

SJTU CCOE Annual Report and Renew Request



James LIN and Yizhong GU

Center for HPC, Shanghai Jiao Tong University

<http://hpc.sjtu.edu.cn>

15th February 2014

1. About SJTU CCOE

SJTU was awarded a CCOE in late 2011 and has been **the most active CCOE in China** since then. James Lin and Yizhong Gu, Directors of Center for HPC, are the co-PIs of SJTU CCOE, <http://ccoe.sjtu.edu.cn>.

2. Achievements in Year 2013

1) **Supercomputer: NO.1 Kepler-based Supercomputer in China**

- Π , the supercomputer of SJTU

Build in April 2013, the supercomputer in SJTU, currently **the fastest** supercomputer among the universities in China and also the **fastest one** in Shanghai, is named π . π , with a peak performance of 263 TFLOPS, ranked 204th of the TOP500 supercomputers in Nov 2013. It

utilizes "CPU+GPU+MIC+FAT" hybrid architecture, with 332 CPU nodes, **50 GPU nodes with 100 Kepler K20m**, 5 MIC nodes and 20 FAT nodes. In comparison with pure CPU



computing architecture, hybrid architecture is much more energy-efficient and can provide greater computational capabilities in some applications. π is also supported by a high-speed Infiniband 56G FDR network and 720TB shared storage system, which enable the carriage of data-intensive applications.

- Opening Ceremony

We invited NVIDIA China PSG GM Ashok Pandey to our opening ceremony in Oct 2013, after whole system test for 4 months including the hottest summer in Shanghai for recent 100 years.



- π is Open and GPU is well used.

π is the first supercomputer in China that puts real time information including system load online, <http://pi.sjtu.edu.cn>. Even surprise to us, SJTU users fully used the hybrid supercomputer, including Kepler GPU. Here is the snapshot of system load in 17th Dec, 2013. **We found 79% of 100 Kepler K20m have been used, by GROMACS and two in-house CUDA codes, however, 0% of Xeon Phi have been used.** We believe this is partly because we had promoted GPU and CUDA to our users so hard for so many years.

GPU Nodes (physical view)	
CPU's Total:	800
Hosts up:	50
Hosts down:	0
Current Load Avg (15, 5, 1m):	
79%, 79%, 79%	
Avg Utilization (last hour):	
79%	
Localtime:	
2013-12-17 11:43	

Mic Nodes (physical view)	
CPU's Total:	32
Hosts up:	5
Hosts down:	0
Current Load Avg (15, 5, 1m):	
0%, 0%, 0%	
Avg Utilization (last hour):	
0%	
Localtime:	
2013-12-17 11:43	

2) Promotion: Free Kepler Test System for NVIDIA Customers

- Shanghai Supercomputer Center in summer

Shanghai Supercomputer Center (SSC) hosted the most powerful supercomputer in China for 10 years before Tianhe-1 comes out. SSC now is planning for its 4th generation supercomputers target for year 2015. As a close partner of SSC, SJTU has promote GPU and CUDA to them for a long time. During early test of π last summer, SJTU help SSC test their GPU versions of FDTD code on 100 Kepler K20 of π for 2 weeks in September. The result will be present by SSC in GTC2014.

- India IISc in summer

Indian Institute of science (IISc), a customers of NVIDIA India, requested to extensively use the our supercomputer π for a month to test various aspects and show GPU acceleration for two critical apps in Molecular Dynamics and Quantum Chemistry – LAMMPS and QE. According to Ananda Sekhar Bhattacharjee of NVIDIA India, *“giving us access to the system was very critical at that juncture as it helped us to show the value proposition of GPU’s to one of India’s best research universities who also has lot of influence overall in the Indian scientific community”*.

- GTD program among China Universities since Nov

Because our supercomputer π is connected to major hub of China Education Network, so students and users in other China universities can access the π very smoothly, even they test large amount of data. Since last November, we have help at lest 7 students from other universities, and each of them have been awarded 10 Kepler K20 on π for 2 weeks for free usage.

3) Teaching: 1st HPC course taught in English in China

- SJTU CS075: the 1st HPC course in English, taught by James, Eric, and Jianwen (for team members, see Appendix 5.2). ~50% content is CUDA related.

Courseware is available online: <http://hpc.sjtu.edu.cn/Education/Courseware.htm>




- “GPU on π ” Seminar: We host this series seminar in SJTU every two-month.



4) Research: Optimization PIC on Kepler with an APS Fellow

- Particle-in-Cell code on GPU (Minghua Wen, James Lin, and Zhenming Sheng)

Original PIC code for Laser Plasmas Interaction Physics was developed by Zhenming Sheng, a distinguished Professor in SJTU and member of user committee for our center. He is a **Fellow of APS (America Physics Society)** since year 2012 and there are only two APS Fellows in China. We have been working with him since CCOE was awarded and now helping him optimize GPU code on Kepler K20 of π . Some kernels have been speedup to 30X compare to single thread CPU. The latest progress was published as a paper in HPC China 2013, see Appendix 5.1.




- National Computing Grid from MOST (James Lin)

Our Center becomes the 15th National Computing Grid (NCG) node of China in year 2013 and gets funding support from Ministry of Science and Technology (MOST) for our contribution the GPU Power on π into the national Grid. This program is similar to XSEDE in US, so the researchers in other universities who can access NCG will be able to use the 100 Kepler K20 on π .

5) Outreach: Largest student cluster contest in Asia, ASC13

Collaborating with INSPUR, vendor for π , our center hosted the largest student cluster contest finals in Asia in May2013. ASC was one of three biggest students cluster contest in the world. The other two are ISC in June and SC in Nov. 10 teams from 6 different Asia countries have attended this contest.



















Month	Events	Attendees	Activities
Feb	PPMM within PPoPP2013	30+	Organizer
Mar	GTC2013	1	Attend
Apr	ASC13 (Asia Student Cluster Challenge)	200+	Host
Oct	HPC China 2013	4	Paper and booth
Nov	SC13	2	Booth



3. Plan for Year 2014

1) **ICSC2014: HPC world leaders' summit in SJTU, 5~9 May**

We are so exciting to have so many HPC big names over the world, US, EU, Japan and also China. It could be the first time in China. <http://icsc2014.sjtu.edu.cn>. We will discuss about the challenge for Extreme Scale Scientific Computing. We will also have 2-day GPU pre-conference tutorial.

Keynote Speaker		Panel Discussion Moderator		
	Jack Dongarra Director Innovative Computing Laboratory		William Harrod Division Director U.S. Department of Energy	
Invited Speakers				
				
Jacqueline H. Chen (SNL)	Shiyi Chen (PKU)	Omar Ghattas (UT-Austin)	Kirk Jordan (CSC)	Petros Koumoutsakos (ETH)
				
Yutong Lu (NUDT)	Satoshi Matsuoka (Tokyo Tech)	Zeyao Mo (IAPCM)	Depei Qian (BHU)	Thomas Schulthess (ETH)
				
Rick Stevens (ANL)	Jinchao Xu (PSU)	Linbo Zhang (CAS)		
Tutorial Instructors				
				
L. Grinberg (IBM)	V. Morozov (ANL)	YH Tang (Brown)		

2) Supercomputer: K40 in production

Unlike some national supercomputer centers in China, we are willing to take some risks to apply the latest technologies into production mode. We have been testing on K40 for some GPU applications from our users for a while, and some performance is quite good compared to K20 on π . We request 10 K40 from NVIDIA and deploy them on π . We believe Both SJTU and NVIDIA will be benefit from this. By using these 10 K40, we can provide faster GPU computing power in production to our users, in the meanwhile, NVIDIA will get valuable feedback from SJTU about how K40 performances in a production mode, and users from NVIDIA GTD program and other VIP users introduced by NVIDIA can use and test these 10 K40 for free. We requested 10 K20 for GTD program last year and we will replace them with these 10 K40 on π .

We will upgrade incrementally our supercomputer and try to put something new into it every year. We have planned to upgrade all accelerators in early 2016, when π is 3 years old, by using Pascal. We will try to allocate 80% of workload onto GPU at that time.

3) Education: Fostering the next generation of China HPC

Current HPC community of China has some gaps with EU's and US's. Very few Persons can be recognized by the world, such as Prof. Satoshi Matsuoka does for Japan. It may partly because of English as 2nd language and the culture. So we believe we need teach our young men in English and help them have international vision and experience in HPC.

- GPU Computing Course with ANU and NCI from Australia

We have developed a HPC course focused on GPU computing with Prof. Alistair Rendell of Australian National University (ANU) and Joseph Antony of National Computational Infrastructure (NCI, the national supercomputing center of Australia and has the fastest supercomputer in Australia) for a while. This joint course in English will be taught in ANU in spring and in SJTU in autumn. As planed, Prof. Alistair Rendell, the vice dean of Research school of computer science, ANU will teach in SJTU for a week this autumn.

- Students oversea exchange program between SJTU and ANU

SJTU and ANU has agreed to exchange one student from each side for three months in year 2014. These two students will work as TA for the joint GPU computing course.

- “GPU on π ” Seminar

We plan to host this series seminar in SJTU every two-month, as last year. Current confirmed Seminars are listed as below.

Months	Topic	Keywords	Instructors	Attendee
Mar	CUDA and OpenACC	2-day, En	DevTech, CAPS	80+
May	Large Scale GPU Computing	2-day, En	ANL, IBM, Brown	100+

4) Research: GPU research collaboration with other CCOEs

- Performance Portability of **OpenACC** (with **Tokyo Tech CCOE** & **Georgia Tech CCOE**)

OpenACC is a directive-based programming standard initialed by NVIDIA, and other three compiler vendors. This research is led by James Lin, PI of SJTU CCOE, and collaborated with Satoshi Matsuoka, PI of Tokyo Tech CCOE and Jeffery Vetter, PI of Georgia Tech. We try to build an analytical performance model of OpenACC to predict the portable performance, by leveraging the OpenARC compiler from Jeffery’ team. Some preliminary results will be present by James during GTC14, see Appendix 5.1.

- **Quantum Espresso** (QE) on Kepler (with **Cambridge U. CCOE**)

QE is a software suite for ab initio electronic-structure calculations and materials modeling distributed under GPL. This research is led by Eric and James of SJTU, and collaborated with Filippo Spiga of Cambridge University CCOE, who is also the executive director of QE board. We try to develop Kepler version of QE together and test it on π of SJTU and Wilkes of Cambridge U.

- **GTC** on Kepler with OpenACC (with **Princeton Plasma Physics Lab** & **Tokyo Tech CCOE**)

Gyrokinetic Toroidal Code (GTC) is a massively parallel, particle-in-cell code for

turbulence simulation. The research is led by Minhua, James and Prof. Zhenmin Sheng, and collaborated with Prof. William Tang of Princeton University, who is the Chief Scientist at the Princeton Plasma Physics Lab. We try to develop an OpenACC version of GTC code on π of SJTU and TSUBAME of Tokyo Tech with portable performance.

5) Outreach: Bringing Students to Key HPC Conferences

Based on 10+ years teaching experience, we find some undergraduate and master students will be interested in HPC research after they attend HPC conferences. Some of them will even apply PhD program for HPC research later. It could be one of the best ways to have more talent young men involved in HPC. Unfortunately, like other China universities, SJTU doesn't support students to attend an academic conference without paper published; so most undergraduates and master students are not qualified to attend. We will use part of CCOE funding to financially support some talent students to attend these three conferences. One master student in SJTU will attend GTC14.

Month	Conference	Attendees (Include Students)	Activities
Mar	GTC14 in San Jose	2	Talk, Poster
Nov	HPC China2014 in Guangzhou	6	Paper, Booth
Nov	SC14 in New Orleans	5	Booth, Poster

4. Requested Support for Year 2014

	Our requests	Spec/Usage
HW	10 K40	For \$3.2, Deploying 10 K40 on π in production mode
SW	PGI Compiler	For \$3.2, X64+Accelerator for Linux-based supercomputer
Funding	\$10K	For \$3.3, Education: <ul style="list-style-type: none">● Local accommodation for ANU teaching course in SJTU● Airfare for SJTU student to ANU and local accommodation for ANU student during her 3-month stay in SJTU● Organization of 6 seminars, including pay for some instructors
	\$30K	For \$3.4, Research collaboration with other CCOEs: <ul style="list-style-type: none">● Salary for SJTU researchers involved● Local accommodation for Tokyo Tech, Georgia Tech, and Cambridge CCOE collaborators● Airfare to visit Georgia Tech and Cambridge● Matching salary for William Tang's Visiting Chair Professorship in SJTU
	\$20K	For \$3.5, for attending GTC14, HPC China 2014 and SC14: <ul style="list-style-type: none">● Airfare, registration fee, and local accommodation

5. Appendix







1) **Publications about GPU in Year 2013**

- [1] *James Lin, Satoshi Matsuoka*, OpenACC vs. OpenMP4: the Strong, the Weak, and the Missing to Develop Performance Portable Applications on GPU and Xeon Phi, **GPU Technology Conference 2014**, San Jose
- [2] *Minhua Wen, James Lin, Simon See*, Parallel PIC based on NVIDIA Kepler, **HPC China 2013** (acceptance rate<30%), Guilin, China







- [3] *Stephen Wang, Simon See and James Lin*, Performance Portability Evaluation for OpenACC on Intel Knights Corner and Nvidia Kepler, **HPC China 2013** (acceptance rate<30%), Guilin, China
- [4] *Minhua Wen, Zhanpeng Yu, Simon See and James Lin*, A NVIDIA Kepler Based Acceleration of PIC Method, **ParCFD 2013**, Changsha, China

2) Members in Center of HPC, SJTU

Staff








	<u>Yizhong Gu</u> Director		<u>James Lin</u> Vice Director
	<u>Minhua Wen</u> Engineering Computing		<u>Jianwen Wei</u> Big Data
	<u>Eric Germaneau</u> Scientific Computing		<u>Song Chen</u> System Support Engineer (Inspur)

Adjunct

	<u>Hong Liu</u> Adjunct Vice Director		<u>Haibing Guan</u> Adjunct Vice Director
	<u>Shenou Cai</u> Adjunct Vice Director		<u>Xianlong Jin</u> Adjunct Vice Director
	<u>Simon See</u> Adjunct Professor		<u>François Bodin</u> Adjunct Professor

The user committee of the center, including 6 professors from different departments of SJTU and James Lin, reviews resource applications from π users under a similar program to INCITE of DOE.

User Committee

	<u>Wei Cai</u>		<u>Zhaohui Jin</u>
	<u>James Lin</u>		<u>Xingzhao Liu</u>
	<u>Hao Lu</u>		<u>Zhengming Sheng</u>
	<u>Dongqing Wei</u>		