

Chapter 3

Deadlocks

- 3.1. Resource
- 3.2. Introduction to deadlocks
- 3.3. The ostrich algorithm
- 3.4. Deadlock detection and recovery
- 3.5. Deadlock avoidance
- 3.6. Deadlock prevention
- 3.7. Other issues

Resources

- Examples of computer resources
 - printers
 - tape drives
 - tables
- Processes need access to resources in reasonable order
- Suppose a process holds resource A and requests resource B
 - at same time another process holds B and requests A
 - both are blocked and remain so

Resources (1)

- Deadlocks occur when ...
 - processes are granted exclusive access to devices
 - we refer to these devices generally as resources
- Preemptable resources
 - can be taken away from a process with no ill effects
- Nonpreemptable resources
 - will cause the process to fail if taken away

Resources (2)

- Sequence of events required to use a resource
 1. request the resource
 2. use the resource
 3. release the resource
- Must wait if request is denied
 - requesting process may be blocked
 - may fail with error code

Resource Acquisition

```
typedef int semaphore;  
semaphore resource_1;
```

```
void process_A(void)  
{  
    down(&resource_1);  
    use_resource_1();  
    up(&resource_1);  
}
```

(a) Acquiring 1 resource

```
typedef int semaphore;  
semaphore resource_1;  
semaphore resource_2;
```

```
void process_A(void)  
{  
    down(&resource_1);  
    down(&resource_2);  
    use_both_resources();  
    up(&resource_2);  
    up(&resource_1);  
}
```

(b) Acquiring 2 resources

Resource Acquisition

```
typedef int semaphore;  
semaphore resource_1;  
semaphore resource_2;
```

```
void process_A(void)  
{  
    down(&resource_1);  
    down(&resource_2);  
    use_both_resources();  
    up(&resource_2);  
    up(&resource_1);  
}
```

```
void process_B(void)  
{  
    down(&resource_1);  
    down(&resource_2);  
    use_both_resources();  
    up(&resource_2);  
    up(&resource_1);  
}
```

(a) Deadlock free code

```
typedef int semaphore;  
semaphore resource_1;  
semaphore resource_2;
```

```
void process_A(void)  
{  
    down(&resource_1);  
    down(&resource_2);  
    use_both_resources();  
    up(&resource_2);  
    up(&resource_1);  
}
```

```
void process_B(void)  
{  
    down(&resource_2);  
    down(&resource_1);  
    use_both_resources();  
    up(&resource_1);  
    up(&resource_2);  
}
```

(b) Code with potential deadlock

Introduction to Deadlocks

- Formal definition :

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

- Usually the event is release of a currently held resource
- None of the processes can ...
 - run
 - release resources
 - be awakened

DEADLOCK ILLUSTRATION



Four Conditions for Deadlock

1. Mutual exclusion condition

- each resource assigned to 1 process or is available

2. Hold and wait condition

- process holding resources can request additional resources

3. No preemption condition

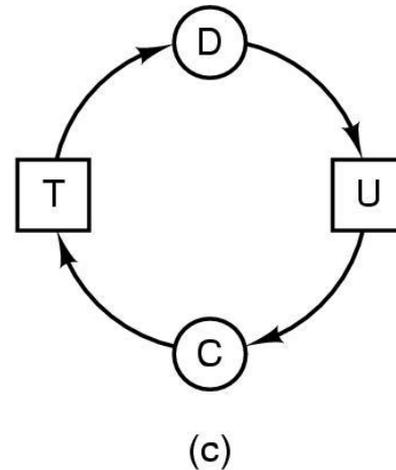
- previously granted resources cannot forcibly taken away

4. Circular wait condition

- must be a circular chain of 2 or more processes
- each is waiting for resource held by next member of the chain

Deadlock Modeling (2)

- Modeled with directed graphs



- resource R assigned to process A
- process B is requesting/waiting for resource S
- process C and D are in deadlock over resources T and U

Deadlock Modeling (3)

Strategies for dealing with Deadlocks

1. just ignore the problem altogether
2. detection and recovery
3. dynamic avoidance
 - careful resource allocation
4. prevention
 - negating one of the four necessary conditions

Deadlock Modeling (4)

A
Request R
Request S
Release R
Release S

(a)

B
Request S
Request T
Release S
Release T

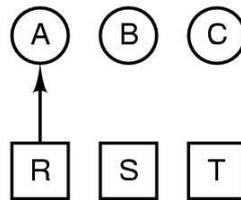
(b)

C
Request T
Request R
Release T
Release R

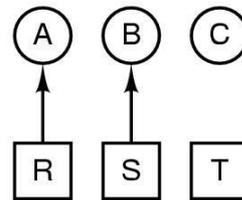
(c)

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

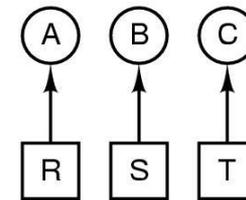
(d)



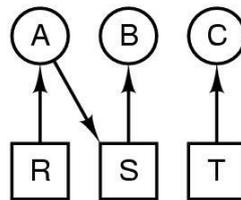
(e)



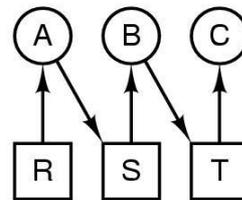
(f)



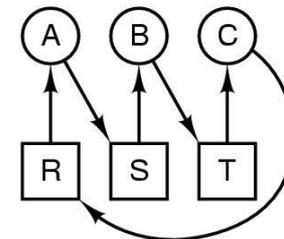
(g)



(h)



(i)



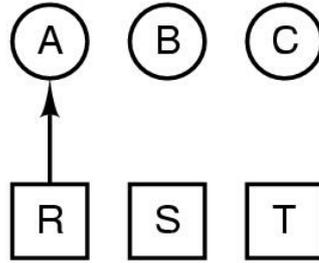
(j)

How deadlock occurs

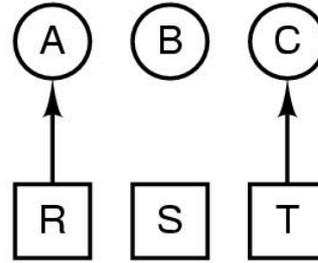
Deadlock Modeling (5)

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

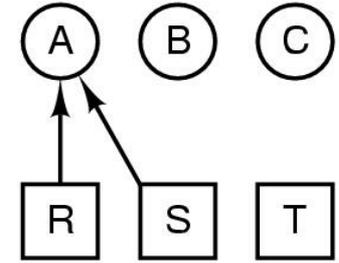
(k)



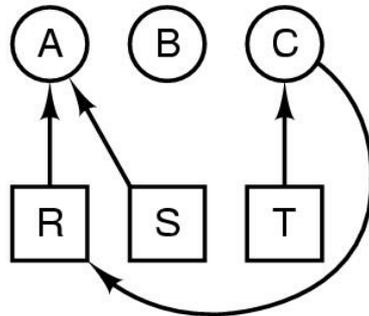
(l)



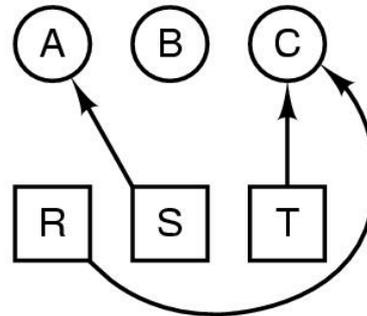
(m)



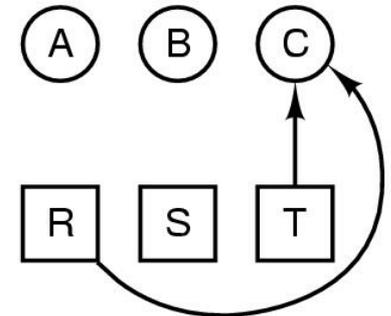
(n)



(o)



(p)



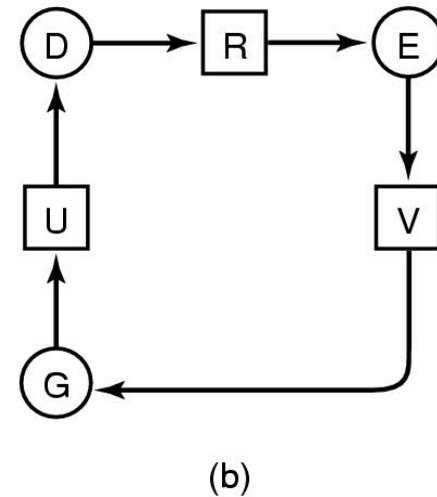
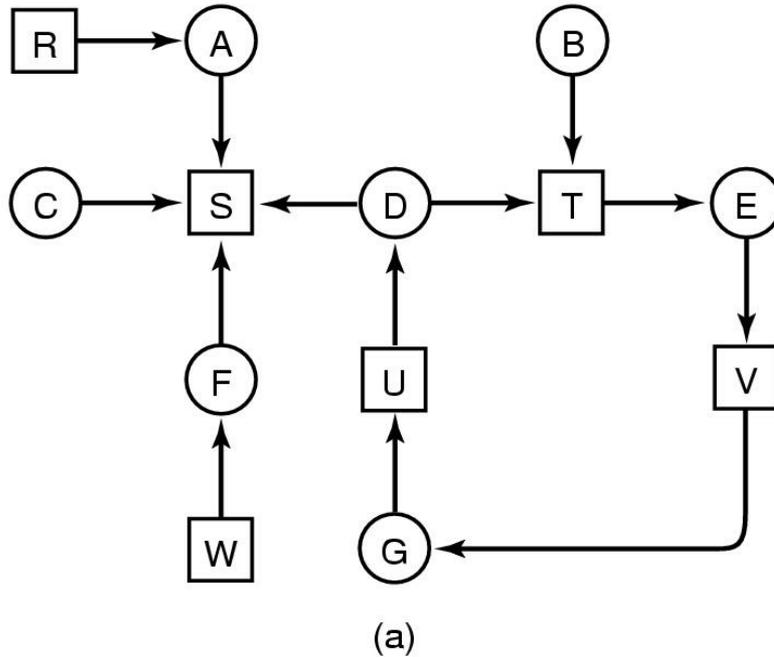
(q)

How deadlock can be avoided

The Ostrich Algorithm

- Pretend there is no problem
- Reasonable if
 - deadlocks occur very rarely
 - cost of prevention is high
- UNIX and Windows takes this approach
- It is a trade off between
 - convenience
 - correctness

Detection with One Resource of Each Type (1)



- Note the resource ownership and requests
- A cycle can be found within the graph, denoting deadlock

Deadlock Detection Algorithm for one resource of each type

1. For each node N in the graph, perform the following 5 steps with N as the starting node
2. Initialize L (i.e., list of nodes) to the empty list and designate all the arcs as unmarked
3. Add the current node to the end of L and check to see if the node now appears in L two times. If it does, the graph contains a cycle (listed in L) and the algorithm terminates
4. For the given node, see if there are 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. We have now reached a dead end. Remove it and go back to the previous node, that is, the one that was current just before this one, make that one the current node, and go to step 3. IF this node is the initial node, the graph does not contain any cycles and the algorithm terminates.

Detection with Multiple Resources of Each Type (2)

Resources in existence
($E_1, E_2, E_3, \dots, E_m$)

Resources available
($A_1, A_2, A_3, \dots, A_m$)

Current allocation matrix

Request matrix

$$\begin{bmatrix} C_{11} & C_{12} & C_{13} & \cdots & C_{1m} \\ C_{21} & C_{22} & C_{23} & \cdots & C_{2m} \\ \vdots & \vdots & \vdots & & \vdots \\ C_{n1} & C_{n2} & C_{n3} & \cdots & C_{nm} \end{bmatrix}$$

$$\begin{bmatrix} R_{11} & R_{12} & R_{13} & \cdots & R_{1m} \\ R_{21} & R_{22} & R_{23} & \cdots & R_{2m} \\ \vdots & \vdots & \vdots & & \vdots \\ R_{n1} & R_{n2} & R_{n3} & \cdots & R_{nm} \end{bmatrix}$$

Row n is current allocation
to process n

Row 2 is what process 2 needs

Data structures needed by deadlock detection algorithm

Deadlock detection with multiple Resources of each type

- Each process is initially unmarked. And the algorithm progresses, processes will be marked, indicating that they are able to complete and are thus not deadlocked
- When the algorithm terminates, any unmarked processes are known to be deadlocked.
- The deadlock detection alg is as below:
 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.

Detection with Multiple Resource 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

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

An example for the deadlock detection algorithm

Recovery from Deadlock (1)

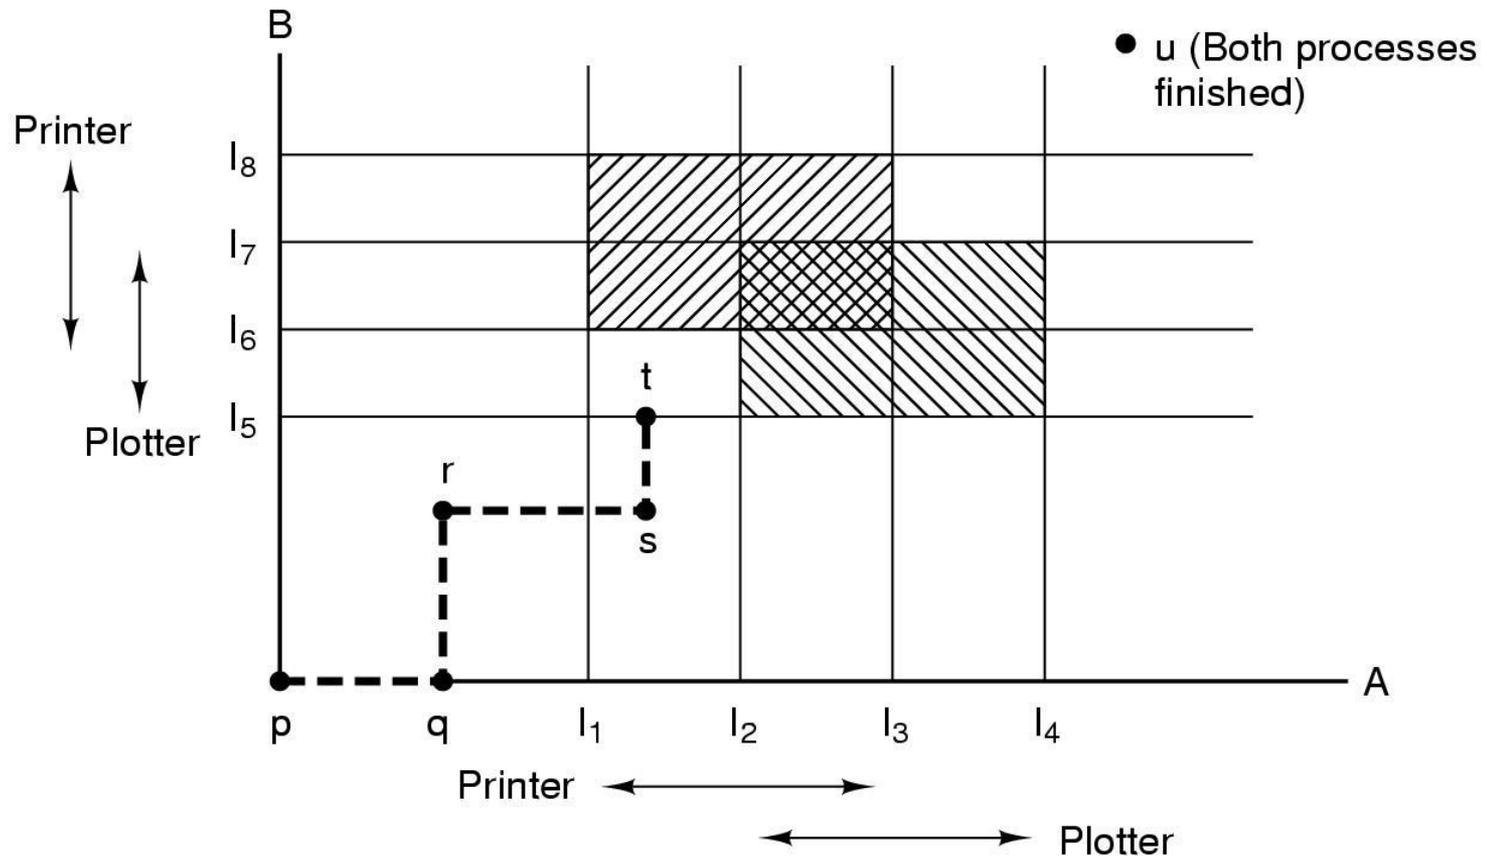
- Recovery through preemption
 - take a resource from some other process
 - depends on nature of the resource
- Recovery through rollback
 - checkpoint a process periodically
 - use this saved state
 - restart the process if it is found deadlocked

Recovery from Deadlock (2)

- Recovery through killing processes
 - crudest but simplest way to break a deadlock
 - kill one of the processes in the deadlock cycle
 - the other processes get its resources
 - choose process that can be rerun from the beginning

Deadlock Avoidance

Resource Trajectories



Two process resource trajectories

Safe and Unsafe states

- A state is called **safe** if it is not deadlocked and there is some scheduling order in which every process can run to completion even if all of them suddenly requests their maximum number of resources immediately.
- Otherwise the state of the system is said to be **unsafe**
- Safe/unsafe states are anonymous to a banker with loans
- Note that an unsafe state is not a deadlocked state, but there is a possibility that the system **MAY** go into a deadlocked state if all processes request their maximum amount of resources at the same time

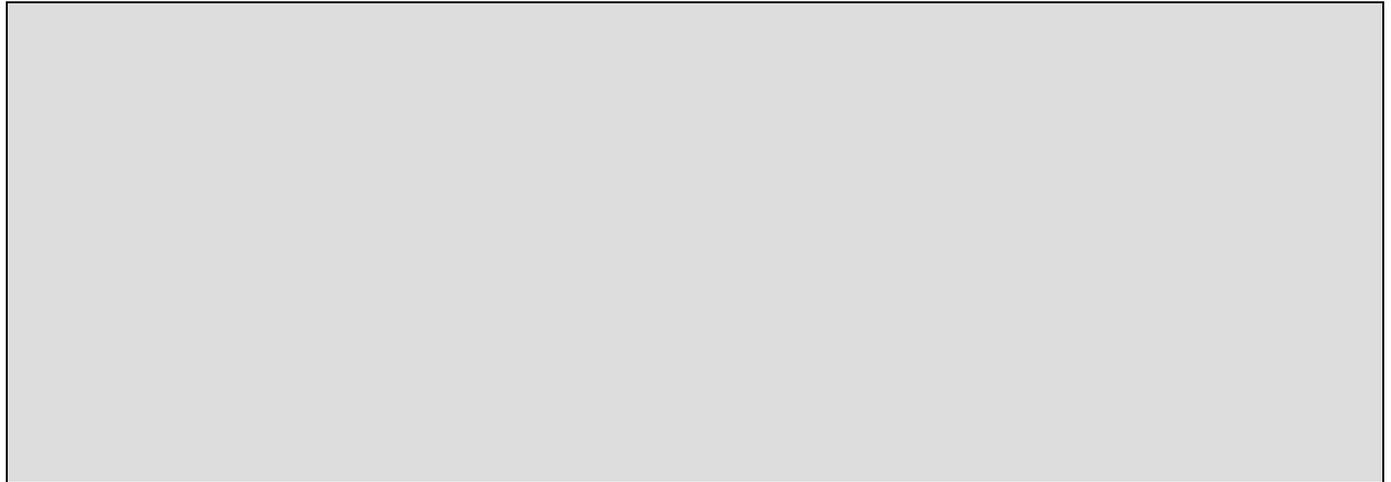
Safe and Unsafe States (1)

Demonstration with one resource where 10 instances of the same resource exists

	Has	Max
A	3	9
B	2	4
C	2	7

Free: 3

(a)



Demonstration that the state in (a) is safe cause deadlock can be avoided with careful scheduling

Safe and Unsafe States (1)

Demonstration with one resource where 10 instances of the same resource exists

	Has	Max
A	3	9
B	2	4
C	2	7

Free: 3
(a)

	Has	Max
A	3	9
B	4	4
C	2	7

Free: 1
(b)

	Has	Max
A	3	9
B	0	–
C	2	7

Free: 5
(c)

	Has	Max
A	3	9
B	0	–
C	7	7

Free: 0
(d)

	Has	Max
A	3	9
B	0	–
C	0	–

Free: 7
(e)

Demonstration that the state in (a) is safe cause deadlock can be avoided with careful scheduling

Safe and Unsafe States (2)

	Has	Max
A	3	9
B	2	4
C	2	7

Free: 3

(a)

	Has	Max
A	4	9
B	2	4
C	2	7

Free: 2

(b)

	Has	Max
A	4	9
B	4	4
C	2	7

Free: 0

(c)

	Has	Max
A	4	9
B	—	—
C	2	7

Free: 4

(d)

Demonstration that the state in b is not safe (after A requests one more resource and the resource is granted to A)

The Banker's Algorithm by Dijkstra for a Single Resource for deadlock avoidance

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

Free: 10

(a)



- Two resource allocation states
 - (a) safe
 - (b) Unsafe
- Grant a request only if it leads to a safe state, otherwise delay the request

The Banker's Algorithm by Dijkstra for a Single Resource for deadlock avoidance

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)

- Two resource allocation states
 - (a) safe
 - (b) Unsafe
- Grant a request only if it leads to a safe state, otherwise delay the request

The Banker's Algorithm by Dijkstra for a Single Resource for deadlock avoidance

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)

- Two resource allocation states
 - (a) safe
 - (b) Unsafe
- Grant a request only if it leads to a safe state, otherwise delay the request

Banker's Algorithm for Multiple Resources

	Process	Tape drives	Plotters	Scanners	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	Scanners	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)

Example of banker's algorithm with multiple resources
E (Existing), P (Possessed), A (Available)

Deadlock Prevention

Attacking the Mutual Exclusion Condition

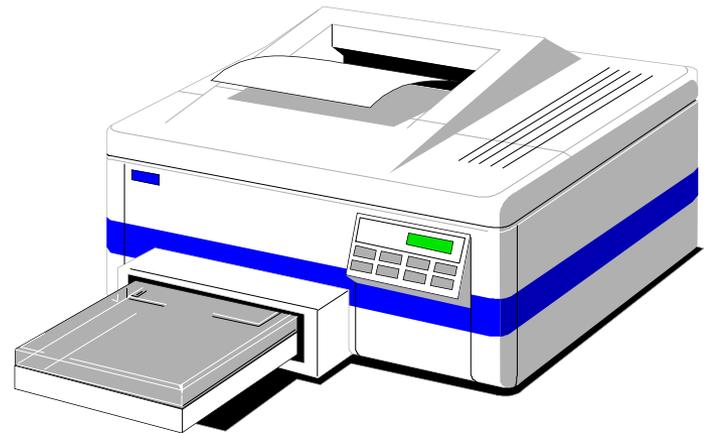
- Some devices (such as printer) can be spooled
 - only the printer daemon uses printer resource
 - thus deadlock for printer eliminated
- Not all devices can be spooled
- Principle:
 - avoid assigning resource when not absolutely necessary
 - as few processes as possible actually claim the resource

Attacking the Hold and Wait Condition

- Require processes to request resources before starting
 - a process never has to wait for what it needs
- Problems
 - may not know required resources at start of run
 - also ties up resources other processes could be using
- Variation:
 - process must give up all resources
 - then request all immediately needed

Attacking the No Preemption Condition

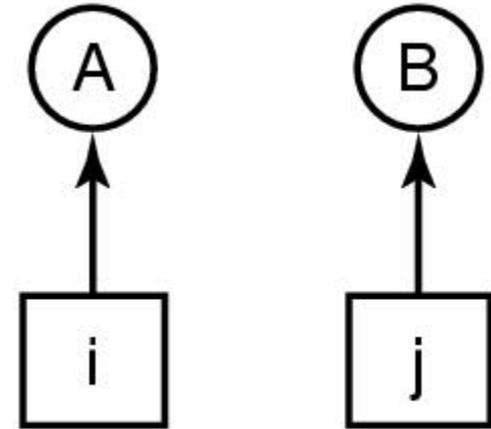
- This is not a viable option
- Consider a process given the printer
 - halfway through its job
 - now forcibly take away printer
 - !!??



Attacking the Circular Wait Condition (1)

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

(a)



(b)

- Normally ordered resources
- A resource graph

Attacking the Circular Wait Condition (1)

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

Summary of approaches to deadlock prevention

Other Issues

Two-Phase Locking

- Phase One
 - process tries to lock all records it needs, one at a time
 - if needed record found locked, start over
 - (no real work done in phase one)
- If phase one succeeds, it starts second phase,
 - performing updates
 - releasing locks
- Note similarity to requesting all resources at once
- Algorithm works where programmer can arrange
 - program can be stopped, restarted

Nonresource Deadlocks

- Possible for two processes to deadlock
 - each is waiting for the other to do some task
- Can happen with semaphores
 - each process required to do a *down()* on two semaphores (*mutex* and another)
 - if done in wrong order, deadlock results

Starvation

- Algorithm to allocate a resource
 - may be to give to shortest job first
- Works great for multiple short jobs in a system
- May cause long job to be postponed indefinitely
 - even though not blocked
- Solution:
 - First-come, first-serve policy