Operating Systems: Internals and Design Principles, 6/E William Stallings

# Chapter 6 Concurrency: Deadlock and Starvation

Patricia Roy Manatee Community College, Venice, FL ©2008, Prentice Hall

## Deadlock

- Permanent blocking of a set of processes that either compete for system resources or communicate with each other
- No efficient solution
- Involve conflicting needs for resources by two or more processes

#### **Deadlock in Traffic**



(a) Deadlock possible

(b) Deadlock

#### Figure 6.1 Illustration of Deadlock

#### Non-deadlock - Joint Progress Diagram



Figure 6.3 Example of No Deadlock [BACO03]

#### Deadlock in a Computer – Fatal Region



Figure 6.2 Example of Deadlock

# **Deadlock Definition**

• 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

- Used by only one process at a time and not depleted by that use
- Processes obtain resources that they later release for reuse by other processes

- Processors, I/O channels, main and secondary memory, devices, and data structures such as files, databases, and semaphores
- Deadlock occurs if each process holds one dedicated resource and requests another held by another process

Process P

Process Q

Step	Action	Step	Action
$\mathbf{p}_0$	Request (D)	$\mathbf{q}_0$	Request (T)
$\mathbf{p}_1$	Lock (D)	$\mathbf{q}_1$	Lock (T)
<b>p</b> <sub>2</sub>	Request (T)	$\mathbf{q}_2$	Request (D)
<b>p</b> <sub>3</sub>	Lock (T)	$q_3$	Lock (D)
$p_4$	Perform function	$\mathbf{q}_4$	Perform function
<b>p</b> <sub>5</sub>	Unlock (D)	$q_5$	Unlock (T)
$\mathbf{p}_6$	Unlock (T)	q <sub>6</sub>	Unlock (D)

Figure 6.4 Example of Two Processes Competing for Reusable Resources

 Space is available for allocation of 200Kbytes, and the following sequence of events occur

P1	P2
• • •	•••
Request 80 Kbytes;	Request 70 Kbytes;
•••	•••
Request 60 Kbytes;	Request 80 Kbytes;

- Deadlock occurs if both processes progress to their second request
- But virtual memory

# Consumable Resources

- Created (produced) and destroyed (consumed)
- 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



# **Conditions for Deadlock**

- Mutual exclusion
  - Only one process may use a resource at a time
- Hold-and-wait
  - A process may hold allocated resources while awaiting assignment of others

# **Conditions for Deadlock**

- No preemption
  - No resource can be forcibly removed from a process holding it
- Circular wait
  - A closed chain of processes exists, such that each process holds at least one resource needed by the next process in the chain

## **Resource Allocation Graphs**

• Directed graph that depicts a state of the system of resources and processes



### **Resource Allocation Graphs**



(c) Circular wait

(d) No deadlock

### **Resource Allocation Graphs**



#### Figure 6.6 Resource Allocation Graph for Figure 6.1b

# **Deadlock Prevention**

- Mutual Exclusion
  - Spooling
- Hold and Wait
  - Require that a process request all of its required resources at one time
  - Requests would be granted/denied simultaneously

# Deadlock Prevention (cont.)

- No Preemption
  - Process must release resource and request again
  - OS may preempt a process and require it to release its resources
- Circular Wait
  - Define a linear ordering of resources
  - Require that processes request resources according to the ordering

# Deadlock Avoidance

- A decision is made dynamically whether the current resource allocation request will, if granted, potentially lead to a deadlock
- Requires knowledge of future process requests

# Two Approaches to Deadlock Avoidance

- Do not start a process if its demands might lead to deadlock
- Do not grant an incremental resource request to a process if this allocation might lead to deadlock

# **Resource Allocation Denial**

- Referred to as the Banker's Algorithm
- State of the system is the current allocation of resources to process
- Safe state is where there is at least one sequence of execution of processes that does not result in deadlock
- Unsafe state is a state that is not safe

### Determination of a Safe State



(a) Initial state

### Determination of a Safe State



(b) P2 runs to completion

#### **Determination of an Unsafe State**



(b) P1 requests one unit each of R1 and R3

### **Deadlock Avoidance Logic**

```
struct state {
    int resource[m];
    int available[m];
    int claim[n][m];
    int alloc[n][m];
}
```

#### (a) global data structures

(b) resource alloc algorithm

### **Deadlock Avoidance Logic**

```
boolean safe (state S) {
   int currentavail[m];
   process rest[<number of processes>];
   currentavail = available;
   rest = {all processes};
   possible = true;
   while (possible) {
      <find a process Pk in rest such that
          claim [k,*] - alloc [k,*] <= currentavail;>
                                          /* simulate execution of Pk */
      if (found) {
          currentavail = currentavail + alloc [k,*];
          rest = rest - {Pk};
      else possible = false;
   return (rest == null);
```

(c) test for safety algorithm (banker's algorithm)

Figure 6.9 Deadlock Avoidance Logic

# Deadlock Avoidance

- Maximum resource requirement must be stated in advance
- Processes under consideration must be independent; no synchronization (order of execution) requirements
- No process may exit/block while holding resources

#### **Deadlock Detection**



#### Figure 6.10 Example for Deadlock Detection

# Strategies Once Deadlock Detected

- Abort all deadlocked processes
- Back up each deadlocked process to some previously defined checkpoint, and restart all process - original deadlock may re-occur
- Successively abort deadlocked processes until deadlock no longer exists
- Successively preempt resources until deadlock no longer exists

## **Dining Philosophers Problem**



Figure 6.11 Dining Arrangement for Philosophers

#### **Dining Philosophers Problem**

```
/* program diningphilosophers */
semaphore fork [5] = {1};
int i;
void philosopher (int i)
{
     while (true) {
          think();
          wait (fork[i]);
          wait (fork [(i+1) mod 5]);
          eat();
          signal(fork [(i+1) mod 5]);
          signal(fork[i]);
     }
}
void main()
{
     parbegin (philosopher (0), philosopher (1), philosopher
(2),
          philosopher (3), philosopher (4));
```

#### Figure 6.12 A First Solution to the Dining Philosophers Problem

#### **Dining Philosophers Problem with Semaphores**

```
/* program diningphilosophers */
semaphore fork[5] = {1};
semaphore room = {4};
int i;
void philosopher (int i)
{
   while (true) {
     think();
     wait (room);
     wait (fork[i]);
     wait (fork [(i+1) mod 5]);
     eat();
     signal (fork [(i+1) mod 5]);
     signal (fork[i]);
     signal (room);
ł
void main()
ł
   parbegin (philosopher (0), philosopher (1), philosopher (2),
          philosopher (3), philosopher (4));
```

#### Figure 6.13 A Second Solution to the Dining Philosophers Problem

#### **Dining Philosophers Problem with Monitor**

Figure 6.14 A Solution to the Dining Philosophers Problem Using a Monitor

#### **Dining Philosophers Problem with Monitor**

```
monitor dining controller;
cond ForkReady[5]; /* condition variable for synchronization */
/* pid is the philosopher id number */
void get forks(int pid)
{
  int left = pid;
  int right = (++pid) % 5;
  /*grant the left fork*/
  if (!fork(left)
                              /* queue on condition variable */
    cwait(ForkReady[left]);
  fork(left) = false;
  /*grant the right fork*/
  if (!fork(right)
    fork(right) = false:
}
void release forks(int pid)
Ł
  int left = pid;
  int right = (++pid) % 5;
  /*release the left fork*/
                         /*no one is waiting for this fork */
  if (empty(ForkReady[left])
    fork(left) = true;
                       /* awaken a process waiting on this fork */
  else
    csignal(ForkReady[left]);
  /*release the right fork*/
  if (empty(ForkReady[right]) /*no one is waiting for this fork */
    fork(right) = true;
                       /* awaken a process waiting on this fork */
  else
    csignal(ForkReady[right]);
```

# **UNIX Concurrency Mechanisms**

- Pipes: circular buffer for two processes like in producer-consumer
- Messages: blocking receive
- Shared memory: shared pages mutual exclusion is not guaranteed
- Semaphores
- Signals

# Linux Kernel Concurrency Mechanism

- Includes all the mechanisms found in UNIX
- Atomic operations execute without interruption and without interference (by blocking the memory bus)

# Linux Atomic Operations

- arithmetic operation plus setting condition code
- spinlocks for mutual exclusion (loop until lock acquired)
- traditional and readers-writer semaphores
- memory barrier operations: limit compiler or CPU in re-ordering instructions