

Feb 14, 11 23:23

I09-handout.txt

Page 1/8

```

1 Handout for CS 372H
2 Class 9
3 15 February 2011
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 14, 11 23:23

I09-handout.txt

Page 2/8

```

40
41
42 2. Example to illustrate interleavings: say that thread A executes f()
43 and thread B executes g(). (Here, we are using the term "thread"
44 abstractly. This example applies to any of the approaches that fall
45 under the word "thread".)
46
47     a.
48
49         int x;
50
51         f() { x = 1; }
52
53         g() { x = 2; }
54
55         What are possible values of x after A has executed f() and B has
56         executed g()?
57
58     b.
59
60         int y = 12;
61
62         f() { x = y + 1; }
63         g() { y = y * 2; }
64
65         What are the possible values of x?
66
67     c.
68
69         int x = 0;
70         f() { x = x + 1; }
71         g() { x = x + 2; }
72
73         What are the possible values of x?
74
75 3. Linked list example
76
77     struct List_elem {
78         int data;
79         struct List_elem* next;
80     };
81
82     List_elem* head = 0;
83
84     insert(int data) {
85         List_elem* l = new List_elem;
86         l->data = data;
87         l->next = head;
88         head = l;
89     }
90
91     What happens if two threads execute insert() at once and we get the
92     following interleaving?
93
94     thread 1: l->next = head
95     thread 2: l->next = head
96     thread 2: head = l;
97     thread 1: head = l;
98

```

Feb 14, 11 23:23

I09-handout.txt

Page 3/8

```

97
98
99 4. Producer/consumer example:
100
101 /*
102 "buffer" stores BUFFER_SIZE items
103 "count" is number of used slots. a variable that lives in memory
104 "out" is next empty buffer slot to fill (if any)
105 "in" is oldest filled slot to consume (if any)
106 */
107
108 void producer (void *ignored) {
109     for (;;) {
110         /* next line produces an item and puts it in nextProduced */
111         nextProduced = means_of_production();
112         while (count == BUFFER_SIZE)
113             ; // do nothing
114         buffer [in] = nextProduced;
115         in = (in + 1) % BUFFER_SIZE;
116         count++;
117     }
118 }
119
120 void consumer (void *ignored) {
121     for (;;) {
122         while (count == 0)
123             ; // do nothing
124         nextConsumed = buffer[out];
125         out = (out + 1) % BUFFER_SIZE;
126         count--;
127         /* next line abstractly consumes the item */
128         consume_item(nextConsumed);
129     }
130 }
131
132 /*
133 what count++ probably compiles to:
134 reg1 <-- count    # load
135 reg1 <-- reg1 + 1 # increment register
136 count <-- reg1    # store
137
138 what count-- could compile to:
139 reg2 <-- count    # load
140 reg2 <-- reg2 - 1 # decrement register
141 count <-- reg2    # store
142 */
143
144 What happens if we get the following interleaving?
145
146 reg1 <-- count
147 reg1 <-- reg1 + 1
148 reg2 <-- count
149 reg2 <-- reg2 - 1
150 count <-- reg1
151 count <-- reg2
152

```

Feb 14, 11 23:23

I09-handout.txt

Page 4/8

```

153 5. Protecting the linked list.....
154
155     Lock list_lock;
156
157     insert(int data) {
158         List_elem* l = new List_elem;
159         l->data = data;
160
161         acquire(&list_lock);
162
163         l->next = head;    // A
164         head = l;        // B
165
166         release(&list_lock);
167     }
168

```

169 6. How can we implement list\_lock, acquire(), and release()?

170  
171 6a. Here is A BADLY BROKEN implementation:  
172

```
173     struct Lock {
174         int locked;
175     }
176
177     void [BROKEN] acquire(Lock *lock) {
178         while (1) {
179             if (lock->locked == 0) { // C
180                 lock->locked = 1; // D
181                 break;
182             }
183         }
184     }
185
186     void release (Lock *lock) {
187         lock->locked = 0;
188     }
189
```

190 What's the problem? Two acquire()s on the same lock on different CPUs  
191 might both execute line C, and then both execute D. Then both will  
192 think they have acquired the lock. This is the same kind of race we  
193 were trying to eliminate in insert(). But we have made a little  
194 progress: now we only need a way to prevent interleaving in one place  
195 (acquire()), not for many arbitrary complex sequences of code.  
196

197 6b. Here's a way that is correct but that is appropriate only in  
198 some circumstances:  
199

200 Use an atomic instruction on the CPU. For example, on the x86,  
201 doing

```
202     "xchg addr, %eax"
203 does the following:
```

```
204
205     (i) freeze all CPUs' memory activity for address addr
206     (ii) temp = *addr
207     (iii) *addr = %eax
208     (iv) %eax = temp
209     (v) un-freeze memory activity
210
```

```
211     /* pseudocode */
212     int xchg_val(addr, value) {
213         %eax = value;
214         xchg (*addr), %eax
215     }
216
```

```
217     struct Lock {
218         int locked;
219     }
220
```

```
221     /* bare-bones version of acquire */
222     void acquire (Lock *lock) {
223         pushcli(); /* what does this do? */
224         while (1) {
225             if (xchg_val(&lock->locked, 1) == 0)
226                 break;
227         }
228     }
229
```

```
230     /* optimization in acquire; call xchg_val() less frequently */
231     void acquire(Lock* lock) {
232         pushcli();
233         while (xchg_val(&lock->locked, 1) == 1) {
234             while (lock->locked) ;
235         }
236     }
237
```

```
238     void release(Lock *lock){
239         xchg_val(&lock->locked, 0);
240         popcli(); /* what does this do? */
241     }

```

242

243 The above is called a \*spinlock\* because acquire() waits in a  
244 busy loop.  
245

246

247 Unfortunately, insert() with these locks is only correct if each  
248 CPU carries out memory reads and writes in program order. For  
249 example, if the CPU were to execute insert() out of order so  
250 that it did the read at A before the acquire(), then insert()  
251 would be incorrect even with locks. Many modern processors  
252 execute memory operations out of order to increase performance!  
253 So we may have to use special instructions ("lock", "LFENCE",  
254 "SFENCE", "MFENCE") to tell the CPU not to re-order memory  
255 operations past acquire()s and release()s. The compiler may  
256 also generate instructions in orders that don't correspond to  
257 the order of the source code lines, so we have to worry about  
258 that too. One way around this is to make the asm instructions  
259 volatile.

260

261 Moral of the above paragraph: if you're implementing a  
262 concurrency primitive, read the processor's documentation about  
263 how loads and stores get sequenced, and how to enforce that the  
264 compiler \*and\* the processor follow program order.

265

266 The spinlock above is great for some things, not so great for  
267 others. The main problem is that it \*busy waits\*: it spins,  
268 chewing up CPU cycles. Sometimes this is what we want (e.g., if  
269 the cost of going to sleep is greater than the cost of spinning  
270 for a few cycles waiting for another thread or process to  
271 relinquish the spinlock). But sometimes this is not at all what we  
272 want (e.g., if the lock would be held for a while: in those  
273 cases, the CPU waiting for the lock would waste cycles spinning  
274 instead of running some other thread or process).

274

```

275
276 6c. Here's an object that does not involve busy waiting; it can work
277 as the list_lock mentioned in #5, above. Note: the "threads" here
278 can be user-level threads, kernel threads, or threads-inside-kernel.
279 The concept is the same in all cases.
280
281 struct Mutex {
282     bool is_held;           /* true if mutex held */
283     thread_id owner;       /* thread holding mutex, if locked */
284     thread_list waiters;   /* queue of thread TCBS */
285     Lock wait_lock;       /* as in 6b */
286 }
287
288 Now, instead of acquire(&list_lock) and release(&list_lock) as
289 in #5, we'd write, mutex_acquire(&list_mutex) and
290 mutex_release(&list_mutex). The implementation of the latter two
291 would be something like this:
292
293 void mutex_acquire(Mutex *m) {
294
295     acquire(&m->wait_lock); /* we spin to acquire wait_lock */
296     while (m->is_held) {    /* someone else has the mutex */
297         m->waiters.insert(current_thread)
298         release(&m->wait_lock);
299         schedule(); /* run a thread that is on the ready list */
300         acquire(&m->wait_lock); /* we spin again */
301     }
302     m->is_held = true;     /* we now hold the mutex */
303     m->owner = self;
304     release(&m->wait_lock);
305 }
306
307 void mutex_release(Mutex *m) {
308
309     acquire(&m->wait_lock); /* we spin to acquire wait_lock */
310     m->is_held = false;
311     m->owner = 0;
312     wake_up_a_waiter(m->waiters); /* select and run a waiter */
313     release(&m->wait_lock);
314
315 }
316
317 [Please let me (MW) know if you see bugs in the above.]
318

```

```

319
320 7. Terminology
321
322 To avoid confusion, we will use the following terminology in this
323 course (you will hear other terminology elsewhere):
324
325 --A "lock" is an abstract object that provides mutual exclusion
326
327 --A "spinlock" is a lock that works by busy waiting, as in 6b
328
329 --A "mutex" is a lock that works by having a "waiting" queue and
330 then protecting that waiting queue with atomic hardware
331 instructions, as in 6c. The most natural way to "use the hardware"
332 is with a spinlock, but there are others, such as turning off
333 interrupts, which works if we're on a single CPU machine.

```