

Monitors con't / Reader/Writer Locks /  
Deadlock (start)

# Changelog

Changes made in this version not seen in first lecture:

2 October: bounded buffer producer/consumer: condition should have been `buffer.full`, not `!buffer.full`

2 October: bounded buffer producer/consumer: signalling only when `buffer.size = capacity - 1` doesn't work correctly

2 October: writer-priority reader/writer lock: condition for signaling writer should have been `waiting_writers != 0`

2 October: simulation of reader/writer lock: correct readers being decremented too early

2 October: simulation of reader/writer lock: condition for signaling writer should have been `waiting_writers != 0`

2 October: rwlock exercise solution?: add "if (need to wait)"

2 October: rwlock exercise solution?: remove extraneous writer IDs

2 October: monitor exercise: make entire code fit on slide

2 October: monitors with semaphore: clarify on slide that this is to

# last time

barriers — wait for everyone else

counting semaphores

- track number of something

- wait if not any

monitors: mutex + condition variables

condition variable: wait and signal/broadcast

- pattern: loop of waiting (spurious wakeup)

- associated mutex lock: check if need to wait safely

producer/consumer solution with semaphores/monitors

- producer: add to queue, wait if full

- consumer: remove from queue, wait if empty

# life HW

life HW is out

checkpoint (Friday): use POSIX barriers

final (week from Friday): write your own barriers

questions?

# unbounded buffer producer/consumer

```
pthread_mutex_t lock;  
pthread_cond_t data_ready;  
UnboundedQueue buffer;
```

```
Produce(item) {  
    pthread_mutex_lock(&lock);  
    buffer.enqueue(item);  
    pthread_cond_signal(&data_ready);  
    pthread_mutex_unlock(&lock);  
}
```

```
Consume() {  
    pthread_mutex_lock(&lock);  
    while (buffer.empty()) {  
        pthread_cond_wait(&data_ready, &lock);  
    }  
    item = buffer.dequeue();  
    pthread_mutex_unlock(&lock);  
    return item;  
}
```

# unbounded buffer producer/consumer

```
pthread_mutex_t lock;  
pthread_cond_t data_ready;  
UnboundedQueue buffer;
```

```
Produce(item) {  
    pthread_mutex_lock(&lock);  
    buffer.enqueue(item);  
    pthread_cond_signal(&data_ready);  
    pthread_mutex_unlock(&lock);  
}
```

```
Consume() {  
    pthread_mutex_lock(&lock);  
    while (buffer.empty()) {  
        pthread_cond_wait(&data_ready, &lock);  
    }  
    item = buffer.dequeue();  
    pthread_mutex_unlock(&lock);  
    return item;  
}
```

rule: never touch buffer  
without acquiring lock

otherwise: what if two threads  
simultaneously en/dequeue?  
(both use same array/linked list entry?)  
(both reallocate array?)

# unbounded buffer producer/consumer

```
pthread_mutex_t lock;  
pthread_cond_t data_ready;  
UnboundedQueue buffer;
```

```
Produce(item) {  
    pthread_mutex_lock(&lock);  
    buffer.enqueue(item);  
    pthread_cond_signal(&data_ready);  
    pthread_mutex_unlock(&lock);  
}
```

```
Consume() {  
    pthread_mutex_lock(&lock);  
    while (buffer.empty()) {  
        pthread_cond_wait(&data_ready, &lock);  
    }  
    item = buffer.dequeue();  
    pthread_mutex_unlock(&lock);  
    return item;  
}
```

check if empty  
if so, dequeue

okay because have lock

other threads **cannot** dequeue here

# unbounded buffer producer/consumer

```
pthread_mutex_t lock;  
pthread_cond_t data_ready;  
UnboundedQueue buffer;
```

```
Produce(item) {  
    pthread_mutex_lock(&lock);  
    buffer.enqueue(item);  
    pthread_cond_signal(&data_ready);  
    pthread_mutex_unlock(&lock);  
}
```

wake one Consume thread  
*if any are waiting*

```
Consume() {  
    pthread_mutex_lock(&lock);  
    while (buffer.empty()) {  
        pthread_cond_wait(&data_ready, &lock);  
    }  
    item = buffer.dequeue();  
    pthread_mutex_unlock(&lock);  
    return item;  
}
```



# unbounded buffer producer/consumer

```
pthread_mutex_t lock;  
pthread_cond_t data_ready;  
UnboundedQueue buffer;
```

```
Produce(item) {  
    pthread_mutex_lock(&lock);  
    buffer.enqueue(item);  
    pthread_cond_signal(&data_ready);  
    pthread_mutex_unlock(&lock);  
}
```

```
Consume() {  
    pthread_mutex_lock(&lock);  
    while (buffer.empty()) {  
        pthread_cond_wait(&data_ready, &lock);  
    }  
    item = buffer.dequeue();  
    pthread_mutex_unlock(&lock);  
    return item;  
}
```

Thread 1

Produce()
...lock
...enqueue
...signal
...unlock

Thread 2

Consume()
...lock
...empty? no
...dequeue
...unlock
return

0 iterations: Produce() called before Consume()  
1 iteration: Produce() signalled, probably  
2+ iterations: spurious wakeup or ...?

# unbounded buffer producer/consumer

```
pthread_mutex_t lock;  
pthread_cond_t data_ready;  
UnboundedQueue buffer;
```

```
Produce(item) {  
    pthread_mutex_lock(&lock);  
    buffer.enqueue(item);  
    pthread_cond_signal(&data_ready);  
    pthread_mutex_unlock(&lock);  
}
```

```
Consume() {  
    pthread_mutex_lock(&lock);  
    while (buffer.empty()) {  
        pthread_cond_wait(&data_ready, &lock);  
    }  
    item = buffer.dequeue();  
    pthread_mutex_unlock(&lock);  
    return item;  
}
```

Thread 1

Thread 2

	Consume()
	...lock
	...empty? yes
	...unlock/start wait
Produce()	waiting for data_ready
...lock	
...enqueue	
...signal	stop wait
...unlock	lock
	...empty? no
	...dequeue
	...unlock
	return

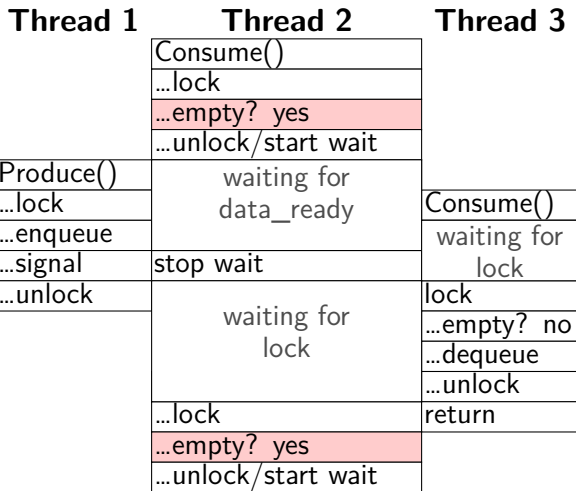
0 iterations: Produce() called before Consume()  
1 iteration: Produce() signalled, probably  
2+ iterations: spurious wakeup or ...?

# unbounded buffer producer/consumer

```
pthread_mutex_t lock;  
pthread_cond_t data_ready;  
UnboundedQueue buffer;
```

```
Produce(item) {  
    pthread_mutex_lock(&lock);  
    buffer.enqueue(item);  
    pthread_cond_signal(&data_ready);  
    pthread_mutex_unlock(&lock);  
}
```

```
Consume() {  
    pthread_mutex_lock(&lock);  
    while (buffer.empty()) {  
        pthread_cond_wait(&data_ready, &lock);  
    }  
    item = buffer.dequeue();  
    pthread_mutex_unlock(&lock);  
    return item;  
}
```



0 iterations: Produce() called before Consume()  
1 iteration: Produce() signalled, probably  
2+ iterations: spurious wakeup or ...?

# unbounded buffer producer/consumer

```
pthread_mutex_t lock;  
pthread_cond_t data_ready;
```

in pthreads: signalled thread not guaranteed to hold lock next

alternate design: signalled thread gets lock next called "Hoare scheduling" not done by pthreads, Java, ...

```
pthread_mutex_lock(&lock);  
while (buffer.empty()) {  
    pthread_cond_wait(&data_ready, &lock);  
}  
item = buffer.dequeue();  
pthread_mutex_unlock(&lock);  
return item;
```

Thread 1

```
Produce()  
...lock  
...enqueue  
...signal  
...unlock
```

Thread 2

```
Consume()  
...lock  
...empty? yes  
...unlock/start wait  
waiting for data_ready  
stop wait  
waiting for lock  
...lock  
...empty? yes  
...unlock/start wait
```

Thread 3

```
Consume()  
waiting for lock  
lock  
...empty? no  
...dequeue  
...unlock  
return
```

0 iterations: Produce() called before Consume()  
1 iteration: Produce() signalled, probably  
2+ iterations: spurious wakeup or ...?

# bounded buffer producer/consumer

```
pthread_mutex_t lock;  
pthread_cond_t data_ready; pthread_cond_t space_ready;  
BoundedQueue buffer;
```

```
Produce(item) {  
    pthread_mutex_lock(&lock);  
    while (buffer.full()) { pthread_cond_wait(&space_ready, &lock); }  
    buffer.enqueue(item);  
    pthread_cond_signal(&data_ready);  
    pthread_mutex_unlock(&lock);  
}
```

```
Consume() {  
    pthread_mutex_lock(&lock);  
    while (buffer.empty()) {  
        pthread_cond_wait(&data_ready, &lock);  
    }  
    item = buffer.dequeue();  
    pthread_cond_signal(&space_ready);  
    pthread_mutex_unlock(&lock);  
    return item;  
}
```

# bounded buffer producer/consumer

```
pthread_mutex_t lock;  
pthread_cond_t data_ready; pthread_cond_t space_ready;  
BoundedQueue buffer;
```

```
Produce(item) {  
    pthread_mutex_lock(&lock);  
    while (buffer.full()) { pthread_cond_wait(&space_ready, &lock); }  
    buffer.enqueue(item);  
    pthread_cond_signal(&data_ready);  
    pthread_mutex_unlock(&lock);  
}
```

```
Consume() {  
    pthread_mutex_lock(&lock);  
    while (buffer.empty()) {  
        pthread_cond_wait(&data_ready, &lock);  
    }  
    item = buffer.dequeue();  
    pthread_cond_signal(&space_ready);  
    pthread_mutex_unlock(&lock);  
    return item;  
}
```

# bounded buffer producer/consumer

```
pthread_mutex_t lock;  
pthread_cond_t data_ready; pthread_cond_t space_ready;  
BoundedQueue buffer;
```

```
Produce(item) {  
    pthread_mutex_lock(&lock);  
    while (buffer.full()) { pthread_cond_wait(&space_ready, &lock); }  
    buffer.enqueue(item);  
    pthread_cond_signal(&data_ready);  
    pthread_mutex_unlock(&lock);  
}
```

correct (but slow?) to replace with:

```
pthread_cond_broadcast(&space_ready);  
(just more "spurious wakeups")  
pthread_cond_wait(&data_ready, &lock);  
}  
item = buffer.dequeue();  
pthread_cond_signal(&space_ready);  
pthread_mutex_unlock(&lock);  
return item;
```

# bounded buffer producer/consumer

```
pthread_mutex_t lock;  
pthread_cond_t data_ready; pthread_cond_t space_ready;  
BoundedQueue buffer;
```

```
Produce(item) {  
    pthread_mutex_lock(&lock);  
    while (buffer.full()) { pthread_cond_wait(&space_ready, &lock); }  
    buffer.enqueue(item);  
    pthread_cond_signal(&data_ready);  
    pthread_mutex_unlock(&lock);  
}
```

```
Consume() {  
    pthread_mutex_lock(&lock);  
    while (buffer.empty()) {  
        pthread_cond_wait(&data_ready, &lock);  
    }  
    item = buffer.dequeue();  
    pthread_cond_signal(&space_ready);  
    pthread_mutex_unlock(&lock);  
    return item;  
}
```

correct but slow to replace  
data\_ready and space\_ready  
with 'combined' condvar ready  
and use broadcast  
(just more "spurious wakeups")



# monitor pattern

```
pthread_mutex_lock(&lock);
while (!condition A) {
    pthread_cond_wait(&condvar_for_A, &lock);
}
... /* manipulate shared data, changing other conditions */
if (set condition B) {
    pthread_cond_broadcast(&condvar_for_B);
    /* or signal, if only one thread cares */
}
if (set condition C) {
    pthread_cond_broadcast(&condvar_for_C);
    /* or signal, if only one thread cares */
}
...
pthread_mutex_unlock(&lock)
```

# monitors rules of thumb

never touch shared data without holding the lock

keep lock held for **entire operation**:

verifying condition (e.g. buffer not full) *up to and including*  
manipulating data (e.g. adding to buffer)

create condvar for every kind of scenario waited for

always write **loop** calling `cond_wait` to wait for condition X

broadcast/signal condition variable **every time you change X**

# monitors rules of thumb

never touch shared data without holding the lock

keep lock held for **entire operation**:

verifying condition (e.g. buffer not full) *up to and including*  
manipulating data (e.g. adding to buffer)

create condvar for every kind of scenario waited for

always write **loop** calling `cond_wait` to wait for condition X

broadcast/signal condition variable **every time you change X**

correct but slow to...

broadcast when just signal would work

broadcast or signal when nothing changed

use one condvar for multiple conditions

# monitor exercise (1)

suppose we want producer/consumer, but...

but change to `ConsumeTwo()` which returns a **pair of values**  
and don't want two calls to `ConsumeTwo()` to wait...  
with each getting one item

what should we change below?

```
pthread_mutex_t lock;  
pthread_cond_t data_ready;  
UnboundedQueue buffer;
```

```
Produce(item) {  
    pthread_mutex_lock(&lock);  
    buffer.enqueue(item);  
    pthread_cond_signal(&data_ready);  
    pthread_mutex_unlock(&lock);  
}
```

```
Consume() {  
    pthread_mutex_lock(&lock);  
    while (buffer.empty()) {  
        pthread_cond_wait(&data_ready,  
                          &lock);  
    }  
    item = buffer.dequeue();  
    pthread_mutex_unlock(&lock);  
    return item;  
}
```

# building semaphore with monitors

```
pthread_mutex_t lock;
```

lock to protect shared state

# building semaphore with monitors

```
pthread_mutex_t lock;  
unsigned int count;
```

lock to protect shared state

shared state: semaphore tracks a count

# building semaphore with monitors

```
pthread_mutex_t lock;  
unsigned int count;
```

```
/* condition, signal when becomes count > 0 */  
pthread_cond_t count_is_positive_cv;
```

lock to protect shared state

shared state: semaphore tracks a count

add cond var for each reason we wait

semaphore: wait for count to become positive (for down)

# building semaphore with monitors

```
pthread_mutex_t lock;  
unsigned int count;  
/* condition, signal when becomes count > 0 */  
pthread_cond_t count_is_positive_cv;  
void down() {  
    pthread_mutex_lock(&lock);  
    while (!(count > 0)) {  
        pthread_cond_wait(  
            &count_is_positive_cv,  
            &lock);  
    }  
    count -= 1;  
    pthread_mutex_unlock(&lock);  
}
```

lock to protect shared state

shared state: semaphore tracks a count

add cond var for each reason we wait

semaphore: wait for count to become positive (for down)

**wait** using condvar; broadcast/signal when condition changes



# building semaphore with monitors

```
pthread_mutex_t lock;
unsigned int count;
/* condition, signal when becomes count > 0 */
pthread_cond_t count_is_positive_cv;
void down() {
    pthread_mutex_lock(&lock);
    while (!(count > 0)) {
        pthread_cond_wait(
            &count_is_positive_cv,
            &lock);
    }
    count -= 1;
    pthread_mutex_unlock(&lock);
}
```

```
void up() {
    pthread_mutex_lock(&lock);
    count += 1;
    if (count == 1) { /* became > 0 */
        pthread_cond_signal(
            &count_is_positive_cv
        );
    }
    pthread_mutex_unlock(&lock);
}
```

lock to protect shared state

shared state: semaphore tracks a count

add cond var for each reason we wait

semaphore: wait for count to become positive (for down)

wait using condvar; **broadcast/signal** when condition changes

# monitors with semaphores: locks

```
sem_t semaphore; // initial value 1
```

```
Lock() {  
    sem_wait(&semaphore);  
}
```

```
Unlock() {  
    sem_post(&semaphore);  
}
```

# monitors with semaphores: cvs (attempt 1)

condition variables are more challenging

start with only wait/signal:

```
sem_t threads_to_wakeup; // initially 0
Wait(Lock lock) {
    lock.Unlock();
    sem_wait(&threads_to_wakeup);
    lock.Lock();
}
Signal() {
    sem_post(&threads_to_wakeup);
}
```

# monitors with semaphores: cvs (attempt 1)

condition variables are more challenging

start with only wait/signal:

```
sem_t threads_to_wakeup; // initially 0
Wait(Lock lock) {
    lock.Unlock();
    sem_wait(&threads_to_wakeup);
    lock.Lock();
}
Signal() {
    sem_post(&threads_to_wakeup);
}
```

annoying: signal wakes up non-waiting threads (in the far future)

# monitors with semaphores: cvs (attempt 2)

condition variables are more challenging

start with only wait/signal:

```
sem_t private_lock; // initially 1
int num_waiters;
sem_t threads_to_wakeup; // initially 0
Wait(Lock lock) {
    sem_wait(&private_lock);
    ++num_waiters;
    sem_post(&private_lock);
    lock.Unlock();
    sem_wait(&threads_to_wakeup);
    lock.Lock();
}
```

```
Signal() {
    sem_wait(&private_lock);
    if (num_waiters > 0) {
        sem_post(&threads_to_wakeup);
        --num_waiters;
    }
    sem_post(&private_lock);
}
```

# monitors with semaphores: cvs (attempt 2)

condition variables are more challenging

start with only wait/signal:

```
sem_t private_lock; // initially 1
int num_waiters;
sem_t threads_to_wakeup; // initially 0
Wait(Lock lock) {
    sem_wait(&private_lock);
    ++num_waiters;
    sem_post(&private_lock);
    lock.Unlock();
    sem_wait(&threads_to_wakeup);
    lock.Lock();
}
```

```
Signal() {
    sem_wait(&private_lock);
    if (num_waiters > 0) {
        sem_post(&threads_to_wakeup);
        --num_waiters;
    }
    sem_post(&private_lock);
}
```

but what if we want to guarantee threads woken up in order?

# monitors with semaphores: cvs (attempt 3)

if we want to make sure threads woken up **in order**

```
ThreadSafeQueue<sem_t> waiters;
Wait(Lock lock) {
    sem_t private_semaphore;
    ... /* init semaphore
         with count 0 */
    waiters.Enqueue(&semaphore);
    lock.Unlock();
    sem_post(private_semaphore);
    lock.Lock();
}

Signal() {
    sem_t *next = waiters.DequeueOrNull();
    if (next != NULL) {
        sem_post(next);
    }
}
```

# monitors with semaphores: cvs (attempt 3)

if we want to make sure threads woken up **in order**

```
ThreadSafeQueue<sem_t> waiters;
Wait(Lock lock) {
    sem_t private_semaphore;
    ... /* init semaphore
         with count 0 */
    waiters.Enqueue(&semaphore);
    lock.Unlock();
    sem_post(private_semaphore);
    lock.Lock();
}

Signal() {
    sem_t *next = waiters.DequeueOrNull();
    if (next != NULL) {
        sem_post(next);
    }
}
```

(but now implement queue with semaphores...)



# reader/writer problem

some shared data

only one thread modifying (read+write) at a time

read-only access **from multiple threads** is safe

# reader/writer problem

some shared data

only one thread modifying (read+write) at a time

read-only access **from multiple threads** is safe

could use lock — but doesn't allow multiple readers

# reader/writer locks

abstraction: lock that distinguishes readers/writers

operations:

- read lock: wait until no writers

- read unlock: stop being registered as reader

- write lock: wait until no readers and no writers

- write unlock: stop being registered as writer

# reader/writer locks

abstraction: lock that distinguishes readers/writers

operations:

- read lock: wait until no writers

- read unlock: stop being registered as reader

- write lock: wait until **no readers and no writers**

- write unlock: stop being registered as writer

# pthread\_rwlock\_t

```
pthread_rwlock_t rwlock;  
pthread_rwlock_init(&rwlock, NULL /* attributes */);  
...  
    pthread_rwlock_rdlock(&rwlock);  
    ... /* read shared data */  
    pthread_rwlock_unlock(&rwlock);  
  
    pthread_rwlock_wrlock(&rwlock);  
    ... /* read+write shared data */  
    pthread_rwlock_unlock(&rwlock);  
  
...  
pthread_rwlock_destroy(&rwlock);
```

# rwlocks with monitors (attempt 1)

```
mutex_t lock;
```

lock to protect shared state

# rwlocks with monitors (attempt 1)

```
mutex_t lock;
```

```
unsigned int readers, writers;
```

state: number of active readers, writers

# rwlocks with monitors (attempt 1)

```
mutex_t lock;  
unsigned int readers, writers;
```

```
/* condition, signal when writers becomes 0 */  
cond_t ok_to_read_cv;  
/* condition, signal when readers + writers becomes 0 */  
cond_t ok_to_write_cv;
```

conditions to wait for (no readers or writers, no writers)



# rwlocks with monitors (attempt 1)

```
mutex_t lock;
unsigned int readers, writers;
/* condition, signal when writers becomes 0 */
cond_t ok_to_read_cv;
/* condition, signal when readers + writers becomes 0 */
cond_t ok_to_write_cv;
```

```
ReadLock() {
    mutex_lock(&lock);
    while (writers != 0) {
        cond_wait(&ok_to_read_cv, &lock);
    }
    ++readers;
    mutex_unlock(&lock);
}
```

```
ReadUnlock() {
    mutex_lock(&lock);
    --readers;
    if (readers == 0) {
        cond_signal(&ok_to_write_cv);
    }
    mutex_unlock(&lock);
}
```

```
WriteLock() {
    mutex_lock(&lock);
    while (readers + writers != 0) {
        cond_wait(&ok_to_write_cv);
    }
    ++writers;
    mutex_unlock(&lock);
}
```

```
WriteUnlock() {
    mutex_lock(&lock);
    --writers;
    cond_signal(&ok_to_write_cv);
    cond_broadcast(&ok_to_read_cv);
    mutex_unlock(&lock);
}
```

~~broadcast~~ — wakeup all readers when no writers

# rwlocks with monitors (attempt 1)

```
mutex_t lock;
unsigned int readers, writers;
/* condition, signal when writers becomes 0 */
cond_t ok_to_read_cv;
/* condition, signal when readers + writers becomes 0 */
cond_t ok_to_write_cv;
ReadLock() {
    mutex_lock(&lock);
    while (writers != 0) {
        cond_wait(&ok_to_read_cv, &lock);
    }
    ++readers;
    mutex_unlock(&lock);
}
WriteLock() {
    mutex_lock(&lock);
    while (readers + writers != 0) {
        cond_wait(&ok_to_write_cv);
    }
    ++writers;
    mutex_unlock(&lock);
}
ReadUnlock() {
    mutex_lock(&lock);
    --readers;
    if (readers == 0) {
        cond_signal(&ok_to_write_cv);
    }
    mutex_unlock(&lock);
}
WriteUnlock() {
    mutex_lock(&lock);
    --writers;
    cond_signal(&ok_to_write_cv);
    cond_broadcast(&ok_to_read_cv);
    mutex_unlock(&lock);
}
```

wakeup a single writer when no readers or writers

# rwlocks with monitors (attempt 1)

```
mutex_t lock;
unsigned int readers, writers;
/* condition, signal when writers becomes 0 */
cond_t ok_to_read_cv;
/* condition, signal when readers + writers becomes 0 */
cond_t ok_to_write_cv;
ReadLock() {
    mutex_lock(&lock);
    while (writers != 0) {
        cond_wait(&ok_to_read_cv, &lock);
    }
    ++readers;
    mutex_unlock(&lock);
}
WriteLock() {
    mutex_lock(&lock);
    while (readers + writers != 0) {
        cond_wait(&ok_to_write_cv);
    }
    ++writers;
    mutex_unlock(&lock);
}
ReadUnlock() {
    mutex_lock(&lock);
    --readers;
    if (readers == 0) {
        cond_signal(&ok_to_write_cv);
    }
    mutex_unlock(&lock);
}
WriteUnlock() {
    mutex_lock(&lock);
    --writers;
    cond_signal(&ok_to_write_cv);
    cond_broadcast(&ok_to_read_cv);
    mutex_unlock(&lock);
}
```

problem: wakeup readers first or writer first?

this solution: wake them all up and they fight! inefficient!

# reader/writer-priority

policy question: writers first or readers first?

writers-first: no readers go when writer waiting

readers-first: no writers go when reader waiting

previous implementation: whatever randomly happens

writers signalled first, maybe gets lock first?

...but non-deterministic in pthreads

can make **explicit decision**

# writer-priority (1)

```
mutex_t lock; cond_t ok_to_read_cv; cond_t ok_to_write_cv;
```

```
int readers = 0, writers = 0;
```

```
int waiting_writers = 0;
```

```
ReadLock() {  
    mutex_lock(&lock);  
    while (writers != 0  
           && waiting_writers != 0) {  
        cond_wait(&ok_to_read_cv, &lock);  
    }  
    ++readers;  
    mutex_unlock(&lock);  
}
```

```
ReadUnlock() {  
    mutex_lock(&lock);  
    --readers;  
    if (readers == 0) {  
        cond_signal(&ok_to_write_cv);  
    }  
    mutex_unlock(&lock);  
}
```

```
WriteLock() {  
    mutex_lock(&lock);  
    ++waiting_writers;  
    while (readers + writers != 0) {  
        cond_wait(&ok_to_write_cv, &lock);  
    }  
    --waiting_writers;  
    ++writers;  
    mutex_unlock(&lock);  
}
```

```
WriteUnlock() {  
    mutex_lock(&lock);  
    --writers;  
    if (waiting_writers != 0) {  
        cond_signal(&ok_to_write_cv);  
    } else {  
        cond_broadcast(&ok_to_read_cv);  
    }  
    mutex_unlock(&lock);  
}
```

# writer-priority (1)

```
mutex_t lock; cond_t ok_to_read_cv; cond_t ok_to_write_cv;
```

```
int readers = 0, writers = 0;
```

```
int waiting_writers = 0;
```

```
ReadLock() {  
    mutex_lock(&lock);  
    while (writers != 0  
           && waiting_writers != 0) {  
        cond_wait(&ok_to_read_cv, &lock);  
    }  
    ++readers;  
    mutex_unlock(&lock);  
}
```

```
ReadUnlock() {  
    mutex_lock(&lock);  
    --readers;  
    if (readers == 0) {  
        cond_signal(&ok_to_write_cv);  
    }  
    mutex_unlock(&lock);  
}
```

```
WriteLock() {  
    mutex_lock(&lock);  
    ++waiting_writers;  
    while (readers + writers != 0) {  
        cond_wait(&ok_to_write_cv, &lock);  
    }  
    --waiting_writers;  
    ++writers;  
    mutex_unlock(&lock);  
}
```

```
WriteUnlock() {  
    mutex_lock(&lock);  
    --writers;  
    if (waiting_writers != 0) {  
        cond_signal(&ok_to_write_cv);  
    } else {  
        cond_broadcast(&ok_to_read_cv);  
    }  
    mutex_unlock(&lock);  
}
```

# writer-priority (1)

```
mutex_t lock; cond_t ok_to_read_cv; cond_t ok_to_write_cv;
```

```
int readers = 0, writers = 0;
```

```
int waiting_writers = 0;
```

```
ReadLock() {  
    mutex_lock(&lock);  
    while (writers != 0  
           && waiting_writers != 0) {  
        cond_wait(&ok_to_read_cv, &lock);  
    }  
    ++readers;  
    mutex_unlock(&lock);  
}
```

```
ReadUnlock() {  
    mutex_lock(&lock);  
    --readers;  
    if (readers == 0) {  
        cond_signal(&ok_to_write_cv);  
    }  
    mutex_unlock(&lock);  
}
```

```
WriteLock() {  
    mutex_lock(&lock);  
    ++waiting_writers;  
    while (readers + writers != 0) {  
        cond_wait(&ok_to_write_cv, &lock);  
    }  
    --waiting_writers;  
    ++writers;  
    mutex_unlock(&lock);  
}
```

```
WriteUnlock() {  
    mutex_lock(&lock);  
    --writers;  
    if (waiting_writers != 0) {  
        cond_signal(&ok_to_write_cv);  
    } else {  
        cond_broadcast(&ok_to_read_cv);  
    }  
    mutex_unlock(&lock);  
}
```

# reader-priority (1)

```
...
int waiting_readers = 0;
ReadLock() {
    mutex_lock(&lock);
    ++waiting_readers;
    while (writers != 0) {
        cond_wait(&ok_to_read_cv, &lock);
    }
    --waiting_readers;
    ++readers;
    mutex_unlock(&lock);
}

ReadUnlock() {
    ...
    if (waiting_readers == 0) {
        cond_signal(&ok_to_write_cv);
    }
}

WriteLock() {
    mutex_lock(&lock);
    while (waiting_readers +
           readers + writers != 0) {
        cond_wait(&ok_to_write_cv);
    }
    ++writers;
    mutex_unlock(&lock);
}

WriteUnlock() {
    mutex_lock(&lock);
    --writers;
    if (waiting_readers == 0) {
        cond_signal(&ok_to_write_cv);
    } else {
        cond_broadcast(&ok_to_read_cv);
    }
    mutex_unlock(&lock);
}
```



# reader-priority (1)

```
...
int waiting_readers = 0;
ReadLock() {
    mutex_lock(&lock);
    ++waiting_readers;
    while (writers != 0) {
        cond_wait(&ok_to_read_cv, &lock);
    }
    --waiting_readers;
    ++readers;
    mutex_unlock(&lock);
}

ReadUnlock() {
    ...
    if (waiting_readers == 0) {
        cond_signal(&ok_to_write_cv);
    }
}

WriteLock() {
    mutex_lock(&lock);
    while (waiting_readers +
           readers + writers != 0) {
        cond_wait(&ok_to_write_cv);
    }
    ++writers;
    mutex_unlock(&lock);
}

WriteUnlock() {
    mutex_lock(&lock);
    --writers;
    if (waiting_readers == 0) {
        cond_signal(&ok_to_write_cv);
    } else {
        cond_broadcast(&ok_to_read_cv);
    }
    mutex_unlock(&lock);
}
```

# choosing orderings?

can use monitors to implement lots of lock policies

want  $X$  to go first/last — add extra variables  
(number of waiters, even lists of items, etc.)

need way to write condition “you can go now”

e.g. writer-priority: readers can go if no writer waiting

# simulation of reader/write lock

writer-priority version

W = writers, R = readers, WW = waiting\_writers

reader 1	reader 2	writer 1	reader 3	W	R	WW
				0	0	0

# simulation of reader/write lock

writer-priority version

W = writers, R = readers, WW = waiting\_writers

reader 1	reader 2	writer 1	reader 3	W	R	WW
				0	0	0

# simulation of reader/write lock

writer-priority version

W = writers, R = readers, WW = waiting\_writers

reader 1	reader 2	writer 1	reader 3	W	R	WW
				0	0	0
				0	1	0

ReadLock

```
mutex_lock(&lock);  
while (writers != 0 && waiting_writers != 0) {  
    cond_wait(&ok_to_read_cv, &lock);  
}  
++readers;  
mutex_unlock(&lock);
```

# simulation of reader/write lock

writer-priority version

W = writers, R = readers, WW = waiting\_writers

reader 1	reader 2	writer 1	reader 3	W	R	WW
				0	0	0
ReadLock				0	1	0
(reading)	ReadLock			0	2	0

# simulation of reader/write lock

writer-priority version

W = writers, R = readers, WW = waiting\_writers

reader 1	reader 2	writer 1	reader 3	W	R	WW
				0	0	0
ReadLock				0	1	0
(reading)	ReadLock			0	2	0
(reading)	(reading)	WriteLock wait		0	2	1

```
mutex_lock(&lock);  
++waiting_writers;  
while (readers + writers != 0) {  
    cond_wait(&ok_to_write_cv, &lock);  
}
```

# simulation of reader/write lock

writer-priority version

W = writers, R = readers, WW = waiting\_writers

reader 1	reader 2	writer 1	reader 3	W	R	WW
				0	0	0
ReadLock				0	1	0
(reading)	ReadLock			0	2	0
(reading)	(reading)	WriteLock wait		0	2	1
(reading)	(reading)	WriteLock wait	ReadLock wait	0	2	1



# simulation of reader/write lock

writer-priority version

W = writers, R = readers, WW = waiting\_writers

reader 1	reader 2	writer 1	reader 3	W	R	WW	
				0	0	0	
ReadLock				0	1	0	
(reading)	ReadLock			0	2	0	
(reading)	(reading)	WriteLock wait		0	2	1	
(reading)	(read	mutex_lock(&lock);	wait	ReadLock wait	0	2	1
ReadUnlock	←	--readers;	wait	ReadLock wait	0	1	1
		if (readers == 0)					
		...					

# simulation of reader/write lock

writer-priority version

W = writers, R = readers, WW = waiting\_writers

reader 1	reader 2	writer 1	reader 3	W	R	WW
				0	0	0
ReadLock				0	1	0
(reading)	ReadLock			0	2	0
(reading)	(reading)	WriteLock wait		0	2	1
(reading)	(reading)	WriteLock wait	ReadLock wait	0	2	1
ReadUnlock	(reading)	Write			1	1
	ReadUnlock	Write			0	1

```

mutex_lock(&lock);
--readers;
if (readers == 0)
    cond_signal(&ok_to_write_cv);
mutex_unlock(&lock);
    
```

# simulation of reader/write lock

writer-priority version

W = writers, R = readers, WW = waiting\_writers

reader 1	reader 2	writer 1	reader 3	W	R	WW
				0	0	0
ReadLock				0	1	0
(reading)	Read			0	2	0
(reading)	(rea			0	2	1
(reading)	(rea			0	2	1
ReadUnlock	(reading)	WriteLock wait	ReadLock wait	0	1	1
	ReadUnlock	WriteLock wait	ReadLock wait	0	0	1
		WriteLock	ReadLock wait	1	0	0

```

while (readers + writers != 0) {
    cond_wait(&ok_to_write_cv, &lock);
}
--waiting_writers; ++writers;
mutex_unlock(&lock);
    
```

# simulation of reader/write lock

writer-priority version

W = writers, R = readers, WW = waiting\_writers

reader 1	reader 2	writer 1	reader 3	W	R	WW
				0	0	0
ReadLock				0	1	0
(reading)	ReadLock			0	2	0
(reading)	(reading)	WriteLock wait		0	2	1
(reading)	(reading)	WriteLock wait	ReadLock wait	0	2	1
ReadUnlock	(reading)	WriteLock wait	ReadLock wait	0	1	1
	ReadUnlock	WriteLock wait	ReadLock wait	0	0	1
		WriteLock	ReadLock wait	1	0	0
		(read+writing)	ReadLock wait	1	0	0

# simulation of reader/write lock

writer-priority version

W = writers, R = readers, WW = waiting\_writers

reader 1	reader 2	writer 1	reader 3	W	R	WW
				0	0	0
ReadLock				0	1	0
(reading)	ReadLock			0	2	0
(reading)	(reading)			0	2	1
(reading)	(reading)		wait	0	2	1
ReadUnlock	(reading)		wait	0	1	1
	ReadUn		wait	0	0	1
		WriteLock	ReadLock wait	1	0	0
		(read+writing)	ReadLock wait	1	0	0
		WriteUnlock	ReadLock wait	0	0	0

```

mutex_lock(&lock);
if (waiting_writers != 0) {
    cond_signal(&ok_to_write_cv);
} else {
    cond_broadcast(&ok_to_read_cv);
}
    
```

# simulation of reader/write lock

writer-priority version

W = writers, R = readers, WW = waiting\_writers

reader 1	reader 2	writer 1	reader 3	W	R	WW	
				0	0	0	
ReadLock				0	1	0	
(reading)	ReadLock			0	2	0	
(reading)	(reading)	<pre>while (writers != 0 &amp;&amp; waiting_writers != 0) {     cond_wait(&amp;ok_to_read_cv, &amp;lock); } ++readers; mutex_unlock(&amp;lock);</pre>					
(reading)	(reading)						
ReadUnlock	(reading)						
	ReadUnlock						
		WriteLock	ReadLock	wait	1	0	0
		(read+writing)	ReadLock	wait	1	0	0
		WriteUnlock	ReadLock	wait	0	0	0
			ReadLock		0	1	0

# simulation of reader/write lock

writer-priority version

W = writers, R = readers, WW = waiting\_writers

reader 1	reader 2	writer 1	reader 3	W	R	WW
				0	0	0
ReadLock				0	1	0
(reading)	ReadLock			0	2	0
(reading)	(reading)	WriteLock wait		0	2	1
(reading)	(reading)	WriteLock wait	ReadLock wait	0	2	1
ReadUnlock	(reading)	WriteLock wait	ReadLock wait	0	1	1
	ReadUnlock	WriteLock wait	ReadLock wait	0	0	1
		WriteLock	ReadLock wait	1	0	0
		(read+writing)	ReadLock wait	1	0	0
		WriteUnlock	ReadLock wait	0	0	0
			ReadLock	0	1	0

# rwlock exercise

suppose there are multiple waiting writers

which one gets waken up first?

whichever gets signal'd or gets lock first

could instead keep in order they started waiting

exercise: what extra information should we track?

hint: we might need an array

```
mutex_t lock; cond_t ok_to_read_cv, ok_to_write_cv;  
int readers, writers, waiting_writers;
```



# rwlock exercise solution?

list of waiting writes?

```
struct WaitingWriter {
    cond_t cv;
    bool ready;
};
Queue<WaitingWriter*> waiting_writers;

WriteLock(...) {
    ...
    if (need to wait) {
        WaitingWriter self;
        self.ready = false;
        ...
        while(!self.ready) {
            pthread_cond_wait(&self.cv, &lock);
        }
    }
    ...
}
```

# rwlock exercise solution?

dedicated writing thread with queue

(DoWrite~Produce; WritingThread~Consume)

```
ThreadSafeQueue<WritingTask*> waiting_writes;
```

```
WritingThread() {
```

```
    while (true) {
```

```
        WritingTask* task = waiting_writer.Dequeue();
```

```
        WriteLock();
```

```
        DoWriteTask(task);
```

```
        task.done = true;
```

```
        cond_broadcast(&task.cv);
```

```
    }
```

```
}
```

```
DoWrite(task) {
```

```
    // instead of WriteLock(); DoWriteTask(...); WriteUnlock()
```

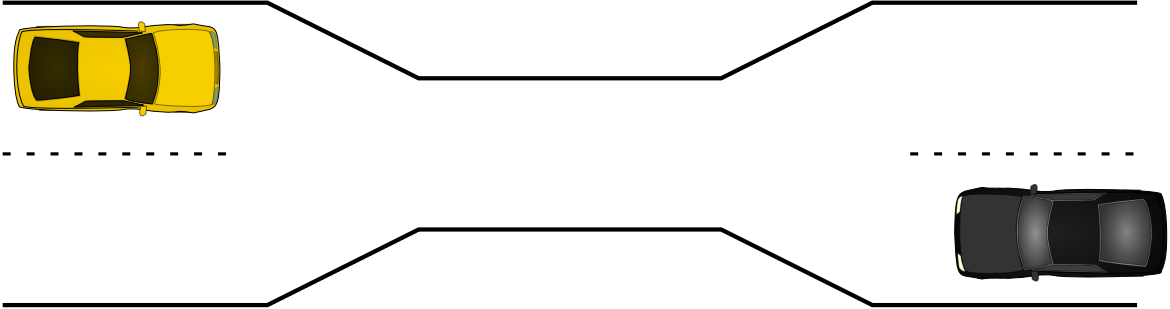
```
    WritingTask task = ...;
```

```
    waiting_writes.Enqueue(&task);
```

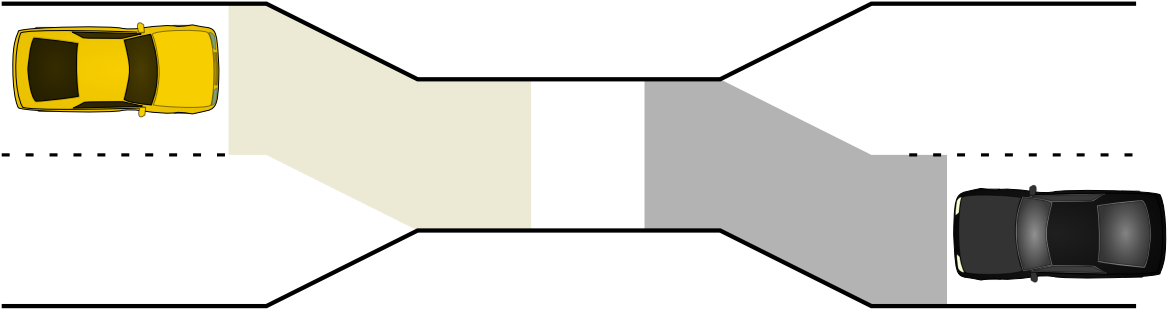
```
    while (!task.done) { cond_wait(&task.cv); }
```

```
}
```

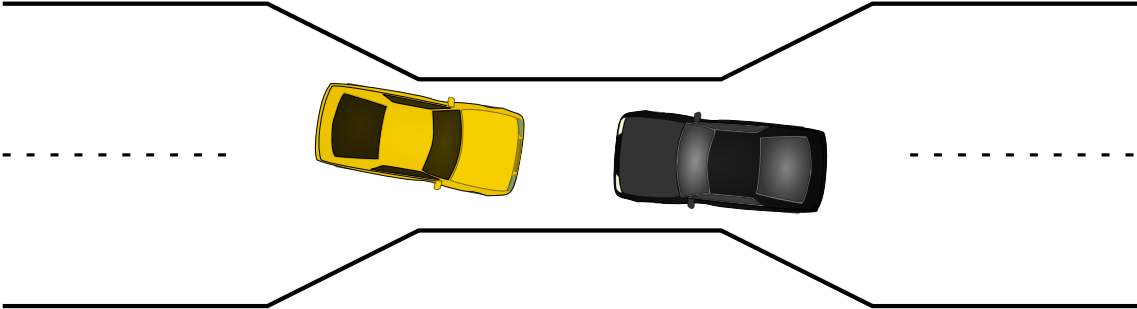
# the one-way bridge



# the one-way bridge



# the one-way bridge



# the one-way bridge

