

# Roadmap

C:

```
car *c = malloc(sizeof(car));
c->miles = 100;
c->gals = 17;
float mpg = get_mpg(c);
free(c);
```

Java:

```
Car c = new Car();
c.setMiles(100);
c.setGals(17);
float mpg =
    c.getMPG();
```

Memory & data  
Arrays & structs  
Integers & floats  
RISC V assembly  
Procedures & stacks  
Executables  
**Memory** & caches  
Processor Pipeline  
Performance  
Parallelism

Assembly  
language:

```
get_mpg(car*):
    lw    a5,0(a0)
    lw    a4,4(a0)
    divw  a5,a5,a4
    fcvt.s.w    fa0,a5
    ret
```



Machine  
code:

```
0111010000011000
100011010000010000000010
1000100111000010
110000011111101000011111
```

OS:



Computer  
system:



# Virtual Memory (VM\*)


- ❖ Overview and motivation
- ❖ VM as a tool for caching
- ❖ Address translation
- ❖ VM as a tool for memory management
- ❖ VM as a tool for memory protection

**Warning:** Virtual memory is pretty complex, but crucial for understanding how processes work and for debugging performance

*\*Not to be confused with “Virtual Machine” which is a whole other thing.*

# Memory as we know it so far... is *virtual*!

- ❖ Programs refer to virtual memory addresses
  - `movq (%rdi), %rax`
  - Conceptually memory is just a very large array of bytes
  - System provides private address space to each process
- ❖ Allocation: Compiler and run-time system
  - Where different program objects should be stored
  - All allocation within single virtual address space
- ❖ But...
  - We *probably* don't have  $2^w$  bytes of physical memory
  - We *certainly* don't have  $2^w$  bytes of physical memory **for every process**
  - Processes should not interfere with one another
    - Except in certain cases where they want to share code or data



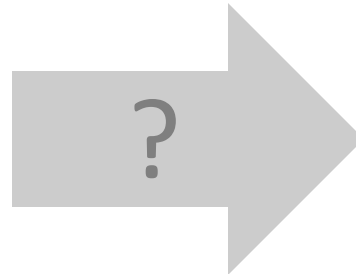
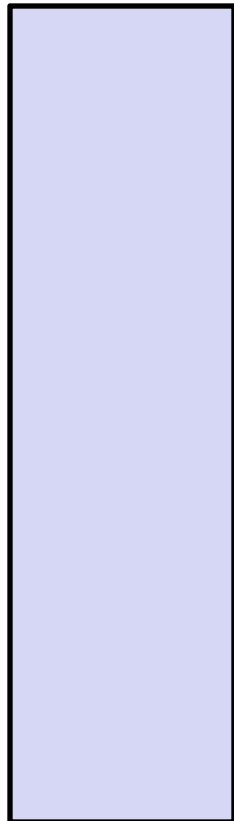
# Problem 1: How Does Everything Fit?

64-bit virtual addresses can address  
several exabytes  
(18,446,744,073,709,551,616 bytes)

16 EiB

Physical main memory offers  
a few gigabytes  
(e.g. 8,589,934,592 bytes)

8 GiB



*(Not to scale; physical memory would be smaller than the period at the end of this sentence compared to the virtual address space.)*

smaller than this!

1 virtual address space per process,  
with many processes...

# Problem 2: Memory Management

We have multiple processes:

Process 1  
Process 2  
Process 3  
...  
Process n

**X**

Each process has...

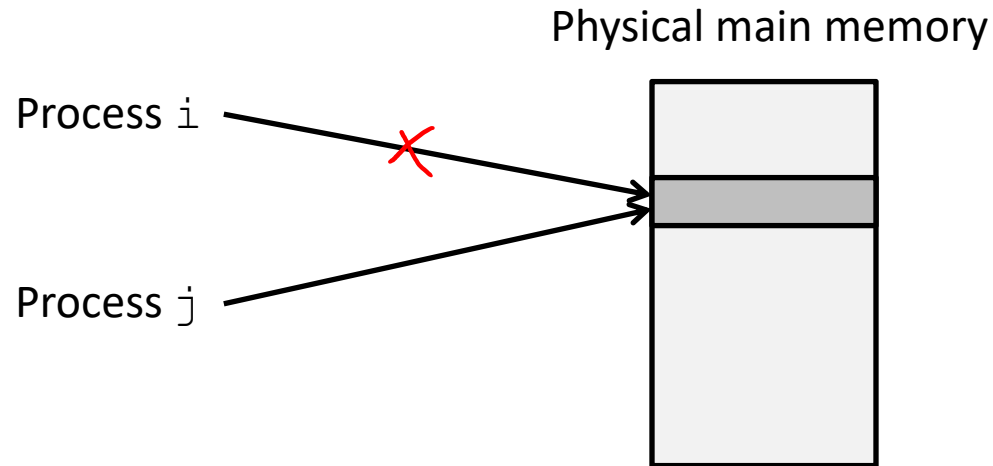
stack  
heap  
.text  
.data

*What goes where?*

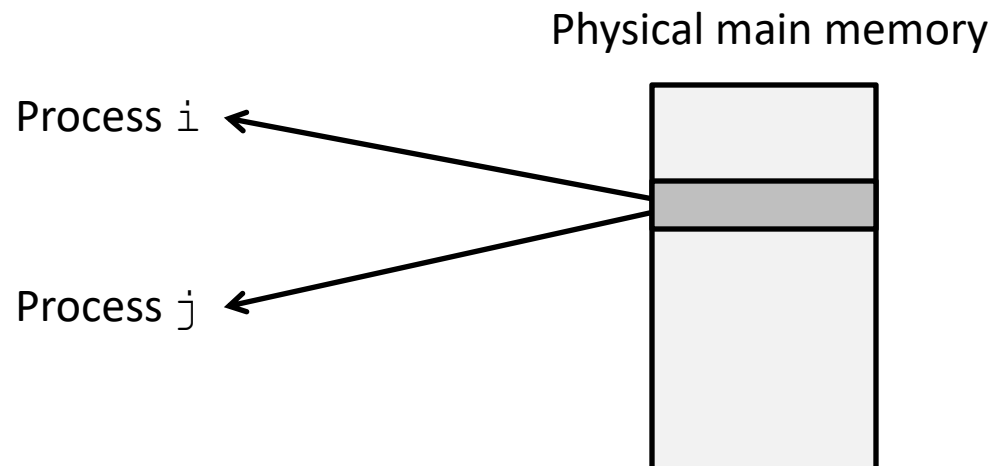
Physical main memory



## Problem 3: How To Protect



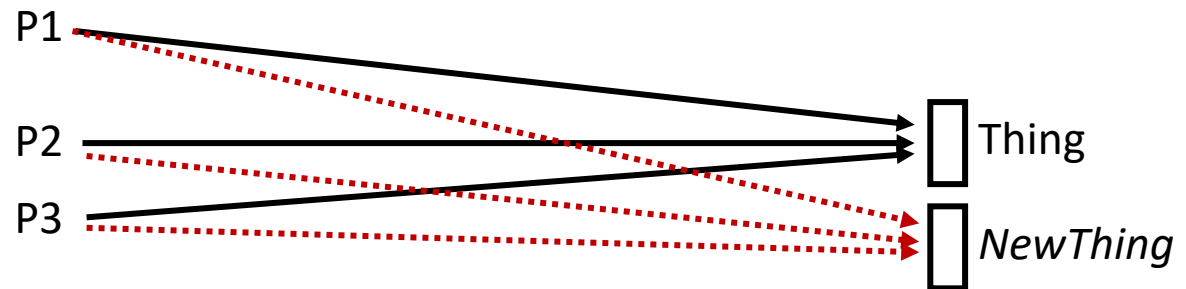
## Problem 4: How To Share?



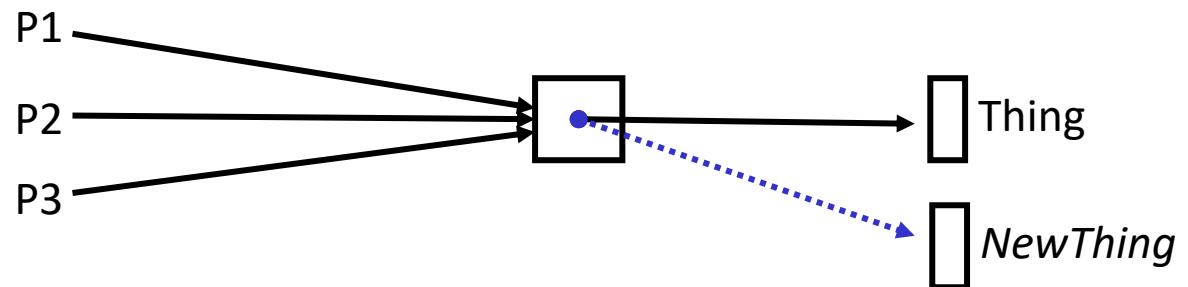
# How can we solve these problems?

- ❖ “Any problem in computer science can be solved by adding another level of **indirection**.” – *David Wheeler, inventor of the subroutine*

- ❖ Without Indirection



- ❖ With Indirection

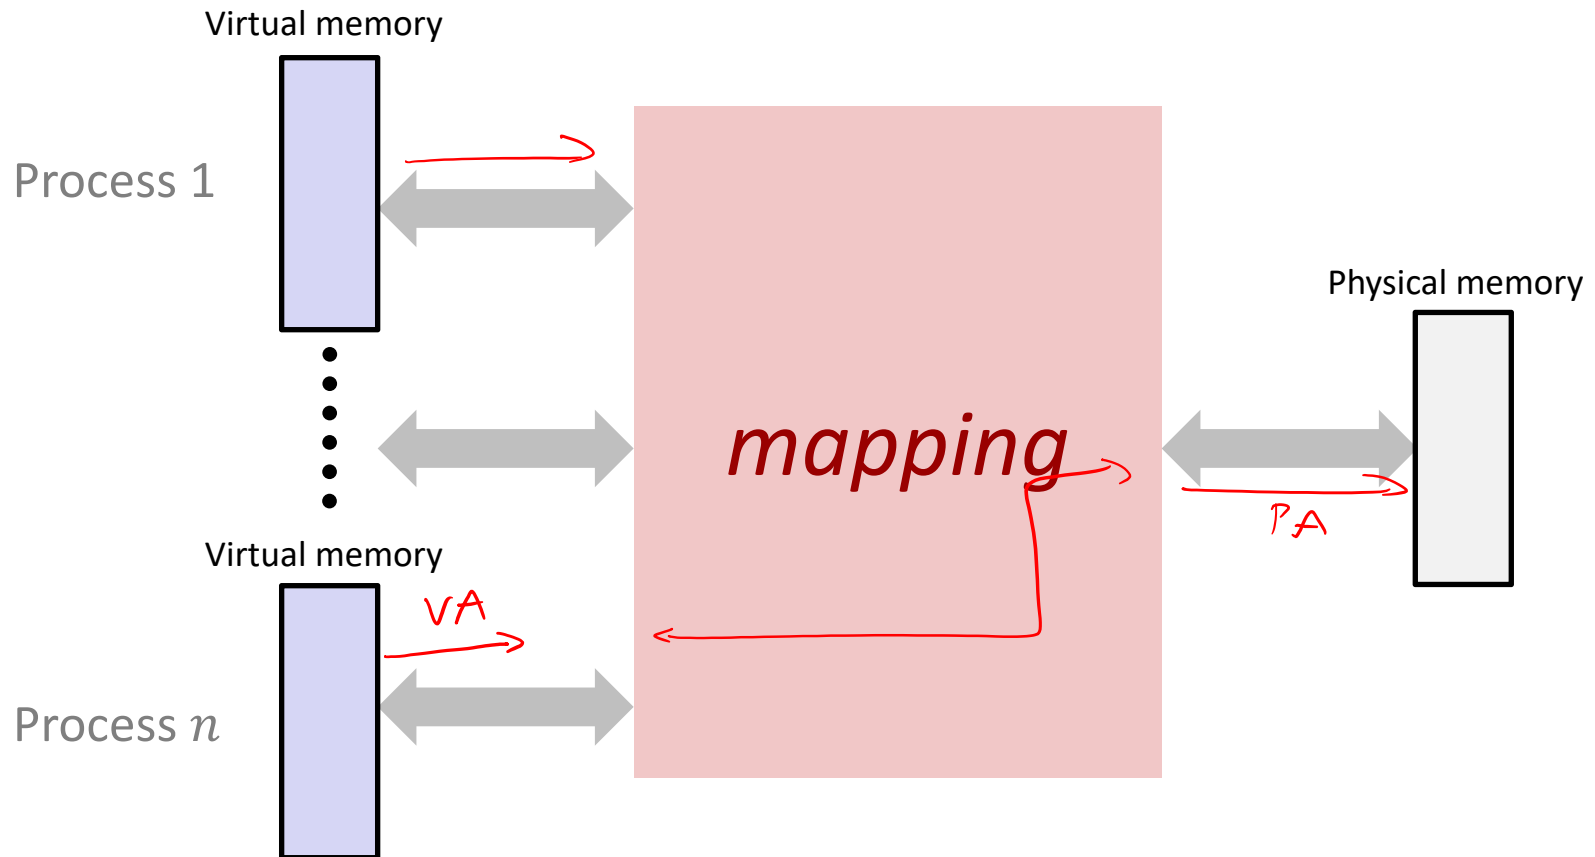


*What if I want to move Thing?*

# Indirection

- ❖ *Indirection*: The ability to reference something using a name, reference, or container instead of the value itself. A flexible mapping between a name and a thing allows changing the thing without notifying holders of the name.
  - ■ Adds some work (now have to look up 2 things instead of 1)
  - + ■ But don't have to track all uses of name/address (single source!)
- ❖ Examples:
  - **Phone system**: cell phone number portability
  - **Domain Name Service (DNS)**: translation from name to IP address
  - **Call centers**: route calls to available operators, etc.
  - **Dynamic Host Configuration Protocol (DHCP)**: local network address assignment

# Indirection in Virtual Memory



- ❖ Each process gets its own private virtual address space
- ❖ Solves the previous problems!

# Address Spaces


- $n = \lceil \log_2 N \rceil$  (bits)  
 $m = \lceil \log_2 M \rceil$  (bytes)
- ❖ **Virtual address space:** Set of  $N = 2^n$  virtual addr
    - $\{0, 1, 2, 3, \dots, N-1\}$
  - ❖ **Physical address space:** Set of  $M = 2^m$  physical addr
    - $\{0, 1, 2, 3, \dots, M-1\}$

- ❖ Every byte in main memory has:
  - one physical address (PA)
  - zero, one, or more virtual addresses (VAs)

unused  
 used by one process  
 used by many processes

# Mapping

- ❖ A virtual address (VA) can be mapped to either **physical memory** or **disk**
  - Unused VAs may not have a mapping
  - VAs from *different* processes may map to same location in memory/disk



# Summary

- ❖ Virtual memory provides:
  - Ability to use limited memory (RAM) across multiple processes
  - Illusion of contiguous virtual address space for each process
  - Protection and sharing amongst processes