

4/24/24

# Paging


→ divide virtual & physical memory into fixed size pages



VAS 1


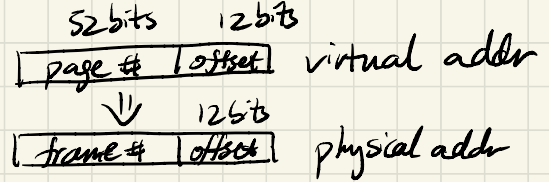


physical memory




VAS 2

→ page level translation

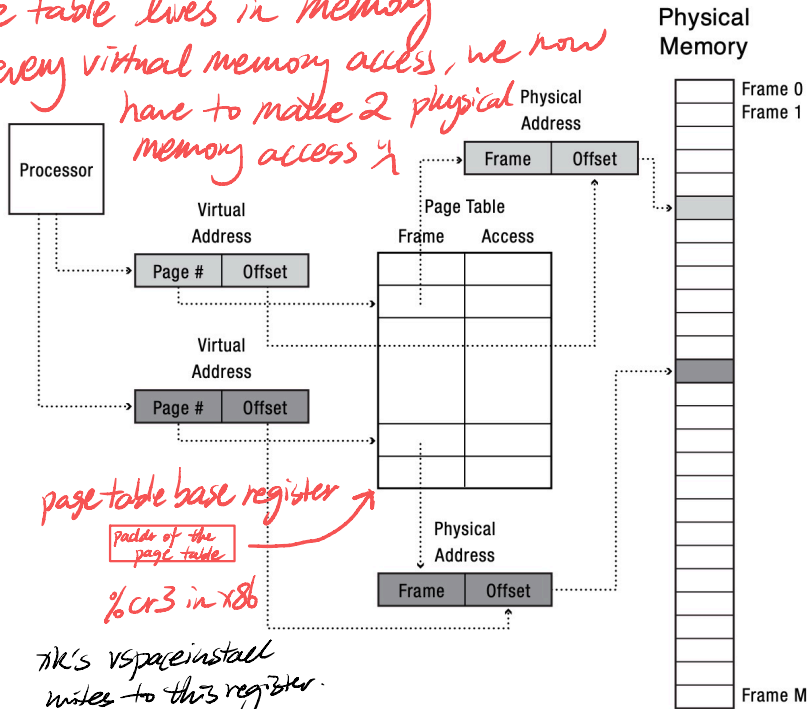


Page Table: stores translations for every page (hw page table walk)

→ kernel sets up the page table, translation performed by hw



\* page table lives in memory  
for every virtual memory access, we now  
have to make 2 physical  
memory access





# The Cost of single array page table

→ per process data structure

→  $2^{52}$  entries take up a lot of space

How to reduce the page table size?

→ Larger page size = 2 MB & 1 GB page size (large / super pages)

↓  
less pages, fewer entries

↓  
smaller page tables

44      20 bits  
page# | offset

$2^{52} \Rightarrow 2^{44}$  entries

but ...

1 byte  
wasted  
///

1 GB

internal fragmentation!

Maybe the problem is too many page tables?

→ Inverted page table = tracks frame mapping instead (global)

indexed  
by frame  
# →

page #, pid

inverted  
page table

→ How to look up given a virtual address?

→ search through each entry of the array until we find a matching page # w/ pid.

works but very slow !!

→ Use a hash function to place pages

→ look up =  $\text{hash}(\text{page \#}, \text{pid}) \% \text{frames} \Rightarrow \text{frame \#}$

what about hash collisions?

shared memory?

cost of unmapped page?

Page  $\Rightarrow$  frame mapping still better for look up!

Multilevel Page Tables = use indirection to only allocate entries for pages in use

