

ICS143A: Principles of Operating Systems

Lecture 15: Locking

Anton Burtsev
November, 2017


Race conditions

- Disk driver maintains a list of outstanding requests
- Each process can add requests to the list

List implementation no locks


```
1 struct list {  
2     int data;  
3     struct list *next;  
4 };  
  
...  
  
6 struct list *list = 0;  
  
...  
  
9 insert(int data)  
10 {  
11     struct list *l;  
12  
13     l = malloc(sizeof *l);  
14     l->data = data;  
15     l->next = list;  
16     list = l;  
17 }
```

Request queue (e.g. incoming network packets)




- Linked list, list is pointer to the first element


CPU1 allocates new request






CPU2 allocates new request




CPUs 1 and 2 update next pointer




CPU1 updates head pointer



CPU2 updates head pointer



State after the race



Mutual exclusion

- Only one CPU can update list at a time

List implementation with locks

```
1 struct list {  
2     int data;  
3     struct list *next;  
4 };  
6 struct list *list = 0;  
    struct lock listlock;  
9 insert(int data)  
10 {  
11     struct list *l;  
13     l = malloc(sizeof *l);  
        acquire(&listlock);  
14     l->data = data;  
15     l->next = list;  
16     list = l;  
        release(&listlock);  
17 }
```

- Critical section

- How can we implement acquire()?

Spinlock

```
21 void
22 acquire(struct spinlock *lk)
23 {
24     for(;;) {
25         if(!lk->locked) {
26             lk->locked = 1;
27             break;
28         }
29     }
30 }
```

- Spin until lock is 0
- Set it to 1

Still incorrect

```
21 void  
22 acquire(struct spinlock *lk)  
23 {  
24     for(;;) {  
25         if(!lk->locked) {  
26             lk->locked = 1;  
27             break;  
28         }  
29     }  
30 }
```

- Two CPUs can reach line #25 at the same time
 - See not locked, and
 - Acquire the lock
- Lines #25 and #26 need to be atomic
 - I.e. indivisible

Compare and swap: xchg

- Swap a word in memory with a new value
 - Return old value

Correct implementation

```
1573 void
1574 acquire(struct spinlock *lk)
1575 {
...
1580     // The xchg is atomic.
1581     while(xchg(&lk->locked, 1) != 0)
1582         ;
...
1592 }
```

xchgl instruction




```
0568 static inline uint
0569 xchg(volatile uint *addr, uint newval)
0570 {
0571     uint result;
0572
0573     // The + in "+m" denotes a read-modify-write
0574     // operand.
0575
0576     asm volatile("lock; xchgl %0, %1" :
0577                 "+m" (*addr), "=a" (result) :
0578                 "1" (newval) :
0579                 "cc");
0580
0581     return result;
0582 }
```

Correct implementation

```
1573 void
1574 acquire(struct spinlock *lk)
1575 {
...
1580     // The xchg is atomic.
1581     while(xchg(&lk->locked, 1) != 0)
1582         ;
1584     // Tell the C compiler and the processor to not move loads or
1585     // stores
1586     // past this point, to ensure that the critical section's memory
1587     // references happen after the lock is acquired.
1588     __sync_synchronize();
...
1592 }
```

Deadlocks


Deadlocks



Lock ordering

- Locks need to be acquired in the same order

Locks and interrupts



Locks and interrupts

- Never hold a lock with interrupts enabled

```
1573 void
1574 acquire(struct spinlock *lk)
1575 {
1576     pushcli(); // disable interrupts to avoid deadlock.
1577     if(holding(lk))
1578         panic("acquire");
1579     // The xchg is atomic.
1580     while(xchg(&lk->locked, 1) != 0)
1581         ;
1582     ...
1583     __sync_synchronize();
1584     ...
1592 }
```

Disabling interrupts

Simple disable/enable is not enough

- If two locks are acquired
- Interrupts should be re-enabled only after the second lock is released
- Pushcli() uses a counter

```
1655 pushcli(void)
1656 {
1657     int eflags;
1658
1659     eflags = readeflags();
1660     cli();
1661     if(cpu->ncli == 0)
1662         cpu->intena = eflags & FL_IF;
1663     cpu->ncli += 1;
1664 }
...
1667 popcli(void)
1668 {
1669     if(readeflags()&FL_IF)
1670         panic("popcli - interruptible");
1671     if(--cpu->ncli < 0)
1672         panic("popcli");
1673     if(cpu->ncli == 0 && cpu->intena)
1674         sti();
1675 }
```

Pushcli()/popcli()

Thank you!