HPC Project2: High Performance Conjugate Gradient

Matthew Fricke

April 2024

1 Project Assistant Meetings

The class project assistant is Ryan Scherbarth. All teams must meet with Ryan at least once before the deadline.

You should also make use of Nicholas Bacon and my office hours.

2 Introduction

In this project you will run the High Performance Conjugate Gradient benchmark (HPCG). This project will allow you to explore GPU vs CPU benchmarking in a multinode environment.
You have three goals:

1. Explore Ahmdal's scaling of HPCG on CPUs (keep problem size fixed and change the number of CPUs).)
2. Explore Gustafson's scaling for HPCG (change the problem size with the number of CPUs).
3. Compare the performance of the CPU-only and GPU accelerated HPCG code as a function of problem size (dimensions of the cube in hpcg.dat). It's fine to include multi-node results.

You have all the tools needed to accomplish this. You have seen how to use the HPCG NVidia Singularity image. You have seen and used slurm job arrays and a template file to perform a parameter sweep (you may use GNU parallel instead if you choose), and you know how to load modules such as the HPCG module to get access to the CPU-only version.

NOTE: Using template files and job arrays will enable you to perform lots of experiments quickly. If you have trouble with the templates, then it might be a good idea to create your parameter sweep by handcoding several Slurm batch scripts and gather less data. Submitting plots with less data is better than getting stuck on adapting the HPL template code to HPCG.

2.1 High Level Overview of HPCG

The following is mostly taken from Jack Dongarra's HPL Lab page: HPL.

The High Performance Conjugate Gradients (HPCG) Benchmark project is an effort to create a new metric for ranking HPC systems. HPCG is intended as a complement to the High Performance LINPACK (HPL) benchmark, currently used to rank the TOP500 computing systems. The computational and data access patterns of HPL are still representative of some important scalable applications, but not all. HPCG is designed to exercise computational and data access patterns that more closely match a different and broad set of important applications, and to give incentive to computer system designers to invest in capabilities that will have impact on the collective performance of these applications. HPCG Benchmark

Note how much simpler this is that our explanation of HPL in Project 1.

3 HPCG Overview

4 Reading

This paper discusses HPCG especially in comparison to HPL. Use it to get an overview of what HPCG is used for. (A full understanding of the mathematical details, while interesting, is not required for this project. Don't get bogged down in them.) HPCG Benchmark: a New Metric for Ranking High Performance Computing Systems

5 Report Format

Your report will be no more than 3 pages. You should illustrate your observations with plots like those from the homeworks and project 1. In particular create three plots to illustrate the work you did to achieve the three goals described above.

I will share a report template with each team.

6 Extra Credit

The performance of each team on the HPCG benchmark will be ranked. Members of the team with the highest verified mean performance across the CPU and GPU benchmarks will receive extra credit for the *final grade of the course*. The second ranked team will receive 2 points, and the 3rd ranked team will receive 1 point.