BOSTON GREEN TEAM

The Team

- The Boston Green Team represents a diverse group of students from the University of Massachusetts Boston.
- The students have joined together to create a team with **diverse academic backgrounds** consisting of students from three majors - Computer Science, Computer Engineering, and Mathematics.
- Green team members come from member institutions of The Massachusetts Green High Performance Computing Center, a LEED Platinum Certified research data center in Massachusetts.
- Our team represents the **diverse backgrounds** of the Boston area with students from many different ethnic, gender, and socio-economic backgrounds.

Team members

- **Evan Donato** is a Mathematics and Computer Science major at the UMass Boston. He is the team captain and is lead on the OpenMC application.
- **Kristen Laird** is a Computer Science major at UMass Boston. She is lead on the Horovod application.
- **James Michaud** is a Computer Science major at UMass Boston. He has been in charge of benchmarking HPL and preparing for possible mystery apps.
- **Cristian Peguero** is a Computer Engineering major at UMass Boston. He is lead on the SeisSol reproducibility application.
- **Rudresh Trivedi** is a Computer Science major at UMass Boston. He is our systems administrator and is cross training on all aspects with the goal of being team captain next year.
- **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 travelling with the team as a backup and has helped with preparation.

Strategy

For benchmarking we are planning on running on the entire cluster. For HPL we will be scaling the wattage on the V100 GPUs to 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. Though the cluster can max out at around 5000 watts we will be limiting power consumption depending on the task to **keep just below 3000 watts.**

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 skill sets 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.

- 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 32GB GPUs
- 8 NVIDIA V100 16GB GPUs
- 1TB DDR4 ECC 2400 MHz memory
- Intel 750 Series 1.2TB PCI-E NVMe SSD
- 1 NVIDIA Jetson TX2

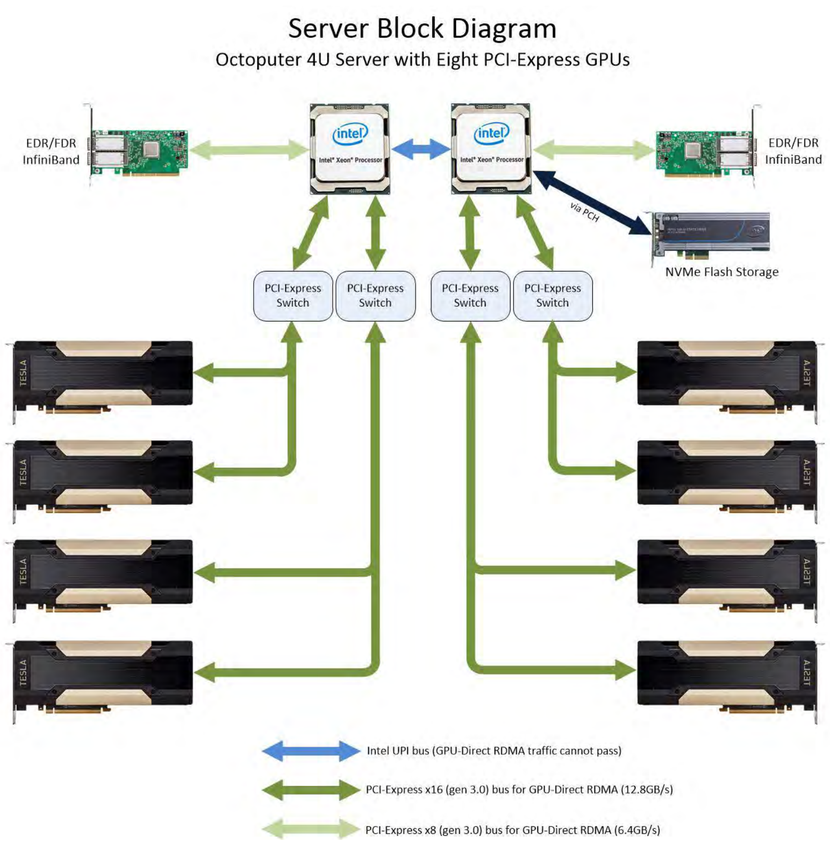

For our software stack, we decided to use CentOS 7 for its high configurability and widespread support in the HPC community. When applicable we used ICC compilers for their integrated support of Intel MKL and native support for CPU instructions. Slurm was chosen as a resource manager to ensure a constant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

- Intel Math Kernel Libraries
- CUDA 9/10
- OpenMPI
- CUDA FFTW Libraries
- CUDA OpenCL Libraries
- Slurm

Applications

- **HPL:** The team's running an optimized version for 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. Since the **optimization** allowed by the ini file is minimal we spent time making sure that CPUs are pinned to close memory and to the least number of hop GPUs.
- **Horovod:** Acting as the communication layer of TensorFlow, Horovod helps simplify the training process for deep learning applications. With build versions and model architecture being provided at the beginning of the competition, our team prepared by experimenting with sample models using synthetic imagenet input vectors, allowing us to get familiar with the application.
- **SeisSol**: This is the reproducibility HPC Software in which we will be trying to reproduce the results published in the "Extreme scare Multi-physics simulations of the Tsunamigeninc 2004 sumatra megathrust Earthquake" paper. The software will be compiled to run on our two node cluster and relies entirely on CPU processing power.
- **OpenMC:** We will be **optimizing OpenMC** by running OpenMP within each NUMA node. We will also bind processes to memory to help prevent cache misses. We will be using our **cloud allotment** for OpenMC allowing us to offload a CPU intensive task to a scalable cluster.
- **Power Shutoff Activity:** The team will leverage built in checkpointing provided by Horovod. Checkpointing for SeisSol is probably unimportant due to run length. We will be running OpenMC in the cloud to leverage the inherent high availability that it provides.

Sponsors

Microway