CS 31: Introduction to Computer Systems 22-23 OS Processes and Parellelism 04-15-2025



OS Big Picture Goals

- OS is a layer of code between user programs and hardware.
- Goal: Make life easier for users and programmers.
- How can the OS do that?

Key OS Responsibilities

- 1. Simplifying abstractions for programs
- 2. Resource allocation and/or sharing
- 3. Hardware gatekeeping and protection

Running multiple programs

More than 200 processes running on a typical desktop!

- Benefits: when I/O issued, CPU not needed
 - Allow another program to run
 - <u>Requires yielding and sharing memory</u>
- Challenges: what if one running program...
 - Monopolizes CPU, memory?
 - Reads/writes another's memory?
 - Uses I/O device being used by another?

OS: Turn undesirable into desirable

- Turn undesirable inconveniences: reality
 - Complexity of hardware
 - Single processor
 - Limited memory
- Into desirable conveniences: illusions
 - Simple, easy-to-use resources
 - Multiple/unlimited number of processors
 - Large/unlimited amount of memory

Virtualization

- Rather than exposing real hardware, introduce a "virtual", abstract notion of the resource
- Multiple virtual processors
 - By rapidly switching CPU use
- Multiple virtual memories
 - By memory partitioning and re-addressing
- Virtualized devices
 - By simplifying interfaces, and using other resources to enhance function

We'll focus on the OS 'kernel'

- "Operating system" has many interpretations
 - E.g., all software on machine minus applications (user or even limited to 3rd party)
- Our focus is the *kernel*
 - What's necessary for everything else to work
 - Low-level resource control
 - Originally called the nucleus in the 60's

The Kernel

- All programs depend on it
 - Loads and runs them
 - Exports system calls to programs
- Works closely with hardware
 - Accesses devices
 - Responds to interrupts (hardware events)
- Allocates basic resources
 - CPU time, memory space
 - Controls I/O devices: display, keyboard, disk, network



Tron, 1982

Kernel provides common functions

- Some functions useful to many programs
 - I/O device control
 - Memory allocation
- Place these functions in central place (kernel)
 - Called by programs ("system calls")
 - Or accessed in response to hardware events
- What should functions be?
 - How many programs should benefit?
 - Might kernel get too big?

OS Kernel

• Big Design Issue: How do we make the OS efficient, reliable, and extensible?

• General OS Philosophy: The design and implementation of an OS involves a constant tradeoff between simplicity and performance.

- As a general rule, strive for simplicity.
 - except when you have a strong reason to believe that you need to make a particular component complicated to achieve acceptable performance
 - (strong reason = simulation or evaluation study)

Main Abstraction: The Process

- Abstraction of a running program
 - "a program in execution"
- Dynamic
 - Has state, changes over time
 - Whereas a program is static
- Basic operations
 - Start/end
 - Suspend/resume



Basic Resources for Processes

- To run, process needs some basic resources:
 - CPU
 - Processing cycles (time)
 - To execute instructions
 - Memory
 - Bytes or words (space)
 - To maintain state
 - Other resources (e.g., I/O)
 - Network, disk, terminal, printer, etc.

Machine State of a Process

- CPU or processor context
 - PC (program counter)
 - SP (stack pointer)
 - General purpose registers
- Memory
 - Code
 - Global Variables
 - Stack of activation records / frames
 - Other (registers, memory, kernel-related state)

Must keep track of these for every running process !

Anatomy of a Process

- Abstraction of a running program
 a dynamic "program in execution"
- OS keeps track of process state
 - What each process is doing
 - Which one gets to run next
- Basic operations
 - Suspend/resume (context switch)
 - Start (spawn), terminate (kill)



Managing Processes

- Processes created by calling fork()
 - "Spawning" a new process
- "Parent" process spawns "Child" process
 - Brutal relationship involving "zombies", "killing" and "reaping".
 (I'm not making this up!)
- Processes interact with one another by sending signals.

Program

Operating System

Computer Hardware

- OS sits between HW and User/Program
 - Manages HW & Makes system easier to use
 - Interrupt Driven: HW or User/Program interrupt it to do something on their behalf

Recap

- HW interrupt, Program system call traps to OS
- OS implements the Process Abstraction
 - Process: a running program
 - Lone view & private virtual address space
 - Why: efficient use of system resources
 - Multiprogramming and Timesharing. (ps $\,$ –A)
 - $\texttt{fork}(\xspace)$: system call to create a new process

Summary: Running a Program

Basic system calls:

- fork: spawns new process
 - Called once, Returns twice (in parent and child process)
- exit: terminates own process
 - Called once, never returns
 - Puts it into "zombie" status
- wait or waitpid: reap terminated children
- execup: runs new program in existing process
 - Called once, (normally) never returns

OS: Turn undesirable into desirable

- Turn undesirable inconveniences: reality
 - Complexity of hardware
 - Single processor
 - Limited memory
- Into desirable conveniences: illusions
 - Simple, easy-to-use resources
 - Multiple/unlimited number of processors
 - Large/unlimited amount of memory

Making Programs Run Faster

- In the "old days" (1980's 2005):
 - Algorithm too slow? Wait for HW to catch up.
- Modern CPUs exploit parallelism for speed:
 - Executes multiple instructions at once
 - Reorders instructions on the fly
- Today, can't make a single core go much faster.
 Limits on clock speed, heat, energy consumption
- Use extra transistors to put multiple CPU cores on the chip.
- Programmer's job to speed-up computation
 - Humans bad at thinking in parallel





Today's Processors are Multi-core

Multiple CPU cores/chip



All CPU cores share same Main Memory

OS manages all cores and memory

RAM contains processes' VM pages

Parallel Abstraction

- To speed up a job, must divide it across multiple cores.
- A process contains both execution information and memory/resources.
- What if we want to separate the execution information to give us parallelism in our programs?

Which components of a process might we replicate to take advantage of multiple CPU cores?

- A. The entire address space (memory not duplicated)
- B. Parts of the address space (memory stack)
- C. OS resources (open files, etc not duplicated.)
- D. Execution state (PC, registers, etc.)
- E. More than one of these (which?)

Which components of a process might we replicate to take advantage of multiple CPU cores?

- A. The entire address space (memory not duplicated)
- B. Parts of the address space (memory stack)
- C. OS resources (open files, etc not duplicated.)
- D. Execution state (PC, registers, etc.)
- E. More than one of these (which?)

Don't duplicate shared resources, duplicate resources where we need a private copy per thread: like execution state, and stack

Threads

- Modern OSes separate the concepts of processes and threads.
 - The process defines the address space and general process attributes (e.g., open files)
 - The thread defines a sequential execution stream within a process (PC, SP, registers)
- A thread is bound to a single process
 - Processes, however, can have multiple threads
 - Each process has at least one thread (e.g. main)

Processes versus Threads

- A process defines the address space, text, resources, etc.,
- A thread defines a single sequential execution stream within a process (PC, stack, registers).
- Threads extract the thread of control information from the process
- Threads are bound to a single process.
- Each process may have multiple threads of control within it.
 - The address space of a process is shared among all its threads
 - No system calls are required to cooperate among threads

Threads



This is the picture we've been using all along:

A process with a single thread, which has execution state (registers) and a stack.





Why Use Threads?

- Separating threads and processes makes it easier to support parallel applications:
 - Creating multiple paths of execution does not require creating new processes (less state to store, initialize Light Weight Process)
 - Low-overhead sharing between threads in same process (threads share page tables, access same memory)
- Concurrency (multithreading) can be very useful

Concurrency?

- Several computations or threads of control are executing simultaneously, and potentially interacting with each other.
- We can multitask! Why does that help?
 - Taking advantage of multiple CPUs / cores
 - Overlapping I/O with computation
 - Improving program structure

Recall: Processes



Scheduling Threads

- We have basically two options
 - 1. Kernel explicitly selects among threads in a process
 - 2. Hide threads from the kernel, and have a user-level scheduler inside each multithreaded process
- Why do we care?
 - Think about the overhead of switching between threads
 - Who decides which thread in a process should go first?
 - What about blocking system calls?

Pthreads Programming

PThreads: The **POSIX** threading interface

- The Portable Operating System Interface for UNIX
- A standard Interface to OS utilities

system calls have same prototype & semantics on all OSes

(e.g.) POSIX compliant code on Solaris will compile on Linux

Pthreads library contains functions for:

- Creating threads (and thread exit)
- Synchronizing threads
 - Coordinating their access to shared state

To compile: gcc myprog.c -lpthread

User-Level Threads


Kernel-Level Threads



Kernel Context switching over threads

Each process has explicitly mapped regions for stacks If you call thread_create() on a modern OS (Linux/Mac/Windows), which type of thread would you expect to receive? (Why? Which would you pick?)

A. Kernel threads

B. User threads

C. Some other sort of threads

If you call thread_create() on a modern OS (Linux/Mac/Windows), which type of thread would you expect to receive? (Why? Which would you pick?)

A. Kernel threads

B. User threads

C. Some other sort of threads

Kernel vs. User Threads

- Kernel-level threads
 - Integrated with OS (informed scheduling)
 - Slower to create, manipulate, synchronize
 - Requires getting the OS involved, which means changing context (relatively expensive)
- User-level threads
 - Faster to create, manipulate, synchronize
 - Not integrated with OS (uninformed scheduling)
 - If one thread makes a syscall, all of them get blocked because the OS doesn't distinguish.

- Code (text) shared by all threads in process
- Global variables and static objects are shared
 - Stored in the static data segment, accessible by any thread
- Dynamic objects and other heap objects are shared
 - Allocated from heap with malloc/free or new/delete
- Local variables should not be shared
 - Refer to data on the stack
 - Each thread has its own stack
 - Never pass/share/store a pointer to a local variable on another thread's stack!!



Ti

Tj

If threads Ti and Tj both execute function *foo* code:

Q1: which variables do they each get own copy of? which do they share?

Q2: which statement can affect values seen by the other thread?

Shared Virtual Address Space: 0×0 foo: **Instructions: Globals:** Heap: Stack: max



- Local variables should not be shared
 - Refer to data on the stack
 - Each thread has its own stack
 - <u>Never pass/share/store a pointer to a local variable on another thread's stack</u>



- Local variables should not be shared
 - Refer to data on the stack
 - Each thread has its own stack
 - <u>Never pass/share/store a pointer to a local variable on another thread's stack</u>



- Local variables should not be shared
 - Refer to data on the stack
 - Each thread has its own stack
 - <u>Never pass/share/store a pointer to a local variable on another thread's stack</u>



Thread-level Parallelism

- Speed up application by assigning portions to CPUs/cores that process in parallel
- Requires:
 - partitioning responsibilities (e.g., parallel algorithm)
 - managing their interaction
- Example: game of life (next lab)



If one CPU core can run a program at a rate of X, how quickly will the program run on two cores? Why?

- A. Slower than one core (<X)
- B. The same speed (X)
- C. Faster than one core, but not double (X-2X)
- D. Twice as fast (2X)
- E. More than twice as fast(>2X)

If one CPU core can run a program at a rate of X, how quickly will the program run on two cores? Why?

- A. Slower than one core (<X) (if we try to parallelize serial applications!)
- B. The same speed (X) (some applications are not parallelizable)
- C. Faster than one core, but not double (X-2X): most of the time: (some communication overhead to coordinate/synchronization of the threads)
- D. Twice as fast (2X)(class of problems called embarrassingly parallel programs. E.g. protein folding, SETI)
- E. More than twice as fast(>2X) (rare: possible if you have more CPU + more memory)

Parallel Speedup

- Performance benefit of parallel threads depends on many factors:
 - algorithm divisibility
 - communication overhead
 - memory hierarchy and locality
 - implementation quality
- For most programs, more threads means more communication, diminishing returns.

Summary

- Physical limits to how much faster we can make a single core run.
 - Use transistors to provide more cores.
 - Parallelize applications to take advantage.
- OS abstraction: thread
 - Shares most of the address space with other threads in same process
 - Gets private execution context (registers) + stack
- Coordinating threads is challenging!



Kernel-Level Threads



Kernel Context switching over threads

Each process has explicitly mapped regions for stacks

Common pthread functions

<u>Creating a thread (starts running start_func w/passed args):</u>

Joining (reaping) a thread (caller waits for thread to exit):

int pthread_join(pthread_t thread, void **retval);

Terminating a thread:

```
void pthread_exit(void *retval)
```

(or just return from thread's main function)

void *

int pthread_create(..., void *args);

void *: a pointer to any type (a generic pointer) can be used for return value and parameter types only

- all memory addresses take up the same number of bytes
 char *cptr; int *ptr; // store 8 byte addresses
- can pass the address of any type as a void *
 pthread_create(..., &x); // addr of an int
 pthread_create(..., &ch); //addr of a char

cannot de-reference a void * pointer directly
*args = 6; // store 6 in 1 byte? 2 bytes? 4 bytes?

– re-cast first before dereference!

*((int *) args) = 6; // store 6 in an int (4 bytes)

Example: hello.c with 2 threads



Concurrent Execution



Example: hello.c with N threads

- Code on slides is subset of full code
 - Minus error detection and handling for space
 - Minus other non-thread important code
- You can copy and try out the full code from here:

\$ cd ~/cs31/WeeklyLabs \$ mkdir week12

\$ cd week12

\$ pwd /home/you/cs31/WeeklyLabs/week12

\$ cp ~chaganti/public/cs31/s25/week12/* ./

[week12]\$./hello 3 Hello! I am thread 0 Hello! I am thread 2 Hello! I am thread 1 Goodbye! I was thread 2 Goodbye! I was thread 0 Goodbye! I was thread 1 count = 1283598 Main thread done [week12]\$./hello 3 Hello! I am thread 0 Hello! I am thread 1 Hello! I am thread 2 Goodbye! I was thread 2 Goodbye! I was thread 1 Goodbye! I was thread 0 count = 1123889

hello.c: main function

static unsigned long long count = 0; // global variable

```
int main(int argc, char *argv) {
    pthread_t *tids; // thread ids
    int ntids, i, *tid_args; //arguments passed to thread funcs.
```

initial var. declaration

hello.c: main function



hello.c: main function



```
ntids = atoi(argv[1]);
tids = (pthread_t *) malloc( sizeof(pthread_t) * ntids);
tid_args = (int *) malloc( sizeof(int) * ntids);
```

for (i=0; i < ntids; i++) { //loop through num. threads
 tid_args[i] = i; //input argument for each thread.
 pthread_create(&tids[i], 0, thread_hello, &tid_args[i]);
 } //create the thread, with tid, func ptr, and input args.</pre>

```
for (i=0; i < ntids; i++) {
    pthread_join(tids[i], 0);
}</pre>
```



Hello.c: main function:

pthread_create:

- function pointer argument (thread_hello)
 - name of the function that the spawned thread will start executing
 - generic function pointer type:
 void *thread main func(void *arg);
- args argument (&tid_args[i])

void *: pass a pointer to any type: int, float, struct, array, ...

pthread_join: like wait in fork-wait

- tid argument: which *pthreads* thread to wait for to exit

hello.c: thread main function



Some runs with 4 threads:

result with 4 threads should be 199999800000

```
./hello 4
count = 793900079488
./hello 4
count = 539879105421
./hello 4
count = 509829883618
./hello 4
count = 581580128846
```

Synchronization

- Synchronize: to (arrange events to) happen such that two events do not overwrite each other's work.
- Thread synchronization
 - When one thread has to wait for another
 - Events in threads that occur "at the same time"
- Uses of synchronization
 - Prevent race conditions
 - Wait for resources to become available (only one thread has access at any time - deadlocks)

Synchronization: Too Much Milk (TMM)

Time	You	Your Roommate
3.00	Arrive home	
3.05	Look in fridge, no milk	
3.10	Leave for the grocery store	
3.15		
3.20	Arrive at the grocery store	
3.25	Buy Milk	
3.30		
3.35	Arrive home, put milk in fridge	Arrive Home
3.40		Look in fridge, find milk
3.45		Cold Coffee (nom)



What mechanisms do we need for two independent threads to communicate and get a consistent view (computer state)?

How many cartons of milk can we have in this scenario? (Can we ensure this somehow?)

Time	You	Your Roommate
3.00	Arrive home	
3.05	Look in fridge, no milk	
3.10	Leave for the grocery store	
3.15		
3.20	Arrive at the grocery store	
3.25	Buy Milk	
3.30		
3.35	Arrive home, put milk in fridge	Arrive Home
3.40		Look in fridge, find milk
3.45		Cold Coffee (nom)



- A. One carton (you)
- B. Two cartons
- C. No cartons
- D. Something else

Synchronization:

Too Much Milk (TMM): One possible scenario

Time	You	Your Roommate
3.00	Arrive home	
3.05	Look in fridge, no milk	
3.10	Leave for grocery	Arrive Home
3.15		Look in fridge, no milk
3.20	Arrive at grocery	Leave for grocery
3.25	Buy Milk	
3.30		Arrive at grocery
3.35	Arrive home, put milk in fridge	
3.40		Arrive home, put milk in fridge
3.45		Oh No!



What mechanisms do we need for two independent threads to communicate and get a consistent view (computer state)?

Synchronization:

<u>Threads get scheduled in an arbitrary manner:</u> bad things may happen: ...or nothing may happen

Time	You	Your Roommate
3.00	Arrive home	
3.05	Look in fridge, no milk	
3.10	Leave for grocery	Arrive Home
3.15		Look in fridge, no milk
3.20	Arrive at grocery	Leave for grocery
3.25	Buy Milk	
3.30		Arrive at grocery
3.35	Arrive home, put milk in fridge	
3.40		Arrive home, put milk in fridge
3.45		Oh No!



What mechanisms do we need for two independent threads to communicate and get a consistent view (computer state)?

Synchronization Example



- Coordination required:
 - Which thread goes first?
 - Threads in different regions must work together to compute new value for boundary cells.
 - Threads might not run at the same speed (depends on the OS scheduler). Can't let one region get too far ahead.
 - Context switches can happen at any time!

Thread Ordering

(Why threads require care. Humans aren't good at reasoning about this.)

- As a programmer you have *no idea* when threads will run. The OS schedules them, and the schedule will vary across runs.
- It might decide to context switch from one thread to another *at any time*.
- Your code must be prepared for this!
 - Ask yourself: "Would something bad happen if we context switched here?"
- hard to debug this problem if it is not reproducible

Example: The Credit/Debit Problem

- Say you have \$1000 in your bank account
 - You deposit \$100
 - You also withdraw \$100
- How much should be in your account?
- What if your deposit and withdrawal occur at the same time, at different ATMs?
```
Thread T<sub>0</sub> Th
Credit (int a) {
    int b;
    b = ReadBalance ();
    b = b + a;
    WriteBalance (b);
}
PrintReceipt (b);
}
```

Thread T_1

```
Debit (int a) {
    int b;
```

```
b = ReadBalance ();
b = b - a;
WriteBalance (b);
```

```
PrintReceipt (b);
```





CONTEXT SWITCH



"Critical Section"



To Avoid Race Conditions



1. Identify critical sections

- 2. Use synchronization to enforce mutual exclusion
 - Only one thread active in a critical section

Critical Section and Atomicity

- Sections of code executed by multiple threads
 - Access shared variables, often making local copy
 - Places where order of execution or <u>thread interleaving will affect the</u> <u>outcome</u>
 - Follows: read + modify + write of shared variable
- Must run atomically with respect to each other
 - <u>Atomicity</u>: runs as an entire instruction or not at all. Cannot be divided into smaller parts.

Which code region is a critical section?



Which code region is a critical section? read + modify + write of shared variable



Large enough for correctness + Small enough to minimize slow down

Which values might the shared s variable hold after both threads finish?



If A runs first



B runs after A Completes



What about interleaving?



One of the threads will overwrite the other's changes.

Is there a race condition?

Suppose count is a global variable, multiple threads increment it: count++;

- A. Yes, there's a race condition (count++ is a critical section).
- B. No, there's no race condition (count++ is not a critical section).
- C. Cannot be determined

How about if compiler implements it as:

movq	(%rdx), %rax	// read count value
addq	\$1, %rax	<pre>// modify value</pre>
movq	<pre>%rax, (%rdx)</pre>	// write count

How about if compiler implements it as:

incq (%rdx) // increment value

Atomicity

- The implementation of acquiring/releasing critical section must be atomic.
 - An atomic operation is one which executes as though it could not be interrupted
 - Code that executes "all or nothing"
- How do we make them atomic?
 - Atomic instructions (e.g., test-and-set, compare-and-swap)
 - Allows us to build "semaphore" OS abstraction

Four Rules for Mutual Exclusion

- No two threads can be inside their critical sections at the same time (one of many but not more than one).
- 2. No thread outside its critical section may prevent others from entering their critical sections.
- No thread should have to wait forever to enter its critical section. (Starvation)
- 4. No assumptions can be made about speeds or number of CPU's.



Railroad Semaphore

 Help trains figure out which track to be on at any given time.



Railroad Semaphore

 Help trains figure out which track to be on at any given time.

O.S. Semaphore:

- Construct that the OS provides to processes.
- Make system calls to modify their value

mutex = 1; //lock and unlock mutex atomically.





- $\mathbf{T}_{\mathbf{0}}\colon$ Wants to execute the critical section
- $\mathbf{T}_0:$ Reads the value of mutex, Changes the value of mutex = 0 (acquires lock) Enters critical section.







Atomicity: run the entire instruction without interruption.



Atomicity: run the entire instruction without interruption.

 T_0 : Completes execution of critical section Updates mutex value = 1. (release lock)







Atomicity: run the entire instruction without interruption.

 $\mathbf{T_{1}}$: Can now acquire lock atomically and Enter the critical section

mutex = 1; //lock and unlock mutex atomically.



- Use a "mutex" semaphore initialized to 1
- Only one thread can enter critical section at a time.
- Simple, works for any number of threads

Atomicity: runs as an entire instruction or not at all.

Semaphores

- Semaphore: OS synchronization variable
 - Has integer value
 - List of waiting threads
- Works like a gate
- If sem > 0, gate is open
 - Value equals number of threads that can enter
- Else, gate is closed
 - Possibly with waiting threads



Semaphores

- Associated with each semaphore is a queue of waiting threads
- When wait() is called by a thread:
 - If semaphore is open, thread continues
 - If semaphore is closed, thread blocks on queue
- Then signal() opens the semaphore:
 - If a thread is waiting on the queue, the thread is unblocked
 - If no threads are waiting on the queue, the signal is remembered for the next thread

Semaphore Operations

```
sem s = n; // declare and initialize
wait (sem s) // Executes atomically(*)
decrement s;
if s < 0:
    block thread (and associate with s);
signal (sem s) // Executes atomically(*)
increment s;
if blocked threads:
    unblock (any) one of them;</pre>
```

(*) With help from special hardware instructions.

Semaphore Operations

```
sem s = n; // declare and initialize
wait (sem s) // Executes atomically(*)
decrement s;
if s < 0:
    block thread (and associate with s);
signal (sem s) // Executes atomically(*)
increment s;
if blocked threads:
    unblock (any) one of them;</pre>
```

Based on what you know about semaphores, should a process be able to check beforehand whether wait(s) will cause it to block?

- A. Yes, it should be able to check.
- B. No, it should not be able to check.

Semaphore Operations

sem s = n;	<pre>// declare and initialize</pre>
<pre>wait (sem s) decrement s; if s < 0: block thread (and associate with s);</pre>	<pre>// Executes atomically(*)</pre>
<pre>signal (sem s) increment s; if blocked threads: unblock (any) one of them;</pre>	<pre>// Executes atomically(*)</pre>

- No other operations allowed
- In particular, semaphore's value can't be tested!
 - No thread can tell the value of s

Synchronization: More than Mutexes

- "I want to block a thread until something specific happens."
 Condition variable: wait for a condition to be true
- "I want all my threads to sync up at the same point."

- Barrier: wait for everyone to catch up.

Barriers

- Used to coordinate threads, but also other forms of concurrent execution.
- Often found in simulations that have discrete rounds. (e.g., game of life)

Barrier Example, N Threads

```
shared barrier b;
```

```
init_barrier(&b, N);
```

```
create_threads(N, func);
```

```
void *func(void *arg) {
  while (...) {
    compute_sim_round()
    barrier_wait(&b)
 }
```



Barrier Example, N Threads

```
shared barrier b;
```

```
init_barrier(&b, N);
```

```
create_threads(N, func);
```

```
void *func(void *arg) {
  while (...) {
    compute_sim_round()
    barrier_wait(&b)
 }
```

Threads make progress computing current round at different rates.



Barrier Example, N Threads

```
shared barrier b;
```

```
init_barrier(&b, N);
```

```
create_threads(N, func);
```

```
void *func(void *arg) {
  while (...) {
    compute_sim_round()
    barrier_wait(&b)
 }
```

Threads that make it to barrier must wait for all others to get there.


Barrier Example, N Threads



Barrier Example, N Threads

```
shared barrier b;
```

```
init_barrier(&b, N);
```

```
create_threads(N, func);
```

```
void *func(void *arg) {
 while (...) {
   compute_sim_round()
   barrier_wait(&b)
}
```

Threads compute next round, wait on barrier again, repeat...

Barrier (0 waiting)

l 2

 T_3

T₀

Time

Synchronization: More than Mutexes

- I want all my threads to sync up at the same point.
 - Barrier: wait for everyone to catch up.
- I want to block a thread until something specific happens.
 - Condition variable: wait for a condition to be true
- I want my threads to share a critical section when they're reading, <u>but still safely write</u>.
 - Readers/writers lock: distinguish how lock is used

Synchronization: Beyond Mutexes Message Passing



- Operating system mechanism for IPC
 - send (destination, message_buffer)
 - receive (source, message_buffer)
- Data transfer: in to and out of kernel message buffers
- Synchronization: can't receive until message is sent

Summary

- We have no idea when OS will schedule or context switch our threads.
 - Code must be prepared, tough to reason about.
- Threads often must synchronize
 - To safely communicate / transfer data, without races
- Synchronization primitives help programmers
 - Kernel-level semaphores: limit # of threads that can do something, provides atomicity
 - User-level locks: built upon semaphore, provides mutual exclusion (usually part of thread library)