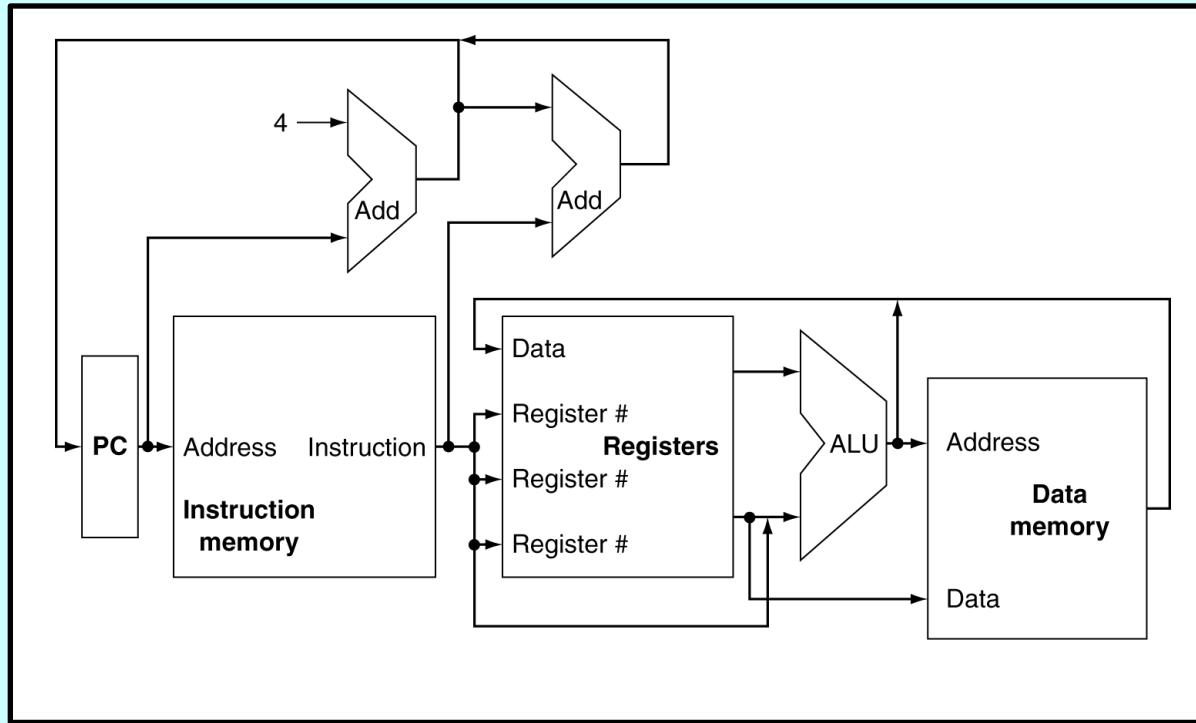


# Operations in Computer Hardware



# Chris Gregg, Based on Slides by Eric Roberts

## CS 208E

### October 7, 2021

# When we compile C code, it turns into Assembly Instructions, which are very basic

- C code:

$$f = (g + h) - (i + j);$$

- Compiled code (for a "MIPS" processor):

```
add t0, g, h  # temp t0 = g + h
```

```
add t1, i, j  # temp t1 = i + j
```

```
sub f, t0, t1  # f = t0 - t1
```

# When we compile C code, it turns into Assembly Instructions, which are very basic

- C code:

$$f = (g + h) - (i + j);$$

- Compiled code (for a "MIPS" processor):

```
add t0, g, h  # temp t0 = g + h
```

```
add t1, i, j  # temp t1 = i + j
```

```
sub f, t0, t1 # f = t0 - t1
```

The processor on your computer interprets these instructions explicitly, based on the underlying encoding of the assembly code to "machine instructions"

# Recognize either of these guys?



# Recognize either of these guys?

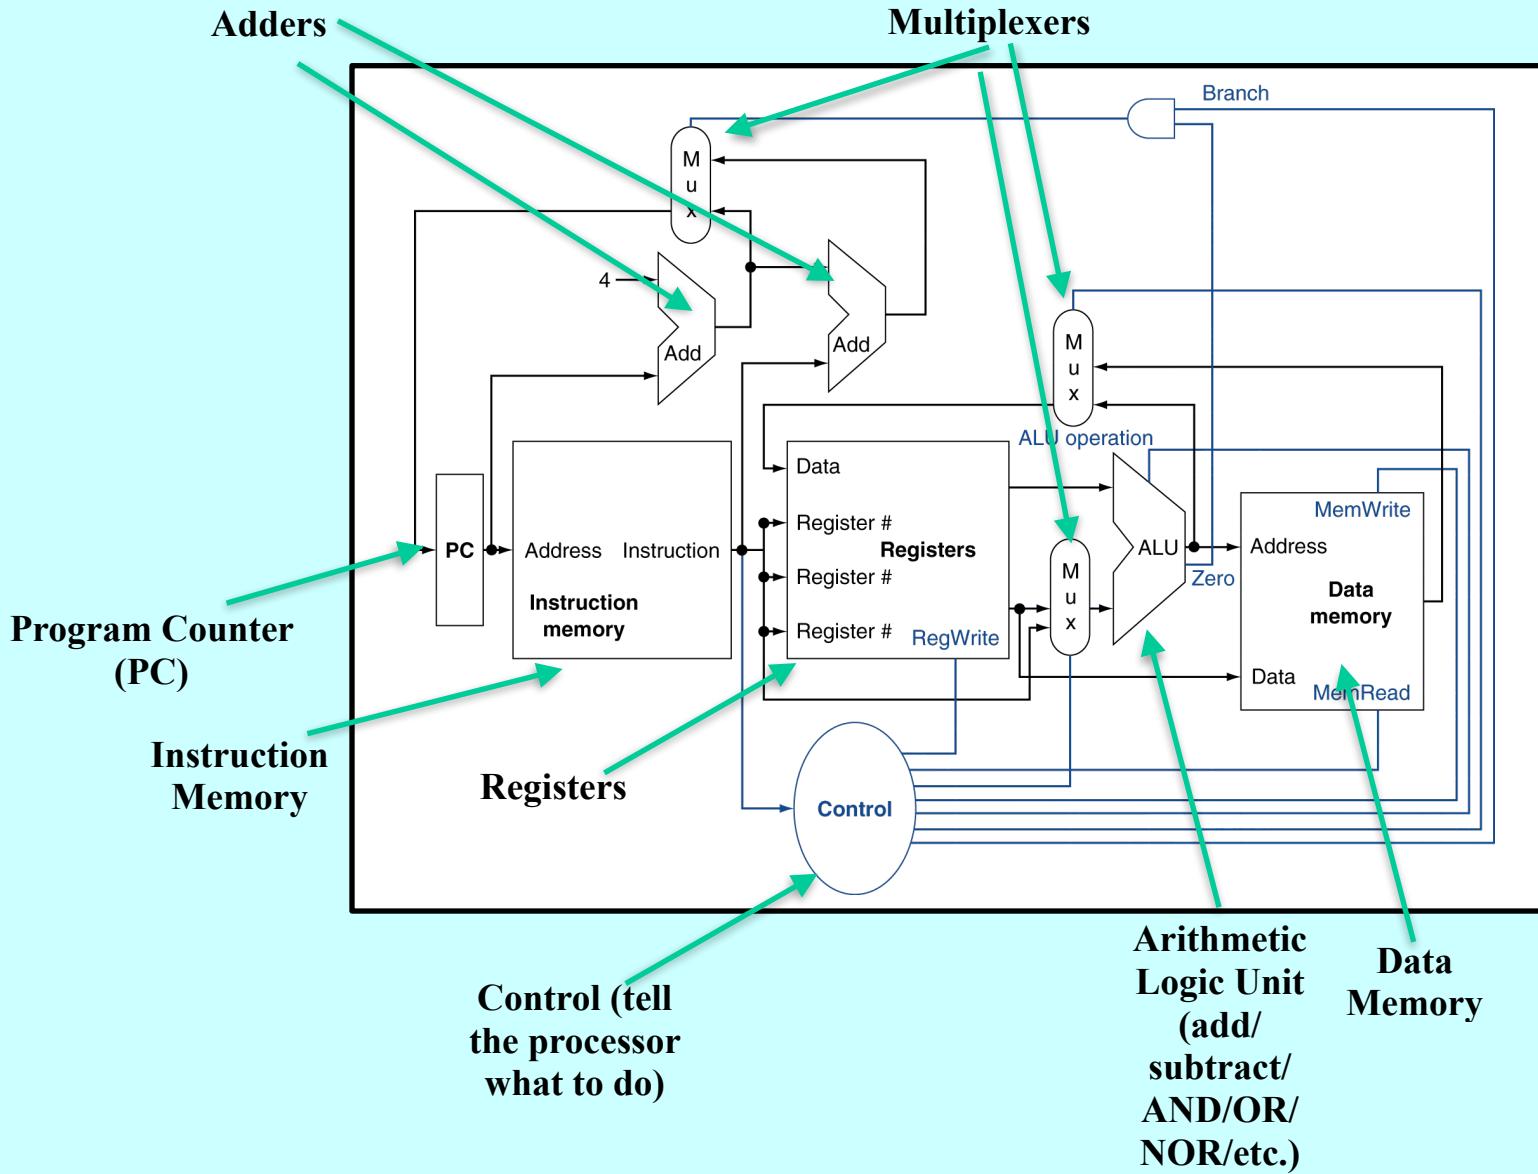
John Hennessy  
— created the MIPS processor, a contemporary RISC to the ARM processor.  
— Stanford Professor in CS and EE, and President of Stanford from 2000-2016.



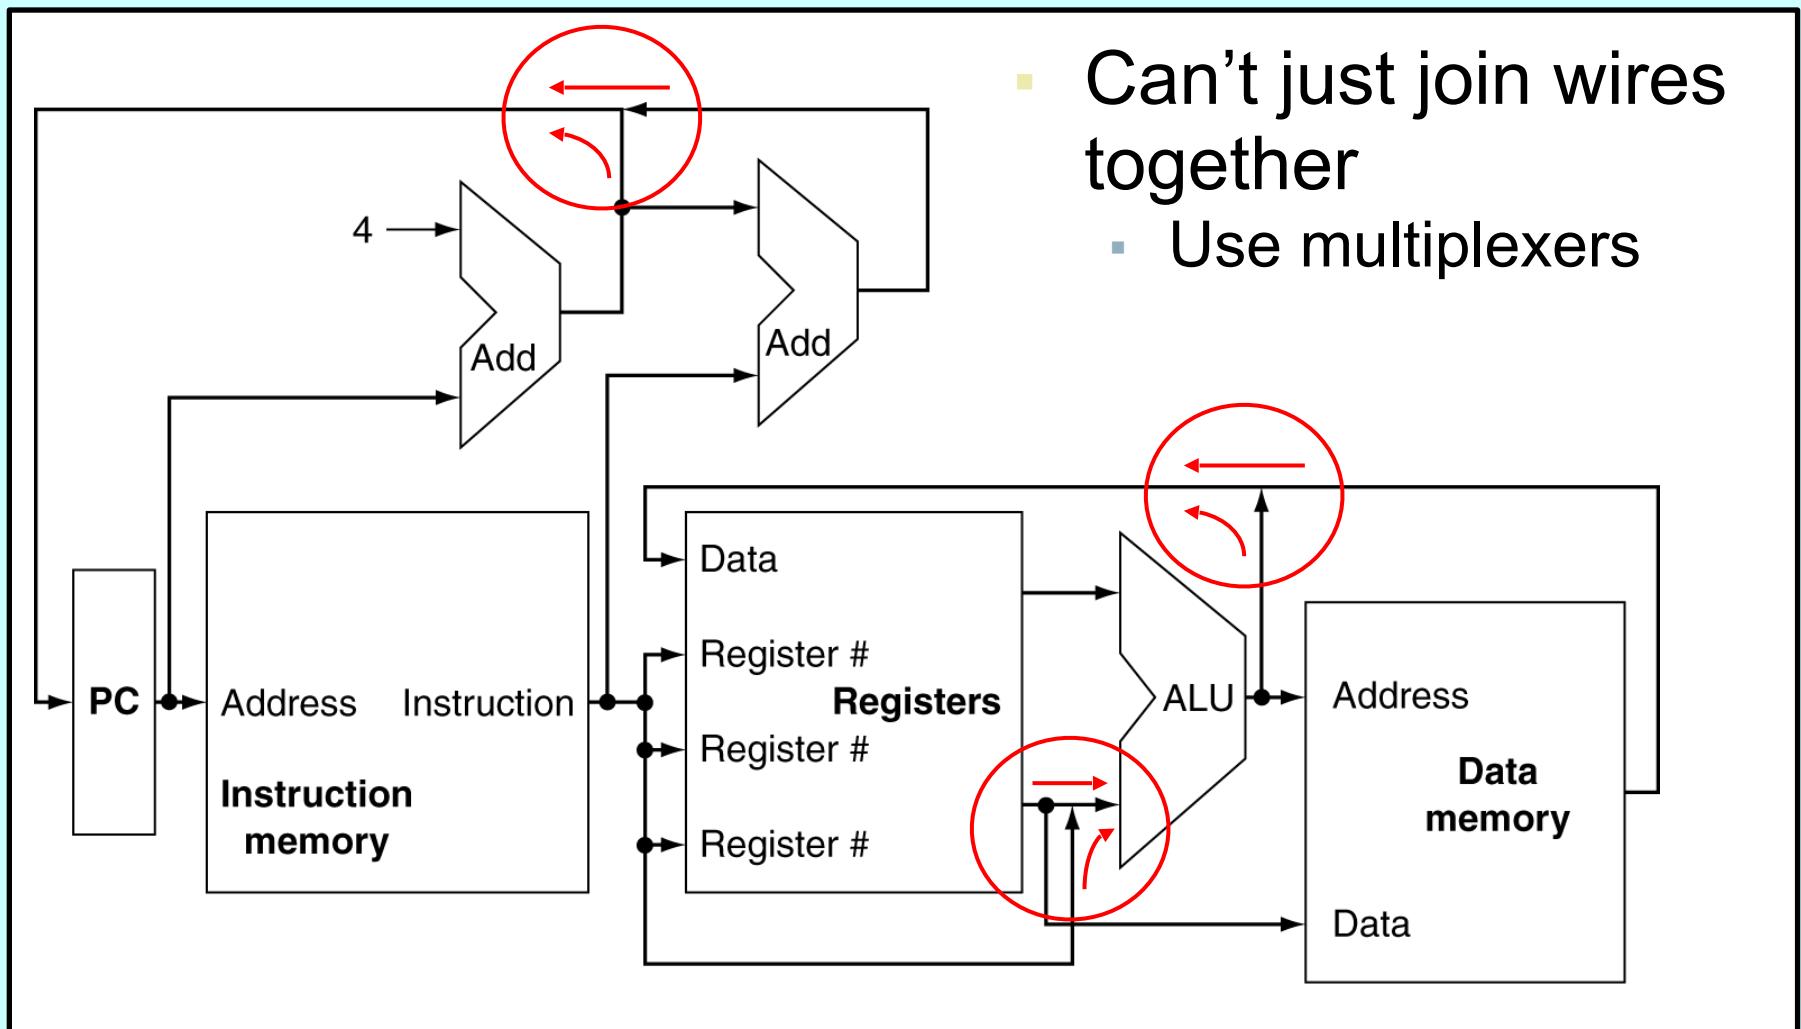
David Patterson  
— coined the term "RISC"  
— Berkeley Professor in EECS.

In 2017, Hennessy and Patterson were named the Turing Award recipients for their work in developing RISC processors. If it wasn't for these two, your phone would have a completely different processor!

# CPU Overview

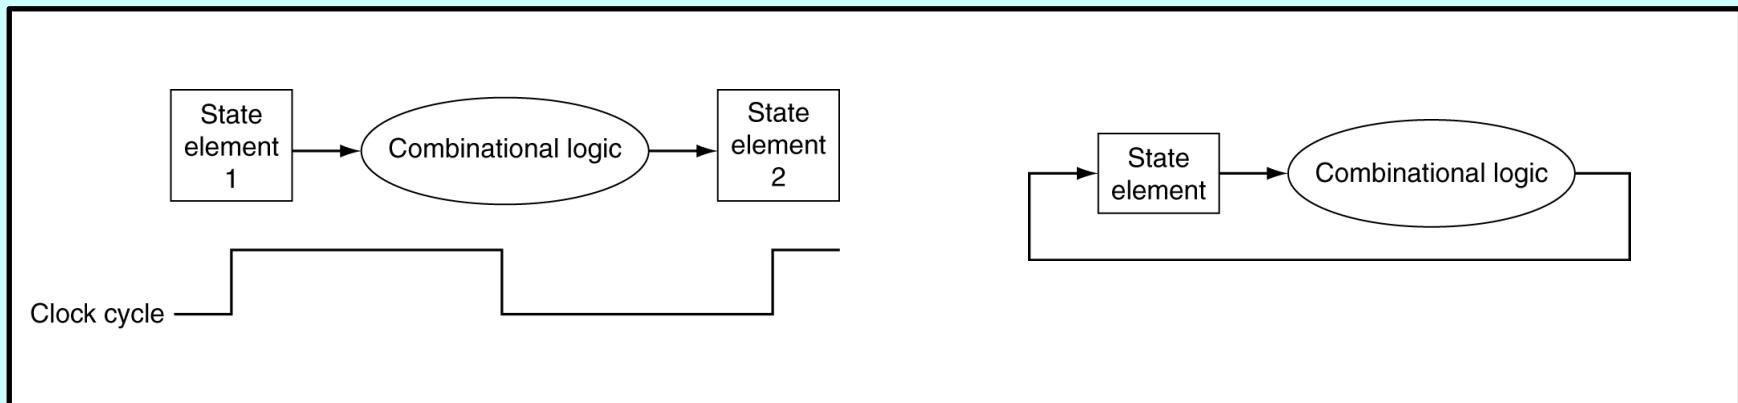


# CPU Overview

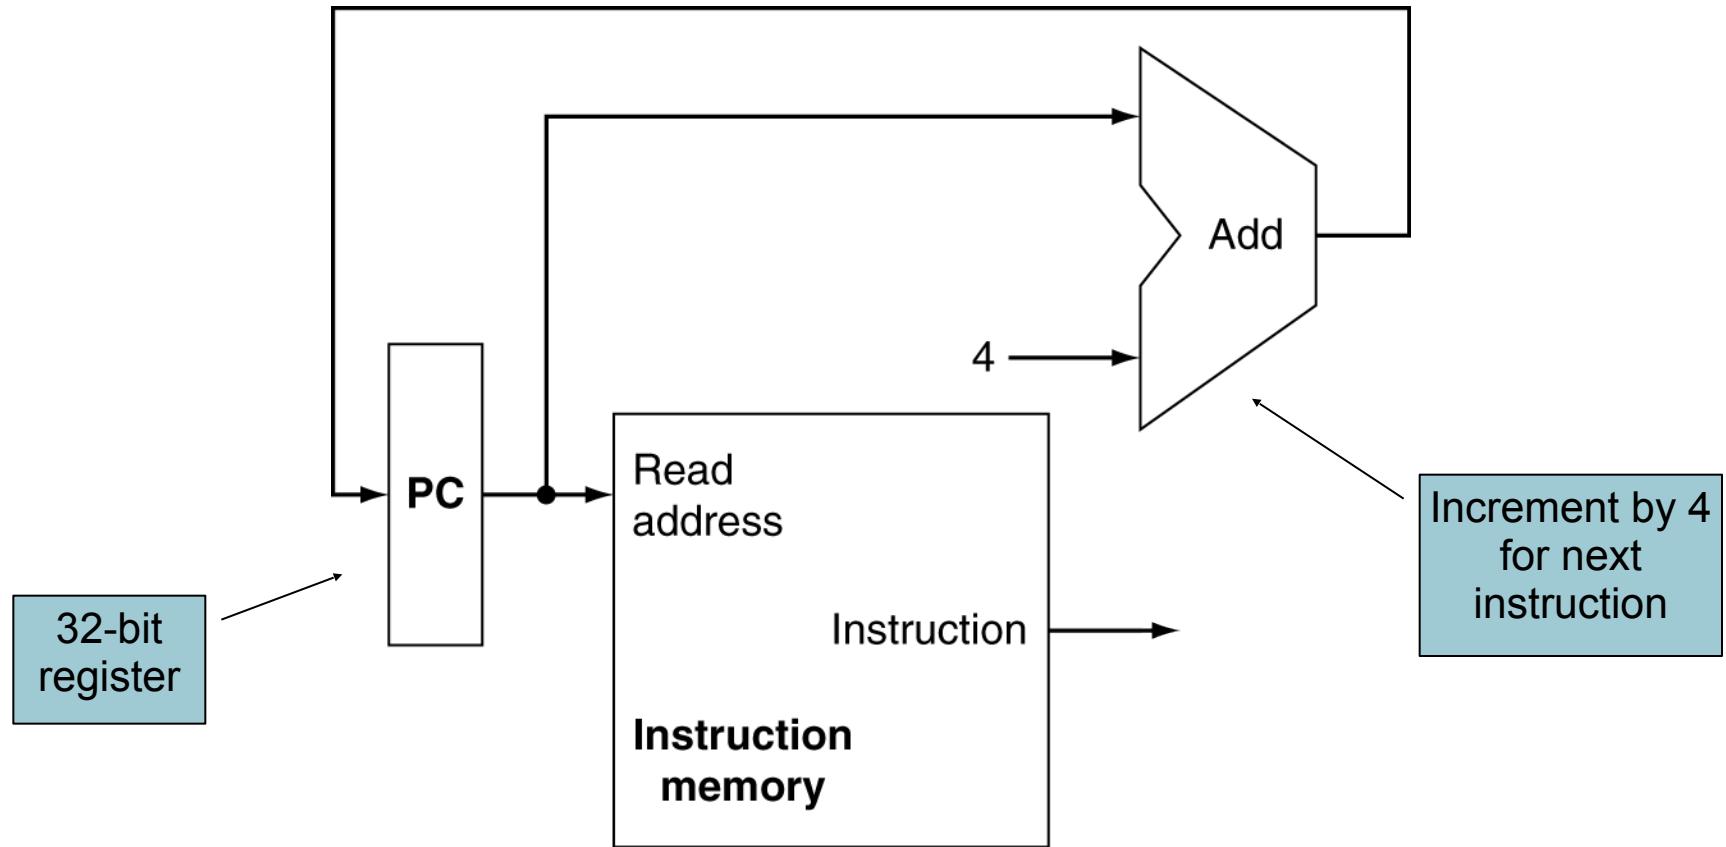


# Clocking Methodology

- Combinational logic transforms data during clock cycles
  - Between clock edges
  - Input from state elements, output to state element
  - Longest delay determines clock period



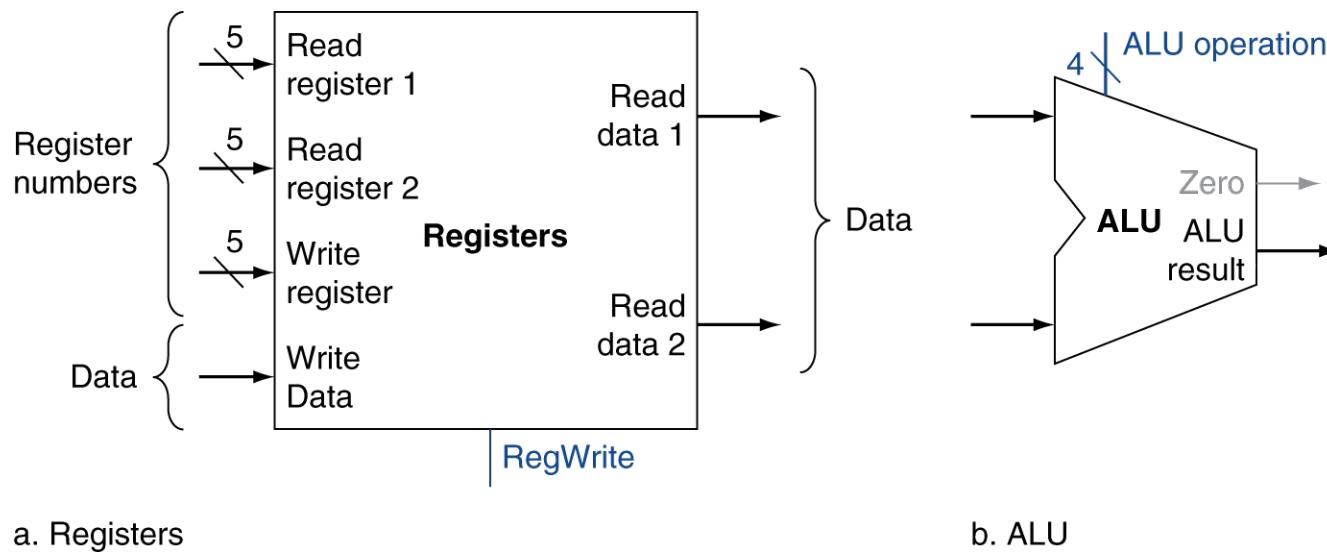
# Instruction Fetch



Instructions in MIPS are 4 bytes long, so each time the clock "ticks", we get the next instruction.

# R-Format Instructions

- Read two register operands
- Perform arithmetic/logical operation
- Write register result



- What does "R-Format" mean??

# Instruction Set

- The repertoire of instructions of a computer
- Different computers have different instruction sets
  - But with many aspects in common
- Early computers had very simple instruction sets
  - Simplified implementation
- Many modern computers also have simple instruction sets

# The MIPS Instruction Set

- Stanford MIPS commercialized by MIPS Technologies ([www.mips.com](http://www.mips.com))
- Large share of embedded core market
  - Applications in consumer electronics, network/storage equipment, cameras, printers, ...
- Typical of many modern ISAs
  - See MIPS Reference Data tear-out card, and Appendixes B and E

# MIPS R-format Instructions

| op     | rs     | rt     | rd     | shamt  | funct  |
|--------|--------|--------|--------|--------|--------|
| 6 bits | 5 bits | 5 bits | 5 bits | 5 bits | 6 bits |

- Instruction fields
  - op: operation code (opcode)
  - rs: first source register number
  - rt: second source register number
  - rd: destination register number
  - shamt: shift amount (00000 for now)
  - funct: function code (extends opcode)

# R-format Example

| op     | rs     | rt     | rd     | shamt  | funct  |
|--------|--------|--------|--------|--------|--------|
| 6 bits | 5 bits | 5 bits | 5 bits | 5 bits | 6 bits |

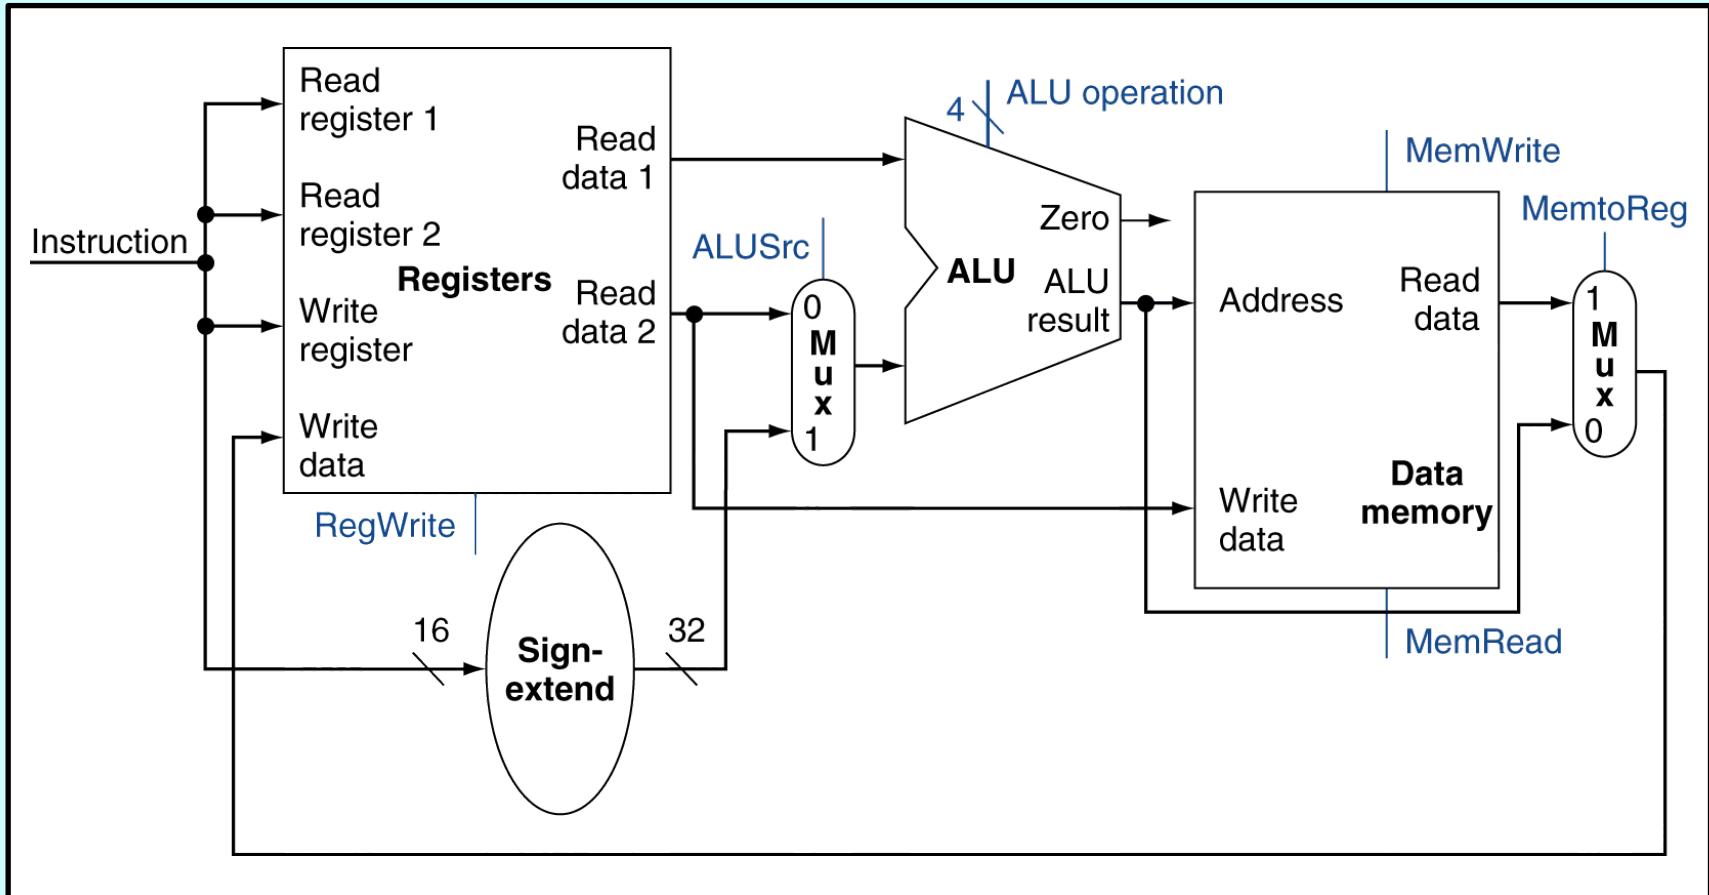
- add \$t0, \$s1, \$s2

This means: compute \$s1 + \$s2 and put the result into \$t0

|         |       |       |       |       |        |
|---------|-------|-------|-------|-------|--------|
| special | \$s1  | \$s2  | \$t0  | 0     | add    |
| 0       | 17    | 18    | 8     | 0     | 32     |
| 000000  | 10001 | 10010 | 01000 | 00000 | 100000 |

0000001000110010010000000100000<sub>2</sub> = 02324020<sub>16</sub>

# R-Type/Load/Store Datapath

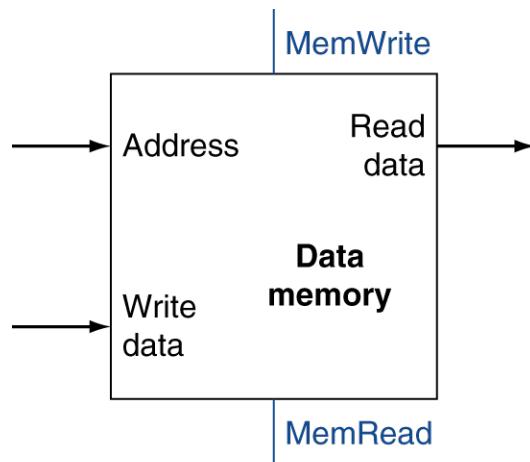


add \$t0, \$s1, \$s2

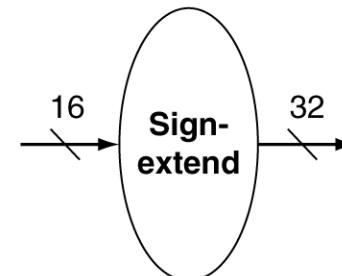
|         |       |       |       |       |        |
|---------|-------|-------|-------|-------|--------|
| special | \$s1  | \$s2  | \$t0  | 0     | add    |
| 000000  | 10001 | 10010 | 01000 | 00000 | 100000 |

# Load/Store Instructions

- Read register operands
- Calculate address using 16-bit offset
  - Use ALU, but sign-extend offset
- Load: Read memory and update register
- Store: Write register value to memory



a. Data memory unit



b. Sign extension unit

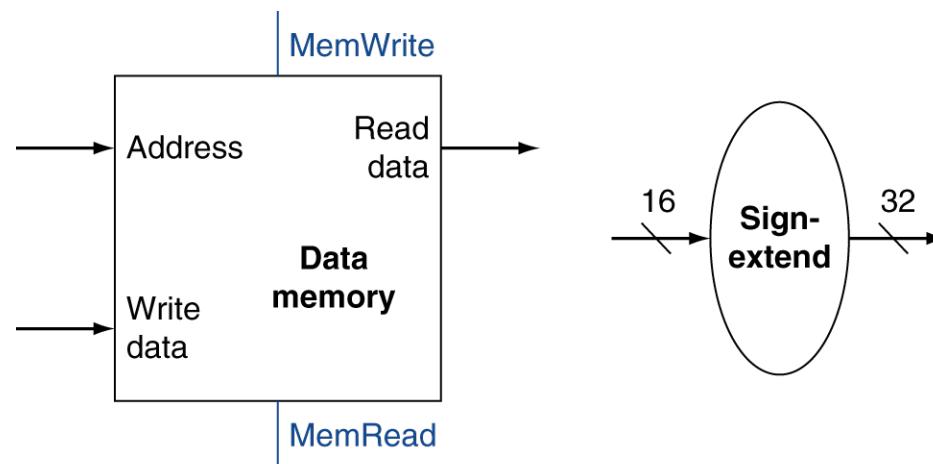
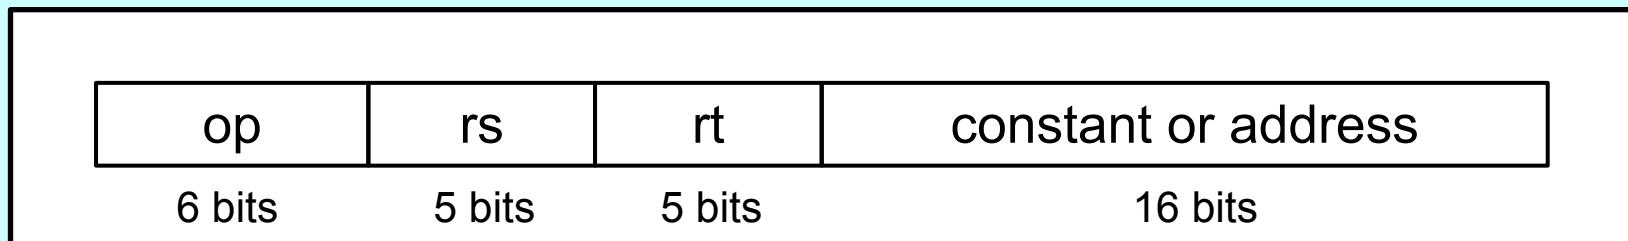
# MIPS I-format Instructions

| op     | rs     | rt     | constant or address |
|--------|--------|--------|---------------------|
| 6 bits | 5 bits | 5 bits | 16 bits             |

- Immediate arithmetic and load/store instructions
  - rt: destination or source register number
  - Constant:  $-2^{15}$  to  $+2^{15} - 1$
  - Address: offset added to base address in rs
-

# Load/Store Instructions

- Load: Read memory and update register
- Store: Write register value to memory



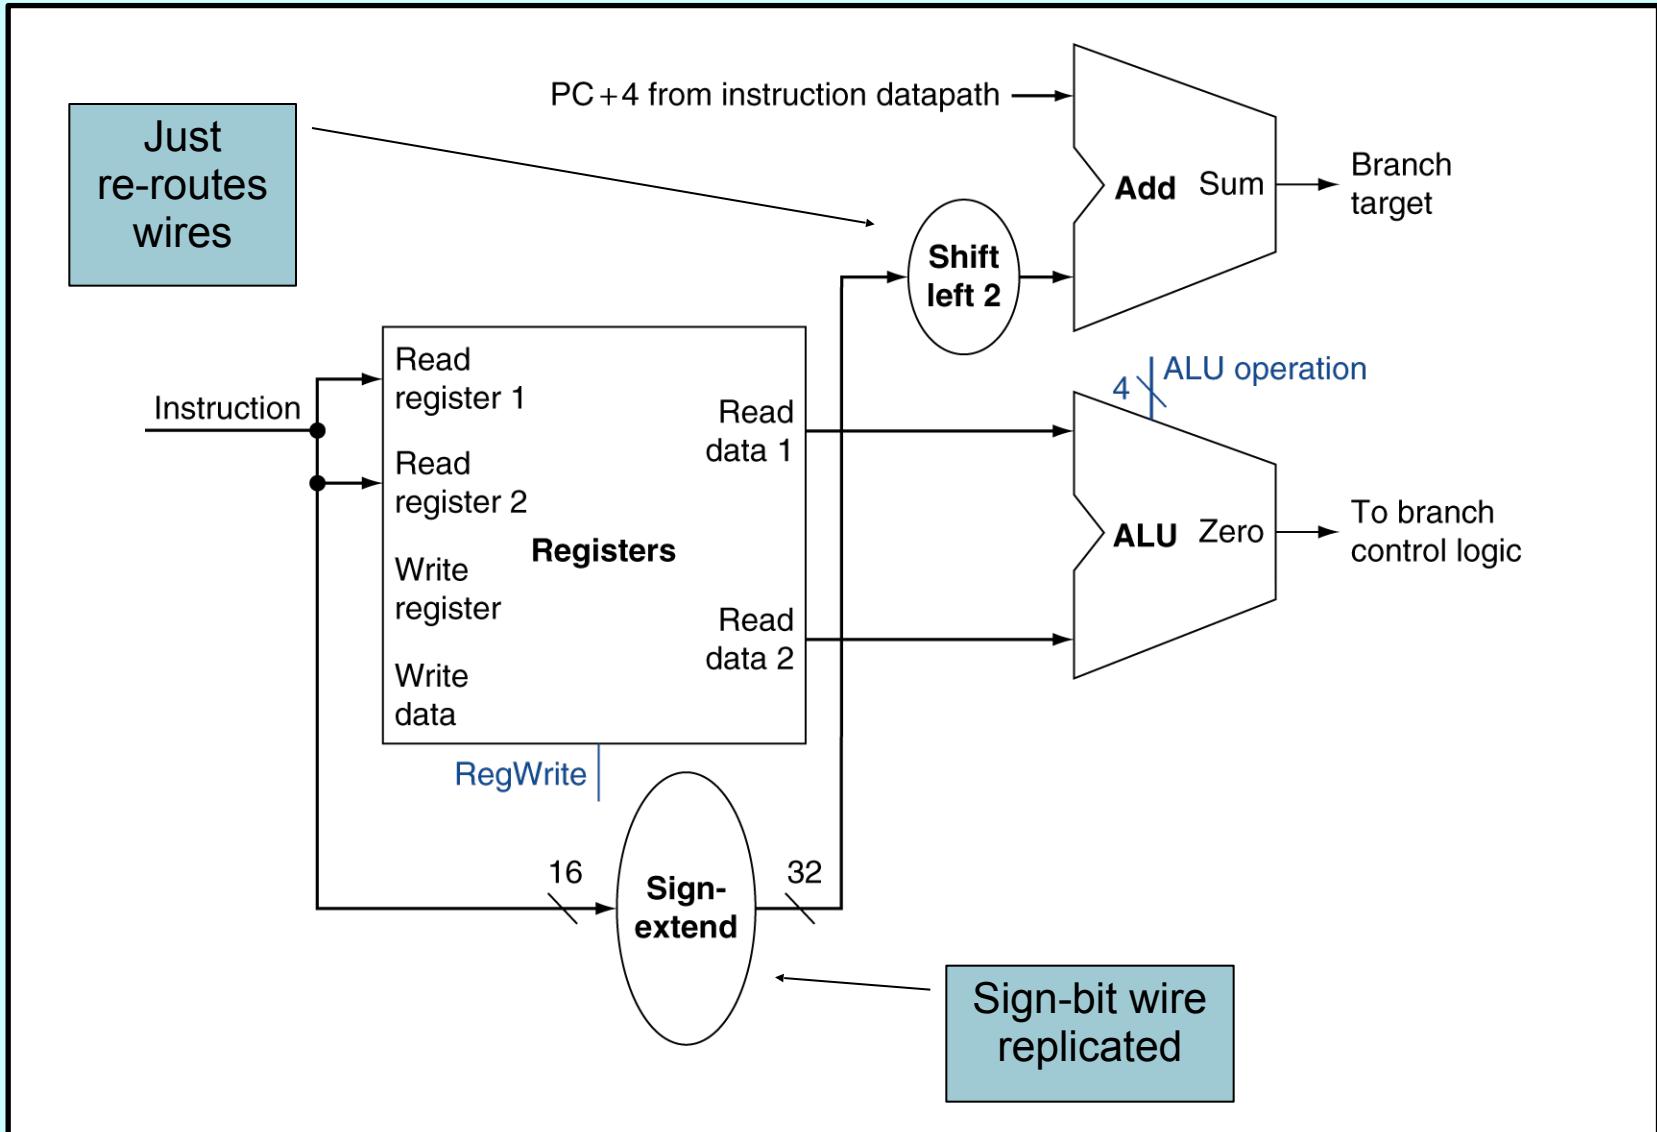
a. Data memory unit

b. Sign extension unit

# Branch Instructions

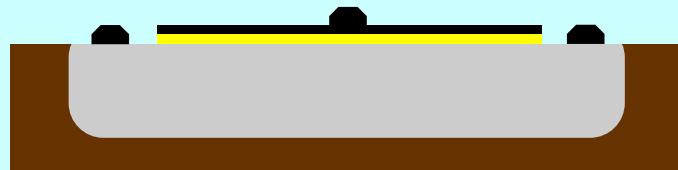
- Read register operands
- Compare operands
  - Use ALU, subtract and check Zero output
- Calculate target address
  - Sign-extend displacement
  - Shift left 2 places (word displacement)
  - Add to PC + 4
    - Already calculated by instruction fetch

# Branch Instructions



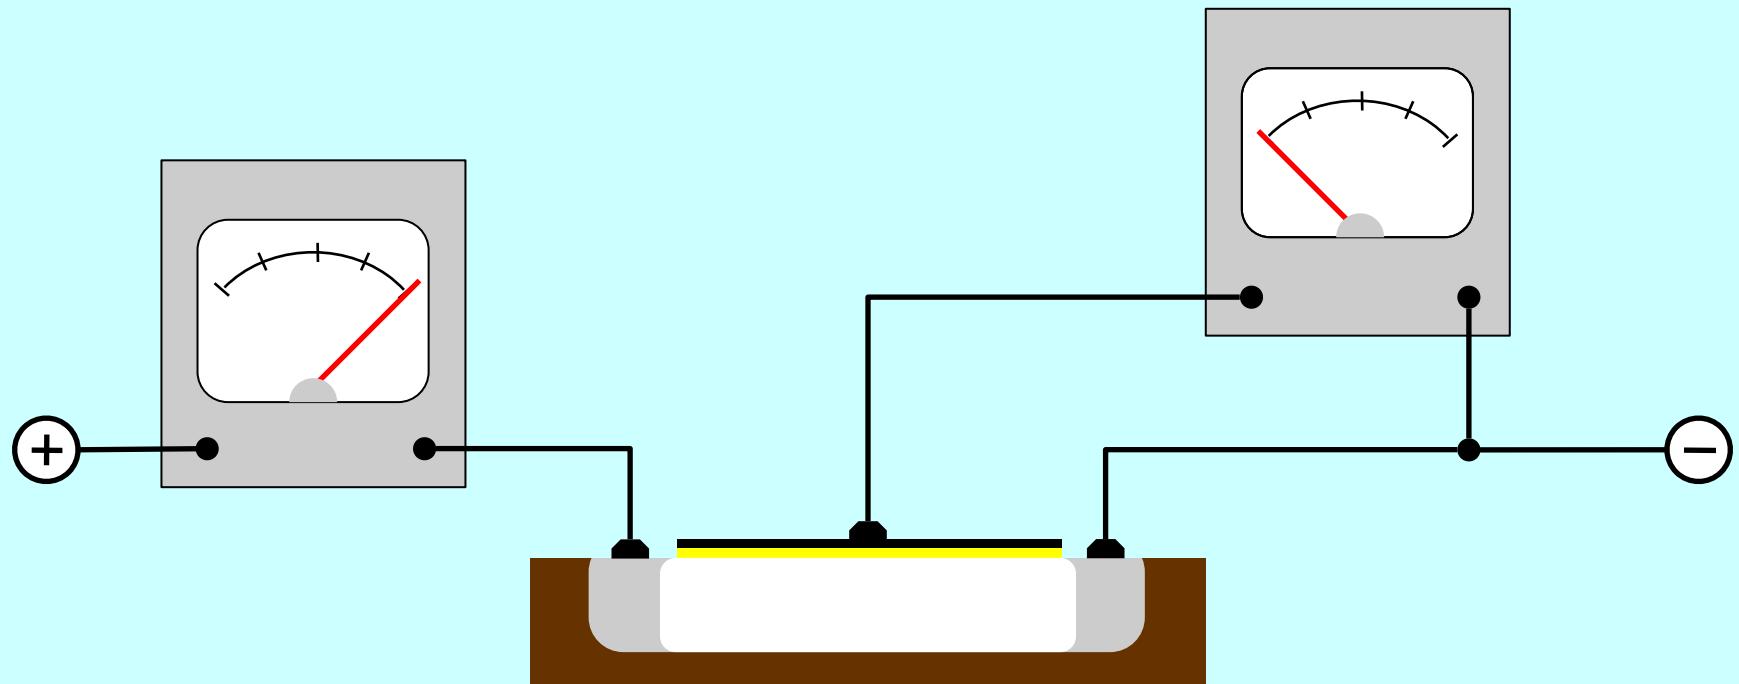
# Implementing Switches in Silicon

- In modern computers, switches are usually implemented in chips made based on a *semiconductor*, which is simply a material (typically some form of silicon) that is conducting in some circumstances and insulating in others.
- One of the most common components in today's computers is the *field-effect transistor* (or *FET*), which you can implement in semiconductor using a structure that looks like this:



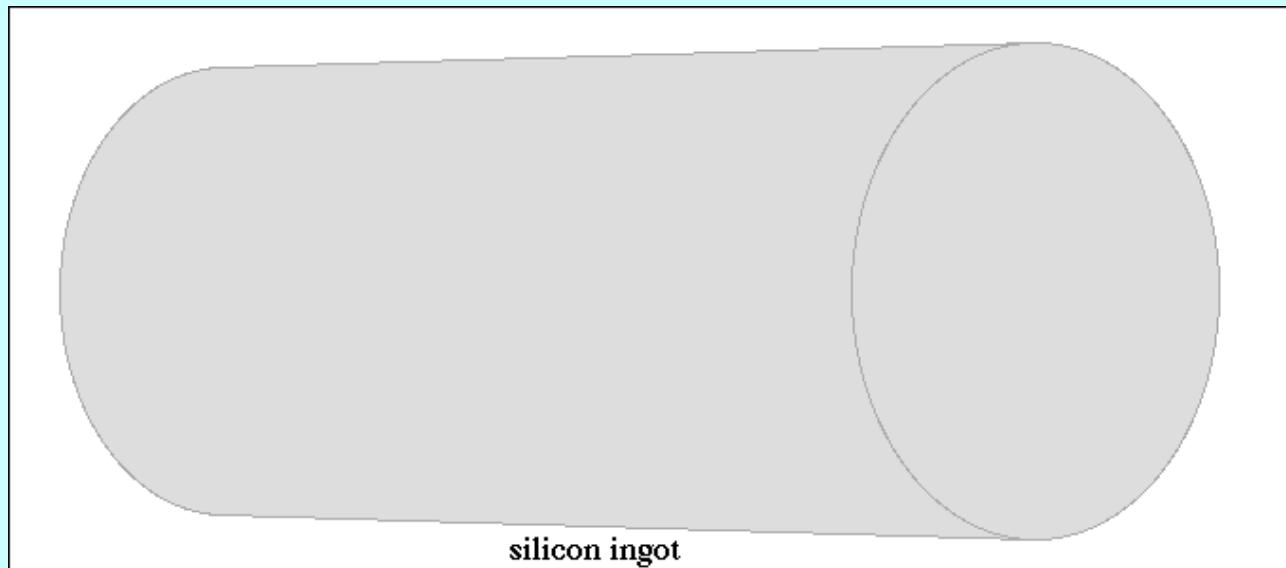
- Field-effect transistors come in two types. In a **n**-type FET, the presence of a charge on the gate allows current to flow across the transistor as a whole. In a **p**-type FET, a charge on the gate inhibits the current flow.

# Field-Effect Transistors



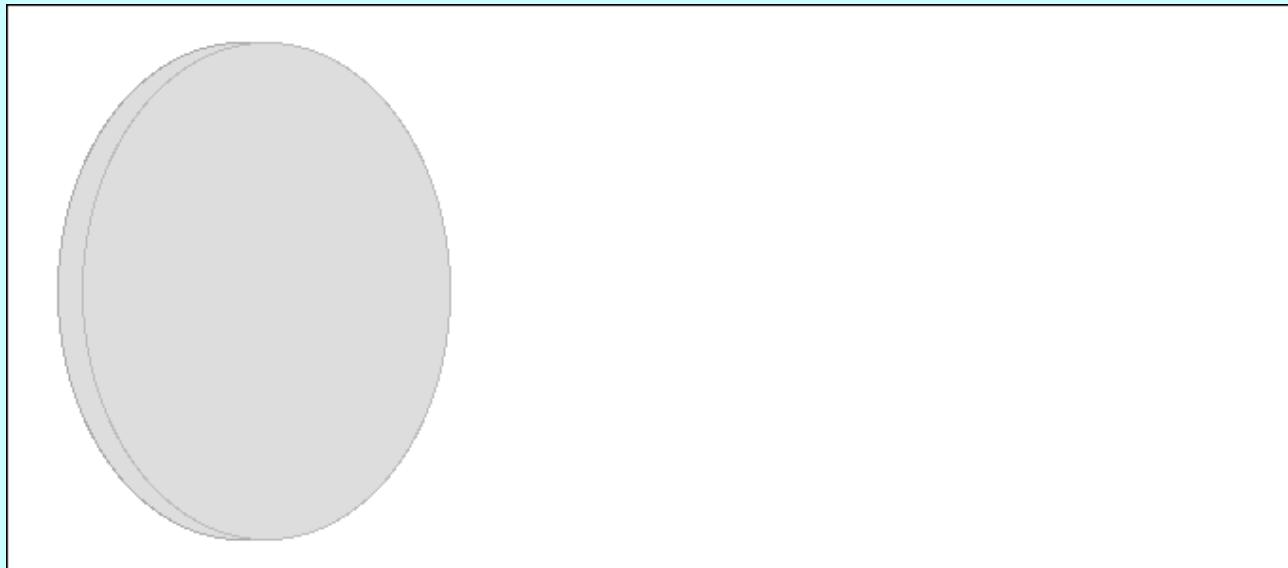
# Chip Fabrication

The manufacture of a semiconductor chip begins with a crystal of extremely pure silicon. The silicon itself is extracted from beach sand and then grown into a cylindrical ingot in such a way that impurities represent no more than one of a trillion atoms in the crystal.



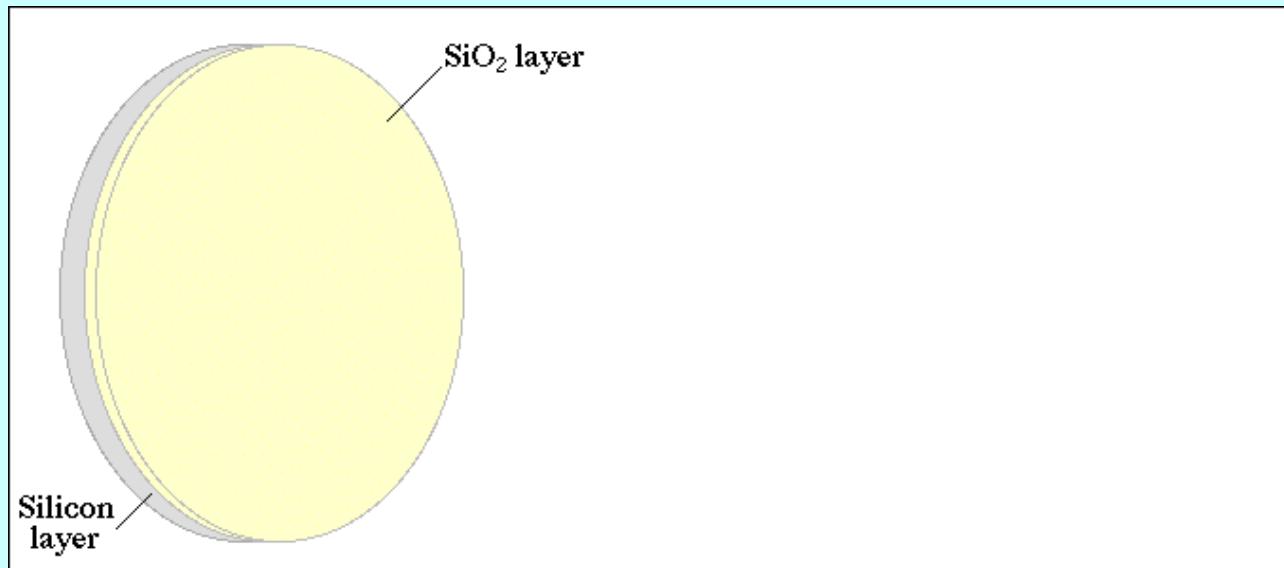
# Chip Fabrication

The silicon ingot is then sliced into individual wafers, which are typically 8-12 inches in diameter and less than a millimeter thick. Each individual wafer is then polished to remove any irregularities in its surface.



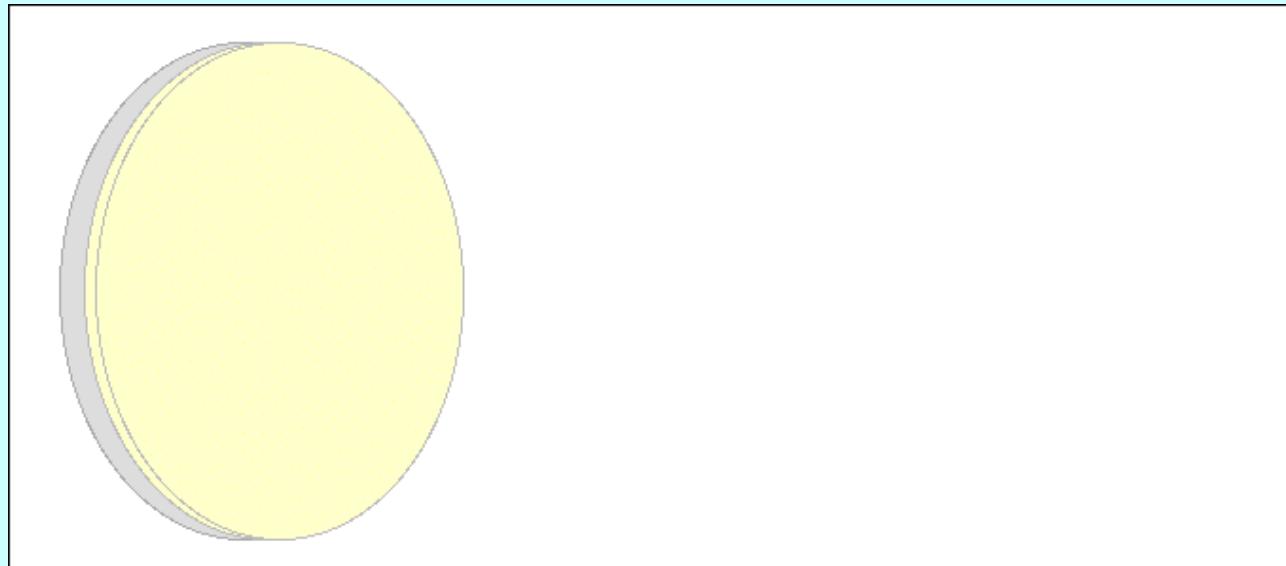
# Chip Fabrication

The polished wafer is then heated in the presence of oxygen to create a thin layer of silicon dioxide ( $\text{SiO}_2$ ). The  $\text{SiO}_2$  layer is nonconducting, and thereby serves to insulate the silicon wafer from the various materials that will be layered above it.



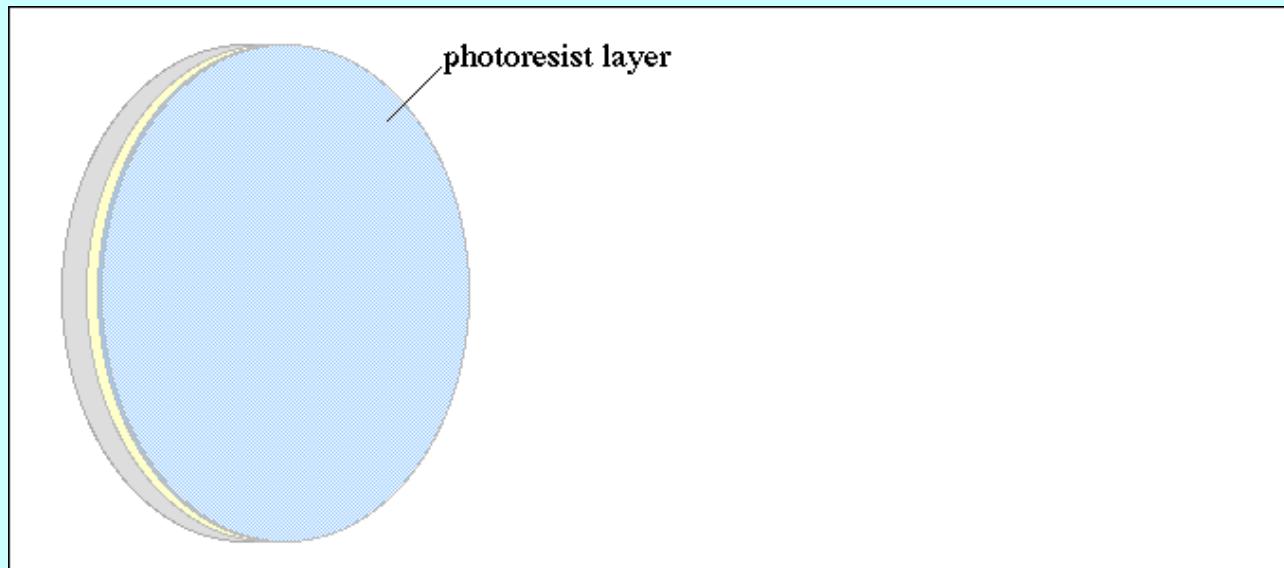
# Chip Fabrication

The problem now is to remove just the right parts of the  $\text{SiO}_2$  layer to expose the silicon underneath, so that these regions of the chip can be used as transistors. This phase of chip fabrication is accomplished using a process called *photolithography*.



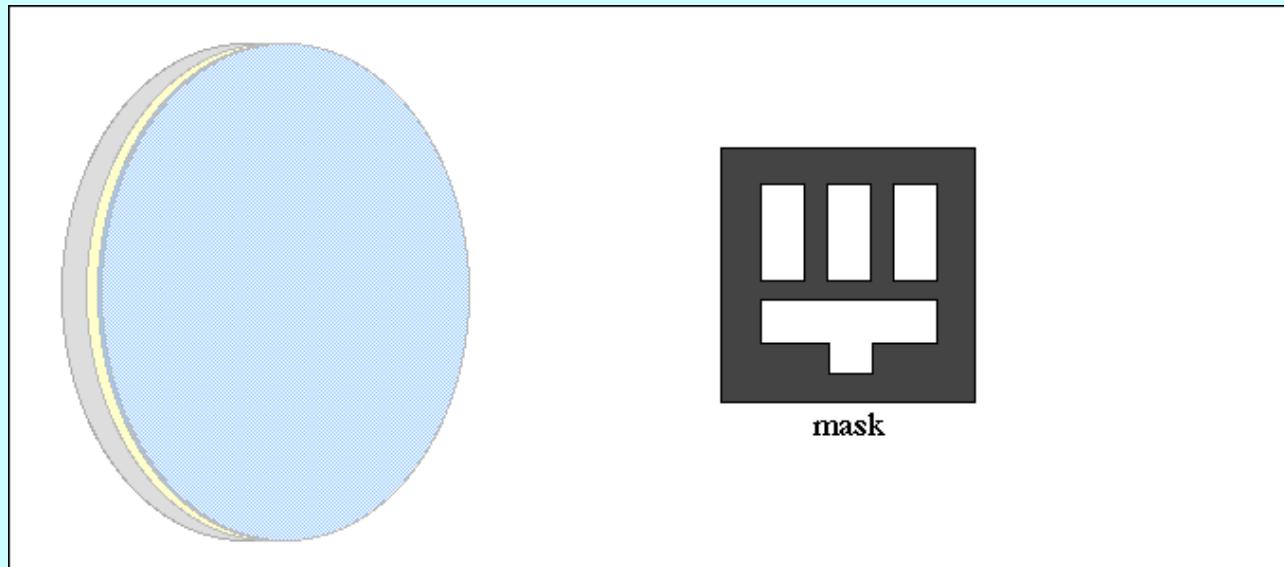
# Chip Fabrication

The first step in the photolithographic process is to coat the insulated wafer with a layer called *photoresist*, which changes its chemical composition when exposed to light.



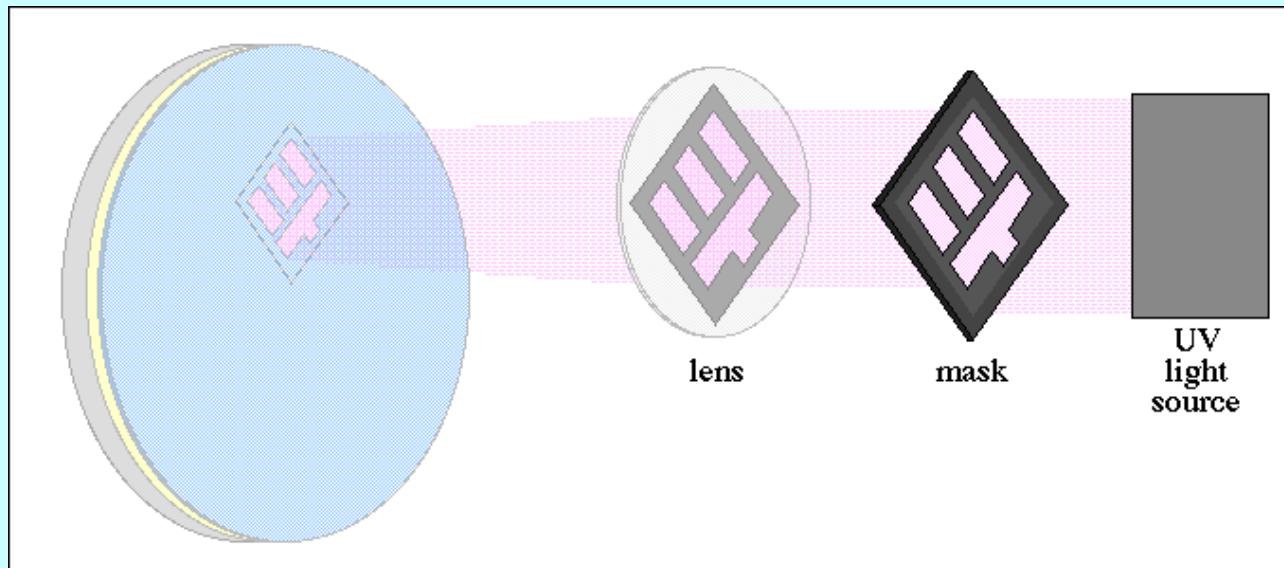
# Chip Fabrication

The pattern to be inscribed on the chip is taken from an optical template called a ***mask***. The mask looks very much like an old photographic negative, with the transparent regions indicating which parts of the chip surface need to be etched away.



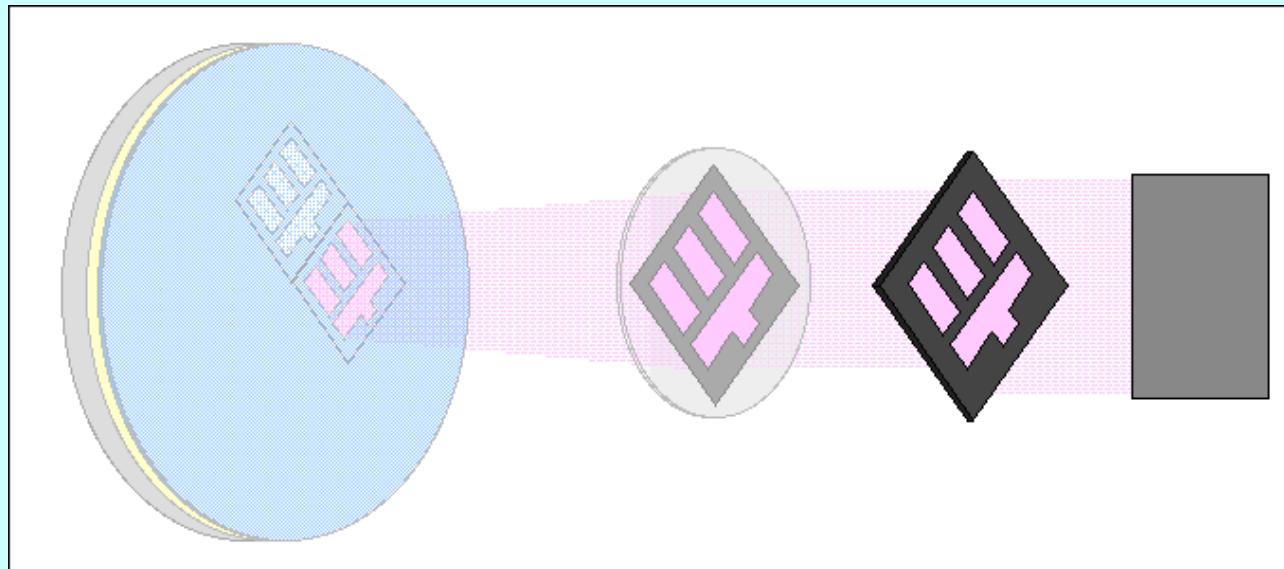
# Chip Fabrication

The image on the mask is transferred to the chip surface by shining ultraviolet light through the mask and then using a lens to focus the image on the desired area of the chip.



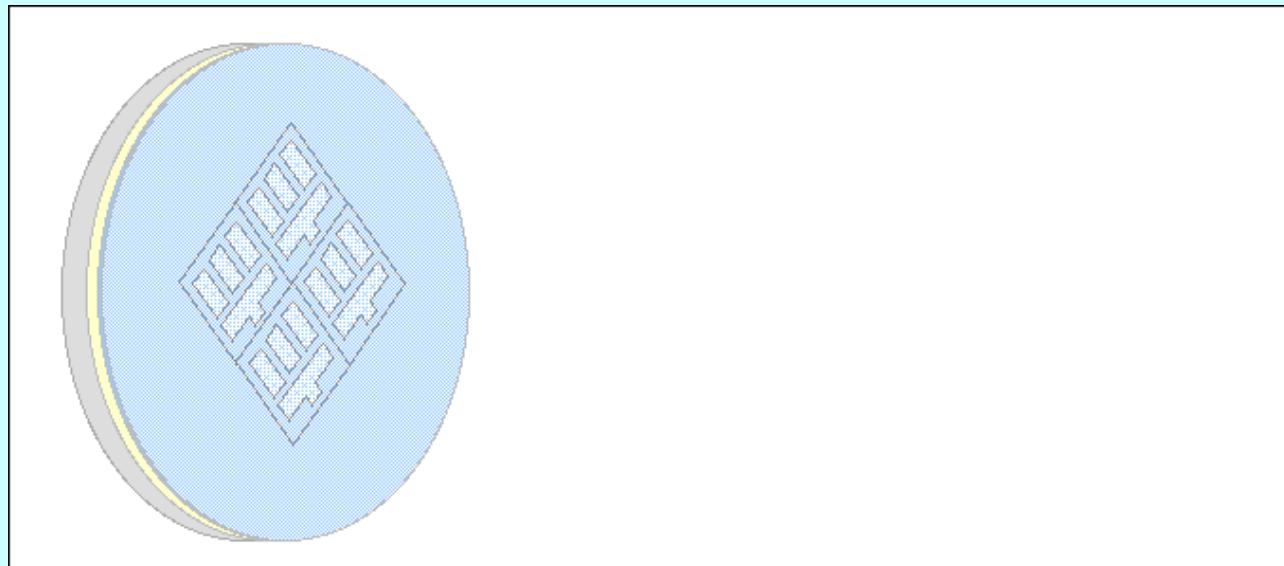
# Chip Fabrication

The same pattern is then created in a new position in other parts of the wafer by repositioning the image.



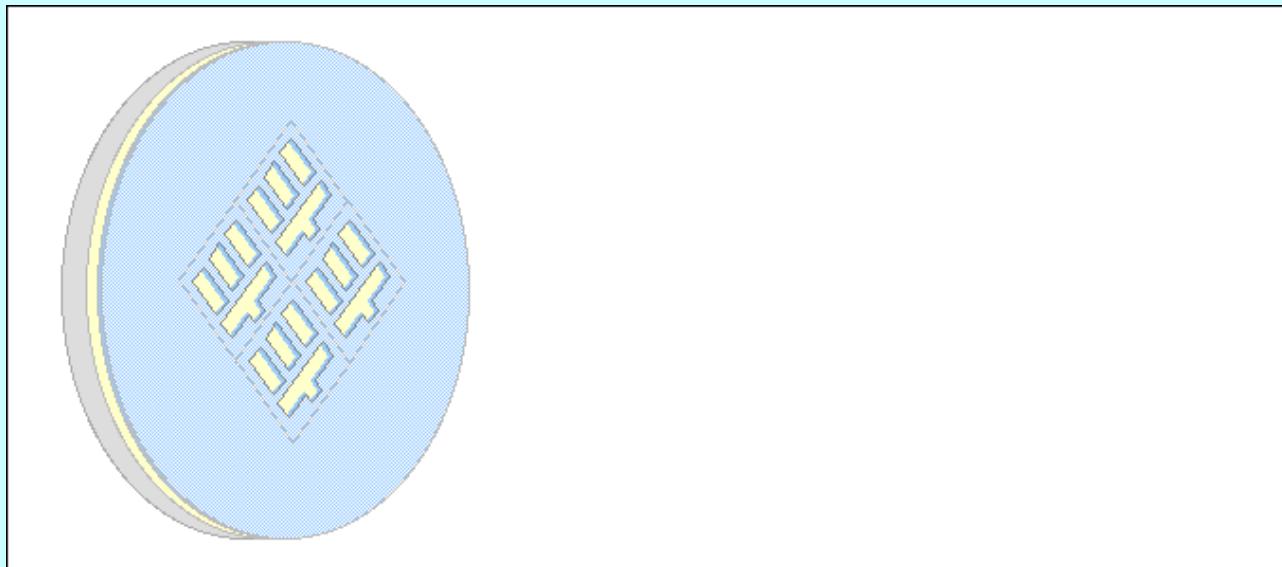
# Chip Fabrication

The same pattern is then created in a new position in other parts of the wafer by repositioning the image. A typical wafer can hold thousands of identical chips, each of which can be etched into the wafer simultaneously.



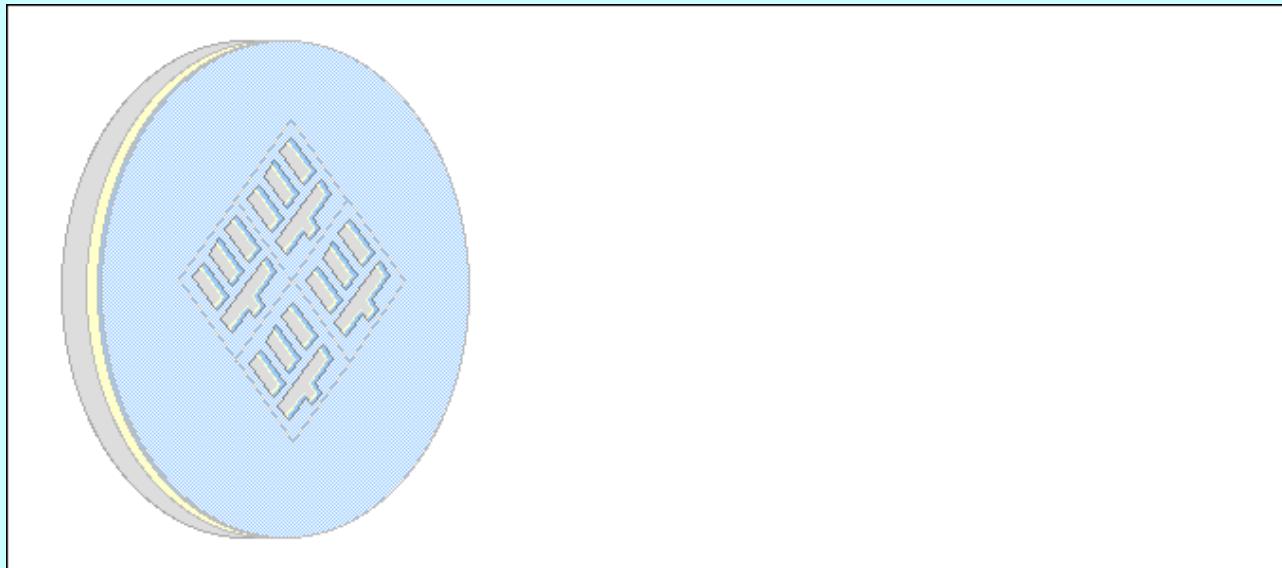
# Chip Fabrication

The exposed photoresist is then removed chemically, exposing the  $\text{SiO}_2$  layer underneath.



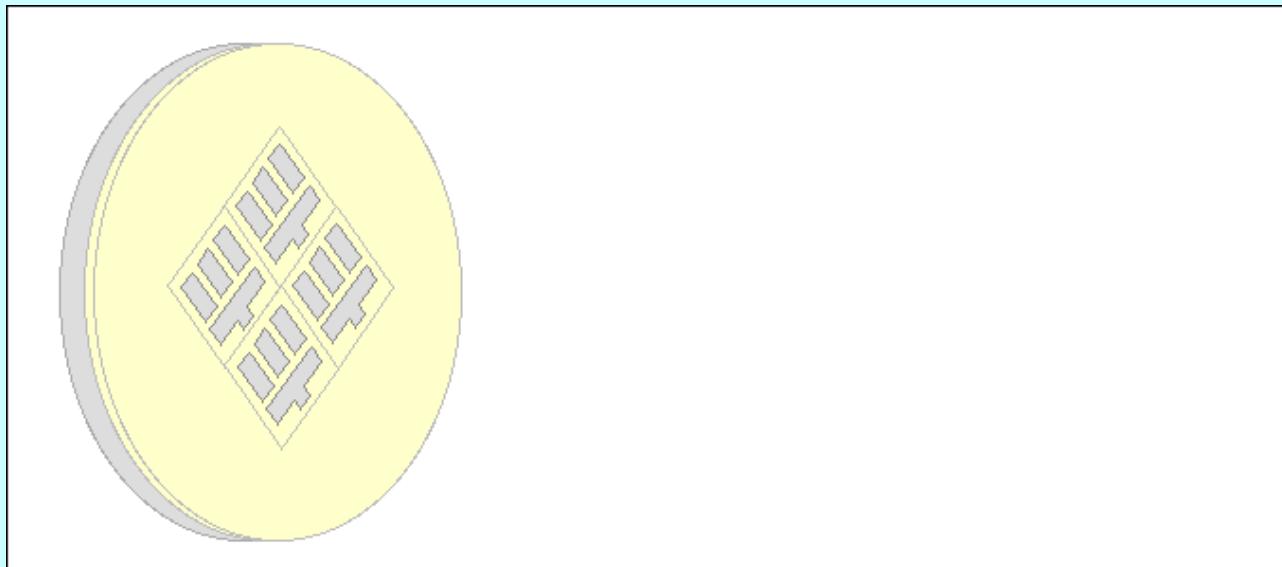
# Chip Fabrication

The exposed photoresist is then removed chemically, exposing the  $\text{SiO}_2$  layer underneath. A further etching step removes the  $\text{SiO}_2$  in the exposed areas, revealing the silicon layer.



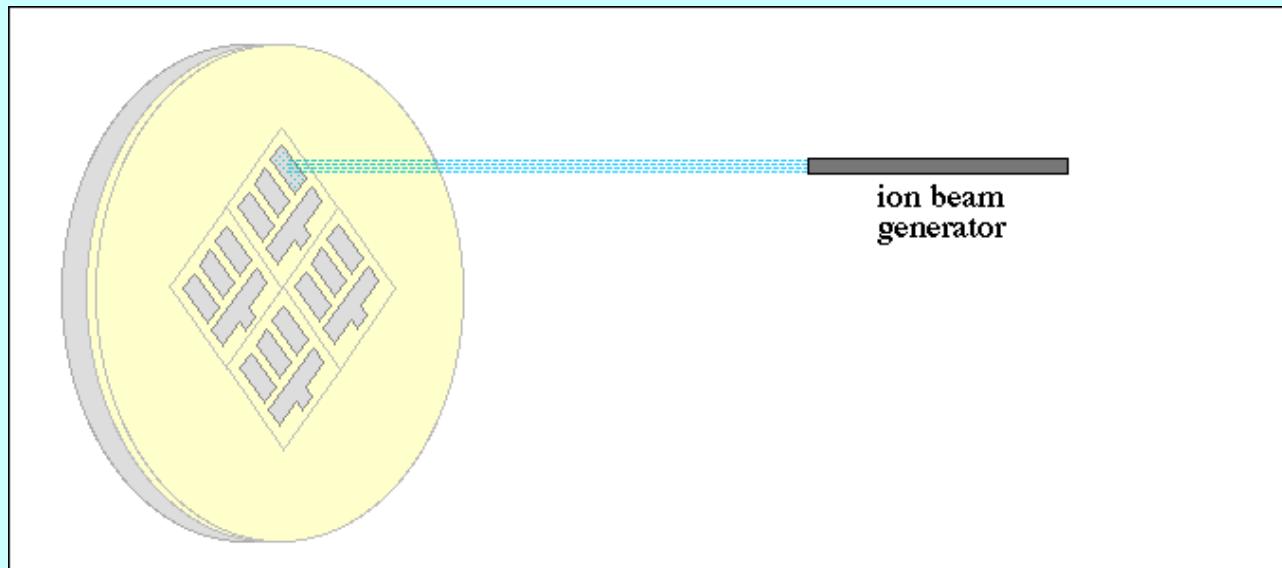
# Chip Fabrication

The exposed photoresist is then removed chemically, exposing the  $\text{SiO}_2$  layer underneath. A further etching step removes the  $\text{SiO}_2$  in the exposed areas, revealing the silicon layer. The final step in the etching process is to remove the unexposed photoresist from the wafer.



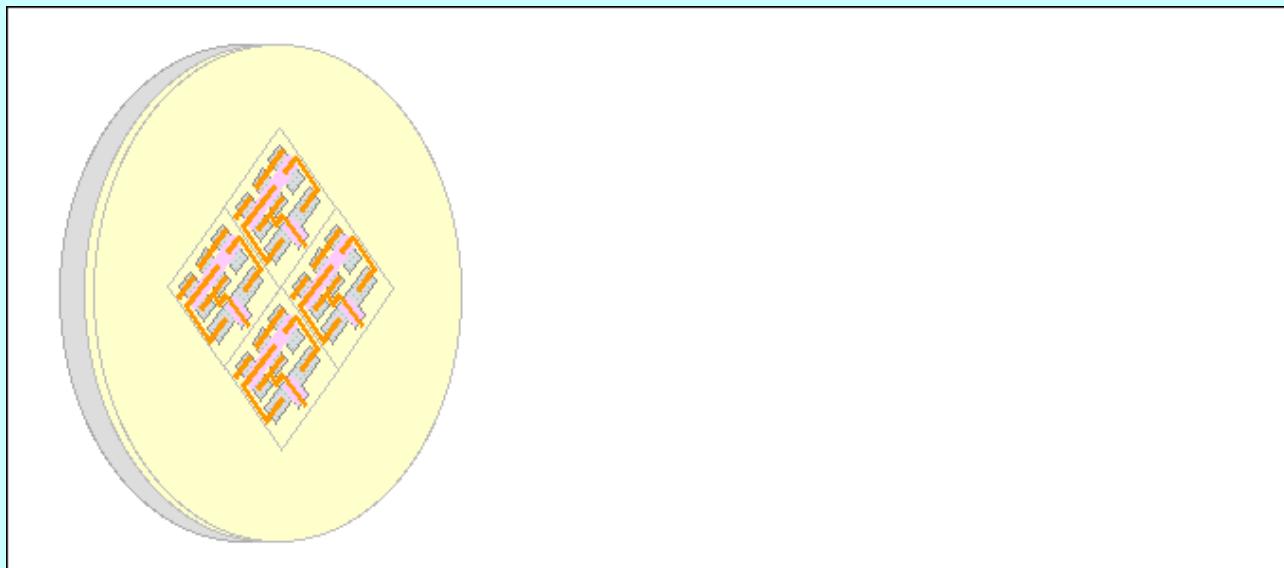
# Chip Fabrication

The next phase in the chip-making process is to introduce new elements into the silicon wafer to make it semiconducting. An ion beam bombards the surface with the doping atoms needed to create **n**-type and **p**-type regions in the exposed silicon wafer.



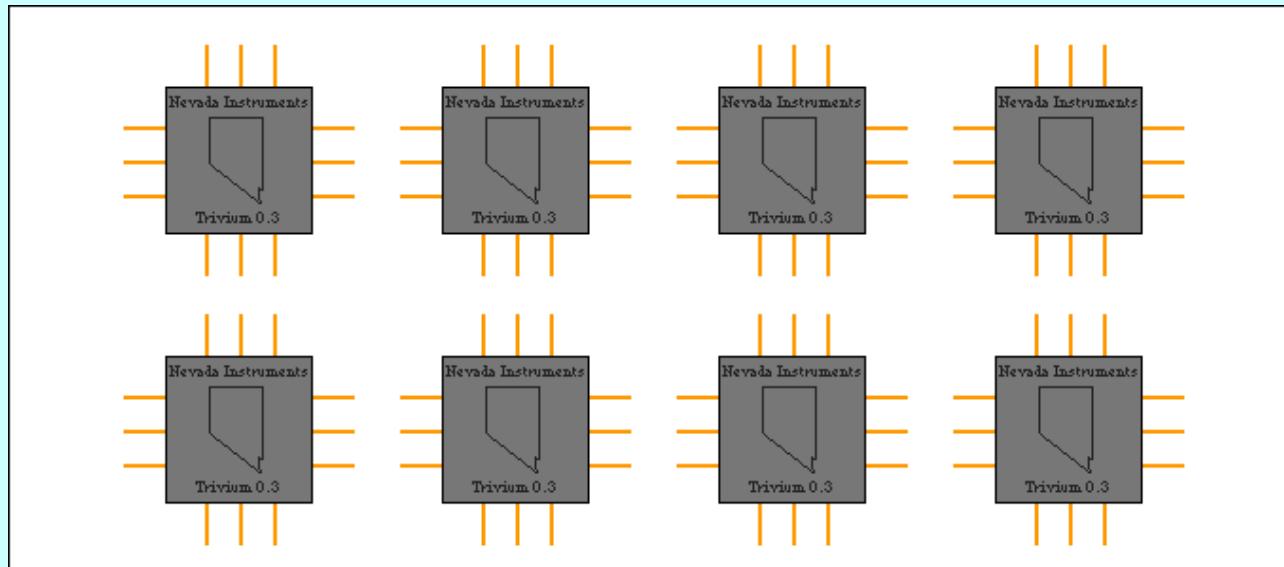
# Chip Fabrication

To complete the fabrication of the chips on the wafer, additional photolithographic steps are used to lay down alternating layers of insulators and conductors. The conductors (typically aluminum and tungsten) create the wiring pattern for each chip.



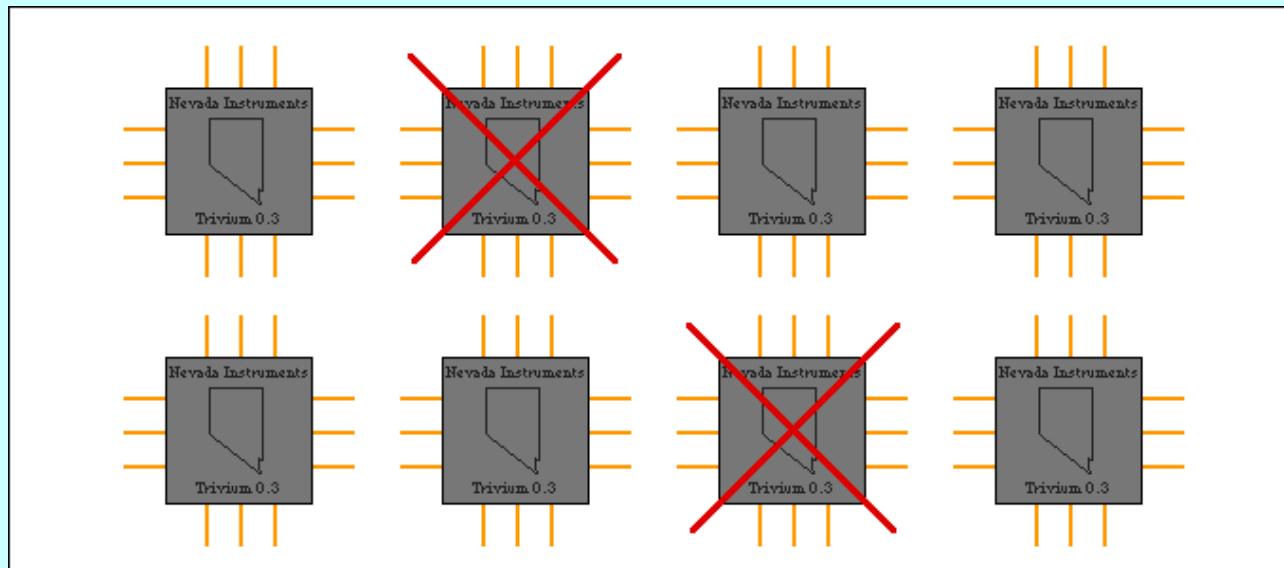
# Chip Fabrication

Once the wiring is complete, diamond saws cut the wafer into the individual chips, which are then encased in plastic, after connecting metallic leads to the appropriate contact positions on the chip.



# Chip Fabrication

The individual chips are then tested to make sure that every internal circuit works correctly. Given the complexity and extremely fine tolerances of the manufacturing process, several of the chips on each wafer typically need to be discarded.



<https://www.youtube.com/watch?v=bor0qLifjz4>

The End