

Feb 12, 13 15:53

I09-handout.txt

Page 1/4

```

1 Handout for CS 439
2 Class 9
3 12 February 2013
4
5 1. Implementing threads
6
7     Per-thread state in thread control block:
8
9     typedef struct tcb {
10         unsigned long esp;    /* Stack pointer of thread */
11         char *t_stack;       /* Bottom of thread's stack */
12         /* ... */
13     };
14
15     Machine-dependent thread-switch function:
16
17     void swtch(tcb *current, tcb *next);
18
19     Machine-dependent thread initialization function:
20
21     void thread_init(tcb *t, void (*fn) (void *), void *arg);
22
23     Implementation of swtch(current, next):
24
25     pushl %ebp; movl %esp, %ebp    # Save frame pointer
26     pushl %ebx; pushl %esi; pushl %edi # Save callee-saved regs
27
28     movl 8(%ebp), %edx             # %edx = current
29     movl 12(%ebp), %eax           # %eax = next
30     movl %esp, (%edx)             # %edx->esp = %esp
31     movl (%eax), %esp            # %esp = %eax->esp
32
33     popl %edi; popl %esi; popl %ebx # Restore callee saved regs
34     popl %ebp                    # Restore frame pointer
35     ret                          # Resume execution
36
37
38 [thanks to David Mazieres]
39

```

Feb 12, 13 15:53

I09-handout.txt

Page 2/4

```

40
41 2. How can we implement locks, acquire(), and release()?
42
43     2a. Here is A BADLY BROKEN implementation:
44
45     struct Lock {
46         int locked;
47     }
48
49     void [BROKEN] acquire(Lock *lock) {
50         while (1) {
51             if (lock->locked == 0) { // C
52                 lock->locked = 1;    // D
53                 break;
54             }
55         }
56     }
57
58     void release (Lock *lock) {
59         lock->locked = 0;
60     }
61
62     What's the problem? Two acquire()s on the same lock on different
63     CPUs might both execute line C, and then both execute D. Then
64     both will think they have acquired the lock. This is the same
65     kind of race we were trying to eliminate to begin with. But we
66     have made a little progress: now we only need a way to prevent
67     interleaving in one place (acquire()), not for many arbitrary
68     complex sequences of code.
69

```

```

70
71 2b. Here's a way that is correct but that is appropriate only in
72 some circumstances:
73
74 Use an atomic instruction on the CPU. For example, on the x86,
75 doing
76     "xchg addr, %eax"
77 does the following:
78
79 (i) freeze all CPUs' memory activity for address addr
80 (ii) temp = *addr
81 (iii) *addr = %eax
82 (iv) %eax = temp
83 (v) un-freeze memory activity
84
85 /* pseudocode */
86 int xchg_val(addr, value) {
87     %eax = value;
88     xchg (*addr), %eax
89 }
90
91 struct Lock {
92     int locked;
93 }
94
95 /* bare-bones version of acquire */
96 void acquire (Lock *lock) {
97     pushcli(); /* what does this do? */
98     while (1) {
99         if (xchg_val(&lock->locked, 1) == 0)
100             break;
101     }
102 }
103
104 /* optimization in acquire; call xchg_val() less frequently */
105 void acquire(Lock* lock) {
106     pushcli();
107     while (xchg_val(&lock->locked, 1) == 1) {
108         while (lock->locked) ;
109     }
110 }
111
112 void release(Lock *lock){
113     xchg_val(&lock->locked, 0);
114     popcli(); /* what does this do? */
115 }
116
117 The above is called a *spinlock* because acquire() spins.
118
119 The spinlock above is great for some things, not so great for
120 others. The main problem is that it *busy waits*: it spins,
121 chewing up CPU cycles. Sometimes this is what we want (e.g., if
122 the cost of going to sleep is greater than the cost of spinning
123 for a few cycles waiting for another thread or process to
124 relinquish the spinlock). But sometimes this is not at all what we
125 want (e.g., if the lock would be held for a while: in those
126 cases, the CPU waiting for the lock would waste cycles spinning
127 instead of running some other thread or process).
128

```

```

129
130 2c. Here's an object that does not involve busy waiting. Note: the
131 "threads" here can be user-level threads, kernel threads, or
132 threads-inside-kernel. The concept is the same in all cases.
133
134 struct Mutex {
135     bool is_held; /* true if mutex held */
136     thread_id owner; /* thread holding mutex, if locked */
137     thread_list waiters; /* queue of thread TCBS */
138     Lock wait_lock; /* as in 2b */
139 }
140
141 The implementation of acquire() and release() would be something like:
142
143 void mutex_acquire(Mutex *m) {
144
145     acquire(&m->wait_lock); /* we spin to acquire wait_lock */
146     while (m->is_held) { /* someone else has the mutex */
147         m->waiters.insert(current_thread)
148         release(&m->wait_lock);
149         schedule(); /* run a thread that is on the ready list */
150         acquire(&m->wait_lock); /* we spin again */
151     }
152     m->is_held = true; /* we now hold the mutex */
153     m->owner = self;
154     release(&m->wait_lock);
155 }
156
157 void mutex_release(Mutex *m) {
158
159     acquire(&m->wait_lock); /* we spin to acquire wait_lock */
160     m->is_held = false;
161     m->owner = 0;
162     wake_up_a_waiter(m->waiters); /* select and run a waiter */
163     release(&m->wait_lock);
164 }
165
166 [Please let me (MW) know if you see bugs in the above.]
167
168 3. NOTE: the above mutex does the right thing only if there are some
169 constraints on the order in which the CPU carries out memory reads and
170 writes. For example, if operations _after_ mutex_acquire() in program
171 order appear to another processor to happen in the opposite order, then
172 they would not be protected by the lock, and the program would be incorrect.
173
174 How do we get the required guarantee? By ensuring that neither the
175 compiler nor the processor reorders instructions with respect to the
176 acquire(). To prevent reordering by the compiler, the programmer can
177 mark the asm instructions as volatile. To prevent the processor from
178 reordering, one must use special assembly instructions. For instance,
179 fences (the "LFENCE", "SFENCE", and "MFENCE" instructions) tell the CPU
180 not to re-order memory operations past the fences. Also, the processor
181 will not reorder instructions with respect to xchg() (and a few others).
182
183 Moral of the above paragraphs: if you're implementing a concurrency
184 primitive, read the processor's documentation about how loads and stores
185 get sequenced, and how to enforce that the compiler *and* the processor
186 follow program order.
187
188 4. Terminology
189
190 To avoid confusion, we will use the following terminology in this
191 course (you will hear other terminology elsewhere):
192
193 --A "lock" is an abstract object that provides mutual exclusion
194
195 --A "spinlock" is a lock that works by busy waiting, as in 2b
196
197 --A "mutex" is a lock that works by having a "waiting" queue and
198 then protecting that waiting queue with atomic hardware
199 instructions, as in 2c. The most natural way to "use the hardware"
200 is with a spinlock, but there are others, such as turning off
201 interrupts, which works if we're on a single CPU machine.

```