inSiDE

inSiDE • Vol. 9 No. 1 • Spring 2011

Innovatives Supercomputing in Deutschland

Editorial

- Prof. Dr.
  H.-G. Hegering (LRZ)
- Prof. Dr. Dr.
  Th. Lippert (NIC)
- Prof. Dr.-Ing. Dr. h.c. Dr. h.c.
  M. M. Resch (HLRS)

Welcome to this new issue of inSiDE the journal on innovative Supercomputing in Germany published by the Gauss Centre for Supercomputing. Good news for High Performance Computing in Europe is ahead.

Two new systems from Germany are going to support PRACE in the coming years. The High Performance Computing Center at Stuttgart will install a 1 PFlop Cray XE6 called Hermit in summer. This will be followed by a 3 PFlop IBM cluster called SuperMUC installed by the Leibniz Rechenzentrum at Garching in 2012. Both systems are part of the German PetaGCS project and are funded by the Federal Ministry of Education and Research (BMBF) and the State Govern-ments of Baden-Württemberg and Bavaria.

Both systems will become part of the European HPC landscape established by PRACE. They join the 1 PFlop IBM Blue Gene/P system called JUGENE, co-funded by the State Government of Nordrhein-Westfalen and being in operation since July 2009. We give a report on the status of PRACE in this issue. Furthermore we can report about the first regular PRACE HPC access grants for research projects. It is obvious that the concept of a European research landscape for simulation takes shape. Bundling its resources Europe again sets out to take a leading role in one of the key technologies of the 21st century.

The wide range of scientific fields that benefit from the newly available hardware is emphasized in this issue of inSiDE. 58 pages of application reports show the impact of new hardware and soft-ware on basic and applied research.

We start with three articles that were awarded the Golden Spike by the steering committee of the HLRS: Climate research from the Leibniz Institute for Ocean Research at Kiel/ Germany, flow over an airplane wing and helicopter simulation from the University of Stuttgart. The two following articles describe the progress in turbulence simulation; both report on research performed by Gauss Large-Scale projects carried out on JUGENE in the previous granting period. Physics and chemistry research are represented by three outstanding and extensive articles. Finally two fascinating articles show what can be done in astro-physics using high end supercomputers.

The success of future applications origi-nates from projects. This issue reports on four of them. LarKC is funded by the European Commission and focuses on high performance semantic web experience. Another European project plugIT works on the integration of busi-ness in High Performance Computing. Monitoring is a big issue for large scale systems. LRZ reports on a project to develop a software package called PerSyst. AstroGrid reports on a break-through in Grid Computing in a project funded by BMBF. Two short communi-cations report on the starting of a new ITEA project called H4H on hybrid programming and on HOPSA, a new joint EU-Russian project for holistic performance analysis.

Among our many activities we highlight the 100 Gigabit Ethernet connection between Research Centre Jülich and Karlsruhe Institute of Technology. All in all inSiDE will again give you a comprehensive overview of HPC not only in Germany but also in Europe.

Enjoy reading!

Contents

News

» New HLRS System HERMIT 4
» Preparing for the new 3 Petaflop/s System at LRZ 6
» PRACE – Full Speed ahead 9
» First regular PRACE HPC Access Grants for Research Projects 12
» Jülich established as CECAM Node 13

Applications

» The Agulhas System as a Key Region of the global oceanic Circulation 14
» Laminar Flow Control in Vortex-deformed Swept-Wing Flows: Pinpoint Suction 18
» Fluid-Structure coupled Flow Simulations of Helicopter Rotors 22
» Joint numerical and theoretical Investigations of fully Developed Fluid Turbulence 26
» A new View on Geometry and conditional Statistics in Turbulence 30
» Applications of Multiferroics: from elementary Particles to new Memory Devices 38
» Nanoparticles as Emulsion Stabilizers 48
» Functional Molecules at Surfaces: The Role of dispersive Interactions 56
» Galaxy Cluster Simulations: The FEARLESS Approach 62
» Simulating Giant Spiral Galaxies in the early Universe 66

Projects

» Towards High Performance Semantic Web – Experience of the LarKC Project 72
» plugIT – Plug Your Business into IT 76
» A scalable Monitoring Tool using Performance Properties 80
» The AstroGrid-D Use Case GEO600: A Breakthrough in Grid Computing 84
» H4H Project launched 87
» HOPSA Project launched 87

Centres 88

Activities 94

Courses 110

New HLRS System HERMIT

Stepping into the Petaflops Era for Engineering and Industry

Petaflops have been around for two years now and the world has seen a number of big installations. The time has come to install a productive system that can sustain its performance in all major areas of High Performance Computing.

The Gauss Center for Supercomputing (GCS) in Germany has launched an ambitious project called PetaGCS in 2008 to roll out petascale systems at its three locations (Garching, Jülich, Stuttgart) over a period of six years. In 2009 Jülich was the first to install such a system and has since then led the GCS into the petascale era. The High Performance Computing Center Stuttgart (HLRS) was scheduled to follow in 2011.

When HLRS started its procurement in 2009 the key issues were defined by its long tradition of a focus on sustained performance, usability and an application set that ranges from classical engineering deep into the advanced new simulation method space. The actual specifications for the system were put together by four main groups. First and foremost it was the user community that had to bring in its input. A working group of 12 outstanding scientists from all over Germany held a two day meeting to collect requirements but also discuss visions for HPC simulation. Another group of scientists from the German national Cluster of Excellence in Simulation Technology (SimTech) held meetings with the HLRS in order to make sure that most advanced methods for simulation were considered in the evaluation process. At the same time HLRS used its regular technical meetings with industry leaders like Porsche to understand the future needs of industry in HPC and consider it for the next installation step. Finally the 40 year experience of HLRS itself in operating and supporting supercomputers was brought together in an internal evaluation group.

As a result of these wide ranging discussions, HLRS again put its focus on balanced systems with a clear predominance on sustained performance for real world applications. The procurement process was started in August 2009 and was finished in April 2010. Finally a contract was signed with Cray on October 26, 2010. The European President of Cray Dr. Ulla Thiel and the Chancellor of the University of Stuttgart Dr. Bettina Buhlmann put their signatures on a contract that would guarantee the delivery of a two step installation of Cray's leading edge systems in the years 2011 and 2013.

Figure 1: Signing-up: Nurcan Rasig, Dr. Bettina Buhlmann, Dr. Ulla Thiel, Prof. Dr. Michael Resch (left to right)

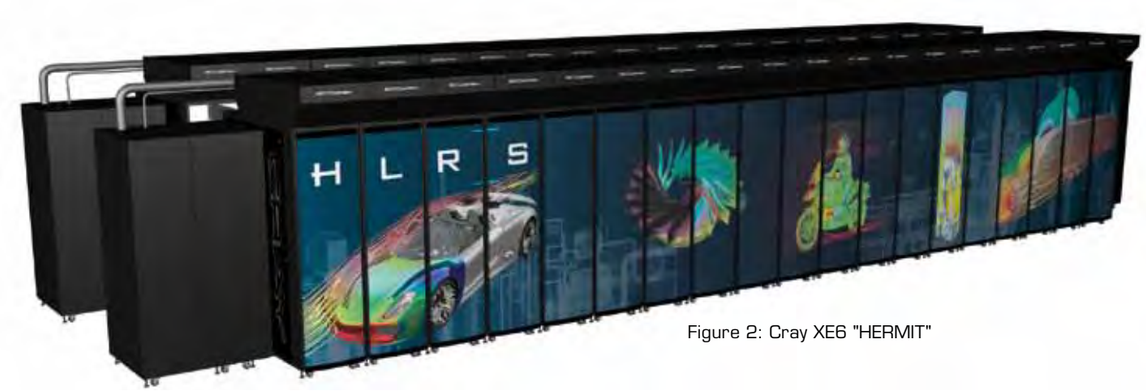

Figure 2: Cray XE6 "HERMIT"

Already in November 2010 Cray installed a transition system at HLRS with 1,344 cores. The system was brought in as early as possible to first allow staff of HLRS to explore and understand the new technology. After a very short period the system was also made available to the international and national users of HLRS in order for them to learn how to port their codes to the new architecture.

Starting in July 2011 Cray will install the first step of the new HLRS HERMIT XE6 system. The roughly 3,500 two socket nodes come with the AMD Interlagos processor. The total number of cores will be at least 100,000. Main memory is 32 GB or 64 GB per node reflecting the variety of needs for applications. This will result in an overall main memory of more than 100 TB. A total of 2.7 PB of disk space will be available for users with an I/O speed in the range of 150 GB/s. The network is the Cray developed Gemini which stands for low latency and high bandwidth. In order to fully support ISV codes on a high end system, Cray's CCM will be used as a basis. This will allow to run ISV codes without any changes in full production mode.

The second step will be installed in 2013. With the contractually agreed level of sustained performance will result in a peak performance of 4 to 5 PFlop/s. Currently, we expect to see about 150,000 cores which might be supplemented by GPGPU nodes depending in the then available technology and the relevance for sustained performance for real applications. The network is going to be the Cray developed Aries which will further reduce latency and increase bandwidth to entirely new levels of performance.

• Michael Resch

University of Stuttgart, HLRS

Spring 2011 • Vol. 9 No. 1 • inSiDE 5

Preparing for the new 3 Petaflop/s System at LRZ

New Building

Less than five years ago Leibniz Supercomputing Centre (LRZ) moved from the city of Munich to its new generous buildings in Garching near Munich. But very soon it became obvious that more space would be needed for all the new services and computers offered by LRZ to the scientific communities in Bavaria, Germany and Europe.

Only two years after planning and construction started, LRZ celebrated the topping out ceremony of the extension of its buildings on October 18th, 2010. The computer building will nearly be doubled making it from a "Computer Cube" to a "Twin Cube". At three floor levels a total of 3,000 square meters is used for tape archives, servers and Linux-clusters and for the new high-end system, respectively. Electricity connections, transformer, UPS devices and air and water conditioning are contained in two additional floor levels. Along with the computer building new office space for staff is established, and finally a huge visualization centre with a powerwall and a cave will enable new dimensions in visualization and interactive steering of large scale simulations on supercomputers.

As part of a massive initiative to enhance HPC resources in Germany, LRZ received 83 million Euros of funding to build and operate its new, "Tier-0" Petascale supercomputer. The funding comes jointly from the State of Bavaria and the Federal Government of Germany - as do the additional 50 million Euros for the extension of LRZ's buildings. In addition, Bavaria will support accompanying projects in the area of High Performance Computing e.g., KONWIHR, the Bavarian Competence Network for Technical and Scientific High Performance Computing.

Based on consultations with its users, LRZ decided to create a cluster that would not only deliver outstanding processing performance but also accommodate a wide range of applications without requiring researchers to port code to a proprietary architecture. At the same time, the new high-end system should deliver exceptional energy efficiency - reducing operational expenses allows to use a greater portion of the funding for hardware purchases.

After conducting a competitive public procurement, the LRZ team finally selected an IBM System x iDataPlex servers with upcoming Intel processors as the hardware foundation for the new supercomputer. Named "SuperMUC" (borrowing the MUC abbreviation for the Munich airport which is just a few kilometers away), the new supercomputer will be comprised of more than 14,000 next-generation Intel Xeon processors with more than 110,000 processor cores and 320 TBytes of memory. When its installation is completed in 2012, SuperMUC will deliver 3 Petaflop/s of peak performance, making it one of the most powerful general purpose supercomputers in the world. A small fraction of it will be delivered as a system for user migration already in June 2011.

The disk storage of SuperMUC consists of 10 PBytes of permanent storage for scratch and work files using the IBM General Parallel File System (GPFS) as its parallel file system infrastructure. Further 2 PBytes of storage intended for user's home directories will be kept

Figure 1: The new "Twin Cube" (yellow) at LRZ which will contain the new Petascale system and other computing systems (Photo: E. Graf).

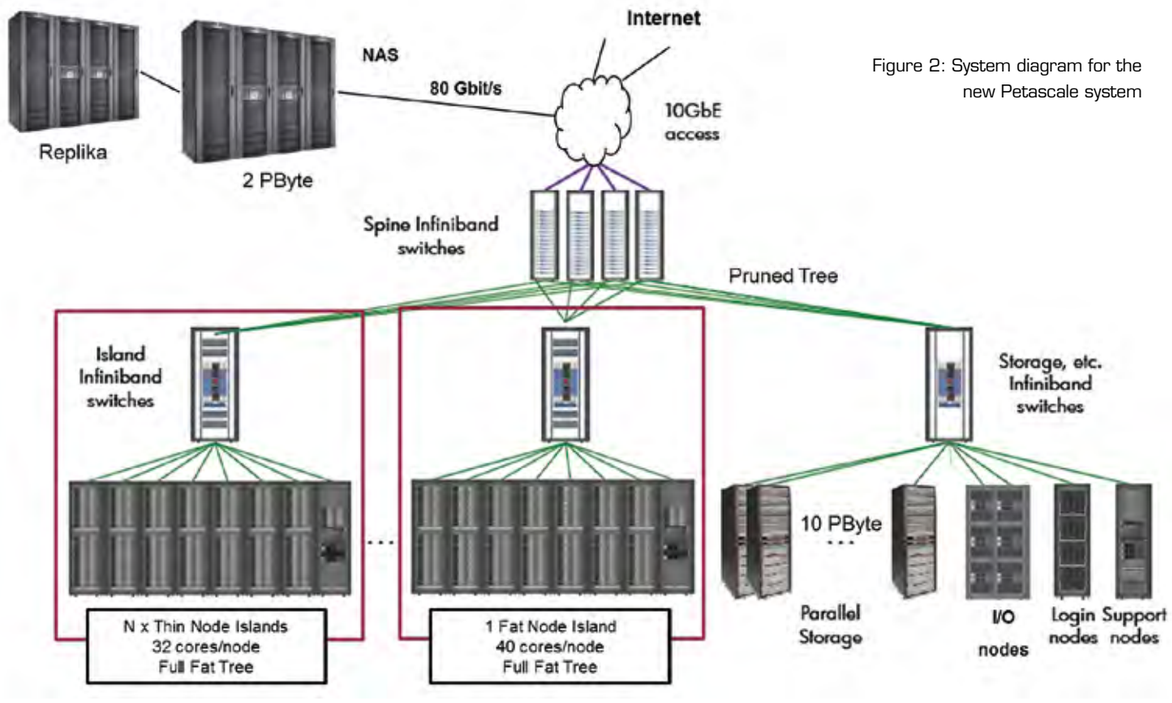

Figure 2: System diagram for the new Petascale system

on NAS (networked attached storage). For security and data integrity all NAS files are replicated to a second system located in a separate room.

The system is built up from a number of Thin Islands and one Fat Node Islands interconnected by an Infiniband-based high performance network; the topology of this network is a full fat tree within each island, and is pruned between islands.

Revolutionary Cooling Concept

The new cooling concept is revolutionary. Active components like processors and memory are directly cooled with water that can have an inlet temperature of up to 45 °C. SuperMUC will capitalize on this innovative hot water-cooling technology that is 4,000 times more efficient than air cooling. Developed by IBM and first implemented in the Aquasar Cluster created for the Swiss Federal Institute of Technology Zurich (ETH Zurich), the chip-level technology will help to dramatically reduce the power required to extract waste heat from the servers, at the same time enabling LRZ to reuse it to heat rooms in the office tracts. The big advantage of that technology is, that by using towers for free cooling no additional chillers are needed. LRZ expects to achieve a power usage effectiveness (thermal PUE) ratio of 1.1 by using the hot water-cooling technology.

It is also planned to employ software tools to help further increase the energy efficiency. Energy-aware scheduling tools will monitor application energy consumption and then create administrative policies that optimize processor frequency for particular applications and power down nodes that are not being used.

A Tier-0 Supercomputer for PRACE

On December 13th, 2010, LRZ and IBM signed the contract for the new system. SuperMUC will continue LRZ's successful strategy to install universally usable supercomputers based on general purpose processors.

German users can apply for large scale projects via the Gauss Centre for Supercomputing (GCS). SuperMUC will also be the third world-class computing facility within PRACE (Partnership for Advanced Computing in Europe), offering unique possibilities for the scientists from 21 European PRACE member states.

Bavaria's Minister of Science Dr. Wolfgang Heubisch named SuperMUC an investment into the future: "Powerful computers and software are today's keys for scientific and technological competitiveness. With this new supercomputer, the Leibniz Supercomputing Centre in Garching will be a pioneer in energetically optimized computer technology."

• Matthias Brehm
• Ludger Palm

Leibniz Supercomputing Centre (LRZ)

Figure 3: Signing the contract for SuperMUC: Bavarian State Minister for Science Dr. W. Heubisch, Prof. Dr. A. Bode, Director LRZ, M. Jetter, CEO IBM Germany, A. Pflieger, Manager IBM Systems Group, Prof. Dr. D. Willoweit, President Bavarian Academy of Sciences and Humanities (left to right)

PRACE – Full Speed ahead

Just one year after its creation the Partnership for Advanced Computing in Europe (PRACE) supports 36 scientific projects and provides more than a billion core hours of compute time on its first production systems. The following achievements mark important milestones for PRACE:

- **The successful completion of the PRACE** project that prepared the legal, financial, administrative, and technical groundwork for Europe's High Performance Computing (HPC) service.

- **The creation** of the PRACE AISBL (Association Internationale Sans But Lucratif) as an international non-profit organization under Belgian law.

- **The start** of a series of regular calls for proposals evaluated through a single European peer review process under the direction of the PRACE Scientific Steering Committee (SSC).

- **The launch** of the First Implementation Phase Project (PRACE-1IP) on July 1, 2010, to accelerate the implementation of the PRACE Research Infrastructure and especially to support users through application scaling and education and training.

- **The completion** of the first 10 projects on the first PRACE Tier-0 on JUGENE, the IBM Blue Gene/P at the Gauss Centre for Supercomputing (GCS) partner site in Jülich, Germany.

- **The deployment** of the second PRACE Tier-0 system CURIE, the 1.6 Petaflop/s Bull system by GENCI installed at CEA in Bruyères-Le-Chatatel, France.

In the first three calls 173 projects involving researchers from many European countries requested nearly six billion core hours resulting in a 6-fold oversubscription of available resources. PRACE issues calls for Tier-0 services every six months. The upcoming systems will increase the available resources since the first call in 2010 more than sevenfold. In the third quarter of 2011 the first installation step called HERMIT at the GCS partner site High Performance Computing Centre of the University of Stuttgart with 1 Petaflop/s will be made available followed by a second step with 4-5 Petaflop/s in 2013. The 3 Petaflop/s SuperMUC at the GCS partner site Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften in Garching, Germany will be available in the middle of 2012, to be followed by installations in Italy and Spain.

Figure 1: The first production system, a 1 Petaflop/s IBM Blue Gene/P (Jugene) at FZJ (Forschungszentrum Jülich) is available for European scientists

In addition to the regular bi-annual project calls PRACE has a continuous call for Preparatory Access. This provides a unique opportunity for scientists to assess, if their application can use the available Tier-0 resources effectively, prior to applying for production runs. There are three flavours in this type of call: Users may test the scalability of their application to determine realistic estimates for required resources prior to submitting a proposal. Alternatively, they may request time on a Tier-0 system to port, scale, or optimize their application to the target architecture. Finally, expert support from PRACE can be requested to assist in moving applications from Tier-1 or Tier-2 systems to Tier-0. Proposals for preparatory access are evaluated every two months and grants range typically in the range of 50-250 thousand core hours, depending on the type of call and the target system.

Figure 2: The second production system, a 1.6 Petaflop/s BULL (CURIE) at CEA-GENCI has been available for European scientists since January 2011.

PRACE continues to educate researchers and students on the use of new systems and novel programming techniques in a series of schools and workshops which include especially hands-on training on the Tier-0 systems or on prototype architectures. All material is available on the PRACE web site as videos or PowerPoint presentations. This service will be enhanced to become a full-fledged training portal in the near future.

PRACE is recognized internationally, not the least through its presence at the key international conferences, especially SC'xx in the US and ISC'xx in Europe.

The next important step will be the Second Implementation Phase project (PRACE-2IP). The consortium has submitted a successful proposal and has been invited to the contract negotiations. PRACE-2IP is expected to start mid 2011 complementing and extending the ongoing work. Selected highlights and expected results of PRACE-2IP are:

Integration of Tier-1 resources and services as currently provided by DEISA until April 2011. This includes a framework for resource interchange between national or regional HPC centres, mechanisms for bartering, and high-speed connection to new partners through GÉANT.

Support for scientific applications to be executed on Tier-1 systems. This will enable a wide range of European researchers from many disciplines to use advanced computing services, even if they do not (yet) require the highest performance levels.

Figure 3:
The PRACE booth at the SC'10 Exhibition in New Orleans, USA.

Community code development that enables developers in the scientific communities to parallelize and scale their codes for execution on future architectures that are likely to be the next generation of production systems. The communities will be selected at the start of the project with the help of the SSC, and experts from PRACE will assist the developers in refactoring the codes.

PRACE will work with researchers in industry to support them to ready their applications to execute efficiently on HPC systems thus increasing their competitive advantage.

An important part of the work is the research in advanced software and hardware technology to position PRACE for the Exascale regime. The hardware research focuses on novel architectures using selected applications as co-design vehicles.

PRACE will investigate novel programming techniques and implement auto-tuned runtime environments for systems with homogenous and accelerator based nodes. It will explore scalable algorithms with higher asynchrony between communication and computation using the hardware prototypes evaluated in parallel.

The European Commission continues to support PRACE through its Seventh Framework Programme (FP7). The expected grant for the Second Implementation Phase project (PRACE-2IP) is 20 million Euro matched by an equal amount from the PRACE membership consortium.

For up-to-date information about the PRACE AISBL and the PRACE-1IP project, the current list of members, and latest news visit: www.prace-ri.eu

References

[1] The PRACE-1IP project receives funding from the European Community's Seventh Framework Programme under grant agreement no. RI-261557. Additional information can be found under: www.prace-project.eu

• Dietmar Erwin

Jülich Supercomputing Centre

First regular PRACE HPC Access Grants for Research Projects

Since August 1, 2010 the Partnership for Advanced Computing in Europe (PRACE) is offering supercomputing resources on the highest level (Tier-0) to European researchers. Jülich Supercomputing Centre (JSC), as one of the three members of the Gauss Centre for Supercomputing (GCS), is dedicating a 35% share of its Blue Gene/P system JUGENE. This is the first Petascale High Performance Computing system available to researchers through PRACE.

Shortly after finishing the "PRACE Early Access Call" reported about in the last issue of inSiDE, proposals for the first regular research projects on JUGENE were solicited in the first Project Access Call, released by PRACE on June 15, with the deadline August 15, 2010. Project Access is intended for individual researchers and research groups, including multinational research groups, and has a one year duration. 59 applications with a volume of about 2,300 Million compute core hours were submitted. After strict technical and scientific evaluation and prioritization, nine research projects have been awarded about 360 Million compute core hours. Two of those research projects are from Germany, two from Spain, and one each from France, Hungary, the Netherlands, Portugal and the UK. Four research projects are from the field of Engineering and Energy, three from Fundamental Physics and one each from Chemistry and Materials and from Medicine and Life Sciences.

More details can be found via the PRACE web pages **http://www.prace-ri.eu/PRACE-1st-Regular-Call**

Calls for Proposals for Project Access are issued twice yearly (May and November). The next PRACE call will open on May 2, 2011 with a deadline of June 30, 2011.

• Walter Nadler

Jülich Supercomputing Centre

Jülich established as CECAM Node

Forschungszentrum Jülich has joined the Centre Européan de Calcul Atomique et Moléculaire (CECAM) as an active Node. Traditionally CECAM is known as a European organization, which promotes activities in the fields of atomistic simulations, e.g. quantum ab initio calculations or molecular dynamics, and development of methods and algorithms for bridging time and length scales in simulations via organizing workshops, tutorials or visitor programs. When CECAM headquarters moved from Lyon to Lausanne in 2008 the structure of CECAM was transformed into an international multi-site structure, where nodes were established in different European countries, supporting activities and running official CECAM programs. Up to now, nodes were established in 7 European countries, i.e. Ireland (Atlantic Centre for Atomistic Modeling), France (Ile de France, Rhône-Alpes), United Kingdom (Hartree Centre Daresbury), Spain (ZCAM Zaragoza), Germany (multi-site node in Berlin-Bremen-Frankfurt-Halle as well as in Berlin - Freie Universität), Netherlands (Amsterdam-Leiden) and Italy (Trieste-Bologna, Pisa), which all have their different thematic focus. Besides promoting and establishing leading research activities in the fields of molecular simulations, multi-scale modeling, and algorithms, the nodes form a network structure, working together towards a European software and knowledge base.

Forschungszentrum Jülich has been a supporting member of CECAM since years and is represented in its Scientific Council. In July 2010, Forschungszentrum Jülich signed an agreement to establish an active CECAM node in Jülich with a thematic focus on Soft Matter Physics, Materials Science, and Computational Science. The Institute for Advanced Simulation is in charge of organizing and running CECAM specific activities like workshops, tutorials, schools, and visitor programs, which all are approved by the CECAM Council and Scientific Advisory Board. Dr. Godehard Sutmann from JSC was appointed director of the Jülich Node and will coordinate local and international activities. This year a workshop on Soft Matter, tutorials on DFT methods and fast algorithms in molecular simulations as well as an international guest student program on scientific computing are organized in Jülich as CECAM activities. In the future, also a limited number of workshops in the field of scientific computing from external groups is envisioned. For further information, please visit the CECAM web pages:

www.cecam.org,
http://www2.fz-juelich.de/jsc/cecam

or send e-mail to:
g.sutmann@fz-juelich.de

• Godehard Sutmann

Jülich Supercomputing Centre

Golden Spike Award by the HLRS Steering Committee in 2010

The Agulhas System as a Key Region of the global oceanic Circulation

The ocean currents around South Africa are an important element in the global ocean circulation. Under present climate conditions the flow of warm and salty waters from the Indian Ocean into the Atlantic Ocean around the southern tip of Africa, the "Agulhas leakage", provides the bulk of the upper limb of the thermohaline circulation in the Atlantic Ocean. Parts of this water later feed into the Gulf Stream system of the North Atlantic that is responsible for the mild climatic conditions in Europe. The understanding of the dynamical factors determining the intensity and variability of Agulhas leakage is still incomplete, and so is its behaviour under a changing climate.

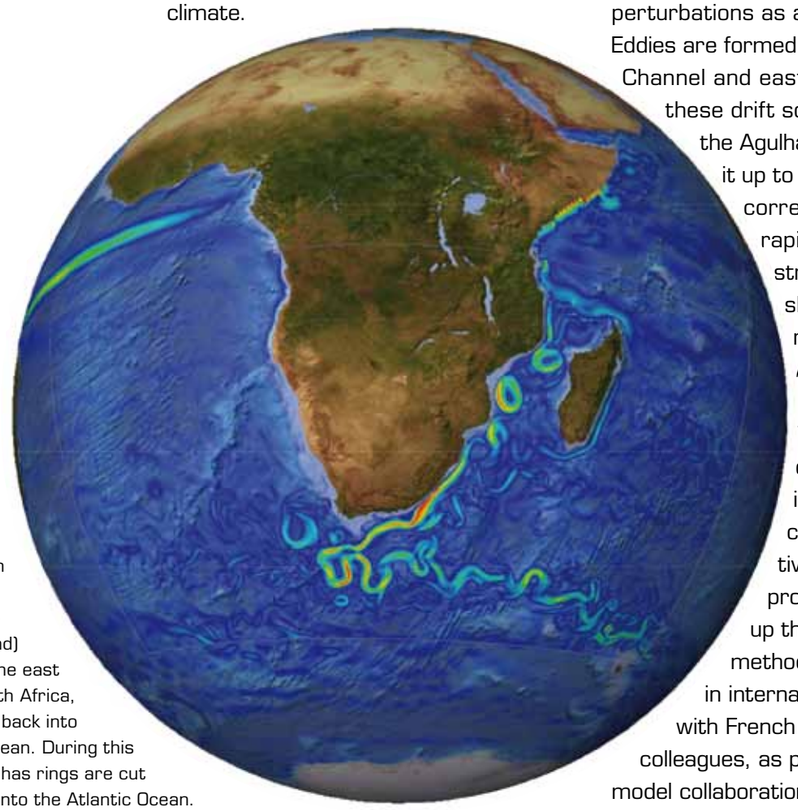

Figure 1: Surface circulation (snapshot) around South Africa. The Agulhas Current (red band) flows along the east coast of South Africa, retroflecting back into the Indian Ocean. During this process Agulhas rings are cut off and drift into the Atlantic Ocean.

In contrast to its large-scale importance the circulation in the Agulhas region is a dynamical mixture of different time and space scales (Fig. 1): A strong western boundary current, the Agulhas Current, transports the warm and salty water southward in the Indian Ocean. South of Africa it overshoots the continental slope and abruptly turns back into the Indian Ocean, while shedding enormous mesoscale rings of several 100 kilometres in diameter and extending over large parts of the water column. These rings transport the heat and salt as pulsating elements into the Atlantic Ocean. The circulation dynamics in the region also includes small-scale upstream perturbations as an important element. Eddies are formed in the Mozambique Channel and east of Madagascar; these drift southwards towards the Agulhas Current, displacing it up to 200 km offshore. The corresponding meanders rapidly progress downstream and trigger the shedding of Agulhas rings and therefore Agulhas leakage.

To examine the role of Agulhas leakage in the global oceanic circulation, an innovative ocean modelling program has been set up that advances new methodologies developed in international cooperation with French and South African colleagues, as part of the European model collaboration DRAKKAR [1].

The model hierarchy is based on the "Nucleus for European Modelling of the Ocean" (NEMO, v.2.3) [2], consisting of a coupled ocean/sea-ice model. The ocean component is a finite-difference discretization of a variant of the Navier-Stokes Equations (the "primitive equations"), stepping three-dimensional velocities, temperatures and salinities forward in time. A free surface formulation (e.g. by a conjugate gradient solver), a high-order polynomial fit of the density equation and lots of parameterizations for different ocean physics and small-scale processes let the complete program appear with a wide range of different numerical methods, though flexible in its use due to the modular formulation. It is written in FORTRAN 90 and has geographical domain decomposition in the horizontal for MPI parallelization. Traditionally the grid space layout and

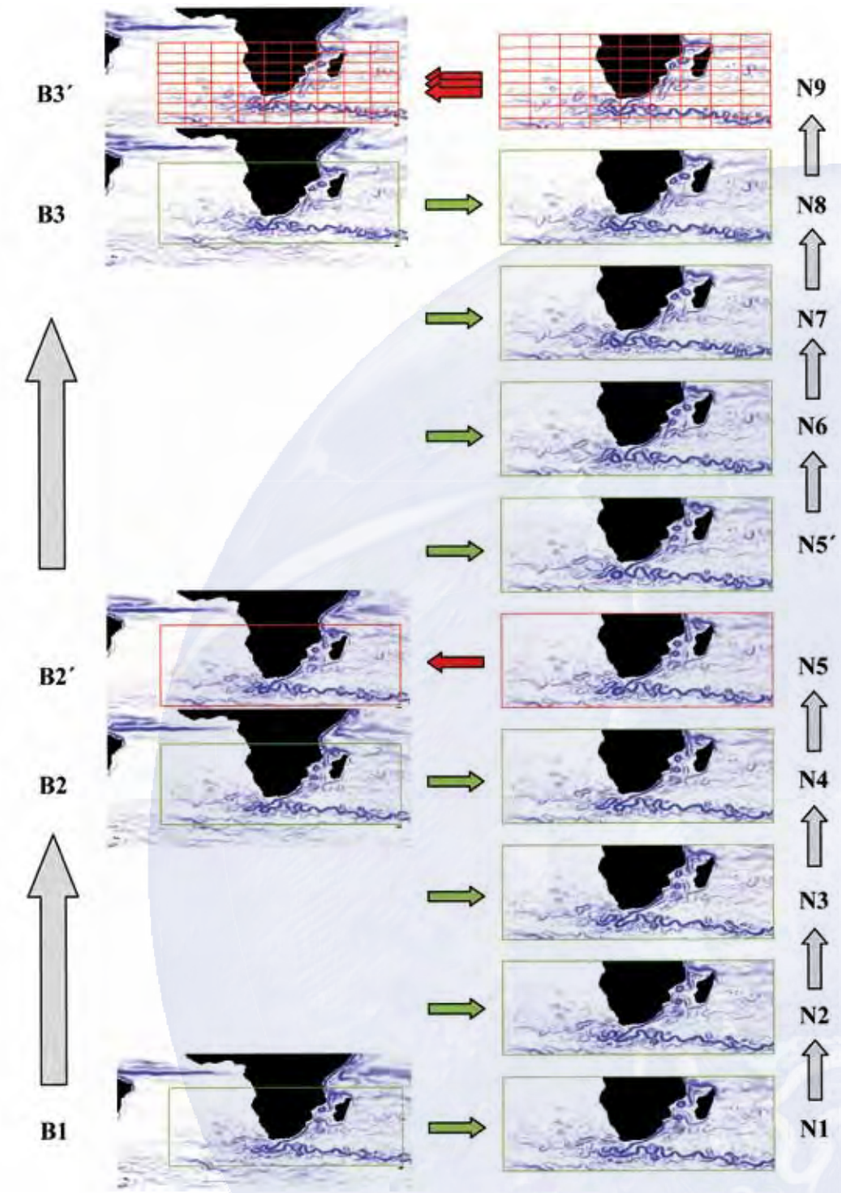

Figure 2: Schematics of AGRIF nesting. Time-stepping of the base (left) and nested (right) grids. The green boxes and arrows indicate an interpolation from the base grid onto the outer boundaries of the nest, the red ones an averaging of the outer and surface boundaries of the nest onto the base grid; the mesh indicates an averaging of the whole nest onto its base grid points in the Agulhas region. Grey arrows and numbers indicate the timesteps of base (Bn) and nest (Nn) and their respective updates (Bn', Nn').

the use of the vertical axis in form of vectors lead to a good performance on vector systems (up to 33% of the peak performance). With about 37 × 10^6 grid points and a high temporal (5-daily) resolution needed the output of a typical 50-year experiment is with more than 5 TB quite large.

The Agulhas model is a combination of a coarse-resolution global base model and a high-resolution nest around South Africa (Fig. 1). With a nominal grid size of 1/2° the base configuration (ORCA05) successfully simulates the large-scale wind-driven and thermohaline circulation [3]. It is forced by observed atmospheric conditions during the period 1958 - 2004. However, for a full representation of the Agulhas dynamics a high spatial resolution with grid scales less than 10 kilometres are needed; this is achieved here by nesting a 1/10° grid into the base model using AGRIF ("Adaptive Grid Refinement In Fortran", [4]).

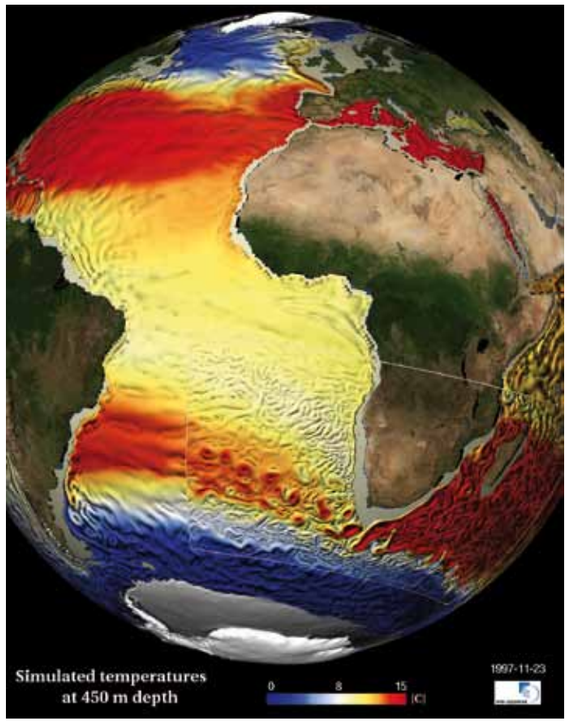

AGRIF recombines the subroutines in the model code via a preprocessing step and provides routines for interpolation and averaging between the two grids (Fig. 2). It allows both models to interact at any given base model time step where (i) the base updates the boundaries of the nest, (ii) the nest updates the coarser grid points of the base. Due to the different resolution of physical processes the nested model has to perform 4-5 timesteps before the base model is stepped forward in time. This effective and novel "two-way" nesting approach assures that this system not only simulates the current system around South Africa with great verisimilitude; it also allows unravelling how the explicitly simulated mesoscale variability in the Agulhas dynamics feeds back to the global ocean.

First analyzes addressed the importance of mesoscale processes, not only in the representation of the circulation around South Africa [5], but also in the net volume transfer between the Indian and Atlantic Ocean (the Agulhas leakage) [6]. Comparison with the coarse-resolution base model alone confirmed that Agulhas leakage is significantly overestimated at coarse resolution, and therefore in current IPCC-type coupled climate models. The explicit simulation of the upstream eddies originating from the Mozambique Channel and east of Madagascar that drift towards the Agulhas Current and do trigger the shedding of Agulhas Rings, however, do not make a significant impact on the variability of Agulhas leakage on timescales of a few years and longer.

Figure 3: Large-scale impact of the Agulhas dynamics. Temperatures and currents at 450 m depth in the high-resolution nest and its embedding in the global model. A similar figure was the basis for the Nature cover page on November 26, 2009 [8].

What is the effect of the Agulhas Current system on the large-scale circulation in the Atlantic Ocean? Comparing the circulation in solutions with and without the high-resolution Agulhas nest allowed identifying an intriguing contribution of the mesoscale Agulhas dynamics on decadal current fluctuations reaching far into the North Atlantic [7]. The dynamical signal originating south of Africa rapidly travels northward by boundary waves. In the tropical and subtropical North Atlantic the Agulhas-induced variability has similar amplitudes as the variability introduced by subpolar deepwater formations events, a mechanism that has been known for its climatic impact and that has been extensively studied in the past.

In addition to the decadal fluctuations by the Agulhas mesoscale another climate-relevant process emerges from the Agulhas dynamics. Observations report on the progressive poleward migration of the Southern Hemisphere westerly winds during the last two-three decades and linked those to anthropogenic forcing. Because of the sparse observational records it has not been possible to determine whether there has been a concomitant response of Agulhas leakage. Results with the nested Agulhas model showed that the transport of Indian Ocean waters into the South Atlantic via the Agulhas leakage has increased during the last decades in response to the change in wind forcing [8]. The increased leakage has contributed to the observed salinification of South Atlantic thermocline waters. Both model and historic measurements of South America suggest that the additional Indian Ocean waters have begun to invade the North Atlantic, with potential implications for a stabilization of the thermohaline circulation.

The findings highlight the importance for studying the Agulhas regime and its associated interoceanic transport as a prominent key region of the global thermohaline circulation.

References

[1] **The DRAKKAR Group**
Eddy-Permitting Ocean Circulation Hindcasts of Past Decades. Clivar Exchanges 12, 8-10, 2007

[2] **Madec, G.**
NEMO = the OPA9 ocean engine. Technical report, Note du Pole de modelisation, Institut Pierre Simon Laplace (IPSL), France, 2008

[3] **Biastoch, A., Böning, C. W., Getzlaff, J., Molines, J.-M., Madec, G.**
Causes of interannual-decadal variability in the meridional overturning circulation of the mid-latitude North Atlantic Ocean, J. Climate 21, 6599-6615, 2008

[4] **Debreu, L., Vouland, C., Blayo, E.**
AGRIF: Adaptive grid refinement in Fortran, Computers and Geosciences 34, 8-13, 2008

[5] **Biastoch, A., Beal, L., Casal, T. G. D., Lutjeharms, J. R. E.**
Variability and coherence of the Agulhas Undercurrent in a high-resolution ocean general circulation model, J. Phys. Oceanogr. 39, 2417-2435, 2009

[6] **Biastoch, A., Lutjeharms, J. R. E., Böning, C. W., Scheinert, M.**
Mesoscale perturbations control inter-ocean exchange south of Africa, Geophys. Res. Lett. 35, L20602, 2008

[7] **Biastoch, A., Böning, C. W., Lutjeharms, J. R. E.**
Agulhas leakage dynamics affects decadal variability in Atlantic overturning circulation, Nature 456, 489-492, 2008

[8] **Biastoch, A., Böning, C. W., Lutjeharms, J. R. E., Schwarzkopf, F. U.**
Increase in Agulhas leakage due to poleward shift of the Southern Hemisphere westerlies, Nature 462, 495-498, 2009

• Arne Biastoch

Leibniz-Institut für Meereswissenschaften (IFM-GEOMAR), Kiel

Laminar Flow Control in Vortex-deformed Swept-Wing Flows: Pinpoint Suction

Improving the fuel efficiency of aircraft has become an important task within the last decades. Not only do airlines benefit from saving increasingly expensive fuel but also the environmental aspect has gained growing interest and it will only be a matter of time until environmental laws limiting greenhouse gas emissions will be approved. Current commercials for newly designed aircrafts show the demand for more efficient airplanes: “The 787 Dreamliner is using 20% less fuel than any other airplane of its size” (Boeing) or “The A380 provides the lowest fuel burn per seat – which allows airlines to substantially reduce CO^2-emissions while achieving profitable, sustainable growth for decades to come” (Airbus).

To date, realized optimizations for new airplanes are limited to enhanced shaping, avoiding too rough surfaces, and engine improvement, but little potential is thought to be left in these fields except surface quality on aerodynamic surfaces. New concepts have therefore to be envisaged that consider the underlying fluid dynamics phenomena in detail. Inflight tests with a kind of shark-skin surface striving for turbulent boundary-layer drag reduction have shown improvements only in the range of very few percents. Laminar flow control (LFC) on the other hand provides a total drag reduction potential of e.g. 16% by realizing 40% laminar boundary-layer flow on wings and control surfaces of a current airliner [5]. Therefore it is the most promising candidate for expedient drag reduction.

Maintaining large regions of laminar boundary-layer flow has been proven for decades now in two-dimensional situations by applying boundary-layer suction which efficiently delays laminar-turbulent transition. A typical airliner wing, however, is swept back (cf. figure 1) to allow for higher cruise speed at acceptable pressure drag. The evolving crossflow component inside the boundary layer causes a new, dominant instability mechanism, and a straightforward implementation of the two-dimensional suction setups is not possible. For a three-dimensional boundary layer it turns out that - typically steady - longitudinal crossflow vortices evolve due to the new primary instability of the flow. While still being laminar these vortices - if grown to large amplitudes - are highly unstable to ubiquitous unsteady background disturbances. Due to the extremely large growth rates of this so-called secondary instability

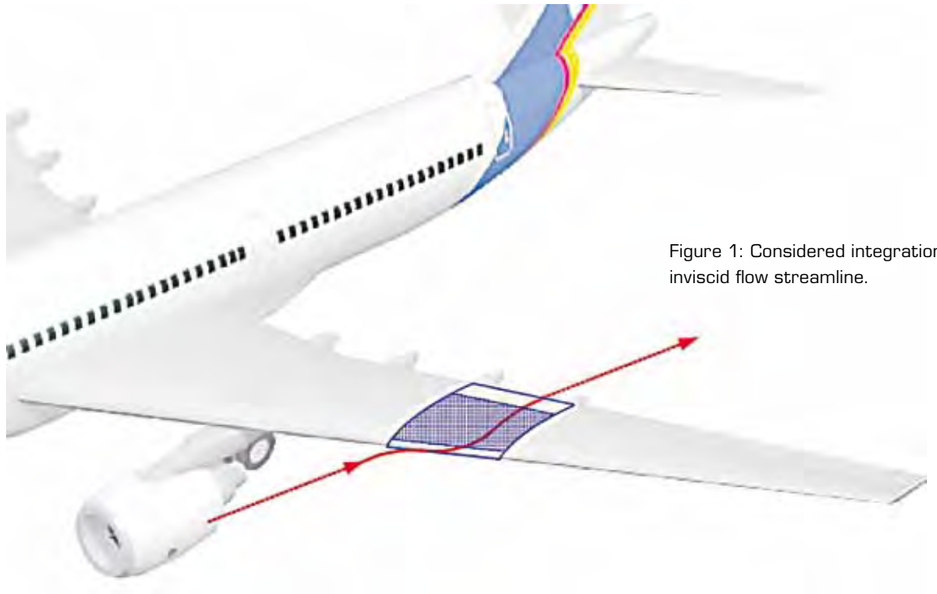

Figure 1: Considered integration domain with inviscid flow streamline.

Applications

laminar-turbulent transition sets in rapidly, typically after only few percents of the airfoil chord length. The crossflow vortices are generated by even minute surface non-uniformity.

Boundary-layer suction diminishes the crossflow by sucking high-momentum fluid to the wall, and thus also attenuates crossflow instability. However, the typically applied discrete suction

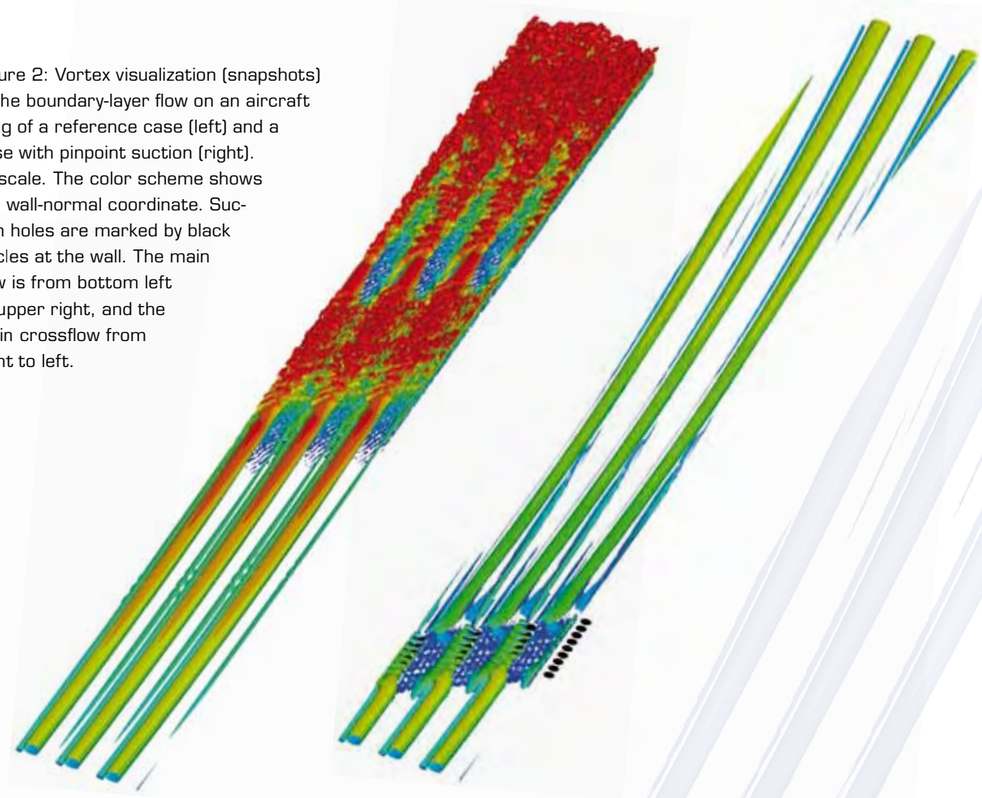

Figure 2: Vortex visualization (snapshots) in the boundary-layer flow on an aircraft wing of a reference case (left) and a case with pinpoint suction (right). To scale. The color scheme shows the wall-normal coordinate. Suction holes are marked by black circles at the wall. The main flow is from bottom left to upper right, and the main crossflow from right to left.

holes can generate, on the other hand, relatively large initial crossflow-vortex disturbances, jeopardizing the LFC.

Improved suction concepts for this kind of three-dimensional wing flows had therefore to be developed. Messing & Kloker [4] proposed an idea called distributed flow deformation (DFD), in particular formative suction. By designing a suitable slot-suction panel useful vortices - with a much closer spanwise spacing than the turbulence-triggering ones - are continuously excited and maintained that are known to be stable with respect to secondary instability and suppress the nocent vortices. Laminar-turbulent transition could be delayed significantly.

A new idea of directly influencing large-amplitude crossflow vortices and secondary instabilities called pinpoint suction is currently developed [1,2,3]. The scenario considered contains the harmful, secondarily unstable crossflow vortices that develop naturally. Localized, strong suction through few holes only at the updraft side of each vortex, i.e. the locally most unstable region with high-shear layers, directly reduces the secondary growth while also reducing the vortex strength.

If the exact position of the vortex is known this suction setup turns out to be very effective and the transition location can be shifted far downstream. Note that this is an inherently non-linear three-dimensional process compared to standard suction with its much lower suction velocity.

Figure 2 shows two snapshots of vortical structures. On the left side a reference case is shown where three steady crossflow vortices can be observed. A pulse-like disturbance is triggered upstream of the shown domain at extremely low amplitudes, eventually undergoing secondary instability. These disturbances grow rapidly and soon finger-like secondary structures can be detected that trigger laminar-turbulent transition. On the right side the development of the same vortices (containing the identical pulse) is shown, now being subject to pinpoint suction (cf. also figure 3). Nine closely spaced holes per vortex are placed at the updraft side of the respective vortex (black holes at the wall) and transition is prevented in the considered domain.

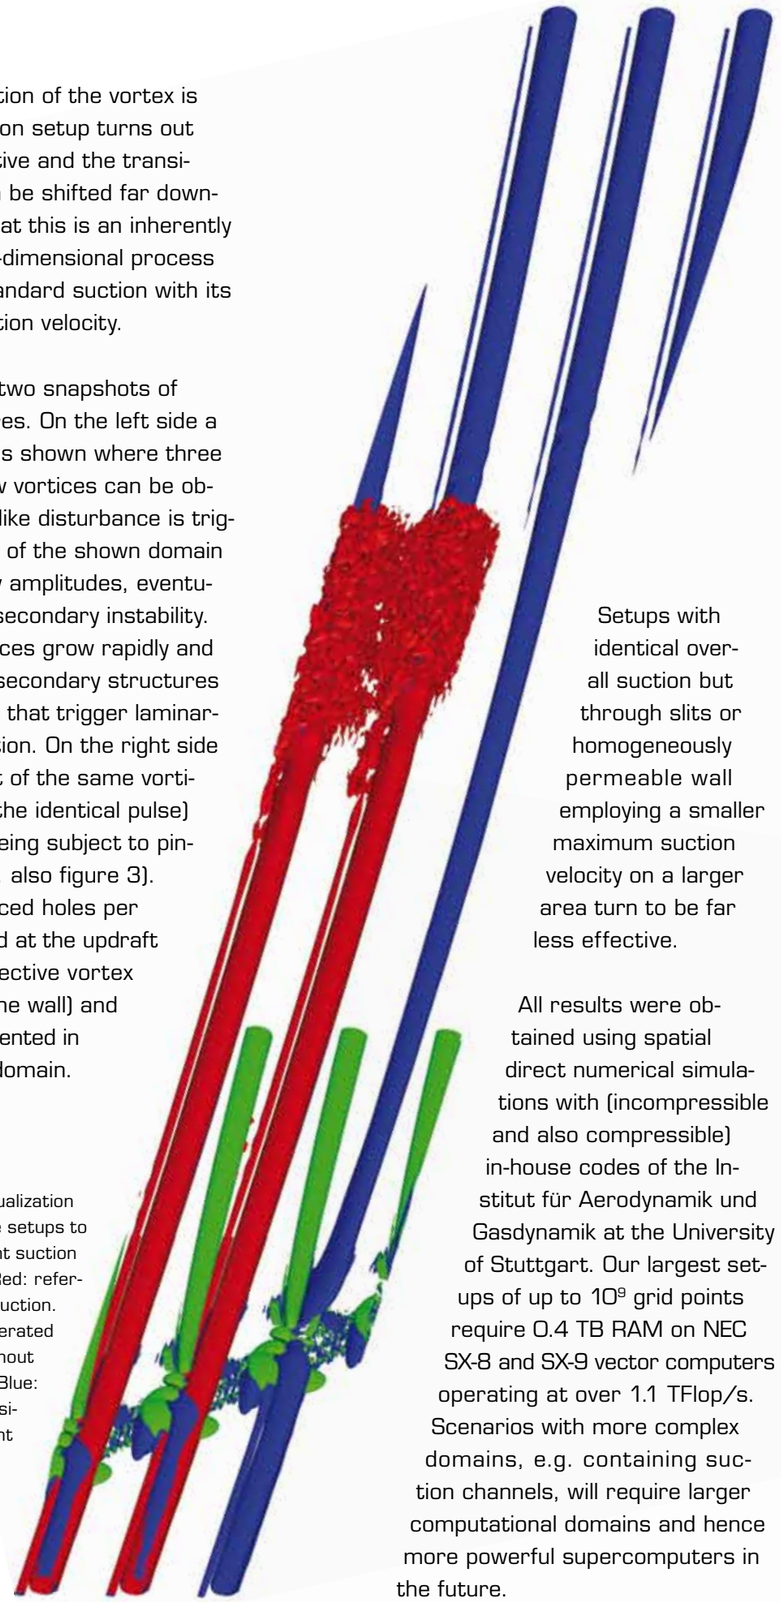

Figure 3: Vortex visualization (snapshots) of three setups to illustrate the pinpoint suction concept. To scale. Red: reference case without suction. Green: vortices generated by suction holes without oncoming vortices. Blue: (non-linear) superposition – applied pinpoint suction.

Setups with identical overall suction but through slits or homogeneously permeable wall employing a smaller maximum suction velocity on a larger area turn to be far less effective.

All results were obtained using spatial direct numerical simulations with (incompressible and also compressible) in-house codes of the Institut für Aerodynamik und Gasdynamik at the University of Stuttgart. Our largest setups of up to 10^9 grid points require 0.4 TB RAM on NEC SX-8 and SX-9 vector computers operating at over 1.1 TFlop/s. Scenarios with more complex domains, e.g. containing suction channels, will require larger computational domains and hence more powerful supercomputers in the future.

References

[1] Bonfigli, G., Kloker, M.
Secondary instability of crossflow vortices: validation of secondary instability theory by DNS, 2007, Journal of Fluid Mechanics, 583, pp. 229-272

[2] Friederich, T., Kloker, M.
Direct Numerical Simulation of Swept-Wing Laminar Flow Control using Pinpoint Suction, in: High Performance Computing in Science and Engineering '10 (eds. Nagel, W.E., Kröner, D.B., Resch, M.M.), Transactions of the HLRS 2010, pp. 231-250, Springer

[3] Friederich, T., Kloker, M.
Numerical Simulation of Crossflow-Transition Control using Pinpoint Suction, 2011, in: New Results in Numerical and Experimental Fluid Mechanics VIII (eds. n.n.), NNFM, reviewed contributions to the 17. STAB/DGLR-Symposium, Nov. 2010, 8 pages, in press, Springer

[4] Messing, R., Kloker, M.
Investigation of suction for laminar flow control of three-dimensional boundary layers, 2010, Journal of Fluid Mechanics, 658, pp. 117-147

[5] Schrauf, G.
Status and perspectives of laminar flow, 2005, The Aeronautical Journal (RAeS), 109, no. 1102, pp. 639-644

• Tillmann A. Friederich
• Markus J. Kloker

Institut für Aerodynamik und Gasdynamik, University of Stuttgart

Spring 2011 • Vol. 9 No. 1 • inSiDE 21

Fluid-Structure coupled Flow Simulations of Helicopter Rotors

Introduction

Today, helicopter technology still poses several unsolved problems in aerodynamics. The accurate numerical prediction of certain fluid-structure interactional flow phenomena of flight dynamic relevance is one example for such problem areas. Prototyping of helicopter aircraft is incrasingly done based on Computational Fluid Dynamics (CFD). Helicopter aeromechanical studies call for a more wholistic approach in the simulation of flow fields in that they require the incorporation of elastic deformations of flexible structures such as the main rotor blades into the computational tool chain. This results in a CFD-CSD (Computational Structural Dynamics) coupled simulation. Additionally, a procedure for trimming the rotor towards some prescribed flight dynamic state has proven essential. Only in fulfillment of these presuppositions can reasonable comparability of simulation and experiment be guaranteed.

Most state-of-the-art simulation environments today use a a so-called structured grid approach, i.e. the fluid volume to be simulated is subdivided into cuboidal elements for which the equations of motion for the fluid can be solved in an nicely ordered manner along the directions of three-dimensional space. Recently, CFD solution methods following a different approach termed unstructured have gained popularity. Here, the elements for the discretization of the volume to be simulated can be of quite arbitrary shape, giving up the nice structure of the system of equations in return of increased geometrical flexibility.

As mentioned above, quantitative numerical investigations of interactional phenomena problems is one of the to date unsolved issues in helicopter aerodynamics. An example within this field is the so-called tail shake phenomenon [1], where in fast forward flight conditions interations of the main rotor wake and the tail boom and fin structure excite a lateral bending close to the frequencies corresponding to the lower elastic modes of the fuselage structure. Figure 1 displays this scenario (figure taken from [1]). This phenomenon exhibits an

Figure 1: Rotor wake / tail interations of the tail shake phenomenon

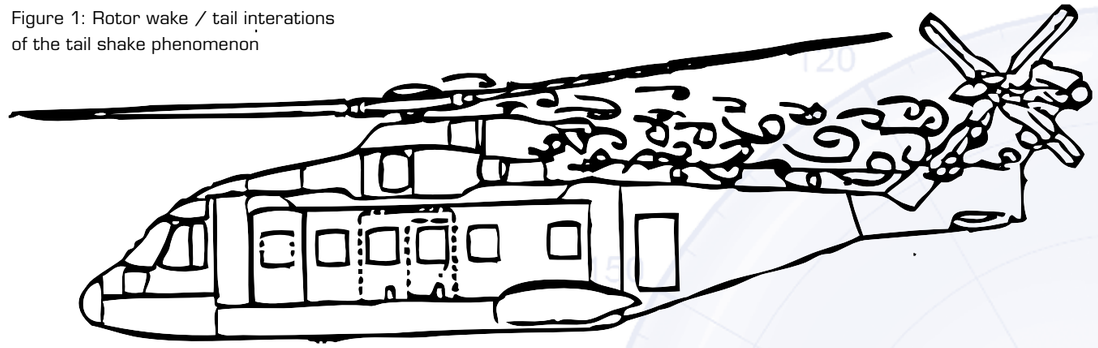

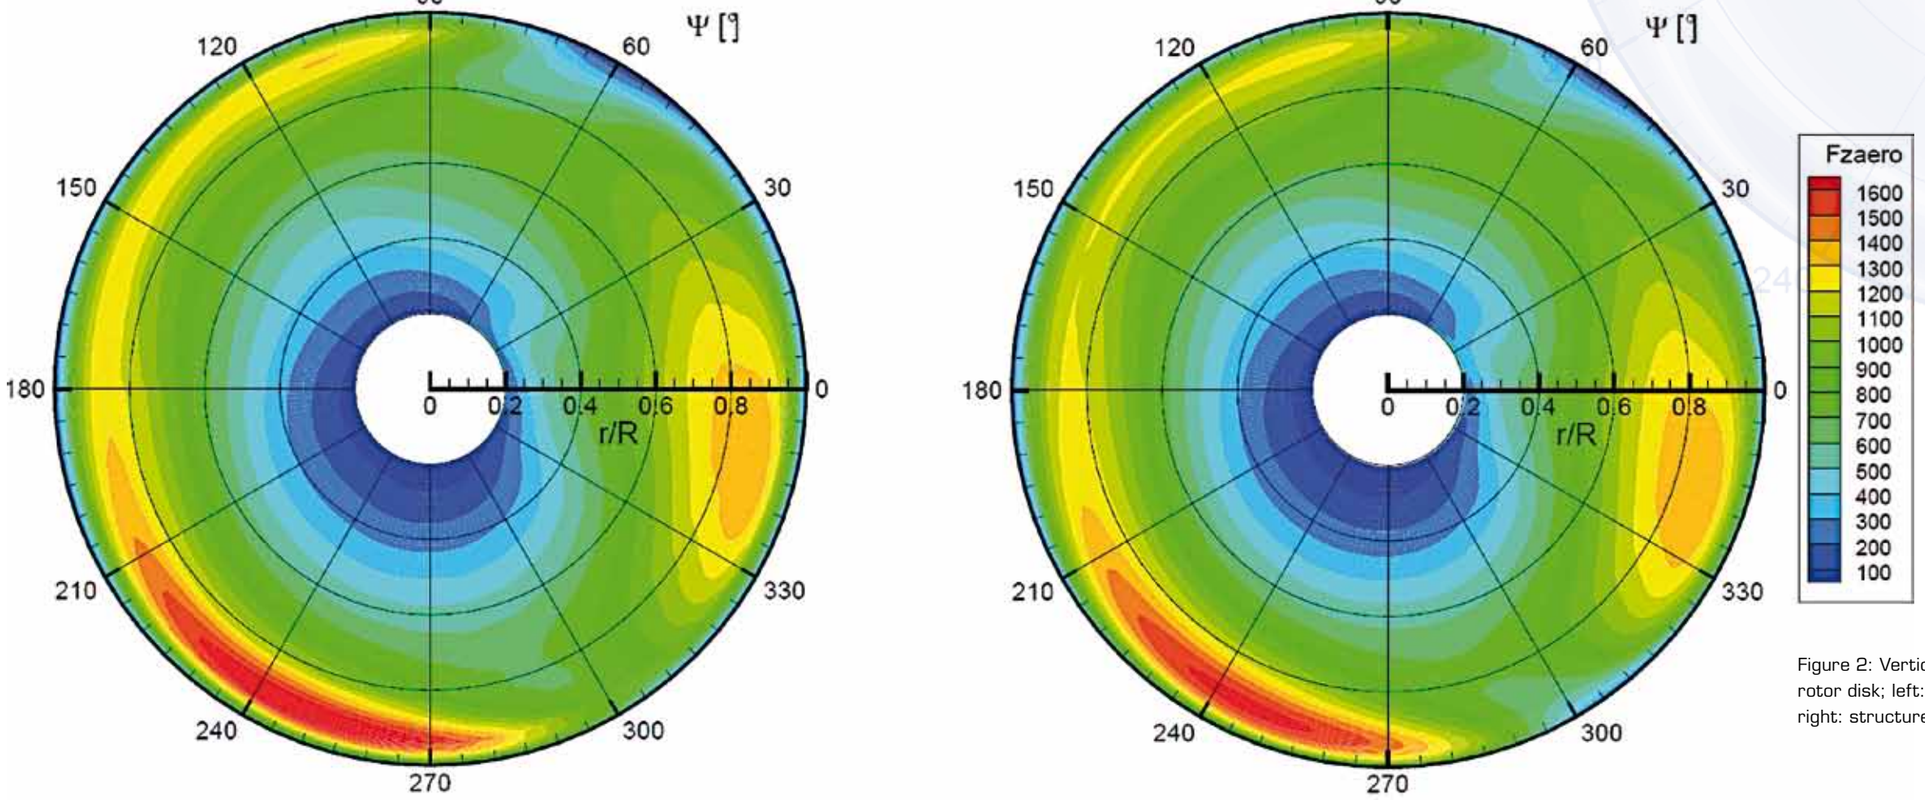

Figure 2: Vertical forces on the rotor disk; left: unstructured, right: structured solver

undesirable random character of unsteady nature which can be felt by the flight crew as lateral “kicks” [1]. To date, CFD methods are incapable of predicting and explaining the character of tail shake even before early prototype flight testing and as a consequence many well-known helicopter types such as the Eurocopter EC135 [2] or the Boeing AH-64D™ Longbow Apache™ [3] showed this effect in early flight testing.

Simulation

The CFD simulation of such phenomena requires a very detailed modelling of the structure of the flow field and thus high geometrical detail especially in the hub area of the main rotor is desired. Here, the above mentioned conventional structured grid approach suffers from the drawback of excessive manual time consumption and eventually becomes impossible. Therefore, a new toolchain based on the unstructured grid approach has been developed. First results of the new simulation environment are compared to the already exiting standard structured-grid based toolchain [4]. In both cases, a fluid-structure coupling procedure termed the weak coupling approach [5] has been employed. Here, periodic data is exchanged between CFD and CSD. In parallel to this, an algorithm for trimming the rotor ensures that specified rotor forces and moments are met. Hereby, control inputs which are usually steered by the pilot such as the collective and cyclic pitch angles are adapted in order to eventually meet these loads also termed trim objectives.

In a present study, helicopter rotor CFD-CSD calculations where performed making use of the standard structured as well as the new unstructured approach. The basis of the current investigation were wind tunnel experiments of a generic helicopter configuration which were performed in the open test section of the German-Dutch wind tunnel (DNW). The simulation contained a four-bladed rotor in slow forward flight conditions.

Figure 2 shows the comparison between the two respective toolchains in the vertical forces on the rotor disc plane where the flight direction is directed from right to left. Both approaches yield a very similar distribution of the forces on the advancing (0..180°) as well as on the retreating side (180..360°). CFD methods also allow for a detailed analysis of the fluid volume. In Figure 3 a vortex visualization of the complete flow field past the rotor is displayed. The vortices shed at the tip and the inner root section of the blades as well as the great rolled-up vortices to the left and right of the entire rotor propagating backwards are clearly visible. The rotor is rotating in a clockwise sense, showing distinct suction areas of low pressure at the outer front parts of the blades (blue color on the blades). Two vertical slices through the vortex system indicate the vortex core locations (blue color).

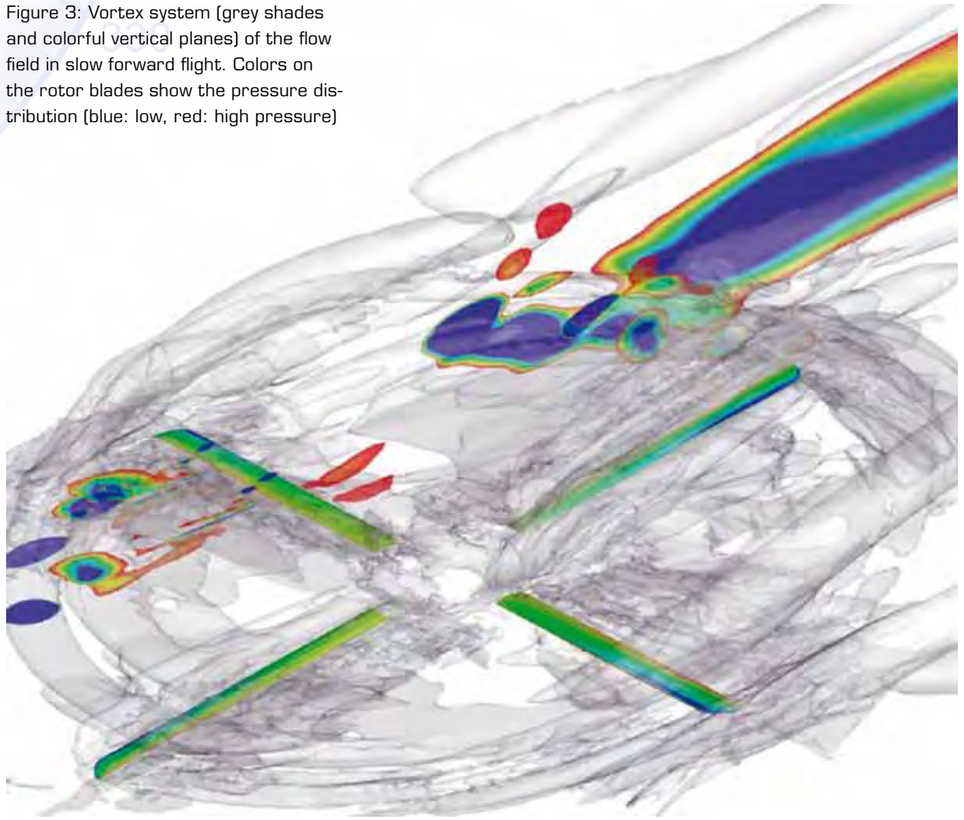

Figure 3: Vortex system (grey shades and colorful vertical planes) of the flow field in slow forward flight. Colors on the rotor blades show the pressure distribution (blue: low, red: high pressure)

Performance

The present calculations were performed on the NEC Nehalem cluster at the High Performance Computing Centre (HLRS) in Stuttgart. While the standard structured toolchain poses limitations on the rescalability of the computational setup to different numbers of computing units, the new unstructured approach allows for easy repartitioning and customization to almost arbitrary numbers of computing processes. Additionally, the new toolchain showed great potential in the scalability to several hundreds of computing processes in comparably small setups. Performance of the code amounted to 11 GFlops which translates into a node-wide peak performance of 12% making use of the NEC Nehalem's Intel Xeon X5560 processor. This flexibility will allow for the computation of extremely challenging simulations in the field of helicopter aerodynamics both in terms of problem size as well as geometrical complexibility in the advent of the access to even larger computational resources.

References

[1] **de Waard, P. G., Trouvé, M.**
Tail shake vibration, 1999, NLR-TP-99505, NLR, The Netherlands

[2] **Kampa, K., Enenkl, B., Polz, G., Roth, G.**
Aeromechanic Aspects in the Design of the EC135, 1999, Proceedings of the 23rd European Rotorcraft Forum, Dresden, Germany

[3] **Hassan, A. A., Thompson, T., Duque, E. P. N., Melton, J.**
Resolution of Tail Buffet Phenomenon for AH-64D™ Longbow Apache™, 1997, Proceedings of the 53rd Annual Forum of the American Helicopter Society, Virginia, USA

[4] **Kroll, N., Eisfeld, B., Bleecke, H. M.**
FLOWer, in: Notes on Numerical Fluid Mechanics Vieweg Braunschweig, 1999, Vol. 71, pp. 58-68

[5] **Altmikus, A., Wagner, S., Beaumier, P., Servera, G.**
A Comparison: Weak versus Strong Modular Coupling for Trimmed Aeroelastic Rotor Simulations, American Helicopter Society, 58th Annual Forum, Montreal, Canada, 2004

• Felix Bensing
• Manuel Keßler

Institut für Aerodynamik und Gasdynamik, University of Stuttgart

Joint numerical and theoretical Investigations of fully Developed Fluid Turbulence

Turbulence research since the days of Leonardo da Vinci, Leonhard Euler and their successors has continuously triggered the advancement of basic science and has contributed to or even laid the foundations of various theoretical concepts like dynamical systems theory, statistical physics of systems far from equilibrium, bifurcation theory and chaos, to name but a few. Apart from being a fascinating theoretical and mathematical problem in its own right, any progress in this field is appreciated in a variety of scientific disciplines ranging from astrophysics over climate system theory to civil engineering and medicine. Still, the statistical description of fully developed turbulence remains one of the most challenging problems of classical physics. In comparison with other fields of science, like e.g. high energy physics, where field theoretic methods allow for calculations of many central quantities with incredible precision, derivations of fundamental statistical properties of turbulent fields directly from the fluid dynamical equations are, with few exceptions, still far ahead of us. The main reason for this is the non-linear and non-local character of the governing equation of motion, the Navier-Stokes equation. These properties are on the one hand responsible for the highly complex spatio-temporal patterns observable in Direct Numerical Simulations (DNS), see Figure 1. On the other hand, the same properties are also the origin of the major obstacle on the way to a statistical description from first principles, namely the closure problem; when deriving evolution equations for moments or probability density functions (PDFs) from the Navier-Stokes equation, the statistical equations contain unclosed terms, which necessitate information from higher-order moments or more points in space. When tackling the problem by analytical means only, the validity of the results depends sensitively on the approximations made, and a too crude approximation can lead to unphysical results. This is one reason why turbulence simulations on modern supercomputers have become an indispensable tool to understand the nature of turbulent flows. Hence one of the goals of the current project is to supplement the analytical treatment of the statistical problem with results from direct numerical simulations. Moreover, temporal correlations of the flow are studied along fluid particle trajectories with special respect to the acceleration acting on the particles.

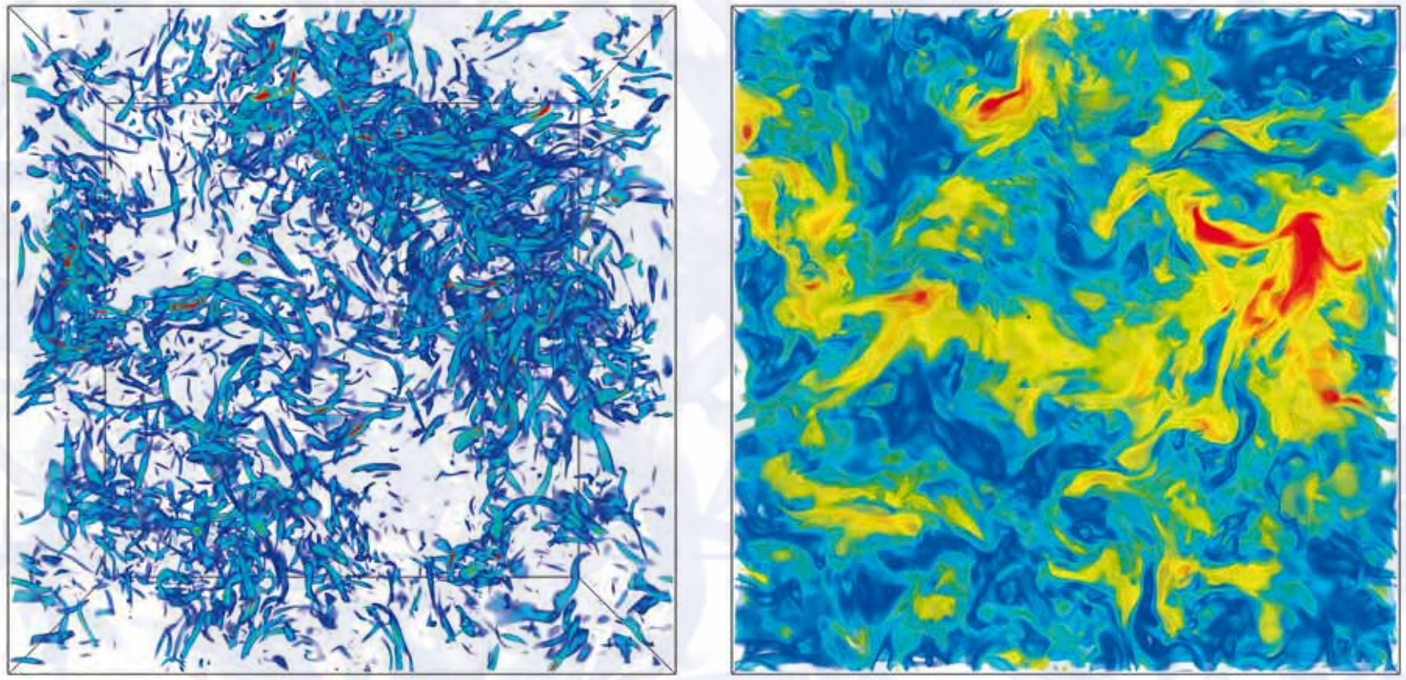

Figure 1: Volume rendering of the magnitude of vorticity (left) and velocity (right) from a simulation with 512^3 grid points. The vorticity tends to organize into thin filaments forming an entangled global structure. The velocity appears less clearly structured but displays long-range correlations. Volume rendering produced by VAPOR: www.vapor.ucar.edu

Numerical Details and conducted Simulations

For the direct numerical simulations we make use of the vorticity formulation of the incompressible Navier-Stokes equation. The method used is a standard pseudospectral method [1], in which most of the computations (evaluation of derivatives etc.) are performed in Fourier space. The nonlinear term is treated in real space, reducing the computational efforts from $O(N^2)$, due to the otherwise arising convolution sums, to O(N log N) needed for the Fourier transforms. Here N is the total number of grid points. Adaptive time stepping is implemented employing a third-order Runge-Kutta scheme [2]. We conduct highly resolved simulations with up to 1024^3 grid points on a periodic domain, integrated for up to tens of thousands of time steps. Our code is MPI parallelized and employs a slab domain decomposition, enabling us to effectively use up to 1,024 cores at a resolution of 1,024 grid points. Parallel IO is incorporated by the use of MPI-IO. The Fourier transforms are performed by the freely available library FFTW [3]. Additionally, we have implemented the possibility to follow trajectories of (so-called Lagrangian) tracer particles, where relevant quantities along the tracer trajectories are interpolated from the turbulent fields with a tricubic interpolation scheme. The HLRB II at LRZ turns out to be an optimal platform for the currently implemented parallelization scheme as the individual cores are comparably performant. The results of a scaling test are presented in figure 2.

A typical simulation basically consists of two stages. First, an artificial large-scale initial condition decays for some large-eddy turnover times, during

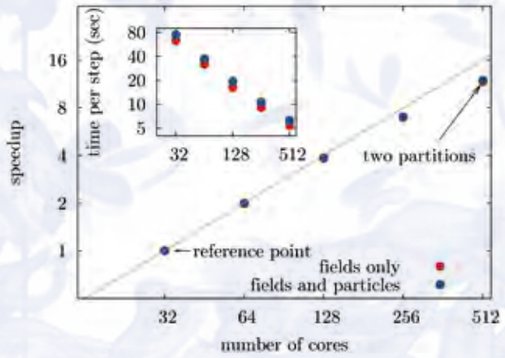

Figure 2:
Scaling results of a 1,024³ run with and without ten million tracer particles. The gray line indicates ideal scaling. The inset displays the time needed for a single simulation time step.

which a turbulent flow develops. Then an external forcing is applied, and the system eventually approaches statistical stationarity. After this preparation of proper initial conditions, the actual simulation is performed. Here, the flow field is advanced in the statistically stationary state. During this period fields (velocity, vorticity, velocity gradients etc.) are stored with a sampling rate sufficient to form a statistical ensemble. The statistical analysis is performed during the post-processing stage. Optionally tracer particles are advected with the flow and stored frequently. In total, a typical 1024^3 run requires several tens of thousands of CPU hours and easily produces a terabyte of data. Within the ongoing project runs with resolutions between 256^3 and 1024^3 grid points with Taylor-based Reynolds numbers ranging from about 75 to 250 have been performed, giving insight into the Reynolds number dependence of the statistical quantities under consideration as well as resolution issues. It has turned out that long simulation durations are important for the constitution of a proper statistical ensemble. Hence, we have performed simulations for more than 100 large-eddy turnover times at a resolution of 512^3 grid points, where special emphasis has been put on an adequate resolution of the small-scale features of turbulence.

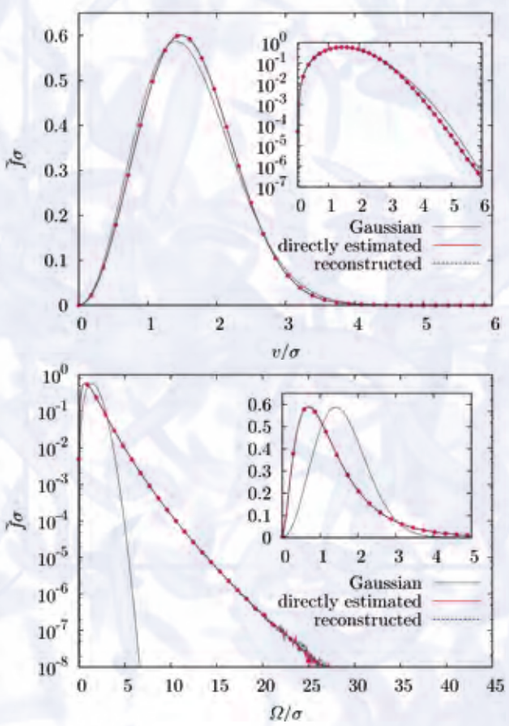

Figure 3: Directly estimated and theoretically reconstructed PDFs of the velocity v (top) and vorticity Ω (bottom) normalized by the standard deviation σ of the corresponding quantity. The theory describes the sub-Gaussian tails of the velocity PDF as well as the strongly super-Gaussian tails of the vorticity PDF.

Scientific Results

Within the ongoing project a number of scientific questions have been addressed so far. The first question on the agenda has been to investigate the statistics of the single-point velocity and vorticity probability density functions within the framework of the Lundgren-Monin-Novikov hierarchy [4,5,6], a theoretical framework for the statistical description of turbulent flows introduced in the late sixties of the last century. By exploiting statistical symmetries, it has been possible to derive relations that express the single-point statistics in terms of local correlations. For example, it has been shown that the single-point velocity PDF may be expressed in terms of the conditional statistics of the pressure gradient, the external forcing and the rate of energy dissipation. For the case of the vorticity statistics, a relation expressing the PDF in terms of vortex stretching and enstrophy dissipation has been established. DNS results then have been used to assess this conditional statistics. Thereby a physical discussion of the unclosed terms has been achieved, finally explaining the often discussed slightly sub-Gaussian shape of the velocity PDF as well as the super-Gaussian shape of the vorticity PDF [7,8]. A comparison between directly estimated PDFs and the results of a consistency check of the theoretical framework is presented in figure 3. Also two-point statistics has been investigated, especially focusing on the interaction of different spatial scales. Furthermore, the simulations have been used to gather acceleration data along Lagrangian tracer particle trajectories. Here, we have analyzed the multi-time PDFs of the acceleration with respect to Markovian properties. As a result, we have found that only for very long time lags the acceleration along the trajectory can be approximated by a Markov process. This especially has implications for the modeling of Lagrangian tracer particle trajectories, which is, for example, relevant in the context of mixing and dispersion.

Outlook

By joint numerical and theoretical efforts the current project has led to a comprehensive characterization of the single-point statistics of the velocity and vorticity field, and insights into the temporal correlations in terms of the Lagrangian acceleration statistics have been gained. One of the goals for future research is to obtain deeper insights into the multi-point statistics of turbulence. This is related to one of the most challenging questions of theoretical turbulence research, namely which aspects of the fine-scale structure of turbulence are responsible for the energy and enstrophy transfer across scales. Furthermore, we aim for a deeper understanding of the Lagrangian description of fully developed turbulence with respect to the origin of Lagrangian intermittency and implications for stochastic models for particle trajectories. Direct numerical simulations of fully developed turbulence on modern supercomputers undoubtedly will continue to foster new theoretical ideas.

Acknowledgments

The authors gratefully acknowledge computational resources granted by LRZ under project h0963 and thank the staff of LRZ for continuous support.

References

[1] **Canuto, C., Hussaini, M. Y., Quarteroni, A., Zang, T. A.**
Spectral Methods in Fluid Dynamics, Springer-Verlag, Berlin, 1987

[2] **Shu, C.-W., Osher, S.**
Efficient implementation of essentially non-oscillatory shock-capturing schemes, J. Comput. Phys., 77(2):439-471, 1988

[3] www.fftw.org

[4] **Lundgren, T. S.**
Distribution functions in the statistical theory of turbulence, Physics of Fluids, 10(5):969-975, 1967

[5] **Monin, A. S.**
Equations of turbulent motion, Prikl. Mat. Mekh., 31(6):1057, 1967

[6] **Novikov, E. A.,**
Kinetic equations for a vortex field, Soviet Physics-Doklady, 12(11):1006-1008, 1968

[7] **Wilczek, M., Friedrich, R.**
Dynamical origins for non-Gaussian vorticity distributions in turbulent flows, Phys. Rev., E 80(1):016316, 2009

[8] **Wilczek, M., Daitche, A., Friedrich, R.**
Theory for the single-point velocity statistics of fully developed turbulence, EPL, 93(3):34003, 2011

• Michael Wilczek
• Oliver Kamps
• Anton Daitche
• Rudolf Friedrich

Institute for Theoretical Physics, Westfälische Wilhelms-Universität Münster

A new View on Geometry and conditional Statistics in Turbulence

Introduction

Turbulent vector and scalar fields are locally subject to sudden and unpredictable changes. Even attempts to describe the statistics of averaged quantities have failed due to the non-linear and non-local nature of the underlying governing equations. A means to overcome this problem is to decompose the turbulent field into smaller structures which are more amenable to statistical treatment. In order to make these statistics meaningful, the geometry of these pieces should, however, be determined by the turbulent fields themselves. This has led us to define so-called dissipation elements, which are space-filling regions within a turbulent scalar field, and to their conditional statistics.

Turbulence

Over more than a century the understanding of turbulence has remained a big challenge of great importance to both science and engineering. From the aspect of applications, turbulence research aims mainly at the improvement of models for the numerical simulation of engineering flows. In recent years a very active stream in this direction is reliable sub-grid models in Large Eddy Simulations (LES). Fundamental research may provide deeper insights into the physics beyond proposed sub-grid models.

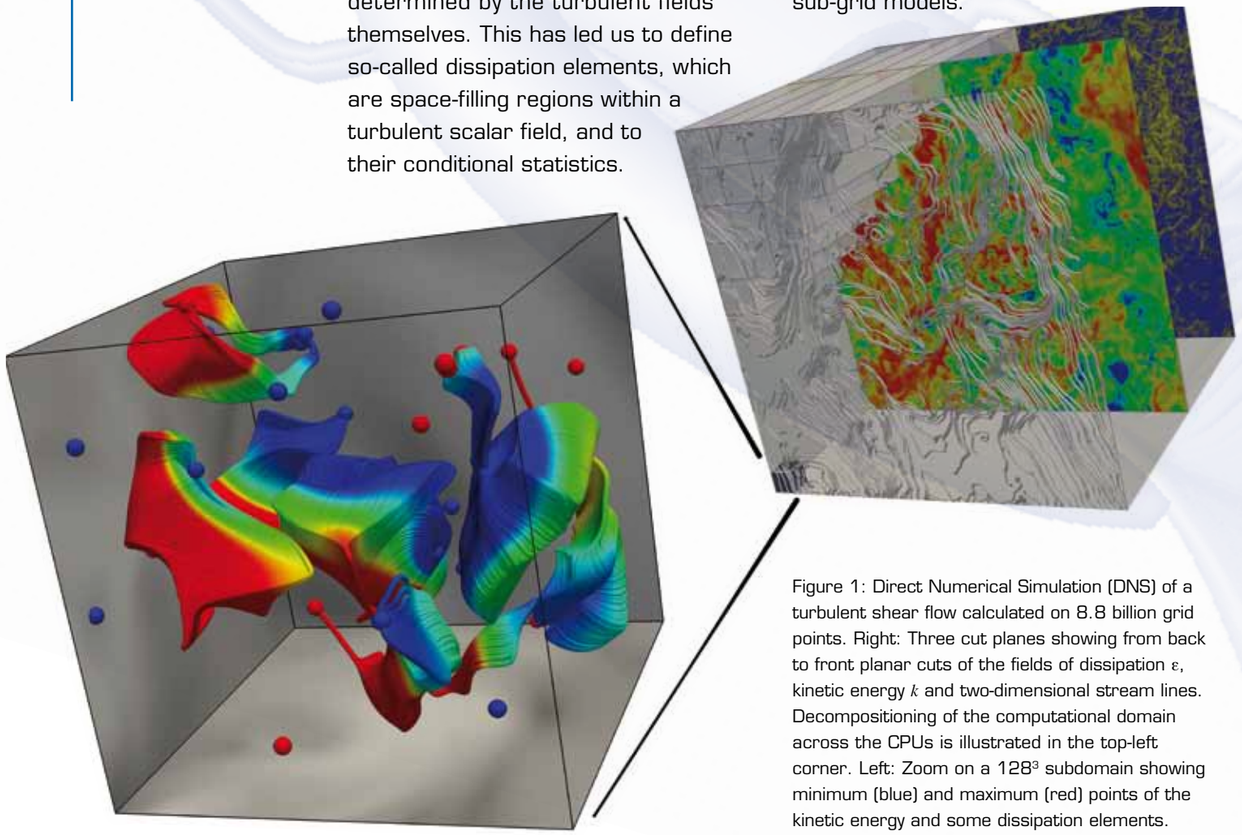

Figure 1: Direct Numerical Simulation (DNS) of a turbulent shear flow calculated on 8.8 billion grid points. Right: Three cut planes showing from back to front planar cuts of the fields of dissipation ε, kinetic energy k and two-dimensional stream lines. Decompositioning of the computational domain across the CPUs is illustrated in the top-left corner. Left: Zoom on a 128^3 subdomain showing minimum (blue) and maximum (red) points of the kinetic energy and some dissipation elements.

As a continuum field phenomenon, turbulence is in principle infinite-dimensional and strongly nonlocal and nonlinear. Although the Navier-Stokes equations are formally deterministic, turbulence dynamics are by no means the same. In addition, nonlocality and perturbations from boundary conditions make averaged equations to be unclosed. These features make mathematics alone inadequate to treat the problem. Meanwhile no appropriate small parameter has yet been identified for perturbation analysis.

Therefore, it is promising to approach turbulence from a different point of view rather than classical mathematical analysis like Fourier analysis alone. One possible idea for solution is to decompose the flow field into relatively small units. The following principle may be taken for granted that, by decomposing an entire field into subunits, the complexity may be reduced and detailed structures can better be understood. If there is any relation among these decomposed units, it will then be possible to reproduce at least statistical properties of the original flow field. In some sense the phenomenological 1941 theory by Kolmogorov (K41), a great milestone in turbulence research, follows the same spirit. In the K41 theory, the flow field is viewed as a set of eddies of different size. These eddies represent kinematic and dynamic features of turbulence, where the smallest ones follow a scaling principle of their lower order moments. For higher order moments unknown features related to geometry seem to limit the use of this concept.

In the following, the work of the authors on geometrical features of turbulence in physical space is briefly reviewed. Although the existing results do improve our knowledge, work in this area is by far incomplete.

Flow Geometry

There have been many attempts to define the structure of geometrical objects in turbulence [1]. For three-dimensional flows geometry can be represented as point, line, surface or volume. In turbulence, it is an important topic to study the kinematics and dynamics of points, especially critical points, which are those with degenerated gradient matrices. These are for instance stagnation points at which the orientation of streamlines is indeterminate or extremal points of scalar fields where the orientation of gradient trajectories is indeterminate. Perry and Chong [2] show the kinematic properties of critical points and to some extent the geometrical features of the flow field that can be presented by those isolated points. Davila and Vassilicos [3] discuss the dynamical properties of stagnation points. Going from critical points to higher dimensional objects as surfaces, results become more informative. An interesting suggestion was made by Wray and Hunt [4], who tried to find a systematic way to subdivide the whole flow field into four types of space-filling

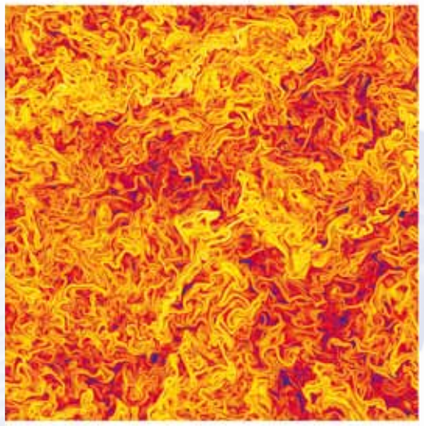

Figure 2: Slice of instantaneous scalar dissipation χ for the 2048^3 case showing very fine filamented structures of the turbulent flow field. The scalar dissipation is a highly intermittent quantity and regions of high χ are confined in localized regions of space.

regions, according to the characteristic values of the second invariant of the velocity derivative tensor Q, as well as the pressure p. The definition of these regions contains, however, an arbitrariness introduced from the setting of the threshold values of the scalars Q and p.

It is evident that the aforementioned structures are more illustrative and descriptive rather than well defined. Relatively early in turbulence research Corrsin [5] asked the following questions: (1) what types (of geometry) are naturally identifiable in turbulent flows? (2) what roles do they play or what properties do they have? (3) what stochastic games can we invent which share some of the difficulties of the turbulent case, but are more treatable? To answer these questions and make the study of geometry more quantitative than qualitative, one needs to construct methods which can identify geometrical elements unambiguously and find a clear definition of the quantities to be sampled.

Dissipation Element Analysis

In the following, we will consider a frozen scalar field, which has the properties of a Morse function and is locally smoothed by diffusion. The scalar can be a passive scalar like temperature or the instantaneous kinetic energy or its dissipation, for example. Starting from any material point, a gradient trajectory can be determined by proceeding in descending and ascending directions until two extremal points are reached which are a minimum and a maximum point, respectively (figure 3, left). The spatial region of all material points whose trajectories share the same pair of minimum and maximum points defines a dissipation element [6]. Each material point belongs to one, and only one dissipation element, so that by definition dissipation elements are space-filling. Examples of dissipation elements from 3D Direct Numerical Simulations (DNS) are shown in figure 3 (right).

	DNS 1	DNS 2	DNS 3	empirical
Re_λ	98.7	125.0	170.0	-
$c_{\varepsilon 1}$	0.425	0.763	1.20	1.44
$c_{\varepsilon 2}$	0.457	0.923	1.64	1.90

Table 1: Model coefficients $c_{\varepsilon 1}, c_{\varepsilon 2}$ for the k-ε model derived using dissipation element analysis.

Dissipation element analysis is not limited to the turbulent velocity field. Taking for example the fluctuating mixture fraction Z as the underlying scalar field allows the physical interpretation of dissipation elements in the context of the flamelet approach in non-premixed combustion [7]. An additional application [8] of this new approach is the definition of the coefficients $c_{\varepsilon 1}, c_{\varepsilon 2}$ of the widely used k-ε model for turbulence. Here, the turbulent kinetic energy k and the dissipation ε are used to close the Navier-Stokes equations. The properties of dissipation ε (figure 1, right) were analysed using dissipation element analysis. The statistics of the entire field were described by the conditioned and volume averaged mean properties of the dissipation elements. Hence, an equation of the temporal evolution of the mean dissipation $\langle\varepsilon\rangle$ was derived. This approach results in well-defined model coefficients of the k-ε model, which were empirically determined until now. The analysis of turbulent fields with various Reynolds numbers reveals the Reynolds number dependency of these coefficients. For high Reynold number they approach the empirical values (cf. table 1).

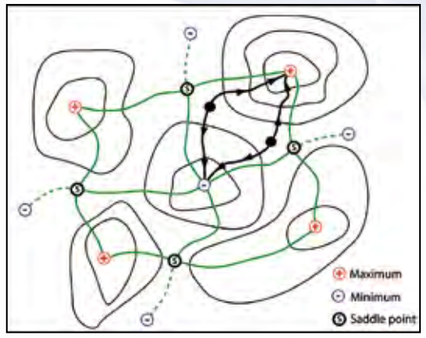

Two-point Velocity Difference along Gradient Trajectories

Conventionally, the two-point velocity difference is considered in the Cartesian coordinate system. Based on Kolmogorov's scaling laws it is expected that the ensemble average $\langle\Delta u\rangle \propto r^{1/3}$, where Δu is the velocity difference in longitudinal direction and r is the separation distance of two points. A corresponding analysis has been performed in two-point statistics along the scalar gradient trajectories of a passive scalar field [9,10]. Strictly derived from the governing equation of the scalar gradient, and based on an observation that at large separation distances the two-point product of the scalar gradient is decorrelated from the velocity difference, it can then be shown that along gradient trajectories of the passive scalar the two-point velocity difference is proportional to the separating arclength l in between, i.e. $\langle\Delta u_n\rangle \propto l$ ($\Delta u_n = u \cdot n$, and therefore the velocity projected on the gradient trajectory). Figure 4 shows on the left the linear increase of the average $\langle\Delta u_n\rangle$ for large separation distances l from

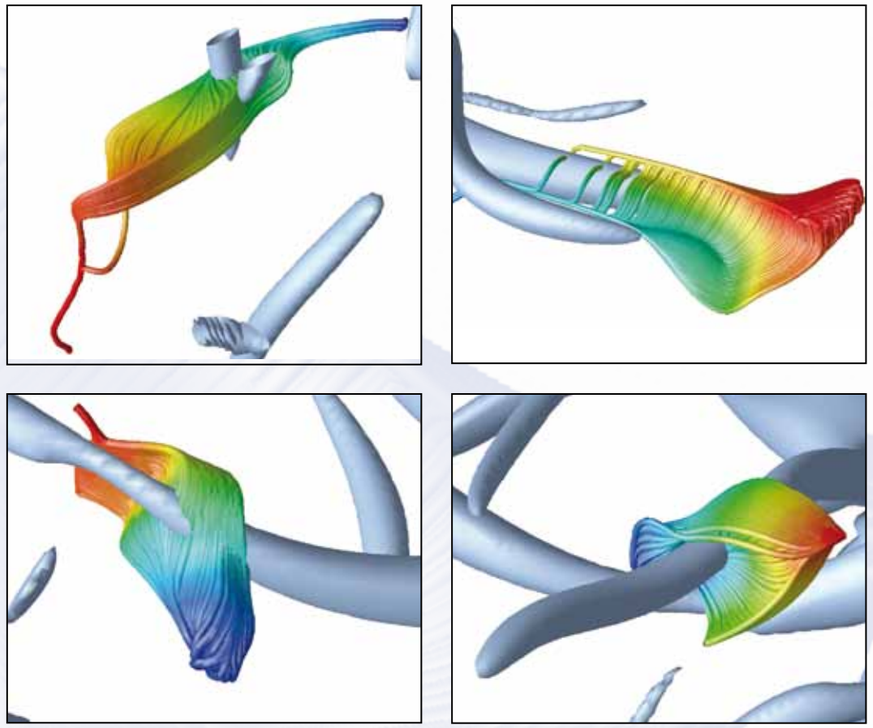

Figure 3:
Left: Two gradient trajectories connecting the same pair of extremal points and therefore belonging to the same dissipation element.
Right: Four dissipation elements interacting with vortex tubes.

different DNS data non-dimensionalized by integral time τ and Taylor length λ. On the right side of figure 4 the same data are shown, but non-dimensionalized by the characteristic strain rate of the corresponding dissipation element a_∞ [11].

Physically this surprising outcome can be explained as follows. By conditioning the statistics on gradient trajectories, regions of large extensive strain are preferentially extracted, because the strain has acted on the scalar fields in a way to allow gradient trajectories to proceed over large distances. On average, the strain is extensive $\langle \Delta u_n \rangle > 0$, except for very short elements where it becomes compressive $\langle \Delta u_n \rangle < 0$.

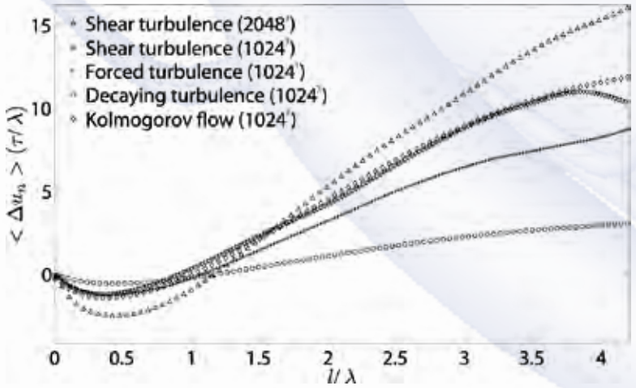

Figure 4: First order velocity structure function $\langle \Delta u_n \rangle$ along a gradient trajectories non-dimensionalized by τ/λ and $a_\infty\lambda$.

In the limit $l \to 0$ this is in agreement with the finding that the scalar gradient is aligned with compressive strain [12].

Direct Numerical Simulations

Direct Numerical Simulation (DNS) is of indispensable importance in the study of turbulence. Compared with experiments, DNS has the advantage of providing a complete set of three-dimensional data with high spatial resolution.Incompressible flows are described by the Navier-Stokes equations: continuity

$$\nabla \cdot U = 0, \qquad (1)$$

momentum

$$\frac{\partial U}{\partial t} + U \cdot \nabla U = -\frac{1}{\rho}\nabla p + \nu \nabla^2 U, \qquad (2)$$

where U, ρ, p and ν denote the velocity field, the fluid density, the pressure field and the kinematic viscosity, respectively. By non-dimensionalizing the Navier-Stokes equations with a reference velocity U_0 and a reference length L, one obtains the Reynolds number Re_L, defined as

$$Re_L = \frac{U_0 \cdot L}{\nu}, \qquad (3)$$

as the only remaining parameter in eq. 2. Thus the Reynolds number Re_L, which indicates the ratio of the inertial to the viscous forces acting on a large scale L, is the most fundamental nondimensional parameter to characterize flows. There exists a critical value of the Reynolds number, above which the flow will transit from laminar to turbulent.

Turbulence is a multiscale phenomenon. The integral scale L describes the size of the largest structures while the Kolmogorov scale η is of the order of the size of the smallest structures of the flow. To fully resolve all properties of a turbulent flow, η must be of the order of the grid size Δx in the calculation domain. The ratio $L/\eta \propto Re_L^{3/4}$ and therefore the total number of grid points needed to resolve the smallest scales in three dimensions is $N^3 \propto Re_L^{9/4}$.

Generally it is believed that the spatial resolution for DNS is about $\Delta x \propto O(\eta)$. Practically, depending on different calculations, this criterion can vary. To study the overall flow parameters, like the mean kinetic energy and energy dissipation it is reasonable to choose $\Delta x \sim 2\eta$ and therefore the grid resolution. This resolution is adequate to obtain first and second-order statistics. But to investigate derivatives or higher order statistics, the above criterion is too loose. More strictly, Sreenivasan [13] pointed out that due to intermittency, scales which are much smaller than η can locally occur in turbulence. To also resolve these local small scales, or spots of violent events, the grid size must be determined in a more stringent way.

To analyze the structures of dissipation elements, the resolution has to conform to this criterion. As shown in figure 3, the geometrical structure of dissipation elements is complicated in 3D space. The number of dissipation elements in a flow field is determined by the number of extremal points. It has been found that extremal points will cluster together with a distance of the order of η ([14]). To resolve extremal points in clusters in a sufficient way, Δx must be smaller than η.

Computational Challenge

With this resolution criterion in mind, a highly resolved DNS for high Reynolds number with a smallest possible numerical error is imperative to study the geometry of turbulence using the gradient trajectory analysis. The required computational power is only accessible on a supercomputer like JUGENE.

Therefore, a new code was developed to meet the requirements of todays supercomputers not compromising in precision and maximum grid size. This led to the pseudo-spectral code "psDNS" solving the incompressible Navier-Stokes equations in a cubic box of size $2\pi^3$ with periodic boundary

DNS case	1	2	3	4	5
Flow type	Shear	Shear	Forced	Kolmogorov	Decaying
No. of grid cells	2048^3	1024^3	1024^3	1024^3	1024^3
Reynolds number Re_λ	295	139	144	188	71
viscosity ν	$9 \cdot 10^{-4}$	$2 \cdot 10^{-3}$	$2.8 \cdot 10^{-3}$	$2.5 \cdot 10^{-4}$	$5 \cdot 10^{-4}$
kinetic energy k	3.510	1.925	3.210	0.115	0.049
dissipation ε	1.160	0.640	1.190	0.010	0.001
integral time $\tau = (k/\varepsilon)$	3.026	3.008	2.697	11.50	49.00
Kolmogorov scale η	0.005	0.011	0.017	0.006	0.010
Taylor length λ	0.165	0.245	0.275	0.170	0.135
resolution $\Delta x/\eta$	0.610	0.558	0.361	0.970	0.610

Table 2: Parameters of the different DNS cases.

conditions. The kernel of the code is a modified version of the efficient MPI-parallelized 2D decompositioning (figure 1) and 3D Fast Fourier Transformation (FFT) library p3dfft [15] to employ FFTs for spatial discretization. Simulations were run on the JUGENE Supercomputer of the Research Center Jülich scaling up to 16,384 CPUs for the 2048^3 case.

DNS of five different types of turbulent flows (table 2), namely of homogeneous shear turbulence (case 1 and 2), homogeneous isotropic forced turbulence (case 3), homogeneous isotropic decaying turbulence (case 4), as well as a Kolmogorov flow (case 5) were performed. In detail, homogeneous shear turbulence has been performed with a mean velocity gradient on 2048^3 and 1024^3 grid points. The temporal advancement is performed by a third-order Runge-Kutta method. The convective term of the Navier-Stokes equations is formulated in the skew-symmetric form in order to reduce aliasing errors and to improve numerical stability and accuracy. Periodic boundary conditions, which are required for the application of spectral methods, cannot be satisfied when a mean gradient is present. To overcome this problem a coordinate transformation of all dependent variables to a moving frame attached to the mean flow is performed. Since the computational frame gets distorted with advancing time a remeshing procedure is applied to keep the distortion in an appropriate range. For case 3-5 the non-linear terms are solved using an explicit Adams-Bashforth method, while the linear terms are solved by an implicit Crank-Nicholson method. Aliasing errors were removed by isotropic truncation applying the 2/3 rule.

Conclusion

For simulation of turbulent flows in science and engineering, model equations for small scale structures are essential. Dissipation element analysis is a technique to improve these model equations, revealing new aspects of turbulent fields, cf. section 5. For instance, using dissipation element analysis a successful contribution has been applied to improve the standard k-ε model as described in section 4.

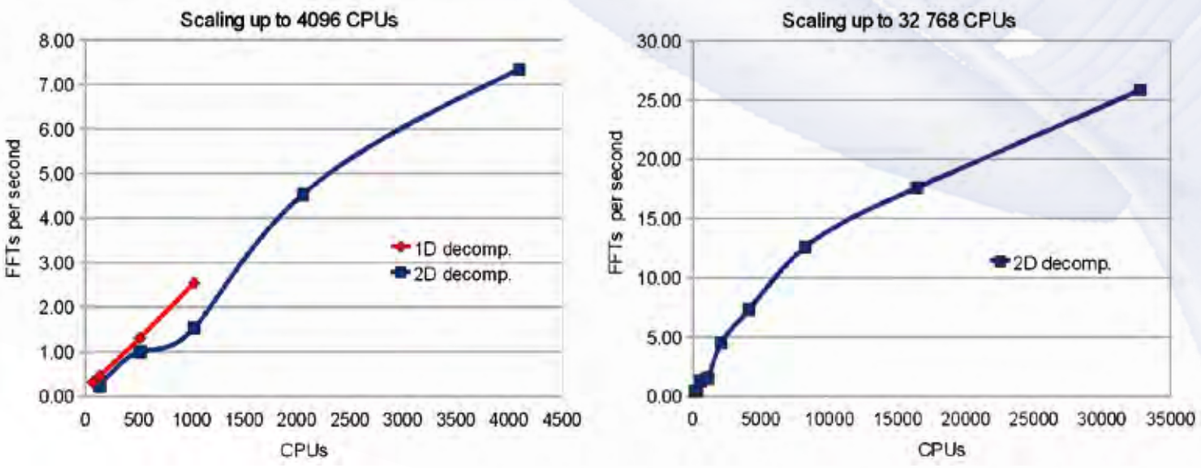

Figure 5: 3D Fast Fourier Transformations per second for a 1024^3 grid. The code takes ~30% of the computational time for 3D FFTs. In the remaining ~70% psDNS scales linear because of no other MPI communication in the iteration loop.

Acknowledgements

This work is founded by Deutsche Forschungsgemeinschaft under Grant Pe 241/38-1. The Jülich Supercomputing Centre, Germany, is appreciated for accessing the Blue Gene supercomputer for the calculations.

References

[1] **Miyauchi, T., Tanahashi, M.**
Coherent fine scale structure in turbulence, IUTAM Sym. on Geometry and Statistics of Turbulence (ed. T. Kambe et al.), pp. 67-75, Kluwer, 2001

[2] **Perry, A. E., Chong, M. S.**
A description of eddying motions and flow patterns using critical-point concepts, Annu. Rev. Fluid Mech., 19:125-155, 1987

[3] **Davila, J., Vassilicos, J. C.**
Richardson's Pair Diffusion on the Stagnation Point Structure of Turbulence, Phys. Rev. Lett. 91, 14:144501, 2003

[4] **Wray, A. A., Hunt, J. C. R.**
Algorithms for classification of turbulent structures, in H. K. Moffat and A. Tsinober, editors, Topological Fluid Mechanics, pp. 95-104, Cambridge University Press, 1990

[5] **Corrsin, S.**
Random geometric problems suggested by turbulence, in M. Rosenblatt and C. van Atta, editors, Statistical Models and Turbulence, volume 12 of Lecture Notes in Physics, pp. 300-316, Springer Verlag, 1971

[6] **Peters, N., Wang, L.**
Dissipation element analysis of scalar fields in turbulence, C. R. Mechanique, 334:493-506, 2006

[7] **Peters, N.**
Multiscale Combustion and Turbulence, 32nd Symposium on Combustion, Montreal 2008, Proc. Combust. Inst., 32:1-25, 2009

[8] **Schaefer, P., Gampert, M., Goebbert, J. H., Wang, L., Peters, N.**
Testing of different model equations for the mean dissipation using Kolmogorov flows, Flow, Turbulence and Combustion, 85:225-243, 2010

[9] **Wang, L.**
Scaling of the two-point velocity difference along scalar gradient trajectories in fluid turbulence, Phys. Rev., E 79, 046325:046325, 2009

[10] **Wang, L., Peters, N.**
Mean velocity increment conditioned on gradient trajectories of various scalar variables in turbulence, Phys. Scr. T, 142:014004, 2010

[11] **Gampert, M., Goebbert, J. H., Schaefer, P., Gauding, M., Peters, N., Aldudak, F., Oberlack, M.**
Extensive strain along gradient trajectories in the turbulent kinetic energy field, New Journal of Physics, 2011

[12] **Ashurst, Wm. T., Kerstein, A. R., Kerr, R. M., Gibson, C. H.**
Alignment of Vorticity and Scalar Gradient with Strain Rate in Simulated Navier-Stokes Turbulence, Phys. Fluids, 30:2343-2353, 1987

[13] **Sreenivasan, K. R., Schuhmacher, J., Yakhot, V.**
Intermittency and direct numerical simulations, 58th Annual Meeting of the Division of Flow Dynamics, APS, Chicago, USA, 2005

[14] **Schaefer, P., Gampert, M., Gauding, M., Peters, N., Trevino, C.**
The secondary splitting of zero-gradient points in a scalar field, Journal of Engineering Mathematics, DOI 10.1007/s10665-011-9452-x, 2011

[15] **Pekurovsky, D.**
Parallel Three-Dimensional Fast Fourier Transforms (p3dfft, v. 2.3), San Diego Supercomputer Center, 2008

• Jens Henrik Göbbert[1]
• Michael Gauding[1]
• Markus Gampert[1]
• Philip Schäfer[1]
• Norbert Peters[1]
• Lipo Wang[2]

[1] Institut für Technische Verbrennung, RWTH Aachen University

[2] University of Michigan-Shanghai Jiao Tong University Joint Institute

Spring 2011 • Vol. 9 No. 1 • inSiDE

Applications

Applications of Multiferroics: from elementary Particles to new Memory Devices

What are Multiferroics?

Multiferroic materials [1] are in a broader sense defined as materials possessing two or more ferroic orders (ferromagnetic, ferroelectric, ferroelastic and ferrotoroidic). In practice, most research is dedicated to the multiferroics which combine ferroelectricity and ferromagnetism, due to their high potential in spintronics applications. In particular, current driven through a magnetic tunnel junction (MTJ) with a multiferroic tunnel barrier could be controlled via the magnetic and/or electric polarization of the barrier. Junctions with a multiferroic tunnel barrier are usually referred to as multiferroic tunnel junctions (MFTJs). Alternatively, MFTJs can be realized by combining a ferroelectric barrier with ferromagnetic leads.

The use of ferromagnetic materials as leads and a multiferroic material as a tunnel barrier in MFTJs would lead to 8 possible resistive states of such junctions (see Fig. 1). Such theoretically conceivable 4-bit memory is in practice still hard to achieve. One of the main problems is the scarcity of ferroelectric materials which possess, at the same time, a net magnetic moment. Most of the known multiferroic materials exhibit either weak ferromagnetism with low ordering temperatures (far below room temperature) or an antiferromagnetic order.

The reason for the rare coexistence of ferroelectric and ferromagnetic order can be found by analyzing the electronic structure of the typical ferroic materials. The most studied ferroelectrics, such as

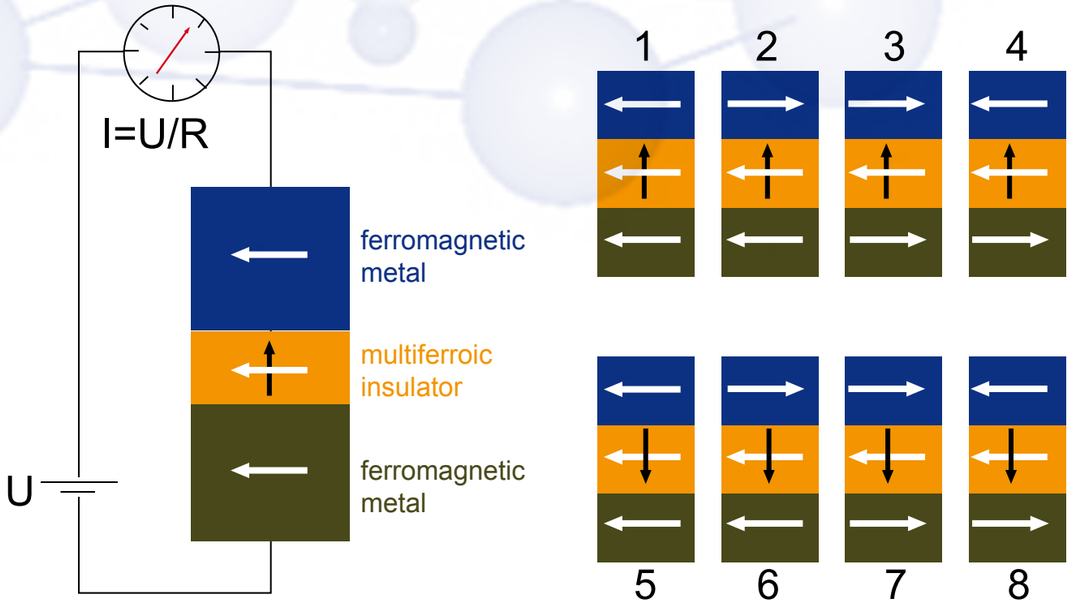

Figure 1: Multiferroic tunneljunction (left) and eight resistive states that it provides (right). The ferroelectric polarization is depicted with the black arrow; the white arrow stands for the magnetization.

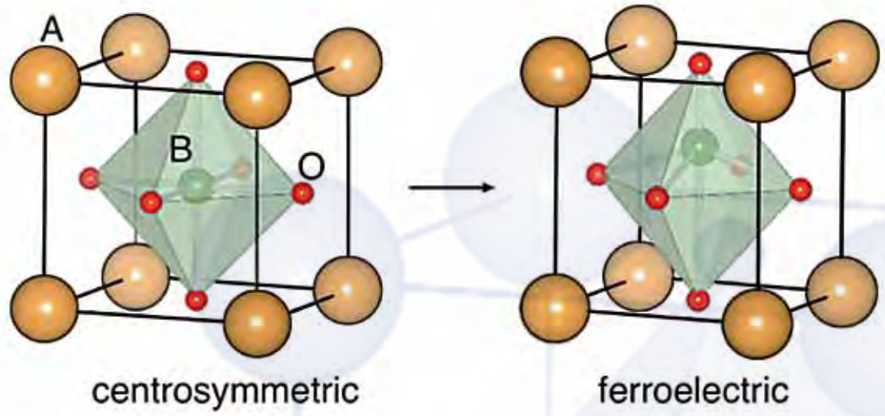

Figure 2: The structure of a perovskite, ABO_3, in the centrosymmetric (left) and the ferroelectric phase (right).

$BaTiO_3$, have a perovskite structure, ABO_3 (see Fig. 2) in which the spontaneous polarization is achieved by the off-center distortion of the small transition metal cation (B) at the center of the octahedron of the oxygen atoms. The existence of such distortion is determined by a balance between the short-ranged repulsions of adjacent electronic clouds and the bonding taking place in the compound; while the former favours the unpolarized, symmetric structure, the latter can induce ferroelectricity via the p-d hybridization where the filled oxygen 2p states donate electrons into the empty d states of the transition metal cation (B) as it moves off-center. Clearly, the transition metal B should have vacant d states for ferroelectricity to occur in the described way. On the other hand, the driving force of magnetism lies exactly in the partly filled transition metal d states. The two conditions oppose each other and in order to induce both ferroic orders in the same material, a "side-way" has to be taken. There are several possible scenarios that could take place [2]. The first one that comes to mind in a multiferroic material design is to simply combine the atoms that induce ferroelectricity and the ones that induce magnetism in one compound. In a perovskite ABO_3 structure that would in practice mean either having a ferroelectricity-inducing cation at the A-site and the magnetic one at the B-site or vice versa.

In the following, we describe a successful computational design of multiferroics following either of these two routes, and their possible applications. In both studies we are using the Density Functional Theory (DFT) and determining the structural and electronic properties of the investigated materials with the use of first-principles calculations, relying on the laws of quantum mechanics.

All the calculations were made on the Jülich JuRoPa cluster.

Room-temperature Multiferroics for Data Storage

While a considerable effort of many research groups aims at designing and synthesizing new materials that would be ferroelectric and would possess a net magnetization at the same time, the main obstacle in the applications of these materials remains their low magnetic ordering temperature. The most-studied room-temperature multiferroic, $BiFeO_3$, is an antiferromagnet, i.e. its net magnetization is zero.

Ferromagnetic oxides in general are scarce, even without imposing the additional request of ferroelectricity.

A recent work, however, reported on high-temperature ferrimagnetism in a double perovskite, Sr_2CrOsO_6 [3]. The compound possesses a net magnetization up to the temperature of about 720 K (but no ferroelectric polarization). The key to such a high magnetic ordering temperature lies in the combination of one element from the 3d-series of the periodic table and one from the 5d-series, at the B and B' lattice sites of the double perovskite structure with the formula $A_2BB'O_6$ (an ordered double perovskite structure resembles the perovskite one, ABO_3, with two kinds of atoms occupying the B site of the lattice in a three-dimensional checkerboard manner). The idea of placing one 3d and one 5d element at the B and B' lattice sites to achieve a high ordering temperature can be further extended: we could also place a cation that induces ferroelectricity at the A-lattice site in hope to obtain a room-temperature multiferroic. One such cation is Bi^{3+}. The presence of this ion in a crystal lattice often leads to polar lattice distortions, due to its "lone pair" of electrons. For example, in the previously mentioned multiferroic, $BiFeO^3$, the $6s^2$ lone pair on Bi stabilizes the ferroelectricity in the compound. Bi cations move in the direction opposite to oxygen-iron cage. This shift is induced by the Bi-O covalence, in contrast to the usual perovskite ferroelectrics where it is the covalence of O with the B-site cation that induces ferroelectricity.

We therefore conducted a computational investigation of the ordered double perovskites $A_2BB'O_6$ with a Bi^{3+} cation at the A-site and several combinations of 3d and 5d cations at the B and B' sites [4]. Among the screened compounds, we found two promising candidates for room-temperature multiferroicity, namely Bi_2MnReO_6 and Bi_2NiReO_6. Their investigation involved, in a first step, a full structural relaxation, in 12 symmetry allowed centrosymmetric structures. Next, polar distortions were allowed and the polar structures newly relaxed. From the comparison of the total energies of all the relaxed structures we identified the ground-state structures of both compounds and continued with the investigation of their magnetic properties.

We found that while in bulk both compounds prefer a centrosymmetric structure, when grown in a thin-film form they both become ferroelectric, with a separation of the positive and negative charges along the diagonal of the (pseudo) cubic lattice cell (i.e. the ferroelectric polarization, P, points along this diagonal). They are also both ferrimagnetic and thus present a net magnetic moment, which is the desirable characteristic for their potential applications. We investigated further their magnetic ordering temperature by extracting the magnetic interactions from the first-principles calculations and using them in a Monte-Carlo calculation. The calculated magnetic ordering temperature for both compounds is significantly higher that the room temperature (330 K for Bi_2MnReO_6 and 360 K for Bi_2NiReO_6). Therefore, if synthesized, these compounds could be used for magnetic data storage.

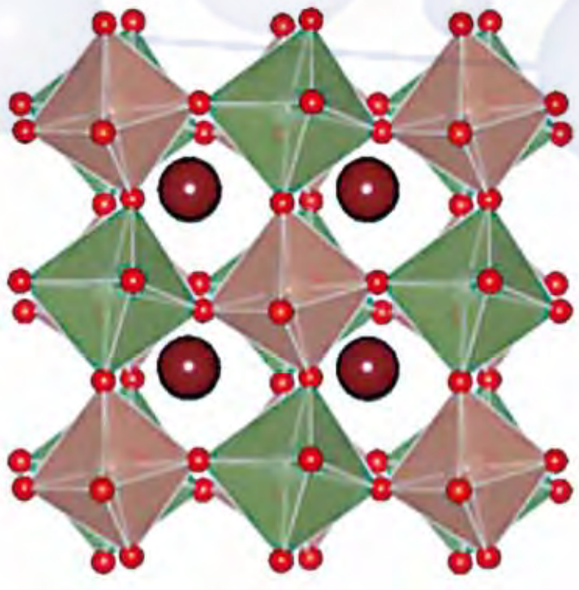

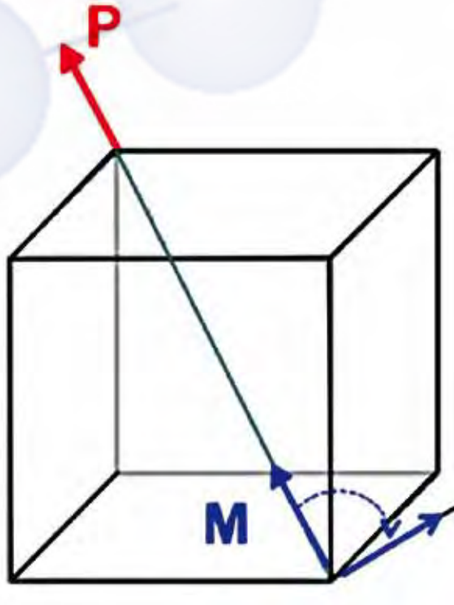

Figure 3: Left: The ferroelectric R3 structure of the double perovskites Bi_2MnReO_6 and Bi_2NiReO_6. Bi^{3+} cations are depicted as the large (brown) spheres and O^{2-} as the small (red) spheres. Re^{4+} and Mn^{2+} (Ni^{2+}) are alternating in the shaded octahedra, in a three-dimensional checkerboard manner (indicated by the color of the octahedra). Right: The ferroelectric polarization (P) and the magnetization (M) of the material are aligned, along the diagonal of the (pseudo) cubic unit cell. The energy needed to rotate the magnetization by 90o is large, of the order of meV/formula unit, making the control of the magnetization direction by an external electric field possible.

While the mere coexistence of the two ferroic orders, magnetism and ferroelectricity, in a single phase of these compounds is already valuable for their potential applications, there is a further aspect one should investigate. Namely, magnetic writing is cumbersome: it is difficult to get focused strong magnetic fields needed for this operation. It would be far easier and more energy-efficient if the information could be encoded into the magnetic medium using an electric field. In a multiferroic, the idea would be to use the magneto-electric coupling for this purpose. However, as in most multiferroics (as well as in the two compounds we investigated) the ferroelectricity and magnetism originate in different constituents of the compound, the coupling of these two ferroic orders is usually insufficiently strong. One should keep in mind though that the ferroelectric polarization deforms the crystal lattice in a specific way, and this deformation changes if an electric field is applied and the direction of P is switched. On the other hand, the direction that the magnetization will assume depends on the shape of the unit cell of the material. Therefore, an external electric field which controls directly the direction of P could indirectly, via the lattice deformation, control also the direction of the magnetization of the material. A necessary condition for such a mechanism would be that the magnetization direction is not too loosely coupled to the lattice. This coupling is a relativistic phenomenon and is here expressed in terms of energy which we have to pay in order to rotate the magnetization from its preferred direction (the so-called "easy axis") to the plane perpendicular to it. We calculated that the magnetic easy axis in the two investigated double perovskite compounds

lies along the direction of the ferroelectric polarization (Fig. 3) and the energy required to rotate the magnetization to the perpendicular plane is very high, approximately 5 meV, which is of the same order of magnitude as in the currently used magnetic recording media. These calculated characteristics (multiferroicity at room temperature, electrically controllable magnetization direction) put the two computationally "designed" compounds right to the top of the list of materials with a high applications potential.

Although an experimental fabrication of computationally designed compounds usually takes a long time and depends on the optimization of many growth parameters, we hope that we will soon receive an experimental confirmation of the exciting predicted properties of these materials.

Multiferroics and elementary Particles: Search for a permanent electric Dipole Moment of the Electron

As previously mentioned, the main obstacle in the applications of multiferroics so far has been their low magnetic ordering temperature. Recently, we investigated the possibility of employing a multiferroic material in a solid-state based search for the permanent electric dipole moment of the electron [5], where the low magnetic ordering temperature was actually a virtue, a necessary condition for application of the material in the experiment designed and conducted in the group of Prof. S. Lamoreaux at the Yale University. We designed computationally a multiferroic compound, (Eu,Ba)TiO^3, according to the experimental needs. The compound was subsequently synthesized and characterized in the groups of Prof. S. Kamba in Prague and Prof. S. Lamoreaux at Yale. The characterization revealed that it had all the experimentally desired properties and the first measurements of the electric dipole moment of the electron are now underway.

An electron has three intrinsic characteristics: its mass, charge and the magnetic dipole moment (spin). If an electric dipole moment of the electron exists, it would be its fourth intrinsic characteristic and it would, for reasons of symmetry, be aligned with its spin.

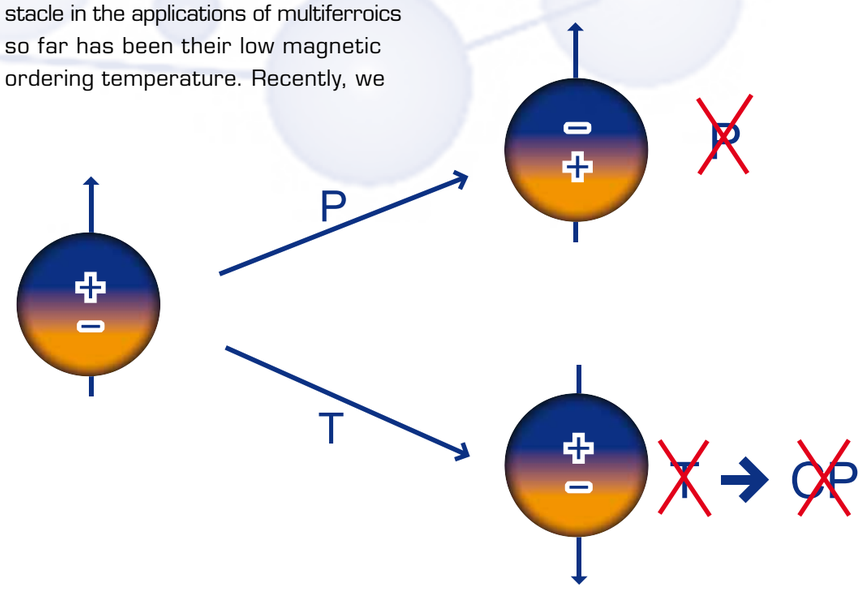

Figure 4: An electron with a magnetic and an electric dipole moment violates the parity symmetry (P) as well as the time reversal symmetry (T); the latter means that the combined symmetry of charge conjugation and parity, CP, is also violated.

The existence of the electronic electric dipole moment (EDM) would, furthermore, violate the discrete symmetries of parity and time reversal. The reason for this can be seen in Fig. 4, where the spin is depicted with an arrow and the electric dipole with the "+" and "-" (the picture of the spatially separated positive and negative charge should be taken only as a way to imagine the electronic electric dipole moment, not as a reality as the electron is a dimensionless particle). The parity symmetry (P) reverses the direction of the electric dipole, but as the spin is an axial vector its direction does not change. Therefore, after the action of parity we end up with a state we did not start from, i.e. the parity symmetry is violated. Similar is valid for the time reversal, T, which flips the direction of spin, but leaves the electric dipole moment intact. Thus, T is also violated. According to the so-called CPT theorem, all the processes in nature should be invariant under a combined action of the three discrete symmetries, the charge conjugation C, the parity P and the time reversal T. Therefore, if the EDM violates the time reversal, it also violates the combined CP symmetry. The CP-symmetry violation is a necessary condition for explaining the observed matter-antimatter disbalance in the Universe.

What are the properties a material should have in order to serve as a "guinea pig" in an experiment searching for the EDM? Firstly, the electrons are situated in atoms and if they possess an EDM, this means that the atom will also have an EDM, if the electronic EDMs are mutually aligned. As the EDM of an electron is aligned with its spin, this means that the spins of the electrons in the atom should also be aligned, i.e. the atom should be magnetic. Furthermore, the atomic EDM is proportional to the electronic one, through the so-called enhancement factor, which is in its turn proportional to the cube of the atomic number. We therefore need a compound containing heavy magnetic atoms, e.g. one of the rare-earth atoms. The idea of the EDM measurement is depicted in Fig. 5: if an external electric field E is applied, the EDMs of most of the magnetic atoms in the solid align with it. This means that also their magnetic moments align with the E-field, giving the whole material a net magnetization. When the electric field direction is flipped, the EDMs, and the magnetization also change sign. This change in the magnetization can be experimentally measured. It is important to note that any linear magneto-electric coupling in the material would mimic the EDM

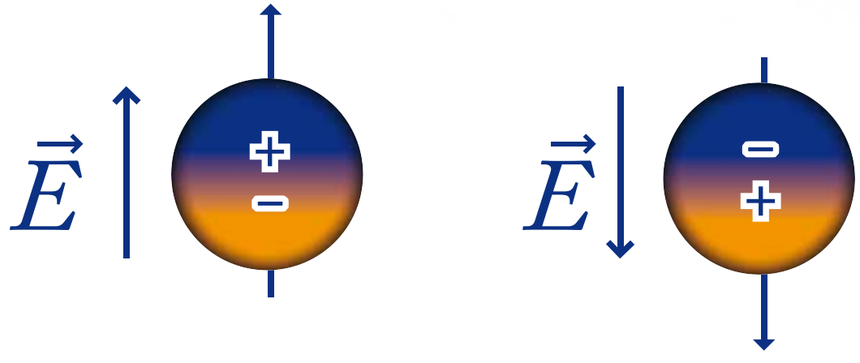

Figure 5: The idea of the experiment searching for the permanent electric dipole moment of the electron: When an external electric field E is applied, the electric dipole moment aligns with it, aligning in this way also the magnetic dipole moment with the E-field. When the direction of the field is reversed, so is the direction of the magnetic dipole moment. In a solid, this induces a tiny, but measurable change in the net magnetization of the material.

signal. In order to avoid this, we use a paramagnet where the linear magnetoelectric effect is forbidden by symmetry. The induced net magnetization is proportional to the density of the magnetic atoms in the solid and the strength of the applied electric field, and inversely proportional to the temperature. For this reason the experiments are conducted at the liquid Helium temperature (4.2 K) and, in order to enhance maximally the externally applied electric field, we use a ferroelectric. We therefore need a material where the ferroelectric polarization is still switchable at the cryogenic temperatures and which is, at these temperatures, paramagnetic but still has a sufficiently high density of the magnetic ions. All these conditions are tough to meet in a single material.

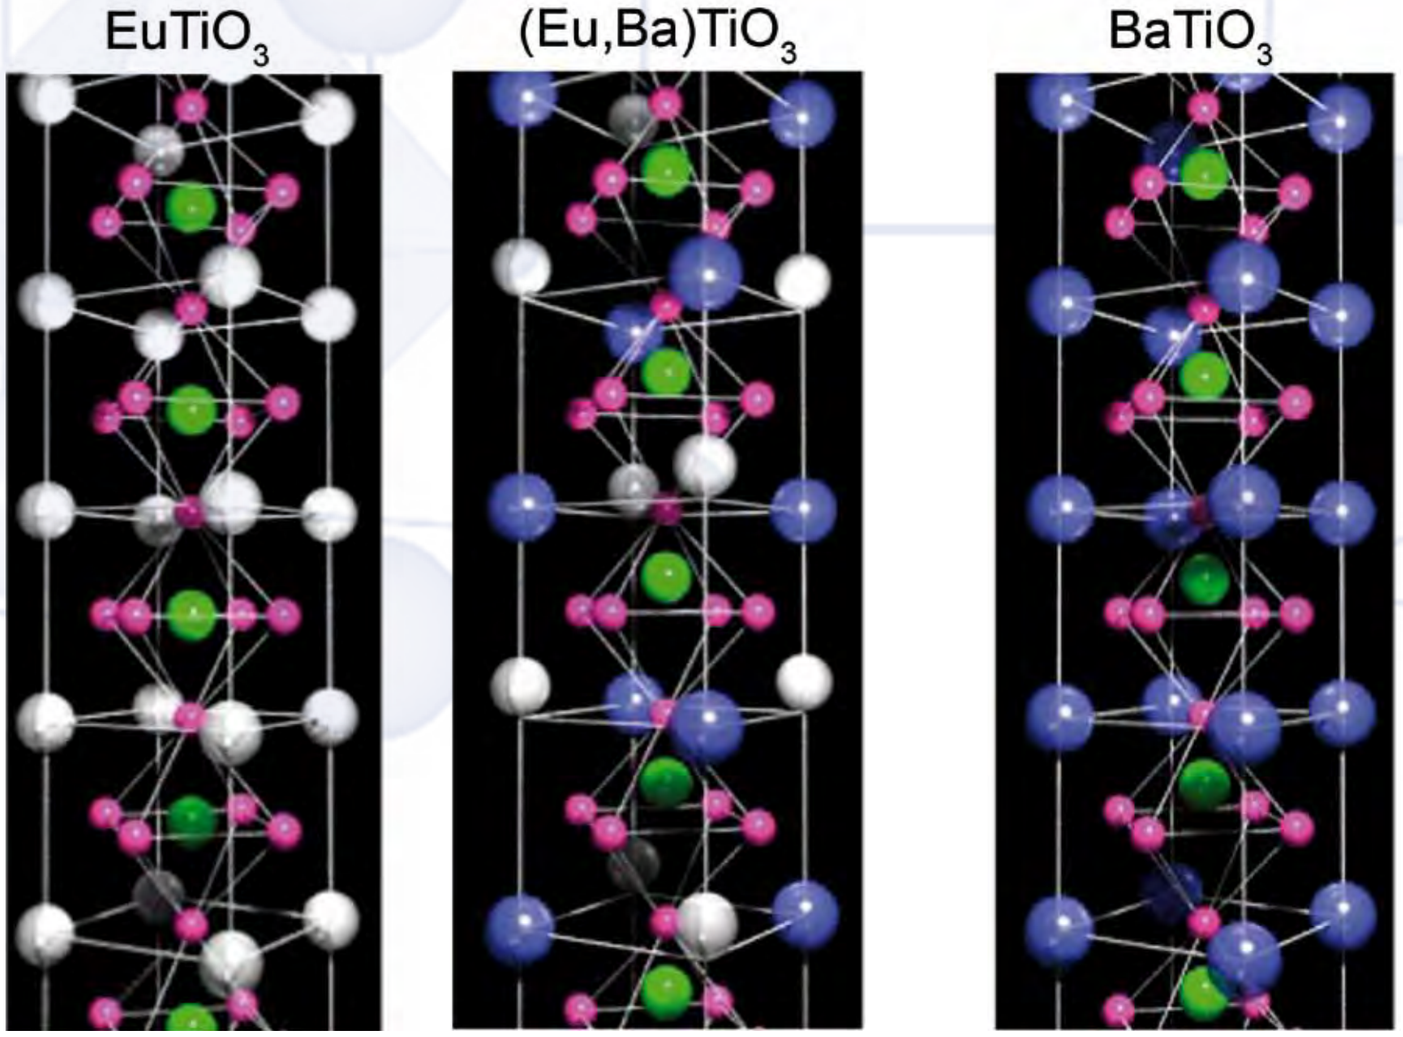

Figure 6: $EuTiO_3$ (left) is a paraelectric antiferromagnet with an ordering temperature of 7 K. $BaTiO_3$ (right) is a non-magnetic ferroelectric. A combination of the two materials, $(Eu,Ba)TiO_3$ (middle) is a multiferroic, with the magnetic ordering temperature of 1.9 K, and the permanent electric dipole moment of the electron. Pink spheres represent oxygen, the green ones titanium, the blue ones barium and the white ones europium atoms.

We start our material design by taking notice of two materials, $EuTiO_3$ and $BaTiO_3$. The former is paraelectric and antiferromagnetic, with magnetic ordering temperature of 7 K - low, but not lower than 4.2 K, which is what we need for the experiment. The latter is a non-magnetic ferroelectric. It is known that some paraelectrics become ferroelectric under the influence of epitaxial or chemical strain (doping by larger cations that enlarge the lattice parameter of the compound). As Ba^{2+} is larger than Eu^{2+}, we make a hypothesis that by alloying $EuTiO_3$ at the Eu site by 50% of Ba atoms, this will result in a ferroelectric $(Eu,Ba)TiO_3$ compound (Fig. 6). Moreover, the alloying should dilute the magnetic Eu sublattice, lowering in this way the magnetic ordering temperature. We investigated computationally the properties of $(Eu,Ba)TiO_3$ (the structural instabilities, electronic structure and the magnetic ordering), concluding that the material should be an ideal candidate for the planned experiment.

Finally, we calculated the total energy the system has in several patterns of magnetic ordering, in order to estimate the likelihood of a magnetic order: the obtained energy differences were tiny (~1 meV per 40-atom supercell), indicating that the magnetic ordering should not set in down to very low temperatures. The compound was subsequently synthesized and characterized by our collaborators; the measured properties confirmed our predictions – the compound is indeed a multiferroic, with a magnetic ordering temperature of 1.9 K. The first measurements searching for the EDM are currently conducted on this material.

Computational Demands

Within the density-functional theory codes that we are using, the Hamiltonian matrix of the system is set up and solved on a grid in the reciprocal space, yielding the eigenstates and the corresponding eigenvectors. The Hamiltonian is diagonalized for each

one of the k points, which sample the Brillouin zone of the material (essentially, the k points are the wave-vectors characterizing the motion of all the electrons in a solid). The size of the matrix depends on the number of basis functions used for the expansion of the potential/charge densities and the wave functions. All the calculations should be reasonably converged with respect to the number of the basis functions and the number of k points. As the energies that characterize the physical phenomena in solids can have very different scales, the necessary precision of a calculation (i.e. the total computational load in terms of time and memory consumption) can be very different from one problem to another. For example, the description of the relativistic phenomena that determine the direction of the magnetization in a solid (such as in the first case described above) are very expensive in terms of the CPU-time, as they rely on tiny differences in energies that strongly depend on the k-point sampling. The calculation of phonon spectrum that is needed in order to evaluate the structural stability of a complex oxide compound is equally demanding, due to the high accuracy needed to evaluate the forces acting on all the atoms contained in the large unit cells of low symmetry (such as in the second described example).

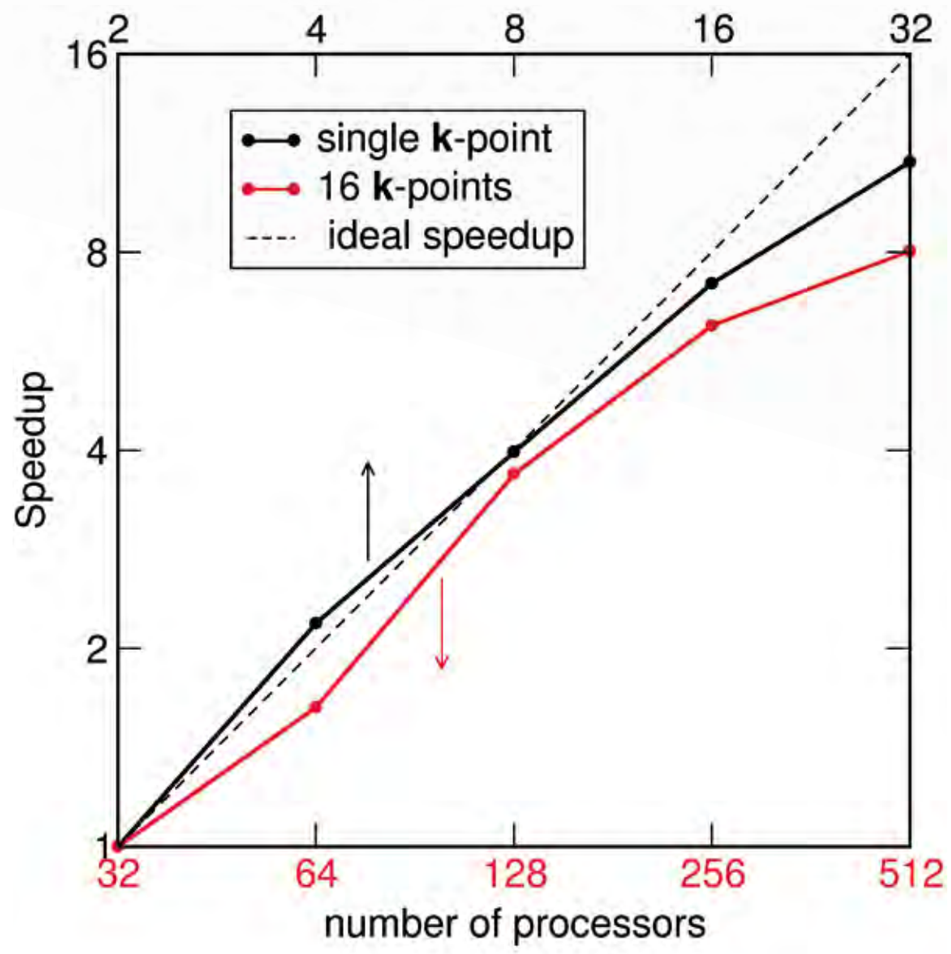

Figure 7: The scaling of the FLEUR code [7] on JuRoPA cluster, for a system with 82 atoms (matrix size 9500x9500). Dashed line indicates the ideal speedup; Black line: the system calculated with 1 *k*-point on 2-32 cores; Red line: the same system calculated with 16 *k*-points, using a combined *k*-points and eigenproblem parallelization on 32-512 cores (courtesy of Dr. G. Bihlmayer).

For an accurate and fast quantum-mechanical description of novel multifunctional oxides, the use of efficient parallelization algorithms on a supercomputer is therefore essential. As an example, in the phonon study of the material designed for the search for the permanent electric dipole moment of the electron, (Eu,Ba)TiO_3, we used supercells containing 270 atoms. Approximately 30 calculations were needed for determining the forces acting on atoms and testing their reliability, using typically up to 256 cores for 12 hours, and requiring up to 8 GBytes of memory on the root node. Additional calculations of the ground-state properties (structural relaxation, magnetic ordering, ferroelectric polarization) required ~10 further calculations with similar CPU-demands. In the determination of the ground-state structure of the double perovskites described here ~50 calculations using typically up to 256 cores for 12 hours and up to 5 GBytes of memory were performed. Our needs, that usually consist of a large number of calculations on several hundreds of cores, are best met by the Jülich JuRoPA cluster.

The two DFT codes used for our studies are the Vienna ab initio Simulation Package (VASP) [6] and the FLEUR code [7]. Both codes are efficiently parallelized. As an example, in Figure 7 the scaling of the FLEUR code on the JuRoPA cluster is shown for an 82-atoms unit cell. The parallelization of the eigenvalue problem is implemented, along with the parallelization of the k-points loop. Figure 7 shows that up to 256 cores the speedup is nearly ideal.

Acknowledgements

The support of the Young Investigators Group Programme of Helmholtz Association, contract VH-NG-409, as well as of the Jülich Supercomputing Centre is gratefully acknowledged.

References

[1] **Spaldin, N. A., Fiebig, M.,** The Renaissance of Magnetoelectric Multiferroics, Science 309, 301, 2005

[2] **Hill, N. A., Filippetti, A.** Why are there any magnetic ferroelectrics? J. Magn. Magn. Mater., 242, 976, 2002

[3] **Krockenberger, Y., Mogare, K., Reehuis, M., Tovar, M., Jansen, M., Vaitheeswaran, G., Kanchana, V., Bultmark, F., Delin, A., Wilhelm, F., Rogalev, A., Winkler, A., Alff, L.** Sr_2CrOsO_6: End point of a spin-polarized metal-insulator transition by 5d band filling, Phys. Rev., B 75, 020404R, 2007

[4] **Ležaić, M., Spaldin, N. A.,** High-temperature multiferroicity and strong magnetocrystalline anisotropy in 3d – 5d double perovskites, Phys. Rev., B 83, 024410, 2011

[5] **Rushchanskii, K. Z., Kamba, S., Goian, V., Vaněk, P., Savinov, M., Prokleška, J., Nuzhnyy, D., Knížek, K., Laufek, F., Eckel, S., Lamoreaux, S. K., Sushkov, A. O., Ležaić, M., Spaldin, N. A.** A multiferroic material to search for the permanent electric dipole moment of the electron, Nature Materials, 9, 649, 2010

[6] http://cms.mpi.univie.ac.at/vasp/

[7] http://www.flapw.de

• Mariana Ležaić

Research Centre Jülich

Applications

Nanoparticles as Emulsion Stabilizers

Particle-laden multiphase Flows

Assume an oil-water mixture where oil is the minority phase forming droplets inside the water. Adding small particles causes these particles to diffuse to the interface which is being stabilized due to a reduced surface energy. Such systems are also referred to as Ramsden-Pickering emulsions [1,2] with a strong potential for applications in the food and cosmetics industry, as well as for drug delivery, crude oil recovery, or waste water treatment [3]. Some of their most interesting properties are related to the complete blocking of Ostwald ripening, which is one of the main processes leading to drop coarsening in emulsions and foams, allowing long-term stabilization, and fabrication of new materials with complex hierarchical structure by using particle-stabilized drops. Additionally, particle-stabilized emulsions show interesting rheological properties, due to the irreversible particle adsorption and to the bridging of the surfaces of the neighbouring droplets by particle monolayers.

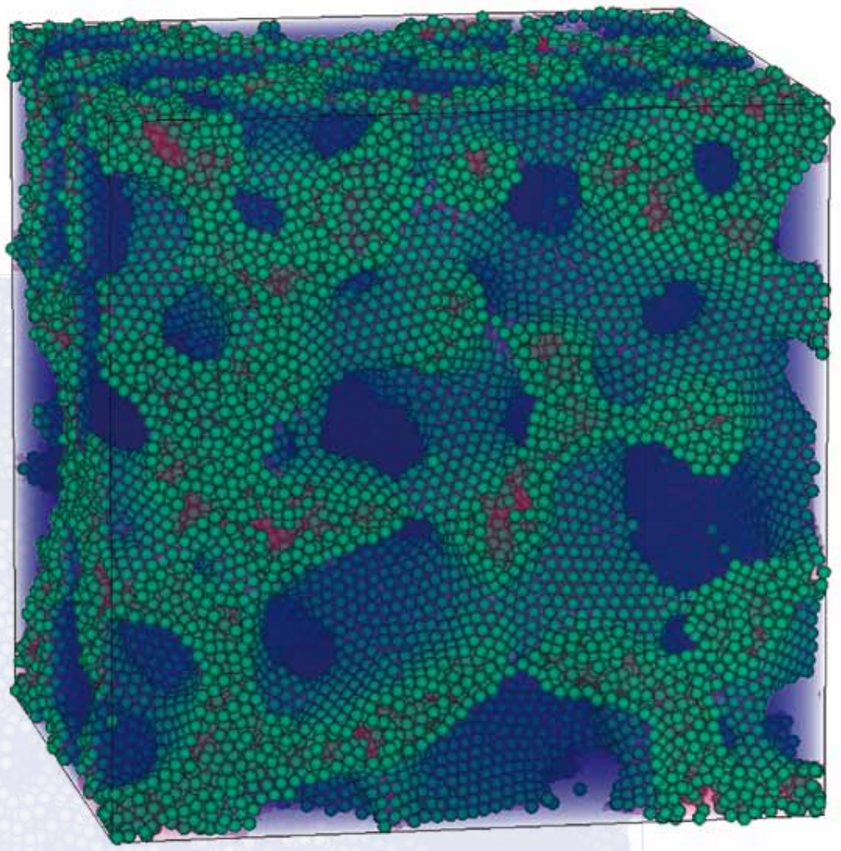

Figure 1: A snapshot from a 3D Bijel simulation: particles (green) self-assemble at the interface between two phase separating immiscible fluids (shaded in red and blue) and stabilize it.

Interestingly, many of the properties of such systems cannot be explained with concepts derived for surfactant stabilized systems due to the larger size of the particles and their missing amphiphilic properties. Theoretical models were developed and underlined by experiments emphasizing the role of specific features of the particle-stabilized dispersions which have no direct analogue in surfactant systems. These include the particle's contact angle, the strong interparticle capillary forces, or the pH value and electrolyte concentration of the solvents [3]. However, the quantitative description of these systems is far from satisfactory. Therefore, we developed a new simulation paradigm combining multicomponent lattice Boltzmann (LB) simulations to describe the solvents with a molecular dynamics solver to model the suspended colloidal particles [4]. The lattice Boltzmann method is a popular mesoscopic simulation method in fluid dynamics. A particular advantage of the LB approach is the availability of established multiphase or multicomponent methods and a straightforward implementation of complex boundary conditions. In recent months, the code was extended to also being able to study anisotropic particles such as ellipsoids in multiphase flow.

Our microscopic studies allow a direct link to macroscopic experimental data and will lead to a substantial improvement of our current understanding of particle-laden multiphase flows. Also we expect an impact on the current "trial and error" based experimental treatment of such systems in industrial processes. However, it has been confirmed during the last year that this is a computationally very challenging task because typical particle diameters need to be at least ten LB length units and typical droplet diameters are around ten times as large. In order to reduce finite size effects and to acquire a statistically relevant number of droplets, the side length of the volume of interest easily reaches 1,000-2,000 LB length units – touching the limit of what is possible on current supercomputers.

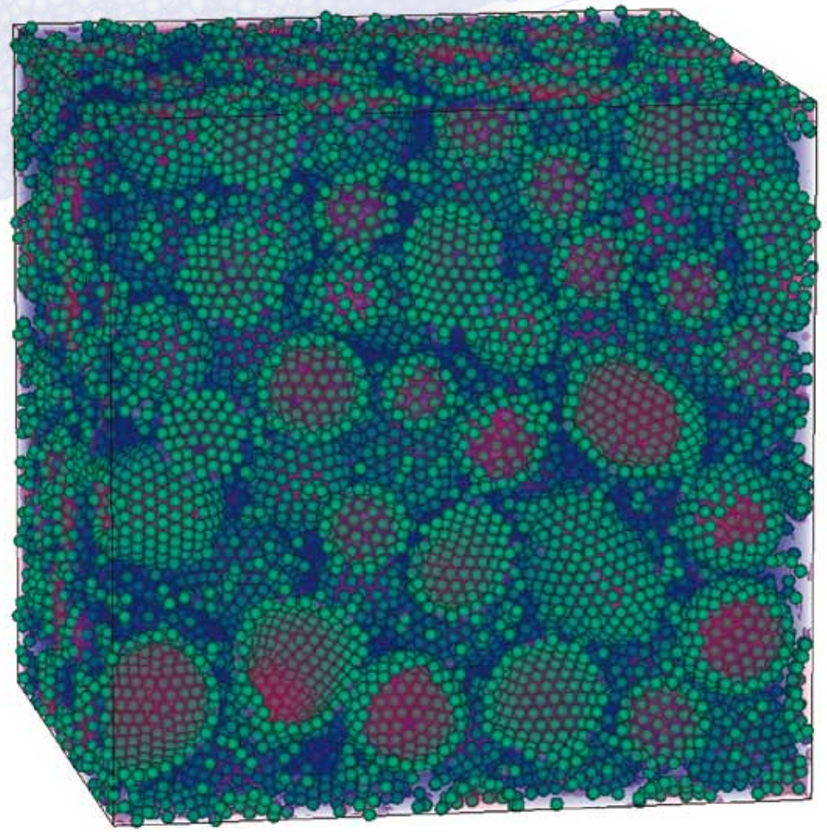

Figure 2:
A snapshot of a Pickering emulsion:
particles (green) self-assemble at the interface between
two immiscible fluids (shaded in red and blue) and stabilize droplets.

A second constraint is given by the very small time step in the simulation. The hydrodynamic interaction between the particles has to be resolved on much smaller scales than the particle motion pushing our time resolution down to the nanosecond scale. In order to still being able to reach steady states, typical simulation runs require some million or tens of million lattice Boltzmann steps. Typically 0.5 to 1 rack month on the Blue Gene/P system JUGENE in Jülich are required for a single simulation.

We studied in detail how particles travel towards the fluid-fluid interface and get jammed there generating fluid-bicontinuous gels (so-called "Bijels" [5]). Figure 1 shows a snapshot from a typical "Bijel" simulation.

Further, it was demonstrated that by modifying the lyophobic/lyophilic properties of the solved particles, the particles travel towards the fluid-fluid interface and accumulate at the surface of a droplet thus forming a "Pickering emulsion" (see Fig. 2).

We have studied in detail the phase behaviour of our system and determined a phase diagram demonstrating for which values of the particle concentration, fluid decomposition or contact angle either a "Bijel" or a "Pickering emulsion" is formed [4].

It can be clearly observed in Figure 3 which shows a representative phase diagram in dependence on particle volume concentration and contact angle that a clear transition line exists. Our simulations do not only contribute to a better understanding of the phase behaviour but also allow to suggest how to move from one state to another. This could for example be useful for drug release applications where a drug is stored in the particle covered "Pickering" droplets. By changing the system parameters, the droplets open and release their content which can spread through the bicontinuous system. Another example could be a fostered transition from "Bijels" to "Pickering emulsion" in order to change the effective viscosity of the emulsion [4].

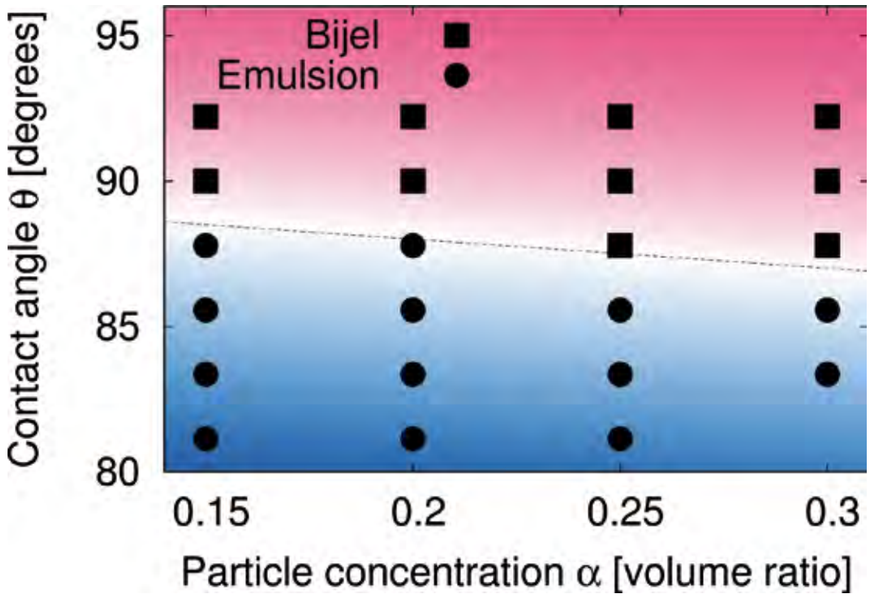

Figure 3: Phase diagram showing the occurence of "Bijels" and "Pickering emulsions" for different particle volume concentration and the contact angle a particle surface forms with the fluid-fluid interface [4].

Figure 4: Average droplet size L versus time for a "Pickering emulsion" showing three different regimes which are hard or impossible to study experimentally, but determine the properties of the final stable emulsion.

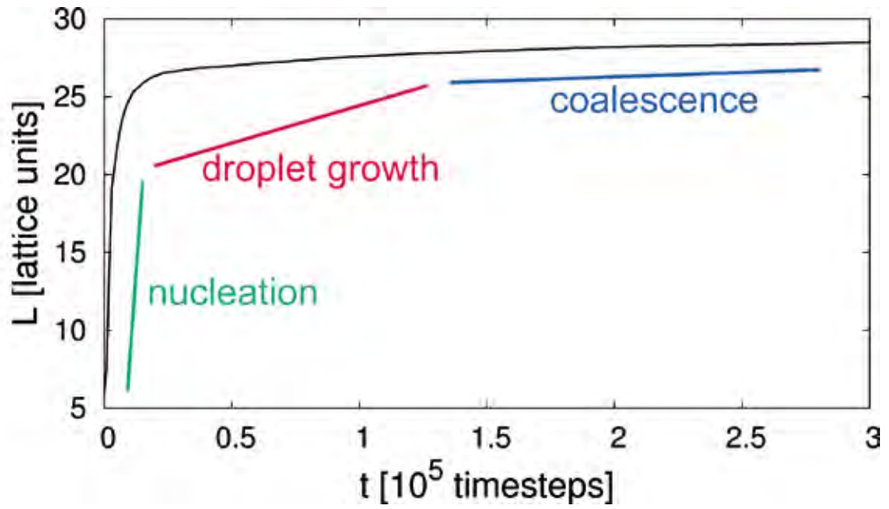

It is our current main focus to understand the rheological properties of "Bijels" and "Pickering emulsions" in detail. For this, further large scale sheared simulations will be performed in order to expede the complex non-Newtonian properties of these systems. Collecting the data from a number of such large scale runs will allow us to quantitatively compare the size distribution of the stabilized droplets as well as the rheological properties of such systems with experimental data.

Large scale parameter studies as well as investigations of the time dependent domain or droplet growth in "Bijels" and "Pickering emulsions" were performed. Here, different regimes were observed in the case of "Pickering emulsions" (see Fig. 4):

all simulations start from a random mixture of fluids and particles modelling a temperature quench. At the beginning of the simulation, small scale droplets nucleate due to ballistic motion of particles and fluids. After the nucleation regime, the droplets grow due to spinodal decomposition. At some point the phase separation comes to a

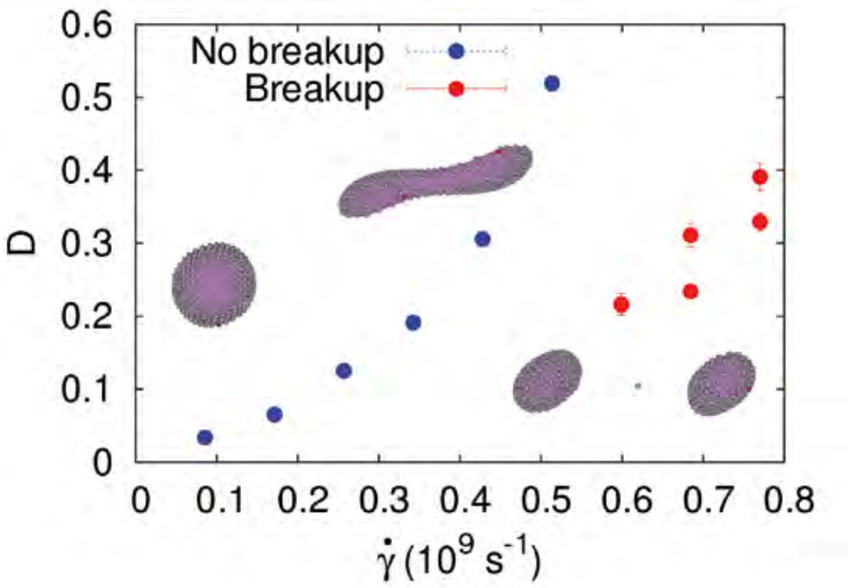

Figure 5:
Droplets stabilized by particles under shear deform and might even break for high shear rates.

halt and droplets can only continue to grow due to coalescence (Ostwald ripening), which is limited due to the particles that were captured at the fluid-fluid interface. This is one of the main processes leading to drop coarsening in emulsions and foams, allowing long-term stabilization. Such time dependent measurements are very hard or even impossible to perform experimentally, but the processes involved determine the final properties of the emulsion. Thus, we expect to be able to contribute to a better understanding how the final product depends on fluid composition, fluid properties, particle concentration, particle size, or the contact angle the fluids form with the particle surface.

A more fundamental study is on systems with only a few droplets covered by particles. We studied the droplet's stability, deformation and breakup behaviour and found a number of interesting phenomena as demonstrated in Figure 5: individual particle covered droplets placed in a shear flow with a given shear rate deform due to the external forcing. The deformation D shows a parabolic behaviour where the detailed shape of the curve depends on the particle properties and their concentration. For sufficiently high shear, the droplets can break up and continue to move individually as shown in the right part of the figure.

Simulation Method and Code Performance

Our massively parallel 3D LB code (LB3D) is based on Shan and Chen's multicomponent LB model, which can be utilized to simulate a number of miscible or immiscible fluids. In addition, amphiphiles were added to the model. Interactions between different fluid species are modelled by a mesoscopic force between the phases. The code was applied to a large number of problems to study for example the behaviour of binary and ternary fluid mixtures under shear, the formation of surfactant mesophases, flow in porous media, or to study fluid boundary interactions in detail [6,7,8]. Mesoscopic methods describing fluid flow can be extended to model (colloidal) suspensions. Here, the individual particles are treated by a molecular dynamics (MD) algorithm and momentum is transferred between them and the fluid after a sufficiently small number of timesteps. Inter-particle interactions which are not due to hydrodynamic forces are usually taken care of within the MD solver. For example, in the case of colloids, we include effective electrostatic interactions and van der Waals attraction, a lubrication force and Hertzian contact forces. Within the last three years our lattice Boltzmann code was combined with a parallel MD code enabling us to simulate (colloidal) suspensions and has further led to the development of a new simulation methodology for blood flow [9,10].

LB3D is a very efficient implementation scaling almost linearly up to 262,144 CPU cores of the Blue Gene/P system JUGENE in Jülich and similarly well on various other supercomputing

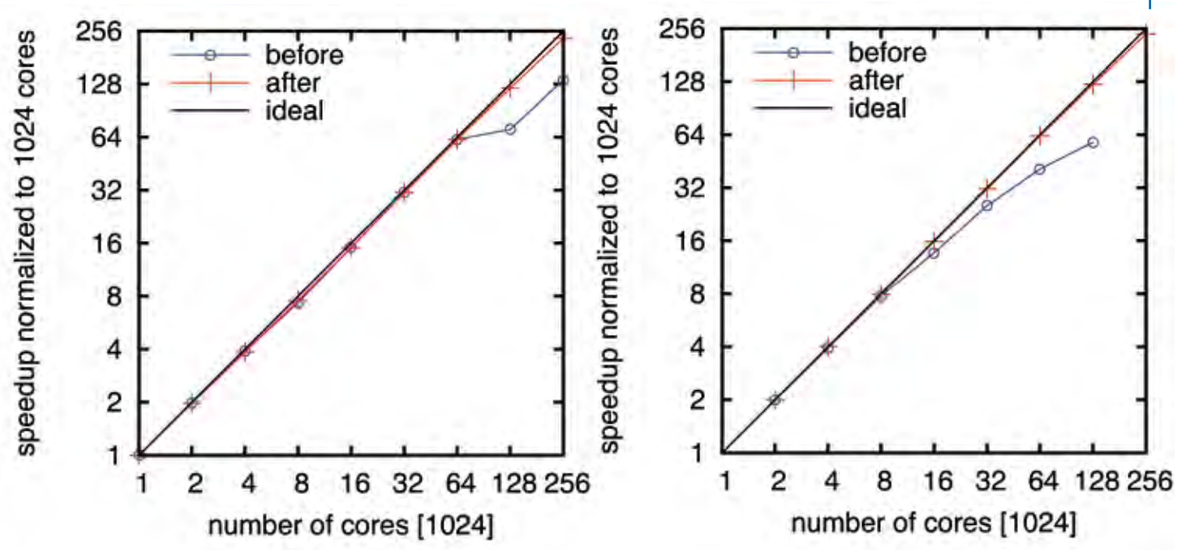

Figure 6: Strong scaling of LB3D on JUGENE: the system size is 1024x1024x2048 lattice units and the speedup was normalized to 1024 cores. The left figure shows the scaling behaviour for the pure lattice Boltzmann code. The optional Molecular Dynamics (MD) code adds 4.000.000 finite size particles to the system (right). The plots show the scaling behaviour before and after the Jülich Blue Gene/P Extreme Scaling Workshop 2011, where it was possible to identify a mismatch of simulation and hardware topology as well as a routine that scaled only mildly serially. Having sorted out these issues, our simulation code is now able to harness the computational power of almost the whole machine.

platforms we have access to. The code was one of the first applications being awarded the "goldstar-rating" for its excellent scalability by the Edinburgh Parallel Computing Centre (EPCC). Please see Figure 6 for details on the scaling behavior of our code on JUGENE and Figure 7 for a comparison of the performance on various other supercomputers.

Outlook

We already extended our simulation model to be able to describe non-spherical particles such as ellipsoids. Ellipsoidal particles are more suitable to describe for example Laponites or certain kind of viruses or other biological systems. We plan to use the improved version of the code to investigate the influence of the particle shape on the properties of the emulsion including the formation of "Bijels" and "Pickering" systems as well as the stability and breakup behaviour of droplets.

The existing lattice Boltzmann code includes a well tested and established mean field model for surfactants. Similarities and differences between surfactants and particles will be studied. It is planned, for example, to compare to data obtained from earlier projects where we investigated the influence of surfactants on the process of spinodal decomposition and the formation of microemulsions. The current work on particle-laden multiphase flows allows to draw conclusions on if and how the known domain growth behaviour can be reproduced using particles by directly comparing the simulation results from the current work to results obtained earlier. These studies are of fundamental relevance since they cannot easily be performed experimentally: while we are able to only modify specific properties of the system, in experiments it is in general not possible to change for example the interactions between particles or surfactant molecules without inducing changes in the solvent-solvent interactions.

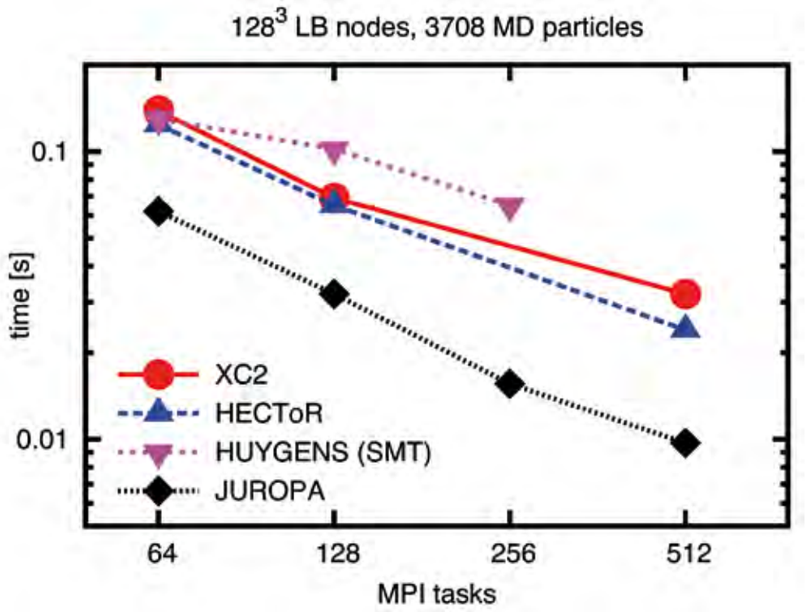

Figure 7: Scaling and performance comparison of LB3D on JUROPA at JSC, the XC2 at SSC Karlsruhe, HECToR at EPCC, and Huygens at SARA. The studied system is comparably small which is the reason for the less good scaling in comparison to JUGENE.

Acknowledgments

We thank the Jülich Supercomputing Centre and the Scientific Supercomputing Centre Karlsruhe for the allocated CPU time and technical support. The project was made possible by a large scale grant of the Gauss Centre for Supercomputing.

References

[1] **Ramsden, W.,**
Separation of solids in the surface-layers of solutions and "suspensions", Proc. R. Soc. London, 72:156, 1903

[2] **Pickering, S. U.,**
Emulsions. J. Chem. Soc. Trans., 91: 2001, 1907

[3] **Binks, B. P., editor**
Colloidal particles at liquid interfaces, Oxford University Press, 2006

[4] **Jansen, F., Harting, J.**
From Bijels to Pickering emulsions: a lattice Boltzmann, study. Phys. Rev., E 83, 046707, 2011

[5] **Stratford, K., Adhikari, R., Pagonabarraga, I., Desplat, J.-C., Cates, M. E.**
Colloidal Jamming at Interfaces: A Route to Fluid-Bicontinuous Gels. Science, 309:2198, 2005

[6] **Harting, J., Harvey, M., Chin, J., Venturoli, M., Coveney, P. V.**
Large-scale, lattice Boltzmann simulations of complex fluids: advances through the advent of computational grids, Phil. Trans. R. Soc. Lond. A, 363:1895–1915, 2005

[7] **Narváez, A., Zauner, T., Raischel, F., Hilfer, R., Harting, J.,**
Quantitative analysis of numerical estimates for the permeability of porous media from lattice-Boltzmann simulations, J. Stat. Mech: Theor. Exp., 2010:P211026, 2010

[8] **Harting, J., Kunert, C., Hyväluoma, J.**
Lattice Boltzmann simulations in microfluidics: probing the no-slip boundary condition in hydrophobic, rough, and surface nanobubble laden microchannels, Microfluidics and Nanofluidics, 8:1, 2010

[9] **Janoschek, F.**
Gekoppelte Molekulardynamik- und Gitter-Boltzmann-Simulation von Suspensionen, Diploma thesis, Universität Stuttgart, 2009

[10] **Janoschek, F., Toschi, F., and Harting, J.**
Simplified particulate model for coarse grained hemodynamics simulations,d Phys. Rev., E, 82:056710, 2010

- Jens Harting
- Fabian Jansen
- Stefan Frijters
- Florian Janoschek
- Florian Günther

Institute for Computational Physics, University of Stuttgart, Stuttgart

Department of Applied Physics, Eindhoven University of Technology, Eindhoven

Functional Molecules at Surfaces: The Role of dispersive Interactions

Introduction

Our work focuses on first-principles simulations of the molecular switch azobenzene, a promising candidate component in a future molecular nanotechnology. Since the advent of semiconductor based microelectronics, the component density of integrated circuits has increased at a steady exponential rate. As component sizes approach the nano-scale, maintaining this pace will require a shift towards alternative materials more versatile than doped silicon. One such approach is to exploit the great variety and adaptability of organic molecules in component design.

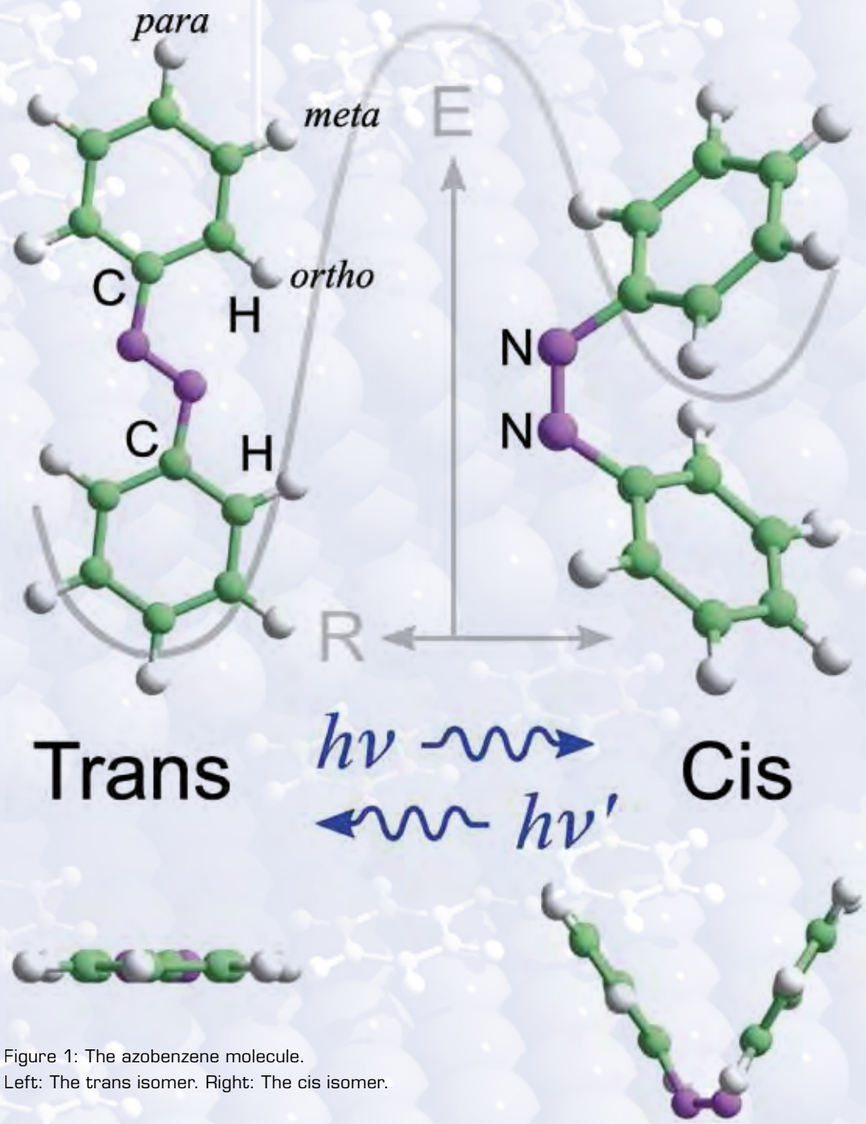

Figure 1: The azobenzene molecule.
Left: The trans isomer. Right: The cis isomer.

In the context of such a molecular nanotechnology, molecules with properties bi-stably and reversibly modifiable by external stimuli, so-called molecular switches, are a research topic of paramount importance [1]. The azobenzene molecule (C_6H_5-N_2-C_6H_5, see Fig. 1) qualifies in this class by undergoing a reversible trans-cis isomerization ("switching"), and has come under intensive experimental scrutiny. Numerous potential applications have been proposed e.g., as a light-driven actuator, or data storage medium.

While currently well characterized in solution or gas phase, properties of switches adsorbed at solid surfaces are more relevant for microelectronics, and have in recent years emerged as an important research field. The switching function being an innate property of the azobenzene electronic structure, the choice of substrate and substrate coupling is in this context non-trivial:

While bonding strong enough to localize and order switches is desirable, significant hybridization of the electronic frontier orbitals, or steric hindrance due to substrate registry is not. Close-packed coinage metal (Cu, Ag, Au) (111) surfaces ostensibly offer a reasonable such balance. However, even at these substrates, the switching yield of pure azobenzene ranges from low at Au, to zero at Ag and Cu. Consequentially, strategies of decoupling the chromophore moiety by functionalization with ligands acting as adsorbate-substrate spacers have been developed. Of this class of azobenzene derivates, the most studied, and most successfully switched, is tetra-tert-butyl-azobenzene (TBA) [2].

While experimental literature on these molecular switches is extensive, theoretical work is much more scarce. This is unsurprising, given that a predictive theoretical description of azobenzene and TBA at coinage metals requires accurate treatment of heavy-element chemistry, metallic systems, extended or periodic systems, large organic molecules and weak, dispersive van der Waals (vdW) interactions typical of aromatic molecules at surfaces.

The latter aspect is particularly difficult, since standard exchange-correlation (xc) functionals for the work-horse usually employed for surface adsorption problems, density functional theory (DFT), do not account for such vdW interactions. This fact, combined with the computational intractability of higher-level theory for most scientifically relevant systems has made weak

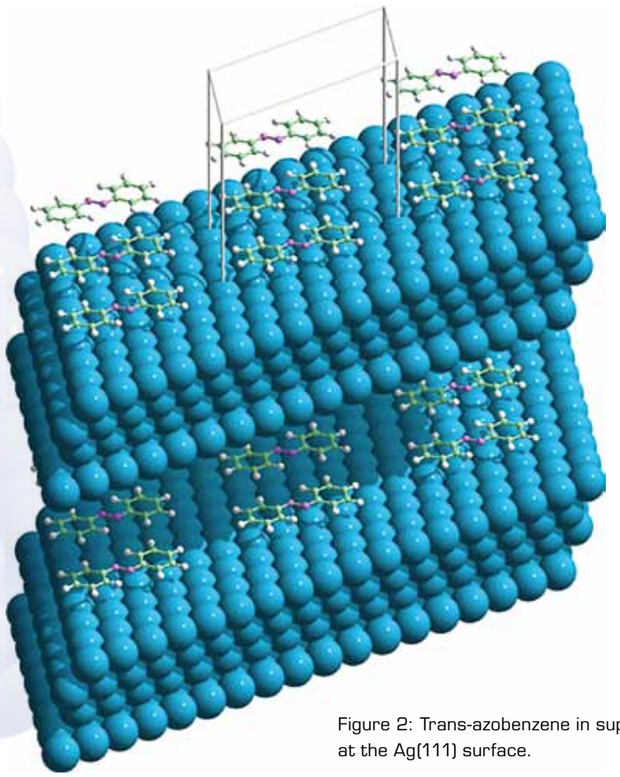

Figure 2: Trans-azobenzene in supercells at the Ag(111) surface.

interactions one of the most-studied topics in contemporary first-principles and ab initio modelling.

Methodology / HLRB Implementation

The approach we elected to explore is that of semi-empirical vdW corrections to DFT. In this method, the missing vdW interactions in DFT are approximated by a sum of simple semi-empirically parametrized two-body potentials derived from the London approximation. Thus, the dispersion energy enters as a simple function of the geometry, the value of which is added to the DFT total energy. The advantage of this method is the negligible computational cost compared to DFT calcultations themselves. The disadvantage of the semi-empiricity might be an in general limited transferability which needs to be carefully scrutinized.

On the computational side, we implemented a number of published semi-empirical vdW correction schemes in the CASTEP DFT code [3]. These include the popular Grimme (G06) [5] correction, and the state-of-the-art scheme due to Tkatchenko and Scheffler (TS) [6], the latter relying on a Hirschfeld analysis [9] that needed to be implemented directly into CASTEP. The rest of the implementation is written as an independent FORTRAN 90 module, licensed under LGPL, that can be easily interfaced to any periodic or finite-system DFT packages alike. For two other high-profile DFT codes, such an interest has already been expressed. CASTEP expands the electronic wave-function in a planewave basis set in periodic boundary conditions, with the wave-function close to the atomic nuclei represented by ultrasoft pseudopotentials.

Figure 3: Runtime of a test calculation at HLRB, as a function of CPUs per k-point. Note the break-point at 96 CPUs.

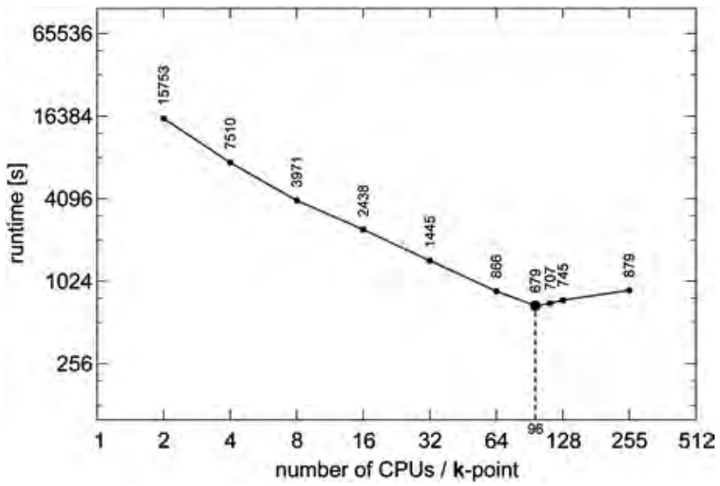

Isolated molecules adsorbed at an extended metal surface were modeled in the so-called supercell approach (see Fig. 2), in which the molecule is adsorbed on a crystal slab of metal atoms. Minimizing adsorbate-adsorbate interactions, its thickness and lateral extent is increased until the property of interest converges.

The electronic wave-function is calculated in Fourier space. The latter is sampled by so-called k-points, which form a natural, commonly employed parallelization strategy for plane-wave DFT codes requiring a minimal amount of communication. However, for systems on the scale of the azobenzene molecule, only a few k-points need to be considered. Therefore, parallel scaling beyond this level (which has been taken into account in the implementation of the Hirshfeld analysis) is crucial to the feasibility of this study.

In the calculations presented below, we have run CASTEP parallelized over 3, 4 and 6 k-points, corresponding to 288, 384 and 576 CPUs, with a minimum of parallel overhead. The robustness of the CASTEP code combined with the stability and the excellent scaling at HLRB have allowed for an unusually high result-per-CPU-hour efficiency.

Results

In our initial study, published in Physical Review B [4], we have analyzed the adsorption of azobenzene at coinage metal (111) surfaces. The near absence of the vdW interaction in common semi-local DFT functionals leads to a dramatic bonding picture, which, if true, would have radical consequences for our understanding of the azobenzene switching mechanism at these substrates: The cis isomer is overstabilized, at the Cu(111) surface to the point of reversing the sign of the gas-phase relative isomer stability. In addition, where the covalent interaction of the azobridge and the surface is strong, the trans isomer is severly distorted by Pauli-repulsion.

Since the vdW interaction is strongest for the phenyl rings, the Pauli repulsion experienced by the rings at short distances is substantially offset by the semi-empirical vdW correction (see Fig. 4). Consequentially, the rings are restored towards their gas phase geometry, and the equilibrium adsorbate-substrate distance is substantially reduced. The cis isomers are less affected by the correction, restoring or even surpassing also the relative cis-trans stability towards the gas-phase value. These results have been summarized in a second publication in Physical Review B [7].

Such radical semi-empirical corrections to state-of-the-art DFT cry for verification by more accurate theory or experiment. For systems of this size, only the latter is presently a viable option. Therefore we initiated a collaboration with the experimental groups of Prof. Dr. Stefan Tautz at Forschungszentrum Jülich, and Prof. Dr. Martin Wolf at the Freie Universität Berlin. Using novel Near-Incidence X-ray Standing Wave (NIXSW) methods, the former were able to measure select geometry parameters of azobenzene at Ag(111) with extreme accuracy. The latter group is measuring the adsorption energy of the same system using Thermal Programmed Desorption (TPD).

Comparison to measured TPD activation energies requires the corrrection of the energetics for the change in vibrational Zero-Point Energy (ZPE) upon desorption. They can be obtained by

Figure 4: The effects of vdW corrections on the DFT azobenzene geometry.

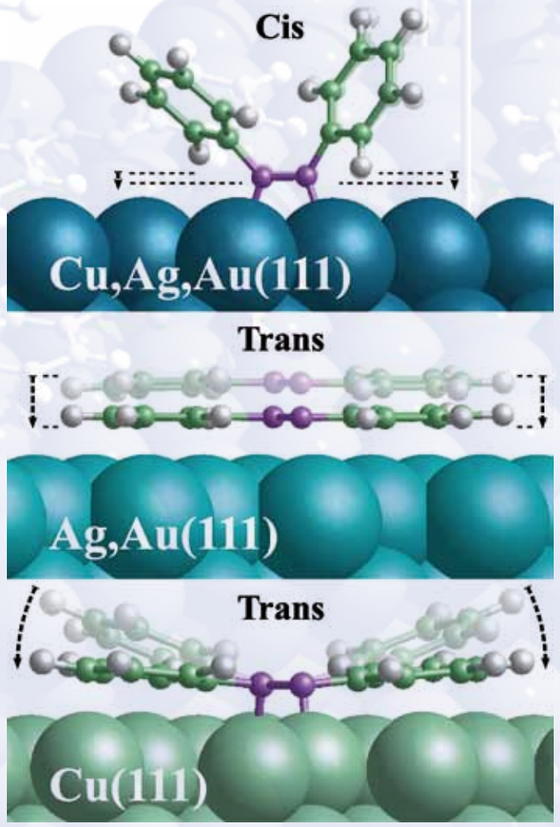

means of the nuclear Hessian also with dispersion corrected DFT, which unfortunately has to be obtained numerically by finite differences. Therefore, the CPU time awarded to us at HLRB was a crucial premise for the two times 3N necessary total energy and force evaluations for large molecules of N constituent atoms. As the force overestimation decays faster with depth than that of the energy, one can expect vdW corrected geometries to be more accurate than the corresponding binding energies. This expectation is confirmed by Figure 5.

First results, published in Physical Review Letters [8], show a remarkable agreement between in particular the TS-corrected DFT results and the experimentally measured geometry: The theoretical error in the equilibrium adsorbate-substrate distance is smaller than 0.1 Ångström [Å] (see Fig. 5).

This is particulary promising, since it implies this approach very likely to be accurate also for related derivates of the azobenzene molecule, including TBA. We have already successfully used HLRB CPU time to fully optimize the structures of TBA in large supercells at Ag(111) and Au(111) surfaces (see Fig. 6), with both vdW corrected DFT as well as uncorrected DFT and also obtained the aforementioned Hessians. In our latest publications, we show that the calculated vibrational properties are in very good agreement with available experimentally measured IR spectroscopy data [10] and point out the important consequences about the relations between bonding to the surface and switching functionality [11], significantly changing the common, previously accepted picture.

Figure 5: ZPE-corrected binding energy curves of trans- azobenzene at Ag(111), using standard DFT (PBE), the G06 correction scheme (PBE+G06), and the TS correction scheme (PBE+TS). Experimental results for the geometry / adsorption energy come from NIXSW / TPD.

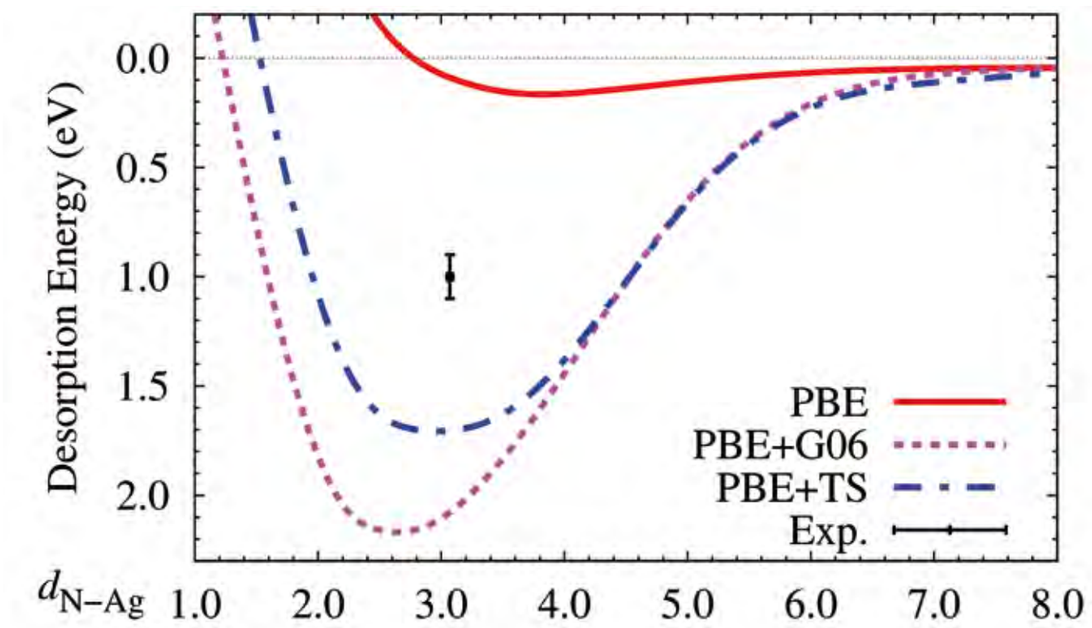

References

[1] **(Eds.) Wolf, M., von Oppen, F.**
Special issue on Elementary Processes in Molecular Switches at Surfaces, Appl. Phys., A 93, 2008

[2] **Wolf, M. & Tegeder, P.**
Reversible molecular switching at a metal surface: A case study of tetra-tertbutyl-azobenzene on Au (111), Surf. Sci., 603, 15061517, 2009

[3] **Clark, S. J., Segall, M. D., Pickard, C. J., Hasnip, P. J., Probert, M. I. J., Refson, K., and Payne, M. C.**
First principles methods using CASTEP, Z. Kristallogr., 220, 567-570, 2005

[4] **McNellis, E., Meyer, J., Dehghan Baghi, A., & Reuter, K.**
Stabilizing a molecular switch at a solid surface: A density functional theory study of azobenzene on Cu(111), Ag(111), and Au(111), Phys. Rev., B 80, 035414, 2009

[5] **Grimme, S.**
Semiempirical GGA-Type Density Functional with a Long-Range Dispersion Correction, J. Comput. Chem., 27, 1787-1799, 2006

[6] **Tkatchenko, A. & Scheffler, M.**
Accurate Molecular van der Waals Interactions from Ground- State Electron Density and Free-Atom Reference Data, Phys. Rev. Lett., 102, 073005, 2009

[7] **McNellis, E., Meyer, J., & Reuter, K.**
Azobenzene at Coinage Metal Surfaces: The Role of Dispersive van der Waals Interactions, Phys. Rev., B, 80, 205414, 2009

[8] **Mercurio, G., McNellis, E. R., Martin, I., Hagen, S., Leyssner, F., Soubatch, S., Meyer, J., Wolf, M., Tegeder, P., Tautz, F. S., Reuter, K.**
Phys. Rev. Lett., 104, 036102, 2010

[9] **Hirshfeld, F.**
Theor. Chim. Acta, 44, 129, 1977

[10] **McNellis, E. R., Bronner, C., Meyer, J., Weinelt, M., Tegeder, P., Reuter, K.**
Phys. Chem. Chem. Phys., 12, 6404 2010

[11] **McNellis, E. R., Mercurio, G., Hagen, S., Leyssner, F., Meyer, J., Soubatch, S., Wolf, M., Reuter, K., Tegeder, P., Tautz, F. S.**
Chem. Phys. Lett., 499, 247, 2010

Figure 6: The TBA molecule at the Ag(111) surface.

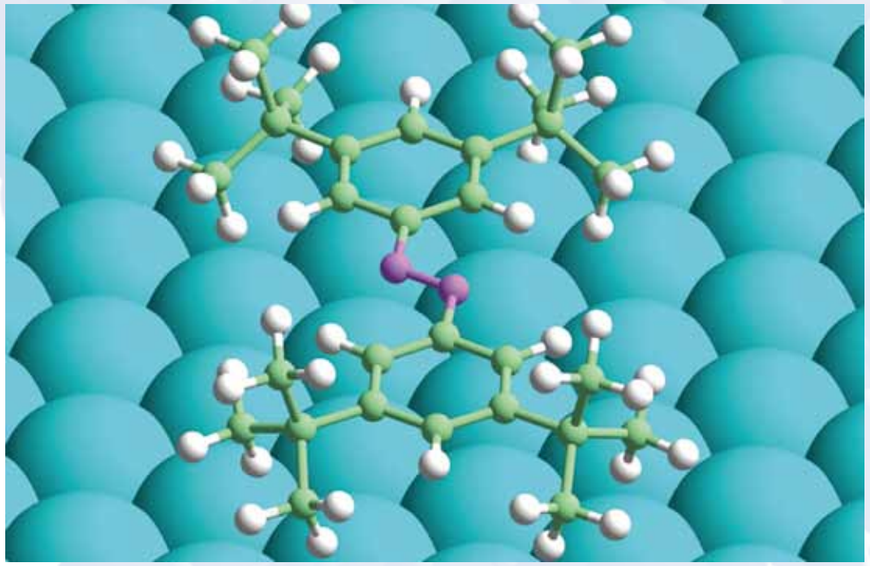

• Erik McNellis
• Jörg Meyer
• Karsten Reuter

Fritz-Haber-Institut der MPG, Berlin

Technische Universität München

Galaxy Cluster Simulations: The FEARLESS Approach

Clusters of galaxies are the biggest gravitationally bound objects in the Universe, and the basic constituents of the cosmic large-scale structure. Their typical size is on the order of some Megaparsec (1 Mpc = 10^6 parsec = 3.26×10^6 light years; for comparison, the diameter of the disk of our Milky Way galaxy is around 50 kpc), and their mass is up to some 10^{15} solar masses. Already in the last century, astronomers found out that the component visible with optical telescopes (the galaxies) is just a minor part of their mass. Nowadays we believe that only ~ 5% of the cluster mass is condensed into galaxies and stars. Another 10% is formed by hot (10^7 – 10^8 K), X-ray emitting diffuse gas, and the biggest part (about 85%) is constituted by a still unknown dark matter, which does not emit light and interacts with the ordinary (baryonic) matter only through the gravitational force. Being large and bright in optical wavelengths and X-rays, galaxy clusters are very important probes of the structure of our Universe, of its constituents and of theories of its evolution.

According to the most accepted mathematical description of our Universe (the so-called standard cosmological model), the cosmic structure has developed from tiny fluctuations of dark matter density, imprinted right after the Big Bang, which evolved by gravitationally attracting the surrounding material, including the baryonic gas. In this process the large-scale structure develops a filamentary structure (Fig. 1), where galaxy clusters form at the intersection of filaments. This process is very complex and cannot be studied in its details with simple analytical formulae: for this reason, hydrodynamical simulations of the evolution of gas and dark matter are the principal tool of theoretical investigation in this field.

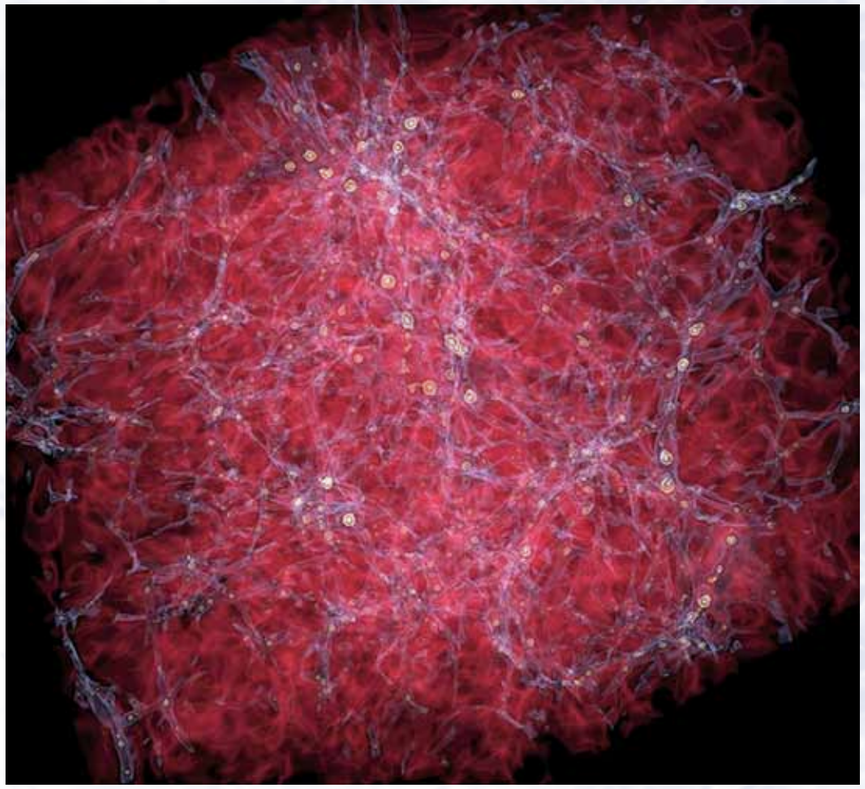

Figure 1: Volume rendering of baryon density at current epoch, in a hydrodynamical simulation of the evolution of the cosmic large-scale structure.

Numerical Tools: the FEARLESS Approach

In many astrophysical fluids, presumably including the intra-cluster medium (ICM), the flow is in a regime of fully developed turbulence (in the language of fluid dynamics, it has a Reynolds number around or larger than 10^3, possibly by many orders of magnitude). As a result, the length scale separation between the integral length scale for turbulence injection and the Kolmogorov scale, where the kinetic motions are dissipated into internal energy by viscous effects, is so large that it cannot be resolved in state-of-the-art direct numerical simulations, even using adaptive mesh refinement (AMR) in grid-based codes. In many fields of computational fluid dynamics, the influence of unresolved turbulence on the resolved scales is therefore modelled by means of heuristic subgrid scale (SGS) models, coupled to the large scales of the system, for which the fluid equations are solved (Large Eddy Simulations, LES for brevity). A new numerical tool, arising from the combined use of AMR and LES, has been recently developed for the study of turbulence in astrophysical fluids. This novel approach (1), called FEARLESS (Fluid mEchanics with Adaptively Refined Large Eddy SimulationS), unites the adaptive refinement of the regions where turbulent flows develop with a turbulence model (2) for the unresolved turbulence energy, governing the production, the diffusion and the dissipation of kinetic energy on subgrid scales. The resulting tool, implemented in the AMR, grid-based ENZO code, is especially useful in simulations of turbulent clumped flows, where turbulence is localized and intermittent, like clusters of galaxies.

Turbulence Modeling and the Physics of the Intracluster Medium

The FEARLESS numerical technique is described in detail in (1), where it is also first applied to numerical simulations of galaxy cluster evolution. Our numerical setup was especially designed for the study of the evolution of a single cluster, evolving in a region where nested static grids and AMR was allowed, in order to have a large dynamical range and reach an effective spatial resolution of 7.8 kpc h^{-1} in the ICM (h = 0.7 is the Hubble parameter). The computational box has a size of (128 Mpc h^{-1})^3, a volume large enough to accommodate the growth of a massive cluster.

We find that the production of turbulence is closely correlated with merger events occurring in the cluster environment, and its dissipation locally affects the cluster energy budget. In minor mergers, the motion of the subclumps accreted by the main cluster induce shearing instabilities in the ICM. The injection of turbulence by merger events is well elucidated in the two panels of Figure 2: for example, one can observe the small subcluster immediately on the left of the cluster core, which is moving downwards around the centre and has stirred the ICM in its “turbulent wake”,

as indicated by the large value of the turbulence energy in that region. This behavior is reminiscent of idealized subcluster simulations, for example (3).

While minor mergers as those shown above stir the ICM only locally, mergers with mass ratio close to unity can deeply perturb the cluster structure and inject volume-filling turbulences (4). These events are extremely interesting for understanding the diffuse radio emission in galaxy clusters (radio halos and relics), and for constraining the acceleration mechanisms of the non-thermal particles responsible for this emission.

Data Visualization and technical Details

Data analysis and visualization are attracting growing interest, because these steps are important for the physical interpretation of the simulation results, and for the creation of simulated observations which are crucial for the design of future observational instruments and strategies. Nowadays, the analysis tools have to cope with the ever increasing complexity and output size of the simulations.

For a better understanding of the turbulence production caused by structure growth in the ICM, two animations (5) have been created from the simulations snapshots shown in Figures 2 and 3.

The visualization has been produced using yt (http://yt.enzotools.org/) a multi-code analysis toolkit for astrophysical simulation data based on Python. The production of these two movies required parallel data analysis, and was made possible by the computational resources available at the LRZ. In order to create high quality images and movies, we produced and stored 550 data dumps, corresponding to 1.6 million AMR grid files (about 0.5 TB of data). The analysis has been performed in parallel, on 220 cores. The final results are a good example of the close link between data handling and High Performance Computing, which is the key for fully exploiting the resources available on current and future machines.

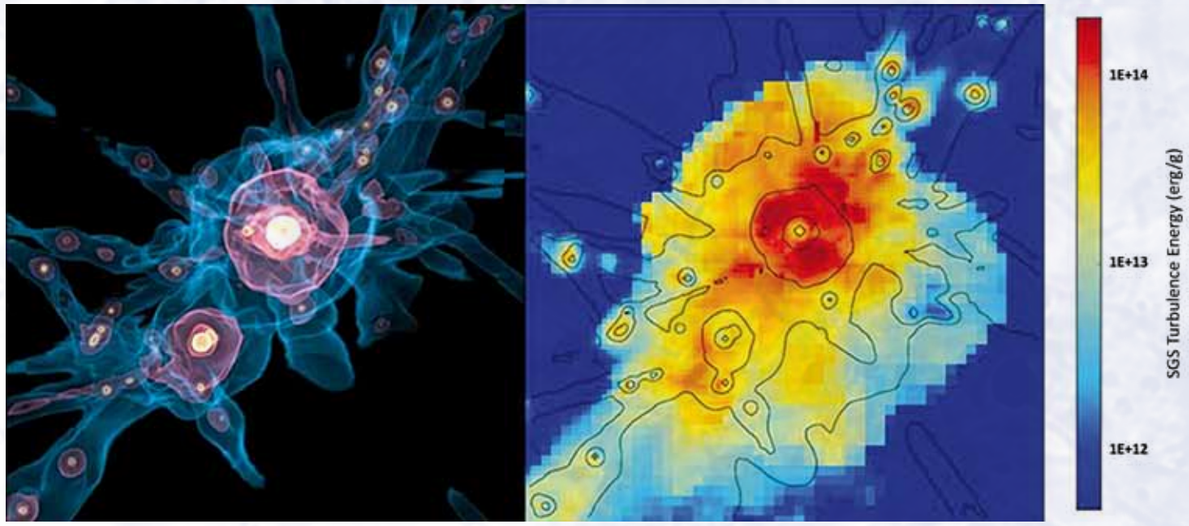

Figure 2: Two different visualizations of the same galaxy cluster. Left: volume rendering of the density; right: projection of the unresolved turbulence energy modeled in FEARLESS by a subgrid scale model.

Acknowledgements

The ENZO code is developed by the Laboratory for Computational Astrophysics at the University of California in San Diego (http://lca.ucsd.edu). Special thanks to Helmut Satzger (LRZ) for support with the data storage and the analysis parallelization.

References

[1] **Maier, A., Iapichino, L., Schmidt, W., Niemeyer, J. C.**
"Adaptively Refined Large Eddy Simulations of a Galaxy Cluster: Turbulence Modeling and the Physics of the Intracluster Medium", ApJ, vol. 707, pp. 40-54, 2009

[2] **Schmidt, W., Niemeyer, J. C., Hillebrandt, W.**
"A localised subgrid scale model for fluid dynamical simulations in astrophysics. I. Theory and numerical tests", A&A, vol. 450, pp. 265-281, April 2006

[3] **Iapichino, L., Adamek, J., Schmidt, W., Niemeyer, J. C.**
"Hydrodynamical adaptive mesh refinement simulations of turbulent flows - I. Substructure in a wind", MNRAS, vol. 388, pp. 1079-1088, 2008

[4] **Paul, S., Iapichino, L., Miniati, F., Bagchi, J., Mannheim, K.,**
"Evolution of Shocks and Turbulence in Major Cluster Mergers", ApJ, vol. 726, p. 17, 2011

[5] The described animations are available at the website: http://www.ita.uni-heidelberg.de/~luigi/movies.html

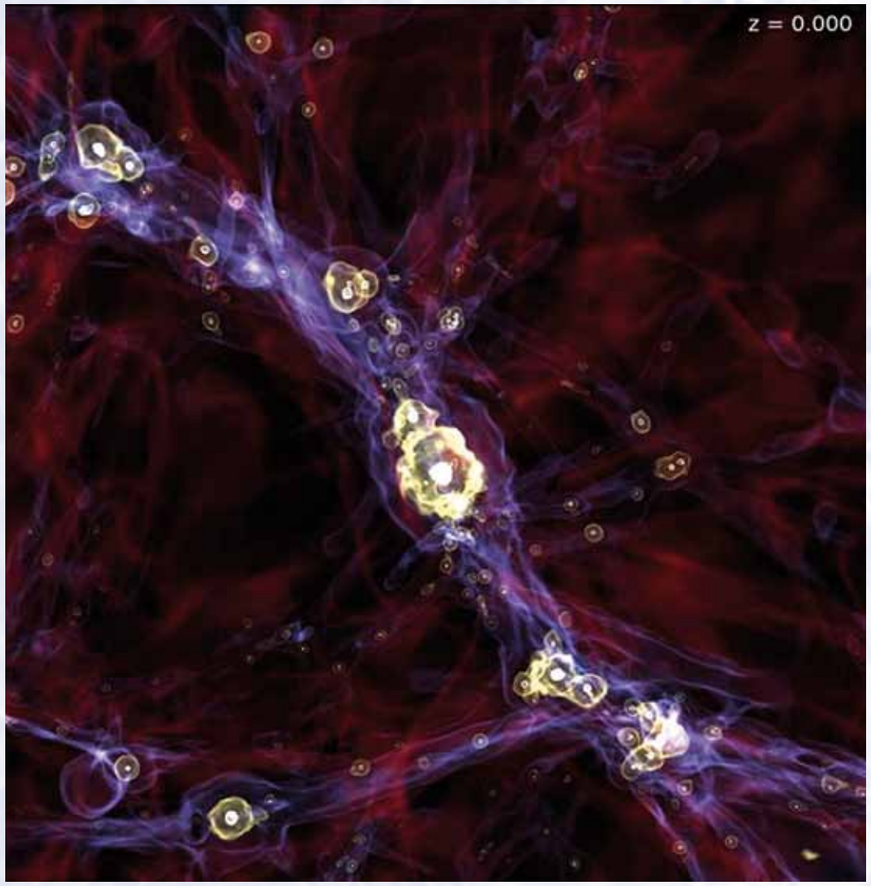

Figure 3: Volume rendering of baryon density centered on a cluster surrounded by the filamentary large-scale structure and growing its mass by accretion of subclumps.

• Luigi Iapichino

Zentrum für Astronomie der Universität Heidelberg, Institut für Theoretische Astrophysik

Simulating Giant Spiral Galaxies in the early Universe

Introduction

In this article we review the formation of clumpy star-forming disk galaxies in the early Universe.

Recent high resolution observations with the SINFONI instrument by the Infrared Group at MPE in Garching [1,2], revealed the existence of a previously unknown class of massive galaxies at a cosmological redshift of z~2, i.e. approximately 3 billion years after the Big Bang. Those galaxies have star-forming disks whose kinematics is dominated by rotation, albeit with exceptionally high velocity dispersions compared to present day spiral galaxies. The star formation in these disks is distributed in large and massive clumps, unfamiliar in the local Universe. The detailed formation mechanism of these galaxies is poorly understood, and a few fundamental questions arise from these observations. How can such high star-formation rates, which in the local Universe exist only in merging starburst galaxies, be sustained over billions of years? What mechanism provides the fuel for such high star-formation rates? What mechanism drives the highly turbulent motions? What is the physics behind the large star-forming clumps - how do they form, what is their internal structure, and how long can they survive?

Recent numerical work led to the hypothesis that those galaxies grow from high-flux "cold streams" of gas arriving from large cosmological distances directly into the inner regions of dark matter halos. Yet, those studies were not able to resolve details of the formation and evolution of the central galaxies. Several "zoom-in" simulations have been performed to study the formation of galaxies at z~2, either with the SPH (smoothed particle hydrodynamics) method or with grid codes such as in the AMR (adaptive mesh refinement) technique. [3] have used the AMR technique on a few "zoomed-in" halos, and have demonstrated that star-forming fragmenting gas disks can form in a cosmological context. However, it is not yet clear whether the intensity of star-formation and the large masses and sizes of some of the observed galaxies can be generally reproduced with the current models. In addition, the model prerequisites for reproducing the observed systems are uncertain and the implications of such a new model for galaxy formation as a whole is unclear.

It is an open question of whether giant clumps survive long enough to become bound stellar systems and/or sink to the center of the galaxy and contribute to the buildup of a bulge. Conflicting theoretical claims have been made, while observations have not been conclusive so far. No cosmological or isolated simulation included strong feedback from star-formation so far. As the "clumpy phase" is ubiquitous in the early Universe, it is of fundamental importance for our understanding of galaxy formation to know whether the giant clumps play an important role in building bulges/spheroids.

Computations

We use the N-body/SPH code Gadget-2 with different physically-motivated recipes for metal dependent radiative cooling of gas, star-formation, as well as feedback from star-formation in form of thermal energy, mass, momentum and production of metals. We wish to test the significance of the different ingredients of the model to the formation of realistic star-forming galaxies, by changing, among other factors, the super-wind efficiency and speed, the cooling rate as a function of gas composition and ionising background and the equation-of-state of the star-forming gas.

Gadget-2 solves the Poisson equation for a Newtonian gravitational potential in an expanding universe for all types of particles in the simulation. Gravitational forces are computed in Gadget-2 using the TreePM method. In addition, gas particles are subject to pressure forces, and so the Euler equations (conservation of mass, momentum and energy) are solved under an assumption of an equation of state for the gas. The hydrodynamics is solved using an entropy-conserving SPH method. Radiative cooling of the gas is computed using tabulated cooling functions. The code also implements a sub-grid Monte-Carlo model for star-formation, which involves integrating a few linear differential equations.

We have implemented a sophisticated scheme for the generation of initial conditions for zoom-in simulations. This scheme follows at high resolution only the minimal number of particles needed to resolve the galaxy of interest, therefore offers strong computing-time efficiency. We use that scheme on smoothly accreting z~2 halos chosen from a large cosmological simulation of dark matter [4].

We use three resolution levels that are described in Table 1. Table 2 gives the required CPU resources for each. Our tests have shown that the scaling in CPU-time between each resolution step (that has 8 times more particles) is slightly worse, but close, to linear, i.e. there is an increase of ~15 in CPU-time. Our tests also show that the required CPU-time increases by ~1.3 with each increase of CPU number by a factor of 2. That is, the gain in wall-time is ~1.5.

Resolution	Dark matter particle mass [solar masses]	Gas particle mass [solar masses]	Gas gravitational smoothing [comoving parsecs]	128
Low	29000000	5900000	400	-
Medium	3600000	740000	200	14
High	450000	92000	100	275

Table 1: Resolution levels

Number of CPUs \ Resolution	16	32	64	128
Low	0.5	0.75	-	-
Medium	7	8.5	11	14
High	-	170	220	275

Table 2: CPU-time (in 1,000 hours)

We run the simulations from z~43 to z=2 with the extended GADGET-2 version of [5]. This version includes treatment of several baryonic processes, such as ionisation and heating by a uniform background radiation in the optically thin limit, atomic cooling down to T=10^4 degrees Kelvin from hydrogen and metals, star-formation and feedback, as well as mass loss and metal enrichment from AGB stars and supernovae (of both types).

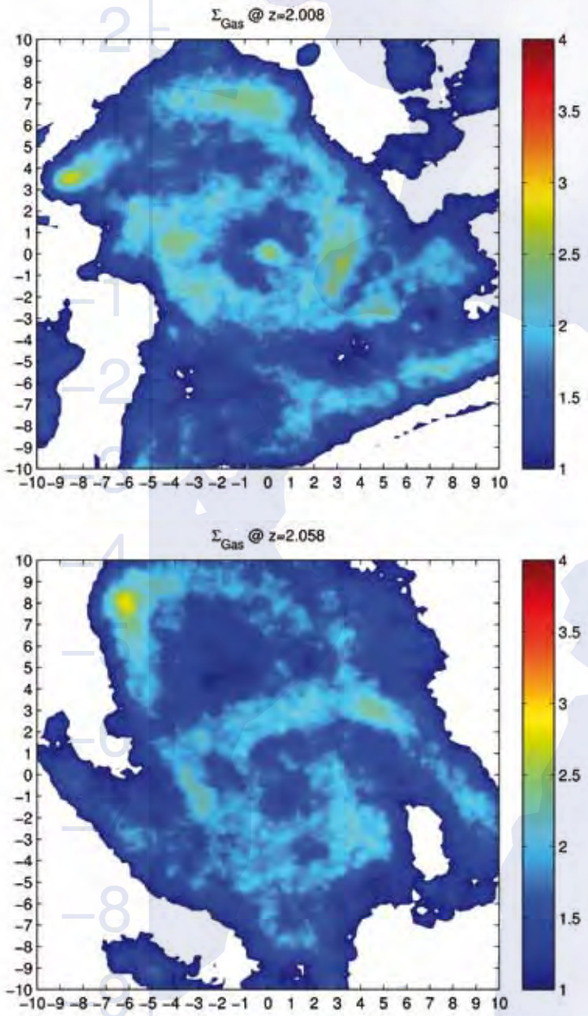

Figure 1, 2: Example of a gas disk

The feedback scheme developed by [5] builds on the kinetic feedback scheme of [6], where star-forming gas particles are stochastically being given velocity kicks. Subsequently, they can decouple from the hydrodynamics for a short time to allow them to propagate out of their star-formation sites, or to be completely removed from their galaxies. [5] used this mechanism with a significant change to the two parameters that control the wind: the magnitude of the kick and the mass loading factor (the ratio between the mass outflow rate and the star formation rate). The stellar velocity dispersion of the galaxy, used to determine the wind velocity, is calculated from its mass based on an on-the-fly Friends-of-friends (FOF) finder. Introducing velocity dispersion-dependencies, these two parameters now scale with the mass of the galaxy, following the theory of momentum-driven winds [7] induced by the radiation pressure from young stars, and in accordance with observational evidence.

Scientific Results

The global properties of our preliminary sample of six halos match observations well. First, in the star-formation rate versus stellar mass plane our simulated galaxies populate the same region as observed z~2 galaxies. The stellar-to-halo mass ratios as a function of halo mass for our simulated halos are low compared to the cosmic baryon fraction, like the case in observed mass functions. Both of these global correlations are a result of the wind model, without which they would compare much worse to observations. We consider them important constraints that simulated galaxies should

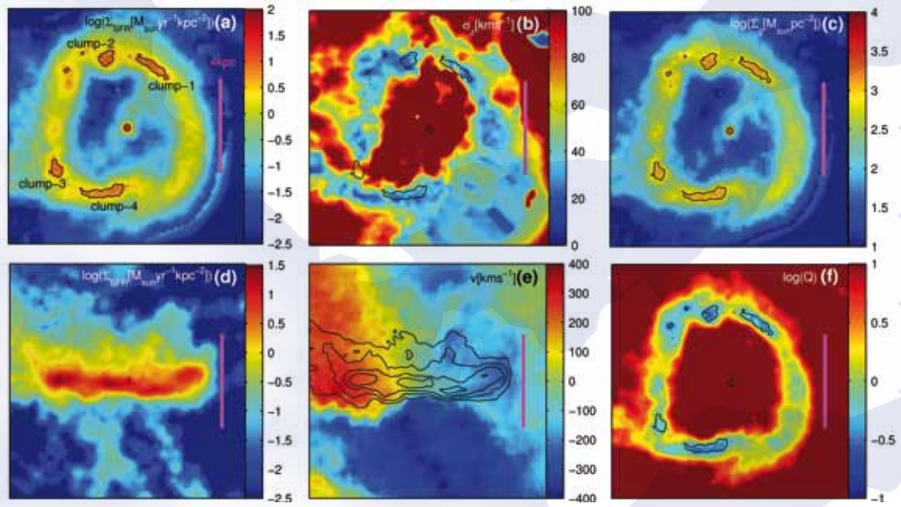

Figure 3: Example of a gas disk, surface density and velocity fields from different viewing angles.

match to be considered realistic. The galaxy gas fractions we obtain are high (~50%) and agree with recent observations [8].

Figures 1 and 2 show gas surface density maps of two exemplary disks. One particular halo, named s224, is shown in Figure 3. At z=2 the gas is distributed in a pronounced ring of ~3 kiloparsec in radius (Figure 3(c)). An edge-on view of the gas disk is shown in Figure 3(d). The gas and stellar masses in the disk region are ~1.3x10^{10} and ~2.6x10^{10} solar masses, respectively, with a star-formation rate of ~50 solar masses per year at z=2, being quite stable since z~3. In many respects this galaxy is typical of our sample.

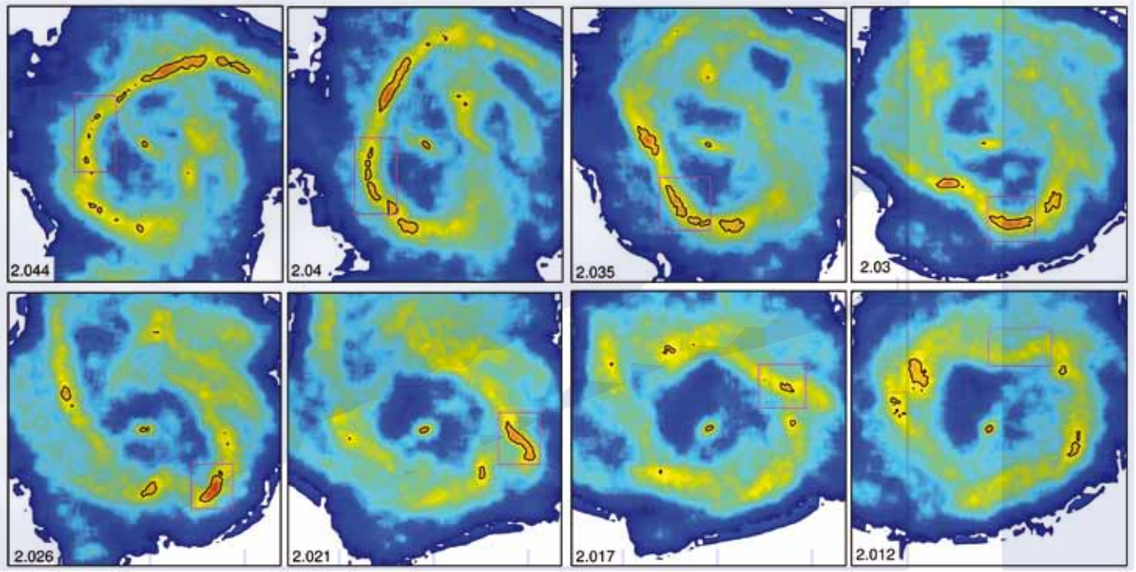

Figure 4: Demonstration of the disruption of a clump

In Figure 3(a) it is seen that several star-forming clumps are embedded in the ring. The clumps, defined here as regions with gas surface density exceeding 1,000 solar masses per square parsec, have Toomre's Q<1 (panel (f)), indicating (marginal) gravitational instability. The clumps' masses are 1-2 orders of magnitude larger than the thermal Jeans mass, being driven to such high values by the large random motions in the disk (Figure 3(b)). The clumps are very gas rich (>~60%). Their stellar populations are comprised of "background stars" whose age distribution is taken from that of the stars in the disk as a whole, and "clump stars" that form over the lifetime of the clump and constitute typically up to ~40% of the clump's stellar mass.

Most notably, the clumps have short lifetimes of about half a disk orbital time, or ~50 million years. The star-formation rate in each clump is typically ~(2-8)% of the total disk star-formation rate. Figure 4 demonstrates the formation and disruption of a clump in s224. The clumps are short-lived because the star-formation in them takes place on a timescale T_{SF}~300Myr, and so the timescale on which the wind drives the gas out of the clumps is approximately T_{SF}/η ~100Myr, where η is the mass loading factor of the wind. This is comparable to the disk orbital time and the gas surface density in the clumps decreases faster than the rate at which it is replenished by the instability inside the disk plane. These outflows from clumps have now been observed for the first time by the IR group at MPE [9].

Several aspects of the dynamics of the disk and clumps can be seen in Figure 3. The vertical velocity dispersion is ~20-100 km/s, in excellent agreement with the values observed at z~2 (Figure 3(b)). Figure 3(e) shows the edge-on velocity field of the disk, which is very regular, like in many observed systems, indicating a quiet merger history. The masses of the most pronounced clumps are >~5×10^8 solar masses and their sizes are about 500 parsecs, therefore their circular velocities equal ~70-100 km/s, which is significantly larger than the velocity dispersion of the gas within them or their rotational velocities.

In Figure 5 we present mock maps of the Hα line, which traces starformation. Shown are line intensity, velocity and velocity dispersion, which are obtained by "placing" galaxy s224 at z=2, convolving it with a FWHM=0.17" resolution, and pasting it onto the real SINFONI datacube of BX502, in order to reproduce realistic resolution and noise properties that correspond to a representative total integration time (~6 hours) for our SINFONI+AO data sets. The clumpiness, smooth velocity field and relatively flat velocity dispersion map (outside the center), which are the characteristics of real SINS clumpy disks, are all reproduced well when "observing" our simulations. The most significant difference to the non-degraded images in Figure 3 (panels (a), (b), (d) and (e) can be directly compared in both figures) is in the velocity dispersion maps. The "beam smearing" increases the apparent velocity dispersion where there are velocity gradients (this is the reason for the diagonal feature in the inclined image in panel (c)), and the clumps are no more seen as clear minima. To compensate for that, the small but present velocity gradients across the clumps get smoothed out. Thus, the combination of the non-virialized state of the clumps and the clumps being minima in velocity dispersion results in that the clumps do not show strong features in observations of the currently available resolution of ~1-2kpc, even if their masses are dynamically significant.

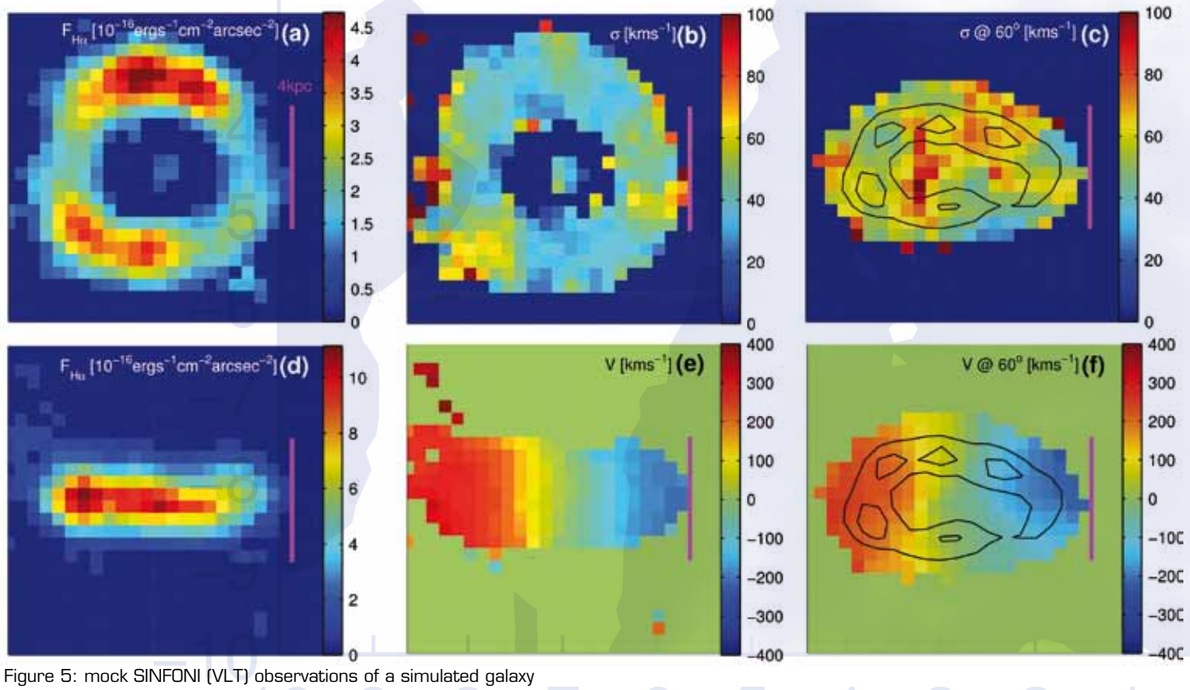

Figure 5: mock SINFONI (VLT) observations of a simulated galaxy

References

[1] Förster-Schreiber, N. M. et al.
SINFONI integral field spectroscopy of z~2 UV-selected galaxies: Rotation curves and dynamical evolution, ApJ, 645, 1062, 2006

[2] Genzel, R. et al.
The rapid formation of a large rotating disk galaxy three billion years after the Big Bang, Nature, 442, 687, 2006

[3] Ceverino, D., Dekel, A. and Bournaud, F.
High-redshift clumpy disks and bulges in cosmological simulations, MNRAS, 404, 2151, 2009

[4] Oser, L., Ostriker, J. P., Naab, T., Johansson, P., and Burkert, A.
The two phases of galaxy formation, ApJ, 725, 2312, 2010

[5] Oppenheimer, B. and Davé, R.
Mass, metal, and energy feedback in cosmological simulations, MNRAS, 387, 577, 2008

[6] Springel, V. and Hernquist, L.
Cosmological smoothed particles hydrodynamics simulations: a hybrid multiphase model for star formation, MNRAS, 339, 289, 2003

[7] Murray, N., Quataert, E., Thompson, T. A.
On the maximum luminosity of galaxies and their central black holes: feedback from momentum-driven winds, ApJ, 618, 569, 2005

[8] Tacconi, L. J. et al.
High molecular gas fractions in normal massive star-forming galaxies in the young Universe, Nature, 463, 7282, 2010

[9] Genzel, R. et al.
The SINS survey of z~2 galaxy kinematics: properties of the giant star forming clumps, arXiv1011.5360G accepted Astrophys. Journal, February 9, 2011

- Thorsten Naab
- Shy Genel
- Reinhard Genzel
- Andreas Burkert
- Peter Johansson
- Ludwig Oser

Max-Planck-Institute for Astrophysics, Garching

Towards High Performance Semantic Web – Experience of the LarKC Project

The essence of Semantic Web is the idea that the Web can exploit techniques from, e.g. formal Knowledge Representation, to make information available in a machine-processable format, so that a more intelligent user support can be achieved on the Web [1]. Such machine-understandable data formats, for instance the Resource Description Framework (RDF), enable novel uses of the Web such as semantic search, data integration, statistical analysis and others. Recent advantages in Semantic Web have forced Web applications to scale up to the requirements of the rapidly increasing amount of interconnected and distributed data such as observed in the Linked Open Data repository for data located across the Web or the Linked Life Data semantic integration platform for the biomedical domain, but also in e-Science and e-Commerce (e.g. Ontoprise).

The massive and tremendously growing amount of data requires effective exploitation and is hence of a great challenge for the modern IT platforms and infrastructures. Another big challenge for achieving the efficiency and web-scaling of Semantic Web applications is the heterogeneous nature of explored data on the Web, which results in data inconsistencies, incompleteness, but also redundancies due to varying methodologies used during the data generation and collection.

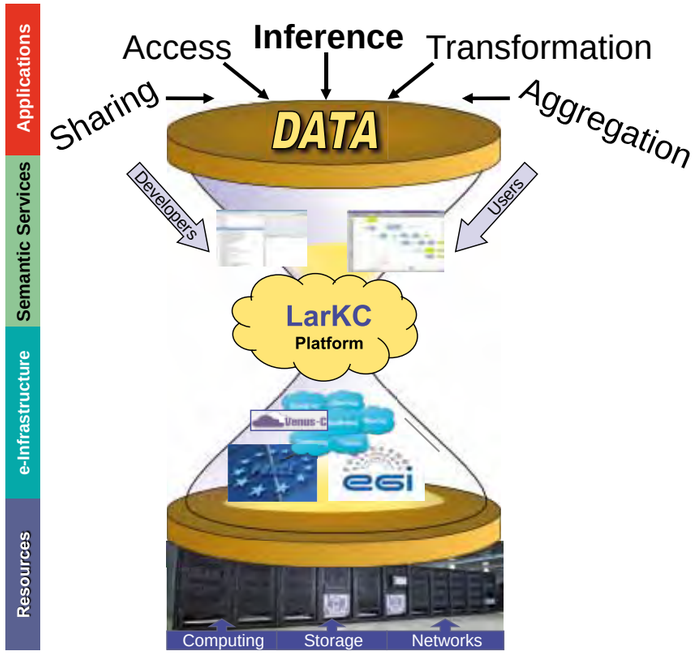

Figure 1: LarKC – a high performance Semantic Web reasoning platform

Given the large problem sizes addressed by Semantic Web and considering the complexity of some data exploration algorithms such as Random Indexing described below, it seems natural to explore the benefits of porting Semantic Web applications for running on High Performance Computing architectures.

Large Knowledge Collider

One of the major practical attempts to build a Semantic Web engine capable of processing billions of structured data, i.e., web-scale data, is performed in the EU FP7 project LarKC (www.larkc.eu). LarKC, which stands for the Large Knowledge Collider, builds an experimental platform for massive distributed incomplete reasoning (see Figure 1), which aims at removing the scalability barriers recognized for most of the currently available reasoning engines. This goal is achieved by means of a number of the original techniques adopted by LarKC, e.g. a highly innovative reasoning approach for merging the retrieval process and the reasoning by means of selection, identification, or transformation [2]. On the other hand, LarKC enables numerous novel IT infrastructure solutions to support those optimization techniques, such as High Performance or Grid Computing e-Infrastructures. The optimal resource provisioning is of especial importance for ensuring the web-scale property of Semantic Web applications. However, since introducing a special e-Infrastructure for Semantic Web, as done in LarKC, processing vast amount of data is not a major bottleneck any more.

Nevertheless, leveraging those resources requires necessary adaptations in the traditional (serial) application codes, i.e. their parallelization. The parallelization becomes thus a major challenge for

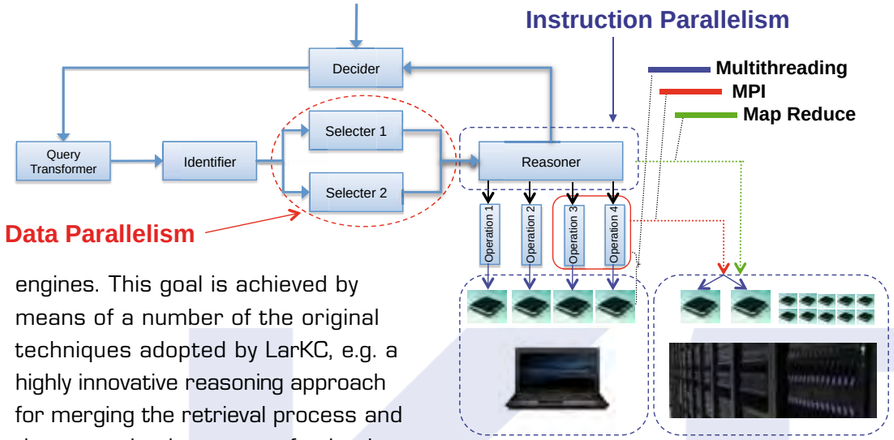

Figure 2: Parallelization in Semantic Web application worklows

the next-generation Semantic Web applications executed in a context of e-Infrastructure.

Parallelization Strategies adopted by LarKC

In solving those issues related to the large-scale Semantic Web applications, LarKC allows a reasoning application to be built on top of numerous lightweight Semantic Web computational blocks (plug-ins, see the actual list on LarKC Market Place at http://www.larkc.eu/plug-in-marketplace), used for identification, selection, transformation, and actual reasoning. When combined in a common workflow, such as one shown in Figure 2, these plug-ins can be efficiently utilized for solving problems of the almost virtual dimensionality.

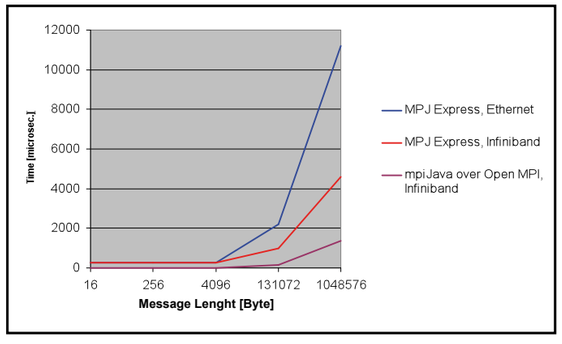

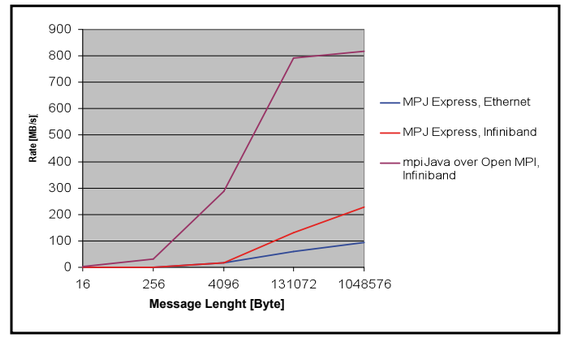

Figure 4: Comparison of time (a) and bandwidth (b) of inter-node communication of different MPI libraries for Java (on the HLRS NEC Nehalem cluster with Ethernet and Infiniband interconnects).

To support this feature, the composition of the plug-ins in a workflow enables parallel execution of the plug-ins on the high performance resources. In terms of LarKC applications, the parallelization suggests the identification of the concurrent regions of the application data - as well as instruction flow, with further mapping them to the independent processor units of a parallel system.

Among the most widely utilized and sustainable in Semantic Web parallelization approaches, such as Multithreading, Map-Reduce, as well as the Message-Passing Interface (MPI), the latter (MPI) is the most promising one in terms of the implementation efforts needed for a serial application as well as in terms of the provided scalability. There have been several initiatives striving to provide HPC support for Java, which is de-facto a default programming language in the LarKC Semantic Web community. One of the most successful MPI implementations for Java has proved to be mpiJava, chosen for adoption in LarKC (Figure 3). The mpiJava framework is currently developed and supported by HLRS.

MPI Realization of Random Indexing

Random Indexing is a novel approach for vector space modelling [3]. The vector space represents the distributional profile of the words in relation to the considered contexts/documents. The main methodical value of this profile is that it enables calculation of the semantic similarity between the words in scope of the document collection (text corpus), based on the cosine similarity function of the given words' context vectors (1).

$$\forall x \in X^{1}, \exists\, v=[f_q(x, c_{j=\overline{1..m}})] \quad (1)$$

where f_q is a co-occurrence function between the word x from the word set X and each of the contexts $c_j \epsilon C^m$, m is a total number of the contexts, n is a total number of the words in all contexts.

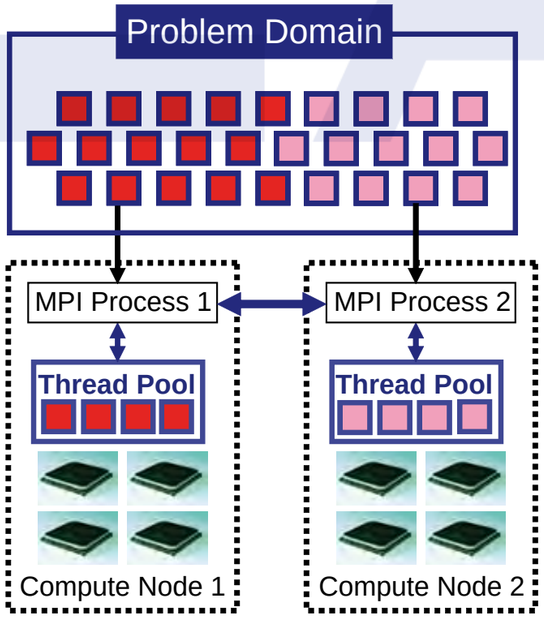

Figure 4: Hybrid MPI + JavaThreads communication pattern.

However, such popular Random Indexing implementation packages, such as Airhead [4], are increasingly ineffective when complexly addressing large data amounts, e.g. as collected by Linked Life Data. LarKC has examined the domain decomposition based parallel implementation of Random Indexing, as depicted in Figure 4.

With regard to the aforementioned Airhead library [5], very promising performance characteristics were obtained for both pure MPI and hybrid MPI-JavaThreads implementations (see Figure 5). The document set based on a selection of the Wikipedia articles (1M high density entries, 16 GB disk space) was used for this performance

Figure 5: Performance characteristics of the parallel Random Indexing realization (a) and comparison of pure MPI vs. MPI + JavaThreads communication performance (b).

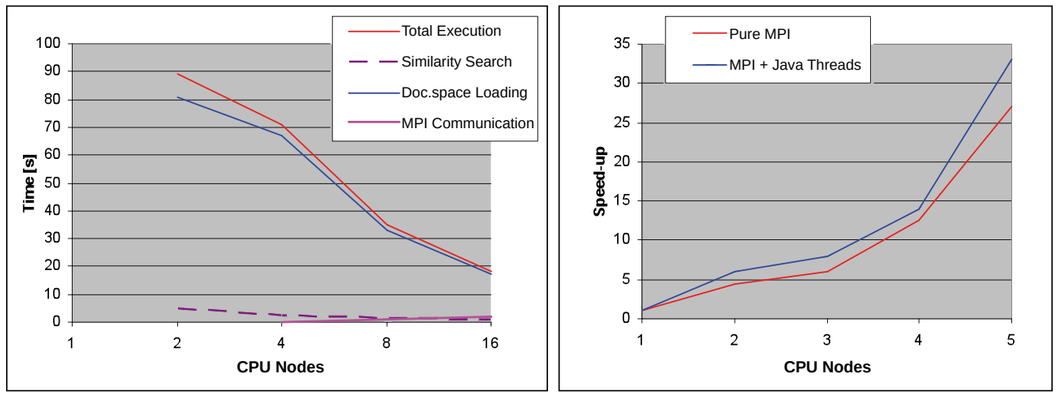

benchmark. Detailed results for different input document sizes as well as cluster configurations are reported in [6].

Outlook

Recent advantages in the Semantic Web require the underlying (Java) applications to scale up to the requirements of the rapidly increasing amount of processed data, such as those coming from millions of sensors and mobile devices, or TB of data volumes conducted during scientific experiments using laboratory equipment. Introducing HPC in Semantic Web domain can greatly support this challenge.

Traditionally, the Semantic Web and the High Performance Computing community have been somewhat disjoint. However, as the needs and capabilities of these two communities continue to converge, it turns to be beneficial for both to leverage their respective technologies. Parallel realization of serial codes is a key enabler of high performance architectures and is therefore a great challenge for the majority of Semantic Web applications. LarKC aims at simplifying development of high performance, parallelized applications, and thus bridging the gap between Semantic Web and HPC.

References

[1] **Daconta, M. C., Smith, K. T., Obrst, L. J.**
The Semantic Web: a Guide to the Future of Xml, Web Services, and Knowledge Management. John Wiley & Sons, Inc., 2003

[2] **Fensel, D., van Harmelen, F.**
Unifying Reasoning and Search to Web Scale. In: IEEE Internet Computing 11 (2), pp. 96-95, 2007

[3] **Sahlgren, M.**
An introduction to random indexing, Methods and Applications of Semantic Indexing Workshop at the 7th International Conference on Terminology and Knowledge Engineering TKE 2005, pp. 1-9, 2005

[4] **Jurgens, D., Stevens, K.**
The S-Space Package, An Open Source Package for Word Space Models. Proceedings of the ACL 2010 System Demonstrations, pp. 30-35, Association for Computational Linguistic, 2010

[5] **Cheptsov, A., Assel, M., Koller, B., Gallizo, G.**
Enabling High Performance Computing for Java Applications using the Message-Passing Interface, P. Iványi and B.H.V. Topping (eds.), Proceedings of the Second International Conference on parallel, distributed, grid and cloud computing for engineering, 2011

[6] **Assel, M., Cheptsov, A., Czink, B., Damljanovic, D., Quesada. J.**
MPI Realization of High Performance Search for Querying Large RDF Graphs using Statistical Semantics. Proceedings of the 1st Workshop on High-Performance Computing for the Semantic Web, collocated with the 8th Extended Semantic Web Conference (ESWC 2011), to appear in 2011

• Alexey Cheptsov
• Matthias Assel

University of Stuttgart, HLRS

Projects

plugIT – Plug Your Business into IT

Business and IT Alignment using a Model-based Plug-in Framework

The alignment of Information Technology and Business is still a highly complex and hard to automate process and remains therefore mainly driven and performed by humans. By nature, the background and knowledge of those humans can differ, depending on their role within their organization. Since different parties often don´t share a common knowledge space, the whole situation is likely to become complex.

plugIT in General

The plugIT project [1] is based on the observation of the necessity to align Business and IT [2] due to the role change of IT from an enabler to an industrial sector. plugIT aims at developing an IT Socket that realizes the vision of "plugging" business into IT in a way similar to the one used to provide electricity via a socket to any device that can be plugged in. This challenge can be taken up by capitalizing on semantic technologies for IT Governance. In plugIT, the Next Generation Modelling Framework [3] is developed which relies on research advances yielding the following benefits:

- A tighter involvement of domain experts is made possible to express formal and semi-formal knowledge via the use of graphical modelling languages
- Different graphical modelling languages for different views and different levels of formal expressiveness can be used
- Domain specific notations for semantics are introduced by merging formal concepts of semantics with graphical notations

An HLRS Use Case

HLRS is one of three use case partners within plugIT. The detailed use case of HLRS covers an Online Proposal Submission process (OPS) in which project applicants can submit a project request to access HLRS computing resources and perform their computational tasks.

Based on the requirements of the project applicant, model translations are used to find the best fitting offer, represented as a set of recommendations for Service Level Agreements (SLAs) [4]. The IT Socket supports the whole process of creating a proposal, analysing the proposal parameters with respect to existing models and finally recommending and generating SLAs. This is enabled thanks to the use of a so-called semantic kernel which uses graphical models combined with semantic information.

In the current production version of the OPS process, HLRS uses a web form based application enabling a project applicant to make requests for computing resources at HLRS. The applicant can enter various pieces of information describing his/her computing resource needs. Once submitted, the request is analyzed by a project approver and approved or, in case modifications are necessary, sent back for updates. So far, all this has been done without any automated supporting processe and has relied heavily on the knowledge of the project approver.

Necessary Enhancements to cover future Developments

Now, with the advent of new paradigms like Cloud Computing, this process needs to be enhanced. Whilst up to now most of the applicants can be assumed to be specialists within their domain, which makes the process simple to manage, the offering of computing resources needs to become more intelligent in the future. In the long term, we need to ensure that also applicants with only moderate knowledge of the underlying system infrastructure need to be able to apply for resources.

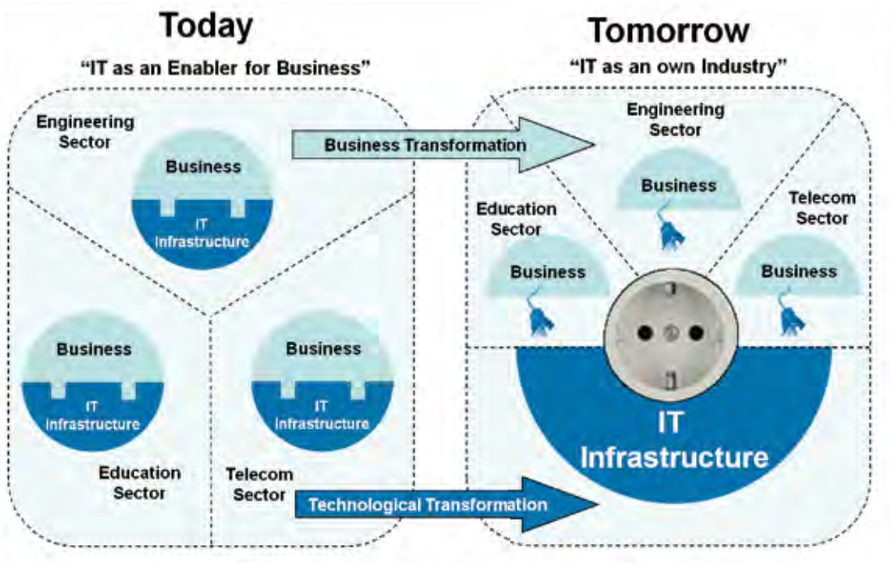

Figure 1: "Plug your Business into IT"

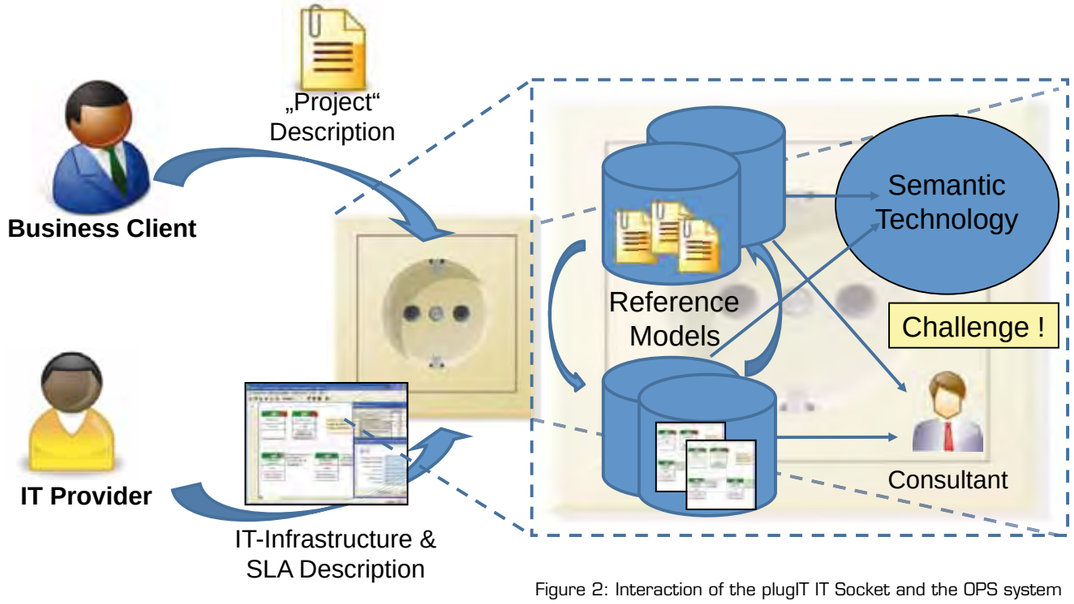

Figure 2: Interaction of the plugIT IT Socket and the OPS system

Moreover, the resources themselves are also getting more and more complex. The acquisition of a new Cray supercomputer at HLRS [5] is just one step towards a new generation of highly complex infrastructures. The role of the project approver thus gets more and more difficult and the benefit of any supporting technology becomes obvious.

The plugIT Enhancements

By means of the plugIT IT Socket, HLRS has concentrated on the provisioning of support for the project approvers. In parallel, the realization of the necessary steps for introducing SLAs into the OPS process have been addressed.

As plugIT follows a model-based approach [6] [7] [8], the first action within the project was to create a number of reference models which were intended to provide the necessary foundations for the project development. This was done via an online model repository containing graphical models related to the IT infrastructure and IT services of HLRS. Modelling was performed with the Next Generation Modelling Framework, one of the deliverables of the project. All the models are linked to each other and they represent the knowledge of HLRS project approvers.

The OPS process is executed now as follows: A project applicant sends a request for computing resources to the HLRS project approver. Upon reception of the request, the approver makes use of the IT Socket to get back from it recommendations on which SLA to offer. This is realized through the automated processing of the project applicant's request by a semantic workflow of the IT Socket. The recommendations are SLA offers that define the category of the SLA that could be proposed to the project applicant. In the current scenario, the categories are bronze, silver or gold, based on the quality of the offer. Each offer relates to a dedicated computing resource. In addition, the recommendations are ranked and visualized in a way simple for the project approver to understand which SLA offer recommendations are the best fitting the project applicant's requirements. The HLRS project approver has the possibility to view the graphical models if he/she requires more information regarding the SLA offer recommendations.

The possibility to plug in the business requirements – the project applicant's request – into the IT domain improves the efficiency and overall performance of the OPS system and allows HLRS to broaden its customer base. The graphical modelling approach also shows obvious advantages in terms of maintenance of information and knowledge transfer.

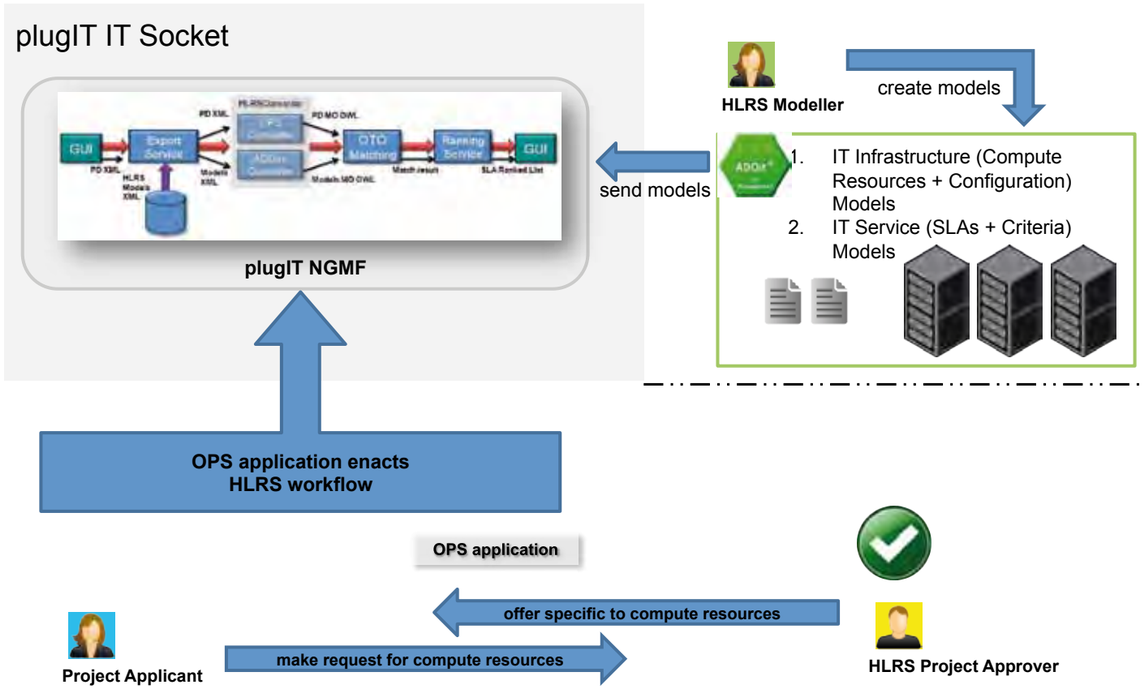

Figure 3: HLRS in plugIT

Facts

plugIT is a project funded by the European Union within the 7th Framework program. The consortium consists of eight project members. plugIT started on the March 1, 2009 and will run until the August 31, 2011.

The Partners

- BOC Asset Management GmbH (AT)
- Telespazio Italia (IT)
- University of Vienna, Department of Knowledge & Business Engineering (AT)
- Foundation for Research and Technology Hellas (GR)
- Fachhochschule Nordwestschweiz (CH)
- CINECA (IT)
- Innovation Technology Group SA (PL)
- University of Stuttgart, HLRS (GER)

Website

http://plug-it.org/

References

[1] plugIT web site, http://plug-it.org

[2] **Woitsch, R., Karagiannis, D., Plexousakis, D., Hinkelmann, K.** Business and IT Alignment: The IT-Socket. E & I Elektrotechnik und Informationstechnik 126, pp. 308-321, 2009

[3] Next Generation Modelling Framework Portal of the 2nd plugIT Prototype: http://83.65.190.84/plugIT/workbench/

[4] **Koller, B.** Enhanced SLA Management in the High Performance Computing Domain, Ph.D. Dissertation, University of Stuttgart, 2011

[5] Cray wins Supercomputer Contract From the University of Stuttgart valued at more than $60 Million, http://investors.cray.com/phoenix.zhtml?c=98390&p=irol-newsArticle&ID=1486975&highlight

[6] **Woitsch, R., Karagiannis, D., Plexousakis, D., Hinkelmann, K.** Plug your Business into IT: Business and IT-Alignment using a Model-based IT-Socket, eChallenges e-2009 Conference, (eChallenges 09), Turkey, IOS Press, 2009

[7] **Bork, D., Sinz, E. J.** Design of a SOM Business Process Modelling Tool based on the ADOxx Meta-modelling Platform. In de Lara, J., Varro, D., Margaria, T., Padberg, J., Taentzer, G., eds.: 4th International Workshop on Graph-based Tools (GraBaTs 2010), Enschede, The Netherlands, pp. 89-101, 2010

[8] **Bézivin, J.** Model Driven Engineering: An Emerging Technical Space, Generative and Transformational Techniques in Software Engineering, Volume 4143, pp. 36-64, 2006

- Axel Tenschert
- Pierre Gilet
- Bastian Koller

University of Stuttgart, HLRS

A scalable Monitoring Tool using Performance Properties

Massive parallel systems are an expensive resource and analyzing the performance of its usage is important to detect and correct bottlenecks. There are several challenges associated with efficient monitoring of large scale parallel systems. The most indispensable ones to tackle are scalability and how to deal with the storage of large quantities of data.

A monitoring system was developed at the Leibniz Supercomputer Centre (LRZ), with successful approaches to face challenges in system monitoring of high performance computers. The monitoring and analysis system PerSyst Monitoring is productive on the Itanium 2 based Altix 4700 platform at LRZ. It is based on hardware counter monitoring and analysis, taking advantage of the rich facility given by the Itanium 2 architecture to provide detailed information on micro-architectural events. The main features of PerSyst Monitoring are:

- Hierarchy of analysis agents
- Hierarchy of performance properties
- Information selection based on thresholds
- Reduction data volume by aggregation.

PerSyst Monitoring provides analysis by gathering performance properties, collectible per job and calculates aggregated information. The closer examination at an application level would be the next step to our monitoring and analysis system by means of instrumentation and the use of a performance tool like Periscope [1], which can use the PerSyst results.

Agent Architecture and Communication

PerSyst Monitoring has been developed as a distributed system with a simple agent hierarchy: A High Level Agent (HLAgent), at the top of the hierarchy tree. The HLAgent communicates to Analysis Agents (AAgents). The agent hierarchy is shown in Figure 1.

Figure 1: Agent hierarchy

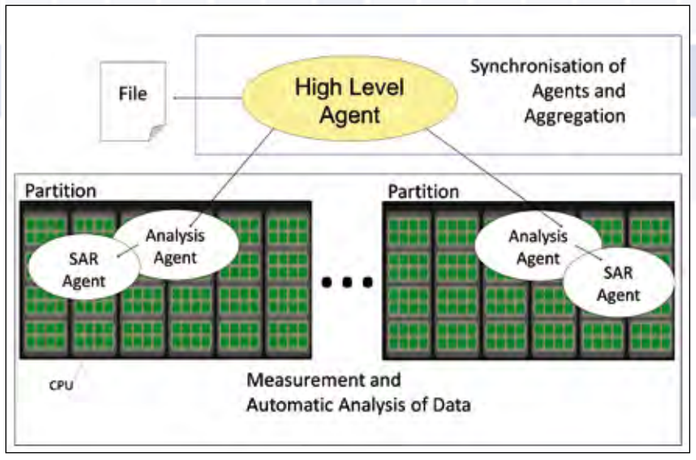

All AAgents are released by the HLAgent when it begins execution. They listen at a given port for instructions from the HLAgent before measuring. The HLAgent itself does not perform any measurements but synchronizes the AAgents and performs data aggregation. An AAgent is in charge of monitoring only one partition of the supercomputer. It can monitor groups of CPUs parallel with the same measurement strategy. The measuring tool used by the AAgents, pfmon [3], collects architecture specific hardware counter samples for the processors.

A simple communication scheme as shown in Figure 2 is used to send small packets of data between the HLAgent and the AAgents. This is done so that the Agents carry out measurements at the same time. Figure 2 also shows one complete cycle of the communication and process scheme. The HLAgent sends the command to start measuring and the agents receive and process. The AAgents, in the ideal scenario, finish the measurement process and send a confirmation that the measurement is done and the HLAgent proceeds to execute the aggregation of information per job. The SAR (system activity report) Agent is released and simply terminates after doing the measurements and writing the result file. Finally the storage is done by the HLAgent and the system is ready for a new monitoring cycle. The time between measurements is configurable. Given the nature of the jobs submitted on the Altix 4700, measurements every 10 minutes are enough to detect inefficiencies in applications.

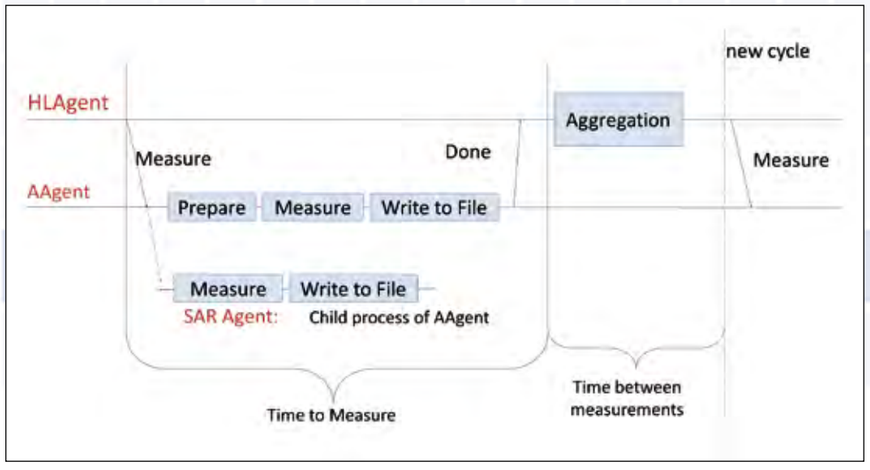

Figure 2: Agent communication scheme

The HLAgent is responsible for the fault tolerance of the rest of individual components of the monitoring system. For this purpose, there are implicit heartbeats with every communication attempt between the HLAgent and the AAgents. There are also cyclical timeouts which are set to all the components. The HLAgent is the only component which would need to be monitored by other tools that probe for continuity of services.

Hierarchical Performance Analysis

A property represents a deficiency or bottleneck, such as an unnecessary idleness of a micro-architectural hardware component. A deficiency can be determined by calculations made to a set of hardware counters. As an example consider the pipeline stalls. The hardware counters indicating a stall in the pipeline can point out very general causes or very specific ones. Every single cause, whether it represents a very general one or a finer-level cause, is computed to produce a property value. For instance, the branch misprediction counter divided by all the cycles of a measured time is the property value for the property value corresponding the branch.

We take the idea of Periscope to differentiate the properties in a hierarchical structure [1,2]. This is what we call a strategy. A strategy has a determined number of properties which in turn

have a determined number of hardware counters involved. Before the measurement is done the AAgent collects all the counters belonging to the strategy in a set and sends it as input to the Performance Script. All the sets of hardware counters involved in a strategy are measured at once or in more measurements via pfmon calls. The AAgent receives the result from the Performance Script and it traverses the strategy tree for each CPU with its corresponding hardware counters to calculate the properties.

Information Reduction

We measure a complete strategy for all the jobs in the supercomputer. From this preliminary information we calculate the property values and severities. The severity is a function of the property value and the threshold. Thus, the first information reduction is not to store the raw hardware counters but the property values. A second important filtering of information is to leave out unnecessary properties. We decided not to store property values whose severity is not indicating a bottleneck. The strategy tree can have several levels of nodes. The absence of an inefficiency problem at a node yields an entire subtree not qualified for further analysis or storage. This also applies to the root node, in case no inefficiency is found the entire strategy tree will not be stored.

The selected and condensed data is written to files or finally to a database. The HLAgent is in charge of reading and processing the files to produce aggregated results over jobs with respect to the property value. This task is easily done by the HLAgent given that jobs mainly run on more than one partition, making it impossible for an AAgent to aggregate for jobs that are distributed over partitions. AAgents would need communication capabilities between themselves in order to aggregate over a job.

Only the aggregated data is stored. The aggregated data consists of meaningful statistics per timestamp in a job, e.g. the average, the minimum and ten percentiles including 50th percentile and 100th percentile. Therefore even for the largest jobs only a definite number of results are stored.

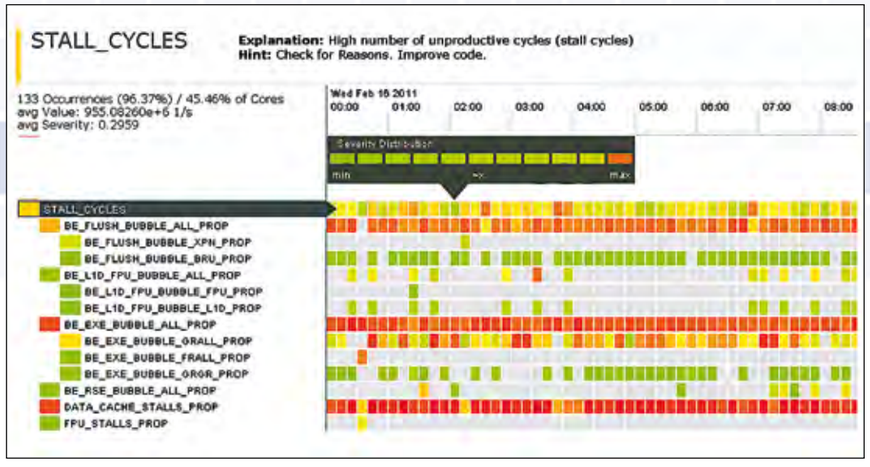

Figure 3: Display of the performance properties of a job. Colour boxes on the left side show the severity of bottlenecks for the whole job, colour boxes in the middle display the severity at a timestamp, and the slider shows the severity distribution at a timestamp.

Results, Discussion and Outlook

Figure 3 shows a display of the stall cycles strategy tree and its properties collected for a job over time. The colour code in the boxes gives the user an idea of how severe the properties are. A purple box has the highest severity average (=1) and green the lowest (=0). Grey boxes are absence of data due to no severity. In the properties shown in Figure 3 we can see that the data cache stalls are the dominating problem in the job.

The reduction strategies produce, on average, about a 90% reduction of the data. Figure 4 shows the aggregate data of a property value plotted over time for a job. For simplicity, not all the collected percentiles are shown in this graph, only the 0th, 50th, 100th percentiles and the average are shown. Due to the wide spectrum of values shown in this graph, it can be readily detected that this application suffers from load imbalance.

The most outstanding advantage of the use of performance properties is that this leads to a better and faster interpretation of results than just reading the raw hardware counters.

Future work will include the porting and testing of PerSyst Monitoring to the new petascale system at LRZ. To achieve this, we plan to develop an interface that allows easy and fast development of further strategies for different processor architectures.

References

[1] **Gerndt, M., Fuerlinger, K.**
Automatic performance analysis with periscope, Concurrency and Computation: Practice and Experience, Wiley InterScience, John Wiley & Sons, Ltd., 2009

[2] **Gerndt, M., Kereku, E.**
Search strategies for automatic performance analysis tools, in: EuroPar 2007, vol. LNCS 4641, pp. 129-138, 2007

[3] **Homepage**
pfmon tool. www.hpl.hp.com/research/linux/perfmon/pfmon.php4

• Carla Guillen
• Wolfram Hesse
• Matthias Brehm

Leibniz Supercomputing Centre (LRZ)

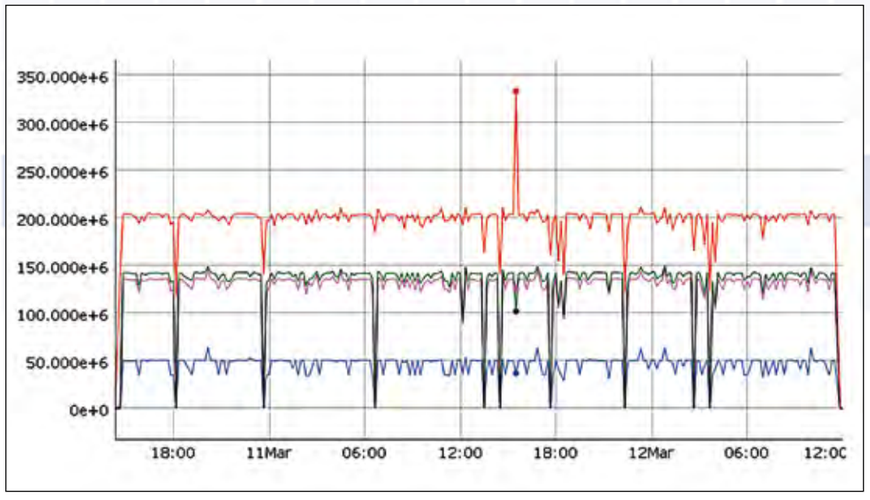

Figure 4: Percentile 0 (blue), 50 (pink), 100 (red) and average (green) of FP_OPS_RETIRED property value plotted over time.

The AstroGrid-D Use Case GEO600: A Breakthrough in Grid Computing

The use case GEO600 was part of the AstroGrid-D project, one of the first five scientific D-Grid community projects. It was contributed by the Max Planck Institute for Gravitational Physics' Albert Einstein Institute (AEI). The goal of the GEO600 (The German-British Gravitational Wave Detector, www.geo600.org) use case was to port the analysis of the gravitational wave data measured with the GEO600 detector near Hannover to the Grid. The program for analyzing these data is Einstein@Home (www.einstein-athome.org). Einstein@Home is based on the BOINC framework (boinc.berkeley.edu).

The principle of the data analysis is as follows. Any user who wants to contribute to the analysis of the gravitational wave data must register at Einstein@Home, download the appropriate Einstein@Home client software to his computer, and then start the client. Whenever the computer idles, the client orders a dataset from the server, and start its analysis on the client computer, and the results are transmitted to the server.

Einstein@Home is an ideal candidate for a Grid application because of multi-platform support, well tested software base, simple resource requirements, build-in checkpoint and recovery methods, fine grained adjustable run time, and linear scaling with node number. The Einstein@Home client program itself was invoked as a black box, but all components which are necessary to bring this use case onto the Grid (the deployment), keep it running in production mode by restarting it after a regular job end and by cleaning recoverable errors, and making job statistics, was developed within the AstroGrid project by the AEI.

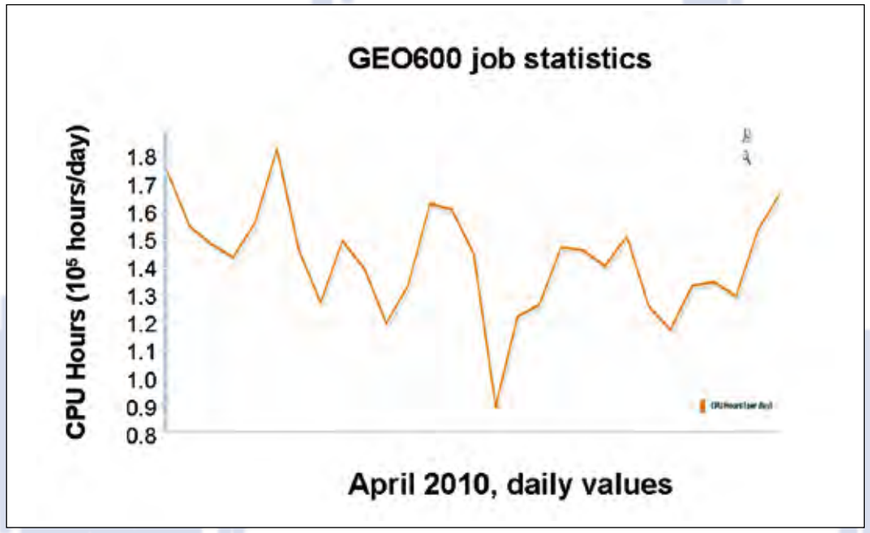

Figure 1: Example for the daily fluctuations of computation time consumed by GEO600.

For the submission of a job to a Grid resource, the Grid middleware Globus is used. The deployment is triggered by a script which is invoked in a Globus job to all Grid machines, where the GEO600 jobs should run. Special software packages which are required by GEO600 are installed automatically during the deployment.

GEO600 itself is started by a Perl script on the submission host, which invokes the submission of one or more Grid tasks to Globus resources. A configuration file enables the setting of certain task submission parameters individual for each Grid resource, e.g. the location of the deployment directory of the GEO600 software on a target machine, the total number of tasks to be submitted to a target resource, the number of tasks which should be submitted at a time, the walltime to be allocated for an Einstein@Home task. All these configurations affect the way, how the Einstein@Home clients should run on the Grid resources, it does not affect the Einstein@Home client itself.

The submission script uses a local MySQL database to control all submitted tasks based on the task identifiers, and to save the exit code when a task has terminated. Depending on the number of currently pending and active tasks and the parameters in the GEO600 resource configuration file, the submission script can automatically determine when to submit new tasks to a Grid resource. To establish a continuous Grid task submission scheme, it is therefore sufficient to invoke the submission script periodically, e.g. as a cron job, on the submission machine.

During the submission and execution of a complex software package like GEO600 to different HPC clusters in the Grid, various Grid related errors can occur which are difficult to track and analyze by hand. The Einstein@Home Client itself does not end up in an error state. An automated handling of typical Grid related errors has been implemented, such that it takes not more than 10 minutes a day to check the job submissions for failures. The GEO600 use case runs now in production mode, and it consumes between 100,000 and 150,000 CPU hours a day on the Globus resources of D-Grid (see Fig. 1).

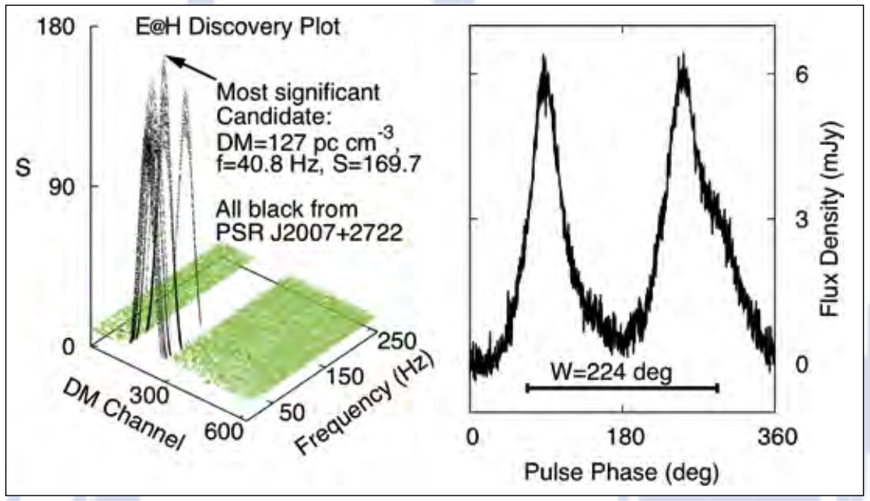

Figure 2:
E@H Pulsar discover plot, from Knispel et al. 2010. Left: significance S as a function of DM and spin frequency (all E@H results for the discovery beam). Right: the pulse profile at 1.5 GHz (GBT). The bar illustrates the extent of the pulse.

Currently Einstein@Home runs on the D-Grid Globus Ressources:

- srvgrid01.offis.uni-oldenburg.de
- udo-gt03.grid.tu-dortmund.de
- lxgt2.lrz-muenchen.de
- emilia.zih.tu-dresden.de
- juggle-glob.fz-juelich.de
- gramd1.d-grid.uni-hannover.de
- gridmon.gwdg.de
- iwrgt4.fzk.de
- gt4-fzk.gridka.de

Two of the HPC clusters that have been used for the data analysis are operated by the GCS partners FZJ and LRZ. In addition to the provision of D-Grid compute resources, both GCS members provide central D-Grid services which are required for the operation of the D-Grid infrastructure and therewith for the submission of jobs to D-Grid compute resources.

The GEO600 use case has not been running properly in production mode in its beginning. In a first approach it was planned to store all the checkpoint files of the Einstein@Home client jobs (the GEO600 tasks) on a central database at the AIP, the so-called "AstroData-Server". However, this approach led to a huge net traffic that ended up in longer times for transferring the checkpoint files from this central database to the execution hosts than the real execution times afterwards. So we decided to store the checkpoint files on the local file systems on the execution hosts.

The "Gridification" of the Einstein@Home brought a substantial breakthrough for Grid computing in Germany. In July 2010, a radio pulsar has been discovered in data recorded with the Arecibo Observatory in Puerto Rico by means of the Einstein@Home project. The data analysis on the D-Grid clusters is worldwide one of the largest contributions to the Einstein@Home project. On the other hand, Einstein@Home is the most successful scientific application of D-Grid. Without the Grid contribution to Einstein@Home, it would not have been possible to analyze sufficiently large amounts data to discover the pulsar. Figure 2 shows the discovery plot of this pulsar, obtained from Knispel et al. 2010.

For the future we are planning to extend the Einstein@Home job submission also to gLite and UNICORE based D-Grid resources. In order to get the job submission control independent of the addressed Grid middleware, larger changes are necessary. We would like to use the Grid Application Toolkit (GAT, see: www.cs.vu.nl/ibis/javagat.html). However, this requires converting the job control mechanism from Perl to Java.

References

[1] **Pulsar Discovery by Global Volunteer Computing: Knispel, B., Allen, B., Cordes, J. M., Deneva, J. S., Anderson, D., Aulbert, C., Bhat, N. D. R., Bock, O., Bogdanov, S., Brazier, A., Camilo, F., Champion, D. J., Chatterjee, S., Crawford, F., Demorest, P. B., Fehrmann, H., Freire, P. C. C., Gonzalez, M. E., Hammer, D., Hessels, J. W. T., Jenet, F. A., Kasian, L., Kaspi, V. M., Kramer, M., Lazarus, P., van Leeuwen, J., Lorimer, D. R., Lyne, A. G., Machenschalk, B., McLaughlin, M. A., Messenger, C., Nice, D. J., Papa, M. A., Pletsch, H. J., Prix, R., Ransom, S. M., Siemens, X., Stairs, I. H., Stappers, B. W., Stovall, K., Venkataraman, A.**
ScienceExpress: www.sciencexpress.org, August 12, 2010, page 1, 10.1126, science.1195253

• Alexander Beck-Ratzka

Max Planck Institute for Gravitational Physics

H4H Project launched

The European ITEA2 project "Hybrid Programming for Heterogeneous Architectures" (H4H), partly funded by the Federal Ministry of Education and Research (BMBF) in Germany, was launched on October 1, 2010. The objective of this project is to provide compute-intensive application developers with a highly efficient hybrid programming environment for heterogeneous computing clusters composed of a mix of classical processors and hardware accelerators such as GPUs.

To meet this challenge, the project will leverage and consistently advance the state of the art in several key software areas: programming models and associated run-time systems, performance measurement and correctness tools, smart translation in particular from OpenMP to CUDA or OpenCL, combined use of MPI and OpenMP, dynamic automatic tuning, and prediction of the execution time of a parallel application on different platforms. To achieve its objective, this project has attracted 25 project partners from France, Germany, Spain and Sweden, including a wide range of HPC users to validate the proposed technology in applications from various domains. For GCS, teams from HLRS and JSC are participating in H4H and contribute their long-term experience in hybrid programming and performance analysis for parallel programs. The project will run for three years and has a total volume of € 15.6 million.

HOPSA Project launched

The new HOPSA project ("HOlistic Performance System Analysis") is funded in the EU 7th framework programme for two years (2011/2012) and as part of a special EU-Russia call, it is also coordinated with a matching Russian project. The objective of this twin project is to create an integrated performance diagnosis infrastructure for a combined system and application tuning. The Russian partners are responsible for the first part, the European for the second. Based on a system-wide basic screening of the performance properties of individual jobs, an automated workflow will route findings on potential performance bottlenecks either to system administrators or application programmers together with recommendations on how to identify their root cause using more powerful diagnostic tools. For this, the European performance tools ThreadSpotter (RogueWave AB, formerly Acumem), Paraver (BSC), Vampir (TU Dresden) and Scalasca (JSC/GRS) will be more tightly integrated. On the Russian side, the HPC computing centres of Moscow State University and the Russian Academy of Science, T-Platforms, and the Southern Federal University in Taganrog are participating in the HOPSA project.

• Bernd Mohr

Jülich Supercomputing Centre

Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities (Leibniz-Rechenzentrum, LRZ) provides comprehensive services to scientific and academic communities by:

- giving general IT services to more than 100,000 university customers in Munich and for the Bavarian Academy of Sciences
- running and managing the powerful communication infrastructure of the Munich Scientific Network (MWN)
- acting as a competence centre for data communication networks
- being a centre for large-scale archiving and backup, and by
- providing High-Performance Computing resources, training and support on the local, regional and national level.

Research in HPC is carried out in collaboration with the distributed, statewide Competence Network for Technical and Scientific High-Performance Computing in Bavaria (KONWIHR).

Contact:
Leibniz Supercomputing Centre

Prof. Dr. Arndt Bode
Boltzmannstr. 1
85478 Garching near Munich
Germany

Phone +49-89-358-31-80 00
bode@lrz.de
www.lrz.de

View of "Höchstleistungsrechner in Bayern HLRB II", an SGI Altix 4700
(Photo: Kai Hamann, produced by gsiCom)

Compute servers currently operated by GCS@LRZ

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
SGI Altix 4700 "HLRB II" Intel IA64 19 x 512-way	9,728 Cores 39 TByte	62.3	Capability Computing	German Universities and Research Institutes, DEISA
Linux-Cluster Intel Xeon EM64T/ AMD Opteron 2-, 4-, 8-, 16-, 32-way	4,438 Cores 9.9 TByte	36.3	Capacity Computing	Bavarian and Munich Universities, D-Grid, LCG Grid
SGI ICE Intel Nehalem 8-way	512 Cores 1.6 TByte	5.2	Capability Computing	Bavarian Universities, PRACE
SGI Altix 4700 & SGI Altix 3700 BX2 Intel IA64 128 & 256-way	384 Cores 1 TByte	2.5	Capability Computing	Bavarian Universities
Linux Cluster Intel IA64 2, 4- and 8-way	220 Cores 1.1 TByte	1.3	Capacity Computing	Bavarian and Munich Universities

A detailed description can be found on LRZ's web pages: www.lrz.de/services/compute

First German National Center
Based on a long tradition in supercomputing at University of Stuttgart, HLRS (Höchstleistungsrechenzentrum Stuttgart) was founded in 1995 as the first German federal Centre for High Performance Computing. HLRS serves researchers at universities and research laboratories in Europe and Germany and their external and industrial partners with high-end computing power for engineering and scientific applications.

Service for Industry
Service provisioning for industry is done together with T-Systems, T-Systems sfr, and Porsche in the public-private joint venture hww (Höchstleistungsrechner für Wissenschaft und Wirtschaft). Through this co-operation industry always has acces to the most recent HPC technology.

Bundling Competencies
In order to bundle service resources in the state of Baden-Württemberg HLRS has teamed up with the Steinbuch Center for Computing of the Karlsruhe Institute of Technology. This collaboration has been implemented in the non-profit organisation HVV.

World Class Research
As one of the largest research centers for HPC HLRS takes a leading role in research. Participation in the German national initiative of excellence makes HLRS an outstanding place in the field.

Contact:
Höchstleistungsrechenzentrum
Stuttgart (HLRS)
Universität Stuttgart

Prof. Dr.-Ing. Dr. hc. Dr. h.c.
Michael M. Resch
Nobelstraße 19
70569 Stuttgart
Germany

Phone +49-711-685-8 72 69
resch@hlrs.de / www.hlrs.de

View of the HLRS Cray XE6 "HERMIT"

Compute servers currently operated by GCS@HLRS

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
Cray XE6 "HERMIT" (1.11.2011)	6500 2-way nodes with 110.000 AMD Interlagos cores	1,050	Capability Computing	European and German Research Organisations and Industry
NEC Hybrid Architecture	12 16-way nodes SX-9 with 8 TByte main memory + 5,600 Intel Nehalem cores 9 TB memory and 64 NVIDIA Tesla S1070	146	Capability Computing	German Universities, Research Institutes and Industry, D-Grid
IBM BW-Grid	3,984 Intel Harpertown cores 8 TByte memory	45.9	Grid Computing	D-Grid Community
Cray XT5m	896 AMD Shanghai cores 1.8 TByte memory	9	Technical Computing	BW Users and Industry

A detailed description can be found on HLRS's web pages: www.hlrs.de/systems

View of the HLRS BW-Grid IBM Cluster (Photo: HLRS)

JÜLICH
FORSCHUNGSZENTRUM

The Jülich Supercomputing Centre (JSC) at Forschungszentrum Jülich enables scientists and engineers to solve grand challenge problems of high complexity in science and engineering in collaborative infrastructures by means of supercomputing and Grid technologies.

Provision of supercomputer resources of the highest performance class for projects in science, research and industry in the fields of modeling and computer simulation including their methods. The selection of the projects is performed by an international peer-review procedure implemented by the John von Neumann Institute for Computing (NIC), a joint foundation of Forschungszentrum Jülich, Deutsches Elektronen-Synchrotron DESY, and GSI Helmholtzzentrum für Schwerionenforschung.

Supercomputer-oriented research and development in selected fields of physics and other natural sciences by research groups of competence in supercomputing applications.

Implementation of strategic support infrastructures including community-oriented simulation laboratories and cross-sectional groups on mathematical methods and algorithms and parallel performance tools, enabling the effective usage of the supercomputer resources.

Higher education
for master and doctoral students in cooperation e.g. with the German Research School for Simulation Sciences.

Contact:
Jülich Supercomputing Centre (JSC)
Forschungszentrum Jülich

Prof. Dr. Dr. Thomas Lippert
52425 Jülich
Germany

Phone +49-24 61-61-64 02
th.lippert@fz-juelich.de
www.fz-juelich.de/jsc

View on the supercomputers JUGENE, JUST (storage cluster), HPC-FF and JUROPA in Jülich (Photo: Research Centre Jülich)

Compute servers currently operated by GCS@JSC

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
IBM Blue Gene/P "JUGENE"	72 racks 73,728 nodes 294,912 processors PowerPC 450 144 TByte memory	1,002.6	Capability computing	European Universities and Research Institutes, PRACE
Intel Linux CLuster "JUROPA"	2,208 SMT nodes with 2 Intel Nehalem-EP quad-core 2.93 GHz processors each 17,664 cores 52 TByte memory	207	Capacity and Capability Computing	European Universities, Research Institutes and Industry, PRACE
Intel Linux CLuster "HPC-FF"	1,080 SMT nodes with 2 Intel Nehalem-EP quad-core 2.93 GHz processors each 8,640 cores 25 TByte memory	101	Capacity and Capability Computing	EU Fusion Community
IBM Cell System "QPACE"	1,024 PowerXCell 8i processors 4 TByte memory	100	Capability Computing	QCD applications SFB TR55, PRACE
Intel GPU Cluster "JUDGE"	54 nodes with 2 Intel Westmere 6-core 2.66 GHz processors each 108 graphic processors (NVIDIA Fermi) 5.1 TByte memory	62.5	Capacity and Capability Computing	selected HGF Projects

ITEA Gold Award for ParMA

GCS is happy to announce that the European ITEA2 research project ParMA (Parallel Programming for Multicore Architectures) received the "ITEA Achievement Award 2010 in Gold". With this award, ITEA rewards high-level technical contributions based on European collaborations providing significant results while promoting ITEA and its aims. The ITEA2 Board was impressed by the innovation and fast exploitation in ParMA which developed advanced technologies to exploit multicore architectures for High Performance Computing. For GCS, teams from HLRS and JSC are participating in ParMA. HLRS was work package leader and acted as German coordinator for the project. An important contribution to ParMa was the UNITE development tool package designed and implemented by JSC which includes a full set of interoperable tools for advanced debugging and performance analysis including the JSC tool Scalasca.

Understanding the Formation of Wait States in parallel Programs

• Bernd Mohr

Jülich Supercomputing Centre

With today's supercomputers featuring tens of thousands of cores, writing efficient codes that exploit all the available parallelism becomes increasingly difficult. Load and communication imbalance, which frequently occurs during simulations of irregular and dynamic domains – typical of many engineering codes – presents a key challenge to achieving satisfactory scalability. Even delays of single processes may spread wait states across the entire machine, and their accumulated duration can constitute a substantial fraction of the overall resource consumption. In general, wait states may propagate across process boundaries along far-reaching cause-effect chains before they materialize at a synchronization point much later in the program.

To better understand how the effects of such imbalanced behaviour slow down program execution, David Böhme, a Ph.D. candidate at the Jülich Supercomputing Centre, and his colleagues developed a scalable technique that analyzes the formation of wait states and attributes their costs in terms of resource waste to their original cause. Building on earlier work by Meira, Jr. et al. [1], the researchers take execution traces of parallel programs and replay the recorded communication. A first replay in forward direction identifies wait states and measures their duration. A second replay, performed in backward direction, traces these wait states back to the imbalance responsible for their occurrence, letting their costs travel along the reversed cause-and-effect chain until they can eventually be mapped onto their root cause. Since the replay occurs in parallel, it was possible to demonstrate the new approach with up to 65,536 processes. An article describing the idea along with experimental results [2] won the Best Paper Award at the International Conference on Parallel Processing (ICPP) 2010 in San Diego, California. To allow more target-oriented optimizations of imbalance phenomena in the daily practice of application tuning at HPC centres, the new method is currently being integrated into Scalasca [3] (Fig. 1), a performance-analysis tool developed at the Jülich Supercomputing Centre and the German Research School for Simulation Sciences in Aachen.

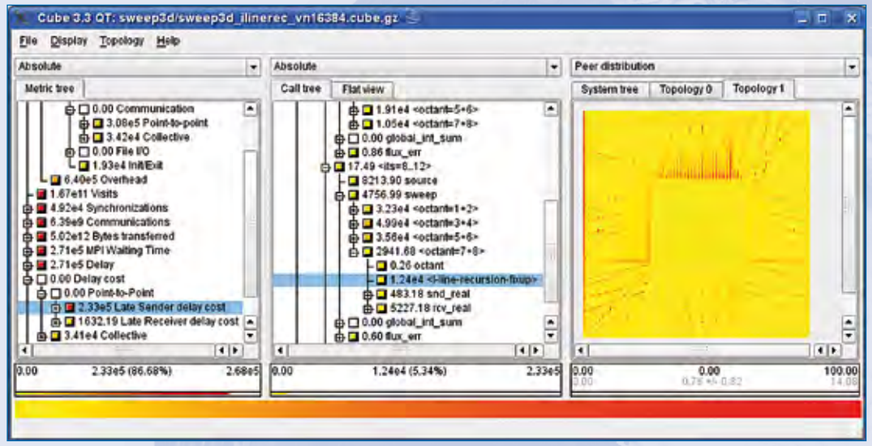

Figure 1: Using the new technique, future versions of the performance-analysis software Scalasca will allow a precise analysis of the sources and costs of wait states that occur in the wake of load imbalance.

David Böhme's dissertation project is funded through a scholarship from the Aachen Institute for Advanced Study in Computational Engineering Science (AICES), a graduate school at RWTH Aachen University established in November 2006 under the auspices of the Excellence Initiative of the German state and federal governments.

References

[1] **Wagner Meira, Jr., LeBlanc, T. J., Virgílio, A. F.** Almeida: Using cause-effect analysis to understand the performance of distributed programs, in: Proc. of the SIGMETRICS Symposium on Parallel and Distributed Tools (SPDT'98), Welches, OR, USA, pp. 101-111, ACM, August 1998

[2] **Böhme, D., Geimer, M., Wolf, F., Arnold, L.** Identifying the root causes of wait states in large-scale parallel applications, in: Proc. of the 39th International Conference on Parallel Processing (ICPP), San Diego, CA, USA, pp. 90-100, IEEE Computer Society, September 2010

[3] Scalasca: www.scalasca.org

• Felix Wolf

German Reseach School for Simulation Sciences, Aachen

From Computational Biophysics to Systems Biology 2011 (CBSB11)

• Walter Nadler

Jülich Supercomputing Centre

Since 2006 this series of successful workshops has been held in both Germany and the USA. Its goal is to bring together scientists from various scientific fields to bridge the gap between biological simulations at the molecular level and approaches to describe biological systems at the cellular level and beyond.

This year's fifth workshop will take place from 20 to 22 July 2011, again at the Research Centre Jülich, jointly organized by the Jülich Supercomputing Centre (JSC), the German Research School for Simulation Sciences (GRS), and Michigan Technological University (MTU). It is dedicated to Harold Scheraga who will celebrate his 90th birthday later this year. Dr. Scheraga pioneered the use of computers in chemistry and biology. His work inspired many of the research areas that are the topic of this meeting. In the spirit of Harold Scheraga's work, this workshop will bring together researchers from physics, chemistry, biology, and computer science to acquaint each other with current trends in computational biophysics and systems biology, to explore avenues of cooperation, and to establish together a detailed understanding of cells at a molecular level. Its main focus will be on:

- Protein folding and aggregation,
- Multi-protein complexes and supramolecular assemblies,
- Cellular environments and interaction networks,

and on Models, Algorithms, and High Performance Computing. The workshop will act as a platform for in-depth discussion of cutting-edge research results obtained by international scientists at all levels in their careers. Invited talks will highlight recent algorithmic developments in and successful applications of HPC to life sciences. A limited number of contributed talks and a poster session will allow participants at an early stage of their career to discuss their ongoing research and to put it into the context of the workshop. Finally, several panel sessions will stimulate the exchange of views between the various scientific fields and the different approaches to understanding the biological systems in question.

For the full program see **http://www2.fz-juelich.de/conference/cbsb11**; there also links to the programs and proceedings of the earlier workshops can be found.

Publications in Computational Science and Engineering

Krause, E., Shokin, Y., Resch, M., Kröner, D., Shokina, N. (Eds.), Computational Science and High Performance Computing IV, The 4th Russian-German Advanced Research Workshop, Freiburg, Germany, October 12-16, 2009. Series: Notes on Numerical Fluid Mechanics and Multidisciplinary Design (NNFM) , Vol. 115, Springer, Berlin – Heidelberg – New York, 2011.

This volume contains 27 contributions to the 4th Russian-German Advanced Research Workshop on Computational Science and HPC presented in October 2009 in Freiburg, Germany. The workshop was organized jointly by the HLRS, the Institute of Computational Technologies of the Siberian Branch of the Russian Academy of Sciences (ICT SB RAS) and the Section of Applied Mathematics of the University of Freiburg (IAM Freiburg). The contributions range from computer science, mathematics and HPC to applications in mechanical and aerospace engineering. They show a wealth of theoretical work and simulation experience with a potential of bringing together theoretical mathematical modelling and usage of high performance computing systems presenting the state of the art of computational technologies.

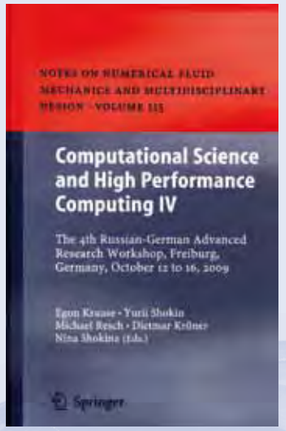

Dimitrakos, T., Martrat, J., Wesner, S. (Eds.), Service-oriented Infrastructures and Cloud Service Platforms for the Enterprise, Springer, Berlin - Heidelberg, 2010.

Service-oriented Infrastructures including Grid and Cloud Computing are technologies in a critical transition to wider adoption by business. Their use may enable enterprises to achieve optimal IT utilization, including sharing-resources and services across enterprises and on-demand utilization of those made available by business partners over the network. This book is an essential reference for researchers and practitioners in service-oriented IT. It analyses a selection of common capabilities (services capturing reusable functionality of IT solutions) that have been applied to tackle challenging business problems and were validated by the BEinGRID consortium in real-life business trials covering most European market sectors.

New 100 Gigabit/s Ethernet Technology successfully tested between FZJ and KIT

In a joint effort, the German Research Network Provider DFN, the Jülich Supercomputing Centre (JSC) at the Forschungszentrum Jülich (FZJ) and the Karlsruhe Institute of Technology (KIT) as well as the industrial partners Cisco Systems, Huawei Technology and GasLINE tested the newly developed 100 Gigabit/s Ethernet technology on a 447 km connection between JSC and KIT.

Facing growing demand for communication bandwidth e.g. in the area of Telecommunication Provider Backbones or Supercomputer I/O, in summer 2010 the IETF and IEEE finalized the IEEE802.3ba standard for 100 GBit/s Ethernet. This standard especially addresses today's limiting factors for high bandwidth communication: constraints concerning the number of parallel links that can be combined into a single logical link (IEEE802.3ad link aggregation) and the number of wavelengths that can be transmitted in parallel over a single fibre with DWDM (dense wavelength division multiplexing).

Figure 1: Huawei DWDM-transmission system with 100GE wavelength conversion board

After several weeks of preparation and the installation of multiple racks of network equipment, partly still in beta-status, the tests between JSC and KIT started in autumn 2010 via a 447 km fibre path with 6 optical regenerators provided by GasLINE. The optical layer consisted of a Huawei DWDM transmission system with 100 Gbit/s optical transceivers (so called 100GBase-LR4 CFPs) at the customer side interface. While during the setup phase some incompatibility issues between different CFP-vendors were found on the optical layer, these problems could be solved by CFP exchanges and rearrangements. During the 3 weeks lasting first test-phase a jointly administered carrier-class router Cisco CRS-3 per site aggregated up to twelve 10 GBit/s Ethernet-streams generated by test nodes. The tests with UDP- and TCP-based traffic patterns were accompanied by continuous monitoring of traffic statistics and the measurement of packet loss, delay and variation via additional dedicated measuring systems. We could show that a sustained throughput of 99.6 GBit/s was achievable via this long range 100 GBit/s Ethernet connection. In a second, 10 days lasting test-phase, a similar test scenario was set up with Huawei NE5000/NE40E routers. With intensive support of Huawei engineers the Huawei routers were able to drive the 100 GBit/s Ethernet connection close to 100% utilization, too.

Figure 2: Cisco CRS-3 and Huawei NE5000

Figure 3: Linux-based test nodes at the JSC

As a result of months of preparation and testing it turned out that the new 100 GBit/s Ethernet standard is already well implemented and supported by network components (routers, DWDM transmission systems) of major vendors and ready to be deployed in real-world production networks.

• Olaf Mextorf

Jülich Supercomputing Centre

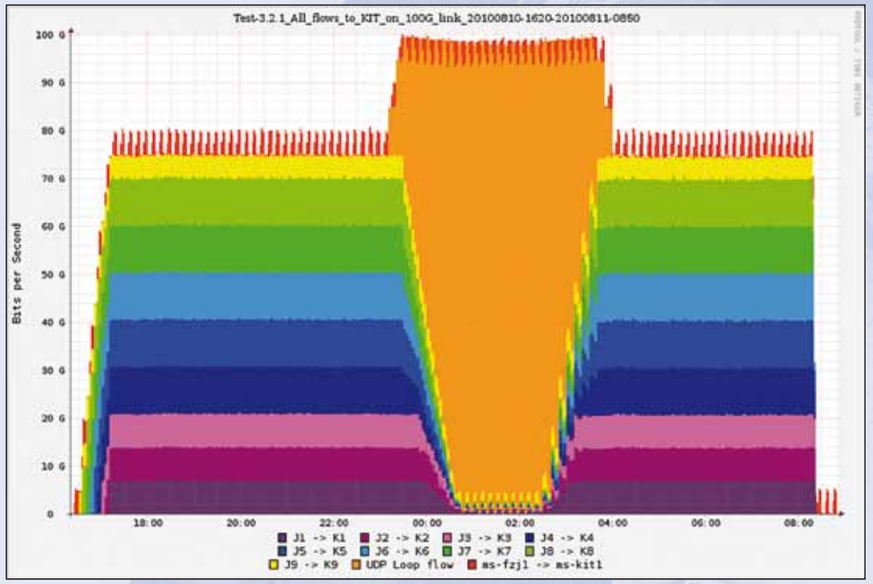

Figure 4: UDP traffic forces multiple TCP-streams on the 100GE-link to adapt and slowdown

10th HLRS/hww Workshop on Scalable Global Parallel File Systems

Representatives from science and industry interested in high performance storage solutions did meet at HLRS from May 2-4, 2011 for the tenth annual HLRS/hww Workshop on Scalable Global Parallel File Systems. More than 100 participants did follow a total of 26 presentations that have been on the workshop agenda. In addition, a panel discussion on Monday evening illuminated the issues when preparing file systems for the exabyte era.

The Director of HLRS, Prof. Michael Resch, opened the workshop with an opening address on Monday morning. Dr. Peter Braam, SVP Software at Xyratex, delivered a keynote speech on Bridging the Peta- to Exascale I/O Gap. He presented the foreseeable issues with exascale file systems from today's perspective and he proposed seven I/O dwarfs. These dwarfs are I/O benchmark kernels which are as small as possible but broad enough to cover the essential points. In the second talk, Björn Andersson, Director HPC Product Marketing of BlueArc, explained how BlueArc's pNFS solution has been architected. An introduction to the OneFS clustered file system was given by David Vitera, Global Systems Engineer, EMC and Isilon. He described the internal architecture especially pointing out the advantages of the different design decisions.

The roadmap to Exascale for Lustre was the topic of the first afternoon talk given by Eric Barton, Whamcloud. He explained in great detail the intended changes in the Lustre file system to enable it for the era of exabyte file systems. Thomas Uhl, GrauData and Rising Tide Systems, showed

Scalable Open Source Storage. He explained how the different open source solutions could work together to form a high-performance cluster file system with direct HSM access at scale.

The first day was concluded by a panel discussion of the speakers moderated by Peter Haas, HLRS. The discussion about the way to exascale for storage showed that the hardware is considered the biggest problem. Co-designing storage hardware and file system software was seen as a solution for this. In addition, programmers will see shifts in the paradigm how data is accessed from applications.

Tuesday morning saw a series of presentations covering the remaining global parallel file systems. Kalyan Gunda, IBM, explained how GPFS information lifecycle management works. He also explained how data can be shared conveniently between clusters in the future.

Brent Welch, Panasas, provided insight into the pNFS usage with the Panasas storage systems and Sven Breuner provided an update to the Fraunhofer File System FhGFS, showing the current developments and future plans. Talks from DataDirect Networks, Crossroads and FusionIO covered the Hardware aspects of storage systems.

In the research part of the workshop, speakers from TU Dresden, T-Systems and Fraunhofer presented

the setup of the 100GBit test bed between Dresden and Freiberg enriched with measurement data from different wide area file system setups. Speakers from CEA, France, provided information about their developments, a policy engine named Robinhood and an open NFSv4.x server.

On Tuesday afternoon, the High Performance Networking Forum (HNF) Europe provided talks about the evolution of the datacenter network and Cisco's view to the HPC market and the concluding product designs. In the parallel OpenFabrics track, high performance storage solutions based on Infiniband networks have been discussed.

Tuesday night's agenda started with an invitation to closely study ancient storage solutions in the medieval town center of Esslingen. In one of the old storage spaces, the products of Germanys oldest sparkling wine maker could be tested. From there, the bus moved on to the Stuttgart Museum of Fine Arts for a very memorable dinner at the Cube Restaurant on fourth floor.

The presentation on Wednesday morning showed the data management solutions in the European

Projects DEISA and PRACE. Bastian Koller, HLRS, explained Service Level Agreements and the role they might play in future storage. As an outlook in the networking area, Klaus Grobe from Adva took a look to the time after 100GBit connections. And it turned out that doing multicore is not only a solution in processor technology but it will also play a role in fibre optics.

HLRS appreciates the great interest it has received from the participants of this workshop and gratefully acknowledges the encouragement and support of the sponsors who have made this event possible.

Sponsored by:

• Thomas Boenisch

University of Stuttgart, HLRS

Blue Gene Extreme Scaling Workshop 2011

From February 14 to 16, JSC organized its 2011 Blue Gene Extreme Scaling Workshop. Similar to the previous workshops in 2009 and 2010, the main focus was on application codes able to be scaled up during the workshop to the full Blue Gene/P system JUGENE, which consists of 72 racks with a total of 294,912 cores – still the highest number of cores available worldwide in a single system.

Interested application teams had to submit short proposals which were evaluated. Selection criteria were the required extreme scaling, the application-related constraints which had to be fulfilled by the JUGENE software infrastructure, and the scientific impact that the codes could produce. The process was very competitive: 8 out of the submitted 15 high-quality proposals were selected. Participating teams came from the Princeton Plasma Physics Laboratory in the US, the Royal Institute of Technology (KTH) in Sweden, the King Abdullah University of Science and Technology (KAUST) in Saudia Arabia, the Mickiewicz University in Poland, University College London in UK, the University of the Basque Country in Spain, RZG MPI, and the University of Heidelberg.

During the workshop, the teams were supported by JSC parallel application experts, the JUGENE system administrators and two IBM MPI and compiler experts. In addition, the participants shared their expertise and knowledge. The workshop was extremely successful: the 8 teams succeeded in submitting one or more successful full 72-rack jobs for 11(!) different applications during the course of the workshop as some teams experimented with more than one code. A total of 308 jobs were launched using 122 rack days of the total 157 rack days reserved for the workshop.

Many interesting results were achieved: The team of the University of the Basque Country uses the time-dependent density functional theory code Octopus to perform first principle simulations of the excited states of chlorophyll molecules combined with other chromophores and proteins which form the different complexes that take part in photosynthetic processes. This type of simulations has a direct impact on our understanding of photo-induced processes in biological systems. The large size of these molecules, of the order of thousands or tens of thousands of atoms, makes them very challenging to model computationally. During the workshop, they performed runs simulating the absorption of light by these large molecular systems. For the first time, they were able to simulate molecules with 2,676 and 5,879 atoms (the whole chlorophyll molecule) using all 72 racks of the machine.

Researchers from the Royal Institute of Technology (KTH) in Sweden investigated the scaling of large-scale neural simulations using an experimental neural simulator (ANSCore) and an experimental code simulating a model of the neocortical network of the brain (BrainCore). Both codes successfully scaled up to the full system scale when applied to large-scale neural simulations. Scaling of core components for building a variety of neural models was featured in ANSCore. There, communication was handled by MPI collective calls while the BrainCore model featured a large homogeneous single network and a straight-forward application in terms of associative memory using point-to-point communication. The results achieved open the path for simulations of neural models of sizes comparable to real mammalian nervous systems with a much higher complexity than so far attempted. They will allow to handle spiking communication in models of the neocortical network as these scale up to sizes of real mammalian brains. Furthermore, their study paves the way for the use of extremely scalable brain network models for information processing of data obtained with e.g., large-scale sensor arrays. The knowledge gained can also be used to investigate the design of dedicated hardware.

Finally, the group of KAUST investigated the scaling of their Billions-Body Molecular Dynamics (BBMD) code which is a highly optimized, parallel C++ MD code for Lennard-Jones particle systems. The code is used to study for the first time the behaviour of large-size structured glasses characterized by tens of billions of particles. This will permit to answer long-standing questions in the field of complex systems concerning the existence and the dynamics of specific wave vibrations of structurally disordered systems like glasses. The BBMD code was successfully run on all 72 racks of the Blue Gene/P system for experiments with 10 billion particles. The scaling from 32,786 cores (8 racks of BG/P) to 294,912 cores (72 racks of BG/P) shows an efficiency of 89%. The amount of communication compared to elapsed time can be seen to be less than 2% for 8 racks and only grows to 7% for 72 racks. This shows that utilizing mainly nearest neighbour communication, with a very limited amount of global communication, is very advantageous for extreme scaling.

All experience gathered during the workshop will be summarized in a technical report. For more information on the workshop as well as on the reports of the previous workshops see:

http://www2.fz-juelich.de/jsc/bg-ws11

• Bernd Mohr

Jülich Supercomputing Centre

LRZ-Workshop 2011 Scientific 3D Animation with Blender

The two-day course gave an introduction to the visualisation of scientific simulation data using the open source 3D animation package Blender (www.blender.org). 27 participants learned how to generate impressive and professional looking still images and animations of the data obtained from their own scientific projects.

The central idea of the workshop is that the scientists get their 3D models from experimental data (e.g. the molecular structure of a molecule) or from computer simulations (e.g. computational fluid dynamics). Apart from analysis software that produces the necessary diagrams and plots for a scientific discussion, there is increasing demand for high quality and artistic renderings of the 3D models.

In the hands-on part of the workshop, the participants learned all the relevant steps to create stunning 3D animations. Based on a real-world example project using a protein molecule, all steps necessary to produce a 3D animation were discussed: importing the data, cleaning-up the geometry, assigning materials, illuminating the scene, setting key frames, rendering and post-production. Since one of the presenters works in the core development team of Blender, even advanced topics such as python scripting could be discussed in-depth. At the end of the workshop, each participant left with his own 3D animation of an enzyme.

A DVD with all the slides and data necessary to follow the workshop at home is available from the LRZ.

• Helmut Satzger

Leibniz Supercomputing Centre

LRZ-Workshop 2011 Molecular Modelling on Supercomputers

The Leibniz Supercomputing Centre held a three day workshop on molecular modelling on supercomputers on January 24-26 in Garching with 12 participants from the life science community. The objective of the workshop was to help life science researchers approach challenging problems in computational chemistry and biology. The workshop gave an introduction and hands-on sessions to five major software packages for molecular modelling. Day one was held by Schrödinger Inc. who presented advanced features of the Schrödinger suite of programs with a special focus on homology modeling. Each participant was provided with a laptop for the extensive hands-on session by Schrödinger Inc. that had the software and data files pre-installed. On day two, LRZ gave hands-on sessions to CPMD, VMD and NAMD. The third day was organized by the EU-funded project Scalalife (www.scalalife.eu, where LRZ is a partner). The aim of the project is to provide scalable software services for the life sciences in High Performance Computing. Two Scalalife developers from GROMACS and DALTON, both from KTH in Sweden, gave an introduction and hands-on sessions to their software packages. There were interesting discussions how to bridge the different important scales from single molecule first principle quantum mechanics over force-field calculations up to coarse-grained molecular dynamics. The presentations covered many interesting topics, such as ligand binding, visualization of biomolecules and trajectories, and automated generation of configuration files for simulation runs as well as the job submission on the supercomputers of the LRZ.

• Helmut Satzger

Leibniz Supercomputing Centre

HLRS Scientific Tutorials and Workshop Report and Outlook

• Rolf Rabenseifner

University of Stuttgart, HLRS

In December 2010, HLRS has installed a first Cray XE6 system with a performance of 10 TFlop/s. The final XE6 system with a performance of 1 PFlop/s is expected to be operational in Q3/2011. We strongly encourage you to use the already existing small system to port your applications to the new architecture as early as possible. To support such effort we invite current and future users to participate in special **Cray XE6 Optimization Workshops**. With these courses, we will give all necessary information to move applications from the current NEC SX-9, the Nehalem cluster, or other systems to the upcoming Petaflop system. The Cray XE6 will provide our users with a new level of performance. To harvest this potential will require all our efforts. We are looking forward to working with our users on these opportunities. These 1-3 day courses in cooperation with Cray were at Feb. 2-4 and Apr. 19. Next courses will be on Sep. 6-7, and Nov. 2-4.

Additionally, the Virtual Institute for High-Productivity Supercomputing (VI-HPS) offered a workshop for tuning and code optimization specifically on our Cray XE6 pre-installation system. This **7th VI-HPS Tuning Workshop** took place at HLRS on March 28-30.

One of the flagships of our courses is the week on **Iterative Solvers and Parallelization**. Prof. A. Meister and Prof. B. Fischer teach basics and details on Krylov Subspace Methods. Lecturers from HLRS give lessons on distributed memory parallelization with the Message Passing Interface (MPI) and shared memory multi-threading with OpenMP. This course was presented in March 2011 at HLRS in Stuttgart and will take place on September 12-16, 2011 at LRZ in Garching.

Another highlight is the **Introduction to Computational Fluid Dynamics**. This course was initiated at HLRS by Dr.-Ing. Sabine Roller. She is now a professor at the German Research School at RWTH Aachen, and with this, the course was held the first time there. In April 2011, it was presented again in Stuttgart, and it will come back to Aachen on September 26-30, 2011. The emphasis is placed on explicit finite volume methods for the compressible Euler equations. Moreover outlooks on implicit methods, the extension to the Navier-Stokes equations and turbulence modelling are given. Additional topics are classical numerical methods for the solution of the incompressible Navier-Stokes equations, aero-acoustics and high order numerical methods for the solution of systems of partial differential equations.

Tutorials at Scientific Conferences in 2011

Hybrid Message Passing + Shared Memory Programming. Rolf Rabenseifner, Pekka Manninen, et al.; lecture with practical, at PRACE Summer School: Taking the Most out of Supercomputers, CSC, Espoo, Finland, August 29 – September 1, 2011.

Hybrid MPI and OpenMP Parallel Programming. Gabriele Jost, et al.; at DEISA/PRACE Spring School: Tools and Techniques for Extreme Scalability, EPCC, Edinburgh, Scotland (GB), March 29-31, 2011.

Performance-Oriented Programming on Multicore-based Clusters with MPI, OpenMP & Hybrid MPI/OpenMP. Georg Hager, Gabriele Jost, Jan Treibig, Gerhard Wellein. Full-day tutorial at International Supercomputing Conference 2011 (ISC'11), Hamburg, Germany, June 19-23, 2011.

Our general course on parallelization, the **Parallel Programming Workshop**, Oct. 10-14, 2011 at HLRS, will have three parts: The first two days of this course are dedicated to parallelization with the Message Passing Interface (MPI). Shared memory multi-threading is taught on the third day, and in the last two days, advanced topics are discussed. As in all courses, hands-on sessions (in C and Fortran) will allow users to immediately test and understand the parallelization methods. The course language is English.

Several three and four day-courses on **MPI & OpenMP** will be presented at different locations all over the year.

We also continue our series of **Fortran for Scientific Computing** in January and July 2011, mainly visited by PhD students from Stuttgart and other universities in Germany to learn not only the basics of programming, but also to get an insight on the principles of developing high-performance applications with Fortran.

With **Unified Parallel C (UPC) and Co-Array Fortran (CAF)** in July and December 2011, the participants will get an introduction of partitioned global address space (PGAS) languages.
Our series on "**GPU Programming using CUDA**" will also continue in July and December 2011.

In cooperation with Dr. Georg Hager from the RRZE in Erlangen and Dr. Gabriele Jost from TACC, the HLRS also continues with contributions on

hybrid MPI & OpenMP programming at **PRACE Summer Schools** and **Tutorials** at conferences; see the box on the left page.

In the table below, you can find the whole HLRS series of training courses in 2011. They are organized at HLRS and also at several other HPC institutions: LRZ Garching, NIC/ZAM (FZ Jülich), ZIH (TU Dresden), TUHH (Hamburg-Harburg), and GRS/RWTH (Aachen).

2011 – Workshop Announcements
Scientific Conferences and Workshops at HLRS, 2011
14th Teraflop Workshop (March 21- 22)
10th HLRS/hww Workshop on Scalable Global Parallel File Systems (April 11 - 13)
5th ZIH+HLRS Parallel Tools Workshop (September 26 - 27, Dresden)
High Performance Computing in Science and Engineering - The 14th Results and Review Workshop of the HPC Center Stuttgart (October 4 - 5)
IDC International HPC User Forum (October 6 - 7)
Parallel Programming Workshops: Training in Parallel Programming and CFD
GPU Programming using CUDA (HLRS, January 24 - 25, July 4 - 6, and December 5 - 7)
Cray XE6 Optimization Workshops (HLRS, February 2 - 4, Sep. 6-7, and Nov. 2 - 4)
Parallel Programming and Parallel Tools (TU Dresden, ZIH, February 14 - 17)
Iterative Linear Solvers and Parallelization (HLRS, February 28 – March 4)
VI-HPS Tuning Workshop (HLRS, March 28 - 30)
Introduction to Computational Fluid Dynamics (HLRS, April 4 - 8)
Platforms at HLRS and Cray XE6 Optimization Workshop (HLRS, April 18 -19)
Unified Parallel C (UPC) and Co-Array Fortran (CAF) (HLRS, July 7 - 8, and Dec. 8 - 9)
Parallel Programming with MPI & OpenMP (TU Hamburg-Harburg, August 1 - 3)
Iterative Linear Solvers and Parallelization (LRZ, Garching, September 12 - 16)
5th Parallel Tools Workshop (ZIH Dresden, September 26 - 27)
Introduction to Computational Fluid Dynamics (GRS / RWTH Aachen, September 26 - 30)
Message Passing Interface (MPI) for Beginners (HLRS, October 10 -11)
Shared Memory Parallelization with OpenMP (HLRS, October 12)
Advanced Topics in Parallel Programming (HLRS, October 13 -14)
Parallel Programming with MPI & OpenMP (FZ Jülich, JSC, November 28 - 30)
Training in Programming Languages at HLRS
Fortran for Scientific Computing (January 17 - 21, and July 11-15)
URLs: http://www.hlrs.de/events/
http://www.hlrs.de/organization/sos/par/services/training/course-list/
https://fs.hlrs.de/projects/par/events/2011/parallel_prog_2011/

GCS – High Performance Computing

NWChem Training Tutorials at LRZ

Date and Location
June 9 - 10, 2011,
9:00 - 17:00
LRZ Building, Garching/Munich

Day 1
Morning lectures:
- Basic Introduction of Computational Chemistry
- Basic Introduction of NWChem Software
- Ground and Excited States with DFT and TDDFT
- Correlated Methods for Ground and Excited States

Afternoon:
- Reserved for hands-on tutorials (last lecture of morning can be moved to directly after lunch too)

Day 2
Morning lectures:
- Relativity, Spectroscopy and the EMSL Basis Set Library
- Classical Molecular Dynamics and QM/MM
- Ab initio Molecular Dynamics

Afternoon:
- Reserved for hands-on tutorials and getting your science to run

Audience
This workshop is aimed at new and experienced users of NWChem. Basic knowledge of computational chemistry is desirable.

Website
http://www.lrz.de/services/compute/courses

GPU Programming using CUDA

Dates and Location
July 4 - 6, 2011
(repeated December 5 - 7, 2011)
Stuttgart, HLRS

Contents
The course provides an introduction to the programming language CUDA, which is used to write fast numeric algorithms for NVIDIA graphics processors (GPUs). Focus is on the basic usage of the language, the exploitation of the most important features of the device (massive parallel computation, shared memory, texture memory) and efficient usage of the hardware to maximize performance. An overview of the available development tools and the advanced features of the language is given. Hands-on sessions (in C and Fortran) are included.

Webpage
http://www.hlrs.de/events/

Unified Parallel C (UPC) and Co-Array Fortran (CAF)

Dates and Location
July 7 - 8, 2011
(repeated December 8 - 9, 2011)
Stuttgart, HLRS

Contents
Partitioned Global Address Space (PGAS) is a new model for parallel programming. Unified Parallel C (UPC) and Co-Array Fortran (CAF) are PGAS extensions to C and Fortran. PGAS languages allow any processor to directly address memory/data on any other processors. Parallelism can be expressed more easily compared to library-based approches as MPI. Hands-on sessions (in UPC and/or CAF) will allow users to immediately test and understand the basic constructs of PGAS languages.

Webpage
http://www.hlrs.de/events/

Fortran for Scientific Computing

Date and Location
July 11 - 15, 2011
Stuttgart, HLRS

Contents
This introduction to C++ is taught with lectures and hands-on sessions. This course is organized by HLRS and Institute for Computational Physics. This course is dedicated for scientists and students to learn (sequential) programming of scientific applications with Fortran. The course teaches newest Fortran standards. Hands-on sessions will allow users to immediately test and understand the language constructs.

Webpage
http://www.hlrs.de/events/

Courses and Tutorials

Parallel Programming with MPI, OpenMP and PETSc

Date and Location
August 1 - 3, 2011
TU Hamburg-Harburg

Contents
The focus is on programming models MPI, OpenMP, and PETSc. Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the basic constructs of the Message Passing Interface (MPI) and the shared memory directives of OpenMP. This course is organized by the TUHH in collaboration with HLRS.

Webpage
http://www.hlrs.de/events/

JSC Guest Student Programme: Introduction to Parallel Programming with MPI and OpenMP

Date and Location
August 2 - 5, 2011
JSC, Research Centre Jülich

Contents
The course provides an introduction to the two most important standards for parallel programming under the distributed- and shared memory paradigms: MPI, the Message Passing Interface and OpenMP. While intended mainly for the JSC Guest Students, the course is open to other interested persons upon request.

Webpage
http://www.fz-juelich.de/ias/jsc/events/mpi-gsp

Cray XE6 Optimization Workshop

Date and Location
September 6 - 7, 2011
Stuttgart, HLRS

Contents
In December 2010, HLRS has installed a first Cray XE6 system with a performance of 10 TFlop/s. The final XE6 system with a performance of 1 PFlop/s is expected to be operational in Q3, 2011. We strongly encourage you to use the already existing small system to port your applications to the new architecture as early as possible. To support such effort we current and future users to participate in a special course Cray XE6 Optimization Workshop. With this course, we will give all necessary information to move applications from the current NEC SX-9, the Nehalem cluster, or other systems to the upcoming Petaflop system. The Cray XE6 will provide our users with a new level of performance. To harvest this potential will require all our efforts. We are looking forward to working with you on these opportunities.

Webpage
http://www.hlrs.de/events/

Iterative Linear Solvers and Parallelization

Date and Location
September 12 - 16, 2011
LRZ Building, Garching/Munich

Contents
The focus is on iterative and parallel solvers, the parallel programming models MPI and OpenMP, and the parallel middleware PETSc. Different modern Krylov Subspace Methods (CG, GMRES, BiCGSTAB ...) as well as highly efficient preconditioning techniques are presented in the context of real life applications.

Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the basic constructs of iterative solvers, the Message Passing Interface (MPI) and the shared memory directives of OpenMP.
This course is organized by LRZ, University of Kassel, HLRS, and IAG.

Webpage
http://www.lrz.de/services/compute/courses/

5th Parallel Tools Workshop

Date and Location
September 26 - 27, 2011
ZIH, TU Dresden

Contents
Developing for current and future processors will more and more require parallel programming techniques for

GCS – High Performance Computing

application and library programmers. This workshop offers to the industrial and scientific user community, as well as the tools developers itself an in-depth workshop on the state-of-the-art of parallel programming tools, ranging from debugging tools, performance analysis and best practices in integrated developing environments for parallel platforms. Participants and tools developers itself will get the chance to see the strengths of the various tools. Therefore, this workshop is focused on persons who already know about parallel programming, e.g. with MPI or OpenMP. Hands-on sessions will give a first touch and allow to test the features of the tools. This workshop is jointly organized by the Center for Information Services and High Performance Computing Dresden (ZIH) and HLRS.

Webpage
http://tools.zih.tu-dresden.de/

Introduction to Computational Fluids Dynamics

Date and Location
September 26 - 30, 2011
RWTH Aachen

Contents
Numerical methods to solve the equations of Fluid Dynamics are presented. The main focus is on explicit Finite Volume schemes for the compressible Euler equations. Hands-on sessions will manifest the content of the lectures. Participants will learn to implement the algorithms, but also to apply existing software and to interpret the solutions correctly. Methods and problems of parallelization are discussed. This course is based on a lecture and practical awarded with the "Landeslehrpreis Baden-Württemberg 2003" and organized by German Research School for Simulation Sciences.

Webpage
http://www.hlrs.de/events/

GPGPU Programming

Date and Location
October 10 - 12, 2011,
10:00 - 17:00
LRZ Building, Garching/Munich

Day 1
Introduction to NVIDIA CUDA, CUDA Parallel Programming, Hands-on

Day 2
HMPP (LRZ), Advanced@CUDA (nVidia)

Day 3
PGI Accelerator Compilers, GPU-accelerated Software (LRZ)

Prerequisites
Participants should have a fair understanding of programming in general and should have knowledge of C/C++

Website
http://www.lrz.de/services/compute/courses

Message Passing Interface (MPI) for Beginners

Date and Location
October 10 - 11, 2011
Stuttgart, HLRS

Contents
The course gives an full introduction into MPI-1. Further aspects are domain decomposition, load balancing, and debugging. An MPI-2 overview and the MPI-2 one-sided communication is also taught. Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the basic constructs of the Message Passing Interface (MPI). Course language is ENGLISH (if required).

Webpage
http://www.hlrs.de/events/

Shared Memory Parallelization with OpenMP

Date and Location
October 12, 2011
Stuttgart, HLRS

Contents
This course teaches shared memory OpenMP parallelization, the key concept on hyper-threading, dual-core, multi-core, shared memory, and ccNUMA platforms. Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the directives and other interfaces of OpenMP. Tools for performance tuning and debugging are presented. Course language is English (if required).

Courses and Tutorials

Webpage
http://www.hlrs.de/events/

Advanced Topics in Parallel Programming

Date and Location
October 13 - 14, 2011
Stuttgart, HLRS

Contents
Topics are MPI-2 parallel file I/O, hybrid mixed model MPI+OpenMP parallelization, OpenMP on clusters, parallelization of explicit and implicit solvers and of particle-based applications, parallel numerics and libraries, and parallelization with PETSc. Hands-on sessions are included. Course language is English (if required).

Webpage
http://www.hlrs.de/events/

Introduction to the Usage of High Performance Systems, Remote Visualization and Grid Facilities at LRZ

Date and Location
October 20, 2011,
10:00 - 17:00
LRZ Building, Garching/Munich

Content
The National High-End System HLRB II provides nearly ten thousand cores and the Linux cluster systems at LRZ more than 4,000 cores to scientists and students. Furthermore, powerful remote visualization facilities are available. Based on hands-on examples, an easy-to-follow introduction to basic Linux usage, specific information on the hard- and software of the LRZ cluster systems, the visualization systems and usage of the Grid middleware (Globus Toolkit) is given. Grid certificates can be provided if needed.

Prerequisites
For obtaining a grid certificate participants are required to show a valid ID card or passport.

Website
http://www.lrz.de/services/compute/courses

Introduction to Molecular Modeling on Supercomputers

Date and Location
October 24 - 26, 2011,
10:00 - 17:00
LRZ Building, Garching/Munich

Content
The course gives an introduction into the simulation of molecules based on several software packages on the supercomputers at LRZ Garching. (Maestro, Desmond, VMD, NAMD, Gromacs) This also includes an introduction to the remote visualization services at LRZ as well as hands-on sessions. The course focuses on biomolecules and targets the life science community. Hands-on sessions featuring example applications are given.

Day 1
Advanced feature of the graphical user interface Maestro (Schrodinger)

Day 2
Remote visualization services, VMD/NAMD (LRZ)

Day 3
Gromacs and Dalton (ScalaLife)

Prerequisites
Participants should have some basic knowledge in programming and life science software.

Website
http://www.lrz.de/services/compute/courses

GPU Programming

Date and Location
October 2011
JSC, Research Centre Jülich

Contents
An introduction to CUDA, OpenCL, and multi-GPU programming is given using examples of increasing complexity. The focus will be on optimization and tuning of scientific applications.

Prerequisites
Knowledge in C

Webpage
http://www.fz-juelich.de/ias/jsc/events/gpu

GCS – HPC Courses and Tutorials

Cray XE6 Optimization Workshop

Date and Location
November 2 - 4, 2011
Stuttgart, HLRS

Contents
In December 2010, HLRS has installed a first Cray XE6 system with a performance of 10 TFlop/s. The final XE6 system with a performance of 1 PFlop/s is expected to be operational in Q3, 2011. We strongly encourage you to use the already existing small system to port your applications to the new architecture as early as possible. To support such effort we current and future users to participate in a special course Cray XE6 Optimization Workshop. With this course, we will give all necessary information to move applications from the current NEC SX-9, the Nehalem cluster, or other systems to the upcoming Petaflop system. The Cray XE6 will provide our users with a new level of performance. To harvest this potential will require all our efforts. We are looking forward to working with you on these opportunities. Course language is English (if required).

Webpage
http://www.hlrs.de/events/

Introduction to the Programming and Usage of the Supercomputing Resources in Jülich

Date and Location
November 24 - 25, 2011
JSC, Research Centre Jülich

Contents
This course gives an overview of the supercomputers JUROPA/HPC-FF and JUGENE. Especially new users will learn how to program and use these systems efficiently. Topics discussed are: system architecture, usage model, compiler, tools, monitoring, MPI, OpenMP, performance optimization, mathematical software, and application software.

Webpage
http://www.fz-juelich.de/ias/jsc/events/sc

Parallel Programming with MPI, OpenMP and PETSc

Date and Location
November 28 - 30, 2011
JSC, Research Centre Jülich

Contents
The focus is on programming models MPI, OpenMP, and PETSc. Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the basic constructs of the Message Passing Interface (MPI) and the shared memory directives of OpenMP.

This course is organized by JSC in collaboration with HLRS. Presented by Dr. Rolf Rabenseifner, HLRS.

Webpage
http://www.fz-juelich.de/ias/jsc/events/mpi

Parallel Computing with R

Date and Location
December 5 - 6, 2011,
09:00-17:00
LRZ Building, Garching/Munich

Content
Primary drivers for the increased focus on parallel computing are new hardware trends (multi-core), larger data sets, and increased computational requirements stemming from more sophisticated methodologies. This course demonstrates the efficient use of R in parallel computing. In the beginning advanced R programming skills (vectorization, apply functions, profiling) will be repeated. After a short theoretical course for parallel computing the parallel program design, methods and techniques for parallel computing, and parallel thinking are communicated with several examples. Main part of the course is the practical application of the R packages snow, multicore and foreach. Nevertheless, there are exercises for using batch systems and other R packages for parallel computing (snowfall, NWS).

Prerequisites
Basic knowledge in R and first R programming skills (e.g. implementing your own R functions). Basic knowledge in using Linux systems.

Website
http://www.lrz.de/services/compute/courses

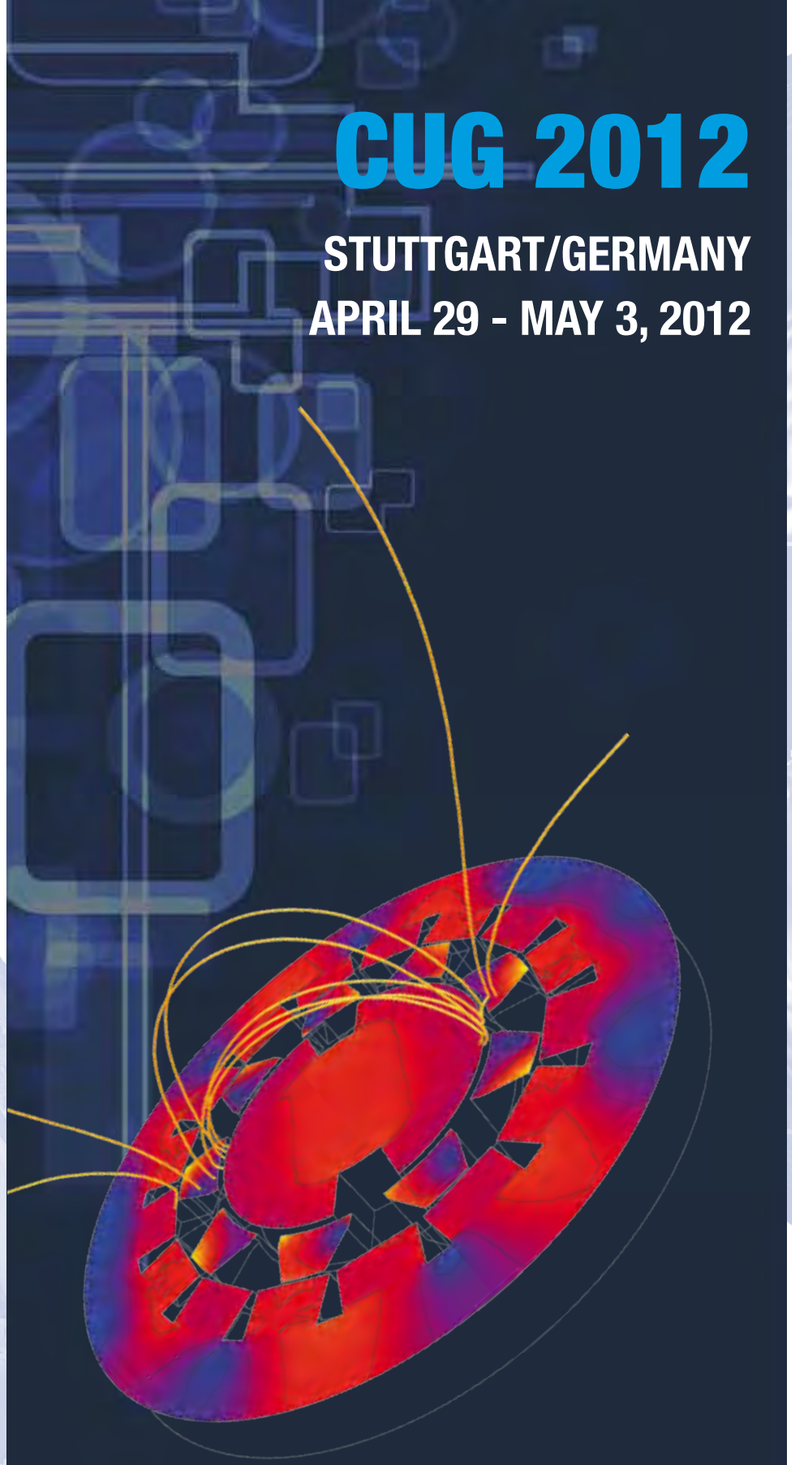

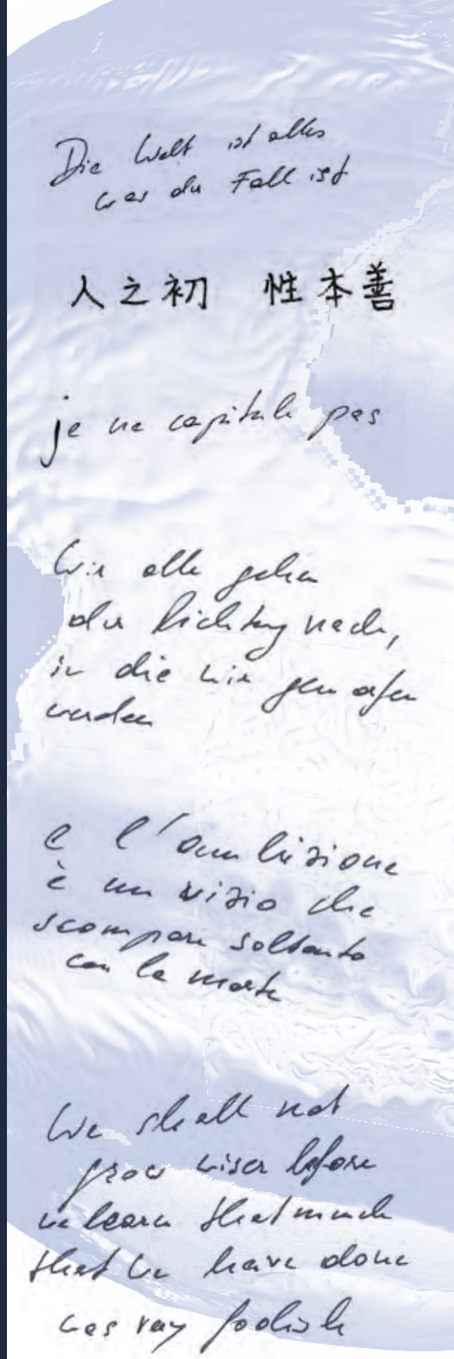

inSiDE is published two times a year by The GAUSS Centre for Supercomputing (HLRS, LRZ, JSC)

Publishers

Prof. Dr. H.-G. Hegering | Prof. Dr. Dr. T. Lippert | Prof. Dr. M. M. Resch

Editor & Design

F. Rainer Klank, HLRS klank@hlrs.de
Carina Möhlig, HLRS moehlig@hlrs.de

Authors

Alexander Beck-Ratzka
alexander.beck-ratzka@aei.mpg.de
Felix Bensing
bensing@iag.uni-stuttgart.de
Arne Biastoch
abiastoch@ifm.geomar.de
Thomas Boenisch
boenisch@hlrs.de
Matthias Brehm
Matthias.Brehm@lrz.de
Andreas Burkert
Anton Daitche
a.daitche@googlemail.com
Dietmar Erwin
d.erwin@fz-juelich.de
Tillmann A. Friederich
friederich@iag.uni-stuttgart.de
Rudolf Friedrich
fiddir@uni-muenster.de
Stefan Frijters
s.c.j.frijters@tue.nl
Markus Gampert
m.gampert@itv.rwth-aachen.de
Michael Gauding
m.gauding@itv.rwth-aachen.de
Shy Genel
Shy.Genel@campus.lmu.de
Reinhard Genzel
Pierre Gilet
gilet@hlrs.de
Jens Henrik Göbbert
jh.goebbert@itv.rwth-aachen.de
Florian Günther
f.s.guenther@tue.nl
Carla Guillen
Carla.Guillen@lrz.de
Jens Harting
jens@icp.uni-stuttgart.de
Wolfram Hesse
Wolfram.Hesse@lrz.de
Luigi Iapichino
luigi@ita.uni-heidelberg.de
Florian Janoschek
florian@icp.uni-stuttgart.de
Fabian Jansen
fabian@icp.uni-stuttgart.de
Peter Johansson
Oliver Kamps
okamp@uni-muenster.de
Manuel Keßler
kessler@iag.uni-stuttgart.de
Markus J. Kloker
kloker@iag.uni-stuttgart.de
Bastian Koller
koller@hlrs.de
Mariana Ležaić
m.lezaic@fz-juelich.de
Olaf Mextorf
o.mextorf@fz-juelich.de
Jörg Meyer
joerg.meyer@ch.tum.de
Bernd Mohr
b.mohr@fz-juelich.de
Thorsten Naab
joerg.meyer@ch.tum.de
Walter Nadler
w.nadler@fz-juelich.de
Eric Mc Nellis
erikmcnellis@gmail.com
Ludwig Oser
Ludger Palm
Ludger.Palm@lrz.de
Norbert Peters
n.peters@itv.rwth-aachen.de
Rolf Rabenseifner
rabenseifner@hlrs.de
Michael M. Resch
resch@hlrs.de
Karsten Reuter
karsten.reuter@ch.tum.de
Helmut Satzger
Helmut.Satzger@lrz.de
Philip Schäfer
p.schaefer@itv.rwth-aachen.de
Godehard Sutmann
g.sutmann@fz-juelich.de
Axel Tenschert
tenscher@hlrs.de
Lipo Wang
lipo.wang@sjtu.edu.cn
Michael Wilczek
mwilczek@uni-muenster.de
Felix Wolf
f.wolf@grs-sim.de

If you would like to receive inSiDE regularly, send an email with your postal address to klank@hlrs.de