# Lecture 05: Advanced Concurrency Control

# CS343 – Operating Systems Branden Ghena – Spring 2022

Some slides borrowed from: Stephen Tarzia (Northwestern), and Shivaram Venkataraman (Wisconsin)

Northwestern

# Today's Goals

- Understand how we can apply locks to gain correctness and maintain performance
  - Counter
  - Data Structures

- Signaling between threads to enforce ordering
  - Condition Variables
  - Semaphores

#### Review: Locks/Mutexes

Simple mutual exclusion primitive

Init(), Acquire(), Release()

- Implementations trade complexity, fairness, and performance
  - Spinlocks
  - Ticket locks
  - Yielding locks
  - Queueing locks

#### Outline

Applying Locks

Concurrent Data Structures

• Ordering with Condition Variables

• Semaphores

#### Review: Need to enforce mutual exclusion on critical sections

```
#include <stdio.h>
#include <pthread.h>
```

}

```
static volatile int counter = 0;
static const int LOOPS = 1e9;
```

```
void* mythread(void* arg) {
  printf("%s: begin\n", (char*)arg);
  for (int i=0; i<LOOPS; i++) {
    counter++;
}</pre>
```

```
printf("%s: done\n", (char*)arg);
return NULL;
```

```
int main(int argc, char* argv[]) {
   pthread_t p1, p2;
   printf("main: begin (counter = %d)\n", counter);
   pthread_create(&p1, NULL, mythread, "A");
   pthread_create(&p2, NULL, mythread, "B");
```

```
// wait for threads to finish
pthread_join(p1, NULL);
pthread_join(p2, NULL);
printf("main: done with both (counter = %d, goal
was %d)\n", counter, 2*LOOPS);
return 0;
}
```

#### Naively locked counter example

```
#include <stdio.h>
#include <pthread.h>
```

```
static volatile int counter = 0;
static const int LOOPS = 1e9;
static pthread_mutex_t lock;
```

```
void* mythread(void* arg) {
  printf("%s: begin\n", (char*)arg);
  for (int i=0; i<LOOPS; i++) {
    pthread_mutex_lock(&lock);
    counter++;
    pthread_mutex_unlock(&lock);
  }
  printf("%s: done\n", (char*)arg);</pre>
```

```
int main(int argc, char* argv[]) {
   pthread_t p1, p2;
   pthread_mutex_init(&lock, 0);
   printf("main: begin (counter = %d)\n", counter);
   pthread_create(&p1, NULL, mythread, "A");
   pthread_create(&p2, NULL, mythread, "B");
```

```
// wait for threads to finish
pthread_join(p1, NULL);
pthread_join(p2, NULL);
printf("main: done with both (counter = %d, goal
was %d)\n", counter, 2*LOOPS);
return 0;
}
```

```
}
```

return NULL;

### Problem: locking overhead decreases performance

When iteratingSingle-threaded counter: 3.850 secondsone billion times:Multithreaded no-lock counter: 4.700 seconds (Broken!)Naïve-locked counter:80.000 seconds (Correct...)

- Formerly loop contained 3 instructions (mov, add, mov)
- Now it has
  - Two function calls
  - Multiple instructions inside of those
  - Possibly even interaction with the OS...
  - 3 instructions -> 60 instructions

### Simple mutual exclusion: one big lock

- Simple solution "one big lock"
  - Find all the function calls that interact with shared memory
  - Lock at the start of each function call and unlock at the end
- Essentially, no concurrent access
  - Correct but poor performance
  - If you've forgotten all of this years from now, "one big lock" will still work

#### Counter example with big lock technique

}

```
#include <stdio.h>
#include <pthread.h>
```

```
static volatile int counter = 0;
static const int LOOPS = 1e9;
static pthread_mutex_t lock;
```

```
void* mythread(void* arg) {
   pthread_mutex_lock(&lock);
   printf("%s: begin\n", (char*)arg);
   for (int i=0; i<LOOPS; i++) {
      counter++;
   }
   printf("%s: done\n", (char*)arg);
   pthread_mutex_unlock(&lock);
   return NULL;
}</pre>
```

```
int main(int argc, char* argv[]) {
   pthread_t p1, p2;
   pthread_mutex_init(&lock, 0);
   printf("main: begin (counter = %d)\n", counter);
   pthread_create(&p1, NULL, mythread, "A");
   pthread_create(&p2, NULL, mythread, "B");
```

```
// wait for threads to finish
pthread_join(p1, NULL);
pthread_join(p2, NULL);
printf("main: done with both (counter = %d, goal
was %d)\n", counter, 2*LOOPS);
return 0;
```

#### Problem: locking decreases performance

Single-threaded counter: 3.850 seconds Multithreaded no-lock counter: 4.700 seconds (Broken!) Naïve-locked counter: 80.000 seconds

Big lock counter: 3.895 seconds

- Big lock technique basically returned us to single-threaded execution time (and single-threaded implementation)
- Why is the no-lock multithreaded version so slow?
  - Not 100% certain
  - Likely something to do with hardware memory/cache consistency

#### Reducing lock overhead

- We want to enable parallelism, but deal with less lock overhead
  - Need to increase the amount of work done when not locked
  - Goal: lots of parallel work per lock/unlock event
- "Sloppy" updates to global state
  - Keep local state that is operated on
  - Occasionally synchronize global state with current local state
- Counter example
  - Keep a local counter for each thread (not shared memory)
  - Add local counter to global counter periodically

#### Sloppy counter example

```
#include <stdio.h>
#include <pthread.h>
```

```
static volatile int counter = 0;
static const int LOOPS = 1e9;
static pthread_mutex_t lock;
```

```
void* mythread(void* arg) {
    int sloppy_count = 0;
    printf("%s: begin\n", (char*)arg);
    for (int i=0; i<LOOPS; i++) {
        sloppy_count++;
        if (i%1000 == 0) {
            pthread_mutex_lock(&lock);
            counter += sloppy_count;
            pthread_mutex_unlock(&lock);
            sloppy_count = 0;
        }
</pre>
```

```
int main(int argc, char* argv[]) {
   pthread_t p1, p2;
   pthread_mutex_init(&lock, 0);
   printf("main: begin (counter = %d)\n", counter);
   pthread_create(&p1, NULL, mythread, "A");
   pthread_create(&p2, NULL, mythread, "B");
```

```
// wait for threads to finish
pthread_join(p1, NULL);
pthread_join(p2, NULL);
printf("main: done with both (counter = %d, goal
was %d)\n", counter, 2*LOOPS);
return 0;
}
```

Offscreen Tail condition: don't forget to update "counter" again when the for loop is complete!

#### Problem: locking decreases performance

Single-threaded counter: 3.850 seconds Multi-threaded no-lock counter: 4.700 seconds (Broken!) Naïve-locked counter: 80.000 seconds

Big lock counter: 3.895 seconds

Sloppy lock (synchronize every 100): 2.150 seconds Sloppy lock (synchronize every 10000): 1.472 seconds Sloppy lock (synchronize every 1000000): 1.478 seconds Sloppy lock (synchronize every 100000000): 1.500 seconds

• Optimal for this counter example will be synchronizing once, when entirely finished with the local sum

#### Outline

• Applying Locks

#### Concurrent Data Structures

#### • Ordering with Condition Variables

#### • Semaphores

#### Thread-safe data structures

- "Thread safe" works even if used by multiple threads concurrently
  - Can apply to various libraries, functions, and data structures
- Simple data structures implementations are usually not thread safe
  - Some global state needs to be shared among all threads
  - Need to protect critical sections
- Challenge: multiple function calls each access same shared structure
  Need to identify the critical section in each and lock it with shared lock

### Linked List

```
void List_Insert(list_t *L, int key) {
  node_t *new = malloc(sizeof(node_t));
  if (new == NULL) {
    perror("malloc");
    return; // fail
  }
  new->key = key;
  new->next = L->head;
  L->head = new;
  return; // success
}
```

#### Concurrent Linked List – Big lock approach

```
void List_Insert(list_t *L, int key) {
  pthread_mutex_lock(&L->lock);
  node t *new = malloc(sizeof(node_t));
  if (new == NULL) {
    perror("malloc");
    pthread_mutex_unlock(&L->lock);
    return; // fail
  new->key = key;
  new->next = L->head;
  L->head = new;
  pthread_mutex_unlock(&L->lock);
  return; // success
```

Most important part of this example. Don't forget to unlock if returning early.

 Much better than counter example, because we are only serializing the list itself. Hopefully the rest of the code can run concurrently.

#### Better Concurrent Linked List – Only lock critical section

```
void List_Insert(list_t *L, int key) {
  node t *new = malloc(sizeof(node t));
  if (new == NULL) {
    perror("malloc");
    return; // fail
  }
  new->key = key;
  new->next = L->head;
  L->head = new;
  return; // success
}
```

# **Check your understanding:**

Where is the critical section here?

#### Better Concurrent Linked List – Only lock critical section

```
void List_Insert(list_t *L, int key) {
  node t *new = malloc(sizeof(node t));
  if (new == NULL) {
    perror("malloc");
    return; // fail
  }
  new->key = key;
  new->next = L->head;
  L->head = new;
  return; // success
}
```

# **Check your understanding:**

Where is the critical section here?

#### What about malloc? Is that safe to use??

```
void List_Insert(list_t *L, int key) {
  node_t *new = malloc(sizeof(node_t));
  if (new == NULL) {
    perror("malloc");
    return; // fail
  }
  new->key = key;
  new->next = L->head;
  L->head = new;
  return; // success
```

- Thread-safe functions
  - Capable of being called concurrently and still functioning correctly
  - (Because they use locks!)
- How would we know if malloc is threadsafe?

Must check the library documentation to determine thread safety

<u>https://man7.org/linux/man-pages/man3/malloc.3.html</u>

• Malloc (and free) is indeed thread-safe

#### ATTRIBUTES top

For an explanation of the terms used in this section, see attributes(7).

Interface	Attribute	Value
<pre>malloc(), free(), calloc(), realloc()</pre>	Thread safety	MT-Safe

• If it wasn't, we would have to consider it another shared resource that needs to be locked

#### Better Concurrent Linked List – Only lock critical section

```
void List_Insert(list_t *L, int key) {
  node_t *new = malloc(sizeof(node_t));
  if (new == NULL) {
    perror("malloc");
    return; // fail
  }
  new->key = key;
  pthread_mutex_lock(&L->lock);
  new->next = L->head;
  L->head = new;
  pthread mutex unlock(&L->lock);
  return; // success
}
```

- Now new node is created locally in parallel
- Only actual access to the linked list is serialized

# Concurrent Queue

- Separate head & tail locks
- Allows concurrent add & remove
  - Up to 2 threads can access without waiting

21 22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

45

46

```
typedef struct ___node_t {
1
                             value;
        int
2
3
        struct node t
                            *next;
    } node t;
4
5
    typedef struct __queue_t {
6
        node_t
                            *head;
7
        node t
                            *tail;
8
        pthread_mutex_t
                            headLock;
9
        pthread_mutex_t
                            tailLock;
10
    } queue_t;
11
12
13
    void Queue_Init(queue_t *q) {
        node_t *tmp = malloc(sizeof(node_t));
14
        tmp->next = NULL;
15
        q->head = q->tail = tmp;
16
        pthread_mutex_init(&q->headLock, NULL); 44
17
        pthread_mutex_init(&q->tailLock, NULL);
18
19
```

```
void Queue_Enqueue(queue_t *q, int value) {
    node_t *tmp = malloc(sizeof(node_t));
    assert(tmp != NULL);
    tmp->value = value;
    tmp->next = NULL;
  pthread_mutex_lock(&q->tailLock);
    q->tail->next = tmp;
    q \rightarrow tail = tmp;
  __pthread_mutex_unlock(&q->tailLock);
int Queue_Dequeue(queue_t *q, int *value) {
  pthread_mutex_lock(&q->headLock);
    node_t *tmp = q->head;
    node_t *newHead = tmp->next;
    if (newHead == NULL) {
      pthread_mutex_unlock(&q->headLock);
        return -1; // queue was empty
    *value = newHead->value;
    q->head = newHead;
  pthread_mutex_unlock(&q->headLock);
    free(tmp);
    return 0;
```

23

# Concurrent Queue

• "tailLock" controls adding elements

21

22

23

24

25 26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

45

46

Looks similar to ListInsert

```
typedef struct ___node_t {
1
                             value;
        int
2
        struct node t
                            *next;
3
    } node t;
4
5
    typedef struct __queue_t {
6
        node_t
                            *head;
7
        node t
                            *tail;
8
        pthread_mutex_t
                            headLock;
9
        pthread_mutex_t
                             tailLock;
10
    } queue_t;
11
12
13
    void Queue_Init(queue_t *q) {
        node_t *tmp = malloc(sizeof(node_t));
14
        tmp->next = NULL;
15
        q->head = q->tail = tmp;
16
        pthread_mutex_init(&q->headLock, NULL); 44
17
        pthread_mutex_init(&q->tailLock, NULL);
18
19
```

```
void Queue_Enqueue(queue_t *q, int value) {
    node_t *tmp = malloc(sizeof(node_t));
    assert(tmp != NULL);
    tmp->value = value;
    tmp->next = NULL;
  pthread_mutex_lock(&q->tailLock);
    q->tail->next = tmp;
    q \rightarrow tail = tmp;
  pthread_mutex_unlock(&q->tailLock);
int Queue_Dequeue(queue_t *q, int *value) {
  pthread_mutex_lock(&q->headLock);
    node_t *tmp = q->head;
    node_t *newHead = tmp->next;
    if (newHead == NULL) {
      pthread_mutex_unlock(&q->headLock);
        return -1; // queue was empty
    *value = newHead->value;
    q->head = newHead;
  pthread_mutex_unlock(&q->headLock);
    free(tmp);
    return 0;
```

# Concurrent Queue

• Head lock controls removing elements from front

21

22

23

24

25 26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

45

46

Needs to lock almost entire function

```
typedef struct ___node_t {
1
                              value;
        int
2
        struct node t
                             *next;
3
    } node t;
4
5
    typedef struct __queue_t {
6
        node_t
                             *head;
7
        node_t
                             *tail;
8
        pthread_mutex_t
                            headLock;
9
        pthread_mutex_t
                             tailLock;
10
    } queue_t;
11
12
13
    void Queue_Init(queue_t *q) {
        node_t *tmp = malloc(sizeof(node_t));
14
        tmp->next = NULL;
15
        q->head = q->tail = tmp;
16
        pthread_mutex_init(&q->headLock, NULL); 44
17
        pthread_mutex_init(&g->tailLock, NULL);
18
19
```

```
void Queue_Enqueue(queue_t *q, int value) {
    node_t *tmp = malloc(sizeof(node_t));
    assert(tmp != NULL);
    tmp->value = value;
    tmp->next = NULL;
  pthread_mutex_lock(&q->tailLock);
    q->tail->next = tmp;
    q \rightarrow tail = tmp;
  pthread_mutex_unlock(&q->tailLock);
int Queue_Dequeue(queue_t *q, int *value) {
  pthread_mutex_lock(&q->headLock);
    node_t *tmp = q->head;
    node_t *newHead = tmp->next;
    if (newHead == NULL) {
      pthread_mutex_unlock(&q->headLock);
        return -1; // queue was empty
    *value = newHead->value;
    q->head = newHead;
    pthread_mutex_unlock(&q->headLock);
    free(tmp);
    return 0;
```

#### Concurrent Hash Table

- Each bucket is implemented with a Concurrent List
  - We don't have to define any locks!
  - (Locks are in the lists)
- A thread can access a bucket without blocking other threads' access to *other* buckets.
- Hash tables are ideal for concurrency.
  - Hash (bucket id) can be calculated without accessing a shared resource.
  - Distributed hash tables are used for huge NoSQL databases.

```
#define BUCKETS (101)
1
2
    typedef struct ___hash_t {
3
        list_t lists[BUCKETS];
    } hash_t;
5
6
    void Hash_Init(hash_t *H) {
7
        int i;
8
        for (i = 0; i < BUCKETS; i++) {
9
             List_Init(&H->lists[i]);
10
11
12
13
    int Hash_Insert(hash_t *H, int key) {
14
        int bucket = key % BUCKETS;
15
        return List_Insert(&H->lists[bucket], key);
16
17
    }
18
    int Hash_Lookup(hash_t *H, int key) {
19
        int bucket = key % BUCKETS;
20
        return List_Lookup(&H->lists[bucket], key);
21
22
```

#### Lock-free data structures

- In our original example, we put a lock around counter++
  - We could have instead used atomic\_fetch\_and\_add to update counter
  - Lock-free and *still* atomic!!
- This is possible with more complex data structures as well
  - Often based on a compare-and-swap (CAS) approach
  - <u>https://www.cs.cmu.edu/~410-s05/lectures/L31\_LockFree.pdf</u>
- Warning: these are not to be taken lightly
  - Atomic instructions have performance costs on processors
  - Getting this correct involves really understanding hardware
  - <u>https://abseil.io/docs/cpp/atomic\_danger</u>

#### Break + Question: Where is the critical section for vector?

```
typedef struct {
   size_t size;
   size_t count;
   int** data;
} vector_t;
```

```
void vector_add(vector_t* v, int* item) {
    if (v->count == v->size) {
        v->size *= 2;
        v->data = realloc(v->data, sizeof(int*)*v->size);
    }
    v->data[v->count++] = item;
```

#### Break + Question: Where is the critical section for vector?

```
typedef struct {
   size_t size;
   size_t count;
   int** data;
} vector_t;
```

```
void vector_add(vector_t* v, int* item) {
    if (v->count == v->size) {
        v->size *= 2;
        v->data = realloc(v->data, sizeof(int*)*v->size);
    }
    v->data[v->count++] = item;
```

### Outline

• Applying Locks

• Concurrent Data Structures

Ordering with Condition Variables

• Semaphores

#### Requirements for sensible concurrency

#### Mutual exclusion

- Prevents corruption of data manipulated in critical sections
- Atomic instructions  $\rightarrow$  Locks  $\rightarrow$  Concurrent data structures
- Ordering (B runs after A)
  - By default, concurrency leads to a lack of control over ordering
  - We can use mutex'd variables to control ordering, but it's inefficient:
    - while(!myTurn) sleep(1);
  - We would like cooperating threads to be able to signal each other.
    - Park/unpark and futex could be used solve this problem
    - But we want a higher-level abstraction

### Barriers for all-or-nothing synchronization

- Barriers create synchronization points in the program
  - All threads must reach barrier before any thread continues
- pthread\_barrier\_init(barrier\_t)
- pthread\_barrier\_wait(barrier\_t)
- Use case: neural network processing
  - Spawn a pool of threads
  - Each thread handles a portion of the input data
  - Collect results from all threads at the end of the layer
  - Distribute results to appropriate threads for next layer

Basic Signaling with Condition Variable (condvar)

- Queue of waiting threads
  - Combine with a **flag** and a **mutex** to synchronize threads
- wait(condvar\_t, lock\_t)
  - Lock must be held when wait() is called
  - Puts the caller to sleep and releases lock (atomically)
  - When awoken, reacquires lock before returning
- signal(condvar\_t)
  - Wake a single waiting thread (if any are waiting)
  - Do nothing if there are no waiting threads
  - Called while holding the lock; takes the action after lock is released

#### Waiting for a thread to finish

```
pthread_t p1, p2;
```

```
// create child threads
pthread create(&p1, NULL, mythread, "A");
pthread create(&p2, NULL, mythread, "B");
. . .
// join waits for the child threads to finish
thr join(p1, NULL);
thr join(p2, NULL);
                                 How to implement
return 0;
                                 join?
```

# CV for child wait

• Must use mutex to protect "done" flag and condvar

```
int done = 0;
1
    pthread_mutex_t m = PTHREAD_MUTEX_INITIALIZER;
2
    pthread_cond_t c = PTHREAD_COND_INITIALIZER;
3
    void thr_exit() {
5
        Pthread_mutex_lock(&m);
6
        done = 1;
7
        Pthread_cond_signal(&c);
8
        Pthread_mutex_unlock(&m);
9
10
    }
11
    void *child(void *arg) {
12
        printf("child\n");
13
        thr_exit();
14
        return NULL;
15
16
    }
17
    void thr_join() {
18
        Pthread_mutex_lock(&m);
19
        while (done == 0)
20
             Pthread cond wait(&c, &m);
21
        Pthread_mutex_unlock(&m);
22
    }
23
24
    int main(int argc, char *argv[]) {
25
        printf("parent: begin\n");
26
        pthread_t p;
27
        Pthread_create(&p, NULL, child, NULL);
28
        thr_join();
29
        printf("parent: end\n");
30
        return 0;
31
                                                35
32
```

# CV for child wait

• Must use mutex to protect "done" flag and condvar

Parent calls thr\_join()

wait()'s until done==1

```
int done = 0;
1
    pthread_mutex_t m = PTHREAD_MUTEX_INITIALIZER;
2
    pthread_cond_t c = PTHREAD_COND_INITIALIZER;
3
    void thr_exit() {
5
        Pthread_mutex_lock(&m);
6
        done = 1;
7
        Pthread_cond_signal(&c);
8
        Pthread_mutex_unlock(&m);
9
10
    }
11
    void *child(void *arg) {
12
        printf("child\n");
13
        thr_exit();
14
        return NULL;
15
16
    }
17
    void thr_join() {
18
        Pthread mutex lock (&m);
19
        while (done == 0)
20
             Pthread cond wait(&c, &m);
21
        Pthread_mutex_unlock(&m);
22
23
24
    int main(int argc, char *argv[]) {
25
        printf("parent: begin\n");
26
        pthread_t p;
27
        Pthread_create(&p, NULL, child, NULL);
28
        thr_join();
29
        printf("parent: end\n");
30
        return 0;
31
                                                36
32
```

# CV for child wait

• Must use mutex to protect "done" flag and condvar

- Parent calls thr\_join()
  - wait()'s until done==1
- Child calls thr\_exit()
  - sets done to 1
  - calls signal()
  - unlocks mutex

```
int done = 0;
1
    pthread_mutex_t m = PTHREAD_MUTEX_INITIALIZER;
    pthread_cond_t c = PTHREAD_COND_INITIALIZER;
3
    void thr_exit() {
5
        Pthread_mutex_lock(&m);
        done = 1;
7
        Pthread_cond_signal(&c);
8
        Pthread_mutex_unlock(&m);
9
10
11
    void *child(void *arg) {
12
        printf("child\n");
13
        thr_exit();
14
        return NULL;
15
16
    }
17
    void thr_join() {
18
        Pthread_mutex_lock(&m);
19
        while (done == 0)
20
             Pthread cond wait(&c, &m);
21
        Pthread_mutex_unlock(&m);
22
23
24
    int main(int argc, char *argv[]) {
25
        printf("parent: begin\n");
26
        pthread_t p;
27
        Pthread_create(&p, NULL, child, NULL);
28
        thr_join();
29
        printf("parent: end\n");
30
        return 0;
31
                                                37
32
```

# Buggy attempts to wait for a child, no flag

#### **Incorrect Code**

```
void thr_exit() {
   1
   2
            Pthread_mutex_lock(&m);
Child
            Pthread_cond_signal(&c);
   3
           Pthread mutex_unlock(&m);
   5
   6
       void thr_join() {
   7
Parent
            Pthread_mutex_lock(&m);
   8
   9
            Pthread_cond_wait(&c, &m);
           Pthread_mutex_unlock(&m);
  10
  11
```

#### **Correct Code**

```
void thr_exit() {
5
         Pthread_mutex_lock(&m);
6
         done = 1;
7
         Pthread_cond_signal(&c);
8
         Pthread_mutex_unlock(&m);
9
10
    void thr_join() {
18
        Pthread_mutex_lock(&m);
19
        while (done == 0)
20
            Pthread_cond_wait(&c, &m);
21
        Pthread_mutex_unlock(&m);
22
23
```

1) Without *done* variable, the child could run first and signal before the parent starts waiting for the child.

# Buggy attempts to wait for a child, no mutex

#### Incorrect Code

```
Correct Code
```

```
void thr_exit() {
5
         Pthread_mutex_lock(&m);
6
         done = 1;
7
         Pthread_cond_signal(&c);
8
         Pthread_mutex_unlock(&m);
9
10
    void thr_join() {
18
        Pthread_mutex_lock(&m);
19
        while (done == 0)
20
            Pthread_cond_wait(&c, &m);
21
        Pthread_mutex_unlock(&m);
22
23
```

 Without a lock, the parent could see done==0, then the child could finish and signal, then the parent would start waiting (after missing the signal).

# Spurious (fake) wakeups

- Pthreads allows wakeup to return not just when a signaled, but also when a *timer expires* or for *no reason at all!*
- Spurious wakeups were included in the specification because they may allow some implementations be more efficient.
- There is no guarantee that the condition you've been waiting for is true when you are awoken
- So, we must also use a "predicate loop." (*while*, not *if*)

```
int done = 0;
    pthread_mutex_t m = PTHREAD_MUTEX_INITIALIZER;
    pthread_cond_t c = PTHREAD_COND_INITIALIZER;
    void thr_exit() {
        Pthread_mutex_lock(&m);
        done = 1;
        Pthread_cond_signal(&c);
        Pthread_mutex_unlock(&m);
9
10
11
    void *child(void *arg) {
12
        printf("child\n");
13
        thr_exit();
14
15
        return NULL;
16
17
    void thr_join() {
18
        Pthread_mutex_lock(&m);
19
        while (done == 0)
20
            Pthread_cond_wait(&c, &m);
21
        Pthread_mutex_unlock(&m);
24
    int main(int argc, char *argv[]) {
25
        printf("parent: begin\n");
26
        pthread_t p;
27
        Pthread_create(&p, NULL, child, NULL);
28
        thr_join();
29
        printf("parent: end\n");
30
        return 0;
31
                                               40
32
```

#### Another Example: Produce/Consumer Problem

- We have multiple producers and multiple consumers that communicate with a shared queue (FIFO buffer).
  - Concurrent queue allows work to happen asynchronously.
  - Buffer has finite size (does not dynamically expand)
- Two operations:
  - *Put,* which should block (wait) if the buffer is **full**.
  - *Get,* which should block (wait) if the buffer is **empty**.
- This is more complex than a (linked-list-based) concurrent queue because of the finite size and waiting.
- Example scenario: request queue in a multi-threaded web server.

### Managing the buffer

```
int buffer[MAX];
1
    int fill
               = 0;
2
    int use
               = 0;
3
    int count = 0;
4
5
    void put(int value) {
6
        buffer[fill] = value;
7
        fill = (fill + 1)  % MAX;
8
        count++;
9
10
11
    int get() {
12
13
         int tmp = buffer[use];
        use = (use + 1) % MAX;
14
15
        count--;
16
        return tmp;
17
```

- A simple implementation of a circular buffer that stores data in a fixed-size array.
- *fill* is the index of the tail
- *use* is the index of the head

This simple implementation assumes:

- Concurrency is managed elsewhere
- It will overwrite data if we try to put more than MAX elements.

```
cond_t empty, fill;
1
    mutex_t mutex;
2
3
    void *producer(void *arg) {
4
        int i;
5
        for (i = 0; i < loops; i++) {
6
            Pthread_mutex_lock(&mutex);
7
             while (count == MAX)
8
                 Pthread_cond_wait(&empty, &mutex);
9
10
            put(i);
            Pthread_cond_signal(&fill);
11
            Pthread_mutex_unlock(&mutex);
12
13
14
    }
15
    void *consumer(void *arg) {
16
        int i;
17
        for (i = 0; i < loops; i++) {
18
            Pthread_mutex_lock(&mutex);
19
             while (count == 0)
20
                 Pthread_cond_wait(&fill, &mutex);
21
             int tmp = get();
22
            Pthread_cond_signal(&empty);
23
            Pthread_mutex_unlock(&mutex);
24
            printf("%d\n", tmp);
25
26
27
```

- Always acquire *mutex* 
  - Must use same mutex in both functions
- Use *two condvars*

```
cond_t empty, fill;
1
    mutex_t mutex;
2
3
    void *producer(void *arg) {
4
        int i;
5
        for (i = 0; i < loops; i++) {
6
             Pthread_mutex_lock(&mutex);
7
             while (count == MAX)
8
                 Pthread_cond_wait(&empty, &mutex);
9
             put(i);
10
             Pthread_cond_signal(&fill);
11
             Pthread_mutex_unlock(&mutex);
12
13
14
15
    void *consumer(void *arg) {
16
        int i;
17
        for (i = 0; i < loops; i++) {
18
             Pthread_mutex_lock(&mutex);
19
             while (count == 0)
20
                 Pthread_cond_wait(&fill, &mutex);
21
             int tmp = get();
22
             Pthread_cond_signal(&empty);
23
             Pthread_mutex_unlock(&mutex);
24
            printf("%d\n", tmp);
25
26
27
```

- Always acquire *mutex* 
  - Must use same mutex in both functions
- Use *two condvars*
- Producer waits on **empty** while the buffer is full
  - Producer signals fill after put

```
cond_t empty, fill;
1
    mutex_t mutex;
3
    void *producer(void *arg) {
        int i;
5
        for (i = 0; i < loops; i++) {
6
             Pthread_mutex_lock(&mutex);
7
             while (count == MAX)
8
                 Pthread_cond_wait(&empty, &mutex);
9
             put(i);
10
             Pthread_cond_signal(&fill);
11
             Pthread_mutex_unlock(&mutex);
12
13
14
15
    void *consumer(void *arg) {
16
        int i;
17
        for (i = 0; i < loops; i++) {
18
             Pthread_mutex_lock(&mutex);
19
             while (count == 0)
20
                 Pthread_cond_wait(&fill, &mutex);
21
             int tmp = get();
22
             Pthread_cond_signal(&empty);
23
             Pthread_mutex_unlock(&mutex);
24
             printf("%d\n", tmp);
25
26
27
```

- Always acquire *mutex* 
  - Must use same mutex in both functions
- Use *two condvars*
- Producer waits on **empty** while the buffer is full
  - Producer signals fill after put
- Consumer waits on **fill** while the buffer is empty
  - Consumer signals empty after get

```
cond_t empty, fill;
1
    mutex_t mutex;
3
    void *producer(void *arg) {
        int i;
5
        for (i = 0; i < loops; i++) {
6
             Pthread_mutex_lock(&mutex);
7
             while (count == MAX)
8
                 Pthread_cond_wait(&empty, &mutex);
9
             put(i);
10
             Pthread_cond_signal(&fill);
11
             Pthread_mutex_unlock(&mutex);
12
13
14
15
    void *consumer(void *arg) {
16
        int i;
17
        for (i = 0; i < loops; i++) {
18
             Pthread_mutex_lock(&mutex);
19
             while (count == 0)
20
                 Pthread_cond_wait(&fill, &mutex);
21
             int tmp = get();
22
             Pthread_cond_signal(&empty);
23
             Pthread_mutex_unlock(&mutex);
24
             printf("%d\n", tmp);
25
26
27
```

- Always acquire *mutex* 
  - Must use same mutex in both functions
- Use *two condvars*
- Producer waits on **empty** while the buffer is full
  - Producer signals fill after put
- Consumer waits on **fill** while the buffer is empty
  - Consumer signals empty after get
- Loops re-check count condition after breaking out of wait, to handle spurious wakeups.

Broadcast makes more complex conditions possible

- Recall that signal wakes one waiting thread (FIFO)
  - But there are times when threads are not all equivalent
  - The signal may not be serviceable by any of the threads
- For example, consider memory allocation/free requests
  - An allocation can only be serviced by free of >= size
- pthread\_cond\_broadcast wakes all threads
  - This approach may be inefficient, but it may be necessary to ensure progress

#### Condition Variable: rules of thumb

- Shared state determines if condition is true or not
  - Check the state in a while loop before waiting on condvar
- Use a mutex to protect:
  - The shared state on which condition is based, and
  - Operations on the condvar
- Remember to acquire the mutex before calling cond\_signal() and cond\_broadcast()
- Use different condvars for different conditions
  - Sometimes, cond\_broadcast() helps if you can't find an elegant solution using cond\_signal()

#### Break + Administrivia

- Get started on PCLab
  - Hard part isn't the implementation, but the *debugging*
  - PCLab discussion on Friday

- Midterm exam is coming soon!
  - Two weeks from now
  - We'll have a review session during discussion next week Friday
  - I'll distribute a practice exam as well

### Outline

• Applying Locks

• Concurrent Data Structures

• Ordering with Condition Variables

Semaphores

### Generalizing Synchronization

- Condvars have no state or lock, just a waiting queue
  - The rest is handled by the programmer
- Semaphores are a generalization of condvars and locks
  - Includes internal (locked) state
  - A little harder to understand and use, but can do everything

# Semaphores (by Edsger Dijkstra, 1965)

- Keeps an internal integer value that determines what happens to a calling thread
- Init(val)
  - Set the initial internal value
  - Value cannot otherwise be directly modified
- Up/Signal/Post/V() (from Dutch *verhogen* "increase")
  - Increase the value. If there is a waiting thread, wake one.
- Down/Wait/Test/P() (from Dutch proberen "to try")
  - Decrease the value. Wait if the value is negative.



Dijkstra invented Dijkstra's Algorithm!

Also Semaphores and the entire field of Concurrent Programming

https://en.wikipedia.org/ wiki/Edsger W. Dijkstra

# Semaphores vs Condition Variables

- Semaphores
- Up/Post: increase value and wake one waiting thread
- Down/Wait: decrease value and wait if it's negative

- Condition Variables
- Signal: wake one waiting thread
- *Wait*: wait

- Compared to CVs, Semaphores add an integer value that controls when waiting is necessary
  - Value counts the quantity of a shared resource currently available
  - Up makes a resource available, down reserves a resource
  - Negative value -X means that X threads are waiting for the resource

# Check your understanding

• How would we build a mutex out of a semaphore?
typdef struct {
 sem\_init(sem\_t\*,

```
sem_t sem;
} lock_t;
init(lock_t* lock){
```

```
}
acquire(lock_t* lock) {
```

```
}
release(lock_t* lock) {
```

}

# Check your understanding

• How would we build a mutex out of a semaphore?
typdef struct {

```
sem_t sem;
} lock_t;
init(lock_t* lock){
   sem_init(&(lock->sem), 1);
}
acquire(lock_t* lock) {
   sem_wait(&(lock->sem));
}
```

```
}
release(lock_t* lock) {
   sem_post(&(lock->sem));
```

# Implementing a lock with a semaphore

- Choose an appropriate initial value for the semaphore
- To implement a *Lock:* 
  - Initialize to 1 (access to the critical section is the one shared resource)
  - **Lock**  $\rightarrow$  **Down**: (decreases the value and waits if negative)
    - Will decrease the value to 0 if it lock *is not* already taken
    - Will decrease the value to -1 and wait if the lock *is* taken (value was 0)
  - **Unlock**  $\rightarrow$  **Up**: (increases the value and wakes one waiting thread)
    - If value was 0, then no thread was waiting, and no thread is woken
    - If value was -1, then one thread was waiting, and it is woken
    - If value was -x, then x threads are waiting, one is woken, value becomes -(x-1).
  - If value is already 1, Up should not be called. (Unlock before lock?!)

# Semaphores reduce effort for numerical conditions

Semaphore

#### **Condition Variable**

```
void thr exit() {
       void thr_exit() {
   5
            Pthread_mutex_lock(&m);
                                                       sem post(&s);
   6
Child
            done = 1;
                                                   }
            Pthread_cond_signal(&c);
   8
            Pthread_mutex_unlock(&m);
   9
  10
       void thr_join() {
  18
                                                   void thr join() {
           Pthread_mutex_lock(&m);
  19
Parent
                                                       sem_wait(&s);
           while (done == 0)
  20
                                                   }
               Pthread_cond_wait(&c, &m);
  21
           Pthread_mutex_unlock(&m);
  22
                                                   sem_init(&s, 0);
  23
```

- Want parent to wait immediately so initialize to 0
- If child thread finishes first, semaphore increments to 1

#### **Readers-Writers Problem**

 Some resources don't need strict mutual exclusion, especially if they have many *read-only* accesses. (eg., a linked list)

- Any number of readers can be active simultaneously, but
- Writes must be mutually exclusive AND cannot happen during read

• API:

- acquire\_read\_lock(), release\_read\_lock()
- acquire\_write\_lock(), release\_write\_lock()

#### Reader-writer Lock

• "lock" semaphore used as a mutex

```
typedef struct _rwlock_t {
1
                        // binary semaphore (basic lock)
      sem_t lock;
2
      sem_t writelock; // used to allow ONE writer or MANY readers
3
      int
           readers;
                        // count of readers reading in critical section
4
    } rwlock_t;
5
6
    void rwlock_init(rwlock_t *rw) {
7
      rw->readers = 0;
8
      sem_init(&rw->lock, 0, 1);
9
      sem_init(&rw->writelock, 0, 1);
10
    }
11
12
    void rwlock_acquire_readlock(rwlock_t *rw) {
13
      sem_wait(&rw->lock);
14
      rw->readers++;
15
      if (rw->readers == 1)
16
        sem_wait(&rw->writelock); // first reader acquires writelock
17
      sem post(&rw->lock);
18
19
20
    void rwlock_release_readlock(rwlock_t *rw) {
21
      sem_wait(&rw->lock);
22
      rw->readers--;
23
      if (rw->readers == 0)
24
        sem_post(&rw->writelock); // last reader releases writelock
25
      sem_post(&rw->lock);
26
27
28
    void rwlock_acquire_writelock(rwlock_t *rw) {
29
      sem_wait(&rw->writelock);
30
31
    }
32
    void rwlock_release_writelock(rwlock_t *rw) {
33
      sem_post(&rw->writelock);
34
35
```

### Reader-writer Lock

- "writelock" must be held during read to block writes or during write to block reads.
- During reads
  - Number of active readers is counted.
  - First/last reader handles acquiring/releasing writelock.

```
typedef struct _rwlock_t {
1
                       // binary semaphore (basic lock)
      sem_t lock;
2
      sem_t writelock; // used to allow ONE writer or MANY readers
3
      int
           readers;
                        // count of readers reading in critical section
4
      rwlock_t;
5
6
    void rwlock_init(rwlock_t *rw) {
7
      rw->readers = 0;
8
      sem_init(&rw->lock, 0, 1);
9
      sem_init(&rw->writelock, 0, 1);
10
11
12
    void rwlock_acquire_readlock(rwlock_t *rw) {
13
      sem_wait(&rw->lock);
14
      rw->readers++;
15
      if (rw->readers == 1)
16
        sem_wait(&rw->writelock); // first reader acquires writelock
17
      sem post(&rw->lock);
18
19
20
    void rwlock_release_readlock(rwlock_t *rw) {
21
22
      sem_wait(&rw->lock);
      rw->readers--;
23
      if (rw->readers == 0)
24
        sem_post(&rw->writelock); // last reader releases writelock
25
      sem_post(&rw->lock);
26
27
28
    void rwlock_acquire_writelock(rwlock_t *rw) {
29
      sem_wait(&rw->writelock);
30
31
32
    void rwlock_release_writelock(rwlock_t *rw) {
33
      sem_post(&rw->writelock);
34
35
```

### Classical concurrency problems

- Note that this solution could starve writers
  - There might always be readers in the critical section

- Full solution to readers-writers problem with progress guarantee
  - <u>https://en.wikipedia.org/wiki/Readers%E2%80%93writers\_problem</u>

- Generally: try to map your problem to one of these solved problems
  - Producers/Consumers
  - Readers/Writers

### Outline

• Applying Locks

• Concurrent Data Structures

• Ordering with Condition Variables

• Semaphores