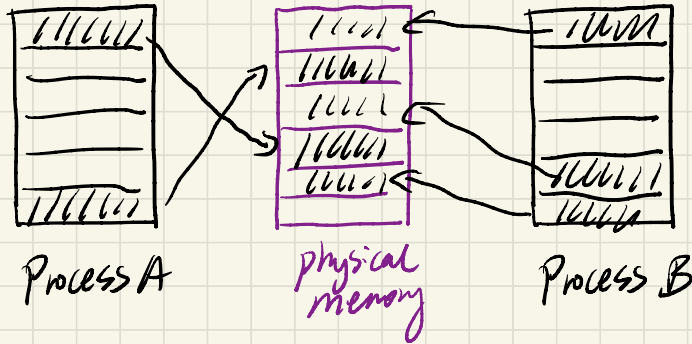


11/1/23

# Paging

VAS  
still  
contiguous




→ base & bound (contiguous translation)

→ paging: translation for each page

page # → frame #


4KB

- fixed sized pages (4096)
- same for physical memory (frames, physical pages)



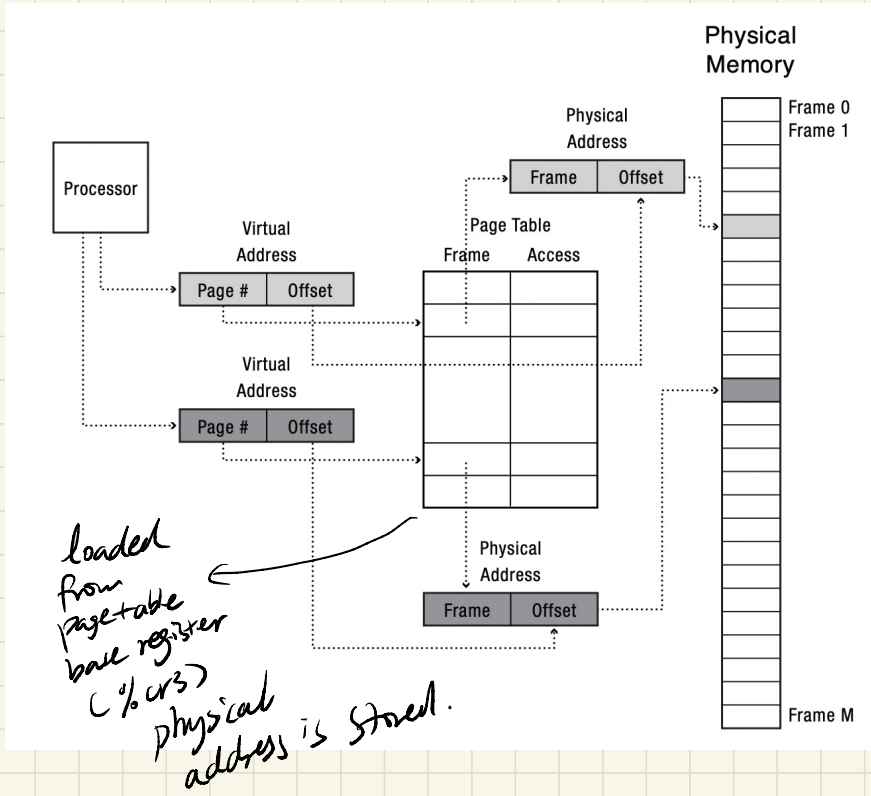
0x1000 - 0x1fff page 1

0x2000 - 0x2fff page 2



# Page Table

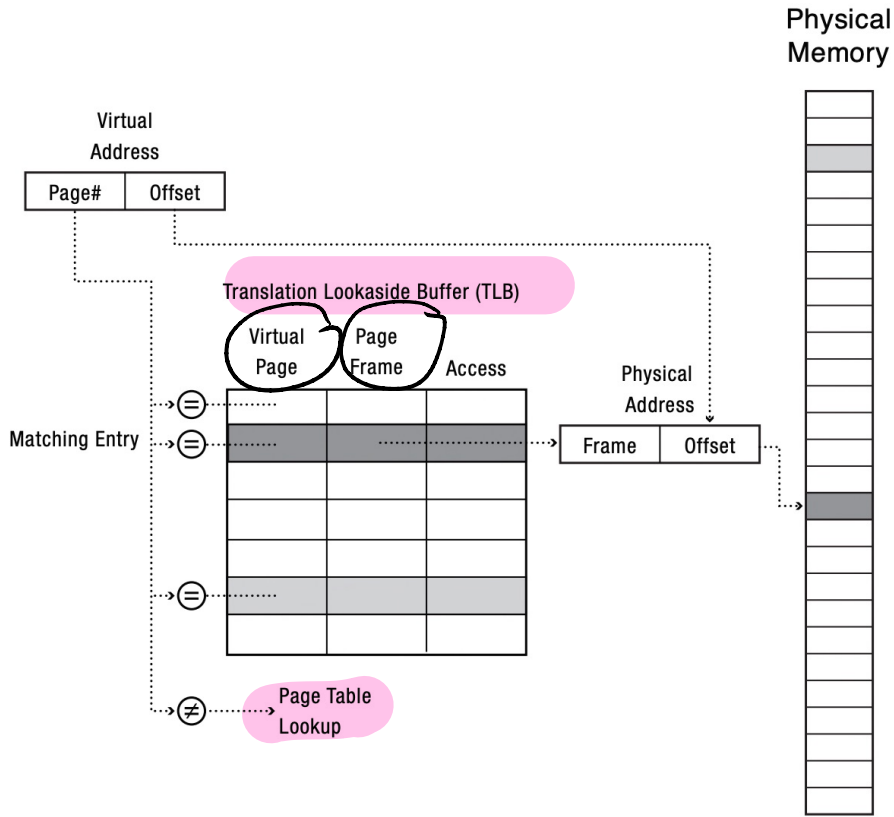
→ a translation array, indexed by page #



- page table walk (translation) done by the hardware
- page table structure defined by architecture

How to speed up translation?  
→ cache the result!

Translation Lookaside Buffer (TLB)  
( $pid + page \#$ )  $\Rightarrow$  frame #



→ Single array: many pages, lots of entries (space), one per process

How to reduce cost?


→ Larger Pages: bigger page size  $\Rightarrow$  less pages & less translation overhead  
(2MB & 1GB supported)

★ problem: internal fragmentation

minimum allocation  
unit is too large waste

→ Inverted Page Table: frame #  $\Rightarrow$  page #, global table

hash function (pid + page #)



takes time, hash collision

→ Multilevel Page Table