# inside

## Innovatives Supercomputing in Deutschland

**10 YEARS**

# Editorial

Welcome to this new issue of InSiDE, the journal on innovative Supercomputing in Germany published by the Gauss Centre for Supercomputing (GCS). We are well into our 10th year and continue to look back a little bit. This time we have invited the Chairman of the Board of GCS Prof. Dr. Heinz-Gerd Hegering for an interview on High Performance Computing in Germany and Europe.

Prof. Hegering not only looks back on the history of High Performance Computing. He and the Leibniz Computing Centre (LRZ) also have good reason to celebrate and look positively towards the future. LRZ inaugurated its new system called "SuperMUC" in July. Following the GCS strategy SuperMUC introduces a new architecture to users in Europe and Germany. Emphasizing Germany's role as a leader in European High Performance Computing LRZ and GCS were also proud to present the SuperMUC to be Europe's number one system in the TOP 500 list. So the spiral of performance in Germany is working as designed and the Jülich Supercomputing Centre (JSC) and the High Performance Computing Center Stuttgart (HLRS) are both preparing for their next step.

The work of GCS is a major part of the European PRACE initiative. Results of the 4th Pan-European call for user proposals are presented in this issue. Showing off these results and getting in touch with a wider community is part of two further contributions: A report on ISC'12 shows the strong interaction of GCS with the international supercomputing community. Introducing the newly created SICOS GmbH we show how small and medium sized enterprises can be supported in their usage of HPC systems.

Again we present a rich section on applications running on the GCS systems. Ranging from Computational Fluid Dynamics to astrophysical simulations the eight contributions clearly show the need for compute power and the potential of HPC simulation for both basic and applied science. The growing strength of GCS in HPC research is highlighted by the seven project descriptions that show how various groups at Garching, Jülich and Stuttgart contribute to the progress in HPC research.

GCS has made big steps over the last years. The co-operation with PRACE has extended its reach into Europe and has allowed us to contribute to European research and development substantially. GCS is committed to continuing these efforts. In the field of HPC research GCS is moving forward both in European and national large scale research programs.

With this, we look forward to a bright future for supercomputing in Germany and hope you enjoy reading.

- Prof. Dr.
  A. Bode (LRZ)

- Prof. Dr. Dr.
  Th. Lippert (JSC)

- Prof. Dr.-Ing. Dr. h.c. Dr. h.c.
  M. M. Resch (HLRS)

# Contents

# "SuperMUC", Europe's fastest Supercomputer, inaugurated at the 50th Anniversary of LRZ



Figure 1: Prof. Dr. Arndt Bode, Chairman of LRZ, Martina Koederitz, General Manager IBM Germany, Federal Minister for Science and Education Prof. Dr. Annette Schavan and Bavaria's Minister of State for Research, Science and the Arts Dr. Wolfgang Heubisch.

"SuperMUC", LRZ's new supercomputer is ranked no. 1 in Europe and no. 4 in the world as announced on the International Suprecomputing Conference at Hamburg on June, 18th, 2012, by Prof. Dr. Hans Meurer, co-founder of the Top500 list of supercomputers (www.top500.org).

SuperMUC, as described in detail in our last issue, delivers 3 PetaFlop/s peak performance to Germany's and Europe's

sciencists. On July 20, 2012, LRZ celebrated the inauguration of SuperMUC and its 50th anniversary. Several hundred guests attended the event together with Prof. Dr. Annette Schavan, Germany's Federal Minister for Science and Education, and Dr. Wolfgang Heubisch, Bavaria's Minister of State for Research, Science and the Arts.

In 1962, the Bavarian Academy of Sciences founded a "Commission for Electronic Computing", which was later renamed to "Commission for Informatics". This Commission for Informatics established and runs the Leibniz-Rechenzentrum LRZ (Leibniz Supercomputing Centre). Starting in a residential building in Munich, it moved to a new and larger building near the famous Karlsplatz "Stachus" in the centre of Munich a few years later. In 2006, LRZ had to move again, now to a much larger facility on the campus in Garching, north of Munich, which had to be extended further already in 2009 to 2011 to its current impressive size.

While High Performance Computing was in the beginning and is at the heart of LRZ, LRZ also is the general IT service provider for Munich's universities and serves more than 100,000 customers in the Munich area. Prof. Dr. Heinz-Gerd Hegering, past Chairman of LRZ for more than twenty years, at the celebration of the 50th anniversary gave an overview on the evolution of the many different services offered by LRZ.

LRZ always has been a pioneer in information technology e.g., communication networks, long term archiving, High Performance Computing in general and today with a special focus on energy efficiency. Thus, SuperMUC, manufactured by IBM, is the first computer of its class that is completely cooled with hot water. LRZ is a leading expert in the field of warm water cooling of computers. The extension of the computer building has been equipped with the necessary infrastructure for this innovative and highly energy-efficient way to run computers from SuperMUC to Linux clusters.

Federal Minister Prof. Dr. Annette Schavan as well as Minister of State Dr. Wolfgang Heubisch pointed out the importance of High Performance Computing for Germany and Bavaria



Figure 4: The first LRZ building.



Figure 5: The LRZ building in the centre of Munich.



Figure 6: The LRZ building after the extension 2009-2011 in Garching near Munich.

which is stressed by the fact that the Federal Government of Germany and the Free State of Bavaria jointly spent more than 80 Million Euro for the present configuration of SuperMUC and about 50 Million Euro for the extensions of the buildings of LRZ. SuperMUC is now available for academia in Germany as part of the Gauss Centre for Supercomputing GCS and for many researchers in 24 member states via the Partnership for Advanced Computing in Europe, PRACE.

• Ludger Palm

Leibniz Supercomputing Centre (LRZ)



Figure 2: Prof. Dr. Hans Meuer, co-founder of the Top500 list, handing the certificates to Martina Koederitz, General Manager IBM Germany, and to Prof. Dr. Arndt Bode, Chairman of LRZ.



Figure 3: Celebrating the 50th anniversary of LRZ and the inauguration of "SuperMUC" with Ministers Prof. Dr. Annette Schavan and Dr. Wolfgang Heubisch.



Figure 7: Prof. Dr. Heinz-Gerd Hegering, LRZ's Chairman for more than 30 years, at the ceremony.



Figure 8: Ministers Schavan and Heubisch putting their hands on the warm-water cooling of SuperMUC.

# Interview with Prof. Hegering

Prof. Hegering, you have been the chairman of the board of the Gauss Centre for Supercomputing (GCS) for the last 4 years. What have been the major changes over that time?

We all are witnesses to the fact that GCS has grown into one virtual supercomputing center that is responsible for the German tier 0 and tier 1 systems of the HPC pyramid in the national as well as the European context:

- In all 3 member centers of GCS (Garching, Jülich, Stuttgart) Petaflop-systems were installed belonging to the worldwide leading systems in the TOP500 list. Being of different but complementary architectures these systems serve a broad range of sciences. GCS now offers the largest and most powerful supercomputing infrastructure in Europe and a vast range of industrial and research activities in various disciplines will benefit from it.

- Access to GCS's resources is granted by a common GCS review and access committee through a unified evaluation process based on purely scientific criteria.

- The GCS governance implies clear and transparent coordination processes with regards to procurements, educational programs, outreach, as well as defining focal and main thematic emphases.

- GCS is member of PRACE and has a leading role in the European HPC sce-nario: it provides the most extensive system base in the Petaflop-range, has the role of PMO in the PRACE projects, and held until just recently also the role of the council chairman in PRACE AISBL.

- GCS is also active in the continuation of the German national HPC concept. Last year GCS proposed a concept paper that could serve as a replacement of the "Reuter-Papier" from 2006 and served as input for the latest HPC-related recommendations of the German Wissenschaftsrat which is the advisory committee of the German government.

The world of HPC is discussing the roadmap to Exaflops intensively. How realistic do you see the chances that we reach this target within the next 8 to 10 years?

It seems to be certain that there will be an Exascale-system within the next decade. And Germany should participate in the development of such a system. Within GCS especially JSC is active in the related field. But I personally believe that within the next 6 years such systems will rather show more experimental characteristics and be still subject to related research than serve as common production platform for a broader scientific usage in an effective and efficient way. Think of the energy consumption and the lack of suitable languages, algorithms, and tools to reach a really high sustained performance.

HPC systems are never stand-alone installations but are enmeshed in an infrastructure of networks and data handling systems. How do you see the development of this overall infrastructure when we look at Exaflops?

The answer is related to what I mentioned before. I think it is important for the overall HPC infrastructure that the HPC pyramid in Europe as well as in the various countries is somewhat well balanced throughout all tiers. For the next round of HPC installations I expect that the leading HPC platforms for real use will be HPC systems with perhaps hundreds of Petaflop/s but not yet Exascale machines.

It is also a question of how much money we should spend for an experimental system that could be detrimental to top production machines. In any case we have to invest in people to improve usage efficiency.

With the European initiative PRACE, HPC has become more international in Europe than ever before. How do you see the role of PRACE in Europe and its impact on German HPC?

The success of PRACE has been to bring the European High Performance Computing community closer together and to provide a system platform of considerably more capacity and capability. This could not have been reached without the European FP7-call in 2007. In Germany this call helped to found GCS in a very short time frame and to allo-cate that significant amount of funding for the project PetaGCS.

When we look at the national level we find that Germany has a well-established pyramid of performance and has organized the medium sized centers in the Gauss Alliance. How do you see the future development of HPC in Germany as a whole?

GCS proved to become a success story that is well recognized in Germany by science and politics. If you take a look at the positive statements in the latest recommendations of the Wissenschaftsrat, it is clearly stated that HPC is of national importance and that GCS is the adequate organization to form the top of the HPC pyramid in Germany with respect to the tiers 0 and 1. The situation for the tiers 2 and 3 is more complicated due to our national constitution, since most of the tier 2 and 3 centers underlie the governance of the local German states. Possible changes of our constitution are under discussion, but this is politically a very complicated field. From a pragmatic standpoint of view it was wise to separate concerns and to found Gauss Centre (GCS) and Gauß-Allianz (GA) as separate associations. The preconditions for a good cooperation were already designed in the statutes of both associations since GCS is a member of GA and a board member of GCS is also member of the GA board. I think that this is the best approach to build a powerful HPC pyramid in Germany facing the different political challenges of

GCS and the Gauß-Allianz. We already have a strong cooperation for joint software research and development projects. I am fully convinced that the overall approach taken so far will lead to a successful structure of the whole HPC pyramid in Germany.

After several decades in HPC, what do you personally see as the most outstanding achievement?

Especially during the last 20 years the realization that simulation besides theory and experiment is an increasingly important and indispensable method for reaching scientific findings has gained momentum. This awareness leads to supporting the concept of national supercomputing centers in 1999. In 2006 the German Federal Ministry for Education and Research, BMBF, gave the inducement to found GCS as an organization that became obligatory for the realization of the top

of the national HPC concept. GCS was selected to carry out the PetaGCS project and to provide the German contributions to the European PRACE project. We succeeded in achieving the leading position in Europe in less than one decade
thanks to the great support of BMBF and the states of Baden-Württemberg, Bavaria, and North Rhine-Westphalia. Now there is the challenge to continue this success story and make the achievements sustainable at a high level.

Prof. Hegering, thank you for the interview. The Interview was conducted by the inSiDE team

**Prof. Heinz-Gerd Hegering was the Managing Director of the Leibniz Rechenzentrum (LRZ) at Garching until 2008. He is one of the founders of GCS and has been the chairman of the Board of Directors of GCS since May 2008.**

# PRACE: Results of the 4th Regular Call

The Partnership for Advanced Computing in Europe (PRACE) is offering super-computing resources on the highest level (tier-0) to European researchers.

The Gauss Centre for Supercomputing (GCS) is currently dedicating shares of its IBM Blue Gene/P system JUGENE in Jülich (currently being upgraded to the IBM Blue Gene/Q system JUQUEEN, see also the corresponding article in this edition), of its Cray XE6 system Hermit in Stuttgart, and of its IBM iDataPlex system SuperMUC in Garching.

The 4th call for proposals, this time for computing time on CURIE, hosted by GENCI in TGCC/CEA, Bruyères-Le-Châtel, France, MareNostrum, hosted by BSC in Barcelona, Spain, FERMI, hosted by CINECA in Casalecchio di Reno, Italy, in addition to the above mentioned GCS systems, closed January 10, 2012.

Eleven research projects have been awarded a total of about 300 million compute core hours on JUGENE, four have been awarded a total of about

140 million compute core hours on Hermit, and ten have been awarded a total of about 200 million compute core hours on SuperMUC for the allocation time period May 2012 to April 2013. Six of those research projects are from Italy, five are from France, four are from Germany, two are from Spain and the UK, each, while one each are from Belgium, Denmark, Finland, Ireland, Portugal, and Switzerland. The research projects awarded computing time cover all scientific areas, from Astrophysics to Medicine and Life Sciences. More details, also on the projects granted access to the machines in France, Spain, and Italy, can be found via the PRACE web pages www.prace-ri.eu/PRACE-4th-Regular-Call.

Evaluation for the 5th call for proposals that closed May 30, 2012 is still under way, as of this writing. Details on the calls, also on the current 6th call, can be found on www.prace-ri.eu/Call-Announcements.

• Walter Nadler

Jülich
Supercomputing
Centre (JSC)

# GCS @ ISC'12

At this year's ISC, the Gauss Centre for Supercomputing came up with a novelty: For the first time in its existence, GCS presented itself on a conjoint booth uniting the three national GCS centres HLRS, JSC and LRZ on a single presentation platform. The brand new 64m² large booth attracted countless like-minded HPC researchers, technology leaders, scientists, IT-decision makers as well as high tech media representatives who sought the interchange with the directors of the three GCS centres Prof. Bode (LRZ), Prof. Lippert (JSC), and Prof. Resch (HLRS) as well as Dr. Claus Axel Müller (Managing Director, GCS) and further GCS representatives.

One of the first visitors to the GCS booth was Dr. George Liang-Thai Chiu of the IBM Thomas J. Watson Research Center. Earlier that day, Dr. Chiu had been awarded the GCS Award in recognition of his scientific paper "Blue Gene/Q: By Co-Design", submitted to the ISC'12 Research Paper Sessions.
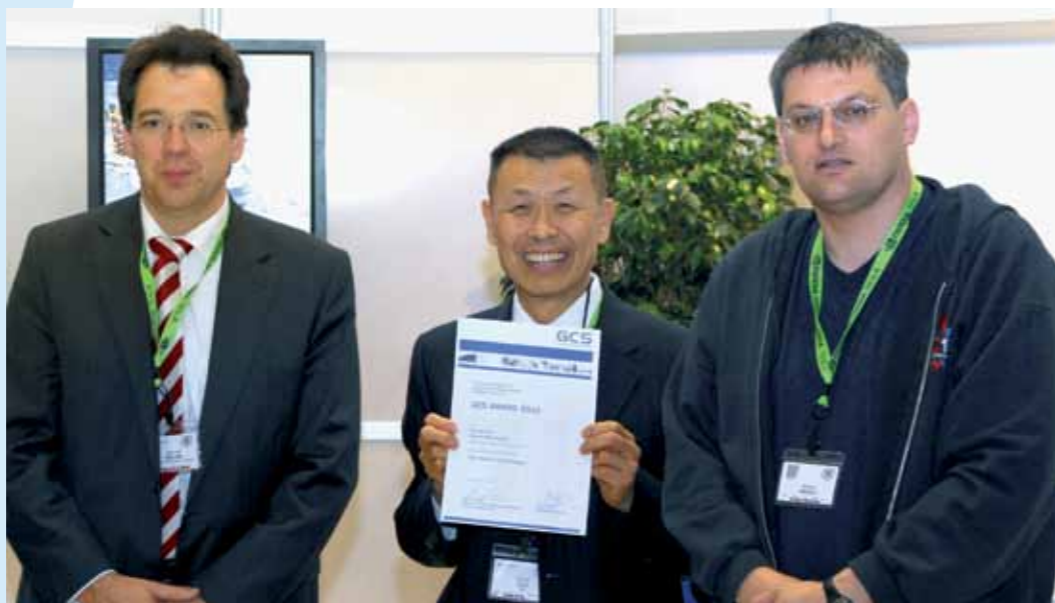
## Two GCS Systems in Top Ten of TOP500

For GCS, undisputed highlight of this year's ISC was the fact that two GCS systems made it into the Top Ten of the latest TOP500: With a Linpack peak performance of about 3 Petaflops, LRZ's SuperMUC ranks as Europe's fastest HPC system to date and came in 4th on the TOP500 world wide listing while JSC's new JUQUEEN (Linpack peak: 1.6 Petaflops) took position 8 in the over-all ranking. HLRS's flagship computer Hermit successfully defended its title as the world's fastest industrially used supercomputer. With a peak performance of more than 1 Petaflops, Hermit continues to be number 1 on the TOP500 Industrial Sub-List.

## Four GCS Workshops, Thursday Keynote on Super-MUC

At ISC'12, GCS hosted a total of four well-attended workshops. The Directors of the three GCS centres and other GCS representatives addressed various subjects related to supercomputing and HPC challenges in the fields of science and industry, underscoring Germany's efforts to maintain a leading role in the supply and use of HPC systems and services in Europe and beyond. Session titles included: "HPC Strategies in Germany", "Supercomputing Projects at LRZ", "Simulation Laboratories at JSC: Community-Oriented Software Support and Development for HPC", and "Scalability Matters for Industrial Applications?!" The GCS workshops were complemented by Thursday's Keynote delivered by Prof. Arndt Bode (LRZ) on the subject "Extreme Energy Efficiency with SuperMUC".

## GCS Booth Highlights

On the GCS booth, the three GCS centres presented their wide-ranging HPC activities. HLRS concentrated on interactive 3D visualizations of simulation results on a 2x2 3D display wall. Its highlight was the visualization of a pumped hydro power station showing large scale terrain rendering combined with visualization of the power house and the actual water turbine/pump and the visualization of water flow through the draft tube of that turbine. JSC exhibited a broad spectrum of scientific results obtained with its super-computers JUGENE and JUROPA in presentation videos and animations. In particular, JSC showcased LLview, the comprehensive interactive monitoring software for supercomputers developed in-house, in live demonstrations on supercomputers worldwide. Apart from demonstrating simulations on its current HPC systems, LRZ had its focus once again on SuperMUC and provided details about its newly installed, highly energy-efficient 3 Petaflop/s system to the vast number of visitors and journalists.

## About ISC

The ISC has seen an extremely positive feedback by the scientific community and has proven to be the number one event in Supercomputing together with the US Supercomputing Conference. GCS as a main supporter of the conference and fair has contributed in a variety of activities to this success and continues to endorse ISC.

• Regina Weigand

Gauss Centre for
Supercomputing
(GCS)



Figure 1: GCS Award winner Dr. Chiu with Claus Axel Müller (Managing Director, GCS, left) and Prof. Michael M. Resch (GCS Board Member)



Figure 2: The brand new GCS booth at ISC'12

# SICOS: HPC and more for Industry

The competitive pressure for enterprises grows continuously. Especially for, but not limited to companies in the field of manufacturing (automotive, aerospace, machinery, etc.) the importance of optimal efficiency with respect to their products cannot be overestimated. In times of dramatically increasing energy costs the weight of products plays a key role for example in the transportation context. Being able to decrease the weight of parts while at the same time improving their structural stability is an important advantage over the competition.

Simulation technology has been developed by scientists in the last decades into an extremely powerful tool. Large companies have taken simulation up and integrated it into their development process. Modern cars for example would either be less secure or much more expensive if they had not been developed using crash simulations.

However this take-up, that is not so easy even for large companies, has not yet fully reached the smaller and medium sized enterprises. They are usually not able to invest into the learning time, that every company that wants to use simulation technology needs, until the benefits can be harvested. And even if they are successful in applying this technology: especially smaller companies find it difficult to scale their IT infrastructure to the growing needs, once they want to integrate simulation into their standard development process.

This is the background that led the Karlsruhe Institute of Technology (KIT) and the University of Stuttgart to found the SICOS GmbH more than one year ago. SICOS stands for SImulation, Computing and Storage, which outlines the technology fields that SICOS will address primarily. Its main task is to facilitate the access to the resources of its founders for industry, with focus on small and medium sized enterprises (SMEs); regionally SICOS concentrates on Baden-Württemberg, but is not limited to it.

Parts of the activities of SICOS are also oriented towards the internal structures of its founders' computing centres, the HLRS in Stuttgart and SCC in Karlsruhe. In order to be able to offer an attractive platform for industrial users, SICOS will work with the two computing centres on optimizing the internal processes to meet the needs of the customers (which in turn also helps the academic users in many aspects). Currently a common security concept is under development and the first steps towards an integrated user management between the two centres have been made.

Of course these activities are fully aligned with the hww - Höchstleistungsrechner für Wissenschaft und Wirtschaft GmbH. Since more than 15 years KIT, the University of Stuttgart and the State of Baden-Württemberg closely cooperate with industry (today: Porsche and T-Systems) in the field of High Performance Computing as partners in hww. hww is naturally integrated in the projects that SICOS pursues regarding common security and user management in order to provide a professional and integrated service to the outside.

SICOS will be developed into a one-stop-shopping institution for industrial users that are looking for a way to utilize numerical simulation or large scale data facilities for their needs. This ranges from novices, who must be led through the initial phases of the use of simulation technology, over occasional users, who just need some help now and then, up to high end users, who can use the resources of the shareholders in order to be able to improve their offering. With the power of its shareholders in the back and a wide spread network with highly competent partners, SICOS will be able to identify the best solutions and it supports the companies on their way to apply them successfully.

www.sicos-bw.de



• Andreas Wierse

SICOS GmbH


Figure 1: Dr. Andreas Wierse, Managing Director


Figure 2: Julia Englert, Communication


Figure 3: Individual Consulting

# Simulation of Nasal Cavity Flows for Virtual Surgery Environments

The anatomy and the functionality of the human respiration system is well understood, whereas physical processes such as the details of the gas transport within the respiration cycle lack insight. The influence of the air flow on physiological functions like the sense of smell and taste is only one example which is of great interest. Furthermore, the impact of pathologies in the human nasal cavity and the lung are still not fully explained and are the topic of intensive scientific research in the fields of medical treatment, engineering and computer science. An impaired respiration capability can lead to a reduction of the respiratory efficacy and can also induce diminished olfactory and degustation capabilities. Such pathologies are in general caused by allergies, mal- or deformations of e.g. the nasal septum, where their impact on the flow field in the nose plays a

major role in the analysis and understanding of the patients complaints. In addition, the nasal cavity plays an important role in the protection of the human lung. The large surface of the human nasal cavity allows the moisturization and heating of the air to attain optimal conditions for the lower airways. These functions can be impaired by e.g., congenital malformations or rhino-anaplasty. The nasal cavity is also responsible for filtering out particles and aerosols from the inhaled air. Diesel aerosols and respirable dust particles are proofed to cause damage to the human lung and induce the growth of cancer. Particle deposition in the respiration tract is a topic of intensive research to further understand the impact of such pathologies. To analyze the functions of the nasal cavity it is almost impossible to do in-vivo investigations, although some parameters like

the pressure and the temperature can be investigated with a rhinomanometry and temperature measurements in the pharynx. These aspects only scratch the surface of the knowledge, which is required to fully understand the physics of the human respiration and starve for methods of higher detail.

## "Rhinomodel", an interdisciplinary Research Project

Computer-based methods have been developed over the last decades for the analysis of fluid mechanical problems and are well established in air and space science. They are, however, not very common in medical fields due to the high complexity of human organs and the physical aspects of the respiration, i.e. the particle and gas transport, temperature and moisturization. Within the scope of several research projects funded by the German Research Foundation (DFG) the Institute of Aerodynamics of RWTH Aachen University (AIA) investigates the human respiration system experimentally with the help of Particle Image Velocimetry methods (PIV), a laser- and photography-based particle tracking system, and numerically by applying Computational Fluid Dynamics methods (CFD). The flow field in the complex shape of the nasal cavity and the human lung is simulated on High Performance Computers (HPC) like the Cray XE6 Hermit System at the Höchstleistungsrechenzentrum (HLRS) in Stuttgart. These projects are highly interdisciplinary as they also require the expertise of scientists from the medical and computer science field. Therefore, a strong cooperation exists with radiologists
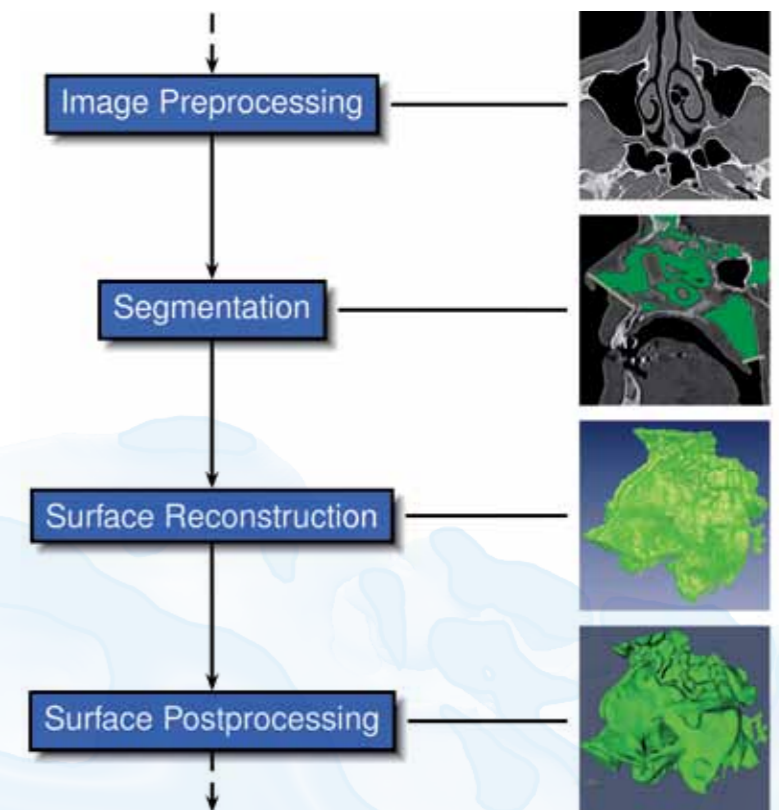


Figure 2: Surface extraction pipeline of the human nasal cavity surface from CT-data. The CT-images are preprocessed by a sharpening filter before a segmentation takes place, which splits the cavity into volumes of tissue and air. The surface is extracted and smoothed to obtain an accurate geometry of the nasal cavity.

from the University Hospital in Aachen (UKA), rhinology experts from the Institute of Statistics and Medical Informatics Cologne (IMSIE) and the Virtual Reality (VR) group of the Computing Center of the RWTH Aachen University (RZ RWTH). A key aspect of this research is the relation between the physical properties of the flow and the comfort of the patient, which is raised in a patient survey.

The study cycles through an analysis procedure as shown in Fig. 1, where in a first step it is required to obtain the geometrical representation of the nasal cavity from Computer Tomography (CT) images obtained from the UKA. The CT-images contain density intensities
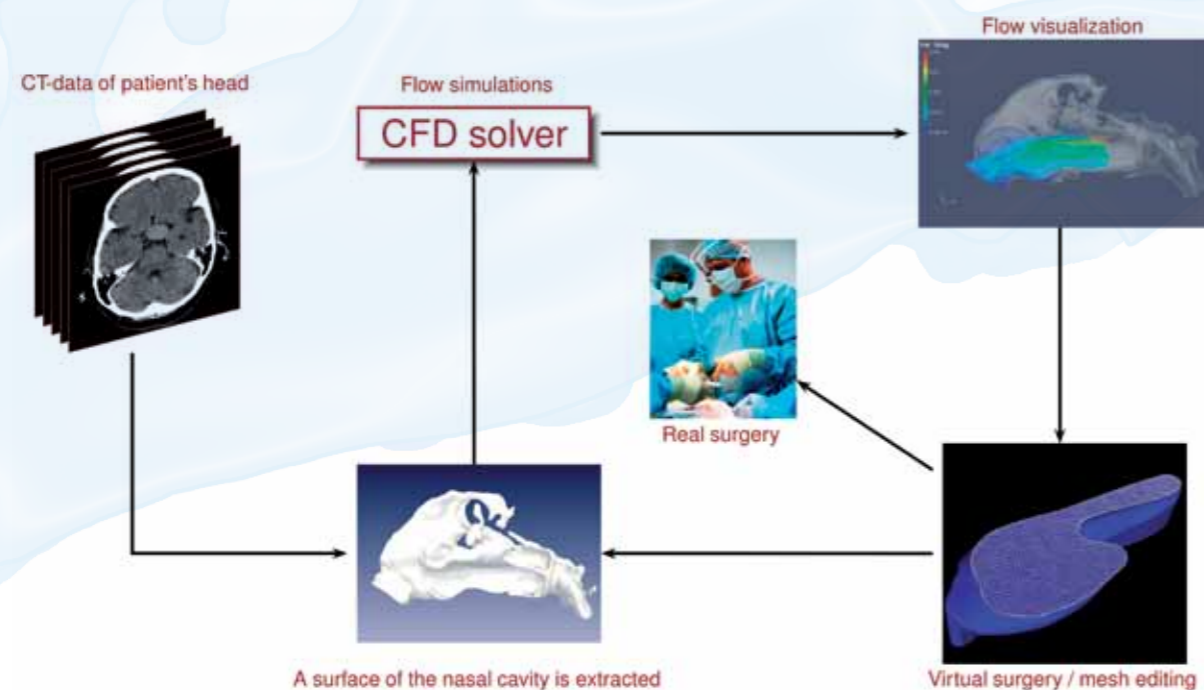


Figure 1: Virtual surgery environment for treatment of individual pathologies in rhinology. A surface of the nasal cavity is extracted from CT-images and is used for flow simulations to analyze the patient's complaints. A virtual surgery is performed and the resulting nasal cavity is analyzed and evaluated until a satisfactory treatment is found and a real surgery is carried out.
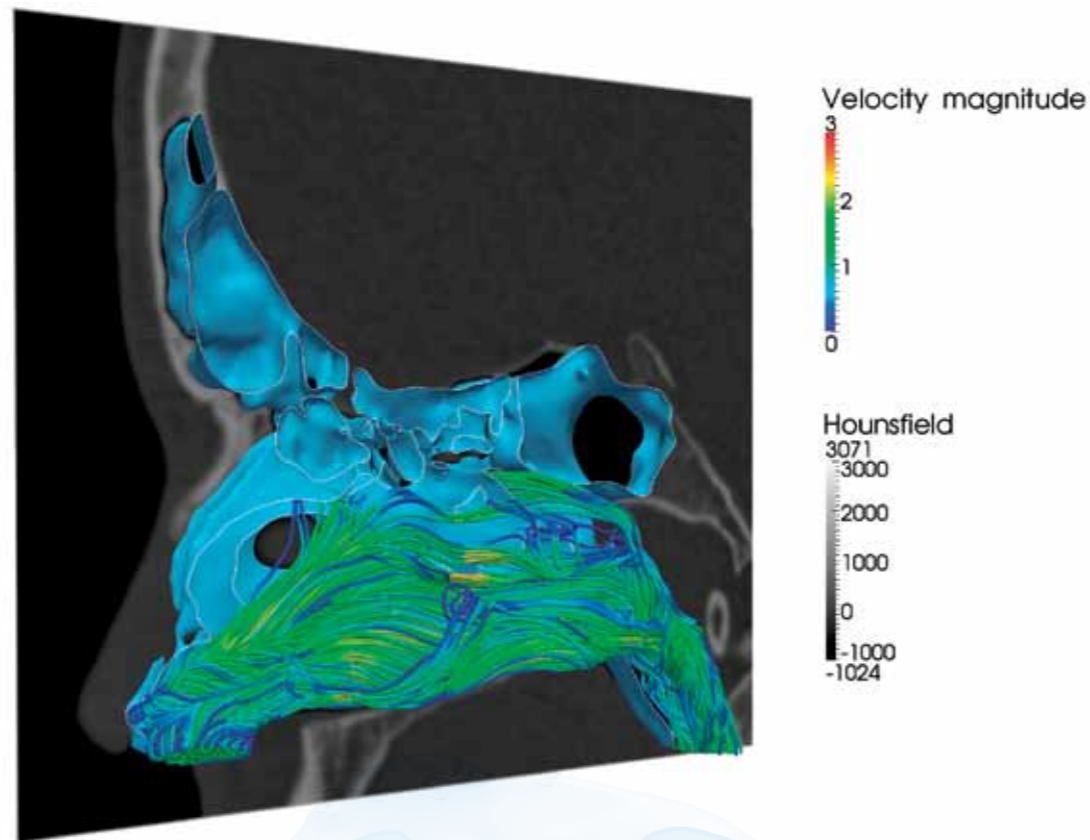
Figure 3: Surface of the nasal cavity based on CT-image data. A saggital slice of the CT-data shows the relative position of the extracted surface. Streamlines, colored by their velocity visualize the flow field.

(see Fig. 4) the pipeline enters an optimization cycle, where at the beginning a flow simulation with CFD methods is performed to obtain the raw fluid mechanical properties, which are then analyzed by surgeons from the UKA and engineers from the AIA. Based on the gained knowledge, the surgeons perform a virtual surgery in the virtual environment "CAVE" provided by the VR group of the RZ RWTH and edit the extracted surface by virtually ablating or remodeling tissue. This results in a new triangular representation of the nasal cavity, which again enters the optimization cycle to detect improvements to the patient's comfort based on the fluid mechanical properties of the flow. This cycle is repeated until a satisfactory modification of the nasal cavity is obtained which can then be applied in a real surgery. This way, the

patient's complaints can be explored and treated individually. The simulation of the complete respiratory cycle can be simulated on the Cray XE6 Hermit system at HLRS, depending on the accuracy of the solution, within hours and days, which makes this method being within reach for the application in real world clinical environments.

## From Billions of Cubes to a Simulation

The flow solver ZFS (Zonal Flow Solver) for the simulation of the respiration cycle is developed at the AIA and works on Cartesian grids. This kind of computational meshes allows a fully automatic grid generation, which is a great advantage over grid generation methods producing boundary fitted meshes. Such meshes have to be set up manually and their configuration

stored in a three-dimensional mesh. The extraction process involves several extraction procedures itself and is depicted in Fig. 2. In a preprocessing step, it is sometimes required to apply an image enhancement filter to the CT-data. Unsharp boundaries between tissue and air complicate the extraction of the surface and hence need to be sharpened with a convolution filter, which applies a filter mask to each voxel involving the surrounding neighboring voxels. To detect the volume of interest, a seeded region growing algorithm recursively marks voxels depending on their density value as provided by the CT. A triangular representation is then generated from the segmentation by using the Marching Cubes algorithm, which produces a water-tight surface. The segmentation of the CT-data operates on binary data and hence the resulting surface contains unsmooth stairsteps. The surface is then smoothed by applying filter methods such as a windowed sinc function smoothing, which is a low-pass filter reducing the high frequent surface variations deflected by the stairsteps. Based on this surface
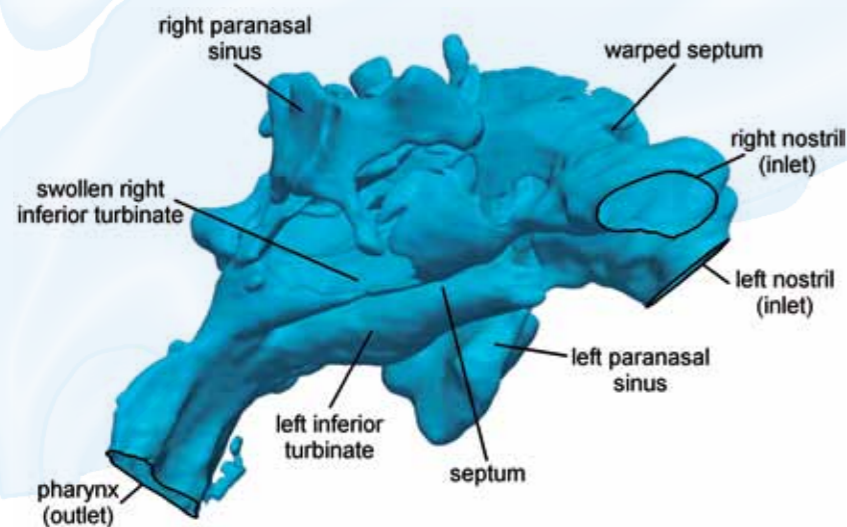


Figure 4: Surface of the nasal cavity extracted from CT-data. This patient has a warped septum and a swollen inferior turbinate on the right side.
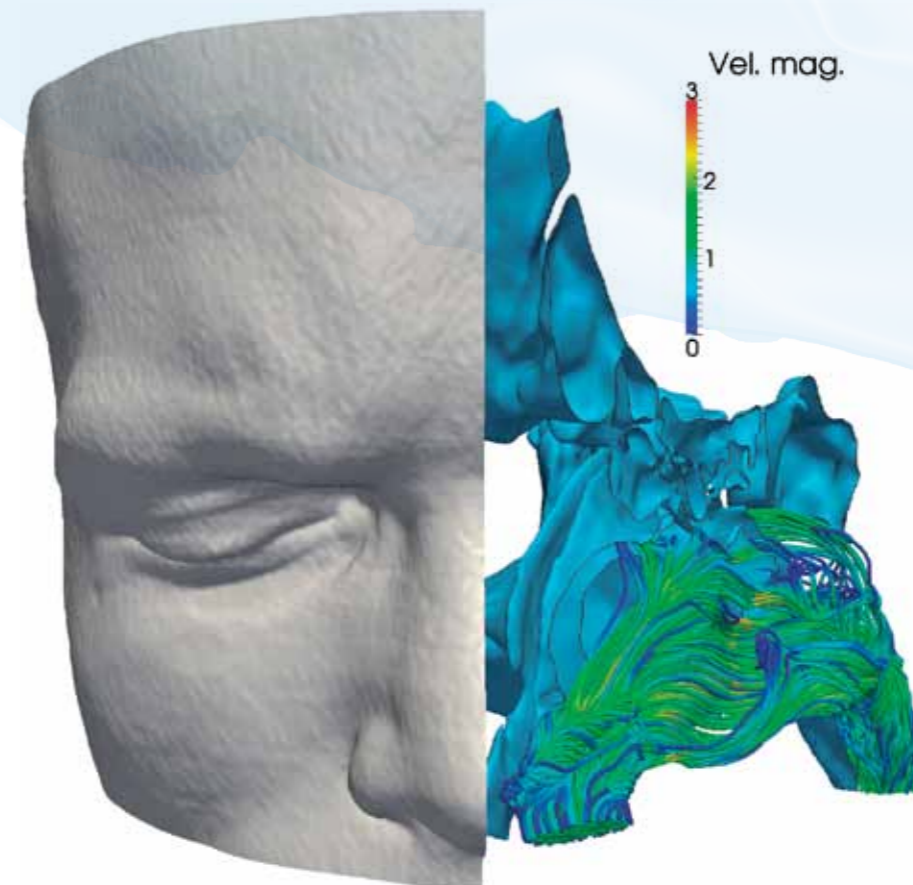


Figure 5: Streamlines of the nasal cavity flow colored by the local flow velocity. The bending in front shows the flow characteristics in the region of the inferior turbinate. A mixing region is found in the pharynx, where the left and right cavity merge.
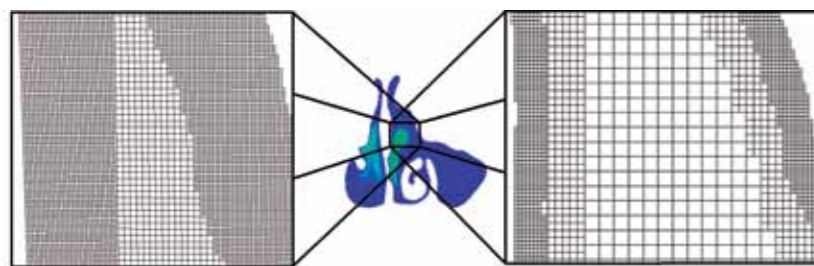
Figure 6: Local boundary refinement in the human nasal cavity. The mesh is refined based on the distance to the boundary. In the center, a frontal slice through the nasal cavity shows the location of the mesh detail.

for complicated geometries such as the human respiratory system would take several days and is hence not feasible. ZFS combines a parallel grid generator, a Finite Volume (FV) and a Lattice-Boltzmann Method (LBM) solver for the simulation of the fluid flow, a Lagrangian particle tracking and post-processing functionalities. As a base for the flow solvers the space needs to be split into small cubes to discretize the equations of fluid motion, i.e. the Navier-Stokes equations for the FV solver and the Boltzmann equation for the LBM solver, respectively. The grid generator is efficiently parallelized to be executed on thousands of processors on HPC systems, allows a resolution of billions of cells per nasal cavity and takes only seconds to generate on the
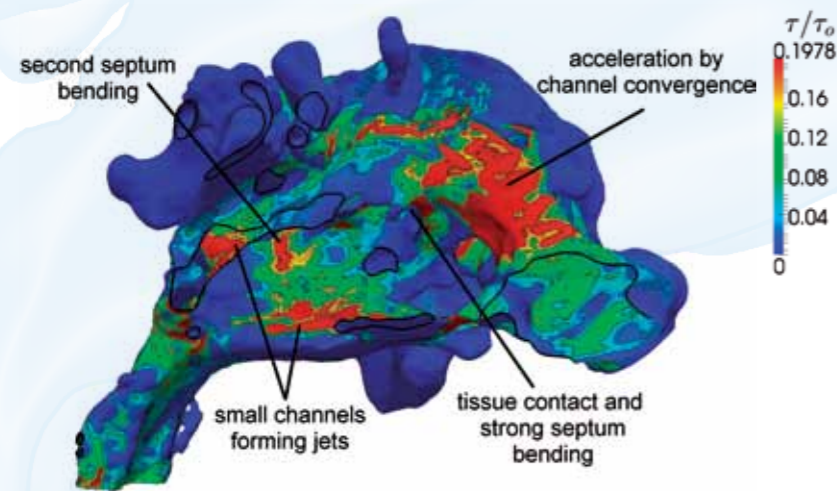


Figure 7: A mapping of the wall-shear stresses on the nasal cavity surface. Red regions indicate high shear stresses, which appear in convergent channels, generating higher pressure losses.

Cray XE6 Hermit system at HLRS. The parallelism is a major advantage over serial grid generators, which are strongly limited to the available memory on only one system and may take up to several days to generate a grid with a sufficient resolution for the simulations in the respiratory tract. Local grid refinement, as shown in Fig. 6, allows a high resolution in regions, where high gradients are expected, i.e. in wall-bounded shear layers and mixing zones and increase the accuracy of wall-shear stress calculations, which are of great importance in the analysis of nasal cavity flows. Local flow phenomena, such as shear layers also need a high resolution for an accurate prediction. Such refinement is obtained by dynamic solution adaptive refinement initiated during the simulation. For the simulation, the LBM solver is used since it is more efficient for the low Mach number flow as it appears in the respiration system. In the LBM, the Lattice-BGK equation, which is the discretized BGK-equation, a simplified form of the Boltzmann equation, is solved iteratively for so called particle probability density functions (PPDFs). In contrast to the Navier-Stokes equations, the Boltzmann equation is a statistical approach for the description of the rate of change of the number of particles in an infinitesimal small volume of fluid. The LBM solver contains two simple steps, the propagation and collision step, and can easily be parallelized on thousands of processors. The algorithm computes the new PPDFs in each cell for a certain number of discrete directions, i.e. for 14, 18 or 26 neighboring cells in three dimensions. After this collision step the information is exchanged between neighboring cells and is again used in the computation of the new PPDFs. The LBM is valid for

quasi-incompressible flows and hence only in the low Mach numbers regime. This leads to a decoupling of the energy equation from the Navier-Stokes equations. To additionally solve for the temperature field, a one-directional coupled scalar transport equation is integrated, in which the convection of the scalar is controlled by the velocity. With this procedure all primitive variables, velocity, density, pressure and temperature can be obtained in each time step in every grid cell.

## The nasal Cavity, a highly complex and individual Organ

Several aspects play a major role in specifying the quality of a nasal cavity. One of the properties governing the respiration efficacy is the static pressure loss resulting from the shape of the cavity. Warped septums and swollen turbinates are typical diagnoses of genetic endowments and allergies, respectively. Within the DFG-funded project "Rhinomodel", simulations of fluid flows in the nasal cavity have shown to be responsible for an increased pressure loss in such pathologies. The formation of recirculation zones inhibiting the direct access to the pharynx is analyzed by considering the characteristics of streamlines as shown in Fig. 3 and Fig. 5. Here, the streamlines are colored by the velocity of the fluid and show strong curvatures in regions of high vorticity. This form of visualization also supports the investigation of the olfactory capability as fluid not directed along the olfactory organ suggests a diminished sense of smell since the olfactory receptors cannot be reached. The forces acting on the nasal cavity tissue generating high pressure losses are investigated by mapping the wall-shear stresses, which are highest in regions of strong velocity gradients
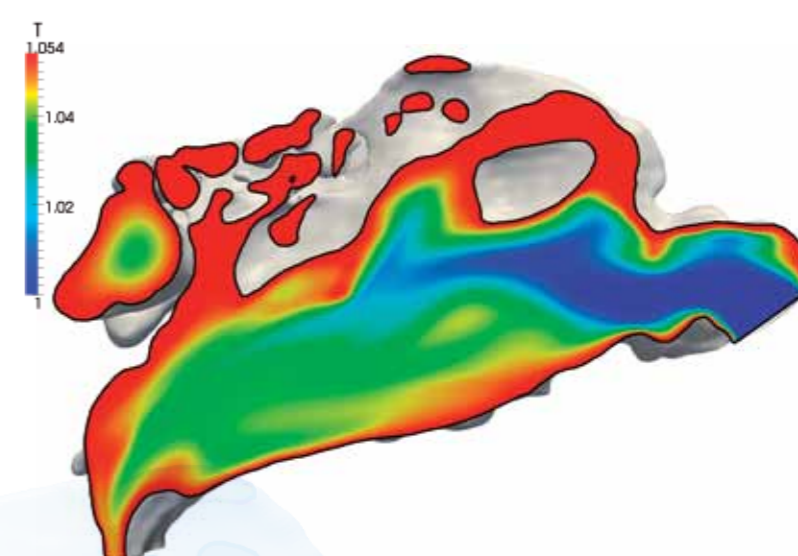


Figure 8: Simulation of the temperature distribution in the human nasal cavity. The inflow condition is set to 20° C and the wall temperature to 37° C body temperature. The inhaled air is heated up depending on the heating capability of the nasal cavity and reaches almost body temperature in the pharynx under optimal conditions.

near the wall, onto the geometry. Fig. 7 shows an example of such a visualization, where the stresses are highest (red) in highly converged channels and in zones of increased vorticity. To accurately compute these stresses, a high grid resolution is required, which is why local grid refinement plays a fundamental role in the estimation of the wall shear-stresses. The temperature increase is important for attaining optimal conditions for the lower airways and is considered by investigating the passive scalar transport from the nostrils, with a boundary condition of 20°C ambient temperature and heated up by the 37°C tissue temperature by diffusion and convection, to the pharynx. Healthy cavities show an increase of the temperature to almost body temperature while pathological cavities often exhibit a reduced heating capability. An example of the temperature distribution in a nasal cavity is shown in Fig. 8. Another aspect not being verified evidently in literature is whether the flow in the nasal cavity is
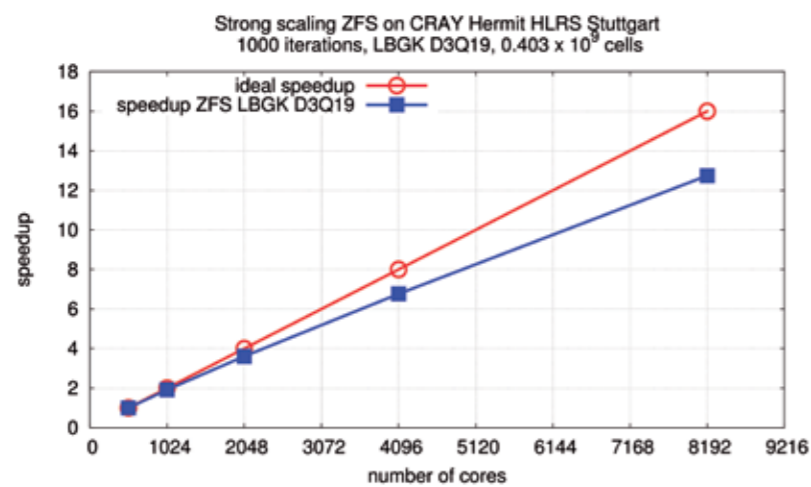
Figure 9: Speedup results for a stong scaling test for a total grid size of 0.403 x 10$^9$ cells for the LBGK D3Q19 algorithm on the Cray XE6 Hermit system at HLRS Stuttgart.
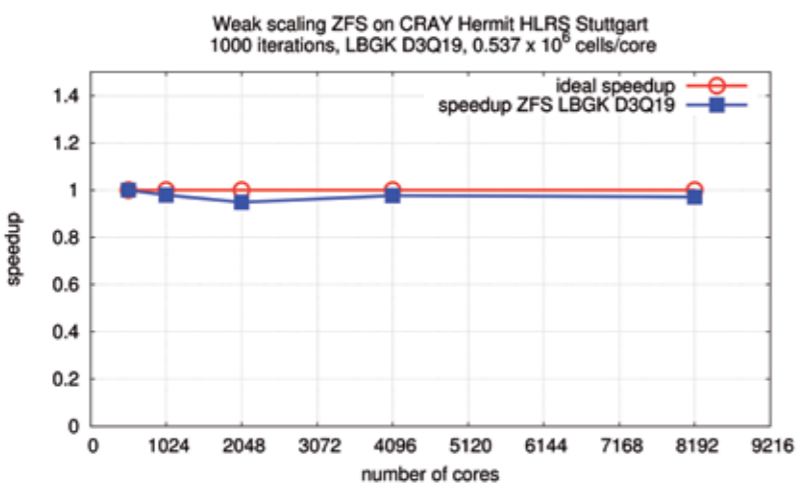


Figure 10: Speedup results for a weak scaling test for a constant local grid size of 0.537 x 10$^6$ cells for the LBGK D3Q19 algorithm on the Cray XE6 Hermit system at HLRS Stuttgart.

completely laminar. The type of flow is especially important for investigations in the lower airways, where at the transition from the pharynx to the larynx accurate boundary conditions have to be set. Recent investigations at the AIA show that the flow can indeed become transitional, depending on the local shape of the geometry and the local properties of the flow. Fig. 11 shows a strong production of vortices after the unification of the nasal cavity chan-

nels near the pharynx. Such zones are potential candidates of transitional to turbulent flow and are currently under detailed investigation.

## Demand for highly optimized Code and future Requirements

The solution algorithm for the simulation of the nasal cavity flow should produce results within a reasonable time, especially when it is used in a virtual surgery environment. Therefore, it is important to have a highly optimized code, which is implemented on HPC systems. In cooperation with computer science specialists from HLRS Stuttgart and CRAY, ZFS has recently been optimized to yield a higher performance on the High Performance Computing systems of HLRS, with a special focus on Cray XE6 Hermit. It is important to note, that the emphasis is not only placed on the reduction of the computation time, but also on the reduction of the communication time. Fig. 9 shows the results of a strong scaling experiment on Cray XE6 Hermit, where for a constant problem size of 0.403 x 10$^9$ cells the number of processes is continuously increased from 1,024 to 8,192 and the speedup is measured. The small decrease in speed for 8,192 processes is due to the increase of the communication overhead. Instead of keeping the overall problem size constant, the problem size per process can be kept constant which leads to a weak scaling, for which the results in Fig. 10 show an almost perfect scalability. For this experiment a maximum number of cells of 10$^9$ has been reached. Currently, the scalability to even higher processor numbers up to the complete Cray XE6 system is in-
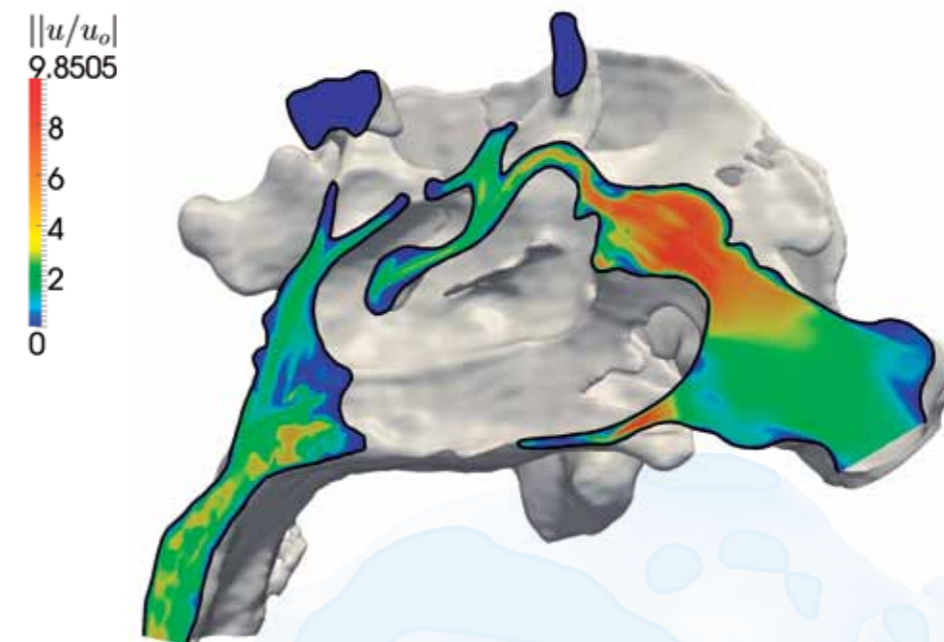


Figure 11: The velocity distribution in the right nasal cavity shows a strong production of vorticity where two nasal cavity channels merge and the flow is directed towards the pharynx.

vestigated and first tests show very promising results. Future research at the AIA includes investigations of the aeroelastic problem which occurs during sleep apnea and also snoring. In this case an interaction of the elastic tissue with the flow field occurs, where the force exerted by the flow leads to a deformation of the tissue which in turn impacts the flow. Such problems require a prediction of unsteady flow fields coupled with a finite element solver for the tissue structure and increases the problem size and the computational effort considerably. These simulations will therefore be ideal candidates for the HPC systems of HLRS.

## Acknowledgments

• Andreas Lintermann

Chair of Fluid Mechanics and Institute of Aerodynamics RWTH Aachen University

# Numerical Investigation of VFE-2 Delta-Wing

# the Flow Field about the

The industrial application of Delta-Wings is manifold and reaches from the classical aerospace engineering, e.g. highly agile aircraft, aerodynamic devices or control surfaces, to unique environmental technologies, such as devices for snow clearance. In all cases the development of leading edge vortices is exploited.

However, steadiness and stability of these leading edge vortices is essential for controllability, particularly for highly agile aircraft. It is well known that vortices can undergo a sudden expansion often related to vortex break-down [5].
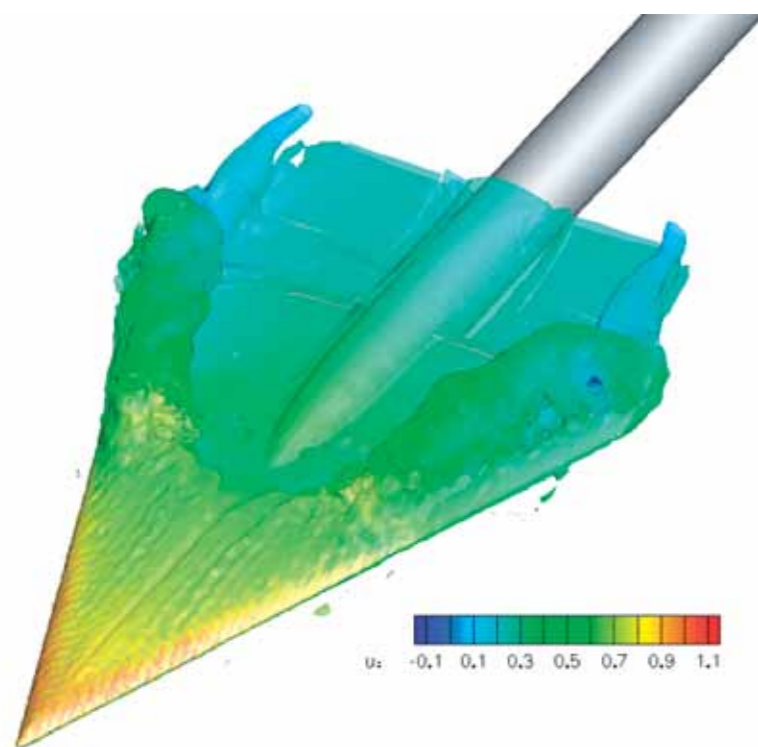


Figure 1: Isosurface of streamwise vorticity colored by streamwise velocity for the VFE-2 Delta Wing at an angle of attack of 13° and a Reynolds number of 0.5 million.

The occurrence of unsteady vortex breakdown is critical for aircraft. It is physically not fully understood, thus further investigation is required. The ongoing investigation is funded by the DFG project "Numerische Untersuchung der instationären Strömung um generische schlanke Deltaflügel" (DFG-B-506/2).

The international Vortex Flow Experiment 2 (VFE-2) Delta-Wing is taken as a generic aerodynamic configuration for which small angles of attack already lead to the development of leading edge vortices.

## Underlying Theory

A profound understanding of vortex formation and breakdown requires a comprehensive insight into the complete unsteady flow field. This insight can only be obtained from time-accurate simulations accompanied by experiments.

The most complete description of flows in continuum mechanics is given by the Navier-Stokes equations, which describe the exchange of momentum in the fluid considering friction.

Solving the Navier Stokes equations requires very high spatial and temporal resolution. A Direct Numerical Simulation (DNS) is still not feasible for complex turbulent flows in industrial applications due to the required tremendous computational resources. For simulating the turbulent flow about the VFE-2 Delta-Wing several hundred CPU years using 1 Terra flops would be needed, making

a simplification of the Navier Stokes equation inevitable.

Commonly used in industry is the RANS (Reynolds-Averaged Navier Stokes) simulation with appropriate statistical turbulence models. This simplified approach often fails to accurately predict separated and reattached flows. Also in the case of the VFE-2 Delta-Wing results do not compare well with the existing experimental results [9].

Better results are expected from Large-Eddy Simulation (LES). In LES only the large flow structures are resolved while small, stochastic structures are modelled with the help of SubGrid Scale (SGS) models. In Implicit Large Eddy Simulation (ILES) the truncation error of the discretization of the convective terms is deliberately tailored to act as a SGS-model which is therefore implicit to the discretization. One implementation of ILES is the Adaptive Local Deconvolution Method (ALDM) [1,4] which has shown considerable potential for the efficient representation of physically complex flows in generic configurations, such as isotropic turbulence, turbulent channel flow with periodic constrictions, turbulent boundary layer separation and turbulent cylinder flow [7]. In this project the ILES approach will be used for the essential numerical investigations.

## Numerical Method

The Adaptive Local Deconvolution Method (ALDM) turbulence model is

incorporated in a solver for the incompressible Navier-Stokes equations with constant density. Continuity is ensured by the pressure-Poisson equation. The equations are discretized on a staggered Cartesian mesh allowing for an easy control of the truncation error and offering superior computational efficiency compared to body-fitted grids [6]. Bounding surfaces of the flow that are not aligned with the grid are accounted for by the Conservative Immersed Interface (CIIM) approach [8].

For time advancement an explicit third-order Runge-Kutta scheme is used. The pressure-Poisson equation and diffusive terms are discretized by second-order centred differences.

Further improvement of efficiency is achieved by modelling the turbulent boundary layer using a wall model [2] and by locally adapting the mesh resolution with Local Mesh Refinement, such as used in [7]. In this context new criteria for the Local Mesh Refinement algorithm based on physical criteria are applied. A further considerable reduction of the number of computational cells is achieved, rendering the simulation much more effective.

## Results

The results in this reporting period have been obtained for a Reynolds number Re = 0.5 million based on root chord length and an angle of attack of 13°. The objective is to reach a qualitative
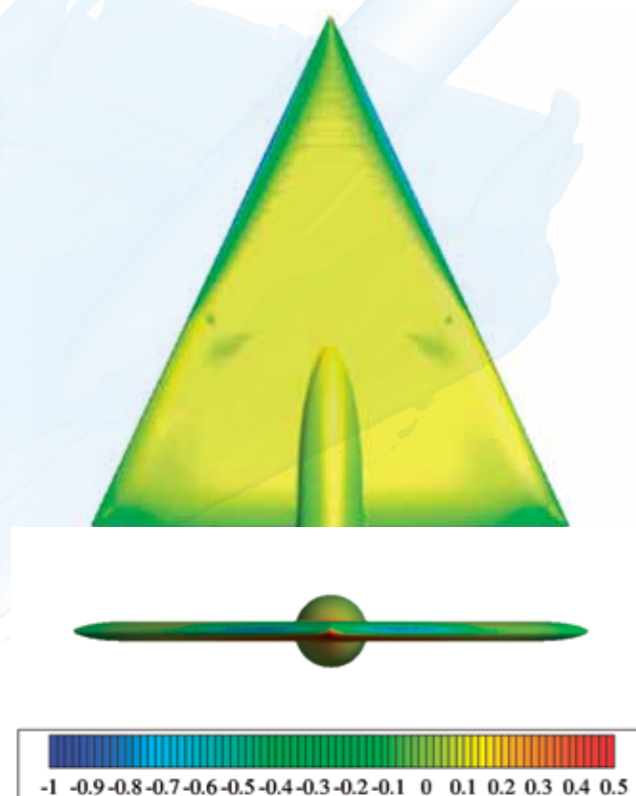
Figure 2: Cp distribution on the upper (top) and front (bottom) surface of the VFE-2 Delta Wing at an angle of attack of 13° and a Reynolds number of 0.5 million.
Qualitative results for this specific configuration look very promising.
A profound understanding is expected after results have been quantitatively validated with the experimental data. The necessary simulations are on-going.

| AoA [°] | Re[-]in mio | LE | Ncells in mio | NCPU | CPUh in mio |
|---|---|---|---|---|---|
| 13 | 0.5 | MR | 20 | 1024 | 0.9 |
| 13 | 1 | MR | 20 | 1024 | 0.9 |
| 18 | 1 | MR | 40 | 1024 | 1.2 |
| 18 | 1 | S | 40 | 1024 | 1.2 |
| 23 | 1 | MR | 40 | 1024 | 1.2 |
| 23 | 1 | S | 40 | 1024 | 1.2 |

Table 1: Survey of conducted simulations (black) and planned simulations (grey). (AoA = Angle of Attack, Re= Reynolds number, LE=Leading Edge, number of cells, number of CPUs, number CPU hours)

agreement with the respective experiments of Furman and Breitsamter [3].

For this angle of attack both, experiment [3] and the current simulation, show vortex formation over half chord length, see Fig. 1. Close to the apex the boundary layer flow accelerates over the leading edge and undergoes laminar-turbulent transition. This is also reflected by the high suction levels visible in Fig. 2. Severe pressure gradients in lateral direction provoke boundary layer separation further downstream. The separation region is indicated by the low pressure regions on the upper surface of the wing (see Fig. 2).

## Computational Details
For this test case again a SGI-ALTIX with ItaniumII processors have been used. CIIM and the wall modelling con-

sume 2% of the overall computational time, the flux calculation with ALDM accounts for 7% and the Poisson solver uses 83%.

Table 1 summarizes further computational details for the present and planned simulations.

Table 1: survey of conducted simulations (black) and planned simulations (grey). (AoA = Angle of Attack, Re= Reynolds number, LE=Leading Edge, number of cells, number of CPUs, number CPU hours)

## On-going Research/Outlook
Since the objective is to investigate vortex bursting process, further investigation at the angles of attack of 18° and 23° is with varying leading edge geometries, i.e. MR= Medium Round and S= Sharp, are planned. As a next step an investigation is planned for actively controlling and preventing the vortex burst with leading edge devices, i.e. oscillating control surfaces.

## References
[1] **Adams, N.A., Hickel, S., and Franz, S.**
Implicit subgrid-scale modeling by adaptive local deconvolution J. Comput. Phys. 200 (2), 412-431, 2004

[2] **Chen, Z., Devesa, A., Meyer, M., Hickel, S., Lauer, E., Stemmer, C., and Adams, N.A.**
Wall Modelling for Implicit Large Eddy Simulation of Favourable and Adverse Pressure Gradient Flows, Proceedings of WALLTURB Workshop, Lille, Springer, Heidelberg, 2009

[3] **Furman, A., and Breitsamter, C.**
Experimental investigations on the vfe-2 configuration at TU Munich, Germany, RTO-TR-AVT RTO-TR-AVT- 113, 21, 1-47, 2008

[4] **Hickel, S., Adams, N.A., and Domaradzki, J.A.**
An adaptive local deconvolution method for implicit LES J. Comput. Phys. 213 (1), 413-436, 2006

[5] **Ludwieg, H.**
Zur Erklärung der Instabilität der bei angestellten Deltaflügeln auftretenden freien Wirbelkerne, Z. Flugwiss. 10 (3), 242-249, 1962

[6] **Meyer, M., Hickel, S., and Adams, N.A.**
Computational Aspects of Implicit LES of Complex Flows, in: High Performance Computing in Science and Engineering, Garching/Munich, Springer, Heidelberg, 2010

[7] **Meyer, M., Hickel, S., and Adams, N.A.**
Asessment of Implicit Large-Eddy Simulation with a Conservative Immersed Interface Method for Turbulent Cylinder Flow, Int. J. Heat Fluid Flow 31 (3), 368-377, 2010

[8] **Meyer, M., Devesa, A., Hickel, S., Hu, X.Y., and Adams, N.A.**
A Conservative Immersed Interface Method for Large-Eddy Simulation of Incompressible Flows, J. Comput. Phys. 229 (118), 6300-6317, 2010

[9] **Morton, S., Forsythe, J., Mitchell, A., and Hajek, D.**
Detached-eddy simulations and Reynolds-Averaged Navier-Stokes simulations of delta wing vortical flow field, Trans. ASME 124 (4), 924-932, 2001

• Michael Meyer
• Stefan Hickel
• Nikolaus A. Adams

Institute of Aerodynamics and Fluidmechanics, Technische Universität München

# Simulating Blood Cells and Blood Flow

Blood performs a multitude of functions on its way through our body, from the transport of oxygen to the immune response after infections. In addition, the circulatory system may be also affected by injuries which cause bleeding, by the formation of plaques in arteries which cause coronary heart disease, and it provides the pathway for the organism invasion by bacteria or viruses. Thus, modeling of blood flow and its functions is an important challenge with many medical implications, but also with many interesting physical phenomena.

Blood is a complex fluid, which consists of blood plasma – a liquid comprised of water, small molecules and various proteins – and blood cells. There are several types of blood cells, which include red blood cells (RBCs), white blood cells (WBCs) and platelets, as shown in Fig. 1, with typical sizes of about

8µm, 10µm, and 2µm, respectively. The blood cells together take up about 45% of the blood volume. However, they appear in very different concentrations: one microliter of blood contains of about 5 million RBCs, 5,000 WBCs, and 250,000 platelets. Therefore, the flow properties of blood are dominated by the highly abundant RBCs [1,2].

## Mesoscopic Modeling of Blood

The numerical modeling of blood flow in the circulatory system is very challenging due to the large range of length scales involved, ranging from the nanometer size of proteins and the micrometer size of blood cells to the centimeter size of large vessels (arteries and veins) and the decimeter size of the heart. Therefore, no single simulation technique is able to cover all required length scales. On the macroscopic level, appropriate for studying blood flow in the heart and in large vessels, Computational Fluid Dynamics (CFD) – which describes blood as a continuum fluid using constitutive equations for the fluid properties – has been very successful. However, constitutive equations are not always accurate. Moreover, this approach fails when the particulate nature of blood becomes important, for instance, when modeling the blood flow in small vessels (arterioles, capillaries, and venules). Therefore, mesoscopic models and simulation approaches, which describe blood as a suspension of soft, deformable particles or cells have been developed in recent years.
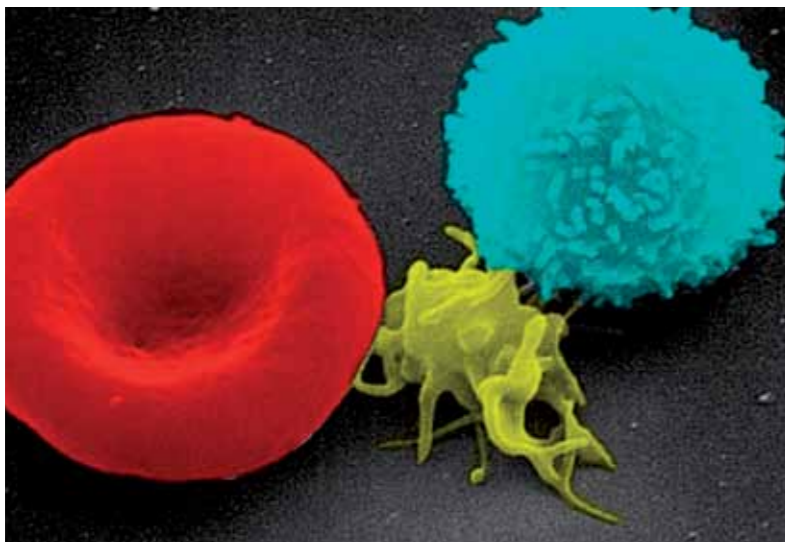


Figure 1: A scanning electron micrograph of blood cells. From left to right: red blood cell (red), activated platelet (yellow), and white blood cell (cyan). Source: The National Cancer Institute at Frederick (NCI-Frederick).
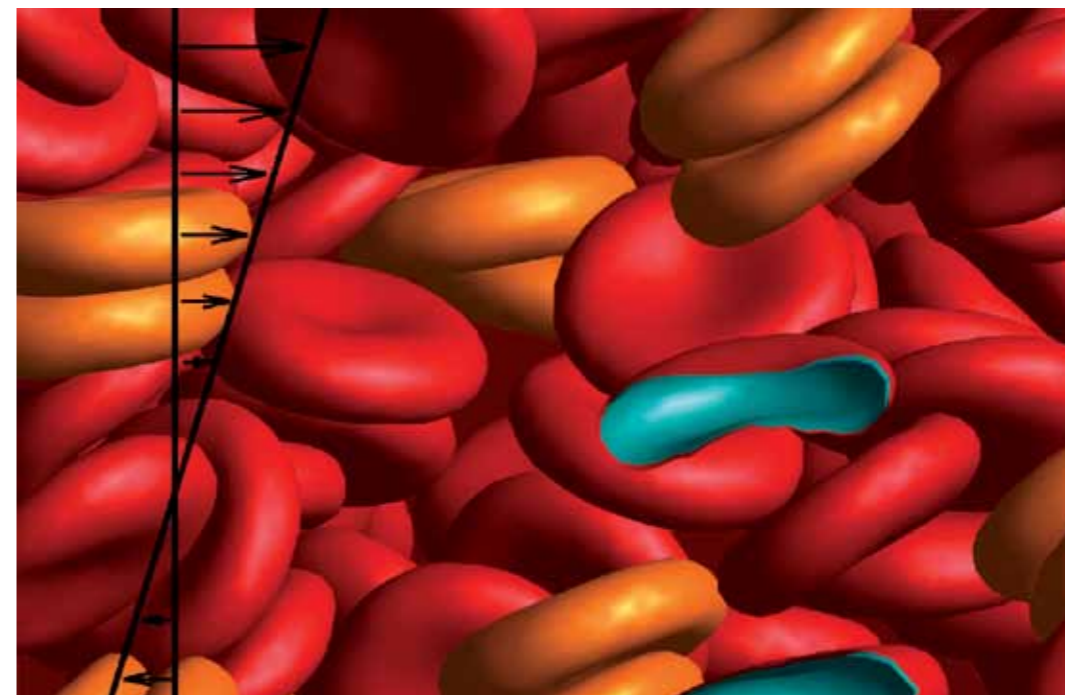


Figure 2: Simulation of blood under shear flow. RBCs are shown in red and in orange, where orange color depicts the rouleaux structures formed due to aggregation interactions between RBCs. The image also displays cut RBCs with the inside drawn in cyan to illustrate RBC shape and deformability.

Two main ingredients are necessary for the mesoscopic modeling of blood cells under flow, the membrane and the embedding fluid. The membrane of a RBC consists of a fluidic lipid bilayer, to which a biopolymer (spectrin) network is attached. This composite membrane can be described by a triangulated network of springs with bending and stretching elasticity [3,4]. For the fluid, mesoscale hydrodynamics simulation techniques – such as Dissipative Particle Dynamics (DPD) [5], Multi-Particle Collision Dynamics (MPC) [6], and the Lattice Boltzmann method (LBM) – have proven to be very advantageous, because these hydrodynamics techniques contain thermal fluctuations and can be easily coupled to the membrane. The mesoscale hydrodynamics approaches are also very suitable for running on massively parallel supercomputers, such as JUROPA and JUGENE at the Jülich Supercomputing Centre, due to the simplicity of parallel-

ization of the above methods. In order to highlight the strengths of mesoscale modeling to study blood flow, we focus further on the two recent examples: blood rheology and the margination of WBCs in microchannel flows.

## Blood Rheology

Rheology is the science of the behavior and properties of fluids under flow. The aim of modern rheology is the prediction of the macroscopic flow properties, such as the fluid viscosity, from the sizes, deformability, and interactions of the constituent molecules and particles. For complex fluids, the rheological properties can be very rich. A characteristic feature is that under shear flow, the fluid viscosity may not be constant, but depend on the shear rate. Such fluids are called "non-Newtonian". Blood is non-Newtonian and exhibits shear-thinning, i.e. with increasing shear rate, the viscosity of blood decreases.
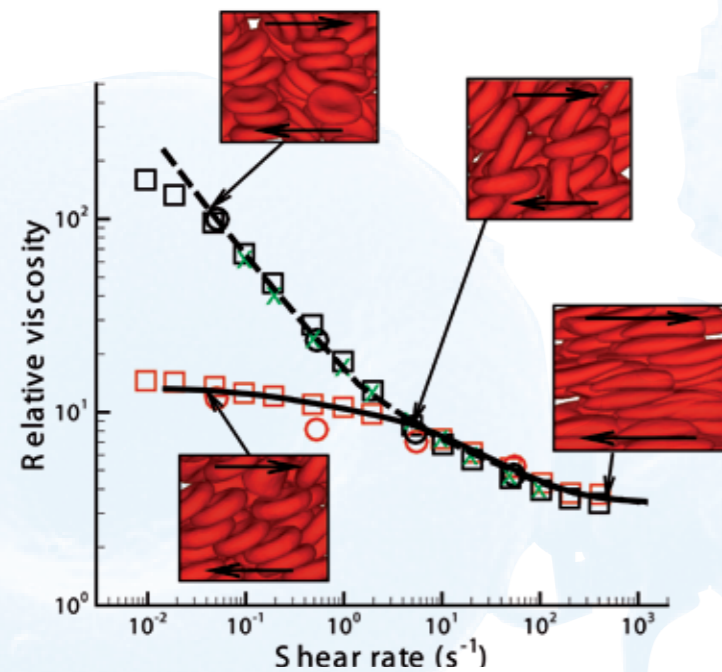
Figure 3: non-Newtonian relative viscosity (the cell suspension viscosity normalized by plasma viscosity) of whole blood as a function of shear rate at $H_t$ =0.45 and 37°C. Experimental data are shown for whole blood (green crosses, black circles, and black squares) and for non-aggregating RBC suspension (red circles and red squares). Simulation results [7] are shown for the both cases with black dashed and solid lines, respectively.

In order to predict the shear-thinning behavior, we study RBC suspensions in shear flow illustrated in Fig. 2, which are characterized by the RBC volume faction $H_t$, called the "hematocrit", and the shear rate $\dot{\gamma}$. The simulation results for the relative viscosity $\eta/\eta_s$ (where $\eta_s$ is the plasma viscosity without cells) are displayed in Fig. 3. The two curves in Fig. 3 are for whole blood, for which proteins in the blood plasma (e.g., fibrinogen) mediate attractive interactions which at low shear rates lead to the formation of cylindrical stacks of RBCs called "rouleaux", and for RBCs in Ringer solution, where these proteins have been removed and the interactions are purely repulsive. In both cases, blood is found to be a shear-thinning fluid, an important property to reduce the load on the heart. However, the shear thinning is much

more pronounced for whole blood due to the formation of rouleaux structures at low shear rates, which break up into smaller aggregates and finally single cells with increasing shear rates. Above a shear rate of approximately 5s$^{-1}$, the viscosity curves for whole blood and Ringer solution merge, and therefore we can conclude that rouleaux are absent at higher shear rates. In Fig. 3, the comparison of simulation and experimental results shows excellent agreement.

What can we now learn from simulations, which has not been already known from experiments? First, the good agreement between simulations and experiments for the viscosity of whole blood is obtained by tuning the strength of the aggregation interactions between RBCs. Thus, we are able to predict the attractive force between RBCs to be in the range of 3 to 7 pico-Newton, 10 million times smaller than the weight of a mosquito! This force is below the resolution of current experimental techniques. Second, Fig. 3 shows a second shear-thinning regime above a shear rate of about 5s$^{-1}$, where aggregation interactions are not relevant. This regime is due to deformations and orientations of RBCs, as revealed by a more detailed analysis of the simulated RBC shapes. In the simulations, the shapes are characterized by the eigenvalues $\lambda_j$ of the gyration tensor, from which the asphericity $\alpha = (\Sigma_{i,j}(\lambda_i - \lambda_j)^2)/(2\Sigma_i\lambda_i)^2$ can be calculated (which yields $\alpha = 0$ for a sphere, $\alpha = 1$ for a long and thin rod). The simulation results for the asphericity distributions, displayed in Fig. 4, show that at low shear rates all RBCs maintain their equilibrium discocyte shape; at higher shear rates, in the second

shear-thinning regime, the distribution broadens and shifts to smaller $\alpha$ values, indicating strongly deformed and more rounded RBCs; finally, at very high shear rates, the distribution shifts to large values of $\alpha$, indicating prolate-shaped RBCs, which are elongated and oriented by the flow. Thus, the simulations provide new insights into the deformation and dynamics of single cells, but also their correlations and interactions, under flow.

## White Blood Cell Margination

WBCs in our organism take part in the defense against various infections. Experimental observations revealed an interesting behavior of WBCs in blood flow such that they migrate towards the vessel walls, a process called "margination". This phenomenon appears to be biomedically very important, since WBCs need to firmly adhere to vessel walls to perform their function, and thus first they have to closely approach or touch the walls, which is facilitated by their margination.

WBC margination is governed by hydrodynamic interactions of cells with the vessel walls and with each other in blood flow. In fact, RBCs play a major role in WBC margination, since RBCs experience a stronger hydrodynamic lift force [8] away from the walls – due to their discoidal shape and deformability – than WBCs, which are nearly spherical in shape and resistant to deformations. As a consequence, RBCs migrate to the center of a vessel and effectively push out WBCs to the vessel walls.

WBC margination strongly depends on blood cell deformability, hematocrit $H_t$, local flow rate, and RBC aggregation. Therefore, numerical modeling of blood
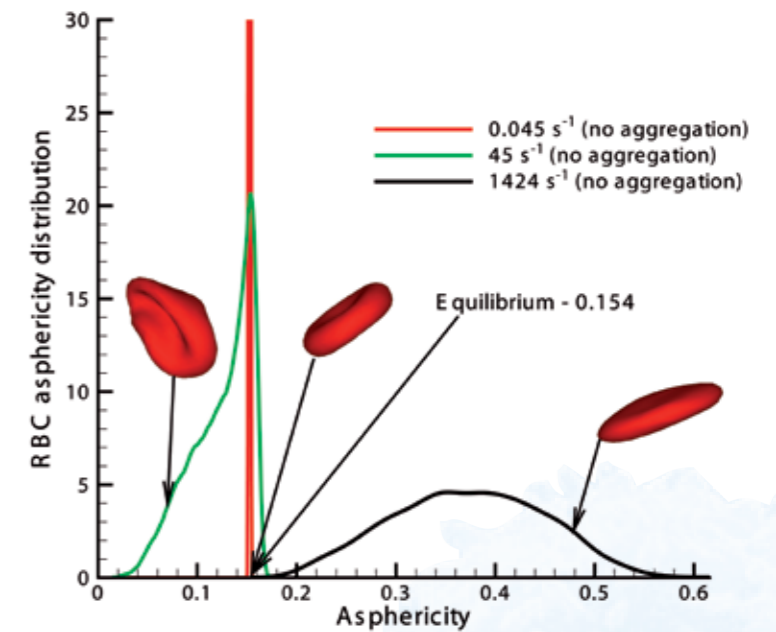


Figure 4: RBC asphericity distributions to describe cell deformations through the deviation from a spherical shape. The asphericity for a RBC in equilibrium is $\alpha$ =0.154.

flow provides an excellent opportunity to study WBC margination in detail for various flow and cell properties. To explore the effect of many parameters on WBC margination, we employ a two-dimensional (2D) blood flow model shown in Fig. 5, which significantly reduces computational cost, however, captures essential physics. In 2D, the cells are modeled using a bead-spring representation illustrated in Fig. 5 [9]. The flow rate in the channel is characterized by a dimensionless shear rate $\dot{\gamma}^* = \bar{\dot{\gamma}}\tau$, where $\bar{\dot{\gamma}}$ is the average shear rate and $\tau$ is a characteristic RBC relaxation time.

Fig. 6 shows a WBC margination diagram for various hematocrits and flow rates. The margination probability is defined as a probability for a WBC to be closer than 500nm to the channel walls. Clearly, efficient WBC margination occurs only within an intermediate range of hematocrits $H_t$=0.2-0.5 and flow rates $\dot{\gamma}^*$= 1 - 10 .

The simulation results are consistent with experimental observations that WBC adhesion is found mainly in venules (but not in arterioles) in the organism. The characteristic values of $\dot{\gamma}^*$ in venules are low enough, while in arterioles $\dot{\gamma}^* \gtrsim 30$ [2]. Thus, efficient WBC margination and consequent adhesion are mainly expected in the venular part of microcirculation. Our simulations also showed that RBC aggregation enhances WBC margination, while WBC deformability attenuates this effect [9].

of other cell and capsule suspensions, blood flow in diseases (e.g., malaria, sickle-cell anemia), biomedical applications (e.g., lab-on-chip, microfluidics), etc. The main advantages of accurate modeling of such suspensions in comparison to experimental tests are robustness and low cost of numerical simulations. The future growth of supercomputing resources and further advances in cell modeling will make blood simulations a versatile and widely available tool in biophysical and biomedical research and applications.
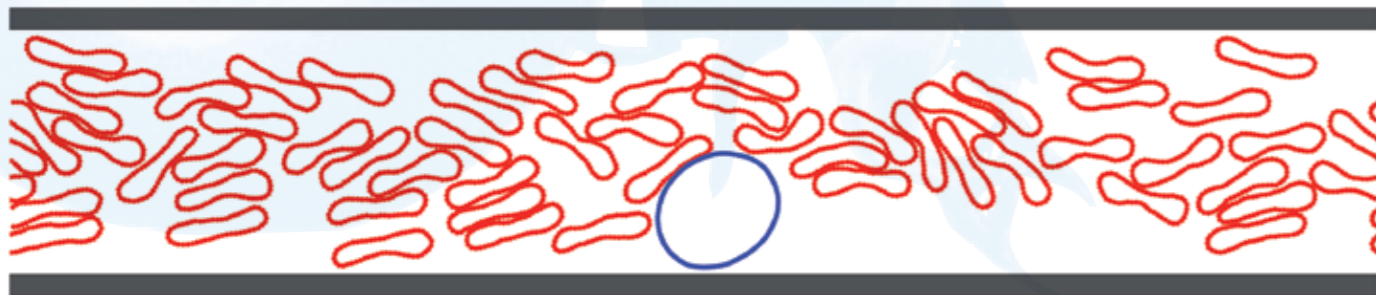
Figure 5: Simulation snapshot of the flow (from left to right) at $H_t$=0.35 and dimensionless shear rate $\dot{\gamma}^*$= 1.59 . The large quasi-spherical particle is the WBC.

## Conclusions

The above models and results demonstrate a promising role of numerical modeling to understand and predict blood flow and its related processes. The potential of these models extends not only to RBC suspensions under various conditions, but also to modeling

## Acknowledgement

## References

[1] **Pries, A.R., Secomb, T.W., and Gaehtgens, P.**
Biophysical aspects of blood flow in the microvasculature, Cardiovasc. Res., 32:654–667, 1996

[2] **Popel, A.S., and Johnson, P.C.**
Microcirculation and hemorheology, Annu. Rev. of Fluid Mech., 37: 43–69, 2005

[3] **Noguchi, H., and Gompper, G.**
Shape transitions of fluid vesicles and red blood cells in capillary flows, Proc. Natl. Acad. Sci. USA, 102:14159–14164, 2005

[4] **Fedosov, D.A., Caswell, B., and Karniadakis, G.E.**
A multiscale red blood cell model with accurate mechanics, rheology, and dynamics, Biophys. J., 98:2215-2225, 2010

[5] **Pivkin, I.V. , Caswell, B., and Karniadakis, G.E.**
Dissipative particle dynamics, in: K.B. Lipkowitz, editor, Reviews in Computational Chemistry, volume 27, pages 85–110, John Wiley & Sons, Inc., Hoboken, NJ, USA, 2010

[6] **Gompper, G., Ihle, T., Kroll, D.M., and Winkler, R.G.**
Multi-particle collision dynamics: a particle-based mesoscale simulation approach to the hydrodynamics of complex fluids, Adv. Polym. Sci., 221:1–87, 2009

[7] **Fedosov, D.A., Pan, W., Caswell, B., Gompper, G., and Karniadakis, G.E.**
Predicting human blood viscosity in silico, Proc. Natl. Acad. Sci. USA, 108:11772–11777, 2011

[8] **Messlinger, S., Schmidt, B., Noguchi, H., and Gompper, G.**
Dynamical regimes and hydrodynamic lift of viscous vesicles under shear, Phys. Rev. E, 80:011901, 2009

[9] **Fedosov, D.A., Fornleitner, J., and Gompper, G.**
Margination of white blood cells in micro-capillary flow. Phys. Rev. Lett., 108:028104, 2012
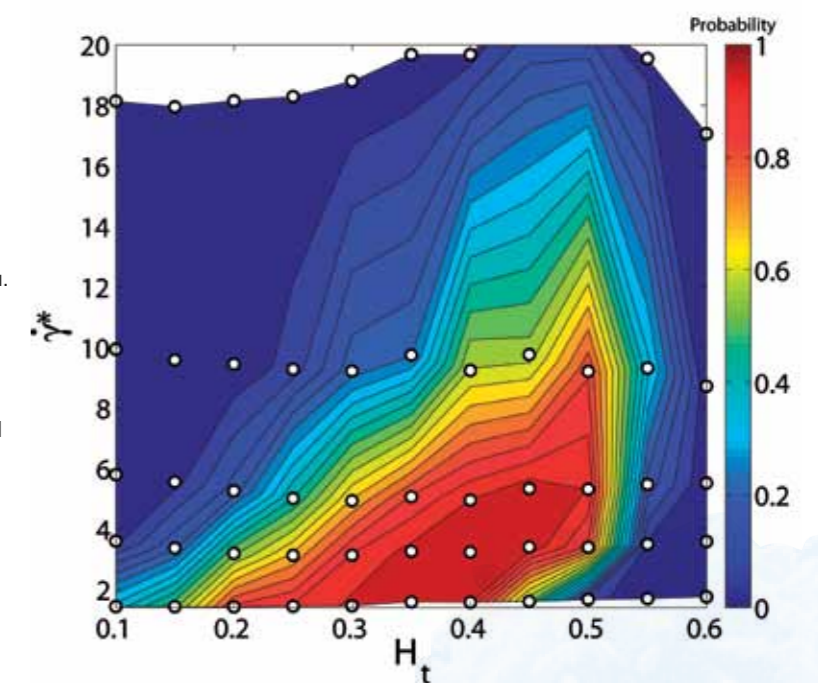
Figure 6: Probability diagrams of WBC margination with respect to $\dot{\gamma}^*$ and $H_t$. Symbols (o) indicate the values of $H_t$ and $\dot{\gamma}^*$ for which simulations were performed.

• Dmitry A. Fedosov*
• Gerhard Gompper

Theoretical
Soft Matter and
Biophysics
Institute of
Complex Systems
and Institute for
Advanced Simulation

Forschungszentrum
Jülich

* Dmitry Fedosov has recently received a Sofia Kowalewskaja Award by the Alexander von Humboldt Foundation for numerical studies of blood flow in health and disease.

# Direct Numerical Simulations of turbulent Rayleigh-Bénard Convection

The large structures occurring in the fluid flow on the Sun's surface, in the atmosphere and oceans of planets, including our Earth, are primarily driven by convection. Their actual shape but also the efficiency of the heat transport is however significantly influenced by the Coriolis force due to rotation. Crystal growth and the ventilation of buildings and aircrafts originate within the same physical framework. Understanding these fundamental processes is thus not only utterly important for geo- and astrophysics, but also in industry.

That is, where the idealization, the so-called Rayleigh-Bénard convection comes into play: The above mentioned highly complex phenomena are simplified to a fluid heated from below and cooled from above, with solely gravity, and hence buoyancy, acting on it.

In most of the theoretical and numerical investigations of natural convection the Oberbeck-Boussinesq (OB) approximation is considered. This means that all physical properties of the fluid are assumed to be independent of temperature and pressure, except the density in the buoyancy term which is considered linearly dependent on the temperature. These assumptions are evidently never fulfilled in reality and the deviations of the flow characteristics due to their violation are called non-Oberbeck-Boussinesq (NOB) effects.

Thus the objective of our studies is to investigate the influence of rotation and NOB effects on turbulent thermal convection, also, but not exclusively, at very high Rayleigh numbers. We address these issues by means of Direct Numerical Simulations (DNS).

## Numerical Method

We performed high-resolved direct numerical simulations of Rayleigh-Bénard convection making use of the well-tested fourth order finite volume code FLOWSI [2]. The code solves the incompressible Navier-Stokes equations and has shown an almost perfect scalability on the HLRB II cluster (Fig. 1).

For the purpose of investigating NOB effects, the code has been advanced by taking temperature-dependent material properties into account. Furthermore, we incorporated a module to model the effects of rotation.
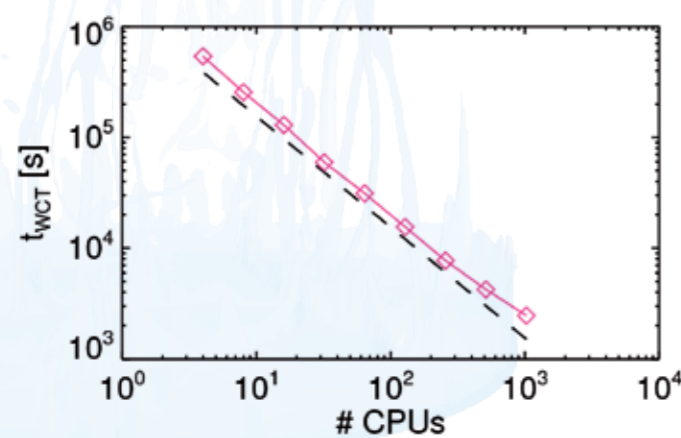


Figure 1: Scaling behaviour of the FLOWSI code (pink diamonds) on HLRB II. The perfect scalability of an ideal code is shown as well (black dashed line).
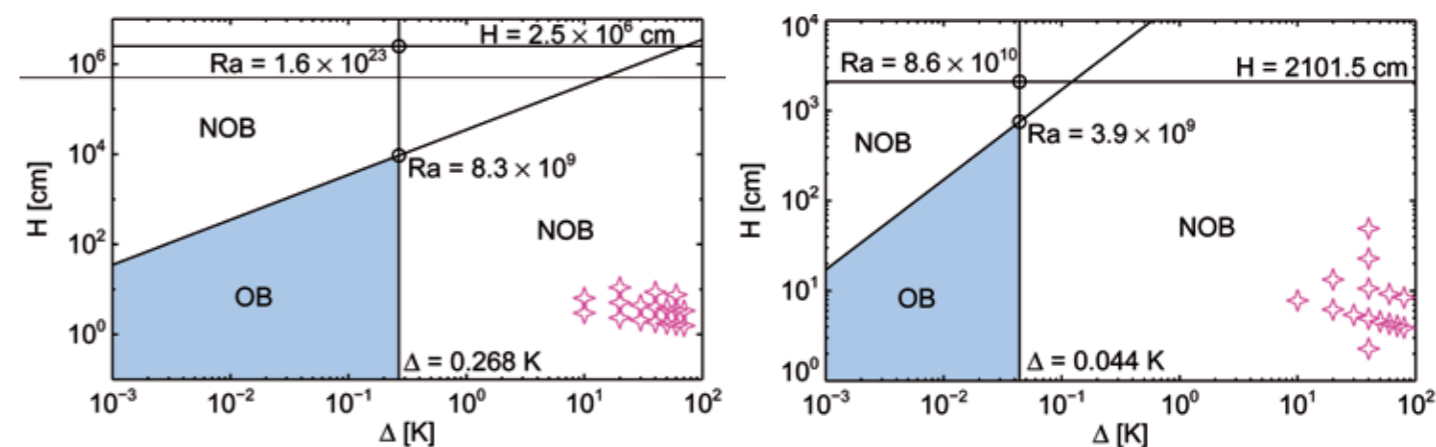


Figure 2: Validity range of the OB approximation according to Gray & Giorgini [2] for water (top) and glycerol (bottom). The pink stars mark the parameters for our NOB simulations.

The mesh size was chosen in a way to fulfil the criterion by Shishkina et al. [5], which guarantees the resolution of the smallest relevant scales, i.e. the Kolmogorov and the Batchelor scale.

So far, our simulations focussed on two fluids, water ($Pr = 4.38$) and glycerol ($Pr = 2547.9$). Their material properties were adopted from experimental data by Ahlers et al. [1]. The parameter ranges for our performed NOB simulations are given in Fig. 2, together with the estimation of the validity range of the OB approximation according to Gray & Giorgini [2]. In sum, we covered the range of Rayleigh numbers $10^5 \leq Ra \leq 10^9$ and NOB conditions up to a temperature difference $\Delta$ between the top and bottom plate of 80 K.

Self-evidently, we performed the corresponding OB simulations for the purpose of comparison as well. Furthermore, we also conducted DNS for different rotation rates under OB and NOB conditions.

However, not only the required resolution but also the necessary computation time to obtain reliable statistics, that is

at least several months, pushes us to the limits of nowadays supercomputers.

## Results

Our NOB simulations revealed that the temperature dependencies of the material properties are able to significantly influence the global flow structures. In general they lead to a breakdown of the top-bottom symmetry typical for OB simulations. The NOB effects that we observe, include, but are not limited to, different thermal and viscous boundary layers, asymmetric plume dynamics and an increase of the bulk temperature Tc (Fig. 3, 4). In glycerol for the NOB case $\Delta = 80$ K we obtain a Tc that is 15 K higher than the arithmetic mean temperature between the plates, whereas in contrast, in the case of water the observed increase of Tc is only about 5 K.

Nevertheless, the Nusselt number Nu and the Reynolds number Re and their scaling with Ra show only a slight deviation from the OB case.

In rotating convection another physical phenomenon plays an important role: the formation of columnar vorti-

ces, as seen in Fig. 3. These so-called Ekman vortices are able to extract hot and cold fluid from the bottom and top boundary layers, respectively, and thereby increase the heat transport. They are also responsible for a persisting mean temperature gradient in the bulk. Their size, and related to that their efficiency, is determined by the heat diffusivity and viscosity, thus they are sensitive to NOB effects. Indeed, while Tc is the same as in the non-rotating case, the absolute value of the temperature gradient is diminished. In line with this, the enhancement of Nu

under rotation is stronger under NOB than under OB conditions.

## On-going Research/Outlook

A next step on studying NOB effects will be the inclusion of compressible convective flows, i.e. almost all gases, where the Low Mach Number (LMN) approximation will be used, which generally admits the working fluid to be a non-perfect gas. And since in an astronomical and geophysical context Ra is typically rather in the order of $10^{20}$, we plan to go to higher Ra, which are of course computational more expensive,



Figure 4: Instantaneous temperature isosurfaces of glycerol (Pr = 2547.9, Ra = $10^9$) under OB (left) and the NOB conditions (right).

because of the mesh size and the required time for sufficient statistics. Additionally, we plan to test existing and as appropriate develop subgrid scale models and perform large-eddy simulations (LES).

## Acknowledgments

## References

[1] Ahlers, G., Brown, E., Fontenele Araujo, F., Funfschilling, D., Grossmann, S., Lohse, D.
J. Fluid Mech., 2006

[2] Gray, D.D., Giorgini, A.
Int. J. Heat Mass Transfer 19, 1996

[3] Horn, S., Shishkina, O., Wagner, C.
J. Phys.: Conf. Ser. 318, 082005., 2011

[4] Shishkina, O., Wagner, C.
C. R. Mecanique 333, 2005

[5] Shishkina, O., Stevens, R.J.A.M., Grossmann, S., Lohse, D.
New J. Physics 12, 2010

## Links

http://scart.dlr.de/site/research-projects/non-oberbeck-boussinesq-effects-in-natural-convection/

http://scart.dlr.de/site/research-projects/highly-resolved-numerical-simulations-of-turbulent-rayleigh-benard-convection/

http://scart.dlr.de/site/research-projects/rotating-turbulent-rayleigh-benard-convection/

• Susanne Horn
• Matthias Kaczorowski
• Olga Shishkina

Institute of Aerodynamics and Flow Technology German Aerospace Center (DLR), Göttingen



Figure 3: Instantaneous temperature isosurfaces of water (Pr = 4.38, Ra = $10^8$, magenta indicates the warm, cyan the cold fluid), shown are the non-rotating OB (upper left) and NOB case (upper right), as well as a moderately rotating OB (lower left) and NOB case (lower right).

# Development and Validation of
# for Li-Ion Batteries in Hybrid

# Thermal Simulation Models
# and Pure Electric Vehicles

For the design and integration of alternative propulsion components such as Li-ion battery modules for electrified vehicles efficient and well suited simulation tools are needed [1]. The tools are needed in the scope of the CAE-determined virtual development process to assure performance, durability, reliability and cost-related requirements for the battery module on vehicle, module and cell level. The simulation level determines the degree of geometric complexity covered by the CAE model. As an example, a cell level CAE model of a Li-ion pouch cell covers the entire electrochemical active area. Thus, modeling strategies are needed that combine traditional, well established, three-dimensional multi-physics CAE tools with electrochemical models as demonstrated in [2]. These modeling strategies should allow the investigation of requirement-related questions with different physical depth aligned with the respective simulation levels of the virtual development process. Since these modeling strategies are considered as pre-competitive by the OEM

consortium of the asc(s [3] the necessity of strategies has led to several intensive discussions and development activities.

The prediction of the electro-thermal behavior of a single Li-ion battery cell under different electric loads can be covered by various modeling approaches described in literature. One group of models, called physics-based models [4], uses the porous-electrode and concentrated solution theory from Newman [5,6], Ohm's law and mass as well as energy balances. Thus, this group of models is computational intensive but covers the highest depth of physical understanding. Another group of models [7], called semi-empirical models, replace basically the porous-electrode and concentrated solution theory by the semi-empirical correlation from Newman and Tiedemann [8,9,10] to describe the current density as a function of the potential difference between the positive and negative electrodes and the current collectors of the battery cell as a func-
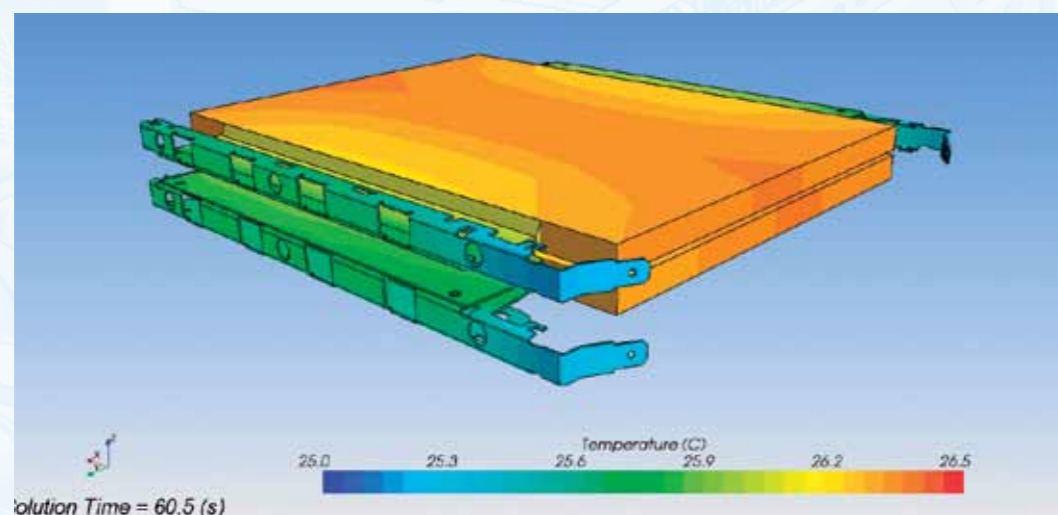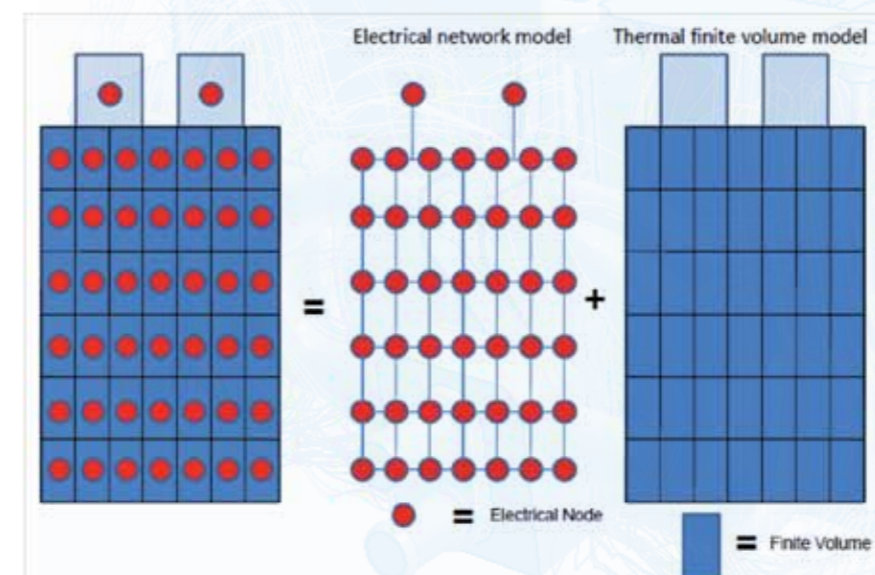


Figure 1: Coupling of electrical network model and thermal finite volume model.

tion of depth of discharge. Hereby, this modeling group becomes less computational intensive, is still on cell level but loses some degree of physical capability in the porous electrodes. The third group of models [11], called equivalent circuit models, describes the current-voltage behavior of a battery cell by an equivalent system, an electric circuit, which is the largest physical simplification in comparison to

both physics-based and semi-empirical models. Thereby, this group of models takes the least computational effort and can also be effectively coupled with multi-physical CFD tools.

The project "Development and Validation of Thermal Simulation Models for Li-Ion Batteries in Hybrid and Pure Electric Vehicles" is a cooperation of Adam Opel AG, asc(s, Battery Design LLC,
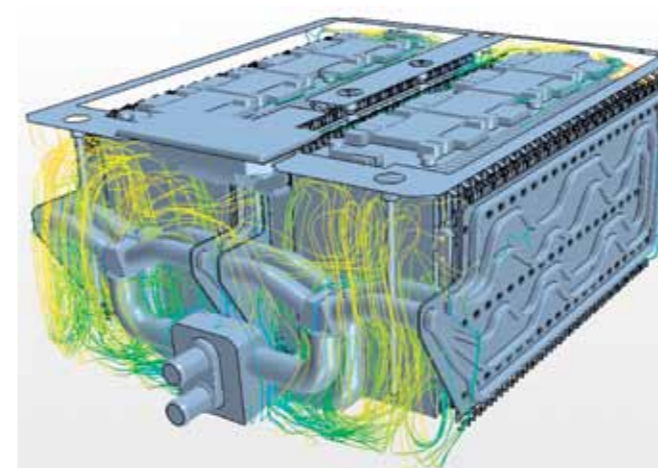




Figure 2: Investigations at different kinds of sub models: two cell model including conduction part between and inside air volume model (all pictures with courtesy of the Behr group).

CD-adapco, Daimler AG, Dr.-Ing. h.c. F. Porsche AG, as well as the Behr group as a subcontractor. It concentrates on the development of a simulation strategy on battery module level and introduces a coupling concept between an equivalent circuit cell model with a multi-physics CFD code (Fig. 1). The

The entire lithium ion battery module consists of approx. 16 million mesh cells and includes multiple high power pouch cells, a cooling circuit and the housing. In order to approve the developed concept for electrothermal battery simulation and evaluate the quality of the simulation model and

is needed to perform battery simulations. But those cores require 3-4 GB RAM each. A common battery simulation within the scope of the project allocates e.g. four 64 GB nodes at the Cray XE6 Hermit but uses only every second core.



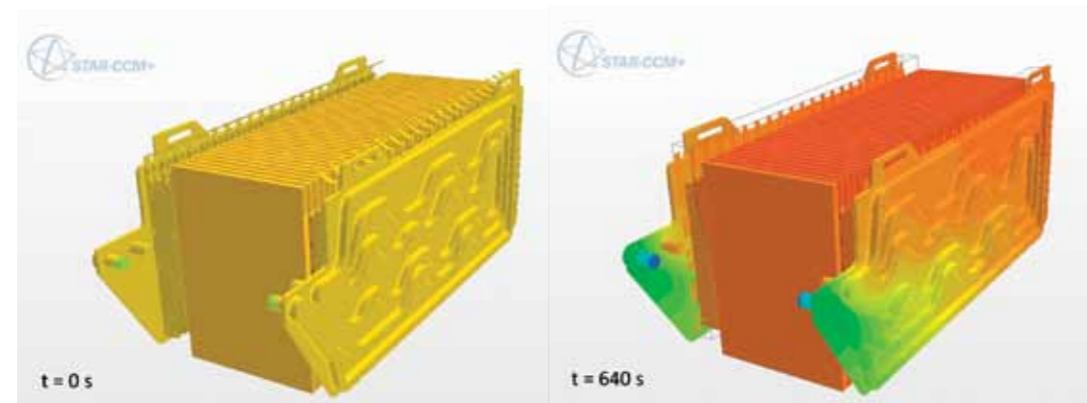Figure 3: Transient behavior of a simplified battery module. Temperature distribution in the battery cells and cooling plates at different times

overall model can be used to investigate heat and energy management and to optimize the layout of the battery cooling system. Knowledge about the temperature and charge distribution between cells within the battery module is crucial, since potential inhomogeneities can have a negative impact on battery life.

In order to develop a best practice simulation strategy theoretical thoughts as well as experimental results have been applied on the basis of several kinds of sub models. Module parts such as a two cell model, a single cooling plate or the inside air volume have been studied extensively, e.g. to find the optimal thermal and electrical mesh resolution of the stack in in- and through-plane direction, or to examine the impact of natural convection and radiation effects on the cell temperature.

numerical set up on cell and pack level, several verification and validation tests have been performed. Fig. 3 demonstrates the interaction between the battery cells and cooling plates for an electrothermal test case. The battery cells are electrically cycled while the coolant plates switch on and off according to defined temperature ranges for the battery cells.

The battery simulations were performed at the NEC Nehalem cluster as well as the CRAY XE6 Hermit. The simulation time quickly amounts to a week in order to run a 2-3 hours lasting electrothermal test case for the full battery pack. While the CFD part of the solution is fully parallelized the electrical calculations work in serial for each battery cell. Thus just a few couple of nodes resp. cores (related to the number of battery cells in the module)
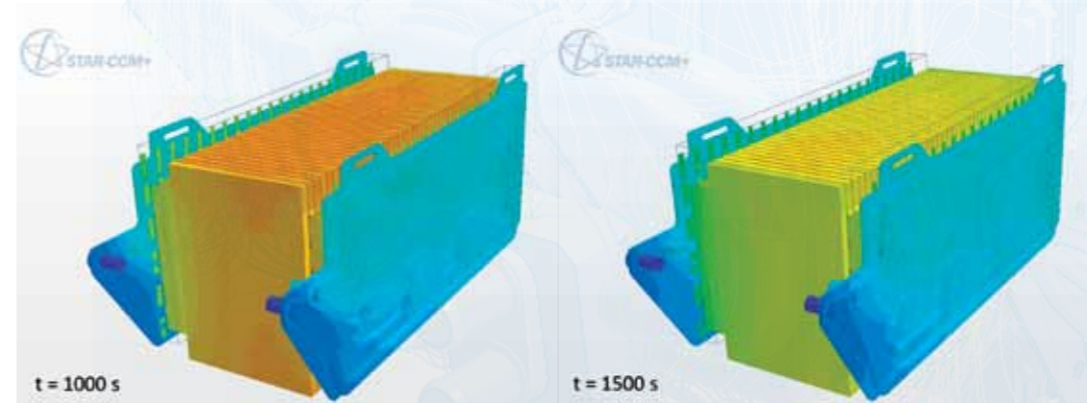
## References

[1] **Fell, S.**
CAE Methods and Test Requirements for Battery System Design, 3. VDI-Fach-'kongress Elektromobilität, 2011

[2] **Damblanc, G., Hartridge, S.,Spotnitz, R., Imachi, K.**
Validation of a new simulation tool for the analysis of electrochemical and thermal performance of lithium ion batteries, JSAE Annual Congress, 2011

[3] **Automotive Simulation Center Stuttgart e.V.**
http://www.asc-s.de

[4] **Doyle, M., Newman, J.**
The use of mathematical modeling in the design of lithium/polymer battery systems, Electrochimica Acta, Volume 40, Issues 13–14, October 1995, pp. 2191-2196

[5] **Newman, J.S. and Tobias, C.W.**
Theoretical Analysis of Current Distribution in Porous Electrodes, J. Electrochem. Soc., Volume 109, Issue 12, pp. 1183-1191, 1962

[6] **Newman, J.S.**
Electrochemical Systems, Prentice-Hall, Englewood Cliffs, New Jersey, 1991

[7] **Kim, U.S., Shin, C.B., Kim, C.S.**
Modeling for the scale-up of a lithium-ion polymer battery, Journal of Power Sources 189 (2009) 841–846

[8] **Tiedemann, W., Newman, J.**
in: Gross, S. (Ed.), The Electrochemical Society, Princeton, NJ, 1979, pp. 39–49

[9] **Newman, J., Tiedemann, W.**
J. Electrochem. Soc. 140 (1993) 1 961–1968

[10] **Gu, H.**
J. Electrochem. Soc. 130 (1983) 1459–1464

[11] USABC Linearized Battery Model, Freedom-CAR Battery Test Manual for Power-Assist Hybrid Electric Vehicles, October 2003

contact: Jenny Kremser,
jenny.kremser@asc-s.de

• Jenny Kremser et al.

Adam Opel AG
asc[s
Battery Design LLC
CD-adapco
Daimler AG
Porsche AG

# Matter-Antimatter Asymmetry and the Search for fundamental Laws of Physics with JUGENE

## The Matter-Antimatter Asymmetry

Our visible universe today consists of matter with only some tiny trace amounts of antimatter, e.g. in cosmic rays. Although this statement seems to be rather trivial given our everyday experience, it has some far-reaching consequences for the origin of our universe. Unless we are willing to accept that the universe started out in an asymmetric state, with more matter than antimatter, some mechanism has to exist that gives preference to the production of matter over antimatter in the course of evolution of our universe.

In 1967 Sakharov has listed the necessary ingredients for such a mechanism. Obviously one ingredient is a process that preferentially transforms matter into antimatter, and indeed such processes have been found in nature. The most prominent of these processes (see Fig. 1) is the so-called "kaon mixing", where a neutral particle (the kaon or $K^\circ$) is transformed into its antiparticle (the anti-kaon or $\bar{K}^\circ$)

The reverse process, the transformation of an anti-kaon $\bar{K}^\circ$ into a kaon $K^\circ$, does also exist, but it progresses at a slightly different rate, thus leaving a slight imbalance between matter and antimatter.

The obvious question now is whether this slight imbalance is able to explain the observed matter dominance of the universe. This question, however, is not so straightforward to answer because the particles involved, the kaons and anti-kaons, are not simple elementary particles. They are what we call mesons, bound states of one quark and one antiquark held together by the strong nuclear force. Because the effects of the strong nuclear force are very difficult to account for precisely, it is not straightforward to deduce from the properties of the (anti-)kaons the fundamental properties of their constituent (anti-)quarks that are ultimately relevant for explaining the observed mater-antimatter asymmetry.

Broadly speaking, there are two possible scenarios: either the experimentally observed neutral kaon mixing can be explained by the known properties of quarks within the current standard model of particle physics, the so-called CKM (Cabibbo-Kobayashi-Maskawa) theory, or not. In the first case one would get a nontrivial confirmation of the standard model and a hint that the matter-antimatter asymmetry has its origin at a much higher energy scale, because the standard model alone is known not to provide enough of it. In the second case one would obtain clear evidence for the existence of "new physics" - new particles, forces, or mechanisms (e.g. extra dimensions) that go beyond the established standard model. In this scenario the "new physics" would provide some part of the experimentally observed kaon mixing and a potential answer for the mater-antimatter asymmetry in the universe.

## QCD

In order to clearly distinguish between these two scenarios, the properties of kaons and their constituent quarks have to be related numerically in a precise way. This is not a straightforward

task because the quarks that form a kaon are bound by the strong nuclear force that we can describe with a theory called quantum chromodynamics (QCD). QCD, however, does not allow for the existence of single, free quarks; quarks have to occur in bound states only. This essential feature of QCD, called quark confinement, is caused by the specific "color" interaction between quarks. In addition to the electrical charge, quarks do carry a "color" charge. This "color" has nothing to do with real color – it was labeled "color" because of the peculiar mixing behavior of the fundamental charges that is reminiscent of the mixing between the primary colors red, green, and blue.

While an electrical charge can only be positive or negative and have a different magnitude, color charge has three components that we may label red, green, and blue. This has surprising and far reaching consequences: let us first imagine placing two electrical charges of the same magnitude but
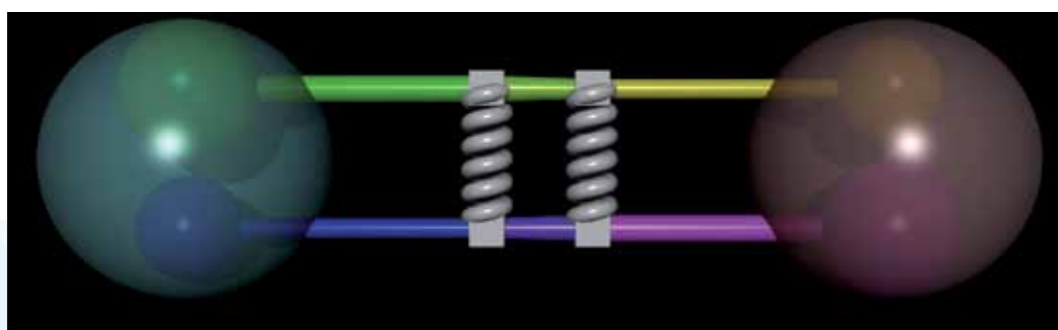


Figure 1: An illustration of neutral kaon mixing. Both the kaon and its antiparticle contain one quark and one antiquark bound together by the strong nuclear force and they can transform into one another (grey).
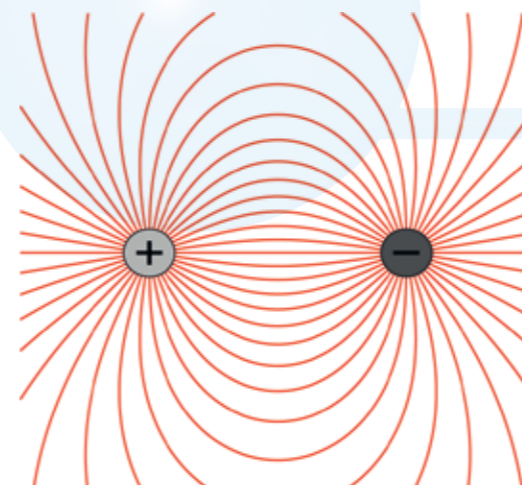


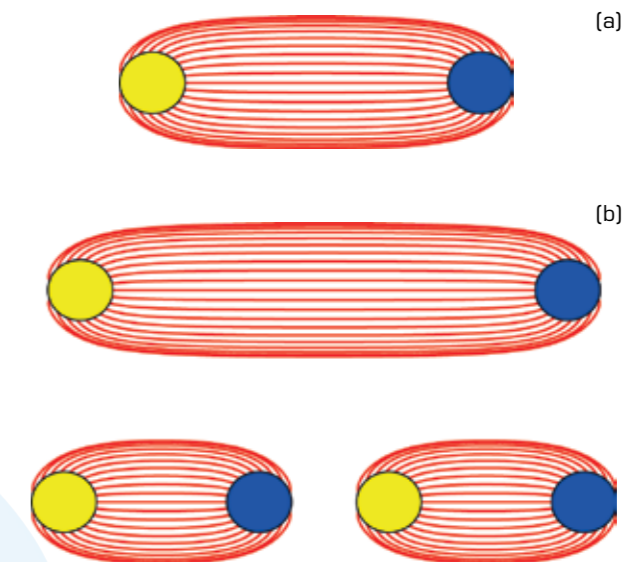Figure 2: Field lines of a pair of opposite electric charges.



(a)

(b)

Figure 3: (a) The color flux lines between a pair of opposite color charges. (b) When separating the two charges, the flux tube gets elongated and ultimately breaks up by producing a quark-antiquark pair out of the vacuum.

opposite sign into otherwise empty space (see Fig. 2). The electrical field, or its field quanta, the photons, can spread out freely into space. This is essentially due to the fact that the photons themselves are electrically neutral.

We now turn to the color field between two opposite color charges or quarks. The field quanta of the color interaction, the gluons, must relate the different color components and therefore cannot all be color neutral. As a consequence, the field lines themselves are color charged and do interact with each other. This results in a squeezing of the color flux tube between the two opposite charges (see Fig. 3a). Trying to pull the charges apart will result in the flux tube containing more and more energy. At a certain point there is enough energy in the flux tube to pair-produce a quark-antiquark pair out of the vacuum (see Fig. 3b). Instead of having split up the quark-antiquark bound state, or meson, into its constituents, we end up with two mesons eventually.

## Lattice QCD

It is now evident that QCD (in the low energy regime that we are interested in here) is very hard to treat numerically: the fundamental degrees of freedom of the theory (quarks, gluons) are very different from the states that occur in nature (bound states like mesons or protons and neutrons). Usual perturbative techniques that deal with small corrections to the free fundamental degrees of freedom are evidently insufficient.

A more direct way of approaching this problem is lattice QCD. One discretizes the QCD equations on a space-time lattice and then proceeds to perform

the quantum mechanical path integration stochastically. This approach is of course extremely demanding numerically. The integration has to be performed on a space with a dimension given by the number of lattice points times the internal degrees of freedom. Since the lattice has to be both fine enough to reproduce all essential details and big enough to fit the entire system being investigated, the typical dimension is of order $10^8$ -$10^{10}$.

In addition to the large dimensionality D of the stochastic integral that needs to be performed, it also needs to be performed over quarks, which are fermions. Fermionic degrees of freedom, however, are anti-commuting numbers and their inclusion ultimately leads to the need of computing determinants of square matrices of dimension $D^2$. Though these matrices are typically sparse, it is nonetheless a challenging task even for petaflop-class installations like JUGENE to perform this stochastic integration.

## Extracting Physics

In order to extract final physics results from a stochastic integration of the QCD path integral, a number of further steps are necessary. The first one of these is reaching the "physical point". This step is rather unique to lattice QCD, as in most other computer simulations, the fundamental parameters of the theory are known beforehand, whereas the interesting observables are connected to the nontrivial behavior of large systems.

In lattice QCD we do not know beforehand the fundamental parameters of the theory – the masses of the quarks and the coupling constant. Free quarks

do not occur in nature and therefore their masses and the strength of their interactions cannot be measured directly. We therefore have to first solve the inverse problem and to determine which parameters of the theory allow us to reproduce the experimentally observed phenomena of the strong nuclear force.

In practice, we looked at the masses of three different bound states of quarks (two mesons, the kaon and the pion, and one heavier particle or baryon, that is a bound state of three quarks) and defined the physical point as the point where the ratio between the three masses assume their experimentally observed values. In our lattice simulations we had three parameters to tune: quark masses and the coupling constant. From the three bound state masses we use to identify the physical point, we obtain two independent mass ratios that constrain the space of three parameters to a one–parameter line of "physical points".

We can further differentiate the physical points along the line by comparing any one mass, e.g. that of the pion, from experiment to the dimensionless mass of the same bound state as measured on the lattice. This procedure is called "scale setting" as it determines the fundamental length scale of the lattice, which is also not known beforehand. The second step which is then necessary to obtain physical results is to extrapolate along the line of physical points to the point where the lattice scale vanishes. This step is called "continuum extrapolation" as it removes the discretization mesh and therefore leads to continuum results.
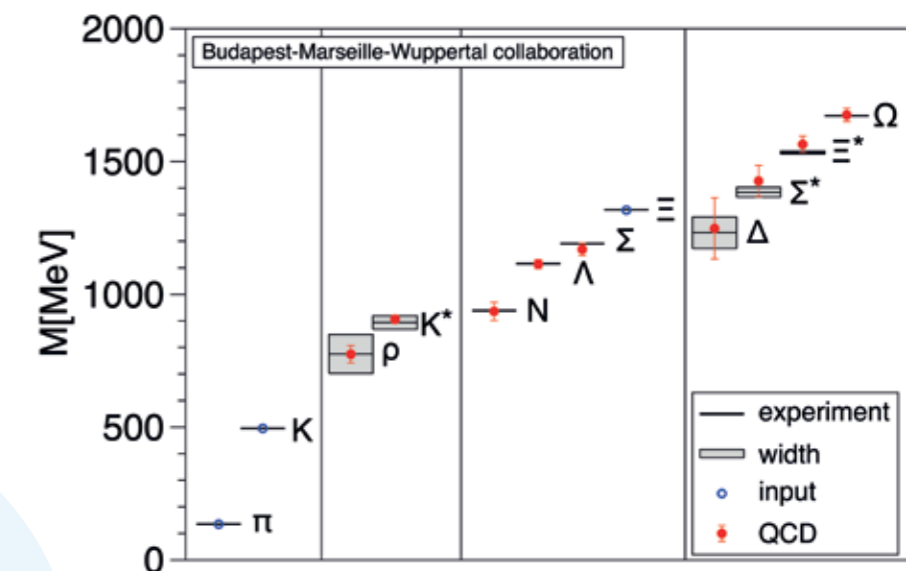


Figure 4: Comparison of some experimentally observed masses of quark bound states to our QCD prediction.

Knowing how to obtain continuum results, we still have to deal with the fact that our lattice represents only a finite volume in space and a finite extent in time. The third step towards obtaining physical predictions is therefore the infinite volume extrapolation. Fortunately most of the finite volume corrections are exponential in the lattice extent so that a large lattice extent renders them negligibly small.

## Controlled Errors

Once we have carried out this procedure of obtaining physical predictions, we might wonder what would happen if we used three other masses initially to construct our line of physical points. Would we have captured physics exactly with our discretized lattice, we would obtain the same line. In reality there are errors due to the truncation however, so one can trace out different lines of "physical points" depending on the input quantities used. If the theory and our methods of solving it are correct, however, all these lines must extrapolate to the same continuum value. We can
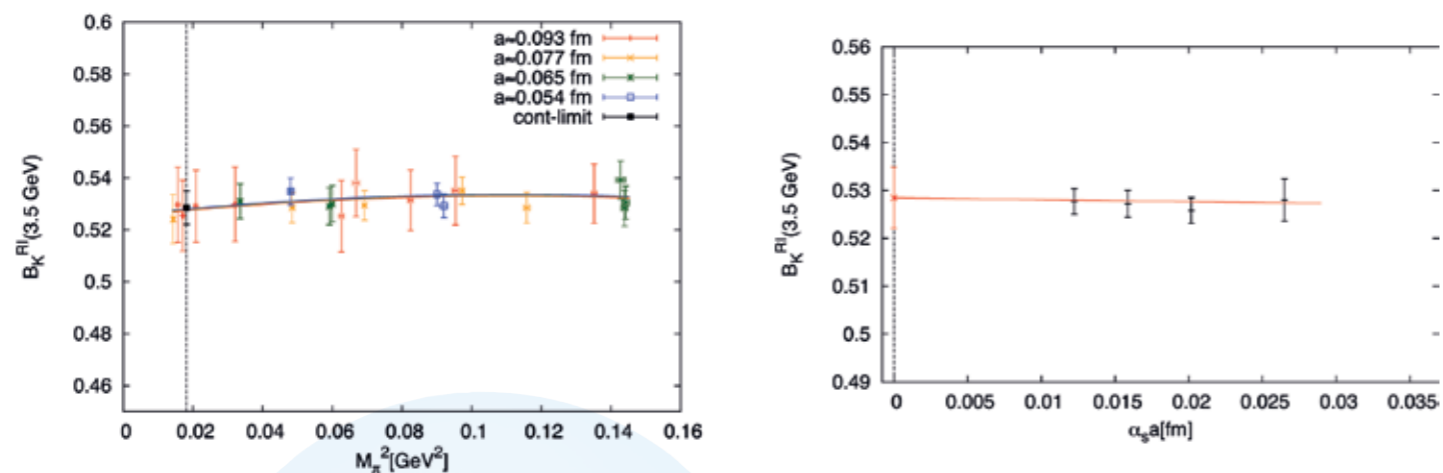
Figure 5: Value of the parameter $B_K$ as measured in our lattice calculations. (a) One interpolation of the lattice values to the physical point and (b) one extrapolation of the lattice values to the continuum.

therefore obtain a strong crosscheck of QCD itself and our ability to solve it by checking that not only our input masses but also the masses of other bound states are correctly reproduced in the continuum limit at the physical point. As shown in Fig. 4 we have successfully performed such a crosscheck in 2008 using mainly the computing resources at the Jülich Supercomputing Centre [1].

This result exemplifies what is the most important aspect of our calculations: a

full control over all sources of systematic error. When making a prediction that has the potential to disprove our standard physical theories about our world we need to make sure that we do not underestimate the inaccuracy of our prediction.

In order to achieve this goal, our first aim was trying to avoid any extrapolations. The supercomputing resources provided to us mainly by the Jülich Supercomputing Centre in combination with algorithmic and theoretical

advances [2,3,4] in fact allowed us for the first time in this kind of calculations to entirely eliminate the extrapolation to the physical point and replace it by an interpolation (see Fig. 5a). The unavoidable continuum extrapolation is very flat (see Fig. 5b) and was performed with different extrapolation functions. The spread between different results allowed us to reliably estimate the corresponding systematic uncertainty.

Following this general strategy, we in fact performed a total of 5,760 analyses resulting in a reliable estimate of the total uncertainty of our results.

## Consistency of the Standard Model

As a final result of our calculation, we obtain the so-called bag parameter of the kaon $B_K$. This number parameterizes the amount of matter-antimatter asymmetry in neutral kaon mixing and previous calculations were placing it in the region of 0.72–0.75 with a typical uncertainty of about 5%. The standard model is favoring a larger number, albeit with a relatively large error, and some authors were seeing this as a possible indication of new physics [5].

Our new result (see Fig. 6) is 0.773(12) and has a precision of 1.5%, which is entirely consistent with the standard model expectations. It is also far more precise than the standard model prediction, which implies that an improved standard model prediction for $B_K$ would increase the accuracy of this test to the percent level.

In conclusion, our calculation provided strong evidence that the CKM mechanism of matter-antimatter asymmetry in the standard model does correctly

explain the observed neutral kaon mixing. The origin of the matter-antimatter asymmetry in our universe is still a mystery of fundamental physics that needs to be explored in the future.

## Acknowledgment

## References

[1] **Durr, S. et al.**
(Budapest-Marseille-Wuppertal collaboration),
Ab-initio Determination of Light Hadron Masses, Science 322 (2008) 1224

[2] **Durr, S. et al.**
(Budapest-Marseille-Wuppertal collaboration),
Lattice QCD at the physical point: light quark masses, Phys. Lett. B701 (2011) 265

[3] **Durr, S. et al.**
(Budapest-Marseille-Wuppertal collaboration),
Lattice QCD at the physical point: Simulation and analysis details, JHEP 1108 (2011) 148

[4] **Durr, S. et al.**
(Budapest-Marseille-Wuppertal collaboration),
Precision computation of the kaon bag parameter, Phys. Lett. B705 (2011) 477

[5] **Brod, J. and Gorbahn, M.**
Next-to-Next-to-Leading-Order Charm-Quark Contribution to the CP Violation Parameter epsilon_K and Delta M_K, Phys. Rev. Lett. 108 (2012) 121801

• Christian Hoelbling
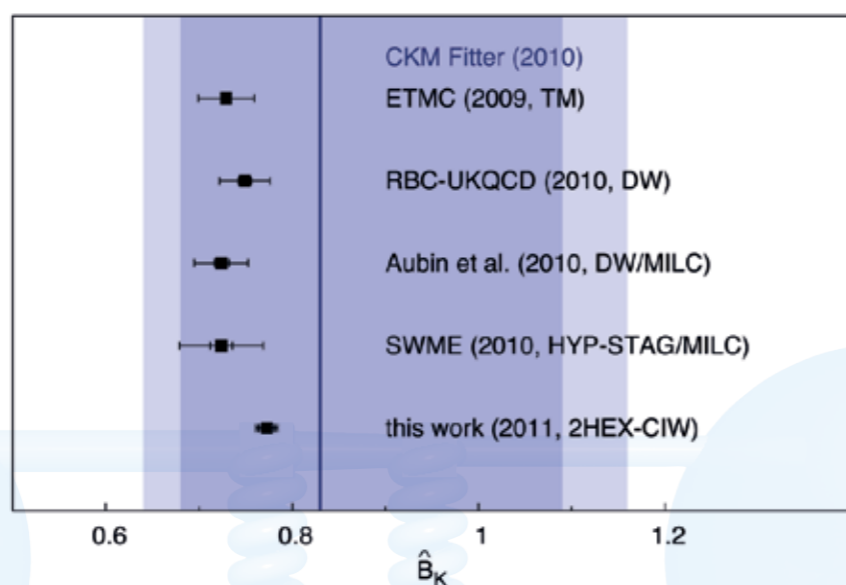
Bergische Universität Wuppertal



Figure 6: Our final value of the parameter $B_K$ compared to previous determinations (black) and a very conservative standard model prediction (blue). For more details see [4].

# Theoretical Mechano-chemistry: Stressing Molecules in the "Virtual Lab"

## What is "Mechanochemistry"?

Most chemical reactions must be activated by applying some form of excess energy. Certainly the oldest approach consists of using fire, i.e. thermal energy, leading to what is commonly known as thermochemistry. Other well–known "chemistries" are triggered by light or electricity, i.e. photochemistry or electrochemistry. A general observation is that different ways of activating chemical reactions might lead to different reaction path-ways – including different products in extreme cases. In contrast and despite pioneering macroscopic experiments by Matthew Carey Lea at the end of the 19th century, the field of mechani-cally induced covalent chemistry, which is dubbed here "mechanochemistry" in analogy to photochemistry, electro-chemistry etc., still remains largely un-explored when it comes **to a theoreti-cal understanding at the fundamental molecular level.**

Clearly this lack of understanding is at odds with the fantastic current **experimental possibilities** to use me-chanical nano–Newton forces as a **tool** to induce, alter, or control chemical reactions by manipulating atoms or olecules [1,2]. This is typically achieved within setups relying on atomic force microscopy (AFM), using individual mol-ecules that are covalently anchored to tips, or sonication experiments carried out in solution, using an ensemble of mechanophores with long polymer chains attached that act as force transducers.

On the other hand it is obvious that mechanical forces almost certainly will lead to different reaction pathways and thus different products much alike what is know from e.g. photochemistry or electrochemistry. This implies a growing gap between the ever increasing potential of experimental nanomanipu-lation techniques and a missing theo-retical framework to understand such experimental observations. Finally, it is noted that understanding mechano-chemistry will not only be important within the realm of chemistry as such, but it will impact as well on getting insights into mechanical aspects of such diverse issues as molecular friction, nanoscale machines, enzymatic reactions, or targeted delivery to name but a few.

The aim of the project "Mechano-chemistry of Covalent Bond Breaking from First Principles" is to complement experiments in the real laboratory by those carried out in the "Virtual Lab" [3], thus advancing the emerging field of theoretical mechanochemistry [4]. In particular, mechanically-induced ring-opening reactions have been in the focus of these investigations so far. Relying on ab initio simulations [5] at its heart, this study yielded unprecedented insights into the molecular details of the influence of constant external me-chanical forces on reaction mechanisms. Before presenting our representative results on ring-opening of cyclopropane derivatives a few words on our simulation approach are in order.

## Isotensional Stretching of Molecules in the "Virtual Lab"

Most quantum-chemical calculations or ab initio simulations of stretched molecules are carried out using so-called "isometric" conditions. Such setups are easily accessible in any conventional program package that has holonomic constraints available. This is either done by fixing an internal distance be-tween two atoms in a molecule, or by constraining the two atoms to specified positions along a space-fixed axis. The latter approach is mostly used in periodic codes, where the former in quantum chemistry codes that have been optimized to treat finite systems. A comprehensive review on theoretical mechanochemistry can be found in Ref. [4]. In contrast to this approach, the mechanochemical simulations presented in the following draw on our previously devised conceptual framework based on force-transformed potential energy surfaces (FT–PES) [6]. In this "isotensional" formalism, the FT–PES, given an ex-ternal constant force $\mathbf{F}_0$, is rigorously defined as

$$V_{EFEI}(\mathbf{x}, \mathbf{F}_0) = V_{BO}(\mathbf{x}) - \mathbf{F}_0 \, \mathbf{q}(\mathbf{x}),$$

where $V_{BO}(\mathbf{x})$ is the usual Born–Oppen-heimer (BO) PES as a function of all nuclear cartesian coordinates $\mathbf{x}$, and $\mathbf{q}$ is the mechanical coordinate, i.e. a structural parameter (a generalized coordinate in terms of $\mathbf{x}$) on which the force acts. By locating the stationary points of this function, in which the "External Force is Explicitly Included" (EFEI), one can evaluate, without invok-ing any approximation, properties such as reactant and transition state (TS) structures as functions of $\mathbf{F}_0$. It is certainly instructive to analyze the FT–PES, which provides a static zero–temperature perspective of mechano-

chemical reaction mechanisms [6,7,8,9], but this approximation clearly lacks the inclusion of thermal, entropic, dynamical and solvation effects except for crude modeling of these.

Using the power of capability platforms such as the Blue Gene/P machine JUGENE at Forschungszentrum Jülich allows us to go a major step beyond computing the FT–PES. The natural generalization is to carry out finite-tem-perature simulations using the FT–PES, i.e. $V_{EFEI}(\mathbf{x}, \mathbf{F}_0)$, as input instead of the usual BO–PES, $V_{BO}(\mathbf{x})$. Clearly, this calls for ab initio simulations [5] where the electronic structure and thus the forces acting on the nuclei is computed "on the fly" to move the nuclei under the influence of a constant external mechanical force, instead of pre-calculating and fitting $V_{BO}(\mathbf{x})$ and thus $V_{EFEI}(\mathbf{x}, \mathbf{F}_0)$. This combination enables us to map multi-dimensional free energy landscapes at constant force, which we call force-transformed free energy surface (FT–FES) [10].

However, following this avenue requires the computation of free energy surfaces for several constant forces, which is a daunting task. This is made possible using the powerful ab initio meta-dynamics technique developed by Laio and Parrinello as reviewed in Ref. [11]. At the technical level, using the so-called "multiple walker" sampling in conjunction with the CPMD software package [12] greatly improves the efficient use of the Blue Gene architecture, see our inSiDE contribution on prebiotic peptide synthesis [13] for technical background and scaling assessment. Together, this package of methods allows us to carry out efficiently ab initio simulations [5] of mechanochemical reactions on Blue Gene platforms.

# Enforced Ring-Opening of Cyclopropanes

Cyclobutene-based mechanophores have received a lot of attention recently, whereas cyclopropane systems, such as **gem**-dihalogencyclopropane derivatives, are under-researched in the realm of covalent mechanochemistry [2,4]. Yet, the force-induced electrocyclic ring-opening of these compound classes has furnished striking results [2,4]. Indeed, the mechanochemical activation of cis benzocyclobutenes has been shown, both experimentally and computationally, to promote a thermally forbidden disrotatory ring-opening process. The application of a transient tensile force on **gem**-difluorocyclopropanes, in its turn, has been shown experimentally to lead to an unexpected isomerization of trans species into their less stable cis isomers via a mechanochemical trapping of a diradical TS. Even more enigmatic are **gem**-dichlorocyclopropane (**gDCC**) systems, which feature a counterintuitive

lack of selectivity in the mechanically assisted ring-opening reactions of cis versus trans isomers: while one expects the external forces to promote the ring-opening of cis **gDCC** more efficiently (due to a better coupling between the mechanical coordinate and the reaction coordinate), the experimental observations indicate that both isomers undergo force-induced ring-opening processes with approximately the same probability.

In our project the mechanochemical reactivity and force-induced stereochemistry of **gDCCs** has been scrutinized based on isotensional ab initio metadynamics simulations [10]. Our thermodynamic approach is supplemented by ab initio trajectory shooting simulations operating on the FT-PESs in order to dissect **genuinely dynamical effects** on branching ratios as a function of force (the details of this technique are provided in the Supporting Information of Ref. [10]). Based on these methods

we have unveiled the mechanisms of force-induced ring-openings of cis versus trans **gDCCs**, which rationalizes puzzling experimental findings. Even more importantly, we have discovered an unprecedented complex mechano-**stereo**-chemical behavior, whereby the ring-opening of trans isomers of 1,3-disubstituted **gDCCs** can lead to two different diastereomers, with the probabilities of obtaining them featuring an intricate dependence on the force exerted onto the system.

The model system chosen to explore the mechanochemistry of **gDCCs** is 1,1-dichloro-2,3-dimethylcyclopropane: its cis and trans isomers and the four possible distinct reaction products of the corresponding ring-opening processes are depicted in Fig. 1. In the first stage of our study, the thermal chemistry (i.e. the chemistry at zero force) of **gDCCs** was dissected. Our simulations (both at 0 K and at 300 K) have revealed that the ring-opening of these molecules, which yield the corresponding 2,3-dichloroalkenes, proceeds via a concerted disrotatory mechanism, whereby the breaking of the C-C bond takes place in concert with the C-Cl bond cleavage and the subsequent Cl migration. For cis **gDCC** there are two possible pathways as indicated in Fig. 1: the "outward pathway" which passes through TS-I and the "inward pathway" via TS-III. On the basis of the difference between the activation energies of both pathways (about 5 kcal/mol, see Fig. 1), it has been concluded that the thermal ring-opening of cis **gDCC** occurs via a "disrotatory outward mechanism", whose BO-PES features a TS of $C_s$ symmetry (TS-I) and a bifurcation point along the intrinsic reaction coordinate (IRC) after the TS is left behind. By virtue of

this topological feature, the migrating Cl atom can move either to the C atom on the right side (thus yielding the Z,R-alkene) or to the left (leading to Z,S-alkene), see Fig. 1. Given the topology of the underlying PES, the ring-opening of cis **gDCC** is expected to yield the two enantiomeric alkenes with equal probability. The disrotatory ring-opening of trans **gDCC** at zero force, in its turn, implies either the TS-II or the TS-IV according to Fig. 1, with neither of these two TSs featuring symmetry. Trajectory shooting simulations initiated from TS-II have brought to light the fact that the ring-opening of trans **gDCC** can yield either the E,S-alkene or the Z,S-alkene in the absence of external forces. The probabilities of obtaining these two diastereomers were calculated to be 0.76 and 0.24, respectively. The branching ratio of obtaining E,R-versus Z,R-alkenes via TS-IV are identical because of symmetry. The computed activation free energies were found to be lower (by about 4 kcal/mol) for the ring-opening of cis **gDCC** than trans **gDCC**, as shown in Fig. 1.

After having set the stage by describing the thermal chemistry of **gDCCs**, we will now focus on the mechanochemical behavior of these molecules. As explained in Ref. [10], of all the reaction pathways depicted in Fig. 1 only two are relevant for fully describing the mechanochemistry of **gDCCs**: the disrotatory outward mechanism of cis **gDCC** and the disrotatory ring-opening of trans **gDCC** passing through TS-II. As shown in Fig. 1, the force-dependence featured by the activation free energy, $\Delta A^{\ddagger}$, and the activation energy, $\Delta E^{\ddagger}$, for the two considered pathways follows a similar general trend. The values of $\Delta A^{\ddagger}$ ($F_0$) have been obtained from the FT-FESs displayed in Fig. 2. Despite the
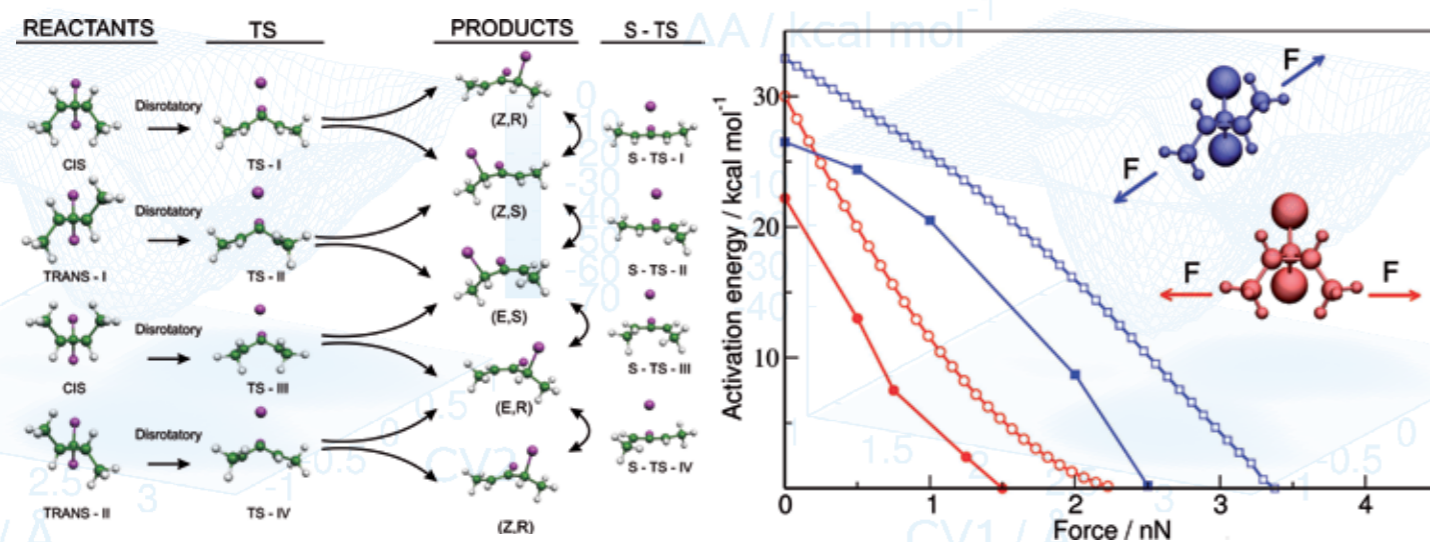


Figure 1: Left: Scheme showing the involved chemical species, i.e. all reactants (cis; trans–I and trans–II being enantiomers), transition states (TS–I to TS–IV; S–TS–I to S–TS–IV), and products (Z,R; Z,S; E,S; E,R). The arrows connecting the reactants with the distinct products via the corresponding TSs represent the reaction paths obtained from IRC mapping and ab initio trajectory shooting starting from the TSs (see text). The second set of TSs (S–TS) belongs to interconversion reactions between selected products as indicated. For simplicity all structures correspond to the stationary points at zero force, the Z,R product is reproduced twice for clarity, and the Cl atoms are colored violet. Right: Force–dependence of activation energies ΔE‡ (F_O) (open symbols) and free energies ΔA‡ (F_O) at 300 K (filled symbols) of the disrotatory ring–opening of cis (red circles for the "outward" pathway) and trans (blue squares) 1,1–dichloro–2,3–dimethylocyclopropanes. The stretching force is applied to the C atoms of the two terminal methyl groups as indicated in the inset.
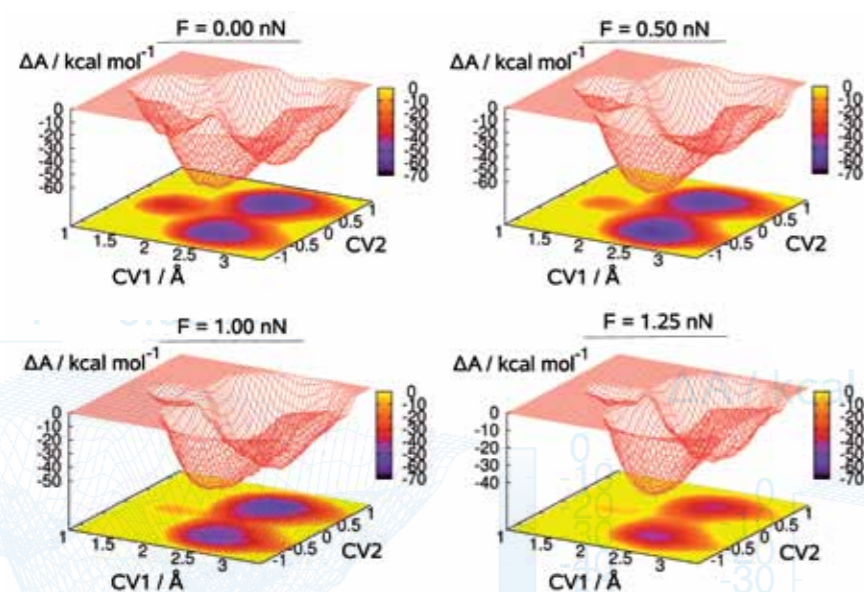
Figure 2: Force-transformed free energy surfaces (FT-FESs) for ring-opening for the cis reactant of 1,1-dichloro-2,3-dimethylcyclopropane in the absence of any external force and with three different external forces ($F_0$ = 0.5, 1.0, and 1.25 nN as indicated). These FT-FESs have been obtained by means of isotensional ab initio metadynamics simulations performed in a reaction subspace spanned by two collective variables: CV1 is the C C distance associated with the bond that yields upon the ring-opening process. CV2 is the difference CN1 - CN2 of the coordination numbers of both chlorine atoms with respect to the left (CN1) or right (CN2) carbon atom in the cyclopropane ring.

similar trends, the values of $\Delta A^{\ddagger}$ ($F_0$) are much lower by about 6-7 kcal/mol over the whole range of forces, which results in a dramatic rate acceleration due to finite-temperature and entropy effects not accounted for when computing only $\Delta E^{\ddagger}$ ($F_0$). With reference to the selectivity with which the tensile load might promote preferentially the ring-opening of one isomer, both the qualitatively different shapes of the energy curves associated with the cis and trans reactants and the corresponding rupture forces (1.5 nN and 2.5 nN, respectively, at 300 K) reflect the fact that the external forces enhance the ring-opening of cis **g**DCCs more efficiently. This is somehow contradictory to the experimentally found lack of selectivity. As argued in Ref. [10], the most plausible suggestion for resolving such apparent contradiction would

be to speculate that sufficiently large forces (forces on the order of 2 nN would suffice, according to the data in Fig. 1) were generated in the sono-chemical experiments as to reach the barrierless regime for the ring-opening reactions of both isomers.

Last but not least, our exploration of the mechanochemistry of **g**DCCs has unveiled a most surprising feature: the force-dependent selectivity of the ring-opening of **trans** reactant to yield either the Z- or the E-diastereomer of the corresponding dichloroalkene. As clearly displayed in Fig. 3, the corresponding branching ratio changes dramatically and **non-monotonically** as a function of $F_0$. This intricate switching behavior obviously implies that a ratio of about 50:50 is expected to be observed only close to some specific critical forces (0.7 nN, 1.9 nN and 2.2 nN). Importantly, in the limit of large forces the majority product is the Z,S-product, which is just opposite to the situation encountered at zero force. To the best of our knowledge, this is the first time that such an intricate mechanochemical behavior of a reaction resulting from intrinsically dynamical effects but determining its stereochemistry is reported.

## Outlook

The theoretical mechanochemistry presented so far has been carried out in the gas phase. Clearly, this is only a first step as each and every experiment is carried out in the condensed phase - be it a liquid phase in sonication or a protein matrix in AFM experiments. Although initial work along these lines has been carried out on JUGENE, we are looking very much forward to using the Blue Gene/Q platform JUQUEEN at Forschungszentrum Jülich to push theoretical mechanochemistry

in much more realistic environments such as solutions and biomolecules.

## Thanks

Figure 3: Force-dependence of the probability of obtaining the E,S-products (red) versus Z,S-products (black) upon ring-opening of trans reactant as computed from dynamical trajectory shooting simulations. The range of forces with yellow and white backgrounds correspond to the forces at which the majority product of the ring-opening process is the E,S-alkene or the Z,S-alkene, respectively.

## References

[1] **Beyer, M.K., Clausen–Schaumann, H.**
Chem. Rev. 105:2921-2948, 2005

[2] **Caruso, M.M., Davis, D.A., Shen, Q., Odom, S.A., Sottos, N.R., White, S.R., Moore, J.S.**
Chem. Rev. 109:5755-5798, 2009

[3] **Marx, D.**
Theoretical Chemistry in the 21st Century: The Virtual Lab", in: Proceedings of the "Idea–Finding Symposium: Frankfurt Institute for Advanced Studies, (Greiner, W.; Reinhardt, J., Eds.), EP Systema, Debrecen, pp 139–153, 2004. http://www.theochem.rub.de/go/cprev.html

[4] **Ribas-Arino, J., Marx, D.**
Chem. Rev. (online at http://pubs.acs.org/doi/abs/10.1021/cr200399q)

[5] **Marx, D., Hutter, J.**
Ab-initio Molecular Dynamics: Basic Theory and Advanced Methods, Cambridge University Press, 2009

[6] **Ribas-Arino, J., Shiga, M., Marx, D.**
Angew. Chem. Int. Ed. 48:4190–4193, 2009

[7] **Ribas-Arino, J., Shiga, M., Marx, D.**
Chem. Eur. J. (Communication) 15:13331–13335, 2009

[8] **Ribas-Arino, J., Shiga, M., Marx, D.**
J. Am. Chem. Soc. 132:10609–10614, 2010

[9] **Dopieralski, P., Anjukandi, P., Rückert, M., Shiga, M., Ribas-Arino, J., Marx, D.**
J. Mater. Chem. 21:8309–8316, 2011

[10] **Dopieralski, P., Ribas-Arino, J., Marx, D.**
Angew. Chem. Int. Ed. 50:7105–7108, 2011

[11] **Laio, A., Parrinello, M.**
Lect. Notes Phys. 703:315–347, 2006

[12] **CPMD**
J. Hutter et al., http://www.cpmd.org

[13] **Nair, N.N., Schreiner, E., Marx, D.**
inSiDE 6:30–35, 2008

• Przemyslaw Dopieralski
• Padmesh Anjukandi
• Jordi Ribas-Arino
• Dominik Marx

Lehrstuhl für Theoretische Chemie Ruhr-Universität Bochum

# Global MHD Simulations of protostellar Jets with radiative Effects

The subject of the work I carried out on the HLRB 2 was the investigation of the effect of radiative cooling on the propagation of protostellar jets and the consequent effect on the observable emissions from these objects.

These highly supersonic jets of gas are ejected from accreting protostars in the early stages of their development, the first 30,000 years or so, labelled the Class 0 phase. The jets at this stage typically extend to some tens of thousands of AU (astronomical unit, 1AU is the distance from the earth to the sun).

Such an outflow interacts with the molecular cloud core in which the protostar is forming, and the resulting super-heated gas emits light at particular wavelengths according to the chemical composition, allowing astronomers to observe the traces of the outflow with telescopes, see Fig. 1.

The protostar itself is hidden within the dark centre of the molecular cloud, making theoretical investigation of the jet ejection process difficult. However, the appearance of the jet on such large scales gives observational clues which, in conjunction with simulations, can provide constraints for theoretical models.

## Modelling Jets with Simulations

Our approach is to model the observed protostellar outflow system as a hydro-
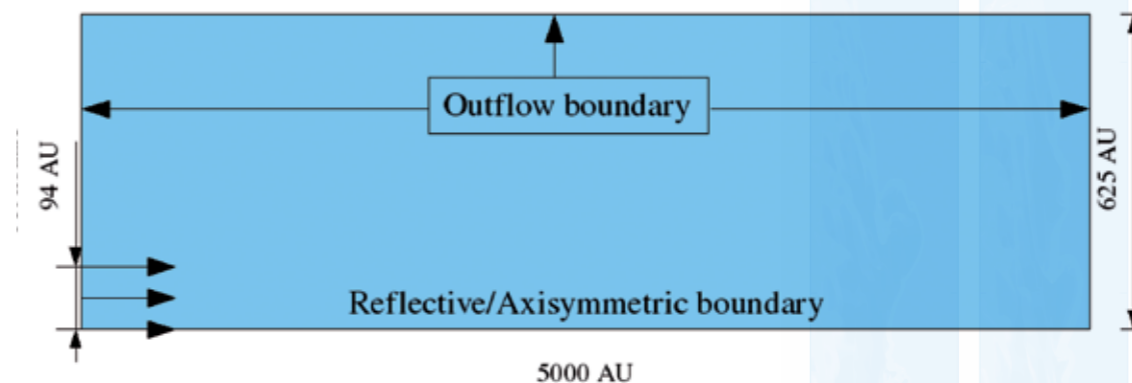


Figure 2: Schematic of the simulation domain setup [2].

dynamic or magnetohydrodynamic domain, representing the ambient medium of the molecular cloud, into which a collimated pulsating flow of gas is injected, representing the beam of the jet.

The geometry of jets lends itself quite well to simulation in a cylindrically axisymmetric 2-dimensional domain, which enabled us to carry out our simulations without the performance cost of full 3-dimensional computation. The setup is as shown in Fig. 2, where initial conditions are set for the gas of both the ambient medium and the jet.

Our simulations examined the effect of key parameters on the jet morphology and emission characteristics. Such parameters include the relative densities of the jet and its surrounding medium, the presence/strength of a magnetic field in the surrounding medium, the chemical composition of the molecular cloud and the level of ionisation of the gas in the jet beam.

We used the PLUTO astrophysical simulation code [1], with an additional module we have written to compute the chemistry and corresponding cooling losses at each grid point for every time step. PLUTO, a grid based finite-volume code, solves the HD/MHD equations in

parallel via several possible integration methods, while the integration method used for the chemistry module was a simple but effective semi-implicit backwards difference method.

The chemical network consisted of the main chemical species and the dominant set of reactions involved in the production of molecular $H_2$, a key coolant and emission source in the gas.

In order to be able to carry out simulations of such large scale objects, while still managing to resolve small-scale shock regions, we used the adaptive mesh refinement capability provided in PLUTO via the Chombo AMR libraries.



Figure 3: Increasing load corresponding to increasing wall-time as the jet propagates into the domain [2].

## Performance and Scalability

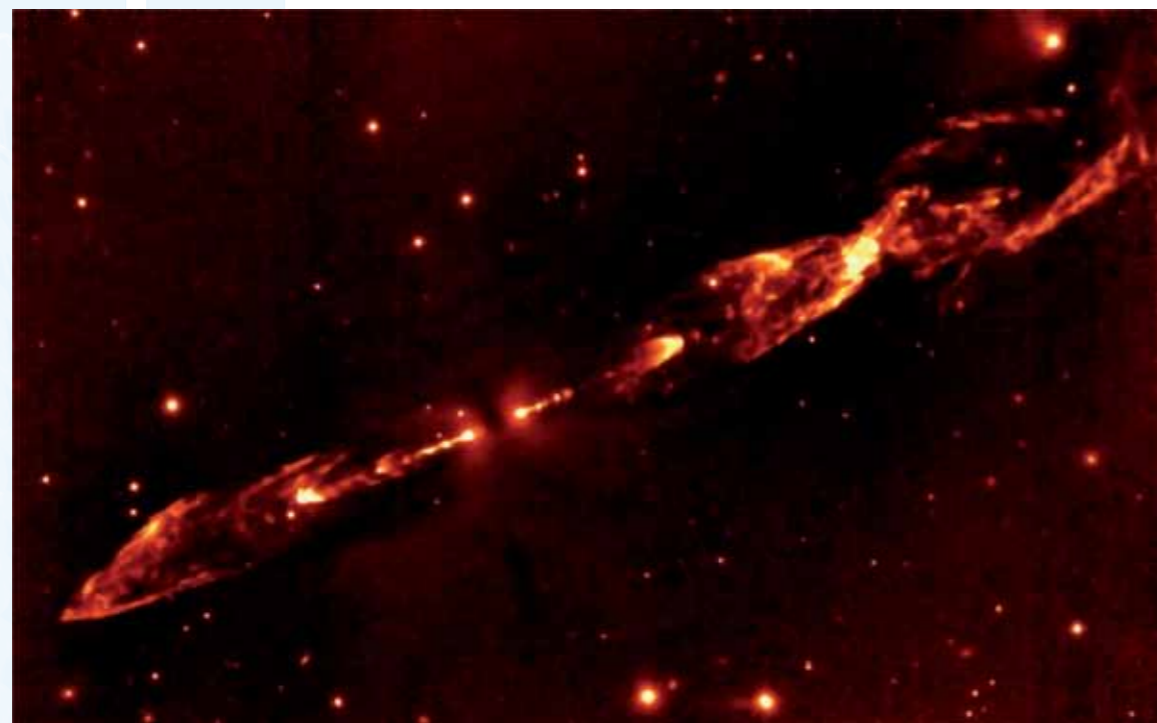The resolution chosen for the production simulations was a base grid of



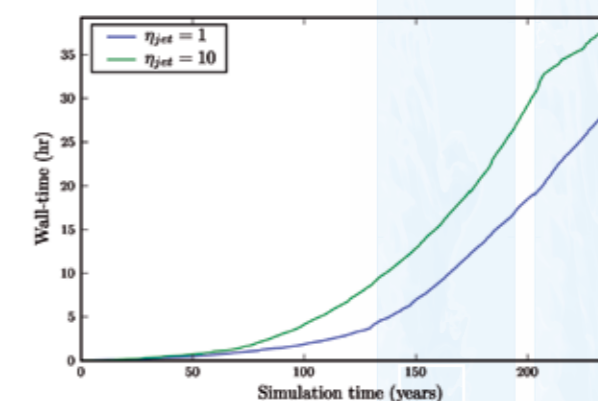Figure 1: HH212, a stunning example of a bipolar prostellar jet emitting in molecular hydrogen lines [3].
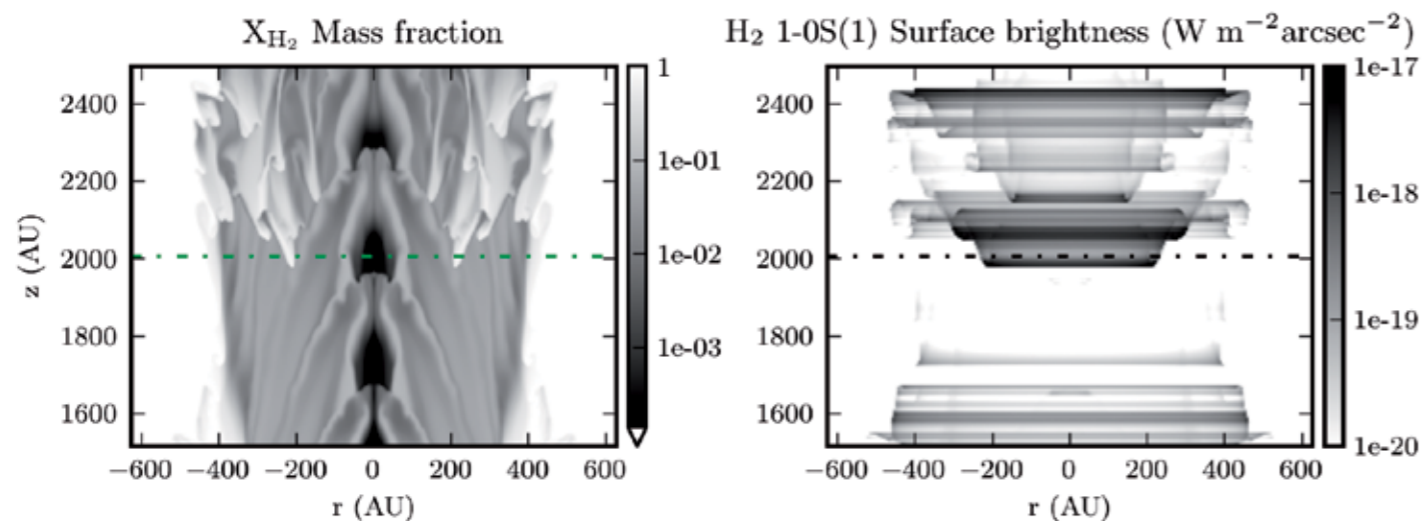
Figure 4: Section of the jet showing the entrainment of molecular ambient material and subsequent emission [2].
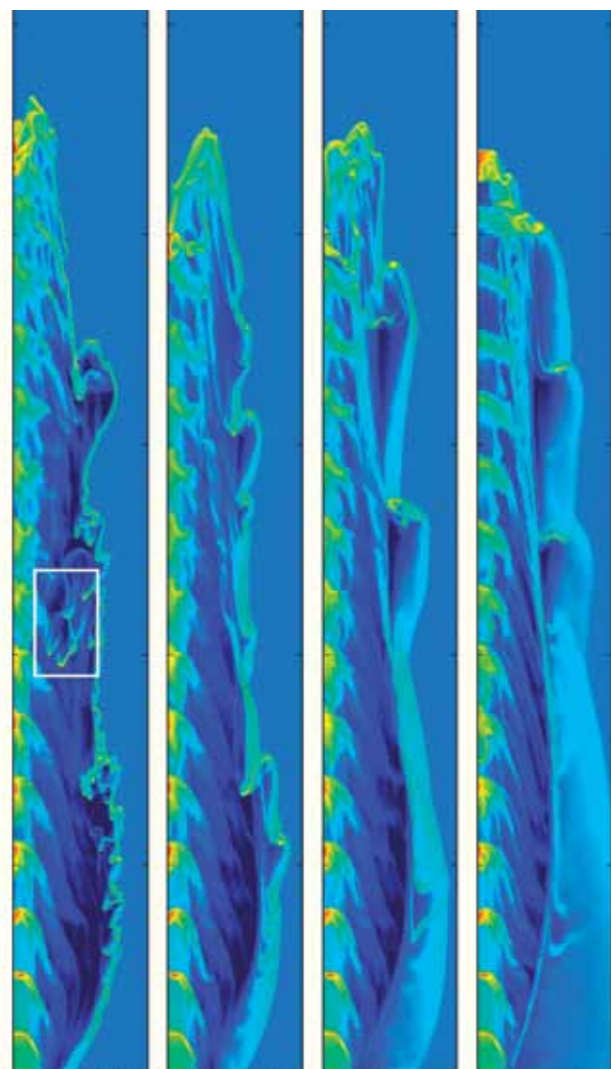


Figure 5: Showing the effect of the magnetic field, seen here increasing from 0µG on the left, through 30, 60 and 118 µG on the right. The white box shows matter from the unstable bowshock breaking up and entering the inner part of the jet in the absence of magnetic fields.

512x64 grid points, with 4 levels of binary mesh refinement, corresponding to an equivalent resolution of 8192x1024 grid points. For the domain of 5000AU this corresponds to a resolution of 0.61AU per grid zone, which is sufficient to resolve the chemical dissociation and cooling zones behind the shock front.

As the jet propagates through the domain and the area of interaction between the jet and the ambient medium increases, the computational load of the simulation also increases, as shown in Fig. 3.

For the production runs we used the HLRB II's SGI Altix 4700 supercomputer. Some thirty runs were carried out in order to probe the effect of different values for the physical and chemical parameters on the propagation and observable emissions.

For early parts of the simulation, where the computational load is small, the performance benefit of having a large number of processors is offset by the communication cost between the processors, so we used a varying number of processors at different stages of the runs, typically starting them with 32 or 64 processors and continuing at 128 or 256 processors. The typical CPU consumption per run was between 5 and 10 thousand CPU hours in total.

## Results

The data from the simulations provided a wealth of information on some of the factors affecting the jet propagation, giving insight into how the observed jets work.One interesting result was in illustrating the role of the magnetic field in maintaining the stability of the bow-shock. In the simulation run with no magnetic field, the cooling instabilities in the bowshock caused it to break up, allowing ambient molecular material to enter the jet "cocoon". The molecular material was then heated by internal shocks from the jet and caused to emit in the ro-vibrational lines of $H_2$, as shown in Fig. 4. This supports the "entrainment hypothesis" where the emissions are thought to arise from ambient matter swept up by the jet, as opposed to (or in addition to) matter in the jet beam itself.

In the case where magnetic fields are present, this is not seen to occur (see Fig. 5), and the emissions in these cases are seen to occur directly at the shock, a situation known as "prompt entrainment".

## Conclusion

The work described here formed part of my PhD work, which I concluded in 2009. Further details are given in my thesis [2].

## References

[1]   **Mignone, A. et al.**
      PLUTO: a Numerical Code for
      Computational Astrophysics, ApJS (2007),
      170, 228, 2007

      http://adsabs.harvard.edu/
      abs/2007ApJS..170..228M

[2]   **O'Sullivan, J.**
      Molecular Cooling and Emissions in Large
      Scale Simulations of Protostellar
      Jets http://www.ub.uni-heidelberg.de/
      archiv/10164, 2009

[3]   **Zinnecker, A. et al.**
      HH212: a prototype molecular hydrogen
      jet from a deeply embedded protostar, in:
      B. Reipurth & C. Bertout, eds., Herbig-Haro
      Flows and the Birth of Stars, Vol. 182 of
      IAU Symposium, 198, 1997

• Jamie O'Sullivan
• Max Camenzind

  Landessternwarte,
  University of
  Heidelberg

# LRZ Validation-Suite for Life Science

The potential of High Performance Computing (HPC) is increasingly being recognized in the fields of life science with particular focus on modelling and simulations. Life science organizations are developing into one of the largest e-Infrastructure users in Europe, and face many technical challenges in their quest to increase collaboration among researchers, streamline and replicate processes, enhance service levels and accelerate processing for faster time-to-market. This development will accelerate tremendously, and puts high demands on simulation software and support services.

This article will review and present a new Validation-Suite framework developed at the Leibniz Supercomputing Centre (LRZ-VS), which is an extended version of the DEISA benchmark-suite, developed by the Jülich Super-computing Centre [1,2]. The suite has been developed and used in the ScalaLife project [3], in order to effec-

tively validate the quality and scalability of several life science related software packages (GROMACS, DALTON and DISCRETE). To achieve this, realistic test cases were provided by alpha users of the Scalalife project, which were then run and analysed on a wide variety of target HPC platforms. This is important to understand and quantify the systems' theoretical and actual performance for the constantly updated releases of the software packages included in the project.

The test cases of the ScalaLife alpha users as well as the software packages have been integrated into the LRZ-VS. Utilizing the LRZ-VS, the most recent releases of the respective software packages are automatically compiled across the various HPC platforms. Furthermore, the application of the LRZ-VS allows automated execution, running and verification of the input datasets provided in the ScalaLife project.
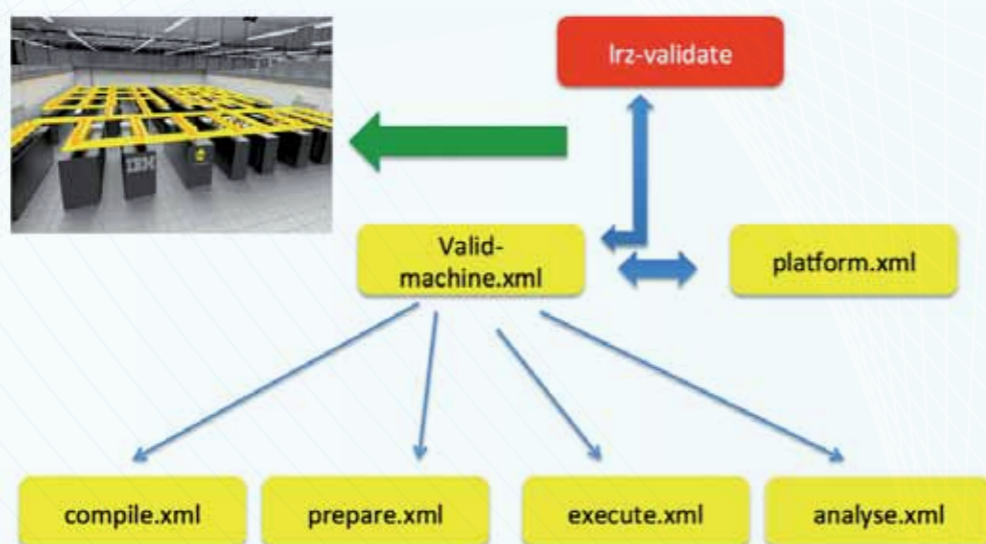


Figure 1: Automated compilation, execution and verification of validation runs - example for SuperMUC.

## Setup of the Validation Suite

For each platform a set of variables can be defined in the file "platform.xml" (see Fig. 1); this file can be found in the directory "platform/".

Once a platform is defined, each individual application that should be used in the LRZ-VS needs to be configured for this platform. At the LRZ, for example, the molecular dynamics application GROMACS can be compiled and submitted on the SuperMUC system with the following commands (see Fig. 1): $ cd ScalaLife/VS/applications/GROMACS $ ../../validation/lrz-validate Valid-SuperMUC.xmlThe LRZ-VS will compile a general binary for all node/task/thread combinations defined. If necessary, these parameters can be statically compiled into the binary. In the next step a so-called sandbox subdirectory is created for each job, ensuring conflict free operation of the individual applications at run time; the input files are prepared automatically.

## The Test Cases

In the following table a brief overview of all the parameters for the test cases of the individual applications is given:

| Software package | Name of test case | Number of atoms | Number of bonds |
|---|---|---|---|
| DALTON | Cytiac | 35 | 36 |
| | Spin-Label | 71 | 75 |
| | Porp | 72 | 79 |
| | | | |
| DISCRETE | 1 KIM | 4,681 | 4,735 |
| | | | |
| GROMACS | Lipid bilayer | 40,380 | 35,020 |
| | Gly water | 60,000 | 45,000 |
| | Snap | 66,463 | 49,597 |
| | Interface (Growth) | 114,000 | 111,989 |
| | Membrane | 261,412 | 196,630 |
| | Virion | 5,097,548 | 50,197 |

## Overview of the different HPC Systems at LRZ

The current LRZ infrastructure consists of a broad range of HPC systems equipped with nodes based on multi-core processors. Programs designed for parallel execution on a node can take advantage of its shared memory architecture. The systems are summarized in the following table:

| System | Processor | Total number of cores | Memory per node |
|---|---|---|---|
| SuperMUC 3 PetaFlops | Intel Sandy Bridge EP 16 cores per node | 147,456 | 32 GB |
| SuperMIG | Intel Westmere-EX (2.4 Ghz) 40 cores per node | 8,200 | 256 GB |
| SGI Altix UV | Intel Westmere-EX (2.4 Ghz) 20 cores per node (2.4 Ghz) | 2,080 | 3 TB |
| CoolMUC | AMD Opteron 6128HE (2 Ghz) 16 cores per node | 2,800 | 16 GB |
| GPU Cluster | AMD Opteron + 8 NVIDIA Quadro FX 5800 GPGPUs | 64 | 128 GB |

## Results and Discussion

At the LRZ, we have measured and validated the application examples on the systems SuperMIG and CoolMUC for Dalton (Fig. 3) and on the SuperMIG, CoolMUC and SGI Altix UV machines for GROMACS (Fig. 4).

All Dalton results scale well up to approximately 100 cores; currently using more cores does not produce results at reasonable scalability. This trend seems to be independent of the system size. Cytiac is the smallest test system but scales better than Spin-Label, which includes two-fold atoms and bonds. In contrast, the Porp system of similar size seems to scale well. Nevertheless, these trends can partly
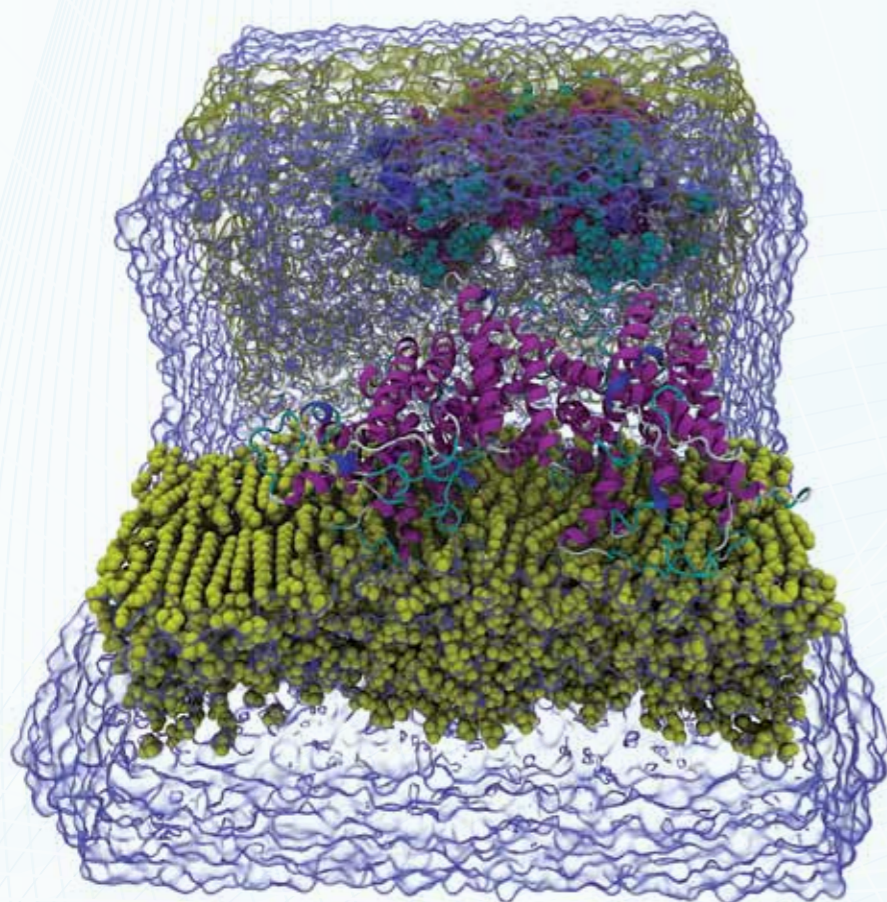
Figure 2: Membrane test case for GROMACS. Two layers of lipid in a water box can be seen. Embedded in the lipid layer is a protein.

be explained by the different operations executed in each test set, where Spin-Label is a large-scale test case for the QM/MM module of Dalton. The Cytiac test case applies advanced quantum mechanics methods, which are already quite expensive for small molecules. In the Porp case the density functional theory from quantum mechanics is applied.

For the GROMACS evaluation, we have compared the scalability of six different test cases on three different architectures. All test cases seem to scale well until approximately 400 cores. This observation is confirmed on the SuperMIG architecture: beyond 400 cores, the first test cases start to produce results with a sub-optimal scalability. The runs

on the CoolMUC and SGI Altix UV machines scale well until the maximum number of cores tested. The best performance is found for the medium-size test case Snap (~66 k atoms), where a special tweak (vsites) is used in order to increase the step size during the simulation by two-fold. However, this feature also seems to make this system less stable and prone to run-time errors; the medium size of the system further limits the maximum number of cores that can be utilised. The test cases with the larger system sizes Membrane (~260 k atoms) and Growth (~110 K atoms) scale up to 800 and 400 cores, respectively. Only the largest test set called Virion, amounting for more than 5 million particles, shows good scalability up to 2,000 cores and possibly more.

In summary, the LRZ Validation Suite provides a valuable means for automated testing, validation and benchmarking a large range of supercomputing codes. Once the system specific adaption has been generated, several program packages can be validated just by adding it to the validation suite. The LRZ-VS is open source and architecture independent and can easily be deployed on a large range of systems as the ScalaLife Computing Centres have already shown. For accessing the validation suite or using it together with your own software projects please visit the ScalaLife website at http://www.scalalife.eu.

## Acknowledgements
We kindly thank all the ScalaLife alpha users for providing the datasets [9].
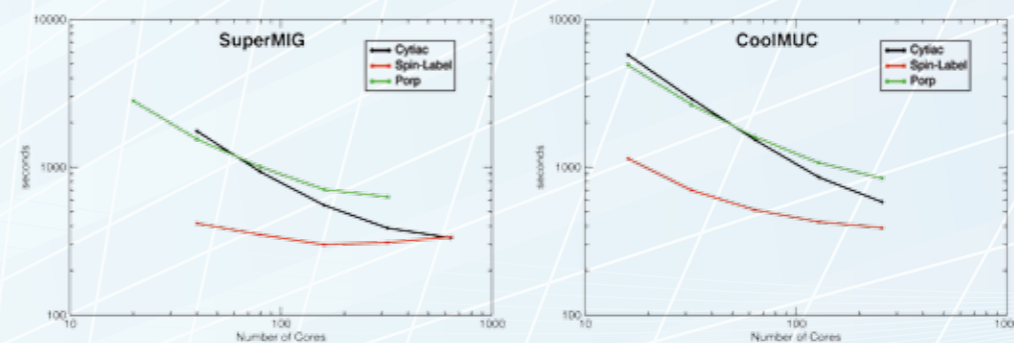


Figure 3: Dalton scaling results for three different test cases: Cytiac (black), Spin-Label (red) and Porp (green). The test runs were performed on two architectures at the LRZ: SuperMIG and CoolMUC.
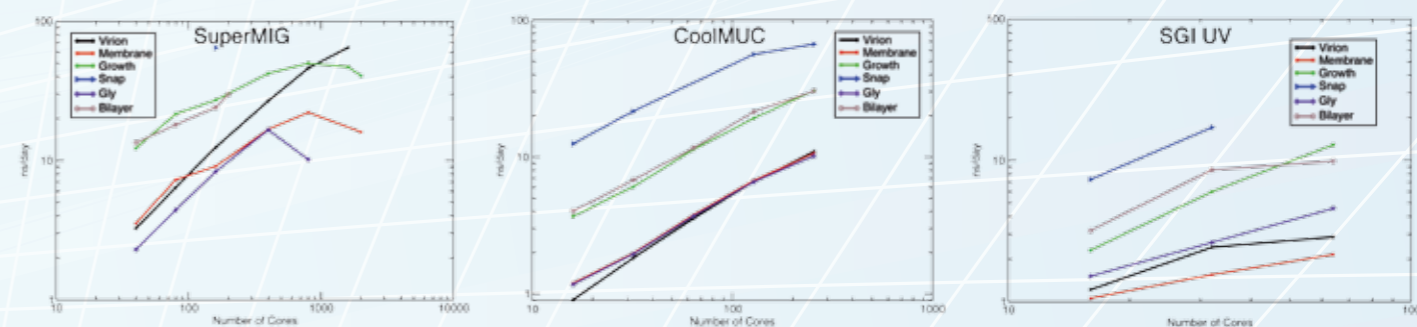


Figure 4: GROMACS scaling results for six different test cases: Virion (black), Membrane (red), Growth (green), Snap (blue), Gly (violet) and Bilayer (beige). The test runs were performed on three architectures i.e. SuperMIG, CoolMUC and SGI Altix UV.

## References
[1]   http://www.deisa.eu/science/benchmarking

[2]   http://www.fz-juelich.de/jsc/jube

[3]   http://www.scalalife.eu

[4]   **Frings, W., Schnurpfeil, A., Meier, S., Janetzko, F., and Arnold, L.**
      A Flexible, Application- and Platform-independent Environment for Benchmarking, Accepted for Proceedings Parallel Computing (ParCo), Lyon, 2009

[5]   **Berendsen, H.J. et al.**
      Comp. Phys. Comm. 91: 43-56, 1995

[6]   **Lindahl, E. et al.**
      J. Mol. Model. 7: 306-317, 2001

[7]   **van der Spoel, D. et al.**
      J. Comput. Chem. 26: 1701-1718, 2005

[8]   **Hess, B. et al.**
      J. Chem. Theory Comput. 4: 435-447, 2008

[9]   http:///www.scalalife.eu/validation

• Momme Allalen
• Marc Offman
• Helmut Satzger
• Ferdinand Jamitzky

Leibniz
Supercomputing
Centre (LRZ)

# Optimizing the Energy-to-Solution on SandyBridge Systems

The increasing popularity of Green Computing has qualified energy consumption and optimization as a primary concern. The key of energy efficient applications lies on the definition on energy models that minimize, or at least optimize, the energy consumption at any instant of the runtime of the applications [1]. Deeply related to these energy models are the energy policies provided by the linux kernel cpufreq infrastructure [2]. These policies allow managing the energy consumption based on frequency and processor utilization criteria.

We have evaluated how energy consumption changes while setting different governor and frequency parameters using the cpufreq kernel infrastructure and using the PAPI-RAPL component to measure the energy consumption. We evaluate also, how the performance and runtime of the application changes with the energy consumption of the running application, the energy to solution behaviour.

## CPUFreq Infrastructure

Starting with the 2.6.0 Linux kernel, the processor frequencies can be dynamically scaled through the cpufreq subsystem. This dynamic scaling of the clock speed gives some control in throttling the system to consume less power when not operating at full capacity. The cpufreq infrastructure makes use of governors (policies) and daemons for setting a static or dynamic power policy for the system. The dynamic governors can switch between CPU frequencies based on processor utilization to allow for power savings while not sacrificing performance.

There are five possible kernel governors available for use with the cpufreq subsystem. These governors set the processor frequency based on certain criteria [3]:

• Performance: The performance governor statically sets the processor to the highest frequency available. The range of frequencies available to this governor can be also per parameter adjusted.

• Ondemand: Introduced in the 2.6.10 kernel, the ondemand governor was the first in-kernel governor to dynamically change processor frequency based on processor utilization. The ondemand governor checks the processor utilization and if it exceeds a configurable threshold, the governor will set the frequency to the highest available.

• Conservative: Based on the ondemand governor, the conservative governor is similar in the way it dynamically adjusts frequencies based on processor utilization; however, the conservative governor allows a more gradual increase in power.

• Powersave: On the flip side of the performance governor, the powersave policy statically sets the processor to the lowest available frequency.

• Userspace: This governor allows the selection and setting a frequency manually. Note that the userspace governor itself does not dynamically change the frequency; rather, it allows a user-space program to dynamically select the processor frequency.

## PAPI-RAPL Component

PAPI (Performance Application Programming Interface) [4] aims to provide the tool designer and application engineer with a consistent interface and methodology for use of the performance counter hardware found in most major microprocessors. One of the new components of the PAPI library is the so called PAPI-RAPL component [5] that makes use of the RAPL [6] sensors available at the SandyBridge microarchitectures [7,8], sensors that allow to measure the power consumption of the CPU-level components by examining the MSR (Model Specific Registers).

The now implemented library joins the cpufreq and PAPI-RAPL features allowing to optimize the energy consumption of running instrumented applications by changing the frequency and governor policy of processors on certain instrumented code regions. These code regions are detected by Periscope, a scalable automatic performance analysis tool developed by the TU München [9,10]. Periscope encloses the code regions with the start_region(...) and end_region(...) function calls, functions used by the library to trigger the methods to change, if needed, the governor parameters or frequency of the process where the caller thread or process runs. In this case for example, we can increment the frequency of the processors that store the threads or processes that are running the loop sections and decrement it in I/O regions. But, not only that is possible; suppose for example that we want to optimize the energy consumed by an application that shows parallel load imbalance (see Fig. 1). The instrumenter will determine the code regions where the processors are idle whereas the library will set a lower frequency or a more saving energy policy for those processors, trying to minimize idle time on tasks/threads.

## Test Cases and Results

The library has been tested with three different applications with three different requirements. These applications are a $O^3$ matrix multiplication code developed in C language parallelized with MPI, a SIP (Strongly-implicit procedure) test bench application developed in Fortran language parallelized with OpenMP and the geophysical Seissol application implemented in Fortran code parallelized with MPI. The three applications have been instrumented with Periscope and then linked with the implemented and here explained library. The three applications have been run in a test cluster composed by eight SandyBridge EP Intel® Xeon® CPU E5-2680 at 2.70 GHz. Both applications parallelized with MPI have been run using 32 tasks, whereas the SIP application has been run using eight threads. For each code, we have defined seven different scenarios: criteria to change the governor and frequency of the processors depending on the region type determined by Periscope. These seven scenarios are specified in Table 1.
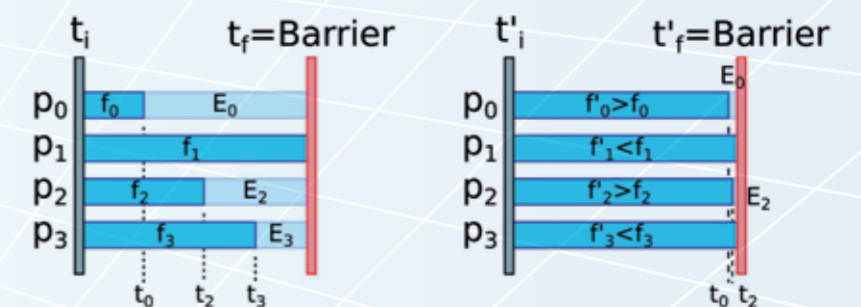


Figure 1: Suppose an application with a high load imbalance with a barrier at the end of a section, as the one on the left side. The processes that finish their code must wait until the barrier, consuming energy. The library will set a different frequency for each parallel process minimizing the idle time until the barrier for each task, as shown in the image of right side.

The runtime, energy consumed and cycles used by the three applications, regarding the seven different scenarios are graphically presented in the figures 2, 3 and 4 respectively.

|  | main region | sub, call, user | loop, for, forall | I/O region | parallel region |
|---|---|---|---|---|---|
| T1 | ondemand 2.7 | ondemand 2.7 | ondemand 2.7 | ondemand 2.7 | ondemand 2.7 |
| T2 | userspace 2.7 | userspace 2.7 | userspace 2.7 | userspace 2.7 | userspace 2.7 |
| T3 | ondemand 1.9 | ondemand 1.9 | ondemand 1.9 | ondemand 1.9 | ondemand 1.9 |
| T4 | userspace 1.2 | userspace 1.2 | userspace 1.2 | userspace 1.2 | userspace 1.2 |
| T5 | ondemand 1.9 | ondemand 1.9 | ondemand 2.7 | ondemand 2.7 | ondemand 2.7 |
| T6 | ondemand 1.9 | ondemand 1.9 | userspace 1.9 | ondemand 2.7 | userspace 1.2 |
| T7 | ondemand 1.9 | ondemand 1.2 | userspace 1.2 | ondemand 1.2 | userspace 1.2 |

Table 1: Governor and frequency constrains for the region type identified by periscope. Each set of constrains have been defined as an optimization profile or scenario test. The frequency specified for the ondemand governor refers to the maximal available one in GHz (the minimal in this case is always the minimal frequency 1.2 GHz), whereas the frequency specified for the userspace governor refers to a certain frequency in GHz to be used and not a maximal one. 2.7 GHz is the maximum frequency of the processors, 1.9 GHz is the medium frequency and 1.2 GHz refers to the minimum one.



Figure 2: Comparison between the energy to solution (Joules) of the three applications, depending on the energy optimization profile $T_i$. There are no results for the Seissol application for scenarios T5, T6 and T7 yet.



Figure 3: Comparison between the elapsed runtime (seconds) of the three applications, depending on the energy optimization profile $T_i$. There are no results for the Seissol application for scenarios T5, T6 and T7 yet.



Figure 4: Comparison between the consumed cycles of the three applications, depending on the energy optimization profile $T_i$. There are no results for the Seissol application for scenarios T5, T6 and T7 yet.

It is very easy to observe, taking into account the achieved results, that each application has its own energy behaviour; we cannot conclude with a golden rule for all the applications but, even so, the three applications have a similar global behaviour: making a comparison between T1 and T2 which use the same frequency but different governors, T1 seems to have a better performance, because it allows the processor to work at lower frequencies, saving energy when the processor really does not need the highest frequency. In the same way, T1 and T3 use the same governor policy but with a different frequency. In that case, T3 saves a little more energy but it also spends more time.

As it has been written before, a lower processor frequency does not necessarily imply reduced energy consumption. This behaviour can also be observed in those tables and figures. For example, T4 uses the same governor pattern as T2 but with the half of frequency. If this principle were right, the energy consumed by the T4 set should be approximately half of the energy consumed by T2, rule not fulfilled by the experiments. What it seems to have better results is setting a mixed combination of governors during the application runtime

(scenarios from T5 to T7), instead a fixed governor for the whole execution, although the procedure to change the governor and to set a new frequency also adds an extra overhead. Although it is not possible to find a direct correspondence between energy and runtime, looking at the results, there is somehow a linear relation that allows to determine a priori the behaviour of the cycles used by the application just from the curve of the energy consumption: the most overall energy is configured and used, the most cycles are needed for the same input dataset.

We can conclude that a same type of region might behave differently and therefore frequency policies will be set depending on the segment code and not per region. The fine tuning of the frequency in the code should then be possible when integrating this library with Periscope within the AutoTune project. AutoTune will extend Periscope, an automatic online and distributed performance analysis tool, with automatic online tuning plugins for performance and energy efficiency tuning. The resulting Periscope Tuning Framework will be able to tune serial and parallel codes with/without GPU kernels and will return tuning recommendations that can be integrated into
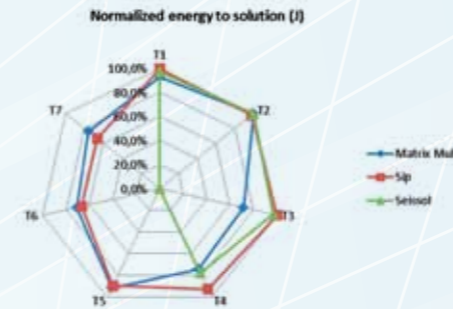
the production version of the code. The whole tuning process, consisting of automatic performance analysis and automatic tuning, will be executed online, i.e., during a single run of the application.

## References

[1]  **Merkel, A., Bellosa, F.**
Balancing power consumption in multi-processor systems, Proceedings of 1st ACM SIGOPS European Conference on Computer Systems, ACM, 2006

[2]  **CPU Frequency Scaling**
https://wiki.archlinux.org/index.php/CPU Frequency Scaling

[3]  **Using the Linux CPUFreq Subsystem for Energy Management.**
http://publib.boulder.ibm.com/infocenter/lnxinfo/v3r0m0/topic/liaai/cpufreq/liaai-cpufreq.htm

[4]  **Browne, S., Deane, C., Ho, G., Mucci, P.**
PAPI: A Portable Interface to Hardware Performance Counters, Proceedings of Department of Defence HPCMP Users Group Conference, 1999

[5]  **Weaver, V., Johnson, M., Ralph, J., Kasichayanula, K., Luszczek, P., Terpstra, D., Moore, S.**
Measuring Energy and Power with PAPI, The 1st International Workshop on Power-Aware Systems and Architectures, 2012

[6]  **David, H., Gorbato, E., Hanebutte, U., Khanna, R., Le, C.**
RAPL: memory power estimation and capping, Proceedings of the 16th ACM/IEEE international symposium on Low power electronics and design, ACM, 2010

[7]  **Intel Corporation**
Intel Xeon Processor, 2012.
http://www.intel.com/xeon

[8]  **Intel Corporation**
Intel 64 and IA-32 Architectures Software Developer Manual, 2012

[9]  **Gerndt, M., Fuerlinger, K., Kereku, E.**
Periscope: Advanced Techniques for Performance Analysis, International Conference ParCo 2005, 2006

[10] **Gerndt, M., Fuerlinger, K.**
Automatic Performance Analysis with Periscope, Concurrency and Computation: Practice and Experience, Wiley InterScience, John Wiley & Sons, Ltd., 2009

• Carmen Navarrete
• Carla Guillen
• Wolfram Hesse
• Matthias Brehm

Leibniz
Supercomputing
Centre (LRZ)

# ParaPhrase – Parallel Patterns for Adaptive Heterogeneous Multicore Systems

The ParaPhrase project aims to produce a new structured design and implementation process for heterogeneous parallel architectures, where developers exploit a variety of parallel patterns to develop component-based applications that can be mapped to the available hardware resources, and which may then be dynamically re-mapped to meet application needs and hardware availability. The project will exploit new developments in the implementation of parallel patterns that will allow us to express a variety of parallel algorithms as compositions of lightweight software components forming a collection of virtual parallel tasks. Components from multiple applications will be instantiated and dynamically allocated to the available hardware resources through a simple and efficient software virtualization layer. In this way, we will promote adaptivity, not only at an application level, but also at a system level.

Finally, we will develop virtualization abstractions across the hardware boundaries, allowing components to be dynamically re-mapped to either CPU/GPU resources on the basis of suitability and availability in a simple and straightforward way.

## Key Features

• Sustainable parallel computing through enhanced programmability and lower power consumption.

• Cost reduction in programmability and implementation of parallel systems.

• Better resource utilisation of parallel heterogeneous CPU/GPU architectures.

The work will enable major progress to be made in programming both current and future (parallel) computer systems. Using ParaPhrase technologies, we anticipate that we will be able to achieve significant parallel speedups for realistic applications and that these results will scale with larger systems.
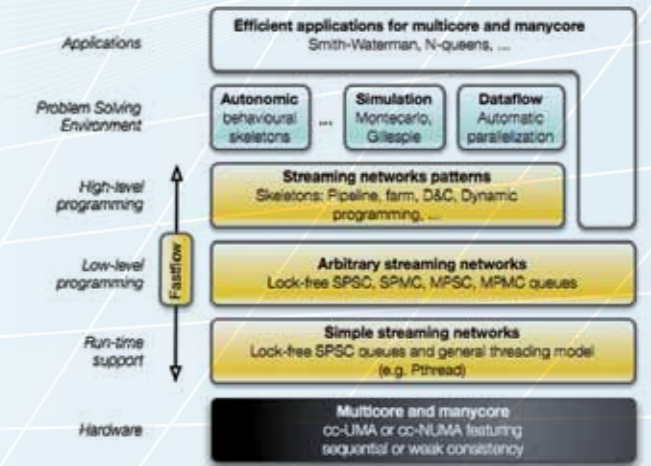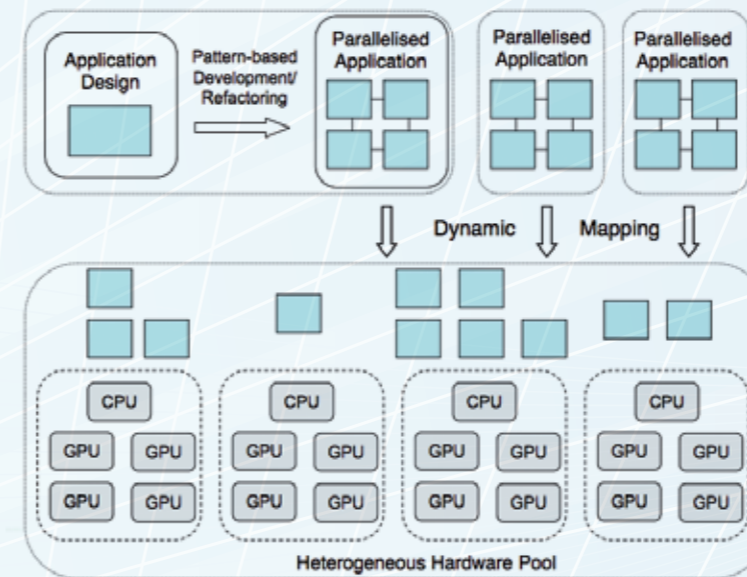
ParaPhrase is based on the FastFlow framework, which is developed, by the University of Pisa and Torino.

HLRS will initially define test kernels and testing frameworks. Together with other end-users, like SCCH, best practices will be adopted for the particular use cases and applications of the end-user. At a later stage, HLRS will port a full scientific application to the ParaPhrase model in order to evaluate it in an HPC production environment. Finally, HLRS will contribute to the optimization of placement of tasks, etc., on various current and future hardware architectures.

## Project Partners

The ParaPhrase Consortium consists of six academic and three industrial partners from five countries.

1. University of St. Andrews, United Kingdom (Coordinator)
2. Robert Gordon University, United Kingdom
3. Universität Stuttgart (HLRS), Germany
4. Universita Degli Studi di Torino, Italy
5. Universita di Pisa, Italy
6. Queen's University Belfast, United Kingdom
7. Mellanox Technologies, Israel
8. Erlang Solutions, United Kingdom
9. Software Competence Center Hagenberg, Austria

• The ParaPhrase Consortium
• Mathias Nachtmann[1]
• José Gracia[1]
• Colin W. Glass[1]

[1] University of Stuttgart, HLRS

# H4H – Hybrid for HPC

H4H – Hybrid for HPC, is an ITEA2 labeled European project co-funded by the German BMBF. It continues in the tradition of the ParMA ITEA2 Project.

## Aim of Project

The objective of H4H is to provide compute-intensive application developers with a highly efficient hybrid program-ming environment for heterogeneous computing clusters composed of a mix of classical processors and hardware accelerators. Future HPC architectures will rely heavily on non-uniformity and heterogeneity to achieve high perfor-mance. The process of incorporating different type of accelerators, each being optimal for a given type of code, has already started. However, we are lacking programming models, methods and tools to take full advantage of such platforms. H4H will leverage and con-sistently advance the state-of-the-art in several key areas: programming models and associated runtimes, per-formance measurement and correctness tools, intelligent mapping of processes / threads to hardware topology, dynamic automatic tuning and prediction of parallel execution time for a given appli-cation on different platforms.

HLRS is contributing to the provisioning of tools for correctness-checking and debugging. The first approach provides a tighter coupling of OpenMPI with the memory checking tools Valgrind and Intel PinTools with the goal to detect runtime race-conditions in the usage of communication buffers, as read-before-write, etc. The flexible implementation allows for different memory checkers to be used as modules on the OpenMPI OPAL layer. As additional benefit the Portable Hardware Locality tool (hwloc) now supports windows-based systems which allows application developers to check and exploit the layout of the un-derlying hardware.

A further correctness-checking tool is also based on Valgrind, This Valgrind extension allows memory checking on CUDA-based architectures. Using the wrapping functionality of the Valgrind framework on the CUDA Driver level allows to go beyond the existing capa-bilities of Valgrind and existing tools like cuda-memcheck.

Last but not least, Temanejo, the graphical debugger for task-parallel programming models developed at HLRS, will be extended and adapted to the latest family of the StarSs pro-gramming model. It allows a visual representation on task dependencies, tasks scheduling, and execution by the runtime environment. A tight integra-tion with conventional debuggers allows the developer to verify parallel tasking and sequential code simultaneously.

## Partners

The project has nine German partners:

- Frauenhofer SCAI
- GNS (Gesellschaft für Numerische Simulation mbH)
- GWT-TUD GmbH
- High Performance Computing Center Stuttgart (HLRS)
- INTES (Ingenieurgesellschaft für Technische Software GmbH)
- Jülich Supercomputing Center (JSC)
- MAGMA Gießereitechnologie GmbH
- RECOM Services GmbH
- ZIH, TU Dresden Zentrum für Informationsdienste und Hochleistungsrechnen

Another 14 partners come from the European countries Sweden, France, Spain.

- Thomas Baumann
- José Gracia
- Shiqing Fan
- Steffen Brinkmann

University of Stuttgart, HLRS

**PROJECT CONSORTIUM**

# The SIOX Project



High Performance Computing has become an essential research tool in many scientific fields. The computing power of HPC systems continues to increase exponentially. At the same time, the capabilities of I/O hardware are not keeping track. To keep up with the application needs, the number of I/O devices used with today's HPC systems is continuously increasing. Furthermore, additional technologies like caching are used to further optimize performance. This growing complexity of the I/O system together with the complexity of the necessary global parallel file systems make it very hard to locate the reason for and understand performance issues in the I/O area.

In this context, a global optimization of the I/O system to the needs of all applications turns out to be very difficult or even impossible. In part due to the disparate nature of the requirements and in part because there is currently no way to identify abnormal I/O behaviour and to trace it back to its source. Moreover, the I/O system is typically not used exclusively by an application. Therefore, the interplay of several applications performing I/O concurrently further complicates the scenario.

Additionally, it is difficult to predict whether the identified I/O-patterns used by a specific parallel application will cause a performance problem or not.

## Goals of the Project

SIOX' main goal is to gain an overview of all I/O activity taking place on a HPC system, relate this I/O activity to the running applications and to use this information to optimize I/O.

Initially, the project's scope spans the development of standardized interfaces and an environment to collect, reduce, relate and store performance data from all relevant software and hardware layers (see Fig. 1). All the created components will be based on open interfaces. This opens for example the possibility to use SIOX with different file systems and different storage hardware.

The collected information will be stored in the SIOX database and it will be analyzed and correlated with previously observed access patterns in order to gain an understanding of the characteristics and causal relationships. The data collection can be controlled according to current requirements. This allows

for collecting extended information when specific I/O issues are investigated.

For I/O optimization, the efficiency of the observed access patterns is to be estimated through an automatic analysis of the data, followed by the proposal of possible improvements based on the knowledge accumulated in the SIOX' knowledge base.

For applications making use of MPI I/O, these improvements will be realized by giving this information back to the Open MPI I/O library. There, the most appropriate I/O algorithm for the request will be chosen accordingly. In addition, the available control parameters will be set to best fitting values.

The obtained data collection gives in addition an overview of the I/O requirements of all applications which have been using the HPC system over time. This will lead to a better understanding of the overall usage of the I/O system and it allows for a better sizing of future HPC systems.

SIOX has started in July 2011 with a duration of three years and is funded by the German Ministry of Research.
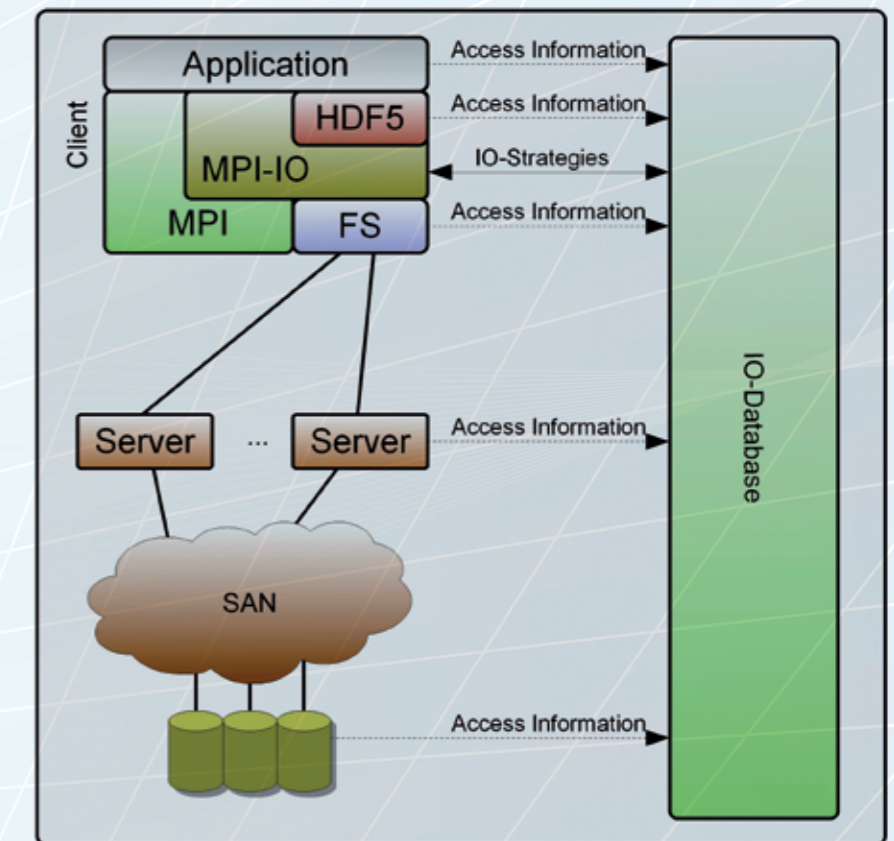
Figure 1: SIOX is collecting data from all relevant software and hardware layers and provides Open MPI I/O with information to optimize parallel I/O requests.

• Thomas Bönisch

University of
Stuttgart, HLRS

## Project Partners
- Center for Information Services and High Performance Computing (ZIH), Technische Universität Dresden
- Universität Hamburg, Chair in Scientific Computing
- High Performance Computing Center

The project is also supported by the following associated partners:

- German Climate Computing Center (DKRZ GmbH), Hamburg
- IBM Germany GmbH, Ehningen

Project web page:
http://www.hpc-io.org

# Performance Dynamics of Massively Parallel Codes

The deployment of adaptive algorithms and sophisticated load-balancing schemes make the execution behavior of simulation programs increasingly dynamic. Effective application optimization therefore requires capturing and analyzing performance data along the dimensions of both space and time. While existing performance analysis tools typically provide detailed information along spatial dimensions like processes and nodes, the aspect of performance dynamics, that is, the evolution of performance metrics along the time dimension, has so far been neglected. Additional insight that can be obtained by taking advantage of performance dynamics is depicted in Fig. 1.

To support the optimization of applications with time-dependent performance characteristics, the LMAC project (Leistungsdynamik massiv-paralleler Codes), funded under the second call of the BMBF program "HPC-Software für skalierbare Parallelrechner", aims to extend the well-established performance-analysis tools Vampir, Scalasca, and Periscope with new functionality to measure and analyze performance dynamics. Vampir is an interactive trace browser whose particular strength is the detailed visualization of the interactions between the different processes of a parallel program, offering highly flexible views to the user. Scalasca, which has been specifically designed for large-scale systems, integrates efficient performance summaries with the ability to automatically identify wait states that occur in simulation codes, for example as a result of unevenly distributed workloads. Whereas the first two tools analyze the performance data postmortem, that is, after the parallel program has been terminated, Periscope characterizes the performance properties of an application and quantifies associated overheads already at runtime.

A major portion of the new capabilities will be integrated in Score-P, a measurement infrastructure shared by all of the above-mentioned tools, which was created in the SILC project (2009-2011) funded under the first call of the same BMBF program. Score-P was designed for scalability and easy of use, supports profiling, event tracing, and online analysis of HPC applications. It enables enhanced interoperability between the end-user tools and will soon replace their proprietary measurement systems (see Fig. 2). The organization of Score-P as a community project is another task in LMAC.

To serve a broad user base, both, within the Gauss Alliance and beyond, most of the performance tools are released free of charge to the community under an open-source license. Only Vampir, due to its sophisticated user interface, is distributed commercially. The software products are accompanied by training- and support offerings through the Virtual Institute – High Productivity Supercomputing, and will be maintained and adapted to emerging HPC architectures and programming paradigms beyond the original project duration.

The academic project partners in LMAC are the German Research School for Simulation Sciences as the coordinator, the Jülich Supercomputing Centre, RWTH Aachen University, TU Dresden and TU Munich. GNS mbH, a private company that specializes in services related to metal forming simulations, such as mesh generation for complex structures and finite element analyses, joined the project as an industrial partner. In addition, the University of Oregon, an associated partner, complements the LMAC objectives with corresponding extensions to the performance tool TAU.

For more information see:
http://www.vi-hps.org/projects/lmac
and http://www.score-p.org

- Christian Rössel [1]
- Bernd Mohr [1]
- Michael Gerndt [2]
- Felix Wolf [3]

[1]  Jülich Supercomputing Centre (JSC)

[2]  Technische Universität München

[3]  German Research School for Simulation Sciences, RWTH Aachen

Figure 1: Runtime-distribution and point-to-point communication times over iterations and processes: (a) Runtime, (b,c) Point-to-point communication. The different colors in (b) correspond to the minimum, the median, and the maximum over all processes of a single iteration. In (c) the color coding denotes the entire distribution over space-time (darker means more). The rise in the maximum communication time in (b) is the cause for the increased iteration-runtime in (a).



Figure 2: Interaction of the tools Persicope, Scalasca, Vampir, TAU and Scalasca's profile browser CUBE via the data exchange formats OTF2 for traces and CUBE4 for profiles.

# PRACE – Third Implementation Phase Project started

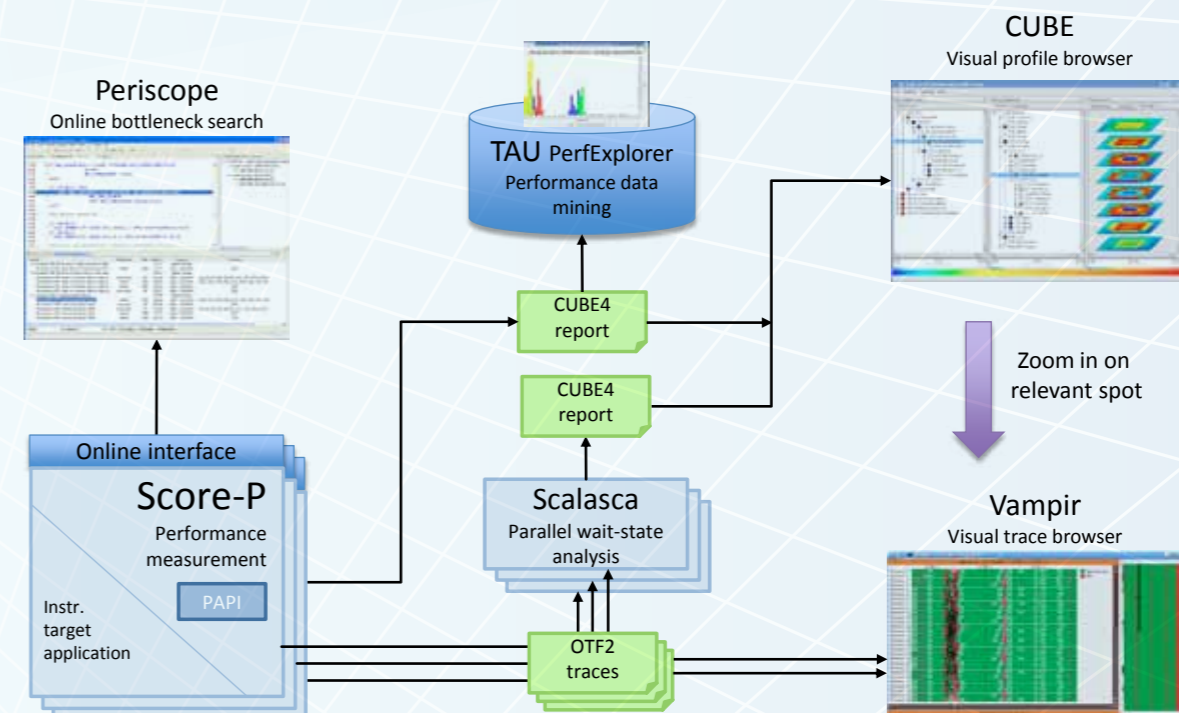To support the accelerated implementation of the Research Infrastructure established by the Partnership for Advanced Computing in Europe (PRACE) the European Commission issued a third call for a proposal in 2011. PRACE partners from 25 countries submitted a successful proposal and started the **Third Implementation Phase** project (PRACE-3IP) on July 1, 2012.

Key objectives of PRACE-3IP are:
• Support for the PRACE AISBL through the development and optimization of processes for collaborations between PRACE and other organizations; management procedures and related tools; impact assessment, management of the DECI (Distributed European Computing Initiative providing access to PRACE Tier-1 resources) call logistics.

• Dissemination and outreach continuing the production of promotional material, presentation of project results through the PRACE web site, electronic news letters, printed digests and at major conferences and exhibitions (e.g. ISC, SC, ICT). Establishing a Summer of HPC programme to encourage students to seek careers in HPC through internships at HPC centres:

• Intensify the successful training of PRACE by transforming the PRACE Advanced Training Centres (PATC) into a permanent service, organizing face to face training classes, and improving the PRACE training portal.

• Provide services for industrial users and small and medium enterprises (SME). Develop models and best practices for industrial users including an integrated access programme and issue pilot projects for SMEs.

• Continue the operation and coordination of the distributed infrastructure through integration of new and upgraded Tier-0 and Tier-1 systems. Support new partners in the deployment of common services and enhance the existing ones.

• Scale and optimize application codes, especially for the Preparatory Access and DECI projects. Investigate and exploit new HPC tools and techniques including those developed in other petascale and exascale projects. Produce best practice guides to support European research communities.

• A completely new undertaking is a pilot on Pre-Commercial Procurement (PCP) which is an approach for procuring R&D (Product Driven Research and Pre-Commercial Development) services. The scope of the PRACE PCP pilot will be a Whole System Design for Energy Efficient HPC. Some PRACE partners, including the AISBL, form the PCP Group of Procurers that manages and co-funds this pilot complementing the EC funding of up to 4.5 Mio Euro with an equal amount. The intention is to demonstrate that PCP can advance innovation in Europe through public procurers. The pilot will also help to understand if PCP is suitable to advance the European HPC industry.

PRACE-3IP is again managed by Forschungszentrum Jülich. It has a budget of nearly 27 Mio Euro including an EC contribution of 19 Mio Euro under grant agreement RI-312763. The duration will be 24 months except for the PCP work that is planned for 48 months due to the complexity of the related multi-stage process.

Over 300 researchers collaborate in PRACE from 25 countries and 45 organizations that receive funding from the European Commission; this includes 26 Beneficiaries and 19 Third Parties from associated universities or centres. 180 collaborators met in Paris from September 5-7, 2012, for the PRACE-3IP kick-off and a joint PRACE-2IP all-hand meeting. The meeting was organized by GENCI and held at the Institut de Physique du Globe de Paris.

## Synopsis of the PRACE Projects

The European Commission supported the creation and implementation of PRACE through four projects with a total EC funding of 67 Mio Euro. The partners co-funded the projects with over 110 Mio Euro in addition to the commitment of 400 Mio Euro by the hosting members to procure and operate Tier-0 systems and the in-kind contribution of Tier-1 resources on systems at presently twenty partner sites. The following table gives an overview of the PRACE projects.

• Dietmar Erwin

Forschungszentrum Jülich



Figure 1: PRACE collaborators at the PRACE-3IP Kickoff.         © PRACE / L. Godart

| Project Id Grant Number | Number Partners | Budget Mio. € | EC Funds Mio. € | Duration | Status |
|---|---|---|---|---|---|
| PRACE-PP RI-211528 | 16 | 20.1 | 10.0 | January 1, 2008 - June 30, 2010 | completed |
| PRACE-1IP RI-261557 | 21 | 28.7 | 20.0 | July 1, 2010 - June 30, 2012 | completed – except for the Evaluation of Prototypes in WP9 till March 31, 2013 |
| PRACE-2IP RI-283493 | 22 | 35.4 | 18.0 | September 1, 2011 - August 31, 2013 | in progress |
| PRACE-3IP RI-312763 | 26 | 26.8 | 19.0 | July 1, 2012 - June 30, 2014 | started PCP to end June 30, 2016 |
| (Total) | | 99.6 | 77.0 | | |

# JUQUEEN: The most recent Blue Gene Architecture at Jülich

Blue Gene/Q is the third generation in a line of highly scalable architectures. It is similar to the previous designs as very power-efficient processors are integrated very densely. But, due to an immense increase of parallelism at the single node level, a significant boost of performance is achieved. The network still being implemented on the chip remains a unique feature of this architecture. This is a key ingredient for being able to scale performance up to a huge number of cores.
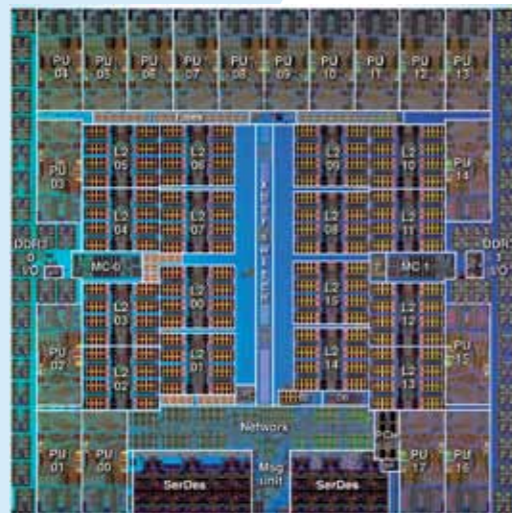
The Blue Gene/Q processor [1] comprises 16 cores for executing the user's application. There is an additional core for operating services. By decoupling the execution of system services, effects of random asynchronous delays of user processes are suppressed. Such noise can significantly deteriorate the scalability of an architecture. Yet another, redundant, processor core is on the chip, but not available in practice; it has primarily been added to increase the manufacturing yield. With each of the cores being 4-way multi-threaded, up to 64 user threads (or processes) can run at the same time.

The processor core architecture is relatively simple and implements a standard 64-bit power instruction set architecture. A particular feature of this core architecture is the support of an auxiliary execution unit. For Blue Gene/Q a Quad Floating-Point Processing Unit (QPU) had been developed. The QPU processes vectors of four 64-bit elements. In each clock cycle it can

perform four fused multiply-add operations in parallel. Like in previous generations of Blue Gene, the vector arithmetic instructions may involve a permutation of the vector elements, which is used to implement complex arithmetics (without the need of separate shuffle operations like in other vector instruction set architectures). With each of the 16 QPUs being able to complete four multiply-add operations per clock cycle at a clock speed of 1.6 GHz, the peak performance is 204.8 GFlop/s.

This huge amount of performance can only be exploited if it is balanced by a powerful memory subsystem. As can be seen in Fig. 1 a large fraction of the

be fetched from L2 cache or external memory, the hardware supports pre-fetching. Like in other processor architectures, such pre-fetching can be triggered by an automatic detection of a sequence of contiguous set of memory addresses being accessed. Beyond this stream pre-fetching scheme there is a new scheme called list pre-fetching. This feature may be exploited by applications where a sequence of (possibly non-contiguous) addresses is accessed many times. Controlled by the user the list pre-fetch engines can record access patterns. When the execution of the relevant code section is repeated the list of recorded addresses is used for pre-fetching the data.
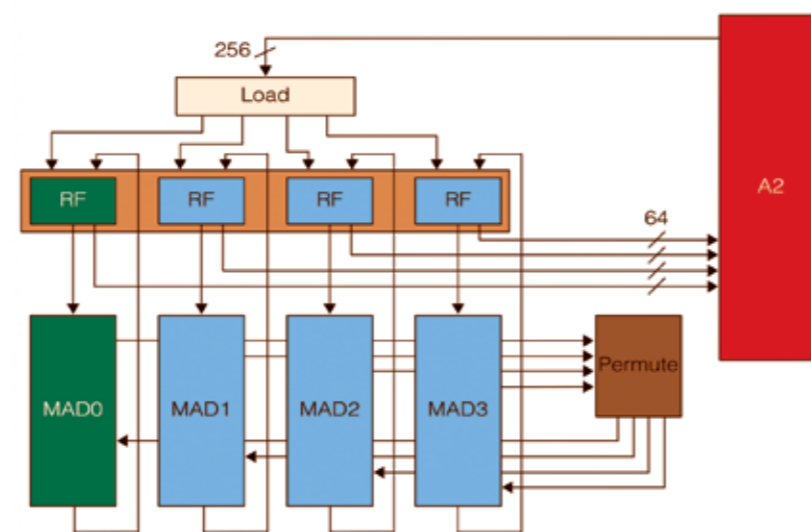
Figure 1: Physical layout of the Blue Gene/Q chip.



Figure 2: The Quad Floating-Point Processing Unit (QPU) is an extension of the A2 processor core to accelerate floating-point computations [1]. Operands stored in the register file (RF) are processed in 4 fused multiply-add dataflow (MAD) pipelines.





Figure 3: A Blue Gene/Q processor (left) is mounted to a heat spreader. A node-card can host 32 nodes. With 16 node-cards per midplane and two midplanes per rack the total number of nodes per rack (right) is 1,024.

die space is occupied by the L2 cache, which has a capacity of 32 MBytes. Data is moved between external memory and this last-level cache by two memory controllers (MC 0 and MC 1). The L2 cache is shared by all processor cores. A central crossbar switch connects it to all cores plus the network subsystem. The cores can read from the L2 cache at an aggregate maximum bandwidth of 409.6 GByte/s.

We would like to highlight several features of the memory hierarchy. To hide the latencies, which occur after an L1 cache miss because data needs to

One feature of the Blue Gene architecture line is the dense integration of a large number of nodes within a single rack (see Fig. 3). Unlike in previous generations of Blue Gene, where air was used to remove the heat generated by the compute nodes, in the new generation of machines the nodes are directly connected to a liquid cooling system. 90% of the heat originating from the compute nodes in the system is directly taken away by the water, the remaining fraction, coming mainly from the power supplies, is still moved out of the rack by air.
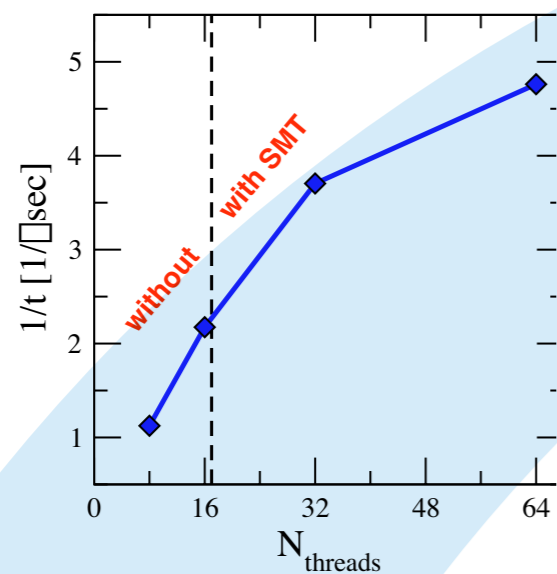
Figure 4: Inverse execution time of a fluid dynamics application running on a single node as a function of the number of threads [6]. The result shows significant gain from Simultaneous Multi-Threading (SMT).

The sustained performance obtained with the High-Performance Linpack (HPL) benchmark is similar to previous generations of Blue Gene machines, i.e. slightly more than 80% of the peak performance. This high level of device utilization in combination with a high power efficiency allowed for a significant leap on the Green500 list [4]. In this list systems are ordered according to the ratio speed achieved with the HPL benchmark versus power consumed by the system during benchmark execution. Previous number one on this list, QPACE [5], a special purpose computer operated at Forschungszentrum Jülich, achieved 0.77 GFlop/s per Watt. With Blue Gene/Q the score went up by almost a factor three to 2.1 GFlop/s per Watt.

The HPL benchmark performance of more than 80% demonstrates that it is at least for some applications feasible to obtain a high level of sustained performance, similar to previous generations of Blue Gene. However, there are challenges for the application developer to efficiently exploit the performance of this architecture. The most notable challenge is the increased parallelism at various levels. Compared to Blue Gene/P the vector units doubled their width, the number of cores and threads per node increased by a factor of four and 16, respectively. Within the Exascale Innovation Centre (EIC) the utilization of performance features has been investigated already at a stage when only prototype hardware had been available.

Fig. 4 shows the inverse execution time (which is proportional to the performance) of a fluid dynamics application based on the Lattice Boltzmann method. In case of ideal scaling the performance increases proportional with the number of threads until each core executes one thread. A significant performance increase is expected until each core executes two threads since one hardware thread can only schedule an instruction every second clock cycle. For this particular application a noticeable further increase of performance could be achieved by placing up to four threads on each core.

The first racks of Blue Gene/Q have been delivered to Forschungszentrum Jülich in April 2012. Only a few weeks later JUQUEEN consisting of eight racks could be made accessible to users for early production runs, see Fig. 5, delivering a peak performance of 1.6 PFlop/s, which is 60% more than JUGENE (72 racks Blue Gene/P, 1PFlop/s peak performance).

Even though JUGENE was an extreme reliable and stable system over the last years - achieving a user job utilization of over 90% - the system had to be shut down end of July and dismounted. Once the space, power, and cooling capacity occupied by JUGENE were freed, the extension of JUQUEEN to its final configuration with 28 racks is taking place. This translates into a more than fivefold increase of compute performance on Blue Gene available for scientific computing in Jülich, Germany and Europe.

A large fraction of the Blue Gene/Q installation at Forschungszentrum Jülich, JUQUEEN, will be available for users of the Gauss Centre for Supercomputing and PRACE.

## References

[1] **Haring, R.A. et al.**
The IBM Blue Gene/Q Compute Chip, IEEE Micron, Vol. 32, No. 2, March/April 2012

[2] **Chen, D. et al.**
The IBM Blue Gene/Q Interconnection Network and Message Unit, SC '11 Proceedings, 2011

[3] **El-Harake, H. et al.**
Juniors: Prototypes for novel exascale I/O concepts, ISC'12 Poster, 2012

[4] http://www.green500.org

[5] **Baier, H. et al.**
QPACE: A QCD parallel computer based on Cell processors, PoSLAT 2009 (2009) 001

[6] **Pozzati, F.**
Porting and optimization of a Lattice Boltzmann D2Q37 code to Blue Gene/Q, Technical Report FZJ-JSC-IB-2011-06, Forschungszentrum Jülich, 2011

• Dirk Pleiter
• Michael Stephan

Jülich Supercomputing Centre (JSC), Forschungszentrum Jülich



Figure 5: 8-rack JUQUEEN at Forschungszentrum Jülich.

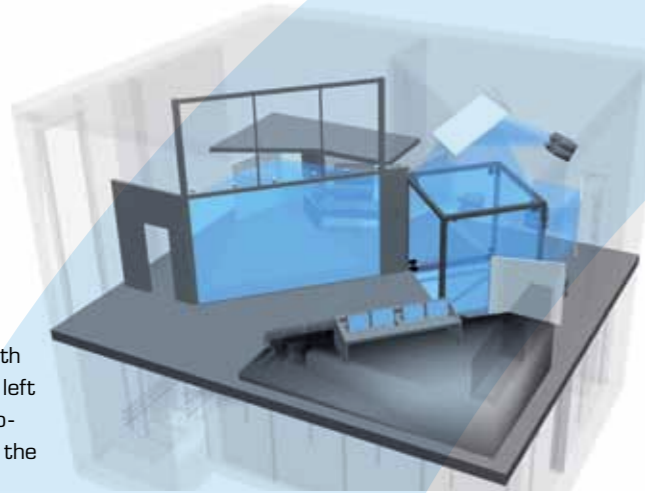# Virtual Reality and Visualization at the Leibniz Supercomputing Centre



Figure 1: The V2C with the powerwall at the left and the 5-sided projection installation at the right.

Due to the growing demand for Virtual Reality (VR) and Visualization services in the recent years the Leibniz Supercomputing Centre (LRZ) has decided to establish a Centre for Virtual Reality and Visualisation (V2C).

Whereas VR as a technology tries to immerse the user by providing highly interactive, real-time, stereoscopic images based on the user's viewpoint, visualization techniques process simulation data, measured data, or abstract data. These data sets can ideally be displayed using VR technology.

## Hardware
The V2C hosts a large scale, high resolution powerwall installation, as well as a 5-sided projection installation, both equipped with optical tracking. Tracking is needed to suffice the basic demands of any VR application and to allow easy and intuitive data exploration by adapting the projection of the scene based on the head position and orientation.

The installations are set up in a room of the size of 12x12x11 m having a raised floor at the height of 5,20 meters. The lower part of the V2C contains the computing equipment, electricity and cooling infrastructure, whereas the upper area is used as visitor and user space.

## Powerwall Installation
The powerwall is typically observed from a cinema-like area consisting of 21 seats at 3 elevation levels. With its seating setup it can be used as a simple presentation medium, while its additional tracking capabilities allow its application for interactive visualizations.

## 5-Sided Projection Installation
Based on the concepts of Carolina Cruz-Neira's CAVE (CAVE Automated Virtual Environment) [1] a 5-sided projection installation, consisting of three back-projected sides, a back-projected floor and ceiling was built. Mirrors are used to reduce the spatial demands.

Whereas the powerwall, although it has the capabilities of interaction with its optical tracking system, is used mainly as a presentation medium, the 5-sided projection medium provides full support for interactive applications especially using finger tracking. With the setup of five display panels it is easily possible to fully cover the field of view of the user and immerse him in the environment.

## Technical Data
An overview on the key specification of the two systems is given in the table.

## Projects
As usual with VR and visualization, the application domains are extremely diverse. Where many centres only focus on the visualization of simulation results, the LRZ supports users of classical virtual environment applications as well. Typical projects come from the HPC users of life sciences but also from arts and humanities.

Some example projects currently worked on in the V2C are terrain visualization, logistics simulation, architecture informatics, information visualization, archaeology, new media art, geo engineering and zoology.

When supporting users and cooperating in their domain specific project, often additional competence is needed and corresponding research is undertaken by the team at LRZ to improve existing visualization techniques, navigation and interaction with the scenes.

Of course the installations are set up in a way that they can communicate via a 10GE network directly with the recently installed LRZ supercomputer SuperMUC (Europe's most powerful supercomputer), but also with other systems world-wide. On the one hand, data display via remote visualization and computation is possible, on the other hand traditional VR and visualiza-



Figure 2: An example from the area of archaeology - visualization of the tomb of Karaburun (Visualization by Andreas Hartmann with photos by Roy Hessing).

tion tools can be used and are able to have an extremely fast access to the computation results of the simulations stored on SuperMUC or elsewhere.

The V2C supports over thirty software packages from the domain of VR and Visualization in order to support the realization of the different project types.

A detailed overview on the installed VR and visualization software, the hardware setup and the running projects can be found at http://v2c.lrz.de

## References
[1]  **Cruz-Neira, C., Sandin, D.J., Defanti, T.A., Kenyon, R.V., and Hart, J.C.**
The cave: Audio visual experience automatic virtual environment, Communications of the ACM, 1992, 35, 64-72

• Christoph Anthes
• Dieter Kranzlmüller

Leibniz
Supercomputing
Centre (LRZ)

| Device | Powerwall | 5-Sided Projection Installation |
|---|---|---|
| **Display** | | |
| Dimensions | 6.00 m x 3.15 m | 5 x 2.70 m x 2.70 m |
| Stereo-separation | Circular Polarisation | Temporal |
| Projectors | 2 x Sony 4K | 10 x Christie Mirage WU3 DLP |
| Resolution | 16.5 MP in stereo | 41.4 MP in stereo |
| Graphics Boards | 2 x nVidia QuadroPlex | 2 x nVidia Quadro 6000 |
| **Compute Hardware** | | |
| Machine | SGI Altix UV10 | SGI Altix XE500 (12 nodes) |
| Memory | 256 GB | 48 GB per render node |
| Tracking | ARTtrack2 (6 cameras) | ART TrackPack4 (8 cameras) |

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

## Compute servers currently operated by **LRZ** are given in the following table

| System | Size | Peak Performance (TFlop/s) | Purpose | User Community |
|---|---|---|---|---|
| IBM System x iDataPlex "SuperMUC" (expected for June 2012) | 155,656 Cores 330 Tbyte | 3,263 | Capability Computing | German Universities and Research Institutes, PRACE Projects |
| Linux-Cluster Intel Xeon EM64T/ AMD Opteron 2-, 4-, 8-, 16-, 32-way | 4,438 Cores 9.9 TByte | 36 | Capacity Computing | Bavarian and Munich Universities, D-Grid, LCG Grid |
| SGI ICE Intel Nehalem 8-way | 512 Cores 1.6 TByte | 5 | Capacity Computing | Bavarian Universities, PRACE |
| SGI UV | 2,080 Cores 6.1 TByte | 20 | Capability Computing | Bavarian Universities, PRACE |
| Megware IB-Cluster | 2,848 Cores 2.8 TByte | 22 | Capability Computing, PRACE Prototype | Bavarian Universities, PRACE |

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

Draft of LRZ's next supercomputer "SuperMUC" by IBM.

# H L R S

## 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 organization SICOS GmbH.

## 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. h.c. 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 HLRS

| System | Size | Peak Performance (TFlop/s) | Purpose | User Community |
|---|---|---|---|---|
| Cray XE6 "Hermit" (Q4 2011) | 3,552 dual socket nodes with 113,664 AMD Interlagos cores | 1,045 | Capability Computing | European and German Research Organizations and Industry |
| NEC Intel Cluster | 4,700 Intel Nehalem cores + 4,600 Sandybridge cores, 10 TB memory and 64 NVIDIA Tesla S1070 | 170 | 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 |

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 Steering Committee of the Gauss Centre for Supercomputing (GCS) and by the John von Neumann Institute for Computing (NIC), the latter 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.

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

**Higher education** of bachelor, master and doctoral students.

### Contact:

Jülich Supercomputing Centre (JSC)
Institute for Advanced Simulation (IAS)
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

Draft of JSC's supercomputer JUQUEEN, an IBM Blue Gene/Q system.

## Compute servers currently operated by JSC

| System | Size | Peak Performance (TFlop/s) | Purpose | User Community |
|---|---|---|---|---|
| IBM Blue Gene/Q "JUQUEEN" | 28 racks 28,672 nodes 458,752 processors IBM PowerPC® A2 448 Tbyte main memory | 5,872 | 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" | 206 nodes with 2 Intel Westmere 6-core 2.66 GHz processors each 412 graphic processors (NVIDIA Fermi) 20.0 TByte memory | 240 | Capacity and Capability Computing | selected HGF Projects |

# New Division "Civil Security and Traffic" at JSC

The last decades have faced a continuously growing urban population and cities. Larger and complex buildings as well as an increasing demand on public transportation systems accompany this growth. Additionally, the frequency and extend of large public events has increased. The study of dense crowds and related safety topics gains even more importance considering the fact that since 2006 half the world population lives in cities. Cities are confronted with dense crowds, dense traffic, and mixtures of traffic types. The potential hazards related to dense crowds become obvious and real, like during the Loveparade 2010 in Duisburg (Germany). It is important to assist urban planners, architects, and safety engineers confronted with the cities' growth and the conception of future cities. Therefore, there is a high demand for new tools which consider the complexity of all involved processes.

In recent years, a working group has formed at the Jülich Supercomputing

Centre (JSC) that is developing simulation models in the areas of fire safety and pedestrian dynamics. To structurally support further developments, a new division "Civil Security and Traffic" has been established at JSC in May 2012. The division is composed of scientists from mathematics, physics, computer sciences, and engineering. This interdisciplinary team develops and applies models and simulations of complex systems in the context of civil security, fire safety, and traffic planning. In combination with high performance computing it is possible to tackle challenges in the simulation of large systems using high fidelity models. The division intensively cooperates with universities around the world, research institutions as well as companies.

At the same time the demand of research on civil security is brought into the focus of the German Government, Federal Ministry of Education and Research (BMBF) and Helmholtz Association (HGF). In cooperation with

the Karlsruhe Institute of Technology (KIT) and German Aerospace Center (DLR), the new division will continue its work in the HGF portfolio project "Security Research". In addition, many

Along with interesting self-organization phenomena there is a multitude of applications, like the evaluation of escape routes in the context of crowd management or the optimization of pedestrian



Figure 2: Simulation of fire propagation in a train compartment. Shown is a three dimensional representation of the flame distribution.

programmes, on European level as well, have been launched with the focus on citizens and their protection. The European Union's current Seventh Framework Programme has many sections directly dedicated to civil security. The next EU programme "Horizon 2020", starting in 2013, will also address the security of citizens. At the national level, the field "Protection and rescue solutions" is one of the funding priorities in the Federal Government's security research programme.

A summary of the three main research areas of the newly founded division is presented in the following.

## Pedestrian Dynamics
The investigation of pedestrian dynamics is a young and lively field of research.

facilities for urban development. Our aim is the quantitative description of pedestrian dynamics by using microscopic models of self-driven particle systems. For model validation the group cooperates with several universities on a systematic enhancement of the empirical data basis. The group has also taken a pioneering role in the conception and execution of large-scale laboratory experiments involving pedestrians. This covers the automatic extraction of information (for instance trajectories) from experiments' video footages using pattern recognition techniques [1] and the analysis and evaluation of the information using high standard methods based on Voronoi decompositions [2]. The field is completed by the development of highly accurate models for pedestrian dynamics [3,4]
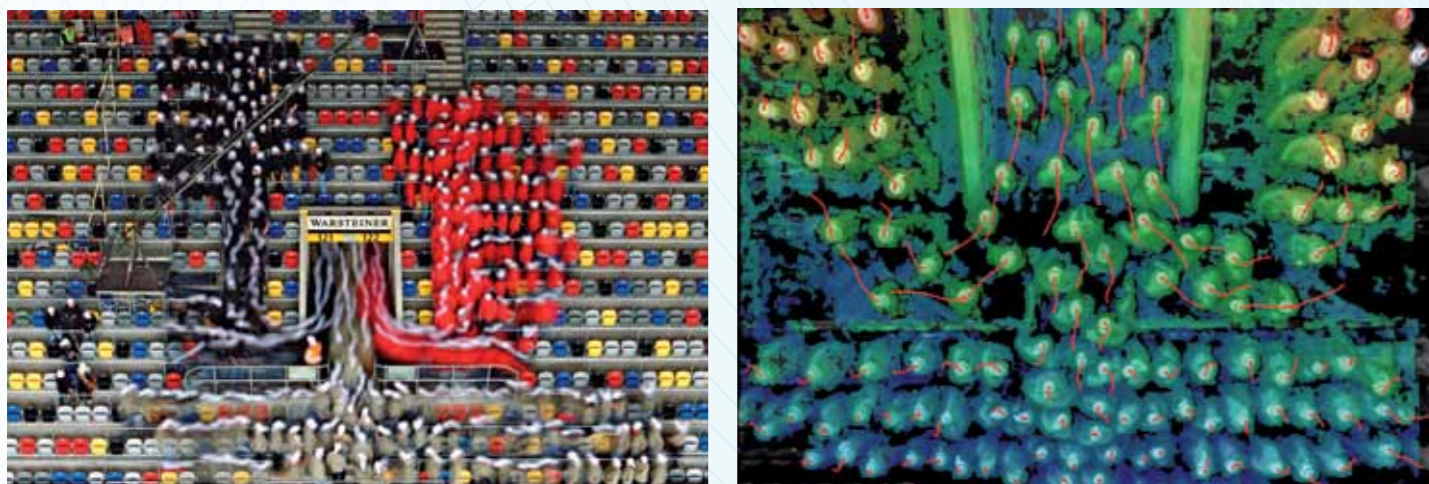


Figure 1: Left: Merging of three pedestrians streams at the gate of the tribune area of a stadium. The right picture shows the automatically extracted trajectories using the software PeTrack [5] developed at JSC.

which are used in simulations to reproduce the observed phenomena. Furthermore, to increase the transparency of the research activities and promote a sustainable development, the models and analysis tools developed so far are available to the scientific

of the available fire safety software to use High Performance Computing facilities is very limited, we investigate ways to improve the parallelization of these codes. This effort will allow us to simulate large structures at an adequate numerical resolution.

bikes increasingly enrich the traffic in cities. However, the consequences of this development for the safety and comfort on the roads have so far not been adequately investigated. In cooperation with the University of Wuppertal, initial laboratory experiments with bicycles have been conducted. Based on the results of such experiments, we compare different means of transport and provide tools to support the planning of routes for mixed traffic. It is planned to extend aforementioned activities from local evacuation of buildings or venues to evacuation of cities or regions. In the latter, the consideration of intermodal traffic is indispensable, for instance by modelling trips starting as pedestrian and using then a combination of various transports means such as cars, trains or bus shuttles.

## Projects
The research activities conducted so far have been supported by two BMBF and DFG-funded projects, which were successfully completed in 2011. The BMBF-funded project Hermes [6] covered the design and implemented an evacuation assistant for large public events. The test venue was the ESPRIT arena in Düsseldorf (Germany). Succeeding Hermes, a new BMBF-funded research project BaSiGo [7] has been launched in March 2012. Beside other topics, this project studies the emergence of critical states in large crowds. For that purpose large experiments are planned.

## References
[1]   **Boltes, M., and Seyfried, A.**
      Collecting Pedestrian Trajectories, Neurocomputing, Special Issue on Behaviours in Video, ISSN 0925-2312, 2012

[2]   **Zhang, J.**
      Experimental study on characteristics of pedestrian traffic with Voronoi-based segmentation method, PhD Thesis, University of Wuppertal, 2012

[3]   **Kemloh Wagoum, A.U.**
      Route choice modeling and runtime optimization for simulation of building evacuation, PhD Thesis, University of Wuppertal, 2012

[4]   **Chraibi, M.**
      Validated force-based modeling of pedestrian dynamics. PhD thesis, Universität zu Köln, 2012

[5]   http://www.fz-juelich.de/jsc/petrack

[6]   **Holl, S., Seyfried, A.**
      Hermes – an evacuation assistant for mass events, inSiDe 2009, 7(1), 60–61

[7]   **Fiedrich, F.**
      Wie im "alten Rom": Sicherheit und Planbarkeit von Großveranstaltungen, Forschungsmagazin Research Bulletin der Bergischen Universität Wuppertal, Nr. 7/ Sommersemester 2012,
      http://www.buw-output.uni-wuppertal.de/p_pics/Output7_web.pdf

Figure 3: Left: mixed and dense traffic involving pedestrians, vehicles and cyclists in the city of Kolkata (India). Right: experiment with bicyclists performed in Wuppertal (Germany). The extraction and analysis of trajectories will give new insights into the development of bicycle traffic models.

community in the form of open source projects combined with databases of experimental results.

## Fire Safety
Modern and complex buildings, like airports and multi-level underground stations, represent novel architectural ideas and pose challenges for fire protection concepts. In general, official fire safety regulations are rigid and hardly applicable to such constructions. Computational fluid dynamics simulations of fire and smoke provide a flexible – and meanwhile also a reliable – tool for the outline of fire protection systems and strategies. Although fires in simple and small compartments and buildings can be reasonably computed, many models involved are based on crude simplifications. As the capability

Current engineering applications in our group are subway vehicles and underground stations. Here we run simulations to study the heat release rates of burning subway cars and the smoke spread in stations. These insights will be used to evaluate new ventilation setups and adaptive operation modes, including modern devices as jet fans and high pressure mist sprinkler. To ease the evacuation of passengers and the access of rescue teams, the creation of zones with a high visibility is of particular interest.

## Traffic
The third area of interest is the modelling of intermodal transport, such as large areas evacuations and the modelling of mixed traffic for the planning of inner city roads. Bicycles and electric

• Armel Ulrich
  Kemloh Wagoum
• Lukas Arnold
• Stefan Holl
• Armin Seyfried

Jülich
Supercomputing
Centre (JSC)

# Greengineering the Future

The High Performance Computing Center Stuttgart (HLRS) at the University of Stuttgart, Germany hosted the annual Cray User Group Conference – CUG 2012.

CUG is an independent organization of users and owners of Cray systems that have come together to support and discuss the usage of Cray systems worldwide. Its annual meeting brings together centers and users on the one hand and Cray staff on the other hand. The meeting serves to exchange experience and discuss issues related to High Performance Computing on Cray systems.

The conference was held at the center of Stuttgart at the Maritim Hotel from April 29 - May 3, 2012. About 150 participants from Europe, Asia and the US enjoyed an interesting technical program.

Keynotes were presented by Wolfgang Nagel from the Supercomputing Center Dresden, who spoke about performance analysis and tools, Richard Kenway from the University of Edinburgh who presented the current status of the European PRACE initiative and Michael Resch from HLRS who discussed the future of High Performance Computing in a braoder sense.

Peter Ungaro – Chief Executive Officer of Cray – gave a tour de horizon of Cray and was open to the public questions of customers for one hour. Ungaro also laid out the future strategy of Cray and emphasized the importance of hardware and software integration in a world that converges technically.

The conference theme "Greengineering the Future" was reflected in a lot of discussions about usage of Cray systems to improve technology or simulate climate to better understand the environment in which we live. The motto reflects the fact that Stuttgart is the place where most engineering innovations worldwide have been created in the last 100 years. Meeting the challenge of global climate change is one of the most important issues for the city of cars which is home to companies like Daimler Porsche or Bosch. Turning the world into a greener planet, while pre-

# Cray User Group Meeting 2012

serving the existing standard of living, is one the key challenges of the future in which simulation on supercomputers will play a major role.

It is obvious that High Performance Computing will be necessary to move technology in an age of growing environmental problems. Simulations on such systems are the basis for our understanding of environmental processes. The help us to better grasp the complexity of our climate. On the other hand simulation has become the major driving key for technology imporvements that help to reduce fuel consumption, energy consumption or the pollution of air and water.

Over the last years simulation has also become a driving factor for sustainable energy supply. Wind, water and solar energy allhave to be substantially imporved to contribute to the energy suplly of an industrial nation. Questions of improved power production, reduced prudction and maintenance costs as well as the implementatin of new technologies all rely on simulation. CUG 2012 was a place to get in touch with the people who make all this possible.

CUG 2012 was an excellent forum for Cray users and developers to exchange their expertise, problems, and solutions. Participants had the opportunity to share ideas and experiences, and meet key members of the Cray team.

The open atmosphere fostered intensive discussion among users, providers and Cray staff. At the same time a arge group of students was able to get in touch with professionals in the field and see what their future work place might look like.

Social events allowed communicating in a relaxed atmosphere. A visit to the Mercedes Museum gave participants a view of 125 years of automobile development showing how simulation can become part of an outstanding history of development in individual mobility. A night out at the Stuttgart spring festival and a candle light dinner at the splendid Solitude castle close to Stuttgart made the conference a lasting impression for many.

- Neriman Emre
- Michael Resch

University of Stuttgart, HLRS

# 47th International IDC HPC User Forum

The 47th IDC HPC User Forum was held July 9-10, 2012 at the High Performance Computing Center Stuttgart (HLRS). After a series of autumn events held at HLRS the meeting was held in July in conjunction with a similar meeting at London. Even though ISC and many other meetings were close to the date about 40 international participants joined and discussed over two days the current problems of High Performance Computing. For the sixth time in a row HLRS hosted an IDC user forum at Stuttgart. The collaboration of HLRS and IDC reflects the common interest in supercomputing applications with a link to the industrial usage of High Performance Computing systems.

The program was focused on: Technologies for advancing supercomputing and supercomputer technology evolution: processors, algorithms and new technologies. After a warm welcome by HPC User Forum Chair Steve Finn and Deputy Director Stefan Wesner from HLRS, Earl Joseph and Steve Conway from IDC gave interesting insights into the market situation of HPC. Covering a field from low-end to high-end systems, the market shows extremely positive signs. A big share in the positive numbers is coming from public funding. Projects like PRACE show a positive impact on the market.

Fausto Giunchiglia from the University of Trento gave a fascinating insight into the FuturICT project highlighting the general potential of ICT to improve the quality of life for everyone. Vladimir Voevodin from the Moscow State University gave an update about the Russian situation in supercomputing and about the well-known Lomonosov system. Both talks showed that there are still many fields and countries in which supercomputing can grow and contribute. Stefan Wesner then discussed the potential of Cloud concepts for supercomputing. Presenting the EC-funded Bonfire project he showed how Clouds can help to improve access to large scale systems.

Large scale data was one of the big issues. Jim Kasdorf presented a machine with big memory. Showing the data concept of Pittsburgh Supercomputing Center (PSC) Jim worked out how large scale shared memory can help to solve large scale problems. Edgar Gabriel from the University of Houston showed how supercomputing can handle the problem of fast I/O that comes with the increased size of data sets to be handled and created by supercomputing. Paul Muzio gave a fascinating talk about the application of these technologies in the field of urban analytics. The session made clear that large scale data is already today an important issue in supercomputing and will play an ever larger role in the future.

A number of further talks were addressing issues mainly in industrial usage of supercomputers, new challenges in automotive simulation and supercomputing hardware. New products as well as new methods were discussed. Michael Resch gave an overview of supercomputing in Germany and pointed at the internationally renowned strategy of joining forces at the national level in order to improve the international position of German supercomputing.

• Michael Resch

University of
Stuttgart, HLRS

# Minister Dr. Theresia Bauer to acknowledge Prof. Michael M. Resch as "Visionary Thinker"

# 6th International Parallel Tools Workshop

On Thursday, May 31, HLRS was honored with the visit of Dr. Theresia Bauer, Minister of Science, Research and the Arts in Baden-Württemberg. The reason for Minister Bauer's visit was to present Prof. Michael M. Resch, Director of HLRS, with the "Visionary Thinker" award of the State of Baden-Württemberg. Prof. Resch's project "High Performance Computing Addressing Key Problems of Our Time" had been chosen earlier among more than 700 visionary ideas submitted in a state wide competition, called "Die Übermorgenmacher". The competition was initiated in celebration of the 60th anniversary of the State of Baden-Württemberg. Prof. Resch was nominated by external experts for this position and was one of only six scientists accepted.

In recompense for being chosen as "Visionary Thinker", each awardee was given the opportunity to publicly introduce his project and was granted one wish to be fulfilled by the State of Baden-Württemberg. Guidelines for this one wish were: a) must be reasonable, b) must not exceed 1,000 in expenses, and c) must be related to the "Visionary Thinker Project". Prof. Resch's request was a comparatively easy one to meet: He asked for Minister Bauer to take ample time and meet with and talk to HLRS employees in order to learn what their (scientific) work is all about, comprehend their general ideas and concerns, and hear out their expectations from the government.

Minister Bauer complied with Prof. Resch's wish and spent a good two hours at HLRS. When she eventually left the premises, Minister Bauer had been given a comprehensive overview of HLRS and had engaged herself in meeting and interacting with the people who are behind the scientific projectworks and all the "visionary ideas" flourishing at HLRS.

• Regina Weigand

  Gauss Centre for
  Supercomputing
  (GCS)



Handing over of the "Visionary Thinker Award". Minister T. Bauer and M. Resch.



The International Parallel Tools Workshop, jointly organized by the Center for Information Services and High Performance Computing, University of Dresden (ZIH-TUD), and the High Performance Computing Center Stuttgart (HLRS), is an annual forum for providers of tools for debugging and performance tuning of parallel applications [1]. The workshop serves as discussion forum between the developers and the users of state-of-the-art tools used in High Performance Computing, and to give feedback on mutual requirements and expectations. The last instance of this series was hosted by HLRS from September, 25 - 26 in Stuttgart [2].

Recent advances in the High Performance Computing hardware, such as increased capabilities of a single NUMA node or heterogeneous architectures combining traditional CPU nodes with the GPU accelerators, have fostered the creation of a new generation of highly scalable parallel applications and libraries. However, new capabilities of modern supercomputers have caused an increasing complexity of parallel applications development. Despite numerous efforts to improve and simplify the development process, there is still a lot of manual tuning needed to fully benefit from modern HPC architectures. The tools facilitating the debugging and performance analysis of parallel applications, either OpenMP and MPI, or more modern programming models, such as StarSs, are still of the

highest importance for the development and porting of applications for new HPC architectures.

A dozen technical papers were presented on the performance and debugging tools Threadspotter, Vampir, Scalasca Score-P, MUST, Temanejo, and MemPin, as well as on integration and interoperability of tools, and on visualization of large amounts of performance data. In addition to the usual presentations, this year's workshop featured two half-day tutorials on Likwid and Allinea DDT, respectively, addressing the training of advanced users.

The Allinea DDT tutorial focused on one side on the debugging at scale for large installations as the Cray XE6 Hermit at HLRS. A second focus was the newly introduced support for the OpenACC [3] programming model which targets accelerators. The Likwid tutorial particularly focused of the perculiarities of the Intel Sandybridge processor family which is expected to form the basis of many future HPC systems.

• José Gracia

  University of
  Stuttgart, HLRS



[1]  http://toolsworkshop.hlrs.de/2012/index.
     php/previous-workshops
[2]  http://toolsworkshop.hlrs.de/2012/
[3]  http://www.openacc.org/

# 1ˢᵗ CHANGES Workshop

New and highly scalable algorithms, performance tools, and performance modelling have been the topics of the first CHANGES workshop on High Performance Computing, which took place from 3 to 5 September 2012 at JSC. The workshop was initiated by the Computer Network Information Center (CNIC) of the Chinese Academy of Sciences (CAS), the National Center for Supercomputing Applications (NCSA) at University of Illinois, Urbana-Champaign (UIUC) and the Jülich Supercomputing Centre at Forschungszentrum Jülich. These global players in supercomputing together with their mother organizations joined forces and reached an agreement about co-sponsoring a CHinese-AmericaN-German E-Science and cyberinfrastructure workshop series (CHANGES), where the three sites will be local organizers alternately every year. In the past, the three institutions already had a couple of bilateral cooperations and workshops, like the US-American-Chinese series ACCESS and the German-Chinese Workshop on

Supercomputing (CHSC11). CHANGES, however, is the first trilateral cooperation and will not only consider the issues of the partner institutions but will also take into account national topics.

The first CHANGES workshop was a high-level platform dealing with latest trends in supercomputing, sophisticated information techniques and interdisciplinary applications. About 45 well-known experts, among them 21 speakers, came together by invitation and discussed latest results of their research fields. The workshop focussed on the above mentioned topics but covered also several fields in supercomputing, e-Science and its applications in China, Germany and the United States. Besides the presentations the workshop provided a forum for trilateral cooperations on student exchange and mutual research projects. First promising cooperation ideas were developed and are foreseen to be launched in the near future.

- Norbert Attig

  Jülich
  Supercomputing
  Centre (JSC)

# Workshop on Hybrid Particle-Continuum Methods

A main purpose of computational materials physics is to establish a fundamental link between atomic-scale processes and the macroscopic behavior of condensed matter, including composite materials, complex fluids, and materials of technological interest. A common characteristic of these systems is the existence of important features at multiple time or length scales. Typical examples are e.g. crack propagation in solids or protein folding in solution. A precise description of interatomic interactions is needed at the crack tip or when two side chains of the protein come close to each other. However, stress boundary conditions or local electrostatic fields in these examples are strongly affected by long-range interactions. The latter are conveyed by a medium that can be described in terms of continuous fields. Modeling such systems is challenging because the small and the large scale have to be incorporated simultaneously and their underlying constitutive equations differ. This is why sequential multi-scale modeling is not an option, despite its success in the bottom-up development of field theoretical approaches to homogeneous media.

There has been much progress on coupling different descriptions and levels of resolution in the communities interested in complex fluids and complex solids. However, there has been surprisingly little exchange between them. In the tradition of previous workshops organized by research groups of the John von Neumann Institute for Computing (NIC) at Jülich fostering the exchange between various scientific communities,
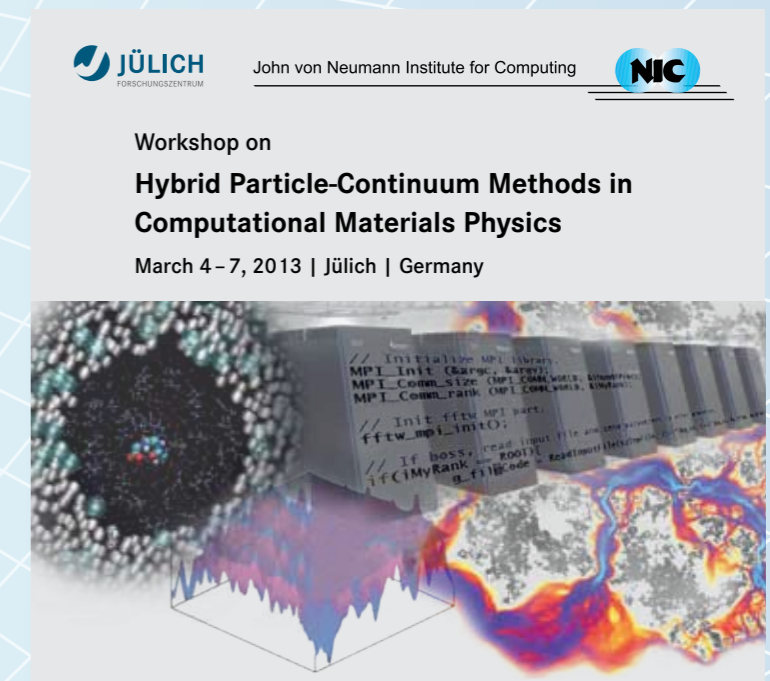
a workshop on "Hybrid Particle-Continuum Methods in Computational Materials Physics" will be held at the Forschungszentrum Jülich from 4 to 7 March 2013. In order to narrow down the focus of the workshop, we will concentrate on methods that couple particles, or discrete degrees of freedom, to continuum fields. The two main topics are:
- Continuum-mediated interactions between particles
- Adaptive and non-adaptive coupling between particle-based and continuum-based descriptions of materials.

So far, 16 internationally recognized scientists have accepted our invitation to present their research. We will also allow for a limited number of contributed talks and give everybody attending the workshop the possibility to present a poster. More details can be found at http://www.fz-juelich.de/jsc/HY-BRID2013

- Martin Müser
- Godehard Sutmann
- Roland G. Winkler

  Forschungszentrum
  Jülich

Workshop on
**Hybrid Particle-Continuum Methods in Computational Materials Physics**

March 4–7, 2013 | Jülich | Germany

# JSC Guest Student Programme on Scientific Computing 2012

As one of the leading HPC centres in Europe, Jülich Supercomputing Centre provides supercomputer resources, IT tools and HPC expertise for computational scientists at German and European universities, research institutions, and in industry. An important part of this mission is to introduce young academics to HPC and its role in scientific research. To this end, JSC hosts regular support activities and educational programmes in the field of Scientific Computing.

The JSC Guest Student Programme on Scientific Computing has already been successfully running for 13 years. Since the first programme in 2000, a total of 133 students have had the opportunity to join research teams from JSC and other institutes at Forschungszentrum Jülich for ten weeks. Working on challenging topical scientific projects, they received training with up-to-date hardware and software as well as HPC-related methods and algorithms. For many students, the programme has been the foundation for a career in HPC and the basis for fruitful continuing cooperations with their advisers, occasionally leading to widely recognized scientific publications, e.g. [1] to give a recent example. Some students also returned to JSC as Master or PhD candidates.

The JSC Guest Student Programme 2012 took place from August 6 -



October 12. Once again it was run under the CECAM framework (Centre Européen de Calcul Atomique et Moléculaire) and organized in cooperation with the German Research School for Simulation Sciences (GRS). Support by IBM Germany through a sponsorship within the IBM University Relations programme is also gratefully acknowledged.

This year's announcement yielded a record response of more than 30 applications from 13 countries, comprising applicants from a particularly wide range of disciplines, including: mathematics, physics, chemistry, computer science and engineering, biomedicine and earth sciences. This meant that competition for the available places was especially strong, and after the final selection process, 13 students were invited to Jülich.

The programme started with ten days of introductory courses on parallel programming, including use of MPI on distributed-memory cluster systems and OpenMP for shared-memory parallelization, as well as CUDA and OpenCL for GPU-accelerated machines. For the remainder of their stay, the participants worked on research projects in a broad area of fields, ranging from applications in atmospheric sciences, fluid and molecular dynamics, particle-in-cell methods, and safety research, to fundamental research in elementary particle physics and mathematical algorithms. Besides using the multi-purpose cluster JUROPA and the GPU system JUDGE, one of the key topics

was the utilization of the newly installed IBM Blue Gene/Q system JUQUEEN in Jülich.

During a concluding colloquium, the participants had the opportunity to present and discuss their work with other students and supervisors. Finally, more detailed reports have been prepared and were compiled into a printed JSC publication, which will be made available online.

Next year's JSC Guest Student Programme will start on August 5, 2013. It will be officially announced in January 2013 and is open to students from the natural sciences, engineering, computer science and mathematics after completing their Bachelor and before reaching their Master's degree. The application deadline has been fixed to April 26, 2013. Additional information as well as the reports of previous years can be found online on www.fz-juelich.de/jsc/gsp
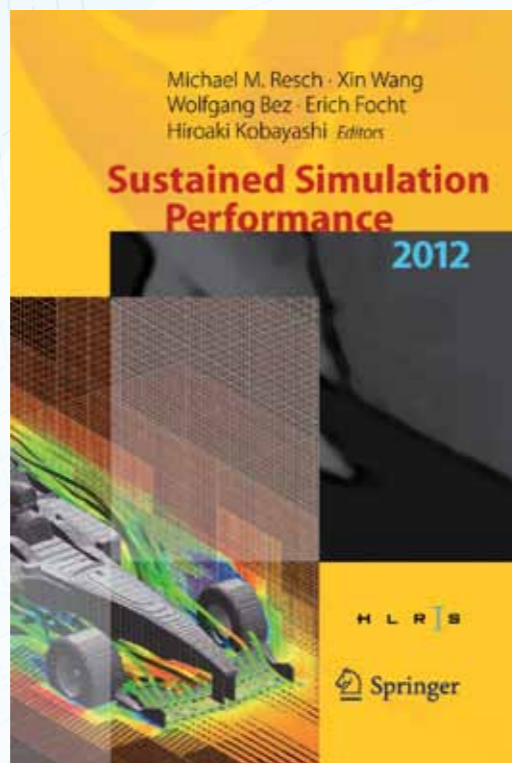
- Mathias Winkel

Jülich Supercomputing Centre (JSC)

## References

[1] **Dapp, W., Lücke, A., Persson, B., Müser, H.** Self-Affine Elastic Contacts: Percolation and Leakage, Phys. Rev. Lett. 108, p. 244301, 2012

# New Books in HPC

## Sustained Simulation Performance 2012

Resch, M., Wang, X., Bez, W., Focht, E., Kobayashi, H. (Eds.)
2013, XII, 182 p. 81 illus., 67 in color.
Proceedings of the joint Workshop on High Performance Computing on Vector Systems, Stuttgart (HLRS), and Workshop on Sustained Simulation Performance, Tohoku University, 2012.

The book presents the state-of-the-art in High Performance Computing and simulation on modern supercomputer architectures. It covers trends in hardware and software development in general and specifically the future of high performance systems and heterogeneous architectures. The application contributions cover Computational Fluid Dynamics, material science, medical applications and climate research. Innovative fields like coupled multi-physics

or multi-scale simulations are presented. All papers were chosen from presentations given at the 14th Teraflop Workshop held in December 2011 at HLRS, University of Stuttgart, Germany and the Workshop on Sustained Simulation Performance at Tohoku University in March 2012.

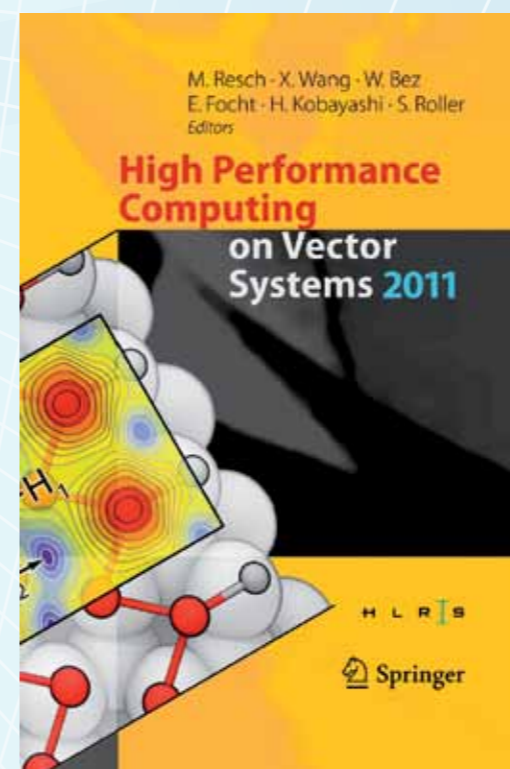## Tools for High Performance Computing 2011
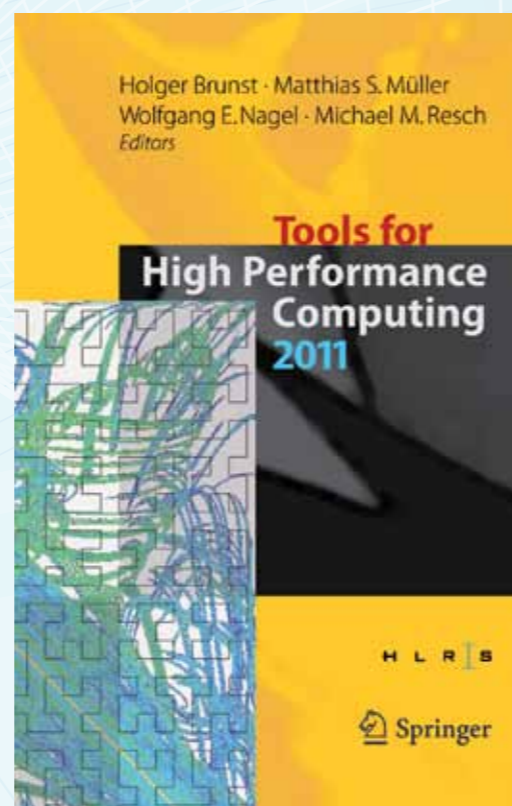
Brunst, H., Müller, M., Nagel, W.E., Resch, M. (Eds.)
2012, XVII, 154 p. 116 illus., 44 in color.
Proceedings of the 5th International Workshop on Parallel Tools for High Performance Computing, September 2011, ZIH, Dresden.

The proceedings of the 5th International Workshop on Parallel Tools for High Performance Computing provide an overview on supportive software tools and environments in the fields of System Management, Parallel Debugging and Performance Analysis. In the pursuit to maintain exponential growth for the performance of high performance computers the HPC community is currently targeting Exascale Systems. The initial planning for Exascale already started when the first Petaflop system was delivered. Many challenges need to be addressed to reach the necessary performance. Scalability, energy efficiency and fault-tolerance need to be increased by orders of magnitude. The goal can only be achieved when advanced hardware is combined with a suitable software stack. In fact, the importance of software is rapidly growing.

## High Performance Computing on Vector Systems 2011

Resch, M., Wang, X., Focht, E., Kobayashi, H., Roller, S. (Eds.)
2012, VIII, 184 p. 89 illus., 68 in color.

The book presents the state-of-the-art in High Performance Computing and simulation on modern supercomputer architectures. It covers trends in hardware and software development in general and specifically the future of vector-based systems and heterogeneous architectures.

The application contributions cover Computational Fluid Dynamics, material science, medical applications and climate research. Innovative fields like coupled multi-physics or multi-scale simulations are presented.

All papers were chosen from presentations given at the 13th Teraflop Workshop held in October 2010 at Tohoku University, Japan. For 8 yers now this workshop brings together experts in high performance computing that put the emphasis of their work on solutions for real world problems focussing on the sustained performance achievable on modern hardware architectures.

With the advent of Petaflops systems and the predictable arrival of Exascale systems in the future, sustained performance is getting more and more important. this book serves as a guide through the jungles of architectures and simulation software systems.

# HLRS Scientific Tutorials and Workshop Report and Outlook

HLRS has installed **Hermit**, a Cray XE6 system with AMD Interlagos processors and 1 PFlop/s peak performance. We strongly encourage you 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 to this 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 4-day courses in cooperation with Cray and multi-core optimization specialists are April 23 - 26, 2013 and also in autumn 2013.

Programming of Cray XK6 clusters with GPUs is taught in **OpenACC Programming for Parallel Accelerated Supercomputers** on April 29 - 30, 2013.

These Cray XE6 and XK6 courses are also presented to the European community in the framework of the **PRACE Advanced Training Centre** (PATC). GCS, i.e., HLRS, LRZ and the Jülich Supercomputer Centre together, serve as one of the first six PATCs in Europe.

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 will be presented twice, in March 2013 at HLRS in Stuttgart and September 2013 at LRZ.

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 moving to Siegen. It is again scheduled in February 2013 in Stuttgart. 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 modeling 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.

Our general course on parallelization, the **Parallel Programming Workshop**, September 2 - 6, 2013 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 March 2013, 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 May 2013, the participants will get an introduction of partitioned global address space (PGAS) languages.

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 with **tutorials** at conferences; see the box on the left page.

In the table, you can find the whole HLRS series of training courses in 2013. 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).

• Rolf Rabenseifner

University of Stuttgart, HLRS

**2013 – Workshop Announcements**

**Scientific Conferences and Workshops at HLRS**

11th HLRS/hww Workshop on Scalable Global Parallel File Systems (March/April 2013)

7th ZIH+HLRS Parallel Tools Workshop (date and location not yet fixed)

High Performance Computing in Science and Engineering - The 16th Results and Review Workshop of the HPC Center Stuttgart (October 2013)

IDC International HPC User Forum (October 2013)

**Parallel Programming Workshops: Training in Parallel Programming and CFD**

Parallel Programming and Parallel Tools (TU Dresden, ZIH, February 4-7)

Introduction to Computational Fluid Dynamics (HLRS, February 11-15)

Iterative Linear Solvers and Parallelization (HLRS, March 4-8)

Cray XE6 Optimization Workshops (HLRS, April 23-26)

Unified Parallel C (UPC) and Co-Array Fortran (CAF) (HLRS, May 2-3)

Parallel Programming with MPI & OpenMP (TU Hamburg-Harburg, July 29-31)

Iterative Linear Solvers and Parallelization (LRZ, Garching, September 2-4)

Message Passing Interface (MPI) for Beginners (HLRS, September 23-24)

Shared Memory Parallelization with OpenMP (HLRS, September 25)

Advanced Topics in Parallel Programming (HLRS, September 26-27)

Parallel Programming with MPI & OpenMP (FZ Jülich, JSC, November 25-27)

**Training in Programming Languages at HLRS**

Fortran for Scientific Computing (March 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/2013/parallel_prog_2013/

# GCS – High Performance Computing    Courses and Tutorials

## Introduction to the Programming and Usage of the Supercomputer Resources at Jülich

### Date and Location
November 22 - 23, 2012
JSC, Research Centre Jülich

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

### Web page
http://www.fz-juelich.de/ias/jsc/events/sc-nov

## Parallel Programming with MPI, OpenMP and PETSc

### Date and Location
November 26 - 28, 2012
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. Course language is English. This course is organized by JSC in collaboration with HLRS. Presented by Dr. Rolf Rabenseifner, HLRS.

### Web page
http://www.fz-juelich.de/ias/jsc/events/mpi

## Node-Level Performance Engineering

### Date and Location
December 06, 2012
December 07, 2012
LRZ Building,
University Campus Garching,
near Munich, Boltzmannstr. 1.

### Contents
This course teaches performance engineering approaches on the compute node level. "Performance engineering" as we define it is more than employing tools to identify hotspots and bottlenecks. It is about developing a thorough understanding of the interactions between software and hardware. This process must start at the core, socket, and node level, where the code gets executed that does the actual computational work. Once the architectural requirements of a code are understood and

correlated with performance measurements, the potential benefit of optimizations can often be predicted. We introduce a "holistic" node-level performance engineering strategy, apply it to different algorithms from computational science, and also show how an awareness of the performance features of an application may lead to notable reductions in power consumption:

- Introduction
- Practical performance analysis
- Microbenchmarks and the memory hierarchy
- Typical node-level software over heads
- Example problems:
  - The 3D Jacobi solver
  - The Lattice-Boltzmann Method
  - Sparse Matrix-Vector Multiplication
  - Backprojection algorithm for CT reconstruction
- Energy & Parallel Scalability.

Between each module, there is time for Questions and Answers!

### Web page
http://www.lrz.de/services/compute/courses

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

### Dates and Locations
December 13 - 14, 2012
May 02 - 03, 2013
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.

### Web page
www.hlrs.de/events/

## Parallel Programming with MPI, OpenMP, and Tools

### Date and Location
February 04 - 07, 2013
Dresden, ZIH

### 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. The last day is dedicated to tools for debugging and performance analysis of parallel applications. This course is organized by ZIH in collaboration with HLRS.

### Web page
www.hlrs.de/events/

## Programming with Fortran

### Date and Location
February 04 - 08, 2013
LRZ Building,
University Campus Garching,
near Munich, Boltzmannstr. 1.

### Contents
This course is targeted at scientists with little or no knowledge of the Fortran programming language, but need it for participation in projects using a Fortran code base, for development of their own codes, and for getting acquainted with additional tools like debugger and syntax checker as well as handling of compilers and libraries. The language is for the most part treated at the level of the Fortran 95 standard; features from Fortran 2003 are limited to improvements on the elementary

level. Advanced Fortran features like object-oriented programming or co-arrays will be covered in a follow-on course in autumn.
To consolidate the lecture material, each day's approximately 4 hours of lecture are complemented by 3 hours of hands-on sessions.

### Prerequisites
Course participants should have basic UNIX/Linux knowledge (login with secure shell, shell commands, basic programming, vi or emacs editors).

### Web page
http://www.lrz.de/services/compute/courses

## Introduction to Computational Fluids Dynamics

### Date and Location
February 11 - 15, 2013
Stuttgart, HLRS

### 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 HLRS, IAG, and University of Kassel.

**Web page**
www.hlrs.de/events/

## Iterative Linear Solvers and Parallelization

**Date and Location**
March 04 - 08, 2013
Stuttgart, HLRS

September 02 - 06, 2013
Garching, LRZ

**Contents**
The focus is on iterative and parallel solvers, the parallel programming models MPI and OpenMP, and the parallel middleware PETSc. Thereby, 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 University of Kassel, HLRS, and IAG.

**Web page**
www.hlrs.de/events/

## Parallel Programming of High Performance Systems

**Date and Location**
March 04 - 10, 2013
LRZ Building,
University Campus Garching,
near Munich.
RRZE building,
University Campus Erlangen,
Martensstr. 1:
Via video conference if there is sufficient interest.

**Contents**
This course, a collaboration of Erlangen Regional Computing Centre (RRZE) and LRZ, is targeted at students and scientists with interest in programming modern HPC hardware, specifically the large scale parallel computing systems available in Jülich, Stuttgart and Munich.

Each day is comprised of approximately 4 hours of lectures and 3 hours of hands-on sessions:

• Introduction to High Performance Computing

• Technical aspects of software engineering
• Basic parallel programming with MPI and OpenMP
• Processor Architectures: Register, Cache, Locality, Performance metrics
• Principles of code optimization
• Advanced MPI and OpenMP programming
• Parallel algorithms
• Advanced OpenMP programming
• Parallel Architectures: multi-core, multi-socket, ccNUMA
• Architecture-specific optimization strategiesPrerequisites:
  Good working knowledge of at least one of the standard HPC languages Fortran 95, C or C++.

**Web page**
http://www.lrz.de/services/compute/courses

## Fortran for Scientific Computing

**Date and Location**
March 11 - 15, 2013
Stuttgart, HLRS

**Contents**
This course is dedicated for scientists and students to learn (sequential) programming scientific applications with Fortran. The course teaches newest Fortran standards. Hands-on

sessions will allow users to immediately test and understand the language constructs.

**Web page**
www.hlrs.de/events/

## Advanced Topics in High Performance Computing

**Date and Location**
March 18 - 21, 2013
LRZ Building,
University Campus Garching,
near Munich

**Contents**
In this add-on course to the parallel programming course special topics are treated in more depth, in particular performance analysis, I/O and PGAS concepts. It is provided in collaboration of Erlangen Regional Computing Centre (RRZE) and LRZ. Each day is comprised of approximately 5 hours of lectures and 2 hours of hands-on sessions:

• Intel Tracing Tools
• Introduction to Scalasca
• Parallel input/output with MPI-IO
• I/O tuning on high performance file systems
• Introduction to the PGAS languages Coarray Fortran and UPC

**Prerequisites**
Good MPI and OpenMP knowledge as presented in the course "Parallel programming of High Performance Systems" (see above).

**Web page**
http://www.lrz.de/services/compute/courses

## Parallel I/O and Portable Data Formats

**Date and Location**
March 18 - 20, 2013
JSC, Research Centre Jülich

**Contents**
This course will introduce MPI parallel I/O and portable, self-describing data formats, such as HDF5 and NetCDF. Participants should have experience in parallel programming in general, and either C/C++ or Fortran in particular.
This course is a PATC course (PRACE Advanced Training Centres).

**Web page**
http://www.fz-juelich.de/ias/jsc/events/parallelio

## Eclipse: C/C++/Fortran Programming

**Date and Location**
March 28, 2013
LRZ Building,
University Campus Garching,
near Munich

**Contents**
This course is targeted at scientists who wish to be introduced to programming C/C++/Fortran with the Eclipse C/C++ Development Tools (CDT), or the Photran Plugin. Topics covered include:

• Introduction to Eclipse IDE
• Introduction to CDT
• Hands-on with CDT
• Short introduction and demo of Photran

**Prerequisites**
Course participants should have basic knowledge of the C and/or C++/Fortran programming languages.

**Web page**
http://www.lrz.de/services/compute/courses

## GPU Programming

**Date and Location**
April 15 - 17, 2013
JSC, Research Centre Jülich

# GCS – High Performance Computing     Courses  and Tutorials

**Contents**

Many-core programming is a very dynamic research area. Many scientific applications have been ported to GPU architectures during the past four years. We will give an introduction to CUDA, OpenCL, and multi-GPU programming using examples of increasing complexity. After introducing the basics the focus will be on optimization and tuning of scientific applications.

This course is a PATC course (PRACE Advanced Training Centres).

**Web page**

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

## Advanced GPU Programming

**Date and Location**

April 18 - 19, 2013
JSC, Research Centre Jülich

**Contents**

Today's computers are commonly equipped with multicore processors and graphics processing units. To make efficient use of these massively parallel compute resources advanced knowledge of architecture and programming models is indispensable. This course builds on the introduction to GPU programming. It focuses on finding and eliminating bottlenecks using profiling and advanced pro-

gramming techniques, optimal usage of CPUs and GPUs on a single node, and multi-GPU programming across multiple nodes. The material will be presented in the form of short lectures followed by in-depth hands-on sessions.

**Web page**

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

## Cray XE6 Optimization Workshop

**Date and Location**

April 23 - 26, 2013
Stuttgart, HLRS

**Contents**

HLRS installed Hermit, a Cray XE6 system with AMD Interlagos processors and a performance of 1 PFlop/s. We strongly encourage you 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 this course, we will give all necessary information to move applications from the current NEC SX-9, the Nehalem cluster, or other systems to Hermit. Hermit provides 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. From Monday to Wednesday, specialists from Cray will support you in your effort porting and optimizing your application on our Cray XE6. On Thursday, Georg Hager and Jan Treibig from RRZE will present detailed information on optimizing codes on the multicore AMD Interlagos processor. Course language is English (if required).

**Web page**

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

## Cray XK at HLRS, Stuttgart

**Date and Location**

April 29 - 30, 2013
Stuttgart, HLRS

**Contents**

This workshop will cover the programming environment of the Cray XK6 hybrid supercomputer, which combines multicore CPUs with GPU accelerators. Attendees will learn about the directive-based OpenACC programming model whose multi-vendor support allows users to portably develop applications for parallel accelerated supercomputers.

The workshop will also demonstrate how to use the Cray Programming Environment tools to identify CPU application bottlenecks, facilitate the OpenACC porting, provide accelerat-

ed performance feedback and to tune the ported applications. The Cray scientific libraries for accelerators will be presented, and interoperability of OpenACC directives with these and with CUDA will be demonstrated. Through application case studies and tutorials, users will gain direct experience of using OpenACC directives in realistic applications.

Users may also bring their own codes to discuss with Cray specialists or begin porting.

**Web page**

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

## Introduction to the Programming and Usage of the Supercomputer Resources at Jülich

**Date and Location**

May 16 - 17, 2013
JSC, Research Centre Jülich

**Contents**

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

**Web page**

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

## Guest Student Programme: Education in Scientific Computing

**Date and Location**

August 05 - October 11, 2013
JSC, Research Centre Jülich

**Contents**

Guest Student Programme "Scientific Computing" to support education and training in the fields of supercomputing. Application deadline is April 30, 2013. Students of Computational Sciences, Computer Science and Mathematics can work ten weeks in close collaboration with a local scientific host on a subject in their field.

**Web page**

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

# Authors

**Nikolaus A. Adams**

Nikolaus.Adams@tum.de

**Momme Allalen**

momme.allalen@lrz.de

**Padmesh Anjukandi**

padmesh.anjukandi@theochem.rub.de

**Christoph Anthes**

christoph.anthes@lrz.de

**Lukas Arnold**

l.arnold@fz-juelich.de

**Norbert Attig**

n.attig@fz-juelich.de

**Thomas Baumann**

baumann@hlrs.de

**Thomas Bönisch**

boenisch@hlrs.de

**Matthias Brehm**

matthias.brehm@lrz.de

**Steffen Brinkmann**

brinkmann@hlrs.de

**Max Camenzind**

M.Camenzind@lsw.uni-heidelberg.de

**Przemyslaw Dopieralski**

przemyslaw.dopieralski
@theochem.rub.de

**Neriman Emre**

emre@hlrs.de

**Dietmar Erwin**

d.erwin@fz-juelich.de

**Shiqing Fan**

fan@hlrs.de

**Dmitry A. Fedosov**

d.fedosov@fz-juelich.de

**Michael Gerndt**

gerndt@in.tum.de

**Colin W. Glass**

glass@hlrs.de

**Gerhard Gompper**

g.gompper@fz-juelich.de

**José Gracia**

gracia@hlrs.de

**Carla Guillen**

carla.guillen@lrz.de

**Wolfram Hesse**

Wolfram.Hesse@lrz.de

**Stefan Hickel**

sh@tum.de

**Christian Hoelbling**

hch@physik.uni-wuppertal.de

**Stefan Holl**

st.holl@fz-juelich.de

**Susanne Horn**

susanne.horn@dlr.de

**Ferdinand Jamitzky**

ferdinand.jamitzky@lrz.de

**Matthias Kaczorowski**

matthias.kaczorowski@dlr.de

**Armel Ulrich**

**Kemloh Wagoum**

u.kemloh@fz-juelich.de

**Dieter Kranzlmüller**

kranzlmueller@ifi.lmu.de

**Jenny Kremser**

jenny.kremser@asc-s.de

**Andreas Lintermann**

A.Lintermann@aia.rwth-aachen.de

**Dominik Marx**

dominik.marx@theochem.rub.de

**Michael Meyer**

michael.meyer@aer.mw.tum.de

**Bernd Mohr**

b.mohr@fz-juelich.de

**Martin Müser**

m.mueser@fz-juelich.de

**Mathias Nachtmann**

nachtmann@hlrs.de

**Walter Nadler**

w.nadler@fz-juelich.de

**Carmen Navarrete**

carmen.navarrete@lrz.de

**Marc Offman**

Marc.Offman@lrz.de

**Ludger Palm**

Ludger.Palm@lrz.de

**Dirk Pleiter**

d.pleiter@fz-juelich.de

**Rolf Rabenseifner**

rabenseifner@hlrs.de

**Michael M. Resch**

resch@hlrs.de

**Jordi Ribas-Arino**

jordi.ribas@theochem.rub.de

**Christian Rössel**

c.roessel@fz-juelich.de

**Helmut Satzger**

Helmut.Satzger@lrz.de

**Armin Seyfried**

a.seyfried@fz-juelich.de

**Olga Shishkina**

olga.shishkina@dlr.de

**Michael Stephan**

m.stephan@fz-juelich.de

**Jamie O'Sullivan**

josulliv@lsw.uni-heidelberg.de

**Godehard Sutmann**

g.sutmann@fz-juelich.de

**Regina Weigand**

r.weigand@gauss-centre.eu

**Andreas Wierse**

wierse@sicos-bw.de

**Mathias Winkel**

m.winkel@fz-juelich.de

**Roland G. Winkler**

r.winkler@fz-juelich.de

**Felix Wolf**

f.wolf@grs-sim.de

## Erratum