

\*315 exam 2 Review: This is just answer  
NOTE → to your questions and is not  
1) Deadlock avoidance / Banker's alg / Safety alg a  
Detection alg  
replacement  
for your  
Readings!

2) Difference bw scheduling techniques & how they work

3) Strategy for avoiding circular wait

4) Deadlock Outline map. Deadlock → avoidance  
prevention  
Detection

5) How to implement preemption?

6) What is a trap?

7) Optimal / FIFO / LRU

8) Resource allocation Graph

9) Segmentation vs fragmentation

10) example of paging & linking physical memory addr from logical. ↑  
paging activity

11) Difference b/w segmentation & paging

12) VM  $\rightarrow$  switching pages around ?!?

# Deadlock

- In a multiprogramming environment, processes compete for resources. if a process  $i$  requests a resource that is being held by another process  $j$  that is waiting on another resource held by process  $i$   
 $\Rightarrow$  Deadlock

- Conditions for a deadlock? The following 4 should hold simultaneously

1. Mutual Exclusion

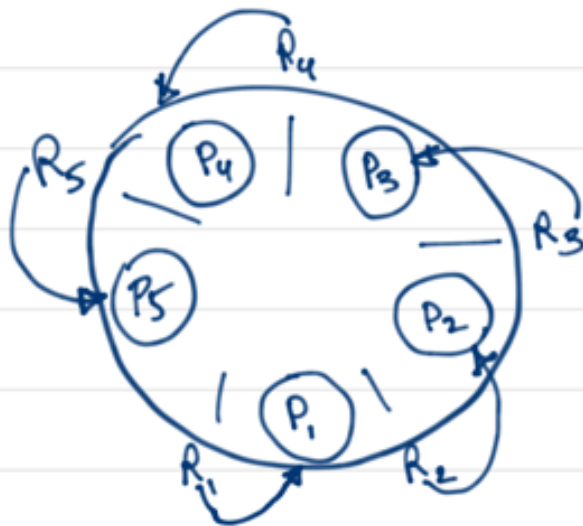
2. Hold and wait: a process is holding at least one resource and waiting to acquire a resource that is being held by another process

3. No Preemption: resources can be only released voluntarily

4. Circular wait: A set of waiting processes  $\{P_0, P_1, \dots, P_n\}$  of waiting processes must exist such that  $P_0$  is waiting on a resource held by  $P_1$  ...  $P_n$  is waiting for a resource held by  $P_0$ .

# Dining Philosophers

1. ME
2. Hold & Wait
3. Circular Wait
4. No Preemption




Models of system resource state:

1) Resource - Allocation Graph: nodes: Processes & Resources

edges: Resource  $\rightarrow$  Process: allocation edge

Process  $\rightarrow$  Resource: Request edge

example:



2) State of the resource-allocation system (Matrix Notation) <sup>vector</sup>  
for  $m$  resource types +  $n$  processes

Available, number of available resource (vector size  $m$ )

Max:  $n \times m$  matrix, maximum demand of each process.

Allocation:  $n \times m$  matrix, number of resources of each type allocated to each process.

Need = Max - Allocation

\*How to handle deadlock?

- 1) avoid or prevent (so we never enter a deadlock situation)
- 2) allow deadlock & then detect & recover
- 3) Ignore it!

1) avoid or prevent:

1-A) Deadlock Prevention: Ensure at least one of the 4 conditions does not hold by some protocols

1-A-i) Mutual Exclusion: not always possible

1-A-ii) Hold & wait: not allow a process which is holding a resource to request new resource. (a) request all resources at beginning (b) release all resources before new request

1-A-iii) No preemption: (a) if a process request some resource



that is not available  $\Rightarrow$  all resources the process is holding currently is preempted. (b) when a process request a resource that is held by some other process, we check if other process is waiting on another resource  $\rightarrow$  preempt the resource.

1-A-iv) Circular Wait: apply some ordering on how resources can be requested. Dining Philosophers: increasing order.

1-B) Deadlock Avoidance: Based on additional info about how resources are to be requested. (Assumes that OS knows what resources will be requested by each process).

+ Safe State: Define a safe state where system can allocate resources to each process (up to its maximum) in "some order" and avoid a deadlock.

Avoidance algorithms:

Request  $\rightarrow$  Safe(?)  $\xrightarrow{\text{Yes}}$  Grant

$\searrow$  No  $\rightarrow$  Don't grant edge

1-B-i) Resource allocation: Possible future requests: chain  
 $\rightarrow$  after each request  $\rightarrow$  safe? (using cycle detection)

(yes  $\rightarrow$  allocate  
No  $\rightarrow$  don't allocate

1-B-ii) Banker's Algorithm: Before granting each request, evaluate and see if the system stays in a safe state after this allocation or not.   
 Yes → Allocate   
 No → don't allocate.

Safe state: Considers the worst case scenario, where all processes request all their needed resources & see if there is one sequence of execution or not.

Yes → safe   
 No → unsafe

example:

	Allocation			Max			Need			Available		
	A	B	C	A	B	C	A	B	C	A	B	C
P <sub>0</sub>	0	1	0	7	5	3	7	4	3	3	3	2
P <sub>1</sub>	2	0	0	3	2	2	1	2	2	5	3	2
P <sub>2</sub>	3	0	2	9	0	2	6	0	0	7	4	3
P <sub>3</sub>	2	1	1	2	2	2	0	1	1	7	4	5
P <sub>4</sub>	0	0	2	4	3	3	4	3	1	10	4	7

System currently in a

(P<sub>1</sub>, P<sub>3</sub>, P<sub>4</sub>, P<sub>2</sub>, P<sub>0</sub>) ⇒ Safe State

1) Can a request (1, 0, 2) be granted? (Available: 3, 3, 2)   
 Yes safe

2) Can a request  $(0, 2, 0)$  by  $P_0$  be granted?  
after previous request?

$P_1: (1, 0, 2)$

	Allocation			Max			Need			Available		
	A	B	C	A	B	C	A	B	C	A	B	C
$P_0$	0	1	0	7	5	3	7	4	3	2	3	0
$P_1$	3	0	2	3	2	2	0	2	0			
$P_2$	3	0	2	9	0	2	6	0	0			
$P_3$	2	1	1	2	2	2	0	1	1			
$P_4$	0	0	2	4	3	3	4	3	1			

Safe  $(P_1, P_3, P_4, P_0, P_2)$

$P_0: (0, 2, 0)$

	Allocation			Max			Need			Available		
	A	B	C	A	B	C	A	B	C	A	B	C
$P_0$	0	3	0	7	5	3	7	2	3	2	1	0
$P_1$	3	0	2	3	2	2	0	2	0			
$P_2$	3	0	2	9	0	2	6	0	0			
$P_3$	2	1	1	2	2	2	0	1	1			
$P_4$	0	0	2	4	3	3	4	3	1			

Safe? No  $\Rightarrow$  Revert back

## 2) Deadlock Detection & Recovery:

- Detection
- (A) 2-A-i) Cycles in a Resource allocation graph.  
wait-for-graph: all for single resource
- 2-A-ii) Deadlock detection algorithm (Similar to Banker's)

[ Available  
Allocation  
Request ] if there is a sequence of execution for the  
current set of requests ~~Yes~~ No deadlock  
No ~~Yes~~ Deadlock

Key Question: What is the difference bw this alg & Banker's?

## 2-B.: Deadlock Recovery: How to recover from it?


2-B-i) Process termination: (a) all processes: cost?  
(b) one by one: cost?

2-B-ii) Resource Preemption: First choose a victim process based on some cost, second: roll-back the process to a safe state (or abort & restart), third watch out for starvation how?



## Key Questions:

- 1) How can we detect a deadlock? (single vs multi resource)
- 2) What happens to the system if a deadlock happens?



Can grow very quickly!

- 3) How can we prevent a deadlock?

4) if we have a priori info about all processes & their future requests, can we avoid deadlock? How? (single resource vs. multiple)

# CPU Scheduling

1) FCFS

2) SJF (Preemptive & non-preemptive)


3) RR (with diff quantum size)

4) Priority (Preemptive & non-preemptive)

- Multi level feedback queue : diff schedulers for diff queues.

# Memory Management :

Each Computer System has a limited amount of memory which should be shared among different processes. Processes enter/exit the system & hence the need for a manager to keep track of & manage the memory space is crucial.



## 1) Contiguous Memory Allocation :

1-A: Best fit

1-B: First fit

1-C: Worst fit

## 2) Segmentation :

decrease external fragmentation by breaking down each process address space into smaller pieces (segments).

3) Paging : Divide process's addr space into equally sized pieces called pages and fit those into free frames available throughout the memory.

Very powerful technique that can help us to implement virtual memory.

# Virtual Memory:

- There is no need to load the whole addr space into memory  $\Rightarrow$  why not just load the pages that are being accessed by the process (Demand Paging)?

- Also, what if a process is larger than the physical memory?

- Higher degree of multiprogramming  $\Rightarrow$  Higher CPU utilization  
Higher throughput

$\Rightarrow$  OS always want to achieve the highest throughput  $\Rightarrow$

if the utilization is low  $\Rightarrow$  increase the degree of multiprogramming

$\Rightarrow$  what can go wrong?  $\Rightarrow$  high page fault rate  $\Rightarrow$

long queue for backing store  $\Rightarrow$  low CPU utilization

$\Rightarrow$  Thrashing: process is busy swapping pages in & out!

Why VM worked?  $\uparrow$  if swapping too much  $\Rightarrow$

$\Rightarrow$  How much frames shall we allocate each process?

Global vs. Local allocation

Priority  $\leftarrow$

$\rightarrow$  Fixed allocation

$\rightarrow$  result in Thrashing where more processes in memory than optimal.


One process paging in & out  $\Rightarrow$  swap device slow the system  $\Rightarrow$  busy



\* What was a page fault?

No Free Frame!

Process is running & its tries to access a page that is not loaded into memory yet.



Probability of page fault:  $0 \leq p \leq 1 \rightarrow$  hit  $(1-p)$  <sup>nsec</sup> 8,000,000

Effective access time? Page fault overhead =  $\Delta = 8 \text{ msec}$   
 memory access =  $\delta = 200 \text{ nsec}$

$EAT = (1-p)\delta + p\Delta$  if  $p = 0.001$  0.1% <sup>Huge gap</sup>

acceptable EAT  $\Rightarrow$  way lower  $p$   
 EAT 8200ns  $\rightarrow$  41 times slower than memory access

Why does this work?  $\leftarrow$  1 every 400,000 time!

# Virtual Memory :

Now that we have the underlying paging mechanism, we can take advantage of that and separate the process view from the actual limitations of Physical Memory.



how many 512 MB  
processes? 16



if just use 128 MB  
for each process? 80 processes!



higher throughput, higher  
CPU utilization

As long as the process has access to what it needs for execution  $\Rightarrow$  it does not have to know what is going on at low level.


1) we don't need to load the whole process

2) Using page fault mechanism we can give the process illusion of having all its logical addr space.

System with 44 bits logical addr has 16 GB RAM. How many bits necessary to address all frames in RAM? Show the Mapping from logical to physical by MMU if page size is 16 KB. ( $2^{14}$ )

- How big can a page table be?  $\rightarrow$  Make sense? How can we fit it into RAM?

30  
2 page table : 1GB : Hierarchical



or 10 and 20.

- How many frames should we allocate to each process?

1) proportional to its size

2) , , , priority