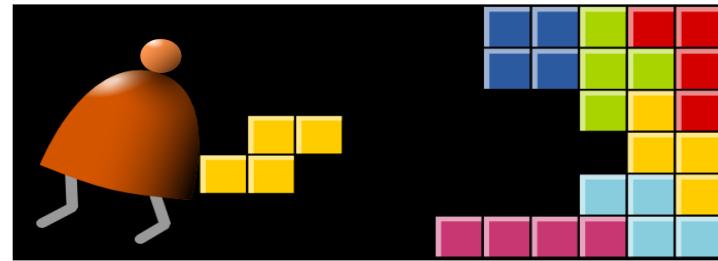


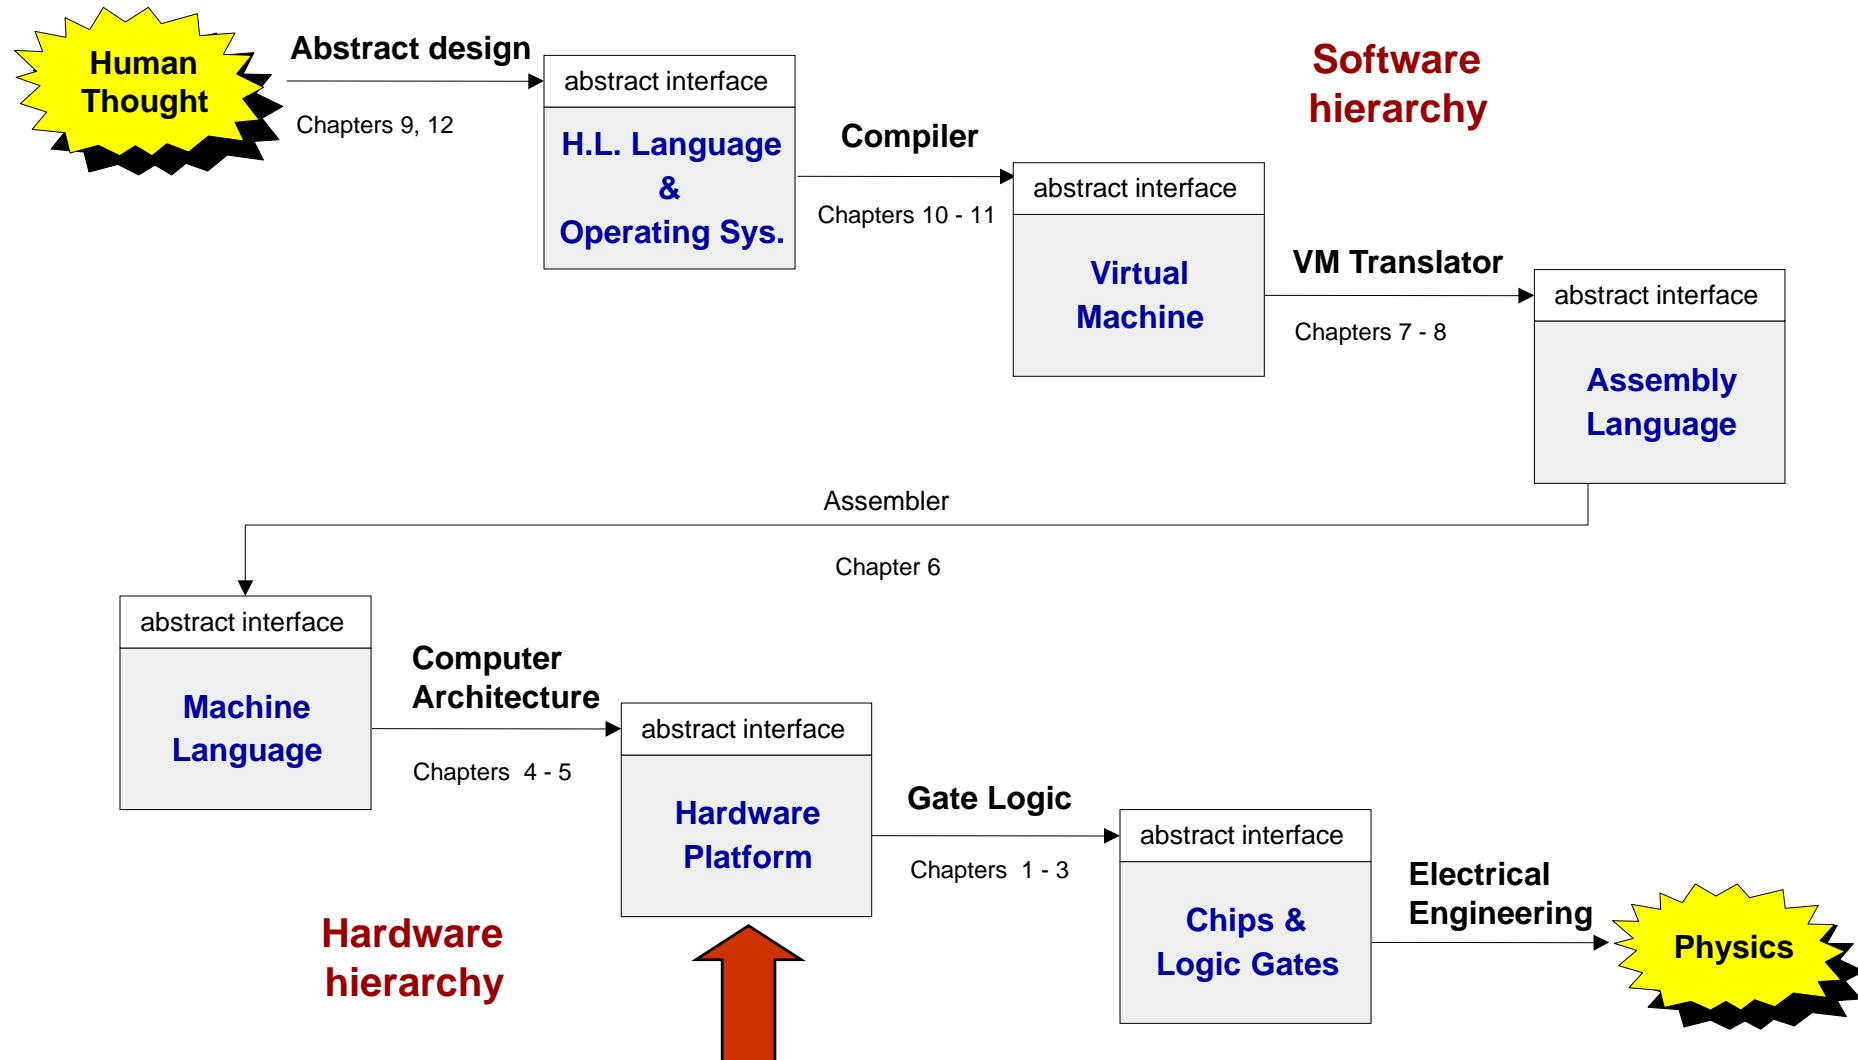
Computer Architecture



Building a Modern Computer From First Principles

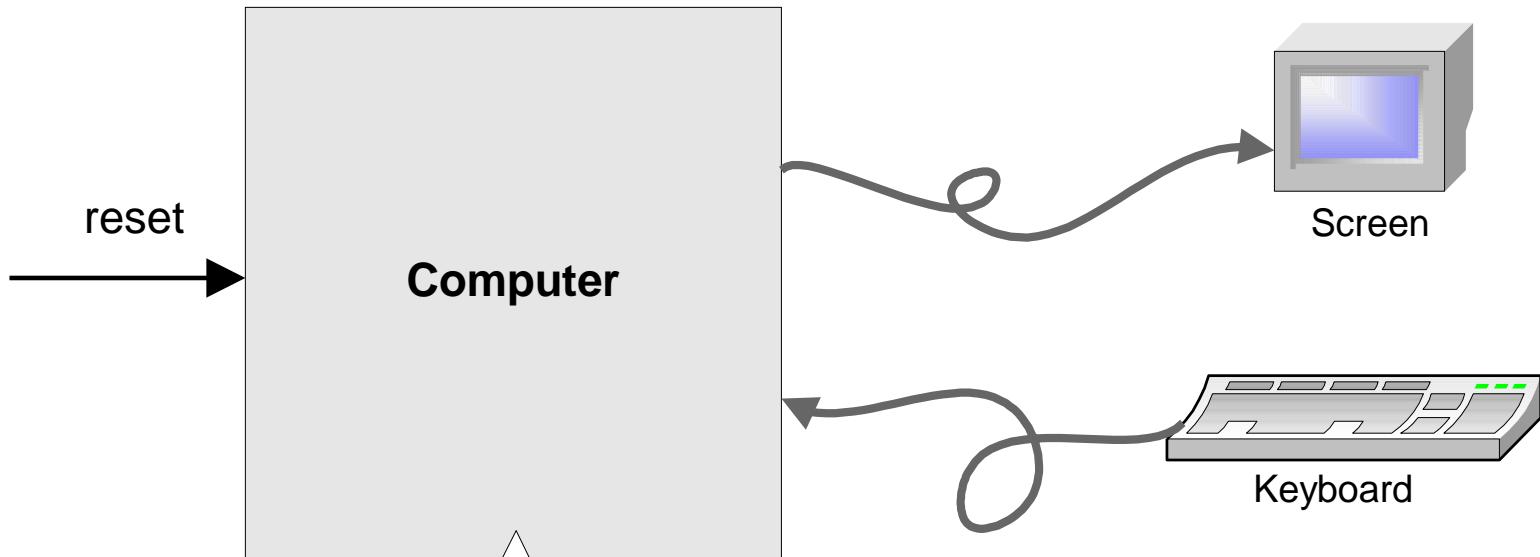
www.nand2tetris.org

Where we are at:



The Hack computer

A 16-bit machine consisting of the following elements:



The program is stored in a ROM.



The A-instruction

symbolic

`@value`

- *value* is a non-negative decimal number $\leq 2^{15}-1$ or
- A symbol referring to such a constant

binary

`0value`

- *value* is a 15-bit binary number

Example

`@21`

`0000 0000 0001 0101`

The C-instruction

symbolic

dest = comp ; jump

binary

111A C₁C₂C₃C₄ C₅C₆ D₁D₂ D₃J₁J₂J₃



The C-instruction

111A C₁C₂C₃C₄ C₅C₆ D₁D₂ D₃J₁J₂J₃

	comp			dest		jump	
(when a=0) comp	c ₁	c ₂	c ₃	c ₄	c ₅	c ₆	(when a=1) comp
0	1	0	1	0	1	0	
1	1	1	1	1	1	1	
-1	1	1	1	0	1	0	
D	0	0	1	1	0	0	
A	1	1	0	0	0	0	M
!D	0	0	1	1	0	1	
!A	1	1	0	0	0	1	!M
-D	0	0	1	1	1	1	
-A	1	1	0	0	1	1	-M
D+1	0	1	1	1	1	1	
A+1	1	1	0	1	1	1	M+1
D-1	0	0	1	1	1	0	
A-1	1	1	0	0	1	0	M-1
D+A	0	0	0	0	1	0	D+M
D-A	0	1	0	0	1	1	D-M
A-D	0	0	0	1	1	1	M-D
D&A	0	0	0	0	0	0	D&M
D A	0	1	0	1	0	1	D M

The C-instruction

111A C₁C₂C₃C₄ C₅C₆ D₁D₂ D₃J₁J₂J₃

comp

dest

jump

A D M

dest

d d d

effect: the value is stored in:

null	0 0 0	the value is not stored
M	0 0 1	RAM[A]
D	0 1 0	D register
DM	0 1 1	D register and RAM[A]
A	1 0 0	A register
AM	1 0 1	A register and RAM[A]
AD	1 1 0	A register and D register
ADM	1 1 1	A register, D register, and RAM[A]

The C-instruction

111A C₁C₂C₃C₄ C₅C₆ D₁D₂ D₃J₁J₂J₃

comp

dest

jump

< = >

jump j j j effect:

null	0	0	0	no jump
JGT	0	0	1	if <i>comp</i> > 0 jump
JEQ	0	1	0	if <i>comp</i> = 0 jump
JGE	0	1	1	if <i>comp</i> ≥ 0 jump
JLT	1	0	0	if <i>comp</i> < 0 jump
JNE	1	0	1	if <i>comp</i> ≠ 0 jump
JLE	1	1	0	if <i>comp</i> ≤ 0 jump
JMP	1	1	1	Unconditional jump

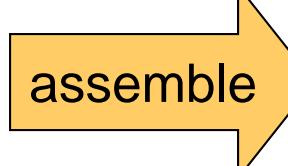
Hack assembly/machine language

Source code (example)

```
// Computes 1+...+RAM[0]
// And stored the sum in RAM[1]
    @i
    M=1      // i = 1
    @sum
    M=0      // sum = 0
(LOOP)
    @i      // if i>RAM[0] goto WRITE
    D=M
    @R0
    D=D-M
    @WRITE
    D;JGT
    @i      // sum += i
    D=M
    @sum
    M=D+M
    @i      // i++
    M=M+1
    @LOOP // goto LOOP
    0;JMP
(WRITE)
    @sum
    D=M
    @R1
    M=D    // RAM[1] = the sum
(END)
    @END
    0;JMP
```

Target code

```
0000000000010000
1110111111001000
0000000000010001
1110101010001000
0000000000010000
1111110000010000
0000000000000000
1111010011010000
0000000000010010
1110001100000001
0000000000010000
1111110000010000
0000000000010001
1111000010001000
0000000000010000
1111110111001000
0000000000001000
1110101010000111
0000000000010001
1111110000010000
0000000000000001
1110001100001000
00000000000010110
1110101010000111
```



assemble

Hack assembler
or CPU emulator

assembly code v.s. machine code

The Hack computer

- A 16-bit stored program platform
- The *instruction memory* and the *data memory* are physically separate
- Screen: 512 rows by 256 columns, black and white
- Keyboard: standard
- Designed to execute programs written in the Hack machine language
- Can be easily built from the chip-set that we built so far in the course

Main parts of the Hack computer:

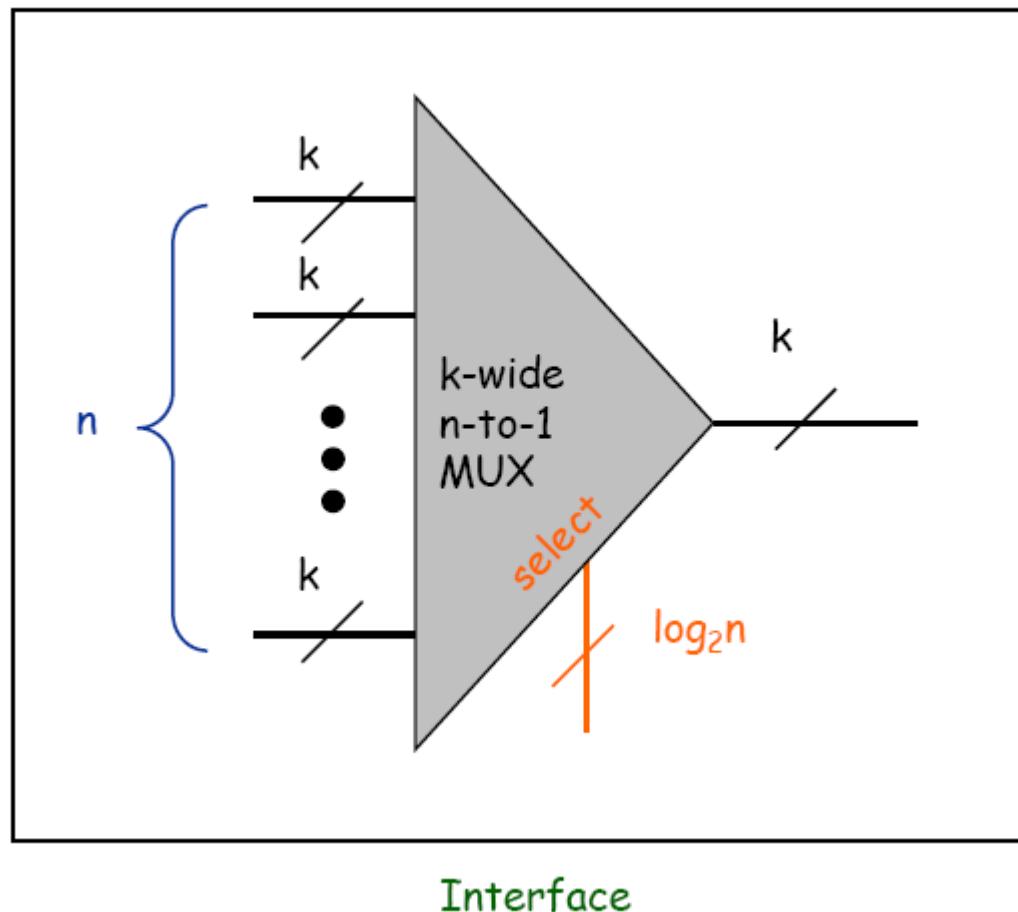
- Instruction memory (ROM)
- Memory (RAM):
 - Data memory
 - Screen (memory map)
 - Keyboard (memory map)
- CPU
- Computer (the logic that holds everything together).



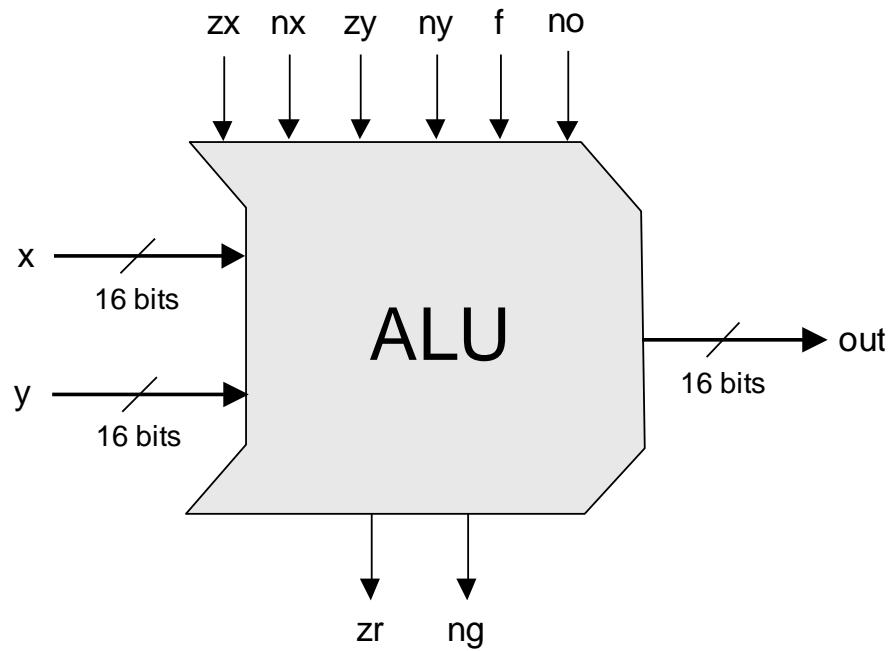
Multiplexer

Goal: select from one of n k -bit buses

- Implemented by layering k n -to-1 multiplexer



Hack ALU



out(x, y, control bits) =

x+y, x-y, y-x,
0, 1, -1,
x, y, -x, -y,
x!, y!,
x+1, y+1, x-1,
y-1,
x&y, x|y

Hack ALU

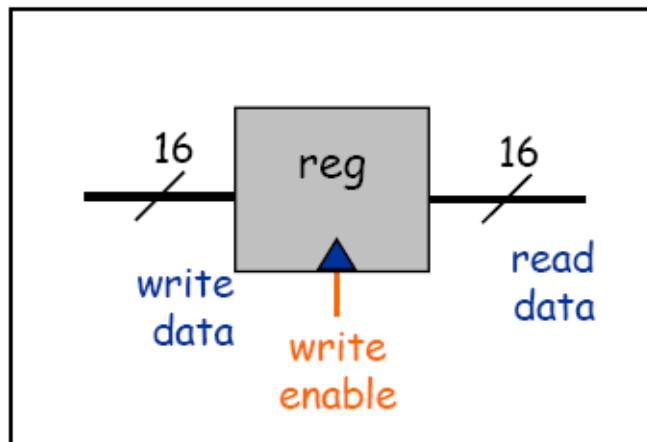
These bits instruct how to preset the x input		These bits instruct how to preset the y input		This bit selects between + / And	This bit inst. how to postset out	Resulting ALU output
zx	nx	zy	ny	f	no	out=
if zx then x=0	if nx then x=!x	if zy then y=0	if ny then y=!y	if f then out=x+y else out=x&y	if no then out=!out	f(x,y)=
1	0	1	0	1	0	0
1	1	1	1	1	1	1
1	1	1	0	1	0	-1
0	0	1	1	0	0	x
1	1	0	0	0	0	y
0	0	1	1	0	1	!x
1	1	0	0	0	1	!y
0	0	1	1	1	1	-x
1	1	0	0	1	1	-y
0	1	1	1	1	1	x+1
1	1	0	1	1	1	y+1
0	0	1	1	1	0	x-1
1	1	0	0	1	0	y-1
0	0	0	0	1	0	x+y
0	1	0	0	1	1	x-y
0	0	0	1	1	1	y-x
0	0	0	0	0	0	x&y
0	1	0	1	0	1	x y

Registers

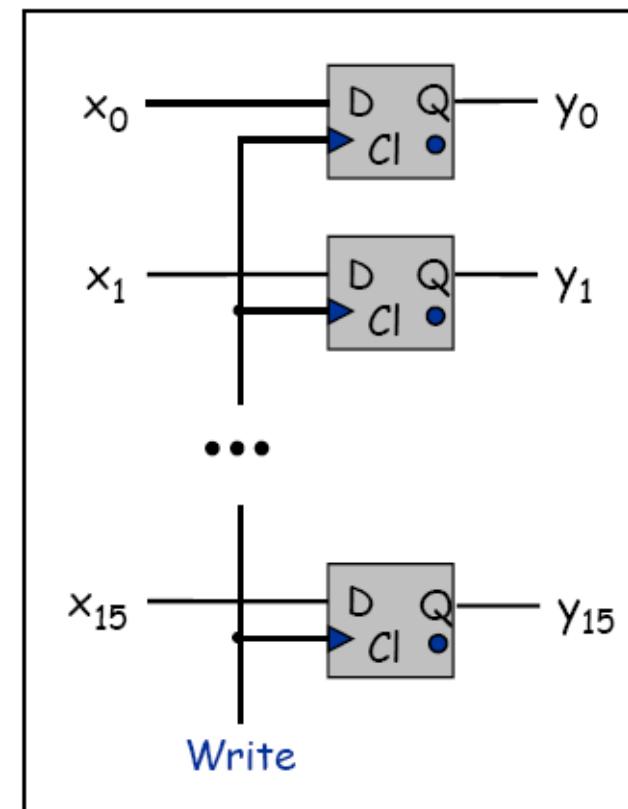
k-bit register.

- Stores k bits.
- Register contents always available on output.
- If write enable is asserted, k input bits get copied into register.

Ex: Program Counter, 16 TOY registers,
256 TOY memory locations.

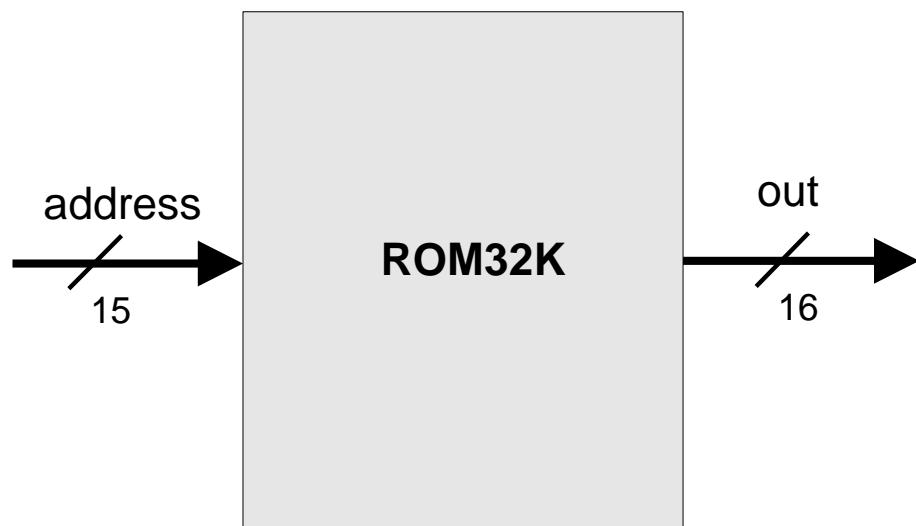


16-bit Register Interface



16-bit Register Implementation

ROM (Instruction memory)



Function:

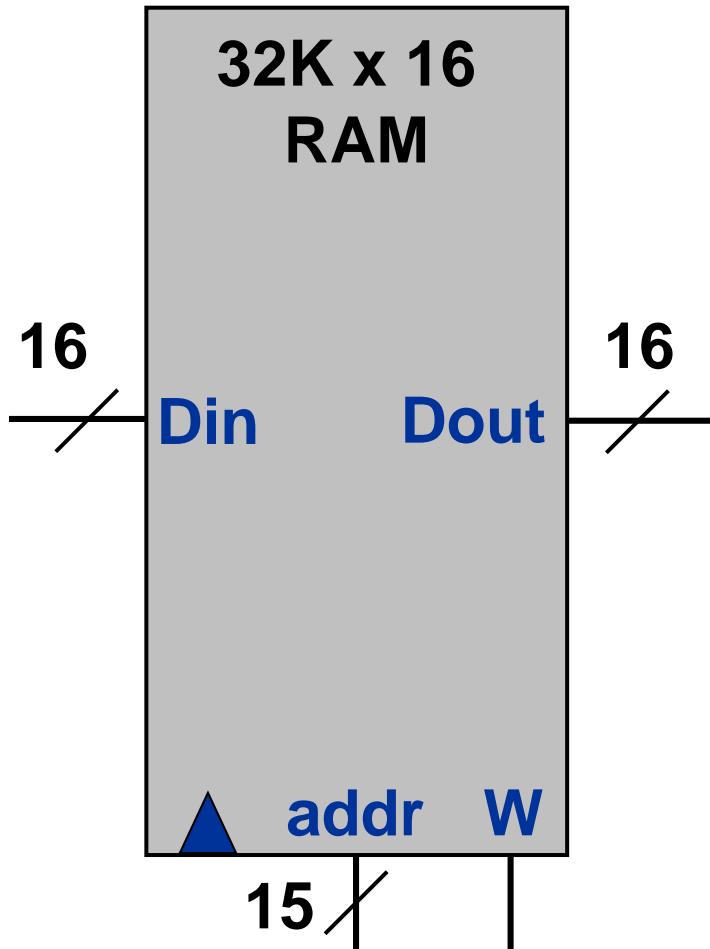
- The ROM is pre-loaded with a program written in the Hack machine language
- The ROM chip always emits a 16-bit number:

```
out = ROM32K[address]
```

- This number is interpreted as the *current instruction*.

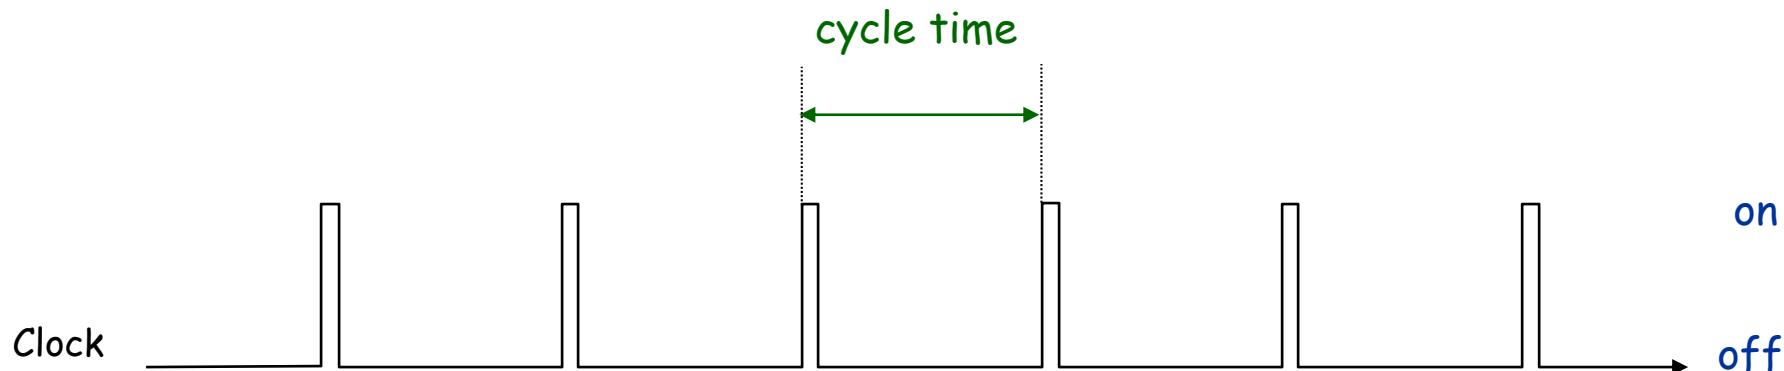
RAM (data memory)

- We will discuss the details for Hack's data memory later.



■ Clock.

- Fundamental abstraction: regular on-off pulse.
 - on: fetch phase
 - off: execute phase
- External analog device.
- Synchronizes operations of different circuit elements.
- Requirement: clock cycle longer than max switching time.



Design a processor

■ How to build a processor (Hack, this time)

→ • Develop instruction set architecture (ISA)

■ 16-bit words, two types of machine instructions

• Determine major components

■ ALU, registers, program counter, memory

• Determine datapath requirements

■ Flow of bits

• Analyze how to implement each instruction

■ Determine settings of control signals

Hack programming reference card

Hack commands:

A-command: @**value** // A<-value; M=RAM[A]

C-command: **dest = comp ; jump** // **dest =** and **;jump**
// are optional

Where:

comp =

0 , 1 , -1 , D , A , !D , !A , -D , -A , D+1 , A+1 , D-1 , A-1 , D+A , D-A , A-D , D&A , D|A ,
M , !M , -M , M+1 , M-1 , D+M , D-M , M-D , D&M , D|M

dest = M, D, A, MD, AM, AD, AMD, or null

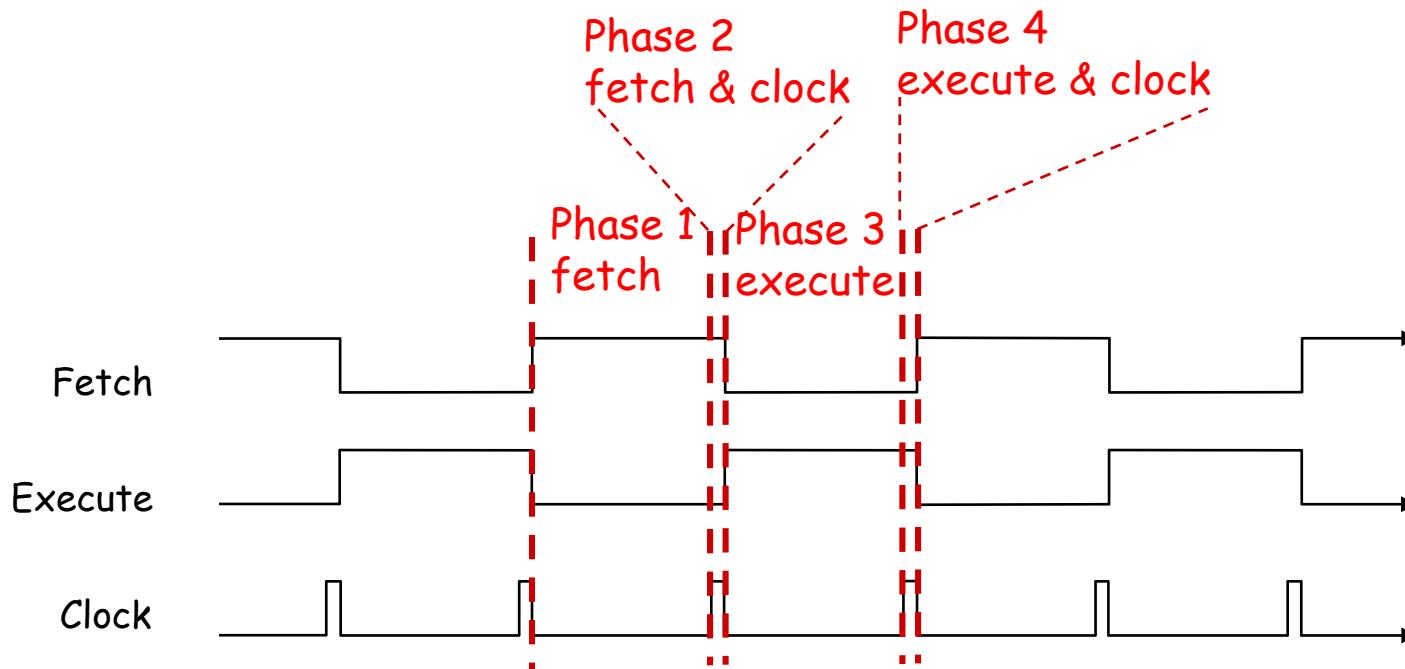
jump = JGT , JEQ , JGE , JLT , JNE , JLE , JMP, or null

In the command **dest = comp; jump**, the jump materializes ($PC < -A$) if (**comp** **jump** 0) is true. For example, in $D=D+1, JLT$, we jump if $D+1 < 0$.

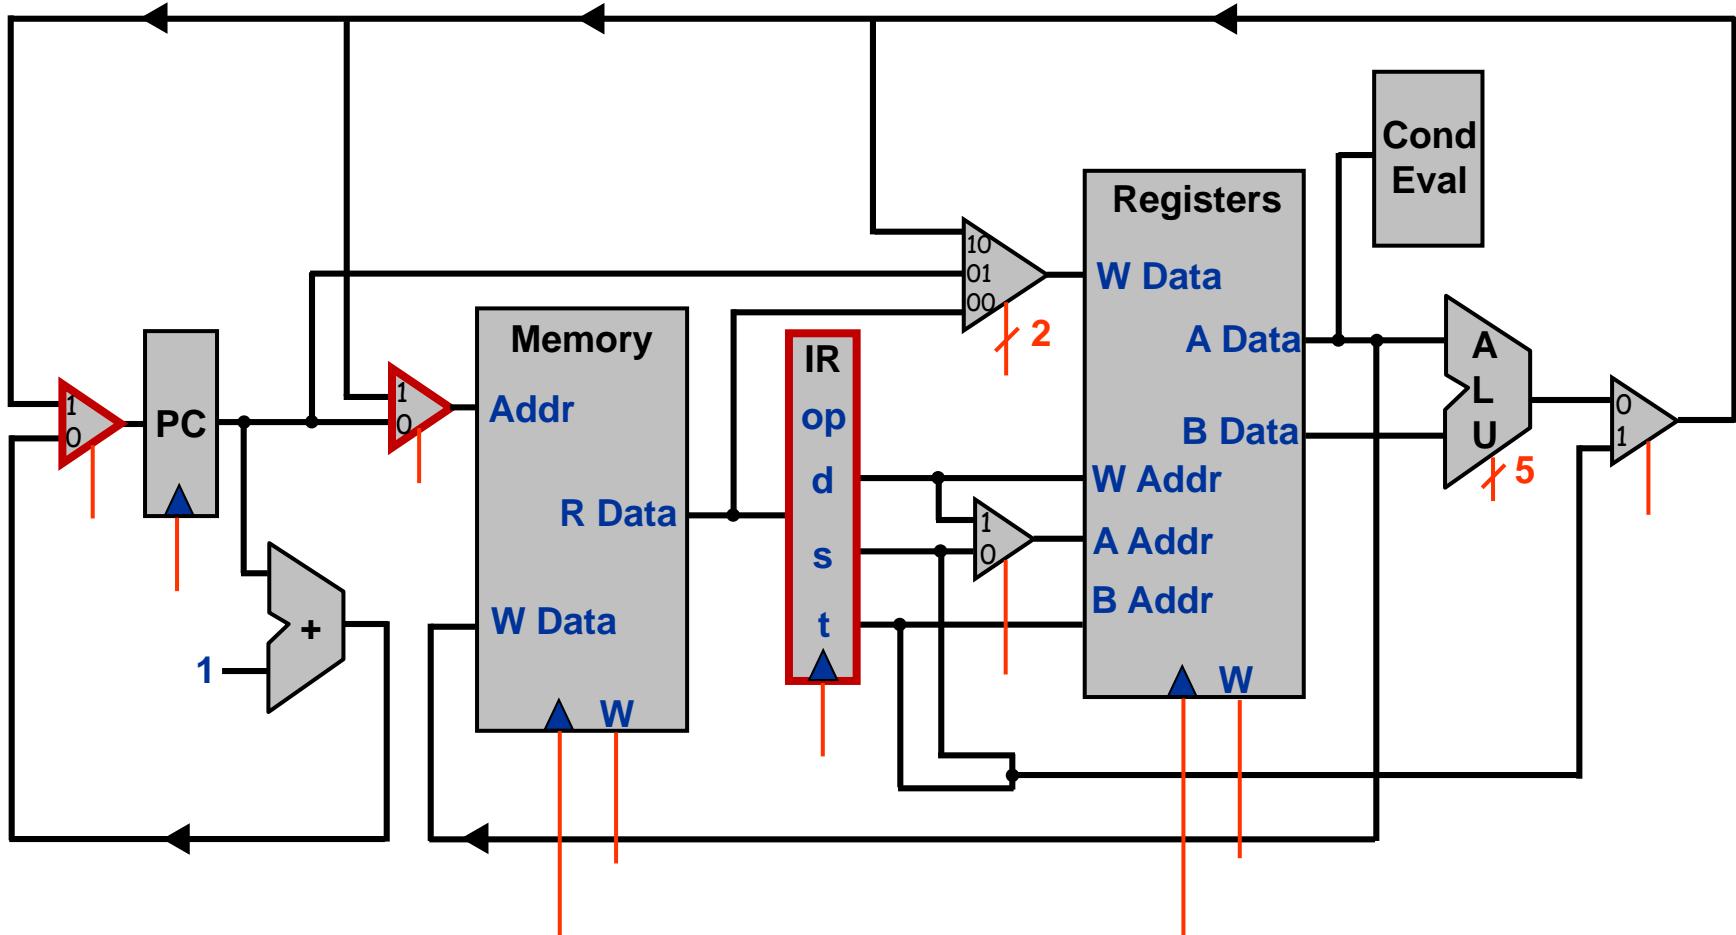
Fetch and execute

- In Toy, we have two phases: fetch and execution .
- We use two cycles since fetch and execute phases each access memory and alter program counter.

- fetch [set memory address from pc]
- fetch and clock [write instruction to IR]
- execute [set ALU inputs from registers]
- execute and clock [write result of ALU to registers]



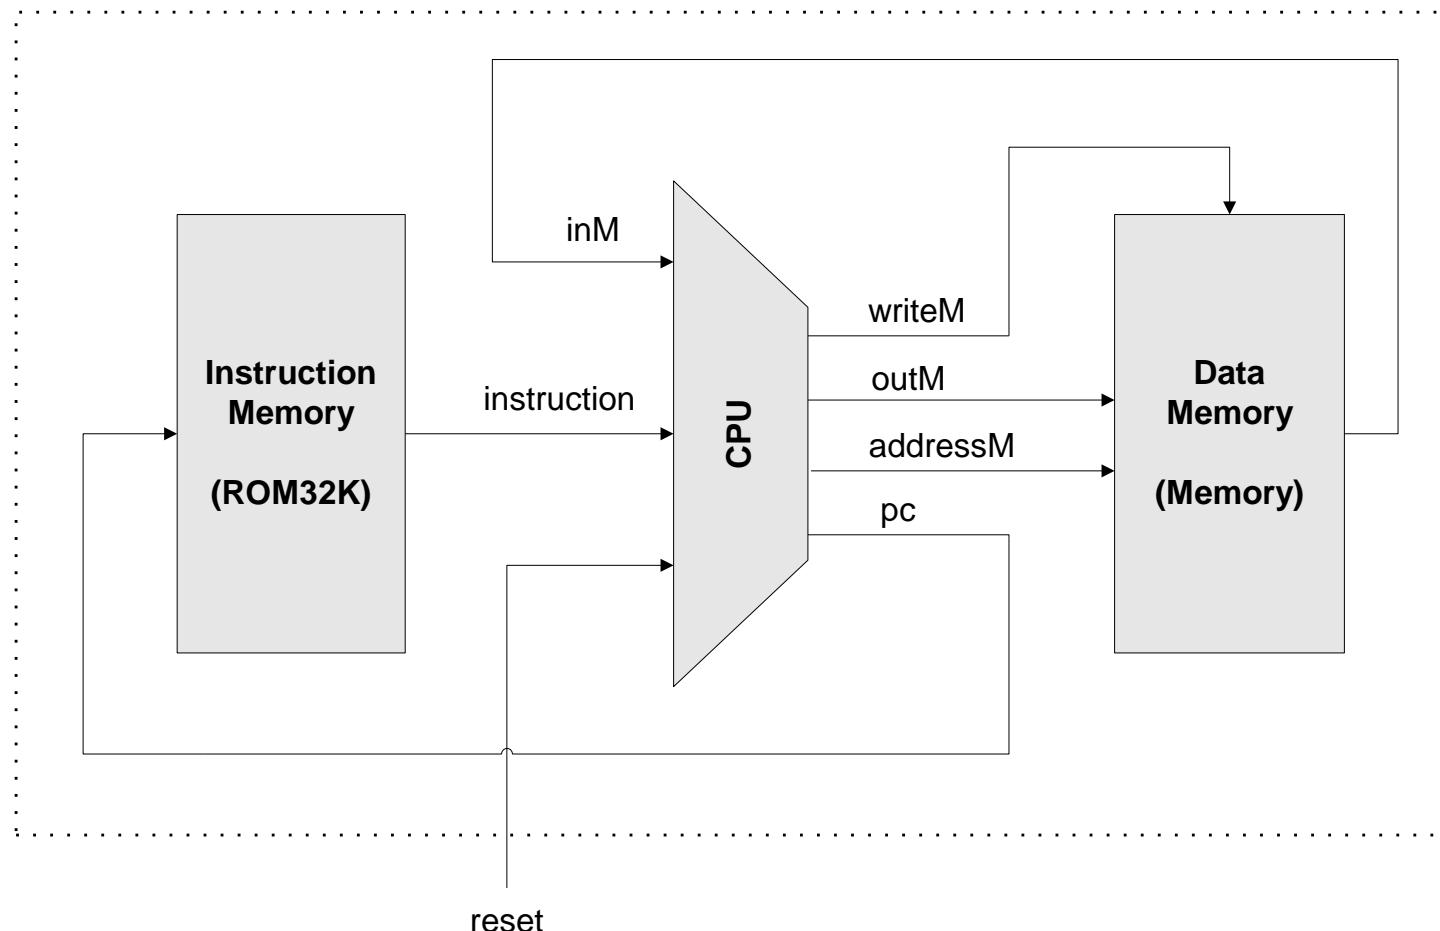
Toy architecture



- Both fetch and execute would access memory. To avoid conflict, we add a MUX. Similar for PC.
- In addition, we need a register IR to store the instruction.

Fetch and execute

- In Hack, we avoid two cycles and IR by using two separate memory chips, one for data and the other for instruction.



Design a processor

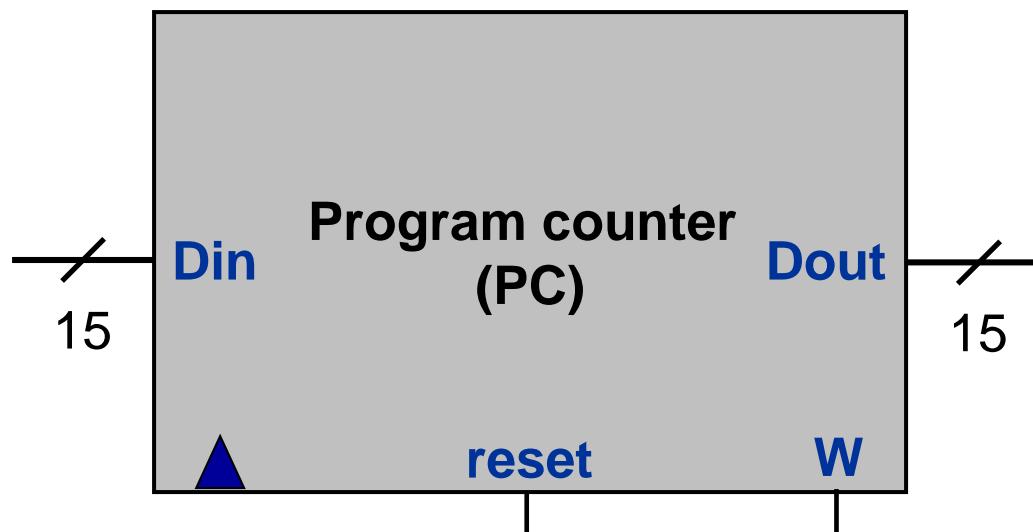
■ How to build a processor (Hack, this time)

- Develop instruction set architecture (ISA)
 - 16-bit words, two types of machine instructions

- 
- Determine major components
 - ALU, registers, program counter, memory
 - Determine datapath requirements
 - Flow of bits
 - Analyze how to implement each instruction
 - Determine settings of control signals

Program counter

- Program counter emits the address of the next instruction.
 - To start/restart the program execution: $PC=0$
 - No jump: $PC++$
 - Unconditional jump: $PC=A$
 - Conditional jump: if (cond.) $PC=A$ else $PC++$



Note that the design is slightly different from your project #3.

Program counter

if (reset) PC=0

else if (W) PC=Din

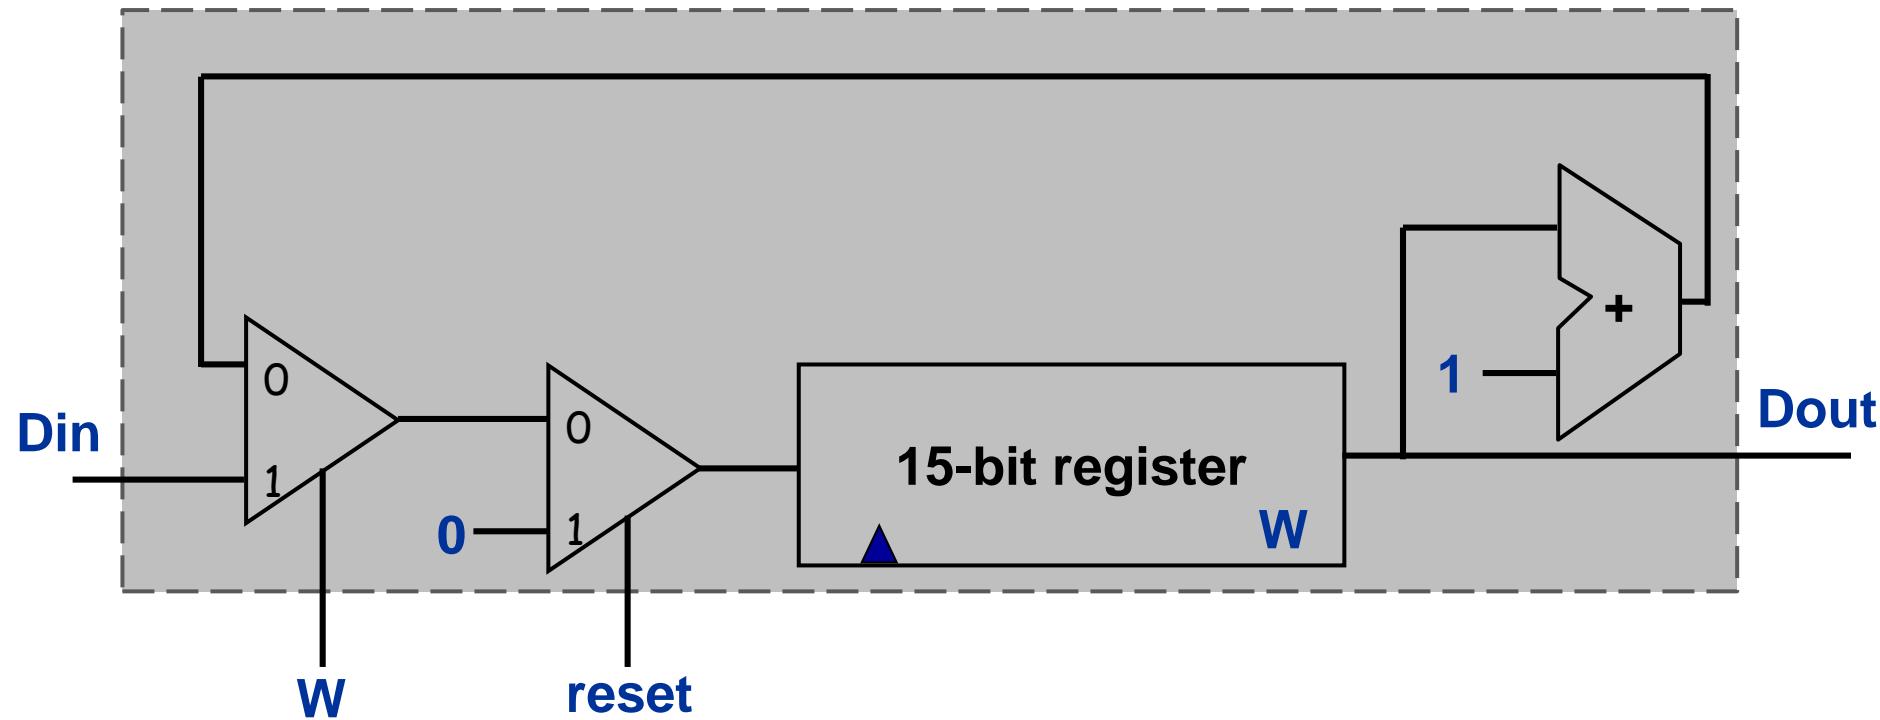
else PC++

Program counter

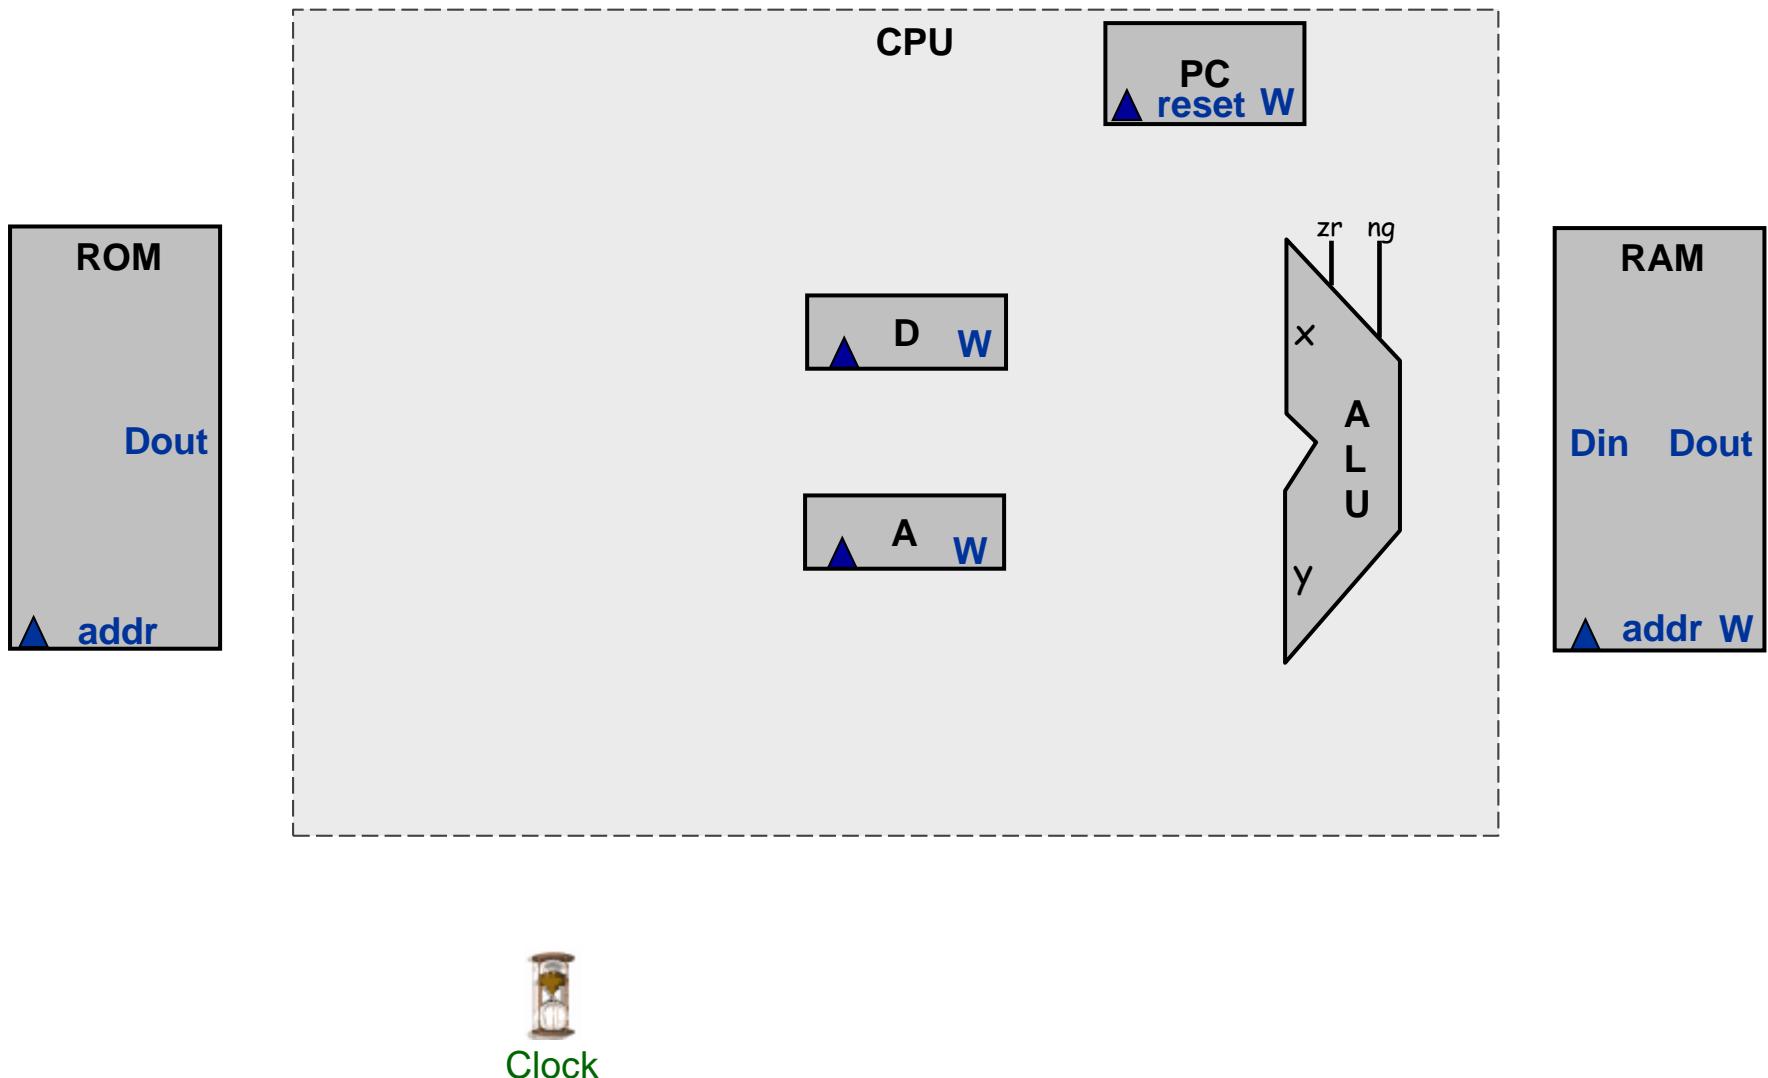
if (reset) PC=0

else if (W) PC=Din

else PC++



Hack architecture (component)

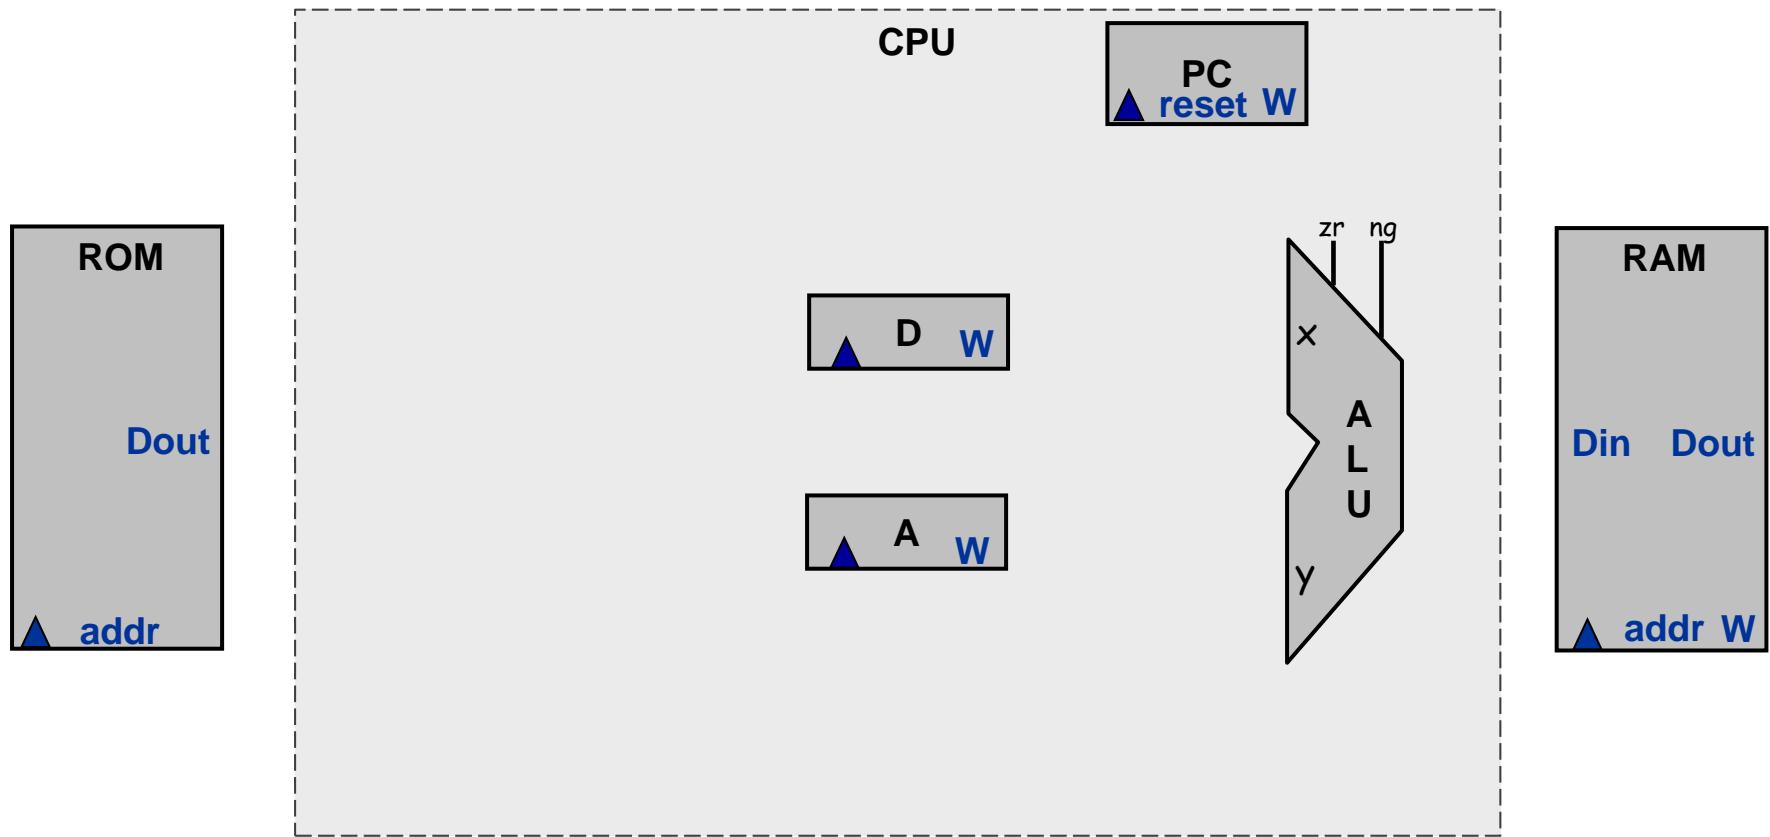


Design a processor

■ How to build a processor (Hack, this time)

- Develop instruction set architecture (ISA)
 - 16-bit words, two types of machine instructions
- Determine major components
 - ALU, registers, program counter, memory
- ● Determine datapath requirements
 - Flow of bits
- Analyze how to implement each instruction
 - Determine settings of control signals

Hack architecture (data path)

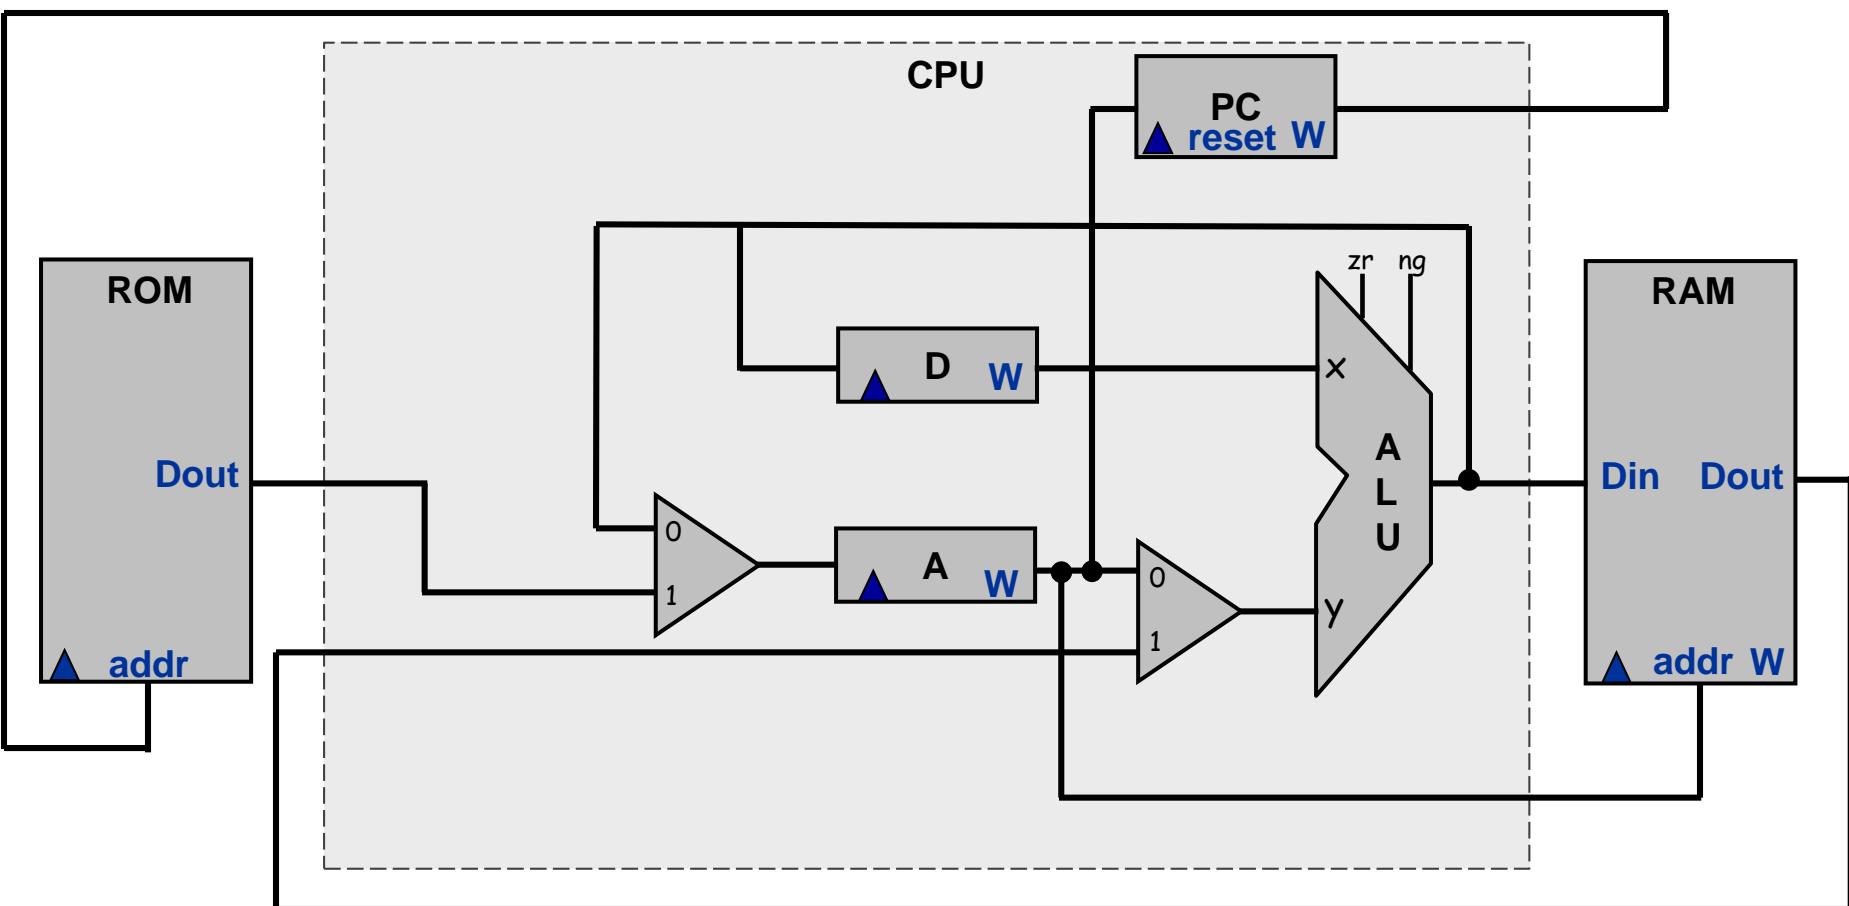


Fetch: instruction=ROM[PC]

@**value** // A<-value; M=RAM[A]

[ADM] = x op y; jump // x=D; y=A or M; if jump then PC<-A

Hack architecture (data path)



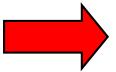
Fetch: instruction=ROM[PC]

@**value** // A<-value; M=RAM[A]

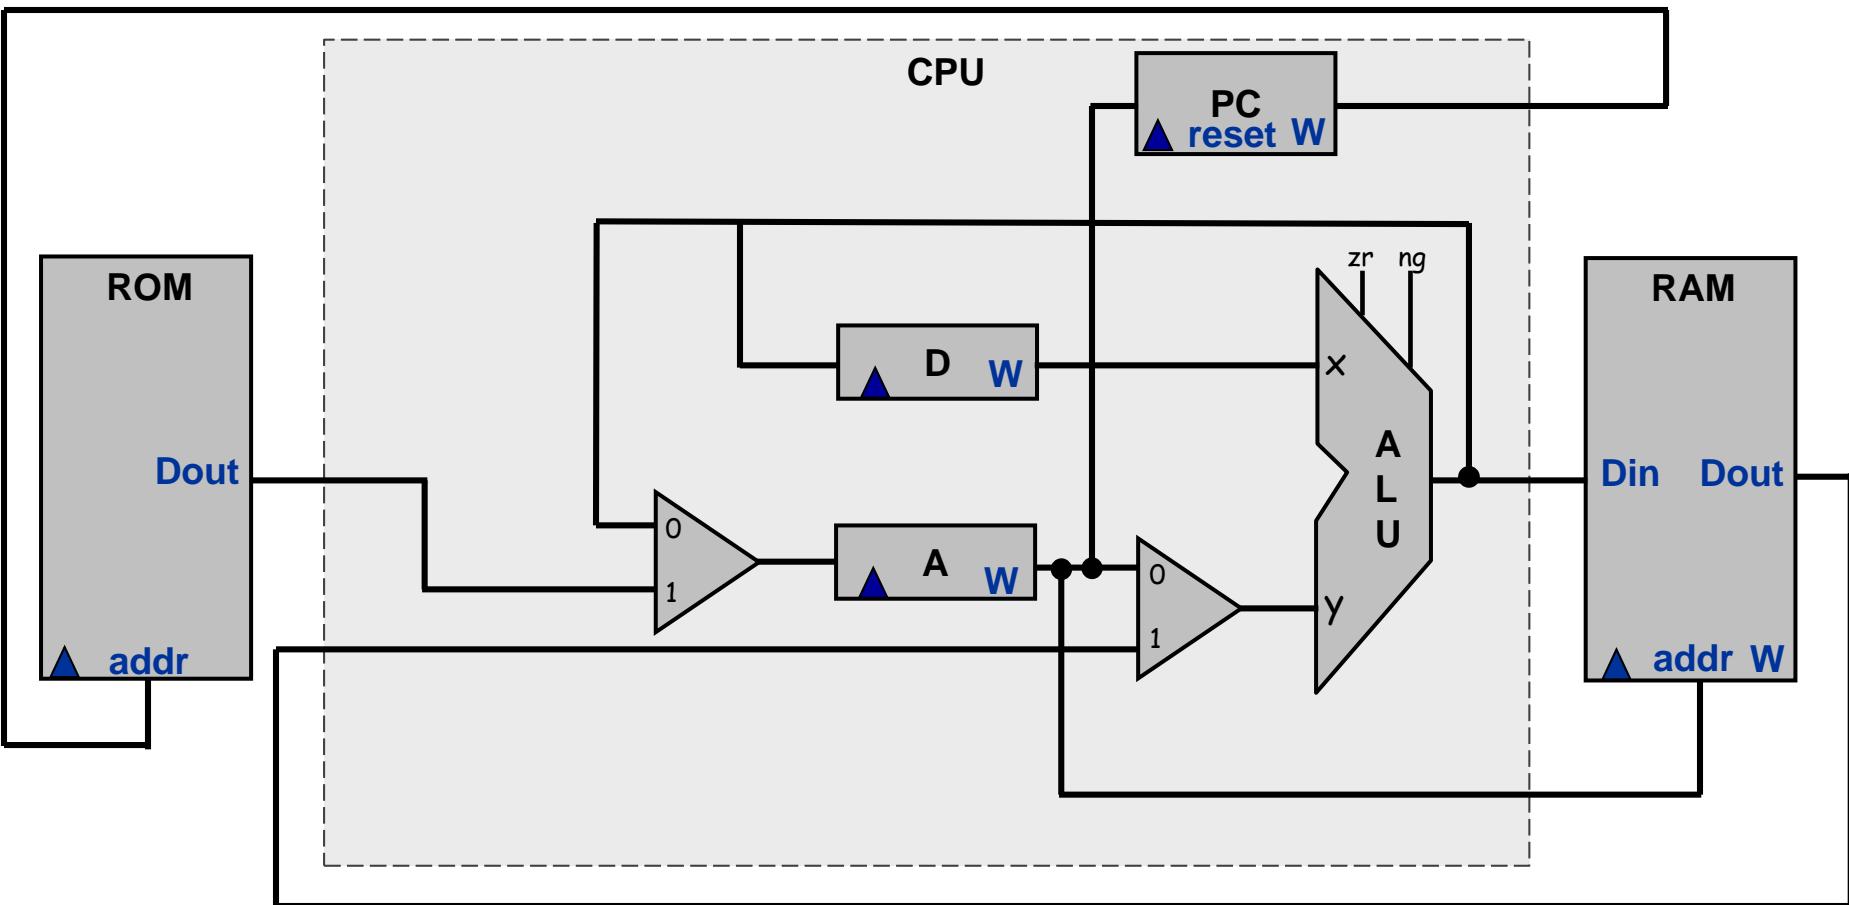
[ADM] = x op y; jump // x=D; y=A or M; if jump then PC<-A

Design a processor

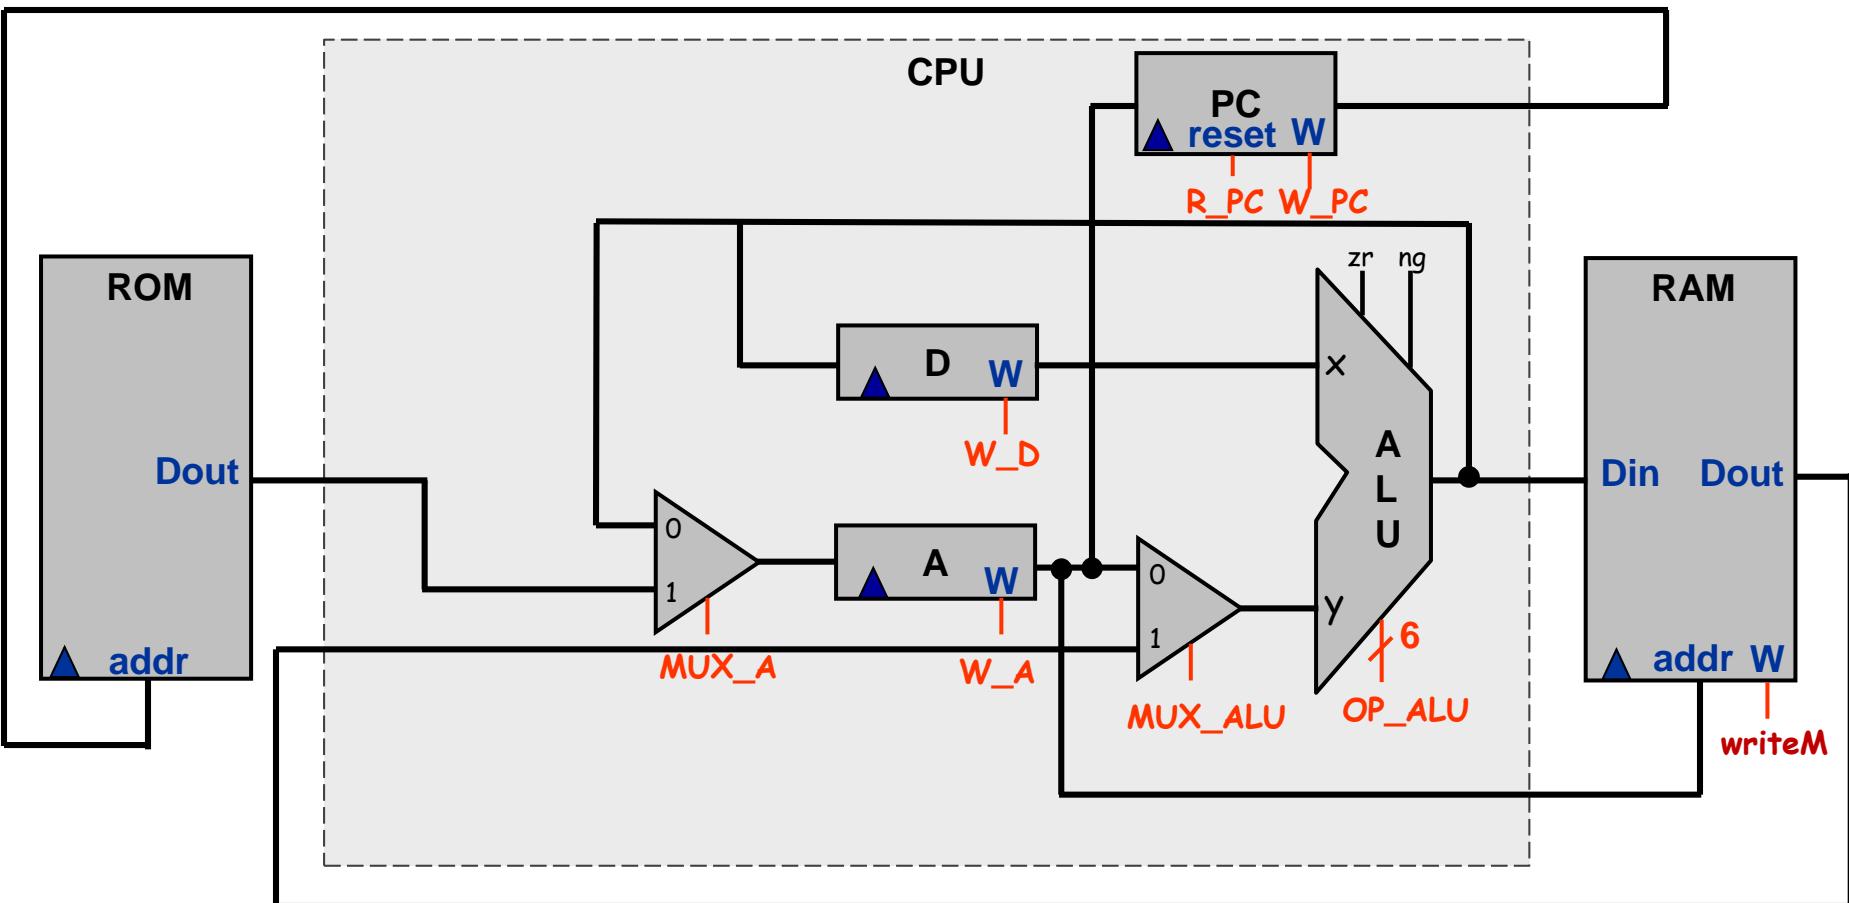
■ How to build a processor (Hack, this time)

- Develop instruction set architecture (ISA)
 - 16-bit words, two types of machine instructions
 - Determine major components
 - ALU, registers, program counter, memory
 - Determine datapath requirements
 - Flow of bits
-
- 
- Analyze how to implement each instruction
 - Determine settings of control signals

Hack architecture (data path)

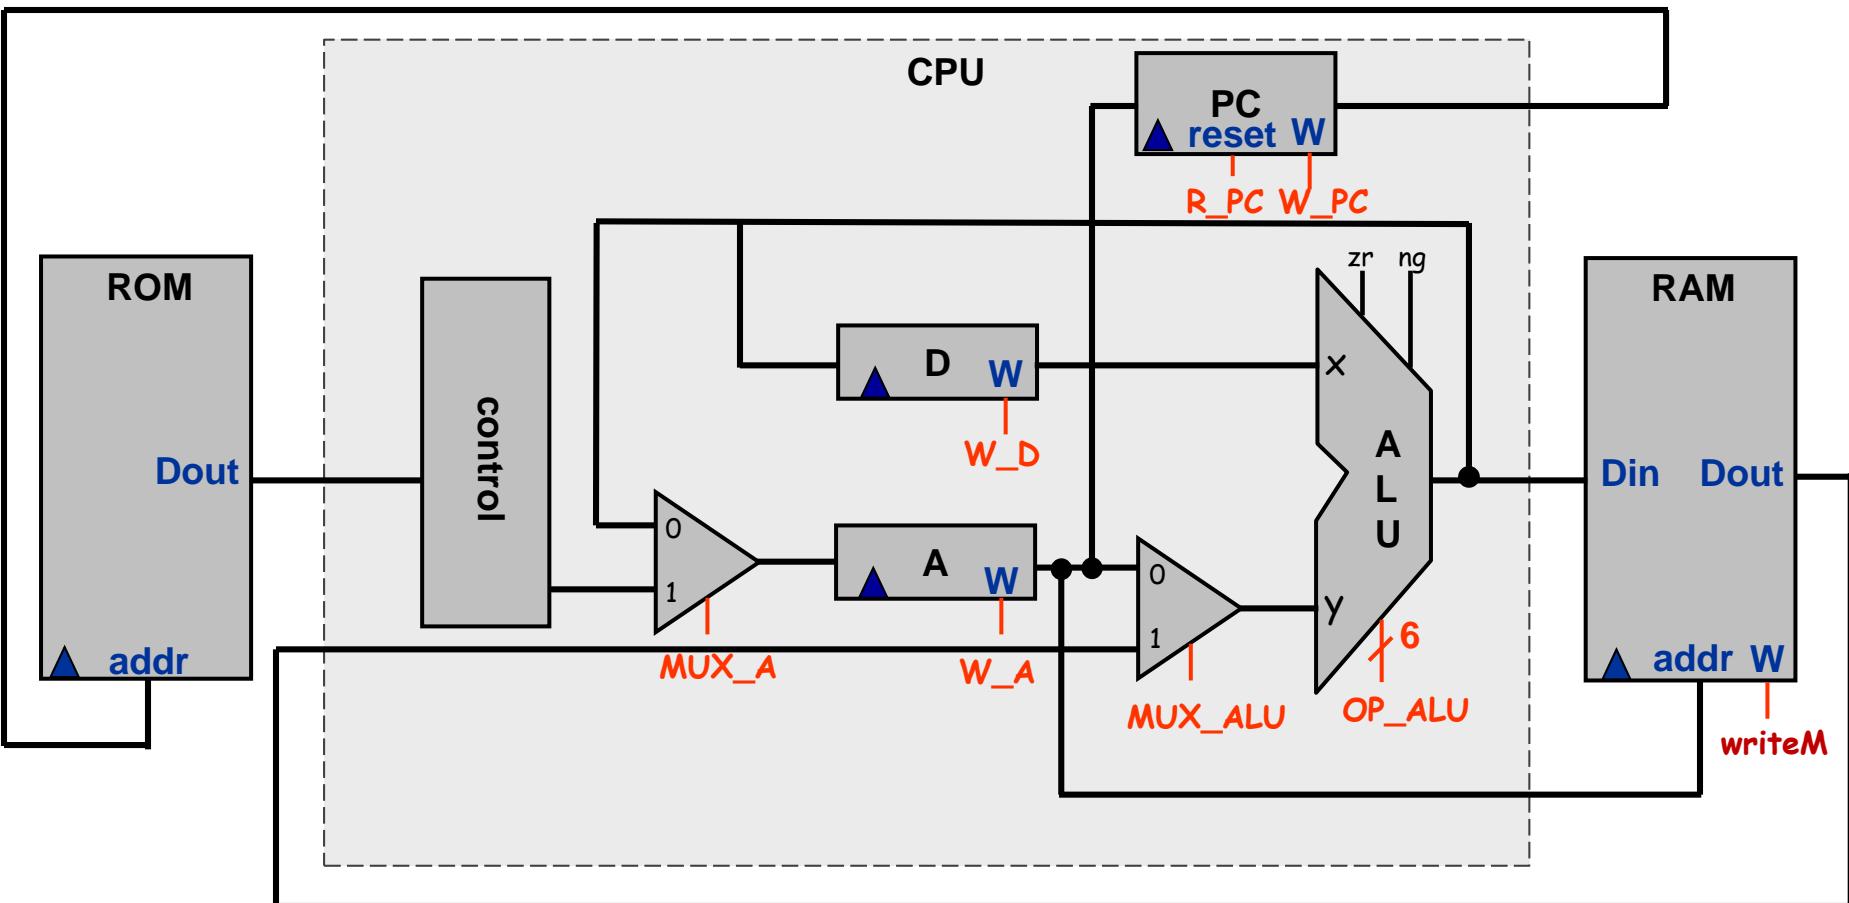


Hack architecture (control)

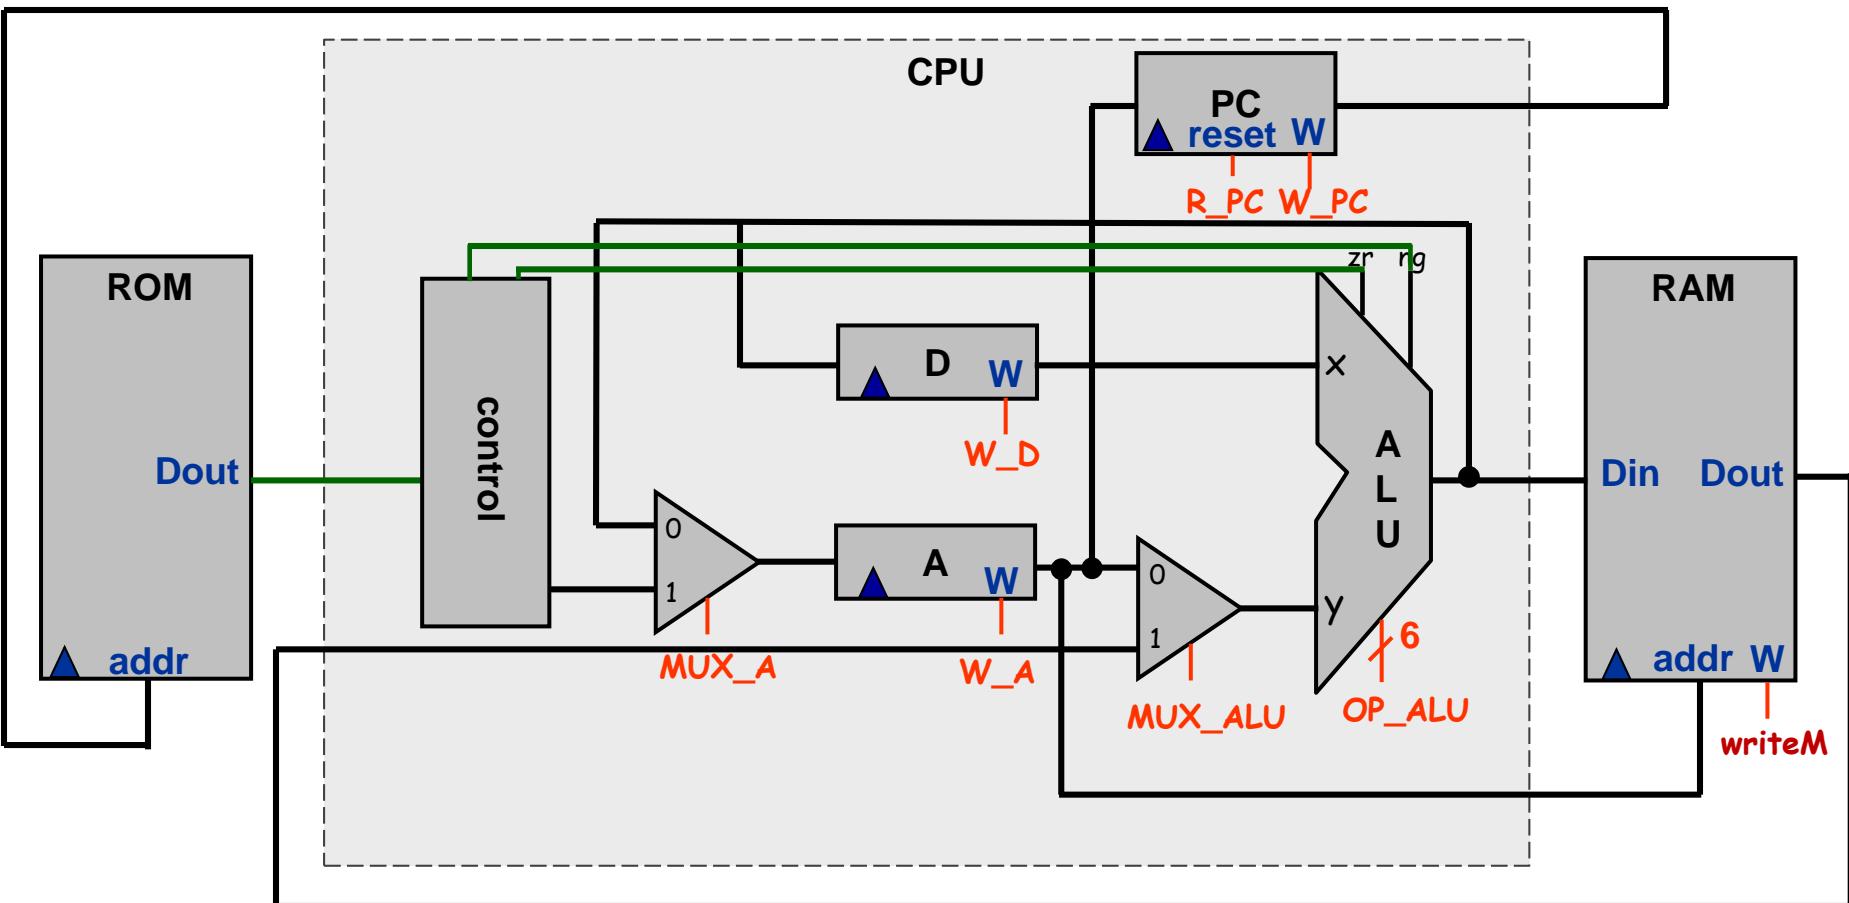


A total of 13 control signals

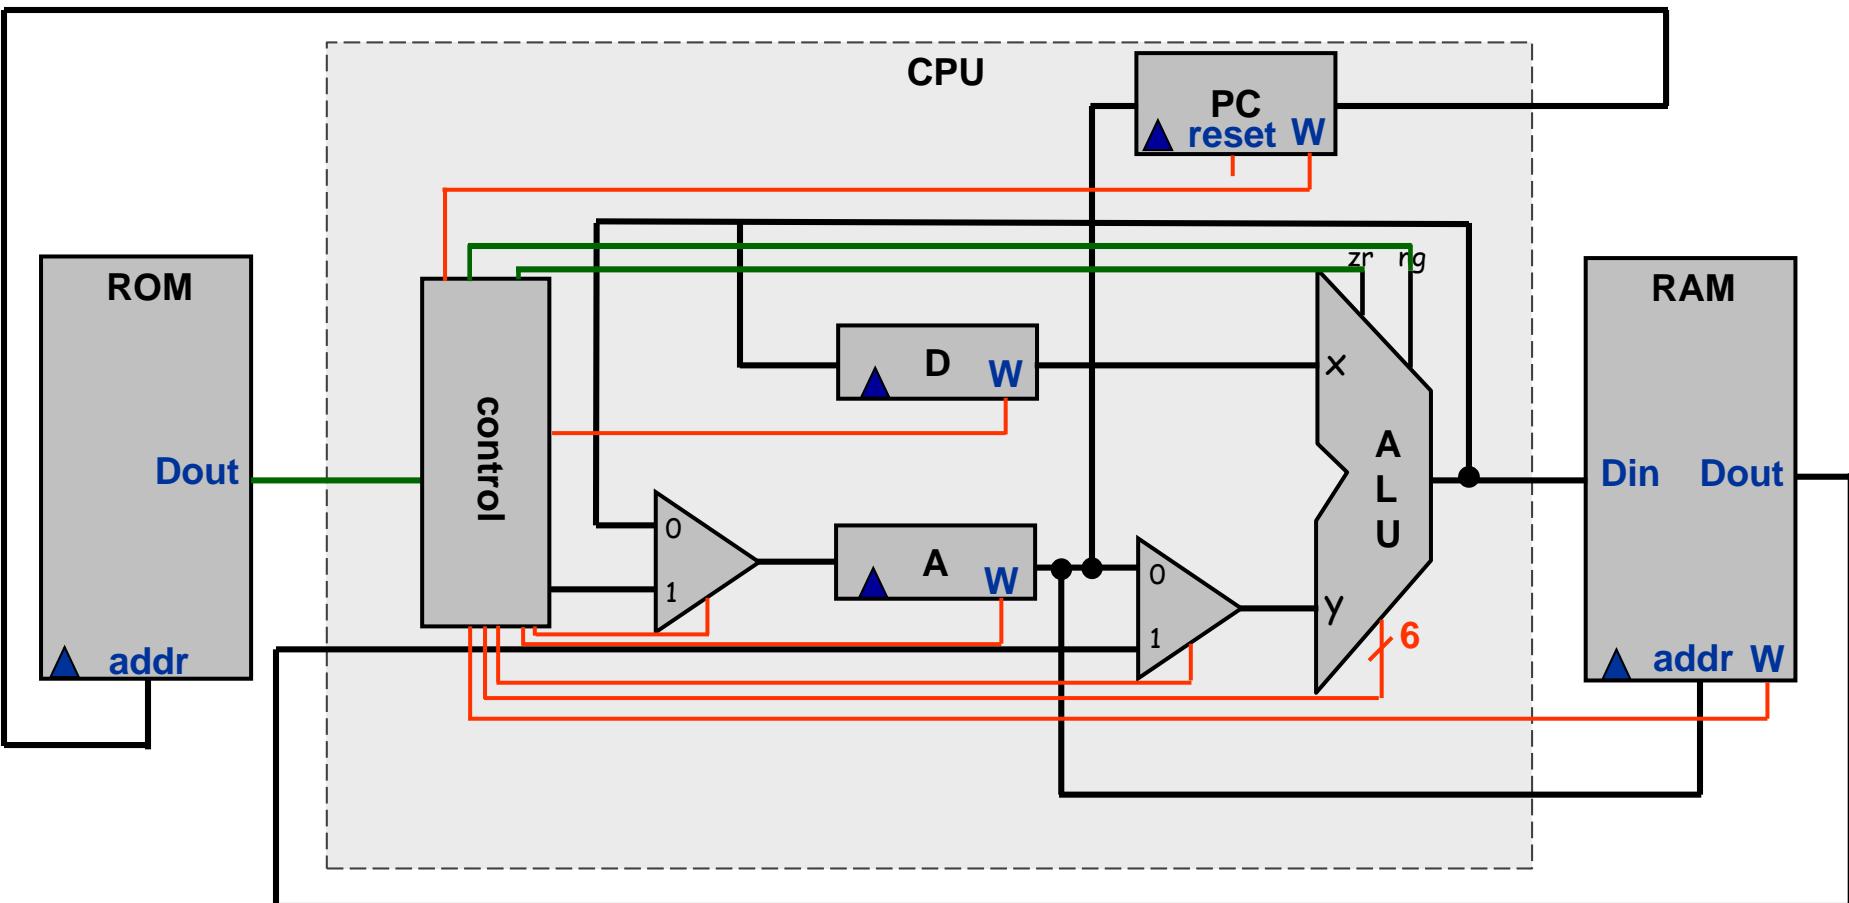
Hack architecture (control)



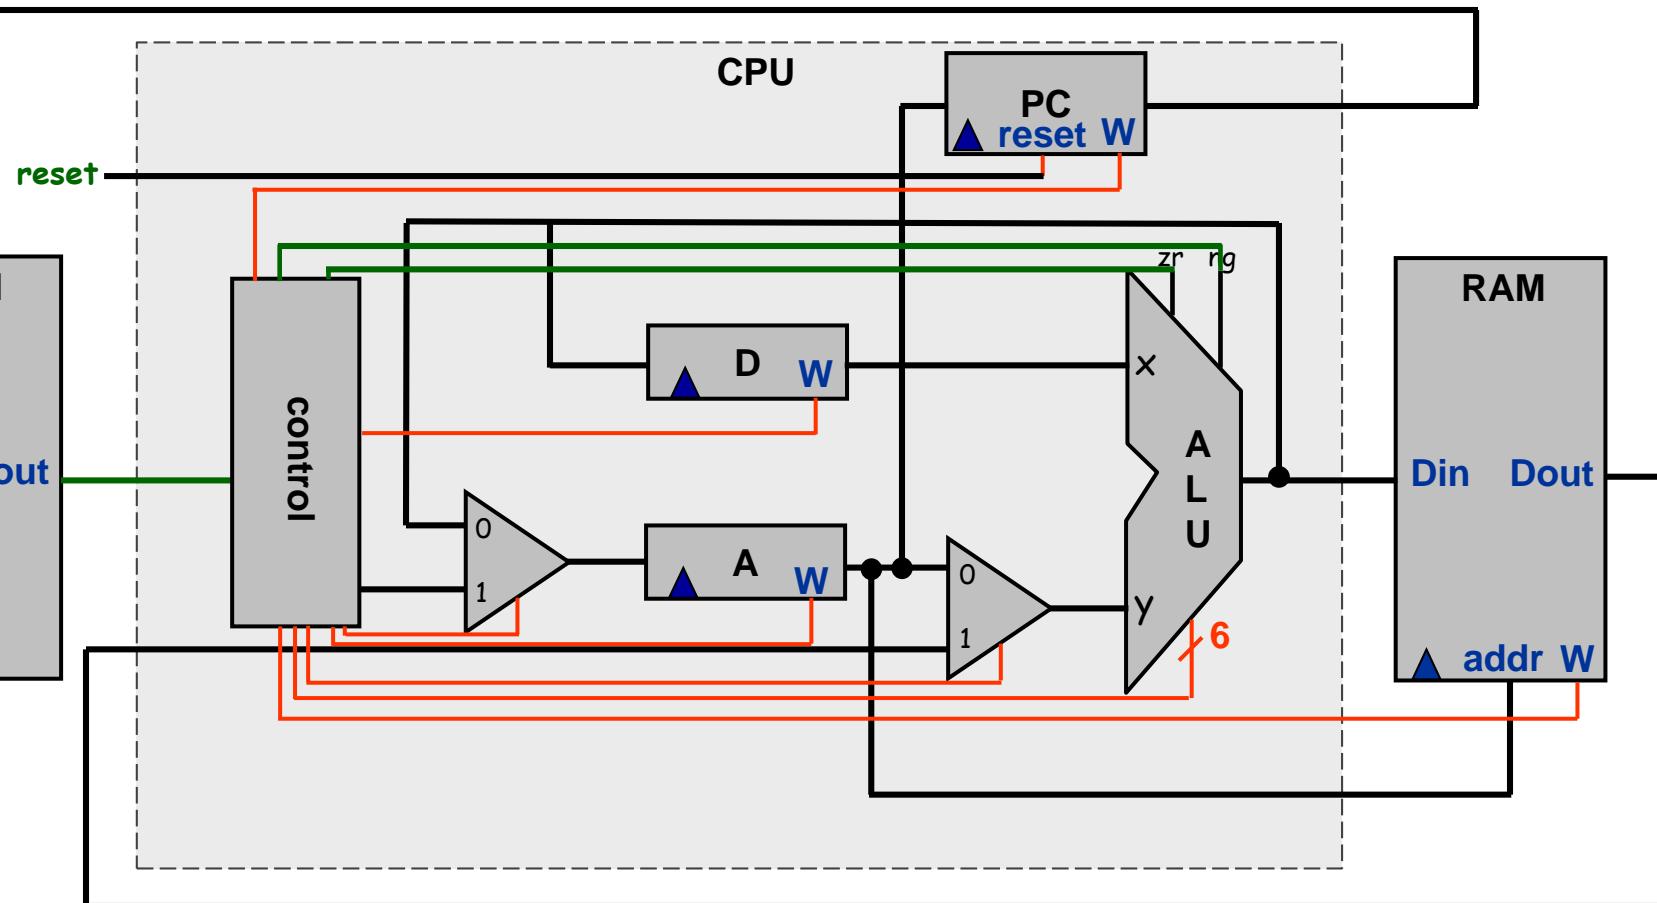
Hack architecture (control)



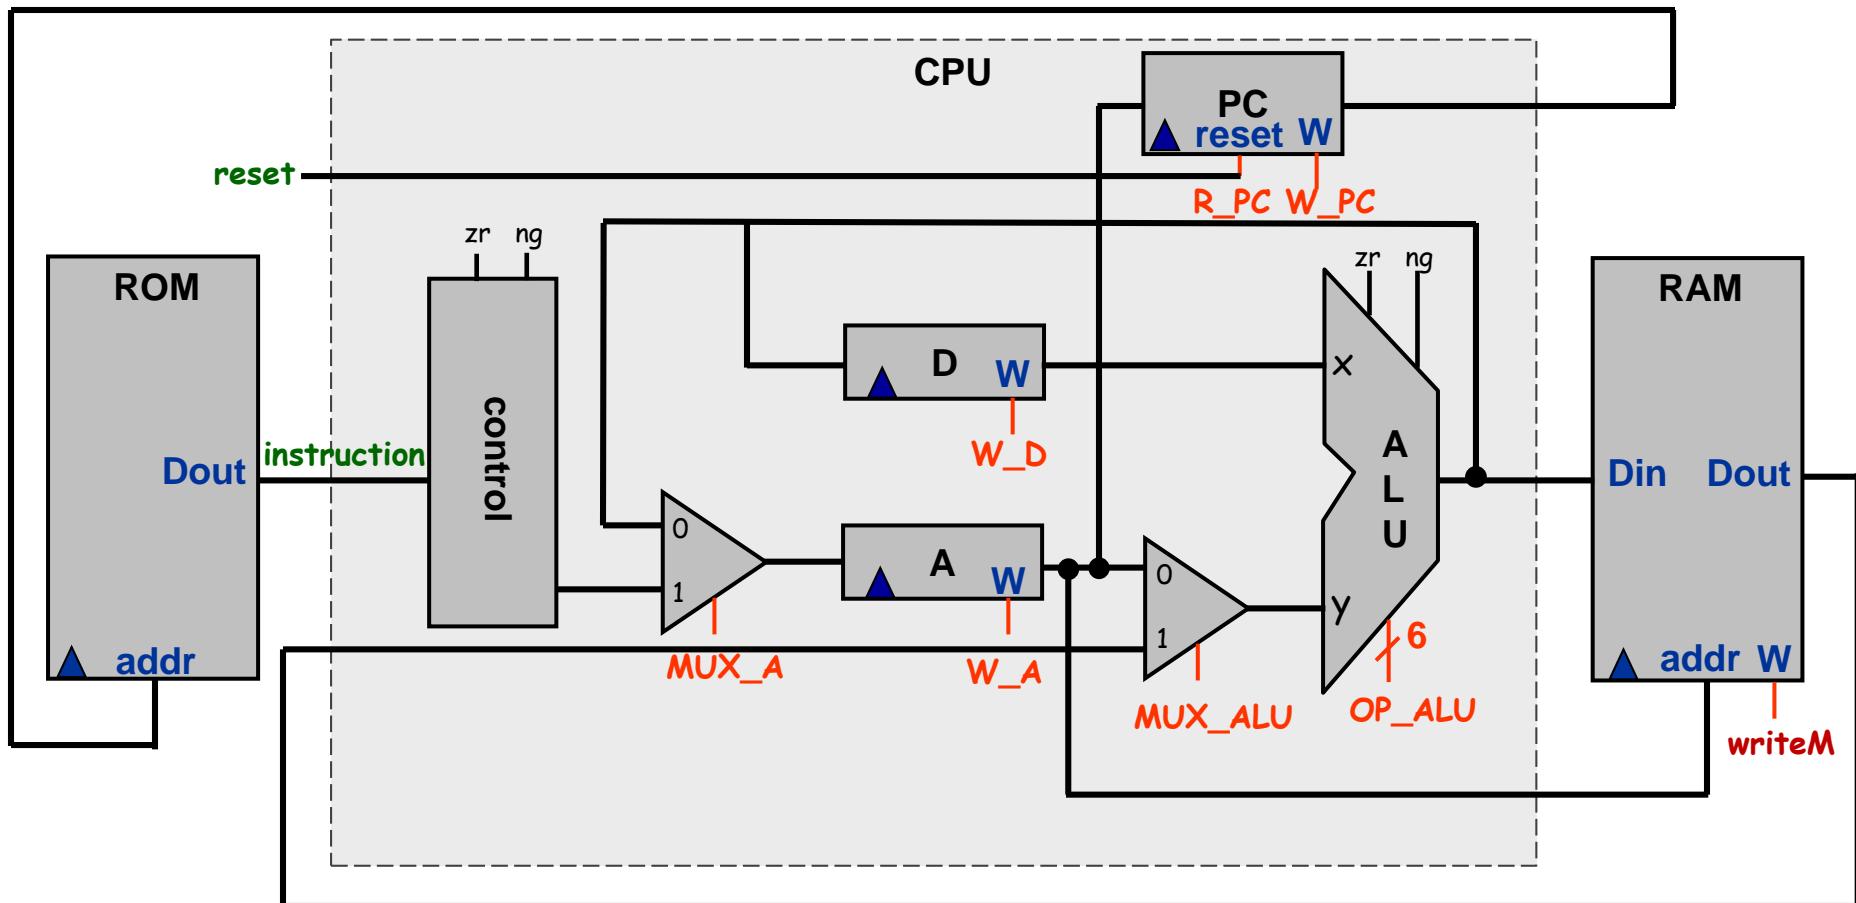
Hack architecture (control)



Hack architecture (control)



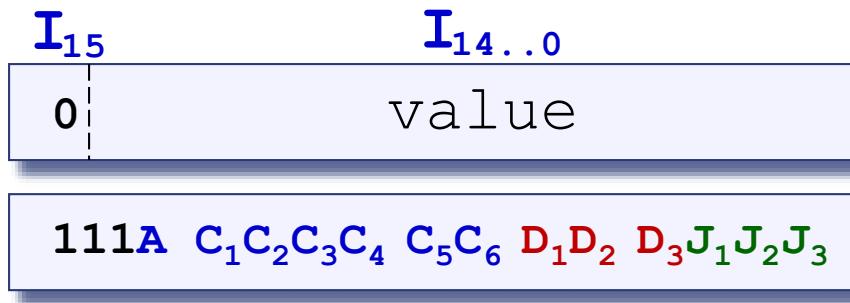
Hack architecture (control)



Hack architecture (control)

■ Inputs: instruction, zr, ng

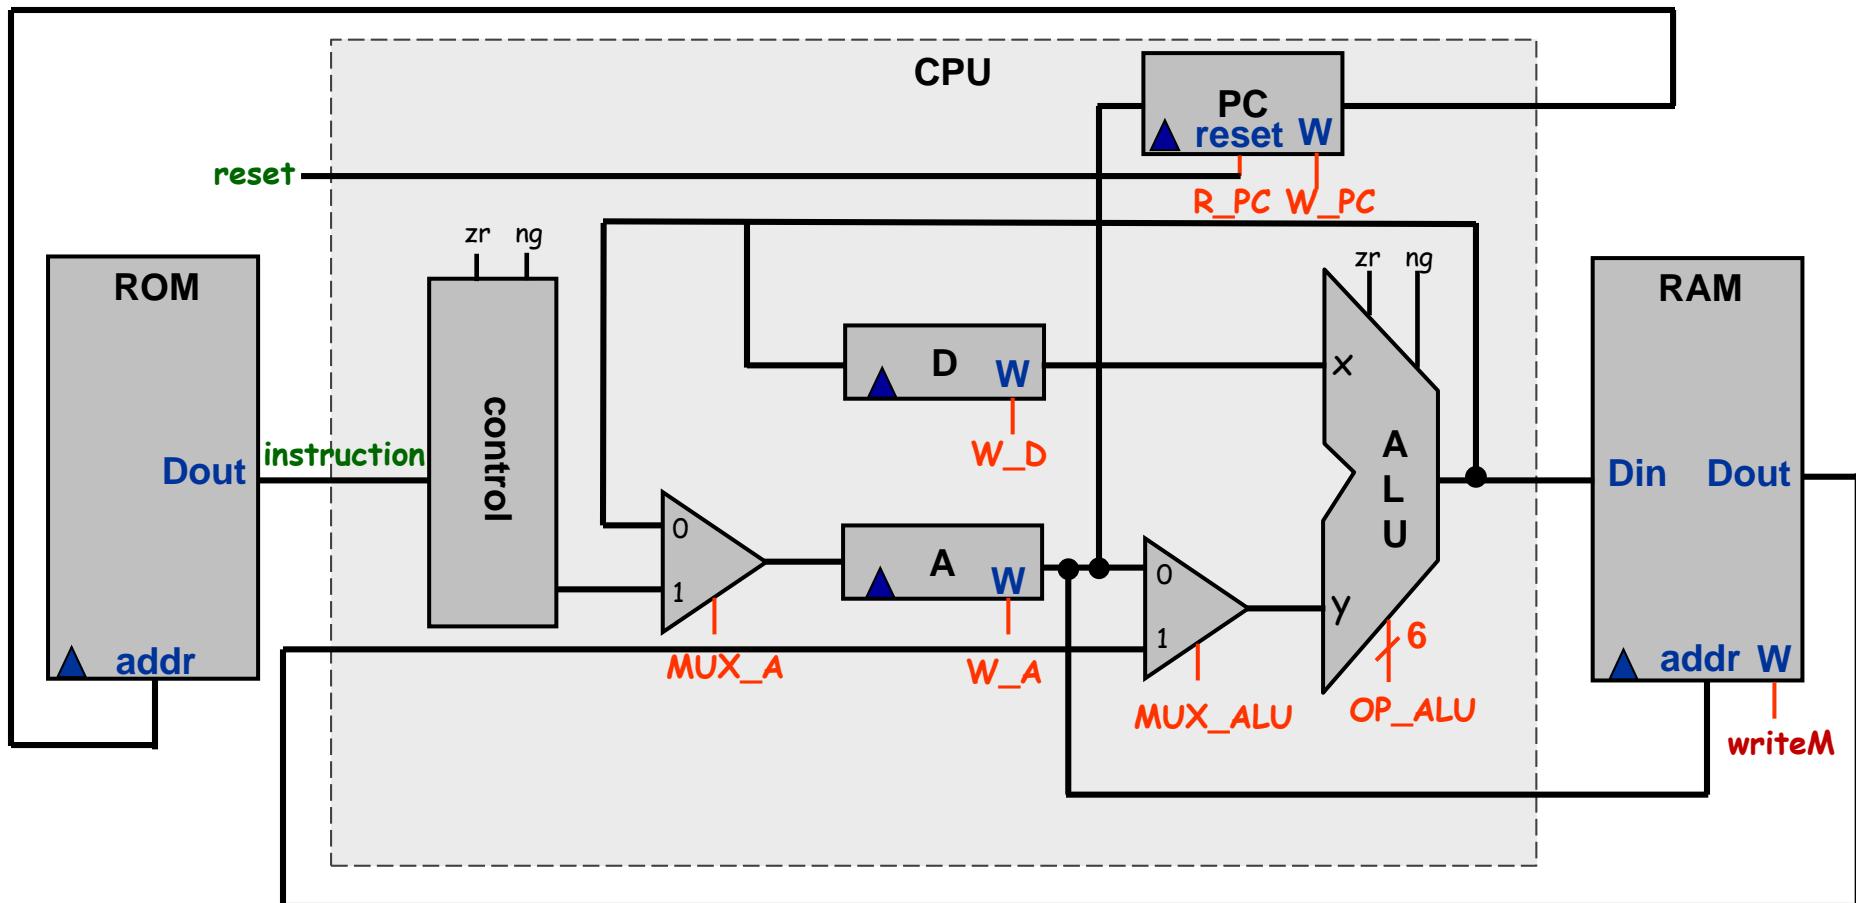
- instruction



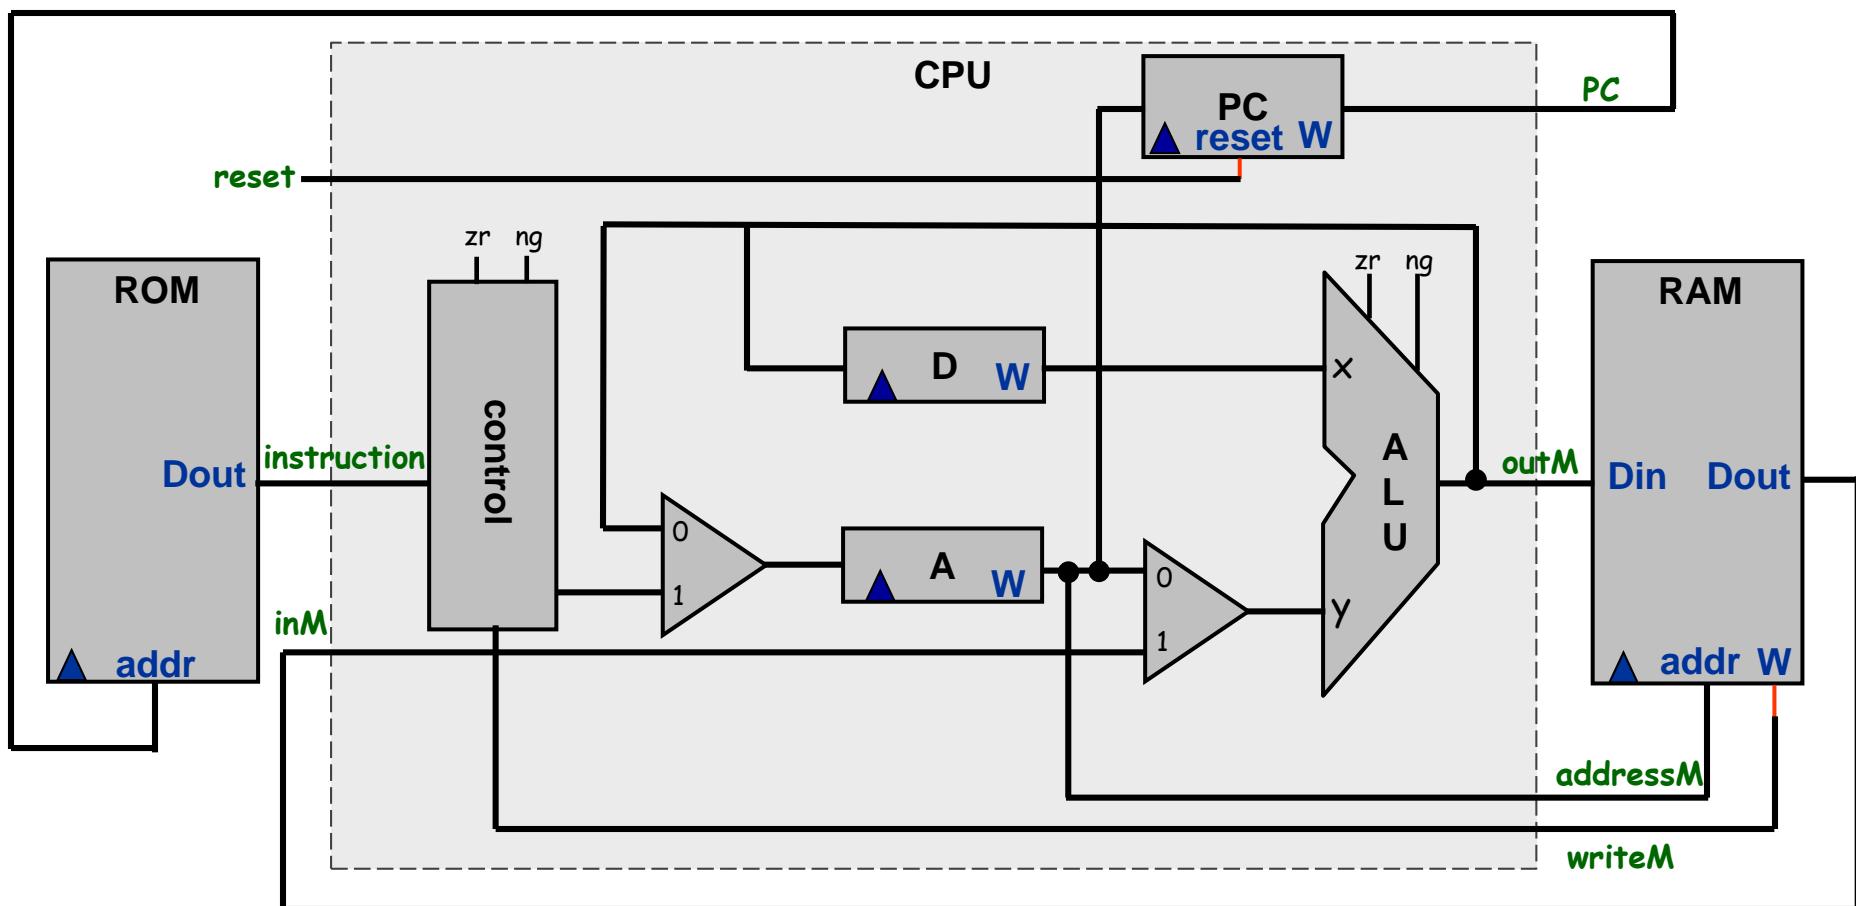
■ Outputs:

- $OP_ALU = \text{C}_1\text{C}_2\text{C}_3\text{C}_4\text{C}_5\text{C}_6$
- $MUX_A = \overline{\text{I}_{15}}$
- $MUX_ALU = \text{A}$
- $W_A = (\text{I}_{15} \& \text{D}_1) + \overline{\text{I}_{15}}$
- $W_D = \text{I}_{15} \& \text{D}_2$
- $writeM = \text{I}_{15} \& \text{D}_3$
- $W_PC = \text{I}_{15} \& ((\text{J}_1 \& \text{ng}) + (\text{J}_2 \& \text{zr}) + (\text{J}_3 \& \text{gt}))$; $\text{gt} = \overline{\text{ng}} \& \overline{\text{zr}}$

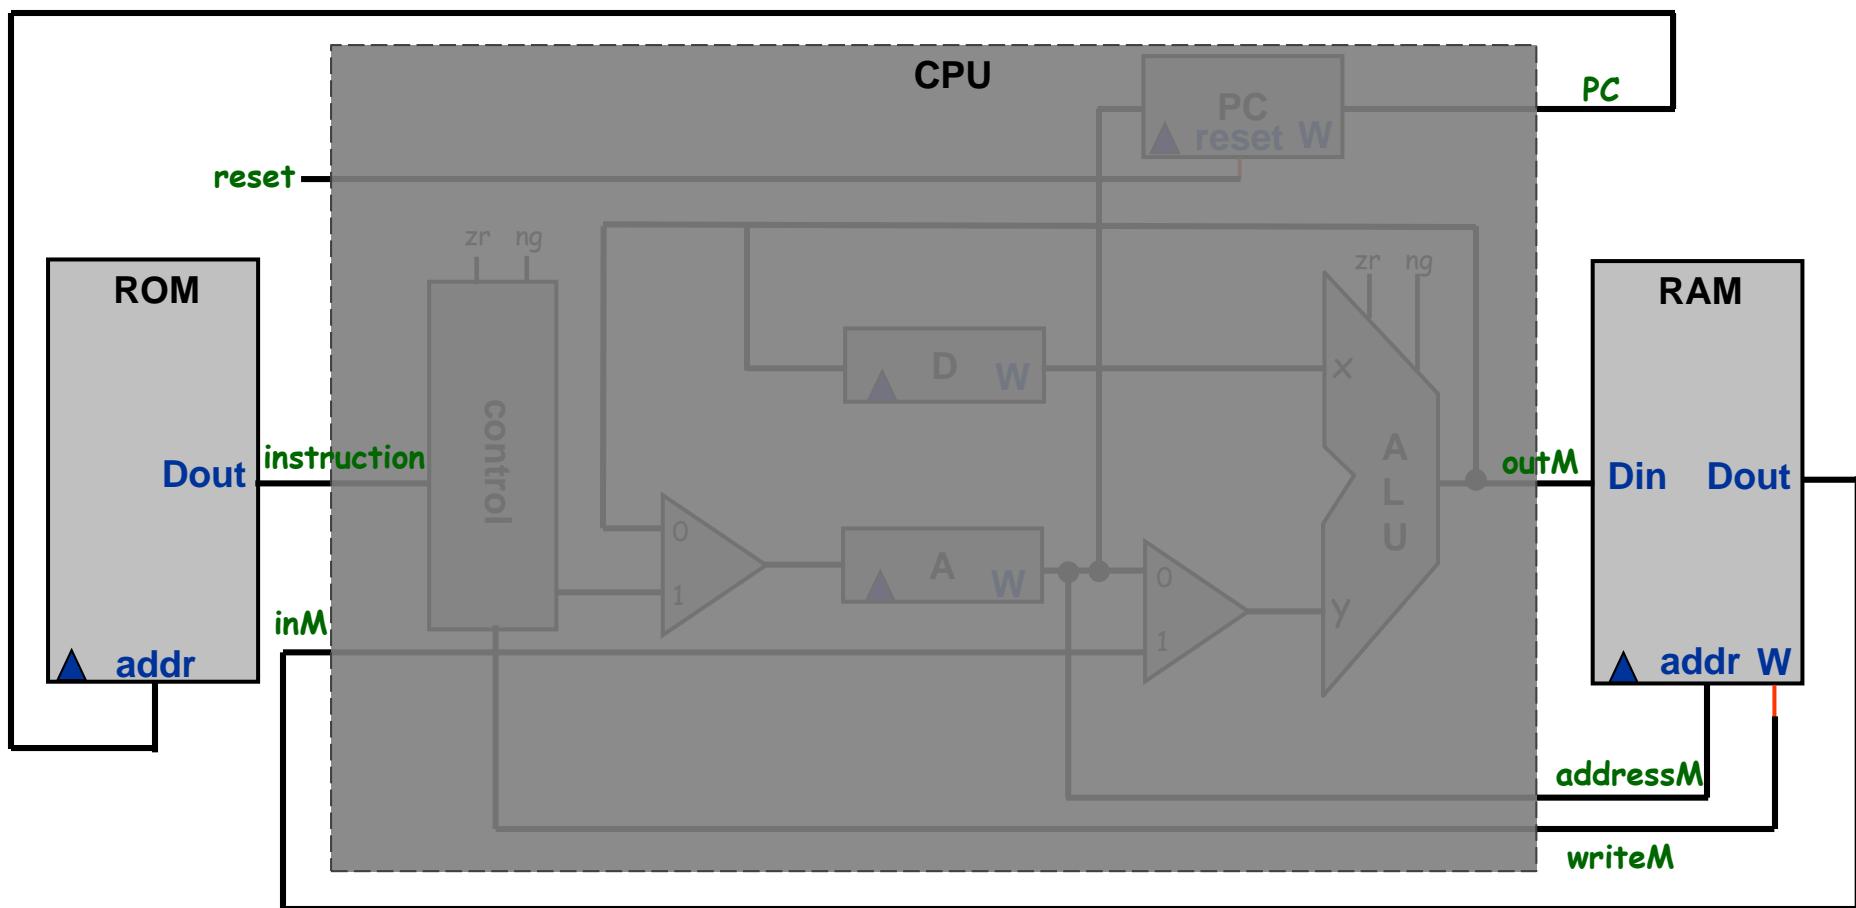
Hack architecture (trace @10 / D=M+1;JGE)



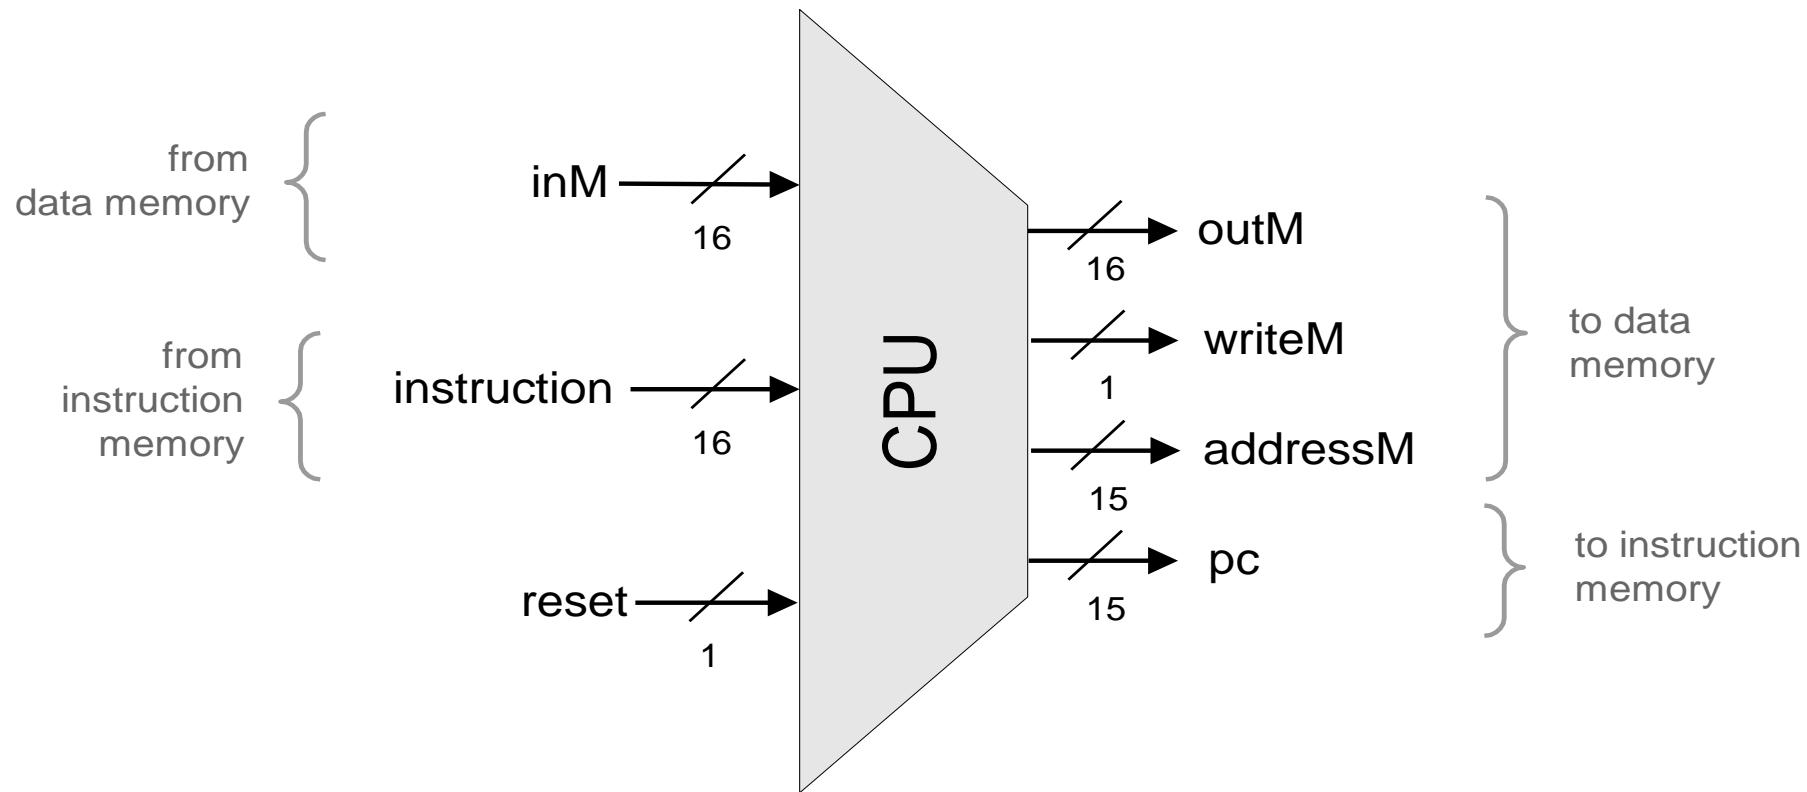
Hack architecture (CPU interface)



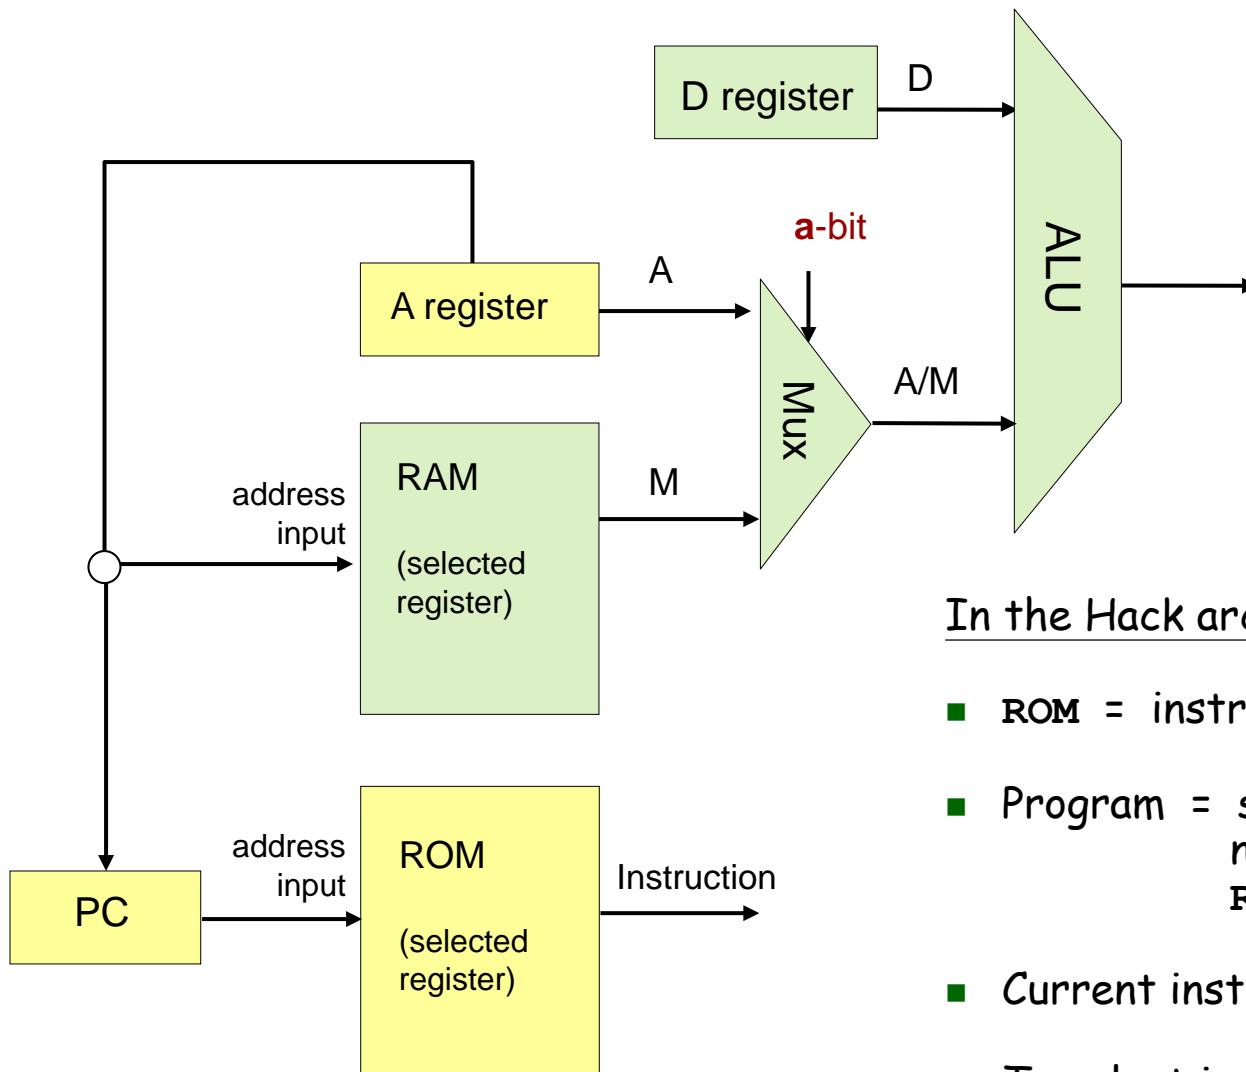
Hack architecture (CPU interface)



Hack CPU



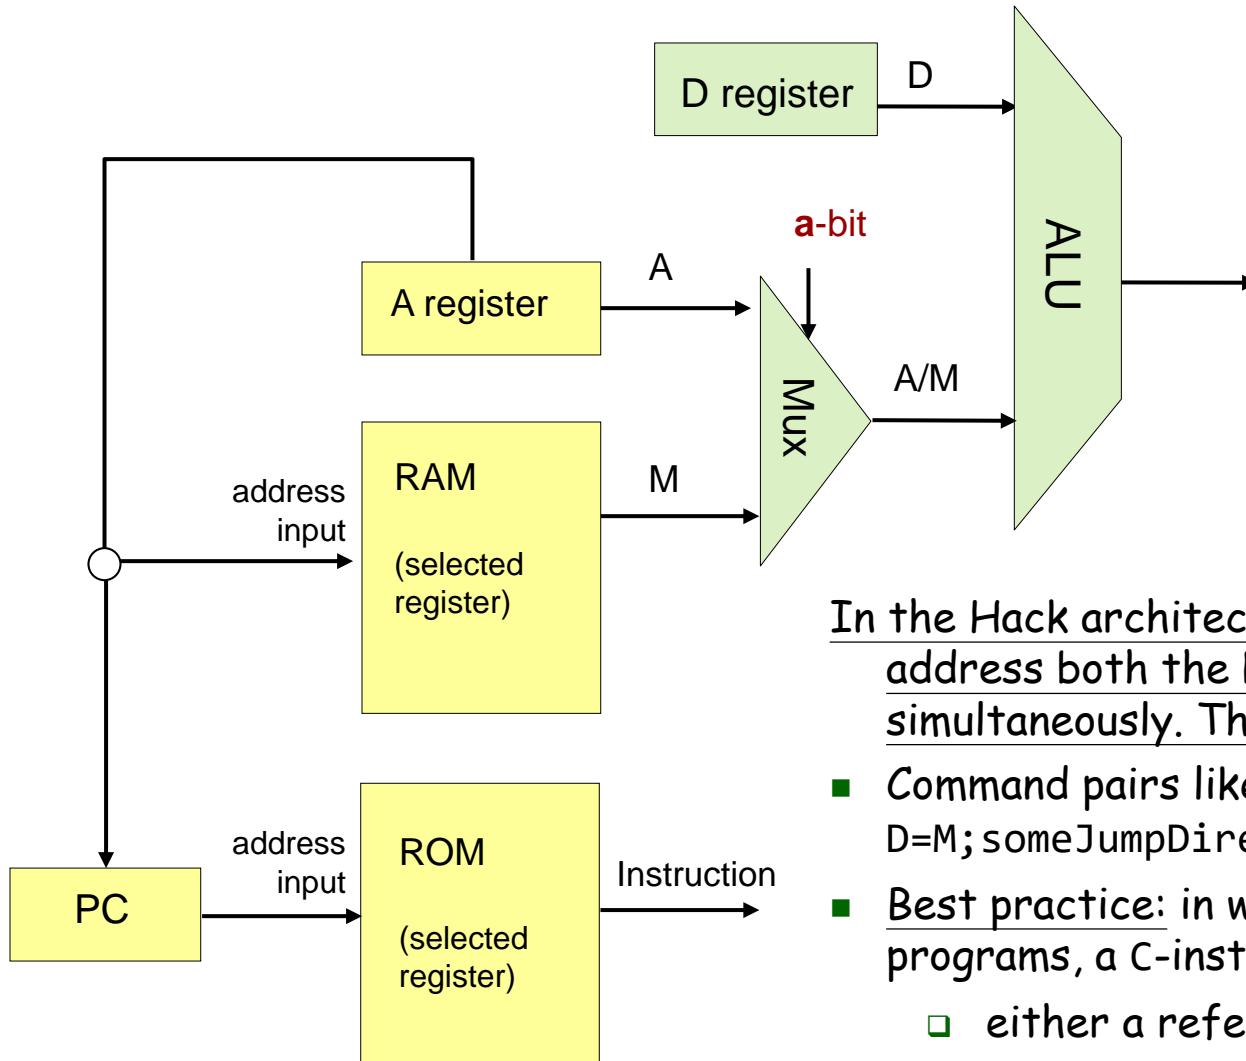
Control (focus on the yellow chips only)



In the Hack architecture:

- ROM = instruction memory
- Program = sequence of 16-bit numbers, starting at ROM[0]
- Current instruction = ROM[PC]
- To select instruction n from the ROM, we set A to n , using the instruction en

Side note (focus on the yellow chip parts only)

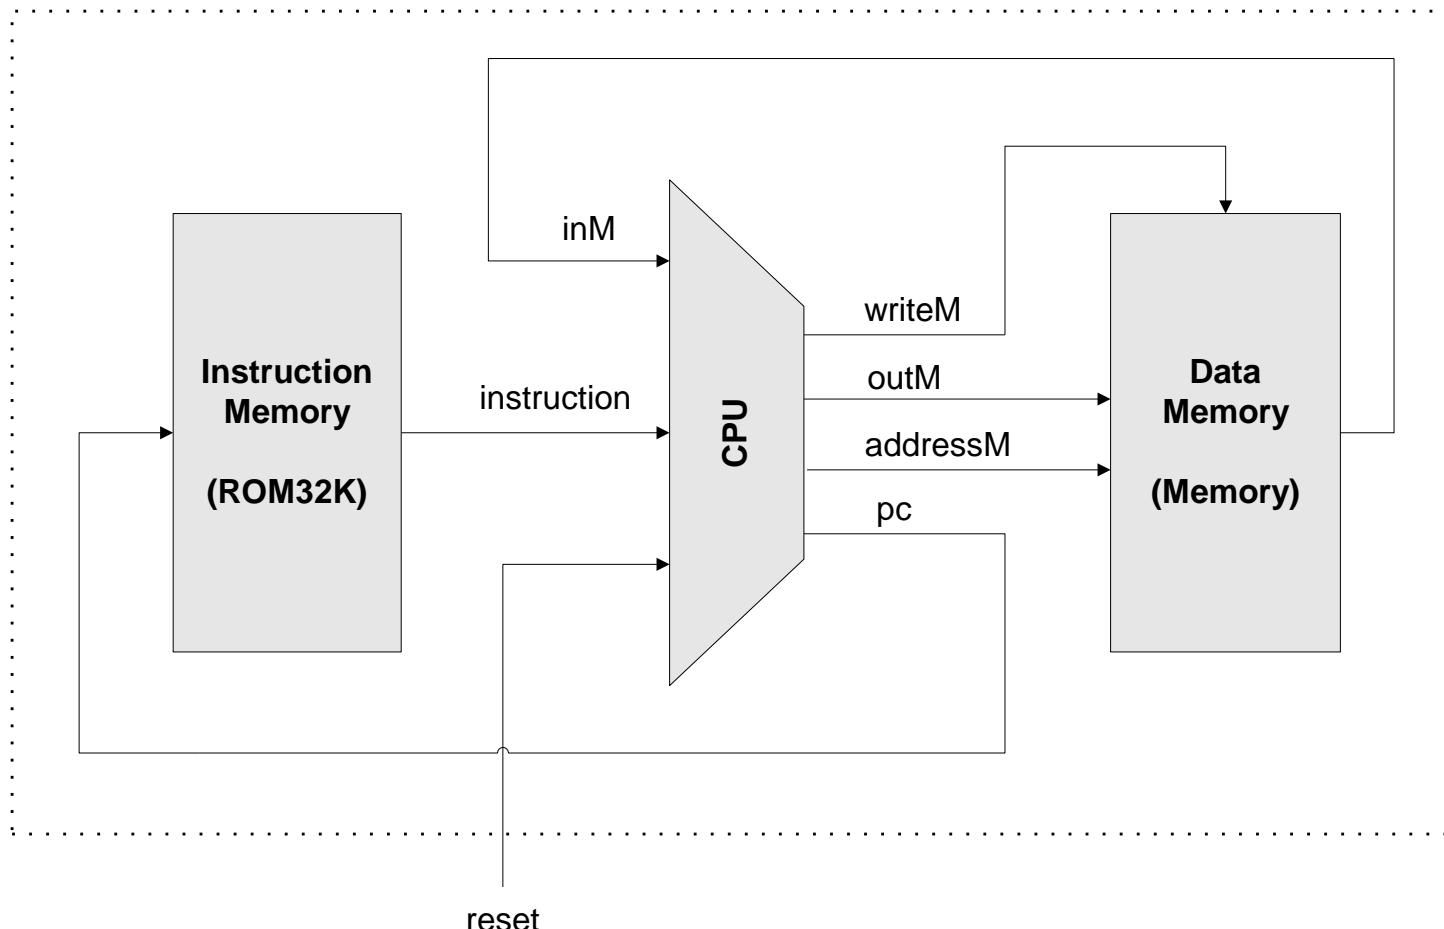


In the Hack architecture, the A register could address both the RAM and the ROM, simultaneously. Therefore:

- Command pairs like @addr followed by D=M; someJumpDirective make no sense
- Best practice: in well-written Hack programs, a C-instruction should contain
 - either a reference to M, or
 - a jump directive,
but not both.

The Hack computer (put together)

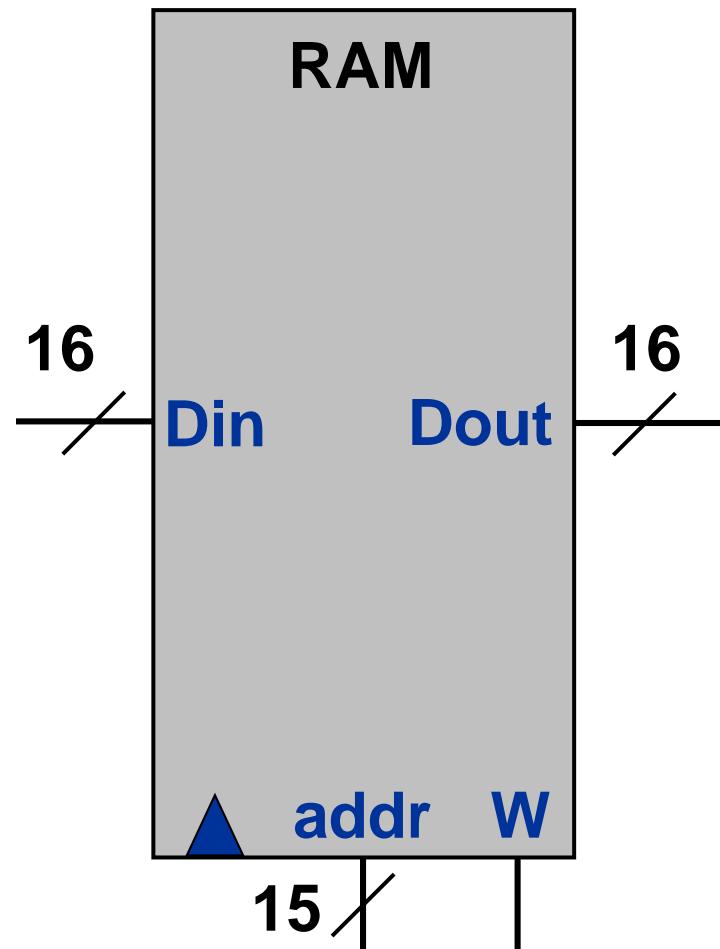
A 16-bit machine consisting of the following elements:



Both memory chips are 16-bit wide and have 15-bit address space.

RAM (data memory)

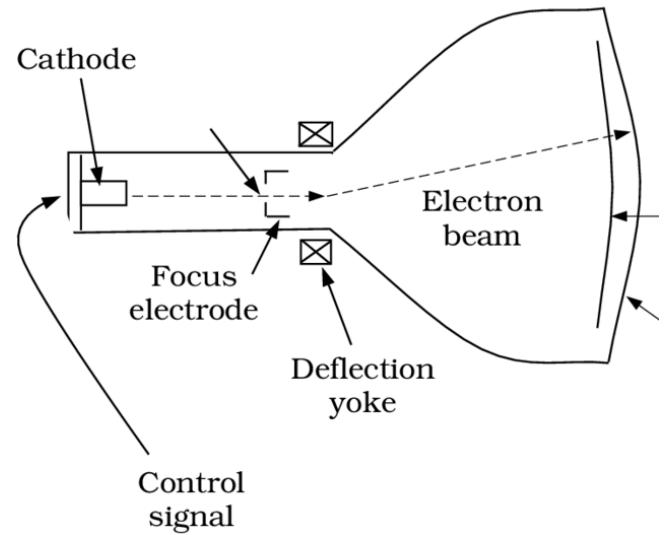
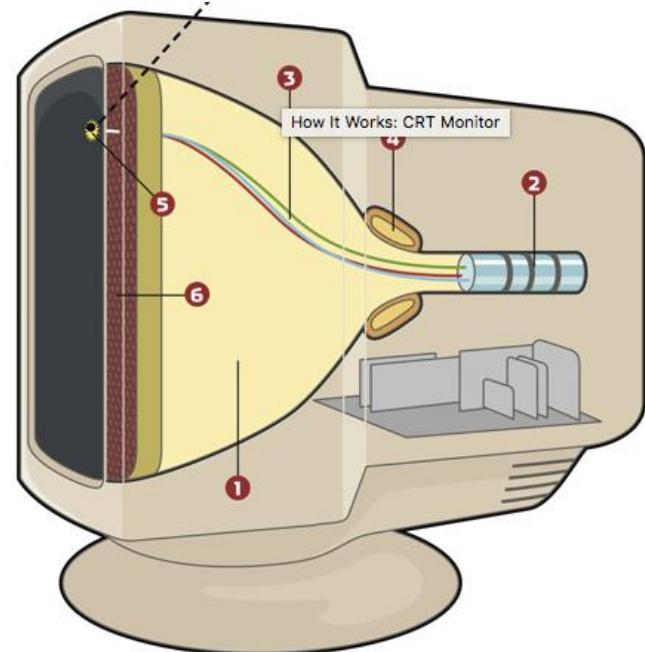
- The RAM used in Hack is different from a normal RAM. It also plays the role for I/O.
- Programmers usually use high-level library for I/O, such as printf, drawline.
- But, at low-level, we usually need to manipulate bits directly for I/O.



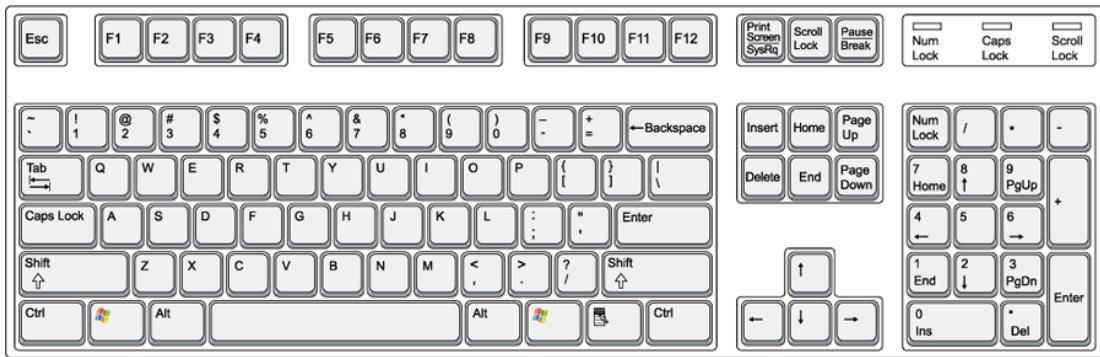
Displays

CRT displays

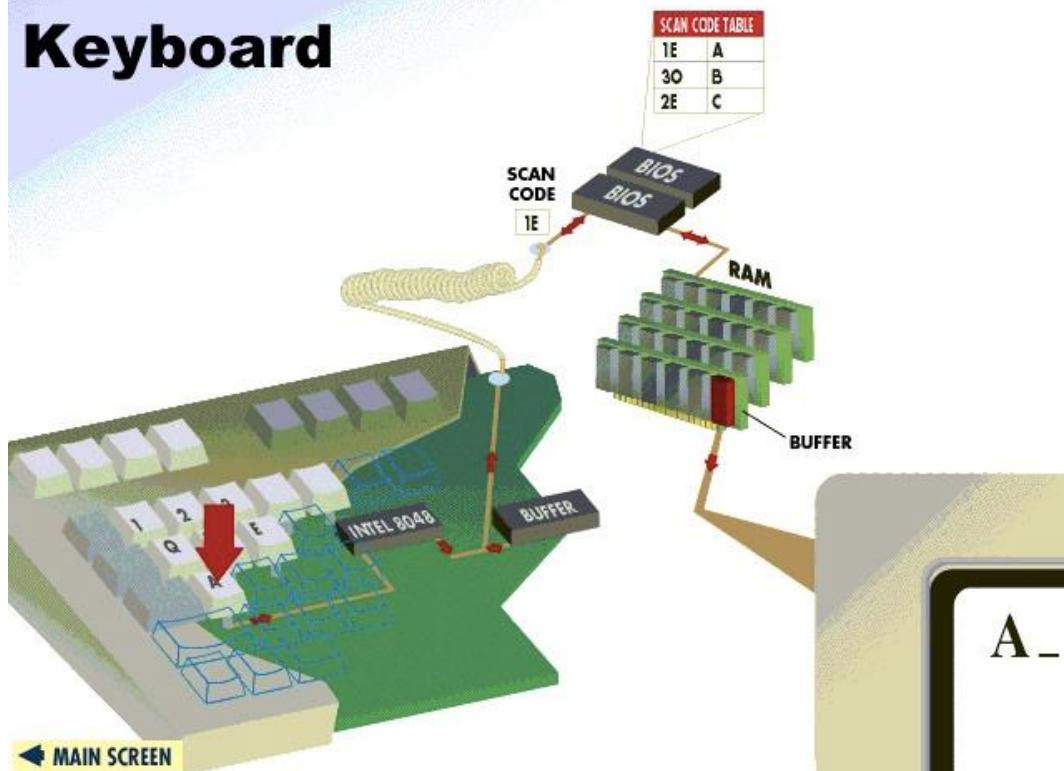
- resolution
- refresh rate



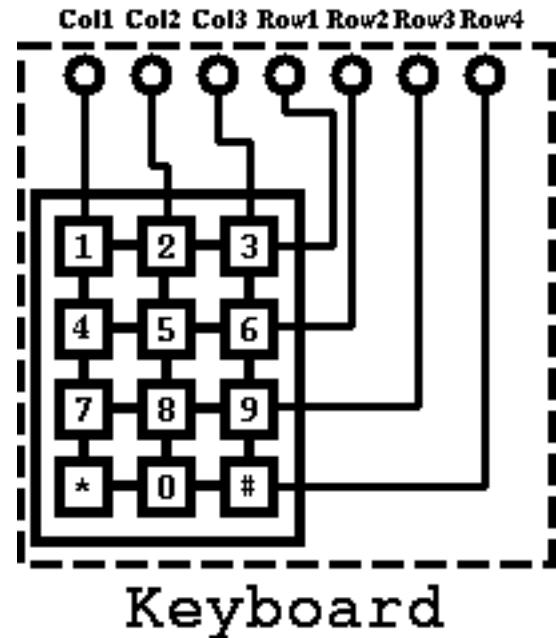
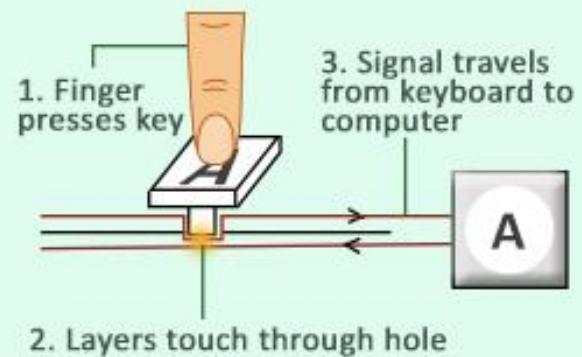
keyboard



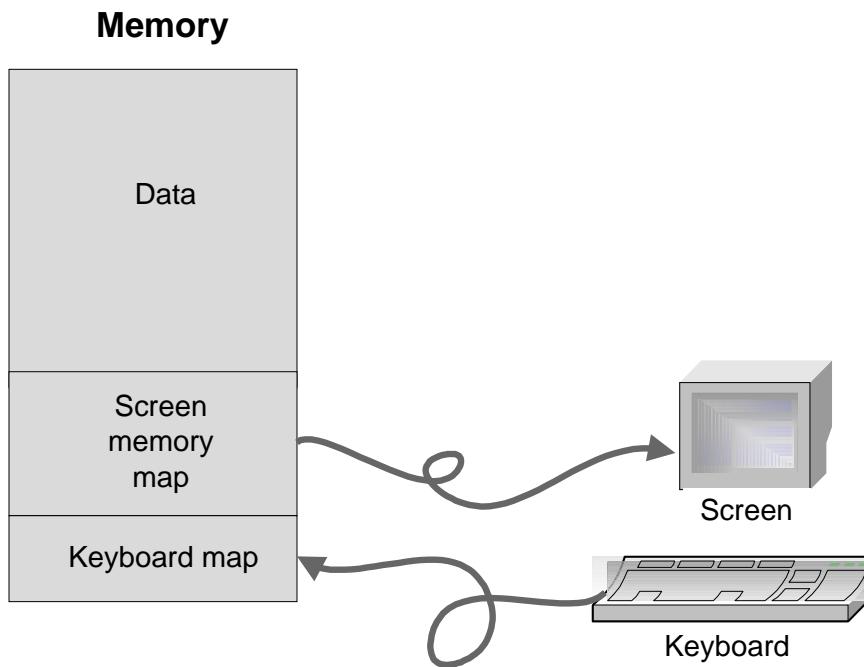
Keyboard



Working of the Keyboard



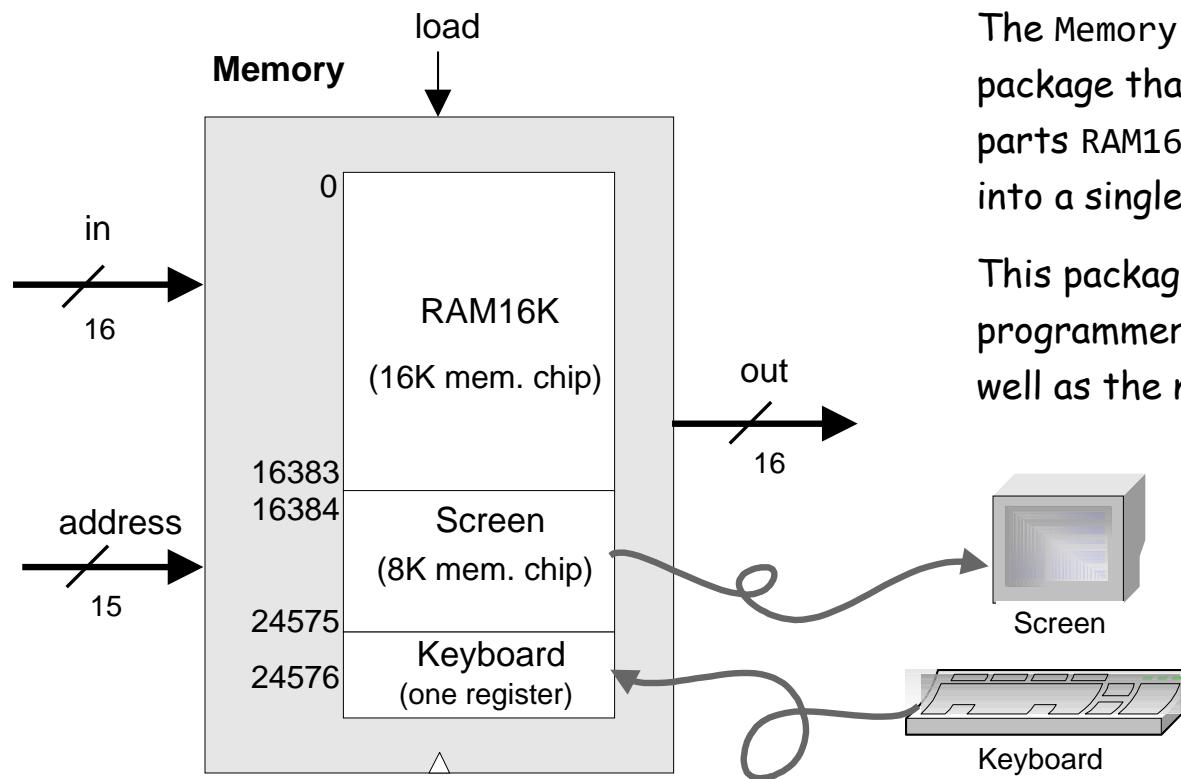
Memory: conceptual / programmer's view



Using the memory:

- ❑ To record or recall values (e.g. variables, objects, arrays), use the first 16K words of the memory
- ❑ To write to the screen (or read the screen), use the next 8K words of the memory
- ❑ To read which key is currently pressed, use the next word of the memory.

Memory: physical implementation



The Memory chip is essentially a package that integrates the three chip-parts RAM16K, Screen, and Keyboard into a single, contiguous address space.

This packaging effects the programmer's view of the memory, as well as the necessary I/O side-effects.

Access logic:

- Access to any address from 0 to 16,383 results in accessing the RAM16K chip-part
- Access to any address from 16,384 to 24,575 results in accessing the Screen chip-part
- Access to address 24,576 results in accessing the keyboard chip-part
- Access to any other address is invalid.

Data memory

Low-level (hardware) read/write logic:

To read $\text{RAM}[k]$: set address to k ,
probe out

To write $\text{RAM}[k]=x$: set address to k ,
set in to x ,
set load to 1,
run the clock

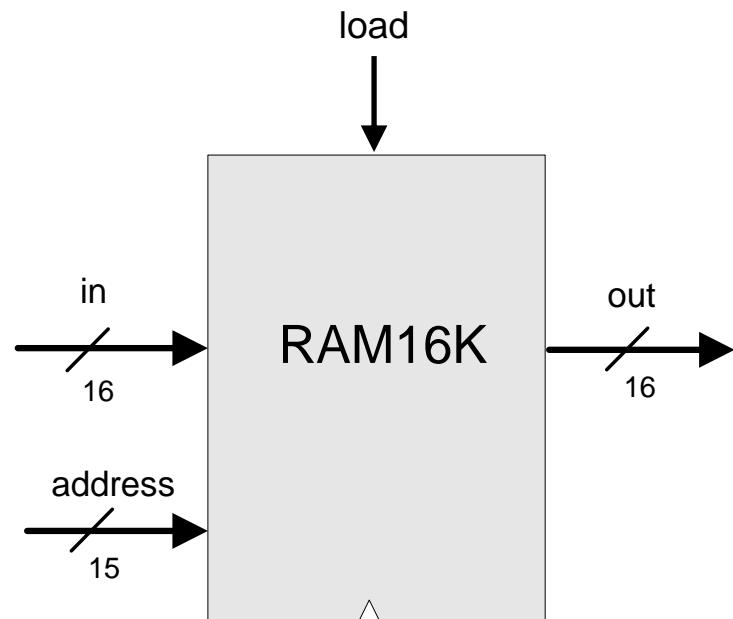
High-level (OS) read/write logic:

To read $\text{RAM}[k]$: use the OS command $\text{out} = \text{peek}(k)$

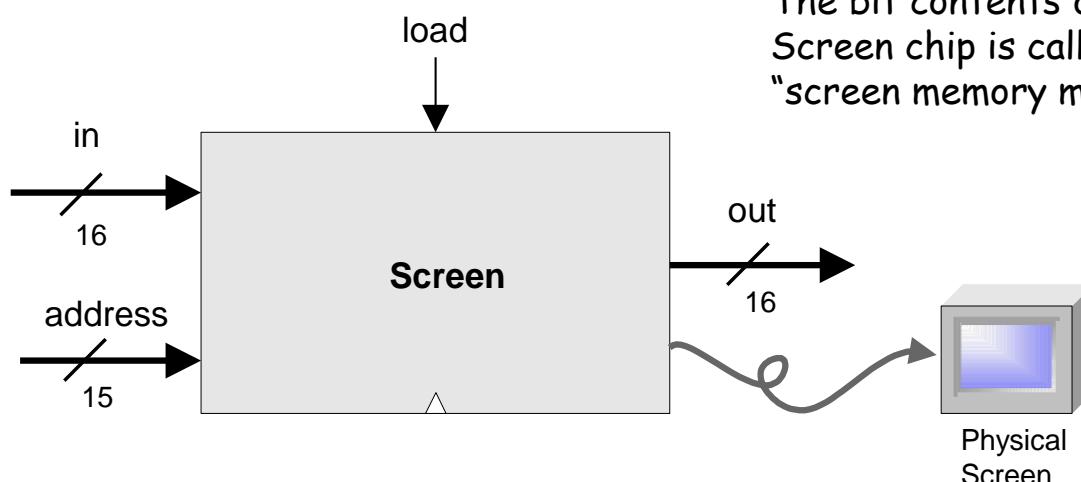
To write $\text{RAM}[k]=x$: use the OS command $\text{poke}(k, x)$

peek and poke are OS commands whose implementation should effect the same behavior as the low-level commands

More about peek and poke this later in the course, when we'll write the OS.



Screen



The bit contents of the Screen chip is called the "screen memory map"

In the Hack platform, the screen is implemented as an 8K 16-bit RAM chip with a side effect of refreshing.

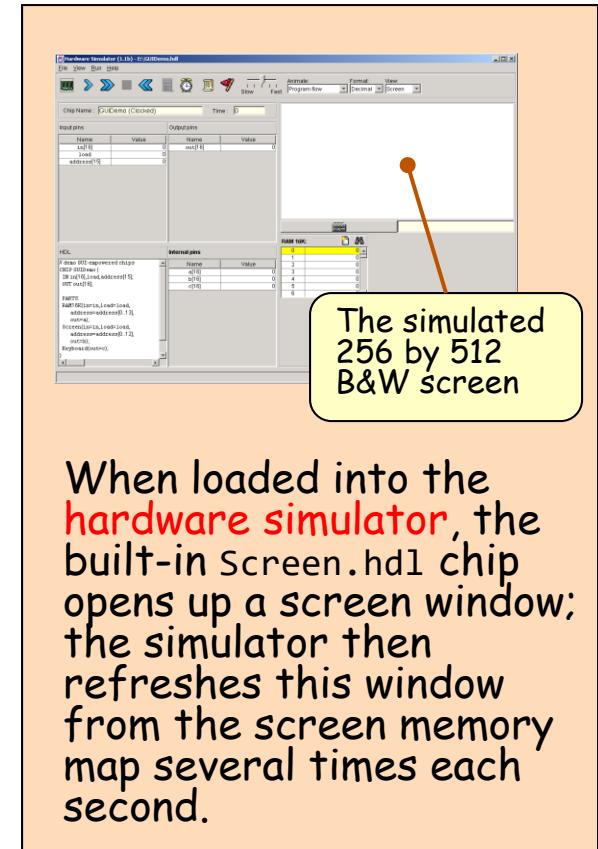
The Screen chip has a basic RAM chip functionality:

- read logic: $out = \text{Screen}[address]$
- write logic: if load then $\text{Screen}[address] = in$

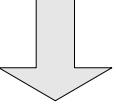
Side effect:

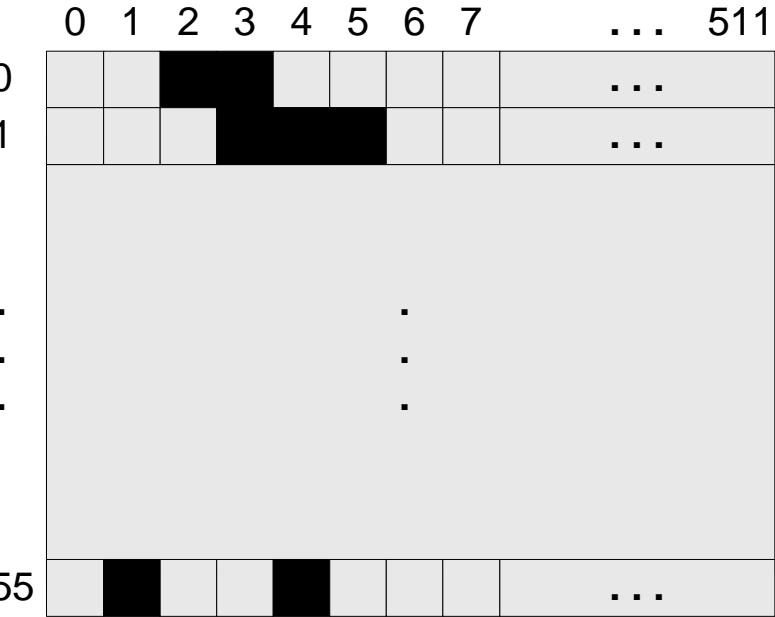
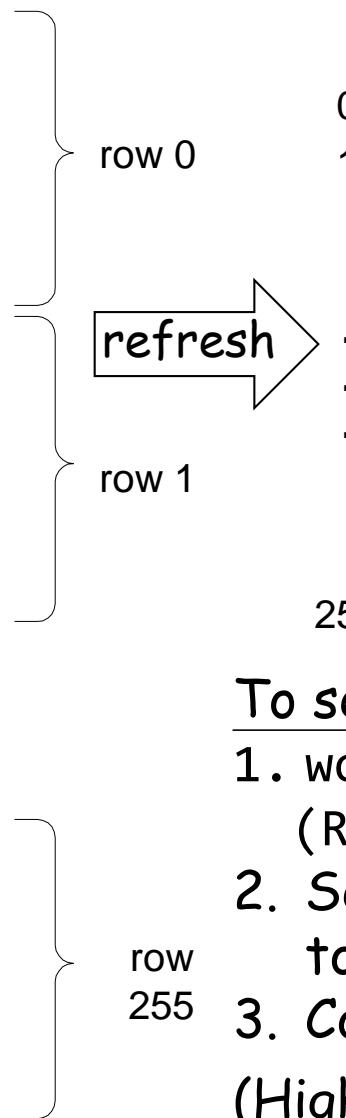
Continuously refreshes a 256 by 512 black-and-white screen device

Simulated screen:



Screen memory map

(16384)	0	0011000000000000
	1	0000000000000000
	⋮	
31	0000000000000000	
32	0001100000000000	
33	0000000000000000	
	⋮	
63	0000000000000000	
		
8129	0100100000000000	
8130	0000000000000000	
	⋮	
8160	0000000000000000	



To set pixel (row, col) black

1. word=Screen[32*row + col/16]
(RAM[16384 + 32*row + col/16])
2. Set the (col%16)-th bit of word to 1
3. Commit word to the RAM
(High-level: use the OS command drawPixel(row,col))

keyboard

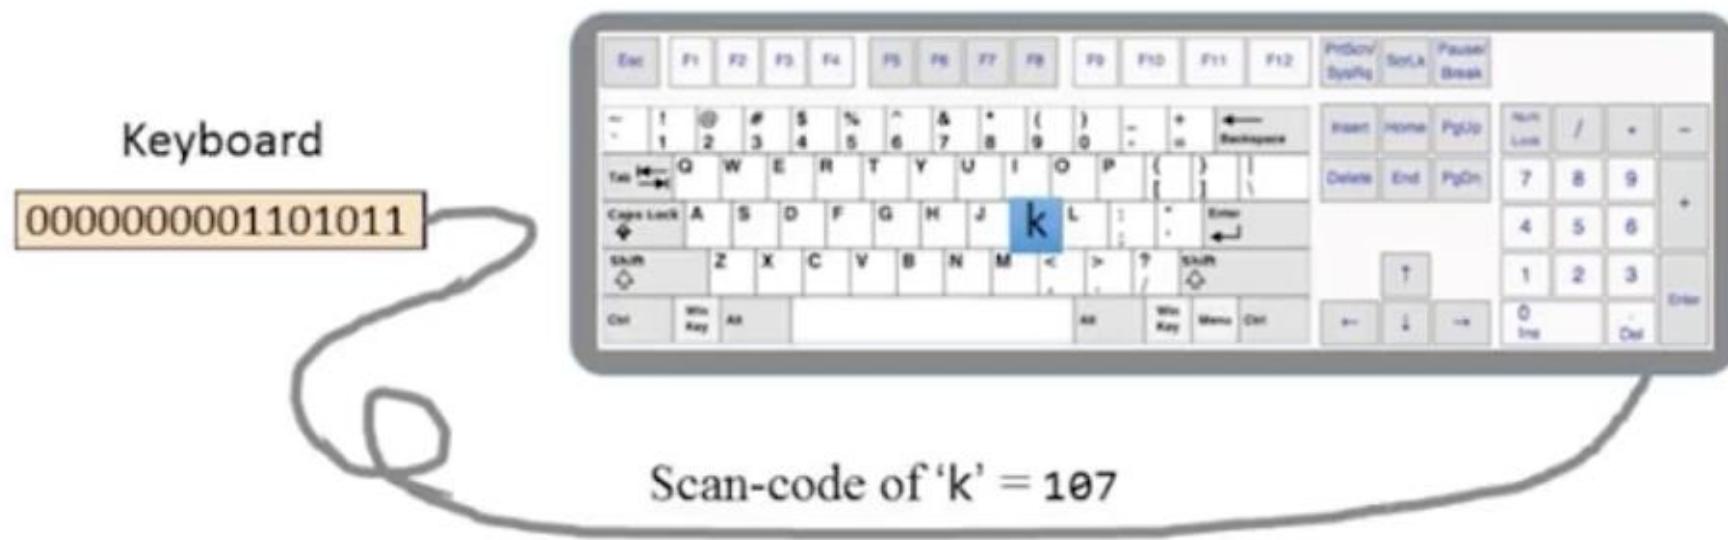
- A 16-bit register is used to keep the key stroke.



When a key is pressed on the keyboard, the key's scan code appears in the keyboard memory map .

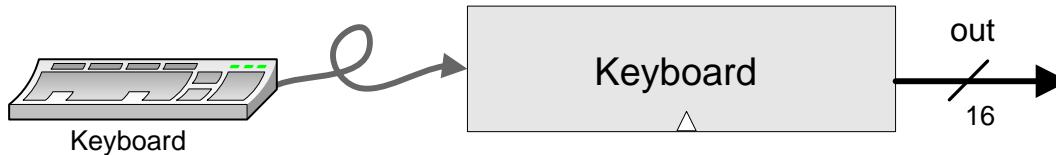
keyboard

- A 16-bit register is used to keep the key stroke.



When a key is pressed on the keyboard, the key's scan code appears in the keyboard memory map .

Keyboard



Keyboard chip: a single 16-bit register

Input: scan-code (16-bit value) of the currently pressed key, or 0 if no key is pressed

Output: same

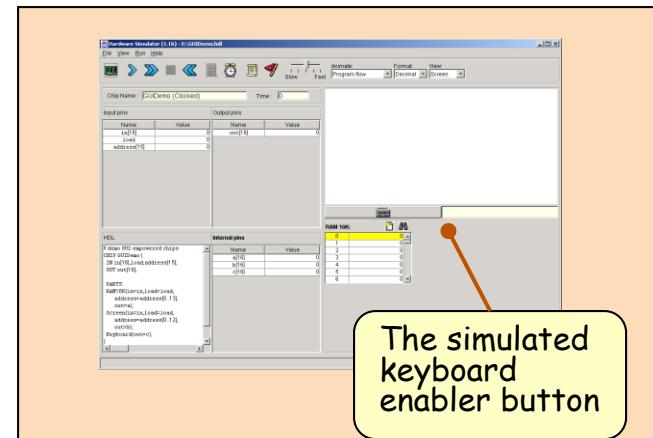
Special keys:

Key pressed	Keyboard output	Key pressed	Keyboard output
newline	128	end	135
backspace	129	page up	136
left arrow	130	page down	137
up arrow	131	insert	138
right arrow	132	delete	139
down arrow	133	esc	140
home	134	f1-f12	141-152

How to read the keyboard:

- ❑ Low-level (hardware): probe the contents of the Keyboard chip
- ❑ High-level: use the OS command `keyPressed()`
(effects the same operation, discussed later in the course, when we'll write the OS).

Simulated keyboard:



The simulated keyboard enabler button

The keyboard is implemented as a built-in `Keyboard.hdl` chip. When this java chip is loaded into the simulator, it connects to the regular keyboard and pipes the scan-code of the currently pressed key to the keyboard memory map.

Some scan codes

key	code
(space)	32
!	33
“	34
#	35
\$	36
%	37
&	38
‘	39
(40
)	41
*	42
+	43
,	44
-	45
.	46
/	47

key	code
0	48
1	49
...	...
9	57
:	58
;	59
<	60
=	61
>	62
?	63
@	64

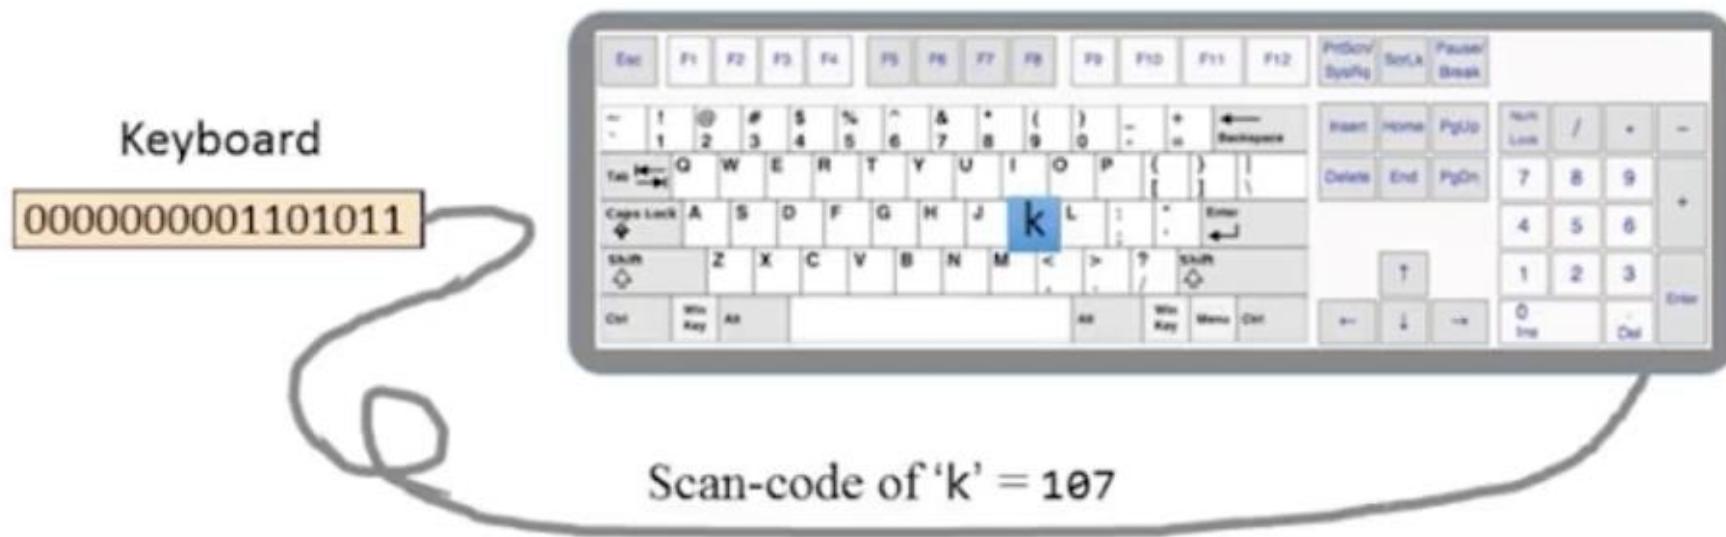
key	code
A	65
B	66
C	...
...	...
Z	90
[91
/	92
]	93
^	94
-	95
`	96

key	code
a	97
b	98
c	99
...	...
z	122
{	123
	124
}	125
~	126

key	code
newline	128
backspace	129
left arrow	130
up arrow	131
right arrow	132
down arrow	133
home	134
end	135
Page up	136
Page down	137
insert	138
delete	139
esc	140
f1	141
...	...
f12	152

(Subset of Unicode)

Keyboard memory map

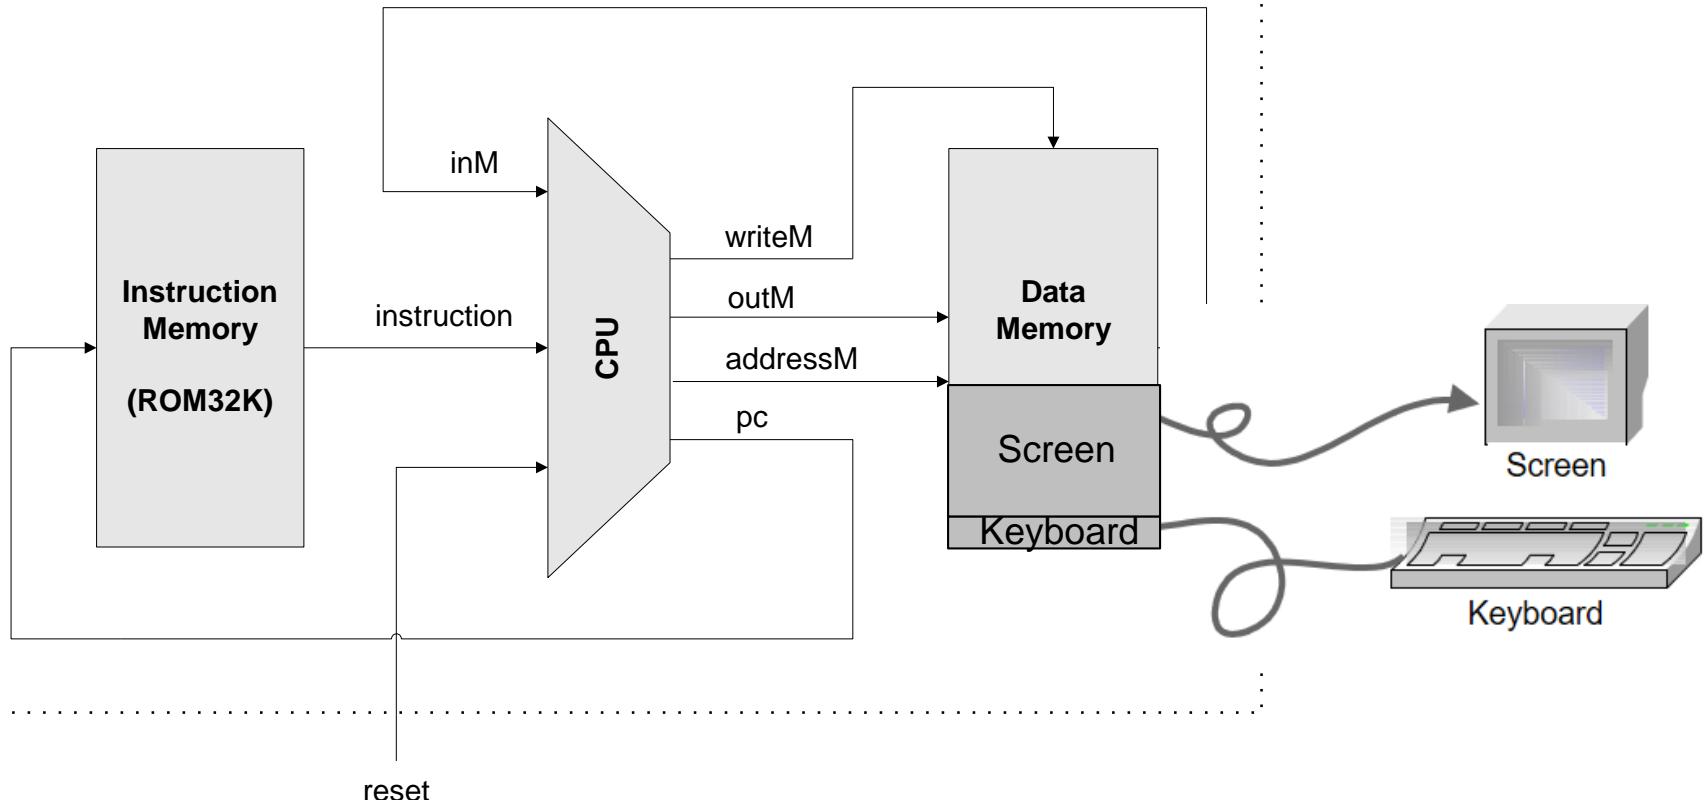


To check which key is currently pressed:

- Probe the content of the Keyboard chip
- In the Hack computer, probe the content of RAM[24576]
- If the register contains 0, no key is pressed.

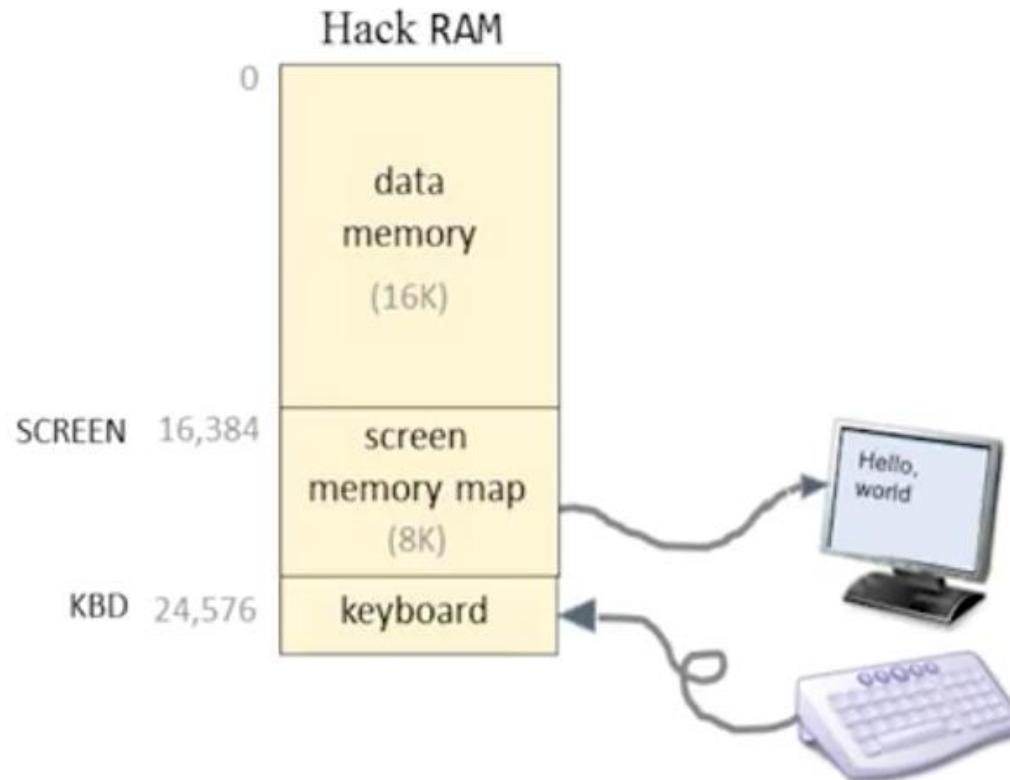
The Hack computer (put together)

A 16-bit machine consisting of the following elements:



Both memory chips are 16-bit wide and have 15-bit address space.

Assembly programming with I/O

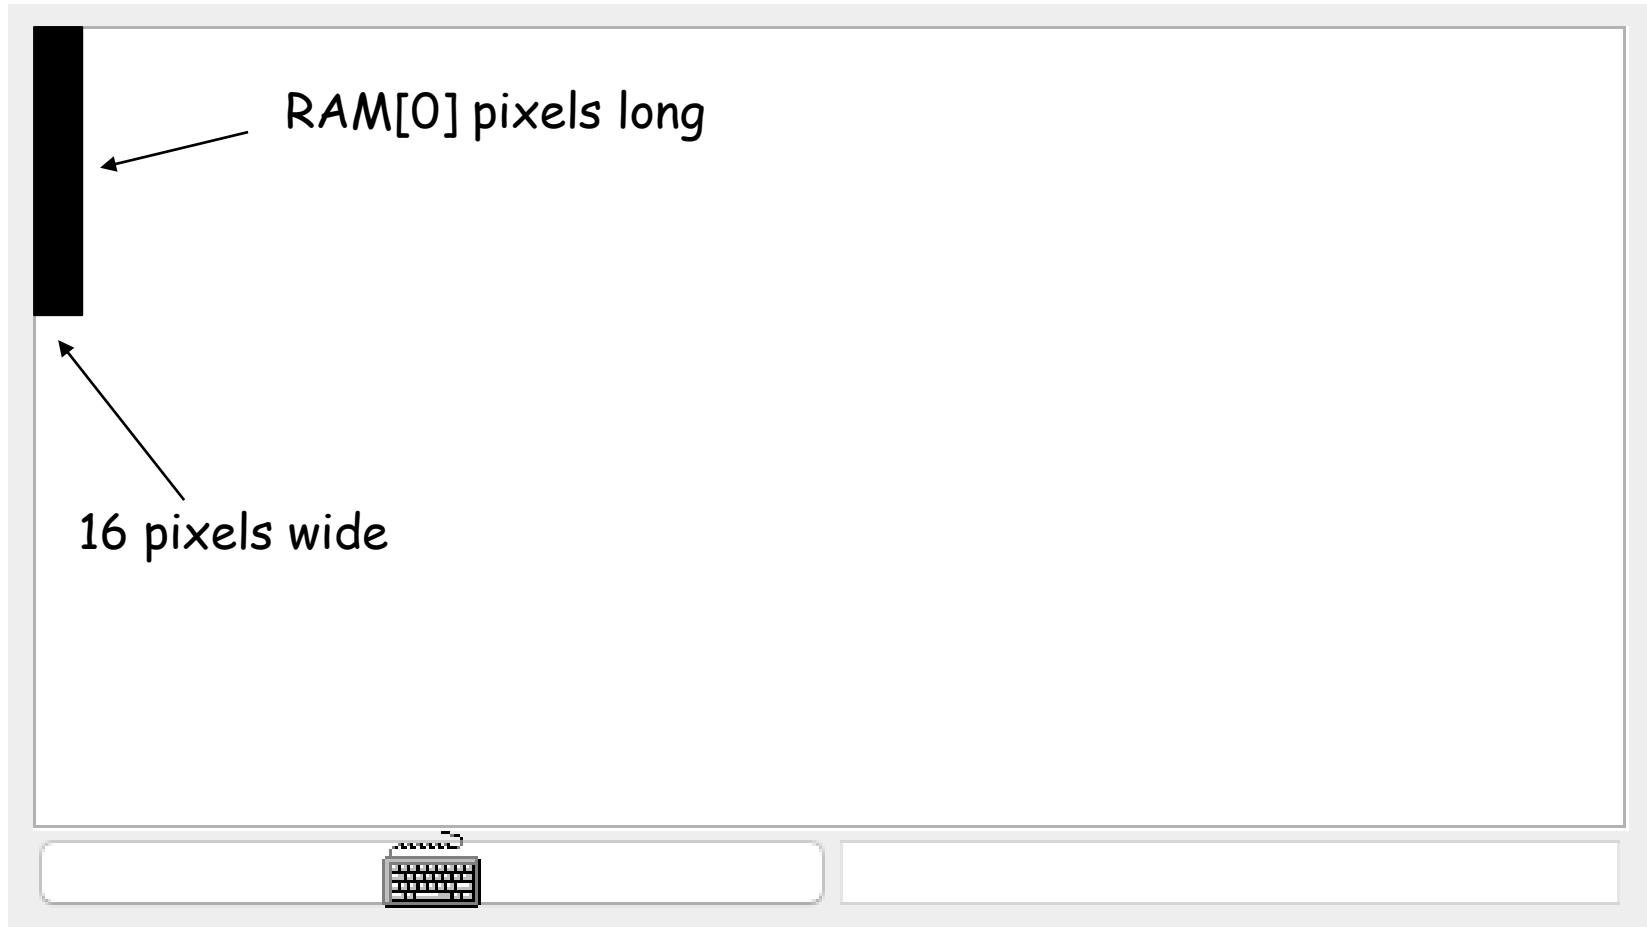


Hack language convention:

- **SCREEN**: base address of the screen memory map, 16,384.
- **KBD**: address of the keyboard memory map, 24,576.

Example: draw a rectangle

- Draw a filled rectangle at the upper left corner of the screen, 16 pixels wide and RAM[0] pixels long. ([demo](#))



Example: draw a rectangle (pseudo code)

```
// for (i=0; i<n; i++)
//     draw 16 black pixels at the beginning of row i

addr = SCREEN
n = RAM[0]
i = 0

LOOP:
    if (i>=n) goto END
    RAM[addr] = -1 // 1111 1111 1111 1111
    addr = addr+32 // advances to the next row
    i++;
    goto LOOP

END:
    goto END
```

Example: draw a rectangle (assembly)

```
@SCREEN  
D=A  
@addr  
M=D      // addr = SCREEN  
  
@0  
D=M  
@n  
M=D      // n = RAM[0]  
  
@i  
M=0      // i=0
```

```
addr = SCREEN  
n = RAM[0]  
i = 0  
  
LOOP:  
    if (i>=n) goto END  
    RAM[addr] = -1  
    addr = addr+32  
    i++;  
    goto LOOP  
END:  
    goto END
```

Example: draw a rectangle (assembly)

(LOOP)

@i

D=M

@n

D=D-M

@END

D; JGE

@addr

A=M

M=-1

addr = SCREEN

n = RAM[0]

i = 0

LOOP:

if (i>=n) goto END

RAM[addr] = -1

addr = addr+32

i++;

goto LOOP

END:

goto END

Example: draw a rectangle (assembly)

(LOOP)

@i

D=M

@n

D=D-M

@END

D ; JGE

@addr

A=M

M=-1

addr = SCREEN

n = RAM[0]

i = 0

LOOP:

if (i>=n) goto END

RAM[addr] = -1

addr = addr+32

i++;

goto LOOP

END:

goto END

Example: draw a rectangle (assembly)

```
@32  
  
D=A  
  
@addr  
  
M=D+M    // addr = addr+32  
  
  
@i  
  
M=M+1    // i++  
  
  
@LOOP  
  
0; JMP   // goto LOOP  
  
  
(END)  
  
@END  
  
0; JMP
```

```
addr = SCREEN  
  
n = RAM[0]  
  
i = 0  
  
  
LOOP:  
  
    if (i>=n) goto END  
  
    RAM[addr] = -1  
  
    addr = addr+32  
  
    i++;  
  
    goto LOOP  
  
END:  
  
    goto END
```

Example: draw a rectangle (assembly)

```
@32  
  
D=A  
  
@addr  
  
M=D+M    // addr = addr+32  
  
  
@i  
  
M=M+1    // i++  
  
  
@LOOP  
  
0; JMP   // goto LOOP  
  
  
(END)  
  
@END  
  
0; JMP
```

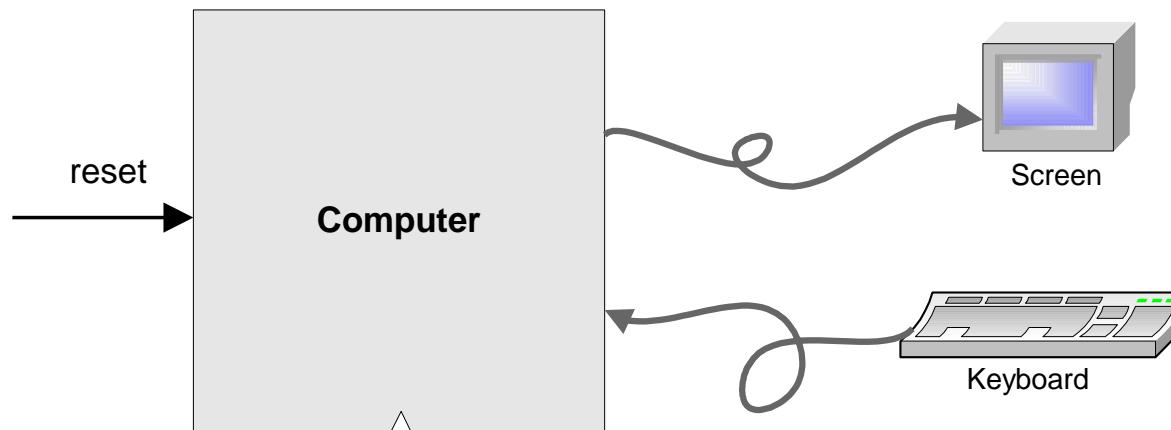
```
addr = SCREEN  
  
n = RAM[0]  
  
i = 0  
  
  
LOOP:  
  
    if (i>=n) goto END  
  
    RAM[addr] = -1  
  
    addr = addr+32  
  
    i++;  
  
    goto LOOP  
  
END:  
  
    goto END
```

Example: draw a rectangle (assembly)

```
@32  
  
D=A  
  
@addr  
  
M=D+M    // addr = addr+32  
  
  
@i  
  
M=M+1    // i++  
  
  
@LOOP  
  
0; JMP   // goto LOOP  
  
  
(END)  
  
@END  
  
0; JMP
```

```
addr = SCREEN  
  
n = RAM[0]  
  
i = 0  
  
  
LOOP:  
  
    if (i>=n) goto END  
  
    RAM[addr] = -1  
  
    addr = addr+32  
  
    i++;  
  
    goto LOOP  
  
END:  
  
    goto END
```

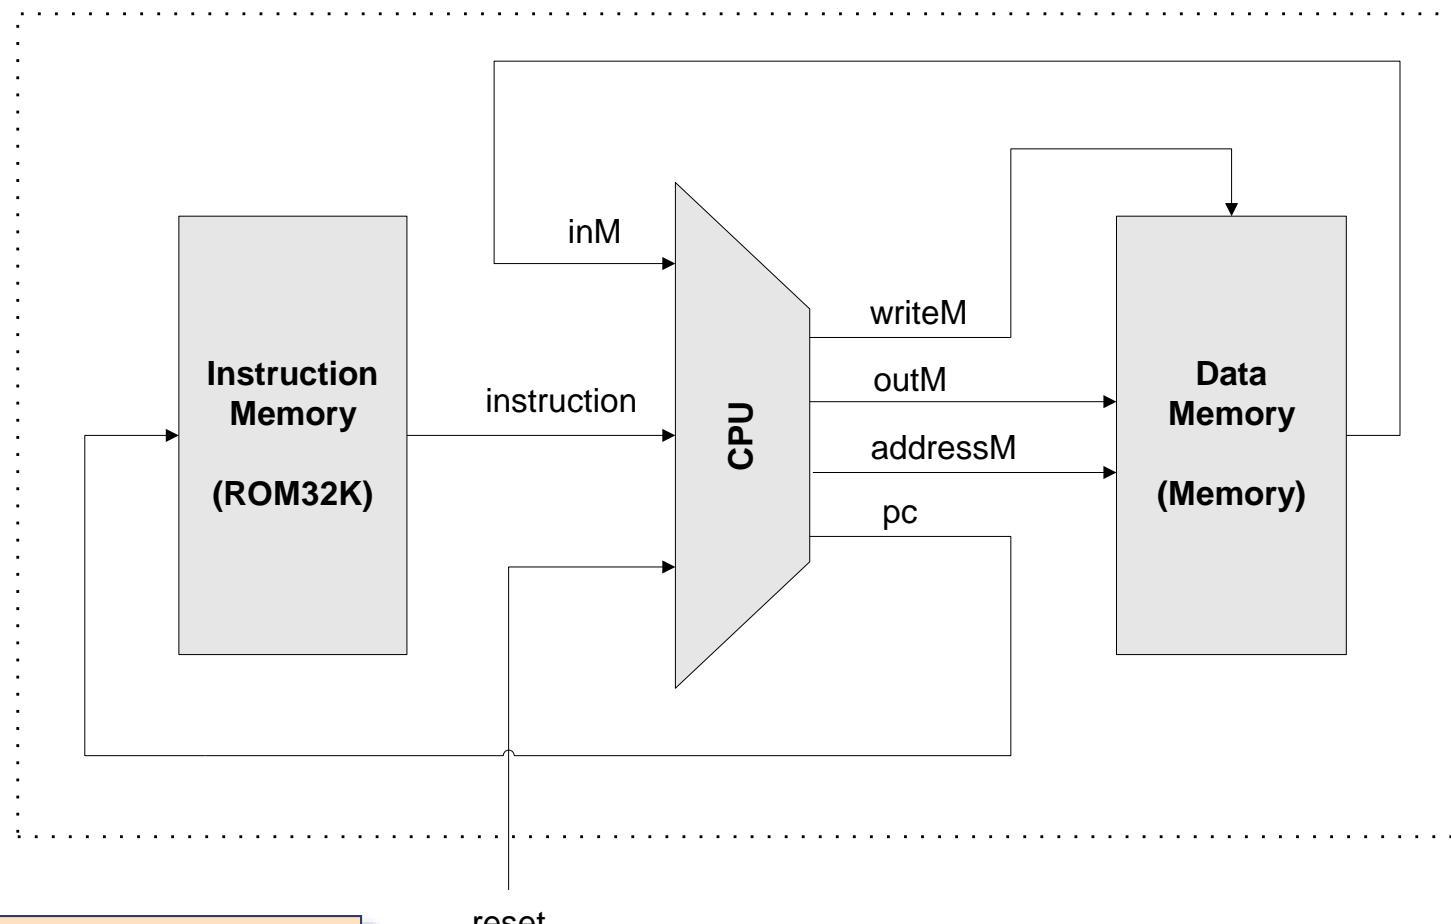
Project #5: Computer-on-a-chip interface



Chip Name: Computer // Topmost chip in the Hack platform
Input: reset
Function: When reset is 0, the program stored in the computer's ROM executes. When reset is 1, the execution of the program restarts. Thus, to start a program's execution, reset must be pushed "up" (1) and "down" (0).

From this point onward the user is at the mercy of the software. In particular, depending on the program's code, the screen may show some output and the user may be able to interact with the computer via the keyboard.

Computer-on-a-chip implementation



```
CHIP Computer {
    IN reset;
    PARTS:
        // implementation missing
}
```

Implementation:

- You need to implement Memory and CPU first.
- Simple, the chip-parts do all the hard work.

Perspective: from here to a “real” computer

- Caching
- More I/O units
- Special-purpose processors (I/O, graphics, communications, ...)
- Multi-core / parallelism
- Efficiency
- Energy consumption considerations
- And more ...