Lectures on Operating Systems (Mythili Vutukuru, IIT Bombay) Lab: Pthreads Synchronization

In this lab, you will learn the basics of multi-threaded programming, and synchronizing multiple threads using locks and condition variables. You will use the pthreads (POSIX threads) API in this assignment.

Before you begin

Please familiarize yourself with the pthreads API thoroughly. Many helpful tutorials and sample programs are available online. The chapter on pthreads in OSTEP is also an excellent guide. Practice writing simple programs with multiple threads, using locks and condition variables for synchronization across threads. Remember that you must include the header file <pthread.h> and compile code using the -lpthread flag when using this API. For example, below is how you will compile a multi-threaded program.

```
$gcc threads.c -lpthread
$./a.out
```

Warm-up exercises

You may write the following simple programs before you get started with this lab.

- 1. Write a program that has a counter as a global variable. Spawn 10 threads in the program, and let each thread increment the counter 1000 times in a loop. Print the final value of the counter after all the threads finish—the expected value of the counter is 10000. Run this program first without using locking across threads, and observe the incorrect updation of the counter due to race conditions (the final value will be slightly less than 10000). Next, use locks when accessing the shared counter and verify that the counter is now updated correctly.
- 2. Write a program where the main default thread spawns N = 10 threads. Thread *i* should print the message "I am thread *i*" to screen and exit. The main thread should wait for all N threads to finish, then print the message "I am the main thread", and exit.

Note that you must learn how to pass arguments (say, thread number) to the thread start function in order to solve this lab. When you pass a variable as an argument, care should be taken to ensure that the value of the variable is retained until the thread actually runs. So, you must use different variables for different threads, and not reuse the same variable to store different values for different threads.

3. You are given the file sequence.c, which creates N = 3 threads. Each thread *i* prints the string "I am thread *i*" to the screen in an infinite loop. You must add synchronization between the threads using locks and conditon variables so that the threads print one after the other in the order 0, 1, 2,

0, 1, 2, and so on. Please note that threads must NOT simply do busy-waiting, but must wait in a condition variable queue until their turn to print comes.

A sample execution of the code is shown below.

```
I am thread 0
I am thread 1
I am thread 2
I am thread 0
I am thread 1
I am thread 2
I am thread 2
I am thread 0
I am thread 1
I am thread 2
...
```

You can try to make your code with larger values of N as well. You can also try different ways of designing your solution, e.g., using only one CV vs. using multiple CVs.

Part A: Speeding up computation using multi-threading

We will consider a program with a long computation, and improve its run time by using multi-threading, in order to understand how real systems parallelize and speedup compute-intensive code.

You are given a program goldbach.c that verifies the Goldbach conjecture up to N=100K numbers (you can try this out for larger values of N as well). The Goldbach conjecture states that every even number greater than 2 can be expressed as a sum of two prime numbers. The program given to you does the following. It first computes all primes up to N using the "Sieve of Eratosthenes" algorithm, and stores this information in a boolean array primes. That is, primes[i] is true if i is prime, and false otherwise.

Now the program begins its task of verifying the Goldbach conjecture. It iterates over all even numbers, and counts the number of "Goldbach pairs", i.e., primes p_1 and p_2 that add up to the even number. It stores this count in an array gb_count of size N/2. The value of gb_count[i] is the number of prime pairs that add up to 2i, for even number 2i. The program computes the time taken for this calculation of the counts, by recording timestamps at the start and end of this calculation, and prints out the total time elapsed. Most systems should take a few seconds to finish this calculation. Finally, the program writes the values of the number of prime pairs to a file output.txt so that we can verify the conjecture.

Now, your job is to speed up this calculation of the number of Goldbach pairs for each even number using multithreading. Your code must create 10 threads, assign different, disjoint ranges of numbers to these threads (there are many ways to do this; the choice is left to you), so that the threads can calculate the number of Goldbach pairs for different even numbers in parallel. You must create threads after we have recorded the starting timestamp, and the threads must finish their calculation and join before we record the ending timestamp. In this way, we can measure the amount of time taken to compute the number of Goldbach pairs in parallel across 10 threads, and compare it with the time taken in the single-threaded program.

The calculation of the number of Goldbach pairs for each even number should be done by exactly one thread. You can distribute the work between the threads in any way you see fit. You must only parallelize

the filling up of the gb_count array, and you need not parallelize the calculation of the primes itself (that happens before the timer is started). You can move around the code you have to parallelize to a separate function, and provide it as the start function for the various threads. You must provide suitable arguments to the threads so that each thread knows which range of numbers it must compute on. You may need to make some of the data structures like the primes and gb_count arrays as global variables so that they can be accessed easily by all threads. Think carefully on whether you need locking if the different threads access different parts of the array corresponding to different ranges of numbers. Please do not tamper with other parts of the code not relevant to you, like the time calculation.

We expect you to write a parallel version of the same program goldbach.c which runs a few times faster than the single-threaded version. A sample execution of both the single threaded code given to you, and a multi-threaded version you must write, and shown below. Note that we have to use the "-lm" flag when compiling to use the math library, and "-lpthread" flag to use the pthread library. We have also given you the output your code must generate in expected-output.txt.

Single threaded execution:

```
$ gcc goldbach.c -lm
$ ./a.out
Computed primes upto 100000, count = 9592
elapsed time: 3448984 microseconds
$ tail output.txt
99982 608
99984 1216
99986 603
99988 736
99990 1855
99992 638
99994 651
99996 1303
99998 605
100000 810
$ diff output.txt expected-output.txt
Ś
```

Multi-threaded execution:

```
$ gcc goldbach.c -lpthread -lm
$ ./a.out
Computed primes upto 100000, count = 9592
elapsed time: 931387 microseconds
$ diff output.txt expected-output.txt
$
```

Note that in the solution written by us, the multi-threaded version with 10 threads runs about 4 times faster, and produces identical output, i.e., correctly computes the number of Goldbach pairs. The exact run times and speedup you may see may differ.

Part B: Master-Worker Thread Pool

In this part of the lab, you will implement a simple master and worker thread pool, a pattern that occurs in many real life applications. The master threads produce numbers continuously and place them in a buffer, and worker threads will consume them, i.e., print these numbers to the screen. This simple program is an example of a master-worker thread pool pattern that is commonly found in several real-life application servers. For example, in the commonly used multi-threaded architecture for web servers, the server has one or more master threads and a pool of worker threads. When a new connection arrives from a web client, the master accepts the request and hands it over to one of the workers. The worker then reads the web request from the network socket and writes a response back to the client. Your simple master-worker program is similar in structure to such applications, albeit with much simpler request processing logic at the application layer.

You are given a skeleton program master-worker-skeleton.c. This program takes 4 command line arguments: how many numbers to "produce" (M), the maximum size of the buffer in which the produced numbers should be stored (N), the number of worker threads to consume these numbers (C), and the number of master threads to produce numbers (P). The skeleton code spawns P master threads, that produce the specified number of integers from 0 to M - 1 into a shared buffer array. The main program waits for these threads to join, and then terminates. The skeleton code given to you does not have any logic for synchronization.

You must write your solution in the file master-worker.c. You must modify this skeleton code in the following ways. You must add code to spawn the required number of worker threads, and write the function to be run by these threads. This function will remove/consume items from the shared buffer and print them to screen. Further, you must add logic to correctly synchronize the producer and consumer threads in such a way that every number is produced and consumed exactly once. Further, producers must not try to produce when the buffer is full, and consumers should not consume from an empty buffer. While you need to ensure that all C workers are involved in consuming the integers, it is not necessary to ensure perfect load balancing between the workers. Once all M integers (from 0 to M - 1) have been produced and consumed, all threads must exit. The main thread must call pthread_join on the master and worker threads, and terminate itself once all threads have joined. Your solution must only use pthreads condition variables for waiting and signaling: busy waiting is not allowed.

Your code can be compiled as shown below.

gcc master-worker.c -lpthread

If your code is written correctly, every integer from 0 to M-1 will be produced exactly once by the master producer thread, and consumed exactly once by the worker consumer threads. We have provided you with a simple testing script (test-master-worker.sh which invokes the script check.awk) that checks this above invariant on the output produced by your program. The script relies on the two print functions that must be invoked by the producer and consumer threads: the master thread must call the function print_produced when it produces an integer into the buffer, and the worker threads must call the function print_consumed when it removes an integer from the buffer to consume. You must invoke these functions suitably in your solution. Please do not modify these print functions, as their output will be parsed by the testing script.

Please ensure that you test your case carefully, as tricky race conditions can pop up unexpectedly. You must test with up to a few million items produced, and with a few hundreds of master/worker threads. Test for various corner cases like a single master or a single worker thread or for a very small buffer size as well. Also test for cases when the number of items produced is not a multiple of the buffer size or the number of master/worker threads, as such cases can uncover some tricky bugs. Increasing the number of threads to large values beyond a few hundred will cause your system to slow down considerably, so exercise caution.

Part C: Reader-Writer Locks

Consider an application where multiple threads of a process wish to read and write data shared between them. Some threads only want to read (let's call them "readers"), while others want to update the shared data ("writers"). In this scenario, it is perfectly safe for multiple readers to concurrently access the shared data, as long as no other writer is updating it. However, a writer must still require mutual exclusion, and must not access the data concurrently with any other thread, whether a reader or a writer. A reader-writer lock is a special kind of a lock, where the acquiring thread can specify whether it wishes to acquire the lock for reading or writing. That is, this lock will expose two separate locking functions, say, *ReaderLock()* and *WriterLock()*, and analogous unlock functions. If a lock is acquired for reading, other threads that wish to read may also be permitted to acquire the lock at the same time, leading to improved concurrency over traditional locks.

There are two flavors of the reader-writer lock, which we will illustrate with an example. Suppose a reader thread R1 has acquired a reader-writer lock for reading. While R1 holds this lock, a writer thread W and another reader thread R2 have both requested the lock. Now, it is fine to allow R2 also to simultaneously acquire the lock with R1, because both are only reading shared data. However, allowing R2 to acquire the lock may prolong the waiting time of the writer thread W, because W has to now wait for both R1 and R2 to release the lock. So, whether we wish to permit more readers to acquire the lock when a writer is waiting is a design decision in the implementation of the lock. When a reader-writer lock is implemented with *reader preference*, additional readers are allowed to concurrently hold the lock with previous readers, even if a writer thread is waiting for the lock. In contrast, when a reader-writer lock is implemented with *writer preference*, additional readers are not granted the lock when a writer is already waiting. That is, we do not prolong the waiting time of a writer any more than required.

In this part of the lab, you must implement both flavors of the reader-writer lock. You must complete the definition of the structure that captures the reader-writer lock in rwlock.h. The following functions to be supported by this lock are also defined in the header file:

- The function InitalizeReadWriteLock () must initialize the lock suitably.
- The function ReaderLock() is invoked by a reader thread before entering a read-only critical section, and the function ReaderUnlock() is invoked by the reader when exiting the critical section.
- The function WriterLock() is invoked by a writer thread before entering a critical section with shared data updates, and the function WriterUnlock() is invoked by the writer when exiting the critical section.

You must write the code to implement these functions. You must write code that implements readerwriter locks with reader preference in rwlock-reader-pref.cpp. Similarly, you must implement reader-writer locks with writer preference in rwlock-writer-pref.cpp. Note that both implementations of the lock must share the same header file, including the definition of the reader-writer lock structure. Therefore, your lock structure may have some fields that are only used in one version of the code and not the other. As part of the autograding scripts, you are given two tester programs that test each variant of your reader-writer lock, and an autograding script that runs both in one go. The tester program takes two command line arguments, say R and W. The program then spawns R reader threads, followed by W writer threads, followed by R additional reader threads. Each thread, upon creation, tries to acquire the same reader-writer lock as a reader or writer (as the case may be), holds the lock for a long period of time, and finally releases the lock. The program judges the correctness of your implementation by observing the relative ordering of the acquisitions and releases of these locks. By invoking this tester with different values of R and W, one can test the reader-writer lock code reasonably thoroughly. Of course, you are encouraged to write your own test programs that use the reader-writer lock as well.

Part D: Semaphores using pthreads

In this part of the lab, you will implement the synchronization functionality of semaphores using pthreads mutexes and condition variables. Let's call these new userspace semaphores that you implement as zemaphores, to avoid confusing them with semaphores provided by the Linux kernel. You must define your zemaphore structure in the file zemaphore.h, and implement the functions zem_init, zem_up and zem_down that operate on this structure in the file zemaphore.c. The semantics of these zemaphore functions are similar to those of the semaphores you have studied in class.

- The function zem_init initializes the specified zemaphore to the specified value.
- The function zem_up increments the counter value of the zemaphore by one, and wakes up any one sleeping thread.
- The function zem_down decrements the counter value of the zemaphore by one. If the value is negative, the thread blocks and is context switched out, to be woken up by an up operation on the zemaphore at a later point.

Once you implement your zemaphores, you can use the program test-zem.c to test your implementation. This program spawns two threads, and all three threads (the main thread and the two new threads) share a zemaphore. Before you implement the zemaphore logic, the new threads will print to screen before the main thread. However, after you implement the zemaphore correctly, the main thread will print first, owing to the synchronization enabled by the zemaphore. You must not modify this test program in any way, but only use it to test your zemaphore implementation.

Next, you are given a simple program with three threads in test-toggle.c. In its current form, the three threads will all print to screen in any order. Modify this program in such a way that the print statements of the threads are interleaved in the order thread0, thread1, thread2, thread0, thread1, thread2, ... and so on. You must only use your zemaphores to achieve this synchronization between the threads. You must not directly use the mutexes and condition variables of the pthreads library in the file test-toggle.c.

The script test-zem.sh will compile and run these test programs for you and can be used for testing. Please see the compilation commands in the script to understand how to compile code that uses pthreads using the -lpthread flag.