

8. Deadlocks

Sungyoung Lee

*College of Engineering
KyungHee University*

Contents

- n *System Model*
- n *Deadlock Characterization*
- n *Methods for Handling Deadlocks*
- n *Deadlock Prevention*
- n *Deadlock Avoidance*
- n *Deadlock Detection*
- n *Recovery from Deadlock*
- n *Combined Approach to Deadlock Handling*

The Deadlock Problem

n A set of blocked processes each holding a resource and waiting to acquire a resource held by another process in the set

n Example

ü System has 2 tape drives

ü P_1 and P_2 each hold one tape drive and each needs another one

n Example

ü semaphores A and B , initialized to 1

P_0	P_1
<i>wait (A);</i>	<i>wait(B)</i>
<i>wait (B);</i>	<i>wait(A)</i>

n Resources

- ü Can only be used by one process at a time
- ü Can be both hardware and software
 - § e.g., tape drives, printers, system tables, database entries, memory locations
- ü Request → Allocate → Use → Release


n Preemptable resources

- ü Sharable resources
- ü Can be taken away from a process with no ill effects

n Nonpreemptable resources

- ü Nonsharable resources
- ü Will cause the process to fail if taken away
- ü Generally deadlocks involve nonpreemptable resources

Bridge Crossing Example



- n Traffic only in one direction
- n Each section of a bridge can be viewed as a resource
- n If a deadlock occurs, it can be resolved if one car backs up (preempt resources and rollback)
- n Several cars may have to be backed up if a deadlock occurs
- n Starvation is possible

System Model

- n Resource types R_1, R_2, \dots, R_m
CPU cycles, memory space, I/O devices

- n Each resource type R_i has W_i instances

- n Each process utilizes a resource as follows:
 - ü request
 - ü use
 - ü release

Deadlock Characterization

Deadlock can arise if four conditions hold simultaneously

- n Mutual exclusion:** only one process at a time can use a resource
- n Hold and wait:** a process holding at least one resource is waiting to acquire additional resources held by other processes
- n No preemption:** a resource can be released only voluntarily by the process holding it, after that process has completed its task
- n Circular wait:** there exists a set $\{P_0, P_1, \dots, P_n, P_0\}$ of waiting processes such that P_0 is waiting for a resource that is held by P_1 , P_1 is waiting for a resource that is held by P_2 , ..., P_{n-1} is waiting for a resource that is held by P_n , and P_n is waiting for a resource that is held by P_0

Resource-Allocation Graph

A set of vertices V and a set of edges E

n V is partitioned into two types:

ü $P = \{P_1, P_2, \dots, P_n\}$, the set consisting of all the processes in the system

ü $R = \{R_1, R_2, \dots, R_m\}$, the set consisting of all resource types in the system

n Request edge


ü directed edge $P_i \rightarrow R_j$

n Assignment edge


ü directed edge $R_j \rightarrow P_i$

Resource-Allocation Graph (Cont'd)


n Process




n Resource Type with 4 instances



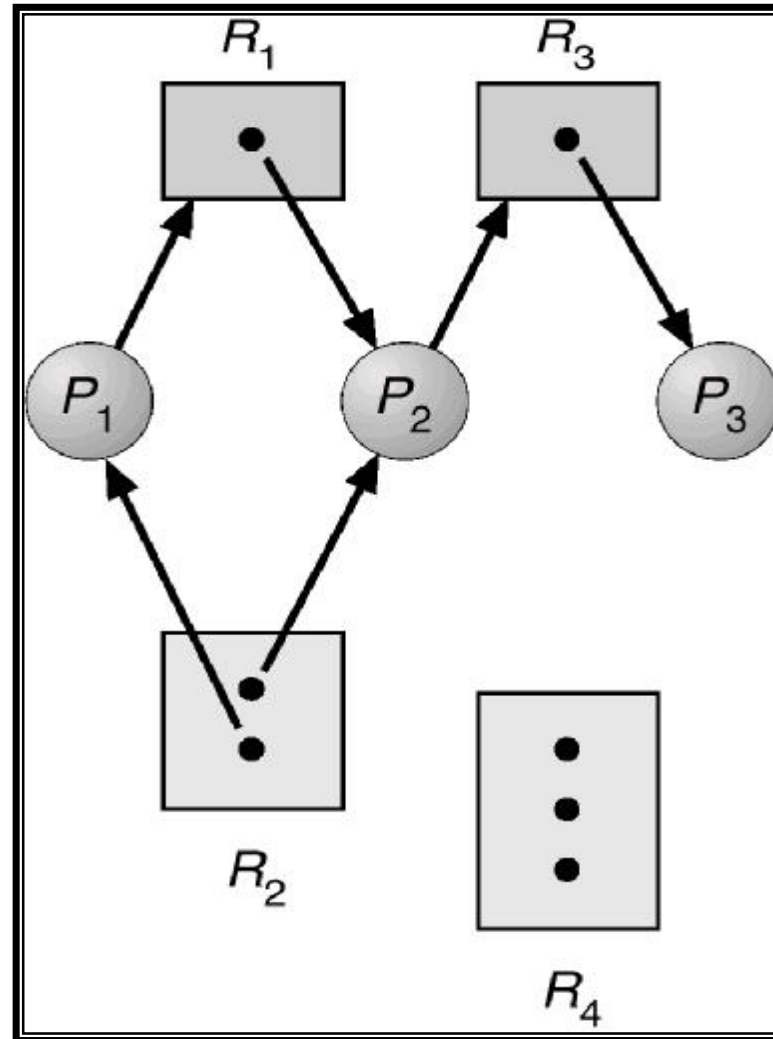
n P_i requests instance of R_j



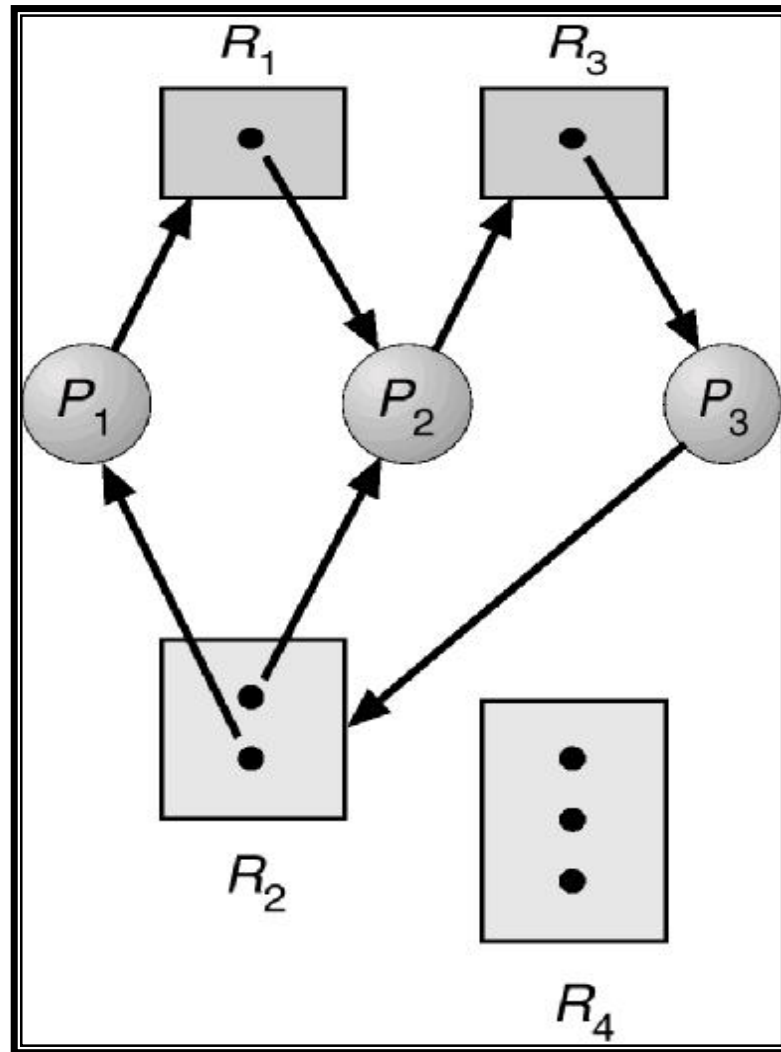
n P_i is holding an instance of R_j



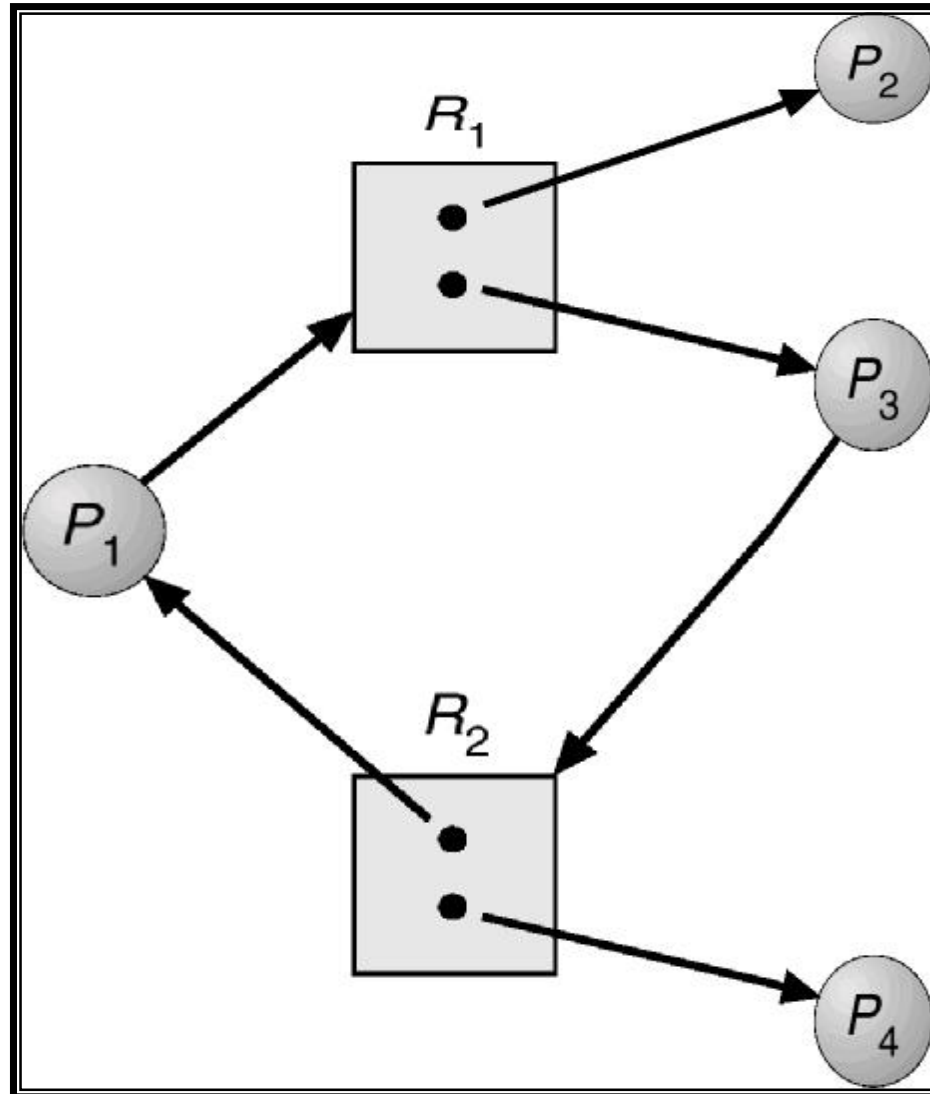
Example of a Resource Allocation Graph



Resource Allocation Graph With A Deadlock



Resource Allocation Graph With A Cycle But No Deadlock



Basic Facts

- n If graph contains no cycles \Rightarrow no deadlock
- n If graph contains a cycle \Rightarrow
 - ü If only one instance per resource type, then deadlock
 - ü If several instances per resource type, possibility of deadlock

Methods for Handling Deadlocks

- n Ensure that the system will *never* enter a deadlock state
- n Allow the system to enter a deadlock state and then recover
- n Ignore the problem and pretend that deadlocks never occur in the system
 - ü used by most operating systems, including UNIX

Handling Deadlocks

n Deadlock prevention

- ü Restrain how requests are made
- ü Ensure that at least one necessary condition cannot hold

n Deadlock avoidance

- ü Require additional information about how resources are to be requested
- ü Decide to approve or disapprove requests on the fly

n Deadlock detection and recovery

- ü Allow the system to enter a deadlock state and then recover

n Just ignore the problem altogether!

Deadlock Prevention

Restrain the ways request can be made

n Mutual Exclusion

- ü Not required for sharable resources
- ü Must hold for nonsharable resources
- ü Make as few processes as possible actually claim the resource

n Hold and Wait

- ü Must guarantee that whenever a process requests a resource, it does not hold any other resources
- ü Require process to request and be allocated all its resources before it begins execution, or allow process to request resources only when the process has none
- ü Low resource utilization
- ü Starvation possible

Deadlock Prevention (Cont'd)

n No Preemption

- ü If a process that is holding some resources requests another resource that cannot be immediately allocated to it, then all resources currently being held are released
- ü Preempted resources are added to the list of resources for which the process is waiting
- ü Process will be restarted only when it can regain its old resources, as well as the new ones that it is requesting

n Circular Wait

- ü Impose a total ordering of all resource types, and require that each process requests resources in an increasing order of enumeration
- ü $F: R \rightarrow N$, where $R = \{R_1, R_2, \dots, R_n\}$ is the set of resource types and N is the set of natural numbers
- ü Whenever a process requests an instance of R_i , it has released any resources R_j such that $F(R_i) >= F(R_j)$

Deadlock Avoidance

Requires that the system has some additional *a priori* information available

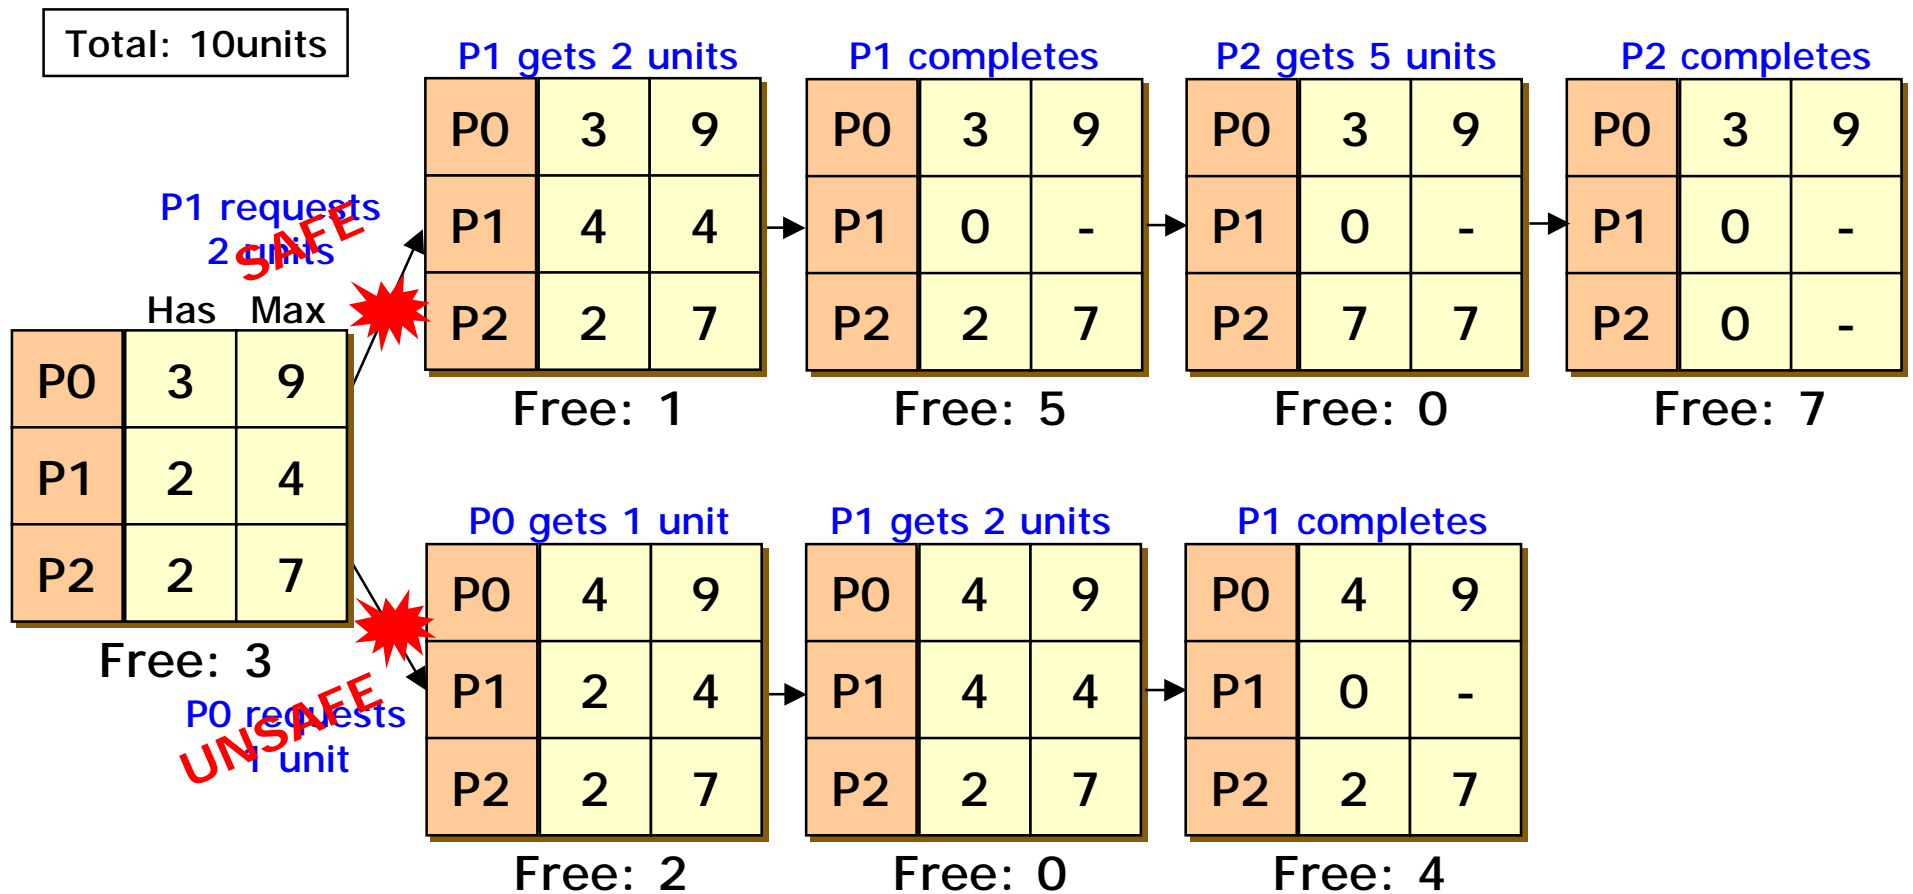
- n Simplest and most useful model requires that each process declare the *maximum number* of resources of each type that it may need
- n The deadlock-avoidance algorithm dynamically examines the resource-allocation state to ensure that there can never be a circular-wait condition
- n Resource-allocation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

Safe State

- n When a process requests an available resource, system must decide if immediate allocation leaves the system in a *safe state*
- n System is in safe state if there exists a safe sequence of all processes
- n 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 $j < i$
 - ü If P_i resource needs are not immediately available, then P_i can wait until all P_j have finished
 - ü When P_j is finished, P_i can obtain needed resources, execute, return allocated resources, and terminate
 - ü When P_i terminates, P_{i+1} can obtain its needed resources, and so on

Safe vs. Unsafe State


n Examples



Basic Facts

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


Safe, Unsafe , Deadlock State




Resource-Allocation Graph Algorithm

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

Resource-Allocation Graph For Deadlock Avoidance



Unsafe State In Resource-Allocation Graph



Banker's Algorithm

- n Multiple instances
- n Each process must a priori claim maximum use
- n When a process requests a resource it may have to wait
- n When a process gets all its resources it must return them in a finite amount of time

Data Structures for the Banker's Algorithm

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

- n** *Available*: Vector of length m . If available $[j] = k$, there are k instances of resource type R_j available
- n** *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
- n** *Allocation*: $n \times m$ matrix. If $Allocation[i,j] = k$ then P_i is currently allocated k instances of R_j .
- n** *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

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

Safety Algorithm

1. Let *Work* and *Finish* be vectors of length m and n , respectively. Initialize:
 $Work := Available$
 $Finish[i] := false$ 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

Resource-Request Algorithm for Process P_i

n $Request_i$: request vector for process P_i

n 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 P the resources are allocated to P_i
- If unsafe P P_i must wait, and the old resource-allocation state is restored

Example of Banker's Algorithm

n 5 processes P_0 through P_4

ü 3 resource types A (10 instances), B (5 instances), and C (7 instances)

n 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	

Example (Cont'd)

n 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

n 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

Example (Redrawn)

Total Resources			Allocated Resources			Max Resources			Needed Resources					
A	B	C	P0	A	B	C	P0	A	B	C	P0	A	B	C
10	5	7	P0	0	1	0	P0	7	5	3	P0	7	4	3
			P1	2	0	0	P1	3	2	2	P1	1	2	2
			P2	3	0	2	P2	9	0	2	P2	6	0	0
			P3	2	1	1	P3	2	2	2	P3	0	1	1
			P4	0	0	2	P4	4	3	3	P4	4	3	1

Available Resources		
A	B	C
3	3	2

Currently safe: $\langle P1, P3, P4, P0, P2 \rangle$ is a safe sequence

Example P_1 Request (1,0,2) (Cont'd)

n 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	

n Executing safety algorithm shows that sequence $\langle P_1, P_3, P_4, P_0, P_2 \rangle$ satisfies safety requirement

n Can request for (3,3,0) by P_4 be granted?

n Can request for (0,2,0) by P_0 be granted?

Banker's Algorithm

n Notes

- ü Safety checking algorithm requires $O(mn^2)$ operations, where
 - § m is the number of resource types
 - § n is the number of processes
- ü Processes rarely know in advance what their maximum resource needs will be
- ü The number of processes is not fixed, but dynamically varying as new users log in and out
- ü Resources that were thought to be available can suddenly vanish
 - § e.g., tape drives or disk drives
- ü In practice, few, if any, existing systems use the banker's algorithm for avoiding deadlocks

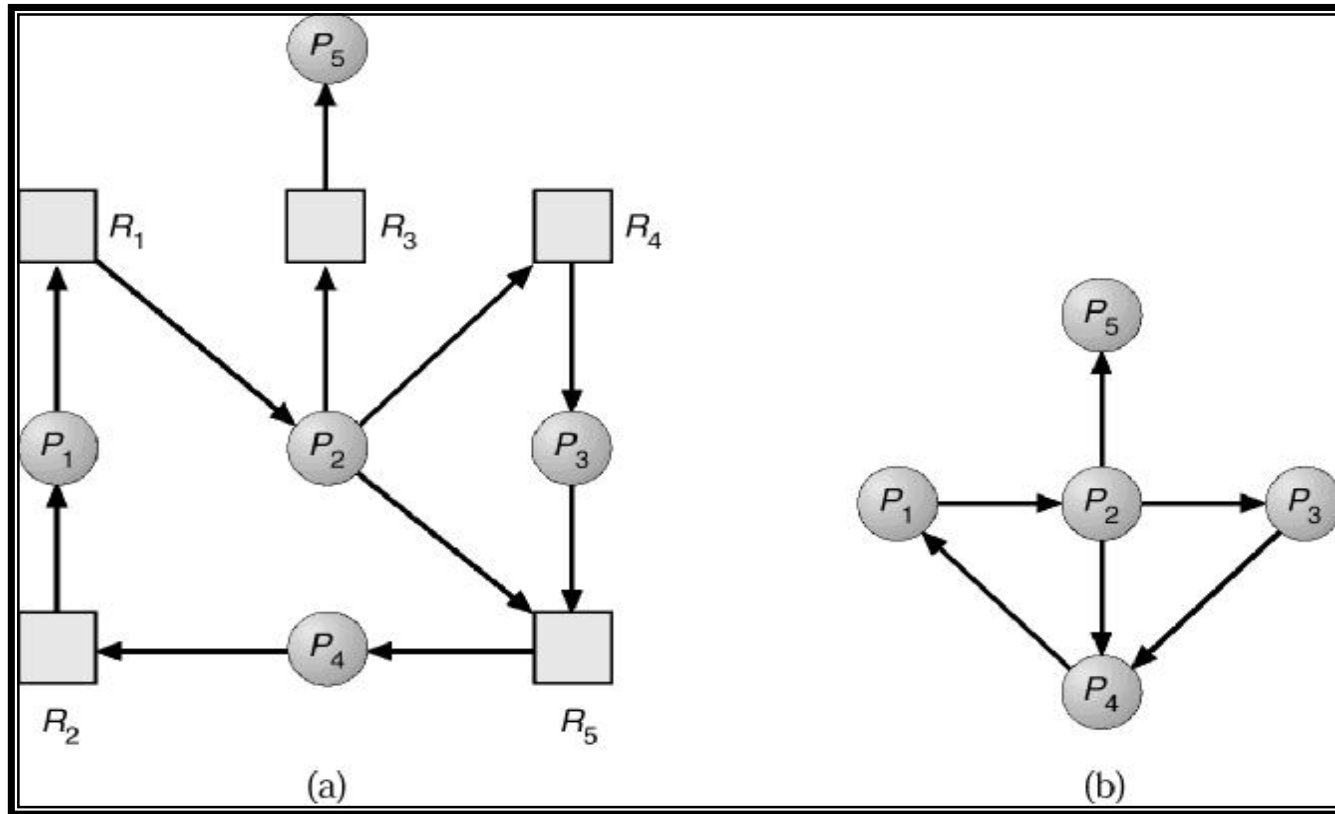
Deadlock Detection

- n Allow system to enter deadlock state
- n Detection algorithm
- n Recovery scheme

Single Instance of Each Resource Type

- n Maintain *wait-for* graph
 - ü Nodes are processes
 - ü $P_i \rightarrow P_j$ if P_i is waiting for P_j
- n Periodically invoke an algorithm that searches for a cycle in the graph
- n An algorithm to detect a cycle in a graph requires an order of n^2 operations, where n is the number of vertices in the graph

Resource-Allocation Graph and Wait-for Graph



Resource-Allocation Graph

Corresponding wait-for graph

Several Instances of a Resource Type

- n** *Available*: A vector of length m indicates the number of available resources of each type
- n** *Allocation*: An $n \times m$ matrix defines the number of resources of each type currently allocated to each process
- n** *Request*: An $n \times m$ matrix indicates the current request of each process. If $Request[i,j] = k$, then process P_i is requesting k more instances of resource type, R_j

Detection Algorithm

1. Let *Work* and *Finish* be vectors of length *m* and *n*, respectively initialize:
 - (a) *Work* := *Available*
 - (b) For $i := 1, 2, \dots, n$, if $Allocation_i \neq 0$, then
Finish[*i*] = false; otherwise, *Finish*[*i*] = true

2. Find an index *i* such that both:
 - (a) *Finish*[*i*] = false
 - (b) $Request_i \leq Work$

If no such *i* exists, go to step 4

Detection Algorithm (Cont'd)

3. $Work := Work + Allocation;$
 $Finish[i] = true$
go to step 2
4. If $Finish[i] = false$, for some $i, 1 \leq i \leq n$, then the system is in deadlock state.
Moreover, if $Finish[i] = false$, then P_i is deadlocked

Algorithm requires an order of $O(m \times n^2)$ operations to detect whether the system is in deadlocked state

Example of Detection Algorithm

- n Five processes P_0 through P_4
- n Three resource types A (7 instances), B (2 instances), and C (6 instances)
- n Snapshot at time T_0 :

	<u>Allocation</u>	<u>Request</u>	<u>Available</u>
	A B C	A B C	A B C
P_0	0 1 0	0 0 0	0 0 0
P_1	2 0 0	2 0 2	
P_2	3 0 3	0 0 0	
P_3	2 1 1	1 0 0	
P_4	0 0 2	0 0 2	

- n Sequence $\langle P_0, P_2, P_3, P_1, P_4 \rangle$ will result in $Finish[i] = \text{true}$ for all i

Example (Redrawn)

Total Resources

A	B	C
7	2	6

Available Resources

A	B	C
0	0	0

Allocated Resources

	A	B	C
P0	0	1	0
P1	2	0	0
P2	3	0	3
P3	2	1	1
P4	0	0	2

Requested Resources

	A	B	C
P0	0	0	0
P1	2	0	2
P2	0	0	0
P3	1	0	0
P4	0	0	2

What if this becomes 1?

Currently not in a deadlock: $\langle P0, P2, P3, P1, P4 \rangle$ will work

Example (Cont'd)

- n P_2 requests an additional instance of type C

	<u>Request</u>		
	A	B	C
P_0	0	0	0
P_1	2	0	2
P_2	0	0	1
P_3	1	0	0
P_4	0	0	2

- n State of system?

- ü Can reclaim resources held by process P_0 , but insufficient resources to fulfill other processes' requests
- ü Deadlock exists, consisting of processes P_1 , P_2 , P_3 , and P_4

Detection-Algorithm Usage

- n When, and how often, to invoke depends on:
 - ü How often a deadlock is likely to occur?
 - ü How many processes will need to be rolled back?
 - § one for each disjoint cycle

- n If detection algorithm is invoked arbitrarily, there may be many cycles in the resource graph and so we would not be able to tell which of the many deadlocked processes “caused” the deadlock

Recovery from Deadlock: Process Termination

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

Recovery from Deadlock: Resource Preemption

n Selecting a victim

- ü Minimize cost

n Rollback

- ü Return to some safe state, restart process for that state

n Starvation

- ü Same process may always be picked as victim, include number of rollback in cost factor

Combined Approach to Deadlock Handling

- n Combine the three basic approaches

 - ü Prevention

 - ü Avoidance


 - ü Detection

allowing the use of the optimal approach for each of resources in the system

- n Partition resources into hierarchically ordered classes

- n Use most appropriate technique for handling deadlocks within each class

Traffic Deadlock for Exercise 8.4



n The Ostrich algorithm

- ü Just put your head in the sand and pretend there is no problem at all
- ü Reasonable if
 - § Deadlocks occur very rarely
 - § Cost of prevention is high
- ü UNIX and Windows take this approach
- ü It is a trade-off between
 - § Convenience
 - § Correctness