



SC22

TEAM SUPERNOVA

Nanyang Technological University, Singapore

NTU HPC

Singapore

Students Aurelio Jethro Prahara, Di Wang, Jia Qing, Ngoc Yen Duong, Ruisi Wang, Shaojie Zhong

Advisors Assoc Prof Lee Bu Sung (NTU), Dr. Loke Yuan Ren (NTU), Mr. Siyuan Liu (NSCC), Mr. Jianxiong Yin (NVIDIA), Dr. James Chen (NSCC), Mr Yang Shenghao

MEET THE TEAM AND SCHOOL

Nanyang Technological University, Singapore

Inaugurated in 1991, NTU has grown to become a research-intensive and comprehensive university. In NTU, the High Performance Computing Centre provides HPC resources to the entire NTU research community. NTU also offers HPC specialization to computer science and computer engineering students.

A Truly Diverse Team

Our HPC team is formed as part of the HPC club, a student-run club that promotes HPC adoption and awareness on campus. We organise weekly training sessions and provide help for students spanning a wide range of schools.

With diversity comes innovation and success.

• Diversity of gender:

2 female and 4 male competent team members

• Diversity of home country and culture:

members from 5 different countries

• Diversity of personal interest:

Different majors: Electrical and Electronic Engineering, Computer Science, Physical and Mathematical Sciences

Interest in different fields:

Artificial Intelligence, Data Science, Integrated Circuit Design, Hardware Acceleration, Computer Vision

Team Members in SC22 Student Cluster Competition



Aurelio Jethro

- algorithm, AI, math
- SC20, ISC21, ISC22
- reproducibility, HPL, HPCG




Di Wang

- system admin, hardware
- ISC22
- IO500



Jia Qing

- compiling, design
- SC20, ISC21, ISC22
- IO500



Ngoc Yen Duong

- compiling, algorithm
- ISC22
- cloud, PHASTA



Ruisi Wang

- algorithm, computer vision
- ISC22
- LAMMPS, mystery



Shaojie Zhong

- algorithm
- new member
- reproducibility, HPL


WHY WE WILL WIN?

Passion for HPC



Good team dynamics and collaboration

Learning from past SCC experience



Powerful hardware and well designed architecture

Preparation

• Wide breadth of practical experience

- experienced in administering and designing small-scale HPC systems
- in charge of managing clusters in the Parallel and Distributed Computing Laboratory
- attended conferences and events to accumulate experience
- some have attended other HPC competitions

• Collaboration

- applications are allocated to multiple members based on experience
- diversity and complementarity of ideas

• Regular meetings and training

- meetings are regularly held on campus for members to train and work together
- sharing and brainstorming of ideas

Support

• Prepared everything in cluster-init

- to save time in competition

• Compared compute optimised, memory optimised and HPC series

- on HPC series we can use RDMA


• Conducted sufficient beforehand testing

- so that we are able to choose the most suitable virtual machine type based on the performance/cost ratio after the budget is announced onsite

HARDWARE AND SOFTWARE

Cluster Overview

- **Scaling Nodes x2**
 - Good balance between CPU and GPU computational capabilities
 - Uniform system config
- **DDR4 & SSD**
- **Infiniband EDR**
 - NVIDIA Networking Connect X-6 EDR InfiniBand adapters per node
 - Low latency with high bandwidth
- **Gigabit Ethernet**
 - Stable for management and monitoring



Hardware Choice

- **High-density GPU cluster**
 - Improve the throughput of GPU applications (HPL, HPCG, LAMMPS)
- **CPU application**
 - tackle PHASTA with cloud component
 - reproduce DaCe with 128 cores
- **AMD 7763, Intel E4-2699, NVIDIA A100**
 - AMD EPYC 7763 (64 cores * 2) and Intel E5-2699 V4 (Broadwell, 22 cores * 2)
 - 40G A100 and 80G A100



NVIDIA A100 x12

- 1555 GB/s memory bandwidth
- 250 W power consumption

Software Choice

- **Gigabit Fedora 35 with 5.15 LTS Kernel**
 - high-quality Linux distribution that gives HPC complete control
- **NVIDIA CUDA Toolkit**
 - Highly optimized GPU-accelerated libraries
- **Intel OneAPI**
 - Intel compiler performs better and is relatively easier to use than GCC
 - Intel MPI is a high performance MPI implementation
- **Mellanox HPC-X**
 - optimised message communications based on OpenMPI for Mellanox InfiniBand
- **MPICH, Intel MPI and OpenMPI MPI libraries**
 - for performing parallel computation
- **Lmod**
 - ease of management for multiple versions of libraries
- **Ansible**
 - to automate server configuration & software installation

APPLICATION OPTIMIZATION

General Methodology


- **Compilation**
 - compare different compilers with various optimization flags
- **Libraries**
 - compare different version of dependencies to avoid incompatibility
- **Profiling**
 - profile with Arm Forge and Intel Vtune in order to find hotspot and find a direction for optimisation
- **Domain Knowledge**
 - study relative researches and bring doubts to professors and industry engineers in respective field

Benchmarks

HPL solves a dense system of linear equations with LU factorization.

HPCG solves a sparse one with preconditioned conjugate gradient.

- ✓ **prepared tuning scripts** including automatic GPU & CPU frequency
- ✓ **fully understood input parameters** e.g. N, TRSM
- ✓ **use latest binaries** provided by NVIDIA
- ✓ **carefully tuned** parameters with power monitors



Production-level Molecular Dynamics simulator that has been optimized parallel as well as single node implementation.

• biomolecular simulation capabilities

• spatial decomposition technique

divide the physical geometry to be simulated into small boxes, one per processor.

• neighbour-list technique

keep track of neighbouring atoms
balance between storage and communication

• bonded and non-bonded force calculations

Compute-intensive: non-bonded interaction between atoms

PHASTA

Computational fluid dynamics application, which supports modeling compressible or incompressible, laminar or turbulent, steady or unsteady flows in 3D, using unstructured grids stabilized finite-element formulation for fluid dynamics. In this competition, our goal is to optimise the running time of PHASTA. To accomplish this, we follow the following roadmap:

- Carefully study the related papers and consult the expert related fields to understand the application clearly.



DaCe

Python library which provides extensions to Python for High-Performance Computing purposes

RUNNING STRATEGIES

Cloud Component

- **Prepared everything in cluster-init**
 - to save time in competition
- **Compared compute optimised, memory optimised and HPC series**
 - on HPC series we can use RDMA
- **Conducted sufficient beforehand testing**
 - so that we are able to choose the most suitable virtual machine type based on the performance/cost ratio after the budget is announced onsite

Power Consumption

- **Built up expertise in monitoring the power both when idle and when more precision is needed**
- **More fine-grained control over hardware so as to control the power needed**

Runtime Schedule

- At the beginning of the 48-hour competition, we'll run DaCe locally in order to get the data output for writing report.
- Then we'll run LAMMPS and PHASTA on our 2 local nodes and compile mystery locally.
- We will decide whether to run mystery locally or on cloud onsite. If we run mystery locally, it will be run on second day and we'll balance mystery and LAMMPS based on the scoring scheme.
- We will use a job scheduler to optimise the execution of tasks

Our Cluster

