

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

C Threads

Eric C. Cooper
Richard P. Draves

June 1988

CMU-CS-88-154

Abstract

The C Threads package allows parallel programming in C under the Mach operating system. The package provides multiple threads of control within a single shared address space, mutual exclusion locks for protection of critical regions, and condition variables for thread synchronization.

This research was sponsored by the Defense Advanced Research Projects Agency (DoD), ARPA order 4864, monitored by the Space and Naval Warfare Systems Command under contract N00039-87-C-0251.

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or of the U.S. Government.

1 Introduction

Mach provides a set of low-level, language-independent primitives for manipulating threads of control [1]. The C Threads package is a run-time library that provides a C language interface to these facilities [4]. The constructs provided are similar to those found in Mesa [5] and Modula-2+ [6]: forking and joining of threads, protection of critical regions with mutex variables, and synchronization of threads by means of condition variables.

2 Naming Conventions

An attempt has been made to use a consistent style of naming for the abstractions implemented by the C Threads package. All types, macros, and functions implementing a given abstract data type are prefixed with the type name and an underscore. The name of the type itself is suffixed with `_t` and is defined via a C typedef. Documentation of the form

```
typedef struct mutex {...} *mutex_t;
```

indicates that the `mutex_t` type is defined as a pointer to a *referent type* `struct mutex` which may itself be useful to the programmer. (In cases where the referent type should be considered *opaque*, documentation such as

```
typedef ... cthread_t;
```

is used instead.)

Continuing the example of the `mutex_t` type, the functions `mutex_alloc()` and `mutex_free()` provide dynamic storage allocation and deallocation. Initialization and finalization of the referent type are accomplished with the functions `mutex_init()` and `mutex_clear()`. These are useful if the programmer wishes to include the referent type itself (rather than a pointer) in a larger structure, for more efficient storage allocation. They should not be called on objects that are dynamically allocated via `mutex_alloc()`. Type-specific functions such as `mutex_lock()` and `mutex_unlock()` are also defined, of course.

3 Initializing the C Threads Package

3.1 cthreads.h

```
#include <cthreads.h>
```

The header file `cthreads.h` defines the C Threads interface. All programs using C Threads must include this file.

3.2 `pthread_init`

```
void  
pthread_init()
```

The `pthread_init()` procedure initializes the C Threads implementation. It is called by the Mach version of the C start-up code before `main()` is entered, so the programmer does not need to call it explicitly.

4 Threads of Control

4.1 Creation

When a C program starts, it contains a single thread of control, the one executing `main()`. The thread of control is an active entity, moving from statement to statement, calling and returning from procedures. New threads are created by *fork* operations.

Forking a new thread of control is similar to calling a procedure, except that the caller does not wait for the procedure to return. Instead, the caller continues to execute in parallel with the execution of the procedure in the newly forked thread. At some later time, the caller may rendezvous with the thread and retrieve its result (if any) by means of a *join* operation, or the caller may *detach* the newly created thread to assert that no thread will ever be interested in joining it.

4.2 Termination

A thread terminates when it returns from the top-level procedure it was executing.¹ If the thread has not been detached, it remains in limbo until another thread either joins it or detaches it; if it has been detached, no such rendezvous is necessary.

4.3 `pthread_fork`

```
typedef ... any_t;  
typedef ... pthread_t;
```

The `any_t` type represents a pointer to any C type. The `pthread_t` type is an abstract “handle” that uniquely identifies a thread of control.² Values of type `pthread_t` will be referred to as thread identifiers.

```
pthread_t  
pthread_fork(func, arg)  
    any_t (*func)();  
    any_t arg;
```

¹The Mach version of the C start-up code also arranges for this to be true of the initial thread executing `main()`: a call to `pthread_exit()` occurs when `main()` returns, allowing detached threads to continue executing. The programmer may explicitly call `exit()` to terminate all threads in the program.

²Uniqueness is guaranteed only for the lifetime of the thread. When a thread exits, the implementation is free to re-use its handle.

The `pthread_fork()` procedure creates a new thread of control in which `func(arg)` is executed concurrently with the caller's thread. This is the sole means of creating new threads. Arguments larger than a pointer must be passed by reference. Similarly, multiple arguments must be simulated by passing a pointer to a structure containing several components. The call to `pthread_fork()` returns a thread identifier that can be passed to `pthread_join()` or `pthread_detach()` (see below). Every thread *must* be either joined or detached exactly once.

4.4 pthread_exit

```
void
pthread_exit(result)
    any_t result;
```

This procedure causes termination of the calling thread. An implicit `pthread_exit()` occurs when the top-level function of a thread returns, but it may also be called explicitly. The result will be passed to the thread that joins the caller, or discarded if the caller is detached.

4.5 pthread_join

```
any_t
pthread_join(t)
    pthread_t t;
```

This function suspends the caller until the thread `t` terminates. The caller receives either the result of `t`'s top-level function or the argument with which `t` explicitly called `pthread_exit()`.

4.6 pthread_detach

```
void
pthread_detach(t)
    pthread_t t;
```

The detach operation is used to indicate that the given thread will never be joined. This is usually known at the time the thread is forked, so the most efficient usage is the following:

```
pthread_detach(pthread_fork(procedure, argument));
```

A thread may, however, be detached at any time after it is forked, as long as no other attempt is made to join it or detach it.

4.7 `pthread_yield`

```
void  
pthread_yield()
```

This procedure is a hint to the scheduler, suggesting that this would be a convenient point to schedule another thread to run on the current processor. Calls to `pthread_yield()` are unnecessary in an implementation with preemptive scheduling, but may be required to avoid starvation in a coroutine-based implementation.

4.8 `pthread_self`

```
pthread_t  
pthread_self()
```

This function returns the caller's own thread identifier, which is the same value that was returned by `pthread_fork()` to the creator of the thread. The thread identifier uniquely identifies the thread, and hence may be used as a key in data structures that associate user data with individual threads. Since thread identifiers may be re-used by the underlying implementation, the programmer should be careful to clean up such associations when threads exit.

4.9 `pthread_set_data`, `pthread_data`

```
void  
pthread_set_data(t, data)  
    pthread_t t;  
    any_t data;
```

```
any_t  
pthread_data(t)  
    pthread_t t;
```

These functions allow the user to associate arbitrary data with a thread, providing a simple form of thread-specific "global" variable. More elaborate mechanisms, such as per-thread property lists or hash tables, can then be built with these primitives.

After a thread exits, any attempt to get or set its associated data is illegal, so any deallocation or other cleanup of the data must be done before the thread exits. It is always safe to access the data associated with the caller's own thread (`pthread_self()`), or with a thread that has not yet been joined or detached.

5 Synchronization

```
typedef struct mutex {...} *mutex_t;
```

```
typedef struct condition {...} *condition_t;
```

This section describes mutual exclusion and synchronization primitives, called *mutexes* and *condition variables*. In general, these primitives are used to constrain the possible interleavings of thread execution streams. They separate the two most common uses of Dijkstra's $P()$ and $V()$ operations into distinct facilities. This approach basically implements monitors [3,5] without the syntactic sugar.

Mutually exclusive access to mutable data is necessary to guarantee consistency of that data. As an example, consider concurrent attempts to update a simple counter. Suppose two threads each read the current value into a (thread-local) register, increment it, and write the value back in some order. The counter will only be incremented once, so the work of one of the threads will be lost. A mutex solves this problem (called a *race condition*) by making the read-increment-write action *atomic*—indivisible with respect to the execution of other threads. Before reading a counter, a thread locks the associated mutex. After depositing a new value, the thread unlocks the mutex. If another thread tries to use the counter while the mutex is held, its attempt to lock the mutex will not succeed until the first thread releases it. If several threads try to lock the mutex at the same time, the C Threads package guarantees that exactly one will succeed; the rest will block.

In general, mutex variables allow the programmer to protect *critical regions*—operations on shared data that must be performed indivisibly. This is sufficient to prevent competing threads from conflicting with one another, but threads do more than just compete: presumably, they also cooperate in a larger computation. In a typical multi-threaded application, one thread will rely on the results produced by another. When these results take the form of updates to a shared data structure, the problem of thread synchronization arises.

The condition variables provided by the C Threads package allow one or more threads to wait until another has finished updating a shared data structure. The definition of when an update is finished depends on the application; it typically requires the computation of some boolean-valued function of the shared data, which we will call its *status*. (The simplest example is a boolean variable whose status is simply its value.) As before, the shared data must be protected by a mutex, so that the update and the status computation exclude one another. For example, an update of a complicated linked structure might temporarily introduce cycles or invalid pointers. A status computation that traverses the structure could fail horribly if it observes such an inconsistent state.

Condition variables are used to indicate that the status of some shared data structure has changed. The association between the condition variable and this status change must be maintained entirely by the application. A thread waits for a status change by performing a `condition_wait()` operation on the associated condition variable. Whenever a thread changes the status of the shared data, it signals the associated condition to wake up any thread waiting for that status change. Unlike Hoare's original monitors [3], there is no guarantee that the awakened thread is the first to execute after the condition is signaled. As soon as the signaling thread releases the mutex, other threads may modify

the data. A waiting thread must always check the status of the shared data after being awakened, and wait again if necessary.

Special care must be taken with data structures that are dynamically allocated and deallocated. In particular, if the mutex that is controlling access to a dynamically allocated record is itself part of the record, one must be sure that no thread is waiting for the mutex before freeing the record.

Attempting to lock a mutex that one already holds is another common error. The offending thread will block waiting for itself. This can happen when a thread is traversing a complicated data structure, locking as it goes, and reaches the same data by different paths. Another instance of the problem occurs when a thread is locking pairs of elements in an array and fails to check for the special case of the elements being identical.

In general, one must be very careful to avoid *deadlock*, in which one or more threads are permanently blocked waiting for one another. The above scenarios are special cases of deadlock. The easiest way to avoid deadlock with mutexes is to impose a total ordering on them, and then ensure that they are locked in increasing order only.

The programmer must decide what granularity to use in protecting shared data with mutexes and conditions. At one extreme, one can use a single mutex protecting all shared memory, and one condition that signifies any change to that shared memory. At the other extreme, one can associate every object with its own mutex and its own condition variables, one for each possible change in the status of that object. Finer granularity normally increases the possible parallelism, because fewer threads are waiting for mutexes or conditions at any time, but it also increases the overhead due to mutual exclusion and synchronization, and increases the possibility of deadlock.

5.1 mutex_lock

```
void
mutex_lock(m)
    mutex_t m;
```

The `mutex_lock()` procedure attempts to lock the mutex `m` and blocks until it succeeds. If several threads attempt to lock the same mutex concurrently, one will succeed, and the others will block until `m` is unlocked. The case of a thread attempting to lock a mutex it has already locked is *not* treated specially; deadlock will result.

5.2 mutex_try_lock

```
int
mutex_try_lock(m)
    mutex_t m;
```

The `mutex_try_lock()` function attempts to lock the mutex `m`, like `mutex_lock()`, and returns `TRUE` if it succeeds. If `m` is already locked, however, `mutex_try_lock()` immediately returns `FALSE` rather than blocking. For example, a busy-waiting version of the `mutex_lock()` procedure could be written in terms of `mutex_try_lock()` as follows:


```

void
mutex_lock(m)
    mutex_t m;
{
    for (;;)
        if (mutex_try_lock(m))
            return;
}

```

5.3 mutex_unlock

```

void
mutex_unlock(m)
    mutex_t m;

```

The `mutex_unlock()` procedure unlocks the mutex `m`, giving other threads a chance to lock it.

5.4 condition_signal

```

void
condition_signal(c)
    condition_t c;

```

The `condition_signal()` procedure should be called when one thread wishes to indicate that the status change represented by the condition variable has occurred. If any other threads are waiting (via `condition_wait()`), then at least one of them will be awakened. If no threads are waiting, then nothing happens.

5.5 condition_broadcast

```

void
condition_broadcast(c)
    condition_t c;

```

The `condition_broadcast()` procedure is similar to `condition_signal()`, except that it awakens *all* threads waiting for the condition, not just one of them.

5.6 condition_wait

```
void
condition_wait(c, m)
    condition_t c;
    mutex_t m;
```

The `condition_wait()` procedure unlocks `m`, suspends the calling thread for some period of time, and then locks `m` again before returning. The application should guarantee that the status change associated with `c` can only occur while `m` is locked.

The caller will be awakened if `c` is signaled by another thread. Since other threads may execute between the time that `c` is signaled and the time that the caller re-acquires `m`, this operation is typically used as follows:

```
mutex_lock(m);
while (/* status of shared data is not OK */)
    condition_wait(c, m);
/* use shared data */
mutex_unlock(m);
```

6 Management of Synchronization Variables

A mutex or condition variable can be allocated dynamically from the heap, or the programmer can take an object of the referent type, initialize it appropriately, and then use its address.

6.1 Allocation

```
mutex_t
mutex_alloc()

condition_t
condition_alloc()
```

These functions provide dynamic allocation of mutex and condition objects.

6.2 Deallocation

```
void
mutex_free(m)
    mutex_t m;

void
condition_free(c)
    condition_t c;
```

These functions allow the programmer to deallocate mutex and condition objects that were allocated dynamically. Before deallocating such an object, the programmer must guarantee that no other thread will reference it. In particular, a thread blocked in `mutex_lock()` or `condition_wait()` should be viewed as referencing the object continually, so freeing the object "out from under" such a thread is erroneous, and can result in bugs that are extremely difficult to track down.

6.3 Initialization

```
void
mutex_init(m)
    struct mutex *m;

void
condition_init(c)
    struct condition *c;
```

These functions allow the programmer to initialize an object of the `struct mutex` or `struct condition` referent type, so that its address can be used wherever an object of type `mutex_t` or `condition_t` is expected. For example, the `mutex_alloc()` function could be written in terms of `mutex_init()` as follows:

```
mutex_t
mutex_alloc()
{
    mutex_t m;

    m = (mutex_t) malloc(sizeof(struct mutex));
    mutex_init(m);
    return m;
}
```

Initialization of the referent type is most often used when the programmer has included the referent type itself (rather than a pointer) in a larger structure, for more efficient storage allocation. For instance, a data structure might contain a component of type `struct mutex` to allow each instance of that structure to be locked independently. During initialization of the instance, the programmer would call `mutex_init()` on the `struct mutex` component. The alternative of using a `mutex_t` component and initializing it using `mutex_alloc()` would be less efficient.

6.4 Finalization

```
void
mutex_clear(m)
    struct mutex *m;

void
condition_clear(c)
    struct condition *c;
```

These functions allow the programmer to finalize an object of the struct `mutex` or struct condition referent type. For example, the `mutex_free()` procedure could be written in terms of `mutex_clear()` as follows:

```
void
mutex_free(m)
    mutex_t m;
{
    mutex_clear(m);
    free((char *) m);
}
```

7 Shared Variables

All global and static variables are shared among all threads: if one thread modifies such a variable, all other threads will observe the new value. In addition, a variable reachable from a pointer is shared among all threads that can dereference that pointer. This includes objects pointed to by shared variables of pointer type, as well as arguments passed by reference in `pthread_fork()`.

When pointers are shared, some care is required to avoid dangling reference problems. The programmer must ensure that the lifetime of the object pointed to is long enough to allow the other threads to dereference the pointer. Since there is no bound on the relative execution speed of threads, the simplest solution is to share pointers to global or heap-allocated objects only. If a pointer to a local variable is shared, the procedure in which that variable is defined must remain active until it can be guaranteed that the pointer will no longer be dereferenced by other threads. The synchronization functions can be used to ensure this.

Unless a subroutine library has been designed to work in the presence of multiple threads, and specifies that fact in its interface, the programmer must assume that the operations provided by the library make unprotected use of shared data. Subroutines that are documented as returning pointers to static data areas are obvious culprits, but other routines may share data privately and so be equally guilty. The programmer must therefore use a mutex that is locked before every library call (or sequence of library calls) and unlocked afterwards.

7.1 Dynamic Allocation

Dynamic allocation and freeing of user-defined data structures is typically accomplished with the standard C functions `malloc()` and `free()`. The C Threads package provides versions of these functions that work correctly in the presence of multiple threads.

8 Using the C Threads Package

All of the functions described have been implemented for the Mach multiprocessor operating system. Three implementations of threads are provided, in the form of libraries. Programs need not be recompiled to use a different thread implementation, only relinked. To compile a program that uses C Threads, the user must include the file `cthreads.h`. To link a program that uses C Threads, the user must specify on the `cc` command line one of the three libraries described below, followed by the `-lmach` library.

8.1 The Coroutine Implementation

The first implementation, `-lco_threads`, uses coroutines within a single Mach task (UNIX process). Scheduling of these threads is non-preemptive, hence `cthread_yield()` should be called within loops that do not otherwise call synchronization procedures. The programmer will typically use this version while debugging.

This implementation includes versions of the Mach interprocess communication primitives `msg_receive()`, `msg_send()`, and `msg_rpc()`, and a version of the UNIX `select()` system call, that can be called from one thread without blocking the other threads in the program. The other forms of UNIX I/O have not been redefined for use with `-lco_threads`, however. For example, one thread's call to `getchar()` may block all threads in the program. To work around this, the programmer should first call `select()` on the relevant file descriptor to guarantee that the subsequent input operation will not block.

8.2 The Mach Thread Implementation

The second implementation, `-lthreads`, uses one Mach thread per C thread. These threads are preemptively scheduled, and may execute in parallel on a multiprocessor. This is the implementation of choice for the production version of a C Threads program.

The current version of the `-lthreads` implementation affords the programmer limited control over how threads wait for mutex and condition variables.

```
extern int mutex_spin_limit;
```

This variable controls the number of iterations of busy waiting before a thread begins to yield the processor when waiting for a mutex.

```
extern int condition_spin_limit;  
extern int condition_yield_limit;
```

These variables control the number of iterations of busy waiting and processor yielding, respectively, before a thread suspends itself when waiting for a condition.

A thread suspends itself via a Mach `msg_receive()` on a per-thread synchronization port; another thread wakes it up using a Mach `msg_send()` to the suspended thread's synchronization port. Allowing each synchronization port to buffer one message eliminates the need for a wakeup-waiting indication.

8.3 The Mach Task Implementation

The third implementation, `-ltask_threads`, uses one Mach task (UNIX process) per thread, and uses the Mach virtual memory primitives to share memory between threads. In most circumstances, the `-lthreads` implementation should be used instead of this one. An exception is when the programmer wishes to use the Mach virtual memory primitives to provide a specialized pattern of memory sharing between C threads.

Users of the `-ltask_threads` implementation should note that capabilities such as Mach ports and UNIX file descriptors are private to the task that creates them, and so cannot be shared. The current `-ltask_threads` implementation also makes stack segments private to each task, so automatic (stack-allocated) variables cannot be shared.

The Mach operating system currently limits the number of tasks (and hence the number of C threads in the `-ltask_threads` implementation) that a user may create. Applications that create large numbers of threads will encounter run-time errors when they exceed this limit. It may be the case that concurrent execution is required to avoid deadlock (for example, in a multi-stage pipeline). For applications with largely independent threads, however, a limited degree of parallelism may still allow correct execution. The following function can be used in such applications.

```
void
cthread_set_limit(n)
    int n;
```

This function limits the number of active threads to `n`. If a newly created thread of control exceeds this limit, it will not begin execution until another thread terminates.

8.4 Controlling Thread Stack Sizes

All C thread stacks are the same size, which is determined by `cthread_init()` based on the value of the UNIX stack resource limit. Since resource limits are inherited, the easiest way to change the thread stack size is to use the shell's `limit` command before running the multi-threaded application, either interactively or in a "wrapper" shell script.

Thread stacks are created full-size, not grown incrementally. The implementations rely on the Mach virtual memory system to allocate physical memory only as needed by the thread, and to manage the resulting sparsely populated address space efficiently.

9 Debugging

It is strongly recommended that the coroutine-based implementation (`-lco_threads`) be used for debugging, for the following reasons:

- The order of thread context switching is repeatable in successive executions of the program, so obvious synchronization bugs may be found easily.
- Since the program is a single Mach task, existing UNIX debuggers can be used.
- The user need not worry about concurrent calls to library routines.

9.1 Low-Level Tracing

```
int pthread_debug;
```

If this variable is TRUE, diagnostic information is printed when each C Threads primitive is executed. Trace output appears on stdout.

9.2 Associating Names with C Thread Objects

```
void
pthread_set_name(t, name)
    pthread_t t;
    string_t name;
```

```
string_t
pthread_name(t)
    pthread_t t;
```

```
void
pthread_mutex_set_name(m, name)
    pthread_mutex_t m;
    string_t name;
```

```
string_t
pthread_mutex_name(m)
    pthread_mutex_t m;
```

```
void
pthread_cond_set_name(c, name)
    pthread_cond_t c;
    string_t name;
```

```
string_t
pthread_cond_name(c)
    pthread_cond_t c;
```

These functions allow the user to associate a name with a thread, mutex, or condition. The name is used when trace information is displayed (see above). The name may also be used by the programmer for application-specific diagnostics.

9.3 Pitfalls of Preemptively Scheduled Threads

The standard storage allocation functions (`malloc()` and `free()`) and most of the standard I/O functions have been modified for safe use with preemptively scheduled threads, but there are still portions of the standard C library that are not safe.³

³At CMU, one can use the command `setpath -ib /usr/mach /usr/mach/parallel` to change the appropriate search paths so that the modified versions of `stdio.h` and `libc.a` will be used by the compiler and linker.

Most of the debuggers available under Mach cannot yet be used on programs linked with `-lthreads` or `-ltask_threads`, although an enhanced version of *gdb* for Mach threads is available [2]. Furthermore, the very act of turning on tracing or adding print statements may perturb programs that incorrectly depend on thread execution speed. One technique that is useful in such cases is to vary the granularity of locking and synchronization used in the program, making sure that the program works with coarse-grained synchronization before refining it further.

10 Examples

The following example illustrates how the facilities described here can be used to program Hoare monitors [3]. The program would be compiled and linked by the command

```
cc hoaremonitor.c -lthreads -lmach
```

```
/*
 * Producer/consumer with bounded buffer.
 *
 * The producer reads characters from stdin
 * and puts them into the buffer. The consumer
 * gets characters from the buffer and writes them
 * to stdout. The two threads execute concurrently
 * except when synchronized by the buffer.
 */
#include <stdio.h>
#include <cthread.h>

typedef struct buffer {
    char *chars;      /* chars[0..size-1] */
    int size;
    int px, cx;      /* producer and consumer indices */
    int count;       /* number of unconsumed chars in buffer
 */
    mutex_t lock;
    condition_t non_empty, non_full;
} *buffer_t;

buffer_t
buffer_alloc(size)
    int size;
{
    extern char *malloc();
    buffer_t b = (buffer_t) malloc(sizeof(struct buffer));
    b->size = size;
    b->chars = malloc((unsigned) size);
    b->px = b->cx = b->count = 0;
    b->lock = mutex_alloc();
    b->non_empty = condition_alloc();
    b->non_full = condition_alloc();
    return b;
}
```

```

void
buffer_free(b)
    buffer_t b;
{
    free(b->chars);
    mutex_free(b->lock);
    condition_free(b->non_empty);
    condition_free(b->non_full);
    free((char *) b);
}

void
buffer_put(ch, b)
    char ch;
    buffer_t b;
{
    mutex_lock(b->lock);
    while (b->count == b->size)
        condition_wait(b->non_full, b->lock);
    ASSERT(0 <= b->count && b->count < b->size);
    b->chars[b->px] = ch;
    b->px = (b->px + 1) % b->size;
    b->count += 1;
    mutex_unlock(b->lock);
    condition_signal(b->non_empty);
}

char
buffer_get(b)
    buffer_t b;
{
    char ch;

    mutex_lock(b->lock);
    while (b->count == 0)
        condition_wait(b->non_empty, b->lock);
    ASSERT(0 < b->count && b->count <= b->size);
    ch = b->chars[b->cx];
    b->cx = (b->cx + 1) % b->size;
    b->count -= 1;
    mutex_unlock(b->lock);
    condition_signal(b->non_full);
    return ch;
}

```

```

void
producer(b)
    buffer_t b;
{
    int ch;

    do buffer_put((ch = getchar()), b);
    while (ch != EOF);
}

void
consumer(b)
    buffer_t b;
{
    int ch;

    while ((ch = buffer_get(b)) != EOF)
        putchar(ch);
    buffer_free(b);
}

void
main()
{
    buffer_t b = buffer_alloc(100);
    pthread_detach(pthread_fork(producer, b));
    pthread_detach(pthread_fork(consumer, b));
}

```

The following example shows how to structure a program in which a single master thread spawns a number of concurrent slaves and then waits until they all finish. The program would be compiled and linked by the command

```
cc masterslave.c -lthreads -lmach
```

```
/*
 * Master/slave program structure.
 */
#include <stdio.h>
#include <threads.h>

int count = 0;      /* number of slaves active */
mutex_t lock;      /* mutual exclusion for count */
condition_t done;  /* signaled each time a slave finishes */

/*
 * The master spawns a given number of slaves
 * and then waits for them all to finish.
 */
void
master(nslaves)
    int nslaves;
{
    int i;

    for (i = 1; i <= nslaves; i += 1) {
        void slave();
        mutex_lock(lock);
        /*
         * Fork a slave and detach it,
         * since the master never joins it individually.
         */
        count += 1;
        pthread_detach(pthread_fork(slave, random() % 1000));
        mutex_unlock(lock);
    }
    mutex_lock(lock);
    while (count != 0)
        condition_wait(done, lock);
    mutex_unlock(lock);
    printf("All %d slaves have finished.\n", nslaves);
}
```

```

/*
 * Each slave just counts up to its argument,
 * yielding the processor on each iteration.
 * When it is finished, it decrements the global count
 * and signals that it is done.
 */
void
slave(n)
    int n;
{
    int i;

    for (i = 0; i < n; i += 1)
        pthread_yield();
    mutex_lock(lock);
    count -= 1;
    printf("Slave finished %d cycles.\n", n);
    mutex_unlock(lock);
    condition_signal(done);
}

void
main()
{
    lock = mutex_alloc();
    done = condition_alloc();
    master((int) random() % 16); /* create up to 15 slaves */
}

```

References

- [1] Michael J. Accetta, Robert V. Baron, William Bolosky, David B. Golub, Richard F. Rashid, Avadis Tevanian, Jr., and Michael W. Young.
Mach: a new kernel foundation for UNIX development.
In *Proceedings of the Summer 1986 USENIX Conference*, July 1986.
- [2] David L. Black, David B. Golub, Richard F. Rashid, Avadis Tevanian, Jr., and Michael W. Young.
The Mach Exception Handling Facility.
Technical Report CMU-CS-88-129, Computer Science Department, Carnegie Mellon University, April 1988.
- [3] C. A. R. Hoare.
Monitors: an operating system structuring concept.
Communications of the ACM, 17(10):549–557, October 1974.
- [4] Brian W. Kernighan and Dennis M. Ritchie.
The C Programming Language.
Prentice-Hall, 1978.
- [5] Butler W. Lampson and David D. Redell.
Experience with processes and monitors in Mesa.
Communications of the ACM, 23(2):105–117, February 1980.
- [6] Paul Rovner, Roy Levin, and John Wick.
On Extending Modula-2 for Building Large, Integrated Systems.
Research Report 3, DEC Systems Research Center, January 1985.