


# CSE113: Parallel Programming

Jan. 21, 2022

- **Topics:**

- Thread sanitizer
- Atomic operation properties
- Mutex implementations



# Announcements

- Homework 1 is due today!
  - Due at midnight
  - Teaching support will stop around 5
  - Teaching support will not be available during the weekend
- Homework 2 will be assigned today
  - Assigned by midnight (I get the same deadlines as you 😊 )
  - Part 1 is possible after today
  - Part 2 and 3 will be possible after Wednesday

# Announcements

- Homework 1 notes:
  - No assigned speedup required. You should get a noticeable speedup from ILP
    - AMD processors are being a little strange on part 2
    - Please note the processor in your write-up
  - You can start to share results now. Everyone's results will be slightly different
  - Sometimes you cannot account for small differences
    - We should be running the code for more iterations
- Turn-in
  - two files:
    - 1 PDF
    - 1 zip

# Announcements

- Last week remote:
  - I will aim to continue recording lectures
  - But the synchronous lecture **will not be** hybrid
  - As noted in the original outline: in-person attendance is part of your grade
  - But do not come if you are sick (we will make accommodations)

# Today's Quiz

- Please do it!
- Due before class on Monday

# Previous quiz

It is possible to interleave the load and store operations of RMW atomic operations; however, it is so rare that it does not matter in practice.

Modify these programs to use atomic RMWs

Tyler's coffee addiction:

```
atomic_fetch_add(&tylers_account, -1);
```

Tyler's employer

```
atomic_fetch_add(&tylers_account, 1);
```

time



time



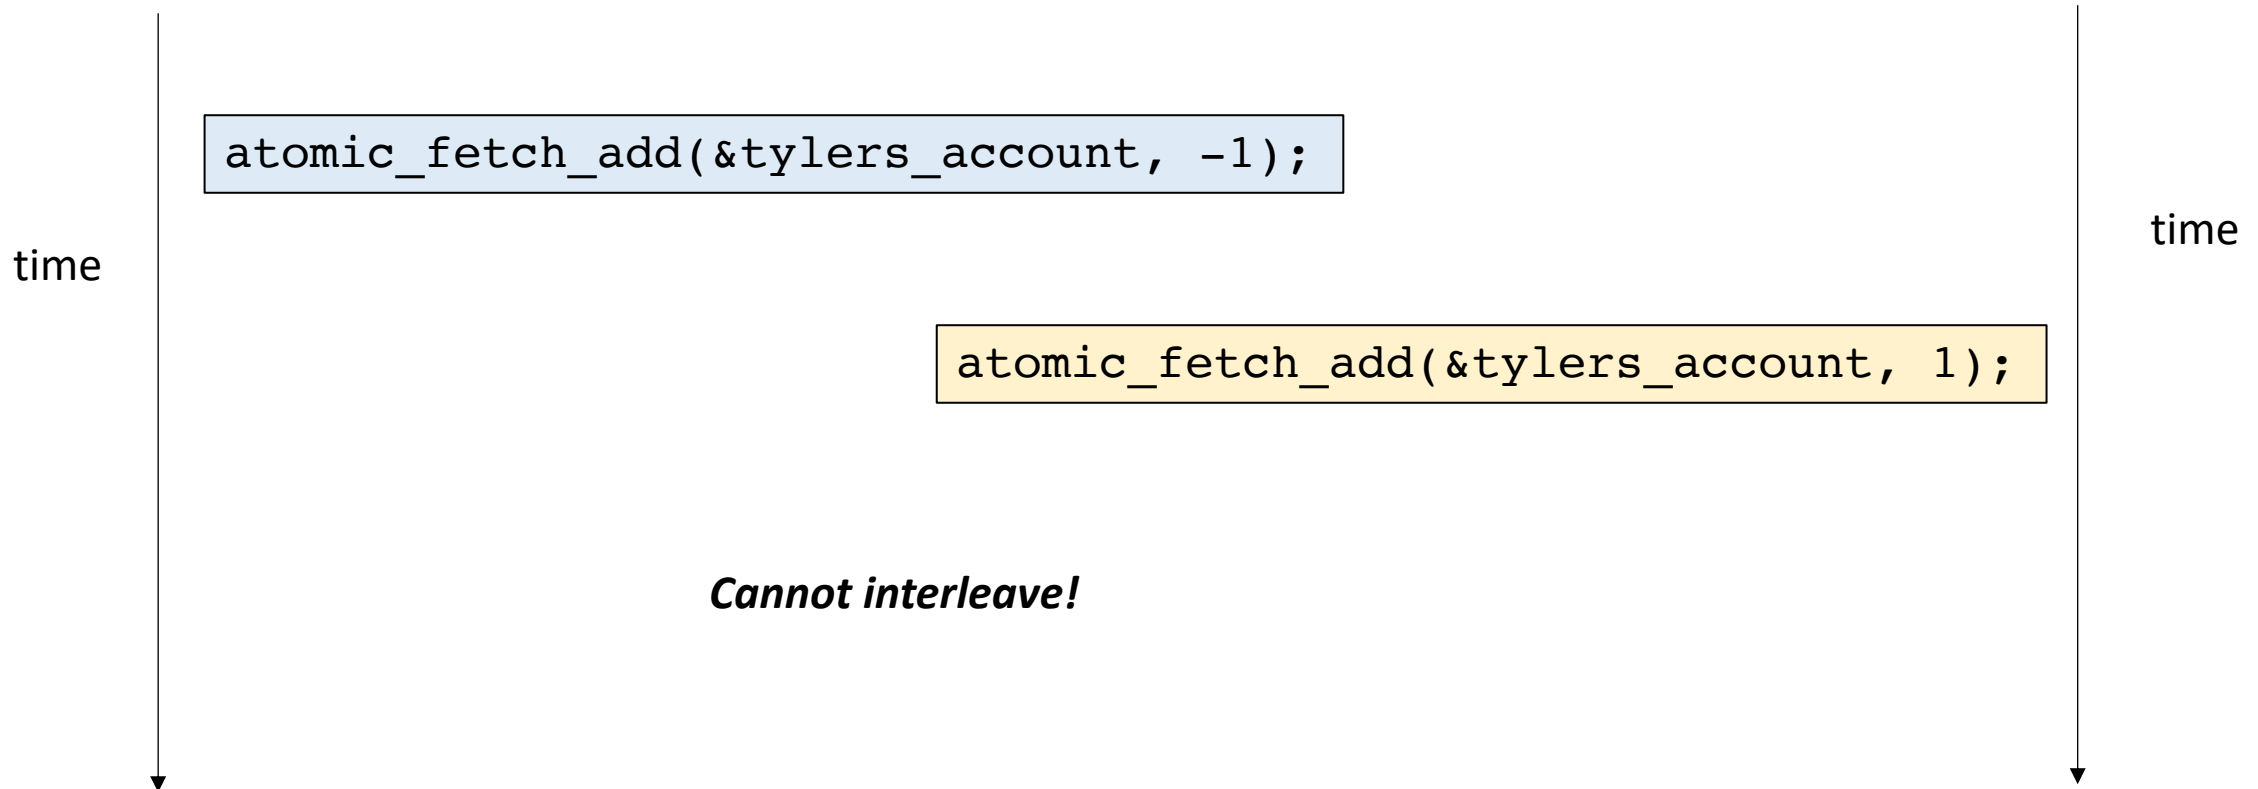
# Modify these programs to use atomic RMWs

Tyler's coffee addiction:

```
atomic_fetch_add(&tylers_account, -1);
```

Tyler's employer

```
atomic_fetch_add(&tylers_account, 1);
```





# Previous quiz


We should aim to make mutual exclusion regions as short as possible because of the caching overhead of locks.

# Mutex Performance

Try to keep mutual exclusion sections small! Protect only data conflicts!

Code example with overhead

*Long periods of waiting in the threads*

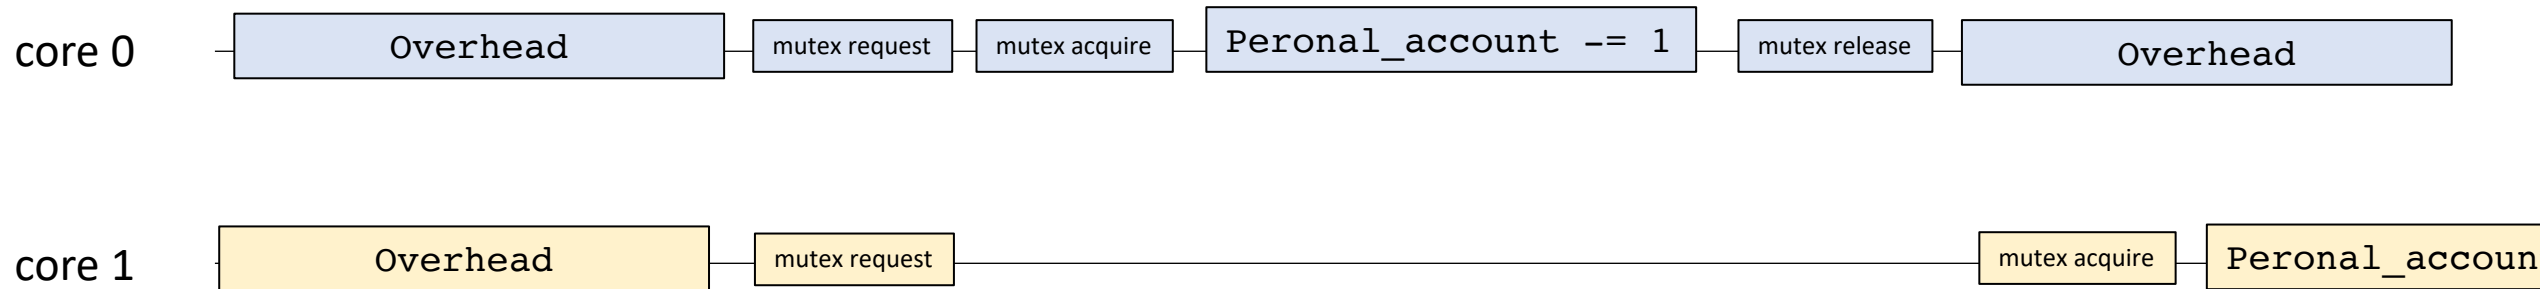


*Long periods of waiting in the threads*

# Mutex Performance

Try to keep mutual exclusion sections small! Protect only data conflicts!

Code example with overhead



*overlap the overhead (i.e. computation without any data conflicts)*

# Previous quiz


In a few sentences explain the trade-offs between using a single mutex to protect all data vs. using multiple mutexes to protect different data locations.

# Multiple mutexes

Mutexes are objects. We can create multiple versions of them to protect different shared data.

MutexP for personal account  
MutexB for business account

Critical sections across different mutexes can overlap

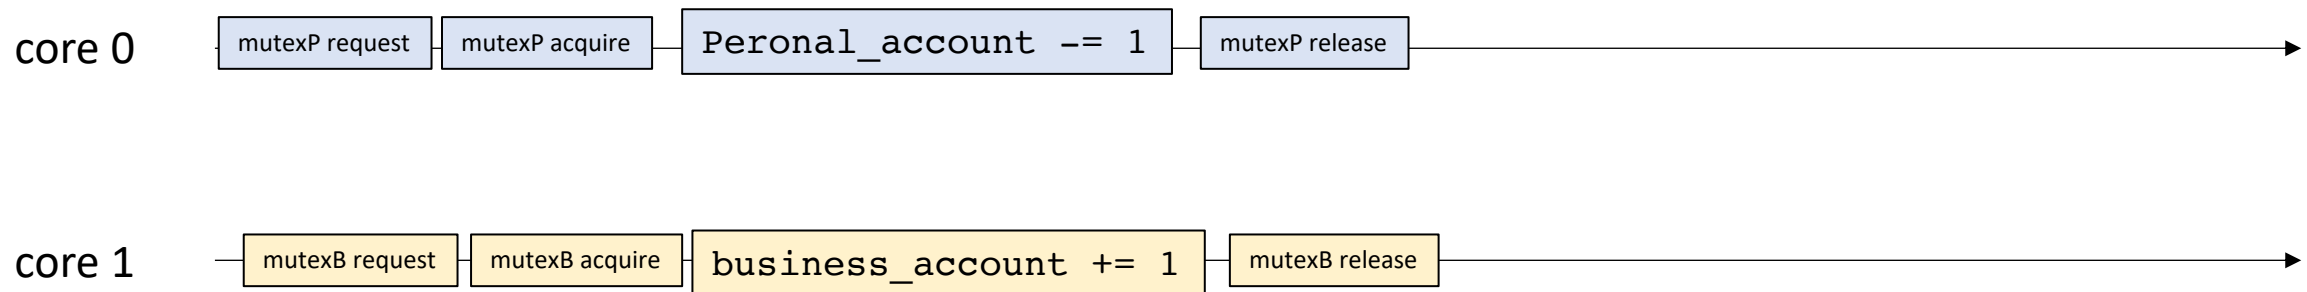


# Multiple mutexes

Mutexes are objects. We can create multiple versions of them to protect different shared data.

MutexP for personal account  
MutexB for business account

Critical sections across different mutexes can overlap



# Programming with mutexes can be HARD!

make sure all data conflicts are protected with a mutex

keep critical sections small

balance between having many mutexes (provides performance) but gives the potential for deadlocks

# There is help: *Thread Sanitizer*

Add compile line:

- `-fsanitize=thread`

What it does

- Checks for data-races
- Checks for deadlocks (*currently it is not working for me*)
- *At an overhead*



# Previous quiz

Which one of the answers is NOT a property of mutexes?

**Deadlock Freedom**

Mutual Exclusion

Deterministic Execution

Starvation Freedom

# Properties of mutexes

Recap: three properties

- **Mutual Exclusion:** Two threads cannot be in the critical section at the same time
- **Deadlock Freedom:** If a thread has requested the mutex, and no thread currently holds the mutex, the mutex must be acquired by one of the requesting threads
- **Starvation Freedom** (*optional*): A thread that requests the mutex must eventually obtain the mutex.

# Properties of mutexes

## Three properties

- **Deadlock Freedom** - If a thread has requested the mutex, and no thread currently holds the mutex, the mutex must be acquired by one of the requesting threads

concurrent execution



time

# Properties of mutexes

## Three properties

- **Deadlock Freedom** - If a thread has requested the mutex, and no thread currently holds the mutex, the mutex must be acquired by one of the requesting threads

Program cannot hang here  
Either thread 0 or thread 1 must acquire the mutex

concurrent execution



time


# Properties of mutexes

## Three properties

- **Deadlock Freedom** - If a thread has requested the mutex, and no thread currently holds the mutex, the mutex must be acquired by one of the requesting threads

Program cannot hang here  
Either thread 0 or thread 1 must acquire the mutex

concurrent execution




# Properties of mutexes

## Three properties

- **Deadlock Freedom** - If a thread has requested the mutex, and no thread currently holds the mutex, the mutex must be acquired by one of the requesting threads

Program cannot hang here  
Either thread 0 or thread 1 must acquire the mutex

concurrent execution



**also allowed**

time


# Properties of mutexes

## Three properties

- **Starvation Freedom** (*Optional*) - A thread that requests the mutex must eventually obtain the mutex.

*Thread 1 (yellow) requests the mutex but never gets it*

concurrent execution



time


# Properties of mutexes

## Three properties

- **Starvation Freedom** (*Optional*) - A thread that requests the mutex must eventually obtain the mutex.

*Thread 1 (yellow) requests the mutex but never gets it*

concurrent execution



time

Difficult to provide in practice and timing variations usually provide this property naturally



On to the lecture!

# Lecture schedule

- Atomic operations
- Mutex implementations

# Building blocks

- Memory reads and memory writes
  - later: read-modify-writes
- We need to guarantee that our reads and writes actually go to memory.
  - And other properties we will see soon
- To do this, we will use C++ atomic operations

# A historical perspective

- Adding concurrency support to a programming language is hard!
- The memory model defines how threads can safely share memory
- Java tried to do this,

wikipedia

The original Java memory model, developed in 1995, was widely perceived as broken, preventing many runtime optimizations and not providing strong enough guarantees for code safety. It was updated through the [Java Community Process](#), as Java Specification Request 133 (JSR-133), which took effect in 2004, for [Tiger \(Java 5.0\)](#).<sup>[1][2]</sup>

Brian Goetz (2019)

It is worth noting that **broken** techniques like double-checked locking are still **broken** under the new memory model, a

# A historical perspective

- How is C++?
- Has issues (imprecise, not modular)
  - but at least considered safe
  - Specification makes it difficult to reason about all programs
  - Open problem!
- Luckily mutexes (and their implementations) avoid the problematic areas of the language!

# Our primitive instructions

- Types: `atomic_int`
- Interface (C++ provides overloaded operators):
  - `load`
  - `store`
- Properties:
  - loads and stores will always go to memory.
  - compiler memory fence
  - hardware memory fence

# Atomic properties

- loads and stores will always go to memory
- Compiler example, performance difference

# Atomic properties

- loads and stores will always go to memory
- Compiler example, performance difference
- Compiler makes reasoning about parallel code hard, but big performance improvements:
  - Example:  $O(2048)$  vs.  $O(1)$



# Atomic properties

- Compiler Fence
- Compiler can be aggressive with memory operations:
  - For non-atomic memory locations, the following optimizations are valid

# Atomic properties

- Compiler Fence
- Compiler can be aggressive with memory operations:
  - For non-atomic memory locations, the following optimizations are valid

```
a[i] = 0;  
a[i] = 1;
```

can be optimized to:

```
a[i] = 1;
```

# Atomic properties

- Compiler Fence
- Compiler can be aggressive with memory operations:
  - For non-atomic memory locations, the following optimizations are valid

```
a[i] = 0;  
a[i] = 1;
```

can be optimized to:

```
a[i] = 1;
```

```
x = a[i];  
x2 = a[i];
```

can be optimized to:

```
x = a[i];  
x2 = x;
```

# Atomic properties

- Compiler Fence
- Compiler can be aggressive with memory operations:
  - For non-atomic memory locations, the following optimizations are valid

```
a[i] = 0;  
a[i] = 1;
```

can be optimized to:

```
a[i] = 1;
```

```
x = a[i];  
x2 = a[i];
```

can be optimized to:

```
x = a[i];  
x2 = x;
```

```
a[i] = 6;  
x = a[i];
```

can be optimized to:

```
x = 6;
```


# Atomic properties

- Compiler Fence
- Compiler can be aggressive with memory operations:
  - For non-atomic memory locations, the following optimizations are valid
- And many others... especially when you consider mixing with other optimizations
  - Very difficult to understand when/where memory accesses will actually occur in your code

# Atomic properties

- Compiler Fence

Compiler cannot keep `personal_account` in a register past the mutex




because this thread needs to see the updated view

# Atomic properties

- Compiler Fence

what can go wrong if the compiler doesn't write values to memory?




# Atomic properties

- Compiler Fence

what can go wrong if the compiler doesn't write values to memory?

*initially personal\_account is 0*





# Atomic properties

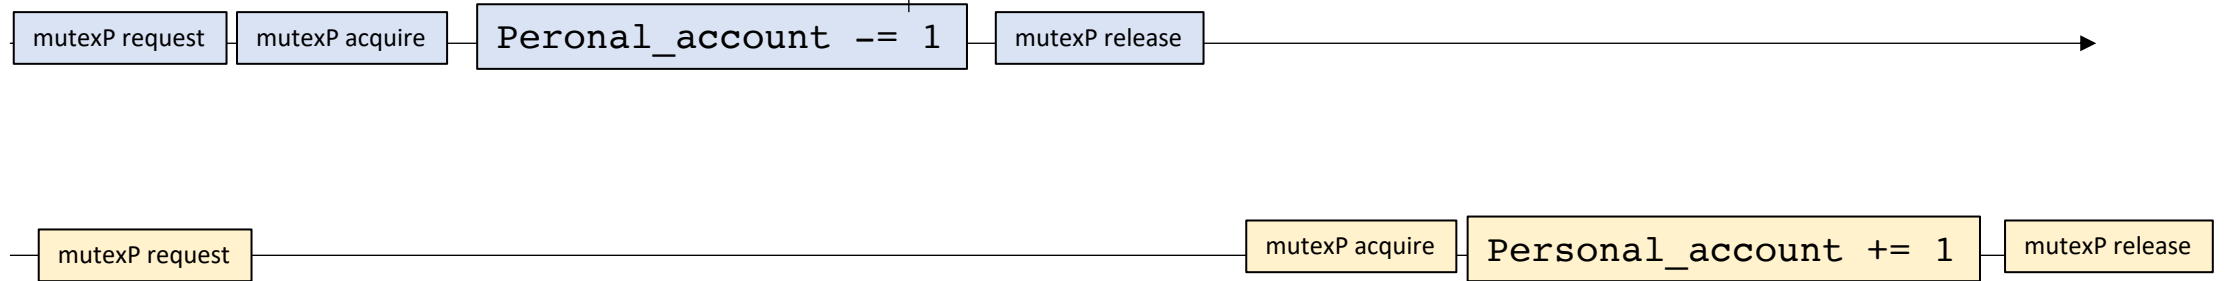
- Compiler Fence

what can go wrong if the compiler doesn't write values to memory?

*initially personal\_account is 0*

*loads 0*

`reg = *personal_account - 1;`



# Atomic properties


- Compiler Fence

what can go wrong if the compiler doesn't write values to memory?

*initially personal\_account is 0*

*loads 0*

reg = \*personal\_account - 1;



# Atomic properties

- Compiler Fence

what can go wrong if the compiler doesn't write values to memory?

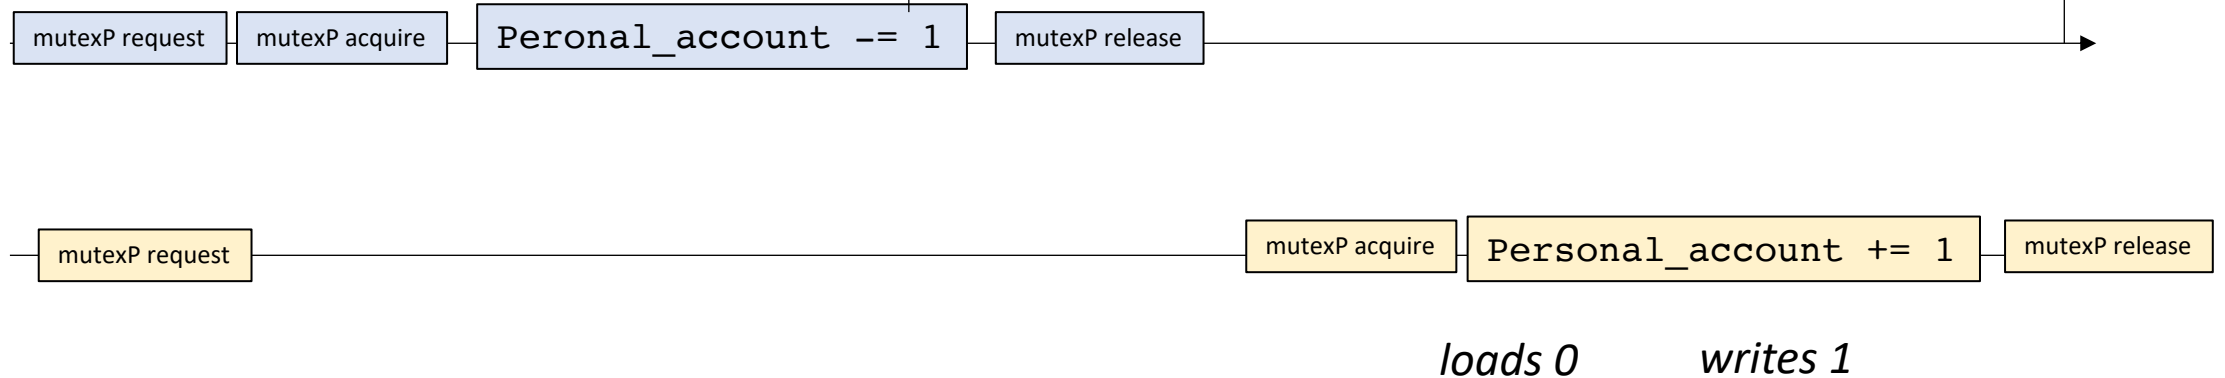
*initially personal\_account is 0*

*loads 0*

***personal\_account is -1***

`reg = *personal_account - 1;`

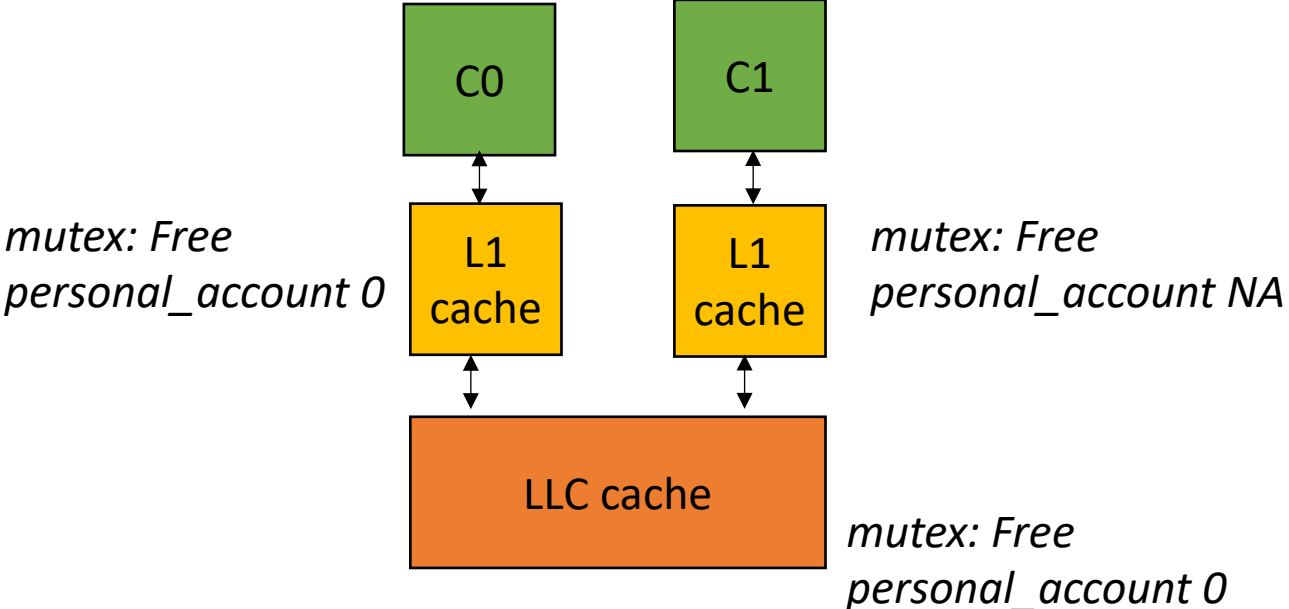
`*personal_account = reg;`



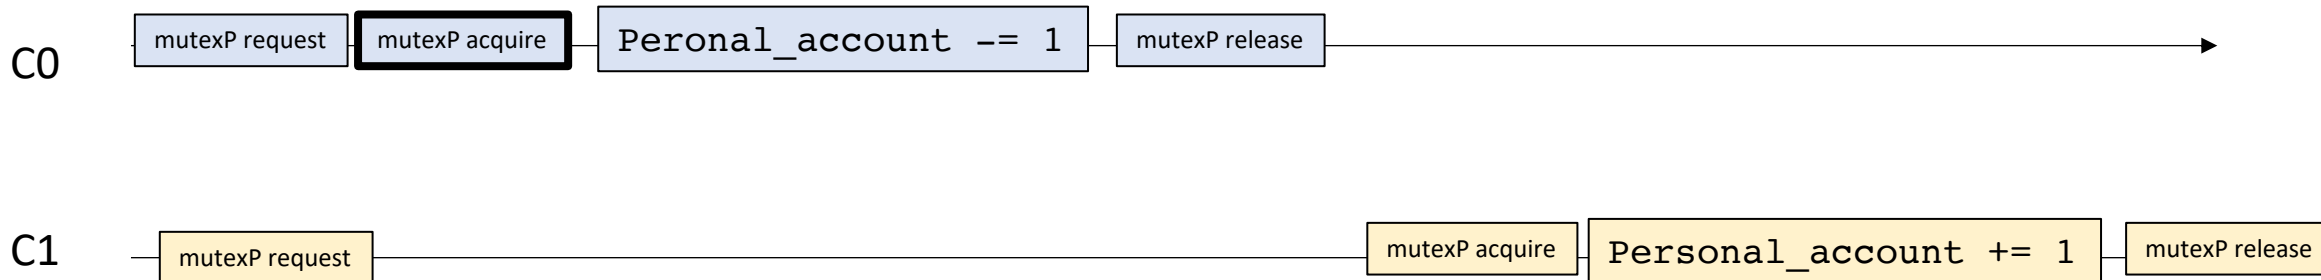

# Atomic properties

- Also provides a memory barrier

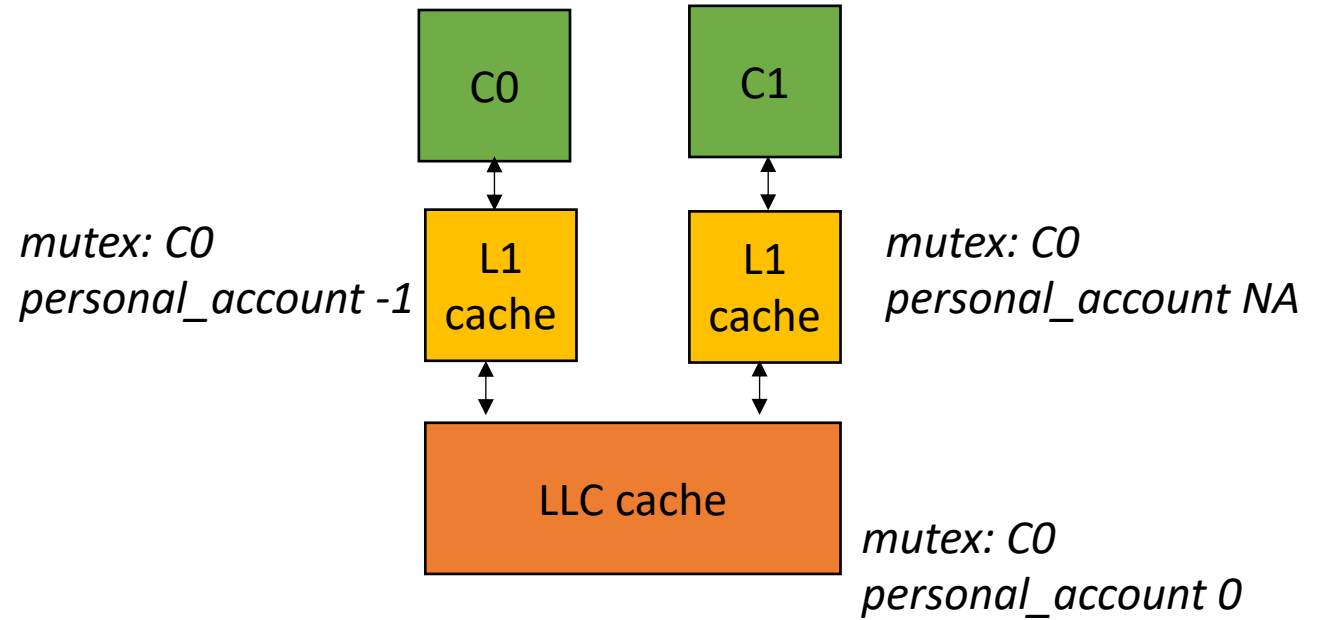
- Memory Fence (or Memory Barrier)





- Memory Fence (or Memory Barrier)



- Memory Fence (or Memory Barrier)





- Memory Fence (or Memory Barrier)







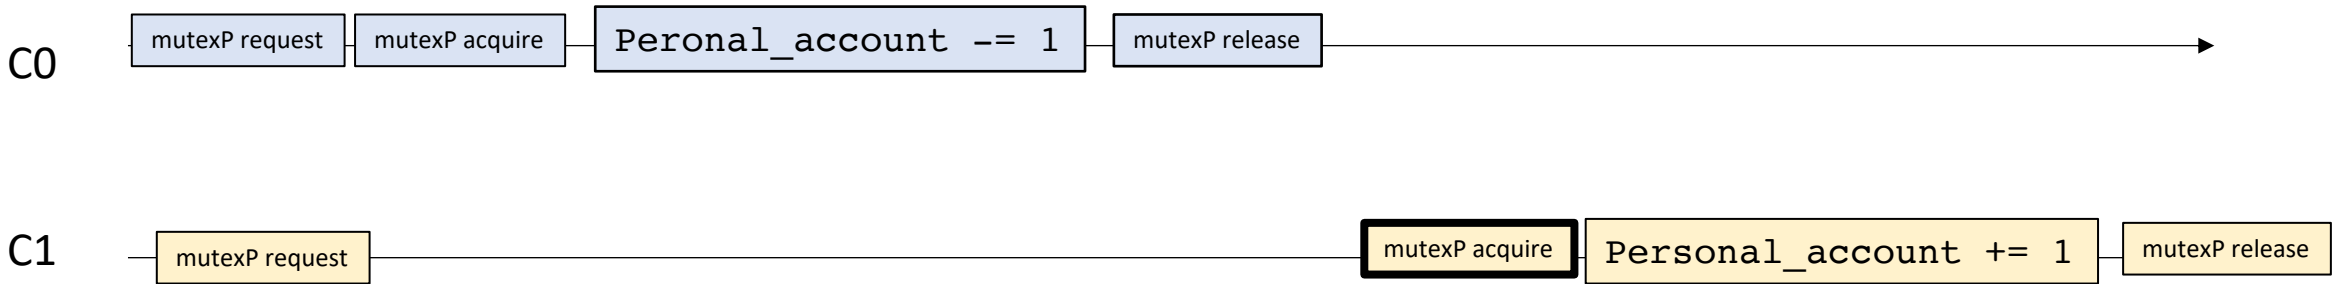
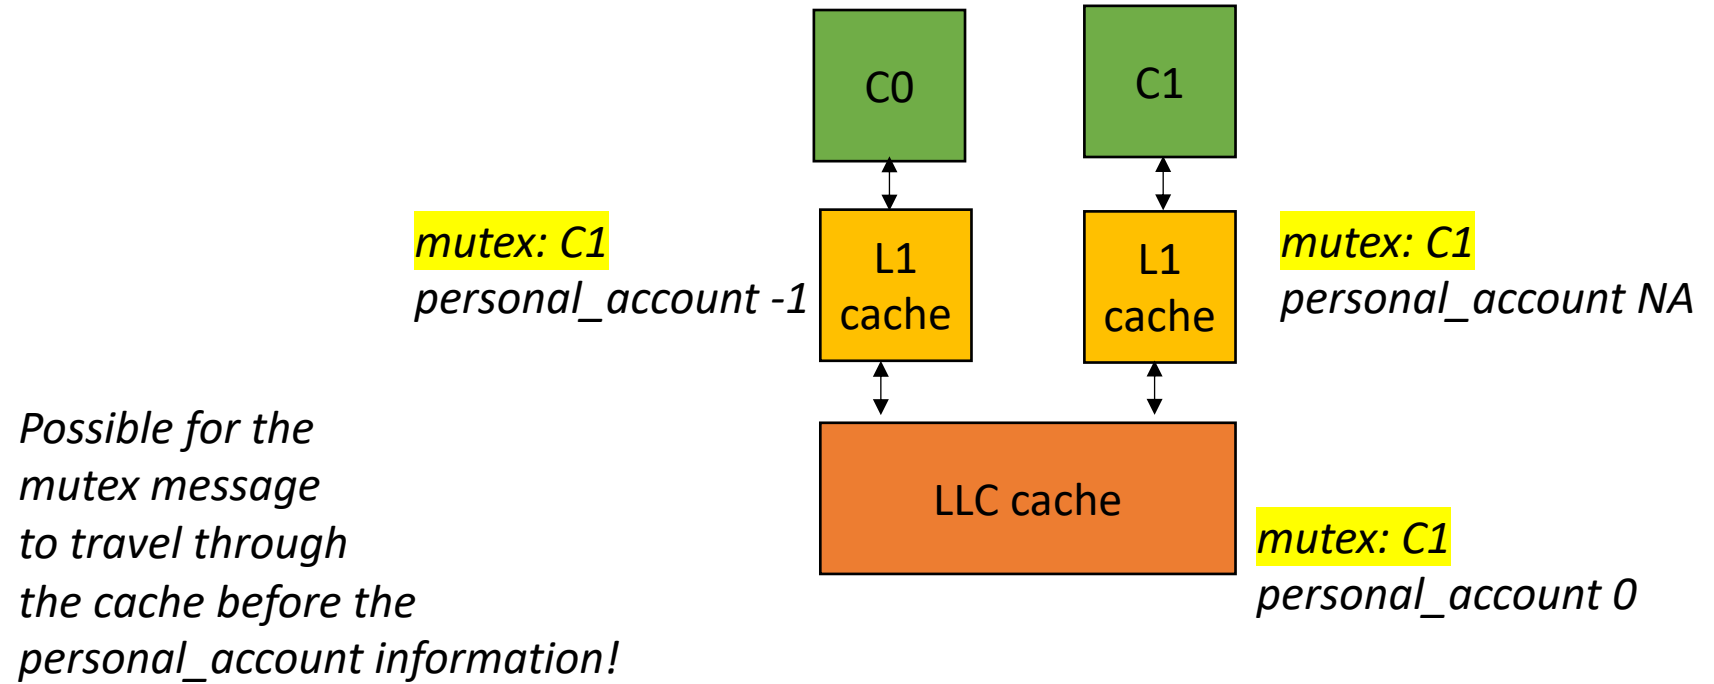
- Memory Fence (or Memory Barrier)



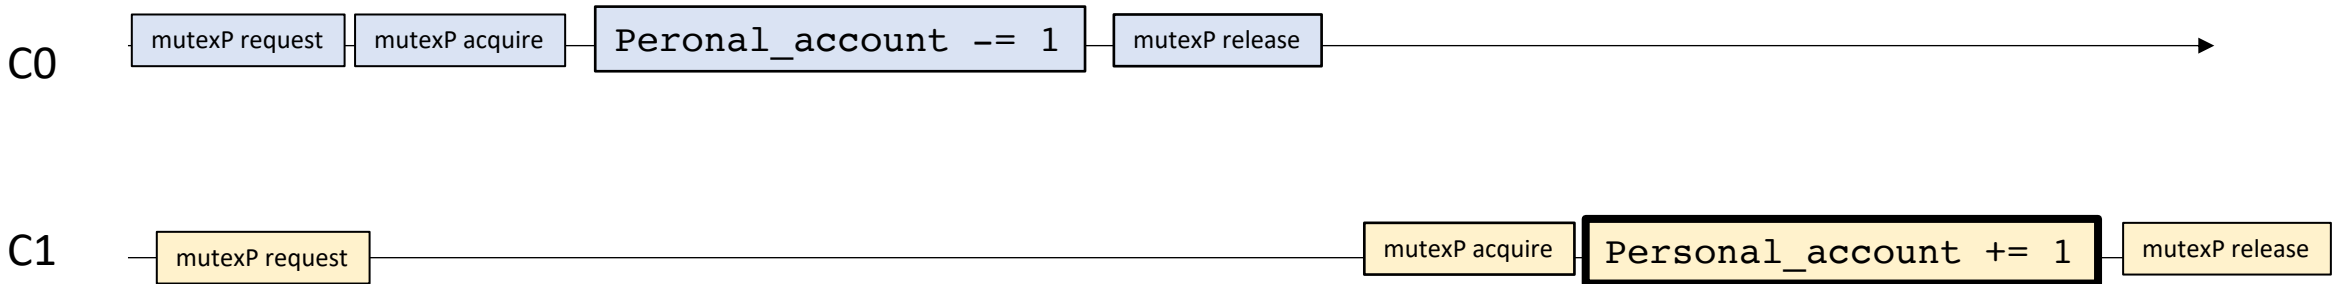
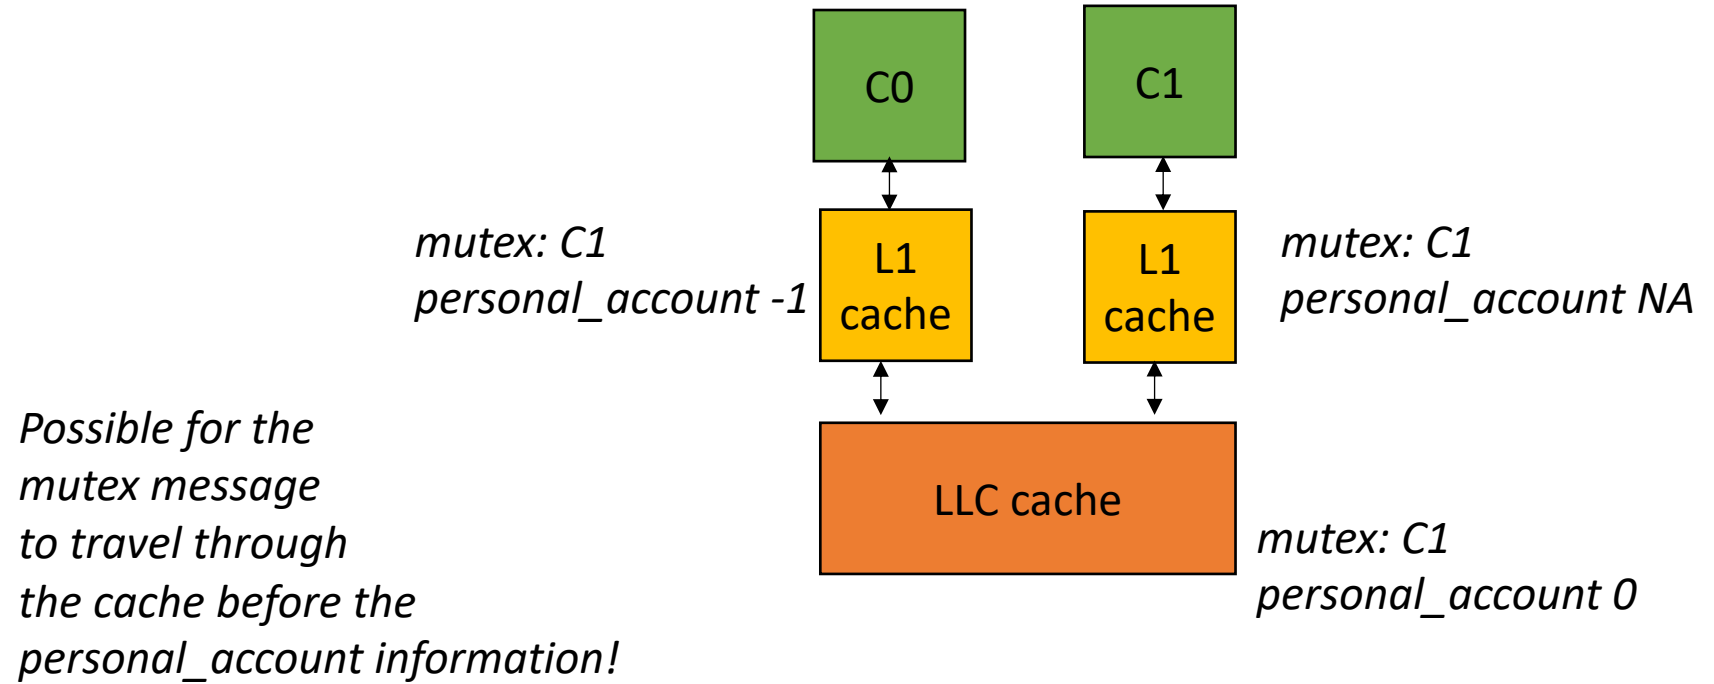
- Memory Fence (or Memory Barrier)





- Memory Fence (or Memory Barrier)



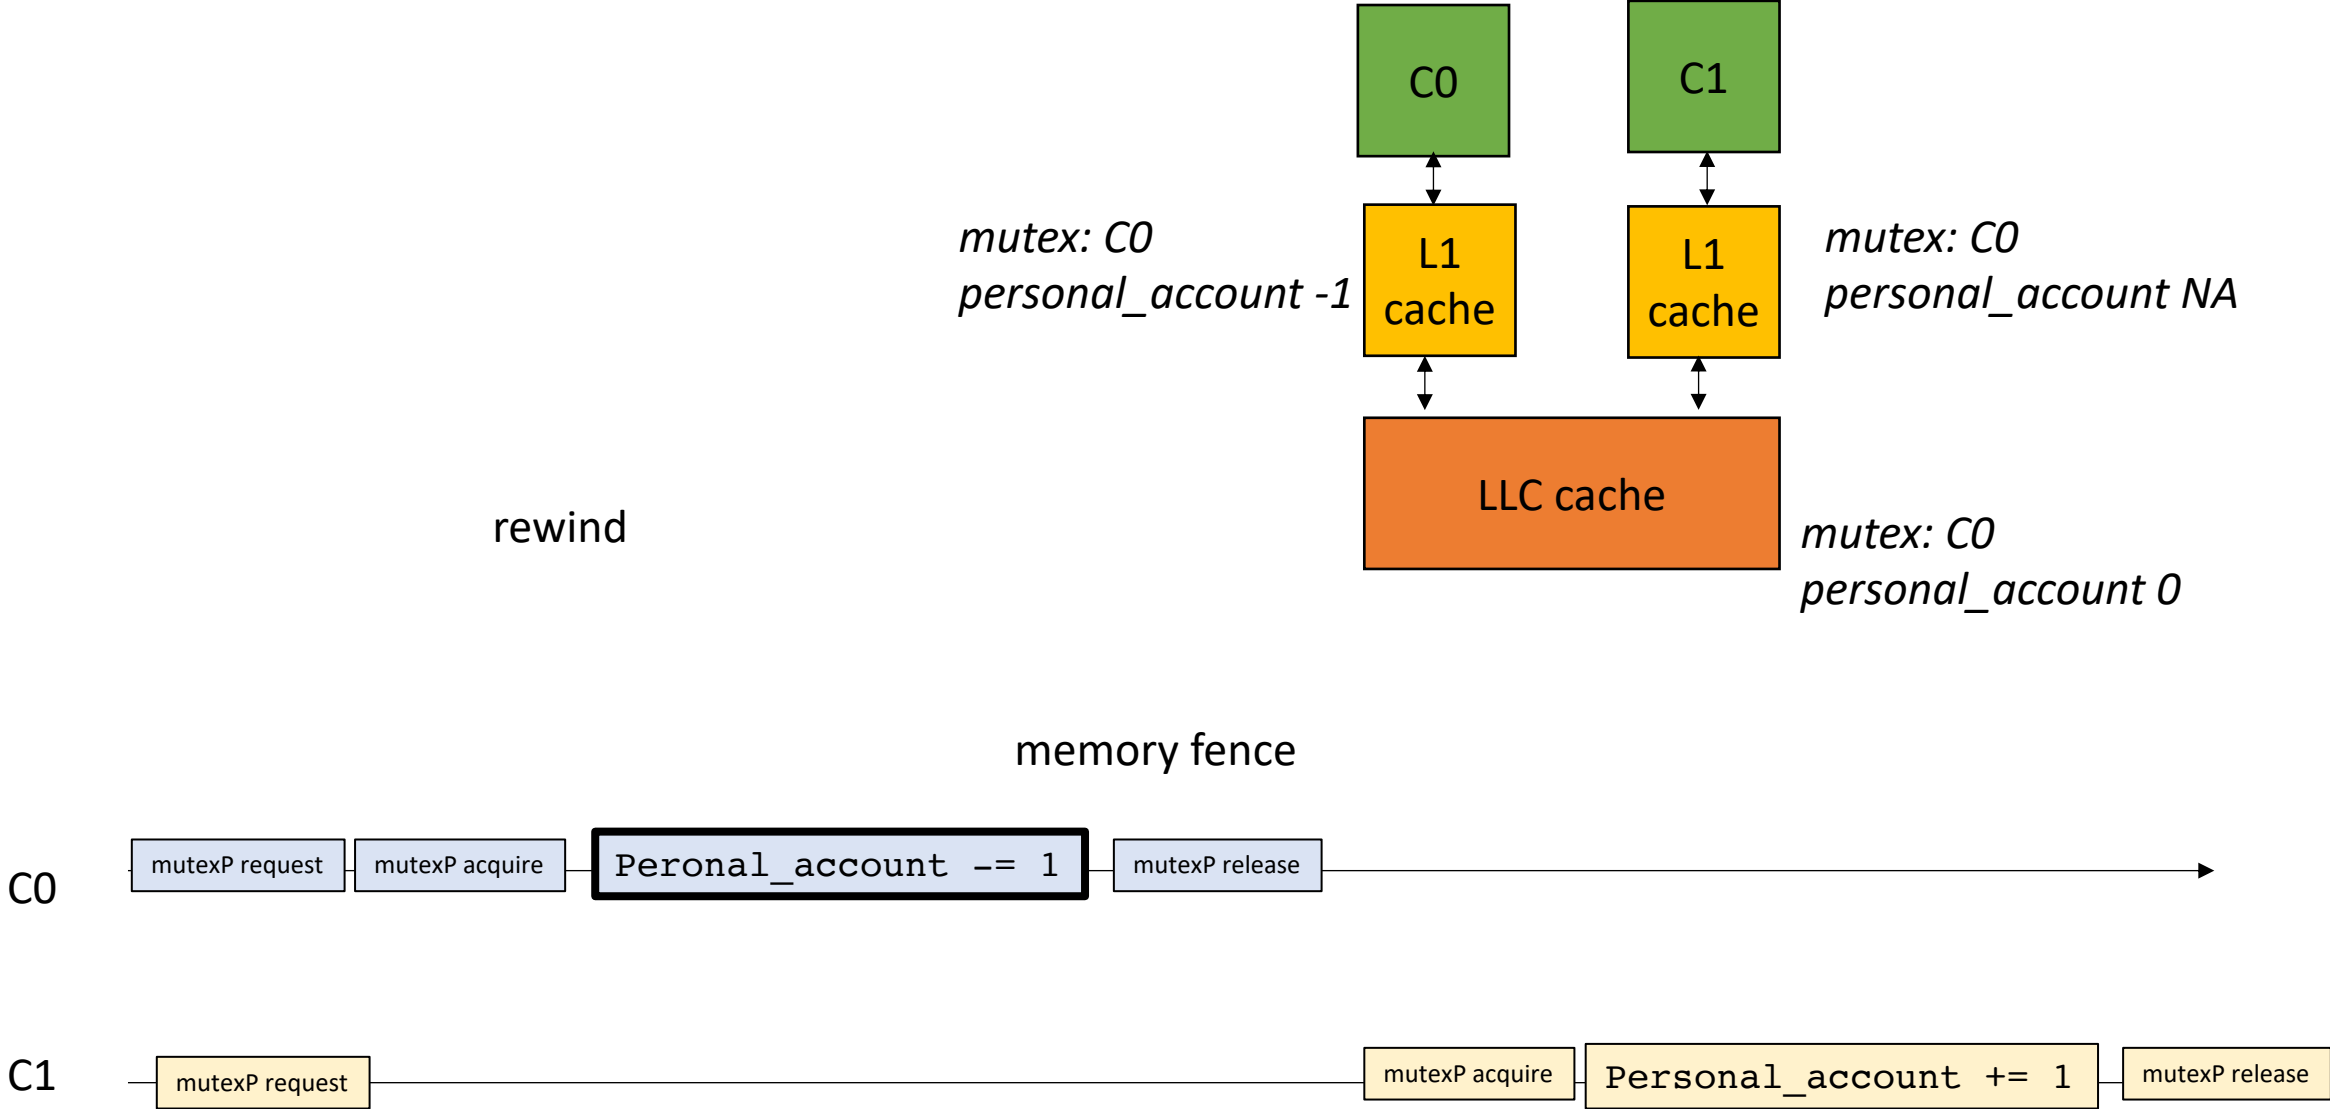
- Memory Fence (or Memory Barrier)





- Memory Fence (or Memory Barrier)





- Memory Fence (or Memory Barrier)



- Memory Fence (or Memory Barrier)





- Memory Fence (or Memory Barrier)


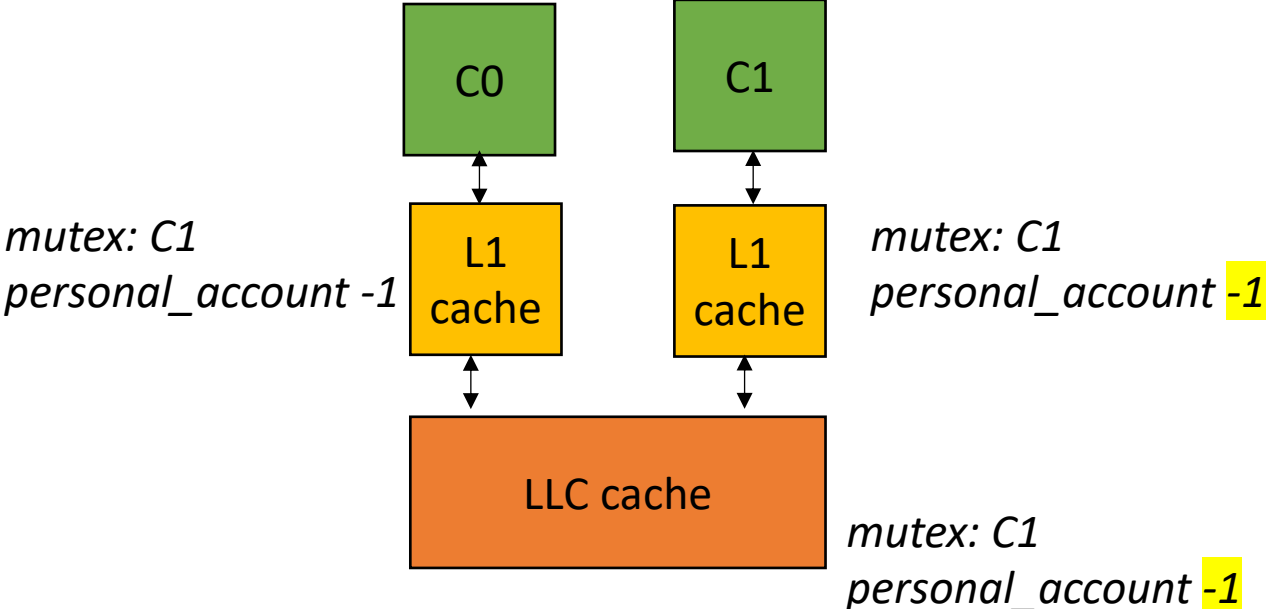




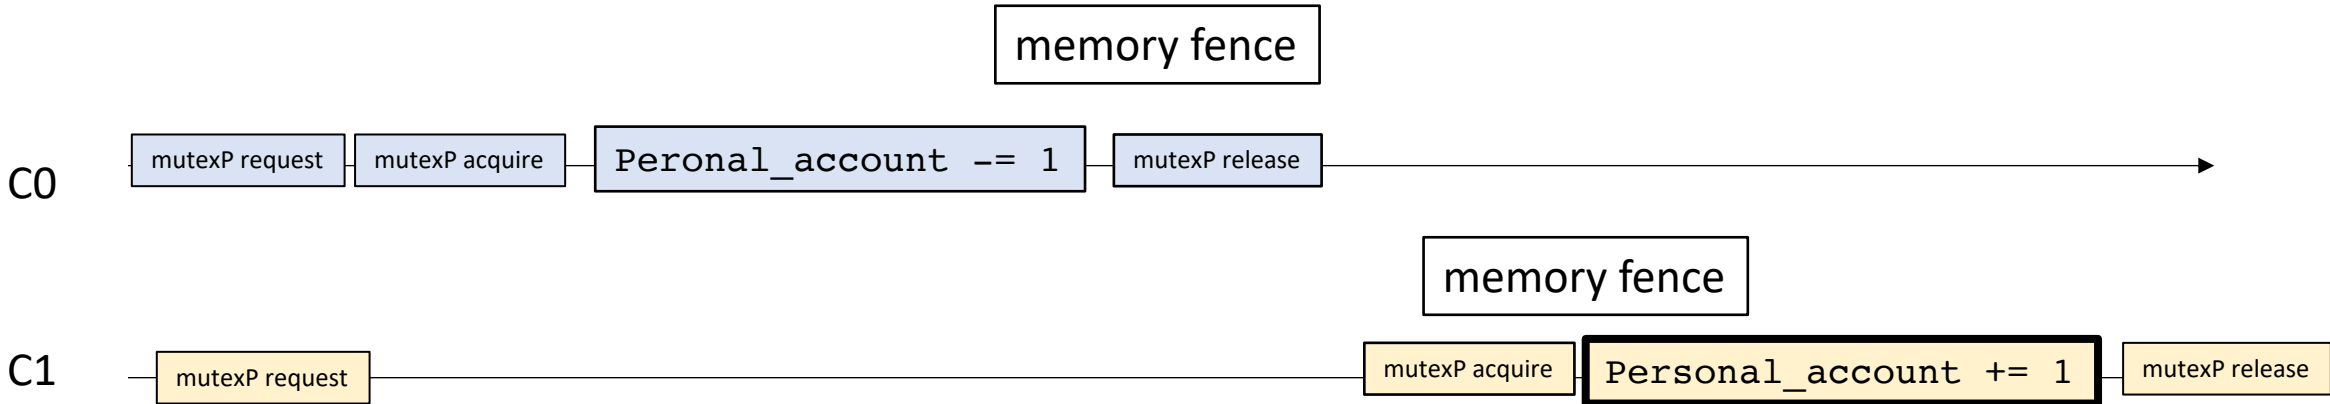
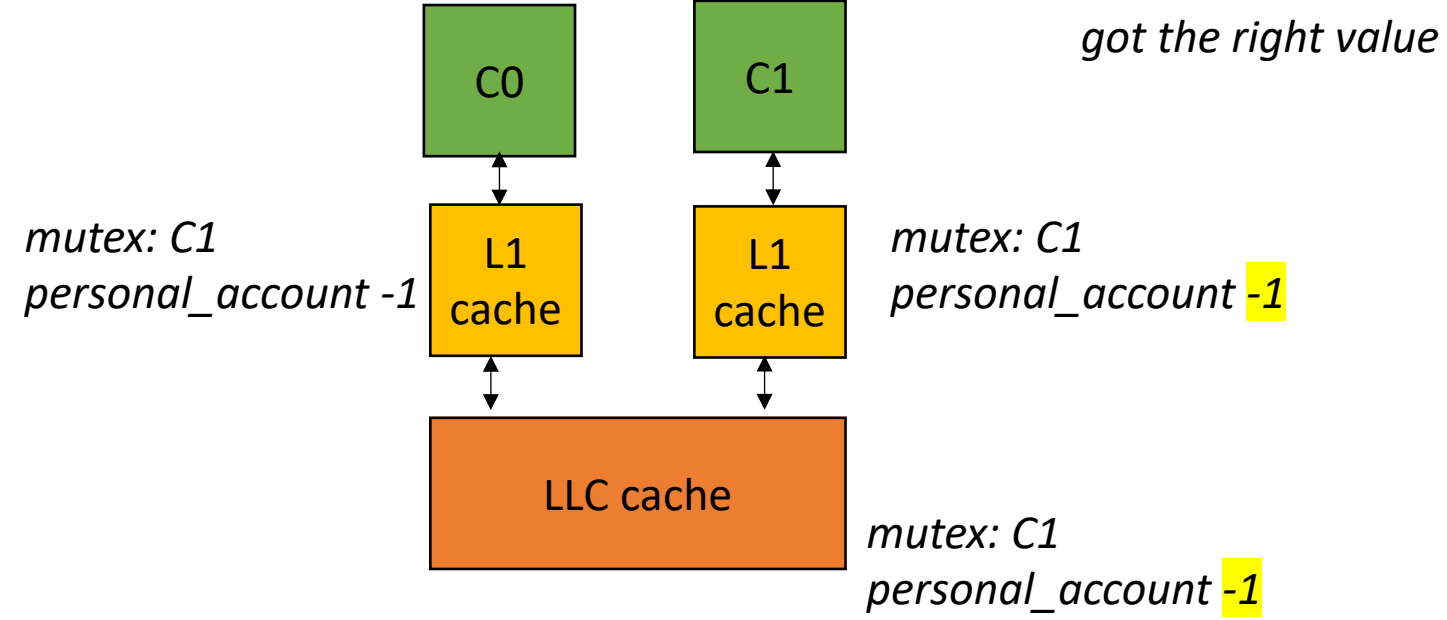
- Memory Fence (or Memory Barrier)



- Memory Fence (or Memory Barrier)



- Memory Fence (or Memory Barrier)



- **Memory Fence (or Memory Barrier)**

different architectures have different memory barriers

Intel X86 naturally manages caches in order

ARM and PowerPC let cache values flow out-of-order

GPUs let caches flow out-of-order

RISC-V has two models:


- more like x86: easier to program

- more like ARM: faster and more energy efficient

*For mutexes, atomics will naturally handle the memory fences for us!*


# Atomics

- What do those fences (compiler and memory) give us?
- Atomics were designed so that we can implement things like mutexes!



# Atomics

- What do those fences (compiler and memory) give us?
- Atomics were designed so that we can implement things like mutexes!



# Lecture schedule

- Atomic operations
- Mutex implementations

# Mutex Implementations

- We will just consider two threads for now, with thread ids 0, 1
- A first attempt:
  - A mutex contains a boolean.
  - The mutex value set to 0 means that it is free. 1 means that some thread is holding it.
  - To lock the mutex, you wait until it is set to 0, then you store 1 in the flag.
  - To unlock the mutex, you set the mutex back to 0.



# Mutex Implementations

```
#include <atomic>
using namespace std;

class Mutex {
public:
    Mutex() {
        flag = 0;
    }
    void lock();
    void unlock();
private:
    atomic_bool flag;
};
```

mutex is initialized to “free”

atomic\_bool for our memory location

# Mutex Implementations

```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```

While the mutex is not available (i.e. another thread has it)

Once the mutex is available, we will claim it

# Mutex Implementations

```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```

While the mutex is not available (i.e. another thread has it)

Once the mutex is available, we will claim it

*Whats up with this while loop?*

# Mutex Implementations

```
void unlock() {  
    flag.store(0);  
}
```

To release the mutex, we just set it back to 0 (available)

# Analysis

```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```

```
void unlock() {  
    flag.store(0);  
}
```


Thread 0:  
m.lock();  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();

core 0



core 1



# Analysis

```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```

```
void unlock() {  
    flag.store(0);  
}
```

Thread 0:

```
m.lock();  
m.unlock();
```


Thread 1:

```
m.lock();  
m.unlock();
```

core 0

m.request

core 1



# Analysis

```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```

```
void unlock() {  
    flag.store(0);  
}
```


Thread 0:

```
m.lock();  
m.unlock();
```

Thread 1:

```
m.lock();  
m.unlock();
```

Mutex request



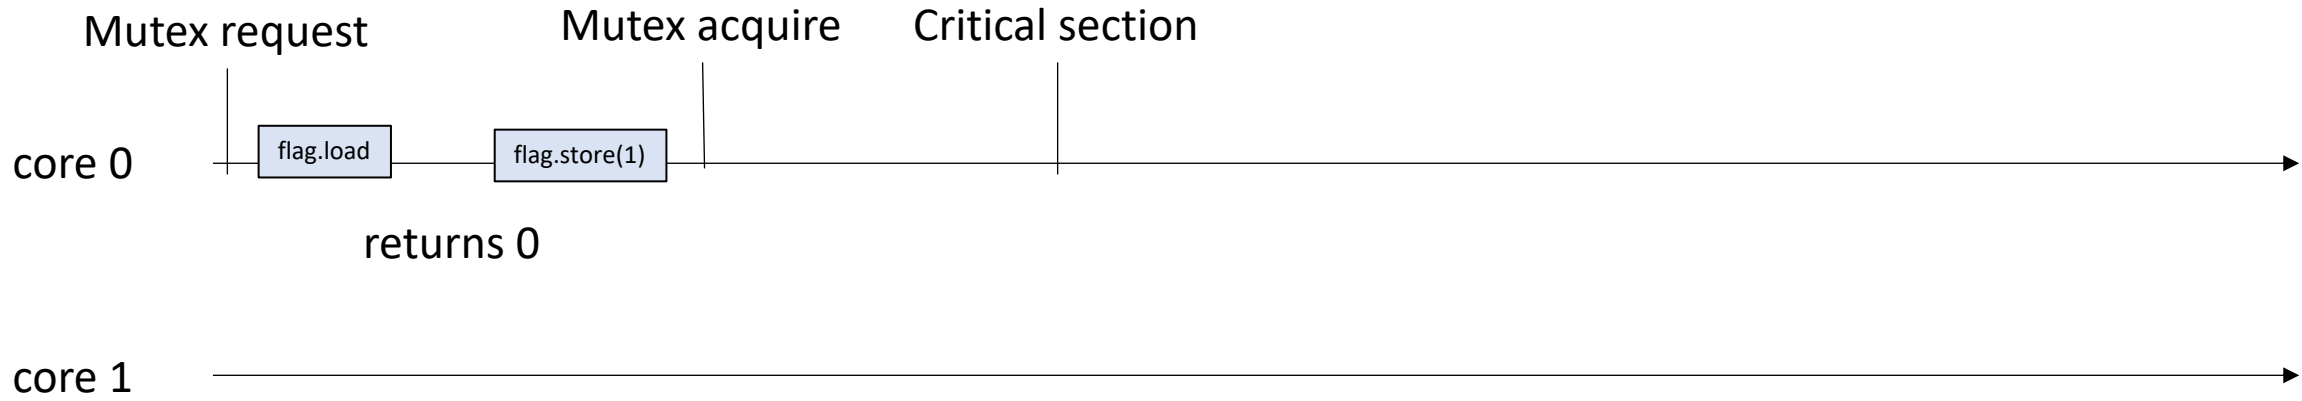
# Analysis

```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```

```
void unlock() {  
    flag.store(0);  
}
```

Thread 0:  
m.lock();  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();






# Analysis

```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```

```
void unlock() {  
    flag.store(0);  
}
```

Thread 0:  
m.lock();  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();




# Analysis

```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```

```
void unlock() {  
    flag.store(0);  
}
```

Thread 0:  
m.lock();  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();



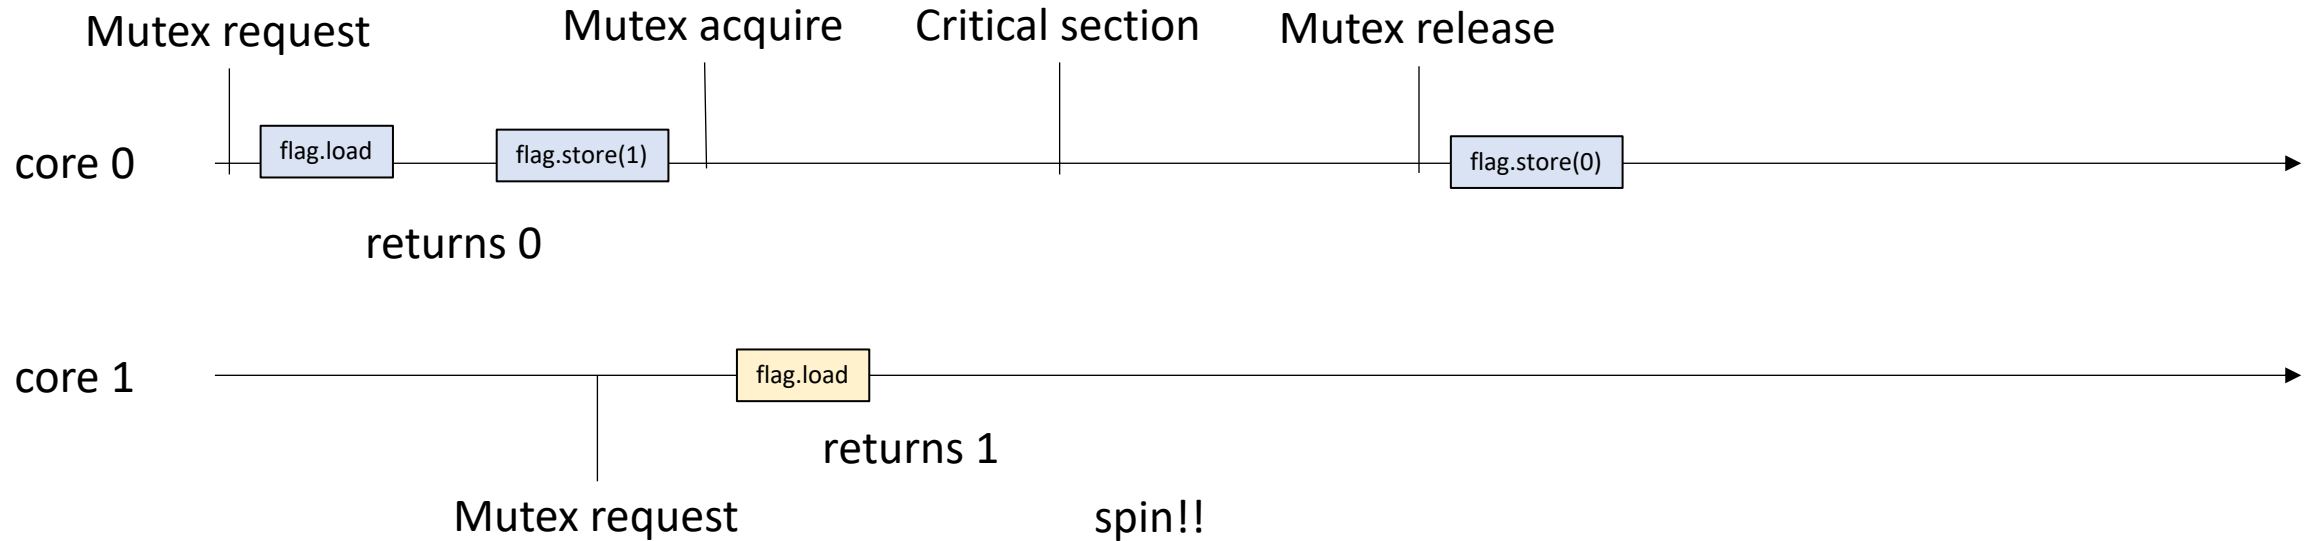
# Analysis

```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```

```
void unlock() {  
    flag.store(0);  
}
```

Thread 0:  
m.lock();  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();




# Analysis

```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```

```
void unlock() {  
    flag.store(0);  
}
```

Thread 0:  
m.lock();  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();




# Analysis

```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```

```
void unlock() {  
    flag.store(0);  
}
```

Thread 0:  
m.lock();  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();



# Analysis


```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```

```
void unlock() {  
    flag.store(0);  
}
```

Thread 0:  
m.lock();  
m.unlock();  
m.lock();  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();

**Mutual Exclusion property!**  
**critical sections do not overlap!**



# Analysis

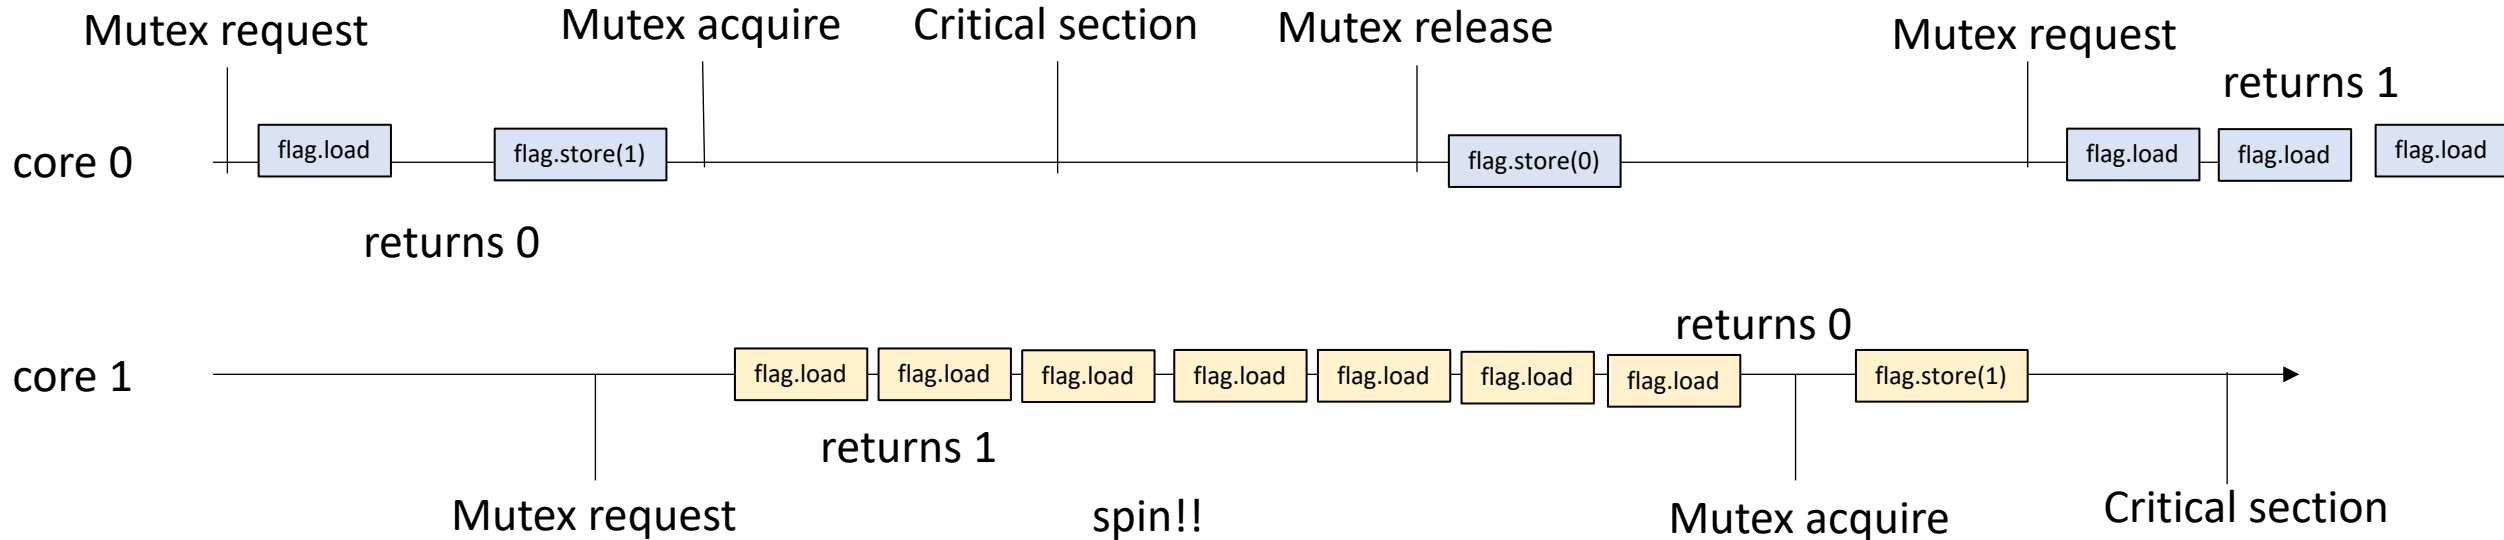
```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```



```
void unlock() {  
    flag.store(0);  
}
```

Thread 0:  
m.lock();  
m.unlock();  
m.lock();  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();


***Mutual Exclusion property!***  
***critical sections do not overlap!***








core 0



core 1



# Analysis


```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```

```
void unlock() {  
    flag.store(0);  
}
```

Thread 0:  
**m.lock();**  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();

*Lets try another interleaving*



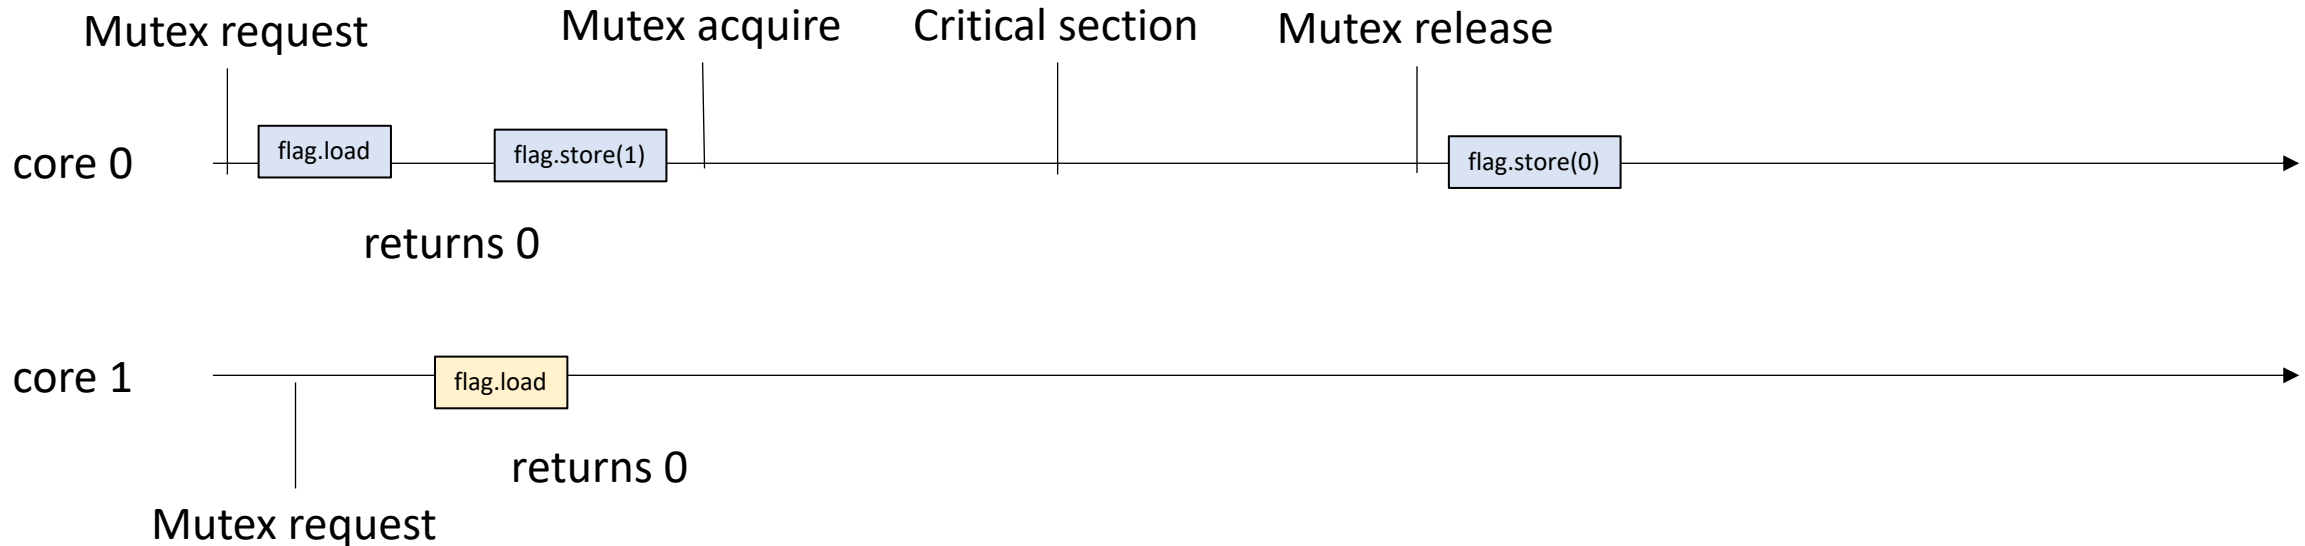
# Analysis

```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```

```
void unlock() {  
    flag.store(0);  
}
```

Thread 0:  
**m.lock();**  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();



# Analysis

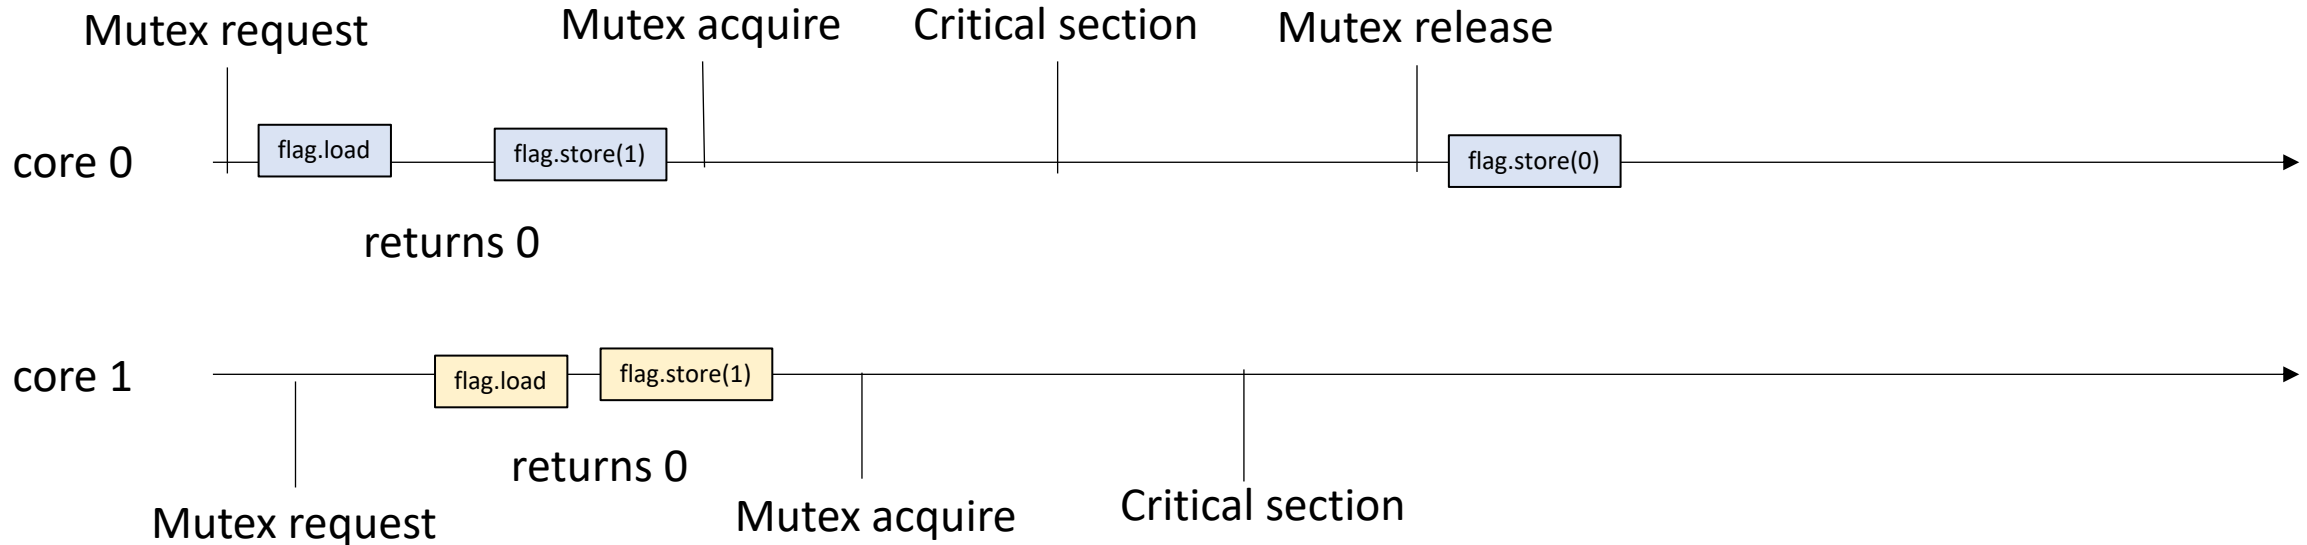
```
void lock() {  
    while (flag.load() == 1);  
    flag.store(1);  
}
```

```
void unlock() {  
    flag.store(0);  
}
```

Thread 0:  
`m.lock();`  
`m.unlock();`

Thread 1:  
`m.lock();`  
`m.unlock();`

*Critical sections overlap! This mutex implementation is not correct!*



# Mutex Implementations

- Second attempt:
  - A flag for each thread (2 flags)
  - If you want the mutex, set your flag to 1.
  - Spin while the other flag is 1 (the other thread has the mutex)
  - To release the mutex, set your flag to 0

# Mutex Implementations

```
#include <atomic>
using namespace std;

class Mutex {
public:
    Mutex() {
        flag[0] = flag[1] = 0;
    }

    void lock();
    void unlock();

private:
    atomic_bool flag[2];
};
```

both initialized to 0

two flags this time

# Mutex Implementations

```
void lock() {  
    int i = thread_id;  
    flag[i].store(1);  
    int j = i == 0 ? 1 : 0;  
    while (flag[j].load() == 1);  
}
```

Thread id (0, or 1)

Mark your intention to take the lock

Wait for other thread to leave the  
critical section

# Mutex Implementations

```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread id (0, or 1)

Mark your flag to say you have left the critical section.



# Analysis

```
void lock() {  
    int i = thread_id;  
    flag[i].store(1);  
    int j = i == 0 ? 1 : 0;  
    while (flag[j].load() == 1);  
}
```

```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread 0:  
m.lock();  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();

core 0



core 1




# Analysis

```
void lock() {  
    int i = thread_id;  
    flag[i].store(1);  
    int j = i == 0 ? 1 : 0;  
    while (flag[j].load() == 1);  
}
```

```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread 0:  
**m.lock();**  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();




# Analysis

```
void lock() {  
    int i = thread_id;  
    flag[i].store(1);  
    int j = i == 0 ? 1 : 0;  
    while (flag[j].load() == 1);  
}
```

```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread 0:  
m.lock();  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();



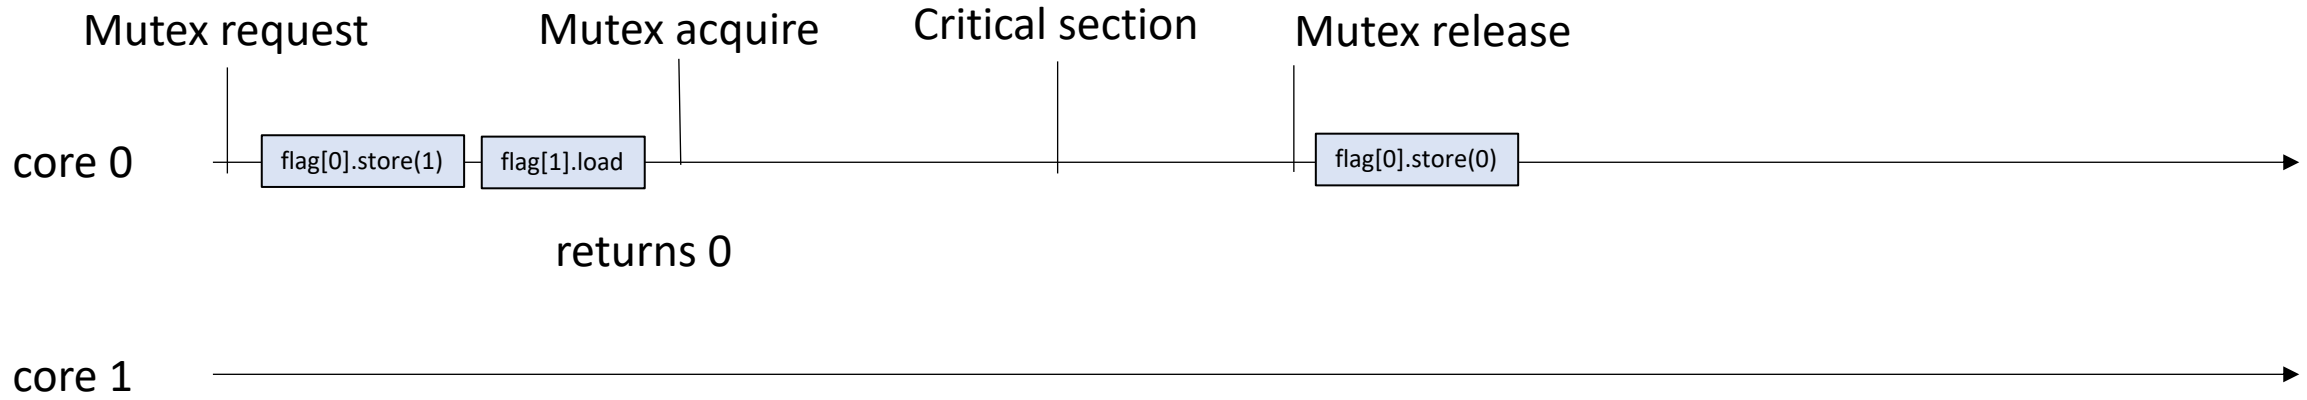
# Analysis

```
void lock() {  
    int i = thread_id;  
    flag[i].store(1);  
    int j = i == 0 ? 1 : 0;  
    while (flag[j].load() == 1);  
}
```

```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread 0:  
m.lock();  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();



# Analysis

```
void lock() {  
    int i = thread_id;  
    flag[i].store(1);  
    int j = i == 0 ? 1 : 0;  
    while (flag[j].load() == 1);  
}
```


```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread 0:  
m.lock();  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();

*critical sections do not overlap!*

*proof?*



# Analysis

```
void lock() {  
    int i = thread_id;  
    flag[i].store(1);  
    int j = i == 0 ? 1 : 0;  
    while (flag[j].load() == 1);  
}
```


```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread 0:

```
m.lock();  
m.unlock();
```

Thread 1:

```
m.lock();  
m.unlock();
```



# Analysis

```
void lock() {  
    int i = thread_id;  
    flag[i].store(1);  
    int j = i == 0 ? 1 : 0;  
    while (flag[j].load() == 1);  
}
```


```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread 0:

```
m.lock();  
m.unlock();
```

Thread 1:

```
m.lock();  
m.unlock();
```



# Analysis

```
void lock() {  
    int i = thread_id;  
    flag[i].store(1);  
    int j = i == 0 ? 1 : 0;  
    while (flag[j].load() == 1);  
}
```

```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread 0:

```
m.lock();  
m.unlock();
```

Thread 1:

```
m.lock();  
m.unlock();
```





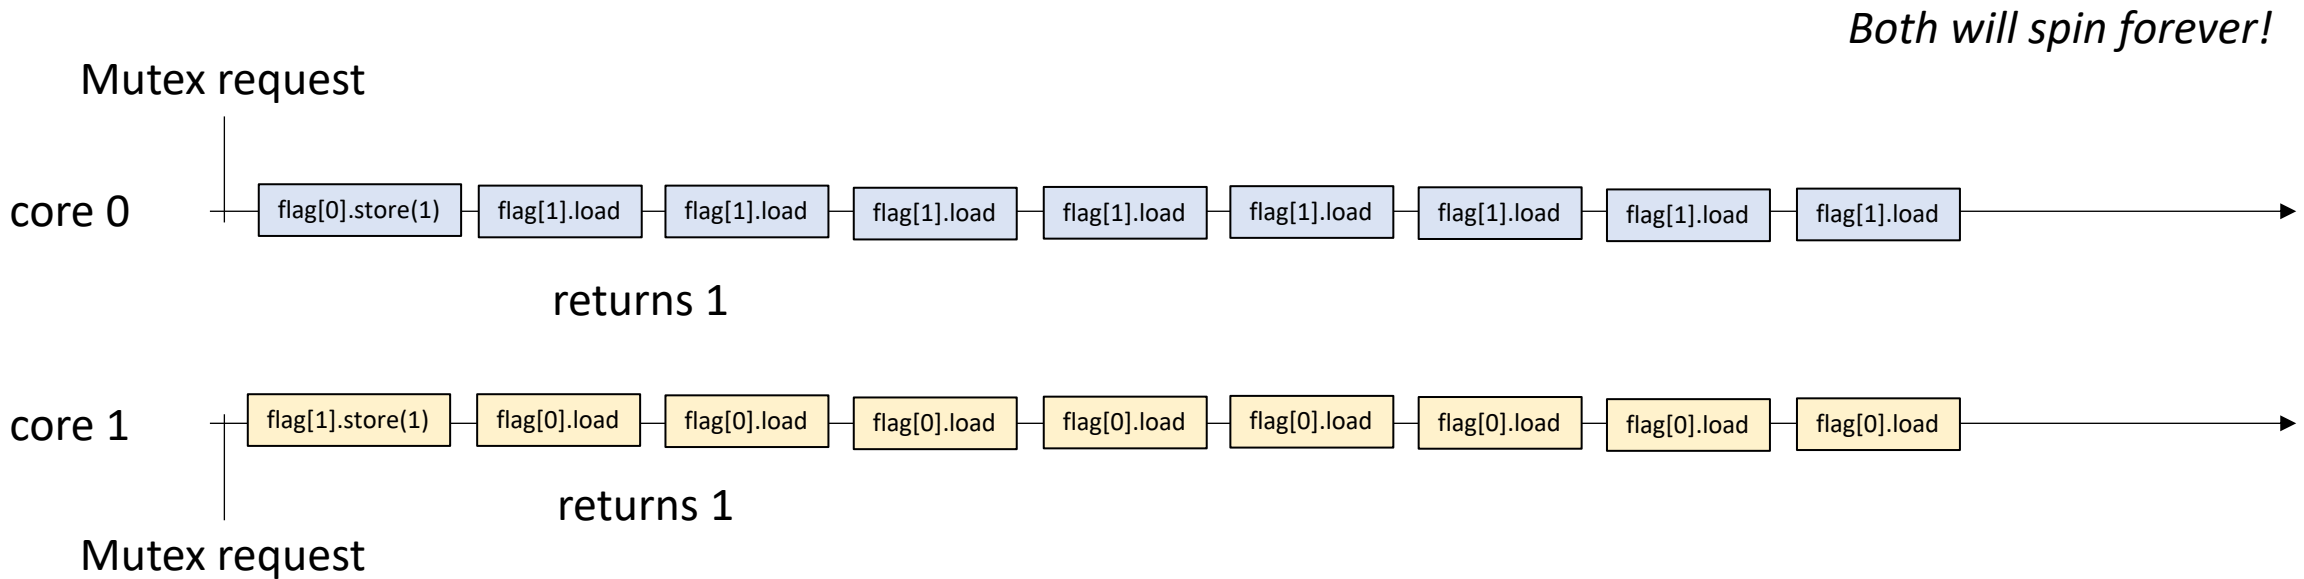
# Analysis

```
void lock() {  
    int i = thread_id;  
    flag[i].store(1);  
    int j = i == 0 ? 1 : 0;  
    while (flag[j].load() == 1);  
}
```

```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread 0:  
`m.lock();`  
`m.unlock();`

Thread 1:  
`m.lock();`  
`m.unlock();`



# Properties of mutexes

## Three properties

- **Deadlock Freedom** - If a thread has requested the mutex, and no thread currently holds the mutex, the mutex must be acquired by one of the requesting threads

Program cannot hang here  
Either thread 0 or thread 1 must acquire the mutex

concurrent execution



time

# Mutex Implementations

Third attempt

# Mutex Implementations

```
class Mutex {  
public:  
    Mutex() {  
        victim = -1;  
    }
```

initialized to -1

```
    void lock();  
    void unlock();
```

```
private:  
    atomic_int victim;  
};
```

back to a single variable

# Mutex Implementations

```
void lock() {  
    victim.store(thread_id);  
    while (victim.load() == thread_id);  
}
```

Volunteer to be the victim

Victims only job is to spin

# Mutex Implementations

```
void unlock() {}
```

No unlock!

```
void lock() {  
    victim.store(thread_id);  
    while (victim.load() == thread_id);  
}
```

```
void unlock() {}
```

Thread 0:

```
m.lock();
```

```
m.unlock();
```

Mutex request

core 0



```
void lock() {  
    victim.store(thread_id);  
    while (victim.load() == thread_id);  
}
```

```
void unlock() {}
```

Thread 0:

`m.lock();`

`m.unlock();`

Mutex request





```
void lock() {  
    victim.store(thread_id);  
    while (victim.load() == thread_id);  
}
```

```
void unlock() {}
```

Thread 0:

`m.lock();`


`m.unlock();`

spins forever if  
the second thread  
never tries to take the mutex!

Mutex request

returns 0

core 0



```
void lock() {  
    victim.store(thread_id);  
    while (victim.load() == thread_id);  
}
```

```
void unlock() {}
```

Thread 0:  
`m.lock();`  
`m.unlock();`

Thread 1:  
`m.lock();`  
`m.unlock();`




```
void lock() {
    victim.store(thread_id);
    while (victim.load() == thread_id);
}
```

```
void unlock() {}
```

Thread 0:  
m.lock();  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();



```

void lock() {
    victim.store(thread_id);
    while (victim.load() == thread_id);
}

```


```

void unlock() {}

```

Thread 0:  
m.lock();  
m.unlock();

Thread 1:  
**m.lock();**  
m.unlock();




```
void lock() {
    victim.store(thread_id);
    while (victim.load() == thread_id);
}
```

```
void unlock() {}
```

Thread 0:  
m.lock();  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();



# Mutex Implementations

Finally, we can make a mutex that works:

Use flags to mark interest

Use victim to break ties

Called the **Peterson Lock**

# Mutex Implementations

```
class Mutex {  
public:  
    Mutex() {  
        victim = -1;  
        flag[0] = flag[1] = 0;  
    }  
  
    void lock();  
    void unlock();  
  
private:  
    atomic_int victim;  
    atomic_bool flag[2];  
};
```

Initially:

No victim and no threads are interested in the critical section

flags and victim

# Mutex Implementations

```
void lock() {  
    int j = thread_id == 0 ? 1 : 0;  
    flag[thread_id].store(1);  
    victim.store(thread_id);  
    while (victim.load() == thread_id  
           && flag[j] == 1);  
}
```

j is the other thread

Mark ourself as interested

volunteer to be the victim in case of a tie

Spin only if:

there was a tie in wanting the lock,  
and I won the volunteer raffle to spin



# Mutex Implementations

```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

mark ourselves as uninterested



# Tie breaking with victim

```
void lock() {  
    int j = thread_id == 0 ? 1 : 0;  
    flag[thread_id].store(1);  
    victim.store(thread_id);  
    while (victim.load() == thread_id  
           && flag[j] == 1);  
}
```

```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread 0:


```
m.lock();  
m.unlock();
```

Thread 1:

```
m.lock();  
m.unlock();
```

Mutex request

only one of the stores will be in victim (one will overwrite the other)



# Tie breaking with victim

```
void lock() {  
    int j = thread_id == 0 ? 1 : 0;  
    flag[thread_id].store(1);  
    victim.store(thread_id);  
    while (victim.load() == thread_id  
           && flag[j] == 1);  
}
```

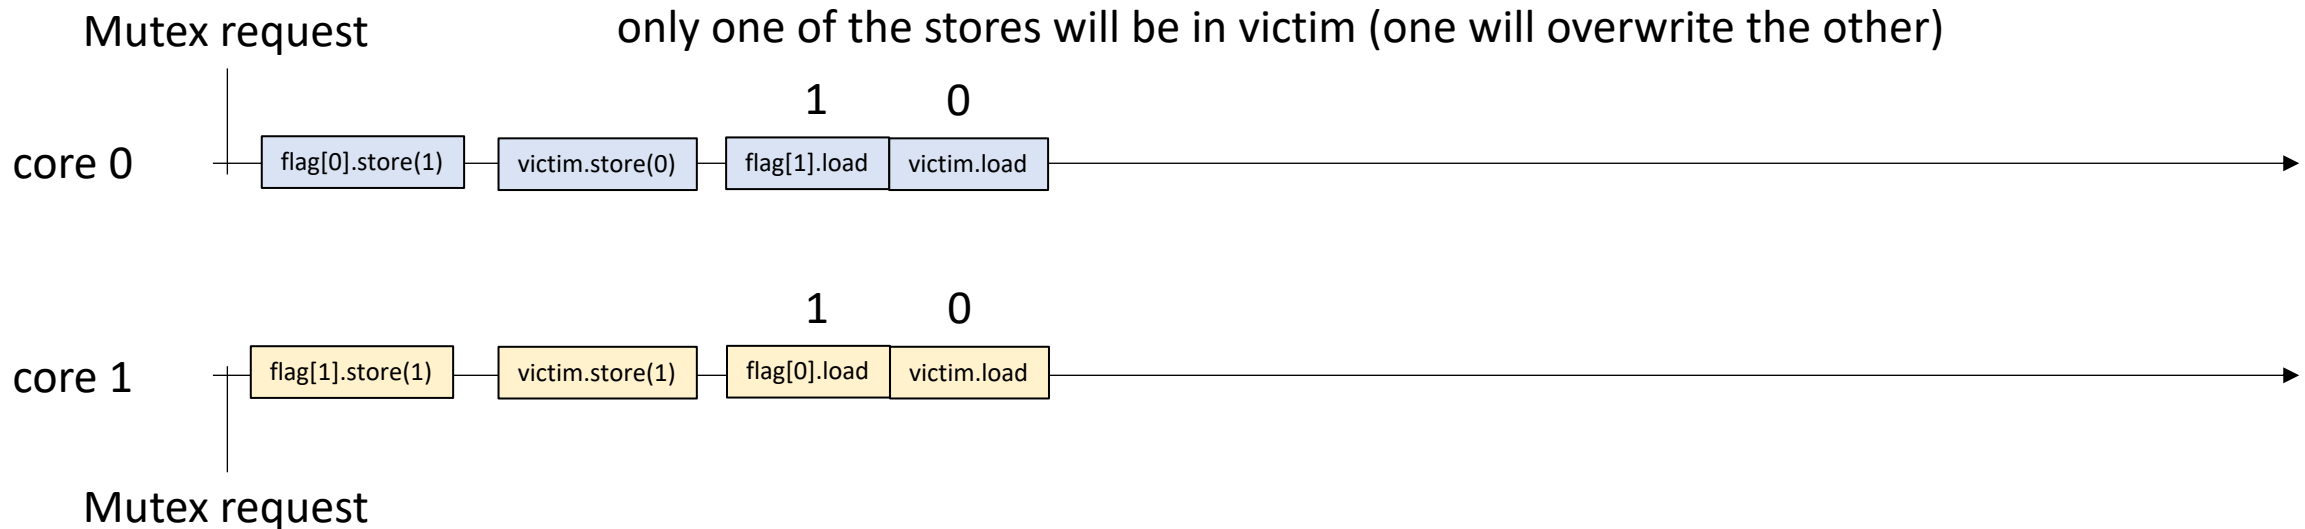
```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread 0:

```
m.lock();  
m.unlock();
```

Thread 1:

```
m.lock();  
m.unlock();
```




# Tie breaking with victim

```
void lock() {  
    int j = thread_id == 0 ? 1 : 0;  
    flag[thread_id].store(1);  
    victim.store(thread_id);  
    while (victim.load() == thread_id  
           && flag[j] == 1);  
}
```

```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread 0:  
**m.lock();**  
m.unlock();

Thread 1:  
m.lock();  
m.unlock();



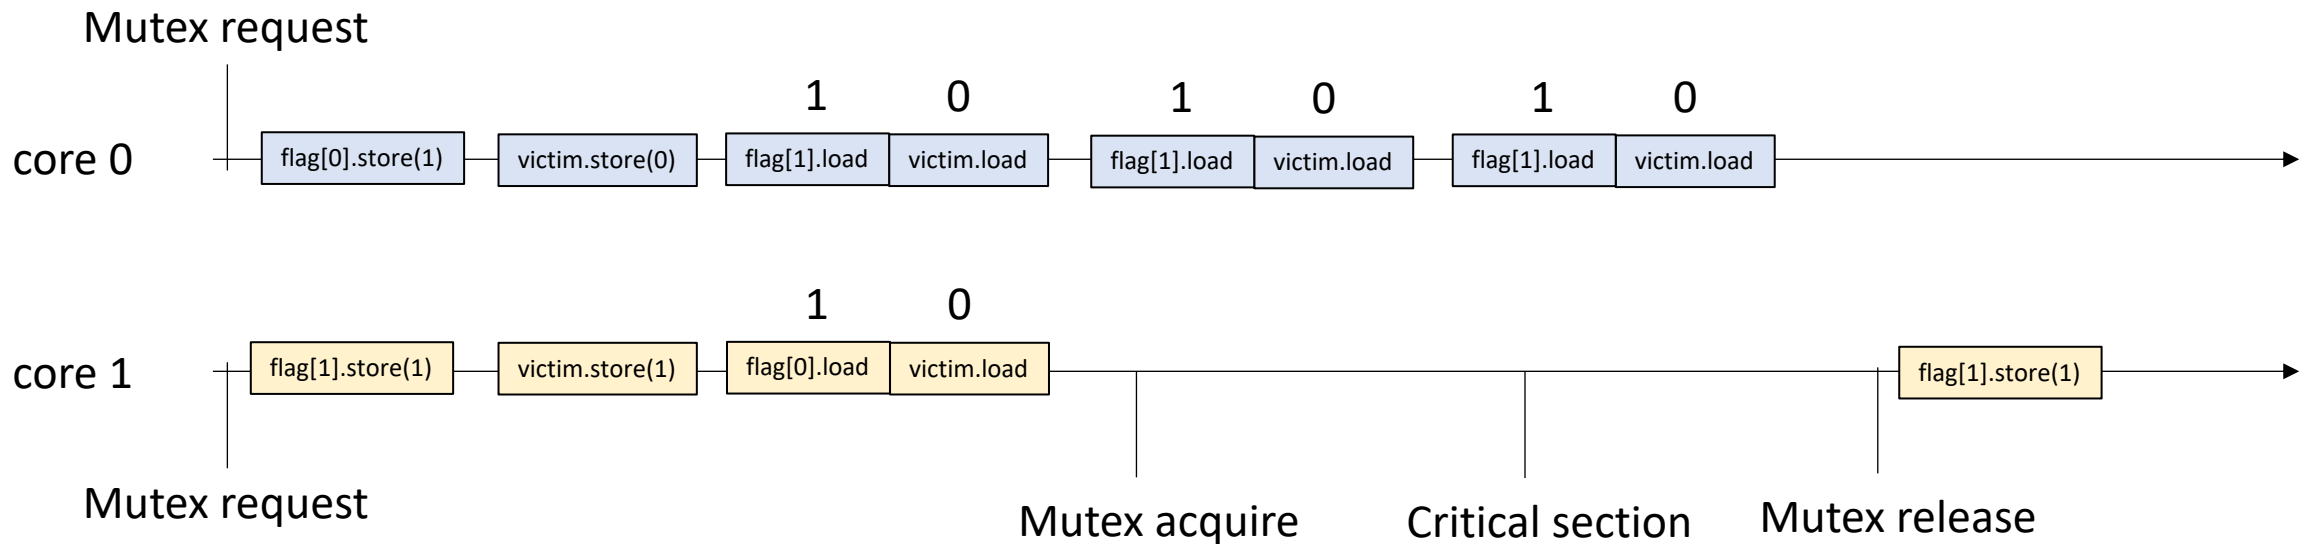
# Tie breaking with victim

```
void lock() {  
    int j = thread_id == 0 ? 1 : 0;  
    flag[thread_id].store(1);  
    victim.store(thread_id);  
    while (victim.load() == thread_id  
           && flag[j] == 1);  
}
```

```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread 0:  
`m.lock();`  
`m.unlock();`

Thread 1:  
`m.lock();`  
`m.unlock();`



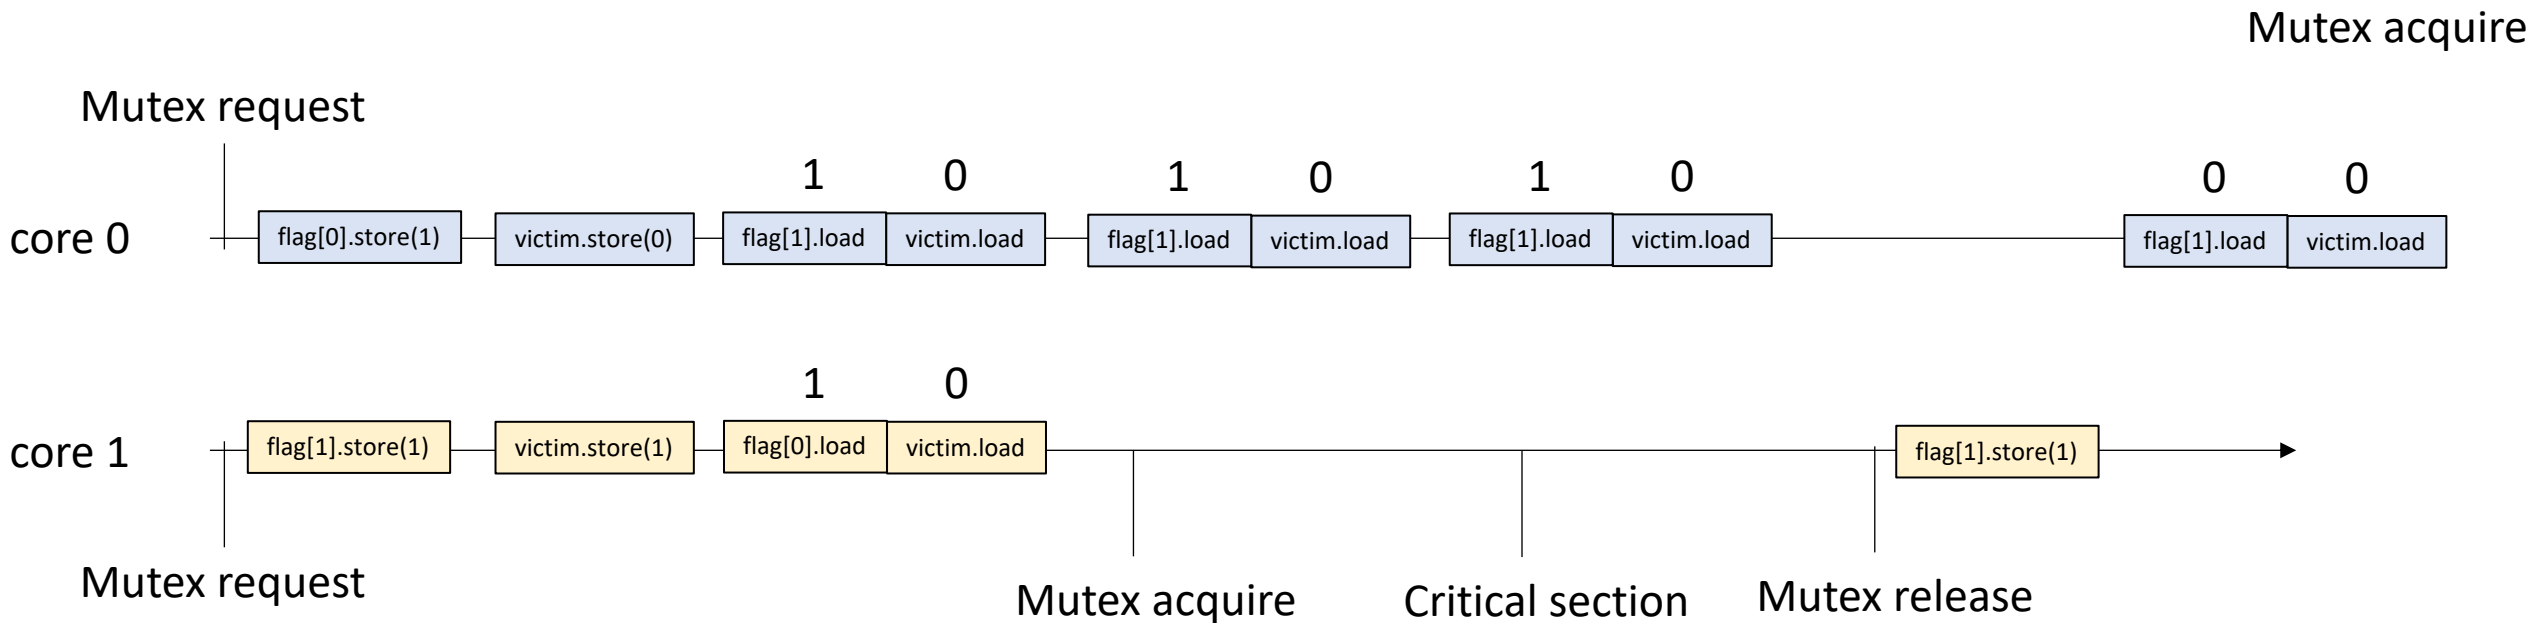
# Tie breaking with victim

```
void lock() {
    int j = thread_id == 0 ? 1 : 0;
    flag[thread_id].store(1);
    victim.store(thread_id);
    while (victim.load() == thread_id
           && flag[j] == 1);
}
```

```
void unlock() {
    int i = thread_id;
    flag[i].store(0);
}
```

Thread 0:  
 m.lock();  
 m.unlock();

Thread 1:  
 m.lock();  
 m.unlock();



# previous victim issue

```
void lock() {  
    victim.store(thread_id);  
    while (victim.load() == thread_id);  
}
```

```
void unlock() {}
```

Thread 0:


`m.lock();`

`m.unlock();`

Mutex request

*will spin forever!*

core 0





# previous flag issue

```
void lock() {  
    int j = thread_id == 0 ? 1 : 0;  
    flag[thread_id].store(1);  
    victim.store(thread_id);  
    while (victim.load() == thread_id  
           && flag[j] == 1);  
}
```

```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread 0:  
m.lock();  
m.unlock();




# previous flag issue

```
void lock() {  
    int j = thread_id == 0 ? 1 : 0;  
    flag[thread_id].store(1);  
    victim.store(thread_id);  
    while (victim.load() == thread_id  
           && flag[j] == 1);  
}
```

```
void unlock() {  
    int i = thread_id;  
    flag[i].store(0);  
}
```

Thread 0:  
m.lock();  
m.unlock();



we can enter critical section because the other thread isn't interested

# This lock satisfies the two critical properties


- Mutual exclusion
- Deadlock freedom
- *More formal proof given in the textbook*
- *How might we test it?*

# What about starvation

recall the starvation property:

*Thread 1 (yellow) requests the mutex but never gets it*

concurrent execution




time

# What about starvation

```
void lock() {  
    int j = thread_id == 0 ? 1 : 0;  
    flag[thread_id].store(1);  
    victim.store(thread_id);  
    while (victim.load() == thread_id  
           && flag[j] == 1);  
}
```

at this point, C1 is the victim and is spinning




# What about starvation

```
void lock() {  
    int j = thread_id == 0 ? 1 : 0;  
    flag[thread_id].store(1);  
    victim.store(thread_id);  
    while (victim.load() == thread_id  
           && flag[j] == 1);  
}
```

at this point, C1 is the victim and is spinning

at this point, C0 volunteers to be the victim

concurrent execution



time


# What about starvation

```
void lock() {  
    int j = thread_id == 0 ? 1 : 0;  
    flag[thread_id].store(1);  
    victim.store(thread_id);  
    while (victim.load() == thread_id  
           && flag[j] == 1);  
}
```

at this point, C1 is the victim and is spinning

at this point, C0 volunteers to be the victim

concurrent execution



time

# What about starvation


Threads take turns in petersons algorithm. It is starvation free

```
void lock() {  
    int j = thread_id == 0 ? 1 : 0;  
    flag[thread_id].store(1);  
    victim.store(thread_id);  
    while (victim.load() == thread_id  
           && flag[j] == 1);  
}
```

at this point, C1 is the victim and is spinning

at this point, C0 volunteers to be the victim

concurrent execution



time



# Mutex Implementations

Peterson only works with 2 threads.

Generalizes to the Filter Lock (Read chapter 2 in the book, part 1 of your homework!)

# Thanks!

- Next time:
  - practical mutual exclusion
- Finish homework 1 and look out for homework 2!
  - use office hours, piazza and tutors
- Do the quiz!