



COMP2017 / COMP9017 Week 11 Tutorial

Synchronisation and Atomics

Question 1: Diagnosing deadlocks

There are four threads in the following program, labelled T_1 to T_4 , which are (amongst other operations) locking and unlocking three mutexes labelled A , B , and C . The order of locking is shown below:

$T_1 : lock(A), lock(B), lock(C), unlock(A), unlock(B), unlock(C)$

$T_2 : lock(A), lock(B), lock(C), unlock(C), unlock(B), unlock(A)$

$T_3 : lock(B), lock(C), unlock(B), unlock(C)$

$T_4 : lock(C), lock(A), unlock(A), unlock(C)$

1. Which threads have the potential to deadlock here?

Show two possible interleavings of instructions that cause at least two threads to go into a deadlock.

Question 2: Dining philosophers

Assume that there are N philosophers sitting at a round table. A single chopstick is placed between two adjacent philosophers. Every philosopher is either thinking or eating. However, a philosopher needs both chopsticks (to the left and to the right) to start eating. They are not allowed to acquire chopsticks that are not immediately adjacent to them. Complete the following program so that each philosopher is able to eat.

```
#include <stdio.h>
#include <stdlib.h>
#include <stdbool.h>
#include <pthread.h>

#define THINKERS 5

static pthread_mutex_t chopsticks[THINKERS];

void* dine(void* arg) {
```

```
const unsigned id = *((unsigned *) arg);

while (true) {
    // TODO: Acquire two chopsticks first
    // the ith philosopher can only reach
    // the ith and (i + 1)th chopstick
    printf("Philosopher %u is eating\n", id);
}

return NULL;
}

int main(void) {

    unsigned args[THINKERS];
    pthread_t thinkers[THINKERS];

    // create the chopsticks
    for (size_t i = 0; i < THINKERS; i++) {
        if (pthread_mutex_init(chopsticks + i, NULL) != 0) {
            perror("unable to initialize mutex");
            return 1;
        }
    }

    // launch threads
    for (size_t i = 0; i < THINKERS; i++) {
        args[i] = i;
        if (pthread_create(thinkers + i, NULL, dine, args + i) != 0) {
            perror("unable to create thread");
            return 1;
        }
    }

    // wait for threads to finish
    for (size_t i = 0; i < THINKERS; i++) {
        if (pthread_join(thinkers[i], NULL) != 0) {
            perror("unable to join thread");
            return 1;
        }
    }

    // remove the chopsticks
    for (size_t i = 0; i < THINKERS; i++) {
        if (pthread_mutex_destroy(chopsticks + i) != 0) {
            perror("unable to destroy mutex");
            return 1;
        }
    }
}
```

```
        }  
    }  
  
    return 0;  
}
```

Question 3: Semaphore philosophers

We have previously solved the dining philosophers problem by using a locking hierarchy. This time, use a semaphore for the table that only allows $N/2$ philosophers to eat at a time.

Question 4: Dancing threads

In following program uses threads to simulate dancers in a ballroom. Each thread is represents a person in the ballroom wearing a coloured dress. In the ballroom there is a main stage. However, dancers are only allowed onto the main stage if they are in a group of three. The second requirement to entering the main stage is that two of the dancers must be wearing a red dress and the other dancer must be wearing a white dress.

In the following program, each thread calls the function that corresponds to its colour after it spawns. i.e. a thread with a red dress calls the `red` function, a thread with a white dress calls the `white` function.

Your task is to add synchronisation code to `red` and `white` such that they block until all three can enter the main stage together, then the function should return. Suppose two red threads are blocked on `red`, and then a white thread calls `white`, the third thread should wake up the other two threads and they should all return.

```
#include <stdio.h>  
#include <pthread.h>  
  
#define NWHITE 1000  
#define NRED    (NWHITE * 2)  
  
void red(void);  
void white(void);  
  
void* red_thread(void* args) {  
    // TODO  
    red();  
    return NULL;  
}  
  
void* white_thread(void* args) {  
    // TODO
```

```
white();  
return NULL;  
}  
  
int main(void) {  
    pthread_t ids[NWHITE + NRED];  
  
    for (size_t i = 0; i < NWHITE; i++) {  
        for (size_t j = 0; j < 2; j++) {  
            pthread_create(ids + i * 3 + j, NULL, red_thread, NULL);  
        }  
  
        pthread_create(ids + i * 3 + 2, NULL, white_thread, NULL);  
    }  
  
    for (size_t i = 0; i < NWHITE + NRED; i++) {  
        pthread_join(ids[i], NULL);  
    }  
}
```

Question 5: Atomics (Extension)

In this exercise we will examine some ways to write concurrent programs using atomic instructions directly built into the processor. Like the SIMD instructions, they are exposed as compiler intrinsics. In `gcc`, they are available as part of the [atomic builtins](#). These builtins are now declared legacy with the advent of C11 atomics. However you need the latest versions of `gcc` and `clang` to use them with the new names defined in `stdatomic.h`. Since `clang` maintains compatibility with `gcc`, these builtins are also available for `clang`.

These basic building blocks allow us to create faster data structures and synchronisation mechanisms.

1. The following example compares the speed of atomic increments with a mutex to protecting the same global variable. Remove the comments around `#define USEATOMIC` and compare differences in performance. Try varying the thread count to something much larger the number of available cores. How does the mutex version scale in comparison to the atomics version.

```
#include <stdio.h>  
#include <pthread.h>  
  
#define NTHREADS 2  
#define REPEATS 1000000  
  
// #define USEATOMIC  
  
#ifndef USEATOMIC  
pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;  
#endif
```

```
unsigned counter = 0;

void* worker(void* arg) {

    for (unsigned i = 0; i < REPEATS; i++) {
        // NOTE: this code is here to benchmark the synchronisation
        // mechanism you should not attempt to atomically increment
        // a highly contended global counter if you can avoid it.
        // A much faster way would be save the results of the sum
        // on the stack and then add it atomically to the final value.

#ifdef USEATOMIC
        __sync_fetch_and_add(&counter, 1);
#else
        pthread_mutex_lock(&mutex);
        counter += 1;
        pthread_mutex_unlock(&mutex);
#endif
    }

    return NULL;
}

int main(void) {

    pthread_t thread_ids[NTHREADS];
    for (size_t i = 0; i < NTHREADS; i++) {
        pthread_create(thread_ids + i, NULL, worker, NULL);
    }

    for (size_t i = 0; i < NTHREADS; i++) {
        pthread_join(thread_ids[i], NULL);
    }

    printf("%u\n", counter);

    return 0;
}
```

2. The following code uses a barrier provided by pthread. Remove it and implement your own using the atomic builtins. You may need to use a spinlock in your implementation. Compare the performance.

```
#include <stdio.h>
```

```
#include <pthread.h>

#define NTHREADS 4
#define REPEATS 1000000

pthread_barrier_t barrier;

void* worker(void* arg) {

    for (unsigned i = 0; i < REPEATS; ++i) {
        pthread_barrier_wait(&barrier);
    }

    return NULL;
}

int main(void) {

    pthread_t thread_ids[NTHREADS];

    pthread_barrier_init(&barrier, NULL, NTHREADS);

    for (size_t i = 0; i < NTHREADS; ++i) {
        pthread_create(thread_ids + i, NULL, worker, NULL);
    }

    for (size_t i = 0; i < NTHREADS; ++i) {
        pthread_join(thread_ids[i], NULL);
    }

    pthread_barrier_destroy(&barrier);

    return 0;
}
```