

```

1 Handout for CS 372H
2 Class 10
3 17 February 2011
4
5 1. Protecting the linked list.....
6
7     Lock list_lock;
8
9     insert(int data) {
10         List_elem* l = new List_elem;
11         l->data = data;
12
13         acquire(&list_lock);
14
15         l->next = head;        // A
16         head = l;            // B
17
18         release(&list_lock);
19     }
20
21 2. How can we implement list_lock, acquire(), and release()?
22
23     2a. Here is A BADLY BROKEN implementation:
24
25     struct Lock {
26         int locked;
27     }
28
29     void [BROKEN] acquire(Lock *lock) {
30         while (1) {
31             if (lock->locked == 0) { // C
32                 lock->locked = 1;   // D
33                 break;
34             }
35         }
36     }
37
38     void release (Lock *lock) {
39         lock->locked = 0;
40     }
41
42     What's the problem? Two acquire()'s on the same lock on different CPUs
43     might both execute line C, and then both execute D. Then both will
44     think they have acquired the lock. This is the same kind of race we
45     were trying to eliminate in insert(). But we have made a little
46     progress: now we only need a way to prevent interleaving in one place
47     (acquire()), not for many arbitrary complex sequences of code.
48

```

```

49     2b. Here's a way that is correct but only sometimes appropriate:
50     Use an atomic instruction on the CPU. For example, on the x86,
51     doing
52         "xchg addr, %eax"
53     does the following:
54
55     (i) freeze all CPUs' memory activity for address addr
56     (ii) temp = *addr
57     (iii) *addr = %eax
58     (iv) %eax = temp
59     (v) un-freeze memory activity
60
61     /* pseudocode */
62     int xchg_val(addr, value) {
63         %eax = value;
64         xchg (*addr), %eax
65     }
66
67     struct Lock {
68         int locked;
69     }
70
71     /* bare-bones version of acquire */
72     void acquire (Lock *lock) {
73         pushcli(); /* what does this do? */
74         while (1) {
75             if (xchg_val(&lock->locked, 1) == 0)
76                 break;
77         }
78     }
79
80     /* optimization in acquire; call xchg_val() less frequently */
81     void acquire(Lock* lock) {
82         pushcli();
83         while (xchg_val(&lock->locked, 1) == 1) {
84             while (lock->locked) ;
85         }
86     }
87
88     void release(Lock *lock){
89         xchg_val(&lock->locked, 0);
90         popcli(); /* what does this do? */
91     }
92
93     The above is called a *spinlock* because acquire() spins.
94
95     Unfortunately, insert() with these locks is correct only if each
96     CPU carries out memory reads and writes in program order. For
97     example, if the CPU were to execute insert() out of order so
98     that it did the read at A before the acquire(), then insert()
99     would be incorrect even with locks. Many modern processors
100    execute memory operations out of order to increase performance!
101    So we may have to use special instructions ("lock", "LFENCE",
102    "SFENCE", "MFENCE") to tell the CPU not to re-order memory
103    operations past acquire()'s and release()'s. The compiler may
104    also generate instructions in orders that don't correspond to
105    the order of the source code lines, so we have to worry about
106    that too. One way around this is to make the asm instructions
107    volatile.
108
109    Moral of the above paragraph: if you're implementing a
110    concurrency primitive, read the processor's documentation about
111    how loads and stores get sequenced (chapter 8 in current
112    architecture manual).
113
114    The spinlock above is great for some things, not so great for
115    others. The main problem is that it *busy waits*: it spins,
116    chewing up CPU cycles. Sometimes this is what we want (e.g., if
117    the cost of going to sleep is greater than the cost of spinning
118    for a few cycles waiting for another thread or process to
119    relinquish the spinlock). But sometimes this is not at all what we
120    want (e.g., if the lock would be held for a while: in those
121    cases, the CPU waiting for the lock would waste cycles spinning

```

122 instead of running some other thread or process).

123

124

125 2c. Here's an object that does not involve busy waiting; it can work
126 as the list_lock mentioned in #1, above. Note: the "threads" here
127 can be user-level threads, kernel threads, or threads-inside-kernel.
128 The concept is the same in all cases.

129

130

131

132

133

134

135

136

137

138 Now, instead of acquire(&list_lock) and release(&list_lock) as
139 in #1, we'd write, mutex_acquire(&list_mutex) and
140 mutex_release(&list_mutex). The implementation of the latter two
141 would be something like this:

142

143

144

145

146

147

148

149

150

151

152

153

154

155

156

157

158

159

160

161

162

163

164

165

166 [Please let me (MW) know if you see bugs in the above.]

167

168

169

169 3. Terminology

170

171 To avoid confusion, we will use the following terminology in this
172 course (you will hear other terminology elsewhere):

173

174

175 --A "lock" is an abstract object that provides mutual exclusion

176

177

178 --A "spinlock" is a lock that works by busy waiting, as in 6b

179

180

181 --A "mutex" is a lock that works by having a "waiting" queue and
182 then protecting that waiting queue with atomic hardware
183 instructions, as in 6c. The most natural way to "use the hardware"
184 is with a spinlock, but there are others, such as turning off
185 interrupts, which works if we're on a single CPU machine.

186

187

188

189

190

191

192

193

194

195

196

197

198

199

200

201

202

185

186

187

188

189

190

191

192

193

194

195

196

197

198

199

200

201

202

203

204

205

206

207

208

209

210

211

212

213

214

215

216

217

218

219

220

221

222

223

224

225

226

227

228

229

230

231

232

233

234

235

236

237

238

239

240

241

186 4. Producer/consumer example [also known as bounded buffer]

187

188

189

190

191

192

193

194

195

196

197

198

199

200

201

202

203

204

205

206

207

208

209

210

211

212

213

214

215

216

217

218

219

220

221

222

223

224

225

226

227

228

229

230

231

232

233

234

235

236

237

238

239

240

241

186 4. Producer/consumer example [also known as bounded buffer]

187

188

189

190

191

192

193

194

195

196

197

198

199

200

201

202

203

204

205

206

207

208

209

210

211

212

213

214

215

216

217

218

219

220

221

222

223

224

225

226

227

228

229

230

231

232

233

234

235

236

237

238

239

240

241

186 4. Producer/consumer example [also known as bounded buffer]

187

188

189

190

191

192

193

194

195

196

197

198

199

200

201

202

203

204

205

206

207

208

209

210

211

212

213

214

215

216

217

218

219

220

221

222

223

224

225

226

227

228

229

230

231

232

233

234

235

236

237

238

239

240

241

Feb 17, 11 16:22

I10-handout.txt

Page 5/9

```

242
243 4b. Producer/consumer [bounded buffer] using mutexes
244
245     Mutex mutex;
246
247     void producer (void *ignored) {
248         for (;;) {
249             /* next line produces an item and puts it in nextProduced */
250             nextProduced = means_of_production();
251
252             acquire(&mutex);
253             while (count == BUFFER_SIZE) {
254                 release(&mutex);
255                 yield(); /* or schedule() */
256                 acquire(&mutex);
257             }
258
259             buffer [in] = nextProduced;
260             in = (in + 1) % BUFFER_SIZE;
261             count++;
262             release(&mutex);
263         }
264     }
265
266     void consumer (void *ignored) {
267         for (;;) {
268
269             acquire(&mutex);
270             while (count == 0) {
271                 release(&mutex);
272                 yield(); /* or schedule() */
273                 acquire(&mutex);
274             }
275
276             nextConsumed = buffer[out];
277             out = (out + 1) % BUFFER_SIZE;
278             count--;
279             release(&mutex);
280
281             /* next line abstractly consumes the item */
282             consume_item(nextConsumed);
283         }
284     }
285

```

Feb 17, 11 16:22

I10-handout.txt

Page 6/9

```

286
287 4c. Producer/consumer [bounded buffer] using mutexes and condition
288     variables
289
290     Mutex mutex;
291     Cond nonempty;
292     Cond nonfull;
293
294     void producer (void *ignored) {
295         for (;;) {
296             /* next line produces an item and puts it in nextProduced */
297             nextProduced = means_of_production();
298
299             acquire(&mutex);
300             while (count == BUFFER_SIZE)
301                 cond_wait(&nonfull, &mutex);
302
303             buffer [in] = nextProduced;
304             in = (in + 1) % BUFFER_SIZE;
305             count++;
306             cond_signal(&nonempty);
307             release(&mutex);
308         }
309     }
310
311     void consumer (void *ignored) {
312         for (;;) {
313
314             acquire(&mutex);
315             while (count == 0)
316                 cond_wait(&nonempty, &mutex);
317
318             nextConsumed = buffer[out];
319             out = (out + 1) % BUFFER_SIZE;
320             count--;
321             cond_signal(&nonfull);
322             release(&mutex);
323
324             /* next line abstractly consumes the item */
325             consume_item(nextConsumed);
326         }
327     }
328
329
330     Question: why does cond_wait need to both release the mutex and
331     sleep? Why not:
332
333         while (count == BUFFER_SIZE) {
334             release(&mutex);
335             cond_wait(&nonfull);
336             acquire(&mutex);
337         }
338

```

Feb 17, 11 16:22

I10-handout.txt

Page 7/9

```

339 4d. Producer/consumer [bounded buffer] with semaphores
340
341 Semaphore mutex(1);          /* mutex initialized to 1 */
342 Semaphore empty(BUFFER_SIZE); /* start with BUFFER_SIZE empty slots */
343 Semaphore full(0);           /* 0 full slots */
344
345 void producer (void *ignored) {
346     for (;;) {
347         /* next line produces an item and puts it in nextProduced */
348         nextProduced = means_of_production();
349
350         /*
351          * next line diminishes the count of empty slots and
352          * waits if there are no empty slots
353          */
354         sem_down(&empty);
355         sem_down(&mutex); /* get exclusive access */
356
357         buffer [in] = nextProduced;
358         in = (in + 1) % BUFFER_SIZE;
359
360         sem_up(&mutex);
361         sem_up(&full); /* we just increased the # of full slots */
362     }
363 }
364
365 void consumer (void *ignored) {
366     for (;;) {
367
368         /*
369          * next line diminishes the count of full slots and
370          * waits if there are no full slots
371          */
372         sem_down(&full);
373         sem_down(&mutex);
374
375         nextConsumed = buffer[out];
376         out = (out + 1) % BUFFER_SIZE;
377
378         sem_up(&mutex);
379         sem_up(&empty); /* one further empty slot */
380
381         /* next line abstractly consumes the item */
382         consume_item(nextConsumed);
383     }
384 }
385
386 Semaphores *can* (not always) lead to elegant solutions (notice
387 that the code above is fewer lines than 1c) but they are much
388 harder to use.
389
390 The fundamental issue is that semaphores make implicit (counts,
391 conditions, etc.) what is probably best left explicit. Moreover,
392 they *also* implement mutual exclusion.
393
394 For this reason, you should not use semaphores. This example is
395 here mainly for completeness and so you know what a semaphore
396 is. But do not code with them. Solutions that use semaphores in
397 this course will receive no credit.
398

```

Feb 17, 11 16:22

I10-handout.txt

Page 8/9

```

399 5. Example of a monitor: MyBuffer
400
401 // This is pseudocode that is inspired by C++.
402 // Don't take it literally.
403
404 class MyBuffer {
405     public:
406         MyBuffer();
407         ~MyBuffer();
408         void Enqueue(Item);
409         Item = Dequeue();
410     private:
411         int count;
412         int in;
413         int out;
414         Item buffer[BUFFER_SIZE];
415         Mutex* mutex;
416         Cond* nonempty;
417         Cond* nonfull;
418     }
419
420 void
421 MyBuffer::MyBuffer()
422 {
423     in = out = count = 0;
424     mutex = new Mutex;
425     nonempty = new Cond;
426     nonfull = new Cond;
427 }
428
429 void
430 MyBuffer::Enqueue(Item item)
431 {
432     mutex.acquire();
433     while (count == BUFFER_SIZE)
434         cond_wait(&nonfull, &mutex);
435
436     buffer[in] = item;
437     in = (in + 1) % BUFFER_SIZE;
438     ++count;
439     cond_signal(&nonempty, &mutex);
440     mutex.release();
441 }
442
443 Item
444 MyBuffer::Dequeue()
445 {
446     mutex.acquire();
447     while (count == 0)
448         cond_wait(&nonempty, &mutex);
449
450     Item ret = buffer[out];
451     out = (out + 1) % BUFFER_SIZE;
452     --count;
453     cond_signal(&nonfull, &mutex);
454     mutex.release();
455     return ret;
456 }
457

```

Feb 17, 11 16:22

l10-handout.txt

Page 9/9

```
458 int main(int, char**)
459 {
460     MyBuffer buf;
461     int dummy;
462     tid1 = thread_create(producer, &buf);
463     tid2 = thread_create(consumer, &buf);
464     thread_join(tid1);
465
466     // never reach this point
467     return -1;
468 }
469
470 void producer(void* buf)
471 {
472     MyBuffer* sharedbuf = reinterpret_cast<MyBuffer*>(buf);
473     for (;;) {
474         /* next line produces an item and puts it in nextProduced */
475         Item nextProduced = means_of_production();
476         sharedbuf->Enqueue(nextProduced);
477     }
478 }
479
480 void consumer(void* buf)
481 {
482     MyBuffer* sharedbuf = reinterpret_cast<MyBuffer*>(buf);
483     for (;;) {
484         Item nextConsumed = sharedbuf->Dequeue();
485
486         /* next line abstractly consumes the item */
487         consume_item(nextConsumed);
488     }
489 }
490
491 Key point: *Threads* (the producer and consumer) are separate from
492 *shared object* (MyBuffer). The synchronization happens in the
493 shared object.
```