Team Gig 'Em Bytes

High Performance Research Computing
DIVISION OF RESEARCH

Lucian Chauvin, Tanya Chandra, Aakash Jain
Mihir Kalvakaalva, Andrii Kryvenko, and Yuzu Shen

SC25 St. Louis, MO | hpc ignites.

Meet the Team!

Tanya Chandra
Major(s): Electrical Engineering, Physics
Skills: Physics, Mathematics, Python
Roles: SST, Mystery

Lucian Chauvin
Major(s): Mathematics, Computer Science
Skills: Linux, Bash, Git
Roles: Team Lead

Aakash Jain
Major: Aerospace Engineering, Mathematics
Skills: Software Development, Optimization, Linux
Roles: Reproducibility, Exascale Climate Emulator

Mihir Kalvakaalva
Major: Materials Science and Engineering
Skills: Statistics, Machine Learning
Roles: Reproducibility, SST

Andrii Kryvenko
Major: Computer Science
Skills: Machine Learning, Software Development, Linux

Yuuzen Shen (Ioezu)
Major: Computer Science
Skills: NLP, LLM, ML, Distributed Computing
Roles: MLPerf, Mystery

Team Overview

The Gig 'Em Bytes team is a diverse and multidisciplinary set of students from different academical, cultural, and professional backgrounds. The team represents different areas of STEM, from Computer Science, Mathematics, Physics, and Engineering. The team's wide-ranging experiences allows for unique approaches to HPC problems as well as a technical prowess necessary to excel. GIven a multidisciplinary approach, each team member works on multiple applications, allowing for Gig 'Em Byte's strength in teamwork to shine through. Our members have engaged in researches and internships in computing-heavy fields, and ready for what the Student Cluster Competition offers.

Applications & Benchmarks

Exascale Climate Emulator

The Exascale Climate Emulator statistically reproduces the outputs of high-resolution climate models, drastically reducing computational and storage costs while maintaining consistency with full simulations.

Structural Simulation Toolkit (SST)

SST is a tool kit enabling the design of next generation computing. It is a discrete-element simulator, and coupling multiple layers of architecture in a fully modular framework. SST supports MPI and multithreading, but notably does not make use of GPU acceleration.

Performance Benchmarks

We evaluate our cluster using three established benchmarks to assess its different computational capabilities.

- **HPL (High Performance Linpack)** is a standard and widely used benchmark that evaluates dense linear algebra performance through LU decomposition, providing insight into sustained computational performance of a cluster.
- **HPL-MxP** assesses the system's ability to perform mixed-precision linear algebra computations. These use lower-precision arithmetic for computational speed while maintaining solution accuracy through refinement, enabling dramatic speed ups while keeping high precision for sensitive operations.
- **MLPerf Inference** measures machine learning inference performance across diverse model architectures and data types, providing insights into real-world ML deployment scenarios.

By utilizing HPRC's resources, we were able to compile and run these benchmarks to gain experience with their implementations and optimizing their performance.

Reproducibility

We reproduced the results of the SC24 paper, "Asynchronous Distributed-Memory Parallel Algorithms for Influence Maximization". As this paper focuses on massive Distributed Memory applications, a test on our 2 node architecture offers a good worst case test to verify the paper's speed up results on a small scale.

Mystery Application and Benchmark

Based on IndySCC 21's Devito and SC24's cat recognition, we decided to emphasis problems focused on compiler, machine learning, and computer vision frameworks. The team established training plans to get familiar with multiple open source resource-heavy applications and performance profiling for them.

Software Configuration

OS: Rocky Linux 9.6

Software Management: Spack and EasyBuild

Job Scheduling: Slurm

File Management: rsync and BeeOND

Hardware Configuration

Hardware Configuration

2 Dell PowerEdge XE7745 nodes, each with:

- 2 AMD EPYC 9555 3.20GHz, 64C/128T, 350 W TDP per CPU
- 4 NVIDIA H200 NVL GPUs with NVLink, 141 GB HBM3e per GPU
- 1.5 TB DDR5-6400 RDIMM memory (24 × 64 GB)
- 2 × 480 GB M.2 NVMe (BOSS-N1, RAID 1) for OS/boot
- 1.6 TB U.2 NVMe data drive

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

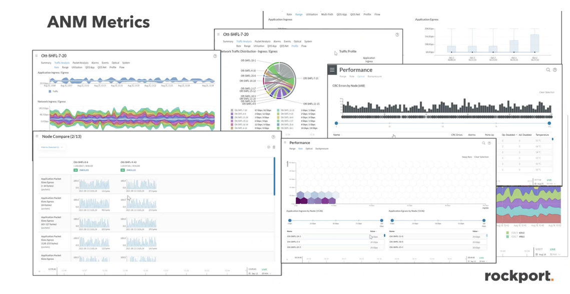

Hardware Motivation

We opted for a high core count, GPU-heavy system with two nodes that fits within the power restrictions when run efficiently. This configuration reduces load on the parallel storage system and improves overall usability. The abundance of GPUs per node also enables the team to be more prepared for Machine Learning and Computer Vision tasks. The Rockport Switchless Network by Cerio 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.

Preparation & Strategy

Overview

The team has access to HPC clusters such as ACES so they can gain experience working with and running applications on clusters. The team has bi-weekly meetings, and receives support from the team advisor, Dr. Lisa M. Pérez.

Time Management

Our team has a diverse set of technical skills, allowing members to specialize on particular tasks. Smart job scheduling using slurm is used to coordinate running computations on the cluster, and enables focusing on problem solving rather than monitoring running jobs. We have also assigned delegated tasks so each member is responsible for a specific application, benchmark, or system task, reducing overlap.

Optimization Strategy

The team has access to HPC clusters like ACES, FASTER and Grace with H100 and A100 GPUs and Launch with AMD CPUs, enabling optimization efforts for both applications and benchmarks prior to having access to the hardware. Optimization with the competition hardware starts from these configurations, with additional run flags and compilation options tried for the specific hardware and to account for power, noise, and other constraints.

Power Limit

The maximum power draw of cluster is well over the competition maximum of 10000W. To limit this power draw, the nvidia-smi tool is used to limit the power consumption of the H200 GPUs. The AMD CPUs power draw will be controlled using the cpupower utility for highly parallel tasks. 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.

Sponsors

DELL Technologies · High Performance Research Computing, Division of Research · NVIDIA · AMD · cerio