


# Lecture 17: Virtual Memory (cont'd)

---

CS 105

Spring 2026

# Review: Address Translation



Virtual Address


MMU

invalid

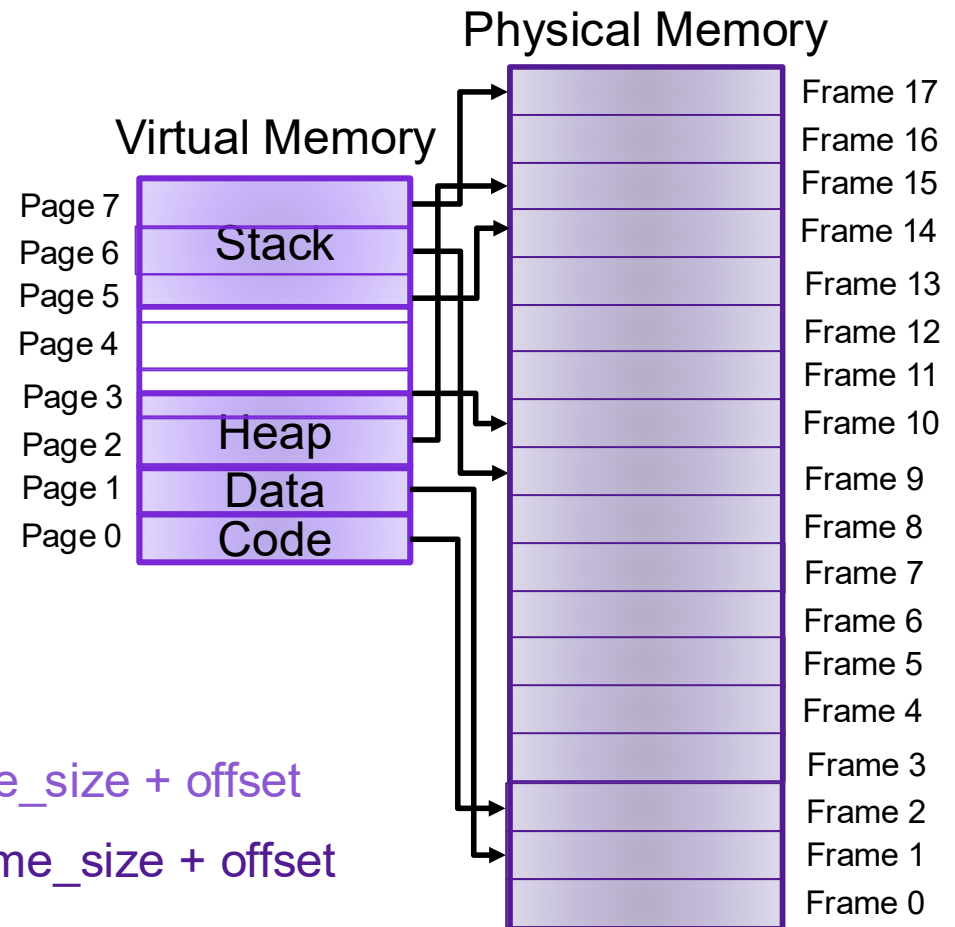
Exception

Physical Address

Data



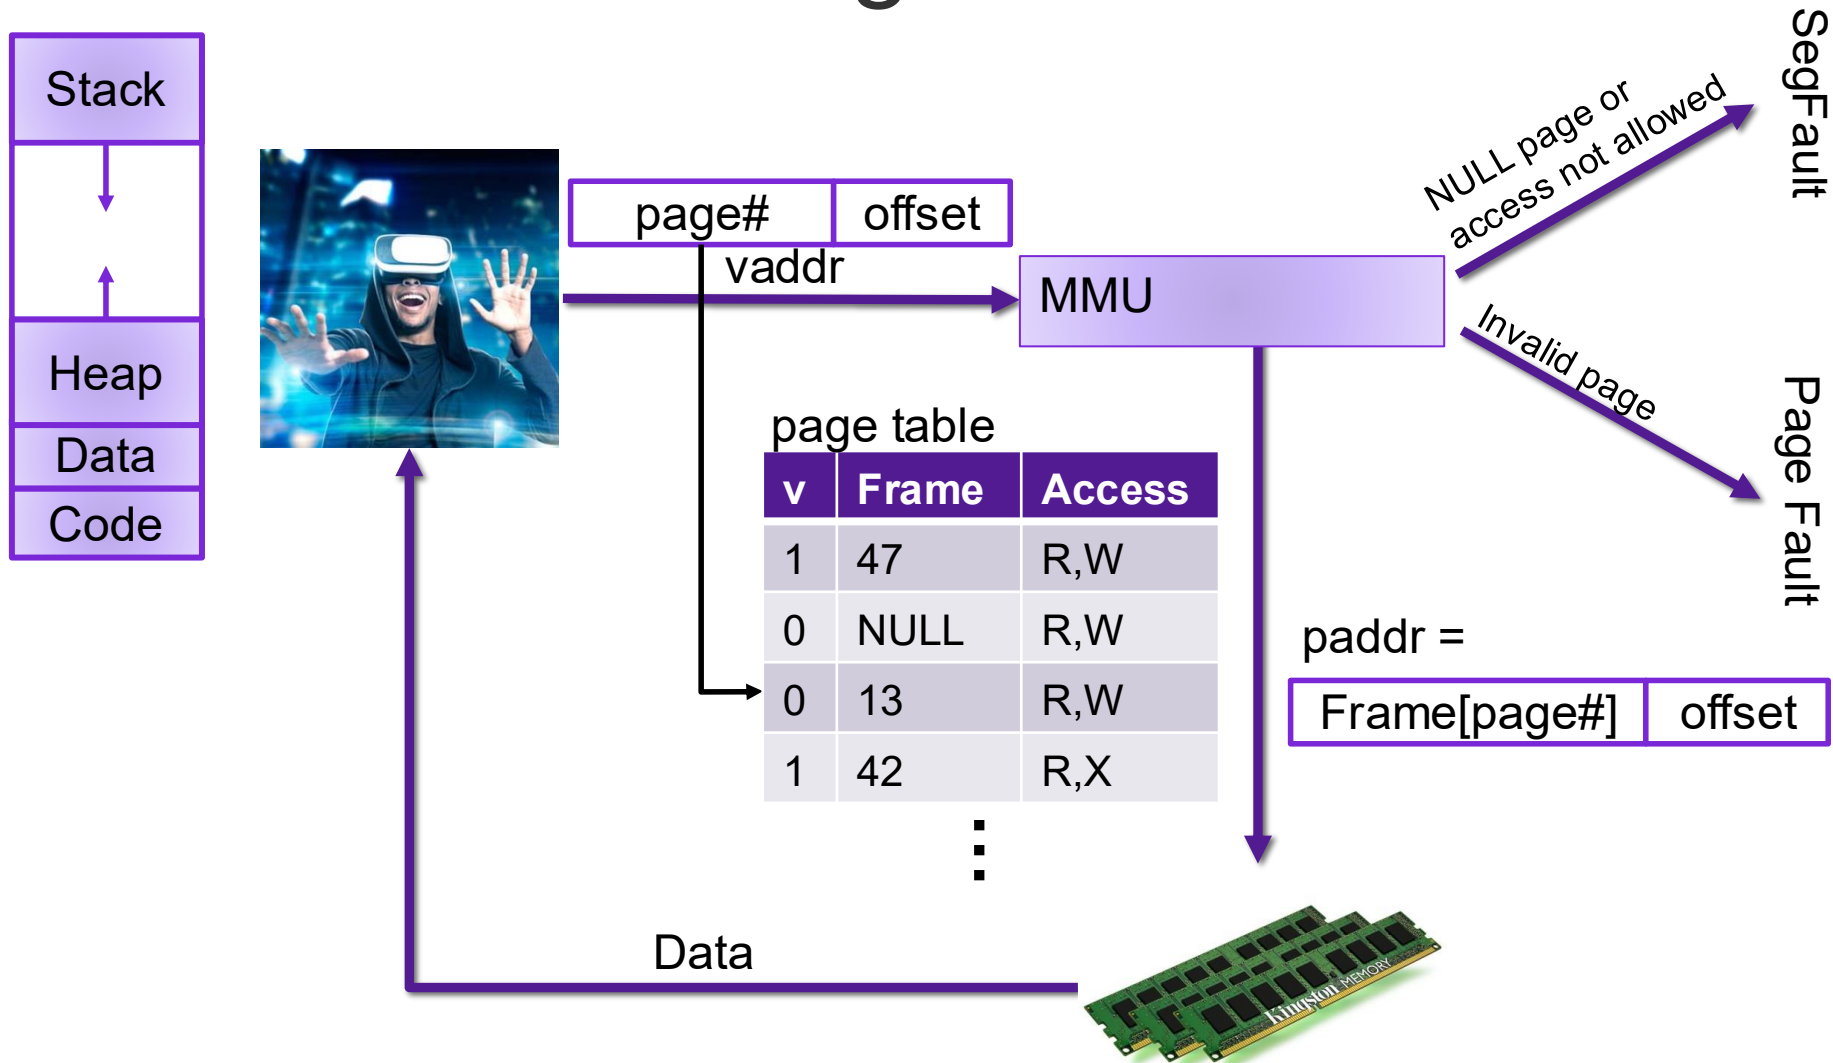
# Review: Paging



$vaddr = page\_num * page\_size + offset$

$paddr = frame\_num * frame\_size + offset$

# Review: Virtual Pages



# Review: Evaluating Paging

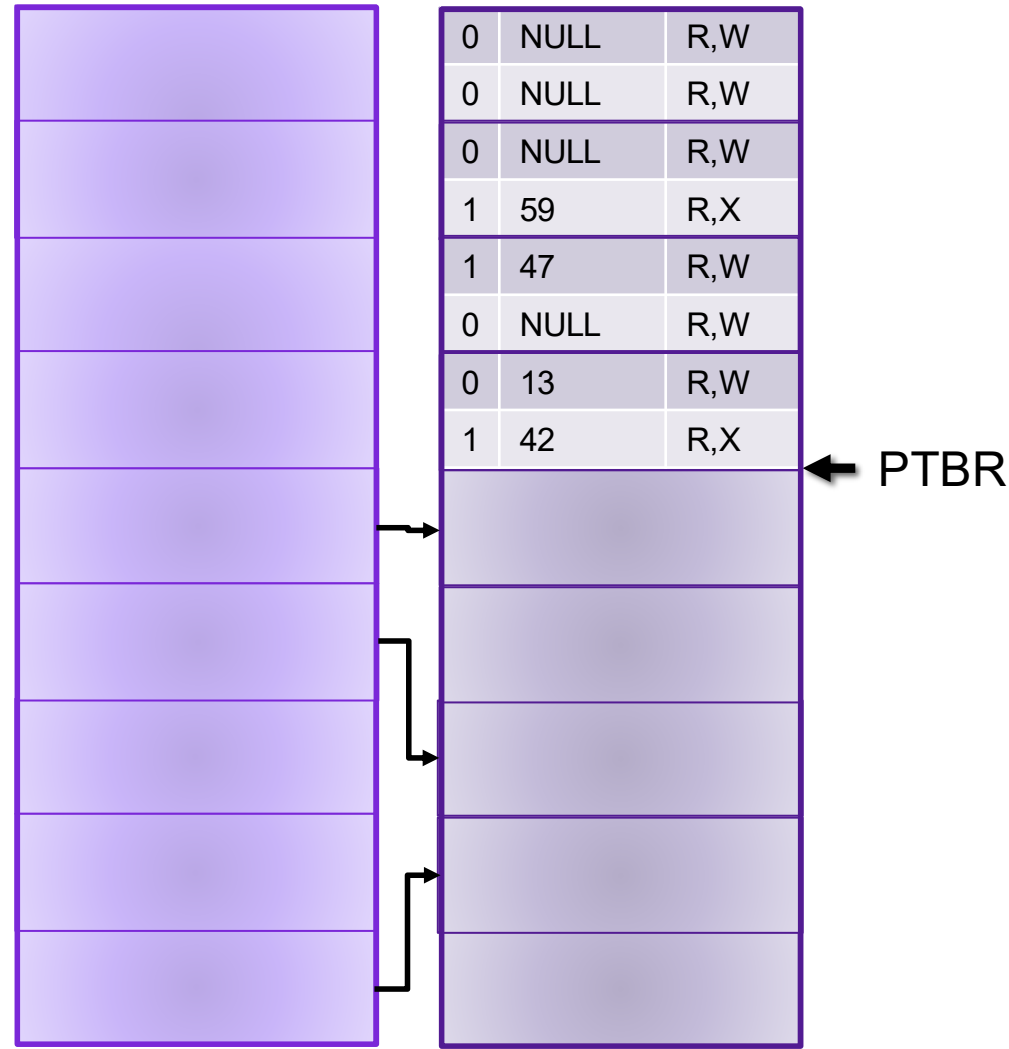


- **Isolation:** don't want different process states collided in physical memory
- **Efficiency:** want fast reads/writes to memory
- **Sharing:** want option to overlap for communication
- **Utilization:** want best use of limited resource
- **Virtualization:** want to create illusion of more resources



# Traditional Paging

- page table is stored in physical memory
- implemented as array of page table entries
- Page Table Base Register (PTBR) stores physical address of beginning of page table
- Page table entries are accessed by using the page number as the index into the page table

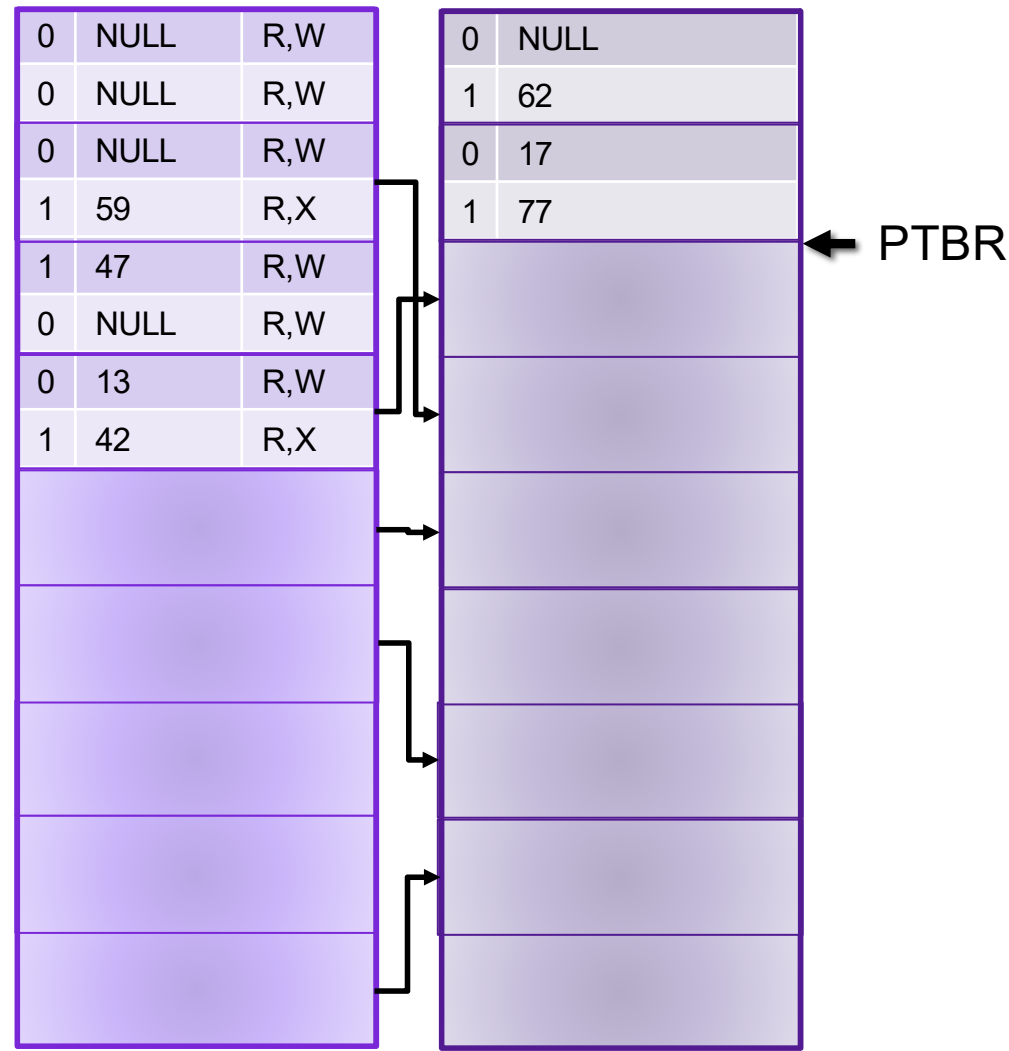


# Problems with Paging

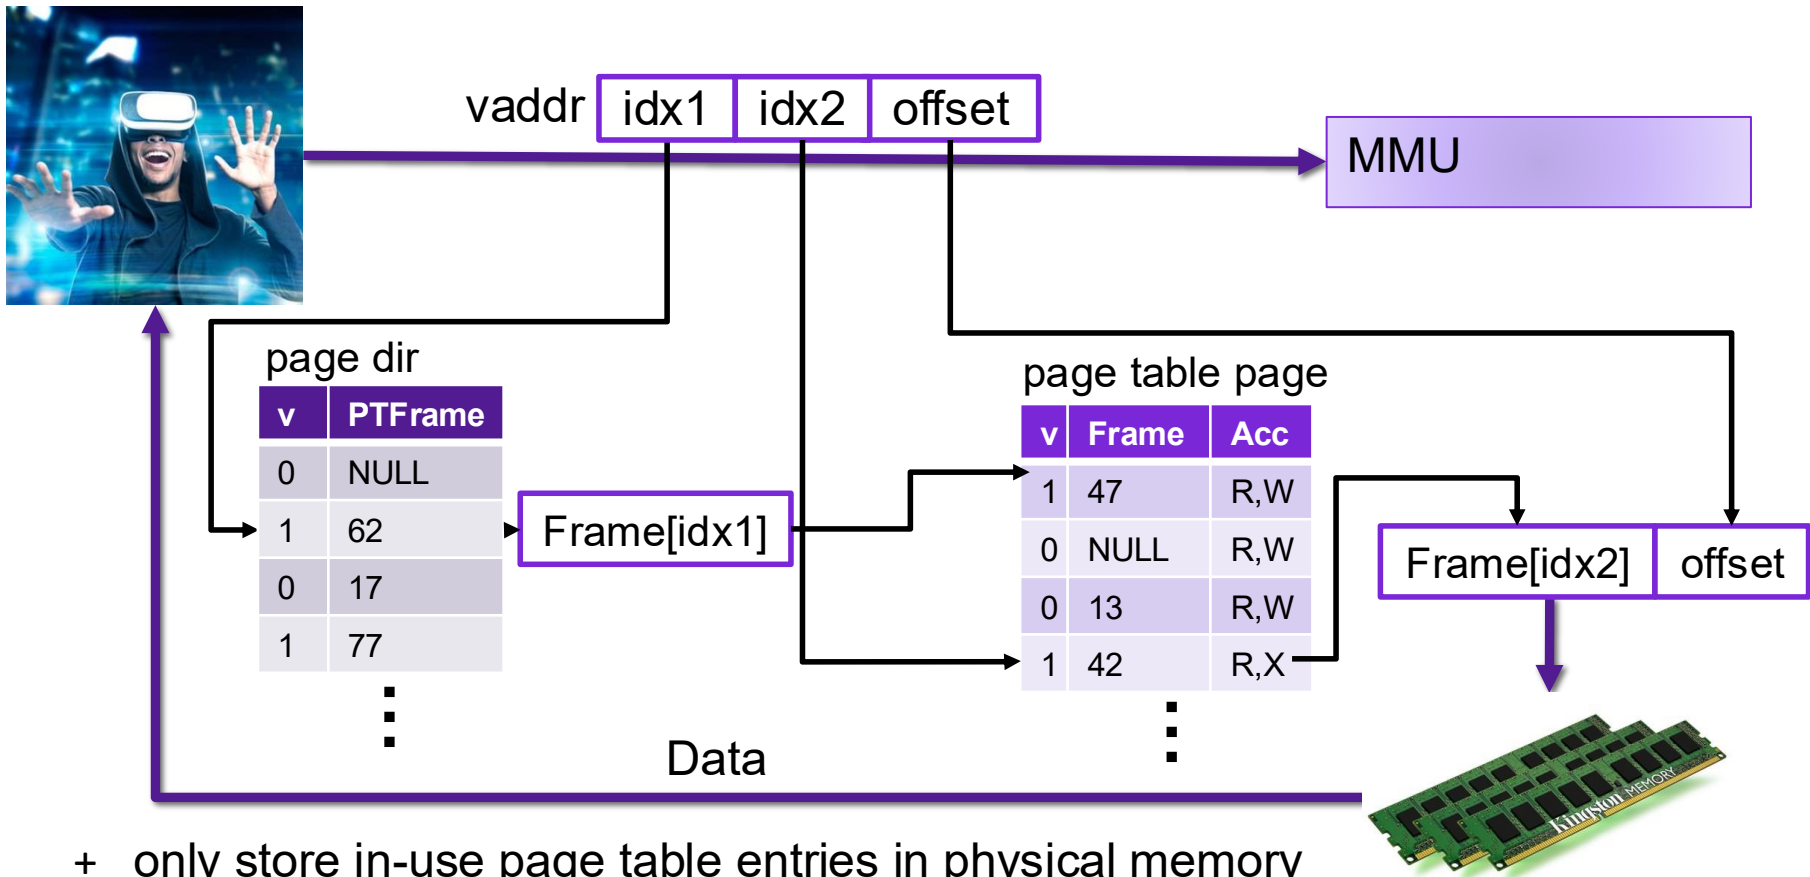
- **Memory Consumption:** page table is really big
  - Example: consider 48-bit address space, 4KB ( $2^{12}$ ) page size, assume each page table entry is 8 bytes.
  - Larger pages increase internal fragmentation
- **Performance:** every data/instruction access requires *two* memory accesses:
  - One for the page table
  - One for the data/instruction

# Two-level Page Tables

- page table is stored in virtual memory pages
- page directory is stored in physical memory (page table for the page table)
- Implemented as array of page directory entries
- Page Table Base Register (PTBR) stores physical address of beginning of page directory



# Two-level Page Tables




- + only store in-use page table entries in physical memory
- + easier to allocate page table
- more memory accesses

# Example: Two-level Page Tables

Assume you are working on an architecture with a 32-bit virtual address space in which each page is 64 KB and a page table entry is 16 bytes.


- How many bits will be in the offset?  **$2^{16}$  bytes = 64 KB**  
**16 bits**
- How many bits will be in idx2? **12 bits**
- How many bits will be in idx1? **4 bits**






# Multi-level Page Tables

- Problem: How big does the page directory get? **1 GB**
  - Assume you have a 48-bit address space
  - Assume you have 4KB pages
  - Assume you have 8 byte page table entries/page directory entries



- Goal: Page Table Directory should fit in one frame
- **Multi-level page tables:** add additional level(s) to tree

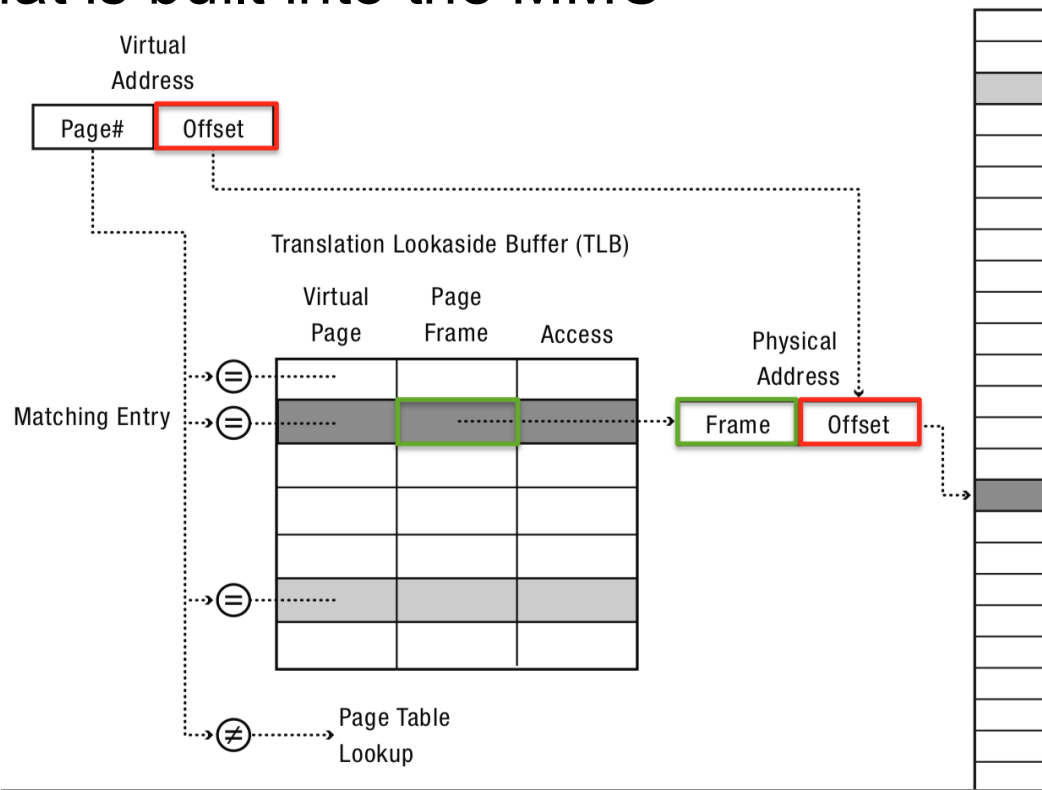


# Review: Problems with Paging

- **Memory Consumption:** page table is really big
  - Example: consider 64-bit address space, 4KB ( $2^{12}$ ) page size, assume each page table entry is 8 bytes.
  - Larger pages increase internal fragmentation
- **Performance:** every data/instruction access requires ~~two~~<sup>five</sup> memory accesses:
  - One for ~~the page table~~ each of the four levels of page table
  - One for the data/instruction

# Translation-Lookaside Buffer (TLB)

- General idea: if address translation is slow, cache some of the answers
- **Translation-lookaside buffer** is an address translation cache that is built into the MMU



# Exercise: TLB

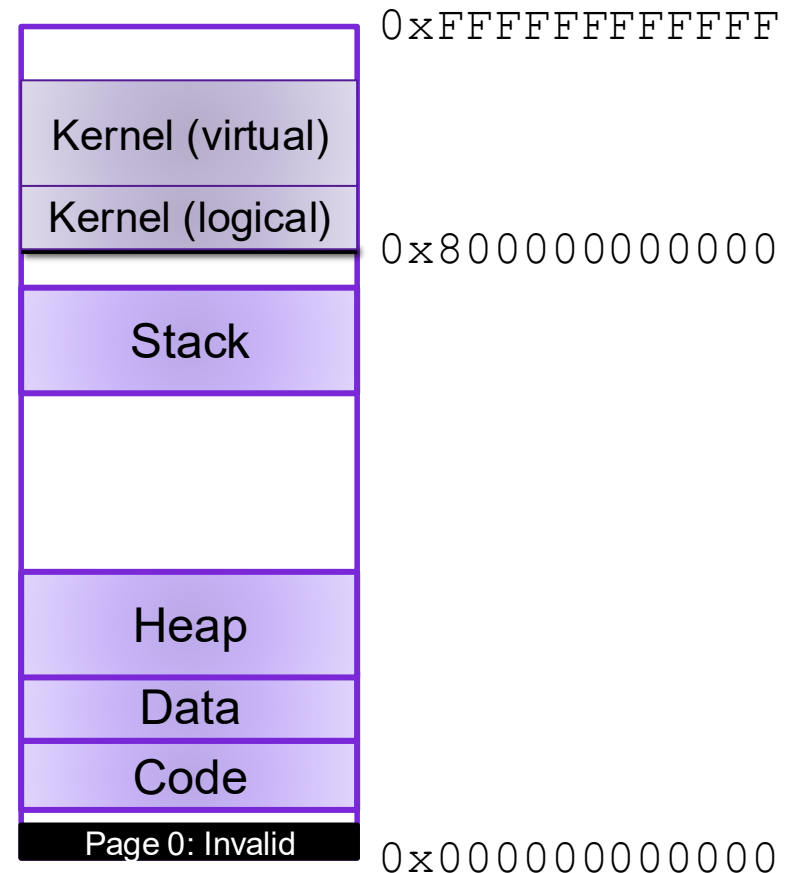
TLB												
idx	v	tag	PPN	v	tag	PPN	v	tag	PPN	v	tag	PPN
0	1	03	B	0	07	6	1	28	3	0	01	F
1	1	31	0	0	12	3	1	3E	4	1	0B	1
2	0	2A	A	0	11	1	1	1F	8	1	07	5
3	1	07	3	0	2A	A	0	1E	2	0	21	B

Assume you are running on an architecture with a one-level page table with 4096 byte pages. For each of the following virtual addresses, determine whether the address translation is stored in the TLB. If so, give the corresponding physical address

- 0x7E37C
- 0x16A48

# Example: The Linux x86 Address Space

- Use "only" 48-bit addresses (top 16 bits not used)
- 4KB pages by default
  - supports larger "superpages"
- Four-level page table
- Physical memory stores memory pages, memory-mapped files, cached file pages
- Page eviction algorithm uses variant of LRU called 2Q
  - approximates LRU with clock
  - maintains two lists (active/inactive)
- Stack is marked non-executable
- Virtual address of stack/heap start are randomized each time process is initialized



# Example: Core i7 Memory Accessing

