BOSTON GREEN TEAM

The Team

- A fun fact - The Boston Green Team has 6 students with 6 different nationalities.
- The Boston Green Team represents a diverse group of students from **two Massachusetts universities, UMass Boston, and UMass Lowell.**.
- The students have joined together to create a **diverse team** consisting of students from three majors - Computer Science, Computer Engineering, and Mathematics.
- All of our team members come from member institutions of The Massachusetts Green High Performance Computing Center (MGHPCC) a LEED Platinum Certified research data center in western Massachusetts.

Team members

- **Taha Azzaoui** is a Computer Science and Mathematics Major from the University of Massachusetts, Lowell. He is lead on porting Tersoff to the Jetson TX2 platform.
- **Evan Donato** is a Mathematics and Computer Science major at the UMass Boston. He is the team captain and is lead on the Born application and has split duties in administering the system.
- **Michiele Ocbagabir** is a Computer Science major at UMass Boston. He is lead on Tersoff.
- **Cristian Peguero** is a Computer Engineering major at UMass Boston. He is lead on the MrBayes application and split duties in administering the system.
- **Hanfei Xu** is a Computer Science major at UMass Boston. She is lead on HPCG benchmark.
- **Alexander Zhurkevich** is a Computer Science major at the UMass Boston. He is lead on optimizing and running HPL benchmark.

Strategy

For benchmarking we are planning on running on the entire cluster. For HPL we will be scaling the wattage on the V100s to bring the clocks in line with the P100s and will take advantage of the power savings afforded by non-linear scaling of FLOPs/watt. For HPCG we will run all 16 GPUs at full wattage as HPCG is far less power intensive. We feel that this approach will give us a huge advantage in the HPCG benchmark.

For the applications we will leverage both our CPU and GPU capabilities as well as our large pool of ram to limit potential paging.

We feel that this combined strategy for benchmarks and applications along with our diverse skillsets will give us the edge to win the competition.

Hardware / Software

The team chose a Dual Node Microway Octoputer cluster. We feel that there is a good balance of CPU and GPU performance in our configuration to allow us to excel in the varied applications of the competition. We are using EDR Infiniband for low latency RDMA between nodes. For maximum fun and excitement we decided to include an NVIDIA Jetson TX2 to liven up our Tersoff benchmarks.

- Intel Xeon E5-2699Av4 22-core 2.4 GHz CPU (2 per node for a total of 88 cores/176 threads)
- Infiniband EDR Interconnect in 4x link mode
- 8 NVIDIA V100 16GB GPUs
- 8 NVIDIA P100 16GB GPUs
- 1TB DDR4 ECC 2400 MHz memory
- Intel 750 Series 1.2TB PCI-E NVMe SSD

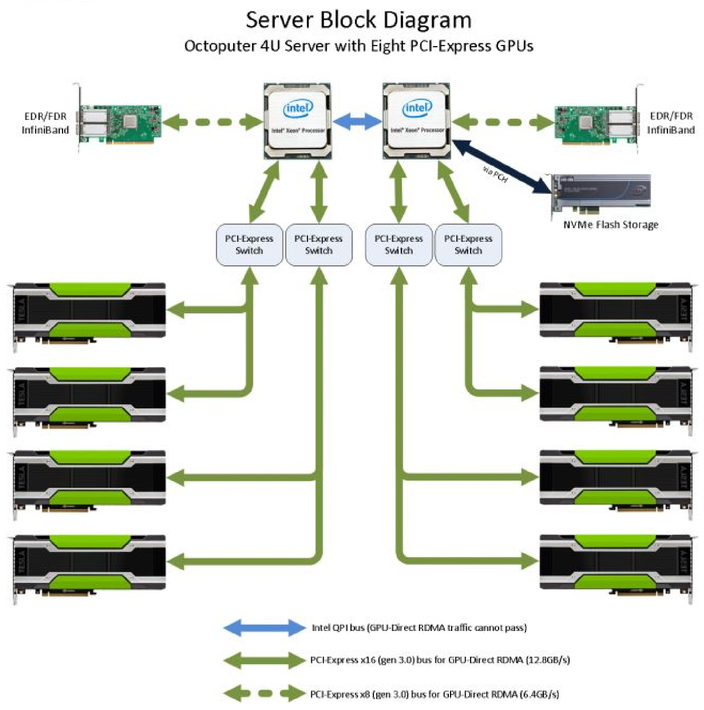

For our software stack, we decided to use CentOS 7 for its high configurability and widespread support in the HPC community. When applicable we used ICC compilers for their integrated support of Intel MKL and native support for CPU instructions. Slurm was chosen as a resource manager to ensure a constant workload to make the most of our power budget.

- Intel Math Kernel Libraries
- CUDA 9
- OpenMPI version 1.10.2
- CUDA FFTW Libraries
- CUDA OpenCL Libraries
- Slurm

Applications

- **HPL:**: The team's running an optimized version of the NVIDIA Volta architecture. For configurations, we experimented with the size of the problem using small deviations of 80% of our total RAM. Our benchmarking yielded a reasonable percentage of our theoretical peak performance. A key part of our experimentation was adjusting block sizes. Block sizes too small limited our performance because there was low data reuse whereas block sizes too large wasted space and extra computation.
- **HPCG**: a metric for ranking HPC systems driven by a multigrid preconditioned Conjugate Gradient algorithm that exercises the key kernels on a nested set of coarse grids. The team's running a CUDA optimized version. We experimented profusely with different matrix sizes each for 60 seconds to get a rough estimate on potential performance. For each result that seemed fruitful. we would run HPCG for 35 minutes and used the matrix size that would produce the most performance for our cluster.
- **MrBayes** is an application used for Bayesian inference in phylogenetic and evolutionary models. MrBayes uses Markov chain Monte Carlo (MCMC) methods to estimate the posterior distribution of model parameters. We have tried many different compiler options such as using GCC, ICC, mpicc and mpiicc. From our testing experience we found that using ICC compiler with OpenMPI 1.10.2 library gave us the best performance. We experimented with using the beagle library but found it not beneficiary on performance or time restrictions.
- **Tersoff**: The main task of the Tersoff multi-body potential was reproducibility of the results in the paper. So the software was compiled as per the recommendation of the paper on our architecture. Outside the recommendations of the paper, we ported Tersoff to the NVIDIA Jetson TX2 after substantial efforts in minimizing code changes. We think this will serve as a unique benchmark of reproducibility for Tersoff.
- **Born:** As each Born shot is a discrete task we will leverage Slurm to run shots when extra resources are available. We will also leverage our cloud allotment for Born.
- **Power Shutoff Activity:** The team will leverage built in checkpointing provided by MrBayes. LAMMPS provides checkpointing but the computation time appears to be low. We will be running Born in the cloud to leverage the inherent high availability that it provides.

Sponsors

Microway