Team SupercompUtes

Qixiang Chao, Hannah Eyre, Emerson Ford, Janaan Lake, Kevin Parker, and Kincaid Savoie

Alternates Paul Fischer and Robben Migacz

ABOUT OUR TEAM

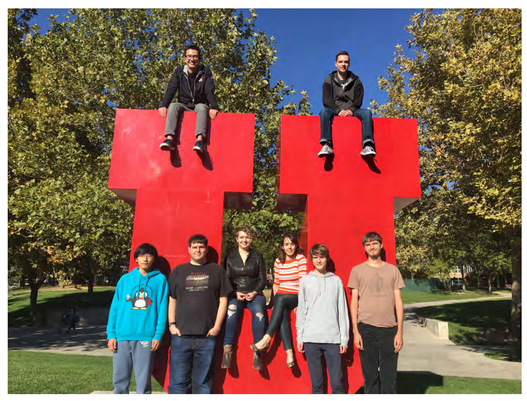

- Our team consists of returning members from last year, who were recruited from a special topics class taught by two professors active in the SC community, and several new recruits
- To recruit students for the class, the female professor advertised in classrooms and an ACM chapter meeting and asked other professors for suggestions to encourage diverse involvement
- Of our team of six, two are female; in contrast, our undergraduate CS/CE majors are represented by only 13% females. We also have one international student on our team
- The team, including its two alternates, is comprised of Computer Science, Computer Engineering, Mathematics, and Mechanical Engineering undergraduates with expertise in the arts, business, and finance; several also have significant work experience at the University of Utah and in industry
- At the beginning of the semester we worked in groups to build and run the applications, but over time the tasks have become individualized to leverage expertise and interest

COMPETITION PREPARATION

- In a special topics course in the fall, we acquired broad knowledge of HPC foundations, scaling and performance, and hardware and software technologies
- HPC-oriented faculty and Center for High Performance Computing (CHPC) staff guided us on underlying concepts and competition preparation
- We learned about the science behind the applications and gave presentations for the other students in class
- We discussed how to improve our performance from last year and read the journal articles from last year's Reproducibility Test
- We could often be found at CHPC's offices, where the cluster was housed during preparations. CHPC was critical in our preparation for the conference, providing space to work and around-the-clock assistance




SOFTWARE AND HARDWARE SYSTEM

The "SaltFlats" cluster

bluebird

- Dell PowerEdge R730
- 2.4Ghz Broadwell Intel CPUs (28 cores)
- 256GB RAM
- 2TB local drive for OS and scratch space, 1TB mirrored SSDs for / home

compute1–compute4

- Dell PowerEdge R740
- 2.1GHz Skylake Intel CPUs (32 cores)
- 192GB RAM
- 2TB local drive for OS and scratch space

GPUs

- bluebird: 1 NVIDIA Tesla P100
- compute1: 3 NVIDIA Tesla P100s
- compute2–compute4: 3 NVIDIA Tesla V100s

Interconnect

- bluebird: 56GBit/s Fourteen Data Rate (FDR) InfiniBand (IB)
- compute1–compute4: 100Gbit/s Enhanced Data Rate (EDR) IB

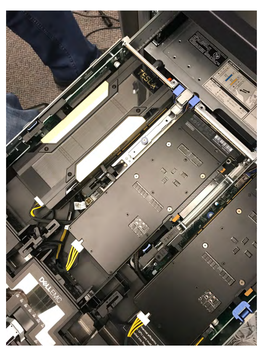

Utah's Bonneville Salt Flats are the home of many land-speed records. The SaltFlats cluster is just as fast!

Hardware Diagram

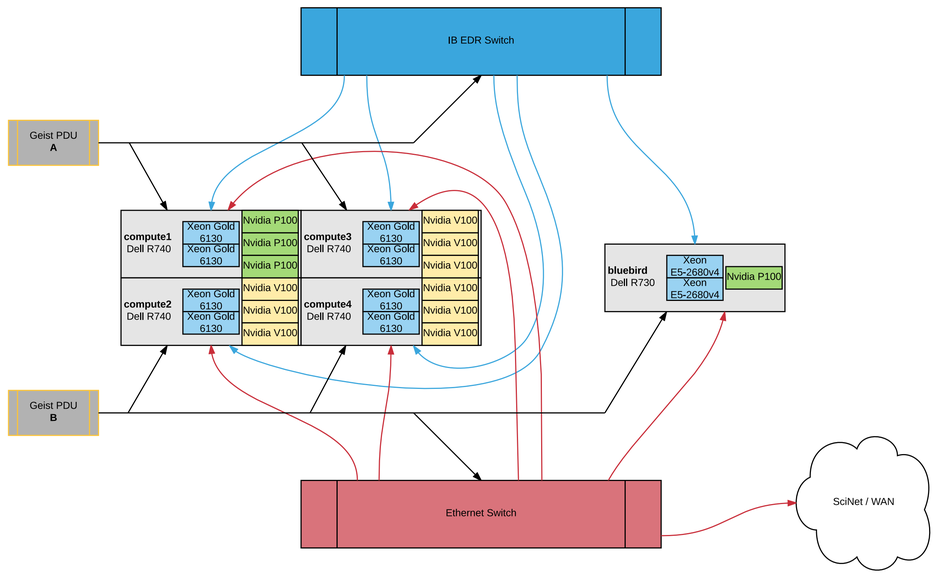

Justification of Hardware Configuration

Our goal is to be able to maximize throughput while abiding by the power limit. We chose a cluster with many GPUs because they offer us higher throughput for the same power limit (versus CPUs).

- HPL, which is compute-bound, uses GPUs to increase parallelism and achieve better performance
- HPCG, which is memory-bound, utilizes the higher memory bandwidth of GPUs
- MrBayes and Born have GPU implementations
- The hybrid cluster (with different hardware generations) will allow us to test the performance portability of Tersoff Multi-Body Potential

Power

- The Dell PowerEdge R740 nodes are configured for a maximum power draw of about 1130 W per node, including 3 NVIDIA Tesla V100 GPUs
- For the competition, a custom monitoring script records power usage while programs are running and will alert when sudden spikes in power occur or when power usage approaches the designated limit

Software

We adopted most of the software stack already installed on the CHPC clusters to make it easier for us to get help with configuration issues, and so those who were already familiar (from work or research experience) would not need to learn new software. This year, we are using Slurm to schedule jobs on the cluster for more efficient scheduling.

- CentOS 7 Linux
- GNU Compiler Collection (GCC) and Intel C++ (ICC), latest versions
- NVIDIA CUDA 8 and 9
- OpenMPI 1.10.7 and 3.0.0, Intel MPI 5.1.3, MVAPICH 2.3b
- OpenMP 3.0
- Intel Math Kernel Library (MKL) 11.3.2
- Lmod from Texas Advanced Computing Center (TACC) for software configuration management
- Slurm Workload Manager, Version 17.02
- Allinea for profiling
- Latest versions of HPL and HPCG from NVIDIA

APPLICATIONS AND OPTIMIZATION STRATEGY

HPL and HPCG

We used optimized implementations of HPL and HPCG and utilized autotuning and manual configurations (steered by plots and charts and user discretion) to determine the best problem size (among other variables) and MPI implementation for both HPL and HPCG.

Reproducibility

Our strategy is to run the application on our heterogeneous CPUs (Broadwell and Skylake and highlight the portability and performance improvement of the vectorization by recording the results on both the AVX2 and AVX-512 SIMD instruction units.

MrBayes

Our preparation for MrBayes involved calculating evolutionary models for whiteflies using genetic available from Whiteflybase. We tested 1, 2, 4, 8, and 16 MPI processes, with and without the beagle acceleration library for GPU and CPU, across different combinations of nodes to determine the best throughput-to-resource ratio.

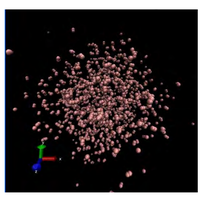

Born

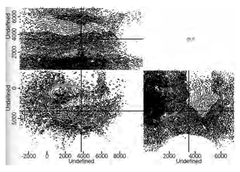

We run Born on the cloud since it scales well and can take advantage of wide SIMD AVX-512 instructions. To optimize Born, we used different profilers and vectorized the most expensive loops. We also tried to access the data in a more cache-friendly manner using tiling. We reduced the number of computations spent in the main loop by reducing the calculated range.

Cloud Utilization and Prioritization

Our team evaluated various workloads and the relative cost of running them in the cloud versus our cluster. We also evaluated the InfiniBand-enabled H-series instances and compute-optimized F-series instances. We leveraged the tooling offered by CycleCloud to build scripts that automatically provision the cloud-based cluster so its software is almost identical to our cluster, allowing us to quickly adapt to new workloads and budgets during the competition.

The reproducibility test will be run on the first day. The remaining time will be used for MrBayes (which is dependent on the number of data sets), Born, and the mystery application. Any gaps on the cloud or cluster will be filled with Born.

WHY WE WILL WIN

- Our team is very diverse with respect to gender, age, skill sets and interests, and backgrounds
- Four members of the previous team are returning with all they learned last year
- State-of-the-art hardware acquired by working with great engineers from three renowned companies
- Optimized HPL and HPCG implementations from NVIDIA
- Input and excellent advice from many experts in different areas

ACKNOWLEDGMENTS

We are grateful to our sponsors Dell, NVIDIA, and Intel; the many people who helped us prepare for this competition, including the professors of our class, Mary Hall and Hari Sundar; and Center for High Performance Computing staff, especially Brian Haymore and including Anita Orendt and Martin Cuma.