

CISC 7310X

C10a Deadlock and Resource Allocation Graph

Hui Chen

Department of Computer & Information Science

CUNY Brooklyn College

Acknowledgement

- These slides are a revision of the slides provided by the authors of the textbook via the publisher of the textbook

Outline

- System Model
- Deadlock Characterization (Necessary Conditions)
- Resource Allocation Graph
- Deadlock in Multithreaded Applications
- Overview of Methods for Handling Deadlocks

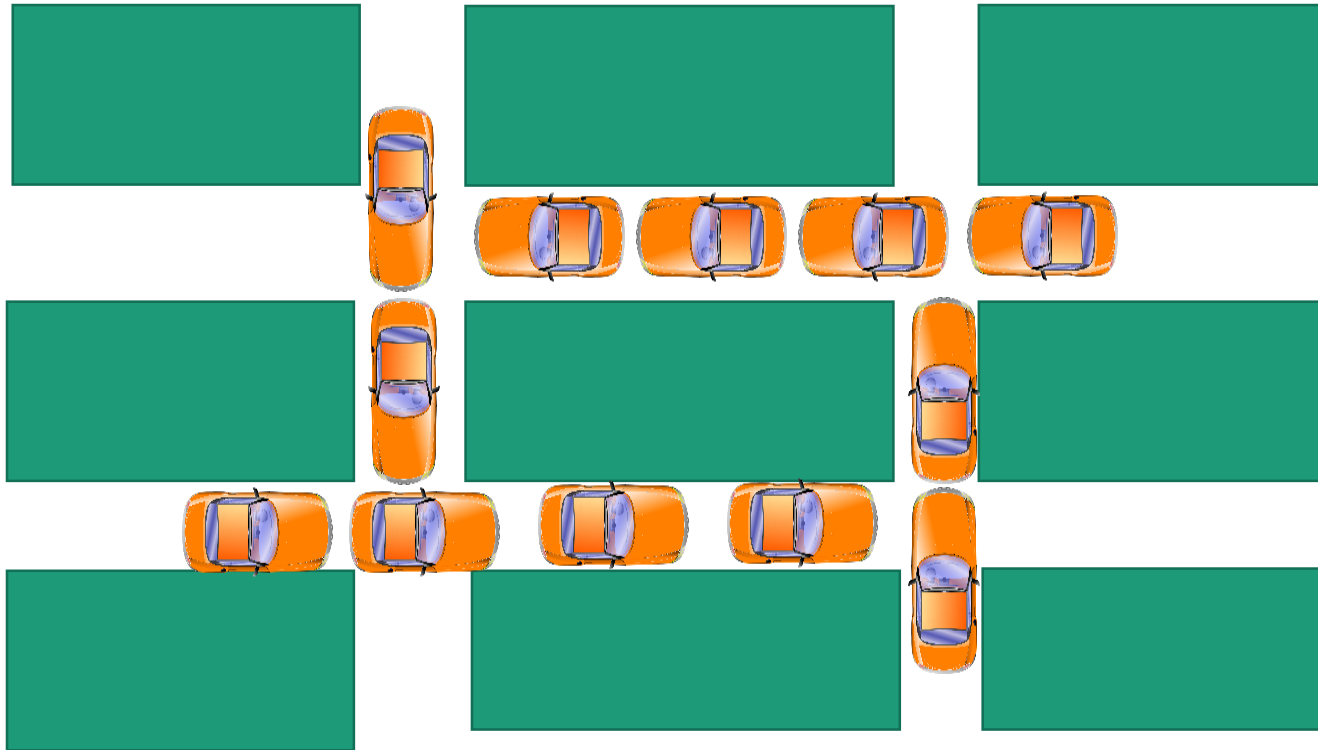
- Deadlock Prevention
- Deadlock Avoidance
- Deadlock Detection
- Recovery from Deadlock

Problem when Sharing Resources

- A proposed legislature in the history by Kansas (Botkin and Harlow, 1953)
 - "When two trains approach each other at a crossing, both shall come to a full stop and neither shall start up again until the other has gone."



This can also happen ...



System Model

- System consists of resources
- Resource types R_1, R_2, \dots, R_m
 - Examples
 - CPU cycles, memory space, I/O devices
- Each resource type R_i has W_i instances.
- A set of processes, and each process utilizes a resource as follows:
 - request
 - use
 - release

Deadlock

- Every process in the set is waiting for an event to be triggered by another in the set (request or release resource)

Deadlock Characterization

- Deadlock can arise if four conditions hold simultaneously. (the 4 necessary conditions for deadlocks)
- **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 .

Questions?

- Concept of deadlock
- Necessary conditions of 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 (drawn in ovals)
 - $R = \{R_1, R_2, \dots, R_m\}$, the set consisting of all resource types in the system (drawn in rectangles)
- **request edge** - directed edge $P_i \rightarrow R_j$
 - P_i requests or waits for R_j
- **assignment edge** - directed edge $R_j \rightarrow P_i$
 - R_j is assigned to or is held by P_i

Resource-Allocation Graph: Example 1


- Can you describe the graphs in English? (Hint: oval: process; rectangle: resource; arrow: Resource \rightarrow Process, Process \rightarrow Resource, i.e., is being held/assigned to or requests by/waiting for)



(a)



(b)



(c)


- Resource allocation graphs. (a) Holding a resource. (b) Requesting a resource. (c) Deadlock. [Figure 6-3 in Tanenbaum & Bos, 2014]

Questions?

- Concept of resource allocation graph
- Examples of simple resource allocation graph
 - Each type of resources has only a single instance
- What if a type of resource has multiple instances?


Resource with Multiple Instances

- A type of resource may have multiple instances
- Notations




Resource Allocation Graph: Example 2

- Can you draw the resource allocation graph for the following scenario?
 - One instance of R1
 - Two instances of R2
 - One instance of R3
 - Three instance of R4
 - T1 holds one instance of R2 and is waiting for an instance of R1
 - T2 holds one instance of R1, one instance of R2, and is waiting for an instance of R3
 - T3 is holds one instance of R3




Is There a Dead Lock?

- Mutual exclusion?
- Hold and wait?
- No preemption?
- Circular wait?




Resource Allocation Graph: Example 3

- Can you draw the resource allocation graph for the following scenario?
 - One instance of R1
 - Two instances of R2
 - One instance of R3
 - Three instance of R4
 - T1 holds one instance of R2 and is waiting for an instance of R1
 - T2 holds one instance of R1, one instance of R2, and is waiting for an instance of R3
 - T3 is holds one instance of R3, and is waiting for an instance of R2




Is There a Dead Lock?

- Mutual exclusion?
- Hold and wait?
- No preemption?
- Circular wait?




Resource Allocation Graph: Example 4

- Can you draw the resource allocation graph for the following scenario?
 - Two instances of R1
 - Two instances of R2
 - T1 holds one instance of R2 and is waiting for an instance of R1
 - T2 holds one instance of R1
 - T3 holds one instance of R1 and is waiting for an instance of R2
 - T4 is waiting for an instance of R2



Is There a Dead Lock?

- Mutual exclusion?
- Hold and wait?
- No preemption?
- Circular wait?



Determine Existence of Deadlocks

- 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

Resource Allocation Graph: Example 5

- What's the resource allocation graph?
 - 2 processes, P1 and P2 share two 2 CD-RW drives (D1, D2)
 - P1 is using D1, P2 is using D2
 - P1 requests D2 before releasing D1; P2 requests D1 before releasing D2
- Is there a deadlock?

Questions?

- Resource allocation graph
- Determine existence of deadlock using resource allocation graph

Resource-Allocation Graph and Scheduling: Example

- Three processes: A, B, C
- Three resources: R, S, T
- Each process requests and release schedule in the sequence below:

A
Request R
Request S
Release R
Release S

(a)

B
Request S
Request T
Release S
Release T

(b)

C
Request T
Request R
Release T
Release R

(c)

Schedule with Deadlock

A
Request R
Request S
Release R
Release S

(a)

B
Request S
Request T
Release S
Release T


(b)

C
Request T
Request R
Release T
Release R


(c)

1. A requests R
 2. B requests S
 3. C requests T
 4. A requests S
 5. B requests T
 6. C requests R
- deadlock


(d)




(e)




(f)




(g)



(h)



(i)



(j)

Schedule without Deadlock

A
Request R
Request S
Release R
Release S

(a)

B
Request S
Request T
Release S
Release T


(b)

C
Request T
Request R
Release T
Release R


(c)

1. A requests R
2. C requests T
3. A requests S
4. C requests R
5. A releases R
6. A releases S
no deadlock


(k)




(l)




(m)




(n)



(o)



(p)



(q)

Semaphores or Mutexes are Resources

- Access non-preemptive resource with semaphore (request, use, release)
 - down/signal/P; up/wait/V

```
typedef int semaphore;  
semaphore resource_1;
```

```
void process_A(void) {  
    down(&resource_1);  
    use_resource_1( );  
    up(&resource_1);  
}
```

(a)

```
typedef int semaphore;  
semaphore resource_1;  
semaphore resource_2;
```

```
void process_A(void) {  
    down(&resource_1);  
    down(&resource_2);  
    use_both_resources( );  
    up(&resource_2);  
    up(&resource_1);  
}
```

(b)

- [Figure 6-1 in Tanenbaum & Bos, 2014]

Deadlock in Multithreaded Application

- Two mutex locks are created and initialized:

```
pthread_mutex_t first_mutex;  
pthread_mutex_t second_mutex;
```

```
pthread_mutex_init(&first_mutex, NULL);  
pthread_mutex_init(&second_mutex, NULL);
```

```

/* 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);
}

```

Is Deadlock Possible in the Example?


- Deadlock is possible.
- If thread 1 acquires `first_mutex` and thread 2 acquires `second_mutex`, Thread 1 then waits for `second_mutex` and thread 2 waits for `first_mutex`.
- Can be illustrated with a **resource allocation graph**:

Illustration using Resource Allocation Graph



Resource-Allocation Graph: Example 6

- Describe the following resource allocation graph?



Resource Allocation is Subtle

Deadlock free

```
typedef int semaphore;
semaphore resource_1;
semaphore resource_2;

void process_A(void) {
    down(&resource_1);
    down(&resource_2);
    use_both_resources( );
    up(&resource_2);
    up(&resource_1);
}

void process_B(void) {
    down(&resource_1);
    down(&resource_2);
    use_both_resources( );
    up(&resource_2);
    up(&resource_1);
}
```

Deadlock

```
semaphore resource_1;
semaphore resource_2;

void process_A(void) {
    down(&resource_1);
    down(&resource_2);
    use_both_resources( );
    up(&resource_2);
    up(&resource_1);
}

void process_B(void) {
    down(&resource_2);
    down(&resource_1);
    use_both_resources( );
    up(&resource_1);
    up(&resource_2);
}
```

- ^(a) [Figure 6-2 in Tanenbaum & Bos, 2014] ^(b)

Remarks

- Whether the deadlock happens or not depends on the result of a race (or scheduling)
 - Difficult to debug because it only happens sporadically
- Difference between deadlock free and deadlocked code is subtle in coding style

Questions?

- Synchronization tools are resources
- Subtle to write deadlock-free code, and difficult to debug

Methods for Handling Deadlocks

- Ensure that the system will *never* enter a deadlock state:
 - Deadlock prevention (by structurally negating one of the four required conditions)
 - Deadlock avoidance (by carefully allocating resources)
- Allow the system to enter a deadlock state and then recover
 - Deadlock detection and recovery (Let deadlocks occur, detect them, and then take action)
- Ignore the problem and pretend that deadlocks never occur in the system.
 - The [Ostrich algorithm](#)

The Ostrich Algorithm

The deadlock in my system happens once in a blue moon ...



Questions?

- System Model
- Deadlock in Multithreaded Applications
- Deadlock Characterization and Resource Allocation Graph
- Methods for Handling Deadlocks
- The Ostrich Algorithm