

- ☐ 1. Last time ONE HANDOUT
- ☐ 2. Practice with concurrent programming
- ☐ 3. Implementation of spinlocks, mutexes
- ☐ 4. Deadlock (Next time)

2 Practice

Example:

- workers interact with a database
- readers never modify
- writers read and modify
- single mutex would be too restrictive
- instead, want:
 - many readers at once OR
 - only one writer (and no readers)

Let's follow the advice:

a. units of concurrency?

b. shared chunks of state?

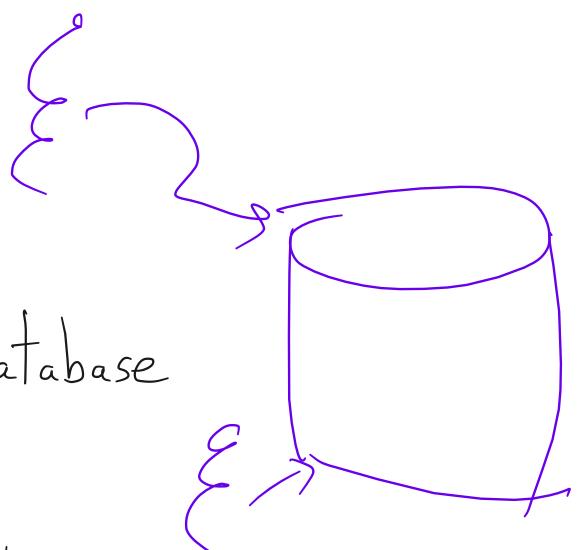
c. what does main function look like?

read()

check in ... wait until no writers

access_DB()

check out ... wake waiting writers, if any



write()

check in ... wait until no one else
access- DB()

check out ... wake up waiting readers
or writers

2. and 3. synchronization constraints and

synchronization objects

mutual exclusion of the state that allows us to
implement checkin() / checkout().

okToWrite writer: no readers or writers. \equiv safe to proceed if no
other thrs in DB

okToRead reader: no writers \equiv safe to proceed if no active writers

4. write the methods

```
int AR = 0; //active readers
```

```
int AW = 0; //active writers
```

```
int AR - 1; //writing readers
```

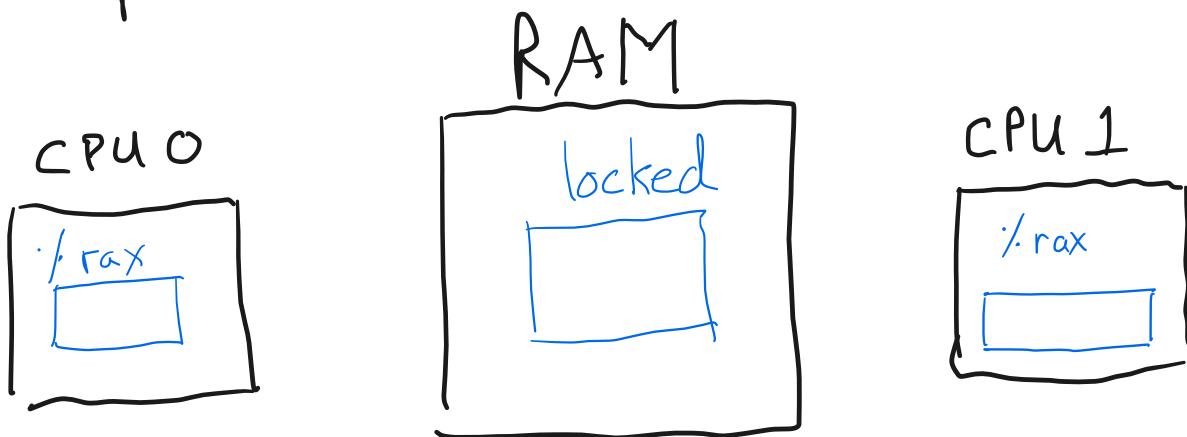
```
int NR = 0; // waiting readers  
int WW = 0; // waiting writers
```

3. Implementation of spinlocks and mutexes

high-level interface: lock() / unlock()

how to provide?

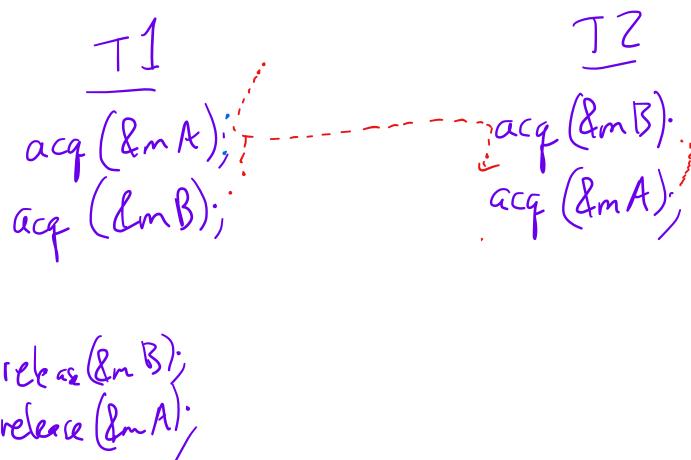
- (a) Peterson's algorithm → busy waiting, static bound
- (b) disable interrupts
- (c) spinlocks



(d) mutexes: spinlock + a queue

- textbook has an implementation
 - handout has another
-

4. Deadlock

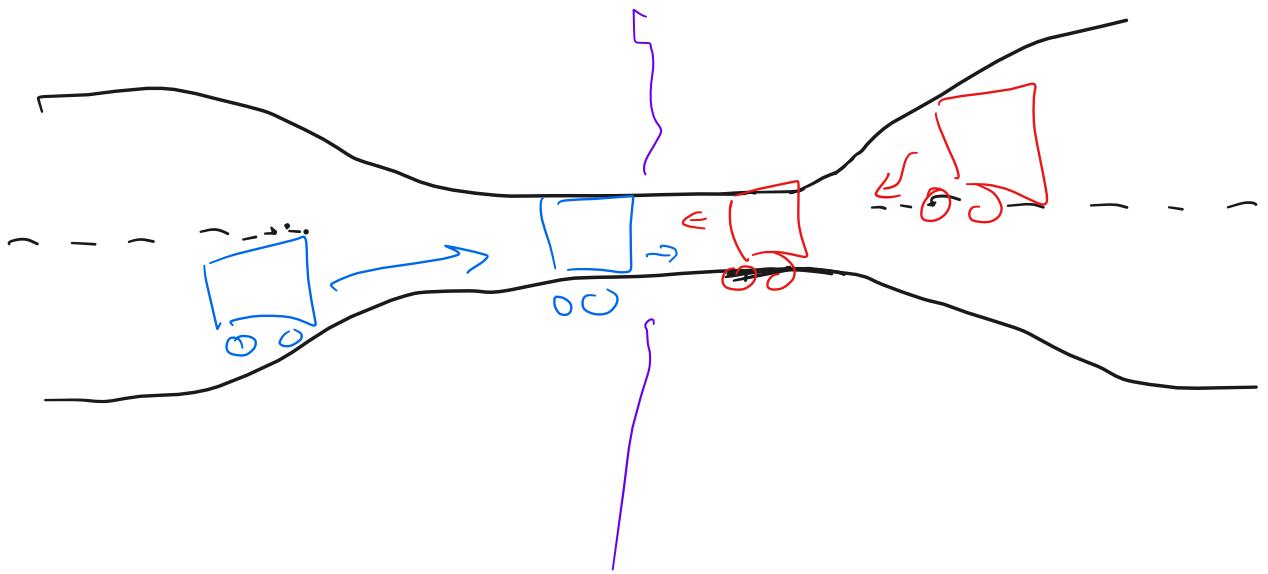


Happens when all four of these conditions are present:

- i. mutual exclusion
- ii. hold and wait
- iii. no pre-emption
- iv. circular wait

What can we do about deadlock?

- (a) ignore it
- (b) detect + recover
- (c) avoid algorithmically
- (d) negate one of the 4 conditions
- (e) static/dynamic detection tools



Feb 09, 25 23:12

handout05.txt

Page 1/4

```

1 CS 202
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).
7
8 1. The bounded buffer as a monitor
9
10 // This is pseudocode that is inspired by C++.
11 // Don't take it literally.
12
13 class MyBuffer {
14     public:
15         MyBuffer();
16         ~MyBuffer();
17         void Enqueue(Item);
18         Item = Dequeue();
19     private:
20         int count;
21         int in;
22         int out;
23         Item buffer[BUFFER_SIZE];
24         Mutex* mutex;
25         Cond* nonempty;
26         Cond* nonfull;
27     };
28
29 void
30 MyBuffer::MyBuffer()
31 {
32     in = out = count = 0;
33     mutex = new Mutex;
34     nonempty = new Cond;
35     nonfull = new Cond;
36 }
37
38 void
39 MyBuffer::Enqueue(Item item)
40 {
41     mutex.acquire();
42     while (count == BUFFER_SIZE)
43         cond_wait(&nonfull, &mutex);
44
45     buffer[in] = item;
46     in = (in + 1) % BUFFER_SIZE;
47     ++count;
48     cond_signal(&nonempty, &mutex);
49     mutex.release();
50 }
51
52 Item
53 MyBuffer::Dequeue()
54 {
55     mutex.acquire();
56     while (count == 0)
57         cond_wait(&nonempty, &mutex);
58
59     Item ret = buffer[out];
60     out = (out + 1) % BUFFER_SIZE;
61     --count;
62     cond_signal(&nonfull, &mutex);
63     mutex.release();
64     return ret;
65 }
66

```

Feb 09, 25 23:12

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

```

Feb 09, 2023 23:12

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     SAccess Data
126     doneRead(); // check self out of system
127 }
128
129 Database::startRead() {
130     acquire(&mutex);
131     while((AW + WR) > 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     SAccess Data
152     doneWrite(); // check out
153 }
154
155 Database::startWrite() {
156     acquire(&mutex);
157     while ((AW + AR) > 0) { // check if safe to write.
158         +WR // 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?

```

not enforced in this line!!

Feb 09, 25 23:12

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.)
```

Feb 12, 25 10:00

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

```

Feb 12, 25 10:00

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). To learn about a remediation of this
88 issue, search the Web for "MCS locks".
89     https://www.cs.rochester.edu/u/scott/papers/1991\_TOCS\_synch.pdf
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.
95

```

Feb 12, 25 10:00

spinlock-mutex.txt

Page 3/3

```

96 3. Mutex implementation
97
98     The intent of a mutex is to avoid busy waiting: if the lock is not
99     available, the locking thread is put to sleep, and tracked by a
100    queue in the mutex. The next page has an implementation.
101
102

```

mutex acq:

Spinlock acquire

critical section protected by spinlock

Spinlock rel

[either wait or continue]

*critical
section the
mutex user
is protecting*

mutex release:

Spinlock acquire

critical section protected by spinlock

Spinlock release

Feb 12, 25 9:50

fair-mutex.c

Page 1/1

```

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_swtch();
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 have held the mutex in lines 39-42
48         // (if we were context-switched out after the spinlock release()),
49         // followed by being run as a result of another thread's release of the
50         // mutex). But if that happens, it just means that we are
51         // context-switched out an "extra" time before proceeding.
52     }
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);
73 }

```

Feb 12, 25 9:50

deadlock.txt

Page 1/3

```

1  Deadlock examples
2
3  1. Simple deadlock example
4
5    T1:
6      acquire(mutexA);
7      acquire(mutexB);
8
9      // do some stuff
10
11     release(mutexB);
12     release(mutexA);
13
14    T2:
15      acquire(mutexB);
16      acquire(mutexA);
17
18      // do some stuff
19
20     release(mutexA);
21     release(mutexB);
22

```

Feb 12, 25 9:50

deadlock.txt

Page 2/3

```

23  2. More subtle deadlock example
24
25  Let M be a monitor (shared object with methods protected by mutex)
26  Let N be another monitor
27
28  class M {
29    private:
30      Mutex mutex_m;
31
32      // instance of monitor N
33      N another_monitor;
34
35      // Assumption: no other objects in the system hold a pointer
36      // to our "another_monitor"
37
38    public:
39      M();
40      ~M();
41      void methodA();
42      void methodB();
43  };
44
45  class N {
46    private:
47      Mutex mutex_n;
48      Cond cond_n;
49      int available;
50
51    public:
52      N();
53      ~N();
54      void* alloc(int nwanted);
55      void free(void*);
56  }
57
58  int
59  N::alloc(int nwanted) {
60    acquire(&mutex_n);
61    while (available < nwanted) {
62      wait(&cond_n, &mutex_n);
63    }
64
65    // peel off the memory
66
67    available -= nwanted;
68    release(&mutex_n);
69  }
70
71  void
72  N::free(void* returning_mem) {
73
74    acquire(&mutex_n);
75
76    // put the memory back
77
78    available += returning_mem;
79
80    broadcast(&cond_n, &mutex_n);
81
82    release(&mutex_n);
83  }
84

```

Feb 12, 25 9:50

deadlock.txt

Page 3/3

```

85     void
86     M::methodA() {
87
88         acquire(&mutex_m);
89
90         void* new_mem = another_monitor.alloc(int nbytes);
91
92         // do a bunch of stuff using this nice
93         // chunk of memory n allocated for us
94
95         release(&mutex_m);
96     }
97
98     void
99     M::methodB() {
100
101         acquire(&mutex_m);
102
103         // do a bunch of stuff
104
105         another_monitor.free(some_pointer);
106
107         release(&mutex_m);
108     }
109
110     QUESTION: What's the problem?

```

