Distributed
European
Infrastructure for
Supercomputing
Applications

DEISA SYMPOSIUM

TOWARDS PETASCALE COMPUTING IN EUROPE

21 - 22 May 2007

Bavarian Academy of Sciences and Humanities, Munich, Germany

PROGRAM

Monday, May 21

13.30	Opening of the symposium, welcome
13:45	Greeting from the Bavarian Minister of Sciences, Research and the Arts, Thomas GOPPEL
14.00	HPC strategies in Europe, Mario CAMPOLARGO
14.40	HPC strategies in the USA, Gary JOHNSON
15.20	Outline of the next generation supercomputer development project in Japan, Kenichi MIURA
16.00	Coffee break
16.30	Existing HPC focused infrastructures: TeraGrid, Pete BECKMAN
17.10	Enabling cooperative extreme computing in Europe, Victor ALESSANDRINI
17.50	PACE – a new European partnership for creating an integrated European high performance computing service, Achim BACHEM
18.30	End of the day
19.00	Cocktail reception

Tuesday, May 22

08.30	Technology trends for petascale computing, Arndt BODE
	Scientific cases for petascale computing in Europe:
09.10	Earth system modelling and high performance computing, Jochem MAROTZKE
09.50	Petascale requirements of ITER, Karl LACKNER
10.30	Coffee break
11.00	Developing multi-scale mathematical models of the Heart, Nicolas SMITH
11.40	DEISA solutions for grand challenge problems in computational cosmology, Gustavo YEPES
12.20	Handy nanostructures for nanoelectronics, Stefan BLÜGEL
13.00	Closing

Abstracts

Enabling cooperative extreme computing in Europe

Victor ALESSANDRINI
DEISA Project Director, IDRIS-CNRS

DEISA's original motivation was to act as a vector of integration of High Performance Computing (HPC) resources at the continental scale. The DEISA services have been tailored to enable seamless access to, and high performance cooperative operation of, a distributed park of leading supercomputing platforms in Europe.
The primordial objective of the DEISA services is to enable capability computing across remote computing platforms and data repositories. Workflows across different platforms, transparent remote I/O, large file transfers, are starting to operate without inducing performance bottlenecks that would invalidate high performance computing. Adoption of these services will grow as the number of scientific users accessing different computing platforms in Europe increases.
After quickly reviewing the existing services and the service provisioning mode of the DEISA research Infrastructure, we will discuss why their persistency is needed to pave the way for new complementary initiatives in the area of HPC. Some comments will be advanced about the relevance of the DEISA environment for the efficient operation of future European supercomputers, and the current vision about the overall role of DEISA in the European HPC ecosystem will be discussed.

Outline of the next generation supercomputer development project in Japan

Kenichi MIURA
National Institute of Informatics, Tokyo

Under the auspices of Ministry of Education, Culture, Sports, Science and Technology (MEXT), a new project has been initiated in Japan to design and develop the next generation supercomputer system. This is a 7 year project starting FY 2006, with an expected total budget of ~1B $. The objectives of this project are development, installation and usage of the multi-petaflops class supercomputing system as a national leadership machine to be utilized by academia, research communities and industries. RIKEN is the main contractor of this project. This project includes such programs as research and development of the system software (OS and grid middleware) , the grand challenge applications in nano-science and life science, development of the hardware system, and establishment and operation of the Advanced Center of Excellence (COE) for Computational Science and Technology based on the said system. The system is scheduled to be completed by the end of March 2012.
The National Research Grid Initiative (NAREGI), which started in FY 2003, as one of the major Japanese national Grid projects, is now funded as a part of the system software development program within the next generation supercomputer project. In this new framework, NAREGI is expected to provide a seamless computational environment to research and academic communities in the forthcoming "peta-scale" era, which is included in the wider concept, called "Cyber-Science Infrastructure".

PACE – a new European partnership for creating an integrated European high performance computing service

Achim BACHEM
Forschungszentrum Jülich

In its roadmap published in 2006, the European Strategy Forum on Research Infrastructures (ESFRI) has identified High Performance Computing as a strategic priority for Europe. Scientists and engineers, e.g. in climate research, earth science, nanotechnology, computational chemistry, high energy physics, life sciences, must be provided with access to capability computers of leadership class in Europe to remain competitive internationally and to maintain or regain leadership.
As a result, the High Performance European Taskforce (HET) group has been established to develop a policy framework for the provision of High Performance Computing within Europe in order to meet the needs of the European academic community, industry, and society.
As an outcome of HET, a consortium called PACE (Partnership for Advanced Computing in Europe) for the ESFRI preparatory phase was build and a Memorandum of Understanding for a European Tier 0 HPC service was signed on April 17th, 2007 by14 European countries.
The presentation will give an overview on the history of HET and PACE and will describe the current plan to create a powerful European High Performance Computing Infrastructure and Service to the needs of all science communities in Europe.

Technology trends for petascale computing

Arndt BODE
Technische Universität München

Petascale computers will be highly parallel and at least partly application specific. Multi-core technologies and application specific accelerators will make it necessary to introduce new programming models and algorithmic support for scalability.
The presentation starts with a summary of architectural choices for petascale computers. An overview on programming models and tool support for such architectures is given. Finally, some examples of applications from LRZ, that have petascale requirements, will be presented.

Earth system modelling and high performance computing

Jochem MAROTZKE
Max-Planck-Institute for Meteorology and German Climate Computing Centre

I will first outline why a perspective spanning the entire Earth system is required to understand and predict the effects of human actions on the Earth's climate. Then, I will show some examples defining the current state-of-the-art of climate and Earth system modelling. Finally, I will discuss the challenges that the needs of Earth system modelling pose for the provision of High Performance Computing.

Petascale requirements of ITER

Karl LACKNER
Max-Planck Institut für Plasmaphysik, Garching

Seven Partners (EU, Japan, Russian Federation, USA, China, Korean Republic, India) are about to start construction of an experimental reactor (ITER), for the study of a largely self-sustaining ("burning") plasma as key element on the path towards commercial energy production by nuclear fusion. The detailed planning of the experimental campaigns, but also the extrapolation of ITER´s results to a DEMO power plant require to accompany construction and operation by the development of a comprehensive computational model for the plasma core and its interaction with the environment. Ab-initio simulations of the key processes (turbulent energy and particle transport, nonlinear behaviour of tolerable, macroscopic plasma instabilities, interaction of superthermal, fusion generated He-ions with the background plasma and the magnetic field…) have progressed very strongly during the past decade, but require petascale facilities to be applied to ITER-type plasma parameters and dimensions. Continuously cross-checked against results of existing experiments and ITER they should be developed and combined into a comprehensive ab-initio model of the core of a fusion power plant ("a virtual tokamak"), ultimately to be used as a design tool for DEMO and future reactors. – The talk describes the key components of such a model, their present status and achievements, and the extrapolated computer requirements.

Developing multi-scale mathematical models of the Heart

Nicolas SMITH
University Computing Laboratory, University of Oxford

Predicting information about human physiology and pathophysiology from genomic data is a compelling, but unfulfilled goal of post-genomic biology. This is the aim of integrative multi-scale physiology and is, undeniably, an ambitious goal. Yet if we can exploit even a small proportion of the rich and varied experimental data currently available, significant insights into clinically important aspects of human physiology will follow. To achieve this requires the integration of data from disparate sources into a common framework. Extrapolation of available data across species, laboratory techniques and conditions requires a quantitative approach. Mathematical models allow us to integrate sub-cellular information into cellular, tissue and organ-level, and ultimately clinically relevant scales. In this talk I will use the Heart as an example to argue that biophysically detailed computational modelling provides the essential tool for this process. Recently developed models of cellular contraction and excitation in disease will be presented. The methods of how these cell models are embedded in spatial representations of tissue will then be outlined. Using these tissue models, combined with anatomical and structural information, simulations of whole heart activation, mechanics and coronary blood flow will be demonstrated. Finally the issues associated with developing an appropriate framework for annotating, data basing and critiquing these types of models will be discussed and their importance for the development of integrative computational biology.

DEISA solutions for grand challenge problems in computational cosmology

Gustavo YEPES

Universidad Autonoma de Madrid

Simulating the formation and evolution of galaxies and other structures in the universe from cosmological initial conditions is one of the most difficult problems in Computational Physics. The non-saturating nature of the gravitational interaction makes impossible to split the computational domain in more or less isolated areas that can be computed in parallel. In order to simulate the evolution of a patch of Universe, one needs to account for the contribution of all matter in the rest of the universe. This is the major bottleneck for the efficient use of cosmological computational codes on grid-based technology, and one needs to have access to the largest parallel supercomputers with low latency interconnection. Nevertheless, there are problems in computational cosmology that can benefit from the access to a grid of supercomputers, like DEISA. In this talk, I will briefly summarize the Grand Challenges of computational cosmology that we expect to accomplish with next generation of Petaflop machines, together with the possibilities of the current DEISA infrastructure in the realization of state-of.the art simulations of the evolution of our Universe.

Handy nanostructures for nanoelectronics

Stefan BLÜGEL

Institut f'ur Festkörperforschung, Forschungszentrum Jülich

Nanoscience is a field which cuts through physics, chemistry, biology, material sience, and engineering working enormous spectrum of different material systems and structure on quite different length scale and different degrees of symmetry. The degree to which new functionalities of nanoscale structures like magnetic clusters, quantum dots, biomolecules or carbon nanowires can be exploited for specific applications depends heavily on our ability to design devices with optimal behavior in response to external stimulation, such as applied voltage. While the basic physical effects are often well understood, quantitative simulations with predictive power that do not rely on empirical models and parameters still pose a major challenge. This is due to the large numerical effort of the calculations, to accurately describe quantum effects at atomic distances. Therefore, modern first-principles simulations in nanoscience depend heavily on efficient algorithms and powerful computer hardware.

During the past ten years, first-principles calculations based on the density-functional theory (DFT) emerged as the most powerful framework to respond to the demands mentioned above on a microscopic level. By first-principles is meant, that the parameters of the theory are fixed by the basic assumptions and equations of quantum mechanics. The overwhelming success of the density-functional theory for the description of the ground-state properties of large material classes including insulators, semiconductors, semimetals, half-metals, simple metals, transition-metals and rare-earths in bulk, at surfaces and as nanostructures makes it the unchallenged foundation of any modern electronic structure based nanoscience. The wide applicability combined with the predictive power of the approach turned it to the standard model in material science.

I discuss two examples: (i) The recently discovered homochirality or single-handedness of magnetic nanostructures and (ii) the simulation of quantum transport properties nanoskopic junctions. These two problems are calculated with two conceptionally different methods on different computer architecture. Conclusions for the method development for massively parallelized Computers are drawn.

Speakers

ALESSANDRINI, Victor

Dr. Victor Alessandrini is Director of IDRIS (since 1993), Research Director at CNRS (since 2002), Director of the DEISA project and Chairman of the DEISA Executive Committee. He obtained a PhD in Theoretical Physics from the University of Buenos Aires in 1965.

BACHEM, Achim

Professor Dr. Achim Bachem has been Chairman of the Board of Directors of Research Centre Jülich since 1 October 2006. He was previously a member of the Board of the German Aerospace Centre (DLR) from 1996 to 2006. In this position he was responsible for space flight and transportation research as well as for in-

formation and communication technology. From 1983 to 1996, he was Full Professor of Applied Mathematics and Director of the Mathematics Institute at Cologne University. In 1988, he additionally became head of the newly founded Institute of Computer Science and, at the same time, was Dean of the Faculty of Science and Mathematics from 1989 to 1991. From 1993 to 1996 he was a founder member and chairman of the board of the Centre for Parallel Computing (ZPR) at Cologne University. His primary scientific interests are discrete mathematics, scientific computing and operations research.

BODE, Arndt

Prof. Dr. Arndt Bode holds the chair for computer technology and computer organization at Technische Universität München. He runs a variety of research projects in the field of computer architecture, tools and applications for parallel and distributed systems. He is member of the board of Leibniz Compute Center, the German Supercomputer Center of the Bavarian Academy of Sciences. Since 1999 he is Vice President and CIO of Technische Universität München and general editor of the widespread professional magazine "Informatik-Spektrum"

BLÜGEL, Stefan

Prof. Dr. Stefan Blügel is Executive Director of the Department of Solid State Research (IFF) at the Forschungszentrum Jülich (FZJ) (since 2006), and Professor (C4) for Theoretical Physics at the University RWTH Aachen (since 2003). He obtained his PhD in Physics at the University RWTH Aachen in 1988.

CAMPOLARGO, Mario

Mario Campolargo is head of the unit for Research Infrastructure in the Directorate-General Information Society and Media of the European Commission, involved in new initiatives in Research Networking, encompassing the deployment of a new backbone network for research in Europe and the launching of testbeds for integration and validation of new technologies. He has degrees in electrical engineering from the University of Coimbra – Portugal and in computing science from the Imperial College, London, a Diploma in Management from the Ecole de Commerce de Solvay, Brussels and received a "Diplôme d'Etudes Européennes" from Université Catholique de Louvain-La-Neuve.

JOHNSON, Gary M.

Within the Office of Science, Dr. Johnson is currently responsible for developing collaborative research activities based on computational science. He was instrumental in structuring DOE's global climate modeling and computational biology research partnerships. He is now exploring the establishment of additional partnerships in advanced nuclear energy systems, subsurface sciences, cyber security for open science, and alternative & renewable energy. Dr. Johnson is also responsible for liaison activities related to high performance computing and its applications, between the Office of Advanced Scientific Computing Research and its Japanese research institution partner, Riken, and other Asian and European organizations. He is a graduate of the U.S. Air Force Academy; holds advanced degrees from Caltech and the von Karman Institute; and has a Ph.D. in applied sciences from the University of Brussels.

LACKNER, Karl

Prof. Dr. Karl Lackner is Co-Director of the Division Tokamakphysik at the Max-Planck Institut für Plasmaphysik (since 2003) and Honorarprofessor at the University of Innsbruck (since 1988). He obtained his PhD from the University of Innsbruck in 1966.

MAROTZKE, Jochem

Prof. Jochem Marotzke is a Scientific Member and Director of the Max Planck Institute for Meteorology in Hamburg, Germany, where he heads the department "The Ocean in the Earth System" and the International Max Planck Research School for Earth System Modelling. He is the acting scientific director of the German Climate Computing Centre (DKRZ). He has worked on the theory, modelling, and observations of the large-scale ocean circulation and its role in climate dynamics, focusing on the ocean's thermohaline circulation (THC). He received his Ph.D. in Physical Oceanography from the University of Kiel, Germany, in 1990, with a thesis on the instabilities and multiple steady states of the THC.

MIURA, Kenichi

Kenichi Miura is a professor in High-end Computing at the National Institute of Informatics (NII) in Tokyo, and the project leader of the Japanese National Research Grid Initiative Project(NAREGI) since 2003. He received his Ph.D. degree in Computer Science from the University of Illinois, Urbana-Champaign in 1973. Dr. Miura joined Fujitsu in 1973 and has since been engaged in the high-end computing. From 1992 to 1996, he was Vice President and General Manager of the Supercomputer Group at Fujitsu America, Inc., and was responsible for the operation of supercomputing-related activities in the U.S. Since June 2002, he has been a Fellow of Fujitsu Laboratories, Ltd. He also served as a visiting professor at the Computer and Communications Center of Kyushu University from 1990 to 1993, and at the National Astronomical Observatory.

SMITH, Nicolas

Dr. Nicolas Smith is currently a University Lecturer University at the Laboratory of Computing at the University of Oxford in August 2006. He holds a PhD in Bioengineering from the University of Auckland awarded in 1999 and has formally held Post Doctoral and Senior Lecturing Positions at the Universities of Oxford and Auckland respectively. His research is characterised by the development of integrated multi-scale and multi-physics models of the heart which provide the ability to link biophysically detailed experimental data to integrated function at the whole organ level. Within the scope of this work he has developed both the computational techniques to enable model development and specific models to provide insight into cardiac physiology. The work within his group has been focused particularly at the cellular level in cardiac electrophysiology and contraction and at the tissue level in the areas of coronary blood flow, cardiac mechanics and electrical activation. He is currently a leader in the international Physiome Project sponsored by the International Union of Physiological Sciences (IUPS) to develop integrated multi-scale computational models of organ systems.

YEPES, Gustavo

Professor Gustavo Yepes is a specialist in Computational Cosmology. He has been working on numerical simulations of galaxy and structure formation in the Universe for more than 20 years. He has a wide experience in the use of HPC in Astrophysics and Cosmology and has been a user of almost all kind of supercomputers available throughout the years, since the first Cray-1 in late eighties until the very recent IBM supercomputer MareNostrum in Barcelona. He is the leader of an international collaboration called MareNostrum Numerical Cosmology Project (http://astro.ft.uam.es/marenostrum) with the aim of performing Grand Challenge numerical simulations of cosmic structure formation in the Universe using a variety of supercomputers across Europe and the DEISA infrastructure.