

# DEISA Newsletter

DISTRIBUTED EUROPEAN INFRASTRUCTURE FOR SUPERCOMPUTING APPLICATIONS



Vol. 3, 2007

## DEISA Training: First hand experience of the infrastructure

At the DEISA training sessions already over 70 scientists from 13 European countries have learned to exploit the European supercomputing environment.

In March 32 scientists and 15 experts in supercomputing gathered to discuss topics of HPC and Grid within the DEISA infrastructure at the Barcelona Supercomputing Center, Spain, where the Europe's fastest supercomputer MareNostrum is located. This training session was an overall success; it met participants' expectations in creating a positive contribution for their research activities as well as in understanding how the DEISA can aid them in achieving their research goals.

At the training participants were first given a detailed introduction to the DEISA environment, followed by the opportunity to experience it first-hand in the "Access to DEISA" hands-on session. The participants were granted with a one-month-long test user account for

DEISA in order to be able to continue experimenting DEISA also once returned to their institutes.

The special topic of this training session was "Performance measurements and codes optimization", explaining to application developers the methodology of performance analysis, showing how to efficiently use some dedicated performance measurement tools and illustrating it on various real codes from many different scientific fields.

Presentations of use cases for the DEISA infrastructure, coming from some DECI projects run in 2006, concluded the event, giving insight to recent success stories on huge simulations in astrophysics, quantum computer simulations and materials science.

The next opportunity to take part in the training will be on May 30 – June 1, when the 4th DEISA training will be held at CSC- the Finnish



Scientists and experts in supercomputing gathered in March to Barcelona Supercomputing Center, Spain for the 3rd DEISA Training.

IT center for science in Espoo, Finland. In addition to introducing DEISA, this training program will include use cases and the special topic of heterogeneous environments and optimization.

**The registration for the 4th DEISA training will be open until May 8, at [www.deisa.org/training](http://www.deisa.org/training).**

## The DEISA Extreme Computing Initiative: Call for proposals 2007

The purpose of the DEISA Extreme Computing Initiative (DECI) is to enhance the impact of the DEISA research infrastructure on leading European science and technology.

This initiative consists of the identification, enabling, deploying and operating flagship applications from different areas of science and technology.

These leading, ground breaking applications must deal with complex, demanding, innovative simulations that would not be possible without the DEISA infrastructure, and which would benefit from the exceptional resources of the Consortium.

Apart from compute resources, DEISA can also provide increased support in the design and enabling of these applications. In some cases, this support can entail a medium term

commitment of human resources, i.e. an engineer for several months.

Projects supported by DECI will be chosen on the basis of innovation potential, scientific excellence and relevance criteria. Proposals from PIs that have yet to benefit from DECI compute resources will be given preference.

On the 30th of April, the DEISA Consortium will call for proposals for a third generation of Extreme Computing Initiative applications that will be supported and/or deployed from November 2007. The closing date for this third call for proposals is June 30, 2007.

Interested scientists can first email the DEISA Applications Task Force at [ataskf@deisa.org](mailto:ataskf@deisa.org) to obtain the guidance needed to find the best fit between their requirements and the distributed supercomputing environment.

For further information and a proposal template form, please visit [www.deisa.org](http://www.deisa.org).

## DEISA Symposium

Towards Petascale Computing in Europe

21 - 22 May '07, Munich, Germany

- HPC strategies in the USA, Japan, and Europe,
- Existing HPC related e-Infrastructures: DEISA, TeraGrid,
- Technology issues in petascale computing,
- Scientific cases for petascale computing in Europe.

The registration will close on May 7.

[www.deisa.org/symposium](http://www.deisa.org/symposium)

## Performance explosion for DEISA



The upgraded Altix 4700 at LRZ with a LINPACK performance of 56.5 TFlop/s and a memory of nearly 40 TBytes.

Throughout the second half of March and well into April 2007, SGI has been busy upgrading the previously run Madison-based 4096-processor Altix 4700 ("phase 1") to a 9728 core Montecito-based system. Due to the upgrade, the Altix at LRZ is now the second fastest computer in the DEISA infrastructure. With the upgraded system, LRZ is now well prepared to tackle the large simulations planned by the scientific communities for the next few years. A fixed percentage of the compute cycles will also be dedicated for grid computing, and in particular for projects within the purview of DEISA.

This upgrade entailed an increase of system main memory from 17 to nearly 40 TBytes and parallel file system disk space from 300 to 600 TBytes. With a LINPACK performance of now 56.5 TFlop/s, the machine is significantly faster than the Earth Simulator and may well rank among the top 10 in the next Top500 list, to be published in June, 2007. Even though both the NUMalink and the memory subsystem remain at the same capacity as before, it was possible to achieve nearly twice the benchmark performance of the phase 1 installation: The weighted mean of all benchmarks yields an aggregate of 16.2 TFlop/s. For the most part, this is due to the increased L3 cache size and the split L2 data and instruction caches of the Montecito. The machine is now divided into 19 partitions with 512 core single system images; six partitions have blades with four cores per memory interface ("high density"), 13 partitions have blades with two cores per memory interface ("high bandwidth"). Hence, jobs can be suitably scheduled depending on their need for memory bandwidth versus computational density.

## FEARLESS: A new modelling approach for turbulent astrophysical flows

Dr Wolfram Schmidt, University of Würzburg, Germany

Turbulence in engineering applications and atmospheric sciences has frequently been modelled by large eddy simulations (LES). In LES, the dynamics of turbulent eddies is computed on large scales, while a subgrid scale model approximates the influence of smaller eddies. However, in astrophysics, phenomena such as supersonic turbulence in star-forming gas clouds challenge the LES approach. The self-similarity hypothesis employed in LES fails to be applicable over a wide range of disparate scales. Thus, Alexei Kritsuk from the University of California, San Diego, proposed to adopt a method called adaptive mesh refinement (AMR). This method involves inserting computational grids of higher resolution into flow regions where turbulent structures such as eddies or shock fronts are forming. A major problem is to find the criteria for the generation of refined grids based on various fluid dynamic processes.

The computational resources granted by the DEISA Extreme Computing Initiative (DECI) to perform highly resolved AMR simulations of supersonic turbulence were used. Depending on the size of the computational grid, 16 to 126 CPUs of the SGI Altix supercomputer at SARA,

Netherlands, were required for each simulation. The basic idea was to trigger the refinement by monitoring flow properties such as the vorticity (the rotation of the velocity field) and the rate of gas compression (due to shocks or gravity). This is illustrated by the simulation of a three-dimensional visualisation of the isosurfaces of the vorticity (refer to the picture). The refined grids (drawn as yellow boxes) are mostly generated in the vicinity of sheet-like structures arising from shocks. The tube-like structures indicate the centres of eddies. An extensive statistical analysis showed good agreement with the known properties of turbulence inferred from simulations without AMR. This makes us confident that AMR can be carried over to more complex scenarios involving thermal and chemical processes as well as self-gravity.

However, owing to the extreme range of different length scales, it is generally impossible to treat fully developed turbulence by means of AMR only. This is because a prohibitive number of refined grids would be necessary. For this reason, we are presently combining AMR with a subgrid scale model which links the notions of AMR and LES. We call this new method "Fluid mEchanics with Adaptively Refined

Large Eddy SimulationS" (FEARLESS). Upcoming applications will encompass the star formation in the turbulent interstellar medium, the feedback from star formation onto the evolution of spiral galaxies and the dynamics of hot gas in galaxy clusters. We expect that FEARLESS will open new perspectives in astrophysics by the as yet unequalled level of sophistication in the treatment of turbulence.

3D Visualisation:  
Isosurfaces of  
vorticity

© Dr Wolfram Schmidt

