



Team Minuscule

Shanghai Jiao Tong University

Shanghai, China

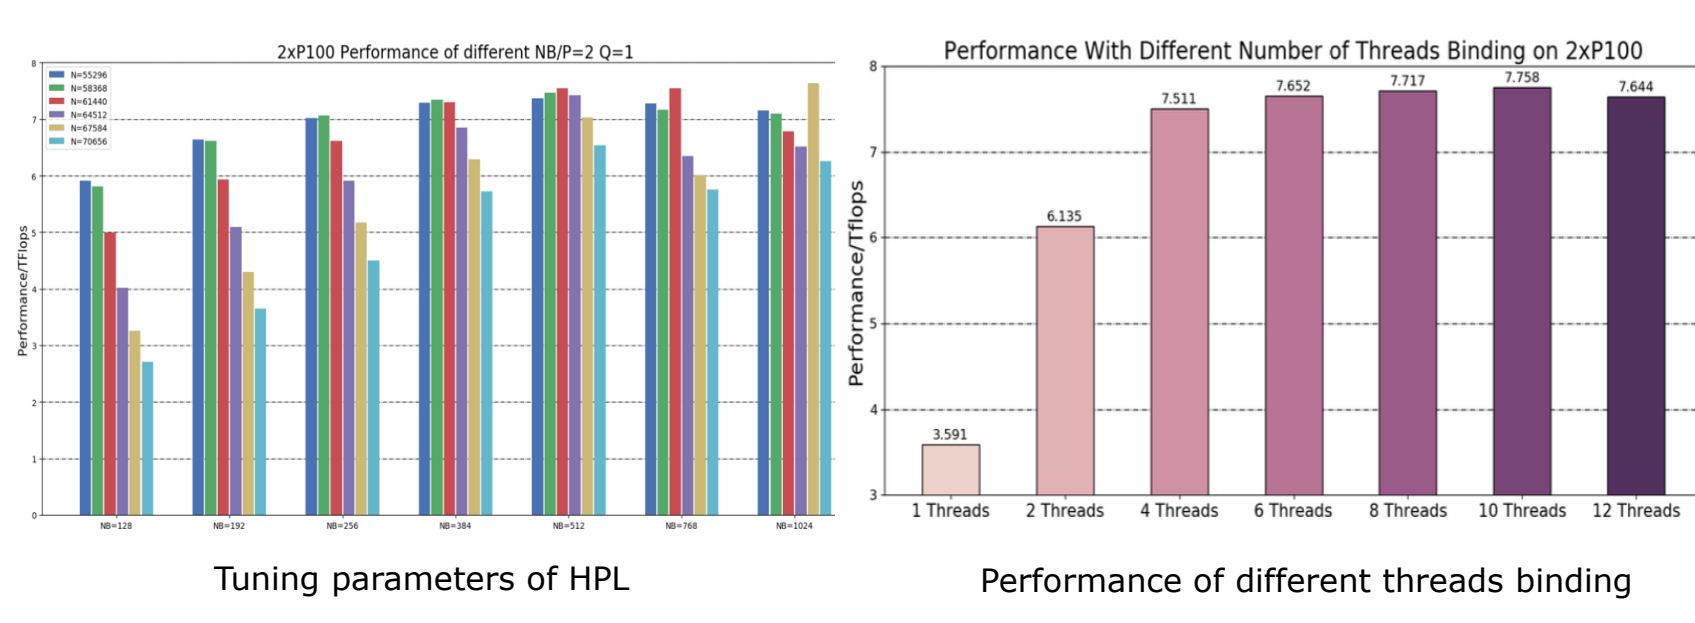


HPL, HPCG & IO-500

Benchmarks

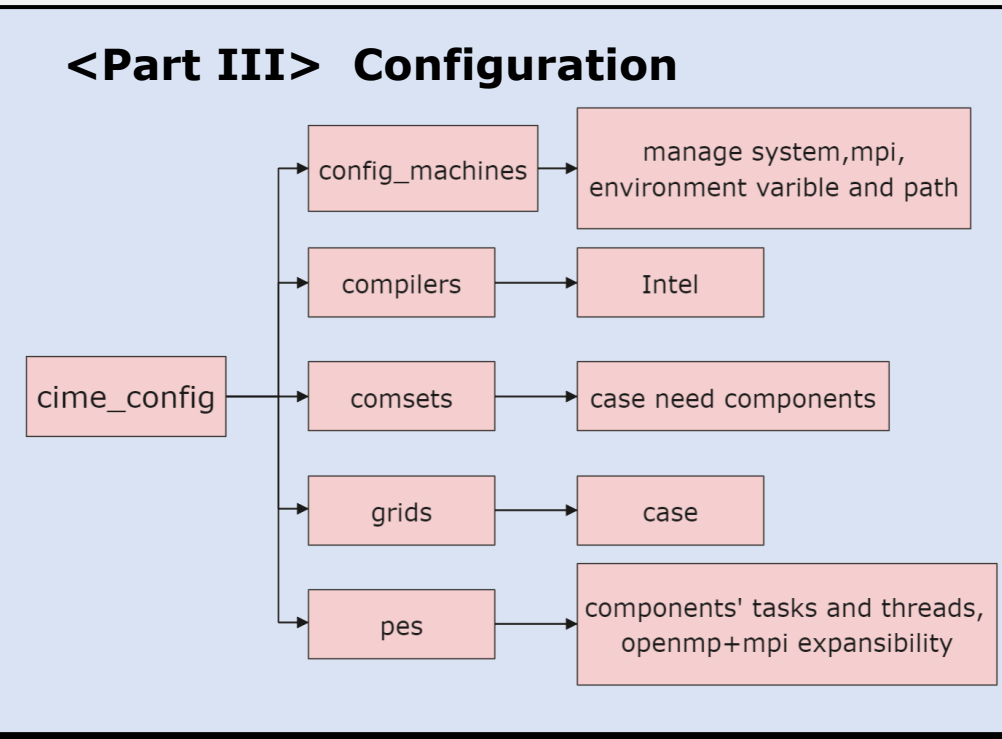
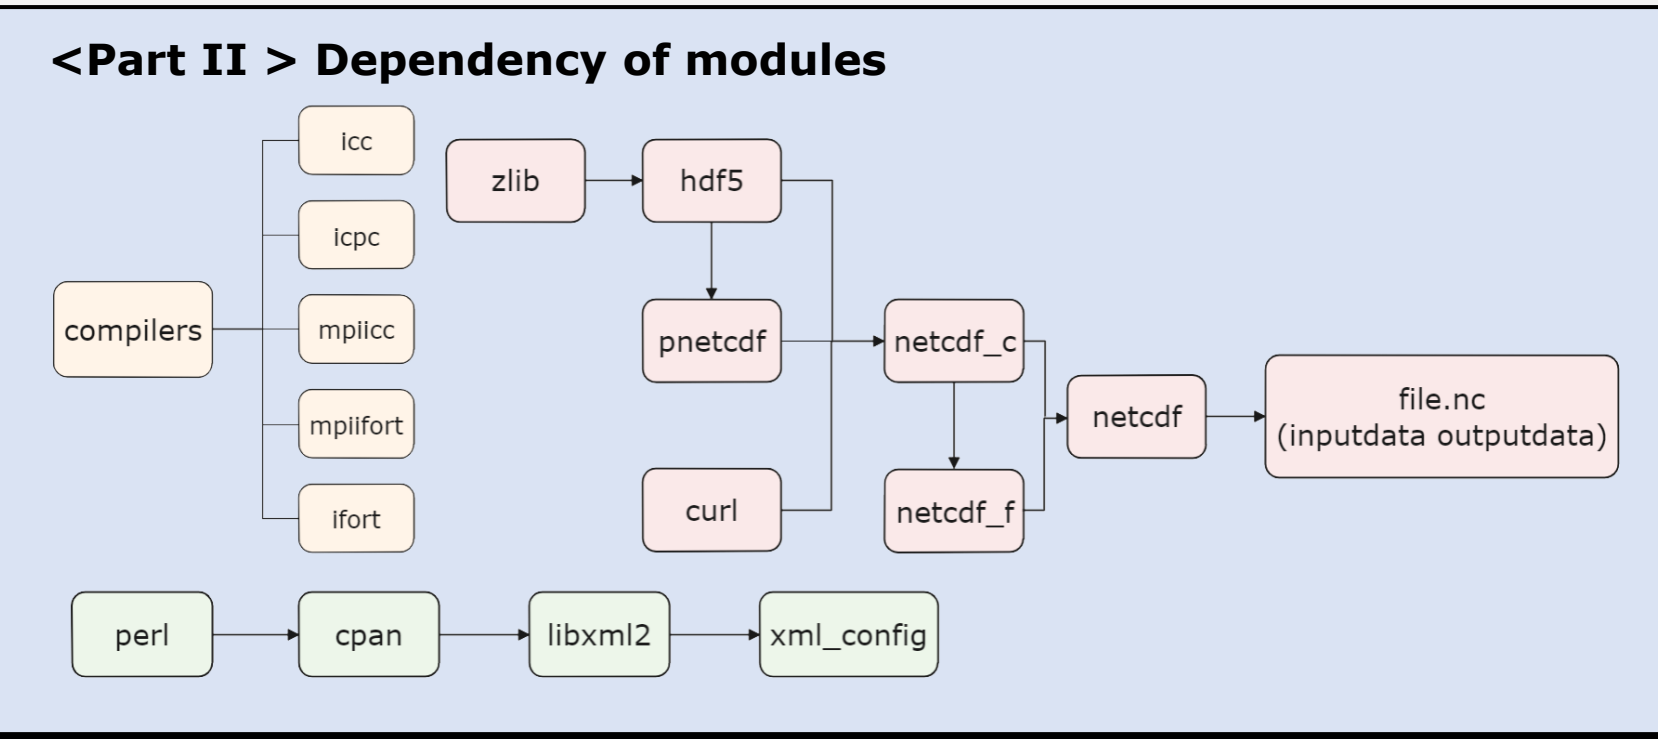
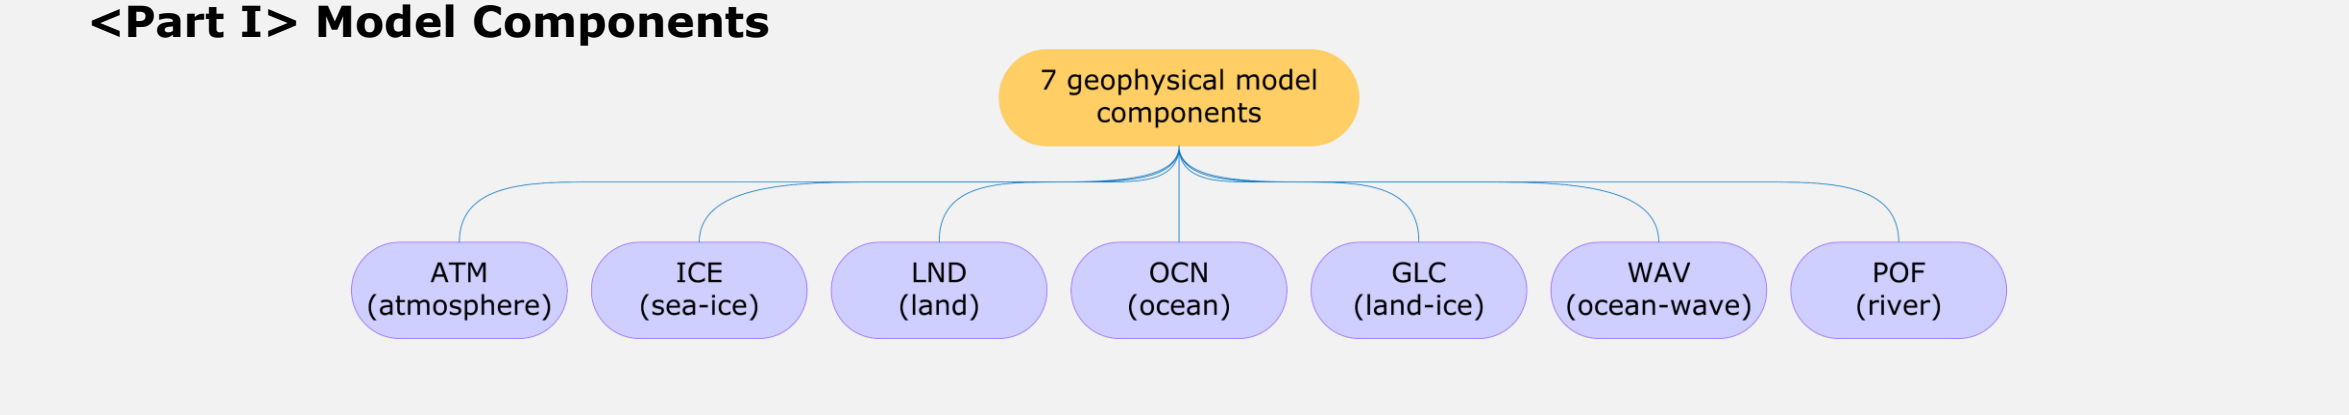
We will run **HPL**, **HPCG**, these linear system solving algorithms, mainly on GPU nodes and **IO-500**, the storage benchmark on memory-optimized SKU(s) of Azure CycleCloud. Before the availability of Azure, we firstly organized our tests on NVIDIA Tesla P100s.

For instance, in HPL, we executed runs with different blocking parameter combinations(setting in the .dat file) on the fixed 2 x P100 node, to gain peak performance. And we also tried optimizing the binding of threads to each process, to maximize the performance on given GPU(s) in the benchmark test.



CESM

Community Earth System Model
a fully-coupled, global climate model that provides state-of-the-art computer simulations of the Earth's past, present, and future climate states.



<Part IV> Running commands

Here are the most basic commands when we try to construct a new case.

```

$. /create_newcase --
case mycase --compset
X --res f19_g16
./case.setup
./case.build
./case.submit

```

<Part V> Evaluation & Optimization

Conclusion:
1. Number of mpi tasks has an effect on both CPL run time and CESM components' run time, while number of OpenMP threads only influences CESM components' run time.
2. PE Configurations affect cesm run time in that it changes the time balance between the components.

Brief Introduction to Our Team

Shanghai Jiao Tong University (SJTU) is a major research university in Shanghai, renowned for many subjects such as ME, CS, Materials Science and so on. It is established in 1896, as one of the oldest universities in China. SJTU has always been actively participating in international cooperation and educational activities.

Our team is supported by the Center for High-Performance Computing (CHPC) of SJTU. Established in 2011, CHPC is delivering a wide range of HPC services across the university. The CHPC's supercomputer is named π -2.0, clustering thousands of CPU nodes with 26,000 cores and an AI HPC cluster organized with 8 DGX-2 AI systems. (<https://www.top500.org/system/179848/>)

Shanghai Jiao Tong University has assembled a team of six undergraduates with a strong diversity of academic backgrounds as follows. And our teammates are advised by three **HPC Engineers** of the CHPC:

- **Key Liao, Minhua Wen, Shenggan Cheng**

Jiarui Lu (Senior) - Chemistry & Mathematics
• Senior
• Reproducibility Task and system manager
• Undergraduate researcher in machine learning on drug discovery

Changrong Li (Senior) - Electronics Engineering
• Senior
• GROMACS Task
• Undergraduate researcher in optimal network

Xuetao Chen (Junior) - Electrical Engineering
• Junior
• GROMACS Task and Benchmark
• CESM Task and Benchmark

Chenhui Wang (Junior) - Electronics Engineering
• Junior
• CESM Task
• Undergraduate researcher in deep learning of images and human action

Zhicheng Fang (Junior) - Computer Science
• Junior
• GROMACS Task
• Undergraduate researcher in deep learning causal inference

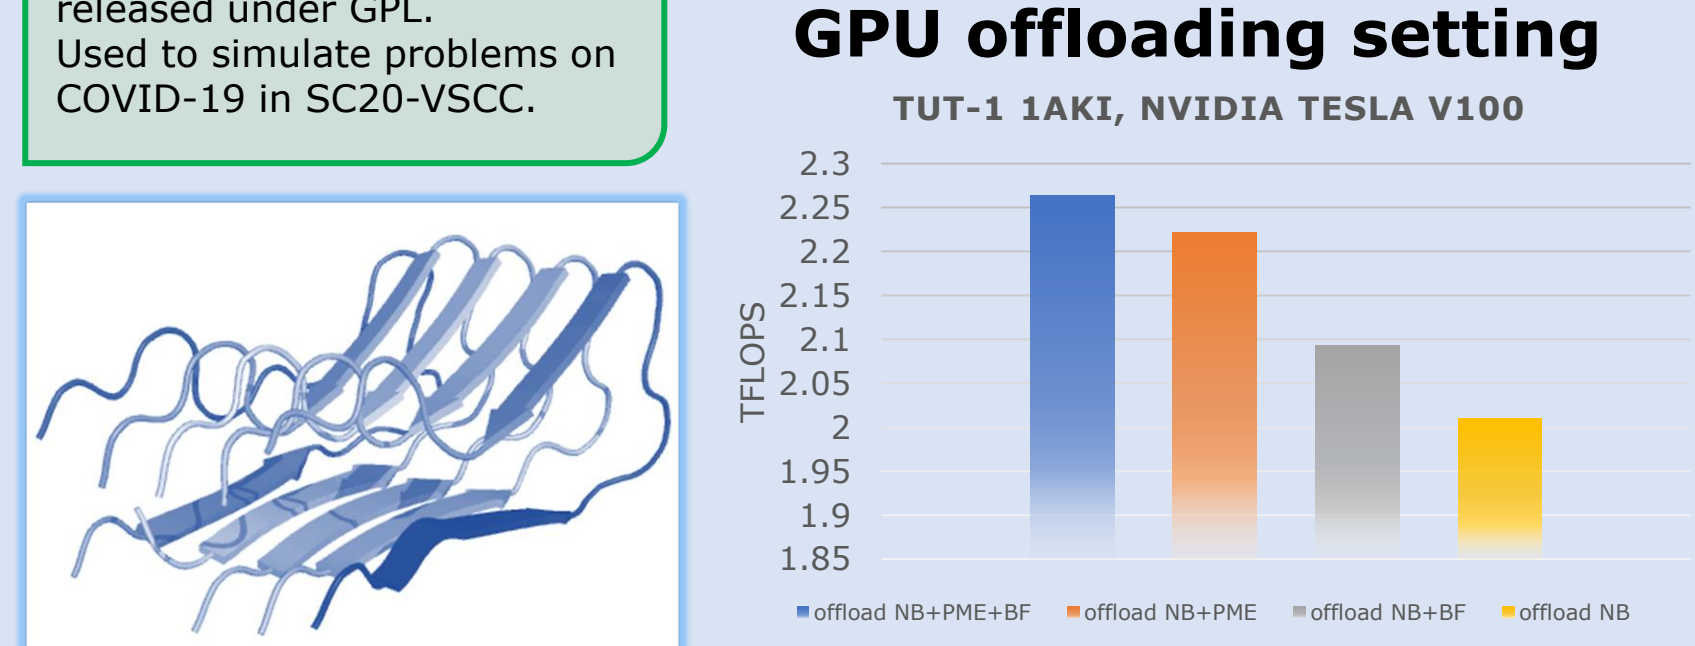
Chen Liang (Junior) - Software Engineering
• Junior
• CESM Task and Benchmark

GROMACS

a versatile package to perform molecular dynamics, designed for biochemical molecules like proteins, lipids and nucleic acids, extremely fast at calculating the nonbonded interactions. Open source released under GPL. Used to simulate problems on COVID-19 in SC20-VSCC.

Case: Tut-1 (Lysozyme 1AKI) Performance

Final mdrun process setting	Time/s
ALL on CPU, 8 tMPI, 5 OMP	655
NB+PME+BF on GPUx1	491



Azure Configuration

This year, we are free to select our Azure cluster configuration. According to the characteristics of each application, we choose SKU(s) for cluster as follows:

VM & SKU(s) Configuration

SKU of CPU Node	\$ / Hour	Memory/GiB	Configuration
E20d_v4	\$ 1.87	160.00	20 * Intel Xeon Platinum 8272CL
E48_v4	\$ 3.63	384.00	48 * Intel Xeon Platinum 8272CL
HB60rs	\$ 2.51	223.52	60 * AMD EPYC 7551
HC44rs	\$ 3.48	327.83	44 * Intel Xeon Platinum 8168
SKU of GPU Node	\$ / Hour	Memory/GiB	Configuration
NC24r_Promo	\$ 2.09	224.00	4 * K80 (with RDMA InfiniBand)

E series are chosen due to its optimization for heavy in memory applications, such as MemXCT. It offers VM(s) configured with high memory-to-core ratios, which is suitable for such computations;

HC series are selected because it is optimized for HPC applications driven by memory bandwidth, such as fluid dynamics, explicit element analysis and weather modeling. Such VMs are fit for the applications like CESM. HB series, similarly, caters to the HPC applications driven by intensive computation, it can assume some computational tasks with its high clock speed.

*GPU nodes are selected temporarily since we only have access to at most K80, with which we are also familiar. If Nvidia V100 is available such as NCsv3 SKU, we prefer that as our GPU node.

System Software & Development Tools

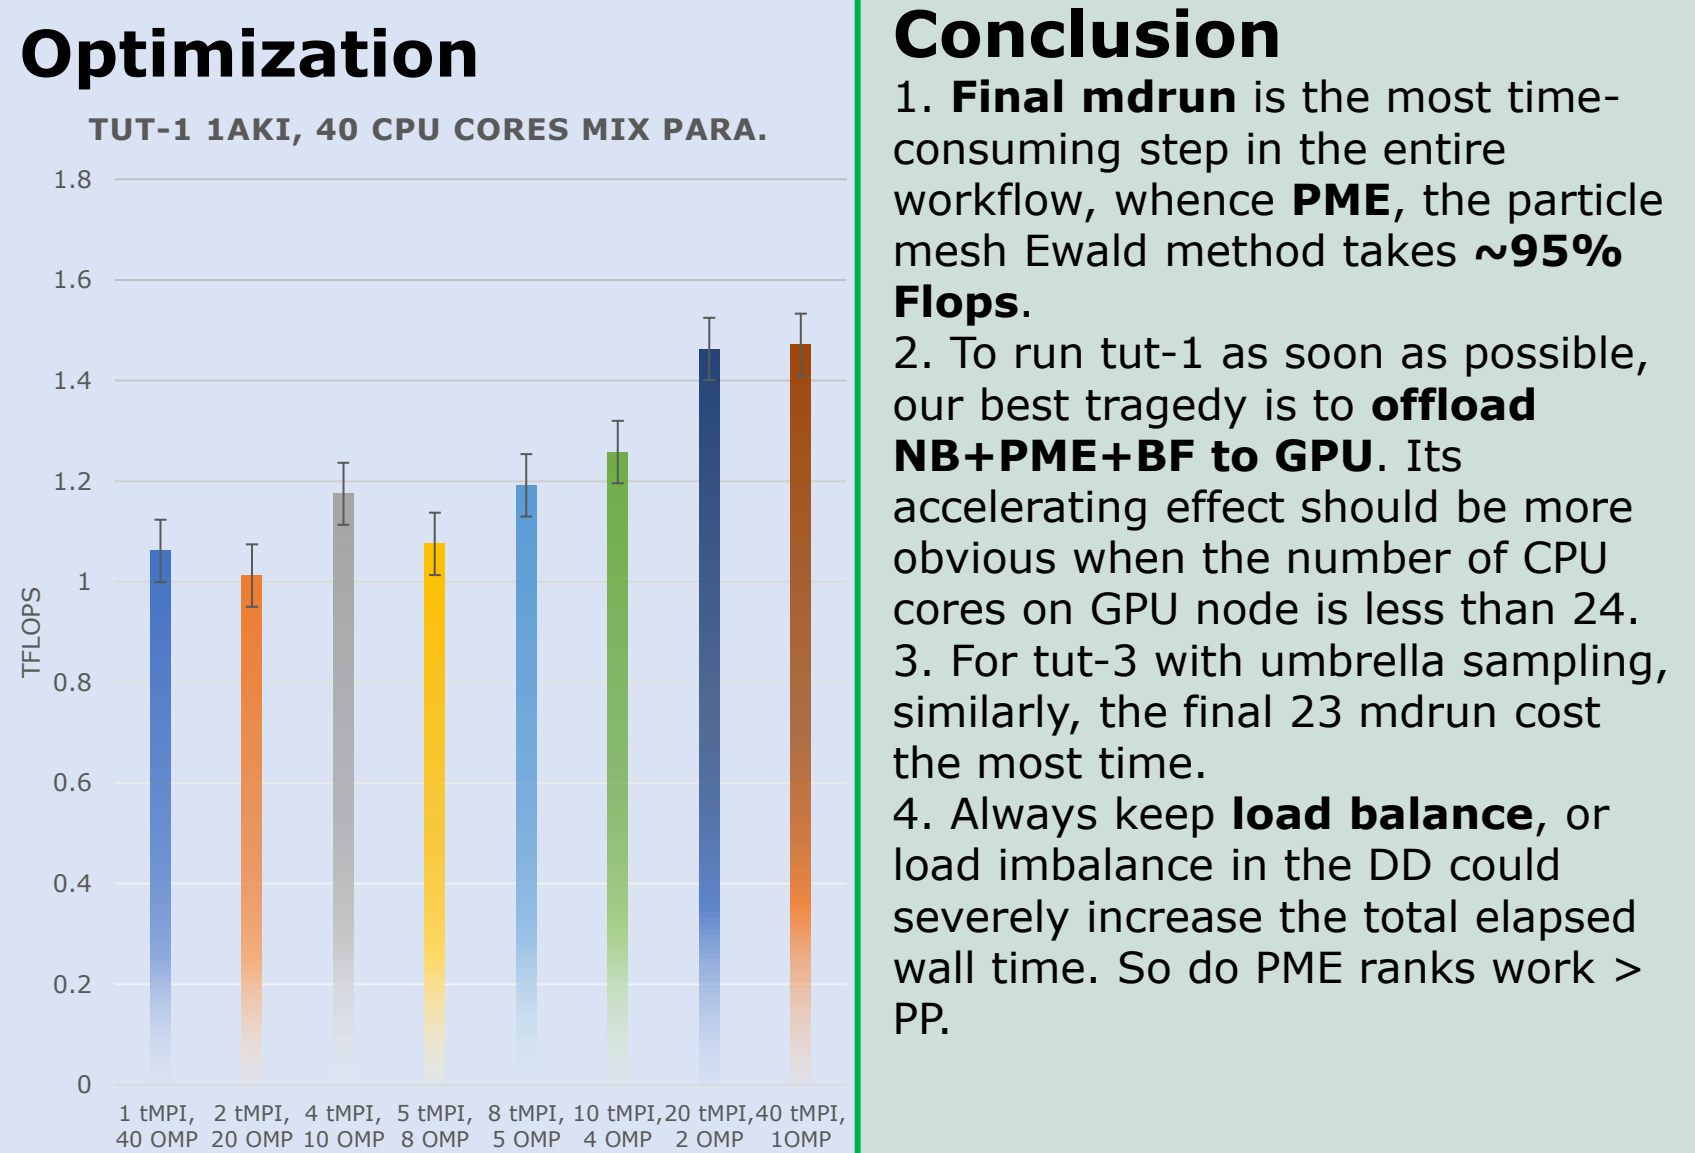
Operating System	Linux CentOS 7, x86_64
Resource Manager	SLURM 19.05.5
Compilers	GNU C/C++ 8.3.0/4.8.5 & Intel Parallel Studio 2019.0.5
Nvidia API	CUDA 10.1/10.2
MPI	Intel MPI Library 2019 Update 5
Optimized Math Library	Intel MKL 2019 Update 5
Profilers	Intel Vtune Amplifier 2019 Update 6
Parallel File System	BeeGFS
Package Manager(Optional)	Module

Why We Will Win?

- Cohesive As a Whole**
We have made it a routine to communicate with each other and hold seminar weekly to share HPC experience among the team.
- Strong Support and Guidance**
We have access to the SJTU π -2.0 supercomputing cluster for our team training. And, Key Liao, as our main advisor who provides mature competition technical guidance, was former team leader and two-time podium team member in ASC16 and ISC16.
- Working with Engineers of HPC**
We have opportunity to be exposed to the development and optimization of some scientific applications in our University. This is a fast lane to mastering the HPC.
- Unquenchable Passion for High-performance Computing**
And the SC20-VSCC competition!

M-flops of each part

Umbrella sampling



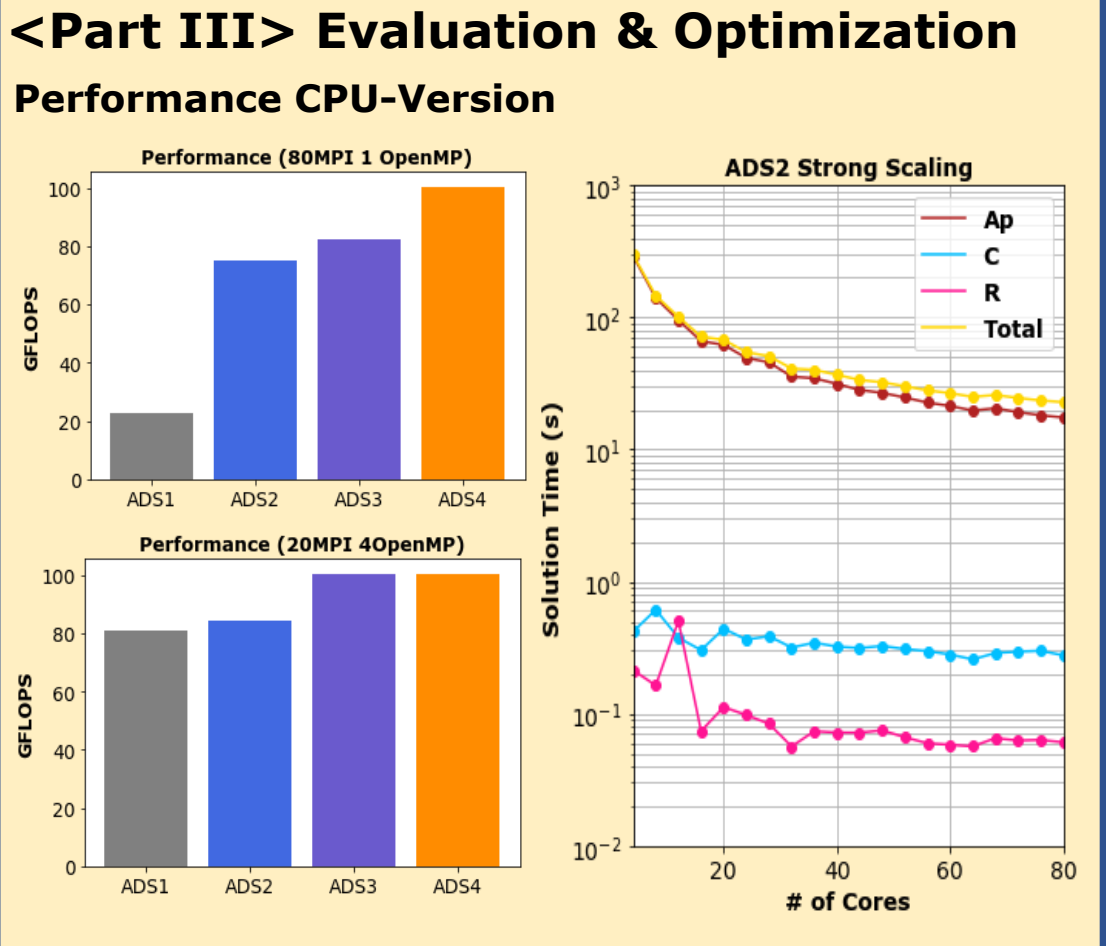
MemXCT - Task

- **Compiler:** mpiicc & nvcc
- **OPT Flags:** -O3 -qopenmp -xhost

<Part I> Procedure & Visualize

Tomogram Visualize (24 Iters)

Results will be verified by Fiji. And our testing results agreed well.



<Part II> Profiling

RunTime Analysis(On ADS2)

Containing Ops:
- Ray-tracing
- Transposition
- Project blocking
- Back blocking
- Fill project matrix
- Fill back matrix

Containing Ops:
- Projection
- Backprojection

Optimized MemXCT Kernel:
By Mert et al.

```

#pragma omp for schedule(dynamic)
for(int block = 0; block < numblocks; block++){
  for(int n = 0; n < blocksize; n++){
    output[n] = 0;
    for(int buff = blockdisp[block]; buff < blockdisp[block+1]; buff++){
      for(int m = 0; m < buffmapnz[buff]; m++){
        input[m] = obj[buffmap[buffmapdisp[buff]+n]];
        start = buff*blocksize;
        for(int n = 0; n < blocksize; m++){
          for(long n = buffdisp[start+m]; n < buffdisp[start+m+1]; n++){
            output[m] += input[(index[n]]*value[n];
          }
        }
      }
    }
  }
}

```

(*Clip from kernels.cpp)

