



Virtual Memory

Process Abstraction, Part 2: Private Address Space

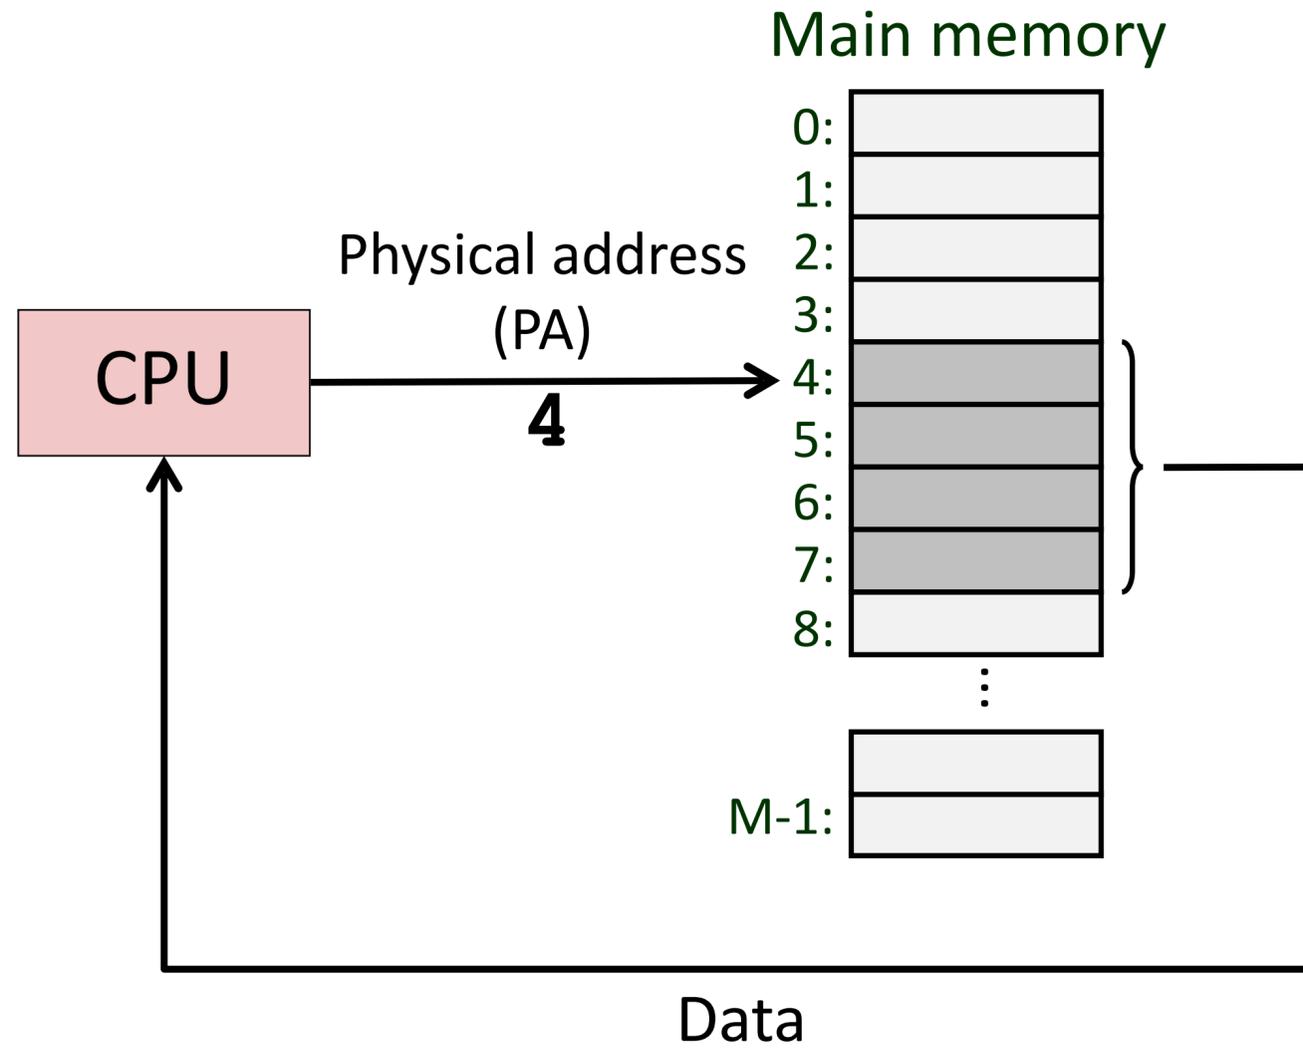
Motivation: why not direct physical memory access?

Address translation with pages

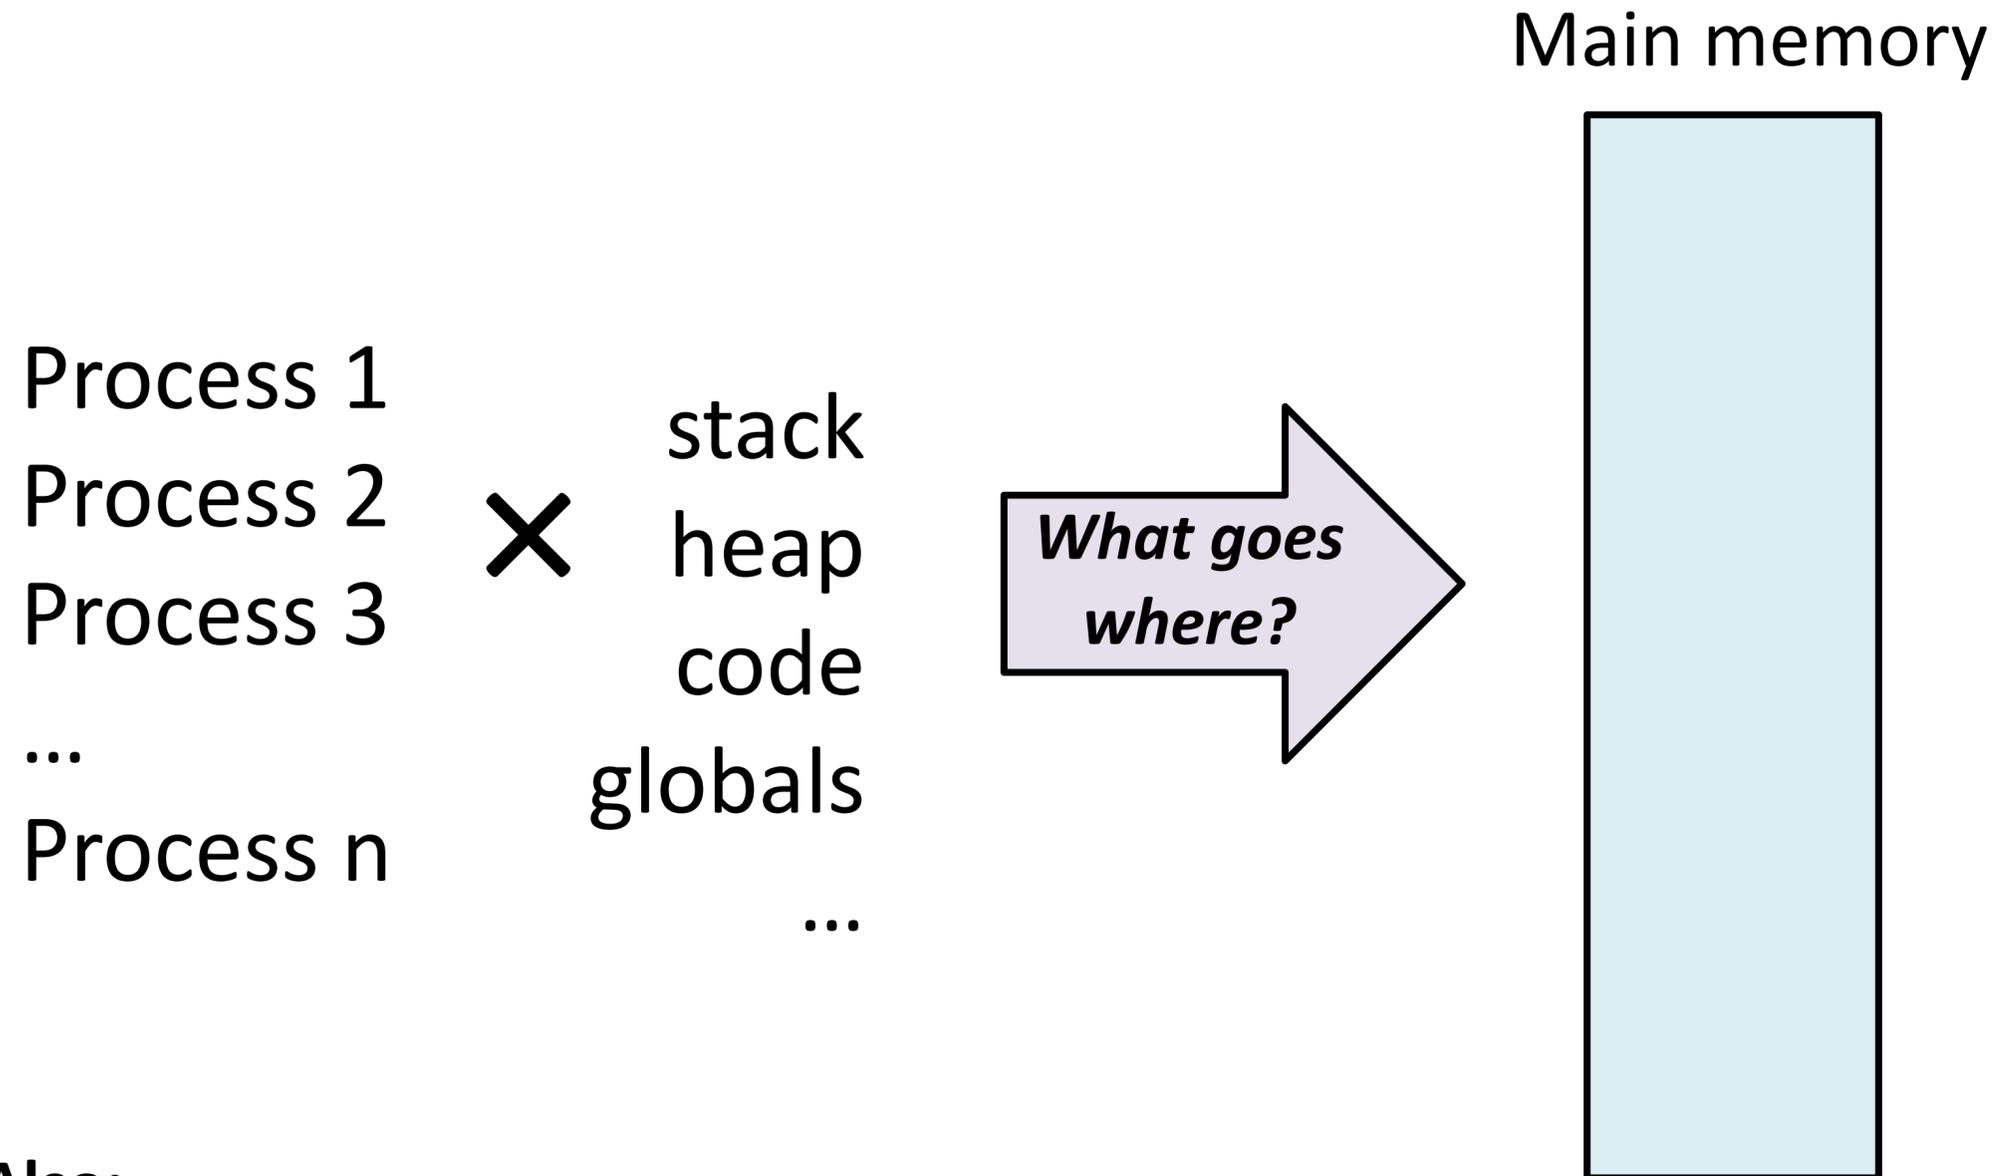
Extra benefits: sharing and protection

Memory as a contiguous array of bytes is a lie! Why?

Problems with physical addressing



Problem 1: memory management



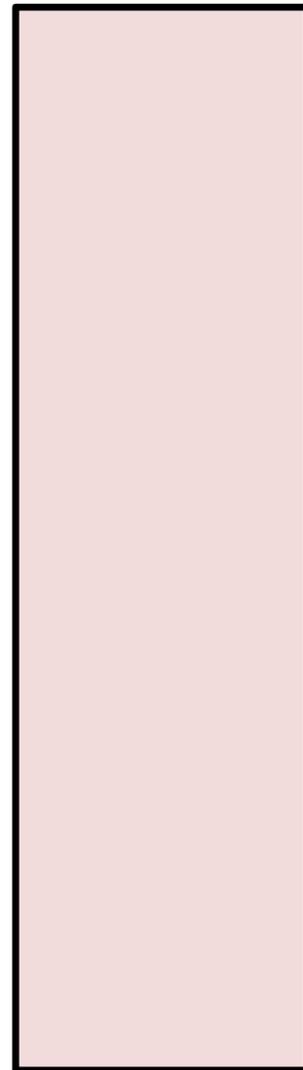
Also:

Context switches must swap out entire memory contents.

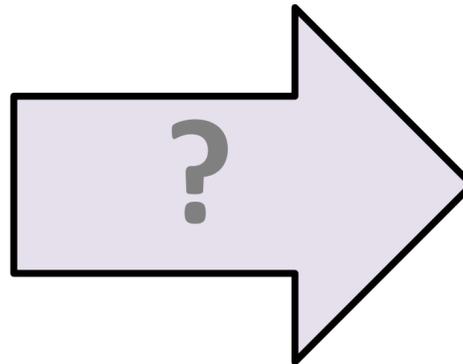
Isn't that **expensive**?

Problem 2: capacity

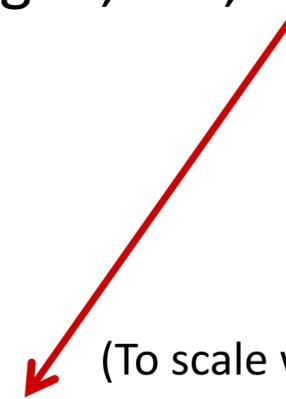
64-bit addresses can address
several exabytes
(18,446,744,073,709,551,616 bytes)



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

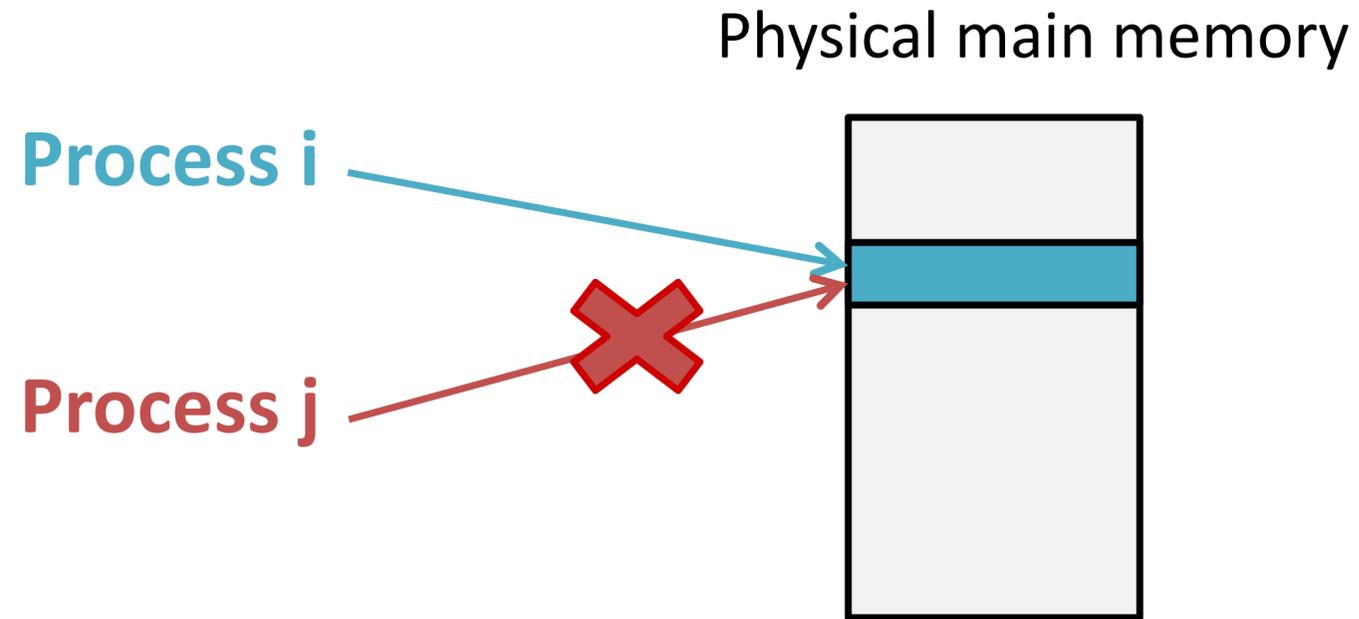


Physical main memory offers
a few gigabytes
(e.g. 8,589,934,592 bytes)

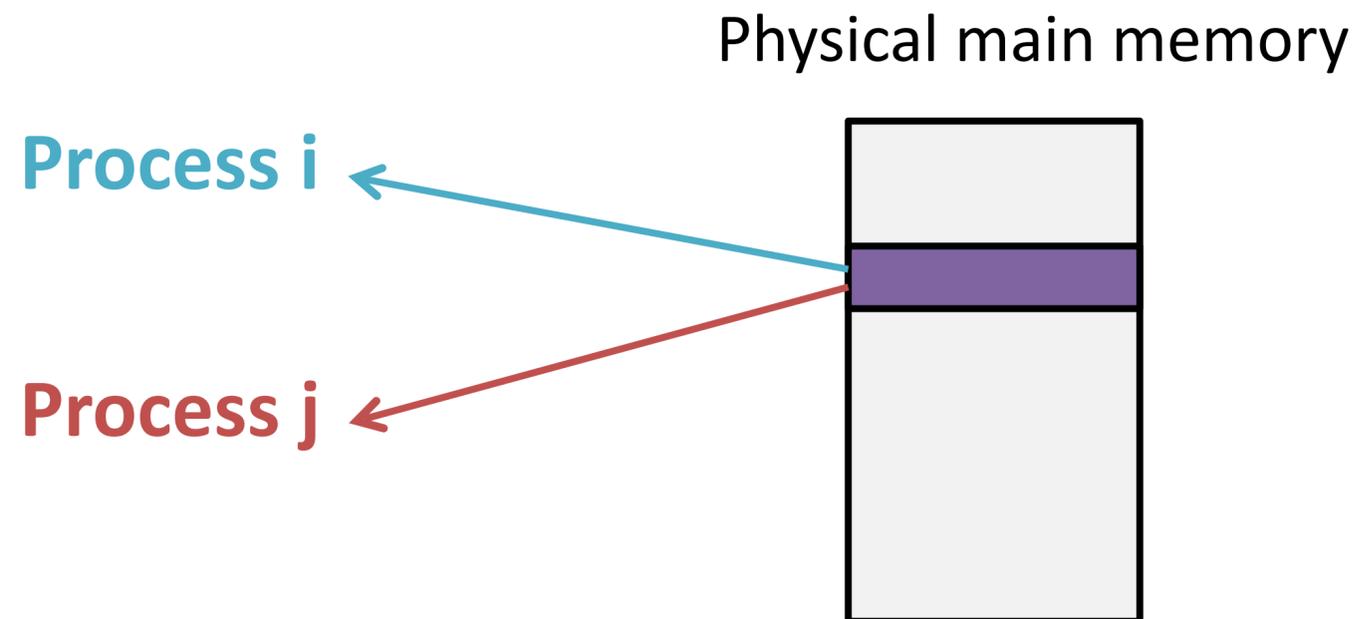


(To scale with 64-bit address space,
you can't see it.)

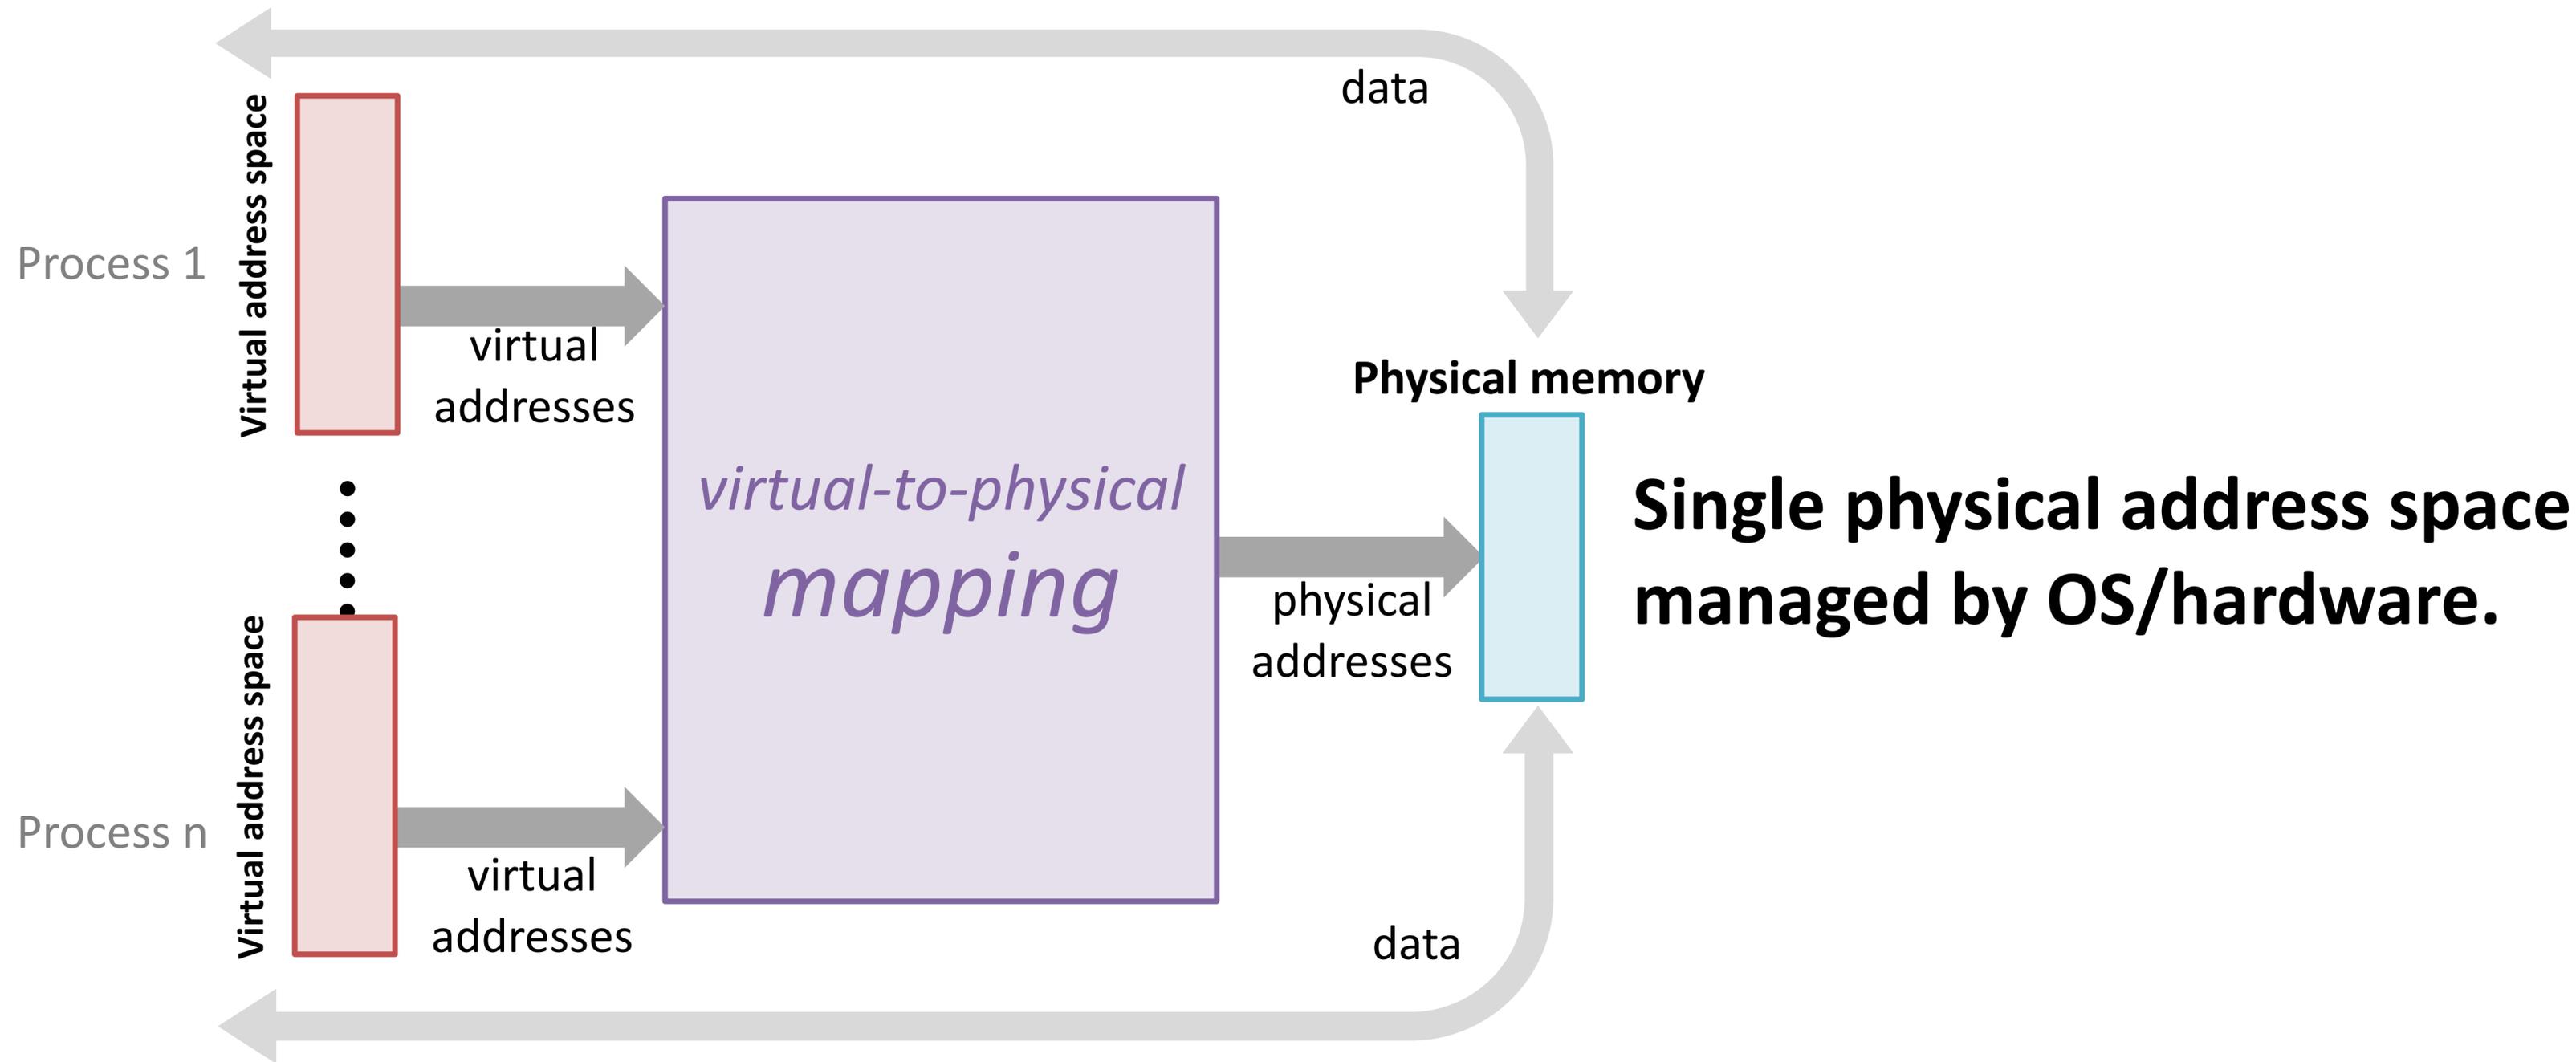
Problem 3: protection



Problem 4: sharing



Solution: Virtual Memory (address *indirection*)

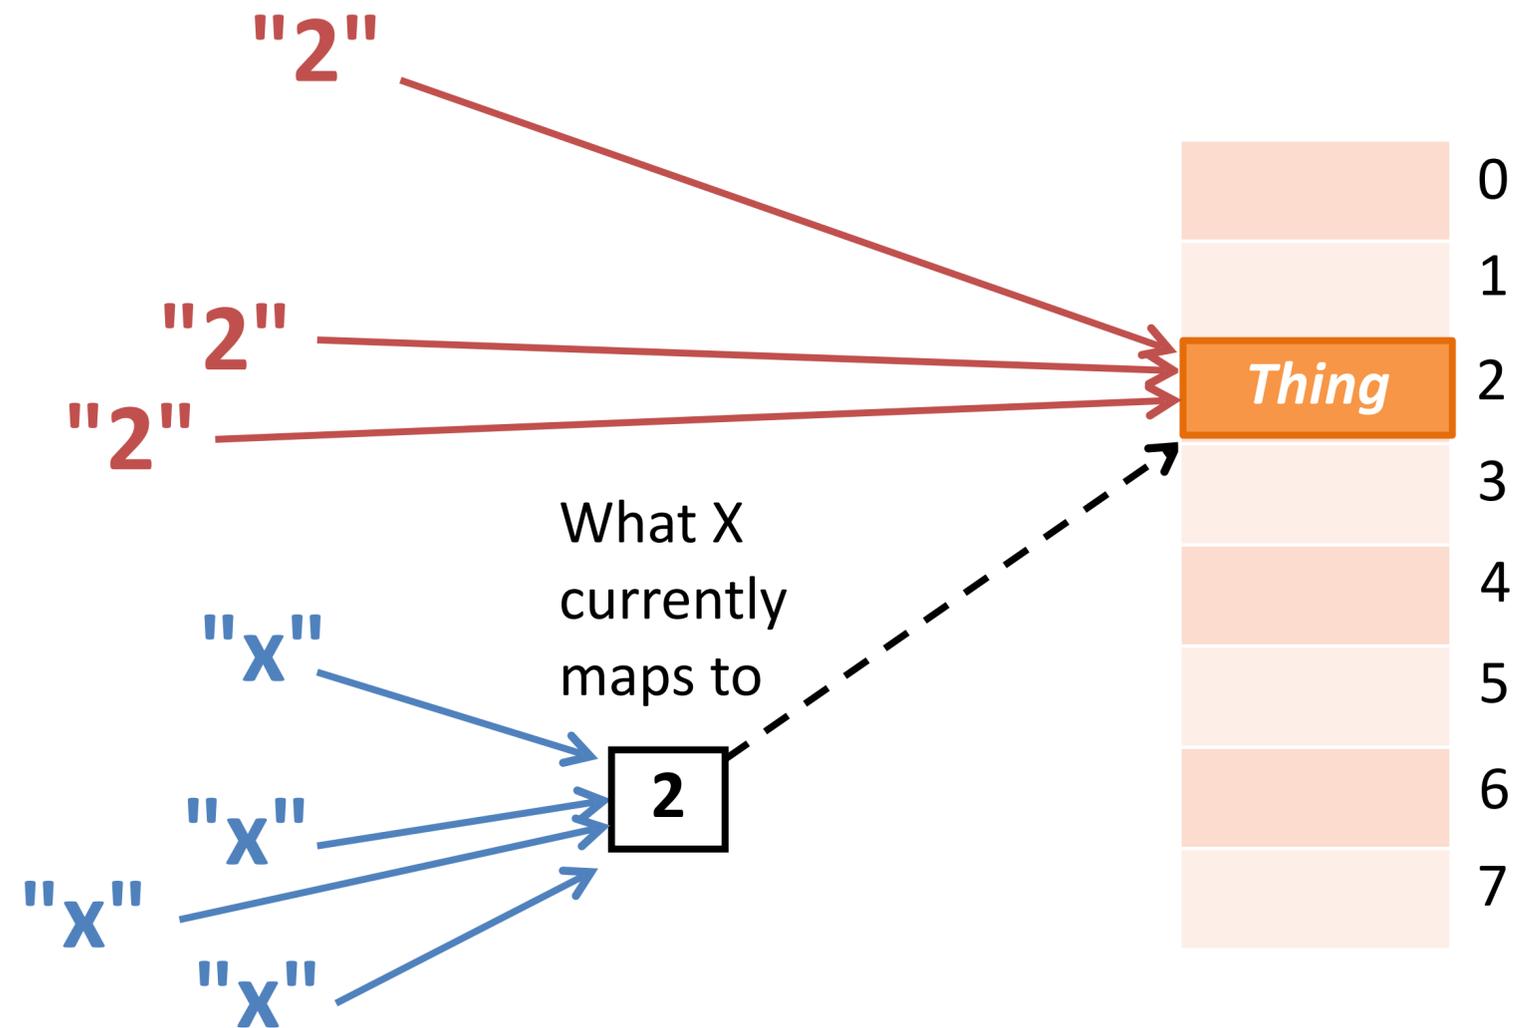


Indirection

(it's everywhere!)

Direct naming

Indirect naming



What if we move *Thing*?

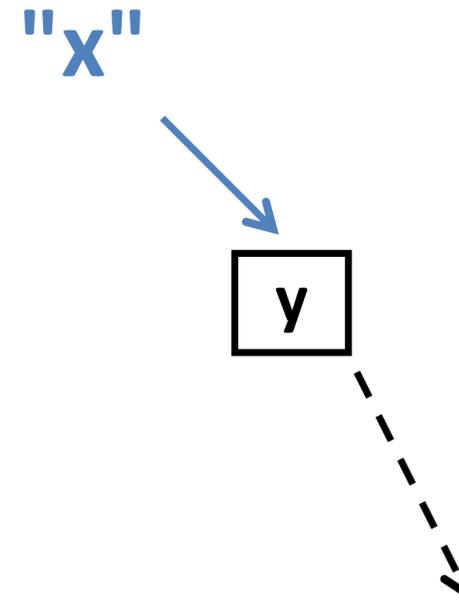
Tangent: indirection everywhere

- Pointers
- Constants
- Procedural abstraction
- Domain Name Service (DNS)
- Dynamic Host Configuration Protocol (DHCP)
- Phone numbers
- 911
- Call centers
- Snail mail forwarding
- ...

“Any problem in computer science can be solved by adding another level of indirection.”

–David Wheeler, inventor of the subroutine, or Butler Lampson

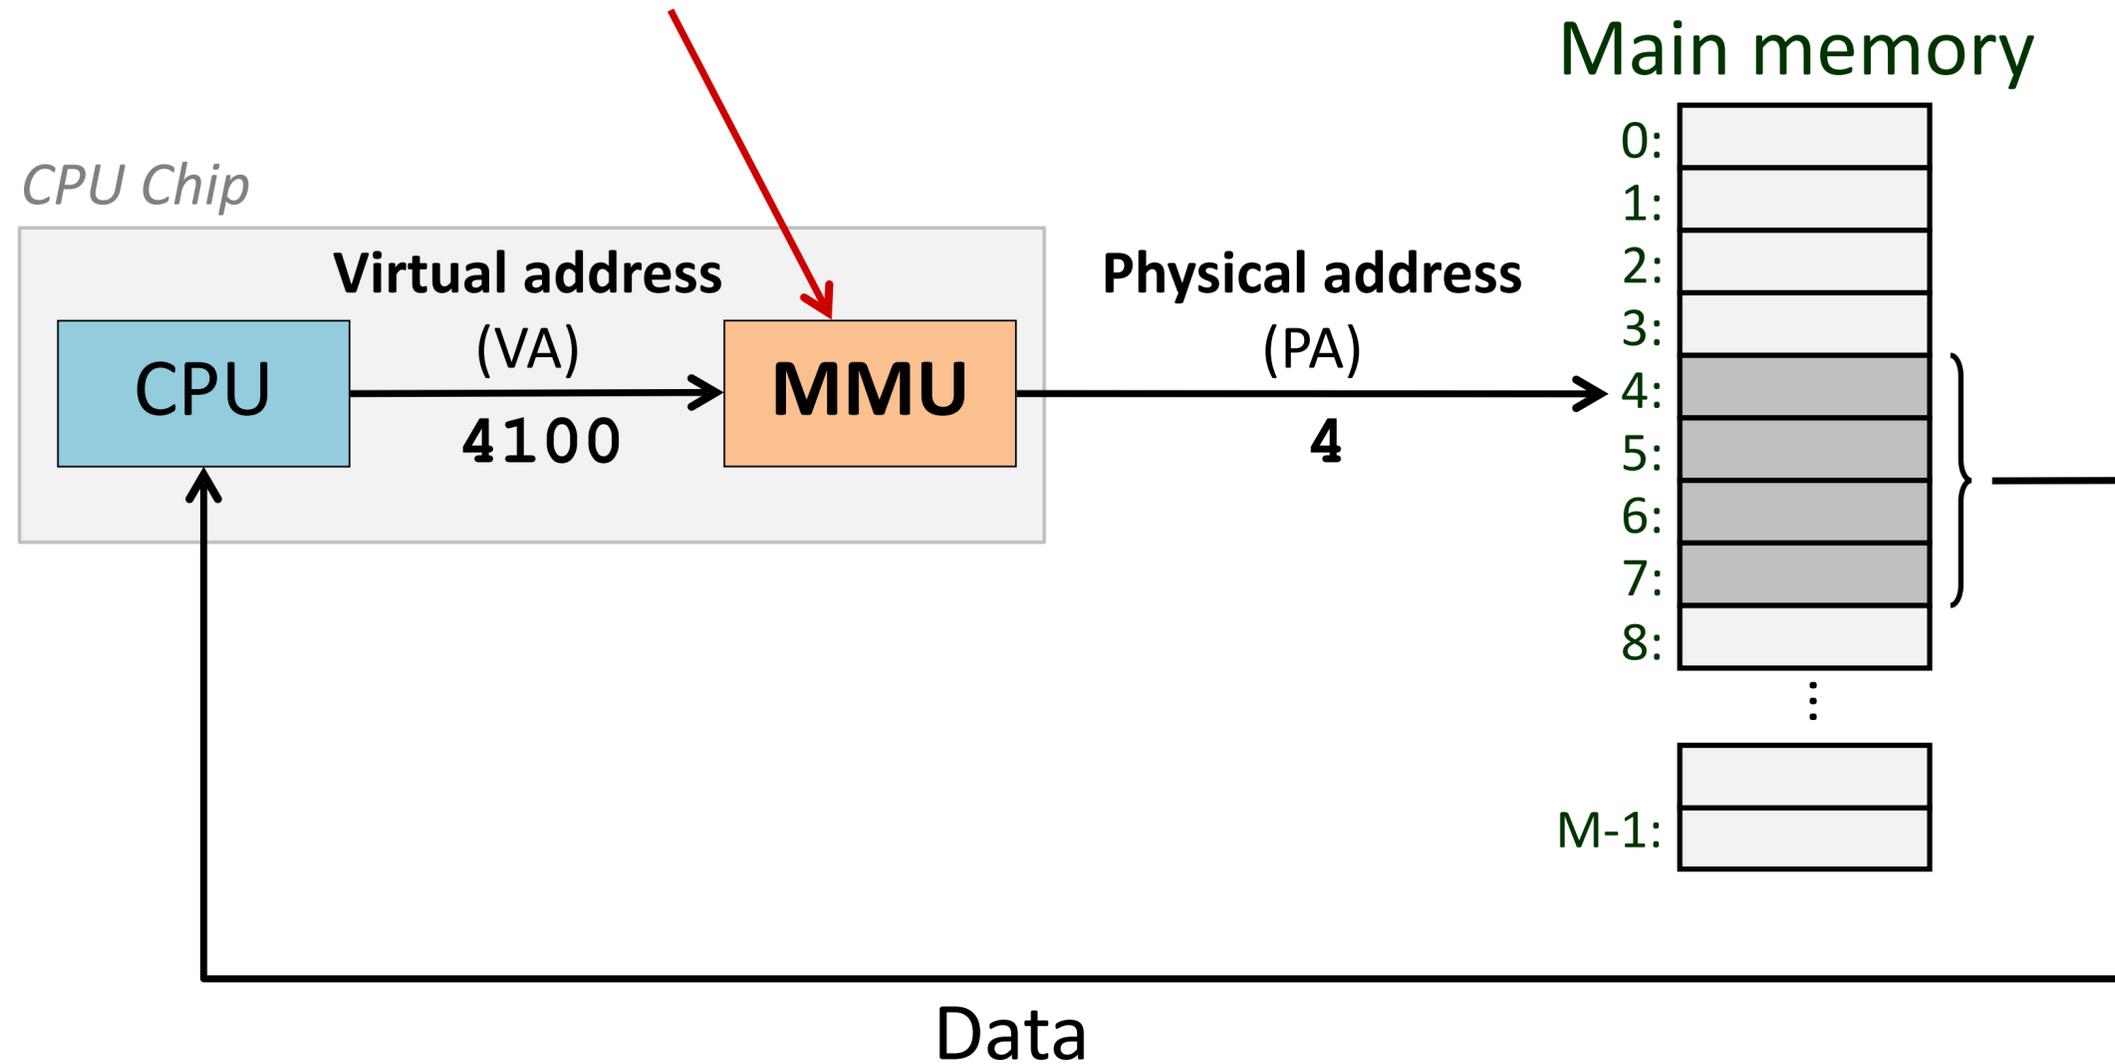
Another Wheeler quote? "Compatibility means deliberately repeating other people's mistakes."



Virtual addressing and address translation

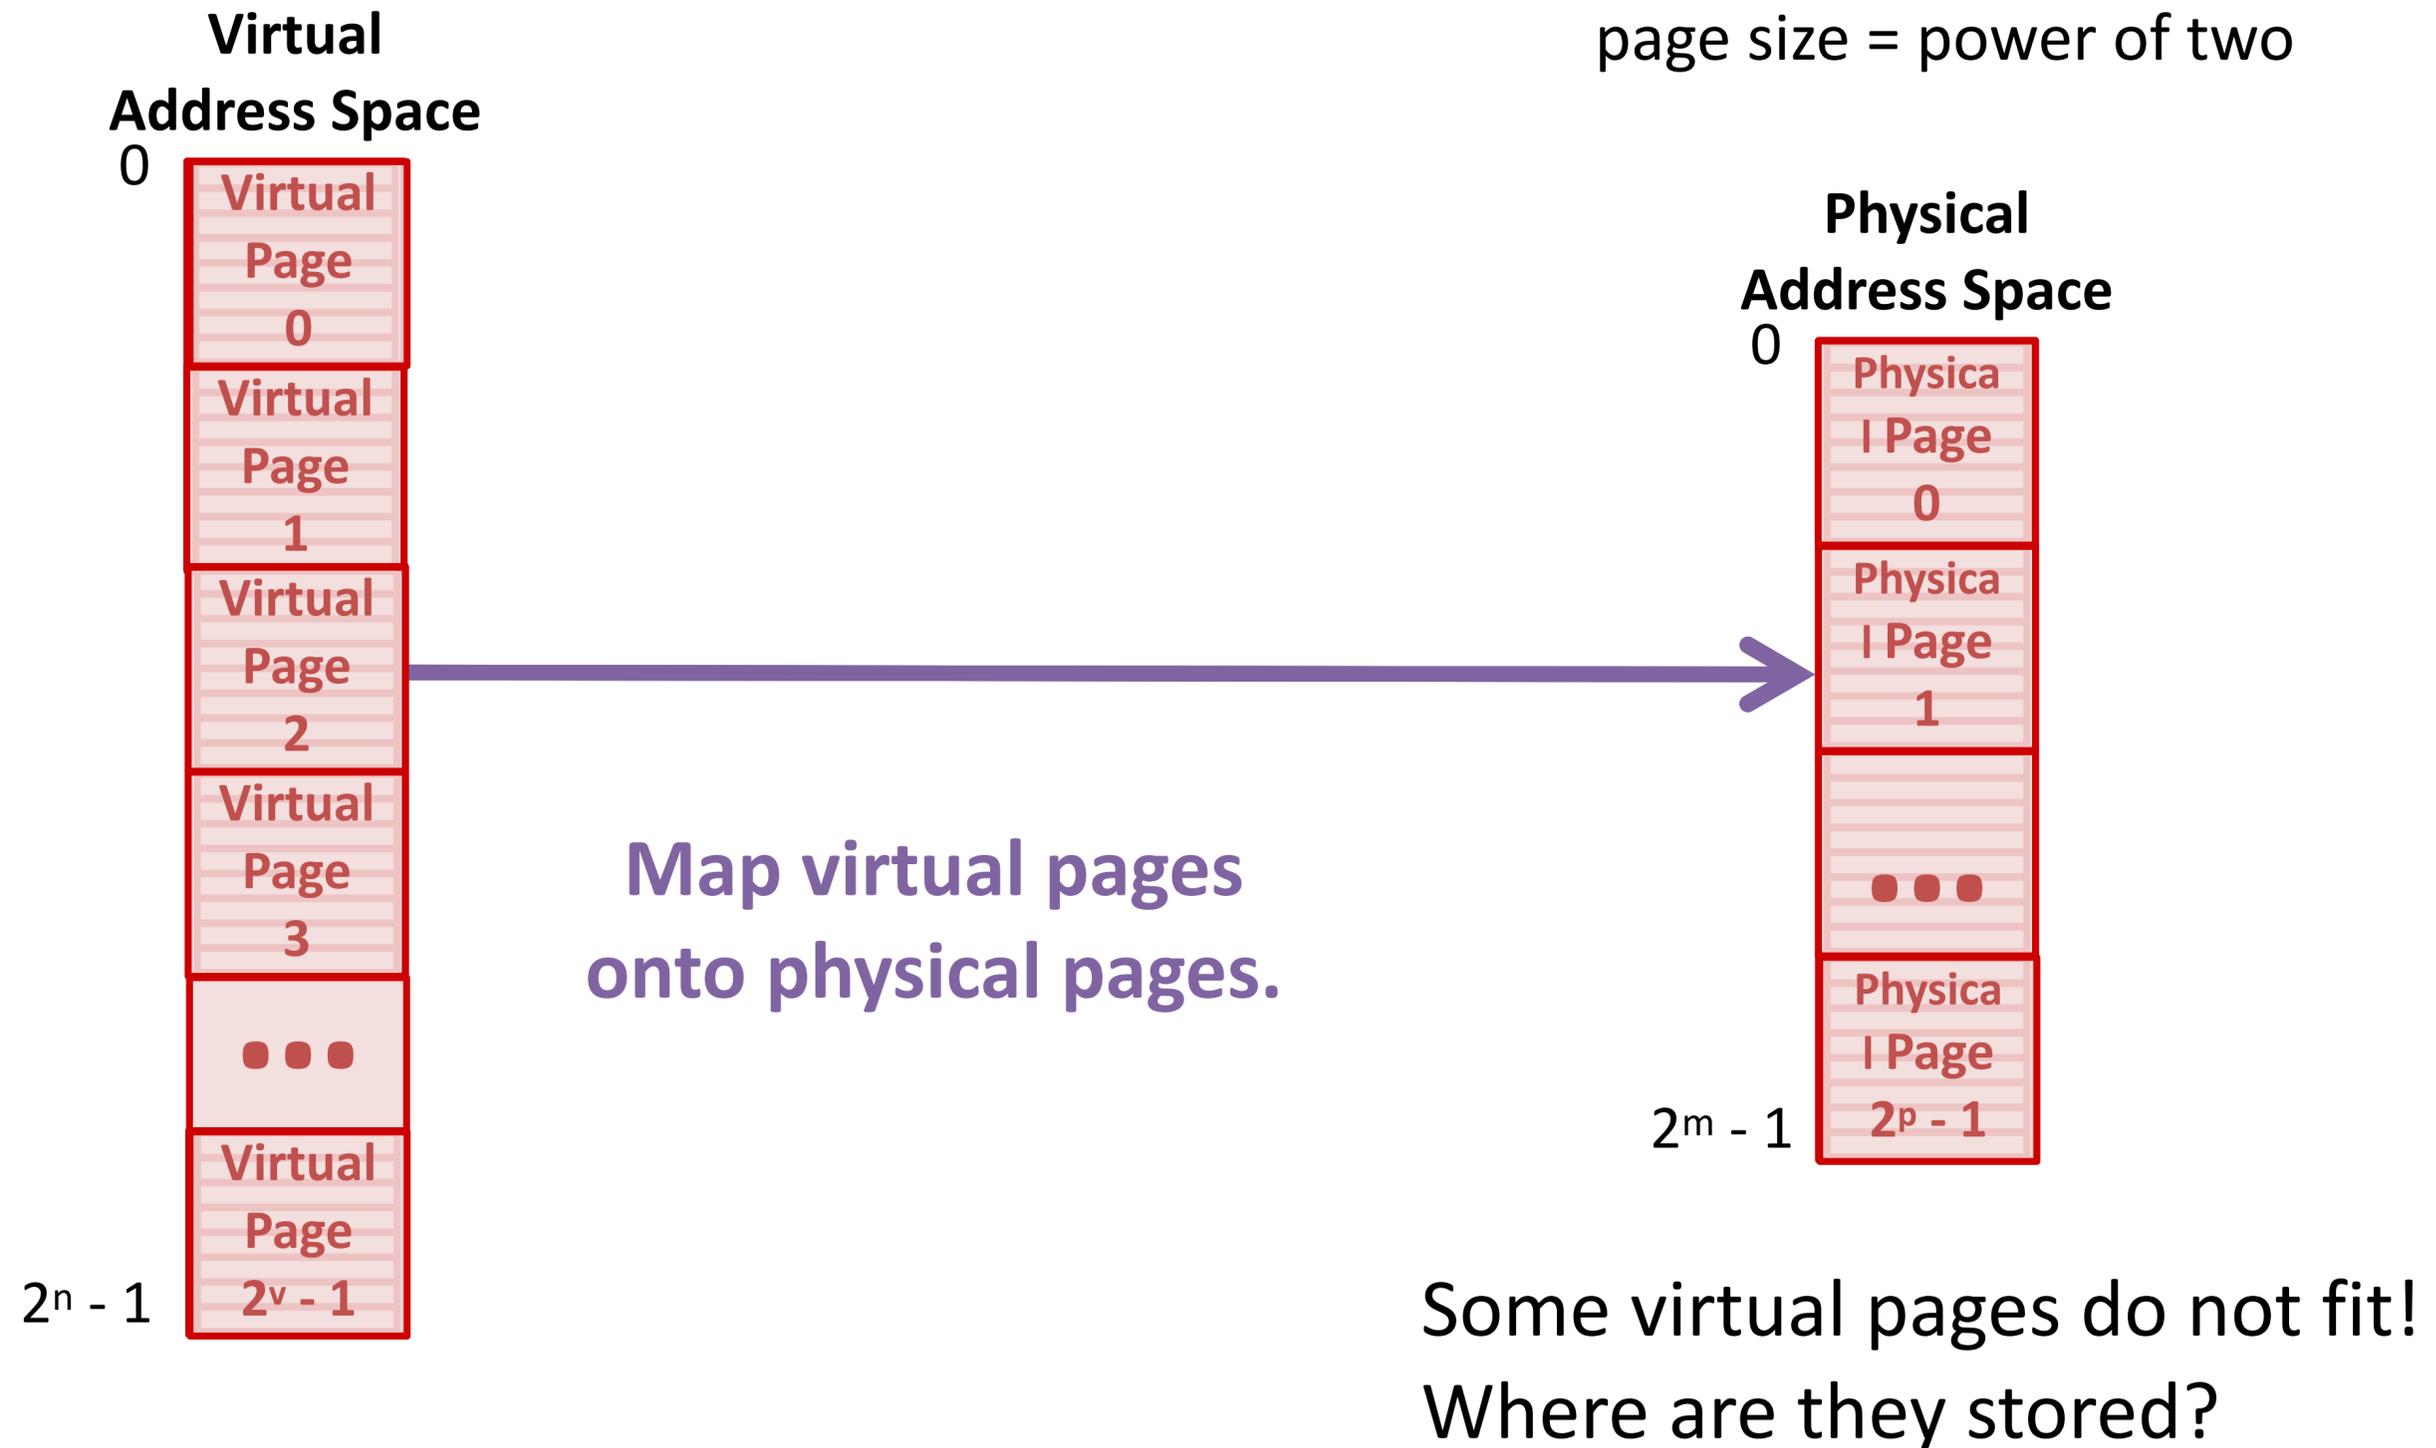
Memory Management Unit

translates virtual address to physical address

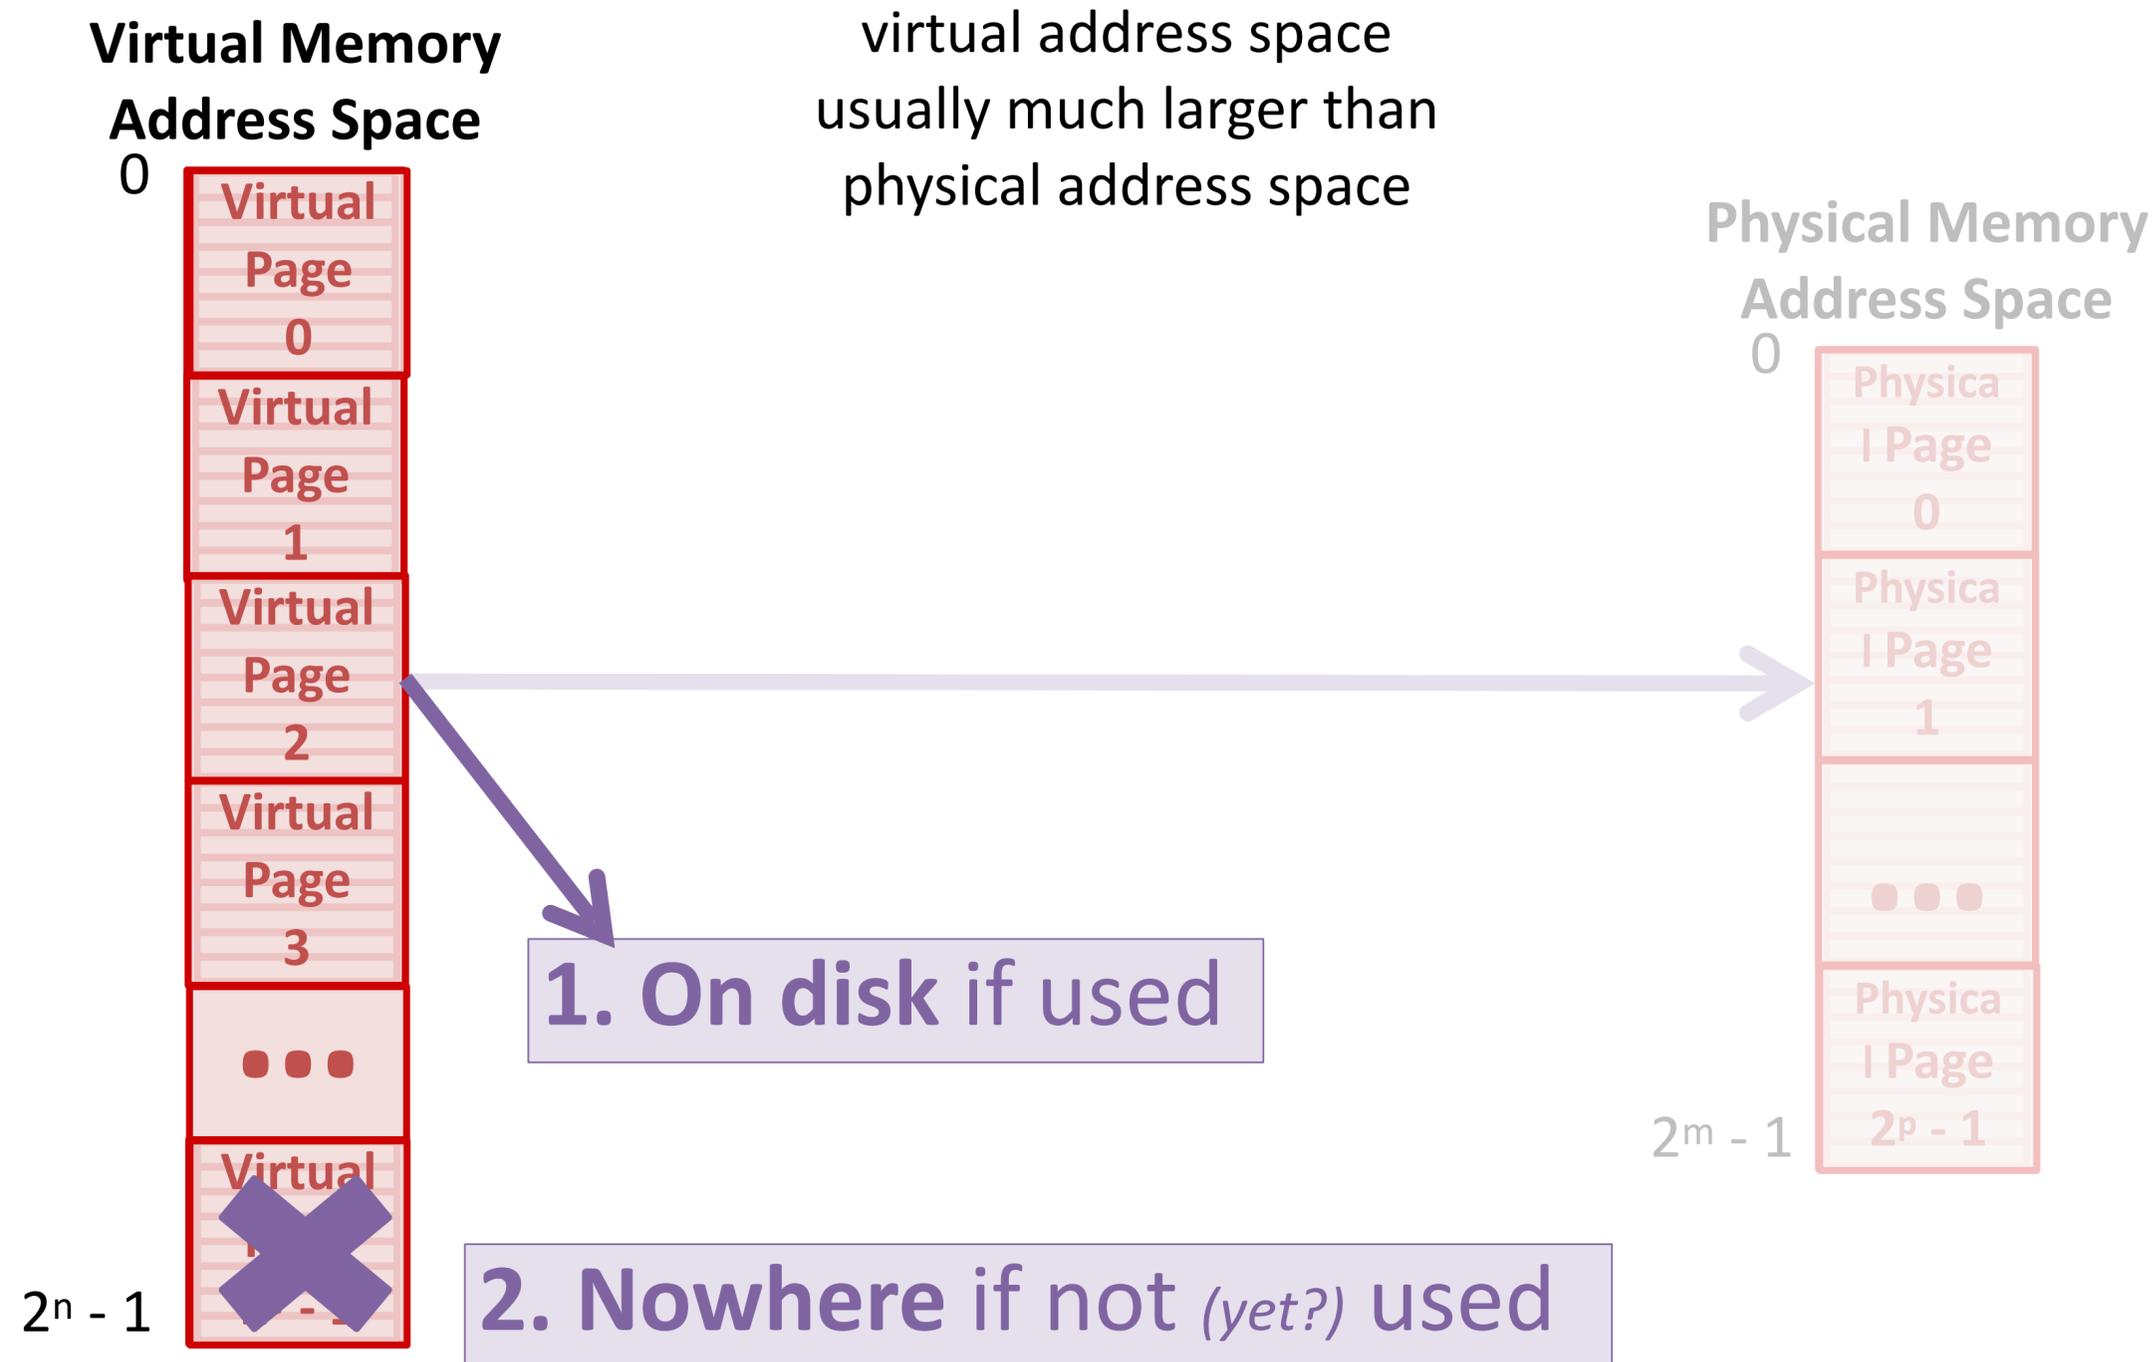


Page-based mapping

fixed-size, aligned *pages*
page size = power of two

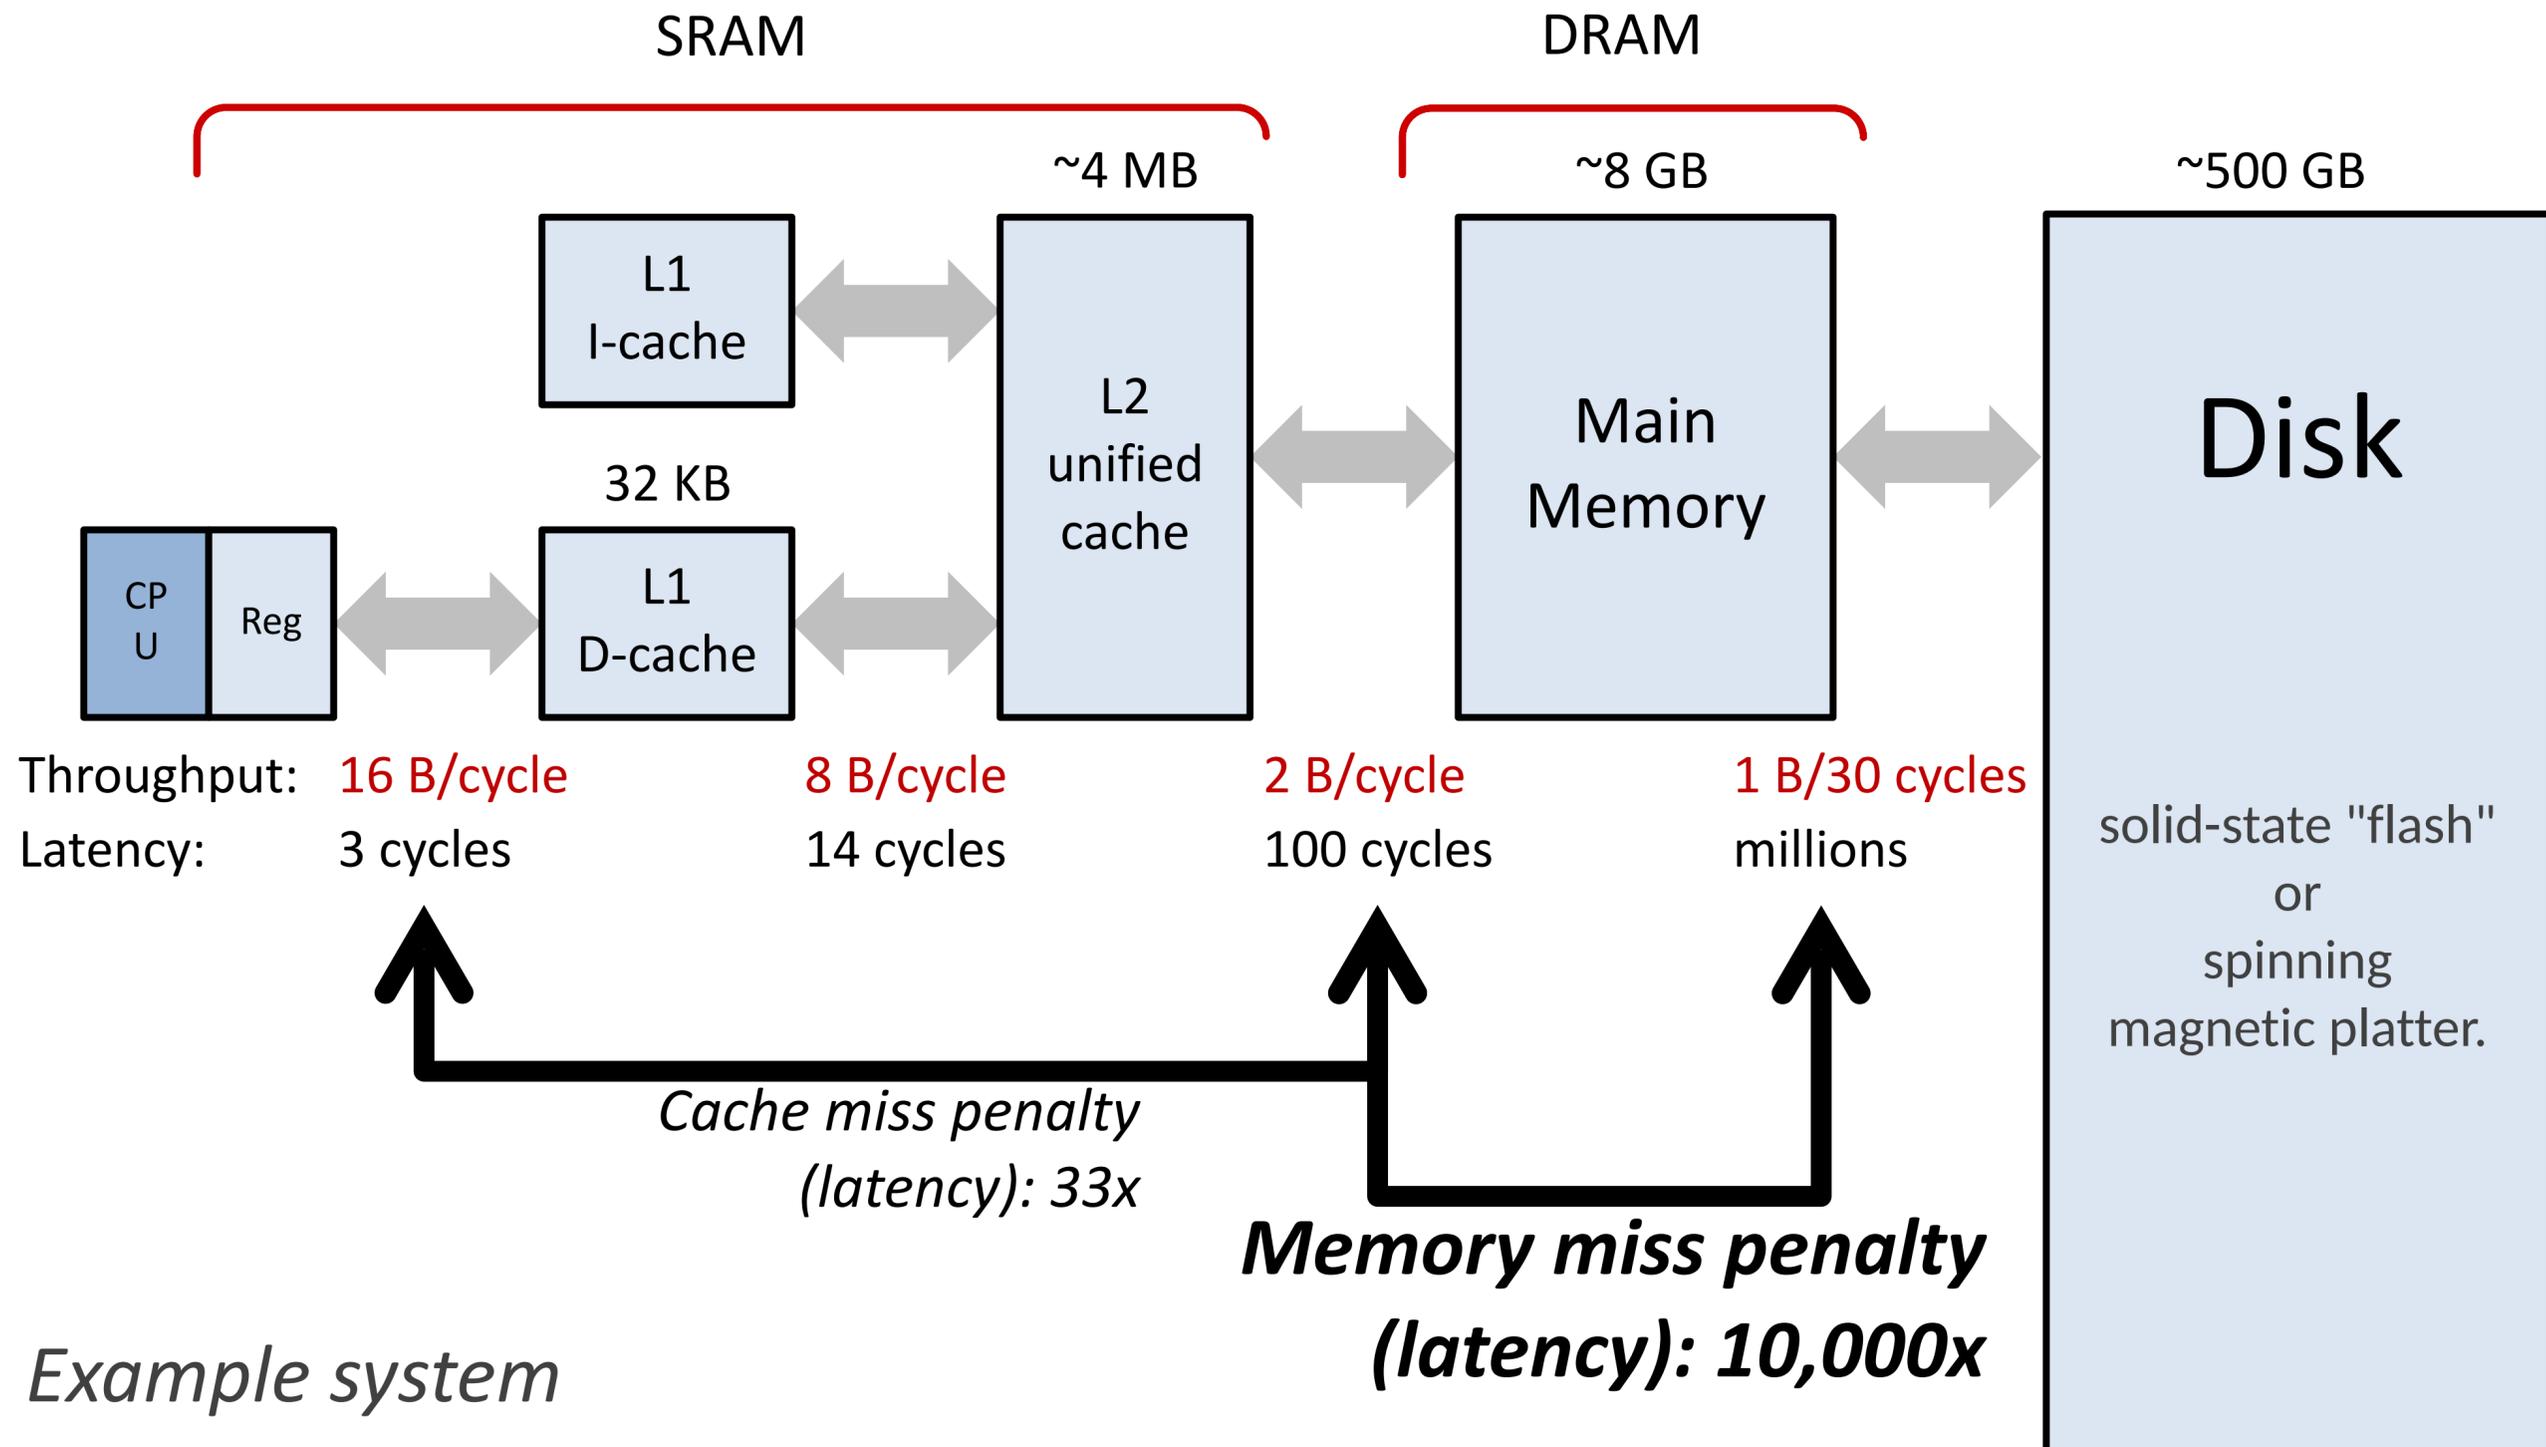


Cannot fit all virtual pages! Where are the rest stored?

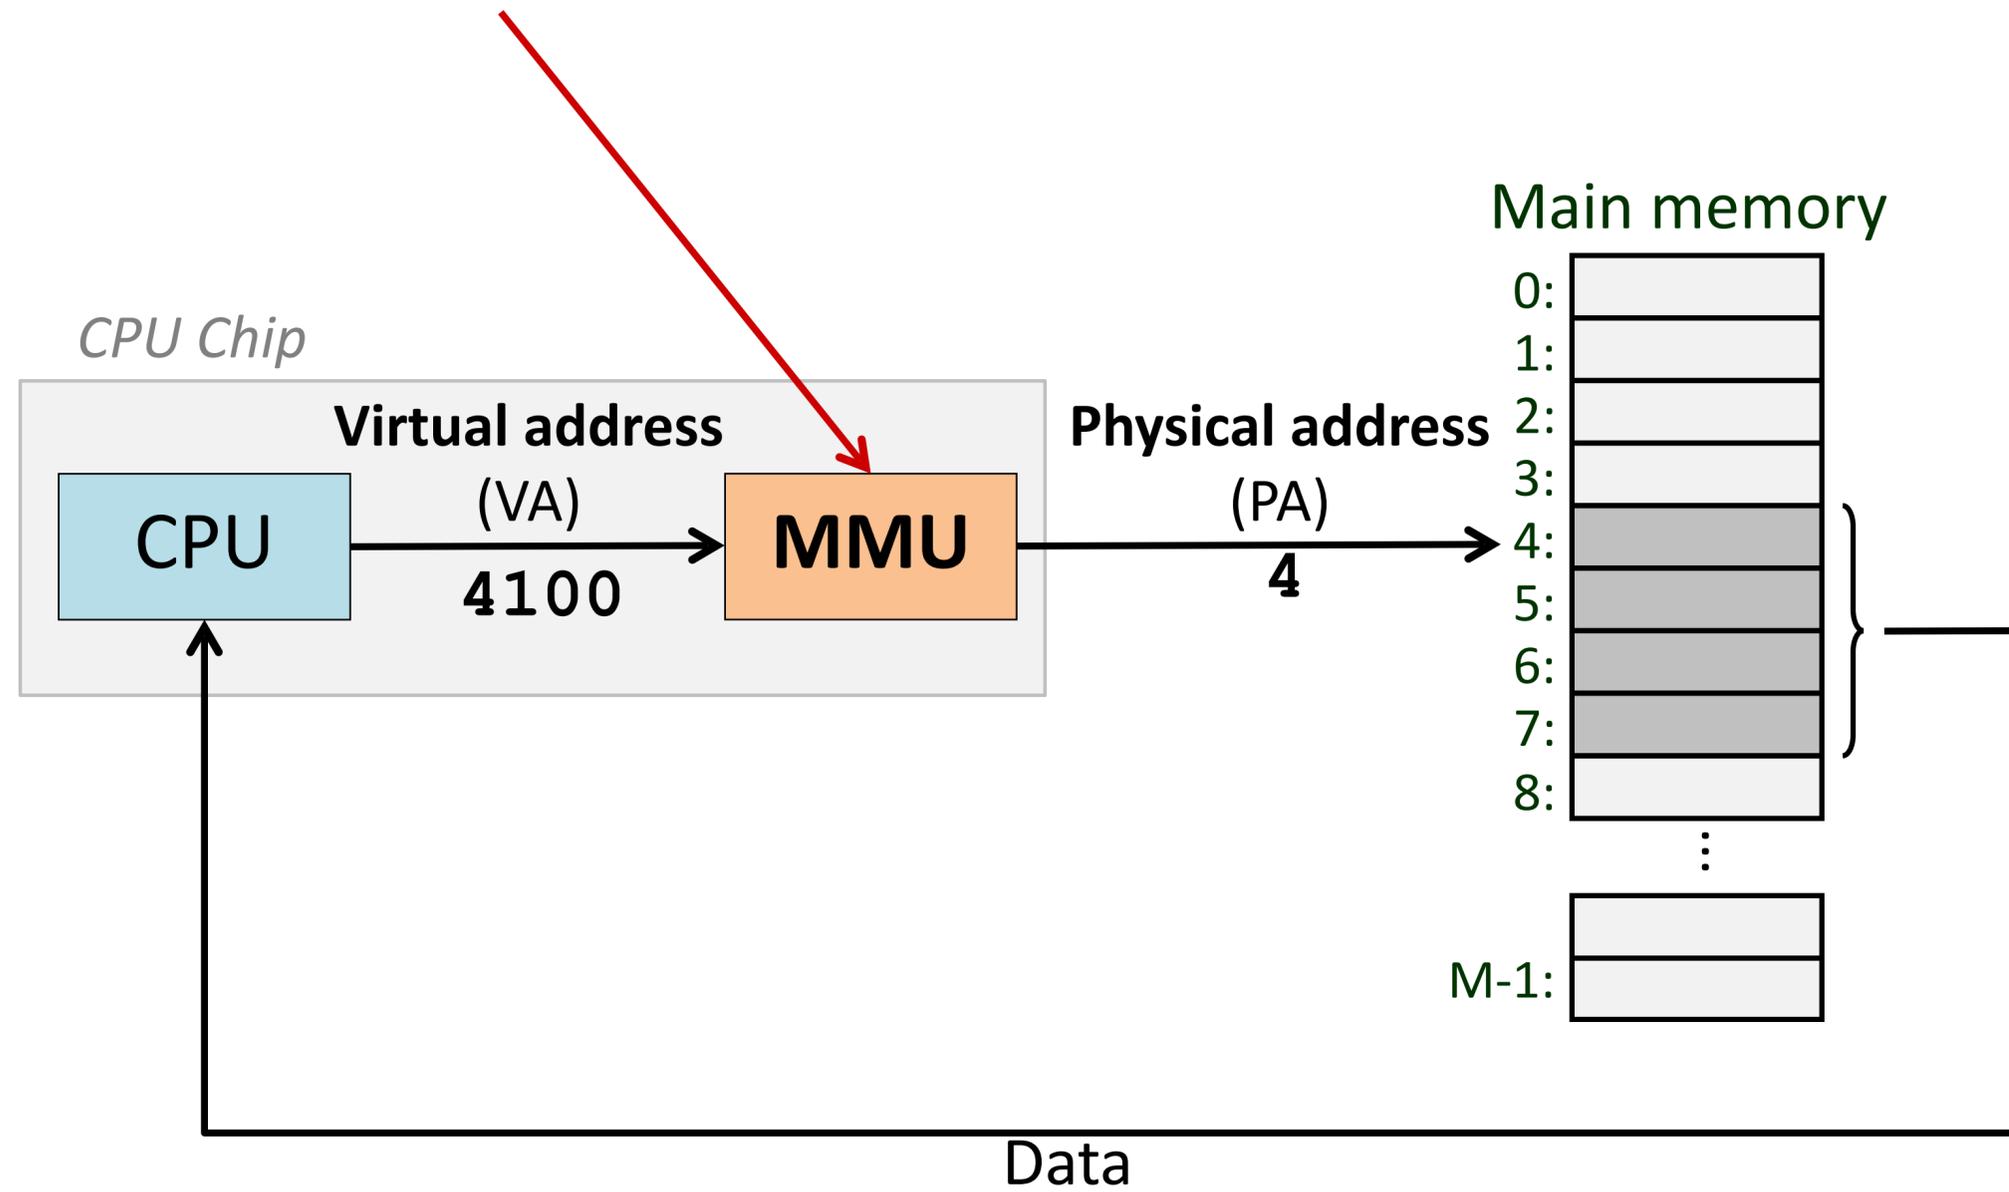


Virtual memory: cache for disk?

Not drawn to scale!

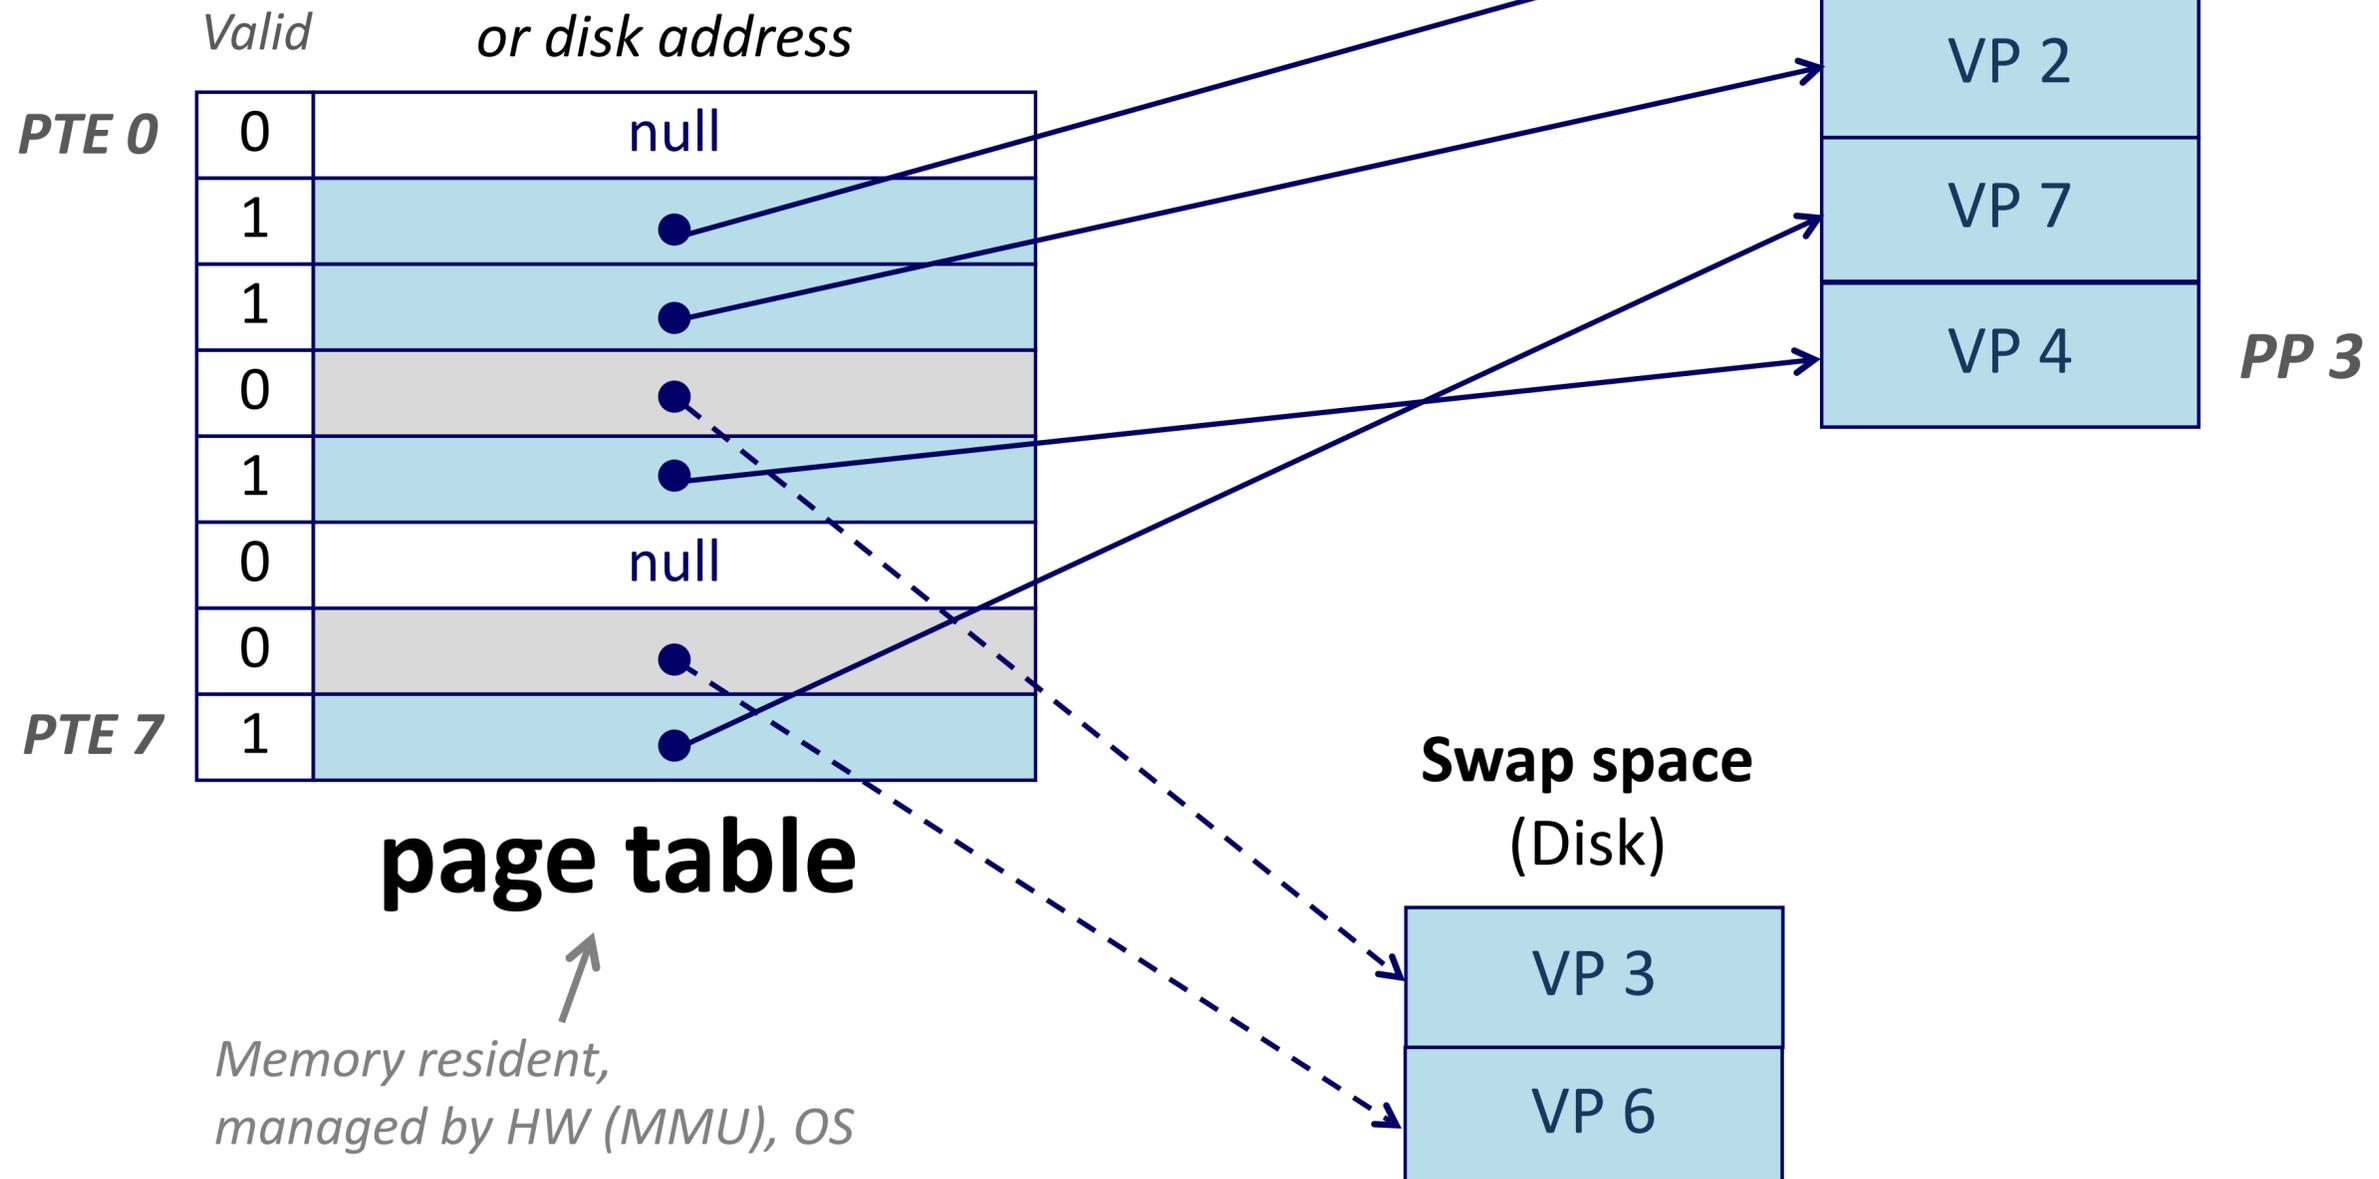


Address translation



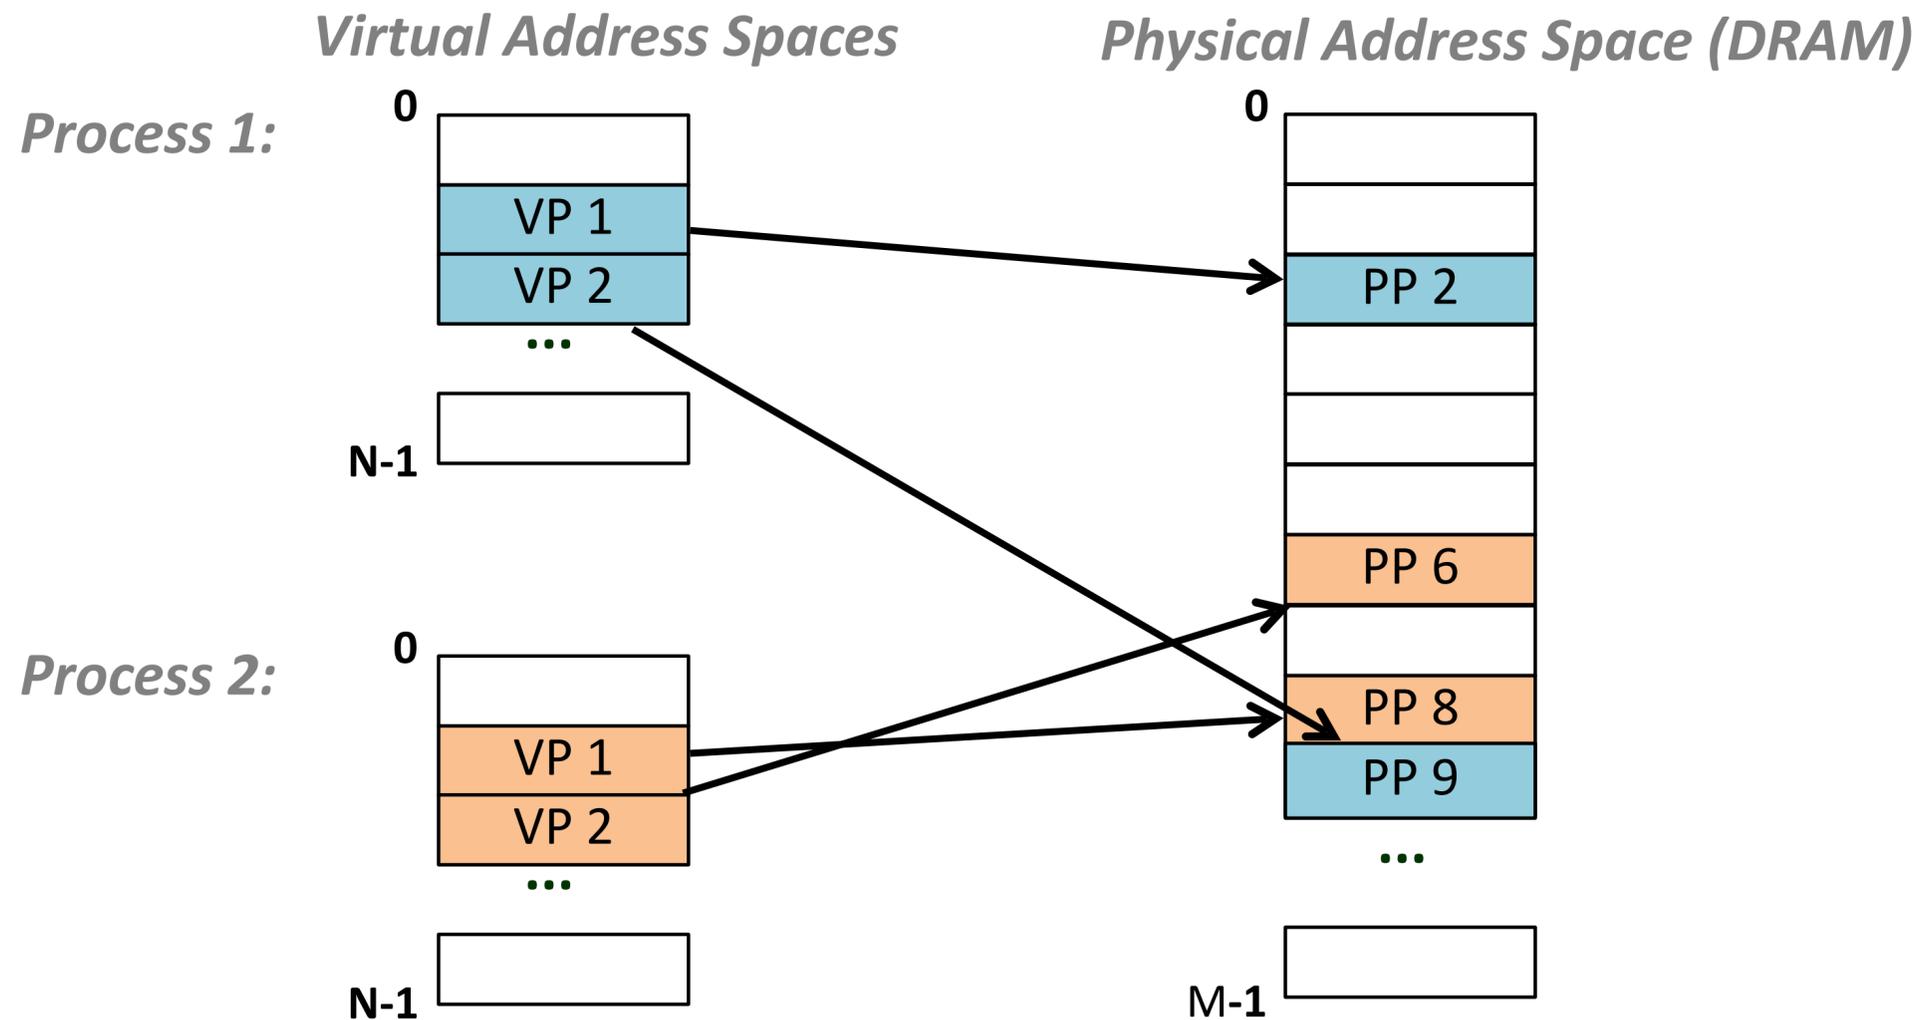
Page table

array of *page table entries* (PTEs)
mapping virtual page to where it is stored
Physical Page Number
or disk address



Virtual memory benefits: Simple address space allocation

Process needs private *contiguous* address space.



Virtual memory benefits:

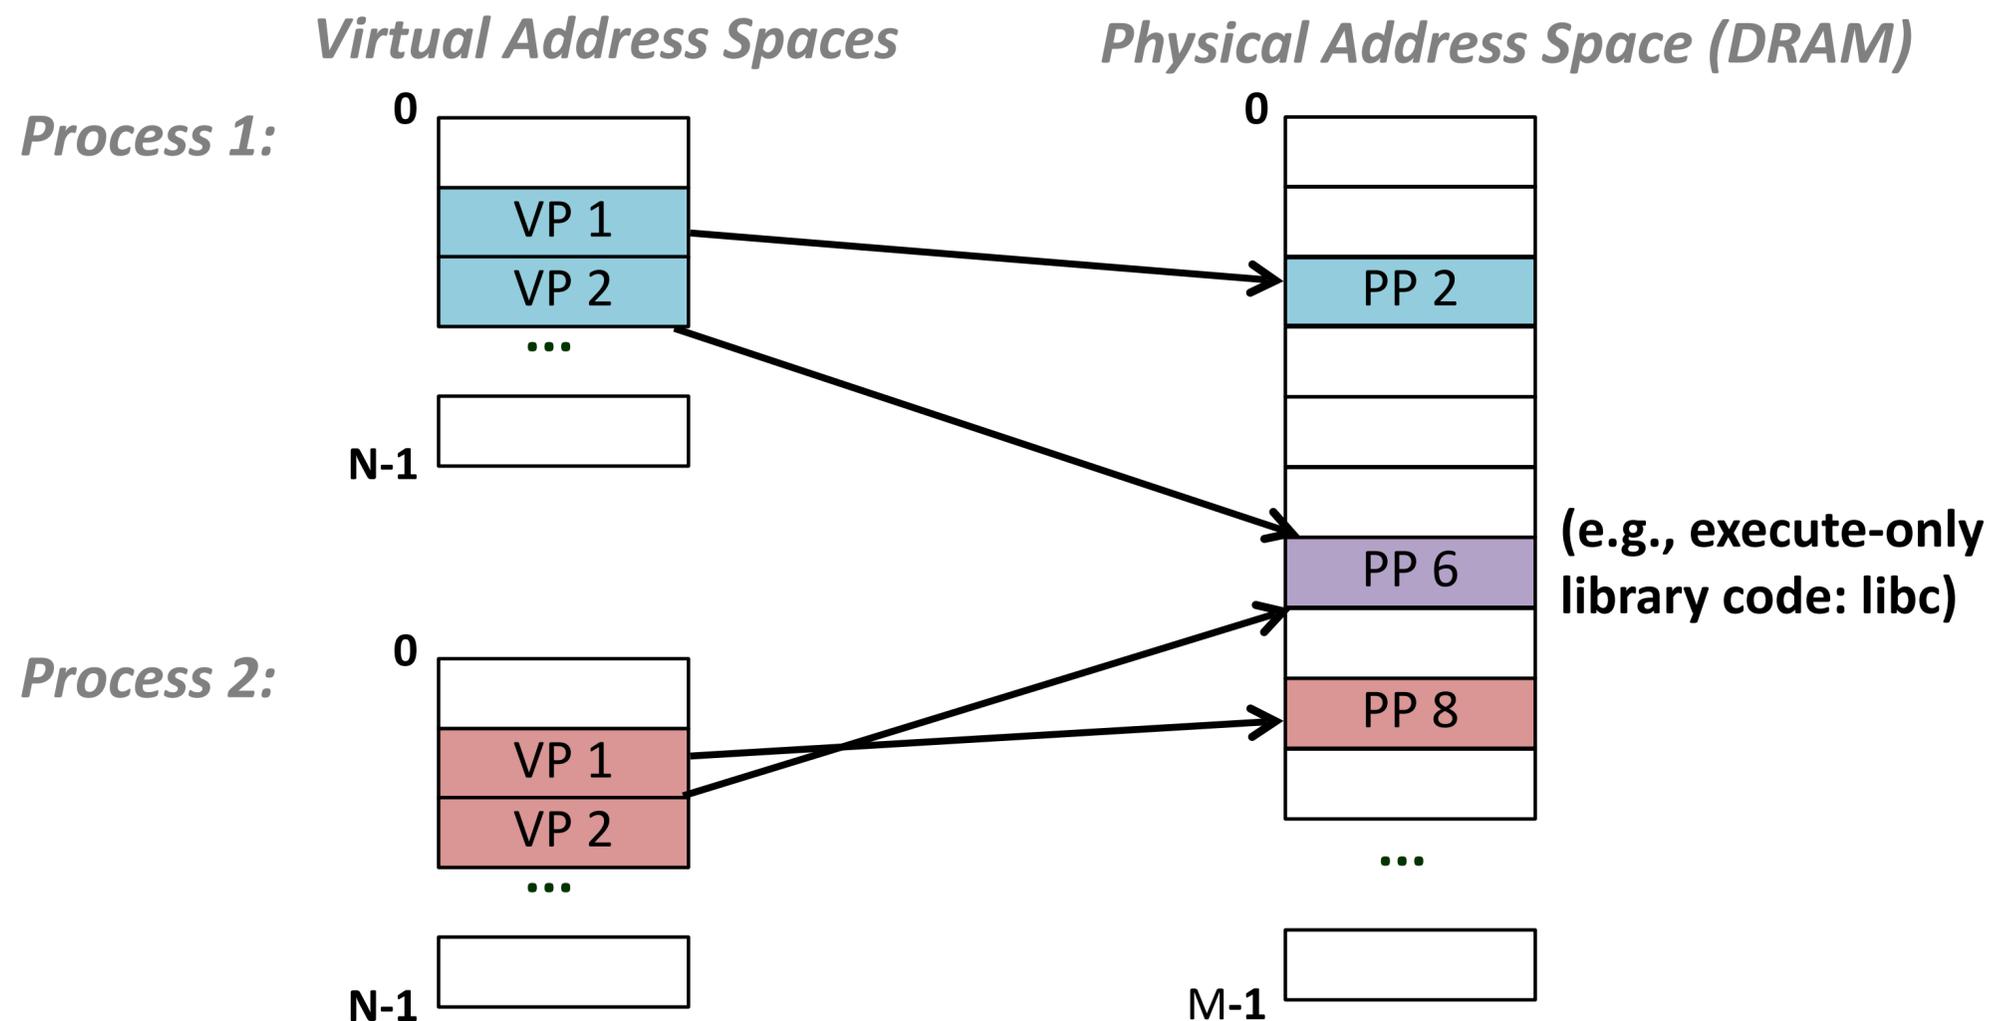
Protection:

All accesses go through translation.

Impossible to access physical memory not mapped in virtual address space.

Sharing:

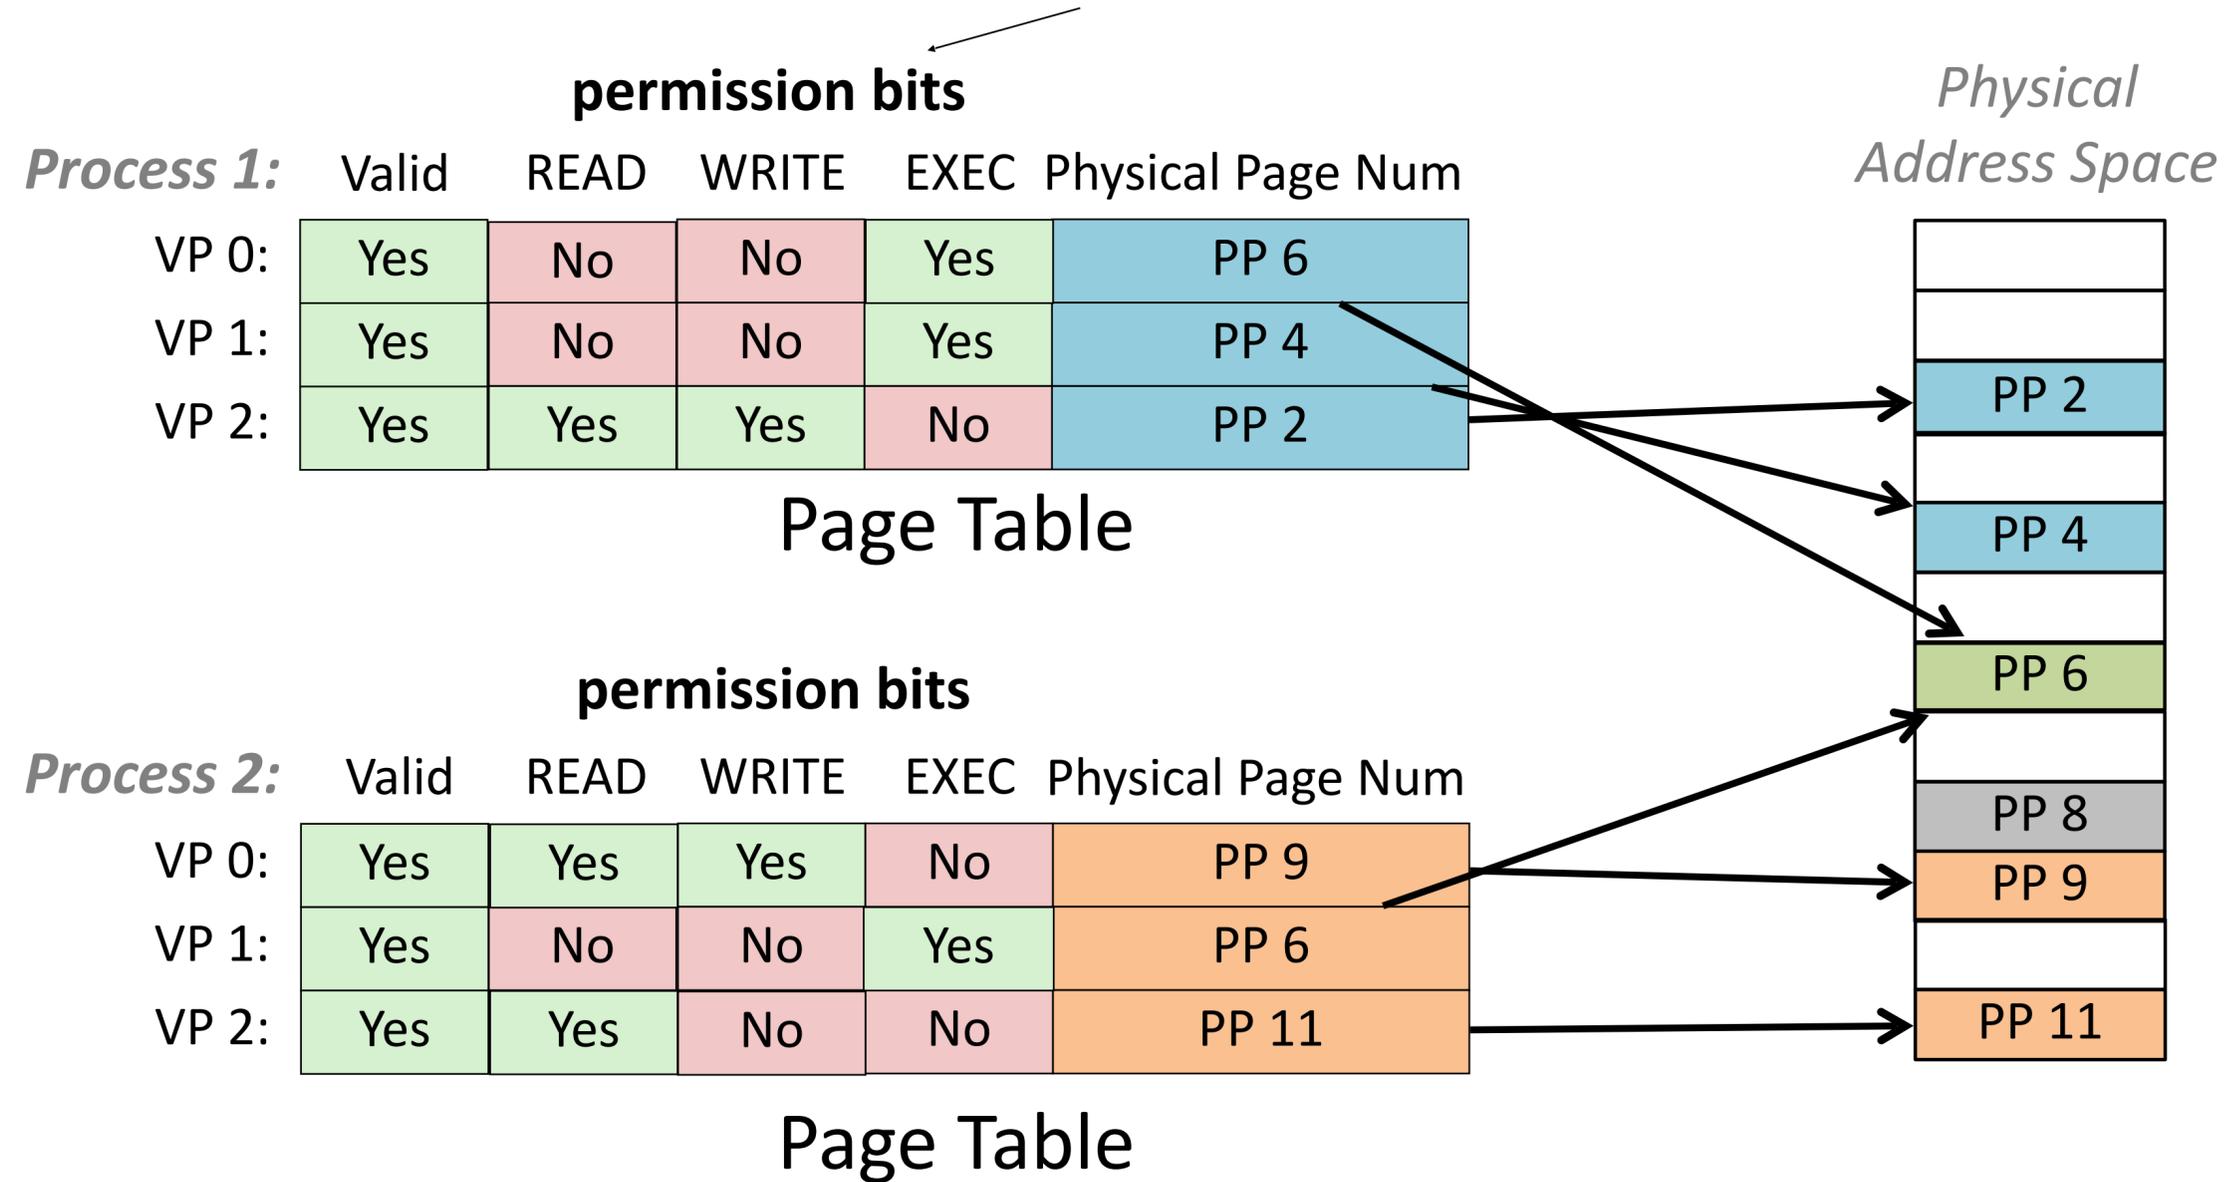
Map virtual pages in separate address spaces to same physical page (*PP 6*).



Virtual memory benefits:

Memory permissions

MMU checks on every access.
Exception if not allowed.



Summary: virtual memory

Programmer's view of virtual memory

Each process has its own private linear address space
Cannot be corrupted by other processes

System view of virtual memory

Uses memory efficiently (due to locality) by caching virtual memory pages

Simplifies memory management and sharing

Simplifies protection -- easy to interpose and check permissions

More goodies:

- Memory-mapped files
- Cheap `fork()` with copy-on-write pages (COW)

