1

Petascale Computing Trends in Europe

MINORU NOMURA
Information and Communications Research Unit

1 Introduction

"High Performance Computing (HPC)" means the calculations with numerous high-end computational complexities used in the simulation of natural phenomena such as weather and climate, astrophysics and plasma analysis, and life-science applications by using supercomputers and grids. Supercomputers are Ultra High Performance Computers with diverse architectures that are adapted according to their utilization in performing highly computationally-intensive calculation tasks in science and technology. Various systems exist that have their performances ranked from high to low. The generic term for such supercomputers is called as 'HPC System' in this report. A Grid is an ecosystem with a large number of computational assets such as computers, storage devices, visualization systems, large-scale experimental observation devices, and data sources distributed throughout a network in which the component units of the system are utilized as a single virtual computer.

According to the TOP500 list,[1] which ranks the performances of the world's top 500 high performance supercomputer systems (HPC systems), Europe is seen to be emerging in its strive for ownership of HPC systems following behind the overwhelming ascendancy of the United States. Recently, further changes in responding to HPC system have been noticed in Europe. These shifts are focused on the government backed research and development of petascale computing (petaflop/s*1 performance supercomputer based HPC) carried out in the United States and Japan. Active deployment of the fastest HPC systems capable of handling large-scale, high-precision, high-speed simulations that are vital for advancing science and technology has also become necessary in Europe. Europe considers that acquiring and utilization of HPC system with petaflop/s performance, hereafter called the 'European Supercomputer System', is necessary and the e-Infrastructure programme of the Seventh Framework Programme for European Research and Technological Development (FP7), active since January 2007, has ranked 'Supercomputer System Deployment' as a new category in addition to the grid infrastructure intensified in FP6.

Relative to this development, the 'High Performance Computing in Europe Taskforce (HET: HPC in Europe Taskforce)' was inaugurated in June 2006. HET is comprised of HPC systems specialists from 11 European countries. HET was set up to propose strategies and actions in order to build a HPC ecosystem (integrated HPC systems) that is sustainable within Europe considering the highest performance HPC systems, infrastructures available in each European country, software developments and the requirements of capacity development in computational science. In January 2007, HET delivered the results of this research.

This article introduces these HET proposals. To begin with, the current situation regarding HPC within Europe is outlined in chapter 2. The proposals submitted by HET are shown in chapter 3; whereas the points that are worth particularly mentioning are described in chapter 4.

2 European high-performance computing (HPC)

2-1 European-owned HPC systems in TOP500 list

The TOP500 list is a list ranking computer power based on the LINPACK Benchmark.*2 Although it does not necessarily reflect the performance of

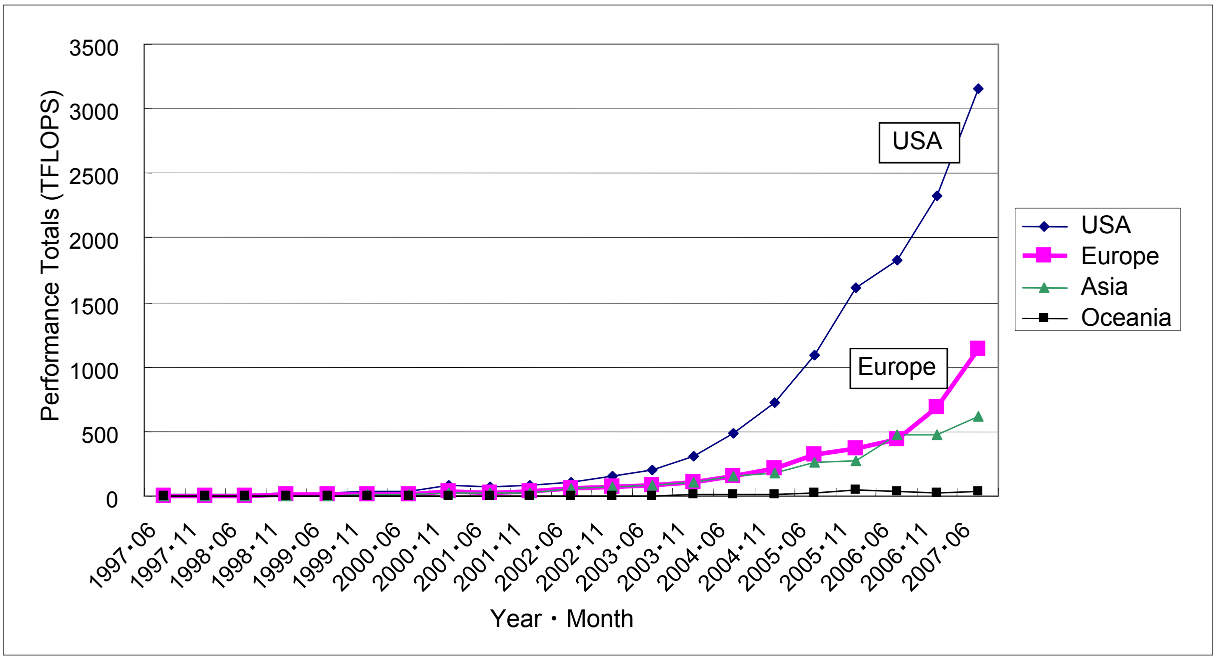

Figure 1 : Transition in performance values for each continent (TOP500 list systems)

Prepared by the STFC based on TOP500 list

a system in its actual operating environment, it has relevance in rating the world's most powerful publicly-known HPC systems. The TOP500 list is published each year in June and November. The situation regarding European HPC systems taken from the TOP500 list published in June 2007 is outlined below.

Figure 1 shows the transition in the LINPACK Benchmark performance values (Rmax) of HPC systems owned by each continent from June 1997 to June 2007. The values are shown as total performance values of all systems ranked in the TOP500 for each continent. From the chart it is clearly evident that Europe is emerging in its strive for ownership of HPC systems following behind the overwhelming predominance of the United States.

Figure 2 shows the performance values of HPC systems owned by each country in Europe with examples for the UK, Germany, France, and Spain where performance growth is particularly noticeable.

Looking at the numbers of systems in the TOP500 list, the UK has 42, Germany 24, France 13, and Spain 6. For reference, the number of systems in Japan is 23, whereas the USA owns 281.

The HPC supercomputer sites listed in each country are UK: the Atomic Weapons Establishment (ranked 24 in the TOP500 list), University of Reading (ranked 36); Germany: Leibniz Rechenzentrum (ranked 10), Forschungszentrum Juelich (FZJ) (ranked 18); France: the Commissariat a I'Energie Atomique (CEA) (ranked 12 and 22); and Spain: the Barcelona Supercomputing Center (BSC) (ranked 9). In addition to these countries, the numbers of HPC systems are increasing in the Netherlands, Switzerland, and Finland.

UK's HECToR (High-End Computing Terascale Resource) is an example of a recently-introduced, large-scale HPC system that is attracting attention. Designed for use by UK's university researchers, it is installed in Edinburgh University's Advanced Computing Facility (ACF). The first-phase system that has top speed of approximately 60TFLOPS (teraflop/s) was opened in October 2007. The second-phase system scheduled for October 2009 will have a top speed of 250TFLOPS. Thereafter, a third-phase with an unspecified top speed is scheduled for 2011.[2]

2-2 European predominance

European HPC systems have their hardware procured almost entirely from overseas. However, software development and HPC system application technologies in Europe are highly sophisticated. Europe's technological developments include contributions to Basic Linear Algebra Subprograms

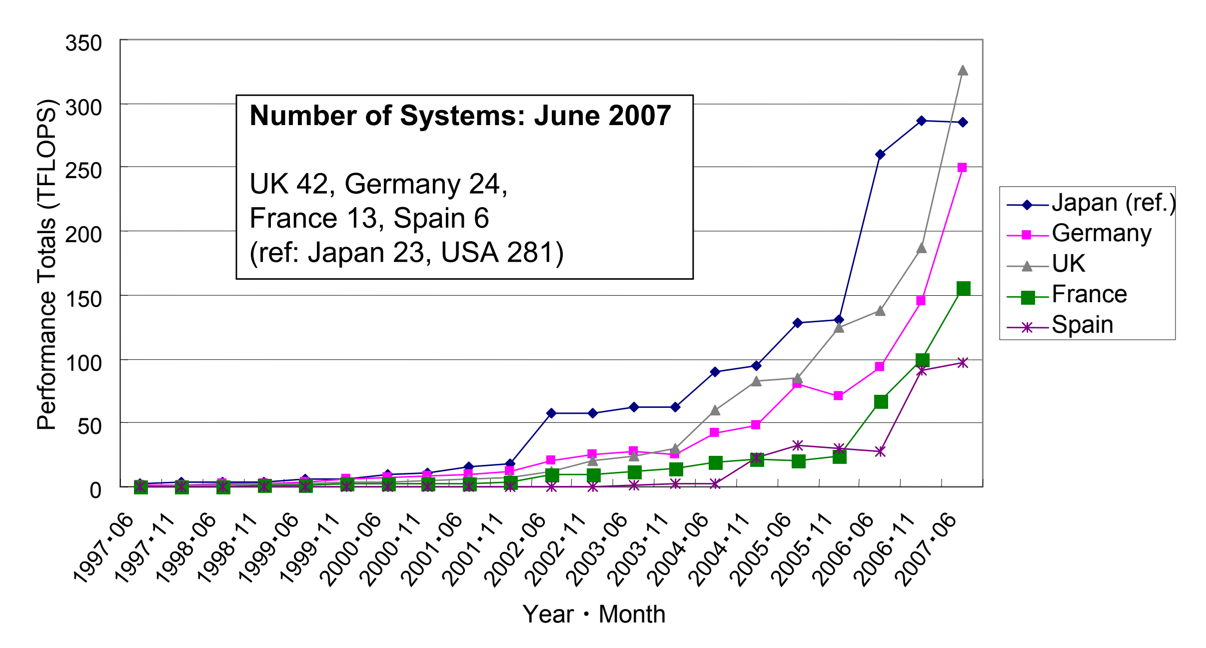

Figure 2 : Transition in performance values for each European country (TOP500 list systems)

Prepared by the STFC based on TOP500 list*3

(BLAS),*4 the development of Numerical Algorithms Group (NAG) library,*5 Computational Fluid Dynamics (CFD) programs, Finite Element Method (FEM) programs, computer chemical packages and statistical packages.

One aspect of this European application predominance is the vigorous progress made in the efficient application of grid-network system of HPC distributed in each country.[3-6]

2-3 Positioning of HPC within R&D frameworks

The e-Infrastructure programme of the Seventh Framework Programme for European Research and Technological Development (FP7) has added ‘Supercomputer System Deployment’ as a new category in addition to the grid infrastructure intensified in FP6. Such a move is possibly an attempt to counter government-backed petascale computing research and development in the USA and Japan. In Europe, activities for ownership and use of supercomputer systems capable of petaflop/s performance have also become essential. In the FP7 programme, Geant2*6 and the grid infrastructure are just to be upgraded. However, new elements such as the Supercomputer, Repository (storage facilities for data, information, software, etc.) and Data Infrastructure have been introduced.[7]

2-4 Transition to petascale computing

Figure 3 shows the progress of the European taskforce’s research on transferring to petascale computing. From August 2005 to April 2006, an International Scientific Panel, put together by the cooperation of scientists from Finland, France, Germany, Italy, the Netherlands, Spain, and the UK, investigated the scientific demands of petascale computing. This study on these scientific demands was launched by identifying a strategic role of high-end HPC (‘Leadership-class supercomputers’) in Europe’s scientific and economic development.

HET was established at a meeting sponsored by e-IRG*7 in June 2006 for the purpose of strengthening cooperation for delivering competitive high-level resources to Europe’s computational science community. Consisting of 11 European countries, it is represented by 22 members from Austria: 2 persons, Finland: 2 persons, France: 2 persons, Germany: 3 persons, Ireland: 2 persons, Italy: 2 persons, the Netherlands: 2 persons, Spain: 2 persons, Sweden: 1 person, Switzerland: 1 person, and the UK: 3 persons. In January 2007, HET presented the findings of its investigation.

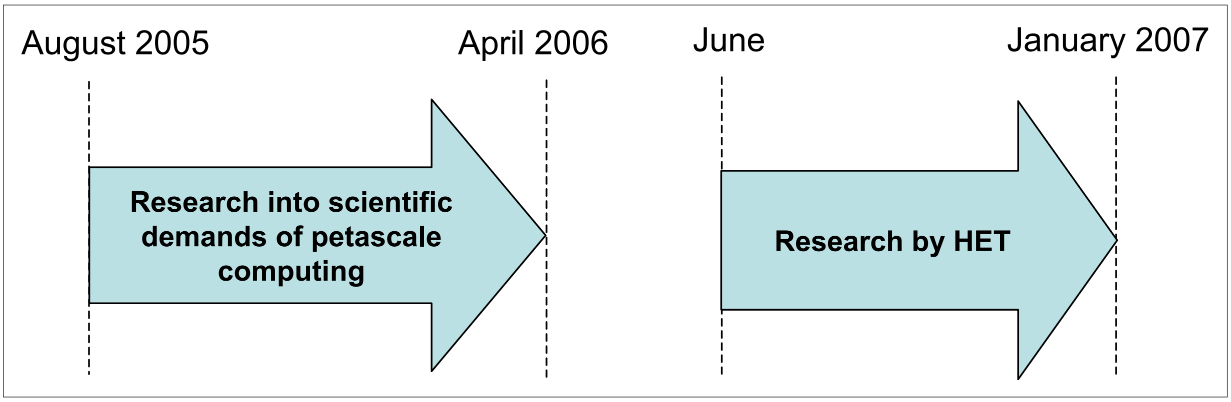

Figure 3 : 'High Performance Computing in Europe Taskforce (HET)' research progress

Prepared by the STFC based on Reference[10]

Table 1 : Estimated Costs of European High-Performance Computing Service

Application		Timescale	Estimated costs
Preparation phase		2006 – 2007	0.16 – 1.6 billion JPY
Medium-level infrastructure	Several medium size installation (5 – 10)	Starting 2007 – 2008	8 – 16 billion JPY every 2 years
High-end (capability) infrastructure	Several installations, where an installation can consist of two different architectures, placed in different locations.	Starting 2008 – 2009	16 – 32 billion JPY every 2 – 3 years
Maintenance and upgrades			8 – 16 billion JPY annually
Support projects such as software development, optimization, and training, etc.			4.8 – 8 billionJPY annually

(Estimated costs: 1 Euro converted to 160 JPY)

Prepared by the STFC based on Reference[8]

2-5 Positioning within the European Strategy Forum on Research Infrastructures (ESFRI) Roadmap[8]

The 'European Strategy Forum on Research Infrastructures (ESFRI)' has prepared a roadmap that will have a great influence on European sciences. The roadmap has reviewed and narrowed down 200 initially proposed projects to 35 projects in seven scientific disciplines (social sciences and humanities; environmental science; energy; biomedical and life sciences; material science; astronomy, astrophysics, nuclear and particle physics; computation and data treatment). Due to the efforts by HET, the 'European High-Performance Computing Service' has been incorporated as one of these projects in the area of computation and data treatment. The estimated costs of this service are presented in Table 1. These cost estimates include all HPC services such as premises for demanding systems and expertise needed for the optimal operation of the HPC systems.

3 "High Performance Computing in Europe Taskforce (HET)" proposals

The proposals presented by HET to the European Commission (EC) in January 2007 are presented in this chapter. Details of these proposals are extracted from references.[9-13]

3-1 Scientific demands for petascale computing

The report[13] created from August 2005 to April 2006 by the International Scientific Panel indicated in section 2-4 is discussed here.

In this report, the scientific demands relative

Table 2 : Scientific fields and applications

Scientific fields	Applications
(1) Weather, climatology, earth sciences	Climate change, oceanography and marine forecasting, meteorology, hydrology and air quality, earth sciences
(2) Astrophysics, high energy physics, plasma physics	Astrophysics, elementary particle physics, and plasma physics
(3) Material science, chemistry, nanoscience	Understanding of complex materials and complex chemistry, and nanoscience
(4) Life sciences	Systems biology, chromatin dynamics, large-scale protein dynamics, protein association and aggregation, supramolecular systems, and medicine
(5) Engineering	Complete helicopter simulation, biomedical flows, gas turbines & internal combustion engines, forest fires, green aircraft, and virtual power plant

Source: Reference[13]

to 'Petascale Computing in Europe' covering the period from 2010 to 2020 are stated. These scientific demands form the basis on which these HET proposals were formulated. The report is put together by five steering committees representing weather, climatology, and earth sciences; astrophysics, high energy physics, and plasma physics; material science, chemistry, and nanoscience; life sciences; and engineering disciplines. Table 2 shows the applications relevant to these five fields. The scientific challenges and potential outcomes of each field are indicated in the Appendices listed at the end of this report.

3-2 Role and scope of HPC in Europe Taskforce (HET)

The HPC in Europe Taskforce (HET) has stated that the lack of powerful facilities within Europe with the capacity to match the computing performance of Japan and the USA represents a significant loss in European computational science competitiveness. HET has attached importance to leadership-class resources that have petaflop/s computing capability and the methods for their effective utilization. HET states that merely having the necessary hardware is not enough to solve the kinds of scientific challenges shown in Table 2; but it will also require software and scalable algorithms adapted to petaflop/s performance computing systems and scientists that have the necessary skills sets. In addition, the integration with national computing centers shall also be required for executing the bulk of processes for which maximum performance computing is not needed. For these reasons, the feasibility of a sustainable, integrated pan-European HPC system (called the HPC ecosystem) was also studied.

3-3 Integration of HPC systems (HPC ecosystem)

(1) Capability and capacity classification

In reference 13, computing by supercomputers is split into the two categories described hereafter. "The requirements for scientific computing can be broadly divided into the categories of capacity and capability. Capacity computing implies high throughput of a large number of program runs of small or medium size on different data sets.

This style of computing can be dealt with in a cost effective manner, by providing adequate computing capacity at the regional or national level, accessed by grid-based computing methods. Capability computing is technologically far more challenging, and is characterized by simulations using a large number of processors in a cooperative mode. Often requiring large amounts of memory and the ability to process extensive data sets, such simulations are dependent on a high-bandwidth, low-latency communication fabric that connects the individual processors. These capabilities imply a single parallel computer that no loosely coupled distributed computing infrastructure can provide. The International Scientific Panel is convinced that the establishment of Europe-wide supercomputer capability is critical for this class of problem."

In reference 9, the difference between capability computing and grid computing is explained by the following example. "Table 3 indicates the broadness of latencies of different components in a grid. Latency is critically dependent on the distance that the signals will have to travel. It explains the very reason why supercomputer (capability computing) performance will never be attainable through grid computing, where by definition systems are included in the grid are spatially (far) apart."

(2) Performance pyramid

The HPC in Europe Taskforce (HET) have indicated the key elements of the HPC ecosystem (HPC integration) by a Performance Pyramid as shown in Figure 4. At the top of the pyramid (Tier 0), they have proposed the construction of a small number of high-end performance HPC systems (referred to here as the 'European Supercomputer Systems') which would be funded through national sources with additional European level funding. HET has ranked systems that perform capability computing as top-of-the-pyramid resources.

The middle layer (Tier 1) represents a number of national and regional level HPC systems, which should have a high enough performance to be capable of running most of the computing load below petaflop and in addition serving as a development platform for the most scalable code directed to the petaflop systems.

The bottom pyramid layer (Tier 2) represents the local level of activity including the development of a strong competence base of scientists in multiple European countries, the renewal and development of new skills through local education and increasing the visibility of scientific computing so to attract new users.

Although different capacity HPC systems exist in different layers, the grid-supported interconnections between layers are included throughout the entire pyramid. This type of service, in addition to the relative computing power, requires efficient storage systems, networks, middleware, and scalable software, etc.

(3) Significance of human factors

In order to support this type of HPC ecosystem, it is necessary to focus on the competence development and other human aspects at various levels. For example, growing investment is required in the following disciplines.

- Code development and optimization skills in all levels of the pyramid
- Research oriented training activities
- Access to expertise and scientific/technical support

Table 3 : Overview of indicative system and network latencies in communication process

Communication system	Latency (nanoseconds)
Supercomputer internal networks HW latency	3 – 5
Supercomputer internal networks SW latency	800 – 2,000
Commercial internal network SW latency	2,000 – 50,000
Speed-of-light in glass fiber(HW) latency per km	10,000
Long distance protocol (SW) latencies	> 1,000,000

(HW:hardware, SW:software) Figures are best estimates, the lower the figure the better.

Source: Reference[9]

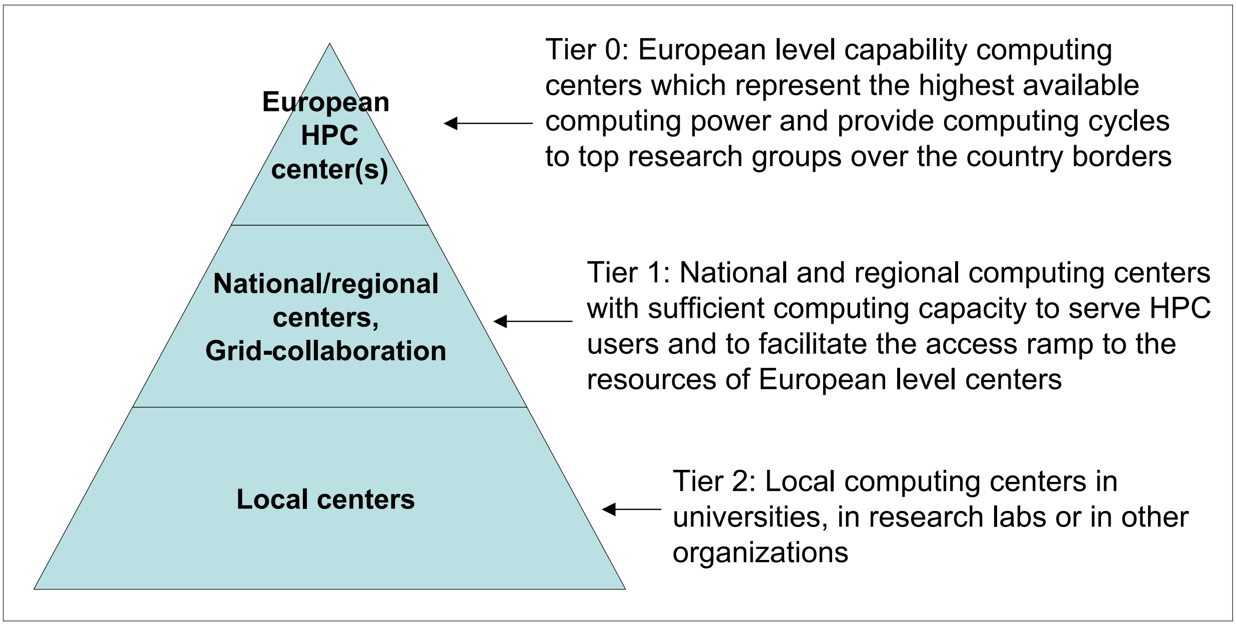

Figure 4 : Performance pyramid

Prepared by the STFC based on Reference[9]

- Access to deliverables such as code libraries
- Collaborative tools promoting sharing of information between different research groups, both within the same discipline and multi-scientific fields

3-4 *Integration with existing infrastructures*

It is essential that all HPC systems in Figure 4 are integrated in the European Research Grid in the same way as other resources regardless of whether they exist in the top or bottom layer of the performance pyramid. The European Research Grid is a highly integrated infrastructure in which every e-resource is integrated at the hardware, software, and communication level. High operational efficiency is achievable through tight integration and optimal utilization of distributed resources and skills sets in each European country. The relationship between HPC systems and the grid is described as follows.

(1) To run all relevant Research Computing to be conducted at the highest competitive level in Europe, supercomputers must be part of any European Research Grid infrastructure.

(2) Grid developments shall support the integration of all e-resources. A scientific research activity is only rarely dependent on just a single IT component, including supercomputing: there may be data pre-processing, a real-time component (with a sensor grid for example), visualization both interactively or as post-processing. This is only achievable by making supercomputing a part of the whole European Research Grid infrastructure. Also, the links between different layers of the computing resources pyramid should be established through grid middleware.

(3) Even at the supercomputing level there is no "one size fits all" for all applications, which means that there should be a flexible environment from which the best resources can be selected. This environment is suitably done within a grid working environment.

Grid integration between HPC systems is currently being achieved by the Distributed European Infrastructure for Supercomputing Applications (DEISA[*8]) project. The grid environment is a means of guaranteeing that all resources within the pyramid can be transparently seen and can be optimally deployed to serve their objective(s) in a single environment. Therefore, grid computing and supercomputing form a complementary relationship.

3-5 *Software development initiatives*

HET have made the following recommendations for addressing the key issues in developing software that allows exploiting the potential of resources having petaflop/s computing performance.

- Top-of-the-pyramid hardware alone will not be sufficient in delivering petascale performance for solving any given scientific problem.
- In order to run such large-scale (and potentially heterogeneous) machines, new operating systems will be required to address the problems related to scheduling, address space management, communication, fault tolerance, scalability, and reconfiguration, etc.
- New program models, compiler techniques, and runtime systems are required in order to process parallel tasks beyond $10.^6$
- Mathematical models of most scientific problems need to be improved in order to scale to the level provided by petaflop machines.
- All these measures are essential for utilizing petascale hardware, and Europe can claim to have a distinct advantage in such software related developments.

3-6 High efficiency, high reliability infrastructure for large-volume data storage

Increasing emphasis on permanent and persistent data repositories as a part of the HPC ecosystem have been recommended for the following reasons.

- The amount of stored scientific data has been doubling every year for the past several decades without indication that this exponential growth will end soon.
- Due to the massive amounts of data involved (many petabytes, approaching the exabyte range), large scale replication is impossible. On the other hand, most of these data cannot be reproduced once they are lost; therefore, systematic measures against failures(including disasters) are necessary.
- To address problems of intellectual property, ownership, and privacy, etc., elaborate security measures must be implemented.

3-7 Development of competent skills sets in computational science

As the computational science may substantially change the way science is performed in basically all fields, HET has proposed that massive educational efforts have to be initiated in order to actively use this new paradigm. In addition, education and training activities which focus on enabling more efficient and higher quality use of the top-end facilities in the long run require enhancement.

In specific terms, activities aimed at computer based modeling, scalable code development, integration work for the existing national infrastructures and training for competence development in computational sciences are cited.

3-8 Funding and utilization models

Building an internationally competitive European HPC infrastructure with a small number of high-end HPC centers requires continuous funding of both hardware procurement and support structures. In describing the design goals for the funding and utilization of Tier 0 European HPC center in Figure 4, HET has proposed the following funding and utilization models for the construction of these.

(1) General goals and design objectives

(a) Regarding utilization, the following conditions are reported to be needed to make the best possible usage of these resources.

- Resources such as computing time, storage, support, etc., should be allocated only to projects that require such capability computing resources and should not be assigned to demands that can be fulfilled by sub-national scale systems.
- Allocation of resources should strictly follow scientific criteria.
- Small-scale partitioning of the Tier 0 resources would compromise capability computing and should therefore be avoided.
- The development of world-class support systems and efficient software for capability computing is a fundamental prerequisite for effective resource utilization.

(b) Regarding funding, in order to accommodate the interests of existing and prospective funding agencies the following matters are proposed.

- Scientists of those EU member states and associated states which provide Tier 0 funding should get access to all European Tier 0 systems according to their countries´ total contribution and that of the EC.
- Scientists of EU member states and associated states not providing Tier 0 funding should get access to all the Tier 0 systems with respect to

the contribution of the EC.

- In accordance with the DEISA project model, the tight integration of top layer (Tier 0) centers with existing national and international middle layer (Tier 1) centers is to be funded.
- While the major part of available EC funding is required for the top-layer (Tier 0) resources, a substantial portion of the EC funding must be reserved for the development of capability computing oriented software and community oriented support structures at the Tier 0 centers and national Tier 1 centers.
- The establishment of pan-European science community should be supported.
- In addition to the core investments by the EC and contributing EU members and associated states, refunds obtained from selling computer time and services to industry as well as science should be considered.

In addition, the funding objectives should in particular include the very strong interest of the EC to include a pre-commercial procurement phase in order to open up European pathways to petaflop prototype systems. Within the pre-commercial procurement phase of about two years, the transfer of part of the R&D activities and production of non-European vendors to Europe and the reflection of Europe's needs in major non-European vendors are to be accelerated.

(2) Mutual utilization model

The introduction of a Mutual Utilisation Model (MUM) has been proposed. The guiding principle of MUM is the establishment of coordinated procurements by two or three EU member states and associated states, following each other within an interval of about 1 to 2 years. In addition to this funding, the EC contributes with a significant share to complement the national funding for every Tier 0 system. Other than by adopting this kind of procurement 'spiral', there is no other way of guaranteeing a sustained world-class Leadership HPC infrastructure in Europe.

A major portion of the resources that corresponds to the contribution of the EC to the costs of the infrastructure will be allocated to projects that are open to all EU member states and associated states. The remaining resources are shared among projects of the states that have committed themselves to carry out the successive procurements. Accordingly, on each of the systems, all the latter EU member states and associated states should get a substantial and fair amount of resources following their contribution.

3-9 Peer review process

The requirements for the top-of-the-pyramid (Tier 0) resources have been defined by considering the principles of subsidiarity with existing national and thematic resources to ensure the most relevant contribution to science. Scientists using the top-of-the-pyramid resources are expected to make contributions in advancing science through numerical simulation. This calls for an evaluation process based on peer review to allocate computer resources. However, since the optimal utilization of highly parallel systems such as the petaflop computers is challenging, a two-level review system that evaluates both scientific excellence and efficiency of the application code is required.

HET have indicated the framework, periodic processes, and eligibility of applications. The following requirements have been cited as eligibility of applications.

- Scientific excellence
- Demonstrated need for Tier 0 level resources
- Proven feasibility of computation, data collection and result analysis. Code must have been validated on suitable equipment available to proposer on a national, thematic or local basis.
- In the case of industrial partners, further examination of benefits is warranted, as examination of property of results, confidentiality etc.

4 Points of special interest

(1) Necessity of capability computing

The necessity of capability computing is described throughout the HET report. The report concerning scientific demands created by the International Scientific Panel states that petascale computing would be a great benefit to the multiple world-class research groups that exist in Europe and it is vital to have immediate access to adequate resources that can reduce the time taken to solve problems (Time to Solution) in order to implement approved science and research at a competitive level.

This International Scientific Panel also states that the procurement of Leadership-class supercomputers is essential for European science to maintain a leading position in multiple domains, whereby without access to capability level computers European science and industry would lose its competitive edge. Areas that are under threat include: climate data usable in international negotiations, construction of ITER and adequate utilization of project results, and evolution in biophysics and systems biology for the understanding of complete cells.

(2) Utilizing the 'European Supercomputer System'

HET has ranked the 'European Supercomputer System' as the system which should be used for capability computing. This is explained in detail by referencing specific latencies (delays occurring in communications) in contrast to the capacity computing concept.

However, in terms of computer systems to be deployed, since it is extremely unlikely that a single design or design concept (architecture) of capability computing will be available to satisfy all requirements in capability computing at the highest level in every academic discipline, the demand for multiple systems with different architectures is emphasized.

Regarding utilization, HET proposes that software for the 'European Supercomputer System' should be used for research goals of highest scientific merit and greatest need, whereby the software to be used are selected by a peer review process.

(3) European competitiveness

HET states that it is necessary to conduct activities that focus on the entire HPC ecosystem or more specifically all layers of the performance pyramid, in order to narrow the gap between Europe and the USA or Japan. In doing so, the numerous important factors that give Europe its competitive edge will be enhanced. In specific terms, investments into scientific software development and code optimization, middleware development, data repositories, and efficient networking are cited. The important factors that have hitherto given Europe its competitive edge should be subsequently exploited.

Section 3-5 indicates the important action assignments to be undertaken in software development fields required in petascale computing. However, in the essential measures cited for petascale computing HET states that Europe maintains an advantage.

(4) Necessity of developing human resources

Section 3-7 indicates the necessity of developing computational science skills sets through enhanced education and training activities. In comparison to the acquisition of facilities, obtaining the human resources to manage the facilities is more difficult, therefore strong appeals for necessary investments are made. It mentions that training of human resources with the required skills sets could take several years; therefore, in order to attract such human resources, investments in building supporting environments shall also be needed.

(5) Responding to non-Europe's HPC vendors

In Europe, there are very few vendors that can supply HPC hardware and processors. Because Europe is challenged to accurately convey its HPC requirements to non-European vendors, measures such as promoting a shift of non-European R&D activities into Europe and incorporating Europe's demands into non-European vendors through a pre-commercial procurement phase are being considered.

(6) Science and technology fields chosen by scientific necessity

The American government's high performance computing policy also mentions the benefits of high performance computing to science and engineering.[14,15] By comparing the American proposals to the demands investigated by the European International Scientific Panel, Europe makes few references to 'national security' but raises detailed points about the many applications to life sciences in addition to including matters of deep interest in engineering fields such as helicopter simulation, forest fires, green aircraft, and virtual power plants, etc., which were not raised by the USA.

5 Conclusions

Areas of keen interest regarding petascale computing trends in Europe have been highlighted from the proposals presented by the HPC in Europe Taskforce (HET).

In order to advance science and technology, there is a universal need in Japan, America, and Europe to respond to changes towards large-scale, high-precision, and high-speed simulation. Presently, in America there are four projects (Department of Energy ASC programs, NLCF program, Department of Defense HPCS program, and the American National Science Foundation's Cyber Infrastructure program) and one project in Japan (development and adoption of state-of-the-art general-purpose computers of high performance) that are progressing in the research and development of capability computing systems. In Europe, the issue of constructing 'European Supercomputer System' facilities as capability computing systems is aimed at sharpening Europe's competitive edge to counter the advances made in the US and Japan. Owing to the actions of HET, the 'European High Performance Computing Service' was incorporated as one of 35 projects in the European Strategy Forum on Research Infrastructures (ESFRI) Roadmap which holds a great influence on European sciences. In addition, in April 2007, the memorandom of understanding for 'Partnership for Advanced Computing in Europe' (PRACE, or formerly referred to as 'PACE')[16] initiative was signed by major European countries. One of the key objectives of PRACE is to achieve a pan-European HPC service by 2009/2010. This implies that Europe has taken a significant first step towards the deployment and utilization of petascale computing. It would be considered worthwhile to follow closely the prospects of Europe where excellent achievements in application and software development technologies have been realized.

Acknowledgements

The author of this article would like to express his sincere gratitude to Mr. Hisashi Nakamura of the Research Organization for Information Science & Technology for his valuable opinions.

Glossary

*1 Flops is a unit indication of computer processing speed. A petaflop/s (or PFLOPS) computer performs 10^{15} FLoating point Operations Per Second.

*2 LINPACK (LINear equations software PACKage) Benchmarks are calculated using software libraries for performing simultaneous linear equations used primarily in calculating floating point numbers. The results of the Benchmark tests, indicated as numbers of floating point operations per second (flops), are listed for numerous computing systems from supercomputers and workstations to personal computers.

*3 Charts 1 and 2 show the performance totals of HPC systems listed in the TOP500. For example, the performances of the UK's 42 listed HPC systems are given as a total value. The performance (Rmax) of the individual highest-ranking HPC system listed in the TOP500 is 280TFLOPS (USA's DOA/NNSA/LLNL). Research and development into petascale computing is aiming at exceeding petaflops in this individual HPC performance.

*4 A de facto application programming interface for performing basic numerical calculations such as linear algebra operations.

*5 Wide-range science, engineering, and statistical application numerical software libraries developed by the NAG research and development team formed by world renowned mathematicians and statisticians.

*6 Geant2 is the seventh generation pan-European research and education network. It provides a state-of-the-art network service to 34 countries connected via the National Research and Education Networks (NRENs) covering the world's largest geographically networked area.

*7 e-Infrastructure Reflection Group: e-IRG is supporting the creation of a framework (political, technological and administrative) for the easy and cost-effective shared use of distributed electronic resources across Europe - particularly for grid computing, storage and networking.

*8 DEISA (Distributed European Infrastructure for Supercomputing Applications):

DEISA is a consortium of leading national supercomputing centers that currently deploys and operates a persistent, production quality, distributed supercomputing environment with continental scope.

Abbreviations

FP7: Seventh Framework Programme

DEISA: Distributed European Infrastructure for Supercomputing Applications

EGEE: Enabling Grid for E-sciencE

ESFRI: European Strategy Forum on Research Infrastructures

GEANT: Seventh generation of pan-European research and education network

HET: HPC in Europe Taskforce

HPC: High Performance Computing

The Challenges and Outcomes in Science and Engineering to be Addressed through Petascale HPC Provision

Area	Application	Science Challenges & Potential Outcomes
Weather, Climatology and Earth Sciences	Climate change	Quantify uncertainties on the degree of warming and the likely impacts by increasing the capability and complexity of 'whole earth system' models that represent in ever-increasing realism and detail the scenarios for our future climate.
	Oceanography and Marine Forecasting	Build the most efficient modelling and prediction systems to study, understand and predict ocean properties and variations at all scales, and develop economically relevant applications to inform policy and develop services for government and industry.
	Meteorology, Hydrology and Air Quality	Predict weather and flood events with high socio-economic and environmental impact within a few days. Understand and predict the quality of air at the earth's surface; development of advanced real-time forecasting systems for allowing early enough warning and practical mitigation in the case of pollution crisis.
	Earth Sciences	Challenges span a wide range of disciplines and have significant scientific and social implications, such as the mitigation of seismic hazards, treaty verification for nuclear weapons, and increased discovery of economically recoverable petroleum resources and monitoring of waste disposal. Increased computing capability will make it increasingly possible to address the issues of resolution, complexity, duration, confidence and certainty, and to resolve explicitly phenomena that were previously parameterized, and will lead to operational applications in other European centres, national centres and in industry.
Astrophysics, HEP* and Plasma Physics	Astrophysics	Deal with systems and structures which span a large range of different length and time scales; almost always non-linear coupled systems of ordinary and partial differential equations have to be integrated, in 3 spatial dimensions and explicitly in time, with rather complex material functions as input. Grand challenges range from the formation of stars and planets to questions concerning the origin and the evolution of the Universe as a whole. Evaluate the huge amount of data expected from future space experiments such as the European Planck Surveyor satellite.
	Elementary Particle Physics	Quantum field theories like QCD (quantum chromodynamics) are the topic of intense theoretical and experimental research by a large and truly international community involving large European centers like CERN and DESY. This research not only promise to yield a much deeper understanding of the standard model of elementary particles and the forces between them, as well as nuclear forces, but is also expected to discover hints for a yet unknown physics beyond the standard model.
	Plasma physics	The science and technology challenge raised by the construction of the magnetic confinement thermonuclear fusion reactor ITER calls for a major theory and modelling activity. Both the success of the experiment and its safety rely on such simulators. The quest to realize thermonuclear fusion by magnetically confining a high temperature plasma poses some of the computationally most challenging problems of nonlinear physics.

Area	Application	Science Challenges & Potential Outcomes
MATERIALS SCIENCE, CHEMISTRY AND NANOSCIENCE	Understanding Complex Materials	The determination of electronic and transport properties central to many devices in the electronic industry and hence progress the understanding of technologically relevant materials. Simulations of nucleation, growth, self-assembly and polymerization central to the design and performance of many diverse materials e.g., rubbers, paints, fuels, detergents, functional organic materials, cosmetics and food. Multiscale descriptions of the mechanical properties of materials to determine the relation between process, conditions of use and composition e.g., in nuclear energy production. Such simulations are central to the prediction of the lifetime of high performance materials in energy technology, such as high-efficiency gas-turbines
	Understanding Complex Chemistry	Catalysis is a major challenge in the chemistry of complex materials, with many applications in industrial chemistry. The knowledge of atmospheric chemistry is crucial for environmental prediction and protection (clean air). Improving the knowledge of chemical processing (from soft chemistry including polymers to the atomistic description of combustion) would improve the durability of chemicals. Supra molecular assemblies open new possibilities for the extraction of heavy elements from spent nuclear fuels. In biochemistry, a vast number of reactions taking place in the human body (for example) are not understood in any detail. A key step in the development of the clean fuels of the future requires the realistic treatment of supported catalytic nanoparticles.
	Nanoscience	The advance of faster information processing or the development of new generations of processors requires the shrinking of devices, which leads inevitably towards nanoelectronics. Moreover, many new devices, such as nanomotors can be envisioned, which will require simulation of mechanical properties at the nanolevel. Composite high performance materials in the fields e.g. adhesion and coatings will require an atomistic based description of nanorheology, nanofluidics and nanotribology. As an example the description of the complex magnetic and mechano-optical properties of nanodevices components is only feasible only on systems in the Petaflop/s range.
LIFE SCIENCES	Systems Biology	The use of increasingly sophisticated models to represent the entire behaviour of cells, tissues, and organs, or to evaluate degradation routes predicting the final excretion product of any drug in any organism. To position Europe in the next 4 years to host the first "in silico" cell.
	Chromatine Dynamics	The organization of DNA in nucleosomes largely modifies the accessibility of transcription factors recognition sites playing then a key role in the regulation of gene function. The understanding of nucleosome dynamics, positioning, phasing, formation and disruption or modifications induced by chemical modifications, or by changes in the environment will be crucial to understand the mechanism of gene regulation mediated by chromatine modelling.
	Large Scale Protein Dynamics	The study of large conformational changes in proteins. Major challenges appear in the simulation of protein missfolding, unfolding and refolding (a key element for the understanding of prionoriginated pathologies).
	Protein association and aggregation	One of the greatest challenges is the simulation of crowded "not in the cell" protein environments. To be able to represent "in silico" the formation of the different protein complexes associated with a signalling pathway opens the door to a better understanding of cellular function and to the generation of new drugs able to interfere in protein-protein interactions.
	Supramolecular Systems	The correct representation of protein machines is still out of range of European groups using current simulation protocols and computers. The challenge will be to analyze systematically how several of these machines work e.g., ribosome, topoisomerases, polymerases.
	Medicine	Genome sequencing, massive genotyping studies are providing massive volumes of information e.g. the simulation of the determinants triggering the development of multigenic-based diseases and the prediction of secondary effects related to bad metabolism of drugs in certain segments of population, or to the interaction of drugs with macromolecules others than their original targets.

Area	Application	Science Challenges & Potential Outcomes
ENGINEERING	Complete Helicopter Simulation	The European helicopter industry has a strong tradition of innovation in technology and design. Computational Fluid Dynamics (CFD) based simulations of aerodynamics, aeroacoustics and coupling with dynamics of rotorcraft already play a central role and will have to be improved further in the design loop.
	Biomedical Flows	Biomedical fluid mechanics can improve healthcare in many areas, with intensive research efforts in the field of the human circulatory system, the artificial heart or heart valve prostheses, the respiratory system with nose flow and the upper and lower airways, and the human balance system. Although experiments have significantly improved the understanding in the field, numerous questions, the answers of which need a high resolution of the flow field, of the surrounding tissue, or of their interactions, require a detailed numerical analysis of the biomedical problem.
	Gas Turbines & Internal Combustion Engines	Scientific challenges in gas turbines or piston engines are numerous. First, a large range of physical scales should be considered from fast chemical reaction characteristics (reaction zone thicknesses of about tens of millimetres, 10^{-6} s), pressure wave propagation (sound speed) up to burner scales (tens of centimetres, 10^{-2} s resident times) or system scales (metres for gas turbines).
	Forest Fires	The development of reliable numerical tools able to model and predict fire evolution is critically important in terms of safety and protection ("numerical fire simulator"), fire fighting and could help in real time disaster management. The social impact is very important and is concerned with land, buildings, human and animal life, agriculture, tourism and the economy.
	Green Aircraft	ACARE 2020 provides the politically agreed targets for an acceptable maximum impact of air traffic on people and the environment, while at the same time allowing the constantly increasing amount of air travel. The goals deal with a considerable reduction of exhaust gas and noise. Air traffic will increase by a factor of 3, accidents are expected to go down by 80%. Passenger expense should drop (50%) and flights become largely weather independent. The "Green Aircraft" is the answer of the airframe as well as engine manufacturing industry. However, it is only by a far more productive high quality numerical simulation and optimisation capability that such a challenging development will be possible. It will be indispensable to be able to compute the real aircraft in operation, including all relevant multi-disciplinary interaction.
	Virtual Power Plant	Safe production of high quality and cost effective energy is one of the major concerns of Utilities. Several challenges must be faced, amongst which are extending the lifespan of power plants to 60 years, guaranteeing the optimum fuel use and better managing waste. These challenges demand access to Petascale machines to perform advanced simulations along with a new generation of codes and simulation platforms.

Note: HEP: High-energy physics

Source: Reference[13]

References

[1] http://www.top500.org/
[2] http://www.hector.ac.uk/
[3] http://www.deisa.org/
[4] http://www.eu-egee.org/
[5] Kenichi Miura, Trends in Grid Computing, J. Plasma Fusion Res. Vol80, No. 9 (2004)
[6] http://www.geant2.net/server/show/conWebDoc.994
[7] Introduction to e-Infrastructure activity under FP7(Mario Campolargo 26th Feb. 2007)
[8] ESFRI European Roadmap for Research Infrastructures Report 2006 (p65)
[9] A sustainable High Performance Computing Ecosystem for Europe (http://www.hpcineuropetaskforce.eu/deliverables)
[10] Towards a Sustainable High-Performance Computing Ecosystem through Enabling Petaflops Computing in Europe (http://www.hpcineuropetaskforce.eu/deliverables)
[11] HET Funding and Utilization Model (http://www.hpcineuropetaskforce.eu/deliverables)
[12] HET Peer Review Process Proposal for Tier 0 applications (http://www.hpcineuropetaskforce.eu/deliverables)
[13] The Scientific Case for a European Super Computing Infrastructure
Petascale Computing in Europe (http://www.hpcineuropetaskforce.eu/deliverables)
[14] Federal Plan for High-End Computing
Report of the High-End Computing Revitalization Task Force(HECRTF) May10,2004
http://www.itrd.gov/pubs/2004_HECRTF/20040702_HECRTF.pdf

[15] Trends in High-End Computing in United States Government, Science & Technology Trends -Quarterly Review No.16 (2005)

[16] Pan-European supercomputer comes a step closer to reality (CORDIS News 2007.04.18)

Minoru NOMURA
Information and Communications Research Unit

Mr. Nomura was engaged in research and development of CAD for computer design purposes and business development of the high performance computing and ubiquitous regions before taking his current position. He is interested in science and technology trends in the information and communications fields, including super computers and LSI design technologies.

(Original Japanese version: published in October 2007)