

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

} Mon.

2 CVs

A. Motivation: last time

B. API

&w

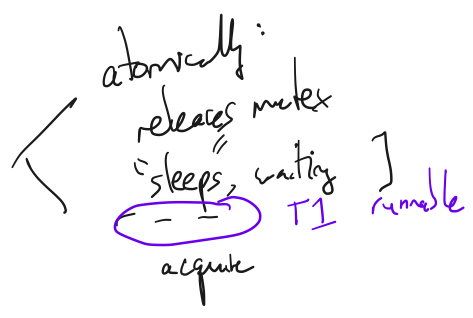
Cond w/

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

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

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

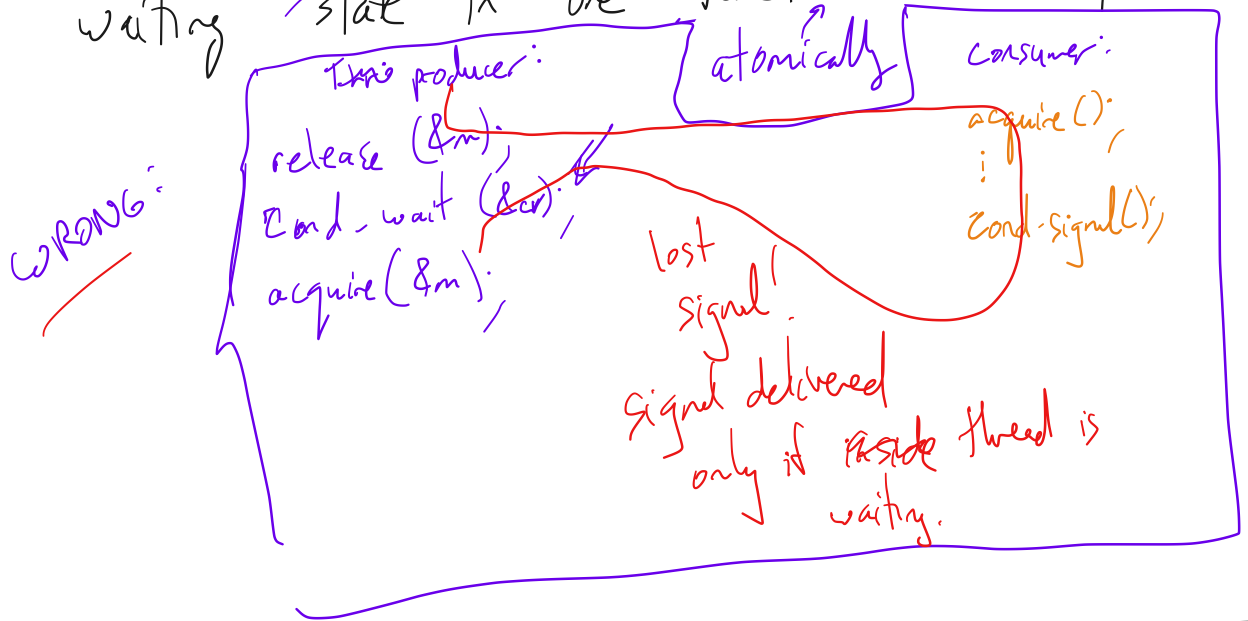
cond-broadcast (Cond* c, ...);



C. Important points

Must use "while" not "if"

Cond-wait() releases mutex and goes into waiting state in one function call; why?



3. Monitors + standards

Monitor: one mutex + 1 or more CVs.



M::f()
(acquire(&mutex);

M::g()
acquire(&mutex);

Mesa

release(&mutex);

CVS

~~semaphores~~

release(&mutex);

Java: synchronized

Comments

Mike D.

Rule: acquire/release at beg/end of a method/function

Rule: hold lock/mutex when doing CV operations wait(), signal(), broadcast();

~~Tony Hoare~~

British Hansen

semantics

wrong

~~mutex.release()
cv.broadcast();~~

Can we replace signal() ~ broadcast()?

Yes

Rule: if you are in wait(), must be prepared to be restarted ^(i.e. woken) at any time NOT JUST when another thread calls signal()/broadcast();




~~wait();~~

Ex: memory allocator

T2 alloc()

T3 alloc()

while (not enough mem) wait(&cv);



+1
alloc();

free();
broadcast();


RULE: NEVER use this pattern: if (something)
wait();
ALWAYS: while (something)
wait();

1 CS 202, Spring 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).

1. Protecting the linked list.....

```

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
    
```





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

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

```

25     Mutex mutex;
26
27     void producer (void *ignored) {
28         for (;;) {
29             /* next line produces an item and puts it in nextProduced */
30             nextProduced = means_of_production();
31
32             acquire(&mutex);
33             while (count == BUFFER_SIZE) {
34                 release(&mutex);
35                 yield(); /* or schedule() */
36                 acquire(&mutex);
37             }
38
39             buffer [in] = nextProduced;
40             in = (in + 1) % BUFFER_SIZE;
41             count++;
42             release(&mutex);
43         }
44     }
45
46     void consumer (void *ignored) {
47         for (;;) {
48
49             acquire(&mutex);
50             while (count == 0) {
51                 release(&mutex);
52                 yield(); /* or schedule() */
53                 acquire(&mutex);
54             }
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
    
```





busy waiting

```

70 2b. Producer/consumer [bounded buffer] with mutexes and condition variables
71
72
73 Mutex mutex;
74 Cond nonempty;
75 Cond nonfull;
76
77 void producer (void *ignored) {
78     for (;;) {
79         /* next line produces an item and puts it in nextProduced */
80         nextProduced = means_of_production();
81
82         acquire(&mutex);
83         while (count == BUFFER_SIZE)
84             cond_wait(&nonfull, &mutex);
85
86         buffer[in] = nextProduced;
87         in = (in + 1) % BUFFER_SIZE;
88         count++;
89         cond_signal(&nonempty, &mutex);
90         release(&mutex);
91     }
92 }
93
94 void consumer (void *ignored) {
95     for (;;) {
96
97         acquire(&mutex);
98         while (count == 0)
99             cond_wait(&nonempty, &mutex);
100
101         nextConsumed = buffer[out];
102         out = (out + 1) % BUFFER_SIZE;
103         count--;
104         cond_signal(&nonfull, &mutex);
105         release(&mutex);
106
107         /* next line abstractly consumes the item */
108         consume_item(nextConsumed);
109     }
110 }
111
112
113 Question: why does cond_wait need to both release the mutex and
114 sleep? Why not:
115
116     while (count == BUFFER_SIZE) {
117         release(&mutex);
118         cond_wait(&nonfull);
119         acquire(&mutex);
120     }
121

```



2 producers
1 consumer

if (count == BUFFER_SIZE)
c.wait();
count == BUFFER_SIZE

2 producers
1 consumer

T1 Producer
T2 Producer
T3 Consumer

This doesn't follow the advice

```

122 2c. Producer/consumer [bounded buffer] with semaphores
123
124 Semaphore mutex(1); /* mutex initialized to 1 */
125 Semaphore empty(BUFFER_SIZE); /* start with BUFFER_SIZE empty slots */
126 Semaphore full(0); /* 0 full slots */
127
128 void producer (void *ignored) {
129     for (;;) {
130         /* next line produces an item and puts it in nextProduced */
131         nextProduced = means_of_production();
132
133         /*
134          * next line diminishes the count of empty slots and
135          * waits if there are no empty slots
136          */
137         sem_down(&empty);
138         sem_down(&mutex); /* get exclusive access */
139
140         buffer[in] = nextProduced;
141         in = (in + 1) % BUFFER_SIZE;
142
143         sem_up(&mutex);
144         sem_up(&full); /* we just increased the # of full slots */
145     }
146 }
147
148 void consumer (void *ignored) {
149     for (;;) {
150
151         /*
152          * next line diminishes the count of full slots and
153          * waits if there are no full slots
154          */
155         sem_down(&full);
156         sem_down(&mutex);
157
158         nextConsumed = buffer[out];
159         out = (out + 1) % BUFFER_SIZE;
160
161         sem_up(&mutex);
162         sem_up(&empty); /* one further empty slot */
163
164         /* next line abstractly consumes the item */
165         consume_item(nextConsumed);
166     }
167 }
168
169 Semaphores *can* (not always) lead to elegant solutions (notice
170 that the code above is fewer lines than 2b) but they are much
171 harder to use.
172
173 The fundamental issue is that semaphores make implicit (counts,
174 conditions, etc.) what is probably best left explicit. Moreover,
175 they *also* implement mutual exclusion.
176
177 For this reason, you should not use semaphores. This example is
178 here mainly for completeness and so you know what a semaphore
179 is. But do not code with them. Solutions that use semaphores in
180 this course will receive no credit.

```

1 CS 202, Spring 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

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

```

class MyBuffer {
public:
  MyBuffer();
  ~MyBuffer();
  void Enqueue(Item);
  Item = Dequeue();
private:
  int count;
  int in;
  int out;
  Item buffer[BUFFER_SIZE];
  Mutex* mutex;
  Cond* nonempty;
  Cond* nonfull;
}

void
MyBuffer::MyBuffer()
{
  in = out = count = 0;
  mutex = new Mutex;
  nonempty = new Cond;
  nonfull = new Cond;
}


void
MyBuffer::Enqueue(Item item)
{
  mutex.acquire();
  while (count == BUFFER_SIZE)
    cond_wait(&nonfull, &mutex);

  buffer[in] = item;
  in = (in + 1) % BUFFER_SIZE;
  ++count;
  cond_signal(&nonempty, &mutex);
  mutex.release();
}

Item
MyBuffer::Dequeue()
{
  mutex.acquire();
  while (count == 0)
    cond_wait(&nonempty, &mutex);

  Item ret = buffer[out];
  out = (out + 1) % BUFFER_SIZE;
  --count;
  cond_signal(&nonfull, &mutex);
  mutex.release();
  return ret;
}

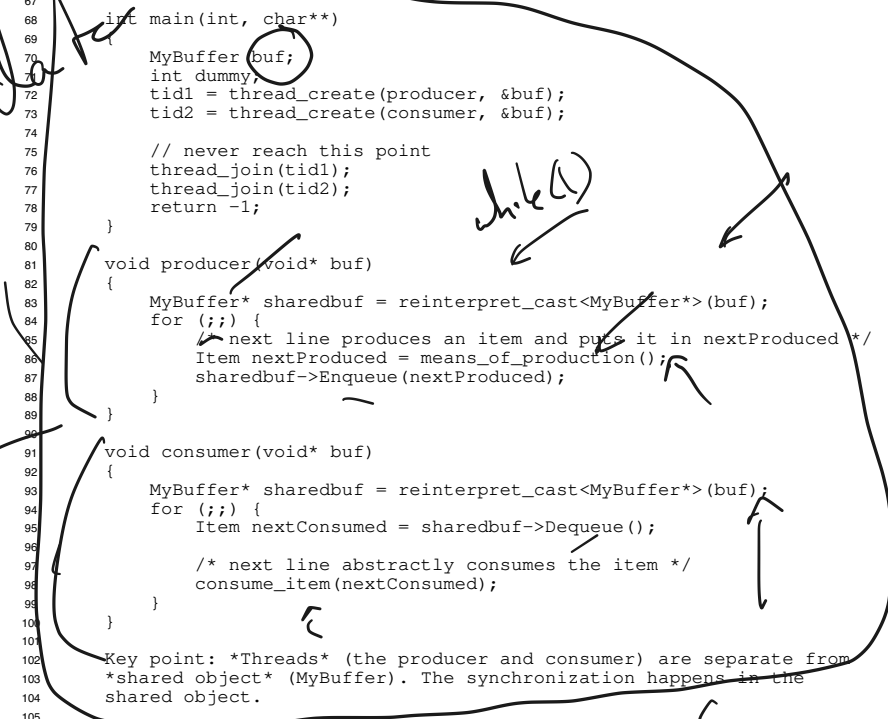
```



```

67 int main(int, char**)
68 {
69     MyBuffer buf;
70     int dummy;
71     tid1 = thread_create(producer, &buf);
72     tid2 = thread_create(consumer, &buf);
73
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

```



1. Getting started

- a. id wnts of memory.
- b. id shared chunks of stack
- c. write in code the main loop of the thread

2. write down synch. constraints on the solution

- id kind: mutual exclusion or scheduling
 or waits for producer

Feb 18, 21 2:30

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?

```

Feb 18, 21 2:30

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
221
222
223
224
225

```

QUESTIONS:

- A. There is a starvation problem here. What is it? (Readers can keep writers out if there is a steady stream of readers.)
- B. How could you use these shared locks to write a cleaner version of the code in the prior item? (Though note that the starvation properties would be different.)


```

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

```

```

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

```

Feb 18, 21 10:43

spinlock-mutex.txt

Page 3/3

95 3. Mutex implementation

96
97
98
99
100
101

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

Feb 18, 21 10:53

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
56 void mutex_release(struct Mutex *m) {
57     // Acquire the spinlock in order to make changes.
58     acquire(&m->splock);
59
60     // Assert that the current thread actually owns the mutex
61     assert(m->owner == id_of_this_thread);
62
63     // Check if anyone is waiting.
64     m->owner = STAILQ_GET_HEAD(&m->waiters);
65
66     // If so, wake them up.
67     if (m->owner) {
68         sched_wakeone(&m->owner);
69         STAILQ_REMOVE_HEAD(&m->waiters, qlink);
70     }
71
72     // Release the internal spinlock
73     release(&m->splock);

```

73 }