

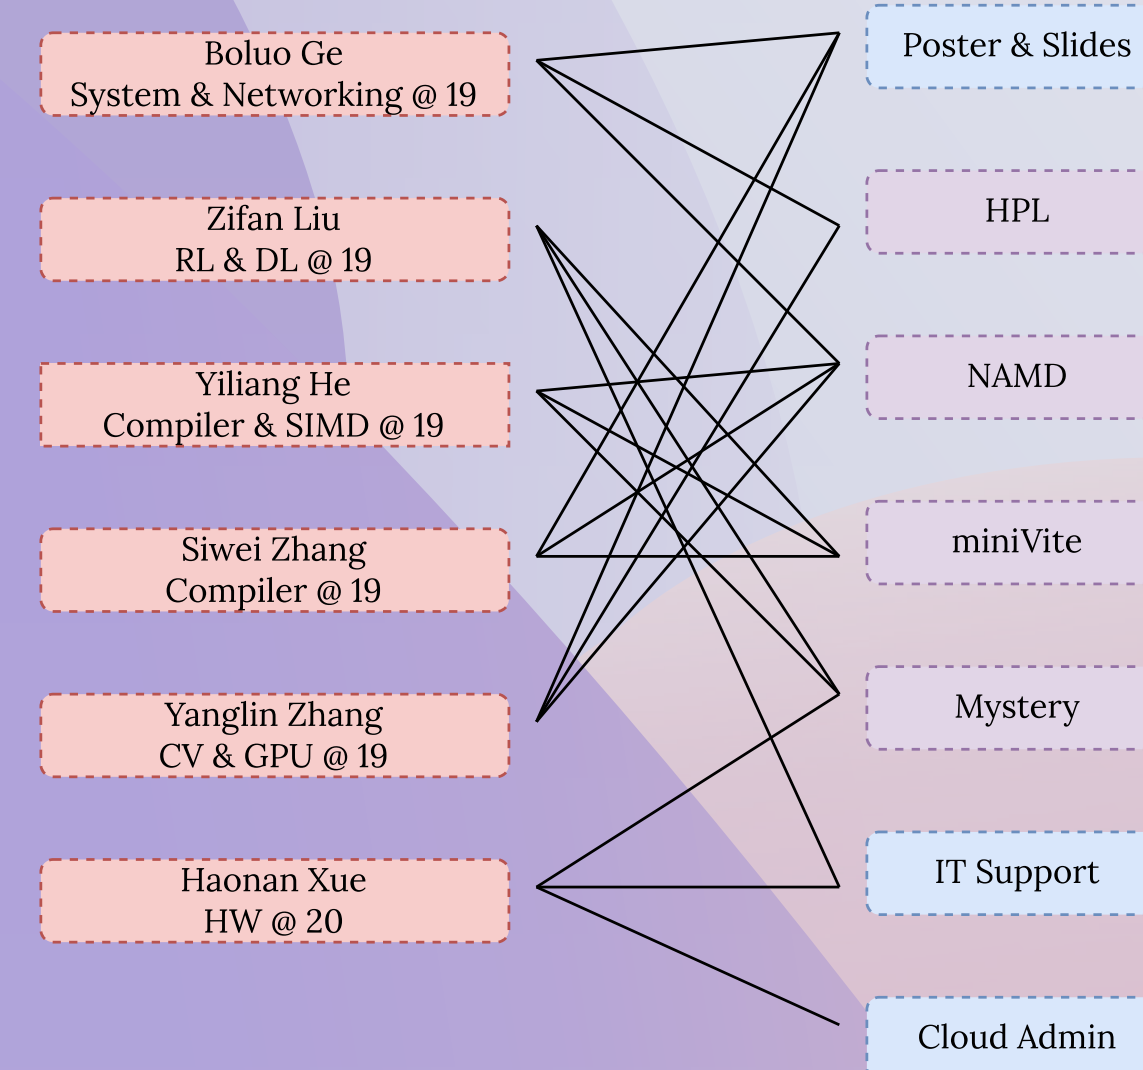
Team Introduction

The Supercomputing Group at CUHK-Shenzhen are sponsored by the Information and Technology Service Office (ITSO) and the School of Data Science (SDS), advised by Dr. Yeh-Ching Chung, the Professor of SDS. In phase I, we have built a mini-cluster of 40 CPU cores, along with 12 dedicated Nvidia GPU cards. We have already performed many benchmarking and testing programs on the mini-cluster, and more CPU and GPU resources are planned in phase II and phase III. Till 2022, we anticipate 400 CPU cores and 20 GPU cards be installed in our cluster. They are all open for our team to practice skills without any limitations. And we also open a portion of computer resources for students interested in our team to do a simple experience.

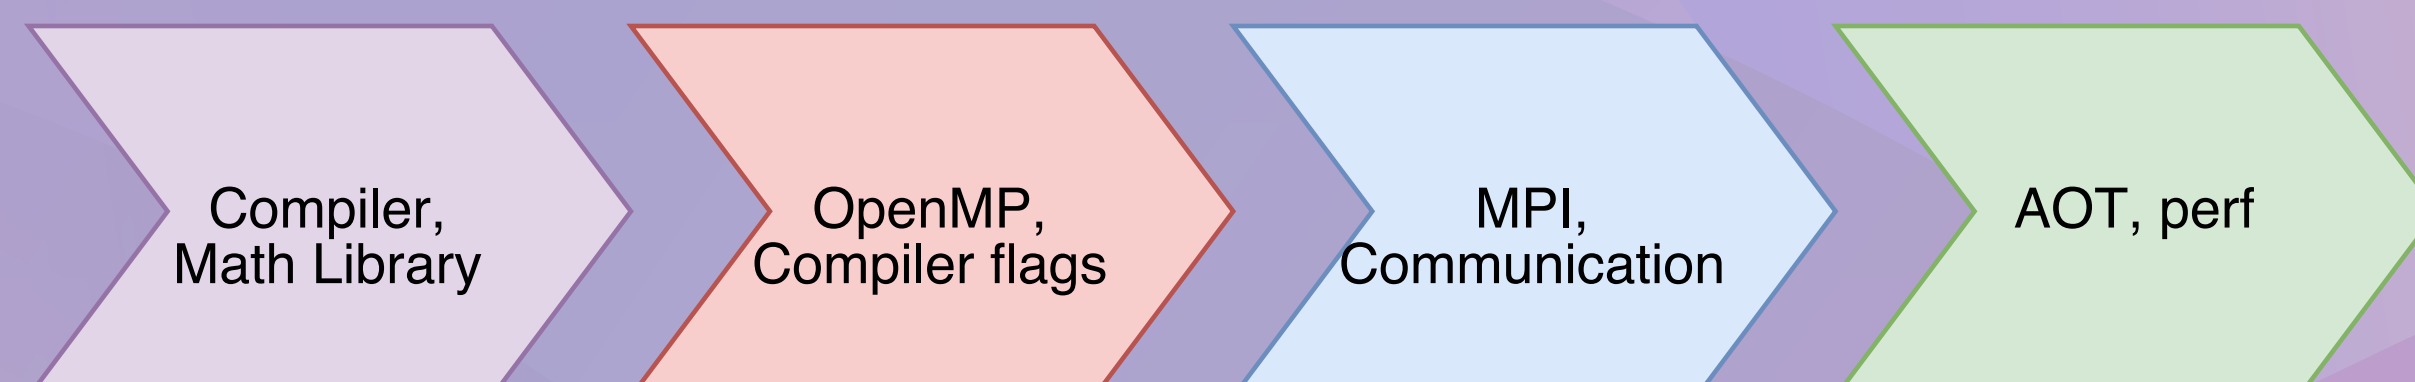
Team Diversity

- **Culture diversity.** Our university is committed to providing top-quality higher education that features an integration of the East and the West and fostering an enriching research environment. We encourage people from different cultural backgrounds to interact and participate in our activities.
- **Gender diversity.** We provide opportunities for both males and females to show their scientific talents and encourage them to cultivate an interest in high-performance computing. At the same time, our professors will also provide some guidance on their career development.
- **Birthplaces diversity.** Our team members come from different provinces covering eastern and western China.

Collaboration



Optimization Approaches



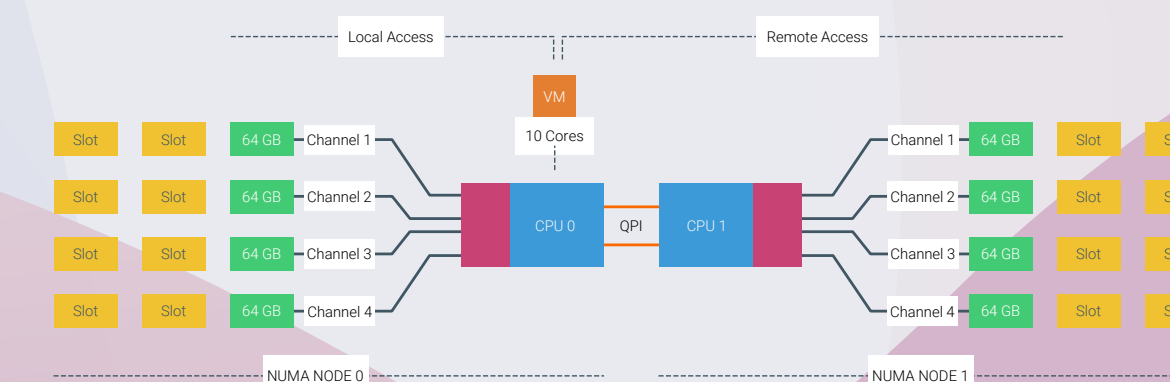
LINPACK Benchmark

- Parameter optimization for theoretical upper limit
- Tune parameters to achieve best and balanced CPU/GPU computing
- Conduct both single node test and cluster test for Linpack benchmark.

NAMD

- Based on the Charm++ parallel programming model
- Great parallel efficiency for molecular dynamics simulation
- Enable interactive simulation for parallel molecular dynamics code

Hardware



- Haswell-EP models with ten and more cores support cluster on die (COD) operation mode, allowing CPU's multiple columns of cores and last level cache (LLC) slices to be logically divided into what is presented as two non-uniform memory access (NUMA) CPUs to the operating system
- Use Turbo Boost and SpeedStep for some instantaneous performance needs. Thermal Monitoring can help us catch. Make good use of NUMA architecture to optimize IO latency. Demand Based Switching also helps us do the power management

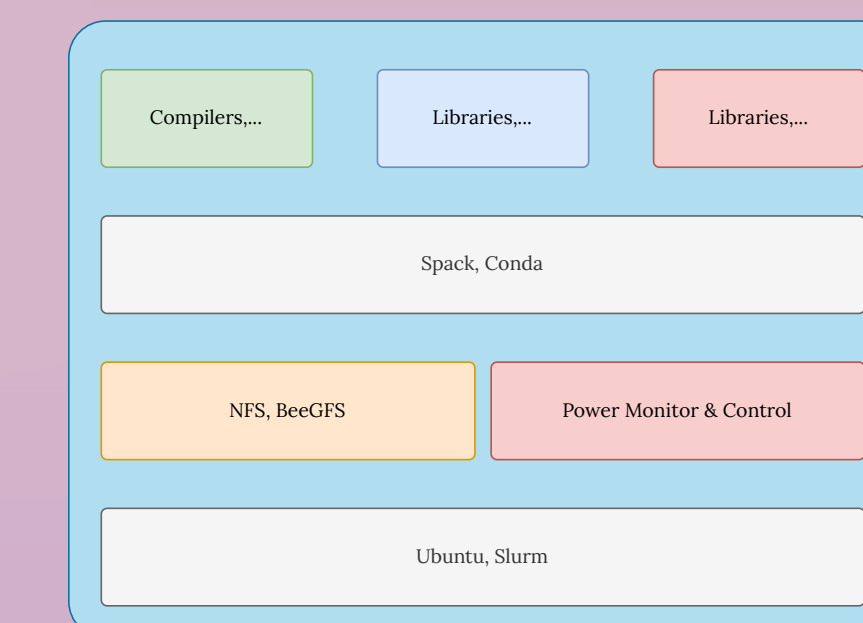
miniVite

- A representative graph analytics benchmark tool for distributed-memory systems
- Strong ability to partition a graph into clusters that are dense within clusters and sparse between clusters
- Represents highly complex computational patterns stressing a variety of system features

Mystery

- Develop with this mystery application quickly
- Compute with this mystery application sufficiently
- Run this mystery application in all mainstream architectures

Software



- Ubuntu 20.04 with tuned kernel across all chassis installed by cloud-init, different hardware is auto-configured on instantiation
- Spack for flexible deployment on different OS and Arch
- NFS on Cheap Chassis to share installed package, BeeGFS on RDMA clusters to share huge datasets
- Slurm, for job schedule