

# Advanced Synchronization and Deadlock

# A house of cards?

- ④ Locks + CV/signal a great way to regulate access to a single shared object...
- ④ ...but general multi-threaded programs touch multiple shared objects
- ④ How can we atomically modify multiple objects to maintain
  - Safety: prevent applications from seeing inconsistent states
  - Liveness: avoid deadlock
    - a cycle of threads forever stuck waiting for one another

# Deadlock

- ⊙ A cycle of waiting among a set of threads, where each thread is waiting for some other thread in the cycle to take some action


```
Producer1() {  
  emptyBuffer.acquire()  
  producerMutexLock.acquire()  
  :  
}
```

```
Producer2() {  
  producerMutexLock.acquire()  
  emptyBuffer.acquire()  
  :  
}
```

# Deadlock

- ⊙ A cycle of waiting among a set of threads, where each thread is waiting for some other thread in the cycle to take some action

Mutually recursive locking




```
lock1.acquire()
...
lock2.acquire()
while (must wait) {
    cv.wait(&lock2)
}
...
lock2.release()
...
lock1.release()
```

```
lock1.acquire()
...
lock2.acquire()
...
cv.signal()
lock2.release()
...
lock1.release()
```

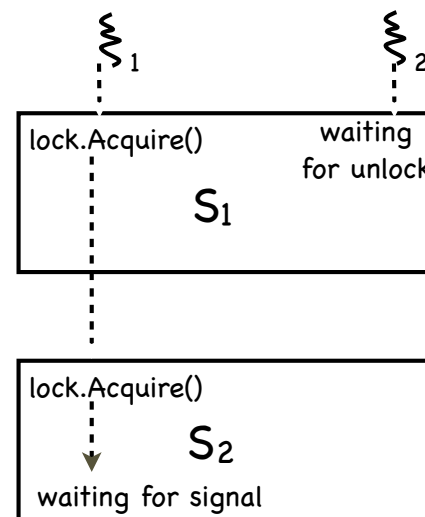
# Deadlock

- ⊙ A cycle of waiting among a set of threads, where each thread is waiting for some other thread in the cycle to take some action

Mutually recursive locking



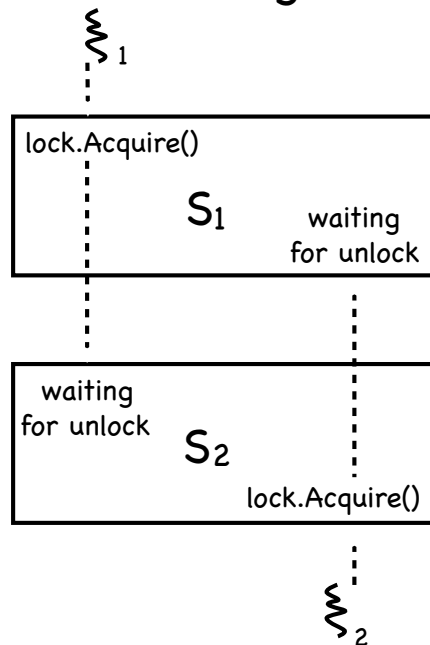
Nested waiting




# Deadlock

- ⊙ A cycle of waiting among a set of threads, where each thread is waiting for some other thread in the cycle to take some action


Mutually recursive locking



Nested waiting



# Dining Philosophers



- ① N philosophers; N plates; N chopsticks
- ② If all philosophers grab right chopstick deadlock!

# Necessary conditions for deadlock

- Deadlock only if the all hold

## Bounded resources

A finite number of threads can use a resource; resources are finite

## No preemption

the resource is mine, MINE! (until I release it)

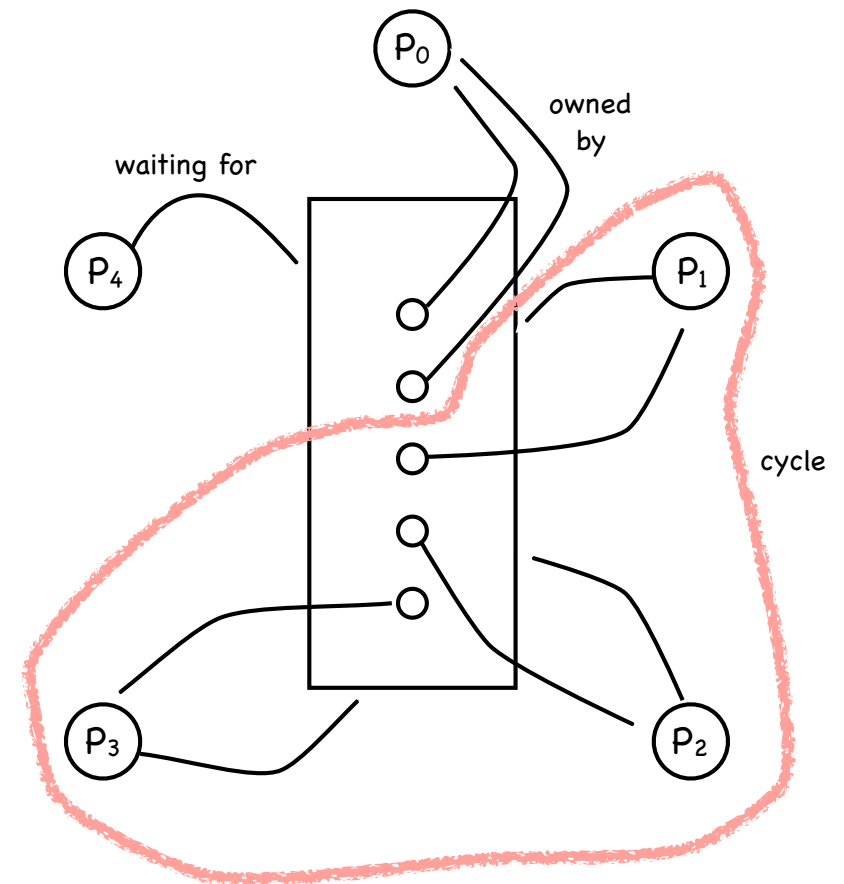
## Wait while holding

holds one resource while waiting for another

## Circular waiting

$T_i$  waits for  $T_{i+1}$  and holds a resource requested by  $T_{i-1}$   
sufficient if one instance of each resource

Not sufficient in general





# Preventing deadlock

- ④ Remove one of the necessary conditions

  - Provide sufficient resources

    - Removes "Bounded resources"

  - Preempt resources

    - Removes "No preemption"

  - Abort requests

    - Removes "Wait while holding"

  - Atomically acquire all resources

    - Removes "Wait while holding"

  - Lock ordering

    - Removes "Circular waiting"

# Lock ordering

## ④ A program code convention

Developers get together, have lunch, plan lock order

Usually reflects static assumptions about the structure of data

lock items in a list in order —what if order changes?

Nothing at compile time or run time prevents violating this convention!

Active research on making it better

Finding locking bugs

Automatically locking things properly

Transactional memory

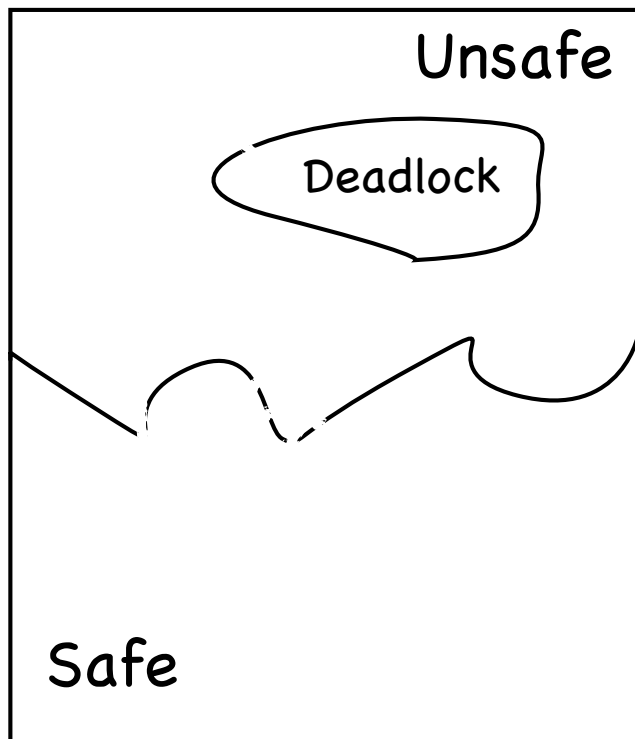
# Avoiding Deadlock: The Banker's Algorithm

E.W. Dijkstra & N. Habermann



- ④ Sum of maximum resources needs can exceed the total available resources
  - if there exists a schedule of loan fulfillments such that
    - all clients receive their maximal loan
    - build their house
    - pay back all the loan
- ④ More efficient than acquiring atomically all resources

# Living dangerously: Safe, Unsafe, Deadlocked



A system's trajectory  
through its state space

- ④ Safe: For any possible set of resource requests, there exists one safe schedule of processing requests that succeeds in granting all pending and future requests
  - no deadlock as long as system can enforce safe schedule
- ④ Unsafe: There exists a set of (pending and future) resource requests that leads to a deadlock, for any schedule in which requests are processed
  - unlucky set of requests can force deadlock
- ④ Deadlocked: The system has at least one deadlock

# The Banker's books

- Max<sub>ij</sub> = max amount of units of resource R<sub>j</sub> needed by P<sub>i</sub>

$$\text{MaxClaim}_i = \text{Max}_{ij}$$

- Alloc<sub>ij</sub> = current allocation of R<sub>j</sub> held by P<sub>i</sub>

$$\text{HasNow}_i = \text{Alloc}_{ij}$$

- Avail<sub>j</sub> = number of units of R<sub>j</sub> available

- A request by P<sub>k</sub> is safe if there is schedule P<sub>1</sub>, P<sub>2</sub>, ..., P<sub>n</sub> such that, for all P<sub>i</sub>, assuming the request is granted,

$$\text{MaxClaim}_i - \text{HasNow}_i \leq \text{Avail} + \sum_{j=1}^{i-1} \text{HasNow}_j$$

# An Example

- 5 processes, 4 resources

Max

Alloc

Avail

0	0	1	2
1	7	5	0
2	3	5	6
0	6	5	2
0	6	5	6

0	0	1	2
1	0	0	0
1	3	5	3
0	6	3	2
0	0	1	4

- Is this a safe state?

# An Example

- 5 processes, 4 resources

Max	Alloc	Avail	MaxRequest
0 0 1 2	0 0 1 2		0 0 0 0
1 7 5 0	1 0 0 0		0 7 5 0
2 3 5 6	1 3 5 3		1 0 0 3
0 6 5 2	0 6 3 2		0 0 2 0
0 6 5 6	0 0 1 4		0 6 4 2

- Is this a safe state?  $P_1, P_4, P_2, P_3, P_5$

While safe sequence does not include all processes:

Is there a  $P_i$  such that  $\text{MaxRequest}_i \leq \text{Avail}$ ?

if no, exit with unsafe

if yes, add  $P_i$  to the sequence and set  $\text{Avail} = \text{Avail} + \text{HasNow}_i$

Exit with safe

# An Example

- 5 processes, 4 resources

	Max				Alloc				Avail				MaxRequest			
	R <sub>1</sub>	R <sub>2</sub>	R <sub>3</sub>	R <sub>4</sub>	R <sub>1</sub>	R <sub>2</sub>	R <sub>3</sub>	R <sub>4</sub>	R <sub>1</sub>	R <sub>2</sub>	R <sub>3</sub>	R <sub>4</sub>	R <sub>1</sub>	R <sub>2</sub>	R <sub>3</sub>	R <sub>4</sub>
P <sub>1</sub>	0	0	1	2	0	0	1	2	1	5	2	0	0	0	0	0
P <sub>2</sub>	1	7	5	0	1	0	0	0	0	7	5	0	0	7	5	0
P <sub>3</sub>	2	3	5	6	1	3	5	3	1	0	0	3	1	0	0	3
P <sub>4</sub>	0	6	5	2	0	6	3	2	0	0	2	0	0	0	2	0
P <sub>5</sub>	0	6	5	6	0	0	1	4	0	6	4	2	0	6	4	2

P2 want to change its allocation to 0 4 2 0

Safe?



# An Example

- 5 processes, 4 resources

	Max				Alloc				Avail				MaxRequest			
	R <sub>1</sub>	R <sub>2</sub>	R <sub>3</sub>	R <sub>4</sub>	R <sub>1</sub>	R <sub>2</sub>	R <sub>3</sub>	R <sub>4</sub>	R <sub>1</sub>	R <sub>2</sub>	R <sub>3</sub>	R <sub>4</sub>	R <sub>1</sub>	R <sub>2</sub>	R <sub>3</sub>	R <sub>4</sub>
P <sub>1</sub>	0	0	1	2	0	0	1	2	2	1	0	0	0	0	0	0
P <sub>2</sub>	1	7	5	0	0	4	2	0	2	1	0	0	1	3	3	0
P <sub>3</sub>	2	3	5	6	1	3	5	3	0	0	0	0	1	0	0	3
P <sub>4</sub>	0	6	5	2	0	6	3	2	0	0	2	0	0	0	2	0
P <sub>5</sub>	0	6	5	6	0	0	1	4	0	6	4	2	0	6	4	2

P2 want to change its allocation to 0 4 2 0

Safe?

# Detecting Deadlock

- 5 processes, 3 resources. We no longer know Max.

	Alloc			Avail			Pending		
	R <sub>1</sub>	R <sub>2</sub>	R <sub>3</sub>	R <sub>1</sub>	R <sub>2</sub>	R <sub>3</sub>			
P <sub>1</sub>	0	1	0	0	0	0	0	0	0
P <sub>2</sub>	2	0	0				2	0	2
P <sub>3</sub>	3	0	3				0	0	0
P <sub>4</sub>	2	1	1				1	0	2
P <sub>5</sub>	0	0	2				0	0	2

- Given the set of pending requests, is there a safe sequence?  
If no, deadlock

# Detecting Deadlock

- 5 processes, 3 resources. We no longer know Max.

	Alloc			Avail			Pending		
	R <sub>1</sub>	R <sub>2</sub>	R <sub>3</sub>	R <sub>1</sub>	R <sub>2</sub>	R <sub>3</sub>			
P <sub>1</sub>	0	1	0	0	0	0	0	0	0
P <sub>2</sub>	2	0	0				2	0	2
P <sub>3</sub>	3	0	3				0	0	1
P <sub>4</sub>	2	1	1				1	0	2
P <sub>5</sub>	0	0	2				0	0	2

- Given the set of pending requests, is there a safe sequence?

If no, deadlock

- Can we avoid deadlock by delaying granting requests?

Deadlock triggered when request formulated, not granted