


University of Illinois at Urbana-Champaign

Roshan Rajan, Daniel Margosian, Mang Yu, Landon Clipp, William Klock, Yan Zhan

Introduction

The University of Illinois at Urbana-Champaign has assembled a talented group of students who are interested in high performance computing and eager to compete in the 2017 Student Cluster Competition (SCC). This is our second time competing and our improved understanding of the competition will make us more competitive. With the support of our distinguished faculty and advisors, as well as the generous support of our sponsors, returning team members recruited a team of exceptional students proficient in high performance computing and system administration, ensuring we have a diverse yet concentrated skill set.

Team Members



- Yan Zhan
- Senior in Computer Engr. Experienced in system administration and hardware.
- Second-time SCC participant




- Roshan Rajan
- Junior in Computer Engr. Experience in optimizing resource management for multi-CPU systems.



- Landon Clipp
- Senior in Computer Engr. Works with the NCSA to develop software for weather processing using the Blue Waters super-computer.



- Mang Yu
- Senior in Elec Engr. Interned with the CSIRO of Australia to develop and optimize Linux network applications.



- Daniel Margosian
- Senior in Comp Sci. Experience developing networking internals for Caterpillar machines and Linux server administration.

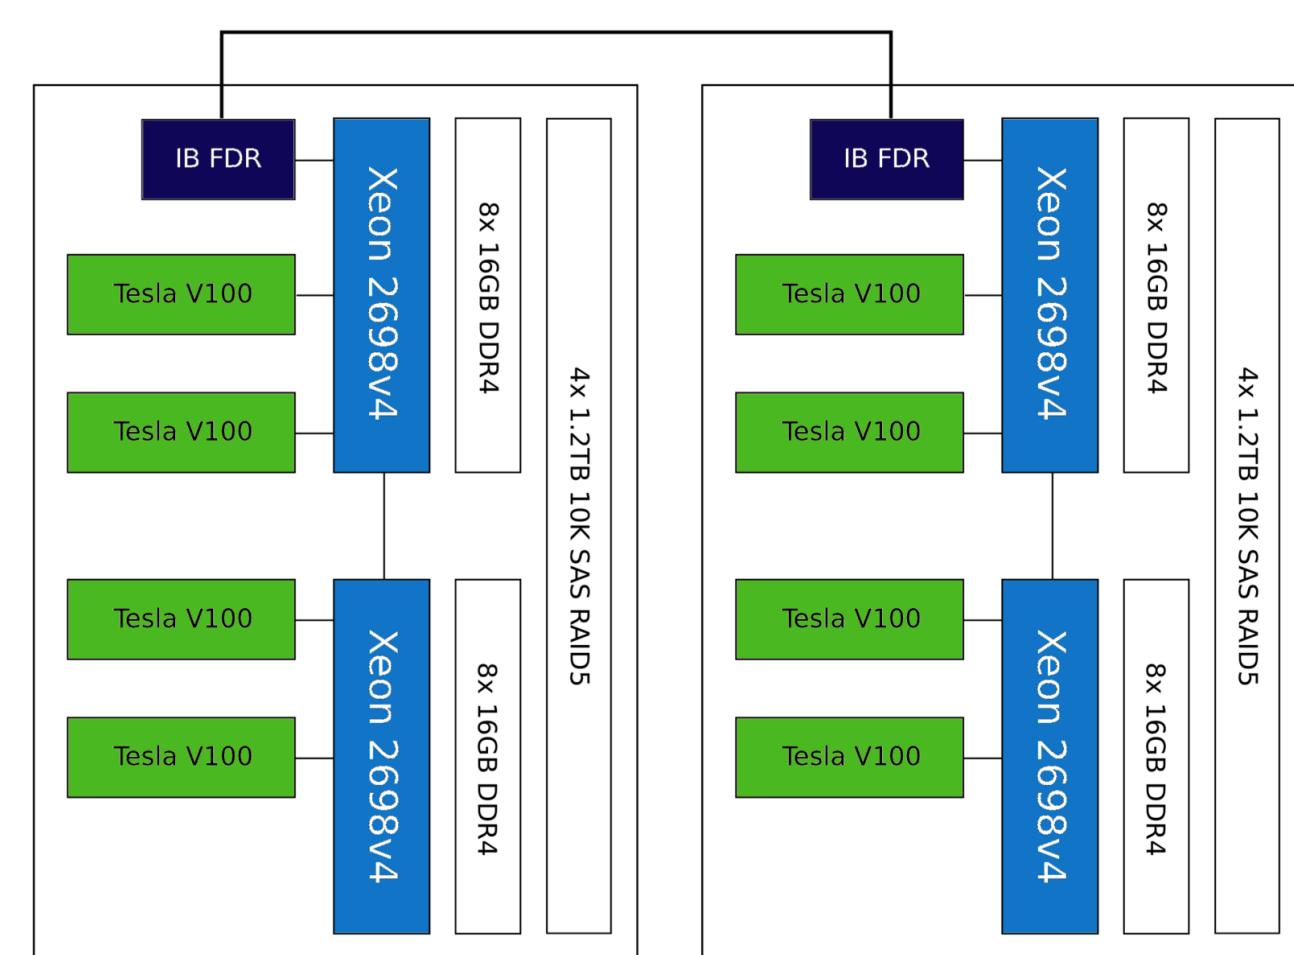


- William Klock
- Senior in Computer Science and Linguistics. Experience in server administration, parallel programming, and signal processing.

Hardware

Dell PowerEdge C4130 (x2)

- Intel Xeon E5-2698 v4 CPUs (x2)
- NVIDIA Tesla V100 GPUs (x4)
- 256 GB DDR4 2133MHz RAM
- 4x1.2TB 10k SAS Hard Drives in RAID 5
- Mellanox ConnectX-3 VPI Infini-Band FDR Adapter



This hardware configuration was chosen to maximize GPU utilization because we determined that General Purpose GPUs are now mature enough for many scientific computing applications. Xeon processors offer powerful performance across many cores, and Tesla V100 GPUs have incredible parallel floating-point performance. InfiniBand offers extremely high bandwidth and low latency for inter-node communication, essential for a modern HPC system. Contrary to common HPC systems, we decided to directly connect the two nodes to each other without a switch in between, which improves link performance and reduces power consumption. By design, our system runs strictly below the power limit specified by the competition rules, so do not expect to overshoot the limit during the competition.

In addition to our physical cluster, we will also make use of the Microsoft Azure compute cloud. This resource is controlled through the CycleCloud interface. Because our physical cluster greatly emphasizes GPU computation, we tailor our cloud towards CPU-centric applications. This affords us the flexibility to fully utilize GPU and CPU paradigms for applications that benefit from them.

Software

Our software stack includes: CentOS 7.4 · Intel Parallel Studio XE 2018 · NVIDIA CUDA Toolkit 9 · OpenMPI 2.1.2 · Mellanox HPC-X 1.9.7 · Torque Resource Manager.

Our software stack was carefully chosen by our team in advance. We chose CentOS due to its good stability, well-written documentation, support for the GPU and interconnect hardware, and wide use at NCSA. We are using Intel's Parallel Studio XE suite, which includes Intel MPI, Intel Math Kernel Library (MKL) and the ICC compiler. Intel MKL provides optimized BLAS routines and other math functions for Intel hardware, and the ICC compiler can produce binaries better optimized for Intel hardware than the standard GCC compiler. The NVIDIA CUDA toolkit is important for its CUDA compiler and relevant math libraries targeting NVIDIA GPUs.


Applications

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

- For HPL, we obtained the latest, GPU-optimized version of the benchmark. We also ran a series of tests to find the optimal inputs for the highest performance.

- We obtained the latest version of HPCG benchmark that is optimized for GPU processing. To maximize the performance, we experimented with different problem sizes and various of MPI runtime parameters to determine the optimal configuration.


HPCG



- For the Tersoff Multi-body potential, we tested many different software configurations to reproduce as much of the work presented in the academic paper as possible.



- Using a combination of Beagle and OpenMPI we focused on maximizing resource utilization and performance on the cluster for MrBayes. Using this configuration, we experimented with many different model parameters for precise and accurate results.



- We optimized Born to optimally utilize our Intel CPUs while maintaining precision and accuracy in the calculation of the reverse time migration.



- 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.

Why We Will Win

We have used the significant academic resources of both the University of Illinois and the National Center for Supercomputing Applications (NCSA) to apply knowledge about high-performance, heterogeneous and parallel computing to making our system run as optimally as possible. Aside from weekly meetings to discuss progress of our system, members attended a graduate-level independent study course. Since the SCC is a complex event with many different point-bearing facets, we managed our time and effort by partitioning the various competition components amongst our competition and understudy members. Our 6 male and 3 female understudy members work closely with our competing members to prepare and expand their own skill set for next year.

