



# CSCI [4 | 6]730 Operating Systems

## Deadlock



Maria Hybinette, UGA


# Chapter 8: Deadlock Questions?

- What is a deadlock?
- What causes a deadlock?
- How do you deal with (potential) deadlocks?

Maria Hybinette, UGA

2

## Deadlock: What is a deadlock?



Deitel & Deitel anecdote

- All are waiting for a resource that is held by another waiting entity. Since all are waiting, none can provide any of the things being waited for.
- Example: narrow bridge (resource) --
  - » if a deadlock occurs, resolved if one car back up (preempts resource and rollback).

Maria Hybinette, UGA

3

## Example: Two Threads?

- Two threads access two shared variables, A and B
  - » Variable A is protected by lock a
  - » Variable B by lock b
- How to add lock and unlock statements?

Thread Maria

```
lock (a) ;
A += 10 ;
lock (b) ;
B += 20 ;

A += B ;
unlock (b) ;
A += 30 ;
unlock (a) ;
```

Does this work?

Thread Tucker


```
lock (b) ;
B += 10 ;
lock (a) ;
A += 20 ;

A += B ;
unlock (a) ;
B += 30 ;
unlock (b) ;
```

Maria Hybinette, UGA

4

## Example: Maria & Tucker



Thread Maria

```
lock (a) ;
A += 10 ;
lock (b) ;
B += 20 ;

A += B ;
unlock (b) ;
A += 30 ;
unlock (a) ;
```

Thread Tucker

```
lock (b) ;
B += 10 ;
lock (a) ;
A += 20 ;

A += B ;
unlock (a) ;
B += 30 ;
unlock (b) ;
```


Maria Hybinette, UGA

5


## Representing Deadlock

- Two common ways of representing deadlock:
  - » Vertices (circles or rectangles)
    - threads (or processes) in system
    - resources [types] (e.g., locks, semaphores, printers)
  - » Edges: indicates ('waiting for' or 'wants') or 'held by' direction of the edge determines the type.

Wait-For Graph



Resource Allocation Graph



Maria Hybinette, UGA

6

## Conditions for Deadlock

- **Mutual exclusion:**
  - » Resource cannot be shared
  - » Requests are delayed until resource is released
- **Hold and wait:**
  - » Thread holds one resource while waits for another
- **No preemption:**
  - » previously granted resources cannot forcibly be taken away
- **Circular wait:**
  - » Circular dependencies exist in "waits-for" or "resource-allocation" graphs
  - » Each is waiting for a resource held by next member of the chain.

*All for conditions must hold simultaneously*

Maria Hyönnelä, UGA

7

## Handling Deadlock

1. **Ignore**
  - » Easiest and most common approach (e.g., UNIX).
2. **Deadlock prevention**
  - » Ensure deadlock does not happen
  - » Ensure at least one of 4 conditions does not occur
    1. Hold&Wait, No Preemption, Circularity, Mutual Exclusion
    2. System build so deadlock cannot happen
3. **Deadlock detection and recovery**
  - » Allow deadlocks, but detect when occur
  - » Recover and continue
4. **Deadlock avoidance**
  - » Ensure deadlock does not happen
  - » Use information about resource requests to dynamically avoid **unsafe** situations



Ostrich algorithm

Maria Hyönnelä, UGA

8

## Deadlock Prevention

- **Approach**
  - » Ensure 1 of 4 conditions cannot occur
  - » Negate each of the 4 conditions
- **No single approach is appropriate (or possible) for all circumstances**

Mutual exclusion  
Hold and wait  
No preemption  
Circular wait

Maria Hyönnelä, UGA

9

## Deadlock Prevention: Mutual Exclusion

- **No mutual exclusion --> Make resource sharable ; examples:**
  - » Read-only files
  - » Printer daemon needs exclusive access to the printer, there is only one printer daemon -- uses spooling.

Mutual exclusion  
Hold and wait  
No preemption  
Circular wait

Maria Hyönnelä, UGA

10

## Deadlock Prevention Hold and Wait

- **Two Approaches:**
  1. Only request resources when it does not hold other resources
    - release resources before requesting new ones

Thread Maria

```
lock(a);
A += 10;
unlock(a);
lock(b);
B += 20;
unlock(b);
lock(a);
A += 30;
unlock(a);
```

Thread Tucker

```
lock(b);
B += 10;
unlock(b);
lock(a);
A += 20;
unlock(a);
lock(b);
B += 30;
unlock(b);
```

Mutual exclusion  
Hold and wait  
No preemption  
Circular wait

Maria Hyönnelä, UGA

11

## Deadlock Prevention Hold and Wait

- **Two Approaches:**
  2. Atomically acquire all resources at once
    - » Example: Single lock to protect all (other variations - e.g., release access to one variable earlier)

Thread Maria

```
lock(AB);
A += 10;
B += 20;
A += 30;
unlock(AB);
```

Thread Tucker

```
lock(AB);
B += 10;
A += 20;
B += 30;
unlock(AB);
```

Mutual exclusion  
Hold and wait  
No preemption  
Circular wait

Maria Hyönnelä, UGA

12

## Deadlock Prevention

### Hold and Wait

- **Summary the Two Approaches:**
  1. Only request resources when it does not hold other resources
  2. Atomically acquire all resources at once
- **Problems:**
  - » Low resource utilization: ties up resources other processes could be using
  - » May not know required resources before run
  - » **Starvation:** A thread that need popular resources may wait forever

Mutual exclusion  
Hold and wait  
No preemption  
Circular wait

Maria Hyönnelle, UGA

13

## Deadlock Prevention

### No Preemption

- **Two Approaches:**
  1. Preempt **requestors** resource
    - **Example:** B is holding some resources and then requests additional resources that are held by other threads, then B releases all its resources (and start over)
  2. Preempt **holders** resource
    - **Example:** A waiting for something held by B, then take resource away from B and give them to A (B starts over).
- **Not possible if resource cannot be saved and restored**
  - » Can't take away a lock without causing problems
- **Only works for some resources (e.g., CPU and memory)**

Mutual exclusion  
Hold and wait  
No preemption  
Circular wait

Maria Hyönnelle, UGA

14

## Deadlock Prevention

### Circular Wait Condition

- **Impose ordering on resources**
  - » Give all resources a ranking; must acquire highest ranked resource first

Mutual exclusion  
Hold and wait  
No preemption  
Circular wait

Maria Hyönnelle, UGA

15

## Deadlock Detection & Recovery

1. Allow system to enter deadlock state
2. Detection algorithm
3. Recovery scheme


Maria Hyönnelle, UGA

16

## Deadlock Detection

### Single Instance of Each Resource Type

- Maintain **wait-for** graph
  - » Nodes are processes.
  - » removes resource nodes and collapsing edges
  - »  $P_i \rightarrow P_j$  if  $P_j$  is waiting for  $P_i$ .
- Periodically invoke an algorithm that searches for a cycle in the graph.
- 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.



Maria Hyönnelle, UGA

17

## Depth first search (example)

For each node in the graph:

```
L = {empty list} and Nodes = {unvisited};
current node = initial node ;
while( current node is not the initial node twice )
  L.enqueue(current node); // add to node to end of L
  if( current node is in L twice )
    there is a cycle => cycle and return
  if( there is an unmarked arc )
    mark the arc as visited and use destination as new
    current node
  else
    go back to previous node
Back to initial node there is no cycle
```

Maria Hyönnelle, UGA

18

# Deadlock detection (1 resource of each)


- Do a depth-first-search on the resource allocation graph

D, E, G ?

are deadlocked

A, C, F ?

are not deadlocked because S can be allocated to either and then the other two can take turn to complete




Maria Hyömette, UGA

19

# Example: Deadlock Detection

- Do a depth-first-search on the resource allocation graph

Initialize a list to the empty list, designate arcs as 'unvisited'




Maria Hyömette, UGA

20

# Example: Deadlock Detection

- Do a depth-first-search on the resource allocation graph




Maria Hyömette, UGA

21

# Example: Deadlock Detection

- Do a depth-first-search on the resource allocation graph




Maria Hyömette, UGA

22

# Example: Deadlock Detection

- Do a depth-first-search on the resource allocation graph




Maria Hyömette, UGA

23

# Deadlock Detection with Multiple Resources

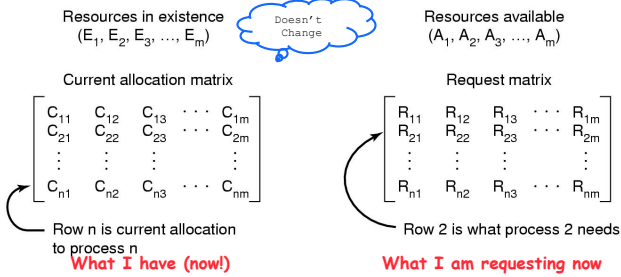
- Theorem:** If a graph does not contain a cycle then no processes are deadlocked
  - A cycle in a RAG is a necessary condition for deadlock
  - Is it a sufficient condition?



Maria Hyömette, UGA

24



# Deadlock Detection Algorithm: Multiple Resource Instances



- **Available:** Indicates the number of available resources of each type (m)
  - **Allocation:** Number of resources of each type currently allocated (nxm)
  - **Request:** current requests of each thread (nxm)
- » If Request  $[j] = k$ , then process  $P_j$  is requesting k more instances of type  $R_j$ .

# Example

- Is there a possible allocation sequence of resources so that each process can complete?




# Detection algorithm

A marked process means it can run to completion


1. Look for an unmarked process  $P_i$ , for which the  $i$ th row of R (need) 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

*If all marked, no deadlock*


# Detection algorithm



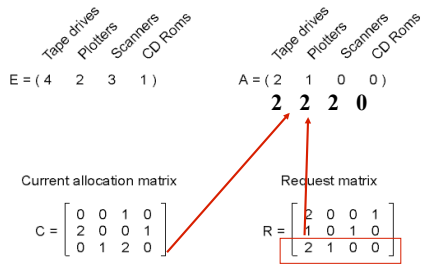
# Detection algorithm



# Detection algorithm



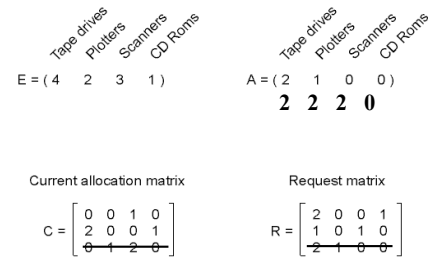
## Detection algorithm



Maria Hyömette, UGA

31


## Detection algorithm



Maria Hyömette, UGA

32


## Detection algorithm



Maria Hyömette, UGA

33

## Detection algorithm



Maria Hyömette, UGA

34

No deadlock!

## Deadlock detection issues

- How often should the algorithm run?
  - » After every resource request?
  - » Periodically?
  - » When CPU utilization is low?
  - » When we suspect deadlock because some thread has been asleep for a long period of time?

Maria Hyömette, UGA

35

## Recovery from deadlock

- What should be done to recover?
  - » Abort deadlocked processes and reclaim resources
  - » Temporarily reclaim resource, if possible
  - » Abort one process at a time until deadlock cycle is eliminated
- Where to start?
  - » Low priority process
  - » How long process has been executing
  - » How many resources a process holds
  - » Batch or interactive
  - » Number of processes that must be terminated

Maria Hyömette, UGA

36

## Other deadlock recovery techniques

- **Recovery through rollback**
  - » Save state periodically
    - take a checkpoint
    - start computation again from checkpoint
  - » Done for large computation systems

Maria Hyönnelie, UGA

37

## Review: Handling Deadlock

- **Ignore**
  - » Easiest and most common approach (e.g., UNIX).
- **Deadlock prevention**
  - » Ensure deadlock does not happen
  - » Ensure at least one of 4 conditions does not occur
- **Deadlock detection and recovery**
  - » Allow deadlocks, but detect when occur
  - » Recover and continue
- **Deadlock avoidance**
  - » Ensure deadlock does not happen
  - » Use information about resource requests to dynamically avoid unsafe situations



Ostrich algorithm

Maria Hyönnelie, UGA

38


## Deadlock avoidance

- **Detection vs. avoidance...**
  - » **Detection** – “optimistic” approach
    - Allocate resources
    - “Break” system to fix it
  - » **Avoidance** – “pessimistic” (conservative) approach
    - Don’t allocate resource if it may lead to deadlock
    - If a process requests a resource...
      - ... make it wait until you are sure it’s OK
  - » Which one to use depends upon the application

Maria Hyönnelie, UGA

39

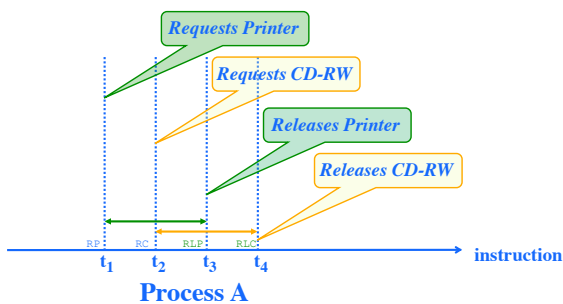
## Process-resource trajectories



Maria Hyönnelie, UGA

40

## Process-resource trajectories



Maria Hyönnelie, UGA

41

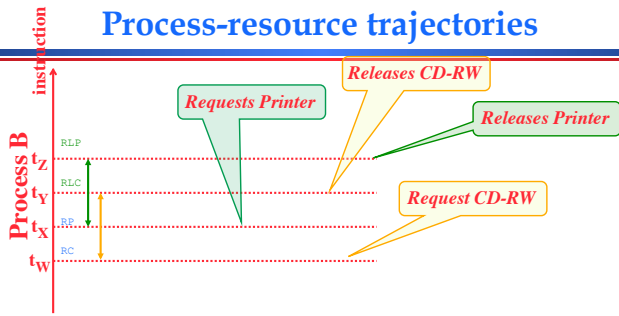
## Process-resource trajectories




Maria Hyönnelie, UGA

42

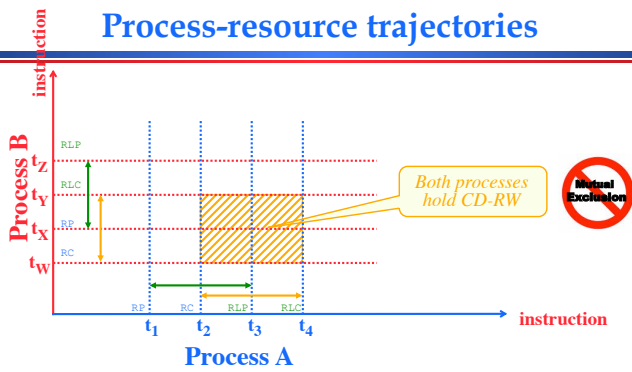
## Process-resource trajectories




## Process-resource trajectories



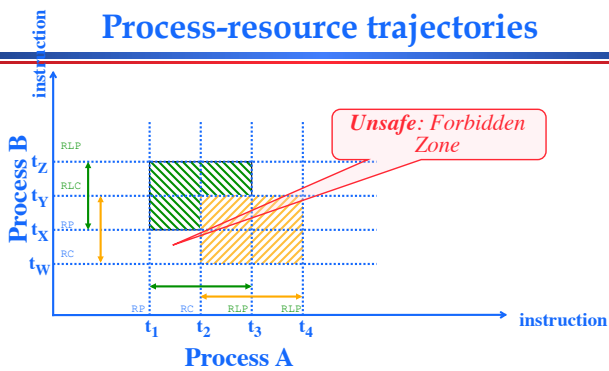
## Process-resource trajectories




## Process-resource trajectories




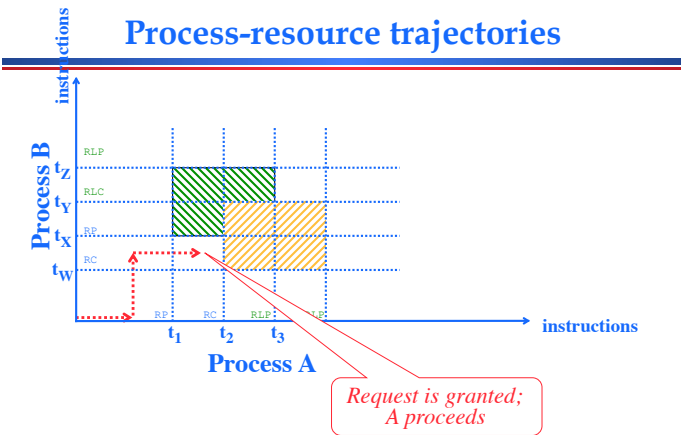
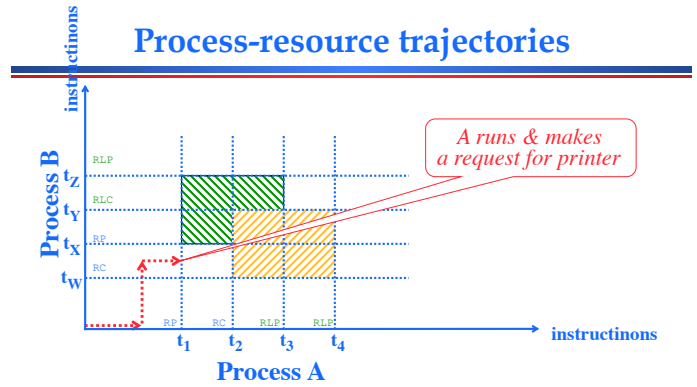
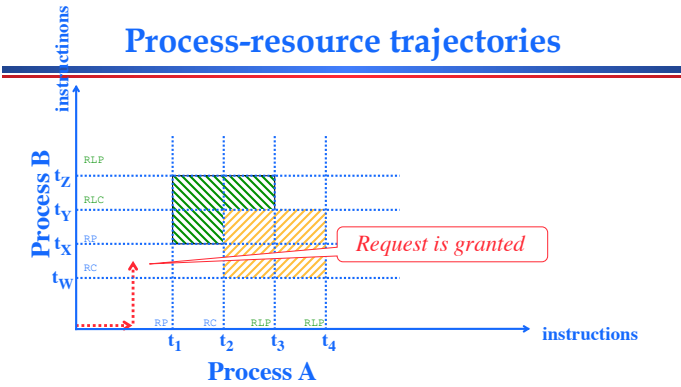


## Process-resource trajectories

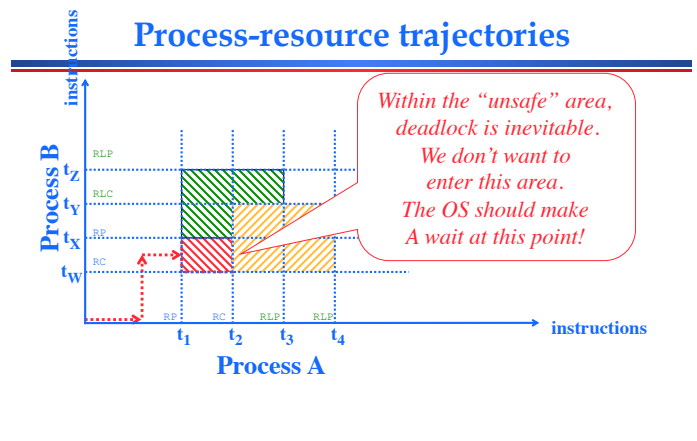




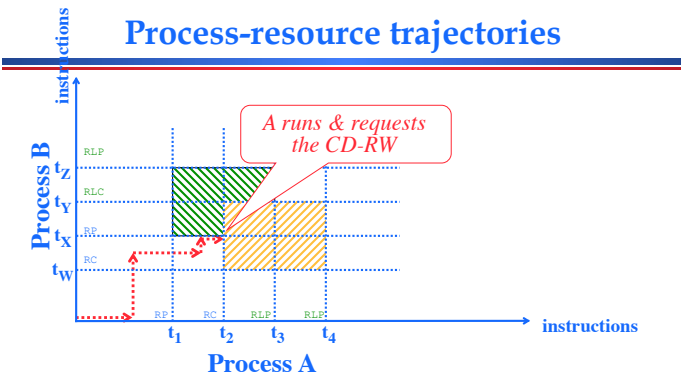


## Process-resource trajectories

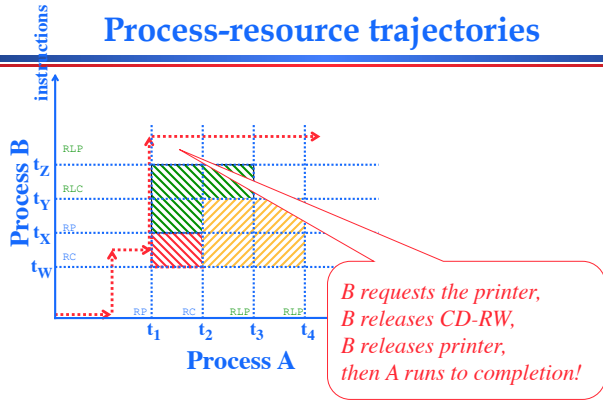








## Process-resource trajectories



Maria Hyönnelä, UGA

61

## Safe states

- The current state: "which processes hold which resources"
- A "safe" state:
  - » No deadlock, and
  - » There is some scheduling order in which every process can run to completion even if all of them request their maximum number of units immediately
- The Banker's Algorithm:
  - » Goal: Avoid unsafe states!!!
  - » Question: When a process requests more units, should the system grant the request or make it wait?

Maria Hyönnelä, UGA

62

## Deadlock Avoidance

- Dijkstra's Banker's Algorithm
- Idea: Avoid unsafe states of processes holding resources
  - » Unsafe states might lead to deadlock if processes make certain future requests
  - » When process requests resource, only give if doesn't cause unsafe state
  - » Problem: Requires processes to specify all possible future resource demands

Maria Hyönnelä, UGA

63

## The Banker's Algorithm

- Assumptions:
  - » Only one type of resource, with multiple units.
  - » Processes declare their maximum potential resource needs ahead of time (total sum is 22 units of credit but only has 10)
- When a process requests more units should the system make it wait to ensure safety?

Example: One resource type with 10 units

	Has	Max		Has	Max		Has	Max		Has	Max		Has	Max	
A	3	9	6	A	3	9	A	3	9	A	3	9	A	3	9
B	2	4	2	B	4	4	B	0	-	B	0	-	B	0	-
C	2	7	5	C	2	7	C	2	7	C	7	7	C	0	-
	Free: 3			Free: 1		Free: 5	Free: 0		Free: 7						

Maria Hyönnelä, UGA

64

## Safe states

- Safe state – "when system is not deadlocked and there is some scheduling order in which every process can run to completion even if all of them suddenly request their maximum number of resource immediately"

10 total

	Has	Max		Has	Max		Has	Max		Has	Max		Has	Max	
A	3	9	6	A	3	9	A	3	9	A	3	9	A	3	9
B	2	4	2	B	4	4	B	0	-	B	0	-	B	0	-
C	2	7	5	C	2	7	C	2	7	C	7	7	C	0	-
	Free: 3			Free: 1		Free: 5	Free: 0		Free: 7						

Maria Hyönnelä, UGA

65

## Unsafe/Safe state?

The difference here is A possesses 1 more resource

	Has	Max		Has	Max		Has	Max		Has	Max
A	3	9	6	A	4	9	5	A	4	9	
B	2	4	2	B	2	4	2	B	4	4	
C	2	7	5	C	2	7	5	C	2	7	
	Free: 3			Free: 2			Free: 0		Free: 4		

Safe Unsafe

Maria Hyönnelä, UGA

66

## Avoidance with multiple resource types

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

Maximum # Needed

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

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

Row n is current allocation to process n

*Note: These are the max. possible requests, which we assume are known ahead of time*

Process 2 needs

Maria

67

## Banker's algorithm for multiple resources

- Look for a row,  $R$ , whose unmet resource needs are all smaller than or equal to  $A$ . If no such row exists, the system will eventually deadlock since no process can run to completion
- Assume the process of the row chosen requests all the resources that it needs (which is guaranteed to be possible) and finishes. Mark that process as terminated and add all its resources to  $A$  vector
- Repeat steps 1 and 2, until either all process are marked terminated, in which case the initial state was safe, or until deadlock occurs, in which case it was not

Maria Hyömette, UGA

68

## Avoidance modeling

Total resource vector

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

Available resource vector

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

Current allocation matrix

Maximum Request Vector

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

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

Row n is current allocation to process n

Row 2 is what process 2 might need

**RUN ALGORITHM ON EVERY RESOURCE REQUEST**

Maria Hyömette, UGA

69

## Avoidance algorithm

Tape drives  
Plotters  
Scanners  
CD Roms  
 $E = (4 \ 2 \ 3 \ 1)$

Tape drives  
Plotters  
Scanners  
CD Roms  
 $A = (2 \ 1 \ 0 \ 0)$

Current allocation matrix

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

More needed matrix

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

Maria Hyömette, UGA

70

## Avoidance algorithm

Tape drives  
Plotters  
Scanners  
CD Roms  
 $E = (4 \ 2 \ 3 \ 1)$

Tape drives  
Plotters  
Scanners  
CD Roms  
 $A = (2 \ 1 \ 0 \ 0)$

Current allocation matrix

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

More needed matrix

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

Maria Hyömette, UGA

71

## Avoidance algorithm

Tape drives  
Plotters  
Scanners  
CD Roms  
 $E = (4 \ 2 \ 3 \ 1)$

Tape drives  
Plotters  
Scanners  
CD Roms  
 $A = (2 \ 1 \ 0 \ 0)$

Current allocation matrix

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

More needed matrix

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

Maria Hyömette, UGA

72

## Avoidance algorithm

$$E = (4 \quad 2 \quad 3 \quad 1)$$

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

Current allocation matrix

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

More needed matrix

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

Maria Hyönelle, UGA

73

## Avoidance algorithm

$$E = (4 \quad 2 \quad 3 \quad 1)$$

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

Current allocation matrix

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

More needed matrix

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

Maria Hyönelle, UGA

74

## Avoidance algorithm

$$E = (4 \quad 2 \quad 3 \quad 1)$$

$$A = \begin{pmatrix} 2 & 1 & 0 & 0 \\ 2 & 2 & 2 & 0 \\ 4 & 2 & 2 & 1 \end{pmatrix}$$

Current allocation matrix

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

More needed matrix

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

Maria Hyönelle, UGA

75

## Deadlock avoidance



- **Deadlock avoidance is usually impossible**
  - » because you don't know in advance what resources a process will need!



Maria Hyönelle, UGA

76