Spring 2019: Advanced Topics in Numerical Analysis: High Performance Computing Assignment 4 (due Apr. 15, 2019)

Handing in your homework: Hand in your homework as for the previous homework assignments (git repo with Makefile), answering the questions by adding a text or a $\[Mathbb{L}^{T}EX\]$ file to your repo.

- 1. **Matrix-vector operations on a GPU.** Modify the CUDA code from class for reduction to implement an inner product between two given (long) vectors on a GPU. Then, generalize this code to implement a matrix-vector multiplication (no blocking needed here) on the GPU. Check the correctness of your implementation by performing the same computation on the CPU and compare them. Report the memory band your code obtains on different GPUs.¹
- 2. **2D Jacobi method on a GPU.** Implement the 2D Jacobi method as discussed in the 2nd homework assignment using CUDA. The convolution example we'll be showing in the class on April 8 will be helpful for that. Monitor the performance on different GPUs. *Extra Credit:* Implement the Gauss-Seidel smoothing with red-black coloring in CUDA and report the performance.
- 3. Pitch your final project. Summarize your current plan for the final project. Detail *what* you are planning to do, and with *whom* you will be cooperating. The preferred size of final project teams is two, but if this makes sense in terms of the size of the project, teams of three or doing a project by yourself is fine as well. Each team is expected to give a 10 minute presentation about the problem they have worked on and their results and experience during the finals' week (likely May 20 and 21) and hand in a report as well as their code in a repo. We assume you have already talked to us about your project ideas when this homework is due, so a short summary is sufficient. We will discuss the project expectation during the next class, and are available for discussions on April 9, 5-6pm and April 11, 11-12:30 in office #1111, and over Slack. We are also posting a list of example final projects, but you are encouraged to work on a project that is motivated by your own research. Note that we will request frequent updates on the progress you are making over the next weeks and will help if you get stuck.

¹The cuda{1-5}.cims.nyu.edu compute servers at the Institute have different Nvidia GPUs, for an overview see the list of compute servers available at the Institute: https://cims.nyu.edu/webapps/content/systems/ resources/computeservers.