

Team Supernova



Students: Li Shenggui, Dong Yunxing, Pang Jin Hui, Tan Jia Qing, Aurelio Jethro Prahara, Do Xuan Long
Advisors: Prof. Bu-Sung Lee (NTU), Mr. Terry Yin (NVIDIA AI Tech Center), Dr. James Chen (AMD), Mr. Liu Siyuan, Mr. Jeffery Tay (NTU)

Team & School

Inaugurated in 1991, Nanyang Technological University (NTU) has grown to become a comprehensive and research-intensive university. Our university is consistently ranked as one of the top universities in Asia. In NTU, the High Performance Computing Centre (HPCC) provides HPC resources to the entire NTU research community. What's more, NTU's School of Computer Science and Engineering offers a HPC specialization to computer science and computer engineering students.

ISC16

First appearance in ISC and ranked 3rd in HPL benchmarking

ISC17

Won Deep Learning Excellence Award and ranked 4th overall.

SC18

Won highest LINPACK award again with another overall second place

ASC15

Broke the ASC LINPACK record under 3kW power budget

SC16

Ranked 1st in secret application (GROMACS) and 5th overall.

SC17

Won both Overall Championship and Highest LINPACK Award (breaking the world record with 51.77 TFLOPS/s)

SC19

Won highest LINPACK award

Our team is a truly diverse team. We have a mix of nationalities in the team, reflecting the diversity of NTU. The six members also have various background in computer science and mathematical science, including algorithm design, operating systems, artificial intelligence, and high performance computing. We together work hard to prepare for the competition.



From left to right: Dong Yunxing, Li Shenggui, Pang Jin Hui, Tan Jia Qing, Aurelio Jethro Prahara, Do Xuan Long

Cloud Configuration

SKU	Cores	Memory	\$/Hour	RDMA	Storage
HB60rs	60	233.52GB	\$2.51	Yes	700GB
HC44rs	44	327.83GB	\$3.48	Yes	700GB
NC24r	24	224.00GB	\$4.32	yes	1.41TB

We will be using the HB and HC series which are designed to deliver superior performance, MPI scalability and cost-efficiency for HPC applications. HB60rs is optimized for applications which heavily rely on memory bandwidth and it can provide more than 260 GB/sec of memory bandwidth. The HC44rs is more suitable for applications with dense computation with its powerful 44 Intel Xeon Platinum 8168 processor cores. These two VM configurations should provide enough resources for our applications. We may also use GPU for applications such as MemXCT. Therefore, the NC24r is a good choice as it has high memory bandwidth. However, the cost factor should be carefully considered since NC24r VM is relatively more expensive.

On the software side, we are using Cycle CentOS 07 with NVIDIA CUDA Toolkit 10.0, Intel Parallel Studio 2019, Mellanox HPC-X, etc. These software suites are well established and enable us to gain the most performance out of our cluster. We use Azure BeeGFS for storage, which is optimized for high performance computing. We also use utilities such as Singularity, Slurm and PBS to better manage our cluster.

Applications Optimization

MemXCT

MemXCT is a memory-centric approach to X-ray CT reconstructions which allows for scalability with parallelisation. It has been tested on several supercomputer architectures involving KNL and GPU. In this competition, we are to reproduce the benchmarks and experiment with its speedup and scalability. Our approach is to

1. carefully study the paper and consult experts in the field of bioinformatics to better understand the content
2. prepare scripts for running experiments and analysing results
3. design experiments on different system architecture and analyse the difference in performance



CESM stands for the Community Earth System Model (CESM), and is a fully-coupled, global climate model that provides state-of-the-art computer simulations of the Earth's past, present, and future climate states. Our approach is to

1. Build the model by following the documentation and test it on various sample input
2. Experiment on different configurations and compare the results
3. Find out the system architecture that maximizes performance



GROMACS is a versatile package to perform molecular dynamics, i.e. simulate the Newtonian equations of motion for systems with hundreds to millions of particles. Calculating the non-bonded operations is the most computationally intensive part in molecular dynamics simulation. As such, in this competition, the goal of this application is to explore the optimization of GROMACS. To achieve this, we will follow the roadmap below

1. Experiment on different option flags in run configuration file and mdrun command
2. Explore system architectural modification that can improve the performance



IO500 is a benchmark that provides comparison between different high performance storage systems. It includes workloads like metadata, I/Os and finding task tests. It can be optimized through the filesystems such as BeeGFS and NFS in Azure CycleCloud and utilizing the system call transition to reduce latency.

LINPACK benchmarks are introduced to measure a system's floating point computing power. It solves a random system of linear equations in 64-bit arithmetic on distributed-memory computers. Linpack can be measured on a specific instance of Azure cloud to determine performance across different cluster sizes and its efficiency using different scaling.



HPCG benchmark is a supercomputing benchmark test. It models the data access pattern to illustrate system performance on more communication intensive sparse matrix patterns. The benchmark is able to be run on Azure cloud with different number of instances in a certain running time.

Execution Strategies

1. Our goal is to maximize the utilization of our cluster throughout the competition under the Azure CycleCloud budget limit.
2. We carefully weigh the pros and cons of using different virtual machine series for each application, and closely monitor the budget consumption and adjust our workloads accordingly to avoid going over the budget.

Why We Will Win

1. Passion is the key! We are enthusiastic computer science students who are willing to work together to solve challenging problems.
2. We have clear teamwork allocation and organized competition plan to allow for frequent discussion and make sure we reach our goal.
3. We consulted domain experts from the academia as well as the industry to gain deep understanding of the theoretical knowledge and common practices.
4. We participated in previous competitions such as ISC 20 and APAC HPC-AI 20 to accumulate hands-on experience in HPC.

Special
Thanks

