

- 1. Last time
 - 2. Condition variables
 - 3. Monitors and standards
 - 4. Advice -
 - 5. Practice w/ concurrent programming -
 - 6. (preview) Implementation of locks: spinlocks, mutexes
-

2. CVs

A. Motivation

cond-init(CV)

B. API

cond-init (Cond *, ..);

cond-wait (Cond *, Mutex* m);

cond-signal (Cond *, ...);

cond-broadcast (Cond *, ...);

C. Important points

13 Must use "while" not "if"

```
while ( _ ) {  
    signal();  
}
```

wait():
atomically rel
sleep
.....
acquire();

14 Cond.wait() releases mutex and goes into waiting state atomically; why?

alternative: producer:

```
release(&m);  
cond_wait(&cv);  
acquire(&m);
```

consumer:

```
acquire(&m);  
cond_signal(&cv);  
nonnull
```

15 Can we replace signal() w/ broadcast()? **Yes**

16 Can we replace broadcast() w/ signal()? **No**

Monitors and standards

5. Monitor

Monitor = one mutex + one or more CVs

M::f()

{ acquire(&m);



release(&m);
}

class M {

private:

Mutex

CV cv1;

.. cv2;

public

f();

g();

}

M::g()

{ acquire(&m);



release(&m);

}

Commandments:

Rule: acquire/release at beginning/end of method or function.

Rule: hold lock when doing CV operations


Rule: a thread in `wait()` must be prepared to be restarted any time, not just when another thread calls `signal()`; \Rightarrow while ^{not safe to} ~~prevent~~ `wait()`;

Rule: don't call `sleep()`;

alloc/free example



alloc (1KB)
wait()
free



T. Hoare

B. Hansen

Sep 22, 21 1:15

handout04.txt

Page 1/4

1 CS 202, Fall 2021
2 Handout 4 (Class 5)

3
4 The handout from the last class gave examples of race conditions. The following
5 panels demonstrate the use of concurrency primitives (mutexes, etc.). We are
6 using concurrency primitives to eliminate race conditions (see items 1
7 and 2a) and improve scheduling (see item 2b).

8
9 1. Protecting the linked list.....

```
10
11     Mutex list_mutex;
12
13     insert(int data) {
14         List_elem* l = new List_elem;
15         l->data = data;
16
17         acquire(&list_mutex);
18
19         l->next = head;
20         head = l;
21
22         release(&list_mutex);
23     }
24
```

Sep 22, 21 1:15



handout04.txt

Page 2/4

25 2. Producer/consumer revisited [also known as bounded buffer]

26
27 2a. Producer/consumer [bounded buffer] with mutexes

```
28     Mutex mutex;
29
30     void producer (void *ignored) {
31         for (;;) {
32             /* next line produces an item and puts it in nextProduced */
33             nextProduced = means_of_production();
34
35             acquire(&mutex);
36             while (count == BUFFER_SIZE) {
37                 release(&mutex);
38                 yield(); /* or schedule() */
39                 acquire(&mutex);
40             }
41             buffer[in] = nextProduced;
42             in = (in + 1) % BUFFER_SIZE;
43             count++;
44             release(&mutex);
45         }
46     }
47
48     void consumer (void *ignored) {
49         for (;;) {
50             acquire(&mutex);
51             while (count == 0) {
52                 release(&mutex);
53                 yield(); /* or schedule() */
54                 acquire(&mutex);
55             }
56             nextConsumed = buffer[out];
57             out = (out + 1) % BUFFER_SIZE;
58             count--;
59             release(&mutex);
60
61             /* next line abstractly consumes the item */
62             consume_item(nextConsumed);
63         }
64     }
65
66     }
67
68
69
```



Sep 22, 21 1:15

handout04.txt

Page 3/4

2b. Producer/consumer [bounded buffer] with mutexes and condition variables

```

70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121

```

Mutex mutex;
Cond nonempty;
Cond nonfull;

void producer (void *ignored) {
for (;;) {
/* next line produces an item and puts it in nextProduced */
nextProduced = means_of_production();
acquire(&mutex);
while (count == BUFFER_SIZE)
cond_wait(&nonfull, &mutex);
buffer[in] = nextProduced;
in = (in + 1) % BUFFER_SIZE;
count++;
cond_signal(&nonempty, &mutex);
release(&mutex);
}

void consumer (void *ignored) {
for (;;) {
acquire(&mutex);
while (count == 0)
cond_wait(&nonempty, &mutex);
nextConsumed = buffer[out];
out = (out + 1) % BUFFER_SIZE;
count--;
cond_signal(&nonfull, &mutex);
release(&mutex);
/* next line abstractly consumes the item */
consume_item(nextConsumed);
}

Question: why does cond_wait need to both release the mutex and sleep? Why not:

```

while (count == BUFFER_SIZE) {
    release(&mutex);
    cond_wait(&nonfull);
    acquire(&mutex);
}

```

Sep 22, 21 1:15

handout04.txt

Page 4/4

2c. Producer/consumer [bounded buffer] with semaphores

```

122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
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
167
168
169
170
171
172
173
174
175
176
177
178
179
180

```

Semaphore mutex(1); /* mutex initialized to 1 */
Semaphore empty(BUFFER_SIZE); /* start with BUFFER_SIZE empty slots */
Semaphore full(0); /* 0 full slots */

void producer (void *ignored) {
for (;;) {
/* next line produces an item and puts it in nextProduced */
nextProduced = means_of_production();
/* next line diminishes the count of empty slots and
* waits if there are no empty slots
*/
sem_down(&empty);
sem_down(&mutex); /* get exclusive access */
buffer[in] = nextProduced;
in = (in + 1) % BUFFER_SIZE;
sem_up(&mutex);
sem_up(&full); /* we just increased the # of full slots */
}

void consumer (void *ignored) {
for (;;) {
/* next line diminishes the count of full slots and
* waits if there are no full slots
*/
sem_down(&full);
sem_down(&mutex);
nextConsumed = buffer[out];
out = (out + 1) % BUFFER_SIZE;
sem_up(&mutex);
sem_up(&empty); /* one further empty slot */
/* next line abstractly consumes the item */
consume_item(nextConsumed);
}

Semaphores *can* (not always) lead to elegant solutions (notice that the code above is fewer lines than 2b) but they are much harder to use.

The fundamental issue is that semaphores make implicit (counts, conditions, etc.) what is probably best left explicit. Moreover, they *also* implement mutual exclusion.

For this reason, you should not use semaphores. This example is here mainly for completeness and so you know what a semaphore is. But do not code with them. Solutions that use semaphores in this course will receive no credit.

Sep 26, 21 21:09

handout05.txt

Page 1/4

```

1 CS 202, Fall 2021
2 Handout 5 (Class 6)
3
4 The previous handout demonstrated the use of mutexes and condition
5 variables. This handout demonstrates the use of monitors (which combine
6 mutexes and condition variables).

```

1. The bounded buffer as a monitor

```

10 // This is pseudocode that is inspired by C++.
11 // Don't take it literally.

```

```

12 class MyBuffer {
13     public:
14         MyBuffer();
15         ~MyBuffer();
16         void Enqueue(Item);
17         Item Dequeue();
18     private:
19         int count;
20         int in;
21         int out;
22         Item buffer[BUFFER_SIZE];
23         Mutex* mutex;
24         Cond* nonempty;
25         Cond* nonfull;
26 }
27
28 void
29 MyBuffer::MyBuffer()
30 {
31     in = out = count = 0;
32     mutex = new Mutex;
33     nonempty = new Cond;
34     nonfull = new Cond;
35 }
36
37 void
38 MyBuffer::Enqueue(Item item)
39 {
40     mutex.acquire();
41     while (count == BUFFER_SIZE)
42         cond_wait(&nonfull, &mutex);
43
44     buffer[in] = item;
45     in = (in + 1) % BUFFER_SIZE;
46     ++count;
47     cond_signal(&nonempty, &mutex);
48     mutex.release();
49 }
50
51 Item
52 MyBuffer::Dequeue()
53 {
54     mutex.acquire();
55     while (count == 0)
56         cond_wait(&nonempty, &mutex);
57
58     Item ret = buffer[out];
59     out = (out + 1) % BUFFER_SIZE;
60     --count;
61     cond_signal(&nonfull, &mutex);
62     mutex.release();
63     return ret;
64 }
65
66

```

Sep 26, 21 21:09

handout05.txt

Page 2/4

```

67
68 int main(int, char**)
69 {
70     MyBuffer buf;
71     int dummy;
72     tid1 = thread_create(producer, &buf);
73     tid2 = thread_create(consumer, &buf);
74
75     // never reach this point
76     thread_join(tid1);
77     thread_join(tid2);
78     return -1;
79 }
80
81 void producer(void* buf)
82 {
83     MyBuffer* sharedbuf = reinterpret_cast<MyBuffer*>(buf);
84     for (;;) {
85         /* next line produces an item and puts it in nextProduced */
86         Item nextProduced = means_of_production();
87         sharedbuf->Enqueue(nextProduced);
88     }
89 }
90
91 void consumer(void* buf)
92 {
93     MyBuffer* sharedbuf = reinterpret_cast<MyBuffer*>(buf);
94     for (;;) {
95         Item nextConsumed = sharedbuf->Dequeue();
96
97         /* next line abstractly consumes the item */
98         consume_item(nextConsumed);
99     }
100 }
101
102 Key point: *Threads* (the producer and consumer) are separate from
103 *shared object* (MyBuffer). The synchronization happens in the
104 shared object.
105

```


Sep 26, 21 21:09

handout05.txt

Page 3/4

```

106 2. This monitor is a model of a database with multiple readers and
107 writers. The high-level goal here is (a) to give a writer exclusive
108 access (a single active writer means there should be no other writers
109 and no readers) while (b) allowing multiple readers. Like the previous
110 example, this one is expressed in pseudocode.
111
112 // assume that these variables are initialized in a constructor
113 state variables:
114     AR = 0; // # active readers
115     AW = 0; // # active writers
116     WR = 0; // # waiting readers
117     WW = 0; // # waiting writers
118
119     Condition okToRead = NIL;
120     Condition okToWrite = NIL;
121     Mutex mutex = FREE;
122
123 Database::read() {
124     startRead(); // first, check self into the system
125     Access Data
126     doneRead(); // check self out of system
127 }
128
129 Database::startRead() {
130     acquire(&mutex);
131     while((AW + WW) > 0){
132         WR++;
133         wait(&okToRead, &mutex);
134         WR--;
135     }
136     AR++;
137     release(&mutex);
138 }
139
140 Database::doneRead() {
141     acquire(&mutex);
142     AR--;
143     if (AR == 0 && WW > 0) { // if no other readers still
144         signal(&okToWrite, &mutex); // active, wake up writer
145     }
146     release(&mutex);
147 }
148
149 Database::write(){ // symmetrical
150     startWrite(); // check in
151     Access Data
152     doneWrite(); // check out
153 }
154
155 Database::startWrite() {
156     acquire(&mutex);
157     while ((AW + AR) > 0) { // check if safe to write.
158         // if any readers or writers, wait
159         WW++;
160         wait(&okToWrite, &mutex);
161         WW--;
162     }
163     AW++;
164     release(&mutex);
165 }
166
167 Database::doneWrite() {
168     acquire(&mutex);
169     AW--;
170     if (WW > 0) {
171         signal(&okToWrite, &mutex); // give priority to writers
172     } else if (WR > 0) {
173         broadcast(&okToRead, &mutex);
174     }
175     release(&mutex);
176 }
177
178 NOTE: what is the starvation problem here?

```

Sep 26, 21 21:09

handout05.txt

Page 4/4

```

179
180 3. Shared locks
181
182 struct sharedlock {
183     int i;
184     Mutex mutex;
185     Cond c;
186 };
187
188 void AcquireExclusive (sharedlock *sl) {
189     acquire(&sl->mutex);
190     while (sl->i) {
191         wait (&sl->c, &sl->mutex);
192     }
193     sl->i = -1;
194     release(&sl->mutex);
195 }
196
197 void AcquireShared (sharedlock *sl) {
198     acquire(&sl->mutex);
199     while (sl->i < 0) {
200         wait (&sl->c, &sl->mutex);
201     }
202     sl->i++;
203     release(&sl->mutex);
204 }
205
206 void ReleaseShared (sharedlock *sl) {
207     acquire(&sl->mutex);
208     if (!--sl->i)
209         signal (&sl->c, &sl->mutex);
210     release(&sl->mutex);
211 }
212
213 void ReleaseExclusive (sharedlock *sl) {
214     acquire(&sl->mutex);
215     sl->i = 0;
216     broadcast (&sl->c, &sl->mutex);
217     release(&sl->mutex);
218 }
219
220 QUESTIONS:
221 A. There is a starvation problem here. What is it? (Readers can keep
222    writers out if there is a steady stream of readers.)
223 B. How could you use these shared locks to write a cleaner version
224    of the code in the prior item? (Though note that the starvation
225    properties would be different.)

```

Sep 26, 21 21:09

spinlock-mutex.txt

Page 1/3

```

1 Implementation of spinlocks and mutexes
2
3 1. Here is a BROKEN spinlock implementation:
4
5     struct Spinlock {
6         int locked;
7     }
8
9     void acquire(Spinlock *lock) {
10         while (1) {
11             if (lock->locked == 0) { // A
12                 lock->locked = 1;    // B
13                 break;
14             }
15         }
16     }
17
18     void release (Spinlock *lock) {
19         lock->locked = 0;
20     }
21
22     What's the problem? Two acquire()s on the same lock on different
23     CPUs might both execute line A, and then both execute B. Then
24     both will think they have acquired the lock. Both will proceed.
25     That doesn't provide mutual exclusion.
26

```

Sep 26, 21 21:09

spinlock-mutex.txt

Page 2/3

```

26
27 2. Correct spinlock implementation
28
29     Relies on atomic hardware instruction. For example, on the x86-64,
30     doing
31         "xchg addr, %rax"
32     does the following:
33
34         (i) freeze all CPUs' memory activity for address addr
35         (ii) temp <-- *addr
36         (iii) *addr <-- %rax
37         (iv) %rax <-- temp
38         (v) un-freeze memory activity
39
40     /* pseudocode */
41     int xchg_val(addr, value) {
42         %rax = value;
43         xchg (*addr), %rax
44     }
45
46     /* bare-bones version of acquire */
47     void acquire (Spinlock *lock) {
48         pushcli(); /* what does this do? */
49         while (1) {
50             if (xchg_val(&lock->locked, 1) == 0)
51                 break;
52         }
53     }
54
55     void release(Spinlock *lock){
56         xchg_val(&lock->locked, 0);
57         popcli(); /* what does this do? */
58     }
59
60
61     /* optimization in acquire; call xchg_val() less frequently */
62     void acquire(Spinlock* lock) {
63         pushcli();
64         while (xchg_val(&lock->locked, 1) == 1) {
65             while (lock->locked) ;
66         }
67     }
68
69     The above is called a *spinlock* because acquire() spins. The
70     bare-bones version is called a "test-and-set (TAS) spinlock"; the
71     other is called a "test-and-test-and-set spinlock".
72
73     The spinlock above is great for some things, not so great for
74     others. The main problem is that it *busy waits*: it spins,
75     chewing up CPU cycles. Sometimes this is what we want (e.g., if
76     the cost of going to sleep is greater than the cost of spinning
77     for a few cycles waiting for another thread or process to
78     relinquish the spinlock). But sometimes this is not at all what we
79     want (e.g., if the lock would be held for a while: in those
80     cases, the CPU waiting for the lock would waste cycles spinning
81     instead of running some other thread or process).
82
83     NOTE: the spinlocks presented here can introduce performance issues
84     when there is a lot of contention. (This happens even if the
85     programmer is using spinlocks correctly.) The performance issues
86     result from cross-talk among CPUs (which undermines caching and
87     generates traffic on the memory bus). If we have time later, we will
88     study a remediation of this issue (search the Web for "MCS locks").
89
90     ANOTHER NOTE: In everyday application-level programming, spinlocks
91     will not be something you use (use mutexes instead). But you should
92     know what these are for technical literacy, and to see where the
93     mutual exclusion is truly enforced on modern hardware.
94

```

Sep 26, 21 21:09

spinlock-mutex.txt

Page 3/3

95 3. Mutex implementation

96
 97 The intent of a mutex is to avoid busy waiting: if the lock is not
 98 available, the locking thread is put to sleep, and tracked by a
 99 queue in the mutex. The next page has an implementation.
 100
 101

Sep 26, 21 21:09

fair-mutex.c

Page 1/2

```

1  #include <sys/queue.h>
2
3  typedef struct thread {
4      // ... Entries elided.
5      STAILQ_ENTRY(thread_t) qlink; // Tail queue entry.
6  } thread_t;
7
8  struct Mutex {
9      // Current owner, or 0 when mutex is not held.
10     thread_t *owner;
11
12     // List of threads waiting on mutex
13     STAILQ(thread_t) waiters;
14
15     // A lock protecting the internals of the mutex.
16     Spinlock splock; // as in item 1, above
17 };
18
19 void mutex_acquire(struct Mutex *m) {
20
21     acquire(&m->splock);
22
23     // Check if the mutex is held; if not, current thread gets mutex and returns
24     if (m->owner == 0) {
25         m->owner = id_of_this_thread;
26         release(&m->splock);
27     } else {
28         // Add thread to waiters.
29         STAILQ_INSERT_TAIL(&m->waiters, id_of_this_thread, qlink);
30
31         // Tell the scheduler to add current thread to the list
32         // of blocked threads. The scheduler needs to be careful
33         // when a corresponding sched_wakeup call is executed to
34         // make sure that it treats running threads correctly.
35         sched_mark_blocked(&id_of_this_thread);
36
37         // Unlock spinlock.
38         release(&m->splock);
39
40         // Stop executing until woken.
41         sched_swch();
42
43         // When we get to this line, we are guaranteed to hold the mutex. This
44         // is because we can get here only if context-switched-TO, which itself
45         // can happen only if this thread is removed from the waiting queue,
46         // marked "unblocked", and set to be the owner (in mutex_release()
47         // below). However, we might actually have held the mutex in lines 39-42
48
49         // (if we were context-switched out after the spinlock release(),
50         // followed by being run as a result of another thread's release of the
51         // mutex). But if that happens, it just means that we are
52         // context-switched out an "extra" time before proceeding.
53     }
54
55 void mutex_release(struct Mutex *m) {
56     // Acquire the spinlock in order to make changes.
57     acquire(&m->splock);
58
59     // Assert that the current thread actually owns the mutex
60     assert(m->owner == id_of_this_thread);
61
62     // Check if anyone is waiting.
63     m->owner = STAILQ_GET_HEAD(&m->waiters);
64
65     // If so, wake them up.
66     if (m->owner) {
67         sched_wakeone(&m->owner);
68         STAILQ_REMOVE_HEAD(&m->waiters, qlink);
69     }
70
71     // Release the internal spinlock
72     release(&m->splock);

```

Sep 26, 21 21:09

fair-mutex.c

Page 2/2

73 }