

Slide 1

---

**Deadlock**  
**COMP3231 Operating Systems**  
**2005 S2**

Slide 2

---

**DEADLOCK**

What is a **deadlock**?


- **Permanent blocking** of a set of processes that either
  - compete for system resources or
  - communicate with each other (message as resource)
- Resources:
  - preemptable
  - nonpreemptable resources
- Deadlocks involve conflicting needs for nonpreemptable resources by two or more processes
- Deadlocks can occur on many levels in the system
  
- ✗ Unfortunately, there is no efficient method to prevent a deadlock in the general case

Let's look at some examples and at the conditions for deadlock

---

Slide 3

Danger of deadlock in continental driving rules:



Slide 4

---

**REUSABLE VERSUS CONSUMABLE RESOURCES**

- **Reusable resource**: used by one process at a time and **not depleted by that use**
  - **Consumable resource**: created (produced) and destroyed (consumed) by a process
- Reusable Resources:**
- Processes obtain resources that they later release for reuse by other processes
  - Examples are processors, I/O channels, main and secondary memory, files, databases, and semaphores
  - In case of two processes and two resources, deadlock occurs if each process holds one resource and requests the other

Typical deadlock with reusable resources:

Process P		Process Q	
Step	Action	Step	Action
P <sub>0</sub>	Request (D)	q <sub>0</sub>	Request (T)
P <sub>1</sub>	Lock (D)	q <sub>1</sub>	Lock (T)
P <sub>2</sub>	Request (T)	q <sub>2</sub>	Request (D)
P <sub>3</sub>	Lock (T)	q <sub>3</sub>	Lock (D)
P <sub>4</sub>	Perform function	q <sub>4</sub>	Perform function
P <sub>5</sub>	Unlock (D)	q <sub>5</sub>	Unlock (T)
P <sub>6</sub>	Unlock (T)	q <sub>6</sub>	Unlock (D)

Slide 5

The following sequence leads to a deadlock:

P<sub>0</sub>, P<sub>1</sub>, q<sub>0</sub>, q<sub>1</sub>, P<sub>2</sub>, q<sub>2</sub>

► Should this really be the problem of the OS designer?

Another example of deadlock with reusable resources:

→ Space is available for allocation of 200K bytes and the following sequence of events occur

P <sub>1</sub>	P <sub>2</sub>
...	...
Request 80kB;	Request 70kB;
...	...
Request 60kB;	Request 80kB;

Slide 6

→ Deadlock occurs if both processes progress to their second request  
 → In this case, the problem can be solved by using virtual memory (this is an example of **resource preemption**)

Consumable Resources:

- Interrupts, signals, messages, and information in I/O buffers
- Deadlock may occur if a Receive message is blocking
- May take a rare combination of events to cause deadlock

Example of deadlock:

→ Deadlock occurs if receive is blocking

Slide 7

P <sub>1</sub>	P <sub>2</sub>
...	...
Receive(P <sub>2</sub> );	Receive(P <sub>1</sub> );
...	...
Send(P <sub>2</sub> , M <sub>1</sub> );	Send(P <sub>1</sub> , M <sub>2</sub> );

## CONDITIONS FOR DEADLOCK


How can we accurately characterise the conditions that lead to a deadlock?

Necessary conditions for deadlock:

- ① **Mutual exclusion**: only one process may use a resource at a time
- ② **Hold-and-wait**: a process holds a resource while awaiting assignment of others
- ③ **No preemption of resources**:
  - A process that is denied a request must not release the resources it already has
  - When one process requests a resource held by another, the second one is not preempted by the OS
- ④ **Circular wait**: we have a closed chain of processes, such that each process holds at least one resource needed by the next in the chain, e.g.,

Slide 8

Slide 9



## DEADLOCK PREVENTION

What is deadlock prevention?

Make it impossible that one of the four conditions for deadlock arise

- ① mutual exclusion
- ② hold-and-wait
- ③ no preemption
- ④ circular wait

Mutual exclusion:

- we can't generally exclude it
- we can avoid assigning resources when not absolutely necessary
- as few processes as possible should claim the resource

Slide 11

## STRATEGIES TO DEAL WITH DEADLOCKS

- ① The Ostrich Algorithm
- ② Prevention
- ③ Avoidance by careful resource allocation
- ④ Detection and Recovery: let them occur, detect them and take action

Slide 10

The Ostrich Algorithm:

Stick your head in the sand and pretend there is no problem at all!

- Unix & Windows
- Avoid deadlock in the kernel!

Hold-and-wait:

- Can we require processes to request all resources at once?
- Most processes do not statically know about the resources they need
- Used in some mainframe batch systems
- Wasteful, but works
- Variation: before requesting new resource, temporarily release other resources

Slide 12

---

No preemption:

Preemption is feasible for some resources (e.g., processor and memory), but not for others (state must be saved and restored)

Circular wait:

Slide 13

- order resources by an index:  $R_1, R_2, \dots$
- requires that resources are always requested in order
- $P_1$  holds  $R_i$  and requests  $R_j$ , and  $P_2$  holds  $R_j$  and requests  $R_i$  is impossible, as it implies
$$i < j \text{ and } i > j$$
- is sometimes a feasible strategy, but not generally efficient

---

### DEADLOCK AVOIDANCE

What is deadlock avoidance?:

- We don't exclude any of the four conditions for deadlock per se
- Instead we decide on a per case basis whether a process is deemed likely to deadlock
- Thus, we have to possess some knowledge about future allocation requests of processes

Slide 14

Generally, we can distinguish two approaches to deadlock avoidance:

- **Process initiation denial:** we just don't start a process if it might deadlock
- **Resource allocation denial:** we deny allocation requests, which are likely to lead to deadlock in the future

---

### PROCESS INITIATION DENIAL

Consider a system of  $n$  processes and  $m$  types of resources:

→ **Resource vector:**  $(R_1, R_2, \dots, R_n)$

→ **Available vector:**  $(V_1, V_2, \dots, V_n)$

→ **Matrices:**

**Claim matrix:**

**Allocation matrix:**

Slide 15

$$\begin{pmatrix} C_{11} & C_{12} & \dots & C_{1m} \\ C_{21} & C_{22} & \dots & C_{2m} \\ \vdots & \vdots & \vdots & \vdots \\ C_{n1} & C_{n2} & \dots & C_{nm} \end{pmatrix} \quad \begin{pmatrix} A_{11} & A_{12} & \dots & A_{1m} \\ A_{21} & A_{22} & \dots & A_{2m} \\ \vdots & \vdots & \vdots & \vdots \\ A_{n1} & A_{n2} & \dots & A_{nm} \end{pmatrix}$$

→  $C_{ij}$  requirement for process  $i$  for resource  $j$

→  $A_{ij}$  allocation of resource  $j$  to process  $i$

---

**Example:** We have two processes  $P_1$  and  $P_2$  and three resources  $R_1, R_2$  and  $R_3$ . Each of the three resources can be allocated to only a single process at each point in time

- $P_1$ 
  - holds  $R_1$
  - requires  $R_1, R_2$

Slide 16

- $P_2$ 
  - holds no resource
  - requires  $R_2, R_3$

→ **Resource vector:**  $(1, 1, 1)$

→ **Available vector:**  $(0, 1, 1)$

**Claim matrix:**

**Allocation matrix:**

$$\begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix} \quad \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

The following relationships hold:

- ①  $R_i = V_i + \sum_{k=1}^n A_{ki}$  : all resources are either available or allocated
- ②  $C_{kj} \leq R_j$ : no process can hold more than the total amount of resources in the system
- ③  $A_{kj} \leq C_{kj}$ : no process is allocated more than it originally claimed to need

**Slide 17** Deadlock avoidance policy:

→ Start a new process  $P_{n+1}$  **only if**, for all  $i$ ,

$$V_i \geq C_{n+1,i} + \sum_{k=1}^n C_{ki}$$

- Unfortunately, this strategy is very wasteful!
- Assumes all processes make their claims together

**RESOURCE ALLOCATION DENIAL**

- At any request of a resource, it is tested whether granting this request bears the **potential** of deadlock
- The standard algorithm to execute this test is due to Dijkstra and known as the **banker's algorithm**

Banker's algorithm:

- Slide 18**
- Resource and available vector & claim and allocation matrix as before
  - The algorithm passes out resources to processes if it has enough on hand to meet potential future demand
  - Whenever we can guarantee that future demand can be met, we are in a **safe state**
  - A request for resources is granted **only if** the state **after** the resource is granted is safe

How do we know whether a state is safe?

- A state is safe if there is at least one sequence of resource allocations that does not result in deadlock
- Pick a process whose outstanding resource claim can be met and run it to completion
- Repeat until either all process have completed, or the system locks up

**Slide 19**

Check that this state is safe:

	R1	R2	R3
P1	3	2	2
P2	6	1	3
P3	3	1	4
P4	4	2	2

Claim Matrix

	R1	R2	R3
P1	1	0	0
P2	6	1	2
P3	2	1	1
P4	0	0	2

Allocation Matrix

R1	R2	R3
9	3	6

Resource Vector

R1	R2	R3
0	1	1

Available Vector

**Slide 20**

P2 runs to completion:

	R1	R2	R3
P1	3	2	2
P2	0	0	0
P3	3	1	4
P4	4	2	2

Claim Matrix

	R1	R2	R3
P1	1	0	0
P2	0	0	0
P3	2	1	1
P4	0	0	2

Allocation Matrix

R1	R2	R3
6	2	3

Available Vector

P1 Runs to Completion:

	R1	R2	R3
P1	0	0	0
P2	0	0	0
P3	3	1	4
P4	4	2	2

Claim Matrix

	R1	R2	R3
P1	0	0	0
P2	0	0	0
P3	2	1	1
P4	0	0	2

Allocation Matrix

R1	R2	R3
7	2	3

Available Vector

Slide 21

P3 Runs to Completion:

	R1	R2	R3
P1	0	0	0
P2	0	0	0
P3	0	0	0
P4	4	2	2

Claim Matrix

	R1	R2	R3
P1	0	0	0
P2	0	0	0
P3	0	0	0
P4	0	0	2

Allocation Matrix

R1	R2	R3
9	3	4

Available Vector

Example of a request leading to an unsafe state:

	R1	R2	R3
P1	3	2	2
P2	6	1	3
P3	3	1	4
P4	4	2	2

Claim Matrix

	R1	R2	R3
P1	1	0	0
P2	5	1	1
P3	2	1	1
P4	0	0	2

Allocation Matrix

R1	R2	R3
9	3	6

Resource Vector

R1	R2	R3
1	1	2

Available Vector

Slide 22

P1 requests R1 & R3:

	R1	R2	R3
P1	3	2	2
P2	6	1	3
P3	3	1	4
P4	4	2	2

Claim Matrix

	R1	R2	R3
P1	2	0	1
P2	5	1	1
P3	2	1	1
P4	0	0	2

Allocation Matrix

R1	R2	R3
0	1	1

Available Vector

Disadvantages of the Banker's algorithm:

- Maximum resource requirement must be stated in advance
- Processes under consideration must be independent; no synchronization requirements
- There must be a fixed number of resources to allocate
- No process may exit while holding resources

Slide 23

### DEADLOCK DETECTION

- An alternative to deadlock avoidance is **deadlock detection**
- However, for this to be useful, we require to be able to either
  - roll processes back (in the extreme case, kill them) or
  - preempt resources

Modification of Banker's algorithm for deadlock detection:

- We need a **request matrix Q** (outstanding requests) instead of the claim matrix
- Disregard processes without any allocation (not holding resources)
- Consider process completed if outstanding requests are satisfied
- Checks can be made each time a resource is allocated
  - early deadlock detection
  - expensive

Slide 24

Algorithm:

Initially, all processes are unmarked

Slide 25

- ① mark each process with zero-row in Request matrix
- ② set temporary vector  $W$  to Available vector
- ③ find  $i$  such that process  $i$  is unmarked,  $Q_{ik} \leq W_k$  for  $1 \leq k \leq n$   
- no such process  $\Rightarrow$  terminate
- ④ mark process  $i$ , add row of allocation matrix to  $W$ , go to step 3

Slide 26

	R1	R2	R3	R4	R5
P1	0	1	0	0	1
P2	0	0	1	0	1
P3	0	0	0	0	1
P4	1	0	1	0	1

Request Matrix Q

	R1	R2	R3	R4	R5
P1	1	0	1	1	0
P2	1	1	0	0	0
P3	0	0	0	1	0
P4	0	0	0	0	0

Allocation Matrix A

R1	R2	R3	R4	R5
2	1	1	2	1

Resource Vector

R1	R2	R3	R4	R5
0	0	0	0	1

Available Vector

Recovery:

Slide 27

- ① Abort all deadlocked processes (most common solution)
- ② Rollback each deadlocked process to some previously defined checkpoint and restart them (original deadlock may reoccur)
- ③ Successively abort deadlocked processes until deadlock no longer exists (invoke deadlock detection algorithm each time)
- ④ Successively preempt some resources from process until deadlock no longer exists  
- a process that has a resource preempted must be rolled back prior to its acquisition