

# Cache Read

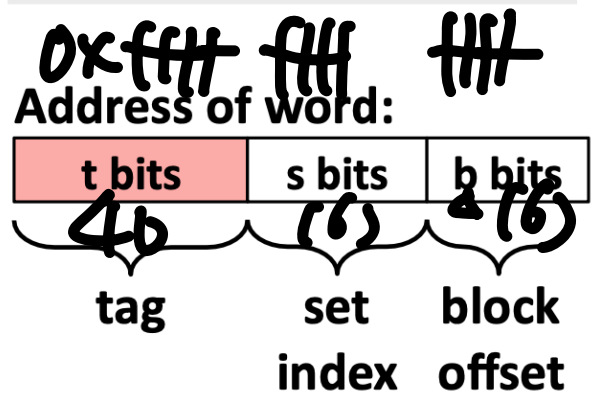
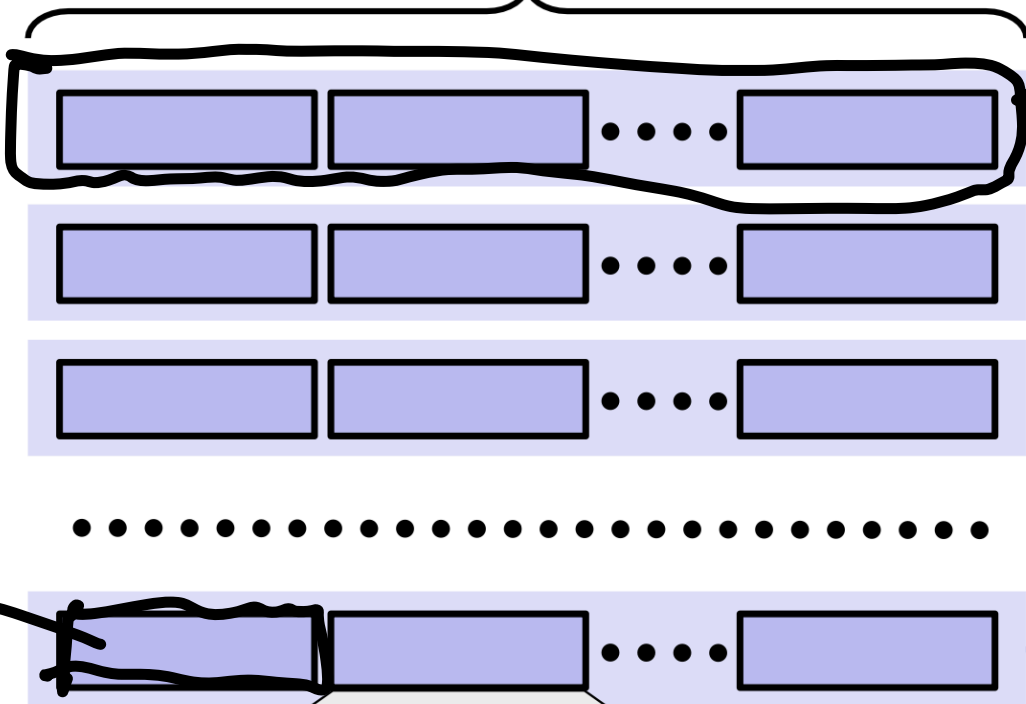
$(E, S, B)$  <sup>fixed</sup> (4B)

$E = 8$ -way  
 $E = 2^e$  lines per set

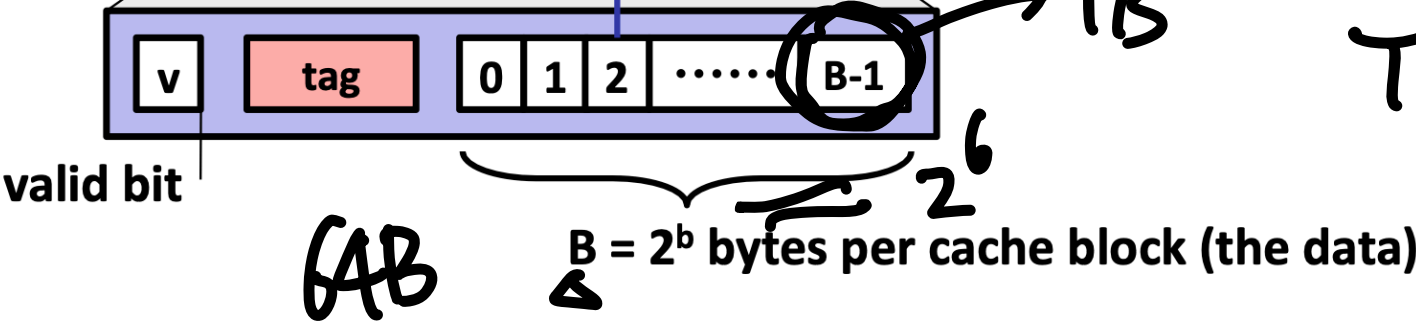
- Locate set
- Check if any line in set has matching tag
- Yes + line valid: hit
- Locate data starting at offset

$64 = 2^6$   
 $S = 2^s$  sets

Cache line

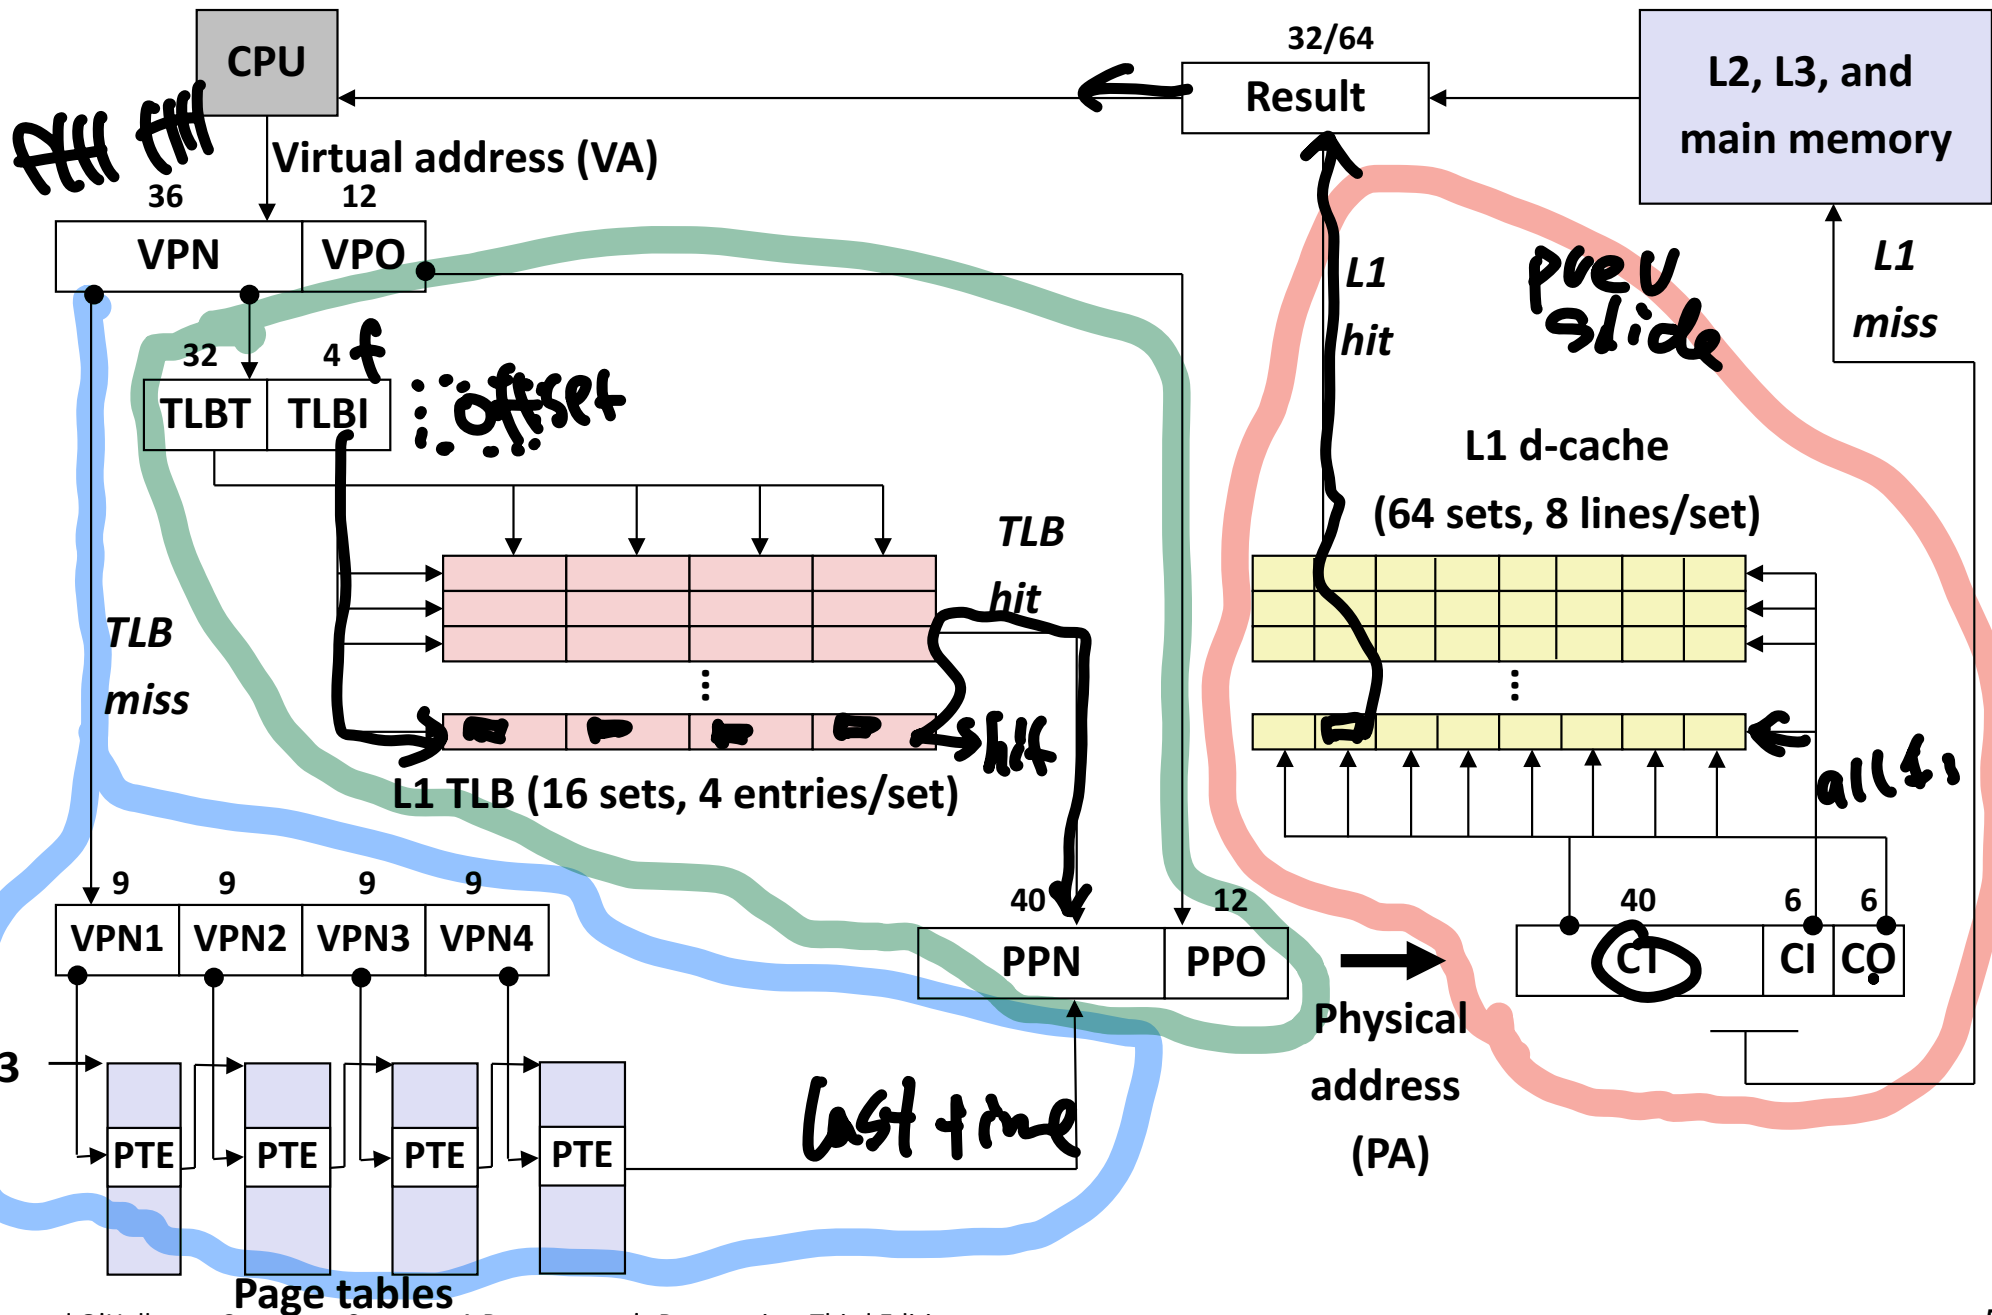


all 1s.  
 all 1s.  
 data begins at this offset

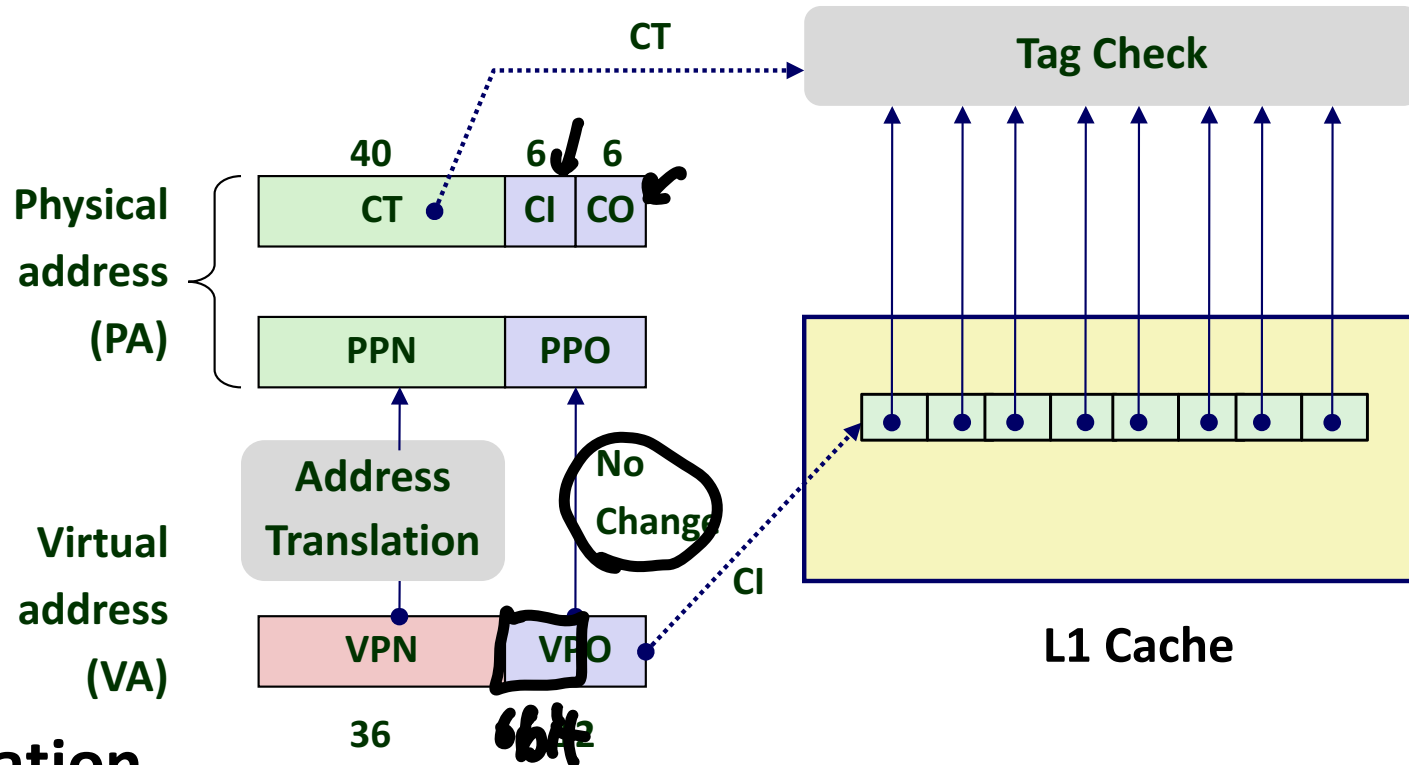


TUB  
 SZMZAR

# End-to-end Core i7 Address Translation



# Cute Trick for Speeding Up L1 Access

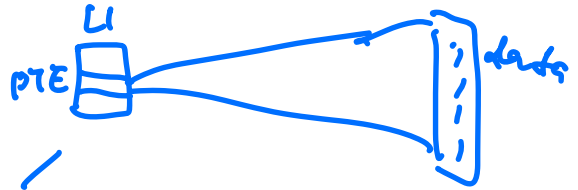


## ■ Observation

- Bits that determine CI identical in virtual and physical address
- Can index into cache while address translation taking place
- Cache carefully sized to make this possible: 64 sets, 64-byte cache blocks
- Means 6 bits for cache index, 6 for *cache* offset
- That's 12 bits; matches *VPO*, *PPO* → One reason pages are  $2^{12}$  bytes = 4 KB

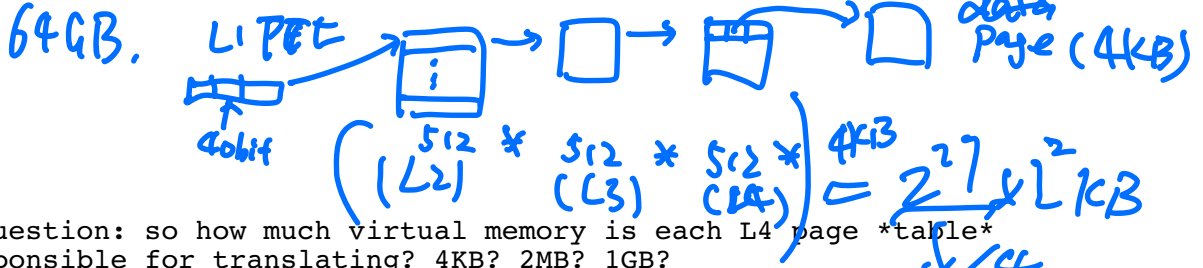


Practice

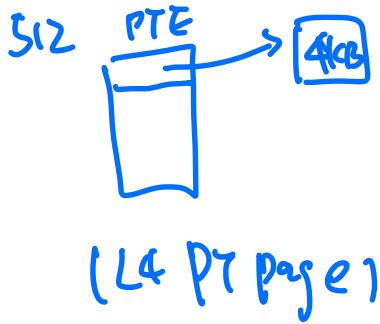


A: memory that different PTEs can address

--Question: how much memory can one L1 page entry address?



--Question: so how much virtual memory is each L4 page \*table\* responsible for translating? 4KB? 2MB? 1GB?



$$512 \times 4KB = 2MB$$

$$\begin{aligned} \# \text{ PTE} &= 2^{29} \text{ KB} \\ &= 2^{19} \text{ MB} \\ &= 2^9 \text{ GB} \\ &= 512 \text{ GB} \end{aligned}$$

B. Allocating memory

What is the minimum number of physical pages required on x86-64 to allocate the following allocations?



(b)  $2^9$  allocations of size of  $2^{12}$  bytes of memory each ⇒

(c)  $2^9 + 1$  allocations of size of  $2^{12}$  bytes of memory each  
 $2^{18} + 1$  allocations of size  $2^{12}$  bytes of memory each

(b)  $L1: 1 + L2: 1 + L3: 1 + L4: [512 \text{ PTEs}] + 512 = 4 + 512$   
 (1 page) (data)

(c) :

$$1 + 1 + 1 + 2 + 513 (\text{data}) = 5 + 513$$

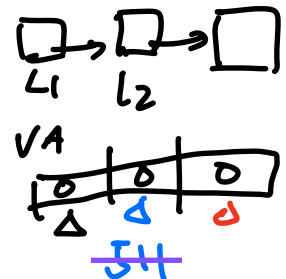
C. page table walk (x86-32) VA 

10	10	12
----	----	----

 (32 bits) data

`%cr3: 0xffff1000 (CPU)`

0xf0f02ffc	0xf00f3007	0xff005ffc	0xbebeeb
	...		...
0xf0f02800	0xff005007	0xff005800	0xf00f8000
	...		...
0xf0f02000 (0)	0xffff5007 (L2)	0xff005000	0xc5201000
	...		...
0xffff1fff (102)	0xd5202007	0xffff5ffc	0xdeadbeef
	...		...
0xffff1800 (5)	0xef005007	0xffff5800	0xff005000
	...		...
0xffff1000 (0)	0xf0f02007 (L1)	0xffff5000 (0)	0xc5202000 (data)



[update 03/21] "511" is a typo; it should be "512" (men)

[update 03/21] "511" is a typo; it should be "512"

- What's the output of the following C excerpt?

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

4B

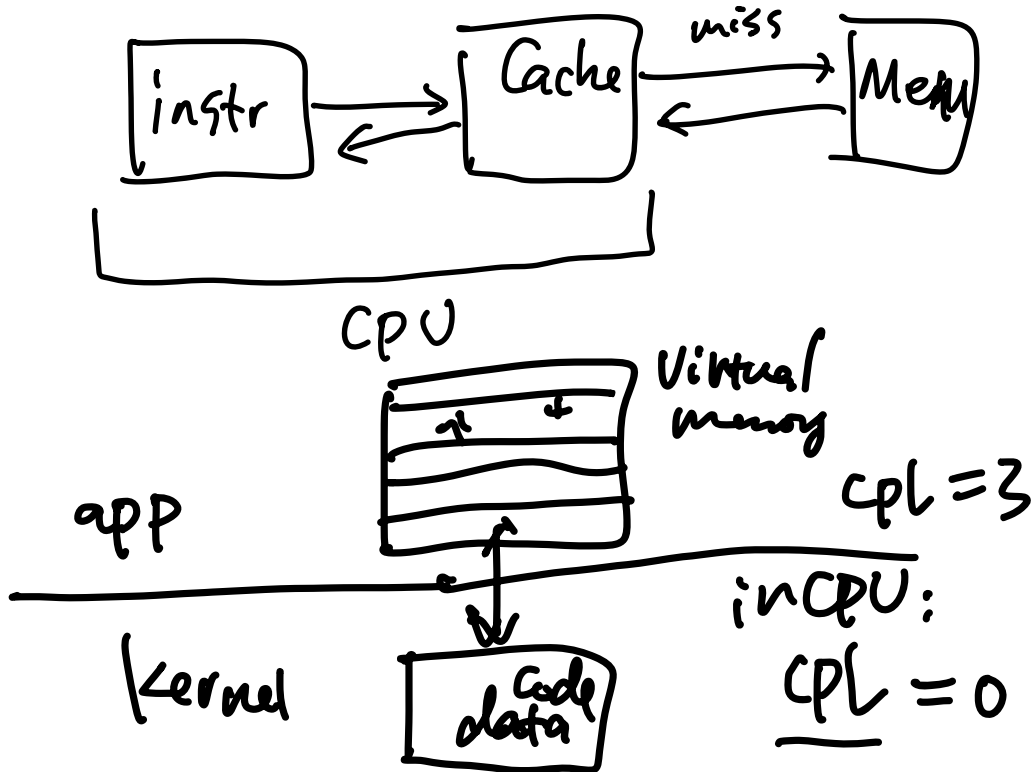
(addr)

→ ~~0x520f000~~  
0xc5202000

TLB: translation lookaside buffer

↳ Cache. mapping: VPN → PPN

CPU Cache



Option 1: kernel own space  
↳ Lab4. ( %cr3 )

⇒ Option 2: kernel app share  
the same space

# Virtual Address Space of a Linux Process

PTE:  
V/S = 0  
⇒

Identical for each process

