

Our university

NTU was inaugurated in 1991, when its predecessor institution, the Nanyang Technological Institute (NTI) merged with the National Institute of Education (NIE). NTU has since grown to become a full-fledged comprehensive and research-intensive university, with over 33,000 undergraduate and postgraduate students. Our university is consistently ranked as one of the top universities in Asia.



Our team

There are six members in our team, Weichao, Siyuan, Yiyang, Daxuan, Hailin, and Ying Hao.

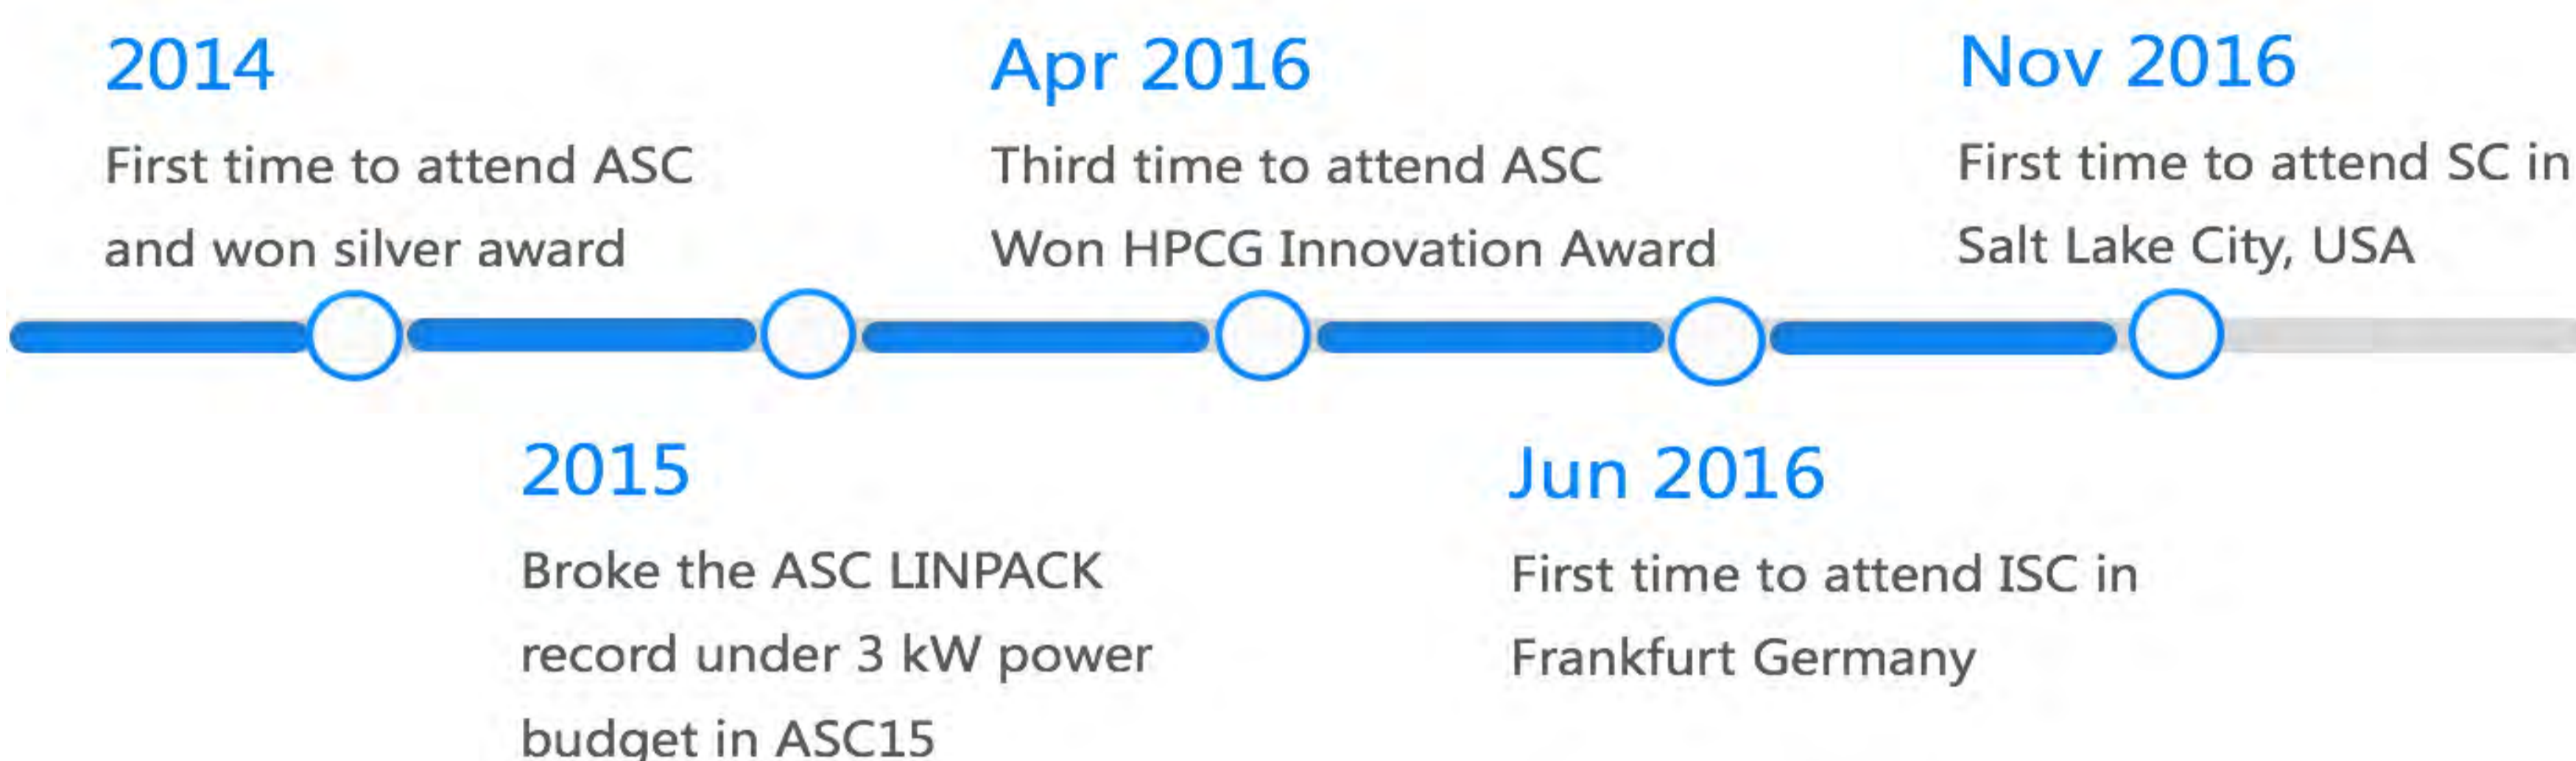
Our team members have a diverse background in computer science, including algorithm design, operating system, artificial intelligence, and high performance computing.

Why our team will win

- We have a very rich experience in high performance computing competitions and feature experts in optimizing for HPL and HPCG.
- We are very well prepared and the way we distribute our tasks have allowed each of us in the team to specialize and optimize for a specific application.
- The performance of our computing cluster which consists of a very well balanced number of CPU cores and graphic cards together with high bandwidth(100G) and low latency interconnect.

Our preparation

- Started preparation for the competition since August
- Two members of the team is assigned to focus on each competition application.
- Each of the application is optimized by the assigned members using a training cluster at NTU Parallel & Distributed Systems Lab
- Regular meetings are conducted twice weekly to discuss progress and solicit optimization ideas from each other
- Regular visit to A*CRC for hands-on training and advice from domain experts




Applications

- **Paraview:** We use Nvidia OpenGL, EGL and Intel MPI to fully utilize CPU and GPU computational resources across all our four unified nodes.
- **ParConnect:** Equipped with 176 Intel CPU cores and Intel optimized compiler, we can achieve best utilization of hardware within power budget. Together with Mellanox EDR InfiniBand Switch, communication overhead can also be minimized.
- **Password recovery:** We choose Hashcat to utilize GPU and CPU computational power using OpenCL runtime. With variety of input parameters and formats, we can tune the application specifically for the competition.
- **HPL:** With our rich experiences and deep understanding in tuning HPL, we can tune HPL specifically for our hardware configuration to achieve the highest result.
- **HPCG:** Different parameters has been experimented, and best combination has been chosen to maximize the performance.


Innovation

- 3D printed wind tunnel as custom cooling solution for NVIDIA Tesla card in workstations



System design

- Good balance between CPU and GPU computation capabilities
- Uniform system configuration
- Infiniband EDR interconnect for maximum bandwidth and minimum latency
- Fits under the 3,120 W power limit for SCC
- SLURM as job scheduler and environment module for ease of management
- Grafana, InfluxDB and Telegraf for cluster statistics monitoring



Item	Spec	Quantity
Servers	Dell Precision Tower 7910 XCTO	4
Ethernet Switch	Netgear Switch	1
Infiniband Switch	Mellanox EDR Switch	1
CPU	Dual Intel Xeon Processor E5-2699 v4 (22C, 2.2GHz, 3.6GHz Turbo, 55MB, 145W)	2 * 4
GPU	NVIDIA Tesla P100	2 * 4
RAM	128 GB DDR4	1 * 4
Storage	128 GB SSD	1 * 4

Item	Spec
Operating System	CentOS 7
Compiler	Intel Compiler & GNU Compiler
MPI	Intel MPI & OpenMPI
Profiler	Alinea
Scheduler	Slurm
Environment Management	Module
File System	NFS
System Monitoring	Grafana

Acknowledgement

Special thanks to National Supercomputing Centre(Singapore), A*STAR Computational Resource Centre, NVIDIA and Mellanox for their support

