

⊛ `mov ebx, (0x abcdef23)`

| in real mode: — physical address  
 | in protected mode — virtual address  
 | w/ paging — mmu <sup>of</sup> on CPU

performs page table

look pgdir ← walk.

— use offset from VA on pgdir to  
check pte.

if pte valid — move to next level  
~~else~~ else page fault.

check pte flags — v, p etc.  
for validity.

& recurse on page table levels.

⊛ on valid / large page fault  
or on enthusiastic mapping of v2p

— find free physical page for mapping.  
 — if no free page, evict + swap!

# OS is overhead

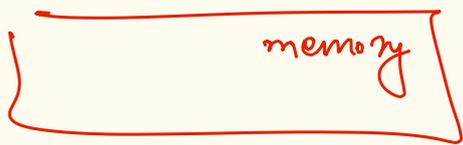
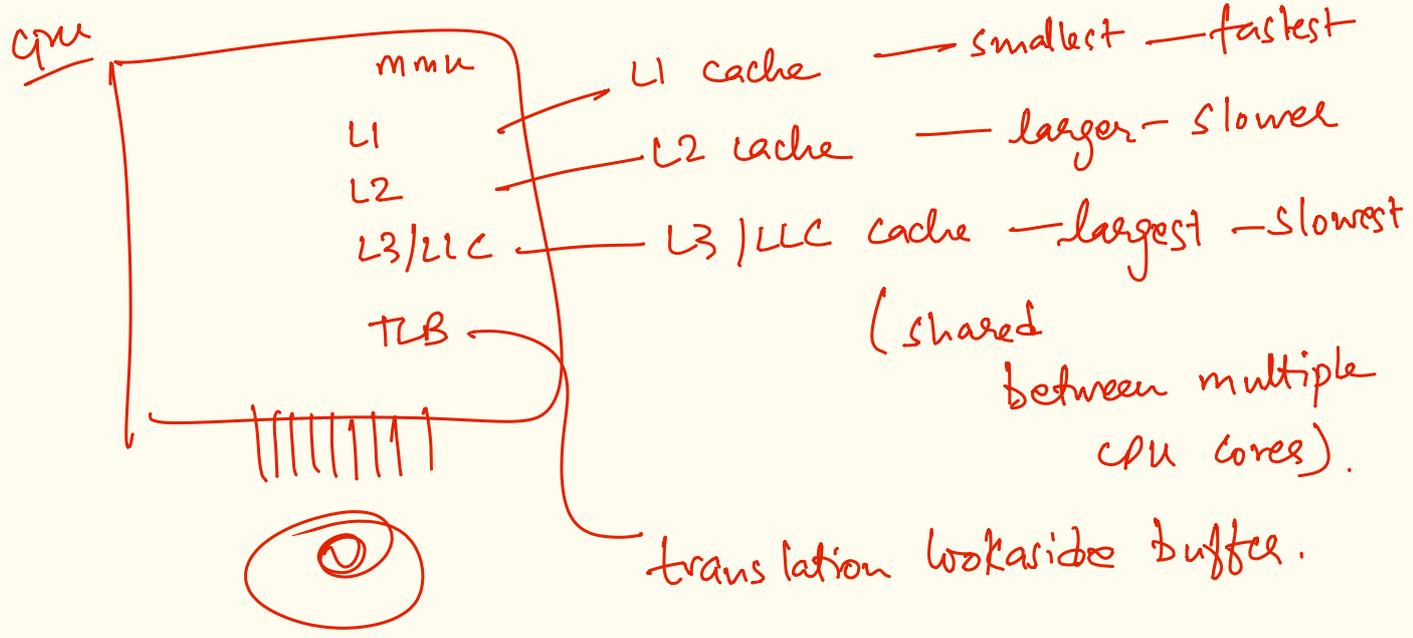
paging is an overhead!

(i) page table walk?

(ii) pages tables need memory!

(i) ISA/h/w has a set of caches

to reduce mmu translation latency.  
potentially



L1 } content caches  
 L2 }  
 L3 } < mem. addr, value >  
           content

TLB - cache of

VPN → PPN

virtual page

physical page

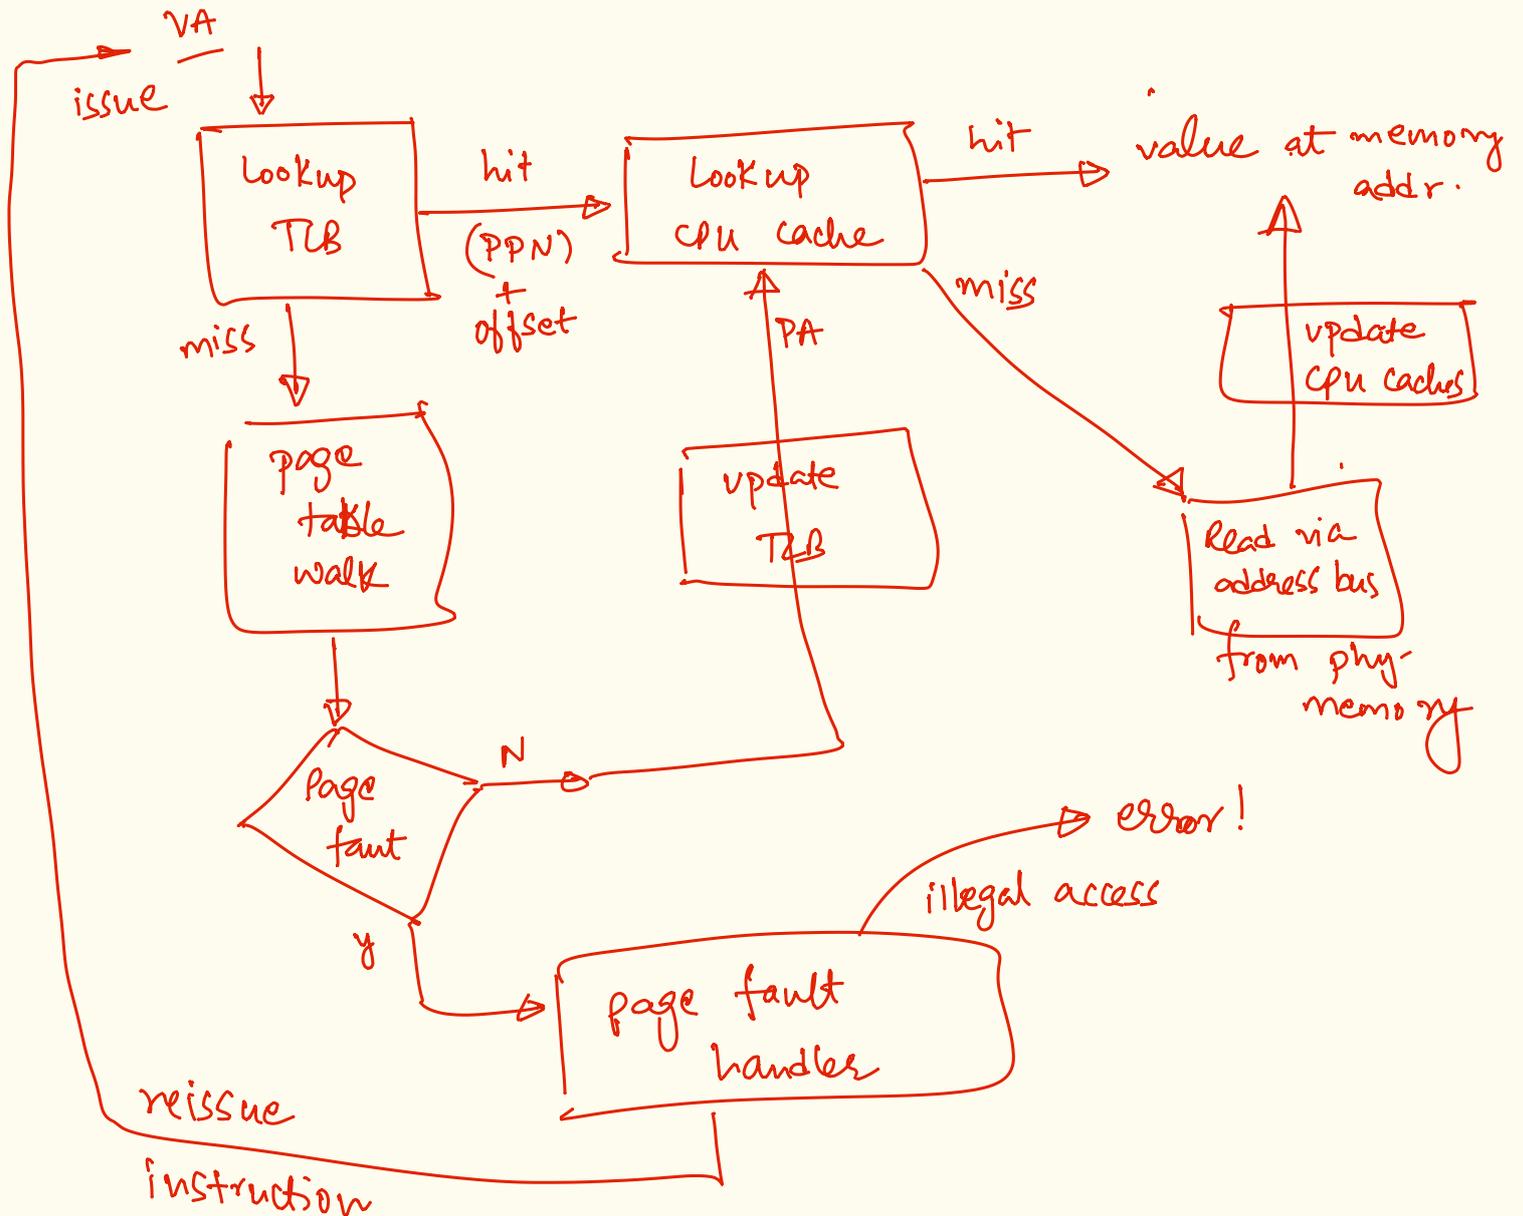
translation cache

(VA) → VPN → PPN

~ TLB's hold VPN  $\rightarrow$  PPN

need to be flushed on process switch!

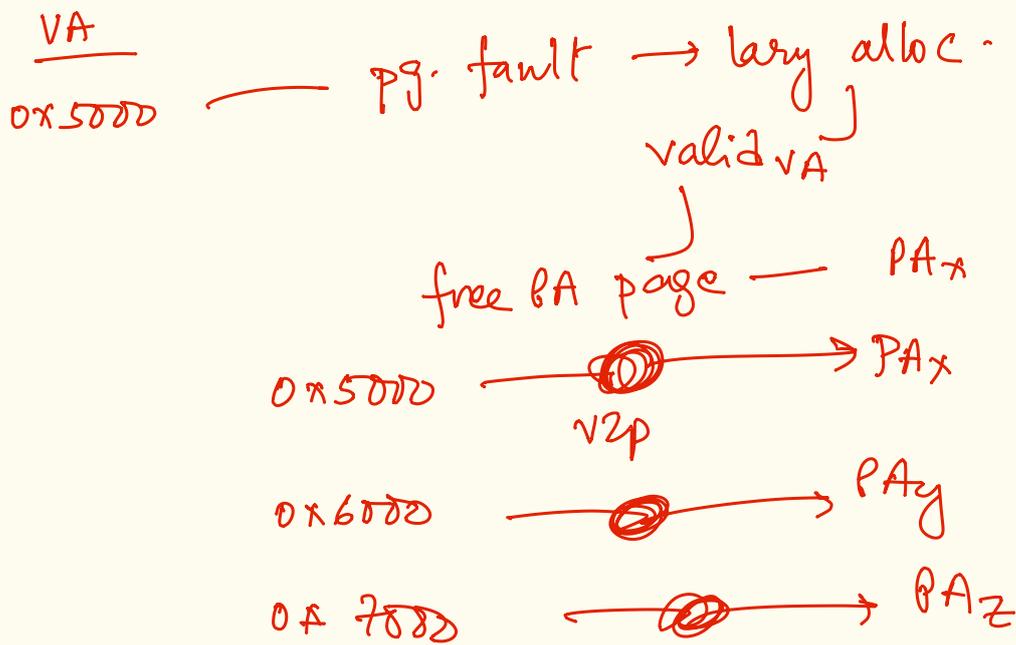
to avoid stale / incorrect page table mappings.



⊗ why is TLB no update by the page fault handler?

(i) ~ pg. fault handler is slow — code  
unless CPU/ISA exposes an instruction for TLB update — this cannot happen.

(ii) handler can do many things!



⊗ why the re-issue?

- note story starts with a VA access!
- pg. fault handler knows the  $\frac{PPN}{PA}$

⊗ how to access the PA?

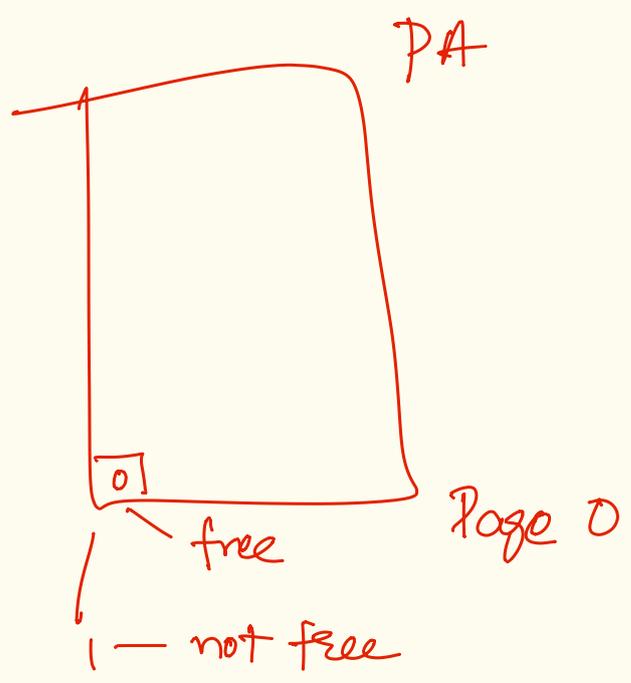
↳ needs of VA.

↳ on reissue use the original VA that was issued.

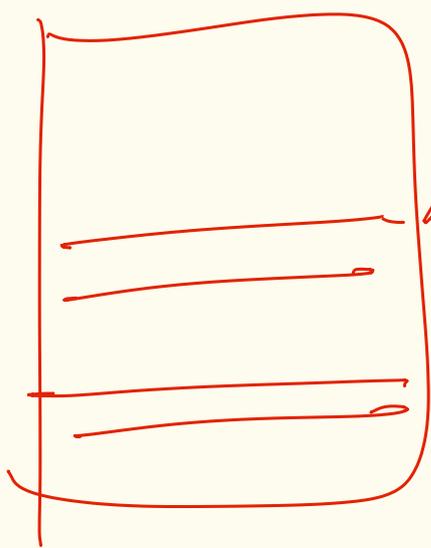
never used (0x) — a Kva that maps to this address!

Kva imp

issue  
~~is~~ reads  
flag in  
PA page



VA  
range



malloc

malloc

range  
VA space

— which  
VA range  
address  
to return?