DEISA

ADVANCING SCIENCE IN EUROPE

Distributed
European
Infrastructure for
Supercomputing
Applications

DEISA
Advancing Science in Europe

EU FP6 eInfrastructure Project.

Acknowledgments

The DEISA Consortium thanks the European Commission for support through contracts FP6-508830 and FP6-031513.

Introduction to DEISA

DEISA (Distributed European Infrastructure for Supercomputing Applications) is a consortium of eleven leading national Supercomputing Centres from seven European countries, strongly focused on high performance distributed supercomputing. The consortium has designed, deployed and is currently operating a high performance, distributed supercomputing environment at a continental scale with a high level of heterogeneity and complexity. The consortium activity started as a European FP6 Integrated Infrastructure Initiative project in 2004.

In the DEISA environment, eleven leading supercomputing platforms in Europe – extending from Italy and Spain through Central and Western Europe to the UK and Finland – are connected by a priority high speed network empowered by GEANT2 and the national research networks (NRENs).

The fundamental objective of DEISA is to provide services that enable high performance operation of remote computing platforms on remote data sets. Some outstanding features of the DEISA integrated infrastructure are:

- High speed end-to-end network connecting the supercomputing platforms
- Common global high performance file system at continental scale to greatly facilitate data management across Europe
- Uniform infrastructure access through the UNICORE middleware system
- Harmonization of the manifold heterogeneous software environments through the DEISA Common Production Environment (DCPE)
- Portals and internet interfaces for transparent access to complex supercomputing environments
- Enabling of cooperative extreme computing in Europe

The design, deployment and operation of the distributed supercomputing infrastructure is carried out by a number of trans-national expert teams providing services that cover a substantial number of areas in distributed computing, like network connectivity, middleware and security, global file systems and data management, application enabling and user support. In addition, a number of research activities have been deployed to closely work with the early adopters of the infrastructure from different scientific communities to address specific needs, e.g., from Materials Science, Cosmology, Plasma Physics, Life Sciences, Engineering & Industry.

Infrastructure access for the Life Science and Materials Science communities has additionally been facilitated by specific science gateways, completely hiding the complexity of the distributed infrastructure and providing application-specific enhancements.

Application and user support specialists have enabled and optimised applications for usage in DEISA and have developed the DEISA Common Production Environment (DCPE). Examples of application enabling include: scaling of parallel programs for the efficient usage on thousands or tens of thousands of processor-cores; design of coupled applications; determination of best suited architecture(s) in DEISA; adaptation of applications to the DEISA infrastructure and architecture dependent optimizations. Through the DCPE, a uniform DEISA environment with the same appearance and functionality is provided to the users, whatever the underlying architecture, thereby enabling a homogeneous access to a highly heterogeneous supercomputing environment. Since late 2005, DEISA has been used for conducting the most challenging European projects in computational sciences requiring the most advanced supercomputing resources available.

Obviously, high performance operation of remote computing platforms on remote data sets is a necessary condition for the efficient operation of the forthcoming shared petascale supercomputers to be deployed in the context of the PRACE FP7 project. The groundbreaking work of DEISA is advancing computational sciences in Europe, is contributing to pave the way towards an integrated European HPC ecosystem and to enhance Europe's visibility in the supercomputing domain.

The DEISA Extreme Computing Initiative

The DEISA Extreme Computing Initiative (DECI) was launched in 2005 to enhance DEISA's impact on science and technology.

The DECI consists of the identification, enabling, deploying and operation of "flagship" applications in selected 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. Projects supported by DECI are chosen on the basis of innovation potential, scientific excellence and relevance criteria. Multi-national proposals are especially encouraged.

For the advancement of computational sciences in Europe in the supercomputing area, application enabling is of key importance. As a prerequisite to handle and master this challenging task, DEISA created the Applications Task Force (ATASKF), constituted as a team of leading experts in high performance and grid computing. ATASKF deals with all the different aspects of Extreme Computing projects from leading European scientists in its full complexity at all the necessary levels, thus enhancing the impact on computational sciences.

DECI also greatly benefits from the DEISA Common Production Environment (DCPE). Not only that the heterogeneous DEISA hardware and software environments show the same appearance and functionality: DCPE also provides a set of preinstalled and configured important applications which are maintained and regularly upgraded. And a European team of system operation specialists has been handling all aspects of the coordination and synchronisation of system services, maintenance measures and failure situations, in support of DECI.

Applications enabling

Application specialists have enabled and optimised applications for usage in DEISA. Tasks included: scaling of applications, workflows, coupled applications, determination of best suited architecture(s) in DEISA, adaptation of applications to the DEISA infrastructure and architecture dependent optimizations.

Scaling of parallel programs for the efficient usage on thousands of processor-cores is a very challenging task faced with modern supercomputer architectures. Design, deployment and optimization of workflow applications allows for the chaining of simulation and pre- and post-processing steps, involving successive operations on large data sets performed on several platforms. The set-up and support of coupled applications is another important field, e.g. for climate system modelling where separate components for e.g. ocean, atmosphere, sea ice, soil, atmospheric chemistry and aerosols etc. are used.

Extreme computing projects

The multitude of the described services, activities, measures and pan-European co-operations altogether have enabled the success of DEISA and of DECI. DECI has supported and continues to support the most challenging supercomputing projects in Europe which require the special resources and skills of DEISA.

A European Call for Extreme Computing Proposals was published annually in spring. After receiving around 40 to 50 proposals in 2005 and 2006, respectively, over 60 proposals from the 2007 call were competing for DEISA resources in 2008, asking for more than 70 million computing hours.

By selecting the most appropriate supercomputer architectures for each project, DEISA is opening up the respective most powerful HPC architectures available in Europe for the most challenging projects, mitigating 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.

So far scientists from 15 different European countries with collaborators from four other continents (North and South America, Asia and Australia) have benefited.

This booklet reports on scientific achievements through DEISA, especially by Extreme Computing projects, but also by the scientific Joint Research Activities.

The scientific projects reported cover major areas of science including astrosciences, Earth sciences, engineering, life sciences, materials science, plasma physics, and quantum chromodynamics.

As one of the highlights, results from a German/British group with Principal Investigator Kurt Kremer were published as the cover story of NATURE (May 24, 2007): For almost two decades, physicists have been on the track of membrane mediated interactions. Simulations in DEISA have now revealed that curvy membranes make proteins attractive.

New modelling approach for turbulent astrophysical flows

A research group at University of Würzburg has used DEISA's computational resources to perform highly resolved adaptive mesh refinement simulations of supersonic turbulence. The basic idea was to trigger the refinement by monitoring flow properties such as the rotation of the velocity field and the rate of gas compression. A new method called "Fluid mEchanics with Adaptively Refined Large Eddy SimulationS", FEARLESS, has been introduced.

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.

Method called adaptive mesh refinement adopted

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, the Netherlands, were required for each simulation.

Simulations of more complex scenarios possible

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 (Figure 1). 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.

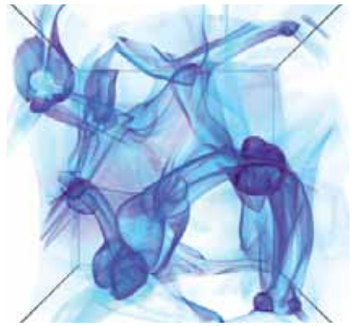

Figure 2: The image depicts the magnitude of the vorticity in an AMR simulation of forced supersonic turbulence. The sheet-like structures that can be seen indicate the formation of strong shock fronts while turbulence is still developing.

Fluid mechanics with adaptively refined large eddy simulations

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.

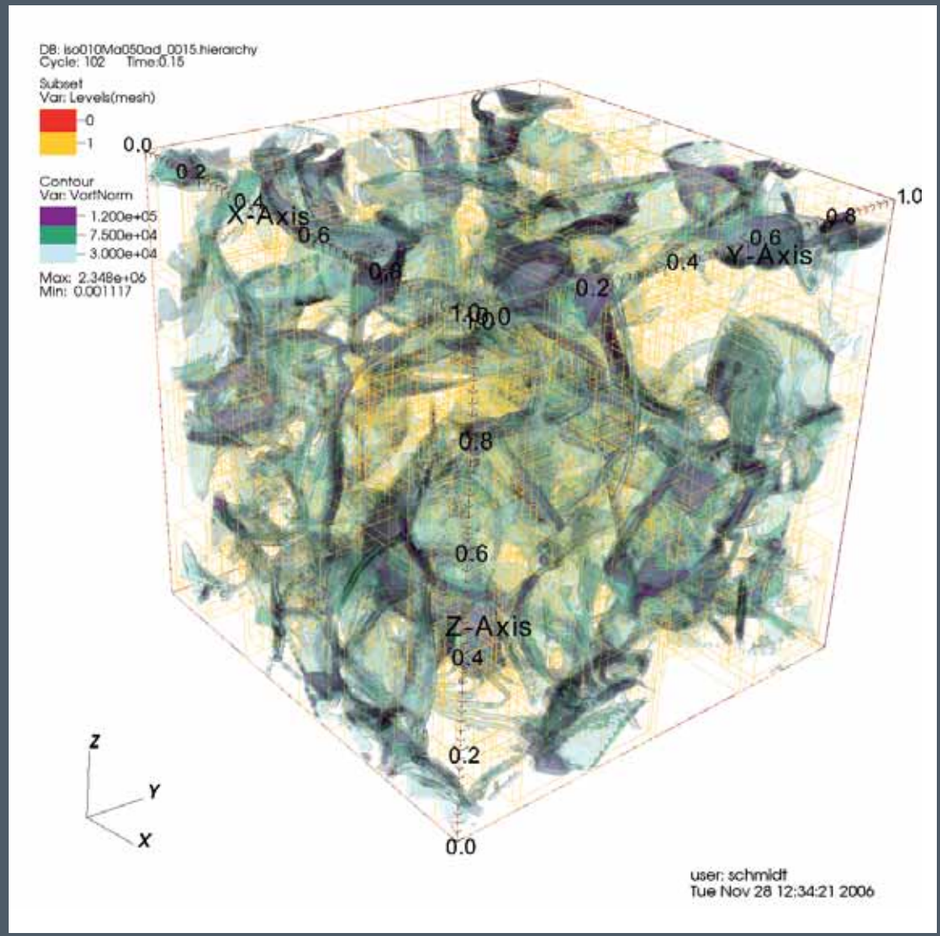

Figure 1: Simulation of a three-dimensional visualisation of the isosurfaces of the vorticity.

Simon White and Carlos Frenk, Max Planck Institute for Astrophysics, Germany, and University of Durham, UK,
a project of the Virgo Consortium

Galaxies–intergalactic medium interaction

These simulations provide unprecedented insight in how galaxies form, and also allow a detailed analysis into how galaxies interact with the surrounding intergalactic medium.

The distribution of galaxies in the Universe displays structure on very large scales. Spiral galaxies tend to agglomerate in small groups, which themselves are arranged in a filamentary pattern, with most elliptical galaxies living in the dense clusters of galaxies that occur at the intersection of filaments. The segregation of galaxy types within the large-scale structure implies that somehow the environment affects the properties of a growing galaxy.

However, most of the mass in the Universe is in the form of dark matter, an as of yet unidentified form of matter that does not interact with light. The stars and gas we see in the galaxies are just tracing the underlying dark matter distribution: they are the icing on the cosmic cake. Computer simulations can follow the growth of structure in the dark matter in great detail. For example, the Millennium simulation studied the formation of structure in a large region of the Universe with unprecedented resolution. It formed millions of dark matter halos, each of which could contain many galaxies. Would these galaxies resemble the observed ones?

The simplest galaxy formation model assumes that gas shock heats as it falls into a dark matter halo, cools radiatively, and then turns into stars. However, what happens in real galaxies is more complex. The observed star formation rate in spiral galaxies is tightly connected to the thickness of their disk – a relation called the Kennicutt law, and is much less than predicted by the simple model. Clearly, some feedback process is operating that regulates star formation.

Energy released by super nova, and by forming black holes, can quench star formation in galaxies. A strong burst of star formation, for example, triggered by a collision between two galaxies, may result in a large fraction of the gas in the galaxies being flung into intergalactic space, powered by hundreds of exploding stars. Such “galactic winds” are observed in some nearby galaxies, and also in galaxies at higher redshifts. Metals produced by stars in the galaxies are also flung into space.

Figure 1: The figure depicts structure formation in a computer simulated Universe covering a dynamic range of a factor of 10 000 in linear scale. The left most image shows the Millennium simulation (Springel at al 2005) which models the distribution of dark matter on very large scales. The central image shows the results of a simulation of a particular region taken from of the Millennium simulation which has been resimulated at higher resolution than the Millennium simulation, and includes baryonic matter. This simulation was run as part of the GIMIC project, the aim of which is to model galaxy formation and its affects on the intergalactic medium. The GIMIC project used time on HPCx and was run under the auspices of the Deisa Extreme Computing Initiative. The right most image shows one example out of many of a disc galaxy forming within the GIMIC high resolution region. © Virgo Consortium

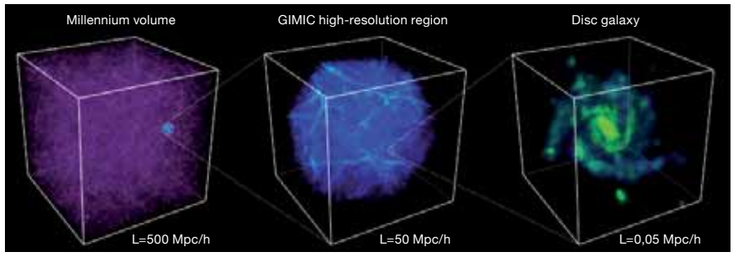

Observations of distant quasars can trace these metals in great detail, constraining galaxy formation models.

Simulations of galaxy formation need to be able to model these processes to produce realistic galaxies. Unfortunately, this means they need to resolve individual galaxies in great detail, which makes it much harder to study a cosmologically representative volume as well.

Full hydrodynamics simulations

GIMIC (Galaxies-Intergalactic Medium Interaction Calculation) is a project of the Virgo consortium to study galaxy formation in a cosmological setting. It simulates the formation of galaxies using full hydrodynamics, in five regions picked from the Millennium simulation. These regions include a high density region, which will form a massive cluster, but also mean, and low density regions (voids). This allows us to investigate galaxy properties as a function of their large-scale environment. Crucially, the GIMIC simulations also have high enough numerical resolution to follow all the physics of gas cooling, star formation, feedback and galactic winds in great detail.

Figure 1 shows a slice through the full Millennium simulation. The purple background depicts the large-scale distribution of dark matter, with its characteristic pattern of filaments surrounding low-density voids. The whole simulation box is 500 Mpc on a side, containing millions of galaxies. The blue region is one of our chosen high-resolution regions, where we follow the formation of galaxies in much more detail, including full hydrodynamics. Two regions in the high resolution region are shown in more detail to the right. One contains a proto-cluster region, with gas heated to very high temperatures, and the forming galaxies engulfed in a cocoon of metal enriched gas flung into space by galactic winds. A lower density region contains a spiral galaxy, with a well developed spiral disk. Clearly, the different environmental properties, low density region versus high density proto-cluster, are indeed affecting the properties of the forming galaxies.

These simulations provide unprecedented insight in how galaxies form, and also allow a detailed analysis into how galaxies interact with the surrounding intergalactic medium. It is possible to compare these simulations to observations at high-redshifts, where the intergalactic medium can be probed in great detail using quasars. At lower redshifts, the intergalactic medium in the filaments gets hotter and can be observed using X-ray telescopes. It is thought that most of the gas in the Universe is contained in this “Warm-Hot Intergalactic Medium”.

A new simulation code developed

Virgo (http://www.virgo.dur.ac.uk/new/index.php) developed a new simulation code to be able to perform these demanding simulations. This code, based on Gadget-2, provides superior load-balance and speed to its predecessor, and has also been enhanced with new physics modules. Virgo is a leading international consortium of numerical cosmologists, with members in the UK, Germany (Max-Planck Institut für Astrophysik), the Netherlands (Leiden), the US and Canada. These simulations were performed within the DEISA Extreme Computing Initiative (DECI) of the Distributed European Infrastructure for Supercomputing Applications.

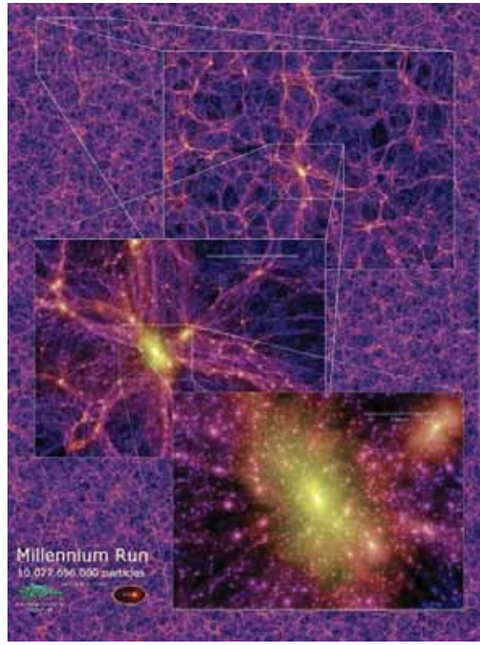

Figure 2: The image shows the projected density field of a slice through the Millennium simulation. The overlaid panels show successive zooms by a factor of 4, centred on a large dark matter halo. © Volker Springel, Virgo Consortium

Fabio Pasian, Istituto Nazionale di Astrofisica - Osservatorio Astronomico di Trieste, Italy, Hannu Kurki-Suonio, University of Helsinki, Finland, and Samuel Leach, International School of Advanced Studies, Trieste, Italy

Planck satellite simulations

DEISA's supercomputing framework was used to simulate several times the whole mission of the Planck spacecraft's low frequency instrument, LFI. The project lead by Fabio Pasian concentrated on understanding the effects derived from the optical behavior of Planck receivers.

The LFI is an array of 22 tuned radio receivers that will operate at –253°C on board the Planck spacecraft. The LFI will image the sky at three frequencies between 30 and 70 GHz. ESA is preparing Planck to be launched in July 2008 with a mission to collect and characterize radiation from the Cosmic Microwave Background (CMB) using sensitive radio receivers operating at extremely low temperatures.

Data on how the Universe began

CMB is relic radiation from the Big Bang, and ever since the detection of small fluctuations in the temperature of this radiation, announced in late 1992, astronomers have used the fluctuations to understand both the origin of the Universe and the formation of galaxies.

Planck will measure the temperature variations across this radiation background and provide a map of the Cosmic Microwave Background field at high angular resolution, covering at least 95% of the sky over a wide frequency range.

The LFI-sim proposal aimed at using the supercomputing framework provided by DEISA to simulate several times the whole mission of Planck's LFI instrument, on the basis of different scientific and instrumental hypotheses, and to reduce, calibrate and analyze the simulated data down to the production of the final products of the mission, in order to evaluate the impact of possible LFI instrumental effects on the quality of the scientific results, and consequently to refine appropriately the data processing algorithms.

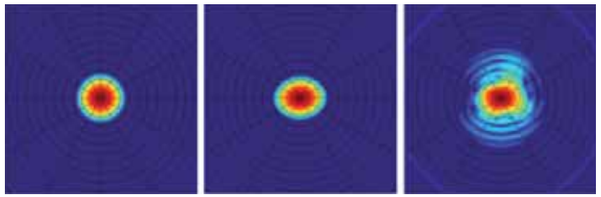

Figure 1: Simulated response of the LFI-28 beam, observing the sky at the frequency of 30 GHz. From left to right: ideal beam response (circular); response non-ideal at the first order (elliptical); realistic response (on the basis of laboratory measurements). © Fabio Pasian

The LFI-sim project within DECI concentrated on understanding the effect of optical systematic effects, i.e., the effects derived from the optical behavior of Planck receivers. In particular, one of the critical aspects previously impossible to evaluate was tackled, i.e., a complete study of the effect of observing the sky with realistic beams on the scientific results of the mission. This was successfully achieved.

Simulating the LFI data

The logical sequence for simulating and processing LFI data has been the following:

- from cosmological parameters, generate ideal CMB sky; optionally, add foregrounds to obtain an ideal reference sky at all LFI frequencies;
- "observe" the ideal sky with a numerical model of the LFI instrument (in this case with realistic beams and noise) and obtain time series of observed data; process, optionally removing systematic effect(s), the time series and obtain frequency maps;
- separate and remove foregrounds to obtain the CMB "observed" map;
- build the "observed" CMB power spectrum and compare with the predicted one.

Results

Several simulation runs were made, and part or all of the above steps have been performed for each run. The basic approach has been to work on ideal CMB maps and to evaluate the impact of realistic beams on the "observed" CMB maps, and on the resulting power spectrum: this has allowed to single out the effect of distorted beams on the scientific results of the mission, isolating possible interactions with other data processing steps.
This approach was followed also when the first laboratory measurements of LFI beam shapes became available, with accurate values for the sidelobes, which mostly contribute to distortion. It was found that the measured effect of beam distortion on the scientific results (i.e., the power spectrum) is of the third order, but is not completely negligible and should be considered when the most refined results are to be produced.

For this case, however, the beam deconvolution code available at the moment is not mature enough to consider realistic noise, and needs further refinement. Some simulation runs were performed to build and process full-fl edged reference skies in various frequencies, with the purpose of understanding possible inter-dependencies of the various data processing steps. It was found that the processing steps are separable up to the numerical precision of the simulation.

Some interesting findings were achieved on the algorithmic aspects of the pipeline. First, the results of the optimal IGLS map-making algorithm are obtainable (within a level of error smaller that the intrinsic instrumental error) using hybrid codes, which combine concepts from both destriping and maximum-likelihood map-making, but are much more efficient to run. IGLS map-making could be run on the Data Processing Centre, DPC, hardware only to produce the "final" most accurate results.

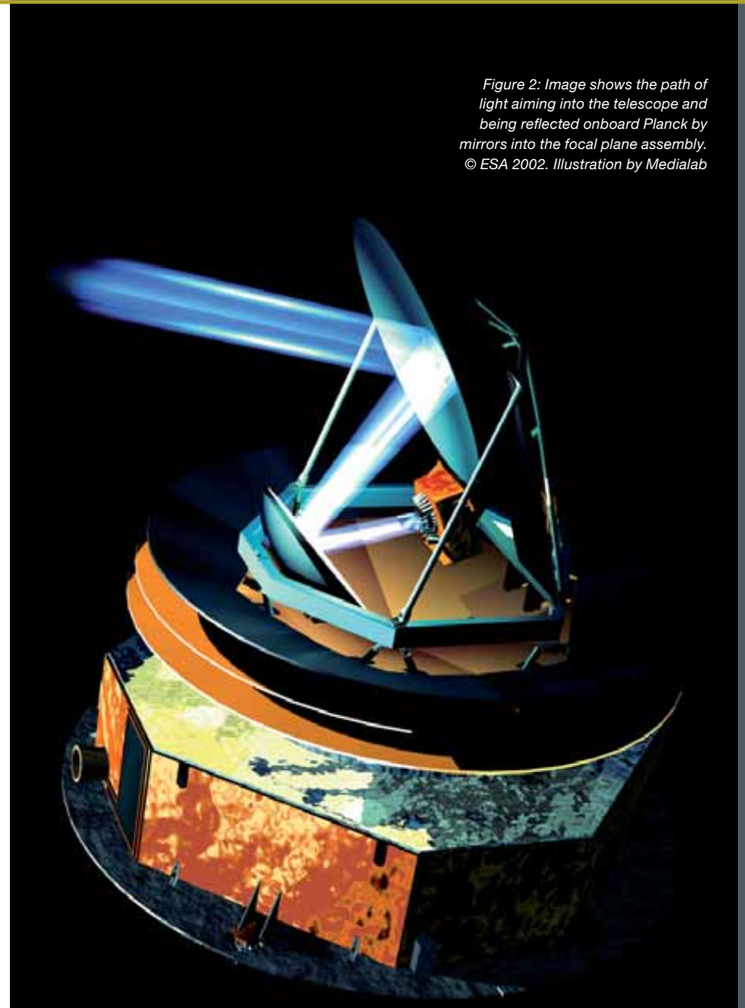

Figure 2: Image shows the path of light aiming into the telescope and being reflected onboard Planck by mirrors into the focal plane assembly. © ESA 2002. Illustration by Medialab

Gustavo Yepes, Universidad Autónoma de Madrid, Spain. Stefan Gottlöber, Astrophysikalisches Institut Potsdam, Germany.
Anatoly Klypin, New Mexico State University, USA. Yehuda Hoffman, Racah Institute of Physics, Israel

Simulating the Local Universe

THE SIMU-LU PROJECT HAS BEEN ABLE TO SIMULATE THE FORMATION AND EVOLUTION OF A SYSTEM PRETTY SIMILAR TO OUR OWN GALAXY AND ITS CLOSEST NEIGHBOUR: THE ANDROMEDA GALAXY.

We have carried out the most accurate representation of the formation of the Local Universe performed to date, starting from cosmological initial conditions compatible with the most recent astronomical observations. To this end, we have been running a set of high-resolution N-body cosmological simulations from initial conditions which include observational constraints on the mass distribution and velocity fields derived from nearby galaxies and clusters around us. We were able to run these simulations with up to 1 billion dark matter particles in different computational volumes. The simulations reproduce the main features of the mass distributions we observe in our local neighbourhood such as the Coma and Virgo clusters, and the Great Attractor and Local Superclusters.

By zooming in to the region where we are supposed to live, the Local Group, we have been able to simulate the formation and evolution of a system pretty similar to our own galaxy and its closest neighbour: the Andromeda Galaxy. The outcome of these numerical experiments will provide us with a deep insight into the dynamics of our local environment and will constitute the starting point to do more realistic simulations, in which ordinary matter (i.e. gas and stars) will be included together with the more exotic, yet dominant, components of the universe: Dark Matter and Dark Energy.

Advanced simulation methods

To simulate the evolution of a self-gravitating fluid in an expanding universe, we need to solve a coupled system of the 3D Poisson and Vlasov equations. These are the standard equations of motion for particles interacting via gravity, simultaneously solving Poisson's PDE for the gravitational potential and Newton's ODE for the acceleration of each particle with the particularity that the phase space variables are implicit functions of time due to the expansion of the universe.

We have used two of the most advanced cosmological simulation N-body codes developed till now. The Adaptive Refinement Tree (ART) code employs standard particle-mesh techniques to compute particle accelerations and advance their coordinates and velocities in time. A regular cubic grid covers the entire computational volume and defines the minimum resolution of the simulation. This grid is then refined to form additional higher resolution meshes in the regions of interest. The GADGET code, on the other hand, uses a combination of Particle Mesh algorithms to compute the long range gravitational forces and a Tree algorithm with monopole expansion of the potential to account for the short scale force. Both codes use domain decomposition techniques to split the particle and mesh information among the different processors and use MPI communication among different nodes. They have been extensively tested in large distributed memory, SMP and NUMA architectures such as the MareNostrum (BSC) and HLRB2 (LRZ) supercomputers where these simulations have been done.

Studies of the dynamics of the Local Supercluster

The purpose of our DECI project was to perform very large cosmological simulations of 1 billion particles in a computational cubic volumes ranging from 64 to 160 Mpc (i.e. 200 to ~500 million light years) on a side. We need to resolve objects that are expected to be formed if dark matter is in the form of cold, weakly interacting massive particles (WIMPS) with differences in mass over many orders of magnitude: from the largest superclusters of galaxies with masses of up to 10^{16} solar masses or more, to the tiniest dwarf galaxies orbiting around the normal ones (with masses less than 10^{10} solar mass).

The novelty of our approach resides in the generation of the initial conditions. We impose observational constraints onto the random realizations of the initial density fluctuation field. Thus, we force the formation of structures similar to those we see today. We have the advantage of dealing with known, well-studied astronomical objects. This significantly simplifies the comparison of the model with the observational data. For instance, we are able to study many details of the dynamics of the Local Supercluster to which our Local Group of galaxies belongs. The results of these simulations will constitute an excellent database to study in detail many aspects of the distribution of matter in our backyard. Now, the next step is to include the ordinary matter into the simulation so that we can witness the birth of stars inside the dark matter halos and make realistic comparison with the most recent astronomical observations of the real objects in our universe.

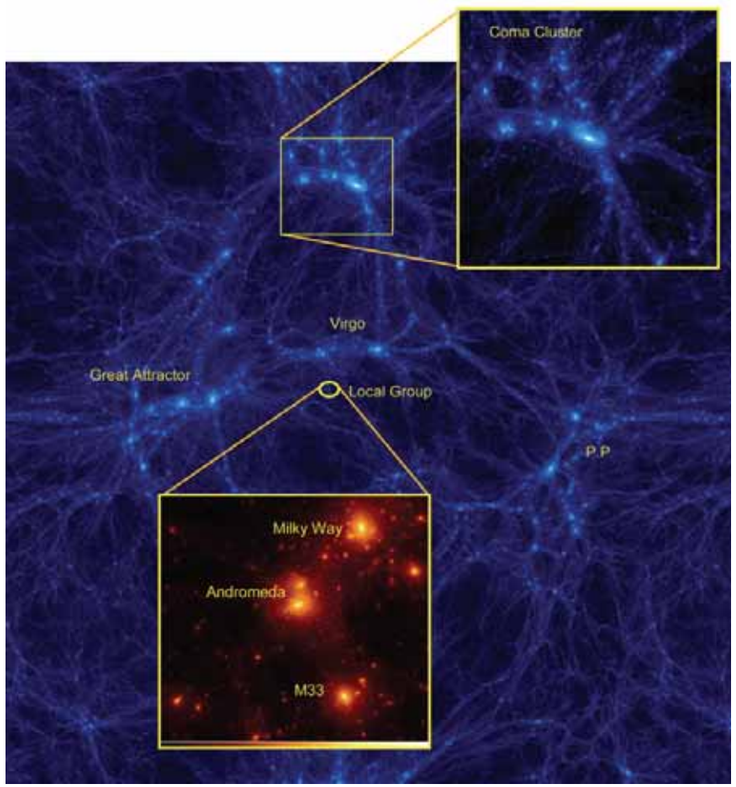

Figure 1: Dark matter density distribution of the Local Universe simulation. The boxsize is 160 Mpc across. Several objects of the real universe are identified. The circle shows the position of our Local Group. A blown up panel shows the detailed structure of the simulated Local Group. The resolution of this small region of 2 Mpc radius is equivalent of having a total of 4096^3 (70 billion) particles in the whole box, which translates in a dynamical mass range of more than 10^6.

Allan Sacha Brun, CEA-Saclay, France

Turbulent, active, and rotating stars

Allan Sacha Brun has studied the interaction between convection, rotation, and magnetic field in stars. Within the DECI project STARS it has been possible to appreciate the difficulty in computing a highly nonlinear and turbulent 3D MHD solar dynamo model at low magnetic Prandtl number. For the first time the dynamo action in a turbulent convective sphere at low magnetic Prandtl number was achieved in the solar context.

The STARS project aims at modelling in a self-consistent and three-dimensional way the complex, time dependent, and nonlinear dynamics present in the Sun and stars. In particular we wish to understand stellar magnetic activity, that depending on the spectral type of the star considered can be cyclic (solar type stars), irregular (very low mass stars, with spectral type later than M3), or even for stars with stellar mass greater than 2 solar mass, without any activity or simply possessing a modulated signal (probably due to the presence of a fossil field in their stably stratified, radiative envelope).

The mechanism thought to be at the origin of the magnetism seen in solar type stars or in low mass stars is likely to be linked to dynamo action in the upper convective layers of such stars. The simultaneous existence of convective turbulent motions (that could even possess helicity), of rotation and its associated differential rotation and shear layers in stars, favour the emergence of a small and/or large scale magnetic field through induction. For more massive stars, possessing a convective core, understanding the interaction between the dynamo generated magnetic field and the probable fossil magnetic field of their radiative envelope constitute a major challenge in stellar fluid dynamics.

To study in great details the interaction between convection, rotation, and magnetic field in stars is the main scientific goal of this project.

Solar dynamo model at low magnetic Prandtl number

This DECI project has been extremely fruitful in terms of understanding the complex interplay between turbulent convection and magnetism and to appreciate how difficult it is to compute a highly nonlinear and turbulent 3D MHD solar dynamo model at low magnetic Prandtl number. It is the first time that dynamo action in a turbulent convective sphere at low magnetic Prandtl number is achieved in the solar context.

We can already say that the level at which dynamo action is successful against Ohmic decay is higher than in the corresponding case with high Pm number published in Brun et al. 2004.

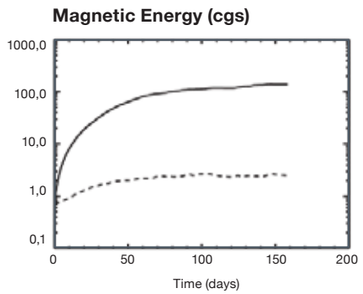

Figure 1: Temporal evolution over the 160 days of simulated time of the total and mean (dashed line) magnetic energy in a turbulent DECI simulation of the solar convection at a Prandtl number Pm of 0.8. ©Allan Sacha Brun

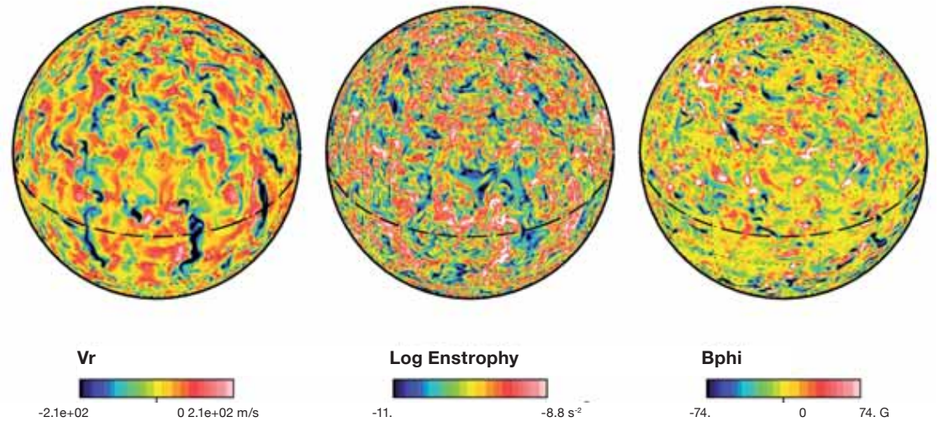

Figure 2: Snapshot of the radial velocity, log of enstrophy and toroidal magnetic field in the bulk of the highly turbulent convection zone of the DECI run. Highly intermittent convection and magnetic fields are observed in this first low Pm simulation of the solar convective envelope. We note the high degree of vorticity present in the down flows. ©Allan Sacha Brun

We have computed several models in order to reach the dynamo threshold while keeping a solar-like differential rotation profile. The first model had a magnetic Reynolds number of around 300 (resolution Nr=256, Ntheta=512 and Nphi=1024), the threshold determined by our previous study. This model did not succeed and the seed magnetic finally decayed away. We then progressively increased the level of turbulence while keeping Pm=0.8 and we had to reach about Rm=400 (and a resolution of 256×784×1568) to get a successful dynamo. This seems to indicate an increase of about 30 per cent of the dynamo threshold with respect to the high Pm cases, and confirms that turbulence is actually making it harder to get a successful dynamo rather than easier (see for example Ponty et al. 2005).

Evolution of the magnetic energies

We display on Figure 1 the temporal evolution of the total (solid line) and mean (dashed line) magnetic energies of our latest successful DECI run. We can notice the linear phase (exponential growth) of the magnetic energy and then its nonlinear saturation via the feed back of the Lorentz forces on the flow in particular in region of high vorticity through Maxwell stresses. The mean (axisymmetric) magnetic energy is found to be small (~ 1%) confirming that turbulent magnetized convection generates mostly non axisymmetric and highly intermittent fields, characteristic of a small scale dynamo process.

To illustrate the richness of the simulations, we show in Figure 2 a snapshot of the convective radial velocity, of the $\log_{10}$ of the enstrophy (square of the vorticity) and the associated dynamo toroidal magnetic field in the bulk of the modelled convection zone. We see how turbulent the convective patterns are and how small scale and intricate the magnetic field can be. This indicates that the magnetic fields generated by dynamo processes in stellar plasma are likely to be disorganized (independently of the magnetic Prandtl number considered or possible inverse cascade processes), and that in order to get a large-scale organized field (mostly toroidal in nature), one needs to include an omega effect not only in the convection zone (as it is already the case in this simulation) but also in a sheared stable layer like the tachocline. A preliminary low resolution study seems to indicate that it is indeed the case; i.e., a tachocline of shear is efficient at making stronger axisymmetric fields.

Friedrich Röpke and Wolfgang Hillebrandt, Max-Planck-Institüt für Astrophysik, Garching, Germany

Simulating Type Ia supernovae

A sound theoretical understanding of the explosion mechanism is one of the great challenges of astrophysics. A researcher group at the Max Planck Institute for Astrophysics has pursued a promising approach. By constructing numerical models of Type Ia supernovae explosions and by comparing the results with detailed observations of nearby Type Ia supernovae, new insights into the explosion mechanism have been gained.

Only about once in a century, a "new star" might appear on the night sky, visible to the naked eye, which fades away over some weeks. A sub-group of these events – classified as Type Ia supernovae (SNe Ia) – are among the most energetic explosions in the universe. Being rare in single galaxies such as the Milky Way, the majority of SNe Ia occurs far away.

Due to their enormous intrinsic brightness, outshining billions of stars, large telescopes facilitate the observation of supernovae that went off when the universe was only about half its present age. Therefore, they can be used as "lighthouses" and "standard candles" to determine the geometry and the expansion rate of the universe. Here, one exploits the fact that SNe Ia are (at least by astronomical standards) remarkably uniform. Knowing the intrinsic luminosity, their apparent brightness tells us their distance.

These measurements indicated that the universe undergoes an accelerated expansion at present – a result which fits into the theoretical framework of General Relativity, given that a yet unknown form of "dark" energy is assumed to dominate the universe today. Thus, SNe Ia may have given birth to a major revision of our understanding of physics and of the universe.

Theoretical understanding of the explosion mechanism

However, a close look to the sample of well-observed nearby SNe Ia reveals an intrinsic scatter in their brightnesses which is evidently connected to other properties. Such correlations are used to calibrate the cosmological distance determinations. But, as yet, these are established only empirically.

A sound theoretical understanding of the explosion mechanism is therefore one of the great challenges of astrophysics. Only models that tackle the underlying physics in a self-consistent way are predictive enough to provide answers concerning the reliability of SNe Ia as cosmic distance indicators. A promising approach is pursued at the Max Planck Institute for Astrophysics. By constructing numerical models of SN Ia explosions and by comparing the results with detailed observations of nearby Type Ia supernovae, new insights into the explosion mechanism are gained.

Although a consensus about the astrophysical scenario exists, the exact explosion mechanism is unclear. SNe Ia are attributed to the thermonuclear explosion of White Dwarf stars, of about 1.4 solar masses, which consist of carbon and oxygen. The behavior of such a final event of the star life implies a very complex variety of phenomena that make SN Ia explosions a problem of turbulent combustion much like the burning of fuel in car engines.

Challenging numerical simulations

The challenge of implementing this scenario into numerical simulations lies in the vast range of relevant length scales. While the radius of the exploding White Dwarf star amounts to ~2000 km, the width of a flame is typically less than a millimeter. In order to model turbulence effects consistently, three-dimensional simulations are inevitable, and it is impossible to resolve the full range of scales with a brute approach. Thus different correlated methods and models have been used to code the solution and the resultant simulation was carried out within the DEISA framework on HPCx, the IBM cluster in the UK using 512 processors. The evolution of the explosion process was obtained using 640^3 cells. The results are illustrated in Figure 1.

Results compared with actual observations

Based on this simulation, observables can be predicted and these will then be compared with actual observations to assess the validity of the SN Ia model. This pipeline of post-processing steps is a pan-European effort including research groups from Germany, Italy, Sweden, and Russia. The distributed DEISA infrastructure will continue to provide the resources for this task. Indeed, DEISA is proud to host this project, particularly since its relevance in cosmology is documented by the choice of the Scientific American newsletter to use its results as the cover story ("How to Blow Up a Star") for the October 2006 issue.

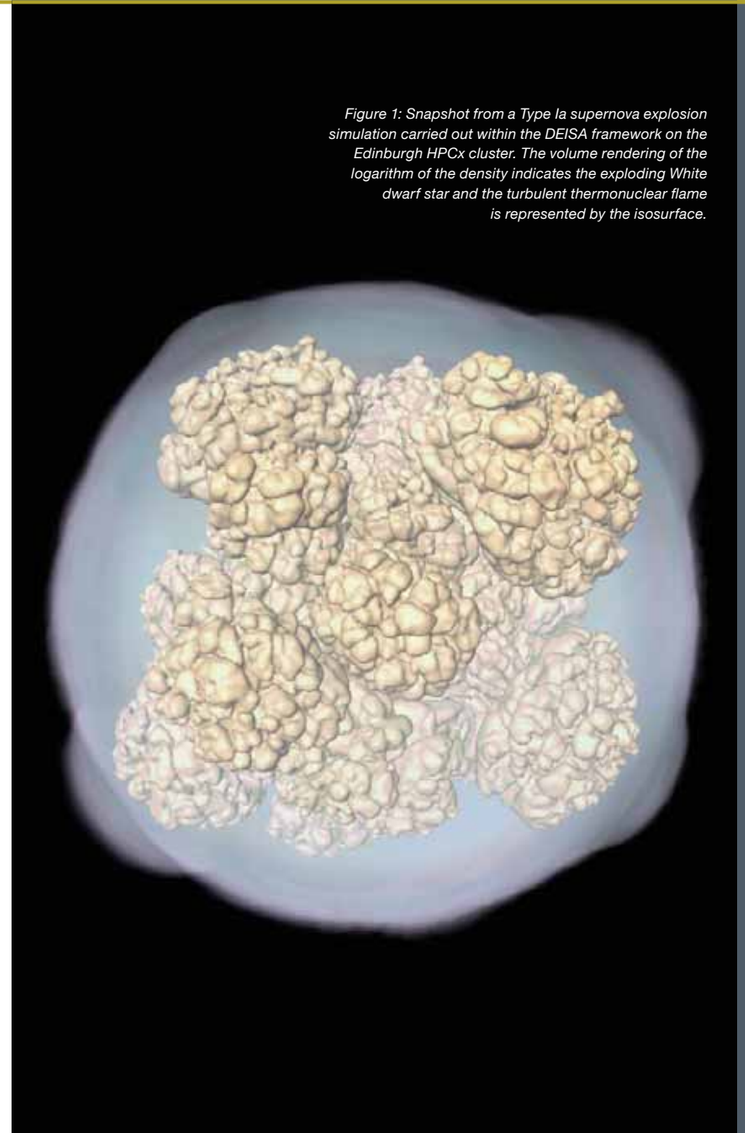

Figure 1: Snapshot from a Type Ia supernova explosion simulation carried out within the DEISA framework on the Edinburgh HPCx cluster. The volume rendering of the logarithm of the density indicates the exploding White dwarf star and the turbulent thermonuclear flame is represented by the isosurface.

Bernd Brügmann, Jose Gonzalez, Mark Hannam, Sascha Husa, Ulrich Sperhake
and Gerhard Zumbusch, University of Jena, Germany

Gravitational wave templates from simulations of binary black holes

This study concerns the inspiral of binary black hole systems in full general relativity, solving the Einstein equations as partial differential equations with finite difference methods. This work has produced gravitational wave templates for non-spinning binary black hole coalescences, and we have obtained astrophysically relevant results on the remnant black holes of mergers. Numerical simulations of the strong field regime have been linked to the regime where perturbation calculations are valid, which yields "complete" waveforms of the inspiral process. The most spectacular result has been that the black hole which forms the product of the merger may recoil with velocities of more than 2000 km/h, which is sufficient to eject black holes even from the most massive elliptical galaxies.

Figure 1: "Plus"polarization of the gravitational wave strain versus time, plotted in geometric units. The model corresponds to two equal mass non-spinning black holes and spans approximately nine orbits of inspiral, the merger and ringdown of the resulting spinning black hole.

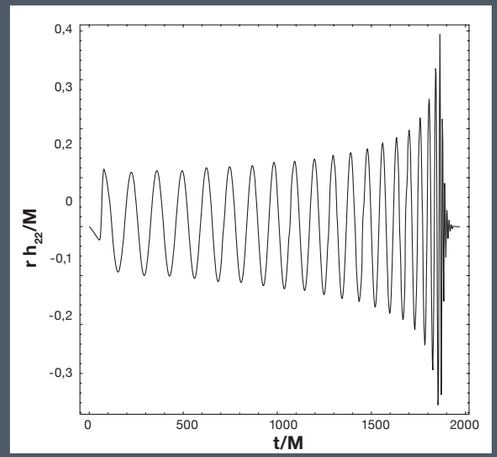

General relativity predicts the emission of gravitational waves in astrophysical processes, which is analogous to the production of electromagnetic waves by accelerated charges in electromagnetism. All the current knowledge about astrophysics and cosmology is based on electromagnetic observations, but observations of gravitational waves will open a new window into the universe and provide information about phenomena hitherto not accessible to direct observation, such as black holes, dark matter, or the very early universe. The detection of gravitational waves has not yet been accomplished, but a growing network of gravitational wave detectors is currently collecting data at design sensitivity. In order to actually extract information on the sources from observations, accurate signal templates are needed for various types of sources.

Until recently, only templates based on analytical approximation methods, in particular post-Newtonian expansion methods, have been available for the inspiral of binary systems of compact objects, such as black holes formed from a radiation reaction. Computations of the inspiral with numerical methods in full general relativity are required to benchmark existing search algorithms based on post-Newtonian templates and extend the validity of the search methods into the strongly non-linear merger regime. Furthermore, accurate modelling of this non-linear regime is necessary to predict the properties of the final black hole, such as the angular momentum and velocity, which provide important inputs for astrophysical modelling.

Moving puncture

We have used the "moving puncture" approach to solve the full Einstein equations numerically for the initial data corresponding to astrophysically relevant systems of inspiraling black holes with negligible eccentricity. The "moving puncture" approach loosely defines a particular setup of the continuum problem corresponding to the Einstein equations as a set of nonlinear hyperbolic partial differential equations. We discretized these equations with fourth and sixth order accurate finite difference operators and solved them with a simple box-based adaptive mesh refinement algorithm, where the refinement boxes tracked the orbital motion of the black holes. The code was parallelized using domain decomposition with MPI.

Black hole recoils

A careful analysis of simulations of a single black hole resulted in the description, for the first time, of the geometrical behaviour of numerical solutions in the now-standard "moving puncture" method of evolving black hole spacetimes. All moving-puncture simulations to date have used a particular type of initial data, Bowen–York puncture data. These data make the physically unrealistic assumption of conformal flatness. Significant progress has been made in removing this assumption for binaries made up of spinning black holes.

Astrophysicists have for many years been waiting for numerical simulations to accurately calculate the recoil when two black holes of unequal mass collide. In the largest parameter study to date of black-hole binaries requiring approximately 150 000 CPU hours, the maximum recoil of unequal-mass non-spinning binaries was calculated to be 175.2±11 km/s. Far more spectacularly, we found that extremely large recoils are possible for spinning black holes, and estimated the kick from one configuration as 2 500 km/s, large enough to eject the remnant black hole even from a giant elliptical galaxy.

Template banks

A joint project between the MPI for Gravitational Physics, the University of Jena and the University of the Balearic Islands has been launched to provide numerical waveforms for the gravitational-wave data-analysis community through the construction of template banks. Further work is in progress on testing the current implementation of the LSC data analysis pipeline by injecting numerical waveforms and recovering them using current template banks.

Long accurate evolutions have been rendered possible by extending our methods to sixth order accurate finite differencing operators, which allowed simulations of nine orbits with an overall phase error of less than 0.13%.

Non-eccentric binaries

A long standing problem in the field has been to establish physically realistic initial data for the non-eccentric inspiral of binary black holes. A major step forward was to show that, at least for the non-spinning equal-mass case, post-Newtonian results can be used to specify the momentum parameters of inspiraling black holes with sufficient accuracy so that the resulting binary has no significant eccentricity.

A detailed comparison of long gravitational wave signals for the non-spinning equal mass case was performed, demonstrating that at least five orbits are necessary if the amplitude disagreement between numerical and post-Newtonian waveforms is to be within numerical error at the matching point of the two approaches.

S. Raugei, A. Magistrato, and P. Carloni SISSA and INFM Democritos Center, Trieste, Italy; M. Sulpizi and M. Sprik, University of Cambridge, UK; J. VandeVondele, University of Zurich, Switzerland; M. Cascella and U. Rothlisbeger, EPFL SB ISIC LCBC, Lausanne, Switzerland

Biological electron transfer simulated

A researcher group lead by SISSA and INFM Democritos Center, Italy, has used state-of-the-art first-principles techniques to investigate electron transfer in the context of iron- and copper-based proteins. For the first time ab initio investigation of entire electron-transfer proteins in aqueous solution was performed. This pioneering study was made possible with the large computational resources granted by DEISA.

Electron-transfer (ET) proteins are crucial for the life of organisms. They serve as electron carrier and/or as catalyst in biochemical ET reactions. As such they are responsible, for instance, of the respiration and the photosynthesis. A detailed understanding of the involved molecular mechanisms is thus of prominent interest from a fundamental point of view as well as for the potential design of highly efficient bio-inspired optoelectronic devices and solar cells.

Studying electron-transfer proteins at full atomic level

The study of ET proteins at full atomic and electronic level represents a new challenge for computational biochemistry. These proteins contain one or more metal ions in the active site, such as iron and copper. During the ET process the metal changes oxidation state (i.e., Fe(III/II) and Cu(II/I)) and this is usually accompanied by a reorganization of the environment surrounding it. The importance of a molecular description of the solvent and the proper inclusion of the polarization medium around the metal center has been largely recognized for a long. As yet, these effects limited the applicability of standard computational schemes.

We have used state-of-the-art first-principles electronic-structure techniques to perform the first *ab initio* investigation of entire electron-transfer proteins in aqueous solution. We exploited a combined approach, which used empirical force field based molecular dynamics coupled to full quantum mechanics (linear scaling density functional theory) and hybrid quantum mechanics/molecular mechanics (QM/MM) calculations, to get insights on the nature of the drastic decrease of reorganization energy for metal ions in the protein frame. This pioneering study has been made possible with the large computational resources made available through the DEISA initiative.

Iron- and copper-based proteins on focus

We focused on a copper protein (azurin for *Pseudomonas Aeruginosa*) and two iron-sulfur proteins (rubredoxin from *Clostridium pasteurianum* and from *Pyrococcus furiosus*). The calculations were carried out at IDRIS in Paris and at CINECA Supercomputer Centre in Bologna with the codes CPMD and CP2K. The full *ab initio* (mixed Gaussian and Plane Waves) CP2K calculations performed on rubredoxin, which represented the most expensive part of our study, took explicitly into account the entire protein, 700 water molecules, and 9 counter ions, for a total of 2825 atoms (the Gaussian basis set for the wave functions included 11 334 primitives). For this system, each electronic wave function optimization required 4.5 hours on 128 processors of the IBM SP4 at IDRIS. To have reliably converged quantities, hundreds of such calculations were required.

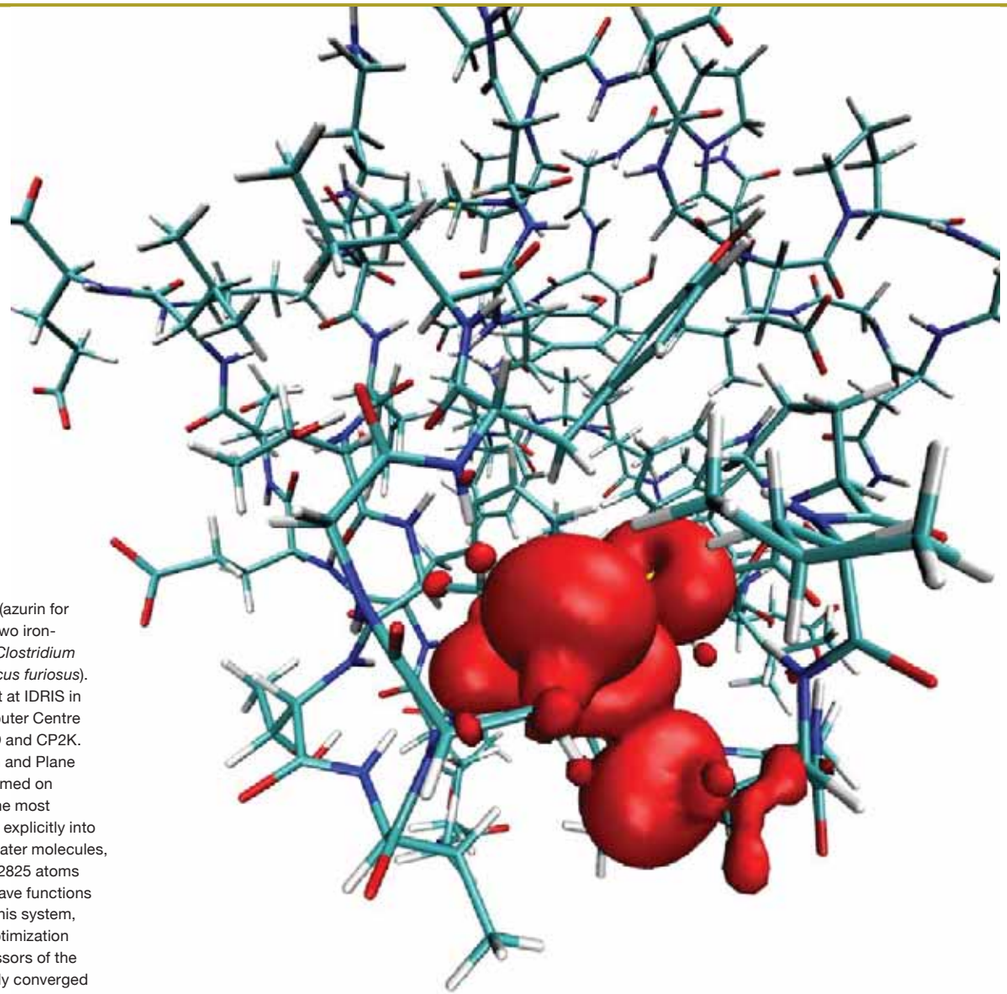

Results

We were able to reproduce the experimentally determined tendency of the proteins to acquire electrons (i.e., the redox potential) and to give the first estimate for the reorganization free energies. In particular, in rubredoxin, we gave an important contribution to unravel the intricate way the protein modulates the reorganization free energy in a balance between shielding of the high dielectric medium and preserving some water accessibility to the metal center for a fast electron transfer to the surface. We could rationalize the magnitude of the reorganization free energy in terms of different dielectric response of the solvent around the active site.

Figure 1: Isosurface enclosing 95 per cent of the total spin density (difference of the up and down electron densities) in the oxidized form of the rubredoxin from Clostridium pasteurianum. All of the atoms (included the solvent which is not show for sake of clarity) were explicitly taken into account at ab initio (DFT) level.

Marie-Pierre Durrieu and Marc Baaden, CNRS France. Peter Bond and Mark Sansom, University of Oxford, UK.
Elmar Krieger, Radboud University Nijmegen, the Netherlands. Richard Lavery, CNRS UMR / Université de Lyon, France.
Laurent Leger, IDRIS, Orsay, France. Nada Taib and Michel Laguerre, Université de Bordeaux, France.

Membranes under tension

This research group has modelled a protein complex that acts as a catalyst in the fusion of two cell membranes. Understanding the guiding principles of membrane fusion will open new opportunities for pharmaceutical development.

In the DEISA Extreme Computing Initiative, more than 200 000 CPU hours were allocated for the atomistic modelling of the membrane-embedded synaptic fusion complex. Membrane fusion is a ubiquitous process in biology. Simulations allow us to better understand the guiding principles of these biological processes and form part of the grand challenge applications, for which HPC resources are key in being able to develop realistic models. This fusion complex, a target for the study of several pathologies, such as botulism and tetanus, requires exceptional computing resources due to both the size of the system and the complexity of its environment. The model requires not only access to huge computing times but also to large amounts of memory and disk space, such as the facilities available to the European scientific community through the DEISA infrastructure.

SNARE proteins drive membrane fusion

Membrane fusion is essential for a number of biological processes: intracellular traffic, infection by enveloped viruses, fertilization, secretion of hormones and neurotransmitters. Membranes do not fuse spontaneously due to the many forces opposing their approach. Fusion is achieved via specialised protein complexes involving one key family of proteins, SNAREs (soluble N-ethylmaleimide-sensitive factor attachment protein receptors), in particular. A detailed understanding of the molecular mechanisms involved is of key interest from a fundamental point of view, as well as for the potential design of new drugs.

One particularly intriguing point is how the SNARE complex inserts itself into the two adjacent vesicular and target membranes to catalyze their fusion. A detailed atomic view of this system is needed in order to understand the fine interactions between lipids, and protein and solvent environments. Studying the membrane-embedded SNARE complex is a challenging task due to the complexity of its environment. Taking into account the transmembrane parts of the complex and the lipid bilayers to which they are anchored requires important computing resources and long simulation times.

All-atom molecular dynamics simulations

We have used long classical molecular dynamics simulations of a full atomic model of the membrane embedded SNARE complex using Gromacs software. The biggest computational challenge was the initial modelling and construction of the 340 000 atom simulation system and the ability to equilibrate and simulate it for a substantial amount of time (~50 ns). For the initial construction we used a combined multi-step approach, which required a workflow based on interactive molecular dynamics and several equilibrium and non-equilibrium simulations.

New insights gained into the SNARE complex

Simulations were carried out using 96 processors on an IBM SP Power 4 system at the national DEISA site IDRIS and at the Rechenzentrum Garching. The latest stable and development versions of Gromacs software were used. The simulated timescale was 40 ns corresponding to 20 million iterations. This equals a real timescale of 99 days and a total of 218 000 CPU hours. The atomistic model comprised 339 792 atoms consisting of four proteins (346 amino acids), two charged POPC:POPS lipid bilayers (1008:123 lipids), 296 Na+ ions, 166 Cl– ions and 92 217 water molecules.

Our study shows that all-atom molecular dynamics simulations of complex biomolecular assemblies under physiological conditions are now possible by making use of today's large-scale supercomputing facilities, such as those made available through the DEISA initiative. New insights into the SNARE complex have been gained and a detailed analysis of these results is currently under way. In particular, the role of electrostatics and direct SNARE-membrane interactions could be observed.

Membranes do not fuse spontaneously due to the many forces opposing their approach.

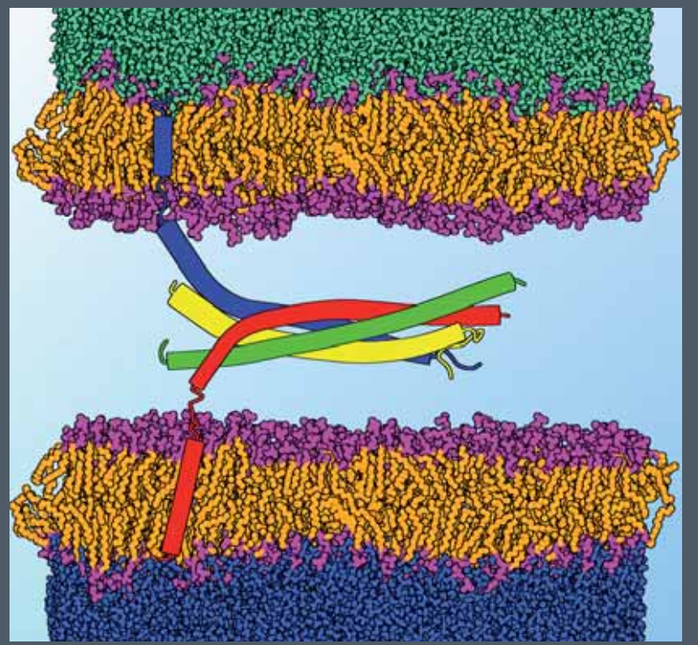

Figure 1: Schematic view of the SNARE complex embedded in two membranes. The upper membrane represents the vesicle, the lower membrane is the target with which fusion occurs. The four-helical SNARE bundle sits in between those membranes with synaptobrevin (blue) inserted into the vesicle and syntaxin (red) inserted into the target membrane. The third SNARE protein, SNAP-25, is represented in green and yellow. Water molecules in between the membranes are omitted for clarity.

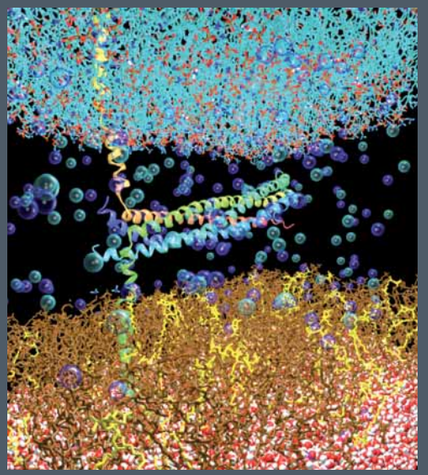

Figure 2: A perspective view on the membrane embedded SNARE system, highlighting the counterions Na+ and Cl- as bubbles. They are required to neutralize the high charge density of the mixed POPC (brown) / POPS (yellow) membrane. Water in between the membranes is omitted for clarity.

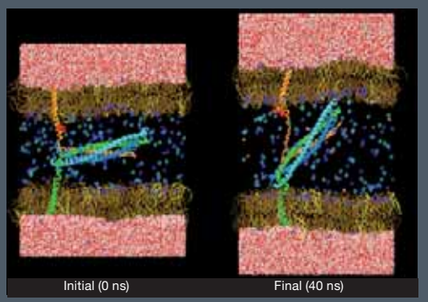

Figure 3: Initial and final structures of the simulations on the membrane embedded SNARE complex. Tryptophanes 89 and 90 of synaptobrevin (orange) are highlighted in red. The composition of the membranes is visible with neutral POPC in brown and charged POPS in yellow.

Roman Affentranger and Xavier Daura, Universitat Autònoma de Barcelona, Spain; Guido Scarabelli and Giorgio Colombo, Consiglio Nazionale delle Ricerche, Italy; Martin Zacharias, Jacobs University Bremen, Germany; Tsjerk Wassenaar and Alexandre MJJ Bonvin, Utrecht University, Bijvoet Center, the Netherlands

Molecular simulation to accelerate vaccine development

A large three-year project involving pharmaceutical and bioinformatics companies and Universities has exploited the DEISA resources to render vaccine development simpler, faster and therefore less expensive for pharmaceutical companies.

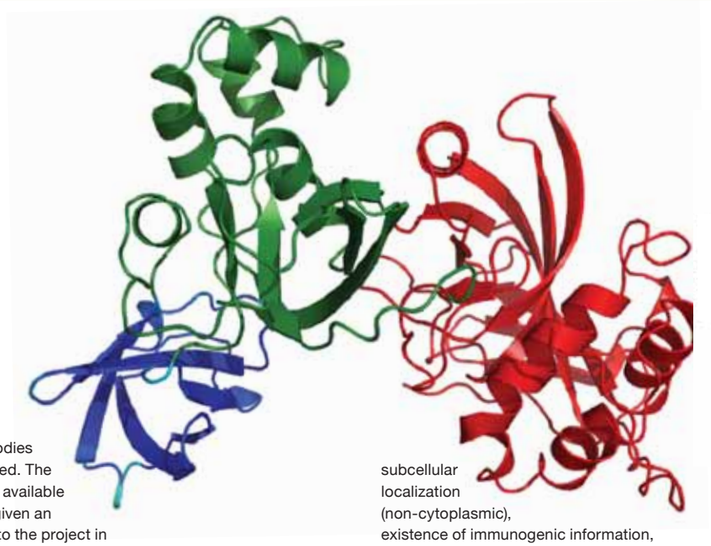

The BacAbsMS project accounts for the computationally intensive part of the FP6 STREP project BacAbs (www.bacabs.org). BacAbs aims at discovering correlations between the bactericidal potential of antibodies and properties of the proteins (antigens) that elicit their production, such as sequence, structure and dynamics. The access of BacAbsMS to the DEISA infrastructure has enabled the analysis of a significant number of protein targets from bacteria of medical interest (e.g., *Neisseria meningitidis*) using molecular dynamics simulation methods.

Dynamics of antigens eliciting protective and non-protective antibodies

High throughput cloning and expression of large sets of genomic ORFs has become a preferred industrial strategy for genome-wide searches of new vaccine candidates. For invasive infections in particular, the aim is to find proteins eliciting antibodies capable of binding to the bacterial cell surface and, through interaction with the complement system, effectively kill the bacteria. However, current data accumulating from reverse vaccinology studies (targeting of possible vaccine candidates starting from genomic information) shows that only a small fraction of surface-exposed proteins elicit antibodies with bactericidal activity. As a consequence, current (genomics-based) approaches to vaccine development are costly and inefficient. Importantly, no structural data is incorporated into the selection of antigens. Thus, there is a clear opportunity to merge experimental and computational methods to enhance the efficiency with which appropriate antigens can be identified. The BacAbsMS objective has been to provide information on the dynamics of antigens and antigenic epitopes eliciting protective and non-protective antibodies, which can be later used in establishing correlations between structural features and protection.

Although support from DEISA was for 2007, the BacAbs project extends over the period 2007–2009. From the various computational tasks of the project, during this first year emphasis has been put on the most CPU intensive ones. For the correlation analysis to be valid, a statistically significant number of bacterial proteins eliciting protective and non-protective antibodies need to be studied. The resources made available by DEISA have given an important push to the project in this direction (see Table 1).

Computational Approach

To study protein dynamics, classical molecular-dynamics simulation in combination with empirical biomolecular force fields have been used. Periodic boundary conditions have been applied, together with algorithms for constant temperature (300 K) and pressure (1 atm). The AMBER and GROMACS software packages have been used, with the param99-SB and GROMOS-45A3 force fields, respectively.

Project description and results

Table 1 shows a summary of the simulations performed. All calculations were carried out at BSC in Barcelona and CINECA in Bologna, using existing installations of the above-mentioned codes. The 12 proteins were chosen based on organism of interest, predicted subcellular localization (non-cytoplasmic), existence of immunogenic information, and availability of structure or of a homologue with known structure. To characterize medium-amplitude motions in these proteins 30 ns simulations were devised. Statistical relevance was taken care of by using two different force fields and set ups. In addition, for each system 5 simulations were performed starting from different velocity distributions. Simulations were run with 16 to 64 CPUs depending on system size and code. System sizes ranged from 22 000 to 221 000 atoms, including the protein, salt and solvent water. Over 2.5 TB of data have been generated, which will be next analyzed.

Understanding the structural factors behind antigen recognition requires the analysis of a broad set of proteins by different experimental and computational techniques. At the simulation level, this challenge can only be addressed if supercomputing resources such as those put together by DEISA are available.

Figure 1: Preliminary analysis of NMB0033 simulations. Root-mean-square fluctuation (RMSF) per residue, calculated for the full-length protein and for each domain separately. The large difference between the two RMSF profiles is indicative of stable domains that move relative to each other. Only domain B shows a relatively high internal motion, acting as hinge domain.

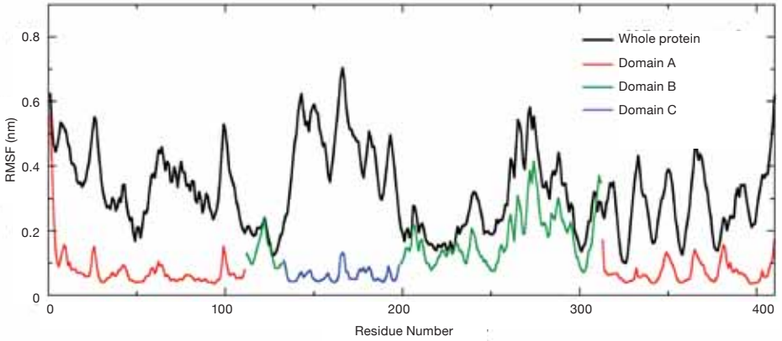

Table 1: Target proteins and simulations performed with DEISA resources.

Protein (gene name)	Organism	Immunity	Sequence length	AMBER 9	GROMACS 3.3.1
NMB0033	Neisseria meningitidis	protective	410	5×30ns (BSC)	5×30ns (BSC) 5×30ns (CINECA)
NMB0382	Neisseria meningitidis	protective	140	5×30ns (BSC)	5×30ns (CINECA, in progress)
NMB0419	Neisseria meningitidis	non-protective	193	5×30ns 5×30ns (BSC, in progress)	(CINECA)
NMB0663	Neisseria meningitidis	protective	155	5×30ns (BSC)	5×30ns (CINECA, in progress)
NMB1053	Neisseria meningitidis	protective	249	5×30ns (BSC)	5×30ns (CINECA, in progress)
NMB1870	Neisseria meningitidis	protective	124	5×30ns (BSC)	5×30ns (BSC)
NMB1946	Neisseria meningitidis	protective	247	5×30ns (BSC)	5×30ns (CINECA)
NMB1969	Neisseria meningitidis	protective	299	5×30ns (BSC)	5×30ns (CINECA, in progress)
NMB1985	Neisseria meningitidis	protective	888	—	5×30ns (CINECA)
CPn0800	Chlamydia pneumoniae	protective	428	5×30ns (BSC)	5×30ns (BSC)
CT043	Chlamydia trachomatis	non-protective	143	5×30ns (BSC, in progress)	5×30ns (CINECA)
SPy2010	Streptococcus pyogenes	protective	926	5×30ns (BSC)	5×30ns (BSC)

Arrigo Calzolari, Stefano Corni and Rosa Di Felice, INFM-CNR S3, Modena, Italy. Alessandra Catellani, CNR IMEM, Parma, Italy. Giancarlo Cicero, Politecnico di Torino, Italy. Francesco Iori and Elisa Molinari, University of Modena and Reggio Emilia, Italy.

Protein–surface interactions mediated by water

The molecular recognition capability of proteins can be specifically oriented toward inorganic surfaces. This project aims to elucidate the mechanisms of interaction between a surface and a protein, including the role played by water.

The dynamics of gold atoms are important for explaining the interaction with the protein and with water.

The interaction of a hydrophilic protein ß-sheet with a gold surface in water at room temperature has been simulated via ab initio molecular dynamics for the first time. The mechanisms of protein-surface interactions that were unknown before have been evidenced. The outcome of the project will help us in elucidating how proteins can selectively recognize inorganic surfaces, such as metals.

Protein–gold–water system

Recent combinatorial biotechnologies have shown that the molecular recognition capability of proteins can be specifically oriented toward inorganic surfaces. However, at present the principles regulating protein-surface interactions are poorly understood. What features of the surface and of the proteins determine which protein is able to bind to a given surface and how? In addition, the exact role of water in the process of protein-surface interaction is unknown. The understanding of such mechanisms would foster several technological applications based on the rational design of protein-surface interactions, ranging from biomaterials to nanobioelectronics. The aim of this project is to elucidate the mechanisms of interaction between a surface and a protein, including the role played by water. In particular, we have chosen to focus on the interaction of a polypeptide with one of the experimentally most studied surfaces, namely Au(111), about which very little is understood from the theoretical viewpoint, and even the wetting behaviour without the protein has not been studied yet via ab initio molecular dynamics (AIMD). As for the protein specimen, an anti-parallel ß-sheet structure has been used, which is a system that joins realism with computational convenience and which is the proposed structure for an experimentally identified gold-binding protein.

Ab initio molecular dynamics simulations

We have performed AIMD simulations at a finite temperature using plane-wave density functional theory on a system that is composed of an Au(111) slab and an anti-parallel ß-sheet of poly-Serine (a natural amino acid occurring in experimental gold-binding peptides), immersed in explicit water molecules and replicated by 3D periodic boundary conditions. The Car–Parrinello approach has been used in particular to calculate the time evolution of the system. The initial set up was performed with preliminary classical molecular dynamics simulations.

New mechanisms of protein–surface interactions found

The AIMD of the protein–gold–water system was carried out using the Mare Nostrum supercomputer at the Barcelona Supercomputing Center, with the ab initio molecular dynamics code included in the Quantum-ESPRESSO suite. The fine-tuning and optimization of the code for the Mare Nostrum was undertaken with the support of DEISA staff. The simulated system is composed of 4 layers of Au (112 atoms), 130 protein atoms and 115 water molecules, giving a total of 587 atoms and approximately 2500 electrons. We simulated the time evolution of the system using the Car–Parrinello method for 15 ps (plus 2 ps of initial thermalization), a duration long enough to obtain meaningful averaged quantities (also considering the large number of atoms in the system). Each ps of dynamics required ~100 h with 200 processors of the Mare Nostrum supercomputer.

The results of our simulations show the interaction mechanisms between the three components of the system, i.e., between water and gold, poly-Serine and gold, and poly-Serine and water. Especially the analogies and differences in the behaviour of water with respect to gold and the related side-chain of the Serine (hydroxyl group, OH) could be identified. Remarkably, some of the mechanisms identified were not known before, and required AIMD for a correct description. The dynamics of gold atoms are especially important for explaining the interaction with the protein and with water. In addition, our calculations demonstrate for the first time that such large systems, including metals, can be simulated using Car–Parrinello AIMD for a meaningful duration, which yield scientifically sound results when the extreme computational power provided by initiatives such as DEISA is available. The full analysis of the trajectories and of the electronic structure will eventually allow us to give the quantitative description of the phenomena.

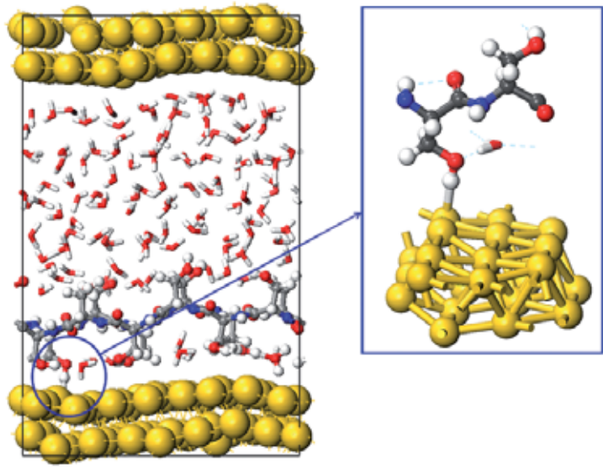

Figure 1: A snapshot from the ab initio molecular dynamics of the protein–gold–water system. The atoms represented in the left panel of the figure are repeated by 3D periodic boundary conditions, creating an infinite gold slab and an infinite poly-Serine ß-sheet. The right panel highlights the instantaneous interaction between the Serine side-chain and the gold surface, which takes place (in this example) via the hydrogen atom of the hydroxyl group. Similar interactions can be seen, in the other snapshot, between gold and water.

Pierre Bahurel, GIP Mercator Océan, France, Alvarez Fanjul, Establecimiento de un Sistema Español de Oceanografia Operationa (ESEOO), Spain

Global to regional oceanographic modelling

A French-Spanish research team has developed new oceanographic models that will help us to understand the evolution of the climate over the last decades and to provide an ocean analysis and forecasting service on a daily basis. A supercomputing infrastructure is essential for oceanographic modelling in order to generate ever more realistic simulations of the ocean's behaviour.

The Global to Regional Oceanographic Modelling (GROM) project was realised in collaboration with two European operational oceanographic projects, Mercator Océan in France and ESEOO in Spain. This project was a part of the research and development activities realised during 2006 for operational oceanography. The main objective was to build and validate new ocean model configurations from regional to global scales which are able to simulate eddy permitting to eddy resolving physical processes.

Global eddy resolving model

Several steps were defined during this project to reach the final goal of producing a global eddy resolving model, (1/12°), and a higher resolution model, (1/36°), which include high frequency physical processes. The regional model covers the North East Atlantic, in particular the European coast from North Africa to the south of Scandinavia. This objective has not yet been reached, but the main components of this future system have been built and the intermediate configurations are based on the global eddy permitting model, (1/4°), and the eddy resolving (1/12°) North Atlantic Ocean and North East Atlantic regional model.

The ocean model NEMO (Nucleus for a European Model of the Ocean) has been used during the GROM project to carry out several simulations useful for the development of these future systems, and for understanding the numerical representation of the physical processes of the ocean and sea ice.

Modelling the ocean requires high computation capacities

The DEISA Extreme Computing Initiative, which has supported the GROM project, enabled the realisation of the inter-annual simulation of the ocean at several horizontal resolutions, over several geographic domains and with a representation of different physical processes. The computational cost of these ocean simulations is very high. For example, one year of the North Atlantic and Mediterranean Sea model (NATL12) configuration used 8 200 CPU hours and one year of the global ORCA12 configuration used 100 000 CPU hours on the ECMWF IBM SP4. In total around 400 000 CPU hours of DEISA resources were used.

Results

The main results obtained in the simulations carried out during the GROM project were:

- An interannual (1993–2006) global eddy permitting simulation, which reproduces the main features of the thermo dynamic processes in the ocean (Kuroshio, Gulf Stream, Antarctic circumpolar current, Aghulas rings, El Niño events in the Pacific ocean, and monsoon seasonal cycle in the Indian ocean) and in the sea ice.
- A "realistic" simulation of the North Atlantic with a 1/12° horizontal resolution: The main currents of the North Atlantic (Gulf Stream, the North Atlantic current, Azores current, and Greenland current) were simulated, and these results can be compared to observations in terms of position, intensity, and energy.
- A regional model of the North East Atlantic embedded in a larger geographical model to the open boundaries conditions: The diagnosed eddy kinetic energy allowed us to show the positive impact of the high frequencies on refreshing boundaries conditions.

The final goal that remains is the development of a global eddy-resolving model with a resolution comparable to that realised in the regional models (1/12°). This model is now under development but it is still too large for the currently available computational facilities.

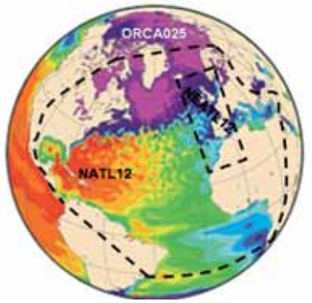

Figure 1: Three models were developed and evaluated during the GROM project. The global ORCA025 model simulates the global ocean and sea ice physics. It provides boundaries conditions for the North Atlantic and Mediterranean Sea model (NATL12) which provides boundaries conditions for the North East Atlantic model (NEATL12). During the GROM project an improved ocean configuration was in development, a global 1/12° model, which will replace ORCA025 and a very high resolution (1/36°) model covering the North East Atlantic to replace the NEATL12 model.

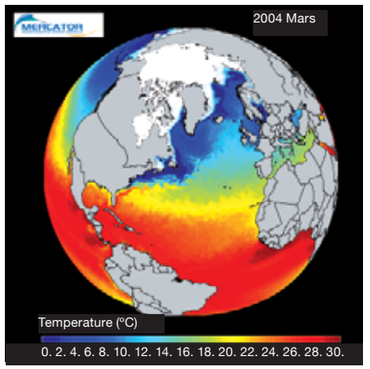

Figure 2: Temperature of the sea surface (coloured regions) and sea ice extension (white regions) in March 2004, as simulated by the global high resolution model developed by Mercator-Ocean. Note the trajectory of the Gulf Stream, one of the globe's largest western boundary currents, which draws warm water from the Caribbean Sea and the Gulf of Mexico and flows into the North East Atlantic.

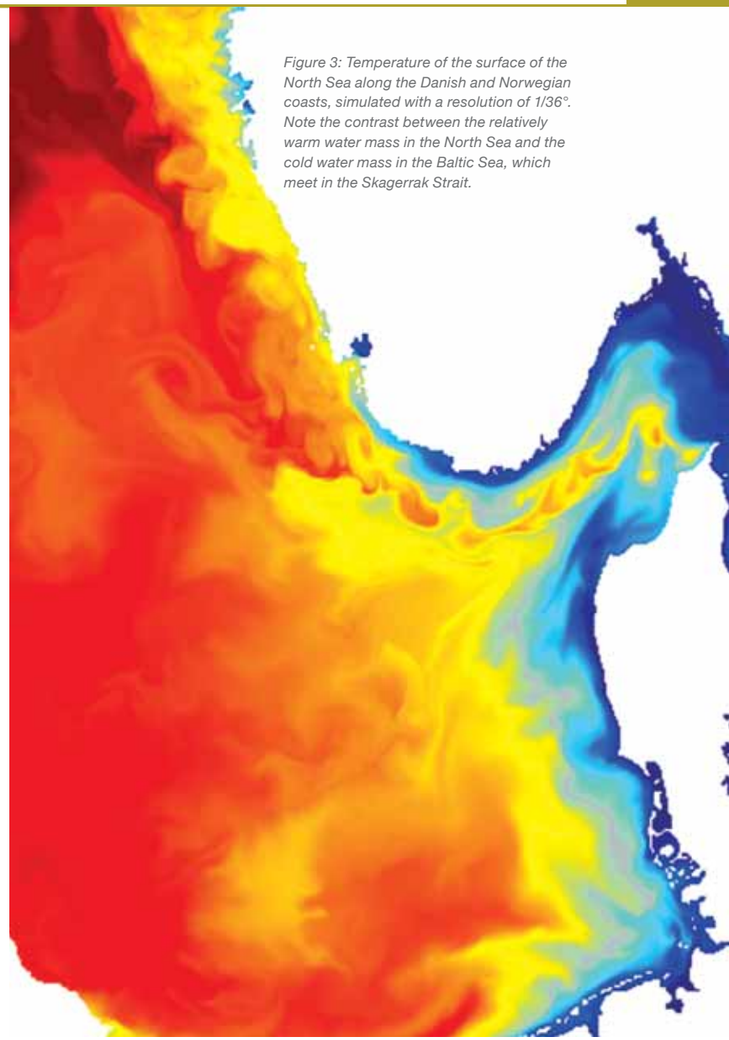

Figure 3: Temperature of the surface of the North Sea along the Danish and Norwegian coasts, simulated with a resolution of 1/36°. Note the contrast between the relatively warm water mass in the North Sea and the cold water mass in the Baltic Sea, which meet in the Skagerrak Strait.

Henk Dijkstra, Utrecht University, the Netherlands; Wilco Hazeleger, KNMI, the Netherlands

Statistics of climate variability

THE ENSEMBLE SIMULATIONS OF EXTREME WEATHER EVENTS UNDER NONLINEAR CLIMATE CHANGE (ESSENCE) PROJECT WAS SET UP, UNDER THE DECI, TO STUDY CLIMATE TRENDS.

The main aim of the ESSENCE project was to compute an adequate estimate of the statistics of internal climate variability and hence be able to obtain a good signal-to-noise ratio for the forced signal due to the increase of greenhouse gases. In the project, a 17-member ensemble simulation of climate change in response to the SRES-A1b scenario was carried out using the ECHAM5/MPI-OM climate model. The relatively large size of the ensemble indeed enabled us to better distinguish the forced signal from internal variability. We showed that in large parts of the world the observed warming over the last 60 years is statistically indistinguishable from the warming forced by increased greenhouse gas concentrations.

Earliest detection times were found

A great advantage of a large ensemble is the large noise reduction that can be achieved by averaging over all ensemble members. We were able to determine the year in which the forced signal (i.e., the trend) in the atmospheric (2 meter) temperature emerges from the noise. A student t-test, in which the trend over a particular period is compared with the standard deviation of the noise, was used.

The earliest detection times are found off the equator in the western parts of the tropical oceans, where the signal emerges as early as around 2000 (and for some regions even earlier) from the noise. In these regions the internal variability is extremely low while the trend is only modest. A second region with an early detection is the Arctic, in which the trend is very large due to the decrease of the sea-ice. The longest detection times are found along the equatorial Pacific where, due to El Nino, the variability is very high, as well as in the Southern Ocean and the North Atlantic, where the trend is very low.

Large amount of data

The ESSENCE simulations were successfully performed on the DEISA infrastructure, using the NEC SX-8 system of HLRS. Each run was performed on one node (8 cpus) of the SX-8; 6 cpus were needed for the atmosphere model, 1 for the ocean model, and one for the interface between the two. Up to 20 runs were done in parallel, although due to I/O limitations the bulk of the runs were done with only 8–10 in parallel. In total about 6 000 model years have been simulated.

The 50 TB dataset of the ensemble simulations, comprising of 151 years time series and nearly 200 variables, is a unique source of information for many researchers, and its analysis will provide clear answers on the changes in statistics of extreme weather events. This large amount of data obtained from an ensemble makes it possible to determine weather extremes and their possible change more accurately. In addition, data analysis is facilitated by using high-resolution tiled panel displays at SARA.

In large parts of the world the observed warming over the last 60 years is statistically indistinguishable from the warming forced by increased greenhouse gas concentrations.

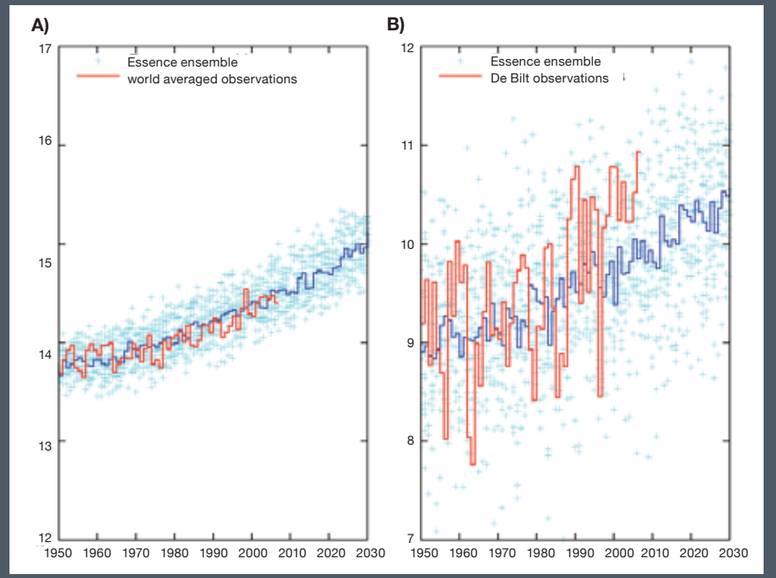

Figure1: Annual-mean surface temperature for the 17 ensemble members (light blue crosses), their mean (blue line) and observations (red line) for (a) the global average and (b) station De Bilt (the Netherlands).

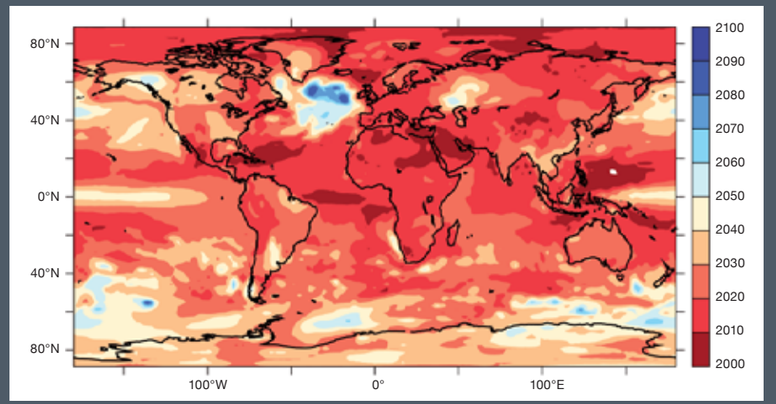

Figure 2: Year in which the trend (measured from 1980 onwards) of the annual-mean 2-meter temperature emerges from the weather noise at the 95%-significance level.

José M. Baldasano Recio and Pedro Jiménez, BSC-CNS, Barcelona, Spain

Global and regional interactions of air quality and climate

The environmental climate changes over the coming decades will include the exceedance of critical levels that will result in changes in weather patterns, extreme events, modifications in marine circulations, and high levels of atmospheric pollutants, etc. The GLORIA project deals with high-resolution three-dimensional global circulation models (GCMs) combined with regional climate models (RCMs) in order to analyze the impacts of climate change on air quality in the Mediterranean region. The project's main achievement is that it covers the integrated quantification of climate impacts on air quality on a regional scale, applied at a high-resolution scale to problematic areas, such as the Mediterranean.

The main aim of this project is to study the links between climate, the control of atmospheric pollution, the transboundary and intercontinental transport of air pollutants and their mitigation. The results of this project are expected to contribute to the analysis of climate and its relationships with diverse agents, especially air pollutants, and which may provide a basic tool for controlling future air pollution.

High-resolution simulations

The computational kernel of the simulations is concentrated on high-resolution simulations using different models. The three applications deployed require the use of large supercomputing facilities, mainly due to the size of the grid needed for the simulation and the amount of information managed. A long simulation period is required in order to assess climate issues. It is estimated that the requirements with an Intel Xeon Processor (3.2 GHz installed on a HP Workstation xw6000) are in the order 0.5 processor hours per simulated hour for the full production run, so a requirement of around 450 000 CPU hours is obtained.

Domain in study

The effects of climate change on air quality may affect long-term air quality planning. However, policies aimed at improving air quality have not taken into account variations in the climate. Furthermore, the coarse horizontal resolution of current global climate-chemistry simulations does not allow the effects of climate change on tropospheric photo-oxidant distributions to be estimated on a regional scale.

The work undertaken in GLORIA has so far been based on determining how concentrations of atmospheric pollutants (mainly focusing on ozone and particulate matter) respond to changes in the climate over the Mediterranean by using the regional modelling system WRF-CMAQ-DREAM installed in the MareNostrum supercomputer. The domain of the study covers an area of 4940 × 2640 km with a horizontal resolution of 20 km and a vertical resolution of 31 layers in the troposphere. Three simulations corresponding to past climate variations under summertime conditions (August of 1960, 1980, and 2000) have been performed and compared. In addition, two future scenarios corresponding to the year 2030, the SRES A1B and B1 IPCC scenarios following Unger et al. (2006), were investigated with meteorology corresponding to the control year. The results for the control year where evaluated against background station data from the EMEP network, which showed an accurate behaviour in the entire domain of the study.

Regional patterns revealed

The maximum 1-hour ozone concentration variation ranges from –20 to 70 μg m^{-3} in the months of the study, depending on the region of the domain. The most significant increases are achieved in northern Italy, while the largest decreases in the ozone mixing ratios are found over Greece and the Aegean Sea. The average summertime concentration shows a marked gradient from the central Mediterranean to the extremes of the domain, where no noticeable increases or even slight decreases are observed. The number of exceedances of the maximum 1-hour ozone threshold has progressively increased in the central Mediterranean in the months of August during the years 1960–1980–2000; on the other hand, the number of exceedances has decreased over the eastern coast and the Aegean Sea (Figure 1). A very similar pattern is observed for the summertime sulphate concentrations.

Figure 1: Number of exceedances of the ozone information threshold (180 μg/m^3) in the Mediterranean area with the WRF-EMEP-CMAQ modelling system: (up) difference year 2000–1960; and (down) difference year 2000–2030 (IPCC SRES A1B Scenario).

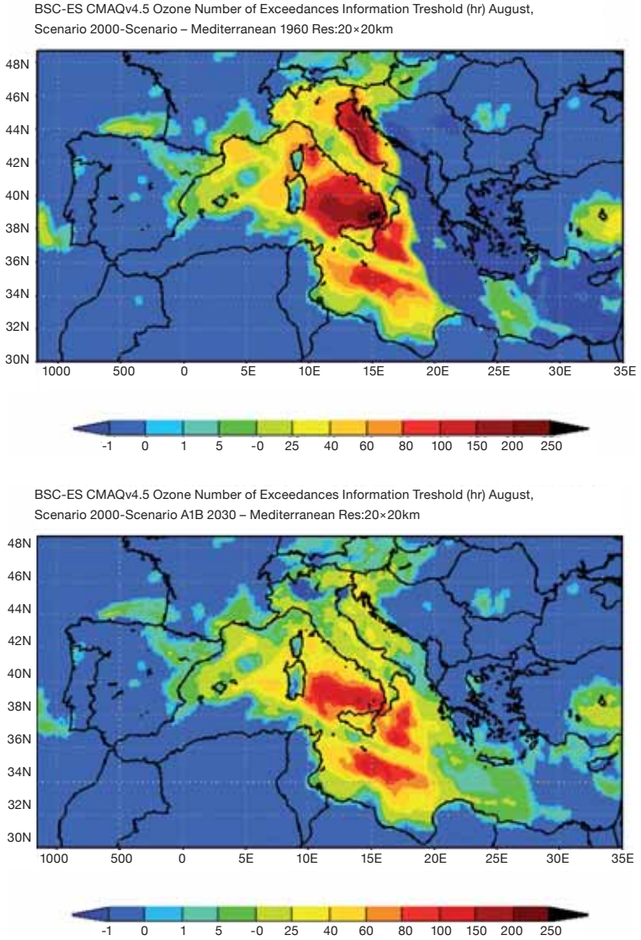

Lastly, how air quality reacts to the future evolution of emissions has been studied by using the SRES A1B and B1 scenarios for the year 2030. The 1-hour maximum concentrations in the Mediterranean decrease in the whole domain for the B1 scenario; and, as a consequence, the number of hours when the 1-hour threshold is exceeded substantially decreases in the entire Mediterranean but mostly in the central part of the domain.

Therefore, the complex topography of the Mediterranean causes pronounced regional patterns and differentiated behaviour in the Western, Central, and Eastern basins. The assumption of unchanged emissions leads to an increase in the mean concentrations of pollutants in most regions, which may be driven by an enhanced secondary production as a consequence of the temperature increase in the period of study and the higher pressures extending towards the eastern regions observed in the summertime of 1960, 1980, and 2000.

Martin Käser, Ludwig-Maximilians-Universität München, Germany
Michael Dumbser, Universität Stuttgart, Germany

Seismic wave propagation solutions for realistic 3D media

THE AIM OF THIS PROJECT IS TO PROVIDE REALISTIC SIMULATIONS OF EARTHQUAKE SCENARIOS BY RUNNING THE HIGHLY ACCURATE AND POWERFUL SIMULATION CODE **SEISSOL**. THE CODE IS ABLE TO INCORPORATE COMPLEX GEOLOGICAL MODELS AND TO ACCOUNT FOR A VARIETY OF GEOPHYSICAL PROCESSES THAT AFFECT THE SEISMIC WAVE PROPAGATION, SUCH AS STRONG MATERIAL HETEROGENEITIES, VISCOELASTIC ATTENUATION, AND ANISOTROPY. THE SYNTHETIC SEISMOGRAMS FOR THE MODELS, DESCRIBING REAL EARTHQUAKE SCENARIOS OR VOLCANO-TRIGGERED SEISMIC EVENTS, ARE USED IN A VARIETY OF APPLICATION FIELDS.

The objective of the SEISSOL project has been to perform highly accurate large-scale simulations of realistic wave propagation problems, such as real earthquakes or volcano-triggered seismic events. The models include a high degree of complexity involving heterogeneous material properties, real free-surface topography from digital elevation models and real source receiver locations. In this study, the sources consist of extended earthquake rupture models derived by inversion from real strong-ground motion data. On the one hand, the creation of highly accurate seismograms provides key datasets to investigate the physical properties of the earthquake rupture process. On the other hand, such synthetic datasets are used in volcano seismology to benchmark the moment tensor inversion routines and to determine their limits and applicability to real data in detail.

Performance improvements

As the underlying numerical method of SEISSOL is the Discontinuous Galerkin approach, it uses an explicit one-step time integration scheme to solve the governing partial differential equations. Therefore, a large amount of calculations are carried out on a rather small amount of data, i.e., the numerical algorithm is programmed in a cache-oriented manner. To achieve a performance improvement, the implementation of p-adaptation and local-time stepping on unstructured tetrahedral meshes has been investigated in detail. In particular, a more adequate mesh partitioning strategy, the use of optimised linear algebra library routines, and the use of optimised standard sparse matrix algorithms have been applied.

The simulations were performed on the HLRB 2 at the Leibniz-Rechenzentrum in Munich and typically used 256 to 512 processors. However, a few simulations were also carried out using 1024 processors, when the accuracy requirements for the results and therefore the number of elements or the approximation order was accordingly large.

Lancey Earthquake and volcano-triggered earthquakes

The project enabled us to create highly accurate data sets of synthetic seismograms for a variety of large-scale applications. Due to the incorporation of complex material properties, p-adaptation and local time stepping, an estimation of the peak ground motion in the alpine valley of Grenoble could be performed. It was possible to analyse the influence of the topographic effect of the surrounding alpine terrain in detail, and a map of absolute velocity values of the ground motion along the Belledonne fault after the Lancey Earthquake of 8th September 1995 was produced (see Figure 1). Such data provides invaluable information for civil engineers when constructing earthquake resistant structures.

Furthermore, a comprehensive accuracy study was carried out. The influence of the geometrical source complexity on extended kinematic earthquake rupture models was also studied in great detail.

In addition, synthetic seismograms for the Merapi volcano in Indonesia have been calculated for a high-precision topographical model (see Figure 2), allowing the effect on the seismic wave field to be determined. Different subsurface models of the volcano's interior structure have been assumed, which deviate from the homogeneous reference model. The distortion of the velocity structure in the modified models enabled the influence on the seismic signals recorded at the volcano's surface to be analysed. In this way, the performance of techniques to invert for the source location as well as for the seismic moment tensor that describes the volcano-triggered earthquake mechanism could be tested. This is an important issue, especially when searching for precursor signals for early warning systems that aim at the short-term forecast of future volcano eruptions.

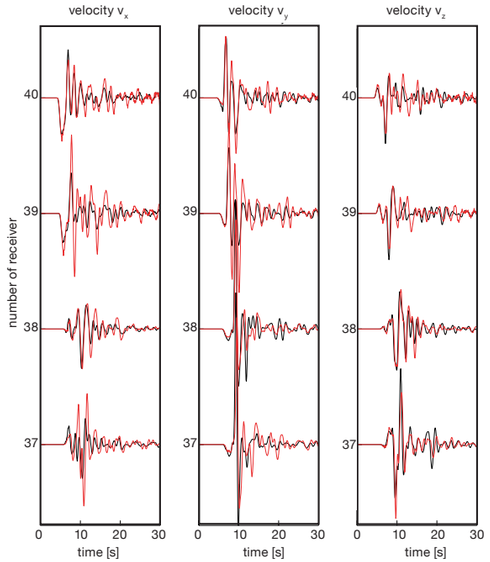

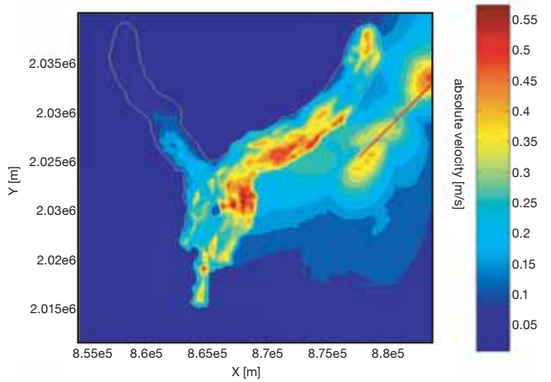

Figure 1: Comparison of seismograms calculated for different receivers in the valley of Grenoble (left). Seismograms for a model without topography are shown in black, whereas seismograms with stronger amplitudes due to topography are shown in red. A map of the absolute velocity of the ground motion for the area of Grenoble shows the endangered regions (right). The thick red line shows the surface projection of the Belledonne fault, the thin yellow line represents the extension of the sedimentary basin below Grenoble.

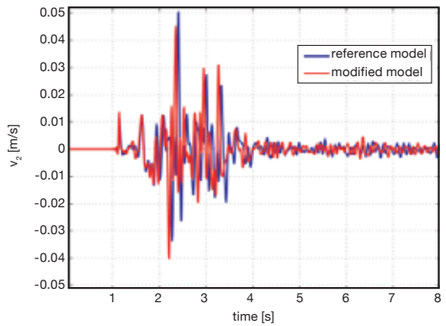

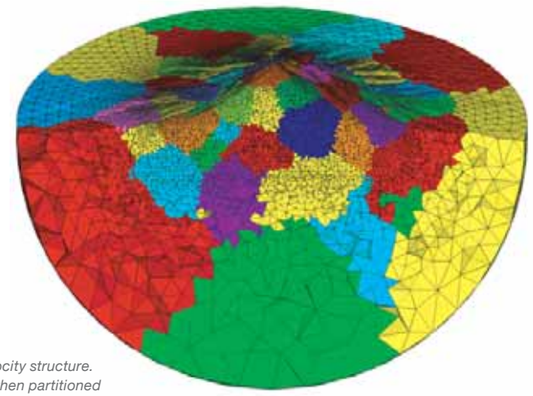

Figure 2: Synthetic seismograms calculated for the Merapi volcano in Indonesia (left) comparing the reference model with a modified model with a distorted velocity structure. The volcano is discretized by an unstructured tetrahedral mesh (right), which is then partitioned into subdomains to optimize parallel efficiency.

Computing wall turbulence at experimental Reynolds numbers

The main goal of this grand challenge application was to better understand the physics of turbulent boundary layers at high Reynolds numbers. This simulation is now the world's largest one in the field of wall turbulence, with 2×10^{10} grid points.

It is important to note the multiscale character of the intermediate layer present in turbulent flows. This is where the longest flow structures are found and where the range of scales is the widest. The purpose of this experiment was to extend the available simulations to a Reynolds number twice as high as has been previously studied. Wall-bounded turbulence has been studied for a long time, but it is only recently – in fact, only with this experiment – that we were able to represent the logarithmic layer linking the near-wall and outer flow regions. This is the location of the self-similar energy and momentum cascades, which are the essence of turbulent flows.

Based on previous simulations and experiments done by other researchers, we have some understanding of the nonlinear dynamics of the near-wall structures and the largest outer scales. Both are essentially single-scale phenomena. However, the dynamics of the intermediate range of scales are more poorly understood, because up to now it has been impossible to simulate a Reynolds number large enough for a multiscale range.

Application areas

The problem is not just of scientific interest. Wall-bounded turbulence is the interface between the ambient fluid, water or air, and moving vehicles, and between flows and pipes or channels. It is also a key component of the atmospheric boundary layer. The large scales play an important role, for example, in the dispersion of chemical agents in the atmosphere. Due to their associated low frequencies, their pressure fluctuations are responsible for undesirable aerodynamic loads which may result in structural fatigue and long-range noise in vehicles. More than half of the friction drag in vehicles resides in the near-wall and intermediate layers. The hopes of controlling and alleviating these phenomena reside with the understanding of their dynamics.

Statistically converged properties of the flow

The program we developed performs a turbulence simulation using the Navier–Stokes equations for an incompressible fluid in a plane channel between two infinite parallel plates. No modelling is used. The computation was carried out in a doubly periodic domain in the two wall-parallel directions, which was chosen large enough to minimize artefacts due to the spurious periodicity.

The main objective of the simulation was to run for an enough long time to get statistically converged properties of the flow. From our experience with lower Reynolds numbers we chose a domain of size 25 h in the stream-wise direction (x) and 10 h in the span-wise direction (z), where h is the channel half-thicknesses. This procedure – although both memory and compute intensive – allowed us to compute any property of the flow with greater accuracy than would be obtainable using experimental measurements in laboratory flows. In particular, it allowed us to access variables, such as time and space resolved velocity gradients and vorticity, which cannot be measured in the laboratory. The Reynolds number used in the simulation is comparable to those used in many laboratory flows and higher than those used in most of the wall-bounded turbulent flows that have been measured in any detail beyond mean and fluctuating velocities.

Very costly initiative

The program code was based on previous versions developed by our group during the past fifteen years. It broadly follows the standard spectral code developed at the end of the 1980s by groups at Göttingen and at NASA Ames, although the wall-normal discretization uses compact finite differences. It was written in C and Fortran, and was parallelized using MPI, with an excellent speed-up.

In order to run this code in a supercomputer with more than 2048 processes with optimal performance, we did some specific optimization, such as tuning the code to take an advantage of the vectorial instruction set of the PPC970-FX CPU. This was achieved using a version of the FFTW library provided by the Barcelona Supercomputing Center's Deep Computing research group. In addition, the most expensive I/O operations were assigned to dedicated processors to reduce the global file system contention. The latter approach allowed us to overlap the computing stage of 2048 processors with the massive writing operations, using only 60 extra processors. The aim was that the code was be able to follow the normal execution while the dedicated processors saved to the disk a huge amount of information produced by the code – up to 25 TB for the entire simulation.

This was a very costly initiative. The code ran for about 6 000 000 CPU hours, 800 000 of them were assigned trough the DEISA Extreme Computing Initiative. This simulation was performed on the Barcelona Supercomputing Center's supercomputer MareNostrum and is now the world's largest simulation in the field of wall turbulence, with 2×10^{10} grid points.

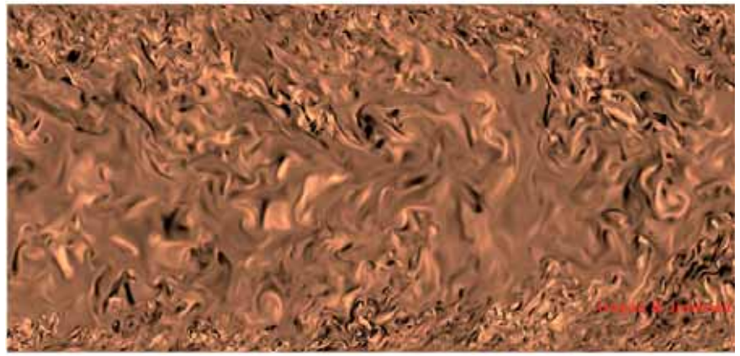

Figure 1: Spanwise component of vorticity.

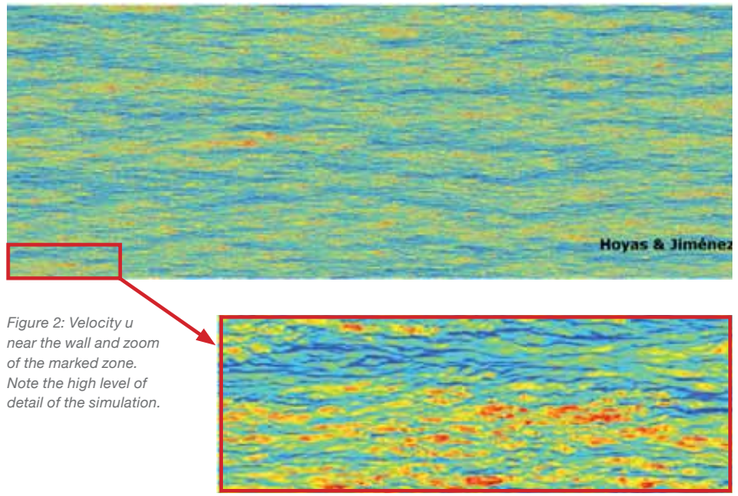

Figure 2: Velocity u near the wall and zoom of the marked zone. Note the high level of detail of the simulation.

Olivier Gicquel, Ecole Centrale Paris – CNRS, France

Numerical simulations of turbulent reacting flows

Taking into account radiative heat transfer in reactive large eddy simulations greatly modifies the flame dynamics, the temperature and fraction species fields, in spite of the limited impact on total heat release transfer. New large 3D calculations have been performed, proving the importance and viability of this kind of coupled simulation.

Huge 3D simulations were considered in order to accurately investigate the impact of radiation on the flame dynamics. The new 3D numerical models included a better description of the turbulent motions and the radiation properties of the combustion chamber walls compared with 2D models. These results will be helpful for gas turbine or furnace burner designers so they can avoid the occurrence of combustion instabilities in the combustion chamber.

Coupling combustion and radiation

Numerical simulations of turbulent reacting flows, including pollutant formation and radiative heat transfers, are mandatory for practical applications but require adapted models and large computational resources. Large eddy simulations (LES), where larger turbulent motions of the flow field are explicitly resolved when only the effects of small ones are modelled, appear to be a promising tool for such simulations. The physical phenomena involved in combustion and radiative heat transfers are also very different. Flow fields are generally described through balance over small volumes (finite volume context), whereas radiative heat transfers involve long distance

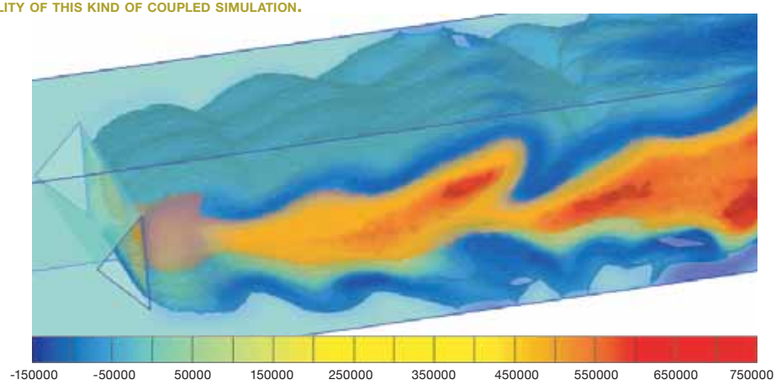

Figure 1: Instantaneous field of the radiative power (in kWm^{-3}). Positive (respectively negative) radiative powers are lost (absorbed) by gases.

interactions. Accordingly, reacting flow and radiative heat transfer codes have a very different structure. The proposed approach, realized in collaboration with the DEISA Coupled Application JRA, benefits from an efficient coupling between an LES solver and codes specifically for radiative heat transfers in which data exchanges occur at time intervals controlled by the physical times of each phenomenon.

Simulations of turbulent propane-air mixture

In the following, we present the results of a turbulent premixed propane-air flame simulations stabilized in the downstream of a triangular-shaped obstacle (see Figure 1). This configuration, called the diedra case, makes it possible to compare the 3D simulations with the 2D ones previously performed and also with a large number of experimental results already available at the EM2C laboratory (Laboratoire d'Energetique Moléculaire et Mascroscopique, Combustion – Ecole Centrale de Paris).

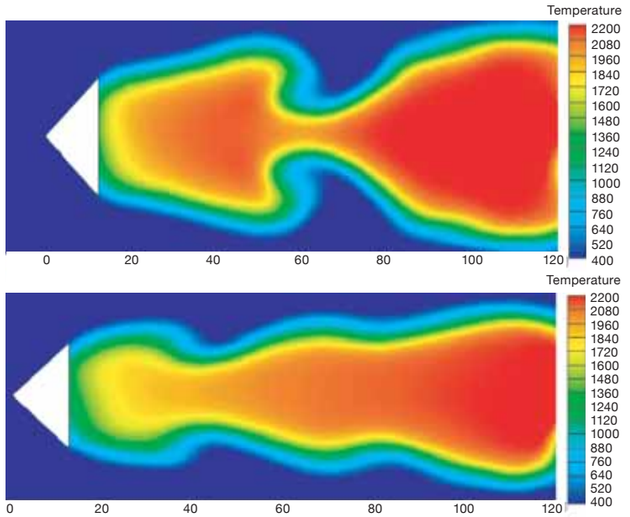

Figure 2: Instantaneous resolved temperature field in the central plane of the burner (z = 0) without (left) and with (right) radiative heat transfer at the physical time t = 0.55 s.

The modifications to the temperature field in the recirculation zone slightly downstream of the flame holder induce major changes in the flame behaviour and to its interaction with turbulence motions.

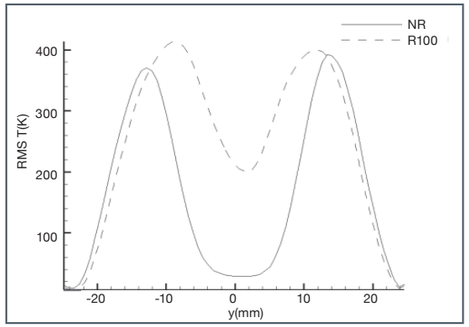

Figure 3: Profiles of the filtered temperature RMS in the central plane of the burner, (z=0), x = 0.05 m downstream of the flame holder without (NR) and with (R100) radiative heat transfer.

First, the maximum mean temperature is decreased by about 100 K because of the increase in energy transfers from hot to cold zones and to the walls due to radiation. These temperature differences directly impact the reaction products and in particular the formation of pollutants, such as carbon monoxide (CO) and nitric oxides (NO_x), which are known to be very sensitive to local temperatures.

Furthermore, the modifications to the temperature field in the recirculation zone slightly downstream of the flame holder are known to play a key role in the flame stabilization mechanism, and to induce major changes in the flame behaviour and its interaction with turbulence motions (see Figure 2).

On a more quantitative level, the filtered temperature RMS (root-mean-square) evidences the modification of the flame dynamics when radiative heat transfer is taken into account (see Figure 3). The RMS values are clearly larger when the radiative heat transfer is taken in account, denoting stronger flame movements. These findings confirm the strong influence of radiative heat transfer, in spite of its limited contribution to the energy transfer (less than 2% of the overall heat).

Radiative heat transfer affects flame dynamics

Radiative heat transfer plays an important role in turbulent combustion, but is often neglected in simulations because of its complexity and the related numerical cost. An original approach has been used to perform numerical simulations of unsteady 3D turbulent combustion which includes radiation.

Results from the three-dimensional simulation of the diedra case show that taking into account radiative heat transfer in reactive large eddy simulations greatly modifies the flame dynamics, the temperature and fraction species fields, in spite of the limited impact on the total heat release transfer.

These new 3D calculations with a more realistic description than the 2D model have been performed to prove the importance and viability of this kind of coupled simulation.

Heavy particles in turbulent flows

A research group at the National Research Council (CNR) in Italy has studied statistical properties of small droplets transported by an incompressible homogeneous and isotropic turbulent flow. Understanding the evolution of impurities and micro-droplets in a turbulent environment is of great interest in a variety of applications ranging from health preservation to engineering and atmospheric sciences.

We produced a rich database of Lagrangian trajectories useful in studying the statistical properties of small particles or droplets transported by an incompressible homogeneous and isotropic turbulent flow. Billions of particles have been evolved by means of a state-of-the-art pseudo-spectral fully-dealiased Direct Numerical Simulation (DNS) of the Navier–Stokes equations.

Our database will be used both as a reference for laboratory experiments and for testing models in applications. Data produced will constitute an outstanding database to systematically study the effects of inertia on the preferential concentration and dynamics, with high resolution on events like vortex trapping. After an initial period of data analysis, we plan to share the raw data obtained from the simulations with the entire international scientific community.

Understanding the dynamics of droplets transported by a turbulent flow

The evolution of suspended particles and micro-droplets in a turbulent environment is an issue of great interest for a variety of applications ranging from health preservation to engineering and atmospheric sciences. In the latter context, micro-droplets' growth by condensation or evaporation is a phenomenon of paramount importance for the early stages of cloud evolution. It is believed to be a crucial factor in determining the cloud reflectivity and thus the intensity of the greenhouse effect.

From the physicist's point of view, understanding the dynamics and the statistics of small particles or droplets transported by an incompressible turbulent flow is a great challenge. This problem is definitely more complicated than the motion of a fluid tracer (i.e., a small volume element of the fluid itself), due to both the size and the density of the particles, which is different from that of the carrier fluid. The experimental investigation of particle transport in turbulence has lagged behind due to technological difficulties; in this context the numerical approach represents an extremely valuable tool, especially as massively parallel computing permits parameter values (i.e., Reynolds number) approaching those of experimentally accessible ones.

Computational Approach

We used a modified version of the numerical codes employed in the past for 10 243 numerical simulations. The same code has also been used at smaller resolution to study inertial particles, whose modelization is a further simplification with respect to Maxey and Riley (Physics of Fluids, 1983), with only the Stokes drag term retained. The present project, together with previous runs at lower resolutions, will allow us to perform a systematic study of varying the value of the Reynolds number, Re, and of the Stokes term, i.e., the two most relevant parameters.

Very few groups have performed DNS of the Eulerian dynamics of a forced turbulent flow at a resolution of 20 483, and no one has previously studied the motion of particles in such a flow. This lack is due to the huge computing power needed for these simulations and the expertise needed to plan and manage all the different theoretical and technical issues related to this kind of project.

Database available for the entire scientific community

We numerically integrated the evolution of 2 139 648 000 particles divided into 21 different classes characterized by different Stokes numbers. We were interested in the space localization in correlation with the Eulerian flow properties of these particles and their time evolution (e.g., velocity autocorrelation along trajectories). Given the huge number of particles a frequent temporal sampling for the totality of particles was impossible. Hence we opted for a reasonable compromise: we stored the full configurations for a few particles (synchronised with the full Eulerian field) and we stored the temporal evolution of a subset of particles.

Figure 1: Particles with different inertia, as they evolve after being released from within a small scale vortex filament. A Lagrangian tracer (red) is strongly sensitive to small scale vortices from the advecting turbulent flow. Particles with higher and higher inertia (green, blue, yellow, etc) are less and less sensitive to small scale vorticity. For a relatively small Stokes value 0.16 (green trajectory) the effect of the turbulent vortex is already barely felt by the particle.

In this way we could keep the data output from the simulation within 7 TB. The simulation was performed on a SGI Altix 4700 named HLRB II at LRZ in Garching, Germany.

The numerical code initially developed by Federico Toschi in 2000 is a pseudo-spectral, dealiased by means of the 2/3 rule, fully parallelized (MPI and FFTW 2.1.5) code capable of evolving particles by means of an out-of-grid tri-linear interpolation. The total run duration took 60 000 time steps for a total of 400 000 CPU hours. Each time step (at production) took on average 41 seconds of which 21 were spent calculating Fourier Transforms. The total waste of CPU-time (due to failed runs, etc.) was less than 20 per cent, which is a good figure considering the fact that the machine, at that time, was in its earlier productive runs. The full run was split into 155 5-hour-long batches.

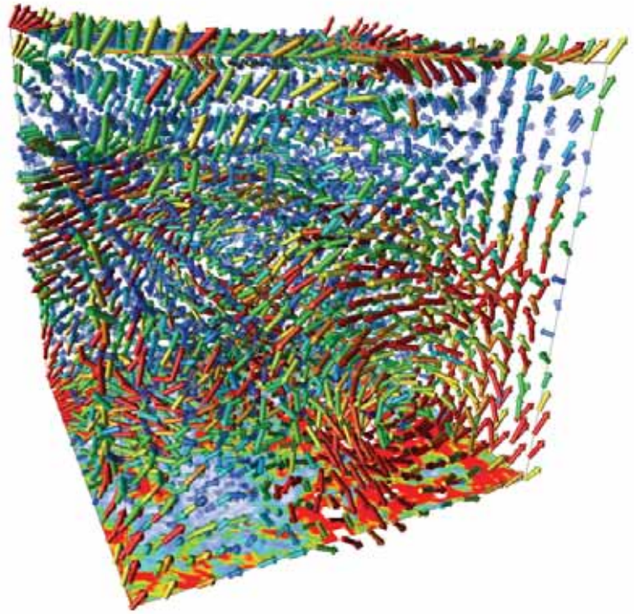

The raw data produced during the study has been divided into Eulerian configurations of the velocity field (20 483 configurations). A dump of all the particles' configurations (2G particles) was made with all information of their position, velocity, fluid velocity at each particles' position, and the fluid velocity gradients at the particle position and its trajectories. All this data has been converted into the HDF5 format. The database produced will soon be made available on the web site of the international CFD database, i.e., the iCFDdatabase at http://cfd.cineca.it.

Figure 2: A snapshot of an Eulerian turbulent velocity field from the Direct Numerical Simulation of the Navier–Stokes equation. As it can be seen, the turbulent field is characterized by the presence of vortex and vorticity at all scales. In the figure we can see a large vortical structure. Particles are evolved in such velocity fields.

Jan Schulze, Christoph Schaupp, and Jörn Sesterhenn, Technische Universität München, Germany

Supersonic jet and its acoustic field simulation

A researcher team at Technische Universität München has used direct numerical simulation to study the physical processes responsible for producing shock-induced noise generated by supersonic jets. The main challenge was to simultaneously resolve the small-scale nonlinear turbulent structures and at the same time the large-scale, small amplitude acoustic waves they produce.

Shock-induced noise is generated by supersonic jets which are not perfectly expanded. This means that their nozzle exit pressure is above or below the ambient pressure. The flow adapts to the ambient pressure by a series of oblique shocks, compression and expansion waves. These interact with the shear layers, producing shock-induced or shock-associated noise. To this end direct numerical simulation was performed to study the physical processes responsible, many of which have not yet been clearly identified. The main challenge was to simultaneously resolve the small-scale nonlinear turbulent structures and at the same time the large-scale, small amplitude acoustic waves they produce.

Numerical methods of high order of accuracy

In this project, the method of direct numerical simulation was used to compute a three-dimensional supersonic rectangular jet that is not perfectly expanded, as it is found at the nozzle exit of jet engines for aircraft. Numerical methods of high order of accuracy were chosen for the direct solution of the compressible Navier–Stokes equations. This gave us the possibility to compute the sound field that was generated by the supersonic jets, directly.

Shock-induced noise plays an important role in jet screech

Consider a supersonic jet, e.g., at the exit of a jet engine, in the over- or under-expanded case. A regular pattern of compression and expansion waves will be found within the supersonic part of the jet flow. A compression wave incident on the sonic line will be reflected as an expansion wave, and vice versa. At the location of interaction between the compression wave and the turbulent mixing layer, acoustic waves are generated. This shock-induced noise also plays an important role in what is called jet screech. This phenomenon manifests itself by a strong and sharp peak in the sound pressure level spectrum, corresponding to a tonal noise at high amplitude (up to 160 dB), which may cause damage to parts that are near the nozzle due to the high dynamic loads. Screech is induced by shock-induced acoustic waves travelling upstream and forcing the “young” shear-layer at the nozzle exit. At this point Kelvin–Helmholtz instabilities are growing to vortices, transported downstream and interacting with the shock tips which are emanating noise again and closing a feedback loop.

Code parallelized using MPI

The compressible Navier–Stokes equations were solved, based on a characteristic-type formulation on an orthogonal grid (approx. 300 million grid points) stretched in both the stream-wise and the transverse directions. Along the span-wise direction, periodicity and statistical homogeneity were assumed.

To capture the sound generation and propagation processes, spatial discretization was done using a finite difference compact scheme of sixth order and a spectral like method in the periodic direction. The code is parallelized using the Massage Passing Interface (MPI). For the current setup 1020 CPU’s were used on a SGI-Altix 4700 with Itanium2 Madison 9M processors. Approximately 12 TB of data were written to disk and 0.5 TB of main memory were used.

Results

In Figure 1 an isosurface of the vorticity (|rot(u)|) is shown in a three-dimensional sketch with a plane of the dilertation field (div(u)). A dominant noise source seems to be close to the nozzle exit (between the first and second shock-cell). The absolute velocity is presented in Figure 2 which visualizes the persisting shock-cell structure in the jet core.

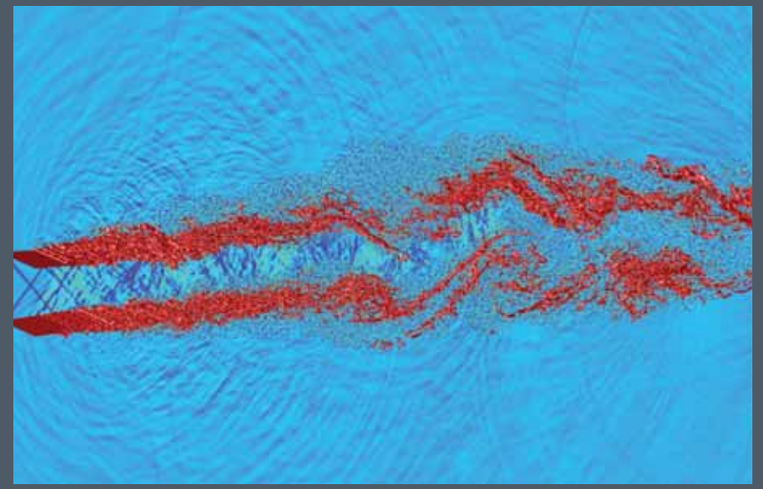

Figure 1: Vorticity (isosurface, red) and dilertation (plane, blue)

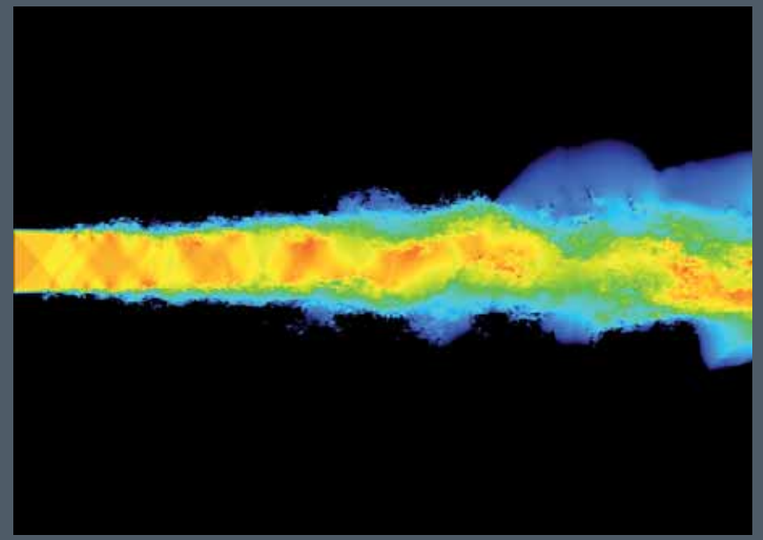

Figure 2: Absolute velocity (|u|)

Jean-Philippe Laval and Matthieu Marquillie, Laboratoire de Mécanique de Lille, France. Christian Stemmer, Technische Universität Munchen, Germany. Paolo Orlandi, Università di Roma La Sapienza, Italy. Lars Davidson and William K. Georges, Chalmers University of Technology, Gothenburg, Sweden. Philip Hancock, University of Surrey, UK.

Direct numerical simulation of flow over a bump

Over the past half a century, a great deal has been learned about boundary-layer turbulence; however, all the existing models of the near-wall region are still empirical. The current project is linked to the objective of generating a database of high-quality DNS and conducting challenging experiments to improve the knowledge on and the modelling of near-wall turbulence.

A direct numerical simulation (DNS) of flow with an adverse pressure gradient was performed with a high Reynolds number, thanks to the large-scale supercomputing facilities available through the DEISA initiative. This type of simulation creates new opportunities for investigating three-dimensional wall turbulent structures. A more detailed characterization of these structures will be used to evaluate and improve turbulent models.

Understanding near-wall turbulence

Over the past half a century, a great deal has been learned about boundary-layer turbulence, both from the statistical and the structural points of view. Using the statistical approach, a wide range of boundary layers submitted to various favourable and adverse pressure gradients have been characterized experimentally. Nevertheless, all the existing models of the near-wall region are still empirical, and even the most sophisticated are not very successful in adverse pressure gradient situations. The current project is linked to the objectives of the European Project WALLTURB, which aims to generate a database of high-quality DNS and to conduct challenging experiments to improve the knowledge on and the modelling of near-wall turbulence. A DNS at a large Reynolds number can elicit information on 3D turbulent structures which is not directly available from the experiment.

Numerical simulations

The numerical code was used for a DNS of a similar flow at a lower Reynolds number. The incompressible Navier–Stokes equations are not solved in the curvilinear coordinate, but the partial differential equations are transformed using mapping. This leads to a modified system of equations which can be solved efficiently with smooth two-dimensional mapping. The code is parallelized in the spanwise direction using MPI routines. For space discretization, fourth-order central finite differences are used for the second derivatives in the streamwise direction. All the first derivatives of the flow quantities appear explicitly in the time-advancing scheme and the first derivatives in the streamwise direction are discretized using eighth-order finite differences. Chebyshev-collocation is used in the wall-normal direction. The transverse direction is assumed to be periodic and is discretized using a spectral Fourier expansion, the nonlinear coupling terms being computed using the conventional de-aliasing technique.

Large amount of data recorded for future analysis

In this project, a single DNS of channel flow over a smooth profile was performed at a Reynolds number $Re_\tau = 600$ at the inlet of the simulation domain. The inlet conditions were generated from a precursor DNS of flat channel flow. The smooth profile at the lower wall was duplicated from a similar experiment at the Laboratoire de Mécanique de Lille.

The computations were performed at HLRS (Germany) with 8 nodes (64 processors) of NEC-SX8 and required a total of 380 GB of memory for a physical space resolution of 2304×384×576. With these parameters, the averaged performances of the code reached 9 GFlops per processor. In order to adapt the mesh to the dissipation scales, the grid was densified in the downstream part of the bump. After reaching a stationary state, a total of 5 TB of raw data were recorded in order to compute high order statistics of the turbulence along both the flat wall and the curved wall. The time evolution of coherent structures, such as vortices, was also recorded and will be analyzed in the separation region.

A DNS at a large Reynolds number can elicit information about 3D turbulent structures which is not directly available from the experiment.

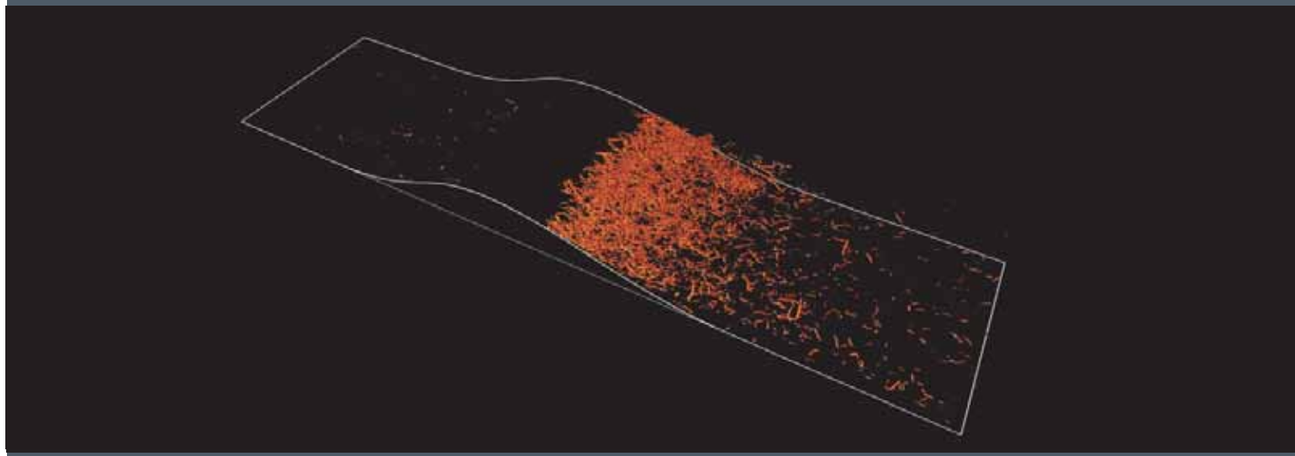

Figure 1: Positive isocontours of the Q-criterion ($Q = ||\Omega||^2 - ||S||^2$, where Ω and S are respectively the vorticity tensor and the strain rate tensor) in the detection of coherent vortices on the lower wall of the simulation. Intense vortices are generated near the summit of the profile near the separation region of the turbulent boundary layer. The separation mechanisms will be investigated with respect to turbulent statistics.

Michele Parrinello, ETH Zurich, Switzerland in partnership with Chemical Research Centre of the Hungarian Academy of Sciences (Institute of Chemistry), Hungary

Catalysis by *ab initio* metadynamics in parallel

There is an ongoing discussion concerning the processes that have contributed significantly to the formation of life on early Earth. Several possible scenarios and theories have been suggested. Their proof, however, often requires performing experiments under extreme conditions, where many complex reactions are simultaneously taking place. These are conditions that exist under the deep-sea, in volcanoes or at the interfaces of extremely hot, pressurized water and iron-sulfur minerals.

This is where high-performance computing and modelling come in: since the laws of physics are the same now as they were on Earth four billion years ago, it is possible to study the early conditions by simulating different scenarios and testing hypothesis with computer modeling.

Ammonia for prebiotic life

When considering the formation of a reducing environment favorable for the production of the organic precursors to life, an essential step is the formation of ammonia. Particularly pyrite (FeS_2) has been suggested to act as the key catalyst in the synthesis of prebiotic molecules in the early history of Earth, since there is experimental evidence that pyrite exhibits high catalytic activity under hydrothermal conditions similar to those present in the prebiotic world, i.e., under the deep-sea, near volcanoes, and at the interfaces between extremely hot, pressurized water and iron-sulfur minerals. Experimental evidence shows, that pyrite catalysis could have provided a significant amount of the required ammonia to the early life forms in the prebiotic world.

Focus on pyrite surface chemistry

The goal of the project was to achieve a better understanding of the pyrite surface and its chemistry by means of *ab initio* calculations. We modeled a defective pyrite surface by a periodic orthorhombic supercell containing a pyrite slab made of 18 atomic layers, each containing eight iron or eight sulfur surface atoms. The calculations were performed with the efficiently parallelized CPMD program package.

We created a surface sulfur vacancy by removing an originally three-coordinated surface sulfur atom, which was subsequently bonded to a surface iron atom. This results in two defects, which could be also obtained by a heterolytic breaking of the surface sulfur dimers as it would occur during pyrite fracturing. For such a defective surface structure, also open shell spin states have to be considered as the possible ground state and consequently

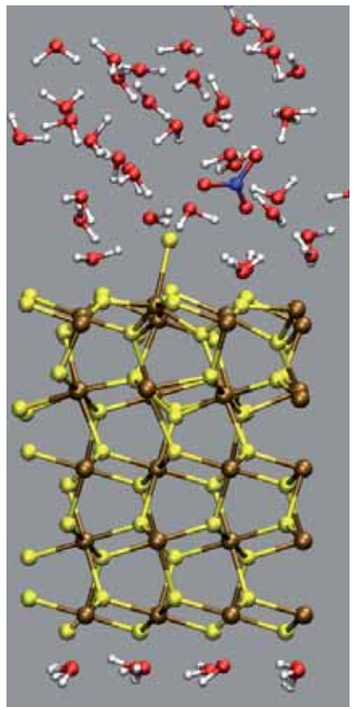

Figure 1: The results of the all-electron and pseudopotential calculations for bulk pyrite using different Gaussian basis sets.

a series of possible multiplicities was investigated. Finally, a spin triplet state turned out to be the ground state for this defect system. This state was also stable against the wetting of the surface. A triplet state implies the presence of two unpaired electrons. Such biradicals are known to be very reactive, i.e., they are able to lower the activation barriers of chemical reactions significantly. In this way, the presence of such radical species presumably facilitates the redox reactions needed for the reduction of nitrate anions to ammonia. This would explain the high concentration of ammonia found in the experiments and thus the availability of ammonia to early life forms in the prebiotic world under similar conditions.

Moreover, the formation of the sulfur defect site was accompanied by a change in the oxidation state of the iron from Fe(II) to Fe(IV). This observation is confirmed by a good agreement between the calculated sulfur core level shifts of different sulfur species and the experimental photoemission spectra.

Surface calculations from two simulations programs

Parallel to the investigation of the defective pyrite surface using Car–Parrinello molecular dynamics simulations, we started the investigation of the same surface employing the QUICKSTEP module of the CP2K program package. QUICKSTEP can perform all-electron calculations, which means that all-electron calculations for bulk pyrite model structures of different sizes could be performed as a reference for the pseudopotential calculations employing large and, thus, accurate all-electron basis sets, which were optimised for the use with QUICKSTEP.

The accuracy of the pseudopotential calculations employing relatively small Gaussian basis sets are excellent compared to the all-electron reference calculations and the experimental data. Encouraged by these results, we started the simulation of the nitrate reduction using an enlarged model system compared to the Car–Parrinello molecular dynamics simulations. However, these simulations are not yet finished.

Results of the project

The short time scale of ab initio simulations is often a serious problem. Therefore a new method was developed, which combines Car–Parrinello and Born–Oppenheimer molecular dynamics in order to accelerate the density functional theory based *ab initio* simulations. This means that simulations as long as tens or even hundreds of picoseconds can be routinely performed, thus making completely new phenomena accessible to *ab initio* simulations.

The computing resources provided by

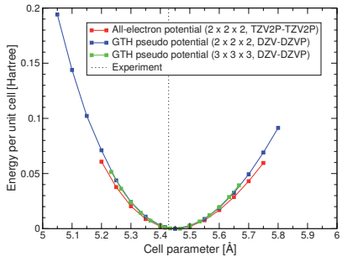

Figure 2: Simulation of the nitrate reduction using an enlarged model system.

DEISA were used for the evaluation of the new method. We could demonstrate that the dynamics are correctly reproduced and that high accuracy can be maintained throughout for systems ranging from insulators to semiconductors and even to metals in condensed phases. This development considerably extends the scope of *ab initio* simulations. The new method has been implemented into the QUICKSTEP code.

Finally, a web application plug-in for the CP2K program package has been implemented in collaboration with the RZ Garching, which will facilitate the use of CP2K within the DEISA computing infrastructure.

Claudio Zannoni, Universita' di Bologna, Italy

On the molecular working of a twisted nematic cell

A research group lead by Claudio Zannoni has studied for the first time, the switching of a pixel in a liquid crystal display at a molecular level resolution using a very large-scale Monte Carlo computer simulation of an off-lattice model. In particular, the group has concentrated on investigating yet poorly understood fundamental features of the twisted nematic cell.

One of the most successful stories in advanced materials is that of liquid crystal (LC) displays. The basic concept underlying the most classic among these devices, the twisted nematic (TN) display, is that a pixel is activated by a change of molecular organization in a few micron thick cell. According to this concept, an initial configuration of the local preferred direction (the director) is established between two orienting surfaces (rubbed glass or polyimide), rotated 90 degrees from one another, that confine the LC. An experimental fact is that polarized light is going through the pixel in this "rest" state and this is compatible with a microscopic helical configuration. If the chosen LC has a positive dielectric anisotropy and a suitable voltage is applied across the cell in correspondence of the pixel, then polarized light is not rotated and light does not go through the device, compatibly with a monodomain organization. When the field is switched off the original organization is re-established.

Huge number of degrees of freedom to be considered

The classic textbook picture which explains how a TN-LCD works is that of molecules in uniformly twisted layers. To the best of our knowledge there is little evidence that the molecular organization at rest is a uniform helix at molecular level. Moreover, the way the organization is established is not obvious. For instance, as the organization at rest is re-established after an aligning cycle, does the reorganization take place from the centre of the cell or from the surface? Uniformly or not? Is a uniform helix really formed? Or how helical is the structure?

Attempting to explain the microscopic working of the TN display is a particularly challenging problem. It necessitates considering a huge number of degrees of freedom, which need to be followed over a very long time period (from a microscopic point of view), as the average TN–LCD response times vary in the range of 8–15 milliseconds, while the typical time scales we are able to access to date through conventional computer simulations is of the order of nanoseconds.

Molecular resolution model of a twisted nematic cell

Within the DEISA project, we have tackled this problem by setting up a molecular resolution model of a TN cell containing O (106) model particles, simulated using the Monte Carlo (MC) method. For this purpose, we developed an MPI parallel MC code using a replicated data scheme, by modifying the canonical Markov chain of the Metropolis algorithm to which allows for multiple simultaneous moves to be performed at the same time by different processors. These kinds of moves are not possible in conventional MC algorithms due to their intrinsic non deterministic nature, but in our case, the big sample size allows every processor to pick an energetically independent particle in suitably chosen cells in which the whole sample is subdivided.

We modeled the LC rod-like molecules contained in the LCD cell as Gay–Berne single ellipsoidal interaction sites, discarding all the intramolecular degrees of freedom. Our TN cell is arranged to simulate a 0.1 μm × 0.1 μm × 0.05 μm display cell. Although in a fully realistic model one would have to consider an order of magnitude larger sample, this is currently unfeasible with the available resources.

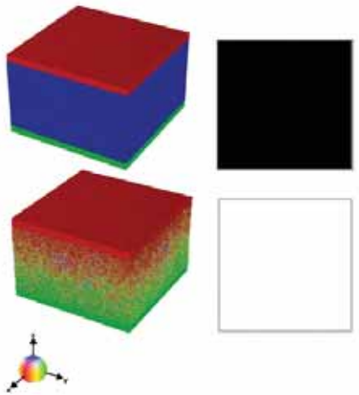

Figure 1: Snapshots of the molecular configurations (left) and corresponding computed pixel images (right) for: a) the initial dark state (top) and b) the final states (bottom), after 150 000 hours of Monte Carlo simulations. Molecules are colou r coded according to their orientations.

The initial configuration (see Figure 1a), is characterized by the LC molecules uniformly aligned perpendicular to the display surfaces to give the dark pixel state, while the display boundaries are modeled with layers of fixed particles oriented along the incoming and outcoming polarizer directions (rotated 90 degrees from one another). The system is then allowed to relax through the MC simulation to the "rest" equilibrium state, which is reached after only 150 000 hours of computing time (see Figure 1b).

Mechanism of clearing of the pixel

Following the time evolution of this equilibration process we notice that the mechanism of clearing of the pixel proceeds by the local induction of order by the two confining surfaces which induce an almost parallel alignment to the display. This is a very slow process, characterized by the progressive alignment of the molecular layers, starting from those closer to the two surfaces. The resulting distribution of local directors at this point is far from being that of a perfect linear helix. On the contrary, the molecular layers align almost parallel to the closest display surface, and are not affected by the farthest surface. Some sort of helicoidal ordering starts to form only when the two fronts of perpendicular oriented particles come into contact in the middle of the sample. However, the temperature dependent fluctuations of molecular orientations in the middle of the sample make the conventional picture of a uniformly helicoidal configuration too simplicistic, and evidence shows that this requirement is not necessary to achieve good optical behavior.

Having reached an equilibrated configuration corresponding to a light pixel, we switched on an electric field in the central region of the display. Our aim was to investigate the mechanics underlying the disruption of helicoidal order to give a black pixel. The interesting result was that the dark region of the pixel starts to grow from the centre of the square area affected by the field, expanding in concentric spherical shells, instead of uniformly all over that area (see Figure 2).

The calculations were carried out at CSC, the Finnish IT center for science, using 128 processors and at the CINECA Supercomputer Centre in Italy, using 64 processors, for a global amount of 200 000 CPU hours.

Figure 2. Pixel image computed after 50 000 cycles of a MC simulation with an electric field switched on in the middle of the display.
© Claudio Zannoni

Mikko Hakala, Toma Susi, Maria Ganchenkova, Timo Vehviläinen, and Adam Foster
Helsinki University of Technology, Finland

Impurities on insulating surfaces and in carbon nanostructures

THREE DIFFERENT PROJECTS WERE CARRIED OUT APPLYING FIRST PRINCIPLES CALCULATIONS IN ORDER TO EXAMINE THE IMPURITIES ON INSULATING SURFACES AND IN CARBON NANOSTRUCTURES.

Our DEISA effort was focused on three different projects, but where very similar first principles methods were applied to each project. The first project concerned the growth and properties of metallic nanoclusters on insulating surfaces. The second project considered the role that iron could play in increasing the magnetic coupling of intrinsic carbon defects in carbon nanotubes. The third project studied the potential of fullerenes for hydrogen storage.

Diffusion of metal atoms on insulating surfaces

The possibilities of fabricating nanoscale systems atom by atom is already a reality, and now researchers are looking for more practical applications of the technique, especially on insulating surfaces. One clear example is in studies of nanoparticle catalysis, where metal nanoparticles have been shown to be surprisingly reactive compared to the bulk. This reactivity is strongly coupled to the interaction of the metal particle with the substrate, and can even be experimentally controlled by the injection of charge carriers. In this study we considered a prototypical system, metal atoms on alkali halide surfaces, and considered how the bonding and diffusion of the atoms depends on their charge. We found that the mobility of Au and Ag is relatively insensitive to charge state, but that Pb can be more readily controlled, on any of the surfaces considered.

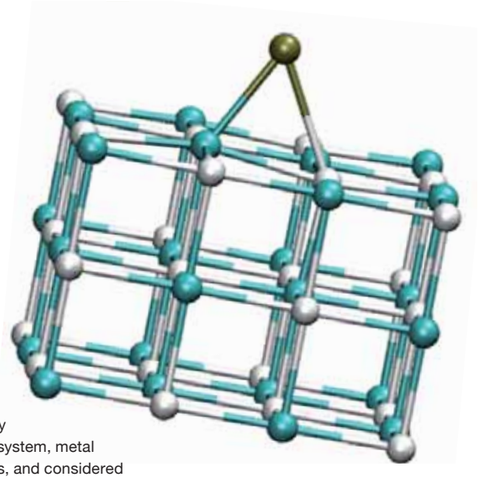

Figure 1: Snapshot of diffusion of a lead anion diffusing on the NaCl (001) surface.

Magnetism in carbon nanostructures

Traditionally, ferromagnetism has been thought to exist only in systems that contain ferromagnetic elements in sufficient quantities. Quite recent experimental results of magnetism in pure carbon systems have altered this view, although the mechanisms remain a topic of intense discussion. One possible explanation is that trace amounts of ferromagnetic impurities could "catalyze" the magnetism of intrinsic carbon defects, and encourage a ferromagnetic ground state. In order to study this, we investigated the role of iron adatoms on the magnetism of a (4,4) semiconducting carbon nanotube. We considered the coupling of carbon adatoms as a function of separation, and presence and location of iron adatoms. We found significant coupling between carbon adatoms in isolation on the tube, but this effect is not enhanced by the presence of iron, unless the iron adatom lies directly between the carbon adatoms – a statistically unlikely event.

The mobility of Au and Ag is relatively insensitive to charge state, but Pb can be more readily controlled.

Fullerene structures as a novel material for hydrogen storage

In this part of the project we used the VASP and CPMD *ab initio* codes to first calculate the electronic and kinetic properties of hydrogen interacting with C20 and C60 fullerenes. The adsorption energies, activation barriers for hydrogen surface diffusion as well as barriers for H penetration through the fullerene surface were calculated using both molecular static and molecular dynamics methods. Next, several different C20-based 3D structures were considered, and their electronic and elastic properties were studied. A new carbon structure, which we have named a quasi-graphite phase (QGP), was proposed. The QGP together with FCC type structures was found to be the most energetically favourable. The study of the hydrogen interaction with the proposed C20-based polymers is planned for the future.

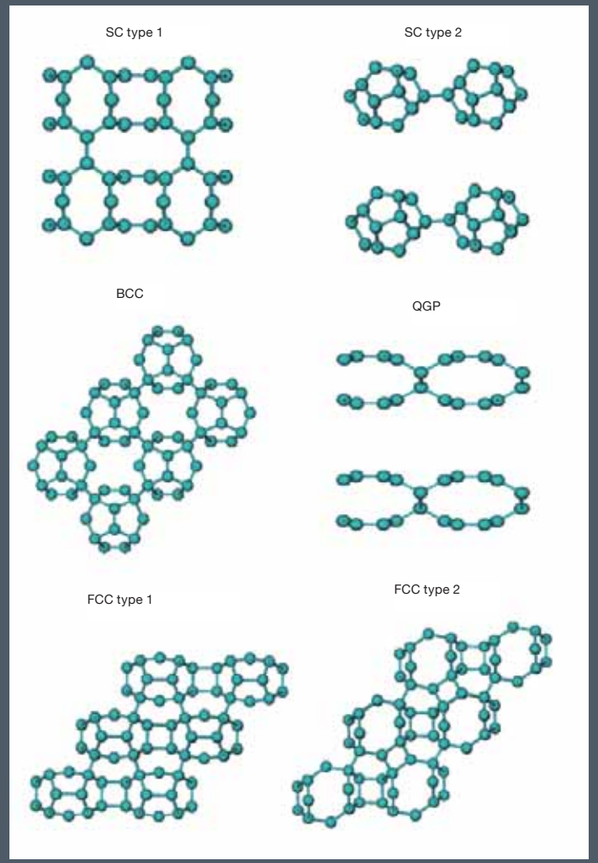

Figure 2: Representation of C20-based 3D structures considered within the project: simple cubic (SC), body-centred cubic (BCC), and face-centred cubic (FCC).

Benedict J. Reynwar, Gregoria Illya, Vagelis A. Harmandaris, Martin M. Müller, Kurt Kremer, and Markus Deserno, Max Planck Institute for Polymer Research, Germany

Curvy membranes make proteins attractive

When proteins bind to cell membranes and, in doing so, bend them, they can attract each other indirectly due to the membrane deformations they cause. With enough proteins available this may lead to a membrane invagination. Using large-scale computer simulations a research team at Max Planck Institute for Polymer Research has for the first time verified this physical model for the initial steps of vesiculation in cells. The research results were published in Nature 447, 461–464 (2007).

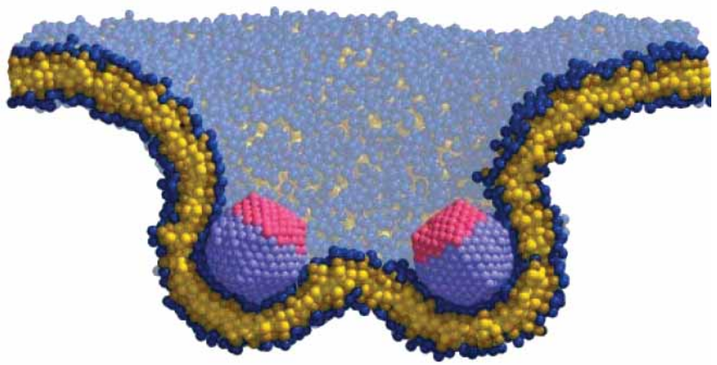

Figure 1: Cross-section of a membrane that is curved by two symmetrically adhering particles. In such a simulation the force between the particles can be "measured".

The cells of our body are highly complex biochemical factories in which thousands of substances are created, processed and decomposed. To control this metabolism, all eukaryotic cells (such with a cell nucleus) possess various distinct organelles that are responsible for specialized tasks: Our genome is stored and read in the nucleus, proteins and lipids are synthesized in the endoplasmic reticulum, and the Golgi apparatus takes for instance the task of sorting proteins.

All these organelles are formed by lipid membranes. These are flexible double layers of lipid molecules, only five nanometers thin, which also make up the exterior envelope of each cell. In order to enable transport of material between the organelles – as well as into and out of the cell – they can change their shape. Particularly important is a process called "vesiculation" (see Figure 1), during which a membrane bud forms that is later cut off the membrane. Proteins enclosed in the interior of such a "vesicle" can be transported to a different location inside the cell without getting mixed up along the way with other substances. How exactly all this happens is an important question in cell biology that is currently being studied with great intensity.

Evidence for a physical mechanism leading to vesiculation

Since such membrane deformations cost energy, the cell drives them using special proteins. Today we know in many cases their identity. Yet, how they actually create a vesicle is much less understood. We have now used computer simulations to provide evidence for a physical mechanism that can lead to vesiculation. Remarkably, it is not necessary that the involved proteins interact with each other, for instance by mutual specific binding. Rather, proteins influence each other indirectly by the deformation of the lipid membrane which they cause by adhering to it.

To create curved membrane structures, each protein has to bend the membrane a little bit. This local curvature spreads around a protein like a little "halo". When two proteins approach, the overlap of their halos may lead to an indirect interaction. One may think of the attraction between two balls lying on a tense rubber membrane. Usually this picture serves as an illustration of Einstein's theory of gravity by space-time curvature. And indeed these two seemingly different phenomena are formally closely related and can be described by similar mathematics.

Since almost two decades physicists have been on the track of membrane mediated interactions, yet the phenomenon remained confusing: on the one hand experiments documented attractions between membrane-bound objects. On the other hand all available theories indicated repulsion – at least if the membrane is curved uniformly. Since neither experiment nor theory are free of potential artifacts, the existence of curvature-mediated protein attractions remained elusive. Nevertheless, cell biologists started to be interested in the effect, as it promised a clear physical model of vesicle creation – provided it worked! The simulations performed in our group show that under suitable conditions (e.g., a minimum curvature imprint of each protein) the mechanism indeed leads to an attraction and with enough proteins available to cooperative vesicle formation. For very strong curvatures the force can even be "measured" directly in the simulation (see Figure 2).

Computer simulations support experimental techniques

The creation of transport vesicles is not just of fundamental importance for all life processes in cells, such as signal transduction or transport of nutrients. The involved mechanisms should also be at work in many other shape-determining tasks of membrane organelles. They furthermore play a role in the interaction of cell membranes with other objects, such as viruses or drugs. Since all this happens on sub-optical scales (about 100 nm), the experimental investigation of such events is a great challenge. Suitable computer simulations therefore support and complement available experimental techniques quite considerably.

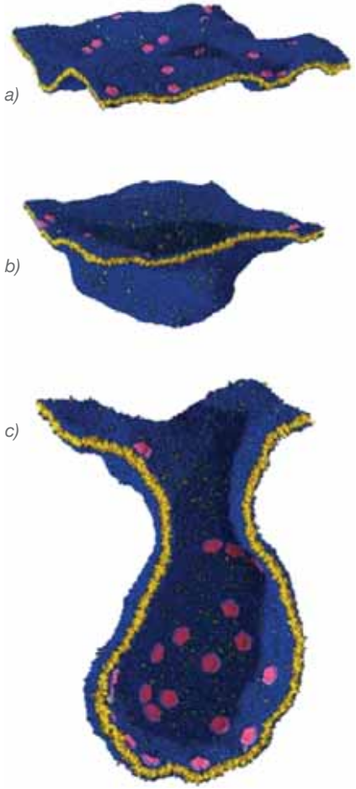

Figure 2: Vesiculation, as seen in a model simulation: a) proteins (red) adhere on a membrane (blue/yellow) and locally bend it; b) this triggers a growing invagination. In c) a cross-section through an almost complete vesicle is shown.

Peter V. Coveney, James Suter, Mary-Ann Thyveetil, Shunzhou Wan, and Stefan Zasada, University College London, UK. Chris Greenwell and Rick Anderson, Bangor University, Gwynedd, UK.

Nanocomposites and bio-inorganic composites

The goal of the British research group is to be able to simulate "life sized" clay mineral platelets that are not constrained by periodic boundary conditions, and thus can reproduce previously unexplored behaviour. The scale of simulation run in the DEISA programme now allows the group to calculate the material properties that are hard to obtain experimentally due to the disordered nature of the platelets.

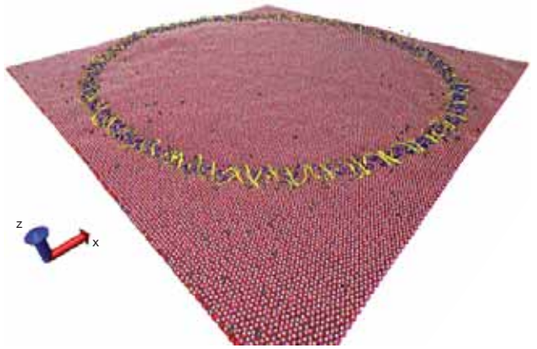

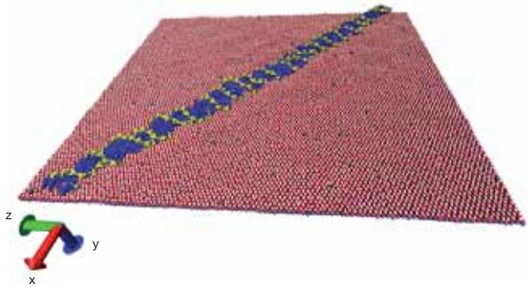

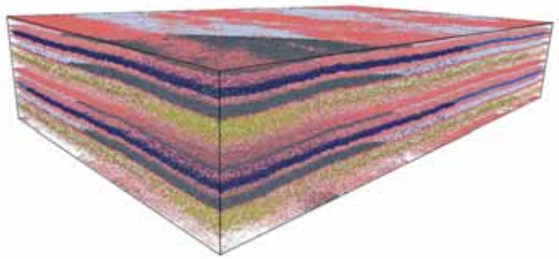

Our study shows that large-scale molecular dynamics simulations of layered silicate nanocomposites and bio-inorganic composites involving full intermolecular interactions are now possible at sizes representative of real clay mineral platelets using today's large-scale supercomputing facilities, such as those made available through the DEISA initiative. From the simulations in our study, we determined emergent properties such as silicate (clay) sheet undulations, which are necessary for the calculation of material properties; this is an important step in constructing a multiscale model of silicate nanocomposites that would capture many, if not all, of the properties of these systems allowing new composite materials to be designed *in silico*. We also investigated the structure and stability of linear and plasmid DNA intercalated between sheets of layered double hydroxides, which is important for understanding the use of these systems in gene therapy and drug delivery, and also has direct relevance to origins of life studies. On a technical level, we were able to demonstrate and effectively use the interoperability that exists between the NGS, TeraGrid, and DEISA using the Application Hosting Environment.

Mesoscopic phenomena occur only in large systems

Our goal in this study is to be able to simulate "life sized" clay platelets that are not constrained by periodic boundary conditions, and thus can reproduce previously unexplored behaviours, such as the unconstrained dynamics of clay sheets in water, or when constrained within a polymer matrix. Thus far, using "life sized" models under periodic boundary conditions has allowed, for the first time on multilayer systems, the calculation of material properties that are hard to obtain experimentally due to the disordered nature of the platelets. The size of the simulation we will undertake is more akin to a "mesoscopic" scale simulation but with the increased detail of full interatomic interactions. This is vitally important for nanocomposites, such as clay filler based composites, where one dimension of the filler is in the nanometre range. Without the detail of atomistic simulation, important features will not be included, such as clay–polymer interactions at the atomic scale that may be "smeared" out. However, large simulation sizes are required to capture the "mesoscopic" phenomena, such as undulations of the clay sheets, which occur over length scales of tens of nanometres. As yet, no simulation has been large enough and accurate enough to capture the motions and interactions of a clay sheet over such long and short length scales simultaneously to capture all the features of a layered silicate nanocomposite.

Two types of clay simulated

We have used the state-of-the-art molecular dynamics code LAMMPS. LAMMPS possesses almost linear scaling with the number of processors and system size. We used several efficient classical molecular dynamics algorithms that reduce wall clock time, which are required for the large-scale simulations. We also used the Application Hosting Environment (AHE), a lightweight hosting environment for running unmodified applications on grid infrastructures, to integrate DEISA resources with the US TeraGrid at the 2007 Supercomputing Conference and TeraGrid 2007 Conference. We focused on two different types of clay: montmorillonite, a cationic clay commonly used as filler in clay–polymer nanocomposites, and layered double hydroxides (LDHs), an anionic clay which can accommodate polar organic compounds between its layers and forms a variety of intercalation compounds.

Properties of systems extracted

We simulated several models: montmorillonite hydrated and surrounded by long chain poly(ethylene) glycol molecules, hydrated Layered Double Hydroxides with linear and plasmid 480 base pair DNA intercalated between the layers. From the simulations, we extracted the materials properties of clay platelets, which are required in the calculation of the elastic properties of the clay–polymer nanocomposites. These new types of composites have recently received much attention due to their enhanced materials properties with the addition of a small filler weight. We also addressed the intercalation of DNA into layered double hydroxides (LDHs), which has various applications including drug delivery for gene therapy and origins of life studies. The large size of the simulation allows the DNA to be simulated in three forms: double stranded, linear and plasmid up to 480 base pairs. From these simulations, we identified the structural changes the DNA undergoes when in contact with the LDH surface and suggested reasons for the enhanced stability at high temperatures and pressure, which may have enabled life as we know it to have emerged on Earth. The work has recently been published in the *Journal of the American Chemical Society*, the world's premier chemistry journal.

Large simulation sizes are required to capture the "mesoscopic" phenomena, which occur over length scales of tens of nanometres.

Giovanni Barcaro and Alessandro Fortunelli, IPCF-CNR, Pisa, Italy.
Paolo Giannozzi, University of Udine, Italy.

Metal adsorption on nanostructured ultrathin films

This project shows that the use of supercomputing facilities allows one to predict structure and properties of nanostructured ultrathin films, fully clarifying unresolved experimental issues.

Metal adsorption on two nanostructured titania (TiO_x) ultrathin films grown on the Pt(111) surface is investigated via a first-principles computational approach. These films present very regular nanopatterns extending over mesoscopic distances, with compact regions intercalated by point defects which act as nucleation and trapping centers, and are thus of great interest as nanotemplates for the growth of ordered supra-assemblies of equally sized (monodisperse) metal nanoclusters. However, due to the large size of their unit cells, their theoretical investigation has been so far hindered by computational limitations.

Our study shows that the use of supercomputing facilities, like those made available through the DEISA initiative, allows one to predict structure and properties of these systems, fully clarifying unresolved experimental issues such as a metal-induced phase transformation. The obtained results are also discussed as paradigmatic examples of the interaction of metals with polar ultrathin films of oxides grown on metal supports, a novel and promising field in materials science.

Metal on oxide on metal

Metal clusters adsorbed on oxide surfaces are the subject of considerable attention due to their scientific interest and technological applications, and this has fuelled a series of fundamental studies on model systems obtained by metal deposition on oxide single crystals in UHV.

More recently, attention has shifted from single crystals to ultrathin films grown on metal supports, as they allow STM measurements to be performed. The underlying metal support can, however, actively participate in the adsorption process, producing a wealth of novel phenomena and posing new issues and challenges. In this field, oxide polar nanophases represent a pioneering and even less unexplored topic, particularly promising for their templating characteristics.

Computational approaches are expected to play an important role in the study of these systems, but their application is limited by the large size of the unit cells. In the present study we focused on two such oxide polar ultrathin phases, namely the z- and z'-TiO_x (titania) phases grown on the Pt(111) surface exhibiting a zig-zag-like habitus in their STM images, and we investigated via a first-principles approach the characteristics of metal adsorption on them. The TiO_x/Pt(111) system is chosen as it presents a variety of phases with different topographies providing a fairly complete representation of the various possibilities, while the dimensions of the unit cells allow one to perform systematic investigations.

Computational approach

We used state-of-the-art first-principles electronic-structure techniques to investigate the structure of the titania films and metal adsorption on them. In particular, we used two plane-wave codes: PWscf for energy minimizations and wave function analysis, and CP (Car–Parrinello) for dynamics simulations, both taken from the ESPRESSO software package. The calculations, totalling 350 000 CPU hours, were carried out at the Louhi (Cray XT4) supercomputer of CSC in Espoo, Finland.

The project was mainly articulated in two steps: (1) potential energy landscape for the adsorption of single metal atoms; (2) small cluster epitaxy and influence of metal adsorption on the structure of the films.

Adsorption of single metal atoms

After confirming previous structural assignments via dynamic (CP) simulations on the bare titania phases, the first step of the project consisted in exploring the potential energy surface for the adsorption of single metal atoms. Au, Pd, and Fe atoms were selected as prototypical examples. A very different landscape was found to occur as a function of the electron affinity and the chemical properties of the adsorbed metal atom, with a peculiar surface reconstruction induced by a Au atom (see Figure).

Adsorption of metal clusters

Calculations were made in two steps. First, density-functional global optimisation (DF-GO) calculations were performed to elucidate the structure, epitaxy and energetics of adsorbed Au clusters. This was the first time that DF-GO calculations were performed on systems of this size and corresponded to the most computationally intensive part of the project. The simulated STM images of cluster growth were in agreement with experimental results. Then, the influence of the metal adsorption on the structure of the z'-film was studied as a function of temperature. It was found that Au (but not, e.g., Fe) induces a surface reconstruction from a rectangular to a hexagonal phase (one of the first examples of a metal- and heat-induced phase transformation in oxide nanophases). This metal-induced phase transformation was proved to be favourable at both the static (DF-GO) and dynamic (CP) level.

It was found that Au induces a surface reconstruction from a rectangular to a hexagonal phase.

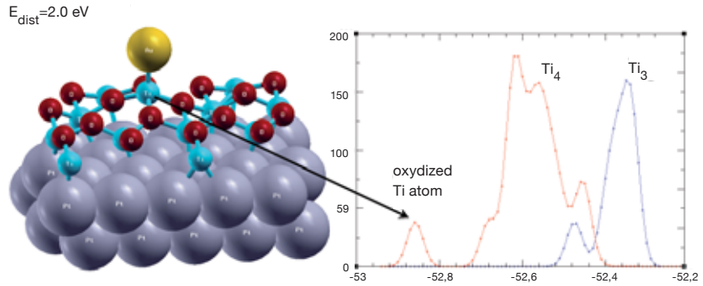

Figure 1: Surface reconstruction (left) after the adsorption of a gold atom on the z'-TiO_x/Pt(111) phase (gold, titanium, oxygen and platinum as yellow, light blue, red and grey, respectively) and Density of States (PDOS, right) projected on the Ti 3s (semi-core) states. The PDOS of the Ti atom interacting with adsorbed gold appears as a peak shifted at higher binding energies.

Karsten Reuter, Fritz-Haber-Institut der Max-Planck-Gesellschaft, Berlin, Germany

First-principles statistical mechanics for molecular switches at surfaces

THIS PROJECT CONCENTRATED ON ACHIEVING MICROSCOPIC INFORMATION FOR THE MOLECULAR SWITCH AZOBENZENE, WHEN ADSORBED AT THE TECHNOLOGICALLY RELEVANT CLOSE-PACKED SURFACES OF THE COINAGE METALS CU, AG, AND AU.

The large-scale density-functional theory calculations enabled by this DEISA project have provided unprecedented atomic-scale insight into the structural and electronic properties of the prototypical molecular switch azobenzene adsorbed at coinage metal surfaces. The computations reveal an interaction that is far more complex than the anticipated weak physisorption, particularly for copper, with the final bonding resulting from quite different interactions of different molecular moieties with the solid surface. The understanding emerging from the trend study of Cu, Ag, and Au surfaces clearly identifies Ag as a preferred system where a controlled reversible switching should be possible.

Azobenzene on copper, silver, and gold surfaces

In view of the rapidly advancing miniaturization in microelectronics and sensing, molecules are envisioned as fundamental building blocks in a future "molecular nanotechnology". Since controlled switching between defined states is a crucial basis component for storage and logic, molecules offering this functionality (e.g., through externally stimulated changes between conformational isomers) attain a central importance. Considering contacting and defined integration into a larger framework, it is more precisely the molecular function when the molecule is stabilized at a solid surface that is of key interest.

A necessary prerequisite for an atomic-scale understanding of this function of the adsorbed switch is a detailed structural and electronic characterization of the stable (or meta-stable) molecular states. The project thus concentrated on achieving this crucial microscopic information for the molecular switch azobenzene, when adsorbed at the technologically relevant close-packed surfaces of the coinage metals Cu, Ag, and Au. Azobenzene is a prototypical molecular switch known to undergo conformational changes between a cis- and a more stable trans-isomer in solution. The main objective was therefore to identify and characterize (meta)stable cis- and trans-adsorption geometries at the three metal surfaces.

Many DEISA sites employed in parallel

The CASTEP code was executed on many DEISA sites in parallel to minimize the project turn-around time, benefitting greatly from DEISA's united computing power, unparalleled in Europe. Groups of tasks were distributed over different sites, each employing 128 processors at the IBM Power4/5 sites RZG (Garching, Germany), IDRIS (Orsay, France), FZJ (Jülich, Germany), and CINECA (Bologna, Italy), and using up to 510 processors on the SGI Altix system HLRB-II at LRZ (Garching, Germany). The multitude of total energy and force evaluations required for the project involved computing the explicit electronic structure of up to one thousand electrons, expanded in a basis approaching a million plane waves. With each such evaluation taking of the order of several hours, parallel and efficient usage of all available DEISA sites was crucial to the success of the project.

Controlled reversible switching should be possible on Ag surfaces.

Best molecular switch: azobenzene on silver

Our study reveals an essentially zero net binding of both conformational isomers (cis- and trans) at Ag(111) and Au(111). This is significantly different at Cu(111), where the bonding, particularly of the cis-isomer, is even strong enough to reverse the gas phase energetic order of the two isomers. We are able to rationalize these findings as a competition between covalent bonding of the central azo (-N=N-) bridge to the substrate on the one side, and the surface interaction of the two closed-shell phenyl ($-C_6H_5$) rings on the other side, with steric effects due to the specific molecular geometry as an additional important factor. Deficiencies in the description of either of these interactions at the employed generalized gradient approximation level will therefore crucially affect the deduced energetic order of the (meta)stable adsorption modes. This concerns, notably, van der Waals contributions to the interaction of the $-C_6H_5$ moiety with the coinage metal surface. Currently, the identified correlation of the cis-isomer stability with the N-substrate bond strength already permits a first prediction of suitable candidate materials for a controlled reversible switching, namely those that offer an intermediate N-substrate bonding leading to roughly energetically degenerate cis- and trans-bound states. Corresponding experimental work was already initiated with silver as one material that fulfils this property within the local collaborative research centre (Sonderforschungsbereich) SFB-658 of the German Research Council.

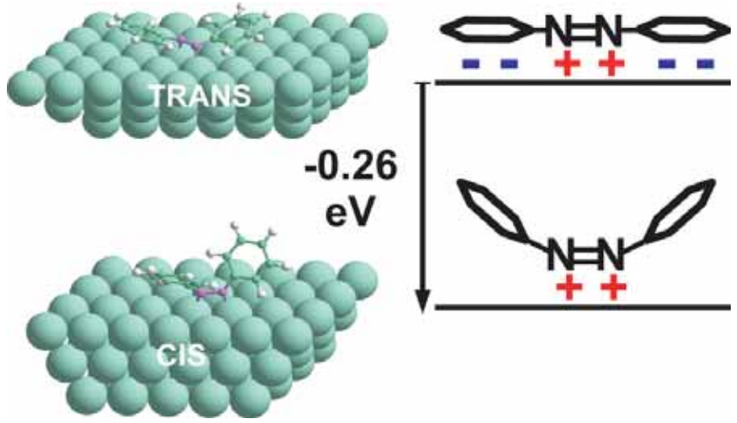

Figure 1: Left panels: Perspective views of adsorbed trans- (top) and cis- (bottom) azobenzene at Cu(111). Cu = green spheres, N = small violet spheres, C = small green spheres, H = small white spheres. Right panels: Corresponding schematic views of adsorbed trans- (top) and cis- (bottom) azobenzene, depicting the different interaction of the molecular moieties with the surface. In the two-dimensional trans-configuration the short covalent N-Cu bond lengths conflict with the strong Pauli repulsion of the closed-shell phenyl rings at these molecule-surface distances, eventually rendering the three-dimensional cis-configuration energetically more favourable.

K. Lackner, J. Nührenberg, and L. Villard, Max Planck Institute for Plasma Physics, Garching/Greifswald, Germany

Multiscale effects in plasma core turbulence

THE PRESENT SIMULATIONS CONSIDER HOW THE TURBULENT FLUCTUATIONS DRIVEN AT DIFFERENT SCALES INTERACT WITH EACH OTHER NONLINEARLY. GLOBAL FULL RADIUS PARTICLE-IN-CELL SIMULATIONS OF TOKAMAK TURBULENCE REVEAL THE NATURE OF THE HEAT TRANSPORT IN A MAGNETIC ISLAND OF THE PLASMA.

The main goal of the simulations performed with the gyrokinetic plasma turbulence code GENE was to better understand the role of turbulent fluctuations on very small spatio-temporal scales – smaller than the ion gyroradius and the ion transit time, respectively. These scales are usually neglected in plasma turbulence computations, but there has been increasing evidence both from theory and experiment over the last few years that they can play an important role and should be retained.

The present GENE simulations were targeted at understanding how turbulent fluctuations driven at different scales interact with each other nonlinearly. Some key results are shown in Figure 1.

It was found that no "superposition principle" exists for different turbulence types. For certain parameters, electron-temperature-gradient (ETG) turbulence contributes significantly to the transport and has a wide range of influence. They also affect scales with $k\rho_i \sim 1$. It turned out that there often tends to be a scale separation between ion and electron thermal transport. While the former is usually carried more or less exclusively by long wavelength fluctuations, a substantial fraction of the latter can be carried by much smaller scales.

Implications for full-torus simulations

This finding represents an important new insight into the physics of turbulent transport in magnetized plasmas, and it has significant implications for future full-torus simulations of large fusion devices like ITER. Including sub-ion scales will make the simulations more realistic, but it will also make them much more expensive than thought so far. Comprehensive gyrokinetic simulations of a device like ITER will require petascale or even exascale supercomputers.

In this context, the long-term goal is to provide the fusion community with reliable predictive simulations of the turbulent transport in fusion plasmas. Since this is crucial in determining the performance of a fusion device, it is one of the most important research areas in theoretical plasma physics.

Global full-radius simulations

Global full-radius particle-in-cell simulations of tokamak turbulence were conducted with the ORB5 code. The nature of the heat transport in the island and island stability were particularly investigated. Results are shown in Figure 2.

In this project, the gyrokinetic plasma turbulence codes ORB5 and GENE were used primarily at LRZ. The codes were optimised through the DEISA Joint Research Activity JRA3 and then ported to the SGI Altix system.

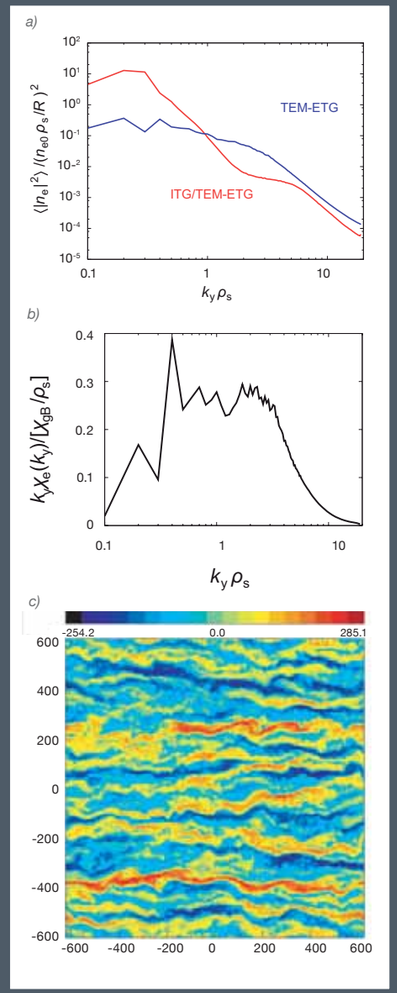

Figure 1 (left): Studies with GENE. a) Density spectrum for different turbulence mixtures – the superposition principle is violated. b) Turbulent electron heat flux – more than 75% of transport is in the high k regime ($k\rho_i > 0.5$). c) Snapshot of electrostatic potential contours in the outboard midplane.

Figure 2 (right): Studies with ORB5: Electrostatic turbulence in the presence of a magnetic island. The island structure is included in ORB5 by adding a magnetic perturbation in the parallel motion. a) Flattening of the temperature profile inside the island: reduction of the turbulence drive. b) Heat transport limited to island separatrix. c) Poloidal cross section: Electrostatic potential perturbation (solid lines: magnetic surfaces).

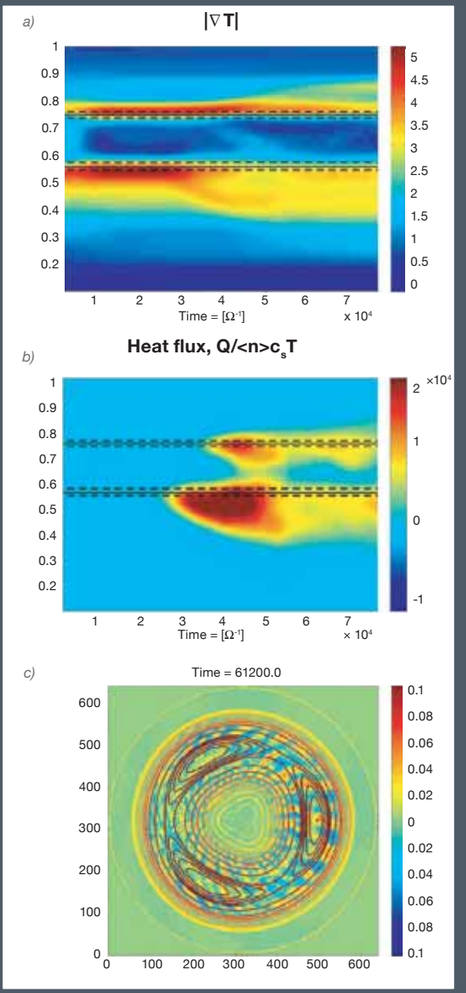

Jukka Heikkinen, VTT (Euratom-Tekes Association), Espoo, Finland.
Timo Kiviniemi, Helsinki University of Technology (Euratom-Tekes Association), Finland.

Edge turbulence in fusion plasmas

A tokamak is a fusion device where the fuel plasma, consisting of the hydrogen isotopes deuterium and tritium, is contained inside a toroidal vessel by strong magnetic fields. Using DEISA resources, simulation of the edge of the middle-sized ASDEX Upgrade tokamak with H-mode parameters has been launched with the ambitious goal of clarifying the mechanisms behind the confinement transitions. The study shows that simulation of both neoclassical and turbulence physics in order to study transport barriers in medium size tokamaks using the global 5D full-f gyrokinetic particle simulation code ELMFIRE for the edge plasma is possible with present large-scale supercomputing facilities.

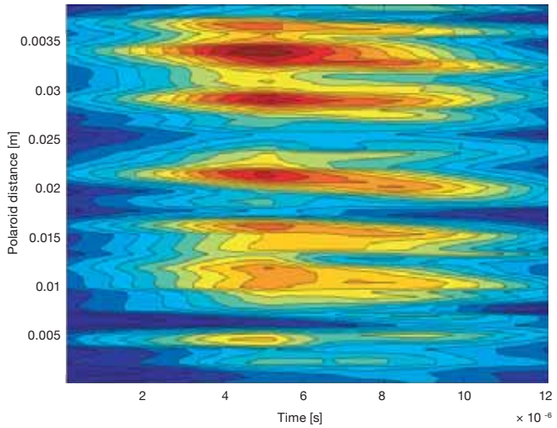

Figure 1: Modest poloidal rotation is observed from the correlation analysis of the density data for the ASDEX Upgrade, indicating an L-mode condition.

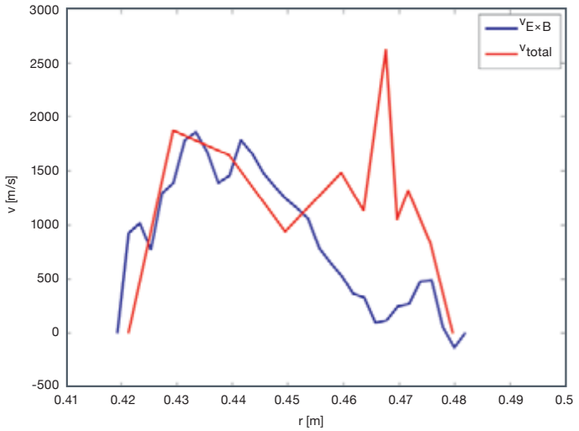

Figure 2: The rotation shown in Figure 1 is mostly explained by the E×B velocity but there is a discrepancy near the outer edge where both physics effects and numerical effects affect the results due to the boundary conditions. Shear in poloidal rotation is still below the threshold for strong turbulence suppression.

It is the plasma edge that, to a large extent, will decide if a magnetically confined in-vessel fusion is feasible. The transport of particles and energy is mainly determined by the strength of the plasma turbulence but can be significantly reduced in so-called transport barriers. In the L–H transition of the edge transport, a transport barrier is created by external heating, which is presently considered an important part of the reactor plasma operation. The enhanced confinement regimes resulting from the flow shear stabilization of turbulence are of considerable scientific interest; systems seldom self-organize to a higher energy state, with reduced turbulence and transport, when an additional source of free energy is applied. Such self-organization together with ion orbit effects may well be responsible for the L–H transition, but the very mechanism behind it is still not known.

The study of transport barriers with the chosen full-f gyrokinetic (GK) approach includes both neoclassical and turbulence physics. The numerical techniques used are valid even in the plasma edge where the gradients are steep and distributions can significantly deviate from Maxwellian. Extending such a self-consistent simulation from small to medium-size tokamaks is a computational challenge.

The ELMFIRE code breaks new ground

In the present work, a global 5D full-f GK particle simulation code, ELMFIRE, is used. This type of implicit GK particle solution method for full-f plasma quasi neutrality provides a rigorous treatment of such global and dynamic transport phenomena as transport barrier generation, intermittence, and wide-orbit effects. Since a full-f technique is used instead of a delta-f technique, the simulation of transport phenomena involving wide-orbit effects, steep gradients, and rapid dynamic changes in profiles becomes possible, but at the same time the requirement for CPU resources increases to reduce noise to an acceptable level. The particle orbits are solved in time with a toroidal configuration in a 5-dimensional phase space, and a 3D electrostatic potential solver is included to capture the turbulence that arises from the E×B convective cells in the presence of a pressure gradient and the toroidicity in the plasma, resulting in enhanced transport. ELMFIRE, as a full-f particle code, breaks new ground in gyrokinetic turbulence modelling.

Simulation of medium-sized tokamaks

The calculations were mostly carried out at RZG in Germany using the ELMFIRE code. Currently, after latest code optimizations, ELMFIRE can run with 500 processors in a 100×600×32 grid using 500 particles per cell, which is acceptable for turbulence saturation studies in reasonably sized annular volumes inside the ASDEX Upgrade toroidal plasma. However, for typical parameter scans, 66 million ions and electrons in a 30×600×4 grid with a 0.1 μs time step (running up to 100 μs in 48 hours with 128 processors) were used for a tentative analysis, although longer simulations (1 ms or more) are important in order to see the orbit effect with more realistic parameters. Our study shows that it is possible to simulate both the formation of an electric field and its impact on turbulence in the same simulation for medium-sized tokamaks as well. The importance of a proper heating model was highlighted in determining the time behaviour of transport coefficients. As real transport time-scale simulations are not yet possible, the results may be sensitive to the given initial profiles. More DEISA runs are scheduled for 2008 to test more realistic outer boundary conditions. In addition, more careful checking of the effect of numerical parameters and initialization is needed.

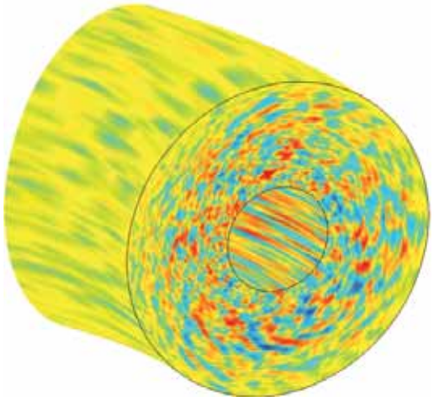

Figure 3: The poloidal cross section of saturated turbulence from an ELMFIRE simulation.

Kari Rummukainen, University of Oulu, Finland; Ari Hietanen, University of Helsinki, Finland

Hot and dense nuclear matter in particle physics and cosmology

Kari Rummukainen's researcher team at the University of Oulu studies the properties of quark–gluon plasma by means of extensive computer simulations. The team has developed a theory enabling scientists to apply computer simulations for studying the properties of plasma at very high temperatures. They have also produced a simulation program for studying events during the initial moments of heavy ion experiments.

During its initial moments, the universe was made up of extremely hot and dense matter. In these conditions, the structural elements of atomic nuclei – protons and neutrons – disintegrate and form a new, exotic state, quark–gluon plasma.

Kari Rummukainen's researcher team at the University of Oulu studies the properties of quark-gluon plasma by means of extensive computer simulations. The team has developed a theory enabling scientists to apply computer simulations for studying the properties of plasma at very high temperatures. These simulations produce much more precise data than simulations of the original physical theory of quarks and gluons, i.e., quantum chromodynamics. The team has also developed a simulation program for studying events during the initial moments of heavy ion experiments. In its work, Rummukainen's team utilized the DEISA resources at CINECA in Italy and CSC in Finland. The total CPU-time used is 125 000 processor hours, 32 or 64 processors at a time.

Focus on thermodynamic properties of hot quark–gluon plasma

Quark–gluon plasma is studied experimentally in heavy ion collisions. In these experiments, two heavy atomic nuclei are collided together at very high energy. If the energy is great enough, the collision yields a microscopic droplet of quark-gluon plasma. An RHIC (Relativistic Heavy Ion Collider) experiment is currently under way in Brookhaven, USA, and the largest heavy ion collision experiment designed so far is being constructed in the LHC particle accelerator at CERN.

In its recent simulations, Rummukainen's team has focused on two themes: the thermodynamic properties of hot quark–gluon plasma, such as pressure and particle number fluctuations, and the stormy initial moments of heavy ion collisions when the quark–gluon plasma is just in the process of formation.

Effective theory of quantum chromodynamics developed

Many researcher teams all over the world are studying the thermodynamics of quark–gluon plasma by means of computer simulations. For this purpose, Rummukainen's team has developed an effective theory of quantum chromodynamics, which enables the team to use computer simulations for studying plasma properties at very high temperatures much more precisely than by simulating the original quantum chromodynamics. In particular, the effective theory is useful in making a clear distinction between the perturbative and non-perturbative elements of thermodynamic quantities, such as pressure. This distinction is very important in terms of theory: the perturbative element can – at least in principle – be calculated analytically, whereas direct computer simulations are the only means known for calculating the non-perturbative element. Some preliminary findings have already been published about the pressure of quark–gluon plasma and about fluctuations in the number of quarks.

Simulation program for studying events during the initial moments of heavy ion experiments

Rummukainen's second project concentrates on the initial moments of heavy ion experiments. The collision of two nuclei with each other is immediately followed by a state of enormous energy density. This state is far from thermal equilibrium and it develops rapidly. It is not yet known for certain how this development towards thermal equilibrium takes place, but findings from heavy ion experiments indicate that it is very fast.

Rummukainen's team developed a simulation program that is suited for studying the above state of non-equilibrium. Simulations have shown that when plasma is in a state of non-equilibrium at the initial stage, it can generate a phenomenon known as plasma instability: waves that grow exponentially. According to the team's latest simulation findings, when this instability has increased sufficiently, it can push the system very rapidly towards thermal equilibrium. Plasma instability thus plays a crucial role in determining the physics of heavy ion collisions.

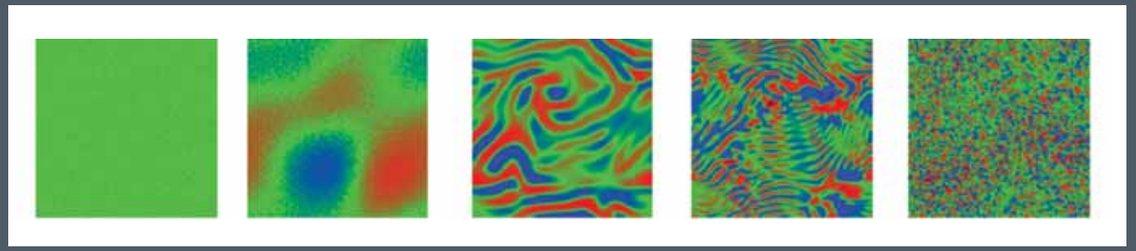

Figure 1: Plasma instability in a heavy ion collision experiment. Instability generates a rapidly growing wave in the gluon field; as it grows, the wave also moves towards shorter wavelengths. The final state has an almost complete thermal distribution.

Karl Jansen, John von Neumann Institute for Computing, Zeuthen, Germany

QCD simulations with light quark flavors

In the framework of non-perturbative computations for the theory of the strong interaction of the fundamental particles, a researcher group lead by Karl Jansen has explored the regime of small values of the pseudoscalar mass. DEISA resources helped significantly to accomplish this ambitious goal.

The project "QCD with twisted-mass fermions" (TMQCD) was conducted within the DEISA Extreme Computing Initiative (DECI). The framework of the project is non-perturbative computations for our theory of the strong interaction of the fundamental particles, the quarks and gluons. The theory describing this interaction is Quantum Chromodynamics (QCD). By discretizing space and time on a 4-dimensional space-time lattice it becomes possible to simulate QCD with numerical simulations which are, however, extremely demanding.

Exploring the regime of small values of the pseudoscalar mass

The aim of the project was to explore the regime of small values of the pseudoscalar mass, $m_{PS} \approx 250$ MeV. When started, this constituted a very challenging project since in the past it has been unthinkable that lattice QCD calculations could reach such small pseudo scalar masses and obtain precise results.

However, within this DECI project this ambitious goal could be fully accomplished. Although in total it has not only been the DEISA resources that led to these results, it is certainly so that the DEISA resources helped significantly to successfully complete the project.

Results

As an example for the outcome of the project, in Figure 1 the dependence of the so-called pseudo scalar decay constant f_{PS} as a function of the quark mass is shown. In the plot, the numerical data are fitted against the prediction of chiral perturbation theory which is an analytical, effective low energy description of QCD. As can be seen, the fit is excellent and the very precise numerical simulation data allow determining a number of the low energy constants of chiral perturbation theory with a high, world record precision. The low energy constants in turn allow computing other physical quantities such as the scalar and the tensor scattering lengths of the pion, which can serve as a severe test of QCD as the theory of the strong interactions. The results of this DECI project establish the most precise constraints on these scattering lengths so far.

Positive experience with the DEISA global file system

This DECI project was one of the pilot projects within DEISA. It is thus not too surprising that the project encountered a number of difficulties in the beginning. However, many of these difficulties could be solved with the excellent help of the DEISA staff people at the supercomputer centers involved.

A very positive aspect of the project was the experience with the DEISA global file system GPFS-MC. In combination with the usage of UNICORE, it was very helpful and easy to define, submit and control the jobs and, in particular, the workflow, since the evaluation of physical observables on the raw data, the configurations, could be done on different machines of the DEISA infrastructure in a flexible and heterogeneous manner.

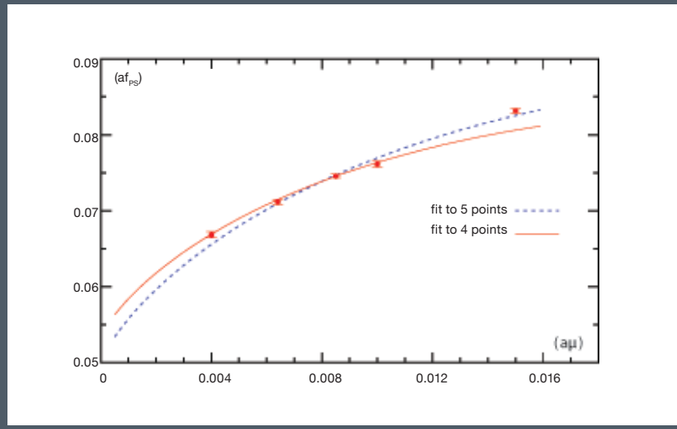

Figure 1: We show af_{PS} as a function of the quark mass $a\mu$ together with fits to chiral perturbation theory formulae. We present two fits, one taking all data and one leaving out the point at the largest value $a\mu = 0.015$.

Guido Arnold, Marcus Richter, and Binh Trieu, Research Centre Jülich, Germany. K. Michielsen, European Marketing and Business Development bvba, Wommelgem, Belgium. H. De Raedt, University of Groningen, the Netherlands. Thomas Lippert, Research Centre Jülich, Germany.

Improving quantum computer simulations

THE IMPACT OF GATE IMPERFECTIONS AND DECOHERENCE ERRORS AT GATE LEVEL WERE ANALYZED IN THIS PROJECT. AS A RESULT THE DEPENDENCE OF THIS SENSITIVITY TO BOTH ERROR SOURCES ON THE SYSTEM SIZE WAS QUANTIFIED.

Operational imperfections and decoherence errors are approximately modelled at gate level as an extension of our massively parallel quantum computer simulator. Decomposing a universal set of basic gates into plane rotations and phase shifts allows Gaussian distributed angle and phase errors to be introduced as effective imperfections of steering pulses acting on a qubit in a physical system. Combined with a simple decoherence model, we investigated the impact and the interplay of operational and decoherence errors.

Massively parallel gate-level simulator

As a first step towards the simulation of realistic many-qubit quantum computer devices, we implemented a massively parallel gate-level simulator, which allowed us to simulate quantum systems up to 37 qubits and which required 3 TB of memory on a high end system such as the IBM Regatta p690+. The simulation of quantum computers is clearly memory-bounded. Highly optimised memory access and communication patterns allow efficient simulation using up to 1024 processors. Within the idealized framework of gate level quantum computer simulation it is possible to approximately model the impact of gate imperfections and decoherence by introducing a simple error model. Given a certain confidence level, we ascertained the thresholds for σ (parametrizing the operational error) and p (as a measure of decoherence) in real applications, such as the Quantum Fourier Transformation or Grover's search algorithm. We computed the error norm of the resultant state vector $e^2(\sigma,p) = |\Psi - \Psi_{corr}|^2$. We used DEISA resources at the Jülich Research Centre to perform the very demanding stochastic simulations on system sizes up to 32 qubits.

Quantum Fourier Transformation

We have tested the robustness of the Quantum Fourier Transformation (QFT) circuit since it is the quantum kernel of Shor's factorization algorithm using $n = \log_2 N$ qubits for system sizes n = 8, 16 with 100 000 repetitions and n = 24 with 10 000 repetitions per experiment.

The results suggest a critical behaviour of the system depending on σ. The system is very robust against operational errors, since we found identical curves of $e^2(\sigma,p)$ for all $\sigma \leq 10,^{-2}$ even for the largest system investigated. Larger operational errors increase the error norm, the larger the system. To quantify the dependency of decoherence errors on the system size, we calculated the simulation results using a 32-qubit system.

Grover's search algorithm

In order to search for an element in an unstructured database of $N=2^n$ entries, we supposed the oracle function $f_k: \{0,1\}^n \rightarrow \{0,1\}$ given as $f_k = \delta_{kx}$ marking the searched element at position k. Using an ancillary qubit, we can implement Grover's quantum search algorithm on system sizes of 8+1, 16+1 and 23+1 qubits. We investigated the amplitude Ψ(k) of the database element we were searching for, expecting its (first) undisturbed maximum after l_{max} =12, 201, 2274 Grover iterations, respectively. In case of non-vanishing decoherence, we can see damping of the value of the amplitude $|\Psi_{max}(p,\sigma = 0)| < |\Psi_{max}| \approx (1 - 1/N)^{1/2}$. The superposed decoherence process leads to a maximum shifted towards $l < l_{max}$. The shift increases with higher n. In the case of operational errors, we saw a more robust behaviour but switching to an appropriate deviation level we also saw a clear shift and damping of the maximal amplitude. In contrast to the QFT algorithm, the σ-threshold is lower and very sensitive to the system size.

Larger operational errors increase the error norm the larger the system is.

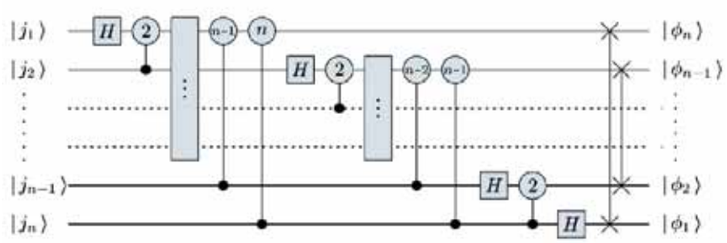

Figure 1: Circuit of the quantum Fourier transform for an n-qubit system.

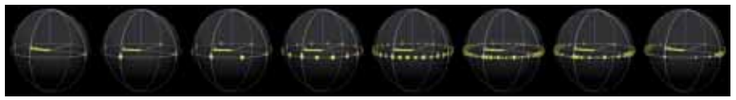

Figure 2: Sketch of the ensemble average qubit vector and cloud of end points measured at the end of each single experiment run.

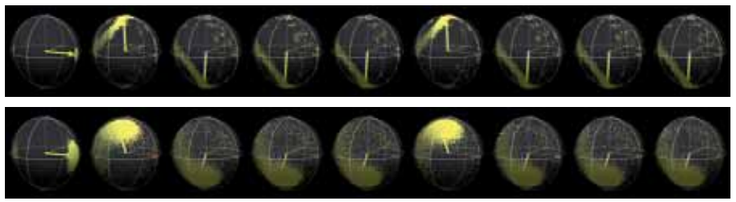

Figure 3: Sketch of the ensemble average qubit vector for the 8+1 qubit system after l=12 Grover iterations with operational errors of subcritical (upper picture) and overcritical values (lower picture).

Conclusion and outlook

We have analyzed the impact of gate imperfections and decoherence errors at gate level. We can quantify the dependence of this sensitivity to both error sources on the system size. We have seen that the quantum Fourier circuit is more robust against disturbance than Grover's algorithm on comparable system sizes.

These results should be compared with future calculations from realistic (=dynamic) simulations of quantum computer devices, taking into account the full time evolution according to a time dependent Hamiltonian describing both the system and the environment.

In the next step we will use our massively parallel quantum computer simulator to numerically investigate the characteristics of different error correction schemes.

Thomas Soddemann, Johannes Reetz, Andreas Schmidt, Hermann Lederer,
RZG, Max Planck Society, Germany

Materials science

THE JOINT RESEARCH ACTIVITY JRA1 HAS SUPPORTED AND DEPLOYED IMPORTANT MATERIALS SCIENCE SIMULATION CODES FOR SIMPLE AND EFFICIENT USAGE WITHIN THE DEISA ENVIRONMENT. THE TEAM HAS DEVELOPED A DEDICATED PORTAL FOR MATERIALS SCIENCE THAT OFFERS A STRAIGHTFORWARD ACCESS TO DEISA RESOURCES.

Physical, chemical, and biological processes for many problems in computational physics, biology, and materials science span length and time scales of many orders of magnitude. For example, on the microscopic level, bond distances between atoms are typically in the order of Angstroms, and atoms vibrate at a frequency of approximately 10^{13} Hz. However, on the other hand, phenomena and applications of practical interest occur on a time scale of seconds, and system sizes can be microns or larger. A Grand Challenge in computational physics, biology, and materials science is to link these vastly different time and length scales.

Efficient and user-friendly usage of materials science simulation codes

Objectives of the JRA1 team are the support and deployment of important materials science simulation codes for efficient and user-friendly usage within the DEISA grid infrastructure. In close collaboration with leading computational scientists in Europe, important applications have been selected for special support and integration into the DEISA environment.

Seamless, straightforward access to applications and simple job support

The work has been focused on the development of a dedicated DEISA Materials Science Portal, on the optimization of important materials science applications, and their integration into the DEISA environment as well as the DEISA Materials Science Portal. The selection of codes was triggered by requirements from the DEISA Extreme Computing Initiative (DECI) and the work was carried out in close collaboration with authors of the respective codes and with the DEISA user Support Activity SA4.

The DEISA Materials Science Portal was developed and deployed, and seven materials science codes have been supported therein: NAMD, LAMMPS, GROMACS, CPMD, CP2K, PWScf, and Wien2K.

The DEISA Materials Science Portal provides simple, seamless, and straightforward access to these important off-the-shelf applications and offers simple job support as well as components tailored specifically for the targeted materials science applications.

The DEISA Materials Science Portal

The DEISA Materials Science Portal offers a straightforward access to DEISA resources, following a threefold strategy:

- A rich internet application based access independent of the location of the user
- DEISA Materials Science Portal Web Services, enabling integration into the user's own applications
- UNICORE Client with application-specific plug-ins

Rich Internet application

The DEISA Materials Science Portal offers a rich internet application based access independent of the location of the user. Accessing compute jobs and their data from anywhere anytime with just a web browser becomes more and more important for DEISA users. Portals leveraging Web 2.0 technologies provide the necessary means to allow users to create, submit, and manage jobs securely from any location with Internet access. The DEISA Materials Science Portal is based on the Portlet Specifications (JSR 168, 286, 303) and makes extensive use of AJAX in order to provide the best usability experience. Furthermore, all components are directly accessible and can hence be used in third party created content mesh-ups.

Access via UNICORE client

Many users are already familiar with the UNICORE 5 client. It is a rich client application used for assembling, submitting, and managing jobs in UNICORE based Grid infrastructures, such as DEISA. The UNICORE client offers the user a variety of plug-ins to enhance its basic functionality e.g. by supporting specific applications. The JRA1 team developed and enhanced UNICORE client plug-ins for CPMD and Wien2k which are available for download via the DMAS portal and from the UNICORE project's web site (only CPMD).

SOAs with DEISA Materials Science Portal web services

As a side effect of the DEISA Materials Science Portal architecture, all functional components can be accessed utilising Web Service interfaces. Hence, users can create own applications by making use of the versatile Web Service API. Therefore, composite applications in terms of Service Oriented Architectures (SOAs) can be created. Here, interoperation between different grid infrastructures has been established, as already shown for AHE (Application Hosting Environment) from UK's RealityGrid.

Supported applications

Currently, DEISA Materials Science Portal (DMAS) supports various applications at various levels of integration. The following table gives an overview of the status as of February 2008.

Field	Application	UNICORE Client	DMAS Portal	DMAS Web Service
Classical Molecular Dynamics				
	NAMD	*	Supported	Supported
	LAMMPS	*	Supported	Supported
	GROMACS	*	Planned*	Planned*
DFT MD				
	CPMD	Supported	Supported	Planned
	CP2K	*	Supported	Supported
	PWScf	*	Supported	Supported
	Wien2K	Supported	*	*

*Supported via Simple Job Submission

Table 1: Supported application specific functionalities include: Input file creation support, syntax validation and semantics verification, management of dependent files like molecule conformations, structure information and pseudo-potential descriptions.

Figure 1: DEISA Materials Science Portal: Snapshot of a molecule view with the help of the embedded Jmol applet.

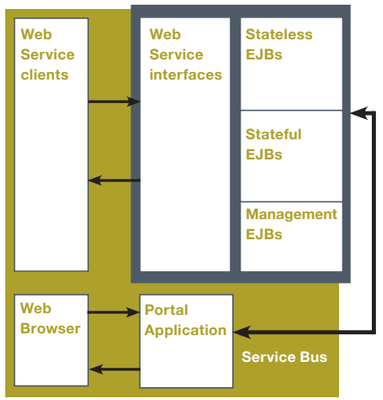

Figure 2: Sketch of the DEISA Materials Science Portal architecture and integration of the web service enhanced JMEA implementation.

Gavin J. Pringle, EPCC, Edinburgh, UK

Cosmological applications

Within the joint research activity JRA2, two of the Virgo Consortium's cosmological simulations, were developed to take advantage of DEISA. A number of serial pre- and post-processing tools have been developed for the Virgo Consortium's collection of simulation codes.

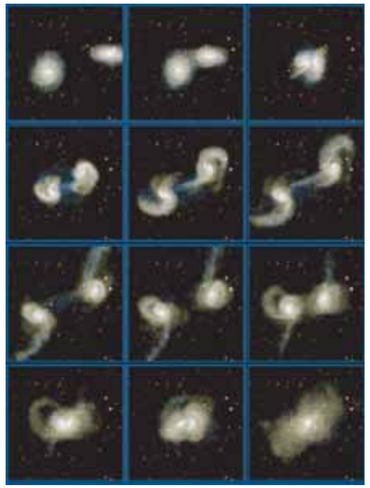

Figure 1: A GADGET simulation shows what happens when two disk galaxies collide and coalesce.

The Virgo Consortium formed 1995, is a large international collaboration of astrophysicists dedicated to carrying out the largest and most precise simulations of the formation of cosmic structure. Within the joint research activity JRA2, EPCC and the Virgo Consortium worked in close partnership to develop two of Virgo's cosmological simulations, namely GADGET-2 and FLASH, along with the associated pre- and post-processing tools, to take advantage of DEISA. This work laid the foundations for Virgo's Theoretical Virtual Observatory (TVO), complementing the approach of the European Virtual Observatory (EuroVO). The TVO represents one of the ultimate aims of the Virgo Consortium: to develop an integrated environment providing uniform access to the consortium's results through a fully functional Web-based TVO. JRA2 was a 2-year-project jointly funded by both DEISA and Virgo's initial TVO project VirtU.

GADGET as a coupled simulation

GADGET is an N-body/gas dynamic simulation and is ideal for simulations of, for example, individual dark halos or galaxies in which there are only a limited number of very high density regions, but is also employed to model the evolution of the entire Universe.

GADGET-2 was ported to many DEISA platforms, where it was used to investigate if simulations could be run in a loosely coupled manner, whereby we utilise more than one platform simultaneously. This might be considered metacomputing, however, GADGET-2 involves highly-localised inter-processor communications. To this end, PACX-MPI was introduced to the hosts and code, and the existing load-balancing aspect of the simulation was utilised.

Initial tests using a small, local three-node Grid, gave promising results, however, running across DEISA platforms proved inefficient. This was ascribed to both employing an old version of PACX-MPI and the code's own load-balancing algorithm. In order to compensate for DEISA's heterogeneous platforms, the load-balancing routine was extended to include a measure of actual execution times, to force more work to more powerful resources. Unfortunately, our initial implementations were found to be unstable, and further work to stabilise this method is required.

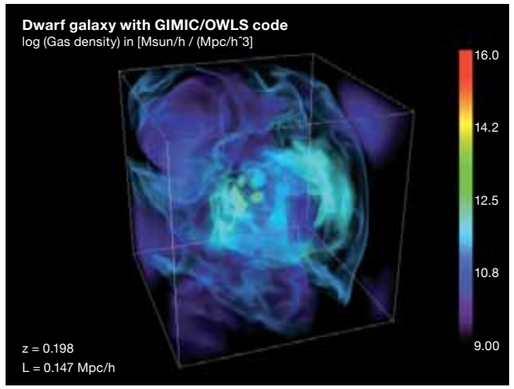

Figure 2: A high resolution simulation showing the formation of a small galaxy. Exploding stars in this galaxy create a wind which blows material out of the galaxy. This wind disturbs the intergalactic medium surrounding the galaxy and transports heavy elements into intergalactic space. The simulation uses the same code as the GIMIC simulations which were run under DECI but the numerical resolution of this simulation is much higher. © Rob Crain and Jim Geach, ICC, Durham and the Virgo Consortium.

FLASH optimised and launched via workstation scripts

FLASH is an AMR code and was in part developed by the DOE-supported ASC/Alliance Center for Astrophysical Thermonuclear Flashes at the University of Chicago. The code employed was a version which had been adapted by the Virgo Consortium and the FLASH team to model the evolution of the Universe. FLASH is a grid-based code and, as such, is more suited to capturing shocks than GADGET. Further, employing FLASH will enable a more robust validation processes for Virgo's huge simulations.

The code was ported to DEISA architectures of different vendors to perform basic profiling. This indicated which routines were computationally expensive and permitted the introduction of platform-independent optimisations. Working closely with Virgo, the iterative Multi-Grid dark matter Poisson solver was replaced with a non-iterative FFT-based version, which ensured that this routine would scale to a large number of processors, thus allowing very large simulations. The optimised version of the Virgo Simulation is now between 6-9 times faster.

Further, additional routines were introduced to permit so-called Code Migration, wherein a job reaching the batch clock limit on one DEISA platform would then resubmit itself, migrating to any another DEISA platform if necessary, to continue the simulation in question. Allowing codes to migrate in this way ensures the fastest time to solution. This is not to be confused with Job Migration, where a job is moved from one batch queue and submitted to another platform's batch queue. Code Migration was implemented [9] in both UNICORE and the DEISA Services for the Heterogeneous management Layer (DESHL), where the DESHL is both an API and a command line interface to UNICORE.

The Code Migration bash script, which utilised the DESHL, has since been and can now be used to employ all the DEISA platforms as a vast Task Farm, where each task can be a Simulation Workflow, and each Workflow stage can be restarted/migrated as necessary.

Grid-enabled toolset developed

A number of serial pre- and post-processing tools have been developed for the Virgo Consortium's collection of simulation codes. These tools include so-called Group Finders, tools to determine Halo Merger Trees and tools for evaluating basic properties of the dark matter, such as its density field, correlation and power spectra.

This work incorporated an investigation into the performance of different portable, binary data formats, to ensure efficiency and functionality within DEISA's heterogeneous Grid. It was found that the performance of HDF5 exceeded all other formats tested, including the default, vendor specific binary I/O. Furthermore, HDF5 can interface with seamlessly EuroVO's data format, namely VOTable.

The tools were parallelised, where necessary, and converted to employ HDF5. Since both GADGET and FLASH employ HDF5, Virgo's simulations, and their associated tools, now take advantage of the DEISA infrastructure, as well as operating within a conventional distributed environment for interoperability outside of the DEISA framework.

Roman Hatzky, R. Tisma and H. Lederer, RZG, Max Planck Society, Germany

Plasma physics

A team within the Joint Research Activity JRA3 has enabled and optimised important European plasma physics simulation codes for adequate usability on state-of-the-art supercomputers available in DEISA. All readily enabled codes were selected for successful projects in the DEISA Extreme Computing Initiative.

Research on magnetic confinement fusion has undergone large changes during the last decade, moving away from the semi-empirical, predominantly experiment-driven approach to one accompanied and supported in all areas by first-principle based modelling. This development has been particularly dramatic in the area of turbulence, which is held responsible for energy and particle transport in magnetically confined plasmas. Supercomputers are absolutely essential as an instrument of numerical modelling.

The objectives of the JRA3 are the enabling and optimisation of important European plasma physics simulation codes for adequate usability on state-of-the-art supercomputers available in DEISA. Simulation results will allow the exploitation of existing and planned large experiments more efficiently, including, in particular, the tokamak and stellarator experiments ITER and Wendelstein 7-X, and to proceed more rapidly to a practicable fusion power plant.

Road to numerical plasma physics devices

The focus of this JRA is turbulence in fusion experiments in Europe. The numerical methods for such simulations are state-of-the-art and still subject of research. Consequently, different numerical approaches are necessary for evaluation purpose of the simulation results. Rigorous simulations from first principles are still out of scope with even the fastest parallel supercomputers. Hence, the class of codes focuses on different aspects of the complete problem by either reducing the complexity in geometry or certain aspects of physics. However, the different approaches are complementary and milestones on a road which leads to a tool box of numerical plasma physics devices. All code modifications occurred in close collaboration with the code developing scientists.

Enabling salient plasma physics simulation codes

The following simulation codes have been enabled and optimised: TORB/ORB5, GENE, EUTERPE, GEM and GYGLES. All enabled codes (but the most recently enabled one, GYGLES) were selected for three successful projects in the DEISA Extreme Computing Initiative (DECI) from the calls in 2005, 2006, and 2007.

TORB/ORB5

The ORB code family consisting of the global codes TORB (theta-pinch) and ORB5 (tokamak) uses a particle-in-cell (PIC), time evolution approach, and takes advantage of all the recent techniques of noise reduction and control in PIC simulations. It bases on a Monte Carlo method using hundreds of millions of particles. The ongoing code development is made under a close collaborative effort of CRPP, Lausanne and IPP, Garching.

Both the memory consumption and the performance have been significantly improved by e.g., replacing the originally used band solver by a sparse format matrix solver. Consequently, much larger physical domains can now be simulated with the same number of processors. In addition, the scalability of ORB5 has been significantly improved by applying the domain cloning concept. As a result, ORB5 can now be used efficiently up to several thousands of processors (hyperscalability).

Speedup

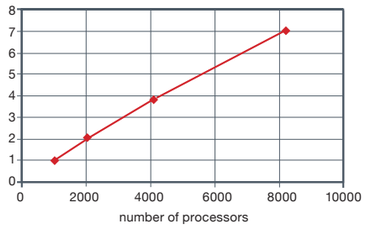

Figure 1. Strong scaling of ORB5 for an ETG simulation with 8 × 108 particles (on BlueGene/L at the IBM Watson Research Center in co-processor mode; normalisation on the 1024 processor result).

GENE

The gyrokinetic continuum code GENE employs a combination of spectral and finite difference methods on very large grids in five-dimensional phase space. GENE uses an efficient flux-tube approach and can deal with arbitrary toroidal geometries (tokamaks or stellarators), retaining more or less all physical effects which are presently considered relevant. The GENE code has also been enabled for hyperscalability. The overall data structure has been adapted, domain decomposition has been extended to four of the five phase space coordinates, and FFT usage has been optimised. As a result, the original limit of 64 parallel processes could be extended by three orders of magnitude.

Speedup

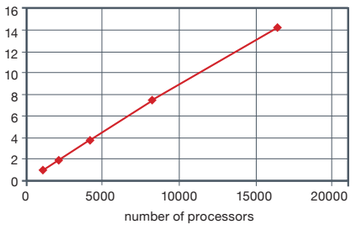

Parallel efficiency

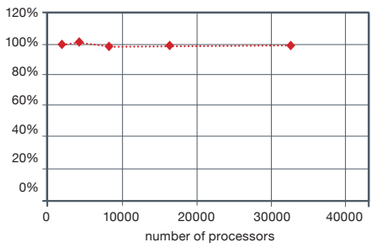

Figure 2. Scaling of GENE on the BlueGene/L at the IBM Watson Research Center.
Top:
Strong scaling with fixed problem size ~ 0.5 TB (in co-processor mode, normalisation to 1024 processors).
Bottom:
Weak scaling with increasing problem sizes from 0.5 TB to 8 TB (in virtual node mode; results normalised to 2048 processors).

EUTERPE

The EUTERPE code solves the nonlinear electrostatic gyrokinetic equation for ions (the electrons are assumed to be adiabatic) in the whole plasma domain (full radius) for a real three-dimensional stellarator geometry. It is presently the only code with such capabilities. The code was developed at CRPP-EPFL (Lausanne) and substantially upgraded at IPP, Garching.

Originally, the code was limited to a maximum usage of typically 64 processors. Through the enabling work, this restriction has been removed by application of the domain cloning concept. For weak scaling, a parallel efficiency of 98% could be achieved by increasing the processor number from 64 to 512 processors. In addition, the code was optimised for large sparse matrices in several ways: the matrix construction time has been decreased by an order of magnitude and all the processors can now take part in the parallel matrix solve of the PETSc library.

GEM

The GEM (Gyrofluid ElectroMagnetic) code solves nonlinear gyrofluid equations for electrons and one or more ion species in tokamak geometry. It has a central role in collaborations involving IPP/Garching, CFN/Lisboa, the University of Innsbruck and FZ-Jülich.

The single processor performance could be optimised by more than 50%. In addition, the scaling properties could be enhanced significantly by an elaborated parallel transpose. As a result, the GEM code can use now 32 times more processors than the original code version which was typically restricted to just 16 processor runs. Very good strong scaling property with a parallel efficiency from 64 to 512 processors of 89% has been proven. The new GEM code (GEMR) with enhanced geometrical functionality now opens up the possibility to run large physically relevant cases such as for the Joint European Torus (JET) and ITER.

GYGLES

The GYGLES (GYrokinetic Global LinEar Solver) code solves the linear gyrokinetic equation for electrons and one or more ion species in tokamak geometry in the whole plasma domain (full radius) within an electromagnetic model.

Enabling of the GYGLES code aimed at the efficient usage on several hundreds of processors instead of the typical 32 processors runs performed in the past. Such good scaling property is needed to make Magnetohydrodynamic (MHD) simulations in a full kinetic model possible.

Isabelle Dupays and Olivier Glorieux, IDRIS, Orsay, France

Life sciences

In its attempt to stimulate and ease the exploitation of the sheer supercomputing power by the Life Sciences community, DEISA has reached a significant new level. An application team within the Joint Research Activity JRA4 has guided DEISA's middleware group into the conception of a web based portal for new HPC users. These users would not spontaneously try to benefit from the supercomputer solutions even if their challenges were well beyond the capacities of their laboratory's facilities.

The DEISA Life Sciences portal conceals the complexity of the supercomputing Grid allowing a user to submit stereotypical jobs through a simple web form. This web site is currently in a pre-production phase. It will be available to all the community after a screening by the Scientific Evaluation Committee. At present, eight DEISA sites are integrated into the portal.

Three bioinformatics' applications available

In the first version of the portal there are three main bioinformatics' applications: BLAST, NAMD and RaxML. They are derived into a variety of services, one for each DEISA site, where they were installed. A regular user is only allowed to see and use a limited number of those services.

The interface is based on a classic web form through which a user selects all the available options of a software package, along with some basic job characteristics. Once the job form is filled out, a user may submit his job through a single click. The interface then provides a very simple page allowing the monitoring of the job and its results, where the process is seamless and hides the actual complexity. An efficient way for retrieving the output data has also been set up. The output data and the logs of the job are available for a determined amount of time.

Initial users very satisfied

The portal is at the moment in a final test phase and seven pilot projects were selected to test it. The first version of the portal received a very warm welcome from its test users from the Life Sciences community. Since the testing phase has been going well, the portal will soon enter the production phase. In the next phase, the portal will be extended by deploying new software and features for the interface, taking into account the needs and suggestions of the community. Some reviews of the portal:

"This access allows us to have a calculation capacity which we could not have had internally. To date, we are very satisfied with using the server –, their availability and their speed. They are very appropriate for our analyses. We are eager to see some evolution on the way calculations are submitted and the results downloaded on our computers. One point of paramount importance for us was the security and confidentiality of all data submitted to the Bioportal. As a private company and as an end-user of Bioportal, we are particularly confident and optimistic about the interest that such a tool brings to insilico studies of new therapeutic molecules."

Xavier Gallet, Bioinformatician,
Research Scientist Nautilus Biotech, Evry, France.

"The portal granted me access to a run power that allowed me to carry out more than 25 ns of simulation of my protein in a few days whereas it would have taken months with the computer system at my laboratory. [...] The NAMD job submission through downloading the configuration file is simple. [...] Once, the jobs through retrieving the results were conducted without any difficulties even for files over 250 MB. [...] I am very satisfied with the service offered and I hope to benefit from the portal again."

Dr Lorraine Brillet, CEA, Grenoble, France.

"The portal's conception seems to be effective and it is easy to use. Finally, the downloading of compressed file results as an option is a good asset and the transfer is rapid."

Dr Bruno Maigret,
CNRS Research Director, LORIA, Nancy, France.

There are three main bioinformatics' applications, BLAST, NAMD and RaxML, in the first version of the portal. Test users from the Life Sciences community have warmly welcomed this new service.

Roberto Tregnago and Enrico Ribaldone, Centro Ricerche FIAT (CRF), Italy

Industrial CFD and CAA applications

Within the Joint Research Activity JRA5, Centro Ricerche FIAT performed different Computational Aero Acoustic (CAA) activities on the DEISA research infrastructure. Results have been very promising. Exploiting advanced scientific computing on the state-of-the art supercomputing environment can provide significant improvements in the area of Computational Fluid Dynamics (CFD) and Computational Aero Acoustic (CAA) for industrial applications.

Figure 1: Fiat Grande Punto

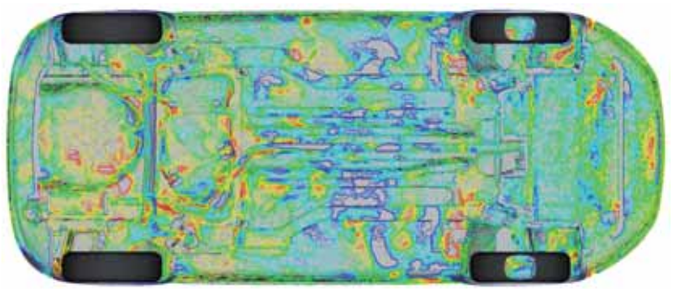

Figure 2: Contours of acoustic pressure on Fiat Grande Punto underbody.

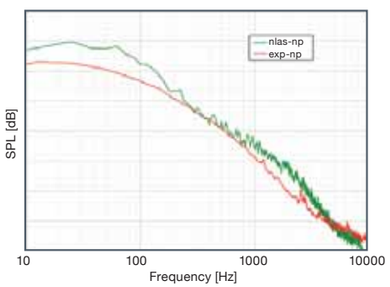

Figure 3: Comparison of numerical and experimental results for point 20.

Today the reduction of noise plays an important role for human well-being and health. Consequently noise generation has become environmentally significant issue for a wide range of applications, including the automotive industry.

Computational Fluid Dynamics (CFD) and Computational Aero Acoustic (CAA) have the potential to identify the fluid dynamics features that contribute to the generation of sound, and to guide the way to noise control techniques – and thereby to the reduction of noise radiation. The assessment of the aero acoustic properties in the new car's design can be largely anticipated through the use of CAA. Several layout alternatives can be tested numerically: a practice that is at the moment unfeasible in physical testing.

These new design processes allow automotive companies to anticipate the realization of design targets and to reduce design and experimental costs. However, at the same time, these new deisgn processes are highly demanding in computing power. The need of extreme computational resources renders a partnership with research infrastructures such as DEISA very important for automotive companies.

Streamlined design process and improved performance of the final product

Exploiting advanced scientific computing on a research infrastructure such as the DEISA supercomputing environment, may offer significant improvements in the area CFD and CAA for industrial applications.

Centro Ricerche FIAT's use experience with the DEISA infrastructure has demonstrated that appropriate numerical techniques, together with suitably structured high performance multiprocessor systems, do represent powerful tools for industry. They enable tackling problems that would be impossible to face with standard resources within the time constraints of a product development process. Thanks to the increasing performance of the supercomputers it is likely that we will see in the near future comprehensive simulations of these designs that encompass multidisciplinary optimisation aspects, too.

The competitiveness of the company has been increased through the achievement of a streamlined design process and the improved performance of the final product.

Raising temporal limit of industrial simulations' capabilities

The main objective of the partnership between Centro Ricerche FIAT and DEISA has been the raising of the temporal limit of the capabilities of industrial simulations and, in particular, the extension of actual and complex geometries of aeroacoustic simulations with Navier–Stokes schemes. Three different CAA activities, with growing complexity, were performed:

- Cavity simulation for tonal noise prediction;
- Simulation of car cabin air distribution ducts for broad band noise prediction;
- Shape noise of fully detailed car underbody simulation for broad band noise prediction, which will be extended to the whole car (both upper and under body in the same simulation).

The underbody of Fiat Grande Punto (Figure 1) was simulated.

The simulation was performed with a time step of 5.0×10^{-5} s in order to increase the accuracy of pressure waves' resolution. This corresponds to a Nyquist frequency of 10 000 Hz. The resolution of these high frequencies is needed in order to achieve results that can be considered a relevant breakthrough in the industry simulation capability and, above all, to make simulation a viable alternative to physical testing.

Physical time history was accumulated up to 0.6 seconds, with a global resource consumption of roughly 30 000 CPU hours (in DEISA normalized units), distributed on 64 and – above all – 128 CPUs.

The numerical pressure signal has been recorded at every time step, in 138 different positions on the underbody of the car, in order to obtain accurate noise maps, which show the noise power spectra over the whole geometry for a given frequency.

For the validation of these numerical results a number of experiments were performed in FIAT Auto's aeroacoustic wind tunnel. Twenty flat microphones were placed in the same positions of a subset of the numerical probes and numerical, and then experimental power spectra of pressure signal at single locations were compared. Some relevant results of the simulation are presented in the Figure 2.

Globally, comparison of the results of the simulation with the experimental results, show excellent agreement (Figure 3). The level of accuracy reached has driven design solutions for noise emission, which is a critical area of car design.

Gilles Grasseau, IDRIS-CNRS, Orsay, France

Coupled applications

The Joint Research Activity JRA6 offers the opportunity to research projects using cou-pled simulations to take advantage of HPC facilities, such as the DEISA infrastructure. As a result of this work, research teams have the possibility to tackle new classes of numerical investigations.

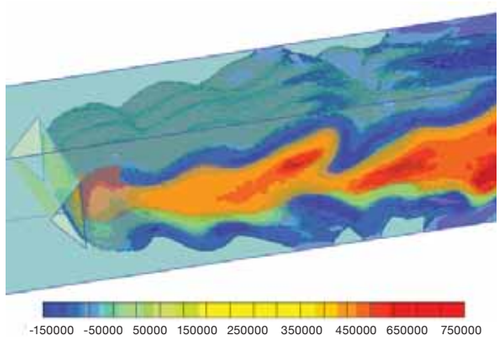

Figure 1: A 3D flame numerical simulation in a diedra flame-holder case (EM2C laboratory). Instantaneous field of the radiative power (in kWm^{-3}) is here shown. Positive (respectively negative) radiative powers are lost (absorbed) by gases. Up to 256 processors for the Radiative Heat Transfer and 120 processors for the Combustion part have been deployed on one site of the DEISA Infrastructure.

Many complex systems can be seen as made of components that obey their own physical laws, and these components interact weakly with one another through boundary conditions, for example. Code-coupling which deals with multiphysics, multimodels and multiscale numerical simulations, is especially well-suited for interdisciplinary research projects in which different legacy codes cooperate within the whole coupled application.

Generally, due to the complexity of these applications, coupled codes run with a moderate efficiency on parallel machines as available in the DEISA infrastructure. The major contribution of the JRA6 team has been to enhance the coupled applications in HPC context, such as DEISA infrastructure, and, as a result, provide research teams with the possibility to tackle new classes of numerical investigations.

The main scientific and technical results of the nine coupling projects driven by this JRA are presented in the following.

Combustion in 3D

Radiative heat transfer plays an important role in the turbulent combustion, but is often neglected in simulations because of its complexity and the related heavy numerical cost. DEISA Extreme Computing Initiative (DECI) allowed French researchers at the Energetique Moléculaire et Mascroscopique Combustion laboratory to validate the 2D results with a more realistic description, such as an LES, 3D radiative code, etc., and validate the coupling approach for large configurations.

The results obtained in the DECI context demonstrate the importance of radiation heat transfer in turbulent combustion (see Figure 1).

Deploying of a coupled architecture for aeroacoustic simulations

The research from the Institut für Aerodynamik und Gasdynamik in Stuttgart proposes a smart way to treat aeroacoustic problems by applying different Navier–Stokes numerical models with different refinement according to the domain complexity topology and the obstacle proximity (see Figure 2).

The JRA6 contribution has mainly consisted of greatly enhancing the parallelism level and thus opening access to totally new problem classes. Furthermore, the benefit of deploying the application on heterogeneous configurations (vector and scalar architecture) has been demonstrated.

>>>

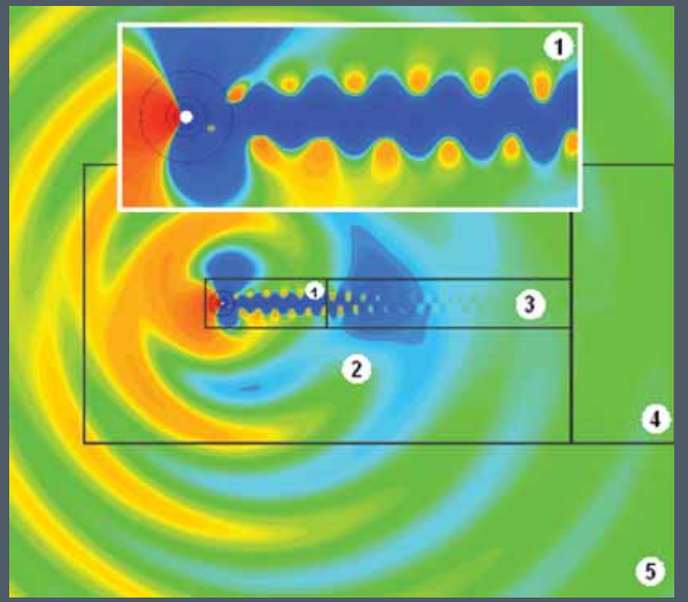

Figure 2: Visualisation of the pressure field generated by a von Karman Vortex Street decomposed into (1) unstructured Navier–Stokes (overlapped circles represent different mesh resolutions), (2) structured nonlinear Euler, (3) structured Navier–Stokes, (4) structured nonlinear Euler and (5) structured linearized Euler.

Coupled applications (continued)

High added value to the design of inhibitors

In order to design better inhibitors involved in many diseases, such as Alzheimer or AIDS, the Chemistry Department at Ecole Normale Supérieure in Lyon is studying the whole enzymatic process using a coupling approach: The reactive centre is described using a quantum method (CPMD), whereas the rest of the solvated protein is described by a classical molecular dynamics method (GROMACS). The JRA6 team has brought high added value to this research by enhancing the initial coupled application (CPMD/GROMACS) to explore much bigger molecular configurations.

Driving molecular dynamics simulations

Large-scale computer simulations of biological systems provide valuable insight into molecular processes as diverse as enzymatic catalysis and membrane fusion. By adding external forces (with the graphic interface), the simulation systems can be driven towards states of particular interest (see Figure 3), which are unlikely to be observed spontaneously. Furthermore, the mechanical properties of (macro)molecules and their assemblies can be examined in this way by researchers of the Theoretical biochemistry laboratory at Institut de Biologie Physico-Chimique in Paris.

Driving GROMACS MD simulations with the well-known sophisticate visualization tool VMD has been the first step of our work in the HPC context.

Fluid-structure coupling for medical applications

Some of the modern medical numerical investigations deal with simulating blood flow in veins, including the interactions on the lung assisted by artificial ventilation. Due to the complexity of the considered biological models, such simulations lead by Lehrstuhl für Numerische Mathematik in Munich, require taking into consideration of a wide range of different physical properties and more and more details.

In order to exploit efficiently all these components, the JRA6 activity tackled global optimization problems and, in particular, the component mapping onto the heterogeneous DEISA infrastructure.

Coupled application between 3D natural convection and radiation

The JRA6 contribution, in collaboration with Laboratoire d'Informatique pour la Mécanique et les Sciences de l'Ingénieur in Orsay, enabled deploying the numerical coupled application between 3D natural convection and radiation for the thermal efficiency of buildings. The numerical tools have been optimised by providing two parallelism levels which offers the possibility to simulate cavities with high Rayleigh numbers. The middleware capability for heterogeneous configuration (vector architecture and scalar architecture) over the DEISA infrastructure has been validated.

Combustion in 2D

Combustion, radiation heat transfers and pollutant formation are the three main physical phenomena involved in this coupling project. Numerical investigations driven by Energetique Moléculaire et Macroscopique Combustion laboratory and made possible by the technical work achieved in JRA6, demonstrate that taking into account the radiative heat transfers not only modify the instantaneous and mean temperature fields as expected, but also the flame structure itself and its dynamics, which was little expected.

Investigations of water cycles over dry and wet regions

Exploiting a multimodels architecture that couples three components: atmosphere, vegetation and hydrologic processes, permits very innovative and relevant investigations lead researchers at Laboratoire d'Etudes des Transferts en Hydrologie et Environnement in Grenoble about the water cycle over dry regions as well wet ones. Switching from dry to wet hydrological components offered the possibility to collaborate with researchers of AMMA – the African Multidisciplinary Monsoon Analysis international project. These results are now part of the international AMMA project.

Observational predictions on the very first objects in the universe

Dark matter evolution, baryonic gas dynamics and chemical reactions inside the baryonic gas are the different physical phenomena involved and coupled together in this project propose by Laboratoire de l'Univers et de ses Théories in Meudon, France.

In order to exploit large physical configurations, substantial enhancements of the coupled application (parallelism level, load balancing and general optimisations), have been achieved to run efficiently on the DEISA infrastructure. This numeric tool, which considers the dark energy effects, provides precious observational predictions about the very first objects in the universe.

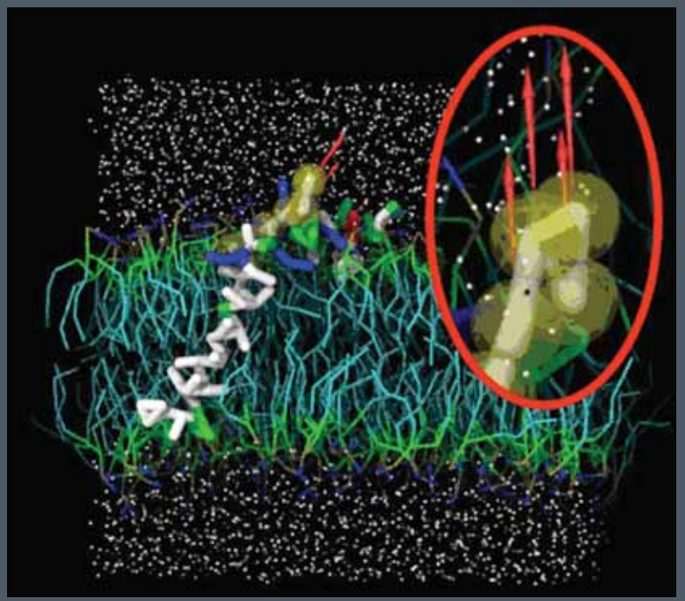

Figure 3: A single snapshot of a coarse-grained Molecular Dynamics simulation of a protein (Synaptobrevin transmembrane domain) in a lipid bilayer. The protein is shown as licorice, lipids as lines and water as points. The red arrows (see figure zoom) show the user external forces applied interactively in the simulation. Here, the membrane anchoring of protein is probed by pulling on one particular side (on tryptophane residues).

Terry Sloan, EPCC, Edinburgh, UK

Access to resources in heterogeneous environments

A joint research activity (JRA7) within DEISA has used emerging modern Grid standards to develop a uniform means for accessing and exploiting a heterogeneous supercomputing Grid environment. The DEISA Services for the Heterogeneous management Layer (DESHL) has proven popular with researchers. Scientific users are now able to build middleware and hardware independent task farms, workflows and code migration/resubmission suites.

Figure 1: The DESHL installation wizard.

Since the dawn of simulation researchers have wanted to exploit computational resources without the overhead of shoe-horning their hypotheses into buggy, unintuitive computer code. The fragmented nature of the computational landscape into different architectures, operating systems and now grid middleware has further frustrated those researchers who would rather be doing computationally enabled science than computer science.

A small step on the way to fulfilling this desire has been the standardisation efforts of the Open Grid Forum. Here various bodies have been actively working on agreeing the standards that can form the software foundations for easier utilisation of computational resources independent of architecture, operating system and middleware. If these standards are widely adopted then it is hoped that this will help in achieving the long-held dream of access to computation being as ubiquitous as access to electricity.

Uniform access regardless of the underlying hardware and software

The DEISA JRA7 activity has been using these emerging modern Grid standards to develop a uniform means for accessing and exploiting a heterogeneous supercomputing Grid environment. To this end, the DEISA Services for the Heterogeneous management Layer, better known as the DESHL, has been developed by JRA7. EPCC and ECMWF from the UK, FZJ from Germany and CINECA from Italy are the participants in this research activity.

DESHL allows users and their applications to manage batch jobs and data across a computing Grid in a uniform manner regardless of the underlying hardware and software on the various nodes of the Grid. The DESHL employs emerging Grid standards (SAGA – Simple API for Grid Applications and JSDL – Job Submission Description Language) as well as the Open Group Batch Environment Services specification so that a user and their applications are not affected by the differences or changes in the hardware and software configuration on Grid nodes.

Reliance on well-established standards also helps ensure future interoperability with other Grid environments. Figure 2 illustrates how the DESHL can be used to access the DEISA heterogeneous supercomputing Grid infrastructure.

Command line and scripting capabilities proven popular with researchers

The DESHL provides users and their applications with a command line tool and application programming interfaces for job and data management in a UNICORE-based Grid such as the DEISA heterogeneous supercomputing Grid infrastructure.

The DESHL also includes a GUI-based installer (see Figure 1), an installation and user manual, a design description and API documentation.

On the DEISA infrastructure, the DESHL command line and scripting capabilities have proven popular with researchers. These capabilities have allowed scientific users to build middleware and hardware independent task farms, workflows and code migration/resubmission suites. In addition, the prototype DEISA Life Sciences portal utilizes the DESHL SAGA API thus demonstrating that this API can also be successfully used by third party application developers.

Away from DEISA, the DESHL has been used by the German D-Grid and the European Immunogrid project where it has been used in a prototype simulator capable of submitting jobs using JSDL to both OMII and UNICORE-based environments. At Supercomputing 2007 as part of the Open Grid Forum's Grid Interoperability Now demonstrations, the DESHL was employed in the DEISA-GridAustralia interoperation demonstration on drug simulation for HIV.

Important contribution to global Grid computing

Finally, it is worth noting that the DESHL was the first publicly-available Java-based implementation of file and job management using the SAGA Grid standard. More significantly, it was the first SAGA implementation to be actively deployed, tested against and used by scientific researchers within a production-level, continental Grid infrastructure and thus is an important contribution to global Grid computing.

The DESHL is available from **http://deisa-jra7.forge.nesc.ac.uk/.**

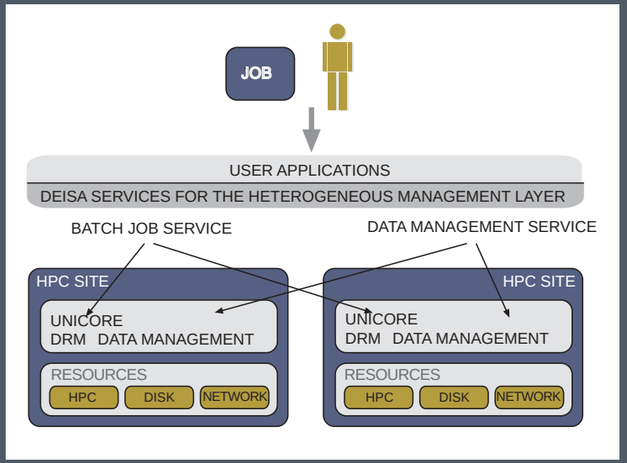

Figure 2: Using the DESHL: At a local site a user wants to run a job on the DEISA heterogeneous environment. Standards-based interface allows user applications access to heterogeneous sites.

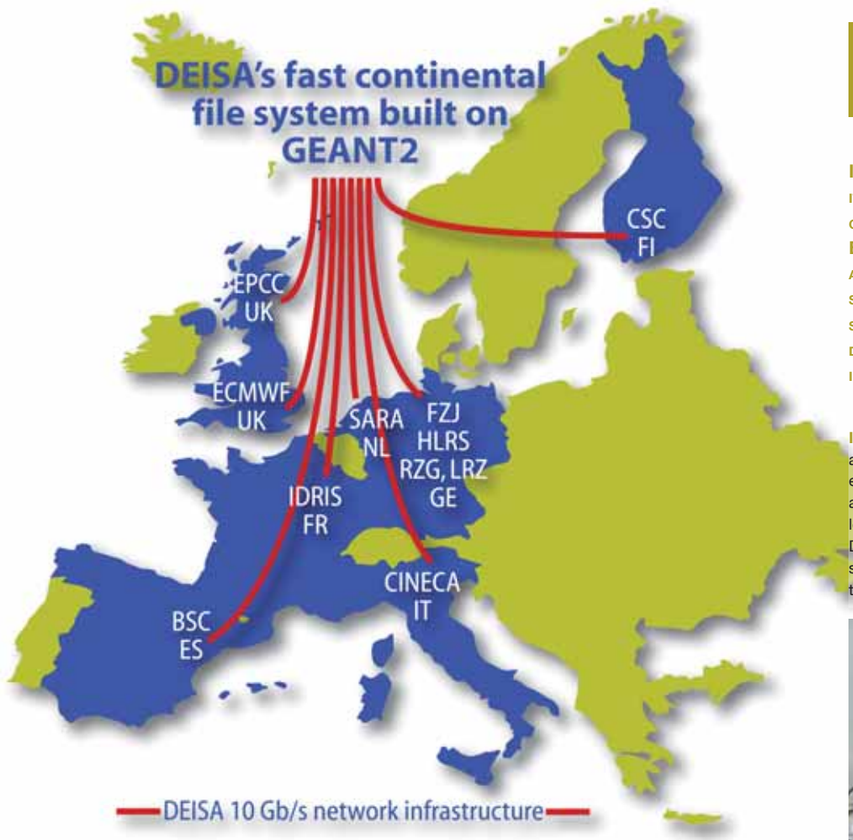

Figure 1: Eleven leading European Supercomputing Centres from seven European countries have coordinated their actions and combined their competences and jointly built-up and put into production state the Distributed European Infrastructure for Supercomputing Applications – extending from Italy and Spain through Central and Western Europe to the UK and Finland – connected by a priority high speed network empowered by GEANT /GEANT2 and the NRENs.

Outlook

In spring 2002, the idea has manifested itself to overcome the fragmentation of supercomputing resources in Europe both in terms of system availability and in the necessary skills for efficient supercomputing support, and the establishment of a distributed European supercomputing infrastructure was proposed.

In May 2004 the DEISA project was started as a EU FP6 Integrated Infrastructure Initiative by eight leading European supercomputing centres and expanded in 2006 by three additional leading centres. Through the joint efforts, DEISA reached production quality soon after to support leading edge capability computing for the European scientific community.

Figure 2: DEISA organises training sessions to enable fast development of user skills and know-how needed for the efficient utilisation of the DEISA infrastructure.

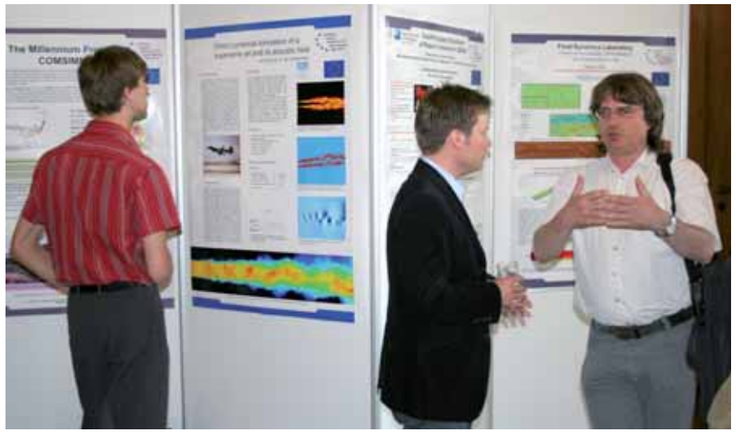

Figure 3: DEISA Symposium is organised annually to provide a forum where scientists from around the world can discuss HPC-eInfrastructures in general and for DEISA users to share their experiences and results.

DEISA has also contributed to raising awareness of the need for a persistent European HPC infrastructure. The huge continental demand for integrated capability computing resources in Europe has been documented by the waves of proposals in response to the annual DECI calls. The importance of a persistent European HPC ecosystem was recognised in the ESFRI, resulting in very concrete recommendations in the 2006 report.

In line with this ESFRI roadmap, the DEISA Consortium will continue to support and further develop the distributed high performance computing infrastructure and its services in EU FP7 through the DEISA2 project. Activities and services relevant for Applications Enabling, Operation, and Technologies will be continued and further enhanced, as these are indispensable for the effective support of computational sciences in the HPC area. The service provisioning model will be extended from one that supports single projects to one supporting Virtual European Communities. Collaborative activities will be carried out with new European and other international initiatives.

Of strategic importance is the cooperation with the PRACE project which is preparing for the installation of a limited number of leadership-class Tier-0 supercomputers in Europe.

Further important projects to cooperate with include Europe's GEANT2, EGEE, and EGI, as well as others world-wide.

The key role and aim will be to deliver a turnkey operational solution for a future persistent European HPC ecosystem, as suggested by ESFRI. The ecosystem will integrate national Tier-1 centres and the new Tier-0 centres.

Editorial

DEISA ADVANCING SCIENCE IN EUROPE

www.deisa.org

ISBN 978-952-5520-32-3

Editor-in-Chief
Saara Värttö, CSC

Copy editors
Juha Fagerholm, CSC
Petri Nikunen, CSC
Kirsti Turtiainen, CSC
Leena Jukka, CSC
Vera Hansper, CSC
Gavin J. Pringle, EPCC

Editorial board
Victor Alessandrini, IDRIS-CNRS
Hermann Lederer, RZG
Achim Streit, FZJ
Saara Värttö, CSC
Gavin J. Pringle, EPCC

Graphic Design and layout
Up-to-Point Oy

Printed by
TT-Urex Oy

DEISA – Distributed European Infrastructure for Supercomputing Applications is a consortium of leading national supercomputing centres that currently deploys and operates a persistent, production quality, distributed supercomputing environment with continental scope. The purpose of this FP6 funded research infrastructure is to enable scientific discovery across a broad spectrum of science and technology, by enhancing and reinforcing European capabilities in the area of high performance computing. This becomes possible through a deep integration of existing national high-end platforms, tightly coupled by a dedicated network and supported by innovative system and grid software. The DEISA project receives funding from the European Community's Sixth Framework Programme (FP6/2002-2006) under the grant agreement n° RI-508830 and n° RI-031513.

DEISA

DEISA IS A CONSORTIUM OF LEADING NATIONAL SUPERCOMPUTING CENTRES THAT CURRENTLY DEPLOYS AND OPERATES A PERSISTENT, PRODUCTION QUALITY, DISTRIBUTED SUPERCOMPUTING ENVIRONMENT WITH CONTINENTAL SCOPE.

The purpose of this FP6 funded research infrastructure is to enable scientific discovery across a broad spectrum of science and technology, by enhancing and reinforcing European capabilities in the area of high performance computing.

Grid technologies are used to strongly integrate national supercomputing platforms, and to provide to scientific users transparent access to a European pool of computing resources. The joint and coordinated operation of this environment is tailored to provide enhanced computing power and resources to end users, and to enable new, ground breaking research activities in science and technology.

DEISA operates as a virtual European supercomputing centre. Human competences are also pooled, to provide first class, substantial added value services to computational sciences.

A number of research activities are supported, aimed at reinforcing the impact of the infrastructure on leading applications and facilitating the incorporation of new, emerging Grid technologies.

WWW.DEISA.ORG

The DEISA Consortium receives funding from the European Community's Sixth Framework Programme (FP6/2002-2006) under the grant agreement n° RI-508830 and n° RI-031513.