

Lab 3: Finish Accumulators and NQueens Problem

Instructor: Mack Joyner

Course Wiki: <http://comp322.rice.edu>

Staff Email: comp322-staff@mailman.rice.edu

HJlib Documentation: <https://wiki.rice.edu/confluence/display/PARPROG/API+Documentation>

Goals for this lab

- Implement a parallel version of NQueens using Finish Accumulators, and evaluate its performance using abstract metrics (`Lab3NQueensAbstractMetricsCorrectnessTest`)
- Extend the implementation for improved real performance (`Lab3NQueensHjPerformanceTest`) by using a “cutoff” strategy.

Downloads

As with previous labs, the provided template project is accessible through your private GitHub repo at <https://classroom.github.com/a/XKe2q2HQ>

The below instructions will assume that you have already checked out the lab3 folder, and that you have imported it as a Maven Project if you are using IntelliJ.

1 The N-Queens Problem

The lecture on Finish Accumulators introduced the N-Queens problem: how can we place N queens on an $N \times N$ chessboard so that no two queens can capture each other? This problem was also presented in the demonstration video for topic 2.3.

The lab code already contains a sequential implementation for solving the N-Queens problem (`NQueensSequential`). The first goal of this lab is to create a parallel version of an N-Queens solver using HJlib and finish-accumulators. These changes should be made in the `NQueensHjLib` class, primarily in the `nqueensKernel` method. The sequential solution has been inserted into `NQueensHjLib`, with TODOs specifying the steps to transform it to a parallel solution. Your code for this section should be tested by only running `Lab3NQueensAbstractMetricsCorrectnessTest`.

The implementation of the `okToPlace()` method in the `NQueensRunner` class performs one call to `doWork(1)` for each pair of squares that is tested, to track abstract metrics.

2 A Note on Real Performance and the Cutoff Strategy

2.1 Real Performance

In this lab, we move from abstract metrics to real multi-threaded performance for the first time. Real performance is messier than abstract performance because it is affected by its environment. Many students will have varying hardware, varying software, and possibly other applications running at the same time when doing this lab on their laptops. Even your laptop’s power manager can throttle the number of hardware cores your HJlib program gets to use, limiting your speedup when not plugged into a wall socket.

```
void foo (int depth) { // Assume WORK decreases as depth increases
    if (n == LIMIT ) {
        // Process base case
    }

    if (n < CUTOFF) { // PARALLEL VERSION
        . . .
        async foo(n+1);
        . . .
    }
    else { // SEQUENTIAL VERSION
        . . .
        foo(n+1);
        . . .
    }
}
```

Figure 1: Code schema for parallel divide-and-conquer algorithm with cutoff

For today’s lab, you will see what real-world speedup you can achieve on your local laptop. If you do not see any speedup in real performance (i.e., if you cannot pass `Lab3NQueensHjPerformanceTest`), you can try closing down any other expensive applications that might be running, or plugging your laptop into a power supply. You will not be penalized if your code fails `Lab3NQueensHjPerformanceTest` on your laptop due to environmental factors.

2.2 Cutoff Strategy

A common way to reduce the overheads seen in real performance when creating large numbers of tasks (each of which does very little work) is for the programmer to add a *cutoff test*. Figure 1 shows a code schema for a parallel divide-and-conquer algorithm with a `CUTOFF` value. While conceptually `async`s do not do any actual application work, the creation of an `async` still consumes both CPU and memory resources. As a result, while creating an excessive amount of `async`s might maximize the abstract parallelism of your application, it may actually lead to your code running slower than a sequential implementation.

In this section, we will use the cutoff strategy to implement a more efficient parallel implementation that can also pass the real world performance tests in `Lab3NQueensHjPerformanceTest`. We can implement the cutoff strategy by saying that for any `depth` greater than or equal to a certain cutoff, we will run the remaining computation sequentially. You can access a recommended cutoff value for the current test by calling `getCutoff()` inside of `NQueensHjLib`, or you can use a suitable constant that you choose as the cutoff.

Once you have implemented the cutoff strategy, try re-running the tests in `Lab3NQueensHjPerformanceTest` on your laptop to see if the results have improved from your initial, maximally parallel solution.

3 Demonstrating and submitting in your lab work

For this lab, you will need to demonstrate and submit your work, as follows.

1. Show your work to an instructor or TA to get credit for this lab. You have until Thursday, March 4th at 2pm to get your lab checked off.
2. Commit, push, and check that all the work for today’s lab is in your `lab3` directory by opening <https://classroom.github.com/a/XKe2q2HQ> in your web browser and checking that your changes

have appeared.