Northeastern University

Advisor: David Kaeli
Students: Jason Booth, Spencer Hance, Maddy Leger, Zachary Marcus, Carter McCardwell, Tristan Sweeney

Meet the Team

Our team is made up of students from Prof. Kaeli's computer architecture research lab.

Jason Booth
Third Year CE/CS
Software Eng.
Mr. Bayes

Spencer Hance
Fourth Year CE
CPU Architecture
HPCG

Maddy Leger
Fourth Year CE
Product Mgmt.
Team Admin / HPL

Zachary Marcus
Fourth Year CE/CS
Algorithm Dev.
Tersoff / HPL

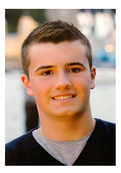

Carter McCardwell
Fourth Year CE
Systems Security
Born

Tristan Sweeney
Fourth Year CE
Systems Eng.
System Admin / HPL

Our team has an active recruitment program leveraging underclassmen participating in research to train next year's team. We are currently working with two students who will be joining us at SC this year. Our focus with recruitment is creating a more diverse team, so that we can bring a wider base of knowledge and experience to future competitions!

Our System

Pantheon's 4 Nodes

- 2 AMD EPYC 7551 CPUs (32 cores/64 threads, 2 GHz)
- 2 GPUs (2 AMD nodes, 2 NVIDIA nodes)
  - AMD Radeon MI25
  - NVIDIA Tesla V100
- Single Port, PCIe-16 FDR Mellanox IB
- 256 GB DDR4 RAM, 126 GB SSD, 2 TB HDD

General Hardware

- Fully connected GigE network/LAN via GigE switch
- IB network via an EDR IB switch
- RaspberryPi serves as router, headnode, and dns/dhcp/tftp/ftp server for that LAN

Software

- Ubuntu Server 16.04 LTS for familiar environment
- Custom ROCm kernel, AMD's compute environment
- System config maintained by Ansible
- Slurm used as the scheduler

Acknowledgements

Prof. Kaeli, our advisor, has been an invaluable educator, motivator, and mentor to our team over the years, and has been pivotal in our presence at SC.
Kurt Keville previously served as our advisor, and provided support to our team this year.
AMD's ROCm team, chiefly Greg Stoner and Guy Ludden, has provided us constant guidance in the use of their software.
Our presence at this competition was made possible by the generous donations of hardware by AMD, Supermicro, NVIDIA, and Mellanox.

Why We'll Win

Experience: Our team brings years of AMD-backed experience to SCC, and is ready to blow away our competition with the power of our Zen-based cluster.
Hardware: Our cluster has 128 threads per node (64 per Ryzen CPU), trading a lower system clock for an incredible amount of cores and lower compute to memory clock ratio. This allows us to exploit massive task-level parallelism or run applications in parallel efficiently. This computational power is backed by a 100 Gbps IB interconnect and fully connected GigE network.
Schedulers: Using Slurm and a custom tool for scheduling Born runs on the cloud, we will be able to fully occupy our cluster over the course of our the course of the competition.
Learning from Past Failures: Ansible has safeguarded us against the dreaded in-transit drive failure, ensuring that from an unimaged drive we can have our system reconfigured in time for the competition start. Ansible also gives us the flexibility to explore last-minute system configuration options with vendors on-site, and roll back as needed.
Power Consumption: Utilizing the IPMI ports on our nodes, CPU frequency adjustment, and the SMI tools for our GPUs we will be able to dynamically adjust the power consumption of our cluster to keep power consumption under but near the allocated 3 kW.
Hybrid System: AMD's MI25 cards don't have double precision floating point units, and so for LINPACK and HPCG, which are double precision benchmarks, we will scale down power as low as possible on the MI25 nodes and overclock the V100 nodes. We can use chassis IPMI to overdrive our fans and dissipate the heat the V100s generate.

2015: Used Symmetric's Distributed Memory System. Driver issues forced us to run standard Linux instead of their custom kernel.

2016: Our cluster got lost in shipping and we hacked together things we flew on the plane, as well as parts borrowed from AMD's booth!

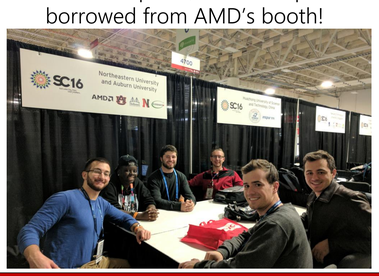

2017: Back to win with new recruits and lots of learning from previous years experiences.

The student groups we're a part of & co-ops/internships our team has done have helped prepare us for SCC17!

Applications & Strategies

LINPACK

HPL will take advantage of the superior double precision floating point performance of the NVIDIA Tesla V100s. We tuned it to fully occupy the V100s' VRAM, maximizing computational efficiency across the two NVIDIA nodes.

HPCG

HPCG's sparse matrix operations have been optimized for our V100s and studied across a variety of runtime parameters. The low communication overhead between nodes facilitates scaling across the 4 GPUs.

Tersoff

We are accelerating Tersoff using the GPU optimizations from the paper's artifact description. The vectorization is shown to provide speedups by isolating the computation to just the GPU for all computational precisions.

Mr. Bayes

We are customizing Mr. Bayes to take advantage of our NVIDIA V100s and AMD EPYC CPUs, as well as the large amount of RAM on our system used for the Bayesian modeling.

Born

Born processes seismic tests to refine it's result, presenting task-level parallelism which we exploit through CycleCloud. The discrete "shots" let us schedule the work to take advantage of times when the local cluster isn't busy.

Mystery App

Our system has a properly configured environment, and we're going to have rested team members ready to tackle the mystery app once announced. With our architecture, we're ready for anything GPU, CPU, or network bound.