

Lab 5: Data-Driven Tasks

Instructor: Mack Joyner, Zoran Budimlić

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

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

Goals for this lab

- Experimentation with Data-Driven Tasks with Abstract Metrics

Lab

Today, we will use a Cholesky factorization example to learn about using Data Driven Tasks. We will compute performance for this part of the lab using abstract metrics.

The Maven project for this lab is located in the following GitHub classroom repository:

<https://classroom.github.com/a/wrqqv0vU>

For instructions on checking out this repository through IntelliJ or through the command-line, please see the Lab 1 handout. The below instructions will assume that you have already checked out the lab5 folder, and that you have imported it as a Maven Project if you are using IntelliJ.

1 Parallelization of Cholesky (using Data-Driven Tasks)

In linear algebra, the Cholesky factorization is a decomposition of a positive-definite matrix into the product of a lower triangular matrix and its conjugate transpose. This decomposition is useful to efficiently obtain numerical solutions, e.g., for a system of linear equations of the form $A \times x = b$.

We have provided a sequential version of Cholesky Factorization in the `CholeskyFactorization.java` and `CholeskyFactorizationSequential.java` files. Your goal will be to parallelize this computation using data-driven tasks and data-driven futures, and to evaluate your parallelization using abstract metrics. The fact that the sequential version uses a generic `dataStore` container can be leveraged to simplify this conversion.

Your assignment is to create a parallel version of `CholeskyFactorizationSequential.runComputation()` that implements the Cholesky Factorization algorithm by using data-driven tasks:

1. Write a parallel version of the `runComputation()` method of the `CholeskyFactorizationParallel.java` file using data-driven tasks with calls to `asyncAwait()`. The provided file has helpful hints to guide you in this process.
2. For abstract metrics, note that `CholeskyFactorizationParallel.java` is already set up with zero overhead for creating tasks.
3. Run the unit tests locally to verify whether your parallel computation achieves the required abstract metrics.

2 Turning in your lab work

For lab5, you will need to turn in your work before Wednesday, Feb 16, 2022 at 4:30 pm, as follows:

1. Show your tests passing locally to an instructor or TA to get credit for this lab.
2. Commit and push your work to your `lab5` turnin folder. The only changes that must be committed are your modifications to `CholeskyFactorizationParallel.java`. Check that all the work for today's lab is in your `lab5` directory by opening <https://classroom.github.com/a/wrqqv0vU> in your web browser and checking that your changes have appeared.