University of Illinois at Urbana Champaign

Chung Ting Huang, Edward Hutter, Jonathan Pan-Doh,
Wei Ren, Nihar Sheth, Yan Zhan

Introduction

The University of Illinois is a world-class research facility located in Urbana-Champaign. While many of our faculty and associates frequent the Supercomputing Conference, this is our first time participating in the Student Cluster Competition. With the support of our distinguished faculty and advisors, as well as the generous support of our sponsors, we're hoping to leave a lasting impression at our first cluster competition.

Our team is diverse in background and education, united under the goal of winning the competition. Our students are juniors and seniors majoring in Computer Engineering, Electrical Engineering, Computer Science, and Engineering Physics. We have a broad range of skills, including system administration, parallel programming, research experience, system programming, and visualization. Our complementary skills and experience make us a force to be reckoned with.

Team Members

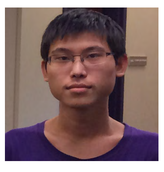

- Yan Zhan
- Junior in Computer Engineering with a minor in Computational Science and Engineering. Engaged in research at NCSA. Experienced in system administration and hardware.

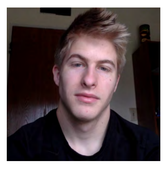

- Edward Hutter
- Senior in Computer Engineering. Involved in research in the Parallel Programming Lab. Experienced in implementing parallel algorithms.

- Jonathan Pan-Doh
- Senior in Electrical Engineering specialized in Operating Systems with interests in Computer Architecture.

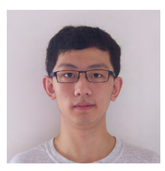

- Wei Ren
- Double major in Engineering Physics and Computer Engineering with research experience on LSST scientific visualization.

- Nihar Sheth
- Junior in Computer Science with experience in multithreaded environments and systems programming.

- Chung Ting Huang
- Senior in Computer Engineering with a minor in Computational Science and Engineering. Experienced in CUDA programming.

Why We Will Win

With the support of both the University of Illinois and the National Center for Supercomputing Applications, we've learned a lot and applied our new knowledge to making our system run as optimally as possible. Aside from weekly meetings to discuss progress of our system, we also meet separately every week to discuss HPC concepts in a classroom setting, as we review lectures from a graduate-level class offered on campus. We also worked closely with our vendors and sponsors to improve the performance of our applications. Our ability to work together, coupled with the aid of our incredible advisors, gives us our competitive edge and will distinguish our performance from the other teams.

Software

Our software stack includes: CentOS 7 · Intel Parallel Studio XE 2017 · Nvidia CUDA Toolkit 8 · Allinea Forge & Performance Reports · OpenMPI 1.10 and 2.0 · John the Ripper · Torque Resource Manager.

Our software stack was carefully chosen by our team in advance. We chose CentOS 7 because it's widely used for clusters and has a thriving community and support. The Intel Parallel Studio XE suite offers a variety of wonderful tools including a compiler and an MPI implementation. Nvidia CUDA toolkit 8 is important for its CUDA compiler and relevant math libraries. John the Ripper was chosen as our password cracking tool as it is one of the most popular tools available, is open source, and offers GPU support for several hashing algorithms. Lastly, we're grateful to Allinea for offering a license to use their performance profiling toolset for this competition.

Hardware

- Dell PowerEdge R7 Server (x5)
  - Intel Xeon E5-2698v4 CPUs (x2)
  - Nvidia Tesla P100 GPUs
  - 256 GB DDR4 RAM
  - 400 GB SSD
  - Mellanox ConnectX-4 InfiniBand EDR
  - Intel X520 10Gb Ethernet
- 6x 800GB SSD on head node, RAID-5
- Mellanox SB7890 Infiniband EDR Switch

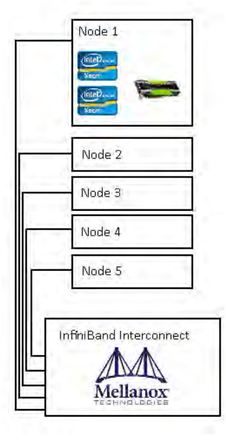

This hardware configuration was chosen to provide a balance between CPUs and GPUs. Xeon processors offer powerful performance across many cores, and Tesla P100 GPUs have incredible parallel floating-point performance. EDR Infiniband offers extremely high bandwidth and low latency for inter-node communication, essential for a modern HPC system. By design, our system runs strictly below the power limit specified by the competition rules, so we will not overshoot the limit during the competition.

Applications

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

HPCG

ParaView
Parallel Visualization Application

• For HPL, we obtained a highly optimized implementation of the benchmark for our GPU accelerators to ensure best performance. We also used a series of tests to determine the optimal configuration for the benchmark.
• We improved our HPCG score by obtaining updated binaries for the GPU implementation and determining an appropriately large problem size that maximizes the benchmark score.
• We tested with various flags and compiled ParaView from source to fully utilize the CPU and GPU hardware. Instead of creating a GLX context for rendering process we use the new EGL interface to bypass the requirement of X in our cluster.
• For ParConnect, we have prepared a program to show our cluster's capability to reproduce the results achieved in the corresponding paper. Our program launches many jobs, taking advantage of all nodes and placing varying numbers of processes on each node so as to balance the CPU workload.
• We chose John The Ripper as our tool for the password recovery task due to its ability to use GPUs. We also prepared scripts to evenly distribute the workload across CPUs and GPUs.
• We prepared for this year's mystery application by building and running previous mystery applications. We improved our ability to quickly familiarize ourselves with new applications through this process.

NCSA JumpLabs DELL NVIDIA ILLINOIS UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN