



國立清華大學

NATIONAL TSING HUA UNIVERSITY




Introduction to NTHU

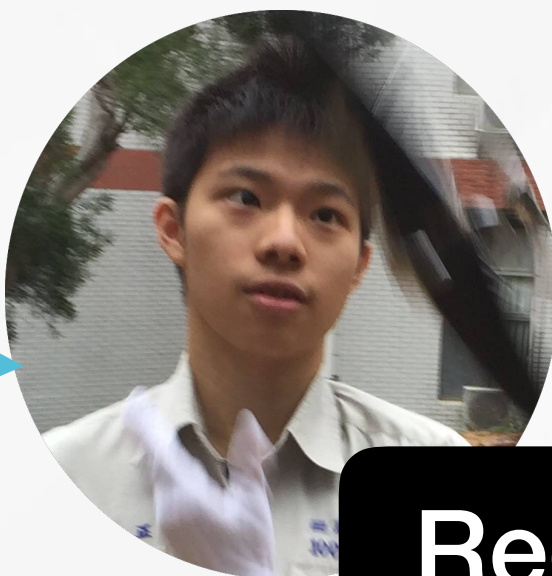
Our University

Located in Hsinchu, Taiwan, National Tsing Hua University is a comprehensive research university offering a full range of degree program in science, technology, engineering, humanities, social science, and management. It consistently ranks as one of the premier universities in East Asia.


Our Team




afe
Responsible for Born and system setup. Team leader and system administrator. Familiar with Linux and excellent in Python, C++ and Go. Has experience with implementing profilers. Participated in SCC16.




Reese
Responsible for MrBayes and Born. Has abundant experience in vectorization and optimization on several computing architecture. Participated in ASC17 and ported MASNUM on Sunway TaihuLight.




Scott
Responsible for MrBayes and cloud platform. Double major in mathematics and computer science. Familiar with C/C++, Android Java and Web Development. Interested in explaining numerical equations from applications with a high level view of advanced mathematics skills to team members.



lupiter
Responsible for reproducibility challenge. Internship in TCS (TA TA Consultancy Service), Bank-Core System project. Not only familiar with COBOL and utilizing database, but also keep acquiring new knowledge of parallel computing.



ron
Responsible for reproducibility challenge. Participated in open-source project NTHUOJ. Care about code quality and communication between team members. Familiar with C and Python. Embrace creative thoughts, enjoy working with smart people.

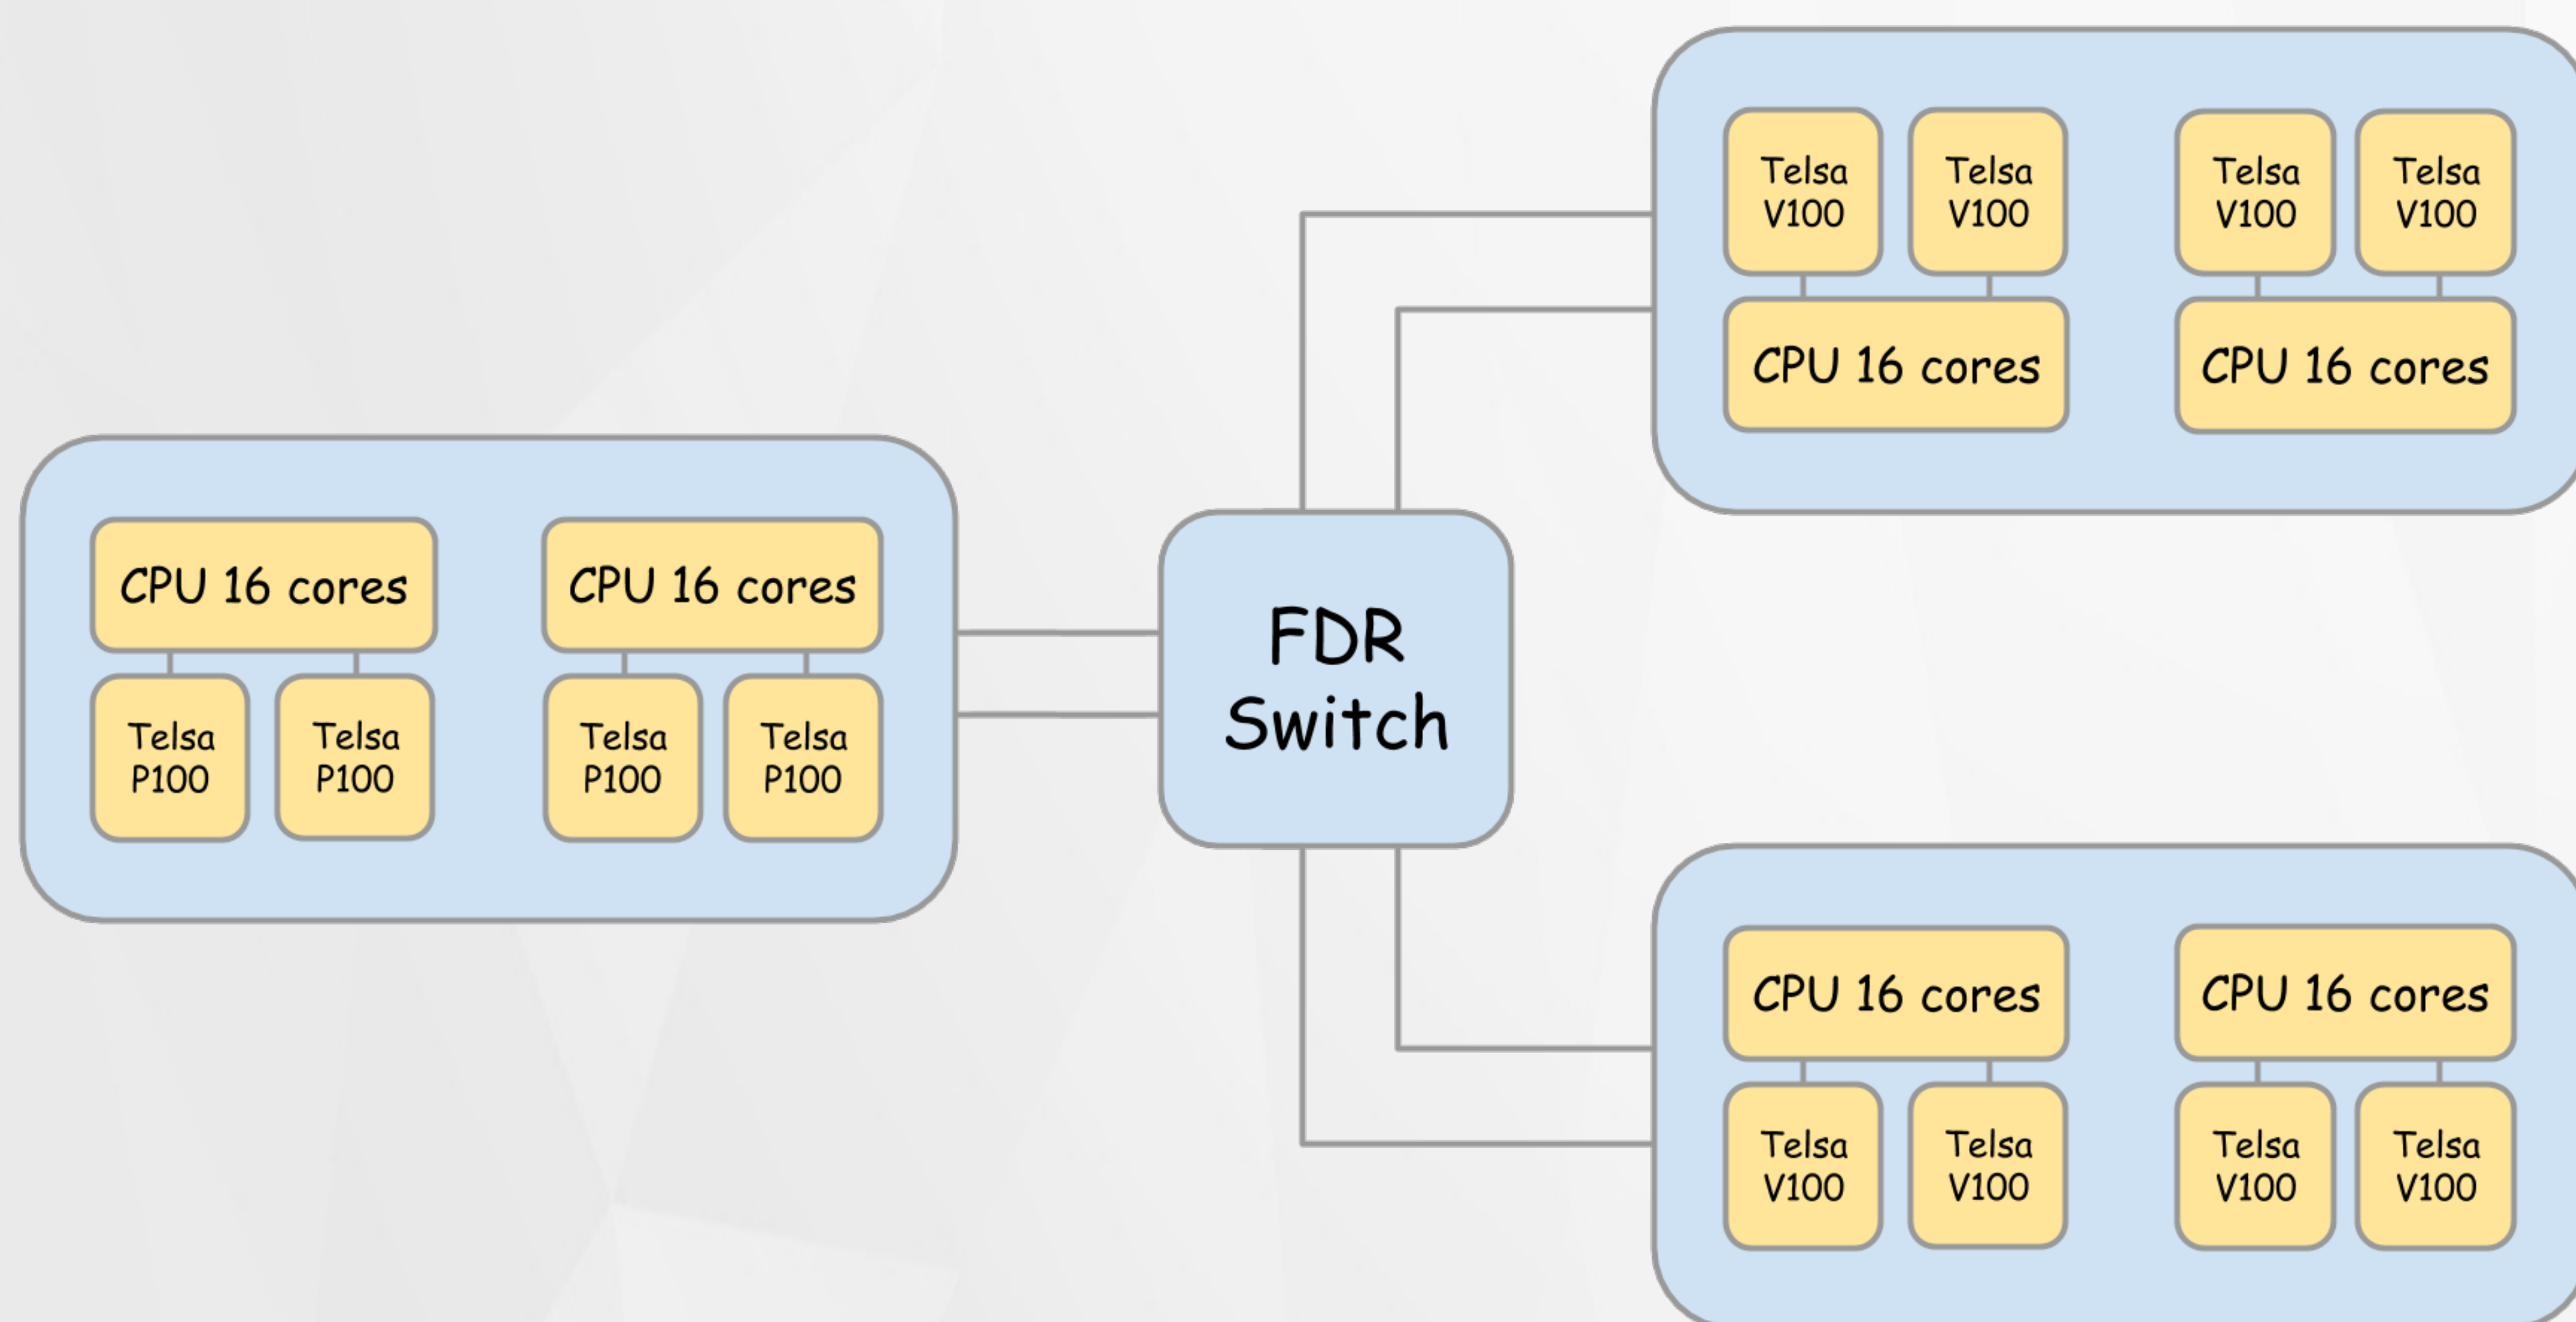


Katrina
Responsible for MrBayes. Internship in ITRI (Industrial Technology Research Institute). Familiar with parallel architecture and C. Also the director of student association. In this competition I will put my heart and soul into championship.

Hardware Configuration

Our server is sponsored by QCT, which is very energy efficient and ideal for technical computing and mission-critical applications. With its enhanced thermal and power design, we are confident that it is suitable for the competition.

We choose Intel Xeon Gold 6130 Processor, since it has both great single-core and multi-core performance, and its low power consumption makes it ideal for this competition. We also use V100 and P100, since they give excellent performance in both single and double precision floating point operation.



Server	S5BV 4 nodes
CPU	Intel Xeon Gold 6130 Processor 2.1GHz
Memory	384GB DDR4 2400GHz
Accelerator	NVIDIA Tesla V100, NVIDIA Tesla P100
Disk	Intel SSD 800GB (SATA)
Interconnect	Mellanox Dual Port FDR with Switch

Software Configuration

We use CentOS as the operating system because of its stability and there are plenty of resources on the Internet. After several experiments of different MPI libraries with different applications, we decide to choose Intel MPI as our main MPI solution.

OS	CentOS 7.4
Compiler	GCC 4.9.4 Intel Parallel Studio XE 2018/2017/2016
GPU software	CUDA 8.0 / 9.0, driver 384.81
MPI	Intel MPI from psxe OpenMPI 3.0.0 / 1.10.2 MVAPICH2 2.2

Preparation for SCC17

We met our professors and coaches on weekly basis to learn the basics, and discussed on the problems we encountered during preparation for SCC17 chosen application. Several additional HPC applications are selected to be installed and optimized every few weeks, in order to get familiar with the cluster environment. Abundant preparation and frequent meeting boosted our confidence. We are ready for the fight!

Strategy & Optimization

MrBayes

MrBayes uses Markov Chain Monte Carlo (MCMC) to inference the evolutionary and phylogenetic model, while BEAGLE library accelerates the core calculation of Bayesian inference with Intel SSE and OpenMP. For better performance, we extend the core calculation with Intel AVX instructions. To achieve better cache accessing, we change the pattern of matrix multiplication from linear combination to computing the dot product of the rows and columns.

Born

We found that the image kernels are suitable for GPUs, so we ported them to CUDA and verified its correctness. Our experiments show that our GPU version of Born is significantly faster than the CPU code.

LAMMPS

The same experiments stated in the paper are tested on both our machine and Cloud platform. Due to different installation path, customized build scripts for different platforms are prepared, so that we can set up the environment faster and avoid typos. Our strategy is preparing scripts of reproducing and processing the data needed for the figures, to minimize the time parsing from the raw data manually. To parallelize our work time, Lammmps will run first, so that Ron and lupiter can start working on the report, while our machine is working on another application.

CycleCloud™

Cyclecloud™ platform provides an opportunity for extra computing power and avoids data loss in the power shutoff activity. Once the "West US 2" region is accessible, MrBayes can take advantage of a Intel Xeon Platinum cluster (Fv2); on the other hand, Born can benefit from a NVIDIA Tesla P100 instance type (NCv2).

Power Consumption

To control the power consumption, we dynamically adjust CPU and GPU clock frequency and the cooling system. We also did a power consumption estimation for each application.

Why We will Win ?

- The passion toward parallel programming supports us during the long preparation.
- We have known each other for a long time, we formed a tight bound between each other!
- We not only work on performance tuning, but also study the background knowledge and the architecture of these applications.
- In spare time, we take courses and learn about the knowledge of supercomputing.
- We just got the championship of Taiwan Student Cluster Competition.
- Our team members, Scott and Afg majoring in Mathematics and Power Mechanical Engineering, can use their knowledge in science domain to help us analyze problems and offer scientific insights.
- By modifying the CPU code of Born, we enabled GPU acceleration, which has a significant improvement.

