

sync-rwlocks

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 rwlocks

```
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);
```

rwlock effects exercise

```
pthread_rwlock_t lock;
void ThreadA() {
    pthread_rwlock_rdlock(&lock);
    puts("a");
    ...
    puts("A");
    pthread_rwlock_unlock(&lock);
}
void ThreadB() {
    pthread_rwlock_rdlock(&lock);
    puts("b");
    ...
    puts("B");
    pthread_rwlock_unlock(&lock);
}
```

```
void ThreadC() {
    pthread_rwlock_wrlock(&lock);
    puts("c");
    ...
    puts("C");
    pthread_rwlock_unlock(&lock);
}
void ThreadD() {
    pthread_rwlock_wrlock(&lock);
    puts("d");
    ...
    puts("D");
    pthread_rwlock_unlock(&lock);
}
```

exercise: which of these outputs are possible?

1. aAbBcCdD
2. abABcdDC
3. cCabBAdD
4. cdCDaAbB
5. caACdDbB

Backup slides

rwlocks with monitors (attempt 1)

```
mutex_t lock;
```

rwlocks with monitors (attempt 1)

```
mutex_t lock;
```

rwlocks with monitors (attempt 1)

```
mutex_t lock;
```

```
unsigned int 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;
```

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);] {.fragment
fragment-index=5 .custom .myem-only}]{.fragment
fragment-index=6 .custom .myem-only}
    [[cond_broadcast(&ok_to_read_cv);] {.fragment
fragment-index=4 .custom .myem-only}]{.fragment
fragment-index=6 .custom .myem-only}
    mutex_unlock(&lock);
}
```

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);]]{.fragment
fragment-index=5 .custom .myem-only}}{.fragment
fragment-index=6 .custom .myem-only}
    [[cond_broadcast(&ok_to_read_cv);]]{.fragment
fragment-index=4 .custom .myem-only}}{.fragment
fragment-index=6 .custom .myem-only}
    mutex_unlock(&lock);
}
```

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);]]{.fragment
fragment-index=5 .custom .myem-only}}{.fragment
fragment-index=6 .custom .myem-only}
    [[cond_broadcast(&ok_to_read_cv);]]{.fragment
fragment-index=4 .custom .myem-only}}{.fragment
fragment-index=6 .custom .myem-only}
    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);
}
```

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 (readers == 0 && 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 (readers == 0 && 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