

Eliminating External Fragmentation: Paging

- Allocate VA & PA memory in **chunks of the same, fixed size** (**pages** and **frames**, respectively)

- Adjacent pages in VA (say, within the stack) need not map to contiguous frames in PA!

- Free frames can be tracked using **a simple bitmap**

- ▶ **0011111001111011110000** one bit/frame

- No more external fragmentation!

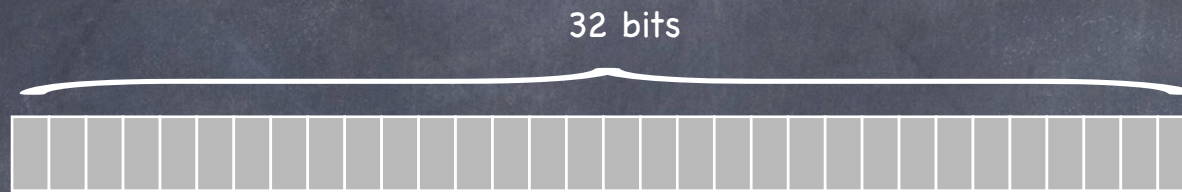
- But now **internal** fragmentation (you just can't win...)

- when memory needs are not a multiple of a page

- typical size of page/frame: 4KB to 16KB

How can I reference
a byte in VA space?

Virtual address



- Interpret VA as comprised of two components
 - **page:** which page?
 - **offset:** which byte within that page?

Virtual address



- Interpret VA as comprised of two components
 - **page:** which page?
 - ▶ no. of bits specifies no. of pages are in the VA space
 - **offset:** which byte within that page?

Virtual address



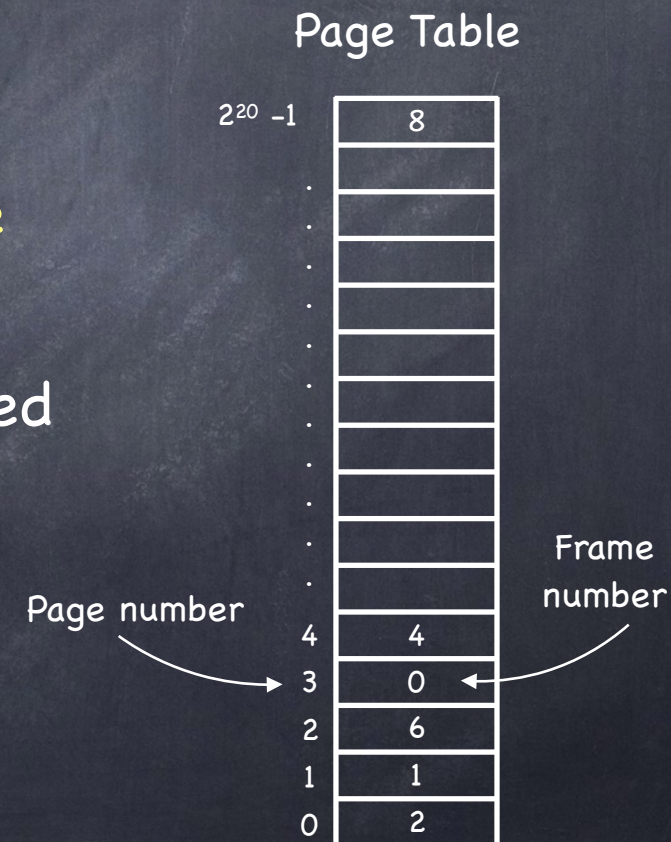
- Interpret VA as comprised of two components
 - **page:** which page?
 - ▶ no. of bits specifies no. of pages are in the VA space
 - **offset:** which byte within that page?
 - ▶ no. of bits specifies size of page/frame

Virtual address

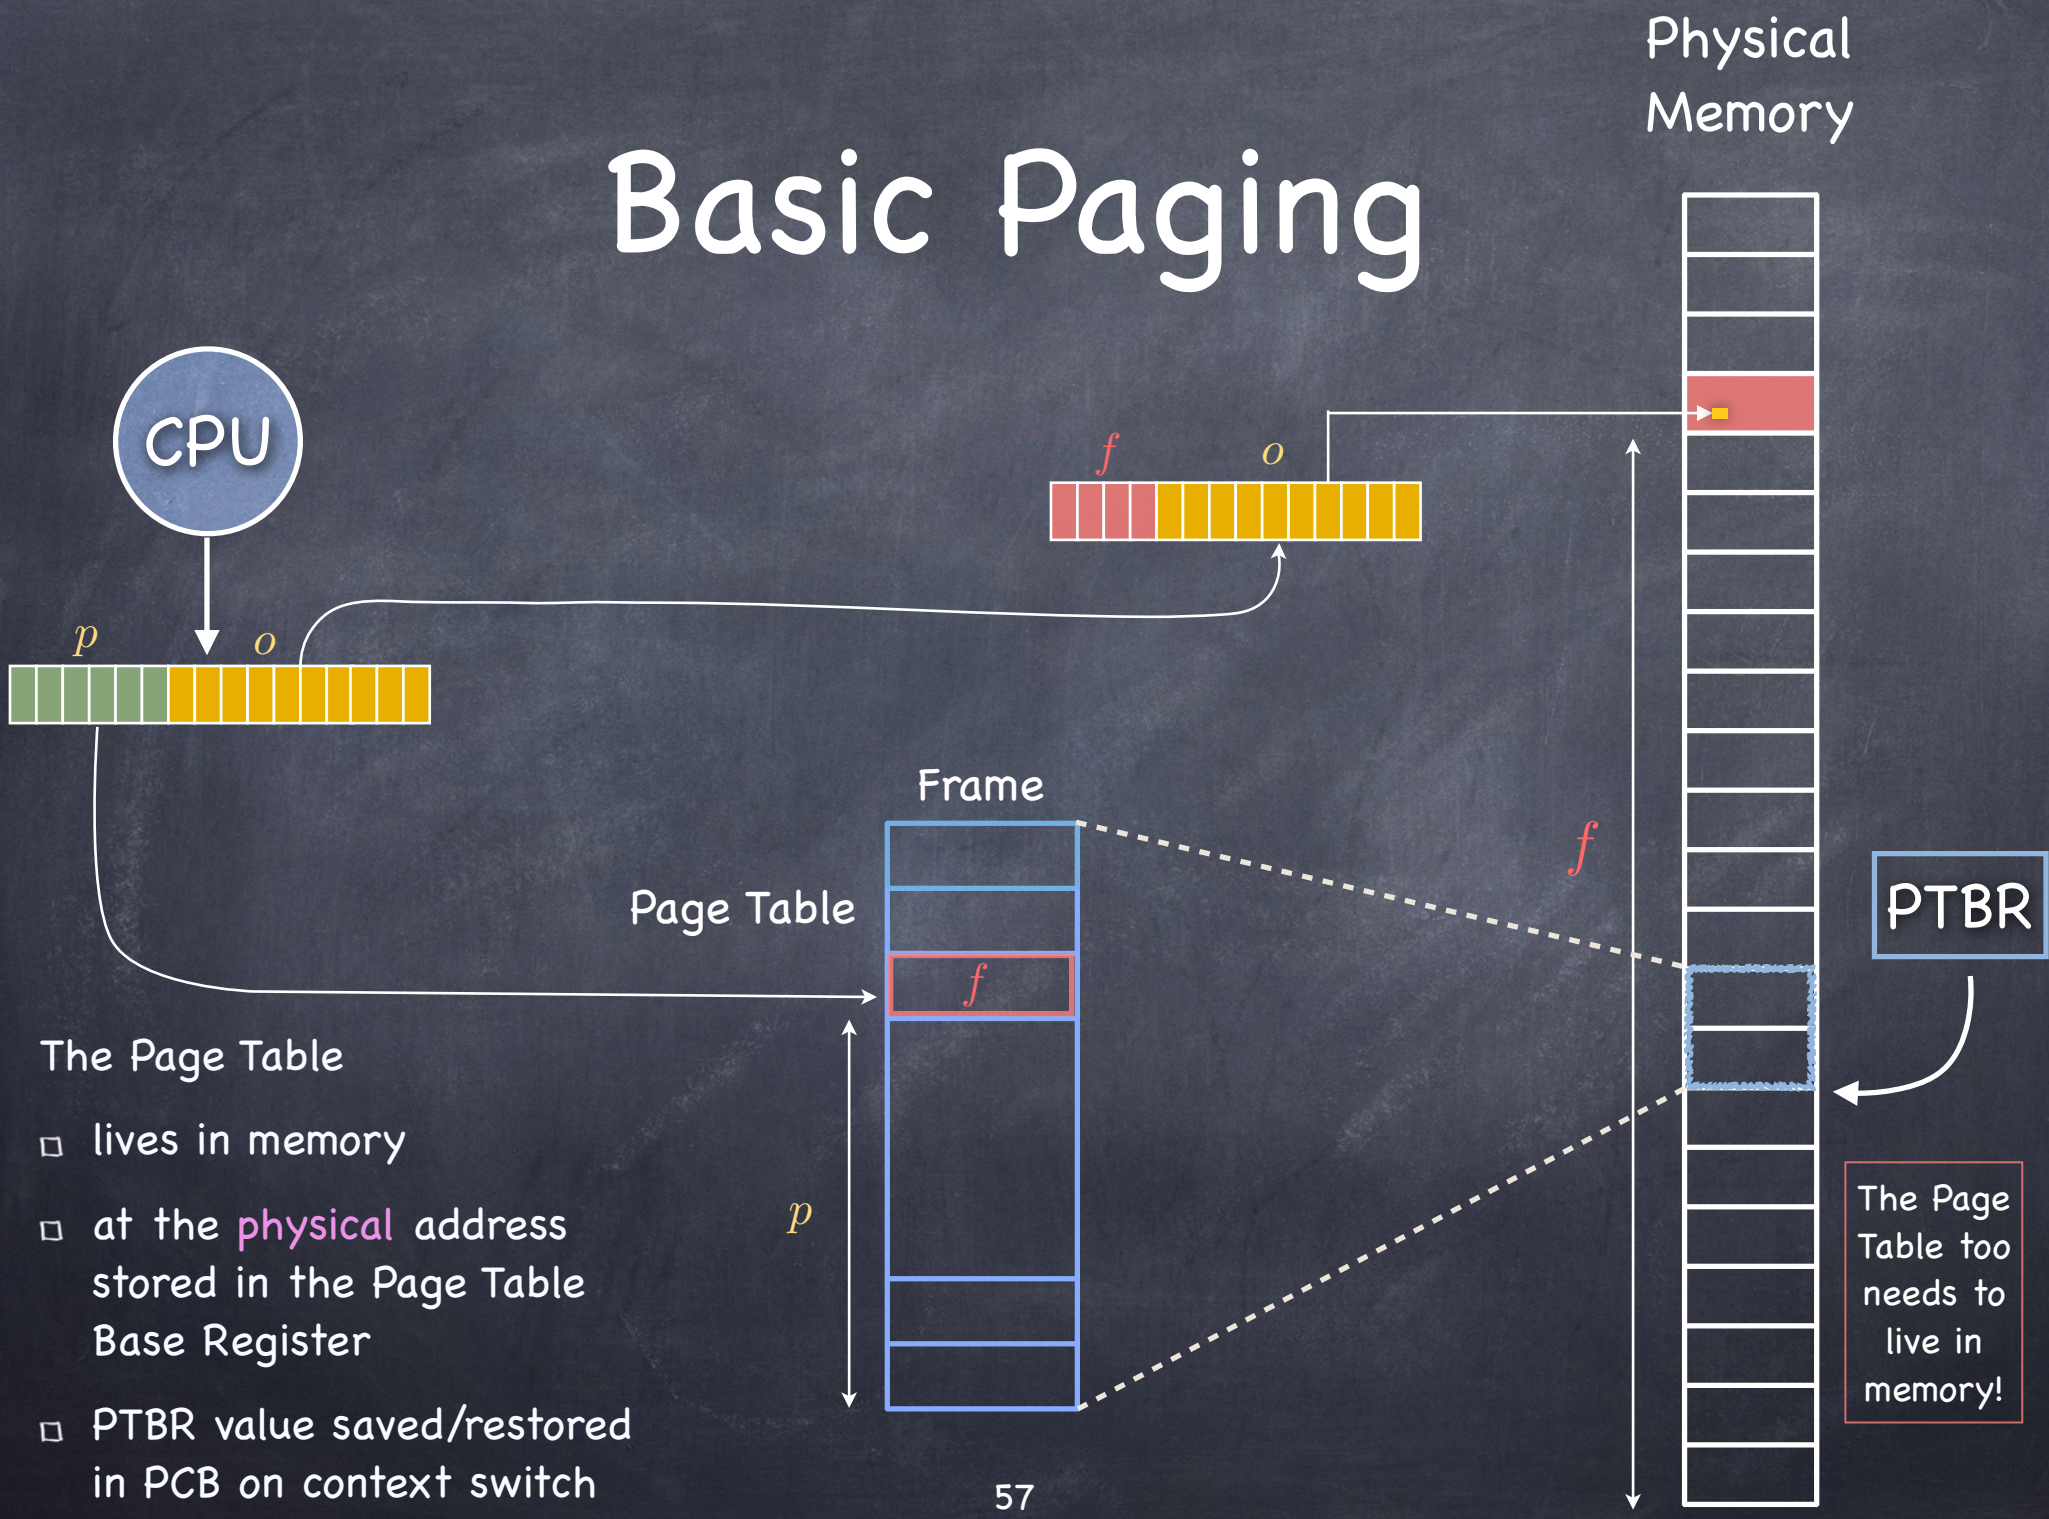


To access a byte

- extract page number
- map that page number into a frame number using a page table
 - ▶ **Note:** not all pages may be mapped to frames
- extract offset
- access byte at offset in frame



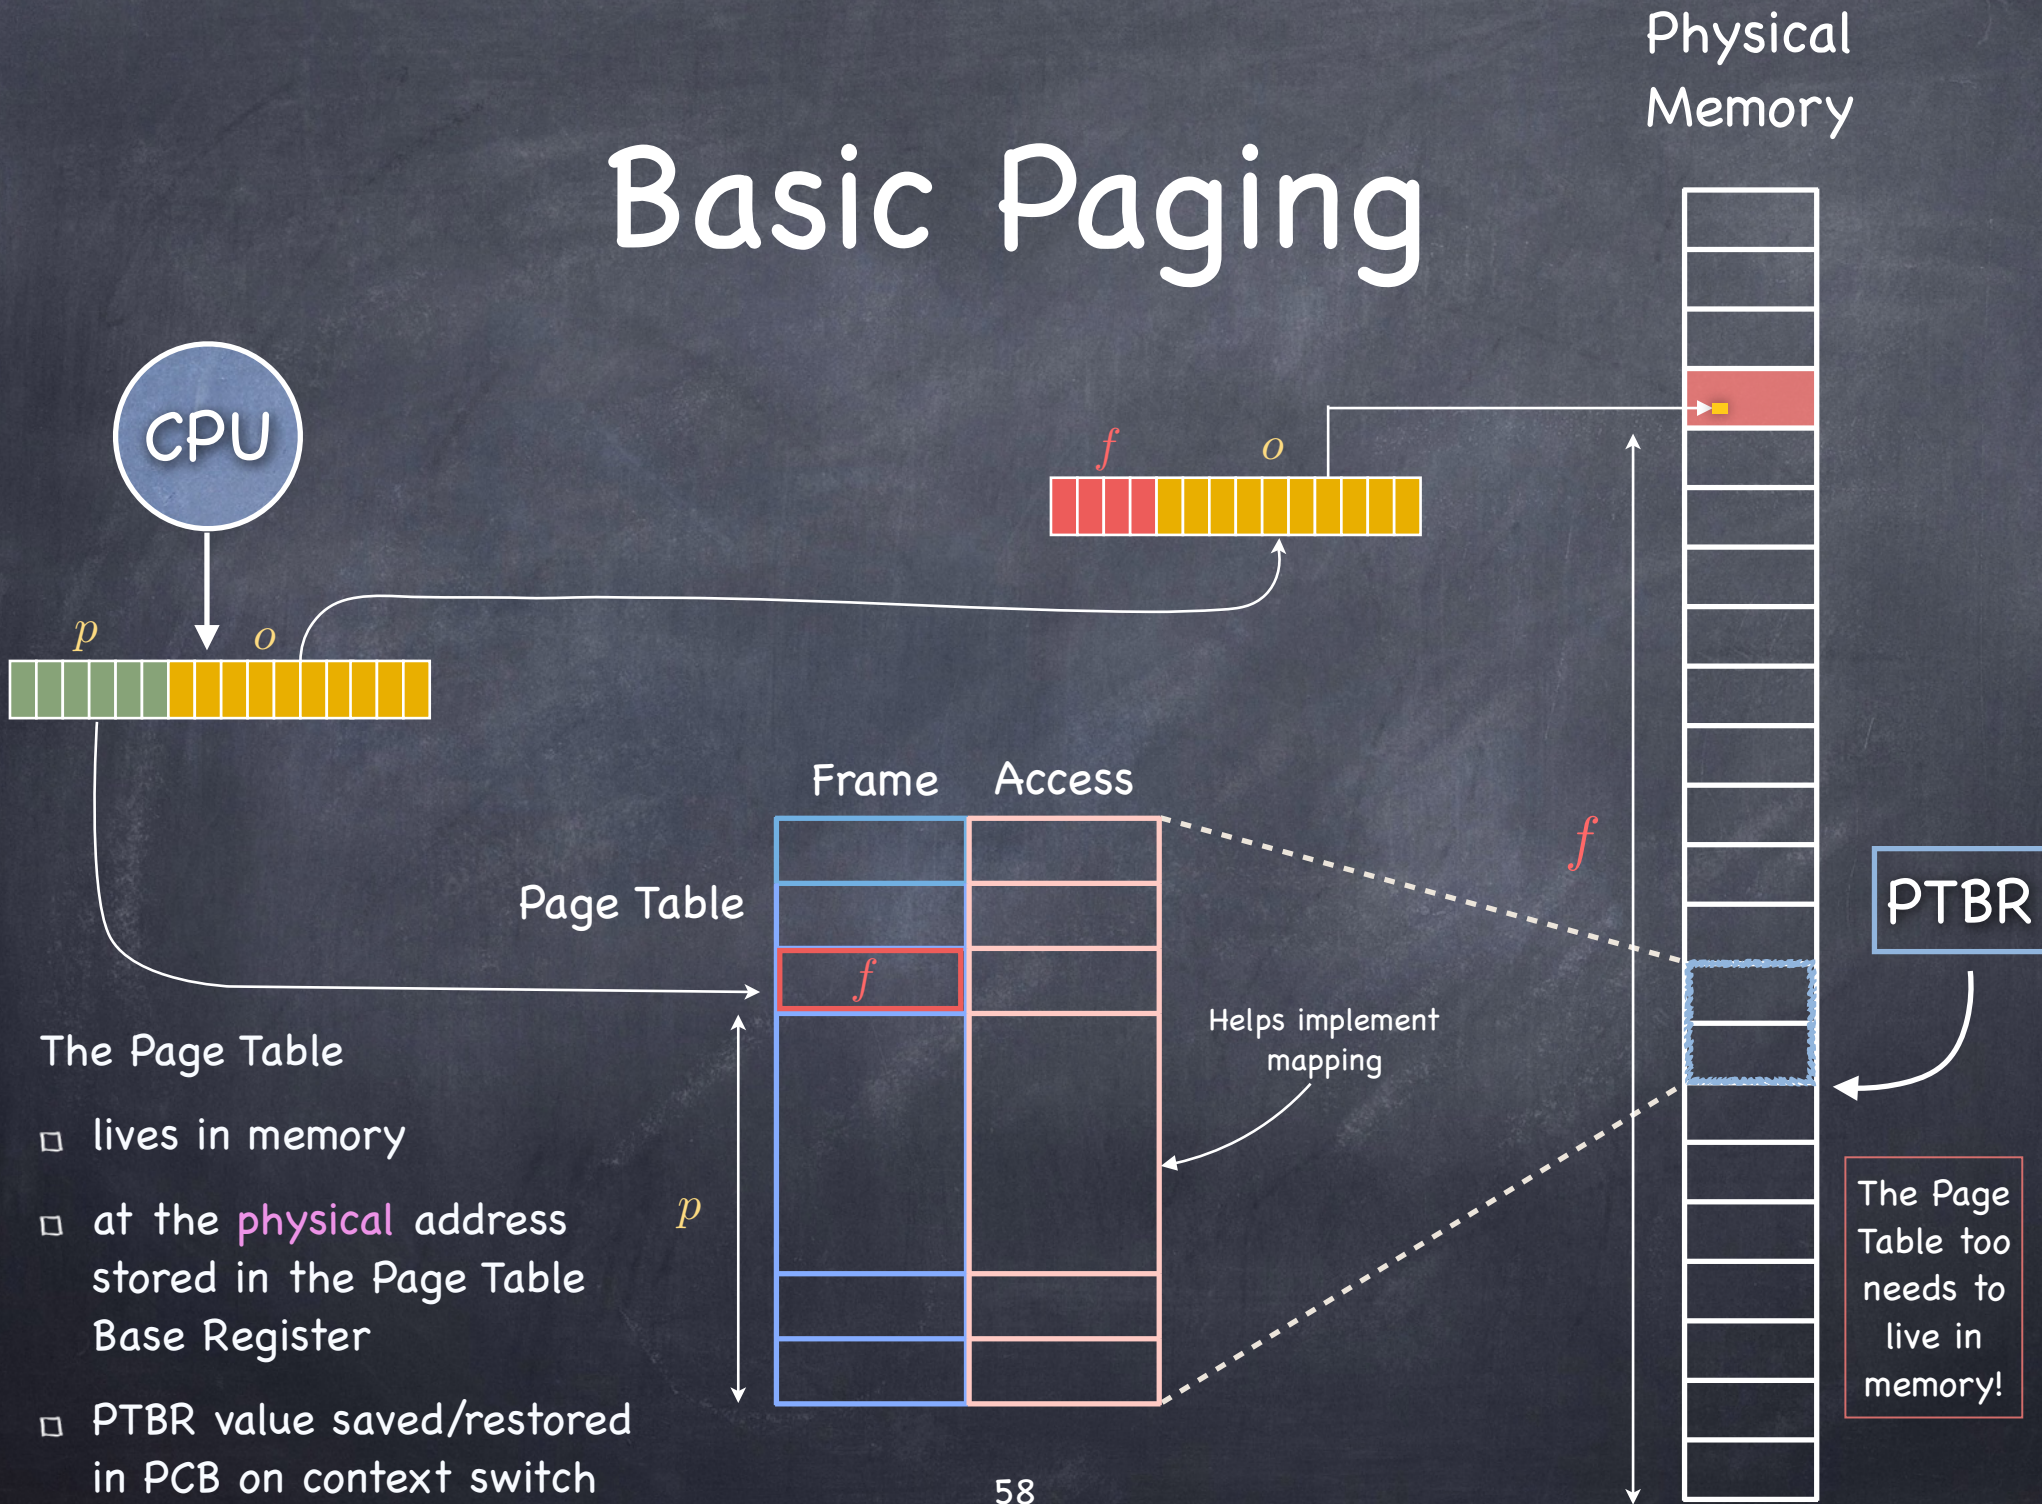
Basic Paging



The Page Table

- lives in memory
- at the **physical** address stored in the Page Table Base Register
- PTBR value saved/restored in PCB on context switch

Basic Paging



The Page Table

- lives in memory
- at the **physical** address stored in the Page Table Base Register
- PTBR value saved/restored in PCB on context switch

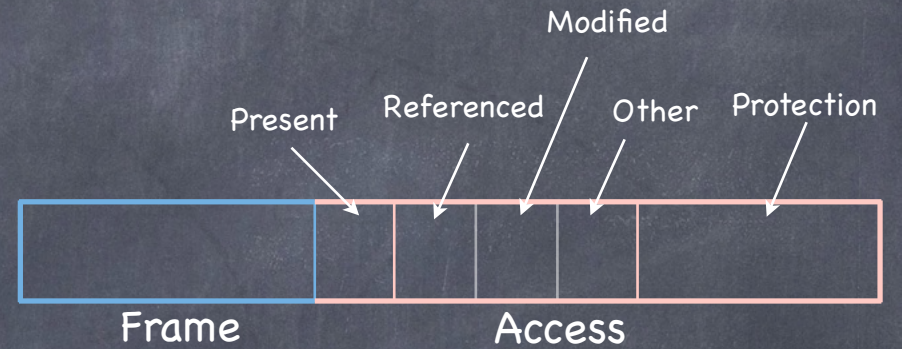
PTBR

The Page Table too needs to live in memory!

Helps implement mapping

Page Table Entries

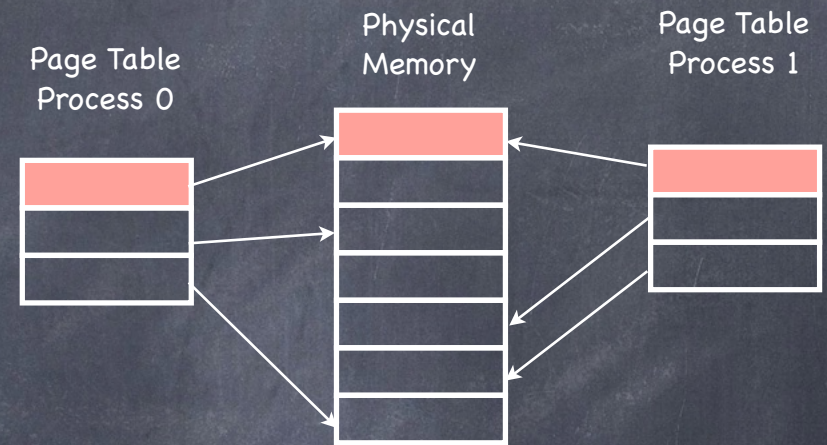
- **Frame number**
- **Present (Valid/Invalid) bit**
 - Set if entry stores a valid mapping. If not, and accessed, page fault
- **Referenced bit**
 - Set if page has been referenced
- **Modified (dirty) bit**
 - Set if page has been modified
- **Protection bits (R/W/X)**



	Page table		Protection bits (R/W/X)	Physical memory
15	4	0		11
14	7	0		2
13	2	0		9
12	0	0		4
11	7	1		5
10	6	0		4
9	5	1		5
8	4	0		0
7	2	0		1
6	0	0		3
5	3	1		0
4	4	1		1
3	0	1		3
2	6	1		
1	1	1		
0	2	1		

Sharing

- Processes share a page by each mapping a page of their own virtual address space to the same frame
 - Fine tuning using protection bits (RWX)
- We can refine COW to operate at the granularity of pages
 - on fork(), mark all pages in page table Read only



- on write:
 - page fault
 - allocate new frame
 - copy page
 - mark both pages R/W

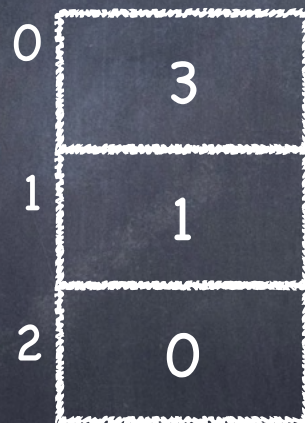
Example

Page size: 4 bytes

VA
Space



Page
Table



4
PA
Space



Space overhead

- Two sources, in tension:
 - data structure overhead (the Page Table itself)
 - fragmentation
 - How large should a page be?

Overhead for paging:

$$\begin{aligned} & (\#entries \times sizeofEntry) + (\# \overset{\text{sets of contiguous pages}}{\text{"segments"}} \times pageSize/2) = \\ = & ((VA_Size/pageSize) \times sizeofEntry) + (\# \text{"segments"} \times pageSize/2) \end{aligned}$$

- What determines sizeofEntry?
 - enough bits to identify physical page ($\log_2 (PA_Size / \text{page size})$)
 - should include control bits (present, dirty, referenced, etc)
 - usually word or byte aligned

Computing paging overhead

- ⑥ 1 MB maximum VA, 1 KB page, 3 segments (program, stack, heap)
 - $((2^{20} / 2^{10}) \times \text{sizeofEntry}) + (3 \times 2^9)$
 - If I know PA is 64 KB then $\text{sizeofEntry} = \text{sizeofFrameNo} + \text{\#ofAccessBits} = 6$ (since we have 2^6 frames) + \#ofAccessBits
 - ▶ if 7 access bits, byte aligned size of entry: 16 bits

What's not to love?

• Space overhead

- ▶ With a 64-bit address space, size of page table can be huge!

• Time overhead

- Accessing data now requires two memory accesses
 - ▶ must also access page table, to find mapped frame

...and, like most times, space and time are in tension...

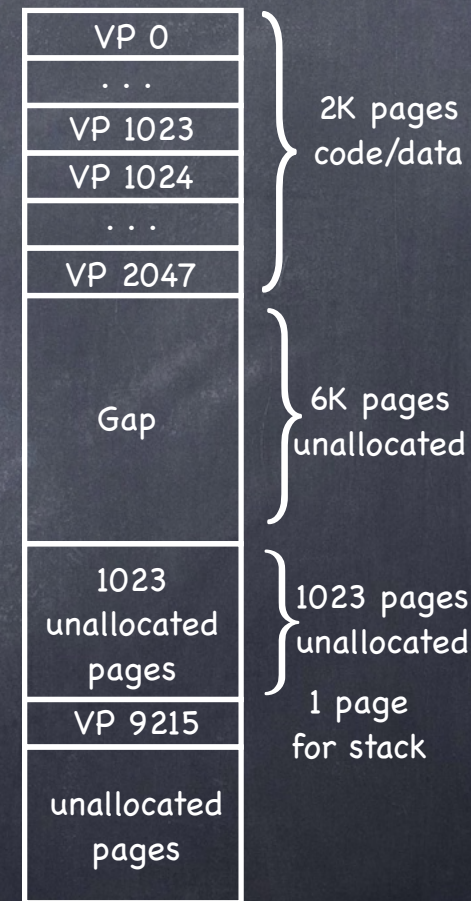
Reducing the Storage Overhead of Page Tables

- Size of the page table for a machine with 64-bit addresses and a page size of 4KB?
 - an array of 2^{52} entries!
- Good news
 - most space is unused
- Use a better data structure to express the Page Table
 - a tree!

Example

- 32 bit address space
- 4Kb pages
- 4 bytes PTE

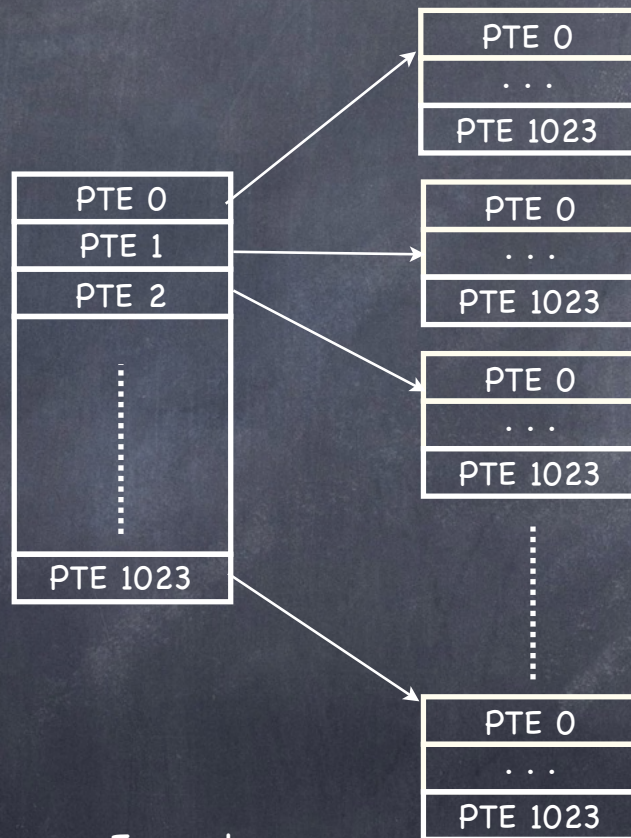
65



Page Table

Reducing the Storage Overhead of Page Tables

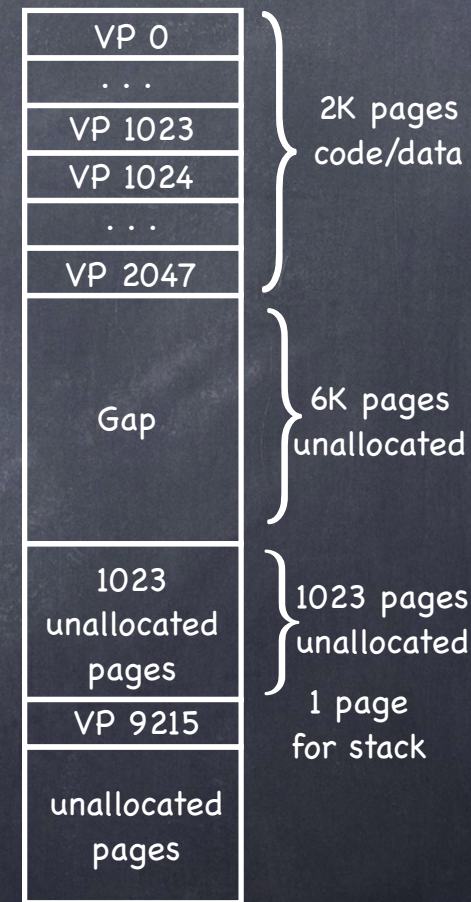
- Size of the page table for a machine with 64-bit addresses and a page size of 4KB?
 - an array of 2^{52} entries!
- Good news
 - most space is unused
- Use a better data structure to express the Page Table
 - a tree!



Example

- 32 bit address space
- 4Kb pages
- 4 bytes PTE

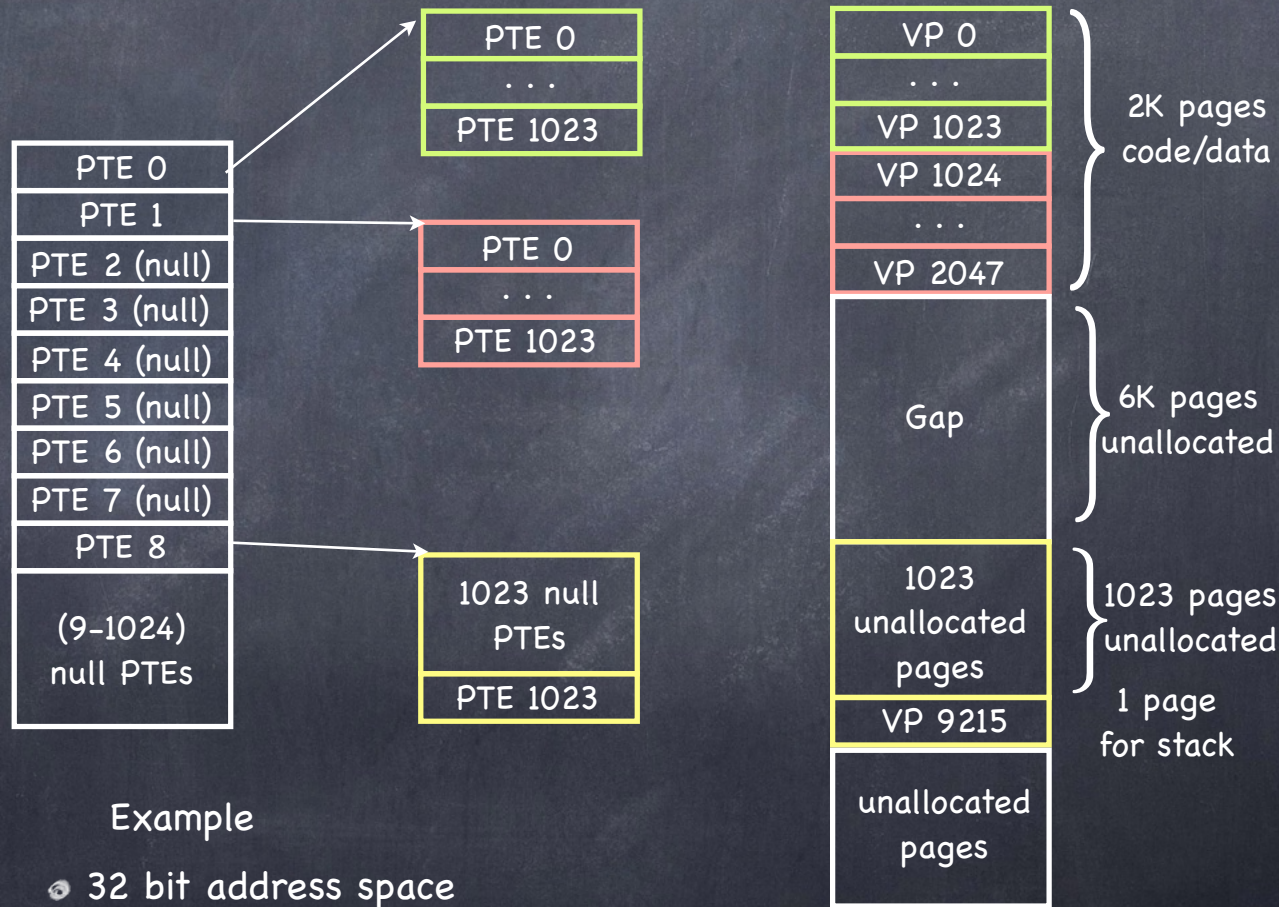
66



Page Table

Reducing the Storage Overhead of Page Tables

- Size of the page table for a machine with 64-bit addresses and a page size of 4KB?
 - an array of 2^{52} entries!
- Good news
 - most space is unused
- Use a better data structure to express the Page Table
 - a tree!



Example

- 32 bit address space
- 4Kb pages
- 4 bytes PTE