

Feb 18, 10 9:32

I09-handout.txt

Page 1/6

```

1 Handout for CS 372H
2 Class 9
3 16 February 2010

```

1. Say that thread A executes `f()` and thread B executes `g()`. (Here, we are using the term "thread" abstractly. This example applies to any of the approaches that fall under the word "thread".)

1a.

```

11 int x;
12
13 f() { x = 1; }
14
15 g() { x = 2; }

```

What are possible values of `x` after A has executed `f()` and B has executed `g()`?

1b.

```

21 int y = 12;
22
23 f() { x = y + 1; }
24 g() { y = y * 2; }

```

What are the possible values of `x`?

1c.

```

29 int x = 0;
30 f() { x = x + 1; }
31 g() { x = x + 2; }

```

What are the possible values of `x`?

2. Linked list example

```

36 struct List_elem {
37     int data;
38     struct List_elem* next;
39 };
40
41 List_elem* head = 0;
42
43 insert(int data) {
44     List_elem* l = new List_elem;
45     l->data = data;
46     l->next = head;
47     head = l;
48 }

```

What happens if two threads execute `insert()` at once and we get the following interleaving?

```

54 thread 1: l->next = head
55 thread 2: l->next = head
56 thread 2: head = l;
57 thread 1: head = l;

```

Feb 18, 10 9:32

I09-handout.txt

Page 2/6

3. Producer/consumer example:

```

61 /*
62 "buffer" stores BUFFER_SIZE items
63 "count" is number of used slots. a variable that lives in memory
64 "out" is next empty buffer slot to fill (if any)
65 "in" is oldest filled slot to consume (if any)
66 */
67
68 void producer (void *ignored) {
69     for (;;) {
70         /* next line produces an item and puts it in nextProduced */
71         nextProduced = means_of_production();
72         while (count == BUFFER_SIZE)
73             ; // do nothing
74         buffer [in] = nextProduced;
75         in = (in + 1) % BUFFER_SIZE;
76         count++;
77     }
78 }
79
80 void consumer (void *ignored) {
81     for (;;) {
82         while (count == 0)
83             ; // do nothing
84         nextConsumed = buffer[out];
85         out = (out + 1) % BUFFER_SIZE;
86         count--;
87         /* next line abstractly consumes the item */
88         consume_item(nextConsumed);
89     }
90 }

```

```

91
92 /*
93 what count++ probably compiles to:
94 reg1 <-- count      # load
95 reg1 <-- reg1 + 1   # increment register
96 count <-- reg1     # store
97
98 what count-- could compile to:
99 reg2 <-- count      # load
100 reg2 <-- reg2 - 1   # decrement register
101 count <-- reg2     # store
102 */

```

What happens if we get the following interleaving?

```

106 reg1 <-- count
107 reg1 <-- reg1 + 1
108 reg2 <-- count
109 reg2 <-- reg2 - 1
110 count <-- reg1
111 count <-- reg2

```

Feb 18, 10 9:32

I09-handout.txt

Page 3/6

113 4. Protecting the linked list.....

```

114
115     Lock list_lock;
116
117     insert(int data) {
118         List_elem* l = new List_elem;
119         l->data = data;
120
121         acquire(&list_lock);
122
123         l->next = head;    // A
124         head = l;        // B
125
126         release(&list_lock);
127     }

```

128 5. How can we implement list\_lock, acquire(), and release()?

129 5a. Here is A BADLY BROKEN implementation:

```

130
131     struct Lock {
132         int locked;
133     }
134
135     void [BROKEN] acquire(Lock *lock) {
136         while (1) {
137             if (lock->locked == 0) { // C
138                 lock->locked = 1;    // D
139                 break;
140             }
141         }
142     }
143
144     void release (Lock *lock) {
145         lock->locked = 0;
146     }

```

147  
148  
149  
150 What's the problem? Two acquire()s on the same lock on different CPUs  
151 might both execute line C, and then both execute D. Then both will  
152 think they have acquired the lock. This is the same kind of race we  
153 were trying to eliminate in insert(). But we have made a little  
154 progress: now we only need a way to prevent interleaving in one place  
155 (acquire()), not for many arbitrary complex sequences of code.  
156

Feb 18, 10 9:32

I09-handout.txt

Page 4/6

157 5b. Here's a way that is correct but that is appropriate only in  
158 some circumstances:

```

159
160     Use an atomic instruction on the CPU. For example, on the x86,
161     doing
162         "xchg addr, %eax"
163     does the following:

```

```

164
165     (i) freeze all CPUs' memory activity for address addr
166     (ii) temp = *addr
167     (iii) *addr = %eax
168     (iv) %eax = temp
169     (v) un-freeze memory activity

```

```

170
171     /* pseudocode */
172     int xchg_val(addr, value) {
173         %eax = value;
174         xchg (*addr), %eax
175     }

```

```

176
177     struct Lock {
178         int locked;
179     }

```

```

180
181     void acquire (Lock *lock) {
182         pushcli(); /* what does this do? */
183         while (1) {
184             if(xchg_val(&lock->locked, 1) == 0)
185                 break;
186         }
187     }

```

```

188
189     void release(Lock *lock){
190         xchg_val(&lock->locked, 0);
191         popcli(); /* what does this do? */
192     }

```

193  
194 The above is called a \*spinlock\* because acquire() waits in a  
195 busy loop.

196  
197 Unfortunately, insert() with these locks is only correct if each  
198 CPU carries out memory reads and writes in program order. For  
199 example, if the CPU were to execute insert() out of order so  
200 that it did the read at A before the acquire(), then insert()  
201 would be incorrect even with locks. Many modern processors  
202 execute memory operations out of order to increase performance!  
203 So we may have to use special instructions ("lock", "LFENCE",  
204 "SFENCE", "MFENCE") to tell the CPU not to re-order memory  
205 operations past acquire()s and release()s. The compiler may  
206 also generate instructions in orders that don't correspond to  
207 the order of the source code lines, so we have to worry about  
208 that too. One way around this is to make the asm instructions  
209 volatile.

210  
211 Moral of the above paragraph: if you're implementing a  
212 concurrency primitive, read the processor's documentation about  
213 how loads and stores get sequenced, and how to enforce that the  
214 compiler \*and\* the processor follow program order.  
215

216  
217 The spinlock above is great for some things, not so great for  
218 others. The main problem is that it \*busy waits\*: it spins,  
219 chewing up CPU cycles. Sometimes this is what we want (e.g., if  
220 the cost of going to sleep is greater than the cost of spinning  
221 for a few cycles waiting for another thread or process to  
222 relinquish the spinlock). But sometimes this is not at all what we  
223 want (e.g., if the lock would be held for a while: in those  
224 cases, the CPU waiting for the lock would waste cycles spinning  
225 instead of running some other thread or process).

Feb 18, 10 9:32

I09-handout.txt

Page 5/6

```

226 5c. Here's a lock that does not involve busy waiting. Note: the
227 "threads" here can be user-level threads, kernel threads, or
228 threads-inside-kernel. The concept is the same in all cases.
229
230 struct Mutex {
231     bool is_locked;           /* true if locked */
232     thread_id owner;         /* thread holding lock, if locked */
233     thread_list waiters;     /* queue of thread TCBS */
234     spinlock wait_lock;     /* exactly as in 5b */
235 }
236
237 Now, mutex.acquire() looks something like this:
238
239     wait_lock.acquire()
240     while (is_locked) {
241         waiters.insert(current_thread)
242         wait_lock.release()
243         schedule(); /* run a thread that is on the ready list */
244         wait_lock.acquire();
245     }
246     is_locked = 1;
247     owner = self;
248     wait_lock.release();
249
250 And mutex.release() looks something like this:
251
252     wait_lock.acquire()
253     is_locked = 0;
254     owner = 0;
255     wake_up_a_waiter(); /* selects a waiter and runs it */
256     wait_lock.release()
257
258 [Please let me (MW) know if you see bugs in the above.]
259

```

Feb 18, 10 9:32

I09-handout.txt

Page 6/6

```

260 6. Terminology
261
262 To avoid confusion, we will use the following terminology in this
263 course (you will hear other terminology elsewhere):
264
265 --A "lock" is an abstract object that provides mutual exclusion
266
267 --A "spinlock" is a lock that works by busy waiting, as in 5b
268
269 --A "mutex" is a lock that works by having a "waiting" queue and
270 then protecting that waiting queue with atomic hardware
271 instructions, as in 5c. The most natural way to "use the hardware"
272 is with a spinlock, but there are others, such as turning off
273 interrupts, which works if we're on a single CPU machine.

```