

Team Texas Longcats

HARDWARE AND SOFTWARE CONFIGURATION

HARDWARE

- ★ 8 x Dell PowerEdge R740 nodes
- ★ 2 x Intel Xeon Platinum 8160 2.1Ghz per node
- ★ 192GB RAM per node
- ★ 1 x Nvidia Tesla V100 per node
- ★ 1 x Mellanox Connect-X5 single port adapter
- ★ 2 x Intel P4500 Series 4.0TB NVMe SSDs (Striped)
- ★ The nodes are connected by a Mellanox InfiniBand EDR switch capable of transmitting 100Gb/s



SOFTWARE

- ★ CentOS 7.4
- ★ Puppet
- ★ GCC, Intel compiler
- ★ Python 2.7
- ★ OpenMPI, Intel MPI
- ★ Intel MKL/TBB
- ★ CUDA 9.0
- ★ Allinea
- ★ Lmod
- ★ OpenHPC

Given that we have a mix of applications that will be run with CPU-only and GPU-only configurations, we decided on a hardware configuration that will allow us to do both. We have outfitted our cluster with the most powerful CPUs and GPUs available within our power budget and have optimized for performance per watt. This also gives us flexibility for the mystery application. We will be running CentOS 7.4 as our operating system because it is the latest release and is what it currently used on TACC systems. OpenHPC is a convenient suite of compilers and HPC performance libraries. Eight nodes was the maximum we could fit in our power budget while also providing a healthy amount of dedicated resources (cores & RAM) to each of our scientific applications.

GENERAL STRATEGY

- ★ Depending on the given cloud budget, we will look to run as much as we can on cloud compute nodes, especially Born and MrBayes. We think and hope the cloud component will be an important part of the competition and we are prepared to capitalize on it.
- ★ In order to keep under the power budget, we will be downclocking any CPUs and GPUs not being used, especially for the HPL benchmarking application, the most power demanding of the applications. We tuned BIOS settings to help keep power consumption under control.

RUNNING AND OPTIMIZATION

HPCG The HBM2 on NVIDIA's V100 GPUs, along with the NVIDIA-provided optimized CUDA 9.0 HPCG binary allows us to achieve significant performance gains over previous generation hardware and software. We spent many hours iterating over problem sizes, GPU clocks, and exploiting any possible benefits from CUDA aware OpenMPI.

HPL The high count of TFLOPS that the NVIDIA V100s are capable of, along with an optimized CUDA 9.0 HPL binary provided by NVIDIA will allow us to get a high score. We tested many different HPL.dat parameter combinations, frequencies, and affinity settings to get the best run configuration for our cluster under the power budget.


MrBayes We will use the MPI version of MrBayes to farm out each Markov chain. In order to use MPI efficiently, we are running only as many cores as there are chains multiplied by runs. We may use a starting tree in order to limit the search space or vary the heating temperature to more quickly find isolated peaks. The ability to checkpoint our runs will allow for us to fully stop and restart if necessary. We will also be able to merge the output trees of various runs and use this starting tree for a new run. We have also compiled using the Beagle library in order to speed up calculating likelihood of sequence evolution. If we have a significant cloud budget, we intend to run MrBayes on the cloud.

Born Born's processor and memory demands have led us to experiment with many different processor affinities and TBB vs OMP threading strategies. Our choice of CPUs provide a good mix of compute and memory speed, as well as enough capacity to run multiple shots at once. The NVMe SSD RAID 0 also provides exceptional IOPS and space for Born's gigantic workload.

Reproducibility Challenge We have built and are able to run the vectorized Tersoff potential for LAMMPS using various CPU options or with GPU acceleration enabled. We are ready to run it on whichever architecture is most convenient or required.

THE TEAM

- ★ Our team is composed of students from two great Texas universities, the University of Texas at Austin and Texas State University in San Marcos. This is the second year of collaboration among the two universities. All team members are first time competitors.
- ★ Recruitment for the team was done at both universities and aimed at both CS and non-CS students. The recruitment team strove to compile an SCC team representative of the makeup of our diverse central Texas communities.
- ★ Our team members are diverse in our education, ethnicity, gender, orientation, and motivations. Half of the team consists of members traditionally underrepresented in STEM.
- ★ We are proud to have a collective of individuals that each bring something different to the table.



Austin Reichert
Computer Science
HPCG, Born



Vivian Nguyen
Electrical & Computer
Engineering
Mr Bayes




Jamie Sullivan
Astronomy, Physics, &
Mathematics
Reproducibility Challenge



Collin Weir
Electrical Engineering
Born, MrBayes



Jeremie Gallegos
Mechanical Engineering
HPL, HPCG



Carlos Hurtado
Computer Science
HPL, Cloud

PREPARATION

- ★ We started attending preparatory meetings in April where we learned the basics of building and running the applications.
- ★ Due to the close relationship between UT and TACC (Texas Advanced Computing Center), we had access to their supercomputer systems which allowed us to run, test, and optimize the applications even before building our own cluster.
- ★ As time progressed, we split up the applications among ourselves and dove into a deeper level of understanding of them.
- ★ For each application, we performed power analysis to create a budgeting system to prevent going over power, tested different node and accelerator configurations using TACC systems, varied number of threads and tasks launched with OMP and MPI plus affinity, varied CPU and GPU frequencies each run
- ★ We spent many hours learning new skills to prepare us for the competition.
- ★ Our mentors were very knowledgeable and helpful throughout the whole process.
- ★ Once our cluster had shipped to the competition, we directed our full attention to Cyclecloud, the cloud component of the competition. It was necessary that we learn how to run the applications on the cloud in the best way possible.



WHY WE WILL WIN

- ★ We have:
 - a diverse team with unique manners of thinking and broad expertise
 - massive support and guidance from our very knowledgeable advisors
 - access to TACC resources that we could utilize 24/7 for realistic application testing and optimization on hardware comparable to that at the competition
 - a good understanding of and ability to utilize Cyclecloud resources
 - powerful cluster able to run GPU- and CPU-dependent applications
 - put countless hours of work into learning and practicing