

電機資訊學院 2024 作 BRAIN PLUS HAND 實作專題競賽

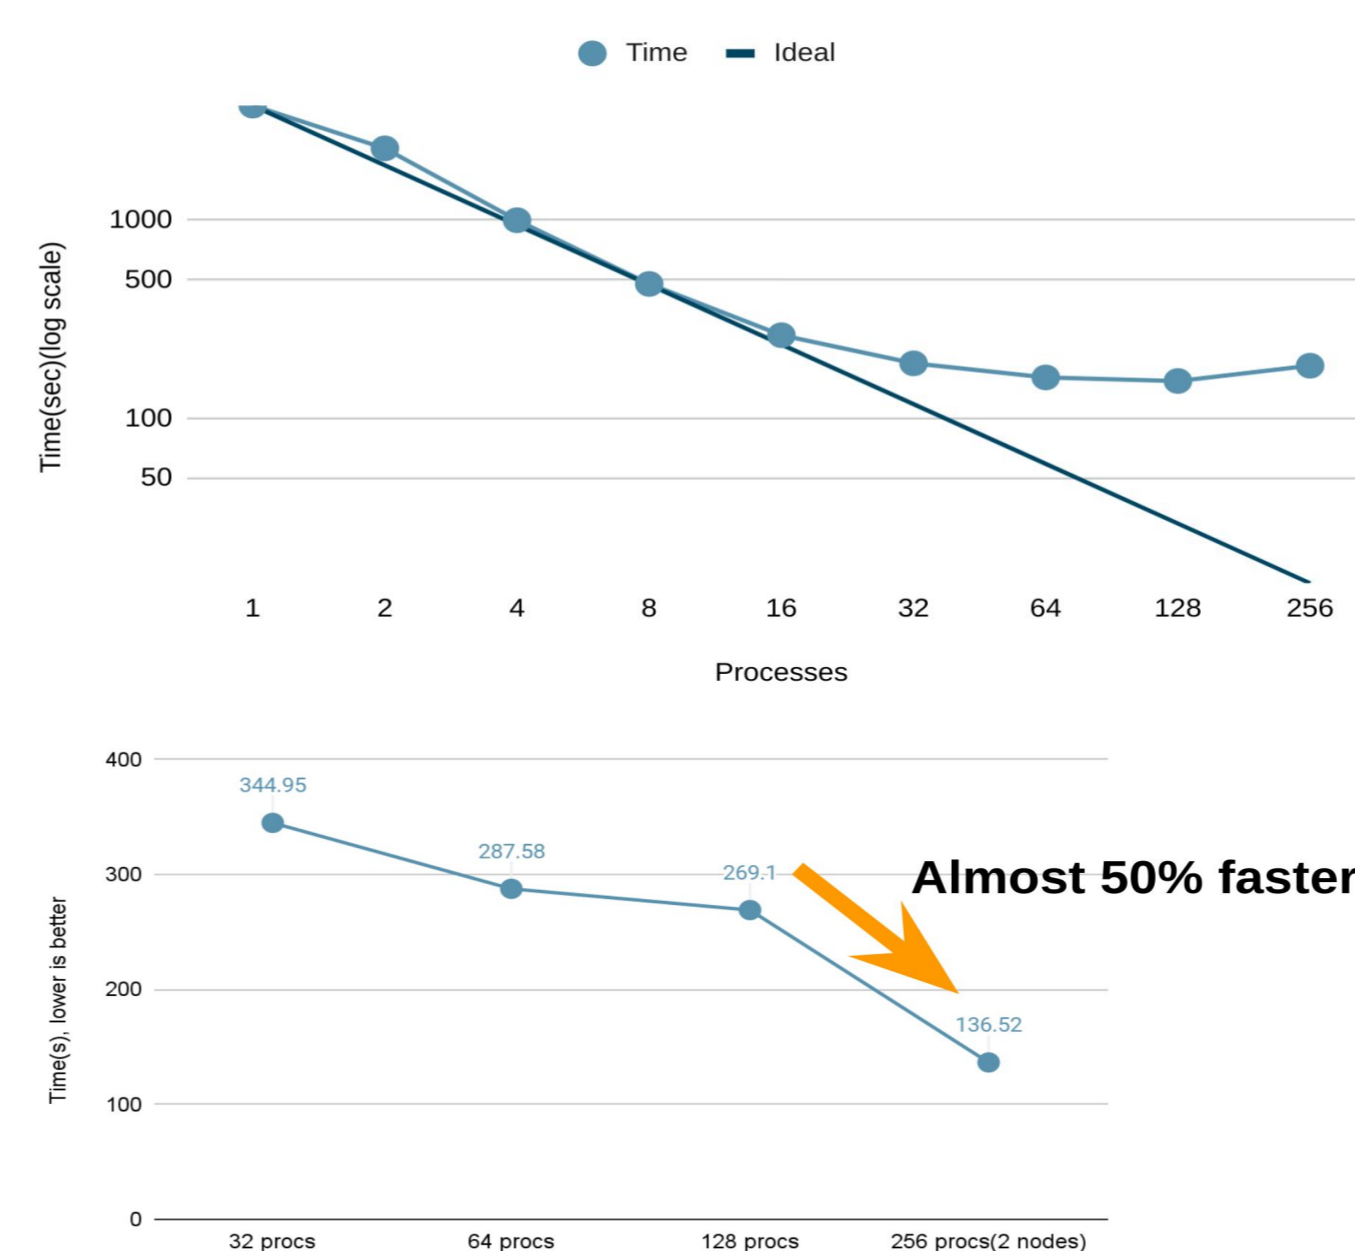
Team EECS005: Student Cluster Competition at SC22

隊伍成員: 吳邦寧、郭品毅

Applications

PHASTA

- Vectorizing: utilize hardware features
- CFD: have a scientific idea of how PHASTA works
- Automated Scripts: run & resume from checkpoints
- Profiling: get hints for optimization and compare different build configurations
- Linear System Solvers: experiment with different solvers to achieve the best performance



Benchmarks

HPL

- N: Problem size
 - o The HPL workload creates a matrix with a total size equal to $(N^2 * 8)$ bytes. In general, the best performance comes from a matrix that uses about 80% of total memory available.
- P, Q: Process grid
 - o $P * Q =$ number of process used
 - o Typically, determining P and Q with the following rules would result in better performance
 - P, Q are even
 - $Q/4 \leq P \leq Q$
 - $P = 2^n$
- NB: Block size
 - o The best value for this parameter varies. We would have to try different sizes to find the best NB value.
 - o Common sizes of NB are 96 - 512, and should be a multiple of the number of processes used.
 - o Lower values would result in higher efficiency, but the bandwidth would become the bottleneck
- **Result: 49.14 TeraFLOPs**

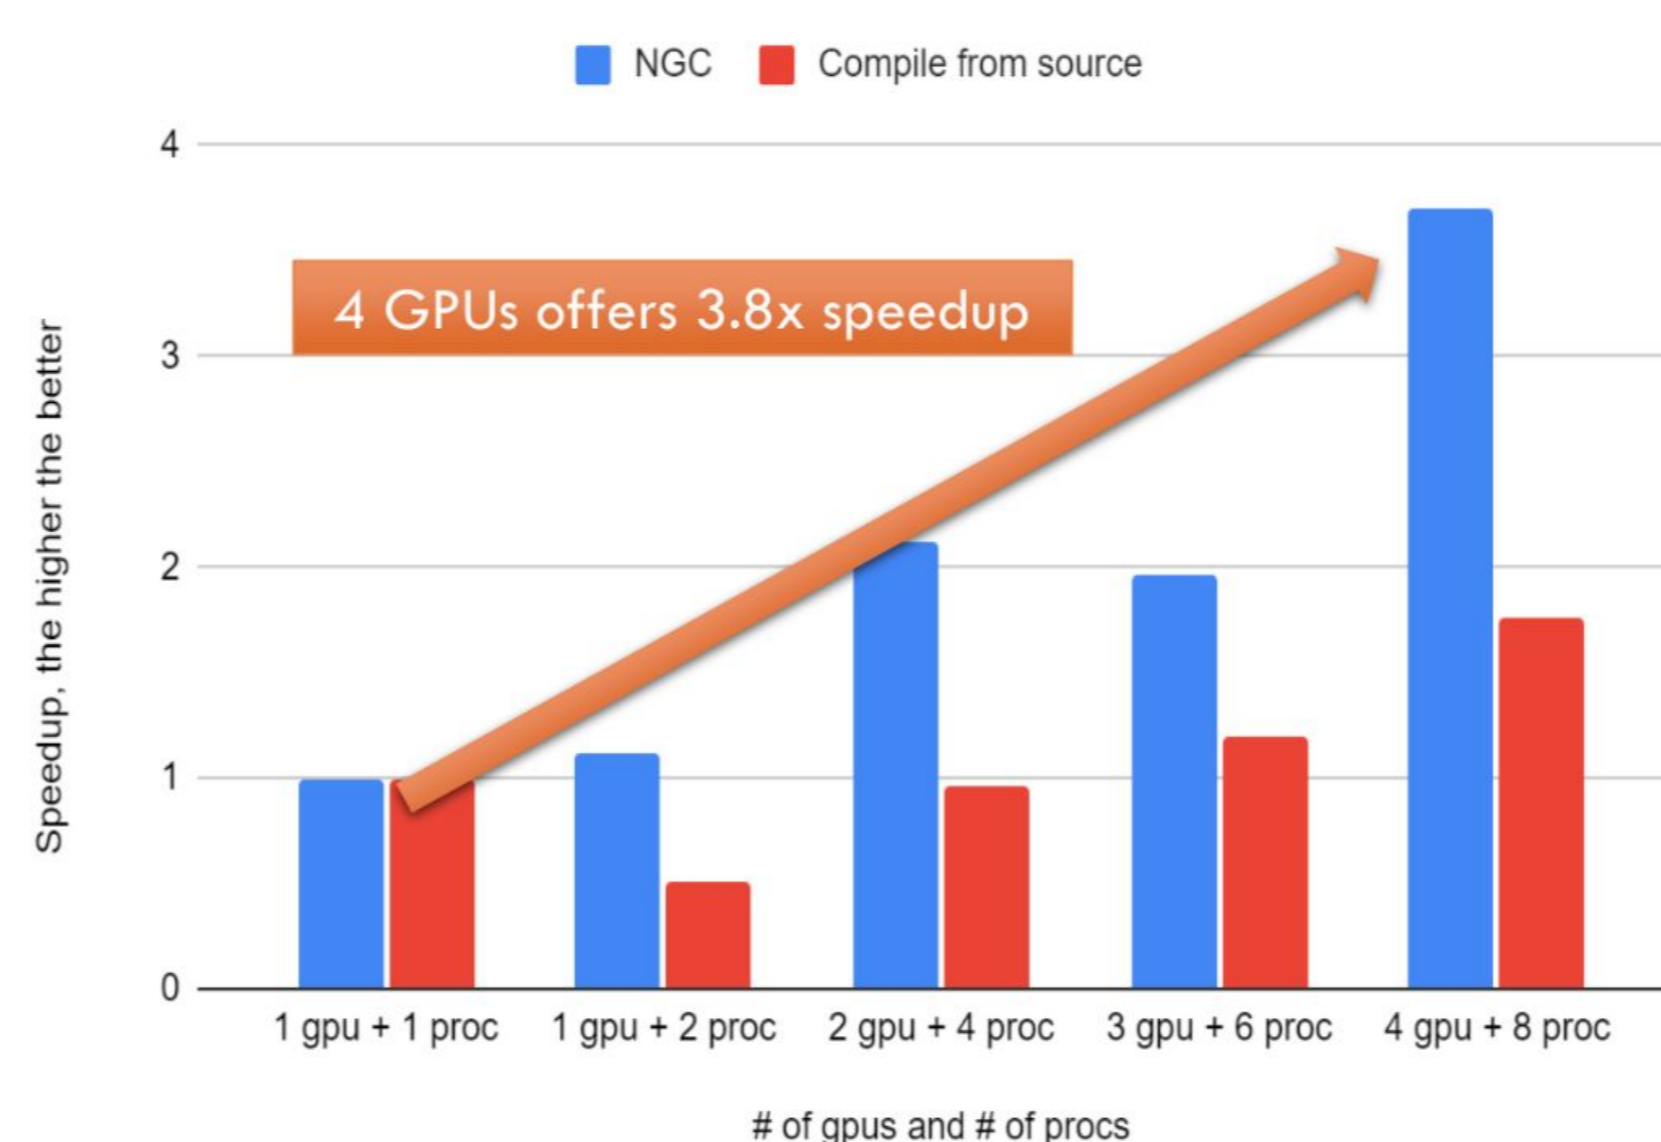
LAMMPS

Paper survey

- Single node: employ GPUs
- Multiple nodes: CPU only for better scalability

Experimental result

- Hyper-threading should not be applied on LAMMPS
- Use KOKKOS acceleration package for GPU speedup
- Scalability on OCI (Oracle Cloud Infrastructure) is poor for LAMMPS



HPCG

Tuning

- Initiate runs with shorter runtime (around 30 seconds)
- Explore N_x, N_y, N_z between 64 to 1024
- Identify the optimal value and run with a valid runtime (30 minutes)

Result

- $(N_x, N_y, N_z) = (128, 256, 512)$
- Run for 1800 seconds (30 minutes)
- Run with 6 A100 GPUs
- **962.156 GigaFLOPs**

SeisSol

- Mystery Application: Announce the task during the competition.
- Successfully built a CPU-only application that completed the first task with modest performance, but it's too slow for the other two tasks.
- Got great score by completing the second task that requires GPU optimization, which **other team were not able to build the GPU task.**
- Successfully built the unique GPU version by self-building dependencies and resolving linking errors with custom linker flags.

IO500

Tuning

- Hardware RAID
- Linux RAID (mdadm)
- ZFS
- Parallel File Systems

Result

- Linux RAID 0 ext4 > ZFS and Parallel File Systems
- EXT4 on Samsung U.2 NVMe SSD x2
- Bandwidth = 2.691 GiB/s
- IOPS = 195.9909

Educational Efforts

Training Camp

- Host an annual training camp
- Passing knowledge to juniors & Provide them with practical experiences
- Topics evolve year by year to meet new challenges and emerging issues

Time	2/6	2/7	2/8
10:00-12:00	SCC Introduction & Linux Command Review	Build your own cluster	Container Introduction
13:00-14:30	MPI Basic Knowledge	Compile OpenMPI by yourself	Computer Architecture & Parallel Systems
14:40-16:00	MPI Basic coding	Executables & Build System module & spack	
16:10-17:40	Multi node Scheduler		Homework release

Summer Camp

- Collaboration with NCHC (National Center for High-Performance Computing)
- Over 80 students participated



Team Composition

Team Member	Year	Working Items
牟展佑	4	System Design, Benchmarking, Mystery Application, Poster Design
張富強	4	Reproducibility Challenge (DaCe), MLPerf
丁緒慈	4	PHASTA, Paper Survey
郭品毅	3	PHASTA, Mystery Application, System Setup, Benchmarking, Monitoring Tools
吳邦寧	3	LAMMPS, IO500, Lightning Talk
黃恩明	3	Reproducibility Challenge (DaCe), Presentation