

Deadlocks

Contents of Lecture:

- ❖ System Model
- ❖ Deadlock Characterization
- ❖ Methods for Handling Deadlocks
- ❖ Deadlock Prevention

References for This Lecture:

- ✓ *Abraham Silberschatz, Peter Bear Galvin and Greg Gagne, Operating System Concepts, 9th Edition, Chapter 7*

System Model

- ❖ A system consists of a finite number of resources to be distributed among a number of competing processes.
- ❖ The resources may be partitioned into several types (or classes (R1, R2, . . ., Rm)), each consisting of some number of identical instances.
 - ✓ CPU cycles, files, and I/O devices (such as printers and DVD drives) are examples of resource types.
 - ✓ If a system has two CPUs, then the resource type CPU has two instances. Similarly, the resource type printer may have five instances
- ❖ Under the normal mode of operation, a process may utilize a resource in only the following sequence:
 1. **Request:** The process requests the resource. If the request cannot be granted immediately (for example, if the resource is being used by another process), then the requesting process must wait until it can acquire the resource.
 2. **Use:** The process can operate on the resource (for example, if the resource is a printer, the process can print on the printer).
 3. **Release:** The process releases the resource.

Deadlock Characterization

- ❖ In a deadlock, processes never finish executing, and system resources are tied up, preventing other jobs from starting.

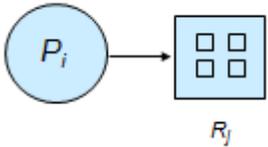
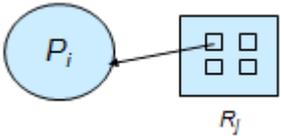
Necessary Conditions

- ❖ A deadlock situation can arise if the following four conditions hold simultaneously in a system:
 1. **Mutual exclusion:** At least one resource must be held in a nonsharable mode; only one process at a time can use the resource. If another process requests that resource, the requesting process must be delayed until the resource has been released.
 2. **Hold and wait:** A process must be holding at least one resource and waiting to acquire additional resources that are currently being held by other processes.
 3. **No preemption:** Resources cannot be preempted; a resource can be released only voluntarily by the process holding it, after that process has completed its task.
 4. **Circular wait:** There exists a set $\{P_0, P_1, \dots, P_n\}$ of waiting processes such that P_0 is waiting for a resource held by P_1 , P_1 is waiting for a resource held by P_2 , ..., P_{n-1} is waiting for a resource held by P_n , and P_n is waiting for a resource held by P_0 .

Resource-Allocation Graph

- ❖ Deadlocks can be described more precisely in terms of a directed graph called a **system resource-allocation graph**.
- ❖ This graph consists of a set of **vertices V** and a set of **edges E**.
 - ✓ The **set of vertices V** is partitioned into two different types of nodes:
 - $P = \{P_1, P_2, \dots, P_n\}$, the set consisting of all the **active processes** in the system.
 - $R = \{R_1, R_2, \dots, R_m\}$, the set consisting of all **resource types** in the system.
 - ✓ A directed edge from process P_i to resource type R_j is denoted by $P_i \rightarrow R_j$; it signifies that process P_i has requested an instance of resource type R_j and is currently waiting for that resource.
 - A directed edge $P_i \rightarrow R_j$ is called a **request edge**.
 - ✓ A directed edge from resource type R_j to process P_i is denoted by $R_j \rightarrow P_i$; it signifies that an instance of resource type R_j has been allocated to process P_i .
 - A directed edge $R_j \rightarrow P_i$ is called an **assignment edge**.
- ❖ **Representation of graph:**
 - ✓ Represent each process P_i as a circle
 - ✓ Each resource type R_j as a rectangle.
 - ✓ Since resource type R_j may have more than one instance, represent each such instance as a dot within the rectangle.

- ✓ Note that a request edge points to only the rectangle R_j , whereas an assignment edge must also designate one of the dots in the rectangle.

| | |
|---|---|
| <ul style="list-style-type: none"> ▪ Process |  |
| <ul style="list-style-type: none"> ▪ Resource Type with 4 instances |  |
| <ul style="list-style-type: none"> ▪ P_i requests instance of R_j |  |
| <ul style="list-style-type: none"> ▪ P_i is holding an instance of R_j |  |

- ❖ The resource-allocation graph shown in next figure (Figure 7.1) depicts the following situation:

- ✓ The sets P, R, and E:
 - $P = \{P_1, P_2, P_3\}$
 - $R = \{R_1, R_2, R_3, R_4\}$
 - $E = \{P_1 \rightarrow R_1, P_2 \rightarrow R_3, R_1 \rightarrow P_2, R_2 \rightarrow P_2, R_2 \rightarrow P_1, R_3 \rightarrow P_3\}$
- ✓ Resource instances:
 - One instance of resource type R1
 - Two instances of resource type R2
 - One instance of resource type R3
 - Three instances of resource type R4
- ✓ Process states:
 - Process P1 is holding an instance of resource type R2 and is waiting for an instance of resource type R1.
 - Process P2 is holding an instance of R1 and an instance of R2 and is waiting for an instance of R3.
 - Process P3 is holding an instance of R3.

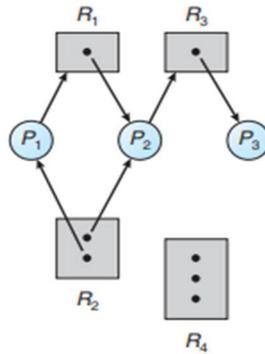


Figure 7.1 Resource-allocation graph.

❖ **Basic Facts:**

- ✓ If graph contains no cycles → no deadlock
- ✓ If graph contains a cycle →
 - if only one instance per resource type, then deadlock
 - if several instances per resource type, possibility of deadlock

❖ **Example of graph with a deadlock:** Two cycles exist in the next figure (Figure 7.2):

- ✓ $P1 \rightarrow R1 \rightarrow P2 \rightarrow R3 \rightarrow P3 \rightarrow R2 \rightarrow P1$
- ✓ $P2 \rightarrow R3 \rightarrow P3 \rightarrow R2 \rightarrow P2$

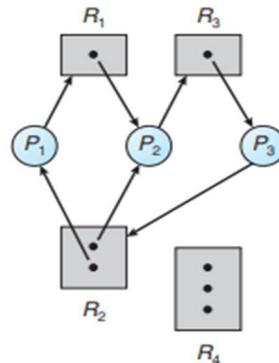


Figure 7.2 Resource-allocation graph with a deadlock.

❖ **Example of graph without a deadlock:** Consider the resource-allocation graph in next figure (Figure 7.3). also have a cycle:

- ✓ $P1 \rightarrow R1 \rightarrow P3 \rightarrow R2 \rightarrow P1$

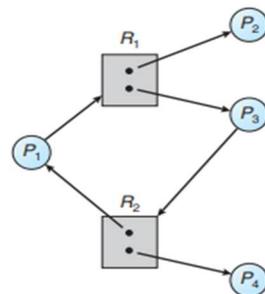


Figure 7.3 Resource-allocation graph with a cycle but no deadlock.

Methods for Handling Deadlocks

- ❖ Generally speaking, we can deal with the deadlock problem in one of three ways:
 - ✓ Can use a protocol to prevent or avoid deadlocks, ensuring that the system will **never** enter a deadlocked state.
 - Deadlock **prevention**
 - Deadlock **avoidance**
 - ✓ Can allow the system to enter a deadlocked state, **detect** it, and recover.
 - ✓ Can ignore the problem altogether and pretend that deadlocks never occur in the system (used by most operating systems, including UNIX).

Deadlock Prevention

- ❖ To occur a deadlock, there are four necessary conditions must hold. By ensuring that at least one of these conditions cannot hold, can prevent the occurrence of a deadlock.
 1. **Mutual Exclusion:** not required for sharable resources (e.g., read-only files); must hold for non-sharable resources
 2. **Hold and Wait:** must guarantee that whenever a process requests a resource, it does not hold any other resources
 - ✓ Require process to request and be allocated all its resources before it begins execution, or allow process to request resources only when the process has none allocated to it.
 - ✓ **Low resource utilization; starvation possible**
 3. **No Preemption:** If a process that is holding some resources requests another resource that cannot be immediately allocated to it, then all resources currently being held are released
 - ✓ Preempted resources are added to the list of resources for which the process is waiting
 - ✓ Process will be restarted only when it can regain its old resources, as well as the new ones that it is requesting
 4. **Circular Wait:** impose a total ordering of all resource types, and require that each process requests resources in an increasing order of enumeration

❖ Next figure (Figure 7.4) show deadlock example

```

/* thread_one runs in this function */
void *do_work_one(void *param)
{
    pthread_mutex_lock(&first_mutex);
    pthread_mutex_lock(&second_mutex);
    /**
     * Do some work
     */
    pthread_mutex_unlock(&second_mutex);
    pthread_mutex_unlock(&first_mutex);

    pthread_exit(0);
}

/* thread_two runs in this function */
void *do_work_two(void *param)
{
    pthread_mutex_lock(&second_mutex);
    pthread_mutex_lock(&first_mutex);
    /**
     * Do some work
     */
    pthread_mutex_unlock(&first_mutex);
    pthread_mutex_unlock(&second_mutex);

    pthread_exit(0);
}

```

Figure 7.4 Deadlock example

❖ Next figure (Figure 7.5) show deadlock example with lock ordering

```

void transaction(Account from, Account to, double amount)
{
    mutex lock1, lock2;
    lock1 = get_lock(from);
    lock2 = get_lock(to);

    acquire(lock1);
    acquire(lock2);

    withdraw(from, amount);
    deposit(to, amount);

    release(lock2);
    release(lock1);
}

```

Figure 7.5 Deadlock example with lock ordering.