SC22 Intramural Competition of Southeast University

April 4, 2022

1 Competition Introduction

The SC22 Intramural Competition of Southeast University uses three types of HPC benchmark tests as the problems, namely HPL, HPCG and IO-500, from which participants can choose any one of them. The contestants are encouraged to do more than one test, and extra points will be given to those who complete extra tests, but the weights of the three problems are different from each other in our evaluation criteria.

Your task is to **install the benchmark program, run it, and adjust it to the best performance.** Contestants will be required to submit a proposal including the test plan, process and results, as well as the hardware and software configuration used. Additional points will be awarded for the full English Proposal.

2 Submission Requirements

- Deadline: 2022/4/22, 23:59
- Submit to: [redacted]@seu.edu.cn
- Email subject: SC22-studentID-YourName
- File format: "SC22-studentID-YourName.pdf", for example:"SC22-12345678-张三.pdf"

3 Problems

3.1 HPL

Introduction HPL is a software package that solves a (random) dense linear system in double precision (64 bits) arithmetic on distributed-memory computers. It can thus be regarded as a portable as well as freely available implementation of the High Performance Computing Linpack Benchmark.

The algorithm used by HPL can be summarized by the following keywords: Two-dimensional block-cyclic data distribution - Right-looking variant of the LU factorization with row partial pivoting featuring multiple look-ahead depths - Recursive panel factorization with pivot search and column broadcast combined - Various virtual panel broadcast topologies - bandwidth reducing swap-broadcast algorithm - backward substitution with look-ahead of depth 1.

The HPL package provides a testing and timing program to quantify the accuracy of the obtained solution as well as the time it took to compute it. The best performance achievable by this software on your system depends on a large variety of factors. Nonetheless, with some restrictive assumptions on the interconnection network, the algorithm described here and its attached implementation are scalable in the sense that their parallel efficiency is maintained constant with respect to the per processor memory usage.

The HPL software package requires the availability on your system of an implementation of the Message Passing Interface MPI (1.1 compliant). An implementation of either the Basic Linear Algebra Subprograms BLAS or the Vector Signal Image Processing Library VSIPL is also needed. Machine-specific as well as generic implementations of MPI, the BLAS and VSIPL are available for a large variety of systems.

Download Address https://netlib.org/benchmark/hpl/hpl-2.3.tar.gz

3.2 HPCG

Introduction 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 is a complete, stand-alone code that measures the performance of basic operations in a unified code:

- Sparse matrix-vector multiplication.
- Vector updates.
- Global dot products.
- Local symmetric Gauss-Seidel smoother.
- Sparse triangular solve (as part of the Gauss-Seidel smoother).
- Driven by multigrid preconditioned conjugate gradient algorithm that exercises the key kernels on a nested set of coarse grids.
- Reference implementation is written in C++ with MPI and OpenMP support.

Download Address https://github.com/hpcg-benchmark/hpcg/

3.3 IO-500

Introduction Benchmarking of HPC storage systems is a complex task. Parallel I/O is not only influenced by CPU performance for latency and the networking stack but also on the underlying storage technology and software stack. With the IO-500, we have defined a comprehensive benchmark suite that enables comparison of high-performance storage systems. Similar to the TOP500 list for compute architectures, IO500 will allow tracking performance growth over the years and analyze changes in the storage landscape. The IO-500 will not only provide key metrics of expected performance, but serve as a repository for foster- ing sharing best practices within the community.

Download Address https://github.com/IO500/io500

4 Additional Requests

The proposal should include descriptions of the software environment, (operating system, compiler, math library, MPI software, software version, etc.), the performance optimization and testing methods, performance measurement, problem and solution analysis, etc. In-depth analysis on HPL algorithms and the respective source codes would be a plus. It would be great if you could edit the code to improve the performance.

It is recommended to run verification and optimization of HPL and HPCG benchmarks on GPU platforms, not only on CPU. If other hardware platforms are used, you are welcomed to submit the related analysis and results that demonstrate adequate performance.

You can also use other binary packages of the three benchmarks instead of those provided by the download address above, which may result in higher performance. If you do so, please clearly describe the version of the benchmark and where you get it.

Note that there is no hardware limit as long as you can run the benchmarks. Please attach the edited code to the Appendix of your proposal, if any.

It's okay if you submit a proposal with none of the benchmark being done. Life is full of surprise, isn't it? :)