

# CS 33

## Multithreaded Programming II

## Problem

```
pthread_create(&thread, 0, start, 0);  
  
...  
  
void *start(void *arg) {  
    long BigArray[128*1024*1024];  
    ...  
    return 0;  
}
```

Here we are creating a thread that has a very large local variable that, of course, is allocated on the thread's stack. How can we be sure that the thread's stack is actually big enough? As it turns out, the default stack size for threads in Linux is two megabytes.

## Thread Attributes

```
pthread_t thread;
pthread_attr_t thr_attr;

pthread_attr_init(&thr_attr);

...

/* establish some attributes */

...

pthread_create(&thread, &thr_attr, startroutine, arg);

...
```

A number of properties of a thread can be specified via the **attributes** argument when the thread is created. Some of these properties have default values set as part of the POSIX specification, others are left up to the implementation. By burying them inside the attributes structure, we make it straightforward to add new types of properties to threads without having to complicate the parameter list of **pthread\_create**. To set up an attributes structure, one must call **pthread\_attr\_init**. As seen in the next slide, one then specifies certain properties, or attributes, of threads. One can then use the attributes structure as an argument to the creation of any number of threads.

Note that the attributes structure only affects the thread when it is created. Modifying an attributes structure has no effect on already-created threads, but only on threads created subsequently with this structure as the attributes argument.

Storage may be allocated as a side effect of calling **pthread\_attr\_init**. To ensure that it is freed, call **pthread\_attr\_destroy** with the attributes structure as argument. Note that if the attributes structure goes out of scope, not all storage associated with it is necessarily released — to release this storage you must call **pthread\_attr\_destroy**.

## Stack Size

```
pthread_t thread;
pthread_attr_t thr_attr;

pthread_attr_init(&thr_attr);
pthread_attr_setstacksize(&thr_attr, 130*1024*1024);

...

pthread_create(&thread, &thr_attr, startroutine, arg);

...
```

Among the attributes that can be specified is a thread's **stack size**. The default attributes structure specifies a stack size that is probably good enough for “most” applications. How big is it? While the default stack size is not mandated by POSIX, in Linux it is two megabytes. To establish a different stack size, use the **pthread\_attr\_setstacksize** function, as shown in the slide.

How large a stack is necessary? The answer, of course, is that it depends. If the stack size is too small, there is the danger that a thread will attempt to overwrite the end of its stack. There is no problem with specifying too large a stack, except that, on a 32-bit machine, one should be careful about using up too much address space (one thousand threads, each with a one-megabyte stack, use a fair portion of the address space).

What happens if a thread uses more stack space than was allotted to it? It would probably clobber memory holding another thread's stack, which could lead to some rather difficult to debug problems. To guard against such happenings, The lowest-address page of a thread's stack (recall that stacks grow downwards) is made inaccessible, meaning that any reference to it will generate a fault. Thus, if the thread references just beyond its allotted stack, there will be a fault which, though not good, makes it clear that this thread has exceeded its stack space.

## Execution Speed

```
// thread 1  
lengthy_func();  
x = 1;
```

```
// thread 2  
x = 0;  
assert(x==0);
```

Assume we have two threads that execute the code in the slide, both starting at about the same time. While it's unlikely that thread 2's assertion will fail, it could happen. In general, we can't make conclusions about what happens based on our assumptions about the speed of thread execution.

There may be lurking bugs in program that are unlikely to happen, but that doesn't mean they won't happen.

## Mutual Exclusion



The mutual-exclusion problem involves making certain that two things don't happen at once. A non-computer example arose in the fighter aircraft of World War I (pictured is a Sopwith Camel). Due to a number of constraints (e.g., machine guns tended to jam frequently and thus had to be close to people who could unjam them), machine guns were mounted directly in front of the pilot. However, blindly shooting a machine gun through the whirling propeller was not a good idea — one was apt to shoot oneself down. At the beginning of the war, pilots politely refrained from attacking fellow pilots. A bit later in the war, however, the Germans developed the tactic of gaining altitude on an opponent, diving at him, turning off the engine, then firing without hitting the now-stationary propeller. Today, this would be called **coarse-grained synchronization**. Later, the Germans developed technology that synchronized the firing of the gun with the whirling of the propeller, so that shots were fired only when the propeller blades would not be in the way. This is perhaps the first example of a mutual-exclusion mechanism providing **fine-grained synchronization**.

# Threads and Mutual Exclusion

## Thread 1:

```
x = x+1;  
/*  
    movl x,%eax  
    incr %eax  
    movl %eax,x  
*/
```

## Thread 2:

```
x = x+1;  
/*  
    movl x,%eax  
    incr %eax  
    movl %eax,x  
*/
```

Here we have two threads that are reading and modifying the same variable: both are adding one to  $x$ . Although the operation is written as a single step in terms of C code, it might take three machine instructions, as shown in the slide. If the initial value of  $x$  is 0 and the two threads execute the code shown in the slide, we might expect that the final value of  $x$  is 2. However, suppose the two threads execute the machine code at roughly the same time: each loads the value of  $x$  into its register, each adds one to the contents of the register, and each stores the result into  $x$ . The final result, of course, is that  $x$  is 1, not 2.

## Quiz 1

Suppose gcc produces the following code. Will it still be the case that x's value might not be incremented by 2?

- a) yes
- b) no

**Thread 1:**

```
x = x+1;  
/*  
  incr x  
*/
```

**Thread 2:**

```
x = x+1;  
/*  
  incr x  
*/
```

In this example, gcc generates an instruction that directly increments the memory location holding x.



## POSIX Threads Mutual Exclusion

```
pthread_mutex_t m =  
    PTHREAD_MUTEX_INITIALIZER;  
    // shared by both threads  
int x; // ditto  
  
pthread_mutex_lock(&m);  
  
x = x+1;  
  
pthread_mutex_unlock(&m);
```

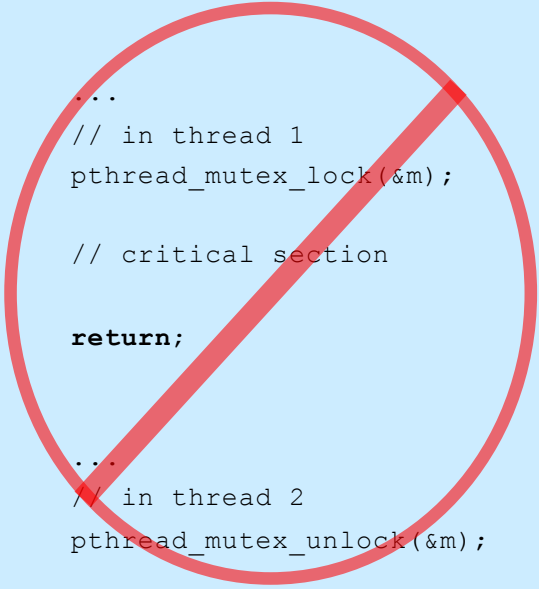
To solve our synchronization problem, we introduce **mutexes** — a synchronization construct providing mutual exclusion. A mutex is used to insure either that only one thread is executing a particular piece of code at once (code locking) or that only one thread is accessing a particular data structure at once (data locking). A mutex belongs either to a particular thread or to no thread (i.e., it is either locked or unlocked). A thread may lock a mutex by calling **pthread\_mutex\_lock**. If no other thread has the mutex locked, then the calling thread obtains the lock on the mutex and returns. Otherwise, it waits until no other thread has the mutex, and finally returns with the mutex locked. There may of course be multiple threads waiting for the mutex to be unlocked. Only one thread can lock the mutex at a time; there is no specified order for who gets the mutex next, though the ordering is assumed to be at least somewhat “fair.”

To unlock a mutex, a thread calls **pthread\_mutex\_unlock**. It is considered incorrect to unlock a mutex that is not held by the caller (i.e., to unlock someone else’s mutex). However, it is somewhat costly to check for this, so most implementations, if they check at all, do so only when certain degrees of debugging are turned on.

Like any other data structure, mutexes must be initialized. This can be done via a call to **pthread\_mutex\_init** or can be done statically by assigning `PTHREAD_MUTEX_INITIALIZER` to a mutex. The initial state of such initialized mutexes is unlocked. Of course, a mutex should be initialized only once! (I.e., make certain that, for each mutex, no more than one thread calls **pthread\_mutex\_init**.)

## Correct Usage

```
pthread_mutex_lock(&m);  
  
// critical section  
  
pthread_mutex_unlock(&m);  
  
...  
// in thread 1  
pthread_mutex_lock(&m);  
  
// critical section  
  
return;  
  
...  
// in thread 2  
pthread_mutex_unlock(&m);
```



An important restriction on the use of mutexes is that the thread that locked a mutex should be the thread that unlocks it. For a number of reasons, not the least of which is readability and correctness, it is not good for a mutex to be locked by one thread and then unlocked by another.

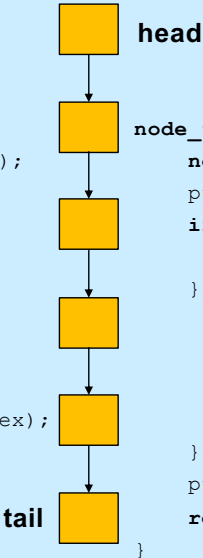
## A Queue

```

void enqueue(node_t *item) {
    pthread_mutex_lock(&mutex);
    item->next = NULL;
    if (tail == NULL) {
        head = item;
        tail = item;
    } else {
        tail->next = item;
    }
    pthread_mutex_unlock(&mutex);
}

node_t *dequeue() {
    node_t *ret;
    pthread_mutex_lock(&mutex);
    if (head == NULL) {
        ret = NULL;
    } else {
        ret = head;
        head = head->next;
        if (head == NULL)
            tail = NULL;
    }
    pthread_mutex_unlock(&mutex);
    return ret;
}

```



Here we have **enqueue** and **dequeue** functions that can be called by multiple threads to add and remove items from a queue. We employ a single mutex to make certain that at most one thread is performing a queue operation at a time. This could result in a bottleneck if there are lots of threads calling both of the functions. Thus, we might seek a solution that employs separate mutexes for callers to enqueue and for callers to dequeue.

Since enqueue modifies one end of the queue and dequeue modifies the other, one might think that we could have a two-mutex solution, with one mutex protecting one end of the queue and another mutex protecting the other. But this becomes difficult in certain edge conditions. For example, suppose the queue contains just one item. A thread is calling dequeue to remove that item, but at the same time another thread is calling enqueue to add another item. With one mutex protecting the tail of the queue and another protecting the head, we could have a situation in which the next field of the node being enqueued refers to the node being dequeued. If we have a single mutex ensuring mutually exclusive access to the entire queue, this is easy to prevent (the situation can't happen in the code of the slide). But if callers to enqueue use a different mutex than callers to dequeue, we can't easily prevent the problem (if at all).

## Removing a Freelist Block

```
void pull_from_freelist(fblock_t *fbp) {  
    ...  
    fbp->blink->flink = fbp->flink;  
    fbp->flink->blink = fbp->blink;  
    ...  
}
```

We review some code that should be familiar to you from the malloc assignment. We assume, for the sake of the example, that there are at least three blocks in the freelist. (Why it's three and not two comes up later.)

## Parallelizing It

- **Coarse grained**
  - one mutex for the heap
  - threads lock the mutex before doing any operation
  - unlock it afterwards
  - only one thread at a time
- **Fine grained**
  - one mutex for each block
  - threads lock mutexes of only the blocks they are using
  - multiple threads at a time

However, we'd now like to modify all the malloc code so that multiple threads can use it at the same time.

One way of doing this, which is pretty straightforward, is to employ a single mutex, which threads must lock before doing any heap operation, and then unlock when they're finished. However, this approach allows only one thread to be manipulating the heap at a time. This approach to using mutexes is called **coarse-grained**.

To get more advantage (in terms of speed) from using multiple threads, we'd like it to be possible for multiple threads to be manipulating the heap at once. To accomplish this, each block has a mutex inside of it. To do anything to this block (such as pull it from the free list or add it to the free list), the thread doing so must lock the block's mutex, then, of course, unlock it when done. This approach to using mutexes is called **fine-grained**.

## Removing a Freelist Block: Coarse Grained

```
void pull_from_freelist(fblock_t *fbp) {  
    pthread_mutex_lock(&heap_mutex);  
    ...  
    fbp->blink->flink = fbp->flink;  
    fbp->flink->blink = fbp->blink;  
    ...  
    pthread_mutex_unlock(&heap_mutex);  
}
```

Here's the modified code for `pull_from_freelist`, using coarse-grained synchronization.

## Removing a Freelist Block: Fine Grained (1)

```
void pull_from_freelist(fblock_t *fbp) {  
    pthread_mutex_lock(&fpp->mutex);  
    ...  
    fbp->blink->flink = fbp->flink;  
    fbp->flink->blink = fbp->blink;  
    ...  
    pthread_mutex_unlock(&fpp->mutex);  
}
```

Let's try a similar approach for the fine-grained approach.

However, we can see an immediate problem: the thread executing this code will not only modify `*fbp`, but also its predecessor and successor blocks in the freelist.


## Removing a Freelist Block: Fine Grained (2)

```
void pull_from_freelist(fbblock_t *fbp) {  
    pthread_mutex_lock(&fpp->mutex);  
    ...  
    pthread_mutex_lock(&fpp->blink->mutex);  
    fbp->blink->flink = fbp->flink;  
    pthread_mutex_lock(&fpp->flink->mutex);  
    fbp->flink->blink = fbp->blink;  
    ...  
    pthread_mutex_unlock(&fpp->blink->mutex);  
    pthread_mutex_unlock(&fpp->flink->mutex);  
    pthread_mutex_unlock(&fpp->mutex);  
}
```

So, we add code to lock and unlock the mutexes of the adjacent blocks.




## Multiple Pulls



Pictorially, our thread first locks the mutex on `*fbp` (the middle block), then its predecessor and then its successor.

## Multiple Pulls



But suppose other threads are calling `pull_from_freelist` at the same time. Let's say that thread 1 is pulling the middle block, thread 2 is pulling the upper block, and thread 3 is pulling the lower block.

All three threads have locked the mutex on the block they're pulling. Thread one tries to lock the mutex on the upper block (once it successfully locks it, then it will try to lock the mutex on the lower block). However, thread 2 has already locked the mutex on the upper block, and won't unlock it until after it locks the mutex on the middle block. But thread 1 won't unlock it until it locks the mutex on the upper block (not to mention the lower block).


So, we're stuck. threads 1 and 2 (as well as 3) won't ever be able to lock all the mutexes they need, and, because of this, can't make any further progress. This phenomenon is called **deadlock**.

## Taking Multiple Locks

```
func1( ) {  
    pthread_mutex_lock(&m1);  
    /* use object 1 */  
    pthread_mutex_lock(&m2);  
    /* use objects 1 and 2 */  
    pthread_mutex_unlock(&m2);  
    pthread_mutex_unlock(&m1);  
}  
  
func2( ) {  
    pthread_mutex_lock(&m2);  
    /* use object 2 */  
    pthread_mutex_lock(&m1);  
    /* use objects 1 and 2 */  
    pthread_mutex_unlock(&m1);  
    pthread_mutex_unlock(&m2);  
}
```

In this example our threads are using two mutexes to control access to two different objects. Thread 1, executing **func1**, first takes mutex 1, then, while still holding mutex 1, obtains mutex 2. Thread 2, executing **func2**, first takes mutex 2, then, while still holding mutex 2, obtains mutex 1. However, things do not always work out as planned. If thread 1 obtains mutex 1 and, at about the same time, thread 2 obtains mutex 2, then if thread 1 attempts to take mutex 2 and thread 2 attempts to take mutex 1, we have a **deadlock**.

## Preventing Deadlock



Deadlock results when there are circularities in dependencies. In the slide, mutex 1 is held by thread a, which is waiting to take mutex 2. However, thread b is holding mutex 2, waiting to take mutex 1. If we can make certain that such circularities never happen, there can't possibly be deadlock.

## Taking Multiple Locks, Safely

```
proc1( ) {  
    pthread_mutex_lock(&m1);  
    /* use object 1 */  
    pthread_mutex_lock(&m2);  
    /* use objects 1 and 2 */  
    pthread_mutex_unlock(&m2);  
    pthread_mutex_unlock(&m1);  
}  
  
proc2( ) {  
    pthread_mutex_lock(&m1);  
    /* use object 1 */  
    pthread_mutex_lock(&m2);  
    /* use objects 1 and 2 */  
    pthread_mutex_unlock(&m2);  
    pthread_mutex_unlock(&m1);  
}
```

If all threads take locks in the same order, deadlock cannot happen.

How can we modify our `pull_from_freelist` code to use this approach (of all threads taking locks in the same order)?

## Dining Philosophers Problem



The problem we've been looking at is a special case of what's known as the "dining philosophers problem", posed by Edsger Dijkstra in EWD310, first published as Hierarchical Ordering of Sequential Processes in Operating Systems Techniques, C.A.R. Hoare and R.H. Perrot, Eds., Academic Press, New York, 1972. The idea is that we have five philosophers sitting around a table. At the center of the table is a plate of spaghetti. Between each pair of philosophers is a single chopstick (Dijkstra's original formulation used forks, but chopsticks make more sense). The algorithm of a philosopher is:

```
while (1) {  
    think();  
    when available  
        grab chopstick from one side();  
    when available  
        grab chopstick from the other side();  
    eat some spaghetti();  
    put chopsticks down();  
}
```

How long each operation takes varies. Which chopstick is grabbed first is not specified, but if each philosopher grabs their right chopstick first, they may starve to death. There are many subtle issues involved in its solution. (It has many, none of which are as interesting as the problem itself.)

Philosophers clockwise from top: Laozi, Swami Vivekananda, Aristotle, Mary Wollstonecraft, Zara Yacob.

## Practical Issues with Mutexes

- **Used a lot in multithreaded programs**
  - **speed is really important**
    - » shouldn't slow things down much in the success case
  - **checking for errors slows things down (a lot)**
    - » thus errors aren't checked by default

## Set Up

```
int pthread_mutex_init(pthread_mutex_t *mutexp,  
    pthread_mutexattr_t *attrp)  
  
int pthread_mutex_destroy(pthread_mutex_t *mutexp)  
  
int pthread_mutexattr_init(pthread_mutexattr_t *attrp)  
  
int pthread_mutexattr_destroy(pthread_mutexattr_t *attrp)
```

The functions **pthread\_mutex\_init** and **pthread\_mutex\_destroy** are supplied to initialize and to destroy a mutex. (They do not allocate or free the storage for the mutex data structure, but in some implementations they might allocate and free storage referred to by the mutex data structure.) As with threads, an attribute structure encapsulates the various parameters that might apply to the mutex. The functions **pthread\_mutexattr\_init** and **pthread\_mutexattr\_destroy** control the initialization and destruction of these attribute structures, as we see a few slides from now. For most purposes, the default attributes are fine and a NULL **attrp** can be provided to the **pthread\_mutex\_init** routine.

Note that, as we've already seen, a mutex that's allocated statically may be initialized with `PTHREAD_MUTEX_INITIALIZER`.



## Stupid (i.e., Common) Mistakes ...

```
pthread_mutex_lock(&m1);  
pthread_mutex_lock(&m1);  
    // really meant to lock m2 ...  
  
pthread_mutex_lock(&m1);  
    ...  
pthread_mutex_unlock(&m2);  
    // really meant to unlock m1 ...
```

In the example at the top of the slide, we have mistyped the name of the mutex in the second call to **pthread\_mutex\_lock**. The result will be that when **pthread\_mutex\_lock** is called for the second time, there will be immediate deadlock, since the caller is attempting to lock a mutex that is already locked, but the only thread who can unlock that mutex is the caller.

In the example at the bottom of the slide, we have again mistyped the name of a mutex, but this time for a **pthread\_mutex\_unlock** call. If m2 is not currently locked by some thread, unlocking will have unpredictable results, possibly fatal. If m2 is locked by some thread, again there will be unpredictable results, since a mutex that was thought to be locked (and protecting some data structure) is now unlocked. When the thread who locked it attempts to unlock it, the result will be even further unpredictability.

## Runtime Error Checking

```
pthread_mutexattr_t err_chk_attr;
pthread_mutexattr_init(&err_chk_attr);
pthread_mutexattr_settype(&err_chk_attr,
    PTHREAD_MUTEX_ERRORCHECK);

pthread_mutex_t mut1;
pthread_mutex_init(&mut1, &err_chk_attr);

pthread_mutex_lock(&mut1);


if (pthread_mutex_lock(&mut1) == EDEADLK)
    fprintf(stderr, "error caught at runtime\n");

if (pthread_mutex_unlock(&mut2) == EPERM)
    fprintf(stderr, "another error: you didn't lock it!\n");
```

Checking for some sorts of mutex-related errors is relatively easy to do at runtime (though checking for all possible forms of deadlock is prohibitively expensive). However, since mutexes are used so frequently, even a little bit of extra overhead for runtime error checking is often thought to be too much. Thus, if done at all, runtime error checking is an optional feature. One “turns on” the feature for a particular mutex by initializing it to be of type “ERRORCHECK,” as shown in the slide. For mutexes initialized in this way, **pthread\_mutex\_lock** checks to make certain that it is not attempting to lock a mutex that is already locked by the calling thread; **pthread\_mutex\_unlock** checks to make certain that the mutex being unlocked is currently locked by the calling thread.

Note that mutexes with the error-check attribute are more expensive than normal mutexes, since they must keep track of which thread, if any, has the mutex locked. (For normal mutexes, just a single bit must be maintained for the state of the mutex, which is either locked or unlocked.)

# Producer-Consumer Problem



In the **producer-consumer problem** we have two classes of threads, producers and consumers, and a buffer containing a fixed number of slots. A producer thread attempts to put something into the next empty buffer slot, a consumer thread attempts to take something out of the next occupied buffer slot. The synchronization conditions are that producers cannot proceed unless there are empty slots and consumers cannot proceed unless there are occupied slots.

This is a classic, but frequently occurring synchronization problem. For example, the heart of the implementation of UNIX pipes is an instance of this problem.


## Guarded Commands

```
when (guard) [  
    /*  
        once the guard is true, execute this  
        code atomically  
    */  
  
    . . .  
  
]
```

Illustrated in the slide is a simple pseudocode construct, the **guarded command**, that we use to describe how various synchronization operations work. The idea is that the code within the square brackets is executed only when the guard (which could be some arbitrary boolean expression) evaluates to true. Furthermore, this code within the square brackets is executed atomically, i.e., the effect is that nothing else happens in the program while the code is executed. Note that the code is not necessarily executed as soon as the guard evaluates to true: we are assured only that when execution of the code begins, the guard is true.

Keep in mind that this is strictly pseudocode: it's not part of POSIX threads and is not necessarily even implementable (at least not for the general case).

# Semaphores



- **P(S) operation:**

```
when (S > 0) [  
    S = S - 1;  
]
```

- **V(S) operation:**

```
[S = S + 1;]
```

Another synchronization construct is the semaphore, designed by Edsger Dijkstra in the 1960s. A semaphore behaves as if it were a nonnegative integer, but it can be operated on only by the semaphore operations. Dijkstra defined two of these: P (for **prolagen**, a made-up word derived from **proberen te verlagen**, which means “try to decrease” in Dutch) and V (for **verhogen**, “increase” in Dutch). Their semantics are shown in the slide.

We think of operations on semaphores as being a special case of guarded commands — a special case that occurs frequently enough to warrant a highly optimized implementation.

## Quiz 2

```
semaphore S = 1;  
int count = 0;
```

```
void func( ) {  
    P(S);  
    count++;  
    ...  
    count--;  
    V(S);  
}
```

The function func is called concurrently by n threads. What's the maximum value that count will take on?

- a) 1
- b) 2
- c) n
- d) indeterminate

- **P(S) operation:**

```
when (S > 0) [  
    S = S - 1;  
]
```

- **V(S) operation:**

```
[S = S + 1;]
```

## Producer/Consumer with Semaphores

```
Semaphore empty = BSIZE;
Semaphore occupied = 0;
int nextin = 0;
int nextout = 0;

void Produce(char item) {
    P(empty);
    buf[nextin] = item;
    if (++nextin >= BSIZE)
        nextin = 0;
    V(occupied);
}

char Consume( ) {
    char item;
    P(occupied);
    item = buf[nextout];
    if (++nextout >= BSIZE)
        nextout = 0;
    V(empty);
    return item;
}
```

Here's a solution for the producer/consumer problem using semaphores — note that it works only with a single producer and a single consumer, and only one item at a time is produced or consumed, though it can be generalized to work with multiple producers and consumers.

## POSIX Semaphores

```
#include <semaphore.h>

int sem_init(sem_t *semaphore, int pshared, int init);
int sem_destroy(sem_t *semaphore);
int sem_wait(sem_t *semaphore);
    /* P operation */
int sem_trywait(sem_t *semaphore);
    /* conditional P operation */
int sem_post(sem_t *semaphore);
    /* V operation */
```

Here is the POSIX interface for operations on semaphores. (These operation names are not typos — the “pthread\_” prefix really is not used here, since the semaphore operations come from a different POSIX specification — 1003.1b. Note also the need for the header file, **semaphore.h**) When creating a semaphore (**sem\_init**), rather than supplying an attributes structure, one supplies a single integer argument, **pshared**, which indicates whether the semaphore is to be used only by threads of one process (**pshared = 0**) or by multiple processes (**pshared = 1**). The third argument to **sem\_init** is the semaphore’s initial value.

All the semaphore operations return zero if successful; otherwise, they return an error code. The function **sem\_trywait** is similar to **sem\_wait** (and to the P operation) except that if the semaphore’s value cannot be decremented immediately, then rather than wait, it returns -1 and sets `errno` to `EAGAIN`.



## Producer-Consumer with POSIX Semaphores

```
sem_init(&empty, 0, BSIZE);
sem_init(&occupied, 0, 0);
int nextin = 0;
int nextout = 0;

void produce(char item) {
    sem_wait(&empty);
    buf[nextin] = item;
    if (++nextin >= BSIZE)
        nextin = 0;
    sem_post(&occupied);
}

char consume() {
    char item;
    sem_wait(&occupied);
    item = buf[nextout];
    if (++nextout >= BSIZE)
        nextout = 0;
    sem_post(&empty);
    return item;
}
```

Here is the producer-consumer solution implemented with POSIX semaphores.

## Quiz 3

**Does the POSIX version of the producer-consumer solution work with multiple producers and consumers?**

- a) Yes**
- b) No, but it can be made to work by using mutexes to make sure that only one thread is executing the producer code at a time and only one thread is executing the consumer code at a time**
- c) It can't easily be made to work**

# Start/Stop



- Start/Stop interface

```
void wait_for_start(state_t *s);
```

```
void start(state_t *s);
```

```
void stop(state_t *s);
```

We'd like to design a “start-stop” interface. A thread calling **wait\_for\_start** waits for the **start** button to be pressed. Once it's been pressed, those waiting will be released and subsequent threads calling **wait\_for\_start** will return immediately. However, once the **stop** button is pressed, then all threads calling **wait\_for\_start** will wait until the **start** button is pressed again.

## Start/Stop



- Start/Stop interface

```
void wait_for_start(state_t *s){
    if (s->state == stopped)
        sleep();
}

void start(state_t *s) {
    state = started;
    wakeup_all();
}

void stop(state_t *s) {
    state = stopped;
}
```

Here's a possible implementation. Callers of **sleep** don't return from sleep until **wakeup\_all** has been called.

However, calls to **wakeup\_all** merely wakeup all who are currently in **sleep**. They have no effect on subsequent calls to sleep. Thus, there could be a problem in the above code if a thread calls start while another thread has just checked the state in **wait\_for\_start**, but hasn't yet called **sleep**.

# Start/Stop



- Start/Stop interface

```
void wait_for_start(state_t *s) {
    pthread_mutex_lock(&s->mutex);
    if (s->state == stopped) {
        pthread_mutex_unlock(&s->mutex);
        sleep();
    } else pthread_mutex_unlock(&s->mutex);
}

void start(state_t *s) {
    pthread_mutex_lock(&s->mutex);
    state = started;
    wakeup_all();
    pthread_mutex_unlock(&s->mutex);
}
```

Here's one attempt to fix the problem of the previous slide using mutexes. It clearly doesn't help – the thread calling start might get the mutex and call wakeup\_all just before the other thread calls sleep.

# Start/Stop



- Start/Stop interface

```
void wait_for_start(state_t *s) {  
    pthread_mutex_lock(&s->mutex);  
    if (s->state == stopped) {  
        sleep();  
    }  
    pthread_mutex_unlock(&s->mutex);  
}  
  
void start(state_t *s) {  
    pthread_mutex_lock(&s->mutex);  
    state = started;  
    wakeup_all();  
    pthread_mutex_unlock(&s->mutex);  
}
```

This code is perhaps worse, the thread waits in sleep with the mutex locked, preventing any thread from calling wakeup\_all.

# Start/Stop



- Start/Stop interface

```
void wait_for_start(state_t *s){
    pthread_mutex_lock(&s->mutex);
    while(s->state == stopped)
        pthread_cond_wait(&s->queue, &s->mutex);
    pthread_mutex_unlock(&s->mutex);
}

void start(state_t *s) {
    pthread_mutex_lock(&s->mutex);
    s->state = started;
    pthread_cond_broadcast(&s->queue);
    pthread_mutex_unlock(&s->mutex);
}
```

This code actually works; it uses a POSIX threads construct known as the **condition variable**. The thread in `wait_for_start` first locks the mutex, then checks the state. If it's stopped, it calls **pthread\_cond\_wait**, which, all at once, put the calling thread to sleep, enqueues it on the queue (known as a condition variable, and unlocks the mutex.

A thread calling `start` can't proceed until it has locked the mutex, thus ensuring that no thread is in the midst of checking the state and then calling **pthread\_cond\_wait** in `wait_for_start`. Once the thread calling `start` has the mutex, it sets state to started and calls **pthread\_broadcast**, waking up all threads who are waiting on the queue (the condition variable). It then unlocks the mutex.

The thread that was waiting within **pthread\_cond\_wait** is woken up, but it doesn't return from the call to **pthread\_cond\_wait** until it locks the mutex. Thus, it enters **pthread\_cond\_wait** with the mutex locked and exits it with the mutex locked. While its inside **pthread\_cond\_wait**, it does not have the lock on the mutex (though some other thread might).

Thus, this code is a correct implementation of the start/stop interface.

## Condition Variables

```
when (guard) [  
    statement 1;  
    ...  
    statement n;  
]  
  
// code modifying the guard:  
...  
  
pthread_mutex_lock(&mutex);  
while(!guard)  
    pthread_cond_wait(  
        &cond_var, &mutex);  
statement 1;  
...  
statement n;  
pthread_mutex_unlock(&mutex);  
  
pthread_mutex_lock(&mutex);  
// code modifying the guard:  
...  
pthread_cond_broadcast(  
    &cond_var);  
pthread_mutex_unlock(&mutex);
```

**Condition variables** are another means for synchronization in POSIX; they represent queues of threads waiting to be woken by other threads and can be used to implement guarded commands, as shown in the slide. Though they are rather complicated at first glance, they are even more complicated when you really get into them.

A thread puts itself to sleep and joins the queue of threads associated with a condition variable by calling **pthread\_cond\_wait**. When it places this call, it must have some mutex locked, and it passes the mutex as the second argument. As part of the call, the mutex is unlocked and the thread is put to sleep, **all in a single atomic step**: i.e., nothing can happen that might affect the thread between the moments when the mutex is unlocked and when the thread goes to sleep. Threads queued on a condition variable are released in first-in-first-out order. They are released in response to calls to **pthread\_cond\_signal** (which releases the first thread in line) and **pthread\_cond\_broadcast** (which releases all threads). However, before a released thread may return from **pthread\_cond\_wait**, it first relocks the mutex. Thus, only one thread at a time actually returns from **pthread\_cond\_wait**. If a call to either function is made when no threads are queued on the condition variable, nothing happens — the fact that a call had been made is not remembered.

So far, though complicated, the description is rational. Now for the weird part: **a thread may be released from the condition-variable queue at any moment**, perhaps spontaneously, perhaps due to sun spots. Thus, it's extremely important that, after **pthread\_cond\_wait** returns, that the caller check to make sure that it really should have returned. The reason for this weirdness is that it allows a fair amount of latitude in implementations. However, the Linux implementation behaves rationally, i.e., as in the



first two paragraphs. (But don't depend on this behavior — it could change tomorrow!)

## Set Up

```
int pthread_cond_init(pthread_cond_t *cvp,  
    pthread_condattr_t *attrp)  
  
int pthread_cond_destroy(pthread_cond_t *cvp)  
  
int pthread_condattr_init(pthread_condattr_t *attrp)  
  
int pthread_condattr_destroy(pthread_condattr_t *attrp)
```

Setting up condition variables is done in a similar fashion as mutexes: The functions **pthread\_cond\_init** and **pthread\_cond\_destroy** are supplied to initialize and to destroy a condition variable. They may also be statically initialized by setting them to `PTHREAD_COND_INITIALIZER` in their declarations. As with mutexes and threads, default attributes may be specified by supplying a zero. The functions **pthread\_condattr\_init** and **pthread\_condattr\_destroy** control the initialization and destruction of their attribute structures.

## PC with Condition Variables (1)

```
typedef struct buffer {  
    pthread_mutex_t m;  
    pthread_cond_t  more_space;  
    pthread_cond_t  more_items;  
    int             next_in;  
    int             next_out;  
    int             empty;  
    char            buf[BSIZE];  
} buffer_t;
```

Here we begin a producer-consumer solution using condition variables and mutexes; this solution, unlike the previous, allows multiple producers and consumers. We define a struct **buffer** to represent a buffer, associated synchronization variables, and other associated variables. In our example, producers wait for empty slots to become available, and consumers wait for occupied slots to become available. Waiting producers are queued on the condition variable **more\_space** and waiting consumers are queued on the condition variable **more\_items**.

## PC with Condition Variables (2)

```
void produce(buffer_t *b,
             char item) {
    pthread_mutex_lock(&b->m);
    while (!(b->empty > 0))
        pthread_cond_wait(
            &b->more_space, &b->m);
    b->buf[b->nextin] = item;
    if (++(b->nextin) == BSIZE)
        b->nextin = 0;
    b->empty--;
    pthread_cond_signal(
        &b->more_items);
    pthread_mutex_unlock(&b->m);
}

char consume(buffer_t *b) {
    char item;
    pthread_mutex_lock(&b->m);
    while (!(b->empty < BSIZE))
        pthread_cond_wait(
            &b->more_items, &b->m);
    item = b->buf[b->nextout];
    if (++(b->nextout) == BSIZE)
        b->nextout = 0;
    b->empty++;
    pthread_cond_signal(
        &b->more_space);
    pthread_mutex_unlock(&b->m);
    return item;
}
```

Here we have the remaining code of our solution. A producer, if there is at least one empty slot, fills the one at location **nextin**, increments **nextin** (taking wraparound into account), calls **pthread\_cond\_signal** to notify any waiting consumers that there is now an occupied slot in the buffer, and releases the mutex. If there are no empty slots in the buffer, the producer calls **pthread\_cond\_wait** to wait for one.

As discussed previously, this call has a fairly complicated effect: it releases the mutex given as the second argument and puts its caller to sleep, after queuing it on the condition variable given as the first argument. At some point in the future, a consumer should call **pthread\_cond\_signal**, with **more\_space** as the argument.

Note that we've used **pthread\_cond\_signal** rather than **pthread\_cond\_broadcast**. We can do this here since, if, for example, **n** threads are waiting within the call to **pthread\_cond\_wait** in the producer, then there must be **n** calls to **consume** to release them all. If we'd used **pthread\_cond\_broadcast** instead, the solution would still work, but would probably be less efficient, since in many cases waiting threads would return from **pthread\_cond\_wait**, discover that the guard is still false, and have to call **pthread\_cond\_wait** again.

If our producer is the first in the queue associated with **more\_space**, it is released from the queue, but it does not yet return from **pthread\_cond\_wait**. Instead, it continues execution inside that routine, where it effectively makes a call to **pthread\_mutex\_lock** to reacquire the mutex it had when it entered **pthread\_cond\_wait** in the first place. Once it obtains the mutex, it then returns from **pthread\_cond\_wait**. Note that when the thread attempts to reacquire the mutex, other threads might be waiting for the mutex at the entrance of the producer code. One of these other threads might obtain the mutex first — thus there is no guarantee that callers of produce are served in FIFO order.

The order in which threads are released from a condition variable's queue is first-in-first-

out within priority levels. Thus, waiting high-priority threads are released before waiting low-priority threads; threads of the same priority are released in the order in which they called **pthread\_cond\_wait**.