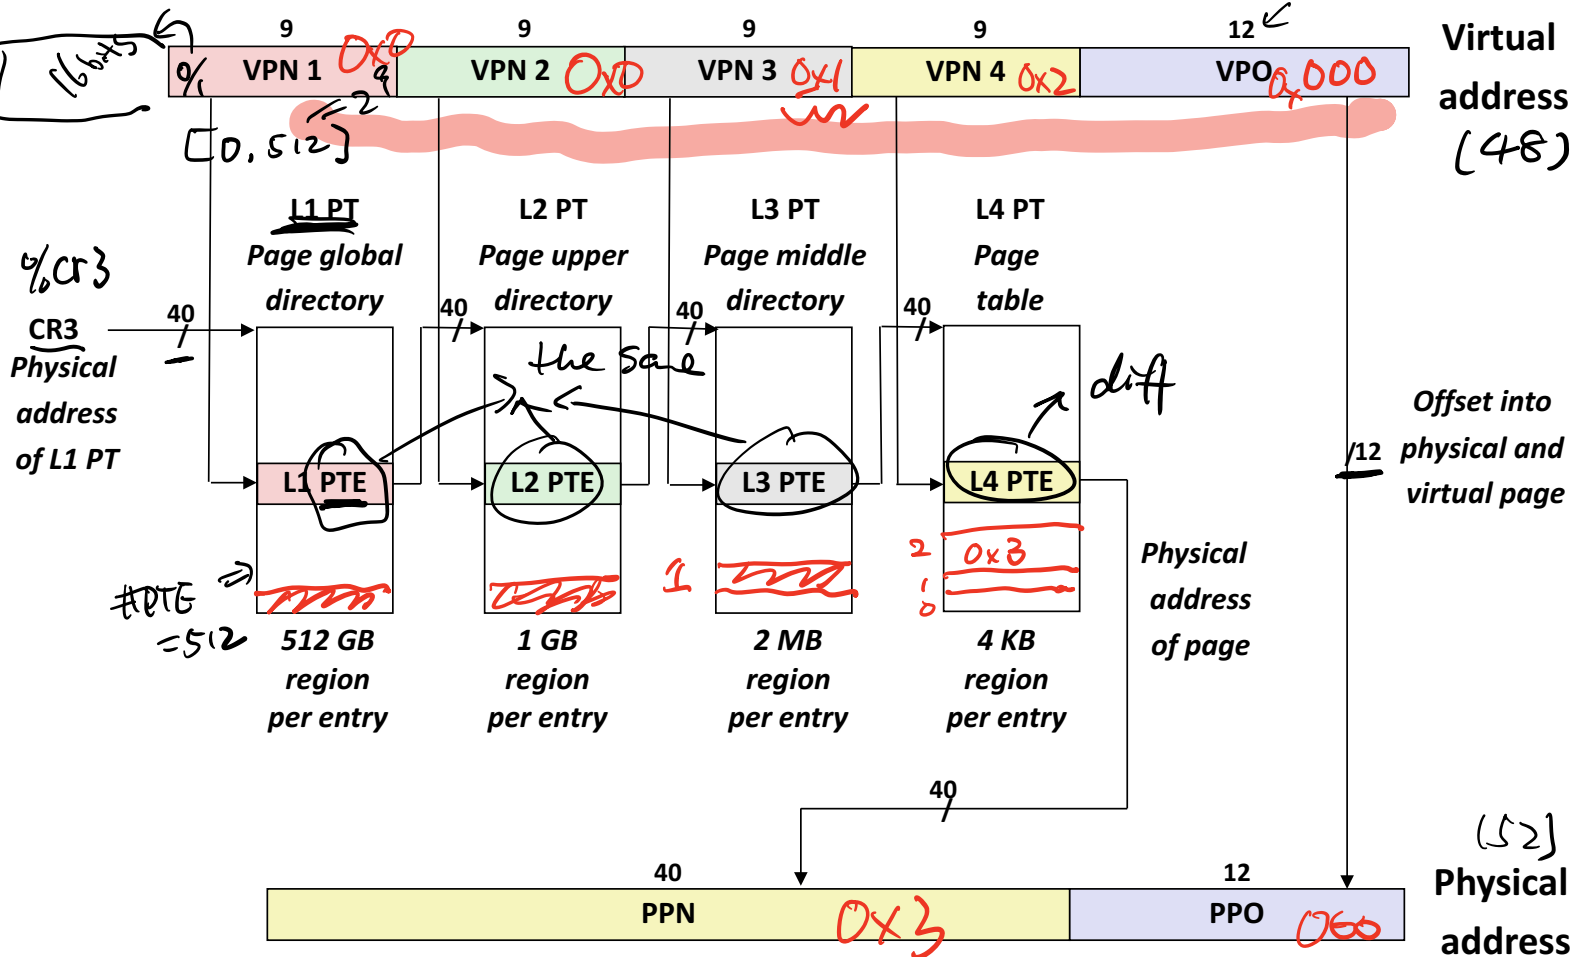


# Core i7 Page Table Translation

0x202 → binary  
 0010 0000 0010

VA = 0x202000<sup>12</sup> 1261  
 PA: 0x3000<sup>12</sup>



# Review of 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)


## ■ Components of the virtual address (VA)

- TLBI: TLB index
- TLBT: TLB tag
- VPO: Virtual page offset
- VPN: Virtual page number

## ■ Components of the physical address (PA)

- PPO: Physical page offset (same as VPO)
- PPN: Physical page number
- CO: Byte offset within cache line
- CI: Cache index
- CT: Cache tag

# Core i7 Level 1-3 Page Table Entries



Each entry references a 4K child page table. Significant fields:

**P:** Child page table present in physical memory (1) or not (0). ←

**R/W:** Read-only or read-write access access permission for all reachable pages. ←

**U/S:** user or supervisor (kernel) mode access permission for all reachable pages. ←

**WT:** Write-through or write-back cache policy for the child page table. ←

**A:** Reference bit (set by MMU on reads and writes, cleared by software).

**PS:** Page size: if bit set, we have 2 MB or 1 GB pages (bit can be set in Level 2 and 3 PTEs only).

**Page table physical base address:** 40 most significant bits of physical page table address (forces page tables to be 4KB aligned)

**XD:** Disable or enable instruction fetcher from all pages reachable from this PTE.

Page Fault  
 movq src, dst  
 Red write  
 Ptw=0

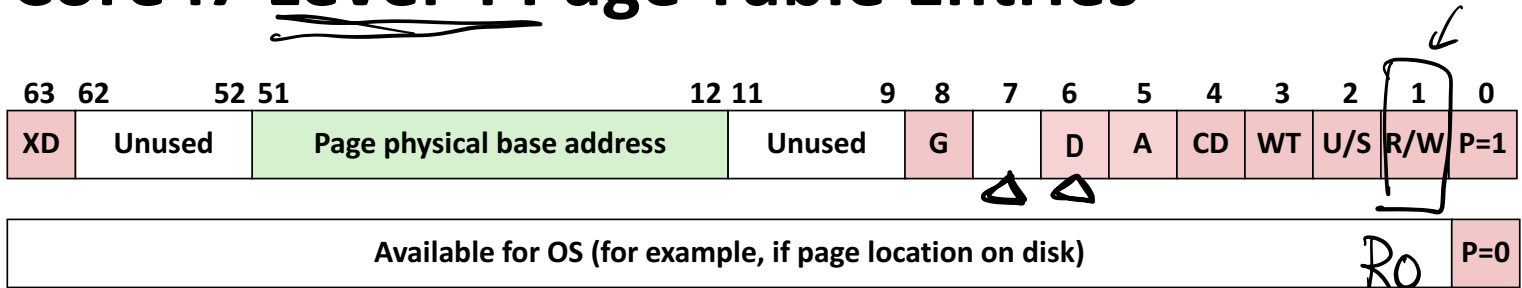
large page

Jump Addr L ⇒ %rip



Code

# Core i7 Level 4 Page Table Entries



Each entry references a 4K child page. Significant fields:

**P:** Child page is present in memory (1) or not (0)

**R/W:** Read-only or read-write access permission for this page

**U/S:** User or supervisor mode access *moving src, dst*

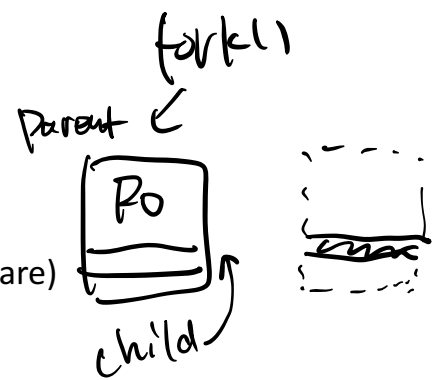
**WT:** Write-through or write-back cache policy for this page

**A:** Reference bit (set by MMU on reads and writes, cleared by software)

**D:** Dirty bit (set by MMU on writes, cleared by software)

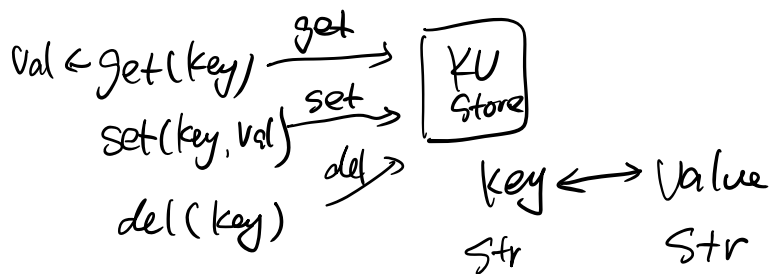
**Page physical base address:** 40 most significant bits of physical page address  
(forces pages to be 4KB aligned)

**XD:** Disable or enable instruction fetches from this page.

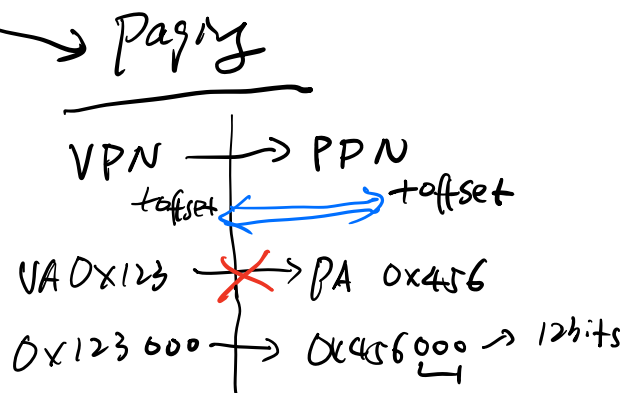



1. Last time
  2. x86-64: addresses
    - virtual
    - physical
  3. x86-64: page table structures
  4. Practice
- 

## Lab 3. Concurrent KUStore




### Virtual memory



Naive proposal

Array





$$\text{Array}[0x123] = 0x456$$

Size:

$$\text{len}(\text{Array}) = \# \text{VPN}$$



$$= 2^{36} \times 8B = 2^1 \cdot 2^3 \cdot 2^{30} B$$

$$48 \text{ bit} \approx 12 \text{ bit}$$

$$= 36 \text{ bit}$$

$$= 512 \text{ GB}$$

### Multi-level PT





Example: x86-64 VA. 48 bit

OS maps  $[0, 2\text{MB}-1] \Rightarrow \text{PAE}$

Q: how many (pages) in PT?


$$[0x0, \text{redacted} \cdot 2^{21}-1]$$



$$1 + 1 + 1 + 1 = 4$$

n

X: 2GB, how many pages in P1?



512 WHY?  $\Rightarrow$  HW  $\frac{4KB}{8B} = 512$

$\frac{144}{516}$



①

[0, 2GB]

$2^{31} - 1$  48bit

$\rightarrow$  38

0060111111



$1 + 1 + 2 + \frac{2 \times 512}{1024} = 1028 \Leftrightarrow$

②



$\frac{2GB}{4KB} = \frac{1}{2} \cdot 2^{20}$   
 $= 2^{19}$  pages

#L4 PT =  $\frac{2^{19}}{2^9}$

$= 2^{10}$   
 $= 1024$  L4

#L3 PT =  $\frac{1024}{512}$   
 $= 2$

X86-64.




$$2^{48} = 2^8 \text{ TB} = 256 \text{ TB}$$

[ 5-level paging

• PA. 52bit

4-level  $\Rightarrow$  5-level

WHY  $\Rightarrow$  HW




$$2^{52} = 4 \text{ PB}$$

48 bits  $\Rightarrow$  57 bits

$$2^{57} = 128 \text{ PB}$$

(16GB)

• VA (48)  $\rightarrow$  PA (52)





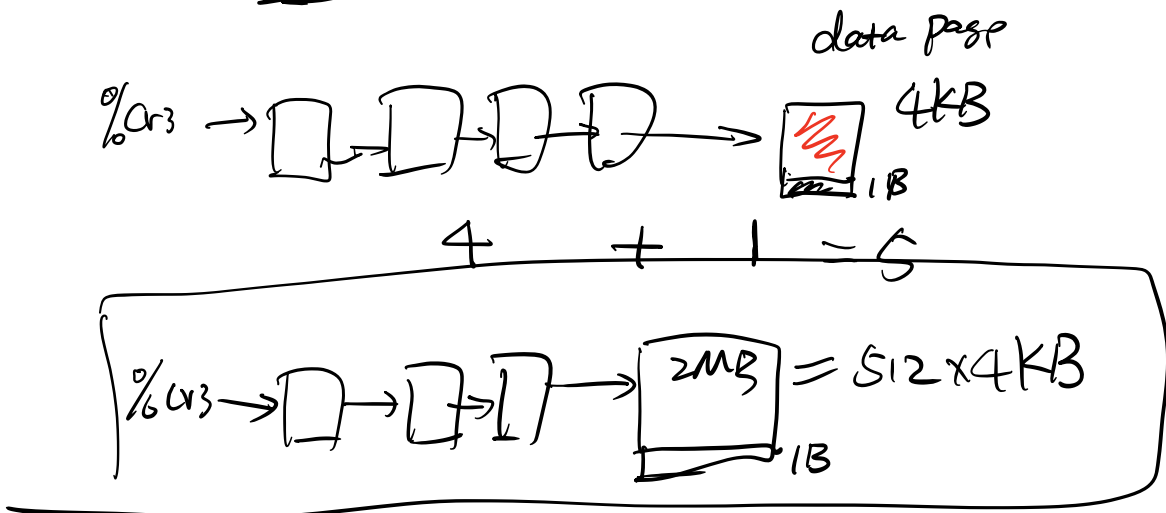
Q: x86-64

OS.

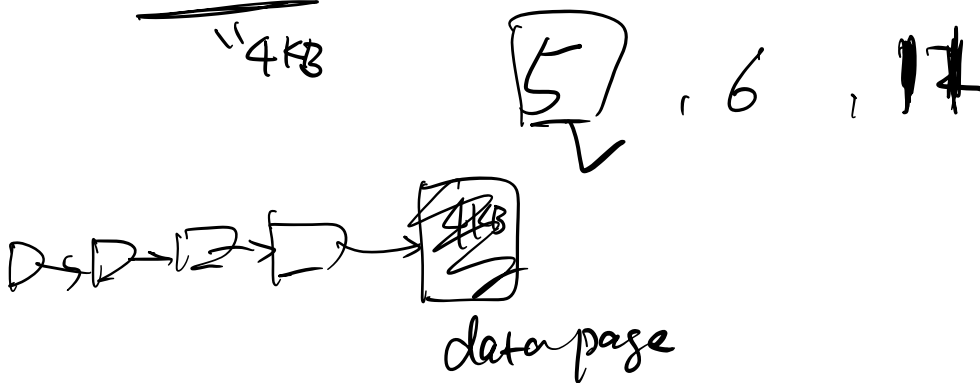
Program allocate Memory min

how many pages in memory?


• allocate 1B :




• alloc  $\frac{2^{12} B}{4KB}$  ?



• alloc  $\frac{2^9 \times 2^{12} \text{ Bytes}}{4KB}$




512 x 4FB



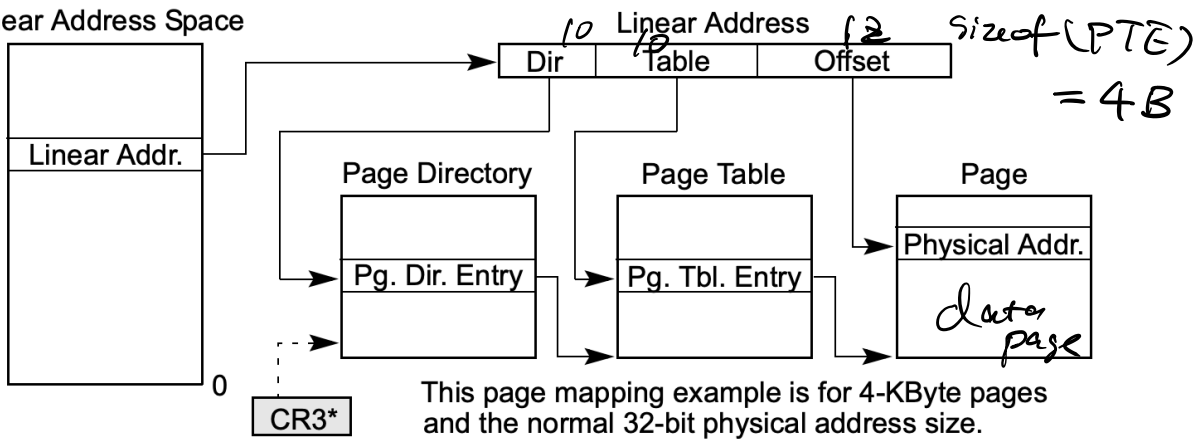
$$1 + 1 + 1 + 1 + 512 = 516$$

x86-32



↳ 1024 PTE

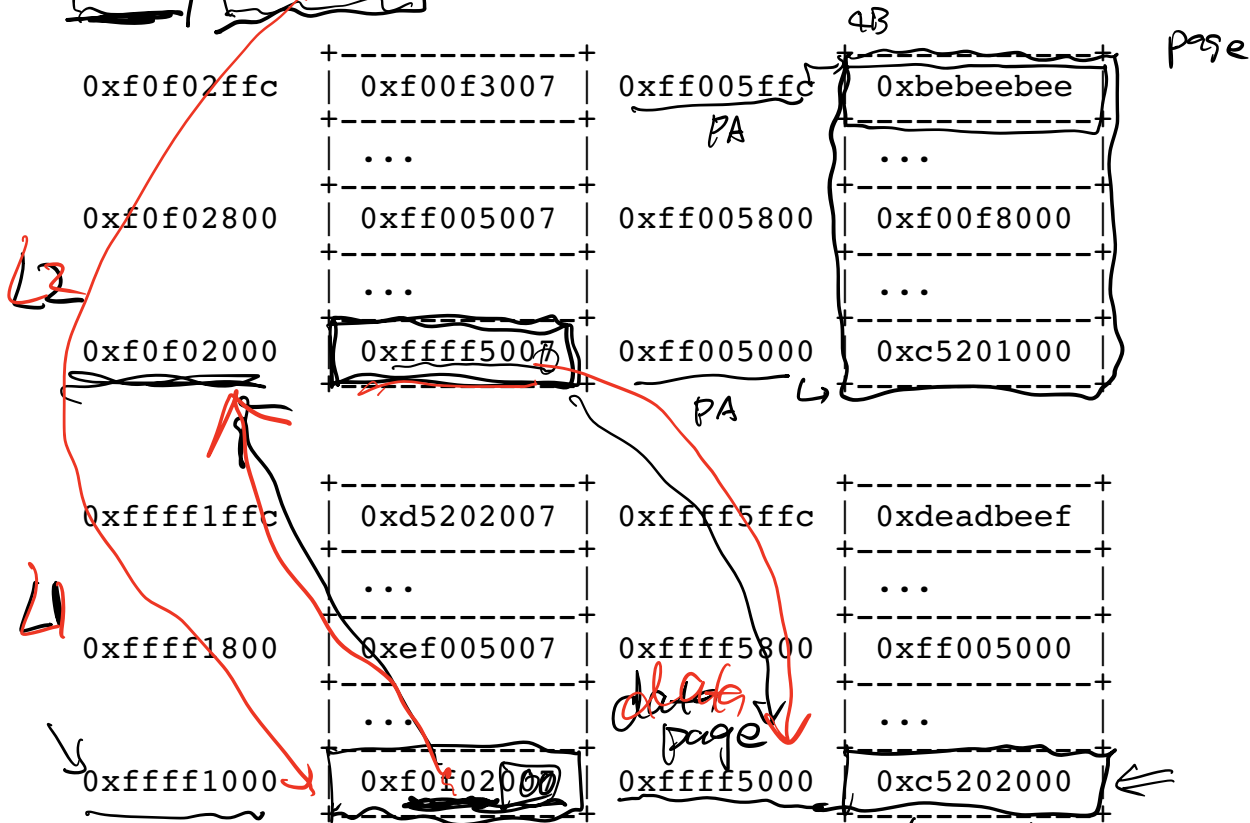
Linear Address Space




This page mapping example is for 4-KByte pages and the normal 32-bit physical address size.

\*Physical Address

`%cr3: 0xffff1000`



```
int *ptr1 = (int *) 0x0;
printf("%x\n", *ptr1);
```



`*ptr1 = 0x0`

-  $\rightarrow$  Удтлтт 000-  
Print (\*ptr1) 0xc5202000  $\leftarrow$