

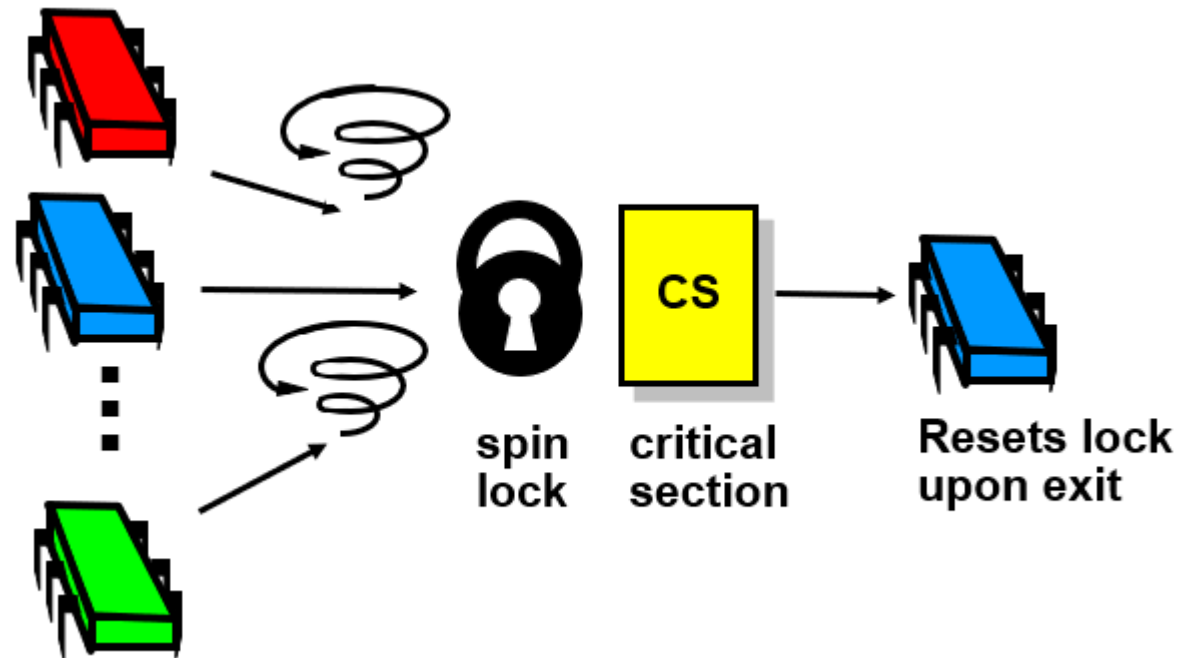
The background features abstract, overlapping geometric shapes in various shades of blue, ranging from light sky blue to dark navy blue. These shapes are primarily located on the left and right sides of the frame, creating a modern, layered effect. The central area is a plain white space where the text is located.

Spin Locks and Contention

The goal

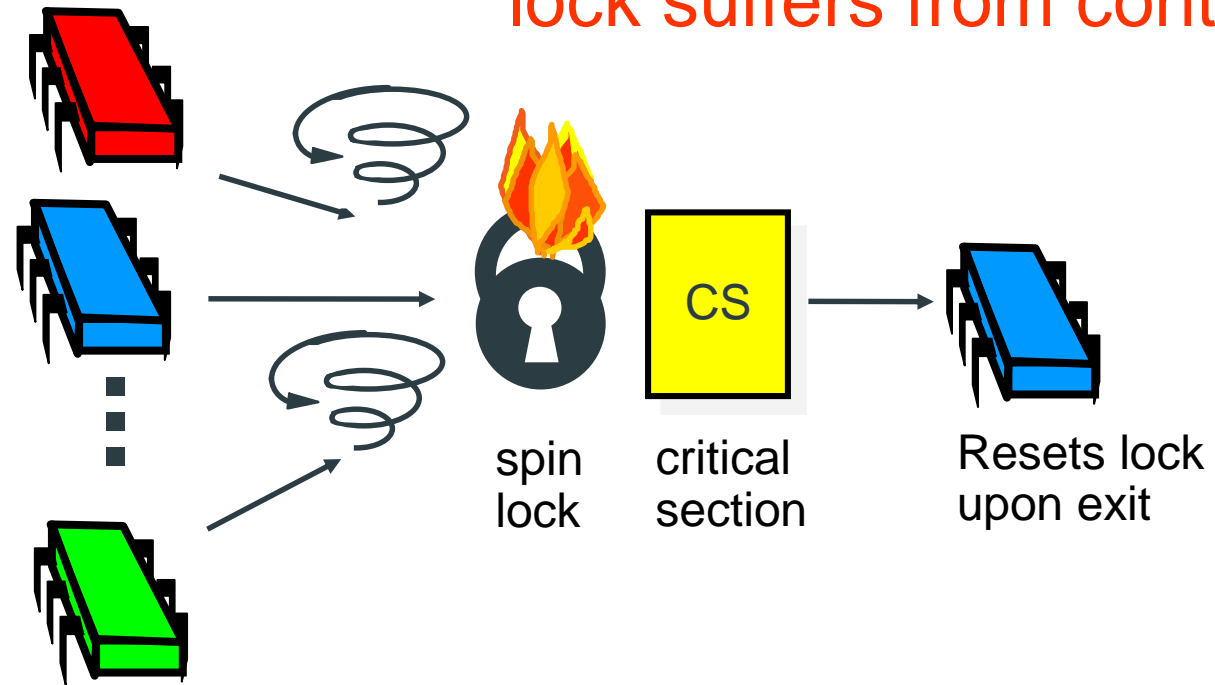
The goal of this lecture is to understand how architecture affects performance, and how to exploit this knowledge to write efficient concurrent programs.

Basic Spin-Lock



Basic Spin-Lock



lock suffers from contention



Spin lock and blocking

- ▶ Spin lock- repeatedly testing the lock. Since the thread remains active but is not performing a useful task, the use of such a lock is a kind of busy waiting.
- ▶ blocking- suspend yourself and ask the operating system's scheduler to schedule another thread on your processor. blocking makes sense if you expect the lock delay to be long.

Mutex lock-code

```
1  Lock mutex = new LockImpl(...); // lock implementation
2  ...
3  mutex.lock(); 
4  try {
5  ... // body
6  } finally { 
7  mutex.unlock();
8  }
```

Peterson algorithm

Lock()

flag[i]=true

turn=1-i

While (flag[1-i] and turn=1-i)



Unlock()

flag[i]=false

Not providing mutual exclusion in practice

- ▶ Running the Peterson algorithm in practice will provide a slightly off value than expected.
- ▶ It must be that both threads are occasionally in the critical section at the same time, even though we have proved that this cannot happen.

“How often have I said to you that when you have eliminated the impossible, whatever remains, however improbable, must be the truth?”

Why?

- ▶ compilers that reorder instructions to enhance performance - It is therefore possible that the order of writes of `flag[i]` and turn by thread i will be reversed by the compiler.
- ▶ writes to multiprocessor memory do not necessarily take effect when they are issued - it is therefore possible that the order of i 's write to turn is delayed and may arrive in memory only after i reads `flag[j]`.

GetAndSet()

- ▶ Atomic instruction.
- ▶ `getAndSet(b)` that atomically replaces the current value with *b*, and returns the previous value.
- ▶ We will now look at two algorithm that uses this method.

Solution- Test-And-Set Locks

```
1 public class TASLock implements Lock {  
2     AtomicBoolean state = new AtomicBoolean(false);  
3     public void lock() {  
4         while (state.getAndSet(true)) {}  
5     }  
6     public void unlock() {  
7         state.set(false);  
8     } }
```

TASLock algorithm


TTASLock lock algorithm

```
1 public class TTASLock implements Lock {  
2     AtomicBoolean state = new AtomicBoolean(false);  
3     public void lock() {  
4         while (true) {  
5             while (state.get()) {};  
6             if (!state.getAndSet(true))  
7                 return;}}
```

TTASLock un-lock algorithm

```
8 public void unlock() {  
9     state.set(false); } }
```


Schematic performance of a TASLock, a TTASLock, and an ideal lock



TASLock vs. TTASLock

- ▶ Both algorithms are the same.
- ▶ Yet they have different performant- TTASLock perform better then TASLock.
- ▶ Both algorithms preform poorly vs. the ideal performance.

Bus architecture



TASLock vs. TTASLock

TASLock:

- ▶ Each test-and-set call go over the bus.
- ▶ Thread that want to release the lock will be delayed.
- ▶ Threads not waiting to lock can be delayed when access memory

TTASLock:

- ▶ Set calls not use bus if value in cache.
- ▶ Thread that want to release the lock will not be delayed.

TTASLock is still far from ideal because when the lock is free all threads try test-and-set call

Solution- Backoff

- ▶ When lock is free → tries to lock it
 - ▶ If succeed great
 - ▶ Else do backoff

- ▶ We will like that the larger the unsuccessful tries the longer the backoff

Backoff

```
1 public void backoff() throws  
  InterruptedException {  
2   int delay = random.nextInt(limit);  
3   limit = Math.min(maxDelay, 2 * limit);  
4   Thread.sleep(delay);  
5 }
```


Backoff lock

```
1 public class BackoffLock implements Lock {
2 private AtomicBoolean state = new AtomicBoolean(false);
3 public void lock() {
4     Backoff backoff = new Backoff(MIN_DELAY, MAX_DELAY);
5     while (true) {
6         while (state.get()) {};
7         if (!state.getAndSet(true)) {
8             return;
9         } else {
10            backoff.backoff();
11        }
12    }
13 }
```

Backoff un-lock

```
12 public void unlock() {  
13     state.set(false);  
14 }
```

Performance graph



Backoff algorithm

Pros:

- ▶ Better than TTASLock in contention (less contention).
- ▶ Easy to implement.


Cons:

- ▶ All threads spin on the same shared location causing cache-coherence traffic on every successful lock access.
- ▶ Threads delay longer than necessary, causing the critical section to be underutilized.
- ▶ Parameters need to be chosen carefully.


Solution- queue lock

- ▶ each thread spin on a different cell in the queue.
- ▶ Thread is notified when it is it's time to go to the critical section through the queue.

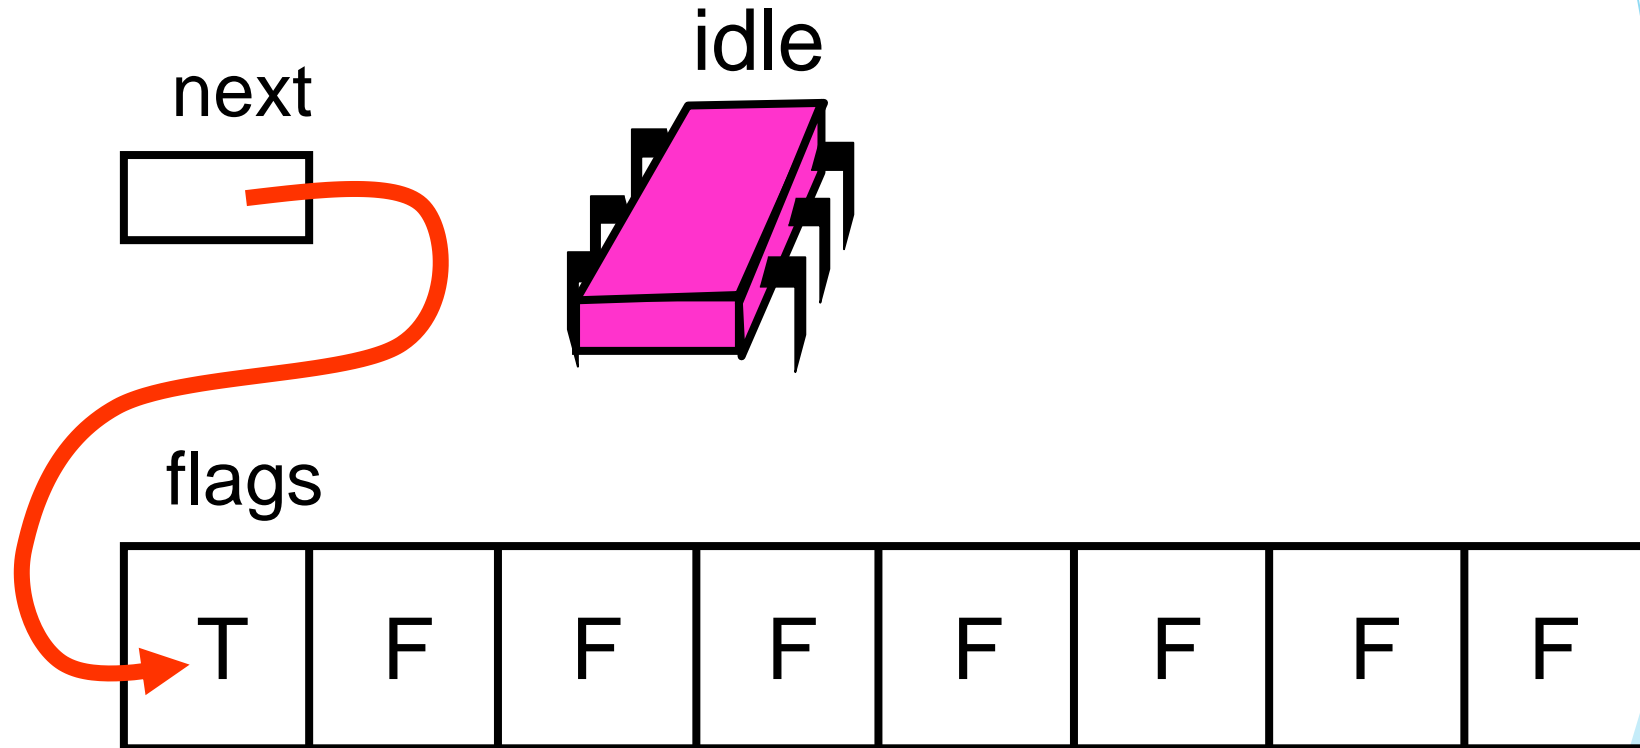
Queue Lock



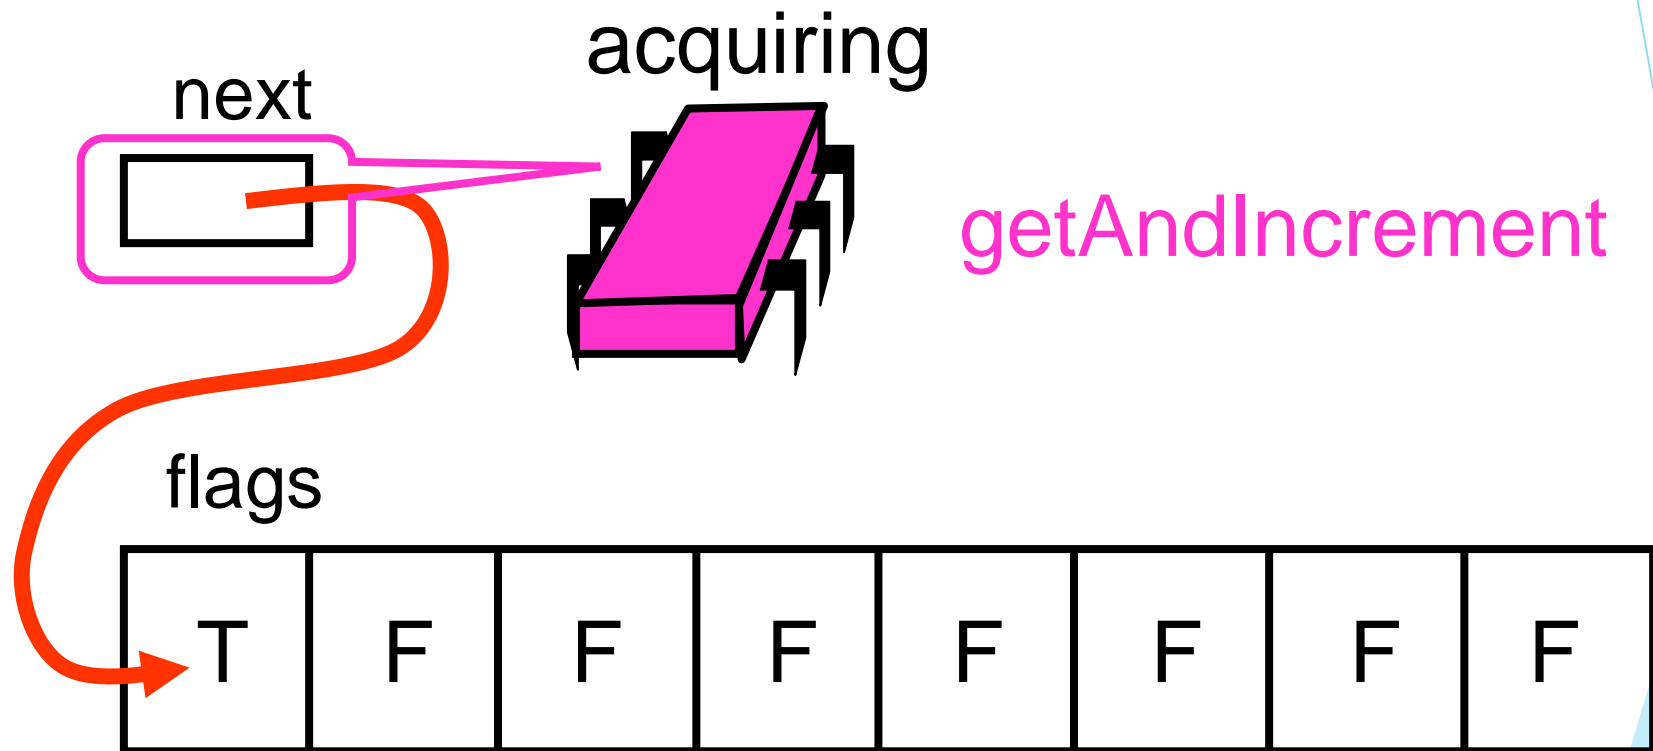
Queue Lock



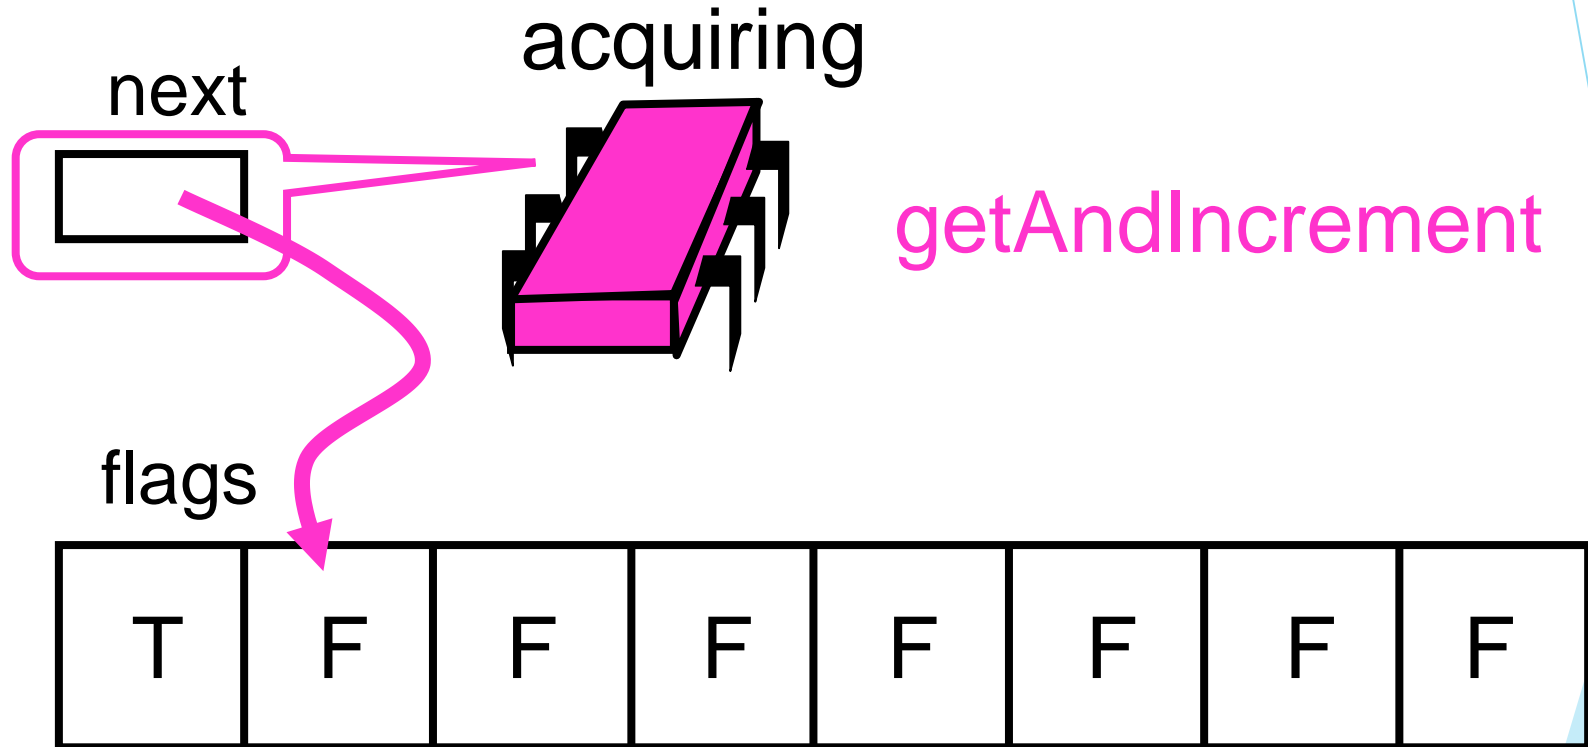
Queue Lock




Queue Lock



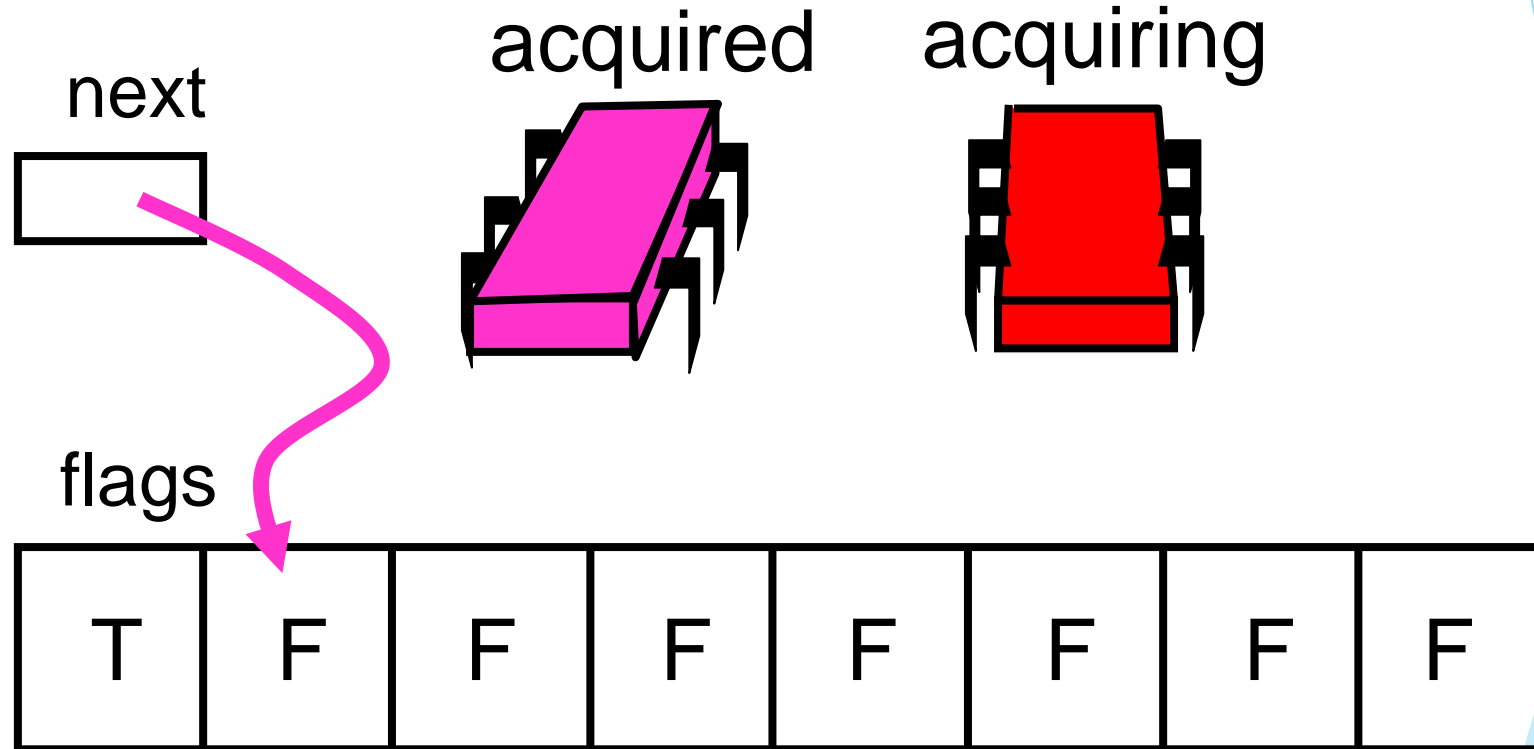
Queue Lock




Queue Lock




Queue Lock




Queue Lock




Queue Lock




Queue Lock



Queue Lock



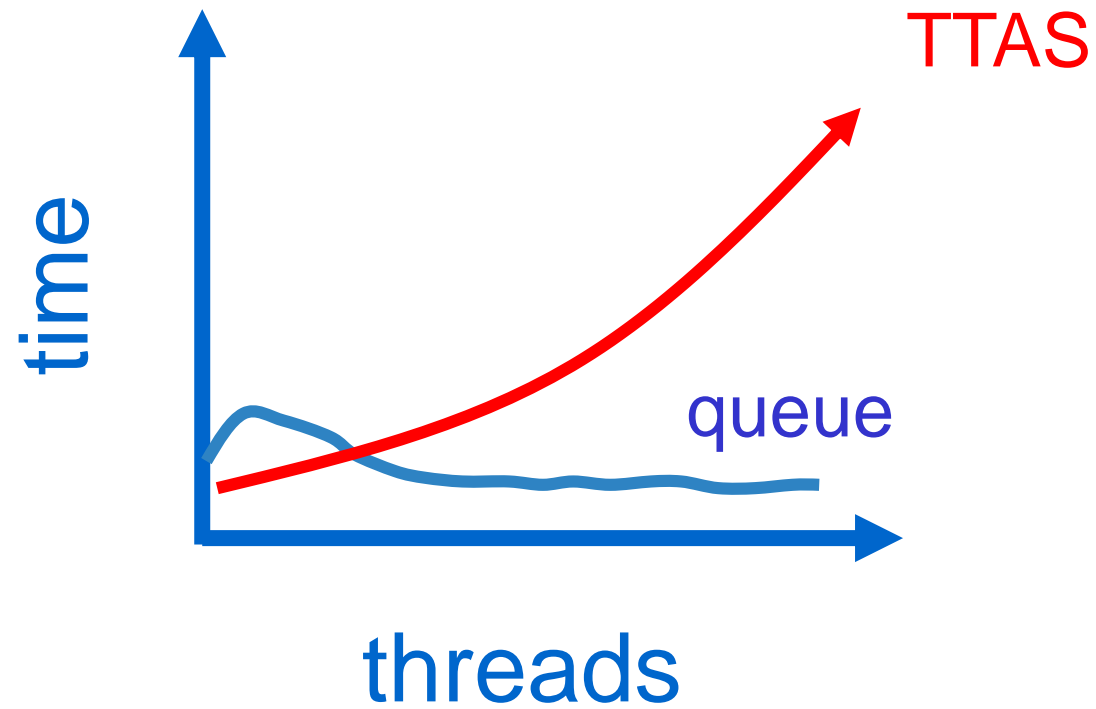
Queue Lock



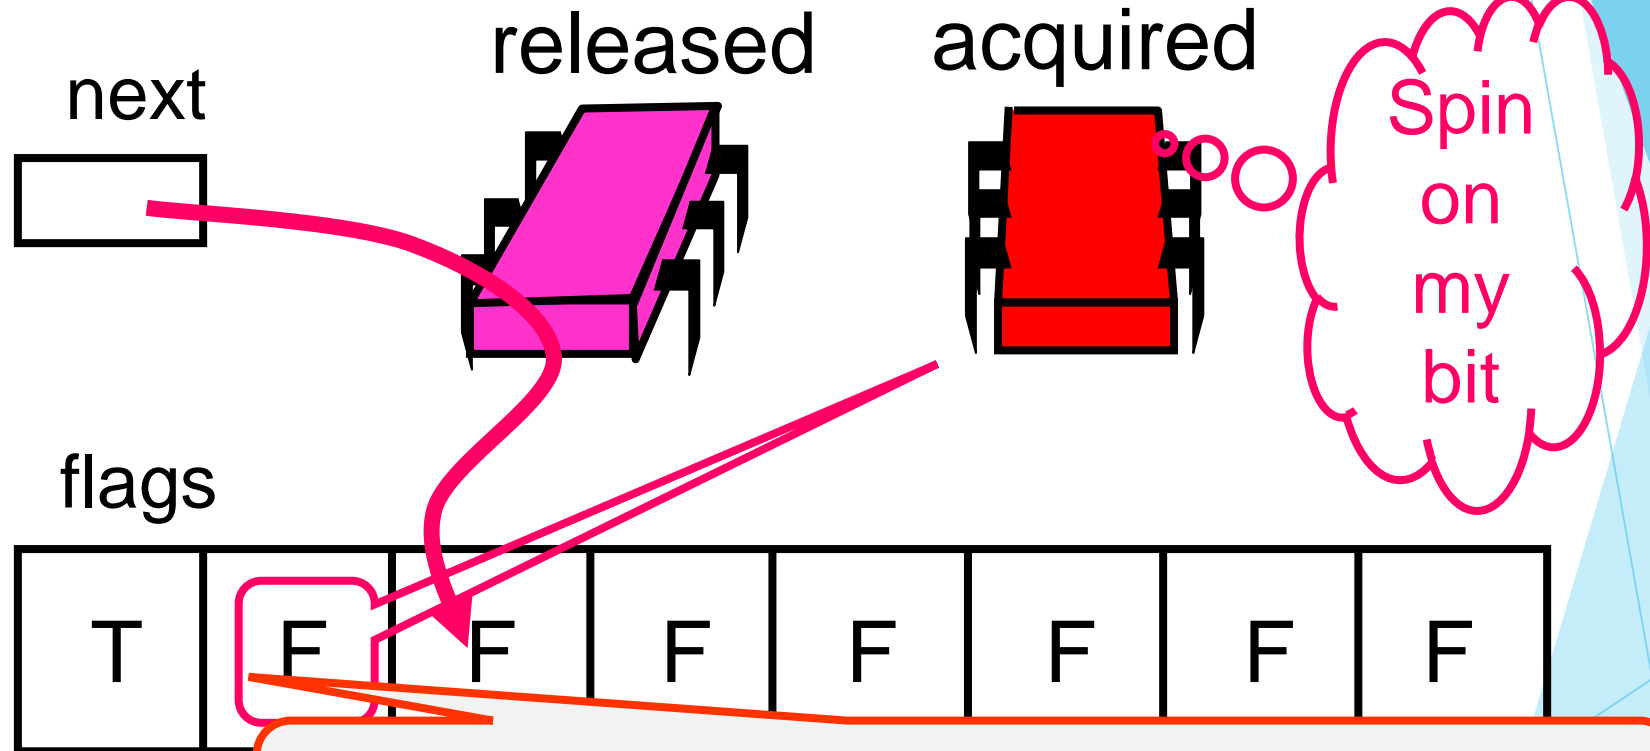
Queue Locks

```
1 public void lock() {  
2     int slot = tail.getAndIncrement() % size;  
3     mySlotIndex.set(slot);  
4     while (! flag[slot]) {};  
5 }  
6 public void unlock() {  
7     int slot = mySlotIndex.get();  
8     flag[slot] = false;  
9     flag[(slot + 1) % size] = true;  
10 }
```

Performance graph

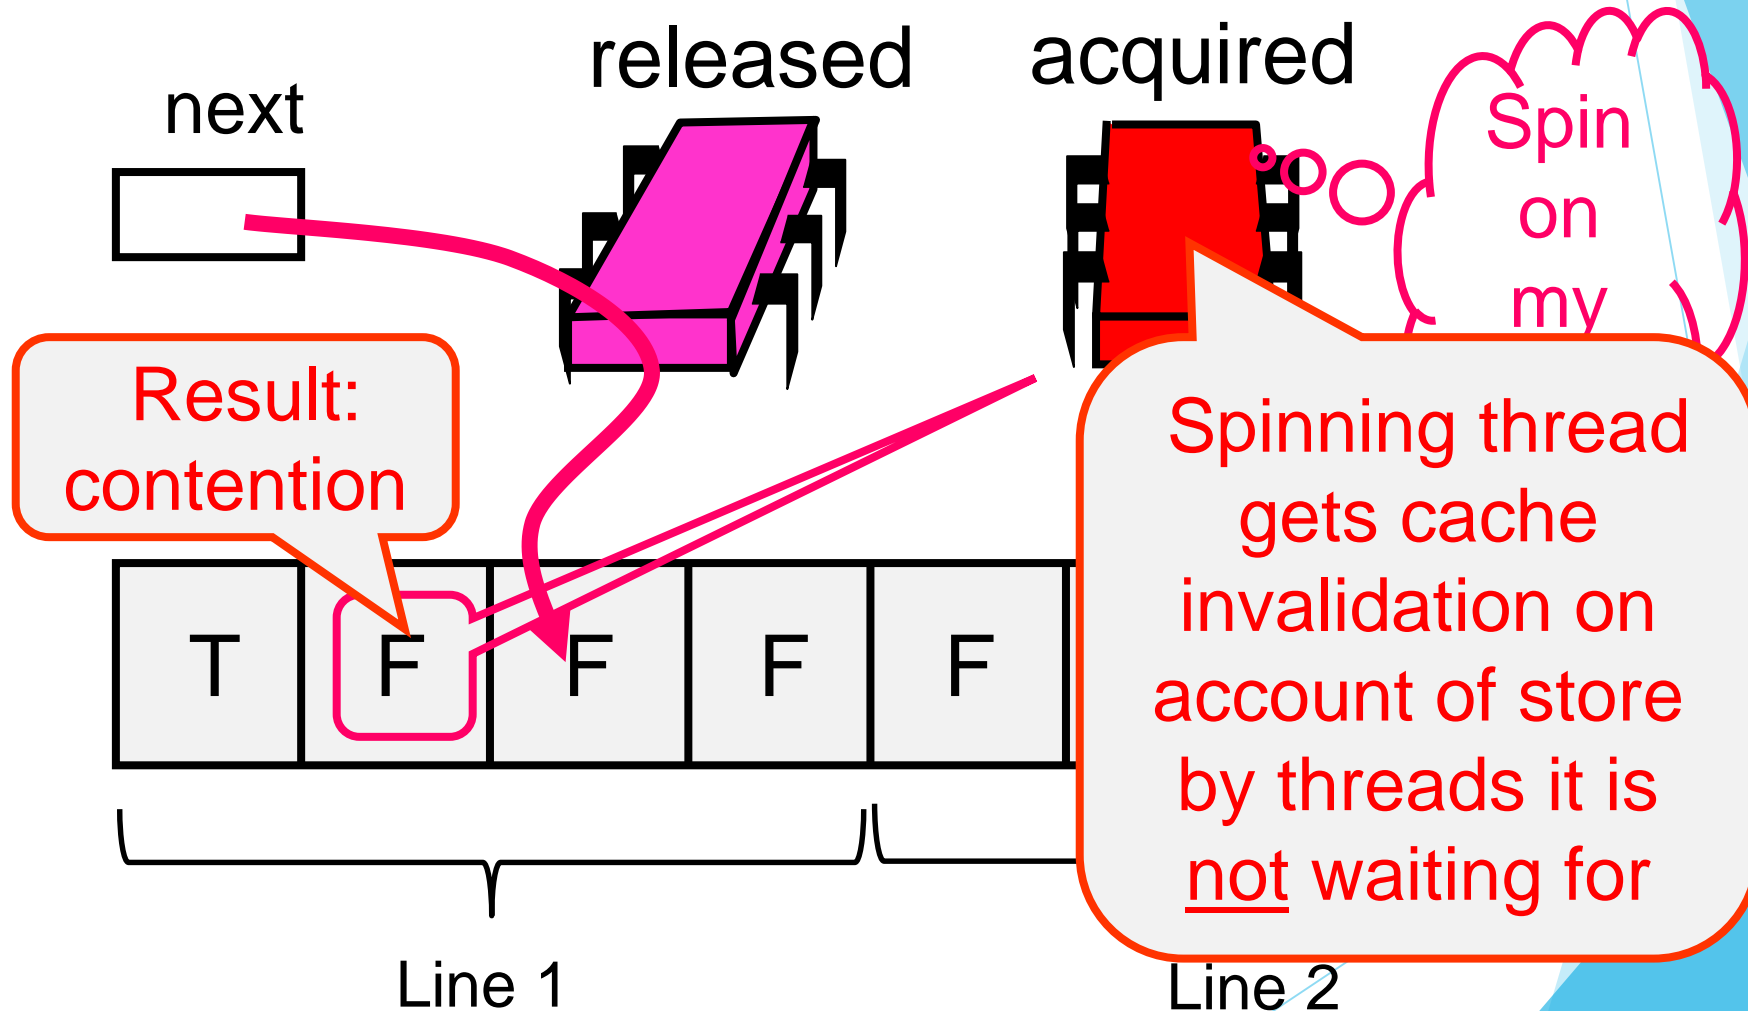


Local Spinning

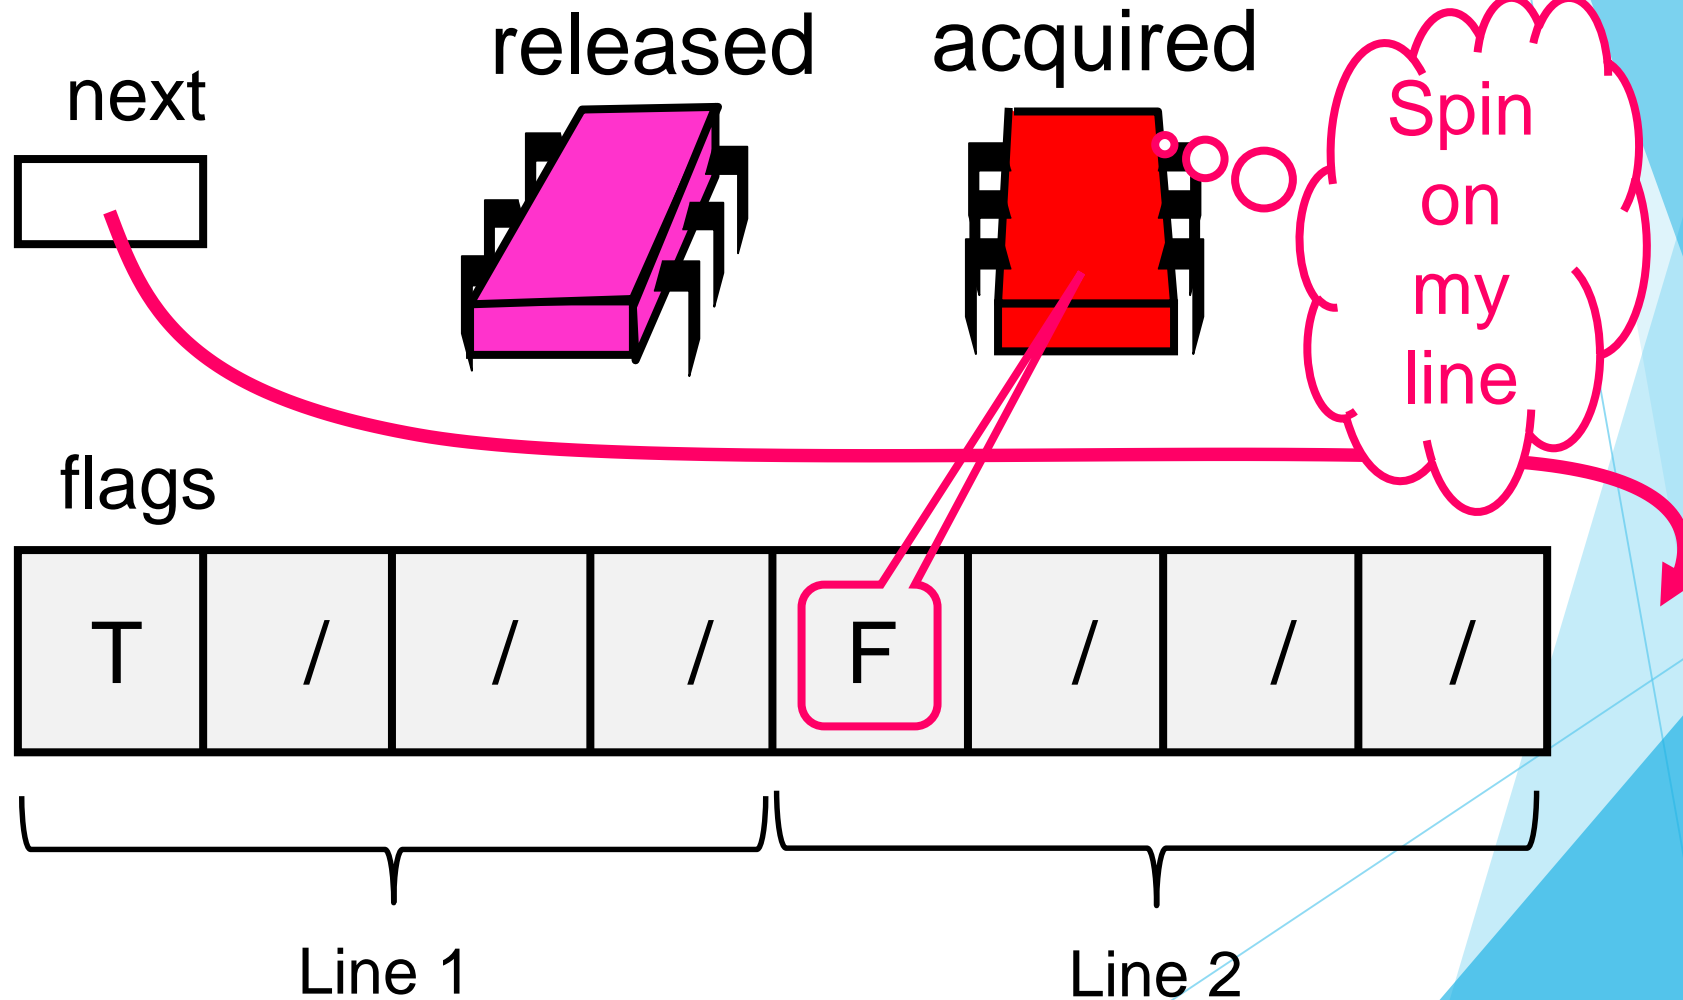


Unfortunately many bits share cache line

False Sharing



Solution: Padding



Queue Locks


Pros:

- ▶ first-come-first-served fairness.
- ▶ Easy to implement.
- ▶ Scalable performance.

Cons:

- ▶ Not space-efficient $O(n)$ **per lock** for n threads.
- ▶ Specifically one cache line per thread.


Initially




Lock is free

Initially

idle




tail




Pink Wants the Lock

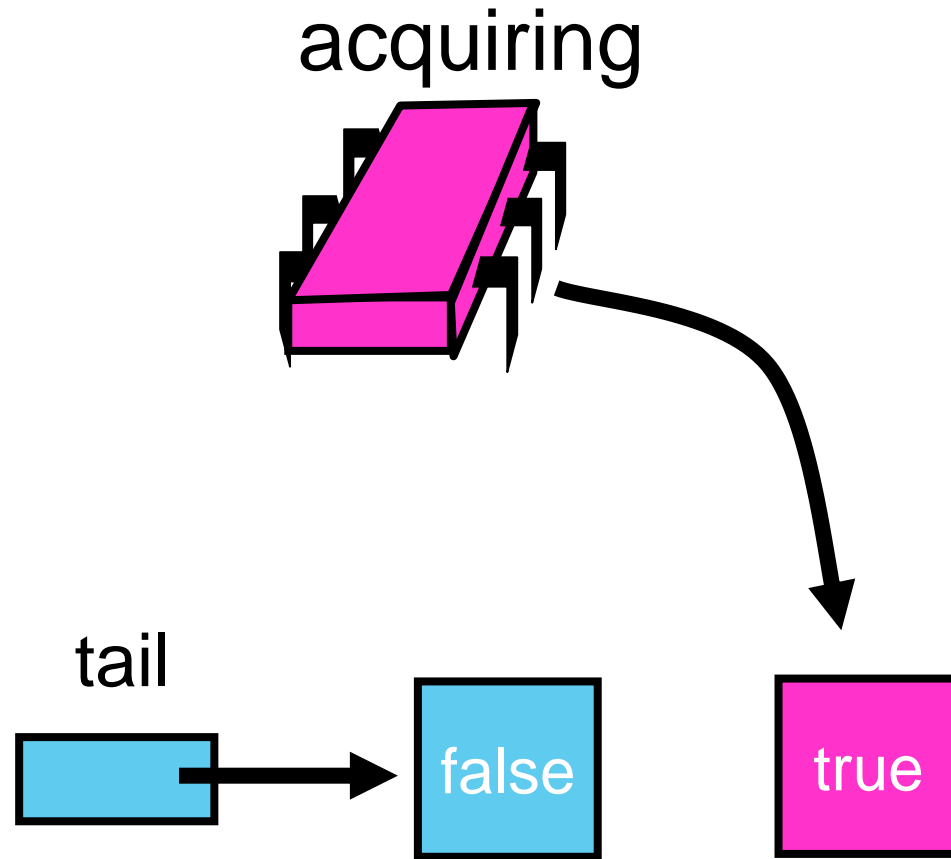
acquiring




tail




Pink Wants the Lock




Pink Wants the Lock




Pink Has the Lock




Red Wants the Lock



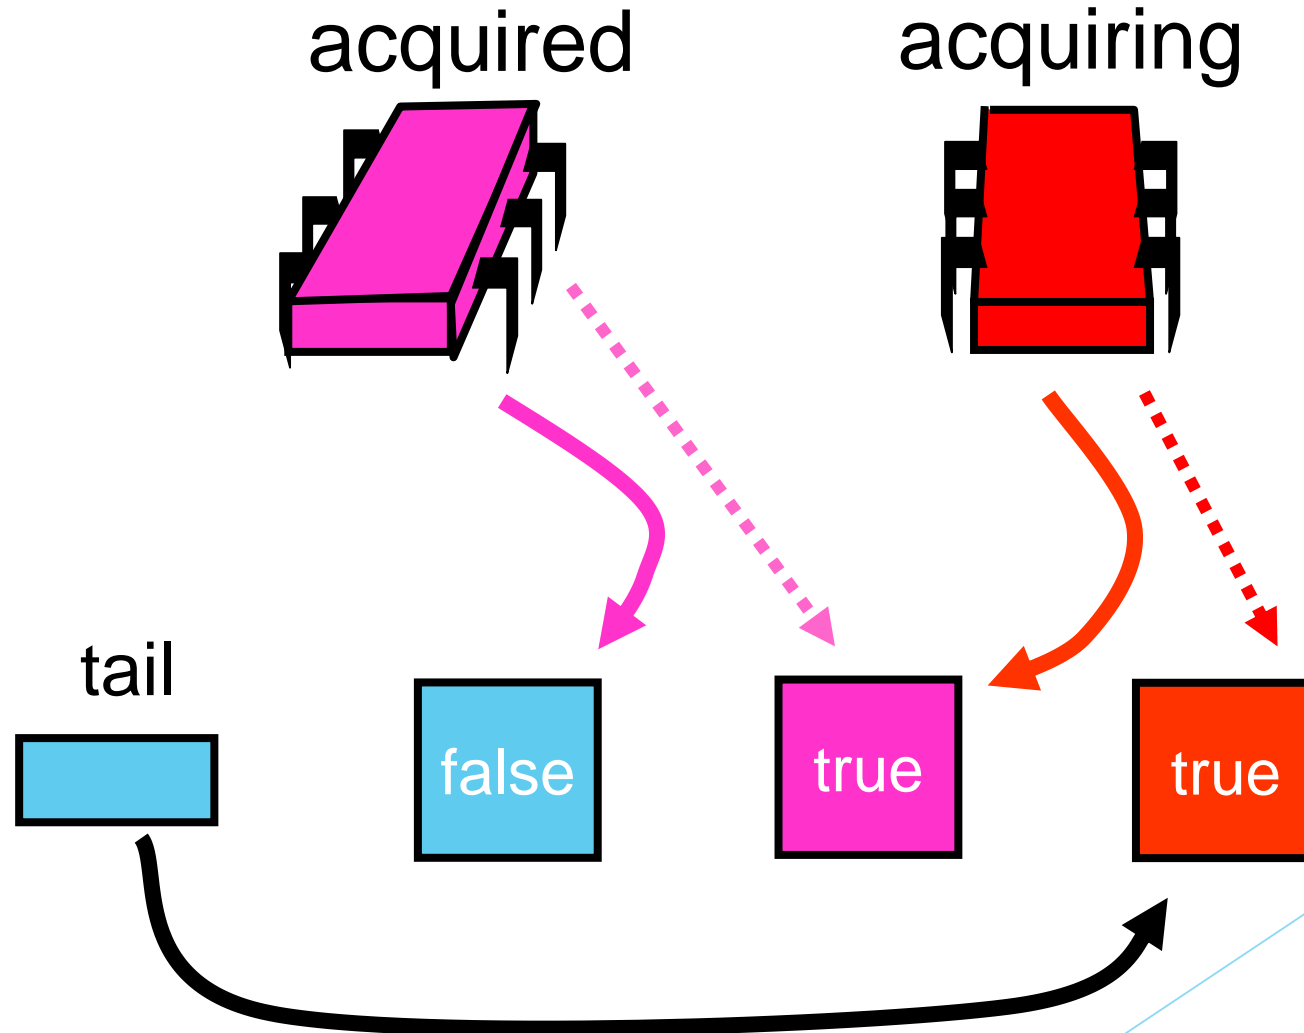
Red Wants the Lock




Red Wants the Lock




Red Wants the Lock



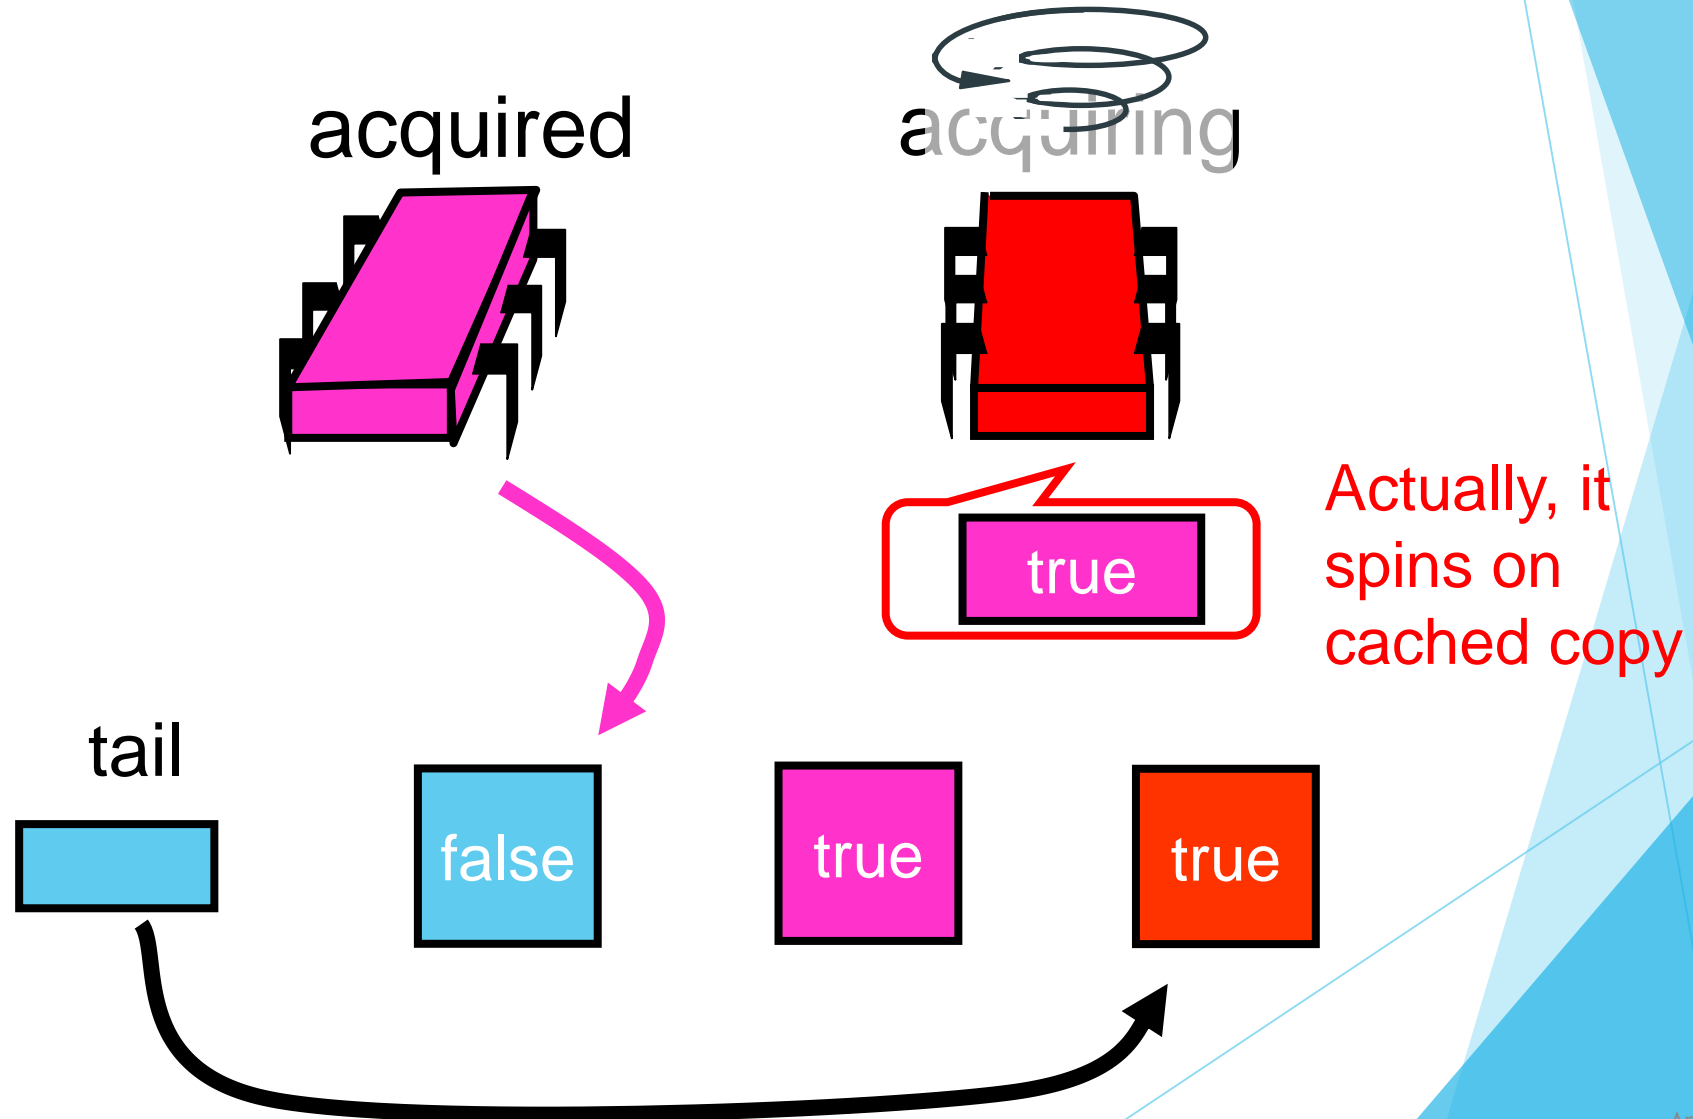
Red Wants the Lock



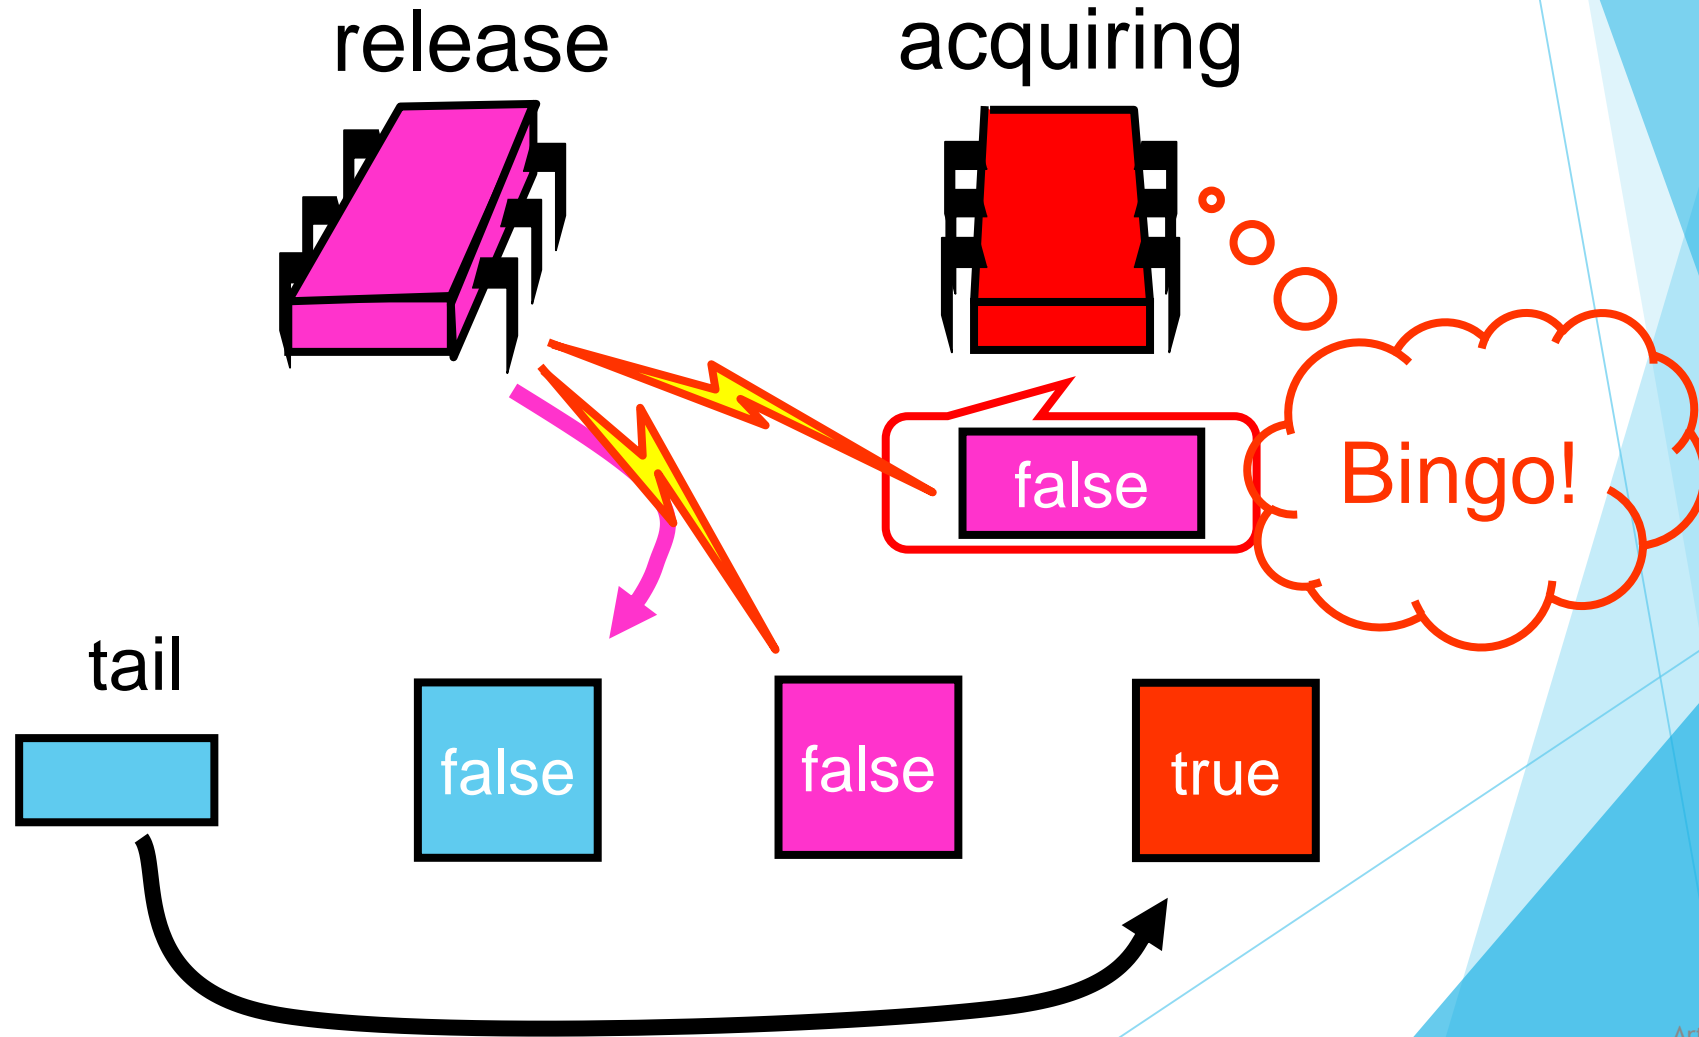
Red Wants the Lock



Red Wants the Lock




Pink Releases




Pink Releases

released




acquired



tail



true



CLH Lock

```
1 public void lock() {  
2   QNode qnode = myNode.get();  
3   qnode.locked = true;  
4   QNode pred = tail.getAndSet(qnode);  
6   while (pred.locked) {}  
7 }  
8 public void unlock() {  
9   myNode.locked.set(false);  
10  myNode=pred
```

CLH Lock

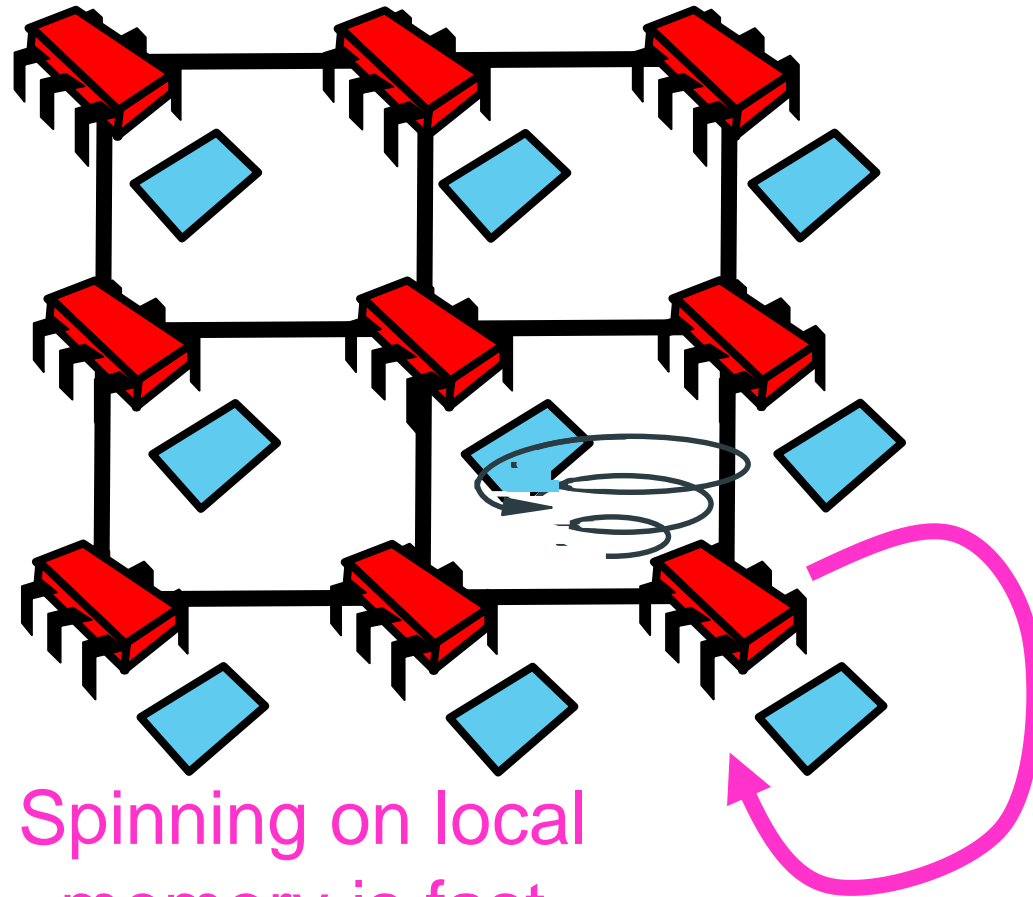
Pros:

- ▶ Constant size space.
- ▶ First-come-first-served fairness.
- ▶ Each thread spin on a distinct location.
- ▶ Does not require a knowledge of number of threads to come.

Cons:

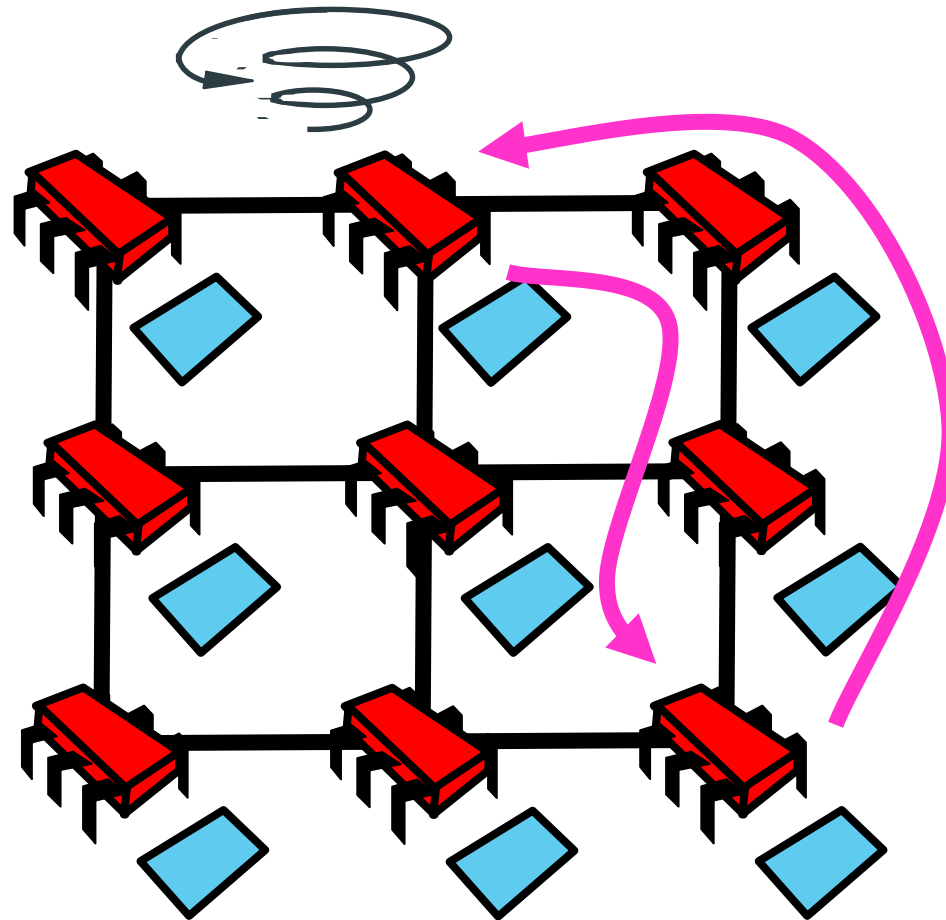
- ▶ Doesn't work for un-cached NUMA architectures.

NUMA Machines



Spinning on local
memory is fast

NUMA Machines



Spinning on remote
memory is slow

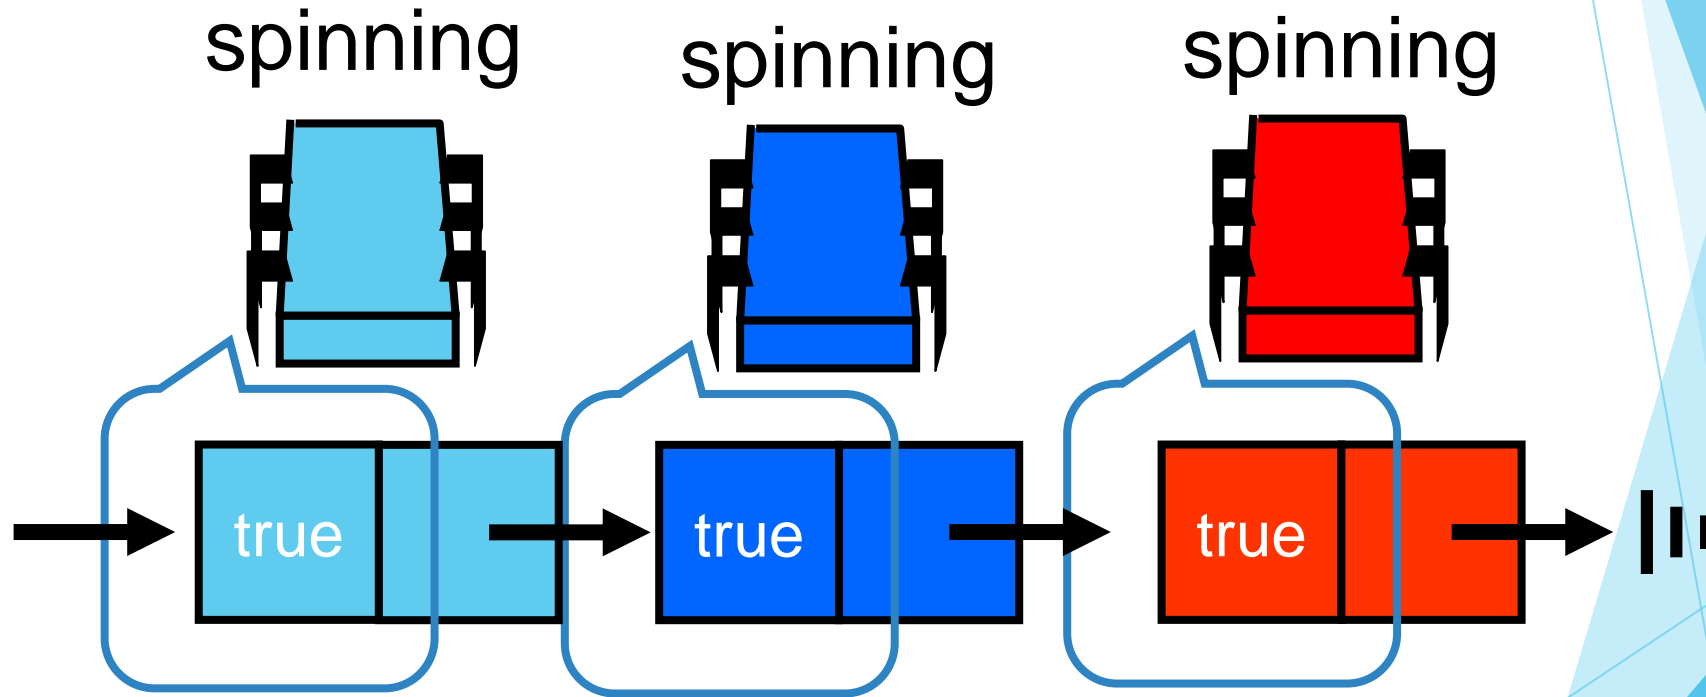
MCS Queue

- ▶ The thread spins on a (local) locked field in its own QNode waiting until its predecessor sets this field to **false**.
- ▶ As a result each thread controls the location on which it spins.


Lock with timeout

- ▶ Up till now we studied first come first serve with some contention algorithm.
- ▶ In real time system sometimes thread need the ability to give up waiting for lock.
- ▶ In Backoff lock the thread can simply return from the lock() function.
- ▶ In queue lock algorithm if a thread just returns the threads queued up behind it will starve.

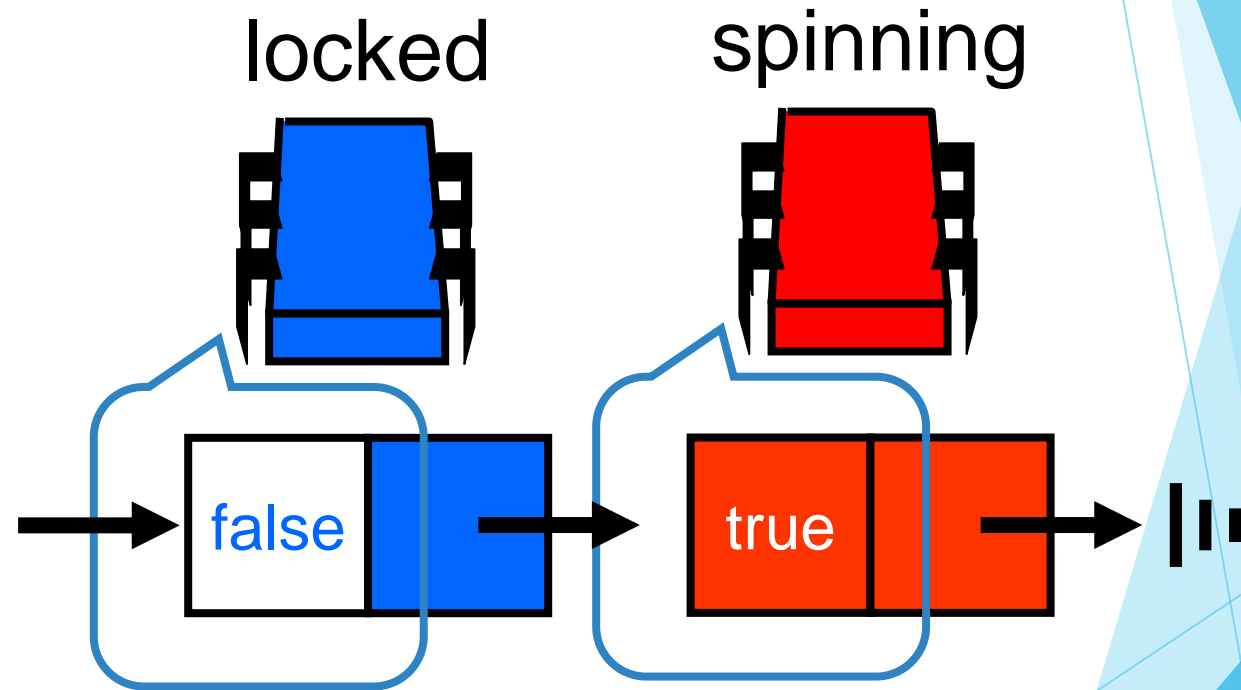
Queue Locks



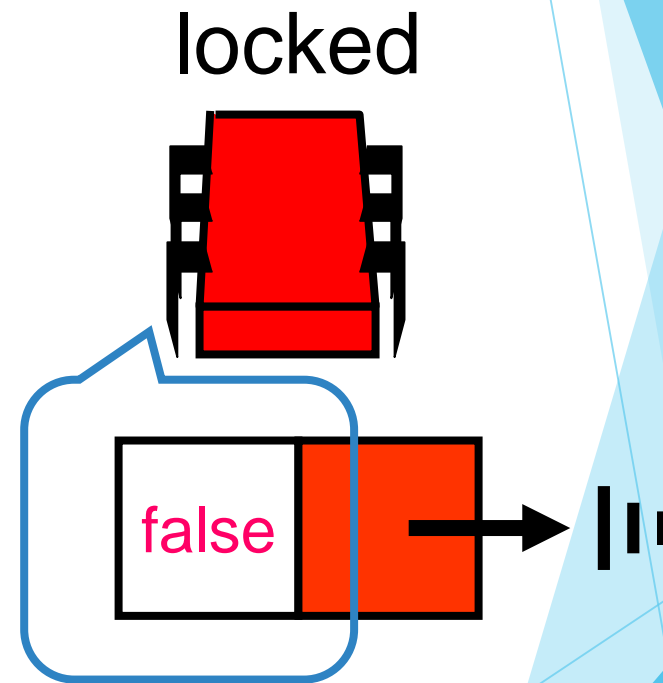
Queue Locks



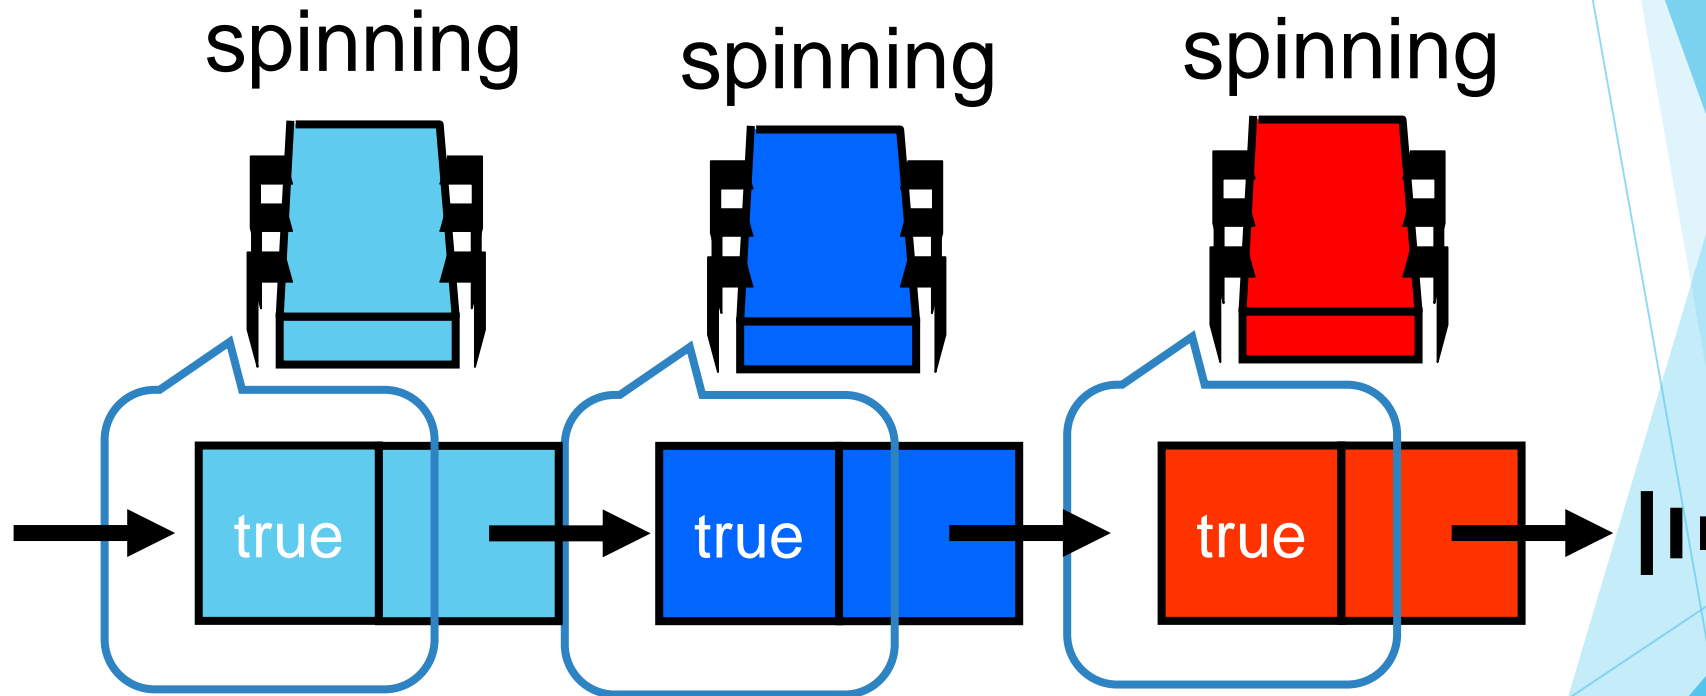
Queue Locks




Queue Locks



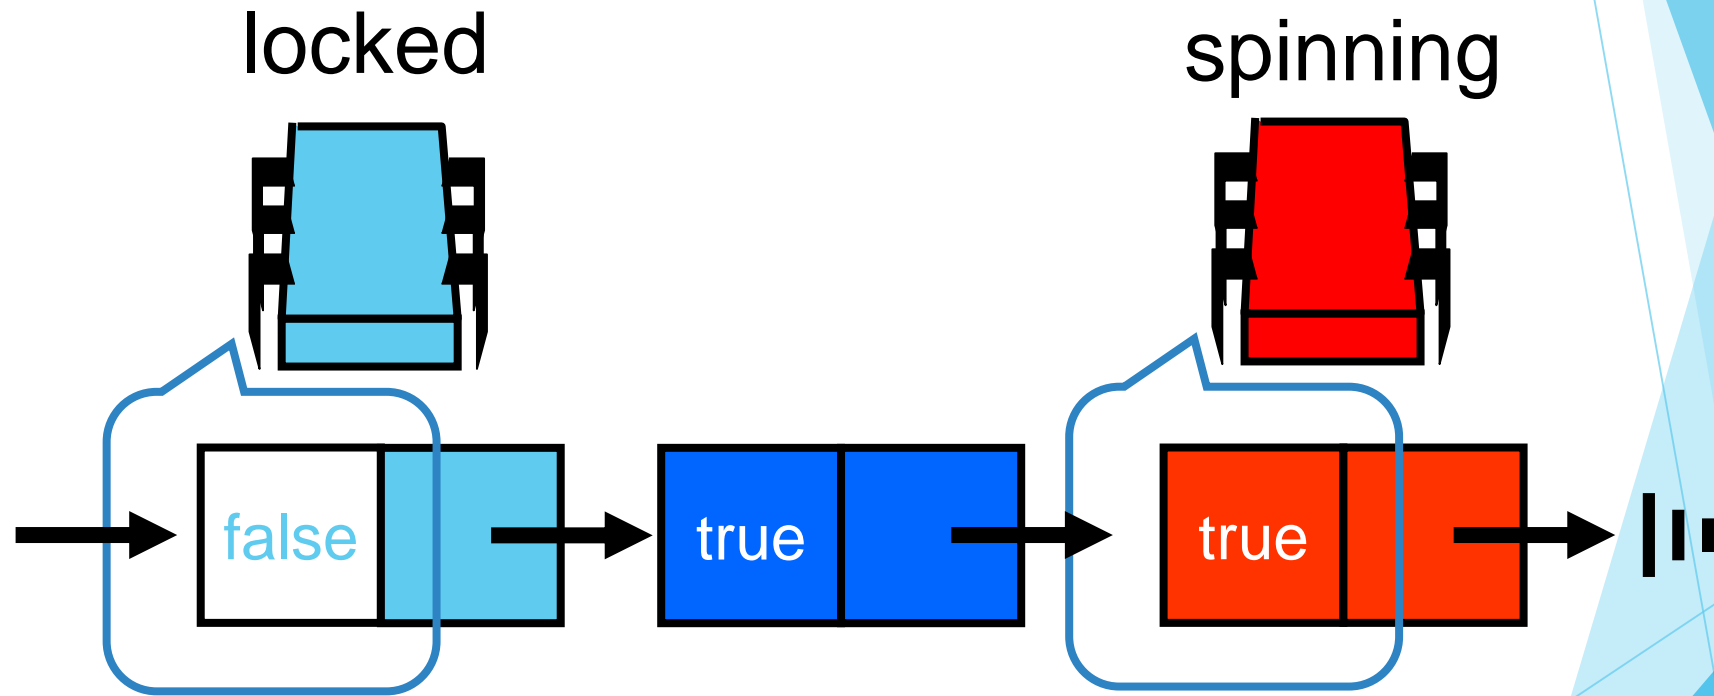
Queue Locks




Queue Locks



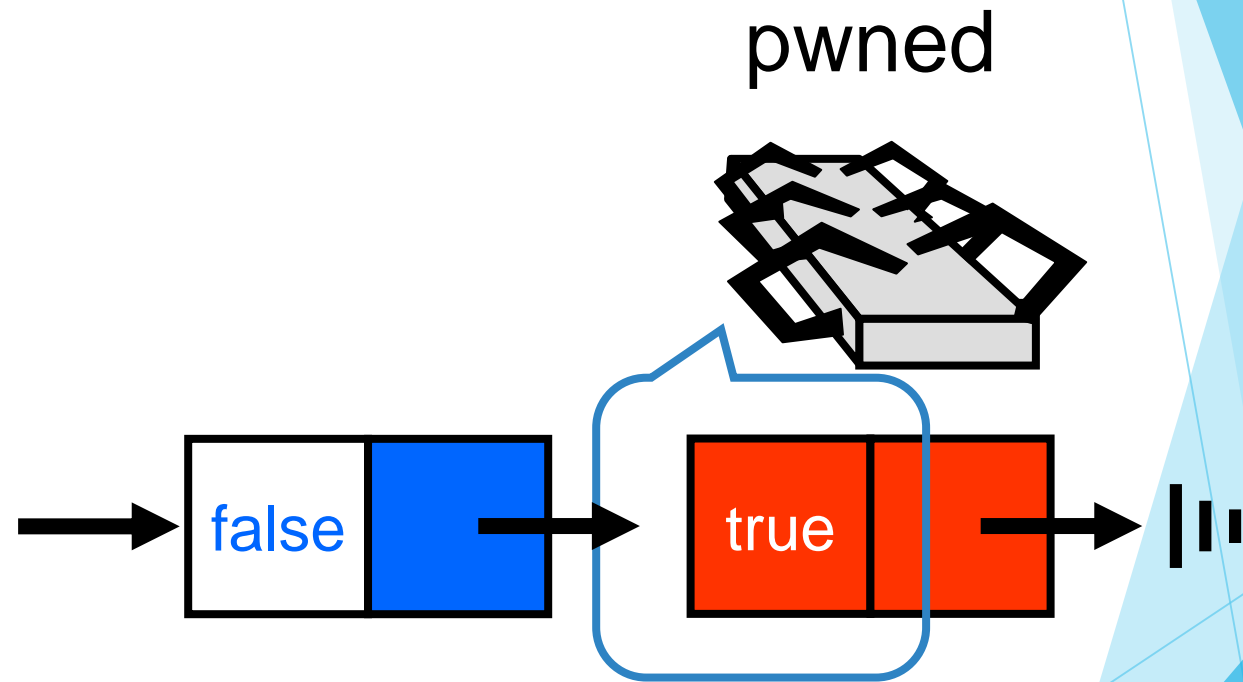
Queue Locks



Queue Locks



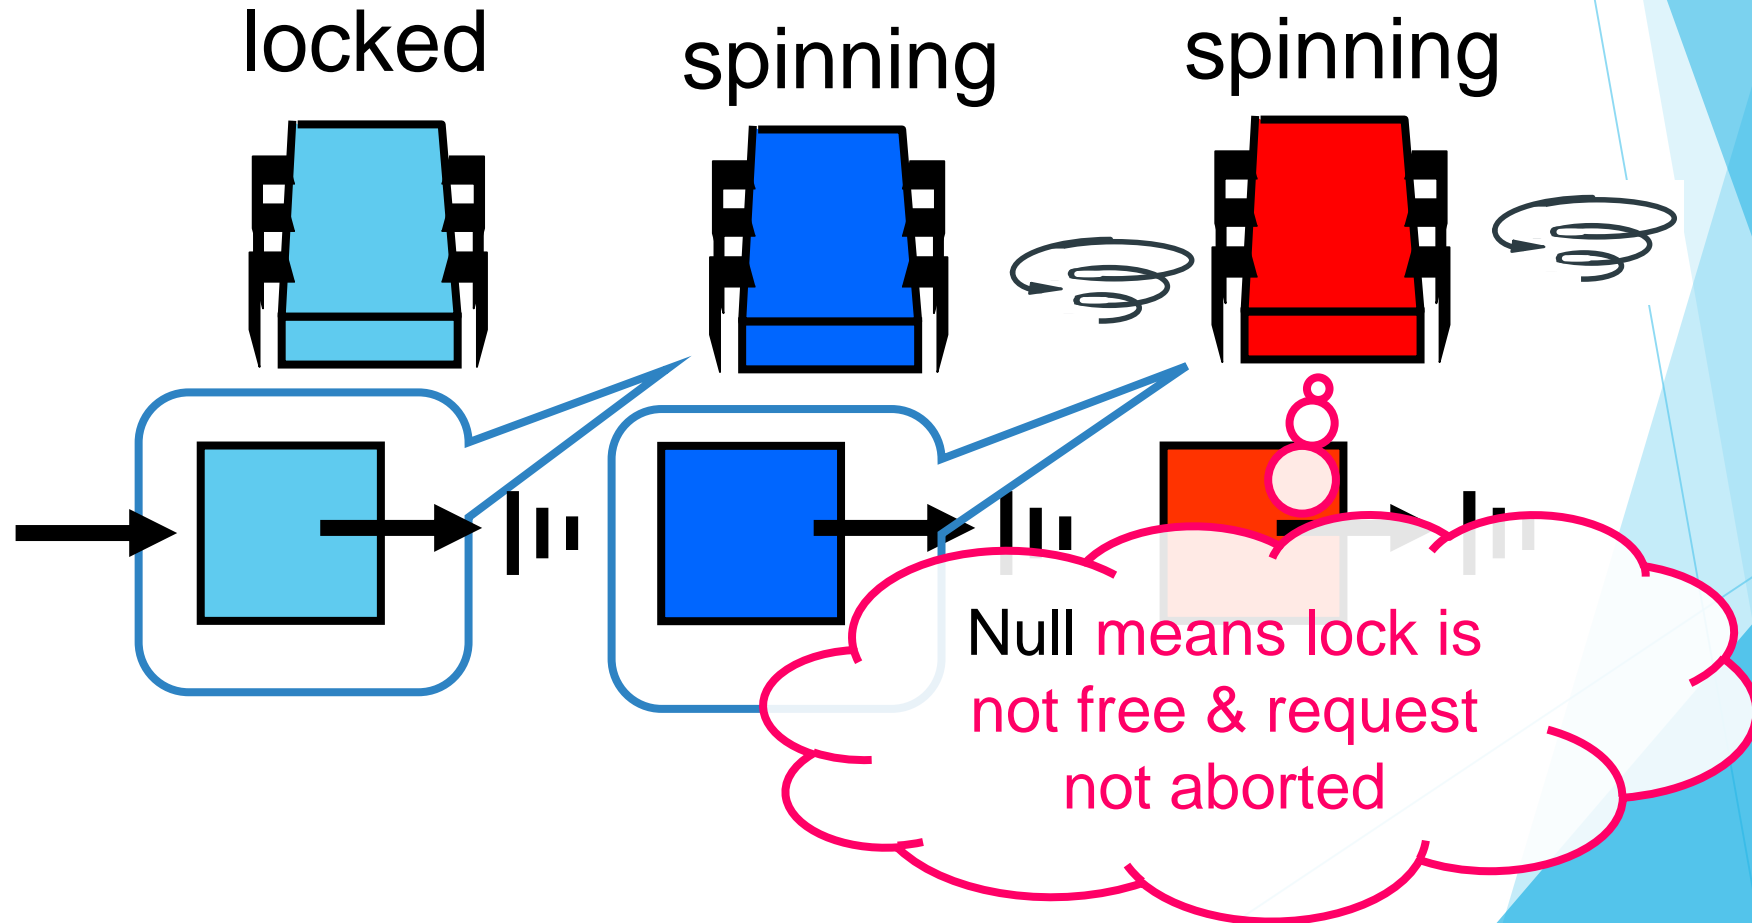
Queue Locks



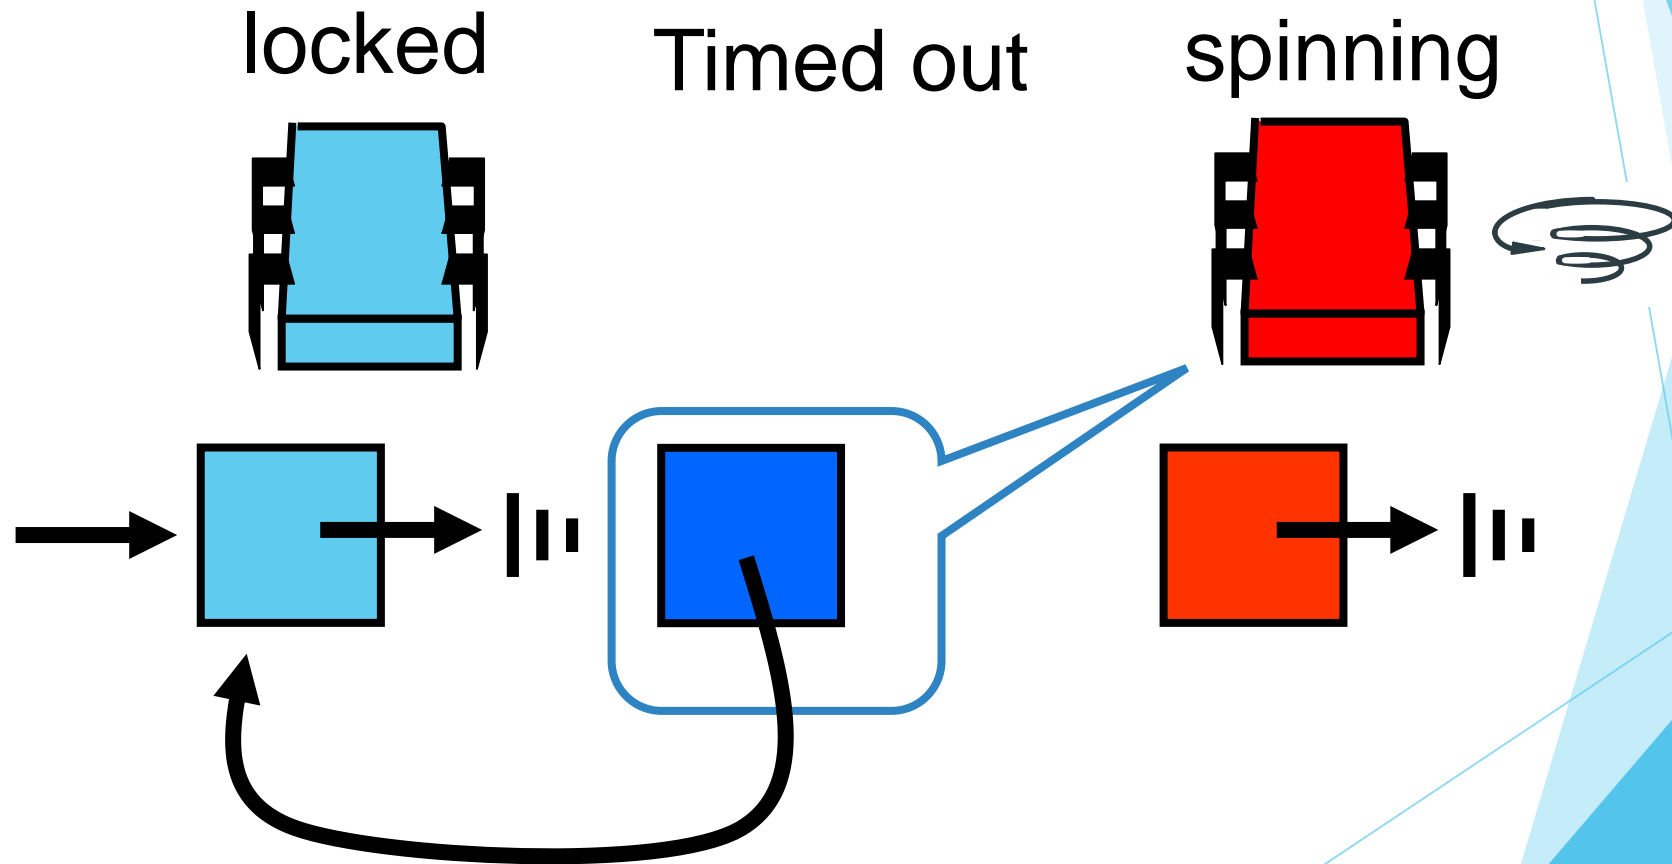
Timeout in CLH lock

- ▶ The idea here is to let the successor know about the abandoned so it will deal with it.
- ▶ The thread successor in the queue, if there is one, notices that the node on which it is spinning has been abandoned, and starts spinning on the abandoned node's predecessor.

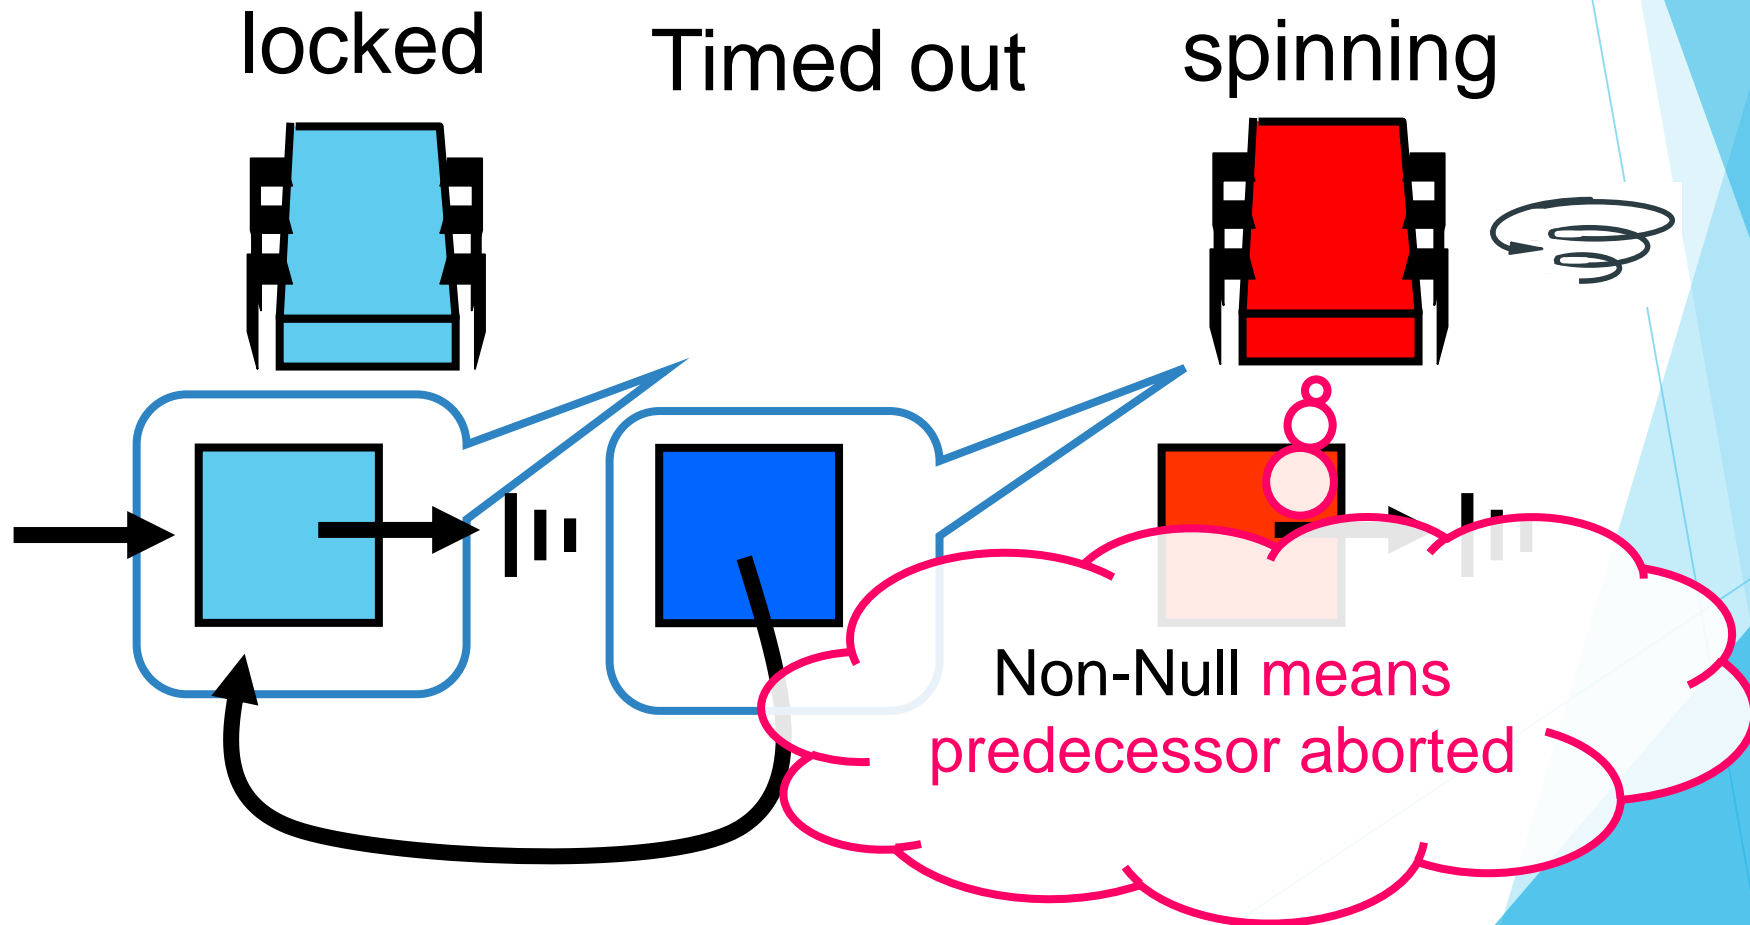
Normal Case



One Thread Aborts




Successor Notices



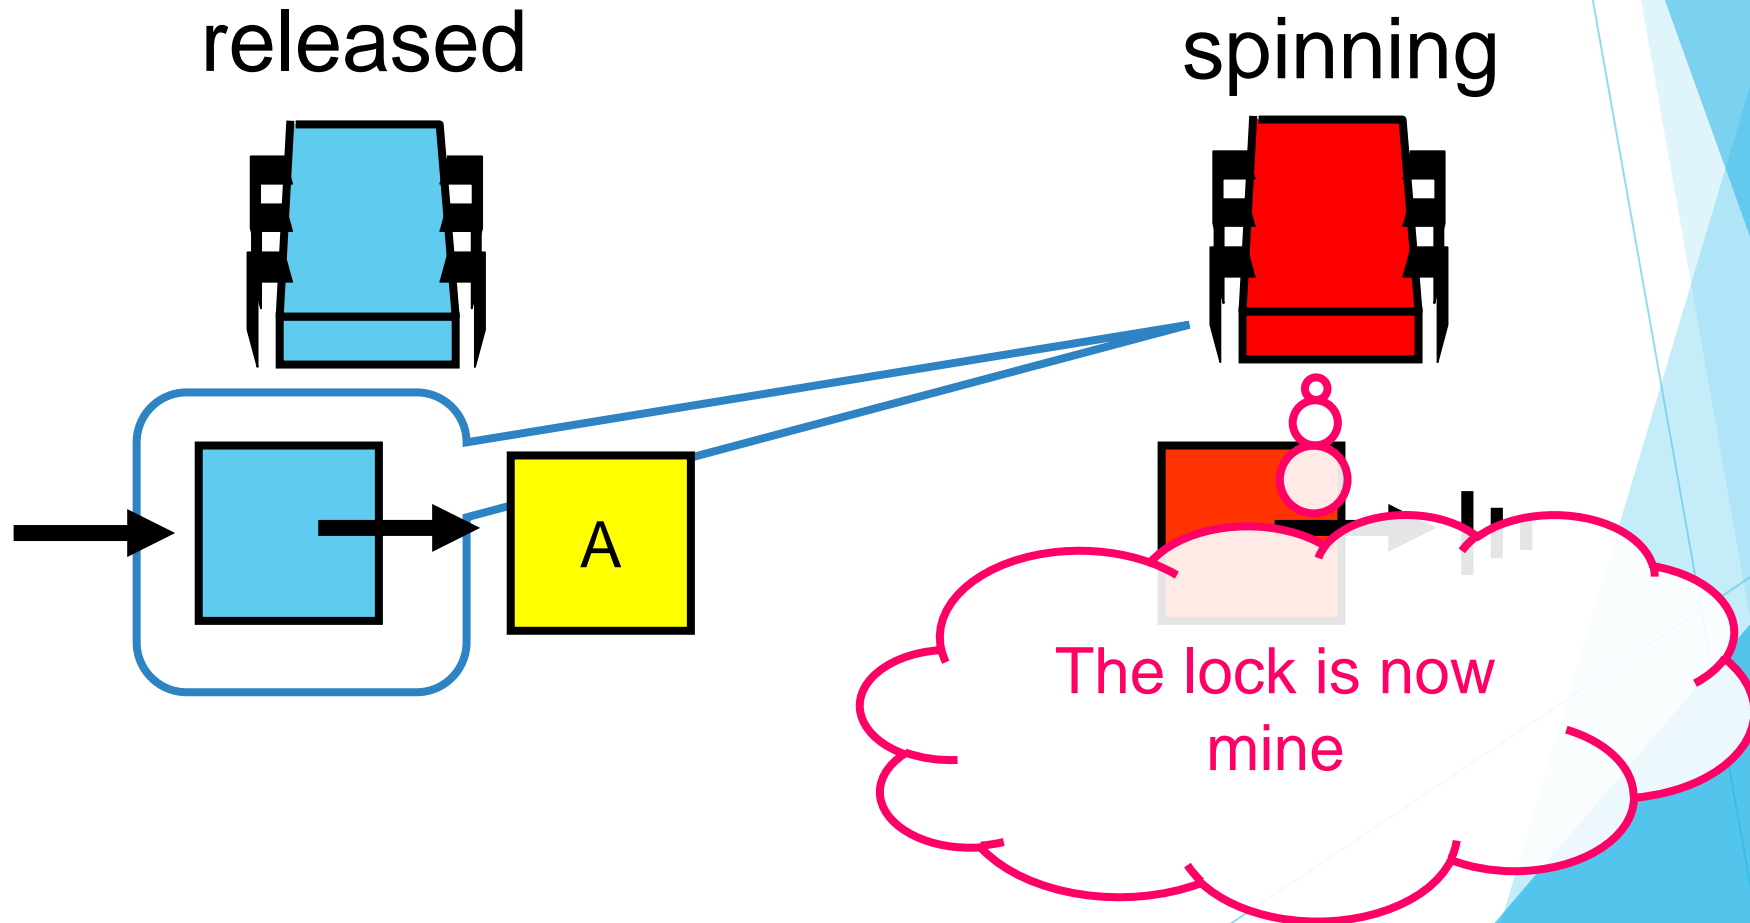
Recycle Predecessor's Node



Spin on Earlier Node



Spin on Earlier Node



Timeout in CLH lock

```
1 public boolean tryLock(long time, TimeUnit unit)
2 QNode qnode = new QNode();
3 myNode.set(qnode);
4 qnode.pred = null;
5 QNode myPred = tail.getAndSet(qnode);
6 if (myPred == null || myPred.pred == AVAILABLE) {
7     return true;
```


Timeout in CLH lock

```
8 long start=now()
9 while (now()- startTime < patience) {
10     QNode predPred = myPred.pred;
11     if (predPred == AVAILABLE) {
12         return true;
13     } else if (predPred != null) {
14         myPred = predPred; } }
```

Timeout in CLH lock

```
15 if (!tail.compareAndSet(qnode, myPred))  
16     qnode.pred = myPred;  
17     Return }
```

If my time is up!

Timeout in CLH un-lock

```
18 public void unlock() {  
19     QNode qnode = myNode.get();  
20     if (!tail.compareAndSet(qnode, null))  
21         qnode.pred = AVAILABLE;  
22 }
```

Composite Lock

- ▶ consider an advanced lock algorithm that combines the best of Queue locks and backoff lock.
- ▶ The CompositeLock class keeps a short, fixed-size array of lock nodes.
- ▶ Each thread that tries to acquire the lock selects a node in the array at random.
- ▶ If: that node is in use then backoff.
- ▶ Else: The thread spins on the preceding node, and when that node's owner signals it is done, the thread enters the critical section.

Composite Lock

Pros:

- ▶ When threads back off, they access different locations, reducing contention.
- ▶ Abandoning a lock request is easy.
- ▶ For L locks and n threads, the CompositeLock class, requires only $O(L)$ space.

Cons:

- ▶ Complex code.
- ▶ No guarantee first-come-first-served.

Conclusion

- ▶ We have seen a variety of spin locks that vary in characteristics and performance.
- ▶ Each spin lock algorithm will work best in different situations.
- ▶ This depends on the application, hardware and which properties are important.

Thank you for listening!