

Perspectives of China's HPC system development: a view from the 2009 China HPC TOP100 list

Yunquan ZHANG (✉)^{1,2}, Jiachang SUN^{1,2}, Guoxing YUAN³, Linbo ZHANG⁴

¹ Laboratory of Parallel Computing, Institute of Software, Chinese Academy of Sciences, Beijing 100190, China

² State Key Laboratory of Computer Science, Chinese Academy of Sciences, Beijing 100190, China

³ Institute of Applied Physics and Computational Mathematics, Beijing 100088, China

⁴ Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing 100190, China

© Higher Education Press and Springer-Verlag Berlin Heidelberg 2010

Abstract The China HPC TOP100 list, an annual report of the 100 most powerful high performance computing (HPC) systems installed in mainland China, has traced the rapid growth of HPC technology in China since its first publication in 2002. This paper introduces the China HPC TOP100 list and reviews the current status of HPC systems in China in terms of system features, manufactures, and areas of application using the data reported in the most recent list, published on November 1st, 2009. We provide further analysis, prediction of future trends, and directions of the development of HPC systems in China referencing historical data accumulated through archived TOP100 lists and other publically available information. We predict that the aggregated Linpack performance of the top 100 HPC systems will reach 10 PFlops in 2011, a single system with 10 PFlops peak performance will appear between 2012 and 2013, the aggregated performance of the top 100 systems will reach 100 PFlops in 2014, and a single system with 100 PFlops peak performance will appear around 2015.

Keywords high performance computer (HPC), China HPC TOP100 list, Linpack performance, trend analysis

1 Introduction

Since its first publication in 1993, the TOP500 list [1–3], which ranks the 500 most powerful supercomputers installed worldwide, based on measured High Performance Linpack (HPL) [4] benchmark performance, has become one of the most important references to assess the HPC market and application space in terms of technology development, application deployment, and analysis of future trends and directions. But until 2002, the TOP500 list included no high performance computer (HPC) systems installed in mainland China because no such systems published their HPL benchmark data.

This situation changed in November 2002 when China's own TOP50 list [5,6] was released by the Specialty Association of Mathematical and Scientific Software (SAMSS) [5] of the Chinese Software Industry Association (CSIA)¹⁾.

We organized this list, following the practices of the worldwide TOP500 list by ranking the top 50 most powerful supercomputers in mainland China using measured HPL benchmark performance data. Three weeks after its publication, the TOP500 list of November 2002 included the No. 1 system from China's TOP50 list, namely the Lenovo DeepComp 1800 system installed at the Chinese Academy of Science in Beijing, China, as the 43rd most powerful supercomputer in the world. The next edition of the TOP50 list dated November 2003 was expanded to include the 100 most powerful HPC systems in China and thus was renamed the TOP100 list [5,6]. In

Received April 5, 2010; accepted August 12, 2010

E-mail: zyz@mail.rdcps.ac.cn

1) China Software Industry Association, <http://www.csia.org.cn/home/index/index.htm>

November 2009, the 8th edition of the TOP100 list was published.

The China HPC TOP100 (henceforth TOP100) list has gained broad recognition ever since its first publication. Its collection of publically available HPL performance data follows the same standards and procedures adopted by the worldwide TOP500 (henceforth TOP500) list. In many specific cases, the authors traveled to installation sites to audit the benchmark data generation process in order to ensure its accuracy. The TOP100 list has been coordinated with the worldwide TOP500 list since its first release. It is a tradition that the English version of the TOP100 list be submitted to the organizers of the TOP500 list, and the TOP500 website, in addition to regularly maintaining a link pointing to the TOP100 website, frequently publishes special articles highlighting major updates in the TOP100 list.

Eight years since its first publication, the China HPC TOP100 list has become the de facto standard of reference in describing China's HPC market and application, and has been widely cited by government agencies, researchers, HPC users, and hardware/software vendors in technical reports, research publications, and grant proposals. The 2004 Annual Progress Report of China's Computer Science and Technology, edited by China Computer Federation (CCF)¹⁾, dedicated a chapter to describe China's HPC industry referencing data exclusively from the TOP100 list. David Keyes, Principle Investigator (PI) of the SCIDAC project²⁾ in the United

States, delivered a presentation on supercomputing in China in 2004 [7], which was based on statistical data from the TOP100 list.

The TOP100 list has observed several major shifts in both technology and user applications in China's HPC market place. The list has also served the purpose of promoting benchmark driven system development, shifting China away from the tradition of using theoretical peak performance to represent the computing power of a supercomputer. Both HPC customers and system vendors have often used the TOP100 list as one of the key system procurement indices.

The TOP100 list is partially funded by a project of China's National 863 Program on HPC Computer and Kernel Software³⁾.

In Section 2, we review the overall trends of development as reflected in the 2009 list and in Section 3 we discuss the market shares of key system vendors. We analyze the applications sector in Section 4 and finally in Section 5, we discuss perspectives and present our conclusions.

2 Performance analysis

In this section, we provide an in-depth review and analysis of the 2009 edition of the TOP100 list, focusing on the performance aspect of the overall trends of growth. Table 1 summarizes the Top 10 systems in the list [5].

Table 1 Top 10 systems in the 2009 edition of China HPC TOP100 list

Rank	Vendor	System	Install site	Install year	Application area	Processor cores	Linpack (GFlops)	Peak (GFlops)
1	NUDT	Tianhe-1	Tianjin Supercomputing Center	2009	Industry	24576	563100	1206210
2	DAWNING	Dawning 5000A	Shanghai Supercomputer Center	2008	Industry	30720	180600.00	233472
3	Lenovo	DeepComp 7000	SCCAS	2008	Scientific Computing	12160	106500	145293
4	IBM	HS22 Cluster	Network Company	2009	Gaming	7168	38790	72540
5	IBM	HS22 Cluster/	Network Company	2009	Gaming	7168	38790	72540
6	IBM	HS22 Cluster	Network Company	2009	Gaming	7168	38790	72540
7	IBM	HS22 Cluster	Nanjing University	2009	Education	3200	31310	34048
8	DAWNING	Dawning 5000	State Key Laboratory of Computational Physics, Beijing	2009	Scientific Computing	3360	31049	40320
9	IBM	HS22 Cluster	Network Company	2009	Gaming	5376	31030	54410
10	IBM	HS22 Cluster	Network Company	2009	Gaming	5376	31030	54410

1) <http://www.ccf.org.cn>

2) Scientific Discovery through Advanced Computing (SciDAC) program, <http://www.scidac.gov/>

3) National 863 High Performance Computer Evaluation Center, <http://www.hpctest.org.cn/>

The system on the top of the list is Tianhe-1, China's first petaflop system, installed at Tianjin Supercomputer Center. It is also the first and only CPU + GPGPU (general-purpose computing on graphic processing units) heterogeneous system in the TOP100 list. It outperforms the Dawning 5000A, which was No. 1 in the 2008 TOP100 list, by more than 300%. The birth of this petaflop system is one year ahead of our previous prediction [6] and has boosted China to become the second country in the world, after the United States, to achieve the capability to design and manufacture a petascale supercomputer. With its peak performance of over 1.2 PFlops and HPL max performance of 563.1 TFlops, Tianhe-1 is ranked No. 5 on the worldwide TOP500 list of November 2009 and is the only petaflop heterogeneous CPU + GPGPU system on that list.

Following Moore's Law, more and more CPU cores are integrated into the TOP100 HPC systems, the entry level to the list more than doubled in 12 months. The 100th system on the 2009 list reported 6.8 TFlops HPL performance, compared to 3.3 TFlops performance reported by the 100th system on 2008's list. The aggregated HPL performance of all 100 systems on the list has jumped to 2.2 PFlops, which represents a 112% increase over the aggregated performance from a year ago. The average HPL performance per system is 22 TFlops, which represents a 112% increase over that of 2008. In comparison, 22 TFlops was the average HPL performance per system in the worldwide TOP500 list dated June 2008. This is an indication that China's overall HPC market is behind the worldwide market by about 18 months.

The HPL performance of the 10th system in the TOP100 list is 31.03 TFlops, while this position achieved 17.68 TFlops on the 2008 list. All systems in the top 10 are blade clusters; among them nine are quad-core processor based. The average number of CPU kernels is above 3000 while the average peak performance reached 38 TFlops.

Examining from a system architectural perspective, 96 systems on the TOP100 list are clusters, while there were only 92 clusters on the 2008 list. Compared to the worldwide TOP500 list which has 83.4% of systems using cluster architectures, we can see that China has experienced higher cluster architecture penetration.

Examining the processor technology adopted by the TOP100 systems, we find that 4% of the systems run on the latest six-core processors while 77% of the systems run on quad-core processors. For the remainder of the

systems, 11% are built on dual-core processors and 8% run on single-core processors.

Among the TOP100 systems, 77% use Intel processors, a six-percentage point increase from the previous year. AMD captured a 21% share which remains flat compared to its share in 2008. IBM POWER processor based systems are down from 4% last year to only 2% this year. While MIPS-based systems had a 4% share in 2008, no MIPS systems are observed on the list this year.

Systems using gigabit Ethernet have the largest share of 55% on the list despite the significant share decrease from the previous year's 71%. In contrast, Infiniband shows strong momentum with 35% share on the latest list, representing an increase of 18 points from last year's 17%. 10 Gb Ethernet based systems have registered on the list with a 2% share; the remaining 8% are split between various proprietary interconnect technologies and are 4% lower than last year's level.

3 Vendor analysis

In this section, we analyze the statistics of hardware vendors, which can be categorized into 2 groups, domestic and international, as shown in Table 2. The statistics are also displayed in Figs. 1 and 2 showing the number of systems from each vendor, and aggregated Linpack performance, respectively.



There are seven Chinese domestic vendors whose products entered the TOP100 list and the total number of systems manufactured by these seven vendors is 45, representing a small increase from 43 systems reported in 2008. Dawning is no doubt the leading domestic vendor with a 27% share, far ahead of the other six domestic vendors, namely Lenovo, NUDT, Inspur, SUNWAY, Powerleader, and BJCC. IBM and HP maintain a noticeable lead within the group of international vendors, which also includes DELL, SUN and SGI.

Four of the ten most powerful systems on the TOP100 list were manufactured by China's domestic vendors including Tianhe 1, the No. 1 system on the list, designed and manufactured by National University of Defense Technology (NUDT). The other six of the top 10 systems were manufactured by a single international vendor, namely IBM. Two other domestic vendors, DAWNING and Lenovo, were responsible for the remaining three systems in the top 10.

However, if we sort the statistics based on the aggregated Linpack performance as plotted in Fig. 2, we

Table 2 Vendor statistics based on 2009 edition of China HPC TOP100 list

Manufacturer		Share/%	R_{\max} /(GFlops)	R_{peak} /(GFlops)	Efficiency/%	Cores
Domestic vendors	DAWNING	27	450745	636372	69.38	78028
	Inspur	6	53511	78612	70.82	8088
	Lenovo	3	126692	182275	60.83	16128
	SUNWAY	3	50742	64495	80.23	6096
	PowerLeader	3	49018	68142	69.67	5912
	NUDT	2	570155	1215544	60.20	25600
	BJCC	1	7138	8823	80.90	976
Domestic total	45	1308000	2254262	70.29	140828	
International vendors	IBM	26	542541	981515	56.63	96800
	HP	23	293357	510894	58.90	58495
	DELL	4	42369	61322	72.65	5580
	SUN	1	10460	13584	77.00	1200
	SGI	1	7127	8192	87.00	1280
International total	55	895854	1575507	70.44	163355	
Total	100	2203855	3829770	64.13	304183	

**Fig. 1** Vendor number of systems share on 2009 China HPC TOP100**Fig. 2** Vendor Linpack performance share on 2009 China HPC TOP100

find China's domestic vendors collectively contribute 59.3% while international vendors contribute 40.7%. The increase of the domestic vendor share compared with the 2008 list is slightly greater than 3%.

This is the first time in Chinese history that a domestic vendor, namely Dawning, has become the largest hardware vendor in the TOP100 with a total of 27 systems on the list replacing IBM who held the top spot for the past 7 years. With 26 systems on the list, IBM is the No. 2 vendor and HP is the No. 3 vendor with 23 systems on the list. Domestic company Inspur took the No. 4 position with 6 systems and DELL has the No. 5 position with 4 systems on the list. It is worth noting that China's domestic vendors have been steadily increasing their share during the past few years despite the fact that a 49% share is still held by two international vendors; IBM and HP.

Also for the first time in Chinese history, NUDT became the top manufacturer for aggregated Linpack performance in the TOP100, contributing 25.8%. IBM moved its position to the No. 2 spot with 24.64% of aggregated Linpack performance even though this number represents a relatively large increase from 18.06% share reported on 2008's list. Dawning took the third place of aggregated Linpack performance with 20.4% share, which showed a large decline from 32.68% on the 2008 list.

4 Areas of HPC application analysis

In this section, we examine the application areas in which the TOP100 systems are used. Our analysis can be divided

Table 3 Statistics of 2009 China HPC TOP100 list based on areas of applications

Areas of applications	Share/%	$R_{\max}/(\text{GFlops})$	$R_{\text{peak}}/(\text{GFlops})$	Efficiency/%	Cores
Energy	20	254314.86	481548.42	54.39	48176
Online gaming	15	325516.00	592010.00	55.13	62836
Scientific computing	14	269046.67	359545.36	76.79	32144
Meteorology	10	114200.37	152169.64	77.70	16800
General industry	9	813762.46	1535635.16	70.76	66792
Government	9	112323.00	207515.96	55.88	25674
Education	7	95487.47	117343.52	79.56	12108
TeleComm	6	100129.80	181019.78	57.77	19849
Bioinformatics	4	34914.00	58464.40	62.68	5596
Earthquake engineering	2	37372.00	50066.08	76.15	4608
B2C	1	16190.00	35200.00	46.00	3520
Animation rendering	1	12115.26	22131.20	54.70	2080
Financial	1	9287.00	17920.00	51.80	1600
Video computing	1	9196.00	19200.00	47.90	2400
Total	100%	220384.89	3829769.52	64.13	304183

into 2 categories based on the number of systems installed or the aggregated Linpack performance, as shown in Figs. 3 and 4, respectively.

Our analysis shows that the China HPC TOP100 systems support 14 areas of application, as summarized in Table 3. Energy (geophysics applications used by the petroleum industry) is the top area of application with a 20% share. Online gaming accounts for 15% of the system deployment, followed by the No. 3 application, scientific computing, which takes a 14% share. Meteorology applications (including weather modeling and forecast) captures a 10% share, fourth place on the list, while general industrial applications and government each take a 9% share. Education (universities) has a 7% share followed by Telecommunications, which has a 6% share. Bioinformatics and another 5 areas of applications accounts for the final 10%.

Compared to the 2008 list, energy maintains the top position in area of application in terms of number of systems on the list as shown in Fig. 3. The industry remains the largest buyer of large systems although its share has declined sharply from the 2008 level when it had a 35% share. Meanwhile, online gaming has become the fastest growing application with its share doubled since 2008. This is a unique phenomena and a new development direction in China's HPC market. A report, "New Dimensions in HPC," released by Tabor research on April 2008 classified the HPC market into traditional HPC (tHPC) and edge HPC (eHPC) [8]. eHPC leverages the

experience and technology of the traditional HPC market, while introducing new areas for innovation. Most importantly they believe that eHPC is at the cusp of significant market generation and growth [8]. Tabor divides the eHPC market into four major segments: Complex Event and Business Processing, Process Optimization, Virtual Infrastructure and Environments, and Ultra-scale Computing [8]. Online gaming belongs to the Virtual Infrastructure and Environments segment.

However, if we examine the list based on the aggregated Linpack performance, as plotted in Fig. 4, we will find that the general industry, with its 37% share, is the area of application that uses the highest number of HPC systems on the TOP100 list. Online gaming maintains its No. 2 place with almost 15% share, which is roughly the same as percentage number of systems deployed. The No.3 position is held by scientific computing applications with a 12.21% percent share.

The top five areas of applications on China HPC TOP100 list are not identical to those on the worldwide TOP500 list. The top five on China's list are energy, general industry, online gaming, scientific computing, and meteorology while on the worldwide list, we find scientific computing, financial, information services, energy and logistics. Only scientific computing, and energy are among the top 5 applications on both the TOP100 and TOP500 lists.

Figure 5 examines the trend of industrial and commercial adoption of HPC systems based on eight editions of

the TOP100 list. If we reclassify areas of application into the following three segments, namely commercial, higher education, and scientific research, we find the ratio of HPC system deployments in these three segments to be 63:7:30. Systems installed at commercial institutions actually fell by 14% from 2008 level. While installations in higher education have increased slightly, we observe a rapid increase in the area of scientific research, which used to maintain a historical leadership on the TOP100 list. This is evidence of China's national reinvestment in pure HPC technology research and will further promote the development of HPC methodology and application software, strengthening the foundation of commercial HPC applications. We recommend that China's national strategic plan be developed at mid and long-term levels, and national supercomputing research and application centers be established in order to stabilize funding and to coordinate the research and development in computational science, HPC hardware technology, and especially application software and tools.

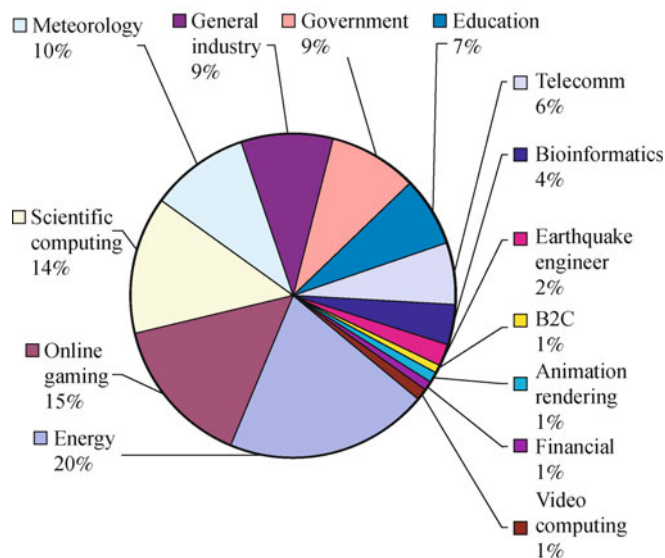


Fig. 3 Number of system share of application area on 2009 China HPC TOP100

5 Perspectives and conclusions

Figure 6 shows an outline of performance trends of China's HPC systems from 1993 to 2009 and also includes our projections from 2010 to 2017. To predict the performance trends, we use the $y = 1.9548e^{0.8205x}$ exponential fitting function for the total performance and the $y = 0.1937e^{0.883x}$ exponential fitting function for the




Fig. 4 Performance share of application area on 2009 China HPC TOP100




Fig. 5 Areas of HPC applications documented by China TOP100 list (2002–2009)

No. 1 peak performance. Historical data publicly available covering the 10 years before the first release of TOP100 was also included in this figure.

China has a long history developing its own high performance computer systems. We collected the statistics data of the largest HPC systems installed in China since 1993, which was nine years ahead of the first publication of the China HPC TOP100 list. Figure 6 summarizes the

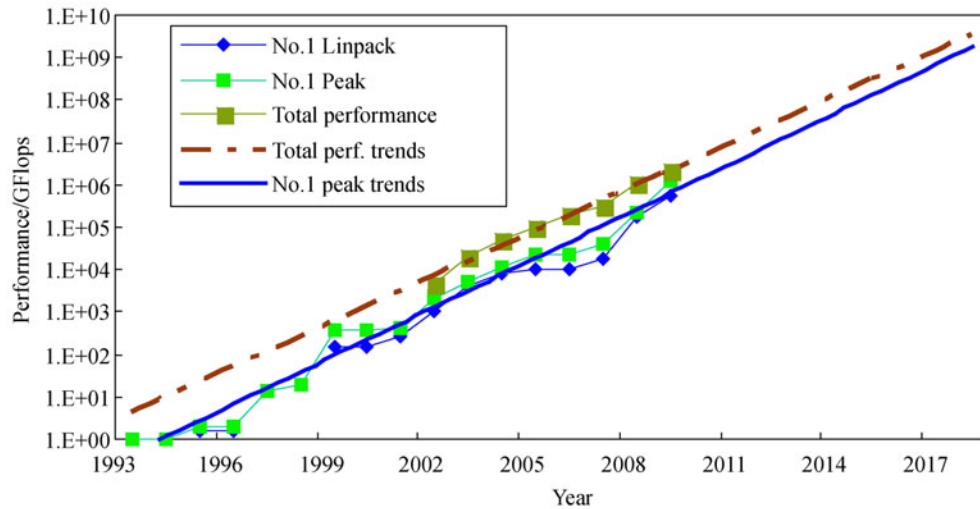


Fig. 6 Performance trends of China's HPC systems from 1993 to 2009 and projections from 2009 to 2017

historical trends of the HPC systems in China covering the 26 years time span between 1993 and 2009, in which we include pre-TOP100 data that has never been published in such an integrated format. Even though the pre-TOP100 data are not as accurate as that collected after 2002 due to large user environment variations and other factors, it still provides insight into the events that shaped the HPC industry and user application space in China.

Figure 6 includes Linpack performance numbers for both the No. 1 HPC system as well as the aggregated performance of all the systems on the TOP100 list. Before 2003, there were only a limited number of HPC systems in China so the aggregated system performance between 1993 and 2002 is roughly the same as that of all the HPC systems in the country.

The growth of HPC systems in China has not been steady during the past 26 years. It remained flat for 3 years between 1993 and 1996 before a major growth spurt was observed during the next three years, represented by the general release of the Dawning 1000 in 1996 and the SUNWAY-I in 1999. This was followed by three years of marginal growth between 1999 and 2001. Beginning in 2001, China's HPC industry entered a new period of rapid growth with the release of the Dawning 3000 in 2001, and the momentum continued for the following 4 years with the follow-up release of both the Dawning 4000A and DeepComp 6800 systems in 2004, before returning to steady development. The incremental increase between the No. 1 systems on the 2006 and 2007 lists indicated the arrival of the next period of growth. 2008 was marked by the birth of two 100 TFlop systems in China, Dawning

5000A and DeepComp 7000, which coincides well with our projection. However, the release of Tianhe-1 in 2009 fast-forwarded China's HPC industry into the PFlop era, which is ahead of our projection by a full year.

Based on the newly collected information, we expect that the aggregated Linpack performance for all TOP100 systems will reach 10 PFlops during 2011; a system with 10 PFlops performance will arrive between 2012 and 2013; the aggregated Linpack performance will reach 100 PFlops/s around 2014; and finally, a system with 100 PFlops performance will appear around 2015.

The China HPC TOP100 list is an annual report of the 100 most powerful HPC systems installed in mainland China and has traced the rapid growth of HPC technology in China since its first publication in 2002. We have briefly reviewed the historical development of the list and described in detail its eighth edition, published in November 2009. In conclusion, we summarize the highlights of the 2009 list, as follows:

- 1) China successfully installed its first petascale HPC system, one year ahead of expectations; the designer and manufacturer of the system, NUDT, captured the No. 1 spot on the TOP100 list for the first time.

- 2) Two HPC systems are able to push Linpack benchmark efficiency beyond the 90% level. They are an IBM blade system installed at Nanjing University and a SUNWAY 4000A system installed by the Jiangnan Institute of Computer Technology.

- 3) For the first time in history, a Chinese domestic vendor, Dawning, captured the highest market share on the TOP100 list.

4) The entry point for the TOP100 list rose to 6.8 TFlops; Linpack performance aggregated over all TOP100 systems exceeded 2 PFlops; average Linpack performance per system exceeded 22 TFlops.

5) Cluster architecture was adopted by 96 systems on the TOP100, making it more urgent to develop a system with a novel architecture. It is extraordinary and refreshing that the No. 1 system on the TOP100 list employs CPU + GPGPU hybrid architecture.

6) Fourteen system vendors contributed HPC systems to the TOP100 list.

7) Among 14 areas of applications, the top five are energy, online gaming, scientific computation, meteorology, and general industry.

Looking at 2010 and beyond, we look forward to the arrival of groundbreaking events such as an HPC system designed and manufactured by a Chinese domestic vendor that delivers petascale Linpack performance using a general purpose CPU designed and made in China.

Acknowledgements This work was supported in partial by the National High Technology Research and Development of China (863 Programme) (Nos. 2006AA01A105, 2009AA01A129, and 2009AA01A134). We would like to thank Dr. David Chen for his valuable suggestions during the preparation of this paper.

References

1. Meuer H, Simon H, Strohmaier E, et al. TOP500 supercomputer list, <http://www.top500.org/list/2010/06/100>
2. Dongarra J. An overview of high performance computing and challenges for the future. San Diego, CA, USA: SIAM. 2008
3. Dongarra J. Performance of Various Computers Using Standard Linear Equations Software, (Linpack Benchmark Report). University of Tennessee Computer Science Technical Report, CS-89–85, 2010
4. Petitet A, Whaley C, Dongarra J, et al. HPL Benchmark 2.0, <http://netlib.org/benchmark/hpl/>
5. Sun J, Yuan G, Zhang L, et al. 2009 China HPC TOP100 rank list. 2009, <http://www.samss.org.cn/sites/shuxue/pcC.jsp?contentId=2502334815421>
6. Zhang Y, Sun J, Yuan G, et al. A brief introduction to China HPC TOP100: from 2002 to 2006. In: Workshop on High Performance Computing in China: Solution Approaches to Impediments for High Performance Computing, Supercomputing 2007. Reno, 2007, 32–36
7. Keyes D. Supercomputing in China. In: Workshop on High Performance Computing with BG/L and QCDOC. Brookhaven National Laboratory, USA, 2004, http://124.16.137.70/bnl_china.pdf
8. HPCwire. New Dimensions in HPC, http://www.hpcwire.com/features/New_Dimensions_in_HPC.html?page=1