pipelining 2: hazards

Changelog

6 October 2020: multiple forward paths (2): HCL corrected to use both reg_outputA and reg_outputB

last time

pipelining latency and throughput

diminsihing returns: register delays + uneven splitting

pipeline stages
- PC update with fetch — need next instruction ASAP

naming convention
- D_...— pipeline register outputs to decode stage
- d_...— pipeline register inputs from decode stage
- ...

adding pipelining to existing CPUs:
- redraw signals to not skip stages
- add registers to pass values from one stage to next

SEQ without stages

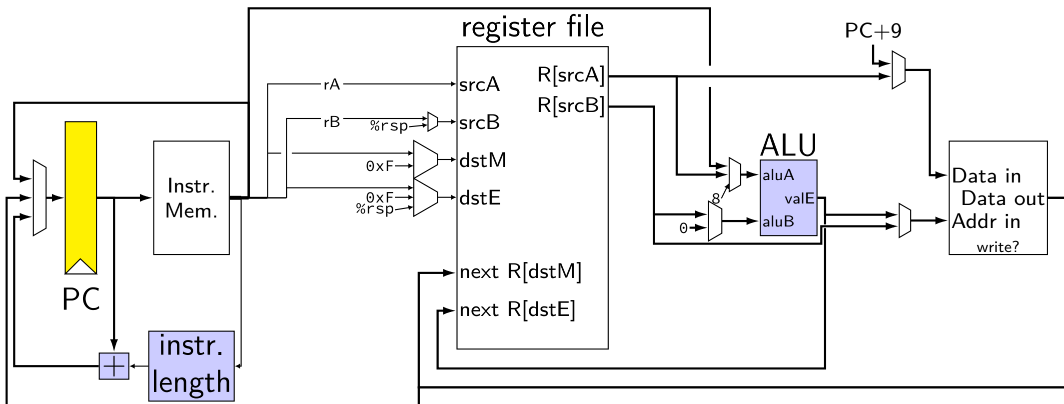

SEQ with stages

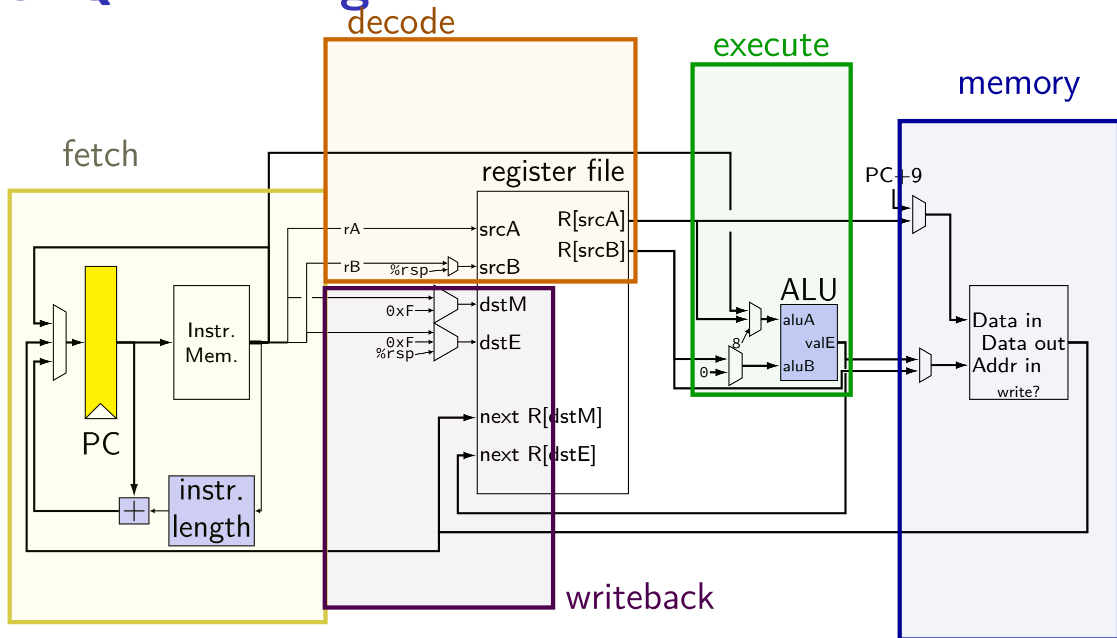

SEQ with stages

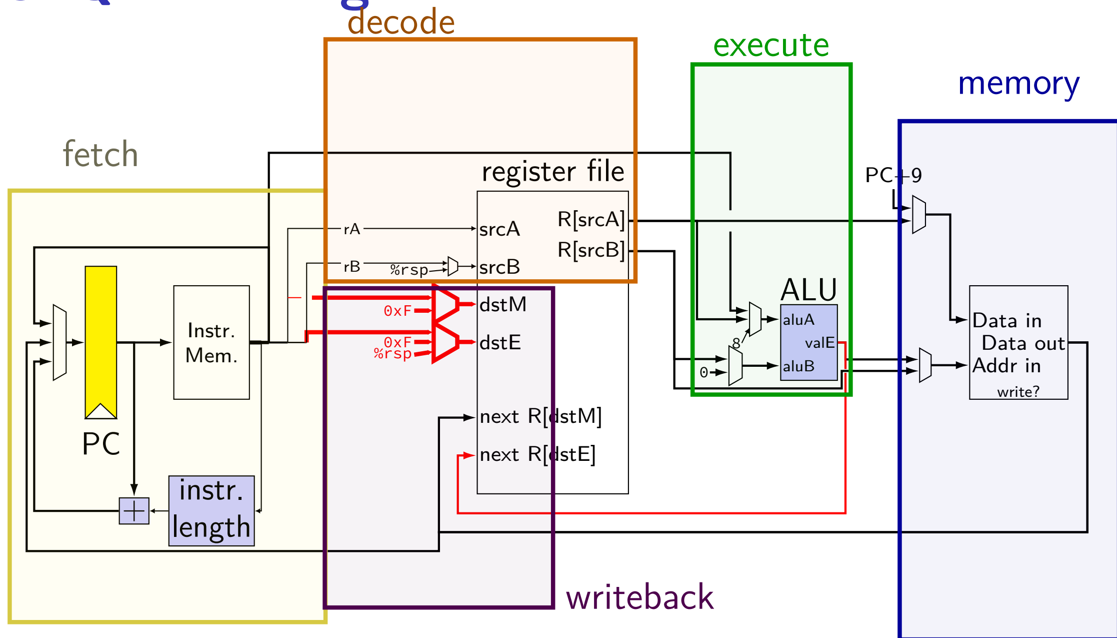

SEQ with stages

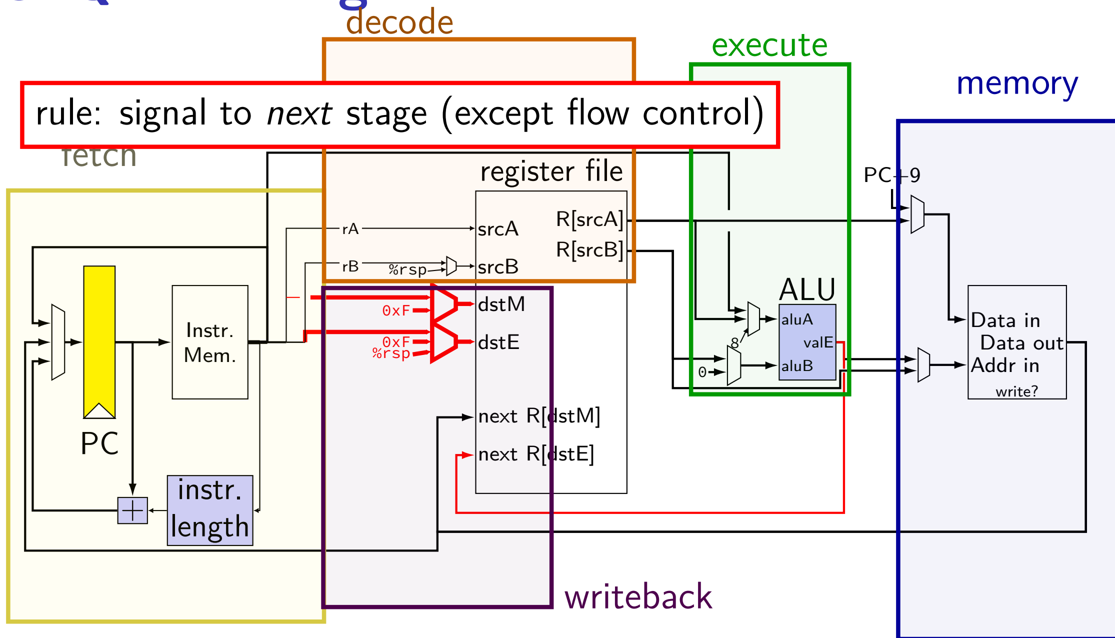

SEQ with stages (actually sequential)

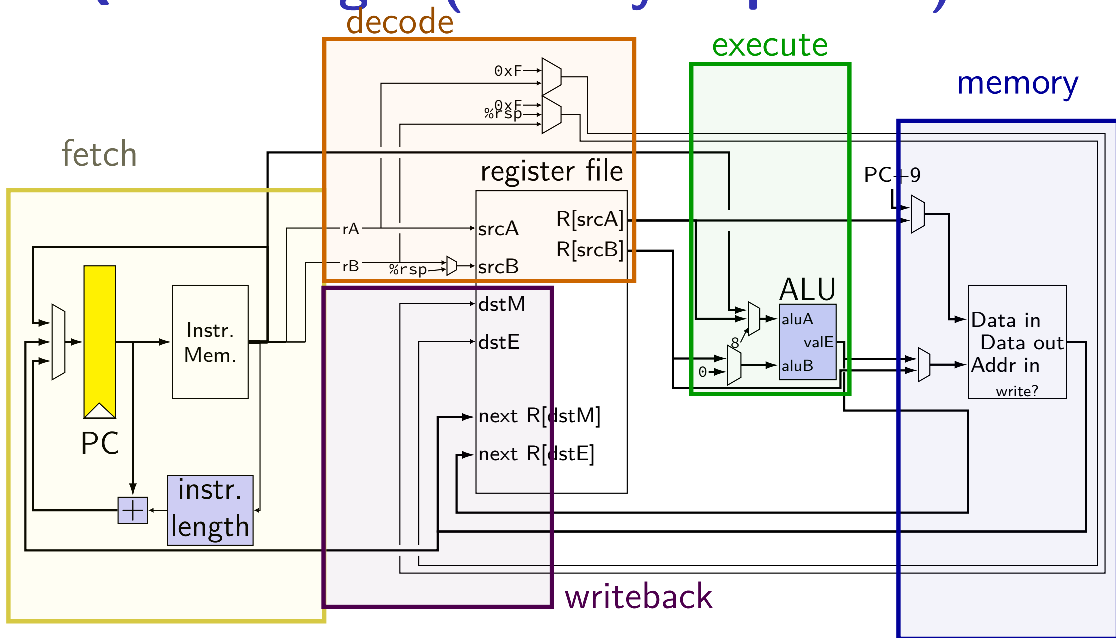

adding pipeline registers

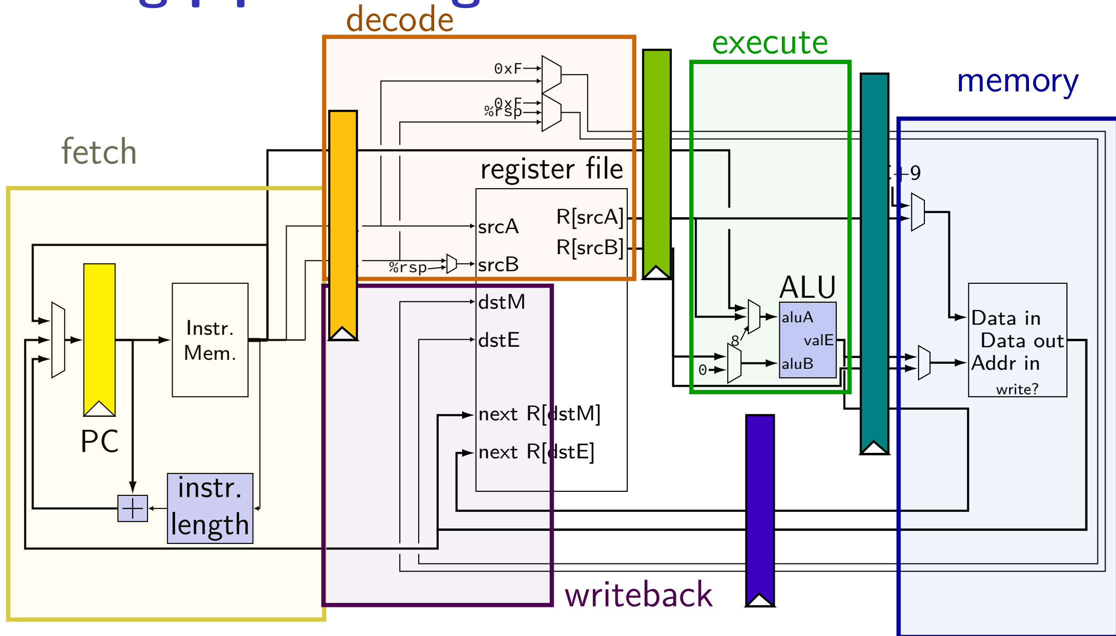

adding pipeline registers

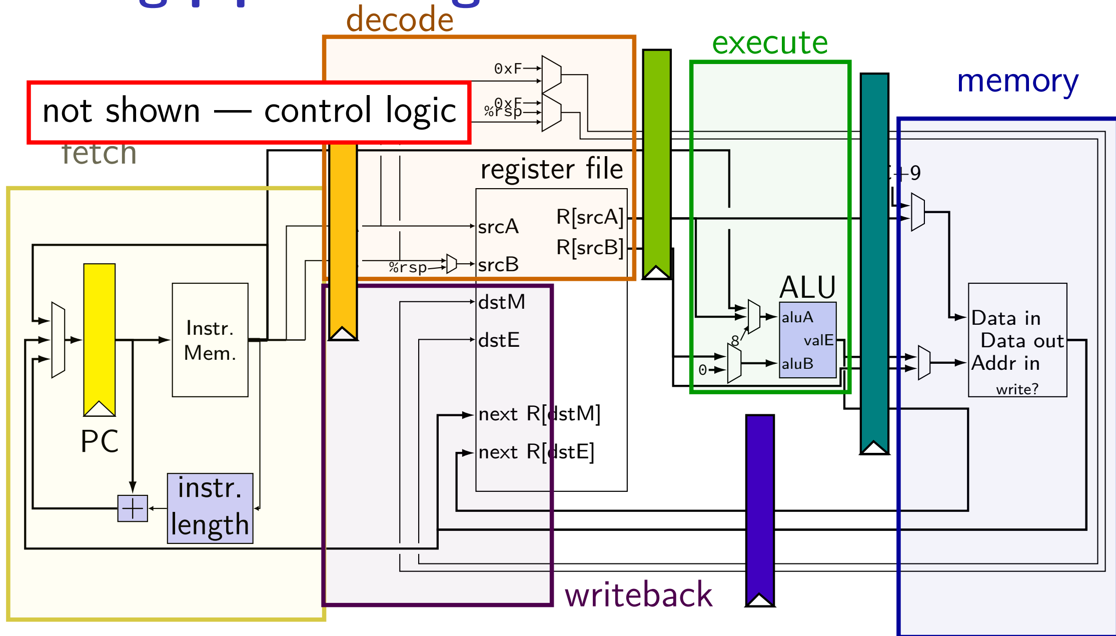

passing values in pipeline

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

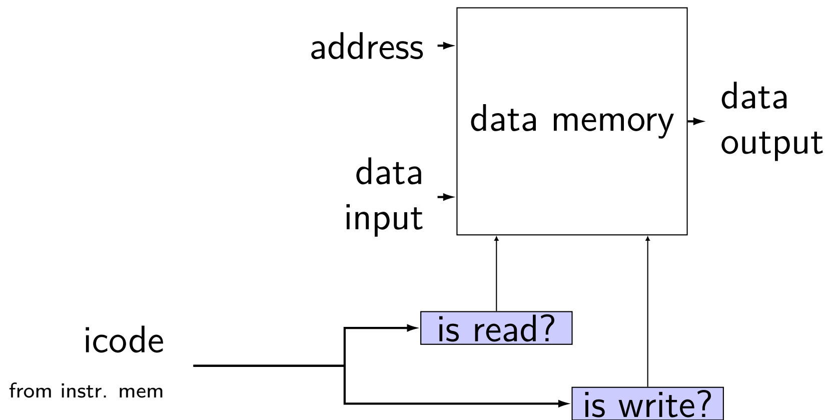

memory read/write logic

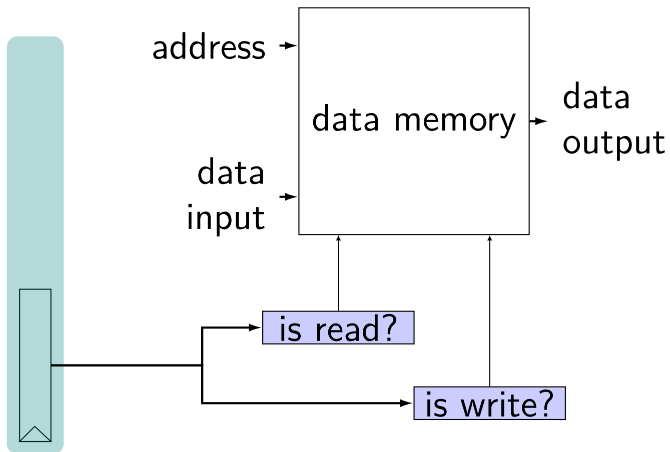

memory read/write logic

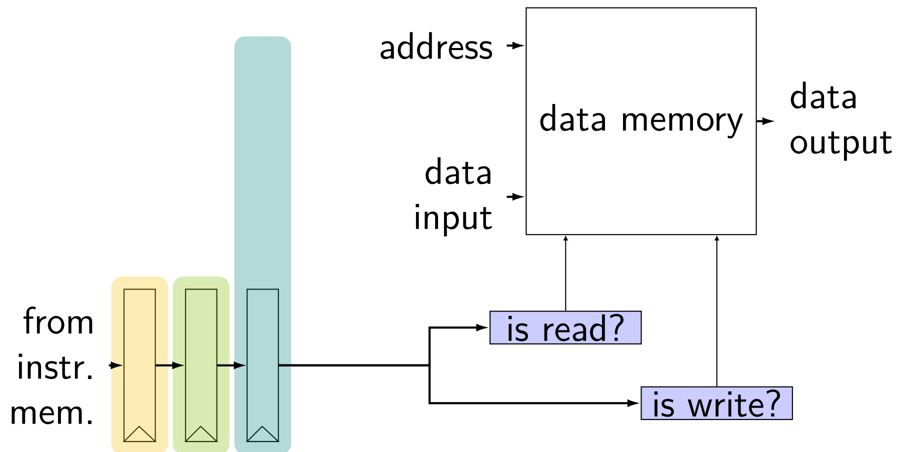

memory read/write: SEQ code

```
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 ...
```

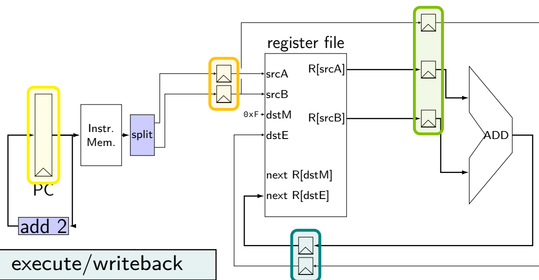

	fetch	fetch/decode		decode/execute			execute/writeback	
cycle	PC	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstE]	dstE
0	0x0							
1	0x2	8	9					
2		9	8	800	900	9		
3				900	800	8	1700	9
4							1700	8

addq processor: data hazard

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

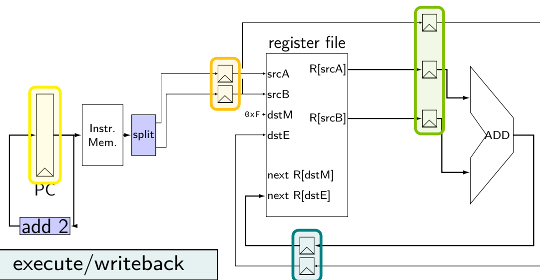

	fetch	fetch/decode		decode/execute			execute/writeback	
cycle	PC	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstE]	dstE
0	0x0							
1	0x2	8	9					
2		9	8	800	900	9		
3				900	800	8	1700	9
4							1700	8

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 `addq`s 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
```

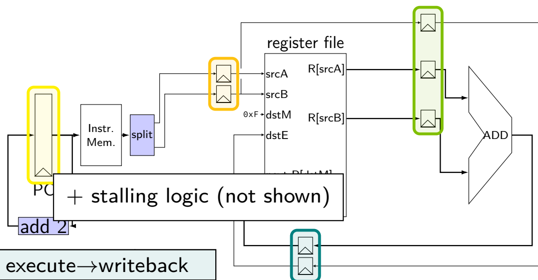

	fetch	fetch→decode		decode→execute			execute→writeback	
cycle	PC	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstE]	dstE
0	0x0							
1	0x2*	8	9					
2	0x2*	F	F	800	900	9		
3	0x2	F	F	---	---	F	1700	9
4	0x4	9	8	---	---	F	---	F
5		10	11	1700	800	8	---	F
6				1000	1100	11	2500	8

addq processor: data hazard stall

```
// initially %r8 = 800,
//           %r9 = 900, etc.
addq %r8, %r9
// hardware stalls twice
addq %r9, %r8
addq %r10, %r11
```

+ stalling logic (not shown)

	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
```

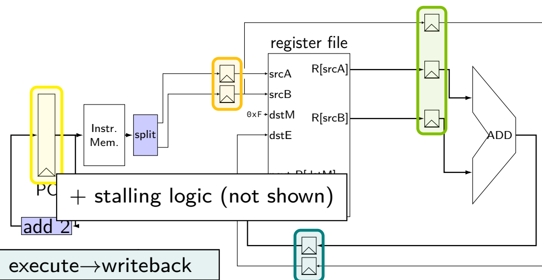

	fetch	fetch→decode		decode→execute			execute→writeback	
cycle	PC	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstE]	dstE
0	0x0							
1	0x2*	8	9					
2	0x2*	F	F	800	900	9		
3	0x2	F	F	---	---	F	1700	9
4	0x4	9	8	---	---	F	---	F
5		10	11	1700	800	8	---	F
6				1000	1100	11	2500	8

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

addq stall

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

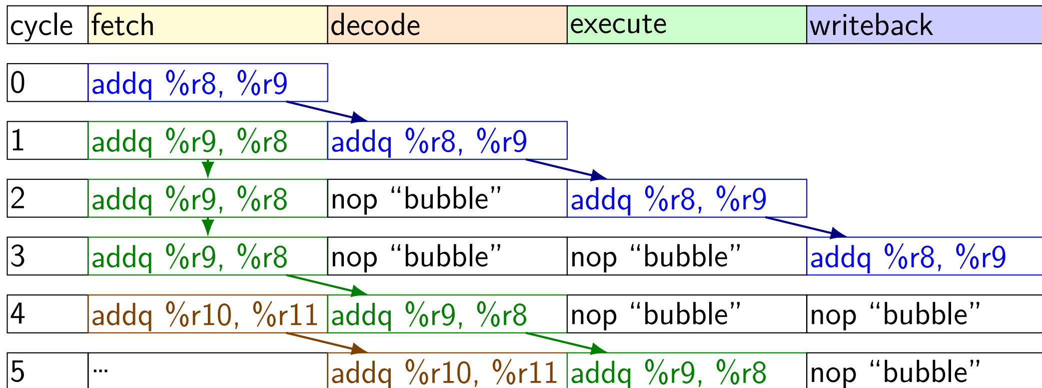

cycle	fetch	decode	execute	writeback
0	addq %r8, %r9			
1	addq %r9, %r8	addq %r8, %r9		
2	addq %r9, %r8	nop “bubble”	addq %r8, %r9	
3	addq %r9, %r8	nop “bubble”	nop “bubble”	addq %r8, %r9
4	addq %r10, %r11	addq %r9, %r8	nop “bubble”	nop “bubble”
5	…	addq %r10, %r11	addq %r9, %r8	nop “bubble”

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 exericse

```
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 exericse solution

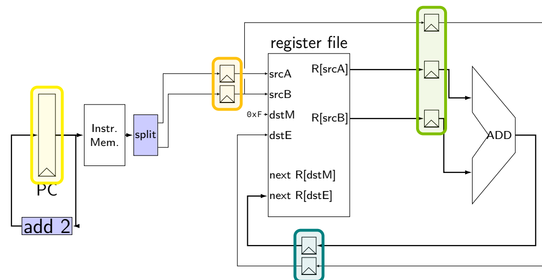

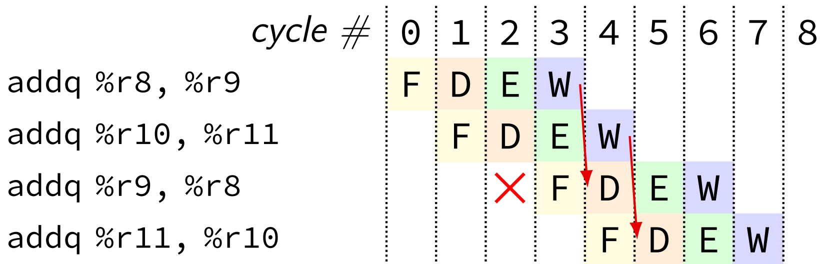

revisiting data hazards

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

location of values during cycle 2:

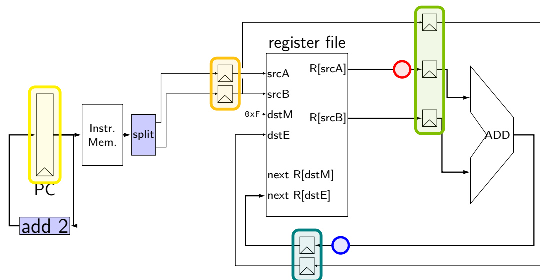


```
// initially %r8 = 800,
//           %r9 = 900, etc.
addq %r8, %r9
addq %r9, %r8
addq ...
addq ...
```

	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

motivation

location of values during cycle 2:

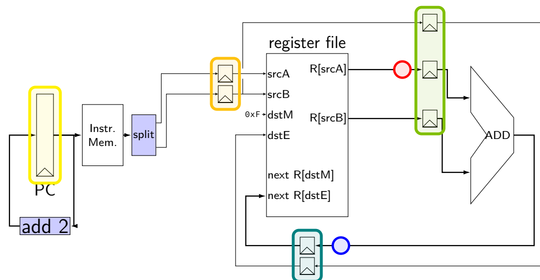


```
// initially %r8 = 800,
//           %r9 = 900, etc.
addq %r8, %r9
addq %r9, %r8
addq ...
addq ...
```

	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

forwarding

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

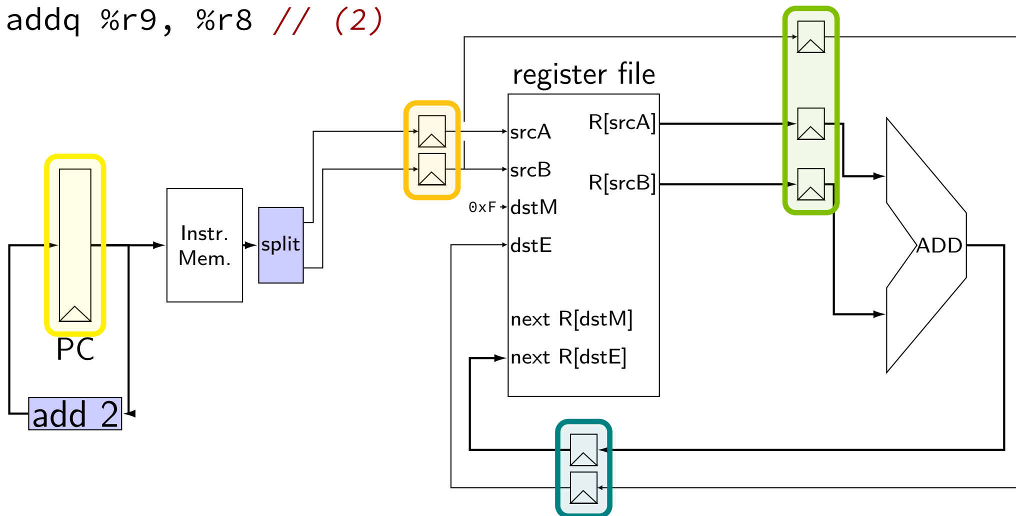

forwarding

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

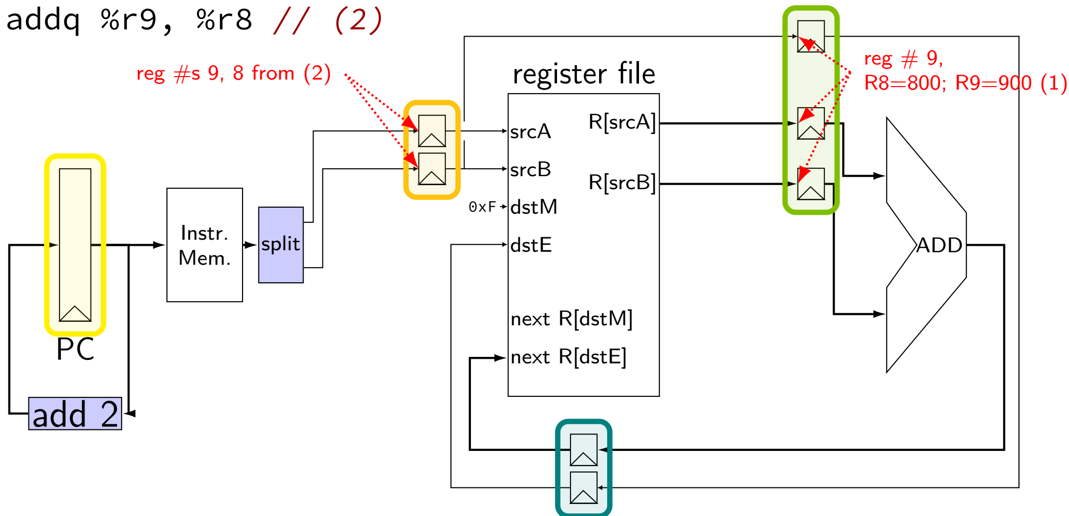

forwarding

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

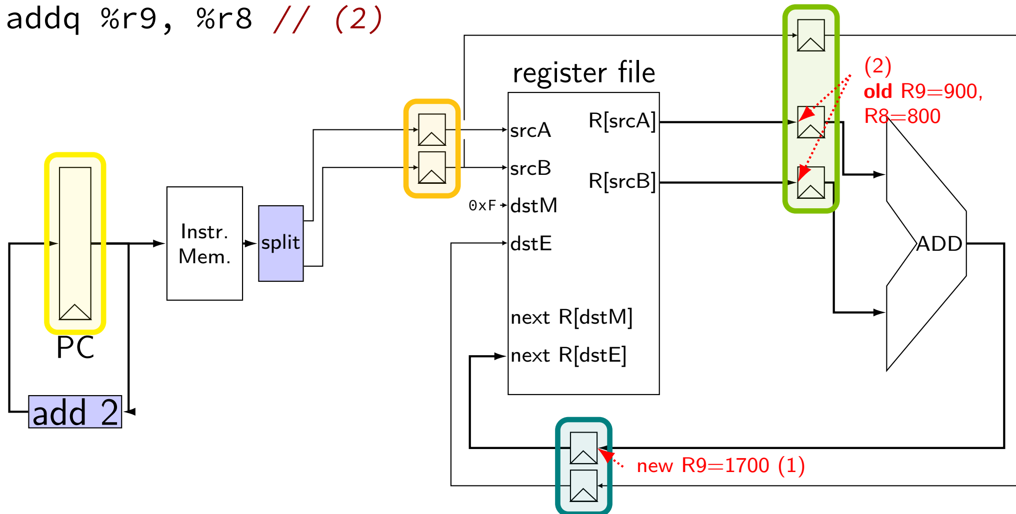

forwarding

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

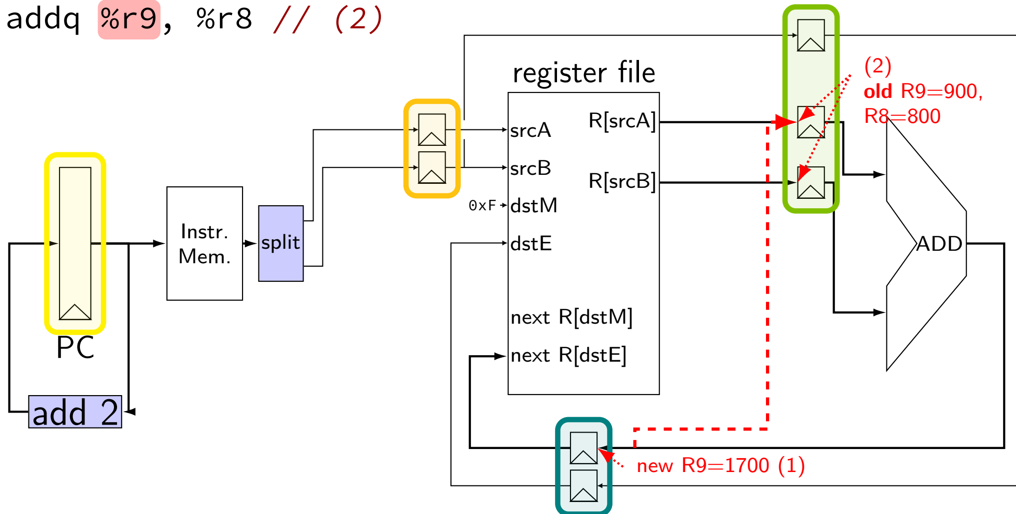

forwarding

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

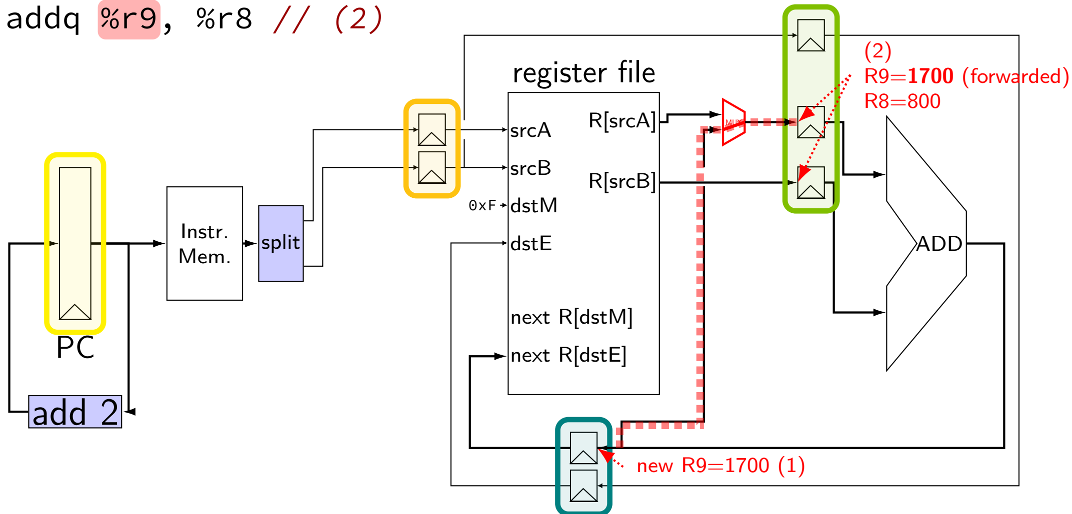

forwarding

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

forwarding: MUX conditions

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

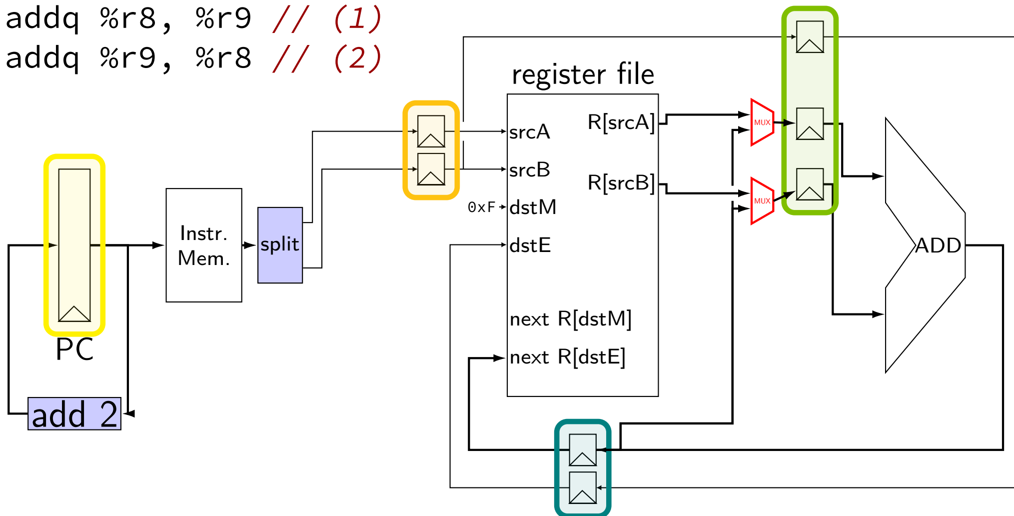

forwarding: MUX conditions

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

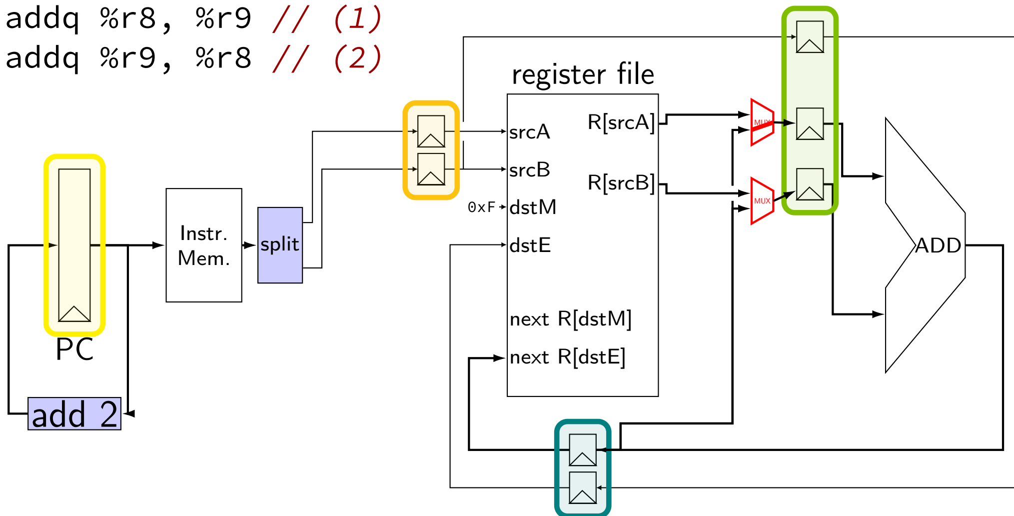

forwarding: MUX conditions

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

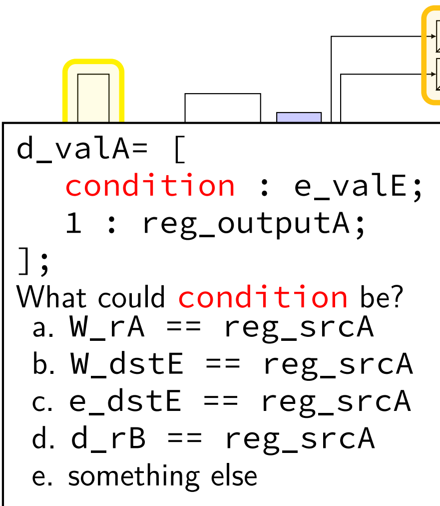

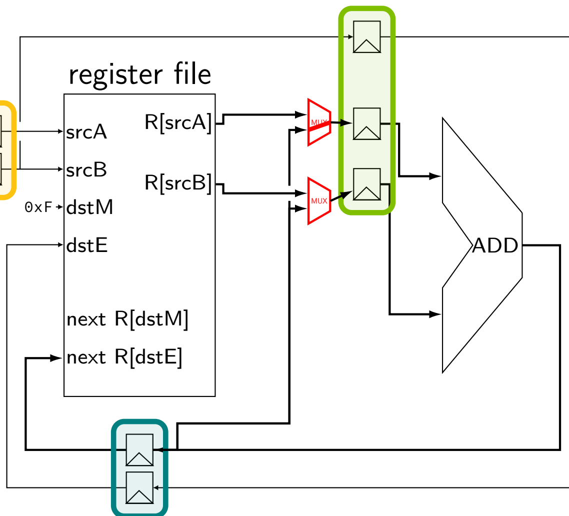

forwarding: MUX conditions

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

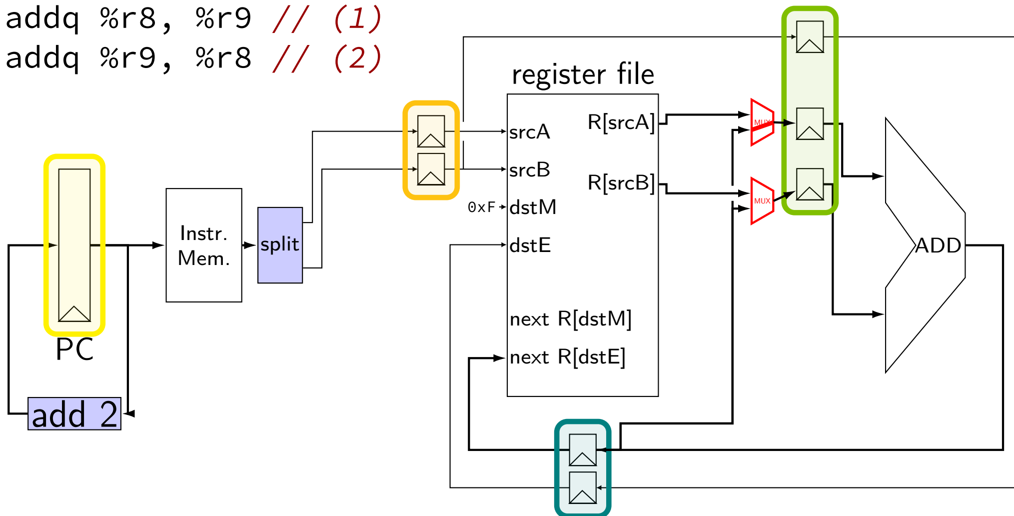

forwarding: MUX conditions

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

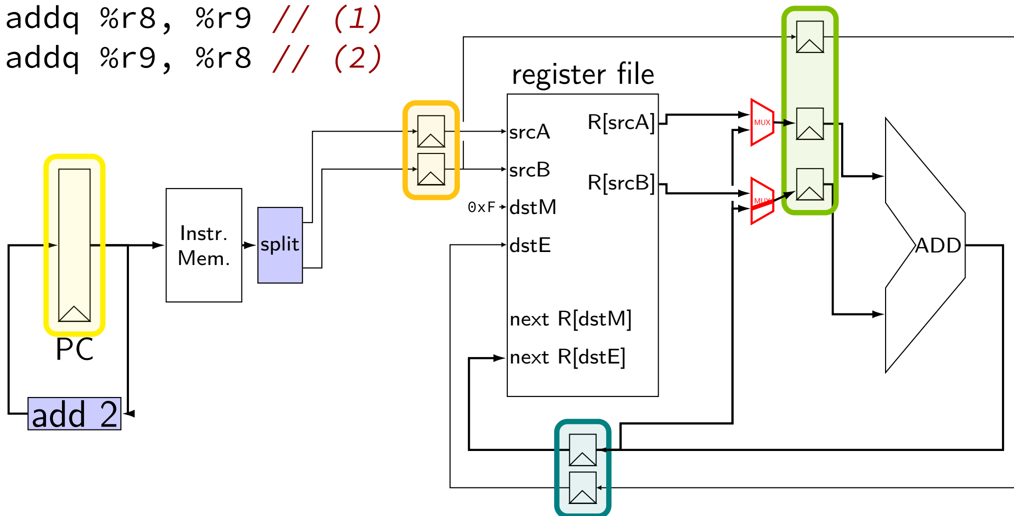

computing destinations early

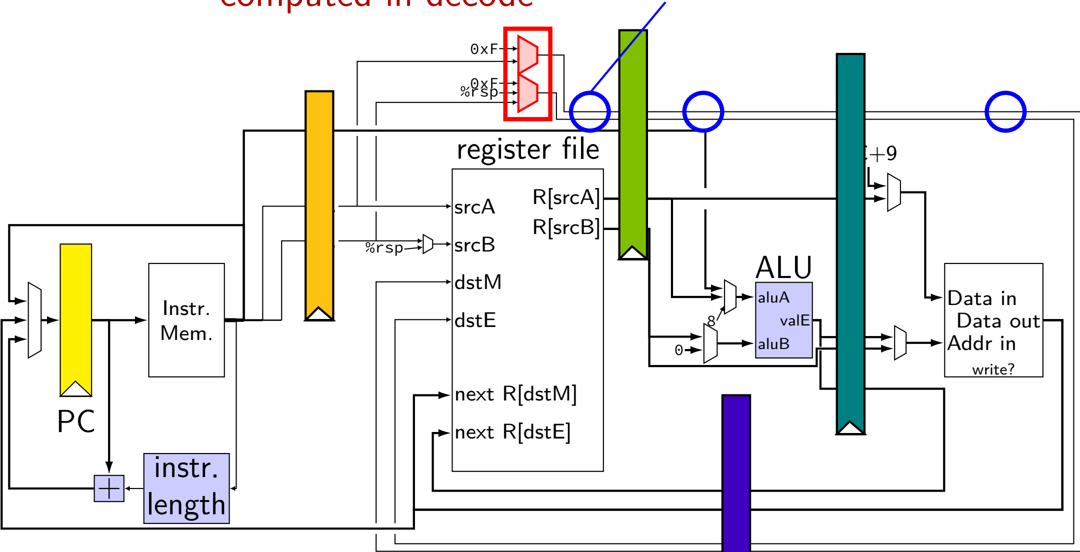

textbook convention on destinations

dstE/dstM computed mostly in decode

- passed through pipeline as d_dstE, e_dstE, ...

valE/valM only set to value to be stored in dstE/dstM

- passed through pipeline as e_valE, m_valE, ...

simplifies forwarding/stalling logic

stalling versus forwarding (1)

with stalling:

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

with forwarding:

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

forwarding more?

location of values during cycle 3:

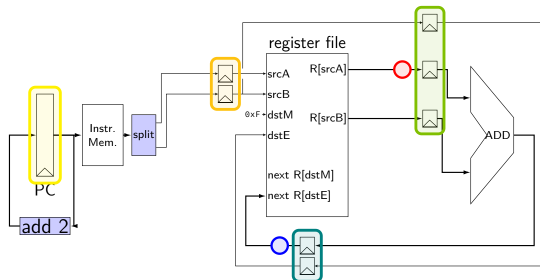


```
// initially %rax = 0,
//           %r8 = 800,
//           %r9 = 900, etc.
addq %r8, %r9
addq %rax, %rax
addq %r9, %r10
```

	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	0	0	800	900	9		
3		9	10	0	0	0	1700	9
4				900	1000	10	0	0
5							1900	10

should be 1700

forwarding two stages?

```
addq %r8, %r9    // (1) R9 <- R8 (800) + R9 (900)
addq %rax, %rax // (2)
addq %r9, %r10  // (3) R10 <- R9 (1700) + R10 (1000)
```

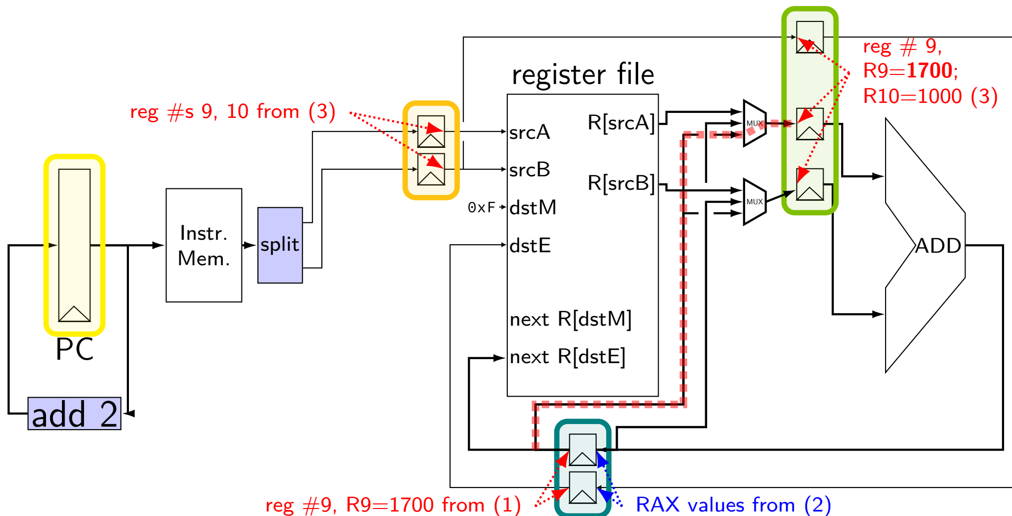

some forwarding paths

	0	1	2	3	4	5	6	7	8
`addq %r8, %r9`	F	D	E	M	W				
`subq %r9, %r11`		F	D	E	M	W			
`mrmovq 4(%r11), %r10`			F	D	E	M	W		
`rmmovq %r9, 8(%r11)`				F	D	E	M	W	
`xorq %r10, %r9`					F	D	E	M	W

(column header: *cycle #*)

some forwarding paths

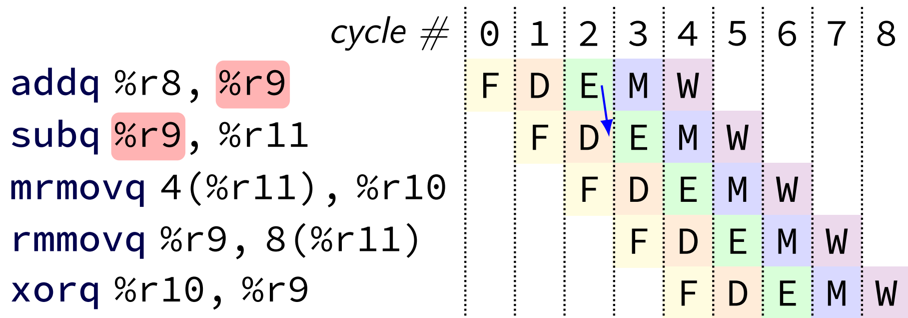

some forwarding paths

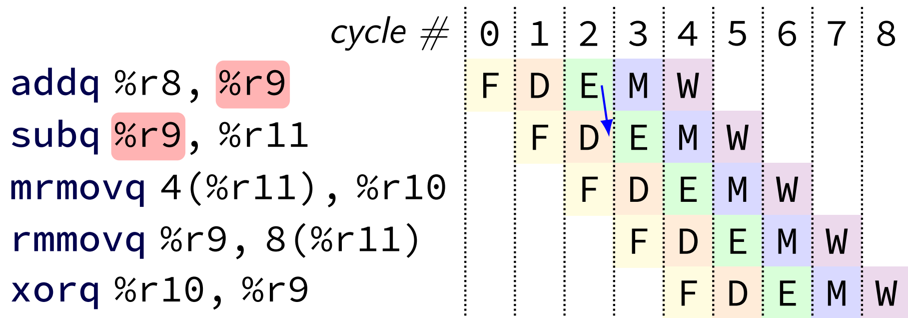

some forwarding paths

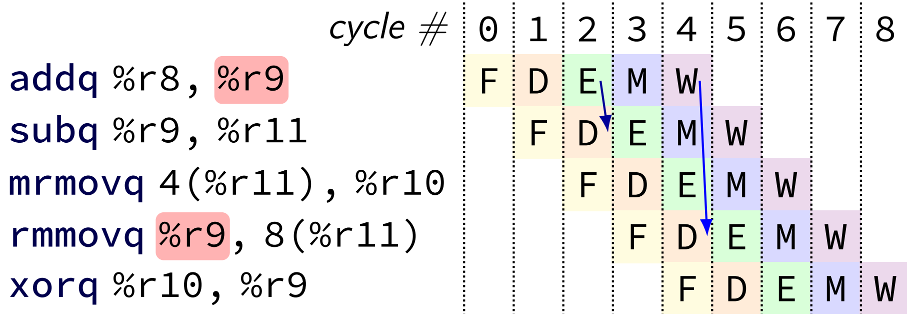

some forwarding paths

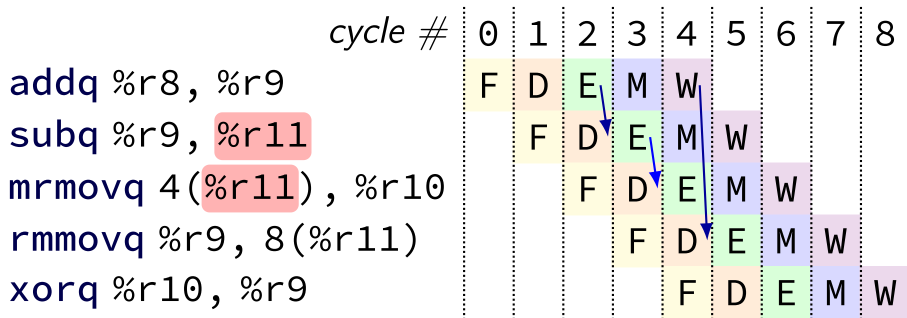

some forwarding paths

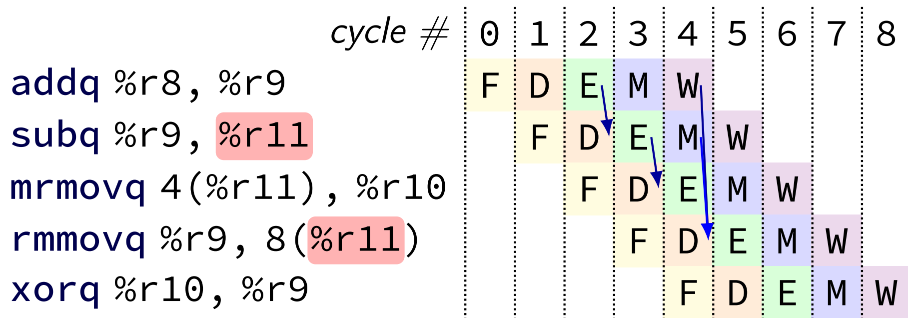

some forwarding paths

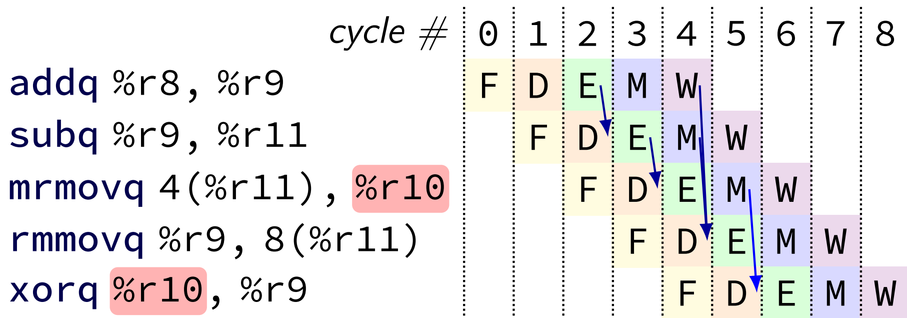

some forwarding paths

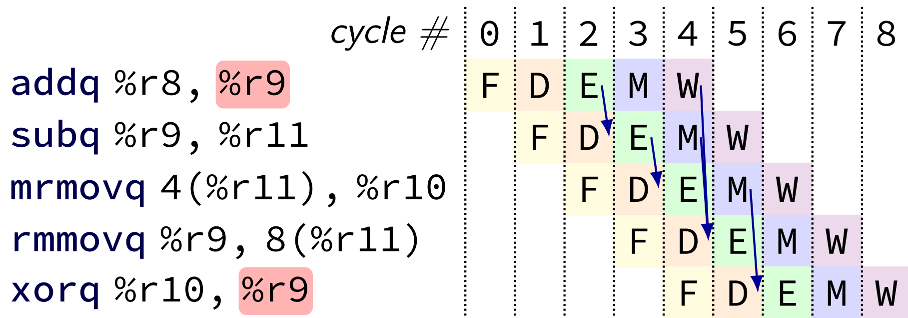

some forwarding paths

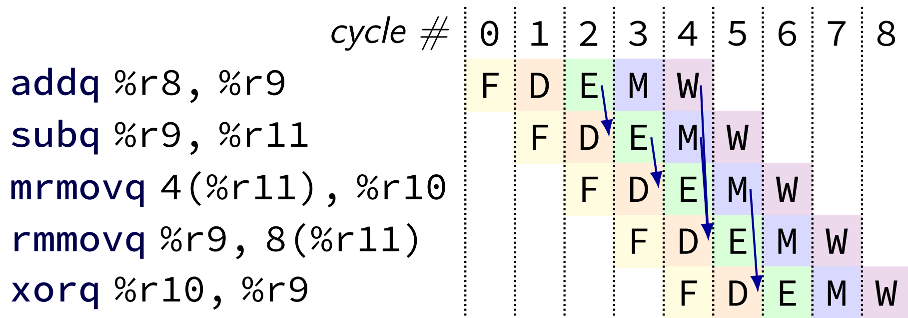

multiple forwarding paths (1)

	cycle #	0	1	2	3	4	5	6	7	8
addq `%r10, %r8`		F	D	E	M	W				
addq `%r11, %r8`			F	D	E	M	W			
addq `%r12, %r8`				F	D	E	M	W		

multiple forwarding paths (1)

	cycle #	0	1	2	3	4	5	6	7	8
addq %r10, %r8		F	D	E	M	W				
addq %r11, %r8			F	D	E	M	W			
addq %r12, %r8				F	D	E	M	W		

multiple forwarding HCL (1)

```
/* decode output: valA */
d_valA = [
    ...
    reg_srcA == e_dstE : e_valE;
        /* forward from end of execute */

    reg_srcA == m_dstE : m_valE;
        /* forward from end of memory */


    ...
    1 : reg_outputA;
];
```

multiple forwarding paths (2)

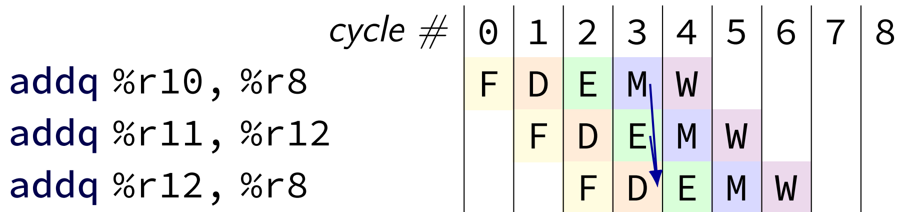

multiple forwarding paths (2)

	cycle #	0	1	2	3	4	5	6	7	8
`addq %r10, %r8`		F	D	E	M	W				
`addq %r11, %r12`			F	D	E	M	W			
`addq %r12, %r8`				F	D	E	M	W		

multiple forwarding paths (2)

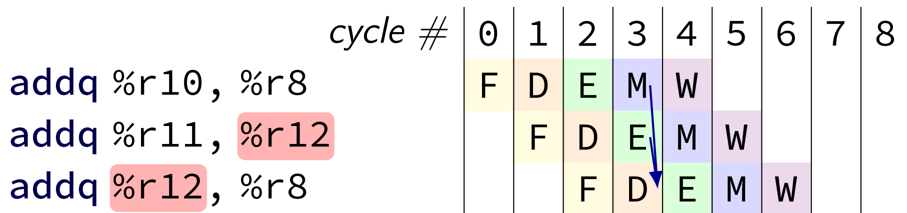

multiple forwarding HCL (2)

```
d_valA = [
    ...
    reg_srcA == e_dstE : e_valE;
    ...
    1 : reg_outputA;
];
...
d_valB = [
    ...
    reg_srcB == m_dstE : m_valE;
    ...
    1 : reg_outputB;
];
```

exercise: forwarding paths

	cycle #	0	1	2	3	4	5	6	7	8
`addq %r8, %r9`		F	D	E	M	W				
`subq %r8, %r10`			F	D	E	M	W			
`xorq %r8, %r9`				F	D	E	M	W		
`andq %r9, %r8`					F	D	E	M	W	

in subq, %r8 is __________ addq.

in xorq, %r9 is __________ addq.

in andq, %r9 is __________ addq.

in andq, %r9 is __________ xorq.

A: not forwarded from

B-D: forwarded to decode from {execute,memory,writeback} stage of

unsolved problem

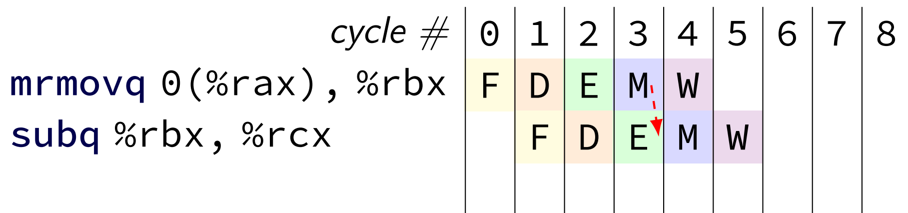

combine stalling and forwarding to resolve hazard

assumption in diagram: hazard detected in `subq`'s decode stage

(since easier than detecting it in fetch stage)

typically what you'll implement

intuition: try to forward, but detect that it won't work

unsolved problem

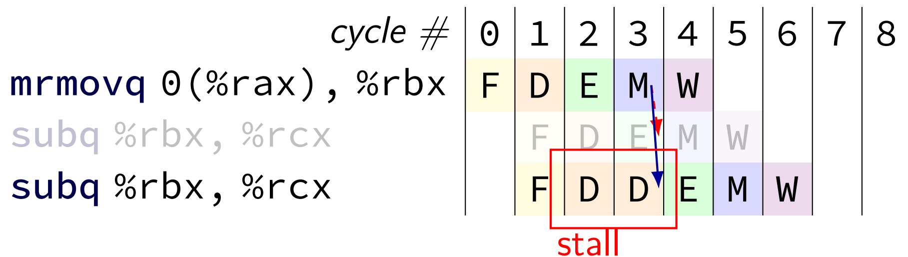

combine stalling and forwarding to resolve hazard

assumption in diagram: hazard detected in `subq`'s decode stage

(since easier than detecting it in fetch stage)

typically what you'll implement

intuition: try to forward, but detect that it won't work

control hazard

```
subq %r8, %r9
je   0xFFFF
addq %r10, %r11
```

	fetch		fetch→decode		decode→execute			execute→writeback	
cycle	PC	SF/ZF	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstE]	dstE
0	0x0	0/1							
1	0x2	0/1	8	9					
2	???	0/1	0xF	0xF	800	900	9		

control hazard

```
subq %r8, %r9
je   0xFFFF
addq %r10, %r11
```

	fetch		fetch→decode		decode→execute			execute→writeback	
cycle	PC	SF/ZF	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstE]	dstE
0	0x0	0/1							
1	0x2	0/1	8	9					
2	???	0/1	0xF	0xF	800	900	9		

`0xFFFF` if R[8] = R[9]; `0x12` otherwise

control hazard: stall

```
addq %r8, %r9
// insert two nops
je   0xFFFF
addq %r10, %r11
```

	fetch		fetch→decode		decode→execute			execute→writeback	
cycle	PC	SF/ZF	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstE]	dstE
0	0x0	0/1							
1	0x2*	0/1	8	9					
2	0x2*	0/1	0xF	0xF	800	900	9		
3	0x2	0/0	0xF	0xF	---	---	0xF	1700	9
4	0x10	0/0	0xF	0xF	---	---	0xF	---	0xF
5			10	11	---	---	0xF	---	0xF
6					1000	1100	11	---	0xF

control hazard: stall

```
addq %r8, %r9
// insert two nops
je   0xFFFF
addq %r10, %r11
```

	fetch		fetch→decode		decode→execute			execute→writeback	
cycle	PC	S[illegible]							
0	0x0	0/1							
1	0x2*	0/1	8	9					
2	0x2*	0/1	0xF	0xF	800	900	9		
3	0x2	0/0	0xF	0xF	---	---	0xF	1700	9
4	0x10	0/0	0xF	0xF	---	---	0xF	---	0xF
5			10	11	---	---	0xF	---	0xF
6					1000	1100	11	---	0xF

wait for two cycles for addq to update SF/ZF

control hazard: stall

```
addq %r8, %r9
// insert two nops
je   0xFFFF
addq %r10, %r11
```

	fetch		fetch→decode		decode→execute			execute→writeback	
cycle	PC	SF/ZF	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstE]	dstE
0	0x0	0/1							
1	0x2*								
2	0x2*	0/1	0xF	0xF	800	900	9		
3	0x2	0/0	0xF	0xF	---	---	0xF	1700	9
4	0x10	0/0	0xF	0xF	---	---	0xF	---	0xF
5			10	11	---	---	0xF	---	0xF
6					1000	1100	11	---	0xF

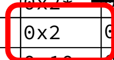

execute je instruction (use SF/ZF)

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

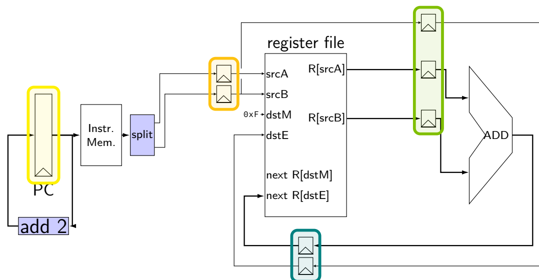

no pipelining: 1 instruction per 550 ps

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)