

CS 3214 Fall 2020 Test 2 Solutions

November 11, 2020

Contents

1	Pass on Threads? (12 pts)	2
2	A River Crossing Puzzle (12 pts)	4
3	Making Random Permutations (15 pts)	9
4	A Deadlock Mystery (8 pts)	16
5	Optimizing Parallel Performance (8 pts)	20
6	Submission Requirements	21

Rules

- This exam is open book, open notes, and open Internet, but in a read-only way.
- You are not allowed to post or otherwise communicate with anyone else about these problems.
- You are required to cite any sources you use, except for lecture material, source code provided as part of the class materials, and the textbook.
- If you have a question about the exam, you may post it as a *private* question on Piazza. If it is of interest to others, I will make it public.
- During the test period, you may not use Piazza to post other questions or reply to questions posted there.
- Any errata to this exam will be published prior to 12h before the deadline.

1 Pass on Threads? (12 pts)

Consider the following program:

```
1  #include <stdio.h>
2  #include <pthread.h>
3  #include <stdint.h>
4  #include <unistd.h>
5  #include <stdlib.h>
6  #include <stdatomic.h>
7
8  #define LIST_SIZE 5
9
10 struct bagOStuff {
11     atomic_int threadID;
12     long * shouldBeHidden;
13     atomic_int * listOStuff;
14 };
15
16 void * thread_start1(void * arg) {
17     struct bagOStuff * myBag = (struct bagOStuff *)arg;
18
19     for(int i = 0; i < LIST_SIZE; ++i) {
20         myBag->listOStuff[i] = myBag->threadID;
21     }
22
23     printf("Thread: %u\n", myBag->threadID);
24
25     return NULL;
26 }
27
28 void * thread_start2(void * arg) {
29     struct bagOStuff * myBag = (struct bagOStuff *)arg;
30     printf("Thread: %u\n", myBag->threadID);
31     printf("Thread's hidden: %lu\n", *myBag->shouldBeHidden);
32
33     for(int i = 0; i < LIST_SIZE; ++i) {
34         printf("%u ", myBag->listOStuff[i]);
35     }
36     printf("\n");
37
38     *myBag->shouldBeHidden = 9035768;
39
40     return NULL;
41 }
42
```

```

43 int main(void) {
44     long hidden = 8675309;
45     pthread_t thread1, thread2;
46     atomic_int * intList = (atomic_int *)calloc(LIST_SIZE, sizeof(atomic_int));
47     struct bagOStuff toPass = {1, &hidden, intList};
48
49     pthread_create(&thread1, NULL, thread_start1, (void *)&toPass);
50
51     toPass.threadID = 2;
52
53     pthread_create(&thread2, NULL, thread_start2, (void *)&toPass);
54
55     pthread_join(thread1, NULL);
56     pthread_join(thread2, NULL);
57
58     printf("Main's hidden: %lu\n", hidden);
59 }

```

Using the atomic data types provided by `stdatomic.h` ensures that even in the presence of data races, the resulting operations on the protected variable are sequentially consistent (i.e., behaves how programmers would expect). While the atomics library ensures that the low-level variable accesses are well-behaved, thread inter-leavings may still impact the result of execution (i.e., sequential consistency is not the same as having one possible total ordering of variable accesses). Keep that in mind while answering the following questions:

- (2 pts) What is/are the possible ID values printed by the first thread created in `main()`? Why? Saying “Because that’s what happened when I ran it” is insufficient.
[Solution] The possible ID values will be 1 or 2, depending on the ordering of `main()` preparing the struct for thread 2 and thread 1’s execution.
- (2 pts) What is/are the possible ID values printed by the second thread created in `main()`? Why? Saying “Because that’s what happened when I ran it” is insufficient.
[Solution] The only possible ID printed is 2, since the change of ID happens before thread 2 is created. Thus the ID value update always happens-before the ID print by thread 2.
- (2 pts) Both threads take an argument `arg` that they store in the variable `myBag`. Assume that instead of a `struct` the parameter was a scalar (e.g., `int`). See line 104 of Problem 3 for an example of how to do this with threads. Does such a variable need to be protected from data races across threads? Why or why not?
[Solution] Threads have their own stack and registers and scalars, such as `ints`, are passed by value. When that value is assigned to a thread’s local variable (`myBag` in this case), it resides in that thread’s stack or registers. Thus, all accesses are private and not subject to data races.
- (2 pts) What does the thread see as the value of the variable `shouldBeHidden` on line 31 and why (especially with respect to the previous question/answer)?

[Solution] 8675309 Because `main()`'s creation of thread 2 happens-before the printing by thread 2. Even though you threads have their own stack (i.e., the previous answer) and `hidden` is allocated on `main()`'s stack, threads share the same, single, virtual address space. Thus, if you have a variable's address, you can access that variable.

5. (2 pts) How many possible list strings does the second thread created print out (e.g., "0 0 0 0 0", "0 0 0 0 1", and "1 0 0 0 0" count as three different list strings)?

[Solution] There is a happens-before relationship that creates two distinct thread scheduling races that we must account for (I use the term scheduling race to indicate when the results of execution depend on thread scheduling. When this is the case, a program has multiple sequentially-consistent outcomes.): by the time thread 2 starts printing, `main()` updated the thread ID in the struct passed to the threads during creation, which thread 1 uses for list updates (as you know from the first question). This means that thread 1 will set the list to some, potentially zero, number of 1's, followed by enough 2's to fill the five-item list. This results in six possible prefixes. For each of these prefixes, the second thread scheduling race determines whether the value of 0 from `main()` is printed or thread 1's ID (which must now be 2, due to the previously-described happens-before relationship). The key idea behind this scheduling race is that the list value updates by thread 1 can be preempted—at any time—by thread 2's printing. Likewise, thread 2's printing can be preempted—at any time—by thread 1's list updates. This makes each index of the array (as printed by thread 2) the subject of an independent scheduling race between the 0 (as set by `main()`) or thread 1's ID. Thus, for the N bits remaining after the leading-1's, there are 2^N possible values. To determine the total number of possible string values, we sum the result for each of the six possible leading-1 prefixes. This results in a total of 63 possible strings. As a "fun" extension to this problem, consider what happens if thread 1 and thread 2 traverse the list from opposite directions.

6. (2 pts) What does `main()` see as the value of its variable `hidden` on line 58 and why?

[Solution] 9035768, because a happens-before relationship is created by the `pthread_join()`. Again shared virtual address space (i.e., no copy-on-write as seen with multiprocessing).

2 A River Crossing Puzzle (12 pts)

The wolf, goat, and cabbage problem is a river crossing puzzle that, according to Wikipedia [URL], dates back centuries.

Once upon a time a farmer went to a market and purchased a wolf, a goat, and a cabbage. On his way home, the farmer came to the bank of a river and rented a boat. But crossing the river by boat, the farmer could carry only himself and a single one of his purchases: the wolf, the goat, or the cabbage.

If left unattended together, the wolf would eat the goat, or the goat would eat the cabbage.

The farmer's challenge was to carry himself and his purchases to the far bank of the river, leaving each purchase intact. How did he do it?

If the man, the wolf, the goat, and the cabbage were threads, a possible solution that denotes their actions could be described by the following program:

```
1  #include <pthread.h>
2  #include <semaphore.h>
3  #include <stdlib.h>
4  #include <stdio.h>
5  #include <assert.h>
6  #include <stdint.h>
7  #include <stdbool.h>
8  #include <string.h>
9  #include <unistd.h>
10
11 // you may define your condition variable(s) here
12 // if you need mutexes too, define them here
13 // as well as any other global variables (if needed)
14 // you may use helper functions.
15 //
16 // You may not use semaphores.
17 //
18
19 // you may change the following four functions, but you may not change
20 // the order of the printf statements, nor the thread that makes those
21 // statements, and you may not add additional printf statements.
22 static void *
23 man(void *_arg)
24 {
25     printf("man enters boat\n");
26     printf("man returns\n");
27     printf("man returns once more\n");
28     return NULL;
29 }
30
31 static void *
32 wolf(void *_arg)
33 {
34     printf("wolf and man cross left to right\n");
35     return NULL;
36 }
37
38 static void *
39 cabbage(void *_arg)
40 {
41     printf("cabbage and man cross left to right\n");
42     return NULL;
43 }
44
45 static void *
```

```

46 goat(void *_arg)
47 {
48     printf("goat and man cross left to right\n");
49     printf("goat and man cross right to left\n");
50     printf("goat arrives on other side for good\n");
51     return NULL;
52 }
53
54 int
55 main(int ac, char *av[])
56 {
57     pthread_t t[4];
58
59     // you may add code here
60     pthread_create(&t[0], NULL, man, NULL);
61     pthread_create(&t[1], NULL, wolf, NULL);
62     pthread_create(&t[2], NULL, cabbage, NULL);
63     pthread_create(&t[3], NULL, goat, NULL);
64
65     for (int i = 0; i < 4; i++)
66         pthread_join(t[i], NULL);
67     // you may add code here
68 }

```

Unfortunately, due to the unpredictable nature of thread scheduling on a preemptive multicore system, this does not output a solution that would preserve all 3 items.

Your task is to use condition variables (not semaphores!) to ensure that the program always outputs the correct solution, which is:

```

man enters boat
goat and man cross left to right
man returns
cabbage and man cross left to right
goat and man cross right to left
wolf and man cross left to right
man returns once more
goat arrives on other side for good

```

- Your program must not use semaphores or other POSIX constructs besides POSIX condition variables and, if necessary mutexes. You may define global variables if that is needed.
- You cannot use functions such `sleep()` or `nanosleep()` because they do not guarantee that a thread will be scheduled right after being woken up.
- Your program should not contain data races and should not busy wait, which you can test with the usual tools.

[Solution]

```

1  #include <pthread.h>
2  #include <semaphore.h>
3  #include <stdlib.h>
4  #include <stdio.h>
5  #include <assert.h>
6  #include <stdint.h>
7  #include <stdbool.h>
8  #include <string.h>
9  #include <unistd.h>
10
11 /*
12  *   Wolf Cabbage Goat Man   --
13  *   --> Goat Man
14  *   Wolf Cabbage           -- Goat Man
15  *   <-- Man
16  *   Wolf Cabbage Man       -- Goat
17  *   --> Cabbage Man
18  *   Wolf                   -- Goat Cabbage Man
19  *   <-- Goat Man
20  *   Wolf Goat Man         -- Cabbage
21  *   --> Wolf Man
22  *   Goat                   -- Wolf Cabbage Man
23  *   <-- Man
24  *   Goat Man              -- Wolf Cabbage
25  *   --> Goat Man
26  *                           -- Wolf Cabbage Goat Man
27  */
28 pthread_mutex_t statelock = PTHREAD_MUTEX_INITIALIZER;
29 pthread_cond_t cond = PTHREAD_COND_INITIALIZER;
30 int state; // 0, 1, 2, 3, ...
31
32 static void signal()
33 {
34     state++;
35     pthread_cond_broadcast(&cond);
36 }
37
38 static void waitfor(int wantstate)
39 {
40     while (state != wantstate)
41         pthread_cond_wait(&cond, &statelock);
42 }
43
44 static void *
45 man(void *_arg)

```

```

46 {
47     pthread_mutex_lock(&statalock);
48     printf("man enters boat\n");
49     signal();
50     waitfor(2);
51     printf("man returns\n");
52     signal();
53     waitfor(6);
54     printf("man returns once more\n");
55     signal();
56     waitfor(7);
57     pthread_mutex_unlock(&statalock);
58     return NULL;
59 }
60
61 static void *
62 wolf(void *_arg)
63 {
64     pthread_mutex_lock(&statalock);
65     waitfor(5);
66     printf("wolf and man cross left to right\n");
67     signal();
68     pthread_mutex_unlock(&statalock);
69     return NULL;
70 }
71
72 static void *
73 cabbage(void *_arg)
74 {
75     pthread_mutex_lock(&statalock);
76     waitfor(3);
77     printf("cabbage and man cross left to right\n");
78     signal();
79     pthread_mutex_unlock(&statalock);
80     return NULL;
81 }
82
83 static void *
84 goat(void *_arg)
85 {
86     pthread_mutex_lock(&statalock);
87     waitfor(1);
88     printf("goat and man cross left to right\n");
89     signal();
90     waitfor(4);
91     printf("goat and man cross right to left\n");

```



```

92     signal();
93     waitfor(7);
94     printf("goat arrives on other side for good\n");
95     pthread_mutex_unlock(&statelock);
96     return NULL;
97 }
98
99 int
100 main(int ac, char *av[])
101 {
102     pthread_t t[4];
103
104     pthread_create(&t[0], NULL, man, NULL);
105     pthread_create(&t[1], NULL, wolf, NULL);
106     pthread_create(&t[2], NULL, cabbage, NULL);
107     pthread_create(&t[3], NULL, goat, NULL);
108
109     for (int i = 0; i < 4; i++)
110         pthread_join(t[i], NULL);
111 }

```

3 Making Random Permutations (15 pts)

The goal of the following program is to output 10 permutations of the string `abcdefghij`, for example:

```

cebdihgjaf
ghfejacibd
djcehigbaf
begafjchdi
ibaegjhdcf
jacfebgihd
ecjaifgdbh
jibegfhcad
idfhjcaegb
ijgahcedfb

```

The program is shown below:

```

1 // making random permutations
2 #include <pthread.h>
3 #include <semaphore.h>
4 #include <stdlib.h>
5 #include <stdio.h>
6 #include <assert.h>

```

```

7  #include <stdint.h>
8  #include <stdbool.h>
9  #include <string.h>
10 #include <unistd.h>
11
12 #define NTHREADS 10
13
14 static pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZER;
15 static sem_t permutation_done;
16 static sem_t permutation_extra;    // not currently used
17
18 // you may change these 2 lines below if necessary
19 const int PERMUTATION_DONE_INITIAL_VALUE = 0;
20 const int PERMUTATION_EXTRA_INITIAL_VALUE = 0;
21
22 static bool use_slow_threads;
23 typedef void (*func_t)(void);
24
25 /*
26  * This function is called by exactly n threads, repeatedly, as
27  * shown below in thread_fun.
28  * Each time it's called by these threads it should output a
29  * permutation of the string 'abcdefghij', following by a newline.
30  * Thread 1 must output 'a', thread 2 must output 'b', and so on
31  * Thread i must output 'a' + i.
32  *
33  * Exactly one thread should invoke the `when_permutation_finished`
34  * function when the permutation has been output and before the
35  * output of the next permutation has been started.
36  */
37 static void
38 make_a_permutation(int n, int id, func_t when_permutation_finished)
39 {
40     assert (0 <= id && id < n);
41     // you may not change this line
42     printf("%c", 'a' + id);
43
44     // you may make changes below in the remainder of this function
45     static int permutation_index = 0;
46
47     pthread_mutex_lock(&lock);
48     int pi = ++permutation_index;
49     pthread_mutex_unlock(&lock);
50
51     if (pi == n) {
52         when_permutation_finished();

```

```

53     sem_post(&permutation_done);
54 }
55
56 sem_wait(&permutation_done);
57 sem_post(&permutation_done);
58
59 pthread_mutex_lock(&lock);
60 pi = --permutation_index;
61 pthread_mutex_unlock(&lock);
62
63 if (pi == 0)
64     sem_wait(&permutation_done);
65 // you may not make changes below this line
66 }
67
68 static void
69 newline(void)
70 {
71     printf("\n");
72     fflush(stdout);
73 }
74
75 static void *
76 thread_fun(void *_arg)
77 {
78     int id = (uintptr_t) _arg;
79     for (int i = 0; i < 10; i++) {
80         if (use_slow_threads) {
81             // slow threads sleep a random amount of time
82             // drawn from [.5ms, 1.5ms]
83             struct timespec ts = {
84                 .tv_sec = 0,
85                 .tv_nsec = 500000 + rand() % 1000000
86             };
87             nanosleep(&ts, 0);
88         }
89         make_a_permutation(NTHREADS, id, newline);
90     }
91     return NULL;
92 }
93
94 int
95 main(int ac, char *av[])
96 {
97     pthread_t threads[NTHREADS];
98     sem_init(&permutation_done, 0, PERMUTATION_DONE_INITIAL_VALUE);

```

```

99     sem_init(&permutation_extra, 0, PERMUTATION_EXTRA_INITIAL_VALUE);
100
101     use_slow_threads = !(ac > 1 && !strcmp(av[1], "-f"));
102
103     for (int i = 0; i < NTHREADS; i++)
104         pthread_create(threads + i, NULL, thread_fun, (void *) (uintptr_t) i);
105
106     for (int i = 0; i < NTHREADS; i++)
107         pthread_join(threads[i], NULL);
108 }

```

Consider this program, and answer the following questions.

1. (3 pts) If you compiled and ran this program on a multiprocessor machine such as our rlogin cluster without specifying any command line switches, it would appear to “work” probably most of the time and indeed output 10 permutations. Analyze the program and briefly explain why this is the case.

[Solution] Under benign conditions (imagine the threads running very slowly, or one at a time), the first $n - 1$ threads to arrive in the `make_a_permutation` function will wait on line 56 until the last thread arrives, which then calls `when_permutation_finished` on line 52, then posts to the semaphore. As a result, one thread will make it past line 56. This thread will post this semaphore, allowing another thread to make it past line 56. The semaphore is thus signaled $n + 1$ times overall; one of the threads (the thread that sees that `pi` has reached 0) will wait on line 64 to consume the extra signal. Thus under benign conditions, each thread waits on a “turnstile” semaphore that will not be signaled until the last thread arrives. This ensures that `when_permutation_finished` is not called until all threads have arrived.

2. (3 pts) Unfortunately, the program ran under conditions that have a tendency to hide concurrency related bugs in it – simulated here by letting each thread sleep for a small amount of time in its main loop (`use_slow_threads` is true by default). The `nanosleep` function moves a thread (and only this thread) into the BLOCKED state for some amount of (wallclock) time. When you run the program instead with the `-f` switch so that each thread will no longer have this sleep period, the program suddenly starts to fail.

Analyze the program for data races. Data races are defined in lecture; or more formally in the C11 memory model [URL] as follows:

When an evaluation of an expression writes to a memory location and another evaluation reads or modifies the same memory location, the expressions are said to conflict. A program that has two conflicting evaluations has a data race unless either

- both conflicting evaluations are atomic operations
- one of the conflicting evaluations *happens-before*¹ another

Which data races does this program contain, if any? Justify your answer, referring to line numbers in the code if this is necessary.

¹Here, *happens-before* refers to the cross-thread relationship between events, see Slide 11.

[Solution] Interestingly, this program does not contain data races because all accesses to shared variables are protected with a lock.

3. (3 pts) Discuss one way in which program fails and explain why it fails. If the failure is due to a data race, explain how the data race contributed to the failure. If the failure is due to another kind of concurrency related problem, describe the problem and how it contributed to the failure.

[Solution] Although the program does not contain data races, it contains other concurrency bugs. Here, there is an ordering violation: the programmer assumed that no thread will invoke `make_a_permutation` again until all threads have left this function (and the `permutation_index` variable has reached 0). If a thread that has returned from this function calls it again before this variable reaches zero, then this thread will increment it again, and the variable may reach n again even though not all threads have had a chance to output their character for the second permutation. Conversely, it is possible for threads to leave this function after a thread has looped around, decrementing the variable, which in turn prevents any thread from executing the if-branch on line 51, leading to a deadlock of all threads being blocked on line 56. You can observe this effect when running the program under `gdb`. Naming one of these 2 scenarios is sufficient.

4. (6 pts) Now fix the program, subject to the following constraints:

- You may make changes only to the marked areas in the `make_a_permutation` function, but you may also change (if necessary) the initial values of the 2 semaphores defined in the program. The section of code to which you may make changes are marked. Only 1 semaphore is currently used, you will need to use the second (marked “extra”) semaphore to make the program work.
- You may not introduce additional mutexes, semaphores, or condition variables, or use other POSIX APIs not related to mutexes and semaphores.
- A correct program should output 10 permutations, separated by a newline, not only when run under `-f`, but under any permissible scheduling of the 10 threads.

[Solution]

It is possible to use the second semaphore in a similar manner as the first one to ensure that no thread leaves this function until `permutation_index` has reached zero. This “extra” semaphore acts as a second turnstile. It is signaled when `pi` reaches 0, then each waiting thread executes a wait/post sequence, letting out thread after thread. Since there will be an extra signal when the last thread leaves the function, this extra signal must be consumed for the next round; this can be done by the last thread to arrive (the thread that starts the process of letting all others threads through the first turnstile. To ensure this works the first time around, this semaphore can be initialized to 1.

```
1 // making random permutations
2 #include <pthread.h>
3 #include <semaphore.h>
4 #include <stdlib.h>
5 #include <stdio.h>
```

```

6  #include <assert.h>
7  #include <stdint.h>
8  #include <stdbool.h>
9  #include <string.h>
10 #include <unistd.h>
11
12 #define NTHREADS 10
13
14 static pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZER;
15 static sem_t permutation_done;
16 static sem_t permutation_extra;    // not currently used
17
18 // you may change these 2 lines below if necessary
19 const int PERMUTATION_DONE_INITIAL_VALUE = 0;
20 const int PERMUTATION_EXTRA_INITIAL_VALUE = 1;
21
22 static bool use_slow_threads;
23 typedef void (*func_t)(void);
24
25 /*
26  * This function is called by exactly n threads, repeatedly, as
27  * shown below in thread_fun.
28  * Each time it's called by these threads it should output a
29  * permutation of the string 'abcdefghij', following by a newline.
30  * Thread 1 must output 'a', thread 2 must output 'b', and so on
31  * Thread i must output 'a' + i.
32  *
33  * Exactly one thread should invoke the `when_permutation_finished`
34  * function when the permutation has been output and before the
35  * output of the next permutation has been started.
36  */
37 static void
38 make_a_permutation(int n, int id, func_t when_permutation_finished)
39 {
40     assert (0 <= id && id < n);
41     // you may not change this line
42     printf("%c", 'a' + id);
43
44     // you may make changes below in the remainder of this function
45     static int permutation_index = 0;
46
47     pthread_mutex_lock(&lock);
48     int pi = ++permutation_index;
49     pthread_mutex_unlock(&lock);
50
51     if (pi == n) {

```

```

52     // last thread to arrive
53     when_permutation_finished();
54     sem_wait(&permutation_extra); // consume extra signal
55     sem_post(&permutation_done); // let in next thread
56 }
57
58 sem_wait(&permutation_done); // wait for first signal
59 sem_post(&permutation_done); // let in next thread
60
61 pthread_mutex_lock(&lock);
62 pi = --permutation_index;
63 pthread_mutex_unlock(&lock);
64
65 if (pi == 0) {
66     // last thread to leave
67     sem_wait(&permutation_done); // consume extra signal
68     sem_post(&permutation_extra); // let out first thread
69 }
70
71 sem_wait(&permutation_extra); // wait to leave function
72 sem_post(&permutation_extra); // let out next thread
73 // you may not make changes below this line
74 }
75
76 static void
77 newline(void)
78 {
79     printf("\n");
80     fflush(stdout);
81 }
82
83 static void *
84 thread_fun(void *_arg)
85 {
86     int id = (uintptr_t) _arg;
87     for (int i = 0; i < 10; i++) {
88         if (use_slow_threads) {
89             // slow threads sleep a random amount of time
90             // drawn from [.5ms, 1.5ms]
91             struct timespec ts = {
92                 .tv_sec = 0,
93                 .tv_nsec = 500000 + rand() % 1000000
94             };
95             nanosleep(&ts, 0);
96         }
97         make_a_permutation(NTHREADS, id, newline);

```

```

98     }
99     return NULL;
100 }
101
102 int
103 main(int ac, char *av[])
104 {
105     pthread_t threads[NTHREADS];
106     sem_init(&permutation_done, 0, PERMUTATION_DONE_INITIAL_VALUE);
107     sem_init(&permutation_extra, 0, PERMUTATION_EXTRA_INITIAL_VALUE);
108
109     use_slow_threads = !(ac > 1 && !strcmp(av[1], "-f"));
110
111     for (int i = 0; i < NTHREADS; i++)
112         pthread_create(threads + i, NULL, thread_fun, (void *) (uintptr_t) i);
113
114     for (int i = 0; i < NTHREADS; i++)
115         pthread_join(threads[i], NULL);
116 }

```

4 A Deadlock Mystery (8 pts)

When Dr. Back tried to come up with a question for this exam, he experimented with a “mystery” program that he hoped would deadlock at some point. Eventually, he caught the program deadlocking in gdb. You can see the resulting debugging session below, with linebreaks inserted for readability.

```

1 GNU gdb (Ubuntu 8.2-0ubuntu1~18.04) 8.2
2 Copyright (C) 2018 Free Software Foundation, Inc.
3 License GPLv3+: GNU GPL version 3 or later <http://gnu.org/licenses/gpl.html>
4 This is free software: you are free to change and redistribute it.
5 There is NO WARRANTY, to the extent permitted by law.
6 Type "show copying" and "show warranty" for details.
7 This GDB was configured as "x86_64-linux-gnu".
8 Type "show configuration" for configuration details.
9 For bug reporting instructions, please see:
10 <http://www.gnu.org/software/gdb/bugs/>.
11 Find the GDB manual and other documentation resources online at:
12   <http://www.gnu.org/software/gdb/documentation/>.
13
14 For help, type "help".
15 Type "apropos word" to search for commands related to "word"...
16 Reading symbols from ./mystery...done.
17 (gdb) run
18 Starting program: /home/gback/cs3214/exams/mystery
19 [Thread debugging using libthread_db enabled]

```



```

20 Using host libthread_db library "/lib/x86_64-linux-gnu/libthread_db.so.1".
21 [New Thread 0x7ffff73b9700 (LWP 6087)]
22 [New Thread 0x7ffff6bb8700 (LWP 6088)]
23 [New Thread 0x7ffff63b7700 (LWP 6089)]
24 [New Thread 0x7ffff5bb6700 (LWP 6090)]
25 ^C
26 Thread 1 "mystery" received signal SIGINT, Interrupt.
27 0x00007ffff79b7d2d in __GI___pthread_timedjoin_ex (threadid=140737341265664,
28     thread_return=0x0, abstime=0x0, block=<optimized out>)
29     at pthread_join_common.c:89
30 89     pthread_join_common.c: No such file or directory.
31 (gdb) thread apply all backtrace
32
33 Thread 5 (Thread 0x7ffff5bb6700 (LWP 6090)):
34 #0  __lll_lock_wait () at ../sysdeps/unix/sysv/linux/x86_64/lowlevellock.S:135
35 #1  0x00007ffff79b9023 in __GI___pthread_mutex_lock (mutex=0x555555755080 <NW>)
36     at ../nptl/pthread_mutex_lock.c:78
37 #2  0x0000555555554973 in car (_info=0x3) at mystery.c:47
38 #3  0x00007ffff79b66db in start_thread (arg=0x7ffff5bb6700) at pthread_create.c:463
39 #4  0x00007ffff76dfa3f in clone () at ../sysdeps/unix/sysv/linux/x86_64/clone.S:95
40
41 Thread 4 (Thread 0x7ffff63b7700 (LWP 6089)):
42 #0  __lll_lock_wait () at ../sysdeps/unix/sysv/linux/x86_64/lowlevellock.S:135
43 #1  0x00007ffff79b9023 in __GI___pthread_mutex_lock (mutex=0x5555557550c0 <SE>)
44     at ../nptl/pthread_mutex_lock.c:78
45 #2  0x0000555555554930 in car (_info=0x2) at mystery.c:40
46 #3  0x00007ffff79b66db in start_thread (arg=0x7ffff63b7700) at pthread_create.c:463
47 #4  0x00007ffff76dfa3f in clone () at ../sysdeps/unix/sysv/linux/x86_64/clone.S:95
48
49 Thread 3 (Thread 0x7ffff6bb8700 (LWP 6088)):
50 #0  __lll_lock_wait () at ../sysdeps/unix/sysv/linux/x86_64/lowlevellock.S:135
51 #1  0x00007ffff79b9023 in __GI___pthread_mutex_lock (mutex=0x555555755100 <SW>)
52     at ../nptl/pthread_mutex_lock.c:78
53 #2  0x00005555555548ea in car (_info=0x1) at mystery.c:33
54 #3  0x00007ffff79b66db in start_thread (arg=0x7ffff6bb8700) at pthread_create.c:463
55 #4  0x00007ffff76dfa3f in clone () at ../sysdeps/unix/sysv/linux/x86_64/clone.S:95
56
57 Thread 2 (Thread 0x7ffff73b9700 (LWP 6087)):
58 #0  __lll_lock_wait () at ../sysdeps/unix/sysv/linux/x86_64/lowlevellock.S:135
59 #1  0x00007ffff79b9023 in __GI___pthread_mutex_lock (mutex=0x555555755040 <NE>)
60     at ../nptl/pthread_mutex_lock.c:78
61 #2  0x00005555555548a4 in car (_info=0x0) at mystery.c:26
62 #3  0x00007ffff79b66db in start_thread (arg=0x7ffff73b9700) at pthread_create.c:463
63 #4  0x00007ffff76dfa3f in clone () at ../sysdeps/unix/sysv/linux/x86_64/clone.S:95
64
65 Thread 1 (Thread 0x7ffff7fbfb80 (LWP 6082)):
66 #0  0x00007ffff79b7d2d in __GI___pthread_timedjoin_ex (threadid=140737341265664,
67     thread_return=0x0, abstime=0x0, block=<optimized out>)
68     at pthread_join_common.c:89
69 #1  0x0000555555554a7c in main () at mystery.c:66

```

```

70 (gdb) info threads
71   Id  Target Id                               Frame
72 * 1   Thread 0x7ffff7fbfb80 (LWP 6082) "mystery" 0x00007ffff79b7d2d
73       in __GI___pthread_timedjoin_ex (threadid=140737341265664, thread_return=0x0,
74       abstime=0x0, block=<optimized out>) at pthread_join_common.c:89
75   2   Thread 0x7ffff73b9700 (LWP 6087) "mystery" __lll_lock_wait ()
76       at ../sysdeps/unix/sysv/linux/x86_64/lowlevellock.S:135
77   3   Thread 0x7ffff6bb8700 (LWP 6088) "mystery" __lll_lock_wait ()
78       at ../sysdeps/unix/sysv/linux/x86_64/lowlevellock.S:135
79   4   Thread 0x7ffff63b7700 (LWP 6089) "mystery" __lll_lock_wait ()
80       at ../sysdeps/unix/sysv/linux/x86_64/lowlevellock.S:135
81   5   Thread 0x7ffff5bb6700 (LWP 6090) "mystery" __lll_lock_wait ()
82       at ../sysdeps/unix/sysv/linux/x86_64/lowlevellock.S:135
83 (gdb) p NE
84 $1 = pthread_mutex_t = {Type = Normal, Status = Acquired, possibly with waiters,
85       Owner ID = 6090, Robust = No, Shared = No, Protocol = None}
86 (gdb) p NW
87 $2 = pthread_mutex_t = {Type = Normal, Status = Acquired, possibly with waiters,
88       Owner ID = 6088, Robust = No, Shared = No, Protocol = None}
89 (gdb) p SW
90 $3 = pthread_mutex_t = {Type = Normal, Status = Acquired, possibly with waiters,
91       Owner ID = 6089, Robust = No, Shared = No, Protocol = None}
92 (gdb) p SE
93 $4 = pthread_mutex_t = {Type = Normal, Status = Acquired, possibly with waiters,
94       Owner ID = 6087, Robust = No, Shared = No, Protocol = None}

```

They say a picture says a thousand words, so to make your job easier, we ask that you draw a sketch that illustrates the cause of this deadlock.

[Solution]

The sketch we were looking for is the classic illustration of gridlock in an intersection, shown in Figure 1.

4 cars (threads) are in the blocked state. For instance, thread 6090 holds lock NE (“North East”) and is blocked in an attempt to acquire lock NW (“North West”), similarly for the others. The program used to produce this gdb output is shown below:

```

1  #include <pthread.h>
2  #include <stdio.h>
3  #include <unistd.h>
4  #include <stdint.h>
5  #include <stdlib.h>
6  #include <time.h>
7
8  enum Direction { NORTH, SOUTH, EAST, WEST };
9  pthread_mutex_t NE = PTHREAD_MUTEX_INITIALIZER;
10 pthread_mutex_t NW = PTHREAD_MUTEX_INITIALIZER;
11 pthread_mutex_t SE = PTHREAD_MUTEX_INITIALIZER;
12 pthread_mutex_t SW = PTHREAD_MUTEX_INITIALIZER;
13
14 static void *

```




Figure 1: Source: Stallings, *Operating Systems: Internals and Design Principles*

```
15 car(void * _info)
16 {
17     struct timespec req = {
18         .tv_sec = 0,
19         .tv_nsec = 1000
20     };
21     enum Direction dir = (enum Direction) _info;
22     for (int i = 0; i < 100000; i++) {
23         switch (dir) {
24             case NORTH:
25                 pthread_mutex_lock(&SE);
26                 pthread_mutex_lock(&NE);
27                 nanosleep(&req, NULL);
28                 pthread_mutex_unlock(&NE);
29                 pthread_mutex_unlock(&SE);
30                 break;
31             case SOUTH:
32                 pthread_mutex_lock(&NW);
33                 pthread_mutex_lock(&SW);
34                 nanosleep(&req, NULL);
35                 pthread_mutex_unlock(&SW);
36                 pthread_mutex_unlock(&NW);
37                 break;
38             case EAST:
39                 pthread_mutex_lock(&SW);
40                 pthread_mutex_lock(&SE);
41                 nanosleep(&req, NULL);
42                 pthread_mutex_unlock(&SE);
43                 pthread_mutex_unlock(&SW);
44                 break;
45             case WEST:
```

```

46     pthread_mutex_lock(&NE);
47     pthread_mutex_lock(&NW);
48     nanosleep(&req, NULL);
49     pthread_mutex_unlock(&NW);
50     pthread_mutex_unlock(&NE);
51     break;
52 }
53 }
54 }
55
56 int
57 main()
58 {
59     pthread_t t[4];
60     pthread_create(&t[0], NULL, car, (void *) NORTH);
61     pthread_create(&t[1], NULL, car, (void *) SOUTH);
62     pthread_create(&t[2], NULL, car, (void *) EAST);
63     pthread_create(&t[3], NULL, car, (void *) WEST);
64
65     for (int i = 0; i < 4; i++)
66         pthread_join(t[i], NULL);
67 }

```

5 Optimizing Parallel Performance (8 pts)

Using a visualization tool, you have obtained a timeline of the execution of a multi-threaded program you are currently optimizing, which can be seen Figure 2. Time moves from the top to the bottom, and different colors are used to denote phases of execution, as described in the caption.

Answer the following questions:

1. (2 pts) Estimate the overall CPU utilization for this workload in percent, where 100% would be full utilization of all 6 cores.

[Solution] The CPU utilization here is 82%, corresponding to the ratio of non-white areas to the area of the entire rectangle.

2. (2 pts) Identify and describe the effect that most affects CPU utilization for this example.

[Solution] This workload shows a so-called “convoy effect” - one thread (thread 5) holds a lock (the black lock) for long periods of time (relative to the time others are holding this lock), causing a convoy of other threads that also wish to acquire this lock to form behind it. When those threads try to acquire this lock they are blocked and the corresponding CPU cannot be used.

3. (2 pts) Which thread and/or lock should be optimized to improve performance, and how?

[Solution] A main optimization goal should be to reduce the time thread 5 holds the black lock (or alternatively, find a way that avoids that threads 2, 3, 4 also compete for this lock).

4. (2 pts) Could this program have deadlocked under a different scheduler?




Figure 2: Timeline showing the execution of 6 threads on 6 CPU cores. Each vertical stripe (or column) represents one thread. Time that is spent executing while holding no lock is shown in green. Time during which the CPU is idle is shown in white. Time spent holding a lock is shown in a color corresponding to the lock, either black or blue. None of the threads engage in I/O or move into the BLOCKED state for any reason other than waiting for a lock.

[Solution] From the information shown, no thread ever holds 2 locks, removing at least one of the necessary conditions for deadlock. Thus, deadlock cannot occur.

Briefly justify each answer.

6 Submission Requirements

Submit a tar file that contains the following files:

- `random_permutation.c` to answer Question 3 part 4. This program should compile.
- `mystery.pdf` to answer Question 4.
- `rivercrossing.c` to answer Question 2. This program should also compile.

- `answers.txt` with answers to Question 1, Questions 3 (parts 1 to 3), and Question 5.