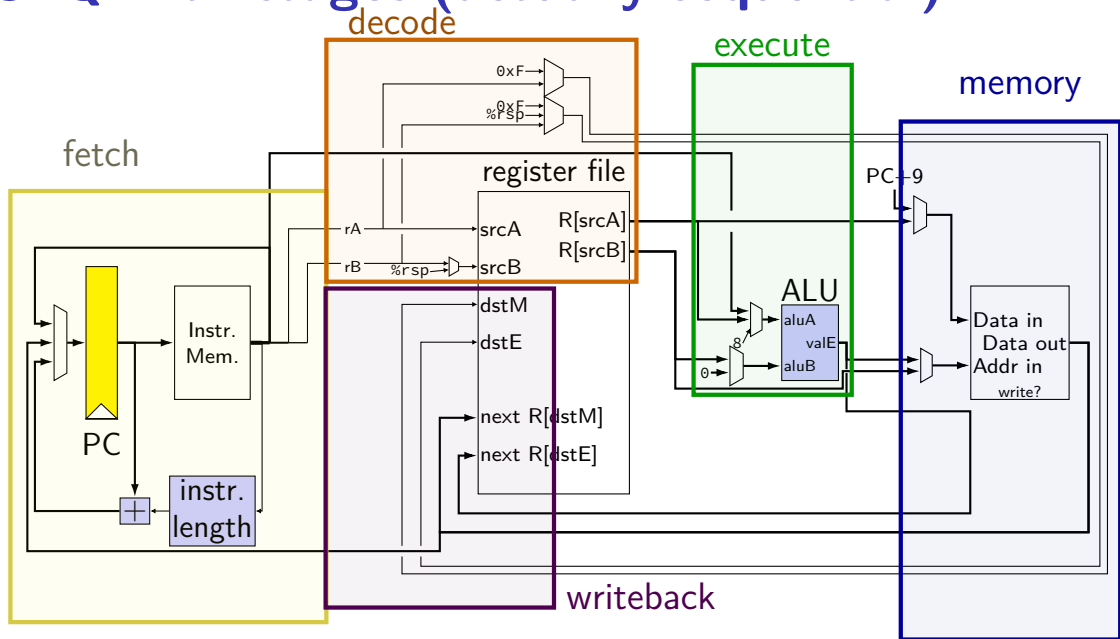
SEQ with stages



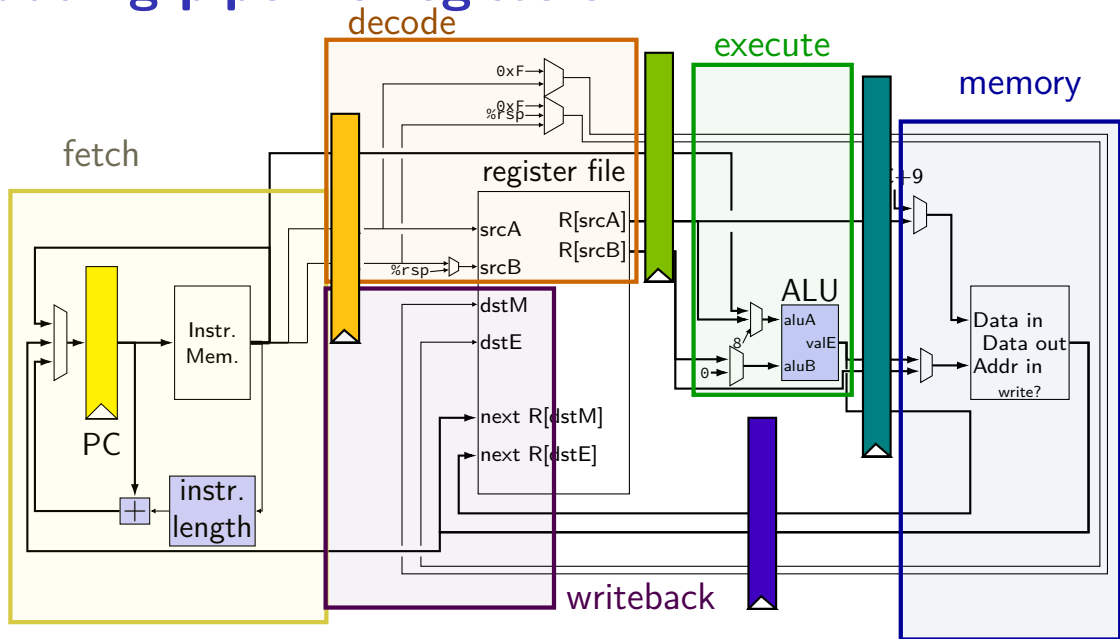
SEQ with stages




SEQ with stages (actually sequential)



adding pipeline registers



adding pipeline registers



passing values in pipeline

read **prior stage's outputs**

e.g. decode: get from fetch via pipeline registers (D_icode, ...)

send **inputs for next stage**

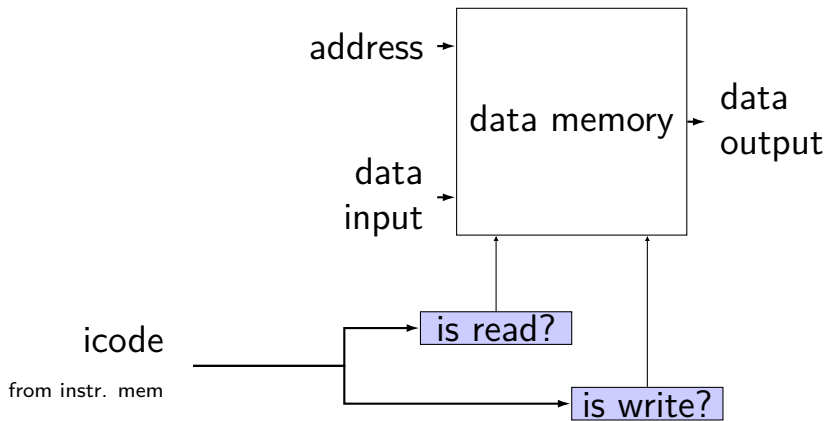
e.g. decode: send to execute via pipeline registers (d_icode, ...)

exceptions: **deliberate sharing** between instructions


via register file/memory/etc.

via control flow instructions

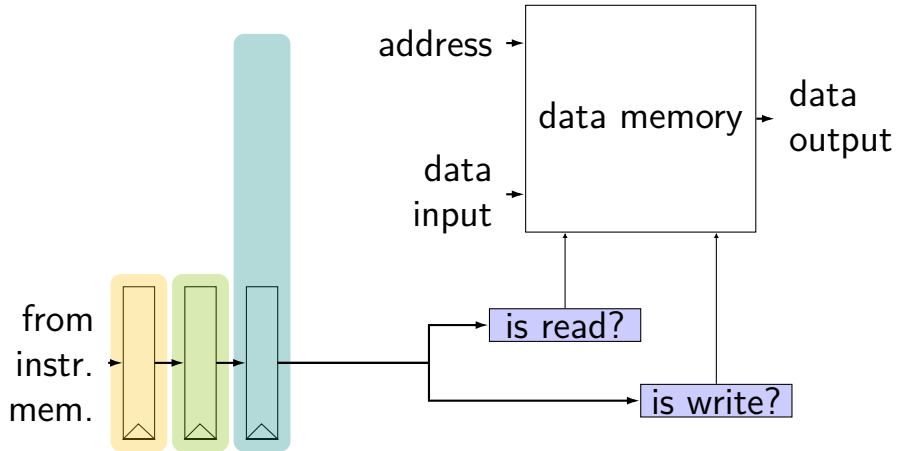
memory read/write logic



memory read/write logic



memory read/write logic



memory read/write: SEQ code

```
icode = i10bytes[4..8];  
mem_readbit = [  
    icode == MRMOVQ || ...: 1;  
    0;  
];
```

memory read/write: PIPE code

```
f_icode = i10bytes[4..8];
register fD { /* and dE and eM and mW */
    icode : 4 = NOP;
}
d_icode = D_icode;
...
e_icode = E_icode;
mem_readbit = [
    M_icode == MRMOVQ || ...: 1;
    0;
];
```

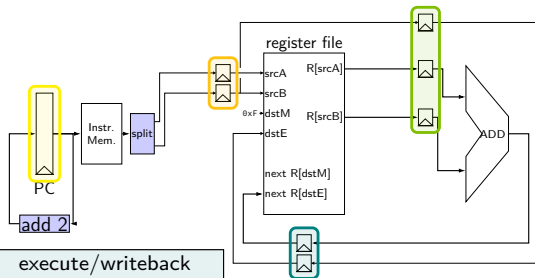
memory read/write: PIPE code

```
f_icode = i10bytes[4..8];  
register fD { /* and dE and eM and mW */  
    icode : 4 = NOP;  
}  
d_icode = D_icode;  
...  
e_icode = E_icode;  
mem_readbit = [  
    M_icode == MRMOVQ || ...: 1;  
    0;  
];
```

addq processor: data hazard

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

```
addq %r8, %r9  
addq %r9, %r8  
addq ...  
addq ...
```




	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		9	8	800	900	9		
3				900	800	8	1700	9
4							1700	8

addq processor: data hazard

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

```
addq %r8, %r9
addq %r9, %r8
addq ...
addq ...
```



	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		9	8	800	900	9		
3				900	800	8	1700	9
4							1700	8

should be 1700

data hazard

addq %r8, %r9 // (1)

addq %r9, %r8 // (2)

step#	pipeline implementation	ISA specification
1	read r8, r9 for (1)	read r8, r9 for (1)
2	read r9, r8 for (2)	write r9 for (1)
3	write r9 for (1)	read r9, r8 for (2)
4	write r8 for (2)	write r8 ror (2)

pipeline reads **older value**...

instead of value ISA says was just written

data hazard compiler solution

```
addq %r8, %r9  
nop  
nop  
addq %r9, %r8
```

one solution: **change the ISA**

all addqs take effect **three instructions later**

make it **compiler's job**

problem: recompile everytime processor changes?

data hazard hardware solution

```
addq %r8, %r9  
// hardware inserts: nop  
// hardware inserts: nop  
addq %r9, %r8
```

how about hardware add nops?

called **stalling**

extra logic:

- sometimes don't change PC

- sometimes put do-nothing values in pipeline registers

addq processor: data hazard stall

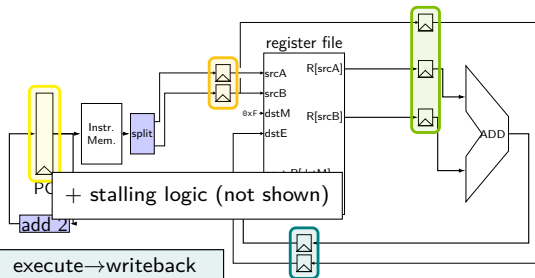
```
// initially %r8 = 800,
//                %r9 = 900, etc.
```

```
addq %r8, %r9
```

```
// hardware stalls twice
```

```
addq %r9, %r8
```

```
addq %r10, %r11
```



	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	0x2*	F	F	800	900	9		
3	0x2	F	F	---	---	F	1700	9
4	0x4	9	8	---	---	F	---	F
5		10	11	1700	800	8	---	F
6				1000	1100	11	2500	8

addq processor: data hazard stall


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

```
addq %r8, %r9
```

```
// hardware stalls twice
```

```
addq %r9, %r8
```

```
addq %r10, %r11
```



	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	0x2*	F	F	800	900	9		
3	0x2	F	F	---	---	F	1700	9
4	0x4	9	8	---	---	F	---	F
5		10	11	1700	800	8	---	F
6				1000	1100	11	2500	8

addq processor: data hazard stall


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

```
addq %r8, %r9
```

```
// hardware stalls twice
```

```
addq %r9, %r8
```

```
addq %r10, %r11
```




	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	0x2*	F	F	800	900	9		
3	0x2	F	F	---	---	F	1700	9
4	0x4	9	8	---	---	F	---	F
5		10	11	1700	800	8	---	F
6				1000	1100	11	2500	8

R[9] written during cycle 3; read during cycle 4

addq stall

```
addq %r8, %r9  
// hardware stalls twice  
addq %r9, %r8  
addq %r10, %r11
```



hazards versus dependencies

dependency — X needs result of instruction Y?


has potential for being messed up by pipeline
(since part of X may run before Y finishes)

hazard — will it not work in some pipeline?

before extra work is done to “resolve” hazards
multiple kinds: so far, *data hazards*

data hazard exercise

```
addq %r8, %r9
addq %r10, %r11
addq %r9, %r8
addq %r11, %r10
```




to resolve data hazards with stalling, how many stalls are needed?

hint: complete timeline

instr	cycle: 0	1	2	3	4	5	6	7
addq R8, R9	F	D	E	W				
addq R10, R11		F						
addq R9, R8								
addq R11, R10								

data hazard exercise solution



	cycle #	0	1	2	3	4	5	6	7	8
addq %r8, %r9		F	D	E	W					
addq %r10, %r11			F	D	E	W				
addq %r9, %r8				×	F	D	E	W		
addq %r11, %r10						F	D	E	W	

revisiting data hazards

stalling worked

but very unsatisfying — wait 2 extra cycles to use anything?!

...or more with 5-stage pipeline

observation: **value** ready before it would be needed

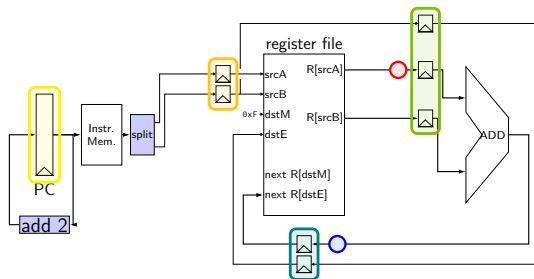
(just not stored in a way that let's us get it)

motivation

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

```
addq %r8, %r9
addq %r9, %r8
addq ...
addq ...
```

location of values during cycle 2:



	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		9	8	800	900	9		
3				900	800	8	1700	9
4							1700	8

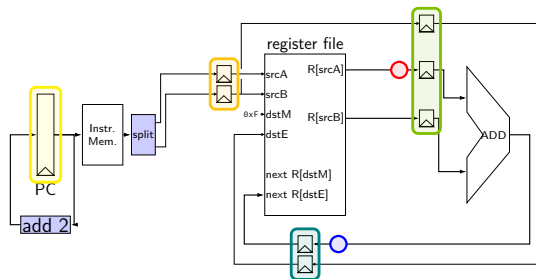
should be 1700

motivation

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

```
addq %r8, %r9
addq %r9, %r8
addq ...
addq ...
```

location of values during cycle 2:




	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		9	8	800	900	9		
3				900	800	8	1700	9
4							1700	8

should be 1700

forwarding

addq %r8, %r9 // (1)

addq %r9, %r8 // (2)




forwarding

addq %r8, %r9 // (1)

addq %r9, %r8 // (2)

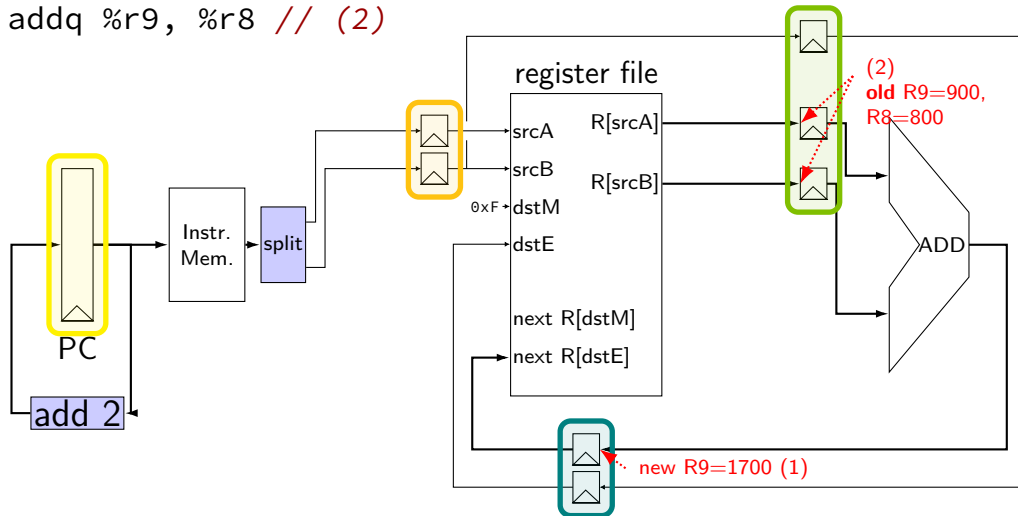
reg #s 9, 8 from (2)



forwarding

addq %r8, %r9 // (1)

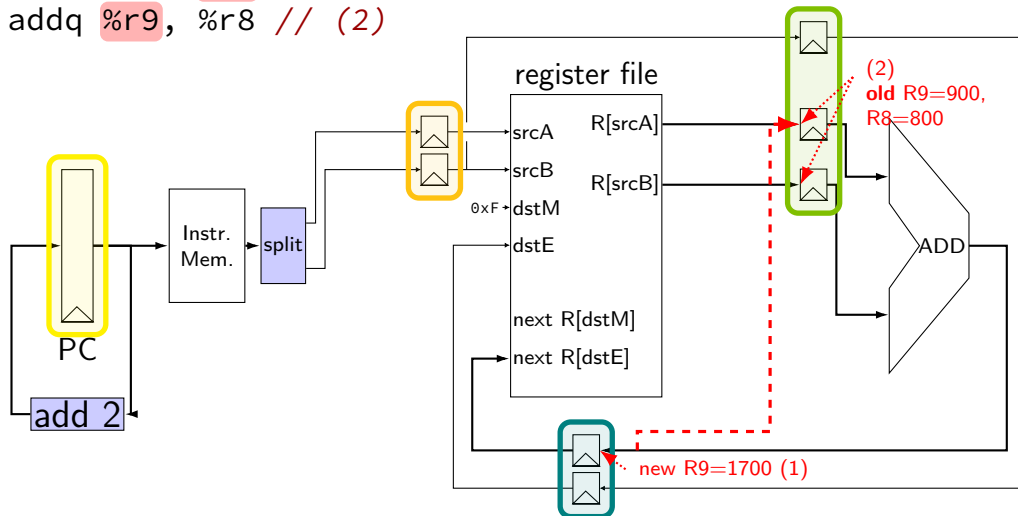
addq %r9, %r8 // (2)



forwarding

addq %r8, %r9 // (1)

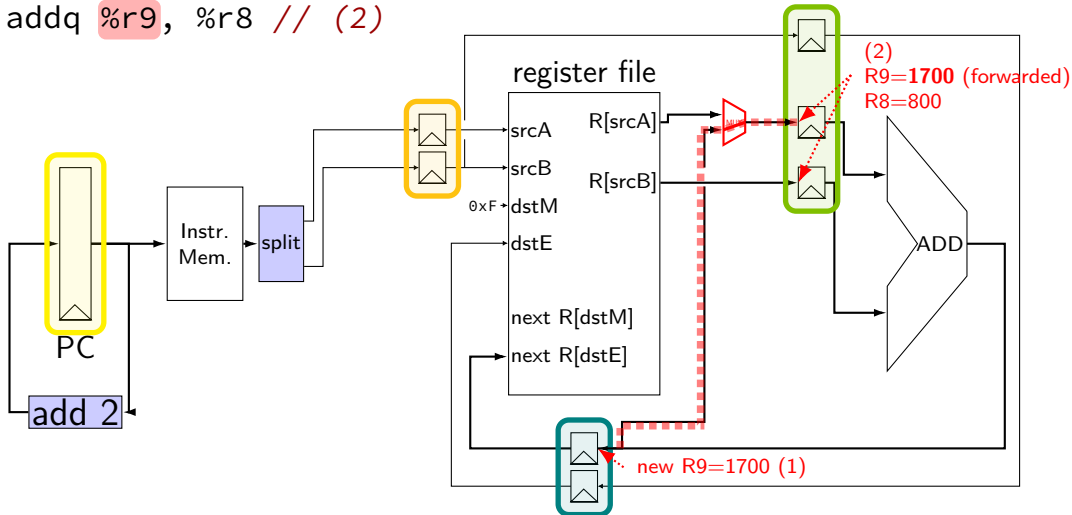
addq %r9, %r8 // (2)



forwarding

addq %r8, %r9 // (1)

addq %r9, %r8 // (2)




forwarding

addq %r8, %r9 // (1)

~~addq %r9, %r8 // (2)~~

addq %r10, %r9 // (2b)




backup slides

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)