

CSCI-GA.3033-017 Special Topic: Multicore Programming

Homework 2

Due October 6, 2016

Please solve the following and upload your solutions to your private GitHub repository for the class as homework2.pdf by 11:59pm on the due date above. If for some reason this poses a technical problem, or you wish to include diagrams that you don't wish to spend time drawing in a drawing application, you may hand in a printed copy (*not* hand-written) at the beginning of class (7:10pm) on the day of the deadline. **Unlike labs, late homeworks will be assigned a grade of 0.**

This homework takes the place of mandatory readings for this week, and give you some extra practice thinking about synchronization and thread safety. It is intended to help you hone the skills you need for Lab 1.

1. Warm up: Running on some specific machine X with set hardware and software, how many threads should your program spawn, given that it has enough parallel tasks to run to occupy many threads? Think about how the hardware (eg, memory and number of cores) and software (ie, programs already occupying some resources) affect the number of threads you should spawn. Don't overthink this one!

2. If multiple on-die caches (eg, L2 and L3 cache) are higher-latency than an L1 cache, why do we have them anyway? (Hint: how big and slow is main memory? How big and slow is the L1 cache?)

3. Consider the following code:

```
1   static int sum_stat_a = 0;
2   static int sum_stat_b = 0;
3   int aggregateStats(int stat_a, int stat_b) {
4       sum_stat_a += stat_a;
5       sum_stat_b += stat_b;
6       return sum_stat_a + sum_stat_b;
7   }
8   void init(void) { }
```

Use a single pthread mutex to make this function thread-safe. Add global variables and content to the `init()` function as necessary.

4. Let's make this more parallelizable. We always want to reduce critical sections as much as possible to minimize the time threads need to wait for a resource protected by a lock. Modify the original code from question 3 to make it thread-safe, but use two mutices this time, one for `sum_stat_a` and one for `sum_stat_b`.

5. What's the difference between `join()`ing and `detach()`ing a pthread? When would we want to use each technique?