

LECTURE - XII  
**DEADLOCKS &  
MAIN MEMORY MANAGEMENT**

Tevfik Koşar

University at Buffalo  
October 11<sup>th</sup>, 2011

## Roadmap

- **Deadlocks**
  - Resource Allocation Graphs
  - Deadlock Detection
  - Deadlock Prevention
  - **Deadlock Avoidance**
  - **Deadlock Recovery**
- Main Memory Management



2

## Deadlock Avoidance

**Deadlock Prevention:** prevent deadlocks by restraining resources and making sure one of 4 necessary conditions for a deadlock does not hold. (system design)

--> possible side effect: low device utilization and reduced system throughput

**Deadlock Avoidance:** Requires that the system has some additional *a priori* information available. (dynamic request check)

i.e. request disk and then printer..

or request at most n resources

--> allows more concurrency

- **Similar to the difference between a traffic light and a police officer directing the traffic!**

3

## Deadlock Avoidance

- Simplest and most useful model requires that each process declare the *maximum number* of resources of each type that it may need.
- The deadlock-avoidance algorithm dynamically examines the resource-allocation state to ensure that there can never be a circular-wait condition.
- Resource-allocation *state* is defined by the number of available and allocated resources, and the maximum demands of the processes.

4

## Example

**P1:**

Request Disk  
Request Printer

....

Release Printer  
Release Disk

**P2:**

Request Printer  
Request Disk

....

Release Disk  
Release Printer

5

## Safe State

- A state is **safe** if the system can allocate resources to each process (upto its maximum) in some order and can still avoid a deadlock.
- When a process requests an available resource, system must decide if immediate allocation leaves the system in a safe state.
- System is in safe state if there exists a **safe sequence** of all processes.

6

## Safe State

- Sequence  $\langle P_1, P_2, \dots, P_n \rangle$  is safe if for each  $P_i$ , the resources that  $P_i$  can still request can be satisfied by currently available resources + resources held by all the  $P_j$ , with  $j < i$ .
  - If  $P_i$  resource needs are not immediately available, then  $P_i$  can wait until all  $P_j$  have finished.
  - When  $P_j$  is finished,  $P_i$  can obtain needed resources, execute, return allocated resources, and terminate.
  - When  $P_i$  terminates,  $P_{i+1}$  can obtain its needed resources, and so on.
- If no such sequence exists, the state is **unsafe!**


7

## Basic Facts

- If a system is in safe state  $\Rightarrow$  no deadlocks.
- If a system is in unsafe state  $\Rightarrow$  possibility of deadlock.
- Avoidance  $\Rightarrow$  ensure that a system will never enter an unsafe state.

8

## Safe, Unsafe, Deadlock State



9

## Example

Consider a system with 3 processes and 12 disks.

At  $t = t_0$ ;

	<u>Maximum Needs</u>	<u>Current Allocation</u>
P1	10	5
P2	4	2
P3	9	2

10

## Example (cont.)

Consider a system with 3 processes and 12 disks.

At  $t = t_1$ ;

	<u>Maximum Needs</u>	<u>Current Allocation</u>
P1	10	5
P2	4	2
P3	9	3


11

## Resource-Allocation Graph Algorithm

- Claim edge*  $P_i \rightarrow R_j$  indicated that process  $P_i$  may request resource  $R_j$ ; represented by a dashed line.
- Claim edge converts to request edge when a process requests a resource.
- When a resource is released by a process, assignment edge reconverts to a claim edge.
- Resources must be claimed *a priori* in the system.


12

### Resource-Allocation Graph For Deadlock Avoidance



13

### Unsafe State In Resource-Allocation Graph



14

### Banker's Algorithm

- Works for multiple resource instances.
- Each process declares maximum # of resources it may need.
- When a process requests a resource, it may have to wait if this leads to an unsafe state.
- When a process gets all its resources it must return them in a finite amount of time.

15

### Data Structures for the Banker's Algorithm

Let  $n$  = number of processes, and  $m$  = number of resources types.

- **Available:** Vector of length  $m$ . If available  $[j] = k$ , there are  $k$  instances of resource type  $R_j$  available.
- **Max:**  $n \times m$  matrix. If  $Max[i,j] = k$ , then process  $P_i$  may request at most  $k$  instances of resource type  $R_j$ .
- **Allocation:**  $n \times m$  matrix. If  $Allocation[i,j] = k$  then  $P_i$  is currently allocated  $k$  instances of  $R_j$ .
- **Need:**  $n \times m$  matrix. If  $Need[i,j] = k$ , then  $P_i$  may need  $k$  more instances of  $R_j$  to complete its task.

$$Need[i,j] = Max[i,j] - Allocation[i,j].$$

16

### Safety Algorithm

1. Let **Work** and **Finish** be vectors of length  $m$  and  $n$ , respectively. Initialize:

$$Work = Available$$

$$Finish[i] = false \text{ for } i = 1, 2, \dots, n.$$

2. Find an  $i$  such that both:
  - (a)  $Finish[i] = false$
  - (b)  $Need_i \leq Work$
 If no such  $i$  exists, go to step 4.
3.  $Work = Work + Allocation_i$   
 $Finish[i] = true$   
 go to step 2.
4. If  $Finish[i] == true$  for all  $i$ , then the system is in a safe state.

17

### Resource-Request Algorithm for Process $P_i$

Let  $Request_i$  be the request vector for process  $P_i$ .

If  $Request_i[j] = k$  then process  $P_i$  wants  $k$  instances of resource type  $R_j$ .

1. If  $Request_i \leq Need_i$  go to step 2. Otherwise, raise error condition, since process has exceeded its maximum claim.
2. If  $Request_i \leq Available$ , go to step 3. Otherwise  $P_i$  must wait, since resources are not available.
3. Pretend to allocate requested resources to  $P_i$  by modifying the state as follows:

$$Available = Available - Request_i;$$

$$Allocation_i = Allocation_i + Request_i;$$

$$Need_i = Need_i - Request_i;$$

- If safe  $\Rightarrow$  the resources are allocated to  $P_i$ .
- If unsafe  $\Rightarrow P_i$  must wait, and the old resource-allocation state is restored

18

## Example of Banker's Algorithm

- 5 processes  $P_0$  through  $P_4$ ; 3 resource types: A (10 instances), B (5 instances), and C (7 instances).
- Snapshot at time  $T_0$ :

	<u>Allocation</u>			<u>Max</u>			<u>Available</u>		
	A	B	C	A	B	C	A	B	C
$P_0$	0	1	0	7	5	3	3	3	2
$P_1$	2	0	0	3	2	2			
$P_2$	3	0	2	9	0	2			
$P_3$	2	1	1	2	2	2			
$P_4$	0	0	2	4	3	3			

19

## Example of Banker's Algorithm

- The content of the matrix. Need is defined to be Max - Allocation.

	<u>Need</u>		
	A	B	C
$P_0$	7	4	3
$P_1$	1	2	2
$P_2$	6	0	0
$P_3$	0	1	1
$P_4$	4	3	1

20

## Example of Banker's Algorithm

- Snapshot at time  $T_0$ :

	<u>Allocation</u>			<u>Max</u>			<u>Available</u>			<u>Need</u>		
	A	B	C	A	B	C	A	B	C	A	B	C
$P_0$	0	1	0	7	5	3	3	3	2	7	4	3
$P_1$	2	0	0	3	2	2				1	2	2
$P_2$	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

21

## Example of Banker's Algorithm

- Snapshot at time  $T_0$ :

	<u>Allocation</u>			<u>Max</u>			<u>Available</u>			<u>Need</u>		
	A	B	C	A	B	C	A	B	C	A	B	C
$P_0$	0	1	0	7	5	3	3	3	2	7	4	3
$P_1$	2	0	0	3	2	2				1	2	2
$P_2$	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

- The system is in a safe state since the sequence  $\langle P_1, P_3, P_4, P_2, P_0 \rangle$  satisfies safety criteria.

22

## Example: $P_1$ Requests (1,0,2)

- Check that Request  $\leq$  Available (that is, (1,0,2)  $\leq$  (3,3,2))  $\Rightarrow$  true.

	<u>Allocation</u>			<u>Need</u>			<u>Available</u>		
	A	B	C	A	B	C	A	B	C
$P_0$	0	1	0	7	4	3	2	3	0
$P_1$	3	0	2	0	2	0			
$P_2$	3	0	1	6	0	0			
$P_3$	2	1	1	0	1	1			
$P_4$	0	0	2	4	3	1			

- Executing safety algorithm shows that sequence  $\langle P_1, P_3, P_4, P_0, P_2 \rangle$  satisfies safety requirement.
- Can request for (3,3,0) by  $P_4$  be granted?
- Can request for (0,2,0) by  $P_0$  be granted?

23

## Recovery from Deadlock: Process Termination

- Abort all deadlocked processes. --> expensive
- Abort one process at a time until the deadlock cycle is eliminated. --> overhead of deadlock detection alg.
- In which order should we choose to abort?
  - Priority of the process.
  - How long process has computed, and how much longer to completion.
  - Resources the process has used.
  - Resources process needs to complete.
  - How many processes will need to be terminated.
  - Is process interactive or batch?

24

## Recovery from Deadlock: Resource Preemption

- Selecting a victim - minimize cost.
- Rollback - return to some safe state, restart process for that state.
- Starvation - same process may always be picked as victim, include number of rollback in cost factor.


25

## Main Memory Management

26

## Memory Management Requirements

- The O/S must fit multiple processes in memory
  - ✓ memory needs to be subdivided to accommodate multiple processes
  - ✓ memory needs to be allocated to ensure a reasonable supply of ready processes so that the CPU is never idle
  - ✓ memory management is an **optimization** task under **constraints**



Fitting processes into memory is like fitting boxes into a fixed amount of space

27

## Memory Allocation


- Fixed-partition allocation
    - Divide memory into fixed-size partitions
    - Each partition contains exactly one process
    - **The degree of multi programming is bound by the number of partitions**
    - When a process terminates, the partition becomes available for other processes
- no longer in use

OS
process 5
process 9
process 10
process 2

28

## Memory Allocation (Cont.)

- Variable-partition Scheme (Dynamic)
  - When a process arrives, search for a hole large enough for this process
  - Hole - block of available memory; holes of various size are scattered throughout memory
  - Allocate only as much memory as needed
  - Operating system maintains information about:
    - a) allocated partitions
    - b) free partitions (hole)



29

## Dynamic Storage-Allocation Problem

How to satisfy a request of size  $n$  from a list of free holes

- **First-fit**: Allocate the *first* hole that is big enough
- **Best-fit**: Allocate the *smallest* hole that is big enough; must search entire list, unless ordered by size. Produces the smallest leftover hole.
- **Worst-fit**: Allocate the *largest* hole; must also search entire list. Produces the largest leftover hole.

First-fit is faster.

Best-fit is better in terms of storage utilization.

Worst-fit may lead less fragmentation.

30


## Example

Given five memory partitions of 100 KB, 500 KB, 200 KB, 300 KB, and 600 KB (in order), how would each of the first-fit, best-fit, and worst-fit algorithms place processes of 212 KB, 417 KB, 112 KB, and 426 KB (in order)? Which algorithm makes the most efficient use of memory?

31

## Summary

- **Deadlocks**
  - Resource Allocation Graphs
  - Deadlock Detection
  - Deadlock Prevention
  - **Deadlock Avoidance**
  - **Deadlock Recovery**
- **Main Memory Management**



32

## Acknowledgements

- “Operating Systems Concepts” book and supplementary material by A. Silberschatz, P. Galvin and G. Gagne
- “Operating Systems: Internals and Design Principles” book and supplementary material by W. Stallings
- “Modern Operating Systems” book and supplementary material by A. Tanenbaum
- R. Doursat and M. Yuksel from UNR

33