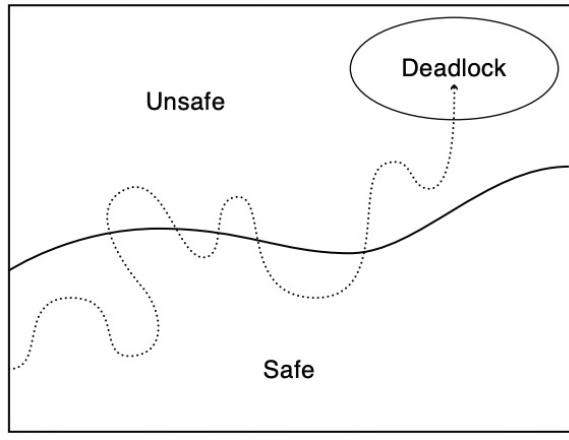


4/22/24

# Deadlock Wrapup & Memory Management

request 1 chopstick  
at a time  
— max 2 requests

## Deadlock Avoidance



3 Chopsticks, 3 philosophers

safe

A	B	C
0 <sup>(2)</sup>	0 <sup>(2)</sup>	0 <sup>(2)</sup>

avail = 3

safe

A ✓	B ✓	C ✓
1 <sup>(1)</sup>	0 <sup>(2)</sup>	0 <sup>(2)</sup>

*can finish*

avail = 2 *starts out w/ avail*

to be avail = 2  
2 - 1 + 2 = 3  
(A request 1 more) (A will finish & return both)

safe

A ✓	B ✓	C ✓
1 <sup>(1)</sup>	0 <sup>(2)</sup>	C <sup>(1)</sup>

avail = 1

to be avail = 1  
1 - 1 + 2 = 2 (A can finish)  
2 - 0 + 2 = 2 (B can finish)  
2 - 1 + 2 = 3 (C can finish)

unsafe

A x	B	C
1 <sup>(1)</sup>	1 <sup>(1)</sup>	1 <sup>(1)</sup>

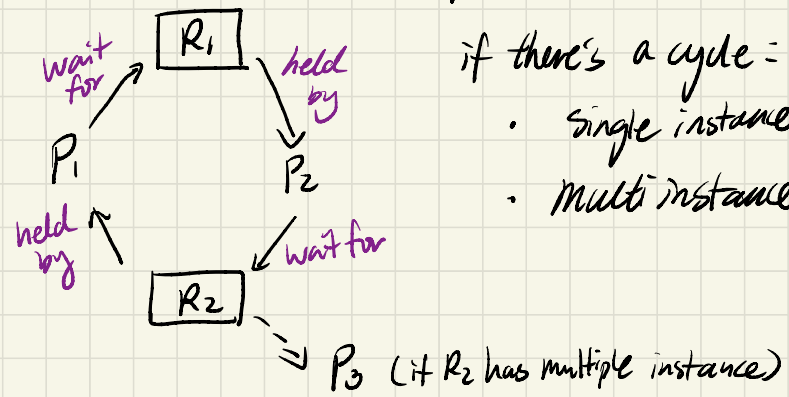
avail = 0

to be avail = 0 no one can finish their future request

# Deadlock Detection

→ rare event, let it happen, detect & recover when it happens

→ Resource Allocation Graph



- if there's a cycle = *\* high false positive*
- single instance resource  $\Rightarrow$  deadlock
  - multi instance resource  $\Rightarrow$  potential deadlock

*Ostrich Algorithm*  
→ pretend nothing's wrong, let user deal w/ it

→ recover from deadlock  
→ abort / terminate a process in the cycle



# Physical Memory Management

- volatile, byte addressable, order of GBs
- ~200 cycles access latency

Problem = many processes, limited physical memory

→ Attempt #1 = let one process use all of physical memory

→ no need for translation

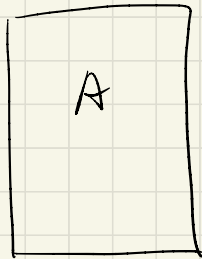
→ context switch

Scheduling overhead

→ write A's VAS to disk

→ load B's VAS into memory

→ how do we enforce kernel/user memory separation



Physical memory

→ Attempt #2: Divide up physical memory among processes

→ low context switch overhead

→ virtual to physical address translation

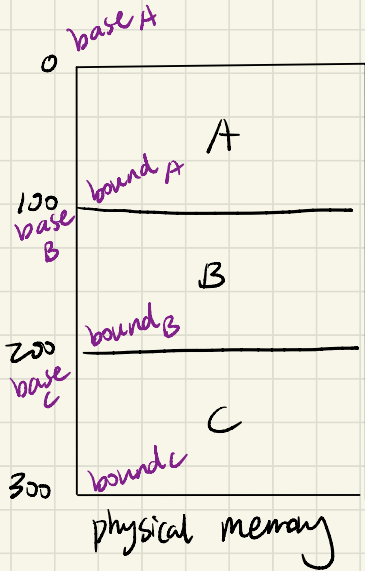
→ base & bound registers (hw supports)

$$PA = VA + base \quad (VA < bound)$$

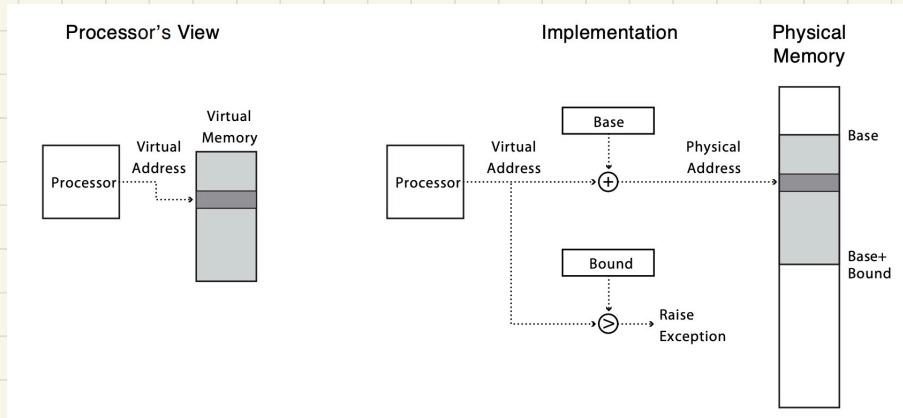
base register 0 for process A  
 bound register 100

→ how do we support memory growth?

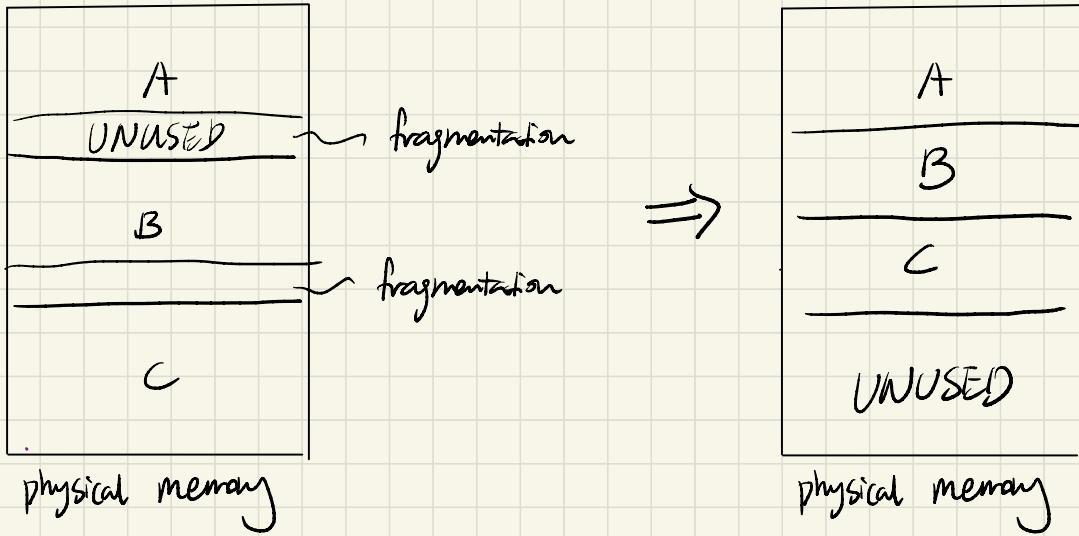
→ variable sized allocation leads to external fragmentation



A  
—  
C D doesn't fit, can't allocate



# Compaction

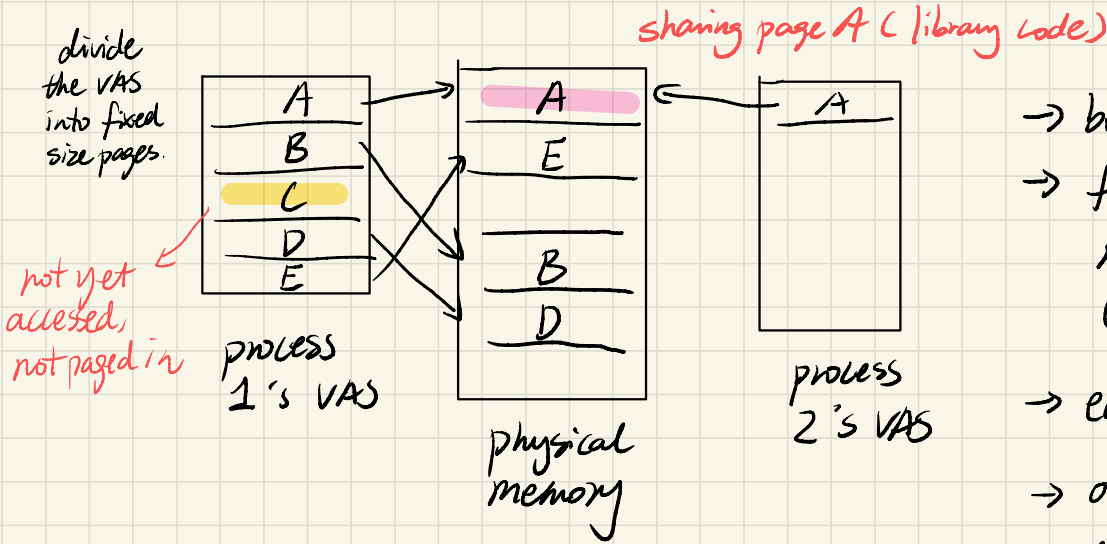


update base & bound registers for each moved (compacted) process

expensive operation, causes variability in performance.

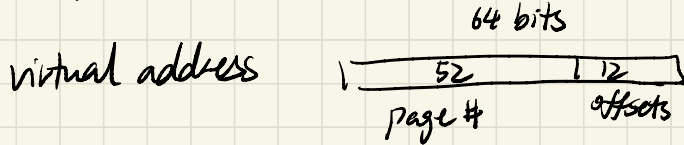
- Can we share library code or support can fork with this approach?
- translation done on the entire VAS, can't share part of it

→ Attempt #3 = finer grained translation & permission



- better control for sharing
- fixed sized allocation (page, 4kB)  
no external fragmentation  
(may see internal ↑ instead)
- easy to support memory growth
- only load pages in use into memory (demand paging)

What should be translated



offset within a page  
= offset within a physical page.  
only need to translate page #