

Students: Bryan Guanrong SHAN, Chanpaisit Nattapol, Danxu ZHANG, Saeng-Nil Nattakan, Woon Yee NG, Ze Ming CHUA


Advisors: Assoc. Prof. Bu Sung LEE (NTU), Dr. Yuan Ren LOKE (NTU), Mr. Siyuan LIU, Dr. James CHEN, Mr. Shenghao YANG, Mr. Terry YIN

Meet the School and Team




Nanyang Technological University

- Established in 1991 in Singapore, is a globally renowned university focused on research and innovation. It excels in areas like artificial intelligence, sustainability, and biomedical sciences, with world-class research centers and an HPC-supported Parallel and Distributed Computing Lab.



A Truly Diverse Team:

- 4 Nationalities: Singapore, Malaysia, Thailand, China
- Diverse majors and field of interests:
Artificial Intelligence, Data Science, Parallel Computing, Quantum Computer, Physics, Computer Network, Cyber Security, Hardware Acceleration, Business
- Languages:
English, Chinese, Malay, Thai, Japanese, Cantonese, Hokkien
- Diverse experiences:
*National Serviceman @ Singapore Army
Intern @ A*STAR - Agency for Science, Technology Research
Intern @ Centre for Climate Research Singapore
Intern @ Agoda, Hewlett-Packard Inc and etc.*



Woon Yee NG
• Computer Science Y4
• ISC24
• ICON, MLPerf



Danxu ZHANG (Captain)
• Computer Science Y4
• ISC24
• ICON, Reproducibility, SysAdmin



Chanpaisit Nattapol
• Computer Engineering Y3
• ISC24
• NAMD, Reproducibility



Ze Ming CHUA
• Computer Science Y2
• New Member
• HPL, SysAdmin, Hardware Management




Saeng-Nil Nattakan
• Business and Computer Engineering Y3
• SC23, ISC24
• NAMD, Mystery Benchmark



Bryan Guanrong SHAN
• Business and Computer Science Y2
• ISC24
• MLPerf, SysAdmin, Hardware Management


Why we will win!

Good Team Dynamics and Collaboration



Learning from Past SCC Experience

Passion for HPC



Powerful Hardware and Well-Designed Architecture

- ✓ Extensive Hands-on Experience
- ✓ Regular Meetings and Training
- ✓ Adequate Preliminary Testing
- ✓ Discussions with Domain Experts
- ✓ Deep Understanding on the background of the Task


Our Cluster: Coffeepot



Ubuntu 22.04





NVIDIA
ConnectX-7
Infiniband



6x AMD EPYC 9654
96-Core CPU



12x NVIDIA
H100 GPU



Coffeepot
Cluster

by coffee enthusiasts



Prometheus
TimeSeries DB



Grafana
Visualization



Telegram Bot
Alerting

Hardware Choice:

- 3x Supermicro GPU A+ Server AS-4125GS-TNRT2 servers
- Dual sockets AMD EPYC 9654 96-Core Processors for each server
- 6x Samsung NVMe SSD
- 12x NVIDIA H100 accelerators with 350W power and 80GB memory
- NVIDIA ConnectX-7 400Gbps InfiniBand Network
- Raspberry Pi for system monitoring

Software Choice:

OS: Ubuntu 22.04 with Linux Kernel 6.5.0-27

Compilers:

- AMD AOCC 4.2, Intel OneAPI 22/23/24, GNU gcc 12/13/14,
- CUDA 11.8/12.4, nvhpc 24.7, clang-LLVM

MPIS:

- MVAPICH, Intel MPI, OpenMPI4/5 with CUDA and OpenUCX

Profiler:

- IPM, ScoreP, NVIDIA's Nsights, Intel Vtune, Darshan, TAU

Tools:

- BeeGFS, NFS, EasyBuild, LMOD, Docker, Singularity, SLURM
- Ansible, Prometheus, Grafana, IPMITools, Terraform

Libraries:

- MKL, OpenBLAS, HDF5, NetCDF, ScaLAPACK, fftw, OpenFOAM...

Application Optimization

General Methodology

- Evaluate compilers with **various optimization flags and compare library versions** to avoid compatibility issues.
- Profile system performance** using TAU, IPM, mpiP, Nsight, and Intel Vtune to identify bottlenecks and guide optimization efforts.
- Engage with relevant research and **consult professors and industry experts** to refine our approach.
- Test GPU performance by **experimenting with different GPU topologies on nodes** to achieve optimal results.
- Utilize **OpenACC with device offloading** to efficiently migrate computational kernels to accelerator cards.

Benchmarks


- HPL – Optimize floating-point performance, ensure efficient memory bandwidth usage, and maximize parallelism for solving dense linear systems with LU factorization across multiple nodes.
- MLPerf - Optimize GPU utilization and multi-GPU scaling, streamline data loading and memory usage, make sure the efficient distributed training and inference.
- Use the latest binaries provided by NVIDIA with careful fine-tuning under power budget.

Mystery Benchmark

- Prepare by running HPCG, OSU Micro-benchmark, and Stream to tune our system for handling general cases

NAMD

- A parallel molecular dynamics code designed for high-performance simulation of large biomolecular systems.
- We will be using profiler to find the hotspots, implement memory optimization to reduce the bottleneck and proceed to finetune MPI for a better performance.




ICON


- A modelling framework for weather, climate, and environmental prediction.
- We will start with profiling, find out memory, computation and network constraints, then leverage our domain knowledge in ICON in ISC24 and other weather models to finetune parameters.

Reproducibility Challenge

- A scientific workflow analyzer for improving data placement and task scheduling.
- We carefully read the paper and familiarize our selves with different workflow setup and scheduling.
- Different number of VMs and containers will be tested on multiple nodes to simulate data movement behaviors of workflows at different scales.



Running Strategy



Power Consumption, Monitoring, Noise Level Control, Automation on Deployment

- Monitored memory, CPU, GPU, and power using Prometheus and Grafana, controlled fan speed with IPMI tools.
- Automated system management across 3 nodes with Ansible Playbooks and Terraform for consistency.
- Implemented fine-grained hardware control using tuning scripts, leveraging NTUHPG Wiki references.
- Noise control with cushion and foam for sound insulation

Runtime Schedule and Scripts Writing

- Before scripting, we'll assess HPC application scalability across nodes and cores to identify bottlenecks in load balancing, communication, and memory usage.
- Next, we'll optimize and profile MPI settings, compiler flags, and other parameters to ensure peak performance and consistency.
- Finally, using our pre-written scripts, we'll run benchmarks to verify correct system setup during the competition.
- SLURM is also used to manage multiple tasks and cool-down between jobs.