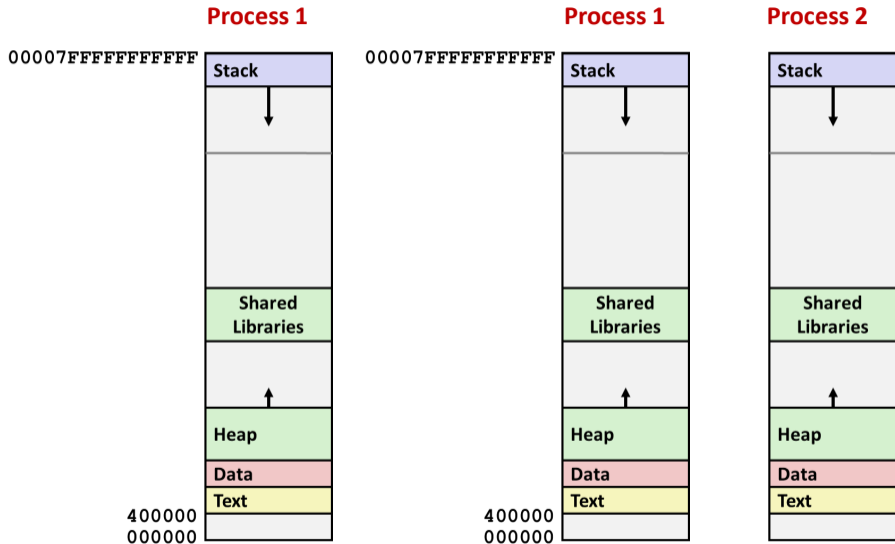


Addressing Memory

- We know that each byte of memory (RAM) in a computer has an *address*.
- We know that memory at an address can contain code (machine instructions) or data (bytes in some data format)
- We know (from previous class) that both code and data *may* be stored in a cache to speed up future accesses (thanks to locality)
- We've seen this first hand with a running program using GDB

How does this work?



How does this work?

Using a crucial system called *Virtual memory*

Let's introduce VM by starting with what we had before: *physical memory addressing*

- **Linear address space:** Ordered set of contiguous non-negative integer addresses:

$\{0, 1, 2, 3 \dots \}$

- **Virtual address space:** Set of $N = 2^n$ virtual addresses

$\{0, 1, 2, 3, \dots, N-1\}$

- **Physical address space:** Set of $M = 2^m$ physical addresses

$\{0, 1, 2, 3, \dots, M-1\}$

Why virtual memory (VM)?

- **Uses main memory efficiently**
 - Use DRAM as a cache for parts of a virtual address space
- **Simplifies memory management**
 - Each process gets the same uniform linear address space
- **Isolates address spaces**
 - One process can't interfere with another's memory
 - User program cannot access privileged kernel information and code

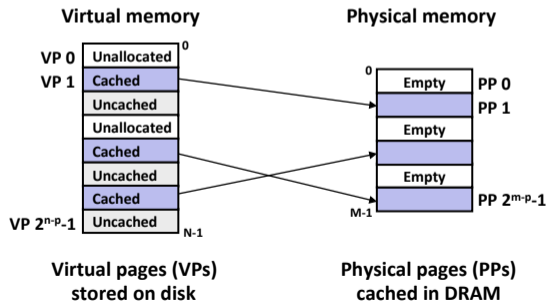
- Class notes

- ① Virtual Memory

- Address spaces
- VM as a tool for caching
- VM as a tool for memory management
- VM as a tool for memory protection
- How it works: address translation and memory mapping
- Summary

VM as a tool for caching

- Conceptually, **virtual memory** is an array of N contiguous bytes stored on disk.
- The contents of the array on disk are cached in **physical memory (DRAM cache)**
 - These cache blocks are called *pages* (size is $P = 2^p$ bytes)

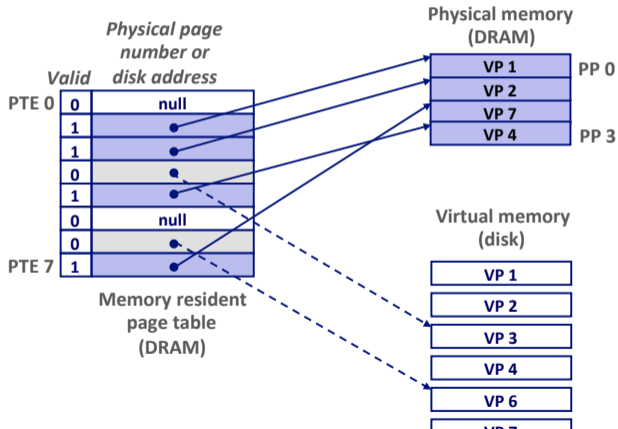


- **DRAM cache organization driven by the enormous miss penalty**
 - DRAM is about **10x** slower than SRAM
 - Disk is about **10,000x** slower than DRAM

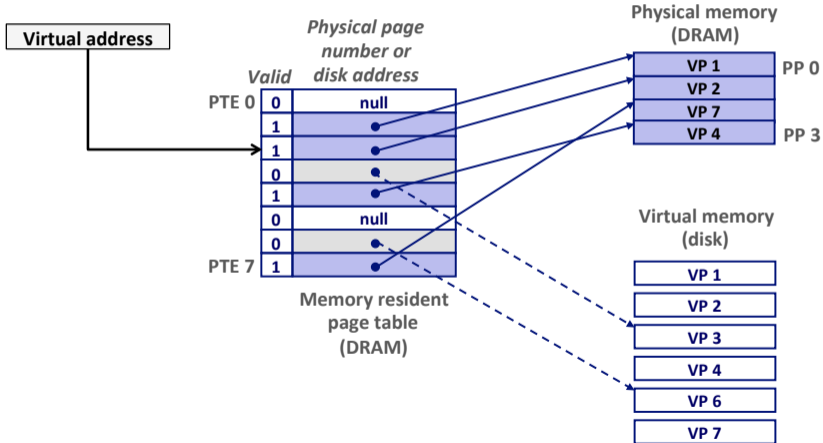
- **Consequences**
 - Large page (block) size: typically 4 KB, sometimes 4 MB
 - Fully associative
 - Any VP can be placed in any PP
 - Requires a “large” mapping function – different from cache memories
 - Highly sophisticated, expensive replacement algorithms
 - Too complicated and open-ended to be implemented in hardware
 - Write-back rather than write-through

Enabling data structure: page table

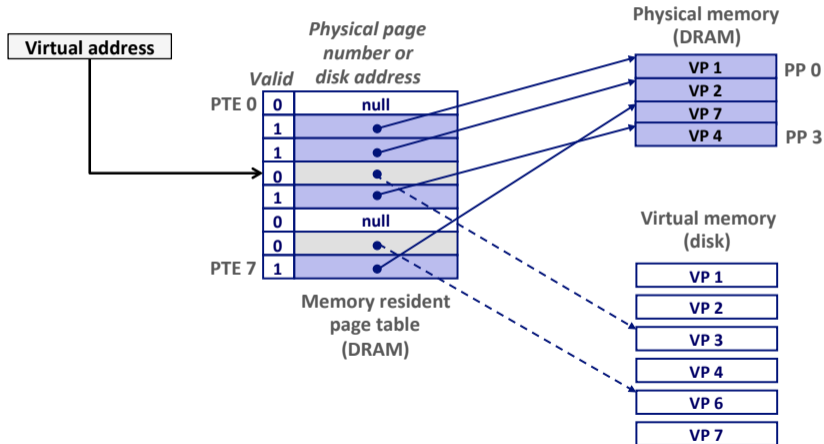
- A **page table** is an array of page table entries (PTEs) that maps virtual pages to physical pages.
 - Per-process kernel data structure in DRAM



- **Page hit:** reference to VM word that is in physical memory (DRAM cache hit)

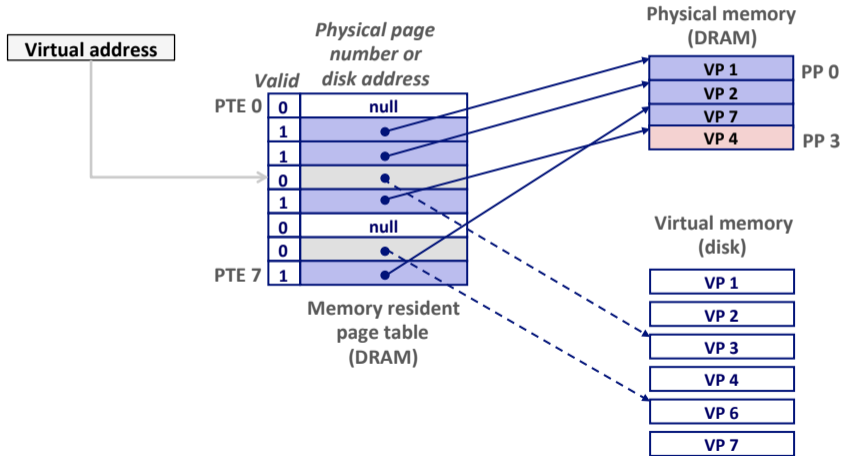


- **Page fault:** reference to VM word that is not in physical memory (DRAM cache miss)



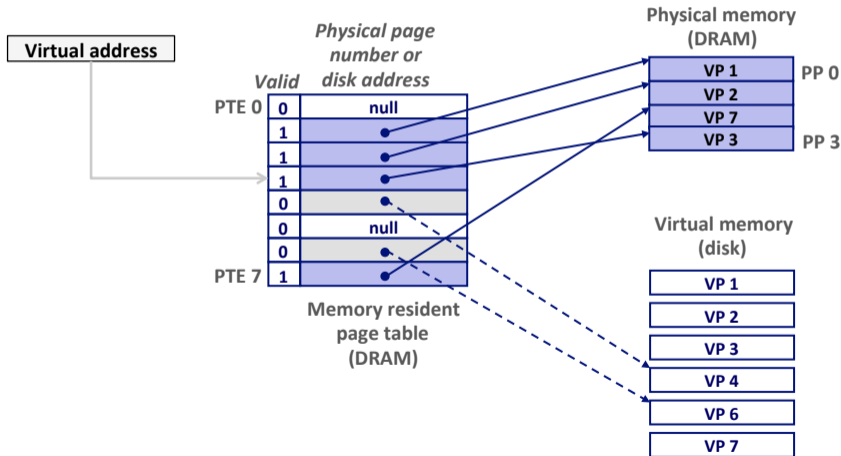
Handling page fault

- Page miss causes page fault (an exception)
- Page fault handler selects a victim to be evicted (here VP 4)



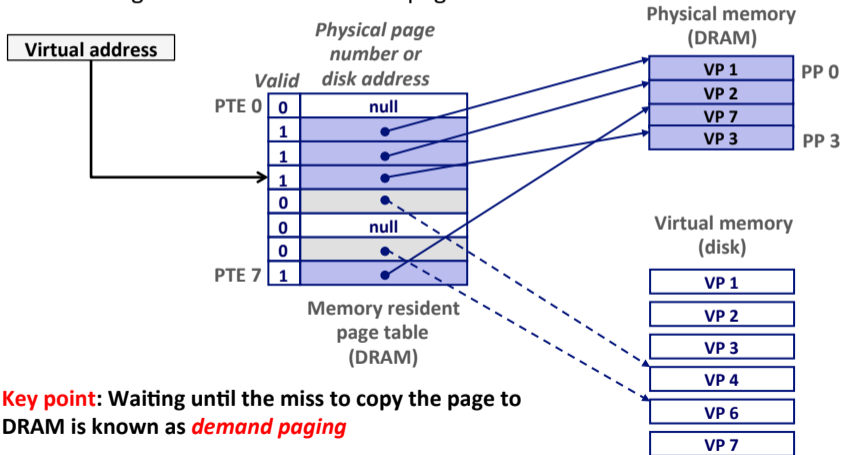
Handling page fault

- Page miss causes page fault (an exception)
- Page fault handler selects a victim to be evicted (here VP 4)



Handling page fault

- Page miss causes page fault (an exception)
- Page fault handler selects a victim to be evicted (here VP 4)
- Offending instruction is restarted: page hit!



Key point: Waiting until the miss to copy the page to DRAM is known as **demand paging**

Locality to the rescue! (again)

- Virtual memory seems terribly inefficient, but it works because of locality.
- At any point in time, programs tend to access a set of active virtual pages called the *working set*
 - Programs with better temporal locality will have smaller working sets
- If (working set size < main memory size)
 - Good performance for one process after compulsory misses
- If ($\text{SUM}(\text{working set sizes}) > \text{main memory size}$)
 - *Thrashing*: Performance meltdown where pages are swapped (copied) in and out continuously

- Class notes


- ① Virtual Memory

- Address spaces
- VM as a tool for caching
- VM as a tool for memory management
- VM as a tool for memory protection
- How it works: address translation and memory mapping
- Summary

VM as a tool for memory management

- **Key idea: each process has its own virtual address space**

- It can view memory as a simple linear array
- Mapping function scatters addresses through physical memory
 - Well-chosen mappings can improve locality




VM as a tool for memory management

■ Simplifying memory allocation

- Each virtual page can be mapped to any physical page
- A virtual page can be stored in different physical pages at different times

■ Sharing code and data among processes

- Map virtual pages to the same physical page (here: PP 6)



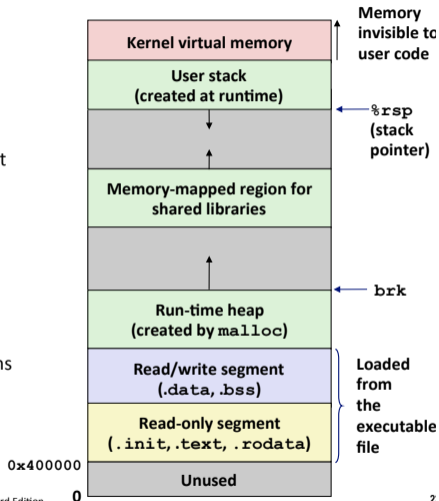
Simplifying linking and loading

■ Linking

- Each program has similar virtual address space
- Code, data, and heap always start at the same addresses.

■ Loading

- `execve` allocates virtual pages for `.text` and `.data` sections & creates PTEs marked as invalid
- The `.text` and `.data` sections are copied, page by page, on demand by the virtual memory system



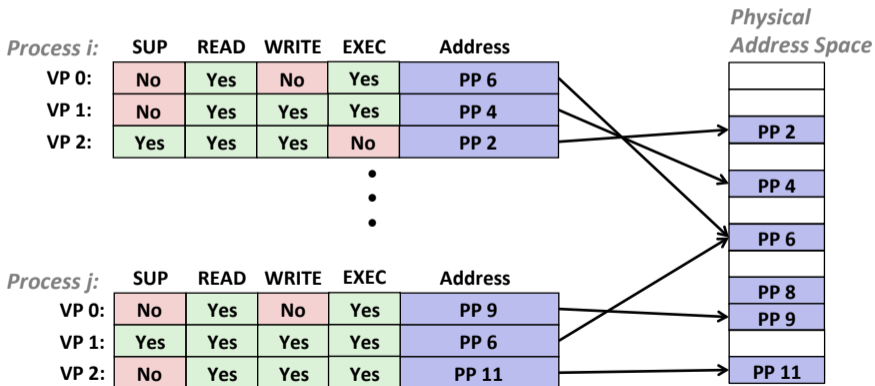
- Class notes

- ① Virtual Memory

- Address spaces
- VM as a tool for caching
- VM as a tool for memory management
- VM as a tool for memory protection
- How it works: address translation and memory mapping
- Summary

VM as a tool for memory protection

- Extend PTEs with permission bits
- MMU checks these bits on each access

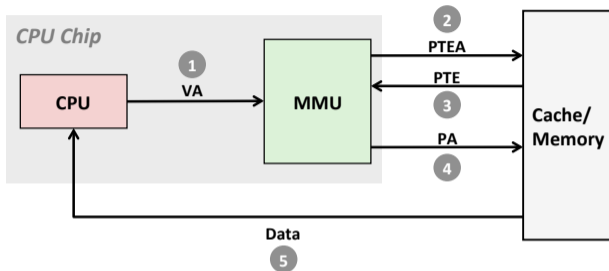


- Class notes

- ① Virtual Memory

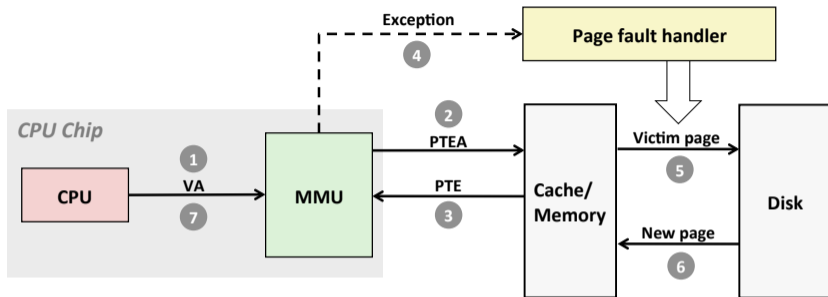
- Address spaces
- VM as a tool for caching
- VM as a tool for memory management
- VM as a tool for memory protection
- How it works: address translation and memory mapping
- Summary

Address translation: page hit



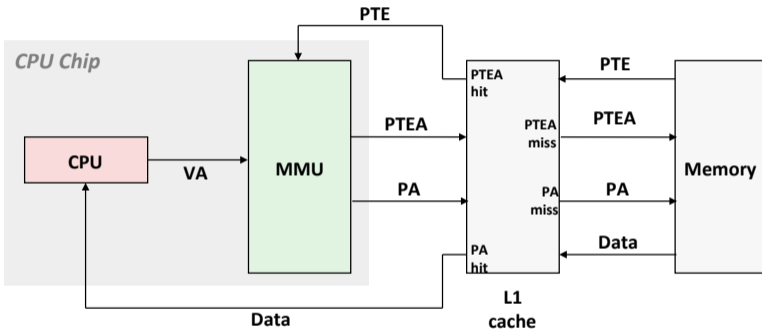
- 1) Processor sends virtual address to MMU
- 2-3) MMU fetches PTE from page table in memory
- 4) MMU sends physical address to cache/memory
- 5) Cache/memory sends data word to processor

Address translation: page fault



- 1) Processor sends virtual address to MMU
- 2-3) MMU fetches PTE from page table in memory
- 4) Valid bit is zero, so MMU triggers page fault exception
- 5) Handler identifies victim (and, if dirty, pages it out to disk)
- 6) Handler pages in new page and updates PTE in memory
- 7) Handler returns to original process, restarting faulting instruction

Integrating VM and cache

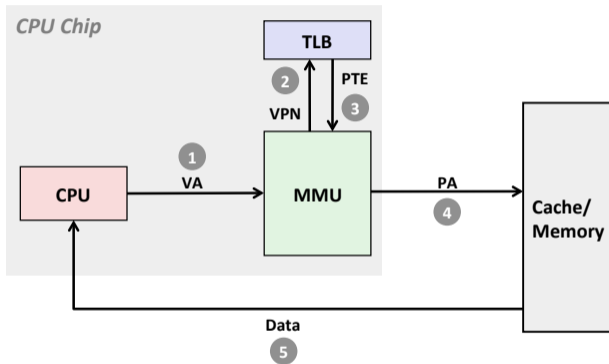


VA: virtual address, PA: physical address, PTE: page table entry, PTEA = PTE address

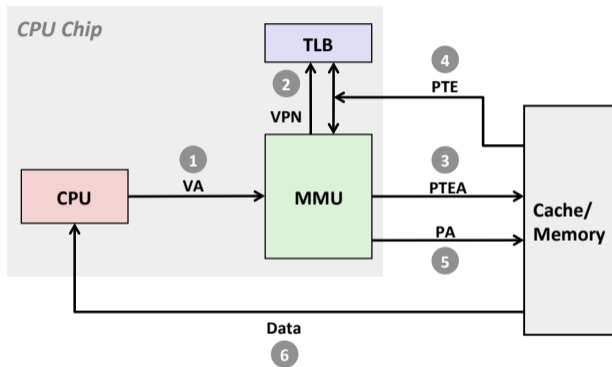
Speeding up translation with a TLB

- **Page table entries (PTEs) are cached in L1 like any other memory word**
 - PTEs may be evicted by other data references
 - PTE hit still requires a small L1 delay
- **Solution: *Translation Lookaside Buffer* (TLB)**
 - Small set-associative hardware cache in MMU
 - Maps virtual page numbers to physical page numbers
 - Contains complete page table entries for small number of pages

TLB Hit



A TLB hit eliminates a memory access



A TLB miss incurs an additional memory access (the PTE)

Fortunately, TLB misses are rare. Why?

■ **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 by caching virtual memory pages
 - Efficient only because of locality
- Simplifies memory management and programming
- Simplifies protection by providing a convenient interpositioning point to check permissions