DEISA Mini-Symposium on Extreme Computing in an Advanced Supercomputing Environment

Wolfgang GENTZSCH and Hermann LEDERER
Rechenzentrum Garching der Max-Planck-Gesellschaft
Max Planck Institute for Plasma Physics, 85748 Garching, Germany
E-mail: {Gentzsch, Lederer} @rzg.mpg.de

Abstract.
The DEISA Mini-Symposium held in conjunction with the ParCo 2009 Conference was devoted to ŞExtreme Computing in an Advanced Supercomputing EnvironmentŤ. For a representative overview of the application related key areas involved, the Mini-Symposium was structured into three sessions devoted to the description of the DEISA Extreme Computing Initiative DECI and the Science Community support, the application oriented support services, and examples of scientific projects performed in the Distributed European Infrastructure for Supercomputing Applications (DEISA) both from the DECI and the fusion energy research related science community behind the EU FP7 project EUFORIA.

Introduction

A consortium of the major supercomputing centres in Europe is operating the Distributed European Infrastructure for Supercomputing Applications [1]. Initiated in 2002 and started as an EU FP6 project in 2004, operation and further development of the HPC infrastructure is continued with EU FP7 support from 2008 to 2011. DEISA2 focuses on the provisioning and operation of infrastructure services which allow its users to work efficiently within a distributed high performance computing environment. Through these services and the continued operation of this infrastructure with global importance, DEISA2 contributes to the effective support of world-leading computational science in Europe [2,3].

1. The DEISA Extreme Computing Initiative and its Science Community Support

The DEISA Extreme Computing Initiative [4,5] was launched in May 2005 by the DEISA Consortium, as a way to enhance its impact on science and technology. The main purpose of this initiative is to enable a number of grand challenge applications in all areas of science and technology. These leading, ground breaking applications must deal with complex, demanding and innovative simulations that would not be possible without the DEISA infrastructure, and which benefit from the exceptional resources provided by the Consortium. The DEISA applications are expected to have requirements that cannot be fulfilled by the national HPC services alone.

A European Call for Extreme Computing Proposals is published regularly. By selecting the most appropriate supercomputer architectures for each project, DEISA is opening up the currently most powerful HPC architectures available in Europe for the most challenging projects. This mitigates the rapid performance decay of a single national supercomputer within its short lifetime cycle of typically about 5 years, as implied by Moore's law.

The number of DECI proposals and accepted projects increases from year to year. Supercomputing resources are awarded for scientific projects on the fastest supercomputers in Europe, accompanied by application enabling services which are distributed over all partner sites, with access to the application specialists' knowledge of each site.

In DEISA2 the consortium has extended its service provisioning model from individual project support, as in the DEISA Extreme Computing Initiative, to persistent service provision to specific European user communities. A call for expressions of interest, published at the end of 2008 [6], resulted in expressions of interest of important communities from the science areas of fusion energy research, cosmology and space science, climate research, and life sciences.

2. Application-Oriented DEISA Infrastructure Services

The applications submitted to DECI consist of complex and innovative simulations requiring the advanced European supercomputer infrastructure DEISA. Therefore DEISA provides a broad spectrum of advanced and integrated services for challenging applications, such as:

- Common global high performance file system, to greatly facilitate data management across Europe;
- Uniform access methods common to the entire infrastructure, such as the UNICORE middleware;
- DEISA Common Production Environment (DCPE) to unify the heterogeneous software environments of the different partner sites;
- A team of European system operation specialists handling the system services operations;
- An Applications Task Force team at the service of the scientists both for user support and enabling of complex applications to utilise DEISA services efficiently.

The various heterogeneous software environments to deal with applications, especially the compilers, libraries, pre-installed applications and tools, have been unified into one homogeneously appearing environment, the DCPE. Together with documentation available to all DEISA users, it dramatically eases application portability from one site to another. DCPE is accessed in an identical way from each site and provides users with a uniform environment, independent of the underlying architecture. There is a particular advantage for the user if the required application is part of the DCPE; in this case the enabling effort required to exploit the DEISA infrastructure will be small because there will already be a compiled and optimized version of the program ready for the user to execute on the target platform. This is often the case in computational chemistry or life sciences where increasing use is made standard codes rather than writing own codes from scratch.

An Applications Task Force team at the service of the scientists has been established both for user support and enabling of complex applications to utilise DEISA services efficiently. Examples of application enabling include scaling of parallel programs for the efficient usage of thousands of processors, architecture dependent code optimizations, management of loosely coupled applications, multi-site task farms and complex work-flows. The process of application enabling for, say a DECI project, involves the researcher who wrote or wishes to use the application, the technical support from the DEISA site which is physically closest to the principle researcher, and the support personnel of the DEISA sites who will provide the computing cycles, who give advice on how best to execute the code on their local architecture. The application enabling effort is an important component of the DEISA infrastructure and is estimated and taken into account for every project proposal or related activity.

3. Science Project: Direct Numerical Simulation of Turbulent Development of a Round Jet at Reynolds Number 11,000

In his talk Christophe Bogey described how a Direct Numerical Simulation (DNS) of a round jet at a Reynolds number of 11,000 has been performed in the framework of the DECI project JetTurb using 755 million mesh points. The turbulent statistics in the region of flow establishment was calculated, and databases of 3-D unsteady fields are now available for further analyses of the jet physics.

The DNS has been performed using an in-house solver of the full 3-D Navier-Stokes equations, based on low-dissipation and low-dispersion explicit finite-difference methods. The jet simulated is round and isothermal, at a Mach number of 0.9 and a Reynolds number of 11,000. The DNS was carried out at the HLRS Centre in Stuttgart, on 36 processors of the NEC SX-8. The overall CPU speed obtained was 212 Gflops (i.e. 5.9 Gflops per processor), and 30,000 CPU hours have been required.

The achievement of the present project is clearly illustrated by preliminary results. The flow development, from the initially laminar to turbulent shear layers, then to the turbulent jet, is spectacular. As expected from a DNS at the Reynolds number considered, fine turbulent scales are visible. First insights into the flow properties are also given by the variations of the terms in the turbulent energy budget. Along the centreline for instance, molecular dissipation, turbulence diffusion and mean flow convection are dominating. In addition, the convergence of the energy terms appears satisfactory, and their sum is nearly nil, which supports that the present DNS is accurately resolved.

4. Science Project: Chemical Characterization of Superheavy Elements on Gold Clusters by 4Component Full Relativistic DFT

The presentation by Leonardo Belpassi was related to the chemical characterization of super heavy elements (SHE), achieved by studying their adsorption on a heavy metal surface, currently being of fundamental importance for the placement of new elements in the periodic table. Relativistic effects on SHE chemistry are known to alter expectations dramatically. Full relativistic 4-component Dirac-Kohn-Sham (DKS) theory has been used in order to gain an accurate understanding of the chemical properties of the element 112 (E112) bound to several gold clusters and extend this approach to the next candidate to experimental characterization: E114. Cluster calculations, analyzing the energetics and electronic structure of the interaction, and assessing convergence with respect to the basis set size and number of atoms have been carried out using the full relativistic density functional theory code BERTHA.

The large numbers of heavy atoms that need be considered, the peculiar structure in the DKS calculations and the large basis set adopted required an extreme computational effort, demanding the development and implementation of highly effective parallelization scheme. The parallelization has been performed using MPI and ScaLAPACK. One peculiar aspect of the approach is that only the master process needs to allocate all the arrays, that is the DKS matrix, the overlap matrix and eigenvectors one. Each slave process allocates only some temporary small arrays when needed.

With this new parallel implementation of the DKS method, the electronic structure and interaction energy of E112 and E114 with several gold clusters up to Au_{34} were accurately computed. The chemical characterization of these SHE has been carried out comparing the energetics, electronic structure and charge transfer with the results obtained for their homologues Hg and Pb and the inert noble gas atom Rn. All the calculations (i.e. 350,000 CPU hours) have been performed, within the DECI project, on the SGI Altix 4700 at Leibniz Rechenzentrum (LRZ), Germany.

5. Science Project: EUFORIA: Exploring E-Science for Fusion

The science community activities related to the EU FP7 project EUFORIA were described by David Coster. In preparing for ITER [7], a number of computational challenges need to be overcome: individual parts of the problem are "grand challenge" problems in their own right, but they also need to be combined to prepare for simulations that encompass all the relevant space and time scales.

ITER is the next generation of fusion devices and is intended to demonstrate the scientific and technical feasibility of fusion as a sustainable energy source for the future. To exploit the full potential of the device and to guarantee optimal operation for the device, a high degree of physics modelling and simulation is needed already in the current construction phase of the ITER project. The detailed modelling tools that are needed for an adequate description of the underlying physics cover both a wide range of time scales and spatial orderings and are in general very demanding from a computational point of view.

The project will enhance the modelling capabilities for ITER sized plasmas through the adaptation, optimization and integration of a set of critical applications for edge and core transport modelling targeting different computing paradigms as needed (serial and parallel grid computing and HPC). Deployment of both a grid service and a High Performance computing service are essential to the project. A novel aspect is the dynamic coupling and integration of codes and applications running on a set of heterogeneous platforms into a single coupled framework through a workflow engine. This strongly enhances the integrated modelling capabilities for fusion plasmas and will at the same time provide new compute tools to the fusion community in general.

Acknowledgments

The authors are deeply grateful to the following colleagues for their contribution to the DEISA Mini-Symposium at ParCo 2009: Alison Kennedy EPCC; Giovanni Erbacci, CINECA; Denis Girou, IDRIS; Gavin J. Pringle, EPCC; Mariano Vazquez, BSC; Juha Fagerholm, CSC; Christophe Bogey and Oliver Marsden, Laboratoire de Mécanique des Fluides et dŠAcoustique, Ecole Centrale de Lyon; Leonardo Belpassi, Loriano Storchi, and Francesco Tarantelli, Department of Chemistry, University of Perugia; David Peter Coster, Max Planck Institute of Plasma Physics, Garching; Par Strand, Chalmers University, Goteborg; Michael Martinez, Vlad Cojocaru, Paolo Mereghetti, and Rebecca C. Wade, Molecular and Cellular Modelling Group, EML Research gGmbH, Heidelberg; Kia Balali-Mood and Mark Sansom, Dept. of Biochemistry, University of Oxford.

The DEISA Consortium thanks the European Commission for support through contracts RI-508830, RI-031513, and RI-222919.

References

[1] DEISA, 2009, see http://www.deisa.org

[2] Lederer, H.: DEISA2: Supporting and developing a European high-performance computing ecosystem, Journal of Physics: Conference Series 125 (2008) 011003; doi:10.1088/1742-6596/125/1/011003

[3] Supercomputing gets its own superhero, see http://cordis.europa.eu/ictresults/index.cfm/section/news/tpl/article/BrowsingType/Features/ID/90477

[4] DECI, 2009, see www.deisa.org/deci/

[5] Lederer, H.: DECI - The DEISA Extreme Computing Initiative, inSiDE, Vol 6 No 2, Autumn 2008, eds. H-G. Hegering, Th. Lippert, M. Resch; Gauss Centre for Supercomputing, 2008

[6] DEISA Announces Virtual Community Support Initiative, HPCwire 2008, see http://www.hpcwire.com/industry/government/DEISA-Announces-Virtual-Community-Support-Initiative-36014839.html

[7] ITER, 2009, see http://www.iter.org