

1 CS 202, Spring 2015
2 Handout 5 (Class 6)

3
4 Implementation of spinlocks and mutexes

5
6 1. Here is a BROKEN spinlock implementation:

```
7
8     struct Lock {
9         int locked;
10    }
11
12    void acquire(Lock *lock) {
13        while (1) {
14            if (lock->locked == 0) { // A
15                lock->locked = 1;    // B
16                break;
17            }
18        }
19    }
20
```

```
21    void release (Lock *lock) {
22        lock->locked = 0;
23    }
24
```

25 What's the problem? Two acquire()s on the same lock on different
26 CPUs might both execute line A, and then both execute B. Then
27 both will think they have acquired the lock. Both will proceed.
28 That doesn't provide mutual exclusion.
29

29

30 2. Correct spinlock implementation

31

32 Relies on atomic hardware instruction. For example, on the x86,
33 doing
34 "xchg addr, %eax"
35 does the following:

36

```
37     (i) freeze all CPUs' memory activity for address addr
38     (ii) temp = *addr
39     (iii) *addr = %eax
40     (iv) %eax = temp
41     (v) un-freeze memory activity
42
```

43 /* pseudocode */

```
44 int xchg_val(addr, value) {
45     %eax = value;
46     xchg (*addr), %eax
47 }
48
```

49 /* bare-bones version of acquire */

```
50 void acquire (Lock *lock) {
51     pushcli(); /* what does this do? */
52     while (1) {
53         if (xchg_val(&lock->locked, 1) == 0)
54             break;
55     }
56 }
57
```

```
58 void release(Lock *lock){
59     xchg_val(&lock->locked, 0);
60     popcli(); /* what does this do? */
61 }
62
```

63 /* optimization in acquire; call xchg_val() less frequently */

```
64 void acquire(Lock* lock) {
65     pushcli();
66     while (xchg_val(&lock->locked, 1) == 1) {
67         while (lock->locked) ;
68     }
69 }
70
```

71
72 The above is called a *spinlock* because acquire() spins. The
73 bare-bones version is called a "test-and-set (TAS) spinlock"; the
74 other is called a "test-and-test-and-set spinlock".
75

76 The spinlock above is great for some things, not so great for
77 others. The main problem is that it *busy waits*: it spins,
78 chewing up CPU cycles. Sometimes this is what we want (e.g., if
79 the cost of going to sleep is greater than the cost of spinning
80 for a few cycles waiting for another thread or process to
81 relinquish the spinlock). But sometimes this is not at all what we
82 want (e.g., if the lock would be held for a while: in those
83 cases, the CPU waiting for the lock would waste cycles spinning
84 instead of running some other thread or process).
85

86 NOTE: the spinlocks presented here can introduce performance issues
87 when there is a lot of contention. (This happens even if the
88 programmer is using spinlocks correctly.) The performance issues
89 result from cross-talk among CPUs (which undermines caching and
90 generates traffic on the memory bus). If we have time later, we will
91 study a remediation of this issue (search the Web for "MCS locks").
92

93 ANOTHER NOTE: In everyday application-level programming, spinlocks
94 will not be something you use (use mutexes instead). But you should
95 know what these are for technical literacy, and to see where the
96 mutual exclusion is truly enforced on modern hardware.
97

```

98 3. Mutex implementation
99
100 The intent of a mutex is to avoid busy waiting: if the lock is not
101 available, the locking thread is put to sleep, and tracked by a
102 queue in the mutex.
103
104 struct Mutex {
105     bool is_held;           /* true if mutex held */
106     thread_id owner;       /* thread holding mutex, if locked */
107     thread_list waiters;   /* queue of thread TCBS */
108     Lock wait_lock;       /* as in item 2, above */
109 }
110
111 The implementation of mutex_acquire() and mutex_release() would
112 be something like:
113
114 void mutex_acquire(Mutex *m) {
115     acquire(&m->wait_lock); /* we spin to acquire wait_lock */
116     while (m->is_held) {    /* someone else has the mutex */
117         m->waiters.insert(current_thread)
118         release(&m->wait_lock);
119
120         /*
121          * NOTE! Right here, mutex_release() could execute. To
122          * avoid "losing the wakeup", we check whether we are
123          * on the scheduler's ready list. If we are, we
124          * shouldn't yield().
125          */
126         yield_if_we_are_not_ready();
127
128         acquire(&m->wait_lock); /* we spin again */
129         m->waiters.remove(current_thread)
130     }
131     m->is_held = true;      /* we now hold the mutex */
132     m->owner = self;
133     release(&m->wait_lock);
134 }
135
136 void mutex_release(Mutex *m) {
137     acquire(&m->wait_lock); /* we spin to acquire wait_lock */
138     m->is_held = false;
139     m->owner = 0;
140
141     /* tell scheduler to run a waiter */
142     place_a_waiter_on_ready_list(m->waiters);
143     release(&m->wait_lock);
144 }
145
146
147
148
149
150
151
152
153
154
155
156

```