DEISA

DEISA DIGEST 2008

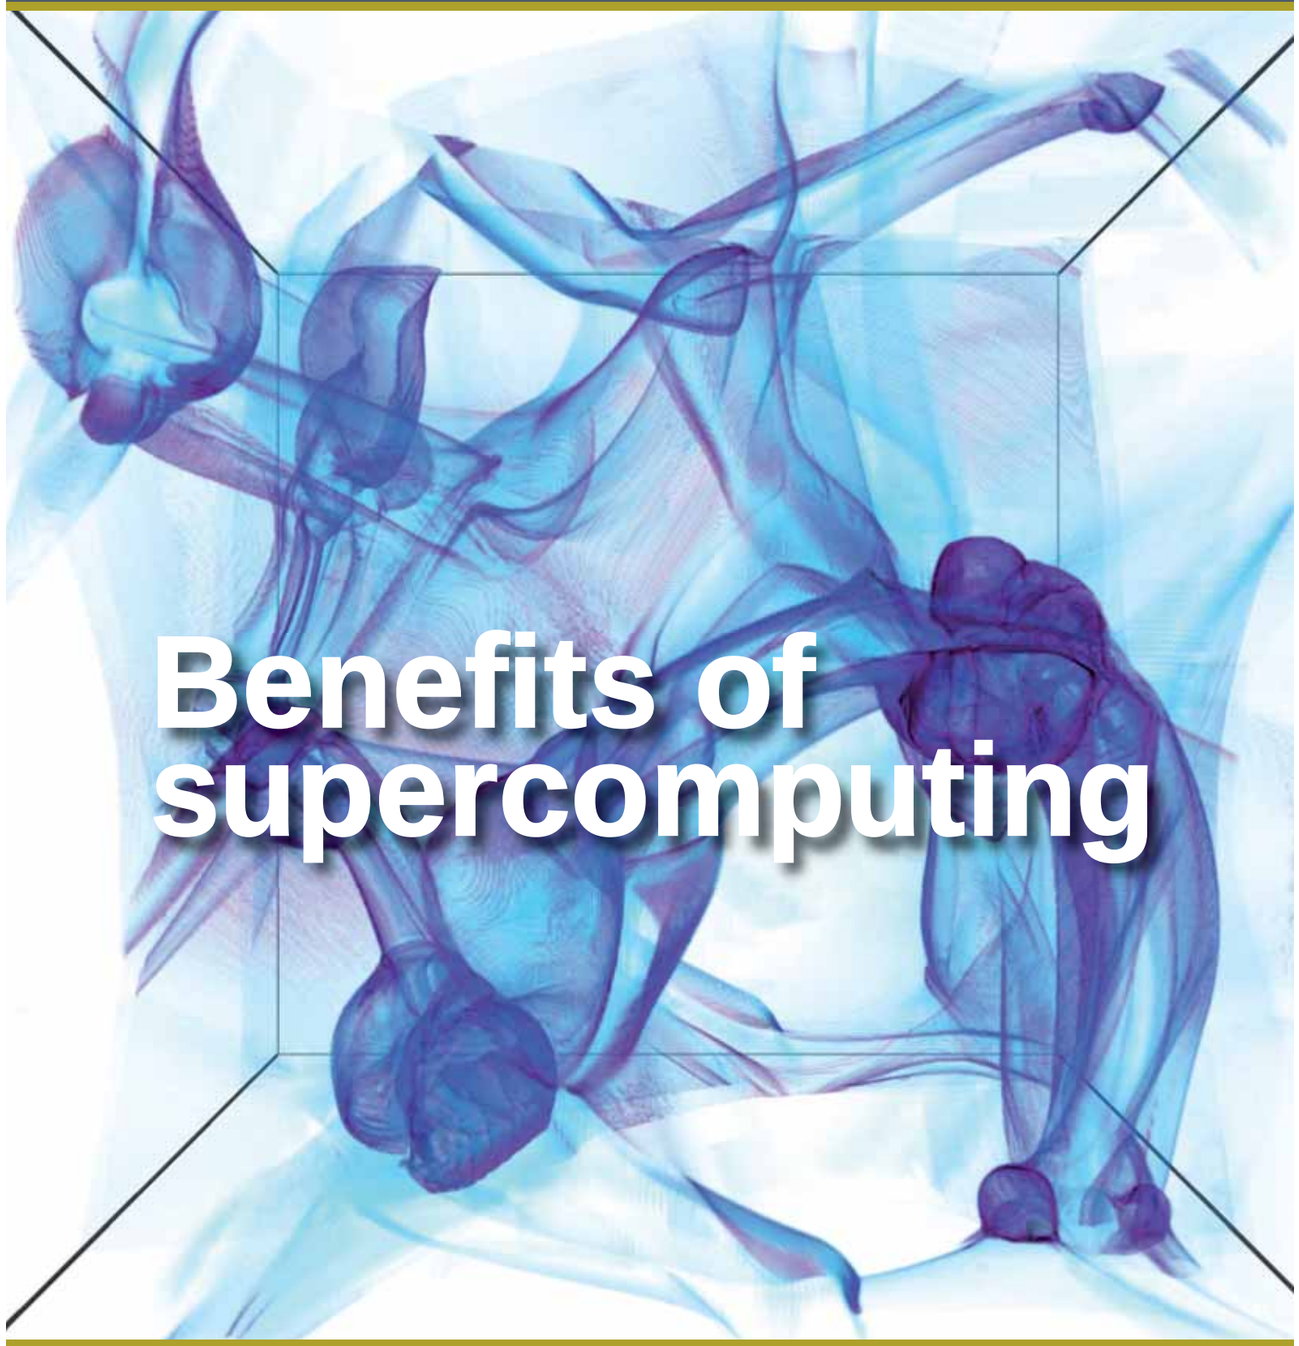

Benefits of supercomputing

DEISA DIGEST
www.deisa.eu

ISBN 978-952-5520-33-0

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 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 Consortium receives funding from the European Community's Sixth Framework Programme (FP6) under the grant agreements RI-508830 and RI-031513, and from the Seventh Framework Programme (FP7) under the grant agreement RI-222919.

Distributed European Infrastructure for Supercomputing Applications

e-infrastructure

CONTENTS

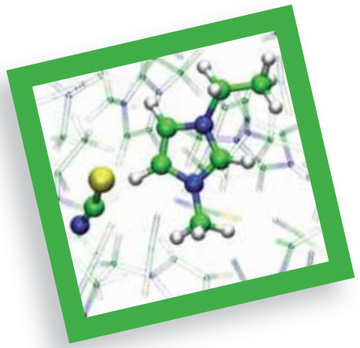

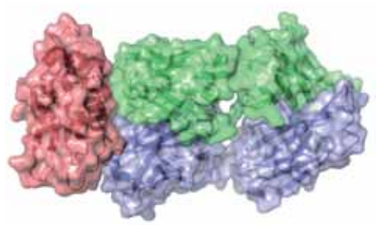

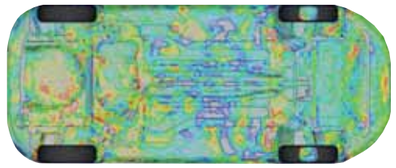

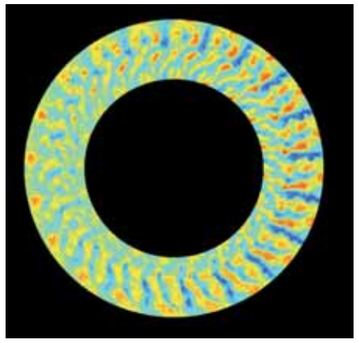

"DEISA provides supercomputing resources and an environment for services. It facilitates access to HPC for European scientists as well as supporting DEISA administrative functions. European scientists and technologists are provided world-class leadership supercomputers with capabilities equal to those available in the USA and Japan", says Professor Dr. Victor Alessandrini, co-ordinator of the DEISA consortium.

Benefits of supercomputing

Christopher Lazou

THE TERM "SUPERCOMPUTER" IS ATTRIBUTED TO A FEW OF THE MOST POWERFUL SCIENTIFIC COMPUTERS AVAILABLE AT A GIVEN TIME. IT IS IMPOSSIBLE TO SAY THAT ONE COMPUTER MODEL IS THE MOST POWERFUL BECAUSE LIKE EVERY COMPLEX SYSTEM THE POWER OF A SUPERCOMPUTER IS NOT LINEAR.

In the last fifty years, Control Data, Cray, IBM, Fujitsu and NEC are examples of vendors who built supercomputers. Often the different architectures of each type of computer are a significant factor in determining which supercomputer is most suitable for a particular application.

Supercomputers are indispensable tools for solving the most challenging and complex scientific and engineering problems including the simulation and modelling of physical phenomena. Their use is of strategic importance to national security and defence, as well as for research and development in many areas of science and engineering. The key virtue of supercomputers is that they can advance knowledge and generate insight that would not otherwise be possible or that could not be captured in time to be actionable. >>>

For example, although simulating the path of say a hurricane after the event and devastation it caused would be worthless, simulating its path days before, would have enormous benefits as it would allow mitigation actions to be initiated and hence minimize damage.

Supercomputer simulations can augment or replace experimentation in cases where experiments are hazardous, expensive or even impossible to perform or to instrument. They can even enable virtual experiments with imaginary worlds to test theories beyond the range of observable events. Further, supercomputers have the potential to suggest entirely novel experiments that can revolutionize our perspective of the world. They enable faster evaluation of design alternatives, thus improving the quality of engineered products essential for industrial prowess and economic competitiveness.

Most important technological activities use supercomputers in fundamental and essential ways. These include physical phenomena such as climate change, energy, environment, weapons, the analysis of data such as national security intelligence, the behaviour of molecules in life sciences, protein folding, genome sequencing, medical treatments, astronomical observations and the intricate design of engineered products such as automotives, aerospace, nano-technology, and new material designs using up to 10 000 atoms. Integrated solutions also need to handle the enormous growth of data, in some cases of over 10 Petabytes size.

As computer power increased and computer use broadened many interesting applications could be ran on capacity cluster computers, as they require only modest amounts of computing. Yet new problems have arisen whose computational demands for scaling and timeliness are so great they exceed the capabilities of our current supercomputers. Many of these problems are fundamental to the government's ability to address important national issues.

In its roadmap published in 2006, the European Strategy Forum on Research Infrastructures (ESFRI) has identified High Performance Computing (HPC) as a strategic priority for Europe. Put simply, supercomputers are of the highest and most pervasive strategic importance. They enable scientists and engineers to solve today's problems and develop the new technology for tomorrow's industry, affecting national employment patterns and national wealth. The success of a new motorcar design, a new aircraft design or a new pharmaceutical process and so on, will determine whether European companies remain competitive in the market place. In short, European scientists and engineers need to have access to world-class supercomputers for their research if they are to have a positive impact on European competitiveness.

In 2002, the Japanese installed the fastest supercomputer at that time named Earth Simulator for study of earth sciences. This feat had a major impact on HPC, similar to that Sputnik had on the space race in the late 1950's. The DARPA productivity project and the hardware Petaflops/s programme (a Petaflop/s is equivalent to a million trillion calculations every second) in the USA are inspired by it. DARPA is funding R&D for IBM and Cray to develop a Petaflops/s system by 2010, with the aim of achieving supercomputing supremacy. The Japanese government initiated their own Petaflops/s hardware and software supercomputer programme, spearheaded by RIKEN, making supercomputer developments a top priority for strategic R&D, culminating in a system delivering 1 Petaflop/s sustained performance by year 2011.

There are still many challenges to be overcome in building Petascale supercomputers, including power consumption and robustness. However, technology trends and known near products in the pipeline, both in the USA and Japan, aspire to deliver Petaflops/s supercomputers within the above timeframe, costing several hundred Million Euros each. These Petascale supercomputers will be highly parallel with a heterogeneous architecture and partly application specific. They will contain components, for fast serial performance, MIMD data level parallelism (e.g. vector), a fine-grained MIMD parallelism (e.g. multi-threading), as well as regular and sparse bandwidth of varying intensities. The overall objective is delivering a significant increase in performance.

It is a known truism that for ease of programming global shared memory is unbeatable. It provides the lowest latency communication and lowest overhead communication. It enables fine-grained overlap of computation and communication and tolerates latency with processor concurrency. In contrast, message passing concurrency is constraining and hard to program. Vectors provide concurrency within a thread, multi-threading provides concurrency between threads. The challenge is to exploit locality, to reduce bandwidth demand. This is achieved by using hierarchical processor architectures, to enhance temporal locality and lightweight thread migration, to exploit spatial locality. Other techniques to reduce network traffic, such as atomic memory operations and single word network transfers when no locality is present, are also used. >>>

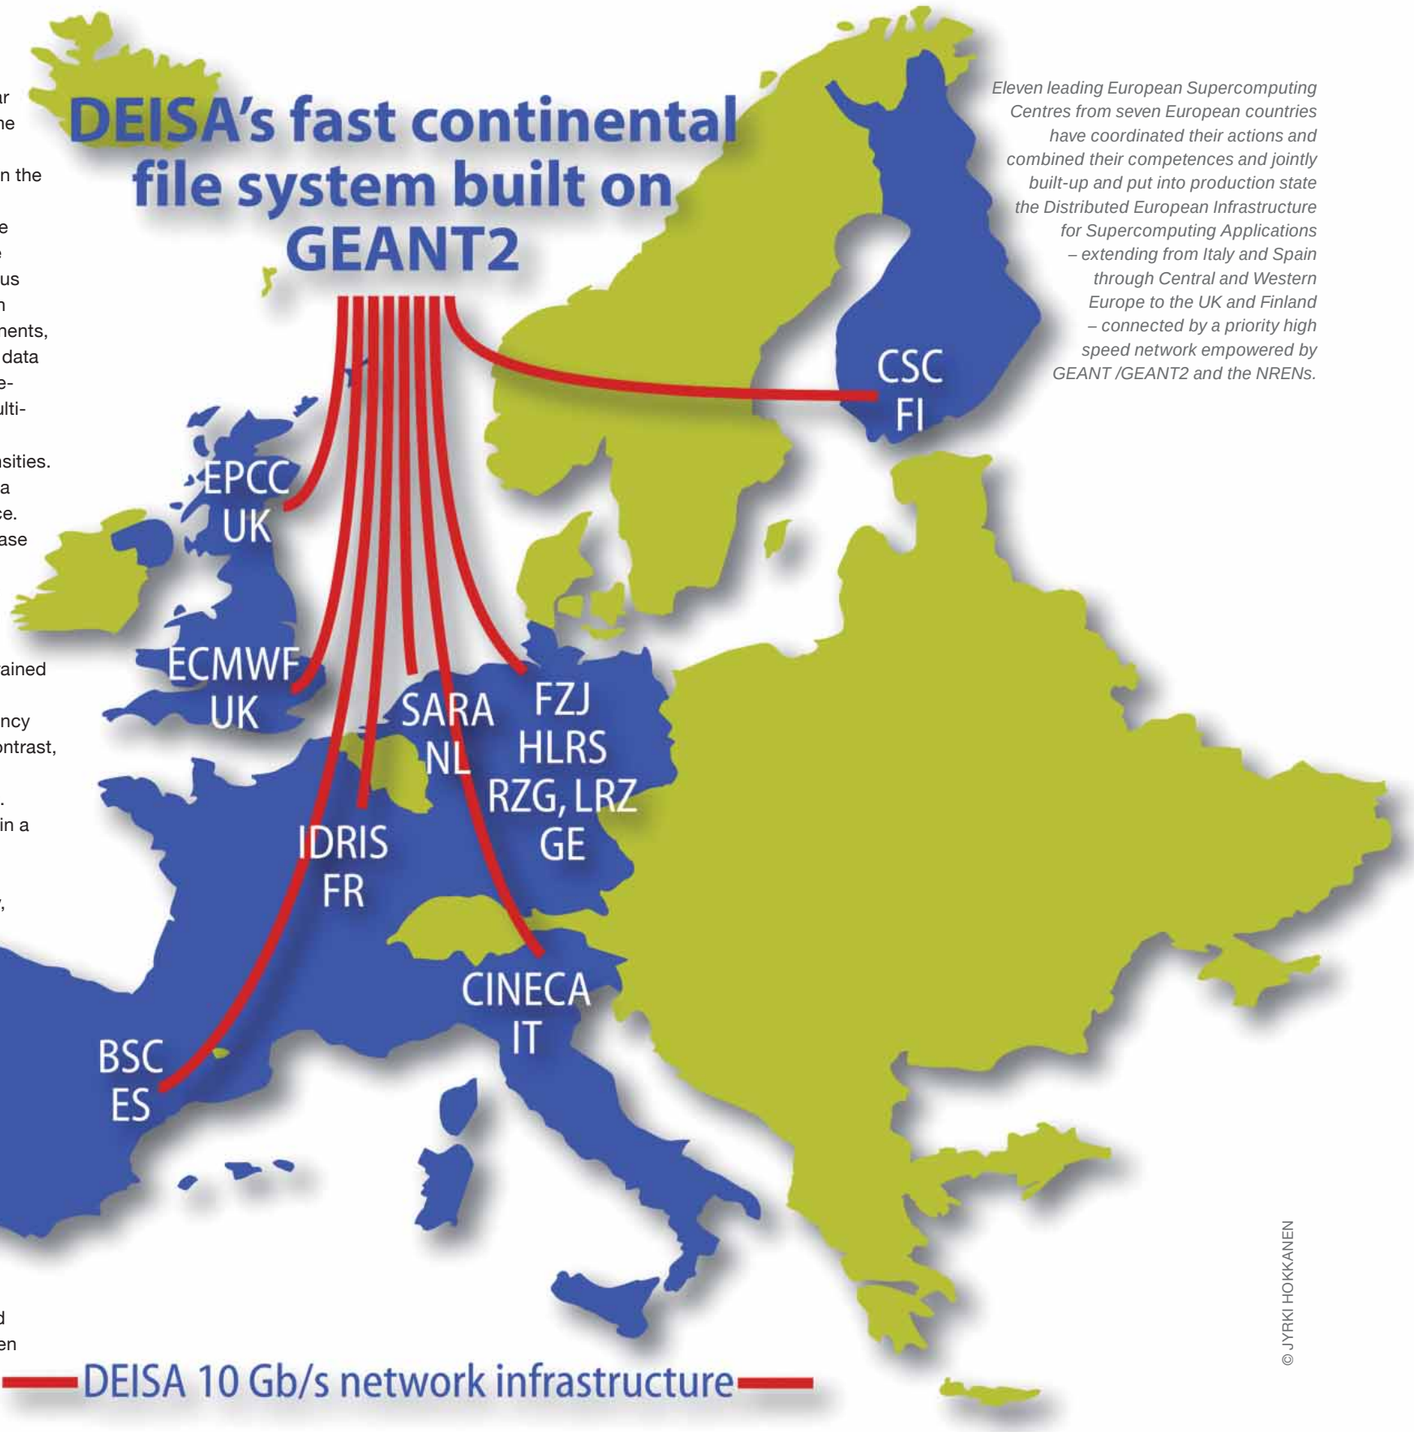

Eleven leading European Supercomputing Centres from seven European countries have coordinated their actions and combined their competences and jointly built-up and put into production state the Distributed European Infrastructure for Supercomputing Applications – extending from Italy and Spain through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

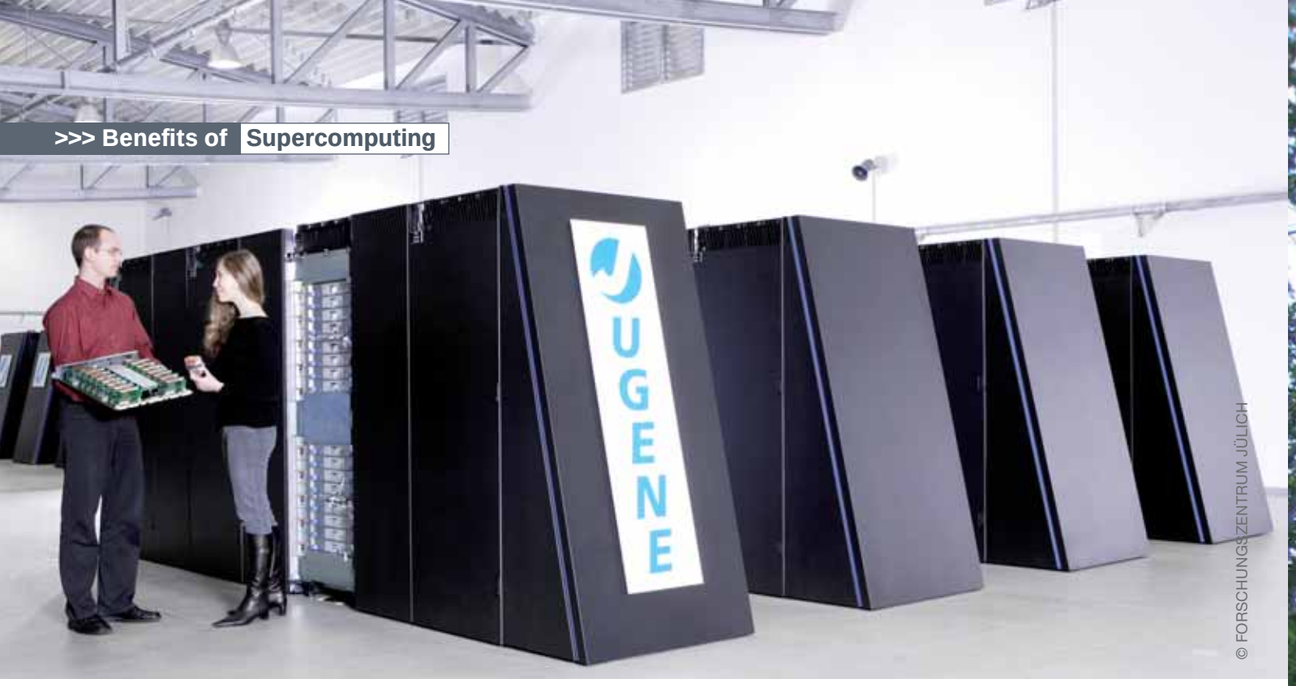

© FORSCHUNGSZENTRUM JÜLICH

Europe's fastest supercomputer JUGENE, located at the Forschungszentrum Jülich – Jülich Supercomputing Centre, Germany and part of the DEISA infrastructure, performs around 167 trillion mathematical calculations per second. In the current global ranking list, the TOP500, it holds the second place.

Multi-core chip technologies and application specific accelerators in the new supercomputers will make it necessary to introduce new programming models and algorithmic support on how to address the technical issues for extreme scaling: For example, existing codes will need to scale to 100 000 processors, or even more.

For access to these eclectic powerful supercomputers the USA has initiatives on infrastructure (TeraGrid) funded by the National Science Foundation (NSF) under the cyber-infrastructure programme and Japan's equivalent is called NAREGI.

Europe has no significant HPC hardware industry, but it has created European research infrastructures for nurturing European research. These include the Distributed European Infrastructure for Supercomputing Applications (DEISA), European Grid Initiative (EGI), HPC-Europa, Open Middleware Infrastructure Institute-Europe (OMII-Europe).

DEISA is a consortium of eleven leading national supercomputing centres that currently deploy and operate a persistent, production-quality, distributed supercomputing environment, within Europe. The purpose of this research infrastructure funded by the EU sixth Framework Programme (FP6) is to enable scientific discovery across a broad spectrum of science and technology by enhancing and reinforcing European capabilities in the area of HPC. 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.

DECI is the acronym of the DEISA Extreme Computing Initiative project. The objective of DECI is to identify, enable and run a limited number of ground-breaking applications in different fields of science and technology. These applications should consist of complex, demanding and innovative simulations that are impossible to implement without the DEISA infrastructure.

Over 50 projects already benefited from the 2005 and 2006 DECI proposal calls. From the DECI call 2007, 45 projects were accepted for operation using DEISA in 2008. These projects were awarded around 30 million processor-core hours in total. An allocation of 1 million processor-core hours could in principle be spent by using 16 processor-cores requiring about 62 000 hours (or 7 years) or by using 1024 processor-cores taking about 1000 hours (or 40 days) in real time. Since DECI is dedicated to capability computing, using large fractions of a big supercomputer (or several supercomputers) simultaneously is mandatory, and only such capability computing projects were considered for execution using the DEISA infrastructure.

The 45 projects accepted for operation in 2008 cover major areas of science including materials science (12 projects), astrophysics sciences (8 projects), engineering (8 projects), life sciences (8 projects), earth sciences (4 projects), plasma physics (3 projects), and informatics (2 projects). The projects to be supported involve scientists from 14 different European countries and collaborators from three more continents "DEISA provides supercomputing resources and an environment for services. It facilitates access to HPC for European scientists as well as supporting DEISA administrative functions. European scientists and technologists are provided world-class leadership supercomputers with capabilities equal to those available in the USA and Japan. The established DEISA infrastructure is a good springboard for the pivotal European HPC project, Partnership for Advanced Computing in Europe (PRACE), which aims to establish a permanent European research infrastructure as a single legal entity, including governance, funding, procurement, and usage strategies, in 2010", says Professor **Dr. Victor Alessandrini**, director of IDRIS-CNRS, France, co-ordinator of the DEISA consortium. ■

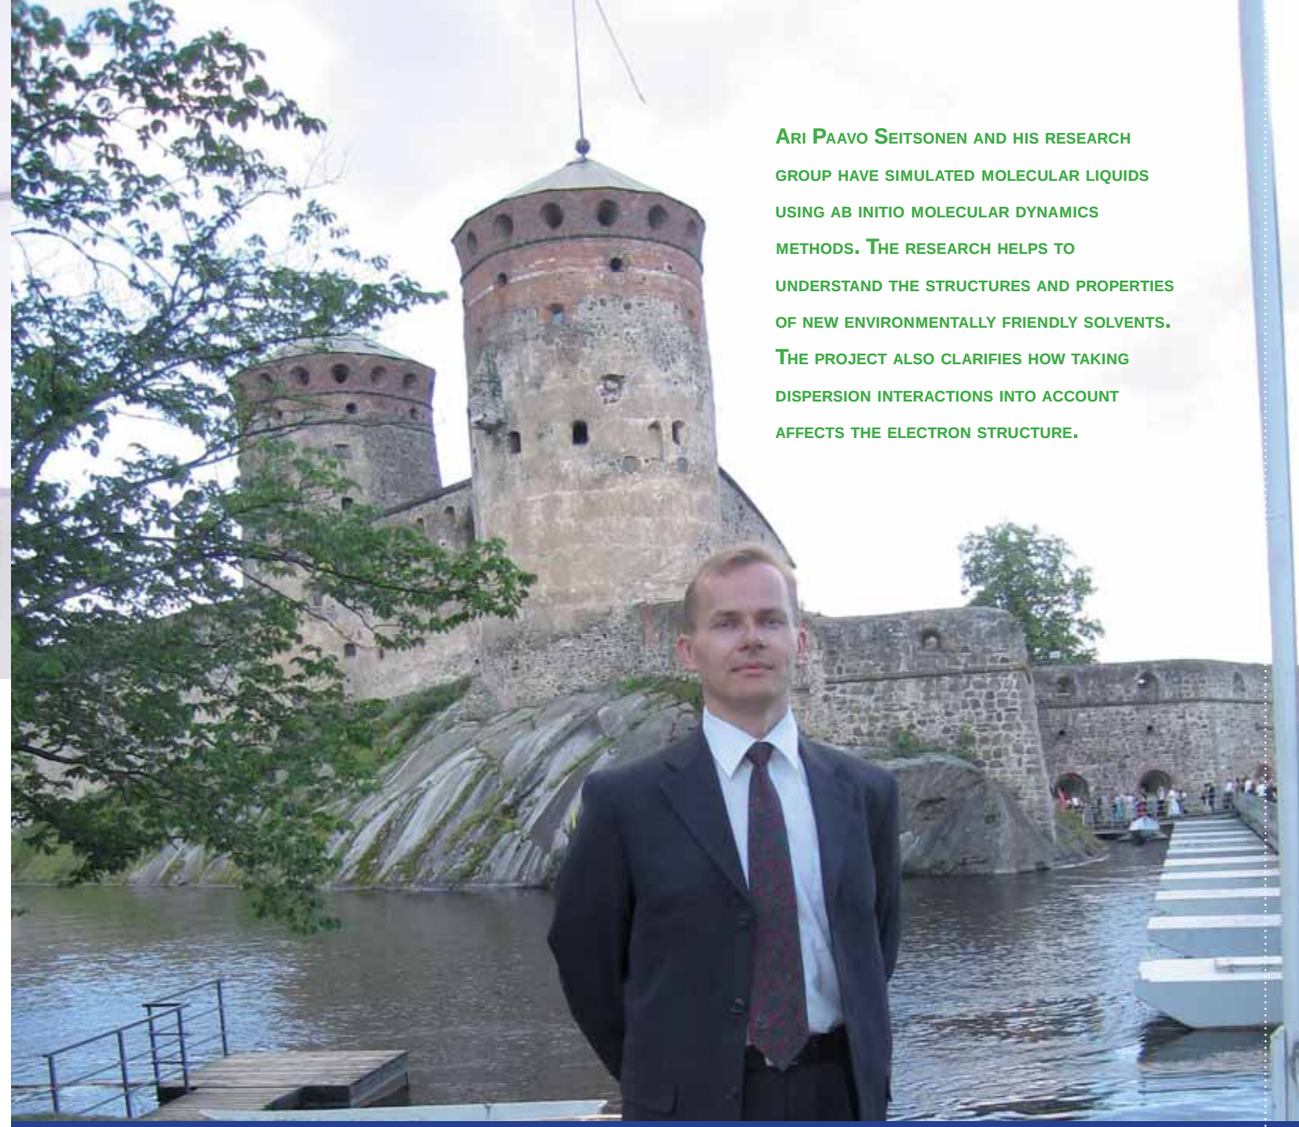

Ari Paavo Seitsonen and his research group have simulated molecular liquids using ab initio molecular dynamics methods. The research helps to understand the structures and properties of new environmentally friendly solvents. The project also clarifies how taking dispersion interactions into account affects the electron structure.

DEISA deployed **for advancing research into behaviour of liquids** >>>

Irja Nurmi-Rättö, Atte Sillanpää

Ari Paavo Seitsonen, research engineer from the Institute of Mineralogy and Physics at Universitè Pierre et Marie Curie

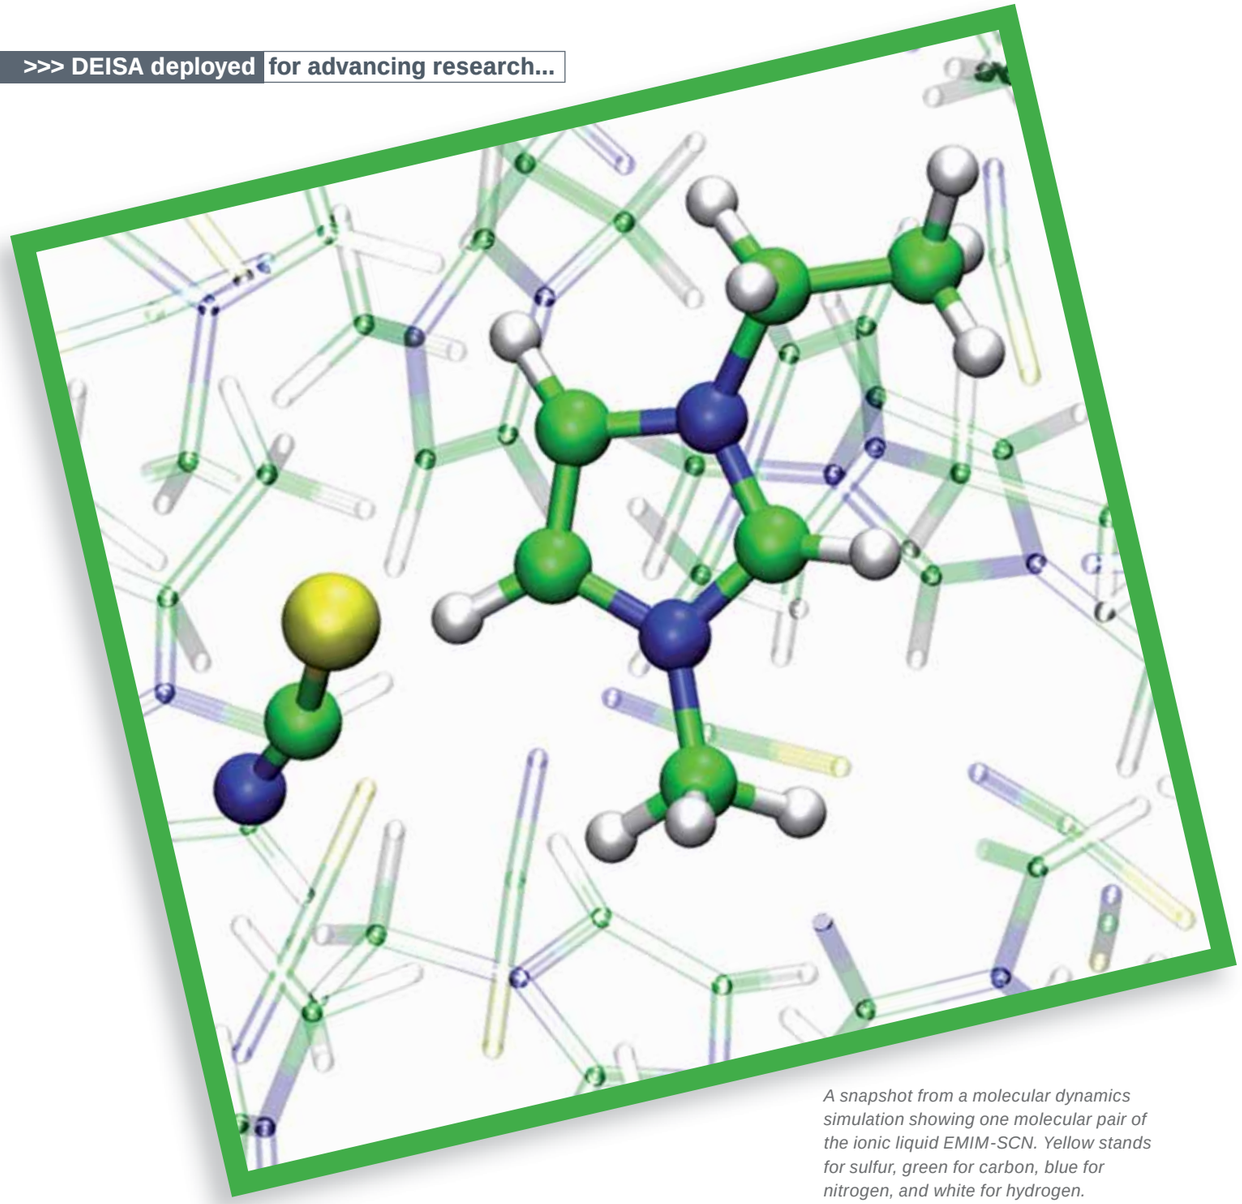

A snapshot from a molecular dynamics simulation showing one molecular pair of the ionic liquid EMIM-SCN. Yellow stands for sulfur, green for carbon, blue for nitrogen, and white for hydrogen.

The computing resources of the EPCC supercomputing centre in Scotland were used when the AIMDliqD (Ab Initio Molecular Dynamics for Liquids and Diamond) project was implemented as part of extensive continuing international collaboration. Molecular dynamics methods were used for simulating two liquids at the atomic level to provide a realistic finite-temperature description.

"In ab initio methods the system's electron structure is explicitly solved. Being a dynamic simulation, describing the structure requires solving the electronic structure consecutively many times during the dynamic simulation. This substantially increases the required computation capacity compared with, for example, simulations performed using parameterized potential energy models. However, the attained results are more reliable and usually more accurate", says **Ari Paavo Seitsonen**, research engineer from the Institute of Mineralogy and Physics at Universitè Pierre et Marie Curie.

Ab initio-based simulation produces the system's electron structure, while traditional parameterized models only provide positions for the nuclei and the system's total energy. Electronic structure analysis is also useful for investigating other properties of the material and reasons behind them. These properties cannot be investigated using classical or semiempirical methods. Ab initio molecular dynamics (AIMD) methods reveal information about the target that might never be gained through other simulations or experimental techniques. It is extremely difficult to conduct empirical studies on, for example, the short-lived hydrogen bonds or other contacts and short-distance range structures in a liquid. Yet these have a significant effect on, for example, solvents' functions. Over the past few years, parameterized energy potentials have been used in several simulations. The results gained in the AIMDliqD project can also be used for evaluating the accuracy of these simulations.

One of the biggest simulations ever

Ari Paavo Seitsonen started his current work at the beginning of 2005. He is involved in several collaboration projects with European researchers in the fields of physical chemistry and surface physics.

Seitsonen, as the responsible project leader collaborated in the AIMDliqD project with Professor **Barbara Kirchner** and her research group at the Wilhelm-Ostwald-Institut, Universität Leipzig, Professor **Jürg Hutter's** group at the Institute of Physical Chemistry, University of Zurich, and Professor **Kari Laasonen's** group at the University of Oulu.

The computational simulations for the AIMDliqD project were performed during the spring and early summer of 2007 at EPCC in Scotland. Ari Paavo Seitsonen was mainly responsible for the computation, but all groups have participated in the actual project planning and results analyses. The results analyses are still ongoing and once they are completed, several scientific publications will follow.

"The project is part of our continuing international collaboration. These research groups involved have worked together before, especially on ionic liquids, and work will continue, in the form of new projects. Only a few ab initio simulations have previously been executed for ionic liquids, and the one performed in the AIMDliqD project was one of the largest", Seitsonen explains, to clarify the significance of the project.

Discovering basic properties of ionic liquids

"In the AIMDliqD project we focused on two targets. The first was the effect of weak forces, van der Waals interactions in particular, on molecular dynamics and consequently, on the properties of the liquids. The second was the ionic liquids, which have become a keen topic of interest in recent years in the field of "green chemistry", says Seitsonen.

To date, very few ab initio molecular dynamic simulations have been carried out to investigate the importance of van der Waals forces, because modelling these forces analytically is difficult. Professor **Ursula Röthlisberger's** group at the EPFL in Lausanne has quite recently developed a computational method that allows the calculation of these forces to be included in the AIMD calculations with only a minor additional computational effort.

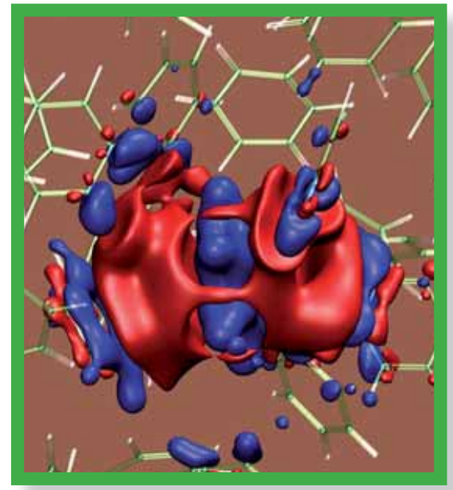

Change in electron density in liquid benzene due to the van der Waals interactions: Red regions, increased electron density; blue regions, decreased electron density.

"We simulated a liquid using an ab initio model with and without the van der Waals forces. The results of these computations enable us to make conclusions on the significance of these forces for the material's properties. The van der Waals forces are especially important in interactions between molecules. They are weaker than hydrogen bonding or electrostatic interactions that occur between charges. However, they are the origin of forces such as those keeping hydrocarbons in a liquid state", says Seitsonen.

The increasing interest in ionic liquids is due to the reduced losses during chemical reactions. Ionic liquids can be effectively recycled, so no chemical waste is produced. The diverse properties of ionic liquids also pave the way for energy-efficient or unmatched process conditions superior to those obtainable with traditional solvents. Many of the basic properties of ionic liquids are not completely understood; the AIMDliqD project serves to fulfil this information gap.

Figure on this page shows how electron density in liquid benzene changes, when van der Waals interactions have been taken into account via the method used in this project. The picture reveals that electron density is shifted from the immediate vicinity of atoms (blue) to the intra-molecular space (red). The electron density helps us to understand the attraction forces between molecules.

Massive tools are needed

The computation time allocated to the AIMDliqD project was essential for carrying out the study. The simulations require heavy-duty computation and they must be performed with massively parallel computers. To cope with these requirements, the resources given by DEISA were needed. The simulation of the ionic liquid consisted of 64 ion pairs (1-ethyl-3-methylimidazolium cation EMIM+ and thiocyanate anion SCN-), i.e. 1048 atoms with 3840 valence electrons to be calculated explicitly.

"For our practical calculations the Unicore Client interface was a very good tool, because it hid the complexities of the grid calculation from the user. Unicore Client is the software that integrates all DEISA centres. Computation jobs are submitted via the Unicore Client to be run by supercomputers in different computing centres. Furthermore, the participants at EPCC also supported us in running our computations, which was important especially at the early phase of the project", says Seitsonen.

The project belongs to basic research, and it does not have direct applications. Usually, temperature is not directly taken into account at the atomic level when reaction barriers are being studied; its effect is calculated afterwards using different approximations. This can comprise a serious flaw, if the solvent participates in the chemical reactions or if it plays a major role in another way. This can be avoided by using ab initio molecular dynamic methods. An aim for the future is to be able to simulate also industrially relevant chemical reactions in ionic liquids. Hence, simulations such as those in the AIMDliqD project could be used to introduce quantitative analyses for studying reaction barriers and processes. ■

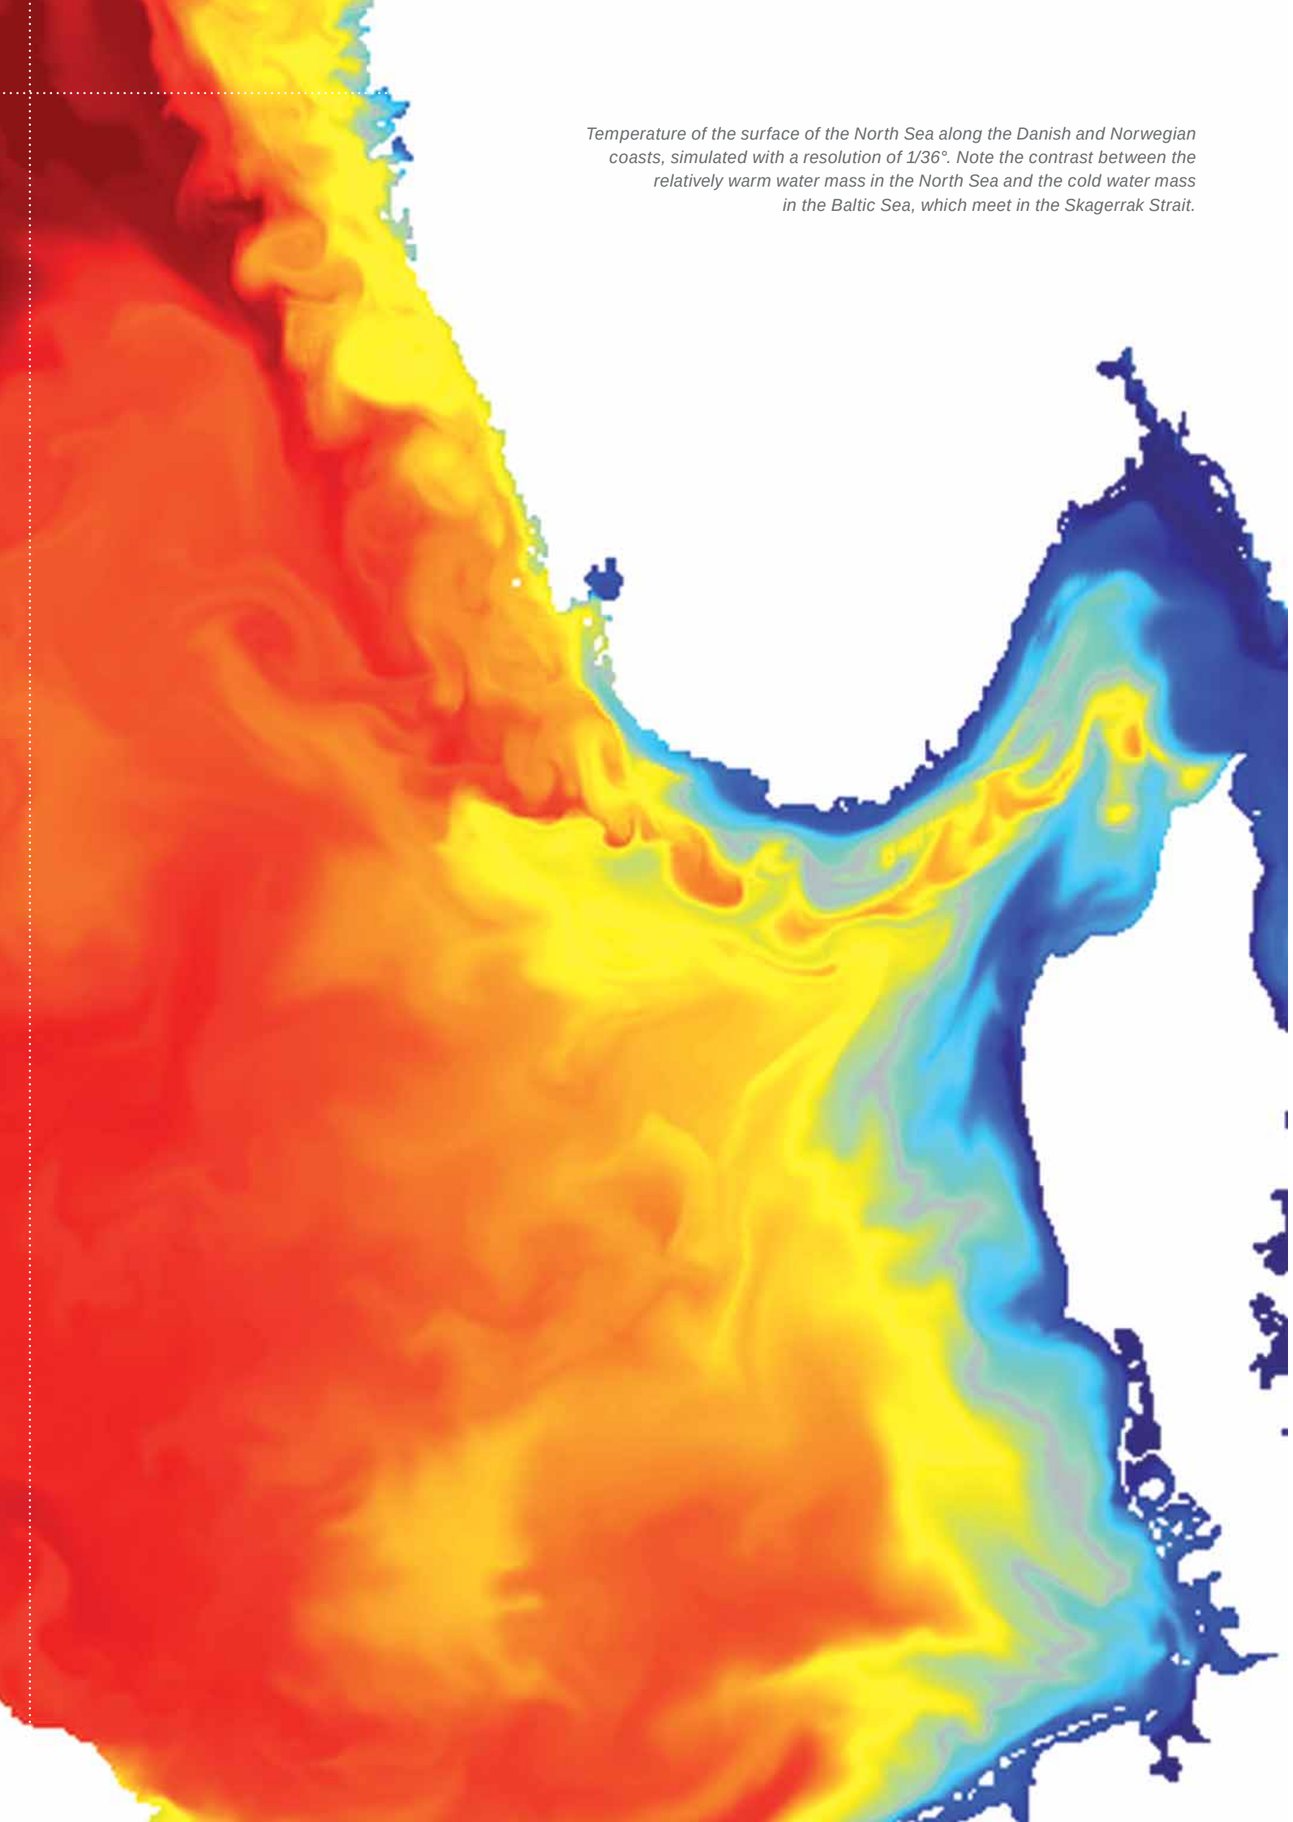

Temperature of the surface of the North Sea along the Danish and Norwegian coasts, simulated with a resolution of 1/36°. Note the contrast between the relatively warm water mass in the North Sea and the cold water mass in the Baltic Sea, which meet in the Skagerrak Strait.

Global to Regional **Oceanographic Modelling**

Damien Lecarpentier

Using DEISA computing resources, Pierre Bahurel and his team from Mercator Ocean in France have, together with ESEOO in Spain, 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 behavior.

Mercator Ocean was founded, as a consortium company, in 2002, with the aim of establishing an operational system for describing the state of the ocean at any given time and place on the planet.

"The ocean is an integral part of our environment, upon which many depend for survival, and it is a basic element of the Earth's climate. It is also an important site of transit for both goods and people", points out **Pierre Bahurel**, head of Mercator Ocean.

"Therefore, understanding both the state of the ocean and the ways in which it might change is crucial. Operational oceanography is essential for many scientific activities in the fields of oceanic physics, marine biology, and meteorology, and it has also practical benefits for the security and the improvement of maritime transportation", continues the French researcher.

Global and regional oceanographic modelling

In collaboration with ESEOO in Spain, Mercator Ocean launched the GROM project in 2006, with the general aim of developing global and regional operational oceanographic models.

"The project we are currently working on is to build and validate new ocean model configurations on a global to regional scale, with which we can simulate mesoscale or sub-mesoscale physical processes such as eddies, meanders, fronts or currents", explains **Yann Drillet**, ocean modeller at Mercator Ocean.

Mercator Ocean's team uses two types of models to conduct its research. The first one is called an "eddy-permitting" model, which allows a realistic representation of the main ocean currents. The second model is called "eddy-resolving", and offers finer representations of a number of mesoscale features of the ocean, essential for operational forecasting and for many activities (such as pollution detection and forecasting, ship routing, offshore fishing and halieutic resource management) for which the impact of these mesoscale processes is crucial.

"Our final goal is to develop a global eddy-resolving model (1/12°)", says Bahurel. Such a model is essential for understanding the evolution of the climate over the last decades and for providing an ocean analysis and forecasting service on a daily basis. We are also working on a better resolution for regional eddy-resolving models (1/36°) in specific areas, capable of representing sub-mesoscale and high frequency physical processes such as tides and storm waves, and providing useful large ocean forcings to scientists or operators involved in coastal oceanography", he explains.

Modelling the ocean requires high computation capacities

Modelling the ocean requires, however, very high computation capacities. In order to validate the simulations both physically and statistically, Mercator Ocean's team needs to know the model's behavior over relatively long periods – and each year simulated requires tens of thousands of computing hours. Physical processes in the ocean, such as mesoscale phenomena or "mixing" are often very complicated to model, and others such as tides and storm waves, considerably increase the computation needs. >>>

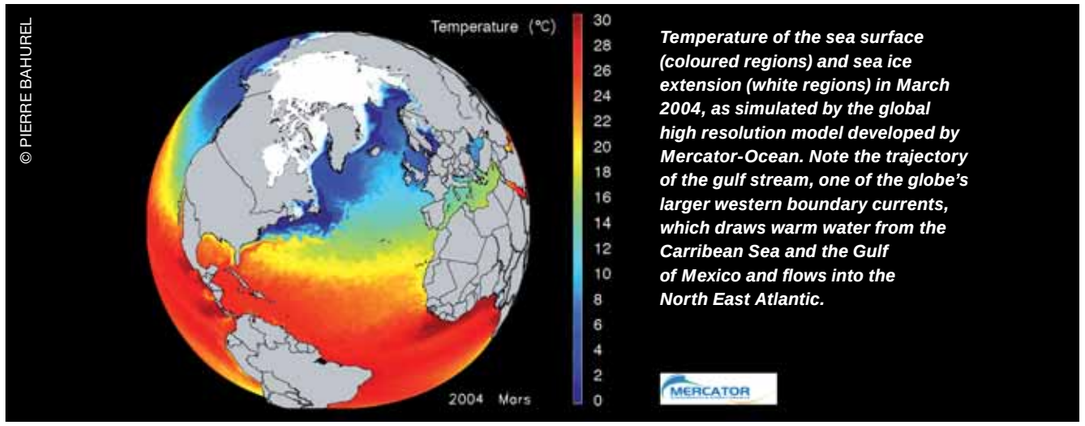

Temperature of the sea surface (coloured regions) and sea ice extension (white regions) in March 2004, as simulated by the global high resolution model developed by Mercator-Ocean. Note the trajectory of the gulf stream, one of the globe's larger western boundary currents, which draws warm water from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

By making available to Mercator Ocean some 450000 CPU hours on the ECMWF IBM, the DEISA Extreme Computing Initiative (DECI) made possible some major achievements within the GROM project:

"These computation capacities have enabled us to carry out interannual simulations, which are crucial for testing the validity of our models, for setting up systems of operational oceanic forecasting, and for deepening our understanding of the ocean more generally. Thanks to DEISA, we have been able to develop new models starting from the global eddy-permitting model to the regional eddy-resolving models", explains Drillet.

The results obtained by Mercator Ocean with the help of DEISA have proved extremely innovative and promising:

"We have been able to develop a global eddy-permitting model (1/4°) that allows us to represent an excellent degree of variability and eddy kinetic energy in all the oceans of the globe", begins Drillet.

"We have also been able to study for the first time the coupling between sea ice and a global eddy-permitting ocean", he adds. "Understanding this coupling is essential to efforts to realistically simulate circulation in the high latitude ocean, which has consequences for large-scale ocean circulation and deep water formation. An important aspect of the ocean forecasting is to provide information on sea ice-free transport pathways particularly in the Artic Ocean and in the Canadian straits."

Considerable achievements have also been made at a more regional level, especially in the North Atlantic area: "We have managed to represent correctly the Gulf Stream pathway and in particular the separation of the current from the coast at Cap Haterras to become a zonal jet in the Atlantic", says Drillet. According to the researcher, "the eddy-resolving simulation obtained (1/12°) is the most realistic simulation ever produced."

Photo of Pierre Bahurel, head of Mercator Ocean (on the right) and Yann Drillet, Ocean Modeller at Mercator Ocean.

"The accuracy of these simulation results in the key area of the North Atlantic region is new for to the entire physical oceanographic modelling community", points out Drillet.

"It allows us to represent the position and the intensity of the Gulf Stream and its North Atlantic current extension, which has a strong influence on oceanic and atmospheric properties", says Bahurel.

"It also enables us to represent the cold and warm eddies that are formed in this current and which transport temperature and salinity properties in the Sargasso Sea in the south and along the North American coast", adds Drillet.

"We have also developed a regional model within the North East Atlantic and the Western Mediteranean Sea areas, which allows us to refine locally the modelling process. It makes it now possible to introduce new physical phenomena, such as tides, into these models, and to get a better representation of the level of fresh water penetration from rivers into the ocean. This vastly improved the quality of our simulations, and created many new potential applications, especially on the continental shelf where human and biological activities are more important", says Bahurel.

First steps towards finer resolutions

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

"The process of improving the horizontal resolution in the ocean modelling is not complete. One of our objectives is to be able to simulate sub-mesoscale processes or coastal phenomena with a resolution of 1km. This will be soon possible with regional models thanks to progress in the level of computational power", says Bahurel.

"At the global scale, ocean observations have to be improved in terms of resolution, global coverage and real time distribution. Given that the atmosphere is the ocean's engine, we will also have to improve our knowledge of this component. It is crucial for the future development of our environment monitoring and forecasting capabilities", concludes Bahurel. ■

Galaxies Intergalactic Medium Interaction Calculation

Christopher Lazou

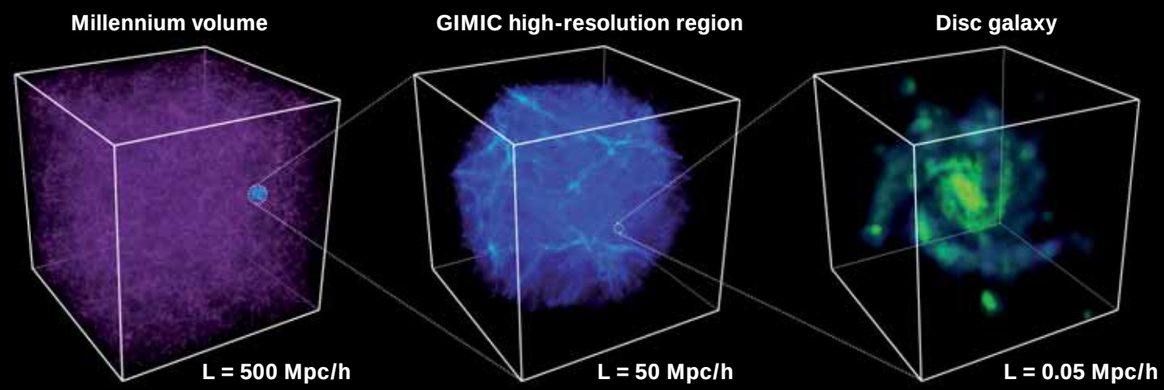

Millennium simulation is the largest ever calculation to follow the formation of the dark haloes that seed galaxies.

The figure depicts structure formation in a computer simulated Universe covering a dyanamic range of a factor of 10 000 in linear scale. The left most image shows the Millennium simulation (Springel et 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.

To understand the properties of the galaxies themselves, it is necessary to simulate how gas cools and forms stars in such haloes, says Professor **Carlos Frenk**, Director of the Institute for Computational Cosmology (ICC) and Principal Investigator of the Virgo consortium. Galaxies Intergalactic Medium Interaction Calculation (GIMIC) is a project of the Virgo consortium. Virgo is an international consortium of cosmologists that performs large numerical simulations of the formation of galaxies and the cosmic large-scale structure. GIMIC simulates the formation of galaxies in five regions selected from the Millennium simulation, but now it includes hydrodynamics. This allows Virgo members to obtain unprecedented insight into how galaxies form on truly cosmological scales.

GIMIC makes full use of DECI's (DEISA Extreme Computing Initiative) common data repository and coordinated scheduling in a work farm approach to computation scheduling and post-processing, thereby facilitating joint international analysis. These simulations were performed within the DECI initiative of DEISA, and were run on HPCx with the assistance of EPCC.

GIMIC is a collaboration project between the Max Planck Institute of Astrophysics in Germany and the Institute for Computational Cosmology in Durham, in the UK. GIMIC includes also colleagues at Leiden in the Netherlands, at Sussex and at Nottingham in the UK. >>>

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.

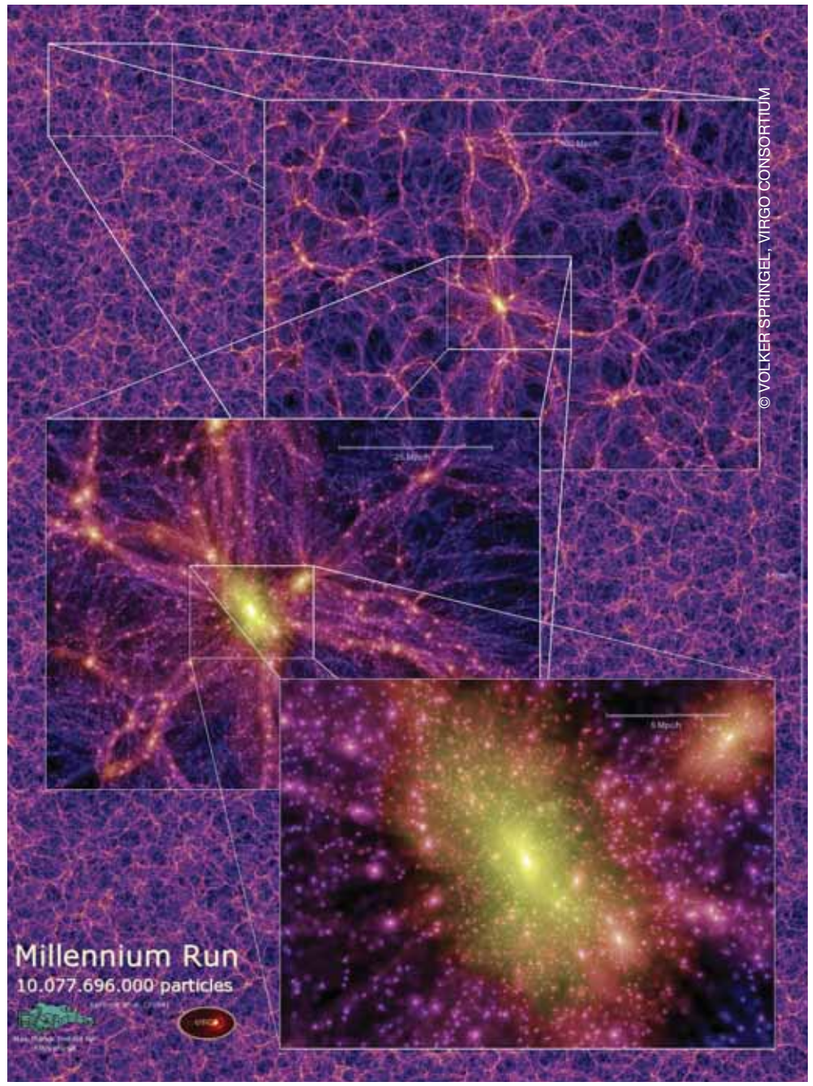

DEISA provides a focus for this large international collaboration, contributing important computing resources and a common data repository, which allows members in different countries to monitor the simulation and access the data.

Simulating the formation and evolution of galaxies and other structures in the universe from cosmological initial conditions, is one of the most difficult problems in computational physics. The non-local nature of the gravitational interaction makes it difficult to split the computational domain into more or less isolated areas that can be computed in parallel. In order to simulate the evolution of a patch of universe, one needs to account for the contribution of all matter in the rest of the universe. This requires subtle parallelization strategies.

The non-locality of the computational domains is a major bottleneck for the efficient use of cosmological computational codes; one needs to have access to the largest parallel supercomputers with low latency interconnection, as available in DEISA.

Provisional Results

"We are still actively analysing the data from the simulation. This is a non-trivial process because of the unprecedented size of the simulation, which requires the development of new analysis techniques. However, we have already been able identify where most of the matter in the universe is expected to reside. At present, astronomers have direct evidence for only about one percent of the mass in the universe", says Frenk. The Millennium simulation had clearly shown that the dark matter is arranged in filamentary structures made up of dark matter clumps – the cosmic web. GIMIC has now revealed that astrophysical processes separate the ordinary or baryonic matter from the dark matter even on large scales. As gas collapses to make a galaxy, the energy liberated by stars can blow powerful winds which heat the surrounding gas and pollute it with the products of nuclear fusion in the centres of stars – what astronomers call "heavy elements." "We now have an inventory of the distribution and thermodynamic state of the baryonic matter in the universe and its heavy element content. This will serve to guide astronomical searches for the currently missing bulk of the mass in the universe", explains Frenk.

In spite of the progress they have made in the GIMIC programme, the problem of galaxy formation remains largely unsolved. "Nobody in the world has yet succeeded in producing a realistic spiral galaxy like the Milky Way in a computer. We do not yet know if the reason for this is our poor understanding of the physics of galaxy formation or if our cosmological model is somehow incomplete. For example, the cosmological model that has been so successfully explored in the Millennium simulation assumes a particular kind of dark matter, the so-called cold dark matter. Since the particles that would make up this cold dark matter have not yet been discovered in the laboratory, we cannot be sure that our assumptions are correct. Petaflop machines will simultaneously allow us to model the physics of galaxy formation with increasing realism and to explore alternative assumptions for the cosmological model, including the nature of the dark matter. Ultimately, we would like to simulate a representative region of the Universe with full gas physics – in short to create a virtual universe", concludes Frenk. ■

Professor Carlos Frenk explains their work on running a Galaxies Intergalactic Medium Interaction Calculation (GIMIC) computer simulation at EPCC under the auspices of DEISA.

DEISA Infrastructure and Virgo Codes ported to it

The Distributed European Infrastructure for Supercomputing Applications (DEISA) is a consortium of eleven leading national supercomputing centres that currently deploy and operate a persistent, production-quality, distributed supercomputing environment, within Europe. The purpose of this research infrastructure funded by the EU sixth Framework Programme (FP6) 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.

DECI is the acronym of the DEISA Extreme Computing Initiative project. The objective of DECI is to identify, enable and run a limited number of ground-breaking applications in different fields of science and technology. These applications should consist of complex, demanding and innovative simulations that are impossible to implement without the DEISA infrastructure.

"We undertook an activity to port and optimise two of Virgo's simulation codes, GADGET and FLASH, onto the DEISA infrastructure", says Frenk.

GADGET is a N-body dark matter/gas dynamic simulation written by Dr **Volker Springel**, Prof. **Simon White** and collaborators at the Max Planck Institute for Astrophysics. GADGET is ideal for simulations of, for example, individual dark haloes or galaxies, where there are only a limited number of very high-density regions, and is the best code available to model the evolution of the entire Universe.

FLASH is an Adaptive Mesh Refinement (AMR) code and was in part developed by the US ASC/Alliance Centre for Astrophysical Thermonuclear Flashes, at the University of Chicago and further developed by Dr **Tom Theuns**, at the Institute for Computational Cosmology, to model more efficiently the evolution of the Universe. FLASH is a grid-based code and as such, is better suited to capturing shocks than the particle-based GADGET code. In addition, employing FLASH enables a more robust validation process for Virgo's huge simulations.

Both the simulation codes were successfully ported and optimised on a number of the DEISA platforms. Indeed, many of these platform-independent optimisations will be incorporated back into the publicly available versions of these codes. Code migration has also been introduced to allow very long FLASH simulations. A job reaching the batch time limit on one DEISA platform would thus resubmit itself, migrate to a different platform if necessary and restart on the infrastructure at the earliest possible opportunity. Allowing codes to migrate in this way ensures the fastest time to solution.

GIMIC used 600 000 CPU hours at the HPCx supercomputer at EPCC in Edinburgh. Supplementary simulations were carried out at the Cosmology Machine in Durham and these utilized a further 500 000 CPU hours. The Millennium simulation from which the GIMIC regions were extracted had generated 40 Terabytes of data. For GIMIC, a smaller dataset, of only a few Terabytes, was necessary. ■

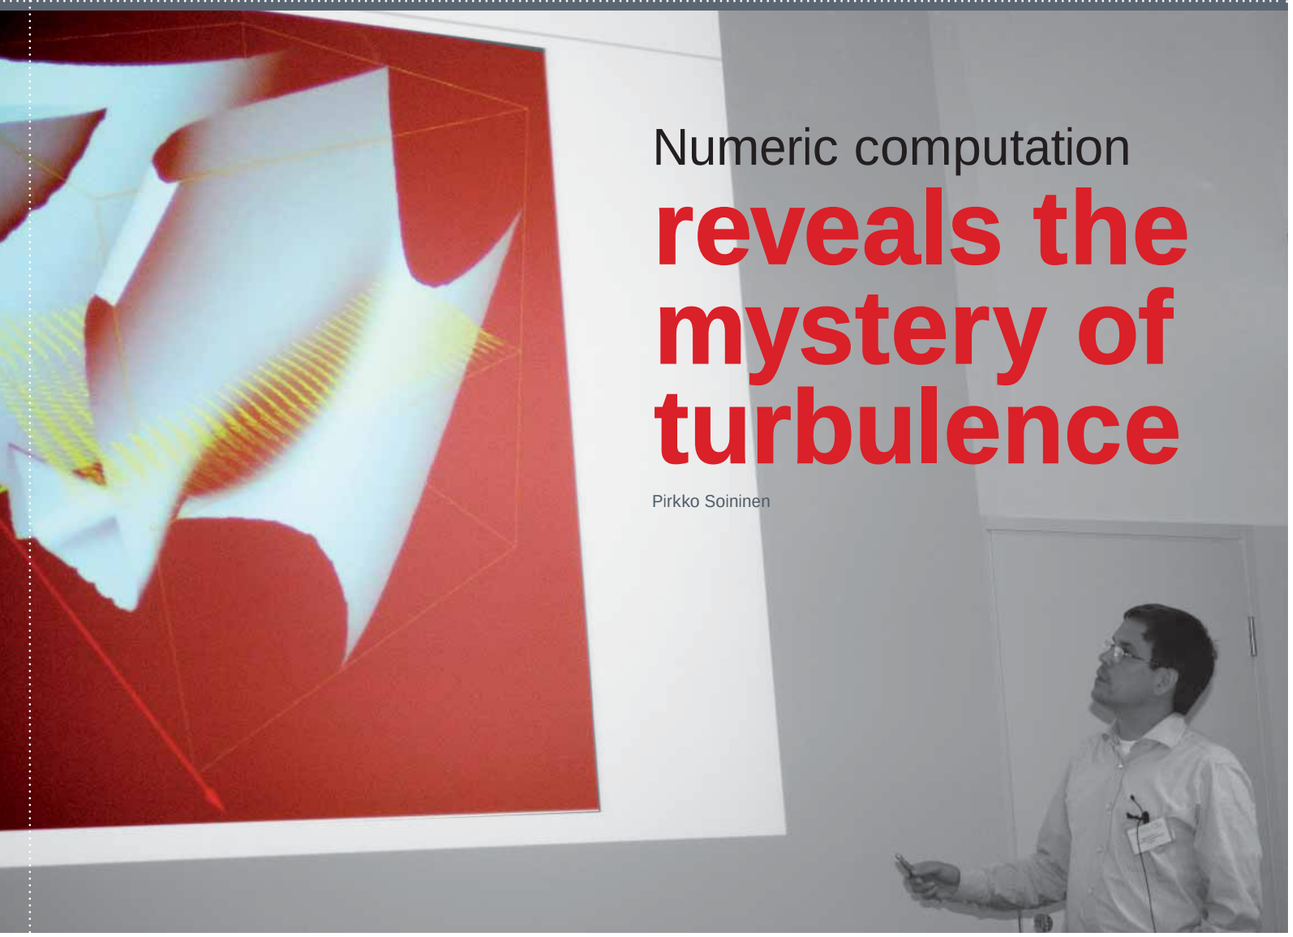

Numeric computation reveals the mystery of turbulence

Pirkko Soininen

Turbulence is a phenomenon that remains a mystery not yet revealed by classical physics. In Germany, Jörg Schumacher has performed numeric simulations on turbulence using the smallest possible scale. Such calculations would not have been possible without the DEISA supercomputers.

Resolving the nature of turbulence is one of the major challenges in physical research. The big question that remains unresolved in computational flow mechanics concerns the modeling of turbulence.

"The problem is not insignificant since turbulent phenomena occur practically everywhere. Consider, for example, the turbulent whirls we can see on the surface of the water in a river. These whirls dissipate into smaller ones to ultimately become the randomly distributed kinetic energy of molecules, in other words, heat", explains **Jörg Schumacher**, Ilmenau Technical University, Germany, who took part in the DEISA conference at the CSC.

Why is it important to study turbulence?

Turbulence is by definition a rapid variation of flow in relation to time, i.e. fluctuation. The definition is, however, not unambiguous. Fluctuation is three-dimensional and increases the internal diffusion of the substance, thus promoting, among other things, heat transfer and homogenization.

Turbulent flows can be seen, for example, in fallen leaves swirling around buildings in autumn. We see a large whirl that is all the time divided into smaller and smaller ones that are invisible to the human eye. Also, we can see the air whirling in the tree-tops, but we cannot distinguish the individual particles moving around.

Turbulence is exploited in many ways by nature and man. Turbulence is present in various types of air and fluid flows, and nearly everywhere. In many cases, it is a desired phenomenon. For example, in the chemical industry or combustion technology, turbulence is used to contribute to the mixing of various substances and their reactions.

On the other hand, in terms of aero- or hydrodynamics, turbulence is not a desired phenomenon. Rather, we strive to either prevent or at least delay its occurrence to avoid the increased flow resistance that is typical of turbulent flow and results in, for example, a higher fuel consumption of airplanes.

"In order to be able to understand turbulence, we should be able to model it correctly. However, highly complex models are needed, since turbulence involves events that are both complicated and irregular. This is why we tried to solve the problem by using computation on the smallest possible scale", says Schumacher.

Direct computation on a small scale

Jörg Schumacher and his colleagues have studied turbulent flows using a direct numerical simulation approach. Traditionally, turbulence has been modeled by means of averaged equations. Now, Schumacher has performed computations involving extremely fine whirls at a level where they are about to dissipate into non-existence.

"For the first time, we were able to really look into the inside of the phenomenon of turbulence. We have seen the powerful events that come into existence and die within turbulence", explains Schumacher.

The Direct Numerical Simulation (DNS) used by Schumacher is currently the most accurate method for numerical investigation of turbulent flow. The flow field is resolved directly from the Navier-Stokes equations without any averaging or turbulence modeling. The method requires immense computing capacity, and therefore, it has not been widely used.

Jörg Schumacher's team performed the simulation using a superfine grid.

"Numerical computation yields accurate and reliable information concerning small-scale phenomena that cannot be modeled", Schumacher points out.

The superfine grid density was determined by means of Reynolds number. With an increasing Reynolds number, the ratio of the largest and smallest longitudinal scales in a turbulent flow increases, in other words, the flow is dissipated into finer and finer whirls. The computation grid had to be dense enough to model the motions on the most minute scale.

"We have been able to confirm the theoretical hypotheses concerning the behavior of the finest turbulence. We have revealed the fundamental structure of turbulence", concludes the German researcher.

The result concerns an extremely fine portion of a turbulent flow, and as the next step, Schumacher would like to explore the intermingling and distribution of turbulent phenomena. Increasing the computation volume would, however, require a supercomputer capacity that is not yet available.

A test to replace wind tunnels

Schumacher's team utilized the DEISA supercomputers in their simulations.

"Data processing was extremely challenging, because the simulation was so extensive. A particularly high resolution was applied in the computation. We used a spatial grid with over eight billion cells, so the simulation generated an enormous volume of data. We had to figure out in advance how to arrange the data in an optimal and efficient manner", says Schumacher.

The data volume was so large that it simply was not possible to store everything on a disk.

"A scientist would, of course, like to save and store all the data generated, but in this case it was not possible. We were able to store only an approximate 15 percent of all the data", Schumacher regrets.

Schumacher's simulation took altogether 45 days and 800 000 CPU hours.

"I believe that, in the future, the traditional flow calculation methods and our direct numerical simulation may converge. This, however, calls for a higher computing capacity", says Schumacher.

By means of computational methods, we have, in recent years, been able to reduce the expensive wind tunnel experiments necessary in, for example, the automotive industry.

"Flow in the wind tunnel differs from the reality, if the parameters are not exactly accurate. The technology is expensive and laborious. So, if we can develop more sophisticated models for turbulence, there will be less need for expensive wind tunnel measurements for cars, trains and airplanes", Schumacher concludes. ■

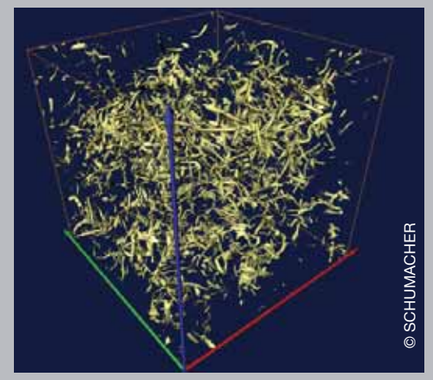

Most intense vortices in a turbulent flow. The figure shows an isovolume plot of the magnitude of vorticity. The maxima are organized in tube-like structures. The flow turbulence is simulated in a box with periodic side walls for a grid resolution of 2048 points in each of the three space dimensions.

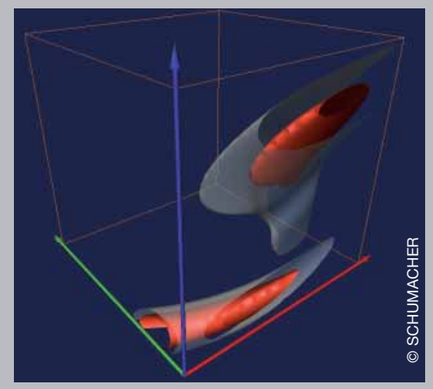

Interaction of two intense vortices. This "dance" around each other is studied in the simulations. Therefore a swarm of tracers goes with the turbulent flow. In their vicinity the vortex structures are monitored with respect to time. The figure shows such an observation cube. We can see two twisted tube segments which are indicated by the nested isosurfaces taken at different vorticity intensity levels. The stretching of the tubes leads to a rapid growth of the local vorticity.

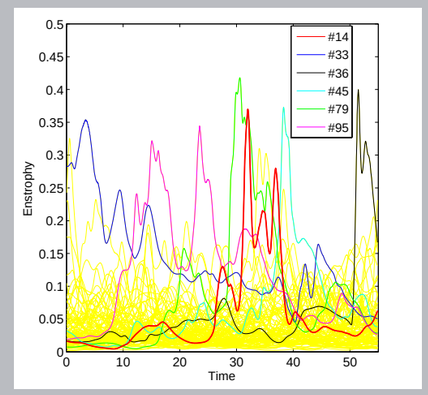

Time traces of the accumulated vorticity in each of the observation cubes. One hundred tracers were moving through different parts of the box. Six of these one hundred tracers monitored a very rapid temporal growth of vorticity. The peak of the thick red line (tracer no. 14) corresponds with the zoom-in of figure above.

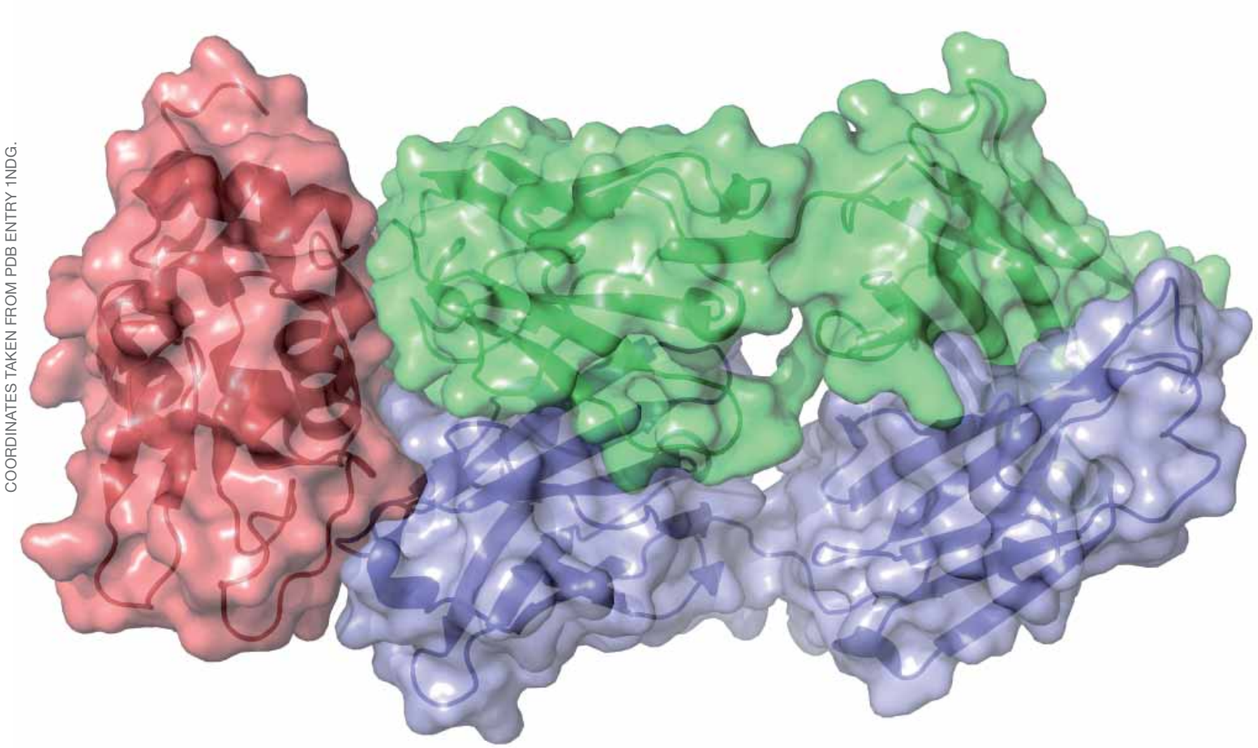

COORDINATES TAKEN FROM PDB ENTRY 1NDG.

Molecular simulation to accelerate vaccine development

Saara Värttö

A large three-year project (BacAbs) involving pharmaceutical and bio-informatics companies and Universities has exploited the DEISA resources to render vaccine development simpler, faster and therefore less expensive for pharmaceutical companies and eventually to the end-users being treated.

The goal of the project is to develop tools that help predict which antigens will be suitable for vaccine development at an initial stage of the development process. The molecular dynamics simulations run on DEISA were finalized at the end of 2007, and the final results of the whole European Union's FP6 funded BacAbs project will be available in 2009.

"In the process of vaccine development, the early stages until the immunogenic testing can take a couple of years: If we'd be able to reduce the number of proteins expressed – say by half – this would be a huge gain, since the initial steps are heavy consisting of hundreds of proteins", tells **Xavier Daura**, who works at Universitat Autonoma de Barcelona.

Daura's group evaluates the structural requirements for viable bactericidal vaccine candidates through a novel multidisciplinary approach and develops bioinformatics tools to predict compliance with such structural requirements.

"Many bacterial proteins actually do elicit production of antibodies when injected into a model animal, but only asmall fraction of them are good antigens in the sense that together with the antibody it will produce a bactericidal response which will lead to immunity", says Daura. He explains that there are many complexes between an antigen and an antibody that do not lead to the killing of the pathogen. "For the cure the antigen and the antibody must bind to one another effectively, and this antigen and antibody complex must be recognized by the so-called complement system. There is a particular protein, C1q, which binds first. We will also try to study this type of interaction."

To this end, a systematic analysis of sequence, structure, dynamics and interactions of potential antigens are analysed systematically using as model system serogroup-B Neisseria meningitidis, a pathogen causing septicemia and meningitis, and for which no effective vaccine exists. This analysis involves the application of several computational techniques, some of which posing extreme demands on CPU time, to the study of complex biomolecular systems.

Protein antigen bound to the Fab (antigen-binding) domain of an antibody in ribbons and molecular surface representation.

Specifically, a list of antigens (proteins) and their complexes with antibodies and the complement component C1q has been studied using Monte–Carlo simulations of coarse-grained protein models to assist structure determination by experimental partners in the consortium. Extensive molecular-dynamics simulations have also been run to characterize the conformational properties of specific protein regions – epitopes.

"We are interested in both global dynamics and in the dynamics of specific regions", says Daura. "One effort in our consortium is to map the regions that interact with monoclonal antibodies. This is called epitope mapping. Epitope is a small part of a macromolecule that is recognized by the immune system, in this case by antibodies." He adds that there are basically two types of recognition by antibodies. The one is purely by sequence and the other one is by structure. This means that if the antibody is able to recognize a peptide independently of its shape, it will indicate that it is identified by the sequence and not by the structure. In some cases the antibody will only recognize a specific peptide if it is in a specific structure. Connected to that, one possibility is that epitopes which adapt to a specific structure are recognized by it.

Continues Daura: "We have hypothesized that epitopes may have similar structural properties when isolated in solution. This is why we do all these studies on epitopes, their sequences separated from the proteins. There is something intrinsic in the structure of the peptide that makes it recognizable by an antibody independently of the protein."

Need of Supercomputers

"Our research is very challenging since there are a lot of degrees of freedom, many interactions between the proteins and such system sizes that must be simplified to be run with currently available resources", says Daura and explains that at the moment it is possible to do a simulation in the range of nanoseconds to microseconds. "However, this is far from the length of any biological process, which normally is in the time frame from milliseconds to seconds. It is likely to take decades before supercomputers will be powerful enough to simulate a whole biological process", points out Daura.

The research group needs to obtain from their simulations information on as many proteins as possible in order to have results that are statistically significant. Supercomputers are essential for the research in order to study significant number of proteins and protein complexes. The research group was granted DEISA resources for one year and they will apply new supercomputing resources for 2008–2009 in order to continue to have new information on proteins.

The BacAbsMS project lead by Xavier Daura accounts for the computationally intensive part of the FP6 funded project BacAbs. The access to the DEISA infrastructure has enabled the analysis of a significant number of protein targets from bacteria of medical interest using molecular dynamics simulation methods.

© SAARA VÄRTTÖ

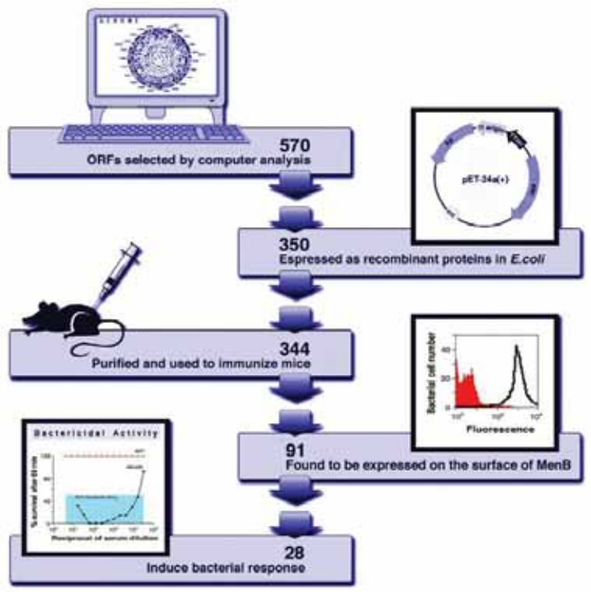

© NOVARTIS VACCINES & DIAGNOSTICS

Initial steps in a vaccine development process, from research at Novartis Vaccines & Diagnostics (partner to this project) on group-B N. meningitidis.

Part of the project is dealing with docking and development of algorithms. Molecular docking is a method which predicts the preferred orientation of one molecule to a second when bound to each other to form a stable complex. This part could be run in a grid type of a system since although thousands of jobs are needed, they can run independently. However the problem is to find a grid system that executes the jobs with the required speed. ■

For more information:
http://www.bacabs.org/
http://www.deisa.org/applications/projects2006-2007/BacAbsMS.php

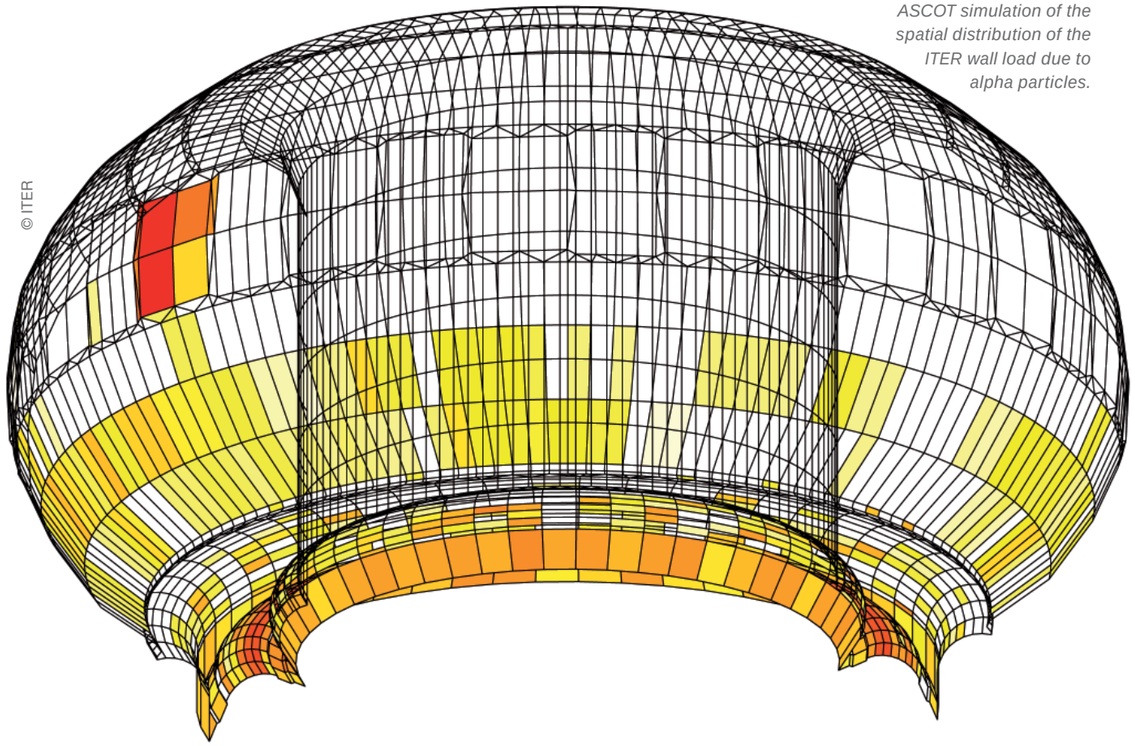

© ITER

ASCOT simulation of the spatial distribution of the ITER wall load due to alpha particles.

Turbulence in FUSION PLASMA

Kaisa Riikilä • Tommi Kutilainen

Timo Kiviniemi's research group is studying plasma confinement in a tokamak-type fusion reactor. The group's researchers were the first in the world to show that it is possible to simulate both the formation of an electric field and its impact on turbulence by using the same software.

Fusion energy is one of the rare alternatives for solving the world's energy problem. It has been estimated that the world's energy consumption will double or triple by the end of the century. Even if the Western world can achieve the goals set for saving energy, it will not be able to compensate for the growing energy needs of the current developing countries.

Fusion energy has several advantages over nuclear fission power or fossil fuels. Fusion energy has minor effects on the environment, it does not generate long-term radioactive waste and the resources are abundant, evenly dispersed around Earth.

"Nuclear fusion is a very challenging research field. At the moment, it involves three main problems: heating, confinement and materials' durability. Our research deals with the confinement of fusion plasma", says **Timo Kiviniemi** who works at the Laboratory of Advanced Energy Systems at Helsinki University of Technology.

To release significant amounts of energy in a fusion reaction under terrestrial conditions requires a temperature of over 100 million °C. Attaining such extreme temperatures requires efficient heating of plasma with minimized thermal losses. The best performance has been attained by using tokamak-type test reactors,

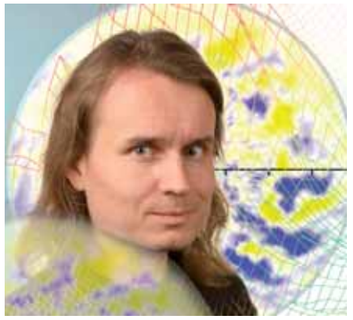

© ERKKI MAKKONEN

Timo Kiviniemi works at the Laboratory of Advanced Energy Systems at Helsinki University of Technology.

where plasma is floated with the help of strong magnets in the chamber to form a toroidal (doughnut-shape) ring. The magnetic field confines the plasma by preventing charged particles from escaping; it also keeps plasma detached from the chamber walls, which reduces thermal loss and stress on the reactor wall structures. There are several experimental devices being tested in Europe and they continuously produce basic information needed to advance fusion research.

Turbulence impedes confinement

The electronically charged particles in plasma, electrons and ions (nuclei), usually adhere to the magnetic flow surfaces of plasma, moving along the magnetic field lines on the flow surfaces. However, heat and particle transport occurs in plasma as they move outwards, perpendicular to the flow surfaces, and finally the particles disappear into the wall structures.

"Most of the particle and heat transport in fusion plasma is due to turbulence. When this process is better understood, we hope to be able to reduce it", says Timo Kiviniemi.

A stepwise improvement in confinement was observed under laboratory conditions as early as in the 1980s, but even today it still cannot be fully explained. One of the essential theories is that the rotational speed of plasma on different flow surfaces, i.e. at different distances from the plasma core, changes so fast that the non-simultaneous timing of the rotation movement tears eddies apart. The assumption is that plasma rotation is due to an electric field formed within plasma.

"Based on our earlier research, we assume that the electric fields are so strong that indeed, they should reduce turbulence", Kiviniemi continues.

Kiviniemi's research group is now investigating how the intensively fluctuating electric field is generated and why it makes turbulence disappear – this knowledge will help in finding a way to suppress the impeding turbulence.

The disappearance of turbulence is fairly well understood, but Kiviniemi's group was the first to use the same software for simultaneous simulation of both the neoclassical electric field formation and turbulence. The simulation was performed using Elmfire software.

Particle simulations require plenty of computing resources

Timo Kiviniemi's research group is one of the heaviest users of CSC's computing resources. Within the DEISA Extreme Computing Initiative the calculations have been mostly carried out at RZG in Germany. Using DEISA resources, the researchers have launched the simulation of the edge of the middle-sized ASDEX Upgrade tokamak with H-mode parameters with the ambitious goal of clarifying the mechanisms behind the confinement transitions. More DEISA runs are scheduled for 2008 to test more realistic outer boundary conditions.

"What we do is massive parallel computing. Usually we utilize more than a hundred processors at a time. By using just one processor, it would take about a year to perform just one of our simulations." Kiviniemi explains: "In other words, an ordinary desktop computer would be tied up with the calculation for a year and could not be used for anything else."

Compared with the continuum approach, the advantage of the particle simulation approach used by Kiviniemi is the fact that particle simulation excels in taking into account plasma edges and their structures where particles may collide. It is also suitable for steep-gradient simulations.

"We have a pool of several tens of millions of test particles. The number of particles in our tokamak is 1020, so every simulation particle we have corresponds roughly to 1000 billion real particles", says Timo Kiviniemi. >>>

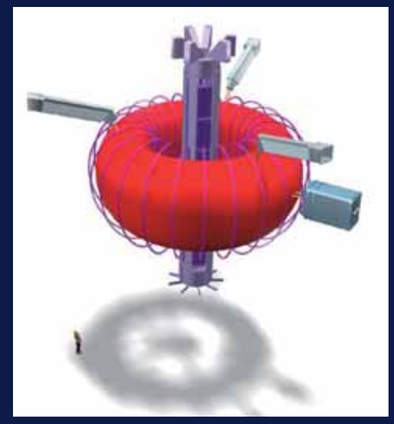

Tokamak

Tokamak is a Russian acronym, derived from the words toroidalnaja kamera magnitnaja katushka, which can be translated as a toroidal chamber inside magnetic coils. The idea of the tokamak was invented by Russian physicists Tamm and Saharov in 1951.

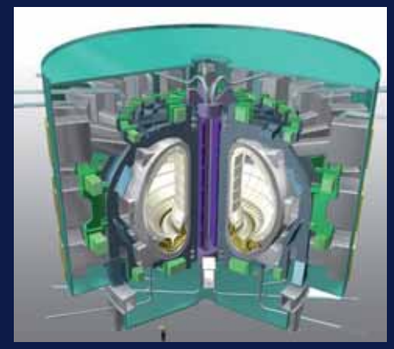

The tokamak fusion reactor is based on the magnetic confinement of fusion plasma. Magnets are used to confine plasma as a toroidal ring in the chamber of the fusion reactor. A helical total magnetic field prevents plasma particles from escaping and keeps plasma detached from the chamber walls. This way the plasma can be heated to over 100 million °C. This is the temperature needed to generate enough kinetic energy for the positively charged nuclei to exceed the repulsion force between the nuclei, hence allowing the fusion reaction. One gram of deuterium-tritium fuel produced in a fusion reaction generates 100 000 kWh of electricity. It would take 8000 kg of coal to produce this amount of energy.

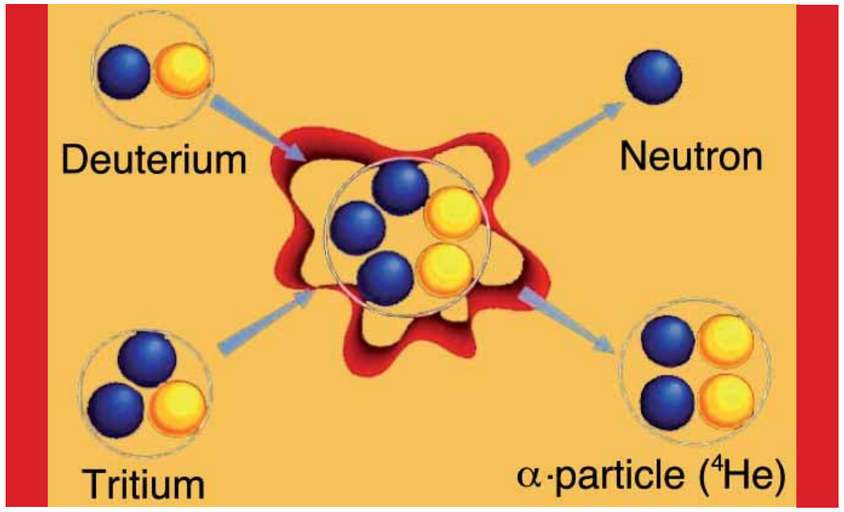

In the fusion reaction that is the most feasible under terrestrial conditions, the heavy isotopes of hydrogen, deuterium and tritium, can be fused into a helium nuclei and a neutron. Exposed to extreme heat, hydrogen isotopes turn into plasma, a form of ionized gas where electrons are completely detached from their atomic nuclei. Approximately 30 g of deuterium can be generated from a cubic meter of sea water. Tritium is very rare on Earth, but it can be produced from lithium, which is abundantly present in the Earth's crust. Of these, deuterium and lithium, and helium as the end product, are not radioactive. Tritium, produced as the interim product, is radioactive, but it is treated and used completely inside the fusion power plant. In the future, fusion reactors may use fuels that contain no radioactive isotopes whatsoever.

Noise in simulation

There is also a dark side in particle simulation compared with continuum simulation: the numerical noise, or statistical error, which is due to every simulation particle representing an extremely large number of real particles. The problem is reduced if the number of test particles is increased, but this requires a corresponding amount more computing resources.

The researchers must optimize the highest acceptable level for numerical noise. Suppressing the noise is expensive, because by increasing the number of particles causes a slow reduction in the noise. Hence, researchers must think carefully and decide the level of noise they can live with.

"We have simulated the smaller tokamak for some three to four years, and now we are starting the simulation experiment with the bigger one. Its simulation is much heavier, since the computation time for simulation is directly proportional to the size of the device."

Elmfire is not, as yet, capable of simulating the ITER test reactor being built in France. But, as Kiviniemi says, computer development is fast. By the time that ITER is finished in 2016, computer efficiency should be sufficient to simulate the turbulence. ITER is being built through an international collaboration project, and it is an important intermediate phase for fusion research.

Fusion power is approaching

The key feature of European fusion research is exact coordination, which has made it possible to allocate the available research resources efficiently to essential areas of research. The fusion plasma research at Helsinki University of Technology is also part of the carefully coordinated program. The national program in Finland is coordinated by the Finnish Funding Agency for Technology and Innovation (Tekes) and the program is strongly integrated to the European fusion program.

The ITER tokamak being built in France will not be connected to the mains network yet. It will be used as a research instrument to demonstrate, among other things, that the power plant can actually generate more energy than it consumes.

"This has never been proved with any device before. However, it has been observed that the ratio improves with the size of the device", says Kiviniemi.

He estimates that in the future, the energy volumes generated by the fusion power plants then in commercial use will roughly match the capacities of the largest currently used power plants, about 1000 megawatts.

"The political decision-making process on the ITER reactor has taken an eternity. Had more money been available for fusion research, the reactor could have been built earlier", says Kiviniemi.

"However, I am convinced that fusion reactors will be used for energy production. Certainly, fusion energy will be expensive at the beginning, but as soon as the technology has been improved, the prices will come down. As other fuel resources become depleted, their price will go up and make fusion energy quite competitive", says Kiviniemi. ■

Snapshot from a GENE simulation depicting turbulent transport in the interior of a fusion experiment.

Towards ITER on supercomputers

Damien Lecarpentier

The GYROKINETICS project was carried out in 2006 and 2007 by researchers from the Max Planck Institute for Plasma Physics at Garching, Germany, and the Ecole Polytechnique Fédérale of Lausanne, in Switzerland. Using DEISA's resources under the DECI and the JRA3 frameworks, the project team further developed the gyrokinetic simulation approach to plasma turbulence, which is expected to help improve the performance of magnetic confinement fusion devices.

>>>

"Magnetic confinement fusion has the potential to provide a substantial proportion of the world's energy needs in the 21st century – and beyond – in a safe and environmentally friendly way", says **Frank Jenko**, researcher at the Max Planck Institute for Plasma Physics and member of the GYROKINETICS project.

"Its realization is, however, hampered by the complex behavior of hot collisionless plasmas (ion gases) in strong magnetic fields. Such plasmas are subject to temperature and density gradient driven microturbulence, which determines the energy and particle transport across flux surfaces and hence the minimum size of a burning plasma. Turbulence, thus, limits the so-called energy confinement time, keeping the plasma from reaching a state in which external heating becomes dispensable", he explains.

"Gaining a thorough theoretical understanding of turbulent transport is, therefore, crucial for improving the performance of magnetic confinement fusion devices, and constitutes one of the key goals in modern fusion research."

Understanding turbulent transport through gyrokinetic simulation

"Simulations are necessary if we are to understand and control plasma microturbulence. However", Jenko continues, "because fusion plasmas are virtually collisionless, a three-dimensional (i.e. in space) fluid description must, in principle, be abandonned, in favor of a six-dimensional (i.e. in phase space) kinetic one."

"Fortunately, several processes on very small spatio-temporal scales – such as the gyrating motion of the particles around magnetic field lines – can be removed, analytically, from the basic equations, thus making the problem five-dimensional. This reduces the computational requirements by many orders of magnitude, without sacrificing accuracy. This approach is called gyrokinetics, which gave the present project its name."

"Over the last few years, several massively parallel gyrokinetic codes, capable of solving nonlinear gyrokinetic equations, have been developed throughout the world, and they have matured to a point at which direct comparisons with experimental data are becoming feasible", Jenko notes.

During the project, Frank Jenko and his colleagues used, in particular, two state-of-the-art gyrokinetic plasma turbulence codes, which were independently developed at the two participating institutions, and which were based on complementary numerical approaches.

These two codes are the "particle-in-cell" ORB5 code, which computes the trajectories of large ensembles of marker particles, and the CFD-type GENE code, which represents the particle distribution functions on a fixed grid in phase space.

"Given that the simulation of turbulence in magnetized plasmas represents a major challenge at the forefront of computational plasma physics, a dual approach like this seemed particularly appropriate. By benchmarking the results of both codes, we expected to both gain more insight into the behaviour of these different numerical methods and to maximize the reliability of the physical results that were obtained", Jenko explains.

Simulations with unprecedented levels of resolution were realised

For the project, the gyrokinetic plasma turbulence codes ORB5 and GENE were adapted to the DEISA environment within the JRA3 framework, and were then ported to the SGI Altix system.

The DEISA infrastructure also provided the researchers with computational resources in order to realise their simulations.

Frank Jenko, researcher at the Max Planck Institute for Plasma Physics and member of the GYROKINETICS project.

"We were granted something in the order of 400 000 CPU hours", notes Jenko. "With that budget, we were able to perform a number of simulations with GENE and ORB5, and thus to assess the role of fluctuations on very small spatio-temporal scales (smaller than the ion gyroradius and the ion transit time)."

"These scales had been neglected in previous studies, 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 in the analysis. Our simulations were some of the largest achieved to date in computational fusion research, with an unprecedented level of resolution", he adds.

"As a result, we were able to show that certain small-scale turbulent processes can make substantial – and even dominant – contributions to the overall heat transport carried by the plasma electrons. It turned out, in particular, 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 proportion of the latter can be carried by much smaller scales."

According to the research team, these findings represent an important new insight into the physics of turbulent transport in magnetized plasmas, and will have important implications for future full-torus simulations of large fusion devices, such as the International Thermonuclear Experimental Reactor ITER.

Virtual fusion experiment: the ultimate goal

For qualitative progress in the theoretical understanding of turbulence in fusion devices, and in the capability to quantitatively predict its consequences, supercomputers – as instruments of numerical modelling – are absolutely essential. Research on magnetic confinement fusion has already moved away from the semi-empirical and predominantly experiment-driven approach to one based on numerical modelling. These changes suggest that experiments might one day be replaced entirely by simulations.

"Our ultimate goal is to conduct virtual fusion experiments", says Jenko, "in which simulations are so realistic that they could almost come to replace experiments."

"The aim of the GYROKINETICS project was to take significant steps in that direction, by extending the scale range far beyond that used in past studies. Virtual fusion experiments remains, however, many years away", Jenko acknowledges.

"To reach this goal, the range of spatio-temporal scales will have to be extended, and a number of additional physical effects will have to be included, gradually establishing computational gyrokinetics as a reliable tool for analyzing and optimizing actual experiments."

"Significant code development will be required and supercomputers will have to become much more powerful still. Comprehensive gyrokinetic simulations of a device like ITER will require Petascale or even Exascale supercomputers", he points out.

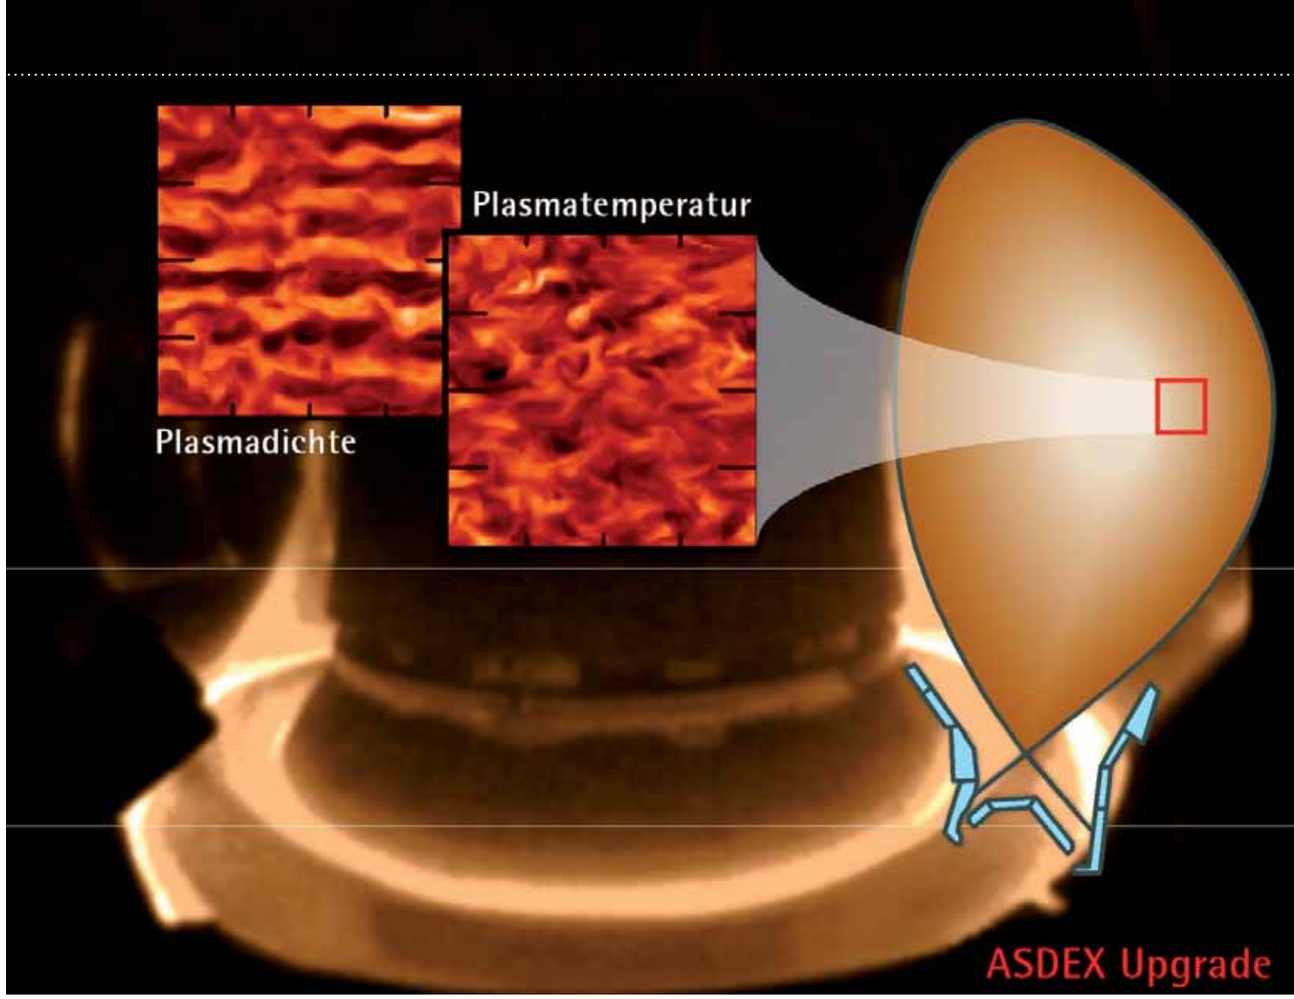

Interior view of the fusion experiment ASDEX-Upgrade during a plasma discharge together with two contour plots from gyrokinetic simulations, demonstrating turbulent behaviour.

"Given their ambitious goals and enormous resource demands, pushing both the software and the hardware they employ to their very limits, computational fusion physicists throughout Europe will always be eager to work with the latest generation of supercomputers. In this context, DEISA provides an excellent framework. It is likely to be instrumental in providing the fusion community with reliable predictive simulations of turbulent transport in fusion plasmas, thus helping to improve the performance of magnetic fusion devices," Jenko concludes. ■

The FOCUS project was carried out between 2005 and 2007 by the EM2C laboratory of the Ecole Centrale of Paris and the IDRIS team at CNRS in France. Using DEISA research infrastructure within the DECI programm, the researchers have developed a new method for simulating combustion processes that will have a real and practical impact on a number of industrial applications.

A new coupling method for simulating combustion processes

Damien Lecarpentier

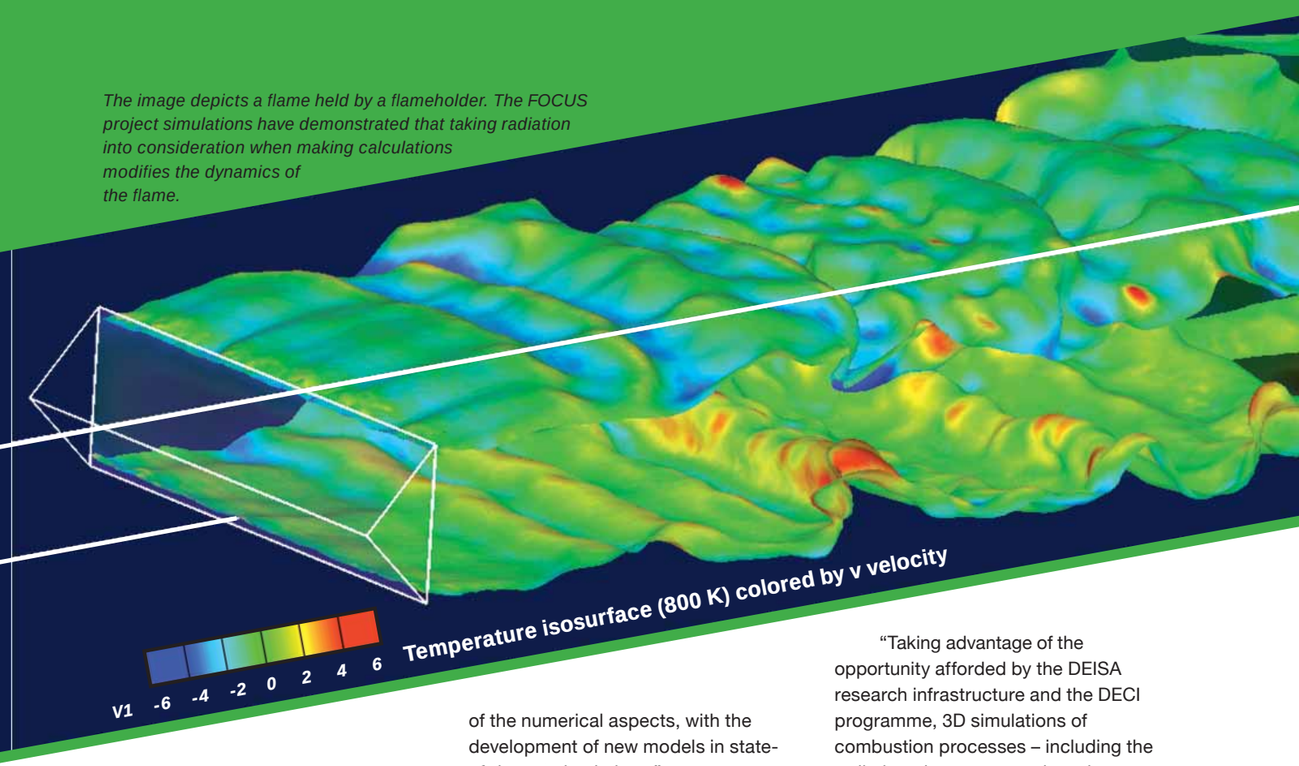

The image depicts a flame held by a flameholder. The FOCUS project simulations have demonstrated that taking radiation into consideration when making calculations modifies the dynamics of the flame.

Combustion is involved in more than 80% of primary energy conversion processes worldwide. It is used, for example, in aeronautical and ground transportation, waste incineration and in various other industries in which burners and engines are required.

For industrial applications, the type of combustion that is most often used is called turbulent combustion, characterized by turbulent flows that aid the mixing process between the combustible and the oxidizer. It remains, however, a complex phenomenon, characterized by many instabilities and side-effects (such as pollutant formation and radiative heat transfers), which need to be better understood.

"Understanding turbulent reacting flows, including elements such as pollutant formation and radiative heat transfers, is crucial for many practical industrial applications", explains **Olivier Gicquel**, a researcher at the EM2C laboratory and one of the leaders of the FOCUS project.

"Optimization of burners has led to increased levels of complexity and design innovation. This very complexity, however, increases the likelihood that the burners will malfunction, for example through combustion instabilities. These instabilities generate limitations in the operating domains of systems, and can significantly increase noise levels. They can even reduce the lifespan of burners or engines."

"Reduction of pollutant emissions is a key issue in the field of energy production. Legal constraints are becoming more and more important for fixed energy sources, such as boilers or gas turbines, and for transport applications, such as automotive engines or air jets", Gicquel adds.

"Understanding and predicting these phenomena is thus a central concern in the development of low emission systems", he concludes.

Numerical simulation of combustion processes is necessary

Because the developments of industrial prototypes can be very expensive, numerical simulation of combustion processes has become compulsory.

Olivier Gicquel, researcher at the EMC2 Laboratory of the Ecole Centrale de Paris and leader of the FOCUS project.

Various different methods exist for performing these simulations. The most recent is called Large Eddy Simulation (LES), and can be viewed as an intermediate technique between direct numerical simulation and classical modelization. Under this approach, larger turbulent motions of the flow field are explicitly computed and resolved, whereas the effects of the smaller ones are only modelled.

"This method appears to be a promising tool for such simulations", says Gicquel "Under this approach, cold fresh gas zones and hot burnt gas zones, which behave very differently in terms of pollutant formation and radiative heat transfers, and which are the most crucial areas in which we need information, are identified at the resolved scale level", he explains.

The physical phenomena involved in combustion and radiative heat transfers are, however, very different: "Flow fields, (e.g. temperature, chemical species, velocity) are generally described in terms of balance over small volumes (finite volume context), whereas radiative heat transfers involve long distance interactions."

"Accordingly, reacting flows and radiative heat transfers codes have a very different structure. Combustion codes are generally parallelized by domain splitting (in which each processor "sees" only a small part of the physical domain), whereas parallelization by wavelength ranges and/or by radiation directions is more appropriate for radiation (where each processor needs to work with the whole domain)."

To gain a better insight into the full process, the researchers of the FOCUS group have developed an original approach: "The proposed approach takes advantage of an efficient coupling between an LES solver and codes devoted to radiative heat transfers, where data exchanges occur at time intervals controlled by the physical times of each phenomenon", explains Gicquel. "This project is innovative both from a theoretical point of view and in terms of the numerical aspects, with the development of new models in state-of-the-art simulations."

Numerical simulations also require large computational resources

Numerical simulations of turbulent reacting flows including pollutant formation and radiative heat transfers require not only well-adapted models, but also large computational resources.

Research infrastructures like DEISA are therefore very much needed in this area. It has played a crucial role in the development of the FOCUS project:

"The DEISA project gave us access to the coupling technology that we required. It also provided us with computational resources to run the simulations", says Gicquel.

"The DEISA infrastructure provided two important resources for the FOCUS project", he continues:

"It allowed for the collaboration of two engineers with computer science backgrounds from the IDRIS team, which meant that we were able to improve the speed of our code by a factor of close to 10. This great improvement was crucial in enabling us to run the final simulations. Within the DECI framework, we were provided 300 000 CPU hours on the IDRIS super-computer, and we had dedicated access to up to 400 processors to run the coupled simulations."

"Taking advantage of the opportunity afforded by the DEISA research infrastructure and the DECI programme, 3D simulations of combustion processes – including the radiation phenomenon – have been carried out. The impact of radiation on the flame dynamics, which was one of our primary concerns, has been clearly investigated and evidenced."

From simulation to practical applications

The results obtained during the FOCUS project will be of benefit to many applications in which combustion processes are involved:

"The results we have obtained are of great interest for designers of gas turbines and furnace burners", notes Gicquel. "It will help lessen the likelihood of combustion instabilities developing within the combustion chamber, meaning that engine components will last longer", he explains.

"The new method that we have developed in this study may also be of relevance in various computational codes. The models will be available to aid the design of high performance – that is, safer, cleaner and more reliable – combustors for various industrial applications."

"Thanks to the DEISA framework, the EM2C laboratory is now one of the leading teams in the international combustion community in the key domain of the impact of radiative heat transfers on flame dynamics", Gicquel concludes. ■

Computer simulation is at the

HEART OF computational materials science

Risto Pitkänen

The specialist field of Academy Professor Risto Nieminen is nanoscience and nanotechnology, or the study of how atoms can be arranged to create interesting properties in materials. In principle, there is an infinite number of combinations, and it would be impossible to study them all empirically in the laboratory.

Risto Nieminen directs the Computational Nanoscience center of excellence at the Laboratory of Physics, Helsinki University of Technology (HUT). The group's name is abbreviated to COMP, which is appropriate, since computing is at the very heart of the research.

Nanoworld systems are extremely complex and nonlinear. It takes modeling and computer simulation to sieve out the items of interest. In this way the number of possible research targets is reduced and research is guided into the right direction.
As opposed to real-world intuition, nanoscience research moves around in the quantum world, but it does not glide into the virtual world. Rather, the research aims to gain something concrete in the physical world. The role of modeling and theoretical research is to survey for possibilities and to guide which way to go.

Computing is also needed for interpreting the results of characterization. Measuring instruments and devices used in empirical research usually produce indirect data on the nanostructures being studied, and interpretation of the data requires powerful computation and theory.
At the borderline between empirical and theoretical research, one of the essential roles of computation is signal and image processing, which converts measurement data into a form that can be interpreted.

Disappearing barriers

According to Nieminen, the traditional borderlines are blurred in computational materials science. Active communication between empirical and computational researchers is of vital importance for both parties. Also the traditional barriers between different disciplines are disappearing. In nanoresearch, whether the researcher is a physicist, chemist, or biologist, has more to do with educational background than with what the scientist is researching.

Furthermore, the difference between basic research and applied science, rather than representing a true vision, is more a reflection of the attitudes of the scientists. The COMP center of excellence led by Nieminen is financed through funds granted for its purely basic research, but the unit also works in collaboration with industry.

"We try to guide young researchers to think that this is not a black-and-white setting. Good research problems may arise from the hottest topics of industrial interests, and yet some of these interests, especially high-tech industrial companies, are keen enough to be involved in the purest basic research that has no immediate applications."

Nieminen's unit has supported the qualification of several doctors who have defended their doctoral theses on pure basic research. Many of them now work as successful entrepreneurs benefiting from ideas that came up as part of their research.

Is this the next megatrend?

A research unit of this size contains numerous groups and even more numerous projects. Nieminen picks three research fields that he feels closest to: ferroelectric nanomaterials, carbon nanotubes and semiconductor materials. HUT has very active experimental, theoretical and computational research into carbon nanotubes. They have coordinated some projects using DEISA resources, in which researchers from other EU countries have participated.

Ferroelectricity was discovered under laboratory conditions as early as in the 1950s. Controlled manufacturing of ferroelectric materials is extremely difficult. Utilizing them has gained speed only during the past ten years, after the discovery of how to handle them at the nanoscale. (One nanometer equals one-millionth of a millimeter.) Advances in empirical research have provided a driving force. Numerous questions and answers, and theory and computation are needed to help in decision-making and to guide research and materials technology in the right direction.

Ferroelectric materials have interesting applications. These materials can sense their environment and they can be used to build sensors that detect ambient temperature, humidity, electric fields, gases, and other properties. According to Nieminen, we are dealing with superior megatrends in science and technology development. The past few decades have represented the triumph of the digital world, of packing and transmitting digital information. The next megatrend might be combining the digital and analogical worlds. In simple terms, we would like to make smarter products.

"Smartness, in this context, refers only to the product being able to detect its environment and communicating with a machine. For this purpose we need to develop "functional" materials, which is a big challenge also for the research and development of nanomaterials."

The promising world of quanta

Nanoresearch represents applied quantum physics. For a long time, researchers have foreseen that quantum-world properties, contradictory to real-world intuition, will be utilized in information processing. The visionary concepts of quantum computers, quantum cryptography and quantum information are founded on the laws of quantum mechanics being parallel or superposition laws. In the quantum world, all things happen simultaneously. In macroscopic objects, information leaks, disappears, decoheres and is averaged. The information unit in current computers is a tiny semiconductor transistor that contains millions of atoms and information is averaged as either 1 or 0. The quantum-world bit is >>>

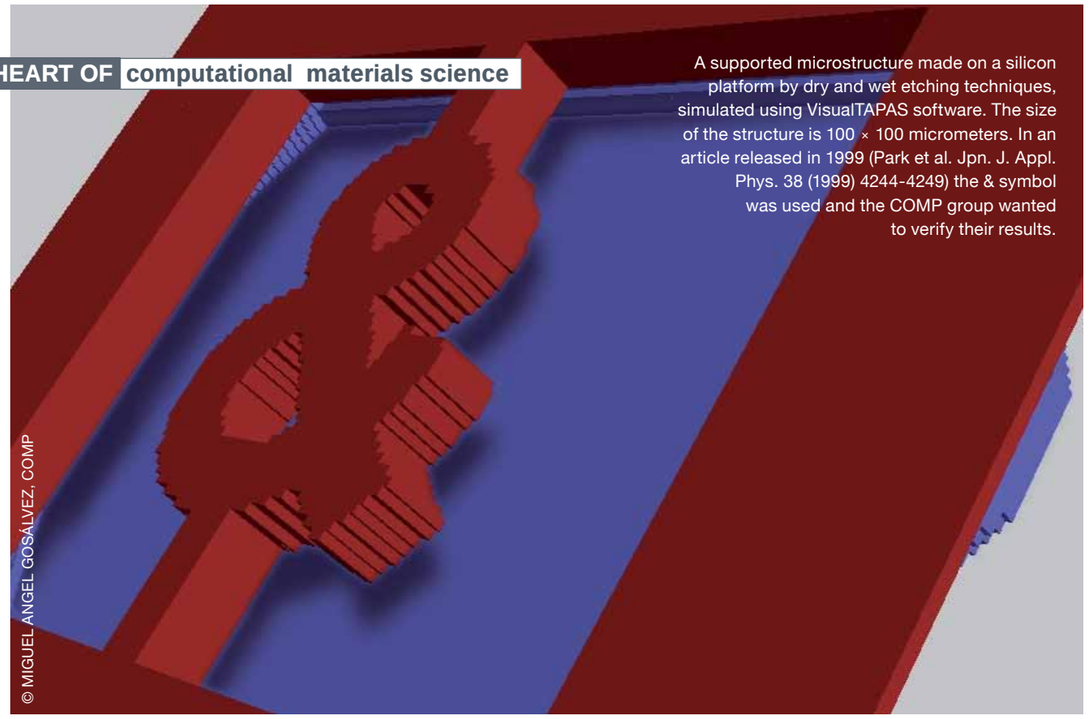

A supported microstructure made on a silicon platform by dry and wet etching techniques, simulated using VisualTAPAS software. The size of the structure is 100 × 100 micrometers. In an article released in 1999 (Park et al. Jpn. J. Appl. Phys. 38 (1999) 4244-4249) the & symbol was used and the COMP group wanted to verify their results.

© MIGUEL ANGEL GOSÁLVEZ, COMP

a qubit; it can be any well-defined number between 0 and 1.

The visionary quantum computer would require us to convert individual quantum-mechanically behaving objects into qubits, in other words, we would need to be able to take information in and out of them. Getting information out from the quantum world into the world of classical physics is not the biggest problem. However, unfortunately a major part of the information is lost, so we cannot record and utilize all possibilities offered by the quantum information. But a bigger problem is a phenomenon called “decoherence”, which means that we need to find ways how to stop the environment from disturbing the quantum computer at too early a stage.

The power of the quantum computer would be based on the fact that it can simultaneously follow numerous paths. This would be substantially beneficial for solving some problems. Breaking encryption algorithms is the best-known concrete example. A task that would take years or decades for the current computers to solve, could be solved as a matter of moments. On the other hand, there are also several computations where a quantum computer would not be any faster than an ordinary computer. At this point, we are still mainly dealing with basic research, but Nieminen believes that the first cryptographic applications are to be expected soon, but actual quantum computation is still decades away.

Is the result right or wrong?

How can you be sure that your computation results are correct? Of course, the classical way is to compare the calculations with test results, the virtual world against the real world. If a materials scientist is asked to calculate, say, a chemical process and determine the temperature at which a certain reaction takes place, the person asking the question would not appreciate an answer that the reaction takes place at 500 plus/minus 200 degrees. He would want the answer to be within an accuracy of one or two degrees.

The computation model is usually tested by using it to calculate a known problem and then checking that the result is as expected. All numerical work, says Nieminen, must always be backed up with an error analysis and therefore, it is important that the codes used do not contain black boxes.

“Every researcher must be able to see into the code and understand which programming languages have been used to compile it and which are the algorithms obtained.”

A numerically correct result may still be wrong, if the theoretical model used is wrong, i.e. it does not describe the chemistry or physics being studied. It is an essential part of computational research to keep making adjustments to the qualification limits of models.

The scientific community has an internal control system that brings the best out of every researcher. Nieminen is annoyed by the fact that many manuscripts published in the branch of computational research do not reveal how the work has been accomplished. The very basic principle in scientific work is to ensure that anyone can verify the result by repeating the study. From personal experience, Nieminen knows that a researcher can be relatively sure he is not the only one calculating on the topic in question.

“You should never be lulled into thinking that no one will ever repeat this computation, considering it too exotic an item. This is the way science works. It arouses trust in the results and it places a burden and responsibility on the scientist to ensure that what he allows to be published really is correct.” ■

Impurities on insulating surfaces and in carbon nanostructures

The research of Professor **Risto Nieminen’s** unit using DEISA resources 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 the research group considered a prototypical system, metal atoms on alkali halide surfaces, and considered how the bonding and diffusion of the atoms depends on their charge. They 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.

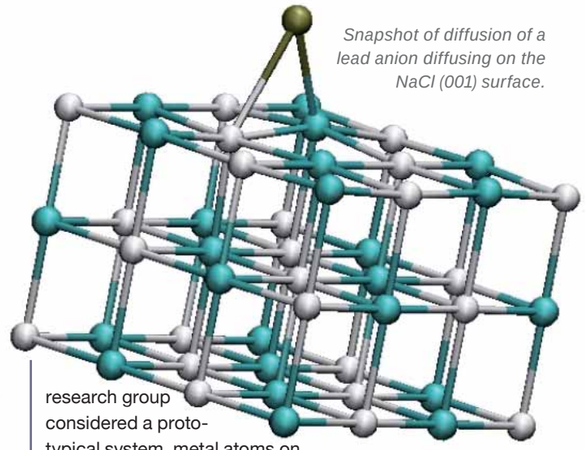

Snapshot of diffusion of a lead anion diffusing on the NaCl (001) surface.

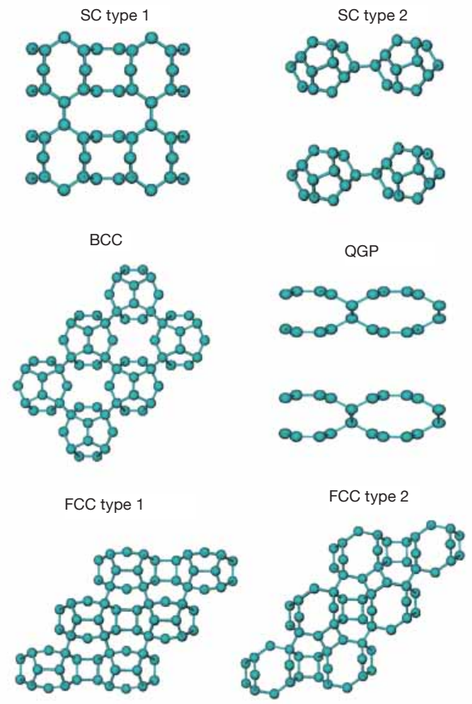

Representation of C20 based 3D structures considered within the project: simple cubic (SC), body-centred cubic (BCC), and face-centred cubic (FCC).

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, the research group investigated the role of iron adatoms on the magnetism of a (4,4) semiconducting carbon nanotube. They considered the coupling of carbon adatoms as a function of separation, and presence and location of iron adatoms. They 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.

Fullerene structures as a novel material for hydrogen storage

In this part of the project the research group used 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 the research group has 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. ■

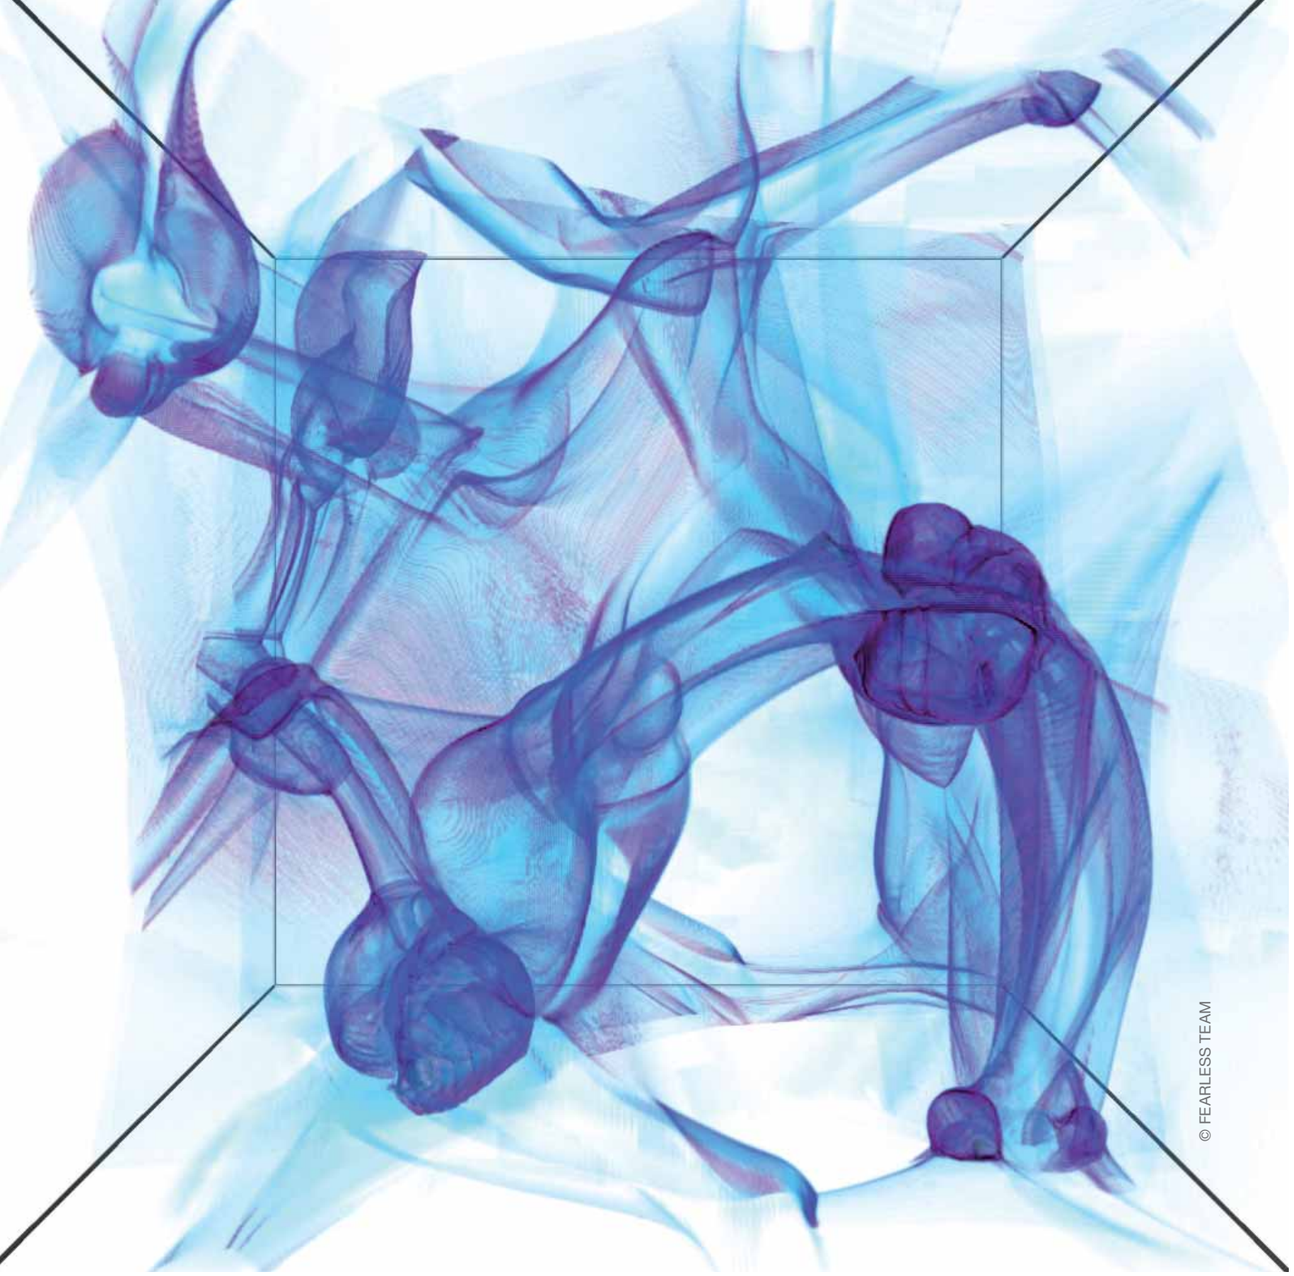

© FEARLESS TEAM

Astrophysical fluid mechanics:
A new method for simulating supersonic turbulence

Damien Lecarpentier

The FEARLESS project was initiated in 2005 by Jens Niemeyer and Wolfram Schmidt, two astrophysicists from the University of Würzburg in Germany. Using DEISA's computational resources within the DECI framework, the project team has developed a new method for simulating turbulent fluids, which will open up new perspectives in the field.

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

Right: In this image, the shock fronts are distorted and broken up into an intricate network of smaller shocklets and vortices as the production of turbulence progresses.

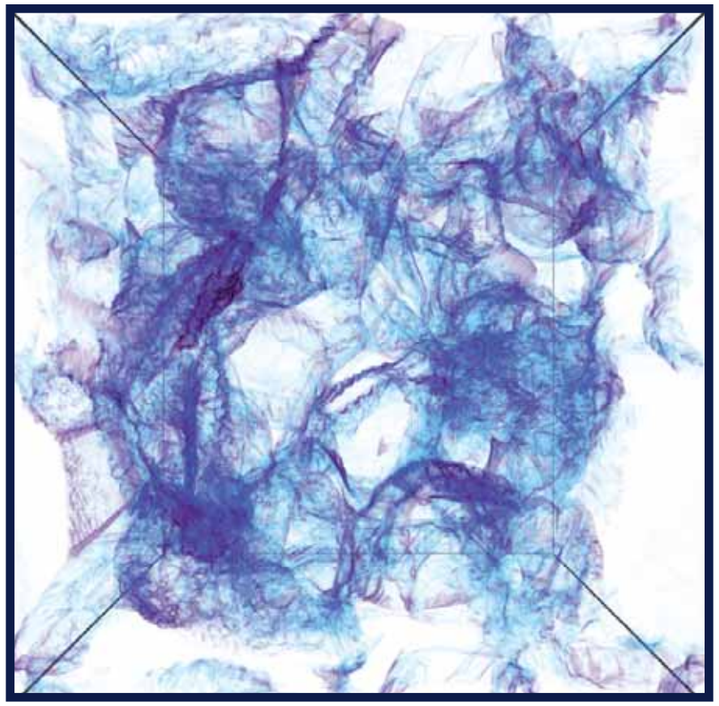

© FEARLESS TEAM

"Fluid mechanics is central to the understanding of many astrophysical phenomena", says **Wolfram Schmidt**, one of the two architects of the FEARLESS project.

"Turbulence, in particular, is worth studying, because of the fundamental role that it plays in the interior of stars, the formation of stars and planets, the evolution of spiral galaxies and, probably, the dynamics of the hot gas that surrounds clusters of galaxies", he explains.

When the FEARLESS project began in 2005, Wolfram Schmidt and his colleague were working on numerical simulations of thermonuclear supernovae – extremely powerful explosions of white dwarf stars.

"In these simulations, we implemented a model for the influence of turbulence on the speed of thermonuclear burning. It occured to us that such a model might be also useful for other applications in astrophysics. The central element of the FEARLESS project was therefore the use of advanced numerical techniques in combination with our turbulence model", explains Schmidt.

Combining models with advanced numerical techniques: Improved simulation of turbulent fluids

Turbulence is frequently modelled in engineering applications and atmospheric sciences by Large Eddy Simulations (LES). In this type of model, the dynamics of turbulent eddies are computed on large scales, while a subgrid scale model approximates the influence of smaller eddies.

"However", Schmidt continues, "in astrophysics, phenomena such as self-gravitating supersonic turbulence in star-forming gas clouds challenge the LES approach, because gravity and thermal processes (such as cooling) break the scale-invariance employed in LES over a wide range of scales."

In order to overcome this problem, the project team used a method called Adaptive Mesh Refinement (AMR): "This method involves inserting computational grids of higher resolution into turbulent flow regions in which strong shock fronts are forming, and the gas is undergoing a process of gravitational collapse", he explains.

Due to the extreme range of different length scales, however, it is generally impossible to treat fully developed turbulence by means of AMR only. This would require too large a number of refined grids. For this reason, the team has developed a new method that combines AMR with a subgrid scale model that links the notions of AMR and LES.

"FEARLESS stands for Fluid mEchanis with Adaptively Refine Large Eddy SimulationS", says Schmidt. "This somewhat complicated title captures the major elements of our concept: We intend to carry out simulations of turbulent fluids using a method that adapts dynamically to the simulated flow by refining the computation in those regions in which turbulence is developing", he explains.

Supercomputing resources and simulations of supersonic turbulence

Supercomputing resources are very much required in order to apply this method and the DEISA infrastructure has played an important role in the development of the project.

According to Schmidt, "Supercomputing resources have been essential for the development of FEARLESS. Using the DEISA infrastructure, elaborate simulations were carried out in order to test the method and to calibrate the turbulence model." >>>

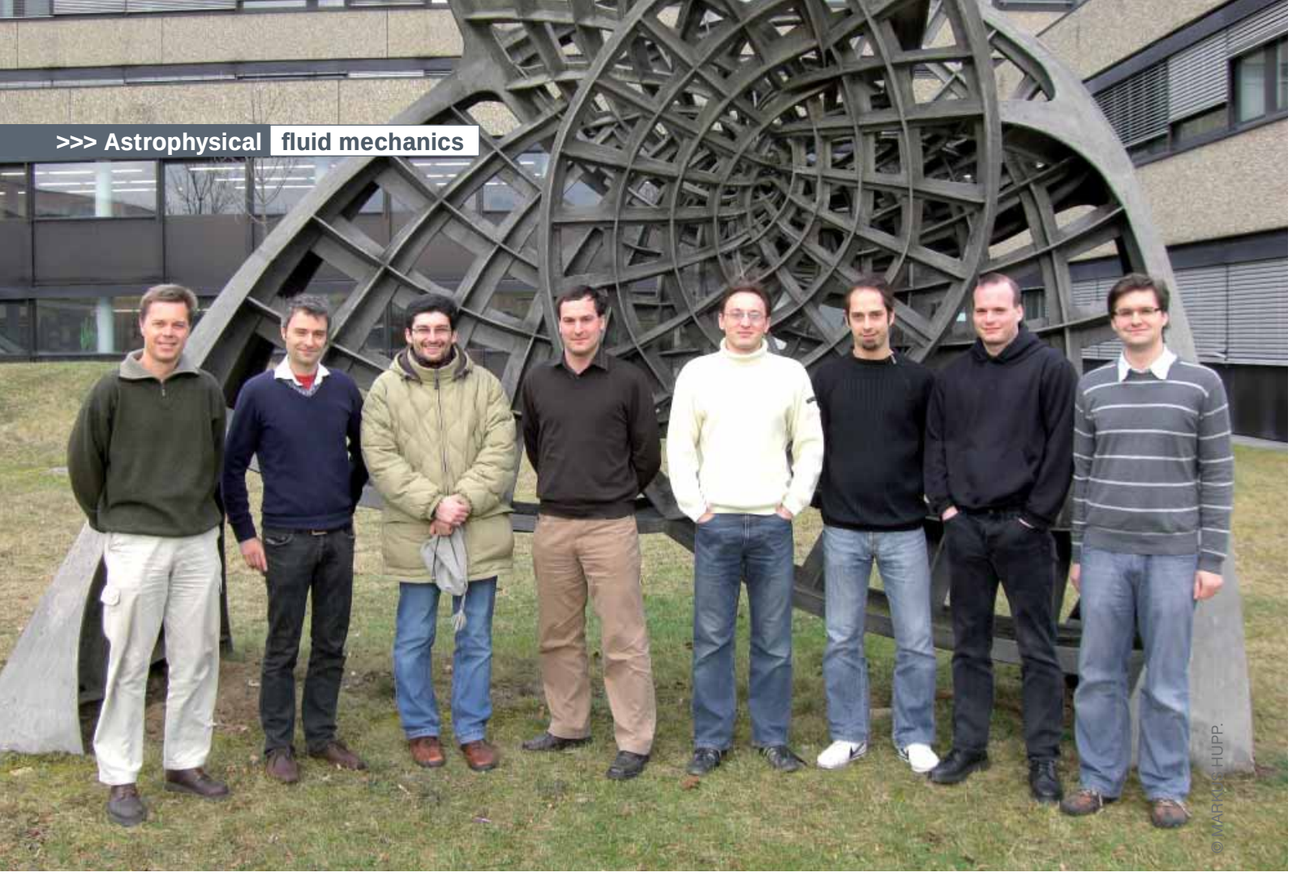

Above: Wolfram Schmidt (second from left) and his colleagues from the FEARLESS team.

"We were granted 300 000 CPU hours for the project. We used all of these computational resources to perform AMR simulations of supersonic turbulence with 126 CPUs in parallel", he adds.

"Depending on the size of the computational grid, 16 to 126 CPUs of the SGI Altix supercomputer (at SARA, in the Netherlands) were required for each simulation. The basic idea was to trigger the refinement by monitoring flow properties such as the vorticity (the rotation of the velocity field) and the rate of gas compression (which varies due to shocks or gravity)."

An extensive statistical analysis has also been carried out by the team, the results of which accorded well with the known properties of turbulence inferred from simulations without AMR.

"This makes the FEARLESS project team confident that AMR can be carried over to more complexe scenarios involving thermal and chemical process, as well as self-gravity", Schmidt concludes.

First steps towards future developments

The numerical study carried out within the DECI framework has laid the foundations for the approach that the project team is seeking to develop.

"The DEISA infrastructure allowed us to perform numerical simulations that were useful in terms of methodological advancement in the early stages of the FEARLESS project", says Schmidt.

"Moreover, results from these simulations are significant for ongoing research into the nature of turbulence in star-forming gas clouds in the Galaxy", he adds.

"FEARLESS is now fully operational. At present, we investigate the role of contemporary star formation in disk galaxies such as the Milky Way and the formation of galaxy clusters in the early history of the Universe. For example, some recently obtained results have demonstrated that the gas in the centre of massive galaxy clusters is significantly more dispersed – due to the effects of small-scale turbulence included in FEARLESS – than was previously thought to be the case."

"Further development of the FEARLESS turbulence model is however required, as it does not yet capture all regimes encountered in astrophysical applications."

"Upcoming applications will include more advanced models of star formation in the turbulent interstellar medium and the feedback from star formation onto the evolution of spiral galaxies. The project team also expects that FEARLESS will generate new perspectives in astrophysics throught the as yet unrivalled levels of sophistication it achieves in the treatment of turbulence. Moreover, methods similar to that adopted in FEARLESS might prove very useful in certain engineering applications, or even in climate research. For example, we have started a collaboration with Emmanuel Lévêque at the École normale supérieure de Lyon, who is a renowned expert on LES working both with engineers and astrophysicists", Schmidt concludes. ■

Federico Toschi's research group at the National Research Council (CNR), Italy, received their by-far largest supercomputing resources from DEISA to study 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.

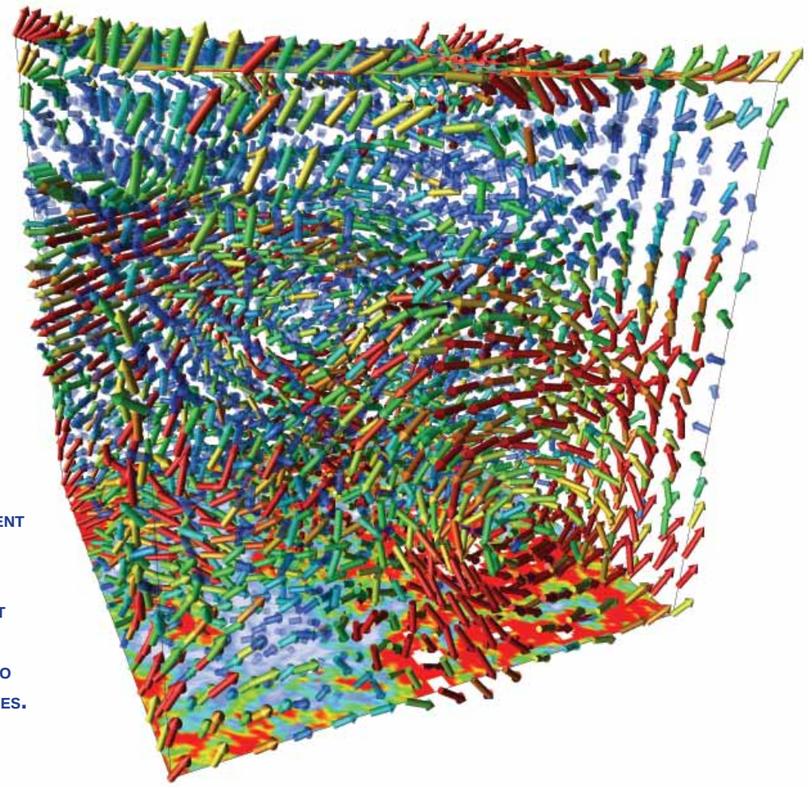

Passage of a droplet

Saara Värttö

"Turbulence is a chaotic state of fluid motion. Not only is the fluid motion complicated, but also the flow. You may change slightly the starting configuration of your flow and after a while the flow configuration is completely different", says **Federico Toschi**. He explains that it is impossible to know the exact starting configuration for the velocity field with infinite accuracy in any experimental or natural situation. It is a property of turbulence that even a tiny error in the knowledge of the starting configuration gets amplified exponentially. "This means that in practice you are not able to describe in a deterministic way the configuration of your system. Therefore a turbulent flow necessarily needs to be described statistically."

Many natural systems are turbulent: they behave in a very complicated, chaotic way with a lot of vortices. "Vorticity is a major feature in dynamic turbulence", continues Toschi. "It can be present from very large scales, like in oceans, all the way down to very small vortex filaments, which are the smallest structure in turbulent flows."

Above: Trajectories of particles with different inertia, released from the same spatial position. The red curve is a particle with no inertia (fluid element, hence following the fluid material line); the yellow curve is a particle with the largest inertia in our simulation. As mass grows, particles are less and less sensitive to small scale vorticity and to rapid velocity variations; particles with high inertia are expelled from high vorticity regions.

Fluids are everywhere and they can transport particles. "There are many possible examples such as the formation of rain droplets in clouds: At the beginning droplets are very small, and then they keep growing on top of condensation nuclei. When they are large enough, they can even collide with other droplets to form larger droplets until they are too heavy. Then they fall and hence the rain begins."

Toschi's HEAVY project run on the DEISA infrastructure addresses this important question to determine the properties of particles transported by a flow, when the flow is turbulent.

>>>

>>> Passage of a droplet

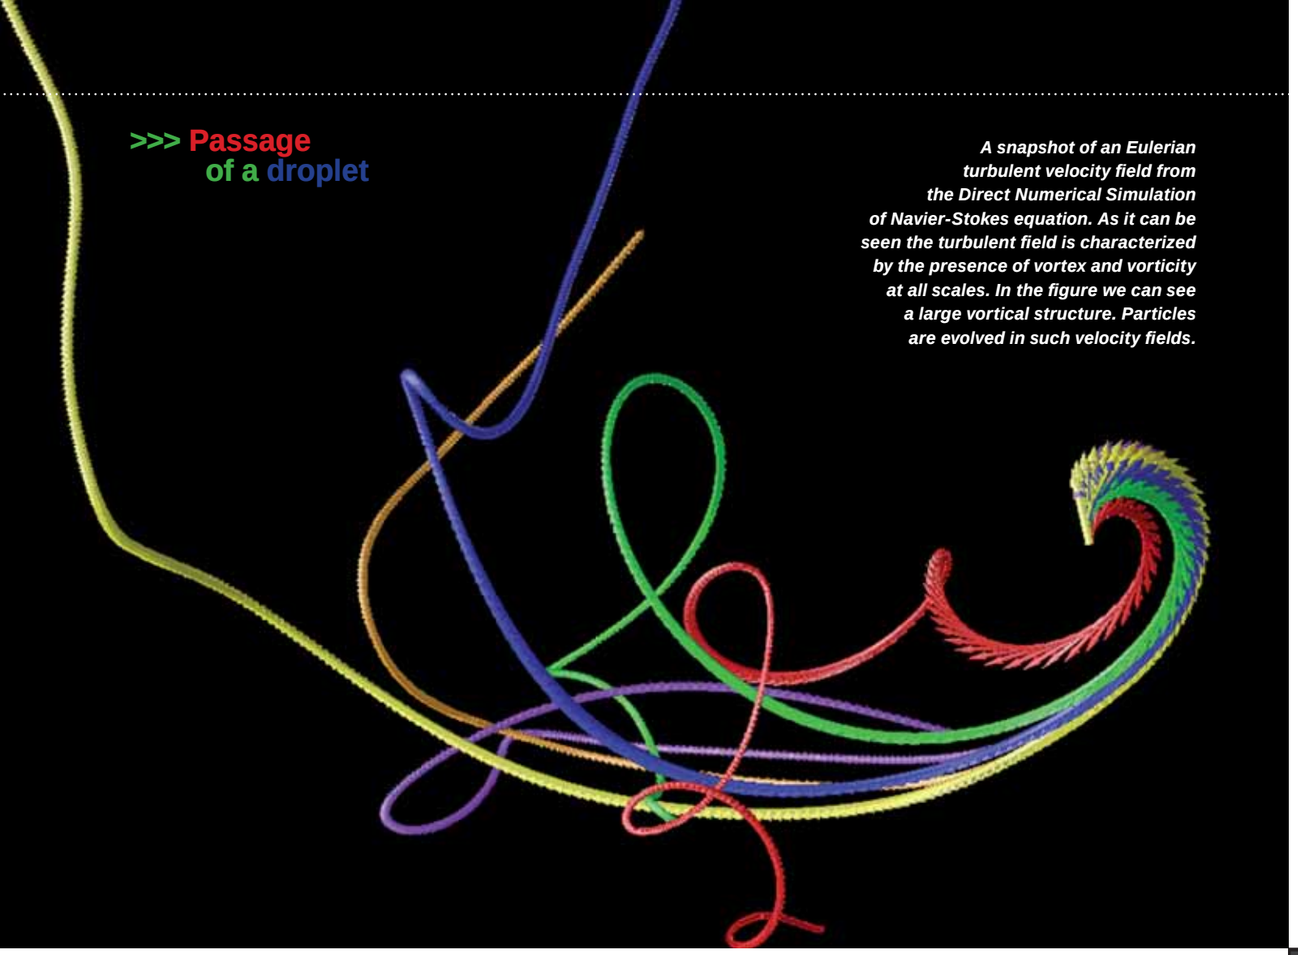

A snapshot of an Eulerian turbulent velocity field from the Direct Numerical Simulation of Navier-Stokes equation. As it can be seen the turbulent field is characterized by the presence of vortex and vorticity at all scales. In the figure we can see a large vortical structure. Particles are evolved in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

"I wanted to study very accurately the statistical properties of particles transported in a turbulent field, in its simplest realization – homogeneous and isotropic turbulence."

Toschi explains that the results of the numerical simulation have already been compared with some state-of-the-art experiments. Also in these experiments the aim is to have a flow field that is the simplest possible – i.e. homogeneous and isotropic. "Inside this turbulent flow experimentalists are tracking the evolution of particles, and a part of our work is to see and validate to which extent we are able to compare our results with experiments", says Toschi.

"These experiments are very challenging, because you have to follow the individual trajectory of particles and that is very complicated. By comparing the results of the experiments with the dynamics of neutral tracers, one is also able to validate the accuracy of the experiment. Vice-versa experiments can be used to validate the model used for the transport of particles."

Possible fields of application

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. "Dispersion of pollutants in a city is one of the possible fields of application", continues Toschi. "Pollutants get transported by wind and its turbulent fluctuations. Of course there is a great deal of factors you will have to take into account, such as the geometry of a city: streets, buildings and boundary layers."

Toschi's study is an idealised situation, where only the basic physical properties of particles transported by turbulence are observed. "However this is already a piece of information and may allow predicting which area of a city can be affected and to which extend."

Another possible field of application is combustion. "The fuel is spread as very small droplets into the combustion chambers. Again there are many complications due to on-going physical and chemical phenomena", says Toschi. "But one of the phenomena is the dynamics of droplets in a turbulent environment. By having a better knowledge of the dynamics of these particles one could for example enhance the efficiency of combustion and diminish the produced pollution."

Reynolds numbers and the need for more powerful supercomputers

Although supercomputers are getting more and more effective, the lack of computing power still limits scientific discoveries. "The present capacity of the fastest supercomputers in the world is not enough in our field of study. The very large simulation now performed deals with the evolution of a particle in a turbulent field. In the turbulent field there are fluctuations from a very large scale to a smaller and smaller scale. There are strong fluctuations, for example, in vorticity", explains Toschi.

The ratio between the largest scale and the smallest scale in a system is connected to a dimensionless number, called the Reynolds number. "At the moment it is impossible to numerically integrate a turbulent flow with realistically large Reynolds numbers", adds Toschi. "For example the Reynolds number of the wind behind of an airplane can easily be in the order of 108. Today with the largest supercomputers it is possible to integrate a Reynolds number, which is at least three orders of magnitude smaller than that. Laboratory experiments can reach Reynolds numbers higher than in the numerical simulations."

This is important when studying a turbulent flow, because in a turbulent flow there are fluctuations all the way from the largest to the smallest scale, and these are coupled in a very complicated way.

Toschi concludes: "If you want to know the statistical properties of particles or of a pair of particles that get separated, or if you want to know how a concentration of pollutants gets dispersed by turbulence, a realistic Reynolds number is a must."

Sharing data

The numerical simulation of the HEAVY project was concluded in April 2007. Now Federico Toschi and his research grou are analysing the results. All the data will be made available for the entire science community in order to enable new research with different perspectives and interests.

The run was concluded in April 2007. Since then the data has been re-organised for future analysis. Toschi's group has started comparing data of numerical simulations with the data from previous experiments. The idea is to gather together findings of different numerical simulations and experiments on this same physical phenomenon. Toschi tells that as soon as they have finished this step of analysis all the data will be made available for the entire science community in order to enable new research with different perspectives and interests. The data will be stored at the International Computational Fluid Dynamics Database, iCFDdatabase (http://cfd.cineca.it) hosted by CINECA, Supercomputing Centre in Bologna. This is very important since the simulations are very demanding from the computational point of view, of the order of half million computing hours. Opportunity to run this kind of simulations is rare. Therefore Toschi's group allow other researcher to exploit all the data.

In the future Toschi plans to study particles that are lighter than the flow, such as air bubbles in a flow. The goal is to consider more and more realistic situations, for example taking into account the fact that a particle may have finite sizes. The fluid dynamics of this kind of particles can be different than that of very small particles. Toschi sees the collision of particles as a very important and interesting research topic, too. ■

Simulations in the HEAVY project

In the HEAVY project a pseudo-spectral code was used with a three dimensional cartesian grid made of 20 483 collocation points. On each of these collocation points the three components of the velocity field numerically integrated the Navier-Stokes equations.

The simulation was performed on 512 processors at Leibniz-Rechenzentrum (LRZ), München, Germany. The computational volume was divided in two slabs. Each one of the 512 slabs was assigned to a different processor.

The major part of the computational effort was spent in evolving the background turbulent velocity field. On top of this there was the numerical integration of particles dynamics. The total number of particles was slightly more than two billion, and there were 16 different kinds of particles.

One particle type was a neutral tracer, a particle just following the flow. Particles had small inertia, medium inertia; and a few of them had a very large inertia. The larger the inertia is, the longer the response time. When a turbulent fluctuation happens in some time scale, a particle with too large inertia is not able to feel the turbulent fluctuation, but it reaches out of this fast turbulent fluctuation.

This has a lot of implications, since particles with different inertia average or reach out the fluctuation of underlying turbulent field and this implies different dynamics. They do not distribute anymore uniformly in the volume, but start to prefer some region to others. In particular, particles with inertia tend to be expelled from regions that have high vorticity.

To quantify particles' preferred location, we will use numerical simulations, measurements and experiments and the proposed theoretical model has to be compared with these data. When integrating the evolution of two billion particles, researchers have to keep track of the position and the velocity of these particles along the trajectory.

This means that information of each particle should be written to a disc, and information of two billion particles fills at once any possible disc space. Toschi's group opted for a compromise writing information on each particle every 10 time steps while turbulence time scale was 1000 steps, i.e. still very frequent with respect to the physical dynamics of the turbulent field.

A compromise was made on particles, too. There were 16 different classes and for each one of those there was a subset of particles, for which accurate time resolved trajectories were stored. Only seldom full snapshots with all 2 billion particles were made. ■

Scientists have long agreed that climate change could have a profound impact on the planet, from melting ice sheets and withering rainforests, to flash floods, droughts, loss of arable land, famines and mass migrations. Policymakers are still debating instead of taking mitigating action to limit greenhouse gases generated from human activities, the likely cause of recent climate change.

ESSENCE: A large model ensemble simulation

Christopher Lazou

Climate experts however, began ranking the most fragile and vulnerable regions on the planet, evaluating the danger of sudden and catastrophic collapse in the 21st century. Although scientists cannot be sure precisely when each region will reach the point of no return, their assessment warns it may already be too late to save Arctic sea ice and the Greenland ice sheet, which they regard as most immediately in peril. By some estimates, there will not be any sea ice in the summer months by the middle of this century and some models predict this could happen as early as 2030 or even earlier.

The atmospheric concentrations of CO_2, CH_4 and other so-called greenhouse-gases (GHG) responsible for global warming have increased rapidly since the beginning of the industrial revolution, leading to an increase of radiative forcing of 2.4 W/m^2 up to the year 2000. Presently, long-term climate change trends are observed everywhere on earth. For example, the global mean surface temperature has increased by 0.2° to 0.6°C during the 20th century. There has been also a widespread retreat of non-polar glaciers, and patterns of pressure and precipitation have changed [Summary for Policymakers, IPCC Fourth Assessment Report].

To determine the causal chain between the increase in radiative forcing and observed climate change, climate simulation models are essential. In the past just one or a few transient coupled climate simulations were performed for a given emission scenario due to the high computational demand of a simulation. The simulation allowed an assessment of the global-mean climate change, but because of the strong natural climate variability, it was not possible to attribute local changes to increased radiative forcing.

To distinguish trends caused by natural climate variability from those induced by increased radiative forcing, a large ensemble of climate simulations is necessary. By averaging over all ensemble members the signal-to-noise ratio is enhanced. The natural climate variability cancels between the members, leaving only the forced signal.

The Ensemble SimulationS of Extreme weather events under Nonlinear Climate change (ESSENCE) project was set up, under the DEISA Extreme Computing Initiative (DECI), to study climate trends. “The main aim of the ESSENCE project was to compute an adequate estimate of the statistics of natural 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”, says Dr. **Andreas Sterl**.

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.

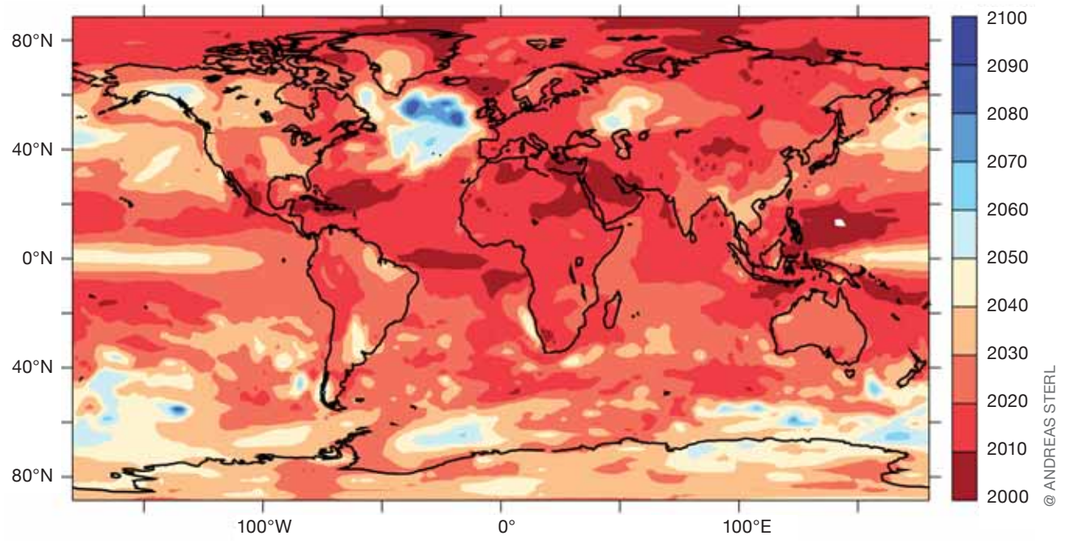

In the ESSENCE project, a 17-member ensemble simulation of climate was carried out using the ECHAM5/MPI-OM coupled climate model of the Max-Planck-Institute for Meteorology in Hamburg. It was chosen because it performed well on a number of criteria during an inter-comparison of models that were considered in the IPCC Fourth Assessment Report (AR4). All computer simulation runs, were driven by the A1b scenario of 21st century greenhouse gas concentrations, from the IPCC Special Report on Emission Scenarios (SRES). The A1b scenario assumes a growth of CO_2 that leads to a concentration of about 650 to 700 ppm by year 2100. Besides a basic ensemble of 17 runs three experimental ensembles were also performed. The baseline experimental period is 1950-2100. For the historical part of this period (1950-2000) the concentrations of greenhouse gases (GHG) and tropospheric sulfate >>>

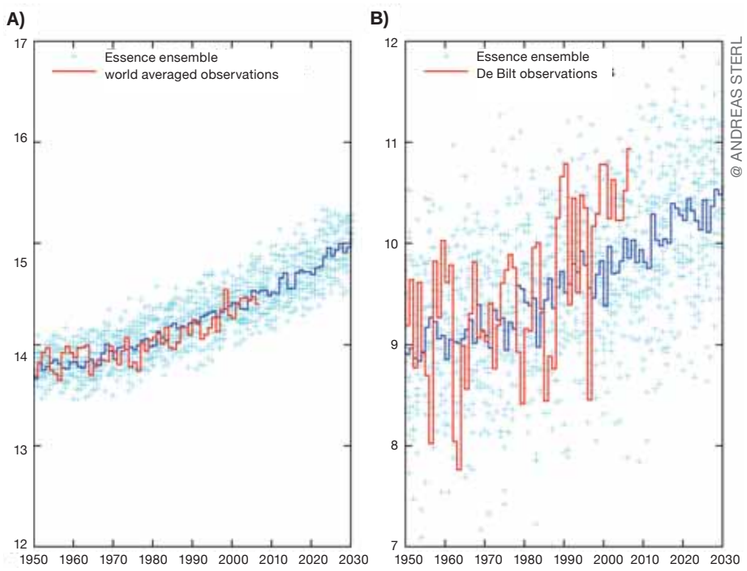

Annual-mean surface temperature for the 17 ensemble members (light blue crosses), their mean (blue line) and observations (red line) for (a) the global average and (b) station De Bilt (the Netherlands).

Dr. Andreas Sterl (right) discussing results from the ESSENCE project with his colleague Dr. Frank Selten (left) in front of a Tiled Panel Display (TPD) from SARA.

ensemble enables us to clearly distinguish the forced signal from natural climate variability. We compared observed and modelled near-surface temperature trends and determined the statistical robustness of projected future trends of some important surface variables. We showed that in large parts of the world observed and modelled warming trends are statistically indistinguishable. We are therefore confident that the model captures the most important driving forces of the climate well, especially the radiative forcing due to increasing greenhouse gas concentrations. Over large parts of Eurasia and Africa the warming signal will be distinguishable from the noise within the next 10 years", says Dr. Sterl.

A great advantage of a large ensemble is the large noise reduction that can be achieved by averaging over all ensemble members. From the above simulations they 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 for this study.

The earliest detection times are found off the equator in the western parts of the tropical oceans, where the signal emerges as early as year 2000 (and for some regions even earlier), from the noise. In these regions the natural climate variability is extremely low while the trend is only modest.

A second region with an early detection is the Arctic, where the trend is very large due to decrease of the sea-ice. The longest detection times are found along the equatorial Pacific where, due to El Niño, the variability is very high, as well as in the Southern Ocean and the North Atlantic, where the trend is very low.

By using a large ensemble under this DECI project the climate change trend studied above delivered invaluable results making it possible for policymakers to initiate informed mitigation actions.

DEISA resources used

Early 2005, the first call of the Distributed European Infrastructure for Supercomputer Applications (DEISA) Extreme Computing Initiative was launched and the CKO* was successful with the ESSENCE (Ensemble SimulationS of Extreme events under Non-linear Climate changE) project. For this project, 170 000 CPU hours were provided on the NEC-SX8 at the High Performance Computing Centre Stuttgart (HLRS).

The DECI-DEISA 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 50TB 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. ■

Summary for Policymakers,
IPCC Fourth Assessment Report,
http://www.ipcc.ch/SPM2feb07.pdf

** CKO is a co-operation between KNMI (Royal Netherlands Meteorological Institute) and IMAU (Institute of Marine and Atmospheric Research Utrecht of Utrecht University).*

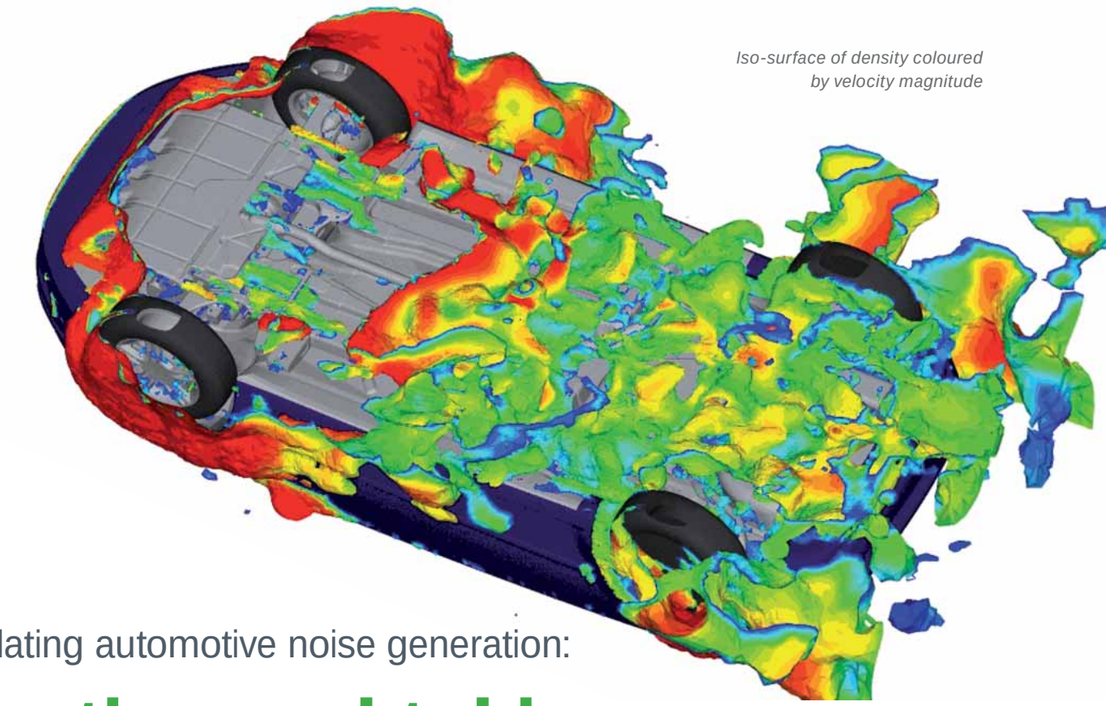

Iso-surface of density coloured by velocity magnitude

Simulating automotive noise generation:

Testing and taking advantage of the DEISA Research infrastructure in an industrial context

Damien Lecarpentier

Using the DEISA Research infrastructure within the Joint Research Activity (JRA5) framework, the Fiat Research Center (CRF) – in Italy – carried out a number of different computational fluid dynamics and aeroacoustic simulations in the important area of automotive noise generation. The project allowed the Center to test the supercomputing infrastructure in order to determine whether industrial problems can be resolved by high quality simulations performed within real time constrains. This experiment reinforced the competitiveness of the group and illustrates the benefits of partnerships between research infrastructures and industry.

"Noise reduction is an important concern in terms of human well-being and health, and is a significant issue for a wide range of applications, including for the automotive industry", says **Roberto Tregnago**, researcher at CRF. "Considerable improvements have been made in automotive acoustics during the past few years but some progress still needs to be made in order to satisfy the increasing demand for quieter vehicles. The research that we are doing at CRF focuses on different types of noise generation, and, in particular, on aerodynamic noise, which results from the interaction between the airflow and the moving vehicle. Aerodynamic noise increases with the speed of the vehicle, and becomes dominant above 100 km/h", he explains.

Computational simulation: an essential tool for studying noise generation and optimizing industrial production

Computer simulations have become increasingly useful for the study of noise generation, as they have the potential to partially substitute numerical simulations for actual physical experiments, thus making the development of the >>>

The simulation was realised on a Fiat Grande Punto.

industrial product in question both considerably faster and cheaper. Computational fluid dynamics (CFD) and computational aeroacoustics (CAA) are the two main numerical applications used by researchers to simulate noise generation.

"These two applications are today widely used in the engineering field, particularly in the aerospace and automotive industries", says Tregnago. "They enable us to study the features of fluid dynamics that contribute to noise generation, and to develop new ways of reducing noise pollution. By using CAA, the aeroacoustic properties of a new car design can be adequately simulated. Several layout alternatives can be numerically tested, a practice that is at the moment unfeasible in physical testing. Computer simulations enable us to study and predict the behaviour and performances of new products at an early design stage, thus reducing considerably development time and costs."

"However", Tregnago continues, "these applications are highly demanding in terms of computing power, particularly where simulated details and expected results are pushed to the level required for new industrial product development, and given the short time-period characteristic of the modern design process, within which they must be delivered. That is why a partnership with a research infrastructure, such as DEISA, can be of genuine interest to automotive industries."

Raising industrial simulation capabilities through the partnership with DEISA

The Joint Research Activity JRA5, which focused on computational fluid dynamics (CFD) and computational aeroacoustics (CAA), provided the framework for such a partnership.

"CRF joined the DEISA project in order to explore – and take advantage of – the possibilities offered by a distributed computational grid in resolving engineering problems in the area of noise generation. Our main objective was, in particular, to test the DEISA infrastructure in an industrial context, and to push back the boundaries of numerical simulation for industrial applications", says Tregnago.

"Practically, we defined four relevant test cases related to simulations that are very important for the automotive design process: The first test was a shape-noise aeroacoustic simulation on a detailed car body; the second focused on open sunroof buffeting; the third on noise generation and propagation inside ventilation ducts; and the final test dealt with shape optimisation of car bodies. Each test case was scheduled with four degrees of increasing complexity, both in terms of model dimension and time/frequency resolution."

"The simulation was realised on a Fiat Grande Punto. It was performed with a time step of 5.0×10^{-5} s in order to increase the quality of the pressure wave resolution. This corresponds to a Nyquist frequency of 10 000 Hz. The resolution level afforded by these high frequencies was necessary to achieve results of the requisite quality – high enough to be considered a genuine breakthrough in terms of industrial simulation capabilities, and, above all, to make simulation a viable alternative to physical testing."

Flat microphones on the underbody of Grande Punto in the wind tunnel.

"Computed time history was accumulated up to 0.6 seconds, with a global resource consumption of roughly 30 000 CPU hours (in DEISA normalized units). The numerical 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 that show the noise power spectra over the whole vehicle for a given frequency."

"In order to validate these numerical results, a number of experiments were also carried out in FIAT Auto's aeroacoustic wind tunnel. Twenty flat microphones were placed in the same positions as a subset of the numerical probes, in order to enable us to compare the experimental power spectra of the pressure signal at each location with that derived from the numerical simulation."

The simulation, according to Tregnago, was genuinely promising: "Globally, the comparison of the results of the simulation with those of the experiment shows them to be largely in accord. In particular, the simulation was able to predict – within acceptable limits – the suppression of the peak in the spectrum due to fluid-acoustic resonance that increases noise within the passenger compartment, and the level of accuracy reached has allowed us to propose some design solutions to the problem of noise emission."

"The high quality results obtained with these simulations enable us to move away from simple trend prediction (which only gives partial support to engineers during the design phase) towards providing detailed measurements within a short period of time (which enables effective decision making during the design process). The competitiveness of the company has been increased through the achievement of a streamlined design process and the improved performance of the final product", he notes.

Comparison of numerical and experimental results for point 20.

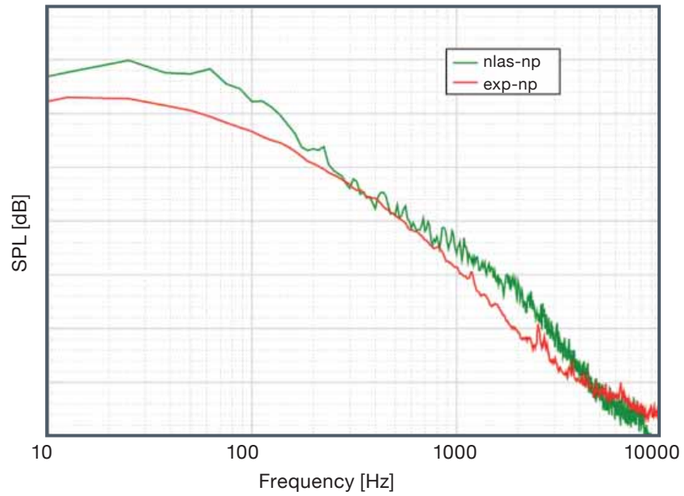

"The main problem that we encountered during the project was related to the licensing of commercial codes that are used for industrial CAA applications. Hopefully, in the near future, a complete solution – combining the availability of both supercomputing resources and code licenses – will be developed, in order to satisfy industrial needs."

"CRF's experience with the DEISA infrastructure has, however, demonstrated that appropriate numerical techniques, together with suitably high performance multiprocessor systems, do represent powerful tools for industry. They enable us to tackle problems that would be impossible to face with standard resources within the time constraints of a product development process. Thanks to the increasing power of supercomputers, it is likely that we will see, in the near future, comprehensive simulations of these designs that also encompass elements of multi-disciplinary optimisation." ■

Instantaneous acoustic pressure distribution on the underbody of the vehicle. The image depicts the computed acoustic pressure at a certain point in time on the surface of the vehicle underbody. In the red regions, the instantaneous fluctuating pressure is higher, while in the blue regions it is lower. This pressure distribution stimulates the underbody panels that generate noise within the cabin.

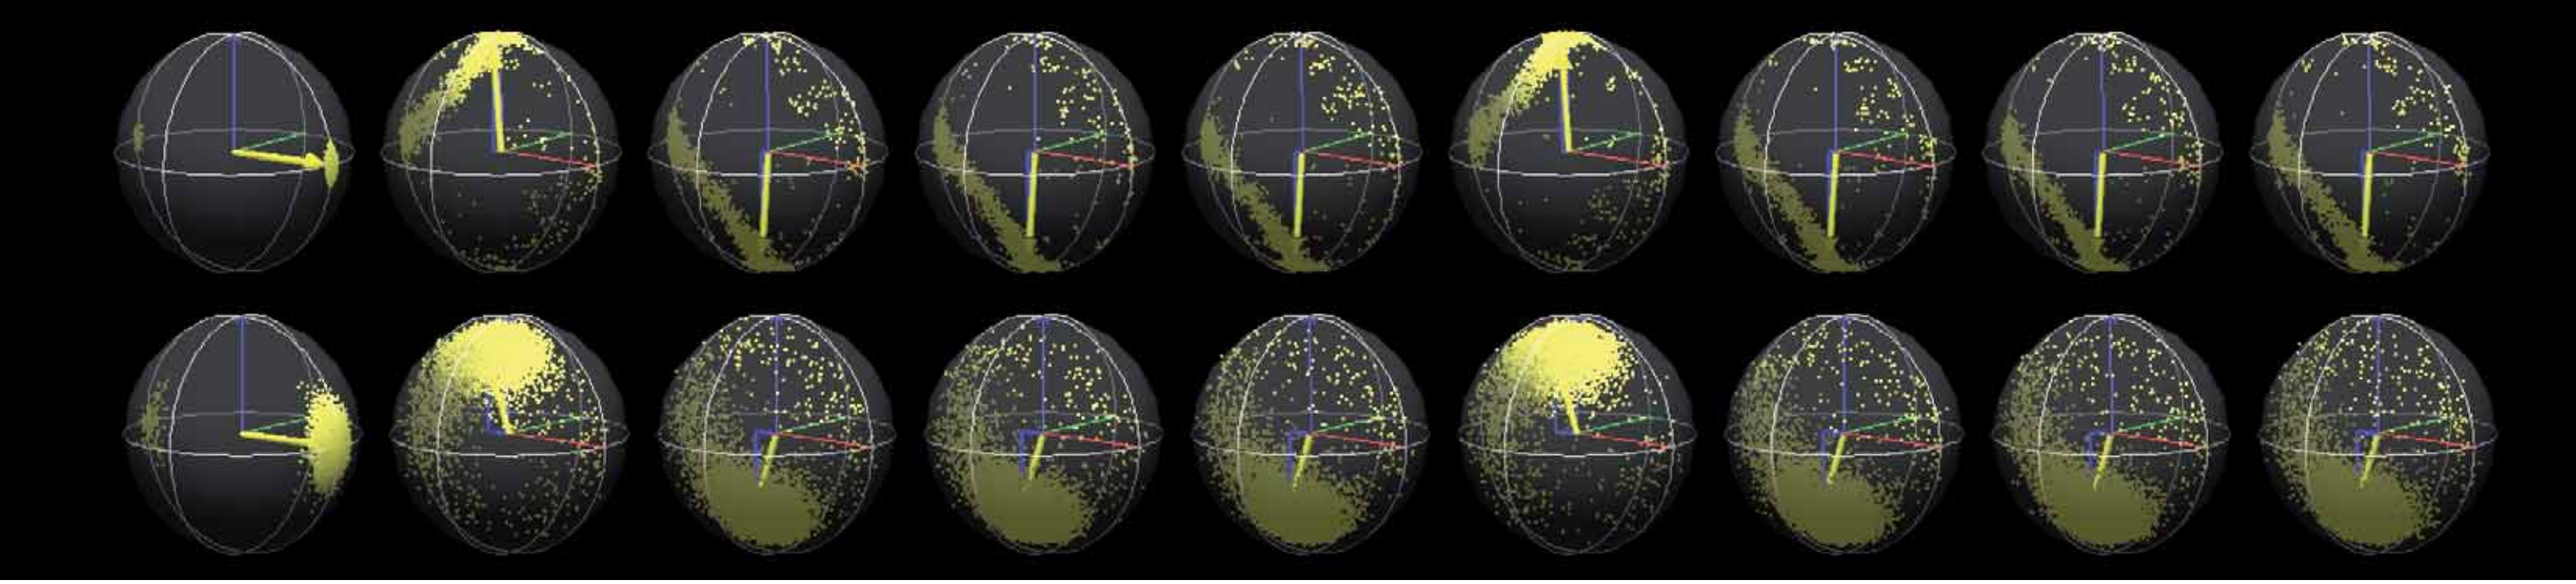

Above: Grover's quantum search algorithm: the 8 qubits encode the searched data base element k=17=10001 (the leftmost qubit is ancillary). The upper line corresponds to an undercritical error distribution while the parameter of the lower sequence is above the threshold.

Improving Quantum Computer Simulations:

Benefits and limits of a distributed supercomputing infrastructure

Damien Lecarpentier

Quantum computing research is one of the newest fields in computational science. Using DEISA's computational resources within the DECI framework, Marcus Richter and his colleagues from the Jülich Supercomputing Center in Germany, together with the Computational Science Group of the University of Groningen in the Netherlands, initiated the IQCS project in 2005, with the aim of improving quantum computer simulations: a first step towards constructing a realistic quantum computer in the near future.

Quantum computing research is a rapidly evolving field. Although still in their infancy, quantum computers already appear as promising devices for the future. Their theoretical ability to perform vast numbers of operations simultaneously – by using quantum bits (qubits) instead of the conventional bits – has the potential to solve certain complex algorithms, far faster than conventional computers. They could therefore be used, for example, to break data encryption codes or to search large databases.

The development of quantum computers is, however, hampered by the fault liability of the quantum bit – the problem of quantum decoherence.

Quantum decoherence: a major problem for quantum computing

"One of the hot topics in quantum computation is how to maintain quantum coherence during quantum operations", says **Marcus Richter**, one of the leaders of the IQCS project and researcher at the Jülich Supercomputing Centre.

"The loss of quantum coherence, which is called decoherence, results from interactions of the quantum bits with the surrounding environment. This leads to decoherence errors."

"This process is very hard to suppress, as the rate of decoherence grows rapidly with the increase in environmental degrees of freedom – the number of which is usually quite large."

Marcus Richter tested the robustness of quantum algorithms under operational errors and decoherence, a prerequisite for building useful quantum information processing devices.

"Therefore, detailed studies of mechanisms to defeat decoherence are a necessary requirement for practical quantum computation", he concludes.

The ability to simulate realistic models of quantum computation devices that interact with their environment is crucial for the successful realization of scalable quantum computers. Within the IQCS project, Marcus Richter and his team sought, in particular, to investigate the impact and the interplay of operational and decoherence errors on quantum algorithms:

"The aim of the project was to test the robustness of the most prominent quantum algorithms, such as Grover's search algorithm and the Quantum Fourier transform (which is the kernel of Shor's factorization algorithm) under operational errors and decoherence", says Richter.

"The goal was to identify error thresholds below which reliable results can be obtained despite gate imperfections and decoherence."

"We used the Massively Parallel Quantum Computer Simulator – a highly efficient parallel code for the simulation of complex quantum algorithms at gate level – as a starting point."

"In order to include operational errors and decoherence (the loss of quantum information due to interaction with the environment), the code has been expanded by an effective quantum error model", he explains.

Quantum computer simulations are memory bounded

Quantum computer simulations, carried out on supercomputers, require a large amount of computational resources.

"We have been granted 144 000 CPU hours under the DEISA Extreme Computing Initiative (DECI) to perform the very demanding stochastic simulations on system sizes up to 32 qubits", says Richter.

"Without the DECI grant, we would not have been able to carry out the simulations in such detail", he adds.

"As a result, we were able to identify error thresholds below which reliable quantum computation is feasible. These thresholds depend on the algorithm considered, as well as on the system size (number of qubits). We found, in particular, that the quantum Fourier circuit is more robust (that is, less error sensitive) to disturbance than Grover's algorithm on comparable system sizes."

"To be of practical use, quantum computers will require at least several tens of qubits and the ability to perform hundreds of gate operations. The identification of reliable error thresholds is therefore an indispensable prerequisite for building useful quantum information processing devices", Richter explains.

While these important achievements were made possible by the DECI framework, the development of quantum computer simulations remains limited, due to the lack of computer memory:

"Quantum computer simulations are memory bounded, since the size of the state vector grows exponentially with the number of qubits. Adding one further qubit means that one has to double the amount of memory used for the simulation. For example, the simulation of 37 qubits requires 3 TB of memory", says Richter.

"Since the state vector of the system is, in such a case, distributed over the whole memory, the communication between the different computation nodes limits the performance of the code. Highly optimized memory access and communication patterns are vital to efficient simulations."

"Unfortunately, given that the network connection between the different DEISA supercomputing centers is much slower than the internode communication, the IQCS project is at present not appropriate for distributed computation", Richter acknowledges.

Future developments

Despite such limitations, these simulations remain of crucial importance, and Richter's team has begun applying its simulator in new research projects.

"At the moment, we are focusing on the simulation of (fault tolerant) error correction schemes. In the next step we will use our massively parallel quantum computer simulator to numerically investigate the characteristics of different error correction schemes", says Richter.

"Another branch is the dynamical (i.e. time-dependent) simulation of ion-traps. At present, ion-traps are the most promising technology for building quantum information processing devices (due to the long coherence time of the trapped ions)."

"Our goal", he concludes, "is to guide experimentalists by carrying out realistic simulations in order to improve the quality of quantum computing devices. Computer simulations are a first step towards building a useful quantum computer in the near future." ■

A parallel FEM code for THERMO-HYDRO-MECHANICAL coupled problems in porous media

Christopher Lazou

Professor Olaf Kolditz, Helmholtz Centre for Environmental Research, Leipzig explains the potential of GeoSys/Rockflow software platform for grand challenge projects, such as those using the DEISA infrastructure.

© ANDRE KÜNZELMANN (UFZ)

"Many applied problems in geo-science require knowledge about complex interactions between thermal, hydraulic and mechanical processes in the subsurface. As a direct experimental investigation is often not possible, numerical simulation is a common approach. The numerical analysis of coupled thermo-hydro-mechanical (THM) problems is computationally very expensive and until recently, the applicability of existing codes was limited to simplified problems. Therefore, high-performance-computation (HPC) and the development of appropriate methods for visualisation of huge complex data sets is very important for advancement in numerical geo-system simulation", says Professor **Olaf Kolditz**.

In a recent report, Professor Kolditz and his colleagues, at Helmholtz Centre for Environmental Research – UFZ, Leipzig, explained how they developed a software platform using a novel implementation of a parallel finite element method (FEM) for the numerical analysis of coupled THM problems in porous media. In addition, for their current HPC efforts, a high-performance-visualization platform TESSIN was established at UFZ.

The FEM scheme is partitioned by a priori domain decomposition and then assigned to the available CPU nodes concurrently. Parallel running is achieved by simultaneously establishing the sub-domain mesh topology, synchronously assemble linear equation systems in sub-domains and obtaining the overall solution with a sub-domain linear solver (parallel BiCGStab method with Jacobi pre-conditioner).

The present parallel method is implemented in an object-oriented way using MPI instructions for inter-processor communication. The parallel code was successfully tested against other codes with a 2D example from the international DECOVALEX benchmarking project. The achieved speedup for a 3D extension of the test example on a CRAY XT3 computer at CSCS Switzerland (see figure on page 49) demonstrates the advantage of the new parallel code over serial schemes. Despite the progress in parallel THM computation on medium size clusters, the observed speedup behaviour reveals limitations of the present method particularly for large cluster configurations. Therefore, as a next step object-oriented data is to be tested on vector computers such as the NEC SX-8 to further improve the parallel scheme for use on supercomputers. Object-orientation with referencing to object data via pointer seems to be a bottleneck for vectorization efficiency on single processors.

Rationale

The numerical analysis of multi-field problems in porous media is an important subject for many geo-engineering tasks such as the management of geo-resources (e.g. engineering of geothermal, oil and gas reservoirs) as well as waste management (e.g. chemo-toxic, nuclear waste and CO_2 sequestration). >>>

A diagram showing the processes involved in turning data into knowledge to gain a better understanding of the interaction of coupled processes in environmental systems

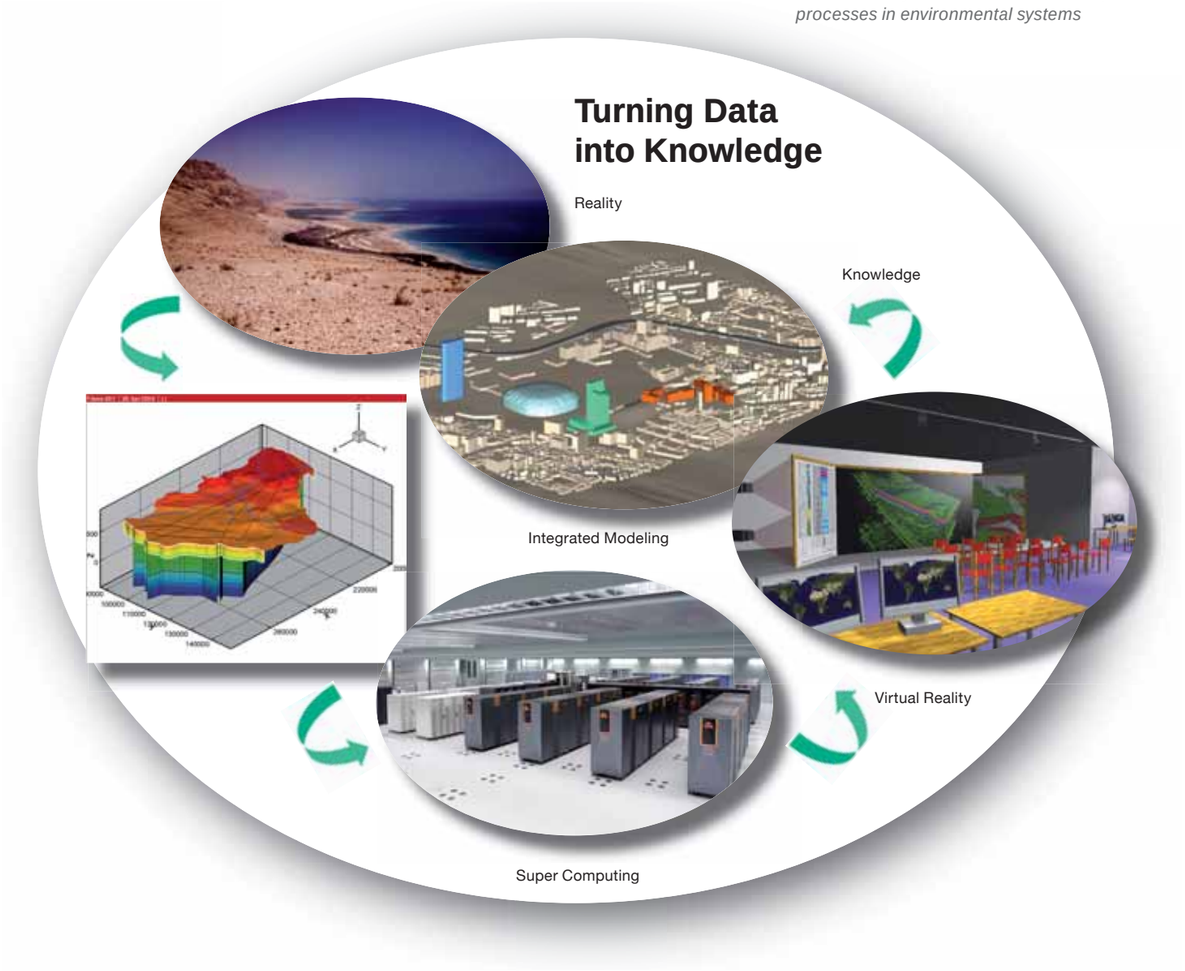

For instance, nuclear waste repositories constructed in the deep geologic underground are exposed to considerable weight loads accompanied by large litho-static stresses. Moreover, radioactive waste emplacements generate heat with temperatures over 100°C for long periods, thus, consideration of thermo-hydro-mechanical phenomena is very important in these geological environments. In addition, geo-chemical long-term effects are important considerations in the process analysis.

Normally, liquid phase flow is negligible, in intact very low-permeable host rocks, such as salt or clay stone. However, even under those conditions, vapour flow can develop due to artificially induced thermal gradients caused by heat emitting waste deposits. Gas phases can be produced due to the local heating of the formation water or as a result of chemical reactions during corrosion processes.

"To gain appropriate understanding of the complex interactions of different physical processes of chemo-toxic and nuclear waste repositories a fully coupled thermo-hydro-mechanical (THM) analysis is required. Numerical modelling tools are used more and more in order to gain a better understanding of the interaction of THM coupled processes in porous media as well as for the analysis of possible scenarios, long-term predictions, and safety assessments. In the past decade, many efforts have been done for solving THM problems by the use of numerical methods by research groups e.g. involved in the DECOVALEX (http://www.decovalex.com) and the BAMBUS projects, and by other researchers", says Kolditz.

From the mathematics point of view, THM processes lead to a coupled initial-boundary-value-problem (IBVP), which needs to be solved numerically. Among the available numerical methods, the finite element method (FEM) is mainly used so far for the solution of THM coupled problems. Irrespective of the specific numerical method the computation of coupled THM problems is very expensive. This is mainly due to two reasons: degree of freedom (i.e. number of field variables) and strong non-linearity. There are several ways to improve the computational efficiency

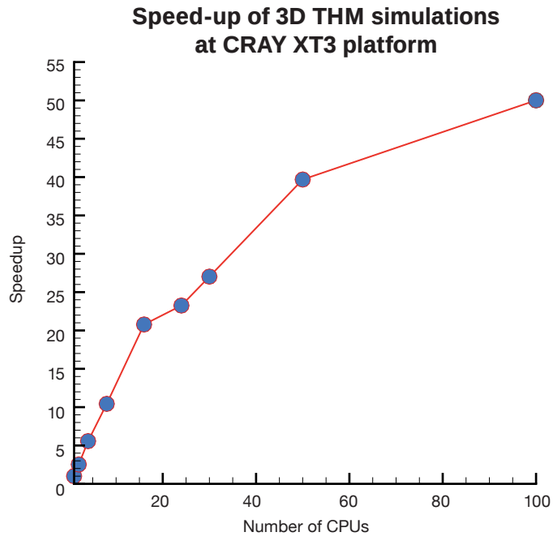

Speed-up of 3D THM simulations on CRAY XT3 platform at CSCS, Switzerland, achieved together with researchers at the Paul Scherrer Institut – PSI.

The 3D domain is discretised into 66 900 tetrahedra with 12 485 element nodes. The mesh is partitioned into 2 to 100 sub-domains for parallel speed-up measurements. With a maximum number of 100 CPU nodes the computation is about 50 times faster than the sequential simulation, i.e. less than optimal speed-up, but nevertheless this means that 50 times more numerical simulations are feasible now during the same wall time for 3D analysis of practical THM problems. The speed-up behaviour is clearly influenced by the coupling and non-linearity of the THM multi-field problem resulting in additional iteration loops. The number of non-linear iterations in sub-domains, which are close to residual or fully saturation, is much larger than in others. As these moisture fronts are moving, those phenomena are difficult to optimise, e.g. by dynamic load balancing.

such as efficient numerical algorithms, optimisation of memory management in the code, and parallel techniques. Parallel computing provides the most powerful speedup for numerical simulations. Thanks to the decreasing hardware cost in the past years, parallel computation is getting more attractive for research in computer science.

Only a few works are known for parallel finite element analysis of THM coupled problems in porous media. The parallel computations are normally conducted on small clusters with 8 CPU nodes maximum. Parallel codes make use of the sub-structuring technique based on following steps: 1) Divide the system into different sub-structuring levels, consisting of several element cells or super-elements, 2) condense the terms associated with internal nodal points in the sub-structure level assembly, 3) parallel assemble local condensed equation systems at sub-structures, 4) assemble a condensed global equation system with unknowns only associated with sub-structure interface nodes and then solve the problem in a so called multi-front manner, 5) recover internal unknowns using the condensed equations.

"In contrast, our novel parallel scheme is entirely based on the domain decomposition concept (DDC). The numerical calculation for each individual T/H/M process is partitioned by the a priori decomposition of the finite element mesh into sub-domains. All, sub-domain related operations are then assigned to the involved CPUs concurrently", says Kolditz.

The parallel object-oriented finite element scheme for each T/H/M process is realized in three major steps:
1) Concurrently establish sub-domain mesh topology for both linear and quadratic interpolations, 2) simultaneously assemble linear equation systems in sub-domains, 3) obtain the overall solution by using sub-domain linear equation solver. For step 3), a Jacobi preconditioned BiCGStab linear solver is implemented. The matrix-vector multiplications, which are the basic computational operations of the solver algorithm, are performed at sub-domain level by all involved CPUs concurrently.

Finally, the overall numerical solutions for the entire domain are updated by inter-processor communications. It is important to note, that this parallel scheme completely avoids the assembly of a global equation system. All equation assembly is restricted to sub-domain level.

The details of the parallel method including aspects of the object-oriented implementation in the framework of the open-source scientific software including numerical examples to demonstrate the speed-up of the parallel method are presented for 2D and 3D THM simulations are found on www.rockflow.net

Apart from clusters and the Cray XT3 at CSCS, in a joint project with HLRS, Stuttgart, GeoSys/RockFlow is being ported on a number of other DEISA platforms, the NEC SX-8, at HLRS, the MareNostrum at Barcelona, JUMP at Jülich Helmholtz supercomputing centre and so on. The test results to-date, are very encouraging and should enhance the value added to scientific research by the DEISA infrastructure. ■

Acknowledgements

We acknowledge the efforts of **Wenqing Wang** and **Chan-Hee Park** (UFZ), **Georg Kosakowski** (PSI) as well as **Andriy Chut** and **Thomas Bönisch** (HLRS) in the GeoSys/RockFlow parallelisation project.

Using supercomputers **to investigate twisted nematic liquid crystal displays**

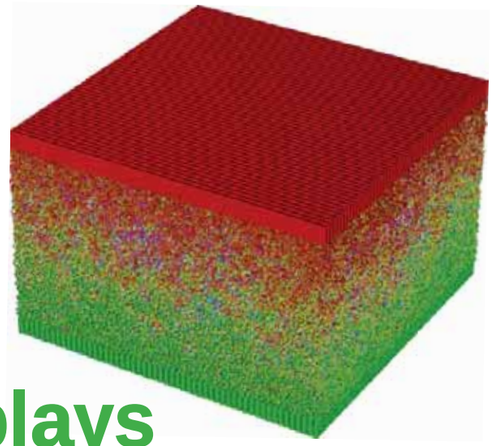

Damien Lecarpentier

The LCDIS project was undertaken in 2005 and 2006 by a research group from the University of Bologna, Italy, under the direction of Claudio Zannoni. Carrying out calculations both at CINECA, in Italy, and CSC in Finland, the team sought to further investigate the functioning of twisted nematic liquid crystal displays.

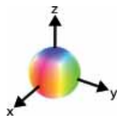

Snapshot of the final molecular configuration produced by the simulation in absence of applied voltage. Molecules are color coded according to their orientation as in the color palette. The color coding of the ellipsoidal particles – adopted to clarify how particles are organized inside the display – is the following: ellipsoids oriented along the x axis are red; along the y axis, green; and along the z axis, blue. Intermediate orientations are given intermediate colors, according the the palette.

"One of the most successful stories in the field of advanced materials is that of liquid crystal displays (LCD) and, in particular, of twisted nematic (TN) displays, which are today the most common of these devices", says **Claudio Zannoni**, coordinator of the LCDIS project, and researcher at the University of Bologna.

"However, although twisted nematic technology has been used by the LCD industry for over 30 years, certain basic features, such as the molecular mechanisms affecting the surface alignment of the pixel molecules, the role of surface pretilt, and the fluid structure at pixel boundaries, are still poorly understood."

"The basic concept underlying the twisted nematic display is that each pixel – which consists of a layer of molecules – is activated by a change of molecular organization in a cell no more than a few microns wide", explains **Matteo Ricci**, who also took part in the project. "According to this concept, the initial configuration of the local preferred direction (the director) is established between two orienting surfaces (rubbed glass or polyimide), rotated at 90 degrees to each another, within which the liquid crystal is confined. It can be shown experimentally that polarized light can pass through the pixel in this "off" state, as is compatible with a microscopic helical configuration.
If the chosen liquid crystal has a positive dielectric anisotropy and a suitable voltage is applied across the cell in correspondence with the pixel, then polarized light is not rotated and light does not pass through the device, as would be expected with a monodomain organization. When the field is switched off, the original molecular organization is re-established."

"The classic textbook depiction of the functioning of a TN-LCD is that of uniformly twisted layers, but to the best of our knowledge there is little evidence that the configuration of the pixel in the "off"-state is a uniform helix at the molecular level. Moreover, the manner in which this configuration is established is not obvious and many questions remain unsolved: As the "off" configuration is re-established after an aligning cycle, does this reconfiguration take place from the centre of the cell or from the surface? Does this occur uniformly or not? Is a uniform helix really formed? If not, how helical is the structure?", Zannoni asks.

"Acquiring a better understanding of the microscopic mechanisms of ordering inside LCDs is fundamental if we are to achieve substantial enhancements in the properties and energy consumption of such devices", point out the two researchers.

Claudio Zannoni, the LCDIS project coordinator

Investigating the microscopic working of the TN-LCD display through computer simulations: a challenging task

Computer simulation of twisted nematic displays is an essential tool for gaining insights into the microscopic details of such devices; it is also, however, a difficult task, due to the huge number of variables that have to be taken into consideration when performing the simulation, and to the very long time-scales involved. Full simulation of a realistic display remains far beyond our reach today.

"Attempting to investigate the microscopic functioning of a TN display is a particularly challenging task, as it necessitates taking a huge number of degrees of freedom into consideration, each of which must be analyzed for what are – from a microscopic perspective – extremely long periods of time. The average TN-LCD response varies within the range of 8–15 milliseconds, while the typical time-scales that we have been able to access to date through conventional computer simulation are of the order of nanoseconds", Zannoni notes.

"In this project, we tackled this problem by setting up a molecular resolution model of a TN cell containing ^{O}(106) model particles, and we simulated it using the Monte Carlo (MC) method. For this purpose, we developed an MPI parallel MC code using a replicated data scheme, by modifying the famous Markov Chain using the Metropolis algorithm to allow 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 the intrinsicly non-deterministic nature of MC. However, in our case, the large sample size allowed every processor to pick an energetically independent particle in suitably chosen cells, into which the whole sample was subdivided", he explains.

"We also modeled the LC rod-like molecules contained in the LCD cell as single ellipsoidal interaction sites based on the Gay-Berne model, discarding all of the intramolecular degrees of freedom", Ricci adds. "Our TN cell was arranged as to model a 0.1 μm × 0.1 μm × 0.05 μm display cell. Although a fully realistic model would have to consider a sample several orders of magnitude larger, this is currently unfeasible with the available resources", he notes. "The initial configuration was characterized by the LC molecules uniformly aligned perpendicular to the display surfaces to give the dark pixel state, while the display boundaries were modeled with layers of fixed particles oriented towards the directions of the incoming and outgoing polarizers (rotated at 90 degrees to each other)."

The calculations were carried out at the CSC in Finland, employing 128 processors, and at CINECA in Italy, using 64 processors, with a total global amount of 200 000 CPU hours. **>>>**

Despite the fact that the project team could not exploit the full power of the DEISA infrastructure, the huge amount of resources afforded by DEISA was, according to the reseachers, a determining factor in the sucess of the project.

Moving away from conventional accounts of the functioning of TN-LCDs

The simulations carried out within the DEISA research infrastructure allowed the research team to answer some of the problems that had been, until then, poorly understood, and provided new insights into the functioning of the TN-LCD devices that take us beyond the conventional wisdom on this topic.

"The system was allowed to relax through MC simulation to the "off" equilibrium state, which was reached only after 150 000 hours of computing time. Over the time-period of this equilibration process, we observed that the process of clearing the pixel proceeds with the local induction of order by the two confining surfaces, which induce an almost parallel alignment to their orientation in the LC molecules filling 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

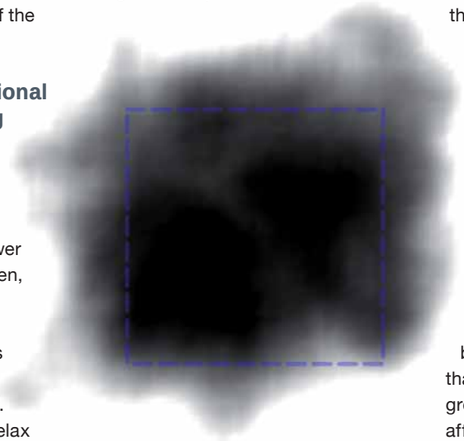

Pixel image computed when voltage is applied, after applying physical formulas relating the passage of light through layers of anisotropic molecules and their orientations.

linear helix", Zannoni notes. "On the contrary, the molecular layers align almost parallel to the closest display surface, and are not affected by the surface farthest away. Some sort of helicoidal ordering begins to form only when the two fronts of perpendicularly-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 render the conventional picture of a uniformly helicoidal configuration too "simplicistic", illustrating, moreover, that this requirement is not necessary to achieve good optical performance."

"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 the helicoidal order that leads to a black pixel. The interesting result is 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."

According to the research team, these findings will help to furnish a deeper understanding of the microscopic behaviour of TN-LCD devices and will have real practical consequences in terms of enhancing certain elements of their functioning and reducing their energy consumption. ■

Matteo Ricci (second from left), member of the LCDIS project, at DEISA training session in Barcelona, March 2007.

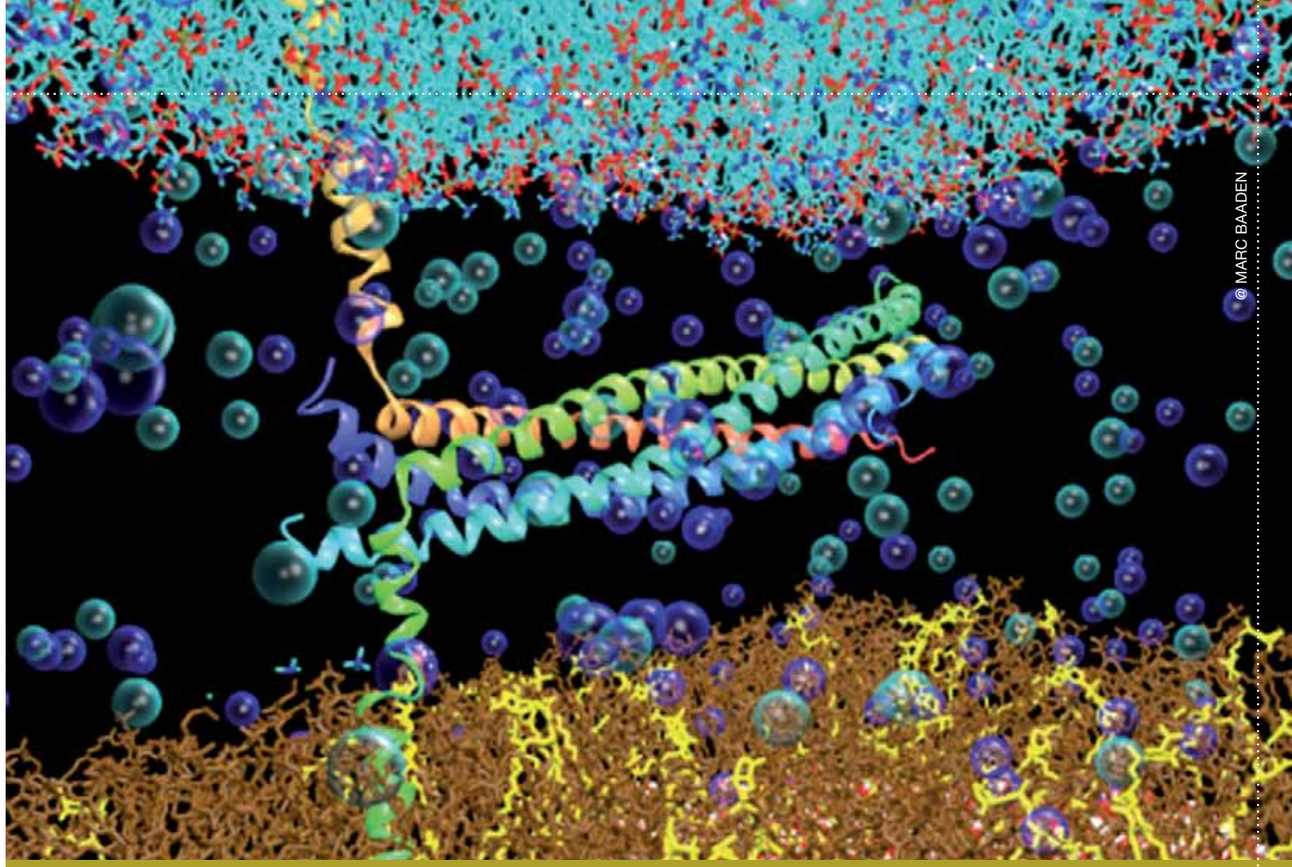

© MARC BAADEN

Membrane fusion and the SNARE protein complex

Pirkko Soininen

Utilizing DEISA computing resources, Marc Baaden, his research team in Paris, France, and colleagues from France, the United Kingdom and Austria, have modeled a protein complex that acts as a catalyst in the fusion of two membranes. This exceptionally complex model could not have been constructed without the distributed European computation capacity offered by DEISA. Understanding the functioning of proteins will open up new opportunities for pharmaceutical development.

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.

"Basic research is essential, since there are several aspects concerning the functioning of proteins and cell membranes that are not yet fully understood. A better nderstanding of these mechanisms will facilitate, for example, the development of new pharmaceutical agents", explains **Marc Baaden** who took part in the DEISA conference.

In medicine, the membrane-embedded proteins are of particular interest, as many of them are potential target sites for pharmaceuticals.

"By examining a phenomenon at the atomic level, we can gain insight into the behavior of cell membranes and proteins in general and on a larger scale", says the French researcher.

"Researchers often describe molecule-level events by sketching schematic models on paper. This kind of "cartoon biology" is, however, easily misleading, because such schematic models are rarely consistent with the actual physical properties of the proteins, such as their size or form", says Baaden.

"Computational molecular modeling, or simulation, takes the physical properties into account in a more realistic manner. A good molecular model enables us to examine the smallest details of the system in >>>

a controlled fashion and under the desired circumstances. We can also easily change any of the properties of the model to test different hypotheses", explains Baaden.

Experimentally determined molecular structures provide a static picture of the phenomena at the atomic level; they lack the dynamics and motion that are a central part of the functioning of, say, proteins. Simulation, on the other hand, facilitates the examination of dynamic events.

Challenging accuracy in modeling

Baaden and his team have modeled the SNARE protein complex that contributes to membrane fusion in cells.

"Traditionally, computational modeling has been applied to simulate a single membrane – this is already a very challenging task. We, however, simulated two lipid membranes held together by a protein complex, so our challenge was even greater", Baaden says.

In his experiments, Baaden has constructed a functioning simulation environment. While examining the SNARE complex, he has also tested various future methods for modeling cell membranes. The new model has also facilitated the testing of the accuracy of the existing models and theories.

"We have been able to glimpse into the very finest details of a biological system to actually see what happens at the molecular level. This would not have been possible with any other method."

© ERKKI MAKKONEN

Marc Baaden has simulated the interaction of the SNARE protein complex with lipid membranes. For the simulations, the GROMACS simulation software was applied to compute classical molecular dynamics based on force fields.

"I believe that, in the future, those methods will succeed that combine highly accurate atomic level modeling with coarse-grained models simulating a more extensive event. Communication between the simulations on various levels is important. This will allow us to simulate a detail of particular significance more accurately, and larger entities by means of a coarse-grained model", Baaden concludes.

Complex modeling demands high computation capacity

Marc Baaden has striven for high accuracy in the modeling. The team used standard molecular modeling methods in their work, but the system subject to simulation was exceptionally extensive and complex. The total number of atoms to simulate in the model was 340 000. And yet, on the cellular scale, the area to be simulated was extremely small.

"Our greatest challenge concerned the construction of the computational model. We ran several trial tests on different models without success until we managed to make the system work. There was a stage when I was about to give up, since it seemed that the model simply would not work. Luckily, an Austrian colleague, who loves challenging cases, came to my aid. Together we were able to construct a model that could be run using GROMACS", Baaden remembers.

"DEISA has helped us greatly. In Paris, we did not have sufficient computation capacity, but the DEISA infrastructure was an optimal suit for our purposes. The only drawback was that the computing hours reserved for us coincided with the start of the academic year, so it was quite hard to recruit students", Baaden says.

First steps toward future applications

The functional disorders of cell membranes are known to be associated with many diseases.

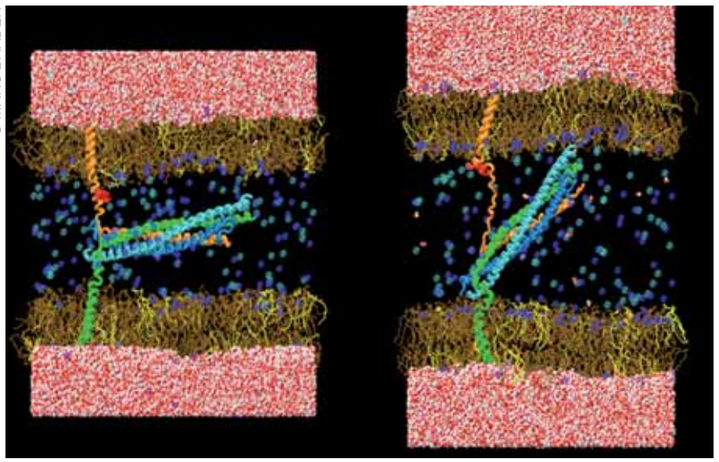

© MARC BAADEN

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.

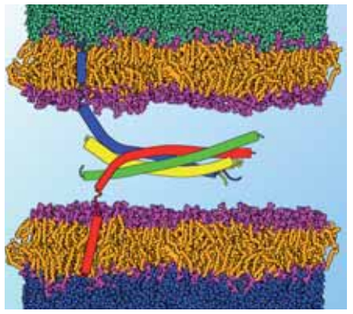

© MARC BAADEN

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.

The fusion of membrane structures, which is catalyzed by the protein complex studied by Baaden, may be disturbed in two ways: either the membranes do not fuse at all, or they fuse too heavily.

"At present, we do not sufficiently understand the details of the cell membrane mechanisms. By modeling these events, we may, in the future, be able to develop more efficacious drugs that are conveyed directly to the target site", says Baaden.

"It is important for a researcher carrying out basic research to be able to visualize the goal. We are ontributing to a long process, even if we are only taking the first steps in the track. We may imagine potential future applications, but it is equally possible that once we learn to better understand the cell functions, the applications will be something completely different", he says.

Apart from medical science, the cosmetics industry and nanotechnology will benefit from a better understanding of protein functioning.

"Knowledge of the functioning of cell membranes will open up opportunities in the field of nanotechnology as well. In technical terms, the proteins we study are clever machines that perform their intended tasks excellently, that is, to fuse two lipid membranes firmly together", Baaden points out. ■

SNARE proteins in cell membrane fusion

Tommi Kutilainen

SNARE proteins are a large protein superfamily. Their primary role is to control membrane fusion of cellular transport vesicles with cell membranes or with target compartment membranes. Such fusion activities are essential for the normal function of the cell.

A vesicle is a compartment separated from the cytoplasm by at least one lipid bilayer. Vesicles transport, store, and digest products or waste produced within the cell. The vesicle membrane enclosing the vesicle is similar to that of the cell membrane. SNARE proteins inside the cell assist in the fusion of the vesicle membrane and the cell membrane. The membrane fusion plays an active role in the vesicle-mediated transport of materials inside the cell and through the cell membrane.

For example, when a particle is removed from a cell in exolysis, the particle is encapsulated by a vesicle that migrates to the cell surface and anchors to the cell membrane with its inside facing the outside of the cell, thus secreting the particle out of the cell. In a similar way, vesicles transport proteins and lipids to the Golgi apparatus, which is responsible for processing and packaging of proteins and lipids.

Mutual identification of SNARE proteins precedes the fusion of the vesicle and target cell membranes. A receptor (v-SNARE, synaptobrevin) protein anchored on the vesicle membrane and a corresponding receptor (t-SNARE, syntaxin) on the target membrane are wound with two SNAP-25 proteins creating a SNARE protein complex in the form of a four-α-helix bundle. This four-protein complex is a necessary, but probably not sufficient, condition for cell membrane fusion to take place.

The scope of SNARE research is not limited to adding to the understanding of normal cellular activities. For example, disturbed functioning of SNARE proteins may initiate the adult-onset diabetes (type 2 diabetes). Hence, understanding the SNARE function may facilitate development of new therapeutic treatments. It is essential to develop new therapies, as diabetes is one of the most significant national diseases in Finland. According to the current WHO figures, approximately three percent of people in Finland develop diabetes, but the number is constantly growing. ■

For more information http://www.baaden.ibpc.fr/projects/snaredeci/

Distributed
European
Infrastructure for
Supercomputing
Applications

DEISA

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 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

e-infrastructure

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 and from the Seventh Framework Programme (FP7) under the grant agreement n° RI-222919.