# CS 105
"Tour of the Black Holes of Computing!"

Virtual Memory

Topics
- **Address translation**
- **Motivations for VM**
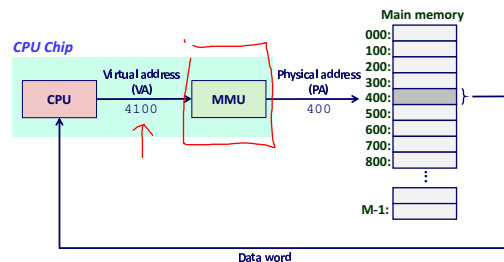- **Accelerating translation with TLBs**

Physically Addressed System

Used in "simple" systems like embedded microcontrollers in devices like car tires, elevators, and digital picture frames

---

Virtually Addressed System

Used in all modern servers, laptops, and smart phones

One of the great ideas in computer science

What Is Virtual Memory?

If you think it's there, and it is there…it's *real*.

If you think it's not there, and it really isn't there…it's *nonexistent*.

If you think it's not there, and it really is there…it's *transparent*.

If you think it's there, and it's not really there…it's *imaginary*.

Virtual memory is imaginary memory: it gives you the illusion of a memory arrangement that's not physically there.

Address Spaces

Linear address space: Ordered set of contiguous non-negative integer addresses:
$\{0, 1, 2, 3 \ldots\}$

Virtual address space: Set of $N = 2^n$ virtual addresses
$\{0, 1, 2, 3, \ldots, N-1\}$

$M \neq N$

Physical address space: Set of $M = 2^m$ physical addresses (may have "holes")
$\{0, 1, 2, 3, \ldots, M-1\}$

Clean distinction between data (bytes) and their attributes (addresses)
Every byte in main memory has *one* physical address and *zero or more* virtual addresses

---

Why Virtual Memory (VM)?

Uses main memory efficiently
- Use DRAM as a cache for parts of a large 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 can't access privileged kernel information and code


VM as 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)

4 Kbytes

Virtual memory   **Physical memory**

VP 0 Unallocated 0	0 Empty PP 0
VP 1 Cached	Empty PP 1
Uncached	
Unallocated	Empty
Cached	
Uncached	Empty
Cached	M-1 PP 2^{m-p}-1
VP 2^{n-p}-1 Uncached N-1	

Virtual pages (VPs) stored on disk   **Physical pages (PPs) cached in DRAM**

DRAM "Cache" Organization

DRAM (main memory) cache organization driven by the enormous miss penalty
- DRAM is about *10x* slower than SRAM   L1
- Hard disk is about *10,000x* slower than DRAM

Consequences
- Large page (block) size: typically 4-8 KB, sometimes 4 MB
- Fully associative
  - Any VP can be placed in any PP
  - Requires a "large" mapping function—different from CPU caches
- 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.
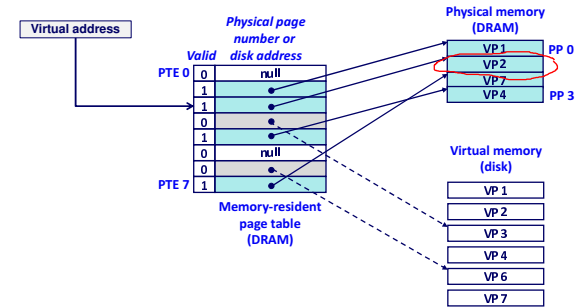
- Per-process kernel data structure in DRAM
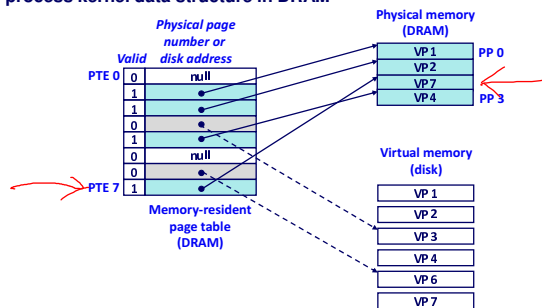
Page Hit

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

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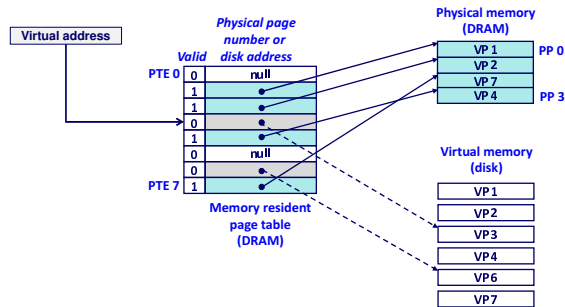
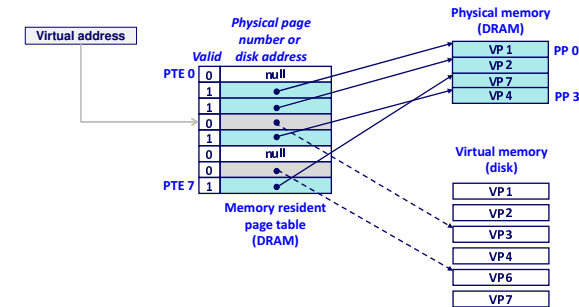*Page fault:* reference to VM word that is not in physical memory (DRAM cache miss)

Handling Page Fault

Page miss causes page fault (an exception)

Handling Page Fault

Page miss causes page fault (an exception)

Page fault handler selects a victim to be evicted (here VP 4)

Virtual address

Physical page number or disk address

Valid

PTE 0

null

PTE 7

Memory resident page table (DRAM)

Physical memory (DRAM)

VP 1 — PP 0
VP 2
VP 7
VP 4 — PP 3

Virtual memory (disk)

VP 1
VP 2
VP 3
VP 4
VP 6
VP 7
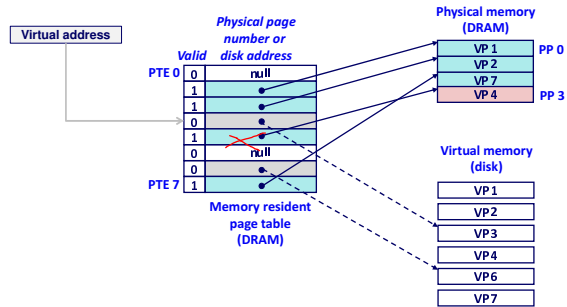
– 13 –

CS 105

Handling Page Fault

Page miss causes page fault (an exception)

Page fault handler selects a victim to be evicted (here VP 4)

Virtual address

Physical page number or disk address

Valid

PTE 0

null

PTE 7

Memory resident page table (DRAM)

Physical memory (DRAM)

VP 1 — PP 0
VP 2
VP 7
VP 3 — PP 3
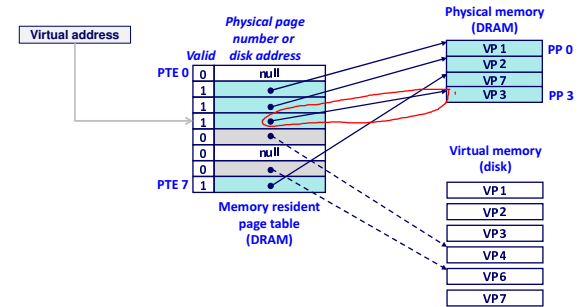
Virtual memory (disk)
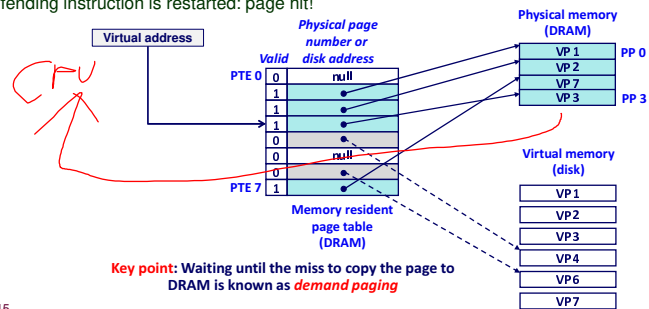
VP 1
VP 2
VP 3
VP 4
VP 6
VP 7

– 14 –

CS 105

Handling Page Fault

Page miss causes page fault (an exception)

Page fault handler selects a victim to be evicted (here VP 4)

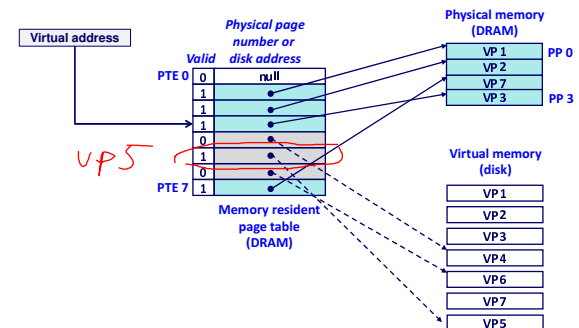Offending instruction is restarted: page hit!

CPU

Virtual address

Physical page number or disk address

Valid

PTE 0

null

PTE 7

Memory resident page table (DRAM)

Physical memory (DRAM)

VP 1 — PP 0
VP 2
VP 7
VP 3 — PP 3

Virtual memory (disk)

VP 1
VP 2
VP 3
VP 4
VP 6
VP 7

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

– 15 –

CS 105

Allocating Pages

Allocating a new page (VP 5) of virtual memory.

Virtual address

Physical page number or disk address

Valid

PTE 0

null

PTE 7

Memory resident page table (DRAM)

VP5

Physical memory (DRAM)

VP 1 — PP 0
VP 2
VP 7
VP 3 — PP 3

Virtual memory (disk)

VP 1
VP 2
VP 3
VP 4
VP 6
VP 7
VP 5

– 16 –

CS 105

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 SUM(working set sizes) > main memory size
- _Thrashing:_ Performance meltdown where pages are swapped (copied) in and out continuously

---

VM as Tool for Memory Management

Key idea: each process has own virtual address space
- Can view memory as a simple linear array
- Mapping function scatters addresses through physical memory
  - Well-chosen mappings can improve locality in L1-L3 caches


VM as Tool for Memory Management

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)
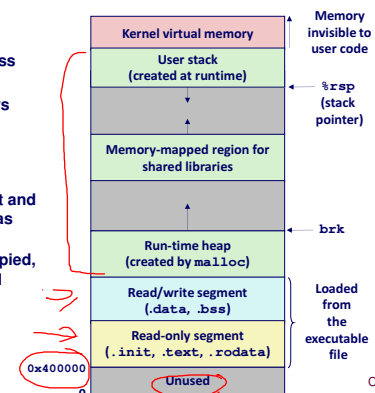

Simplifying Linking and Loading

Linking
- Each program has similar virtual address space
- Code, stack, and shared libraries always start at same virtual address

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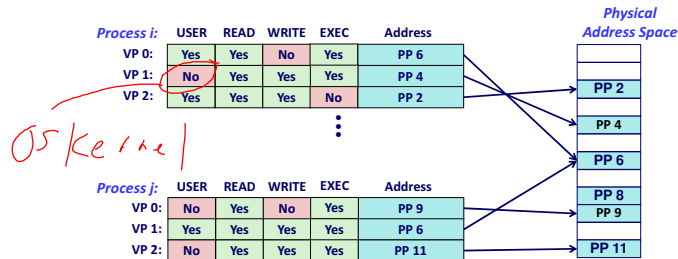- The `.text` and `.data` sections are copied, page by page, on demand by the virtual memory system

Kernel virtual memory — Memory invisible to user code

User stack (created at runtime) — `%rsp` (stack pointer)

Memory-mapped region for shared libraries

Run-time heap (created by `malloc`) — `brk`

Read/write segment (`.data`, `.bss`)

Read-only segment (`.init`, `.text`, `.rodata`)

Loaded from the executable file

0x400000

Unused

0

VM as Tool for Memory Protection

Extend PTEs (page table entries) with permission bits

Page-fault handler checks these before remapping
- If violated, send process SIGSEGV (segmentation fault)

Process i:	USER	READ	WRITE	EXEC	Address
VP 0:	Yes	Yes	No	Yes	PP 6
VP 1:	No	Yes	Yes	Yes	PP 4
VP 2:	Yes	Yes	Yes	No	PP 2

OS Kernel

Process j:	USER	READ	WRITE	EXEC	Address
VP 0:	No	Yes	No	Yes	PP 9
VP 1:	Yes	Yes	Yes	Yes	PP 6
VP 2:	No	Yes	Yes	Yes	PP 11

Physical Address Space

PP 2
PP 4
PP 6
PP 8
PP 9
PP 11

VM Address Translation

Virtual Address Space
- $V = \{0, 1, ..., N–1\}$

Physical Address Space
- $P = \{0, 1, ..., M–1\}$

$N \neq M$

Address Translation
- $MAP: V \rightarrow P \cup \{\varnothing\}$
- For virtual address a:
  - $MAP(a) = a'$ if data at virtual address a is at physical address a' in P
  - $MAP(a) = \varnothing$ if data at virtual address a is not in physical memory
    - » Either invalid or stored on disk

Address-Translation Symbols

Basic Parameters
- $N = 2^n$: Number of addresses in virtual address space
- $M = 2^m$: Number of addresses in physical address space
- $P = 2^p$: Page size (bytes) $4095 = 2^{12}$
- $M' = M / P$: Number of physical pages

Components of the virtual address (VA)
- VPN: Virtual page number
- VPO: Virtual page offset
- TLBI: TLB index
- TLBT: TLB tag

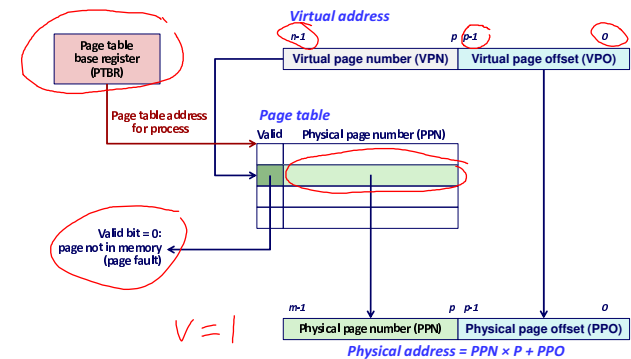**Components of the physical address (PA)**
- PPN: Physical page number
- PPO: Physical page offset (same as VPO)
- PTE: Page table entry
- PTEA: Address of PTE

> There's a bunch of these:
> It'll take time to learn them.
> The highlighted ones are
> the 8 most important.

Address Translation With a Page Table

Virtual address

Page table base register (PTBR)

Page table address for process

Page table

Valid bit = 0: page not in memory (page fault)

$V = 1$

Physical page number (PPN) | Physical page offset (PPO)

Physical address = PPN × P + PPO

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" (reads) 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 but significant L1 delay (3-4 cycles)
  - Net effect is to double time needed to access data in L1 cache!

Solution: *Translation Lookaside Buffer* (TLB)

- Tiny set-associative (or fully associative) hardware cache inside MMU
- Maps virtual page numbers to physical page numbers
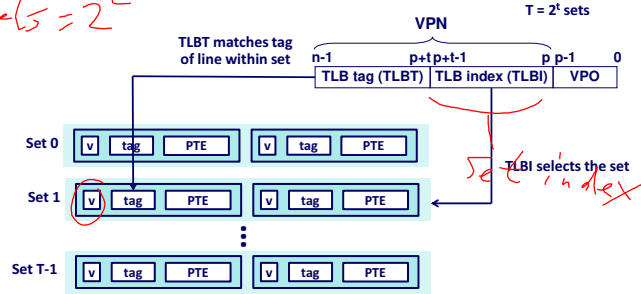- Contains complete page table entries for small number of pages

Accessing the TLB

**MMU uses the VPN portion of the virtual address to access the TLB:**

T sets = 2^t (handwritten)

$T = 2^t$ sets

TLBT matches tag of line within set

VPN

n-1	p+t	p+t-1	p	p-1	0
TLB tag (TLBT)		TLB index (TLBI)		VPO	

Set index (handwritten)

TLBI selects the set

Set 0: | v | tag | PTE | | v | tag | PTE |

Set 1: | v | tag | PTE | | v | tag | PTE |

Set T-1: | v | tag | PTE | | v | tag | PTE |

TLB Hit

CPU Chip

CPU — VA ①→ MMU — PA ④→ Cache/Memory

TLB ② VPN / ③ PTE

Data ⑤

A TLB hit eliminates a cache or memory access to get the PTE

TLB Miss

CPU Chip

CPU — VA ①→ MMU

TLB ② VPN / ④ PTE

MMU — ③ PTEA / ⑤ PA → Cache/Memory

Data ⑥

A TLB miss incurs an additional memory access (the PTE)
Fortunately, TLB misses are rare. Why?

Multi-Level Page Tables

Suppose:
- 4KB (2^{12}) page size, 48-bit virtual address space, 8-byte PTE

Problem:
- Would need a 512 GB page table!
  - $2^{48} * 2^{-12} * 2^3 = 2^{39}$ bytes

Common solution: Multi-level page table

Example: 2-level page table
- Level 1 table (always memory-resident): each PTE points to a page table
- Level 2 table (paged in and out like any other data): each PTE points to a page
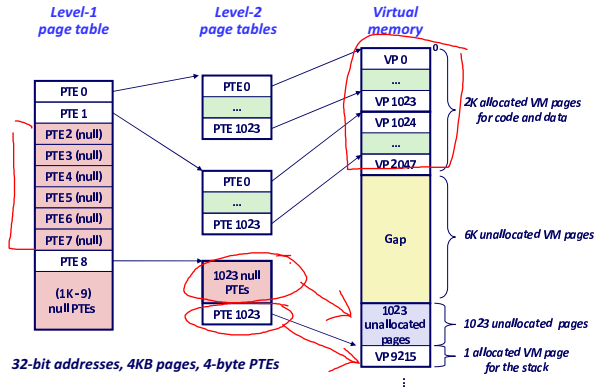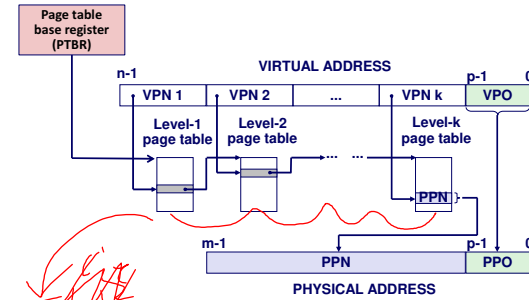
Level-1 Table

Level-2 Tables

A Two-Level Page Table Hierarchy

32-bit addresses, 4KB pages, 4-byte PTEs

– 33 –

CS 105

Translating With a k-level Page Table

– 34 –

CS 105

Summary

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

– 35 –

CS 105