

CS370 Operating Systems

Colorado State University

Yashwant K Malaiya

Fall 2024 L13

Synchronization, Deadlocks



Slides based on

- Text by Silberschatz, Galvin, Gagne
- Various sources

Synchronization Notes

- Producer-consumer with bounded buffer
 - The production and consumption rates may not match.
 - Circular buffer helps.
- Readers-Writers Problem
 - Allow multiple readers to read at the same time
 - Semaphores for mutual exclusion (mutex) and counting
- Synchronization approaches
 - Machine instructions \Rightarrow semaphores \Rightarrow monitor
- Monitor: Implements
 - **mutual exclusion**: only one process may be active at a time
 - **Conditions** with associated queues where processes *wait* until *notified*
 - Our Monitor discussion is generic. Self Exercise for a **Java** example.

Course Notes

- HW3 Due **Oct 3 Th**
 - Must have a working program 2 days earlier.
 - Autograder now working
- Project D1: Was due **Sept 26. Being graded.**
- Midterm: Respondus lockdown browser on laptop. A practice quiz is available.
 - Sec 001 On-campus: **Tues Oct 8** in-class
 - Sec 801 Online: **Wed Oct 9** 12:10 AM-11:50 PM
- D2 progress report: **Oct 31, 2024**

Resuming Processes within a Monitor: Priority

- If several processes queued on condition x , and $x.\text{signal}()$ is executed, which should be resumed?
- FCFS frequently not adequate
- **conditional-wait** construct of the form $x.\text{wait}(c)$
 - Where c is **priority number**
 - Process with lowest number (highest priority) is scheduled next

Single Resource allocation

- Allocate a single resource among competing processes using priority numbers that specify the maximum time a process plans to use the resource

```
R.acquire(t);  
    ...  
access the resource;  
    ...  
R.release;
```

- Where R is an instance of type **ResourceAllocator**
- **A monitor based solution next.**

A Monitor to Allocate Single Resource

```
monitor ResourceAllocator
```

```
{
```

```
    boolean busy;
```

```
    condition x;
```

```
void acquire(int time) {
```

```
    if (busy)
```

```
        x.wait(time);
```

```
    busy = TRUE;
```

```
}
```

```
void release() {
```

```
    busy = FALSE;
```

```
    x.signal();
```

```
}
```

```
    initialization code() {
```

```
        busy = FALSE;
```

```
    }
```

```
}
```

Sleep, Time used
to prioritize
waiting
processes

Wakes up
one of the
processes


Java Synchronization

- For simple synchronization, Java provides the `synchronized` keyword
 - synchronizing methods

```
public synchronized void increment( ) { c++; }
```
 - synchronizing blocks

```
synchronized(this) {  
    lastName = name;  
    nameCount++;  
}
```
- `wait()` and `notify()` allows a thread to wait for an event. A call to `notifyAll()` allows all threads that are on `wait()` with the same lock to be notified.
- `notify()` notifies one thread from a pool of identical threads, `notifyAll()` when threads have different purposes
- For more sophisticated locking mechanisms, starting from Java 5, the package `java.concurrent.locks` provides additional capabilities.

Java Synchronization



Each object automatically has a monitor (mutex) associated with it

- When a method is synchronized, the runtime must obtain the lock on the object's monitor before execution of that method begins (and must release the lock before control returns to the calling code)

`wait()` and `notify()` allows a thread to wait for an event.

- **`wait()`**: Causes the current thread to wait until another thread invokes the `notify()` method or the `notifyAll()` method for this object.
- **`notify()`**: Wakes up a single thread that is waiting on this object's monitor. If any threads are waiting on this object, one of them is chosen to be awakened.
- A call to **`notifyall()`** allows all threads that are on `wait()` with the same lock to be released, they will run in sequence according to priority.

Java Synchronization: Dining Philosophers

```
public synchronized void pickup(int i)
    throws InterruptedException {
    setState(i, State.HUNGRY);
    test(i);
    while (state[i] != State.EATING) {
        this.wait();
        // Recheck condition in loop,
        // since we might have been notified
        // when we were still hungry
    }
}
```

```
public synchronized void putdown(int i) {
    setState(i, State.THINKING);
    test(right(i));
    test(left(i));
}
```

```
private synchronized void test(int i) {
    if (state[left(i)] != State.EATING &&
        state[right(i)] != State.EATING &&
        state[i] == State.HUNGRY)
    {
        setState(i, State.EATING);
        // Wake up all waiting threads
        this.notifyAll();
    }
}
```

Synchronization Examples

- Solaris
- Windows
- Linux
- Pthreads

Solaris Synchronization

- Implements a variety of locks to support multitasking, multithreading (including real-time threads), and multiprocessing
- Uses **adaptive mutexes** for efficiency when protecting data from short code segments
 - Starts as a standard semaphore spin-lock
 - If lock held, and by a thread running on another CPU, spins
 - If lock held by non-run-state thread, block and sleep waiting for signal of lock being released
- Uses **condition variables**
- Uses **readers-writers** locks when longer sections of code need access to data
- Uses **turnstiles** to order the list of threads waiting to acquire either an adaptive mutex or reader-writer lock
 - Turnstiles are per-lock-holding-thread, not per-object
- Priority-inheritance per-turnstile gives the running thread the highest of the priorities of the threads in its turnstile

Windows Synchronization

- Uses interrupt masks to protect access to global resources on uniprocessor systems
- Uses **spinlocks** on multiprocessor systems
 - Spinlocking-thread will never be preempted
- Also provides **dispatcher objects** user-land which may act mutexes, semaphores, events, and timers
 - **Events**
 - An event acts much like a condition variable
 - Timers notify one or more thread when time expired
 - Dispatcher objects either **signaled-state** (object available) or **non-signaled state** (thread will block)

Linux Synchronization

- Linux:
 - Prior to kernel Version 2.6, disables interrupts to implement short critical sections
 - Version 2.6 and later, fully preemptive
- Linux provides:
 - Semaphores
 - atomic operations on integers
 - Spinlocks
 - reader-writer versions of both
- On single-cpu system, spinlocks replaced by enabling and disabling kernel preemption

Pthreads Synchronization

- Pthreads API is OS-independent
- It provides:
 - mutex locks
 - condition variable thus can be used to create a monitor
- Non-portable extensions include:
 - read-write locks
 - Spinlocks
- [A simple example](#)

Alternative Approaches

- Transactional Memory
- OpenMP
- Functional Programming Languages

Transactional Memory

- A **memory transaction** is a sequence of read-write operations to memory that are performed atomically without the use of locks.

```
void update() {  
    atomic{  
        /* modify shared data*/  
    }  
}
```

May be implemented by hardware or software.

OpenMP

- OpenMP is a set of compiler directives and API that support parallel programming.

```
void update(int value)
{
    #pragma omp critical
    {
        count += value
    }
}
```

The code contained within the `#pragma omp critical` directive is treated as a critical section and performed atomically.

CS370 Operating Systems

Colorado State University

Yashwant K Malaiya

Deadlocks



Slides based on

- Text by Silberschatz, Galvin, Gagne
- Various sources

Chapter 8: Deadlocks

- System Model
- Deadlock Characterization
- Methods for Handling Deadlocks
 - Deadlock Prevention
 - Deadlock Avoidance resource-allocation
 - Deadlock Detection
 - Recovery from Deadlock

System Model

- System consists of resources
- Resource types R_1, R_2, \dots, R_m
 - Resource may be CPU cycles, memory space, I/O devices, critical sections*
- Each resource type R_i has W_i instances.
- Each process utilizes a resource as follows:
 - **request**
 - **use**
 - **release**

Deadlock Characterization

Deadlock **can** arise if four conditions hold **simultaneously**.

- **Mutual exclusion**: only one process at a time can use a resource
- **Hold and wait**: a process holding at least one resource is waiting to acquire additional resources held by other processes
- **No preemption**: a resource can be released only voluntarily by the process holding it, after that process has completed its task
- **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 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 P_n is waiting for a resource that is held by P_0 .

Deadlock with Mutex Locks

- Deadlocks can occur via system calls, locking, etc.
- See example
 - Dining Philosophers: each get the right chopstick first
 - we saw this example earlier

Let s and q be two semaphores initialized to 1

P_0	P_1
<code>wait(S);</code>	<code>wait(Q);</code>
<code>wait(Q);</code>	<code>wait(S);</code>
...	...
<code>signal(Q);</code>	<code>signal(S);</code>
<code>signal(S);</code>	<code>signal(Q);</code>

P0 executes wait(s), P1 executes wait(Q)

P0 must wait till P1 executes signal(Q)

P1 must wait till P0 executes signal(S) Deadlock!

Resource-Allocation Graph


A set of vertices V and a set of edges E .

- V is partitioned into two types:
 - $P = \{P_1, P_2, \dots, P_n\}$, the set consisting of all the processes in the system
 - $R = \{R_1, R_2, \dots, R_m\}$, the set consisting of all resource types in the system
- **request edge** – directed edge $P_i \rightarrow R_j$
- **assignment edge** – directed edge $R_j \rightarrow P_i$



Resource-Allocation Graph (Cont.)


- Process




- Resource Type with 4 instances



- P_i requests instance of R_j




- P_i is holding an instance of R_j



Example of a Resource Allocation Graph

P1 holds an instance of R2 and is requesting R1 ..



Does a deadlock exist here?

P3 will eventually be done with R3, letting P2 use it.


Thus, P2 will be eventually done, releasing R1. ...

Answer: No.

Observation: If the graph contains no cycles, then no process in the system is deadlocked.

If the graph does contain a cycle, then a deadlock *may* exist.

Resource Allocation Graph With A Deadlock



Does a deadlock exist?

At this point, two minimal cycles exist in the system:

$P_1 \rightarrow R_1 \rightarrow P_2 \rightarrow R_3 \rightarrow P_3 \rightarrow R_2 \rightarrow P_1$

$P_2 \rightarrow R_3 \rightarrow P_3 \rightarrow R_2 \rightarrow P_2$

Processes P_1 , P_2 , and P_3 are deadlocked.


Graph With A Cycle But No Deadlock

Is there a deadlock?

P_4 will release its instance of resource type R_2 . That resource can then be allocated to P_3 , breaking the cycle. Thus, there is no deadlock.

If a resource-allocation graph does not have a cycle, then the system is **not** in a deadlocked state.

If there is a cycle, then the system may or may not be in a deadlocked state.



Basic Facts

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

Methods for Handling Deadlocks

- Ensure that the system will *never* enter a deadlock state:
 - Deadlock prevention
 - ensuring that at least one of the 4 conditions cannot hold
 - Deadlock avoidance
 - Dynamically examines the resource-allocation state to ensure that it will never enter an unsafe state, and thus there can never be a circular-wait condition
- Allow the system to enter a deadlock state
 - Detection: detect and then recover. Hope is that it happens rarely.
- Ignore the problem and pretend that deadlocks never occur in the system; used by *most* operating systems, including UNIX. However..

Methods for Handling Deadlocks

- **Deterministic:** Ensure that the system will *never* enter a deadlock state at any cost
- **Probabilistic view:** Hope it happens rarely.
Handle if it happens: Allow the system to enter a deadlock state and then recover.

Methods for Handling Deadlocks

Approach	Resource allocation policy	Scheme	Notes
Prevention	Conservative, undercommits resources	Requesting all resources at once	Good for processes with a single burst of activity
		Preemption	Good when preemption cost is small
		Resource ordering	Compile time enforcement possible
Avoidance	midway	Find at least one safe path (dynamic)	Future max requirement must be known
Detection	Liberal	Invoked periodically	Preemption may be needed

Ostrich algorithm

Ostrich algorithm: Stick your head in the sand; pretend there is no problem at all .



Advantages:

- Cheaper, rarely needed anyway
- Prevention, avoidance, detection and recovery
 - Need to run constantly

Disadvantages:

- Resources held by processes that cannot run
- More and more processes enter deadlocked state
 - When they request more resources
- Deterioration in system performance
 - Requires restart

To be fair to the ostriches,
let me say that ...

Deadlock Prevention

For a deadlock to occur, each of the four necessary conditions must hold. By ensuring that at least one of these conditions cannot hold, we can **prevent** the occurrence of a deadlock.

- **Mutual exclusion:** only one process at a time can use a resource
- **Hold and wait:** a process holding at least one resource is waiting to acquire additional resources held by other processes
- **No preemption:** a resource can be released only voluntarily by the process holding it, after that process has completed its task
- **Circular wait:** there exists a set $\{P_0, P_1, \dots, P_n\}$ of waiting processes that are circularly waiting.



Deadlock Prevention: Limit Mutual Exclusion

For a deadlock to occur, each of the four necessary conditions must hold. By ensuring that at least one of these conditions cannot hold, we can **prevent** the occurrence of a deadlock.

Restrain the ways request can be made:

- **Limit Mutual Exclusion** –
 - not required for sharable resources (e.g., read-only files)
 - (Mutual Exclusion must hold for non-sharable resources)

Deadlock Prevention: Limit Hold and Wait

- **Limit Hold and Wait** – must guarantee that whenever a process requests a resource, it does not hold any other resources
 1. Require process to request and be allocated all its resources before it begins execution
 2. Allow a process to request resources when it is holding none.
Ex: Copy data from DVD, sort file, and print
 - First request DVD and disk file
 - Then request file and printer,
 - then start
- Disadvantage: starvation possible

Deadlock Prevention: Limit No Preemption

- **Limit 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

Deadlock Prevention: Limit Circular Wait

- **Limit Circular Wait** – impose a total ordering of all resource types, and require that each process requests resources in an increasing order of enumeration
- Assign each resource a unique number
 - Disk drive: 1
 - Printer: 2 ...
 - Request resources in increasing order
 - *Example soon*

Dining philosophers problem: Necessary conditions **for** deadlock

Relax conditions to
avoid deadlock

- Mutual exclusion
 - 2 philosophers *cannot share* the same chopstick
- Hold-and-wait
 - A philosopher *picks up one* chopstick at a time
 - Will not let go of the first while it *waits for the second* one
- No preemption
 - A philosopher *does not snatch chopsticks* held by some other philosopher
- Circular wait
 - Could happen if each philosopher *picks chopstick with the same hand* first

Deadlock Example: numbering

```
/* 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);
}
```

Allows deadlock. Redesign to avoid.

Assume that thread one is the first to acquire the locks and does so in the order (1) first mutex, (2) second mutex.

Solution: **Lock-order verifier**
“**Witness**” records the relationship that **first mutex must be acquired before second mutex**. If thread two later acquires the locks out of order, witness generates a warning message on the system console.

Deadlock may happen *even* 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);
}
```

Lock ordering:
First *from* lock, then *to* lock

Ex: Transactions 1 and 2 execute concurrently.

Transaction 1 transfers \$25 from account A to account B, and

Transaction 2 transfers \$50 from account B to account A.

Deadlock is possible, even with lock ordering.

Deadlock Avoidance

Manage resource allocation to ensure the system never enters an unsafe state.

Deadlock Avoidance

Requires that the system has some additional *a priori* information available

- Simplest and most useful model requires that each process declare the *maximum number* of resources of each type that it may need
- The deadlock-avoidance algorithm dynamically examines the resource-allocation state to ensure that there can never be a circular-wait condition
- Resource-allocation *state* is defined by the number of available and allocated resources, and the maximum demands of the processes

Deadlock Avoidance: Handling resource requests

- For each resource request:
 - Decide whether or not process should wait
 - To avoid possible future deadlock
- Predicated on:
 1. Currently available resources
 2. Currently allocated resources
 3. *Future requests and releases of each process*

Avoidance: amount and type of information needed

- **Resource allocation state**
 - Number of available and allocated resources
 - Maximum demands of processes
- *Dynamically* examine resource allocation state
 - Ensure circular-wait cannot exist
- Simplest model:
 - Declare maximum number of resources for each type
 - Use information to avoid deadlock

Safe Sequence

System must decide if immediate allocation leaves the system in a safe state


System is in **safe state** if there exists a **sequence** $\langle P_1, P_2, \dots, P_n \rangle$ of ALL the processes such that

- for each P_i , the resources that P_i can still request can be satisfied by
 - currently available resources +
 - resources held by all the P_j , with $j < i$
 - That is
 - If P_i resource needs are not immediately available, then P_i can wait until all P_j have finished and released resources
 - When P_i terminates, P_{i+1} can obtain its needed resources, and so on
- If no such sequence exists: system state is **unsafe**

Deadlock avoidance: Safe states

- If the system can:
 - Allocate resources to each process in some order
 - Up to the maximum for the process
 - Still avoid deadlock
 - Then it is in a **safe state**
- A system is safe **ONLY IF** there is a safe sequence
- A safe state is not a deadlocked state
 - Deadlocked state is an unsafe state
 - Not all unsafe states are deadlock

Safe, Unsafe, Deadlock State



Examples of safe and unsafe states in next 3 slides

Example A: Assume 12 Units in the system

	Max need	Current holding
av		3
P0	10	5
P1	4	2
P2	9	2

At time T0 (shown):

9 units allocated

3 (12-9) units available

*A unit could be a drive,
a block of memory etc.*


- Is the system at time **T0** in a safe state?
 - Try sequence $\langle P1, P0, P2 \rangle$
 - P1 can be given 2 units
 - When P1 releases its resources; there are now 5 available units
 - P0 uses 5 and subsequently releases them (10 available now)
 - P2 can then proceed.
- Thus $\langle P1, P0, P2 \rangle$ is a safe sequence, and at T0 system was in a safe state

More detailed look 

Colorado State University

Example A: Assume 12 Units in the system (timing)

Is the state at T0 safe? Detailed look for instants T0, T1, T2, etc..



	Max need	Current holding	+2 allo to P1	P1 releases all
		T0	T1	T2	T3	T4	T5
av		3	1	5	0	10	3
P0	10	5	5	5	10 done	0	0
P1	4	2	4 done	0	0	0	0
P2	9	2	2	2	2	2	9 done

Thus the state at T0 is safe.

Example B: 12 Units initially available in the system

	Max need	T0	T1 safe?
Av		3	2
P0	10	5	5
P1	4	2	2
P2	9	2	3 Is that OK?

Before T1:

3 units available

At T1:

2 units available

- At time **T1**, P2 is allocated 1 more units. Is that a good decision?
 - Now only P1 can proceed (already has 2, and given be given 2 more)
 - When P1 releases its resources; there are 4 units
 - P0 needs 5 more, P2 needs 6 more. Deadlock.
 - **Mistake** in granting P2 the additional unit.
- The state at **T1** is not a safe state. Wasn't a good decision.

Avoidance Algorithms

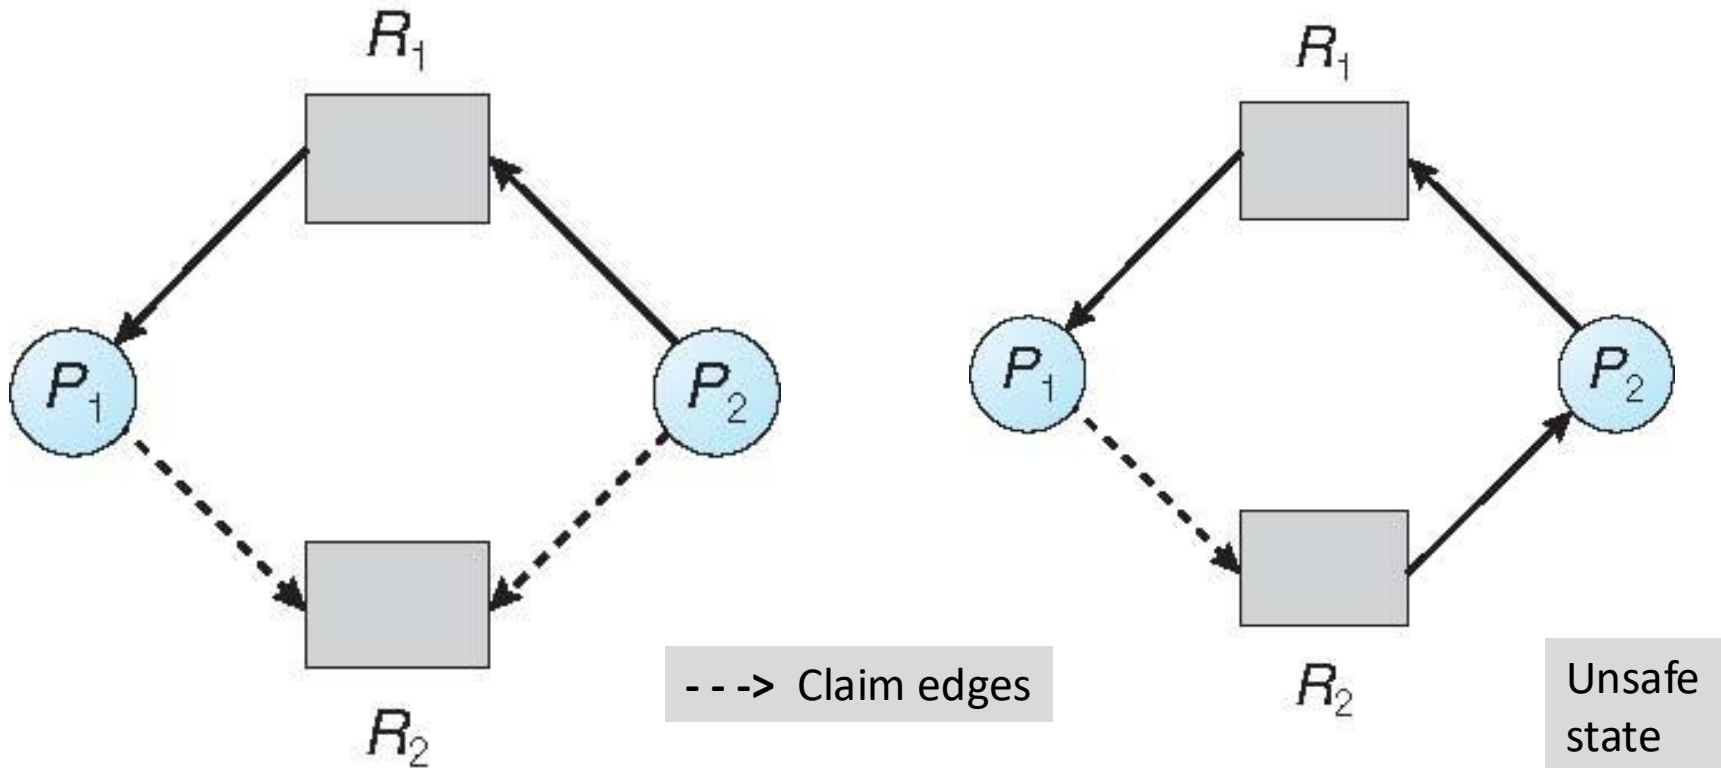
- Dynamic
- Single instance of a resource type
 - Use a resource-allocation graph scheme
- Multiple instances of a resource type
 - Use the banker's algorithm (Dijkstra)



Resource-Allocation Graph Scheme

- **Claim edge** $P_i \rightarrow R_j$ indicated that process P_i may request resource R_j ; represented by a **dashed line**. This is new.
- Claim edge converts to **request edge** when a process **requests** a resource
- Request edge converted to an **assignment edge** when the resource is **allocated** to the process
- When a resource is released by a process, assignment edge reconverts to a claim edge
- Requirement: Resources must be claimed *a priori* in the system

Resource-Allocation Graph



Suppose P_2 requests R_2 . Can R_2 be allocated to P_2 ?

Although R_2 is currently free, we cannot allocate it to P_2 , since this action will create a cycle getting system in an unsafe state.

If P_1 requests R_2 , and P_2 requests R_1 , then a deadlock will occur.

Resource-Allocation Graph Algorithm

- Suppose that process P_i requests a resource R_j
- The request can be granted only if converting the request edge to an assignment edge does not result in the formation of a cycle in the resource allocation graph

Banker's Algorithm: examining a request

- Multiple instances of resources.
- Each process must a priori claim maximum use
- When a process requests a resource,
 - it may have to wait until the resource becomes available ([resource request algorithm](#))
 - Request should not be granted if the resulting system state is unsafe ([safety algorithm](#))
- When a process gets all its resources it must return them in a finite amount of time
- Modeled after a banker in a small-town making loans.

Data Structures for the Banker's Algorithm

Let n = number of processes, and m = number of resources types.

- **Available:** Vector of length m . If available $[j] = k$, there are k instances of resource type R_j available

Processes vs resources:

- **Max:** $n \times m$ matrix. If $Max [i,j] = k$, then process P_i may request at most k instances of resource type R_j
- **Allocation:** $n \times m$ matrix. If $Allocation[i,j] = k$ then P_i is currently allocated k instances of R_j
- **Need:** $n \times m$ matrix. If $Need[i,j] = k$, then P_i may need k more instances of R_j to complete its task

$$Need [i,j] = Max[i,j] - Allocation [i,j]$$

Safety Algorithm: Is this a safe state?

1. Let **Work** and **Finish** be vectors of length m and n , respectively. Initialize:

Work = Initially Available resources

Finish [i] = initially false for $i = 0, 1, \dots, n-1$ (processes done)

2. Find a process i such that both:

(a) **Finish** [i] = false

(b) **Need** _{i} ≤ **Work**

If no such i exists, go to step 4

3. **Work** = **Work** + **Allocation** _{i}

Finish [i] = true

go to step 2

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

n = number of processes,
m = number of resources types
Need _{i} : additional res needed
Work: res currently free
Finish _{i} : processes finished
Allocation _{i} : allocated to i

Resource-Request Algorithm for Process P_i

Notation: $Request_i$ = request vector for process P_i .

If $Request_i[j] = k$ then process P_i wants k instances of resource type R_j

Algorithm: *Should the allocation request be granted?*

1. If $Request_i \leq Need_i$ go to step 2. Otherwise, raise **error condition**, since process has exceeded its maximum claim
2. If $Request_i \leq Available$, go to step 3. Otherwise P_i must wait, since resources are **not available**
3. **Is allocation safe?:** **Pretend** to allocate requested resources to 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 safe \Rightarrow the resources are allocated to P_i Use safety algorithm here
- If unsafe $\Rightarrow P_i$ must wait, and the old resource-allocation state is preserved.

Example 1A: Banker's Algorithm

- 5 processes P_0 through P_4 ;
- 3 resource types: A (10 instances), B (5 instances), and C (7 instances)
- Is it a safe state?

The Need matrix is redundant

Process	Max			Allocation			Need		
type	A	B	C	A	B	C	A	B	C
Currently available				3	3	2			
P0	7	5	3	0	1	0	7	4	3
P1	3	2	2	2	0	0	1	2	2
P2	9	0	2	3	0	2	6	0	0
P3	2	2	2	2	1	1	0	1	1
P4	4	3	3	0	0	2	4	3	1