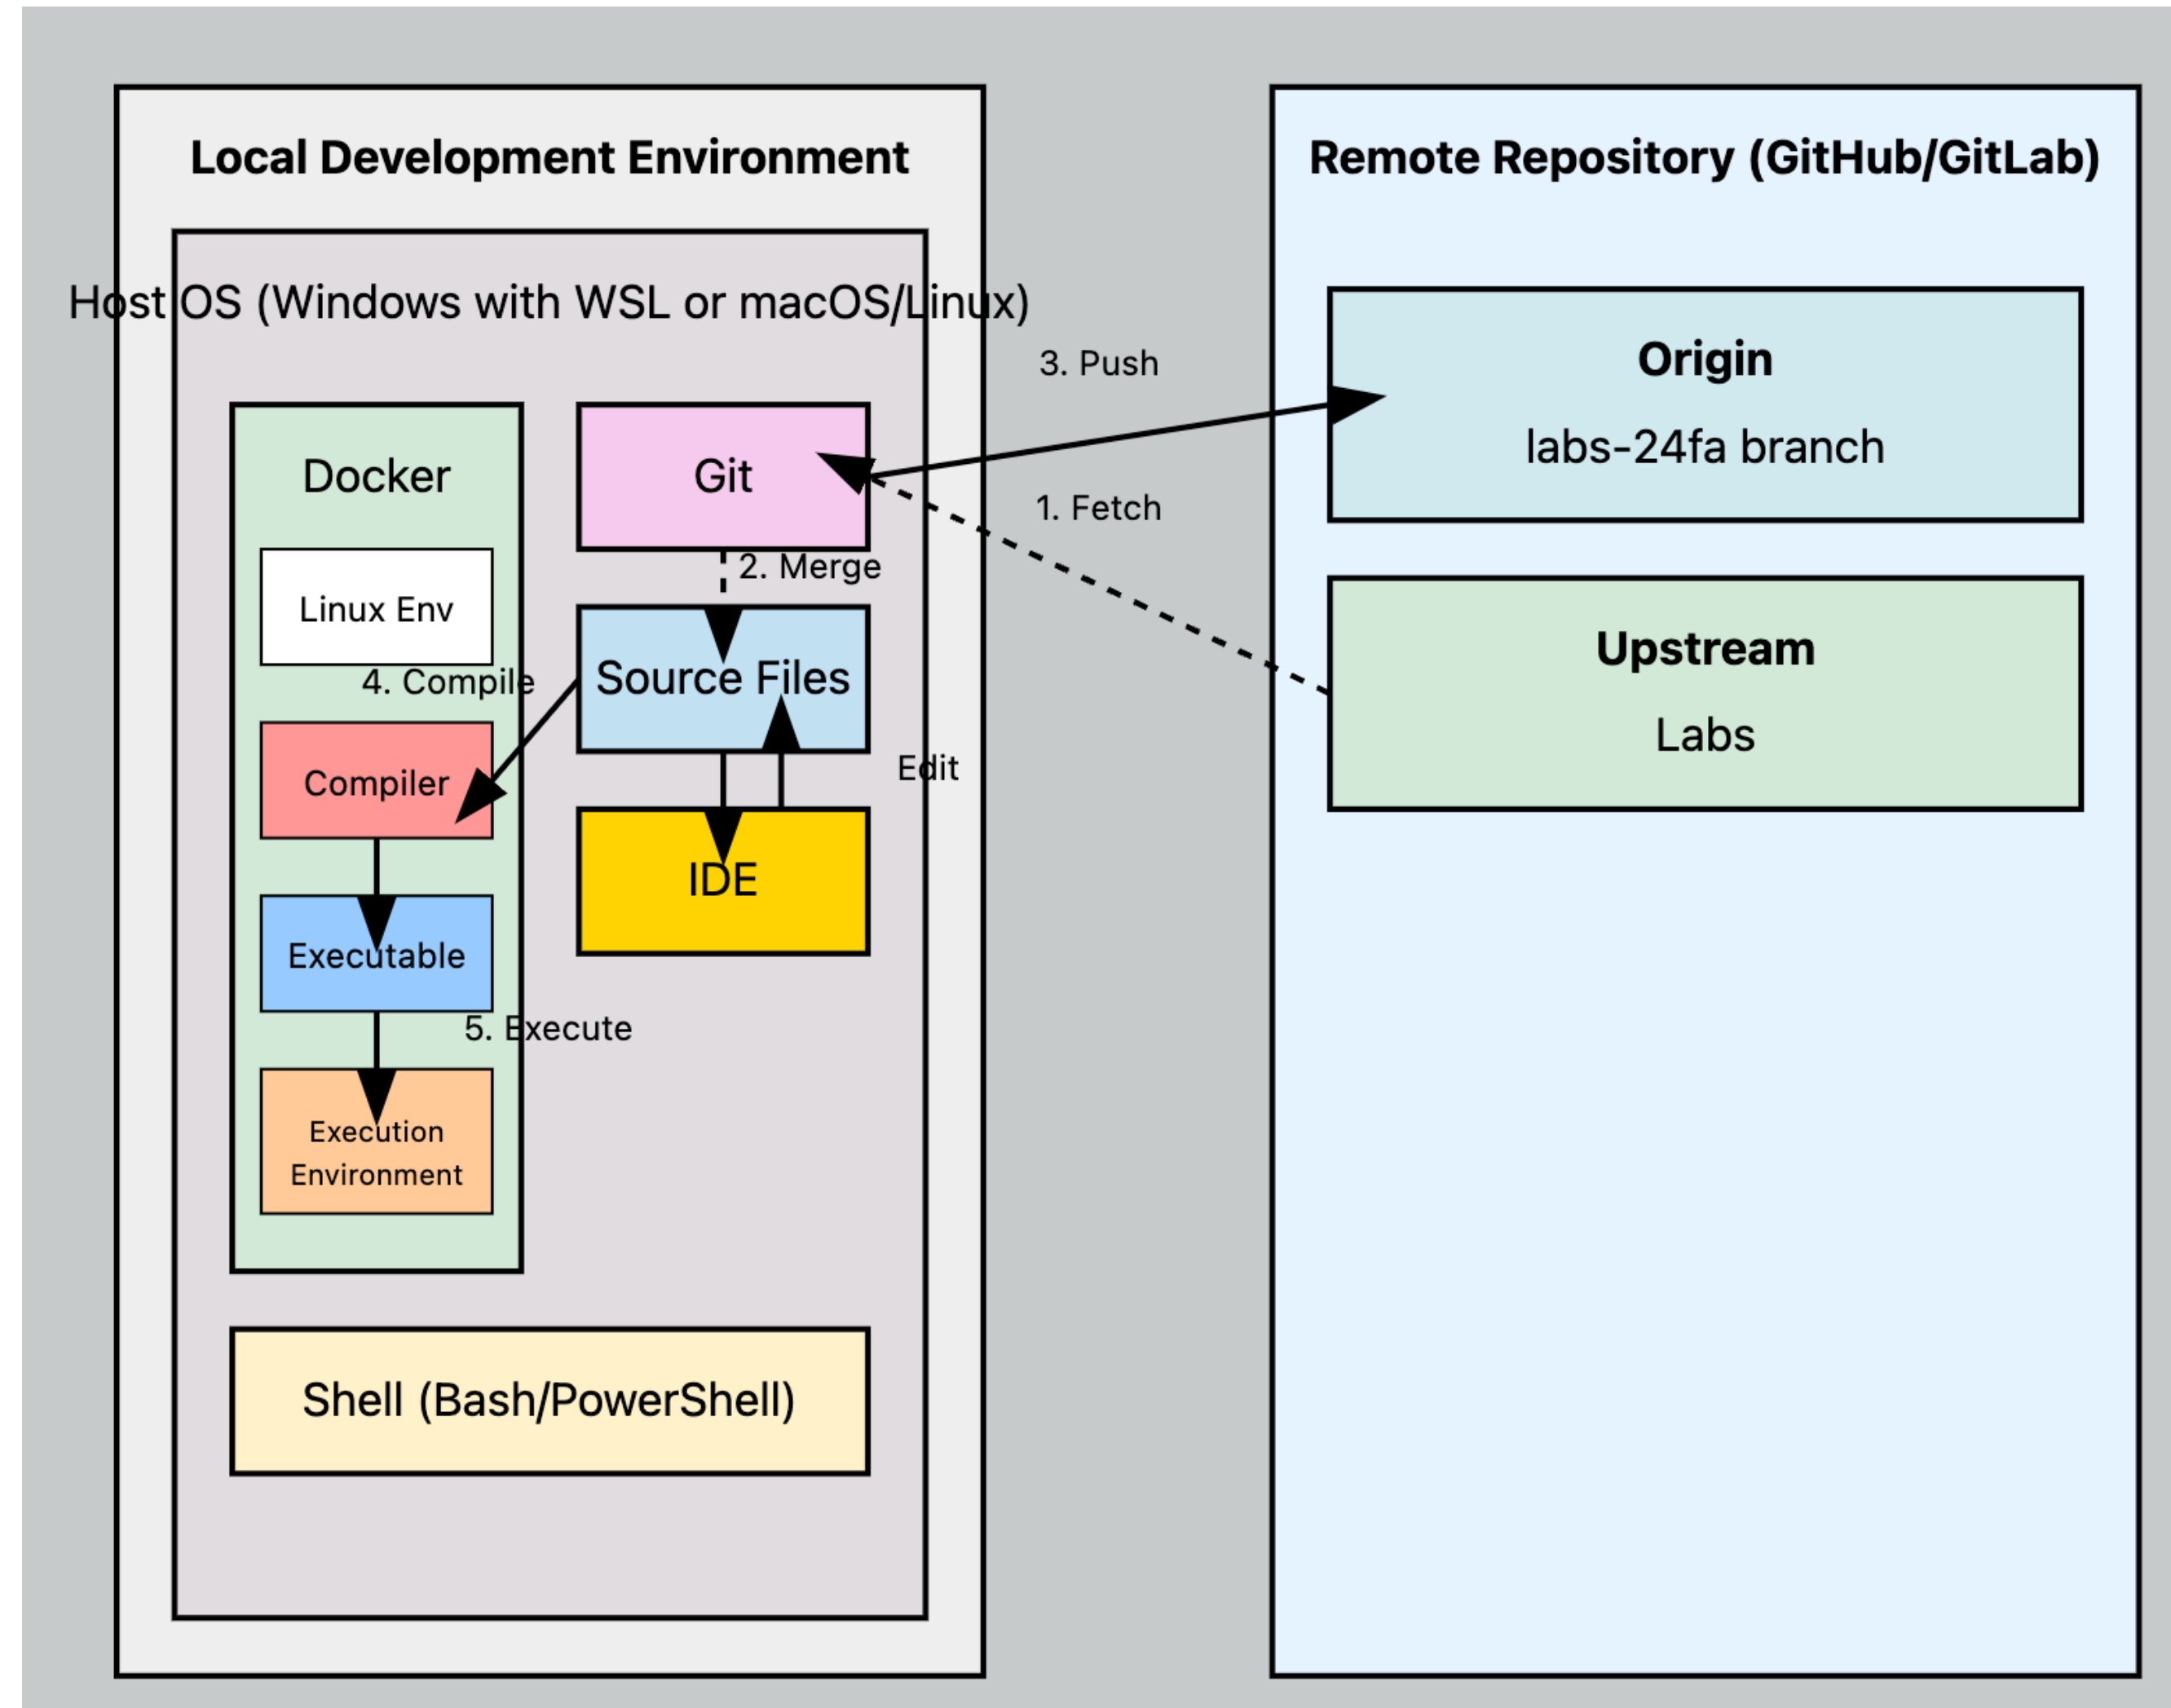


CS202 (003): Operating Systems

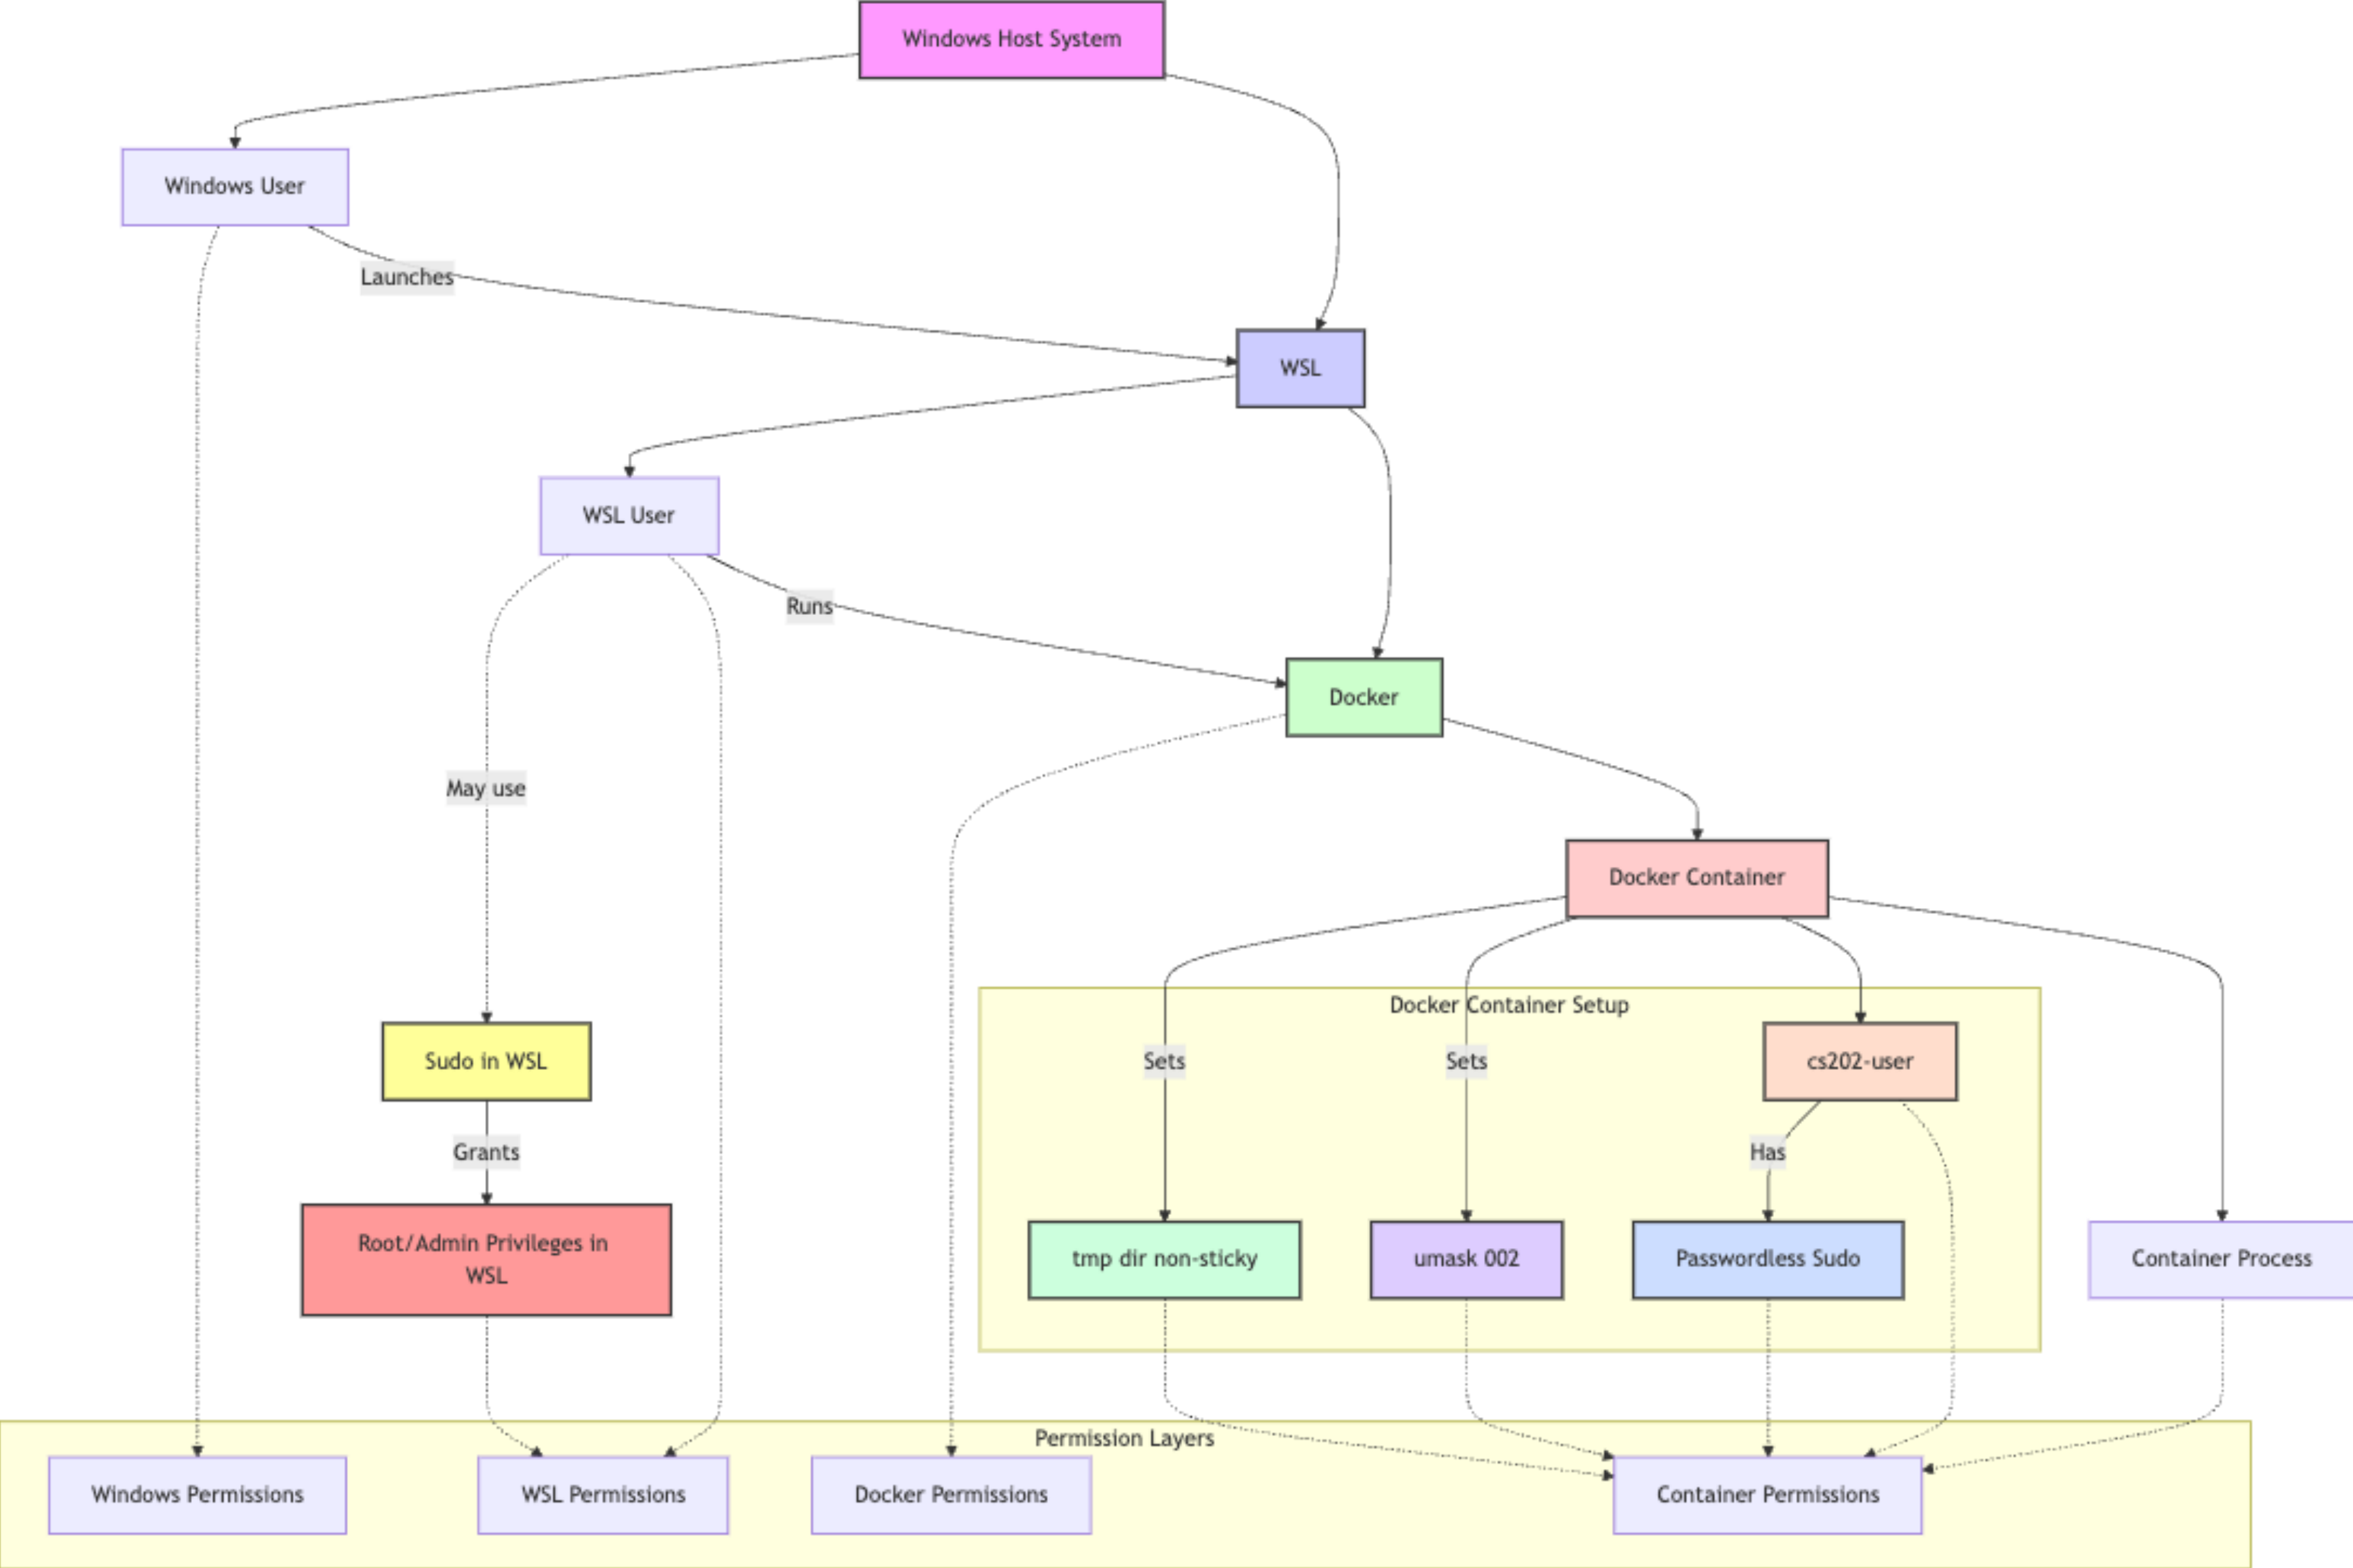
Concurrency IV

Instructor: Jocelyn Chen

Before we start...



A side note on WSL/Docker permission



What makes a good mutex implementation?

Does it provide mutual exclusion?

Does each thread get a shot at acquiring it once it is free?

What is the time overheads added by using the lock?

Spinlock implementation I

```
struct Spinlock {  
    int locked;  
}  
  
void acquire(Spinlock *lock) {  
    while (1) {  
        if (lock->locked == 0) { // A  
            lock->locked = 1; // B  
            break;  
        }  
    }  
}  
  
void release (Spinlock *lock) {  
    lock->locked = 0;  
}
```

What is the problem?

Thread 1 A
Thread 2 A
Thread 2 B
Thread 1 B

Violates mutual exclusion!

Spinlock implementation II

```
/* pseudocode */
int xchg_val(addr, value) {
    %rax = value;
    xchg (*addr), %rax
}

void acquire (Spinlock *lock) {
    pushcli(); /* what does this do? */
    while (1) {
        if (xchg_val(&lock->locked, 1) == 0)
            break;
    }
}

void release(Spinlock *lock){
    xchg_val(&lock->locked, 0);
    popcli(); /* what does this do? */
}
```



- (i) freeze all CPUs' memory activity for address addr
- (ii) temp ← *addr
- (iii) *addr ← %rax
- (iv) %rax ← temp
- (v) un-freeze memory activity

Spinlock implementation II

```
/* pseudocode */
int xchg_val(addr, value) {
    %rax = value;
    xchg (*addr), %rax
}

/* optimization in acquire;
call xchg_val() less frequently */
void acquire(Spinlock* lock) {
    pushcli();
    while (xchg_val(&lock->locked, 1) == 1) {
        while (lock->locked) ;
    }
}

void release(Spinlock *lock){
    xchg_val(&lock->locked, 0);
    popcli();
}
```

Busy waits!

Starvation!

Mutex: spinlock + a queue

```
typedef struct thread {  
    // ... Entries elided.  
    STAILQ_ENTRY(thread_t) qlink; // Tail queue entry.  
} thread_t;
```

qlink is a field that allows each thread_t structure to be part of a singly-linked tail queue. qlink field in each thread_t is what allows these threads to be linked into that queue

```
struct Mutex {  
    // Current owner, or 0 when mutex is not held.  
    thread_t *owner;  
  
    // List of threads waiting on mutex  
    STAILQ(thread_t) waiters;  
  
    // A lock protecting the internals of the mutex.  
    Spinlock splock; // as in item 1, above  
};
```

Mutex: spinlock + a queue

```
typedef struct thread {  
    // ... Entries elided.  
    // Tail queue entry.  
    STAILQ_ENTRY(thread_t) qlink;  
} thread_t;
```

```
struct Mutex {  
    // Current owner  
    //or 0 when mutex is not held.  
    thread_t *owner;  
  
    // List of threads waiting on mutex  
    STAILQ(thread_t) waiters;  
  
    // A lock protecting  
    //the internals of the mutex.  
    Spinlock splock;  
};
```

```
void mutex_acquire(struct Mutex *m) {  
  
    acquire(&m->splock);  
  
    // Check if the mutex is held;  
    // if not, current thread gets mutex and returns  
    if (m->owner == 0) {  
        m->owner = id_of_this_thread;  
        release(&m->splock);  
    } else {  
        // Add thread to waiters.  
        STAILQ_INSERT_TAIL(&m->waiters,  
                           id_of_this_thread,  
                           qlink);  
  
        // Tell the scheduler to add  
        // current thread to the list of blocked threads.  
        sched_mark_blocked(&id_of_this_thread);  
        // Unlock spinlock.  
        release(&m->splock);  
        // Stop executing until woken.  
        sched_switch();  
        // We guaranteed to hold the mutex  
        // when we are here
```

only one thread can modify the mutex's internal state at a time

this thread is waiting and shouldn't be scheduled to run

allowing other threads to access the mutex's internal state

This call switches to another thread

This is because we can get here only if context-switched-TO, which itself can happen only if this thread is removed from the waiting queue, marked "unblocked", and set to be the owner (in mutex_release() below). However, we might have held the mutex in lines 39-42 (if we were context-switched out after the spinlock release(), followed by being run as a result of another thread's release of the mutex). But if that happens, it just means that we are context-switched out an "extra" time before proceeding.

Mutex: spinlock + a queue

```
typedef struct thread {  
    // ... Entries elided.  
    // Tail queue entry.  
    STAILQ_ENTRY(thread_t) qlink;  
} thread_t;
```

```
struct Mutex {  
    // Current owner  
    //or 0 when mutex is not held.  
    thread_t *owner;  
  
    // List of threads waiting on mutex  
    STAILQ(thread_t) waiters;  
  
    // A lock protecting  
    //the internals of the mutex.  
    Spinlock splock;  
};
```

```
void mutex_release(struct Mutex *m) {  
    // Acquire the spinlock in order to make changes.  
    acquire(&m->splock);  
  
    // Assert that the current thread  
    // actually owns the mutex  
    assert(m->owner == id_of_this_thread);  
  
    // Check if anyone is waiting.  
    m->owner = STAILQ_GET_HEAD(&m->waiters);  
  
    // If so, wake them up.  
    if (m->owner) {  
        sched_wakeone(&m->owner);  
        STAILQ_REMOVE_HEAD(&m->waiters, qlink);  
    }  
  
    // Release the internal spinlock  
    release(&m->splock);  
}
```

only one thread can modify the mutex's internal state at a time

safety check to prevent a thread from releasing a mutex it doesn't own

get the first thread from the waiters queue

If there were no waiting threads, the `m->owner` would be NULL, effectively marking the mutex as unheld.

making it ready to run.

The thread is removed from the head of the waiters queue.

What makes a good mutex implementation?

Mechanism	Pros	Cons	Best Use Case
Spinlock + Queue	<ul style="list-style-type: none">- Efficient for both short and long waits- Allows context switching- Fair (FIFO ordering)- Scalable to many threads	<ul style="list-style-type: none">- More complex implementation- Slightly higher overhead for uncontended case	General-purpose locking in multi-threaded environments
Pure Spinlock	<ul style="list-style-type: none">- Very fast for short waits- Simple implementation	<ul style="list-style-type: none">- Wastes CPU cycles for long waits- Starvation and contention	Very short-duration locks with low contention
Disabling Interrupts	<ul style="list-style-type: none">- Simple to implement- Guaranteed mutual exclusion	<ul style="list-style-type: none">- Only works on single-processor systems- Can increase interrupt latency- Can't be used by user-level code	Low-level OS operations on single-processor systems
Peterson's Algorithm	<ul style="list-style-type: none">- Works without hardware support- Guaranteed fairness	<ul style="list-style-type: none">- Limited to two threads- Busy-waiting (similar to spinlock)- Can be less efficient on modern hardware	Educational purposes, simple two-thread synchronization

Deadlock

```
T1:  
acquire(mutexA);  
acquire(mutexB);  
  
// do some stuff  
  
release(mutexB);  
release(mutexA);
```

```
T2:  
acquire(mutexB);  
acquire(mutexA);  
  
// do some stuff  
  
release(mutexA);  
release(mutexB);
```

```
M:  
acquire(&mutex_m);  
n.alloc(nwanted)
```



```
acquire(&mutex_m);
```

```
N:  
  
acquire(&mutex_n)  
navailable < nwanted  
release(&mutex_n)
```

Example 1

Example 2: Code see handout