

Deadlocks

— DM510 Operating Systems

— Lars Rohwedder



Windows



macOS



iOS

Disclaimer

These slides contain (modified) content and media from the official Operating System Concepts slides:
<https://www.os-book.com/OS10/slide-dir/index.html>

Today's lecture

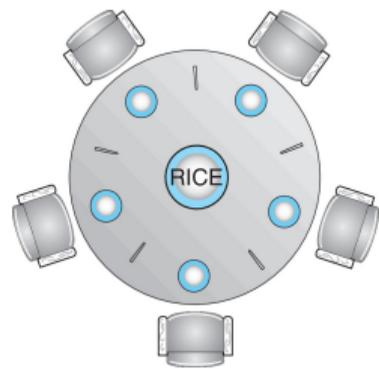
- > Chapter 8 of course book

The Problem

Dining Philosophers

- > 5 philosophers alternately think and eat
- > To eat they need to pick up their left and right chopstick (one at a time)
- > Chopsticks (implemented as mutexes) are shared with the neighbors

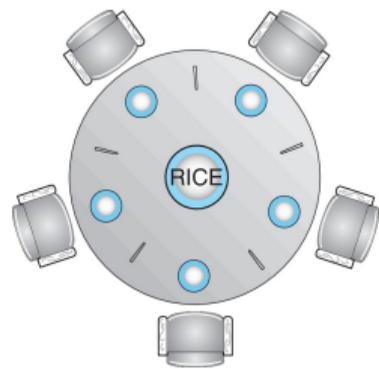
```
/* Algorithm for philosopher i */  
while (true) {  
    acquire(&chopstick[ i ]);  
    acquire(&chopstick[ ( i +1)%5]);  
    eat ();  
    release(&chopstick[ i ]);  
    release(&chopstick[ ( i +1)%5]);  
    think ();  
}
```



Dining Philosophers

- > 5 philosophers alternately think and eat
- > To eat they need to pick up their left and right chopstick (one at a time)
- > Chopsticks (implemented as mutexes) are shared with the neighbors

```
/* Algorithm for philosopher i */  
while (true) {  
    aquire(&chopstick[ i ]);  
    aquire(&chopstick[ ( i +1)%5]);  
    eat ();  
    release(&chopstick[ i ]);  
    release(&chopstick[ ( i +1)%5]);  
    think ();  
}
```



Deadlock

If each philosopher grabs the left chopstick before their neighbor grabs the right one, then they are stuck in a **deadlock**! Deadlocks easily occur in many other situations

Formal model

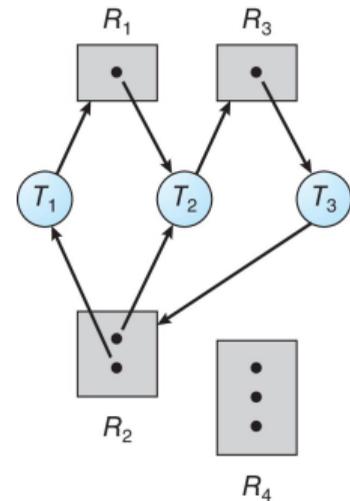
- > Deadlocks can occur with mutexes (last lecture), files, limited resources, etc.
- > Since the core problem is always the same, we consider it in the following abstract model

Nodes

- > Resources R_1, R_2, \dots, R_m with one or more **instances**
- > Threads T_1, T_2, \dots, T_n
- > **Mutual exclusion**: only one thread can hold the same instance at any time

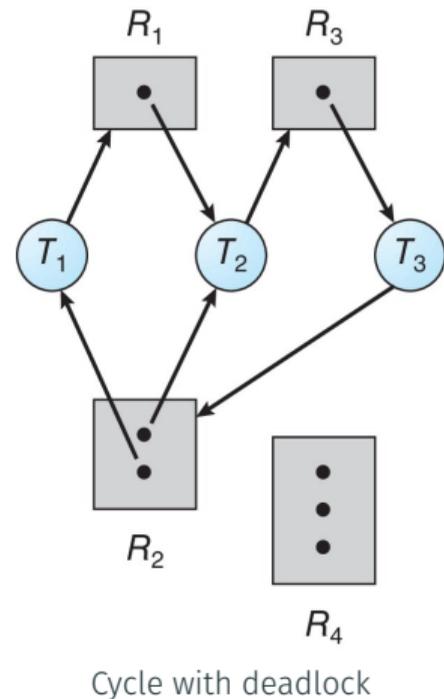
Edges

- > **request edge**: from thread to resource
- > **assignment edge**: from resource instance to thread



Conditions for deadlock

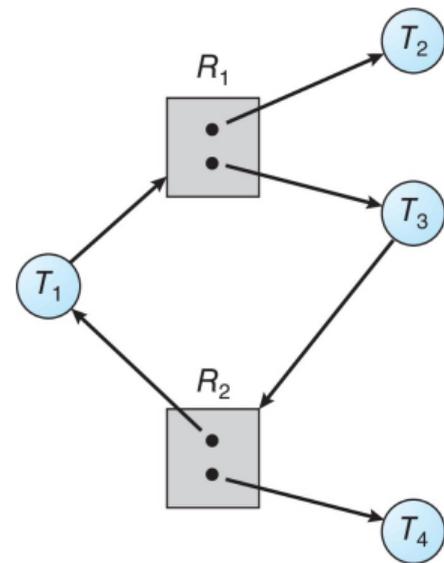
- > **Mutual exclusion** resource instances can only be held by one thread at a time
- > **Hold and wait:** a thread holds one resource instance while waiting for other resources
- > **No preemption:** a resource can only be released voluntarily
- > **Circular wait:** Threads can be ordered as T_1, \dots, T_n such that T_i waits for a resource that T_{i+1} (or T_1 if $i = n$) holds for each $i = 1, 2, \dots, n$
- > The conditions above are necessary (if not fulfilled \Rightarrow no deadlock)



Conditions for deadlock

- > **Mutual exclusion** resource instances can only be held by one thread at a time
- > **Hold and wait:** a thread holds one resource instance while waiting for other resources
- > **No preemption:** a resource can only be released voluntarily
- > **Circular wait:** Threads can be ordered as T_1, \dots, T_n such that T_i waits for a resource that T_{i+1} (or T_1 if $i = n$) holds for each $i = 1, 2, \dots, n$

- > The conditions above are necessary (if not fulfilled \Rightarrow no deadlock)
- > But not sufficient (if fulfilled \Rightarrow maybe deadlock)



Cycle without deadlock

Strategies for handling deadlocks

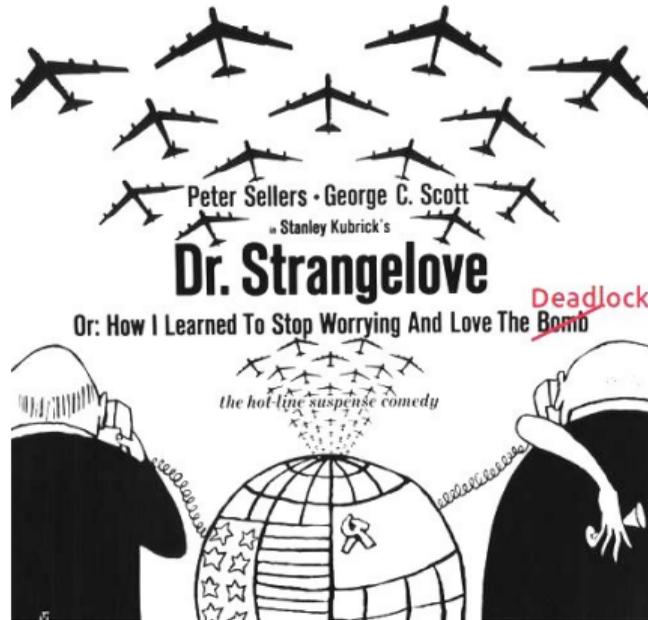
- > Ensure that system never enters deadlocked state by **deadlock prevention** or **deadlock avoidance**

Strategies for handling deadlocks

- > Ensure that system never enters deadlocked state by **deadlock prevention** or **deadlock avoidance**
- > Allow system to enter deadlocked state and perform **deadlock recovery**

Strategies for handling deadlocks

- > Ensure that system never enters deadlocked state by **deadlock prevention** or **deadlock avoidance**
- > Allow system to enter deadlocked state and perform **deadlock recovery**
- > Ignore that there can be deadlocks



Deadlock Prevention

Deadlock prevention

It is enough to remove one of the four necessary conditions for a deadlock

Deadlock prevention

It is enough to remove one of the four necessary conditions for a deadlock

Mutual exclusion

- > Make resource sharable if possible, for example, opening files read-only when possible to safely share them

Deadlock prevention

It is enough to remove one of the four necessary conditions for a deadlock

Mutual exclusion

- > Make resource sharable if possible, for example, opening files read-only when possible to safely share them

Hold and wait

- > Make threads request all resources before execution
- > Make threads request resources only when none are assigned to them
- > **Disadvantage:** lower resource utilization, possible starvation

Deadlock prevention

It is enough to remove one of the four necessary conditions for a deadlock

Mutual exclusion

- > Make resource sharable if possible, for example, opening files read-only when possible to safely share them

Hold and wait

- > Make threads request all resources before execution
- > Make threads request resources only when none are assigned to them
- > **Disadvantage:** lower resource utilization, possible starvation

No preemption

- > Thread waiting for some resource releases the ones it is currently holding
- > Thread is woken up once the new resource and the previously released ones are allocated to it
- > **Disadvantage:** often not applicable, for example: when mutexes are used to avoid race conditions we should not preempt

Deadlock prevention

It is enough to remove one of the four necessary conditions for a deadlock

Mutual exclusion

- > Make resource sharable if possible, for example, opening files read-only when possible to safely share them

Hold and wait

- > Make threads request all resources before execution
- > Make threads request resources only when none are assigned to them
- > **Disadvantage:** lower resource utilization, possible starvation

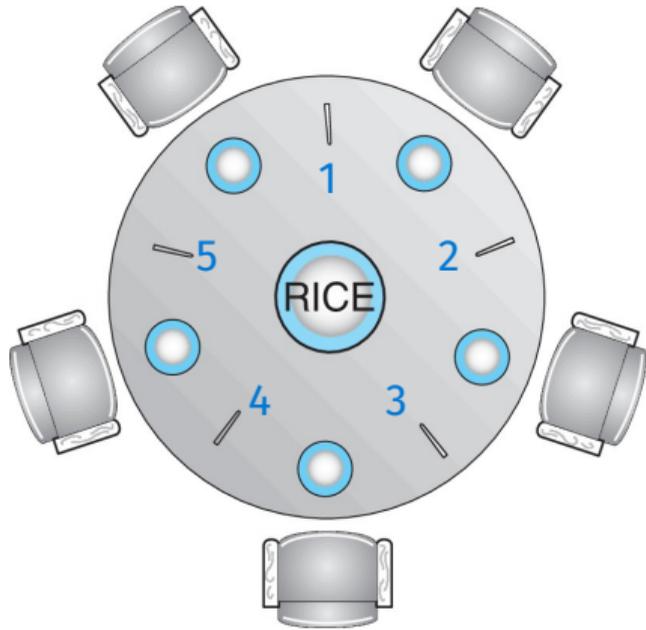
No preemption

- > Thread waiting for some resource releases the ones it is currently holding
- > Thread is woken up once the new resource and the previously released ones are allocated to it
- > **Disadvantage:** often not applicable, for example: when mutexes are used to avoid race conditions we should not preempt

Circular wait

- > Define total order on resources and require threads to request resources in that order

Example: total order



Solution for dining philosophers

```
/* philosopher i=1,2,3,4 */  
while (true) {  
    aquire(&chopstick[i]);  
    aquire(&chopstick[i+1]);  
    eat();  
    release(&chopstick[i]);  
    release(&chopstick[i+1]);  
    think();  
}  
/* for philosopher 5 */  
while (true) {  
    aquire(&chopstick[1]);  
    aquire(&chopstick[5]);  
    eat();  
    release(&chopstick[1]);  
    release(&chopstick[5]);  
    think();  
}
```

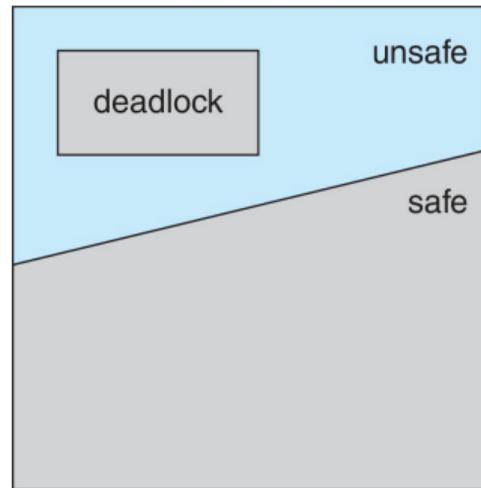
Deadlock Avoidance

Approach

- > **Simplifying assumptions:** we know the resources (and maximum no. instances) a thread can request and no new threads arrive
- > We design an algorithm that grants requested resources to threads in such a way that an unsafe state (= potential deadlock) can never be reached

Safe state

- > Threads can be reordered as T_1, T_2, \dots, T_n such that for each thread T_i the resource that T_1, T_2, \dots, T_{i-1} hold suffice to satisfy T_i 's additional resource need
- > No deadlock can occur because T_1 can continue; once T_1 is done, T_2 can continue; once T_1, T_2 are done, T_3 can continue; etc.



Detecting (un-)safe states

Input

- > **int avail[m]**: no. instances of resource not allocated
 - > **int max[n*m]**: maximum no. instances of resource a thread might request
 - > **int alloc[n * m]**: no. instances of resource allocated to thread
-
- > The algorithm on the right will modify **avail**. It should be called with a copy if **avail** is used outside the algorithm
 - > The idea is that we simulate how threads may finish

initialize **int need[n * m]** with

$$\text{need}[i, j] = \text{max}[i, j] - \text{alloc}[i, j]$$

initialize **int finish[n]** with

$$\text{finish}[i] = \text{false}$$

while true

> let **int i** $\in \{1, 2, \dots, n\}$ with

$$\text{finish}[i] = \text{false} \text{ and } \text{need}[i, j] \leq \text{avail}[j] \forall j$$

> if no such *i* exists: break

> **finish[i]** \leftarrow true

> for **int j** $\in \{1, 2, \dots, m\}$

$$\text{avail}[j] \leftarrow \text{avail}[j] + \text{alloc}[i, j]$$

Result: **safe** if **finish[i]** = true for all threads T_i

Banker's algorithm for deadlock avoidance

Setting: A thread T_i requests some resources and we should check if the request can be granted without entering into an unsafe state

Input

- > **int avail[m]**: no. instances of resource not allocated
- > **int max[n * m]**: maximum no. instances of resource a thread might request
- > **int alloc[n * m]**: no. instances of resource allocated to thread
- > Thread T_i that requests **int req[m]** of additional resources

- > if $\text{alloc}[i, j] + \text{req}[j] > \text{max}[i, j]$ for some j : raise runtime error
- > if $\text{req}[j] > \text{avail}[j]$ for some j : return **declined**
- > store current state in S
- > for each resource R_j :
 - /* give requested resources to T_i */*
 - $\text{avail}[j] \leftarrow \text{avail}[j] - \text{req}[j]$
 - $\text{alloc}[i, j] \leftarrow \text{alloc}[i, j] + \text{req}[j]$
- > check if state is safe
- > if unsafe state detected: restore state S and return **declined**
- > else: return **granted**

If request is declined, try again after resources have been released

Deadlock Recovery

Deadlock detection

Input

- > **int avail[m]**: no. instances of resource not allocated
 - > **int alloc[n * m]**: no. instances of resource allocated to a thread
 - > **int req[n * m]**: no. additional instances of resource requested by a thread
-
- > Periodically check for deadlock
 - > This is by simulating how threads may finish similar to previous unsafe state detection (but with subtle differences)
 - > Algorithm (on the right) requires $O(mn^2)$ operations

initialize **int finish[n]** with

$$\text{finish}[i] = \begin{cases} \text{true} & \text{if } \text{alloc}[i, j] = 0 \forall j \\ \text{false} & \text{otherwise.} \end{cases}$$

while true

- > let **int i** $\in \{1, 2, \dots, n\}$ with

$$\text{finish}[i] = \text{false and } \text{req}[i, j] \leq \text{avail}[j] \forall j$$

- > if no such *i* exists: break
- > **finish[i]** \leftarrow true
- > for **int j** $\in \{1, 2, \dots, m\}$:

$$\text{avail}[j] \leftarrow \text{avail}[j] + \text{alloc}[i, j]$$

Result: every thread *i* with **finish[i]** = false is in deadlock

Recovering: process termination

Terminate threads (processes) to remove deadlock. Variants:

- > Abort all threads (processes) in a deadlock
- > Abort one thread (process) at a time until deadlock is removed

Order for aborting threads

Examples:

- > By priority
- > By how long the thread has been computing or how much longer it needs
- > By resource usage
- > By resource requirements to complete
- > By number of threads that need to be terminated

Recovering: resource preemption

- > pick a victim thread
- > roll back thread to safe state and restart from there
- > can lead to **starvation** if same thread is picked over and over again, to avoid: include number of times picked in victim selection