

Team

We are members of the [KU Supercomputing Club](#) (KSC)! Our team is extremely **diverse**, as we represent **four disciplines** ranging from **computer science** to **neuroscience**, represent **four nationalities**, and have combined experience in HPC from **industry**, **academia**, and a **national lab**.

Abir Haque

- **Major:** Computer Science
- **Skills:** Parallel Programming, Scientific Computing
- **Experience:** Parallel Mesh Algorithm Research, ACM SRC Participant at SC24, SCC23 Participant
- **Responsibilities:** HPL, ICON



Owen Krussow

- **Major:** Computer Engineering
- **Skills:** Systems Administration, Parallel Programming
- **Experience:** System Administrator at I2S, LANL Intern, Cluster Competitions
- **Responsibilities:** MLPerf



Yara Al-Shorman

- **Major:** Computer Engineering
- **Skills:** Digital Systems Design, Computer Architecture, Software Engineering
- **Experience:** Bioinformatics Research, SCC23 Participant
- **Responsibilities:** NAMD, ICON



Wazeen Hoq

- **Major:** Behavioral Neuroscience, Mathematics
- **Skills:** Data Science, Mathematics, Biology, Chemistry
- **Experience:** EEG Data Analysis Research
- **Responsibilities:** HPL, NAMD, ICON



Michael Oliver

- **Major:** Computer Science
- **Skills:** System Administration, Cyber Security
- **Experience:** Cyber-Lab for Cyber Defense Course, Non-Terrestrial Networking Distribution using Bluetooth Mesh, System Administrator at I2S
- **Responsibilities:** MLPerf, ICON

Shad Ahmed Shahul Hameed

- **Major:** Computer Science, Mathematics
- **Skills:** Machine Learning, Software Engineering
- **Experience:** Machine Learning Research, Software Engineer at Microsoft, Data Science
- **Responsibilities:** MLPerf



Amin Mamandipoor (Primary Advisor)

- **Position:** PhD Student
- **Major:** Computer Science
- **Specialization:** Scalable Distributed Computing
- **Skills:** Computer Architecture and Systems



Benchmark Strategies:

Benchmark: *High Performance Linpack (HPL)*

- Solves a dense linear system on a cluster using LU factorization with partial pivoting
- Measures **compute capabilities** of cluster
- Compare **NVIDIA GPU-accelerated**, **multi-core**, and **base** versions of HPL to determine optimal configuration for final HPL runs
- Compare various **BLAS implementations** to yield best performance for matrix and vector operations.
- Focus on determining **best tuning file** based on number of nodes, cores per node, GPU resources per node, total memory, and block size

Benchmark: *MLPerf*

- Benchmark suite for evaluating a cluster's performance in common machine learning tasks
- Experiment with **mixed-precision** training to improve performance
- Compare various deep learning frameworks and optimize them for best performance
- Experiment with **batch-size** and **parameter** tuning
- Test different optimizers such as LAMB or LARS to **improve scalability**

Application Strategies:

Application: *Nanoscale Molecular Dynamics (NAMD)*



- Parallel molecular dynamics package for simulating large biomolecular system
- Compare performance when utilizing various **compilers** (e.g. GCC, ICPC, etc.) to determine best binaries for final application runs
- Compare performance when utilizing **FFTW library** versus other Fourier transform implementations
- Rebuild **Charm++** with various configurations, such as SMP/shared memory, to find the best configuration
- Preliminary experiments show best performance when utilizing **GPU acceleration**

Application: *ICOsehdral Nanohydrostatic Model (ICON)*

- High-performance modeling framework for weather prediction
- Compare performance when compiled with different **MPI implementations** (OpenMPI, MPICH, etc.)
- Experiment with different **dependency packages** and **libraries** to determine best performance
- Compare **BLAS** and **LAPACK implementations** for optimal performance
- Optimize **input grid decomposition** and determine best fit within processor grid

$$A = LU$$

$$LUx = b$$



Software:

Operating System: *Rocky Linux 9*

Scheduler: *SLURM*

- **Scalable** for both small and large clusters
- **Practical** for both limited VM configuration granted for homework tasks and generous configuration granted for competition runs
- Many clusters on Top500 utilize SLURM as their job scheduler
- Team members have **prior experience** setting up and utilizing SLURM

Automation: *Ansible*

- **Quickly configure** important services, such as SLURM, NFS, on **large number of nodes**
- Developing **playbook** to utilize during hero and competition runs
- Team members have **prior experience** setting up and utilizing Ansible

Team Preparation:

- The KSC competition team holds weekly, hour-long meetings that feature **hands-on labs** related to parallel programming, benchmarking, and running applications
- Spent the summer **optimizing benchmarks** and **applications** on **club cluster** generously donated by KU faculty and staff
- Actively **collaborating** as a team when completing **IndySCC homework**
- Rigorously **experiment** with software stack, benchmarks, and applications to deliver **optimal performance** for competition.

Why Will We Win?:

- Having an extremely **diverse team**, such as KU's team, results in **innovative approaches** to benchmark and application optimization.
- We have the **support** and **guidance** of **advisors** who provide direction on the optimization techniques.
- Team members have **previous experience** optimization benchmarks and applications in cluster competitions.