

# CS5460/6460: Operating Systems

## Lecture 11: Locking

Anton Burtsev  
February, 2014


# 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 };
5
6 ...
7
8
9 struct list *list = 0;
10
11 ...
12
13 insert(int data)
14 {
15     struct list *l;
16
17     l = malloc(sizeof *l);
18     l->data = data;
19     l->next = list;
20     list = l;
21 }
```

# Request queue (e.g. incoming network packets)




- Linked list, list is pointer to the first element


# CPU1 allocates new request




# CPU2 allocates new request



# CPU 1 and 2 update next pointer




# CPU1 updates head pointer






# CPU2 updates head pointer



# State after the race



# List implementation with locks

```
1 struct list {
2     int data;
3     struct list *next;
4 };
5
6 struct list *list = 0;
7     struct lock listlock;
8
9 insert(int data)
10 {
11     struct list *l;
12
13     l = malloc(sizeof *l);
14     acquire(&listlock);
15     l->data = data;
16     l->next = list;
17     list = l;
18     release(&listlock);
19 }
20 }
```

# Spinlock

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

# 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

- We switch between processes now

# Correct implementation


```
1473 void
1474 acquire(struct spinlock *lk)
1475 {
...
1480 // The xchg is atomic.
1481 // It also serializes, so that reads after acquire are not
1482 // reordered before it.
1483 while(xchg(&lk->locked, 1) != 0)
1484 ;
1485
...
1489 }
```

# Compare and swap

```
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
        operand.
0574     asm volatile("lock; xchgl %0, %1" :
0575                 "+m" (*addr), "=a" (result) :
0576                 "1" (newval) :
0577                 "cc");
0578     return result;
0579 }
```





# 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

# Disabling interrupts

```
1473 void
1474 acquire(struct spinlock *lk)
1475 {
1476     pushcli(); // disable interrupts to avoid deadlock.
...
1480     // The xchg is atomic.
1481     // It also serializes, so that reads after acquire
are not
1482     // reordered before it.
1483     while(xchg(&lk->locked, 1) != 0)
1484     ;
...
1489 }
```

# 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

# pushcli()

```
1554 void
1555 pushcli(void)
1556 {
1557     int eflags;
1558
1559     eflags = readeflags();
1560     cli();
1561     if(cpu->ncli++ == 0)
1562         cpu->intena = eflags & FL_IF;
1563 }
```

```
1565 void
```

# popcli()

```
1566 popcli(void)
```

```
1567 {
```

```
...
```

```
1570     if(--cpu->ncli < 0)
```

```
1571         panic("popcli");
```

```
1572     if(cpu->ncli == 0 && cpu->intena)
```

```
1573         sti();
```

```
1574 }
```



# Problems with locks

# Problems with locks

- Deadlock
  - Locks break modularity of interfaces, easy to get wrong
- Priority inversion
  - Low-priority task holds a lock required by a higher priority task
  - Priority inheritance can be a solution, but can also result in errors (see What really happened on Mars)

# Problems with locks

- Convoying
  - Several tasks need the locks in roughly the same order
  - One slow task acquires the lock first
  - Everyone slows to the speed of this slow task
- Signal safety
  - Similar to interrupts, but for user processes
  - Can't be disabled, thus can't use locks
  -

# Problems with locks

- Kill safety
  - What if a task is killed or crashed while holding a lock?
- Preemption safety
  - What happens if a task is preempted while holding a lock?

# Optimistic concurrency

# Optimistic concurrency: main idea

- Instead of acquiring a lock try updating a data structure
  - When done, try committing changes
  - If there is a conflict, retry
- Similar to database transactions

# Example: lock-free stack(), aka FIFO queue

```
class Node {  
    Node * next;  
    int data;  
};  
  
// 'head of list'  
Node * head;
```

# Lock-free push()

```
void push(int t) {  
    Node* node = new Node(t);  
    do {  
        node->next = head;  
    } while (!cas(&head, node, node->next));  
}
```



# Lock-free pop()

```
bool pop(int& t) {  
    Node* current = head;  
    while(current) {  
        if(cas(&head, current->next, current)) {  
            t = current->data;  
            return true;  
        }  
        current = head;  
    }  
    return false;  
}
```

# The ABA problem

- The value of a variable is changed from A to B and then back to A
- In our example the variable is a pointer to a stack element
- What if the head gets deallocated with `free()`, and allocated again?
  - There is a good chance that head will have the same pointer value
    - Memory allocators often choose recently deallocated values
  - But really this is a different stack element

# ABA example

Thread 1: pop()  
read A from head  
store A.next `somewhere`

Thread 2:

pop()  
Pops A, discards it  
First element becomes B  
pop(): pops B  
push():  
Memory manager recycles  
A to hold a new variable

cas with A succeeds

# ABA workaround

- Keep an `update counter' along with a pointer
  - Needs a double word CAS
- Don't recycle memory too soon

# Nontrivial lock-free data structures

- For example, a linked list
  - Much more complex
    - Operations on two pointers
  - Insert
    - What if predecessor is removed?

Thank you!