Team Gig'Em Bytes

High Performance Research Computing
DIVISION OF RESEARCH

Patrick Quinn, Curran Watson, Ashwin Kundeti,
Catherine Larson, Luis Diaz, Kenton Romero, Dr. Lisa M. Pérez

SC22 Dallas, TX | hpc accelerates.

Meet the Team!

Patrick Quinn
Major: Applied Math
Skills: Statistics, Cloud Computing
Roles: Team Lead, PHASTA, Linpack

Curran Watson
Major: Chemistry and Physics
Skills: Machine Learning, Parallel Computing
Roles: LAMMPS, HPCG

Ashwin Kundeti
Major: Computer Engineering
Skills: Computer Architecture, Software Development
Roles: IO500, Reproducibility, Mystery

Catherine Larson
Major: Computational Physics
Skills: Software Development
Roles: HPCG, LAMMPS

Luis Diaz
Major: Electrical Engineering
Skills: Computer Architecture, Scripting
Roles: Linpack, IO500, Hardware

Kenton Romero
Major: Electrical & Computer Engineering
Skills: Parallel Computing, Network Design
Roles: Reproducibility, PHASTA

Team Overview

The Gig'Em Bytes team possesses diversity in both the members' areas of study and cultural backgrounds. The team represents different areas of STEM, from computer & electrical engineering to physics and chemistry to statistics, which allows the team as a whole to tackle problems from any STEM discipline. The team's diversity allows for multiple perspectives and ways of thinking, a diversity in perspectives that will cultivate the flow of ideas and problem-solving methods that is perfect for the SCC. Each team member will work on specific applications with at least one overlap so that there is always one other team member that can be referenced when questions or issues arise. In addition to the 6 team members competing, we have an extended team of undergraduate members helping with programs such as LAMMPS, PHASTA and Hardware configurations.

Applications & Benchmarks

LAMMPS

An open-source molecular dynamics code that takes full advantage of A100 GPU acceleration. We have been given the amazing opportunity to receive guidance and support from experts in LAMMPS at Texas A&M's High Performance Research Computing (HPRC). We have access to LAMMPS benchmarking input files allowing us to practice executing and tuning LAMMPS jobs on a wide variety of simulations.

PHASTA

A finite element method fluid solver used in fluid dynamics computations. In many cases it will use a version of the Navier-Stokes equation to calculate the sum of forces on areas dissected by a structured or unstructured mesh. We plan on working with students in the Oceanography department to find and create real-world test cases to familiarize ourselves with the CPU centric PHASTA.

HPCG

A benchmark intended to measure the performance of High Performance Computing (HPC) systems. HPCG will benefit from the 128 GB of DDR4 memory per node and the Rockport interconnect.

Linpack

Linpack solves a dense system of linear equations to measure a computer's floating-point of execution. Through HPRC's system grace, we were able to compile and run Linpack on A100 nodes to familiarize ourselves with the code and learn to optimize the runs by modifying the system size parameters.

IO500

IO500 is a benchmarking software that was introduced in order to compare storage systems. It functions by measuring both the performance of the I/O bandwidth with Interleaved or Random (IOR) and metadata with Mdtest. The Rockport ANM will help us profile and tune our network and BeeGFS filesystem parameters for an optimal IO500 run.

Reproducibility

A productive, portable, and performant python framework across FPGA, CPU, GPU and Piz Daint supercomputer. We will learn about DaCe using the Jupyter Notebook tutorials provided on the DaCe github page and how to achieve optimization with the code on CPU and GPU.

Mystery

Our hardware configuration is prepared for a CPU or GPU centric application. If needed, the mystery application can be sent to the cloud resource.

Software Configuration

OS: Centos 7

Software Management: Spack and EasyBuild

Job Scheduling: slurm

File Management: BeeGFS

Hardware Configuration

Hardware Configuration

5 Dell R7525 servers, each with:

- Dual AMD 7502, 2.5GHz, 32C/64T
- 1 NVIDIA Ampere A100 80GB
- 128GB DDR4
- 480GB SSD
- 1.6GB NVMe

Network: Rockport Switchless Network with Autonomous Network Manager (ANM)

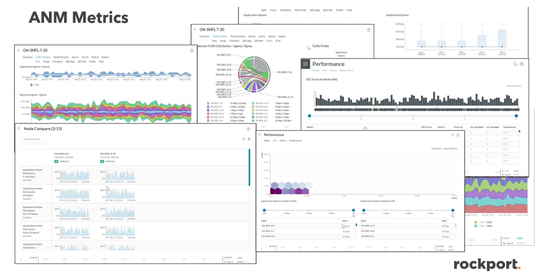

Hardware Motivation

Rockport Switchless Networks provides an amazing low power solution. Power consumption is lower than common HPC networks due to the switch being passive while the power is only consumed by the FPGA card. This network card has 12 bi-directional network connections, uses rockport Network Operating Software (rNOS), with a consumption of 34W at max load. We are working with the vendor to decrease the power usage based on our hardware configuration. The AMD Epyc 7502 CPU provides 32 cores and 64 threads, is extremely power efficient, and features 16MB of L3 cache per 4 cores. The NVIDIA A100 GPU cards provide the best balance of power consumption and performance with support for IEEE-compliant FP64 processing and a large 40MB L2 cache catered for HPC. SSD and NVMe storage offer high performance and low power over traditional spinning disk.

Preparation & Strategy

Overview

The team has access to system administrators and user support specialists at HPRC and attends HPRC training session. The team meets weekly and completes application specific assignments on the HPRC systems.

Time Management

Our team is made up of members with diverse academic skill sets and strong learning abilities so we plan to each specialize on an aspect of the competition so that during the competition, we can all efficiently work. Smart scheduling of jobs via the slurm queueing system will allow the team to focus on problem solving and organization while jobs run. Checkpointing is critical in the event that there is an unexpected interruption or need to kill a running job.

Optimization Strategy

Since we have access to A100 nodes on HPRC system, we have compiled applications and benchmarks with various flags to maximize the performance. We also benchmark various system sizes to use the resources efficiently.

Dynamic Power Limit

To control power usage, the team will utilize the nvidia-smi tool to limit the power consumption of the A100 GPUs. For the AMD Epyc CPUs we will make use of the cpupower utility to set the cpu frequency. Our use of Rockport will conserve an immense amount of energy within the cluster due to its switchless fiber optic design. We will have cron jobs in place to monitor total and node level power consumption and automatically modify cpu and gpu power usage when needed.

Cloud Resources

We plan to use the cloud resource to offload calculations during low power limit periods. We plan on using our TAMU student credits to gain familiarity with Azure CycleCloud so that we can readily utilize any Cloud credits we receive during the competition.

Sponsors

Dell Technologies | NVIDIA | High Performance Research Computing, DIVISION OF RESEARCH | rockport.