



“HPC Infrastructures for Petascale Applications”

Report from the joint DEISA PRACE Symposium May 11 – 13, 2009 in Amsterdam

Wolfgang Gentzsch

For the first time, the two large European supercomputing projects **DEISA** (the **Distributed European Infrastructure for Supercomputing Applications**) and **PRACE** (the **Partnership for Advanced Computing in Europe**) merged their annual science symposia and together organised the DEISA PRACE Symposium in Amsterdam, May 11 – 13, hosted by **SARA** and **NCF**. The Symposium attracted almost 200 participants from more than 20 countries from Europe, America, Asia, and Australia; scientific users, HPC technology experts and vendors, government representatives and industry partners. As the *Primeur HPCN Magazine* pointed out in their feature article about this event, “Amsterdam was supercomputing capital of Europe for three days”.

The Symposium took place at the Royal Tropical Institute (KIT) in the heart of Amsterdam which is a centre of knowledge and expertise in the areas of international and intercultural cooperation, operating at the interface between theory and practice and between policy and implementation.

“Amsterdam was supercomputing capital of Europe for three days”

- *Primeur HPCN Magazine*

Prominent keynote speakers from all over the world gave the global perspectives of High Performance Computing (HPC) in the Petascale era, the next generation of supercomputers for scientific research. Speakers on the first day included **Kostas Glinos** of the European Commission; **Abani Patra** from the National Science Foundation (USA), **Ryutaro Himeno** from RIKEN (Japan), **Horst Simon** from Lawrence Berkeley National Laboratory (USA), **Ben Evans** from the Australian National University (Australia) and **Vladimir Voevodin** from the Moscow State University (Russia). In the following, we sketch a few of their key messages, and we recommend to visit the DEISA website for their detailed presentations.



Global Perspectives

Kostas Glinos, Head of Unit, Géant and e-Infrastructures of the European Commission, was highlighting the ambition to make Europe a leading player in supercomputing, not only from the user perspective, but also for European suppliers of components, software, systems, and services. This can be achieved when Member States and the Commission join forces and pool investments. But it is up to the scientific community to implement a coordinated plan of action. The supercomputing strategy should strengthen the European researchers in their global competitiveness. Areas such as medicine, climate, energy are of worldwide importance.

The European Commission is stimulating the support for eScience with e-Infrastructures consisting of high-speed networking, Grid infrastructures, and data infrastructures. A relative new item is supercomputing. The goal is to increase the number of Europe-based supercomputers in the upper regions of the TOP500 list.

Also the European IT industry should be encouraged. At first, the ideas were a bit hardware system focused using a pyramid with at the top a few tier-0 systems, then a level with tier-1 country systems, etc.

Maria Ramalho Natário from the European Commission explained, that the focus is more on a complete European HPC ecosystem, with different machines and Grids at different levels, but accessible by a scientist through a workspace that hides as much as possible the complexity of the systems and the Grids, such that he can create workflows where different steps can use different system architectures or Grids when needed. The next funding call for projects, that will open in July and closes in November will support this complete ecosystem.

Abani Patra from the National Science Foundation (USA) in his Global perspective talked about Computational Discovery at Scale. One of his key messages was about the “new oracle” Predictive Science when data plus models

>>>

>>> Report from the joint DEISA PRACE Symposium May 11 – 13, 2009 in Amsterdam

plus people together build a systematic integration of observational data, knowledge of physics and best available computing methodology and hardware to extrapolate from available data and settings. He closed his presentation with a set of challenging questions which are worth noting here:

- What are the new applications that are emerging or likely to emerge in the coming decade?
- How can we best stimulate development of software for applications and the hardware architectures to match?
- How can useful software that has been developed be sustained beyond the development period?
- What systems software will be required? Distributed systems support, programming environments, runtime support, data management user tools?
- What application support environments will be needed? Application packages, numeric and non-numeric library packages, problem-solving environments?

• How can we aid or catalyze developments that make it possible to use the same tools, including compilers, debuggers and performance tools, on system scales all the way down to the typical researcher's laptop or desktop?

• What education and training actions should be considered to prepare researchers, students and educators for future cyberinfrastructure?

Horst Simon from Lawrence Berkeley National Laboratory (USA), gave an overview of the current supercomputing developments in the USA. With two supercomputers currently running over a petaflop/s, and studies underway on exascale systems that are a thousand

times faster, many people believe we will get there in a few years (in fact, around 2019). But Horst Simon said this is too optimistic. We do not know how to build (exaflop/s) systems, we do not know how to operate them, and we do not know how to program them. Exascale computer architectures necessitate radical changes to the software used to operate them and the science applications. The change is as disruptive as the shift from vector to distributed memory supercomputers 15 years ago. The current message passing coupled with sequential programming languages will be inadequate for architectures based on many-core chips. The present code development, correctness, and performance analysis tools cannot scale up to millions of threads that will run in an exaflop/s computer. With perhaps millions of cores, exaflop/s computers are bound to fail on at least small parts, very often. But the current checkpointing will be inadequate for reaching fault tolerance at the exascale.

The exaflop/s systems will also produce massive amounts of data. Fundamental changes are necessary to manage and extract knowledge from this tsunami of data.

We certainly cannot devote all of our HPC efforts to exascale systems. Though the first supercomputers have crossed the petaflop/s barrier, we still are not at a level where day-to-day sustained petaflop/s performance is delivered to many types of applications. Horst Simon sees no coherent Petascale software plan across different platforms and different agencies (in the US).

The USA are spending a lot of money each year on HPC. For the fiscal year 2009, Horst Simon estimates that it amounts to probably about \$2B total: High End Computing Infrastructure

and Applications get \$1,142 M and High End Computing R&D \$492 M.

Vladimir Voevodin from the Moscow State University (MSU) presented Russia's way to petaflop/s. Currently the fastest system in Russia is on position 35 in the TOP500 with 70 teraflop/s performance. Vladimir Voevodin, expects that by next year a machine with 0.5 petaflop/s performance will be operational. The new 258,48 teraflop/s MSU supercomputer will be installed in Q3'09, and upgraded early next year to 500 teraflop/s.

Other Global perspectives were presented by **Ryutaro Himeno** from RIKEN (Japan), **Ben Evans** from the Australian National University (Australia), and **John Towns** from TeraGrid USA). **Achim Bachem** from Research Centre Juelich (FZJ) and Coordinator of PRACE, and **Stefan Heinzl** from Garching Computing Centre (RZG) of the Max Planck Society and Coordinator of DEISA, both stressed the importance of a joint European approach to high-performance computing, which is done in a complementary way, with PRACE focusing on new European leadership-class top-tier systems, and with DEISA operating and further improving the existing distributed supercomputing infrastructure.

Achim Bachem presented a detailed map on how Europe is moving ahead towards a persistent high-end computing service for the global scientific community, as a part of the long-term roadmap detailed in the ESFRI European Strategy Forum on Research Infrastructures. Within PRACE, the partners are installing several petaflop/s systems over the next few years. Concerning HPC, 19 European PRACE member states are now talking with one voice. And all member states will need the skills for peta- and exascale computing, and not only those which could afford to build up their own installations.

Stefan Heinzl recognized that, over the last few years, the European HPC infrastructure has successfully been built and operated and is intensively used by the computational science community. Now the focus needs to be more on the grand-challenge applications for petascale computing, running potentially on hundreds of thousands of processor cores. Challenges of dramatic optimisation, multi-threading, new parallelisation strategies, new programming methods, new algorithms, libraries and tools are faced, and new codes will have to be developed through strong science community efforts in cooperation with the supercomputing centres.



>>> Report from the joint DEISA PRACE Symposium May 11 – 13, 2009 in Amsterdam

Science Communities

The symposium also featured speakers from different science communities which are supported by DEISA and are planning to use the PRACE petaflop systems. “High-performance computing is crucial for climate research to understand mechanisms of climate change and predict future climate change perturbed by human activities. The powerful computing is needed to understand and to predict extreme events and assess the regional impacts of the climate change on society and economy”, stated Prof. **Sylvie Jousaume**, researcher at CNRS and an expert in climate modelling. She is also chairing the European Network for Earth System modelling, ENES, which started the new FP7 Infrastructure project IS-ENES to better understand and predict future climate change by high-end simulations.

Peter Coveney from University College London represented the Life Science community with his talk on the Virtual Physiological Human project which is a collaboration of partners from seven European countries. **Frank Jenko** from the European Fusion Research community, researcher at the Max-Planck-Institute for Plasma Physics in Garching, presented an overview of the world-wide ITER project and the HPC needs for its success.

Carlos Frenk from Durham University gave an excellent insight into the world of cosmology and the challenging simulations of cosmic evolution by the VIRGO Consortium, the world-leading research group in this field.

More details about PRACE were presented by **Aad van der Steen** (The Netherlands), on Directions in HPC Technology, **Mark Bull** from EPCC (UK), on a Survey of HPC Systems and Applications in Europe, and **Tim Stitt** from CSCS (Switzerland), on Training and Education for Petascale Computing,

DEISA Extreme Computing

From the DEISA Extreme Computing Initiative ten computational science grand-challenge projects from all over Europe were presented: **Harm Jonker** (Weather and Climate Research, The Netherlands & UK), **Hans-Joachim Bungartz** (Engineering, Germany), **Hannu Häkkinen** (Materials Science, Finland), **Rosa Dominguez-Tenreiro** (Astrophysics, Spain), **Jean-Philippe Laval** (Engineering, France),



Michele Migliori (Computational Neuro Sciences, Italy), **Matthias Krack** (Materials Science, Switzerland), **Luís O. Silva** (Plasma Physics, Portugal), **Rainer Spurzem** (Astrophysics, Germany), and **David van der Spoel** (Computational Bio Sciences, Sweden). Last but not least, **Zoltan Fodor** from the University of Wuppertal reported on the QCD Breakthrough 2008: “The Origin of Mass of the Visible Universe”, trying to answer the question of what is

the source of the mass of the ordinary matter, successfully, in layman’s terms.

The presentations from the symposium are available on the DEISA and PRACE web sites:

http://www.deisa.eu/news_press/symposium
and
<http://www.prace-project.eu/documents>

New Video Documentation on High-Performance Computing

Markus Rampp

DEISA has just released a new video documentation named “DEISA - Towards a European High Performance Computing Infrastructure: Combined Power for Faster Computation”. This professional video production highlights the motivation and relevance of high-performance computing in general and in particular portraits the DEISA infrastructure, its scientific applications and success stories. A selection of high-impact projects conducted within the DEISA Extreme Computing Initiative (DECI) illustrates today’s broad range of high-performance computing applications and documents the scientific relevance of DEISA. The video is 24 min in length and is mainly targeted at researchers, policy makers, funding agencies as well as at the general public audience. A DVD can be requested from any of the DEISA partners or the video can be downloaded from the DEISA website:

www.deisa.eu/publications/deisa-video



Laying the foundations for advances in bone implants

Euan MacDonald

Bioglasses have been used in bone and tooth prostheses since the 1970s. However, the ways in which they interact with living tissue is still not well understood at a molecular level. In order to address this issue, the BIOGLASS project was initiated in 2007 by Professor Piero Ugliengo and Doctor Marta Corno from the University of Turin in Italy, in close collaboration with a team from the University of Modena, using DEISA resources to simulate the structural and vibrational features of different bioglasses at a quantum-mechanical level.

Bioactive glasses – or “bioglasses” – were developed in the late 1960s, when Professor Larry Hench and his team of researchers at the University of Florida discovered that certain silicate glass compositions were able to form strong bonds to bone when implanted in rats’ femurs. Given that metallic or synthetic polymer materials are rejected by the human body, this was an important medical advance in terms of bone implants and prostheses.

“Over time, research has moved on from the study of inert biomaterials to focus on those that are capable of actually stimulating tissue regeneration and growth”, notes Piero Ugliengo, Professor at the University of Turin and the lead researcher on the BIOGLASS project. “However, despite the improvement of various specific properties of bioglasses by

the addition of ‘doping atoms’ to the original composition, the fundamental mechanism governing the interaction between the inorganic material and the human body – known as the ‘Hench Mechanism’ – is still not well understood at a molecular level”.

One crucial point in this process is the formation on the bioglass surface of an hydroxyapatite-like layer (the calcium phosphate of which the bones of mammals are largely formed), which in turn is recognized by the living cells leading to bond formation between the prosthesis and the tissue. This issue has been the subject of ongoing, interdisciplinary collaboration between various research centres throughout Italy during the last six years, involving both experimental work and computer simulations. The BIOGLASS project, undertaken using DEISA resources in collaboration with Professor Cristina Menziani and Dr. Alfonso Pedone from the University of Modena, has taken the simulations of the quantum mechanics (QM) of bioglasses and hydroxyapatite to the next level.

Modelling the quantum-mechanics of bioglasses: A daunting task

“The objective of the BIOGLASS project was to analyze structural and vibrational features of different bioglass models, with variable con-

tent of phosphorous in their unit cells, using computer simulations at a quantum-mechanical level of theory,” according to Doctor Marta Corno, who worked closely with Ugliengo on the project. “Phosphorus conditions the release rate of silica into the body, and this plays a vital role in the integration of the bioactive material with living tissue. However, bioglasses are, by nature, amorphous materials, which makes their quantum mechanical simulation a daunting task”.

Usually, bioglasses are simulated by adopting a very large unit cell, envisaging enough structural disorder to accurately represent the actual material and applying periodic boundary conditions to the calculation to avoid spurious surface effects. However, although this strategy works well for classical molecular simulation, it rapidly becomes prohibitively expensive when applied to a quantum mechanical (QM) model, in which, unlike the classical molecular mechanics (MM) approach, all electrons and nuclei are treated explicitly. Doing so requires the choice of a means of approximating the energy generated by the interactions between electrons and nuclei. These approximations – referred to as “hybrid functionals” – play a key role in density functional theory, on which the project was based; in this case, the hybrid B3LYP functional, which has been proven to give excellent

results in this area, was used. The challenge, then, was to adopt a cell-size that was small enough to be feasible using B3LYP, yet retaining enough structural disorder to closely imitate the real material.

“To tackle the problem we adopted a multi-scale strategy, combining classical techniques with quantum mechanics methods,” notes Corno. “The researchers at the University of Modena have many years of experience in this field, and have developed a strategy for simulating

amorphous glasses. Firstly, a specific composition is fixed and the corresponding number




Fig. 1. Unit cell content of the B3LYP optimized structures of P0 (0% phosphorus, on the left) and P9 (9.5% phosphorus, on the right) bioglasses models. Silico-phosphorous frameworks are here represented as tetrahedra (yellow ones for phosphorus atoms); calcium ions as green spheres; and sodium ions as violet spheres.

>>> Laying the foundations for advances in bone implants

of atoms are located in the unit cell of the model, according to the experimental density; then a classical molecular dynamics simulation is run, starting at high temperature so that a melt-like status for the bioglass is obtained. The third step is the progressive cooling of the system, so that the melt solidifies into a highly disordered bioglass. Lastly, classic techniques of energy minimization are carried out on the most representative structures, in order to generate the proper relaxed structure for full quantum-mechanical periodic calculations. By merging computational techniques in this way, we were able to strike a balance between the large cell units of molecular dynamics and the requirements of quantum simulations”.

In order to test the validity of the strategy, the researchers generated models of 4555 Bioglass composition with 78 atoms in the unit cell following the above procedure, and carefully compared the resulting structures optimized using the QM and MM approaches. Given that QM methods are more accurate than those of MM, any large differences between structures optimized at the two levels would indicate a flaw in the adopted force-field. Fortunately, the comparison revealed a high degree of similarity between the two structures, demonstrating the credibility of the procedure and of the adopted force field. The analysis was then extended to four different bioglass models, each studied at variable phosphorus composition, in order to uncover the precise nature of the phosphate groups within the unit cell – a fact that has important implications for the dissolution properties of bioglass when in contact with body fluids.

Simulating bioglasses could lead to important advances in medical science

The quantum-level simulation of bioglasses using DEISA resources has generated a number of impressive results. According to Ugliengo, “integrating the quantum mechanics approach allowed to us to represent the electronic features of the bioglass, which had been entirely missing from classical simulations, providing important insights on the nature of modifying cations and the bridging oxygen atoms of the silica framework”.

“It is worth stressing that this was the first full quantum-mechanical calculation of the

vibrational spectrum carried out with the hybrid B3LYP functional, which is one of the most accurate. This enabled us to make a detailed assignment of the vibrational modes. Besides being important in its own right, this advance can facilitate enormously the work of spectroscopists in assigning the spectral features of an experimental spectrum recorded for a real bioglass”.

Bioglasses have a wide range of important applications, and, as life expectancies increase, so does the need for implants capable of replacing repairing damaged bone tissues. “In particular, recent work has shown the high potential of bioactive glasses for 3rd-generation tissue-engineering applications, where a highly porous, biodegradable scaffold combined with tissue cells hosts the *in vitro* growth of immature bone like material, then implanted *in vivo* where the tissue-engineered construct adapts to the living environment and stable mature bone is formed. Although most of these achievements have been on the experimental side, classical and quantum mechanical simulations, like those of the BIOGLASS project, can be extremely useful in helping experimentalists design better materials”, Ugliengo explains.

“The project has allowed us not only to gain insight on the vibrational features of bioglasses as a function of phosphorous content, but also to identify new optimized bulk structures that could be used as a starting point for designing new surfaces; the next step is to move from the quantum mechanical simulation of the bulk to that of the surface interface between the bioglass and the body tissues. We will be applying again to the DEISA initiative in the near future in order to be able to continue our work in this field”.




Fig. 2. Doctor Marta Corno and Professor Piero Ugliengo

Supercomputing resources essential to quantum-level simulations of bioglasses

According to Ugliengo, DEISA resources were essential in achieving the targets of the BIOGLASS project, for several reasons. Most importantly, because very large unit cells were adopted, a large number of CPUs (at least 128, or better 256) were needed to complete the tasks in the allotted time. The necessary calculations were made possible by a highly sophisticated programme called CRYSTAL06, developed by a group headed by Professor Roberto Dovesi at the University of Turin, which is able to solve the Schrödinger equation for molecules, polymers, 2D slabs and 3D crystals on an equal footing. CRYSTAL06 was able to exploit the supercomputing resources available within the DEISA initiative by running in a massive parallel fashion using both standard highly optimized BLAS libraries and distributed memory. This is extremely important for treating systems of the size of the bioglasses because the diagonalization of the huge Fock matrix is carried out in a well-balanced and distributed way on all the available CPUs.

All calculations were run using either the IBM BCX/5120 supercomputer provided by CINECA supercomputing center of Bologna, or at the Marenostum as provided by the Barcelona SuperComputing Center, both nodes of DEISA. The best compromise between waiting time and efficiency was achieved by using 128 CPUs for each run lasting 24 elapsed hours; the CRYSTAL06 code has been extremely well engineered on both systems. In effect, DEISA enabled the researchers to carry out a project that was more ambitious by far than anything that would have been possible using only in-house resources.