HACKs: Huskies with Accelerated Computing Kernels

Northeastern University

Advisor: David Kaeli
Students: Harrison Barclay, Jason Booth, Chris Bunn, Jason Fitch, Anton Lazarev, Faridat (Toyin) Yusuf

Meet the Team

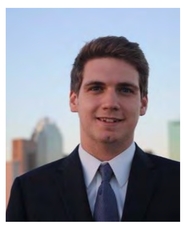

Harrison Barclay
4th Year
ECE
HPCG Benchmark

Jason Fitch
3rd Year
ECE, CS minor
OpenMC

Jason Booth
4th Year
CE/CS
Horovod

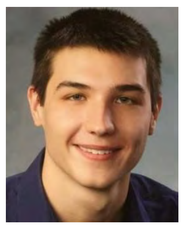

Anton Lazarev
3rd Year
CE/CS
System Admin

Chris Bunn
3rd Year
CE/Economics
SeisSol

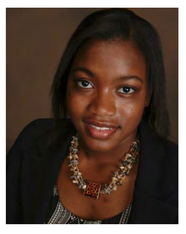

Toyin Yusuf
4th Year
CE (BS/MS)
LINPACK

This year, our team is composed of 5 new members, 4 of which have never attended SC before. We are committed to building a diverse team by recruiting through different engineering groups on campus. By incorporating diversity into our team's values, we will be able to bring different perspectives to the student cluster competition.

Why We'll Win

Backed by Experience: Over the past few years, our AMD-based solutions have performed well during the competition. We are able to draw upon the experience of past team members, who have competed in past cluster competitions for multiple years. Through Northeastern's co-op program, our members also have experience from top organizations such as AMD, Amazon Robotics, Intel, Intuit, MathWorks, MIT Lincoln Laboratory, GE, RKF Engineering, and various startups around Boston.

Industry Support: Through NUCAR, we have developed strong relationships with AMD and cluster vendors. We are able to utilize these relationships not only for top-of-the-line hardware but software support as well. By utilizing new libraries in our applications that are closely optimized for our hardware, we will be able to achieve top performance results.

A Unique System Architecture: Our system utilizes AMD CPUs as well as NVIDIA GPUs. As a result, we are able to balance power and performance through the use of the large number of processing cores in the AMD CPUs and through the lower power consumption of the GPUs.

Our Power Consumption Strategy: Despite the large amount of hardware, we will be able to keep our overall power consumption to below 3kW by scaling the CPU or GPU voltage depending on the resources needed for each application.

Deployment: During the competition, we will be using Ansible for application orchestration. This ensures that usage of our hardware will be as close to 100% as possible throughout the entire competition.

Our System

Our Cluster's 4 Nodes Each Have:

- 2 AMD EPYC 7601 CPUs (32 cores, 2 GHz)
- 3 NVIDIA V100 GPUs per node
- ConnectX-5 Infiniband Card
- 512 GB DDR4 RAM
- 250 GB SSD SATA Drive

General Hardware

- Fully connected GigE network/LAN via GigE switch
- IB network via an EDR Infiniband switch
- A Raspberry Pi will serve as our router, head node, and DNS/DHCP/TFTP/FTP server for that LAN

Software

- Ubuntu Server 18.04 LTS for familiar environment and compatibility with NVIDIA GPUs
- Red Hat's Ansible for application deployment management
- Required libraries and compiler to fully accelerate our applications on our hardware (including CUDA, Docker, OpenMP, OpenMPI, and gcc/g++/gfortran)

Applications & Strategies

LINPACK Benchmark

LINPACK is a software library for solving dense linear equations. This particular application is well suited for our NVIDIA V100s. The floating point performance and large amounts of VRAM in each GPU will ensure that LINPACK is executed as swiftly as possible.

HPCG Benchmark

Used to rank TOP500 systems, this benchmark is representative of important scalable HPC applications. Here, the large number of CPU cores as well as GPUs will produce competitive results.

Deep Learning with Horovod

Horovod is a MPI based training framework for TensorFlow. The three GPUs per node as well as the Infiniband switch will ensure that we are able to accelerate runs with minimal cross node overhead.

OpenMC

An open-source Monte Carlo simulator. As these simulations are large in size and scalable, it is best ran in the cloud. Because of the scalability of this program, we are able to fully scale without exceeding the cloud budget.

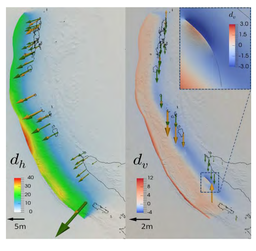

Reproducibility Challenge with SeisSol

SeisSol is a seismic wave simulator. During the competition, our team will replicate the results published from this work. We will obtain similar performance to the published work despite architecture differences by optimizing the libraries used for linear algebra.

Mystery Application

It will be difficult to prepare for this aspect of the competition. We are confident that our knowledge of our cluster and HPC will ensure that we are able to successfully run this application.

Acknowledgments

- Prof. David Kaeli, our team advisor, who has provided us with an immeasurable amount of guidance and support through the last few years
- Kaustubh Shivdikar, our graduate advisor, who has assisted us in cluster development and planning out our software strategy
- Our vendor sponsors at Nvidia, AMD, and Dell EMC for providing us with the necessary hardware