Project 2: User Level Thread Library

15-410 Operating Systems

February 9, 2022

1 Overview

An important aspect of operating system design is organizing computations that run concurrently and share memory. Concurrency concerns are paramount when designing multi-threaded programs that share some critical resource, be it some device or piece of memory. In this project you will write a thread library and concurrency primitives. This document provides the background information and specification for writing the thread library and concurrency primitives.

We will provide you with a miniature operating system kernel (called "Pebbles") which implements a minimal set of system calls, and some multi-threaded programs. These programs will be linked against your thread library, stored on a "RAM disk," and then run under the supervision of the Pebbles kernel. Pebbles is documented by the companion document, "Pebbles Kernel Specification," which should probably be read *before* this one.

The thread library will be based on the `thread_fork` system call provided by Pebbles, which provides a "raw" (unprocessed) interface to kernel-scheduled threads. Your library will provide a basic but usable interface on top of this elemental thread building block, including the ability to join threads.

You will implement mutexes and condition variables based on your consideration of the options provided by the x86 instruction set—including, but not limited to, the `XCHG` instruction for atomically exchanging registers with memory or registers with registers. The lecture material discusses several atomic operations in more detail.

2 Goals

- Becoming familiar with the ways in which operating systems support user libraries by providing system calls to create processes, affect scheduling, etc.
- Becoming familiar with programs that involve a high level of concurrency and the sharing of critical resources, including the tools that are used to deal with these issues.
- Developing skills necessary to produce a substantial amount of code, such as organization and project planning.
- Working with a partner is also an important aspect of this project. You will be working with a partner on subsequent projects, so it is important to be familiar with scheduling time to work, a preferred working environment, and developing a good group dynamic before beginning larger projects.
- Coming to understand the dynamics of source control in a group context, e.g., when to branch and merge.

The partner goal is an important one–this project gives you an opportunity to debug not only your code deliverables but also your relationship with your partner. You may find that some of the same techniques apply.

3 Important Dates

Wednesday, February 9th Project 2 begins

Friday, February 11th You should be able to draw a *very detailed* picture of the parent and child stacks during `thr_create()` at the point when the child does its first `PUSHL` instruction. In your design multiple pictures may be equally plausible, but it is important that you be able to draw at least one case in detail. It is wise for each partner to independently draw this picture before you compare notes and agree on all the details. It is not wise to skip this step (unless you have previously written a thread library).

Monday, February 14th You should have thread creation working well enough to pass the STARTLE test we provide. In particular, you should be able to run:

```
[410-shell]$ misbehave_wrap 0 startle
```

Wednesday, February 16th You should have thread creation, mutexes, and condition variables working well.

Friday, February 18th If you haven't least begun debugging CYCLONE and AGILITY_DRILL by this point, you run the risk of turning in a thread library with serious structural flaws and lacking concurrency debugging experience useful for the kernel project.

Wednesday, February 23rd Project 2 due at 23:59:59

4 Overview of Deliverables

The library you will write will contain:

- System-call stub routines (see the "Pebbles Kernel Specification" document)
- A software exception handler which implements automatic stack growth for legacy programs
- A software exception handler which handles thread crashes for multi-threaded programs
- Thread management calls
- Mutexes and condition variables
- Semaphores
- Readers/writers locks

Unlike system call stubs (see the "Pebbles Kernel Specification" document), thread library routines need not be one-per-source-file, but we expect you to use good judgment when partitioning them (and this may influence your grade to some extent). You should arrange that the Makefile infrastructure you are given will build your library into `libsyscall.a` and `libthread.a` according to the directions found in the `README` and `config.mk` files found in the tarball.

4.1 Automatic stack growth for legacy single-threaded applications

If you examine the `_main()` "main wrapper" code in `crt0.c`, you will find that it receives two "extra" parameters from the kernel's program loader, called `stack_high` and `stack_low`, and passes them to a function called `install_autostack()` before `main()` runs. These two parameters are the highest and lowest virtual address of the initial stack space that the kernel created before launching the program (when `install_autostack()` is called, `%ESP` will naturally be some value between `stack_high` and `stack_low`). You are expected to register a `swexn()` exception handler of your own devising which resolves appropriate page-fault exceptions by performing automatic stack growth according to the venerable Unix tradition (automatic stack growth is discussed further in the "Pebbles Kernel Specification" document). Your handler is not expected to "correct" or otherwise resolve any exceptions other than page faults, and is also not expected to "correct" page faults which are unrelated to automatic stack growth. We have provided a simple test program, `stack_test1`, which you can use to exercise your stack-growth handler.

4.2 Thread-crash handling for multi-threaded applications

If one thread in a multi-threaded application experiences a fatal exception, the application as a whole is unlikely to continue in a useful fashion. The crashed thread may well have been holding locks, may have temporarily broken data-structure invariants, or may have been working to produce a result that other threads will later need. This suggests that a thread library should provide a sensible default behavior which reacts to the untimely death of a thread.

In the other direction, note that automatic stack growth is not a typical feature in multi-threaded environments. There are multiple stack regions, which must be placed within the address space by the thread library. However, only the application knows how many thread stacks it will create and how large they will need to be. For that reason, it is typical for threaded programs to pre-declare the stack-size needs of each thread (see `thr_init()` below) and for thread libraries not to support growing thread stacks once they are created.

This means that threads in a multi-threaded application should probably handle exceptions according to a different plan than that used by the initial thread in a single-threaded legacy application. As a result, you should give careful consideration to how `thr_init()` should transition an application from the legacy single-threaded environment which was set up before `main()` to a multi-threaded environment. What guarantees, if any, that obtained before `thr_init()` was called should remain valid, and how should those guarantees be implemented? What code should implement which policies for which threads?

Threads can also "crash voluntarily" in the absence of a hardware exception. Various library routines and language facilities (e.g., `assert()` and `affirm()`) rely on a function called "panic": `void panic(const char *format, ...)`. We have provided you with a skeletal (i.e., insufficient) implementation of `panic()`, which you should enhance appropriately.

5 Thread Library API

You may assume that programs which use condition variables will include `cond.h`, programs which use semaphores will include `sem.h`, etc.

Please note that all lock-like objects are defined to be "unlocked" when created.

5.1 Return values

You will note that many of the thread-library primitives (e.g., `mutex_unlock()`) are declared as returning void. This is because there are some operations that can't meaningfully "return an error code." Consider what would happen if a program tried to invoke `exit()` and `exit()` "failed." What could the program do? Or consider the `free()` storage-allocator function. If a program called `free()` on an address, and `free()` said "No," what would that mean? Should the program continue using the memory area or not? Could it reasonably expect the next call to `malloc()` to work?

"Returning an error" is sensible when an operation might reasonably fail in plausible, non-emergency circumstances and where higher-level code can do something sensible to recover, or at least has a reasonable chance to explain to a higher authority what went wrong. If an operation *cannot* fail in reasonable circumstances (i.e., a failure means the computational state is irrevocably broken) and there is no reasonable way for higher-level code to do anything reasonable, other approaches are required, and void functions may be reasonable.

Note well that the author of a void function bears the responsibility of designing the implementation in such a way that the code fails *only* in "impossible" situations. This may require the author to design other parts of the code to take on extra responsibilites so the "must work reliably" functions are indeed reliable.

Note further that a void return type is a contractual specification that when the function returns the documented action will have been completed successfully. Said another way, if some circumstance prevents a void function from acting as specified, it cannot return.

Some of the thread-library interface functions below are declared as void functions. In each case, you will need to list possible failure cases and think through them. The function will need to work in all "might reasonably happen" situations and do something reasonable if it discovers that the computation is irretrievably broken. You will generally need to consider and trade off the cost of checking for a particular bad situation against how bad it would be to leave the situation undetected. For more guidance, refer to the "Errors" lecture.

5.2 "Illegal"

At several points in this document you will encounter language similar to this: "It is illegal for an application to ..." The meaning of "illegal" involves both subtlety and judgment. Consider the meaning and implications of a typical real-world illegality situation: it is illegal to litter in a park. That probably means:

- There is general agreement that reasonable people will not litter in the park,
- Any single instance of a person littering in the park may result in a penalty,

- The park management invests some resources in detecting, apprehending, and penalizing littering, and
- The park management also invests some resources in making it easy for visitors to dispose of rubbish *without* littering.

It probably does *not* mean:

- Littering in the park never happens,
- Every single instance of a person littering in the park results in a lifetime prison sentence for the person, or
- The park management focuses so much on litter prevention that all the plants die and the play equipment decays into unusability.

For more guidance, refer to the "Questions" lecture and the "Errors" lecture. Also see Section 8.1.

5.3 Thread Management API

- `int thr_init( unsigned int size )` - This function is responsible for initializing the thread library. The argument `size` specifies the amount of stack space which will be available for each thread using the thread library.

  This function returns zero on success, and a negative number on error.

  The thread library can assume that programs using it are well-behaved in the sense that they will call `thr_init()`, exactly once, before calling any other thread library function (including memory allocation functions in the `malloc()` family, described below) or invoking the `thread_fork` system call. Also, you may assume that all threads of a task using your thread library will call `thr_exit()` instead of directly invoking the `vanish()` system call (and that the root thread will call `thr_exit()` instead of `return()`'ing from `main()`).

- `int thr_create( void *(*func)(void *), void *arg )` - This function creates a new thread to run `func(arg)`. This function should allocate a stack for the new thread and then invoke the `thread_fork` system call in an appropriate way. A stack frame should be created for the child so that the indicated thread-body function is run appropriately. On success the thread ID of the new thread is returned, on error a negative number is returned.

  You should pay attention to (at least) two stack-related issues. First, the stack pointer should essentially always be aligned on a 32-bit boundary (i.e., %esp mod 4 == 0). Second, you need to think very carefully about the relationship of a new thread to the stack of the parent thread, especially right after the `thread_fork` system call has completed.

- `int thr_join( int tid, void **statusp )` -

  This function "cleans up" after a thread, optionally returning the status information provided by the thread at the time of exit.

  The target thread `tid` may or may not have exited before `thr_join()` is called; if it has not, the calling thread will be suspended until the target thread does exit.

If `statusp` is not NULL, the value passed to `thr_exit()` by the joined thread will be placed in the location referenced by `statusp`.

Only one thread may join on any given target thread. Other attempts to join on the same thread should return an error promptly. If thread `tid` was not created before `thr_join(tid)` was called, an error will be returned.

This function returns zero on success, and a negative number on error.

- `void thr_exit( void *status )` - This function exits the thread with exit status `status`. If a thread other than the root thread returns from its body function instead of calling `thr_exit()`, the behavior should be the same as if the function had called `thr_exit()` specifying the return value from the thread's body function.

  Note that `status` **is not a "pointer to a void."** It is frequently not a pointer to anything of any kind. Instead, `status` is a pointer-sized opaque data type which the thread library transports uninterpreted from the caller of `thr_exit()` to the caller of `thr_join()`.

- `int thr_getid( void )` - Returns the thread ID of the currently running thread.

- `int thr_yield( int tid )` - Defers execution of the invoking thread to a later time in favor of the thread with ID `tid`. If `tid` is -1, yield to some unspecified thread. If the thread with ID `tid` is not runnable, or doesn't exist, then an integer error code less than zero is returned. Zero is returned on success.

Note that the "thread IDs" generated and accepted by your thread library routines (e.g., `thr_getid()`, `thr_join()`) are not required to be the same "thread IDs" which are generated and accepted by the thread-related system calls (e.g., `thread_fork`, `gettid()`, `make_runnable()`). If you think about how you would implement an "M:N" thread library,[1] or a user-space thread library, you will see why these two name spaces cannot always be the same. Whether or not you use kernel-issued thread ID's as your thread library's thread ID's is a design decision you will need to consider.

However, you **must not** aggressively recycle thread ID's, as this significantly reduces the utility of, e.g., `thr_yield()`.

5.4 Mutexes

Mutual exclusion locks prevent multiple threads from simultaneously executing *brief* critical sections of code. To implement mutexes you may use the `XCHG` instruction documented on page 3-714 of the Intel Instruction Set Reference. For more information on the behavior of mutexes, feel free to refer to the text, or to the Solaris or Linux `pthread_mutex_init()` manual page.

- `int mutex_init( mutex_t *mp )` - This function should initialize the mutex pointed to by `mp`. It is illegal for an application to use a mutex before it has been initialized or to initialize one when it is already initialized and in use. This function returns zero on success, and a negative number on error.

[1] You probably aren't.

- `void mutex_destroy( mutex_t *mp )` - This function should "deactivate" the mutex pointed to by `mp`. It is illegal for an application to use a mutex after it has been destroyed (unless and until it is later re-initialized). It is illegal for an application to attempt to destroy a mutex while it is locked or threads are trying to acquire it.

- `void mutex_lock( mutex_t *mp )` - A call to this function ensures mutual exclusion in the region between itself and a call to `mutex_unlock()`. A thread calling this function while another thread is in an interfering critical section must not proceed until it is able to claim the lock.

- `void mutex_unlock( mutex_t *mp )` - Signals the end of a region of mutual exclusion. The calling thread gives up its claim to the lock. It is illegal for an application to unlock a mutex that is not locked.

For the purposes of this assignment, you may assume that a mutex should be unlocked only by the thread that most recently locked it.[2]

5.5 Condition Variables

Condition variables are used for waiting, for a while, for mutex-protected state to be modified by some other thread(s). A condition variable allows a thread to voluntarily relinquish the CPU so that other threads may make changes to the shared state, and then tell the waiting thread that they have done so. If there is some shared resource, threads may de-schedule themselves and be awakened by whichever thread was using that resource when that thread is finished with it. In implementing condition variables, you may use your mutexes, and the system calls `deschedule()` and `make_runnable()`. For more information on the behaviour of condition variables, you may refer to the Solaris or Linux documentation on `pthread_cond_wait()`.

- `int cond_init( cond_t *cv )` - This function should initialize the condition variable pointed to by `cv`. It is illegal for an application to use a condition variable before it has been initialized or to initialize one when it is already initialized and in use. This function returns zero on success, and a negative number on error.

- `void cond_destroy( cond_t *cv )` - This function should "deactivate" the condition variable pointed to by `cv`. It is illegal for an application to use a condition variable after it has been destroyed (unless and until it is later re-initialized). It is illegal for an application to invoke `cond_destroy()` on a condition variable while threads are blocked waiting on it.

- `void cond_wait( cond_t *cv, mutex_t *mp )` - The condition-wait function allows a thread to wait for a condition and release the associated mutex that it needs to hold to check that condition. The calling thread blocks, waiting to be signaled. The blocked thread may be awakened by a `cond_signal()` or a `cond_broadcast()`. Upon return from `cond_wait()`, `*mp` has been re-acquired on behalf of the calling thread.

- `void cond_signal( cond_t *cv )` - This function should wake up a thread waiting on the condition variable pointed to by `cv`, if one exists.

[2] Opinions differ, but you might want to wait until after the scheduling lecture(s) before solidifying yours.

- `void cond_broadcast( cond_t *cv )` - This function should wake up all threads waiting on the condition variable pointed to by `cv`.

  Note that `cond_broadcast()` should *not* awaken threads which may invoke `cond_wait(cv)` "after" this call to `cond_broadcast()` has begun execution.[3]

When designing your condition-variable implementation, your first priority should be a solution that correctly and efficiently solves the thread-blocking/thread-awakening problem in an "effectively atomic" or thread-safe fashion, via a mixture of careful design on the one hand and experience with test code on the other hand (at least our test code; ideally some of yours as well).

Once you have achieved that, if time permits, we encourage you to consider what happens if a synchronization primitive other than your condition variables uses the same synchronization system calls. For example, will your condition-variable implementation work correctly if a thread sometimes blocks on a condition variable and at other times blocks on a barrier object, something similar to `java.util.concurrent.CountDownLatch`, implemented in terms of the same system calls? This is a tricky design problem[4] which should be solved for full credit, but we explicitly advise you to attempt it only after getting a condition-variable solution completely done and passing all Project 2 tests.

5.6 Semaphores

As discussed in class, semaphores are a higher-level construct than mutexes and condition variables. Implementing semaphores on top of mutexes and condition variables should be a straightforward but hopefully illuminating experience.

- `int sem_init( sem_t *sem, int count )` - This function should initialize the semaphore pointed to by `sem` to the value `count`. It is illegal for an application to use a semaphore before it has been initialized or to initialize one when it is already initialized and in use. This function returns zero on success and a number less than zero on error.

- `void sem_destroy( sem_t *sem )` - This function should "deactivate" the semaphore pointed to by `sem`. It is illegal for an application to use a semaphore after it has been destroyed (unless and until it is later re-initialized). It is illegal for an application to invoke `sem_destroy()` on a semaphore while threads are waiting on it.

- `void sem_wait( sem_t *sem )` - The semaphore wait function allows a thread to decrement a semaphore value, and may cause it to block indefinitely until it is legal to perform the decrement.

- `void sem_signal( sem_t *sem )` - This function should wake up a thread waiting on the semaphore pointed to by `sem`, if one exists, and should update the semaphore value regardless.

[3] If that sounds a little fuzzy to you, you're right–but if you think about it a bit longer it should make sense.

[4] The ideal situation is for your condition-variable implementation to satisfy the language found in 47 CFR § 15.19(a)(3), namely: "This device complies with part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) This device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation."

5.7 Readers/writers locks

Readers/writers locks allow multiple threads to have "read" access to some object simultaneously. They enforce the requirement that if any thread has "write" access to an object, no other thread may have either kind of access ("read" or "write") to the object at the same time. These types of locking behaviors are often called "shared" (for readers) and "exclusive" (for writers) locks.

The generic version of this problem is called the "readers/writers problem." Two standard formulations of the readers/writers problem exist, called unimaginatively the "first" and "second" readers/writers problems. In the "first" readers/writers problem, no reader will be forced to wait unless a writer has already obtained an exclusive lock. In the "second" readers/writers problem, no new reader can acquire a shared lock if a writer is waiting. You should think through the reasons that these formulations allow starvation of different access types; starvation of writers in the case of the "first" readers/writers problem and starvation of readers in the case of the "second" readers/writers problem.

In addition to a correct implementation of shared and exclusive locking, we expect you to implement a solution that is "at least as good as" a solution to the "second" readers/writers problem. That is, your solution should not allow starvation of writers. Your solution need not strictly follow either of the above formulations: it is possible to build a solution which does not starve any client. No matter what you choose to implement, you should explain what, how, and why.

You may choose which underlying primitives (e.g., mutex/cvar or semaphore) you use to implement readers/writers locks. Once again, you should explain the reasoning behind your choice.

- `int rwlock_init( rwlock_t *rwlock )` - This function should initialize the lock pointed to by `rwlock`. It is illegal for an application to use a readers/writers lock before it has been initialized or to initialize one when it is already initialized and in use. This function returns zero on success and a number less than zero on error.

- `void rwlock_destroy( rwlock_t *rwlock )` - This function should "deactivate" the lock pointed to by `rwlock`.

  It is illegal for an application to use a readers/writers lock after it has been destroyed (unless and until it is later re-initialized). It is illegal for an application to invoke `rwlock_destroy()` on a lock while the lock is held or while threads are waiting on it.

- `void rwlock_lock( rwlock_t *rwlock, int type )` - The `type` parameter is required to be either `RWLOCK_READ` (for a shared lock) or `RWLOCK_WRITE` (for an exclusive lock). This function blocks the calling thread until it has been granted the requested form of access.

- `void rwlock_unlock( rwlock_t *rwlock )` - This function indicates that the calling thread is done using the locked state in whichever mode it was granted access for. Whether a call to this function does or does not result in a thread being awakened depends on the situation and the policy you chose to implement.

  It is illegal for an application to unlock a readers/writers lock that is not locked.

- `void rwlock_downgrade( rwlock_t *rwlock )` - A thread may call this function only if it already holds the lock in `RWLOCK_WRITE` mode at a time when it no longer requires exclusive

access to the protected resource. When the function returns: no threads hold the lock in `RWLOCK_WRITE` mode; the invoking thread, and possibly some other threads, hold the lock in `RWLOCK_READ` mode; previously blocked or newly arriving writers must still wait for the lock to be released entirely. During the transition from `RWLOCK_WRITE` mode to `RWLOCK_READ` mode the lock should at no time be unlocked. This call should not block indefinitely.[5]

Note: as readers/writers locks are a "classic problem" (and widely used in systems-related code), the Internet is full of solutions (good and bad) to various versions of the problem. Please recall that this is a design class, not a copy-and-paste class. We believe it is feasible and very educational for you to design readers/writers locks yourself, "from scratch." If you consult external sources for inspiration, you must do so in compliance with the terms of the syllabus (which you are required to read). Furthermore, the score we assign your readers/writers submission will depend on the insight you demonstrate. A zero-insight-added copy-and-paste from an online source, especially one which does not fulfill the requirements of our version of the problem, is likely to receive a zero (or even negative) score. Even worse, turning in code you don't understand is *very* unlikely to be a good way to prepare for exams.

Note: We *will not grade* your readers/writers implementation unless your thread library passes a specified series of tests; see Section 12.

5.8 Safety & Concurrency

Please keep in mind that much of the code for this project must be thread safe. In particular the thread library itself should be thread safe. However, by its nature a thread library must also be concurrent. In other words, you may *not* solve the thread-safety problem with a hammer, such as using a global lock to ensure that only one thread at a time can be running thread library code. In general, it should be possible for many threads to be running each library interface function "at the same time."

As you design your library, your model should be that some system calls "take a while to run." You should try to avoid situations where "too many" threads are waiting "too long" because of this. This paragraph provides a design hint, not implementation rules: acting on it will require you to think about system calls and the meanings of "too many" and "too long."

5.9 Distribution Files

To begin working on the project, fetch and unpack the tarball posted on the course web page. Please read the README included therein.

6 Documentation

For each project in 15-410, functions and structures should be documented using doxygen. Doxygen uses syntax similar to Javadoc. The Doxygen documentation can be found on the course website. The provided Makefile has a target called `html_doc` that will invoke doxygen on the source files listed in the Makefile.

[5] We do not ask you to implement this function's partner, "`rwlock_upgrade()`"—and for good reason! See if you can figure out why.

While Doxygen allows functions to be commented either in header files or with the source code, it is the observation of the course staff that Doxygen comments that are with the source code are more likely to be up-to-date with the code.

7 Thread Group Library

A commonly used program paradigm involves one or more manager threads overseeing the completion of a large task which has been split into parts assigned to a pool of worker threads. Examples of this model include databases, Apache, and Firefox. Once a worker thread has completed its job, it exits; manager threads dispatch new worker threads based on system load, new requests, and the results obtained by previous worker threads. In this environment it is not convenient for a manager to know which particular worker thread it should next call `thr_join()` on; instead it is convenient to wait until the next thread in the worker pool completes.

We have provided you with a simple library implementing "thread groups." This library essentially provides an abstraction layer above the thread library you will write–a compliant program will use `thrgrp_create()` and `thrgrp_join()` instead of calling `thr_create()` and `thr_join()` directly.

These functions and their requisite data structures are defined in 410user/libthrgrp/thrgrp.c and 410user/libthrgrp/thrgrp.h.

- `thrgrp_group_t`
  A structure representing a thread group.
- `thrgrp_init_group(thrgrp_group_t *tg)`
  This function initializes a thread group. It must be called before the thread group is used for anything. Returns 0 on success, non-zero otherwise.
- `thrgrp_destroy_group(thrgrp_group_t *tg)`
  This function destroys a thread group, cleaning up all of its memory. This should be called if a thread group isn't to be used further. The effects of using a thread group after it has been destroyed are be undefined. Returns 0 on success, non-zero otherwise.
- `thrgrp_create(thrgrp_group_t *tg, void *(*func)(void *), void *arg)`
  This function spawns a new thread (analogous to `thr_create()`) in the threadgroup `tg`. The spawned thread must not call `thr_exit()`. Instead, `func()` should return an exit code (of type `void *`) which will be made available to a manager thread.

  Returns 0 on success, non-zero otherwise.
- `thrgrp_join(thrgrp_group_t *tg, void **statusp)`
  If there are any unreaped threads in the thread group `tg` then it will reap one of them, setting `*statusp` appropriately, and return. If there are no unreaped threads in the group, it will block until one does exit, reap it, and return.

8 The C Library

This is simply a list of the most common library functions that are provided. For details on using these functions please see the appropriate `man` pages.

Other functions are provided that are not listed here. Please see the appropriate header files for a full listing of the provided functions.

Some functions typically found in a C I/O library are provided by `410user/libstdio.a`. The header file for these functions is `410user/libstdio/stdio.h`, aka `#include <stdio.h>`.

- `int putchar(int c)`
- `int puts(const char *str)`
- `int printf(const char *format, ...)`
- `int sprintf(char *dest, const char *format, ...)`
- `int snprintf(char *dest, int size, const char *formant, ...)`
- `int sscanf(const char *str, const char *format, ...)`
- `void lprintf( const char *format, ...)`

Note that `lprintf()` is the user-space analog of the `lprintf_kern()` you used in Project 1.

Some functions typically found in various places in a standard C library are provided by `410user/libstdlib.a`. The header files for these functions are `stdlib.h`, `assert.h`, and `ctype.h`.

- `int atoi(const char *str)`
- `long atol(const char *str)`
- `long strtol(const char *in, const char **out, int base)`
- `unsigned long strtoul(const char *in, const char **out, int base)`
- `void assert(int expression)`

We are providing you with *non-thread-safe versions* of the standard C library memory allocation routines. You are *required* to provide a thread-safe wrapper routine with the appropriate name (remove the underscore character) for each provided routine. These should be genuine wrappers, i.e., do *not* copy and modify the source code for the provided routines.

- `void *_malloc(size_t size)`
- `void *_calloc(size_t nelt, size_t eltsize)`
- `void *_realloc(void *buf, size_t new_size)`
- `void _free(void *buf)`

You may assume that no calls to functions in the "malloc() family" will be made before the call to `thr_init()`.

These functions will typically seek to allocate memory regions from the kernel which start at the top of the data segment and proceed to grow upward. You will thus need to plan your use of the available address space with some care.

Some functions typically found in a C string library are provided by `410user/libstring.a`. The header file for these functions is `410user/libstring/string.h`.

- `int strlen(const char *s)`
- `char *strcpy(char *dest, char *src)`
- `char *strncpy(char *dest, char *src, int n)`
- `char *strdup(const char *s)`
- `char *strcat(char *dest, const char *src)`
- `char *strncat(char *dest, const char *src, int n)`
- `int strcmp(const char *a, const char *b)`
- `int strncmp(const char *a, const char *b, int n)`
- `void *memmove(void *to, const void *from, unsigned int n)`
- `void *memset(void *to, int ch, unsigned int n)`
- `void *memcpy(void *to, const void *from, unsigned int n)`

8.1 Assertions

8.1.1 `assert()`

Our C library includes the well-known `assert()` macro, which—sometimes—verifies that a specified expression evaluates to true and crashes the program if instead it's false. However, in compliance with relevant standards, the `assert()` macro we provide *does nothing* if the `NDEBUG` preprocessor symbol is defined (see `assert.h`). As a convenience, you can adjust the value of `NDEBUG` via the `CONFIG_NDEBUG` directive in `config.mk`, without needing to make changes to your source code.

8.1.2 `contracts.h`

Our C library also includes the "15-122 `contracts.h` macros" (`ASSERT()`, `REQUIRES()`, and `ENSURES()`). According to the practice model of 15-122, the predicates checked by those macros may be so computationally expensive as to preclude their evaluation during normal use of the code base—for example, a module-internal consistency predicate might take $O(N^2)$ time to check invariants of a data structure whose operations take $O(log(N))$ time. Thus the `contracts.h` checks are expected to be active only while *debugging* your code, as denoted by the definition of the `DEBUG` preprocessor symbol (see `contracts.h`). As a convenience, you can adjust the value of `DEBUG` via the `CONFIG_DEBUG` directive in `config.mk`, without needing to make changes to your source code.

8.1.3 "Production code"

Some students arrive in this class with the correct understanding that low-level system code has debugging builds and production builds, but also with the—vastly mistaken—belief that production builds have all error checking disabled. However, "production" use of your code

should include *some* error checking, especially at module boundaries, even if `assert()` and/or the `contracts.h` macros are unused or disabled.

To that end, we have extended `assert.h` with four macros that *always* evaluate the specified expressions, regardless of "debug settings," and tastefully invoke `panic()` as required.

Examples:

- `affirm(3 > 0);`
- `affirm_msg(3 > 0, "arithmetic broken around %d and %d", 3, 0);`
- `reject(0 > 3);`
- `reject_msg(0 > 3, "arithmetic broken around %d and %d", 0, 3);`

In Project 2 you will be reading, and optionally writing, user-space application code, and you will be writing hopefully-robust user-space library code. We expect you to apply the "Errors" lecture, the "Questions" lecture, Section 5.1, and Section 5.2 while selecting between/among `assert()`, the `contracts.h` macros, `affirm()`/`reject()`, and `panic()`.

9 Debugging Support Code

The same `MAGIC_BREAK` macro which you used in Project 1 is also available to user code in Project 2 if you `#include` the `410user/libsimics/simics.h` header file.

The function call `lprintf()` may be used to output debugging messages from user programs. Its prototype is in `410user/libsimics/simics.h`.

Also, user code can be symbolically debugged using the Simics symbolic debugger. **If you restrict yourself to debugging with printf() it may cost you significant amounts of time.**

10 Build Options

As was the case in Project 1, you may select among compilation alternatives by editing the `CC` directive in `config.mk`. As in Project 1, using more compiler variants may result in finding more bugs.

Also in `config.mk`, you can use `CONFIG_DEBUG` to enable/disable the `contracts.h` macros and `CONFIG_NDEBUG` to enable/disable `assert()`.

From time to time, the course staff may release alternate reference-kernel binaries, which will be found in `410kern`. If we release a reference kernel called `xxx`, the binary will be `410kern/kernel_xxx.o` and you can build your thread library against that kernel by placing a line of the form `REFK=xxx` in your `config.mk` file. When we grade your submission we will disable your `REFK` selection if you have made one.

10.1 "SMPathos" reference kernel

At times students wish to investigate how their thread-library code runs on a multi-processor machine. Two options are available.

1. In the hallway outside Professor Eckhardt's office there is a "crash box" upon which you can run your code, documented on the Projects web page. The crash box has two processors, so in theory it is able to perform amazing feats of multithreading (limited sometimes in practice by the fact that the graphics hardware is pretty slow).

2. A special configuration of Simics, `simics46smpathos`, can run your code on a machine with up to eight simulated cores. Debugging may be a little more complicated. You can use `pselect cpu2` to switch the debugger's attention to CPU 2, and you can also prefix many commands with a CPU name, e.g., `cpu1.bt` to examine the stack on CPU 1. You can increase the number of processors above the default of four by setting an environment variable, `SIMICS_NUM_CPUS`, to a number between five and eight.

Running your thread library on a multi-processor machine won't result in different behavior unless the kernel sitting between your thread library and the hardware launches all of the processors and simultaneously schedules threads onto multiple processors. Dr. Michael Sullivan, a former OS TA and CMU CSD alumnus who is a world expert on 15-410 reference kernels, has graciously provided an SMP version of the P2 reference kernel.

You should spend zero time experimenting with running your thread library on multiple processors unless your P2 is essentially done. If you have some spare time for an experiment, you can try it out by:

1. Adding `REFK=smpathos` to your config.mk and running `make`

2. Booting the resulting `bootfd.img` on the crash box or via `simics46smpathos`.

Please do *not* try this before your thread library is solid (as far as you can tell): debugging is easier with a single-processor kernel, and if you run into trouble with SMPathos it is actually possible that the trouble will be due to a bug in it rather than a bug in your thread library.

If you do have time to try SMPathos, we will be thrilled by reports of the form "We thought our thread library was 100% done, but as soon as we tried SMPathos our code blew up and we found a subtle race condition in our lock-free transactional concurrent green-purple splay treap."

But if you send us a report like "With regular Pathos, `agility_drill` was crashing, but when we tried SMPathos it hangs instead—why would that be?" we just won't know; honestly, you probably shouldn't have spent the time to try SMPathos, because you almost certainly should have spent that time debugging the `agility_drill` crash you already had in hand.

11 Submission Inventory

Implement the functions for automatic stack growth, system-call stub routines, the thread library, and concurrency tools conforming to the documented APIs. Hand in all source files that you generate, in ready-to-build form, in accordance with the directions found on the hand-in web page. Be sure to provide a design description in README.dox, including an overview of existing issues and any interesting design decisions you made. **Any use of, or reliance on, outside code must be in accordance with course policy as stated in the syllabus.**

12 Grading Criteria

You will be graded on the completeness and correctness of your project. A complete project is composed of a reasonable attempt at each function in the API. Also, a complete project follows the prescribed build process, and is well documented. A correct project implements the provided specification. Also, code using the API provided by a correct project will not be killed by the kernel, and will not suffer from inconsistencies due to concurrency errors in the library. Please note that there exist concurrency errors that even carefully-written test cases may not expose. Read and think through your code carefully. Do not forget to consider pathological cases.

The most important parts of the assignment to complete are the thread management, mutex, and condition variable calls. These should be well-designed, solidly implemented, and thoroughly tested with `misbehave()` (see below). It is probably unwise to devote substantial coding effort to the other parts of the library before the core is reliable. In particular, we **will not grade** readers/writers implementations for Project 2 submissions which do not pass the "hurdle" subset of the test suite (see the project web page for details).

Because a thread library is designed to support the activities of multiple threads, concurrency (the ability to get more than one thing done at a time) is important. As you do your design, ask yourself which things your thread library can accomplish in parallel. You may tune your code, especially your locking code, based on the assumption that it will run on a uni-processor machine. However, if we were to run your kernel on a multi-processor machine instead, the core of the thread library shouldn't artificially limit concurrency–if there were four processors, threads running on top of your thread library should be able to use all four productively most of the time. Another way to think about this is that on a uni-processor machine the timer and the kernel scheduler determine when thread execution is interleaved; if the structure of your thread library results in much less concurrency (interleaving) than the kernel provides, something isn't right.

Because a thread library is core software for an application, it should behave responsibly with respect to error conditions and should not leak or unwisely over-consume resources.

Finally, code that is robust doesn't randomly refuse to perform its job. It is not really robust for `mutex_lock()` to refuse to lock something because it can't allocate memory, and it is *downright unreasonable* for `cond_wait()` to refuse to block a thread because of a memory-allocation problem: what's the caller supposed to do—keep running? These and similar operations should do their jobs in a prompt and reliable manner.

Important note: if you don't carefully apply the material found in the code-quality and thread-synchronizations lectures, it is possible to **lose up to two letter grades** on this assignment, even if your thread library passes all of the tests we provide. Code that "runs ok" in the expected/common case may completely fail to meet the robustness standards of the class.

12.1 "Spin Locks"

Try to avoid tricking yourself via use of the term "spin lock." It is probably unwise to believe that:

- any locking code which contains a loop is a "spin lock", and
- of course any lock needs to have at least one loop, and
- people excitedly talk about "spin locks" in the hallways,

...so therefore any locking code you write is a "spin lock" and thus fine by definition.

Whether your locking code is looked on favorably or not will depend on what it does, not what you call it—except that if you call code something it genuinely isn't your grader may find a way to take a minor deduction.

13 Debugging

13.1 Requests for Help

Please do not ask for help from the course staff with a message like this:

> The kernel is killing my threads! Why?

or

> Why is my program stuck in `malloc()`?

An important part of this class is developing your debugging skills. In other words, when you complete this class you should be able to debug problems which you previously would not have been able to handle.

Thus, when faced with a problem, you need to invest some time in figuring out a way to characterize it and close in on it so you can observe it in the actual act of destruction. Your reflex when running into a strange new problem should be to start thinking, not to start off by asking for help.

Having said that, if a reasonable amount of time has been spent trying to solve a problem and no progress has been made, do not hesitate to ask a question. But please be prepared with a list of details and an explanation of what you have tried and ruled out so far.

13.2 Warning

When making design decisions, beware of basing them on how long it takes Simics to execute your code. Because Simics is a simulator, it inherently runs some of your code much slower than a real machine would, but in the other direction it employs various heuristics to accelerate other parts of your code so they run much *faster* than a real machine would. Unfortunately, this means that if you wish to use measured execution time to support a design decision *you must measure that execution time when running directly on real hardware, not on any simulation or virtualization framework.*

13.3 Debugging Strategy

In general, when confronted by a mysterious problem, you should begin with a "story" of what you *expect* to be happening and measure the system you're debugging to see where its behavior diverges from your expectations.

To do this your story must be fairly detailed. For example, you should have a fairly good mental model of the assembly code generated from a given line of C code. To understand why "a variable has the wrong value" you need to know how the variable is initialized, where its value is

stored at various times, and how it moves from one location to another. If you're confused about this, it is probably good for you to spend some time with `gcc -S`.

Once your "story" is fleshed out, you will need to measure the system at increasing levels of detail to determine the point of divergence. You will find yourself spending some time thinking about how to pin your code down to observe whether or not a particular misbehavior is happening. You may need to write some code to periodically test data-structure consistency, artificially cause a library routine to fail to observe how your main code responds, log actions taken by your code and write a log-analyzer perl script, etc.

Don't forget about the debugger. In particular, **any time you find yourself "stuck,"** please review the course web site's page listing useful debugger commands. If you are "stuck," it is fairly likely that you should use one or two debugger commands that you have never used before. As you proceed through this class you are likely to encounter problems you have not encountered before; it is unwise to restrict yourself to using only old tools while trying to solve new problems.

When you encounter a fault or exception, you **must** determine three key pieces of information:

1. You must determine which instruction (not "line of code") can't be executed. Processors don't execute "lines of code"; they execute instructions.
2. Based on the surrounding code, determine what that instruction was intended to accomplish. Generally speaking, the instruction was selected by a compiler, based on preconditions expected to be true before the instruction executes and on conditions desired to be true after it's done. It is possible you will need to look up a description of exactly what the instruction does.
3. You will need to determine exactly why the instruction could not be executed. Generally speaking, some precondition isn't true, or some input value is wrong. Depending on the exception, the processor may write down some information about this particular execution failure; you will need to consult appropriate documentation to find what information is available and how to decode it. It is unwise to guess at which precondition/value is the source of the problem.

Please note that the user-space memory allocator we provide you with is very similar to the allocator written by 15-213 students in the sense that errors reported by the allocator, or program crashes which take place inside the allocator, are likely to mean that the user of some memory overflowed it and corrupted the allocator's meta-data. In the other direction, complaints by "lmm" are coming from the kernel's memory allocator, and probably indicate kernel bugs (see below).

13.4 Reference Kernel Panics and Crashes

If the Pebbles kernel tells you something went horribly wrong and drops you into the debugger, don't panic. It probably won't happen to most of you, but we are fully aware that we haven't nailed the last bug yet...

It's probably a good idea for you to tar up your working directory and make a brief note of what you were doing when the kernel ran into trouble. For example, what sequence of test programs had you run since boot? If you have a short repeatable way of getting the kernel to die,

that's excellent, and we'd appreciate a snapshot that lets us reproduce it, even if you then go on to modify your code to make the crash go away.

To send us a snapshot, tar it up somewhere in your group's scratch directory,

```
tar cfz .../mygroup/scratch/kcrash.somename.tgz .
```

create a brief summary of how to reproduce it,

```
$EDITOR .../mygroup/scratch/kcrash.somename.README
```

and send a brief note to the staff mailing list. While such an event will of course attract our attention, it's not likely that we can provide a fix in a small number of minutes...you may need to try to guess what went wrong and work around it temporarily, or work on some other part of your project for a while.

14 Strategy

14.1 Suggestions

First, this may be the first time you have written code with this variety and density of concurrency hazards. If so, you will probably find this code much harder to debug than code you've written before, i.e., you should allocate more debugging time than usual. Of course, the silver lining in this cloud is that experience debugging concurrent code will probably be useful to you after you leave this class.

Second, *several* of the thread library functions are *much* harder then they first appear. It is fairly likely that you will write half the code for a thread library function before realizing that you've never written "that kind of code" before. When this happens the best course of action is probably to come to a complete stop, think your way through the problem, and then explain the problem and your proposed solution to your partner. It may also happen that as you write your fifth function you realize your second must be scrapped and re-written.

Third, the Pebbles kernel offers a feature intended to help you increase the solidity of your code. A special system call, `void misbehave( int mode )`, alters the behavior of the kernel in ways which may expose unwarranted assumptions or concurrency bugs in your library code. Values for `mode` range from zero (the default behavior) to sixty-three (or maybe higher—see the test code), or you may select -1 for behavior which may be particularly challenging. As you experiment with `misbehave()`, you may become able to predict or describe the behavior of a particular `mode`. Each group must keep confidential its own understanding of the meanings of particular `mode` values.

Fourth, we recommend *against* splitting the assignment into two parts, working separately until the penultimate day, and then meeting to "put the pieces together." Instead, we recommend the opposite, namely that you make it a habit to read *and talk about* each other's code every few days. **You may encounter an exam question related to code your partner wrote, so when you "read" your partner's code, the reading should be sufficiently vigorous to find bugs.**

Fifth, we have observed that a particularly bad division of labor is for one person to write system call stubs, linked lists, queues, and maybe semaphores, while the other person writes everything else. This puts the first person at risk of doing poorly on exams.

Sixth, instead of typing linked-list traversal code 100 times throughout your library, thus firmly and eternally committing yourselves to a linear-time data structure, give some consideration to encapsulation. One approach is mentioned in Section 14.3.6.

Seventh, we **strongly** recommend that you use a source-control system to manage the evolution and/or devolution of your code. While the complexity of this project does not outright necessitate the use of source control, this is a good opportunity for you to get used to it and set up a work flow with your partner. No matter how busy you are now, you will be even busier during the kernel project, when source control will be even more important to your success than it is now. As a result, deciding not to use source control for this project may equate to giving up a letter grade on the kernel project.

Eighth, don't forget to do an update when `make` starts beeping at you. If you're in the middle of debugging a problem, you probably don't want to switch kernels, but you generally *do* want to upgrade when we issue new things, because we do so to help. A particularly bad thing to do is to work on your thread library for two weeks using the very oldest kernel and then 15 minutes before the assignment deadline switch to the very newest one and find that one time in a thousand you call `new_pages()` in an improper way which got through before and doesn't any more. So don't do that. The update process gives you the power to decide when to import changes, but that means the responsibility lies with you as well.

Ninth, speaking of the last minute, please avoid breaking your entire thread library in a frantic last-minute merge maneuver. A popular way to do this is for one person to debug some problem up until the last minute while another person rampages through the entire source tree fixing documentation, deleting dead code, and making code tweaks; then everything is dumped together and submitted. **Very frequently** this results in a submission which won't build, or which builds but won't boot, or which builds and boots but fails many tests; meanwhile, most of the documentation "fixes" are neutral or harmful. Untangling this sort of situation is very difficult because what's wrong will be one or two small changes buried among 150 other small changes. Probably the best way to avoid this scenario is to schedule a brief documentation and cleanup session for the *first* 15 to 30 minutes of time on *each work day*; naturally, it is best if your code changes and documentation changes are independent commits. Regardless, it is *known* that last-minute mega-merge mania usually fails; please don't try it to see whether it works out well for you.

Tenth, if you find yourself confused about what is meant when the kernel specification or the thread-library handout says "before," try to imagine what "before" might mean on a multiprocessor machine.

Finally, we have observed that the single most effective decision a group can make is to schedule standing "work meetings" of one or two hours duration two or three times per week. It is important that these be at **fixed times** each week agreed upon **in advance** (just like a class). Groups that do this consistently do better on the thread library and kernel projects than groups who don't.

14.2 Suggested Steps

1. Read the handouts. We believe that you should print both of them out (all the way onto paper!), read the Pebbles kernel specification from start to finish, then read this thread-library handout from start to finish, then most likely go back and read the Pebbles kernel specification again. Whether or not you print the handouts all the way onto paper, we

believe that "reading" them should result in extensive notes, whether those notes take the form of document markup or items on a group todo list.

2. Agree on two to three meeting times per week. An excellent thing to discuss early on is what source control system to use. By the way, make sure you configure it to *not* track changes to large random files such as `bootfd.img`, `bootfd.gz`, `user_apps.S`, the contents of `temp/`, etc., or your disk quota will be consumed very quickly.

3. Be sure to review the syllabus material on collaboration and the use of outside code. Seriously, please *right now* stop reading this list and go read that material, *even if you read it before*.

4. It is probably a good idea to acquire some practice with the "May we assume?" protocol from the "Questions" lecture by applying the protocol to at least a couple of the "may assume" statements found in this document.

5. Carefully review the comments in `config.mk`. That file contains information that will be useful to you while developing your code, and also information that you will need to know for your code to be graded successfully.

6. **Promptly** write system call wrappers for one or two system calls and run a small test program using those system calls. This is probably the best way to engage yourself in the project and to get an initial grasp of its scope. Good system calls to begin with are `set_status()` and `vanish()`, since the C run-time start-up code invokes the `exit()` library routine, which depends on them. A good second step would be `print()`.

7. Write the remaining system call wrappers (with the exception of `thread_fork`, of course).

8. If you're using revision control, make sure your repositories are private (readable by only you, your partner, and optionally the members of the course staff). Be careful: some popular project-hosting web sites *require* all repositories to be public, and they will automatically publish your code on various search engines even if you do not explicitly publish it. Also note that some "web clipboard" services, e.g., `pastebin.com`, default to publicizing any content you share through/with them; search engines will eventually pick up and index that content. Be careful out there! Meanwhile, note that your group's course AFS space contains a `REPOSITORY` directory which can be used as a central point when you're logged in to an Andrew Linux machine... and can even be accessed remotely with a smidgin of hacking, see `410/pub/dotssh-slash-rc`.

9. If you're not using revision control, you should be.

10. Read at least half of the test code we have provided. Doing this early can avoid potentially-costly last-minute discoveries of misunderstandings.

11. Design and make a **draft** version of mutexes and condition variables. In order to do that, you will probably need to perform a hazard analysis of which code sequences in your thread library would suffer if the scheduler switched from executing one of your threads to another. For mutexes in particular, please do **not** "start from infinity," by which we mean designing the best possible mutex you can imagine before you implement anything at all. It is much better to implement "ok mutexes" or "pretty good mutexes" so you can get started on other things. You can revisit your mutex design later.

12. Now would not be a bad time to read the source to the "thread group" library (Section 7). If you read the source code we provide *before* "debugging time," it may help you do a better design, and thus need to do less debugging. Now would also probably be a good time to read the textbook material on the "Producer-Consumer" (aka "Bounded-Buffer") pattern.

13. If you haven't yet, agree on two to three meeting times per week.

14. Now would be a good time to write at least an initial version of your `malloc()` wrappers.

15. What can you test at this point? Be creative.

16. Think hard about stacks. What should the child's stack look like before and after a `thread_fork`? In fact, it is probably a good idea for you to draw every detail of the parent's stack and the child's stack before and after `thread_fork`. You should reach this point by Friday, February 11th.

17. Write and test `thr_init()` and `thr_create()`. Run the STARTLE test. You should reach this point by Monday, February 14th.

18. Write an initial `thr_exit()` that makes threads go away, without solving all of the problems needed to make `thr_join()` work plus all of the other problems.

19. From this point forward, it is *very important* that the code you write is based on case analysis of errors as described in the "Errors" lecture. It is *not a good idea* to "code up" your thread library without attention to error case analysis, then use the test suite to "debug" your thread library, and then think about error handling the last minute. This is known to work out poorly in a grade sense for this project, and also known to set you off on the wrong foot for the next project.

20. Test mutexes and condition variables. Try to reach this point by Wednesday, February 16th.

21. Try all the `misbehave()` flavors.

22. Write and test `thr_join()`. This will likely involve upgrades to `thr_exit()`.

23. This might be a good time to consider all possible negative interactions between/among `thr_create()`, `thr_exit()`, and `thr_join()`.

24. Write a basic software exception handler which implements automatic stack growth for legacy single-threaded applications. This should not be a lot of code, and you can revisit your implementation later if you wish.

25. This might be a good point to relax and have fun writing semaphores.

26. This is a potential point for you to revisit your mutex design.

27. Test. Debug. Test. Debug. Test. Sleep once in a while.

28. Try all the `misbehave()` flavors (again). Note that most of the tests provided to you by the course staff (see `410user/progs/README`) are really multiple tests if you think about it... you probably shouldn't declare a test "passed" until *all* versions pass. Remember that you should be running CYCLONE and AGILITY_DRILL by Friday, February 18th.

29. Design, implement, and test readers/writers locks.

30. Pick the two least-obvious/most-tricky lock placements in your code; talk them through with your partner. For each, document why the lock is acquired in the right place and why it is released in the right place.

31. Revisit your stack-growth exception handler and also decide what should happen when threads in a multi-threaded application run into various kinds of exceptions or panic.

32. If you have time, maybe try out your thread library on the 15-410 "crash box" (see the "Projects" web page).

33. Celebrate! You have assembled a collection of raw system calls into a robust and useful thread library.

14.3 Questions & Challenges

Below we briefly discuss common questions about this assignment and issue several optional challenges. It is very important that your implementation be solid, and you should not be diverted from this primary goal by attempting to solve these challenges. However, we are providing this challenge list as a way for interested students to deepen their understanding and sharpen their design skills.

14.3.1 Questions

From time to time we are asked how many threads must be supported by a library implementation. In general the answer is that the thread library should not be a limiting factor—it should be possible to use all available memory for threads, and of course it could happen one day that Pebbles would run on a machine with more memory. If, however, you feel you *must* impose an a-priori static limit on the number of threads (or some other run-time feature), we will grade less harshly if you document your reasoning.

Sometimes we are asked to state a simple requirement about bounded waiting (e.g., "Are we required to implement the bounded waiting algorithm presented in the lecture slides?"). Since this is a design class, you should give serious consideration to the issue of bounded waiting and the interplay between bounded waiting and the system environment you will be using. Then you should be in a position to evaluate the necessity of ensuring or approximating bounded waiting and how you might go about doing that. Whatever you choose to do should sensibly balance cost against utility. Your project documentation should briefly but convincingly explain your reasoning.

What should happen if a thread is killed by an exception? Solving this problem in the general case is much too difficult for this assignment, but it probably would be a good idea to think about whether there is anything reasonable you could do with "a bit" of code and, if so, to try to do it.

14.3.2 Challenge: efficient `thr_getid()`

There is an easy way to implement `thr_getid()`, but it is woefully inefficient. Can you do better? We have given you a serious hint.

14.3.3 Challenge: `thr_init()`

Is it really necessary that `thr_init()` be called before `malloc()`? How might you build `malloc()` to make that unnecessary?

14.3.4 Challenge: "reaper thread"

If you feel you need a "reaper thread," consider whether it's *really* necessary.

14.3.5 Challenge: memory-efficient `thr_exit()`

Since there is no bound on how much time can pass between a thread exiting and its "parent" or "manager" thread calling `thr_join()`, it is undesirable for a "zombie thread" to hold onto large amounts of memory. Can you avoid this situation? There are multiple approaches, with different tradeoffs.

14.3.6 List Traversal Macros

You may find yourself wishing for a way for a TCB to be on multiple lists at the same time but not relish the thought of writing several essentially identical list traversal routines. Other languages have generic-package facilities, but C does not. However, it is possible to employ the C preprocessor to automatically generate a family of similar functions. If you wish to pursue this approach, you will find a template available in `vq_challenge/variable_queue.h`. It is certainly possible to write a thread library without doing this. However, making this investment now will likely bear fruit during the kernel project.