

11/31/23

# Cost of Address Translation (Paging)

→ single array = page #  $\Rightarrow$  frame #



→ large page = internal fragmentation

→ inverted page table = frame #  $\Rightarrow$  page #

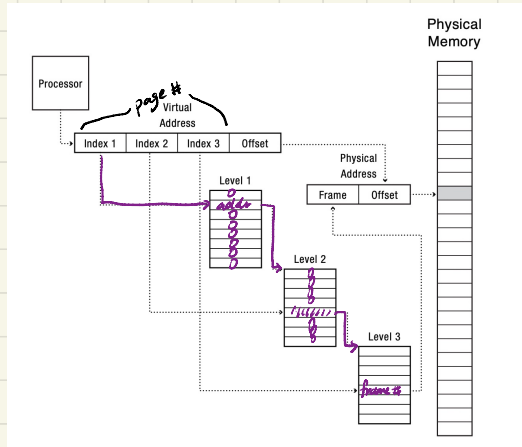
→ multilevel page table

## Page Table

needs to allocate entry for each page



hash collision  
page sharing the same frame



21 bit virtual address

[0x01000 - 0x1fff] code

[0x8f000 - 0x8ffff] stack

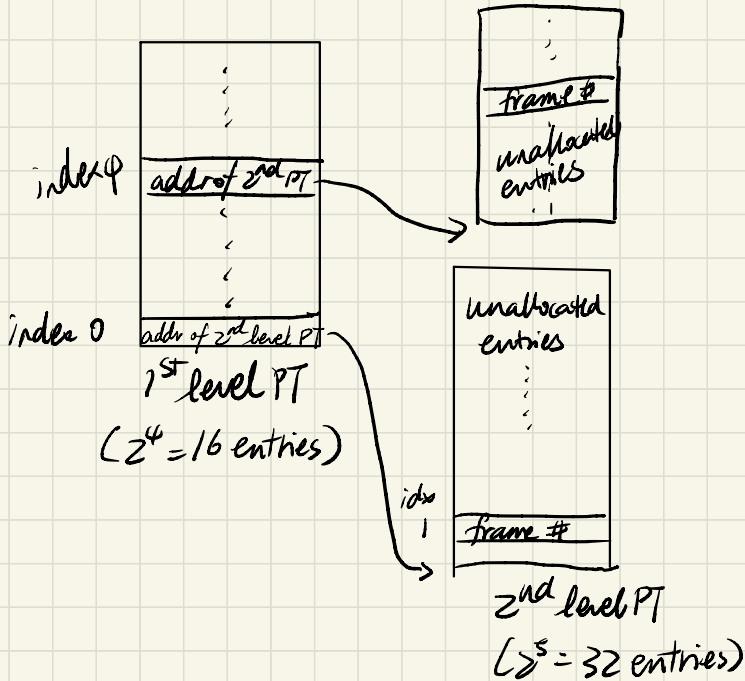
4 bits	5 bits	12 bits
1 <sup>st</sup> level idx	2 <sup>nd</sup> level idx	offset
0000	00001	-----
0100	01111	-----

4KB page size

9 bits for page #

$2^9 = 512$  pages max

single level PT always allocate 512 entries



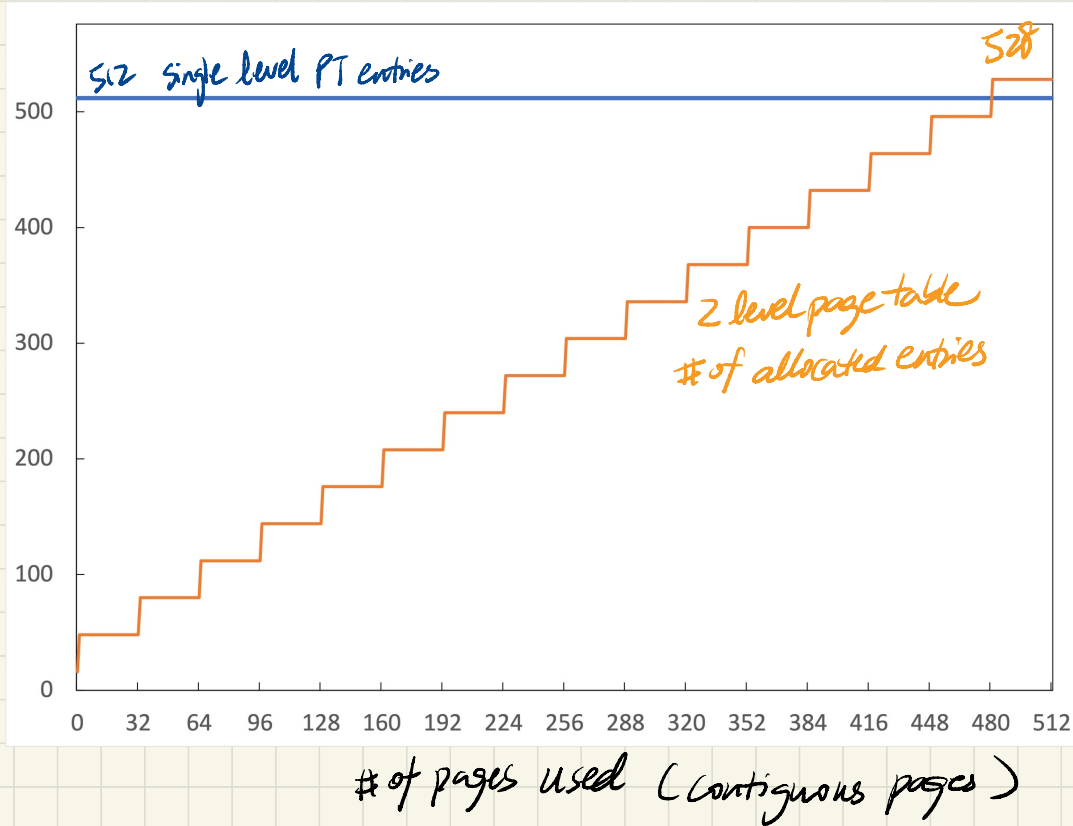
2 page tables needed to map just the code page (one 1<sup>st</sup> level + one 2<sup>nd</sup> level)  
 $16 + 32 = 48$  entries

3 page tables total when we map the stack page (one 1<sup>st</sup> level + two 2<sup>nd</sup> level)  
 $16 + 32 \times 2 = 80$  entries

21 bit virtual address

4 bits	5 bits	12 bits
1 <sup>st</sup> level idx	2 <sup>nd</sup> level idx	offset

1<sup>st</sup> level PT =  $2^4 = 16$  entries  
2<sup>nd</sup> level PT =  $2^5 = 32$  entries



★ Not always  
smaller than  
single level

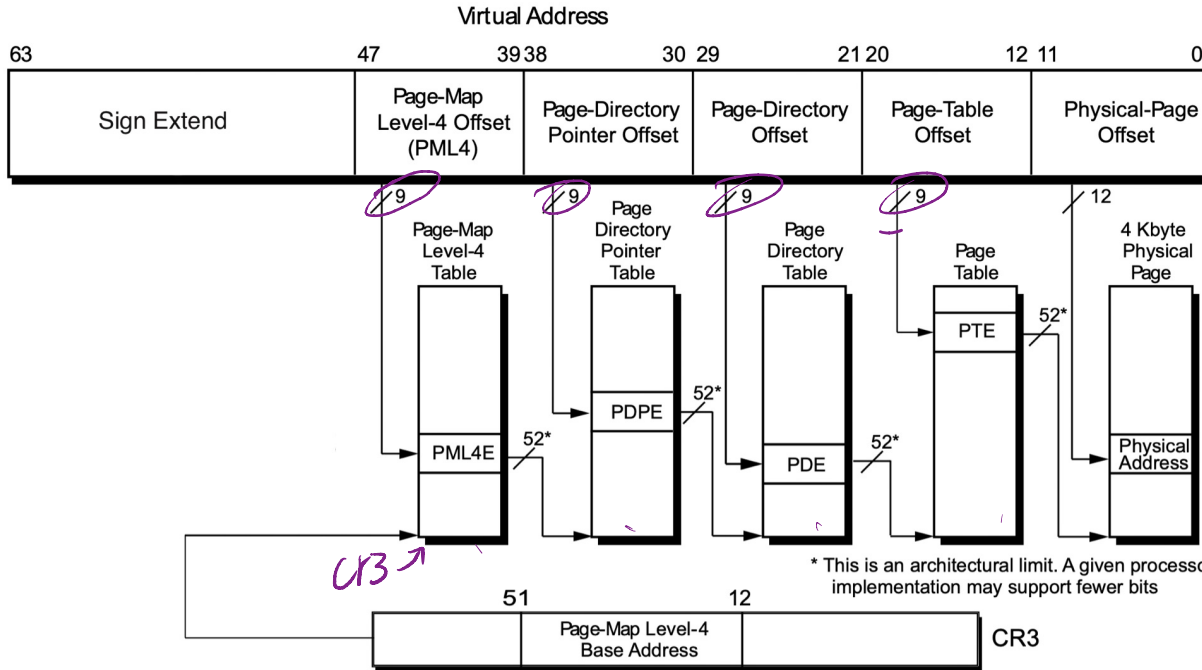
# X86-64 Address Translation

- 4 level page table, each page table takes up a full page in memory

2<sup>9</sup> = 512 entries  
8 byte per entry

$$512 \times 8 = 4096$$


x86-64 VM.C  
(x86)





# Page Fault

- triggered on memory translation error: exception 14
  - ↳ page table walk encounters unallocated page table
  - ↳ page not mapped (page table entry's present bit not set)
  - ↳ mismatch permission (write to readonly page, kernel page accessed in user mode)
- need to identify valid vs. invalid page faults



How does the kernel determine what kind of fault it is?

→ kernel has bookkeeping structures (machine independent) that track information about each page.

\* Machine independent vs. machine dependent page table  
(vspace, vregion, vpage-info) (x86-64 pt)

info for address translation

info for page fault

```
42 struct vregion {
43     enum vr_direction dir; // direction of growth
44     uint64_t va_base; // base of the region
45     uint64_t size; // size of region in bytes
46     struct vpi_page *pages; // pointer to array of page_infos
47 };
48
49 struct vspace {
50     struct vregion regions[NREGIONS]; // the regions for a process' virtual space
51     pml4e_t *pgtbl; // process' page table
52 };
```

first level page table

info about each page

```
17 struct vpage_info {
18     short used; // whether the page is in use
19     uint64_t ppn; // physical page number
20     short present; // whether the page is in physical memory
21     short writable; // does the page have write permissions
22     // user defined fields
23
24 };
```