


# CS110 Lecture 15: Mutexes and Condition Variables

CS110: Principles of Computer Systems

Winter 2021-2022

Stanford University

Instructors: Nick Troccoli and Jerry Cain



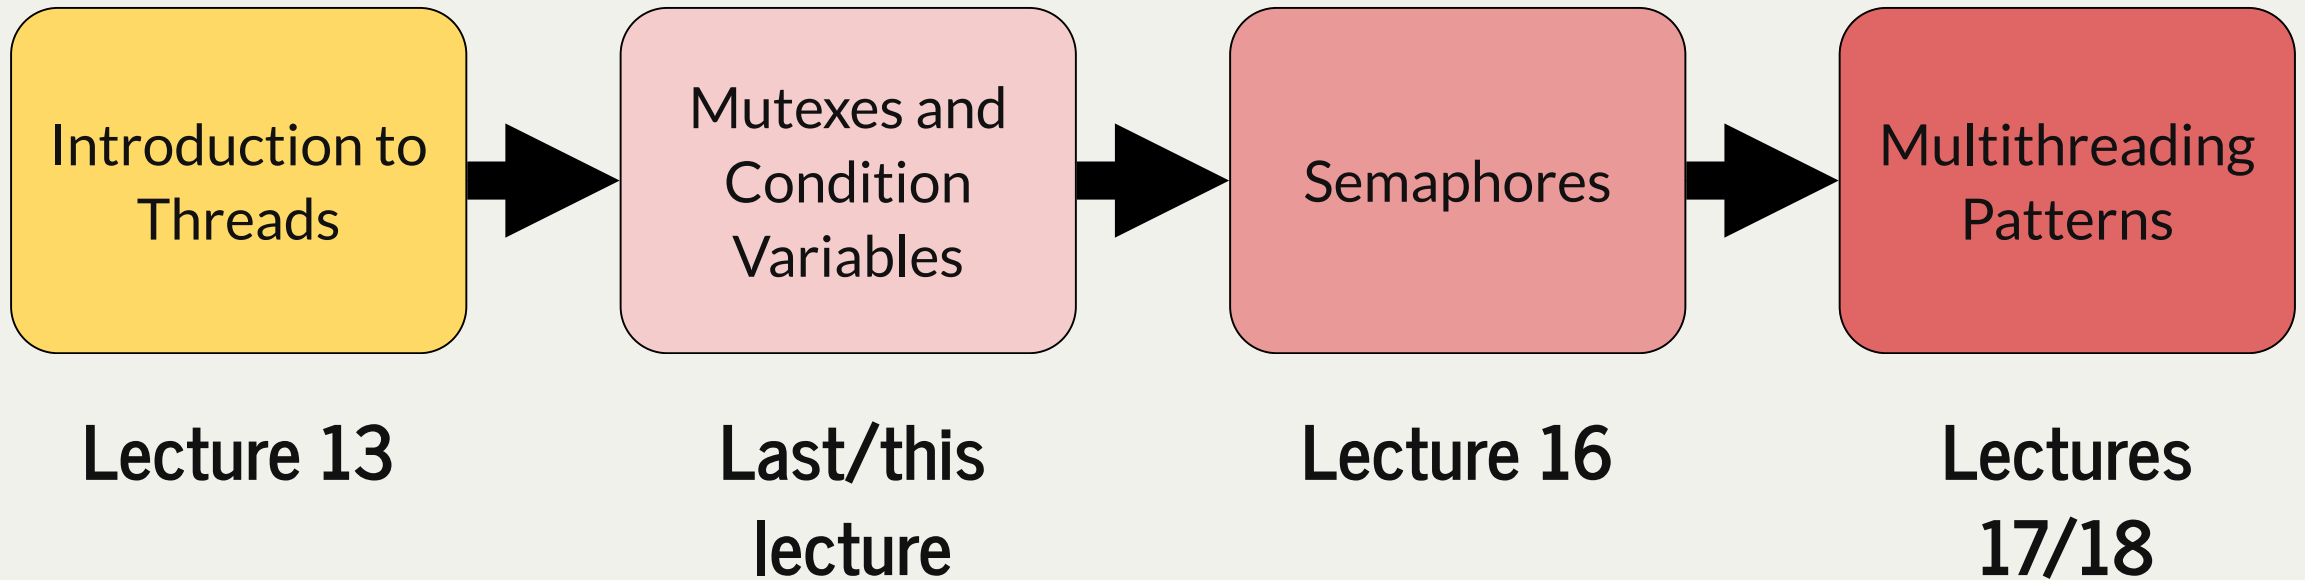
[https://commons.wikimedia.org/wiki/File:An\\_illustration\\_of\\_the\\_dining\\_philosophers\\_problem.png](https://commons.wikimedia.org/wiki/File:An_illustration_of_the_dining_philosophers_problem.png)



[PDF of this presentation](#)

**CS110 Topic 3:** How can we have  
concurrency within a single process?

# Learning About Multithreading



assign5: implement your own multithreaded news aggregator to quickly fetch news from the web!

# Learning Goals

- Get more practice using mutexes to prevent race conditions
- Apply the race condition checklist to eliminate race conditions
- Learn how condition variables can let threads signal to each other
- Get practice with the "available permits" resource model

# Plan For Today

- **Recap:** Race Conditions and Mutexes
- Dining With Philosophers
- Encoding Resource Constraints
- Condition Variables

# Plan For Today

- Recap: Race Conditions and Mutexes
- Dining With Philosophers
- Encoding Resource Constraints
- Condition Variables

# Mutexes

- We can create a lock-and-key combo by creating a variable of type **mutex**.
- A **mutex** is technically a type of lock; there are others, but we focus just on **mutexes**
- When you create a mutex, it is initially unlocked with the key available
- You call **lock()** on the mutex to attempt to lock it and take the key
- You call **unlock()** on a mutex *if you have ownership of it* and wish to unlock it and return the key. That thread continues normally; one waiting thread (if any) then takes the lock and is scheduled to run.

```
// Assume multiple threads share this same mutex
mutex myLock;

...

myLock.lock();
// only one thread can be executing here at a time
myLock.unlock()
```

# Mutexes

When a thread calls `lock()`:

- **If the lock is unlocked:** the thread takes the lock and continues execution
- **If the lock is locked:** the thread blocks and waits until the lock is unlocked
- **If multiple threads are waiting for a lock:** they all wait until it's unlocked, one receives lock (not necessarily one waiting longest)

```
// Assume multiple threads share this same mutex
mutex myLock;

...

myLock.lock();
// only one thread can be executing here at a time
myLock.unlock()
```



# Mutex Usage

1. Identify a critical section; a section that only one thread should execute at a time.
2. Create a mutex and *pass it by reference* to all threads executing that critical section
3. Add a line to lock the mutex at the start of the critical section
4. Add a line to unlock the mutex at the end of the critical section

```
// Assume multiple threads share this same mutex
mutex myLock;

...

myLock.lock();
// only one thread can be executing here at a time
myLock.unlock()
```

# Mutex Usage

1. Identify a critical section; a section that only one thread should execute at a time.
2. Create a mutex and *pass it by reference* to all threads executing that critical section
3. Add a line to lock the mutex at the start of the critical section
4. Add a line to unlock the mutex at the end of the critical section

```
1 static void sellTickets(size_t id, size_t& remainingTickets) {
2     while (true) {
3
4         if (remainingTickets == 0) break;
5         remainingTickets--;
6         sleep_for(500); // simulate "selling a ticket"
7         cout << oslock << "Thread #" << id << " sold a ticket (" << remainingTickets << " remain)." << endl << osunlock;
8
9     }
10    cout << oslock << "Thread #" << id << " sees no remaining tickets to sell and exits." << endl << osunlock;
11 }
```

# Mutex Usage

1. Identify a critical section; a section that only one thread should execute at a time.
2. Create a mutex and *pass it by reference* to all threads executing that critical section
3. Add a line to lock the mutex at the start of the critical section
4. Add a line to unlock the mutex at the end of the critical section

(It turns out that the mutex type *can't* be passed by copy in C++; since it doesn't make sense to make a copy of a mutex).

```
1 int main(int argc, const char *argv[]) {
2     thread ticketAgents[kNumTicketAgents];
3     size_t remainingTickets = 250;
4     mutex counterLock;
5
6     for (size_t i = 0; i < kNumTicketAgents; i++) {
7         ticketAgents[i] = thread(sellTickets, i, ref(remainingTickets), ref(counterLock));
8     }
9     ...
10 }
```

# Mutex Usage

1. Identify a critical section; a section that only one thread should execute at a time.
2. Create a mutex and *pass it by reference* to all threads executing that critical section
3. Add a line to lock the mutex at the start of the critical section
4. Add a line to unlock the mutex at the end of the critical section

```
1 static void sellTickets(size_t id, size_t& remainingTickets, mutex& counterLock) {
2     while (true) {
3         size_t myTicket;
4
5         counterLock.lock();
6         if (remainingTickets == 0) {
7             counterLock.unlock();
8             break;
9         } else {
10            myTicket = remainingTickets;
11            remainingTickets--;
12            counterLock.unlock();
13        }
14
15        sleep_for(500); // simulate "selling a ticket"
16        cout << oslock << "Thread #" << id << " sold a ticket (" << myTicket << " remain)." << endl << osunlock;
17    }
18    ...
19 }
```

# Mutex Usage

1. Identify a critical section; a section that only one thread should execute at a time.
2. Create a mutex and *pass it by reference* to all threads executing that critical section
3. Add a line to lock the mutex at the start of the critical section
4. Add a line to unlock the mutex at the end of the critical section

```
1 static void sellTickets(size_t id, size_t& remainingTickets, mutex& counterLock) {
2     while (true) {
3         size_t myTicket;
4
5         counterLock.lock();
6         if (remainingTickets == 0) {
7             counterLock.unlock();
8             break;
9         } else {
10            myTicket = remainingTickets;
11            remainingTickets--;
12            counterLock.unlock();
13        }
14
15        sleep_for(500); // simulate "selling a ticket"
16        cout << oslock << "Thread #" << id << " sold a ticket (" << myTicket << " remain)." << endl << osunlock;
17    }
18    ...
19 }
```

# Deadlock

**Deadlock** is a situation where a thread or threads rely on mutually blocked-on resources that will never become available.

- **Example**: a thread must acquire a lock before proceeding. We forget to call unlock somewhere, so one thread keeps the lock forever while others are stuck waiting for the lock forever.

```
1 static void sellTickets(size_t id, size_t& remainingTickets, mutex& counterLock) {
2     while (true) {
3         size_t myTicket;
4
5         counterLock.lock();
6         if (remainingTickets == 0) {
7             break;
8         } else {
9             myTicket = remainingTickets;
10            remainingTickets--;
11            counterLock.unlock();
12        }
13        ...
14    }
15    ...
16 }
```

# The Race Condition Checklist

- Identify shared data that may be modified concurrently.** What shared data is used across threads, passed by reference or globally?
- Document and confirm an ordering of events that causes unexpected behavior.** What assumptions are made in the code that can be broken by certain orderings?
- Use concurrency directives to force expected orderings and add constraints.** How can we use mutexes, atomic operations, or other constraints to force the correct ordering(s)?

# The Race Condition Checklist

- Identify shared data that may be modified concurrently.** What shared data is used across threads, passed by reference or globally? *the ticket count.*
- Document and confirm an ordering of events that causes unexpected behavior. What assumptions are made in the code that can be broken by certain orderings?
- Use concurrency directives to force expected orderings and add constraints. How can we use mutexes, atomic operations, or other constraints to force the correct ordering(s)?



# The Race Condition Checklist

Identify shared data that may be modified concurrently. What shared data is used across threads, passed by reference or globally? *the ticket count.*

Document and confirm an ordering of events that causes unexpected behavior.

What assumptions are made in the code that can be broken by certain orderings? ***one thread will check-and-sell a ticket at a time.***

Use concurrency directives to force expected orderings and add constraints. How can we use mutexes, atomic operations, or other constraints to force the correct ordering(s)?

# The Race Condition Checklist

Identify shared data that may be modified concurrently. What shared data is used across threads, passed by reference or globally? *the ticket count.*

Document and confirm an ordering of events that causes unexpected behavior. What assumptions are made in the code that can be broken by certain orderings? *one thread will check-and-sell a ticket at a time.*

**Use concurrency directives to force expected orderings and add constraints.** How can we use mutexes, atomic operations, or other constraints to force the correct ordering(s)? *add a mutex that must be acquired before checking-and-selling a ticket.*

# The Race Condition Checklist

- ☑ **Identify shared data that may be modified concurrently.** What shared data is used across threads, passed by reference or globally? *the ticket count.*
- ☑ **Document and confirm an ordering of events that causes unexpected behavior.** What assumptions are made in the code that can be broken by certain orderings? *one thread will check-and-sell a ticket at a time.*
- ☑ **Use concurrency directives to force expected orderings and add constraints.** How can we use mutexes, atomic operations, or other constraints to force the correct ordering(s)? *add a mutex that must be acquired before checking-and-selling a ticket.*

# Plan For Today

- **Recap: Race Conditions and Mutexes**
- **Dining With Philosophers**
- Encoding Resource Constraints
- Condition Variables

# Dining Philosophers Problem

- This is a [canonical multithreading example](#) of the potential for deadlock and how to avoid it.
- Five philosophers sit around a **circular table**, eating spaghetti
- There is **one fork** for each of them
- Each philosopher **thinks, then eats**, and repeats this **three times** for their three daily meals.
- **To eat**, a philosopher must grab the fork on their left *and* the fork on their right. Then they chow on spaghetti to nourish their big, philosophizing brain. When they're full, they put down the forks in the same order they picked them up and return to thinking for a while.
- **To think**, the a philosopher keeps to themselves for some amount of time. Sometimes they think for a long time, and sometimes they barely think at all.



[https://commons.wikimedia.org/wiki/File:An\\_illustration\\_of\\_the\\_dining\\_philosophers\\_problem.png](https://commons.wikimedia.org/wiki/File:An_illustration_of_the_dining_philosophers_problem.png)

# Dining Philosophers Problem

**Goal:** we must encode resource constraints into our program.

**Example:** how many philosophers can hold a fork at the same time? One.

**How can we encode this into our program?** Let's make a mutex for each fork.

- Each philosopher either holds a fork or doesn't.
- A philosopher grabs a fork by locking that mutex. If the fork is available, the philosopher continues. Otherwise, it blocks until the fork becomes available and it can have it.
- A philosopher puts down a fork by unlocking that mutex.

```
1 static void philosopher(size_t id, mutex& left, mutex& right) {
2     ...
3 }
4
5 int main(int argc, const char *argv[]) {
6     mutex forks[kNumForks];
7     thread philosophers[kNumPhilosophers];
8     for (size_t i = 0; i < kNumPhilosophers; i++) {
9         mutex& left = forks[i];
10        mutex& right = forks[(i + 1) % kNumPhilosophers];
11        philosophers[i] = thread(philosopher, i, ref(left), ref(right));
12    }
13    for (thread& p: philosophers) p.join();
14    return 0;
15 }
```

# Dining Philosophers Problem

A philosopher thinks, then eats, and repeats this three times.

- **think** is modeled as sleeping the thread for some amount of time

```
static void think(size_t id) {
    cout << oslock << id << " starts thinking." << endl << osunlock;
    sleep_for(getThinkTime());
    cout << oslock << id << " all done thinking. " << endl << osunlock;
}

static void eat(size_t id, mutex& left, mutex& right) {
    ...
}

static void philosopher(size_t id, mutex& left, mutex& right) {
    for (size_t i = 0; i < kNumMeals; i++) {
        think(id);
        eat(id, left, right);
    }
}
```

# Dining Philosophers Problem

A philosopher thinks, then eats, and repeats this three times.

- **eat** is modeled as grabbing the two forks, sleeping for some amount of time, and putting the forks down.

```
1 static void eat(size_t id, mutex& left, mutex& right) {
2     left.lock();
3     right.lock();
4     cout << oslock << id << " starts eating om nom nom nom." << endl << osunlock;
5     sleep_for(getEatTime());
6     cout << oslock << id << " all done eating." << endl << osunlock;
7     left.unlock();
8     right.unlock();
9 }
```

*Spoiler:* there is a race condition here that leads to deadlock.



# The Race Condition Checklist

- Identify shared data that may be modified concurrently.** What shared data is used across threads, passed by reference or globally?
- Document and confirm an ordering of events that causes unexpected behavior.** What assumptions are made in the code that can be broken by certain orderings?
- Use concurrency directives to force expected orderings and add constraints.** How can we use mutexes, atomic operations, or other constraints to force the correct ordering(s)?

# The Race Condition Checklist

- Identify shared data that may be modified concurrently.** What shared data is used across threads, passed by reference or globally? *the "forks" (mutexes).*
- Document and confirm an ordering of events that causes unexpected behavior. What assumptions are made in the code that can be broken by certain orderings?
- Use concurrency directives to force expected orderings and add constraints. How can we use mutexes, atomic operations, or other constraints to force the correct ordering(s)?

# The Race Condition Checklist

- Identify shared data that may be modified concurrently. What shared data is used across threads, passed by reference or globally?
- Document and confirm an ordering of events that causes unexpected behavior.**  
What assumptions are made in the code that can be broken by certain orderings?
- Use concurrency directives to force expected orderings and add constraints. How can we use mutexes, atomic operations, or other constraints to force the correct ordering(s)?

# Dining Philosophers Problem

A philosopher thinks, then eats, and repeats this three times.

- **eat** is modeled as grabbing the two forks, sleeping for some amount of time, and putting the forks down.

```
1 static void eat(size_t id, mutex& left, mutex& right) {  
2     left.lock();  
3     right.lock();  
4     cout << oslock << id << " starts eating om nom nom nom." << endl << osunlock;  
5     sleep_for(getEatTime());  
6     cout << oslock << id << " all done eating." << endl << osunlock;  
7     left.unlock();  
8     right.unlock();  
9 }
```

*Discuss with your neighbor:* come up with an ordering of events that causes deadlock.

# Food For Thought

What if: all philosophers grab their left fork and then go off the CPU?

- **deadlock!** All philosophers will wait on their right fork, which will never become available.
- **Testing our hypothesis:** insert a `sleep_for` call on line 3, between getting left fork and right fork
- We should be able to insert a `sleep_for` call anywhere in a thread routine and have no concurrency issues.

```
1 static void eat(size_t id, mutex& left, mutex& right) {
2     left.lock();
3     sleep_for(5000); // artificially force off the processor
4     right.lock();
5     cout << oslock << id << " starts eating om nom nom nom." << endl << osunlock;
6     sleep_for(getEatTime());
7     cout << oslock << id << " all done eating." << endl << osunlock;
8     left.unlock();
9     right.unlock();
10 }
```

# The Race Condition Checklist

- Identify shared data that may be modified concurrently. What shared data is used across threads, passed by reference or globally?
- Document and confirm an ordering of events that causes unexpected behavior. What assumptions are made in the code that can be broken by certain orderings? *we are assuming that someone is always able to pick up 2 forks.*
- Use concurrency directives to force expected orderings and add constraints. How can we use mutexes, atomic operations, or other constraints to force the correct ordering(s)?

# The Race Condition Checklist

- Identify shared data that may be modified concurrently. What shared data is used across threads, passed by reference or globally?
- Document and confirm an ordering of events that causes unexpected behavior. What assumptions are made in the code that can be broken by certain orderings? *we are assuming that someone is always able to pick up 2 forks.*
- Use concurrency directives to force expected orderings and add constraints. How can we use mutexes, atomic operations, or other constraints to force the correct ordering(s)?

# Race Conditions and Deadlock

When coding with threads, you need to ensure that:

- there are **never** any race conditions
- there's **zero** chance of deadlock; otherwise a subset of threads are forever starved
- **Race conditions** can generally be solved with **mutexes**.
  - We use them to mark the boundaries of critical regions and limit the number of threads present within them to be at most one.
- **Deadlock** can be programmatically prevented by implanting directives to limit the number of threads competing for a shared resource. **What does this look like?**



# Plan For Today

- **Recap:** Race Conditions and Mutexes
- Dining With Philosophers
- **Encoding Resource Constraints**
- Condition Variables

# Race Conditions and Deadlock

**Goal:** we must encode resource constraints into our program.

**Example:** how many philosophers can *try to eat* at the same time? **Four.**

- *Alternative:* how many philosophers can *eat* at the same time? **Two.**
- Why might the first one be better? Imposes less bottlenecking while still solving the issue.

**How can we encode this into our program?**

- let's add another shared variable representing a count of "permits" or "tickets" available.
- In order to try to eat (aka grab forks at all) a philosopher must get a permit
- Once done eating, a philosopher must return their permit

What does this look like in code?

- If there are permits available ( $\text{count} > 0$ ) then decrement by 1 and continue
- If there are no permits available ( $\text{count} == 0$ ) then block until a permit is available
- To return a permit, increment by 1 and continue

# Tickets, Please...

- Let's add a new variable in main called **permits**, and a lock for it called **permitsLock**, so that we can update it without race conditions.
- We pass these to each philosopher by reference.

```
1 int main(int argc, const char *argv[]) {
2     // NEW
3     size_t permits = 4;
4     mutex permitsLock;
5
6     mutex forks[5];
7     thread philosophers[5];
8     for (size_t i = 0; i < 5; i++) {
9         mutex& left = forks[i];
10        mutex& right = forks[(i + 1) % 5];
11        philosophers[i] = thread(philosopher, i, ref(left), ref(right), ref(permits), ref(permitsLock));
12    }
13    for (thread& p: philosophers) p.join();
14    return 0;
15 }
```

# Tickets, Please...

- Each philosopher takes two additional parameters as a result.
- The implementation of **think** does not change, as it does not use permits.

```
1 static void philosopher(size_t id, mutex& left, mutex& right,  
2                          size_t& permits, mutex& permitsLock) {  
3     for (size_t i = 0; i < kNumMeals; i++) {  
4         think(id);  
5         eat(id, left, right, permits, permitsLock);  
6     }  
7 }
```

# Tickets, Please...

- The implementation of `eat` changes:
  - Before eating, the philosopher must get a permit
  - After eating, the philosopher must return their permit.

```
1 static void eat(size_t id, mutex& left, mutex& right, size_t& permits, mutex& permitsLock) {
2     // NEW
3     waitForPermission(permits, permitsLock);
4
5     left.lock();
6     right.lock();
7     cout << oslock << id << " starts eating om nom nom nom." << endl << osunlock;
8     sleep_for(getEatTime());
9     cout << oslock << id << " all done eating." << endl << osunlock;
10
11     // NEW
12     grantPermission(permits, permitsLock);
13
14     left.unlock();
15     right.unlock();
16 }
```

# grantPermission

- How do we implement `grantPermission`?
- Recall: "To return a permit, increment by 1 and continue"

```
1 static void grantPermission(size_t& permits, mutex& permitsLock) {  
2     permitsLock.lock();  
3     permits++;  
4     permitsLock.unlock();  
5 }
```

# waitForPermission

- How do we implement `waitForPermission`?
- Recall:
  - "If there are permits available (`count > 0`) then decrement by 1 and continue"
  - "If there are no permits available (`count == 0`) then block until a permit is available"

```
1 static void waitForPermission(size_t& permits, mutex& permitsLock) {
2     while (true) {
3         permitsLock.lock();
4         if (permits > 0) break;
5         permitsLock.unlock();
6         sleep_for(10);
7     }
8     permits--;
9     permitsLock.unlock();
10 }
```

**Problem:** this is busy waiting! We are unnecessarily and arbitrarily using CPU time to check when a permit is available.

It would be nice if someone could let us know when they return their permit. Then, we can sleep until this happens.



# Plan For Today

- **Recap:** Race Conditions and Mutexes
- Dining With Philosophers
- Encoding Resource Constraints
- Condition Variables

# Plan For Today

A **condition variable** is a variable that can be shared across threads and used for one thread to notify to another thread when something happens. Conversely, a thread can also use this to *wait* until it is notified by another thread.

```
class condition_variable_any {
public:
    void wait(mutex& m);
    template <typename Pred> void wait(mutex& m, Pred pred);
    void notify_one();
    void notify_all();
};
```

- We can call **wait** to sleep until another thread signals this condition variable.
- We can call **notify\_all** to send a signal to waiting threads.

# waitForPermission

- How do we implement **waitForPermission**?
- Recall:
  - "If there are permits available ( $\text{count} > 0$ ) then decrement by 1 and continue"
  - "If there are no permits available ( $\text{count} == 0$ ) then block until a permit is available"

## Idea:

- when someone returns a permit and it is the only one now available, signal.
- if we need a permit but there are none available, wait.

# Condition Variables

Now we must create a condition variable to pass by reference to all threads.

```
1 int main(int argc, const char *argv[]) {
2     mutex forks[kNumForks];
3
4     size_t permits = kNumForks - 1;
5     mutex permitsLock;
6
7     // NEW
8     condition_variable_any permitsCV;
9
10    thread philosophers[kNumPhilosophers];
11    for (size_t i = 0; i < kNumPhilosophers; i++) {
12        mutex& left = forks[i];
13        mutex& right = forks[(i + 1) % kNumForks];
14        philosophers[i] = thread(philosopher, i, ref(left), ref(right),
15                                ref(permits), ref(permitsCV), ref(permitsLock));
16    }
17    for (thread& p: philosophers) p.join();
18    return 0;
19 }
```



dining-philosophers-with-cv-wait.cc

# grantPermission

For `grantPermission`, we must signal when we make permits go from 0 to 1.

```
1 static void grantPermission(size_t& permits, condition_variable_any& permitsCV, mutex& pe
2     permitsLock.lock();
3     permits++;
4     if (permits == 1) permitsCV.notify_all();
5     permitsLock.unlock();
6 }
```



# waitForPermission

For `waitForPermission`, if no permits are available we must wait until one becomes available.

```
1 static void waitForPermission(size_t& permits, condition_variable_any& permitsCV, mutex&
2   permitsLock.lock();
3   while (permits == 0) permitsCV.wait(permitsLock);
4   permits--;
5   permitsLock.unlock();
6 }
```

Here's what `cv.wait` does:

1. it puts the caller to sleep *and* unlocks the given lock, all atomically
2. it wakes up when the cv is signaled
3. upon waking up, it tries to acquire the given lock (and blocks until it's able to do so)
4. then, `cv.wait` returns



# Recap

- **Recap:** Race Conditions and Mutexes
- Dining With Philosophers
- Encoding Resource Constraints
- Condition Variables

**Next time:** more condition variables and introducing semaphores