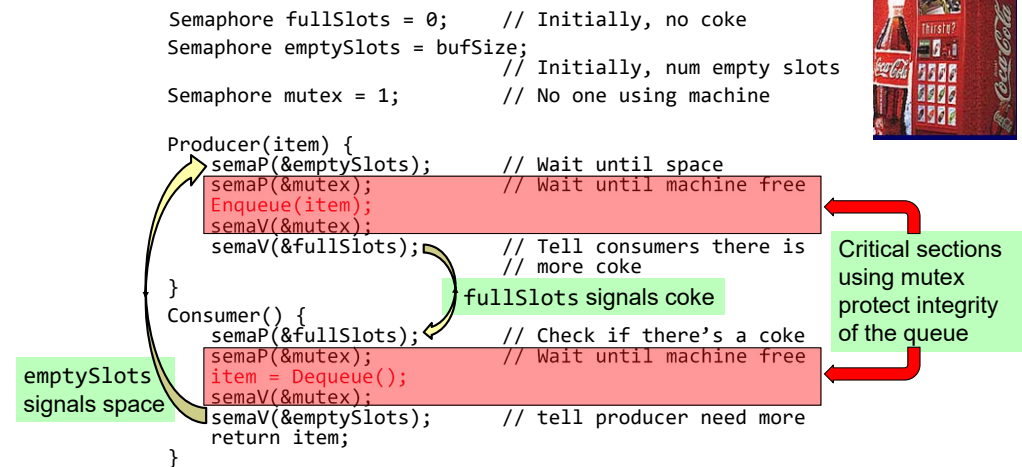


CS162  
Operating Systems and  
Systems Programming  
Lecture 10

Synchronization 4: Readers/Writers  
Scheduling Intro: Pintos Concurrency, FCFS

February 20<sup>st</sup>, 2024  
Prof. John Kubiatowicz  
<http://cs162.eecs.Berkeley.edu>

Recall: Bounded Buffer, 3<sup>rd</sup> cut (coke machine)



2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.2

Recall: Monitors and Condition Variables

- **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
- **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!**

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.3

Recall: Bounded Buffer – 4<sup>th</sup> cut (Monitors, pthread-like)

```

lock buf_lock = <initially unlocked>
condition producer_CV = <initially empty>
condition consumer_CV = <initially empty>

Producer(item) {
    acquire(&buf_lock);
    while (buffer full) { cond_wait(&producer_CV, &buf_lock); }
    enqueue(item);
    cond_signal(&consumer_CV);
    release(&buf_lock);
}

Consumer() {
    acquire(&buf_lock);
    while (buffer empty) { cond_wait(&consumer_CV, &buf_lock); }
    item = dequeue();
    cond_signal(&producer_CV);
    release(&buf_lock);
    return item;
}

```

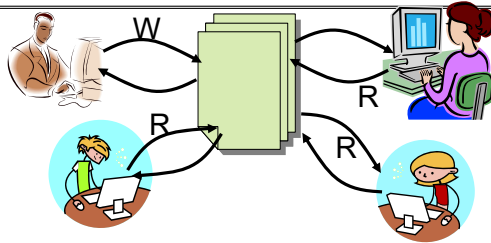
Annotation: **What does thread do when it is waiting? - Sleep, not busywait!**

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.4

## 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

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.5

## Basic Structure of Mesa Monitor Program

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

```
lock
while (need to wait) {
    condvar.wait();
}
unlock

do something so no need to wait

lock

condvar.signal();
unlock
```

} Check and/or update  
state variables  
Wait if necessary

} Check and/or update  
state variables

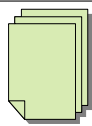
2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.6

## 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



2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.7

## Code for a Reader

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

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.8

## Code for a Writer

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

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.9

## Simulation of Readers/Writers Solution

- Use an example to simulate the solution
- Consider the following sequence of operators:
  - R1, R2, W1, R3
- Initially: AR = 0, WR = 0, AW = 0, WW = 0

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.10

## Simulation of Readers/Writers Solution

- R1 comes along (no waiting threads)
- AR = 0, WR = 0, AW = 0, WW = 0

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond_wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond_signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.11

## Simulation of Readers/Writers Solution

- R1 comes along (no waiting threads)
- AR = 0, WR = 0, AW = 0, WW = 0

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond_wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond_signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.12

## Simulation of Readers/Writers Solution

- R1 comes along (no waiting threads)
- AR = 1, WR = 0, AW = 0, WW = 0

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.13

## Simulation of Readers/Writers Solution

- R1 comes along (no waiting threads)
- AR = 1, WR = 0, AW = 0, WW = 0

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.14

## Simulation of Readers/Writers Solution

- R1 accessing dbase (no other threads)
- AR = 1, WR = 0, AW = 0, WW = 0

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.15

## Simulation of Readers/Writers Solution

- R2 comes along (R1 accessing dbase)
- AR = 1, WR = 0, AW = 0, WW = 0

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.16

## Simulation of Readers/Writers Solution

- R2 comes along (R1 accessing dbase)
- AR = 1, WR = 0, AW = 0, WW = 0

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.17

## Simulation of Readers/Writers Solution

- R2 comes along (R1 accessing dbase)
- AR = 2, WR = 0, AW = 0, WW = 0

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.18

## Simulation of Readers/Writers Solution

- R2 comes along (R1 accessing dbase)
- AR = 2, WR = 0, AW = 0, WW = 0

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.19

## Simulation of Readers/Writers Solution

- R1 and R2 accessing dbase
- AR = 2, WR = 0, AW = 0, WW = 0

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
```

Assume readers take a while to access database  
Situation: Locks released, only AR is non-zero

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.20

## Simulation of Readers/Writers Solution

- W1 comes along (R1 and R2 are still accessing dbase)
- AR = 2, WR = 0, AW = 0, WW = 0

```
Writer() {
    acquire(&lock);
    while ((AW + AR) > 0) { // Is it safe to write?
        WW++; // No. Active users exist
        cond_wait(&okToWrite, &lock); // Sleep on cond var
        WW--; // No longer waiting
    }
    AW++;
    release(&lock);

    AccessDBase(ReadWrite);

    acquire(&lock);
    AW--;
    if (WW > 0) {
        cond_signal(&okToWrite);
    } else if (WR > 0) {
        cond_broadcast(&okToRead);
    }
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.21

## Simulation of Readers/Writers Solution

- W1 comes along (R1 and R2 are still accessing dbase)
- AR = 2, WR = 0, AW = 0, WW = 0

```
Writer() {
    acquire(&lock);
    while ((AW + AR) > 0) { // Is it safe to write?
        WW++; // No. Active users exist
        cond_wait(&okToWrite, &lock); // Sleep on cond var
        WW--; // No longer waiting
    }
    AW++;
    release(&lock);

    AccessDBase(ReadWrite);

    acquire(&lock);
    AW--;
    if (WW > 0) {
        cond_signal(&okToWrite);
    } else if (WR > 0) {
        cond_broadcast(&okToRead);
    }
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.22

## Simulation of Readers/Writers Solution

- W1 comes along (R1 and R2 are still accessing dbase)
- AR = 2, WR = 0, AW = 0, WW = 1

```
Writer() {
    acquire(&lock);
    while ((AW + AR) > 0) { // Is it safe to write?
        WW++; // No. Active users exist
        cond_wait(&okToWrite, &lock); // Sleep on cond var
        WW--; // No longer waiting
    }
    AW++;
    release(&lock);

    AccessDBase(ReadWrite);

    acquire(&lock);
    AW--;
    if (WW > 0) {
        cond_signal(&okToWrite);
    } else if (WR > 0) {
        cond_broadcast(&okToRead);
    }
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.23

## Simulation of Readers/Writers Solution

- R3 comes along (R1 and R2 accessing dbase, W1 waiting)
- AR = 2, WR = 0, AW = 0, WW = 1

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond_wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0) {
        cond_signal(&okToWrite);
    }
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.24

## Simulation of Readers/Writers Solution

- R3 comes along (R1 and R2 accessing dbase, W1 waiting)
- AR = 2, WR = 0, AW = 0, WW = 1

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.25

## Simulation of Readers/Writers Solution

- R3 comes along (R1 and R2 accessing dbase, W1 waiting)
- AR = 2, WR = 1, AW = 0, WW = 1

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    lock.release();

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.26

## Simulation of Readers/Writers Solution

- R3 comes along (R1, R2 accessing dbase, W1 waiting)
- AR = 2, WR = 1, AW = 0, WW = 1

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.27

## Simulation of Readers/Writers Solution

- R1 and R2 accessing dbase, W1 and R3 waiting
- AR = 2, WR = 1, AW = 0, WW = 1

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
```

Status:

- R1 and R2 still reading
- W1 and R3 waiting on okToWrite and okToRead, respectively

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.28

## Simulation of Readers/Writers Solution

- R2 finishes (R1 accessing dbase, W1 and R3 waiting)
- AR = 2, WR = 1, AW = 0, WW = 1

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.29

## Simulation of Readers/Writers Solution

- R2 finishes (R1 accessing dbase, W1 and R3 waiting)
- AR = 1, WR = 1, AW = 0, WW = 1

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.30

## Simulation of Readers/Writers Solution

- R2 finishes (R1 accessing dbase, W1 and R3 waiting)
- AR = 1, WR = 1, AW = 0, WW = 1

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.31

## Simulation of Readers/Writers Solution

- R2 finishes (R1 accessing dbase, W1 and R3 waiting)
- AR = 1, WR = 1, AW = 0, WW = 1

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.32

## Simulation of Readers/Writers Solution

- R1 finishes (W1 and R3 waiting)
- AR = 1, WR = 1, AW = 0, WW = 1

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.33

## Simulation of Readers/Writers Solution

- R1 finishes (W1, R3 waiting)
- AR = 0, WR = 1, AW = 0, WW = 1

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.34

## Simulation of Readers/Writers Solution

- R1 finishes (W1, R3 waiting)
- AR = 0, WR = 1, AW = 0, WW = 1

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.35

## Simulation of Readers/Writers Solution

- R1 signals a writer (W1 and R3 waiting)
- AR = 0, WR = 1, AW = 0, WW = 1

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.36

## Simulation of Readers/Writers Solution

- W1 gets signal (R3 still waiting)
- AR = 0, WR = 1, AW = 0, WW = 1

```
Writer() {
    acquire(&lock);
    while ((AW + AR) > 0) { // Is it safe to write?
        WW++; // No. Active users exist
        cond_wait(&okToWrite, &lock); // Sleep on cond var
        WW--; // No longer waiting
    }
    AW++;
    release(&lock);

    AccessDBase(ReadWrite);

    acquire(&lock);
    AW--;
    if (WW > 0) {
        cond_signal(&okToWrite);
    } else if (WR > 0) {
        cond_broadcast(&okToRead);
    }
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.37

## Simulation of Readers/Writers Solution

- W1 gets signal (R3 still waiting)
- AR = 0, WR = 1, AW = 0, WW = 0

```
Writer() {
    acquire(&lock);
    while ((AW + AR) > 0) { // Is it safe to write?
        WW++; // No. Active users exist
        cond_wait(&okToWrite, &lock); // Sleep on cond var
        WW--; // No longer waiting
    }
    AW++;
    release(&lock);

    AccessDBase(ReadWrite);

    acquire(&lock);
    AW--;
    if (WW > 0) {
        cond_signal(&okToWrite);
    } else if (WR > 0) {
        cond_broadcast(&okToRead);
    }
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.38

## Simulation of Readers/Writers Solution

- W1 gets signal (R3 still waiting)
- AR = 0, WR = 1, AW = 1, WW = 0

```
Writer() {
    acquire(&lock);
    while ((AW + AR) > 0) { // Is it safe to write?
        WW++; // No. Active users exist
        cond_wait(&okToWrite, &lock); // Sleep on cond var
        WW--; // No longer waiting
    }
    AW++;
    release(&lock);

    AccessDBase(ReadWrite);

    acquire(&lock);
    AW--;
    if (WW > 0) {
        cond_signal(&okToWrite);
    } else if (WR > 0) {
        cond_broadcast(&okToRead);
    }
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.39

## Simulation of Readers/Writers Solution

- W1 accessing dbase (R3 still waiting)
- AR = 0, WR = 1, AW = 1, WW = 0

```
Writer() {
    acquire(&lock);
    while ((AW + AR) > 0) { // Is it safe to write?
        WW++; // No. Active users exist
        cond_wait(&okToWrite, &lock); // Sleep on cond var
        WW--; // No longer waiting
    }
    AW++;
    release(&lock);

    AccessDBase(ReadWrite);

    acquire(&lock);
    AW--;
    if (WW > 0) {
        cond_signal(&okToWrite);
    } else if (WR > 0) {
        cond_broadcast(&okToRead);
    }
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.40

## Simulation of Readers/Writers Solution

- W1 finishes (R3 still waiting)
- AR = 0, WR = 1, AW = 1, WW = 0

```
Writer() {
    acquire(&lock);
    while ((AW + AR) > 0) { // Is it safe to write?
        WW++; // No. Active users exist
        cond_wait(&okToWrite, &lock); // Sleep on cond var
        WW--; // No longer waiting
    }
    AW++;
    release(&lock);

    AccessDBase(ReadWrite);

    acquire(&lock);
    AW--;
    if (WW > 0) {
        cond_signal(&okToWrite);
    } else if (WR > 0) {
        cond_broadcast(&okToRead);
    }
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.41

## Simulation of Readers/Writers Solution

- W1 finishes (R3 still waiting)
- AR = 0, WR = 1, AW = 0, WW = 0

```
Writer() {
    acquire(&lock);
    while ((AW + AR) > 0) { // Is it safe to write?
        WW++; // No. Active users exist
        cond_wait(&okToWrite, &lock); // Sleep on cond var
        WW--; // No longer waiting
    }
    AW++;
    release(&lock);

    AccessDBase(ReadWrite);

    acquire(&lock);
    AW--;
    if (WW > 0) {
        cond_signal(&okToWrite);
    } else if (WR > 0) {
        cond_broadcast(&okToRead);
    }
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.42

## Simulation of Readers/Writers Solution

- W1 finishes (R3 still waiting)
- AR = 0, WR = 1, AW = 0, WW = 0

```
Writer() {
    acquire(&lock);
    while ((AW + AR) > 0) { // Is it safe to write?
        WW++; // No. Active users exist
        cond_wait(&okToWrite, &lock); // Sleep on cond var
        WW--; // No longer waiting
    }
    AW++;
    release(&lock);

    AccessDBase(ReadWrite);

    acquire(&lock);
    AW--;
    if (WW > 0) {
        cond_signal(&okToWrite);
    } else if (WR > 0) {
        cond_broadcast(&okToRead);
    }
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.43

## Simulation of Readers/Writers Solution

- W1 signaling readers (R3 still waiting)
- AR = 0, WR = 1, AW = 0, WW = 0

```
Writer() {
    acquire(&lock);
    while ((AW + AR) > 0) { // Is it safe to write?
        WW++; // No. Active users exist
        cond_wait(&okToWrite, &lock); // Sleep on cond var
        WW--; // No longer waiting
    }
    AW++;
    release(&lock);

    AccessDBase(ReadWrite);

    acquire(&lock);
    AW--;
    if (WW > 0) {
        cond_signal(&okToWrite);
    } else if (WR > 0) {
        cond_broadcast(&okToRead);
    }
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.44

## Simulation of Readers/Writers Solution

- R3 gets signal (no waiting threads)
- AR = 0, WR = 1, AW = 0, WW = 0

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.45

## Simulation of Readers/Writers Solution

- R3 gets signal (no waiting threads)
- AR = 0, WR = 0, AW = 0, WW = 0

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.46

## Simulation of Readers/Writers Solution

- R3 accessing dbase (no waiting threads)
- AR = 1, WR = 0, AW = 0, WW = 0

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.47

## Simulation of Readers/Writers Solution

- R3 finishes (no waiting threads)
- AR = 1, WR = 0, AW = 0, WW = 0

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.48

## Simulation of Readers/Writers Solution

- R3 finishes (no waiting threads)
- AR = 0, WR = 0, AW = 0, WW = 0

```
Reader() {
    acquire(&lock);
    while ((AW + WW) > 0) { // Is it safe to read?
        WR++; // No. Writers exist
        cond_wait(&okToRead, &lock); // Sleep on cond var
        WR--; // No longer waiting
    }
    AR++; // Now we are active!
    release(&lock);

    AccessDBase(ReadOnly);

    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond_signal(&okToWrite);
    release(&lock);
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.49

## Administrivia

- Still grading Midterm 1 (Sorry)
  - Finishing soon!
  - Solutions also will be up soon.
- Homework #2 due Thursday, 2/22
- Homework #3: Released on Friday 2/23
  - Option to do this in Rust!
  - Rust crash course on Monday 2/26
- Professor Kubi's office hours changed:
  - Monday (1:00-2:00PM), Thursday (3:00-4:00PM)
  - 673 Soda Hall
- FYI: Next Midterm on 3/14
  - PI Day!

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.50

## Questions

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

```
while ((AW + WW) > 0) { // Is it safe to read?
    WR++; // No. Writers exist
    cond_wait(&okToRead, &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
    cond_signal(&okToWrite); // Wake up one writer
```
- Further, what if we turn the signal() into broadcast()
 

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

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.51

## Use of Single CV: okContinue

<pre>Reader() {     // check into system     acquire(&amp;lock);     while ((AW + WW) &gt; 0) {         WR++;         cond_wait(&amp;okContinue, &amp;lock);         WR--;     }     AR++;     release(&amp;lock);      // read-only access     AccessDBase(ReadOnly);      // check out of system     acquire(&amp;lock);     AR--;     if (AR == 0 &amp;&amp; WW &gt; 0)         cond_signal(&amp;okContinue);     release(&amp;lock); }</pre>	<pre>Writer() {     // check into system     acquire(&amp;lock);     while ((AW + AR) &gt; 0) {         WW++;         cond_wait(&amp;okContinue, &amp;lock);         WW--;     }     AW++;     release(&amp;lock);      // read/write access     AccessDBase(ReadWrite);      // check out of system     acquire(&amp;lock);     AW--;     if (WW &gt; 0) {         cond_signal(&amp;okContinue);     } else if (WR &gt; 0) {         cond_broadcast(&amp;okContinue);     }     release(&amp;lock); }</pre>
--	--

What if we turn okToWrite and okToRead into okContinue (i.e. use only one condition variable instead of two)?

2/20/24

Lec 10.52

## Use of Single CV: okContinue

```

Reader() {
    // check into system
    acquire(&lock);
    while ((AW + WW) > 0) {
        WR++;
        cond_wait(&okContinue, &lock);
        WR--;
    }
    AR++;
    release(&lock);

    // read-only access
    AccessDbase(ReadOnly);

    // check out of system
    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond_signal(&okContinue);
    release(&lock);
}

Writer() {
    // check into system
    acquire(&lock);
    while ((AW + AR) > 0) {
        WW++;
        cond_wait(&okContinue, &lock);
        WW--;
    }
    AW++;
    release(&lock);

    // read/write access
    AccessDbase(ReadWrite);

    // check out of system
    acquire(&lock);
    AW--;
    if (WW > 0) {
        cond_signal(&okContinue);
    } else if (WR > 0) {
        cond_broadcast(&okContinue);
    }
}

```

Consider this scenario:

- R1 arrives
- W1, R2 arrive while R1 still reading → W1 and R2 wait for R1 to finish
- Assume R1's signal is delivered to R2 (not W1)

2/20/24

Lec 10.53

## Use of Single CV: okContinue

```

Reader() {
    // check into system
    acquire(&lock);
    while ((AW + WW) > 0) {
        WR++;
        cond_wait(&okContinue, &lock);
        WR--;
    }
    AR++;
    release(&lock);

    // read-only access
    AccessDbase(ReadOnly);

    // check out of system
    acquire(&lock);
    AR--;
    if (AR == 0 && WW > 0)
        cond_broadcast(&okContinue);
    release(&lock);
}

Writer() {
    // check into system
    acquire(&lock);
    while ((AW + AR) > 0) {
        WW++;
        cond_wait(&okContinue, &lock);
        WW--;
    }
    AW++;
    release(&lock);

    // read/write access
    AccessDbase(ReadWrite);

    // check out of system
    acquire(&lock);
    AW--;
    if (WW > 0 || WR > 0) {
        cond_broadcast(&okContinue);
    }
    release(&lock);
}

```

Need to change to  
broadcast() !

Must broadcast()  
to sort things out!

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.54

## Can we construct Monitors from Semaphores?

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

```

Wait(Semaphore *thesema) { semaP(thesema); }
Signal(Semaphore *thesema) { semaV(thesema); }

```
- Does this work better?
 

```

Wait(Lock *thelock, Semaphore *thesema) {
    release(thelock);
    semaP(thesema);
    acquire(thelock);
}
Signal(Semaphore *thesema) {
    semaV(thesema);
}

```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.55

## 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 *thelock, Semaphore *thesema) {
    release(thelock);
    semaP(thesema);
    acquire(thelock);
}
Signal(Semaphore *thesema) {
    if semaphore queue is not empty
        semaV(thesema);
}

```

  - 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?

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

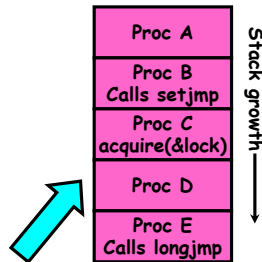
Lec 10.56

## 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() {
    acquire(&lock);
    ...
    if (exception) {
        release(&lock);
        return errReturnCode;
    }
    ...
    release(&lock);
    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!

## Concurrency and Synchronization in C

- Harder with more locks

```
void Rtn() {
    lock1.acquire();
    ...
    if (error) {
        lock1.release();
        return;
    }
    ...
    lock2.acquire();
    ...
    if (error) {
        lock2.release();
        lock1.release();
        return;
    }
    ...
    lock2.release();
    lock1.release();
}
```

- Is `goto` a solution???

```
void Rtn() {
    lock1.acquire();
    ...
    if (error) {
        goto release_lock1_and_return;
    }
    ...
    lock2.acquire();
    ...
    if (error) {
        goto release_both_and_return;
    }
    ...
    release_both_and_return:
        lock2.release();
        release_lock1_and_return:
            lock1.release();
}
```

## 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!

## 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;
    ...
}
```

## Much better: C++ Lock Guards

```
#include <mutex>
int global_i = 0;
std::mutex global_mutex;

void safe_increment() {
    std::lock_guard<std::mutex> lock(global_mutex);
    ...
    global_i++;
    // Mutex released when 'lock' goes out of scope
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.61

## Python with Keyword

- More versatile than we show here (can be used to close files, database connections, etc.)

```
lock = threading.Lock()

...
with lock: # Automatically calls acquire()
    some_var += 1
...
# release() called however we leave block
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.62

## Java synchronized Keyword

- Every Java object has an associated lock:
  - Lock is acquired on entry and released on exit from a **synchronized** method
  - Lock is properly released if exception occurs inside a **synchronized** method
  - Mutex execution of synchronized methods (beware deadlock)

```
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;
    }
}
```

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.63

## Java Support for Monitors

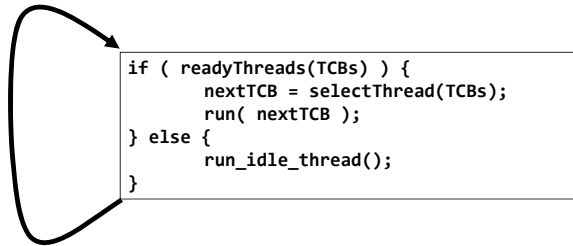
- Along with a lock, every object has a single condition variable associated with it
- To wait inside a synchronized method:
  - void wait();
  - void wait(long timeout);
- To signal while in a synchronized method:
  - void notify();
  - void notifyAll();

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.64

## Goal for Today



- Discussion of Scheduling:
  - Which thread should run on the CPU next?
- Scheduling goals, policies
- Look at a number of different schedulers

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.65

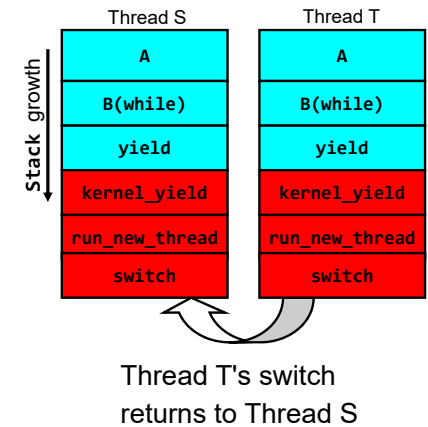
## Recall: Stacks for Yield with Multiple Threads

- Consider the following code blocks:

```

proc A() {
    B();
}
proc B() {
    while(TRUE) {
        yield();
    }
}

```

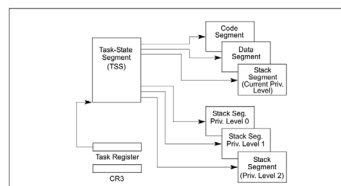


- Suppose we have 2 threads:
  - Threads S and T
  - Assume that both have been running for a while

Lec 10.66

## Hardware context switch support in x86

- Syscall/Intr (U → K)
  - PL 3 → 0;
  - TSS ← EFLAGS, CS:EIP;
  - SS:ESP ← k-thread stack (TSS PL 0);
  - push (old) SS:ESP onto (new) k-stack
  - push (old) EFLAGS, CS:EIP, <err>
  - CS:EIP ← <k target handler>
- Then
  - Handler saves other regs, etc
  - Does all its work, possibly choosing other threads, changing PTBR (CR3)
- iret (K → U)
  - PL 0 → 3;
  - EFLAGS, CS:EIP ← popped off k-stack
  - SS:ESP ← popped off k-stack



pg 2,942 of 4,922 of x86 reference manual

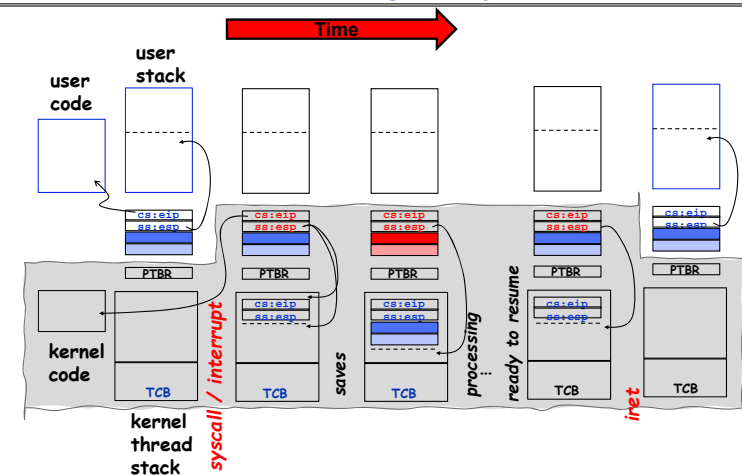
Pintos: tss.c, intr-stubs.S

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.67

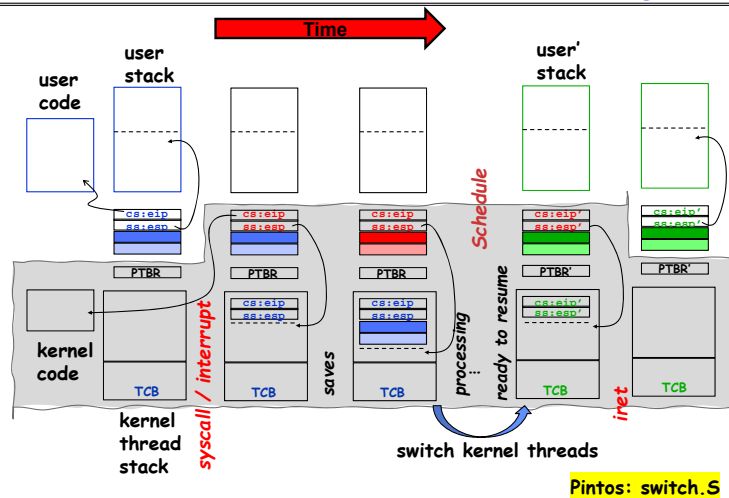
## Pintos: Kernel Crossing on Syscall or Interrupt



Kubiatowicz CS162 © UCB Spring 2024

Lec 10.68

## Pintos: Context Switch – Scheduling

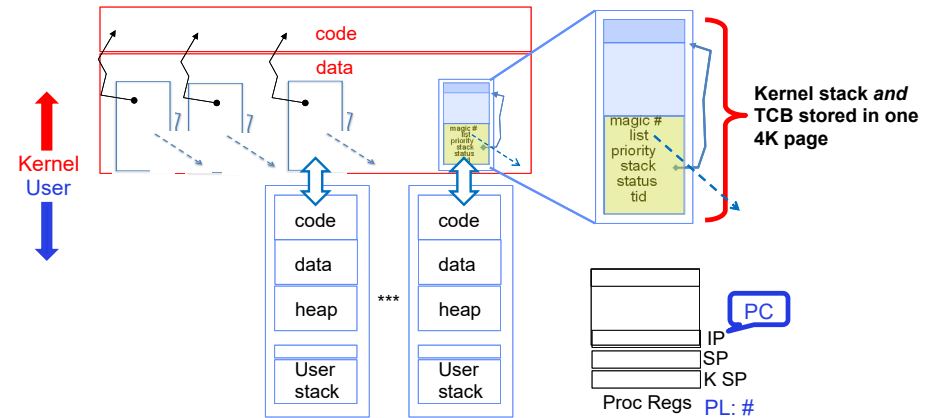


2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.69

## MT Kernel single Thread Process ala Pintos/x86



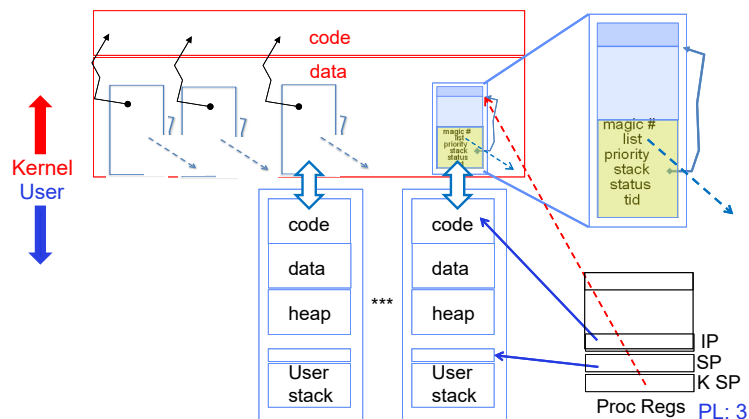
- Each user process/thread associated with a kernel thread, described by a 4KB page object containing TCB and kernel stack for the kernel thread

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.70

## Running User Code with Kernel stack waiting



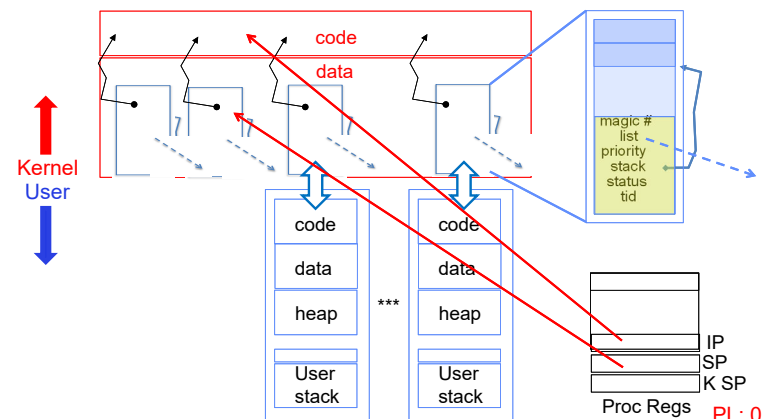
- x86 CPU holds interrupt SP in register
- During user thread execution, associated kernel thread is “standing by”

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.71

## In Kernel Thread: No User Component



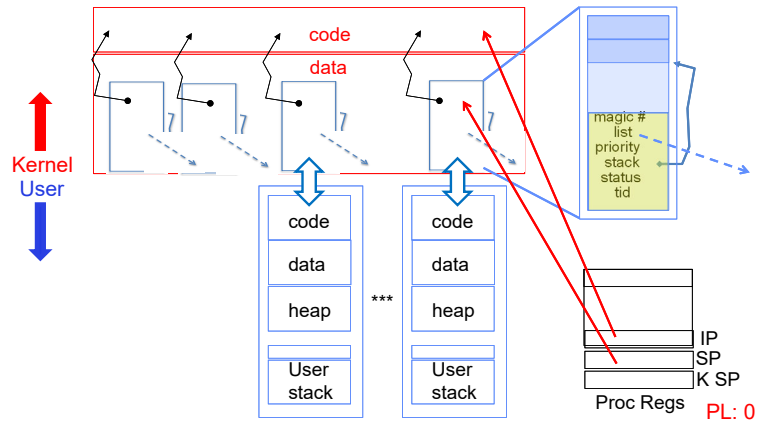
- Kernel threads execute with small stack in thread structure
- Pure kernel threads have no corresponding user-mode thread

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.72

## User → Kernel (interrupts, syscalls)



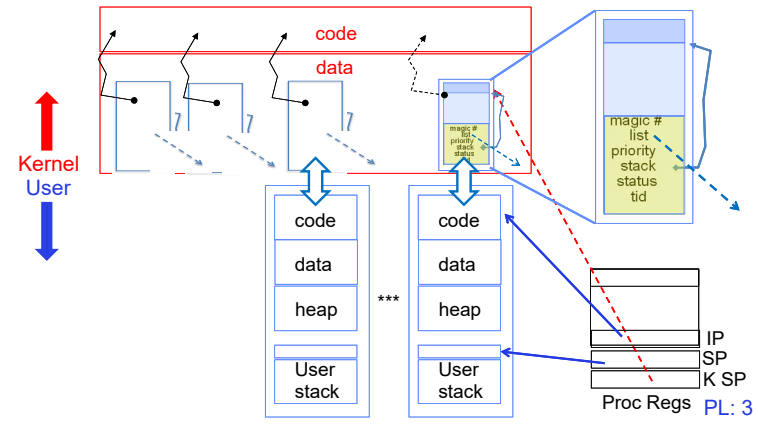
- Mechanism vectors through “interrupt vector”

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.73

## Kernel → User



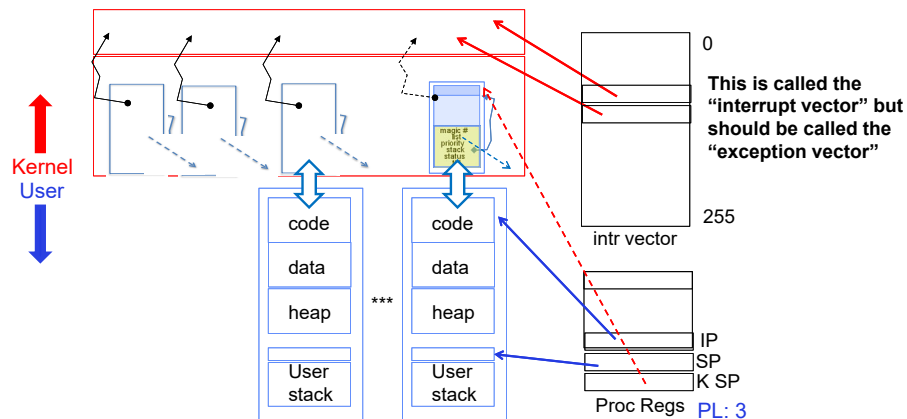
- Interrupt return (iret) restores user stack, IP, and PL

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.74

## User → Kernel via “interrupt vector” (interrupts & traps)



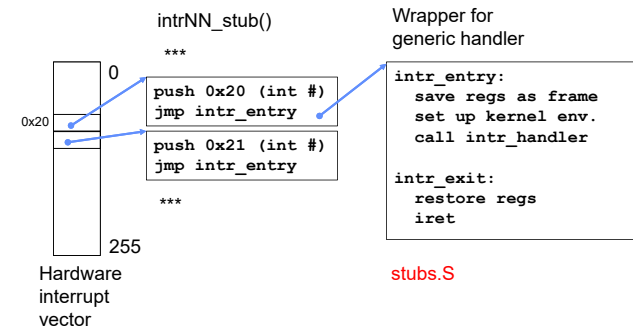
- Interrupts (timer) or trap (syscall, page fault) transfers through interrupt vector (IDT)
  - Each slot for different exception type

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.75

## Pintos Interrupt Processing for Timer (0x20)



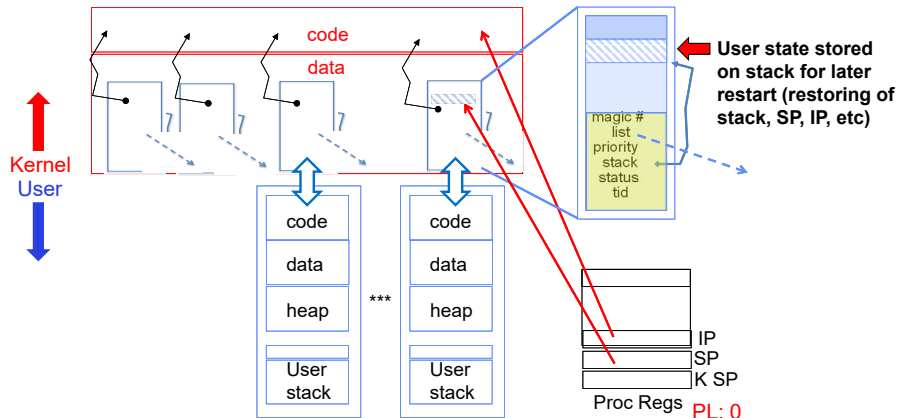
stubs.S

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.76

## Switch to Kernel Stack for Thread



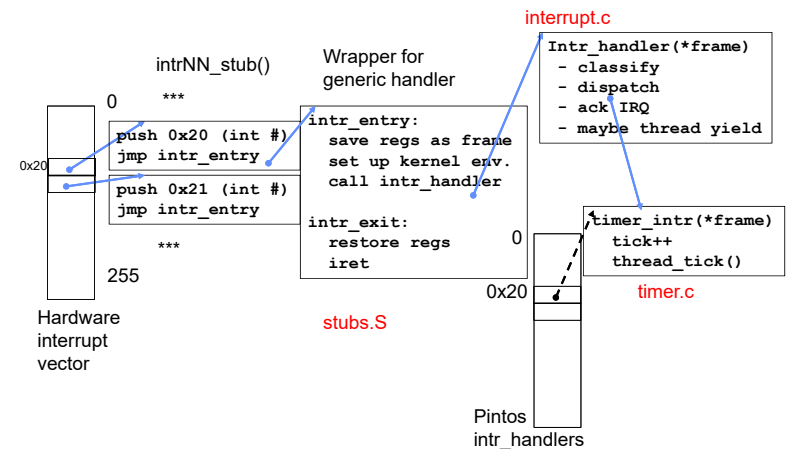
- Information required to restart thread stored on kernel stack
  - Switching becomes simple to different kernel stack and restoring

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.77

## Pintos Interrupt Processing for Timer (0x20)

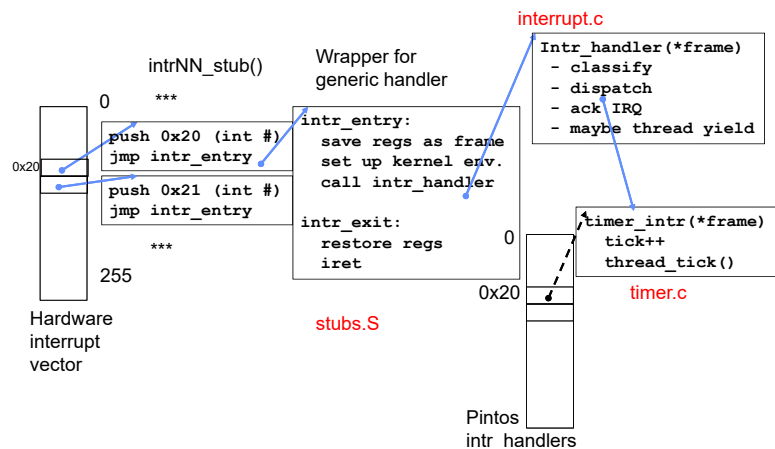


2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.78

## Pintos Interrupt Processing for Timer (0x20)



2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.79

## Timer may trigger thread switch

- thread\_tick
  - Updates thread counters
  - If quanta exhausted, sets yield flag
- thread\_yield
  - On path to rtn from interrupt
  - Sets current thread back to READY
  - Pushes it back on ready\_list
  - Calls schedule to select next thread to run upon iret
- Schedule
  - Selects next thread to run
  - Calls switch\_threads to change regs to point to stack for thread to resume
  - Sets its status to RUNNING
  - If user thread, activates the process
  - Returns back to intr\_handler

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.80

The diagram illustrates the memory layout of a process, showing the relationship between Kernel and User space, and the components of the process memory stack.

**Kernel/User Space:** A vertical double-headed arrow on the left indicates the boundary between **Kernel** (top) and **User** (bottom) space.

**Process Memory Stack:** The stack is divided into two main sections:

- Kernel Space (Top):** Contains **code** and **data** segments. A red box highlights a portion of this stack, with red arrows pointing to the **code** and **data** labels.
- User Space (Bottom):** Contains **code**, **data**, **heap**, and **User stack** segments. A blue box highlights a portion of this stack, with blue arrows pointing to the **code** and **data** labels.

**Process Registers (Proc Regs):** A stack of registers is shown at the bottom right, including **IP** (Instruction Pointer), **SP** (Stack Pointer), and **K SP** (Kernel Stack Pointer). The **SP** register is highlighted in red, and the **K SP** register is highlighted in blue.

**Process ID (PL: 0):** A label at the bottom right indicates the process ID.

- 2/20/24 Kubiawicz CS162 © UCB Spring 2024

2/20/24

The diagram illustrates the interrupt handling flow in Pintos. It starts with the **Hardware interrupt vector** (0 to 255) pointing to `intrNN_stub()`. This function calls `intrNN_stub()`, which pushes the interrupt number and jumps to `intr_entry`. `intr_entry` saves registers, sets up the kernel environment, and calls `intr_handler`. `intr_handler` calls `Intr_handler(*frame)`, which may yield the thread. Then `timer_intr(*frame)` increments the tick and calls `thread_tick()`. Finally, `thread_yield()` is called to schedule the next thread. The flow returns from `intr_exit` to `stubs.S`, which then resumes the thread.

```

graph TD
    IV[Hardware interrupt vector] -- 0 --> intrNN_stub
    subgraph intrNN_stub_code [intrNN_stub()]
        direction TB
        intrNN_stub_code_1[***]
        intrNN_stub_code_2[push 0x20 int #]
        intrNN_stub_code_3[jmp intr_entry]
        intrNN_stub_code_4[push 0x20 int #]
        intrNN_stub_code_5[jmp intr_entry]
        intrNN_stub_code_6[***]
    end
    intrNN_stub_code_3 --> intr_entry
    subgraph intr_entry_exit [intr_entry: intr_exit:]
        direction TB
        intr_entry_1[save regs as frame]
        intr_entry_2[set up kernel env.]
        intr_entry_3[call intr_handler]
        intr_exit_1[restore regs]
        intr_exit_2[iret]
    end
    intr_entry_3 --> intr_handler
    subgraph intr_handler [Intr_handler(*frame)]
        direction TB
        intr_handler_1[- classify]
        intr_handler_2[- dispatch]
        intr_handler_3[- ack IRQ]
        intr_handler_4[- maybe thread yield]
    end
    intr_handler_4 --> timer_intr
    subgraph timer_intr [timer_intr(*frame)]
        direction TB
        timer_intr_1[tick++]
        timer_intr_2[thread_tick()]
    end
    timer_intr_2 --> thread_yield
    subgraph thread_yield [thread_yield()]
        direction TB
        thread_yield_1[- schedule]
    end
    thread_yield_1 --> schedule
    subgraph schedule [schedule()]
        direction TB
        schedule_1[- switch]
    end
    schedule_1 --> stubsS
    subgraph stubsS [stubs.S]
        stubsS_1[Resume Some Thread]
    end
    stubsS_1 --> intrNN_stub_code_1
    style stubsS_1 fill:none,stroke:none

```

Lec 10.82

The diagram illustrates the memory layout and context switch mechanism between User and Kernel space. A vertical double-headed arrow on the left indicates the transition between **Kernel** (top) and **User** (bottom) spaces.

**Kernel Space:**

- Contains **code** and **data** segments.
- A detailed view of the **data** segment shows a **magic #**, **list**, **priority**, **stack**, and **status** field.
- A **Proc Regs** block is shown with fields for **IP**, **SP**, and **K SP**.

**User Space:**

- Contains **code**, **data**, **heap**, and **User stack** segments.
- A detailed view of the **User stack** shows the **IP**, **SP**, and **K SP** fields.

**Context Switch Mechanism:**

- Arrows indicate the flow of data and control between the Kernel and User spaces.
- A red dashed arrow points from the **magic #** field in the Kernel **data** segment to the **IP** field in the User **stack**.
- A blue dashed arrow points from the **list** field in the Kernel **data** segment to the **SP** field in the User **stack**.
- A blue dashed arrow points from the **status** field in the Kernel **data** segment to the **K SP** field in the User **stack**.

- 2/20/24 Kubiatowicz CS162 © UCB Spring 2024

2/20/24

```

2230  /*
2231  * If the new process paused because it was
2232  * swapped out, set the stack level to the last call
2233  * to savu(savu). This means that the return
2234  * value is executed immediately after the call to aretu
2235  * actually returns from the last routine which did
2236  * the savu.
2237  *
2238  * You are not expected to understand this.
2239  */

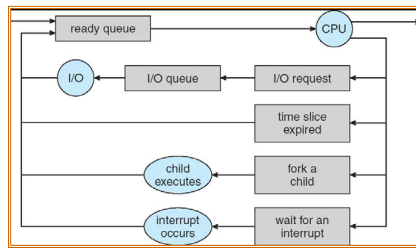
```

*"You are not expected to understand this."*

Included by Ali R. Butt in CS3204 from Virginia Tech

Lec 10.84

## Recall: Scheduling



- Question: How is the OS to decide which of several tasks to take off a queue?
- **Scheduling**: deciding which threads are given access to resources from moment to moment
  - Often, we think in terms of CPU time, but could also think about access to resources like network BW or disk access

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.85

## Scheduling: All About Queues



2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.86

## 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

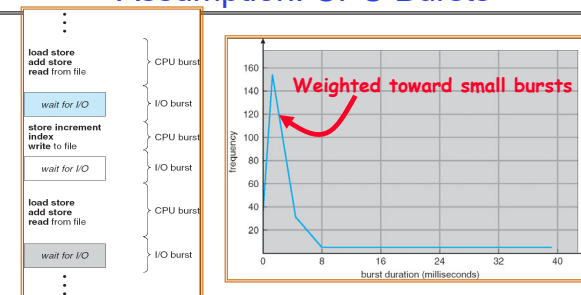


2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.87

## 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

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.88

## 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

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.89

## First-Come, First-Served (FCFS) Scheduling

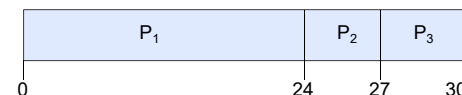
- First-Come, First-Served (FCFS)
  - Also “First In, First Out” (FIFO) or “Run until done”
    - » In early systems, FCFS meant one program scheduled until done (including I/O)
    - » Now, means keep CPU until thread blocks



- 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:

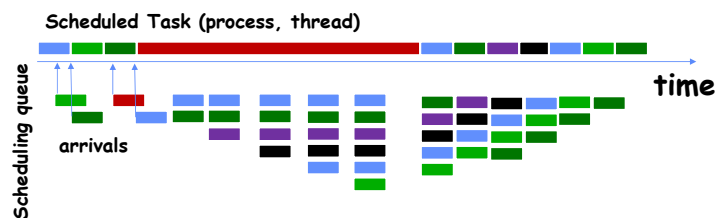


- Waiting time for  $P_1 = 0$ ;  $P_2 = 24$ ;  $P_3 = 27$
- Average waiting time:  $(0 + 24 + 27)/3 = 17$
- Average Completion time:  $(24 + 27 + 30)/3 = 27$
- **Convoy effect:** short process stuck behind long process

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.90

## Convoy effect



- With FCFS non-preemptive scheduling, convoys of small tasks tend to build up when a large one is running.

2/20/24

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.91

## 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 items!
    - Upside: get to read about Space Aliens!

Kubiatowicz CS162 © UCB Spring 2024

Lec 10.92

## Conclusion

---

- **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()**
- Monitors represent the logic of the program
  - Wait if necessary
  - Signal when change something so any waiting threads can proceed
- Readers/Writers Monitor example
  - Shows how monitors allow sophisticated controlled entry to protected code
  - Mesa scheduling allows a more relaxed checking of wait conditions
- Monitors supported natively in a number of languages
- **Scheduling Goals:**
  - Minimize Response Time (e.g. for human interaction)
  - Maximize Throughput (e.g. for large computations)
  - Fairness (e.g. Proper Sharing of Resources)
  - Predictability (e.g. Hard/Soft Realtime)
- **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