



Programming Assignment 1 Experiments with Matrix Addition

Assigned Date

1/23/2018

Due Date

1/30/2018

Overview

- Task: Add two $N \times N$ matrices using CUDA
- What should be done:
 - Dynamic size for the matrices: dynamically allocated memory and add keyboard input statements to specify N
 - Should be LARGE matrices (at least one million)
 - You need to compare with and without striding
 - Different CUDA grid/block structures and sizes – Add keyboard statements to input different values for numbers of threads in a block and number of blocks in a grid
 - Include checks for invalid input
 - Timing -- Add statements to time the execution of the code using CUDA events, both for the host-only (CPU) computation and with the device (GPU) computation, and display results.
 - Compute and graph the appropriate metrics (runtime, speed-up factor, throughput...).

Deliverables

- Two parts:
 - Report:
 - Results: multiple timings of runs of various sizes
 - Appropriate graphs
 - Code **ONLY GITHUB, don't use any library:**

- Sequential C part
- Cuda part
- Same repository and output comparing results
- Have a pdf of your report emailed to Fred Harris, Lee Barford, and Rui Wu.