

# CS202 (003): Operating Systems

## Concurrency III

Instructor: Jocelyn Chen

# Last time

# Condition Variables

Warning: Condition Variable is not really a Variable!

```
void cond_init(Cond *cond, ...);  
void cond_wait(Cond *cond, Mutex *mutex);  
void cond_signal(Cond *cond);  
void cond_broadcast(Cond *cond);
```

```
mutex_lock(&mutex);  
while (!condition_is_met) {  
    cond_wait(&cond, &mutex);  
}  
// Modify shared state  
mutex_unlock(&mutex);
```

Why is this a while?

# Condition Variables

Warning: Condition Variable is not really a Variable!

```
void cond_init(Cond *cond, ...);  
void cond_wait(Cond *cond, Mutex *mutex);  
void cond_signal(Cond *cond);  
void cond_broadcast(Cond *cond);
```

```
mutex_lock(&mutex);  
while (!condition_is_met) {  
    cond_wait(&cond, &mutex);  
}  
// Modify shared state  
mutex_unlock(&mutex);
```

This **MUST** be a while!

# More hypothetical questions...

Why do `cond_wait` releases the mutexes and goes into the waiting state in one function call (see panel 2b of handout 04)?

If those two steps were separate, could get stuck waiting.

```
Producer: while (count == BUFFER_SIZE)
Producer: release()
Consumer: acquire()
Consumer: .....
Consumer: cond_signal(&nonfull)
Producer: cond_wait(&nonfull)
```

**Producer never hears the signal!**

# More hypothetical questions...

Can we replace SIGNAL with BROADCAST, and preserve correctness\*?

Yes, but it might hurt performance

Since while() checks the invariant,  
Only thread satisfying the invariant will make progress

=> this does not affect correctness

But we make needlessly wakeup of threads

=> this might hurt performance

# More hypothetical questions...

Can we replace BROADCAST with SIGNAL, and preserve correctness\*?

No race conditions, but may never make progress

# Monitor: Mutex + Conditional Variables (but in OOP)

**All** method calls of a class are protected by a **mutex**

Synchronization happens with condition variables whose associated mutex is the **mutex that protects the method calls**

“Monitor” can be used to refer to either a *programming convention* or a *method in certain programming languages*\*

\* <https://docs.oracle.com/javase/tutorial/essential/concurrency/syncmeth.html>



# What does monitor enable us to do?

**Encapsulation!**

Separation of program logic inside threads from the shared object

The monitor handles all synchronization internally so threads don't need to worry about locking, unlocking or conditional signaling

**Look at the first page of handout05!**

## Producer/Consumer w/ Monitor

```
int main(int, char**)
{
    MyBuffer buf;
    int dummy;
    tid1 = thread_create(producer, &buf);
    tid2 = thread_create(consumer, &buf);
}

void producer(void* buf)
{
    MyBuffer* sharedbuf = reinterpret_cast<MyBuffer*>(buf);
    for (;;) {
        Item nextProduced = means_of_production();
        sharedbuf->Enqueue(nextProduced);
    }
}

void consumer(void* buf)
{
    MyBuffer* sharedbuf = reinterpret_cast<MyBuffer*>(buf);
    for (;;) {
        Item nextConsumed = sharedbuf->Dequeue();
        consume_item(nextConsumed);
    }
}
```

## Producer/Consumer w/ Mutex & CV

Mutex mutex;

```
void producer (void *ignored) {
    for (;;) {
        nextProduced = means_of_production();

        acquire(&mutex);
        while (count == BUFFER_SIZE) {
            release(&mutex);
            yield(); /* or schedule() */
            acquire(&mutex);
        }
        buffer [in] = nextProduced;
        in = (in + 1) % BUFFER_SIZE;
        count++;
        release(&mutex);
    }
}
```

```
void consumer (void *ignored) {
    for (;;) {
        acquire(&mutex);
        while (count == 0) {
            release(&mutex);
            yield(); /* or schedule() */
            acquire(&mutex);
        }
        nextConsumed = buffer[out];
        out = (out + 1) % BUFFER_SIZE;
        count--;
        release(&mutex);

        consume_item(nextConsumed);
    }
}
```

# Semaphores: Mutex + Conditional Variables (but more general)

```
#include <semaphore.h>
sem_t s;
sem_init(&s, 0, 1);

int sem_wait(sem_t *s) {
    decrement the value of semaphore s by one
    wait if value of semaphore s is negative
}

int sem_post(sem_t *s) {
    increment the value of semaphore s by one
    if there are one or more threads waiting, wake one
}

sem_wait(&m);
// critical section here
sem_post(&m);
```

# Semaphores: Mutex + Conditional Variables (but more general)

Semaphores manage a count, mutex+CV do not inherently do this

Semaphores can allow multiple threads access, unlike a basic mutex

Semaphores can be used for locking, but can also be used for other purpose

**DO NOT USE SEMAPHORE IN THIS CLASS!**

# Monitor: Mutex + Conditional Variables

**All** method calls are protected by a **mutex**

Synchronization happens with condition variables whose associated mutex is the **mutex that protects the method calls**

“Monitor” can be used to refer to either a *programming convention* or a *method in certain programming languages*\*

**Please follow these conventions on Lab 3!**

\* <https://docs.oracle.com/javase/tutorial/essential/concurrency/syncmeth.html>

# Mike Dahlin's "Programming with Threads"

**You are required to follow this document  
(although we don't code in Java)**

**You will lose a lot of points if you don't follow  
(in labs and exams)**

**Do not program concurrency in other ways  
unless you are a concurrency guru**

# Standards for Programming w/ Threads

**Rule I:** acquire/release at beginning/end of methods

**Rule II:** hold lock when doing condition variable operations

**Rule III:** a thread that is in wait() must be prepared to be restarted at any time, not just when another thread calls "signal()"

**Rule IV:** don't call sleep()

# Advice for concurrent programming

**Top-level piece of advice: SAFETY FIRST**

Locking at coarse grain is easiest to get right, so do that

Don't worry about performance at first

Don't view deadlock as a disaster

**MAKE SURE YOU PROGRAM NEVER DOES THE WRONG THING**



# Advice for concurrent programming

## Getting started

1. Identify unit of concurrency
2. Identify chunks of state
3. write down high-level main loop of each thread

**Write down the synchronization constraints, and the type**

**Create a lock or CV for each constraint**

**Implement the methods, using the locks and CVs**

```

1 CS 202, Fall 2024
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

```

```

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

```

```

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

```

- workers interact with a database
- readers never modify
- writers read and modify
- allow:
  - many readers at once
  - OR
  - only one writer (no reader)

Unit of concurrency?

Shared chunks of state?

What does main function look like?

Synchronization constraints and objects?