CS-202 Exercises on Parallelism & Concurrency (L18)

Exercise 1: Parallelism vs. Concurrency [Basic]

1. For each scenario below, explain whether it is an example of parallelism, concurrency, or neither. In some scenarios, there may be more than one correct explanation at different levels of detail (e.g., components A and B run *concurrently*, and component C runs *in parallel* to both).
    - Two cooks work at the same time in a kitchen preparing different dishes.
    - One cook prepares a dish while switching between chopping vegetables and stirring a pot.
    - A multi-core processor runs two threads simultaneously on separate cores.
    - A single-core CPU switches rapidly between two threads to handle user input and a background task.
    - A program performs multiple mathematical computations in a loop.
    - A server handles multiple client connections using asynchronous I/O.
    - A teacher grading one exam, then another, then another, one after the other.
    - A team of movers each carries boxes at the same time to speed up the move.

2. Consider the scenarios described below, composed of tasks lists with time estimates. For each scenario, explain how parallelism and concurrency can help. Calculate the end-to-end time for the scenario to be accomplished sequentially, with concurrency, and with parallelism (using as many people in parallel as you see fit).

Scenario 1: Preparing dinner

Task	Description	Estimated Time
A	Chop vegetables	10 minutes
B	Boil pasta	12 minutes
C	Set the table	5 minutes
D	Make salad dressing	7 minutes
E	Clean up prep area	6 minutes

Scenario 2: Moving out of an apartment

Task	Description	Estimated Time
A	Pack clothes	15 minutes
B	Disassemble furniture	20 minutes
C	Load truck	25 minutes
D	Sweep floors	10 minutes
E	Return keys	5 minutes

Scenario 3: Office Morning Routine

Task	Description	Estimated Time
A	Unlock doors and disable alarm	3 minutes
B	Brew coffee using the machine	20 minutes
C	Boot all computers	7 minutes
D	Check voicemails and emails	10 minutes
E	Refill printer paper	5 minutes

Exercise 2: Variables [Advanced]

Consider the following code that reverses sentences in place.

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

#define MAX_SENTENCE_LEN 256
#define NUM_SENTENCES 3

typedef struct {
   char* sentence;
} ThreadArg;

// Runs in each worker thread
void* reverse_sentences(void* arg) {
   ThreadArg* data = (ThreadArg*)arg;
   char* sentence = data->sentence;
   int sentence_len = strlen(sentence);

   for (int i = 0; i < sentence_len / 2; i++) {
       char temp = sentence[i];
       sentence[i] = sentence[sentence_len - i - 1];
       sentence[sentence_len - i - 1] = temp;
   }

   return NULL;
}

// Runs in the main thread
int main() {
   const char* original_sentences[NUM_SENTENCES] = {
       "Liechtenstein is a long word to reverse",
       "A bad beginning makes a bad ending",
       "A dog! A panic in a pagoda!"
   };

   pthread_t threads[NUM_SENTENCES];
   ThreadArg args[NUM_SENTENCES];

   // Allocate and prepare each sentence
   for (int i = 0; i < NUM_SENTENCES; i++) {
       args[i].sentence = malloc(MAX_SENTENCE_LEN);
       if (args[i].sentence == NULL) {
           (...)
       }
       strncpy(args[i].sentence, original_sentences[i], MAX_SENTENCE_LEN - 1);
       args[i].sentence[MAX_SENTENCE_LEN - 1] = '\0';

       if (pthread_create(&threads[i], NULL, reverse_sentences, &args[i]) != 0) {
           (...)
       }
   }

   // Wait for all threads to finish
   for (int i = 0; i < NUM_SENTENCES; i++) {
```

```
        pthread_join(threads[i], NULL);
    }

    for (int i = 0; i < NUM_SENTENCES; i++) {
        printf("%s\n", args[i].sentence);
        free(args[i].sentence);
    }

    return 0;
}
```

For each of the variables listed below, say where they stay in memory, and which thread (main or worker thread) is their owner.

`original_sentences`,`threads`, `args`, `args[0].sentence`, `sentence`, `temp`.

Main

Stack	Heap	Text/data segment

Thread

Stack	Heap	Text/data segment

Exercise 3: Data races - shared counter [Advanced]

Consider the following program, which uses 5 threads to increment a shared counter. Identify all values of the counter the program may print. For each such value, write down a concrete execution schedule that produces the value.

```
#include <stdio.h>
#include <pthread.h>
int counter = 0;

void *incr(void *arg) {
  counter = counter + 1;
  return NULL;
}

int main(int argc, char *argv[]) {
  pthread_t threads[5];
  // Create two threads T1 and T2
  for (int i=0; i < 5; i++) {
     pthread_create(&threads[i], NULL, incr, NULL);
  }
  for (int i=0; i < 5; i++) {
     pthread_join(threads[i], NULL);
  }
  printf("Counter: %d\n", counter);
  return 0;
}
```

Exercise 4: Data races - dining philosophers [Advanced]

Remember the dining philosophers problem from the lecture.

The dining philosophers problem

- Philosophers think, eat, think, eat…
- They need both forks to eat
- Should not try to grab the same fork simultaneously

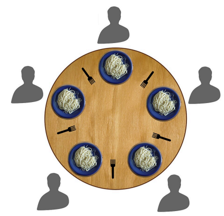

And consider the following bogus “solution”, presented as the first candidate during the lecture.

```
/* problem setup */

void philosopher(size_t i) {
   while (true) {            // repeat forever
       think(i);
       take_forks(i);        // acquire two forks
       eat(i);               // yum-yum, spaghetti
       put_forks(i);         // put both forks back
   }
}

int main() {
   pthread_t t1, t2, t3, t4, t5;

   pthread_create(&t1, NULL, philosopher, 1);
   pthread_create(&t2, NULL, philosopher, 2);
   ...
   pthread_join(t1, NULL);
   pthread_join(t2, NULL);
   ...
}
```

```
/* solution */

bool fork_taken[5] = {false};

void take_forks(i) {
   while (fork_taken[i] || fork_taken[(i + 1) % 5])
{
       // wait until both forks are free
   }
   // take fork to the left
   fork_taken[i] = true;
   // take fork to the right
   fork_taken[(i + 1) % 5] = true;
}

void put_forks(i) {
   // put down the fork to the left
   fork_taken[i] = false;
   // put down fork to the right
   fork_taken[(i + 1) % 5] = false;
}
```

1. Explain, in a few sentences, what causes a data race in this solution, and what problem the data race may lead to.

2. Data races are problematic when a thread gets interrupted at the “wrong time”, and the program executes with a schedule that leads to an unintended result. Write down a concrete problematic execution schedule for the threads executing philosopher 1 and philosopher 2**.** (You can express the execution schedule using a table with a column per thread, where each row corresponds to one atomic operation, like in slides 29-34).