

MODERN OPERATING SYSTEMS

Third Edition
ANDREW S. TANENBAUM

Chapter 6 Deadlocks

Preemptable and Nonpreemptable Resources

- Non-sharable resource: the resource can be used by only one process at a time
- A process may use a resource in only the following sequence:
 1. Request: If the resource cannot be granted, the requesting process must wait until it can acquire the resource.
 2. Use: The process can operate on the resource.
 3. Release: The process releases the resource.

Introduction To Deadlocks

Deadlock can be defined formally as follows:

A set of processes is deadlocked if each process in the set is waiting for an event that only another process in the set can cause.

Conditions for Resource Deadlocks


Necessary conditions for a deadlock to occur

1. Mutual exclusion: The resource is non-sharable.
2. Hold and wait: A process that is holding resources can request new resources.
3. No preemption: A resource can be released only by the process holding it.
4. Circular wait: There is a circular chain of two or more processes, each of which is waiting for a resource held by the next member of the chain.
5. All four conditions must simultaneously hold.


Deadlock Modeling (1)

Resource graph: a directed graph with two types of nodes:


Processes (circles) and resources (squares)



(a)



(b)



(c)

Figure 6-3. Resource allocation graphs. (a) Holding a resource. (b) Requesting a resource. (c) Deadlock.

Deadlock Modeling

Use resource graph to detect deadlocks

An example:

- _ Three processes A, B, and C
- _ Three resources R, S and T
- _ Round robin scheduling

Using resource graph, we can see if a given request/release sequence leads to deadlock:

Carry out the request and release step by step, check if there is any circle after each step.

Deadlock Modeling (2)

A
Request R
Request S
Release R
Release S

(a)

B
Request S
Request T
Release S
Release T


(b)

C
Request T
Request R
Release T
Release R


(c)

1. A requests R
2. B requests S
3. C requests T
4. A requests S
5. B requests T
6. C requests R
deadlock


(d)



(e)



(f)




(g)

Figure 6-4. An example of how deadlock occurs and how it can be avoided.


Deadlock Modeling (3)

1. A requests R
2. B requests S
3. C requests T
4. A requests S
5. B requests T
6. C requests R
deadlock


(d)




(e)




(f)




(g)



(h)



(i)




(j)


Figure 6-4. An example of how deadlock occurs and how it can be avoided.

Deadlock Modeling (4)


1. A requests R
2. C requests T
3. A requests S
4. C requests R
5. A releases R
6. A releases S
no deadlock




(k)




(l)




(m)




(n)



(o)



(p)



(q)

Figure 6-4. An example of how deadlock occurs and how it can be avoided.

Deadlock Modeling (5)

Strategies for dealing with deadlocks:

1. Just ignore the problem.
2. Detection and recovery. Let deadlocks occur, detect them, take action.
3. Dynamic avoidance by careful resource allocation.
4. Prevention, by structurally negating one of the four required conditions.

Ignoring Deadlocks

The Ostrich algorithm:

- Stick your head in the sand and pretend that deadlocks never occur.
- Used by most operating systems, including UNIX.
- Tradeoff between convenience and correctness

An Example in Unix

An example of deadlock in UNIX:

- Process table has 100 slots
- 10 processes are running
- Each process needs to fork 12 subprocesses
- After each forks 9 subprocesses, the table is full
- Each original process sits in the endless loop: fork and fail

Deadlock Detection

- In a system where a deadlock may occur, the system must provide:
- An algorithm than exams the state of the system to determine whether a deadlock has occurred
- An algorithm to recover from the deadlock

Detection

- Every time a resource is requested or released, check resource graph to see if any cycles exist.
- How to detect cycles in a directed graph?
- Depth-first search from each node. See if any repeated node. $O(N)$ algorithm.

Deadlock Detection with One Resource of Each Type (1)

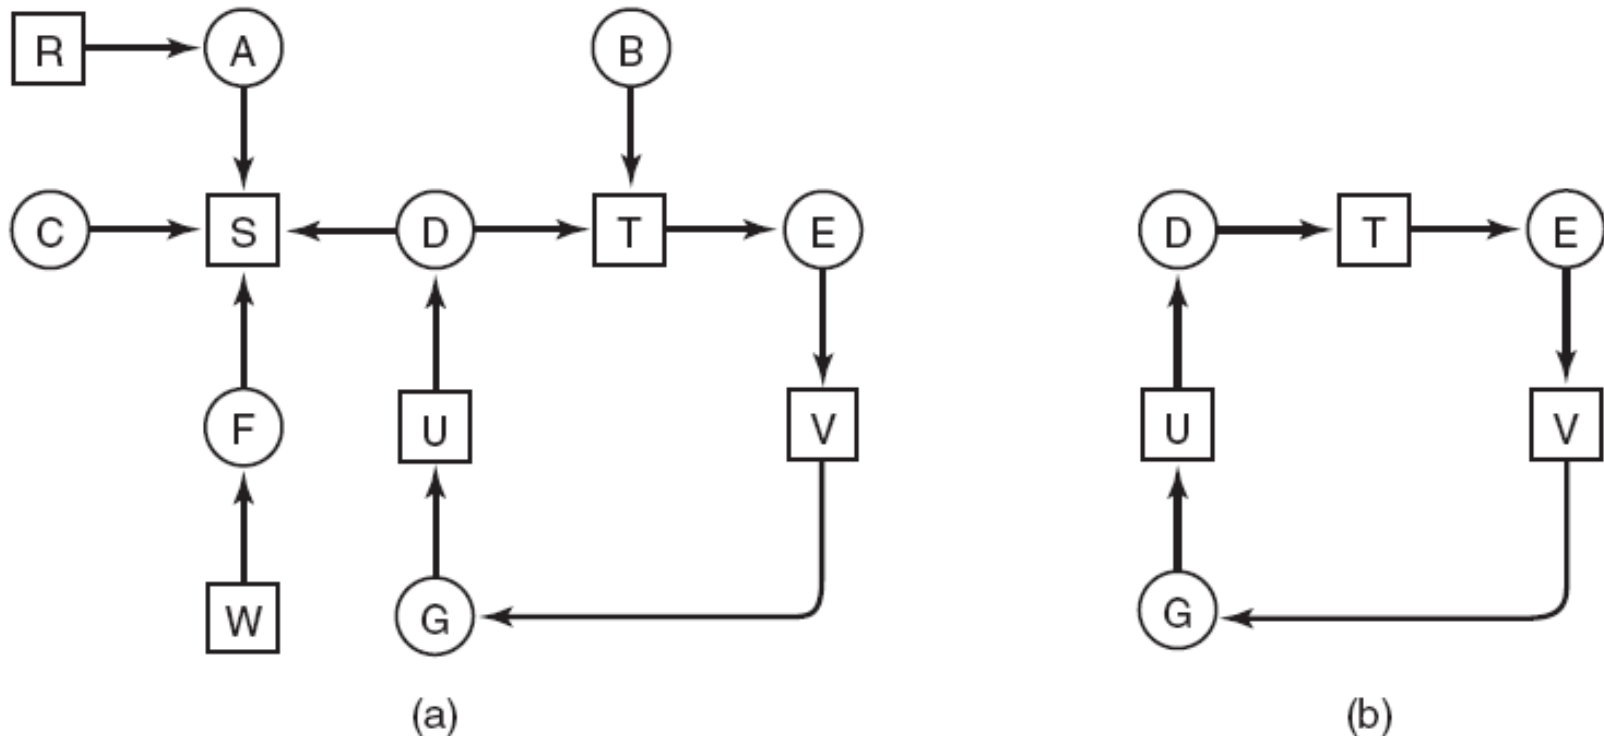


Figure 6-5. (a) A resource graph. (b) A cycle extracted from (a).

Deadlock Detection with One Resource of Each Type (2)

Algorithm for detecting deadlock:

1. For each node, N in the graph, perform the following five steps with N as the starting node.
2. Initialize L to the empty list, designate all arcs as unmarked.
3. Add current node to end of L, check to see if node now appears in L two times. If it does, graph contains a cycle (listed in L), algorithm terminates.

...

Deadlock Detection with One Resource of Each Type (3)

4. From given node, see if any unmarked outgoing arcs. If so, go to step 5; if not, go to step 6.
5. Pick an unmarked outgoing arc at random and mark it. Then follow it to the new current node and go to step 3.
6. If this is initial node, graph does not contain any cycles, algorithm terminates. Otherwise, dead end. Remove it, go back to previous node, make that one current node, go to step 3.

Deadlock Detection with Multiple Resources of Each Type (1)




Figure 6-6. The four data structures needed by the deadlock detection algorithm.

Deadlock Detection with Multiple Resources of Each Type (2)

Deadlock detection algorithm:

1. Look for an unmarked process, P_i , for which the i -th row of R is less than or equal to A .
2. If such a process is found, add the i -th row of C to A , mark the process, and go back to step 1.
3. If no such process exists, the algorithm terminates.

Deadlock Detection with Multiple Resources of Each Type (3)

$$E = \begin{pmatrix} 4 & 2 & 3 & 1 \end{pmatrix}$$

Tape drives
Plotters
Scanners
CD Roms

$$A = \begin{pmatrix} 2 & 1 & 0 & 0 \end{pmatrix}$$

Tape drives
Plotters
Scanners
CD Roms

Current allocation matrix

$$C = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 2 & 0 & 0 & 1 \\ 0 & 1 & 2 & 0 \end{bmatrix}$$

Request matrix

$$R = \begin{bmatrix} 2 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 2 & 1 & 0 & 0 \end{bmatrix}$$

Figure 6-7. An example for the deadlock detection algorithm.

Recovery from Deadlock

Recovery

- Abort one process at a time until the deadlock cycle is eliminated.
- A simpler way (used in large main frame computers):
Do not maintain a resource graph. Only periodically check to see if there are any processes that have been blocked for a certain amount of time, say, 1 hour. Then kill such processes.
- To recover the killed processes, need to restore any modified files. Keep different versions of the file.

Recovery from Deadlock

- Recovery through preemption (take the resource back)
- Recovery through rollback (to the nearest checkpoint)
- Recovery through killing processes

Deadlock Avoidance

Analyzing each resource request to see if it can be safely granted.

Resource trajectories: A model for two processes and two resources

- An example:
- Processes A and B
- Resources: printer and plotter
- A needs printer from I1 to I3
- A needs plotter from I2 to I4
- B needs plotter from I5 to I7
- B needs printer from I6 to I8
- Each point in the diagram is a joint state of A & B
- Can only go vertically or horizontally (one CPU)
- Start at point p, run A to point q, run B to point r, run A to point s, grant printer, run B to point t, request plotter, can only run A to completion.

Deadlock Avoidance




Figure 6-8. Two process resource trajectories.

Deadlock Avoidance

Find a general algorithm that can always avoid deadlock by making right decisions.

Banker's algorithm for a single resource:

- A small town banker deals with a group of customers with granted credit lines.
- The analogy:
 - Customers: processes
 - Units: copies of the resource
 - Banker: O.S.
 - State of the system: showing the money loaned and the maximum credit available

Deadlock Avoidance

Safe state:

- There exists a sequence of other states that lead to all customers getting loans up to their credit lines.

The algorithm:

- For each request, see if granting it leads to a safe state. If it does, the request is granted. Otherwise, it is postponed until later.
- Check a safe state:
 - (1) See if available resources can satisfy the customer closest to his maximum. If so, these loans are assumed to be repaid.
 - (2) Then check the customer now closest to his maximum, and so on.
 - (3) If all loans can be eventually paid, the current state is safe.

Safe and Unsafe States (1)




Figure 6-9. Demonstration that the state in (a) is safe.

Safe and Unsafe States (2)




Figure 6-10. Demonstration that the state in (b) is not safe.

The Banker's Algorithm for a Single Resource

	Has	Max
A	0	6
B	0	5
C	0	4
D	0	7

Free: 10
(a)

	Has	Max
A	1	6
B	1	5
C	2	4
D	4	7

Free: 2
(b)

	Has	Max
A	1	6
B	2	5
C	2	4
D	4	7

Free: 1
(c)

Figure 6-11. Three resource allocation states:
(a) Safe. (b) Safe. (c) Unsafe.

The Banker's Algorithm for Multiple Resources

- Processes must state their total resource needs before executing
- n processes and m types of resources
- Two matrices:
 - Current allocation matrix
 - Request matrix
- Three vectors:
 - Existing resource: $E = (E_1, E_2, \dots, E_m)$
 - Possessed resource: $P = (P_1, P_2, \dots, P_m)$
 - Available resource: $A = (A_1, A_2, \dots, A_m)$

$$A = E - P$$

The Banker's Algorithm for Multiple Resources

	Process	Tape drives	Plotters	Printers	CD ROMs
A	3	0	1	1	
B	0	1	0	0	
C	1	1	1	0	
D	1	1	0	1	
E	0	0	0	0	

Resources assigned

	Process	Tape drives	Plotters	Printers	CD ROMs
A	1	1	0	0	
B	0	1	1	2	
C	3	1	0	0	
D	0	0	1	0	
E	2	1	1	0	

Resources still needed

E = (6342)
P = (5322)
A = (1020)

Figure 6-12. The banker's algorithm with multiple resources.

The Banker's Algorithm for Multiple Resources

Algorithm for checking to see if a state is safe:

1. Look for row, R , whose unmet resource needs all $\leq A$. If no such row exists, system will eventually deadlock since no process can run to completion
2. Assume process of row chosen requests all resources it needs and finishes. Mark process as terminated, add all its resources to the A vector.
3. Repeat steps 1 and 2 until either all processes marked terminated (initial state was safe) or no process left whose resource needs can be met (there is a deadlock).

The Banker's Algorithm for Multiple Resources

An example:

- Row D \leq A, then $A = A + (1101) = (2121)$
- Row A \leq A, then $A = A + (3011) = (5132)$
- Row B \leq A, then $A = A + (0100) = (5232)$
- Row C \leq A, then $A = A + (1110) = (6342)$
- Row E \leq A, then $A = A + (0000) = (6342) = E$
- So, the current state is safe.

The Banker's Algorithm for Multiple Resources

Suppose process B requests a printer

- Now $A = (1010)$
- Row D $\leq A$, then $A = A + (1101) = (2111)$
- Row A $\leq A$, then $A = A + (3011) = (5122)$
- Row B $\leq A$, then $A = A + (0110) = (5232)$
- Row C $\leq A$, then $A = A + (1110) = (6342)$
- Row E $\leq A$, then $A = A + (0000) = (6342) = E$
- So, the request is still safe.

If E requests the last printer.

- $A = (1000)$
- No row $\leq A$, will lead to a deadlock.
- So E's request should be deferred.

Deadlock Prevention

- Use a protocol to ensure that the system will never enter a deadlock state.
- Negating one of the four necessary conditions.
 1. Mutual exclusion
 2. Hold and wait
 3. No preemption
 4. Circular wait

Attacking Mutual Exclusion Condition

- Ensure that no resource is assigned exclusively to a single process. Spooling everything.
- Drawback: not all resources can be spooled (such as process table)
- Competition for disk space for spooling itself may lead to deadlock.

Attacking Hold and Wait Condition

- Process requires all its resources before starting
- Problem: processes may not know how many resources needed in advance; not an optimal approach of using resources (low utilization)
- A variant: a process requesting a resource first temporarily releases all the resources it holds. Once the request is granted, it gets all resource back.

Attacking No Preemption Condition

- Forcibly take away the resource. Not realistic.

Attacking the Circular Wait Condition

- Solution 1: A process is entitled only a single resource at any time.
- Solution 2: Global numbering all resources:
Give a unique number to each resource. All requests must be made in a numerical order


Attacking the Circular Wait Condition

- An example: Two processes and five devices. Number the resources as follows:
 - (a) Imagesetter
 - (b) Scanner
 - (c) Plotter
 - (d) Tape drive
 - (e) CD-ROM drive
- Assume process A holds i and process B holds j ($i \neq j$).
- If $i > j$, A is not allowed to request j .
- If $i < j$, B is not allowed to request i .
- Suitable to multiple processes. At any time, there must be a assigned resource with the highest number. This process will not request other assigned resources, only requests higher numbered resource and finishes. Then releases all resources.

Attacking the Circular Wait Condition

1. Imagesetter
2. Scanner
3. Plotter
4. Tape drive
5. CD-ROM drive

(a)



(b)

Figure 6-13. (a) Numerically ordered resources.
(b) A resource graph.

Approaches to Deadlock Prevention

Condition	Approach
Mutual exclusion	Spool everything
Hold and wait	Request all resources initially
No preemption	Take resources away
Circular wait	Order resources numerically

Figure 6-14. Summary of approaches to deadlock prevention.

Approaches to Deadlock Prevention

Problems:

- (1) Process don't know the maximum resources they need in advance
- (2) The number of processes is not fixed
- (3) Available resources may suddenly break

In summary,

Prevention: too overly restrictive

Avoidance: required information may not be available

Still no good general solution yet.

Communication Deadlocks

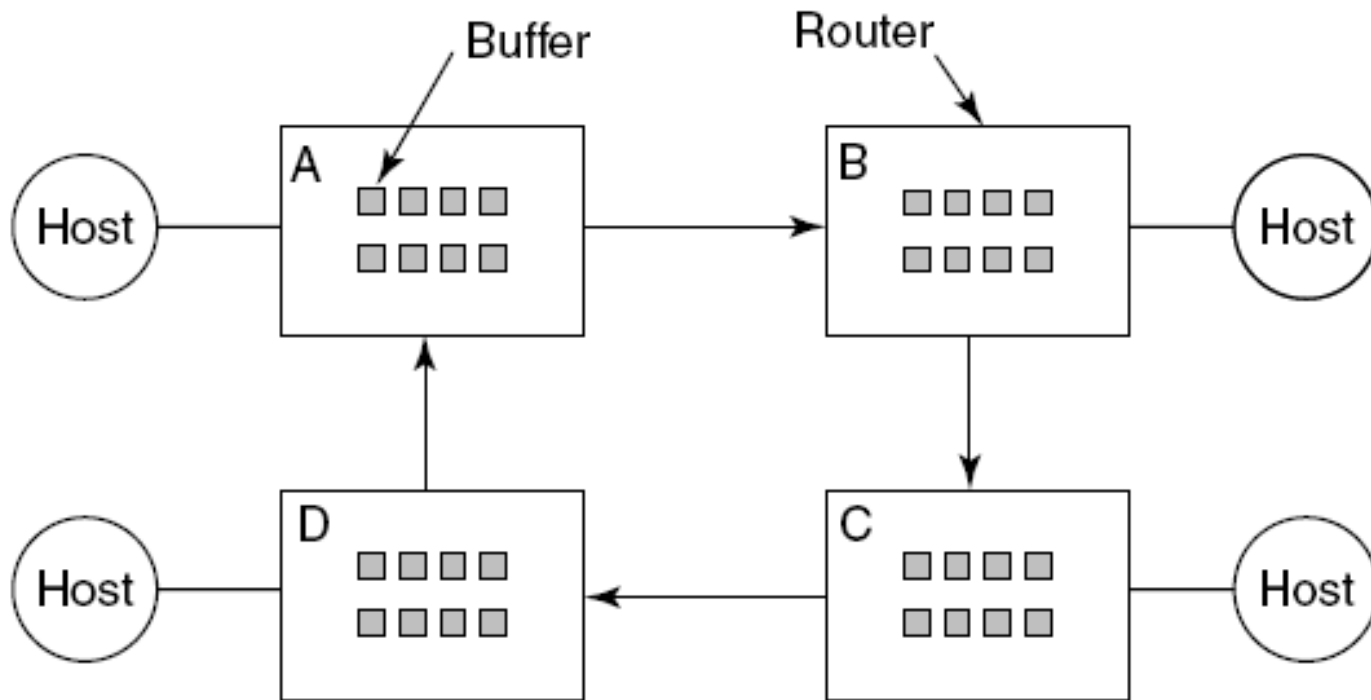


Figure 6-15. A resource deadlock in a network.