

# Spin Locks and Contention

CS 475, Fall 2019

Concurrent & Distributed Systems

# Lock-Free vs. Wait-free


- Wait-Free: each method call takes a finite number of steps to finish
- Lock-free: infinitely often some method call finishes



# Lock-Free Increment

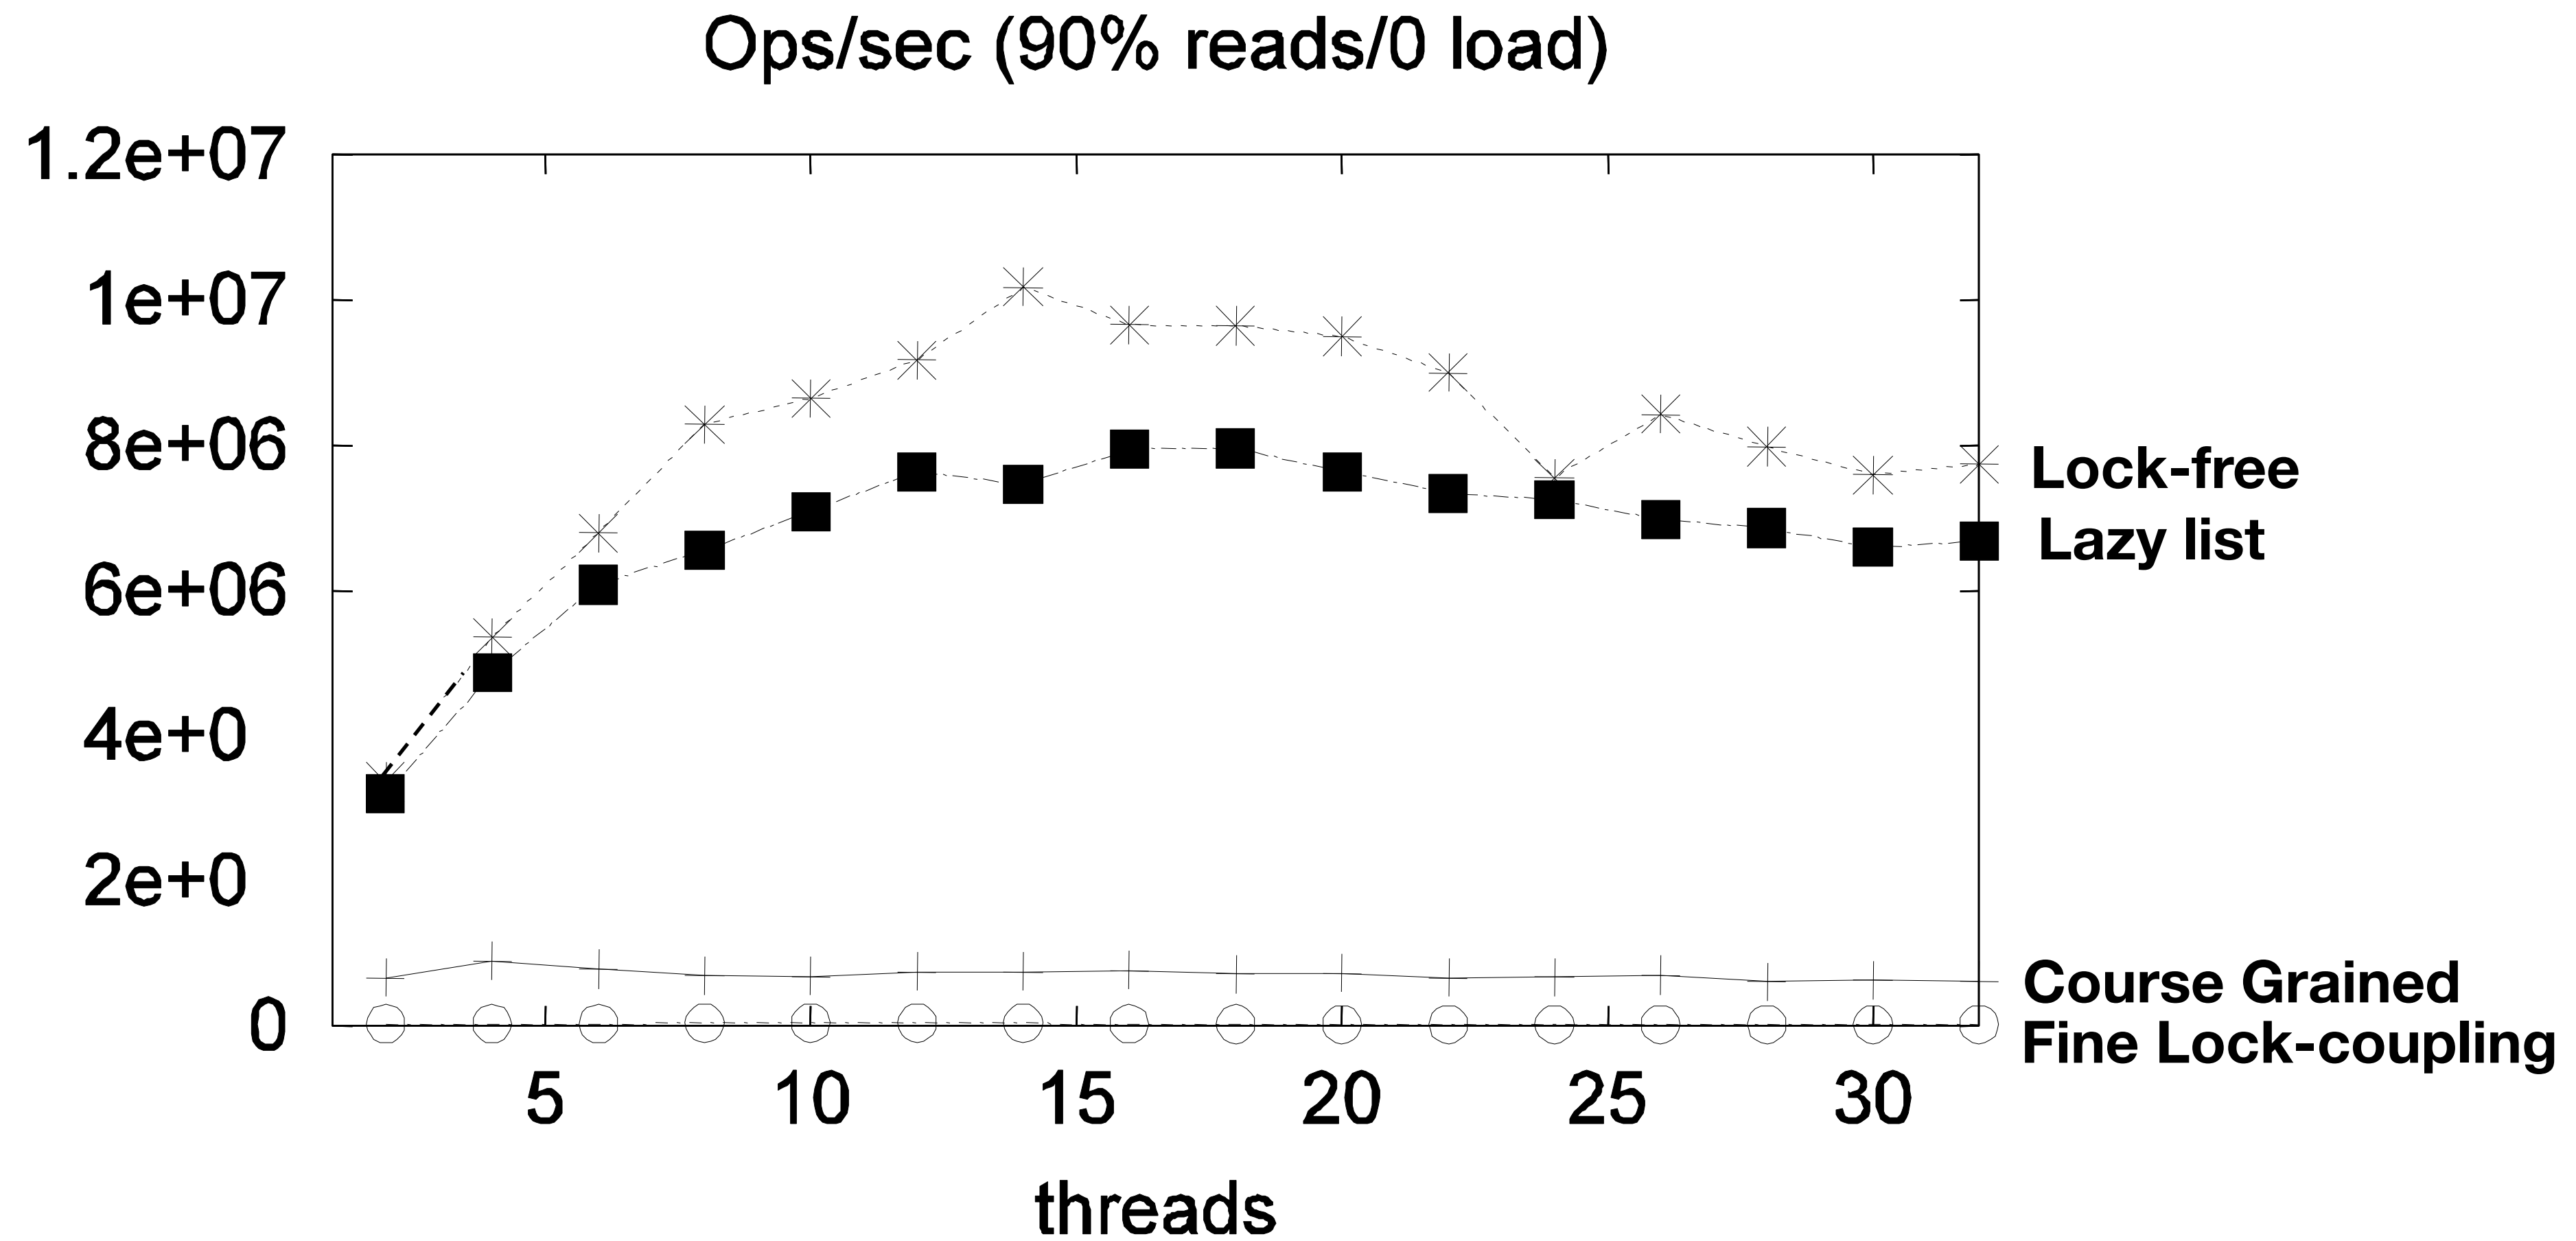
```
public abstract class Increment {  
    private int value;  
  
    public void increment() {  
        boolean success = false;  
        while(!success){  
            int prior = this.value;  
            success = compareAndSet(prior, prior+1);  
        }  
    }  
}
```

**CAS Failure**



time

# High Contains Ratio



# Focus so far: Correctness and Progress

- Models
  - Accurate (we never lied to you)
  - But idealized (so we forgot to mention a few things)
- Protocols
  - Elegant
  - Important
  - But naïve

# New Focus: Performance

- Models
  - More complicated (not the same as complex!)
  - Still focus on principles (not soon obsolete)
- Protocols
  - Elegant (in their fashion)
  - Important (why else would we pay attention)
  - And realistic (your mileage may vary)

# Today: Revisit Mutual Exclusion

- Think of performance, not just correctness and progress
- Begin to understand how performance depends on our software properly utilizing the multiprocessor machine's hardware
- And get to know a collection of locking algorithms...

# What Should you do if you can't get a lock?

- Keep trying
  - “spin” or “busy-wait”
  - Good if delays are short  
(consider cost of switching threads)
- Give up the processor
  - Good if delays are long
  - Always good on uniprocessor



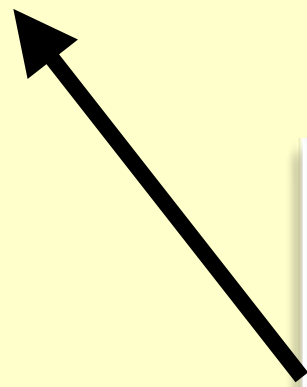
# What Should you do if you can't get a lock?

- Keep trying
  - “spin” or “busy-wait”
  - Good if delays are short  
(consider cost of switching threads)
- Give up the processor
  - Good if delays are long
  - Always good on uniprocessor

**Today**


# Peterson's Algorithm - Reminder

```
public void lock() {  
    flag[i] = true;  
    victim = i;  
    while (flag[j] && victim == i) {};  
}  
public void unlock() {  
    flag[i] = false;  
}
```




**Threads “spin” in this empty loop while waiting**  
**“Busy” waiting - thread keeps running on the CPU**

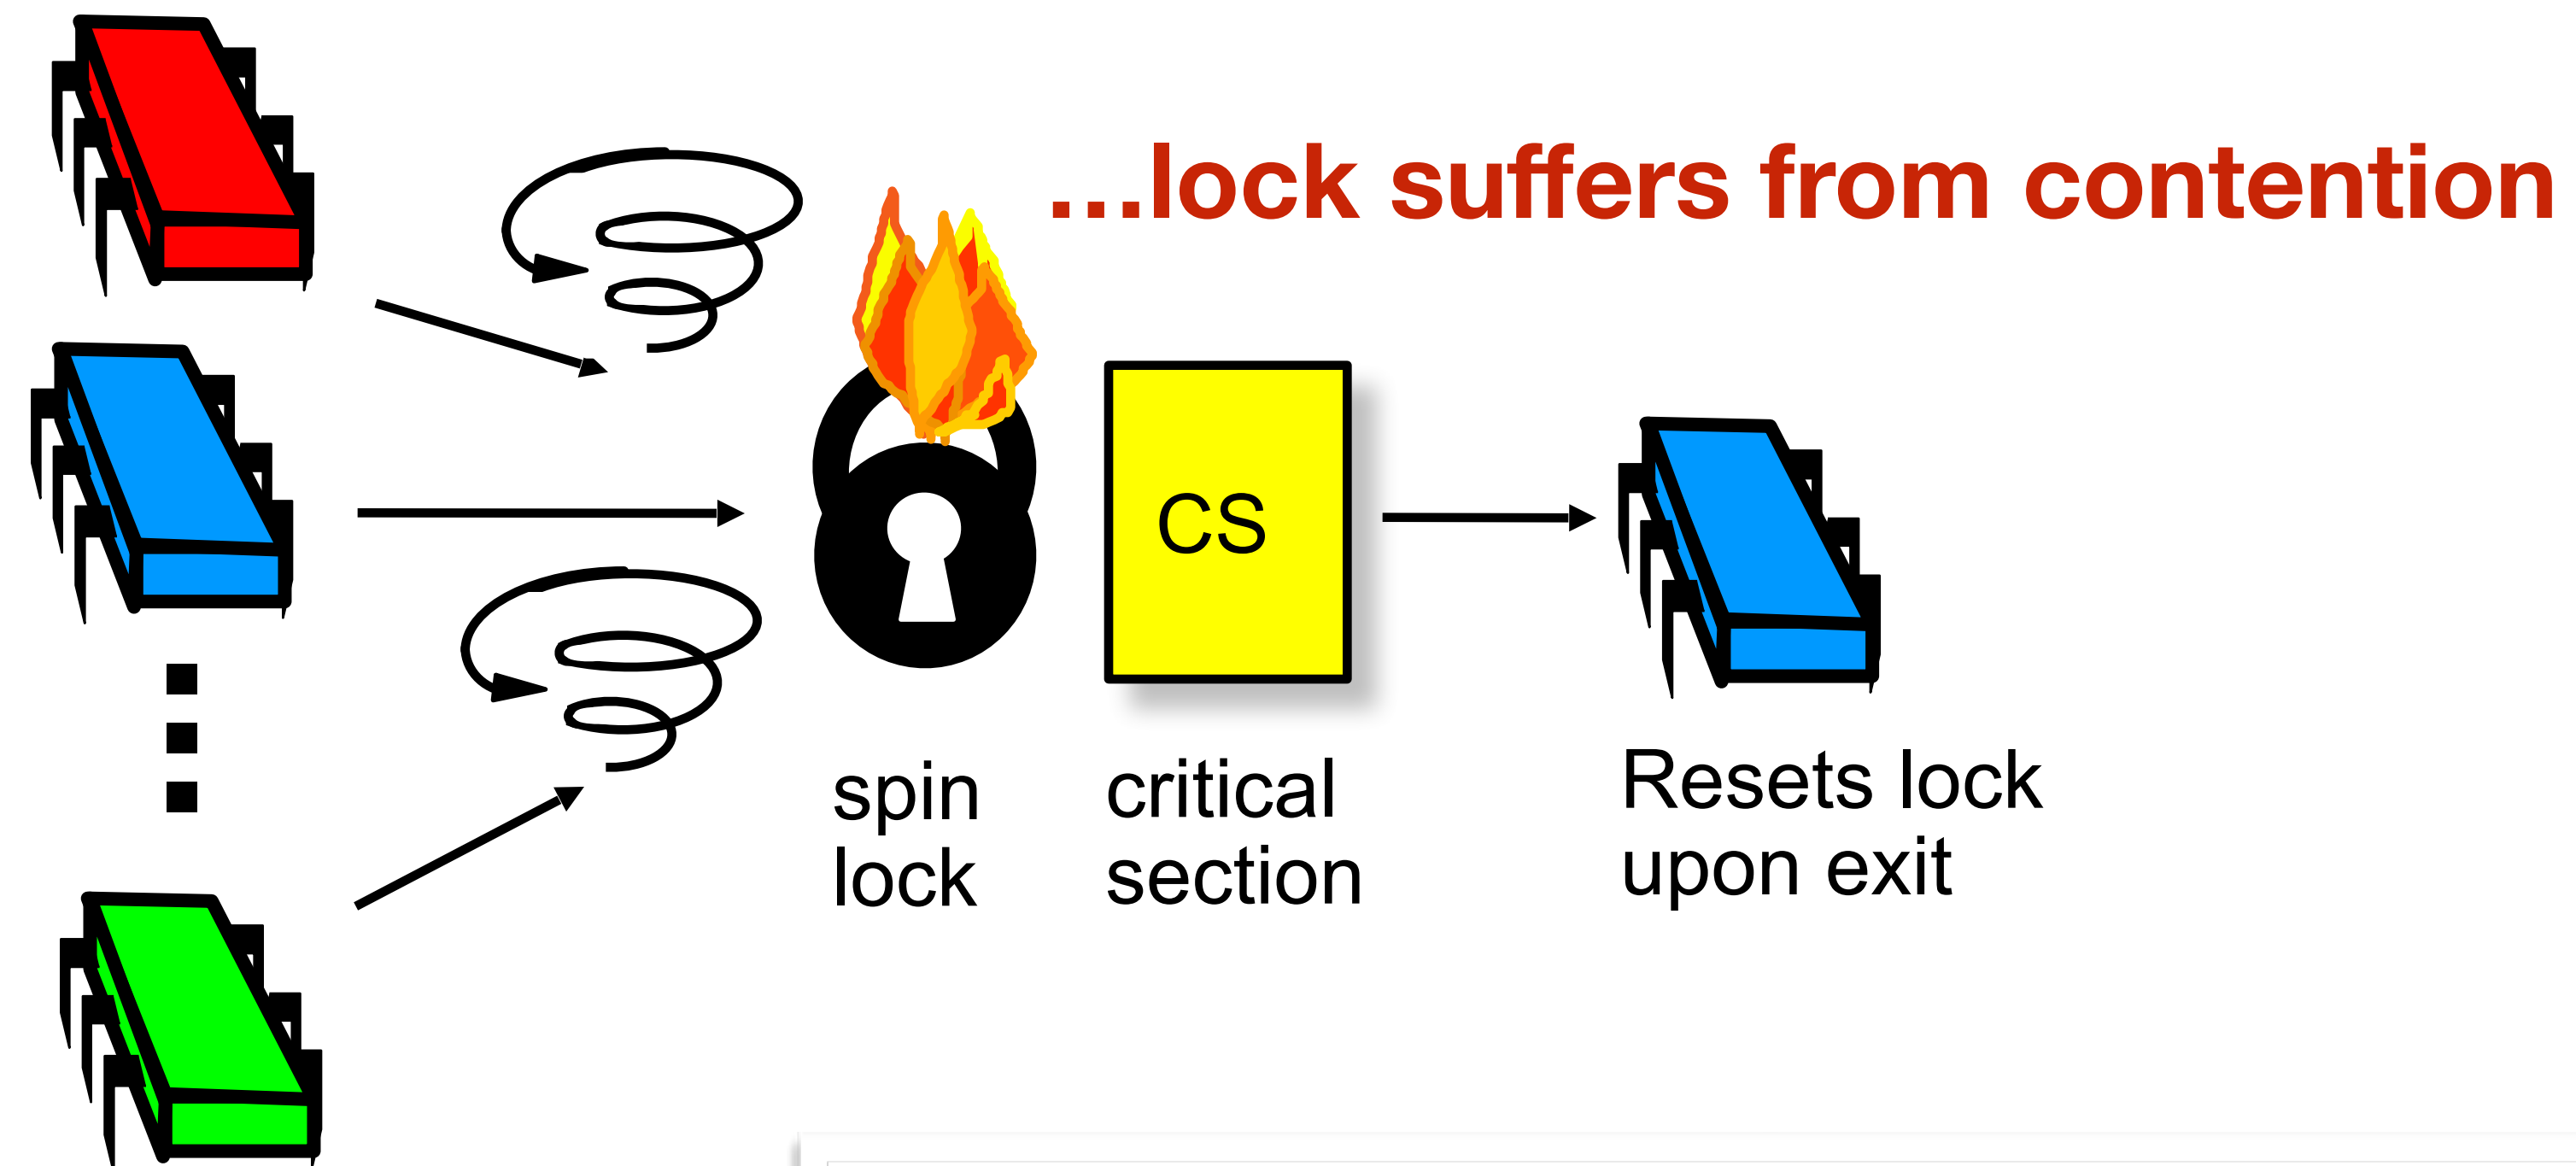
# Basic Spin-Lock



# Basic Spin-Lock

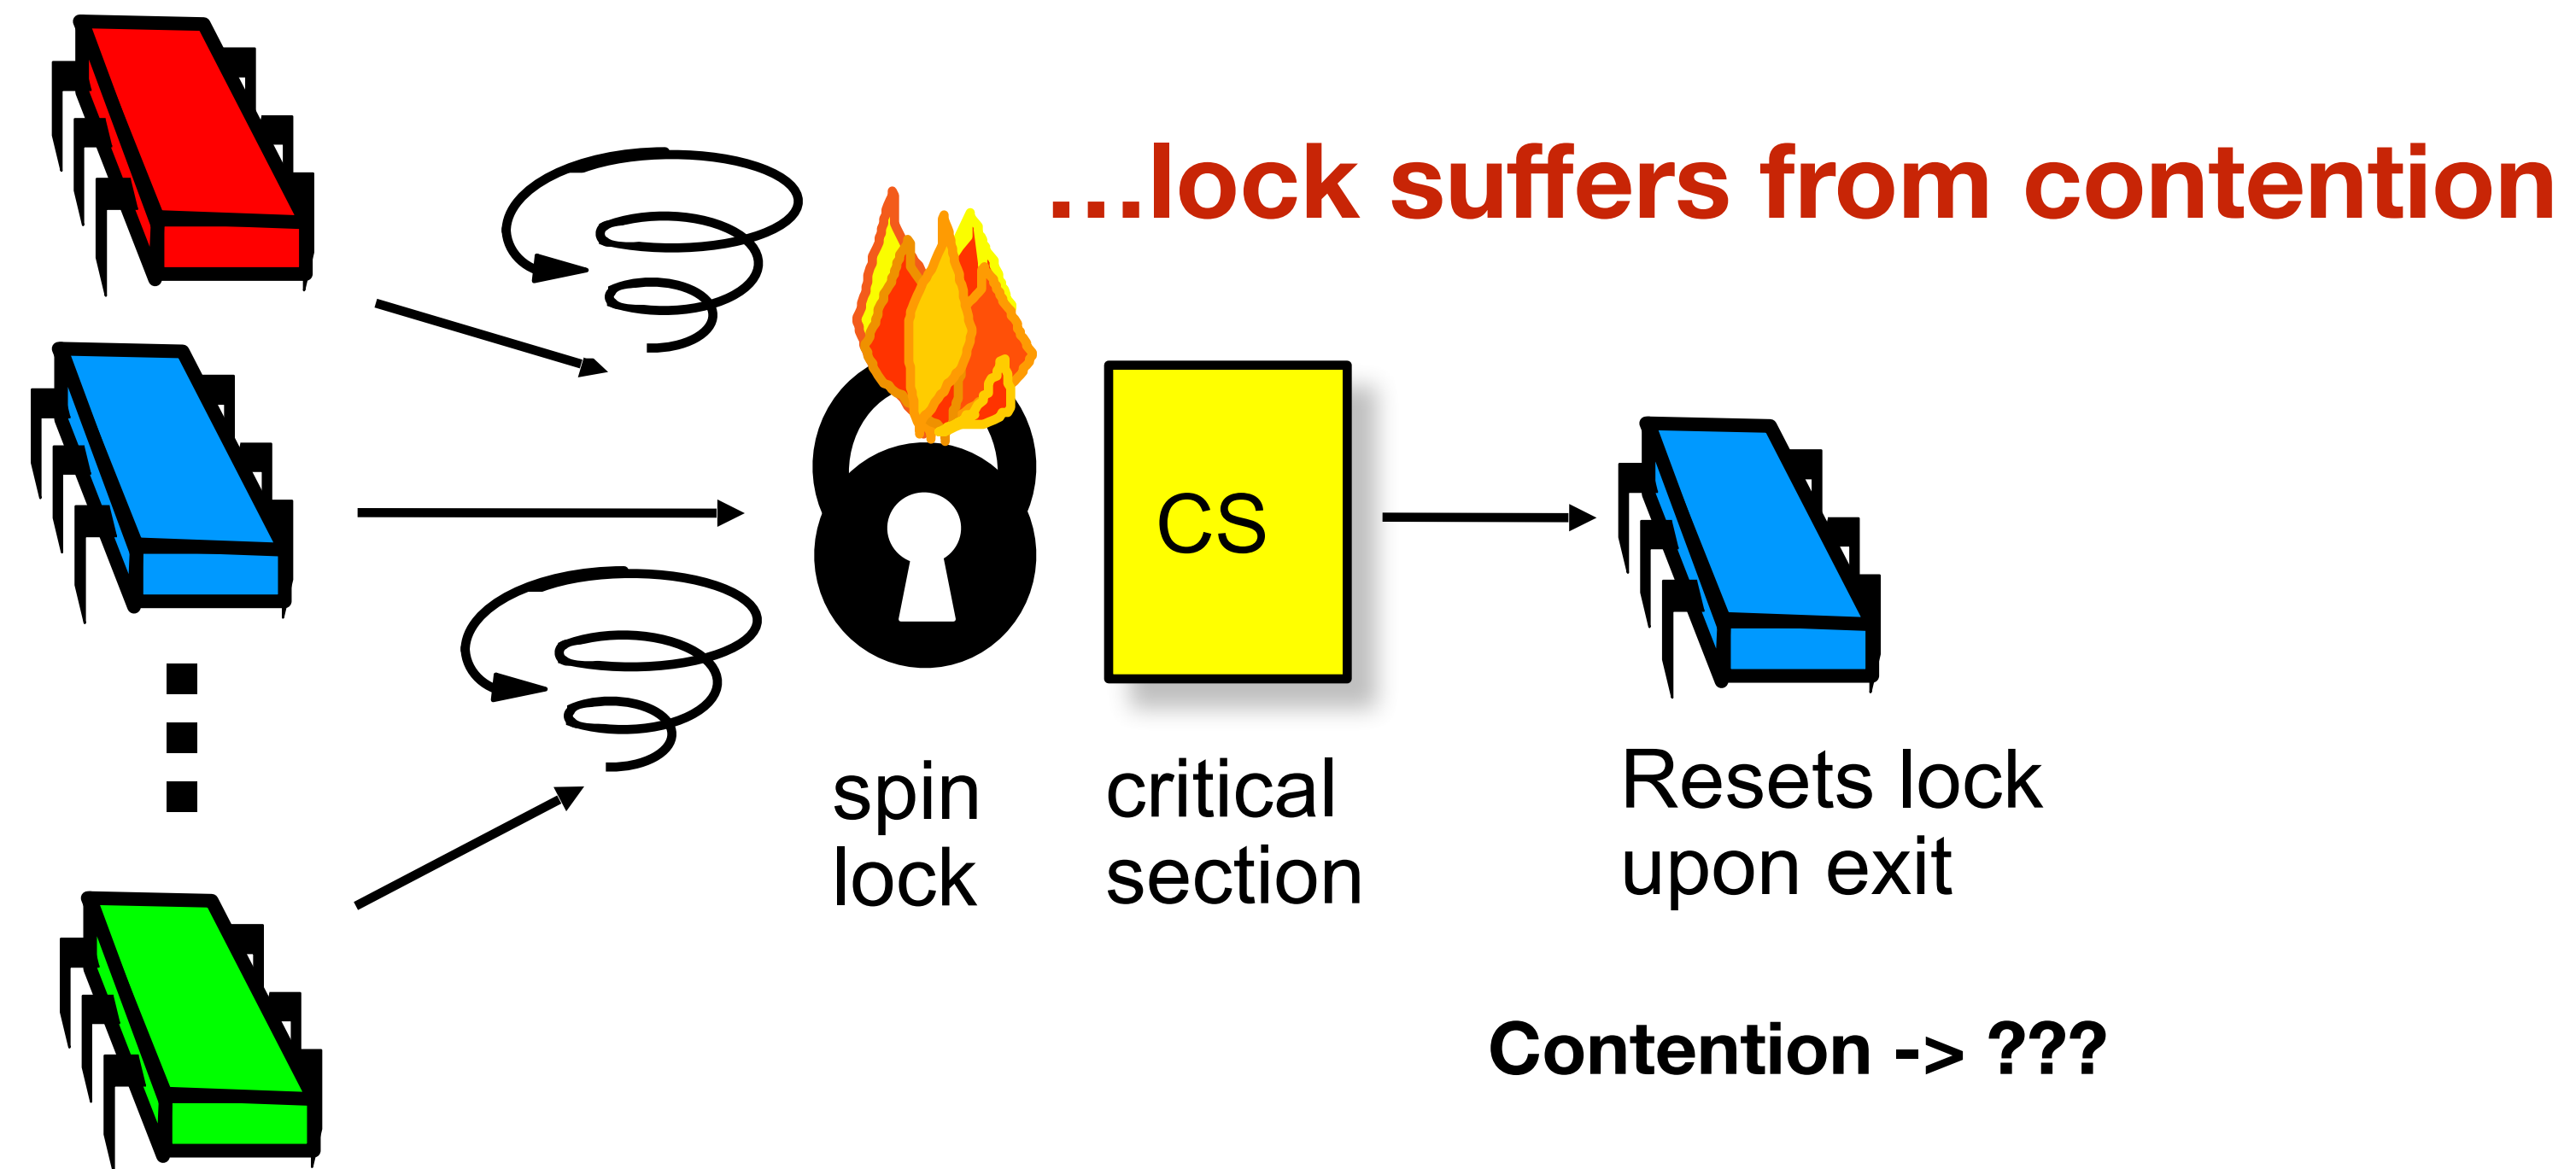


# Basic Spin-Lock



**Note: Contention and bottlenecking are separate phenomena**

# Basic Spin-Lock



# Review: Test-and-Set

- Boolean value
- Test-and-set (TAS)
  - Swap **true** with current value
  - Return value tells if prior value was **true** or **false**
- Can reset just by writing **false**
- TAS aka “getAndSet”

# Review: Test-and-Set

```
public class AtomicBoolean {
    boolean value;

    public synchronized boolean
    getAndSet(boolean newValue) {
        boolean prior = value;
        value = newValue;
        return prior;
    }
}
```



# Review: Test-and-Set

```
public class AtomicBoolean {  
    boolean value;  
  
    public synchronized boolean  
    getAndSet(boolean newValue) {  
        boolean prior = value;  
        value = newValue;  
        return prior;  
    }  
}
```

**Package**  
**java.util.concurrent.atomic**

# Review: Test-and-Set

```
public class AtomicBoolean {  
    boolean value;  
  
    public synchronized boolean  
    getAndSet(boolean newValue) {  
        boolean prior = value;  
        value = newValue;  
        return prior;  
    }  
}
```

**Swap old and new values**

# Review: Test-and-Set

```
AtomicBoolean lock  
  = new AtomicBoolean(false)  
...  
boolean prior = lock.getAndSet(true)
```

# Review: Test-and-Set

```
AtomicBoolean lock  
= new AtomicBoolean(false)  
...  
boolean prior = lock.getAndSet(true)
```

**Swapping in `true` is called “test-and-set”  
or TAS**

# Test-and-Set Locks

- Locking
  - Lock is free: value is false
  - Lock is taken: value is true
- Acquire lock by calling TAS
  - If result is false, you win
  - If result is true, you lose
- Release lock by writing false

# Test-and-set Lock

```
class TASlock {
    AtomicBoolean state =
        new AtomicBoolean(false);

    void lock() {
        while (state.getAndSet(true)) {}
    }

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

# Test-and-set Lock

```
class TASlock {  
    AtomicBoolean state =  
        new AtomicBoolean(false);  
  
    void lock() {  
        while (state.getAndSet(true)) {}  
    }  
  
    void unlock() {  
        state.set(false);  
    }  
}
```

**Lock state is AtomicBoolean**

# Test-and-set Lock

```
class TASlock {  
    AtomicBoolean state =  
        new AtomicBoolean(false);  
  
    void lock() {  
        while (state.getAndSet(true)) {}  
    }  
  
    void unlock() {  
        state.set(false);  
    }  
}
```

**Keep trying until lock acquired**

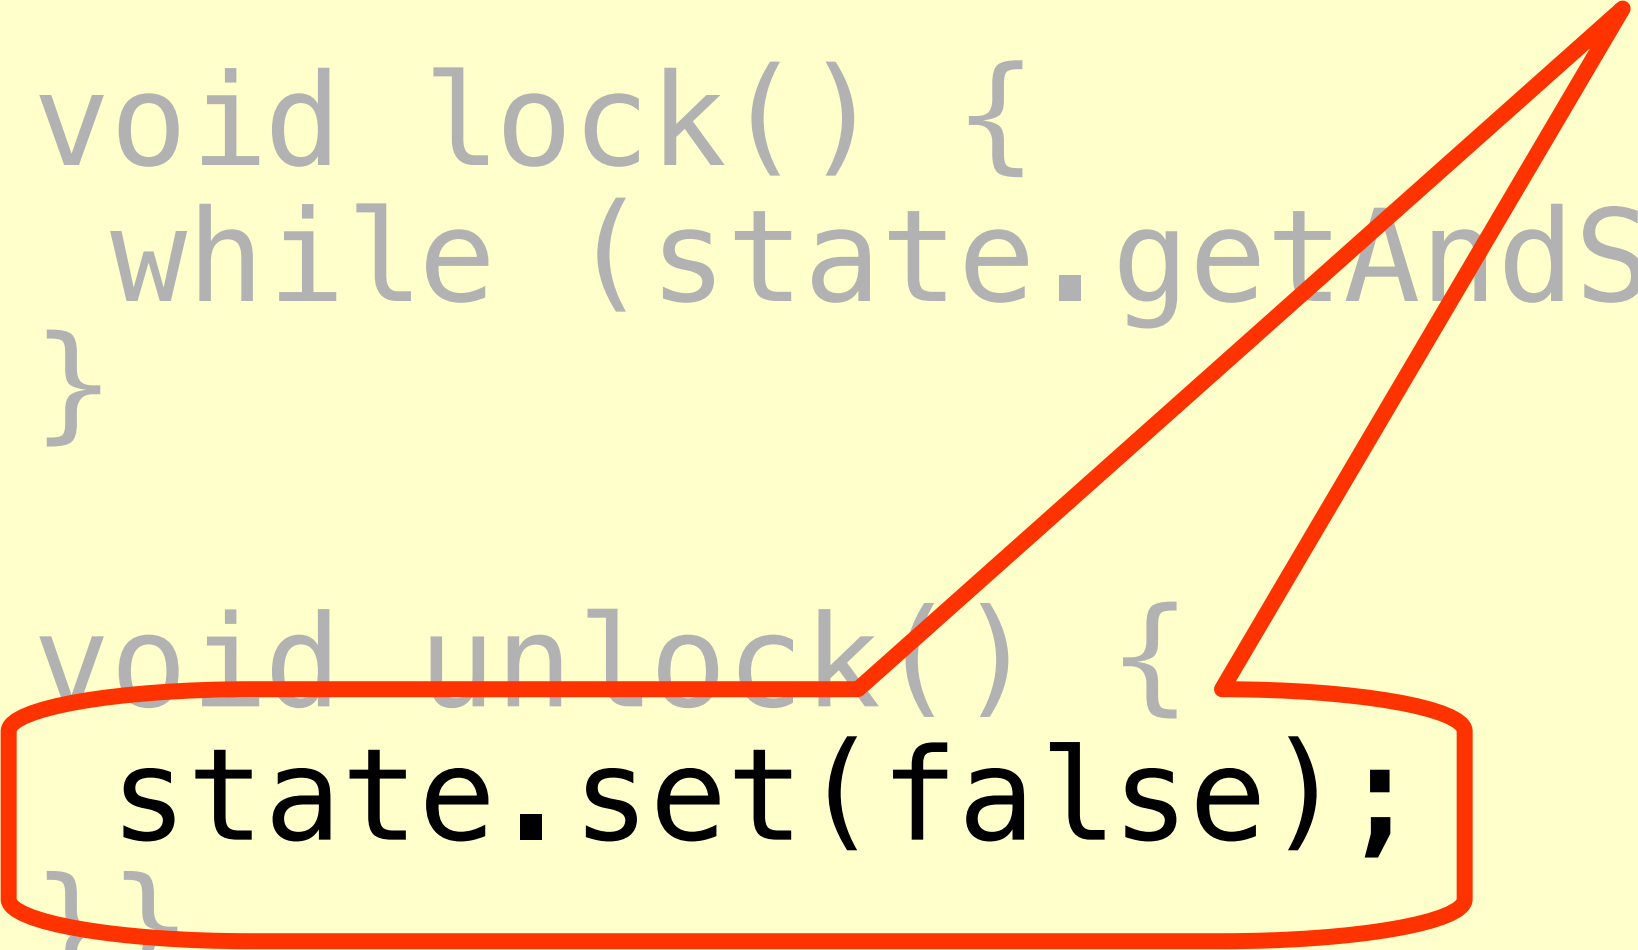


# Test-and-set Lock

```
class TASlock {
  AtomicBoolean state =
  new AtRelease lock by resetting state to false

  void lock() {
    while (state.getAndSet(true)) {}
  }

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



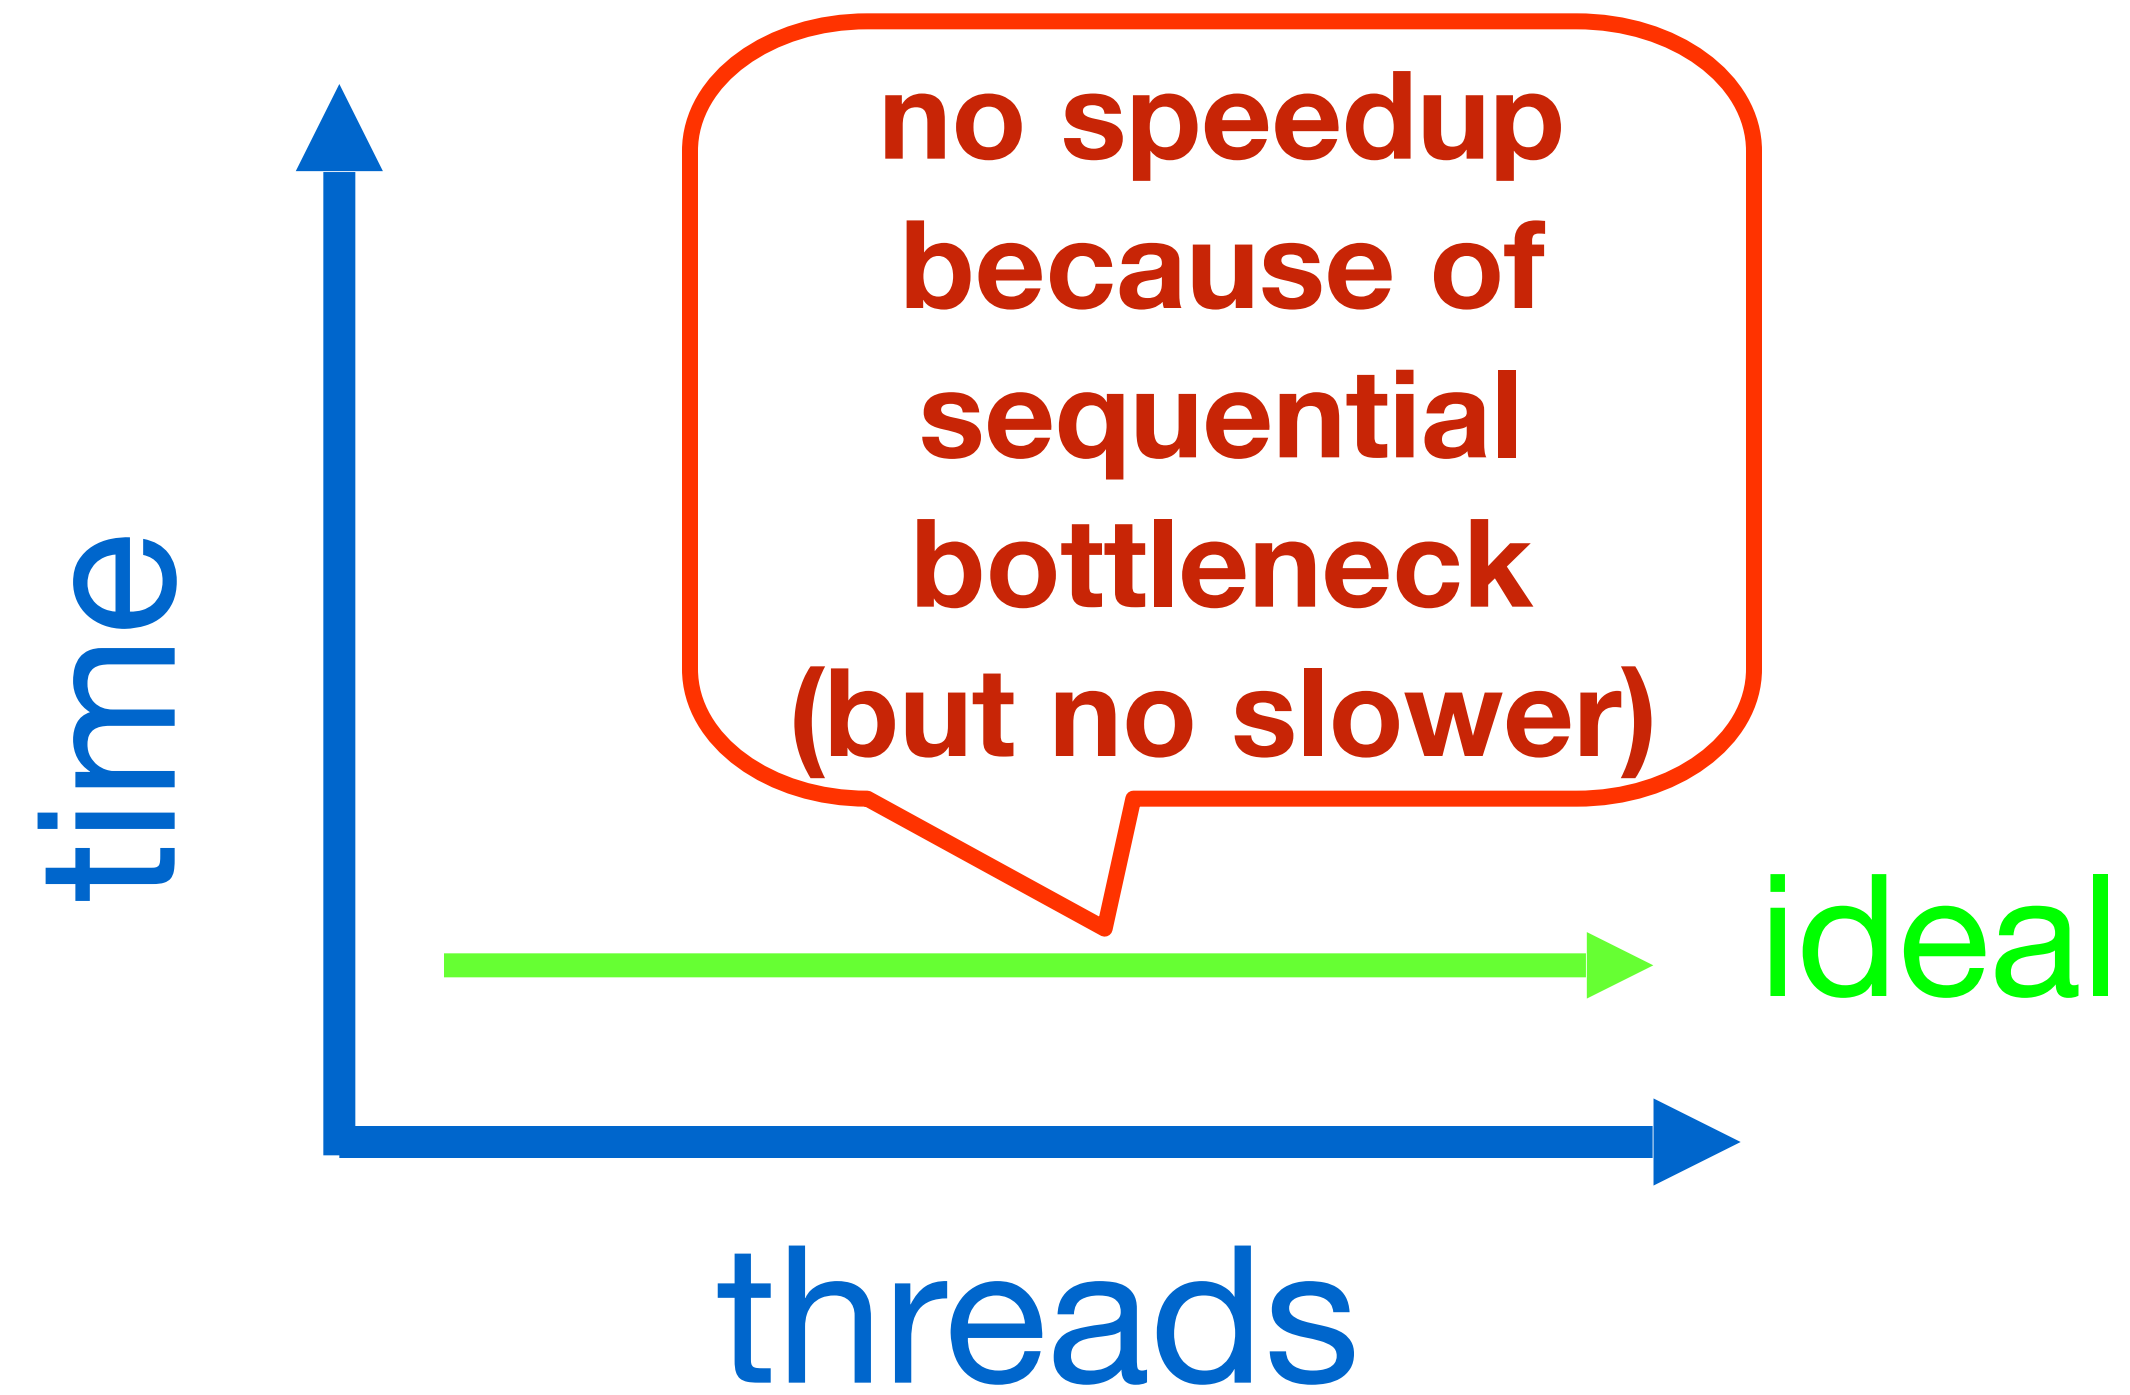
# Space Complexity

- TAS spin-lock has small “footprint”
- N thread spin-lock uses  $O(1)$  space
- As opposed to  $O(n)$  Peterson/Bakery
- How did we overcome the  $\Omega(n)$  lower bound?
- We used an atomic Read/Modify/Write operation
- Under the covers uses some special hardware features, *not* just the synchronized keyword as in the slides (if so, still  $O(n)$ ?)


# Performance

- Experiment
  - $n$  threads
  - Increment shared counter 1 million times
- How long should it take?
- How long does it take?

# Graph



# Mystery #1



**Adding MORE threads makes it SLOWER!**

# Test-and-Test-and-Set Locks

- Lurking stage
  - Wait until lock “looks” free
  - Spin while read returns true (lock taken)
- Pouncing state
  - As soon as lock “looks” available
  - Read returns false (lock free)
  - Call TAS to acquire lock
  - If TAS loses, back to lurking

# Test-and-test-and-set Lock

```
class TTASlock {
    AtomicBoolean state =
        new AtomicBoolean(false);

    void lock() {
        while (true) {
            while (state.get()) {}
            if (!state.getAndSet(true))
                return;
        }
    }
}
```

# Test-and-test-and-set Lock

```
class TTASlock {
  AtomicBoolean state =
    new AtomicBoolean(false);

  void lock() {
    while (true) {
      while (state.get()) {}
      if (!state.getAndSet(true))
        return;
    }
  }
}
```

**Wait until lock looks free**

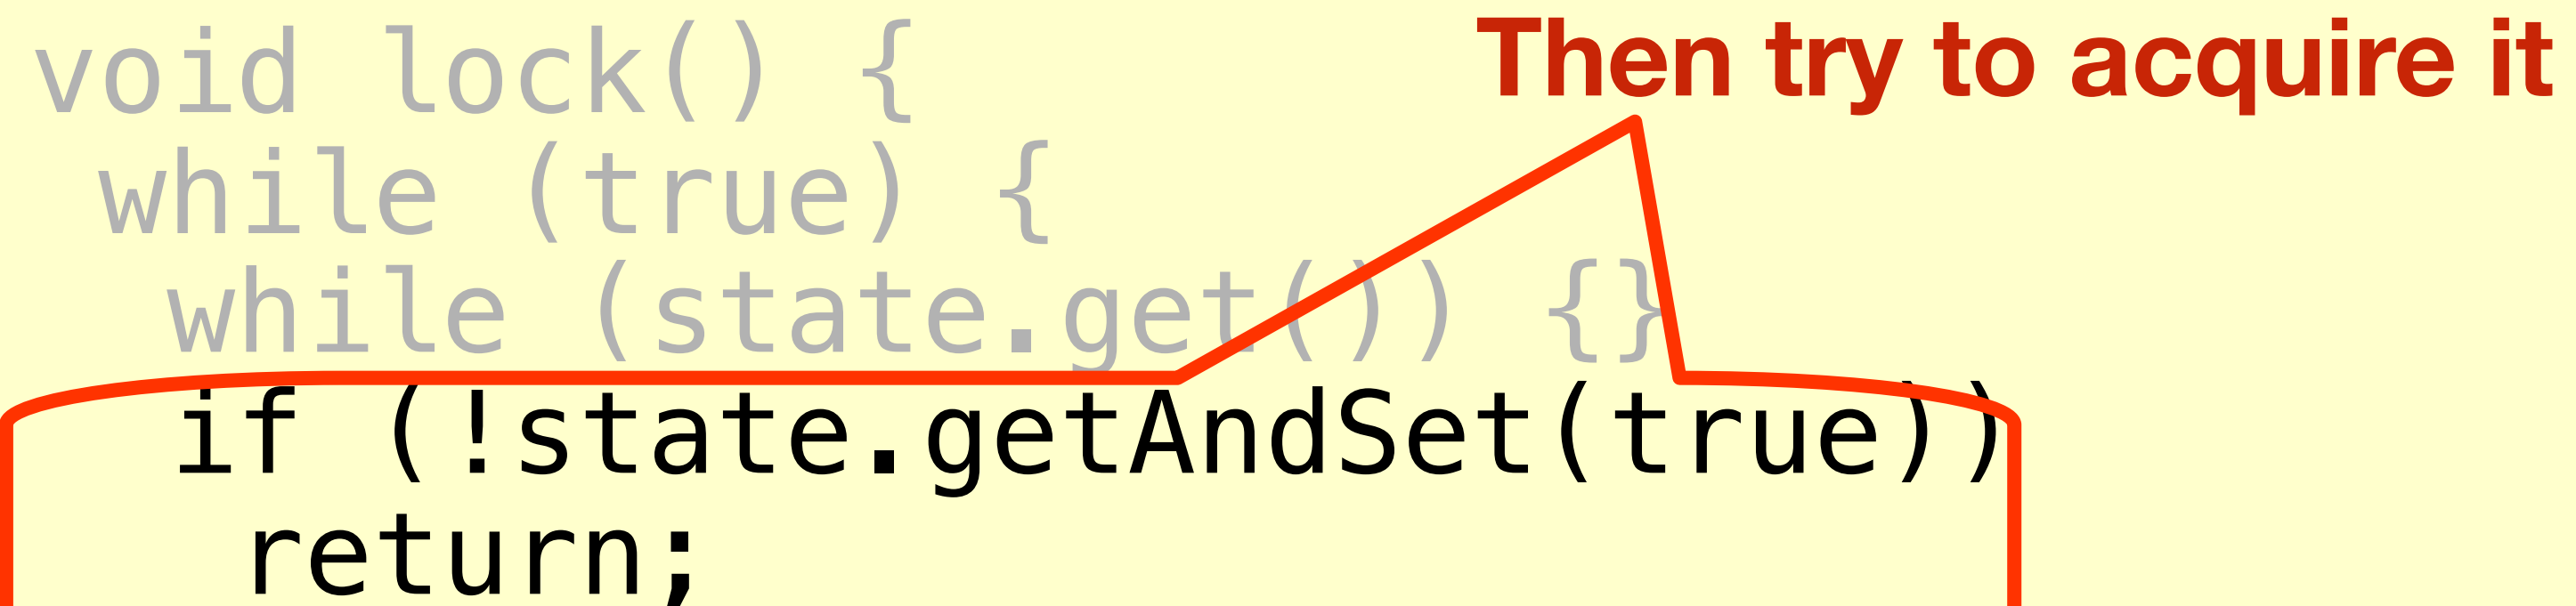


# Test-and-test-and-set Lock

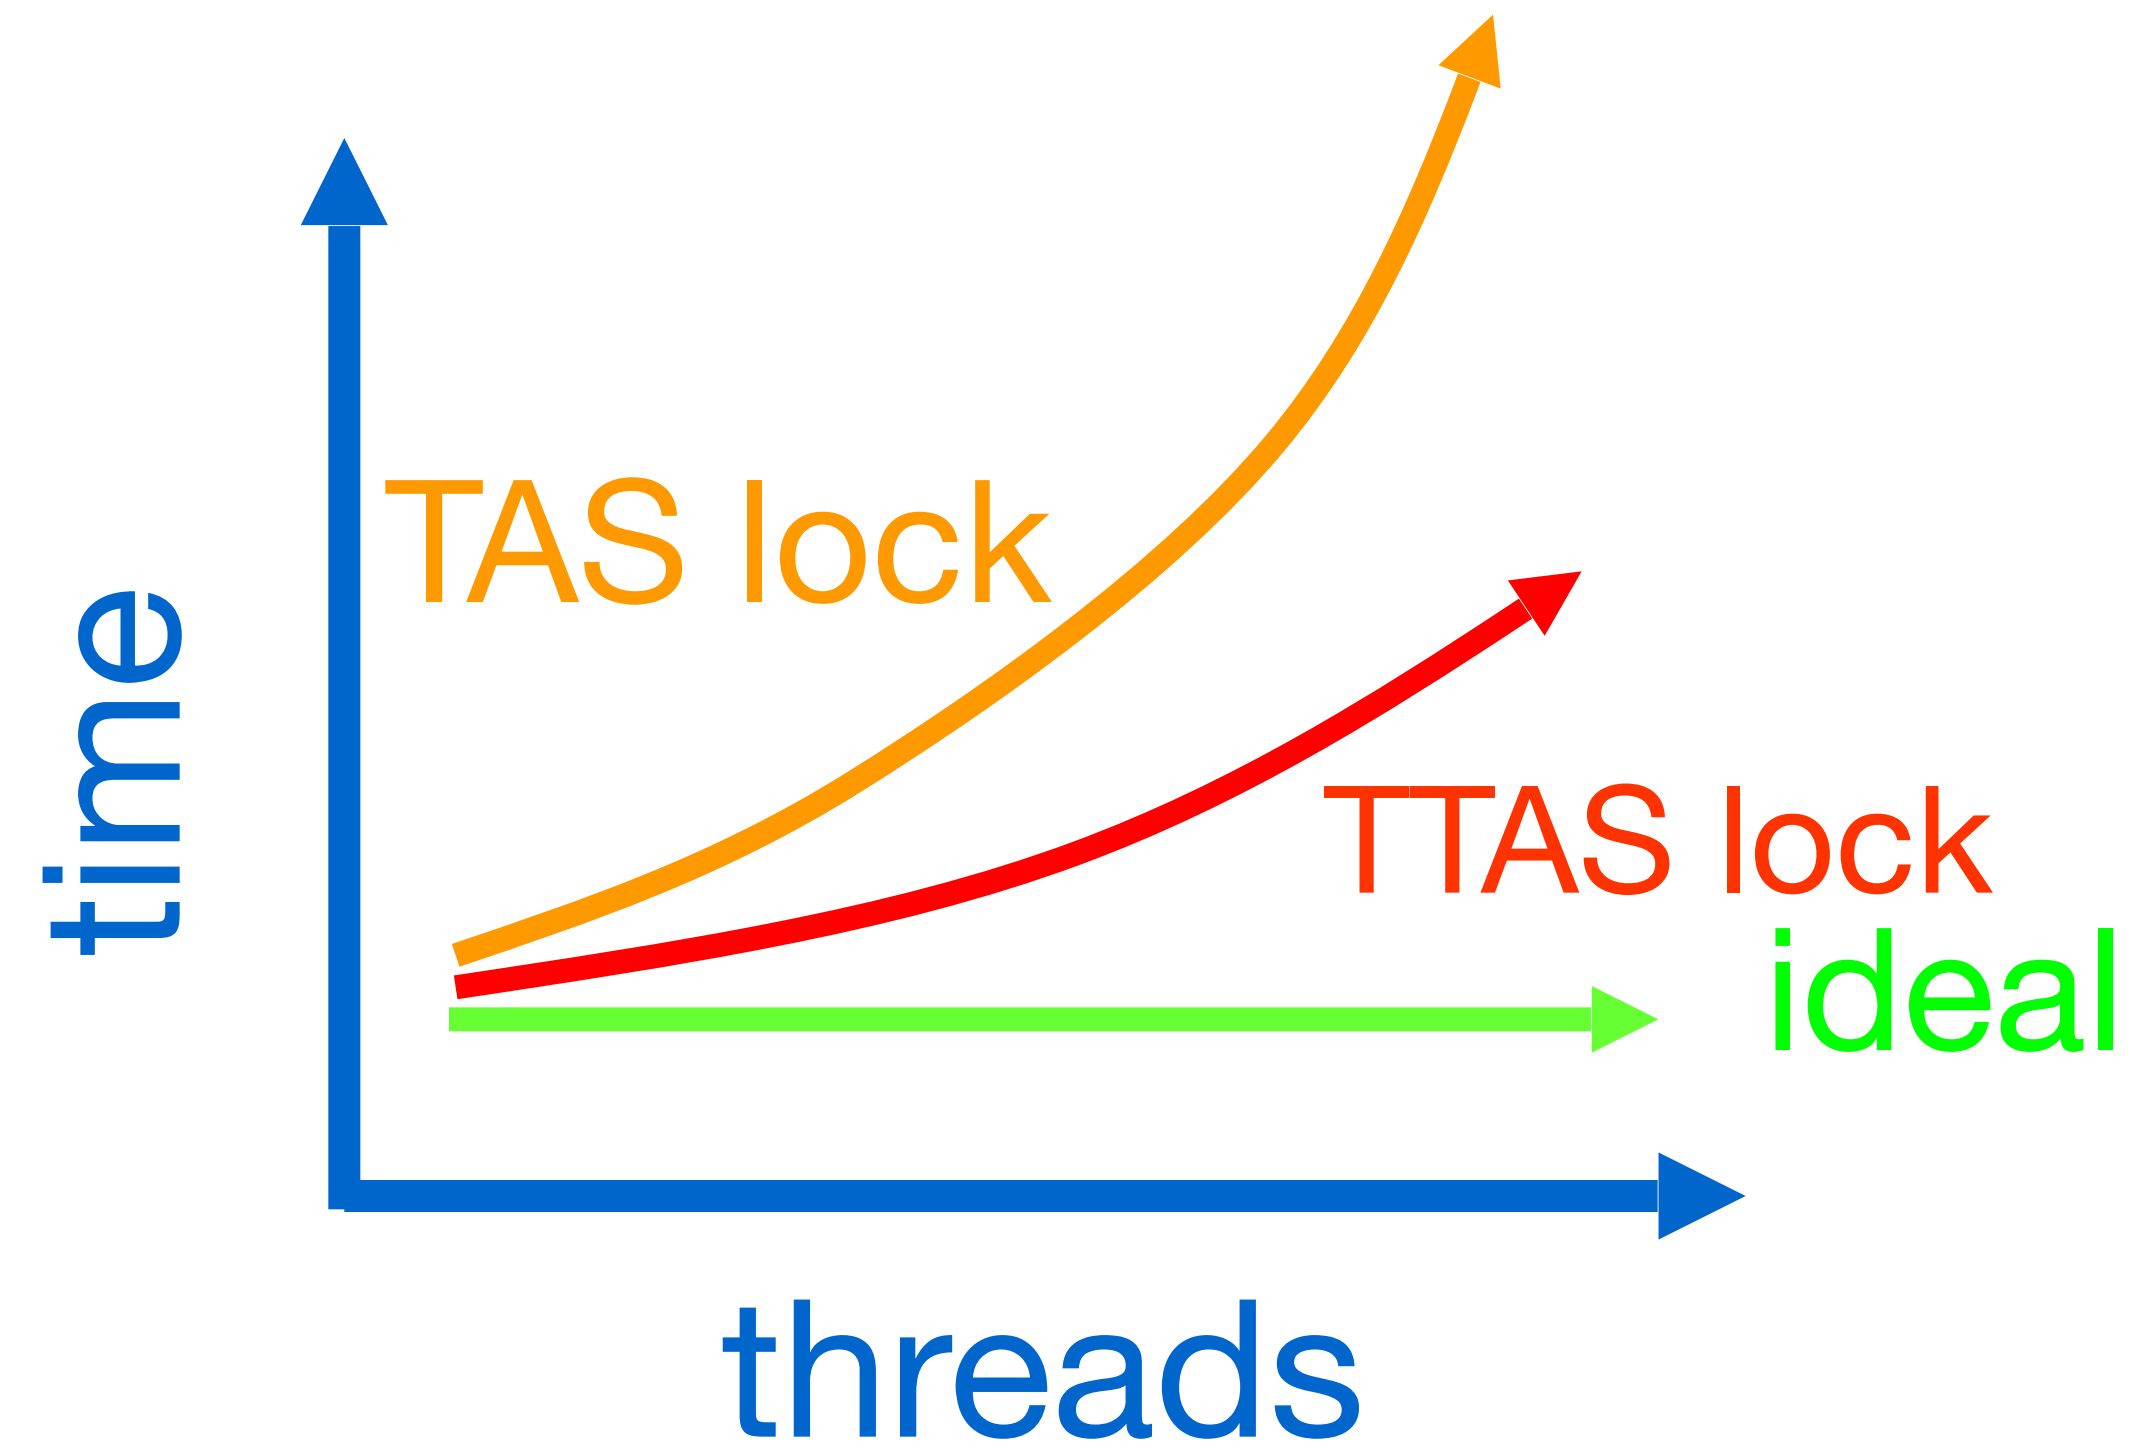
```
class TTASlock {  
    AtomicBoolean state =  
        new AtomicBoolean(false);
```

```
    void lock() {  
        while (true) {  
            while (state.get()) {}  
            if (!state.getAndSet(true))  
                return;  
        }  
    }  
}
```

**Then try to acquire it**



# Mystery #2



# Mystery

- Both
  - TAS and TTAS
  - Do the same thing (in our model)
- Except that
  - TTAS performs much better than TAS
  - Neither approaches ideal

# Opinion

- Our memory abstraction is broken
- TAS & TTAS methods
  - Are provably the same (in our model)
  - Except they aren't (in field tests)
- Need a more detailed model ...



**CPU Architectures**

# Kinds of Architectures

- SISD (Uniprocessor)
  - Single instruction stream
  - Single data stream
- SIMD (Vector)
  - Single instruction
  - Multiple data
- MIMD (Multiprocessors)
  - Multiple instruction
  - Multiple data.

# Kinds of Architectures

- SISD (Uniprocessor)
  - Single instruction stream
  - Single data stream


- SIMD (Vector)
  - Single instruction
  - Multiple data

**Most modern desktop/laptop CPUs**




- MIMD (Multiprocessors)
  - Multiple instruction
  - Multiple data.

# MIMD Bus-Based Architectures





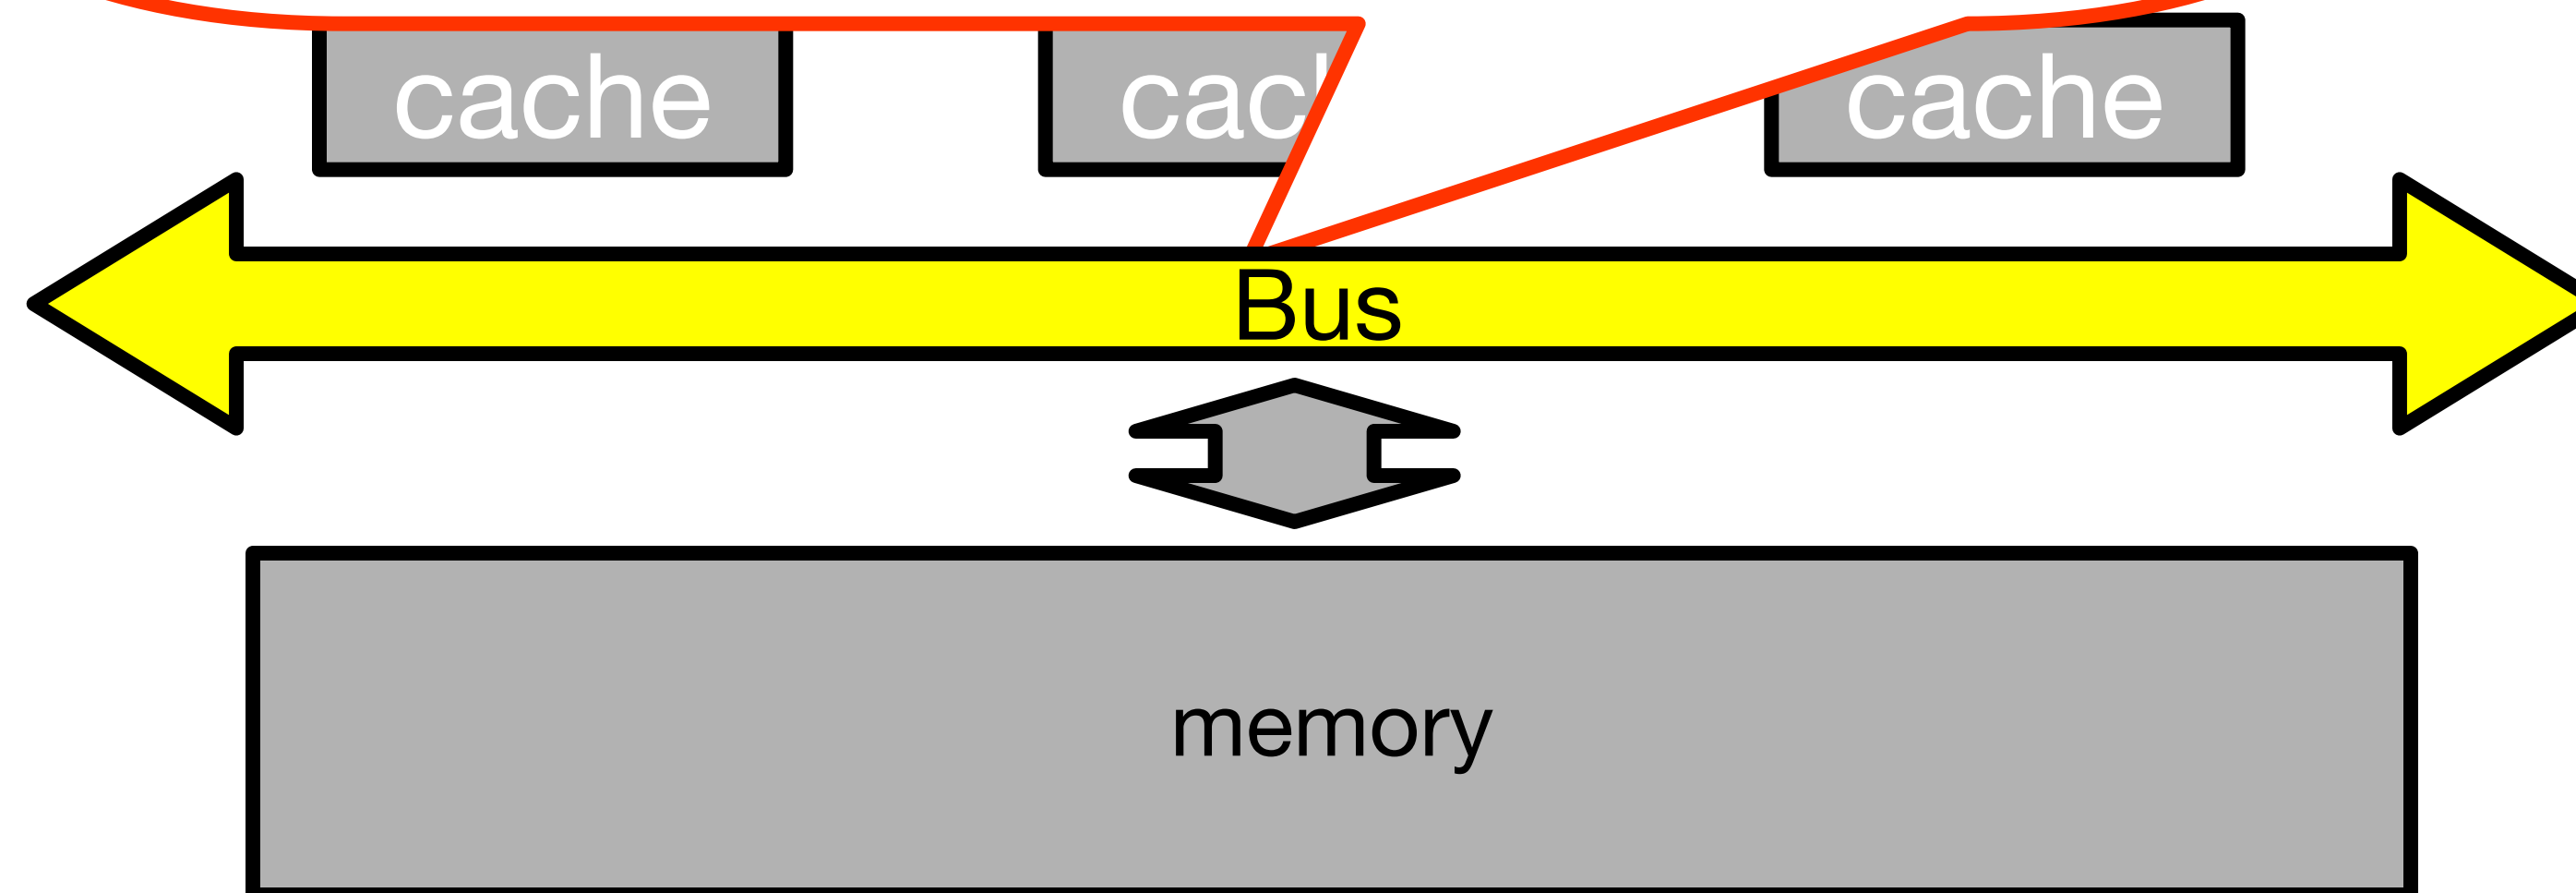
# Bus-Based Architectures



# Bus-Based Architectures

## Shared Bus


- Broadcast medium
- One broadcaster at a time
- Processors and memory all “snoop”



# Bus-Based

## Per-Processor Caches


- Small
- Fast: 1 or 2 cycles
- Address & state information



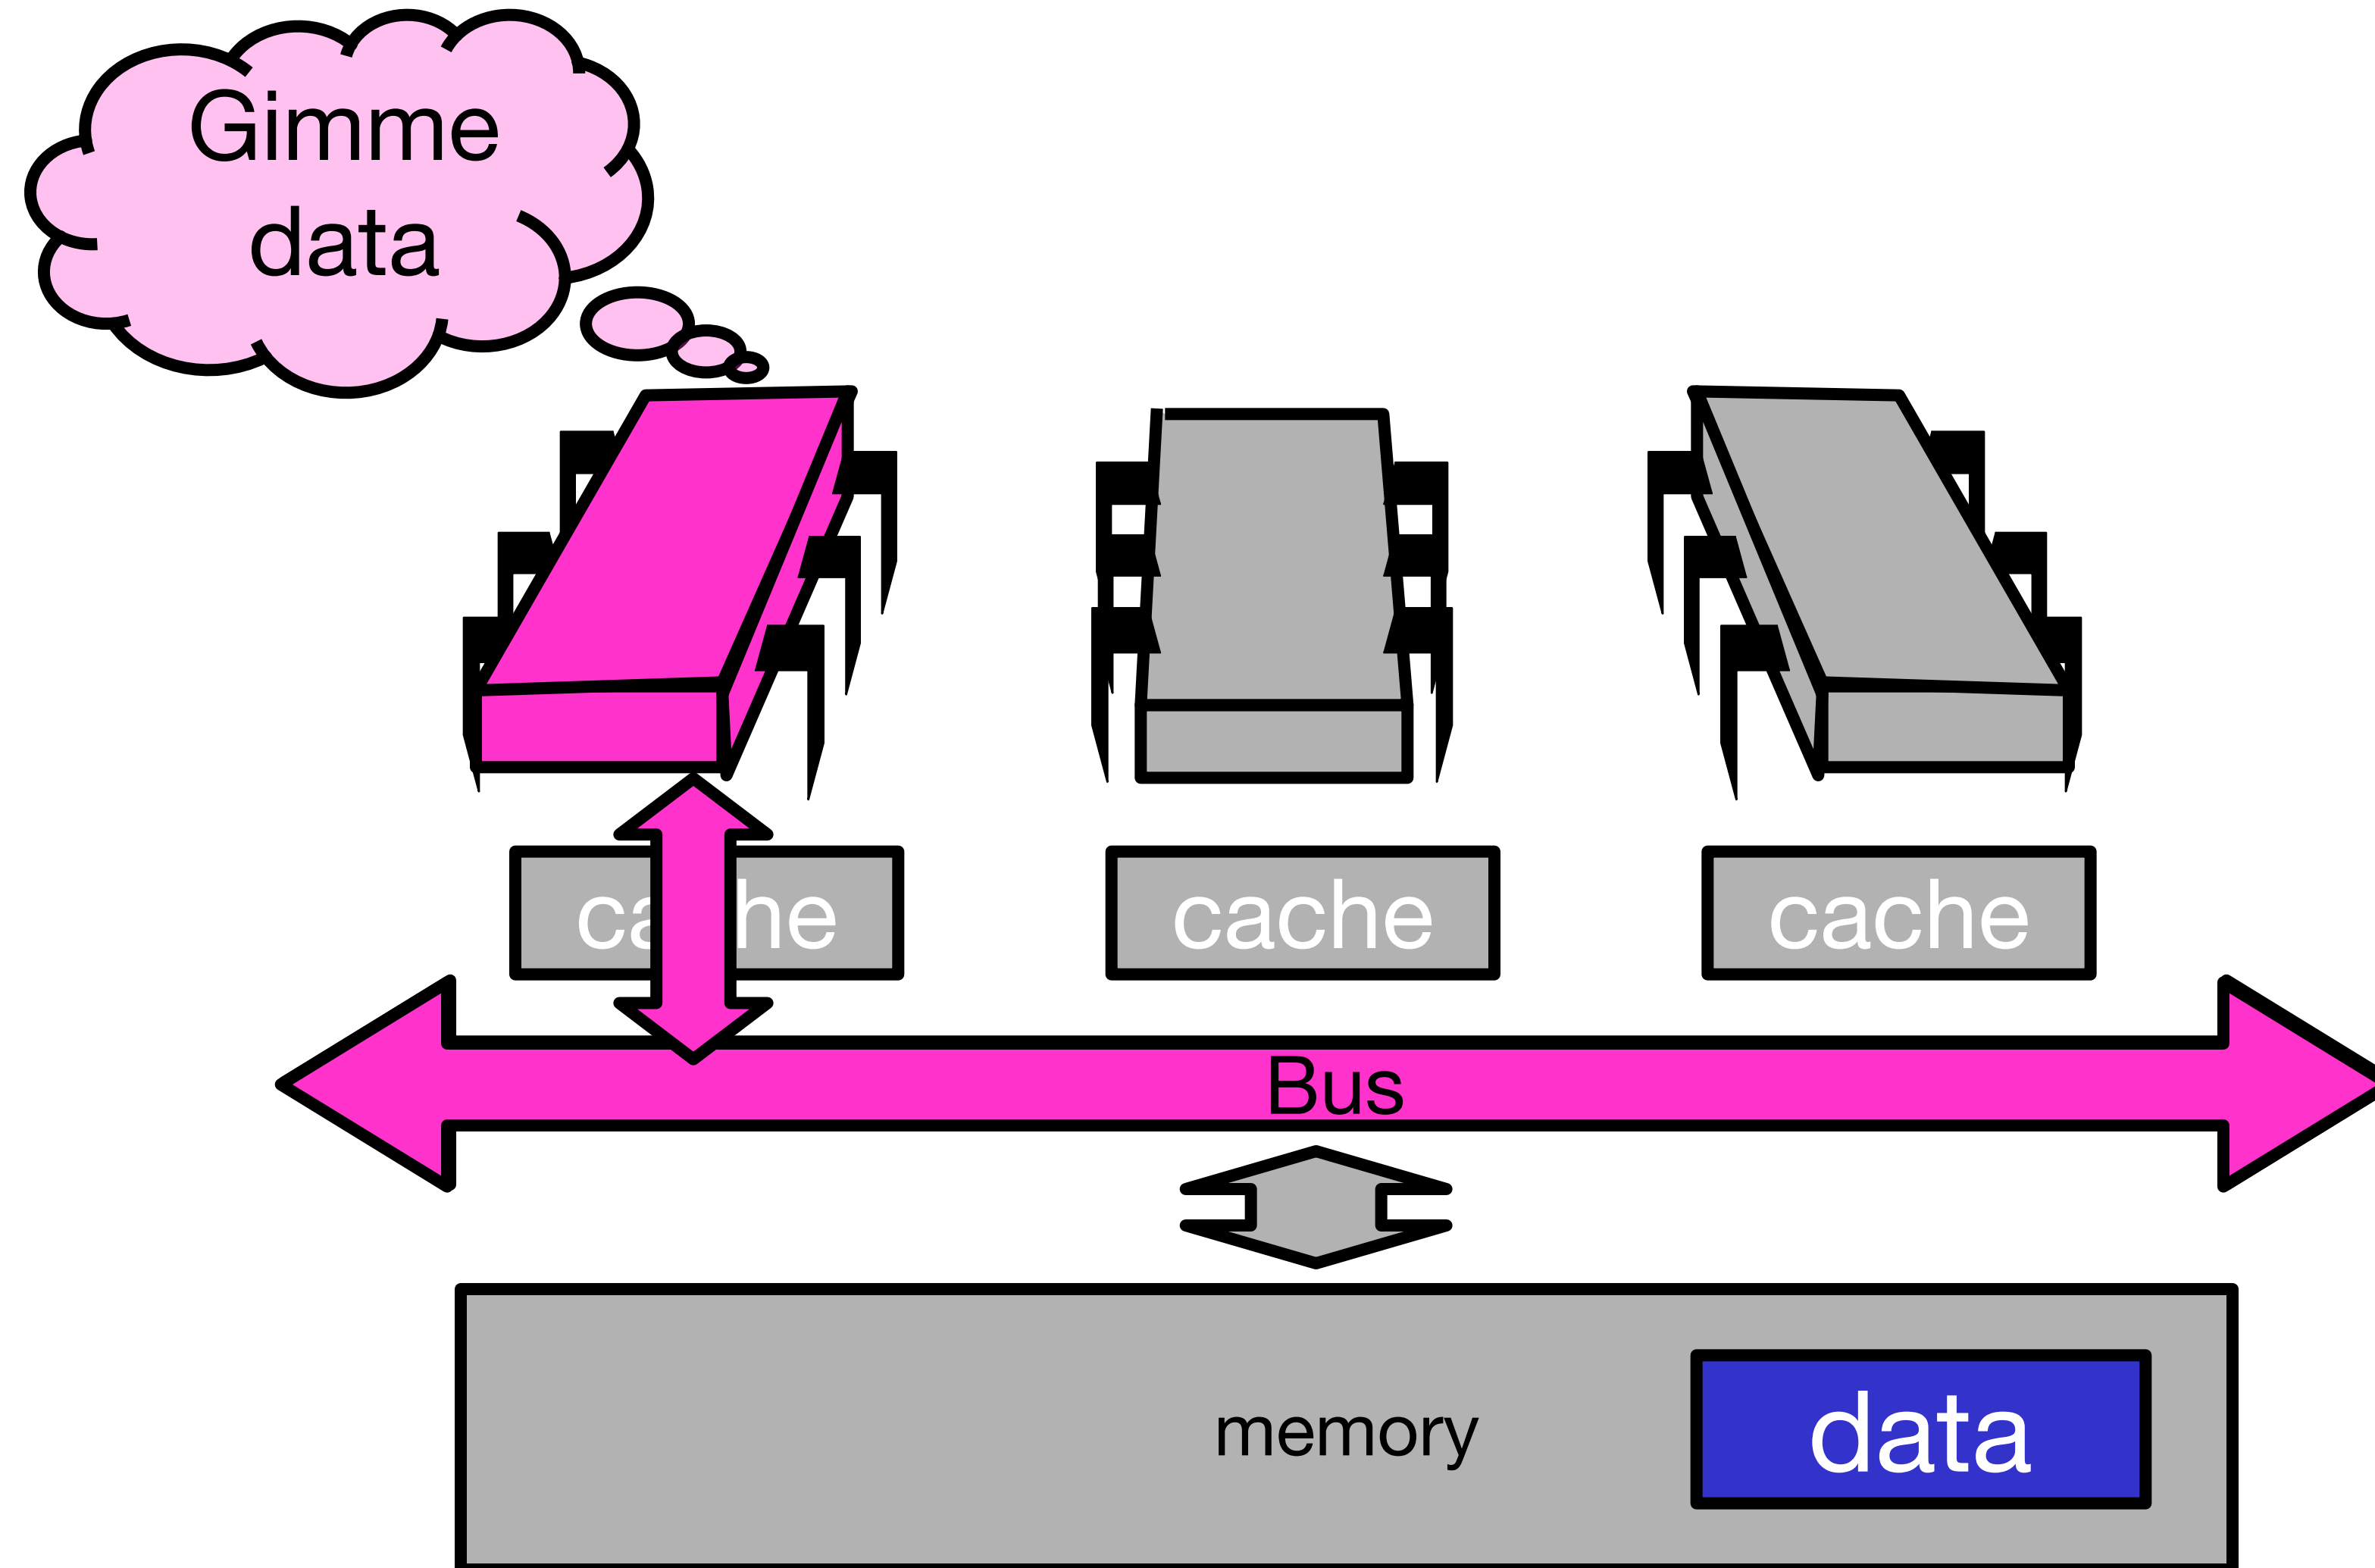
# Jargon Watch

- Cache hit
  - “I found what I wanted in my cache”
  - Good Thing™
- Cache miss
  - “I had to shlep all the way to memory for that data”
  - Bad Thing™

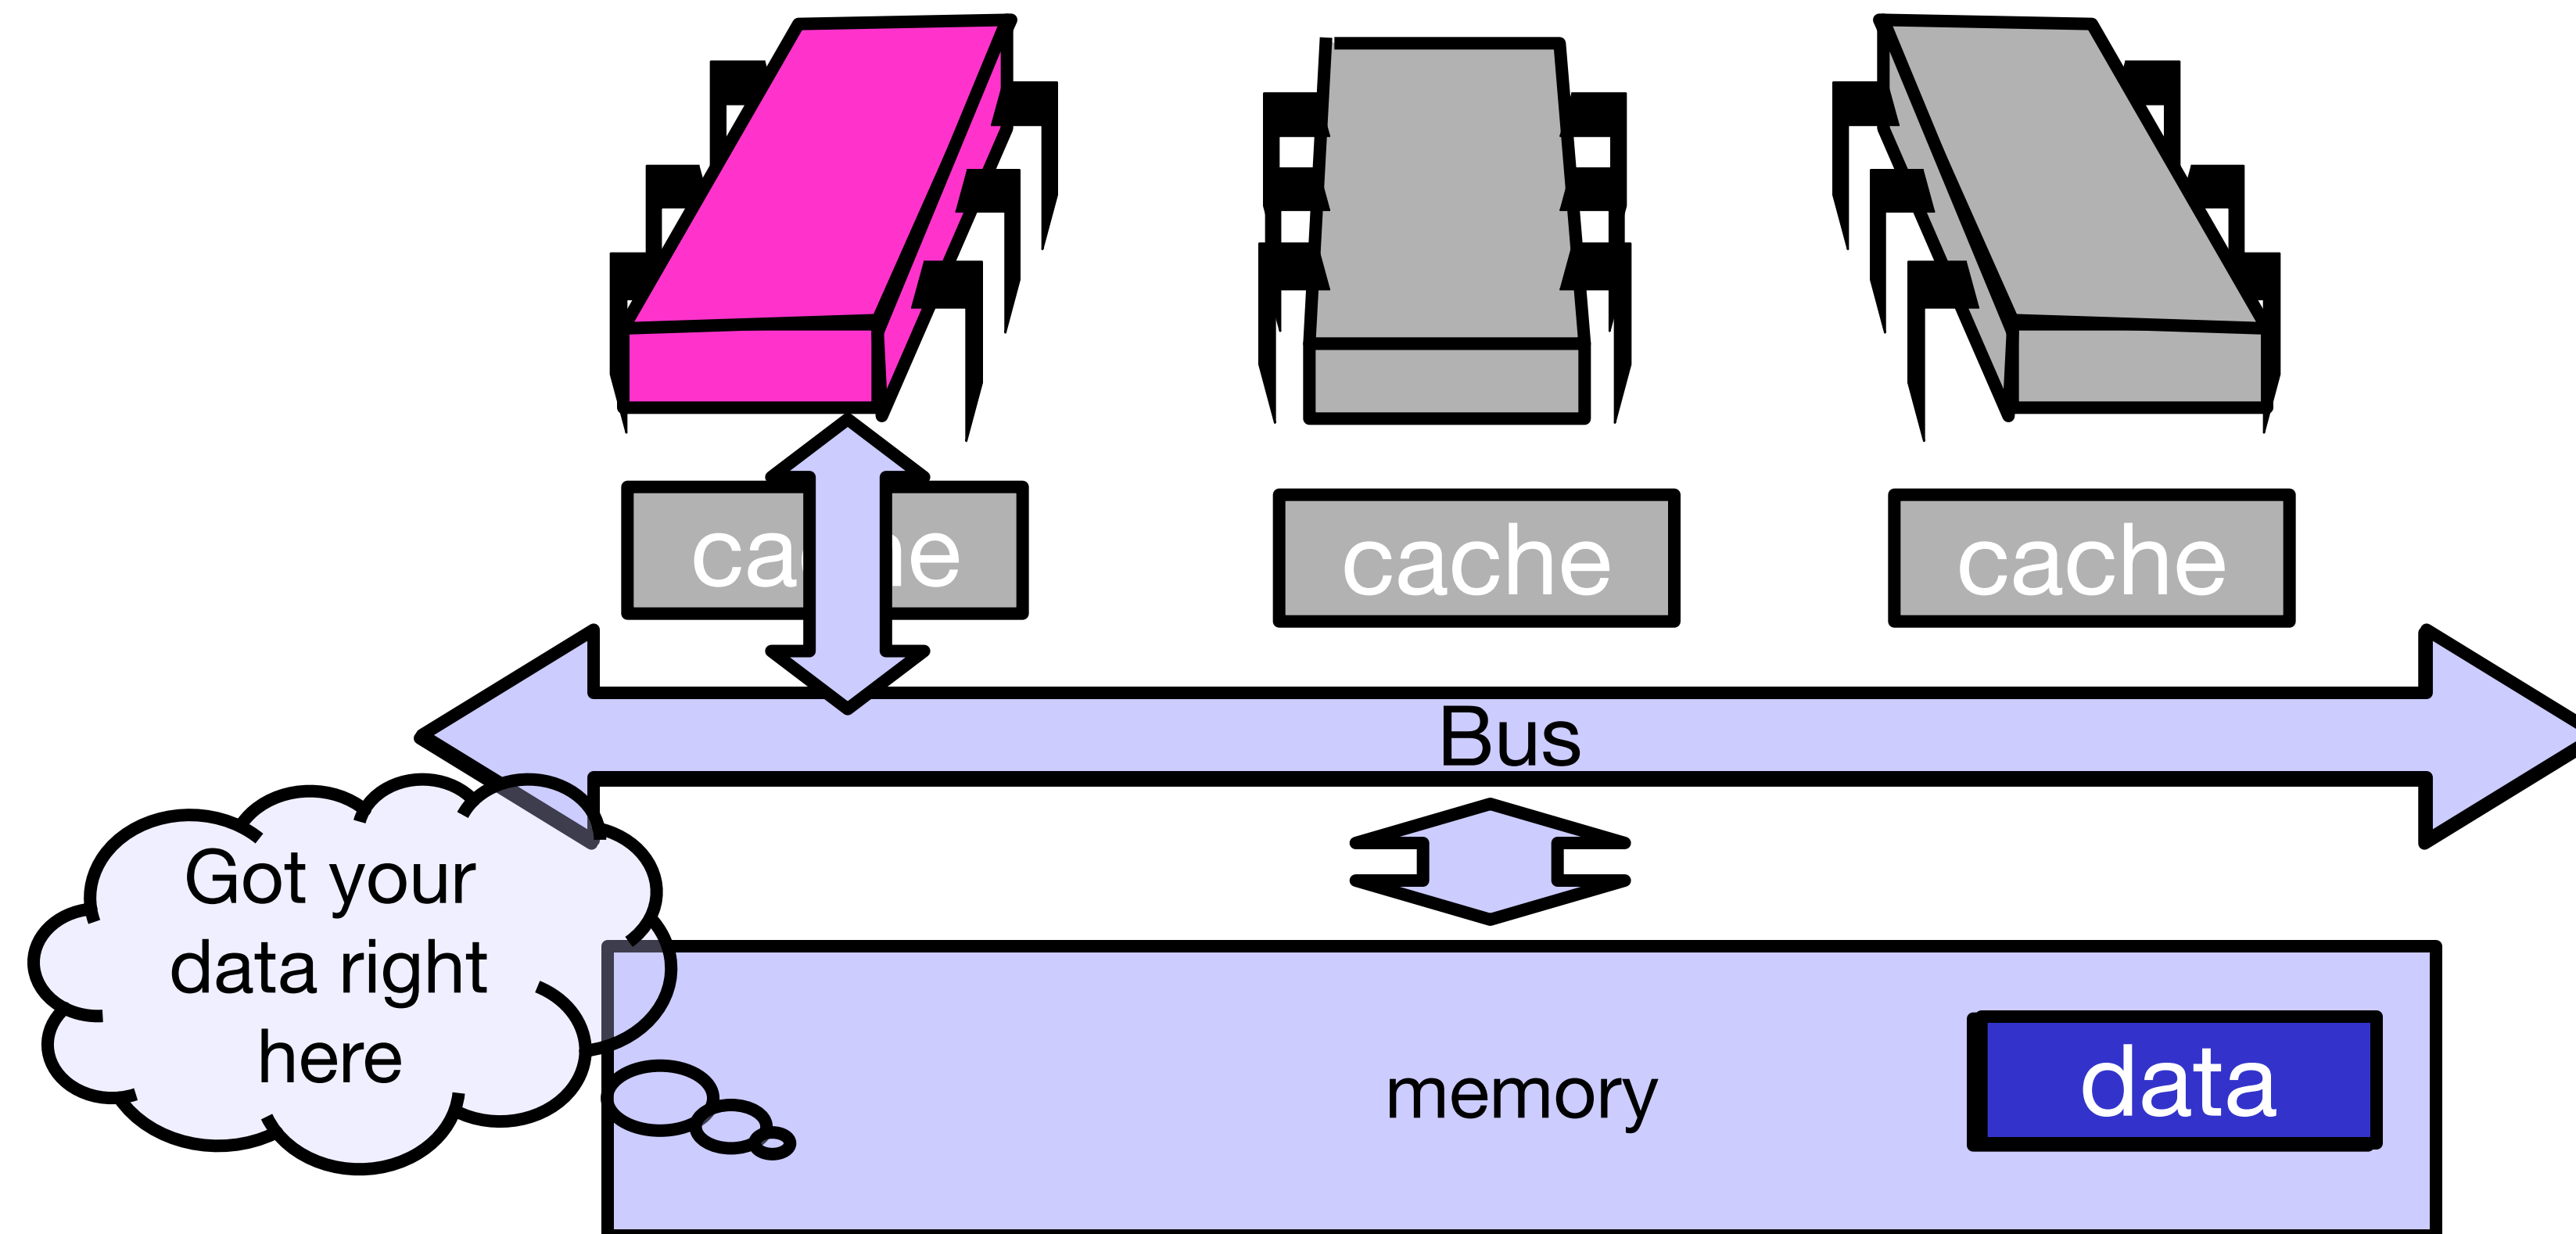
# Processor Issues Load Request



# Processor Issues Load Request




# Memory Responds




Art of Multiprocessor Programming

# Processor Issues Load Request

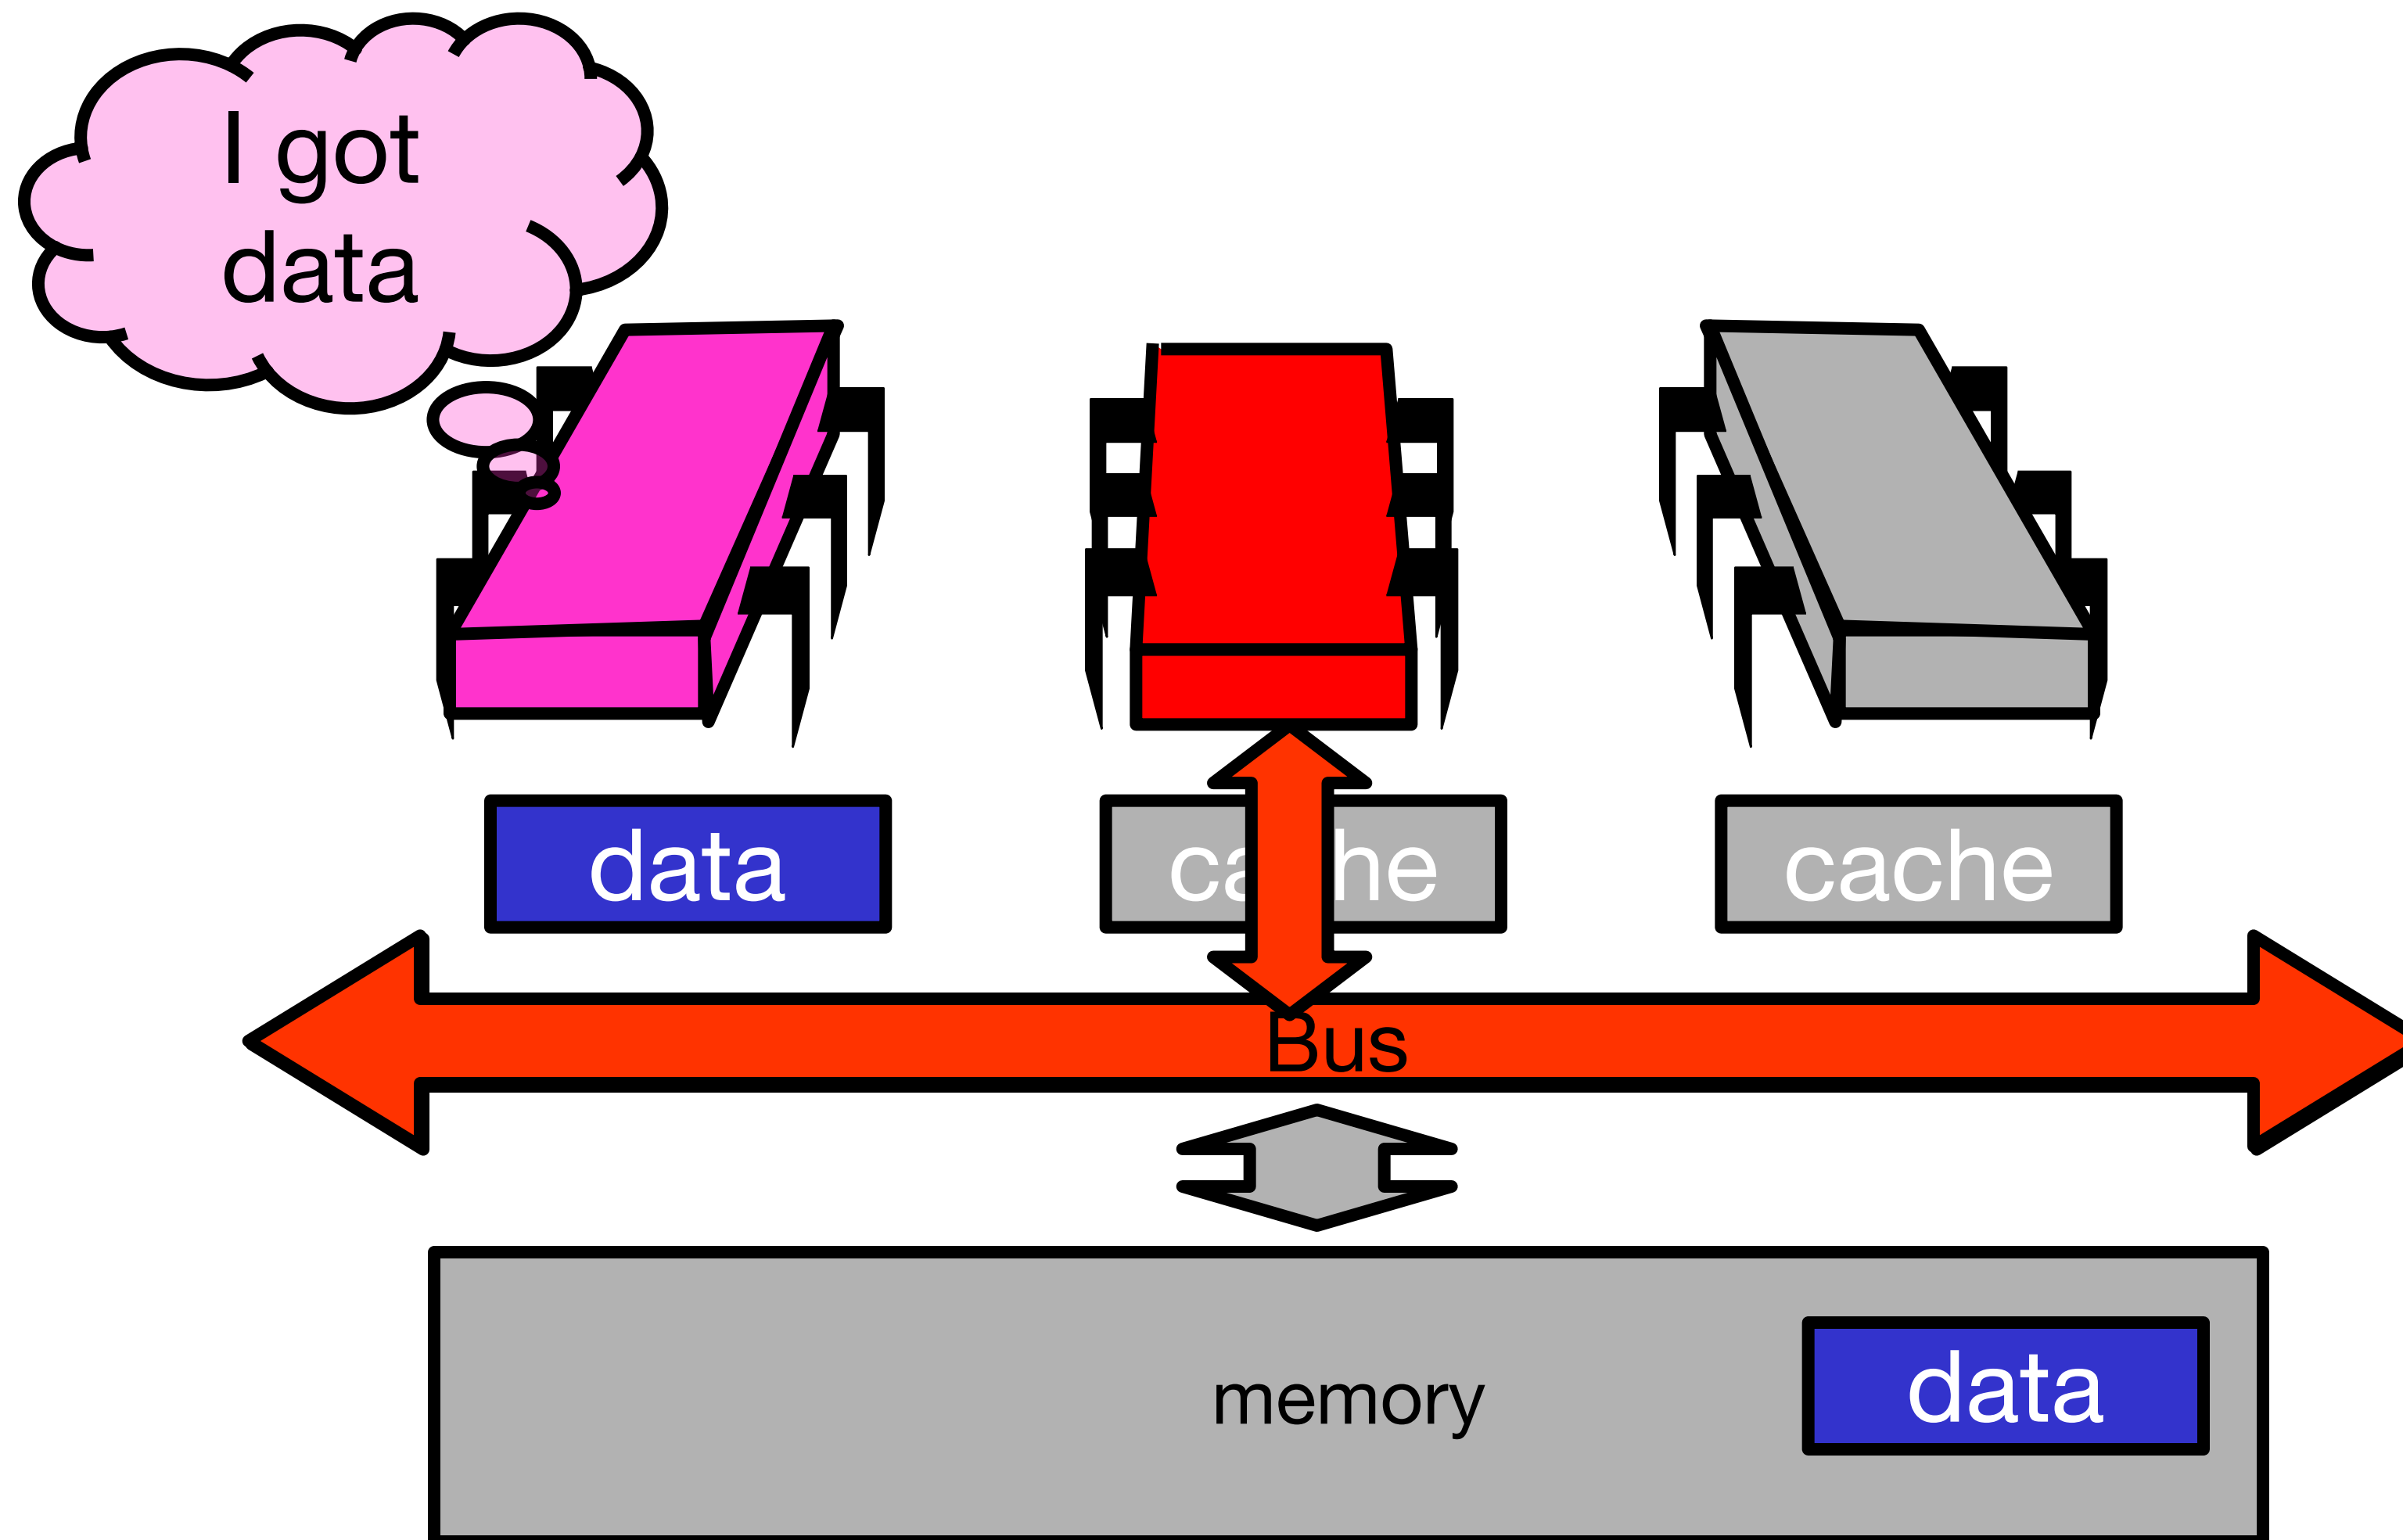




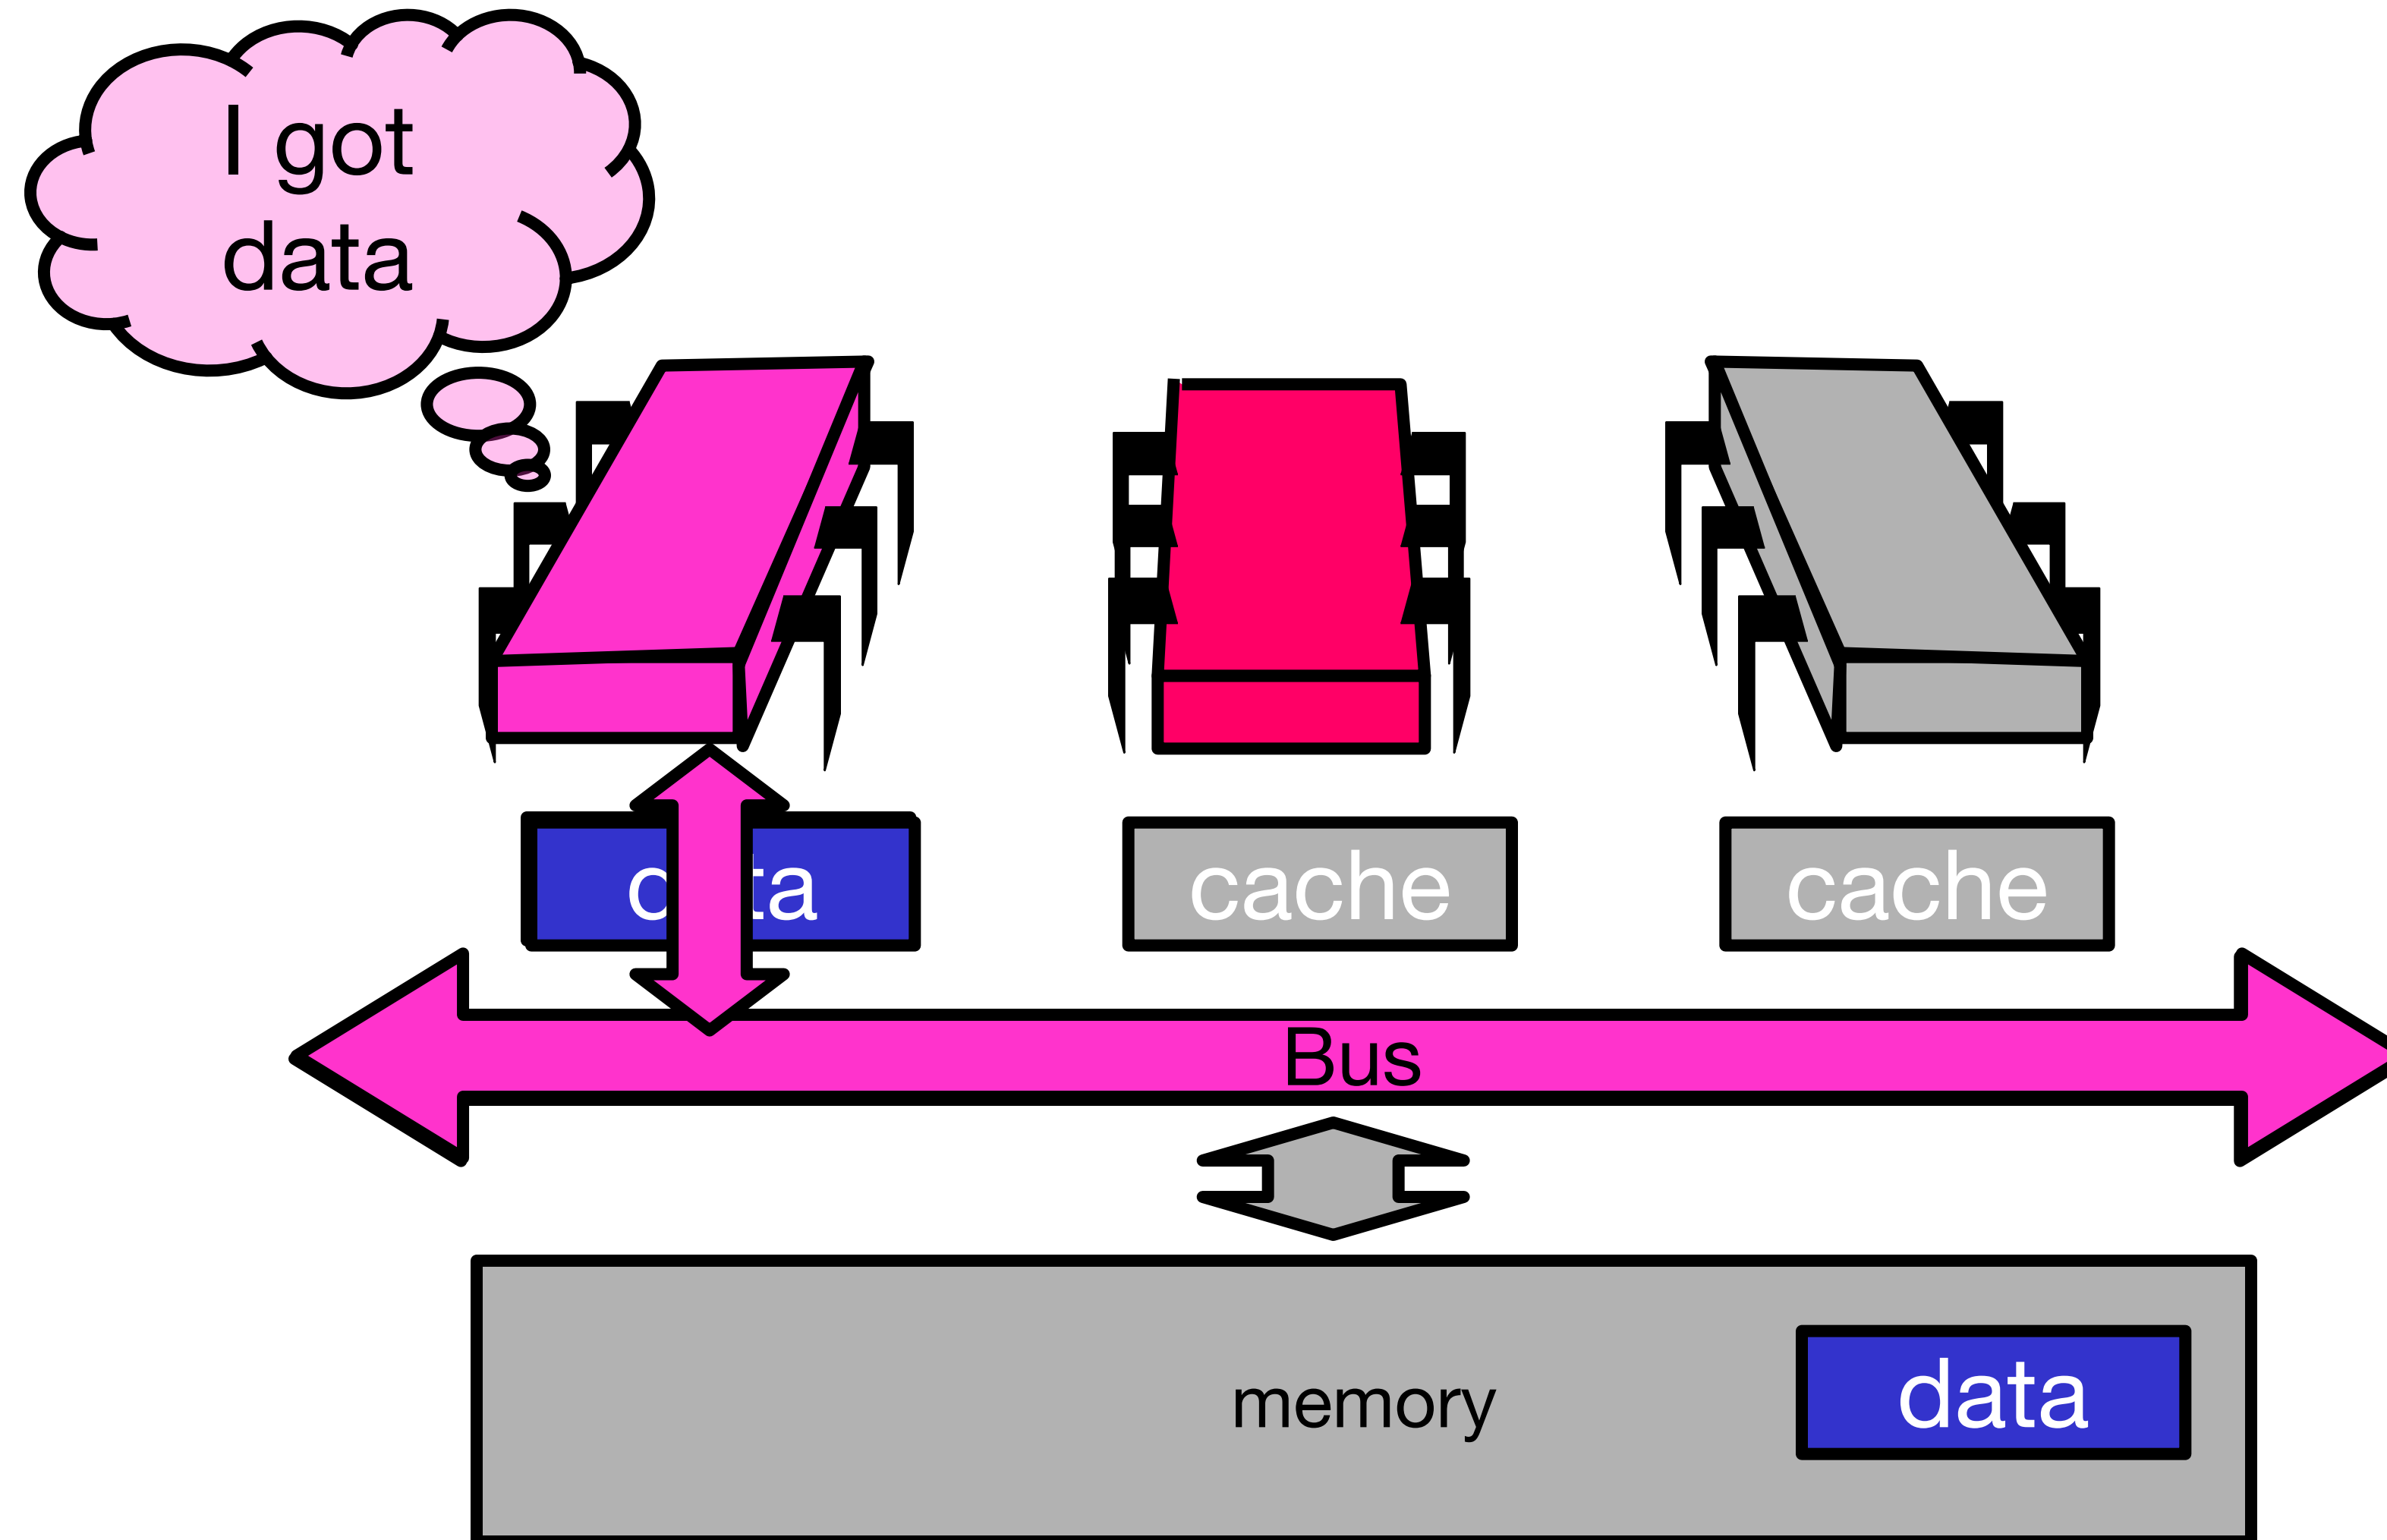
# Processor Issues Load Request




# Processor Issues Load Request




# Other Processor Responds




# Other Processor Responds




# Modify Cached Data




# Modify Cached Data



# Modify Cached Data



# Modify Cached Data





# Cache Coherence

- We have lots of copies of data
  - Original copy in memory
  - Cached copies at processors
- Some processor modifies its own copy
  - What do we do with the others?
  - How to avoid confusion?


# Write-Back Caches

- Accumulate changes in cache
- Write back when needed
  - Need the cache for something else
  - Another processor wants it
- On first modification
  - Invalidate other entries
  - Requires non-trivial protocol ...

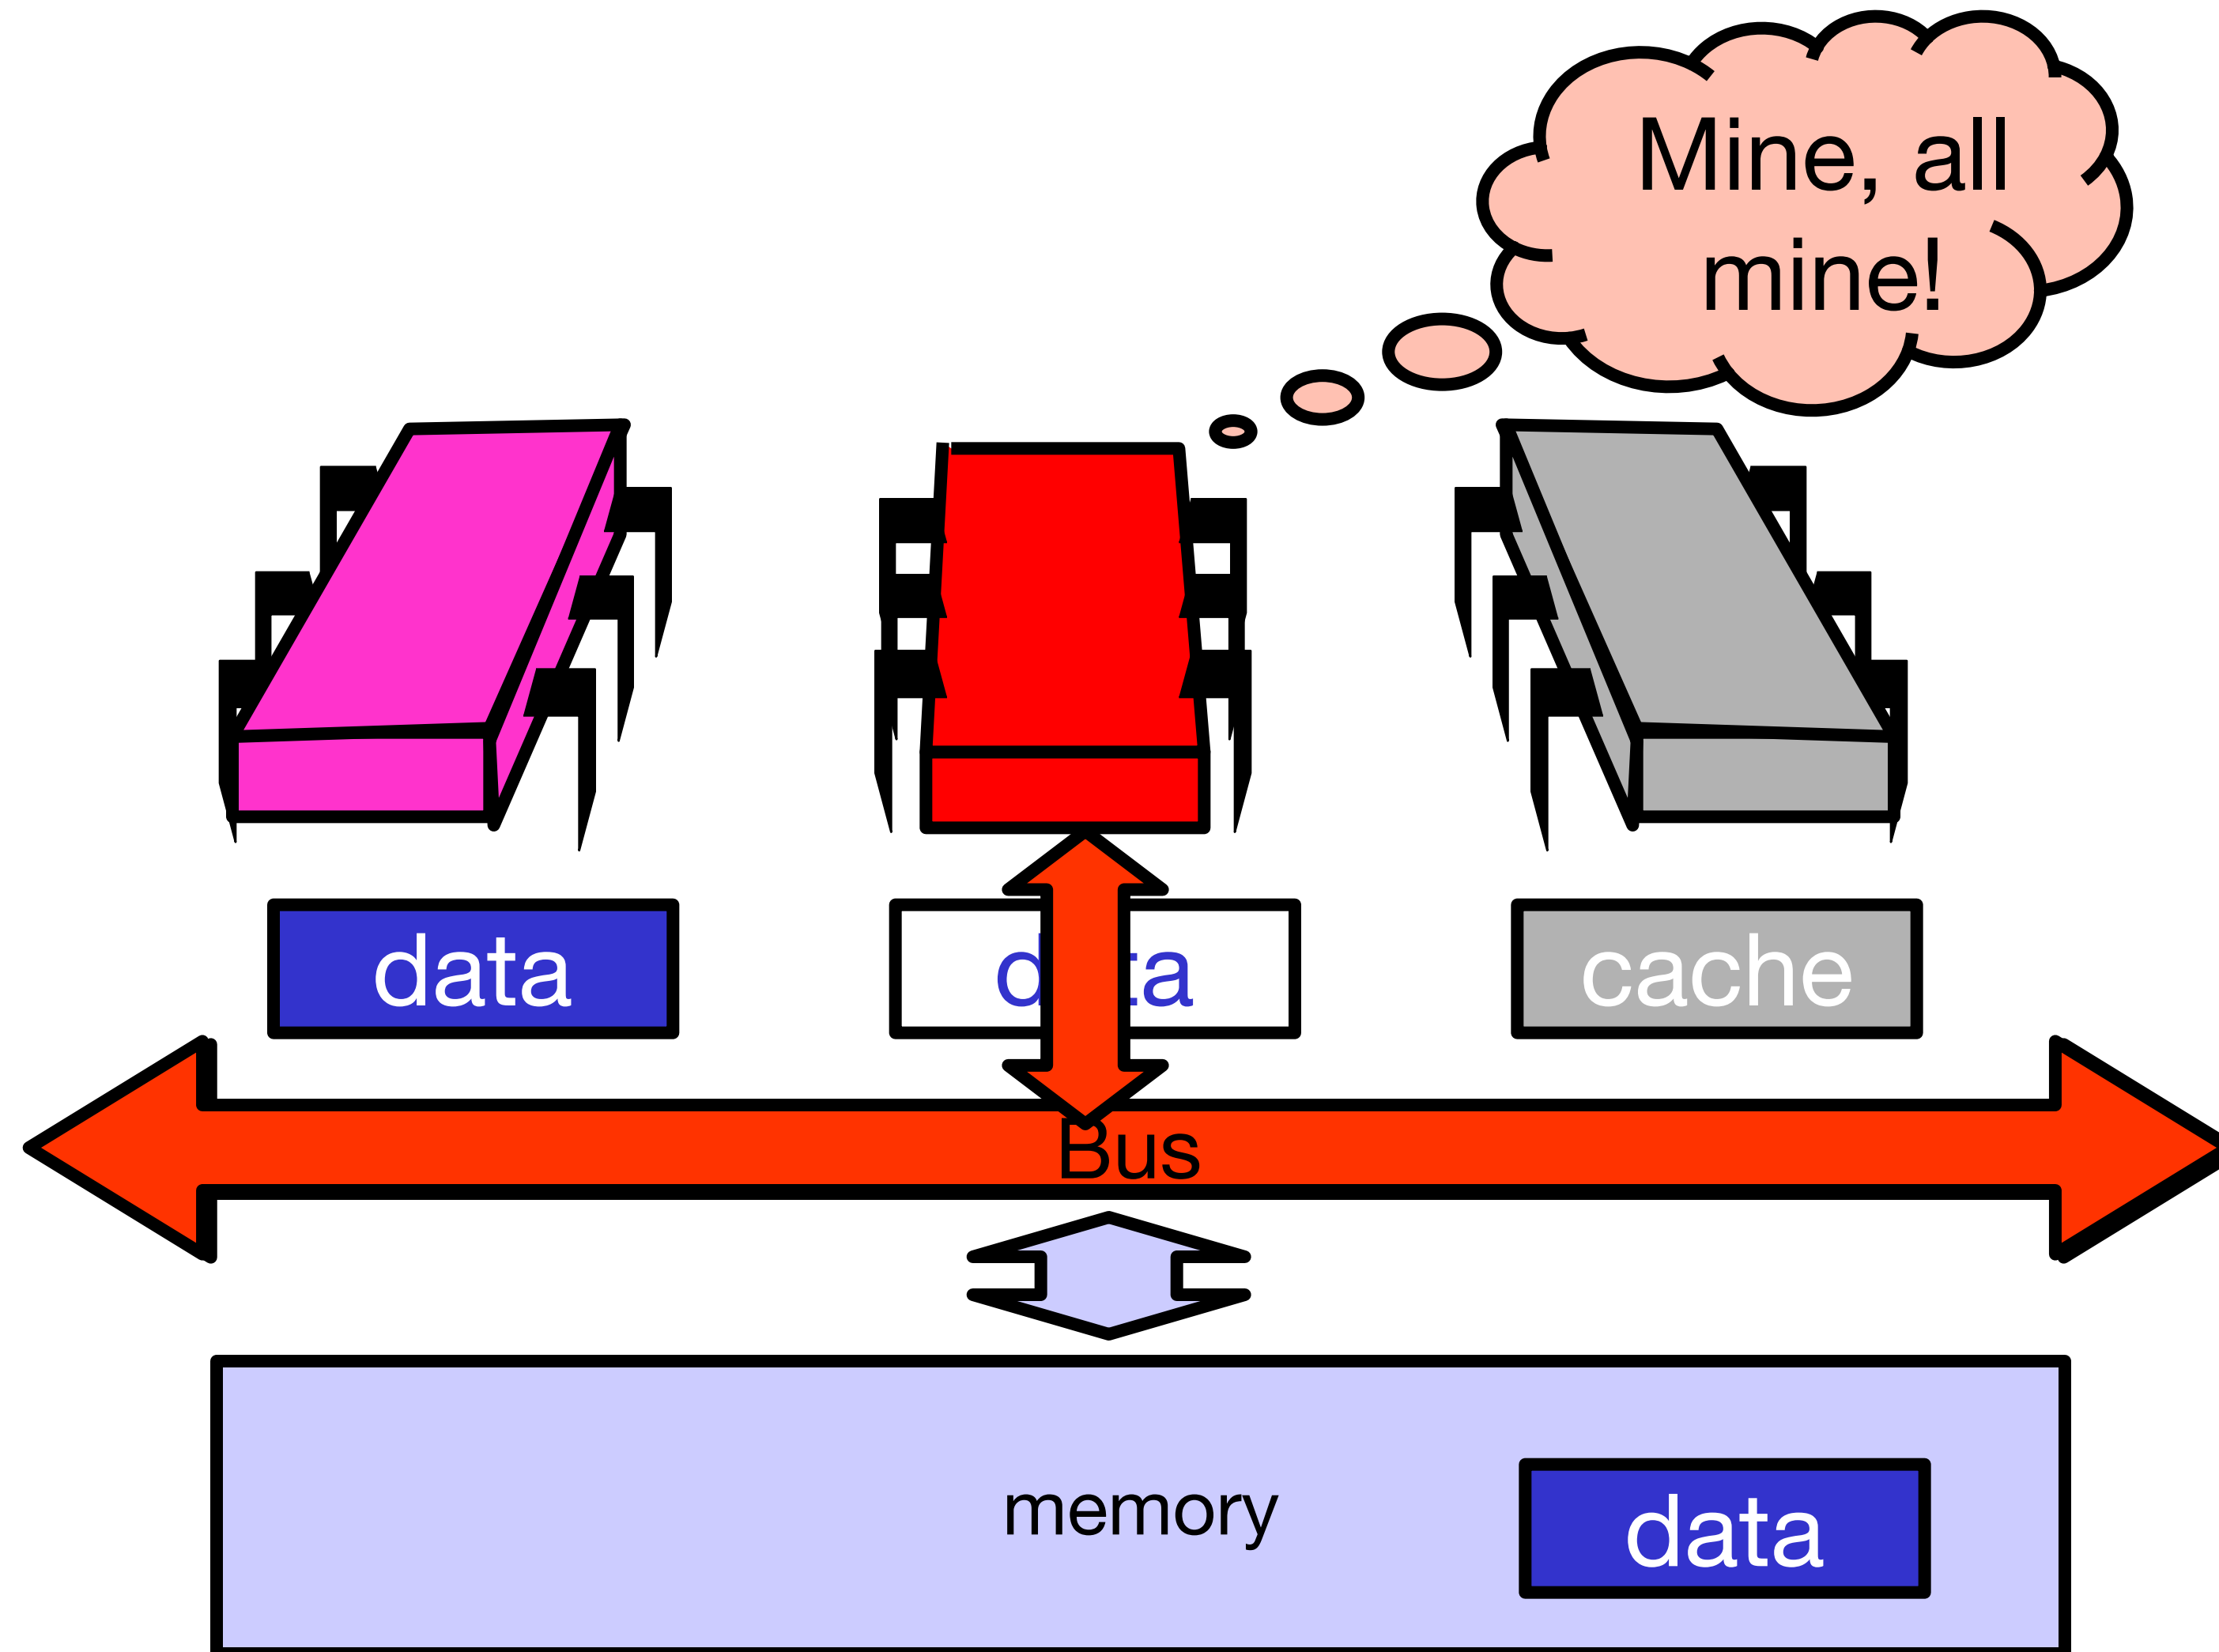
# Write-Back Caches

- Cache entry has three states
  - Invalid: contains raw seething bits
  - Valid: I can read but I can't write
  - Dirty: Data has been modified
    - Intercept other load requests
    - Write back to memory before using cache


# Invalidate



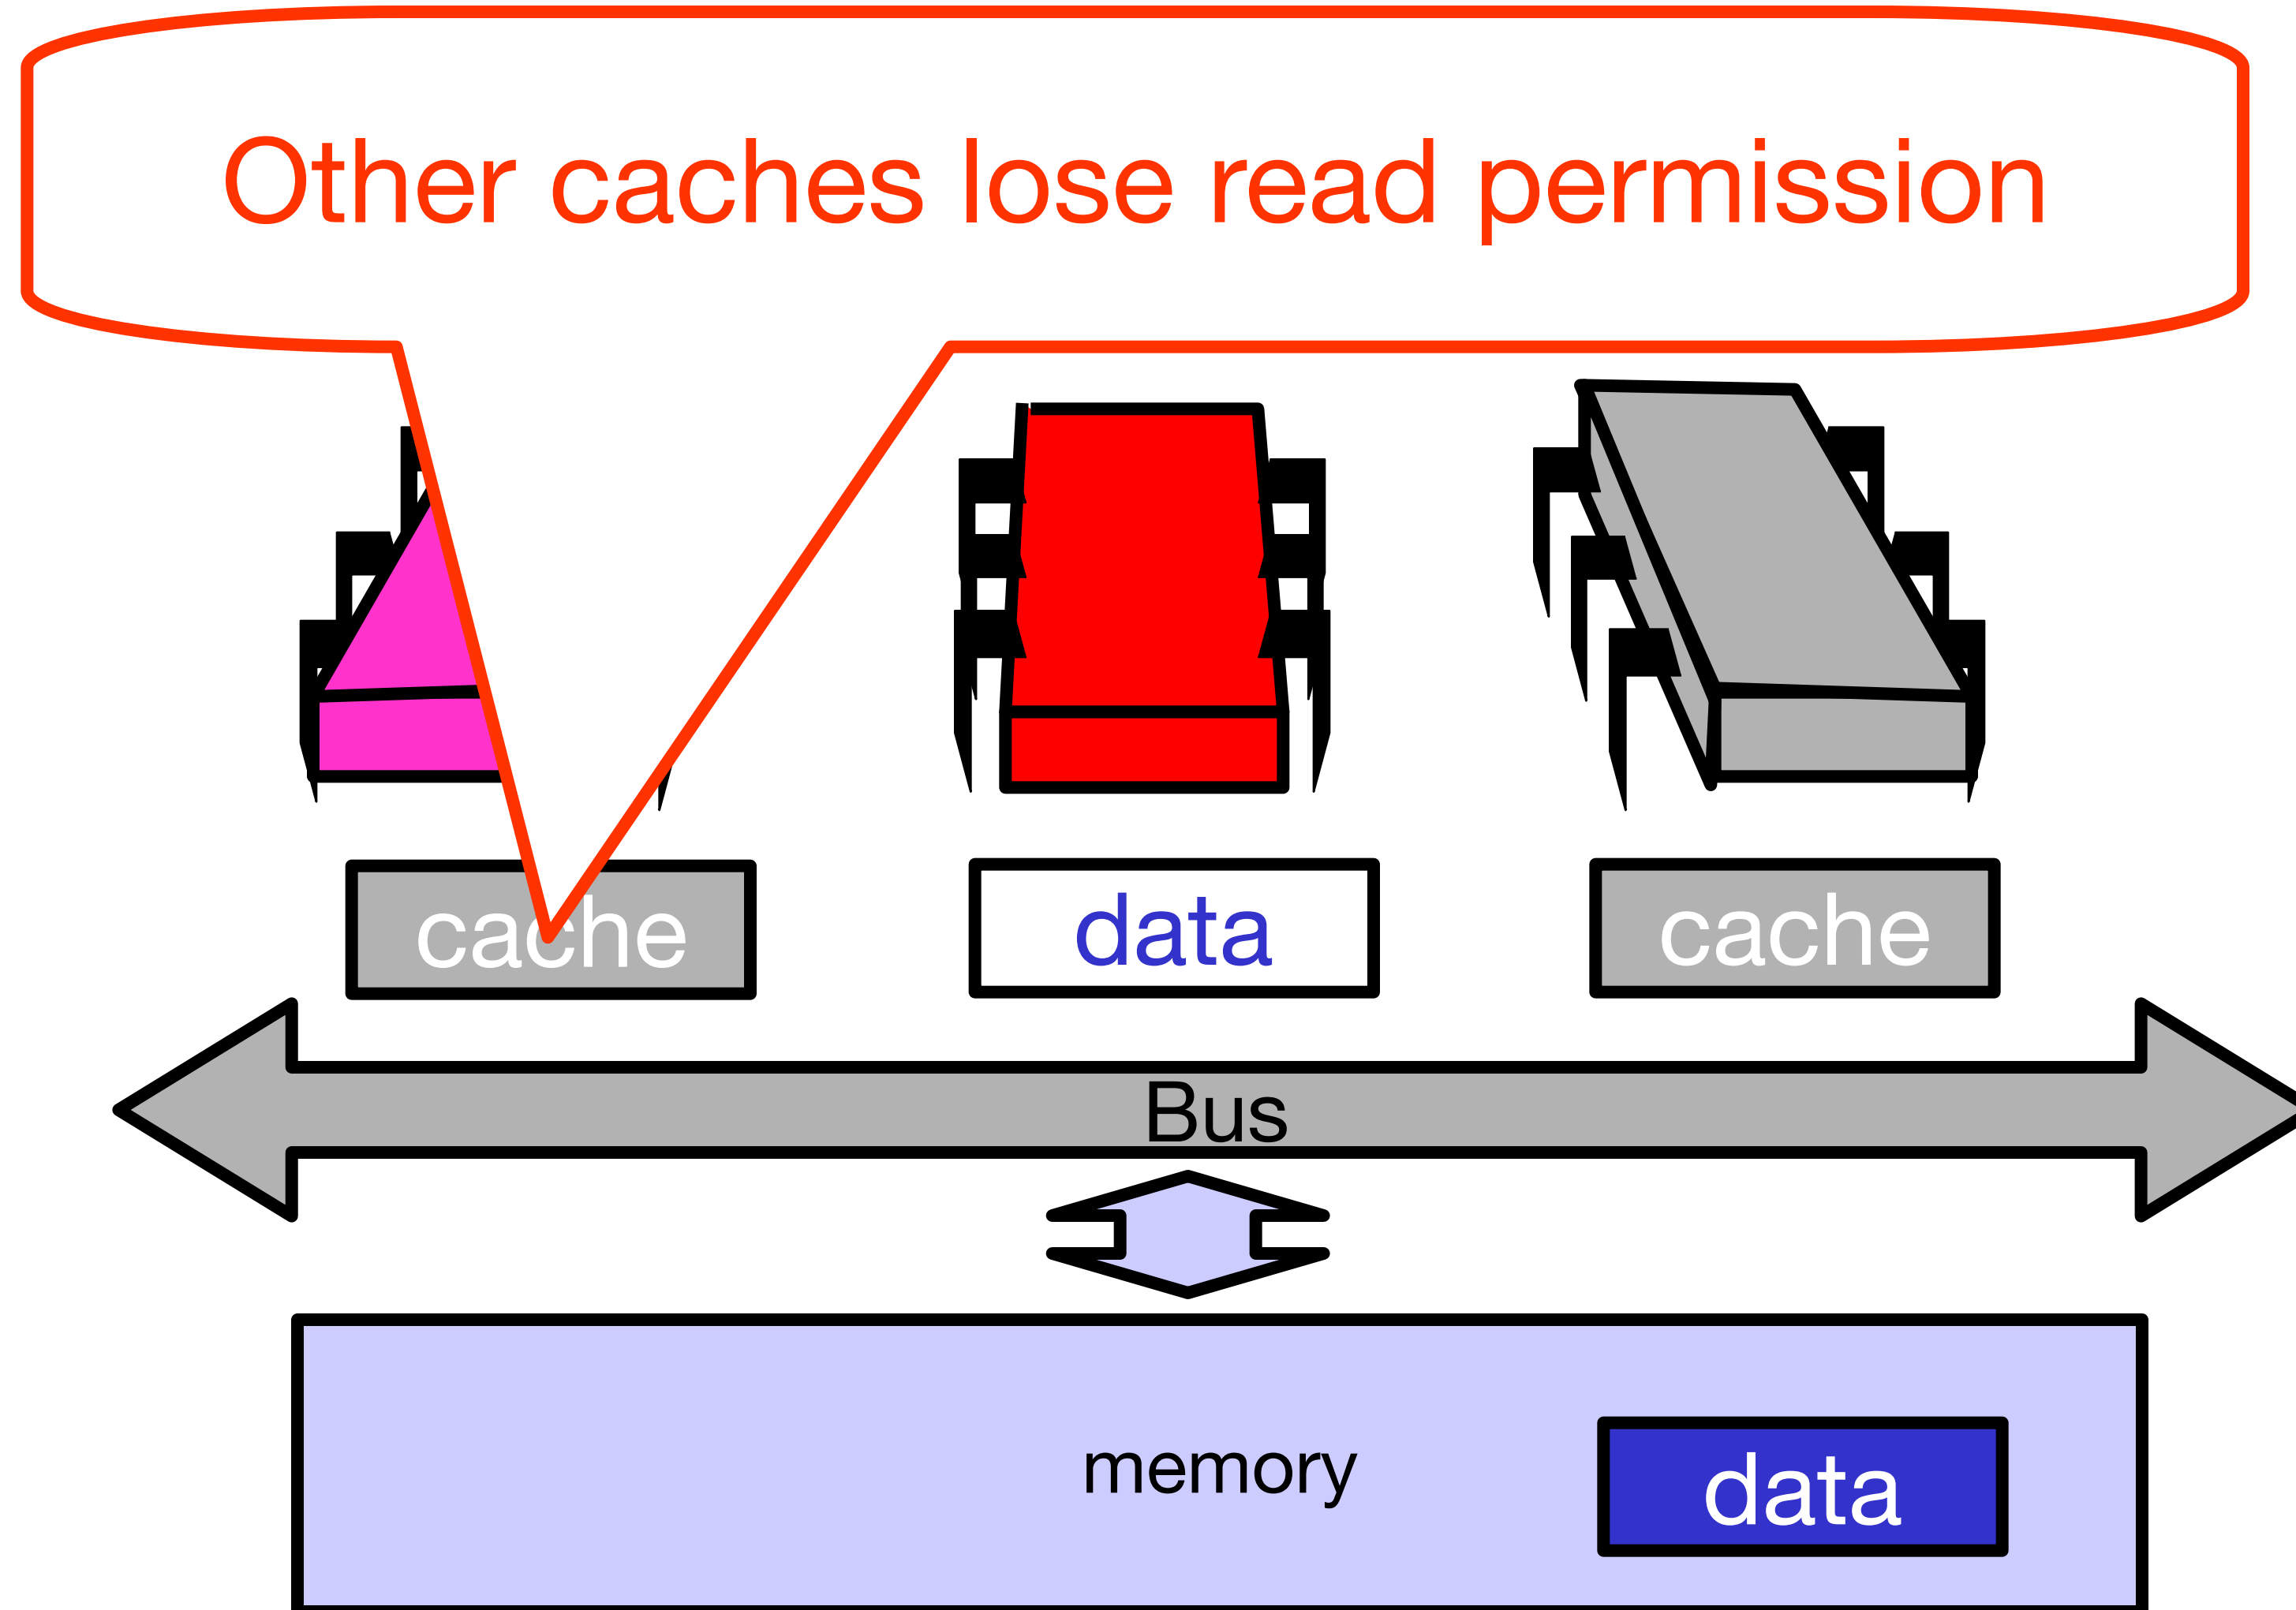
# Invalidate



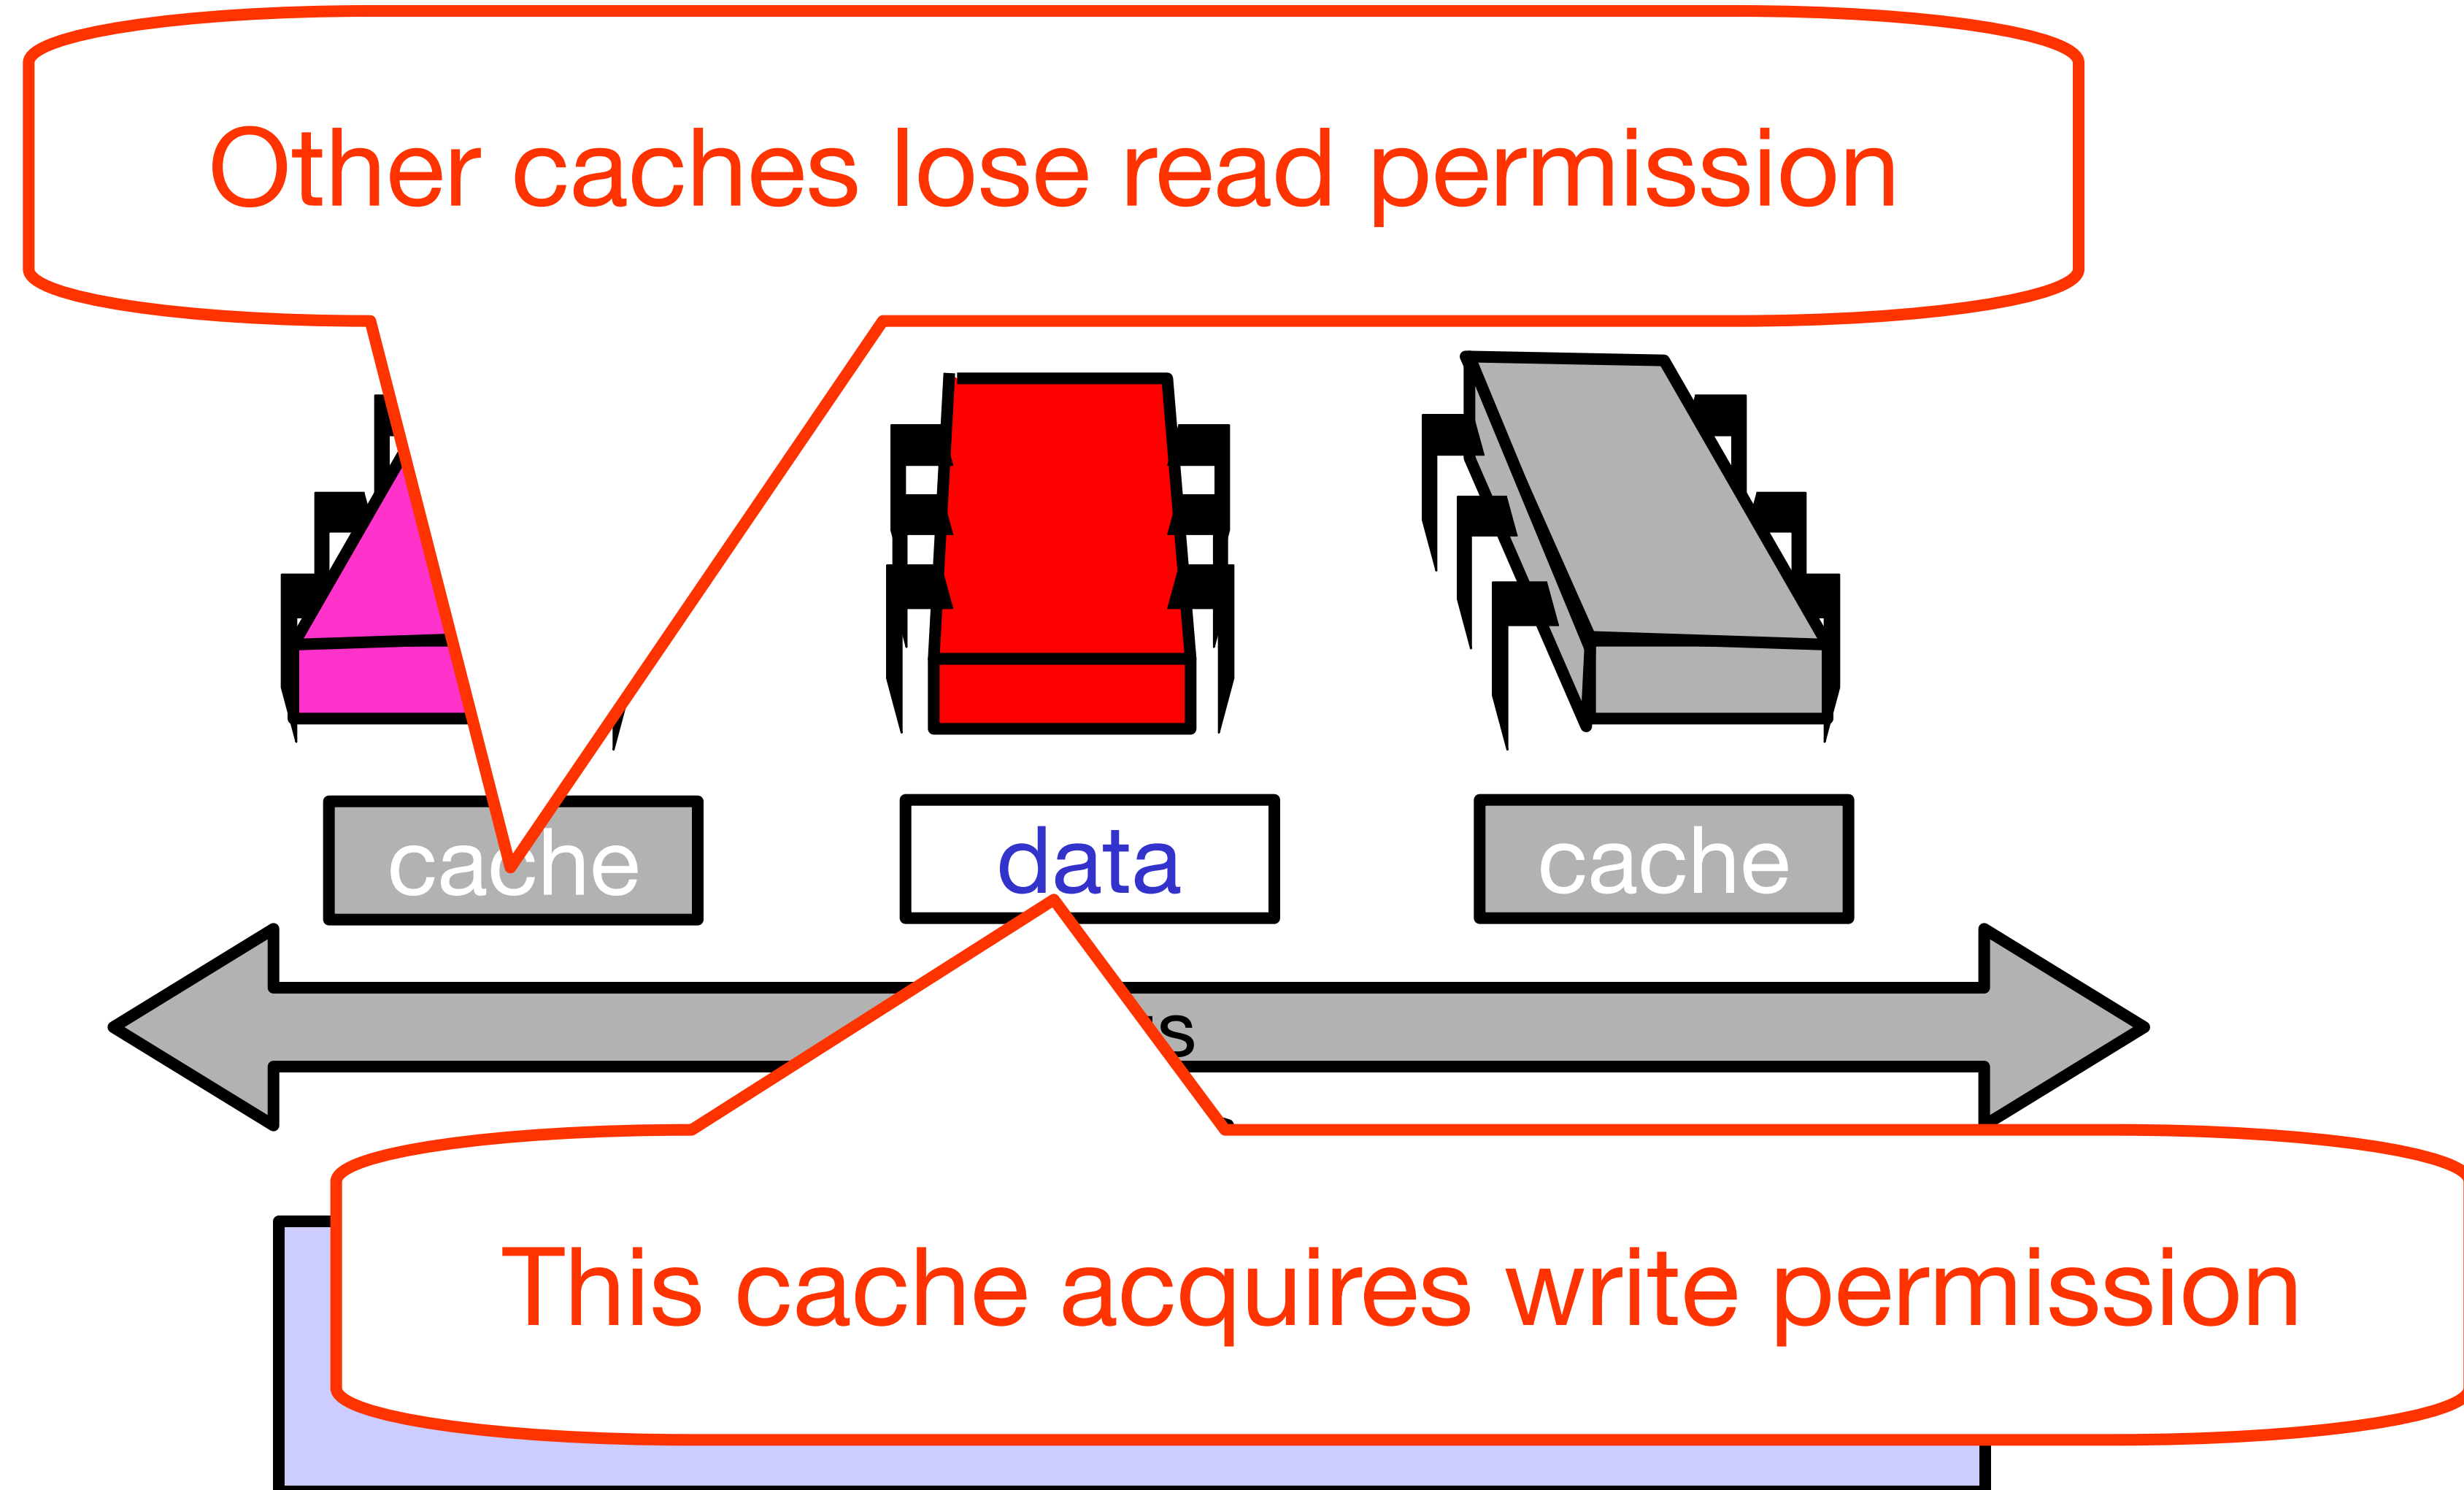
# Invalidate



# Invalidate

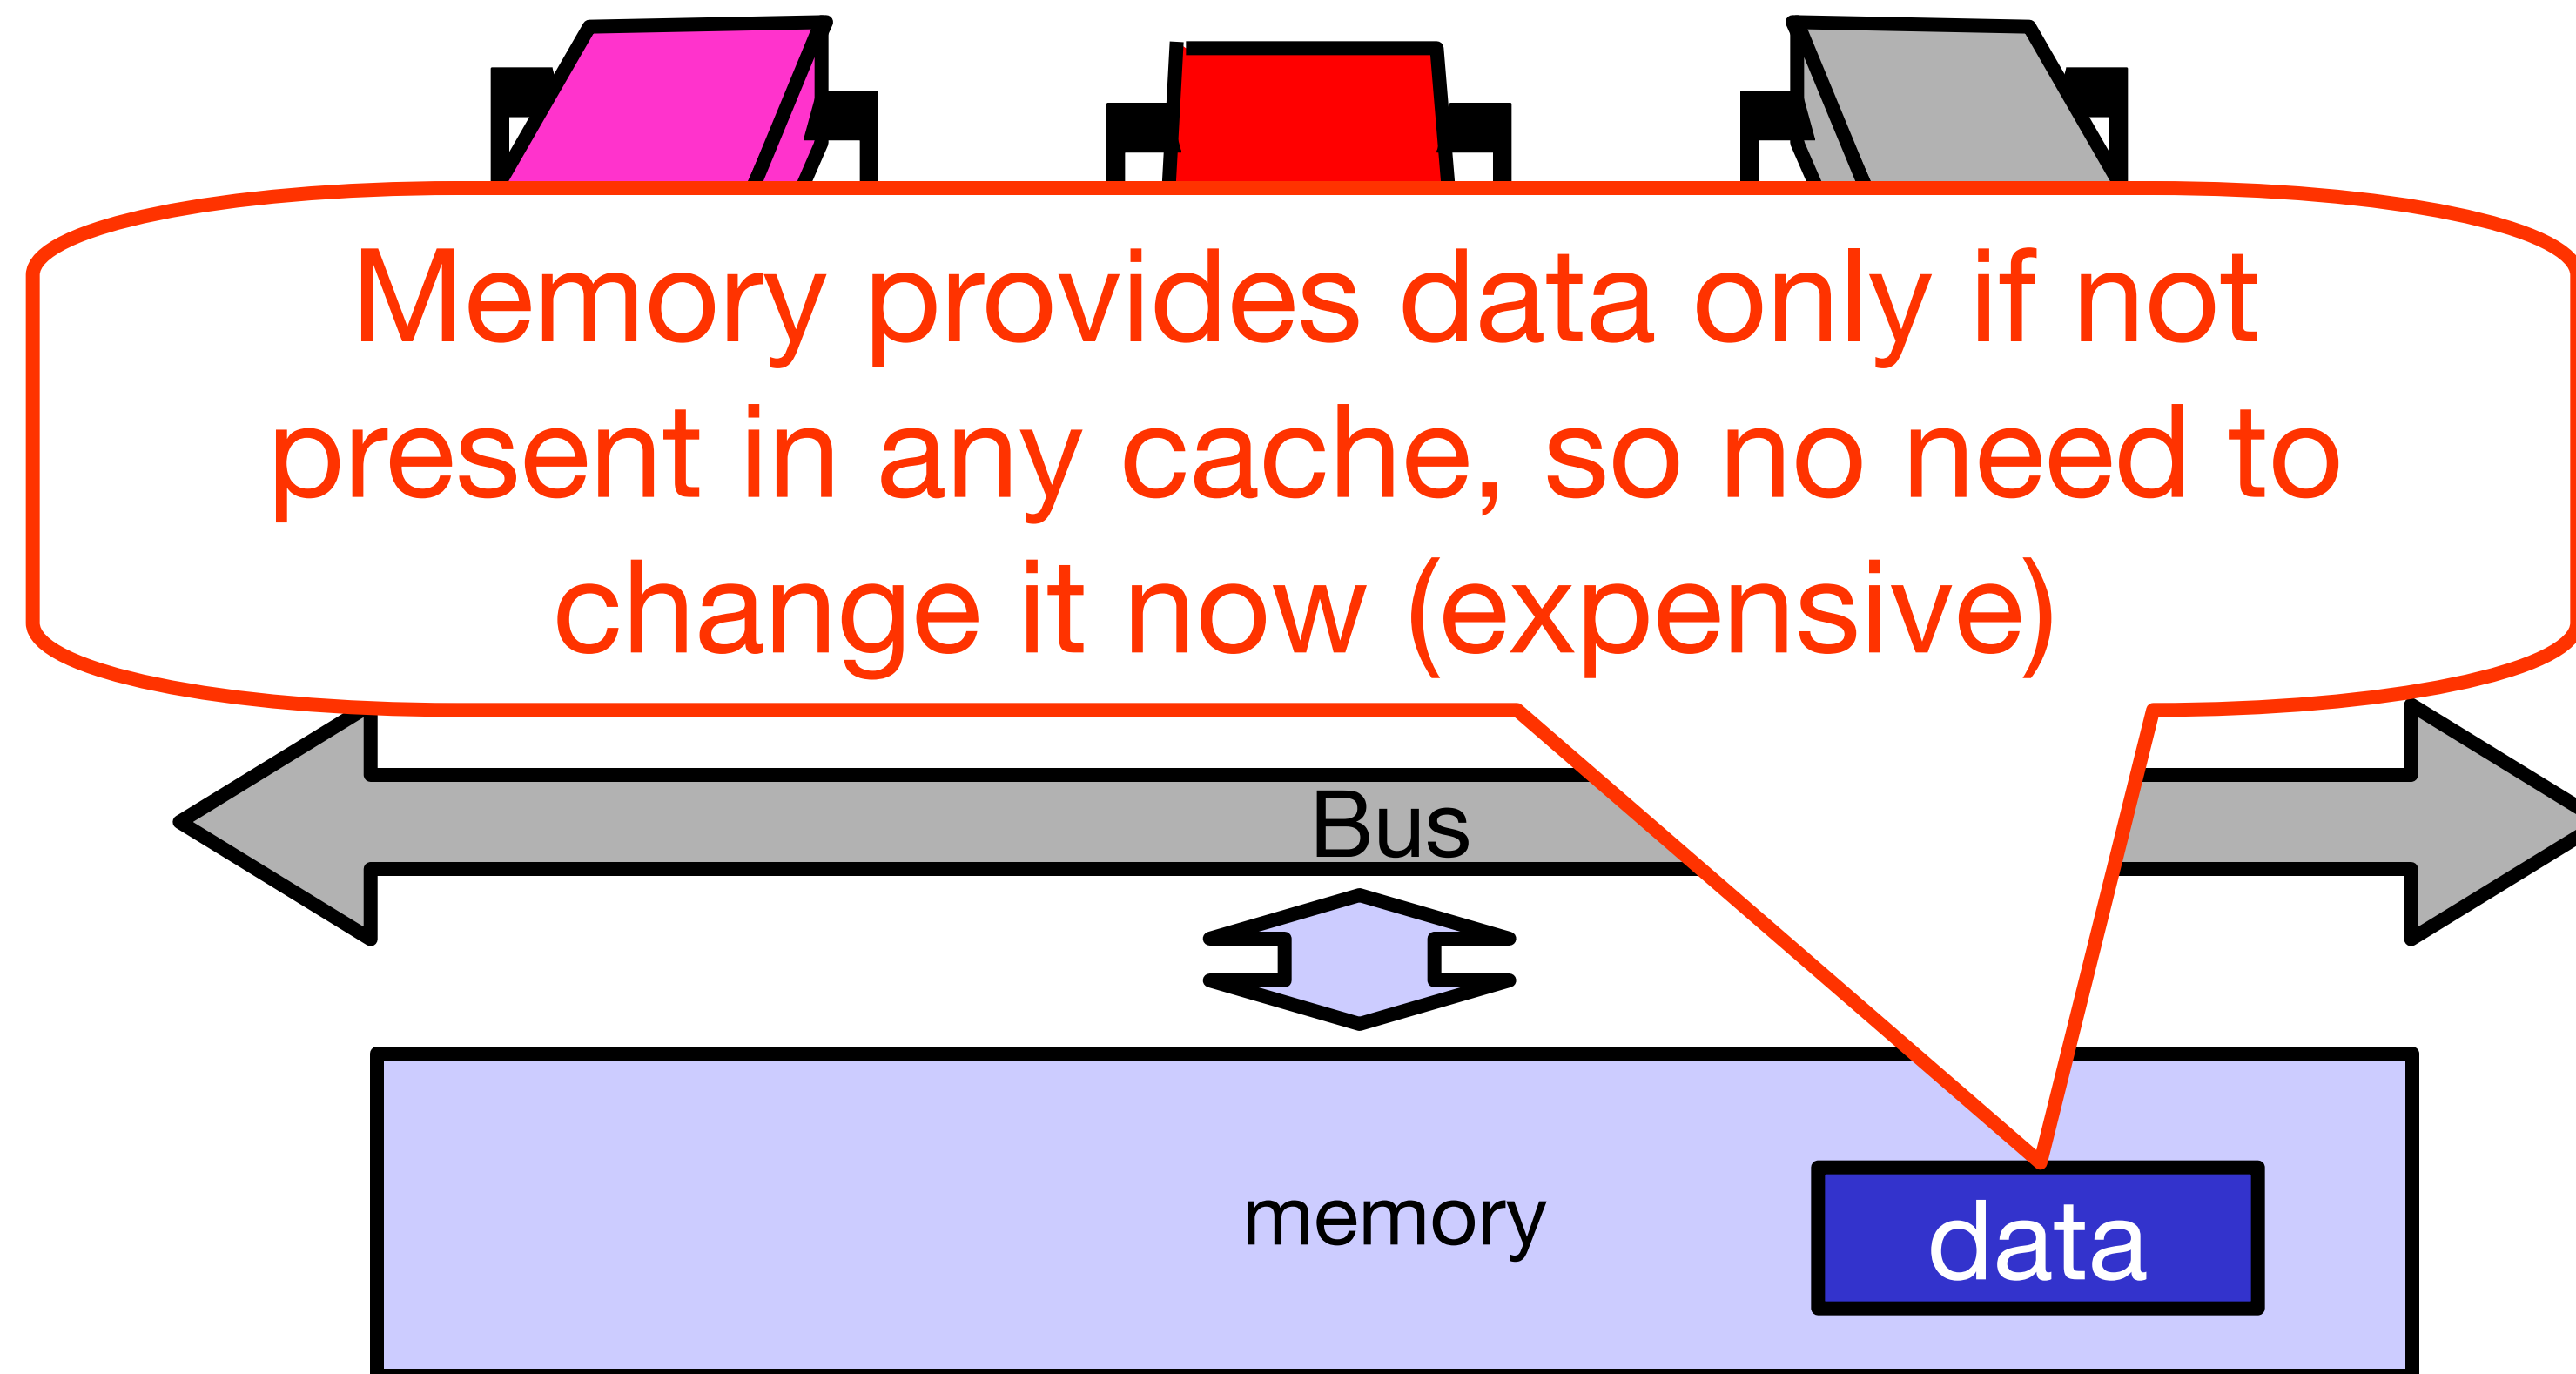


# Invalidate







# Invalidate



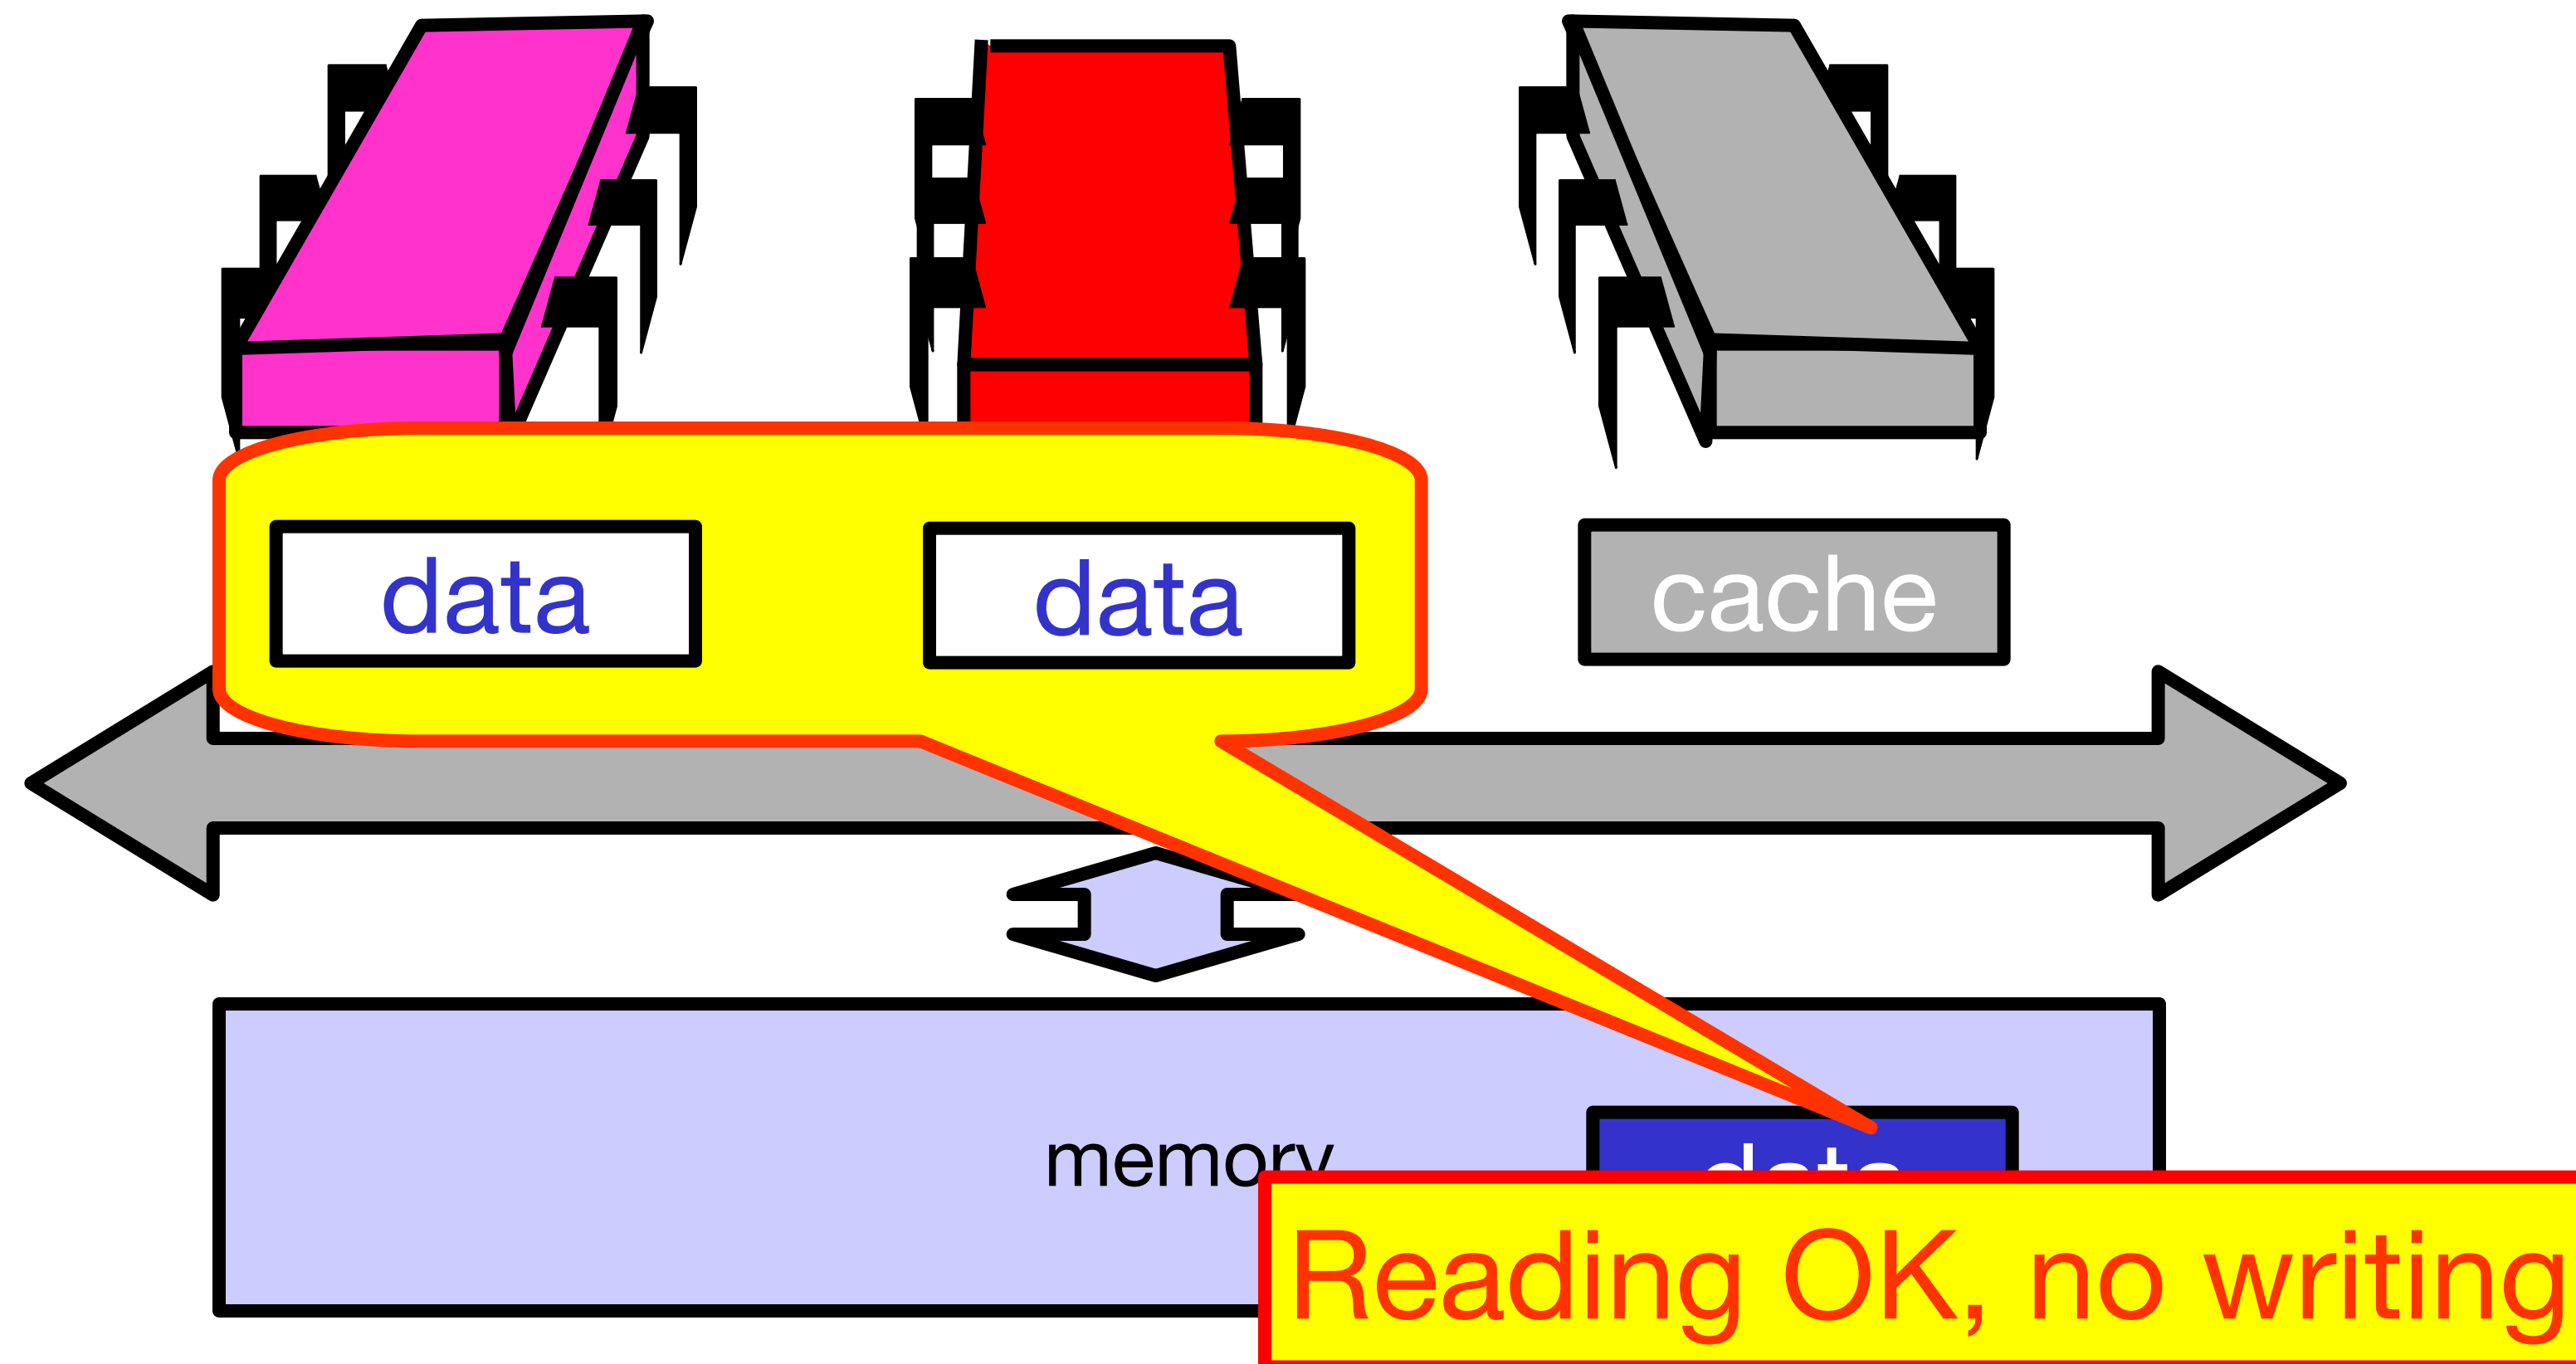
# Another Processor Asks for Data



# Owner Responds



# End of the Day ...



# Mutual Exclusion

- What do we want to optimize?
  - Bus bandwidth used by spinning threads
  - Release/Acquire latency
  - Acquire latency for idle lock


# Simple TASLock

- TAS invalidates cache lines
- Spinners
  - Miss in cache
  - Go to bus
- Thread wants to release lock
  - delayed behind spinners

# Test-and-test-and-set

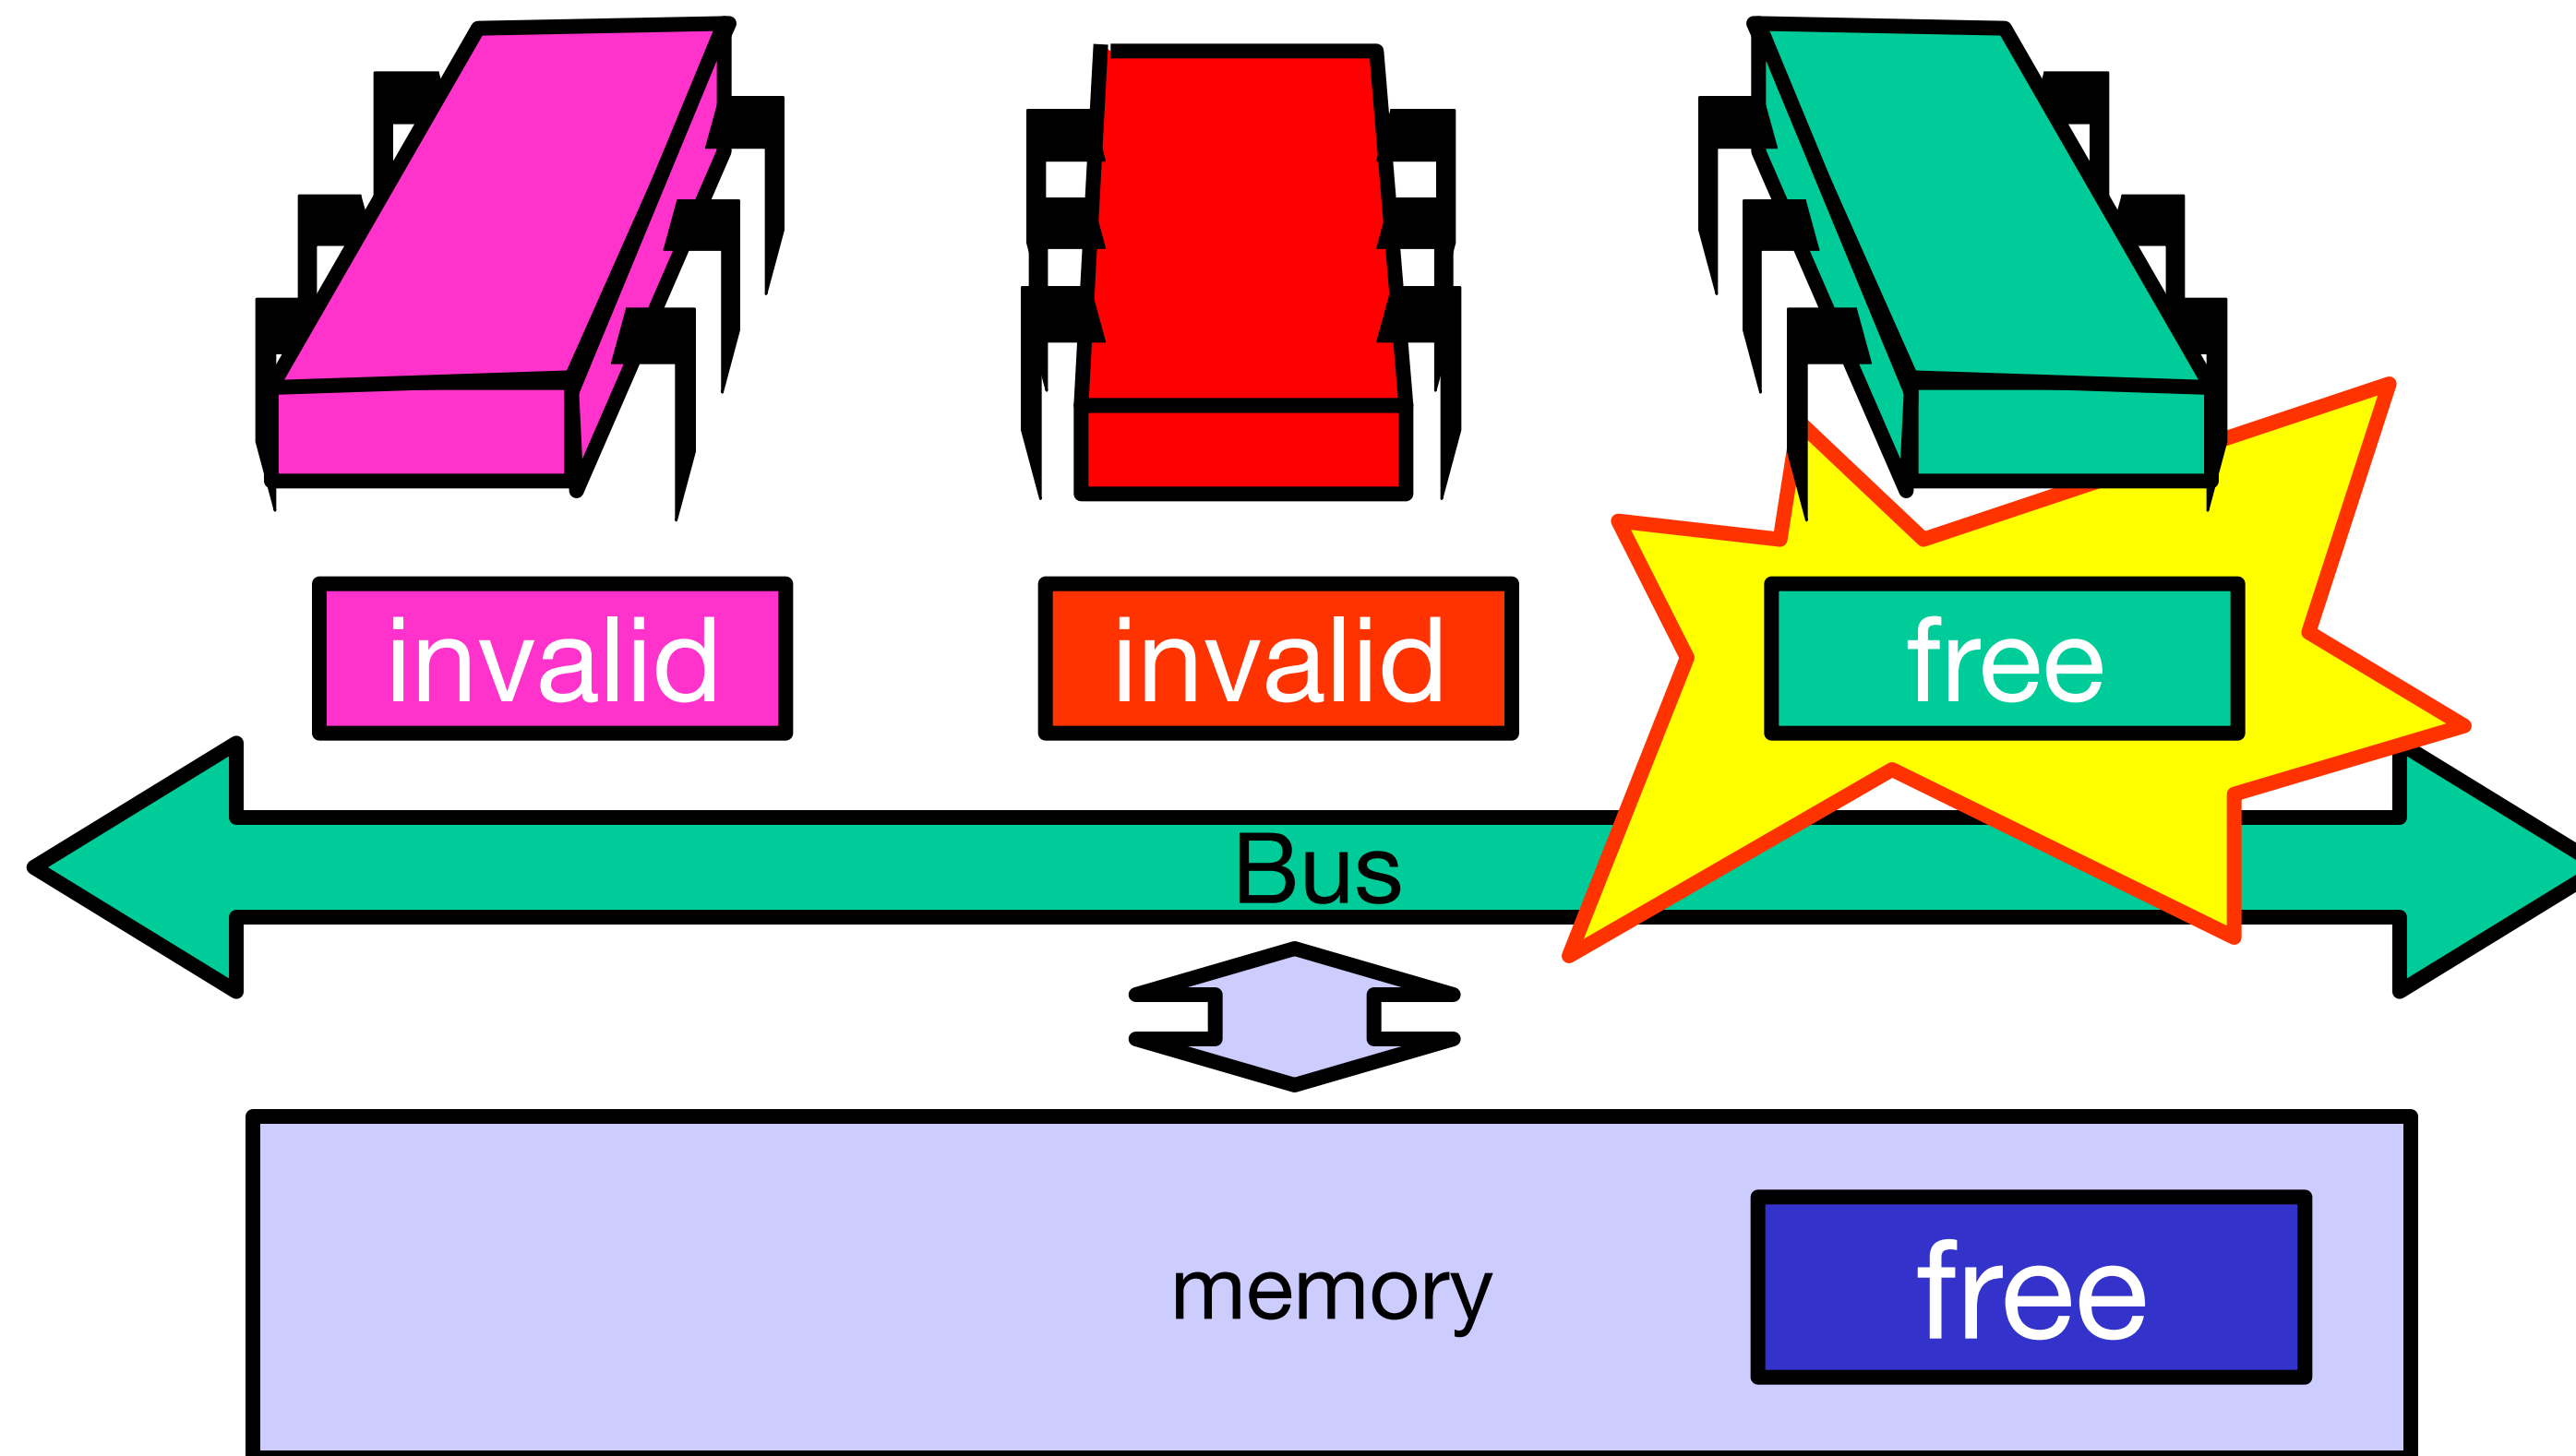
- Wait until lock “looks” free
  - Spin on local cache
  - No bus use while lock busy
- Problem: when lock is released
  - Invalidation storm ...

# Local Spinning while Lock is Busy






# On Release



# On Release


Everyone misses, rereads




(1)

# On Release

Everyone tries TAS



# Mystery Explained




# Problems

- Everyone misses
  - Reads satisfied sequentially
- Everyone does TAS
  - Invalidates others' caches
- Eventually quiesces after lock acquired
  - How long does this take?

# Solution: Introduce Delay


- If the lock looks free
- But I fail to get it
- There must be lots of contention
- Better to back off than to collide again



# Dynamic Example: Exponential Backoff

If I fail to get lock

- wait random duration before
- Each subsequent fail



# Exponential Backoff Lock

```
public class Backoff implements lock {
    public void lock() {
        int delay = MIN_DELAY;
        while (true) {
            while (state.get()) {}
            if (!lock.getAndSet(true))
                return;
            sleep(random() % delay);
            if (delay < MAX_DELAY)
                delay = 2 * delay;
        }
    }
}
```



# Exponential Backoff Lock

```
public class Backoff implements lock {  
    public void lock() {  
        int delay = MIN_DELAY;  
        while (true) {  
            while (state.get()) {}  
            if (!lock.getAndSet(true))  
                return;  
            sleep(random() % delay);  
            if (delay < MAX_DELAY)  
                delay = 2 * delay;  
        }  
    }  
}
```

**Fix minimum delay**

# Exponential Backoff Lock

```
public class Backoff implements lock {
    public void lock() {
        int delay = MIN_DELAY;
        while (true) {
            while (state.get()) {}
            if (!lock.getAndSet(true))
                return;
            sleep(random() % delay);
            if (delay < MAX_DELAY)
                delay = 2 * delay;
        }
    }
}
```

**Wait until lock looks free**

# Exponential Backoff Lock

```
public class Backoff implements lock {
    public void lock() {
        int delay = MIN_DELAY;
        while (true) {
            while (state.get()) {}
            if (!lock.getAndSet(true))
            return;
            sleep(random() % delay);
            if (delay < MAX_DELAY)
                delay = 2 * delay;
        }
    }
}
```

**If we win, return**

# Exponential Backoff Lock

```
public class Backoff implements lock {
    public void lock() {
        int delay = MIN_DELAY;
        while (true) {
            while (state.get()) {}
            if (!lock.getAndSet(true))
                return;
            sleep(random() % delay);
            if (delay < MAX_DELAY)
                delay = 2 * delay;
        }
    }
}
```

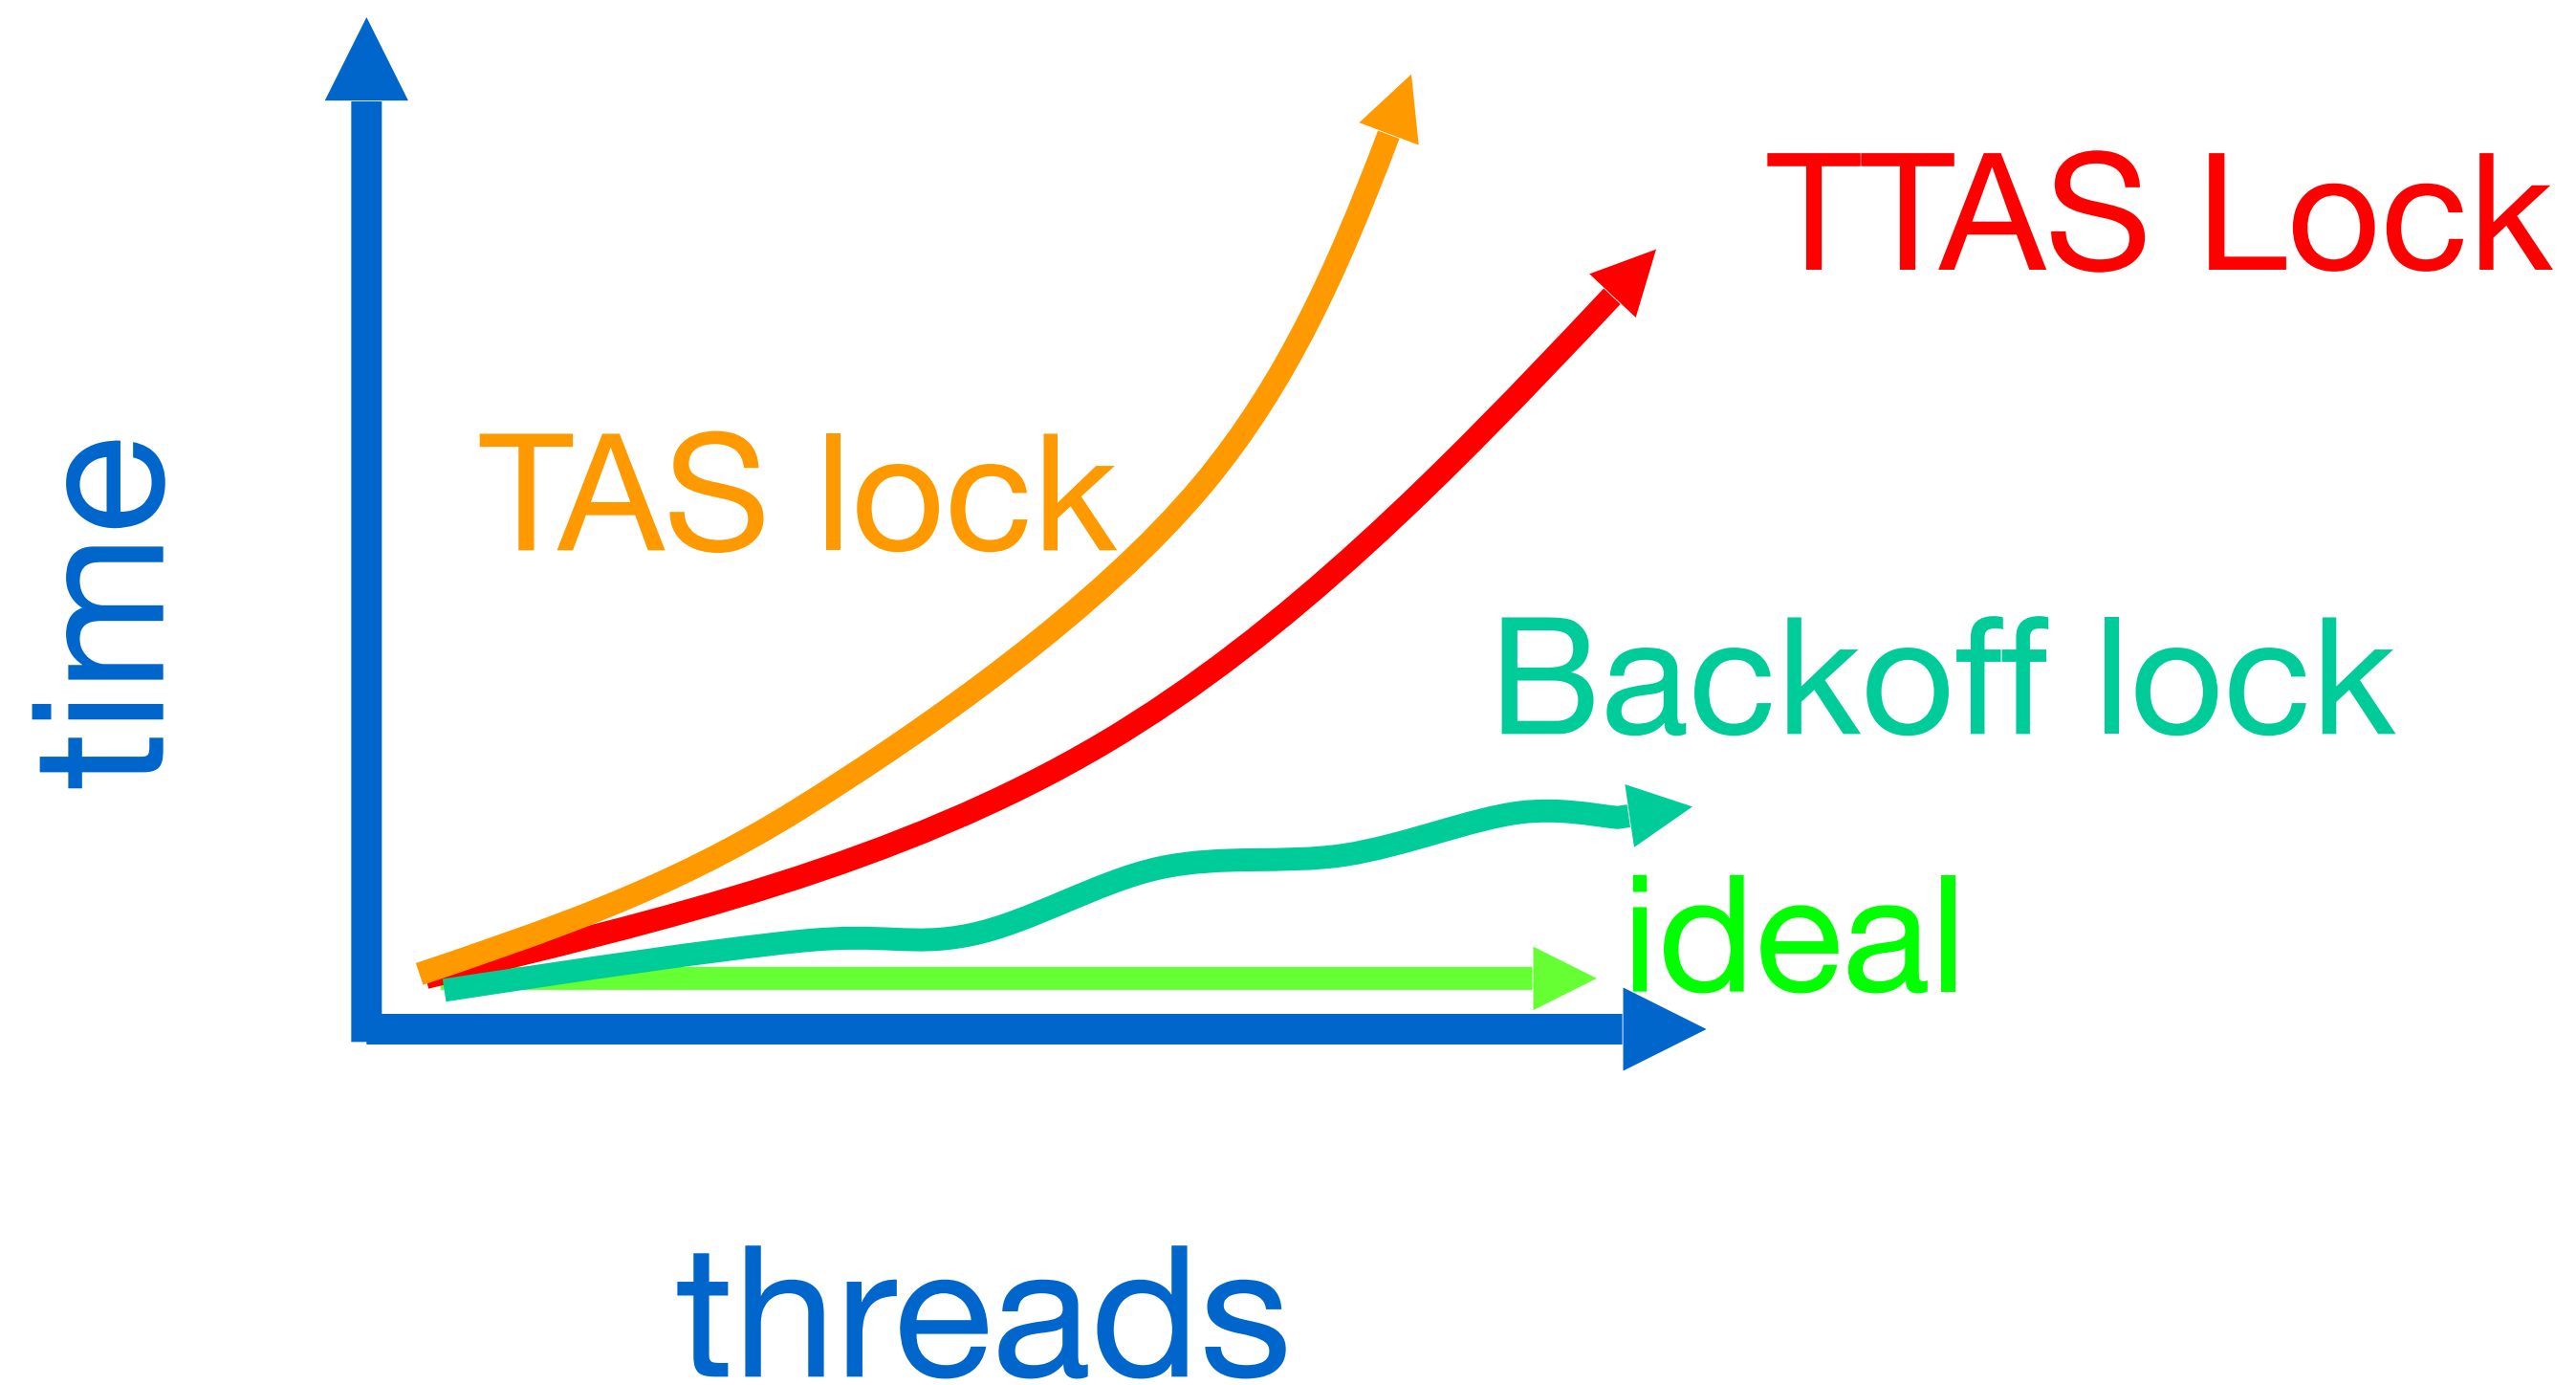
**Back off for random duration**

# Exponential Backoff Lock

```
public class Backoff implements Lock {  
    public void lock() {  
        int delay = MIN_DELAY;  
        while (true) {  
            while (state.get()) {}  
            if (!lock.getAndSet(true))  
                return;  
            sleep(random() % delay);  
            if (delay < MAX_DELAY)  
                delay = 2 * delay;  
        }  
    }  
}
```

**Double max delay, within reason**

# Spin-Waiting Overhead



# Backoff: Other Issues

- Good
  - Easy to implement
  - Beats TTAS lock
- Bad
  - Must choose parameters carefully
  - Not portable across platforms

# Moral of the story?

- EVEN IF we do pretty good with parallelizing most parts our application, we can still see slowdown from **contention** for locks
- For resources infrequently contended, spin locks can be fast because no context switch is necessary
- But, contention in spin locks can have repercussions due to hardware architectures



# This work is licensed under a Creative Commons Attribution-ShareAlike license

- This work is licensed under the Creative Commons Attribution-ShareAlike 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-sa/4.0/>
- You are free to:
  - Share — copy and redistribute the material in any medium or format
  - Adapt — remix, transform, and build upon the material
  - for any purpose, even commercially.
- Under the following terms:
  - Attribution — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.
  - ShareAlike — If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.
  - No additional restrictions — You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.