

# MODULE 3

## CHAPTER 2 – DEAD LOCKS



Prepared By Mr. EBIN PM, AP

49

## DEADLOCK

- In a multiprogramming environment, several processes may compete for a finite number of resources.
- A process requests resources; if the resources are not available at that time, the process enters a wait state.
- Waiting processes may never again change state, because the resources they have requested are held by other waiting processes. This situation is called a **deadlock**.

Prepared By Mr. EBIN PM, AP

EDULINE

50

- A system consists of a finite number of resources to be distributed among a number of competing processes.
- The resources consist of some number of **identical instances**.
- **Memory space, CPU cycles, files, and I/O devices** (such as printers and tape drives) are examples of resource types.
- If a system has two CPUs, then the resource type CPU has two instances.
- Under the normal mode of operation, a process may utilize a resource in only the following sequence:
  - ❖ **Request:** If the request cannot be granted immediately (for example, the resource is being used by another process), then the requesting process must wait until it can acquire the resource.

Prepared By Mr. EBIN PM, AP

EDULINE

51

- ❖ **Use:** The process can operate on the resource (for example, if the resource is a printer, the process can print on the printer).
- ❖ **Release:** The process releases the resource.

Prepared By Mr. EBIN PM, AP

EDULINE

52

## CONDITIONS FOR DEADLOCK

➤ 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 non-sharable mode; that is, 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.

Prepared By Mr. EBIN PM, AP

EDULINE

53

3. **No preemption:** Resources cannot be preempted; that is, a resource can be released only voluntarily by the process holding it, after that process has completed its task.

4. **Circular wait:** A set  $\{P_0, P_1, \dots, P_n\}$  of waiting processes must exist such that  $P_0$  is waiting for a resource that is held by  $P_1$ ,  $P_1$  is waiting for a resource that is held by  $P_2$ , ...,  $P_{n-1}$  is waiting for a resource that is held by  $P_n$  and finally  $P_n$  is waiting for a resource that is held by  $P_0$ .

Prepared By Mr. EBIN PM, AP

EDULINE

54

## RESOURCE ALLOCATION GRAPH

- A tool for recognizing deadlock is —Resource allocation graph (RAG). This graph consists of a set of **vertices V** and a set of **edges E**. Vertices are partitioned in to **processes {P1, P2.....Pn}** and **resources {R1, R2.....Rm}**. Its edges are of two types:

❖ **Request edges:** This edge indicates the request of a process for acquiring a particular resource. It is a directed edge like  $P_i \rightarrow R_j$ .

❖ **Assignment edge:** This edge indicates the allocation of a resource to particular process. It is also a directed edge like  $R_j \rightarrow P_i$ . This edge here indicates that an instance of resource  $R_j$  has been allocated to the process  $P_i$ .


- **Circle** represents **processes** and **rectangles** represent **resources**.

Prepared By Mr. EBIN PM, AP

EDULINE

55

### Resource-allocation graph



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

Prepared By Mr. EBIN PM, AP

EDULINE

56

**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.
  - P3 is holding an instance of R3.
- If a **cycle does not exist in RAG**, then we can say that there is no deadlock exists.
- If a **cycle exists in a RAG**, there must be a **possibility** of deadlock.
- If a resource has only **one instance and also a cycle exist**, then **deadlock must occur**. In the case of multiple instances, the possibility of occurring deadlock is small.

Prepared By Mr. EBIN PM, AP

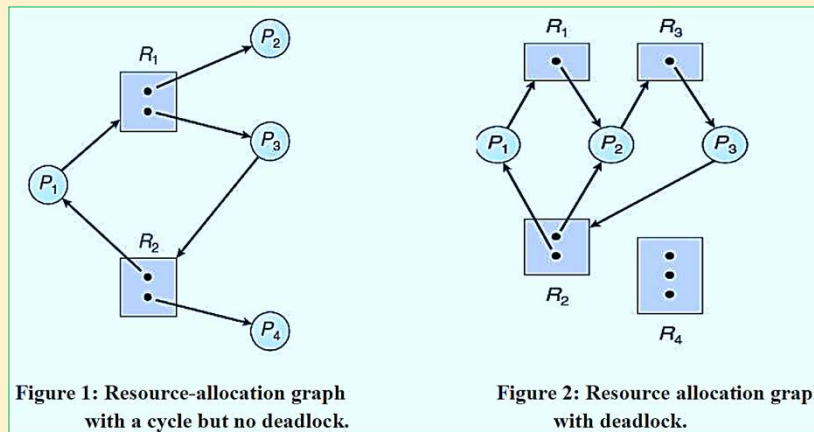
EDULINE

57

- Now consider the following resource-allocation graph. In this example, we also have a cycle:

**$P_1 \rightarrow R_1 \rightarrow P_3 \rightarrow R_2 \rightarrow P_1$**

However, there is no deadlock.



Prepared By Mr. EBIN PM, AP

EDULINE

58

- In figure1 ,P2 holds an instance of R1, and P4 holds an instance of R2. P2 and P4 can release their resources.
- If they do, then R1 and R2 will both have free instances, so there will be no deadlock, as those free instances can be assigned to P1 and P3, respectively, and the arrows will be reversed.
- The request edges will turn into assignment edges. Then the graph will be acyclic.

Multiple instances of resources	One instance of every resource
<ul style="list-style-type: none"> <li>• Cyclic RAG need not always imply deadlock</li> </ul>	<ul style="list-style-type: none"> <li>• Cyclic RAG implies deadlock</li> </ul>
<ul style="list-style-type: none"> <li>• If there is a cycle in the RAG, it does not necessarily imply that a deadlock has occurred. Even though there is a cycle in the RAG, there is no deadlock.</li> </ul>	<ul style="list-style-type: none"> <li>• If the cycle in RAG involves only a set of resource types, each of which has a single instance, then deadlock must have occurred.</li> </ul>

Prepared By Mr. EBIN PM, AP

EDULINE

59

## METHODS FOR HANDLING DEADLOCK

The four methods are:

- Prevention
- Avoidance
- Detection and recovery
- Ignore (used in UNIX, ignore and restart)

### ❖ DEADLOCK PREVENTION

For a deadlock to occur, each of the four necessary conditions (mutual exclusion, Hold and wait, No preemption, Circular wait) must hold. By ensuring that at least one of these conditions cannot hold, we can prevent the occurrence of a deadlock

Prepared By Mr. EBIN PM, AP

EDULINE

60

## 1. Mutual Exclusion

- **Make all the non-sharable resources sharable.**
- The mutual-exclusion condition must hold for non-sharable resources. For example, a printer cannot be simultaneously shared by several processes.
- Sharable resources, on the other hand, do not require mutually exclusive access, and thus cannot be involved in a deadlock.
- **Read-only files are a good example of a sharable resource.** We cannot prevent deadlocks by denying the mutual-exclusion condition because **some resources are intrinsically non-sharable** and mutual exclusion is one of the requirements of the critical section problems solution.

Prepared By Mr. EBIN PM, AP

EDULINE

61

## 2. Hold and Wait

- To prevent Hold and wait condition from happening, we can have a **rule** that says
- A process may not request a resource if it is holding another resource.
- So, to take the print out of the contents of a file, you first request the disk, and then you get it, use it and release it. Then you request the printer, you get it, use it, and then you release it.
- Thus, it implies that a process should have released all its resources before it requests for additional resources.
- **Disadvantage is starvation.**

Prepared By Mr. EBIN PM, AP

EDULINE

62

- Another rule can be there that “a process should request and acquire all the requested resources before its execution begins”.
- For example, to take the print out of the contents of a file, the disk and printer should be requested beforehand.
- **Disadvantage** of this method is **resource utilization will be low**.

Prepared By Mr. EBIN PM, AP

EDULINE

63

### 3. No Preemption

- To ensure that this condition does not happen, **preemption of the resources should be allowed**. Preemption can be added in two manners:

1. If a process is holding some resources and requests another resource that cannot be immediately allocated to it (that is, the process must wait), then all resources currently being held are preempted.

In other words, **these resources are implicitly released**. The preempted resources are added to the list of resources for which the process is waiting. **The process will be restarted only when it can regain its old resources, as well as the new ones that it is requesting.**

Prepared By Mr. EBIN PM, AP

EDULINE

64



2. Alternatively, if a process requests some resources, we first check whether they are available. If they are, we allocate them.

- If they are not available, we check whether they are allocated to some other process that is waiting for additional resources.
- If so, we preempt the desired resources from the waiting process and allocate them to the requesting process.

Prepared By Mr. EBIN PM, AP

EDULINE

65

#### 4. Circular Wait

- For avoiding circular wait we use **numbering scheme**. We can **assign a number for all resources**.
- So each process has requested in a particular order. If the set of resources types R includes tape drives, disk drives and printer, then the function F might be

**F(tape drive)=1;**

**F(disk drive)=5;**

**F(printer)=12;**

Prepared By Mr. EBIN PM, AP

EDULINE

66

- A process can be request any resource based on the condition  $f(R_j) > f(R_i)$ .
- ❖  $R_i$ : The resource number which is currently hold by the process  $P_i$ .
- ❖  $R_j$ : New request from process  $P_i$ .
- If a process holds the disk drive, it cannot be requested for a tape drive, because the number is less, but it can be give a request for printer, which has a greater number.

Prepared By Mr. EBIN PM, AP


EDULINE

67

## DEADLOCK AVOIDANCE

**Safe State:** A state is safe if the system can allocate resources to each process (up to its maximum) in some order and still avoid a deadlock. More formally, a system is in a safe state only if there exists a **safe sequence**. A **deadlock state** is a form of **unsafe state**.

**Safe, unsafe, and deadlocked state spaces**



Prepared By Mr. EBIN PM, AP

EDULINE

68

### ❖ Resource-Allocation Graph Algorithm


- This algorithm is used for only **single instance resource type**.
- In addition to the request and assignment edge, a new type of edge called a **claim edge** is introduced.
- A claim edge  $P_i \text{ ----} \rightarrow R_j$  indicates that process  $P_i$  may request resource  $R_j$  at some time in the future.
- Claim edge is represented by a dashed line.
- When process  $P_i$  requests resource  $R_j$ , the claim edge  $P_i \text{ ----} \rightarrow R_j$  is converted to a request edge.
- Similarly, when a resource  $R_j$  is released by  $P_i$  the assignment edge  $R_j \rightarrow P_i$  is reconverted to a claim edge  $P_i \text{ ----} \rightarrow R_j$

Prepared By Mr. EBIN PM, AP

EDULINE

69

- Consider the resource allocation graph of fig(a). Suppose that  $P_2$  requests  $R_2$ . Although  $R_2$  is currently free, we cannot allocate it to  $P_2$ , Since this allocation will create a cycle in the graph (fig(b)).
- A cycle indicates that the system is in an unsafe state. If  $P_1$  requests  $R_2$ , and  $P_2$  requests  $R_1$ , then a deadlock will occur.



Prepared By Mr. EBIN PM, AP

EDULINE

70

### ❖ Banker's Algorithm

- This algorithm's purpose is deadlock avoidance.
- Several data structures must be maintained to implement the banker's algorithm.
  - **Available:** A vector of length m indicates the number of available resources of each type.
  - **Max:** The maximum demand of each process.
  - **Allocation:** An n x m matrix defines the number of resources of each type currently allocated to each process.
  - **Need:** An n x m matrix indicates the remaining resource need of each process.  $\text{Need}[i,j] = \text{Max}[i,j] - \text{Allocation}[i,j]$

Prepared By Mr. EBIN PM, AP

EDULINE

71

### Example:

	Allocation			Max			Available			Need		
	A	B	C	A	B	C	A	B	C	A	B	C
P0	0	1	0	7	5	3	3	3	2			
P1	2	0	0	3	2	2						
P2	3	0	2	9	0	2						
P3	2	1	1	2	2	2						
P4	0	0	2	4	3	3						

- First we find Need.

$$\text{Need} = \text{Max} - \text{Allocation}$$

- Now compare need and available. First we take process P0. It is not selected because A need 7 but Available of A is 3. Second we take process P1. It is selected. Then the new

Prepared By Mr. EBIN PM, AP

EDULINE

72

Available = (Available + Allocation)

$$= (3 \ 3 \ 2) + (2 \ 0 \ 0) = (5 \ 3 \ 2)$$

- Next we select P2. Compare its need and the new available. ie, between (6 0 0) and (5 3 2)
- Here A wants 6 but only 5 is available. So it is rejected.
- Next take P3 and compare its need and available. ie, between (0 1 1) and (5 3 2). It is accepted.

$$\text{Then new Available} = (5 \ 3 \ 2) + (2 \ 1 \ 1) = (7 \ 4 \ 3)$$

- Next take P4 and compare with its need and available. ie, between (4 3 1) and (7 4 3). It is accepted.

$$\text{Then the new Available} = (7 \ 4 \ 3) + (0 \ 0 \ 2) = (7 \ 4 \ 5)$$

Prepared By Mr. EBIN PM, AP

EDULINE

73

- Next we take P0 and compare with its need and available. ie, between (7 4 3) and (7 4 5). It is accepted.

$$\text{Then new Available} = (7 \ 4 \ 5) + (0 \ 1 \ 0) = (7 \ 5 \ 5)$$

- Now we can take P2 and compare with its need and new available. ie, between (6 0 0) and (7 5 5).
- It is accepted. Then the new final available = (7 5 5) + (3 0 2) = (10, 5, 7)
- So the safe sequence is **P1, P3, P4, P0, P2.**

Prepared By Mr. EBIN PM, AP

EDULINE

74

### ❖ DEADLOCK DETECTION

- If a system does not employ either a deadlock-prevention or a deadlock avoidance algorithm, then a deadlock situation may occur. In this environment, the system must provide:
- An algorithm that examines the state of the system to determine whether a deadlock has occurred
  - An algorithm to recover from the deadlock
  - Detection-and-recovery scheme requires overhead that includes not only the run-time costs of maintaining the necessary information and executing the detection algorithm, but also the potential losses inherent in recovering from a deadlock.

Prepared By Mr. EBIN PM, AP

EDULINE

75

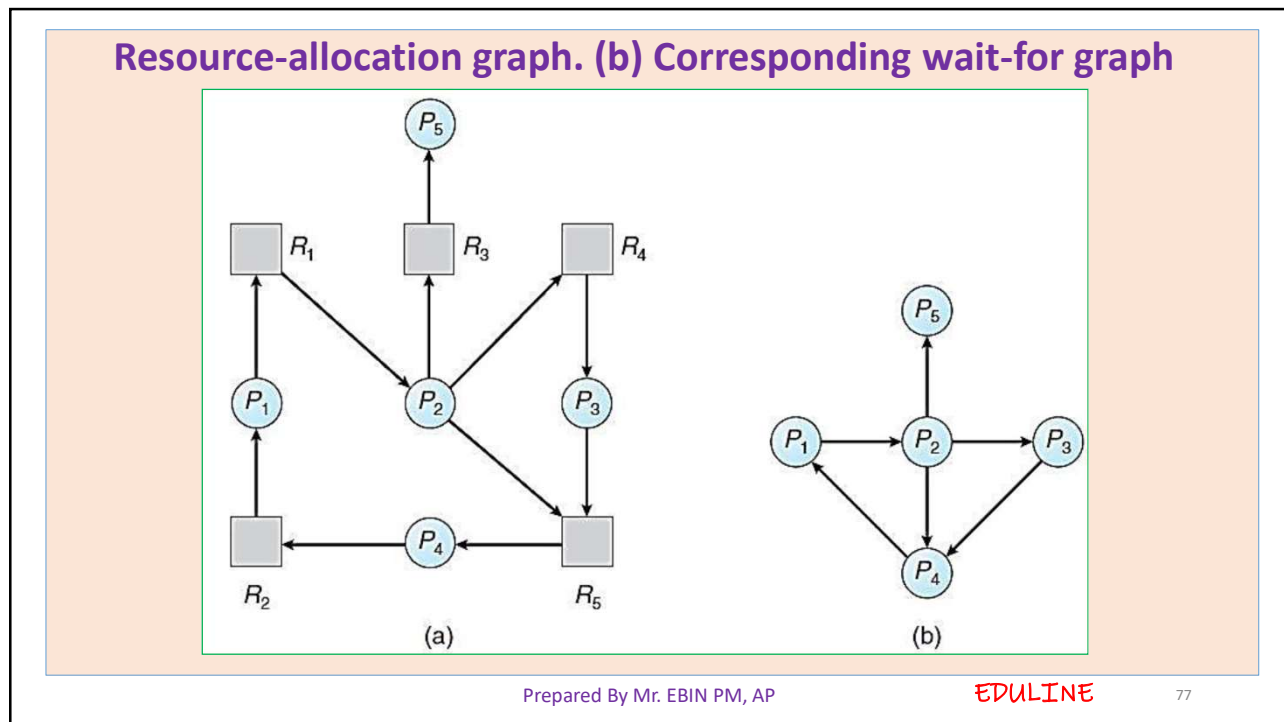
### ➤ Detection in Single Instance of Each Resource Type

- If all resources have only a single instance, then we can define a deadlock detection algorithm that uses a variant of the resource-allocation graph, called a **wait-for graph**.
- We obtain this graph from the resource-allocation graph by removing the nodes of type resource and collapsing the appropriate edges.
- If  $P_i \rightarrow R_j$  and  $R_j \rightarrow P_j$ , then  $P_i \rightarrow P_j$ .
- A deadlock exists in the system iff the wait-for graph contains a cycle. To detect deadlocks, the system needs to maintain the wait-for graph and periodically to invoke an algorithm that searches for a cycle in the graph.

Prepared By Mr. EBIN PM, AP

EDULINE

76



### ❖ RECOVERY FROM DEADLOCK

- Recovery can be handled by manually or automatically.
- There are **two options** for breaking a deadlock. One solution is simply to abort one or more processes to break the circular wait. The second option is to preempt some resources from one or more of the deadlocked processes.

#### ➤ Process Termination

- Here all processes are terminated which is in the deadlock state. Two methods for the process termination methods are:
  - ✓ **Abort all deadlocked processes:** This method clearly will break the deadlock cycle, but at a great expense. If 5 processes are existed, these 5 processes must be terminated. So the used resources are wasted.

✓ **Abort one process at a time until the deadlock cycle is eliminated:**

This method incurs considerable overhead, since, after each process is aborted, a deadlock-detection algorithm must be invoked to determine whether any processes are still deadlocked.

- Aborting a process may not be easy. If the process was in the midst of updating a file, terminating it will leave that file in an incorrect state. Many factors may determine which process is chosen, including:

1. What the priority of the process is
2. How long the process has computed, and how much longer the process will compute before completing its designated task
3. How many and what type of resources the process has used (for example, whether the resources are simple to preempt)

Prepared By Mr. EBIN PM, AP

EDULINE

79

4. How many more resources the process needs in order to complete

5. How many processes will need to be terminated?

6. Whether the process is interactive or batch

➤ **Resource Preemption**

- To eliminate deadlocks using resource preemption, we successively preempt some resources from processes and give these resources to other processes until the deadlock cycle is broken. The three conditions are:

1. **Selecting a victim:** Which resources and which processes are to be preempted?

we must determine the order of preemption to minimize cost.

Prepared By Mr. EBIN PM, AP

EDULINE

80



**2. Rollback:** If we preempt a resource from a process, it cannot continue with its normal execution; it is missing some needed resource. We must roll back the process to some safe state, and restart it from that state. Since, it is difficult to determine what a safe state is; the simplest solution is a total rollback: Abort the process and then restart it.

**3. Starvation:** In a system where victim selection is based primarily on cost factors, it may happen that the same process is always picked as a victim. As a result, this process never completes its designated task, a starvation occurred. Clearly, we must ensure that a process can be picked as a victim only a (small) finite number of times. The most common solution is to include the number of rollbacks in the cost factor.