

CS162 Operating Systems and Systems Programming Lecture 8

Locks, Semaphores, Monitors, and Quick Intro to Scheduling

September 23rd, 2015
Prof. John Kubiawicz
<http://cs162.eecs.Berkeley.edu>

Review: Synchronization problem with Threads

- One thread per transaction, each running:

```
Deposit(acctId, amount) {  
    acct = GetAccount(actId); /* May use disk I/O */  
    acct->balance += amount;  
    StoreAccount(acct);      /* Involves disk I/O */  
}
```

- Unfortunately, shared state can get corrupted:

```
Thread 1                               Thread 2  
load r1, acct->balance                 load r1, acct->balance  
                                       add r1, amount2  
                                       store r1, acct->balance  
  
add r1, amount1  
store r1, acct->balance
```

- **Atomic Operation:** an operation that always runs to completion or not at all
 - It is *indivisible*: it cannot be stopped in the middle and state cannot be modified by someone else in the middle

9/23/15

Kubiawicz CS162 ©UCB Fall 2015

Lec 8.2

Review: Too Much Milk Solution #3

- Here is a possible two-note solution:

```
Thread A                               Thread B  
leave note A;                          leave note B;  
while (note B) {\\X                    if (noNote A) {\\Y  
    do nothing;                          if (noMilk) {  
    }                                       buy milk;  
    }                                       }  
if (noMilk) {                               }  
    buy milk;                               }  
}                                       remove note B;  
remove note A;
```

- Does this work? Yes. Both can guarantee that:
 - It is safe to buy, or
 - Other will buy, ok to quit
- At X:
 - if no note B, safe for A to buy,
 - otherwise wait to find out what will happen
- At Y:
 - if no note A, safe for B to buy
 - Otherwise, A is either buying or waiting for B to quit

9/23/15

Kubiawicz CS162 ©UCB Fall 2015

Lec 8.3

Review: Too Much Milk: Solution #4

- Suppose we have some sort of implementation of a lock (more in a moment).
 - **Acquire(&mylock)** - wait until lock is free, then grab
 - **Release(&mylock)** - Unlock, waking up anyone waiting
 - These must be atomic operations - if two threads are waiting for the lock and both see it's free, only one succeeds to grab the lock
- Then, our milk problem is easy:

```
Acquire(&milklock);  
if (nomilk)  
    buy milk;  
Release(&milklock);
```
- Once again, section of code between Acquire() and Release() called a **"Critical Section"**
- Of course, you can make this even simpler: suppose you are out of ice cream instead of milk
 - Skip the test since you always need more ice cream.

9/23/15

Kubiawicz CS162 ©UCB Fall 2015

Lec 8.4

Goals for Today

- Explore several implementations of locks
- Continue with Synchronization Abstractions
 - Semaphores, Monitors, and Condition variables
- Very Quick Introduction to scheduling

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.

9/23/15

Kubiawicz CS162 ©UCB Fall 2015

Lec 8.5

How to implement Locks?

- **Lock**: prevents someone from doing something
 - Lock before entering critical section and before accessing shared data
 - Unlock when leaving, after accessing shared data
 - Wait if locked
 - » Important idea: all synchronization involves waiting
 - » Should *sleep* if waiting for a long time
- Atomic Load/Store: get solution like Milk #3
 - Looked at this last lecture
 - Pretty complex and error prone
- Hardware Lock instruction
 - Is this a good idea?
 - What about putting a task to sleep?
 - » How do you handle the interface between the hardware and scheduler?
 - Complexity?
 - » Done in the Intel 432
 - » Each feature makes hardware more complex and slow



9/23/15

Kubiawicz CS162 ©UCB Fall 2015

Lec 8.6

Naïve use of Interrupt Enable/Disable

- How can we build multi-instruction atomic operations?
 - Recall: dispatcher gets control in two ways.
 - » Internal: Thread does something to relinquish the CPU
 - » External: Interrupts cause dispatcher to take CPU
 - On a uniprocessor, can avoid context-switching by:
 - » Avoiding internal events (although virtual memory tricky)
 - » Preventing external events by disabling interrupts
- Consequently, naïve Implementation of locks:

```
LockAcquire { disable Ints; }
LockRelease { enable Ints; }
```
- Problems with this approach:
 - **Can't let user do this!** Consider following:

```
LockAcquire();
While(TRUE) {;}
```
 - Real-Time system—no guarantees on timing!
 - » Critical Sections might be arbitrarily long
 - What happens with I/O or other important events?
 - » "Reactor about to meltdown. Help?"



9/23/15

Kubiawicz CS162 ©UCB Fall 2015

Lec 8.7

Better Implementation of Locks by Disabling Interrupts

- Key idea: maintain a lock variable and impose mutual exclusion only during operations on that variable

```
int value = FREE;
```



```
Acquire() {
    disable interrupts;
    if (value == BUSY) {
        put thread on wait queue;
        Go to sleep();
        // Enable interrupts?
    } else {
        value = BUSY;
    }
    enable interrupts;
}

Release() {
    disable interrupts;
    if (anyone on wait queue) {
        take thread off wait queue
        Place on ready queue;
    } else {
        value = FREE;
    }
    enable interrupts;
}
```

9/23/15

Kubiawicz CS162 ©UCB Fall 2015

Lec 8.8

New Lock Implementation: Discussion

- Why do we need to disable interrupts at all?
 - Avoid interruption between checking and setting lock value
 - Otherwise two threads could think that they both have lock

```
Acquire() {
    disable interrupts;
    if (value == BUSY) {
        put thread on wait queue;
        Go to sleep();
        // Enable interrupts?
    } else {
        value = BUSY;
    }
    enable interrupts;
}
```

} Critical Section

- Note: unlike previous solution, the critical section (inside Acquire()) is very short
 - User of lock can take as long as they like in their own critical section: doesn't impact global machine behavior
 - Critical interrupts taken in time!

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.9

Interrupt re-enable in going to sleep

- What about re-enabling ints when going to sleep?

```
Acquire() {
    disable interrupts;
    if (value == BUSY) {
        put thread on wait queue;
        Go to sleep();
    } else {
        value = BUSY;
    }
    enable interrupts;
}
```

Enable Position
Enable Position
Enable Position

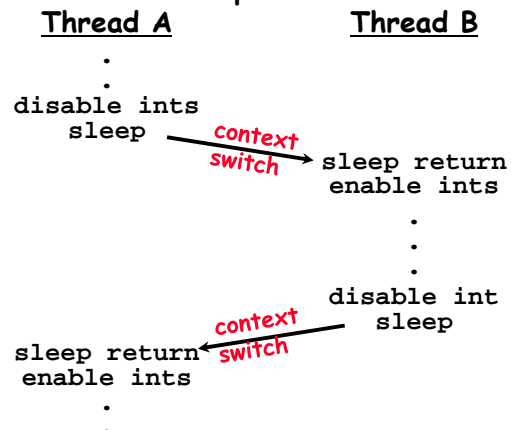
9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.10

How to Re-enable After Sleep()?

- In scheduler, since interrupts are disabled when you call sleep:
 - Responsibility of the next thread to re-enable ints
 - When the sleeping thread wakes up, returns to acquire and re-enables interrupts



9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.11

Administrivia

- First Checkpoint due this Friday 11:59pm PST
 - Yes this is graded!
 - Assume design document is *high level!*
 - » You should think of this as a document for a manager (your TA)
- Design review
 - High-level discussion of your approach
 - » What will you modify?
 - » What algorithm will you use?
 - » How will things be linked together, etc.
 - » Do not need final design (complete with all semicolons!)
 - You *will* be asked about testing
 - » Understand testing framework
 - » Are there things you are doing that are not tested by the tests we give you?
- Do your own work!
 - Please do not try to find solutions from previous terms
 - We will be look out for this...
- Basic semaphores work in PintOS!
 - However, you will need to implement priority scheduling behavior both in semaphore and ready queue

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.12

Atomic Read-Modify-Write instructions

- Problems with previous solution:
 - Can't give lock implementation to users
 - Doesn't work well on multiprocessor
 - » Disabling interrupts on all processors requires messages and would be very time consuming
- Alternative: atomic instruction sequences
 - These instructions read a value from memory and write a new value atomically
 - Hardware is responsible for implementing this correctly
 - » on both uniprocessors (not too hard)
 - » and multiprocessors (requires help from cache coherence protocol)
 - Unlike disabling interrupts, can be used on both uniprocessors and multiprocessors

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.13

Examples of Read-Modify-Write

- `test&set (&address) { /* most architectures */`

```
    result = M[address];
    M[address] = 1;
    return result;
}
```
- `swap (&address, register) { /* x86 */`

```
    temp = M[address];
    M[address] = register;
    register = temp;
}
```
- `compare&swap (&address, reg1, reg2) { /* 68000 */`

```
    if (reg1 == M[address]) {
        M[address] = reg2;
        return success;
    } else {
        return failure;
    }
}
```
- `load-linked&store conditional(&address) {`

```
    /* R4000, alpha */
    loop:
        ll r1, M[address];
        movi r2, 1; /* Can do arbitrary comp */
        sc r2, M[address];
        beqz r2, loop;
    }
```

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.14

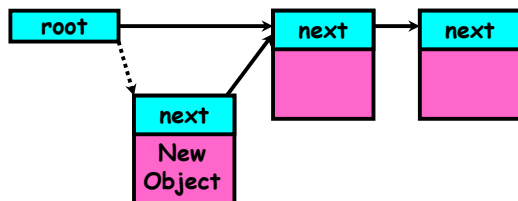
Using of Compare&Swap for queues

- `compare&swap (&address, reg1, reg2) { /* 68000 */`

```
    if (reg1 == M[address]) {
        M[address] = reg2;
        return success;
    } else {
        return failure;
    }
}
```

Here is an atomic add to linked-list function:

```
addToQueue(&object) {
    do {
        // repeat until no conflict
        ld r1, M[root] // Get ptr to current head
        st r1, M[object] // Save link in new object
    } until (compare&swap(&root, r1, object));
}
```



9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.15

Implementing Locks with test&set

- Another flawed, but simple solution:

```
int value = 0; // Free
Acquire() {
    while (test&set(value)); // while busy
}
Release() {
    value = 0;
}
```

- Simple explanation:

- If lock is free, test&set reads 0 and sets value=1, so lock is now busy. It returns 0 so while exits.
- If lock is busy, test&set reads 1 and sets value=1 (no change). It returns 1, so while loop continues
- When we set value = 0, someone else can get lock

- **Busy-Waiting:** thread consumes cycles while waiting

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.16

Problem: Busy-Waiting for Lock

- Positives for this solution
 - Machine can receive interrupts
 - User code can use this lock
 - Works on a multiprocessor
- Negatives
 - This is very inefficient because the busy-waiting thread will consume cycles waiting
 - Waiting thread may take cycles away from thread holding lock (no one wins!)
 - **Priority Inversion**: If busy-waiting thread has higher priority than thread holding lock \Rightarrow no progress!
- Priority Inversion problem with original Martian rover
- For semaphores and monitors, waiting thread may wait for an arbitrary length of time!
 - Thus even if busy-waiting was OK for locks, definitely not ok for other primitives
 - Homework/exam solutions should not have busy-waiting!



9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.17

Multiprocessor Spin Locks: test&test&set

- A better solution for multiprocessors:

```
int mylock = 0; // Free
Acquire() {
    do {
        while(mylock); // Wait until might be free
    } while(test&set(&mylock)); // exit if get lock
}

Release() {
    mylock = 0;
}
```
- Simple explanation:
 - Wait until lock might be free (only reading - stays in cache)
 - Then, try to grab lock with test&set
 - Repeat if fail to actually get lock
- Issues with this solution:
 - **Busy-Waiting**: thread still consumes cycles while waiting
 - » However, it does not impact other processors!

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.18

Better Locks using test&set

- Can we build test&set locks without busy-waiting?
 - Can't entirely, but can minimize!
 - Idea: only busy-wait to atomically check lock value

```
int guard = 0;
int value = FREE;
```



```
Acquire() {
    // Short busy-wait time
    while (test&set(guard));
    if (value == BUSY) {
        put thread on wait queue;
        go to sleep() & guard = 0;
    } else {
        value = BUSY;
        guard = 0;
    }
}

Release() {
    // Short busy-wait time
    while (test&set(guard));
    if anyone on wait queue {
        take thread off wait queue;
        Place on ready queue;
    } else {
        value = FREE;
    }
    guard = 0;
}
```

- Note: sleep has to be sure to reset the guard variable
 - Why can't we do it just before or just after the sleep?

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.19

Higher-level Primitives than Locks

- Goal of last couple of lectures:
 - What is the right abstraction for synchronizing threads that share memory?
 - Want as high a level primitive as possible
- Good primitives and practices important!
 - Since execution is not entirely sequential, really hard to find bugs, since they happen rarely
 - UNIX is pretty stable now, but up until about mid-80s (10 years after started), systems running UNIX would crash every week or so - concurrency bugs
- Synchronization is a way of coordinating multiple concurrent activities that are using shared state
 - This lecture and the next presents a couple of ways of structuring the sharing

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.20

Semaphores



- Semaphores are a kind of generalized lock
 - First defined by Dijkstra in late 60s
 - Main synchronization primitive used in original UNIX
- Definition: a Semaphore has a non-negative integer value and supports the following two operations:
 - **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
 - Note that **P()** stands for "proberen" (to test) and **V()** stands for "verhogen" (to increment) in Dutch

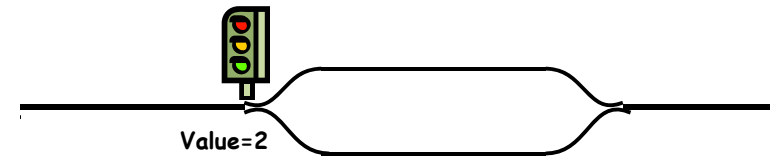
9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.21

Semaphores Like Integers Except

- Semaphores are like integers, except
 - No negative values
 - Only operations allowed are P and V - can't read or write value, except to set it initially
 - Operations must be atomic
 - » Two P's together can't decrement value below zero
 - » Similarly, thread going to sleep in P won't miss wakeup from V - even if they both happen at same time
- Semaphore from railway analogy
 - Here is a semaphore initialized to 2 for resource control:



9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.22

Two Uses of Semaphores

- Mutual Exclusion (initial value = 1)
 - Also called "Binary Semaphore".
 - Can be used for mutual exclusion:

```
semaphore.P();
// Critical section goes here
semaphore.V();
```
- Scheduling Constraints (initial value = 0)
 - Locks are fine for mutual exclusion, but what if you want a thread to wait for something?
 - Example: suppose you had to implement ThreadJoin which must wait for thread to terminate:

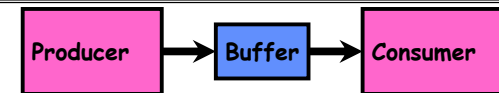
```
Initial value of semaphore = 0
ThreadJoin {
    semaphore.P();
}
ThreadFinish {
    semaphore.V();
}
```

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.23

Producer-consumer with a bounded buffer



- Problem Definition
 - Producer puts things into a shared buffer
 - Consumer takes them out
 - Need synchronization to coordinate producer/consumer
- Don't want producer and consumer to have to work in lockstep, so put a fixed-size buffer between them
 - Need to synchronize access to this buffer
 - Producer needs to wait if buffer is full
 - Consumer needs to wait if buffer is empty
- Example 1: GCC compiler
 - cpp | cc1 | cc2 | as | ld
- Example 2: Coke machine
 - Producer can put limited number of cokes in machine
 - Consumer can't take cokes out if machine is empty



9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.24

Correctness constraints for solution

- **Correctness Constraints:**
 - Consumer must wait for producer to fill buffers, if none full (scheduling constraint)
 - Producer must wait for consumer to empty buffers, if all full (scheduling constraint)
 - Only one thread can manipulate buffer queue at a time (mutual exclusion)
- **Remember why we need mutual exclusion**
 - Because computers are stupid
 - Imagine if in real life: the delivery person is filling the machine and somebody comes up and tries to stick their money into the machine
- **General rule of thumb:**
Use a separate semaphore for each constraint
 - Semaphore fullBuffers; // consumer's constraint
 - Semaphore emptyBuffers; // producer's constraint
 - Semaphore mutex; // mutual exclusion

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.25

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;
}
```

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.26

Discussion about Solution

- **Why asymmetry?**
 - Producer does: emptyBuffer.P(), fullBuffer.V()
 - Consumer does: fullBuffer.P(), emptyBuffer.V()
- **Is order of P's important?**
- **Is order of V's important?**
- **What if we have 2 producers or 2 consumers?**
 - Do we need to change anything?

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.27

Motivation for Monitors and Condition Variables

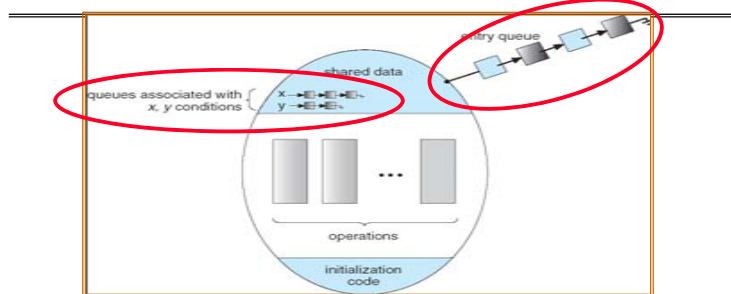
- **Semaphores are a huge step up; just think of trying to do the bounded buffer with only loads and stores**
 - Problem is that semaphores are dual purpose:
 - » They are used for both mutex and scheduling constraints
 - » Example: the fact that flipping of P's in bounded buffer gives deadlock is not immediately obvious. How do you prove correctness to someone?
- **Cleaner idea: Use locks for mutual exclusion and condition variables for scheduling constraints**
- **Definition: Monitor:** a lock and zero or more condition variables for managing concurrent access to shared data
 - Some languages like Java provide this natively
 - Most others use actual locks and condition variables

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.28

Monitor with Condition Variables



- **Lock**: 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
- **Condition Variable**: a queue of threads waiting for something *inside* a critical section
 - Key idea: make it possible to go to sleep inside critical section by atomically releasing lock at time we go to sleep
 - Contrast to semaphores: Can't wait inside critical section

9/23/15

Kubiatowicz CS162 @UCB Fall 2015

Lec 8.29

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
}
    
```

- Not very interesting use of "Monitor"
 - It only uses a lock with no condition variables
 - Cannot put consumer to sleep if no work!

9/23/15

Kubiatowicz CS162 @UCB Fall 2015

Lec 8.30

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)

9/23/15

Kubiatowicz CS162 @UCB Fall 2015

Lec 8.31

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);
}
    
```

9/23/15

Kubiatowicz CS162 @UCB Fall 2015

Lec 8.32

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
```

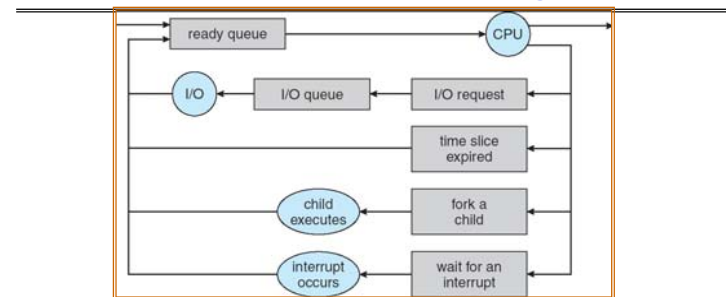
- Answer: depends on the type of scheduling
 - Hoare-style (most textbooks):
 - » Signaler gives lock, CPU to waiter; waiter runs immediately
 - » Waiter gives up lock, processor back to signaler when it exits critical section or if it waits again
 - Mesa-style (most real operating systems):
 - » Signaler keeps lock and processor
 - » Waiter placed on ready queue with no special priority
 - » Practically, need to check condition again after wait

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.33

Recall: CPU Scheduling



- Earlier, we talked about the life-cycle of a thread
 - Active threads work their way from Ready queue to Running to various waiting queues.
- Question: How is the OS to decide which of several tasks to take off a queue?
 - Obvious queue to worry about is ready queue
 - Others can be scheduled as well, however
- **Scheduling**: deciding which threads are given access to resources from moment to moment

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.34

Scheduling Assumptions

- CPU scheduling big area of research in early 70's
- Many implicit assumptions for CPU scheduling:
 - One program per user
 - One thread per program
 - Programs are independent
- Clearly, these are unrealistic but they simplify the problem so it can be solved
 - For instance: is "fair" about fairness among users or programs?
 - » If I run one compilation job and you run five, you get five times as much CPU on many operating systems
- The high-level goal: Dole out CPU time to optimize some desired parameters of system



9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.35

Scheduling Policy Goals/Criteria

- **Minimize Response Time**
 - Minimize elapsed time to do an operation (or job)
 - Response time is what the user sees:
 - » Time to echo a keystroke in editor
 - » Time to compile a program
 - » Real-time Tasks: Must meet deadlines imposed by World
- **Maximize Throughput**
 - Maximize operations (or jobs) per second
 - Throughput related to response time, but not identical:
 - » Minimizing response time will lead to more context switching than if you only maximized throughput
 - Two parts to maximizing throughput
 - » Minimize overhead (for example, context-switching)
 - » Efficient use of resources (CPU, disk, memory, etc)
- **Fairness**
 - Share CPU among users in some equitable way
 - Fairness is not minimizing average response time:
 - » Better *average* response time by making system *less* fair

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.36

First-Come, First-Served (FCFS) Scheduling

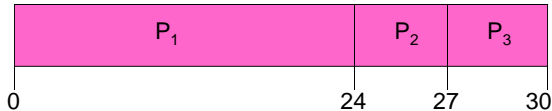
- **First-Come, First-Served (FCFS)**
 - Also "First In, First Out" (FIFO) or "Run until done"
 - » In early systems, FCFS meant one program scheduled until done (including I/O)
 - » Now, means keep CPU until thread blocks



- **Example:**

Process	Burst Time
P_1	24
P_2	3
P_3	3

- Suppose processes arrive in the order: P_1, P_2, P_3
The Gantt Chart for the schedule is:



- Waiting time for $P_1 = 0$; $P_2 = 24$; $P_3 = 27$
- Average waiting time: $(0 + 24 + 27)/3 = 17$
- Average Completion time: $(24 + 27 + 30)/3 = 27$

- **Convoy effect:** short process behind long process

FCFS Scheduling (Cont.)

- **Example continued:**

- Suppose that processes arrive in order: P_2, P_3, P_1
Now, the Gantt chart for the schedule is:



- Waiting time for $P_1 = 6$; $P_2 = 0$; $P_3 = 3$
- Average waiting time: $(6 + 0 + 3)/3 = 3$
- Average Completion time: $(3 + 6 + 30)/3 = 13$

- **In second case:**

- average waiting time is much better (before it was 17)
- Average completion time is better (before it was 27)

- **FIFO Pros and Cons:**

- Simple (+)
- Short jobs get stuck behind long ones (-)
 - » Safeway: Getting milk, always stuck behind cart full of small items. Upside: get to read about space aliens!

Round Robin (RR)

- **FCFS Scheme: Potentially bad for short jobs!**

- Depends on submit order
- If you are first in line at supermarket with milk, you don't care who is behind you, on the other hand...

- **Round Robin Scheme**

- Each process gets a small unit of CPU time (*time quantum*), usually 10-100 milliseconds
- After quantum expires, the process is preempted and added to the end of the ready queue.
- n processes in ready queue and time quantum is $q \Rightarrow$
 - » Each process gets $1/n$ of the CPU time
 - » In chunks of at most q time units
 - » **No process waits more than $(n-1)q$ time units**



- **Performance**

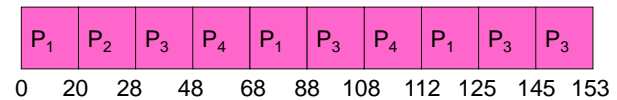
- q large \Rightarrow FCFS
- q small \Rightarrow Interleaved (really small \Rightarrow hyperthreading?)
- q must be large with respect to context switch, otherwise overhead is too high (all overhead)

Example of RR with Time Quantum = 20

- **Example:**

Process	Burst Time
P_1	53
P_2	8
P_3	68
P_4	24

- The Gantt chart is:



- Waiting time for $P_1 = (68-20) + (112-88) = 72$
 $P_2 = (20-0) = 20$
 $P_3 = (28-0) + (88-48) + (125-108) = 85$
 $P_4 = (48-0) + (108-68) = 88$

- Average waiting time = $(72+20+85+88)/4 = 66\frac{1}{4}$
- Average completion time = $(125+28+153+112)/4 = 104\frac{1}{2}$

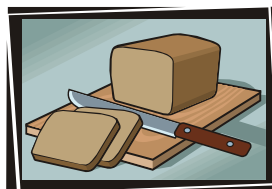
- **Thus, Round-Robin Pros and Cons:**

- Better for short jobs, Fair (+)
- Context-switching time adds up for long jobs (-)

Round-Robin Discussion

How do you choose time slice?

- What if too big?
 - » Response time suffers
- What if infinite (∞)?
 - » Get back FIFO
- What if time slice too small?
 - » Throughput suffers!



Actual choices of timeslice:

- Initially, UNIX timeslice one second:
 - » Worked ok when UNIX was used by one or two people.
 - » What if three compilations going on? 3 seconds to echo each keystroke!
- In practice, need to balance short-job performance and long-job throughput:
 - » Typical time slice today is between **10ms - 100ms**
 - » Typical context-switching overhead is **0.1ms - 1ms**
 - » Roughly **1%** overhead due to context-switching

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.41

Comparisons between FCFS and Round Robin

- Assuming zero-cost context-switching time, is RR always better than FCFS?
- Simple example: 10 jobs, each take 100s of CPU time
RR scheduler quantum of 1s
All jobs start at the same time

Completion Times:

Job #	FIFO	RR
1	100	991
2	200	992
...
9	900	999
10	1000	1000

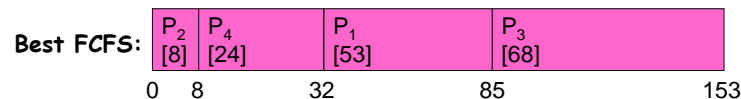
- Both RR and FCFS finish at the same time
- Average response time is much worse under RR!
 - » Bad when all jobs same length
- Also: Cache state must be shared between all jobs with RR but can be devoted to each job with FIFO
 - Total time for RR longer even for zero-cost switch!

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.42

Earlier Example with Different Time Quantum



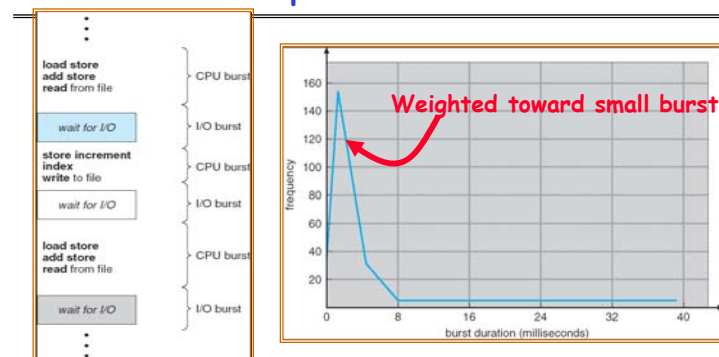
	Quantum	P ₁	P ₂	P ₃	P ₄	Average
Wait Time	Best FCFS	32	0	85	8	31¼
	Q = 1	84	22	85	57	62
	Q = 5	82	20	85	58	61¼
	Q = 8	80	8	85	56	57¼
	Q = 10	82	10	85	68	61¼
	Q = 20	72	20	85	88	66¼
	Worst FCFS	68	145	0	121	83½
Completion Time	Best FCFS	85	8	153	32	69½
	Q = 1	137	30	153	81	100½
	Q = 5	135	28	153	82	99½
	Q = 8	133	16	153	80	95½
	Q = 10	135	18	153	92	99½
	Q = 20	125	28	153	112	104½
	Worst FCFS	121	153	68	145	121¾

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.43

Assumption: CPU Bursts



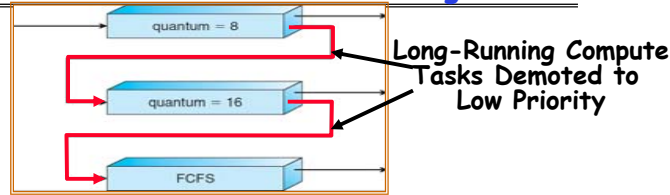
- Execution model: programs alternate between bursts of CPU and I/O
 - Program typically uses the CPU for some period of time, then does I/O, then uses CPU again
 - Each scheduling decision is about which job to give to the CPU for use by its next CPU burst
 - With timeslicing, thread may be forced to give up CPU before finishing current CPU burst

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.44

First peak at responsiveness scheduler: Multi-Level Feedback Scheduling



- A method for exploiting past behavior
 - First used in CTSS
 - **Multiple queues, each with different priority**
 - » Higher priority queues often considered "foreground" tasks
 - **Each queue has its own scheduling algorithm**
 - » e.g. foreground - RR, background - FCFS
 - » Sometimes multiple RR priorities with quantum increasing exponentially (highest: 1ms, next: 2ms, next: 4ms, etc)
- Adjust each job's priority as follows (details vary)
 - Job starts in highest priority queue
 - If timeout expires, drop one level
 - If timeout doesn't expire, push up one level (or to top)

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.45

Summary (1/2)

- Important concept: Atomic Operations
 - An operation that runs to completion or not at all
 - These are the primitives on which to construct various synchronization primitives
- Talked about hardware atomicity primitives:
 - Disabling of Interrupts, test&set, swap, comp&swap, load-linked/store conditional
- Showed several constructions of Locks
 - Must be very careful not to waste/tie up machine resources
 - » Shouldn't disable interrupts for long
 - » Shouldn't spin wait for long
 - Key idea: Separate lock variable, use hardware mechanisms to protect modifications of that variable

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.46

Summary (2/2)

- Semaphores: Like integers with restricted interface
 - Two operations:
 - » P(): Wait if zero; decrement when becomes non-zero
 - » V(): Increment and wake a sleeping task (if exists)
 - » Can initialize value to any non-negative value
 - Use separate semaphore for each constraint
- 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()
- Scheduling: selecting a waiting process from the ready queue and allocating the CPU to it
- FCFS Scheduling:
 - Run threads to completion in order of submission
 - Pros: Simple
 - Cons: Short jobs get stuck behind long ones
- Round-Robin Scheduling:
 - Give each thread a small amount of CPU time when it executes; cycle between all ready threads
 - Pros: Better for short jobs
 - Cons: Poor when jobs are same length

9/23/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 8.47