

1

Main Memory

COS 450 - Fall 2018

2

What is **Main** Memory



...fetch and store from CPU to memory

3

a process in memory...

When does a variable get an address?

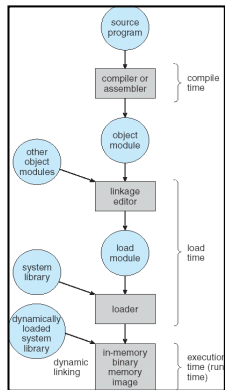
compile-time - absolute

load-time - relocatable

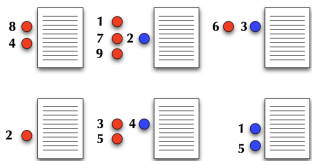
run-time - dynamic



Life of a Program



Multiprocess Environment



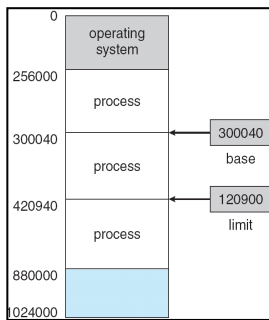
How is memory

organized and protected?

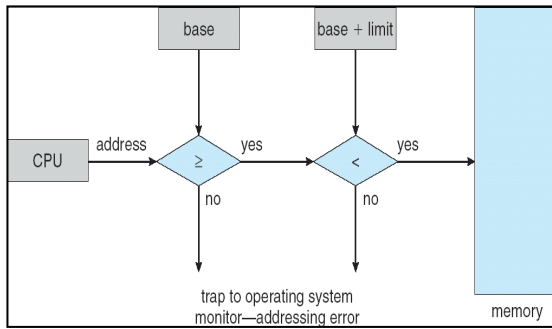
Simple Base + Limit CPU register solution

protection of processes from each other

relocation of code for execution



Base + Limit Registers

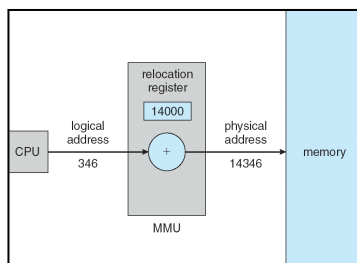


Logical Address Space

address relocation defines a **Logical** address space as well as a **Physical** address space.



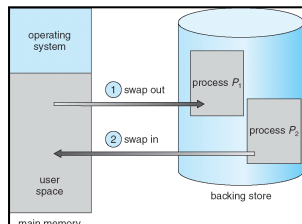
Memory Management Unit



hardware to map logical to physical addresses

Not enough RAM?

- Dynamic Loading
- Dynamic Linking & Shared Libraries
- Swapping



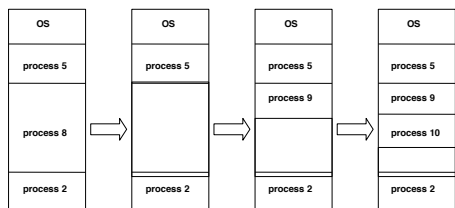
Allocation Strategies

there are two basic strategies;

Contiguous

Non-Contiguous

Contiguous Allocation



Finding Space

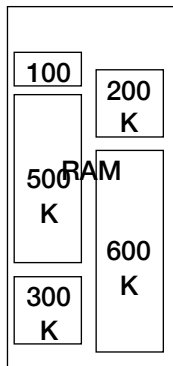
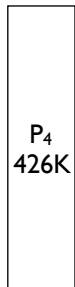
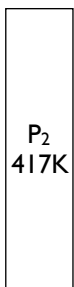
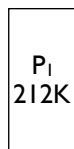
to determine where to place a process we can use;

First-Fit

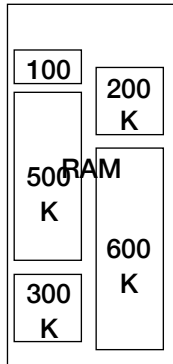
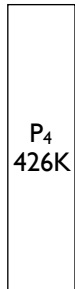
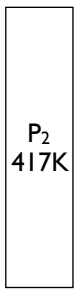
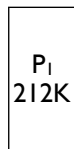
Best-Fit

Worst-Fit

Processes that need Memory...

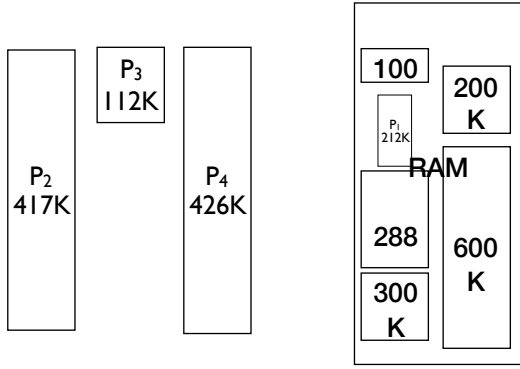


First-Fit Allocation



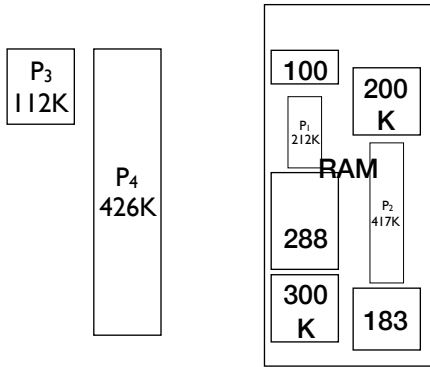
15-2

First-Fit Allocation



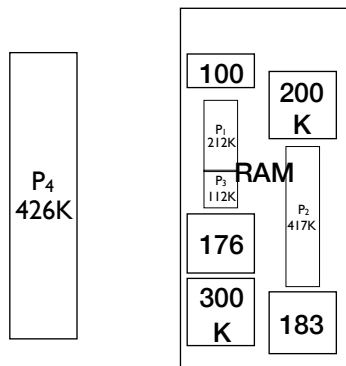
15-3

First-Fit Allocation



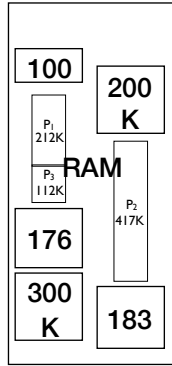
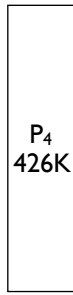
15-4

First-Fit Allocation

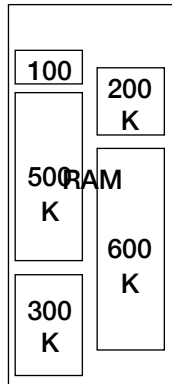
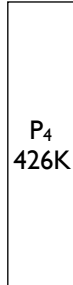
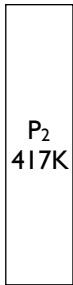
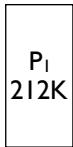


First-Fit Allocation

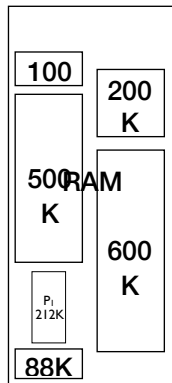
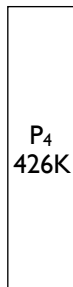
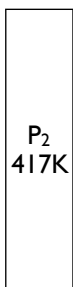
P4 does not fit!



Best-Fit Allocation (work)

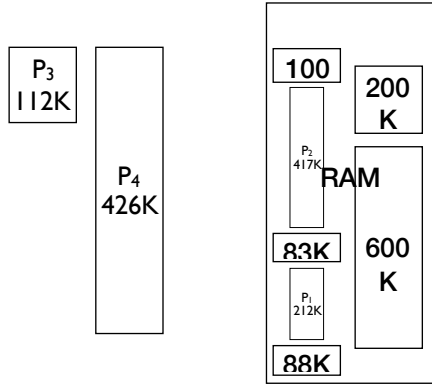


Best-Fit Allocation (work)



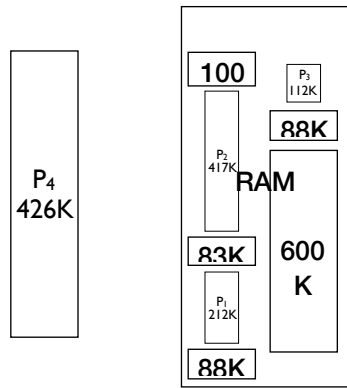
16-3

Best-Fit Allocation (work)



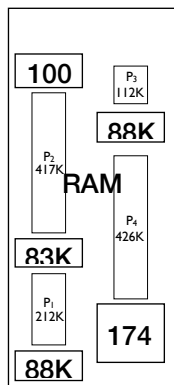
16-4

Best-Fit Allocation (work)

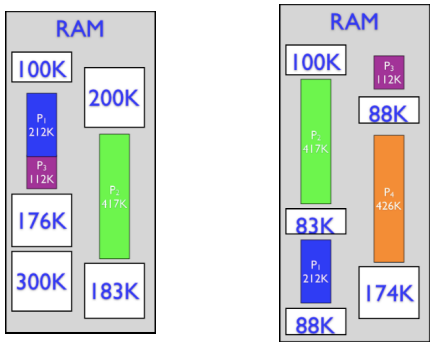


16-5

Best-Fit Allocation (work)



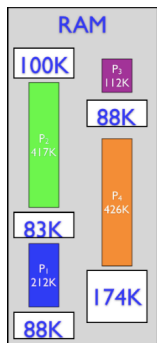
Comparison



Fragmentation

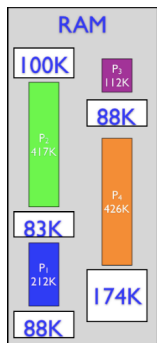
There is enough memory for a 200K process...

...just not all together



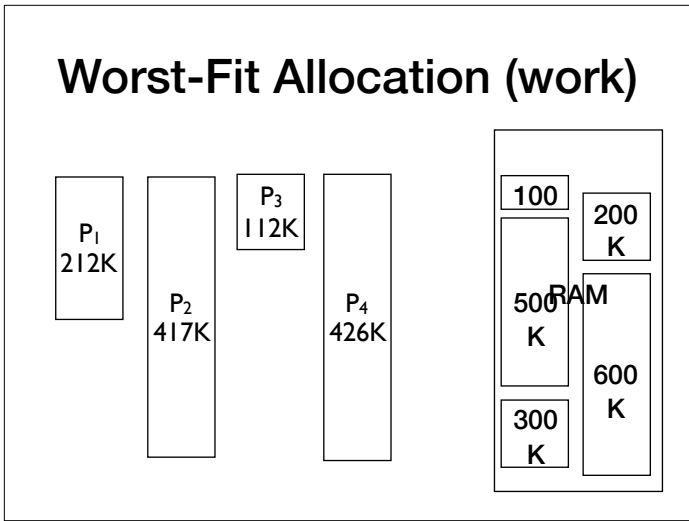
Fragmentation

Little chunks of memory that **cannot be used**



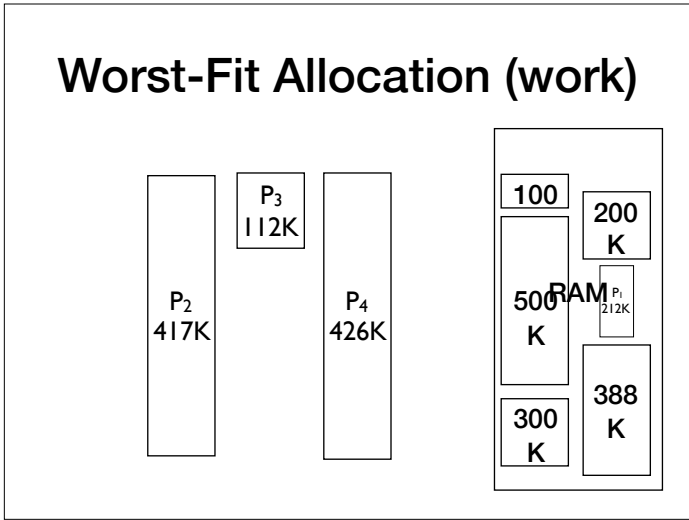
20-1

Worst-Fit Allocation (work)



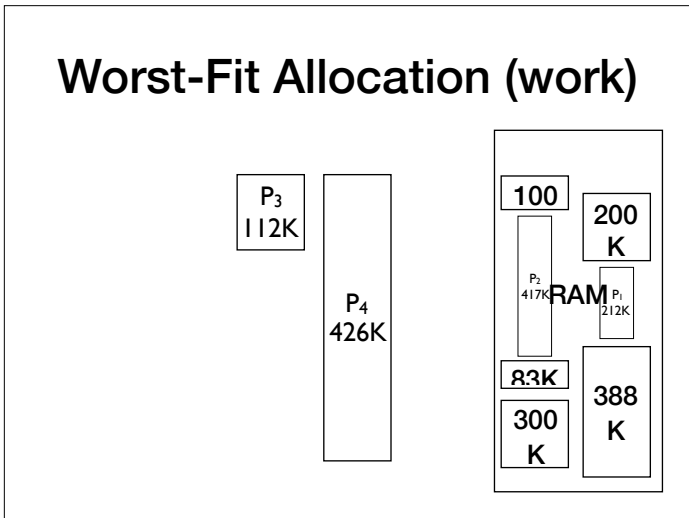
20-2

Worst-Fit Allocation (work)

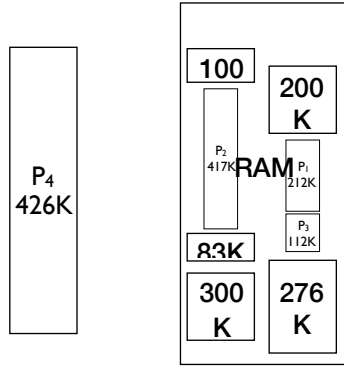


20-3

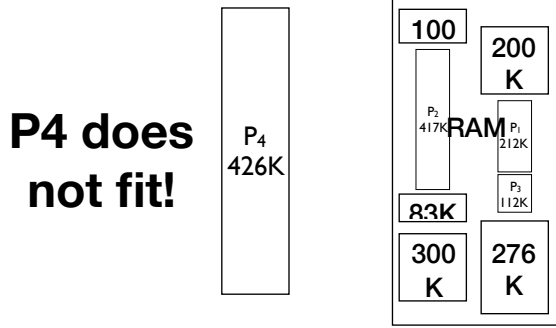
Worst-Fit Allocation (work)



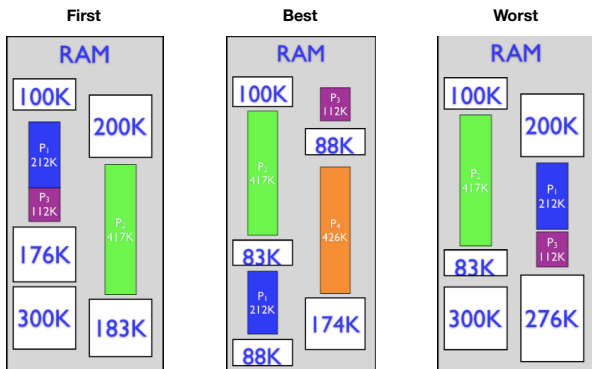
Worst-Fit Allocation (work)



Worst-Fit Allocation (work)



Comparison



Contiguous Allocation

Which is the “best”?

First-Fit - incomplete allocation

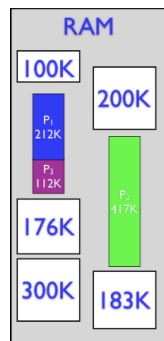
Best-Fit - seems good... expensive

Worst-Fit - incomplete allocation

23-1

Compaction

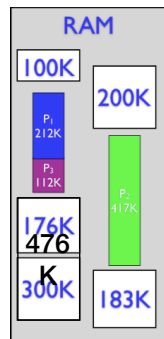
Shuffle allocations
around to remove small bits of
free space... expensive.



23-2

Compaction

Shuffle allocations
around to remove small bits of
free space... expensive.

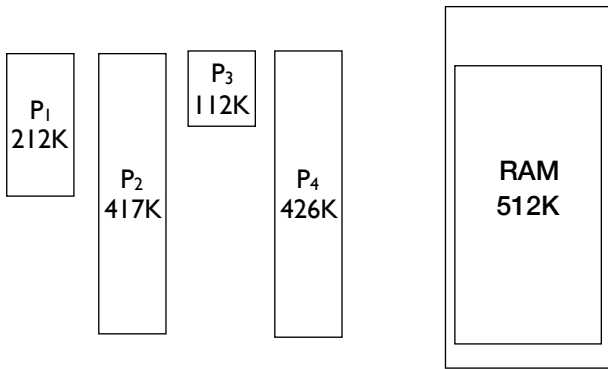


Non-Contiguous Allocation

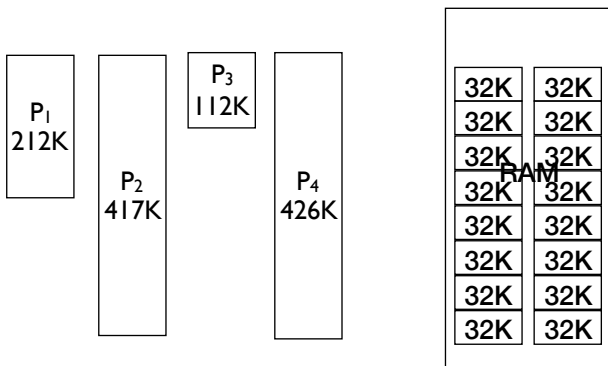
Paged memory

Segmented memory

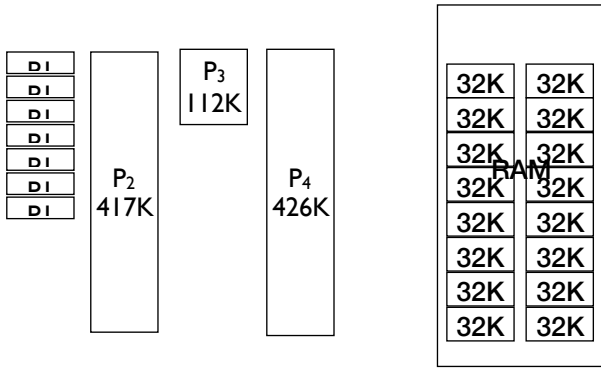
Paged Memory Allocation



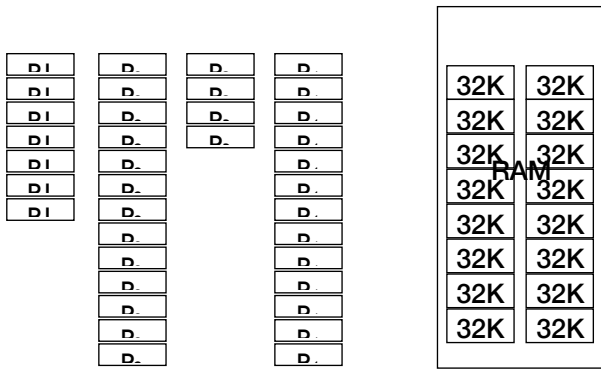
Paged Memory Allocation



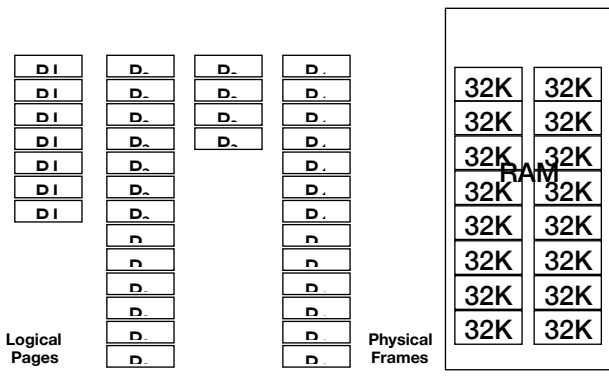
Paged Memory Allocation



Paged Memory Allocation

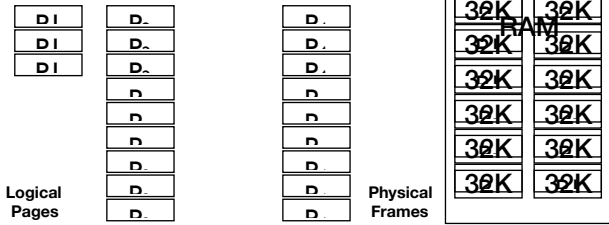


Paged Memory Allocation

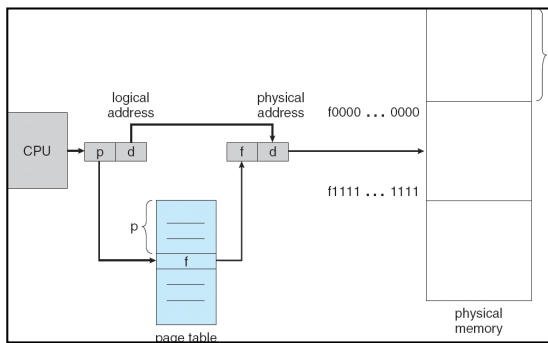


Paged Memory Allocation

Does not all fit
(yet)



Paging Hardware



Paging Details

- Page table per process (PCB pointer)
- Global free frame list
- Shared Pages
- Page tables in main memory

Fragmentation

Little chunks of memory that cannot be used

External Fragmentation (earlier)

Internal Fragmentation

Internal Fragmentation

We can only allocate fixed size chunks of memory (frames)

...on average 1/2 a page/frame per process is lost to fragmentation (so far).

Effective Access Time

To access memory with paging...

we need to access memory **twice!**

$EAT = 2 * \text{memory speed}$

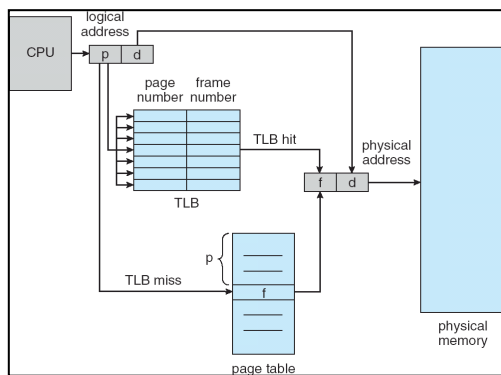
$EAT = 2 * 100ns = 200ns = 200ns$

TLB (cache)

The **Translation Look-aside Buffer**

A **cache** specifically for page table entries

TLB Hardware



TLB Hit Time

When we have a TLB **hit...**

$$\text{EAT} = \text{TLB-time} + \text{memory-speed}$$

$$20\text{ns} + 100\text{ns} = 120\text{ns}$$

TLB Miss Time

When we have a TLB **miss**...

$$\text{EAT} = \text{TLB-time} + \underline{2} * \text{memory-speed}$$

$$20\text{ns} + \underline{2} * 100\text{ns} = \underline{220\text{ns}}$$

Effective Access Time

The TLB hit ratio is key,

$$\text{EAT} = (\text{hit} * 120\text{ns}) + (\text{miss} * 220\text{ns})$$

$$\begin{aligned} \text{EAT} &= (0.8 * 120\text{ns}) + (0.2 * 220\text{ns}) \\ &= 140\text{ns} \end{aligned}$$

$$\begin{aligned} \text{EAT} &= (0.98 * 120\text{ns}) + (0.02 * 220\text{ns}) \\ &= 122\text{ns} \end{aligned}$$

Effective Access Time

The TLB hit ratio is key,

$$\text{EAT} = (\text{hit} * 120\text{ns}) + (\text{miss} * 220\text{ns})$$

$$\begin{aligned} \text{EAT} &= (0.8 * 120\text{ns}) + (0.2 * 220\text{ns}) \\ &= 140\text{ns} \end{aligned}$$

$$\begin{aligned} \text{EAT} &= (0.98 * 120\text{ns}) + (0.02 * 220\text{ns}) \\ &= 122\text{ns} \end{aligned}$$

Effective Access Time

The TLB hit ratio is key,

$$\text{EAT} = (\text{hit} * 120\text{ns}) + (\text{miss} * 220\text{ns})$$

$$\text{EAT} = (0.8 * 120\text{ns}) + (0.2 * 220\text{ns})$$

$$= 140\text{ns}$$

$$\text{EAT} = (0.98 * 120\text{ns}) + (0.02 * 220\text{ns})$$

$$= 122\text{ns}$$

Page Tables are Large

2^{32} (4G) logical address
space with 2^{12} (4K) pages...

over 1 million pages

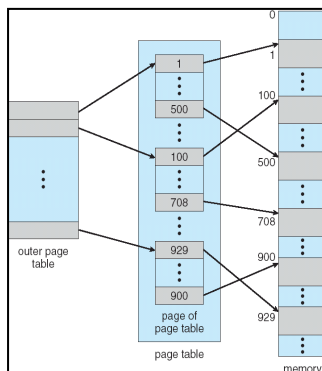
4 bytes each = 4MiB page
table

Page Tables are Large

2^{32} (4G) logical address
space with 2^{12} (4K) pages...

over 1 million pages

4 bytes each = 4MiB page
table

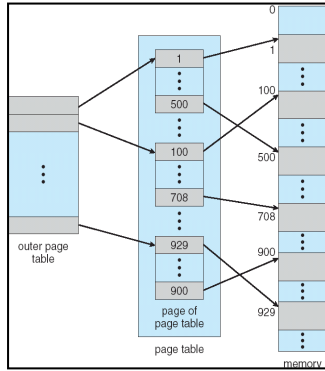


Page Tables are Large

Outer and Inner
pages & tables

page number		page offset
p_1	p_2	d
1	1	1
2	0	0

Outer always in memory
Inner when needed



Segmentation

intel Example