



Reaching for 100TFlops at 3KW Power with Intel Scalable Processors and NVIDIA V100 NVLINK GPUs

+Alex Ballmer, +Brendan Batliner, +Anna Benson, +Blake Ehrenbeck, +Zhen Huang, +Parker Joncus, +Travis Koehring, +Alexandru Orhean, *William Scullin, * Ben Allen, +Ioan Raicu
Department of Computer Science, Illinois Institute of Technology, Chicago, IL
Argonne Leadership Computing Facility, Argonne National Laboratory, Lemont IL
{aballmer,bbatliner,abenson,behrenbe,zhuang38,pjoncus,tkoehring,aorhean}@hawk.iit.edu
{wscullin,bsallen}@alcf.anl.gov, iraicu@cs.iit.edu

Our Team



Alex Balmer is a fourth year coterminial student at IIT studying computer science. He is a returning student to the Student Cluster Competition, having competed in it three times since 2014. His interests include writing (fiction) and experimentation with wireless mesh networking.

Brendan Batliner is a fourth year undergraduate student in CS at IIT. Although new to scientific computing, he has a background in high performance systems programming in C++. He has worked for numerous technology startups, including one of his own. Outside of computer science, he is a member of IIT's track team and is an aspiring DJ.

Anna Benson is a third year CS student at IIT with a minor in applied mathematics. She is primarily interested in mobile applications and web development but has greatly enjoyed learning about and working with high performance systems.

Blake Ehrenbeck is a fourth year student at IIT studying Computer Science with a specialization in Data Science. His interests include Artificial Intelligence and Data Analytics. Over the summer he worked at Argonne National Labs improving the Cobalt job scheduler.

Zhen Huang is a fourth year Computer Science student at IIT, specializing in Data Science. He previously researched benchmarking the Mira supercomputer at Argonne National Laboratory in spring 2018 and continued working there in summer 2018 handling the cobalt data staging with Globus OAuth and Flask API.

Parker Joncus is an applied math and data science coterminial student at IIT. He started work this summer learning Linux and starting to tune the hardware. He has enjoyed learning more about computer systems and broadening his horizons, as well as playing with a couple hundred thousand dollar machine.

Our School

Illinois Institute of Technology (IIT) is a private, technology focused research university offering undergraduate and graduate degrees in engineering, science, architecture, business, design, human sciences, applied technology, and law. IIT is centrally located in Chicago. For the past four years, IIT has worked closely with Argonne National Laboratory to send a team of students to the Student Cluster Competition at the annual Supercomputing Conference.

This year our team is comprised of 6 IIT undergraduates, 1 back up student, and 3 mentors. We have been working since this summer to explore and profile the competition applications in preparation for SCC 18.

Preparation Strategy

The team met at minimum once a week to discuss HPC topics such as CPU architecture, networking, GPU / CPU work loads, and memory management.

Assignments included tuning HPL / HPCG / competition applications on varied hardware, then reporting and analyzing the output.

Sub teams did more in depth studies of each of the applications. This gave insights into the functionality of the non-benchmarking applications and how they are constructed.

Other students were assigned to the cloud platform, learning how to set up a node and decide the best way to gain performance at an efficient cost.

A mock competition was held with judges to prepare the team for the potential on the fly problems to work out and major bugs / inadequacies.

Hardware

Type	Hardware	Quantity
Chassis	SUPER MICRO X11DGQ	3
CPU	Intel Xeon Skylake 2.1 GHz 22 Core	6
Memory	16 GB DDR4 2666 MHz	36
Storage	3.84TB Intel SATA 6 Gb/s	6
Accelerators	NVidia Tesla V100 32GB with NVLINK	12
Interconnect	Mellanox ConnectX-5 100 Gb/s	3



Software

Fedora 27

- More stable than other, more optimized versions of Linux
- Better package availability
- New enough to support modern features and instruction sets

Intel Compilers, MPI, MIKL

- The Intel tools are better optimized for Intel processors than other computing tools

Slurm job scheduler

- SLURM schedules jobs on the cluster
- Extremely configurable but simple to use

Spack package manager

- Allows for rapid deployment of packages using different combinations of libraries and compilers



Why We Will Win

- We have a team member who has been to SCC before, as well as one who has been working on these applications since summer 2018
- HPL and HPCG have been studied by all team members, while the other applications were assigned to teams of two
- We have had excellent support from our advisors, backup team members, teaching assistants, and sponsors
- A steady supply of candy

Acknowledgements

This work is supported in part by NSF OCI-1054974 CAREER award (REU supplement) and NSF CNS-1730689 CRI award (Mystic Testbed). This work has used resources of the ALCF at ANL, which is supported by the DOE contract DE-AC02-06CH11357. Furthermore, the Chicago Fusion team will make use of the NSF-funded Chameleon testbed at Argonne National Laboratory. The team is grateful for the generous support by Intel, NVIDIA, Mellanox, Argonne National Laboratory, Illinois Institute of Technology, and the National Science Foundation. The team would like to especially thank Alexandru Orhean for his assistance.

Power Management

One of the primary constraints in this competition is a total cluster power consumption limit of 3000 Watts. Some of the tools and methods we have used to stay below this limit while maintaining high application performance are:

- Using ipmitool to track instantaneous and average power usage on our nodes
- Controlling fan speeds, turning them off during runs, and turning them on high afterwards to cool the cpus
- Manually capping power by changing clock speed frequency
- Opting not to use network switches to eliminate their idle power usage

Application Optimization

- Experimenting with compiler optimization flags
- Using sysctl to turn off swappiness and tune other Linux kernel settings
- Determining the ideal number of cores per node to run on
- Process pinning to assign specific cores to specific processes
- Investigating how applications scale across multiple nodes