

# CS 3214 Spring 2021 Test 2 Solutions

September 13, 2021

## Contents

<b>1 Async (9 pts)</b>	<b>2</b>
<b>2 Concurrency Bugs (15 pts)</b>	<b>7</b>
2.1 Reading Helgrind (6 pts) . . . . .	7
2.2 A Bug in Mozilla (9 pts) . . . . .	9
<b>3 Map Reduce, Threaded (12 pts)</b>	<b>12</b>
<b>4 Deadlock Detection (12 pts)</b>	<b>18</b>
<b>5 Are we there yet? Roadblocks to increased multithreading utilization. (12 pts)</b>	<b>26</b>

## 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.

## 1 Async (9 pts)

In the handout of project 2, we motivated the need for a fork-join threadpool by demonstrating what results when languages provide an `async` facility without also providing an underlying efficient threadpool implementation. We used this C++ program which performs a parallel sum computation:

---

```
1  #include <iostream>
2  #include <vector>
3  #include <algorithm>
4  #include <numeric>
5  #include <future>
6
7  template <typename RAITer>
8  int parallel_sum(RAITer beg, RAITer end)
9  {
10     auto len = std::distance(beg, end);
11     if (len < 1000)
12         return std::accumulate(beg, end, 0);
13
14     RAITer mid = beg + len/2;
15     auto handle = std::async(std::launch::async,
16                             parallel_sum<RAIter>, mid, end);
17     int sum = parallel_sum(beg, mid);
18     return sum + handle.get();
19 }
20
21 int main()
22 {
23     std::vector<int> v(100000, 1);
24     std::cout << "The sum is " << parallel_sum(v.begin(), v.end())
25               << '\n';
26 }
```

---

A C version of such a program was included in the tests accompanying p2 where it was executed by the threadpool you built. In this question, you are asked to explore how much complexity implementing a threadpool added by developing an equally naive `async` implementation in C. Specifically, you should implement the functions `async` and `get` below such that the resulting program computes a parallel sum in the same way the C++ program does, which is by devoting a new and separate thread to each asynchronous task. Note the restrictions on where you may add code.

Your implementation should be reliable and data race free. For the purposes of this question, you do not need to worry about situations in which creating a new thread may fail due to reaching a per-user limit.

---

```
1  #include <pthread.h>
```

```

2  #include <stddef.h>
3  #include <stdlib.h>
4  #include <stdio.h>
5
6  typedef int (*async_function_t)(void *);
7  /*****
8  // you may make changes below this line
9
10 struct future {
11     // implement this
12 };
13 /* add any additional functions, variables, and data structures here */
14
15 // obtain the result of this asynchronous computation
16 // if necessary, block until the result is available
17 // frees the provided future so no future accesses are
18 // possible
19 int get(struct future *f)
20 {
21     // implement this
22 }
23
24 // start this asynchronous function in a new thread, passing it args
25 // returns a dynamically allocate future to refer to the computation
26 struct future * async(async_function_t fun, void *args)
27 {
28     // implement this
29 }
30
31 // you may not make changes below this line
32 // -----
33
34 struct args {
35     int *beg, *end;
36 };
37
38 int parallel_sum(struct args *args)
39 {
40     int *beg = args->beg, *end = args->end;
41
42     ptrdiff_t len = end - beg;
43     if (len < 1000) {
44         int sum = 0;
45         while (beg < end)
46             sum += *beg++;
47         return sum;

```

```

48     }
49
50     int *mid = beg + len/2;
51     struct args rhalf = { .beg = mid, .end = end };
52     struct future * handle = async((async_function_t) parallel_sum, &rhalf);
53     struct args lhalf = { .beg = beg, .end = mid };
54     return parallel_sum(&lhalf) + get(handle);
55 }
56
57 #define N 1000000
58 int v[N];
59 int main()
60 {
61     for (int i = 0; i < N; i++)
62         v[i] = 1;
63
64     struct args whole = { .beg = v, .end = v + N };
65     printf("The sum is %d\n", parallel_sum(&whole));
66 }

```

---

**[Solution]** The solution involves starting a new thread that receives the future, executes it, and stores the result. This thread is joined in future\_get. The techniques used are very similar to those in project 2.

---

```

1  #include <pthread.h>
2  #include <stddef.h>
3  #include <stdlib.h>
4  #include <stdio.h>
5
6  typedef int (*async_function_t)(void *);
7  /*****
8  // you may make changes below this line
9
10 struct future {
11     pthread_t t;
12     async_function_t task;
13     void *args;
14     int result;
15 };
16
17 static void * run(void *arg)
18 {
19     struct future * f = arg;
20     f->result = f->task(f->args);
21     return NULL;
22 }

```

```

23
24 // obtain the result of this asynchronous computation
25 // if necessary, block until the result is available
26 // frees the provided future so no future accesses are
27 // possible
28 int get(struct future *f) {
29     pthread_join(f->t, NULL);
30     int rc = f->result;
31     free(f);
32     return rc;
33 }
34
35 // start this asynchronous function in a new thread, passing it args
36 // returns a dynamically allocate future to refer to the computation
37 struct future * async(async_function_t fun, void *args)
38 {
39     struct future *f = malloc(sizeof *f);
40     f->task = fun;
41     f->args = args;
42     pthread_create(&f->t, NULL, run, f);
43     return f;
44 }
45
46 // you may not make changes below this line
47 // -----
48
49 struct args {
50     int *beg, *end;
51 };
52
53 int parallel_sum(struct args *args)
54 {
55     int *beg = args->beg, *end = args->end;
56
57     ptrdiff_t len = end - beg;
58     if (len < 1000) {
59         int sum = 0;
60         while (beg < end)
61             sum += *beg++;
62         return sum;
63     }
64
65     int *mid = beg + len/2;
66     struct args rhalf = { .beg = mid, .end = end };
67     struct future * handle = async((async_function_t) parallel_sum, &rhalf);
68     struct args lhalf = { .beg = beg, .end = mid };

```

```
69     return parallel_sum(&lhalf) + get(handle);
70 }
71
72 #define N 1000000
73 int v[N];
74 int main()
75 {
76     for (int i = 0; i < N; i++)
77         v[i] = 1;
78
79     struct args whole = { .beg = v, .end = v + N };
80     printf("The sum is %d\n", parallel_sum(&whole));
81 }
```

---

## 2 Concurrency Bugs (15 pts)

### 2.1 Reading Helgrind (6 pts)

Your p2 team partner received the following output when using the Helgrind race detection checker on their program.

Explain the information you can gather from the output in your own words. What did your teammate likely do wrong? Be sure to translate and include all relevant characteristics (variable names, line numbers, etc.) in your explanation.

Here is the output Helgrind produced:

```
1 ==21540== Helgrind, a thread error detector
2 ==21540== Copyright (C) 2007-2017, and GNU GPL'd, by OpenWorks LLP et al.
3 ==21540== Using Valgrind-3.13.0 and LibVEX; rerun with -h for copyright info
4 ==21540== Command: ./helgrind
5 ==21540==
6 ==21540== ---Thread-Announcement-----
7 ==21540==
8 ==21540== Thread #3 was created
9 ==21540==   at 0x518470E: clone (clone.S:71)
10 ==21540==   by 0x4E4BEC4: create_thread (createthread.c:100)
11 ==21540==   by 0x4E4BEC4: pthread_create@@GLIBC_2.2.5 (pthread_create.c:797)
12 ==21540==   by 0x4C38A27: ??? (in /usr/lib/valgrind/vgpreload_helgrind-amd64-linux.so)
13 ==21540==   by 0x1089A2: main (helgrind.c:37)
14 ==21540==
15 ==21540== ---Thread-Announcement-----
16 ==21540==
17 ==21540== Thread #2 was created
18 ==21540==   at 0x518470E: clone (clone.S:71)
19 ==21540==   by 0x4E4BEC4: create_thread (createthread.c:100)
20 ==21540==   by 0x4E4BEC4: pthread_create@@GLIBC_2.2.5 (pthread_create.c:797)
21 ==21540==   by 0x4C38A27: ??? (in /usr/lib/valgrind/vgpreload_helgrind-amd64-linux.so)
22 ==21540==   by 0x10897F: main (helgrind.c:36)
23 ==21540==
24 ==21540== -----
25 ==21540==
26 ==21540== Lock at 0x309080 was first observed
27 ==21540==   at 0x4C3603C: ??? (in /usr/lib/valgrind/vgpreload_helgrind-amd64-linux.so)
28 ==21540==   by 0x1088A1: thread2 (helgrind.c:22)
29 ==21540==   by 0x4C38C26: ??? (in /usr/lib/valgrind/vgpreload_helgrind-amd64-linux.so)
30 ==21540==   by 0x4E4B6DA: start_thread (pthread_create.c:463)
31 ==21540==   by 0x518471E: clone (clone.S:95)
32 ==21540== Address 0x309080 is 0 bytes inside data symbol "lock2"
33 ==21540==
34 ==21540== Lock at 0x309040 was first observed
35 ==21540==   at 0x4C3603C: ??? (in /usr/lib/valgrind/vgpreload_helgrind-amd64-linux.so)
36 ==21540==   by 0x108849: thread1 (helgrind.c:11)
37 ==21540==   by 0x4C38C26: ??? (in /usr/lib/valgrind/vgpreload_helgrind-amd64-linux.so)
38 ==21540==   by 0x4E4B6DA: start_thread (pthread_create.c:463)
39 ==21540==   by 0x518471E: clone (clone.S:95)
40 ==21540== Address 0x309040 is 0 bytes inside data symbol "lock1"
41 ==21540==
42 ==21540== Possible data race during read of size 4 at 0x3090A8 by thread #3
43 ==21540== Locks held: 1, at address 0x309080
44 ==21540==   at 0x1088AF: thread2 (helgrind.c:24)
45 ==21540==   by 0x4C38C26: ??? (in /usr/lib/valgrind/vgpreload_helgrind-amd64-linux.so)
```

```

46 ==21540==    by 0x4E4B6DA: start_thread (pthread_create.c:463)
47 ==21540==    by 0x518471E: clone (clone.S:95)
48 ==21540==
49 ==21540== This conflicts with a previous write of size 4 by thread #2
50 ==21540== Locks held: 1, at address 0x309040
51 ==21540==    at 0x108860: thread1 (helgrind.c:13)
52 ==21540==    by 0x4C38C26: ??? (in /usr/lib/valgrind/vgpreload_helgrind-amd64-linux.so)
53 ==21540==    by 0x4E4B6DA: start_thread (pthread_create.c:463)
54 ==21540==    by 0x518471E: clone (clone.S:95)
55 ==21540== Address 0x3090a8 is 0 bytes inside data symbol "gvar"
56 ==21540==

```

---

**[Solution]** The Helgrind messages shows that 2 threads access a shared 32-bit variable `gvar` concurrently. These accesses were concurrent because the programmer failed to synchronize access to them with a lock that was held when both threads accessed the variable. Instead, Helgrind reports that Thread #3 held a mutex named `lock2` when accessing the variable on line `helgrind.c:24` whereas Thread #2 held a mutex named `lock1` when accessing the same variable on line 13.

Here is the program that was used to create this Helgrind output:

---

```

1  #include <pthread.h>
2  #include <stdio.h>
3
4  pthread_mutex_t lock1 = PTHREAD_MUTEX_INITIALIZER;
5  pthread_mutex_t lock2 = PTHREAD_MUTEX_INITIALIZER;
6
7  static void *
8  thread1(void * _tn)
9  {
10     int * shared = (int *) _tn;
11     pthread_mutex_lock(&lock1);
12     for (int i = 0; i < 1000000; i++)
13         (*shared)++;
14     pthread_mutex_unlock(&lock1);
15     return NULL;
16 }
17
18 static void *
19 thread2(void * _tn)
20 {
21     int * shared = (int *) _tn;
22     pthread_mutex_lock(&lock2);
23     for (int i = 0; i < 1000000; i++)
24         (*shared)++;
25     pthread_mutex_unlock(&lock2);
26     return NULL;
27 }
28
29 static int gvar;

```



```

30
31 int
32 main()
33 {
34     int N = 2;
35     pthread_t t[N];
36     pthread_create(t + 0, NULL, thread1, &gvar);
37     pthread_create(t + 1, NULL, thread2, &gvar);
38
39     for (int i = 0; i < N; i++)
40         pthread_join(t[i], NULL);
41
42     printf("shared = %d\n", gvar);
43     return 0;
44 }

```

---

## 2.2 A Bug in Mozilla (9 pts)

Concurrency bugs can arise when different threads interact. The following program contains such a concurrency bug. It was reconstructed from a bug that in the past affected systems such as Mozilla's codebase.

---

```

1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <stdbool.h>
4  #include <string.h>
5  #include <unistd.h>
6  #include <pthread.h>
7
8  typedef void *(*worker_func_t)(void *);
9
10 enum worker_state { PENDING, STARTED, FINISHED };
11 struct worker {
12     pthread_t thread_id;
13     enum worker_state state;
14 };
15
16 pthread_mutex_t glock = PTHREAD_MUTEX_INITIALIZER;    // protects worker1 and *worker1
17 struct worker *worker1;
18
19 struct worker *start_worker(worker_func_t wfun, void *wdata)
20 {
21     struct worker *w = malloc(sizeof *w);
22     w->state = PENDING;
23     pthread_create(&w->thread_id, NULL, wfun, wdata);

```

```

24     return w;
25 }
26
27 static void *
28 worker1_fun(void *_ )
29 {
30     pthread_mutex_lock(&glock);
31     worker1->state = STARTED;
32     pthread_mutex_unlock(&glock);
33
34     printf("worker working\n");
35
36     pthread_mutex_lock(&glock);
37     worker1->state = FINISHED;
38     pthread_mutex_unlock(&glock);
39     return NULL;
40 }
41
42 int
43 main()
44 {
45     struct worker * w = start_worker(worker1_fun, NULL);
46     pthread_mutex_lock(&glock);
47     worker1 = w;
48     pthread_mutex_unlock(&glock);
49     pthread_join(worker1->thread_id, NULL);
50     printf("all done\n");
51 }

```

---

1. (3 pts) Determine the bug and explain how the program would fail if the bug manifested itself.
2. (3 pts) Is the bug you found a data race? 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*<sup>1</sup> another

Justify your answer, referring to line numbers in the code if this is necessary.

3. (3 pts) Suggest a way to fix the bug.

---

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

**[Solution]**

1. The program does not guarantee that the global variable 'worker1' is set on line 47 before it is accessed on line 31. If line 31 executes before the assignment on line 47 it would result in a segmentation violation due to accessing a NULL pointer. (In fact, this can be observed when running this code under valgrind.)
2. This is not a data race since the mutex glock serializes these 2 operations. It is an ordering violation.
3. It can be fixed in a number of ways, for instance,
  - (a) switching lines 45 and 46, thus ensuring that line 31 executes after line 48.
  - (b) Moving 'wdata' into 'struct worker', passing 'w' instead of 'wdata' on line 23 and setting 'worker1' in 'worker1\_fun'.
  - (c) Using a semaphore (initialized with 0) that is waited for at the beginning of 'worker1\_fun' and signaled after the 'worker1 = w' assignment.

The original Mozilla bug was both a data race and an ordering violation and was reported in: *Learning from mistakes: a comprehensive study on real world concurrency bug characteristics* by Lu et al at ASPLOS 2008 (Figure 2). [URL]

### 3 Map Reduce, Threaded (12 pts)

In Test 1, you were asked to simulate a MapReduce network using inter-process communication (IPC), specifically Unix pipes. In this problem, we explore how to perform the same simulation using inter-thread communication using what we call “thread pipes.” A thread pipe, like a Unix pipe, supports three operations:

- `pipe_send`: sends a message of `MSG_SIZE` bytes through the pipe. Messages are indivisible (atomic).
- `pipe_receive`: receives the next message of `MSG_SIZE` from the pipe. If multiple threads call `pipe_receive`, any of them may receive the message.
- `pipe_initialize`: initializes a thread pipe object in place

Implement the functions `pipe_initialize`, `pipe_send`, and `pipe_receive` so that the following program works. *Note that the execution result will be the same as was observed in Test 1.*

---

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <stdbool.h>
4  #include <string.h>
5  #include <unistd.h>
6  #include <pthread.h>
7  #include <sys/wait.h>
8  #include "list.h"
9
10 #define THREADS 7
11 #define NPIPES 11
12 #define TOTAL 100
13 #define MSG_SIZE 20
14
15 /* you may not change anything above this line */
16
17 struct thread_pipe {
18     // implement this
19 };
20
21 // define any other data structures here
22
23 /* Initialize a thread pipe object in place */
24 int
25 pipe_initialize(struct thread_pipe *pipe)
26 {
27 }
28
29 /* Send a msg of size MSG_SIZE to a thread pipe
```

```

30  * The pipe does not need to impose a limit on the number of
31  * items currently contained in the pipe.
32  */
33  void
34  pipe_send(struct thread_pipe *pipe, char msg[MSG_SIZE])
35  {
36  }
37
38  /* Receive a msg of size MSG_SIZE to a thread pipe.
39  * Blocks until a msg is in the pipe.
40  */
41  void
42  pipe_receive(struct thread_pipe *pipe, char msg[MSG_SIZE])
43  {
44  }
45
46  /*
47  * You may not make any changes below this line
48  */
49  static struct thread_pipe pip[NPIPES]; //driver-A, A-B, A-C, A-D, B-E, C-E, C-F, D-F, E-G, F-G
50  static int a_to_b, a_to_c, a_to_d, c_to_e, c_to_f, total;
51
52  /* Read `total` messages from read end of pipe pair in `fromfd` and
53  * round-robin distribute them to pipes at `(tofd + i) % stride`
54  *
55  * For simplicity, we send always MSG_SIZE-byte messages.
56  */
57  static void process(const char *name, struct thread_pipe *from, struct thread_pipe *to,
58                    int total, int stride)
59  {
60      char msgi[MSG_SIZE], msgo[MSG_SIZE];
61
62      for (int count = 0; count < total; count++) {
63          pipe_receive(from, msgi);
64          snprintf(msgo, sizeof msgo, "%s%s", name, msgi);
65          pipe_send(&to[count%stride], msgo);
66      }
67  }
68
69  static int divceil(int x, int y)
70  {
71      return (x + y - 1) / y;
72  }
73
74  static void *task_a(void *)
75  {

```

```

76     process("A", &pip[0], &pip[1], total, 3);
77     return NULL;
78 }
79
80 static void *task_b(void *)
81 {
82     process("B", &pip[1], &pip[4], a_to_b, 1);
83     return NULL;
84 }
85
86 static void *task_c(void *)
87 {
88     process("C", &pip[2], &pip[5], a_to_c, 2);
89     return NULL;
90 }
91
92 static void *task_d(void *)
93 {
94     process("D", &pip[3], &pip[7], a_to_d, 1);
95     return NULL;
96 }
97
98 static void *task_e(void *)
99 {
100     process("E", &pip[4], &pip[8], a_to_b, 1);
101     process("E", &pip[5], &pip[8], c_to_e, 1);
102     return NULL;
103 }
104
105 static void *task_f(void *)
106 {
107     process("F", &pip[6], &pip[9], c_to_f, 1);
108     process("F", &pip[7], &pip[9], a_to_d, 1);
109     return NULL;
110 }
111
112 static void *task_g(void *)
113 {
114     process("G", &pip[8], &pip[10], a_to_b + c_to_e, 1);
115     process("G", &pip[9], &pip[10], c_to_f + a_to_d, 1);
116     return NULL;
117 }
118
119 typedef void * (*thread_func_t)(void *);
120 void shuffle_threads(thread_func_t *funcs, size_t n)
121 {

```

```

122     for (int i = 0; i < n - 1; i++) {
123         int j = i + rand() / (RAND_MAX / (n - i) + 1);
124         thread_func_t t = funcs[j];
125         funcs[j] = funcs[i];
126         funcs[i] = t;
127     }
128 }
129
130 int
131 main(int ac, char *av[])
132 {
133     // compute how many messages will be sent
134     total = ac > 1 ? atoi(av[1]) : TOTAL;
135     a_to_b = divceil(total, 3);
136     a_to_c = divceil(total-1, 3);
137     a_to_d = divceil(total-2, 3);
138     c_to_e = divceil(a_to_c, 2);
139     c_to_f = divceil(a_to_c-1, 2);
140
141     // creating all the pipes
142     for (int i = 0; i < NPIPES; i++)
143         if (pipe_initialize(&pip[i])<0) {
144             fprintf(stderr, "Pipe creation error\n");
145             return EXIT_FAILURE;
146         }
147
148     // let's create the threads next
149     pthread_t t[THREADS];
150     thread_func_t tasks[THREADS] = { task_a, task_b, task_c, task_d, task_e, task_f, task_g };
151
152     // for a bit of challenge, shuffle the order in which they are started
153     srand(ac > 2 ? atoi(av[2]) : time(NULL));
154     shuffle_threads(tasks, THREADS);
155
156     for (int i = 0; i < THREADS; i++)
157         pthread_create(t+i, NULL, tasks[i], NULL);
158
159     // send numbers to A thread
160     for (int i = 0; i < total; i++) {
161         char msg[MSG_SIZE];
162         snprintf(msg, sizeof msg, "%d\n", i);
163         pipe_send(&pip[0], msg);
164     }
165
166     // get numbers from last thread
167     for (int i = 0; i < total; i++) {

```

```

168     char msg[MSG_SIZE];
169     pipe_receive(&pip[10], msg);
170     printf("%s", msg);
171 }
172
173 // join threads
174 for (int i = 0; i < THREADS; i++)
175     pthread_join(t[i], NULL);
176 }

```

---

**[Solution]** The solution is a simple application of the producer/consumer queue idea, for instance implemented with a mutex to protect a list of messages and a condition variable to signal availability. We only show the required functions below.

---

```

1  struct thread_pipe {
2      pthread_mutex_t lock;
3      pthread_cond_t msg_available;
4      struct list items;
5  };
6
7  struct pipe_item {
8      char msg[MSG_SIZE];
9      struct list_elem elem;
10 };
11
12 /* Initialize a thread pipe object in place */
13 int
14 pipe_initialize(struct thread_pipe *pipe)
15 {
16     list_init(&pipe->items);
17     int rc = pthread_mutex_init(&pipe->lock, NULL);
18     return rc ? rc : pthread_cond_init(&pipe->msg_available, NULL);
19 }
20
21 /* Send a msg of size MSG_SIZE to a thread pipe
22  * The pipe does not need to impose a limit on the number of
23  * items currently contained in the pipe.
24  */
25 void
26 pipe_send(struct thread_pipe *pipe, char msg[MSG_SIZE])
27 {
28     struct pipe_item *item = malloc(sizeof *item);
29     pthread_mutex_lock(&pipe->lock);
30     list_push_back(&pipe->items, &item->elem);
31     memcpy(item->msg, msg, MSG_SIZE);
32     pthread_cond_signal(&pipe->msg_available);

```



```
33     pthread_mutex_unlock(&pipe->lock);
34 }
35
36 /* Receive a msg of size MSG_SIZE to a thread pipe.
37 * Blocks until a msg is in the pipe.
38 */
39 void
40 pipe_receive(struct thread_pipe *pipe, char msg[MSG_SIZE])
41 {
42     pthread_mutex_lock(&pipe->lock);
43     while (list_empty(&pipe->items))
44         pthread_cond_wait(&pipe->msg_available, &pipe->lock);
45
46     struct pipe_item * item = list_entry(list_pop_front(&pipe->items),
47                                         struct pipe_item, elem);
48     memcpy(msg, item->msg, MSG_SIZE);
49     free(item);
50     pthread_mutex_unlock(&pipe->lock);
51 }
```

---

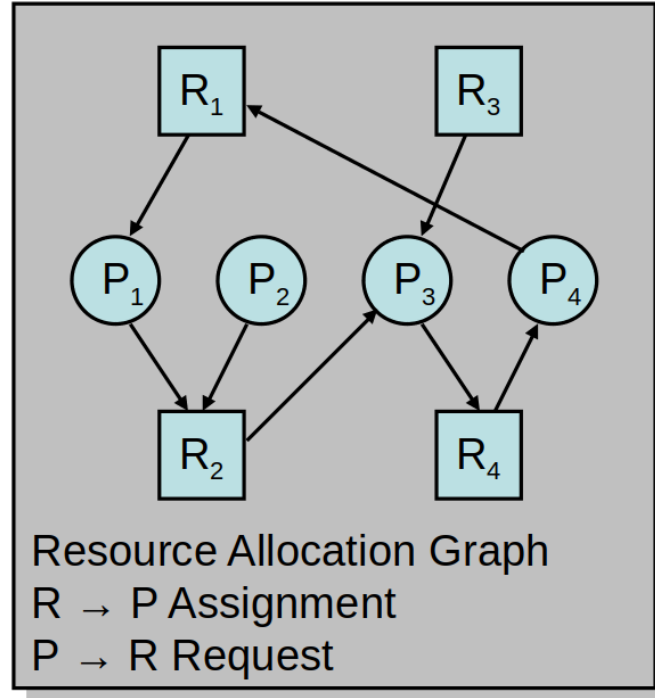


Figure 1: This resource allocation graph shows a deadlock situation that arose with 4 processes and 4 resources.

## 4 Deadlock Detection (12 pts)

Consider the resource allocation graph shown in Figure 1 which is reproduced from the lecture on deadlocks. In this problem, you are asked to completed a multi-threaded program that produces the exact deadlock shown in this graph and that can detect when it deadlocks.

In the following program the 4 involved processes P1, P2, P3, and P4 are represented as threads 1 through 4, executing functions P1fun, P2fun, P3fun, and P4fun, respectively. These threads (once implemented) should acquire resources such that the deadlock situation shown in Figure 1 arises. A barrier divides the actions of each thread in two sequential parts. All threads execute their second part after all threads have executed their first part. However, the threads execute their respective actions within each part concurrently with each other.

- (6 pts) Implement each thread such that the program deadlocks and the resource acquisition graph shown in the Figure results, no matter the order in which the scheduler may schedule those threads.
- (6 pts) Implement the function `check_for_deadlock` such that it detects a cycle in the resource acquisition graph when a process attempts to acquire a resource. Your deadlock detection algorithm should be general and not assume the specific graph shown in the Figure. Follow the required output exactly, which will be explained below.

---

```

1  #include <pthread.h>
2  #include <stdlib.h>
3  #include <stdio.h>
4  #include <unistd.h>
5  #include <semaphore.h>
6
7  struct Process;
8  static _Thread_local struct Process *current;
9
10 pthread_mutex_t glock = PTHREAD_MUTEX_INITIALIZER;
11
12 struct Resource {
13     struct Process *heldby; // process currently holding the resource
14     char *name;
15     sem_t sem;             // represents the resource
16 };
17
18 struct Process {
19     struct Resource *acquires; // resource this process is trying to acquire
20     pthread_t thread;
21     char *name;
22 };
23
24 /* Process `p` is about to request resource `r`.
25  * Check for deadlock.
26  */
27 static void
28 check_for_deadlock(struct Process *p, struct Resource *r)
29 {
30     printf("checking deadlock as process %s wants %s", p->name, r->name);
31     /* if the attempt by process p to acquire resource r would
32      * cause a deadlock, output this and then abort()
33         printf(" - found DEADLOCK\n");
34         abort();
35     */
36     printf(" - no deadlock\n");
37 }
38
39 void acquire(struct Resource *r)
40 {
41     pthread_mutex_lock(&glock);
42     check_for_deadlock(current, r);
43     if (r->heldby == NULL) {
44         r->heldby = current;
45         current->acquires = NULL;

```

```

46     } else {
47         current->acquires = r;
48     }
49     pthread_mutex_unlock(&glock);
50     if (current->acquires)
51         sem_wait(&r->sem);
52 }
53
54 struct Process P1 = { .name = "P1" },
55                    P2 = { .name = "P2" },
56                    P3 = { .name = "P3" },
57                    P4 = { .name = "P4" };
58
59 struct Resource R1 = { .name = "R1" },
60                    R2 = { .name = "R2" },
61                    R3 = { .name = "R3" },
62                    R4 = { .name = "R4" };
63
64 static pthread_barrier_t barrier;
65
66 static void * P1fun(void *process)
67 {
68     current = process;
69     // you may add code here
70     pthread_barrier_wait(&barrier);
71     // you may add code here
72     return NULL;
73 }
74
75 static void * P2fun(void *process)
76 {
77     current = process;
78     // you may add code here
79     pthread_barrier_wait(&barrier);
80     // you may add code here
81     return NULL;
82 }
83
84 static void * P3fun(void *process)
85 {
86     current = process;
87     // you may add code here
88     pthread_barrier_wait(&barrier);
89     // you may add code here
90     return NULL;
91 }

```

```

92
93 static void * P4fun(void *process)
94 {
95     current = process;
96     // you may add code here
97     pthread_barrier_wait(&barrier);
98     // you may add code here
99     return NULL;
100 }
101
102 int
103 main()
104 {
105     sem_init(&R1.sem, 0, 0);
106     sem_init(&R2.sem, 0, 0);
107     sem_init(&R3.sem, 0, 0);
108     sem_init(&R4.sem, 0, 0);
109     pthread_barrier_init(&barrier, NULL, 4);
110     pthread_create(&P1.thread, NULL, P1fun, &P1);
111     pthread_create(&P2.thread, NULL, P2fun, &P2);
112     pthread_create(&P3.thread, NULL, P3fun, &P3);
113     pthread_create(&P4.thread, NULL, P4fun, &P4);
114
115     pthread_join(P1.thread, NULL);
116     pthread_join(P2.thread, NULL);
117     pthread_join(P3.thread, NULL);
118     pthread_join(P4.thread, NULL);
119 }

```

---

Note that since the program is intended to function as an example of a system that is guaranteed to deadlock, a `release()` operation is not implemented.

The expected output of this program will be similar to the following:

```

checking deadlock as process ?? wants ?? - no deadlock
checking deadlock as process ?? wants ?? - no deadlock
...
checking deadlock as process ?? wants ?? - found DEADLOCK
Aborted (core dumped)

```

where the ?? are elided (since you need to find them as part of task 1). The output order does not have to be deterministic.

**[Solution]**

A possible output is shown here:

---

```

checking deadlock as process P1 wants R1 - no deadlock
checking deadlock as process P4 wants R4 - no deadlock
checking deadlock as process P3 wants R3 - no deadlock

```

```
checking deadlock as process P3 wants R2 - no deadlock
checking deadlock as process P3 wants R4 - no deadlock
checking deadlock as process P2 wants R2 - no deadlock
checking deadlock as process P1 wants R2 - no deadlock
checking deadlock as process P4 wants R1 - found DEADLOCK
Aborted (core dumped)
```

---

For part 1, the solution requires 2 phases: in phase 1, every process should acquire the resource that the resource allocation graph shows it holds (e.g., P1 holds R1, P3 holds R2 and R3, and P4 holds R3). In Phase 2, they will be trying to acquire the resource shown in the request edges. The last attempt will lead to detectable deadlock

For part 2, the deadlock detection can be performed with a simple loop following the request and assignment edges. A cycle would lead back to the process currently trying to acquire a resource. The complete solution is shown below:

---

```
1  #include <pthread.h>
2  #include <stdlib.h>
3  #include <stdio.h>
4  #include <unistd.h>
5  #include <semaphore.h>
6
7  struct Process;
8  static _Thread_local struct Process *current;
9
10 pthread_mutex_t glock = PTHREAD_MUTEX_INITIALIZER;
11
12 struct Resource {
13     struct Process *heldby; // process currently holding the resource
14     char *name;
15     sem_t sem;             // represents the resource
16 };
17
18 struct Process {
19     struct Resource *acquires; // resource this process is trying to acquire
20     pthread_t thread;
21     char *name;
22 };
23
24 /* Process `p` is about to request resource `r`.
25  * Check for deadlock.
26  */
27 static void
28 check_for_deadlock(struct Process *p, struct Resource *r)
29 {
30     printf("checking deadlock as process %s wants %s", p->name, r->name);
31     for (;;) {
```

```

32     struct Process *nextp = r->heldby;
33     if (nextp == p) {
34         printf(" - found DEADLOCK\n");
35         abort();
36     }
37     if (nextp != NULL && nextp->acquires) {
38         // printf(" and process %s wants %s", nextp->name, r->name);
39         r = nextp->acquires;
40     } else
41         break;
42 }
43 printf(" - no deadlock\n");
44 }
45
46 void acquire(struct Resource *r)
47 {
48     pthread_mutex_lock(&glock);
49     check_for_deadlock(current, r);
50     if (r->heldby == NULL) {
51         r->heldby = current;
52         current->acquires = NULL;
53     } else {
54         current->acquires = r;
55     }
56     pthread_mutex_unlock(&glock);
57     if (current->acquires)
58         sem_wait(&r->sem);
59 }
60
61 struct Process P1 = { .name = "P1" },
62                   P2 = { .name = "P2" },
63                   P3 = { .name = "P3" },
64                   P4 = { .name = "P4" };
65
66 struct Resource R1 = { .name = "R1" },
67                   R2 = { .name = "R2" },
68                   R3 = { .name = "R3" },
69                   R4 = { .name = "R4" };
70
71 static pthread_barrier_t barrier;
72
73 static void * P1fun(void *process)
74 {
75     current = process;
76     // you may add code here
77     acquire(&R1);

```

```

78     pthread_barrier_wait(&barrier);
79     // you may add code here
80     acquire(&R2);
81     return NULL;
82 }
83
84 static void * P2fun(void *process)
85 {
86     current = process;
87     // you may add code here
88     pthread_barrier_wait(&barrier);
89     // you may add code here
90     acquire(&R2);
91     return NULL;
92 }
93
94 static void * P3fun(void *process)
95 {
96     current = process;
97     // you may add code here
98     acquire(&R3);
99     acquire(&R2);
100    pthread_barrier_wait(&barrier);
101    // you may add code here
102    acquire(&R4);
103    return NULL;
104 }
105
106 static void * P4fun(void *process)
107 {
108     current = process;
109     // you may add code here
110     acquire(&R4);
111     pthread_barrier_wait(&barrier);
112     // you may add code here
113     acquire(&R1);
114     return NULL;
115 }
116
117 int
118 main()
119 {
120     sem_init(&R1.sem, 0, 0);
121     sem_init(&R2.sem, 0, 0);
122     sem_init(&R3.sem, 0, 0);
123     sem_init(&R4.sem, 0, 0);

```



```
124     pthread_barrier_init(&barrier, NULL, 4);
125     pthread_create(&P1.thread, NULL, P1fun, &P1);
126     pthread_create(&P2.thread, NULL, P2fun, &P2);
127     pthread_create(&P3.thread, NULL, P3fun, &P3);
128     pthread_create(&P4.thread, NULL, P4fun, &P4);
129
130     pthread_join(P1.thread, NULL);
131     pthread_join(P2.thread, NULL);
132     pthread_join(P3.thread, NULL);
133     pthread_join(P4.thread, NULL);
134 }
```

---

## 5 Are we there yet? Roadblocks to increased multithreading utilization. (12 pts)

Project 2 inspired a number of you to think about the power of multithreading and how it can affect the program that you write. Specifically, Eric H. posed a timely and interesting question about the current state of the usage of parallel programming, part of which I quote/paraphrase below.

Project 2 specification linked to a specific take by researchers from Berkeley who argue that “parallel programming models, software systems, and a supporting architecture” are key to our incorporation of parallel systems. However, the documentation also highlights the large number of software implementations that do not make use of the latest multi-core processors that have become more common in everyday technology.

**[Solution]** Some of the reasons/explanations are provided below as an example. We will be accepting a broad range of clearly-articulated answers for these questions.

1. (4 pts) What are some of the challenges in accomplishing the goal of widespread parallel programming? (Name at least two.)

**[Solution]** Reasons can range from difficulty of ensuring correctness in parallel programs, complex legacy codes, not every program/application has parallelism, emergence of low-power low-end processors such as in IoT devices, lack of easy-to-use tools, etc.

2. (4 pts) Why do you think there is a mismatch between the software and hardware advancements in parallel computing? (Name at least two reasons.)

**[Solution]** Reasons can range from lack of coordination between software and hardware researchers/developers, lack of powerful tools to exploit the hardware features, application needs, porting of legacy code, difficulty in training new people in hardware-software codesign, etc.

3. (4 pts) What steps do you believe must be taken to expedite the adoption of parallel computing systems?

**[Solution]** Better education and training, use of innovative ways to exploit parallelism such as function-as-a-service, development of new tools and techniques, novel applications, etc. can be some of the ways to expedite the adoption of parallel computing systems.

*Note: This is an open-ended question. Please provide brief answers with well-formed relevant arguments.*