Sun Yat-sen University Team
SC22 Student Cluster Competition

Han Huang, Siran Liu, Zhe Tang, Yang Ye, Tianxing Yang, Tengyang Zheng, Dan Huang and Zhiguang Chen

School of Computer Science and Engineering, Sun Yat-sen University

Our Cloud Approach

Our cloud approach is based on Azure. The following are our hardware and software configurations and settings

Compute Bound Configuration

- Intel tool chain: HC-series
  - HC44rs has 44 Intel® Xeon® Platinum 8168 processors
  - features 100 Gb/sec EDR InfiniBand with support for standard Mellanox OFED drivers and all MPI types and versions
  - the best support Intel's rich ecosystem of software tools
- AMD tool chain: HB-series
  - HB120rsv3 has 120 AMD EPYC™ 7003-series processors
  - features 200 Gb/sec EDR InfiniBand with support for standard Mellanox OFED drivers and all MPI types and versions
  - enhance application performance, scalability, and consistency

Storage Optimized Configuration

- Lsv3-series
  - up to 19.2TB (10x1.92TB) NVMe SSD device for L80sv3
  - up to 3.8M Read IOPS and 20GBps throughput
  - the best choice for I/O intensive applications

GPU Enhanced Configuration

- N-series
  - eight 80 GB NVIDIA Ampere A100 GPUs for ND96amsr A100 v4 with NVLink
  - performance Monster and Budget Drainer
  - the best choice for compute-intensive GPU-accelerated applications

FPGA Task Configuration

- NP-series
  - powered by Xilinx U250 FPGAs for accelerating workloads
  - 40 Intel Xeon 8171M (Skylake) CPUs for NP40s

Diversity and Collaboration

Diversity

- Cultural diversty
  - members speak different languages in our team, with different cultural backgrounds
- Location diversity
  - members from developed provinces and undeveloped provinces
- Approach under represented communities
  - special recruiting program for women
  - weekly lecture for under represented communities

Collaboration

- **Han Huang: Computer Science**
  team leader, Heterogeneous computing
- **Siran Liu:Computer Science**
  machine learning,HPC-AI
- **Zhe Tang:Computer Science**
  computer system and architecture
- **Yang Ye:Computing Science**
  program tuning,cluster management
- **Tengyang Zheng:Computing Science**
  hardware tuning,parallel computing
- **Tianxing Yang:Computing Science**
  algorithm,linpack

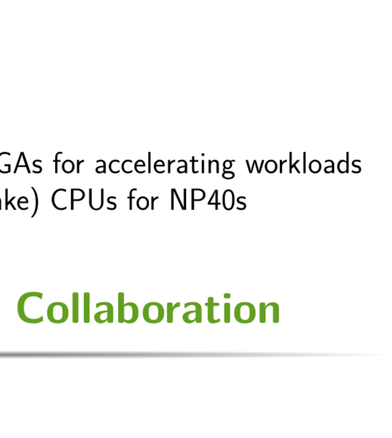

Our Optimizations

General Methodology

- Operating system: a self-optimized Debian 11
- Profiling: find out the hotspots of an application
  - ARM® Allinea MAP, Intel® VTuneTM Amplifier, NVIDIA® Profiling Tools
- Code optimizations
  - parallel computation
  - cache optimization, communication optimization
- Compiling
  - with different compilers and tune options
- Tuning
  - NUMA binding
  - balancing between OpenMP threads and MPI processes

LINPACK&HPCG

- Choose the best SKU from different GPU cards
- Seek for the appropriate amount of GPU nodes
- Manage parameters to get better performance

IO500

- Deploy a file system called kapok
  - developed by National Supercomputer Center in GuangZhou
- Select nodes with NVMe storage to get better performance
- Tune paramaters to make results better
  - file size, file number and so on

PHASTA

- Target
  - compressible and incompressible Navier Stokes equations
- Main functionality
  - perform parallel hierarchic adaptive stabilized transient analysis
- Optimization strategy
  - get ready with leslib and different MPI implementations

LAMMPS

- Backgroud
  - focus on materials modeling
  - a parallel particle simulator at the atomic, meso, or continuum scale
- Tuned performance
  - CPU version:prepare different MPI implementations and MKL
  - GPU version:prepare GPU drivers and Kokkos

Reproducibility Challeng(DaCe)

- deeply understand the research paper and analysing the source code
- prepare different configuration scripts and test the performance

Mystery Application

- Integrated software stack:use Spack
- Resource:balance between the cost and the score
- Analysing:dependent libraries, memory usage and communication

Our Strategies

Dynamic Power Control

- Monitor realtime CPU & GPU workload, throughput and power comsumption via Grafana

- Using RAPL to provide power limiting and power meter capability

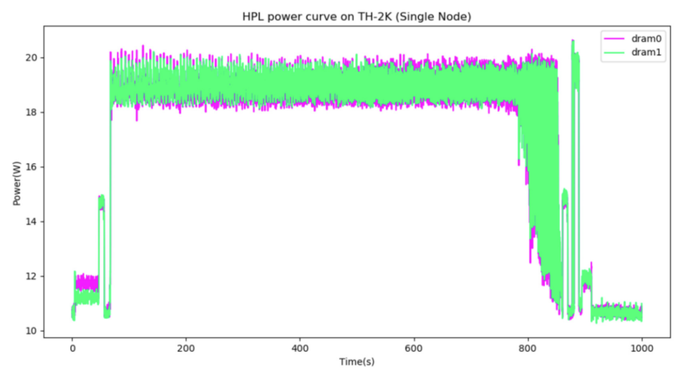

- Using preset script to restrain power under different limits
- An automated alarming program to avoid exceeding power limit

Support

Our team send our gratitude and appreciation to our advisors and sponsors.

Our advisors

- **Dan Huang**
  - research interests: scientific data management, in-memory computing, parallel programming model, and distributed storage systems
- **Zhiguang Chen**
  - research interests: big data storage and processing, parallel and distributed computing, high performance computing and supercomputers