Team Phoenix: Georgia Institute of Technology

Ali Kazmi, Albert Chen, Aman Jain, Marissa Sorkin, Nicole Prindle, Sudhanshu Agrawal
Advisor: Vijay Thakkar (Team 11)

About Us: We are Team Phoenix! We are six computer science students with diverse academic and personal backgrounds. Our experiences include outreach programs that promote STEM to underrepresented youths and fostering inclusion for Women and LGBTQ groups in CS. We come from diverse geographical and ethnic backgrounds, representing five states and four countries. We study seven of the eight CS concentrations Georgia Tech offers.

Application/Benchmark Duos: For a balanced division of labor, we split into specialized pairs.

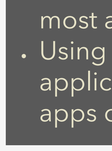
Ali

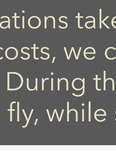
Nicole

IO500
MemXCT

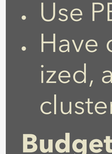
Aman

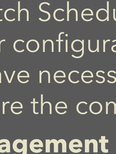
Marissa

HPL
CESM

Albert

Sudhanshu

HPCG
GROMACS

Hardware and Software Configurations

	HPL	HPCG	IO500	MemXCT	GROMACS	CESM
SKUs	Planning to use **NDv2** Series VMs 8 NVIDIA Tesla V100 GPU, 40 vCPU Has NVLINK, high bandwidth connection between GPUs	Use **NC24r** nodes because of RDMA support with high memory bandwidth. If CPU only config is better, use HC-series which has Intel MKL library	1x Management node (**Av2** standard) 2x High-throughput nodes for storage and MDS (**D64d_v4**) These nodes have fast and large SSD storage	CPU: Use HB Nodes (**HBv2** if available, else **HB60**) Memory bandwidth is biggest bottleneck GPU: Use NVIDIA Tesla V100s (**NCv3**) Performance benefits from GPU cache optimizations	Use **NC24s_v3** to leverage compute of V100 GPUs GROMACS works best with several high clock speed cores per GPU	Use HB-series nodes (**HBv2** if available, else **HB60**) Higher core count leads to better performance Leverage RDMA to reduce IPC bottleneck
Software	Use NVIDIA binary, which uses OpenMPI and CUDA Setup NVIDIA Tesla drivers before competition Use Intel MKL for optimized linear algebra	Use either NVIDIA's optimized binary, or Intel's optimized MKL based binary based on performance Use Intel MPI on whichever binary we choose	**BeeGFS** file system for high throughput and ease of use Use Intel Compilers and Intel MPI	Aim to compile using AOCC for better performance Install CUDA for GPU version Tune tile, block, and buffer sizes on PACE to save Azure credits	Configure nodes with **Singularity**, use the GROMACS container available in NVIDIA GPU Cloud registry NVIDIA container contains optimized modules and drivers	Use GNU compilers Build CESM Image on top of Azure CentOS for easy deployment Optimize PE layout on PACE-ICE. This should translate across machines.

Competition Strategy:

Cluster Configuration

- All Clusters based off CentOS 7 with updated compilers
- This is the default image, it will have the best compatibility with Azure Infiniband.
- Use PBS Batch Scheduler.
- Have cluster configurations for each app/benchmark finalized, and have necessary dependencies installed on each cluster before the competition.

Budget Management

- Split budget 20% for benchmarks, 80% for applications. During the competition, we will try to move funds from benchmarks to apps whenever possible.
- Most benchmarks take little time to run if we saturate GPUs, and most applications take many cores and days to run.
- Using SKU costs, we can calculate cost per node-hour for each application. During the competition we can allocate nodes for apps on the fly, while staying under budget.

Preparation Strategy:

Georgia Tech Class on Student Cluster Competition:

- Gained familiarity with HPC concepts
- Weekly meetings made sure we had constant engagement
- We get to know each other, and learn about each other's apps

Application/Benchmark Preparation

We will use PACE-ICE, a Georgia Tech HPC cluster, to tune applications. That way, we can use the majority of our Azure practice credits to tune benchmarks. Benchmarks are sensitive to hard ware and we want to minimize the costs of running benchmarks so we will tune those directly on Azure.

Applications are sensitive to configurations, so we will use PACE-ICE to experiment with how input parameters affect performance. We can extrapolate from our observations during the competition. We have surveyed past SCC apps to prepare for the mystery app.

Why We Will Succeed:

University Support:

We have **world class faculty** helping us prepare
Dr. Richard Vuduc, Dr. Aaron Jezghani, Dr. Jeff Young, Will Powell
Georgia Tech has a **strong SC presence** and past SCC experience
The College of Computing is a **leader in computing education**
We have a **specialized course** for SCC preparation

The Team:

We are **talented** and **motivated** to win
We are **passionate** about computer science
3 of us TA for core computer science courses
4 of us do undergraduate research
All 6 of us have industry experience
We have been **working together** since January

Industry Support:

Penguin Computing supports Georgia Tech and PACE
Intel and **NVIDIA** provide us with HPL binaries
Industry connections help us prepare