Front. Comput. Sci. China 2010, 4(4): 428–436
DOI 10.1007/s11704-010-0118-z

REVIEW ARTICLE

Evolution of supercomputers

Xianghui XIE (✉), Xing FANG, Sutai HU, Dong WU

Jiangnan Institute of Computing Technology, Wuxi 214083, China

Abstract Supercomputers are prevalent and vital to scientific research and industrial fields, and may be used to represent the level of national scientific development. A summary of the evolution of supercomputers will help direct the future development of supercomputers and supercomputing applications. In this paper, we summarize the accomplishments in supercomputing, predict the trend of future supercomputers, and present several breakthroughs in supercomputer architecture research.

Keywords supercomputer, scientific computing, architecture, challenging applications, novel devices

1 Role of supercomputers

Supercomputers are computing systems with the highest performance, largest capability and largest capacity, which provide supercomputing capabilities for challenging issues in a period of time. The research of supercomputers has strategic significance, which is vital to national innovation systems, and represents the national scientific and technological level and capability. Supercomputers have become an international strategic point, especially for developed countries.

All advanced countries include the development of supercomputers in national mid and long-term plans, this continuously advances the development of supercomputers. In the United States and Japan, the government, academia, and commercial companies cooperate in developing supercomputers, focusing on both construction and application. The governments lead the development of high-end supercomputers, and commercial companies manufacture low-mid end productions in large volume. The focus of supercomputing in European countries is directed more towards the application of supercomputers and promotes the sharing of high-performance computing facilities.

Since the birth of supercomputers, they have contributed to many critical computing areas, driven forward issues of national security, and economic and social developments, and also played a critical role in scientific discovery and innovation. Currently, supercomputers are widely used in areas of national security information systems, energy security, climate, aerospace and cosmology, medicine and healthcare, manufacturing, environment protection, and high-energy physics, etc.

2 Evolution of supercomputers

2.1 Evolution of systems and performance

In the past 40 years, the evolution of computer components has promoted the development of supercomputers; supercomputers have undergone four stages: bud stage, booming vector supercomputers stage; flourishing massively parallel processing (MPP) supercomputers stage, and the cluster system stage which is illustrated in Table 1. The performance of supercomputers has improved at a stable rate about a 10-fold per four years: faster than Moore's law. Since the emergence of CDC 6600 in the 1960s, the performance of supercomputers has improved about 10^{10} times. Following this trend, we predict 10 Petaflop supercomputers will emerge in 2012, 100 Petaflop supercomputers will appear in 2016, and 1 Exaflop supercomputers will arrive in 2018.

Received May 31, 2010; accepted September 13, 2010

E-mail: xhxie@ict.ac.cn

Table 1 Evolution of supercomputers

Stages	Budded stage (1960–1975)	Vector stage (1976–1989)	MPP stage (1990–2000)	MPP & cluster stage (2000–2010)
Representive system	CDC6600 STAR-100 ILLIAC-IV	Cray-1/2/3 NEC SX-1/2/3	Intel Paragon TMC CM-5 Cray T3D Intel Option Red	SGI Blue Mountain Compaq ASCI Q NUDT "Tianhe #1" IBM Roadrunner
Performance	Mflops	Gflops	Tflops	Pflops

2.2 Evolution of architecture

Supercomputer architecture is the interface between hardware and software, which directly affects the scalability, programmability, and availability of the system. The evolution of the supercomputer architecture has undergone the following stages: vector machine, multi-processor machine, parallel vector machine, customized MPP based on commodity processors, cluster based on commodity components, and the emerging hybrid architectures, which is illustrated in Table 2. Currently, tightly-coupled customized MPP machines and hybrid architecture machines are still the mainstream of superior supercomputers, such as the Jaguar XT5 and the IBM Roadrunner which held 1st and 2nd place respectively on the top500 list in 2009 November.

2.3 Evolution of software

From the viewpoint of management and operation modes, the system software of supercomputers undergoes four stages: batch process, multi-channel programming systems, operating systems, and high performance systems. The operating system evolves from single-user, single-task to multi-user, multi-task as the system software moves from batch processing to high performance systems.

The programming languages undergo a transition from low-level assembly languages to high-level languages. High-level languages experience a transition from earlier high-level languages to structural languages, from process-oriented to object-oriented. The parallel programming model undergoes the following stages: data parallel programming model, messaging passing programming model, shared-memory programming model, and hybrid programming model. The former three programming models are basic parallel programming models, while hybrid programming models, such as "MPI + OpenMP", are composed of more than two basic programming models and targets are better exploiting the inherent parallelism of hybrid architectures. In the "MPI + OpenMP" hybrid programming model, OpenMP is often used to express the thread-level parallelism within a single node, while MPI is used to express the parallelism across nodes in the whole system. The evolution of supercomputer system software is illustrated in Table 3.

Table 2 Evolution of supercomputer architecture

Year	1980	1990	2000	2010
Architecture	SIMD	SMP + S2MP	MPP	Cluster
	Single processor	MPP	Constellations	MPP
	SMP	Single processor	Cluster	Hybrid architecture

Table 3 Evolution of supercomputers system software

Year	1960–1980	1980–2000	2000–2010
Operating System (OS)	Single-user, Single-task OS	Multi-user, Multi-task OS	New emerging OS
	TSS/360, TOPS-10, VM/CMS, COS	UNIX variants, such as AIX, HPUX, Solaris, BSD, IRIX, Tru64 UNIX, UNICOS, VPP300/VPP700 UXP/V, etc.	Linux variants, such as RedHat, SUSE, Debian, Ubuntu, SUPER-UX, etc.
Programming languages	BASIC, Fortran IV/66/77, Pascal, C, Smalltalk	Object-oriented Languages, such as C++, HPF, Java, Parallel C, etc.	Emerging Languages, such as Charm++, Fortress, Chapel, X10, Brook, PGAS
Programming models	Data parallel programming model	Message passing programming model, share-memory programming model	Hybrid programming model

2.4 Evolution of design principles

Along with the growth of supercomputers, scale, and the transition of supercomputer use from computing centered to service centered, the metrics of supercomputer systems have transitioned from only high performance to high performance, high productivity, and high efficiency. The design principles have transitioned from peak performance centered to available performance or sustained performance centered, from operating efficiency centered to time to solution (TTS) centered, from building cost centered to total ownership cost centered, from machine oriented to network oriented. The operating of supercomputers on network environments has become the future trend. The evolution of design principles is displayed in Table 4.

2.5 Evolution of applications

The application of supercomputers has been widely extended, and passed through the following eras: scientific computing dominant era, coexistence of scientific computing and data processing era, and the network centric era.

Currently, supercomputer application domains are developing in two directions: capability computing and capacity computing. Capability computing is targeted at critical scientific challenges, focusing on computation power to reduce the total time cost to solution. Capacity computing is targeted at communication intensive applications, focusing on throughput, which means finishing as many tasks as possible in the specified time domain.

The use mode of supercomputers exhibits a 20 year cycle of distributed-collective interleaved dominant phenomena. Early supercomputers were domain specific specialist. With the emergence of the notion of high performance computing, supercomputers have penetrated every domain. Since the year 2000, as basic information infrastructures, supercomputers have started to adopt a service oriented architecture (SOA) infrastructure, in the direction of ubiquity, computation center→grid→cloud. The evolution of applications is displayed in Table 5.

3 Challenges for future applications

Supercomputers have provided the capability of prediction, analysis and political evaluation to human society. Whereas, the development of applications provides supercomputers with new challenges.

Table 4 Evolution of design principles

Principles	High performance	High performance/cost ratio	High productivity	High efficiency
Metrics	• Clock frequency • Peak performance (flops)	• Clock frequency • Peak performance • Cost	• Peak performance • Robustness • Programmability • Portability	• Sustained performance • Power efficiency • Volume efficiency • Availability • Cost efficiency (Cost efficiency =(Peak performance*Utilization ratio/Total ownership cost))

Table 5 Evolution of applications

Year	1980–	1990–	2000–	2010–
Application area	Scientific computing	Scientific and engineer computing	Scientific and engineer computing, data processing	Scientific and engineer computing, Data processing, intelligent processing
Property	Capability computing	Capability computing	Capability computing and capacity computing	Capability computing and capacity computing
Scale (nuclear physics)	10^6 operations	10^9 operations; dimension is increased from 2-D to 3-D; computation requirement is increased by 20–100 times.	10^{12} operations, average grid size is reduced by 2–10 times; physical model complexity is increased by 2 times; computation requirement is increased by 100–200 times in 6 years	10^{15} operations, average grid size is reduced by 4–6 times; physical model complexity is increased by more than 100 times; computation requirement is increased by 5000–10000 times in 14 years
Use mode	Domain specialists	Domain specialists + computing center	Computing center + grid	Computing center + grid + cloud computing

3.1 Scientific challenges

Scientific challenges constitute the persistent challenges for supercomputers, such as earth science, climate, materials science, energy science, biology and medical science, academic and industrial applications, and national security domains, etc. Solutions to these critical scientific problems will reveal many scientific principles, and facilitate comprehensive realizations of nature. Since the internal mathematical expression of these scientifically challenging problems is always a complicated correlated equation set, representing many physical processes at different scales, and each physical process requires great computing capability, tightly coupled hardware and software are essential for correlated large volume physical processes.

With the improvements of modern supercomputer capabilities and the requirements of scientific research, scientific computing keeps on evolving in its research goals, simulation scale, problem size and models, etc.

According to discussions in recent scientific computing reports and workshops [1–10], we sum up the following trends in the evolution of scientific challenges:

- Higher resolution, greater dimensions, and larger problem sizes or geometries.
- Coupling of several models.
- Simulation at multiple temporal and spatial scales.
- Advanced equation set solvers.
- More first-principle rules: using more first-principle physical models, chemical models and biological processes to provide higher resolution, higher quality, and accuracy, and to reduce uncertainty.
- More data centric: verification of new theorems and models, simulation result validation, uncertainty quantification, and new property and characteristic prediction, are all based on data analysis and data mining and data synthesis of enormous amounts of simulation and experimental data.

To meet the requirements of scientific challenges, extreme scale computing capable systems are required. Cited as an example, a 2009 workshop sponsored by the U.S. Department of Energy (DOE) Office of Science and DOE Office of Nuclear Energy developed a detailed picture of the computational requirements for nuclear energy modeling, culminating in a 1 exaflop requirement for 3D nuclear reactor design simulation by 2018, and a 10 exaflop requirement for 3D nuclear plant design simulation by the year 2024 [6]. By enabling the modeling and the simulation of critical nuclear power system components with improved geometric fidelity, improved numerical fidelity and improved physics fidelity, exascale computing would change nuclear engineering from a test-based to a science-based discipline.

3.2 Data deluge challenges

Over the past few years, the size and bandwidth of the Internet has grown steadily at a high rate, which promotes the paradigm shift of supercomputing and the emergence of a new supercomputing form that is totally different from conventional supercomputing and sometimes serviced by large-scale data centers, such as the Google translation service and Google web indexing service based on Google file system (GFS) [11] and Google map/reduce programming model [12]. Different from the modeling and simulation of scientific computing applications, these kinds of applications provide services through data analysis, data indexing, data mining and data synthesis of the data captured or collected from the Internet.

Challenges of this kind of application mainly stem from data deluge caused by the rapid increase in the amount of data: especially unstructured data. According to the research of IDC, by 2011, the digital universe will be 10 times the size it was in 2006. In 2011, the amount of digital data produced should be nearly 1800 exabytes, or 10 times that produced in 2006. About 94% of the digital universe is opaque and unstructured, much of which is contributed by Web 2.0 service applications and mobile applications [13]. Data transfer on the Interenet drives Internet flux to grow at rate 42% per year. Searching for meaning in the content of unstructured data such as images, video clips, documents and the numbers and characters in databases is the rocket science of the digital universe [14].

The incredible growth rate of data, especially unstructured data, means more sophisticated and more intelligent information management schemes, security schemes, information indexing schemes, storage schemes, and data center architectures should be introduced into the information infrastructure.

Furthermore, constrained by interconnect bandwidth; transfer of mass-data is expensive, which means computation should be performed as close to the data as possible. So, parallelism scheduling should take network topology into account, which promotes a transition of toward network centric supercomputers. The performance of

these applications is mainly constrained by I/O performance supporting large files transferred between disk and memory, the memory capacity of a single processor, and memory access bandwidth and latency. Throughput and efficiency are main performance metrics for systems targeted at these applications. The evolution of supercomputer applications is illustrated in Table 6.

4 Technology challenges in future supercomputers

As projected by the supercomputer community, an exascale supercomputer will emerge in about 2018. Currently, the research on exascale supercomputers has been carried out across the world [15–17]. Challenges of the exascale supercomputer can be summarized as follows:

4.1 Power

Power will be the biggest problem for the exascale supercomputer. Following current design pattern, along the traditional roadmap, the power of an exascale supercomputer will be approaching several tens of MW [10], as illustrated in Table 6, about more than 10 times the power required for a petascale system even taking the development of greener technology into account, and assuming aggressive circuit and architecture design to reduce power. Currently, it is rare that computing/data centers can support and afford such high power consumption.

4.2 Resiliency

Compared to petascale systems, exascale systems will be confronted with more significant resiliency problems, caused by following:

- More processing cores and chips are integrated into an exascale system, from several millions of memory chips and lots of storage devices.
- With the development of semiconductor technology, the feature size of devices will shrink, which will make devices more sensitive to electromagnetic fields, cosmic radiation and other interference sources.
- Higher interconnect bandwidth requirements and longer interconnects arising from geometric growth will make signal transfer more susceptible to interference sources in an exascale system. Furthermore, signal degradation will also be more severe.
- Operating at low voltage will benefit the power consumption; however it will make the system more susceptible to noise sources.

As predicted, the mean time to error in exascale systems will drop below one hour, however, the time consumption for storing for recovery at a checkpoint or recovering from a checkpoint will increase, and approach half an hour [16]. The former will shorten the interval between two neighboring checkpoints, whereas the latter will increase the invalid time spent on checkpoint recovery, thereby shortening the operating time of the hardware. To remedy this, more hardware devices are required to attain the target performance, which will result in a greater error rate.

4.3 Interconnect and scalability

As predicted by [16], the number of nodes in an exascale system will be 6 to 8 times that of a petascale system, memory bandwidth per processor will be ten times that of a petascale system, and interconnect bisection bandwidth will be about four hundred times that of a petascale system. All of which means it is necessary to greatly increase the amount of memory, or increase the number of I/O pins, or improve the bandwidth of a signal pin. Furthermore, we can see that the growth ratio of

Table 6 Evolution of supercomputer applications

Year	2004	2007	2012	2015	2019
	BG/L	BG/P	25 PF	300 PF	1200 PF
Cores/Node	2	4	8–24	32–128	96–128–500
Clock speed/GHz	0.7	0.85	1.6–4.1	2.3–4.8	2.8–6.0
Flops/Clock/Core	4	4	8–32	8–32	16–64
(Flops/Node)/Gflops	5.6	14	128–640	640–2000	2000–6000
Total cores/Millions	0.13	0.3	0.3–1.2	1–10	4–200
Total RAM/TB	33.6	151	2000–4400	3000–10000	5000–25000
Total power/MW	2.5	4.8	8–20	20–50	30–80
Peak performance/PetaFlops	0.37	1	25	300	1200

interconnect bisection bandwidth is about 70 times the growth rate of system scale, which means not only 6 times the number of switches and routers are needed, but also the bandwidth of a single link should be increased by 70 times. So a high-radix network topology is required to improve bisection bandwidth and reduce the average number of hops, which results in more links, transmitters, receivers and repeaters, etc., thereby greatly increasing the routing cost, size, and difficulty.

4.4 Concurrency and locality

The growth of system scale and the number of cores in a single processor require exploiting more thread-level parallelism, while memory-wall and latency make the problem more difficult. Although flattened clock frequency growth curve alleviate the burden, a significant increase in concurrency means more high-latency events. Furthermore, management, caching, and routing of these high-latency events will induce further latency. Synchronization and interaction between threads will result in exponential growth of the number of waiting independent devices, so the data locality should be increased to mitigate these delays, as "Hardware trend" line illustrated in the Fig. 1.

To meet the requirements of concurrency and scalability, good spatial and temporal locality of applications and algorithms is preferred. However, scientific applications are going to exhibit less data locality, as "Science going like this" line illustrated in Fig. 1 [16], the gap between the locality and the hardware currency level seems to be growing.

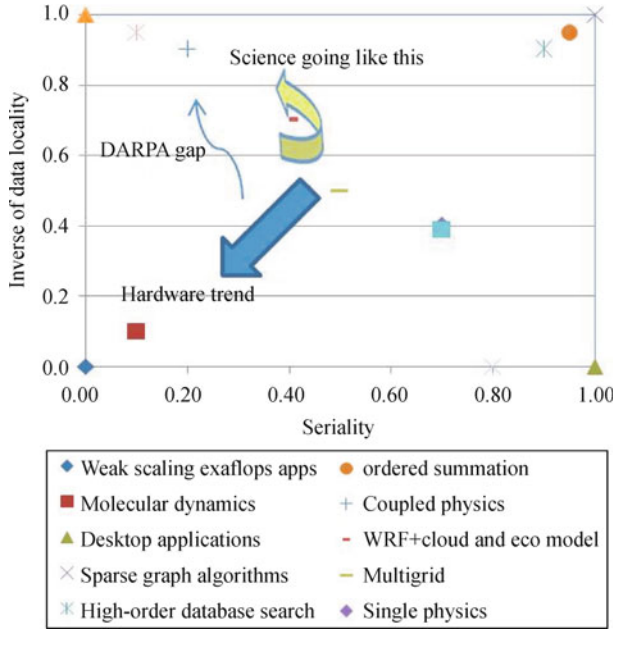

Fig. 1 Projected locality gap between hardware parallelism and locality of applications

4.5 Developing new algorithms

Memory capacity/computation ratio (bytes/flop) continues to descend. In the past decades, this ratio was kept at around 1 bytes/flop; however, it has dropped to 0.3 bytes/flop for current supercomputers constructed from commodity devices. For GPUs, this ratio has even dropped to between 0.002 and 0.01 bytes/flop. In the prototype of exascale systems, this ratio is estimated to be 0.0036 bytes/flop [17], which means algorithms should execute more operations on smaller data sets to maintain the utilization ratio of the system due to the interconnect wall.

Furthermore, since the performance improvement of the supercomputer is becoming more and more dependent on increasing parallelism, fully exploiting the parallelism of the architecture is a big challenge for the development of new algorithms.

4.6 Programmability

Currently, GPUs and multi-core processors are increasingly prevalent in supercomputers; however, GPUs and multi-core processors also present new programmability challenges:

- How to represent heterogeneity and complicate concurrency levels of systems for applications to fully exploit the potential of such architectures.
- How to make the trade off between the programmability and programming productivity to improve programming efficiency.
- How to resolve legacy application code issues.

5 Opportunities with novel IC materials and technologies

In most areas, current technologies and design methods cannot resolve challenges confronted by exascale systems, however emerging novel integrated circuit materials and technologies provide new opportunities for implementing exascale systems, and will promote the innovation of critical devices. Currently, novel IC materials and technologies that have been widely discussed include:

5.1 3D stack

3D stack technology integrates multi-level dies into a single chip with low-latency, high bandwidth and density through-silicon vias (TSV); this can reduce the global wire length, and thereby reduce latency. Furthermore, 3D stack technology provides support for the integration of various technologies, for example, the combination of high-speed CMOS circuits with high-density DRAM.

The main benefits of 3D stack technology are as follows [18,19]:

• Integration of several technologies, for example, combined photonic/electricity circuits, combined DRAM/Logical circuits, combined digital/analog RF circuits, etc. So each circuit can be fabricated in the optimal technology.

• Low-latency on-chip communication and large on-chip memory. Since the height of each die is only tens of micrometers, TSVs are very short, communication through 20 layers will cost only 12 ps, which not only significantly reduces long-wire latency, but also provides support for directly stacking multiple DRAM layers onto the processor cores. At 50 nm technology, a single processor can integrate several tens of GBs on-chip DRAM.

• High-density on-chip communication. Since the size of a TSV is only about 4–10 μm, even assuming a high-end 10 μm, a bus consisting of 1024 bits will consume only 0.32 mm^2, which means more than 300-way 1024-bit bus can be integrated on 1 cm^2 area. As far as multi-layer DRAM is considered, size and latency of the TSV will not be a constraining factor in the following few technology generations.

• Advancement of the 3D stack provides new opportunities to save power based on geometrics. With traditional packaging technology, the power of memory-CPU communication will be approaching 2 pJ/bit, while the 3D stack can reduce this power to 1–20 fJ/bit relying on on-chip interconnects, which will benefit high-performance processors with more on-chip DRAM capacity.

• Stacking of multi-layer dies can better resist radiation, cosmic radioactive rays, radioactive packaging materials, and electromagnetic interference, providing the potential to improve reliability.

Currently, the key problem confronting 3D stack technology is heat emission and is therefore an issue of cooling.

5.2 Optical interconnect

New optical communication technologies, such as dense wavelength division multiplexing and multi-level encoding, etc., greatly increase the number of data channels in a single optical waveguide. Optical interconnect technologies have been widely used in rack-to-rack or wider interconnects, whereas in on-chip interconnects and board-level interconnects, electrical interconnection is still dominant. Since electrical interconnect is confronted more and more by serious power, latency and bandwidth issues, the optical interconnect solution is moving closer to processor cores, such as on-chip optical interconnects and board-level optical interconnects, providing a challenging alternative choice over electrical interconnects.

Compared to the electrical interconnect, the optical interconnect has following benefits [20,21]:

• High-bandwidth, high power efficiency processor-memory interconnect: according to research by MIT and Berkeley, optical interconnects can provide 2 TB/s bandwidth, 250 fJ/bit power. In a system consisting of 16 interconnected DRAM modules and 256 processor cores, an optical/electrical switch can improve the aggregate bandwidth by about 8-10x and power efficiency by about one order of magnitude, compared to an optimal electrical interconnect.

• Due to low degradation rate, an optical interconnect can provide better "bandwidth × length" performance.

• On-chip optical interconnects based on a low-area, low-power front-end CMOS transmitter and CMOS receiver provides the potential for improving interconnect power efficiency to 1 fJ/bit, exceeding 100 fJ/bit in an electrical interconnect.

• High resiliency, immune to crosstalk and electromagnetic interference.

The biggest opportunity provided by optical interconnects is that the performance of interconnects can catch up with the pace of total system performance. The technological maturity of the optical interconnects is the critical point of exascale systems. Currently, the main problems with optical interconnects are cost, high-latency and requirement for high-power transceivers.

5.3 Novel memory technologies

Compared to traditional DRAM, each of the current novel memory technologies, such as NAND Flash, phase change memory (PCRAM), semiconductor oxide nitride

oxide semiconductor (SONOS), and magnetic random access memory (MRAM), etc, have their own advantageous characteristics. MRAM has the largest cell size, but has good persistence performance. A smaller cell size, with good persistence performance and reliability make PCRAM and NAND Flash challenging alternatives to DRAM [18]. Novel memory technologies will influence the memory level of exascale systems in the following aspects:

• DRAM will still be the main choice of main memory in the future, when taking all factors including area, access speed, power, fault tolerance and persistence into account. However, hybrid cache structures composed of small-sized MRAM or DRAM, and fast-accessed SRAM maybe the direction for the cache in future, this can provide higher capacity and a new tradeoff between memory access latency and power consumption. Since future computer systems will be power constrained, this trade off is very important.

• A new level consisting of non-volatile storage can be added between the DRAM level and the disk level to mitigate the performance gap between DRAM and disk, and reduce recovery time from creating and storing system checkpoints.

6 Some implications

The development of applications and design patterns, together with the advancements in network and micro-electronic technology, will promote innovations in super-computer architecture and service patterns, which suggests:

• Large scale network computing based on widespread broadband internet is assumed to become more prominent in supercomputing. With the evolution of scientific computing, complex correlated multi-scale, multi-model applications will be mainstream in scientific computing, which exhibit computing models with intrinsic heterogeneous features, consisting of algorithms of several different types that can be efficiently implemented in various computing environments. Large scale network computing consisting of polymorphous computing resources can integrate the heterogeneity of computing models with the intrinsic distribution of data storage, and support data processing on appropriate machines. It can also ensure the execution of operations occurs on the machines nearest to the target data. This provides a favorable opportunity to fully exploit the potentials of process technologies, reduce communication overheads, fully exploit locality, reduce power consumption, improve reliability, and improve programming productivity. It will exhibit the supercomputing capability with this new form, deeply combining network and computing technology.

• Large scale data processing is becoming the mainstream of HPC in the cyber era and data deluge speeds the paradigm shift of supercomputing toward data-centric computing. Data deluge on cyberspace determines that the processing and analysis of enormous numbers of complex data sets are becoming more and more dominant in affecting computation requirements. Different from traditional scientific computing algorithms, the processing and analysis of many enormous complex data sets requires supercomputing systems to provide efficient, high bandwidth interconnects with efficient memory access, flexible I/O ports, efficient schemes to mitigate memory access latency and interconnect latency, efficient global share file systems for mass data and files, and customized parallel programming models and programming languages, on demand.

• More Intelligent components are necessary for future supercomputers. Currently, microprocessors have been widely used in many chips, components and devices, and are increasingly prevalent. Embedded microprocessor technologies, smart chips and components, are becoming the main impetus for advancements in chips. Intelligent chips assume the potential to improve performance and capability based on their flexibility and capability to easily realize complex logical functions. Intelligent components will benefit supercomputer system scalability, their adaptability to different applications, and help to control their complexity, and enhance their robustness.

Acknowledgements This work was sponsored by the National High Technology Research and Development Program of China (863 program) (2010AA01A141).

References

1. Gerber R A, Wasserman H, eds. Large Scale Computing and Storage Requirements for Biological and Environmental Research. http://www.science.doe.gov/ascr/ProgramDocuments/Docs/BER-NERSCReport.pdf
2. Simon H, Zacharia T, Stevens R. Modeling and Simulation at the Exascale for Energy and the Environment (E3). DOE ASCR Program Technical Report, 2008, http://www.sc.doe.gov/ascr/ProgramDocuments/Docs/TownHall.pdf
3. Finck P, Keyes D, Stevens R, et al. Report on the Workshop on

Simulation and Modeling for Advanced Nuclear Energy Systems. 2006, http://www.sc.doe.gov/ascr/Misc/gnep06-final.pdf

4. McIlroy A, McRae G, Sick V, et al. Basic research needs for clean and efficient combustion of 21st century transportation fuels. 2006, http://www.sc.doe.gov/bes/reports/files/CTF_rpt.pdf
5. Schreck S, Lundquist J, Shaw W, et al. U.S. department of energy workshop report: research needs for wind resource characterization (national renewable energy laboratory technical report NREL/TP-500–43521). 2008, http://www.nrel.gov/docs/fy08osti/43521.pdf
6. Rosner R, Zika M. Multiphysics simulation problems. Scientific grand challenges in national security: the role of computing at the extreme Scale. In: Proceedings of DOE NNSA-ASCR Workshop, October, Washington D. C. USA, 2009
7. Gregurick S, Drell D, Chalk C, et al. Scientific grand challenges: opportunities in biology at the extreme scale of computing. BER/ Biology Workshop Report, August, Chicago, USA, 2009, http:// www.er.doe.gov/ascr/ProgramDocuments/Docs/BiologyReport. pdf
8. Dongarra J, Beckman P, Moore T, et al. International exascale software project roadmap (draft 0.93). 2009, http://www.exascale. org/mediawiki/images/a/a1/Iesp-roadmap-draft-0.93-complete.pdf
9. Berman F. Developing cyberinfrastructure for data-oriented science and engineering. In: Proceedings of Next Steps in Using Combustion Cyberinfrastructure, San Diego, USA, 2007
10. Gaston D, Hansen G, Kadioglu S, et al. Parallel multiplysics algorithms and software for computational nuclear engineering. In: Proceedings of SciDAC Workshop 2009, June, San Diego, USA, 2009
11. Ghemawat S, Gobioff H, Leung S T. The Google file system. In: Proceedings of 19th ACM Symposium on Operating Systems Principles, Lake George, NY, 2003, 29–43
12. Dean J, Ghemawat S. MapReduce: Simplified data processing on large clusters. In: Proceedings of 6th International Symposium on Operating System Design and Implementation, San Francisco, USA, December, 2004, 137–150
13. Gantz J F, Chute C, Manfrediz A, Minton S, Reinsel D, Schlichting W, Toncheva A. The diverse and exploding digital universe. An International Data Cooperation White Paper. 2008, http://www. emc.com/collateral/analystreports/diverse-exploding-digital-universe.pdfMarch
14. Cisco. http://www.cisco.com/cisco crs-3 router
15. Bechtolsheim A. The road from peta to exaflop. In: Proceeding of 2009 International Supercomputing Conference, June, 2009
16. Kogge P, Bergman K, Borkar S, et al. ExaScale computing study: technology challenges in achieving exascale Systems. DARPA IPTO, September, 2008
17. Amarasinghe S, Campbell D, Carlson W, et al. ExaScale software study: software challenges in extreme scale systems. DARPA IPTO, September, 2009
18. Vantrease D, Schreiber R, Monchiero M, et al. Corona: system implications of emerging nanophotonic technology. In: Proceedings of 35th ACM/IEEE International Conference on Computer Architecture, Beijing, China, 2008, 153–164
19. Batten C, Joshi A, Orcutt J, et al. Building Manycore processor-to-DRAM networks with monolithic silicon photonics. In: Proceedings of 16th IEEE Symposium of High Performance Interconnects, 2008, 21–30
20. Miller D. Device requirements for optical interconnects to silicon chips. Proceedings of the IEEE, 2009, 97(7): 1166–1185
21. Parker M, Scott S. The impact of optics on hpc system interconnects. In: Proceedings of the 17th IEEE Symposium on High Performance Interconnects, 2009, 143–147