CS202 (003): Operating Systems Concurrency IV

Instructor: Jocelyn Chen

Most of the materials covered in this slide come from the lecture notes of Mike Walfish's CS202



Last Time

Peterson's Algorithm

volatile int turn;

```
flag[0] = true;
                                                flag[1] = true;
P0:
                                         P1:
                                         P1_gate: turn = 0;
PO_gate: turn = 1;
         while (flag[1] \&\& turn == 1)
                                                  while (flag[0] \&\& turn == 0)
             // busy wait
                                                      // busy wait
         // critical section
                                                 // critical section
         • • •
                                                   • • •
                                                  // end of critical section
        // end of critical section
         flag[0] = false;
                                                  flag[1] = false;
```

- expensive (busy waiting) - requires number of threads to be fixed statically - assumes sequential consistency

```
volatile bool flag[2] = {false, false};
```

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;
```

```
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
};
```

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

Mutex: spinlock + a queue

acquire(&m->splock);

```
// Check if the mutex is held;
// if not, current thread gets mutex and returns
if (m \rightarrow 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_swtch();
// We guaranteed to hold the mutex
// when we are here
```

```
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) {
```

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.

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







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;
1.
ر ک
```

acquire(&m->splock);

```
// 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);
```

```
void mutex release(struct Mutex *m) {
    // Acquire the spinlock in order to make changes.
```

```
// 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);
```

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.

Another implementation is covered in the textbook (https://pages.cs.wisc.edu/~remzi/OSTEP/threads-locks.pdf)









What makes a good mutex implementation?

Mechanism	Pros
Spinlock + Queue	 Efficient for both short and long waits Allows context switching Fair (FIFO ordering) Scalable to many threads
Pure Spinlock	- Very fast for short waits - Simple implementation
Disabling Interrupts	- Simple to implement - Guaranteed mutual exclusion
Peterson's Algorithm	- Works without hardware support - Guaranteed fairness

Cons

More complex implementation
Slightly higher overhead for uncontended case

- Wastes CPU cycles for long waits
- Starvation and contention
- Only works on single-processor systems
- Can increase interrupt latency
- Can't be used by user-level code
- Limited to two threads
- Busy-waiting (similar to spinlock)
- Can be less efficient on modern hardware

Best Use Case

General-purpose locking in multi-threaded environments

Very short-duration locks with low contention

Low-level OS operations on single-processor systems

Educational purposes, simple two-thread synchronization



T1:

- acquire(mutexA);
- acquire(mutexB);

- release(mutexB);
- release(mutexA);

T2:

acquire(mutexB); acquire(mutexA);

// do some stuff // do some stuff

- release(mutexA);
- release(mutexB);

Example 1

```
class M {
                                            class N {
    private:
                                                 private:
        Mutex mutex m;
                                                     Mutex mutex n;
        // instance of monitor N
                                                     Cond cond_n;
        N another_monitor;
                                                     int navailable;
        // Assumption: no other objects
                                                 public:
         // in the system hold a pointer
                                                     N();
        // to our "another_monitor"
                                                     ~N();
                                                     void* alloc(int nwanted);
    public:
                                                     void free(void*);
        M();
        ~M();
                                                                             N:
                                                         M:
        void methodA();
                                                  acquire(&mutex_m);
        void methodB();
                                                  n.alloc(nwanted)
};
                                                                      acquire(&mutex n)
                                                                    navailable < nwanted</pre>
                                                                      release(&mutex_n)
        Example 2: Code see handout
                                                  acquire(&mutex m);
```









Happens when **all four** conditions are present: (1) Mutual exclusion (2) Hold and wait (3) No pre-emption (4) Circular wait



Preventing deadlock

Ignore It!

Detect & Recover

Works in development, not really viable for production

Avoid Algorithmically

Negate Any of the Conditions

There are ways but we don't cover them in this class¹

Mutual exclusion put a queue for accessing resources

> Static: detect potential errors without running the code² Dynamic: detect (potential) error during/after execution³

Static/Dynamic Analysis

Check the following if you are curious: ¹Section 6.5.3 of Modern Operating Systems (Tanenbaum) ²Engler, D. and K. Ashcraft. RacerX: effective, static detection of race conditions and deadlocks. ³Savage, S., M. Burrows, G. Nelson, P. Sobalvarro, and T. Anderson. Eraser: a dynamic data race detector for multithreaded programs.

"admit defeat"

Hold and wait not likely to work

No preemption not likely to work

Circular dependency put partial order on locks (=> no cycles)





Other progress issues

Starvation

Thread waiting indefinitely (if low priority and/or resource is contended)

Priority Inversion

T1:T2:T3:(highest priority)(middle priority)(lowest priority)hold the lockhold the lockstartpreempt T3waiting for lockstart

running

Why does T2 control the CPU?

Solution 1

T1: **T2**: **T3**: (highest priority) (middle priority) (highest priority)

.

hold the lock

start

waiting for lock

finish T3

release the lock

acquire the lock

running

Solution 3

Design the code wisely so that only adjacent priority processes/threads share the lock



.

Don't handle it.



Next lecture: reading is required! (yes, we will quiz you about it at the beginning of the Thursday class)