Massachusetts Green Team
MIT / Boston University / Northeastern University

Po Hao Chen, Carlton Knox, Andrew Nguyen, August Wan, Yiran Yin, Vance Raiti, Kurt Keville

MEMBERS

Po Hao Chen
(BU'23)
Computer Science

- SC20&21 Veteran
- Theory & HPC
- Team Leader
- Benchmarks
- Reproducibility

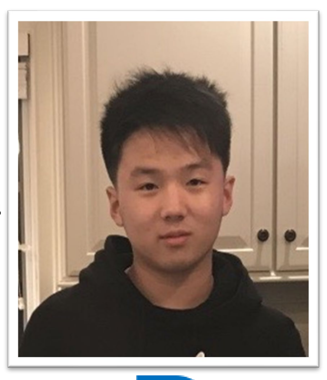

Yida Wang
(BU'24)
Computer Science

- ML & HPC
- Reproducibility

Carlton Knox
(BU'23)
Electrical Computer Engineering

- SC21 Veteran
- Computer Arch
- Benchmarks
- Hardware

Andrew Nguyen
(NEU'24)
Electrical Computer Engineering

- Compilers & OS
- GPUs
- LAMMPS
- Hardware

Vance Raiti
(BU'25)
Electrical Computer Engineering

- Mathematics
- Comp Arch
- PHASTA

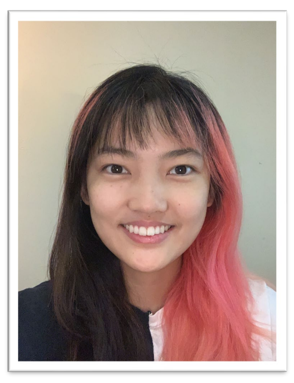

Yiran Yin
(BU'24)
Electrical Computer Engineering

- Mathematics
- HPC
- PHASTA

PHASTA

PHASTA

TEAM INTRODUCTION

Collaboration and Background

- ✓ **Consortium of institutions in Greater-Boston area representing MIT, Boston University, and Northeastern University.**
- ✓ **10-year anniversary of the first appearance of Massachusetts Green Team at SC12. We continue the legacy and partnership between local universities.**

Diversity and Inclusion

- ✓ **BU currently has the only HPC club in Greater Boston area. We strive to create a hub for students of all backgrounds to learn about supercomputing.**
- ✓ **Outreach program. We have fostered a connection such as SHPE to support underrepresented community, hosted recruiting events at local HPC meetups, and contact schools with less of HPC faculty presence.**

Institutional Support

- ✓ **Our efforts are made possible by the support of our universities and advisors: Kurt Keville, Dr. David Kaeli, Dr. Martin Herbordt.**
- ✓ **We represent the values of MGHPCC**

PREPARATION

Training Resources

To investigate the applications, the team will be able to leverage Azure Cyclecloud and remote university-owned clusters to conduct experiments with PHASTA and LAMMPS.

In addition, the team has access to AMD GPU instances through Northeastern University to experiment with benchmarking and practice the HIP compiler in preparation for the mystery challenge.

All of the members of the teams have experiences with hardware and software configurations through volunteering cluster management for the BUHPC club. We 1host training sessions throughout the year to provide hands-on experience to our students.

Furthermore, our advisor, Kurt Keville, will be meeting with us to discuss the details throughout the process.

APPS & BENCHMARKS

< Reproducibility>

Productivity. Portability. Performance – Data-Centric Python

Python is the backbone of scientific computing, and one of the most popular language of choices for researchers. We are very excited to be able to take on the challenges of reproducing the results from SPCL. Computational performance is one of the most important research topics that we are interested in, we believe that there is a lot to learn from this work. The benchmark suite even offers options for AMD GPUs natively, which we hope to leverage to gain a competitive edge.

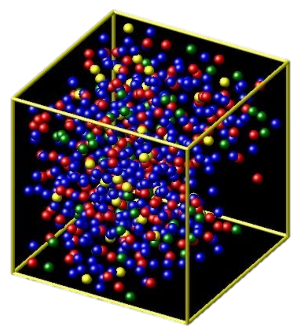

< LAMMPS >

Large-scale Atomic/Molecular Massively Parallel Simulator

Through the incorporation of concepts like Lennard-Jones potential for the interaction between particles, LAMMPS is extremely flexible, which is further shown by how many different platforms and hardware configurations it is compatible with, specifically with GPUs. Multiple hardware optimized packages such as the GPU and Kokkos can be used on the same build, giving us a plethora of options when trying to choose the perfect configuration for our system. To that end, LAMMPS has support for not just CUDA and OpenCL, but also HIP. This will be important as we aim to run it with high-end CDNA GPUs, and we hope to leverage their full potential for this application.

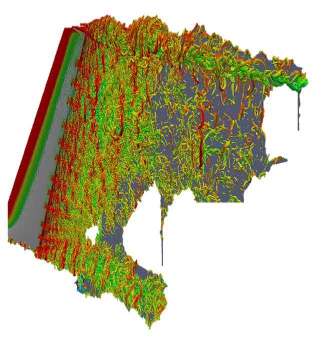

< PHASTA >

Parallel Hierarchic Adaptive Stabilized Transient Analysis

A computational fluid dynamics (CFD) application for supercomputing. PHASTA supports modeling flows of different attributes in 3D using mesh / unstructured grid processes. Our students interested in mathematical modeling have dug into the theory behind the science. We believe that our domain knowledge will yield its benefit when it comes to analyzing and optimizing the problems we are given.

< HPCG >

A set of Conjugate Gradient tests for ranking top systems. This benchmark complements LINPACK as it exercises computational and data access patterns that closely resembles modern scalable applications. Our team has an extensive experience with HPCG, and we hope to leverage our AMD-optimized binaries with its hardware accelerators to break beyond the records we have set in previous years.

< HPL >

The standard for ranking top supercomputers. Our team has extensive experience with HPL from previous years and research studies. We hope to leverage the MI210 GPUs which have demonstrated a competitive performance to its counterparts from other vendors.

< IO500 >

The standard benchmark suite for High-Performance I/O. We have previously worked with traditional HPC distributed filesystems for this challenge. This year, we are hoping to leverage a new modern filesystem, WekaFS, to gain an advantage in the competitions

< MLPerf >

The newest addition to benchmarking in SCC. AI/ML research is one of the fundamental driving forces of research in HPC and distributed systems. This benchmark suite is becoming increasingly relevant, and we believe that we will be able to come up with a competitive solution for MLPerf with our filesystem and hardware accelerator.

OUR SYSTEM

Hardware Configuration:

Our cluster hardware components are based on AMD architecture. We plan to have a total of two nodes; each with AMD infinity fabric supports and containing AMD Instinct MI210 GPUs and EPYC 7003-series CPUs.

In combination with our software stack, we believe that we will be able to fully leverage the accelerators to have an edge in the competition. For example, AMD Matrix Core provides the ability to work large machine learning models and datasets. This optimizes for memory-bound operations and benchmarks such as MLPerf and HPCG. In addition, our architecture is unique and new to the SCC. Thanks to the HIP Compiler and efforts to port applications, we are able to explore the applications with our AMD hardware.

Our storage system will be competitive with the high-bandwidth and low-latency NVMe SSD drives when combined with our filesystem of choice.

Software Configuration:

Ubuntu 18.04: great support for EPYC CPUs and other components in our setup. We will be using an optimized image for HPC workloads.

WekaFS: a modern filesystem that has shown to be able to handle the most demanding I/O workloads.

ROCm: The de-facto software platform for AMD accelerators. It includes optimized compilers (AOCL) and portability tools. In addition, it allows for easy install of scheduler (SLURM), libraries (OpenMPI, BLAS) and package manager (Spack).