

Changelog

24 September 2020: SEQ srcA, srcB: remove “not planned to be included in our assignments” (we didn’t get to this slide in the AM lecture)

24 September 2020: instruction decode (1): add valC wire

24 September 2020: instruction decode (1): add pushq to list of options

SEQ part 3

last time

MUXes (HCLRS: “case expressions”)

N inputs + selector, 1 output

case expression shorthand — assumed translated to HW description

uses MUXes: `nop+halt`, `nop+halt+jmp`

lab: increment PC by right amount to find each instruction

register file operation

set source register number → value on output, same cycle

set destination register number → value on input stored just before next cycle

using register file: `addq`

using HCLRS: debug, interactive mode

aside: HCLRS -q/-quiet

-q option to HCLRS omits all but the last cycle's output

useful for checking “is this okay”

probably should omit it for “what's going on”

some anonymous feedback (1)

“I am very confused by the fact that, while the numerous problems with the presentation of the data were brought up in lecture, none of them have been corrected, and fifty percent of the homework relies on the usage of a table which the staff acknowledged was flawed.”

I assumed that students had dealt with understanding how to extract info in lab, so I was too late for a new version (with own formatting quirks) to be helpful...

Probably not a good assumption

though it seems describing the categories is probably a more important issue

some anonymous feedback (2)

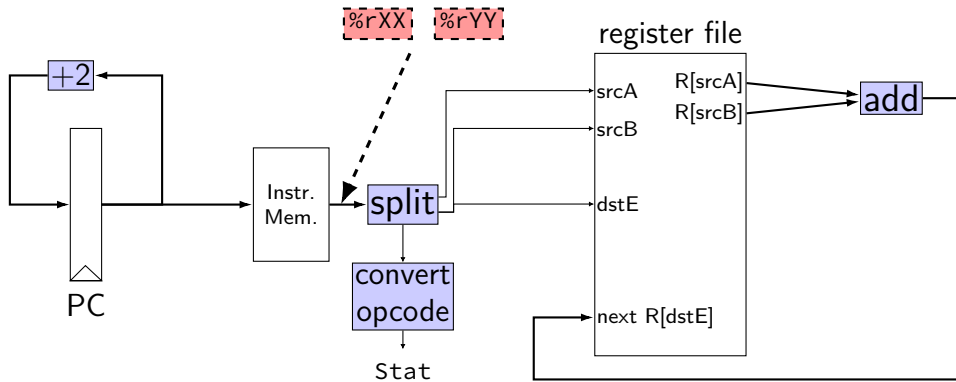
“In the future the benchmark data that we have in text documents, if that could be put in a spread sheet that would help my poor eyeballs read all the lines and numbers”

“The multiple parts within Part 3a and Part 3b seemed repetitive, and I think I would have learned the same idea without inputting the tedious calculations over and over for these. I don't mean this comment in a rude way at all, as I didn't mind doing this, especially if there is a clear learning goal. I just am wondering what the goal of doing these questions over and over is, as I think the lab could be improved if the learning goal could be met without the repetition.”


Yeah, should have supplied as spreadsheets to make this efficient for you
wanted multiple benchmarks because in practice it's important to not draw too many conclusions from just one program
perhaps too much

Also some clarity issues not mentioned in these feedbacks

addq CPU: HCL



addq CPU: HCL




```
register pP {  
    pc : 64 = 0;  
}  
p_pc = P_pc + 2;  
pc = P_pc;
```

```
wire opcode : 4;  
wire rA : 4, rB : 4;  
opcode = i10bytes[4..8];  
rA = i10bytes[12..16];  
rB = i10bytes[8..12];
```

```
reg_srcA = rA;  
reg_srcB = rB;  
reg_dstE = rB;  
reg_inputE =  
    reg_outputA +  
    reg_outputB;
```


addq CPU: HCL




```
register pP {  
    pc : 64 = 0;  
}
```

```
p_pc = P_pc + 2;  
pc = P_pc;
```

```
wire opcode : 4;  
wire rA : 4, rB : 4;  
opcode = i10bytes[4..8];  
rA = i10bytes[12..16];  
rB = i10bytes[8..12];
```

```
reg_srcA = rA;  
reg_srcB = rB;  
reg_dstE = rB;  
reg_inputE =  
    reg_outputA +  
    reg_outputB;
```

addq CPU: HCL

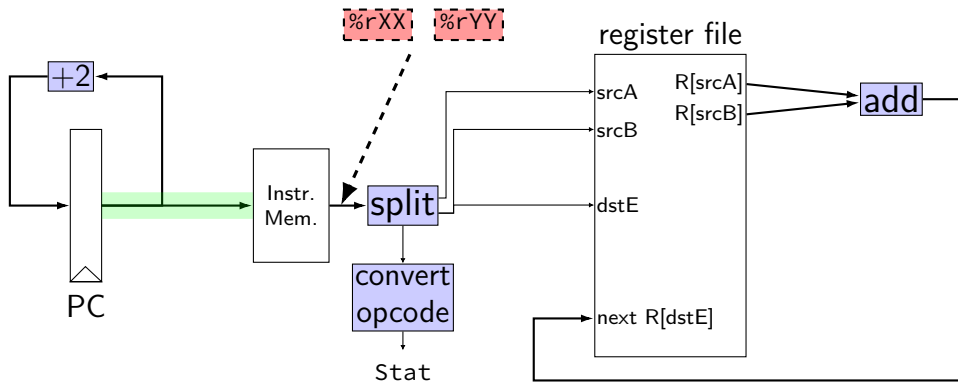


```
register pP {  
    pc : 64 = 0;  
}  
p_pc = P_pc + 2;  
pc = P_pc;
```

```
wire opcode : 4;  
wire rA : 4, rB : 4;  
opcode = i10bytes[4..8];  
rA = i10bytes[12..16];  
rB = i10bytes[8..12];
```

```
reg_srcA = rA;  
reg_srcB = rB;  
reg_dstE = rB;  
reg_inputE =  
    reg_outputA +  
    reg_outputB;
```

addq CPU: HCL




```
register pP {  
  pc : 64 = 0;  
}  
p_pc = P_pc + 2;  
pc = P_pc;
```

```
wire opcode : 4;  
wire rA : 4, rB : 4;  
opcode = i10bytes[4..8];  
rA = i10bytes[12..16];  
rB = i10bytes[8..12];
```

```
reg_srcA = rA;  
reg_srcB = rB;  
reg_dstE = rB;  
reg_inputE =  
  reg_outputA +  
  reg_outputB;
```

addq CPU: HCL




```
register pP {  
    pc : 64 = 0;  
}  
p_pc = P_pc + 2;  
pc = P_pc;
```

```
wire opcode : 4;  
wire rA : 4, rB : 4;  
opcode = i10bytes[4..8];  
rA = i10bytes[12..16];  
rB = i10bytes[8..12];
```

```
reg_srcA = rA;  
reg_srcB = rB;  
reg_dstE = rB;  
reg_inputE =  
    reg_outputA +  
    reg_outputB;
```

addq CPU: HCL

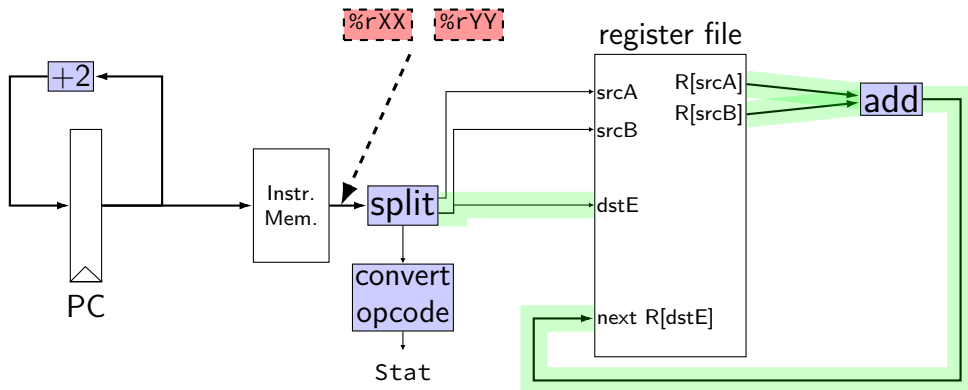


```
register pP {  
  pc : 64 = 0;  
}  
p_pc = P_pc + 2;  
pc = P_pc;
```

```
wire opcode : 4;  
wire rA : 4, rB : 4;  
opcode = i10bytes[4..8];  
rA = i10bytes[12..16];  
rB = i10bytes[8..12];
```

```
reg_srcA = rA;  
reg_srcB = rB;  
reg_dstE = rB;  
reg_inputE =  
  reg_outputA +  
  reg_outputB;
```

addq CPU: HCL




```
register pP {  
    pc : 64 = 0;  
}  
p_pc = P_pc + 2;  
pc = P_pc;
```

```
wire opcode : 4;  
wire rA : 4, rB : 4;  
opcode = i10bytes[4..8];  
rA = i10bytes[12..16];  
rB = i10bytes[8..12];
```

```
reg_srcA = rA;  
reg_srcB = rB;  
reg_dstE = rB;  
reg_inputE =  
    reg_outputA +  
    reg_outputB;
```

ALUs



Operations needed:
add — **addq**, addresses
sub — **subq**
xor — **xorq**
and — **andq**
more?

ALUs not for PC increment

our processor will have one ALU

not used for PC increment (computing next instruction address)

- need to do other computation in same cycle

- don't need a general circuit for it

ALUs in HCLRS

HCLRS doesn't supply an ALU

the HCL the textbook authors use does

...but you can build one yourself

not required — we check functionality

more instructions


we've seen add and nop+halt+jmp

really want all the instructions

combine by adding MUXes

same procedure as adding halt, jmp to existing CPUs

nop/halt CPU



```
register pP {  
    thePc : 64 = 0;  
}  
p_thePc = P_thePc + 1;  
pc = P_thePc;  
Stat = [  
    i10bytes[4..8] == NOP : STAT_AOK;  
    i10bytes[4..8] == HALT : STAT_HLT;  
    1 : STAT_INS; // (default case)  
];
```

exercise: nop/add CPU

Let's say we wanted to make a **add+nop CPU**. Where would we need MUXes? Before...

(modify add CPU to also support the nop instruction)

- A. one or both of the register file 'register number to read' inputs (reg_src...)
- B. the PC register's input (p_pc)
- C. one of the register file 'register number to write' inputs (reg_dst...)
- D. one of the register file 'register value to write' inputs (reg_input...)
- E. the instruction memory's address input (pc)



simple ISA: mov-to-register

```
irmovq $constant, %rYY
```




```
rrmovq %rXX, %rYY
```

```
mrmovq 10(%rXX), %rYY
```





two memories





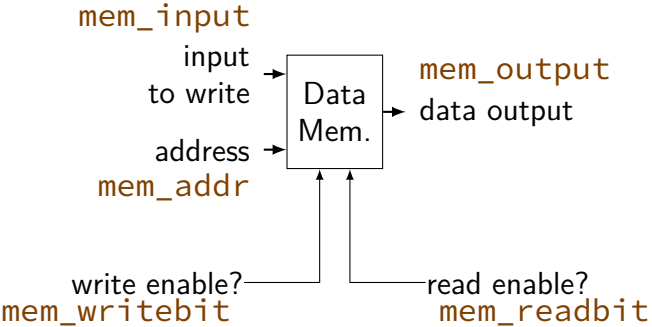

two memories



two memories



two memories



really two memories??

in Y86-64 (and many real CPUs):

writing to address X in data memory:

changes address X in instruction memory

really two memories??

in Y86-64 (and many real CPUs):

writing to address X in data memory:

changes address X in instruction memory

so really just one memory??

we'll explain when we talk about *caches*

exercise: mov-to-register

```
irmovq $constant, %rYY
```

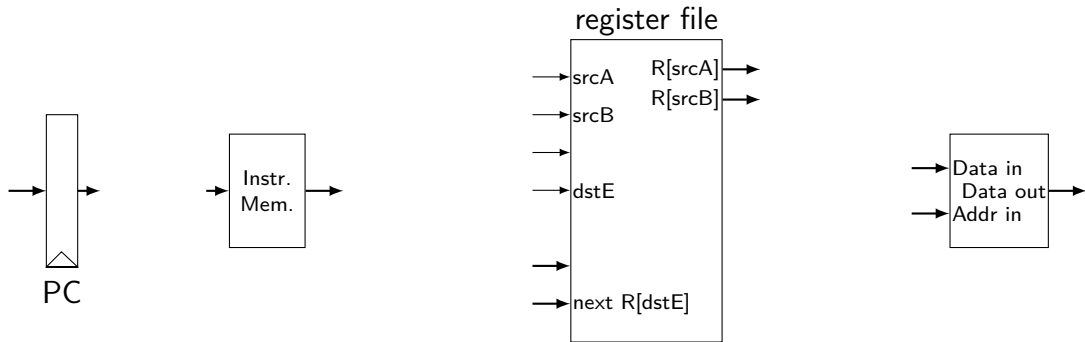
```
rrmovq %rXX, %rYY
```

```
mrmovq 10(%rXX), %rYY
```

for which are these are we going to need MUXes? before...

- A. register file's register number (index) inputs (reg_srcA, reg_srcB, reg_dstE, ...)
- B. register file's value inputs (reg_inputE/M)
- C. PC register's input
- D. instruction memory's address input (pc)


mov-to-register CPU



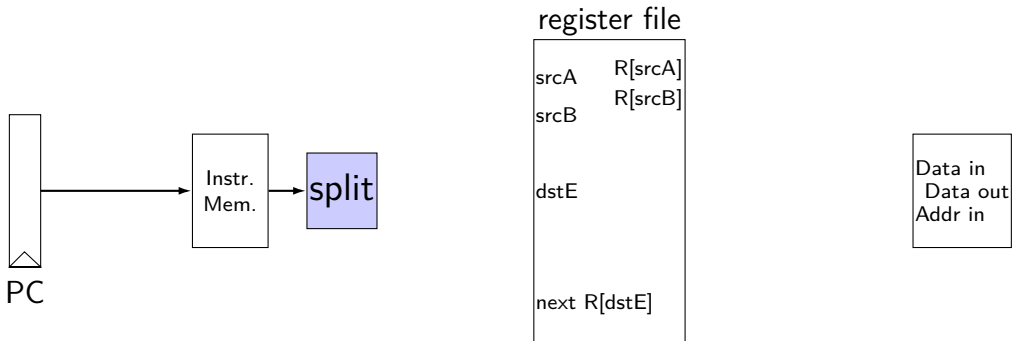
`rrmovq rA, rB`

`irmovq V, rB`

`mrmovq D(rB), rA`




mov-to-register CPU



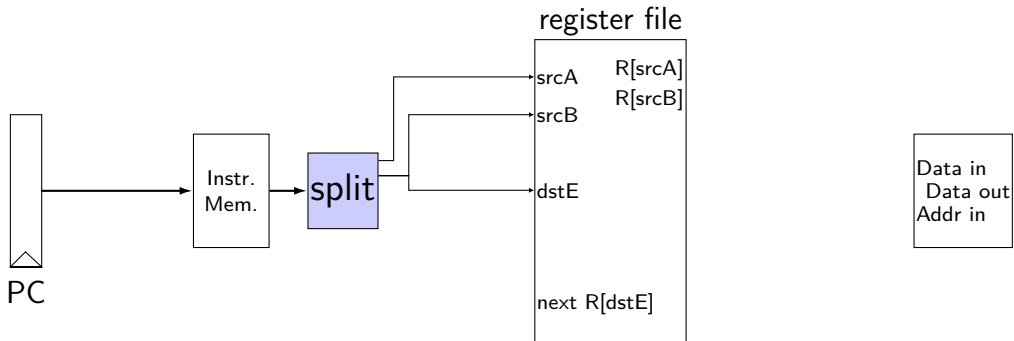
`rrmovq rA, rB`

`irmovq V, rB`

`mrmovq D(rB), rA`




mov-to-register CPU



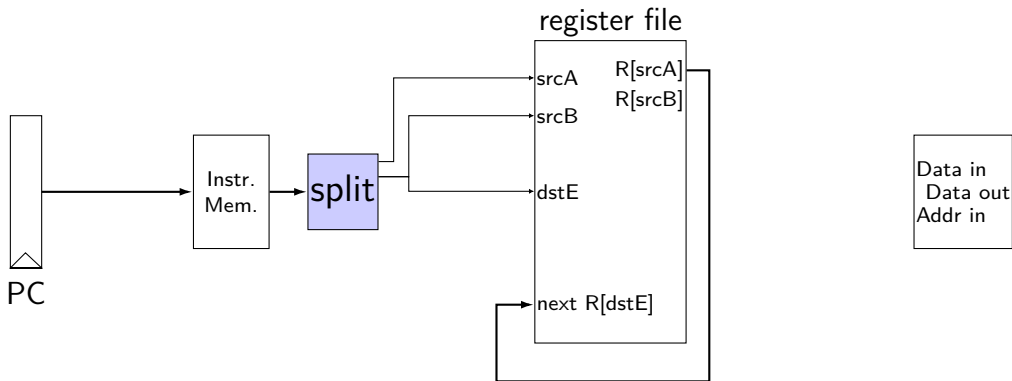
`rrmovq rA, rB`

`irmovq V, rB`

`mrmovq D(rB), rA`



mov-to-register CPU



`rrmovq rA, rB`



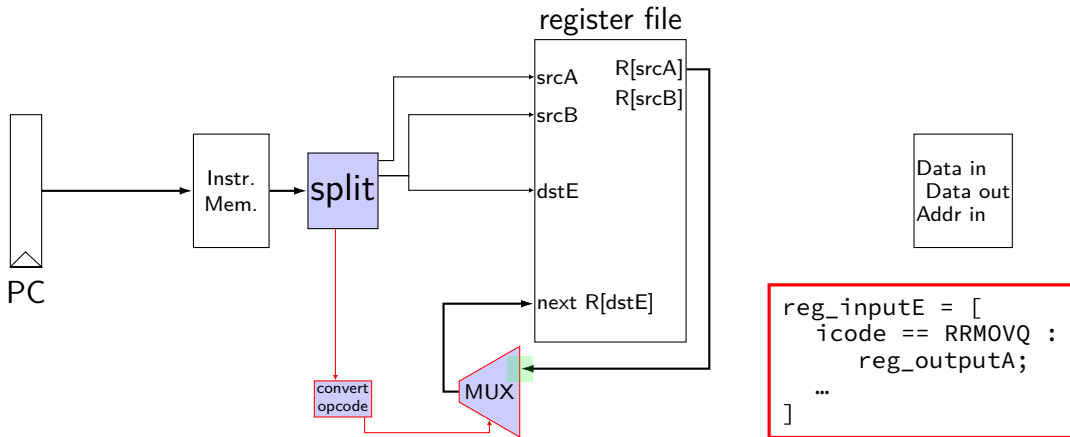
`irmovq V, rB`



`rrmovq D(rB), rA`



mov-to-register CPU



rrmovq *rA*, *rB*



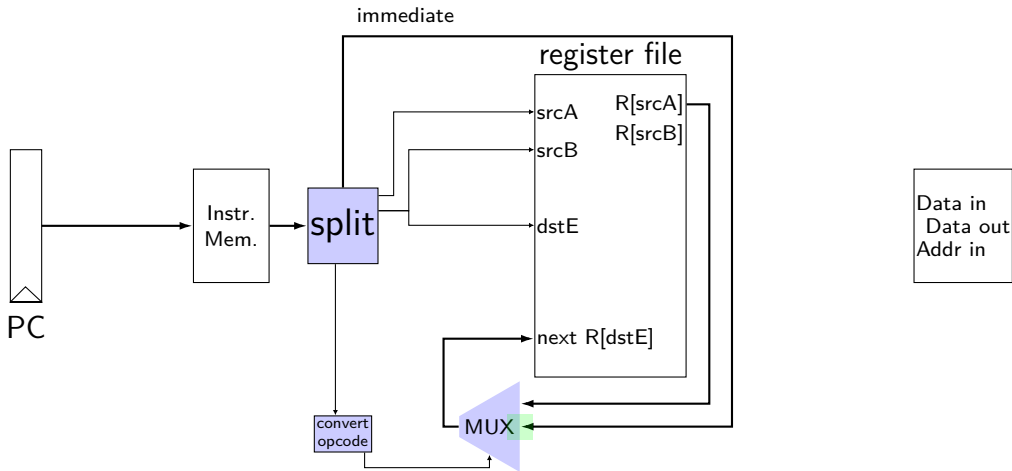
irmovq *V*, *rB*



mrmovq *D(rB)*, *rA*



mov-to-register CPU



Data in
Data out
Addr in


`rrmovq rA, rB`

`irmovq V, rB`

`mrmovq D(rB), rA`

2	0	rA	rB	
3	0	F	rB	V
5	0	rA	rB	D

mov-to-register CPU



`rrmovq rA, rB`



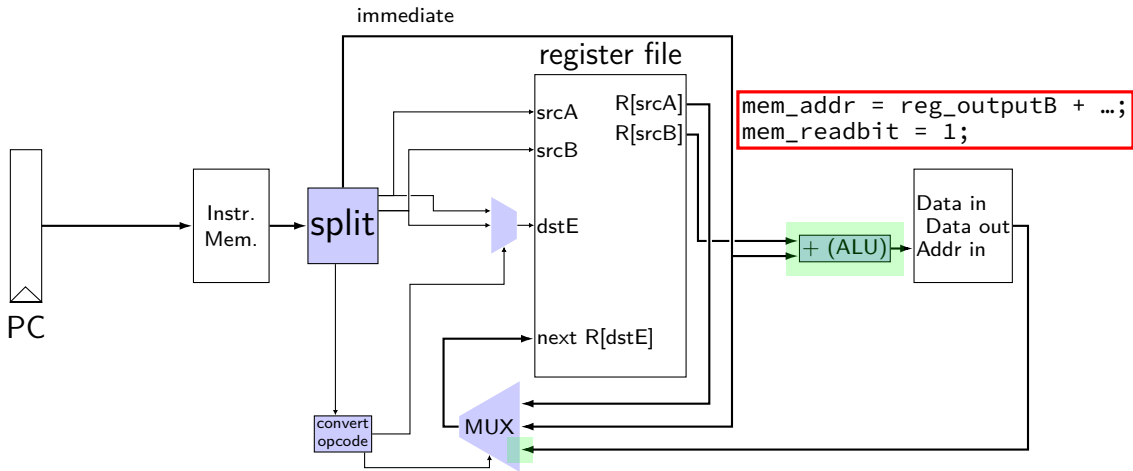
`irmovq V, rB`



`mrmovq D(rB), rA`



mov-to-register CPU



`rrmovq rA, rB`



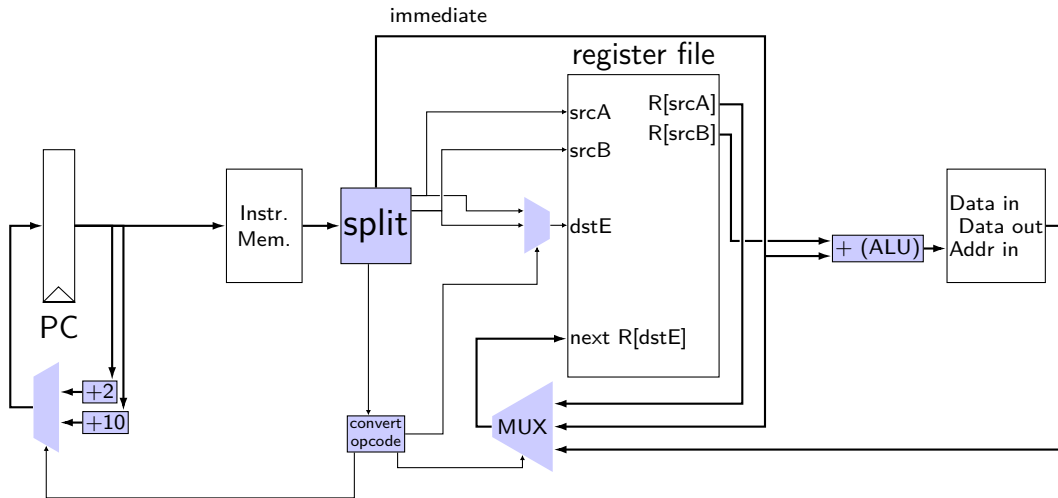
`irmovq V, rB`



`mrmovq D(rB), rA`



mov-to-register CPU



`rmmovq rA, rB`



`irmovq V, rB`



`mrmovq D(rB), rA`



simple ISA: mov (all cases)

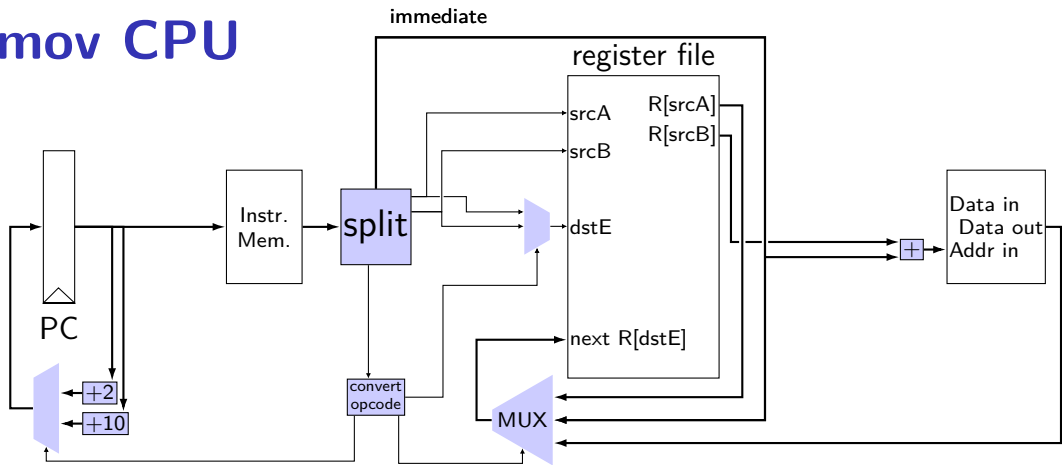
`irmovq $constant, %rYY`

`rrmovq %rXX, %rYY`

`mrmovq 10(%rXX), %rYY`

`rmmovq %rXX, 10(%rYY)`

mov CPU




`rrmovq rA, rB`



`irmovq V, rB`



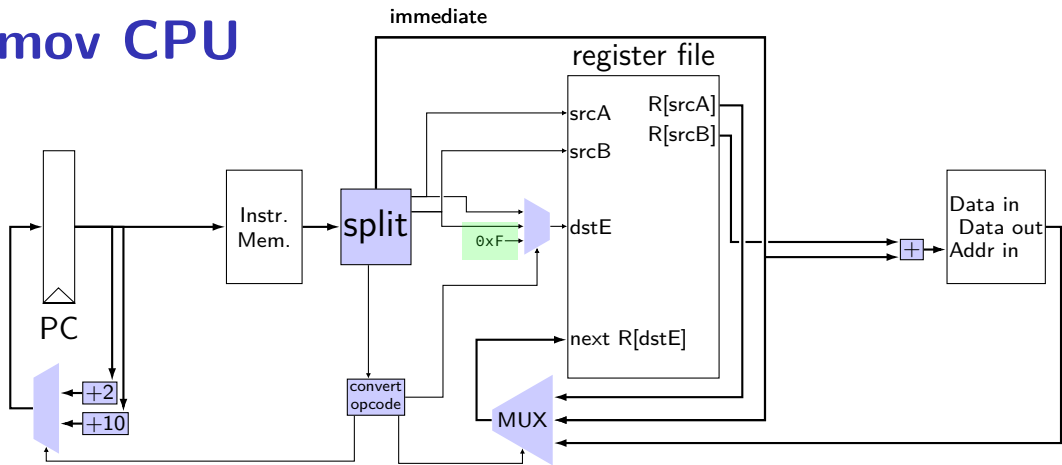
`rrmovq D(rB), rA`



`rmmovq rA, D(rB)`




mov CPU



`rrmovq rA, rB`



`irmovq V, rB`



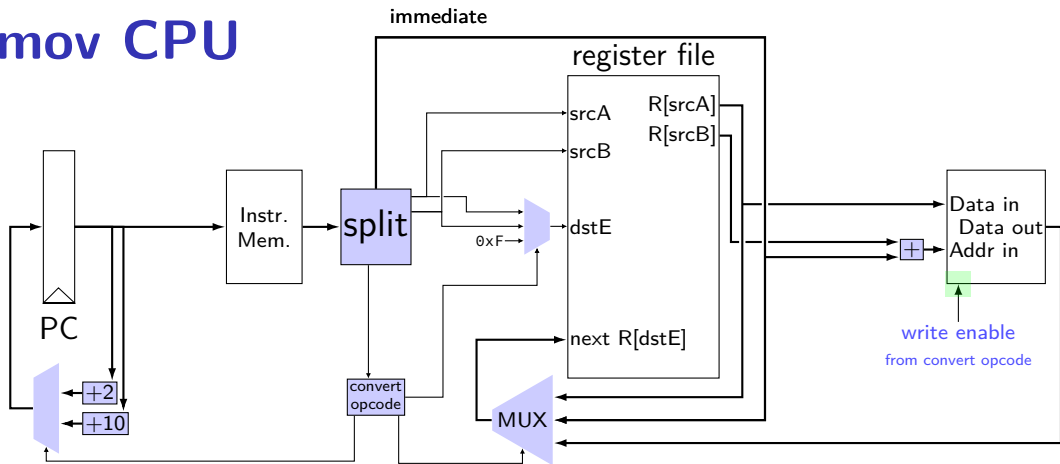
`mrmovq D(rB), rA`



`rmmovq rA, D(rB)`



mov CPU




`rrmovq rA, rB`



`irmovq V, rB`



`rrmovq D(rB), rA`



`rmmovq rA, D(rB)`



data path versus control path

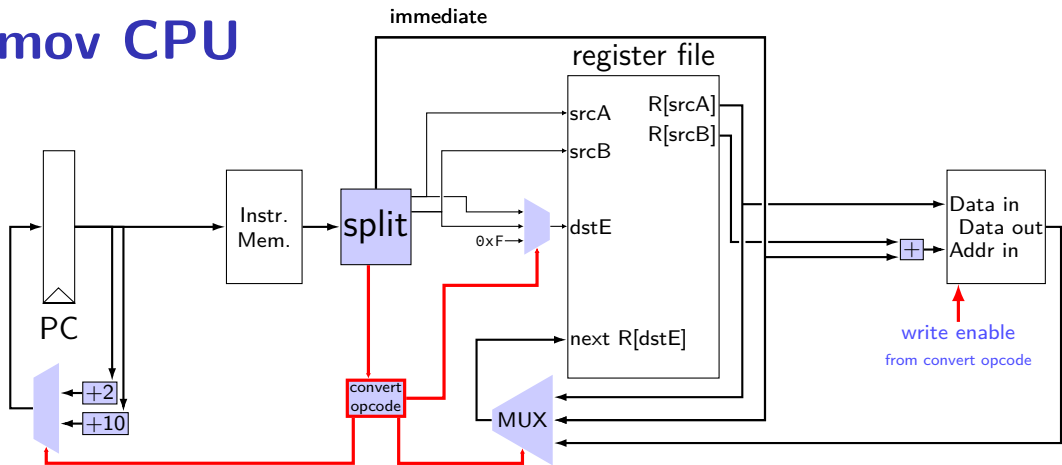
data path — signals carrying “actual data”

control path — signals that control MUXes, etc.

fuzzy line: e.g. are condition codes part of control path?

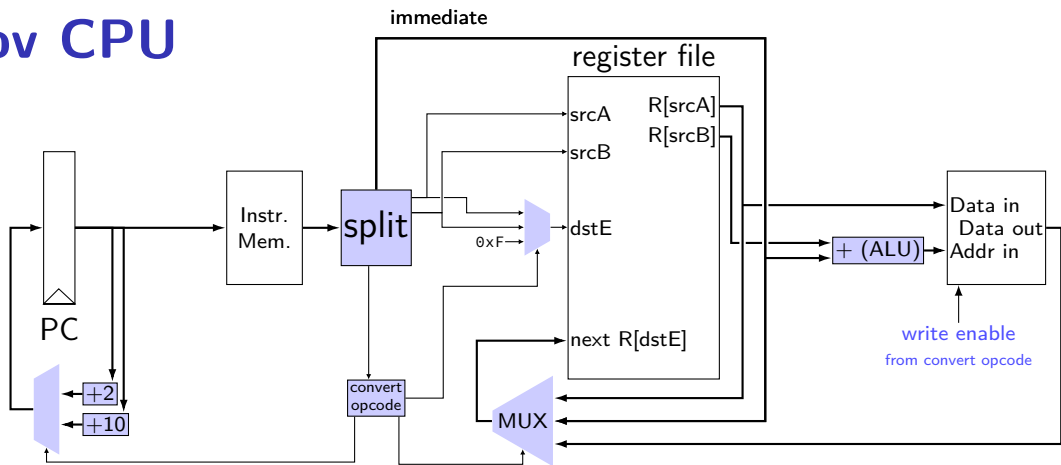
we will often omit parts of the control path in drawings, etc.

mov CPU




<code>rrmovq rA, rB</code>	2	0	rA	rB	
<code>irmovq V, rB</code>	3	0	F	rB	V
<code>mrmovq D(rB), rA</code>	5	0	rA	rB	D
<code>rmmovq rA, D(rB)</code>	4	0	rA	rB	D

mov CPU




mov CPU

PC update




fetch




immediate

decode




writeback

execute



memory



write enable
from convert opcode

Stages

conceptual division of instruction:

fetch — read instruction memory, split instruction, compute length

decode — read register file

execute — arithmetic (including of addresses)

memory — read or write data memory

write back — write to register file

PC update — compute next value of PC

stages and time

fetch / decode / execute / memory / write back / PC update

Order when these events happen pushq %rax instruction:

1. instruction read
2. memory changes
3. %rsp changes
4. PC changes

Hint: recall how registers, register files, memory works

- a. 1; then 2, 3, and 4 in any order
- b. 1; then 2, 3, and 4 at almost the same time
- c. 1; then 2; then 3; then 4
- d. 1; then 3; then 2; then 4
- e. 1; then 2; then 3 and 4 at almost the same time
- f. something else

SEQ: instruction fetch

read instruction memory at PC

split into separate wires:

icode:ifun — opcode


rA, rB — register numbers

valC — call target or mov displacement

compute next instruction address:

valP — $PC + (\text{instr length})$

instruction fetch



register file

srcA	R[srcA]
srcB	R[srcB]
dstM	
dstE	
next	R[dstM]
next	R[dstE]

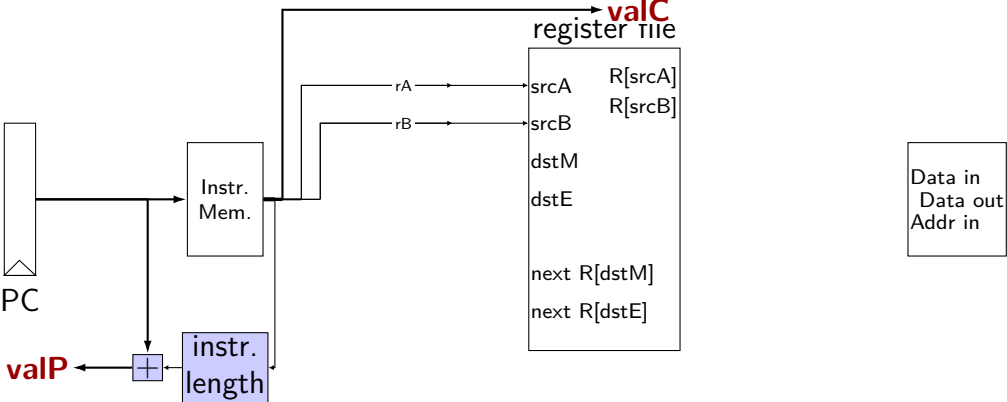
Data in
Data out
Addr in

SEQ: instruction “decode”


read registers

`valA`, `valB` — register values

instruction decode (1)



instruction decode (1)



exercise: which of these instructions can this **not** work for?

nop, addq, mrmovq, rmmovq, jmp, pushq **of these: only pushq**

SEQ: srcA, srcB

always read rA, rB?

Problems:

- push rA

- pop

- call


- ret

book: extra signals: srcA, srcB — computed input register

MUX controlled by icode


SEQ: possible registers to read

instruction	srcA	srcB
halt, nop, jCC, irmovq	none	none
cmovCC, rrmovq	rA	none
mrmovq	none	rB
rmmovq, OPq	rA	rB
call, ret	none?	%rsp
pushq, popq	rA	%rsp




SEQ: possible registers to read

instruction	srcA	srcB
halt, nop, jCC, irmovq	none	none
cmovCC, rrmovq	rA	none
rrmovq	none	rB
rmmovq, OPq	rA	rB
call, ret	none?	%rsp
pushq, popq	rA	%rsp



instruction decode (2)



SEQ: execute

perform ALU operation (add, sub, xor, and)


valE — ALU output

read prior condition codes

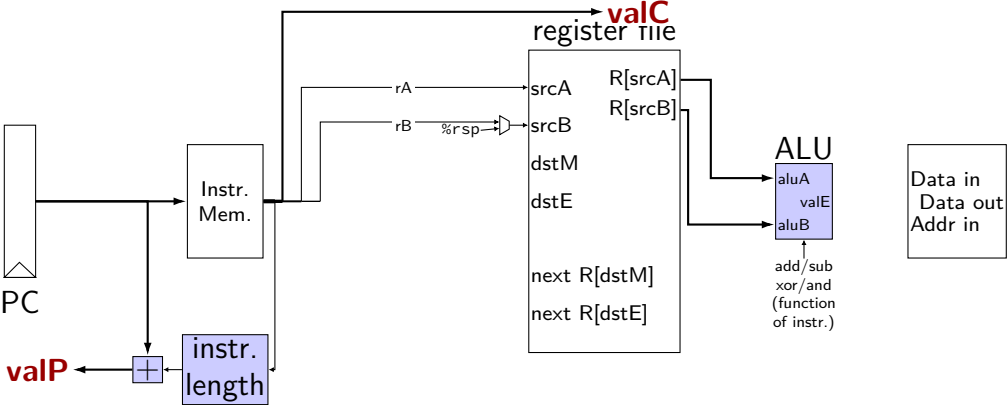
Cnd — condition codes based on ifun (instruction type for jCC/cmouvCC)

write new condition codes

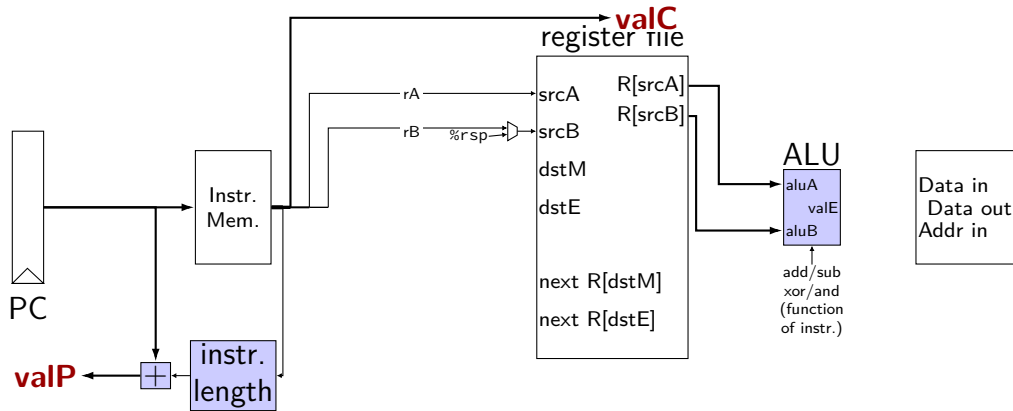
using condition codes: cmov



execute (1)



execute (1)




exercise: which of these instructions can this **not** work for?
nop, addq, mrmovq, popq, call,

SEQ: ALU operations?

ALU inputs always **valA**, **valB** (register values)?

no, inputs from instruction: (Displacement + rB)

mrmovq
rmmovq




no, constants: (rsp +/- 8) (not planned to be in assignments)

pushq
popq
call
ret

extra signals: **aluA**, **aluB**

computed ALU input values

execute (2)

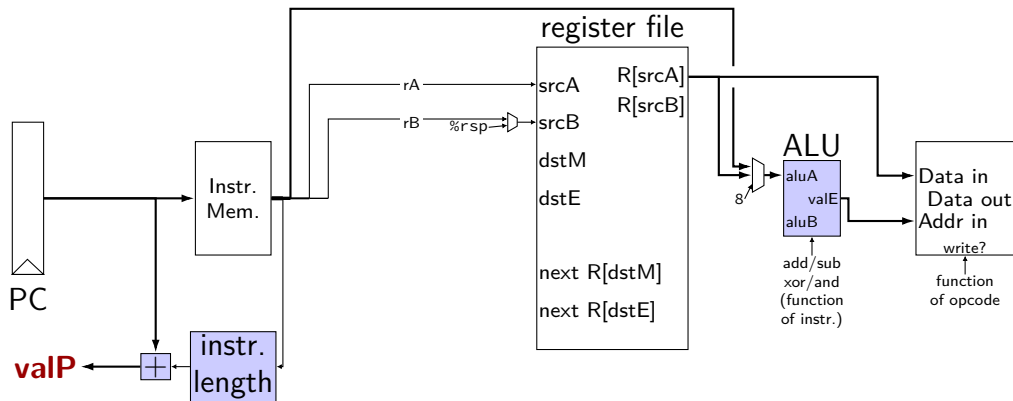


SEQ: Memory

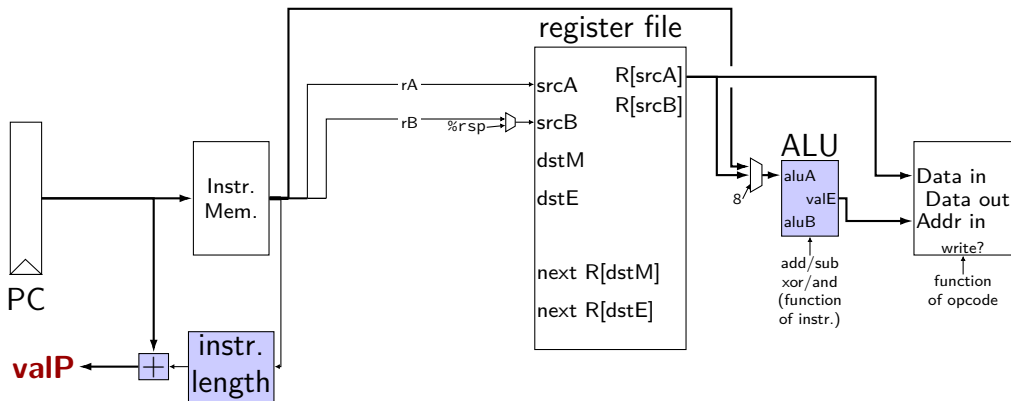
read or write data memory

valM — value read from memory (if any)

memory (1)



memory (1)



exercise: which of these instructions can this **not** work for?
nop, rmmovq, mrmovq, popq, call,

SEQ: control signals for memory

read/write — read enable? write enable?

Addr — address

mostly ALU output


special cases (need extra MUX): `popq`, `ret`

Data — value to write

mostly `valA`

special cases (need extra MUX): `call`


memory (2)




SEQ: write back

write registers

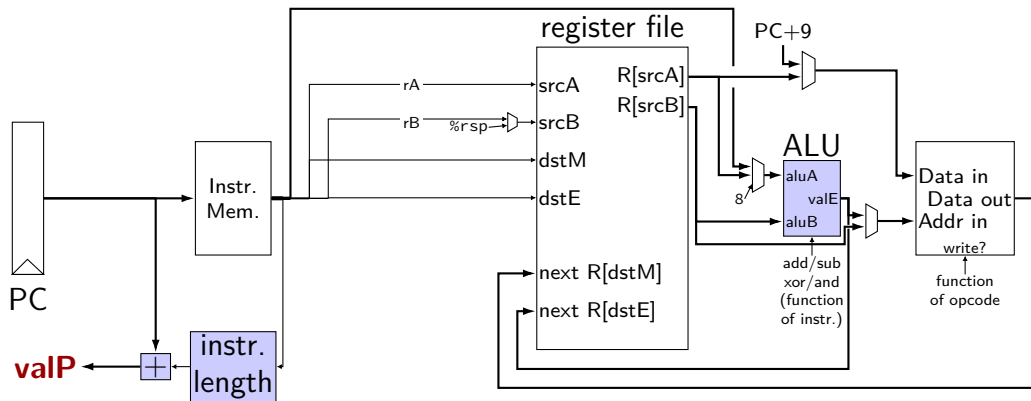
write back (1)



write back (1)



write back (1)




exercise: which of these instructions can this **not** work for?
nop, irmovq, mrmovq, rmmovq, addq

SEQ: control signals for WB

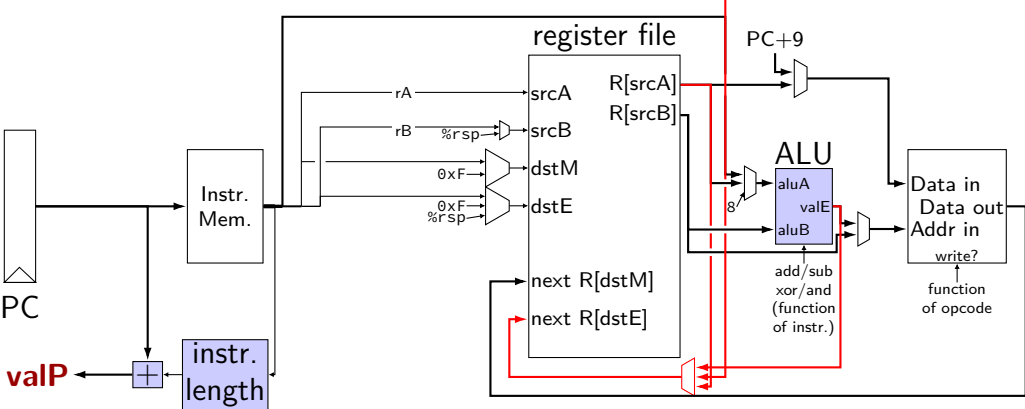
two write inputs — two needed by popq
valM (memory output), valE (ALU output)

two register numbers
dstM, dstE

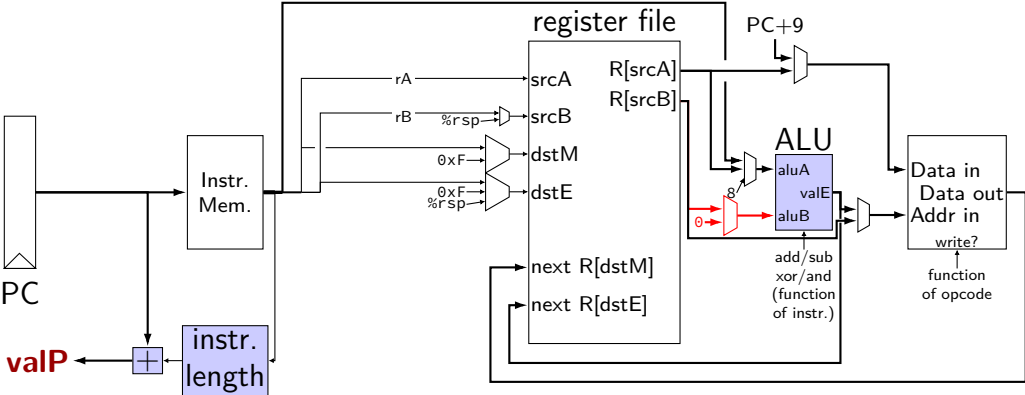
write disable — use dummy register number 0xF



write back (2a)



write back (2b)



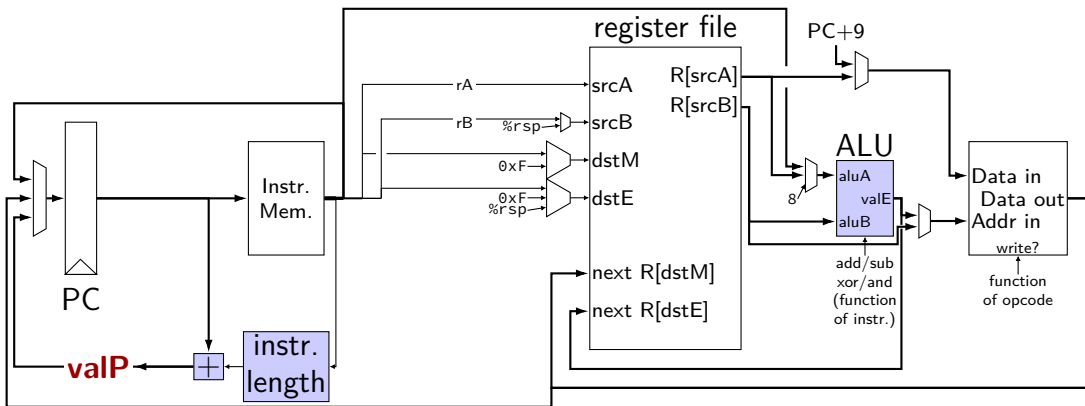
SEQ: Update PC

choose value for PC next cycle (input to PC register)

usually valP (following instruction)

exceptions: `call`, `jcc`, `ret`

PC update



backup slides

differences from book

wire not **bool** or **int**

book uses names like `val C` — not required!

author's environment limited adding new wires

MUXes must have default (`1 : something`) case

implement your own ALU

differences from book

wire not **bool** or **int**

book uses names like `valC` — not required!

author's environment limited adding new wires

MUXes must have default (`1 : something`) case

implement your own ALU

differences from book

wire not **bool** or **int**

book uses names like `val C` — not required!

author's environment limited adding new wires

MUXes must have default (`1 : something`) case

implement your own ALU

comparing to yis

```
$ ./hclrs nopjmp_cpu.hcl nopjmp.yo
```

```
...  
...
```

```
+----- (end of halted state) -----+
```

```
Cycles run: 7
```

```
$ ./tools/yis nopjmp.yo
```

```
Stopped in 7 steps at PC = 0x1e.  Status 'HLT', CC Z=1 S=0 O=0
```

```
Changes to registers:
```

```
Changes to memory:
```

HCLRS summary

declare/assign values to **wires**

MUXes with

```
[ test1: value1; test2: value2; 1: default; ]
```

register banks with **register** i0:

next value on i_name; current value on O_name

fixed functionality

register file (15 registers; 2 read + 2 write)

memories (data + instruction)

Stat register (start/stop/error)