

For example, if  $X = (1,7,3,2)$  and  $Y = (0,3,2,1)$ , then  $Y \leq X$ . In addition,  $Y < X$  if  $Y \leq X$  and  $Y \neq X$ .

We can treat each row in the matrices *Allocation* and *Need* as vectors and refer to them as *Allocation<sub>i</sub>* and *Need<sub>i</sub>*. The vector *Allocation<sub>i</sub>* specifies the resources currently allocated to process  $P_i$ ; the vector *Need<sub>i</sub>* specifies the additional resources that process  $P_i$  may still request to complete its task.

### 6.5.3.1. Safety Algorithm

We can now present the algorithm for finding out whether or not a system is in a safe state. This algorithm can be described as follows:

1. Let *Work* and *Finish* be vectors of length  $m$  and  $n$ , respectively. Initialize *Work* = *Available* and *Finish*[ $i$ ] = *false* for  $i = 0, 1, \dots, n - 1$ .

2. Find an index  $i$  such that both

a. *Finish*[ $i$ ] == *false*

b. *Need<sub>i</sub>*  $\leq$  *Work*

If no such  $i$  exists, go to step 4.

3. *Work* = *Work* + *Allocation<sub>i</sub>*

*Finish*[ $i$ ] = *true*

Go to step 2.

4. If *Finish*[ $i$ ] == *true* for all  $i$ , then the system is in a safe state.

This algorithm may require an order of  $m \times n^2$  operations to determine whether a state is safe.

### 6.5.3.2. Resource-Request Algorithm

Next, we describe the algorithm for determining whether requests can be safely granted. Let *Request<sub>i</sub>* be the request vector for process  $P_i$ .

If *Request<sub>i</sub>* [ $j$ ] ==  $k$ , then

process  $P_i$  wants  $k$  instances of resource type  $R_j$ . When a request for resources is made by process  $P_i$ , the following actions are taken:

1. If  $Request_i \leq Need_i$  , go to step 2. Otherwise, raise an error condition, since the process has exceeded its maximum claim.
2. If  $Request_i \leq Available$ , go to step 3. Otherwise,  $P_i$  must wait, since the resources are not available.
3. Have the system pretend to have allocated the requested resources to process  $P_i$  by modifying the state as follows:

$$Available = Available - Request_i ;$$

$$Allocation_i = Allocation_i + Request_i ;$$

$$Need_i = Need_i - Request_i ;$$

If the resulting resource-allocation state is safe, the transaction is completed, and process  $P_i$  is allocated its resources. However, if the new state is unsafe, then  $P_i$  must wait for  $Request_i$  , and the old resource-allocation state is restored.

**6.5.3.3. An Illustrative Example**

To illustrate the use of the banker’s algorithm, consider a system with five processes  $P_0$  through  $P_4$  and three resource types  $A$ ,  $B$ , and  $C$ . Resource type  $A$  has ten instances, resource type  $B$  has five instances, and resource type  $C$  has seven instances. Suppose that, at time  $T_0$ , the following snapshot of the system has been taken:

	Allocation	Max	Available
	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	

The content of the matrix *Need* is defined to be *Max – Allocation* and is as follows:

	<i>Need</i>		
	<i>A</i>	<i>B</i>	<i>C</i>
<i>P0</i>	7	4	3
<i>P1</i>	1	2	2
<i>P2</i>	6	0	0
<i>P3</i>	0	1	1
<i>P4</i>	4	3	1

We claim that the system is currently in a safe state. Indeed, the sequence  $\langle P1, P3, P4, P2, P0 \rangle$  satisfies the safety criteria. Suppose now that process *P1* requests one additional instance of resource type *A* and two instances of resource type *C*, so  $Request1 = (1,0,2)$ . To decide whether this request can be immediately granted, we first check that  $Request1 \leq Available$ —that is, that  $(1,0,2) \leq (3,3,2)$ , which is true. We then pretend that this request has been fulfilled, and we arrive at the following new state:

	<i>Allocation</i>	<i>Need</i>	<i>Available</i>
	<i>A B C</i>	<i>A B C</i>	<i>A B C</i>
<i>P0</i>	0 1 0	7 4 3	2 3 0
<i>P1</i>	3 0 2	0 2 0	
<i>P2</i>	3 0 2	6 0 0	
<i>P3</i>	2 1 1	0 1 1	
<i>P4</i>	0 0 2	4 3 1	

We must determine whether this new system state is safe. To do so, we execute our safety algorithm and find that the sequence  $\langle P1, P3, P4, P0, P2 \rangle$  satisfies the safety requirement. Hence, we can immediately grant the request of process *P1*.

You should be able to see, however, that when the system is in this state, a request for (3,3,0) by *P4* cannot be granted, since the resources are not available.

Furthermore, a request for (0,2,0) by  $P_0$  cannot be granted, even though the resources are available, since the resulting state is unsafe.

We leave it as a programming exercise for students to implement the banker's algorithm.

## 6.6. 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 may provide:

- An algorithm that examines the state of the system to determine whether a deadlock has occurred
- An algorithm to recover from the deadlock

In the following discussion, we elaborate on these two requirements as they pertain to systems with only a single instance of each resource type, as well as to systems with several instances of each resource type. At this point, however, we note that a 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.

### 6.6.1. 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 resource nodes and collapsing the appropriate edges. More precisely, an edge from  $P_i$  to  $P_j$  in a wait-for graph implies that process  $P_i$  is waiting for process  $P_j$  to release a resource that  $P_i$  needs. An edge  $P_i \rightarrow P_j$  exists in a wait-for graph if and only if the corresponding resource allocation graph contains two edges  $P_i \rightarrow R_q$  and  $R_q \rightarrow P_j$  for some resource  $R_q$ . In Figure 6.7, we present a resource-allocation graph and the corresponding wait-for graph. As before, a deadlock exists in the system if and only if the wait-for graph contains a cycle. To detect deadlocks, the system needs to *maintain* the wait-for graph and periodically *invoke an algorithm* that searches for a