



University of Illinois Student Cluster Team




Jonathan Nativ, Nishant Sheikh, Pranshu Chaturvedi, Freddy Zhang, Rittika Adhikari, Liam Xu

Configurations

Cloud Configuration Overview

- We plan to create **6 clusters with SLURM** to address the following applications: HPL & HPCG, IO-500, GROMACS, CESM, MemXCT CPU and MemXCT GPU.
- Each cluster configuration was decided based on our experimentation with the cloud and scaled to a **\$1250** competition budget.
- One cluster will be reused for the mystery application.

Performance Scaling Comparison



Architectures and Budgeting Estimates



Applications

HPL & HPCG

These benchmarks measure the total computational performance of a cluster. We used **Nvidia binaries** and the most **GPU dense** nodes with Remote Direct Memory Access (**RDMA**) available in order to maximize score while minimizing cost.

IO-500

A holistic benchmark that tests a HPC cluster's I/O capability. We will use **ZFS** which is a free high performance parallel file system to connect the 10 NVMe SSDs in our L80s_v2 node. We decided to use a **single node** for this benchmark in order to **eliminate the overhead of using an RDMA protocol** which would increase latency and overall cost.

GROMACS

GROMACS is a popular open-source molecular dynamics package primarily used to dynamically simulate organic molecules. We chose to use 2 GPU dense nodes as this application shows **strong scaling with GPUs**. Having 2 nodes also gives us the option to run simulations with different inputs simultaneously. We found **OpenMPI** yields the best performance and we are finding that enabling **different runtime flags** significantly alters performance as well.

CESM


CESM is a fully-coupled, community, global climate model that provides state-of-the-art computer simulations of the Earth's past, present, and future climate patterns. As CESM is only runs on the CPU, we chose a CPU core dense AMD powered node which offers the the most cores per dollar at a similar IPC to alternative offerings. Again we found **OpenMPI** to yield the best performance.

Reproducibility Challenge (MemXCT)

MemXCT is a memory-centric X-ray CT reconstruction software repository that dramatically speeds up reconstruction by storing intermediate calculations in memory rather than recomputing results. As we are required to run both the CPU and GPU versions of the repo, we chose an AMD cluster that has a **high CPU cache**, **RDMA**, and a **large amount of RAM** in addition to a GPU dense cluster.

Mystery Challenge

We plan to reuse **1** of the 6 clusters for this challenge so that we can carry over our expertise from working with familiar architecture. In anticipation for a machine learning application, we have gained experience in installing **Tensorflow-GPU** and **PyTorch-GPU**.



Strategy

Competition Preparation

- Thoroughly study applications and profile them to understand performance bottlenecks.
- Create multiple different clusters with different architectures and topologies to maximize performance per dollar on the cloud.
- Create a budget breakdown in advance describing how funds will be allocated in a **worst case scenario** in order to avoid going over budget.

Team Recruitment

- Create team that represents wide range of expertise and backgrounds.
- Recruit understudy members so they can participate in competition preparation and learn and grow into possible future participants.

Vendor Sponsorship

- Team solicited sponsorship from industry leaders.
- Worked with NCSA to identify and approach potential sponsors.
- Worked with sponsors to create our software stack and troubleshoot problems with applications.

Cost Management

- Make use of **checkpointing** in applications that support it like CESM in case a low priority VM gets shut down or we need to switch to a smaller cluster.
- Make use of **spot instances** in case we are comfortably below budget.
- Create **shell scripts** to aid in budget monitoring and in providing alerts.

Software

- CentOS 7.6:** Widely used OS for cluster computing.
- CUDA 11.1:** Development environment for high performance GPU-accelerated applications.
- SLURM:** Widely-used open source cluster management and job scheduling.
- Shell Scripts:** We created shell scripts to monitor cost on each cluster and give real time notifications when approaching limits.
- TensorFlow and PyTorch:** Machine Learning frameworks in anticipation of the mystery challenge
- IPMI and Nvidia SMI:** Tools for measuring and controlling our cluster's power usage.
- OpenMPI 4.0.5:** Open-source high performance Message Passing Interface (MPI) Library.
- GCC:** GNU Compiler Collection for compiling applications
- ZFS:** High performance free parallel file system
- Arm Forge and Performance Reports:** Application profiling tools

Why We Will Win

Our access to the academic and technical resources provided by the University of Illinois' nationally acclaimed Computer Science and Computer Engineering departments along with support from the National Center for Supercomputing Applications ensures that the configuration and utilization of the cloud is as optimized as possible. In addition, our team's diversity and inclusion of understudy members allows us to address all of the competition's requirements to the best of our abilities and bring innovative solutions to both cloud and application specific problems.



Diversity and Collaboration

Meet the team

Pranshu Chaturvedi - Sophomore in CS and Statistics

- Pranshu Chaturvedi is an Indian-American with interests in computational biology and machine learning. His interest in HPC comes from working with distributed systems at NCSA to build AI tools in hackathons and in his current internship. He is working on MemXCT, HPCG, and the poster.




Nishant Sheikh - Sophomore in Computer Science

- Nishant Sheikh is an Indian-American interested in the application of distributed systems in interdisciplinary research. His experience in HPC comes from project work involving constructing and maintaining a small-scale computer cluster, which led him to the competition. He is working on GROMACS.



Rittika Adhikari - Senior in Computer Science

- Rittika Adhikari is an Indian-American with interests in distributed systems and machine learning. Rittika's previous internship at Uber focused on building an automated pipeline for big data processing. She is working on LINPACK and GROMACS.




Freddy Zhang - Junior in Computer Engineering

- Freddy Zhang is a Chinese-American with interests in hardware systems and computer architecture. Freddy's previous work and research experience was from developing soft robots and robotic process automation. Freddy is working on CESM.




Liam Xu - Junior in Electrical Engineering + CS minor

- Liam Xu is a Chinese national with interests in Applied parallel program and General Purpose GPU computing. His prior experiences include working on an engineering simulation for Electromagnetic Interferences at Ansys. He will be working on MemXCT.



Jonathan Nativ - Senior in Computer Engineering

- Jonathan is an American with interests in distributed systems and parallel programming. He is developing an MPI interface for CUDA devices and extending a lightweight communication library for distributed graph analytics and Asynchronous Many-Task programming models. He is working on MemXCT, IO-500, and cloud resource management.



Team Diversity

- Diverse levels of experience:** Team includes sophomores, juniors, and seniors.
- Understudies:** Recruited understudy members to learn and become participants in future competitions.
- Inclusive Participation:** Recruited from Women in ECE (WECE) and Women in CS (WCS).
- Wide range of interests:** Atmospheric Science, System Architecture, Statistics, Cloud Computing, Machine Learning, Economics, Markets, Graphics.
- Industry and academic experts:** Provided input to team for deciding architectures, system administration, and team management.

Sponsorships

Nvidia

- Provided expert support for GPU software

NCSA

- Provided resources for hardware used for practice
- Offered expert advisors in system administration as well as in house support for IO-500 and CESM

Jump Labs

- Provided hardware for previous competitions that are now used for practice

Arm

- Provided application profiling tools and support

