Deliverable D2.3

Report on International Activities

CONTRACT NO EESI 261513
INSTRUMENT CSA (Support and Collaborative Action)
THEMATIC INFRASTRUCTURE
START DATE OF PROJECT: 1 June 2010 DURATION: 18 months

Name of lead contractor for this deliverable: NCF
Authors: Peter Michielse (NCF), Patrick Aerts (NCF)
Name of reviewers for this deliverable:

Abstract: This deliverable reports on progress in exascale projects worldwide. Its main source of information are the IESP meetings. This report will be extended after each IESP meeting.

Release N° 1

Due date of deliverable:	November 30, 2011
Publication date:	January 5, 2011

Release N° 2

Due date of deliverable:	November 30, 2011
Added sections:	Section 5
Publication date:	June 24, 2011

Release N° 3

Due date of deliverable:	November 30, 2011
Added sections:	Section 6
Publication date:	November 30, 2011

Project co-funded by the European Commission within the Seventh Framework Programme (FP7/2007-2013)		
Dissemination Level : PU		
PU	Public	X

Table of Contents

Glossary

Abbreviation / acronym	Description
COTS	Commodity-of-the-shelf
DoE	Department of Energy (US)
EC	European Commission
EESI	European Exascale Software Initiative (Europe)
ESC	Exascale Software Centre (US)
GENCI	Grand Equipement National de Calcul Intensif (France)
HPC	High Performance Computing
HPCI	High Performance Computing Infrastructure (Japan)
IDC	International Data Corporation
IESP	International Exascale Software Project
IPR	Intellectual Property Rights
JSC	Juelich Supercomputing Center
LRZ	Leibniz Rechnung Zentrum
NCF	National Computing Facilities Foundation (the Netherlands)
NNSA	National Nuclear Security Administration (US)
NSF	National Science Foundation (US)
SKA	Square Kilometre Array
WG	Working Group (in EESI)
X-stack	Exascale software stack

1. Executive summary

During the lifetime of the EESI project, there have been three IESP meetings: in Maui, HI, USA, in San Francisco, CA, USA and in Cologne, Germany. During these meetings, various topics have come back. For each of these topics, we briefly state the developments so far, bullet-wise.

- The US, China and Japan have clear intentions and roadmaps on building and installing 100+ Petaflop/s and Exaflop/s systems in the decade ahead. This is not (yet) the case for Europe, as will become clear from the EESI findings and recommendations;
- Co-design is viewed as a key technique to develop hardware, X-stack software and applications in an efficient (iterative) way. The window of opportunity for influencing hardware development, however, is relatively short and is 3-5 years before actual hardware availability. Skeleton applications have been viewed as enough for the co-design process, although new views on this may lead to a combination of kernels and applications itself. The original plans for Co-design centers in the USA have been reduced, at least temporarily, so that only a few Co-design centers have been able to make a start in 2011. Japan and China will adopt co-design as well, and, as in the USA, will focus on predefined application areas;
- The USA plans for establishment of an Exascale Software Center (ESC) have been delayed, but not abandoned. New plans through the DoE for US Congress are due in February 2012, which are currently being prepared by the community in the USA. Requests-for-Information have been issued to both US hardware and software vendors, which will be taken into account when drafting the plans;
- The EESI project has been able to address a huge amount of knowledge through engaging European experts in both scientific, hardware, software, libraries and software engineering fields. This has led to findings and recommendations to the EC, including cost estimates over the period 2012-2020, as presented in Cologne and during the final EESI conference in Barcelona in October 2011. Development of exascale hardware is viewed as a serious option for Europe as well;
- The EC has granted three exascale proposals each with around 8 M€ funding, to develop relevant exascale technology in the hardware, software and applications area. Two of the projects aim to develop prototype hardware, which aim to set the road to exascale, with relevant software components and application vehicles to test. The third project focuses on European-developed and important applications to apply software technology to prepare for eventual exascale execution;
- The involvement of hardware vendors in the development of exascale software is viewed as crucial, but also as difficult as a result of Intellectual Property Rights and Non-Disclosure Agreements between vendors and potential Co-design centers and/or the ESC. Initial models for addressing these difficulties have been established and will need to be further developed;
- International joint governance has been discussed to make sure that X-stack development is done in a coordinated way, without duplications. Further refinement of these models will need to be done, including the addressing questions on steering, delivery and maintenance of the software products. Initial inventories have been made by each of the regions (US, Europe, China, Japan) on which software components could be addressed in which region, with varying levels of commitment. Gap analysis will be needed to cover the full range of what is needed;
- Revolutionary approaches are worthwhile to invest in several areas – if possible in a controlled manner.

2. Introduction

It is widely recognized that High Performance Computing (HPC) will be increasingly important to address global scientific, societal and economic challenges. Although the projected evolution of hardware is a technological challenge in itself, more and more concern is expressed on the ability of scientific software to efficiently use the future hardware architectures. In order to address the expectedly increasing gap between projected exascale-class hardware and efficient software usage, a worldwide effort has been set up at the beginning of 2009, which is known as the International Exascale Software Project (IESP), [1].

With seed funding from key government partners in the United States, European Union and Japan, as well as supplemental contributions from some industry stakeholders, IESP has been formed around the following mission:

The guiding purpose of the IESP is to empower ultra-high resolution and data-intensive science and engineering research through the year 2020 by developing a plan for (1) a common, high-quality computational environment for petascale/exascale systems and (2) catalyzing, coordinating, and sustaining the effort of the international open source software community to create that environment as quickly as possible.

As has been extensively described in [1], there exist good reasons to conclude that such a plan is urgently needed. Firstly and foremost, the scope of the technical challenges for software infrastructure that the novel architectures and extreme scale of emerging systems bring with them is daunting. These challenges, which are already apparent on the leadership-class systems of the US National Science Foundation (NSF) and Department of Energy (DOE), as well as on systems in Europe and Asia, are more than sufficient to require the complete redesign and replacement of the operating systems, programming models, libraries, and tools on which high-end computing necessarily depends. Secondly, the complex web of interdependencies and side effects that exist among such software components means that making sweeping changes to this infrastructure will require a high degree of coordination and collaboration. Failure to identify critical holes or potential conflicts in the software environment, to spot opportunities for beneficial integration, or to adequately specify component requirements will tend to retard or disrupt everyone's progress, wasting time that can not be afforded to be lost. Since creating a software environment adapted for extreme-scale systems (e.g., NSF's Blue Waters) will require the collective effort of a broad community, this community must have good mechanisms for internal coordination. Third, it seems clear that the scope of the effort must be truly international. In terms of its rationale, scientists in nearly every field now depend on the software infrastructure of high-end computing to open up new areas of inquiry (e.g., the very small, very large, very hazardous, very complex), to dramatically increase their research productivity, and to amplify the social and economic impact of their work. It serves global scientific communities who need to work together on problems of global significance and leverage distributed resources in transnational configurations. In terms of feasibility, the dimensions of the task—totally redesigning and recreating, in the period of just a few years, the massive software foundation of computational science in order to meet the new realities of extreme-scale computing—are simply too large for any one country, or small consortium of countries, to undertake on its own.

During 2009 IESP has organized a series of three international workshops, one each in the United States, Europe, and Asia in order to work out a plan. This has led to a technology roadmap to address a critical challenge that now confronts modern science and is produced by a convergence of three factors:

1. the compelling science case to be made, in both fields of deep intellectual interest and fields of vital importance to humanity, for increasing usable computing power by orders of magnitude as quickly as possible;

2. the clear and widely recognized *inadequacy* of the current high end software infrastructure, in all its component areas, for supporting this essential escalation;
3. the near complete lack of planning and coordination in the global scientific software community in overcoming the formidable obstacles that stand in the way of replacing it.

This document is organized as follows. Section 3 will briefly document the four IESP workshops held before the start of the EESI project (three in 2009 and one in 2010). Section 4 will report on the first IESP workshop held during the actual EESI project (in October 2010 in Maui, Hawaii, USA), including a summary of the major findings and actions as a result of the workshops. Section 5 covers the IESP workshop as held in April 2011 in San Francisco, CA, USA. Section 6 covers the IESP workshop as held in October 2011 in Cologne, Germany.

This document has been updated after each IESP workshop, during the course of the EESI project.

3. Overview of IESP workshops until April 2010

3.1 Introduction

As has been documented in [1], IESP develops a plan for producing a new software infrastructure capable of supporting exascale applications. In order to do so, the following sequence of objectives needs to be addressed during the workshops:

1. *Make a thorough assessment of needs, issues and strategies*: A successful plan in this arena requires a thorough assessment of the technology drivers for future peta/exascale systems and of the short-term, medium-term, and long-term needs of applications that are expected to use them. The IESP workshops brought together a strong and broad-based contingent of experts in all areas of HPC software infrastructure, as well as representatives from application communities and vendors, to provide these assessments. Also, leverage has been taken of the substantial number of reports and other material on future science applications and HPC technology trends that different parts of the community have created in the past three years.

2. *Develop a coordinated software roadmap*: The results of the workshop attendees' analysis have been incorporated into a draft of a coordinated roadmap intended to help guide the open source scientific software infrastructure effort with better coordination and fewer missing components.

3. *Provide a framework for organizing the software research community*: Development of an organizational framework to enable the international software research community to work together to navigate the roadmap and reach the appointed destination—a common, high quality computational environment that can support extreme-scale science on extreme-scale systems. The framework will include elements such as initial working groups, outlines of a system of governance, alternative models for shared software development with common code repositories, and feasible schemes for selecting valuable software research and encouraging its translation into usable, production-quality software for application developers. This organization must also foster and help coordinate R&D efforts to address the emerging needs of users and application communities.

4. *Engage and coordinate with the vendor community in cross-cutting efforts*: To leverage resources and create a more capable software infrastructure for supporting exascale science, the IESP is committed to engaging and coordinating with vendors across all of its other objectives. Industry stakeholders have already made contributions to the workshops (i.e.. objectives 1 and 2 above) and IESP expects similar, if not greater participation, in the effort to create a model for cooperation as well as coordinated R&D programs for new exascale software technologies.

5. *Encourage and facilitate collaboration in education and training*: The magnitude of the changes in programming models and software infrastructure and tools brought about by the transition to peta/exascale architectures will produce tremendous challenges in the area of education and training. As it develops its model of community cooperation, the IESP plan must, therefore, also provide for cooperation in the production of education and training materials to be used in curricula, at workshops and on-line. This roadmap document, which focuses objectives 1 and 2 above, represents the main result of the first phase of the planning process. Although some work on tasks 3–5 has already begun, we plan to solicit, and expect to receive in the near future, further input on the roadmap from a much broader set of stakeholders in the computational science community. This version of the roadmap begins that process by including more extensive input from the science application community, international funding agencies, and vendor partners. The additional ideas and information we gather as the roadmap is disseminated are likely to produce changes that need to be

incorporated into future iterations of the document as plans for objectives 3–5 develop and cooperative research and development efforts begin to take shape.

The following sections will give a brief overview of the IESP workshops, held before the start of the EESI project on June 1, 2010. All IESP meetings so far have been by invitation only.

3.2 Workshop in Santa Fe, NM, USA

After an initial meeting during SC08 in Austin, the first workshop has been held on April 7, 8, 2009, in Santa Fe, New Mexico, USA. This first workshop brought together an initial set of representatives from various countries, all active in the area of HPC and software development for demanding scientific applications. Appendix A lists the attendees of all workshops so far.

During this initial meeting, time has been reserved for an introduction to many topics: HPC situations in the USA, Europe and Japan, science drivers, hardware feasibility and challenges, software bottlenecks and funding and organisational issues. It was agreed to develop a version of the IESP roadmap after the third workshop, if possible before SC09 in Portland. This roadmap should consider HPC architectures and hardware, software roadmap and computational science research strategies, but also approaches to funding, on how to collaborate with vendors which also have their own commercial interests. Initial ideas have been raised, brainstorming has been done, which were planned to be developed in detail during the next two workshops in Paris and Tsukuba.

3.3 Workshop in Paris, France

The main goal of the Paris meeting (June 28, 29, 2009) was to initiate development of a coordinated software roadmap and get the participants fully engaged with the process of developing a draft version. Since the roadmap is intended to be a living document, used to guide and coordinate the work of the international open source software community over time, it is intended to serve as the vehicle through which progress on the other objectives in the plan are expressed. For example, the most fundamental set of inputs into the roadmap are the assessments of the short, medium and long term needs, issues, and strategies emerging from an analysis of emerging peta/exascale architectures and the scientific research agendas IESP aspires to serve. Consequently the majority of the effort at the second workshop was concentrated on those analyses and assessments. On the other hand, the other objectives of the plan can be thought of as a derived product of the roadmap, once it is created. These objectives include the following:

- Provide a framework for organizing the software research community;
- Encourage and facilitate collaboration in education and training;
- Engage and coordinate vendor community in crosscutting efforts.

After plenary sessions to establish common context, the participants at the workshop divided into four breakout groups –software, applications, vendors (industrial collaborators), and funders. The division was intended to facilitate the analysis and assessment of the underlying needs, issues and strategies that must feed into the roadmap. The software group, which was the largest, found it natural to divide again in order to attend to intra-processor and inter-processor issues separately; the results of their discussions were then consolidated at the end of the workshop.

In general, the second workshop provided the framework that the project needed to pursue in going forward. More specifically, the key components of roadmap have been identified and their content has begun to be filled in. An initial group of authors who will lead the effort to take the roadmap components forward has been identified, though it is clear many volunteers will be needed to help finish this work in or for the project to stay on schedule.

3.4 Workshop in Tsukuba, Japan

The goal of the workshop in Tsukuba (held from October 18-20, 2009) was to prepare a draft version of the roadmap, for which the framework was set in the Paris meeting. This basically meant that during break-out groups in the workshop, many topics that were identified in the framework roadmap have been discussed in break-out groups, eventually leading to pieces of text to include in the draft roadmap. After the Tsukuba workshop, lots of effort has been put in editing the contributions, presenting a consistent document and avoiding too much overlap.

In fact, the IESP software roadmap is a planning instrument designed to enable the international HPC community to improve, coordinate and leverage their collective investments and development efforts. After the Tsukuba workshop, the software components for exascale computing (X-stack) have been described, together with the principle of co-design, in which hardware and software development will go along with application needs and perspectives.

A key aspect here is a timeline, starting now up to the end of the decade. For each X-stack component, targets and milestones have been defined, which must lead to a full-blown, tested X-stack in the 2018-2019 timeframe.

Another result of the Tsukuba workshop has been that is was recognized that organizational and planning issues need to be discussed as well. This has been an important subject of the next workshop in Oxford, UK.

3.5 Workshop in Oxford, UK

The workshop in Oxford, UK, has been organized from April 12-14, 2010. Important organizational aspects have been discussed, among which:

- Cooperation between IESP and HPC vendor communities;
- IESP organization and governance;

In order to address these aspects, more focused break-out groups have been set-up for the funding agencies on organizational and governance aspects and also for the vendors and facilitators on their considerations.

These discussions have eventually led to version 1.0 of the roadmap, May 30, 2010.[1]

[1] International Exascale Software Project: Roadmap 1.0, www.exascale.org.

4. Workshop in Maui, HI, USA

4.1 Introduction

The fifth IESP workshop has been organised in Maui, HI, USA, on October 18-19, 2010. At the start of the workshop, the goals included the following:

- Refine software roadmap for software and algorithms on extreme-scale systems based on feedback since initial release;
- Refine the prioritized list of software components for Exascale computing as outlined in the Roadmap;
- Refine the assessment of the short-term, medium-term and long-term software and algorithm needs of applications for peta/exascale systems;
- Explore how laboratories, universities, and vendors can work together on coordinated HPC software;
- Explore governance structure and management models for IESP;
- Explore IP/Open Source issues.

The morning session of the first day of the workshop was devoted to updates from the international regions: US, Japan, China and Europe. In the afternoon of the first day, focus has been on the vendor activities and on the co-design proposals the Department of Energy (DoE) in the US is currently evaluating. Sections 4.2 to 4.4 will report on these activities. The second day of the workshop has been organised in discussion sessions in break-out groups, initially focused on hardware, software and applications. Section 4.5 will report on this.

4.2 Update on international efforts

4.2.1 US

Barbara Helland from the Office of Science of the DoE presented the current efforts of the DoE on exascale computing. With respect to the funding of Exascale initiatives in the US, the DoE is active in calls for proposals in various exascale-related areas:

- Applied Mathematics: Uncertainty Quantification (6 proposals funded at $3M/yr);
- Computer Science: Advanced Architectures (6 funded at $5M/yr);
- Computer Science: X-Stack (11 funded at $8.5M/yr);
- Computational Partnerships: Co-Design (21 Proposals requesting ~ $160M/year, decisions expected early 2011).

It has turned out that there is quite some interest in these calls, as the cumulated value of the proposals is much larger that the amount of funding available. Further calls and funding are under development, both from the X-stack perspective and the co-design perspectives. Planning efforts are underway with respect to an exascale software center and five co-design centers. The initial proposals are there, refinement will need to take place, reviewing is expected in spring 2011.

Beckman continued with presenting the planning effort for the aforementioned Exascale Software Center (ESC), with the following initial scope:

- Deliver high quality system software for exascale platforms ~2015, ~2018;
- Identify software gaps, research & develop solutions, test and support deployment;
- Increase the productivity and capability and reduce the risk of exascale deployments;
- Cost:

  - Applied R&D: ~10-20 distributed teams of 3 to 7 people each;
  - Large, primarily centralized Quality Assurance, integration, and verification center;
- Schedule Overview
  - 2010 – Q1 2011: Planning and technical reviews;
  - April 2011: Launch;
  - 2014, 2017: SW ready for integration for 2015, 2018 systems respectively.

Apart from important organizational aspects (like quality assurance and interfacing of software components), the principle of co-design of hardware, taking applications into account during the design process, is regarded as an important way forward. As described above, the DoE is reviewing proposals that apply for co-design centres. The co-design process is where system architects, application software designers, applied mathematicians, and computer scientists work closely together to produce a computational science discovery environment that fully leverages these significant advances in computational capability.

Many important aspects have been identified already, one of the most important is the relation (to be formalized) between the software developers and the vendors, with respect to development, support, risk and acceptance.

It is also recognized that expertise on certain areas of the software stack is concentrated in certain institutes in the US, which makes distributed project staffing a sensible option. Each such “component team” must have a certain mass of development staff, but also with quality assurance and testing manpower available.

4.2.2 Japan

Satoshi Matsuoka and Yutaka Ishikawa described the HPC challenges in Japan, and basically focused on two major aspects:

- The Next-Generation Supercomputer Project

  After the evaluation of the Riken project, the Next-Generation Supercomputer became known as the K (KEI) computer, targeted at 10 Pflop/s to become available during the course of 2012. With more than 640k cores in 80k nodes, more than 1 PB of main memory, stored in 800+ cabinets, the system is massive. Full installation is scheduled for 2011, fine-tuning and further improvements for 2012. Processor cores are based on Fujitisu’s SPARC-8FX technology. The system will be installed in Kobe, around 450 km west of Tokyo.

  Apart from the K-system, other Pflop/s facilities are expected to become available, starting at end 2010, early 2011 (Tokyo Institute of Technology, University of Tsukuba, University of Tokyo and Kyoto University, together T3K alliance). The Tsukuba system will focus on (green) GPU-technology

- The strategic program to promote HPC activities in Japan

  During the third Science and Technology Basic Plan (2006-2010), strategic applications for Japan have been identified for execution on the high-end systems:

  - Life sciences and drug development;
  - Materials science and energy;
  - Global climate change;
  - Engineering and manufacturing;
  - Origin of matter and universe.

  In order to support these research areas, the High Performance Computing Infrastructure (HPCI) has been created. The K-facility at Kobe will act as a centerpiece for HPCI. Many organisations have already joined HPCI, ranging from research institutes to technology providers, but also to

industrial organisations which can apply for compute cycles on the systems. The goals for HPCI are:

- Move from fundamental research levels to actual applications both in academic and industrial sectors;
- Create new mechanisms to let experimentalists use simulations as daily research tools;
- Enable collaboration between computational scientists and computer scientists for effective use of many-node systems.

With respect to further plans, it is expected that there will be around 30 M$ available for system software oriented projects, which will have a runtime of 5 to 7 years. These plans towards exascale and exaflops are under discussion in the fourth Science and Technology Basic Plan (2011-2015).

Finally, Japan and France are funding collaborative calls, including software and algorithmic aspects of HPC. Japan also takes part in the G8 calls, for which many projects have been proposed and which are in the second stage of selection.

4.2.3 China

Xue-bin Chi described the development of HPC in China. After a brief history of the computing resources at the SCCAS (Supercomputing Center of the Chinese Academy of Sciences), details on algorithmic development and applications were given, together with a strong increase in the amount of users of the SCCAS computing systems. Dr. Chi mentioned climate applications, radar models, cosmology, engineering applications in the car manufacturing industry. Not only emphasis is given on standard compute clusters, but also on implementation of applications on GPU-clusters. Several GPU-clusters, each with several hundreds of Tflop/s peak performance are operational at SCCAS.

Since 2001, China operates the China National grid, which connects universities and scientific institutes throughout the country. Currently, the gird is mainly composed of Lenovo and Dawning systems, reaching peak performances in the Pflop/s range. As of November 2010, China currently has two systems which have a LINPACK performance of more than one Petaflop/s. The first one is located in Tianjin, the second one at the Dawning facilities in Shenzen. Both systems are equipped with dual-processor 6-core Intel nodes, extended with Nvidia Fermi GPU cards.

For the period 2011-2015, China expects to operate many petascale systems, including a top system with 50-100 Pflop/s peak performance. In the period 2016-2020, a system in the range of 1 to 10 Exaflop/s is expected. These systems may use China's own development on its low-power processor, based on MIPS-compatible CPUs. The processor, known as Loongson, Godson or Dragon, will be developed at the Institute of Computing Technology, Chinese Academy of Sciences.

4.2.4 Europe

The situation in Europe was covered in three presentations: Berthou, Aerts and Johnson.

Berthou kicked off on European initiatives, the EC vision and funding. He started with an overview of the European Exascale Software Initiative (EESI), which coordinates the European participation in IESP. An important aspect is the concept of Working Groups (WG's), four of which deal with scientific application areas and four of which deal with hardware, system software, numerical algorithms and software engineering. The expected output of EESI is a roadmap and a set of recommendations to the funding agencies, shared by the European HPC community, on software (tools methods and applications) to be developed for exascale supercomputers. For the EC, such output must provide an analysis of European strengths and weaknesses, the willingness of European stakeholders to build exascale systems, in relation to the recently issued IDC report.[2]

[2] "Development of a Supercomputing Strategy in Europe", IDC, October 2010.

Berthou mentioned the EC initiatives on exascale computing in the FP7 program (described in detail by Aerts), and the collaboration with Russia for a total of 6 M€.

With respect to exascale research (hardware and software), there are some initiatives in European countries between universities/research labs/computing centers on one side and industry on the other side. A brief overview:

- Exascale Innovation Center, Germany, JSC (Juelich) and IBM;
- EX@TEC, France, CEA, GENCI, UVSQ and Intel;
- Flanders ExaScience Lab, Belgium, IMEC, five Flemish universities and Intel;
- Exascale Stream Computing Collaboratory, Ireland, IRCSET, four universities and IBM;
- Exascale technology Centre, UK, EPCC and Cray

Berthou concluded with a brief overview of actual projects, installations and plans in the individual European countries.

Aerts continued with an overview of the current EC activities in the FP7 framework program. During his presentation, he raised the question on reciprocity: what should the EC expect from the international cooperations in the exascale domain ? Aerts illustrated this question with the observation that for instance processor design is not a European activity, but that details must be shared with Europe as well to design efficient software.

- Call for Computing Systems

  This 45 M€ call is targeting projects to start mid 2011 and to end in 2014. Most of the call (40 M€) is dedicated to research projects, for which the research vision has been laid out in the European Network of Excellence on High Performance and Embedded Architecture and Compilation. The research topics include:

  - Parallel and concurrent computing;
  - Virtualisation;
  - Customisation;
  - Architecture and technology.

  The expected impacts are:

  - Improved programmability of future systems;
  - Efficient and ubiquitous use of virtualisation for heterogeneous multi-core systems;
  - Accelerated system development and production;
  - Reinforced European excellence in multi-core computing architectures, system software and tools;
  - Strengthened European leadership in cross-cutting technologies.

  The call closes on January 18, 2011.

- Call for exascale computing, software and simulation

  This 25 M€ call is dedicated specifically to exascale computing, which marks the commitment of the EC to support research at the leading edge of HPC. Most of the (24 M€) is dedicated to research projects, which are expected to be collaborations between one or more supercomputing centers, technology and system suppliers and industrial and academic centers. The intention of the call is:

  - To develop a small number of advanced computing platforms (100 petaflop/s in 2014 with potential for exascale by 2020), platforms relying on vendors' proprietary hardware or on COTS hardware;

- To develop optimised application codes driven by the computational needs of science and engineering and of today's grand challenges (e.g. climate change, energy, etc.);
- To address major challenges of extreme parallelism with millions of cores (programming models, compilers, performance analysis, algorithms, power consumption ...).

The expected impact is:

- Put Europe in the frontline of international efforts for the development of HPC system software and tools;
- Strengthen European industry supplying and operating HPC systems: preparing European industry and research organisations to achieve world-leadership in this area;
- European excellence in exascale level simulation codes for the benefit of society, industrial competitiveness and policy making; emergence of EU top-class simulation centres for exascale systems;
- Reinforce cooperation in international endeavours on exascale software and systems.

It is further expected that proposals:

- develop software as open source;
- split the effort roughly 40/60 in applications and simulation vs. systems development;
- demonstrate synergies with efforts under the Capacities programme on the deployment of leadership-class HPC systems;
- include international cooperation components that are essential and complementary to European expertise.

It is clear that international collaboration is needed for implementation of the Exascale Software Roadmap. Not only scientific issues are challenging, also organisational issues like the coordination of efforts, the governance model and the collaboration with vendors and Exascale Software Centers in the US and Asia, etc., are subject for investigation.

This call also closes on January 18, 2011.

Johnson presented the activities in the PRACE project and in the PRACE legal entity, which is responsible for granting access to the PRACE Tier-0 systems. Currently, there is one PRACE Tier-0 system (1 Pflop/s IBM BlueGene/P at JSC in Juelich, Germany). An interim Board of Directors, representing the hosting partners (Germany, France, Spain and Italy), with one representative of the non-hosting partners, has organised the first access calls, including the peer review. The legal entity (AISBL in Belgium) is governed by a council, consisting of all partners (currently 20 countries). A Scientific Steering Committee (SSC) advises the council, an Access Committee (AC) will be recruited from the SSC and will lead the process of assignment of computing time on the Tier-0 system(s). The following schedule of Tier-0 systems is foreseen:

- IBM BlueGene/P at Gauss Center for Supercomputing (GCS), Juelich, Germany;
- Bull system, Intel-based, GENCI, France (initial installation by end 2010, full installation by end 2011);
- 3rd PRACE system at HLRS, Stuttgart, Germany, 2011;
- 4th PRACE system at LRZ, Munich, Germany, 2012;
- 5th and 6th PRACE systems in Italy and Spain, 2013.

The EC intends to finance four PRACE projects: one preparatory phase (has run from January 2008 until June 2010), and three implementation phases (first has started on July 1, 2010). The EC is expected to contribute 70 M€ in total. Partners will match the EC contributions up to 100% maximum. The PRACE projects focus on several aspects:

- Set-up and operational model of the Research Infrastructure;
- Operational aspects, best practices;
- Application usage, optimisation, parallelisation and scalability;

- Benchmarking;
- Future hardware and software developments, including prototyping and novel programming languages;
- Dissemination, education and training;

An important aspect for PRACE is to watch technology and component developments. Since this may interfere with procurement cycles of Tier-0 systems, STRATOS (STRAtegic TechnOlogieS) has been established, which will be able to communicate with vendors on these aspects without being hindered by procurement rules.

4.2.5 Example high-end research infrastructure: SKA

Tim Cornwell presented the Square Kilometre Array (SKA) as an interesting project for co-design of future hardware and software environments. SKA is a 2020 era radio telescope, based on the principle of the Lofar telescope in the Netherlands. SKA will be sited in Australia or South Africa, and will contain 3600 antennas. The Lofar project already uses huge amounts of compute cycles to interpret the received signals, which will need to be order of magnitude more for SKA. It is for this reason that co-design of hardware and software, given the requirements and technical specifications of SKA, could be beneficial.

SKA is expected to need around 100 Tflop/s compute capacity and 100 PB storage capaity a year by 2012, increasing to 1 Pflop/s by 2014. requirements are expected to grow into the exascale area before 2020.

4.3 Vendor perspectives on exascale and co-design

4.3.1 Cray

Peg Williams (Cray) mentioned the key challenges for exascale, together with the focus areas for Cray:

- Power (system infrastructure, network, heterogeneous processors);
- System software (OS, file system scalability, jitter reduction);
- Programming systems (libraries, locality, compilers, tools, languages);
- Resiliency (system, application).

Cray has currently two exascale research initiatives in Europe: at EPCC in Edinburgh, UK, and at CSCS in Lugano, Switzerland. Apart from software life cycle management and the availability of test platforms and test strategies, an important topic for collaboration between vendors and the open source community is the topic of vendor differentiation and IP protection. Definition of Application Programming Interfaces (APIs) could be a valuable vehicle to enable vendors to differentiate without giving full details on their plans.

4.3.2 Fujitsu

Ross Nobes (Fujitsu Laboratories in Europe) presented Fujitsu's activities in capability computing. Fujitsu works with RIKEN on the K-computer in Kobe, Japan. Fujitist looks at:

- SPARC64 processor development;
- 6D mesh/torus interconnects;
- Direct water cooling packaging;
- Programming model: users should not worry on multiple cores on a single chip, Fujitsu's vectorisation experience in compilers will solve that;
- Hybrid task/thread model.

Open initiatives:

- Exascale Application and Data Initiative – XcalableMP: directive-based language extension to avoid using MPI;
- Open Petascale Libraries Project – global collaboration to develop advanced numerical software for supercomputing.

4.3.3 IBM

Robert Wisniewski (IBM) explained IBM's activities on exascale. Internally, there are a lot of projects going on for a few years already, covering both hardware and software developments. Full software ecosystem should be addressed: programming models, performance tools, OS. All aspects need to be re-designed with scalability in mind. A typical example are atomic operations, for which the time to synchronize all threads increases with the number of threads. General messages:

- Exascale software delivery will have challenges, but is tractable;
- Co-design and collaborative development are key components;
- Existing programming models will need to migrate;
- Open Source and community provided software will play a significant role – find models for:
    - Vendor product schedules and delivery;
    - Platform-enabling software must remain proprietary;
    - Define Open Source contributions

4.3.4 Intel

David Lombard (Intel) explains the investments Intel has done in various collaborations on exascale software (Juelich, Leuven, Paris). Various aspects will be covered in these labs. As Intel is a hardware vendor, focus is on hardware enhancements:

- Extreme voltage scaling to reduce core power;
- More parallelism 10x – 100x to achieve speed;
- Re-architecting DRAM to reduce memory power;
- New interconnect to lower power and distance;
- Non-volatile memory to rduce disk power and accesses;
- Resilient design to manage unreliable transistors.

Co-design is viewed as essential, just as collaboration with the Open Source Community, although new approaches will likely be required.

4.3.5 Microsoft

Dan Reed (Microsoft) did not focus on technical challenges, but much more on challenges in business model. Computing is already everywhere, and will increasingly be important in everyday life. Think of all kinds of devices, data and information, but also more and more natural user interfaces. The key question is whether exascale software development for HPC-type of applications will be such that it can be leveraged by the much broader and commercially much more attractive commodity, everyday software environments.

4.3.6 HP

Rob Schreiber (Hewlett-Packard) confirmed that HP is committed to pursue exascale systems for HPC. He mentioned the fact that HP labs is investigating several new technologies, which will be crucial to get to exascale systems: photonics, memory technology (non-volatile, memristor). In the

software area, HP predicts a huge growth of the needs in the commercial arena. Various exascale aspects will need to be solved there as well.

4.3.7 Bull

Jean-Francois Lavignon (Bull) stated that the road to exascale is more of a journey than of a real roadmap, given all challenges to be solved underway. These challenges are not only technological, but also from a business point of view, for instance how to reuse or apply HPC developments for exascale in the more commodity areas. Co-design as a hardware, application, algorithm and software collaboration may help (an example is Bull's Extreme Computing lab, with CEA and others), but no valid business model yet seen for doing this based on a single (set of) application(s). Bull's software strategy is based on Open Source plus added value provided by Bull experts.

4.4 Co-design planning efforts

The Office of Advanced Scientific Computing Research (ASCR) of the Office of Science (SC), U.S. Department of Energy (DoE) has announced a call for proposals from integrated teams of scientific researchers, applied mathematicians, computer scientists and computer architects with the goal of setting up so-called Exascale Co-Design Centers. Co-design refers to a computer system design process where scientific problem requirements influence architecture design and technology and constraints inform formulation and design of algorithms and software. To ensure that future architectures are well-suited for DoE target applications and that major DoE scientific problems can take advantage of the emerging computer architectures, major ongoing research and development centers of computational science need to be formally engaged in the hardware, software, numeric methods, algorithms, and applications co-design process that will be responsible for making key tradeoffs in the design of exascale systems.

At the time of the Maui meeting, the DoE was in the process of evaluating and ranking the received proposals. As said in section 4.2.1, 21 proposals for a total of 160 M$ had been received.

During the Maui meeting, five co-design centers presented themselves. Their presentations have not been made publicly available, as a result of the fact that the DoE is still evaluating the proposals. The presenters have been:

- Robert Harrison, Oak Ridge National Laboratory (ORNL), on the proposal for "Chemistry Exascale Co-design Center" (CECC);
- Andrew Siegel, Argonne National Laboratory (ANL), on the proposal for "Center for Exascale Simulation of Advanced Reactors" (CESAR);
- Alice Koniges, NERSC/Lawrence Berkeley Laboratory, on the co-design proposal for fusion;
- Sriram Swaminarayan, Los Alamos National Laboratory (LANL), on the co-design proposal for materials science;
- Carlo Graziani, The University of Chicago, on the co-desihn proposal for high energy density physics.

4.5 Break-outs on hardware/software/applications

The second day of the workshop focused on the three pillars to get to efficient exascale computing: hardware, software and applications. The discussions have been set up around the following questions:

- What do the software people need from the application and hardware groups ?
- What do the applications people need from the hardware and software groups ?
- What do the hardware people need from the software and application groups ?

In practice, the software and applications group merged, so there were two parallel sessions: hardware and software/applications.

Summary of hardware break-out group

Key question: What do the hardware people need from the software/applications group ?

The key constraints for hardware development and production for exascale computing are power budget, procurement budget and delivery schedules. These constraints get back when considering design trade-offs for the nodes, interconnect, I/O and accelerator hardware. Needed from the software/applications group are answers/insights:

- With respect to node design:
    - Heterogeneous nodes: ratio heavy-weight nodes and light-weight nodes;
    - Threads: importance of single-thread performance and the amount of threads effectively used per coherence domain;
    - Data movement: separate data mover development;
    - Memory hierarchy (cache sizes and hierarchy);
    - Memory design (non-volatile): what about capacity versus latency;
    - Power management of cores, memory, …;
    - Hardware features for resiliency;
- With respect to interconnect:
    - Relative importance of bandwidth, latency and messaging rate;
    - Evaluation of network technologies – how;
    - Importance of inter-node communication with respect to performance;
- With respect to I/O:
    - What I/O rates do we need;
    - New approaches to data analysis and resiliency.

General questions include:

- What types of acceleration can be put in hardware ?
- What hierarchical and hybrid programming models need to be supported ?
- Will there be a standard and stable API to access new features ? Who is responsible ?
- Van HW development use simulation/emulation tools to rapidly characterize application behavior ?

Summary of applications break-out group

Key question: What do the software/applications people need from the hardware group ?

This summary is structured through the question “What Co-Design Vehicles want”. The break-out group not only discussed what answers were needed form the hardware group, but also posed questions to themselves. There have been 21 candidate CDV’s identified, which probably do not represent yet all computing areas, but are a good starting point.

Current dominant findings on applications:

- Most applications are bulk synchronous, MPI-based, some kind of decomposition (domain, particle, …), but not all of them (e.g. electronic structure codes);
- Already running hybrid MPI/OpenMP, OpenMP to several tens of cores;
- Weak-scalability can be done efficiently on current largest systems (300k cores);
- Majority of CDV’s memory bandwidth limited;
- Flop/s per bye of storage is linear or log-linear in weak scaling;
- Intensive I/O in most cases only during start-up and finish, although check-pointing may become more important;
- Global collectives on synchronisation required (time steps, implicit linear algebra);

Algorithms and library dependencies:

- PDEs: equilibrium (implicit) and evolution (explicit);
- FFT;
- FMM;
- Particles pushing;
- Adaptive mesh refinement;
- Sparse and dense linear algebra;
- ODE integrators;
- MPI, GlobalArrays, GasNet;
- ParMetis;
- BLAS, ScaLAPACK;
- UMFPACK, SuperLU, MUMPS;
- PETSc, hypre, Trilinos, SUNDIASL;
- Chombo, SAMRA;
- FFTW;
- GraphLib;
- Vislt, VTK;
- Triangle;
- PALM;
- SILO, ADIOS, HDF5;
- BOOST.

Applications wish list for hardware/software:

- Programming models that optionally expose machine characteristics such as memory hierarchy, cores, etc.
- Memory related
    - Exposure of memory hierarchy;
    - Hardware gather-scatter operations;
    - Programmable caches (not just DMA engines: it would be nice to specify a memory access pattern or to mark regions of memory as write rarely/read often;
    - Access to DMA;
- Compiler related:
    - Single source across different implementations of exascale;
    - Auto vectorisation;
    - Masked SIMD operations;
- Power related:
    - Ability to power on/off pieces of hardware;
    - Ability to hint on power patterns to hardware;
- Tools:
    - Tools for early adopters vs. Tools for general use;
    - Tools for hot spot analysis and bottleneck identification;
    - Hardware counters available to users.
- Implementations of exascale:
    - Just GigaHertz*MegaNode*KiloCore ?
    - Or also other implementations: heterogeneous, GPU, accelerators.

Questions to the software/application group themselves:

- Size of the coherence domain;
- Exploring new programming models on the following aspects:

    - Manycore and heterogeneous nodes;
    - Data locality enhancements;
    - More loose synchronisation models;
    - Data movement becomes relatively more expensive compared to computing;
    - Mixed-precision algorithms;
- Fault tolerance:
    - Categorise type of faults, and identify there severity;
- Reproducibility:
    - Is bit-level reproducibility required for applications results ?
    - What would be the hardware/power advantages when bit-level reproducibility would be relaxed ?
- Big Data:
    - What would the availability of hundreds of TB to tens of PB main memory mean for application design ?

Other aspects:

- Investigate scalability on three levels: weak (one instance of application), strong (one instance of application), workflow scalability (across instances);
- Load balancing becomes extremely critical when using exascale sizes;
- Collective communication aspects;
- Hierarchical algorithms (like multigrid) may suffer on coarse levels.

4.6 Next meeting

The next IESP workshop meeting has been planned in Kyoto, Japan, April 6-7, 2011.

4.7 Conclusions

The Maui IESP meeting was the fifth workshop in a row which started in Santa Fe in April 2009. Building on the findings of the earlier workshops, the discussions during the meeting in Maui focused on the start of the practical implementation of Co-Design Vehicles (CDV's). The main part of the second day has been used for trying to identify as clearly as possible which kind of information was needed for the successful collaboration between application, hardware and software groups. This information has been characterized and should be used for the implementation of Co-Design Centers.

With respect to these Co-Design Centers, the DoE in the US has taken steps forward in selecting Co-Design Centers. Calls for proposals have been made, leading to 21 candidates, which are in the process of being reviewed. It seems that the US has taken the lead here.

With respect to setting up an Exascale Software Center (ESC), there are planning efforts underway, which foresee in delivering high quality system software for exascale platforms in the 2015-2018 range, to enable increased productivity and capability and reduce the risk of exascale deployments. The initial staffing effort of the ESC is anticipated as 10-20 distributed teams of 3 to 7 people each for applied Research & Development, and additional staff to ensure quality assurance and verification. The launch is foreseen in the first half of 2011.

Another important topic which has not been discussed yet in detail is the relationship between vendors on one side and applications/software groups on the other side. There must be a model in which applications/software groups are able to anticipate on new technology developments at vendors, which may be confidential and/or strategic for those vendors. A clearly defined Application Programming Interface (API), including responsibilities on maintaining will be required. An additional dimension to this issue may be the aspect of national security: many vendors are US-based, which could mean limited access to application/software groups outside the US.

After this fifth workshop it has become apparent that IESP is able to build a broad community that is prepared and able to mobilize the stakeholders and others concerned to achieve the goals set forth by the IESP and prepare for a large scale redesign effort for software/applications in close cooperation with technology developers. It is also clear that the US strongly will go ahead whatever happens outside of the US. Only if the EC realises its own stakes in this matter (as up to now seems to be the case if one looks at the present FP7 and IDC efforts), Europe will be ready to deploy next generation HPC tools and equipment to Europe's own case and advantage.

5. Workshop in San Francisco, CA, USA

5.1 Introduction

The sixth IESP workshop has been organised in San Francisco, USA, on April 6-7, 2011. The original plan was to organise the workshop in Kyoto in Japan, but due to the severe earth quake and its consequences, the organisers have concluded to move the meeting. Hopefully, one of the next workshops can be held in Kyoto. At the start of this workshop, the goals included the following:

- Refine software roadmap for software and algorithms on extreme-scale systems based on feedback since initial release;
- Refine the prioritized list of software components for Exascale computing as outlined in the Roadmap;
- Refine the assessment of the short-term, medium-term and long-term software and algorithm needs of applications for peta/exascale systems;
- Explore how laboratories, universities, and vendors can work together on coordinated HPC software;
- Continue to develop co-design process with participating application leaders;
- Explore governance structure and management models for IESP;
- Explore IP/Open Source issues.

The first day of the workshop started with a summary of the previous workshops until now and a summary of a global software stack survey, both reported on in section 5.2. These were followed by updates by Europe on EESI, by the US on the planning of the Exascale Software Center (ESC), by Japan and China on their activities, reported on in section 5.3. The second day of the workshop has started with an update on the US co-design centers (section 5.4), after which break-out sessions continued on earlier break-out sessions in the Maui workshop on the areas of hardware, software an applications, in particular with respect to the information one area needs from the others.

After that, break-out sessions on the following topics were organised:

- Software demands from hardware and applications;
- Hardware demands from software and applications;
- Application demands from hardware and software;
- Actual implementation of co-design;
- Global governance aspects;
- X-stack components and coordinated workplan.

Apart from these, Intellectual Property from a vendor's view was presented as well. Section 5.5 will report on this and on those break-out sessions.

5.2 Summary of previous workshops and software stack survey

5.2.1 Summary of previous workshops

In the EESI project, one of the activities is to report on international efforts with respect to exascale computing. Peter Michielse and Patrick Aerts, both of NCF in the Netherlands, are responsible for this activity in EESI. In this role, Peter Michielse presented a summary on the five earlier IESP workshops so far. Key findings are:

- Attendee statistics: ca. 70% universities/research institutes, ca. 15% vendors/industry, ca. 15% government/funding agencies;

- Main achievements of IESP so far: software roadmap and awareness;
- Current status/approach in IESP:
    - Maintain science drivers;
    - Investigate software stack and prioritise;
    - Set-up co-design centers per application area;
    - Find a proper model to include vendor participation;
    - Pose questions and supply answers on hardware, software and applications aspects;
    - Find a proper funding and organisational model;
    - Timeline: 2010 software roadmap, 2012 budgets, 2015 early technology deliveries, 2018 initial delivery and full X-stack, 2020 production systems.

For more information, check http://www.exascale.org/mediawiki/images/2/2c/EESI-D2_3-report-on-international-activities-05-01-2011.pdf.

5.2.2 Software stack survey

As a preparation to this IESP workshop, a survey on the system software stack in use at the large HPC centers around the globe has been performed. Bernd Mohr has reported on the findings:

- Survey sent to 28 HPC centers around the world, response from 23;
- Not entirely clear how centers filled the survey: wishlist for software, or what is installed, or what is frequently used;
- Not so easy to draw firm conclusions, although some trends can be spotted:
    - Linux as an OS seems clear;
    - GPFS, Lustre, NFS as filesystems;
    - NetCDF, HDF5 and MPI-IO as I/O libraries;
    - Range of batch systems, with LoadLeveler popular in Europe;
    - Programming models: MPI variants with OpenMP;
    - Programming Languages seem to focus on C/C++, F77, F90/95, Python, CAF, UPC and Java;
- Viewed as critical for exascale:
    - Definition of base OS, with APIs for resilience, energy management, performance monitoring and debugging, explicit memory management (including data placement);
    - Customization within I/O;
    - Scalable and fault-tolerant MPI;
    - Heterogeneous node programming model;
    - Load balancing which can deal with faults and lack of resources;
    - Compilers to include these demands;
    - Numerical libraries that include error-tolerance, hierarchical algorithms.

5.3 Global updates

5.3.1 Progress in Europe

As the coordinator of the European Exascale Software Initiative (EESI), Jean-Yves Berthou opened the session on progress in Europe, by providing a quick overview of EESI. The core of the EESI project consists of Work Packages 3 and 4, which deal with the demands of application grand challenges (WP3, led by Stephane Requena) and the enabling technologies for exaflop computing (WP4, led by Bernd Mohr). In some sense, WP3 can be viewed as similar to the applications group in IESP, while WP4 can be viewed as similar to the hardware/software group in IESP.

Requena presented the structure of the applications group in EESI: split in four working groups (WGs: industrial & engineering apps, weather/climatology/earth sciences, fundamental sciences, life sciences & health). The four subgroups in WP3 cover the following aspects:

- Apps drivers for peta and exascale;
- Needs and expectations of scientific applications;
- Economic impact and European competitiveness;
- European vision and roadmap.

Each WG consists of around 15 experts. Meetings are scheduled, working towards an end report by September 2011.

The situation with respect to WP4 is comparable to WP3, as presented by Bernd Mohr. The four WGs in WP4 are: hardware roadmap & links with vendors, software eco-systems, numerical libraries & solvers & algorithms, scientific software engineering. The end reports are scheduled for September 2011.

Both WP3 and WP4 have identified common key issues to structure their research. These are:

- Scientific and technical hurdles;
- Cross cutting issues: Resilience, Power Management, Programmability, Performance and Reproducibility of the results;
- European strengths and weaknesses in the worldwide competition;
- Sources of competitiveness for Europe;
- Needs of education and training;
- Potential collaborations outside Europe;
- Existing funded projects and funding agencies;
- Timeline, needs of HR, provisional costs, ...;
- Building an (or several) exascale prototype(s) in Europe? By when?

For WP3 (application grand challenges), two additional key issues are:

- Description of the scientific and technical perimeters;
- Social benefits, societal, environmental and economical impact.

In subsequent presentations, representatives of the WGs presented the actual status in their WG. These were:

- Jean-Claude Andre, WG3.1 (industrial & engineering apps);
- Giovanni Aloisio, WG3.2 (weather/climatology/earth sciences);
- Godehard Sutmann, WG3.3 (fundamental sciences);
- Ramon Goni, WG3.4 (life sciences & health);
- Riccardo Brunino, WG4,1 (hardware roadmap & links with vendors);
- Franck Capello, WG4.2 (software eco-systems).

5.3.2 Progress in the US

Pete Beckman presented the current status of the Exascale Software Center (ESC), for which a final funding decision will be taken soon. Three main topics were addressed, in the context of operations within the ESC.

1. How to identify which software to work on:

Together with the inventory of the software stack, as presented earlier by Bernd Mohr, a survey has been carried out on various applications for their expected software needs when preparing for exascale systems. This has led to field guides, which typically describe the breadth of the software involved. Apart from the application inventory (bottom up approach), there can also be applied a top down approach, which has science, architecture design and algorithms as starting points to get to software requirements. There will need to be a continuous dialogue between applications, co-design centers, research communities and vendors to identify the right software to work on, and to identify the depth needed.

2. How to organise the work:

Based on the identification of software, work will need to be done to get the software exascale ready. Research will be required for this, which starts the cycle of prototype building, testing and quality assurance, integration, deployment and support on hardware platforms. In the co-design centers and the broader software community, further needs will be identified, which will re-start the cycle again. Within the ESC, the organisation will be split in two: software research and development activities in distributed teams and quality assurance, integration and support activities in centralized teams. Resulting software modules will be deployed on DOE exascale systems. A related aspect here is the status of developed software:

- Open Source Community developed and ESC tested;
- ESC developed and supported software;
- ESC developed and vendor supported software.

3. Tasks and scope of the ESC

Based on items 1 and 2, the following tasks for the ESC can be identified:

- Design and develop open-source software;
- Ensure functionality, stability and performance;
- Partner with platform and software vendors;
- Coordinate outreach.

The ultimate goal is to ensure successful deployment of coordinated exascale software stack on exascale initiative platforms. The current planning is to deliver high quality software for 2015/2016 platforms (100 Pflop/s) and 2018/2019 platforms (1 Eflop/s).

Subsequently, components of the software stack were discussed:

- Operating systems and runtime environment (Ron Brightwell);
- Programming models (Barbara Chapman);
- Tools (Bronis R. de Supinski);
- Mathematical libraries and frameworks (Jack Dongarra);
- Data storage and analysis (Rob Ross);
- Applications inventory (Al Geist).

5.3.3 Progress in Japan

Mitsuhisa Sato presented the Next Generation Supercomputer (NGS) project (the K computer) which is underway to deliver 10 Pflop/s on Linpack in 2012. The system will be manufactured by Fujitsu and the installation at Riken in Kobe of the first out of 800 racks has started late September 2010. The strategic computational science program in Japan has identified five application areas that are expected to create breakthroughs on the K system. These are:

- Life sciences and drug manufacturing
- New materials and energy creation
- Global change prediction for disaster prevention and mitigation
- Manufacturing technology
- Origin of matters and the universe

At Riken, apart from being responsible for running the K computer, research teams in computer science and computational science have been established. For computer science these are:

- System software: OS, communication libraries, runtime environment;
- Programming environment: languages and compiler, runtime aspects, performance tools;

- Processor research: many-core, heterogeneous accelerators;
- Planned: numerical algorithms design, system architectures and interconnects, data-intensive computing and visualization.

For computational sciences, these teams reflect the five application areas, as mentioned before.

Yutaka Ishikawa described the High Performance Computing Infrastructure (HPCI) in Japan, and its key resources: computational, storage, internet and human. Ten universities and research institutes are involved in HPCI, interconnected through a fast network (SINET4). It has been recognised that development of human resources is key to the success of the supercomputing infrastructure in Japan. HPCI has an important role in developing these human resources.

Another aspect that has drawn attention since the earthquake and tsunami disaster, and its consequences for nuclear power installations in Japan, is the power consumption in the greater Tokyo area. This affects the nation's supercomputer installations, which are likely forced to use 30% less peak power consumption.

Satoshi Matsuoka presented the plans Japan has beyond NGS: Post Petascale Computing. More petascale systems will arrive in 2012 or perhaps earlier, with >20 Pflop/s expected in the 2014-2015 timeframe, based on many-core/multi-core architectures. Current system software initiatives have started for post petascale architectures for around 40 M$. Projects consist of:

- Parallel system software for multi-core and many-core (OS oriented);
- System software for post petascale data intensive science (data, filesystems);
- Highly productive high performance application frameworks for post petascale computing (programming, frameworks);
- ppOpen-HPC (open source numerical libraries)
- Development of Eigen-Supercomputing Engine using a post-petascale hierarchical model (numerical library).

Other projects:

- XscalableMP;
- Ultra low power HPC and green supercomputing;
- International collaborations (France, G8, ..).

5.3.4 Progress in China

Zhong Jin presented recent advances in Chinese HPC efforts. Funding of HPC in China is distributed over various agencies, of which CAS is one. The roadmap at CAS continues toward 2015 with 10+ Pflop/s systems. The model in China implements a pyramid model with on top national facilities, supported underneath by regional and local facilities. ScGrid (China Scientific Computing Grid interconnects these facilities. Application areas include drug design, earth science, earthquake, fluid dynamics, astrophysics, high energy physics, oil exploration and material science.

Future plans with respect to hardware are as follows:

- 2012: Pflop/s system and applications, focus on scientific computing, new facility in Beijing, budget 250 MCNY (ca. 40 M$);
- 2015: 10 Pflop/s system, also industry involvement, budget 700 MCNY (ca. 110 M$);

Another funding agency is MOST, which is involved in the five-year plans of the Chinese government:

- 12th five-year plan (2011-2015): several petascale systems, at least one 50-100 Pflop/s system, budget around 4 BCNY (ca. 600 M$);
- 13th five-year plan (2016-2020): 1-10 Eflop/s systems, budget yet unknown.

International collaborations include NCSA/UIUC in the US and Juelich in Germany, while more are under consideration (GENCI in France and Riken in Japan).

Lu Yutong presented the Tianhe-1A system at the National University of Defense Technology (NUDT) in China. The Tianhe-1A system is a heterogeneous system with a proprietary interconnect, delivering a Linpack number of 2.57 Pflop/s. Peak performance is 4.7 Pflop/s, power consumption around 4 MW.

5.3.5 Results of first international call (G8)

A total of 6 proposals have been funded in the G8 call for Exascale Software Applications. These are:

- "Enabling Climate Simulation @ Extreme Scale" – US, Japan, France, Canada, Spain;
- "Icosahedral-Grid Models for Exascale Earth System Simulations" – Japan, UK, France, Germany, Russia;
- "Nuclear Fusion Simulations @ Exascale" – UK, US, Germany, Japan, France, Russia;
- "Using Next-Generation Computers & Algorithms for Modelling Dynamics of Large Biomolecular Systems" -- Japan, UK, France, Germany, Russia;
- "Modeling Earthquakes and Earth's Interior based upon Exascale Simulations of Seismic Wave Propagation" – US, Canada, France;
- "ExArch: Climate Analytics on Distributed Exascale Data Archives" – UK, US, France, Germany, Canada, Italy.

William Tang presented details of the proposal on Nuclear Fusion Simulations. The goal of the project is predictive simulation of fusion experiments at all time and spatial scales. The real experiments will take place at the ITER facility. This is an extremely complex problem, which requires exascale systems to perform large future simulations of more than 10 billion grid points.

5.4 Developments in co-design

Four US Co-design centers (CDC) were presented, for which details follow in the next subsections. All CDCs have a more or less similar iterative approach to co-design. Some important aspects in the co-design process are:

- Development of co-design vehicles (CDV). These cover kernels, skeletons and compact applications which enable fast testing, both on node level and on large-scale behavior. CDVs should reflect a representative workload and include key code components;
- Hardware-based simulation for rapid turnaround of node designs;
- Time management: typically, the window of opportunity for hardware changes is early in the process, while (system) software changes may be implemented during a larger period;
- How to deal with Intellectual Property (IP) in relation to vendors (hardware and software). This may prevent actual co-design, as vendor architectures may be closely guarded secrets.

In the subsequent sections we will briefly list the issues which the CDCs view as the most challenging for their application area.

5.4.1 Reactor Core Physics

Andrew Siegel (ANL) presented CESAR: Center for Exascale Simulation of Advanced Reactors. CESAR deals with three building blocks for which the performance and success for all three depends on:

- Sufficient memory per MPI process;
- Highly efficient MPI collectives;
- Sufficient memory bandwidth;
- Optimized matrix-vector products;
- Performance analysis tools;

- Reproducibility of results;
- Data storage and analysis for massive amounts of data.

5.4.2 Nuclear Fusion

Alice Koniges (LBNL) presented CERF: Co-design for Exascale Research in Fusion. Also here, three building blocks need to be glued together to allow for full simulations. With respect to fusion codes, the following aspects to reach exascale are of highest priority:

- Application coupling;
- Mathematics research (solvers, libraries, algorithms);
- Programming model research;
- Enabling and performance tools;
- Uncertainty quantification.

5.4.3 High-Energy Density Physics

Anshu Dubey (ANL) presented FLASH High-Energy Density Physics Co-Design Center. Instead of discussing the main challenges for HEDP applications, Dubey translated the exascale challenges in co-design opportunities, which offered a useful other angle into the exascale aspects:

Challenge	*Co-design opportunity*
Parallel I/O	Data staging (critical vs. non-critical data and combination with in-situ analysis
More macro-parallelism	New parallel algorithms
Coupling for better physics	Different class of numerical algorithms
Intra-node memory limitations	Aggressive reuse of memory, other programming models
Faults, failures, silent errors	Redundancy and stochastic algorithms

5.4.4 Combustion

Jacqueline Chen (SNL) presented the Combustion Exascale Co-Design Center. Apart from the aspects as mentioned by the other CDCs, the aspect of data management was given attention. In order to minimize data movement, storage requirements and power, in situ data management is considered as a technique to reach these goals.

5.5 Break-out sessions

5.5.1 Vendors and Intellectual Property

Before the break-out sessions got kicked off, Robert Wisniewski (IBM) presented some thoughts on Co-Design and Intellectual Property (IP).

Co-Design and community development seen from a vendor:

- Community development rather than open source;
- Find largest common denominator on bottlenecks across applications;
- Application development by customers and model development by vendors is chicken-egg problem;
- Customers want vendor support, but need to be involved as well;

IP seen from a vendor:

- Some areas will remain proprietary for competitive differentiation;

- Specification of requirements, not solutions;
- Vendors most concerned about contamination;
- Open Source license must be vendor friendly;
- Takes a lot of time to resolve;

How to work together:
- General broad co-design team without NDA;
- Co-design performance teams are vendor specific with full NDA;
- Firewall between co-design performance teams;
- Generalized APIs between general broad design and co-design performance team,
- Make working agreements on this model;
- Vendors need to be fully engaged.

5.5.2 Software break-out group

Software needs application survey results based on actual usage, not on just installation. This should result in:
- Compact/skeleton applications with instructions on usage;
- Programming model details on performance, failures and reproducibility and correctness;
- Taxonomy of main algorithms

Software needs the following details from hardware:
- Fault model, detection, probabilities;
- Parallelism management;
- Relation between usage and power;
- Resource management knobs;
- Plans on solid state devices.

5.5.3 Hardware break-out group

Currently, hardware can provide:
- Abstract machine model, with various levels of information:
    - Hardware specifications, e.g. cache characteristics;
    - Scaling trends, sensitivity to different design parameters;
    - Models with the possibility to trade-off various parameters;
    - Simulators.

Information needed in the exascale hardware design process:
- Detailed specification of application bottlenecks;
- Computational characteristics;
- Memory details (minimum, bandwidth for what ?, NUMA sensitivity);
- Branch prediction, out-of-order execution, etc.
- Communication details, MPI usage;
- Utilization of exascale systems;
- The level/kind of resilience.

5.5.4 Application break-out group

Current observations on hardware and software by applications:
- 100-1000 threads per node;
- 1 or 2 orders of magnitude less memory per thread and less memory bandwidth per thread;
- Weak scaling through MPI, strong scaling through threads.

So, applications think of MPI+X as programming model:
- Still huge amount of recoding required;

- But in a familiar hybrid programming model.

Applications need to provide so-called exaskeletons to hardware and software teams. Exaskeletons should stress the hardware and software in a representative way:
- Quantitative ideas about relative rates, sizes, frequencies and granularities;
- Not just flops and loads/stores, but also branching frequency and sorts of random accesses.

Applications need to gain more insight in:
- Synchronization issues;
- Data locality;
- Arithmetic intensity;
- Precision;
- I/O.

Applications would like to hear more about:
- Also tools for development;
- Tools for dynamic processor allocation;
- Development environment.

5.5.5 Intermezzo: where do we agree

Considering the outcomes of sections 5.5.2 to 5.5.4, we can identify some areas where it is agreed what should be supplied by one and required by the other. These are:

- Application skeletons (“exaskeletons”), to be provided by applications and used by software and hardware;
- Details on “code flow details” (synchronization, data locality, branching, precision, …). Insight to be gained by applications, to be supplied to hardware and software;
- Based on “MPI+X” as the programming model, hardware to supply an abstract machine model with trade-off models on various parameters, to be used by software and applications with their exaskeletons.

5.5.6 Break-out group on implementation of co-design

The break-out group on co-design reconsidered the presented CDC approaches (sections 5.4.1 to 5.4.4). The common approach is:

- Establish common terminology between CDCs;
- If needed, bring coupled codes to the same level of quality;
- Design skeleton applications (repository and usage of new programming models, to be able to experiment with it). If possible, create skeleton codes with tools (ROSE);
- Design a communication structure, to advance in terminology and skeleton applications.

Challenges seen:
- Limited resources, so not all paths can be explored;
- Impact on hardware only during limited timeframe;
- Manage code stack for reproducibility and verification;
- Interaction with programming models people required.

Foreseen time schedule:
- 2012:
    - Common terminology in place;
    - Create skeleton applications and repository;
- To 2015:
    - Evaluate choices using skeletons;
    - Create performance and complexity models;

- To 2018:
    - Rewrite production codes based on experiences gained.

5.5.7 Break-out group on governance aspects

This break-out group first tried to define for what aspects international governance would be required:
- Exascale roadmap updates;
- Monitoring progress in roadmap implementation (review proposals to work and evaluate projects that implement software components);
- Share information, e.g. on co-design progress;
- Contribute to calls from funding agencies;
- Training, education and dissemination;

Open questions (some of those to be answered through the break-out group on international X-stack cooperation (section 5.5.8)):
- Should IESP involve industry ? To what level ?
- Should IESP actually deliver the software stack (which includes QA aspects) ?
- Should IESP steer on what software component is developed where and when ?
- Should IESP facilitate an international testbed for the software stack ?

This break-out group proposed a possible international governance model, for which the key words "joint" and "coordinated" should hold. In this model, there is an IESP-like Executive Committee, which oversees the aspects as defined above. To do this, each geographical region (Asia, Europe, US) sets up its own steering committee, which reports back to the Executive Committee.

5.5.8 Break-out group on X-stack international cooperation

This break-out group started with reporting on the European, Japanese and Chinese projects and funding models with respect to HPC and exascale computing. In Europe, lots of discussion is about the concept of "European IP", as there is not a European hardware industry at large. Further, Europe has to balance between European Framework funding and national initiatives. Japan has issued its first Post-Peta HPC initiatives, while China is currently finalising its next 5-year plan, which includes HPC and computational science developments.

With respect to international cooperation, the following recommendations were made:
- Formalise IESP (for instance through contents of section 5.5.7);
- Exchange information in a standardized way;
- Exascale roadmap:
    - Identify topics for incremental progress and for (expected) disruptive methodologies;
    - Set up a master plan;
    - Define specific exascale challenges to be tackled on national levels;
    - Define working groups which become problem owners of specific topics;
- Use gap analysis to inspire (or steer?) national funding initiatives;
- Open Source preferred by China, Japan and Europe;
- Broaden up G8-type of initiatives.

Together with governance aspects, a realistic operational model could be formulated.

5.6 Next meeting

The next meeting will be organised in Europe: Cologne, Germany, October 5-7, 2011.

5.7 Conclusions

During the first day actual updates were given on various topics:

- For EESI, the progress in especially the Working Groups on application grand challenges and the Working Groups on enabling technology for exaflop computing has been presented;
- For the US, a working model for the Exascale Software Center was presented, together with similar updates on the Co-Design Centers;
- For Japan, the scientific areas, matched upon the provisioned hardware, were presented;
- For China, application focus and enabling systems for that were presented;

During the second day, insight was gained on:

- IP and Co-Design from a vendor's point of view. An initial model for a broad co-design process was formulated, involving all parties (including vendors), and taking care of NDA-aspects in protected areas of the co-design process;
- Co-Design Centers presented their way of working (iterative);
- From the break-out sessions on hardware, software and applications, convergence to the following aspects have been reached:
    - Application skeletons ("exaskeletons"), to be provided by applications and used by software and hardware;
    - Details on "code flow details" (synchronization, data locality, branching, precision, …). Insight to be gained by applications, to be supplied to hardware and software;
    - Based on "MPI+X" as the programming model, hardware to supply an abstract machine model with trade-off models on various parameters, to be used by software and applications with their exaskeletons.
- The Co-Design center break-out group presented a time schedule for developing exascale-ready applications, in whicg skeleton applications are of key importance;
- The break-out group on international governance proposed a possible international governance model, for which the key words "joint" and "coordinated" should hold. Some open questions remain, in particular on the steering of software development and delivery of the software stack – suggestions for answers by international cooperation;
- With respect to international cooperation, some ideas on ensuring coordinated work on the X-stack have been raised. Gap analysis (what is there and what is still lacking) is expected to be a vehicle to inspire or steer national funding initiatives.

6. Workshop in Cologne, Germany

6.1 Introduction

The seventh IESP workshop has been organised in Cologne, Germany, on October 6-7, 2011. At the start of this workshop, much time has been reserved for the progress in Europe, while one of the break-out groups was dedicated to coming up with a distribution of software components over regions (US, Japan, China, Russia and Europe) to work on.

The first day of the workshop was mainly devoted to updates on the current situation and plans in the various countries/regions. The EESI project has presented their findings, recommendations and conclusions from a European perspective, spread over day 1 and day 2. Break-out sessions were devoted to co-design, software development and maintenance, and to revolutionary technologies, also during day 1 and day 2.

One of the explicit goals of this workshop was to get information on paper, which should include the efforts in the various countries/regions and the plans coming from the three break-out groups.

We have structured this chapter as follows. Section 6.2 covers the updates from the various countries/regions, including information from the cartography efforts as carried out in the EESI project. The need for education and training will be covered here as well. Section 6.3 is devoted to the EESI-project, which will end by November 30, 2011. We will also report here on the EC plans and the recently granted exascale projects CRESTA, DEEP and Mont Blanc. Section 6.4 will cover the results of the break-out groups.

As an introduction to the workshop, Thomas Lippert in detail described a few scientific breakthroughs which would have been impossible without simulation on large HPC-systems. As an example of his own institute, he mentioned the 1 Pflop/s IBM BG/P system, in use at FZJ since 2009, including FZJ's dual concept approach: BlueGene/P for extreme scaling applications, cluster systems for other applications. With a total of around 300k cores, it is clear that extreme parallelisation of the application code is required to achieve significant speed-ups. This can be explained by Amdahl's law and its variants and derivatives. The approach of Gustafson in 1988 was such that the runtime is the most important issue for a computational scientist: if the problem size can be increased by applying more core/parallelism, and when the runtime remains constant, the scientist is optimally productive. Taking this concept into a dual hardware approach has lead to the DEEP project, which combines a regular cluster through InfiniBand with a so-called booster cluster, which will be used to offload highly scalable kernels. More information on DEEP will be provided in section 6.3.1. Lippert mentioned energy consumption as one of the major challenges for exascale, and supposed to learn from the human brain, both with respect to energy usage as well as with respect to applying parallelism.

6.2 Global Updates

This section gives a brief overview on the current state of affairs on HPC installations and initiatives, targeting towards multi-petascale and exascale systems, on a global basis.

6.2.1 Progress in China

Zhong Jin (Chinese Academy of Sciences) presented an overview of the Co-design initiatives in China and the plans in China towards beyond Petaflop/s activities. His analysis of China's current situation was high hardware performance but low application scaling. In order to increase the application scaling, funding has been raised to increase code scaling up to 300k cores, with a targeted parallel

efficiency of at least 30% for eight selected application areas: fusion, aircraft design, spaceship design, drug design, animation, mechanics, environment and material design.

On the hardware side, several funding agencies (MOST, CAS, NSFC, Ministries and local government) plan multi-petaflop/s systems, with the top system the system meant for earth system simulation, targeting at 10+ Pflop/s with Godson CPU cores. Application of co-design for these systems will be adopted.

6.2.2 Progress in Japan

In June 2011, the K-computer came on position 1 of the top500 list, with more than 10 Pflop/s peak performance and more than 8 Pflop/s Linpack performance, for 20 MW of power. Focus so far has been on selected application, for which scaling improvement is going through several planned phases. Each of the phases represents a certain level of parallel efficiency of the applications. The stability of the K-system so far has been very good. Further, experiments have been carried out with optimising the interconnect network between the nodes, with the ultimate goal of improved scaling of applications. Current bottlenecks includes C/C++ compiler optimisation, communication library performance and certain numerical libraries (like FFTs).

With respect to the future, a white paper on Strategic Direction/Development of HPC is being prepared (by young researchers), which will lead to calls for proposals for “Feasibility Study of Advanced High Performance Computing”, which address Japanese international competitiveness and national security. It is expected that several computer architectures will be proposed, which will call for further funding by the Japanese government. The white paper is scheduled to be done by the end of 2011, with the start of the selected projects from the feasibility study by the beginning of 2014. In parallel, from 2011, basic research programs for development of software technologies for post-peta scale HPC will be run by the Japan Science and Technology agency (JST).

6.2.3 Progress in the US

Paul Messina gave an overview of the currently funded projects in the US, including some insight in which plans were to be developed. About a year ago, several projects on exascale have been granted, these have been listed already in section 4.2.1. In addition to those, three out of seven proposals for co-design centers have been granted: advanced nuclear reactors, combustion and materials in extreme environments. Details on their activities and workplans can be found in section 5.4.

Unfortunately, not yet funded have been the plans for establishing the proposed Exascale Software Center, as described already in section 5.3.2. On the other hand, the US Senate has explicitly mentioned funding levels for exascale computing of 126 M$. New plans will have to be proposed before February 10, 2012, which include interim milestones towards the exascale target. In order to prepare these plans as good as possible, seven DOE national laboratories (Argonne, Lawrence Berkeley, Lawrence Livermore, Los Alamos, Oak Ridge, Sandia and Pacific Northwest) have formed a consortium to manage a Request-for-Information (RfI) process to deliver exascale computing capability. This information will be provided to the DOE, which has to respond to the US Senate committee.

Among others, it is expected that this information will consist of exascale system goals, and potential partnership models between laboratories and industry. A total of 22 responses have been received, on integrated systems, microprocessors, network and interconnect, memory and solid state, file system and storage and software. Power is seen as a major obstacle, while data management issues were not yet thoroughly investigated.

Apart from the DOE, also the NSF deals with exascale computing. A task force has recommended to investigate new models for partnerships between industry, academia and other agencies, as well as taking into account huge amounts of data in a digital data framework.

6.2.4 Cartography results

Mark Parsons (EPCC) presented cartography results with respect to HPC systems and plans on a global scale. The corresponding report has been delivered as one of EESI deliverables, with authors Xu GUO (EPCC), Damien Lecarpentier (CSC), Per Oster (CSC), Mark Parsons (EPCC), Lorna Smith (EPCC). The actual deliverable is D2.1 and available on www.eesi-project.eu.

6.2.5 Training and education

Vladimir Voevodin (Moscow State University) expressed the need for education and training in the HPC arena. One of the interesting analysis aspects was the fact that students who start their education now on average will be in the Ph.D. stadium in 6 years from now, which is very close to the expected date of the first availability of exascale systems. In other words, when starting now, we can only expect in 2018 students who had HPC classes during their studies.

Voevodin further suggested to include an additional chapter in the IESP roadmap: “Perspectives on Cooperation between IESP and University Communities”, to start with some cartography on current initiatives on HPC education worldwide, finally resulting in a set of recommendations to include HPC curricula in bachelor/master/PhD programs.

6.3 European efforts

Leonardo Flores Anover from the EC presented the current EC observations, questions and potential answers for a European strategy on exascale computing. One of the actual results of this is the fact that three exascale projects have been granted by the EC. Both the EC views and the exascale projects will be discussed in section 6.3.1. Section 6.3.2 will cover the results and recommendations of the EESI project, as they will (have) be (been) presented to the EC and on the final EESI conference in Barcelona. This will be done in a brief and structured way (per Working Group – WG).

6.3.1 EC plans and exascale projects

Recent IDC studies have shown that Europe lags behind the US, Japan and China with respect to actually installed HPC systems. Many other observations are related to this: Japan is larger than all of Europe together, there is fragmentation in European HPC efforts across many countries, and Europe’s Intellectual Property Rights (IPR) are of benefit to others. Recently, the EC has expressed its commitment to support HPC, which has resulted into (among other) PRACE and recent exascale projects.

The key question for the EC is whether Europe should be a global HPC leader, excelling in the application and production of HPC, in all domains (industry, science and society, or a follower, focussing on applications on foreign systems. Apart from the answer to this question, the HPC policy should be European. Although not completely developed yet, the EC is investigating EU native capabilities by its pre-commercial procurement in PRACE and its granting of exascale proposals, which include prototype development of promising exascale hardware.

With respect to answering the key question, many sub-questions arise: what should be done, who are the key players in Europe, why should Europe care, how much will it cost, what will be the benefit, by when should this be done, etc. The European stakeholders seem to answer the question quite clear: there is no choice: a European-wide effort must be engaged to develop autonomous technology (covering the whole spectrum from processor architecture to applications) to build exascale systems in about 10 years from now. Furthermore, action to reach this is required now, according to those stakeholders.

The EC will have to make up his mind on this, in order to, as to speak with Donna Crawford, become self-reliant, capable, independent and responsible for Europe's own future in economic and societal affairs.

Three European exascale projects have been started on either September 1, 2011 or October 1, 2011. These are: CRESTA, DEEP and Mont-Blanc. Each of these will be discussed very briefly.

CRESTA stands for Collaborative Research into Exascale Systemware, Tools and Applications. It has been presented by Mark Parsons from EPCC, with more details at www.cresta-project.eu. The focus of CRESTA lies predominantly in software, not hardware. The project will run for three years, on a total of 12 M€, with 8.57 M€ funding. Co-design is at the heart of the project, with preselected European applications at the core. Both incremental and disruptive approaches on the combination of applications and so-called systemware (the full range of software components needed to exploit future exascale platforms) will be investigated in CRESTA.

DEEP stands for Dynamical Exascale Entry Platform. It has been presented by Norbert Eicker from FZJ. The project will run for three years, on a total of 18.5 M€, with 8.03 M€ funding. DEEP aims to design and develop a hardware architecture which will lead to exascale systems. The concept is to combine a regular cluster through InfiniBand with a so-called booster cluster, which will be used to offload highly scalable kernels. The current idea is to use the Intel MIC architrecture for the Booster nodes, connecting them with the EXTOLL communication network. A set of scientific applications has been selected to be implemented on the new system. Together with energy-aware integration of components, programming environments and tools, the DEEP project aims for a full combination of exascale-promising hardware and software.

The Mont-Blanc project has been presented by Alex Ramirez from BSC, with more information at www.montblanc-project.eu. The project will run for three years, on a total of 14.5 M€, with 8.1 M€ funding. Mont-Blanc aims to develop a prototype HPC system based on currently available energy-efficient embedded technology, to design a next-generation HPC system relying on new embedded technologies, with a small set of representative exascale applications running on the system(s). Energy-efficient building blocks will be used, with hybrid MPI+OmpSs as the programming model. Applications that will be targeted will be recruited from those currently running on PRACE Tier-0 systems.

6.3.2 EESI findings, recommendations and conclusions

During this IESP workshop, a significant amount of time has been reserved for the EESI Working Groups (WGs) to present their findings and recommendations. In total, there have been four WGs on scientific disciplines and four WGs on general hardware and software aspects. Apart from these WG reporting, the EESI project leader has been presenting the general EESI-recommendations, as they will be done to the EC.

Since all reports will be available at the EESI project website (www.eesi-project.eu), we will limit ourselves here to listing the key results per WG and for EESI as a whole.

WG3.1: Industrial Applications: Energy and Transportation

Presented by Jean-Claude Andre, former CERFACS. Chair and vice-chair: Philippe Ricoux (Total) and Jean-Claude Andre.

Directions of research required for efficient usage of exascale technologies:

- At the level of the simulation environment:
    - Unified simulation frameworks, especially with respect to pre- and post-processing;
    - Multi-physics simulations and coupling interfaces;

  - More advanced mesh generation tools, including automatic and adaptive meshing, scalable;
  - Standardization of I/O and data management techniques;

- At the level of application codes:
  - Further development of new numerical methods, algorithms, solvers/libraries, etc.
  - Coupling between stochastic and deterministic models;
  - Uncertainty and risk quantification;
  - Meshless methods and particle simulation;
  - Scalability research on all levels;

- General remarks:
  - Large number of applications in production environment may be “farming” applications;
  - The real goal is zetascale, not exascale;
  - 2012-2020 timeframe: 400 M€ for manpower required.

WG3.2: Weather, Climatology and solid Earth Sciences

Presented by Giovanni Aloisio, CMCC & University of Salento, Italy. Chair and vice-chair: Giovanni Aloisio.

Directions of research required for efficient usage of exascale technologies:

- Efficient and scaling coupling frameworks required between various applications;
- Further investigation of overlapping communication and computation, exploitation of data parallelism to improve scalability up tens of thousands of cores;
- Investigation of different grid systems (i.e. mesh generation);
- Exploration, analysis and manipulation of large datasets, more standardization, better parallel strategies, interactive visualisation, including development of efficient I/O libraries and data compression aspects;
- Quality control of numerical output;
- 2012-2020 timeframe: 300-400 M€ for manpower required;

WG3.3: Fundamental Sciences

Presented by Stephane Requena, GENCI, France. Chair and vice-chair: Godehard Sutmann (FZJ), Jean-Philippe Nomine (CEA).

European strengths on fundamental sciences application codes:

- Many high-quality simulation codes have European origin and still large European development;
- Highly developed methods and algorithms;
- Well-organised communities;

Multiscale modelling is a weakness.

Co-design and European support centers are required to further stimulate exascale application efficiency:

- 7 to 10 domains, 10 fte’s in co-design, 10 fte’s in support, leads to 140-200 fte’s per year;
- 2012-2020 timeframe: 110-160 M€.

WG3.4: Life Sciences and Health

Presented by Ramon Goni, BSC, Spain. Chair/vice-chair: Modesto Orozco/Janet Thornton.

European strengths:

- Elixir, Bioinformatics community;
- EMBL-EBI;
- Research consortia: ENCODE, ICGC, iHEC;
- FET Flagships: personalized medicine, human brain project;
- PRACE;
- Pharmaceutical industry;
- Computational Biology and Bioinformatics Software.

European weaknesses:

- Competitors with privileged access to specialised resources;
- No coherent unified data layer available for bio-data;
- Crucial databases outside Europe;
- Bio-HT techniques are not European;
- Bio-community is separated from the HPC-community;

Main conclusions:

- Not only Exaflop/s are required, exascale as a whole;
- Lack of HPC resources will kill several of the most important projects in Life Sciences;
- Co-development programs to deploy exascale resources in the Life Sciences;
- 2012-2020 timeframe: 200 M€.

WG4.1: Hardware roadmap, links and vendors

Presented by Bernd Mohr, FZJ. Chair and vice-chair: Herbert Huber (LRZ) and Riccardo Brunino, CINECA.

Major deficiencies in current HPC technology providers R&D plans:

- **Future many-Peta to Exascale file systems**
  - Metadata performance;
  - Scalability and reliability;
    - Number of concurrent clients;
    - Number of files;
    - File system capacity
  - Data protection and recovery mechanisms
    - Distributed RAID mechanisms;
    - Analysis and diagnostics database
  - Data mining and visualization tools
- **Future I/O**
  - Mechanisms for end-to-end data integrity;
  - High speed SERDES technology
- **Resilience and fault tolerance**
  - Fault tolerant parallel runtime systems;
  - Mechanisms for error prediction, automatic error detection and recovery
- **Fundamental Research**
  - Novel semiconducting materials;
  - Novel memory technologies;
  - Novel background storage technologies;
  - Integration of optical technology into the standard CMOS process
- **Communication networks with in-interconnect processing capabilities**
- **Operating system enhancements**
  - Coherent inter-node scheduling mechanisms;

  - Micro OS;
  - Data locality
- **Numerical libraries and highly scalable performance analysis tools**

The following opportunities for European HPC Technology providers are recognized:

- Green IT: cooling, energy-efficient data centre infrastructures;
- Data handling and file I/O;
- Fault tolerant runtime environments.

In fact, Europe still has all the necessary knowledge at hand to develop its own HPC technology stack.

WG4.2: Software Ecosystem

Presented by Franck Cappello, INRIA, France and University of Illinois, US. Chair and vice-chair: Franck Cappello and Bernd Mohr, FZJ, Germany.

The software ecosystem covers development environments, tools, system software and some crosscutting dimensions, like resilience and power management. It is clear that exascale challenges cannot be addressed only at the hardware level, the software level will need to be involved too. The most important of those challenges are scalability, heterogeneity, energy, resilience and performance. The following key findings on the software ecosystem has been found:

- Most HPC system software components have to be **substantially adapted or newly developed** to address Exascale challenges
  - A 1000 MY investment in system software development is necessary: 100 persons during 10 years.
- Application developers (in research and in industry) are reluctant to adopt new, not yet fully accepted and standardized programming models
  - **A transition path for HPC applications** to Exascale programming models is required
- The different Exascale system architectures should not be visible at the level of the programming model
  - **The SW ecosystem must be portable** to support various Exascale hardware
- Collaboration between HPC hardware, system vendors, European R&D laboratories and Academic researchers has to be ensured
  - **Co-design processes must be established**
- Significant funding is essential in R&D areas where Europe has technology leadership and critical mass: programming models & runtime, performance, validation and correctness tools.
  - The key players should form alliances to define standards.
- Establish cooperation programs with technology leaders in areas where European technology is not leading: I/O, File systems, OS, etc.
  - To ensure that system components and standard can be influenced from a European perspective
  - To revitalize research in these important areas in Europe
- Consistency and completeness of the software stack as well as portability, robustness, and proper documentation have to be ensures
  - **European Exascale Software Centre** should be established to coordinate research and development of HPC Exascale software ecosystem components developed in National and EU projects.

Another important aspect is the creation of a testbed for developed software from 2015 and beyond. Large hardware facilities are required to test those software aspects which only become visible beyond 100 Pflop/s of peak performance.

An exascale software stack is the ultimate result of these activities. An important organizational issue is the cooperation model – would there be one exascale stack per country ? Who is doing what ? During the breakout session, this topic has been further addressed.

WG4.3: Numerical libraries, solvers and algorithms

Presented by Iain Duff, CERFACS. Chair and vice-chair: Iain Duff and Andreas Grothey, University of Edinburgh, UK.

This group has covered the following areas:

- Dense linear algebra
- Graph and hypergraph partitioning
- Sparse direct methods
- Iterative methods for sparse matrices
- Eigenvalue problems, model reduction
- Optimization
- Control of complex systems
- Structured and unstructured grids

Typical use of library software is hierarchical, so several layers of numerical work, which could be approached by several levels of parallelism. Typical performance and power issues that arise are the costs of data movement. General issues that have to be covered are the lifetime and support of numerical software, floating point issues, programmability, performance optimization, fault tolerance, uncertainty quantification and training.

Apart from technical challenges, there are also organizational challenges: the risk is that research groups work on their own. As an example organization, the Trilinos project may serve, which is based on a framework that federates multiple independent subprojects.

With respect to European activities, it should be recognized that many existing libraries have been developed in Europe or in international projects with significant European input. A funding model that secures long term funding should be in place, including funding for specific research challenges, networking and training and an Exascale Centre would be needed. Total costs expected are around 2000 person years of effort, which adds up to 200 M€ over the period 2012-2030. A co-design centre with extensive visitor facilities will facilitate required collaboration and contacts between the global research groups.

WG4.4: Scientific Software Engineering

Presented by Andrew Jones, NAG, UK. Chair and vice-chair: Mike Ashworth, STCF, UK and Andrew Jones, NAG, UK.

The key objective of this working group is to bring together the scientific and implementation aspects in reliable and sustainable applications codes, ready for exascale execution. The specific themes that have been identified were:

- Application frameworks & workflows
    - End-to-end solutions; coupled models; etc.
- Visualisation & data management
    - Data issues; in-situ visualisation; etc.
- Fault tolerant algorithms
    - Build on API from OS; interface with WG4.2
- Application design
    - Build on prog lang & env; interface with WG4.2
- Software engineering

- Software quality; collaborative working

WG4.4 has formulated a total of 11 recommendations:

- Development of Domain Specific Languages (DSL);
- Development of open standards in coupling software and data exchange;
- Integration of coupling and workflow technologies;
- Development of advanced mesh generation, mesh partitioning and load balancing software;
- Development of a flexible generic I/O layer;
- Development of advanced computational steering technologies;
- Development of techniques for local checkpoint/restart to fast memory;
- Development of an API for error signaling and reporting;
- Development of fault tolerant applications;
- Development of co-design structures in Europe;
- Development of advanced integrated development environments.

The expected total costs (mainly people) is estimated at around 200 M€ for the period 2012-2020.

Europe is strong with respect to open source software and with respect to organization in multi-disciplinary teams, but is not strong in industry commitment and vendor involvement (leading to limited prototype hardware access). With education and training, the European strengths should be maintained and extended with young researchers

General recommendations from EESI

Presented by Jean-Yves Berthou, EDF, France. Project coordinator of EESI.

In his presentation, Berthou has collected the recommendations of the Working Groups. The EESI recommendations are targeted to the funding strategy and organisation, which projects and activities to fund, the required budget and to the European and international governance. As a general recommendation, it is clear that if the recommendations are followed, the effort needs to be sustainable, long term and coordinated. Only then, maximum benefit can be taken from the large investments.

EESI estimates the exacsale challenge to be comparable large international research projects as LHC and ITER, requiring a cross-thematic European program dedicated to exascale, involving several EC departments (DG_INFSO, DG_CLIMA, DG_ENERGY, DG_RESEARCH&INNOVATION, …).

Based on the detailed recommendations, as will be mentioned in detail in the final EESI report, the following main conclusions and recommendations towards the EC can be made:

- Europe needs a **sustainable, long term and coordinated** effort;
- Europe is still well positioned to be part of the few players worldwide deploying and exploiting Exascale technology but action is needed now;
- A 2,5 to 3,5 billions euros (1000 M€) total budget over 10 years, supported by EC, National European funding agencies, industry, … **a several decades** engagement;
- Scientific Computing at Exascale, from a computing and data intensive point of view, are **strategic** for maintaining and developing both **European Scientific Excellence and Industry Competitiveness;**
- **International** collaboration is required;
- Europe should encourage the development of **Open Source solutions** to foster international collaborations and the emergence of international *de facto* standards, enabling commercial exploitation.

6.4 Report on break-out groups

During the wokshop, three parallel sessions have been organised, on co-design, on system software and on revolutionary technologies. Sections 6.4.1 to 6.4.3 report briefly on the discussions.

6.4.1 Break-out group on co-design

The general idea on co-design of hardware, system software and applications is rather well-supported. Several initiatives in the US, Japan, Europe, Russia and China are underway, in various stages of development now. In order to get to exascale performance, application input is needed to develop suitable hardware and software, in a continuous circuit of feedback and adaptations. One of the key questions discussed was the question on how to really influence the design of microprocessors.

From a vendors' point of view, the anticipated microprocessors for 2018 can still be influenced in 2012 and 2013, after which the window of opportunity is away. For architectures and software, it is generally believed that this window is larger, up to 2 to 3 years before actual release. The application vehicle would be ranging from kernels to full applications (with skeleton and compact applications in between).

6.4.2 Break-out group on system software

During this break-out session an attempt has been made to identify the components of the software stack and to distribute the work over the regions. Before doing that, three categories have been defined:

1. For software in this category, a group commits to provide in production quality;
2. Software in this category is seen as experimental;
3. Software in this category is not funded, will rely on others (vendors or other regions/countries).

Software in category 1 is essential and hence it is needed to be supported. This can be done by developers, vendors or open source.

This has led to a tentative distribution of efforts for those regions present at the break-out sessions. In general, each of the regions views programming models/languages, OS kernels, communication libraries as category 1 software, and is able to work on that.

From a gap analysis, components as low level compilers, RAS and system management, batch schedulers and limited power management/fault tolerance do not seem to be covered yet in category 1.

More detailed work distribution will be needed, including organisational aspects.

6.4.3 Break-out group in revolutionary technologies

The break-out group on revolutionary technologies was lead by Thomas Sterling from Indiana University and Bronis de Supinski from LLNL. Quite a few interested people attended the break-out session. As a summary, almost each hardware/software/applications component of current Petaflop/s systems is a candidate for revolutionary approaches, as the challenges to get to exascale are large. More detailed, the break-out group discussed the following opportunities for revolutionary progress:

- Execution model
    - Addresses: Starvation, Latency, Overhead, and Waiting for resolution of contention
    - Impacts co-design across programming models, architectures,

  - Not clear what right one is, some people don't know what it is (informational graph)
  - Come from semantics of application information graph, not just hardware side
- Changing the way we think about system and programming merger
  - Dynamic behavior including aspects of non-determinacy, variability
  - Can't control the *horizontal* or *vertical* – we don't control the machine
  - Change of culture, attitude, not just bottom up – impacts algorithms
- Intelligence
  - Decision making at multiple levels, the decision chain
  - On the fly, control choices and data analysis
  - Spatially aware from data perspective requires hardware support for operation knowledge, predictive
  - Information backplane, FT backplane – protocol between system layers for mutual dynamic introspective behavior
- OS
  - Optimization point for programming model and the architecture
  - Reactive: Driven by revolutions in other areas, can't keep old OS w new architecture
- Programming models
  - Levels of abstraction
    - Where are we falling down?
      - Inertia of comfort zone, familiarity breads complacency
    - Semantics of parallelism
      - Control of sequencing (or not)
      - Forms of parallelism, synchronization constructs
    - Relationship of data structure and control
      - One sided access
      - Dependency for automatic control and data move
  - MPI
    - Provides locality control, permits human intervention of placement and movement of data
    - Ease of reasoning about program execution (or not)
    - Legacy of codes and skill sets
    - Don't want to know where everything is – suffering from tsunami of complexity
  - Intermix programming/execution modes
    - "Right tool for the right job"
- Correctness and debugging
  - Tool for verification and boundedness for quasi reproducibility
  - Program itself knows that it is correct – how does it know what correct is?

Finally, the break-out group addressed the way to incorporate revolutionary approaches in the current way of working and collaborating. This has given the following observations/recommendations:

- Managing revolution
  - Gradual steps, Otherwise cannot plan, Is there a concept of an "incremental revolution"
  - Starting with something credible
  - Don't throw out everything
  - Education and culture change
  - Revolution is being imposed, even for next generation machines
- Can we identify destination before choosing the road
  - Key is to determine the right answer and then worry about how we get there…
  - Often can take a quantum leap or a gradual path to the same destination
- Where do we put intelligence in the systems
  - Revolutionizing the libraries
  - Changing algorithms
- Worry about reliance on compiler technology
  - Abstractions could be too high, doesn't convey to scientific computing
- Persistent storage – a hierarchy that is likely to change; the revolution is coming
  - Out of core calculations

The conclusion from this break-out group is that revolutionary approaches may need to be adopted, but in a controlled way if possible.

6.5 Next meeting

The next meeting is currently being planned and has been preliminary scheduled for Kobe, Japan, on 12-13 April 2012.

6.6 Conclusions

The following conclusions from the Cologne workshop can be drawn:

- The US, China and Japan have clear intentions and roadmaps on building and installing 100+ Petaflop/s and Exaflop/s systems in the decade ahead. This is not (yet) the case for Europe;
- The window of opportunity in the co-design process for influencing hardware development, however, is relatively short and is 3-5 years before actual hardware availability. Skeleton applications have been viewed as enough for the co-design process, although new views on this may lead to a combination of kernels and applications itself. The original plans for Co-Design Centers (CDC) in the USA have been reduced, at least temporarily, so that only a few CDCs have been able to make a start in 2011. Japan and China will adopt co-design as well, and, as in the USA, will focus on predefined application areas;
- The USA plans for establishment of an Exascale Software Center (ESC) have been delayed, but not abandoned. New plans through the DoE for US Congress are due in February 2012, which are currently being prepared by the community in the USA. Requests-for-Information have been issued to both US hardware and software vendors, which will be taken into account when drafting the plans;
- The EESI project has been able to address a huge amount of knowledge through engaging European experts in both scientific, hardware, software, libraries and software engineering fields. This has led to findings and recommendations to the EC, including cost estimates over the period 2012-2020, as presented in Cologne and during the final EESI conference in Barcelona in October 2011. Development of exascale hardware is viewed as a serious option for Europe as well;
- The EC has granted three exascale proposals each with around 8 M€ funding, to develop relevant exascale technology in the hardware, software and applications area. Two of the projects aim to develop prototype hardware, which aim to set the road to exascale, with relevant software components and application vehicles to test. The third project focuses on European-developed and important applications to apply software technology to prepare for eventual exascale execution;
- Initial inventories for X-stack development have been made by each of the regions (US, Europe, China, Japan) on which software components could be addressed in which region, with varying levels of commitment. Gap analysis will be needed to cover the full range of what is needed. Funding commitments over the regions will have to be made;
- Revolutionary approaches are worthwhile to invest in several areas – if possible in a controlled manner.

7. References

[1] *The International Exascale Software Project: A Call to Cooperative Action by the Global High Performance Community*, Dongarra, J., Beckman, P., Aerts, P., Cappello, F., Lippert, T., Matsuoka, S., Messina, P., Moore, T., Stevens, R., Trefethen, A., Valero, M. Volume 23, Number 4, Winter 2009, *International Journal of High Performance Computer Applications*, pp 309-322, ISSN 1094-3420.

[2] International Exascale Software Project Roadmap, version 1.1, www.exascale.org.

8. Appendix A – Attendees of the IESP workshops

8.1 Attendees of previous Workshops

SC08, Austin, TX, USA, November 17-20, 2008
Santa Fe, NM, USA, April 7-8, 2009
Paris, France, June 28-29, 2009
Tsukuba, Japan, October 19-21, 2009
Oxford, UK, April 12-14, 2010

		Affiliation	Country	SC08	Sante Fe	Paris	Japan	Oxford
Patrick	Aerts	NWO	NL		x	x	x	x
Giovanni	Aloisio	Euro-Mediterranean Centre for Climate Change	Italy			x	x	x
Dong	Ahn	LLNL	US	x				
Yutaka	Akiyama	Tokyo Tech	Japan				x	
Jean-Claude	Andre	CERFACS	France			x	x	x
Phil	Andrews	UT	US	x				
Mutsumi	Aoyagi	U Kyushu	Japan			x	x	
Mike	Ashworth	Daresbury	UK			x	x	x
Franck	Barbier	ANR	France			x		
Venkat	Balaji	Princeton	US					X
David	Barkai	Intel	US	x	x	x	x	x
Sanzio	Bassini	CINECA	Italy			x		
Kyriakos	Baxevanidis ●[3]	EC	EU			x		x
Pete	Beckman	ANL	US	x	x	x	x	x
Jean-Yves	Berthou	EDF	France	x	x	x	x	x
Richard	Blake	Daresbury	UK		x			x
Jay	Boisseau	TACC	US	x				
Taisuke	Boku	U of Tsukuba	Japan		x	x	x	x
Bertrand	Braunschweig	ANR	France		x	x	x	
Bill	Camp	Intel	US		x			
Franck	Cappello	INRIA	France	x	x	x	x	x
Charlie	Catlett	ANL	US					x
Ruay-Shiung	Chang	National Dong Hwa University	Taiwan					x
Barbara	Chapman	U of Houston	US		x	x	x	x
Xuebin	Chi	CAS	China				x	x
Alok	Choudhary	NWU	US		x	x	x	x
Iris	Christadler	LRZ	Germany			x		
Almadena	Chtchelkanova	NSF	US		x			

[3] Attendees marked with "•" are official observers.

			Affiliation	Country	SC08	Sante Fe	Paris	Japan	Oxford
Guillaume	Colin de Verdière		CEA	France			x		
Frederica	Darema		NSF	US		x			
Bronis	de Supinski		LLNL	US	x				
Peter	Coveney		U College London	UK					x
David	Dean	●	ORNL/DOE	* US				x	x
Mik	Dewer		NAG	UK					x
Jack	Dongarra		U of Tennessee	US	x	x	x	x	x
Sudip	Dosanjh		SNL	US		x	x	x	x
Thom	Dunning		NCSA	US	x				
Iain	Duff		Rutherford Lab	UK					x
Hugo	Falter		ParTec	Germany		x	x	x	
Teresa	Finchum		U of Tennessee	US				x	x
Leonardo	Flores	●	EC	EU					x
Fabrizio	Gagliardi		Microsoft	US			x		
Alan	Gara		IBM	US		x	x		
Al	Geist		ORNL	US		x			
Luc	Giraud		CERFACS	France		x	x		
Kostas	Glinos	●	EC	EU			x		x
Jean	Gonnord		CEA	France	x		x		x
Robert	Graybill		ISI	US		x			
Bill	Gropp		UIUC	US	x		x	x	x
Jim	Hack		ORNL	US	x				
Jean-Francois	Hamelin		EDF	France		x		x	
Robert	Harrison		ORNL	US				x	x
Bill	Harrod		DAPRA	US	x				
Stefan	Heinzel		Max Planck DEISA	Germany		x	x	x	x
Barb	Helland	●	OS	US	x				x
Mike	Heroux		Sandia	US	x	x	x	x	x
Ryutaro	Himeno		RIKEN	Japan		x		x	x
Kimihiko	Hirao		Riken	Japan				x	
Dan	Hitchcock		OS	US	x				
Thuc	Hoang		NNSA	US		x			
Adolfy	Hoisie		LANL	US				x	
Charlie	Holland		DARPA	US	x				
Koh	Hotta		Fujitsu	Japan				x	x
Herbert	Huber		LRZ	Germany					x
Yuichi	Inoue		MEXT	Japan				x	
Yutaka	Ishikawa		U of Tokyo	Japan		x	x	x	x
Satoshi	Itoh		MEXT	Japan				x	
William	Jalby		U of Versailles	France		x			
Jean-Pascal	Jégu		Teratec	France			x		
Zhong	Jin		CAS	China				x	x
Fred	Johnson		DOE	US	x		x	x	x

		Affiliation	Country	SC08	Sante Fe	Paris	Japan	Oxford
Andrew	Jones	NAG	UK			x	x	
Emma	Jones	EPSRC	UK					x
Laxmilkant	Kale	UIUC	US			x		
Richard	Kenway	EPCC	UK		x		x	x
David	Keyes	Columbia U.	US		x	x		x
Moe	Khaleel	PPNL	US				x	
Kimmo	Koski	CSC	Finland			x		
Bill	Kramer	NCSA	US		x	x	x	x
Dimitri	Kusnezov	NNSA	US	x				
Jesus	Labarta	BSC	Spain		x	x	x	x
Jean-Francois	Lavignon	Bull	France		x	x	x	x
Alain	Lichnewsky	Genci	France		x	x		x
Volker	Lindenstruth	Heidelberg U	Germany			x		
Thomas	Lippert	Juelich	Germany	x		x	x	x
Bob	Lucas	ISI	US		x	x		x
Rusty	Lusk	ANL	US					x
Barney	Maccabe	ORNL	US		x	x	x	x
Satoshi	Matsuoka	TiTech	Japan	x	x	x	x	x
Simon	McIntosh-Smith	U. Bristol	UK					x
Bob	Meisner	NNAS	US	x				
Paul	Messina	ANL	US	x		x	x	
Peter	Michielse	NWO	NL		x	x		x
Kazunori	Mikami	Cray	Japan				x	
Leighanne	Mills	U of Tennessee	US				x	
Bernd	Mohr	Juelich	Germany		x	x	x	x
Terry	Moore	U of Tennessee	US	x	x	x	x	x
Hervé	Mouren	Teratec	France			x		
Jean-Michel	Muller	CNRS	France			x		
Matthias	Müller	Dresden	Germany				x	
Wolfgang	Nagel	Dresden	Germany		x	x	x	x
Kengo	Nakajima	U of Tokyo	Japan				x	x
Hiroshi	Nakashima	Kyoto U.	Japan			x	x	x
Mamoru	Nakono	Cray	Japan				X	
Jeff	Nichols	ORNL	US		x		x	x
Jane	Nicholson	EPSRC	UK				x	x
Jean-Philippe	Nominé	CEA	France			x		x
Nick	Nystrom	PSC	US		x			
Per	Oster	CSC	Finland	x	x			
Mike	Papka	ANL	US					x
Abani	Patra	NSF	US		x	x	x	
Rob	Pennington ●	NSF	US	x	x			x
Serge	Petiton	CNRS	France			x		x
Claude	Puech	INRIA	France		x	x		

			Affiliation	Country	SC08	Sante Fe	Paris	Japan	Oxford
Tracy	Rafferty		U of Tennessee	US	x	x	x	x	x
Dan	Reed		Microsoft	US		x	x		
Michael	Resch		HLRS Stuttgart	Germany		x			x
Catherine	Rivière		GENCI	France		x	x		
Ralph	Roskies		PSC	US	x				
Faith	Ruppert		ANL	US		x			
Christian	Saguez		Teratec	France			x		
Vivek	Sarkar		Rice	US		x			
Stef	Salvini		Oxford	UK					x
Mitsuhisa	Sato		U of Tsukuba	Japan	x	x	x	x	x
Stephen	Scott		ORNL	US		x			
Mark	Seager		LLNL	US		x			
Ed	Seidel		NSF	US	x		x	x	x
Akiyuki	Seki		MEXT	Japan				x	
Satoshi	Sekiguchi		AIST/METI	Japan				x	
Hideo	Sekino		Toyohash Inst Tech	Japan				x	
John	Shalf		LBNL	US			x	x	
Horst	Simon		LBNL	US	x	x			
David	Skinner		LBNL	US		x	x	x	x
Marc	Snir		UIUC	US	x				
Mary	Spada		ANL	US	x				
Thomas	Sterling		LSU	US		x	x	x	
James	Southern		Fujitsu	UK					x
Rick	Stevens		ANL	US	x	x		x	x
Michael	Strayer	●	DOE OS	US	x	x			
Fred	Streitz		LLNL	US				x	x
Bob	Sugar		UCSB	US				x	
Shinji	Sumimoto		Fujitsu	Japan				x	
Makoto	Taiji		Riken	Japan			x	x	x
Toshikazu	Takada		Riken	Japan				x	x
Hiroshi	Takemiya		JAEA	Japan					x
Bill	Tang		PPPL	US				x	x
John	Taylor		CSIRO	AU				x	x
Rajeev	Thakur		ANL	US				x	x
Anne	Trefethen		Oxford	UK		x	x	x	x
Akira	Ukawa		U of Tsukuba	Japan				x	
Mateo	Valero		BSC	Spain	x		x		x
Aad	van der Steen		NCF	NL				x	
Jeffrey	Vetter		ORNL	US		x	x	x	x
Vladimir	Voevodin		Moscow State U	Russia			x		
Andy	White		LANL	US	x	x			
Peg	Williams		Cray	US		x	x	x	x
Robert	Wisniewski		IBM	US				x	x

		Affiliation	Country	SC08	Sante Fe	Paris	Japan	Oxford
Felix	Wolf	Juelich	Germany					x
Kathy	Yelick	LBNL	US	x	x	x		
Akinori	Yonezawa	U Tokyo	Japan				x	
Thomas	Zacharia	ORNL	US	x				x

8.2 Maui, HI, USA, October 18-19, 2010

The attendees of the Maui IESP Workshop were:

Patrick Aerts	NWO	NL
Giovanni Aloisio	Euro-Mediterranea	Italy
Jean-Claude Andre	CERFACS	France
Michael Ashworth	Daresbury	UK
David Barkai	Intel	US
Pete Beckman	ANL	US
Jean-Yves Berthou	EDF	France
Taisuke Boku	U of Tsukuba	Japan
Bertrand Braunschweig	ANR	France
Ron Brightwell	SNL	US
Franck Cappello	INRIA	France
Charlie Catlett	ANL	US
Barbara Chapman	U of Houston	US
Xuebin Chi	CAS	China
Andrew Chien	Self	US
Alok Choudhary	NWU	US
Tim Cornwell	CSIRO	AU
Bronis de Supinski	LLNL	US
David Dean	ORNL/DOE	US
Jack Dongarra	U of Tennessee	US
Sudip Dosanjh	SNL	US
Stephane Ethier	PPPL	US
Hugo Falter	ParTec	Germany
Teresa Finchum	U of Tennessee	US
Karl Fuerlinger	LRZ	Germany
Al Geist	ORNL	US
Carlo Graziani	U Chicago	US
Bill Gropp	UIUC	US
Robert Harrison	ORNL	US
Stefan Heinzel	Max Planck/DEISA	Germany
Barb Helland	DOESC	US
Mike Heroux	Sandia	US
Ryutaro Himeno	RIKEN	Japan
Thuc Hoang	DOE	US

Adolfy Hoisie	PNNL	US
Koh Hotta	Fujitsu	Japan
Yutaka Ishikawa	U of Tokyo	Japan
Zhong Jin	CAS	China
Fred Johnson	DOE	US
Lennart Johnsson	UH	US
Larry Kaplan	Cray	US
David Keyes	Columbia U.	US
Moe Khaleel	PPNL	US
Alice Koniges	LBNL	US
Bill Kramer	NCSA	US
Jesus Labarta	BSC	Spain
Jean-Francois Lavignon	Bull	France
David Lombard	Intel	US
Bob Lucas	ISI	US
Rusty Lusk	ANL	US
Barney MacCabe	ORNL	US
Satoshi Matsuoka	TiTech	Japan
Bronson Messer	ORNL	US
Peter Michielse	NWO	NL
Bernd Mohr	Juelich	Germany
Terry Moore	U of Tennessee	US
Wolfgang Nagel	Dresden	Germany
Hiroshi Nakashima	U of Kyoto	Japan
Jeff Nichols	ORNL	US
Ross Nobes	Fujitsu	UK
Rob Pennington	NSF	US
Walt Polansky	DOESC	US
Tracy Rafferty	U of Tennessee	US
Dan Reed	Microsoft	US
Rob Ross	ANL	US
Faith Ruppert	ANL	US
Robert Schreiber	HP	US
Mark Seager	LLNL	US
John Shalf	LBNL	US
Andrew Siegel	ANL	US
Horst Simon	LBNL	US
David Skinner	LBNL	US
Thomas Sterling	LSU	US
Rick Stevens	ANL	US
Fred Streitz	LLNL	US
Sriram Swaminarayan	LANL	US
Makoto Taiji	Riken	Japan
John Taylor	CSIRO	AU
Rajeev Thakur	ANL	US
Jeffrey Vetter	ORNL	US
Andy White	LANL	US

Peg Williams	Cray	US
Marcus Wilms	DFG	Germany
Bob Wisniewski	IBM	US

8.3 San Francisco, CA, USA, April 6-7, 2011

The attendees of the San Francisco IESP Workshop were:

First Name	Last Name	Affiliation	Country
Patrick	Aerts	NCF	NL
Giovanni	Aloisio	CMCC	IT
Jean-claude	Andre	CERFACS	FR
Mike	Ashworth	STFC	UK
Pete	Beckman	ANL	US
Jean-Yves	Berthou	EDF	FR
Arndt	Bode	LRZ / TU Munich	DE
Ron	Brightwell	SNL	US
Riccardo	Brunino	CINECA	IT
Franck	Cappello	INRIA / UIUC	FR
Barbara	Chapman	U of Houston	US
Jacqueline	Chen	SNL	US
Alok	Choudhary	NWU	US
Tim	Cornwell	CSIRO/CASS	AU
Bronis	De Supinksi	LLNL / CASC	US
David	Dean	DOE	US
Jack	Dongarra	U of Tennessee	US
Sudip	Dosanjh	SNL	US
Anshu	Dubey	U of Chicago / ANL	US
Hugo	Falter	ParTec	DE
Teresa	Finchum	U of Tennessee	US
Al	Geist	ORNL	US
Sergi	Girona	BSC	ES
Ramon	Goni	BSC	ES
William	Gropp	UIUC	US
Stefan	Heinzel	Max Planck / DEISA	DE
Ryutaro	Himeno	RIKEN	JP
Atsushi	Hori	RIKEN / AICS	JP
Koh	Hotta	Fujitsu	JP
Yutaka	Ishikawa	U of Tokyo	JP
Zhong	Jin	CAS	CH
Larry	Kaplan	Cray	US
David	Keyes	KAUST / Columbia U	US
Andreas	Knuepfer	ZIH, TU Dresden	DE

Alice	Koniges	LBNL	US
William	Kramer	NCSA / U of Illinois	US

First Name	Last Name	Affiliation	Country
Jesus	Labarta	BSC	ES
Jean-Francois	Lavignon	Bull	FR
Che-rung	Lee	NTHU	TW
Alan	Lee	AMD	US
David	Lombard	Intel	US
Arthur	Maccabe	ORNL	US
Paul	Mackenzie	Fermilab	US
Satoshi	Matsuoka	Tokyo Tech	JP
Paul	Messina	ANL	US
Peter	Michielse	NCF	NL
Bernd	Mohr	Juelich	DE
Terry	Moore	U of Tennessee	US
Kengo	Nakajima	U of Tokyo	JP
Hiroshi	Nakamura	U of Tokyo	JP
Hiroshi	Nakashima	Kyoto University	JP
Ross	Nobes	Fujitsu	JP
Kathryn	O'brien	IBM Research	US
Tracy	Rafferty	U of Tennessee	US
Stephane	Requena	GENCI	FR
Rob	Ross	ANL	US
Mitsuhisa	Sato	U of Tsukuba / AICS Riken	JP
John	Shalf	LBNL/NERSC	US
Andrew	Siegel	ANL	US
Marc	Snir	U of Illinois	US
Rick	Stevens	ANL	US
Godehard	Sutmann	Juelich	DE
Sriram	Swaminarayan	LANL	US
William	Tang	Princeton	US
Kenjiro	Taura	U of Tokyo	JP
John	Taylor	CSIRO	AU
Rajeev	Thakur	ANL	US
Jeffrey	Vetter	ORNL / Georgia Tech	US
Robert	Wisniewski	IBM Research	US
Kathy	Yelick	LBNL	US

8.4 Cologne, Germany, October 6-7, 2011

The attendees of the Cologne IESP Workshop were:

First name	Name	Affiliation
Patrick	Aerts	Netherlands National Computing Facilities foundation (NCF)
Giovanni	Aloisio	CMCC & University of Salento, Lecce-Italy
Jean-Claude	Andre	ex-CERFACS
Axel	Auweter	Leibniz Supercomputing Centre
Pete	Beckman	Argonne National Laboratory / Exascale Technology & Computing Institute
Jean-Yves	Berthou	EDF
Taisuke	Boku	University of Tsukuba
Ron	Brightwell	Sandia National Laboratories
Franck	Cappello	INRIA and UIUC
Barbara	Chapman	University of Houston
Alok	Choudhary	Northwestern University
Bronis	de Supinski	LLNL/CASC
Ralph	Dieter	BMBF
Jack	Dongarra	University of Tennessee
Sudip	Dosanjh	Sandia National Laboratories
Anshu	Dubey	University of Chicago / Argonne National Laboratory
Iain	Duff	STFC -RAL
Norbert	Eicker	Juelich Supercomputing Centre
Hugo	Falter	ParTec Cluster Competence Center GmbH
Teresa	Finchum	University of Tennessee
Leonardo	Flores	European Commission - DG INFSO-F3
Fabrizio	Gagliardi	Microsoft- EMEA
Al	Geist	Oak Ridge National Laboratory
Sergi	Girona	Barcelona Supercomputing Center
Ramon	Goni	Barcelona Supercomputing Center
Andreas	Grothey	University of Edinburgh
Stefan	Heinzel	Max Planck Society
Ryutaro	Himeno	RIKEN
Atsushi	Hori	RIKEN AICS
Koh	Hotta	Fujitsu
Yutaka	Ishikawa	University of Tokyo
Zhong	Jin	Supercomputing Center, CNIC, Chinese Academy of Sciences
Andrew	Jones	NAG
Larry	Kaplan	Cray Inc.
David	Keyes	King Abdullah University of Science and Technology
Mohammad	Khaleel	Pacifc Northwest National Lab
Axel	Koehler	NVIDIA GmbH Germany
Alice	Koniges	Lawrence Berkeley National Laboratory
William	Kramer	NCSA/University of Illinois
Jesus	Labarta	Barcelona Supercomputing Center
Jean-Francois	Lavignon	Bull
Thomas	Lippert	Juelich Supercomputing Centre

David	Lombard	Intel Corporation
Robert	Lucas	University of Southern California
Arthur	Maccabe	Oak Ridge National Laboratory
Naoya	Maruyama	Tokyo Institute of Technology
Satoshi	Matsuoka	Tokyo Institute of Technology
Paul	Messina	Argonne National Laboratory
Peter	Michielse	Netherlands National Computing Facilities foundation (NCF)
Bernd	Mohr	Juelich Supercomputing Centre
Terry	Moore	University of Tennessee
Matthias	Mueller	ZIH, TU Dresden
Kengo	Nakajima	Information Technology Center, The University of Tokyo
Hiroshi	Nakamura	University of Tokyo
Hiroshi	Nakashima	Kyoto University
Ross	Nobes	Fujitsu Laboratories of Europe
Manish	Parashar	Rutgers University
Mark	Parsons	EPCC, The University of Edinburgh
Thierry	Priol	INRIA
Tracy	Rafferty	University of Tennessee
Alex	Ramirez	Barcelona Supercomputing Center
Stephane	Requena	GENCI
Tetsuya	Sakurai	University of Tsukuba
Mitsuhisa	Sato	University of Tsukuba
Tobias	Schloesser	PR Dept., Forschungszentrum Juelich
Mark	Seager	Intel Corporation
Satoshi	Sekiguchi	AIST
Thomas	Sterling	Indiana University
Shinji	Sumimoto	Fujitsu
Sriram	Swaminarayan	Los Alamos National Laboratory
William	Tang	Princeton University/PPPL
Osamu	Tatebe	University of Tsukuba
Kenjiro	Taura	University of Tokyo
Rajeev	Thakur	Argonne National Laboratory
Anne	Trefethen	Oxford e-Research Centre, University of Oxford
Akira	Ukawa	University of Tsukuba
Mateo	Valero	Barcelona Supercomputing Center
Jeffrey	Vetter	ORNL and Georgia Tech
Vladimir	Voevodin	Research Computing Center, MSU
Marcus	Wilms	DFG
Robert	Wisniewski	IBM Research
Lu	Yutong	National University of Defense Technology, China

