

CS162  
Operating Systems and  
Systems Programming  
Lecture 7

Readers-Writers  
Language Support for Synchronization

February 13, 2008  
Prof. Anthony D. Joseph  
<http://inst.eecs.berkeley.edu/~cs162>

Review: Semaphores

- Definition: a Semaphore has a non-negative integer value and supports the following two operations:
  - Only time can set integer directly is at initialization time
- P(): an atomic operation that waits for semaphore to become positive, then decrements it by 1
  - Think of this as the wait() operation
- V(): an atomic operation that increments the semaphore by 1, waking up a waiting P, if any
  - Think of this as the signal() operation

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.2

Review: Full Solution to Bounded Buffer

```
Semaphore fullBuffer = 0; // Initially, no coke
Semaphore emptyBuffers = numBuffers;
                          // Initially, num empty slots
Semaphore mutex = 1;     // No one using machine

Producer(item) {
    emptyBuffers.P();    // Wait until space
    mutex.P();          // Wait until buffer free
    Enqueue(item);
    mutex.V();
    fullBuffers.V();    // Tell consumers there is
                       // more coke
}

Consumer() {
    fullBuffers.P();    // Check if there's a coke
    mutex.P();          // Wait until machine free
    item = Dequeue();
    mutex.V();
    emptyBuffers.V();  // tell producer need more
    return item;
}
```

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.3

Review: Discussion about Solution

- Is order of P's important?
  - Yes! Can cause deadlock
- Is order of V's important?
  - No, except for scheduling efficiency
- What if we have 2 producers or 2 consumers?
  - Nothing changes!
- Semaphores are a huge step up, but:
  - They are confusing because they are dual purpose:
    - » Both mutual exclusion and scheduling constraints
    - » Example: the fact that flipping of P's in bounded buffer gives deadlock is not immediately obvious
  - Cleaner idea: Use *locks* for mutual exclusion and *condition variables* for scheduling constraints

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.4

## Review: Monitors and Condition Variables

- Definition: **Monitor**: a lock and zero or more condition variables for managing concurrent access to shared data
  - Use of Monitors is a programming paradigm
  - Some languages like Java provide monitors in the language
- The lock provides mutual exclusion to shared data:
  - Always acquire before accessing shared data structure
  - Always release after finishing with shared data
  - Lock initially free

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.5

## Goals for Today

- Continue with Synchronization Abstractions
  - Monitors and condition variables
- Readers-Writers problem and solution
- Language Support for Synchronization
- An Overview of ACID properties in a DBMS

Note: Some slides and/or pictures in the following are adapted from slides ©2005 Silberschatz, Galvin, and Gagne. Many slides generated from my lecture notes by Kubiawicz.

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.6

## Simple Monitor Example (version 1)

- Here is an (infinite) synchronized queue

```
Lock lock;
Queue queue;

AddToQueue (item) {
    lock.Acquire();           // Lock shared data
    queue.enqueue(item);     // Add item
    lock.Release();          // Release Lock
}

RemoveFromQueue () {
    lock.Acquire();           // Lock shared data
    item = queue.dequeue();  // Get next item or null
    lock.Release();          // Release Lock
    return(item);            // Might return null
}
```

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.7

## Condition Variables

- How do we change the RemoveFromQueue() routine to wait until something is on the queue?
  - Could do this by keeping a count of the number of things on the queue (with semaphores), but error prone
- **Condition Variable**: a queue of threads waiting for something *inside* a critical section
  - Key idea: allow sleeping inside critical section by atomically releasing lock at time we go to sleep
  - Contrast to semaphores: Can't wait inside critical section
- Operations:
  - **Wait (&lock)**: Atomically release lock and go to sleep. Re-acquire lock later, before returning.
  - **Signal ()**: Wake up one waiter, if any
  - **Broadcast ()**: Wake up all waiters
- Rule: Must hold lock when doing condition variable ops!
  - In Birrell paper, he says can perform signal() outside of lock - IGNORE HIM (this is only an optimization)

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.8

## Complete Monitor Example (with condition variable)

- Here is an (infinite) synchronized queue

```
Lock lock;
Condition dataready;
Queue queue;

AddToQueue(item) {
    lock.Acquire();           // Get Lock
    queue.enqueue(item);     // Add item
    dataready.signal();      // Signal any waiters
    lock.Release();          // Release Lock
}

RemoveFromQueue() {
    lock.Acquire();           // Get Lock
    while (queue.isEmpty()) {
        dataready.wait(&lock); // If nothing, sleep
    }
    item = queue.dequeue();  // Get next item
    lock.Release();          // Release Lock
    return(item);
}
```

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.9

## Mesa vs. Hoare monitors

- Need to be careful about precise definition of signal and wait. Consider a piece of our dequeue code:

```
while (queue.isEmpty()) {
    dataready.wait(&lock); // If nothing, sleep
}
item = queue.dequeue(); // Get next item
```

- Why didn't we do this?

```
if (queue.isEmpty()) {
    dataready.wait(&lock); // If nothing, sleep
}
item = queue.dequeue(); // Get next item
```

- Answer: depends on the type of scheduling

- Hoare-style (most textbooks):

- » Signaler gives lock, CPU to waiter; waiter runs immediately
- » Waiter gives up lock, processor back to signaler when it exits critical section or if it waits again

- Mesa-style (Nachos, most real operating systems):

- » Signaler keeps lock and processor
- » Waiter placed on ready queue with no special priority
- » Practically, need to check condition again after wait

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.10

## Administrivia


- First design document due *tomorrow* by 11:59pm
  - Good luck!
  - Use the newsgroup for questions (search Google groups)
- Design reviews:
  - Everyone must attend! (no exceptions)
  - 2 points off for one missing person
  - 1 additional point off for each additional missing person
  - Penalty for arriving late (plan on arriving 5–10 mins early)
  - Sign up link will be posted on announcements and projects pages
- What we expect in document/review:
  - Architecture, **correctness constraints**, algorithms, pseudocode, **NO CODE!**
  - **Important: testing strategy, and test case types**

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.11

## Readers/Writers Problem



- Motivation: Consider a shared database

- Two classes of users:

- » Readers - never modify database
- » Writers - read and modify database

- Is using a single lock on the whole database sufficient?

- » Like to have many readers at the same time
- » Only one writer at a time

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.12

## Basic Readers/Writers Solution

- **Correctness Constraints:**
  - Readers can access database when no writers
  - Writers can access database when no readers or writers
  - Only one thread manipulates state variables at a time
- **Basic structure of a solution:**
  - Reader()
    - Wait until no writers
    - Access data base
    - Check out - wake up a waiting writer
  - Writer()
    - Wait until no active readers or writers
    - Access database
    - Check out - wake up waiting readers or writer
  - **State variables (Protected by a lock called "lock"):**
    - » int AR: Number of active readers; initially = 0
    - » int WR: Number of waiting readers; initially = 0
    - » int AW: Number of active writers; initially = 0
    - » int WW: Number of waiting writers; initially = 0
    - » Condition okToRead = NIL
    - » Condition okToWrite = NIL

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.13

## Code for a Reader

```
Reader() {
    // First check self into system
    lock.Acquire();
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        okToRead.wait(&lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    lock.release();
    // Perform actual read-only access
    AccessDatabase(ReadOnly);
    // Now, check out of system
    lock.Acquire();
    AR--; // No longer active
    if (AR == 0 && WW > 0) // No other active readers
        okToWrite.signal(); // Wake up one writer
    lock.Release();
}
```

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.14

## Code for a Writer

```
Writer() {
    // First check self into system
    lock.Acquire();
    while ((AW + AR) > 0) { // Is it safe to write?
        WW++; // No. Active users exist
        okToWrite.wait(&lock); // Sleep on cond var
        WW--; // No longer waiting
    }
    AW++; // Now we are active!
    lock.release();
    // Perform actual read/write access
    AccessDatabase(ReadWrite);
    // Now, check out of system
    lock.Acquire();
    AW--; // No longer active
    if (WW > 0) { // Give priority to writers
        okToWrite.signal(); // Wake up one writer
    } else if (WR > 0) { // Otherwise, wake reader
        okToRead.broadcast(); // Wake all readers
    }
    lock.Release();
}
```

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.15

## Simulation of Readers/Writers solution

- Consider the following sequence of operators:
  - R1, R2, W1, R3
- On entry, each reader checks the following:

```
while ((AW + WW) > 0) { // Is it safe to read?
    WR++; // No. Writers exist
    okToRead.wait(&lock); // Sleep on cond var
    WR--; // No longer waiting
}
AR++; // Now we are active!
```
- First, R1 comes along:  
AR = 1, WR = 0, AW = 0, WW = 0
- Next, R2 comes along:  
AR = 2, WR = 0, AW = 0, WW = 0
- Now, readers make take a while to access database
  - Situation: Locks released
  - Only AR is non-zero

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.16

### Simulation(2)

- Next, W1 comes along:

```
while ((AW + AR) > 0) { // Is it safe to write?
    WW++; // No. Active users exist
    okToWrite.wait(&lock); // Sleep on cond var
    WW--; // No longer waiting
}
AW++;
```
- Can't start because of readers, so go to sleep:  
AR = 2, WR = 0, AW = 0, WW = 1
- Finally, R3 comes along:  
AR = 2, WR = 1, AW = 0, WW = 1
- Now, say that R2 finishes before R1:  
AR = 1, WR = 1, AW = 0, WW = 1
- Finally, last of first two readers (R1) finishes and wakes up writer:

```
if (AR == 0 && WW > 0) // No other active readers
    okToWrite.signal(); // Wake up one writer
```

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.17

### Simulation(3)

- When writer wakes up, get:  
AR = 0, WR = 1, AW = 1, WW = 0
- Then, when writer finishes:

```
if (WW > 0){ // Give priority to writers
    okToWrite.signal(); // Wake up one writer
} else if (WR > 0) { // Otherwise, wake reader
    okToRead.broadcast(); // Wake all readers
}
```
- Writer wakes up reader, so get:  
AR = 1, WR = 0, AW = 0, WW = 0
- When reader completes, we are finished

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.18

### Questions

- Can readers starve? Consider Reader() entry code:

```
while ((AW + WW) > 0) { // Is it safe to read?
    WR++; // No. Writers exist
    okToRead.wait(&lock); // Sleep on cond var
    WR--; // No longer waiting
}
AR++; // Now we are active!
```
- What if we erase the condition check in Reader exit?

```
AR--; // No longer active
if (AR == 0 && WW > 0) // No other active readers
    okToWrite.signal(); // Wake up one writer
```
- Further, what if we turn the signal() into broadcast()

```
AR--; // No longer active
okToWrite.broadcast(); // Wake up one writer
```
- Finally, what if we use only one condition variable (call it "okToContinue") instead of two separate ones?
  - Both readers and writers sleep on this variable
  - Must use broadcast() instead of signal()

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.19

### Can we construct Monitors from Semaphores?

- Locking aspect is easy: Just use a mutex
- Can we implement condition variables this way?

```
Wait() { semaphore.P(); }
Signal() { semaphore.V(); }
```
- Does this work better?

```
Wait(Lock lock) {
    lock.Release();
    semaphore.P();
    lock.Acquire();
}
Signal() { semaphore.V(); }
```

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.20

### Construction of Monitors from Semaphores (con't)

- Problem with previous try:
  - P and V are commutative - result is the same no matter what order they occur
  - Condition variables are NOT commutative
- Does this fix the problem?

```
Wait(Lock lock) {
    lock.Release();
    semaphore.P();
    lock.Acquire();
}
Signal() {
    if semaphore queue is not empty
        semaphore.V();
}
```

  - Not legal to look at contents of semaphore queue
  - There is a race condition - signaler can slip in after lock release and before waiter executes semaphore.P()
- It is actually possible to do this correctly
  - Complex solution for Hoare scheduling in book
  - Can you come up with simpler Mesa-scheduled solution?

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.21

### Monitor Conclusion

- Monitors represent the logic of the program
  - Wait if necessary
  - Signal when change something so any waiting threads can proceed
- Basic structure of monitor-based program:

```
lock
while (need to wait) { } Check and/or update
                        state variables
                        Wait if necessary
unlock

do something so no need to wait

lock

condvar.signal(); } Check and/or update
                    state variables

unlock
```

2/13/08

Joseph CS162 ©UCB Spring 2008


Lec 7.22

BREAK

### C-Language Support for Synchronization

- C language: Pretty straightforward synchronization
  - Just make sure you know *all* the code paths out of a critical section

```
int Rtn() {
    lock.acquire();
    ...
    if (exception) {
        lock.release();
        return errReturnCode;
    }
    ...
    lock.release();
    return OK;
}
```



- Watch out for setjmp/longjmp!
  - » Can cause a non-local jump out of procedure
  - » In example, procedure E calls longjmp, popping stack back to procedure B
  - » If Procedure C had lock.acquire, problem!

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.24

## C++ Language Support for Synchronization

- Languages with exceptions like C++
  - Languages that support exceptions are problematic (easy to make a non-local exit without releasing lock)
  - Consider:

```
void Rtn() {
    lock.acquire();
    ...
    DoFoo();
    ...
    lock.release();
}
void DoFoo() {
    ...
    if (exception) throw errException;
    ...
}
```
  - Notice that an exception in DoFoo() will exit without releasing the lock

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.25

## C++ Language Support for Synchronization (con't)

- Must catch all exceptions in critical sections
  - Catch exceptions, release lock, and re-throw exception:

```
void Rtn() {
    lock.acquire();
    try {
        ...
        DoFoo();
        ...
    } catch (...) { // catch exception
        lock.release(); // release lock
        throw; // re-throw the exception
    }
    lock.release();
}
void DoFoo() {
    ...
    if (exception) throw errException;
    ...
}
```
  - Even Better: `auto_ptr<T>` facility. See C++ Spec.
    - » Can deallocate/free lock regardless of exit method

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.26

## Java Language Support for Synchronization

- Java has explicit support for threads and thread synchronization
- Bank Account example:

```
class Account {
    private int balance;
    // object constructor
    public Account (int initialBalance) {
        balance = initialBalance;
    }
    public synchronized int getBalance() {
        return balance;
    }
    public synchronized void deposit(int amount) {
        balance += amount;
    }
}
```
- Every object has an associated lock which gets automatically acquired and released on entry and exit from a *synchronized* method.

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.27

## Java Language Support for Synchronization (con't)

- Java also has *synchronized* statements:

```
synchronized (object) {
    ...
}
```
- Since every Java object has an associated lock, this type of statement acquires and releases the object's lock on entry and exit of the body
- Works properly even with exceptions:

```
synchronized (object) {
    ...
    DoFoo();
    ...
}
void DoFoo() {
    throw errException;
}
```

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.28

## Java Language Support for Synchronization (con't 2)

- In addition to a lock, every object has a **single** condition variable associated with it
  - How to wait inside a synchronization method or block:
    - » `void wait(long timeout); // Wait for timeout`
    - » `void wait(long timeout, int nanoseconds); //variant`
    - » `void wait();`
  - How to signal in a synchronized method or block:
    - » `void notify(); // wakes up oldest waiter`
    - » `void notifyAll(); // like broadcast, wakes everyone`
  - Condition variables can wait for a bounded length of time. This is useful for handling exception cases:

```
t1 = time.now();
while (!ATMRequest()) {
    wait (CHECKPERIOD);
    t2 = time.now();
    if (t2 - t1 > LONG_TIME) checkMachine();
}
```
  - Not all Java VMs equivalent!
    - » Different scheduling policies, not necessarily preemptive!

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.29

## ACID

- How does a database handle concurrency?
  - It provides A.C.I.D. properties - Atomicity, Consistency, Isolation, Durability
- Key concept: Transaction
  - An **atomic sequence** of database actions (reads/writes)
    - » Actions all happen or none at all
  - Takes DB from one **consistent state** to another



2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.30

## DBMS Consistency Example



- Here, **consistency** is based on our knowledge of banking "semantics"
- In general, up to writer of transaction to ensure transaction preserves consistency
- DBMS provides (limited) automatic enforcement, via **integrity constraints**
  - e.g., balances must be  $\geq 0$

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.31

## Challenge: Concurrent transactions

- **Goal:** execute xacts  $\{T_1, T_2, \dots, T_n\}$ , and ensure a consistent outcome
  - Isolate xact's intermediate state from other xacts
- **One option:** "serial" schedule (one after another)
- **Better:** allow interleaving of xact actions, as long as outcome is equivalent to some serial schedule
- Two possible enforcement methods
  - Optimistic: permit arbitrary interleaving, then check equivalence to serial schedule
  - Pessimistic: xacts set **locks** on data objects, such that illegal interleaving is impossible
    - » More on locking in another lecture...

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.32



## Ensuring Transaction Properties

- DBMS ensures:
    - **Atomicity** even if xact aborted (due to deadlock, system crash, ...)
    - **Durability** (persistence) of **committed** xacts, even if system crashes
  - **Idea:** Keep a **log** of all actions carried out by the DBMS:
    - Record all DB modifications in log, **before** they are executed
    - To abort a xact, undo logged actions in reverse order
    - If system crashes, must:
      - 1) **undo** partially executed xacts (ensures atomicity)
      - 2) **redo** committed xacts (ensures durability)
- *Much trickier than it sounds!*

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.33

## ACID Summary

- **Atomicity:** guarantee that either all of the tasks of a transaction are performed, or none of them are
- **Consistency:** database is in a legal state when the transaction begins and when it ends - a transaction cannot break the rules, or **integrity constraints**
- **Isolation:** operations inside a transaction appear isolated from all other operations - no operation outside transaction can see data in an intermediate state
- **Durability:** guarantee that once the user has been notified of success, the transaction will persist (survive system failure)

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.34

## Summary

- **Monitors:** A lock plus one or more condition variables
  - Always acquire lock before accessing shared data
  - Use condition variables to wait inside critical section
    - » Three Operations: **Wait()**, **Signal()**, and **Broadcast()**
- **Readers/Writers**
  - Readers can access database when no writers
  - Writers can access database when no readers
  - Only one thread manipulates state variables at a time
- **Language support for synchronization:**
  - Java provides **synchronized** keyword and one condition-variable per object (with **wait()** and **notify()**)

2/13/08

Joseph CS162 ©UCB Spring 2008

Lec 7.35