

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>

Acknowledgments: Lecture slides are from the Operating Systems course taught by John Kubiawicz at Berkeley, with few minor updates/changes. When slides are obtained from other sources, a reference will be noted on the bottom of that slide, in which case a full list of references is provided on the last slide.

Review: Synchronization problem with Threads

- One thread per transaction, each running:

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

- Unfortunately, shared state can get corrupted:

Thread 1

```
load r1, acct->balance  
  
add r1, amount1  
store r1, acct->balance
```

Thread 2

```
load r1, acct->balance  
add r1, amount2  
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

Review: Too Much Milk Solution #3

- Here is a possible two-note solution:

Thread A

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

Thread B

```
leave note B;
if (noNote A) {\\Y
    if (noMilk) {
        buy milk;
    }
}
remove note B;
```

- 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

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"**

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 Kubiatoiwicz.

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;  
}
```


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;
}`

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

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!



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!

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;  
}
```

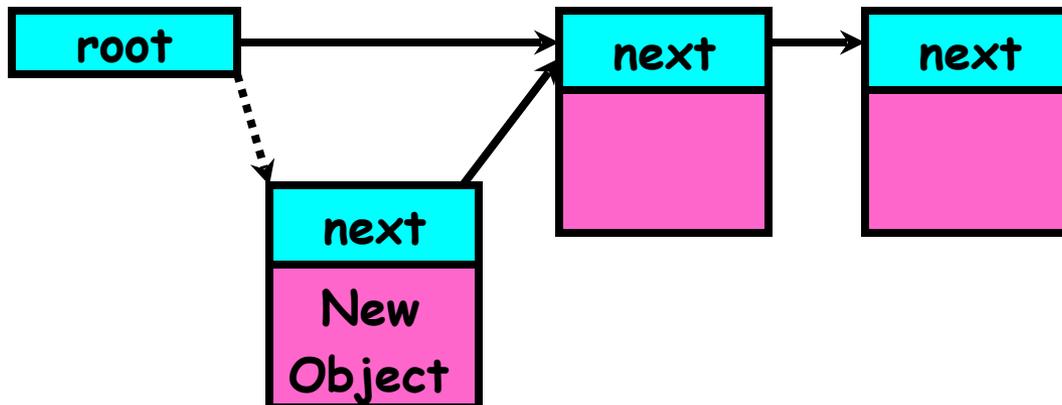
- Note: sleep has to be sure to reset the guard variable
 - Why can't we do it just before or just after the sleep?

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));  
}
```



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

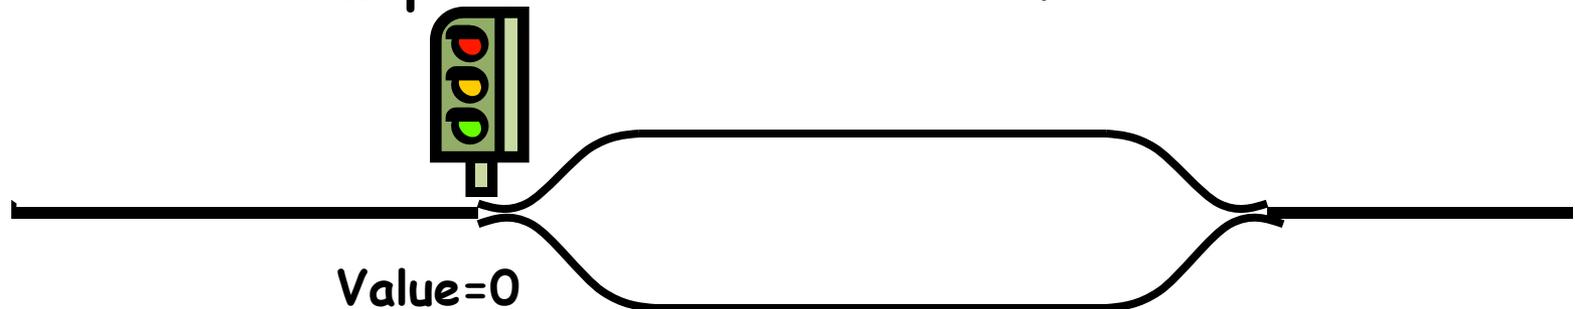
Semaphores



- 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

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:



Two Uses of Semaphores

- **Mutual Exclusion (initial value = 1)**

- Also called "Binary Semaphore".
- Can be used for mutual exclusion:

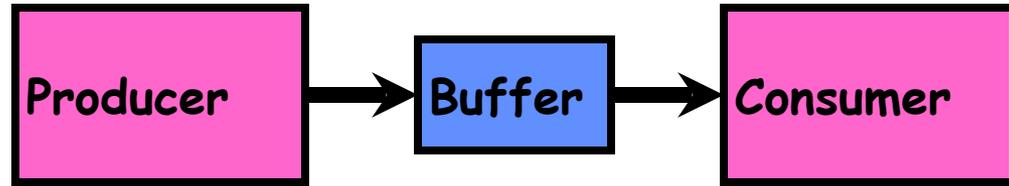
```
semaphore.P();  
// Critical section goes here  
semaphore.V();
```

- **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();  
}
```

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



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

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;
}
```

Discussion about Solution

- Why asymmetry?
 - Producer does: `emptyBuffer.P()`, `fullBuffer.V()`
 - Consumer does: `fullBuffer.P()`, `emptyBuffer.V()`
- Is order of P's important?
 - Yes! Can cause deadlock

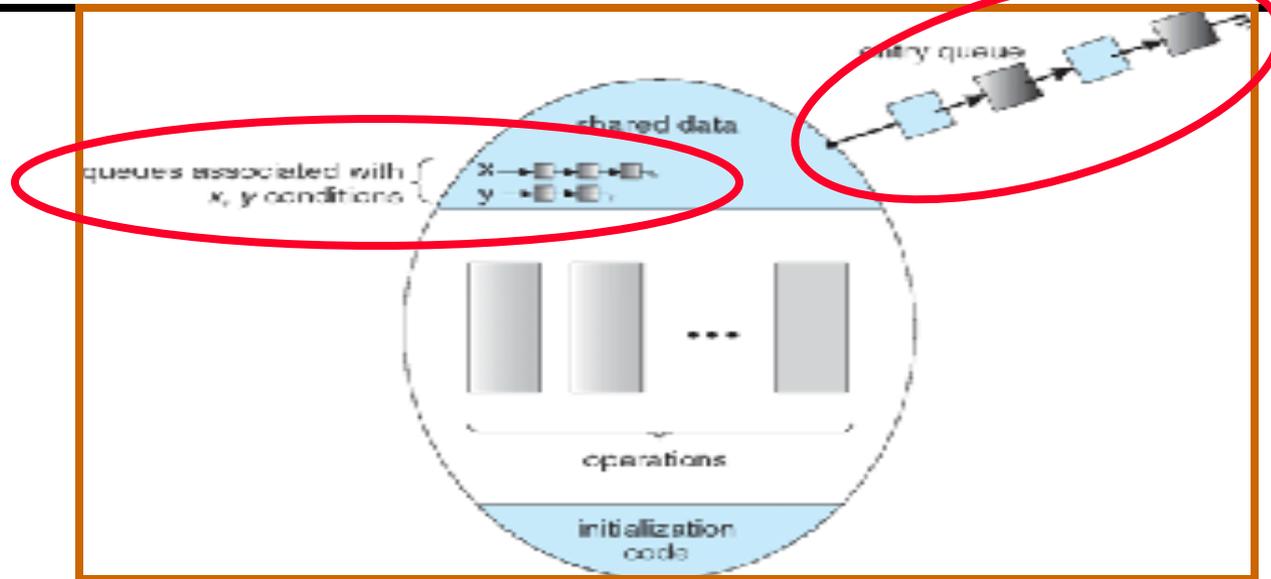
```
Producer(item) {  
    Mutex.P();           // Wait until buffer free  
    emptyBuffers.P();    // Could Wait forever!  
    Enqueue(item);  
    mutex.V();  
    fullBuffers.V();     // Tell consumers more coke  
}
```

- Is order of V's important?
 - No, except that it might affect scheduling efficiency
- What if we have 2 producers or 2 consumers?
 - Do we need to change anything?

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

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

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!

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!

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);
}
```

Mesa vs. Hoare monitors

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

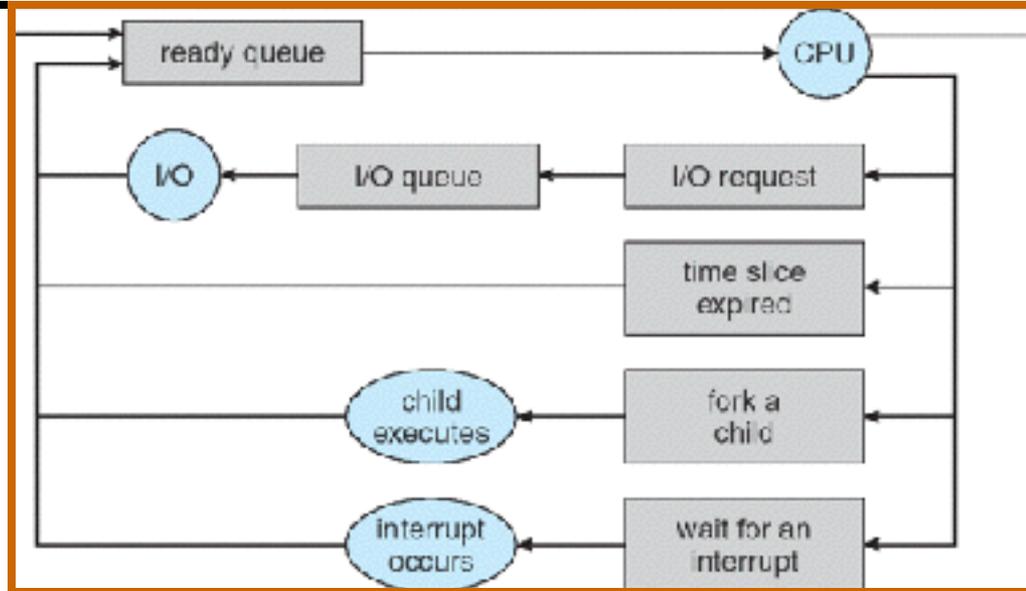
```
while (queue.isEmpty()) {  
    dataready.wait(&lock); // If nothing, sleep  
}  
item = queue.dequeue(); // Get next item
```

- 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

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

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



Time 

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

First-Come, First-Served (FCFS) Scheduling

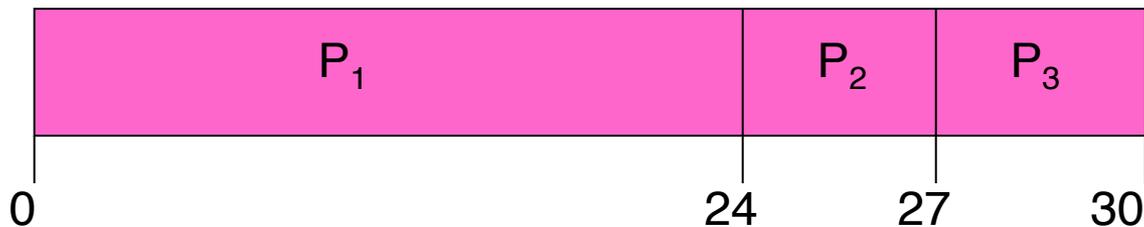
- **First-Come, First-Served (FCFS)**
 - Also "First In, First Out" (FIFO) or "Run until done"
 - » In early systems, FCFS meant one program scheduled until done (including I/O)
 - » Now, means keep CPU until thread blocks



• **Example:**

<u>Process</u>	<u>Burst Time</u>
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:

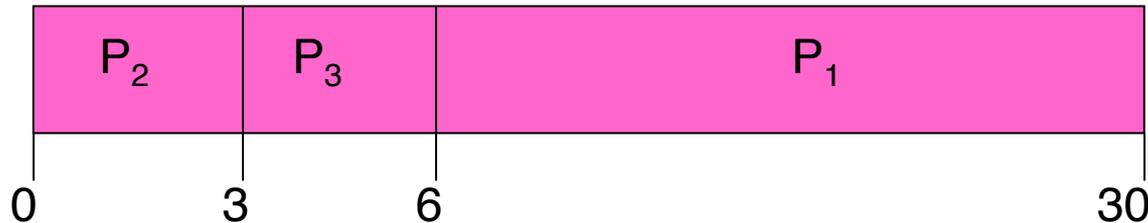


- Waiting time for $P_1 = 0$; $P_2 = 24$; $P_3 = 27$
 - Average waiting time: $(0 + 24 + 27)/3 = 17$
 - Average Completion time: $(24 + 27 + 30)/3 = 27$
- **Convoy effect: short process behind long process**

FCFS Scheduling (Cont.)

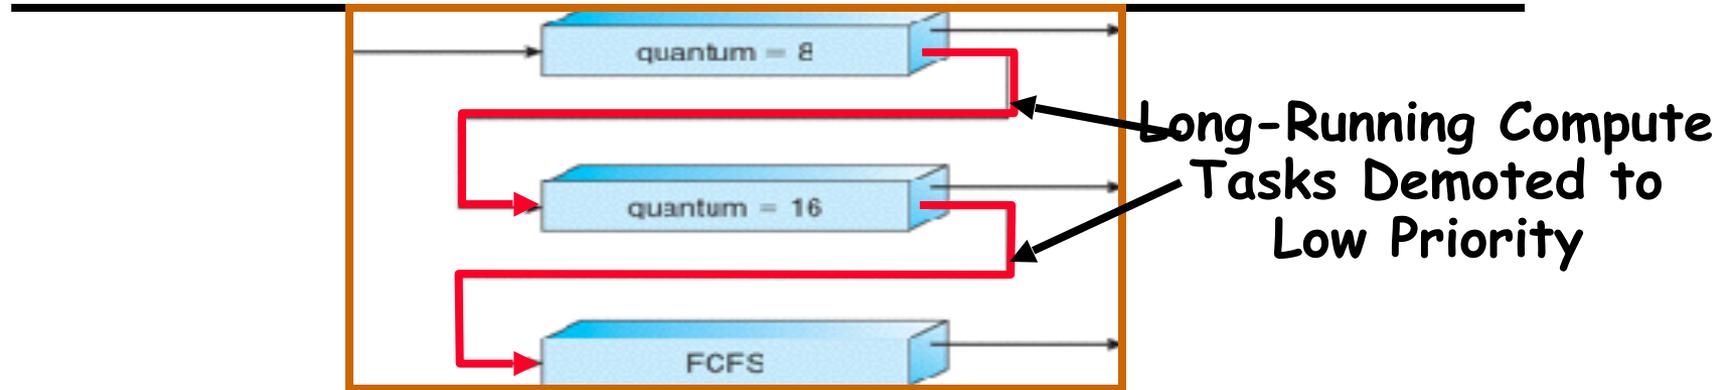
- Example continued:

- 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.

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)

Summary

- **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