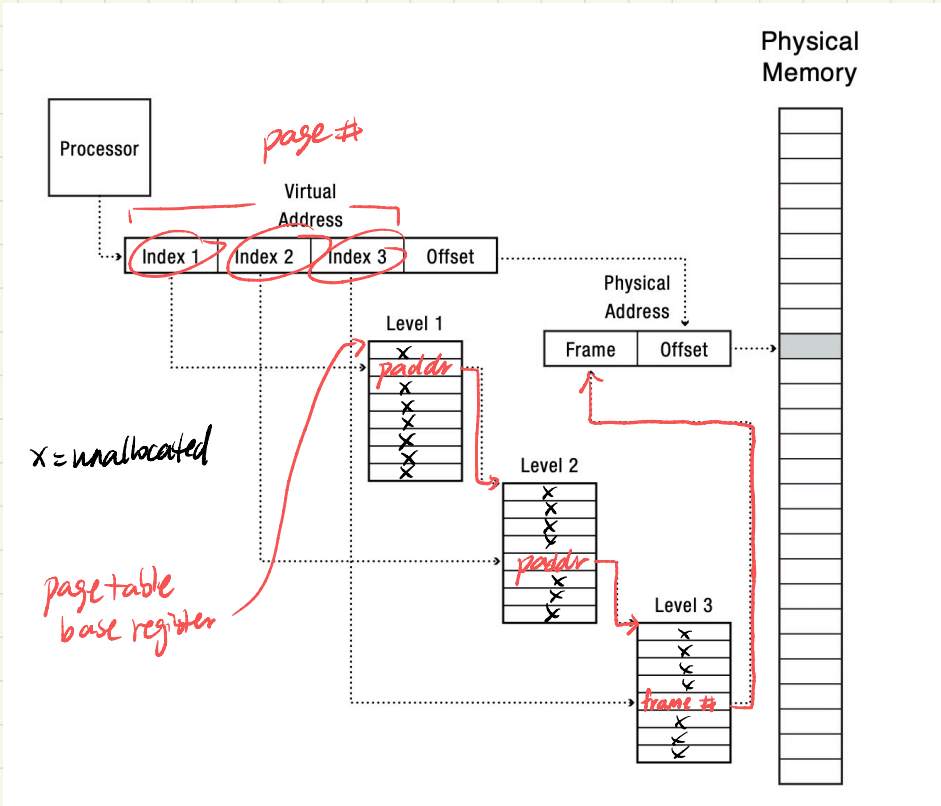


4/26/24

Multi-level Paging



→ each last level page table is a shard of the single array page table

→ indirection takes up more space if most pages are used.

⊗ TLB even more important!

21 bit virtual address

[0x1000 - 0x1fff] Code

[0xfe00 - 0xffff] Stack

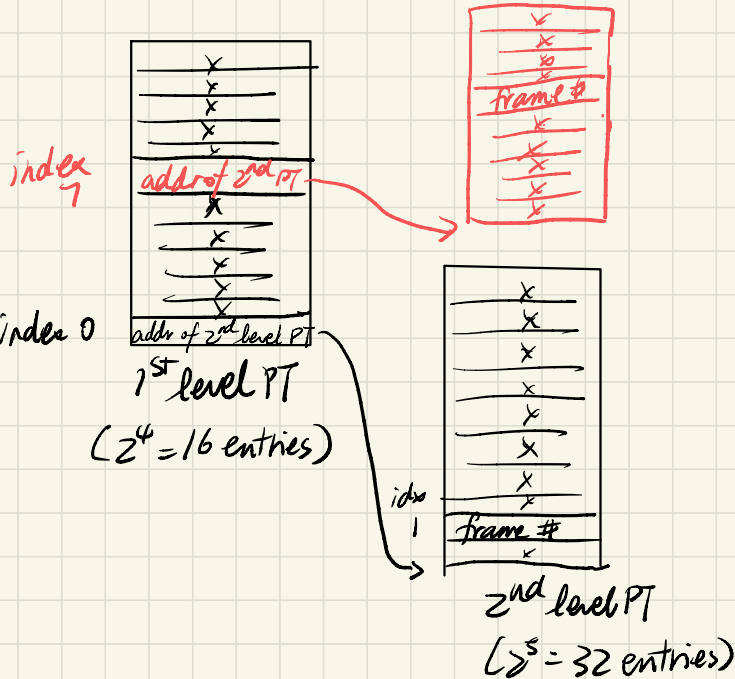
4 bits 1 st level idx	5 bits 2 nd level idx	12 bits offset
0000	00001	-----
0111	11110	-----

4KB
page
size

9 bits for page #

$2^9 = 512$ pages max

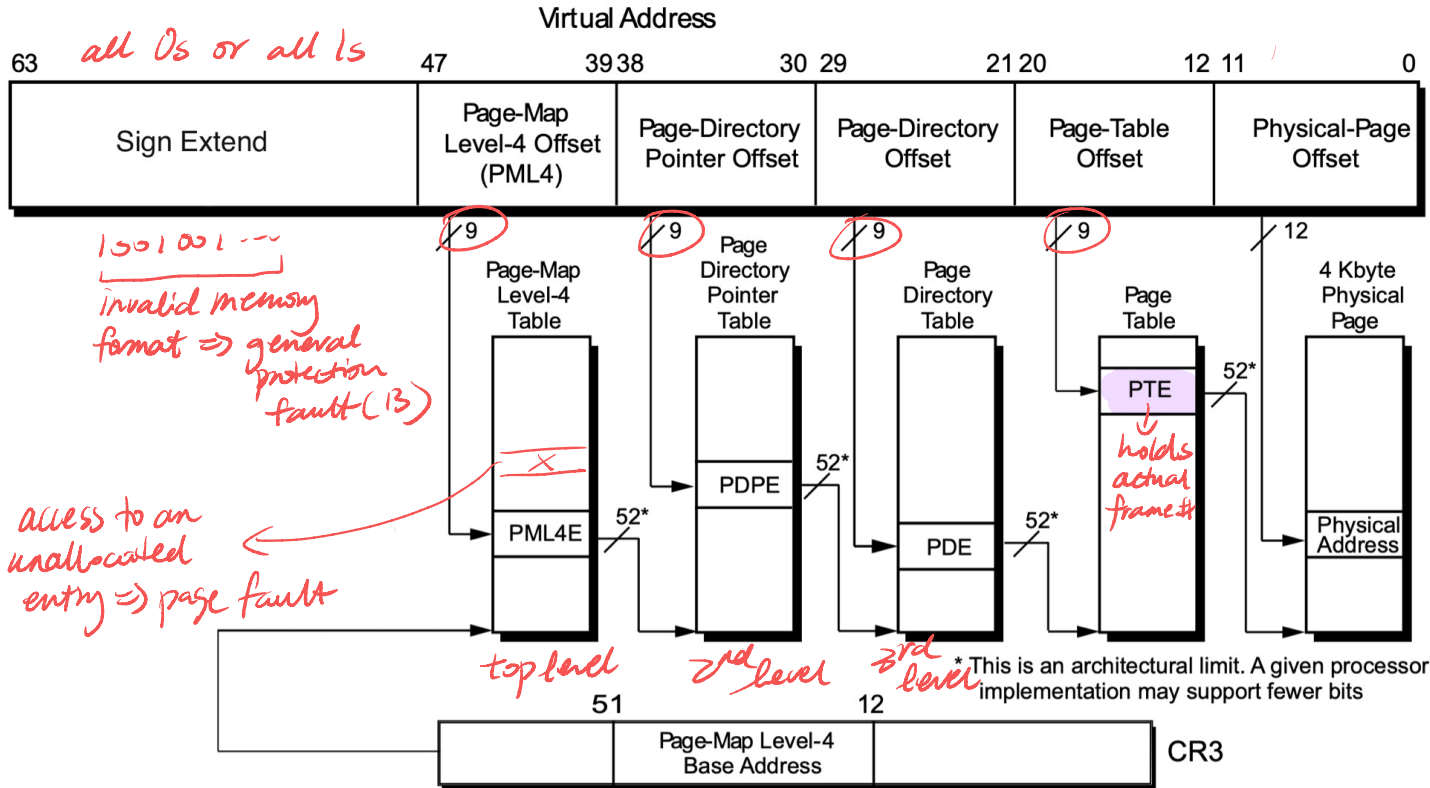
Single level PT always allocate
512 entries



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

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

786-64 Page Table : each level of page table has 2^9 entries & lives in a pages & byte entry



Page Fault (page table walk or TLB)

→ 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 read-only page, kernel page accessed in user mode)

→ types of page fault

→ demand paging (valid)

→ stack growth

→ heap growth

→ memory mapped files

→ Permission Mismatch

→ low access

→ actual violation

(user access kernel addr)
write to read only page.

allocate
new frame,
update PT
& flush TLB

allocate
a frame,
clear out
the contents
before mapping
it.

load
from disk

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

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


xk VM structures

machine dependent

```
struct vspace {  
    struct vregion regions[NREGIONS]; // the regions for a process' virtual space  
    pml4e_t* pgtbl; // process' page table  
};
```

machine independent

vspace_update() updates the machine dependent PT based on machine independent vregion contents




process VAS

```
struct vregion {  
    enum vr_direction dir; // direction of growth  
    uint64_t va_base; // base of the region  
    uint64_t size; // size of region in bytes  
    struct vpi_page *pages; // pointer to array of page_infos  
};
```

```
struct vpi_page {  
    struct vpage_info infos[VPIPPAGE]; // info struct for the given page  
    struct vpi_page *next; // the next page  
};
```

linked list of vpage info array



xk VM structures

```
struct vpage_info {  
    short used;      // whether the page is in use  
    uint64_t ppn;   // physical page number  
    short present;  // whether the page is in physical memory  
    short writable; // does the page have write permissions  
    // user defined fields  
  
};
```

per page metadata