

A common design pattern in systems programming is to use a thread pool to farm out the processing of incoming tasks and scale up to use as many processors as available. Each task is enqueued into a queue and processed by a worker thread.

**Question 1.** Does the program have any data races? If so, explain the data race and fix the code.

```

1 typedef struct Task {
2     void (*func)(void *);
3     void *arg;
4 } Task;
5
6 Task q[QUEUE_SIZE];
7 int in = 0, out = 0, count = 0;
8 bool exitRequested = false;
9 pthread_mutex_t m = PTHREAD_MUTEX_INITIALIZER;
10 pthread_cond_t c = PTHREAD_COND_INITIALIZER;
11
12 void
13 enqueue(void (*func)(void *), void *arg) {
14     Task *t = malloc(sizeof(*t));
15
16     t->func = func;
17     t->arg = arg;
18
19     q[in] = t;
20     in = (in + 1) % BUFFER_SIZE;
21     count++;
22
23     pthread_cond_signal(&c);
24
25 }
26
27 void
28 workerthread(void *ignored) {
29     for (;;) {
30         pthread_mutex_lock(&m);
31
32         if (count == 0)
33             pthread_cond_wait(&c, &m);
34         if (exitRequested)
35             pthread_exit(NULL);
36
37         Task *t = q[out];
38         out = (out + 1) % BUFFER_SIZE;
39         count--;
40         pthread_mutex_unlock(&m);
41
42         t->func(t->arg);
43
44         free(t);
45     }
46 }

```

**Question 2.** Are the condition variables used correctly? If not, explain the bug(s) and fix the code.

**Question 3.** Write a function to terminate the worker pool. Hint: You may need to fix the code to left.