

# Parallel Programming Exercise Session 4

Spring 2020

# Feedback: Exercise 3

# Counter


Let's count number of times a given event occurs

```
public interface Counter {  
    public void increment();  
    public int value();  
}
```

```
// background threads  
for (int i = 0; i < numIterations; i++) {  
    // perform some work  
  
    counter.increment();  
}  
  
// progress thread  
while (isWorking) {  
    System.out.println(counter.value());  
}
```


10 iterations each

Counter




value of the  
shared Counter


Thread 1



Thread 2




Thread 3




number of times  
increment() is called

Counter




value of the  
shared Counter


Thread 1



Thread 2




Thread 3




number of times  
increment() is called

Counter




value of the  
shared Counter


Thread 1




Thread 2



Thread 3




number of times  
increment() is called



value of the shared Counter

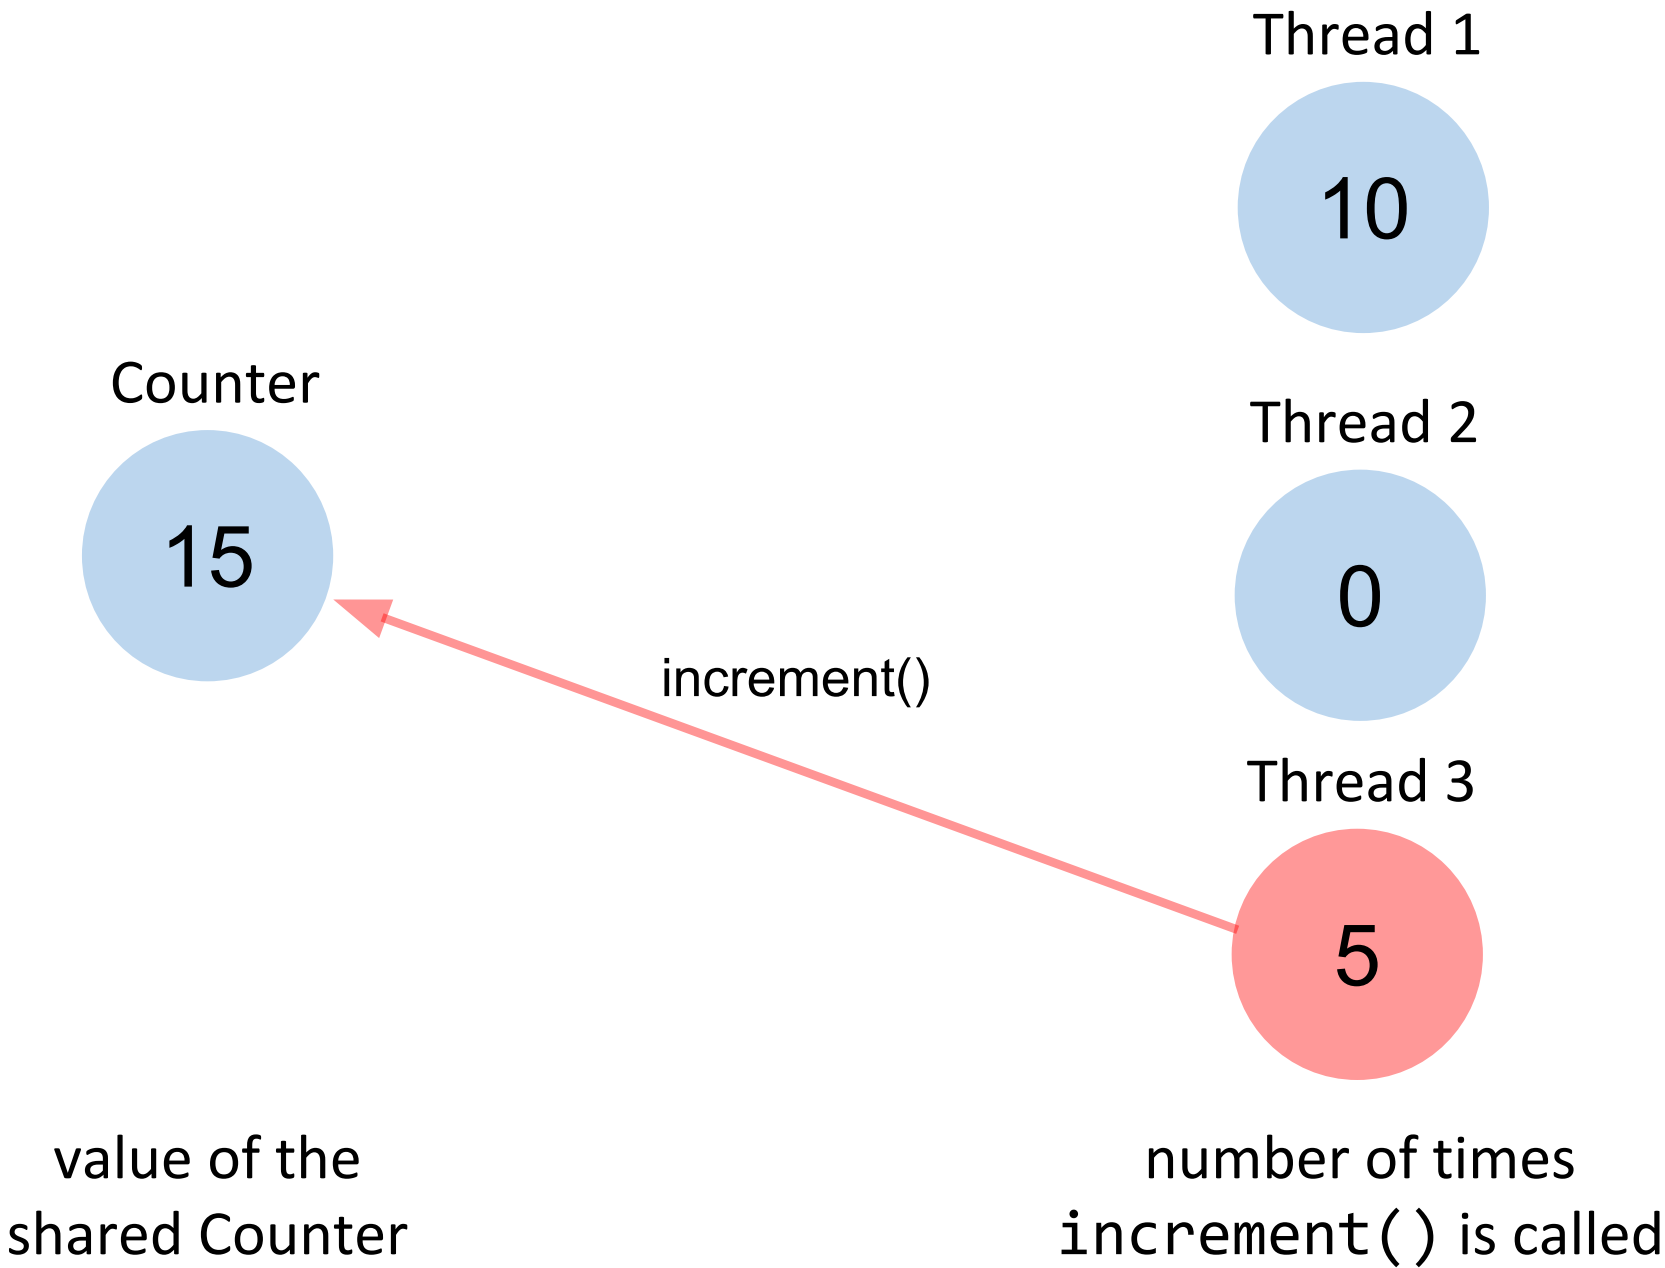
number of times increment() is called

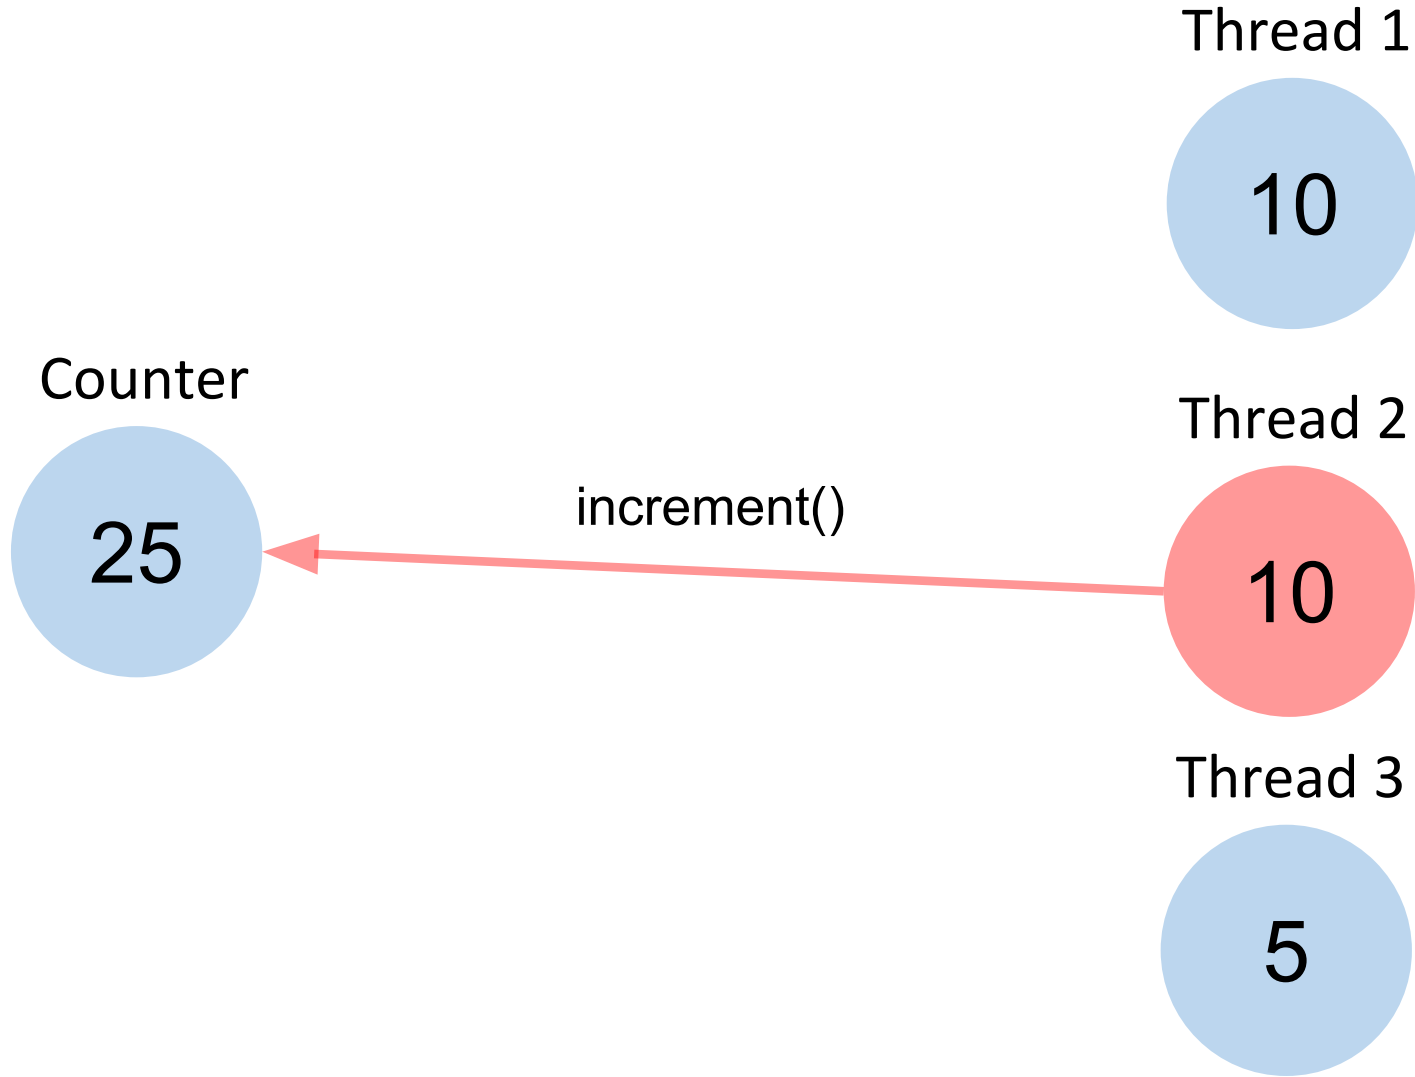


value of the shared Counter

number of times increment() is called









value of the shared Counter



number of times increment() is called



increment()

value of the shared Counter

number of times increment() is called



# Task A: SequentialCounter

```
public class SequentialCounter implements Counter {  
  
    public void increment() {  
        ??  
    }  
  
    public int value() {  
        ??  
    }  
}
```

# Task A: SequentialCounter

```
public class SequentialCounter implements Counter {  
    private int c = 0;  
  
    public void increment() {  
        c++;  
    }  
  
    public int value() {  
        return c;  
    }  
}
```

# Task A: SequentialCounter

Counter  
0

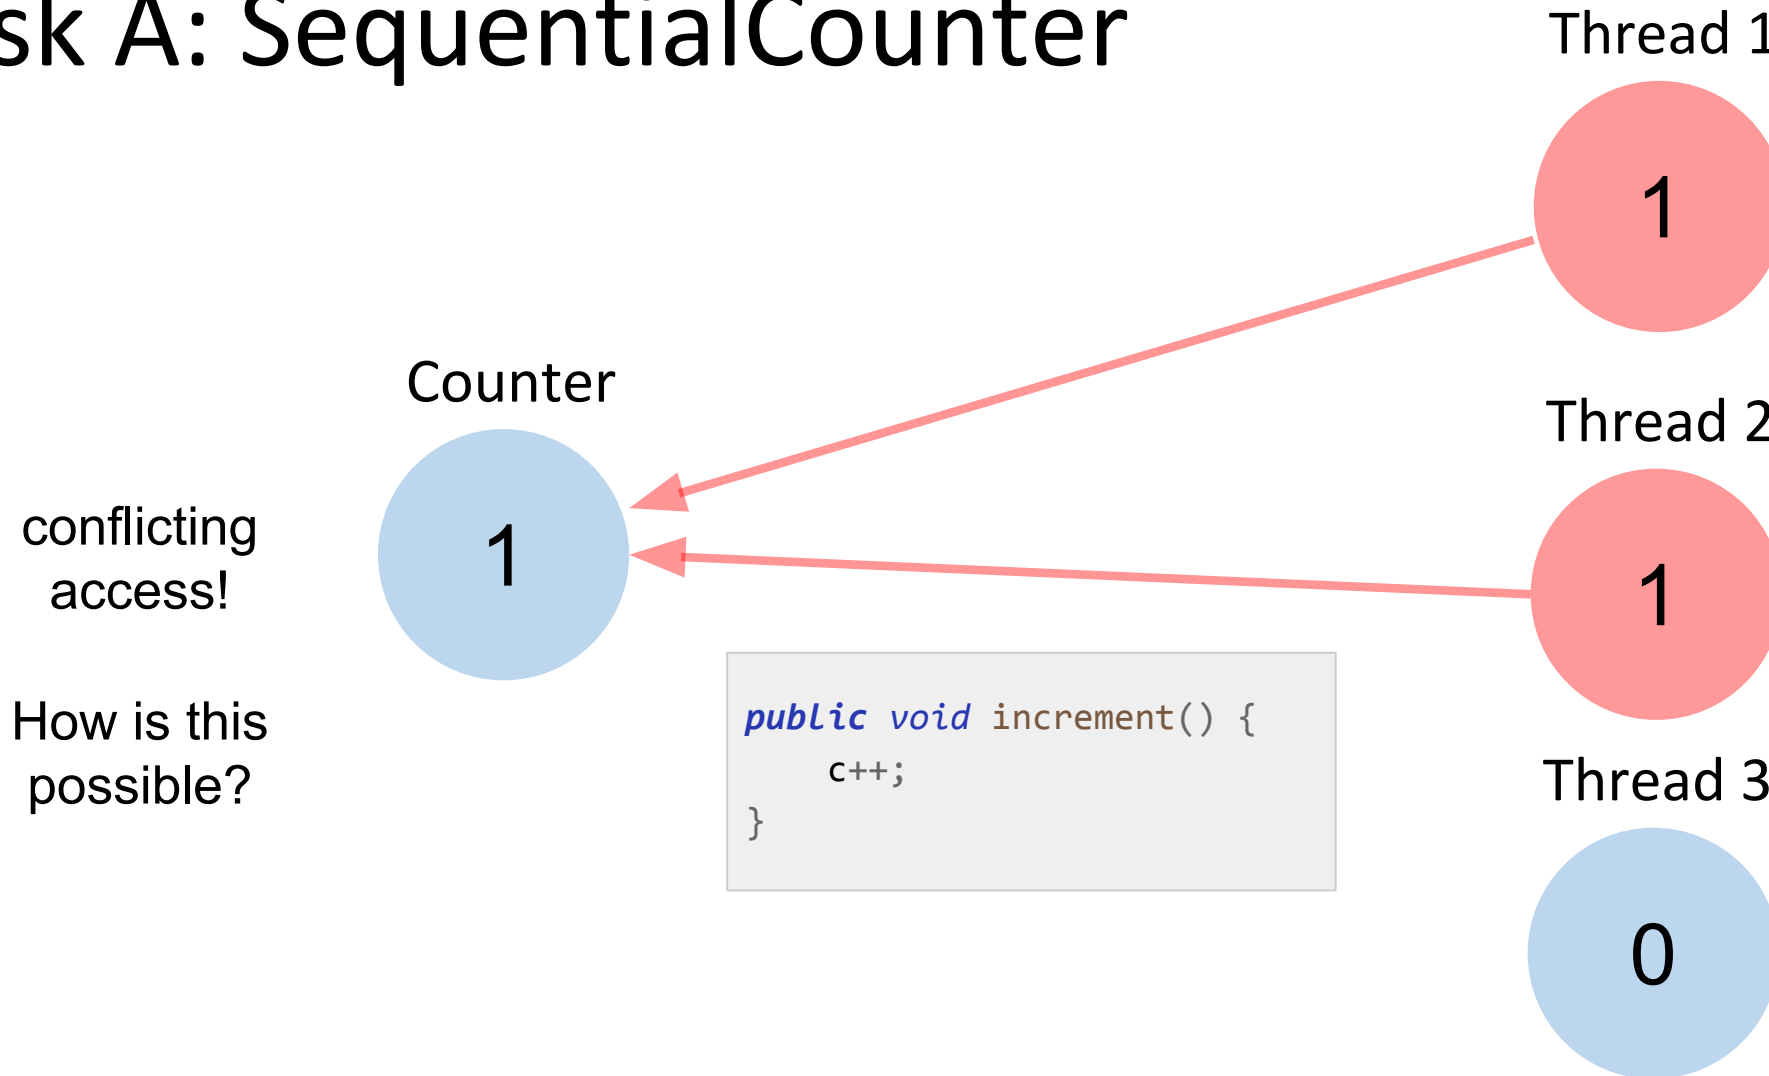
```
public void increment() {  
    c++;  
}
```

Thread 1  
0

Thread 2  
0

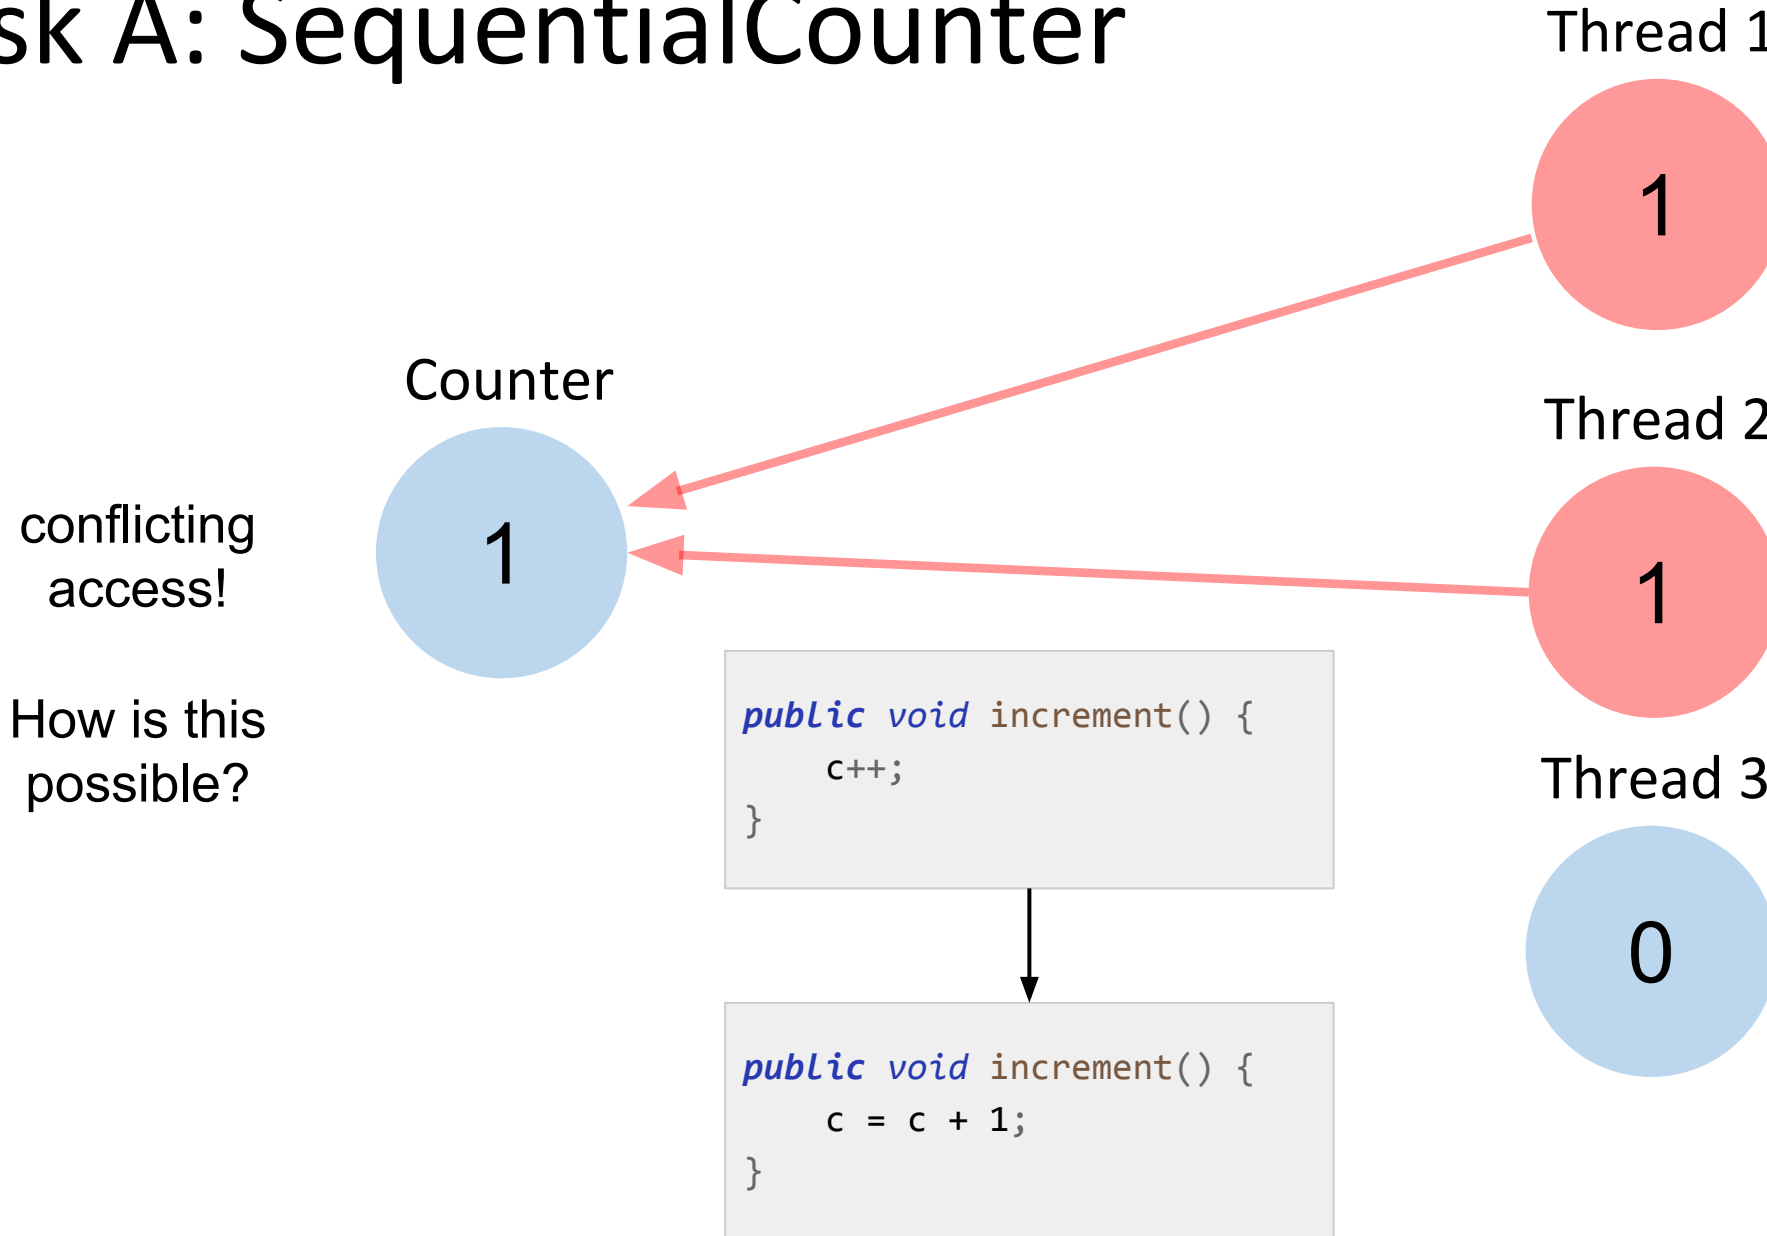
Thread 3  
0

# Task A: SequentialCounter

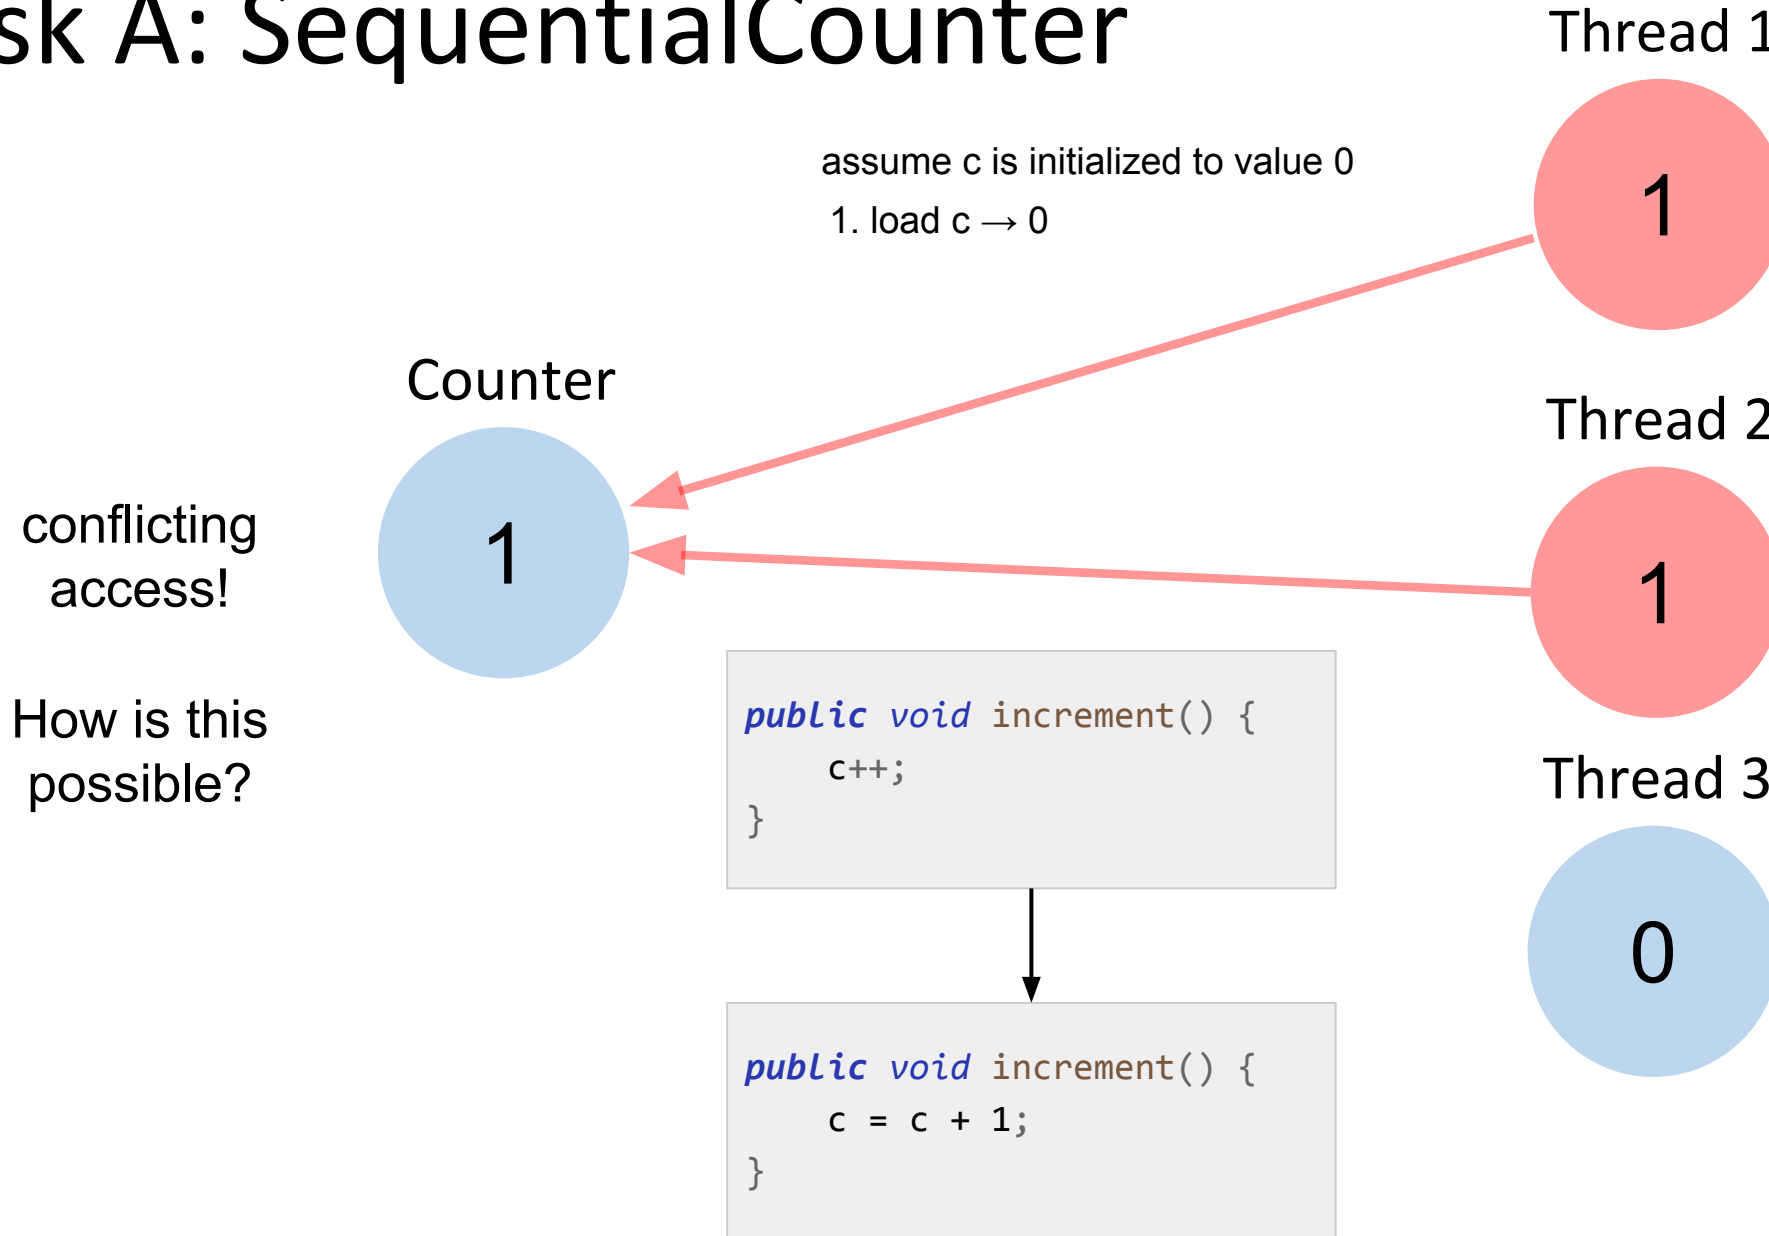




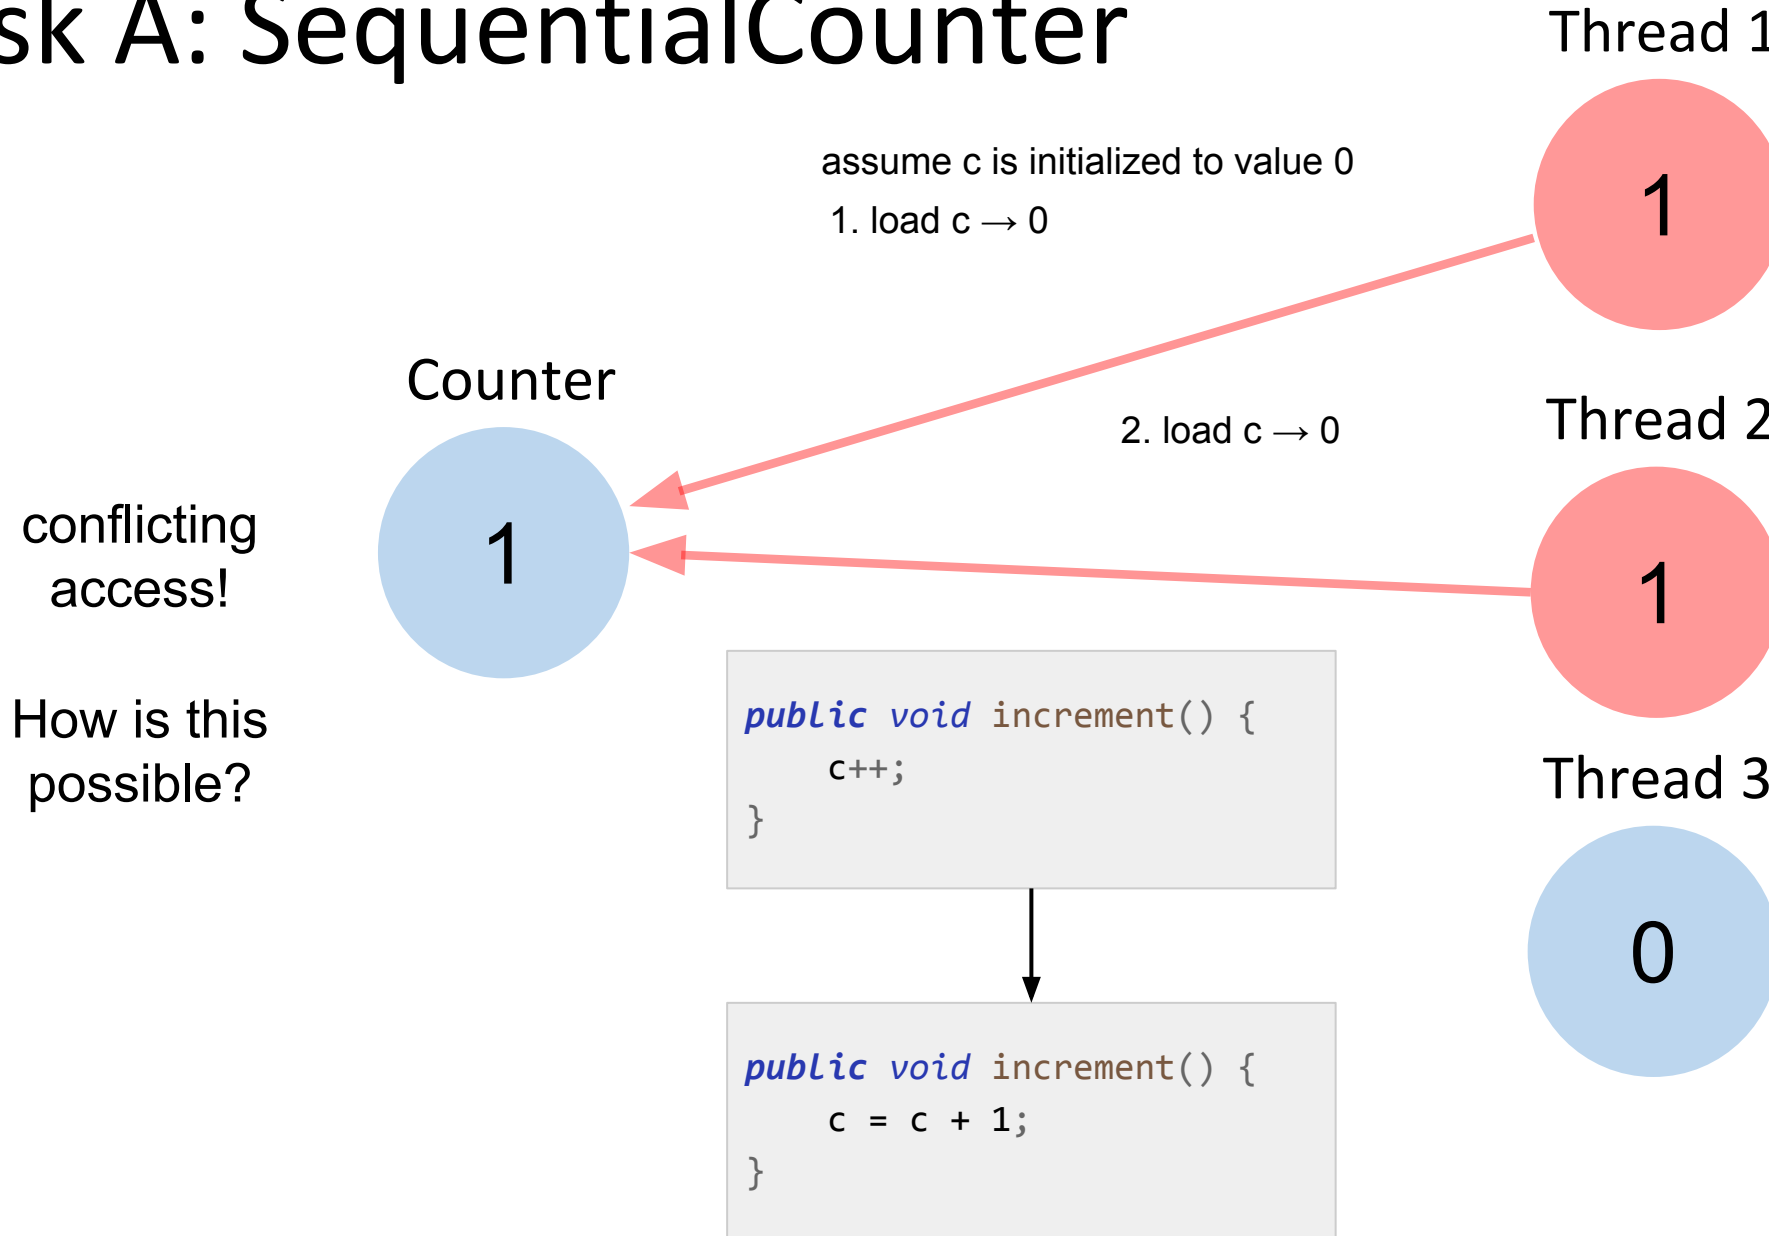
# Task A: SequentialCounter



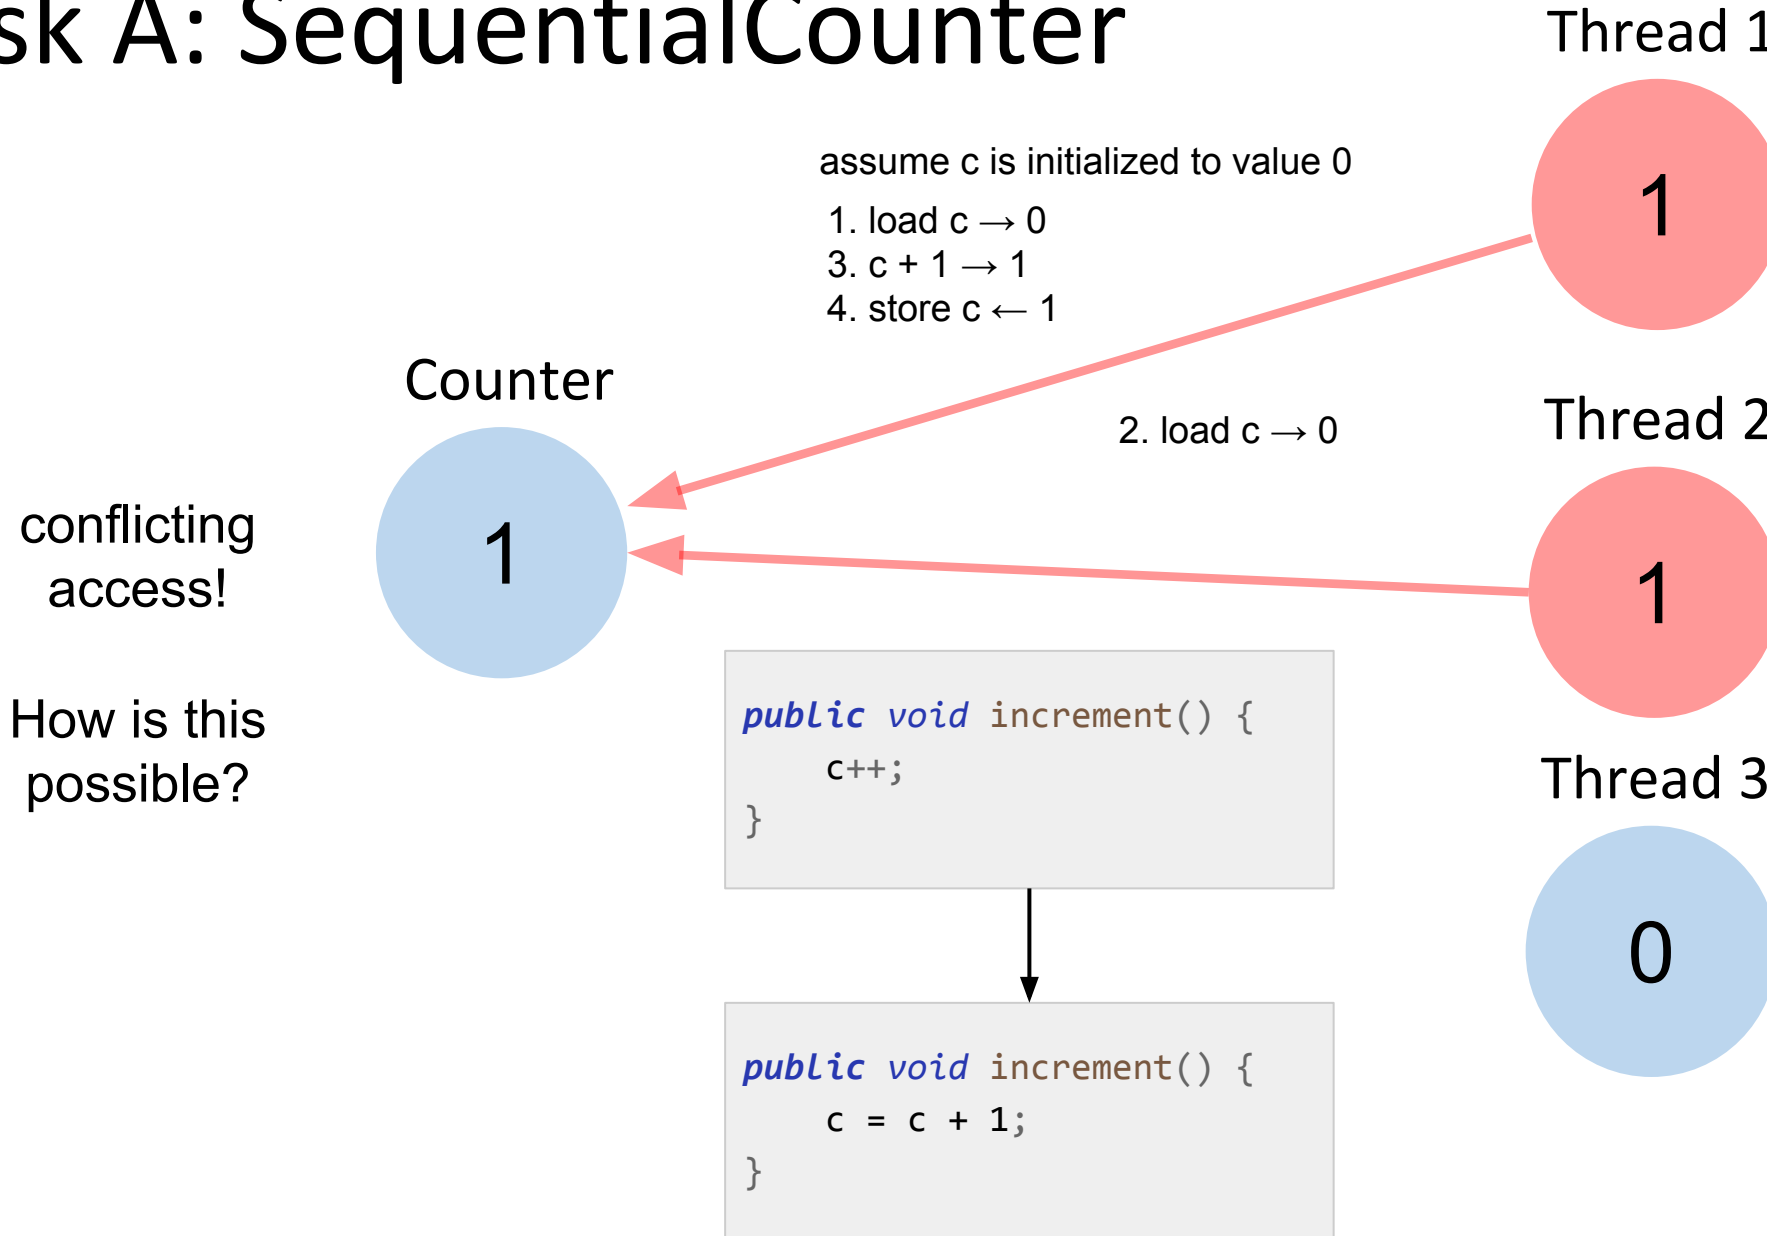
# Task A: SequentialCounter



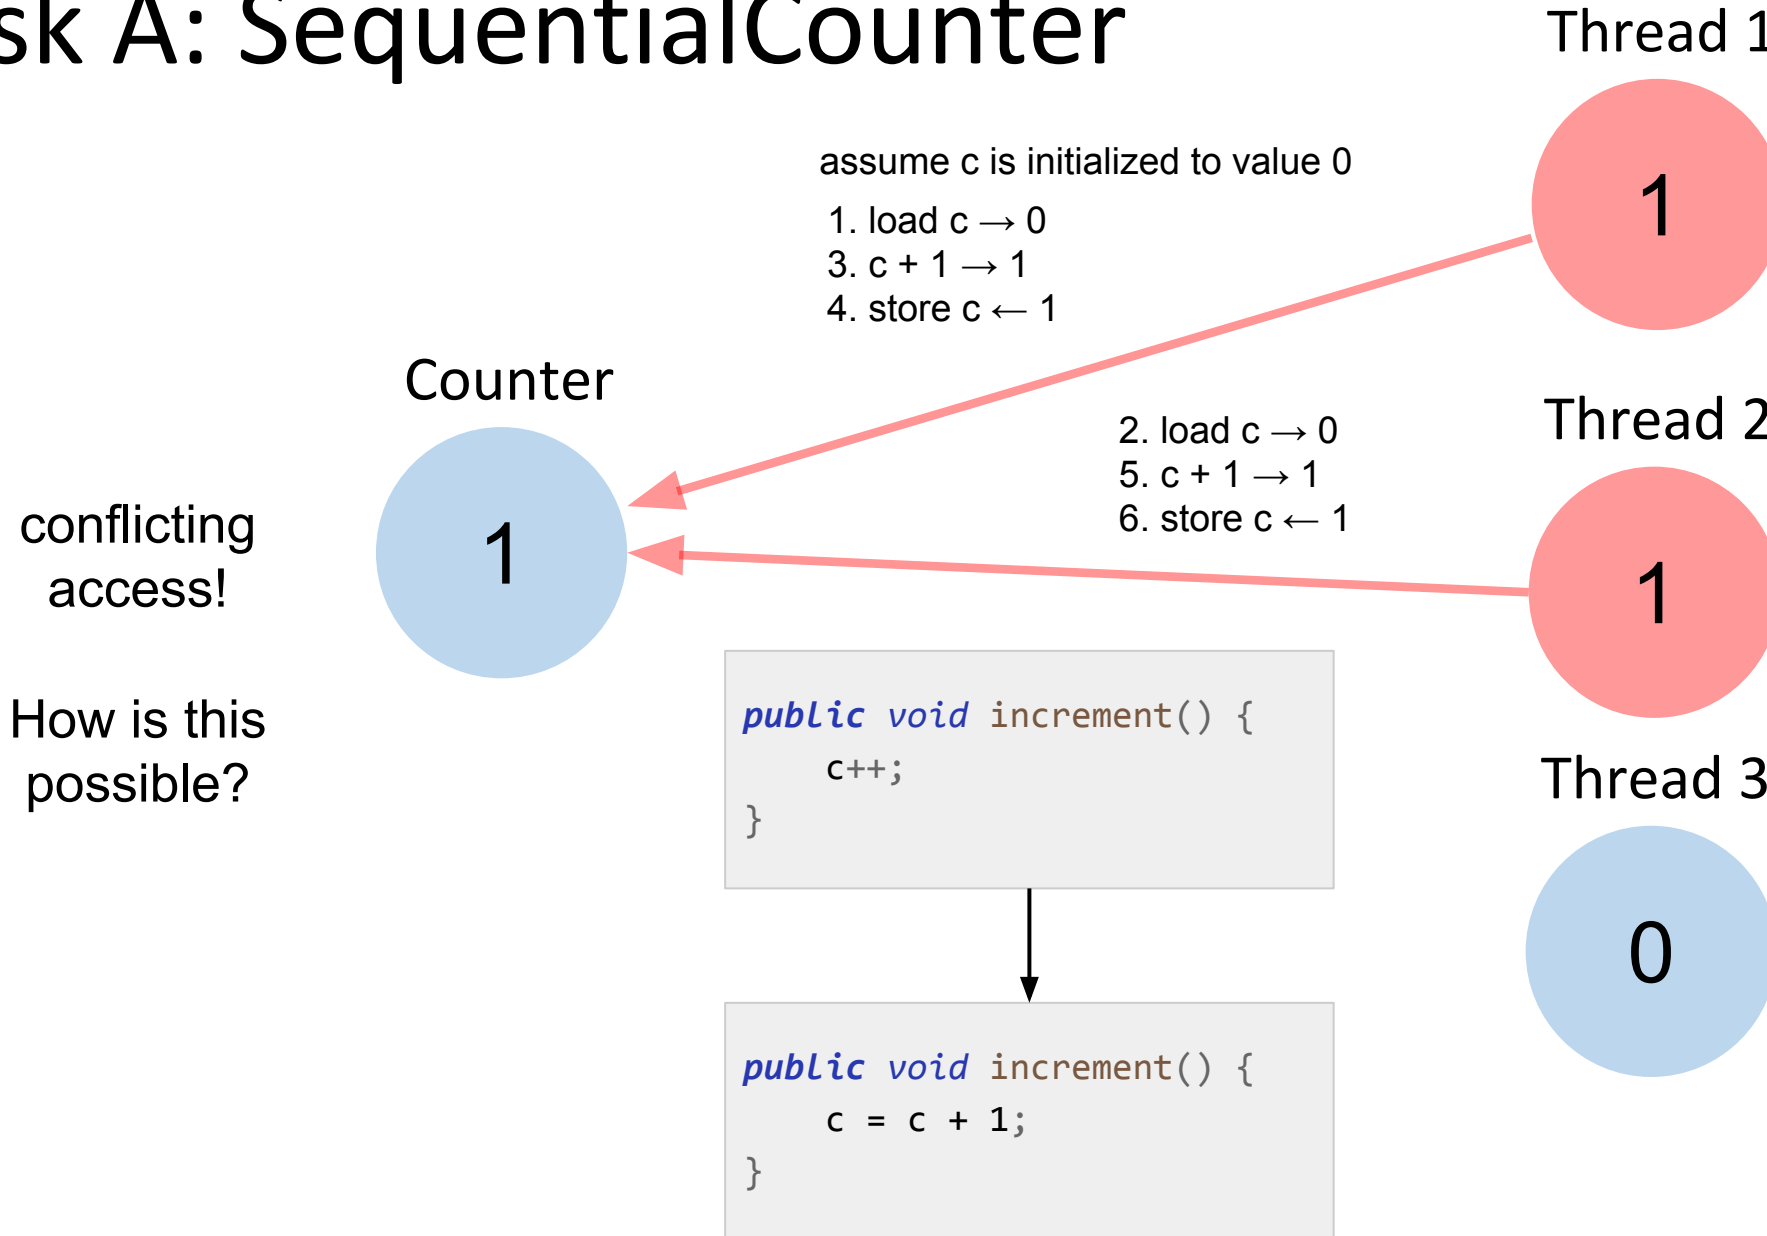
# Task A: SequentialCounter



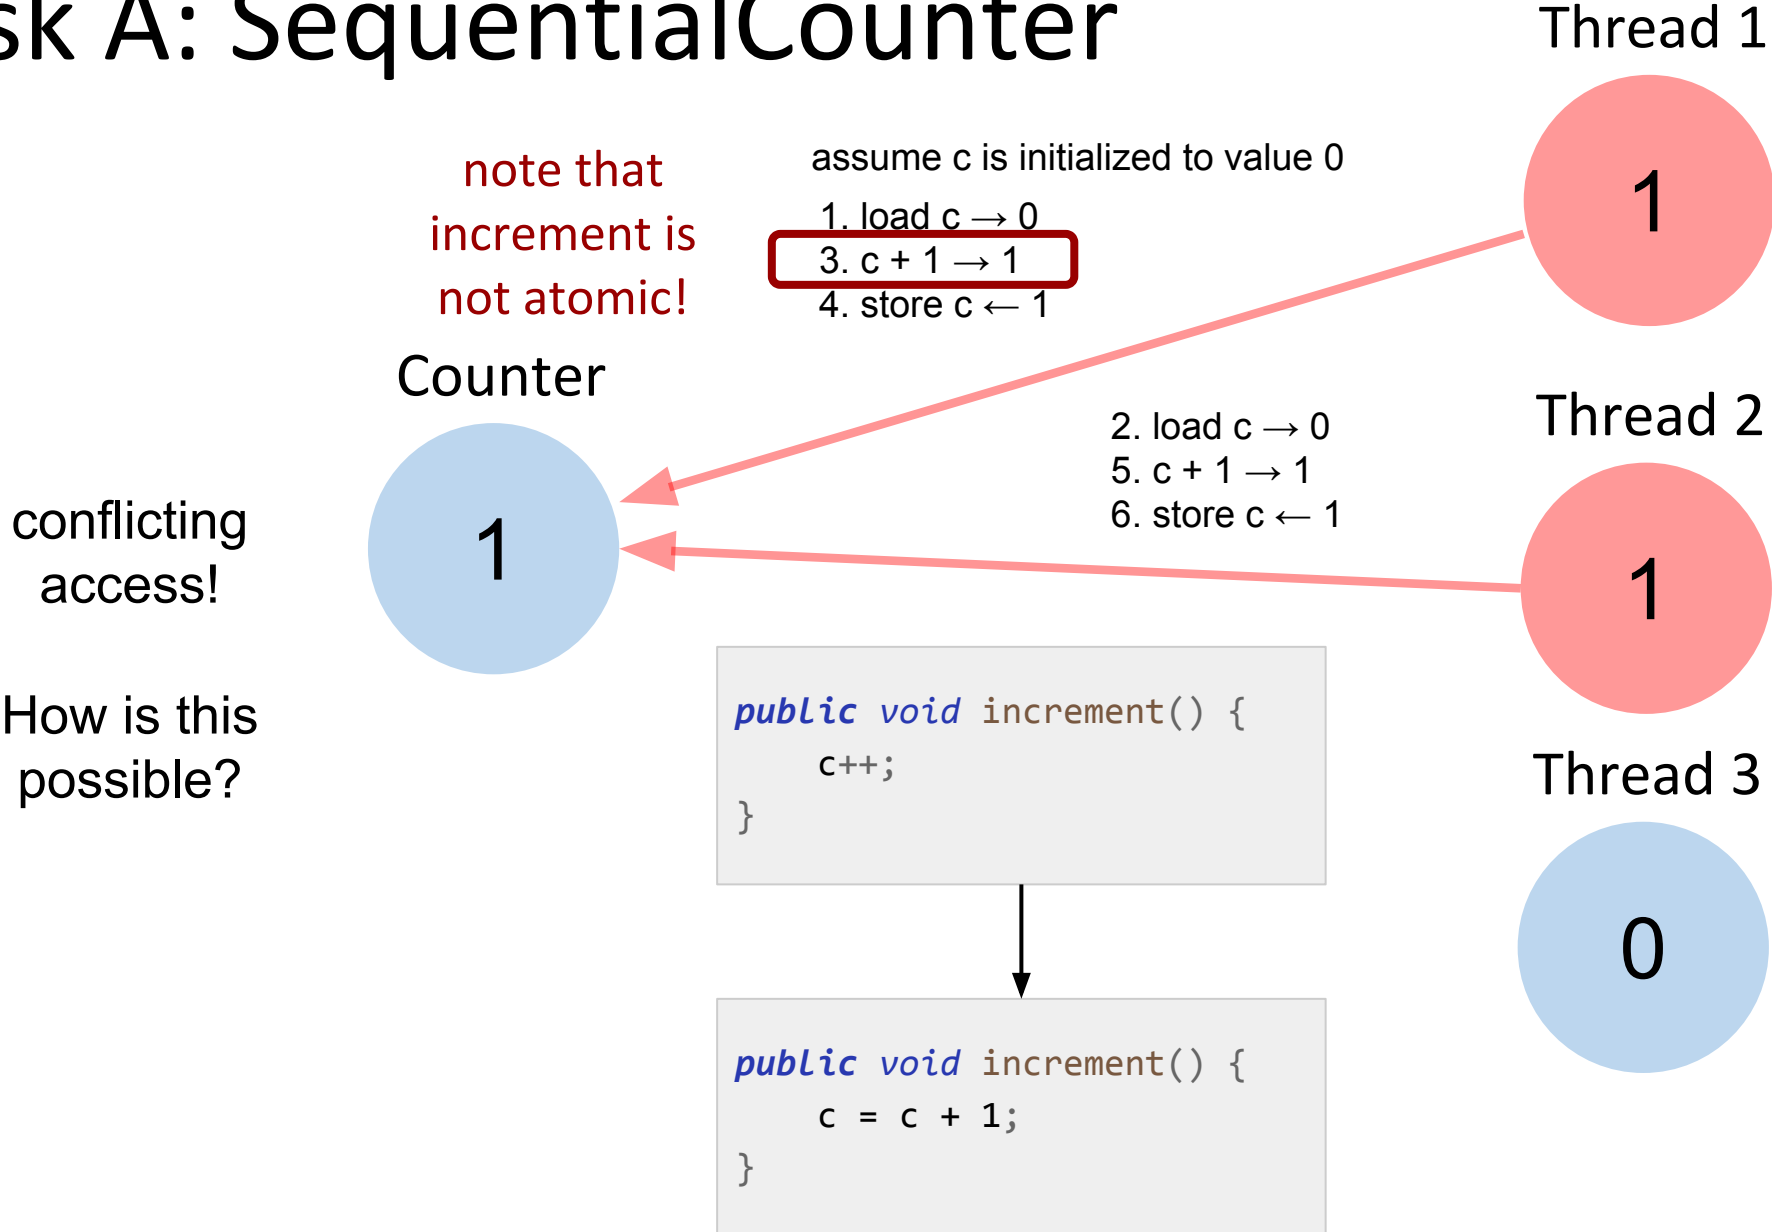
# Task A: SequentialCounter



# Task A: SequentialCounter



# Task A: SequentialCounter



# Task B: SynchronizedCounter

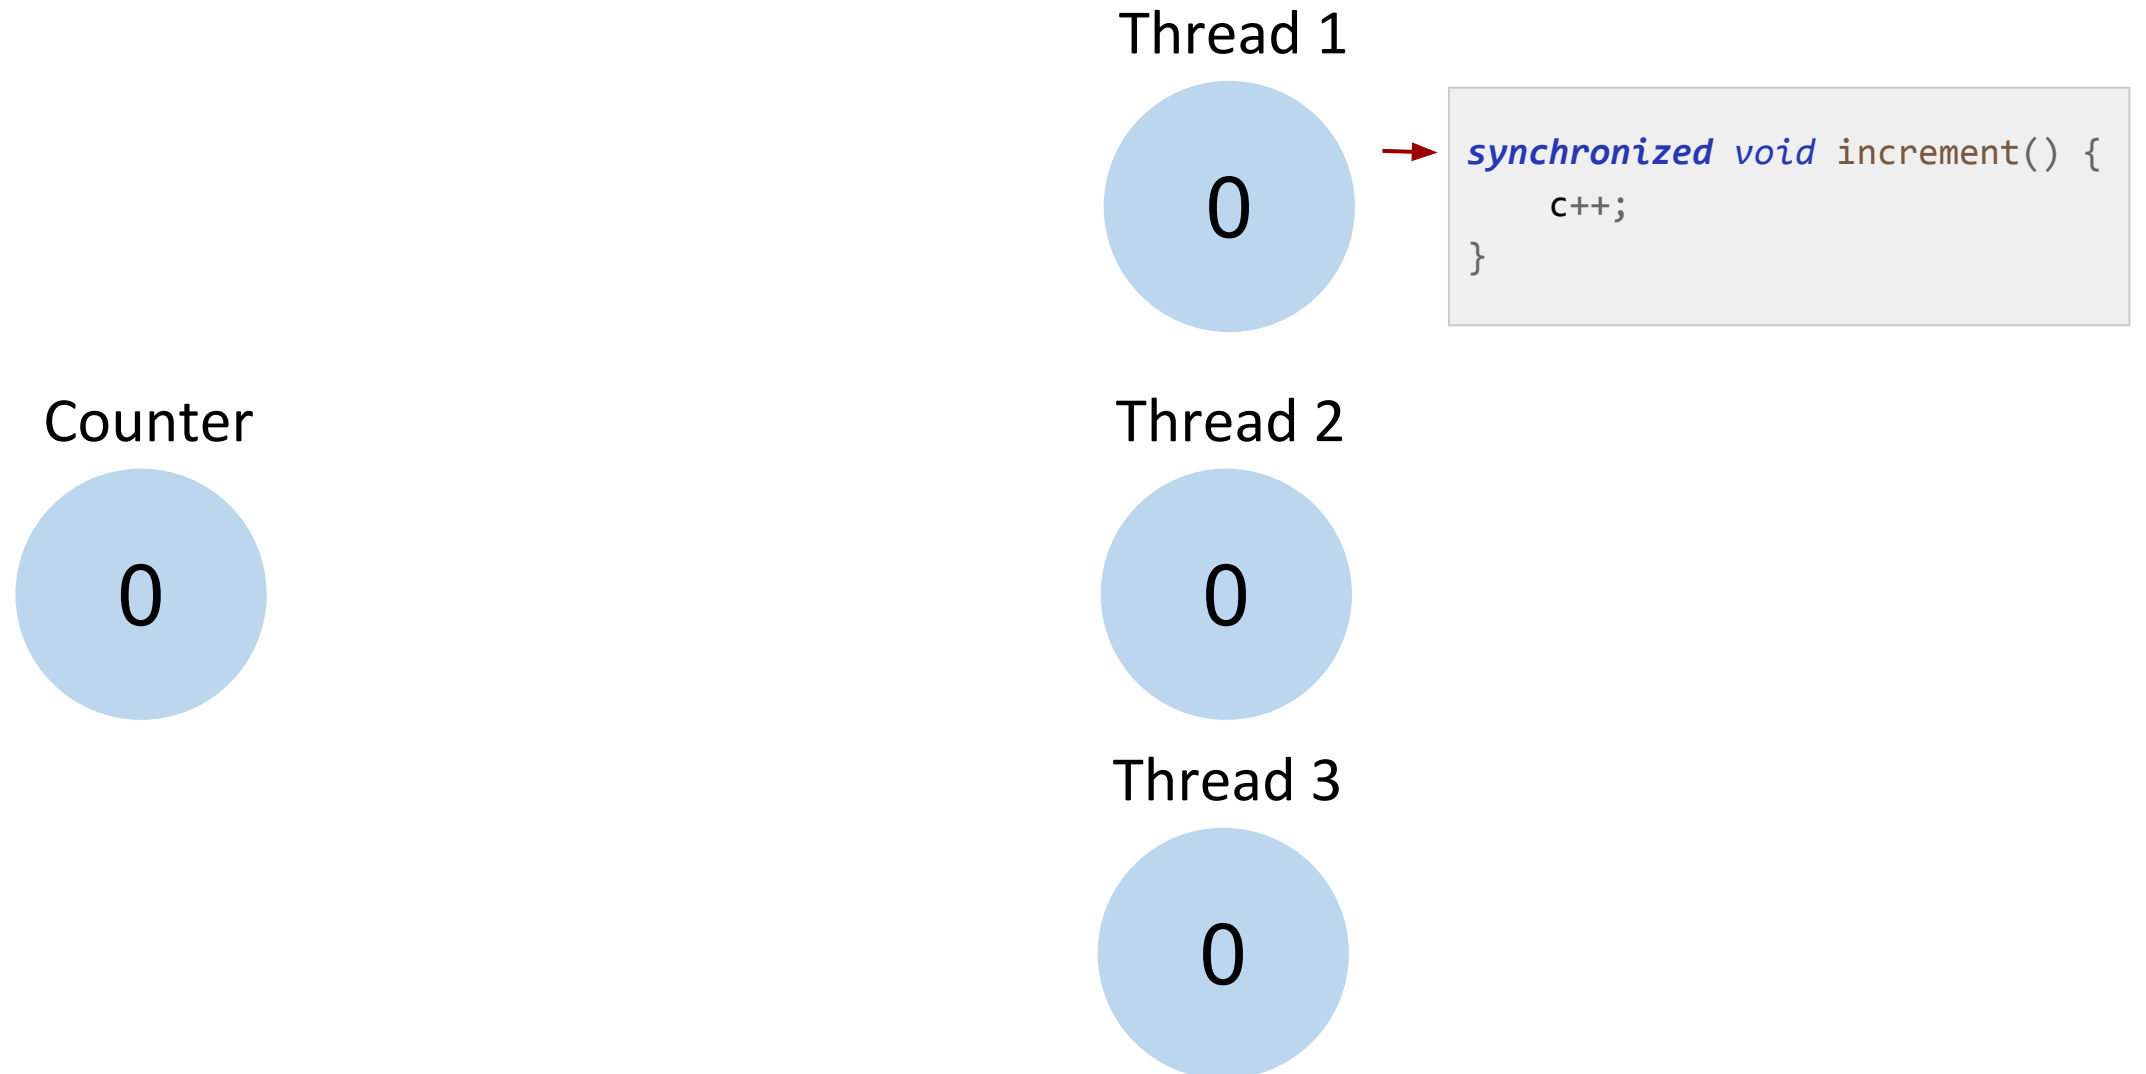
```
public class SynchronizedCounter implements Counter {  
  
    public void increment() {  
        ??  
    }  
  
    public int value() {  
        ??  
    }  
}
```

# Task B: SynchronizedCounter

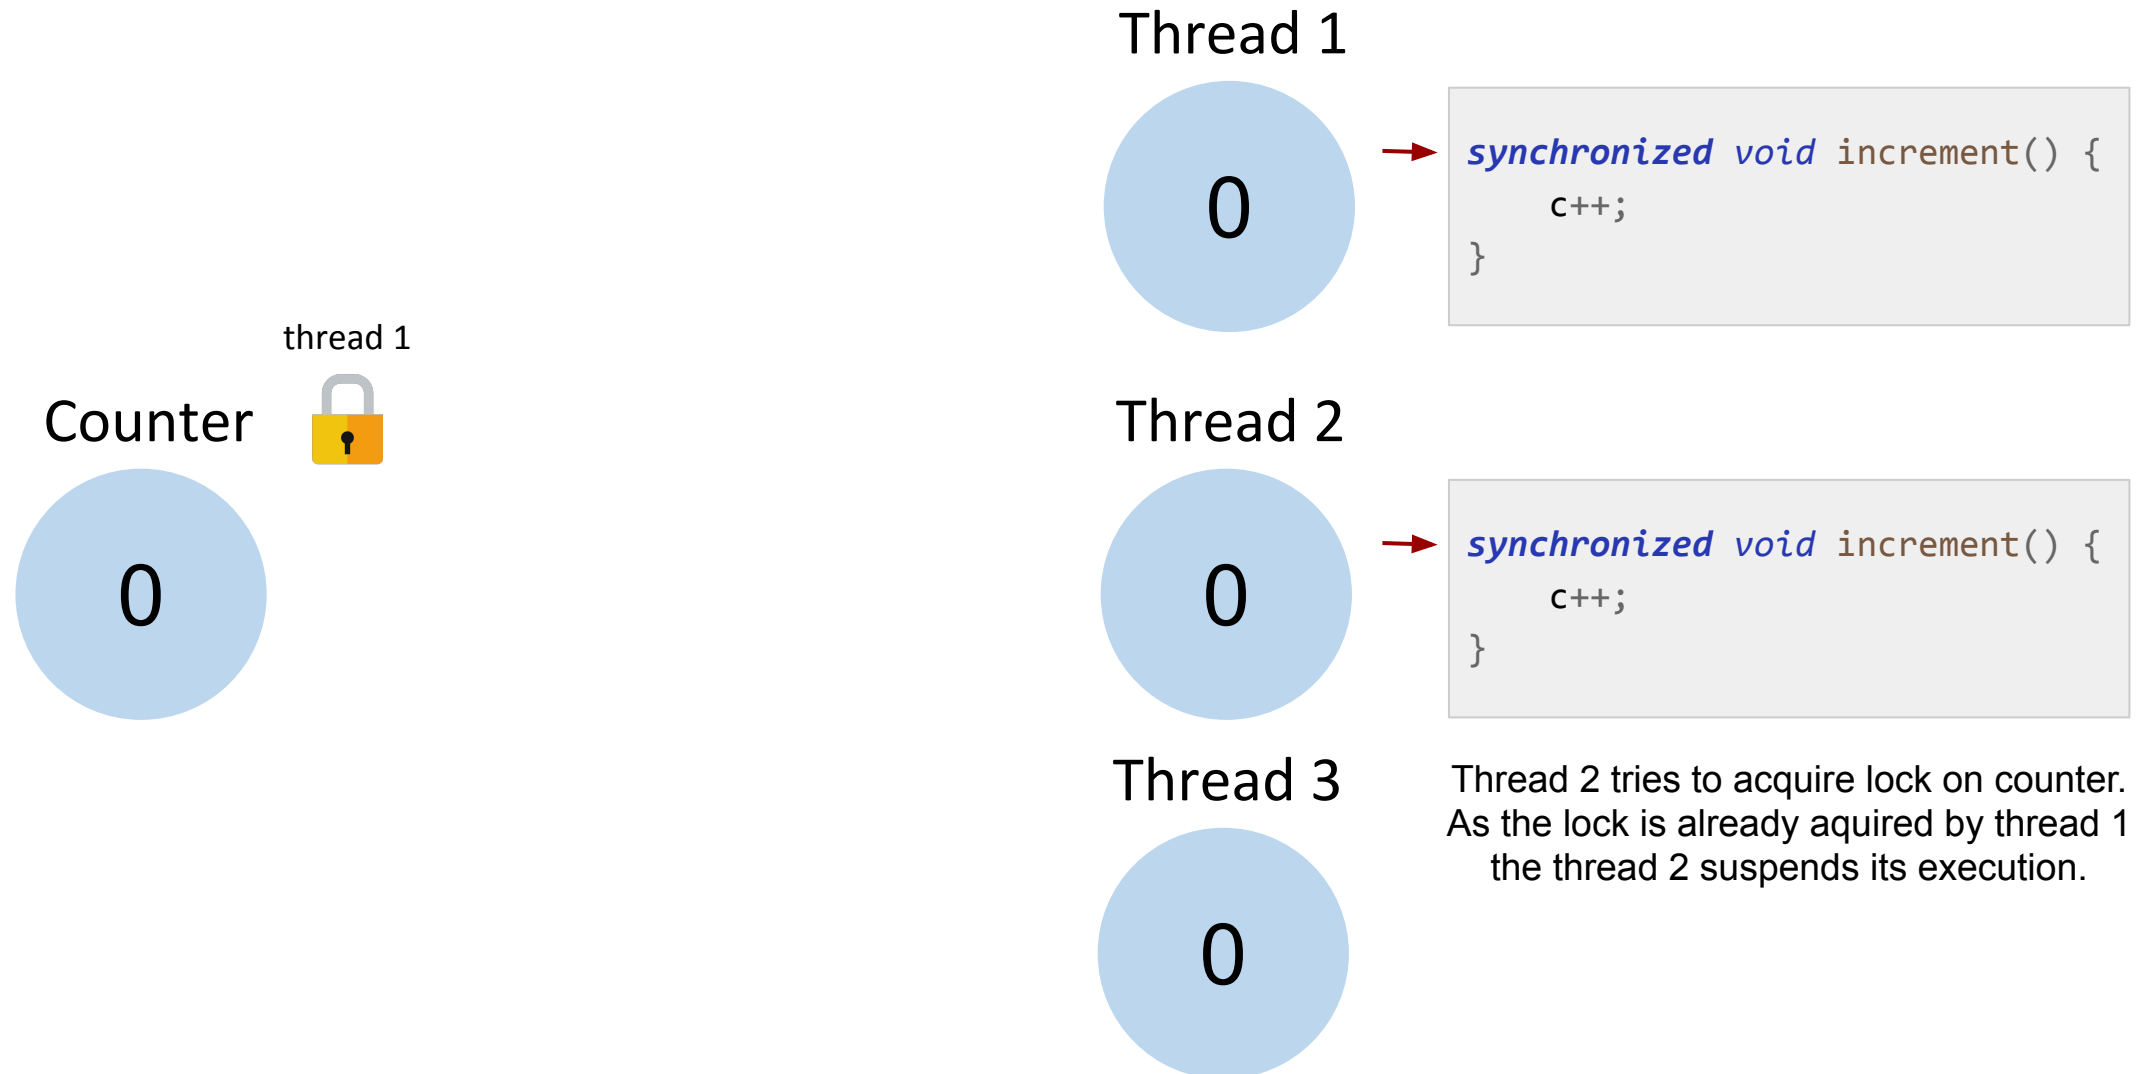
```
public class SynchronizedCounter implements Counter {  
    private int c = 0;  
  
    public synchronized void increment() {  
        c++;  
    }  
  
    public synchronized int value() {  
        return c;  
    }  
}
```



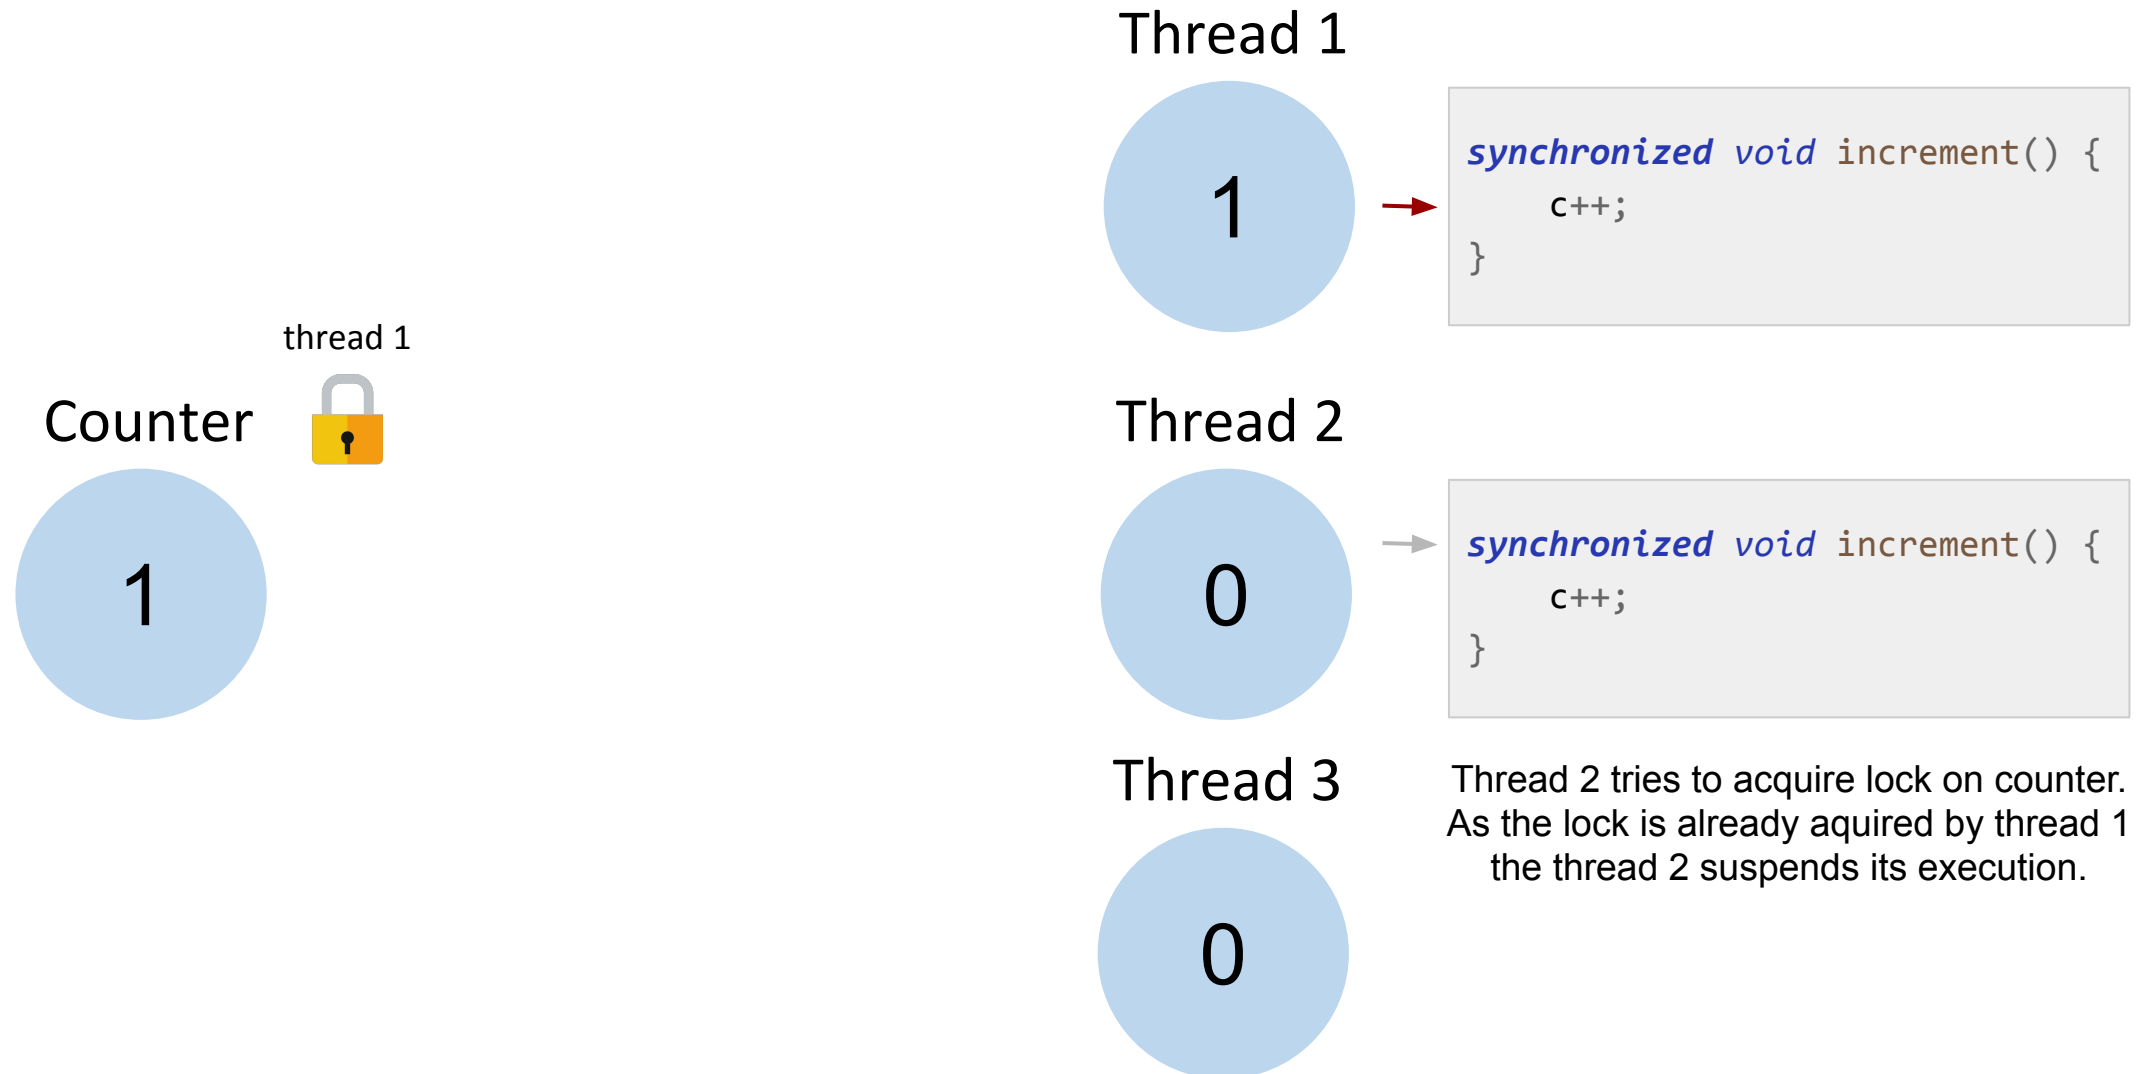
# Task B: SynchronizedCounter



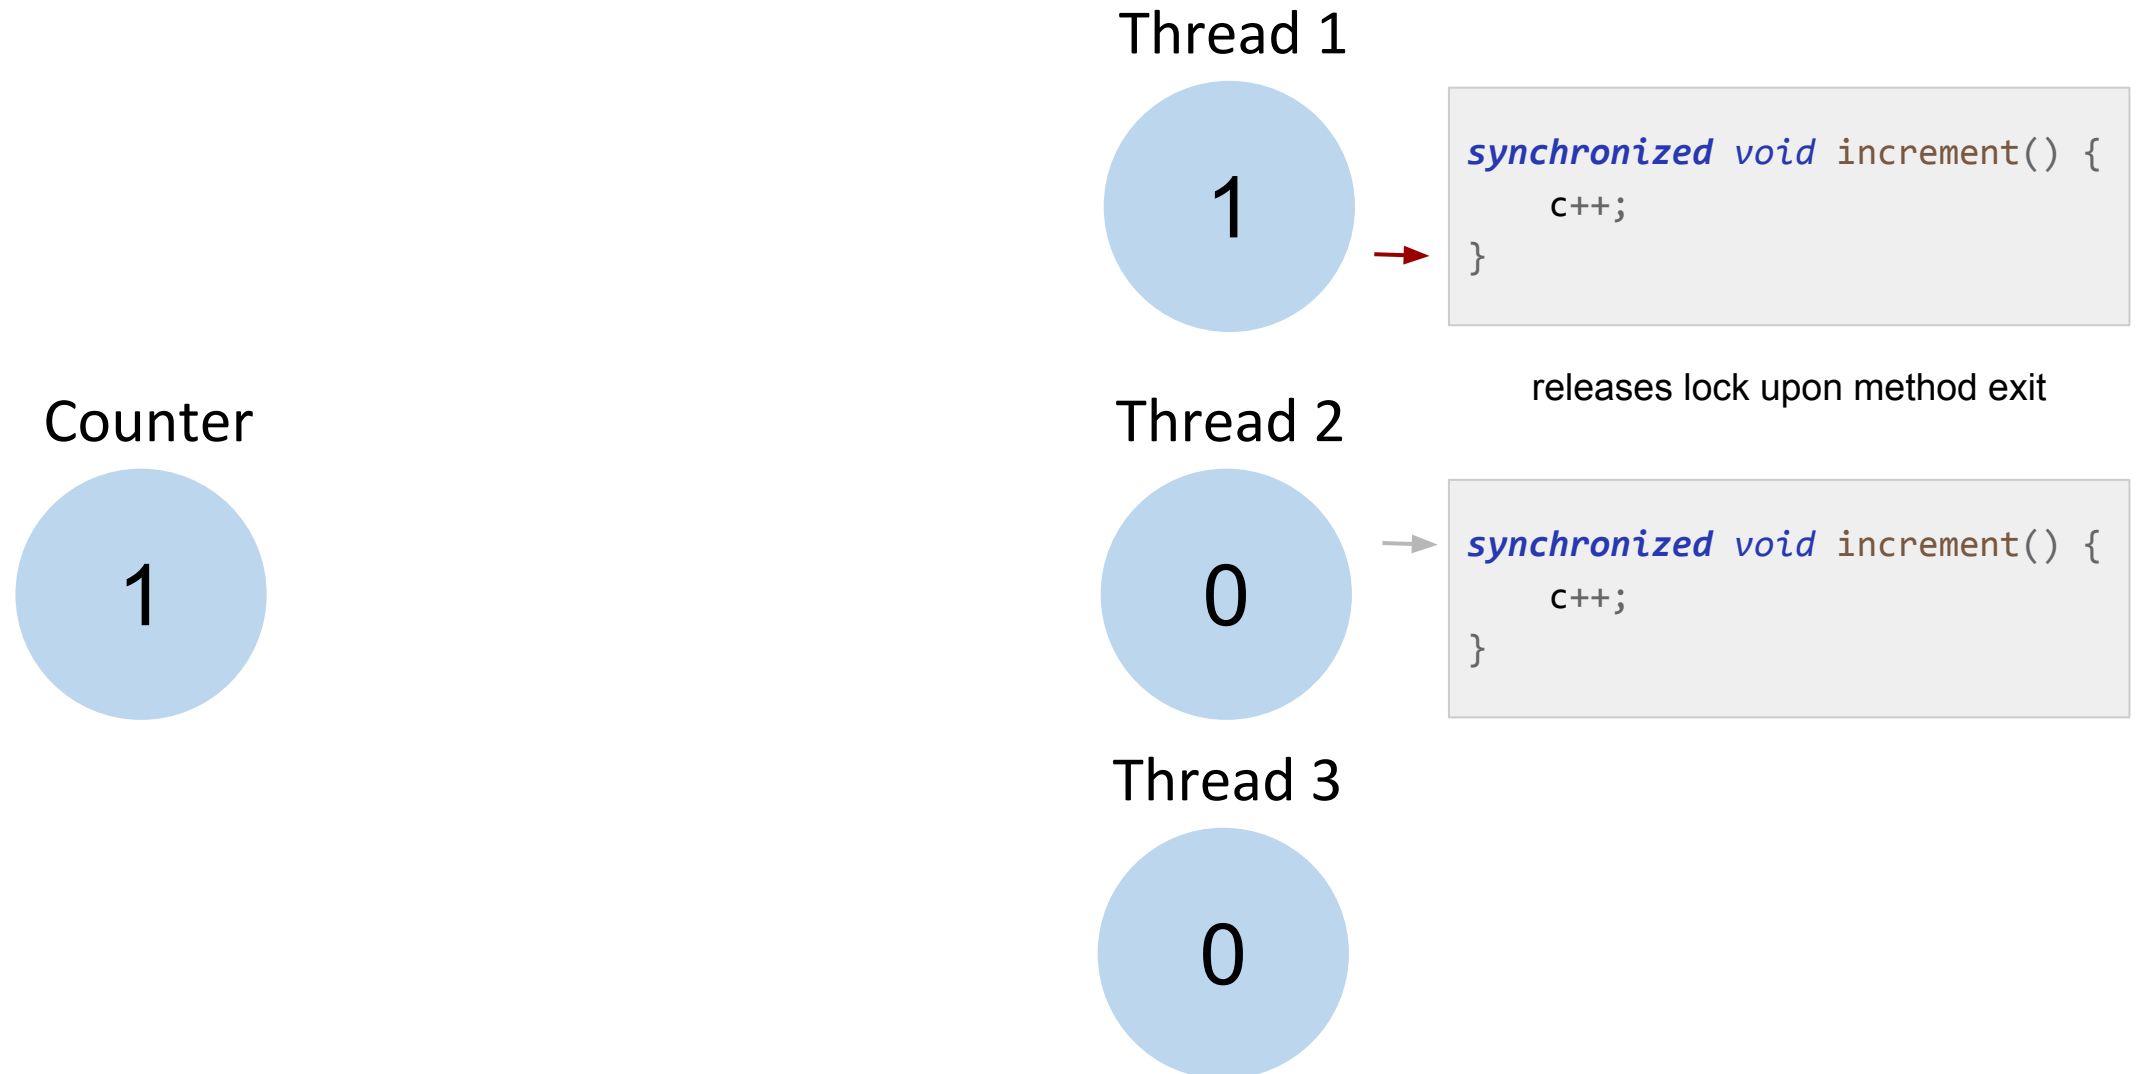
# Task B: SynchronizedCounter



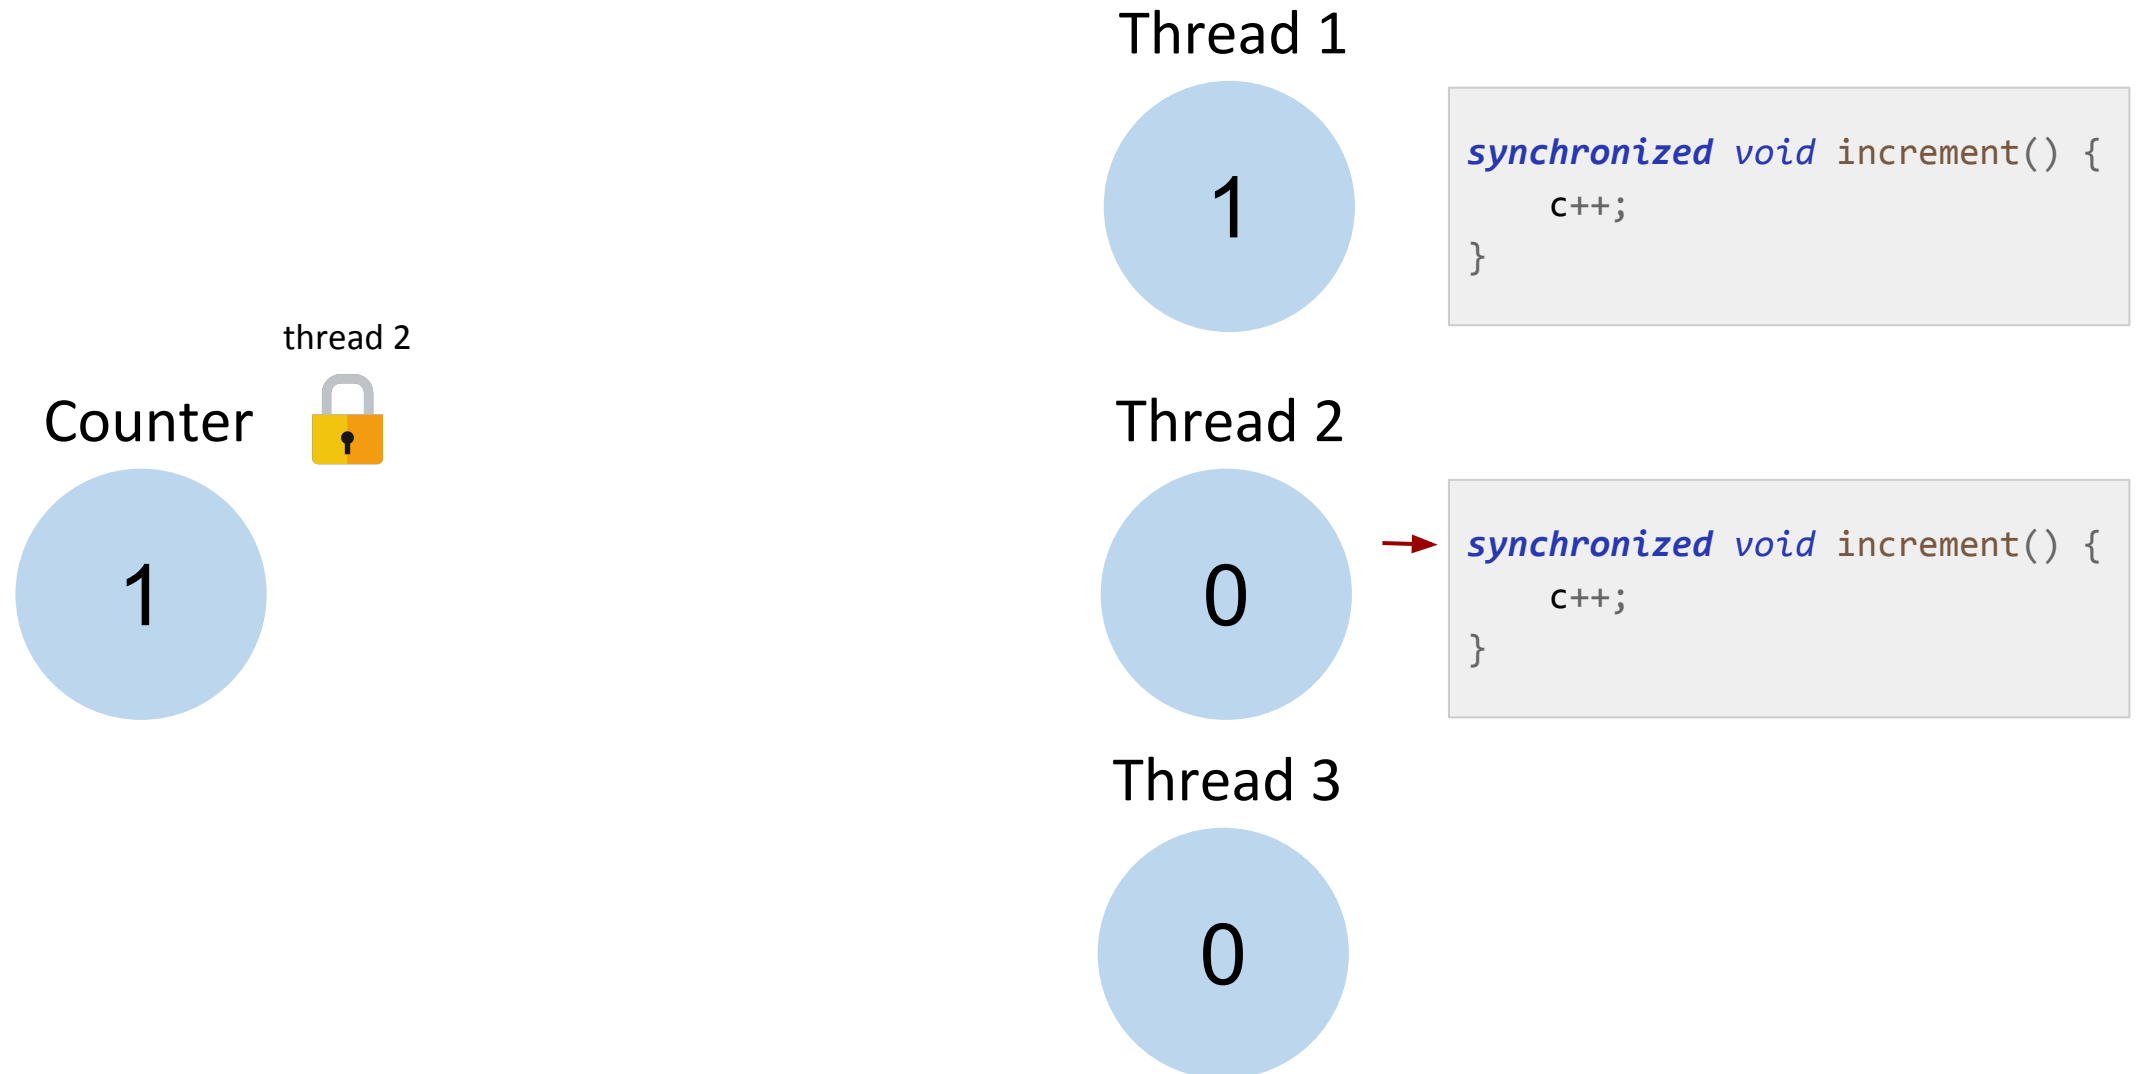
# Task B: SynchronizedCounter



# Task B: SynchronizedCounter



# Task B: SynchronizedCounter



# Task D


- Implement a `FairThreadCounter` that ensures that different threads increment the Counter in an round-robin fashion. That is, two threads with ids 1 and 2 would increment the value in the following order 1, 2, 1, 2, 1, 2, etc. You should implement the scheduling using the `wait` and `notify` methods.
- Can you think of implementation that does not use `wait` and `notify` methods?
- (Optional) Extend your implementation to work with arbitrary number of threads (instead of only 2) that increment the counter in round-robin fashion.

# Wait and Notify Recap


- Object provides `wait()` and `notify()` methods
- To call `wait()` on an object thread must own its lock
- Thread releases the lock and is added to the “waiting list” for that object
- Thread waits until a `notify` method is called on the object
- `notify()` removes one (arbitrary) thread from the object’s “waiting list”
- `notifyAll()` removes all the threads

Thread 1 must increment first!


Counter




Thread 1



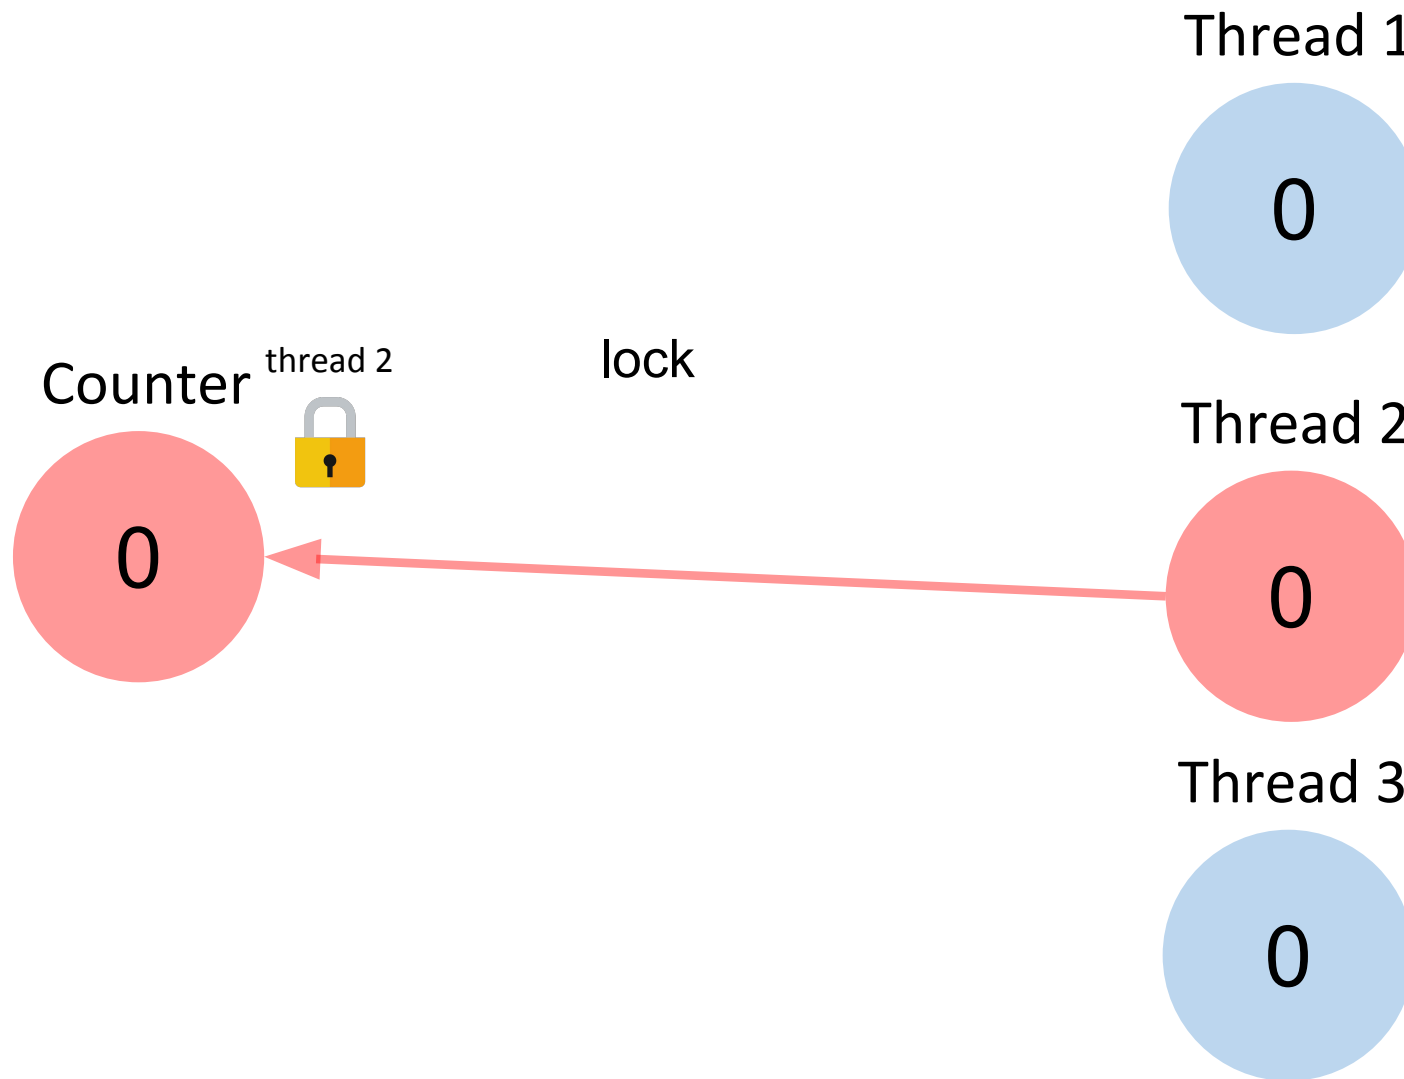
Thread 2




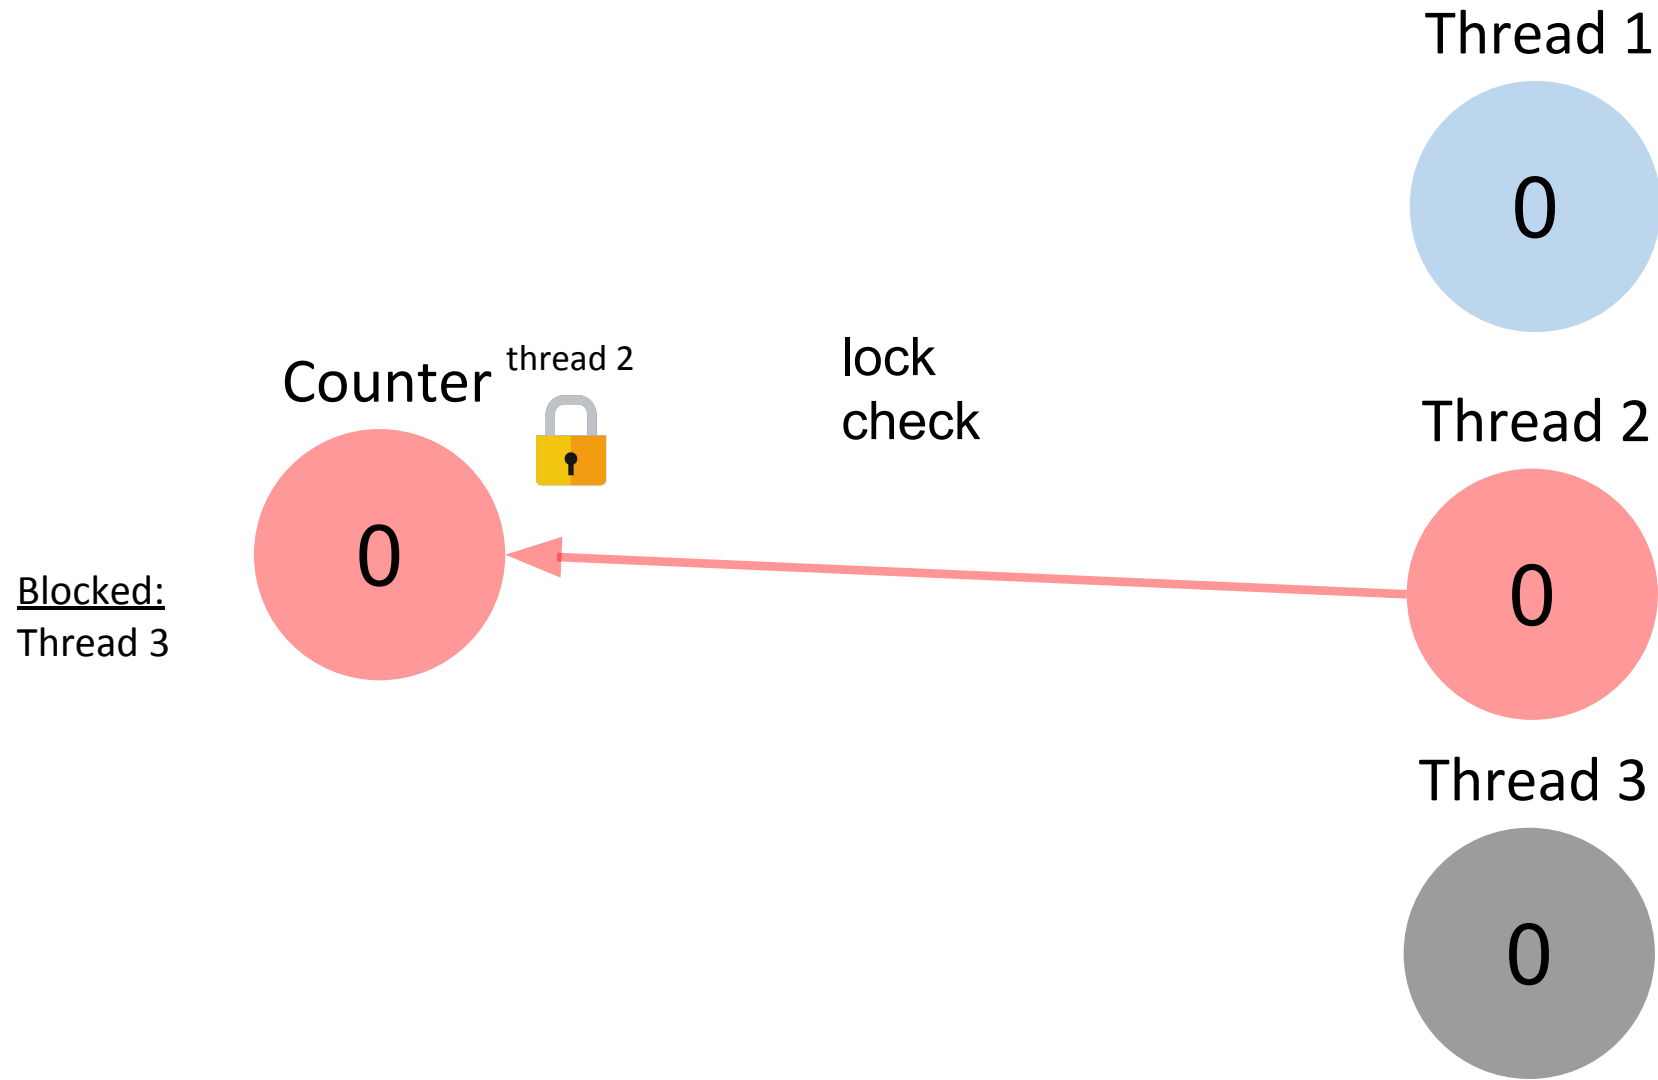
Thread 3

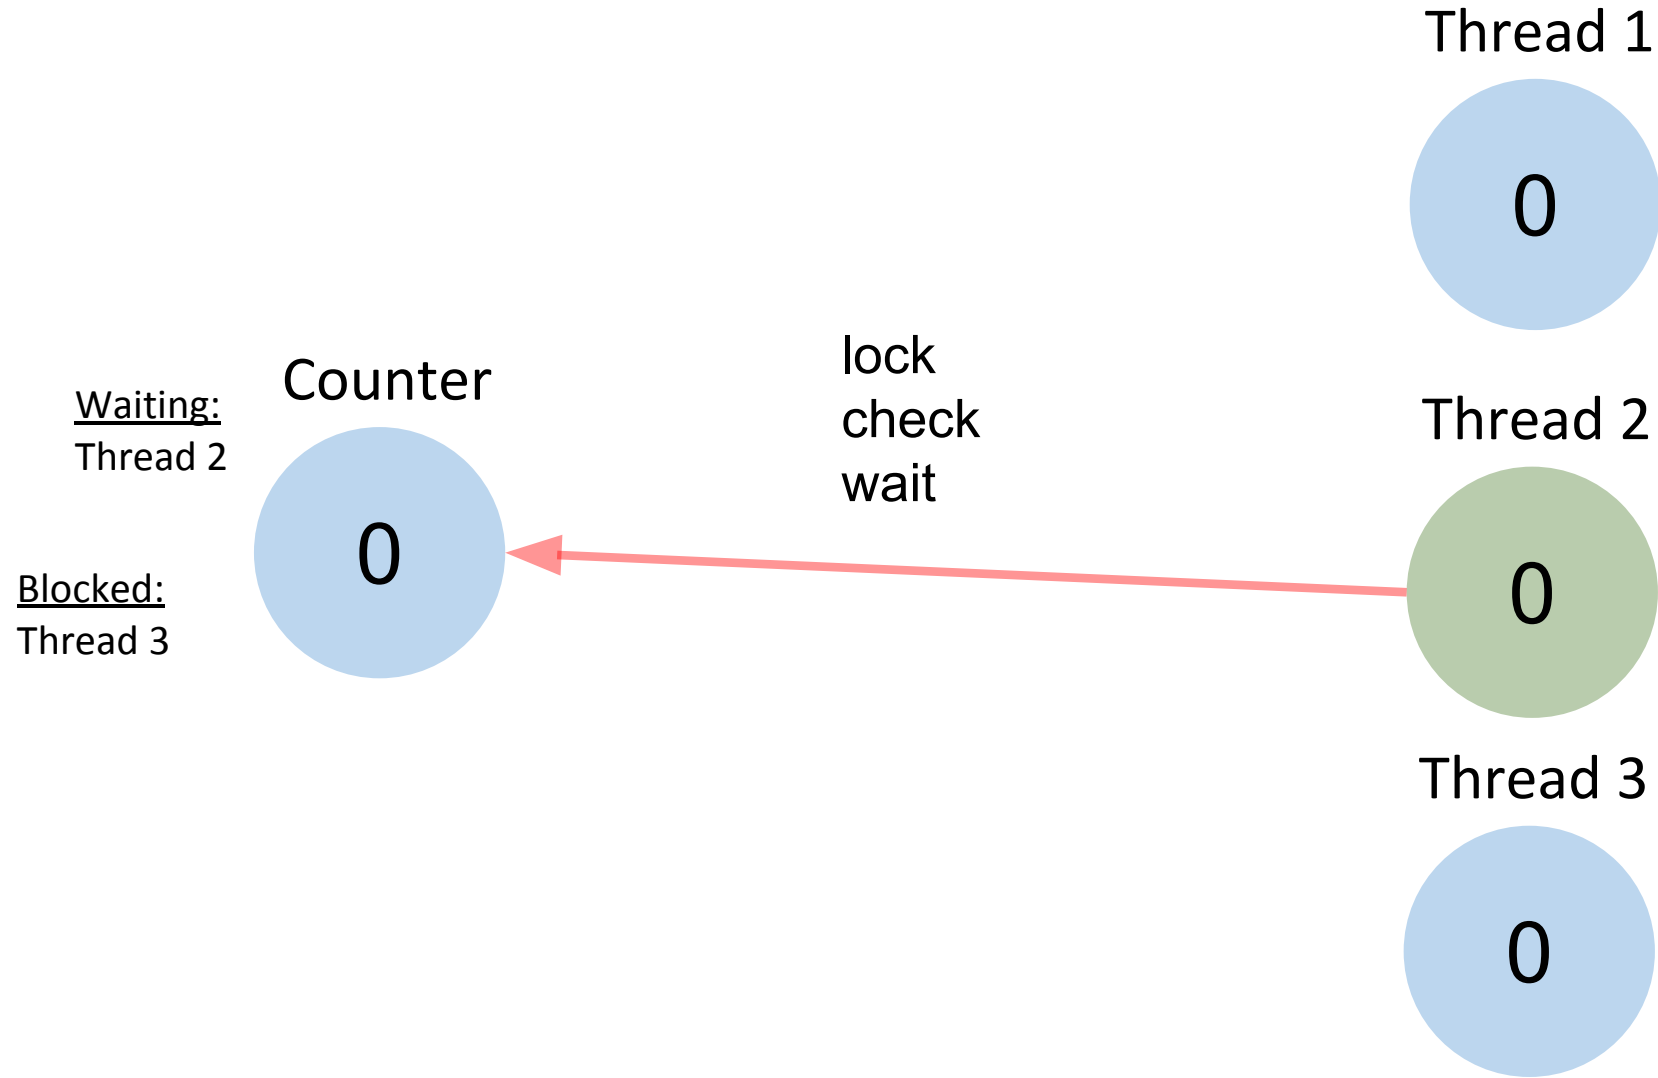


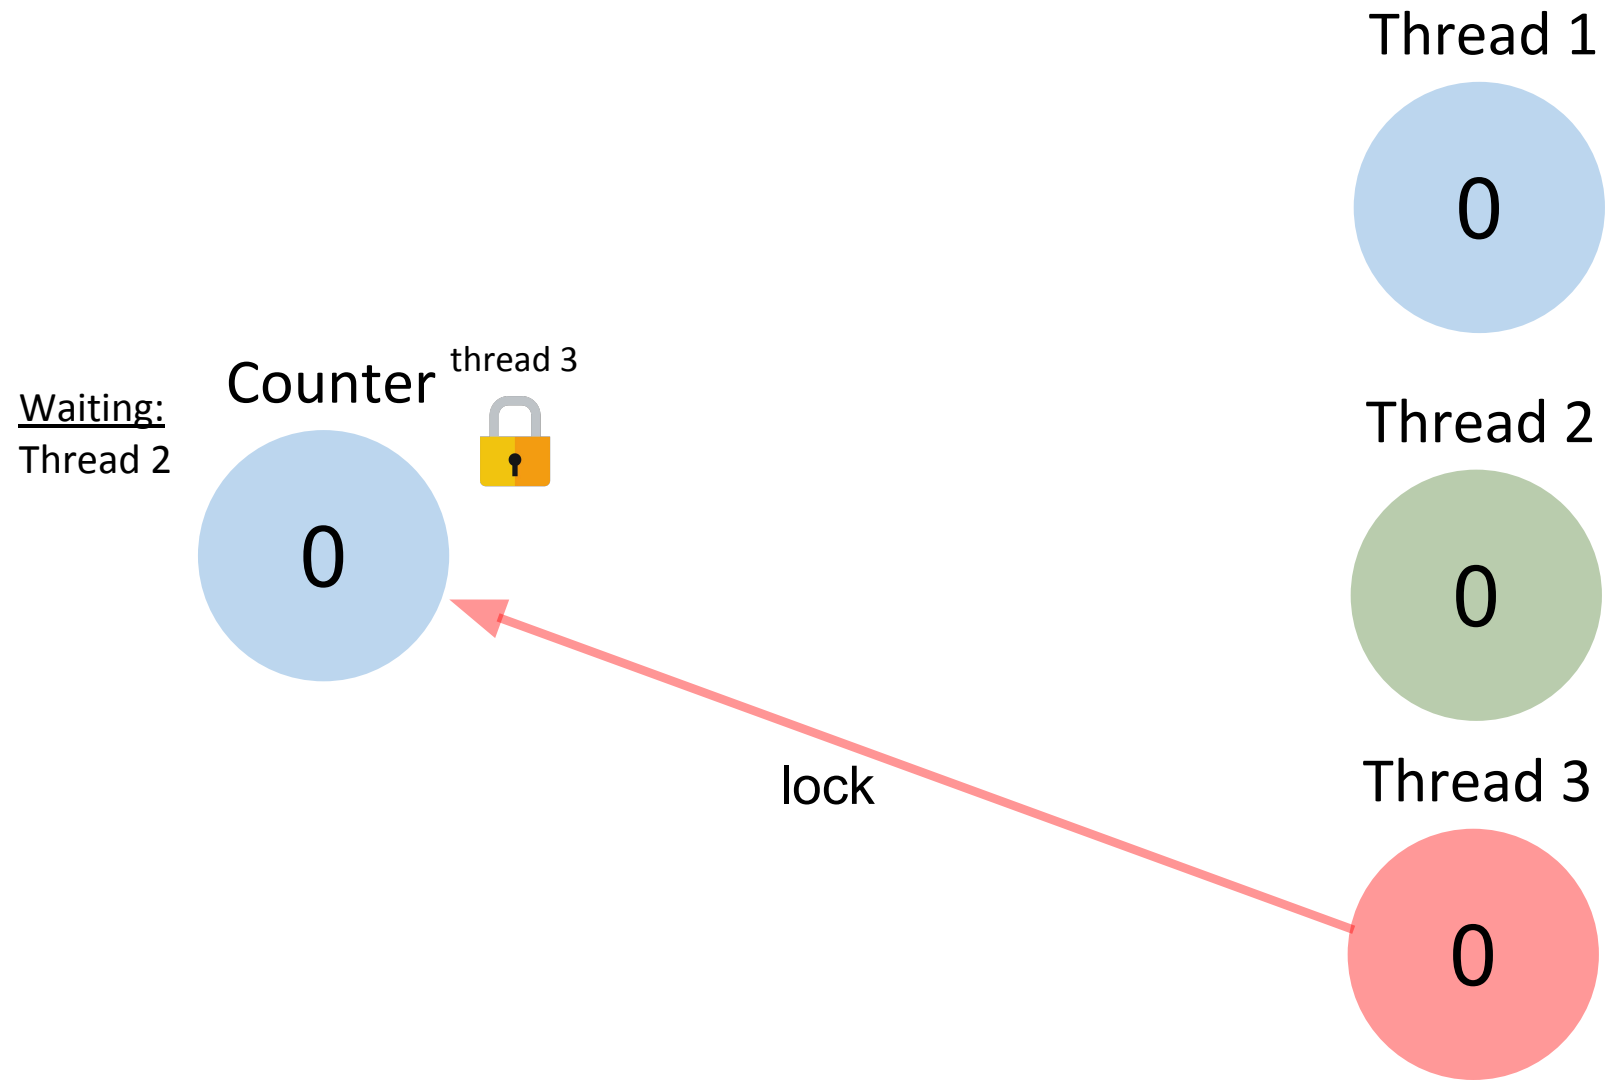


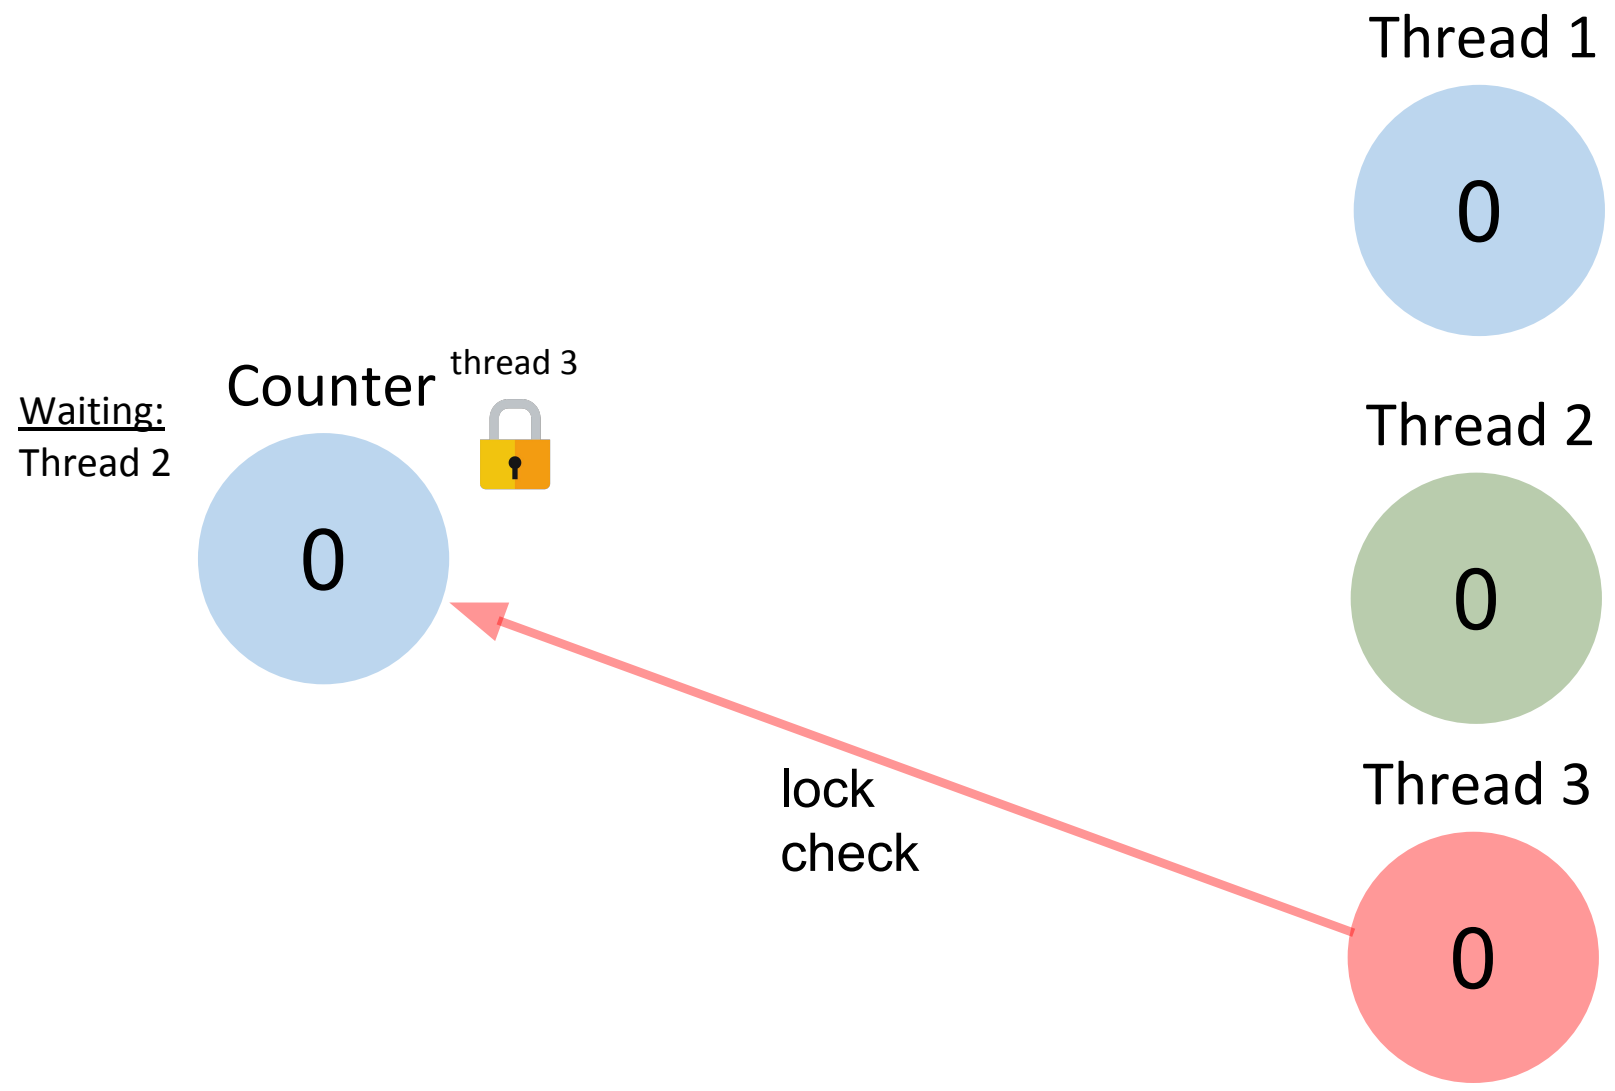


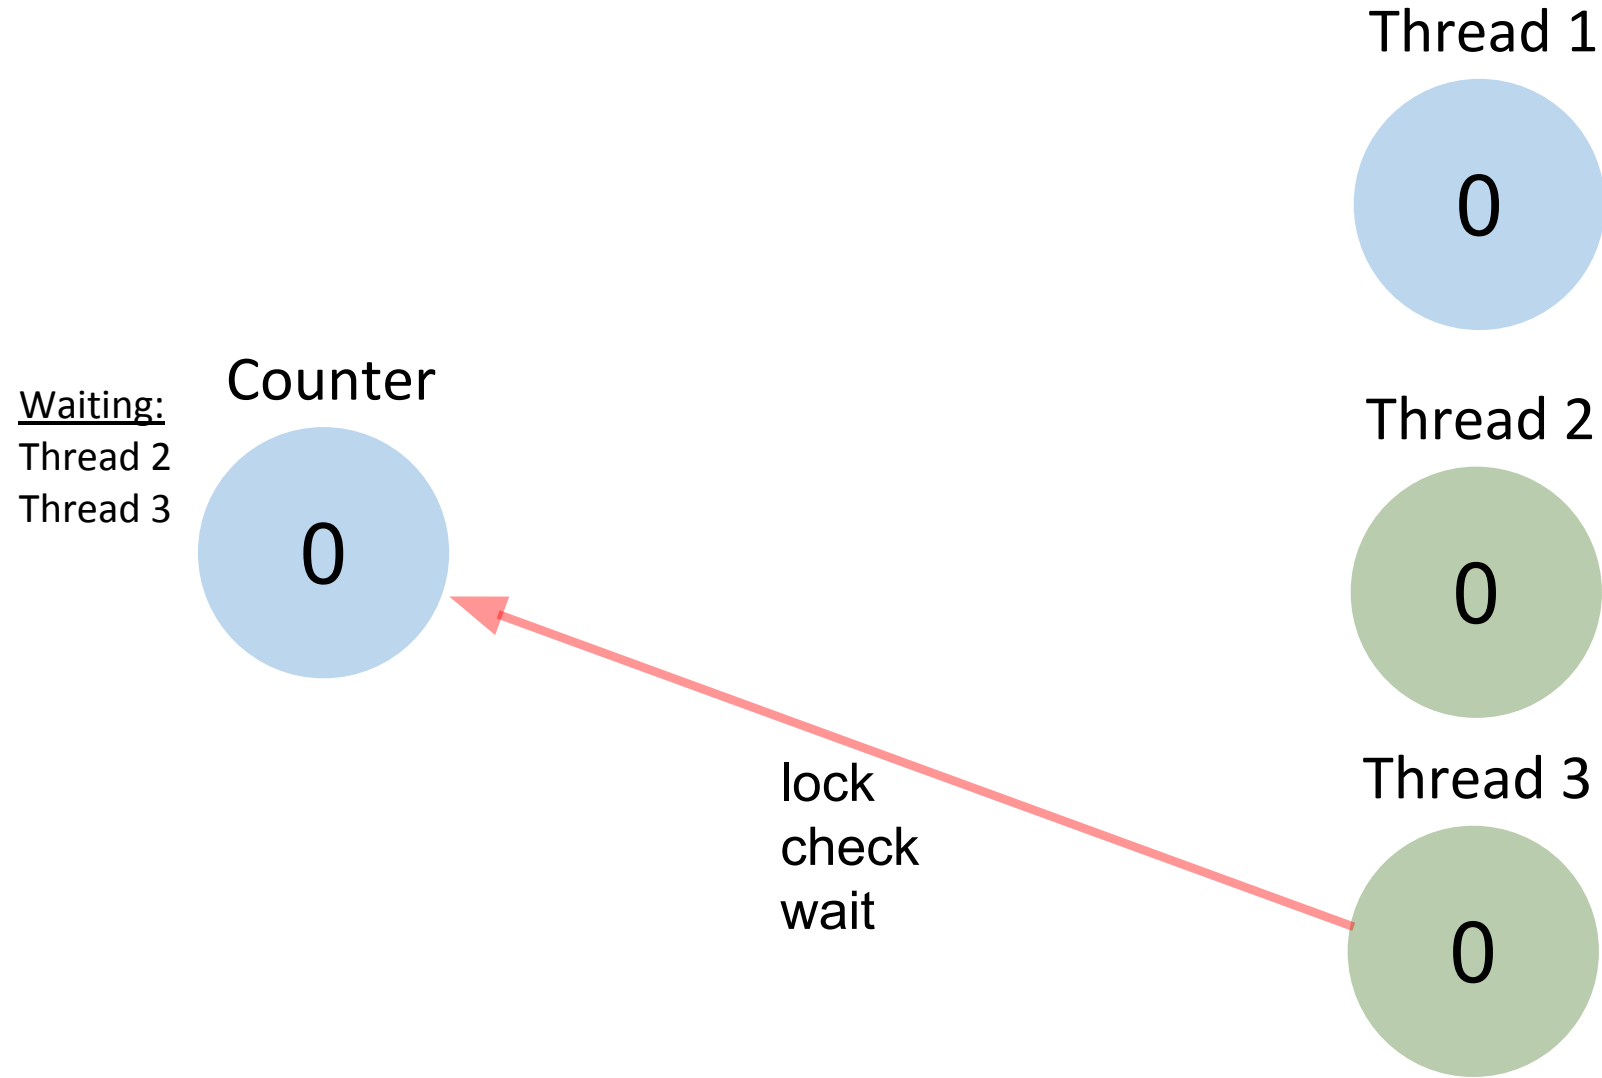


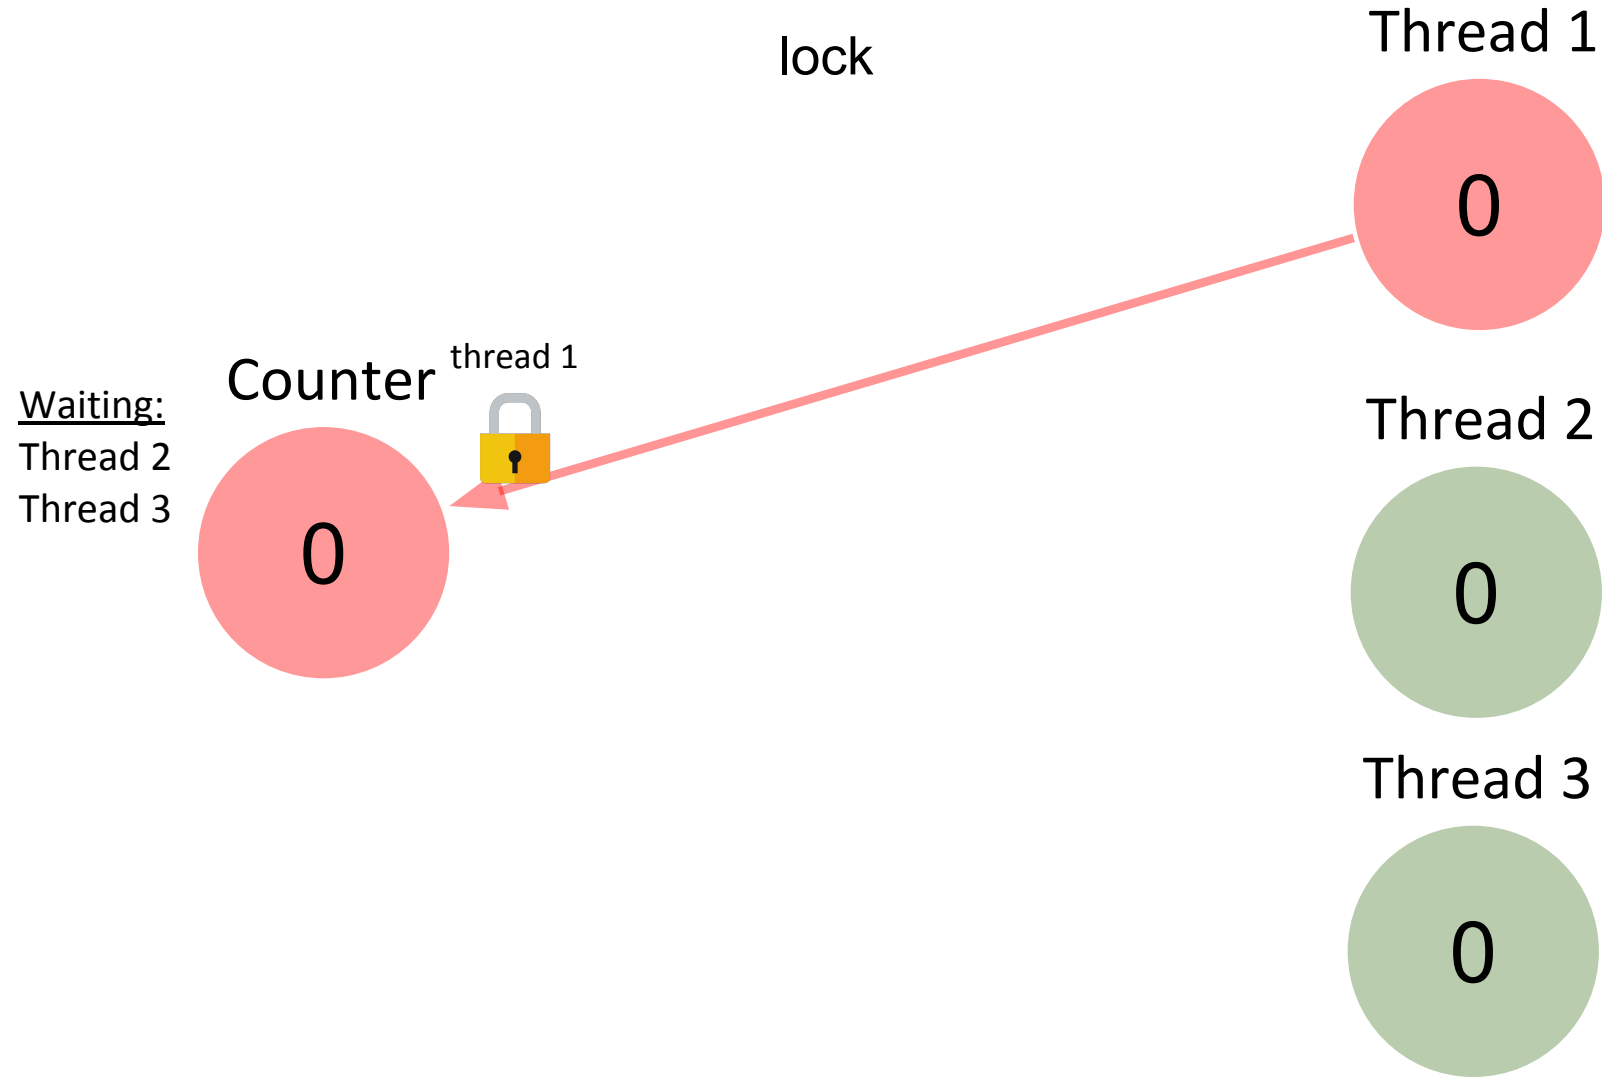




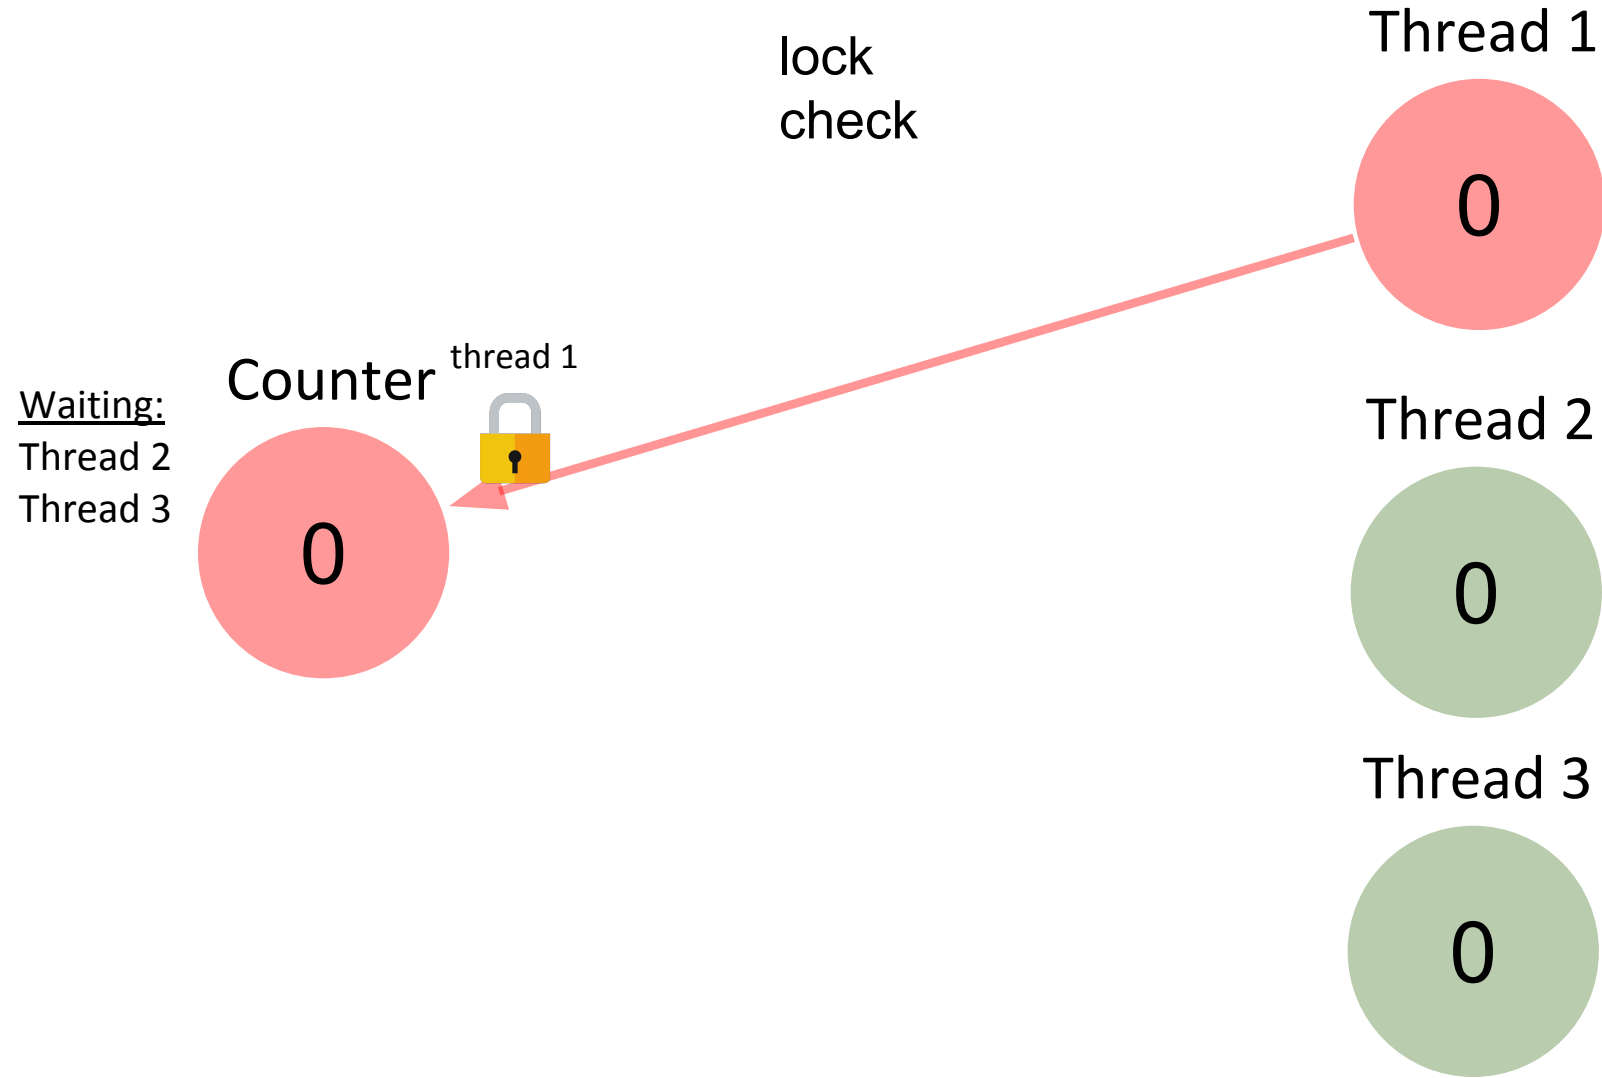


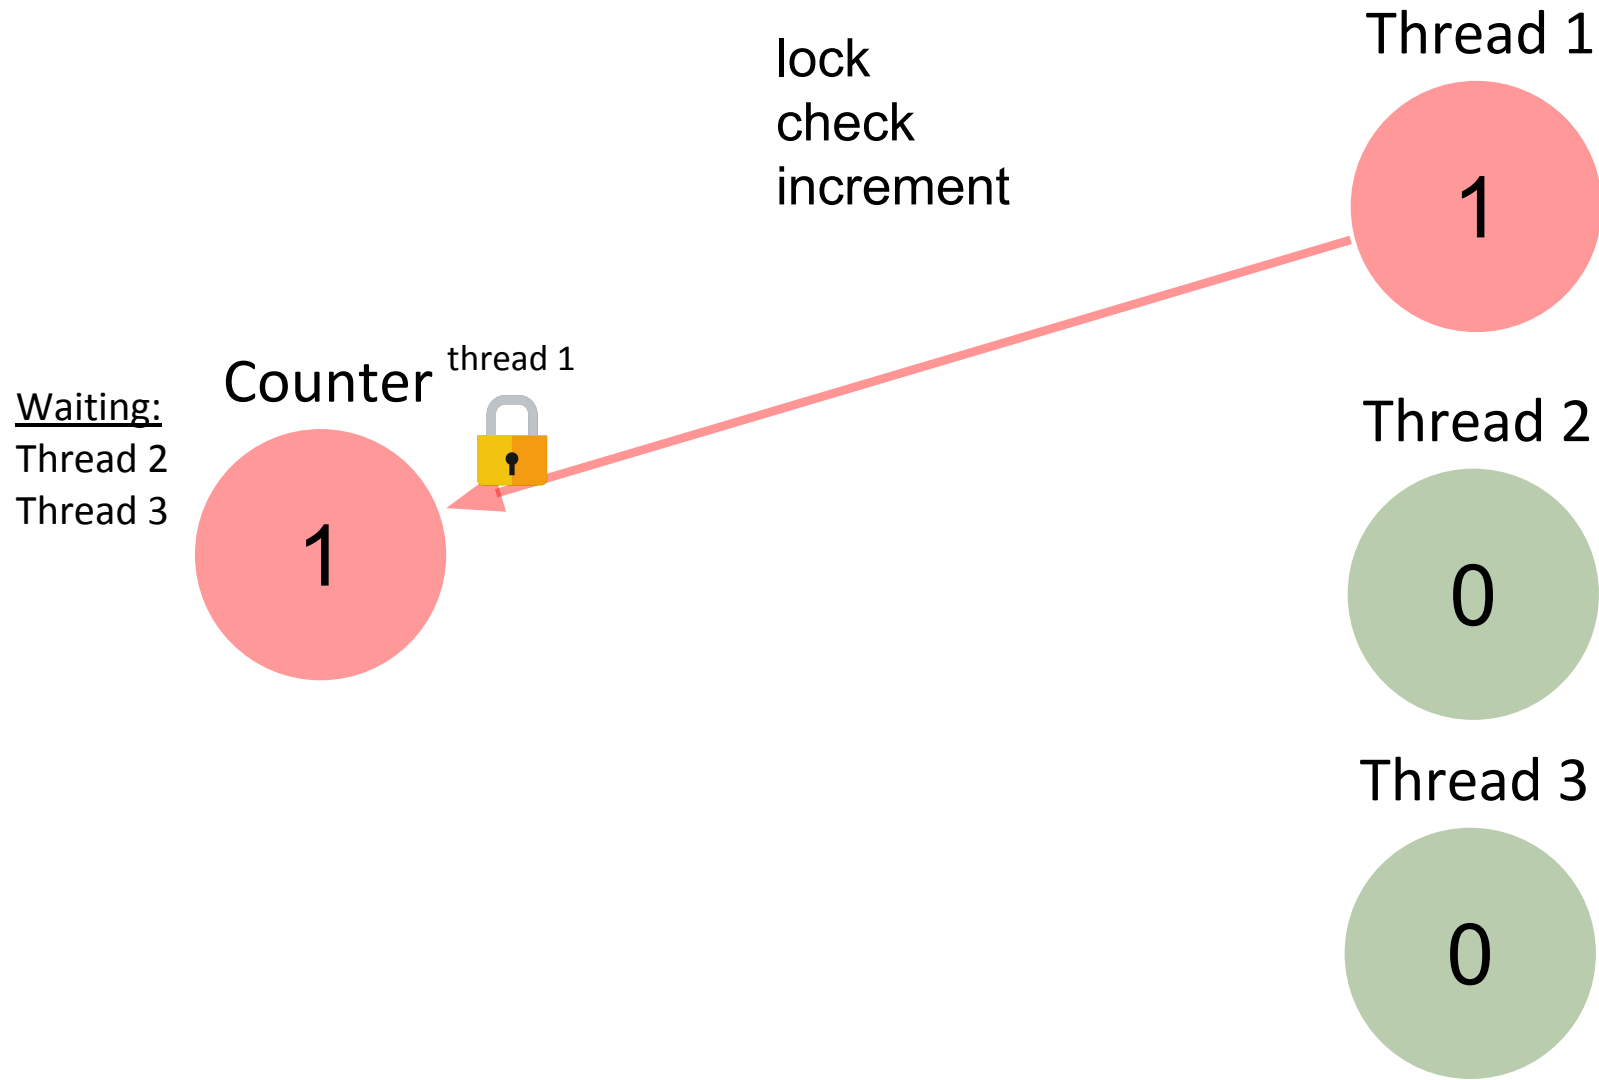


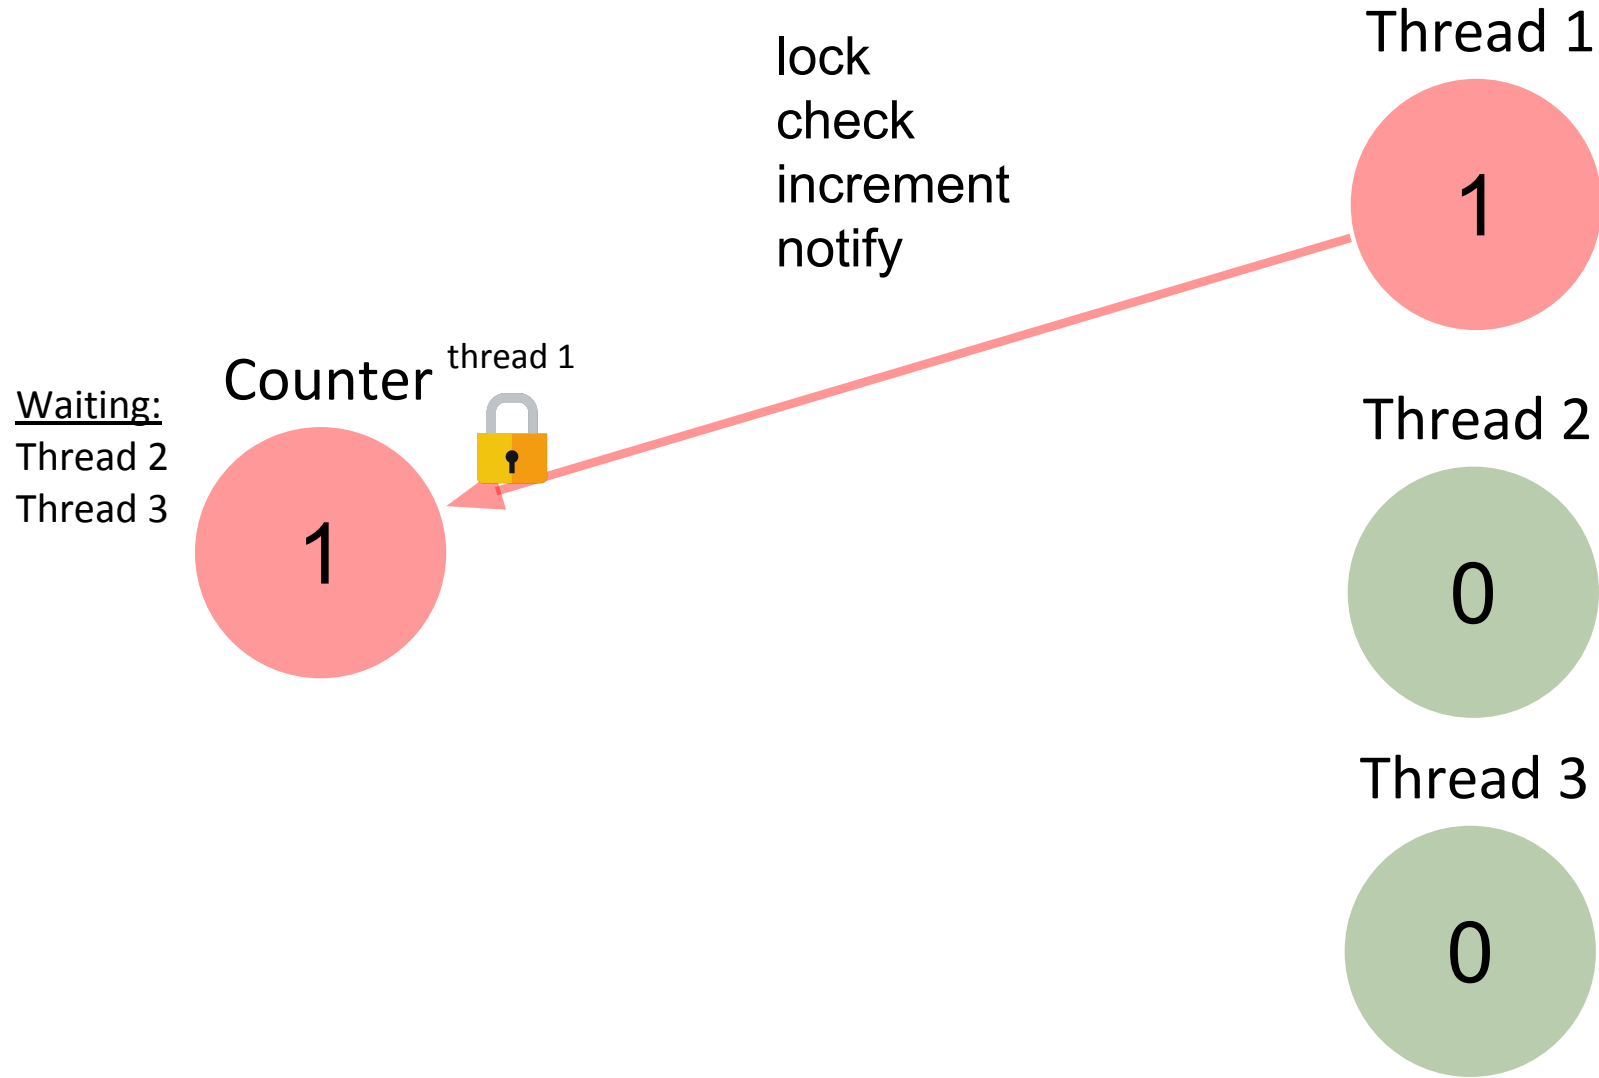


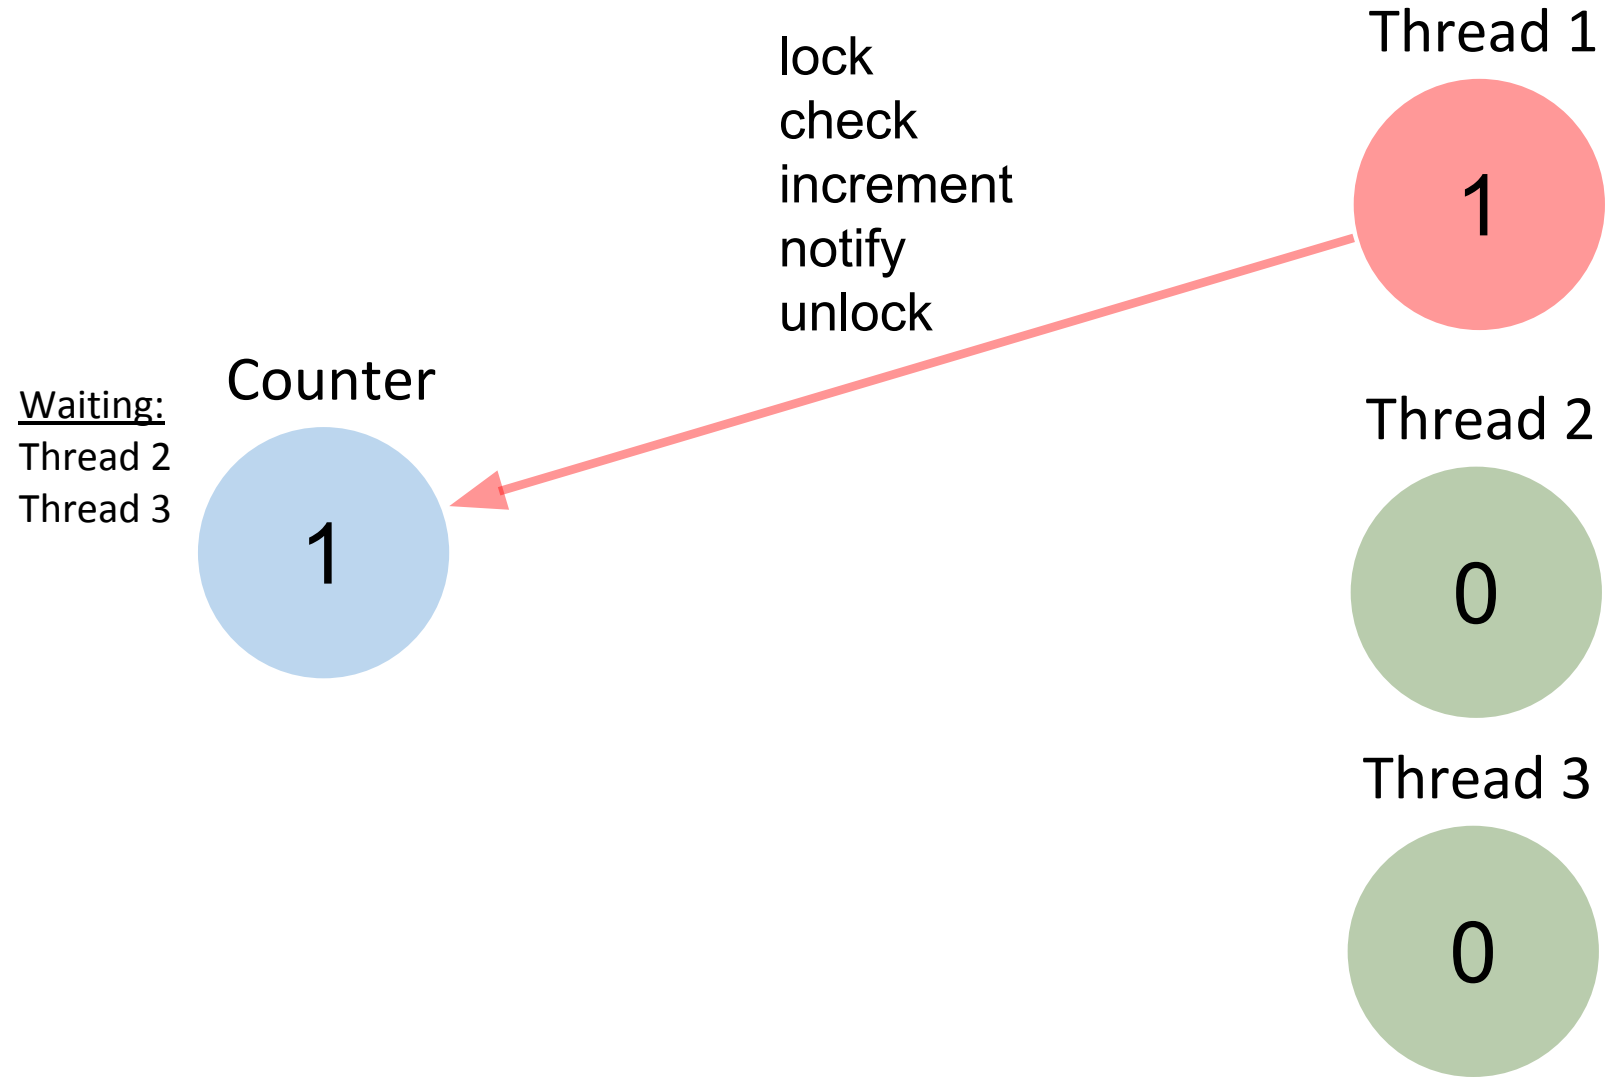


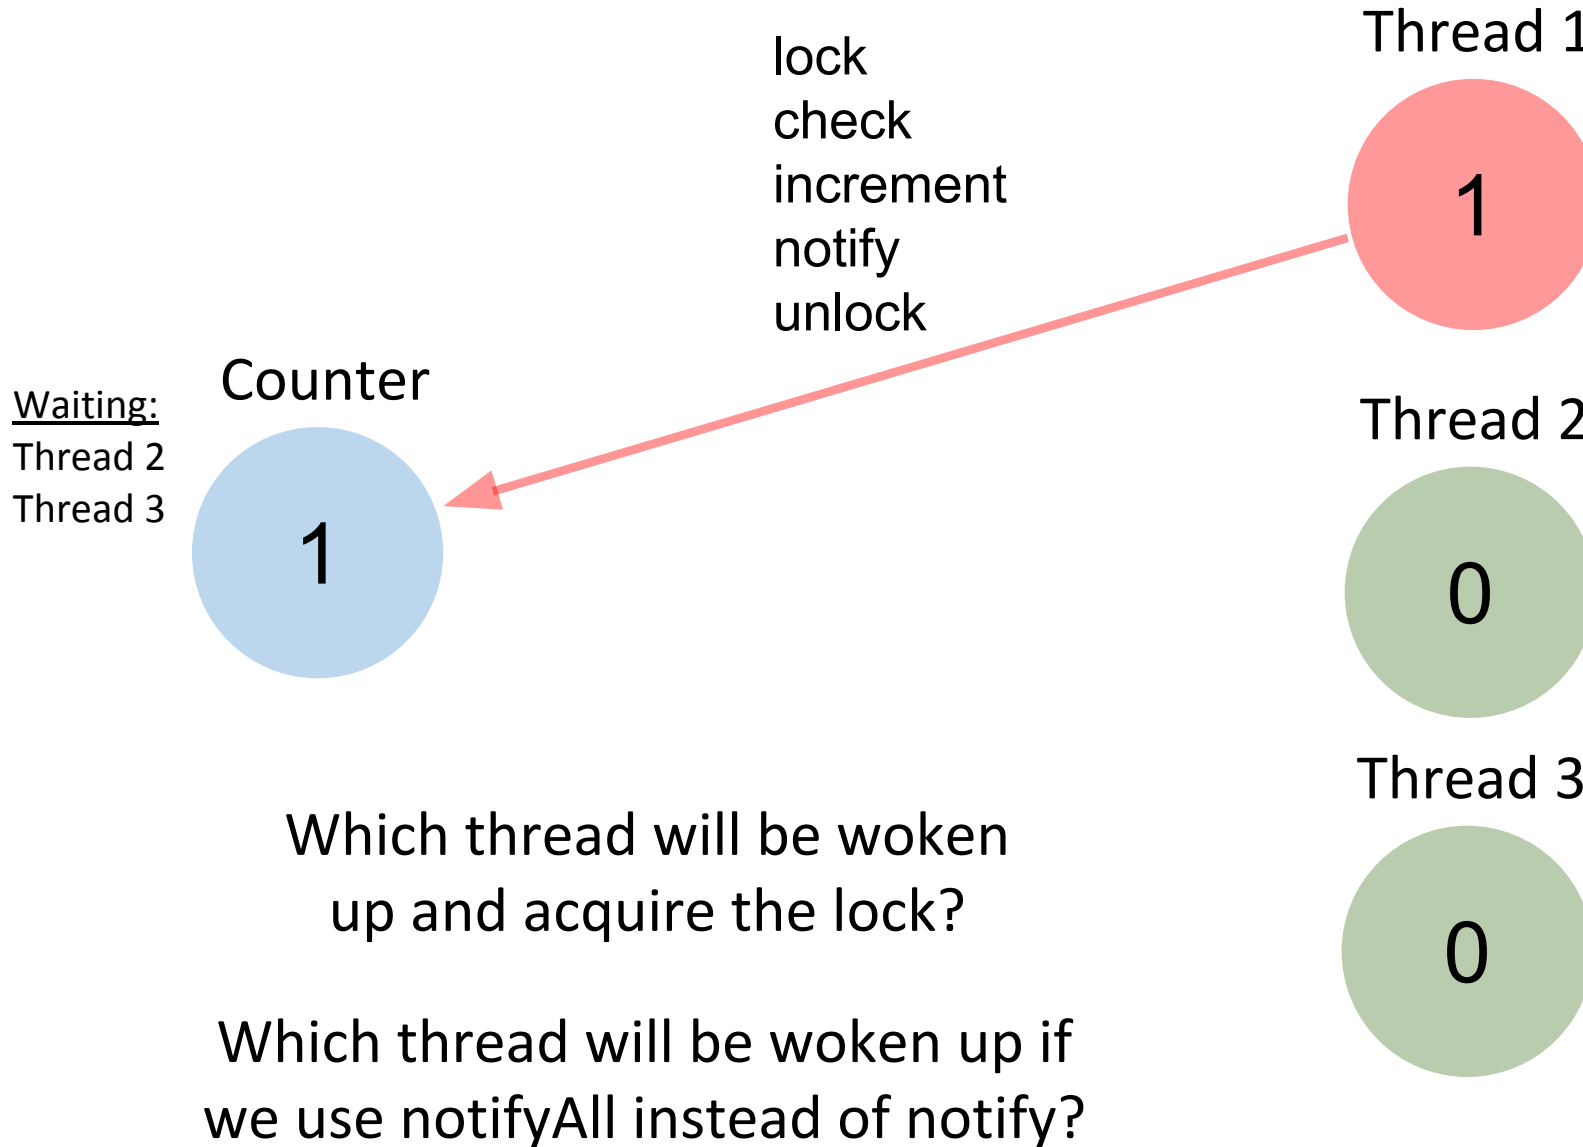












# How to find the difference between notify vs notifyAll?

## notify

```
public final void 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.

The choice is arbitrary and occurs at the discretion of the implementation.

A thread waits on an object's monitor by calling one of the wait methods.

## notifyAll

```
public final void notifyAll()
```

Wakes up all threads that are waiting on this object's monitor. A thread waits on an object's monitor by calling one of the wait methods.

<https://docs.oracle.com/javase/7/docs/api/java/lang/Object.html>

# Task E: AtomicCounter (Optional)

```
public class AtomicCounter implements Counter {  
  
    public void increment() {  
        ??  
    }  
  
    public int value() {  
        ??  
    }  
}
```

# Task E: AtomicCounter (Optional)

```
public class AtomicCounter implements Counter {
    private AtomicInteger c = new AtomicInteger(0);

    public void increment() {
        c.incrementAndGet();
    }

    public int value() {
        return c.get();
    }
}
```



# Task E: AtomicCounter (Optional)

```
public class AtomicCounter implements Counter {
    private AtomicInteger c = new AtomicInteger(0);

    public void increment() {
        c.incrementAndGet();
    }

    public int value() {
        return c.get();
    }
}
```

**What is the difference?**

int

AtomicInteger

c++;

c.incrementAndGet();

# Task E: AtomicCounter (Optional)

```
public class AtomicCounter implements Counter {
    private AtomicInteger c = new AtomicInteger(0);

    public void increment() {
        c.incrementAndGet();
    }

    public int value() {
        return c.get();
    }
}
```

An operation is atomic if no other thread can see it partly executed. Atomic as in “appears indivisible”.

However does not mean it’s implemented as single instruction.

## What is the difference?

int  
1. load c → 0  
2. c + 1 → 1  
3. store c ← 1  
← c++;  
**not atomic**

AtomicInteger  
c.incrementAndGet(); →  
**atomic**

```
incrementAndGet
public final int incrementAndGet()
Atomically increments by one the current value.
Returns:
the updated value
```

# Exercise 4: Pipelining Recap

# Pipelining: Main Concepts Recap

**Latency**

**Throughput**

**Balanced/Unbalanced Pipeline**

# Pipelining: Main Concepts Recap

## **Latency**

time needed to perform a given computation  
(e.g., process a customer)

## **Throughput**

## **Balanced/Unbalanced Pipeline**

# Pipelining: Main Concepts Recap

## **Latency**

time needed to perform a given computation  
(e.g., process a customer)

## **Throughput**

amount of work that can be done by a system in a given period of time  
(e.g., how many customers can be processed in one minute)

## **Balanced/Unbalanced Pipeline**

# Pipelining: Main Concepts Recap

## **Latency**

time needed to perform a given computation  
(e.g., process a customer)

## **Throughput**

amount of work that can be done by a system in a given period of time  
(e.g., how many customers can be processed in one minute)

## **Balanced/Unbalanced Pipeline**

a pipeline is balanced if it has constant latency

# Library

Over at UZH the law students have been tasked with writing a legal essay about the philosophy of Swiss law. In order to write the essay, each student needs to read four different books on the subject, denoted as A, B, C and D (in this order).

This exercise is created by Lasse Meinen and part of the unofficial VIS Prüfungsvorbereitungsworkshop Skripts available at:

<https://vis.ethz.ch/de/services/pvw-scripts/>

Every student takes the exact same amount of time to read a book, concretely:

1) Reading book **A** takes 80 minutes

2) Reading book **B** takes 40 minutes

3) Reading book **C** takes 120 minutes

4) Reading book **D** takes 40 minutes



# Library

Over at UZH the law students have been tasked with writing a legal essay about the philosophy of Swiss law. In order to write the essay, each student needs to read four different books on the subject, denoted as A, B, C and D (in this order).

**Question 1:** Let's assume all law students are a bit too competitive and don't return any books before they're done reading all of them. How long will it take for 4 students until all of them have started writing their essays?

Every student takes the exact same amount of time to read a book, concretely:

1) Reading book **A** takes 80 minutes

2) Reading book **B** takes 40 minutes

3) Reading book **C** takes 120 minutes

4) Reading book **D** takes 40 minutes

# Library



**Question 1:** Let's assume all law students are a bit too competitive and don't return any books before they're done reading all of them. How long will it take for 4 students until all of them have started writing their essays?

Every student takes the exact same amount of time to read a book, concretely:


1) Reading book **A** takes 80 minutes

2) Reading book **B** takes 40 minutes

3) Reading book **C** takes 120 minutes

4) Reading book **D** takes 40 minutes

# Library



**Question 2:** The library introduces a "one book at a time" policy, i.e. the students have to return a book before they can start on the next one. How long will it now take for 4 students until all of them have started writing their essays?

Every student takes the exact same amount of time to read a book, concretely:

1) Reading book **A** takes 80 minutes

2) Reading book **B** takes 40 minutes

3) Reading book **C** takes 120 minutes

4) Reading book **D** takes 40 minutes

# Library

For this pipeline, latency makes sense only if asked for a particular student, not for the whole pipeline.

Latency?



**Question 2:** The library introduces a "one book at a time" policy, i.e. the students have to return a book before they can start on the next one. How long will it now take for 4 students until all of them have started writing their essays?

Every student takes the exact same amount of time to read a book, concretely:

1) Reading book **A** takes 80 minutes

3) Reading book **C** takes 120 minutes

2) Reading book **B** takes 40 minutes

4) Reading book **D** takes 40 minutes

# Library



**Question 2:** The library introduces a "one book at a time" policy, i.e. the students have to return a book before they can start on the next one. How long will it now take for 4 students until all of them have started writing their essays?

Every student takes the exact same amount of time to read a book, concretely:


1) Reading book **A** takes 80 minutes

2) Reading book **B** takes 40 minutes

3) Reading book **C** takes 120 minutes

4) Reading book **D** takes 40 minutes

# Library



The pipeline is not balanced  
since the latency is not constant

Every student takes the exact same amount of time to read a book, concretely:

1) Reading book **A** takes 80 minutes

2) Reading book **B** takes 40 minutes

3) Reading book **C** takes 120 minutes

4) Reading book **D** takes 40 minutes

# Exercise 4

# Task 1 - Pipelining

Bob, Mary, John and Alice



50 min



90 min



15 min

- a) Laundry time using sequential order
- b) Design a strategy with better laundry time
- c) How would the laundry time improve if they bought a new dryer?



# Task 2 - Pipelining II

Assume a processor that can each cycle issue either:

- one multiplication instruction with latency 6 cycles
- one addition instruction with latency 3 cycles

How many cycles are required to execute following loops?

```
for (int i = 0; i < data.length; i++) {  
    data[i] = data[i] * data[i];  
}
```

```
for (int i = 0; i < data.length; i += 2) {  
    j = i + 1;  
    data[i] = data[i] * data[i];  
    data[j] = data[j] * data[j];  
}
```

```
for (int i = 0; i < data.length; i += 4) {  
    j = i + 1;  
    k = i + 2;  
    l = i + 3;  
    data[i] = data[i] * data[i];  
    data[j] = data[j] * data[j];  
    data[k] = data[k] * data[k];  
    data[l] = data[l] * data[l];  
}
```

# Task 3 - Identify Potential Parallelization

Can we parallelize following two loops using parallel for construct?

```
for (int i=1; i<size; i++) { // for loop: i from 1 to (size-1)
    if (data[i-1] > 0) // If the previous value is positive
        data[i] = (-1)*data[i]; // change the sign of this value
} // end for loop
```

```
for (int i=0; i<size; i++) { // for loop: i from 0 to (size-1)
    data[i] = Math.sin(data[i]); // calculate sin() of the value
} // end for loop
```