Team Phoenix: Georgia Institute of Technology

Albert Chen, Aman Jain, Evan Montoya, Marissa Sorkin, Nicole Prindle, Ryan Elliott, Vijay Thakkar

Team Members

Marissa Sorkin
- Cloud Compute, UX/UI, Devices
- SC20 Veteran
- IO500, Cardioid

Nicole Prindle
- Computing Theory
- SC20 Veteran
- Reproducibility, System Admin

Evan Montoya
- Algorithms, ML
- New Member
- HPL, HPCG

Ryan Elliott
- ModSim, HPC
- New Member
- HPL, HPCG

Albert Chen
- HPC, ML
- SC20 Veteran
- IO500, Cardioid

Aman Jain
- HPC SysArch
- SC20 Veteran
- Quantum Espresso

About The Team

University Support
- Georgia Tech is a leader in computing education with a strong SC presence, backed by past SCC experience.
- We attend specialized courses for HPC, including a course to prepare specifically for student cluster competitions.
- We have world class faculty helping us prepare, comprising of Dr. Richard Vuduc, Dr. Aaron Jezghani, Dr. Jeff Young, and Will Powell.

Diversity and Inclusion
- Our experiences include outreach programs that promote STEM to underrepresented youths and fostering inclusion for Women and LGBTQ groups in CS.
- To promote a diverse SCC team composition for the future, we plan to collaborate with different computing clubs with similar values of diversity. We are also targeting students with interdisciplinary interests that may overlap with HPC.

Hardware Configuration

- 4 Dual-Socket Intel Ice Lake Xeon Platinum 32-core 8352Y Nodes
- 128 GB memory per socket in an 8 x 16 GB configuration for higher throughput; a total of 1TB system memory
- 2 x NVIDIA A100 80GB PCIe GPU per node
- 1 x 1.6 TB Intel Optane NVMe per node
- 2 x 100 GB/s OmniPath NIC per node

Hardware Motivation

- An Intel-based system allows us to take advantage of the Intel oneAPI toolkit to optimize CPU performance: ICC for compilation, MKL for BLAS implementation, AVX-512 , and IMPI for MPI .
- NVIDIA A100 will be used for GPU acceleration on Quantum Espresso, HPL, and HPCG.
- Intel Optane NVMe provides 7GB/s sequential write speed and metadata capabilities that will benefit the IO500 benchmark.

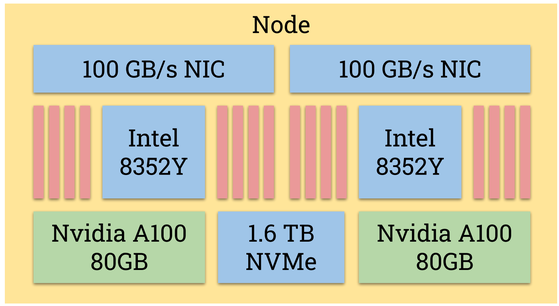

Power Management

- We wanted a Y-series SKU to use Intel SpeedSelect Technology, a command-line tool that gives us granular control over core clock frequency and power draw
- NVIDIA SMI is used to control clock frequency and power draw for GPUs
- Our experiments will investigate the power draw of the system under different loads. Using these results, we can manage power during runtime through shell scripts accordingly.

Software Configuration

- CentOS 8: Familiarity from our practice clusters, Phoenix and PACE ICE
- Slurm: Ease of troubleshooting relative to PBS, as discovered at VSCC20
- Spack: Commonly used package manager
- Singularity: Supports the use of NGC containers optimized for Quantum Espresso, HPL, and HPCG
- GekkoFS: Distributed file system to optimize IO500
- GCC/MVAPICH2: Broad support for mystery application
- Grafana: System monitoring

Preparation

Competition Preparation
- Georgia Tech Class for student cluster competition preparation where we gained knowledge of HPC concepts
- Trained on Georgia Tech clusters for benchmarks and competition apps
- Weekly meetings for discussions on progress, roadblocks, and strategies

Application/Benchmark Preparation
- Used GT HPC clusters, such as PACE-ICE or our personal Team Phoenix cluster to tune benchmarks and applications.
- Used NGC Containers running on Singularity for HPL, HPCG, Quantum Espresso

Mystery App Preparation
- Surveyed past SCC mystery apps
- Reviewed our approach of previous VSCC20 mystery application experience

Overall Assessment

- The top of the line GPUs ensures a good power-to-performance ratio for applications that benefit from GPU acceleration, while also allowing us to control power consumption when needed.
- The Intel Ice Lake Xeon Platinum CPUs offer great performance and an optimized toolkit of libraries with oneAPI
- Potential bottlenecks: memory bandwidth, drive io speeds