

Worked!

Virtual Memory

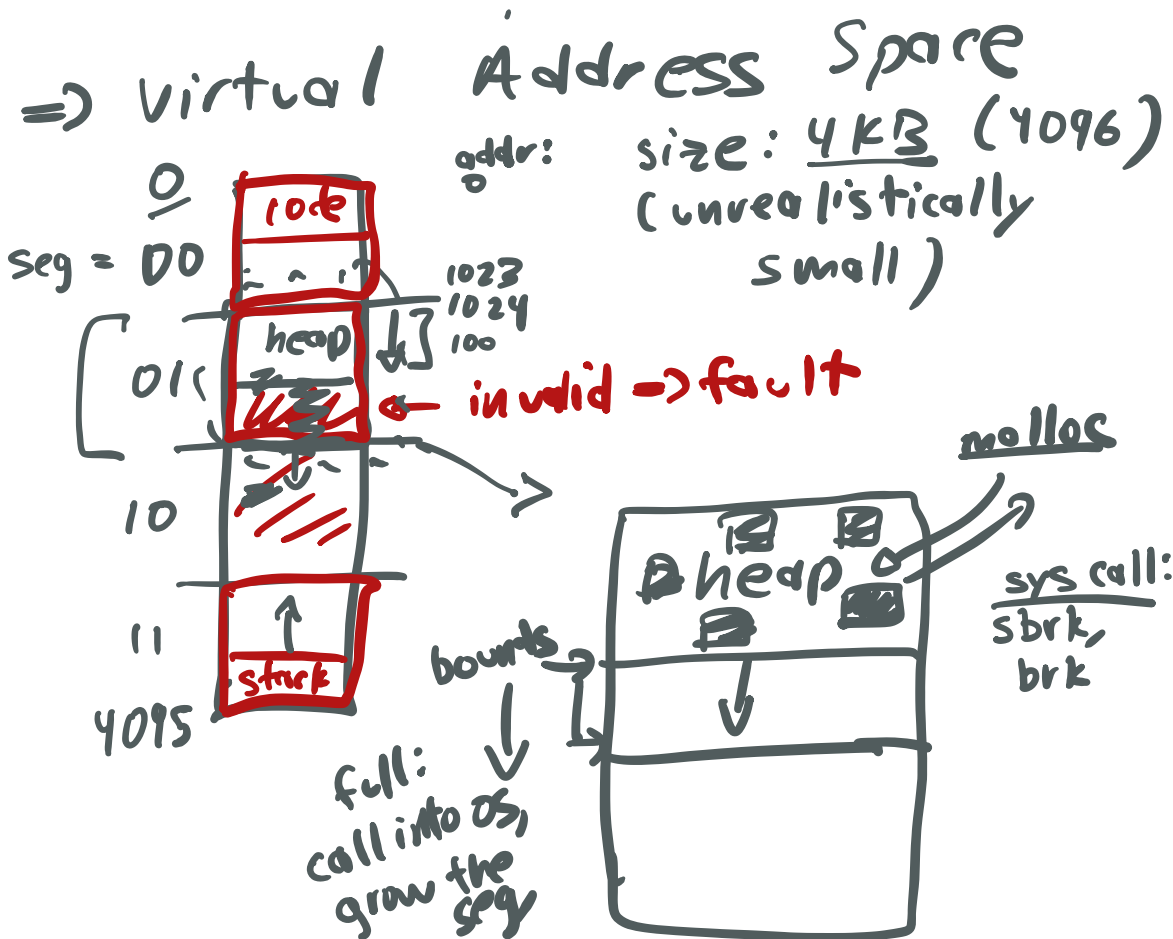
Techniques:

Base / Bounds

Segmentation

Paging

MMU: hardware
per segment
base/bounds



Realistic : 32-bit virtual
addr. space

10 bits → 1 KB (~ 1000)

20 bits → 1 MB (~ million)

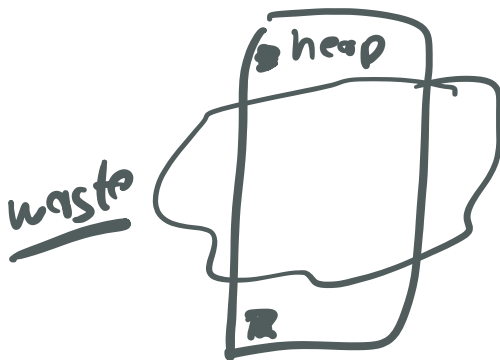
30 bits → 1 GB (~ billion)

Problems w/ segmentation:
(external)

⇒ fragmentation :

free space becomes
full of small used spots
+ free spots intermixed
Phys memory

→ not fully flexible
e.g. "sparse" heap



even though
unused,
must
keep in
memory

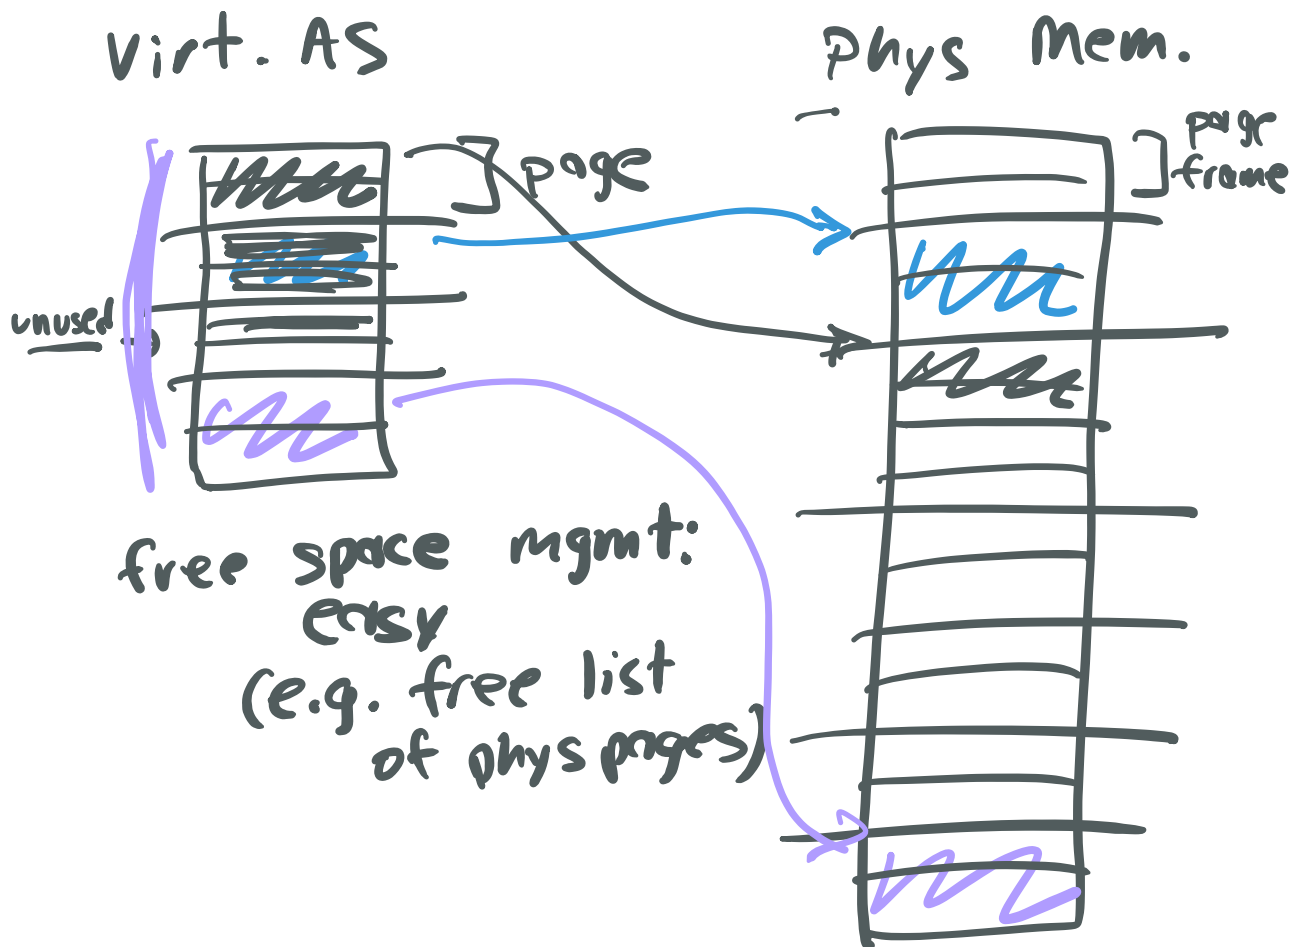


(internal fragmentation)

New approach: Paging

⇒ divide virtual address space
into small, fixed-sized units
(e.g. 4KB)
called pages

⇒ similarly, divide phys mem
into 4KB page frames



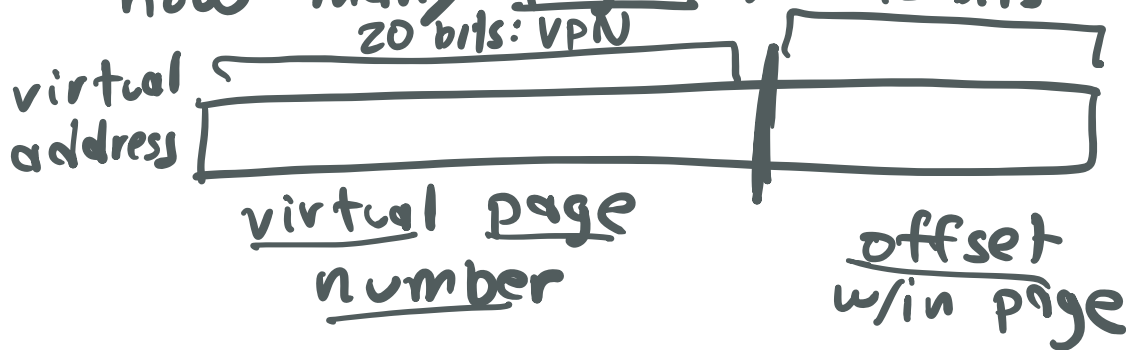
Address

Translation : virtual addr => phys. addr

e.g. virtual address space : 32-bit

page size : [4 KB] ~ 4 GB

how many virtual pages ? => 2^{20} 12 bits





tiny
e.g.

Vpage: 0	0	1
	2	3
VP: 1	4	5
	6	7
VP: 2	8	9
	10	11
VP: 3	12	13
15	14	15

16 bytes : virtual address space size

4 byte page size

virt. Address:



```

0010
0011
0100
0101
0110
0111

```

Problem: ^{#2} too many pages!

e.g. 2^{20} (per process)

(a lot of information)


where to store?

⇒ main memory (managed by OS)

in a data structure

called "page table"

(store virtual → phys
page translations)



page table:
(for process)

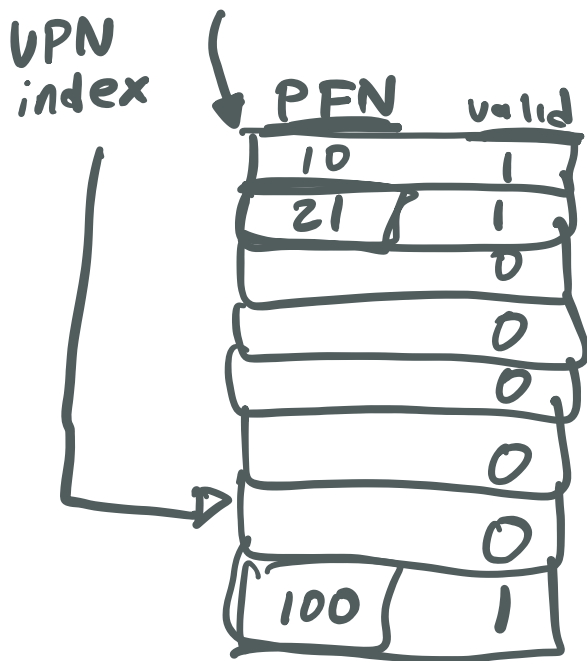
store up to
8 translations

VPN	→ PFN	valid
0	10	1
1	21	1
2		0
3		0
4		0

naturally supports "sparse" addr. space


5	0
6	0
7	100

First page table: linear page table (array)



(array of page table entries)

translation:



hardware: $\left[\text{base address of page table:} \right.$
VPN

$$\underline{B} + (\text{VPN} \times \text{sizeof}(\text{page table entry}))$$

$\left[\text{load PTE from memory} \right.$
extract PFN

form full phys. addr. (PA):

PA

PFN	offset
-----	--------

go to memory for access

=> Slow (extra mem.
ref. per mem ref.)