

```

1 Handout for CS 439
2 Class 10
3 14 February 2013

```

1. How can we implement locks, acquire(), and release()?

1a. Here is A BADLY BROKEN implementation:

```

8
9 struct Lock {
10     int locked;
11 }
12
13 void [BROKEN] acquire(Lock *lock) {
14     while (1) {
15         if (lock->locked == 0) { // C
16             lock->locked = 1;    // D
17             break;
18         }
19     }
20 }
21
22 void release (Lock *lock) {
23     lock->locked = 0;
24 }

```

25
26 What's the problem? Two acquire()s on the same lock on different
27 CPUs might both execute line C, and then both execute D. Then
28 both will think they have acquired the lock. This is the same
29 kind of race we were trying to eliminate to begin with. But we
30 have made a little progress: now we only need a way to prevent
31 interleaving in one place (acquire()), not for many arbitrary
32 complex sequences of code.
33

```

34
35 1b. Here's a way that is correct but that is appropriate only in
36 some circumstances:
37
38 Use an atomic instruction on the CPU. For example, on the x86,
39 doing
40     "xchg addr, %eax"
41 does the following:

```

```

42
43 (i) freeze all CPUs' memory activity for address addr
44 (ii) temp = *addr
45 (iii) *addr = %eax
46 (iv) %eax = temp
47 (v) un-freeze memory activity

```

```

48
49 /* pseudocode */
50 int xchg_val(addr, value) {
51     %eax = value;
52     xchg (*addr), %eax
53 }

```

```

54
55 struct Lock {
56     int locked;
57 }

```

```

58
59 /* bare-bones version of acquire */
60 void acquire (Lock *lock) {
61     pushcli(); /* what does this do? */
62     while (1) {
63         if (xchg_val(&lock->locked, 1) == 0)
64             break;
65     }
66 }
67
68 /* optimization in acquire; call xchg_val() less frequently */
69 void acquire(Lock* lock) {
70     pushcli();
71     while (xchg_val(&lock->locked, 1) == 1) {
72         while (lock->locked) ;
73     }
74 }
75
76 void release(Lock *lock){
77     xchg_val(&lock->locked, 0);
78     popcli(); /* what does this do? */
79 }

```

80
81 The above is called a *spinlock* because acquire() spins.

82
83 The spinlock above is great for some things, not so great for
84 others. The main problem is that it *busy waits*: it spins,
85 chewing up CPU cycles. Sometimes this is what we want (e.g., if
86 the cost of going to sleep is greater than the cost of spinning
87 for a few cycles waiting for another thread or process to
88 relinquish the spinlock). But sometimes this is not at all what we
89 want (e.g., if the lock would be held for a while: in those
90 cases, the CPU waiting for the lock would waste cycles spinning
91 instead of running some other thread or process).
92

```

93
94 1c. Here's an object that does not involve busy waiting. Note: the
95 "threads" here can be user-level threads, kernel threads, or
96 threads-inside-kernel. The concept is the same in all cases.
97
98 struct Mutex {
99     bool is_held;           /* true if mutex held */
100     thread_id owner;       /* thread holding mutex, if locked */
101     thread_list waiters;   /* queue of thread TCBS */
102     Lock wait_lock;       /* as in 1b */
103 }
104
105 The implementation of acquire() and release() would be something like:
106
107 void mutex_acquire(Mutex *m) {
108
109     acquire(&m->wait_lock); /* we spin to acquire wait_lock */
110     while (m->is_held) { /* someone else has the mutex */
111         m->waiters.insert(current_thread)
112         release(&m->wait_lock);
113         schedule(); /* run a thread that is on the ready list */
114         acquire(&m->wait_lock); /* we spin again */
115     }
116     m->is_held = true; /* we now hold the mutex */
117     m->owner = self;
118     release(&m->wait_lock);
119 }
120
121 void mutex_release(Mutex *m) {
122
123     acquire(&m->wait_lock); /* we spin to acquire wait_lock */
124     m->is_held = false;
125     m->owner = 0;
126     wake_up_a_waiter(m->waiters); /* select and run a waiter */
127     release(&m->wait_lock);
128
129 }
130
131 [Please let me (MW) know if you see bugs in the above.]
132

```

```

133
134
135 2. NOTE: the above mutex does the right thing only if there are some
136 constraints on the order in which the CPU carries out memory reads and
137 writes. For example, if operations _after_mutex_acquire() in program
138 order appear to another processor to happen in the opposite order, then
139 they would not be protected by the lock, and the program would be incorrect.
140
141 How do we get the required guarantee? By ensuring that neither the
142 compiler nor the processor reorders instructions with respect to the
143 acquire(). To prevent reordering by the compiler, the programmer can
144 mark the asm instructions as volatile. To prevent the processor from
145 reordering, one must use special assembly instructions. For instance,
146 fences (the "LFENCE", "SFENCE", and "MFENCE" instructions) tell the CPU
147 not to re-order memory operations past the fences. Also, the processor
148 will not reorder instructions with respect to xchg() (and a few others).
149
150 Moral of the above paragraphs: if you're implementing a concurrency
151 primitive, read the processor's documentation about how loads and stores
152 get sequenced, and how to enforce that the compiler *and* the processor
153 follow program order.
154
155
156
157 3. Terminology
158
159 To avoid confusion, we will use the following terminology in this
160 course (you will hear other terminology elsewhere):
161
162 --A "lock" is an abstract object that provides mutual exclusion
163
164 --A "spinlock" is a lock that works by busy waiting, as in 1b
165
166 --A "mutex" is a lock that works by having a "waiting" queue and
167 then protecting that waiting queue with atomic hardware
168 instructions, as in 2c. The most natural way to "use the hardware"
169 is with a spinlock, but there are others, such as turning off
170 interrupts, which works if we're on a single CPU machine.

```