

PIPE

pipelining summary

assembly line for math

divide into pieces

each piece in parallel for different instructions

increase throughput but also increase latency


limited by uneven division of work

limited by dependencies (“hazards”)

limited by register delays

Last time: data hazard stall


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



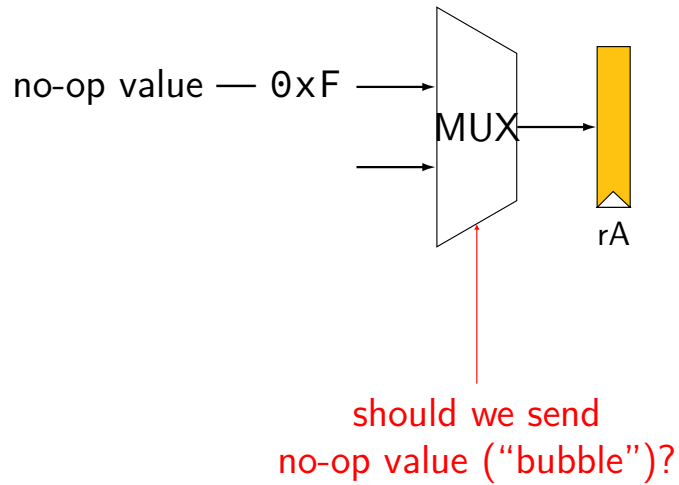
	fetch	fetch/decode		decode/execute			execute/writeback	
cycle	PC	rA	rB	R[srcA]	R[srcB]	dstE	next R[dstE]	dstE
0	0x0							
1	0x2*	8	9					
2	0x2*	F	F	800	900	9		
3	0x2	F	F	---	---	F	1700	9
4		9	8	---	---	F	---	F
5				1700	800	8	---	F
6							2500	8

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

fetch/fetch logic — advance or not



fetch/decode logic — bubble or not



5

preview: HCL2D shortcuts

HCL2D provides these MUXes for you — for **every register bank**

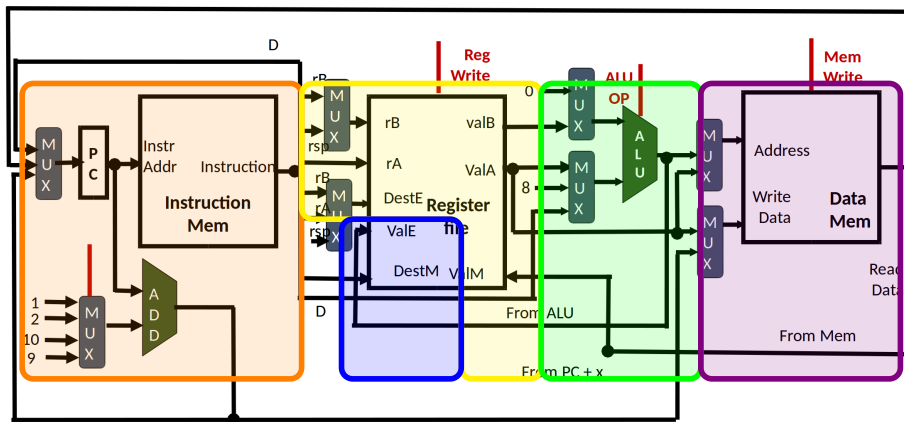
controlled by “**bubble**” and “**stall**” signals

bubble — use default value (stage should do no-op)

stall — keep current value (stage should try again)

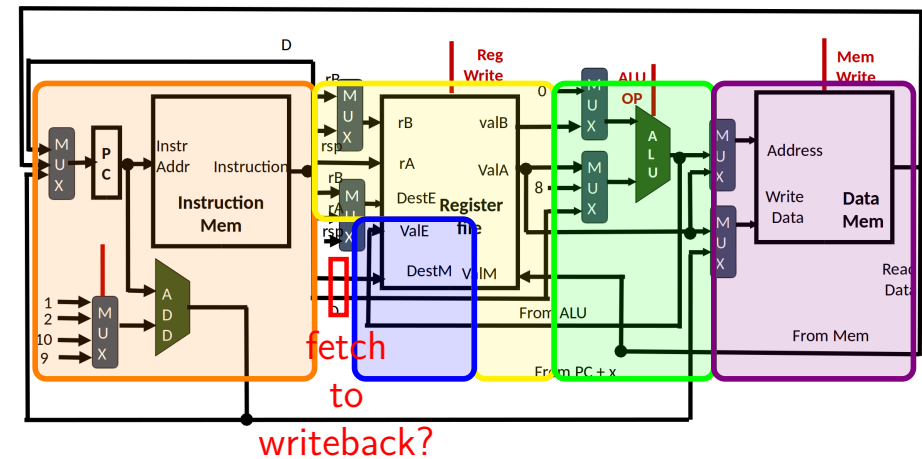
6

old SEQ picture



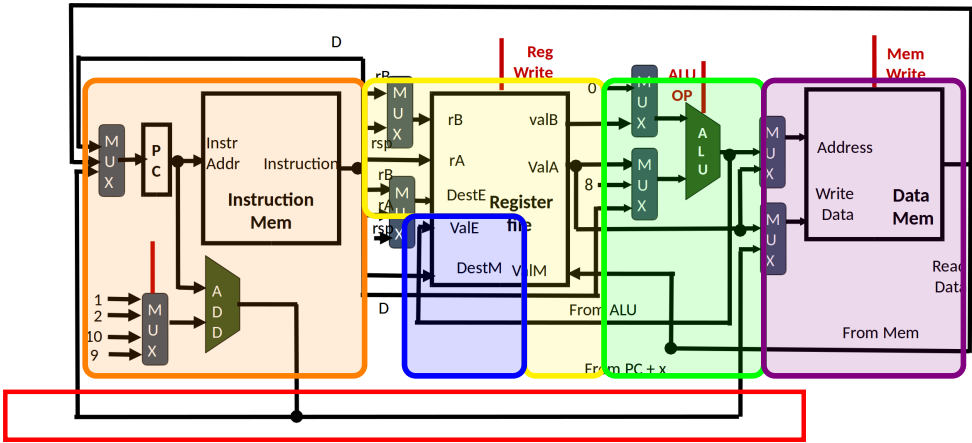
7

old SEQ picture



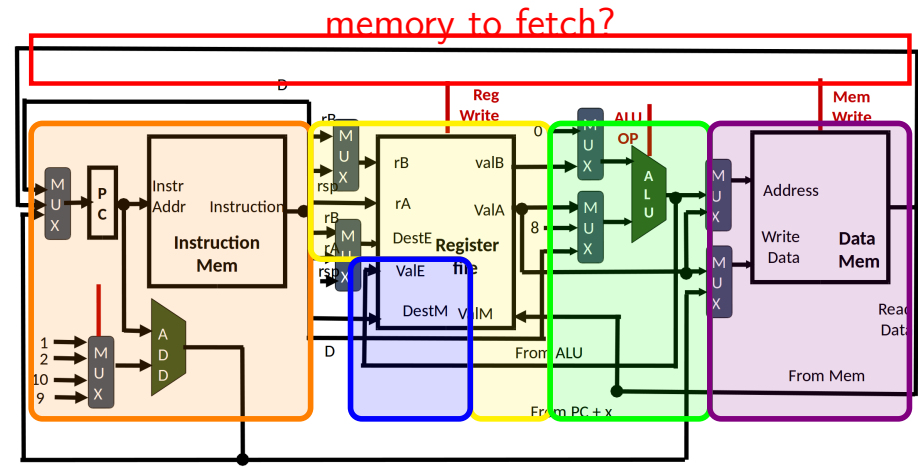
7

old SEQ picture



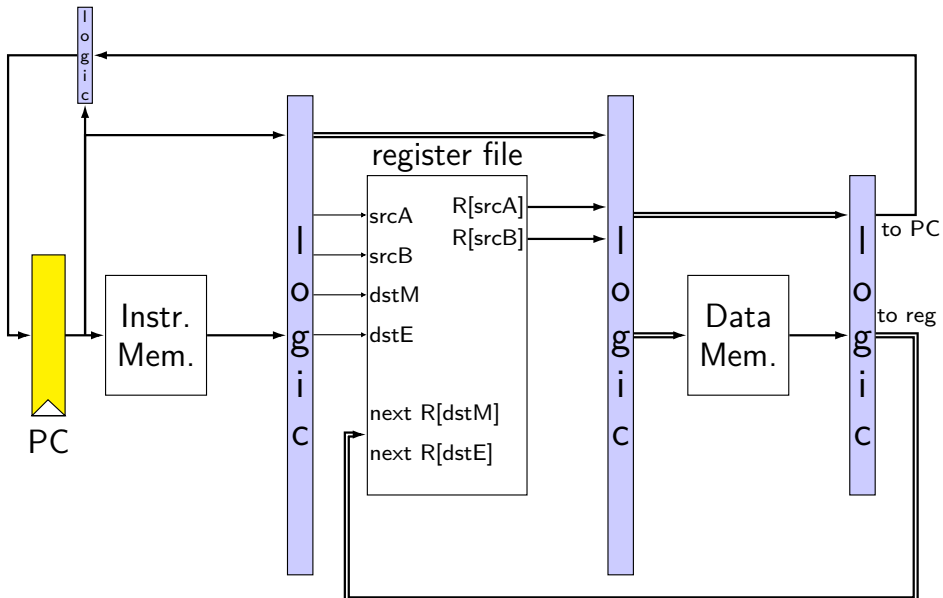
fetch to memory?

old SEQ picture

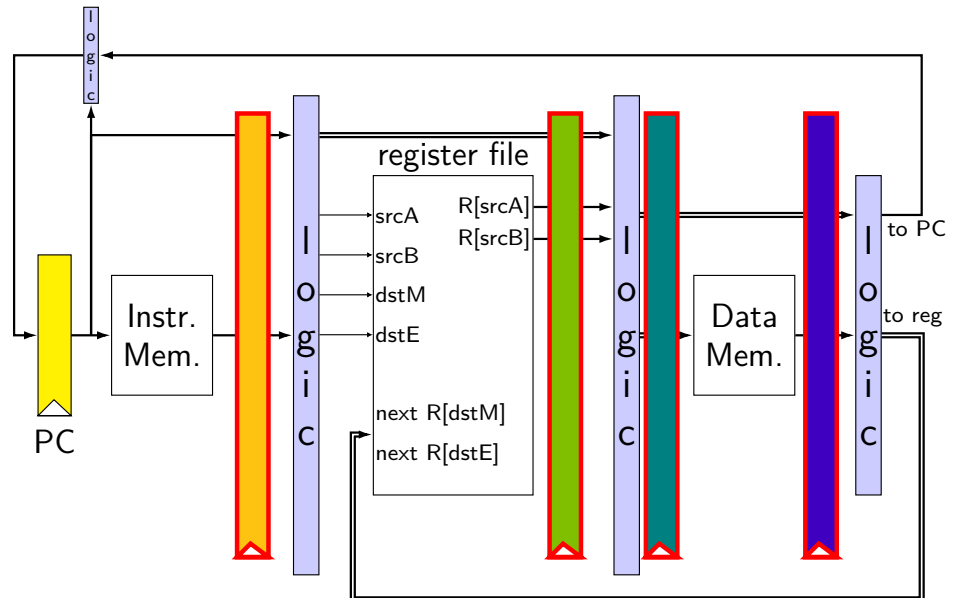


memory to fetch?

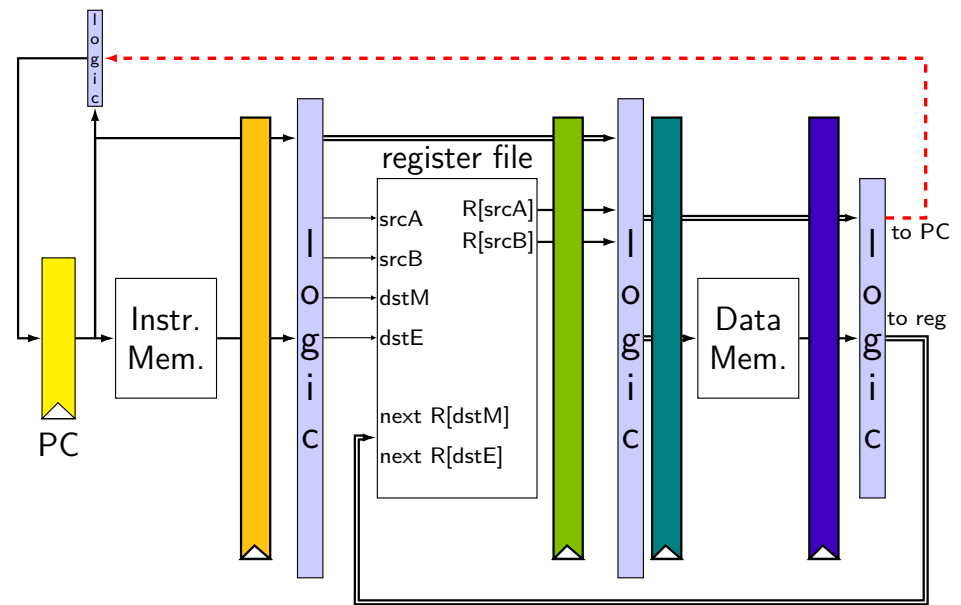
SEQ revised



SEQ + pipeline registers



SEQ + pipeline registers



bad timeline

```
subq %rax, %rbx
addq %rcx, %rdx
...
```

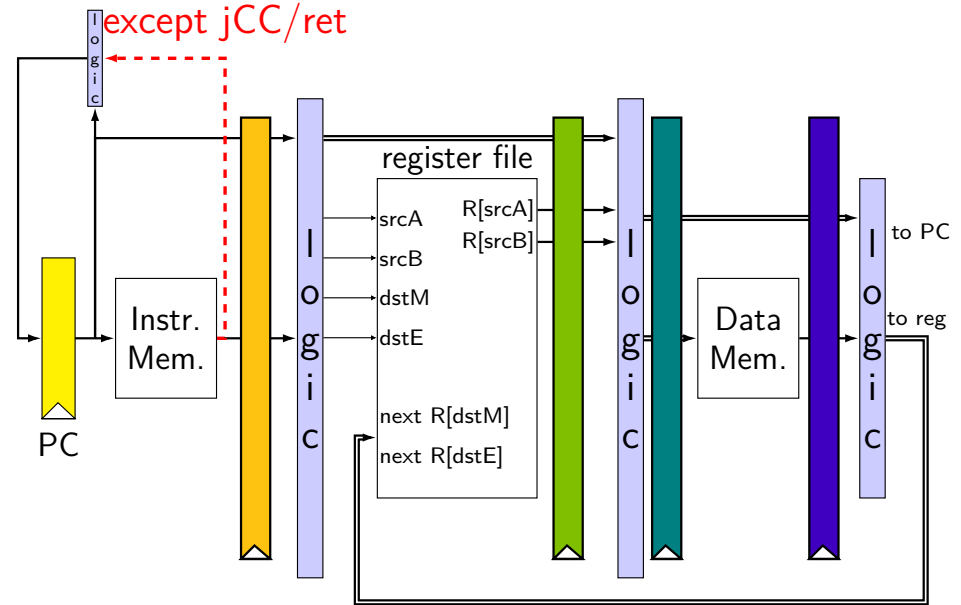
PC update after memory?

time	fetch	decode	execute	memory	writeback
1	subq				
2	waiting	subq			
3	waiting	no-op	subq		
4	waiting	no-op	no-op	subq	
5	addq	no-op	no-op	no-op	subq

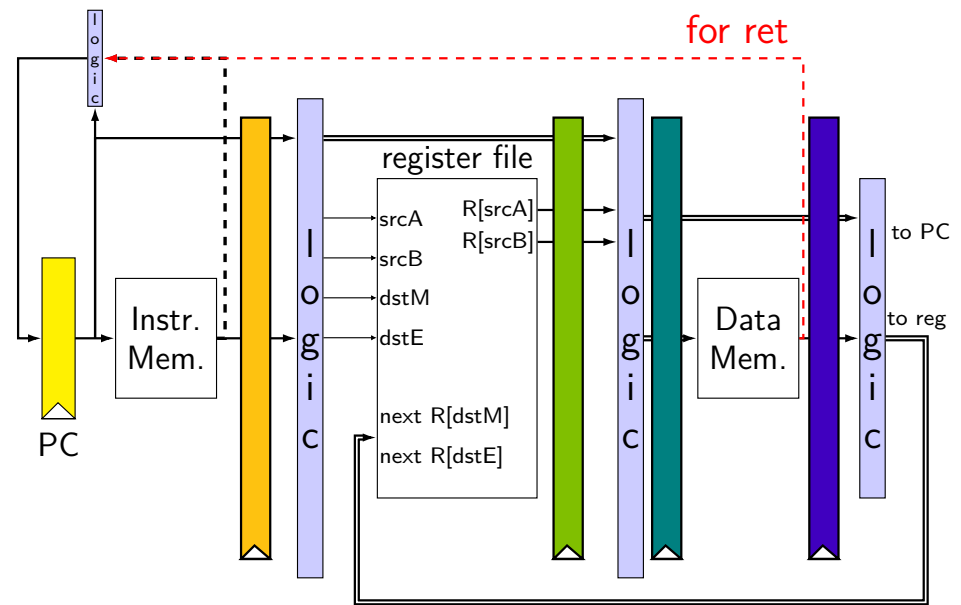
what we want

time	fetch	decode	execute	memory	writeback
1	subq				
2	addq	subq			
3	...	addq	subq		
4	addq	subq	
5	addq	subq

SEQ + pipeline registers



SEQ + pipeline registers



Stages (revised)

fetch — instruction memory, *most* PC computation
decode — reading register file
execute — computation, condition code read/write
memory — memory read/write
writeback — writing register file, writing Stat register

12

Stages (revised)

fetch — instruction memory, *most* PC computation
decode — reading register file
execute — computation, **condition code read/write**
memory — memory read/write
read/write in same stage avoids data hazards
get value updated for prior instruction at register

12

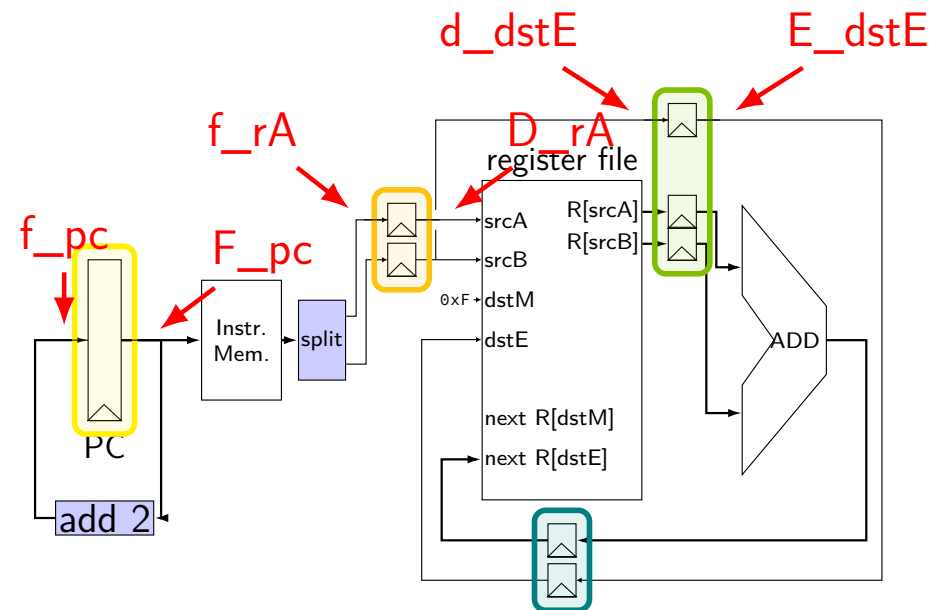
Stages (revised)

fetch — instruction memory, *most* PC computation
decode — reading register file
execute — computation, condition code read/write
memory — memory read/write
writeback — writing register file, **writing Stat register**

don't want to halt until everything else is done

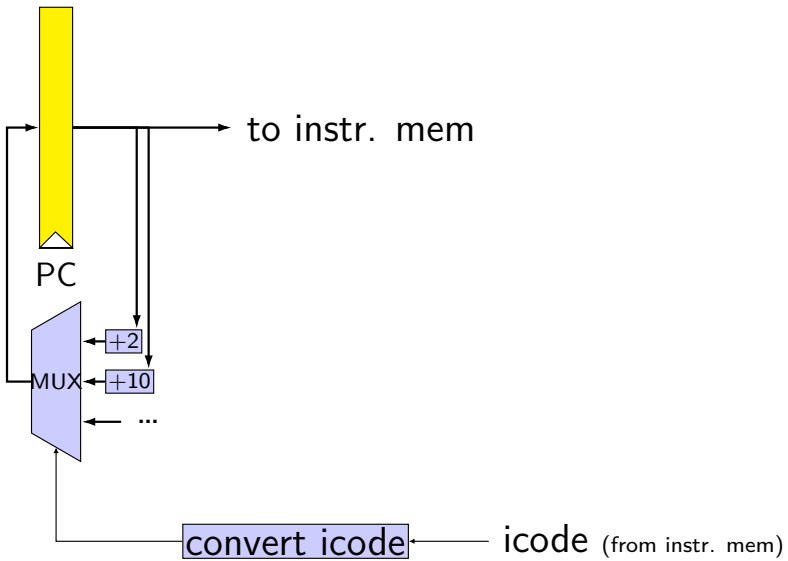
12

pipeline register naming convention



13

normal PC update: logic



14

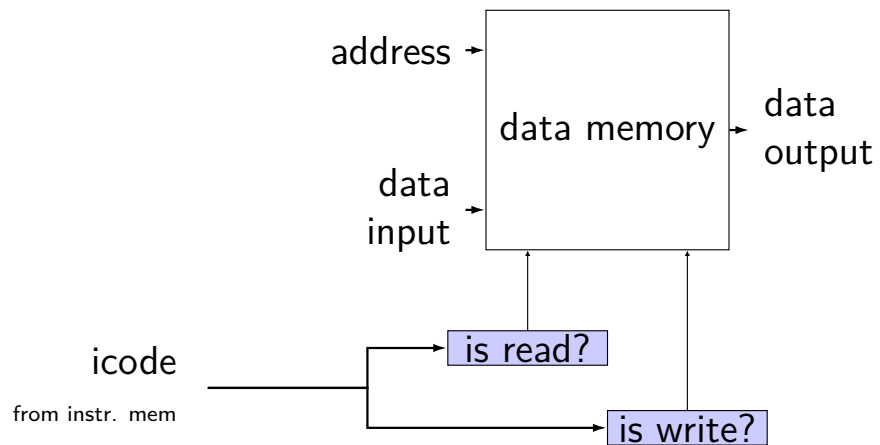
simple PC update

this week's lab...

```
icode = i10bytes[4..8];  
f_pc = [  
    icode == ADD || ...: F_pc + 2;  
    icode == IRMOVQ || ...: F_pc + 10;  
    ...  
];
```

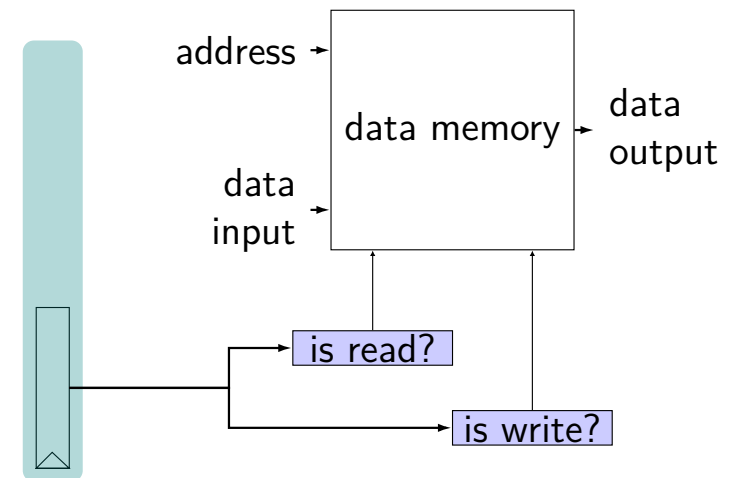
15

memory read/write logic



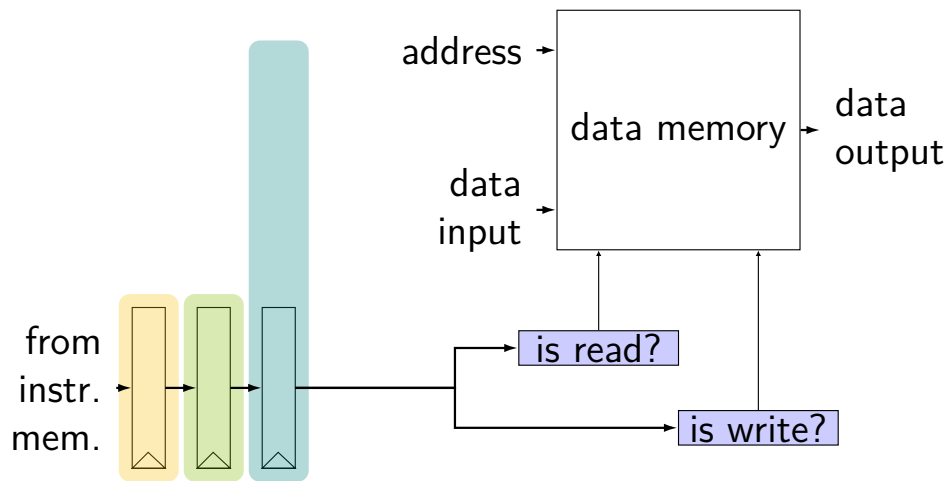
16

memory read/write logic



16

memory read/write logic



16

memory read/write: SEQ code

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

17

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;  
];
```

18

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;  
];
```

18

in general

will always pass **icode** in pipeline registers

control logic (often not drawn) will use it

examples:

register number selection

ALU input selection

stalling

19

coding pipeline stages

use only **prior stage's outputs**

e.g. decode stage: get from fetch (D_icode, \dots)

set only **inputs for next stage**

e.g. decode stage: send to execute (d_icode, \dots)






two exceptions (share between instructions):

what instruction to run next?

data and control hazards






20

pushq pipeline registers

stage	pushq rA	
fetch	$icode : ifun \leftarrow M_1[PC]$ $valP \leftarrow PC + 2$	 PC
PC update	$PC \leftarrow valP$	 icode
decode	$valA \leftarrow R[rA]$ $valB \leftarrow R[\%rsp]$	 icode
execute	$valE \leftarrow valB - 8$	 icode
memory	$M[valE] \leftarrow valA$	 icode
write back		






21

pushq pipeline registers

stage	pushq rA	
fetch	$icode : ifun \leftarrow M_1[PC]$ $valP \leftarrow PC + 2$	 PC
PC update	$PC \leftarrow valP$	 icode, rA, rB
decode	$valA \leftarrow R[rA]$ $valB \leftarrow R[\%rsp]$	 icode
execute	$valE \leftarrow valB - 8$	 icode
memory	$M[valE] \leftarrow valA$	 icode
write back		






21

pushq pipeline registers

stage	pushq rA	
fetch	icode : ifun $\leftarrow M_1[PC]$ valP $\leftarrow PC + 2$	 PC
PC update	PC \leftarrow valP	 icode, rA, rB
decode	valA $\leftarrow R[rA]$ valB $\leftarrow R[\%rsp]$	 icode, valA, valB
execute	valE \leftarrow valB - 8	 icode, valA
memory	$M[valE] \leftarrow$ valA	 icode
write back		






21

pushq pipeline registers

stage	pushq rA	
fetch	icode : ifun $\leftarrow M_1[PC]$ valP $\leftarrow PC + 2$	 PC
PC update	PC \leftarrow valP	 icode, rA, rB
decode	valA $\leftarrow R[rA]$ valB $\leftarrow R[\%rsp]$	 icode, valA, valB
execute	valE \leftarrow valB - 8	 icode, valA, valE
memory	$M[valE] \leftarrow$ valA	 icode
write back		






21

addq pipeline registers

stage	addq rA, rB	
fetch	icode : ifun $\leftarrow M_1[PC]$ rA : rB $\leftarrow M_1[PC+1]$ valP $\leftarrow PC + 2$	 PC
PC update	PC \leftarrow valP	 icode
decode	valA $\leftarrow R[rA]$ valB $\leftarrow R[rB]$	 icode
execute	valE \leftarrow valB + valB	 icode
memory		 icode
write back	$R[rB] \leftarrow$ valE	






22

addq pipeline registers

stage	addq rA, rB	
fetch	icode : ifun $\leftarrow M_1[PC]$ rA : rB $\leftarrow M_1[PC+1]$ valP $\leftarrow PC + 2$	 PC
PC update	PC \leftarrow valP	 icode, rA, rB
decode	valA $\leftarrow R[rA]$ valB $\leftarrow R[rB]$	 icode, rB
execute	valE \leftarrow valB + valB	 icode, rB
memory		 icode, rB
write back	$R[rB] \leftarrow$ valE	






22

addq pipeline registers

stage	addq rA, rB	
fetch	$icode : ifun \leftarrow M_1[PC]$ $rA : rB \leftarrow M_1[PC+1]$ $valP \leftarrow PC + 2$	 PC
PC update	$PC \leftarrow valP$	 icode, rA, rB
decode	$valA \leftarrow R[rA]$ $valB \leftarrow R[rB]$	
execute	$valE \leftarrow valB + valB$	 icode, rB, valA, valB
memory		 icode, rB
write back	$R[rB] \leftarrow valE$	 icode, rB






22

addq pipeline registers

stage	addq rA, rB	
fetch	$icode : ifun \leftarrow M_1[PC]$ $rA : rB \leftarrow M_1[PC+1]$ $valP \leftarrow PC + 2$	 PC
PC update	$PC \leftarrow valP$	 icode, rA, rB
decode	$valA \leftarrow R[rA]$ $valB \leftarrow R[rB]$	
execute	$valE \leftarrow valB + valB$	 icode, rB, valA, valB
memory		 icode, rB, valE
write back	$R[rB] \leftarrow valE$	 icode, rB, valE






22

addq pipeline registers

stage	addq rA, rB	
fetch	$icode : ifun \leftarrow M_1[PC]$ $rA : rB \leftarrow M_1[PC+1]$ $valP \leftarrow PC + 2$	 PC
PC update	$PC \leftarrow valP$	 icode, rA, rB
decode	$valA \leftarrow R[rA]$ $valB \leftarrow R[rB]$ $dstE \leftarrow rB$	
execute	$valE \leftarrow valB + valB$	 icode, dstE, valA, valB
memory		 icode, dstE, valE
write back	$R[dstE] \leftarrow valE$	 icode, dstE, valE

22

addq pipeline registers

stage	addq rA, rB	
fetch	$icode : ifun \leftarrow M_1[PC]$ $rA : rB \leftarrow M_1[PC+1]$ $valP \leftarrow PC + 2$	 PC
PC update	$PC \leftarrow valP$	 icode, rA, rB
decode	$valA \leftarrow R[rA]$ $valB \leftarrow R[rB]$ $dstE \leftarrow rB$	
execute	$valE \leftarrow valB + valB$	 icode, dstE, valA, valB
memory		 icode, dstE, valE
write back	$R[dstE] \leftarrow valE$	 icode, dstE, valE

redundant with rB + icode
but will make handling data hazards easier

22

computing the PC

conditional jmp — instruction and **condition codes**

ret — from **memory**

otherwise — instruction only

23

computing the PC

conditional jmp — instruction and **condition codes**

ret — from **memory**

otherwise — instruction only

instruction only — can be done in fetch stage entirely

23

computing the PC

conditional jmp — instruction and **condition codes**

ret — from **memory**

otherwise — instruction only

must wait till memory stage
worst case: ret immediately follows memory write

23

computing the PC

conditional jmp — instruction and **condition codes**

ret — from **memory**

otherwise — in

must wait till execute stage
worst case: previous instruction sets CCs

23

conditional jmp (w/ stalling)

```
subq %r8, %r8
je label
```

```
label: irmovq ...
```

time	fetch	decode	execute	memory	writeback
1	OPq				
2	jCC	OPq			
3	wait for jCC	jCC	OPq (set ZF)		
4	wait for jCC	nothing	jCC (use ZF)	OPq	
5	irmovq	nothing	nothing	jCC (done)	OPq

24

conditional jmp (w/ stalling)

```
subq %r8, %r8
je label
```

```
label: irmovq ...
```

time	fetch	decode	execute	memory	writeback
1	OPq				
2	jCC	OPq			
3	wait for jCC	jCC	OPq (set ZF)		
4	wait for jCC	nothing	jCC (use ZF)	OPq	
5	irmovq	nothing	nothing	jCC (done)	OPq

24

conditional jmp (w/ stalling)

```
subq %r8, %r8
je label
```

```
label: irmovq ...
```

time	fetch	decode	execute	memory	writeback
1	OPq				
2	jCC	OPq			
3	wait for jCC	jCC	OPq (set ZF)		
4	wait for jCC	nothing	jCC (use ZF)	OPq	
5	irmovq	nothing	nothing	jCC (done)	OPq

ZF sent via register

24

conditional jmp (w/ stalling)

```
subq %r8, %r8
je label
```

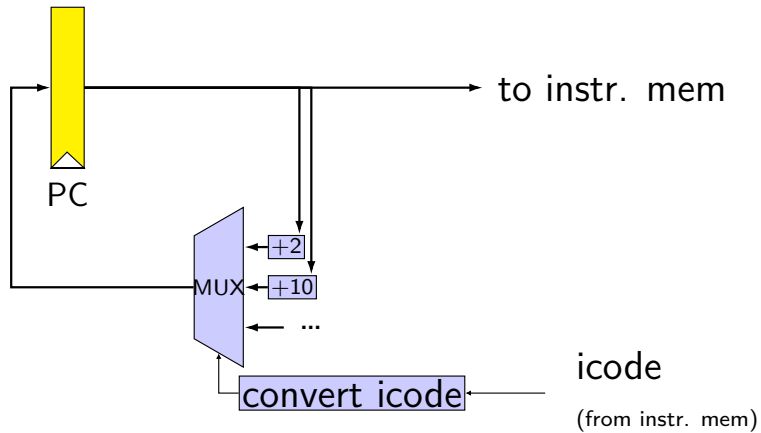
```
label: irmovq ...
```

time	fetch	decode	execute	memory	writeback
1	OPq				
2	jCC	OPq			
3	wait for jCC	jCC	OPq (set ZF)		
4	wait for jCC	nothing	jCC (use ZF)	OPq	
5	irmovq	nothing	nothing	jCC (done)	OPq

"taken" sent from execute to fetch

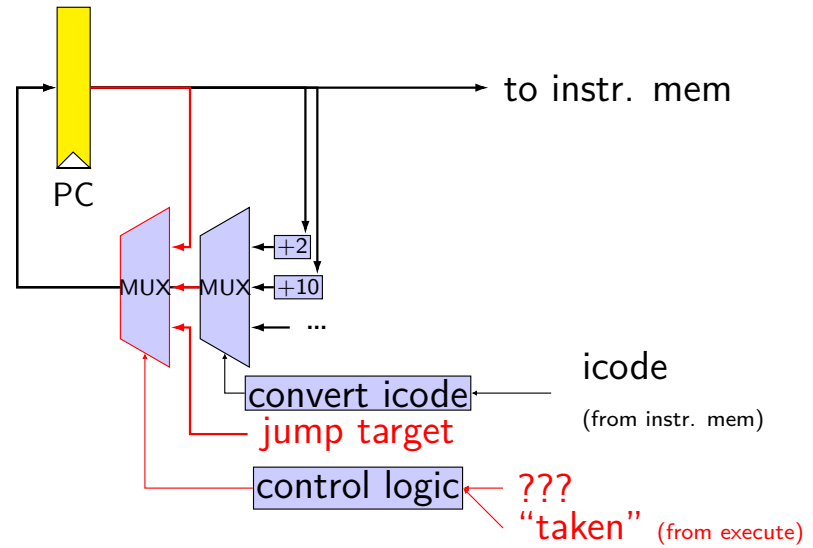
24

PC update (revised)



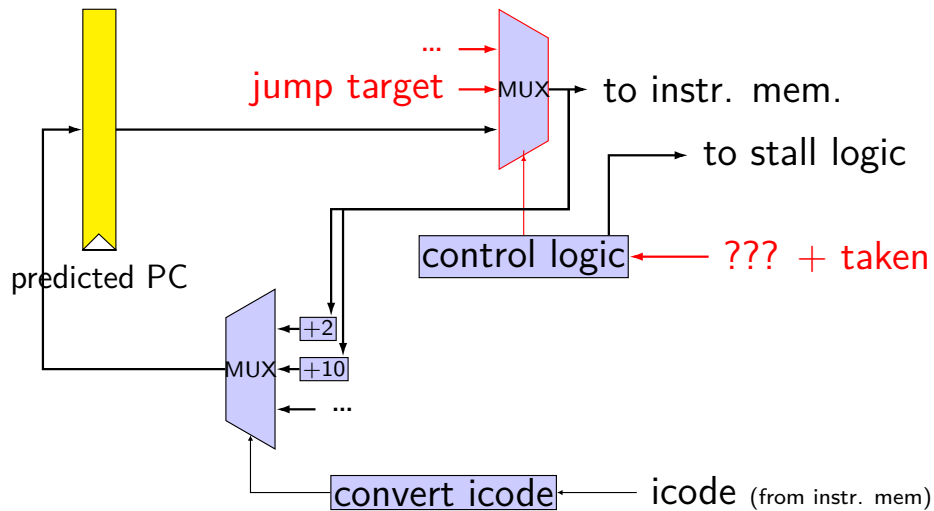
25

PC update (revised)



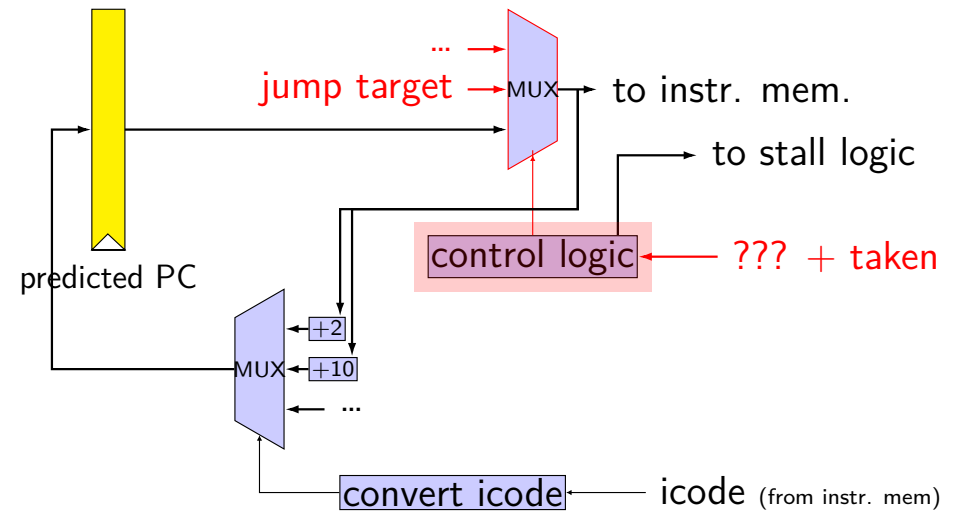
25

PC update (rearranged)



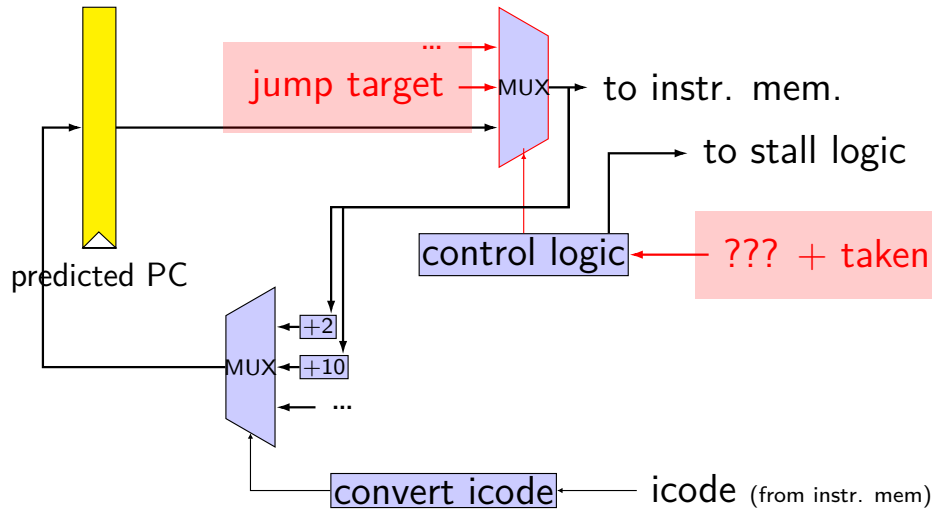
26

PC update (rearranged)



26

PC update (rearranged)



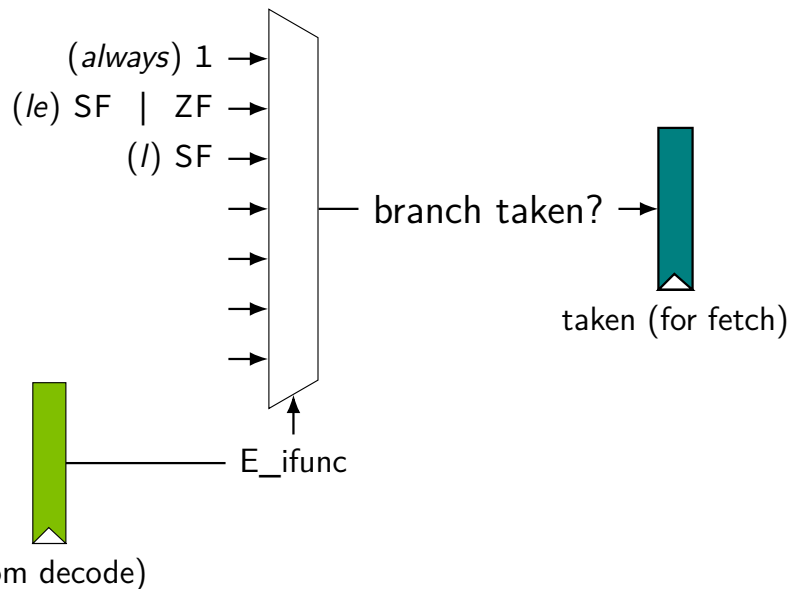
26

HCL2D PC update

```
pc_for_imem = [
  /* where do these come from? */
  conditionCodesSaidTaken : jumpTarget;
  ...
  1: P_predictedPC; /* a register */
];
```

27

"taken" signal (in execute stage)



28

jCC state tracking

time	fetch	decode	execute	memory	writeback
1	OPq				
2	jCC	OPq			
3	wait for jCC	jCC	OPq (set ZF)		
4	wait for jCC	nothing	jCC (use ZF)	OPq	
5	irmovq	nothing	nothing	jCC (done)	OPq

29

jCC state tracking

time	fetch	decode	execute	memory	writeback
1	OPq				
2	jCC	OPq	f_icode = JXX		
3	wait for jCC	jCC	OPq (set ZF)		
4	wait for jCC	nothing	jCC (use ZF)	OPq	
5	irmovq	nothing	nothing	jCC (done)	OPq

29

jCC state tracking

time	fetch	decode	execute	memory	writeback
1	OPq				
2	jCC	OPq			
3	wait for jCC	jCC	OPq (set ZF)	D_icode = JXX	
4	wait for jCC	nothing	jCC (use ZF)	OPq	
5	irmovq	nothing	nothing	jCC (done)	OPq

29

jCC state tracking

time	fetch	decode	execute	memory	writeback
1	OPq				
2	jCC	OPq			
3	wait for jCC	jCC	OPq (set ZF)		
4	wait for jCC	nothing	jCC (use ZF)	OPq	
5	irmovq	nothing	nothing	jCC (done)	OPq

E_icode = JXX

29

jCC state tracking

time	fetch	decode	execute	memory	writeback
1	OPq				
2	jCC	OPq			
3	wait for jCC	jCC	OPq (set ZF)		
4	wait for jCC	nothing	jCC (use ZF)	OPq	
5	irmovq	nothing	nothing	jCC (done)	OPq

M_icode = JXX

29

PC update

```
pc_for_imem = [
    /* when JXX gets to memory stage,
       then we might want to go to
       its jump target */
    M_icode == JXX && M_branchTaken:
        jumpTarget;
    M_icode == JXX && !M_branchTaken:
        predictedPC;
    ...;
]
```

30

ret

```
call empty
addq %r8, %r9
```

```
empty: ret
```

time	fetch	decode	execute	memory	writeback
1	call				
2	ret	call			
3	wait for ret	ret	call		
4	wait for ret	nothing	ret	call (store)	
5	wait for ret	nothing	nothing	ret (load)	call
6	addq	nothing	nothing	nothing	ret

31

ret

```
call empty
addq %r8, %r9
```

```
empty: ret
```

time	fetch	decode	execute	memory	writeback
1	call				
2	ret	call			
3	wait for ret	ret	call		
4	wait for ret	nothing	ret	call (store)	
5	wait for ret	nothing	nothing	ret (load)	call
6	addq	nothing	nothing	nothing	ret

return address stored here

31

ret

```
call empty
addq %r8, %r9
```

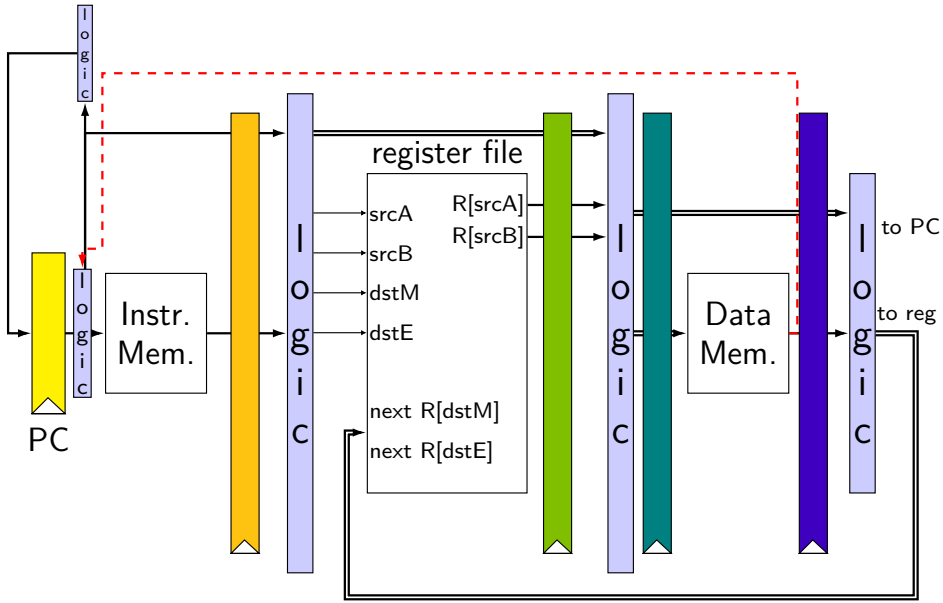
```
empty: ret
```

time	fetch	decode	execute	memory	writeback
1	call				
2	ret	call			
3	wait for ret	ret	call		
4	wait for ret	nothing	ret	call (store)	
5	wait for ret	nothing	nothing	ret (load)	call
6	addq	nothing	nothing	nothing	ret

why not start addq here?

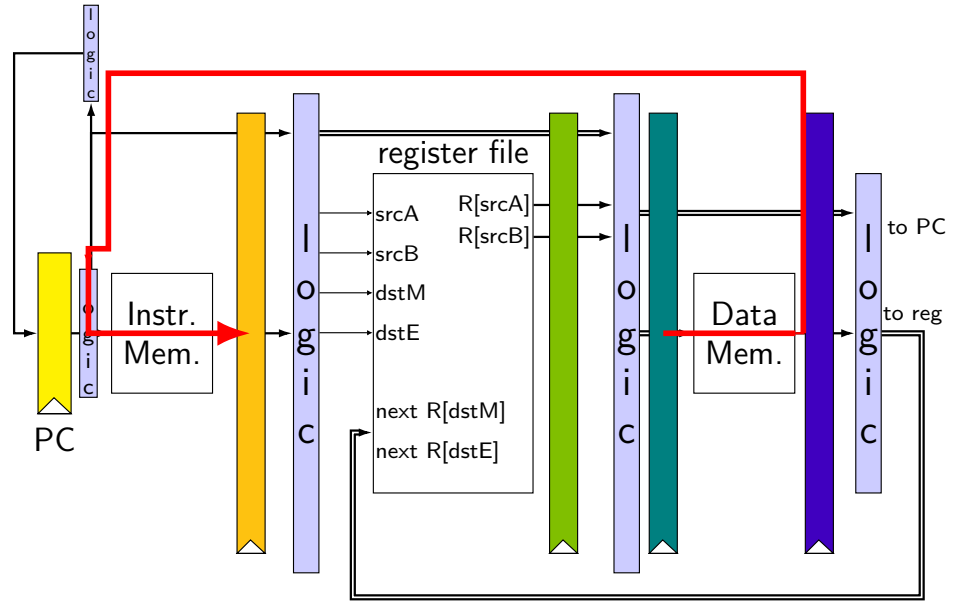
31

why not memory to PC



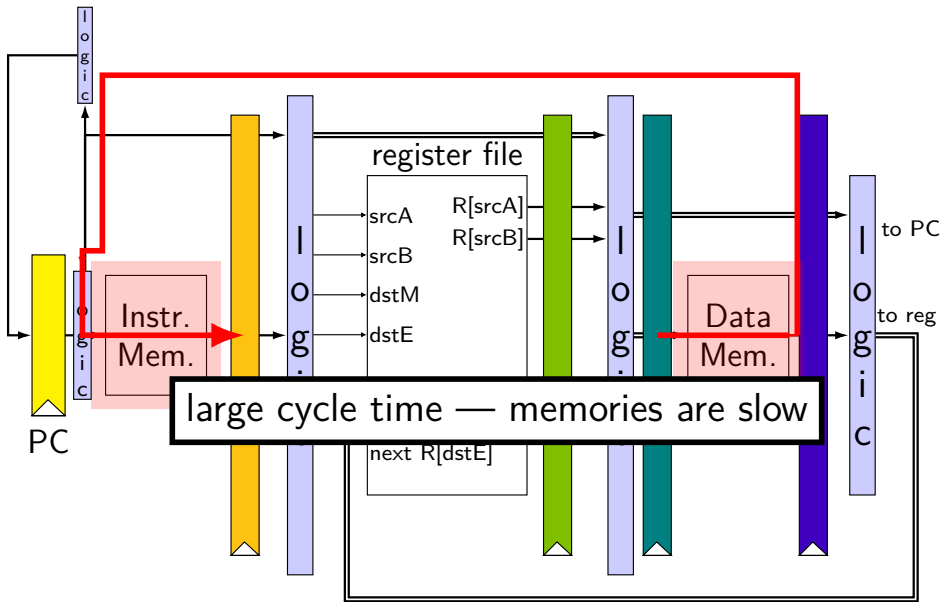
32

why not memory to PC



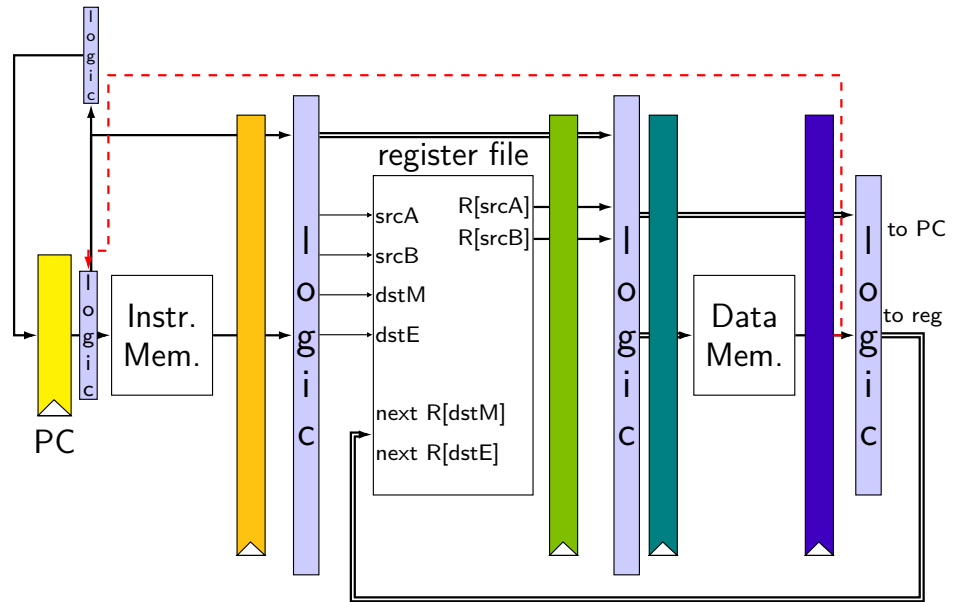
32

why not memory to PC



32

ret wiring



33

when do instructions change things?

... other than pipeline registers/PC:

stage	changes
fetch	(none)
decode	(none)
execute	condition codes
memory	memory writes
writeback	register writes/stat changes

34

when do instructions change things?

... other than pipeline registers/PC:

stage	changes
fetch	(none)
decode	(none)
execute	condition codes
memory	memory writes
writeback	register writes/stat changes

to “undo” instruction during fetch/decode/execute:

suppress condition code update (if any)
forget everything in pipeline registers

34

making guesses

```
subq  %rcx, %rax
jne   LABEL
xorq  %r10, %r11
xorq  %r12, %r13
```

```
LABEL: addq  %r8, %r9
        rmmovq %r10, 0(%r11)
```

speculate: **jne** will goto LABEL

right: 2 cycles faster!

wrong: forget before execute finishes

35

jXX: speculating right

```
subq  %r8, %r8
jne   LABEL
...
```

```
LABEL: addq  %r8, %r9
        rmmovq %r10, 0(%r11)
        irmovq $1, %r11
```

time	fetch	decode	execute	memory	writeback
1	subq				
2	jne	subq			
3	addq [?]	jne	subq (set ZF)		
4	rmmovq [?]	addq [?]	jne (use ZF)	OPq	
5	irmovq	rmmovq	addq	jne (done)	OPq

36

jXX: speculating right

```
subq %r8, %r8
jne LABEL
...
```

```
LABEL: addq %r8, %r9
rmmovq %r10, 0(%r11)
irmovq $1, %r11
```

time	fetch	decode	execute	memory	writeback
1	subq				
2	jne	subq			
3	addq [?]	jne	subq (set ZF)		
4	rmmovq [?]	addq [?]	j were waiting/nothing		
5	irmovq	rmmovq	addq	jne (done)	OPq

36

jXX: speculating wrong

```
subq %r8, %r8
jne LABEL
xorq %r10, %r11
...
```

```
LABEL: addq %r8, %r9
rmmovq %r10, 0(%r11)
```

time	fetch	decode	execute	memory	writeback
1	subq				
2	jne	subq			
3	addq [?]	jne	subq (set ZF)		
4	rmmovq [?]	addq [?]	jne (use ZF)	OPq	
5	xorq	nothing	nothing	jne (done)	OPq

37

jXX: speculating wrong

```
subq %r8, %r8
jne LABEL
xorq %r10, %r11
...
```

```
LABEL: addq %r8, %r9
rmmovq %r10, 0(%r11)
```

time	fetch	decode	execute	memory	writeback
1	subq				
2	jne	"squash" wrong guesses			
3	addq [?]	jne	subq (set ZF)		
4	rmmovq [?]	addq [?]	jne (use ZF)	OPq	
5	xorq	nothing	nothing	jne (done)	OPq

37

jXX: speculating wrong

```
subq %r8, %r8
jne LABEL
xorq %r10, %r11
...
```

```
LABEL: addq %r8, %r9
rmmovq %r10, 0(%r11)
```

time	fetch	decode	execute	memory	writeback
1	subq				
2	jne	subq			
3	addq [?]	jne	subq (set ZF)		
4	rmmovq [?]	addq [?]	jne (use ZF)	OPq	
5	xorq	nothing	nothing	jne (done)	OPq

37

jXX control logic

branch prediction **simplifies** — no stalling logic

... but extra logic to “squash” mispredicted instructions

38

jCC: assume taken?

book's choice: guess all jXXs are taken

empirical observation: most are

intuition:

```
irmovq $100, %rax
irmovq $1, %rbx
LOOP:
...
subq %rbx, %rax
je LOOP // taken 99% of the time
```

39

Performance

hypothetical instruction mix

kind	portion	cycles (predict)	cycles (stall)
not-taken jXX	3%	3	3
taken jXX	5%	1	3
ret	1%	4	4
others	91%	1*	1*

predict: $3 \times .03 + 1 \times .05 + 4 \times .01 + 1 \times .91 =$
1.09 cycles/instr.

stall: $3 \times .03 + 3 \times .05 + 4 \times .01 + 1 \times .91 =$
1.19 cycles/instr.

* — ignoring data hazards

40

stalling/misprediction and latency

case where pipeline **latency** matters

longer pipeline — larger penalty

part of Intel's Pentium 4 problem (c. 2000)

on release: 50% higher clock rate, **2-3x pipeline stages** of competitors

first-generation review quote:

For today's buyer, the Pentium 4 simply doesn't make sense. It's **slower** than the competition in just about every area, it's more expensive, it's using an interface that won't be the flashin

Review quote: Anand Lai Shimpi, "Intel Pentium 4 1.4 & 1.5 GHz", AnandTech, 20 November 2000

41

better branch prediction

forward (target > PC) not taken; backward taken

intuition: loops:

```
LOOP: ...
      ...
      je LOOP

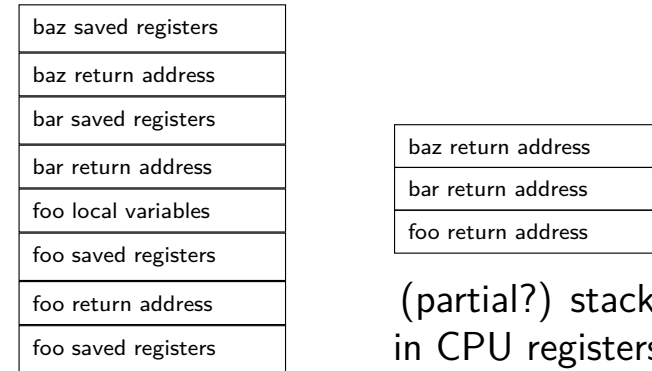
LOOP: ...
      jne SKIP_LOOP
      ...
      jmp LOOP
SKIP_LOOP:
```

42

predicting ret: extra copy of stack

predicting ret — stack in processor registers

different than real stack/out of room? just slower



stack in memory

(partial?) stack
in CPU registers

43

prediction before fetch

real processors can take **multiple cycles** to read instruction memory

predict branches **before reading their opcodes**

how — more extra data structures

44

summary

fetch/decode/execute/memory/writeback

add **pipeline registers**

normal next PC logic in fetch

branch prediction for jXX

assume taken; verify before following 'execute' finishes

stalling for ret

slower due to **branch misprediction** and **stalling**

45

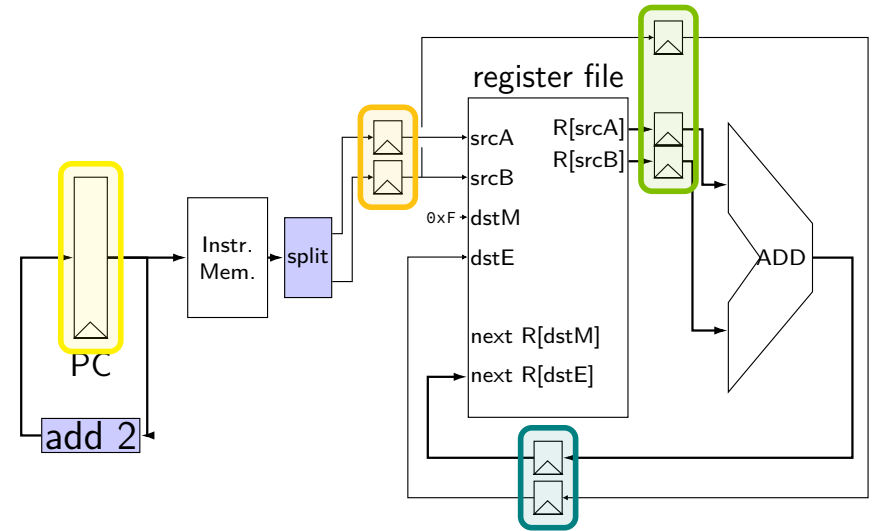
next classes

handling **data hazards** — mostly without stalling

details of stall/cancel logic

46

pipelined addq



47

exercise

```
register aB {
    value : 8 = 0xFF;
};
```

time	a_value	B_value & stall_B & bubble_B
0	0x01 & 0xFF & 0 & 0	
1	0x02 & ___ & 1 & 0	
2	0x03 & ___ & 0 & 0	
3	0x04 & ___ & 0 & 1	
4	0x05 & ___ & 0 & 0	
5	0x06 & ___ & 0 & 0	
6	0x07 & ___ & 1 & 0	
7	0x08 & ___ & 1 & 0	

48

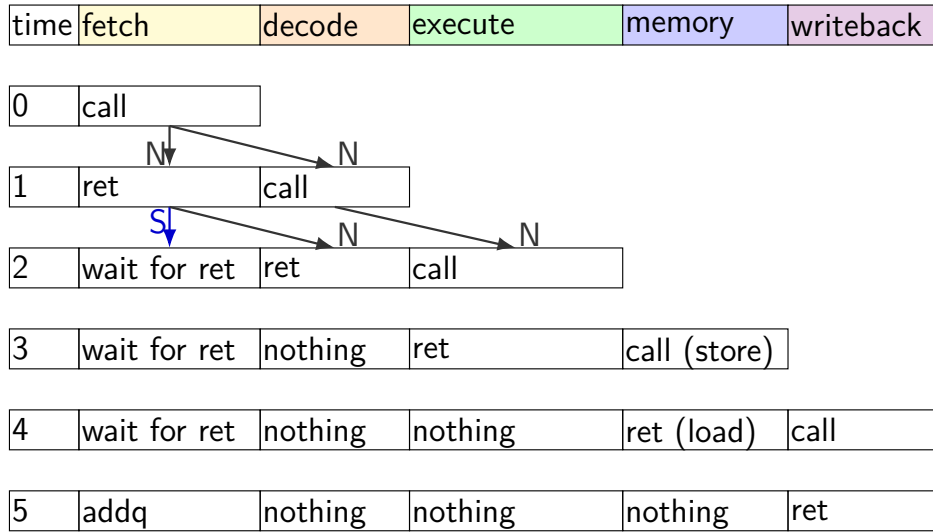
ret stall

time	fetch	decode	execute	memory	writeback
0	call				
1	ret	call			
2	wait for ret	ret	call		
3	wait for ret	nothing	ret	call (store)	
4	wait for ret	nothing	nothing	ret (load)	call
5	addq	nothing	nothing	nothing	ret

stall (S) = keep old value; normal (N) = use new value
 bubble (N) = use default (no-op);

49

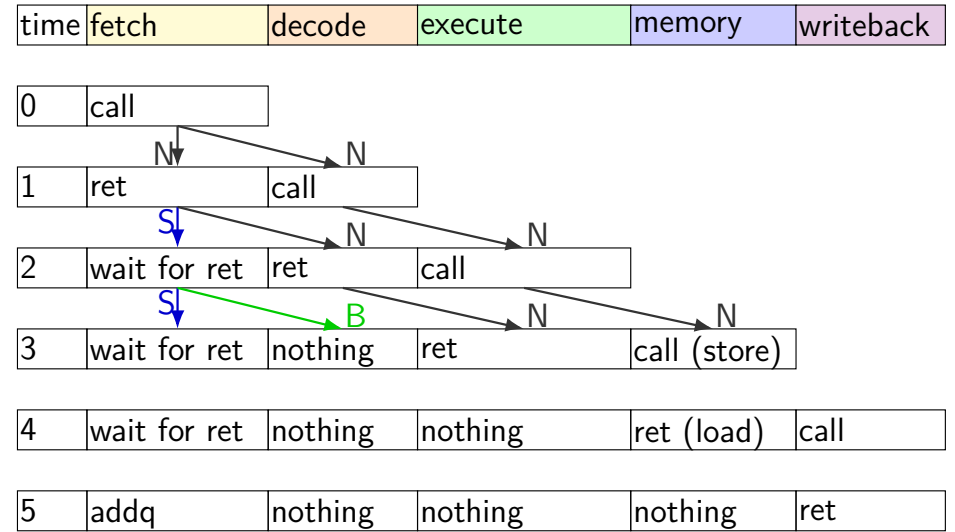
ret stall



stall (S) = keep old value; normal (N) = use new value
 bubble (N) = use default (no-op);

49

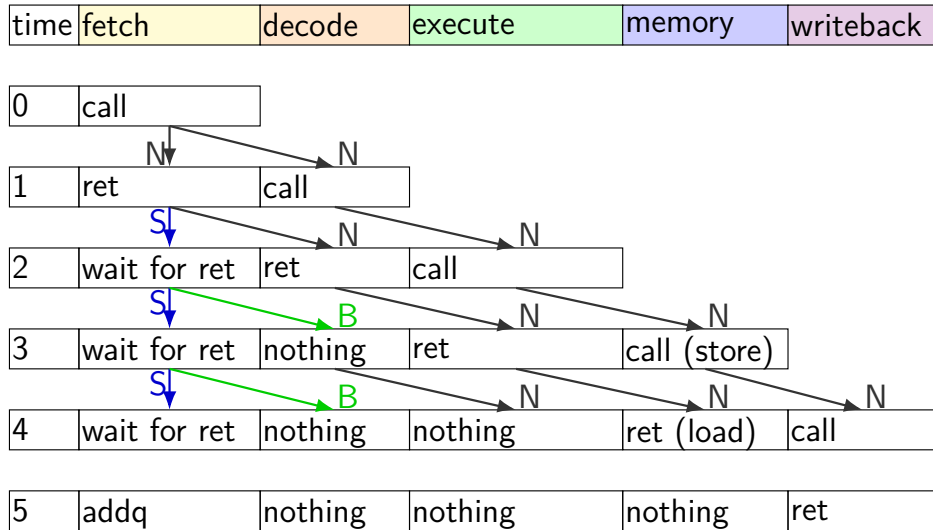
ret stall



stall (S) = keep old value; normal (N) = use new value
 bubble (N) = use default (no-op);

49

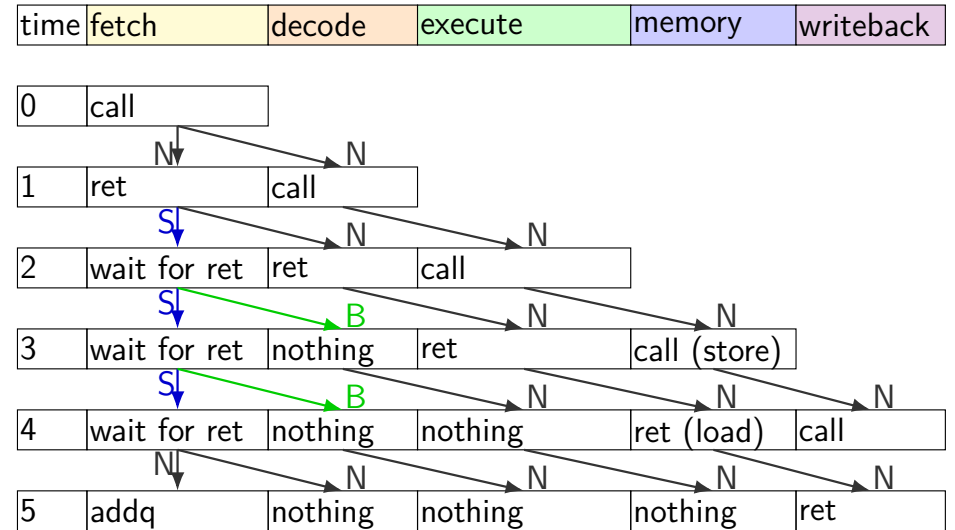
ret stall



stall (S) = keep old value; normal (N) = use new value
 bubble (N) = use default (no-op);

49

ret stall



stall (S) = keep old value; normal (N) = use new value
 bubble (N) = use default (no-op);

49