inside

SPRING 2013

inside • Vol. 11 No. 1

Innovatives Supercomputing in Deutschland

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

Publishers
Prof. Dr. A. Bode | Prof. Dr. Dr. T. Lippert | Prof. Dr.-Ing. Dr. h.c. Dr. h.c. M. M. Resch

Editor
F. Rainer Klank, HLRS klank@hlrs.de

Design
Antje Häusser, HLRS haeusser@hlrs.de
Stefan Steinbach, HLRS steinbach@hlrs.de

Cover: Laminar-turbulent transition on a common dolphin at 1 m/s and 1 % turbulence intensity.
Transition location visualized by turbulent-kinetic-energy contours combined with isolines of pressure coefficient. Streamlines indicate part of oncoming flow forced towards the lower body. (simulation: D. Riedeberger, IAG; University of Stuttgart, Germany, model: V. Pavlov, ITAW, University of Hannover, Germany)

Editorial

Welcome to this new issue of InSiDE, the journal on Innovative Supercomputing in Germany published by the Gauss Centre for Supercomputing (GCS). In this issue, we look at a number of new developments, highlight applications, look at projects in Germany and Europe, and provide a wealth of information about German supercomputing systems, training and activities.

To take a first step is often relatively easy, but to keep walking can be an even bigger challenge, and is often accompanied by a series of failed attempts – every child learning to walk would agree. For GCS this is not the case. Step by step, the strategy of GCS is being implemented.

Firstly, an important step for both Germany and Europe, the Jülich Supercomputing Centre (JSC@GCS) has installed its second hardware step in the common GCS project PetaGCS. The IBM BlueGene/Q was inaugurated on February 14, 2013 in the presence of high-level guests from politics, science and industry. With a performance of nearly 6 PFLOP/s the new system contributes to the growth of computing speed and capacity for German and European scientists. With this system, GCS continues leading in Europe while also continuing its strategic collaboration with IBM as one of the technology leaders in supercomputing worldwide.

In March HLRS signed an agreement with Cray to upgrade its systems to 5 PFLOP/s in 2014 and finally, on April 10, 2013 the contract for the extension of SuperMUC in early 2015 to 6.4 PFLOP/s was signed at LRZ in presence of the Bavarian Minister Dr. Wolfgang Heubisch between IBM and LRZ.

According to its roadmap, GCS will not only install or upgrade a new system about once a year, but also will continue its infrasrtucture development path. After upgrading the computer room facilities at all GCS centers, we are now extending office spaces for a growing number of researchers. Hence, we report on the inauguration of a new research building at the High Performance Computing Center Stuttgart (HLRS@GCS). With its office space for about 60 researchers and its three-floor, five-sided visualization facilities, it is a breading ground for new ideas and methods.

Again, we present a rich section on applications using allocations on the GCS systems. Starting with the winners of the golden spike award of HLRS and ranging from computational fluid dynamics to astrophysical simulations. The contributions clearly show the need for compute power and the potential of HPC simulations for both basic and applied science. The growing strength of GCS in HPC research is highlighted by the research project reports that show the variety of fields in which the GCS centers further push the limits of supercomputing.

Finally we announce a change in leadership at GCS. Michael Resch, director of HLRS, has taken over as the chairman of the board of GCS from Heinz-Gerd Hegering in May this year. The most pressing issues for him will be the positioning of GCS in the future PRACE concept as well as the continued development of GCS towards a center of excellence in computing.

- 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

News

Applications

Projects

Systems

News

Official Inauguration of the Supercomputer JUQUEEN at JSC

On February 14, 2013, the new supercomputer JUQUEEN at the Jülich Supercomputing Centre was officially inaugurated in the presence of representatives from the German Federal Ministry of Education and Research (BMBF) and the Ministry of Innovation, Science and Research of the state of North Rhine-Westphalia (MIWF). After the last four of 28 computer racks had been installed in January, JSC was proud to present JUQUEEN in its final configuration to the public. Prof. Achim Bachem, chairman of the Board of Directors of Forschungszentrum Jülich, welcomed the guests and thanked the ministries for their financial support provided through the Gauss Centre for Supercomputing (GCS).

Prof. Wolf-Dieter Lukas (Fig. 1), head of their Directorate 5 at BMBF, delivered the official greeting on behalf of the Federal Minister. He outlined that the political decision to support supercomputing to the high amount as it had been done for the GCS supercomputers was a good one. However, the supercomputer is not the most important part, but the science that can be done with it. In this spirit, he congratulated Forschungszentrum Jülich for being part of the European FET Flagship "Human Brain Project", which will definitely add its share to the scientific simulations runnning on JUQUEEN. Helmut Dockter, State Secretary at MIWF, addressed JSC in the name of the NRW minister. He expressed his expectation to obtain simulation results relevant to the needs of society in order to facilitate future financial support.

The last speaker in the ceremony was Prof. Thomas Lippert, head of the Jülich Supercomputing Centre. He presented a short history of the rapid development in supercomputing, introduced which scientific problems of high complexity can be tackled with JUQUEEN, and gave an outlook on special simulations suitable to be run on this supercomputer. Afterwards, the guests were invited to visit JUQUEEN in the machine room (Fig. 2) and to celebrate the computer in a reception.

In the afternoon, the event continued with a scientific colloquium "Supercomputing in North Rhine-Westphalia". Four renowned top scientists from the fields of brain research, materials science, fluid dynamics and civil security research presented their simulations. Prof. Katrin Amunts from the C. and O. Vogt-Institute of Brain Research in Düsseldorf gave an overview on how to build a 3-D model of the human brain by using supercomputers. Prof. Heinz Pitsch, Institute for Combustion Technology at the RWTH Aachen University, introduced his research and simulations on turbulent reacting flows in engineering devices, such as aircraft engines and industrial combustors. In his talk "High-Performance Computing in Materials Science: Current Applications and Challenges", Prof. Alexander Hartmaier from the Interdisciplinary Centre for Advanced Materials Simulation (ICAMS), Ruhr-University Bochum, highlighted his development of methods and algorithms to upscale atomistic models. His goal is to unravel the complex interplay and competition of different fundamental mechanisms occurring during the deformation of materials.

In the final talk, Prof. Armin Seyfried from the Department of Computer Simulation for Fire Safety and Pedestrian Traffic at University of Wuppertal presented a decision support system for the evacuation of sport arenas and explained why the multi-scale character of fire simulations demands for large computing resources.

• Sabine Höfler-Thierfeldt

Jülich Supercomputing Centre

Figure 1: Prof. Wolf-Dieter Lukas from the German Federal Ministry of Education and Research gave the official greeting from the minister

Figure 2: Visiting JUQUEEN: Prof. Thomas Lippert (head of JSC), Prof. Achim Bachem (Chairman of the Board of Directors of Forschungszentrum Jülich), Prof. Wolf-Dieter Lukas (German Federal Ministry of Education and Research), Prof. Ernst Schmachtenberg (Rector of RWTH Aachen University), Helmut Dockter (Ministry of Innovation, Science and Research of North Rhine-Westphalia), and Prof. Sebastian Schmidt (member of the Board of Directors of Forschungszentrum Jülich)

Interview with Dr. Claus-Axel Müller

Dr. Müller, you have been the managing director of the Gauss Center for Supercomputing (GCS) since 2009. Why did you decide to move into HPC at that time?

I have been connected to HPC throughout my whole career. I hold a PhD in Computational Fluid Dynamics. Therefore it was not really a move into HPC. My last position in industry was Managing Director of T-Systems Solutions for Research which is specialized to IT-Services for research organisations. It was more a move from industry to public science and research. I appreciated very much when I was asked to join GCS. The challenge for me was to help the three Tier-0 HPC-centres in Germany to build up the most powerful HPC-organisation in Europe.

After 3 years as a managing director, did you see these expectations fulfilled? Has HPC developed as you expected it?

Take a look at the results. We run the most powerful machines in Europe for science and industry and support the largest and most challenging scientific projects. Furthermore, we are the largest provider of High-Performance Computing resources in PRACE. To be honest, at the beginning it was not clear for me if we could reach the goal in such a short time frame, but today all my expectations are really met.

How do you see the cooperation of three large national centres in a single organization like GCS?

You put it into the right words. The most important point for GCS is the successful cooperation of the three centres. GCS is organized as a non-profit association and operates like a management holding company responsible for the overall governance which includes strategy, funding and peer-review-based access to the HPC-resources. Each centre is a member of the association, and has the same rights. The decision process in GCS is purely democratic. On the other hand, each centre remains solely responsible for its own organization and operation. This kind of distribution of power, responsibility, and operation has been proven to be very successful in industry and public organisations. It seems to be the best approach for GCS also. If you think about the good position we reached in the last three years, you will understand that in the long term, the chosen governance will support and develop the strength of the three centres while simultaneously benefitting from the synergy effects of the joint approach.

Europe has started to play a larger role in HPC while GCS was developing its own strategy. How do you see the role of GCS in Europe?

Computer-based simulation plays a very strong role in science and engineering. Therefore HPC is a very important location factor for the national and the European economies. The German economy depends strongly on innovation in science and technology. Therefore it is very important for Germany to have always access to the leading HPC resources and knowledge. The strategy is twofold. On one hand, we have to provide the most suitable HPC services to the national scientists and researchers. On the other hand, the challenges of the upcoming requirements of HPC can only be solved on a European level. We are one of the driving forces of the PRACE initiatives and therefore the exchange of knowledge and resources is very important for us.

And what would you see as the main role that Europe can play in German HPC?

The support of the European Commission by funding the implementation of PRACE already made a positive impact on HPC in Germany. Europe has a wealth of well-trained and excellent students, researchers, and scientists who are looking for good jobs in HPC. A combination of the skills and technologies that are available at the GCS centers and the additional new ideas of this European pool of researchers is fruitful both for Europe and Germany. The fact that the European Commission considers HPC as a vital technology for the 21st century will certainly give a boost to HPC in Germany as well. To realize the potential of HPC on the European level, the commission has to fund not only the implementation of new ideas, but also provide a significant funding for the implemented infrastructure.

The world of HPC is discussing the roadmap to exaflops extensively. How realistic do you see the chances that we reach this target within the next 6 years?

It seems to be certain that there will be an Exaflop system until 2022-2026 in the world. With the knowledge we have today, and from my point of view, the first exaflop system will be very experimental and designed with specific requirements for only a few algorithms. GCS has to support a wide variety of scientific communities, and therefore could not afford to go for a single-shot solution. We will focus on exascale, and not only on exaflop, which means that we will put much more effort into the development and the support of the applications in the future beneath the continuous upgrading of the HPC systems.

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?

With exaflop systems, the amount of data produced will grow at least with the same order of magnitude as the computing power. The simple part of the equation is that we need a lot more of pre- and post-processing power, which includes storage and networking. But as you can see from the discussion about large scale data management, there are other constraints due to content and long-term availability of data. In the future we will need centres more strongly linked together based on the concept of the HPC-Pyramid which are focussed to specific scientific areas, architectures, and data-management requirements.

When we look at the national level, we find that Germany has a well-established pyramid of performance and has organized the medium sized centres in the Gauss Alliance. How do you see the

future development of HPC in Germany as a whole?

The national strategy based on the so-called "HPC-Pyramid" has proven to be very successful, because it addresses the needs and requirements of the scientists and researchers in Germany very well. You have to look at HPC not only from the perspective of the size of the machines. Put yourself in the position of a scientist or a researcher: You started learning how to use a supercomputer at a university when you were a student. The most important thing at that time was a good teacher and direct access to small supercomputers to try, play, fail, and at the end, to win. Now you are a scientist and your goal is to develop new methods. You have to have simple and easy access to mid-size machines at your location to try out the methods and optimize your codes and therefore you need to run jobs of relevant size. At the end you will need a huge amount of computing time to run the full-blown simulation on the biggest machine you can get access to. Therefore it is important for Germany to maintain und strengthen the strategy of the HPC-Pyramid.

You have been a manager in German Telekom industry. Having moved to a public organization like GCS, what do you see as the most important differences?

The differences are smaller than you might think. Especially within GCS we are dealing with big investments and operation of very complex machines. We are mainly driven by the requirements of scientists and researchers. On the other hand, we have to take care of the interests of the public funding agencies. So we also have customer relations, shareholders, operations, and financials. The structure is quite similar, and even the administrative processes are comparablely complex. The main difference is in the key performance indicators.

Would you still make the same decision today?

Sure

Dr. Müller thank you for the interview.

The Interview was conducted by the inside team.

Dr. Claus Axel Müller is Managing Director of GCS since 2009.

PRACE: Results of the 5th and 6th Regular Calls

The Partnership for Advanced Computing in Europe (PRACE) is offering supercomputing 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/Q system JUQUEEN in Jülich, of its Cray XE6 system Hermit in Stuttgart, and of its IBM iDataPlex system SuperMUC in Garching. France, Italy, and Spain are dedicating shares on their systems CURIE, hosted by GENCI at CEA-TGCC in Bruyères-Le-Châtel, FERMI, hosted by CINECA in Casalecchio di Reno, and MareNostrum, hosted by BSC in Barcelona.

The 5th call for proposals for computing time closed on 30 May 2012. Nine research projects have been awarded a total of about 310 million compute core hours on JUQUEEN, ten have been awarded a total of about 160 million compute core hours on Hermit, and eight have been awarded a total of about 200 million compute core hours on SuperMUC for the allocation time period 1 November 2012 to 30 October 2013. Seven of those research projects are from Germany, five are from France, two are from the Czech Republic, Finland, Italy, Netherlands, Spain, Switzerland, and the United Kingdom, each, while one is from Belgium.

The 6th call for proposals for computing time on the above systems closed on 14 November 2012. Six research projects have been awarded a total of about 260 million compute core hours on JUQUEEN, five have been awarded a total of about 160 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 4 March 2013 to 3 March 2014. Five of those research projects are from Germany, three are from France and Italy, each, two are from Switzerland, and one is from Australia, Cyprus, Denmark, Finland, the Netherlands, Portugal, the Russian Federation, and the United Kingdom, each.

The research projects awarded computing time cover many scientific areas, from Astrophysics to Medicine and Life Sciences. More details, also on the projects granted access to the machines in France, Italy, and Spain, can be found via the PRACE web pages www.prace-ri.eu/PRACE-5thRegular-Call and www.prace-ri.eu/PRACE-6thRegular-Call

The 7th call for proposals for the allocation time period 3 September 2013 to 2 September 2014 closed 26 March 2013 and evaluation is still under way, as of this writing.

Details on calls can be found on www.prace-ri.eu/Call-Announcements.

• Walter Nadler

Jülich Supercomputing Centre

Inauguration of the New HLRS Research Building

On October 31, 2012 the High Performance Computing Center Stuttgart / Höchstleistungsrechenzentrum Stuttgart (HLRS) inaugurated its new research building and five sided Cave environment. Together with about 60 guests from politics, administration, science, and participating construction companies, secretary of state of the Ministry of Finance of the state of Baden-Württemberg Ingo Rust officially handed over the new building to the president of the University of Stuttgart Prof. Wolfram Ressel and the director of HLRS Prof. Michael Resch.

Member of German Parliament (MdB) Dr. Stefan Kaufmann, Member of State Parliament Nikolaus Tschenk and secretary of state of the ministry of Finance Ingo Rust (from left to right). Photo: Boris Lehner

Members of the state parliament of the state of Baden-Württemberg and of the German national parliament joined the inauguration and expressed their support for high performance computing simulation at Stuttgart.
In his inaugural address secretary of state Rust highlighted that the University of Stuttgart takes a leading position in Europe in high performance computing. He continued to say that new building will allow HLRS to further develop and extend its research activities.

The president of the University of Stuttgart Prof. Ressel pointed at the virtual workbench that is created in the new building. In his statement he emphasized that this virtual workbench for research, development and design will be visible internationally and will certainly give new impulse to the col-laboration of HLRS with industry. The director of HLRS Prof. Resch said that the new building now allows bringing together all researchers and tools of HLRS in a single building thus substantially im-proving the quality of services HLRS provides to its customers in research, education and industry.

President of the University of Stuttgart Prof. Wolfram Ressel. Photo: Boris Lehner

Director of HLRS Prof. Michael Resch. Photo: Boris Lehner

The new research building for the first time brings together all HLRS researcher in the same building. Office space was increased by 1.400 m^2 totaling now 2.700 m^2. The building furthermore accommodates the newly designed virtual reality Cave environment which allows linking visualization directly to the supercomputers of HLRS in a single building. The five sided virtual reality environment stretches over three floors and allows diving into the visualization of complex simulation results interactively. Supported by a 64 processor cluster it can directly visualize results of large scale simulations done on the Cray XE6 HERMIT system. The Cave is also used for virtual reality support as well as for augmented reality support and hosts a driving simulator that allows for example analyzing drivers' behavior in everyday traffic situations.

The new HLRS research building was planned and constructed over two years. The costs of about 5.6 million Euros were covered by HLRS funds, mainly coming from third party projects. The building was built as an extension of the existing office building inaugurated in 2004 and hence

View of the HLRS 5-sided Cave environment. Photo: Wolfgang Schotte

A view of HLRS in October 2013 – Research building extension on the left.
Photo: Erika Fischer

is directly connected to the computer room of HLRS. The extension became necessary when HLSR nearly tripled its work force over the last 8 years. The existing buildings of HLRS will be complemented by an educa-tion and training center building in 2015. Planning will be finished in 2013 and construction work is expected to begin early in 2014. Funding will come again from HLRS funds based on third party project income. With the new education and training center HLRS will extend its training activities both in number of people trained and in scope of training. A main focus will be on training for industrial users of HLRS.

Piece of Art: A view of the interior of the research building.
Photo: Boris Lehner

• Michael Resch

University of
Stuttgart, HLRS

Happy Anniversary, Hermit! HLRS Supercomputer Hermit Completes First Year of Operation

On February 24, 2012 Hermit, at that point fastest supercomputer in Germany and all of Europe for civil science and research, was officially inaugurated at HLRS Stuttgart. Delivering a peak performance of more than 1 petaflops, Hermit achieved position 12 on the renowned TOP500 list (11/2011). Though Hermit subsequently gave up its initial ranking on the TOP500 list, the HLRS flagship computer is firmly established in the world of industrial supercomputing: to date Hermit still holds the title as the world's fastest supercomputer for use in industrial research (TOP500 11/2012, sub-list Industry).

At the day of its inauguration, Hermit was released for official use with the promise to be of great help in research relating to global challenges such as health, energy, environment and mobility. And Hermit lived up to its prediction. One year down the road, approximately 40 projects have taken advantage of Hermit's petascale computing power providing scientists and researchers with the tools to tackle the big questions of our time. The Cray XE6 system with its 113,664 processing cores has been fully operational and running stably from day one, delivering an average planned system uptime of 99%. With a mean usage capacity rate of 85%, the HLRS HPC system delivers proof that it was designed for sustained performance for real applications. Approximately 60% of scientific users fall into the research realm of engineering. With 30%, physics is the second largest user community relying on Hermit for their research activities.

User Projects

Despite the HLRS workhorse running literally full speed with close to no down time, demand exceeds the available computing time by a factor of 1.5. Users are standing in line waiting for their turn to receive computing hours for their various projects from scientific fields such as materials science, scientific engineering, life sciences, environment, energy and health as well as elementary particle physics and astrophysics. Computing time on Hermit is granted to the researchers by a scientific peer review process. Below are several of the many national and international science projects of breakthrough calibre completed on Hermit:

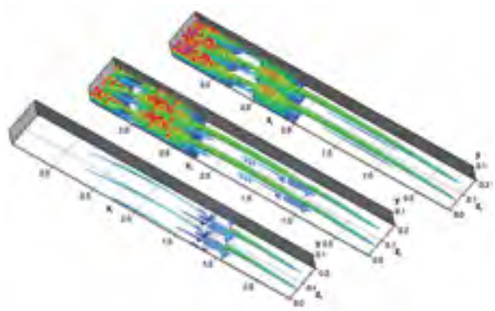

Project LAMTUR - Laminar Turbulent Transition in Aerodynamics Boundary Layers: Prof. Dr.-Ing. Ulrich Rist and Dr.-Ing. Markus Kloker from the Institute of Aerodynamics and Gas Dynamics of University Stuttgart and their teams are doing simulations on Hermit to achieve a comprehensive understanding of three-dimensional dynamic instability processes, which is a prerequisite for successful Laminar Flow Control (LFC).

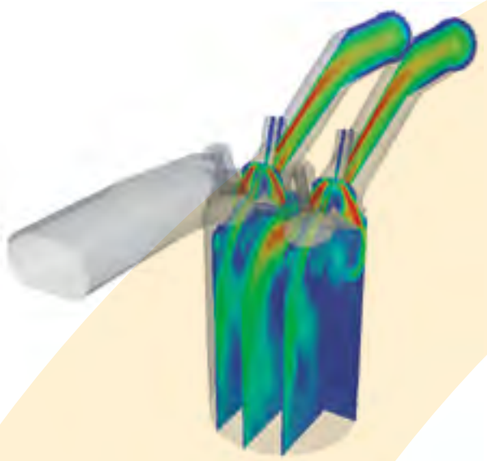

Project Direct Numerical Simulation of the Flow in an Internal Combustion Engine: Prof. Dr.-Ing. Wolfgang Schröder of RWTH Aachen (Fluid Mechanics and Institute of Aerodynamics) and his team use Hermit to study the complex flow field in an internal combustion engine, which has significant influence on the formation of the fuel-air-mixture in the combustion chamber and on the combustion process itself. The scientists aim for a further optimization of modern combustion engines to increase their efficiency and reduce the pollutant emissions.

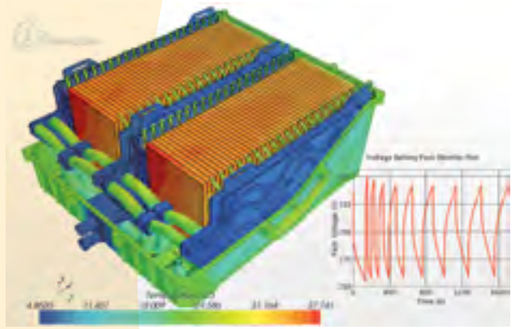

Project Development and Validation of Thermal Simulation Models for Li-Ion Batteries in Hybrid and Pure Electric Vehicles under leadership of Dr. Jenny Kremser, asc(s (Automotive Simulation Center, Stuttgart), concentrates on the development of a simulation environment for the electro-thermal layout of a lithium ion battery module in a vehicle. The project team pursues the development of optimized design concepts for electrified vehicles to fulfill increasing demands on energy consumption, driving range, and durability.

Hermit for European Research

Hermit qualifies as „Tier-0“ system in the European research infrastructure offered through the Partnership for Advanced Computing in Europe (PRACE), therefore the HLRS HPC system also offers computing power for large-scale scientific projects to scientists and researchers from Europe and beyond. As with the national projects, computing time is allocated to applicants based on a single peer-review process. In the last PRACE Regular Call for Proposals, some record-breaking project allocations of computing core hours were allocated to Hermit for the following projects:

Project UPSCALE: Prof. Pier Luigi Vidale and his team of scientists from NCAS-Climate (Department of Meteorology, University of Reading) and from the Met Office (Exeter) in the United Kingdom are using Hermit for compute-intensive simulations in order to increase the fidelity of global climate simulations and provide quantitative information about the frequency of high-impact events and their risks. The research activity comprises a large series of global experiments (an ensemble), with each member of the ensemble dynamically simulating 27 years of both current and future climates.

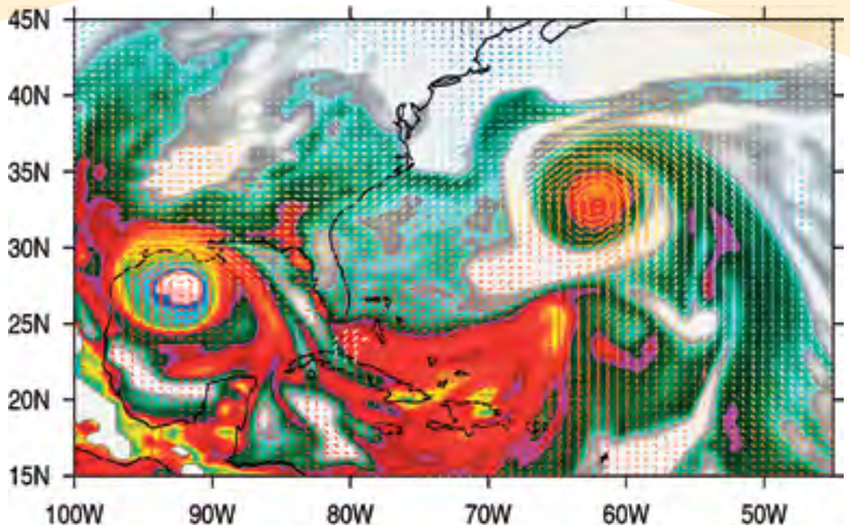

A tropical cyclone making landfall in Louisiana (left) and an extra tropical cyclone (right), both originated in the tropical Atlantic

Project Large Scale Molecular Dynamics Simulations of Nucleation: Professor Jürg Diemand of the Institute for Theoretical Physics of University Zürich (Switzerland) and his team are running very large (up to 8 billion atoms, millions of time-steps) molecular dynamics simulations of homogeneous nucleation from vapor to liquid. The unprecedented size of the simulations allows the formation of significant numbers of droplets (liquid nano-clusters) under realistic conditions: They can resolve nucleation even at relatively low, more realistic supersaturations and there is no significant depletion of the vapor phase during the simulations. The results will allow to test nucleation theories in a previously unexplored parameter range and hopefully lead to an improved understanding and description of this fundamental process.

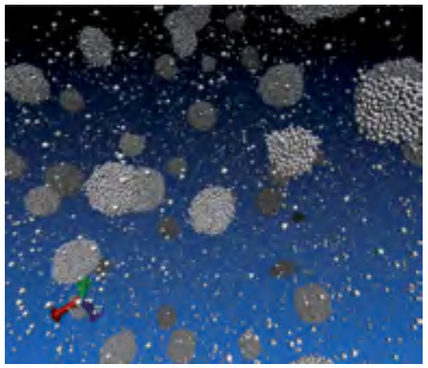

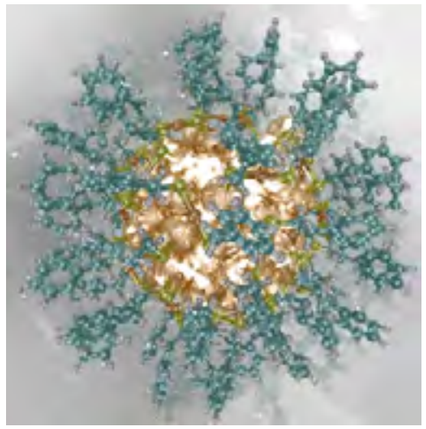

Project Plasmonic Ligand-Stabilized Gold Nanoclusters: Prof. Hannu Häkkinen (University of Jyväskylä, Finland) and his team are employing large-scale time-dependent density functional theory calculations to study

absorption of light by 2-3 nm gold and alloyed gold-silver nanoclusters that are defined to the molecular precision, i.e., by exact composition and structure. The project aims at breakthroughs in microscopic understanding of the "birth of a plasmon" in nanoscale noble metal clusters. This is of a wide scientific interest, since it will answer fundamental questions pertaining to transformation of nanoscale matter and nanoparticles from "molecular" to "metallic" regime with the concomitant change of optical response of the electrons from discrete transitions to collective behavior. This knowledge is important e.g. for designing and controlling plasmonic nanomaterials and nanosensors.

Outlook

Hermit is looking forward to about two more years of operation, but it will soon get support from a follow-up system from Cray. The new Cray XC30 – code named Hornet – will increase performance for HLRS users by a factor of between 4 and 5, and will be operational in 2014.

• Regina Weigand

Gauss Centre for Supercomputing (GCS)

Technical description (installation step 1)	
Peak performance	1.045 PFlops
Cabinets	38 with 96 nodes each
Number of compute nodes	3552
Number of compute cores	per node 2 sockets with 16 cores each: 113 664
Number of service nodes	96
Processor compute nodes	Dual Socket AMD Interlagos @ 2.3GHz 16 cores each
Memory/node	32 GB and 64 GB
Disk capacity	2.7 PB
Node-node interconnect	CRAY Gemini
Special nodes	External Access Nodes, Pre- & Postprocessing Nodes, Remote Visualization Nodes
Power consumption	2 MW maximal

SuperMUC boosts the largest molecular dynamics simulation by 4X in Number of Particles

MD simulation has become a recognized tool in engineering and natural sciences, complementing theory and experiment. Despite its development for over half a century, scientists still quest for ever larger and longer simulation runs to cover processes on greater length and time scales. Due to the massive parallelism MD typically exhibits, it is a preeminent task for high-performance computing.
The simulation code ls1 mardyn has been enhanced within the BMBF-funded project IMEMO [1] in order to provide an efficient tool for large-scale MD simulations of phenomena in inhomogeneous systems in chemical engineering. Targeting condensation processes and flow phenomena on the nanoscale, it supports rigid-body electroneutral molecular models composed out of an arbitrary number of Lennard-Jones (LJ) sites, point charges, point dipoles and point quadrupoles. As these scenarios typically require large particle numbers and also show heterogeneous density distributions of particles, see Fig. 1, sophisticated load balancing algorithms have been incorporated into the program to enable good scalability on large processor counts.
Together with colleagues from LRZ, laboratories of Thermodynamics at TU Kaiserslauten and U Paderborn and HLRS, we optimized our MD code on the micro-architecture level for a specific processor: the Intel Sandy Bridge EP based SuperMUC operated at the Leibniz Supercomputing Centre in Munich. This system features 147456 cores and is at present the biggest x86 system worldwide with a theoretical double precision peak performance of more than 3 PFLOPS or 6 PFLOPS for single-precision, placed #6 on the Top500 list of November 2012. The system was assembled by IBM and features a highly efficient hot-water cooling solution. In contrast to supercomputers offered by Cray, SGI or even IBM's own BlueGene, the machine is based on a high-performance commodity network: a FDR-10 infiniband interconnect by Mellanox. The optimizations targeting SuperMUC include the vectorization for exploiting the new AVX instruction set of the underlying Intel Sandy Bridge architecture, see [2] for details. This results in calculating always four particle interactions with one force calculation kernel call and therefore in drastic reduction of clock cycles. Furthermore, we applied

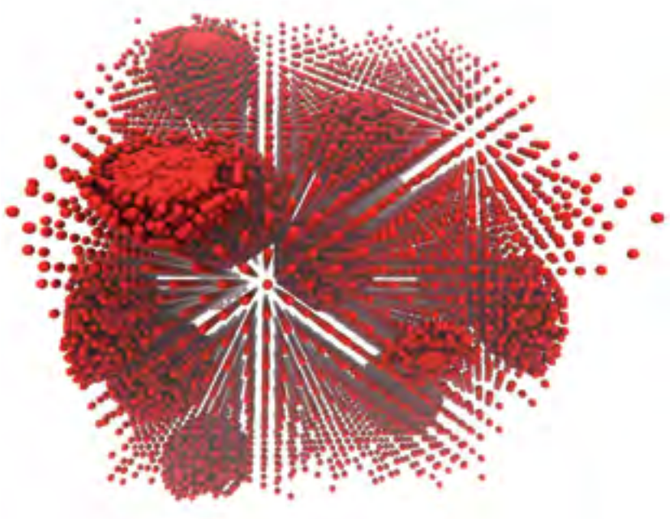

Figure 1.: Scenario with a heterogeneous particle distribution as it occurs in the simulation of nucleation (with help from LRZ application support).

memory optimizations which allow us to reduce the number of bytes needed per particle to 32, please refer to [3]. Due to the utilized LJ potential, we cannot leverage the instruction level parallelism of Sandy Bridge optimally. However, we were able to add a light-weight shared memory parallelization instead which accelerated the code by 12% when using Intel Hyper-Threading Technology on a per core basis.

In order to evaluate the performance of the MD simulation code ls1 mardyn, we executed different tests on SuperMUC. Our test scenario is similar to the initial configuration of nucleation scenarios, where particles are distributed on a regular grid. The single-center Lennard-Jones particles, modeling e.g. Argon, were arranged according to a body-centered cubic lattice, with a number density of $\rho\sigma^3$=0.78 in reduced units. The time step length was set to 1 fs. With respect to strong scaling behavior, we ran a scenario with $N=4.8*10^9$ particles, which fully utilizes the memory available on 8 nodes (128 cores), as 18 GB per node are needed for particle data. Fig. 2 nicely shows that a nearly perfect scaling was achieved for up to 146016 cores using 292032 threads at a parallel efficiency of 42 % comparing 128 to 146016 cores. To better understand this performance, we investigate the influence of the decreasing particle number per core, as it occurs in this strong scaling experiment, in Fig. 3: here we measured achievable GFLOPS depending on the number of particles simulated on 8 nodes. Furthermore, we show the influence of different cutoff-radii, which determine the number of interactions per molecule, as this parameter also severily effects the FLOP rate. To make a fair comparison with preceding publications possible, we conducted our runs with a cutoff radius of 3.5ρ. Already for $N=3*10^8$ particles, i.e. $2.3*10^6$ particles / core (approx. 8% of the available memory) we are able to hit the performance of roughly 550 GFLOPS per 8 nodes, which we also obtained for $N=4.8*10^9$ ($37.5*10^6$ particles per core).

Moreover, we performed a weak scaling analysis with $28.25*10^6$ molecules per core. This allowed us to perform the, to our knowledge, largest MD simulation to date, simulating $4.125*10^{12}$ particles on 146016 cores with one time step taking roughly 40s. For this scenario an absolute performance of 591.2 TFLOPS with a speedup of 133183 X in comparison to a single core was achieved, which corresponds to 9.4% single-precision peak performance efficiency.

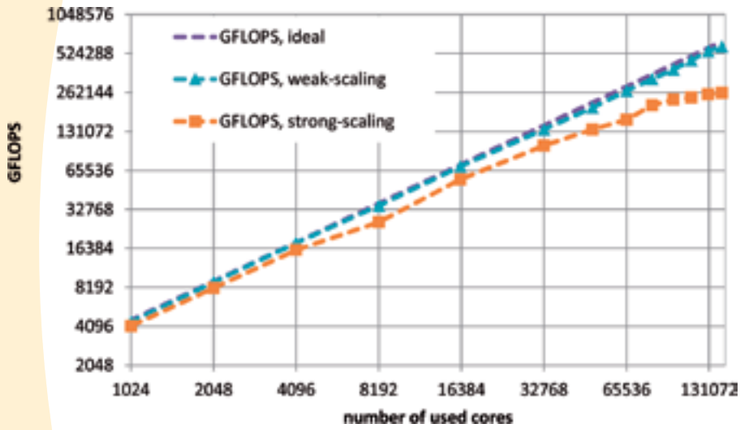

Figure 2a.:FLOPS measured for strong and weak scaling experiment.

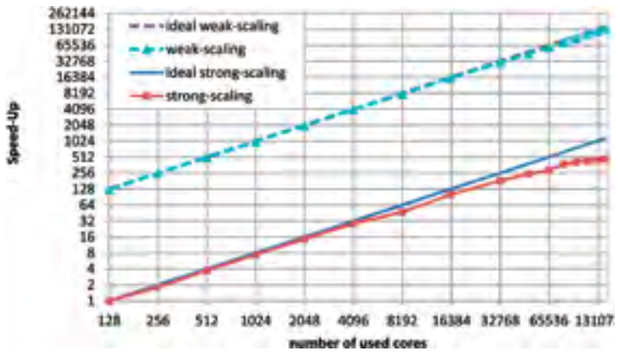

Figure 2b.: Speedup for strong and weak scaling experiment. For strong scaling, the speedup is normalized to 128 processes. For weak scaling, the speedup is relative to a single process.

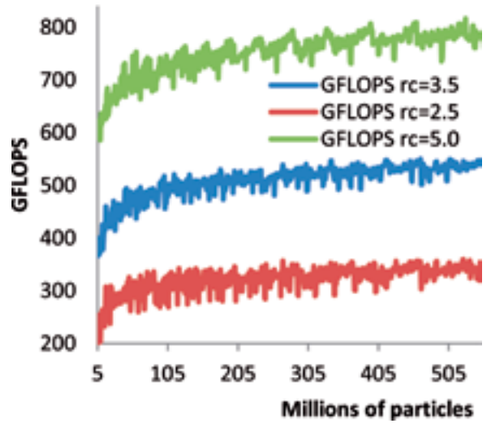

Figure 3.: Flop rate on 128 cores in dependence of the number of particles and the cutoff radius (which determines the number of interaction partners of a particle).

• W. Eckhardt
A. Heinecke

Technische Universität München, Institut für Informatik.

References

[1] **C. Niethammer, C. W. Glass, M. Bernreuther, S. Becker, T. Windmann, M. T. Horsch, J. Vrabec, W. Eckhardt**
Innovative HPC Methods and Application to Highly Scalable Molecular Simulation (IMEMO). In Inside - Innovatives Supercomputing in Deutschland, Volume 10(1), April 2012.

[2] **W. Eckhardt, A. Heinecke**
An efficient Vectorization of Linked-Cell Particle Simulations. In ACM International Conference on Computing Frontiers, p. 241–243, May 2012.

[3] **W. Eckhardt, T. Neckel**
Memory-Efficient Implementation of a Rigid-Body Molecular Dynamics Simulation. In Proceedings of the 11th International Symposium on Parallel and Distributed Computing - ISPDC 2012, p. 103–110. IEEE, Munich, June 2012.

Golden Spike Award by the HLRS Steering Committee in 2012

Numerical simulation of mantle convection with a particle-in-cell method to include the effects of partial melting

Thermo-chemical convection in the silicate mantles of terrestrial planets is one of the most important processes after accretion and core formation, controlling the planetary thermal evolution. The interior heat released during the formation processes of the planet and generated by radioactive decay is mainly transported toward the surface by convective motions due to thermal buoyancy. Of uttermost importance for the mantle dynamics is the strongly temperature-dependent solid creep behavior of the silicate material resulting for most planets in a stagnant surface layer – although tectonically deformed - and a convecting interior. The rigid upper layer is further shaped by volcanoes that witness melting processes in the interior.

The silicate mantle of terrestrial planets is composed of an assemblage of minerals, which start to melt when the mineral with the lowest melting point reaches its melting temperature. This process is called partial melting since in this case the mantle material is not entirely liquid but consists of a mixture of liquid and crystals. Due to its lower density compared to that of the surrounding rock, the liquid can percolate toward the surface, rising through cracks and channels. If the liquid cannot find a path to further rise, it accumulates in the so-called magma chambers. The more magma is supplied, the stronger is the pressure increase in a magma chamber and with time rock fractures around the chamber, resulting in volcanic eruptions at the surface. Surface expressions of the partial melting process have been revealed by planetary missions on all terrestrial bodies in our Solar System. The NASA's spacecraft MESSENGER (MErcury Surface, Space ENvironment, GEochemistry and Ranging) has found clear evidence of a widespread distribution of pyroclastic deposits all over Mercury's surface [1]. In the case of Venus, variations in the thermal emissivity of the surface observed by the Visible and Infrared Thermal Imaging Spectrometer on the ESA's Venus Express spacecraft suggest ages of some lava flows to be younger than 2.5 Ma (Millions of years) [2]. On the Moon, ancient basaltic regions have been found, which are mainly concentrated on the hemisphere facing the Earth. Mars owns the most spectacular volcanic features and also the largest volcanic region in our Solar System, i.e., the Tharsis rise. High resolution images of Martian volcanoes, provided for example by the Mars Express mission, have shown that some of them exhibit nearly uncratered, and hence extremely recent, lava flows, as young as 2 Ma [3].

In our study, we want to understand under which conditions and in which regions melting occurs in the planetary interiorand the effect of partial melting on the subsequent dynamic evolution of a planet. To this end, we apply 2-3D mantle convection simulations that consider partial melting and associated density changes of the mantle material. An example is shown in Fig. 1. The iso-surface of the temperature at 1942 K represents mantle upwellings rising from the core-mantle boundary and undergoing partial melting close to the surface (partial melt region in red color).

Numerical Method

Due to the non-linear nature of the processes in the interior of planetary mantles, analytical solutions can be found only for strongly simplified scenarios. Laboratory experiments, in addition, are restricted to small parameter ranges that are not always relevant for planetary evolution. Over the last years, with the remarkable increase of computational power, the numerical solution of planetary convection problems has grown to be the most powerful approach to model the temporal evolution of the mantle flow with complex rheologies in a 3D spherical geometry. To this end, massive parallel computing has been established as a standard procedure to explore the interior of planetary bodies.

We use the spherical code Gaia to solve the conservation equations of mass, momentum and energy via finite-volume discretization either in 2D cylindrical [4] or 3D spherical [5] geometry on both regular and fully irregular grids. The code is written in C++ and uses its own modules for mesh generation, sparse matrix implementation via the Harwell-Boeing matrix class, iterative solver of a linear system via BiCGStab and does not need any additional libraries. The parallelization via message passing interface (MPI) has been successfully tested with various MPI implementations such as IntelMPI, MVAPICH, HP-MPI. We further use a particle-in-cell method (PIC) to treat thermo-chemical convection [6]. To account for the advection of different material properties that change during the melting process in the mantle, we use massless particles whose motion is first calculated from the velocity field computed on a fixed mesh. A material property field is then interpolated from the distribution and values of the particles (Figure 2). The method has the advantage over classical grid-based methods of being essentially free of numerical diffusion and to enable the natural advection of an arbitrary large number of different properties by solving only one equation. At the moment, only a small number of numerical codes worldwide beside Gaia can handle 3D spherical geometry with a particle-in-cell method to account for material properties. Typical resolutions used in our simulations are: 1.5 x 105 points for the fixed grid and 2.2 x 106 particles for the transport of material properties for the 2D cylindrical geometry whereas in 3D spherical geometry the number of points for the fixed mesh increases to 2.8 x 106 and the number

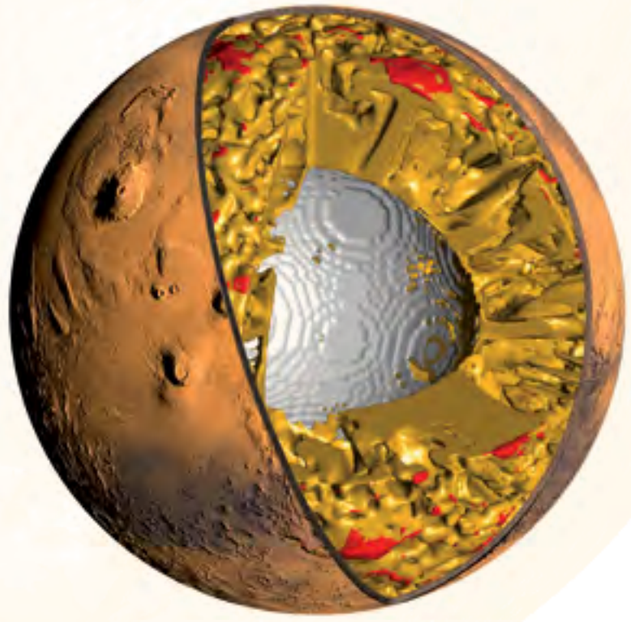

Figure 1: 3D numerical simulation showing a temperature iso-surface at 1942 K and the partial melt regions in red.

of particles to about 22 x 107. Such simulations require high computational power which can only be achieved when using supercomputers like XC4000 at the SCC in Karlsruhe.

Results

We have applied our model to study the formation and stability of mantle heterogeneities in the interior of Mars as suggested by the so-called SNC meteorites (Shergottites, Nakhlites, Chassignites) originating from Mars [7]. Their geochemical analysis suggests the existence of two reservoirs, which are depleted in incompatible lithophile elements relative to chondrites (i.e., primitive, stony meteorites) and of a third reservoir, which is instead enriched in these elements. According to the analysis, the chemical reservoirs have formed very early in the evolution by complex internal processes and have not been remixed since then.

To explain the formation and persistence of the reservoirs but also the observed recent volcanic activity on Mars, specific conditions in the density (compositional) distribution of early Mars are required, i.e., an initial layer of lower density compared to the surrounding homogeneous mantle. Such a layer could be the consequence of large scale melting during the late stage of planetary accretion and core formation but is not consistent with a fractionated mantle due to freezing of a global magma ocean as commonly assumed in the literature [8]. Figure 3 shows density variations in the mantle with time due to the mantle depletion in crustal components during the melting process assuming an initial depleted mantle layer. In general, a variation in the mantle compositionarises by extraction of partial melt, which leaves behind a residuum depleted in incompatible elements and modified in modal mineralogy. This phenomenon, known as mantle depletion, is geodynamically important since a melt-depleted residue is expected to be more buoyant than its parent material. The decrease in residue density derives from both changes in mineral density and changes in the relative proportion of minerals. In our model, the resulting mantle inhomogeneities form a lighter layer close to the surface that prevents the lower mantle from efficient cooling. This layering helps keeping the lower mantle warm and prolongs the partial melt production, which directly impacts the volcanic activity of the planet. The mantle heterogeneities obtained in this case remain stable during the entire thermal evolution – the layering is even stabilized for an initially wet mantle that dries out with time due to volcanism - and could explain the isotopic characteristics of the SNC meteorites.

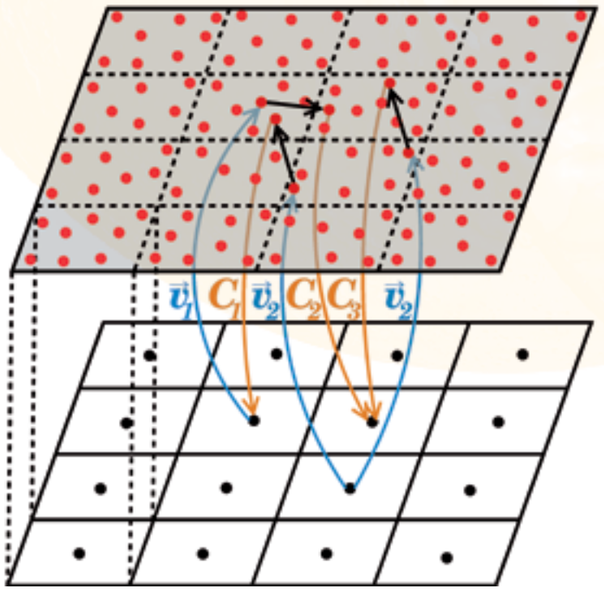

Figure 2:The particle-in-cell method first interpolates the velocity from the field to the particles (blue arrows). The particles move with the given velocity (black arrows) and the material properties are then interpolated from the particles back to the grid (brown arrows).

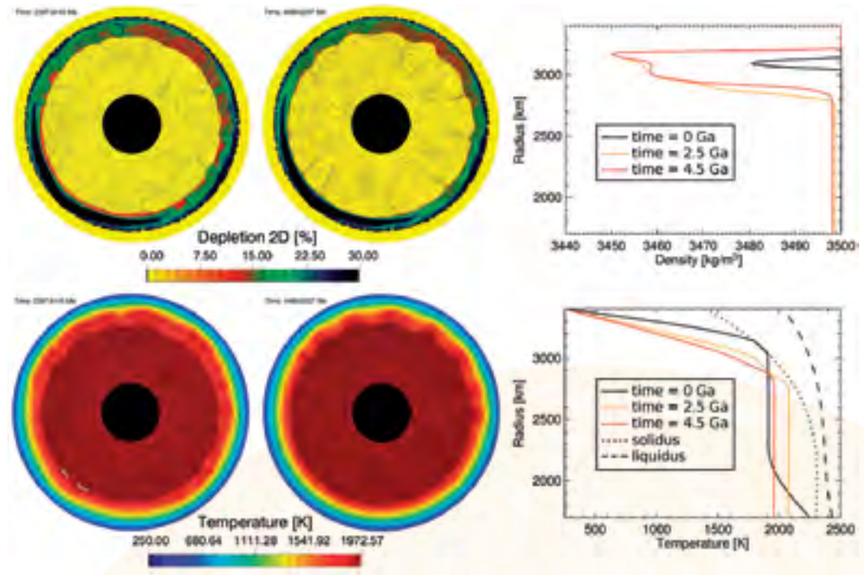

Figure 3: Formation of geochemical reservoirs due to density variations in the mantle caused by partial melting and mantle depletion in crustal components. Two snapshots of the mantle depletion and temperature at 2.5 and 4.5 Ga are shown in the left part of the figure. The right part of the figure shows density and temperature profiles corresponding to the snapshots on the left.

On-going Research/Outlook

To improve our model, future work will include collaboration with the Institute of Mineralogy of the WestfälischeWilhelms-University Münster to assure a proper treatment and interpretation of the mineralogical aspect. This will also help to narrow down the parameter range for which our results might explain the data obtained from the analysis of the SNC meteorites.

Acknowledgements

This research has been supported by the Helmholtz Association through the research alliance “Planetary Evolution and Life”. The project was conducted at the SCC Karlsruhe on the XC4000 supercomputer.

References

[1] **L. Kerber, J. W. Head, D. T. Blewett and S. C. Solomon, L. Wilson S. L. Murchie, M. S. Robinson, B. W. Denevi and D. L. Domingue.**
Planetary and Space Science, 2011.

[2] **S. E. Smrekar, E. R. Stofan, N. Mueller, A. Treiman, L. Elkins-Tanton J. Helbert, G. Piccioni and P. Drossart.**
Science, 2010.

[3] **G. Neukum, R. Jaumann, H. Hoffmann, E. Hauber, J. W. Head, A. T. Basilevsky, B. A. Ivanov, S. C. Werner, S. van Gasselt, J. B. Murray, T. McCord and The HRSC Co-I Team.**
Nature, 2004.

[4] **A.-C. Plesa.**
INFOCOMP 2011: The First International Conference on Advanced Communications and Computation, 2011.

[5] **C. Huettig and K. Stemmer.**
Physics of the Earth and Planetary Interiors, 2008.

[6] **A.-C. Plesa, N. Tosi and C. Huettig.**
Integrated Information and Computing Systems for Natural, Spatial and Social Sciences, 2012.

[7] **J. J. Papike, J. M. Karner, C. K. Shearer and P. V. Burger.**
Geochim and CosmochimActa, 2009.

[8] **L.-T. Elkins-Tanton, S. E. Zaranek, E. M. Parmentier, and P. C. Hess.**
Earth and Planetary Science Letters, 2005.

- A.-C. Plesa[1,2]
- D. Breuer[1]
- T. Spohn[1,2]

[1] Department of Planetary Physics, German Aerospace Center (DLR), Berlin, Germany

[2] Institute of Planetology, Westfälische Wilhelms-University Münster, Germany

Golden Spike Award by the HLRS Steering Committee in 2012

Effects of an oblique roughness on hypersonic boundary-layer transition

Charles Yeager was the first man who officially broke the sound barrier in horizontal flight. On October 14, 1947, he accelerated his rocket-powered aircraft X-1 to supersonic speed. Engineers have a characteristic number, the so-called Mach number, to describe the speed of a vehicle in relation to the speed of sound. Yeager achieved Mach 1.06 (about 1300 km/h) in peak. Since that day engineers have been dreaming of manned flight at even higher speed: For the so-called hypersonic flight an aircraft has to be accelerated to at least about five times the speed of sound. Though there is no discrete "hypersonic barrier" this speed regime is characterized by strong surface heating due to friction and compression. The barrier is thus set by the temperature that a vehicle can stand for a specified time.

In the following we deal with roughness elements on otherwise smooth surfaces in a flow at hypersonic speed. These roughness elements significantly influence the flow along the surface. The flow near the solid wall of an arbitrary body is called a boundary-layer flow. It describes the region where the oncoming air is slowed down to relative zero velocity at the wall itself due to friction effects between the air and the vehicle surface. Boundary-layer flows can exhibit two major states. The laminar state is, roughly spoken, a smooth, steady flow in which wall-parallel layers of fluid slide along each other with only minor interaction. The turbulent flow is dominated by unsteady chaotic movement of the fluid with a large amount of wall-normal interaction and fluid exchange. Naturally, a boundary layer starts laminar. Small disturbances therein experience amplification downstream until they reach a high intensity, and the flow gets turbulent. The changeover from laminar to turbulent flow is called transition. A roughness element can have a substantial influence on the transition process dependent on its height and the mechanisms at work. Hence, it is important to understand how a roughness element changes the natural transition mechanism on a smooth plate. In most cases it promotes transition to turbulence which in general is unwanted. A premature transition and the associated turbulent boundary-layer cause a higher heat flux from the hot boundary-layer flow to the wall. An increased heat flux causes a larger technical effort to "keep the aircraft cool", and, subsequently, leads to an increased structural weight due to an enforced thermal protection system (TPS). Moreover, from the economical point of view, every additional kilogram for the TPS means one kilogram loss in payload.

Figure 1: Ceramic-tile gap filler protruding from the heat shield of the Space Shuttle. Source: www.nasa.gov

Groskopf & Kloker [1] tackle, among few others, the task of understanding the underlying roughness-induced transition mechanisms applying numerical methods to simulate and analyse the boundary-layer flow in the wake of a three-dimensional (3-d) surface roughness. To this end high-performance computers provided by the High Performance Computing Center Stuttgart (HLRS) are employed. The following flow setups resembles a free-flight and a "cold" wind-tunnel experiment. In the former we investigate a Mach-4.8 boundary-layer flow with a freestream temperature of $T_{\{infinity\ symbol\}}$ = 220K. The resulting air speed is about 1430 m/s (exceeding 5100 km/h). The roughness configuration is derived from a Space Shuttle flight experiment (STS-119) that has been conducted after gap fillers that have been found protruding from the heat shield at the shuttle's bottom side (figure 1). The roughness element is slat-like and placed obliquely (Ψ = 45°) with respect to the oncoming flow {u{∞}}, see figure 2. Its height k is 30% of the smooth-wall boundary-layer thickness {greek letter delta}$_{sw}$ at its location. The wall is radiation adiabatic, i.e. the heat flux to the wall is radiated by the wall.

Our investigation of the roughness element's influence on the laminar-turbulent transition starts with a check whether a laminar flow still exists, followed by a look at how the smooth-plate flow is deformed by the 3-d roughness element in this case. The first crossplane in front of the rough-

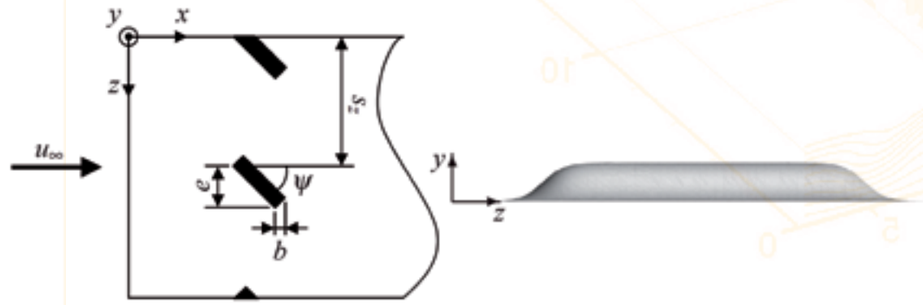

Figure 2: Sketch of the roughness configuration. Left: Top view with element dimensions b and e, spanwise spacing s_z and rotation angle Ψ. Right: Front view.

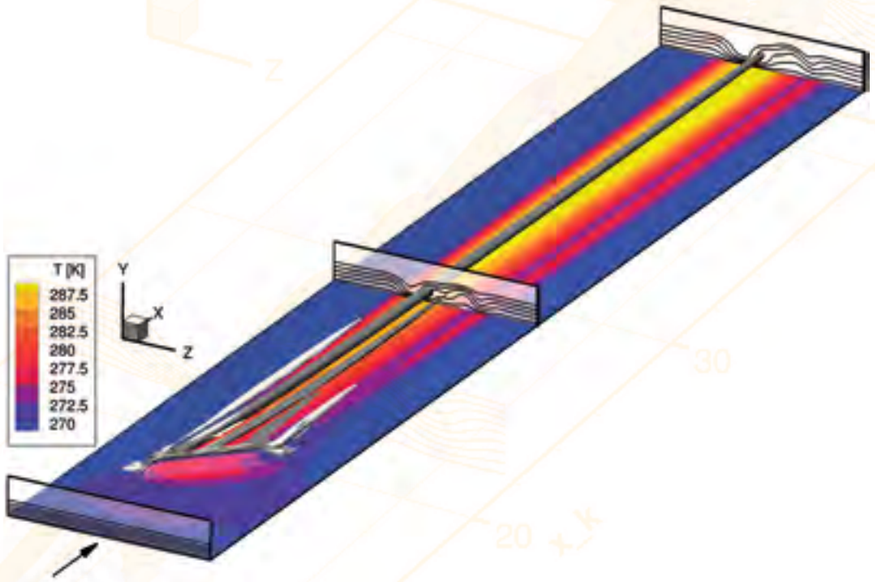

Figure 3: View of roughness-induced steady vortices (white: clockwise, and grey: counter-clockwise rotation sense looking in downstream direction x) and wall-temperature (color shading: blue indicates cool, yellow hot regions). The crossplanes show isolines of streamwise velocity. Oncoming flow is indicated by the arrow. Top: free-flight scenario, radiation-adiabatic wall. Bottom: wind-tunnel conditions, adiabatic wall.

ness (figure 3 top) represents the smooth-wall flow. No variation is evident in spanwise (z) direction. The roughness element deforms the flow downstream. Several steady (time-invariant) longitudinal vortices (white and grey isosurfaces in figure 3 top) are induced with varying strength. Some of the vortices vanish shortly downstream the roughness element, e.g., the white one generated at the trailing edge. Others are strong, and so is their influence on the shape of the flow. Here, the wake flow is dominated by the grey vortex generated at the leading edge of the roughness element. Between neighbouring counter-rotating vortices either fast-flowing air from outside the boundary layer is driven towards the wall (the isolines in the crossplanes getting closer to the surface as well as each other) or the slow fluid near the wall is pushed towards the boundary-layer edge (the isolines bending away from the wall).

The deformation of the flow has a strong impact on the wall-temperature. Whereas the temperature distribution in spanwise direction is uniform for a smooth plate (see figure 3 top upstream of the roughness), the front of the roughness element sticking out from the surface is heated up. But the wall downstream of the element gains even higher temperatures. The highest values are located where fast air comes closest to the wall. Here, the large velocity of the fluid has to be decelerated to zero within a small wall-normal distance due to friction, generating a large amount of thermal energy resulting in the high wall temperature. The spanwise deformation of the flow also alters the stability properties of the flow. The flow is destabilized which might speed up laminar-turbulent transition leading to an even higher heat load at the wall.

Two different methods are applied to investigate the stability of the flow: the linear stability theory (LST) and direct numerical simulation (DNS).
The latter method is a simulation of the complete governing time-dependent flow equations, a rather complex and costly approach. Based on the steady flow shown above, the LST reduces the complexity of the flow problem by making simplifying assumptions, i.a., the 3-d problem is reduced to two dimensions for consecutive crossplanes here. This way information about the complex flow can be gained with less effort. Due to the reduced cost the LST has been established as an engineering tool, albeit for 2-d base flows with significant variation in only one (usually the wall-normal) direction which allows for further simplifications (1-d eigenfunction). With a calibration based on experimental data this theory is applied to roughly predict a transition location though naturally there is no exact location but a finite region. Here, the results of (2-d eigenfunction) LST and (3-d) DNS shall be compared.

LST results reveal eigenmodes of the steady flow, and whether they are amplified. These eigenmodes are assumed here to be waves in streamwise direction and time with a 2-d amplitude distribution in the analyzed crossplane. The amplitude distribution of the most unstable eigenmode is shown in figure 4. This eigenmode is tracked along the streamwise direction to gain information about its downstream evolution.

An unsteady DNS is conducted by introducing disturbances with chosen frequencies into the boundary layer. This is done via blowing and suction through a hole at the wall upstream of the roughness element. In experiments this setup can be implemented based on loudspeakers or a glow discharge beneath the plate. Here, two frequencies are excited, ω_r=5 and ω_r=10, corresponding to about 70 kHz and 140 kHz. Small forcing amplitudes have been chosen to preserve the comparability of DNS and LST. The downstream evolution of the introduced disturbances is influenced by the roughness element and the altered flow in its wake. Once the start-up phenomena vanish from the computational domain the time signal of the DNS data becomes periodic upstream of full turbulence. Thus a Fourier analysis can be applied to gain the growth of disturbance amplitudes as well as their shape.
Results show that the DNS disturbance amplitude shape agrees well with the data from the LST analysis for both frequencies. The 3-d flow deformation caused by the LST eigenmode can be constructed by superposition of the steady flow with the finite-amplitude eigenmode. Though this approach neglects any (non-linear) interaction between those two, the result shows vortex structures similar to those found in the unsteady DNS (figure 5). In particular this comparison shows that the most-amplified eigenmode from

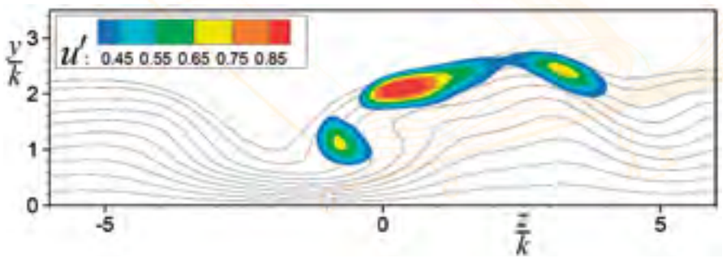

Figure 4: Normalized perturbation amplitude of streamwise velocity for the most unstable eigenmode.

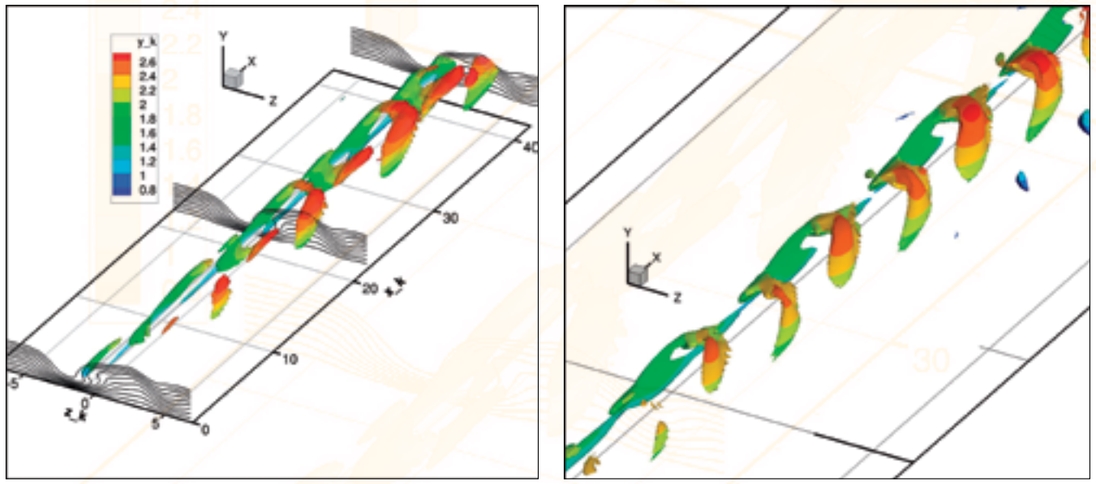

Figure 5: Snapshots of unsteady vortex structures reconstructed from (2-d eigenfunction) LST results (left) and from DNS (right) for frequency ω_r=5.

LST is indeed the dominant disturbance in the unsteady 3-d flow. A comparison of DNS- and LST-disturbance growth is shown in figure 6. The prediction by LST is excellent for the low frequency. Discrepancies are found near the roughness. They are caused by phenomena that can not be reproduced by LST since the underlying assumptions do not apply. For the high frequency, the growth found by LST is somewhat to low. As a reference, figure 6 shows the disturbance evolution on the smooth plate without roughness. Up to a short distance in front of the roughness the behaviour is identical. Approaching the roughness element disturbance amplitudes are boosted (see inset). Then, along the element, they decrease below the smooth-plate values. Downstream the element their growth rates are many times larger than for the reference flow. The high frequency disturbance even experiences amplification only locally on the smooth plate.

In conclusion it can be said that a 3-d oblique roughness in a hypersonic boundary layer induces one dominant ongitudinal vortex at the element's leading edge that deforms the flow in its wake. This, in contrast to a smooth-plate flow, leads to a hot strip at the wall putting higher demands on the heat shield of hypersonic vehicles. Moreover the flow's instability is significantly increased, laminar-turbulent transition is strongly promoted causing even a higher heat load. However, transition is somewhat slower than in subsonic flow, and the local regions of flow reversal near the element play a less significant role. The LST that can be used based on steady base-flow solutions is not able to reproduce all flow phenomena shown by the DNS but the main effects can be computed with sufficient accuracy, provided the base flow is accurately represented, a task often underestimated. However, since LST is based on small (linear) disturbance amplitudes, late (non-linear) stages of the transition process can only be investigated with a DNS.

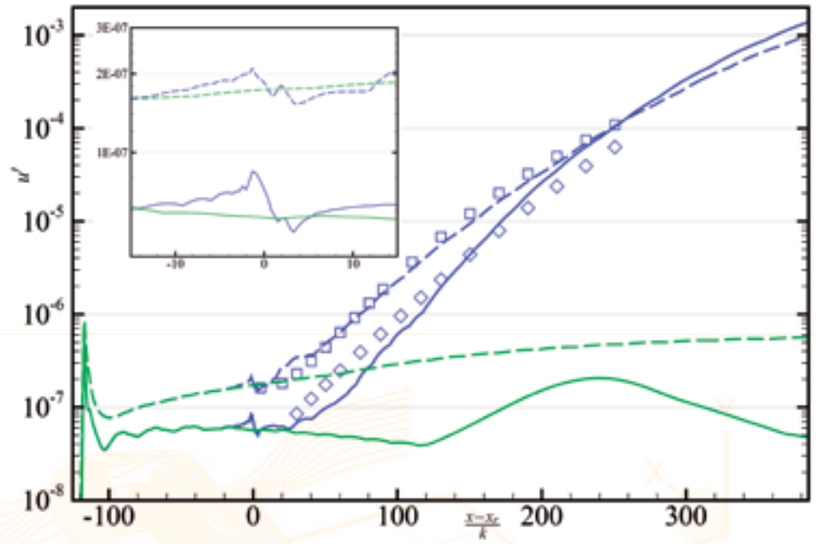

Figure 6: Disturbance amplitude growth for streamwise velocity component u' from DNS (lines) and LST (symbols). Blue lines show flow with roughness. Dashed line, ◊ represent frequency ω_r=5; solid line, □ show ω_r=10. For comparison green lines show disturbance growth for smooth-plate flow. Roughness element is located at $(x-x_r)/k$=0. The inset shows the vicinity of the roughness element in detail.

Ongoing unsteady DNS run on a spatial grid with up to 540 million grid points and take up to about 1.1 million CPU hours, or 127 CPU years, on the CRAY XE6 at HLRS. For comparison the LST analysis takes less than 0.1% of that CPU time, though it requires large amounts of RAM.

References

[1] **G. Groskopf & M. J. Kloker:** Effects of an oblique roughness on hypersonic boundary-layer transition. In: High Performance Computing in Science and Engineering '12 – Transactions of the High Performance Computing Center, Stuttgart (HLRS) 2012, Springer-Verlag 2013.

• Gordon Groskopf
• Markus J. Kloker

Institut für Aerodynamik und Gasdynamik, University of Stuttgart

Golden Spike Award by the HLRS Steering Committee in 2012

Phase Separation in Colloidal Suspensions by Collective Growth of Domains

Colloidal suspensions are fluid materials which contain myriads of micrometer-size solid particles in a solvent. These materials find ubiquitous application in the chemical and food industries (various kinds of paints, ink, cosmetic creams, toothpaste, coatings for various purposes); even bullet-proof vests are produced using such colloidal dispersions. Often one is interested in suspensions that can separate in several phases and are able to coexist with each other. For instance, if the solvent also contains polymers, i.e. long flexible macromolecules, which have a random coil-like structure with a radius comparable to the colloid particles, an attractive force between colloidal particles arises from entropic effects, the so-called "depletion attraction". By a suitable choice of polymer concentration and molecular weight these interactions and hence the state of the suspension (homogenously mixed or separated into colloid-rich and polymer-rich regions) can be manipulated.

Of course, for processing such dispersions one uses containers, and then the particles also interact with the confining walls, which often leads to the formation of colloid-rich "wetting layers" at the walls. Such effects have a particularly pronounced influence in applications of microfluidic devices, where the dispersion is confined in tube-like or slit-like pores, with distances between the walls exceeding the colloid diameter at most by 2 or 3 orders of magnitude. The interplay of wetting layer formation and lateral domain formation (in the directions parallel to the walls of the slit pore) is a challenging problem. Experimental studies of such problems are on-going (e.g. [1]) and clearly show that the kinetics of phase separation in such systems is a multi-stage process. However, in order to be able to experimentally control the resulting domain morphologies, a better theoretical understanding is urgently needed: one would like to know how different effective potentials between the colloids and a wall (which one could modify by suitable wall coating) affect the domain growth, and one also would like to understand the role of hydrodynamic (i.e., flow-mediated) interactions, whose strength could be controlled by changing the solvent viscosity.

However, for such soft matter systems it is impossible to formulate a theoretical first-principles approach, and hence the method of choice to address such problems is large-scale computer simulation. Colloidal suspensions are a true multi-scale problem. Only by a substantial investment of supercomputer resources at the massively parallel supercomputer HERMIT progress could be obtained. Remember that the fluid molecules in the solvent are of nm (10^{-9}m) size. The polymer coils can be approximated as soft spheres and the colloids as (almost) hard spheres, both of size (10^{-9}m). However, the domain structures that one wishes to study are in the range of 10 to 100 μm!

Thus, even on petascale computers a treatment taking full account of molecular details of the solvent molecules is impossible, but it is also not necessary: the physical principle that can be exploited to simplify the problem is the separation of time scales for the dynamics of solvent molecules versus colloids. The latter move many orders of magnitude slower than the solvent molecules. On the typical time scale of colloid motion, all correlations in the motions of solvent particles are fully relaxed, and hence it is possible to replace the solvent by an effective fluid, particles undergoing random collisions, but otherwise behaving like a dense ideal gas. This is achieved by the so-called "multiparticle collision dynamics" method (MPC) [2]. With the appropriate choice of the density and hence the viscosity of this effective fluid, the hydrodynamic interactions in the motion of colloid particles and polymers are properly included.

When we take a colloid diameter as our unit of length, a typical linear dimension of the simulated system is 256 in both directions parallel to the confining walls, which are 10 colloid diameters apart. The walls repel both particles with a smooth potential, decaying with the 12th power of distance, and a range of 50% of the respective particle diameters (polymers were assumed to have a diameter of 0.8 relative to colloids; see [3] for further details of the model). Such a system may contain 236859 colloid particles, about a million polymers, and 52 million effective solvent particles. Using the domain decomposition scheme and message passing (MPI), 4096 CPU cores could be used in parallel to study

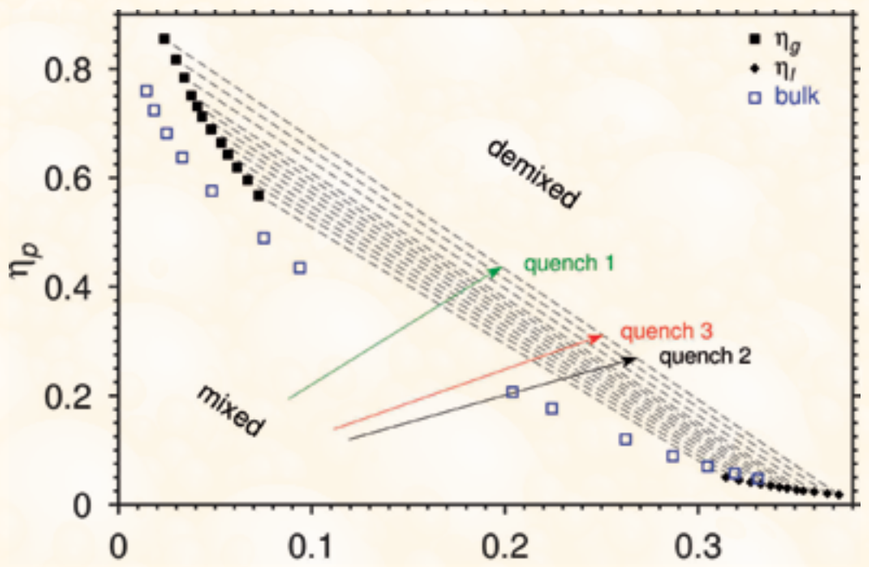

Figure 1: Phase diagrams of the model colloid-polymer mixture in the bulk (open symbols) and confined by two planar repulsive walls a distance D=10 colloid diameters apart (full symbols). The locations of the full squares and diamonds indicate the packing fractions η_c, η_p of the colloids and polymers in the coexisting polymer-rich (squares) and colloid-rich (diamonds) phases, respectively. These coexisting phases are connected by "tie lines" (broken straight lines: a state on such a line is a mixture of the two coexisting phases with volume fraction 1-x, x proportional to the length along the line from the state point to the coexisting phases on the left and right, respectively ("lever rule")). Arrows show three "quenching experiments" (cf. text). From [3].

such systems, with almost no loss of efficiency in comparison to simulations without solvent particles [4].

Unlike in experiments [1], one does not need to worry about small perturbations due to gravitational forces acting on the colloids, and one can obtain the equilibrium phase behavior of the studied system, both with and without confining walls, by separate simulations, using completely different methods, namely Monte Carlo simulations in the grand-canonical ensemble. Of course, the ideal gas-like solvent particles do not affect the static phase diagram, and can in this context simply be omitted. Figure 1 shows the resulting phase diagrams, using the "packing fractions" η_p, η_c of polymers and colloids as variables (the packing fraction gives the percentage of the total volume of the system taken by the particles multiplying the spherical volume of one particle by their number density). One sees that both in the bulk (a bulk system is simulated by choosing a cubic box with periodic boundary conditions in all three directions) and in the confined geometry the system is miscible if either η_p or η_c is sufficiently small (or both). The symbols represent the phase boundary, separating the miscible system (below) from the demixed region (above). One sees already that confinement enhances miscibility on the polymer-rich side, where η_c is small $(\eta_c \leq 0.1)$, but not on the colloid-rich side of the phase diagram (there, the full and open symbols in Fig. 1 almost coincide). This happens because for small η_c the few colloids in the system get almost completely attracted to the walls, forming there enrichment layers with enhanced colloid density, and hence there is no need for unmixing in the remainder of the system any more: it is lateral

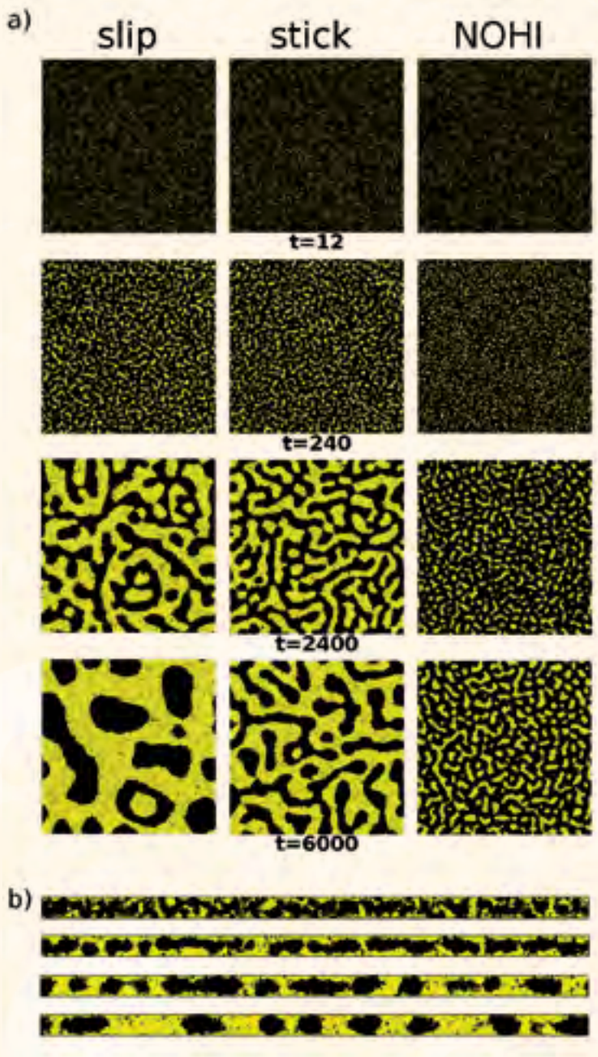

Figure 2: a) Snapshots of the (coarse-grained) domain pattern for the quench #3 in Fig. 1, showing only the (x,y) coordinates of polymers as black dots, for the system of size 256x256x10 at times t=12, 240, 2400 and 6000 Molecular Dynamics time units, as indicated. The coordinates of colloid particles and of the solvent fluid particles are not shown. The data refer to different treatments of hydrodynamic interactions: slip (left) and stick (middle) boundary conditions, and no hydrodynamic interactions ("NOHI", right). b) Side views of the same system with slip boundary conditions, at times t=120, 324, 984 and 1680, from top to bottom. From [3].

phase separation parallel to the wall that matters for the confined system, in the direction normal to the walls some inhomogeneity always occurs. The so-called "tie lines" (shown as dashed straight lines in Fig. 1) indicate the respective two phases that coexist with each other in each case where lateral phase separation has occurred. This complete knowledge about the phases that coexist in equilibrium, in relation to the precise knowledge on interactions, is a feature by which our simulations clearly reach beyond the corresponding experiment [1].

The arrows in Fig. 1 now indicate three "quenching experiment"-simulations, i.e. reduction of the system volume at constant ratio η_c / η_p. The initially homogeneously mixed system becomes thermodynamically unstable, concentration fluctuations grow spontaneously, and various pattern of growing domains form (Fig. 2). These patterns indeed have great similarity to what one finds in corresponding experiments, but we have the bonus that we can isolate the effect of hydrodynamic interactions clearly and thus elucidate their role in domain growth. This is done by an appropriate choice of hydrodynamic boundary conditions for our solvent fluid particles at the walls, all the way from perfect slip to perfect stick (Fig. 2). For perfect slip, hydrodynamic interactions are not affected in the directions parallel to the walls, but for perfect stick, they are partially screened (the screening increases when the distance between the walls decreases). One sees that hydrodynamic interactions clearly speed up the domain growth. For stick boundary condition the pattern at t=6000 Molecular Dynamics time units resembles the pattern at t=2400 if slip boundary conditions are used. (Using the Stokes-Einstein formula for diffusion of colloids in liquids, one estimates that the time unit of the simulation corresponds to at least 1 sec of real time, illustrating once more that a truly molecular simulation, for which the time unit is on the picosecond scale, 10^{-12} s, is completely hopeless.) As an extreme case, one can turn off hydrodynamic interactions completely, and finds then that the growth proceeds considerably slower. Thus, all the factors controlling the domain growth can be "isolated" in our simulations, and hence their role elucidated in detail. Thus, we arrive at a more firm interpretation of the details of these complex processes than is possible in the corresponding experiments, where one always has to consider all factors influencing the domain growth together.

Of course, it is very satisfactory that there is a large similarity between the snapshot pictures recorded in experiment [1] and in our simulation (Fig. 2). But we can go further and enhance the resolution, so that we not only see coarse-grained domain patterns, but visualize individual particles (Fig. 3). In this way, it will become possible to study in more detail (i.e., on the scale of individual colloid particles) what precisely happens e.g. when two domains coalesce and the domain shape changes. Studying such phenomena in more detail is still a challenge for future work.

One aspect that has received considerable attention in the theoretical literature for decades [5] is the question whether in the late stages of domain growth a scaling regime emerges (i.e., the domain

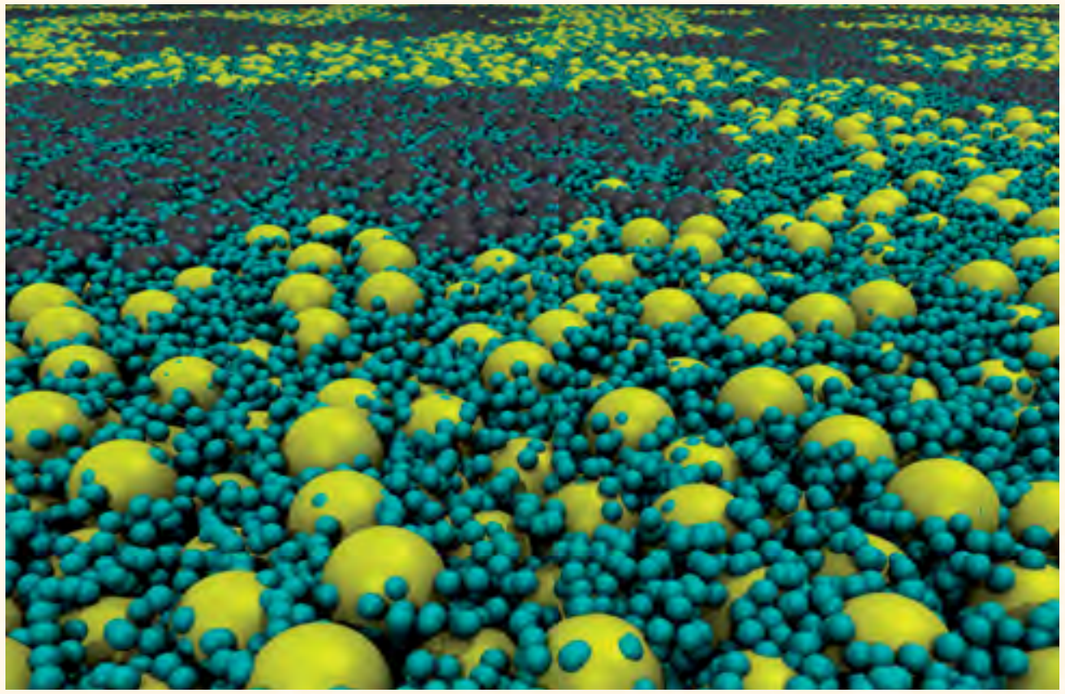

Figure 3: Snapshot of a section of the colloid-polymer mixture during the demixing process, visualizing the three different length scales: the point-like solvent particles (blue), the colloids and polymers (shown in yellow and grey, respectively), as well as the length scales of the polymer-rich and colloid-rich domains. The snapshot corresponds to an ultrathin film of D=1.5, deeply in the two-phase region. For a clearer view, a layer containing many solvent particles on the top is removed.

pattern at different times is "self-similar"), where the characteristic length scale grows as a power law of time: this shows up as straight line on a double-logarithmic plot. Fig. 4 shows that we indeed found some evidence for such a behavior, but the growth exponent does depend on the hydrodynamic boundary condition: for the slip case, the exponent is 2/3 as theoretically predicted [5], while for the other cases the standard Lifshitz-Slyozov-Wagner growth law [5] with exponent 1/3 is recovered.

In conclusions, our study has shed light on the role of hydrodynamic boundary conditions on domain growth in systems that undergo phase separation. We expect that this study will help to better interpret existing experiments, stimulate new experiments, and ultimately help to better control confinement effects on the processing of colloidal suspensions in various applications.

Acknowledgement

We thank the Deutsche Forschungsgemeinschaft (DFG) for support (grant No TR6/projects A4 and A5) and the Jülich Supercomputer Center (JSC/NIC), where this project was begun at the JUROPA supercomputer. We are particularly grateful to the Höchstleistungsrechenzentrum Stuttgart (HLRS) for a generous grant of computer time at the Hermit supercomputer.

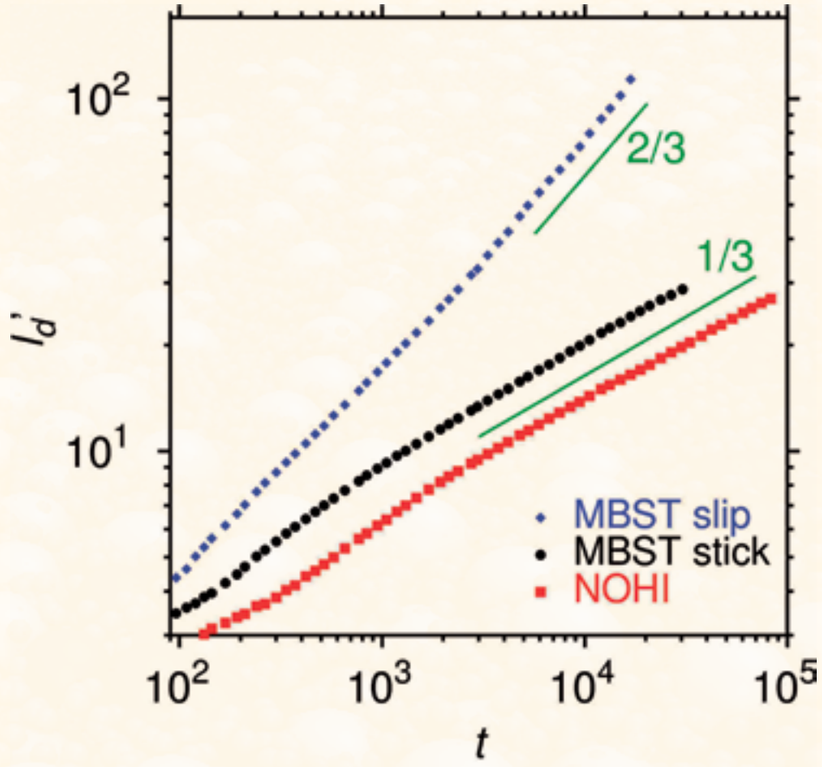

Figure 4: Log-log plot of the reduced domain size versus time, for a deep quench of an ultrathin film (D=1.5). From [3].

References

[1] **Jamie, E.A., Dullens, R.P., Aarts, D.G.**
J. Phys. Condens. Matter **24** (2012) 284120.

[2] **Gompper, G., Ihle, T., Knoll, D.M., Winkler R.G.**
Adv. Polymer Sci. 121 (2009) 1.

[3] **Winkler, A., Virnau, P., Binder, K., Winkler, R.G., Gompper, G.**
EPL **100** (2012) 16003.

[4] **Winkler, A.**
Dissertation, Johannes Gutenberg-Universität Mainz (2012).

[5] **Binder, K., Fratzl, P.**
„Spinodal decomposition", in Phase Transformations in Materials, edited by Kostorz, G. (Wiley-VCH, Weinheim, 2001) p.409.

- Alexander Winkler[1]
- Peter Virnau[1]
- Kurt Binder[1]
- Roland G. Winkler[2]
- Gerhard Gompper[2]

[1] Institut für Physik, Johannes Gutenberg-Universität Mainz

[2] Institute for Advanced Simulation & Institute for Complex Systems, Forschungszentrum Jülich

Structure and Dynamics at a Polymer-Solid Interface

Polymer based composite materials are indispensable in our modern world. We trust our life to them in high-technology applications like airplane wings and in mundane everyday technology like car tires. The aim behind mixing a hard solid and a polymer is always to improve selected properties, most often, and in the above mentioned applications, the mechanical properties like elasticity or yield stress. The composite material is meant to inherit the best of both worlds for its properties. But it not always does so, and if it does, why does it do so? The answer to this question has to lie in the properties of the interface between the solid and the polymer where both worlds meet. To improve our understanding of this interface is the aim of the focused research program SPP 1369, Polymer-Festkörper-Kontakte, funded by the German Science Foundation. The results here presented were obtained in one sub-project of this program and were made possible by a generous grant of computing time on the JUGENE computer of the Jülich Supercomputing Centre.

In engineering applications, these composite materials are typically described by a three-phase model, the polymer, the solid and a so-called interphase between the two with its own properties different from the other two phases. And it turns out that this interphase often has to be assumed to have an extension one to two orders of magnitude larger than the molecular constituents. When the polymeric material is composed of linear macromolecules, the typical extension of the random coil conformations these chains assume in a dense melt varies between a few nanometers (10-9 m) for short chains and about a hundred nanometers for ultra-high molecular weight polymers. To obtain a molecular understanding of the properties of polymer-solid composites, the structure of this interphase region from the size of the single atom to the size of the complete polymer has to be resolved and its dynamic behaviour, which underlies, e.g., the mechanical properties of this phase, has to be studied on all these scales as well.

There exists, of course, an uncountable variety of possible solid materials and polymer materials to mix into a composite, and the exact behaviour at the interface of the respective materials could well be influenced by the very specific chemistries meeting there. However, the large universality in the structural and dynamic behaviour of polymers encourages us to look into simplified model systems and extract from their behaviour qualitative insight which is independent of the specific model. When such a simplified model system can be studied in a chemically realistic way, we have the additional advantage that the computer simulation can predict the results of experiments performed on such a system and be validated this way. Both the airplane wing and the car tire have something in common: they employ a carbon based solid as filler for the polymer. Carbon is a highly versatile filler material, be it as carbon nanotubes or graphene sheets in the most advanced applications or as carbon black (soot) in a car tire. The polymer we studied is 1,4-polybutadiene (PB), one of the polymers making up the rubber material of car tires. This polymer had been studied by us before in its pure form, and a chemically realistic model capable of quantitative property predictions had been established. Our model system therefore consists of a melt of linear 1,4-polybutadiene chains confined by two graphite solids (cf Figure 1).

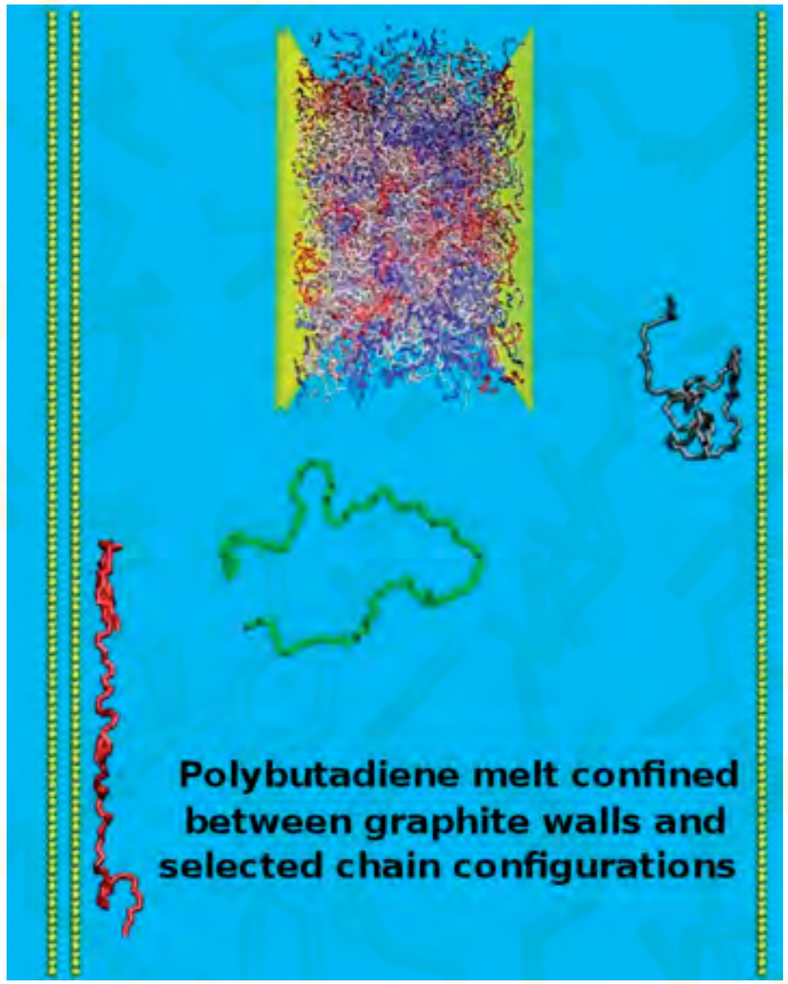

Figure 1: In the interior of the film the chains are assuming random coil conformations in the melt as exemplified by the green chain extracted from a simulation snapshot like the one shown in the inset. At the walls the chains can either adsorb lightly, like the gray chain on the right, or strongly, like the red chain on the left. This adsorption behaviour strongly influences the relaxation properties of the confined polymer film.

Both graphite and PB are apolar materials and inert, so we can perform a classical Molecular Dynamics (MD) simulation of an interacting N-particle system integrating Newton's equation of motion. When one simulates linear polymers, a choice of chain length, n, has to be made in addition. This choice is dictated by computational feasibility.

The longest relaxation time of a polymer chain behaves as $t_{max} = t_0 n^x$ where the segmental time scale t_0 increases from a few picoseconds at high temperature to 100 seconds at the glass transition temperature of the polymer (178 K for PB). The exponent x changes from x=2 for short chains below the so-called entanglement molecular weight to x~3 for long chains. For PB the entanglement molecular weight is around 30 repeat units, so we chose n=29 for our simulations. The integration time step in the MD algorithm, on the other hand, has to be around 1 femtosecond, sufficiently smaller than the time scales of the vibrational degrees of freedom in the model to be able to capture the dynamics on the atomic scale. Taking into account the extension of the chains and the expected extension of the interphase, we opted for simulation box sizes of 15 nm extension parallel to the walls and 10 nm or 20 nm extension between the walls, resulting in 1 - 2 10^5 particles in the simulation volume. The simulations for these system sizes were run in parallel with high efficiency on up to 2048 cores of the Juqueen (formerly Jugene) computer, which delivered a trajectory of 27 ns per day of CPU time for the smaller system. This allowed us to generate trajectories of up to one microsecond duration at lower temperatures. This limits the temperature range over which we can study our model system. We discuss these considerations in detail to illustrate that, especially when one has to study processes covering many decades in time, even the most modern supercomputers can not just be used for brute force solutions of the problem. The modeling has to be a compromise between what is desirable and what is doable. In our case this meant that we, for example, have no possibility to study the dynamics of long entangled chains at solid surfaces in a chemically realistic model and over a relevant temperature range.

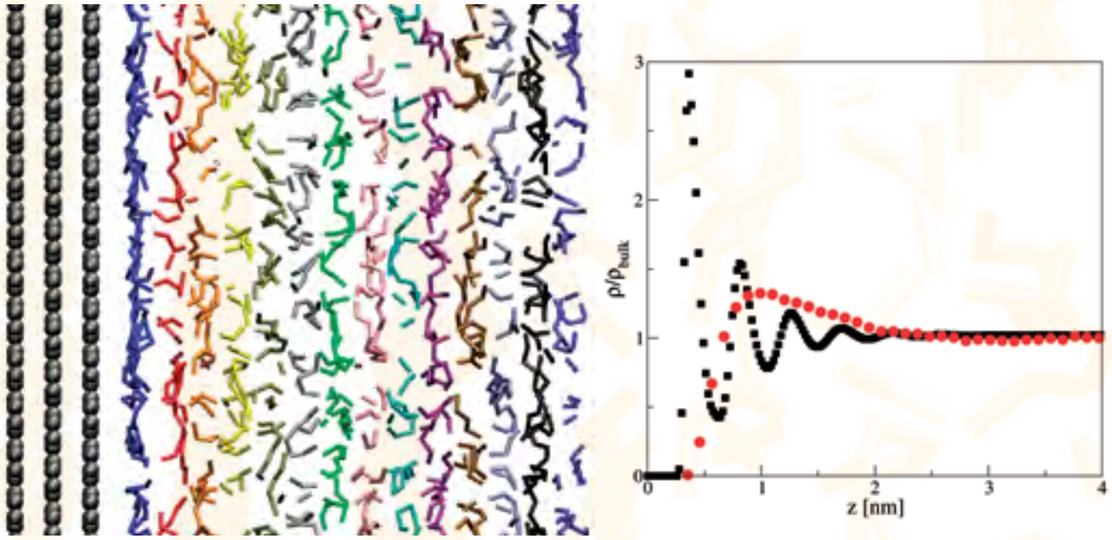

Figure 2: The left part shows a close-up of the distribution of chain segments next to one of the walls. The layering in the segmental distribution is clearly visible, highlighted here by the choice of different colors for different layers. This layering is shown in the right part by the black squares. This figure in addition shows the layering occurring in the chain centre of mass positions by the red circles. Both densities are normalized by the bulk value.

A liquid, like a polymer melt, confined between hard solid walls always exhibits a layering next to the walls as is visible in the configuration snapshot on the left side of Figure 2 and as can be observed in the right part of Figure 2 for the segmental density shown by black squares. The atoms in a solid wall always attract the atoms of the liquid with Van der Waals forces and this leads to a strong layering, as for our case of graphite and PB. This behaviour is well known and it is present in all possible polymer solid composites, so it is one of the universal qualitative features we mentioned above. But additionally, a polymer melt has another structural length scale, the size of the polymer coil as measured, e.g., by its radius of gyration, R_g. When we look at the probability to find the centre of mass of a polymer coil at a distance z from the wall, we observe again a layering phenomenon as shown by the red circles in the right part of Figure 2. Instantaneously, polymer coils in the melt assume the shape of a soap bar, so they have three symmetry axis of different lengths. In the bulk, their average projection on all directions is equal, but close to a wall, the longest axis will always be oriented parallel to the wall. This breaks the rotational symmetry of the coil structure leading also to the breaking of the translational symmetry observable as a layering on the coil scale (for this polymer R_g~1.5 nm in the bulk) in Figure 2. These structural changes in the polymer solid interphase extend roughly over two molecular sizes into the adjacent polymer melt. They are also underlying all effects of the solid walls onto the dynamics in the interphase and therefore on the property modifications occurring in the composite material.

The segmental dynamics as well as the chain dynamics in a polymer melt are accessible to a range of experimental techniques, the most important of these being neutron scattering, nuclear magnetic resonance (NMR) techniques and dielectric spectroscopy (for the chain dynamics this is only applicable for polymers having a net dipole moment along the end-to-end vector of the chain). All these techniques can in principal also be used to gain information on the dynamics in the interphase, however only with increased experimental effort. From our discussion of the melt structure in the interphase it is clear that the dynamics in this region will be heterogeneous, i.e., depending on the distance to the wall, as well as anisotropic, i.e., depending on the direction of motion. For neutron scattering techniques, this means that the experimental signal depends on the direction of the momentum transfer with respect to the wall, for NMR and dielectric experiments it means that the signal depends on the direction of the applied magnetic or electric field, respectively. So, experimentally, one has to realize a well controlled geometry. Furthermore, it would be very desirable to specifically address relaxation processes which occur at a given distance from the wall, which is much harder to realize, as all these techniques are typically obtaining their signals as averages over the complete sample.

In the simulation it is, of course, easy to analyze the dynamics as a function of distance to the wall and as a function of the orientation of applied fields, the direction of momentum transfer or the direction of motion of the segments. We have discussed our findings so far in relation to neutron scattering

Applications

experiments and NMR experiments in two publications [3,4]. Let us discuss the central physical result of these works looking at a quantity which is directly accessible in the simulation and which is a measure of the local directional mobility of the polymer segments – or the chains as a whole – the mean squared displacement (MSD). We will look into the MSD for the center of mass of chains in Figure 3 as its behaviour is a polymer specific result directly related to the layering in the centre of mass density shown in Figure 2.

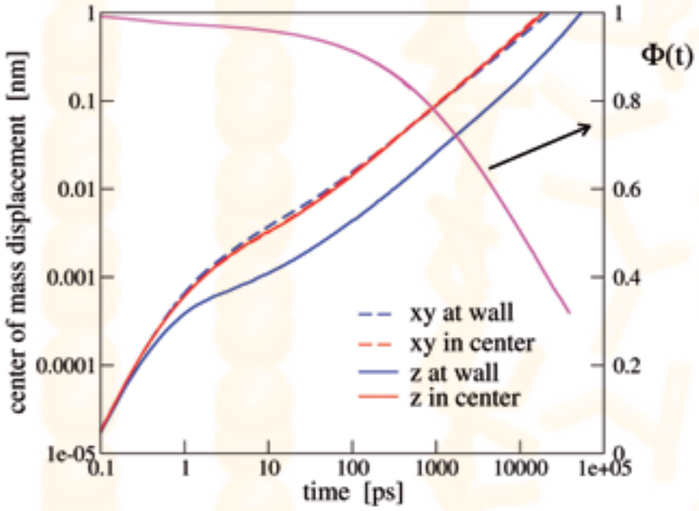

Figure 3: Mean squared centre of mass displacement of polymer chains (left axis) compared to an adsorption autocorrelation function (magenta line and right axis) defined in the text. The blue lines are the displacements for a chain in a layer of thickness R_g at the wall, the red lines are for a chain in the centre of the film. Dashed lines are displacements parallel, full lines are displacements perpendicular to the walls.

The first thing to note is the heterogeneity of the centre of mass MSD. In the centre of the film – shown by the red curves which are basically indistinguishable – it is much faster than close to the wall. For this figure we analyzed the displacements in a region of width R_g next to the wall and in the centre of the film, respectively. At around 1 picosecond, there occurs a crossover from a vibration dominated regime with a slope around two of the MSD as a function of time, to a subdiffusive regime which is typical for the Rouse motion in dense polymer melts of short chains. In short, the centre of the film behaves like the polymer bulk. This is very different at the wall. The motion parallel to the walls (shown by the dashed blue line) is very similar to the motion in the centre of the film, however, the motion perpendicular to the walls is slowed down strongly. The flat regime between the short time motion and the Rouse motion is very reminiscent of what happens at the glass transition in a polymer melt, however, the data shown are taken at T=353 K, twice the glass transition temperature of this material. A glass transition can be described as a time scale separation between vibrational processes and relaxation processes which in a polymer melt is induced by packing effects and the presence of barriers impeding conformational changes of the chains [2]. The slowing down observable in Figure 3 is the indication of such a time scale separation, but it cannot be simply ascribed to the increased density at the wall shown in Figure 2 as the motion parallel to the walls is not slowed down. The time scale separation is induced by a new relaxation process introduced into the confined melt by the presence of the walls. This process is the desorption kinetics of the chains, and the magenta line in Figure 3 is an autocorrelation function which gives the probability that a monomer adsorbed at time zero is still adsorbed at time t. Clearly, the decay of this function and the motion of the whole chain perpendicular to the wall occur on the same time scale.

This new relaxation process will influence different experiments in different ways. For instance, it had been found by neutron spin echo spectroscopy [5] that the motion parallel to the walls is basically unaltered in a tube confinement, whereas secondary ion mass spectrometry [6] showed that the motion perpendicular to the walls is strongly slowed down and that this influence extends far into the polymer melt from the walls. This is in direct agreement with our findings from simulation. But it will also influence different experimental techniques aimed to determine the glass transition temperature in a confined melt in different ways, and this may by one of the reasons why the conclusions about the dependence of the glass transition temperature on film thickness drawn from different experiments have been so contradictory for almost 20 years now. We hope that further detailed analysis of our MD trajectories revealing the molecular motions occurring in the polymer interphase will help resolve this long-standing controversy in the future.

References

[1] **S. G. Advani,**
Processing and Properties of Nanocomposites (World Scientific, Singapore, 2006)

[2] **W. Paul, G. D. Smith,**
Structure and dynamics of amorphous polymers: computer simulations compared to experiment and theory, Rep. Prog. Phys. **67**, 1117 (2004)

[3] **L. Yelash, P. Virnau, K. Binder, W. Paul,**
Slow process in confined polymer melts: Layer exchange dynamics at a polymer solid interface, Phys. Rev. E **82**, 050801 (2010)

[4] **L. Yelash, P. Virnau, K. Binder, W. Paul,**
Three-step decay of time correlations at polymer-solid interfaces, Europhys. Lett. **98**, 28006 (2012)

[5] **J. Martin, M. Krutyeva, M. Monkenbusch, A. Arbe, J. Allgeier, A. Radulescu, P. Falus, J. Maiz, C. Mijangos, J. Colmenero, D. Richter,**
Direct Observation of Confined Single Chain Dynamics by Neutron Scattering, Phys. Rev. Lett. **104**, 197801 (2010)

[6] **X. Zheng, M. H. Rafailovich, J. Sokolov, Y. Strzhemechny, S. A. Schwarz. B. B. Sauer, M. Rubinstein,**
Long-Range Effects on Polymer Diffusion Induced by a Bounding Interface, Phys. Rev. Lett. **79**, 241 (1997).

- Kurt Binder[1]
- Wolfgang Paul[2]
- Mathieu Solar[2]
- Peter Virnau[1]
- Leonid Yelash[1]

[1] Institut für Physik, Johannes Gutenberg-Universität Mainz

[2] Institut für Physik, Martin Luther Universität, Halle

Simulating the Evolution of Binary Stars in Stellar Clusters

Astronomical observations show that most stars are not isolated but are usually part of a pair – so-called binary stars or binaries. Being a pair, binary stars have additional properties compared to single stars. Two important such properties are the relative mass of the two stars and the time they need to revolve around each other, which is denoted as period. This period obviously depends on the distance between both stars.

One aim of binary observations is to understand how likely binaries with certain properties are in the total binary population. One surprising observational finding was that the properties in young binaries can differ significantly from those in older populations. In regions were stars just form more binaries with shorter and as well longer periods than in older populations are observed. This is illustrated in Fig. 1, where the young populations is indicated as 'initial' and the older population as 'final' (Refs. [1-3]). This means that over time a large proportion of the very tight and very wide binaries are destroyed by some mechanism.

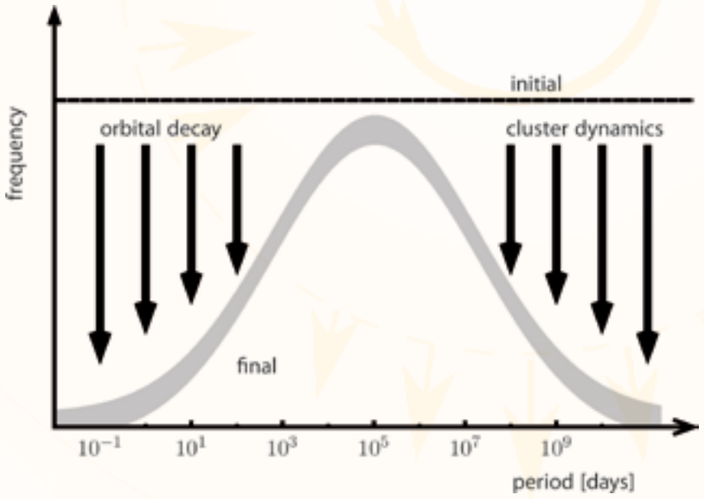

Figure 1: Orbital decay and cluster dynamics convert the initial binary period distribution observed in stellar clusters into the final binary period distribution observed in the field.

The young binary stars are preferentially found in bound groups (clusters) of hundreds to several thousands of stars. After the star formation process has finished these clusters dissolve and become part of the "field" population. Observations of the binary properties in young clusters and the field population give us in some sense the initial and final conditions for our simulations. So ideally the simulations should provide a picture how the binary population is born, changes in the cluster environment and dissolves into a field population.

We use the Orion Nebular Cluster (ONC) as a model star cluster to investigate the evolution of a binary population. The reason is that the ONC is one of the best observed clusters in the solar neighbourhood. So many parameters are well known for the ONC. At an age of one million years it is still relatively young and probably just before it starts to dissolve. It contains approximately 4000 stars which is populous enough to obtain reasonable statistics.

During the first million years of evolution, the cluster population is influenced by two important dynamical processes: First, by gas-induced orbital decay of embedded binary systems, and second, destruction of wide binaries in three-body interactions which will be detailed in the following.

Figure 1 shows the initial period distribution of the binaries and that of the field of a binary population. It can be seen that the gas-induced orbital decay only has an influence on the short-period binaries and that three-body interactions only destruct wide binaries.

The latter processes in stellar clusters have largely been studied (Ref. [4] - [6]) and are relatively well understood. Here the weakly bound wide binaries are separated by the interactions with stars that fly by. By contrast, the process of gas-induced orbital decay has been largely neglected in the past, but will be the focus here.

How does the gas-induced orbital decay work? Stars form from the gas and dust content of large molecular clouds. This natal gas largely stays around until it is expelled from the cluster by various processes. As a result a just formed binary system is still embedded in this natal gas for approximately one million years. The oscillating gravitational potential of the rotating binary system torques the nearby gas. This produces an outgoing wave in the gas, which again leads to a loss of angular momentum (Figure 2) in the binary system. As a consequence the distance between the stars shrinks. In extreme cases it is even possible that the two stars merge.

Steven Stahler (UC Berkeley) was the first to realize that the properties of young binaries might be affected by the surrounding gas. He investigated analytically how an isolated binary system might react to its natal gas (Ref. [8]). Due to the complex nature of the problem he chose a system with special features that are manageable analytically: Two equal-mass stars on a circular orbit with a distance of 100 Astronomical Units (1AU = distance between sun and earth).

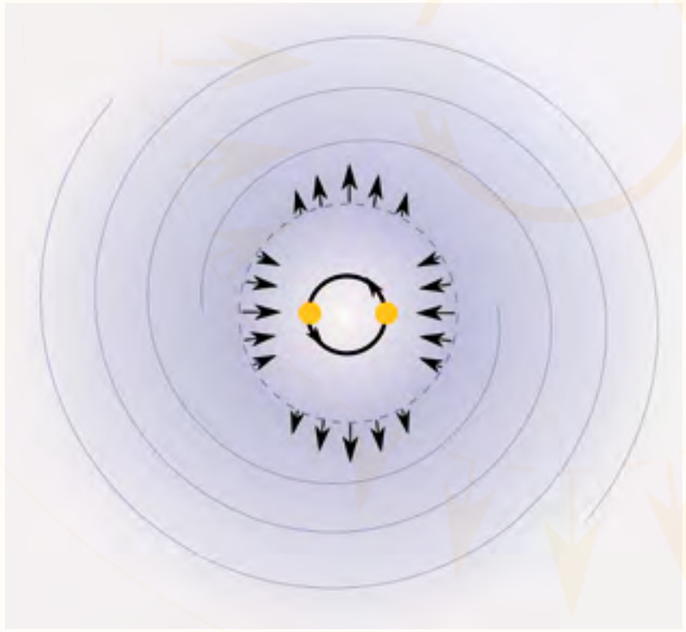

Figure 2: Orbiting binary system embedded in gas. The arrows show the oscillating gravitational quadrupole moment resulting in spiral waves as indicated.

We extended these investigations by simulating a whole binary population on JUROPA at the Jülich Supercomputing Centre. We demonstrated that in combination with dynamical effects this process converts the initial period distribution found in stellar clusters, into the final period distribution, found in the field (Figure 3, Ref [8]).

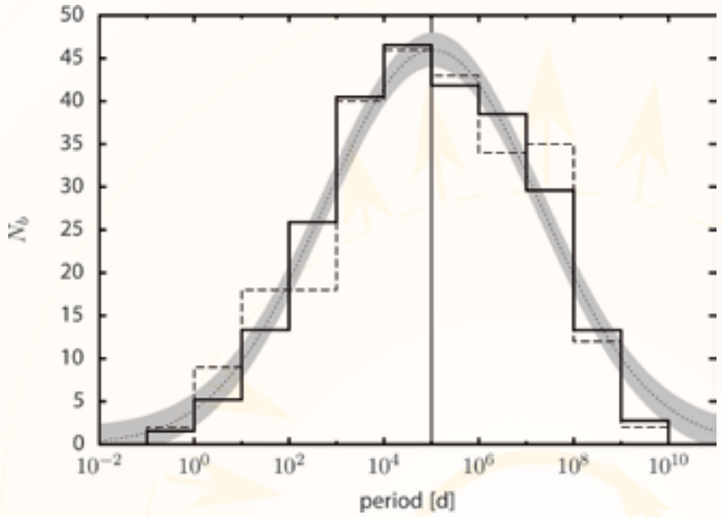

Figure 3: Final period distribution: Observations (solid line), simulation (dashed line), and fit to observations (grey)

As mentioned above, in extreme cases the shrinkage of the binary orbit can even lead to a merger of the two stars. So far it has been largely excluded that massive stars can

merge. The reason was that if one neglects the presence of gas, the only way that two stars merge would be a direct collision. However, that would require much higher stellar densities than usually observed. By contrast, our results suggested that the gas present in young clusters could work as an agent to make the merging of young binaries much more common than previously thought.

We recently developed a hydrodynamic model in order to extend the method from circular orbits to eccentric orbits. In contrast to the analytical calculations, the numerical models have to take special care when simulating the onset of the waves, as this is very sensitive in transporting angular momentum.

The relevant acoustic wave starts only at some distance from the binary [9, 10]. Figure 4 shows our simulations for three different binary systems. In the first row, a circular binary system with equal masses is shown (4.1a). The gravitational interaction with the gas results in waves which form two spiral arms (4.1b). Looking at the density perturbation for two angels separated by 90 degree, we see the expected period shifts by 180 degrees. In the binary system shown in the second row the stars are not of equal mass as before, but one star is twice as heavy as the other one. As a consequence, one of the spiral arms nearly vanishes (4.2b) and the period shifts only by 90 degrees. The binary in Figure 4.3 is on an elliptical orbit (eccentricity = 0.2) which results in an asymmetric acoustic wave, but as the two stars are of equal mass, the period shift remains at 180 degrees.

In the near future we will include the results of the hydrodynamic simulation into an N-body code to model the evolution of the stellar cluster consistently. These simulations will allow us to describe the influence of the cluster of stars and gas on the evolution of the binary population from their initial state to that of the field star population.

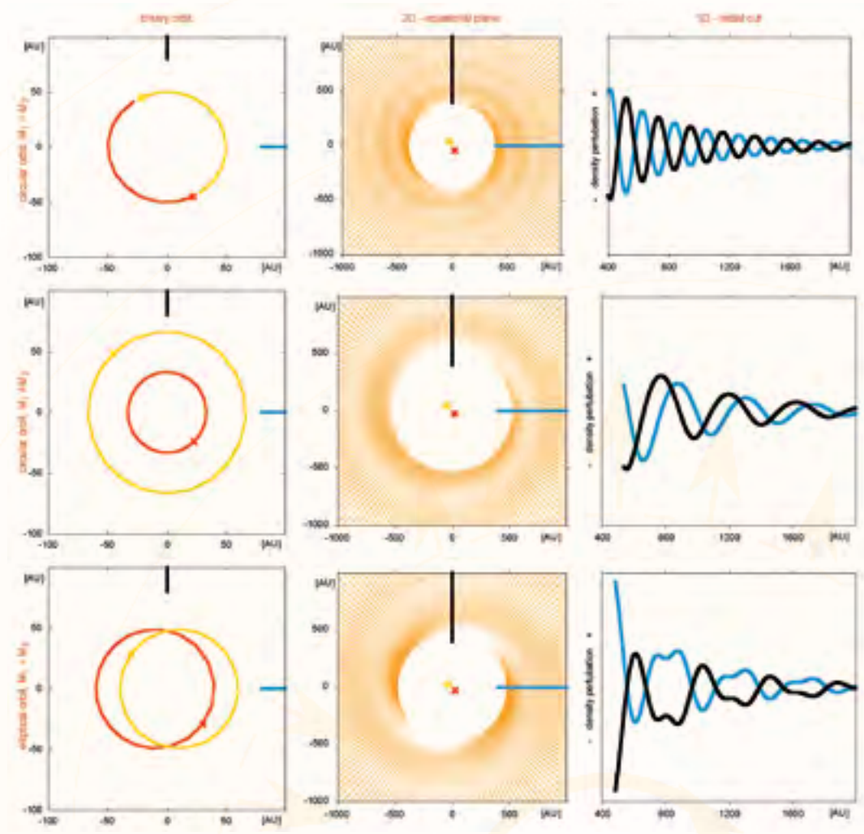

Figure 4: Hydrodynamic simulations of the gas surrounding a stellar binary system.

- Christina Korntreff[1]
- Susanne Pfalzner[2]

[1] JSC Forschungszentrum Jülich

[2] Max-Planck-Institut für Radioastronomie Bonn

References

[1] **Connelley, M. S., Reipurth, B., & Tokunaga, A. T.**
2008, AJ, 135, 2526

[2] **Duquennoy, A., & Mayor, M.**
1991, A&A, 248, 485

[3] **Raghavan, D., McAlister, H. A., Henry, T. J., et al.**
2010, ApJS, 190, 1

[4] **Kroupa, P., & Burkert,**
A. 2001, ApJ, 555, 945

[5] **Pfalzner, S., & Olczak,**
C. 2007, A&A, 475, 875

[6] **Kaczmarek, T., Olczak, C., & Pfalzner,**
S. 2011, A&A, 528, A144

[7] **Stahler, S. W.**
2010, MNRAS, 402, 1758

[8] **C. Korntreff, T. Kaczmarek, and S. Pfalzner**
2012, A&A 543, A126

[9] **Artymowicz, Pawel; Lubow, Stephen H.**
1994, ApJ,421, 651A

[10] **del Valle, Luciano; Escala, Andrés**
2012, ApJ, 761,31D

DNS of Gas Transfer at the Air-Water Interface

The present numerical work considers the problem of mass transfer across the air-water interface driven by two different mechanisms, namely by bottom-shear induced turbulent flow diffusing to the water surface and by buoyant-convective instability at the water surface. In the natural environment, an important example of such a process is the absorption of oxygen from the atmosphere into rivers and quiescent lakes, respectively. Oxygen and many other environmentally important gases have low solubility (high Schmidt number). The interfacial mass transfer of such low-diffusive substances is controlled by the hydrodynamic condition on the liquid side and marked by a thin boundary layer (typically 10 – 1000 µm). Obtaining detailed measurements data within such a thin layer is difficult and so has hindered the progress in gaining fundamental understanding of the problem.

With the present direct numerical simulations, we aim at revealing most details of the gas transfer process to allow a better understanding of the physical mechanisms controlling the process. We use an in-house numerical solver capable of resolving most details of the low-diffusivity scalar transport process [1]. In this code, a fifth-order accurate WENO-scheme [2] and a fourth-order central method were deployed for solving scalar convection and diffusion, respectively. The time-integration of the scalar was carried out using a three-stage Runge-Kutta scheme. For the fluid flow, a fourth-order accurate discretisation of the convection [3] and diffusion was combined with a second-order Adams-Bashforth method for the time-integration. The Poisson equation for the pressure was solved using a conjugate gradient solver with a simple diagonal preconditioning. For higher Schmidt number (Sc) cases, a dual-meshing strategy was employed in order to save computing time in which the scalar was solved on a finer mesh than the flow-field.

The setup of the computational domain was in accordance with the laboratory experiments performed at the Institute for Hydromechanics (IfH), Karlsruhe Institute of Technology (KIT), in which the interfacial gas transfer process was studied using PIV and LIF techniques [4].

Bottom-shear case

In the experiments, the turbulence was generated by an oscillating mesh in the lower part of the tank. To save computing time only a small part of the fluid domain near the surface was modelled and periodic boundary conditions were applied to account for the larger horizontal size of the domain. At the top, the free surface was modelled using a rigid lid approach with a symmetry boundary condition and at the bottom grid turbulence was introduced. The grid turbulence was generated in an LES that ran concurrently with the DNS. Each time step, a snapshot of the turbulence in a plane was extracted, then interpolated to the DNS grid and introduced at the bottom of the DNS. The transfer of the diffused substance into the liquid was forced by keeping the concentration C at the surface saturated, while at the bottom $\partial c/\partial z = 0$.

Fig. 1a shows a snapshot of the concentration isosurface of 50% saturation together with a cross-sectional plane showing the concentration field. Sequences of such images superimposed with their turbulent background flow-fields visualise nicely how eddy structures impinge on the surface from below. Downwelling motions of such eddies initiate the so called peeling process related to surface renewal events, sweeping part of the concentration boundary layer and transporting this oxygen-rich fluid down into the bulk. The upwelling motions, on the other hand, transport oxygen-poor fluid from the bulk up to the surface and at the same time cause a thinning of the boundary layer leading to a higher gas transfer rate. The transfer process was first simulated for Schmidt numbers ranging from 2 to 32 in 5L x 5L x 5L and 20L x 20L x 5L domains using 128 x 128 x 300 and 512 x 512 x 300 grid points respectively. In both simulations the mass boundary layer thickness was found to scale with $Sc^{-0.5}$.

Buoyant-convective case

Fig. 1b shows the results of the three-dimensional DNS run for the buoyant-convective (BC) case with Sc = 500. For this high Schmidt number case, the dual-meshing scheme was employed. The base mesh size was 400 x 400 x 256 grid points and a refinement grid factor of 2 was applied for solving the scalar field. Periodic boundary conditions on the sides and a rigid lid and free slip conditions at the top and bottom, respectively, were applied. In the experiments, the buoyant-convection process was initiated by introducing cold air above the water surface which in turn generated a cold thermal boundary layer on the water side. Similarly, in the present BC simulations, the temperature at the surface was set to be 3°C colder than that in the bulk and random disturbances of 1% were added at t=10 time-units to trigger the instability.

The isosurface in Fig. 1b represents a concentration value of 25% saturation. The simulations show similar flow structures and transport processes as visualized in the PIV-LIF experiments performed at IfH. Shortly after the instability was initiated, the cold upper layer started to sink down in the form of plumes with a typical mushroom shape. These falling plumes transport oxygen-rich fluid from the surface to the oxygen-poor region in the bulk.

On-going Research/Outlook

The first 3D-test case with a high (realistic) Schmidt number for the BC case (see above) shows a good qualitative agreement with experimental results. With the increased computational capacity of the latest generation of high performance computers, we should be able to increase the numerical resolution and perform the necessary grid refinement study by taking further advantage of the dual-mesh approach. Parallel to that, a simulation with Sc=500 for the bottom-shear (BS) case is currently being prepared to run on the new SuperMUC supercomputer. In the BS simulations of the lower Sc cases mentioned above the computational domain was subdivided into 200 to 400 blocks for parallelisation, each block was assigned to one processing core. Communication between processes happened through the standard Message Passing Interface (MPI) protocol. For the BS case with Sc=500 the dual-meshing capability of the code

will be employed by applying a refinement factor of 5. The simulation would require at least 32,000 blocks (owing also to the reduced memory capacity per core on the thin nodes of the new SuperMUC as compared to the SuperMIG migration system).

References

[1] **B. Kubrak, H. Herlina, F. Greve, J.G. Wissink.**
Low-diffusivity scalar transport using a WENO scheme and dual meshing. J. Comput. Phys., 240: 158-173, 2013.

[2] **X. Liu, S. Osher, and T. Chan.**
Weighted essentially non-oscillatory schemes. J. Comput. Phys., 115:200–212, 1994.

[3] **J.G. Wissink.**
On unconditional conservation of kinetic energy by finite-difference discretisations of the linear and non-linear convection equation, Computers and Fluids, 33, 315-343, 2004.

[4] **G.H. Jirka, H. Herlina and A. Niepelt.**
Gas transfer at the air-water interface: experiments with different turbulence forcing mechanisms, Exps Fluids, 49, 319-327, 2010.

- H. Herlina[1]
- J.G. Wissink[2]
- B. Kubrak[2]

[1] Institute for Hydromechanics, Karlsruhe Institute of Technology

[2] School of Engineering and Design, Brunel University London

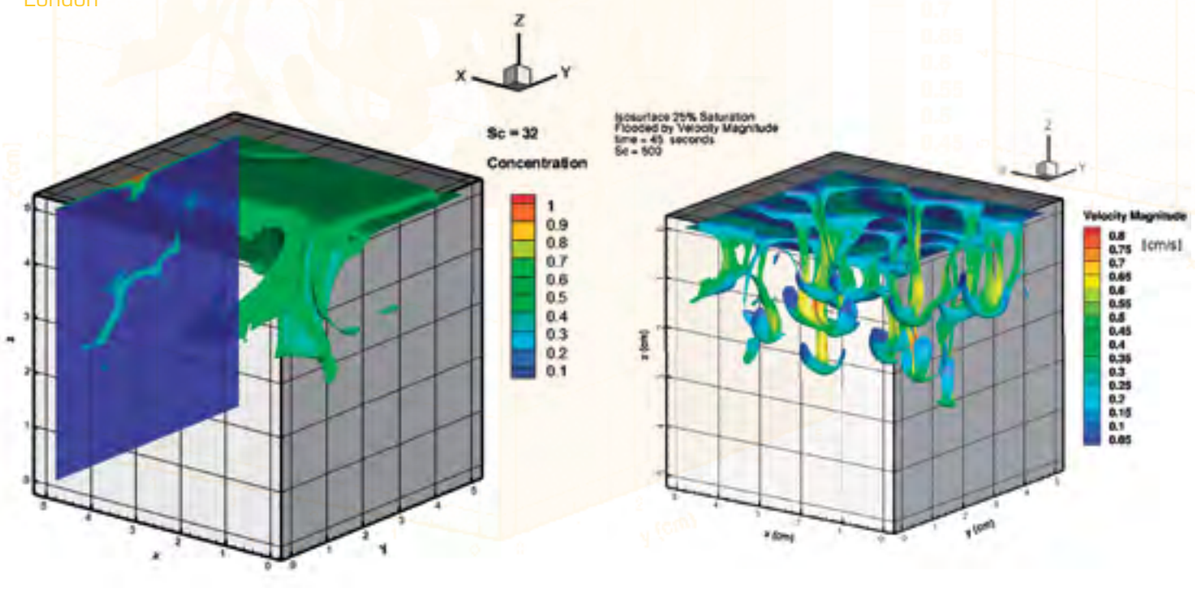

(a) bottom-shear case (BS) (b) buoyant-convective case (BC)

Figure 1: Isosurface of concentration. a) for the bottom-shear case with Sc=32 and b)for the buoyant-convective case with Sc=500. Both domains had a uniform grid size in the x and y directions and a slightly stretched grid distribution in the vertical direction to achieve a denser grid near the interface.

Dynamics and interaction of membrane-spanning protein helices in health and disease

Proteins and Lipid Membranes

Proteins perform countless functions thus enabling living cells to maintain their metabolism, to grow, to communicate with other cells etc. About 25% of all proteins in any organism are embedded within lipid membranes. Membranes define and structure individual cells and are built from a great variety of different lipid molecules. The task of membrane proteins is to mediate regulated fusion of cells, allow for solute exchange with the cell exterior, regulate release of cellular products, mediate lipid translocation from one side of the membrane to the other one, cleave other membrane proteins etc. Most membrane proteins are anchored by helical transmembrane domains (TMDs, Fig. 1). Long regarded as dull membrane anchors, TMDs are increasingly recognized to guide assembly of subunits to larger functional complexes and to be dynamic [1]. Molecular Dynamics (MD) simulations done at single-atom resolution are well-suited to investigate molecular properties of protein helices, surrounding lipids and solvent by computing time series of structures. From these ensembles of structures conformational and dynamical properties are evaluated and related to the underlying molecular interaction forces. For the production runs (typically 100,000 atoms for > 150 ns) we use the description of atomic interactions given by the CHARMM forcefield [2,3] and the MD algorithm as implemented in NAMD [4]. Each simulation project compares a large set of TMDs with different amino acid compositions and sequences in isotropic solvent as well as in membranes of different lipid composition. Validation of the simulations by experimental measurements connects models of exquisite structural detail to physical reality.

One production run uses up to 120 cores at the fat node island of SuperMUC and consumes approximately 100,000 hours of CPU time (three months wall-clock time). In contrast to the LRZ Ultra-Violett Linux cluster and the HLRB II system, the short queue times at the SuperMUC HPC system allowed us to run the simulations in a relative short time. Due to the large number of cores per machine and the high amount of latent memory, the fat nodes are quite perfect for NAMD with its high number of I/O operations and CPU communications, highly depending on low interconnection latencies. First try runs with the Sandy Bridge architecture of SuperMUC's thin nodes showed very similar performance which will allow us to extend our simulations to larger membrane patches and a higher number of embedded transmembrane helices which require more than 1,000 cores.

TMD Helix-Helix Interactions can Support Protein/Protein Interactions

Most membrane proteins form multisubunit complexes and an hitherto unknown number of them do so as a

result of specific TMD helix-helix interactions. MD simulations can provide models of the interfaces and the physical forces that govern the interaction. In one example, we have modeled the sequence-specific interaction of the TMD of sulphhydryl oxidase, a membrane protein that aids folding of other proteins. Experimental work had identified amino acids that are responsible for mutual recognition of these helices and computational work resulted in a homodimeric structure that matches the experimental results (Fig. 2).

Structural Dynamics Supports Protein Function

While the 3-dimensional structures of tens of thousands of different water-soluble proteins have been determined within the last decades, it is becoming increasingly clear that the functional cycles of many proteins require conformational transitions between different structural states, i.e., structural dynamics. Likewisely, membrane proteins are increasingly recognized as dynamic entities. TMDs are particularly important for the function often being located at the structural interface between ligand-sensing extracellular and intracellular signaling domains. TM helices can change their orientations relative to each other on µs time scales; in addition, the helices themselves are intrinsically dynamic structures in a sense that they exhibit local dynamics (like transiently hydrogen-bond opening and side-chain rotations) as well as global motions (like bending, twisting and elongation) within ns (Fig. 3). TMD interaction and dynamics also appear to affect their ability to serve as substrates for intramembrane proteases. Specifically, both factors have been suggested to define the specificity by which the TMD of amyloid precursor protein is cleaved to fragments eliciting Alzheimer's disease.

Transmembrane Helix Dynamics Depends on Amino Acid Sequence

While there is ample evidence for the functional significance of the slow (µs) relative movements of TMDs, only little is known about the role of fast (ns) helix dynamics. MD simulations provide rich molecular detail at the ns time scale and are thus employed by us to study local TM helix dynamics. In a systematic analysis, we have designed a large set of low-complexity model TMDs de novo and modeled their dynamics in isotropic solvent. Interestingly, the results show that the structural fluctuations along the helix strongly depend on their sequences (Fig. 4) Thus, helix dynamics is more specific than just Brownian motion; rather, its sequence-dependence shows that it is a structural property that can be optimized by successive mutations in evolution and thus support the functional diversification of membrane proteins. Our results show [5] that the backbone dynamics i) increases at the termini of the helices and ii) depends on the primary structure. The dependence on sequence is particularly important since the helix-destabilizing amino acid types found here have previously been found to be overrepresented in the TMDs of membrane fusogenic proteins and isolated TMDs are known to induce membrane fusion in vitro. Further, this work yielded mechanistic insights into the molecular determinants of backbone dynamics. Specifically, we found that the nature of the amino acid side chains affects the flexibility of the helix. In essence, those side chains form a molecular 'scaffold' around the central invariant backbone of a helix and the stability of that scaffold is determined by side-chain size, structure and dynamics which determine the mutal interactions.

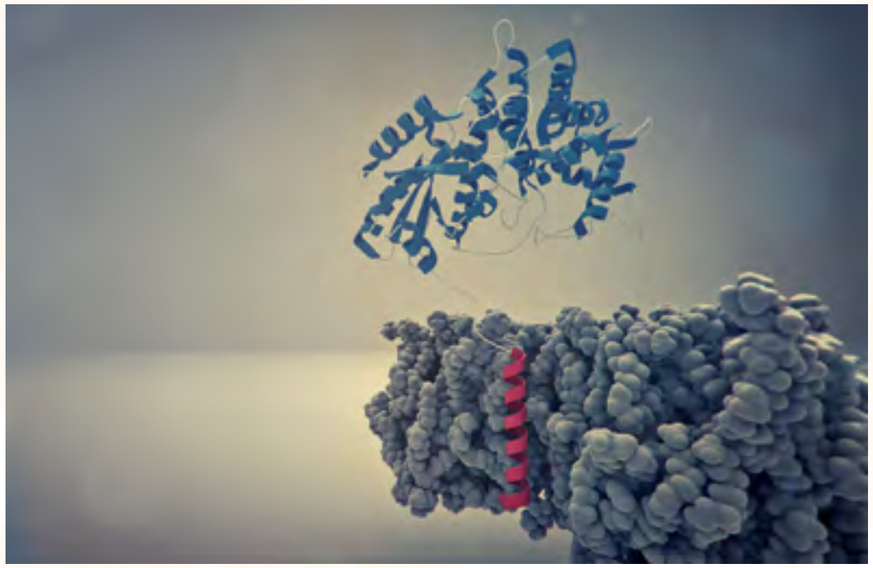

Figure 1: A pictorial view on a membrane protein anchored with a helical transmembrane domain (red) in a lipid patch (grey, only shown partially).

These MD calculations closely fit experimental work (CD spectroscopy, hydrogen/deuterium-exchange kinetics). While the overall time scale for isotope exchange is determined by the slow chemical exchange (time scale of hours), the variations along the sequence depend on local dynamics which is reliably sampled within the 100 ns MD trajectories. Therefore the results from MD could be used to calculate and predict experimental data

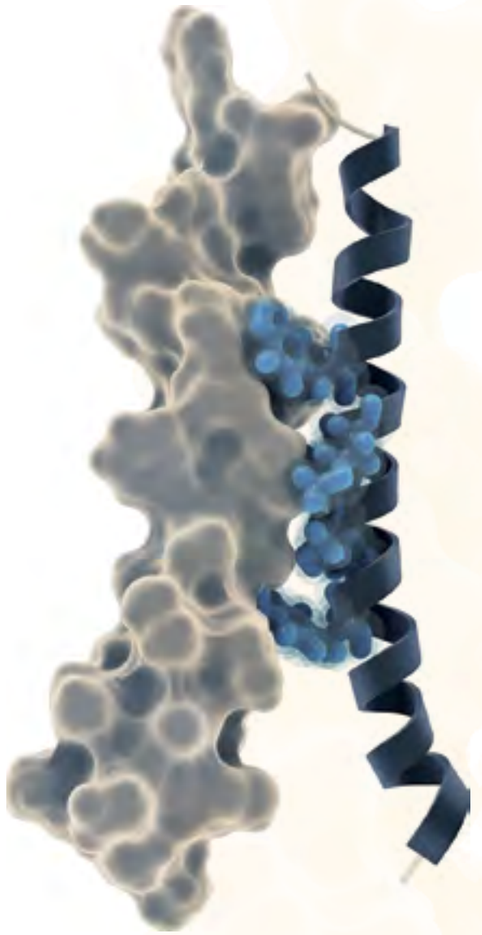

Figure 2: Transmembrane helix-helix interactions stabilize the dimer of sulphhydryl oxidase. The residues of one of the two helices responsible for mutual recognition are shown as blue spheres. Transparent shapes around the atoms represent their van der Waals radii. The second helix is represented by its molecular surface (light grey). The C-terminus of the helices is at the top.

recorded on a much longer time scale with excellent accuracy (Fig. 5).

Helix Dynamics in Lipid Membranes

Another question relates to the potential impact of TMDs on the structure of the membrane and vice versa. Simulating our model TMDs in lipid bilayer patches indicate that the dynamics of membrane-embedded helices are strongly reduced by the membrane. This is expected due to the absence of destabilizing water from the membrane interior. Further, the results nicely demonstrate specific molecular bonding between protein and lipids such that local protein/lipid networks are generated. Further, insertion of the TMDs has profound consequences for the structure of the membrane. Membranes are highly fluid as their constituent lipids are quite flexible themselves. Indeed, the long-range order, i.e., the way lipids are organized, changes drastically around a TMD (Fig. 6). Changes in membrane order underly complex biological processes, such as membrane fusion and lipid exchange. This work is thus expected to shed light on the mechanisms by which membrane proteins catalyze these functions.

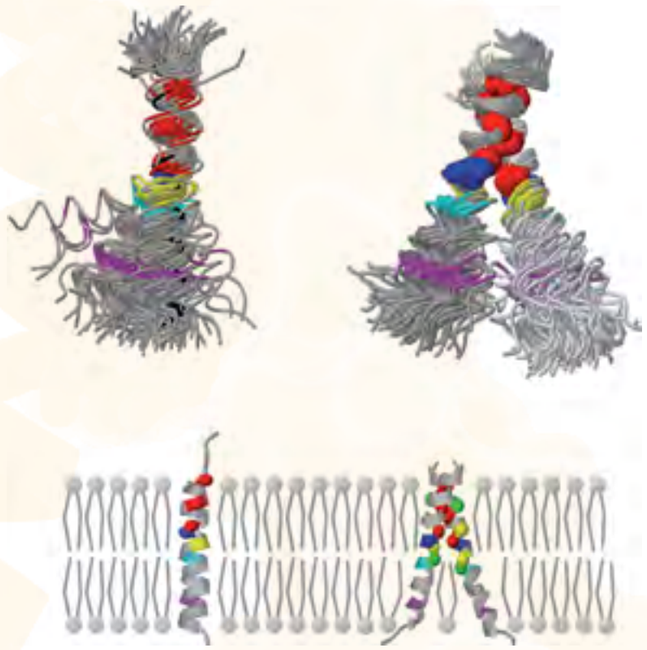

Figure 3: Patch of a plasma membrane and magnifications illustrating backbone dynamics and interactions of monomeric and dimeric transmembrane helices.

Helix Dynamics in Alzheimer's Disease

Alzheimer's disease affects currently ~40 million people on earth and is expected to increase dramatically as a result of increasing life-spans. Although the etiology of the disease is still controversial, it is clear that cleavage of the Amyloid precursor protein (APP) TMD by γ-secretase is an essential step leading to it. It is believed that this TMD forms a dimer which is cleaved at various sites. Successive cleavage leads to the build-up of Aβ peptides in the brain that ultimately kill nerve cells. The toxicity of these peptides depends on their size and size depends on the site of cleavage. Thus, understanding how individual cleavage events are influenced by the structure and dynamics of the APP TMD helix will provide a rationale for understanding a key step in this disease. Simulations of the APP TMD revealed a number of surprising results [6]. For one thing, that part of the helix that engages in forming the dimer is much more dynamic than the actual cleavage domain. Possibly, a highly flexible dimerization domain facilitates movement of the helix within γ-secretase or the final release of cleavage products. The nature of cleavage products, i.e., their length, may be determined by the relative helix dynamics at the site where cleavage begins. Indeed, we find that a number of parameters describing local helix

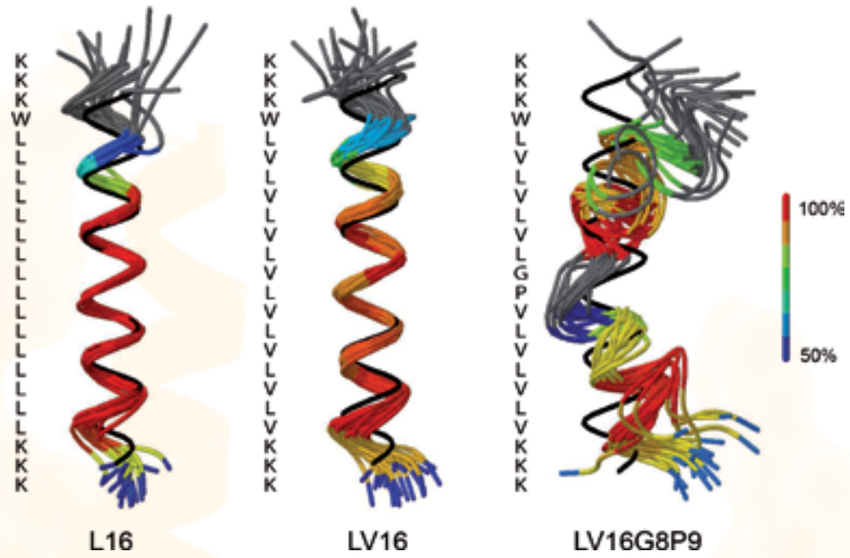

Figure 4: Graphical representation of backbone dynamics for a series of TMD model peptides with an increasing content of helix-destabilizing residues. 20 backbone conformations taken every 5 ns are overlayed with a rigid-body fit. The fraction of intact intrahelical hydrogen-bonds is color coded from red (high hydrogen-bond occupancy, rigid backbone) to blue (high fraction of local unfolding). Residues which cannot form hydrogen-bonds are shown in grey.

unfolding point to enhanced dynamics, viz. cleavage, at the site leading to the less toxic product (Aβ 40) compared to the site prompting the slightly longer and more toxic Aβ 42. Comparing monomeric and dimeric helices reveals another interesting feature. While dimerization rigidifies the dimerization region, as expected, it destabilizes the cleavage domain. A more detailed analysis of the trajectories suggests that this destabilization is due to increased solvation within the dimer. Thus, dimerization could actually promote initial cleavage. Contrary to expectation, the dynamics of the natural substrate, the APP TMD, is similar to that of several non-substrates. However, we find that the APP TMD dynamics is unleashed by certain amino acid exchanges. Some of these exchanges are also known to change the cleavage pattern. Taken together, it appears as if local changes in backbone dynamics can have a pronounced impact on the efficiency of the cleavage reactions. There are several hereditary forms of Alzheimer's disease that are caused by individual amino acid exchanges within the APP TMD and that are characterized by profoundly changed ratios of the Aβ peptides. Elucidating whether such exchanges exert their effects via changing helix dynamics will be a rewarding practical application of this work.

References

[1] Langosch, D., and Arkin, I.T. (2009). Interaction and Conformational Dynamics of Membrane-Spanning Protein Helices. Protein Sci. 18, 1343-1358.

[2] MacKerell, Jr., A. D., Bashford, D., Bellott, M., Dunbrack Jr., R.L., Evanseck, J.D., Field, M.J., Fischer, S., Gao, J., Guo, H., Ha, S., Joseph-McCarthy, D., Kuchnir, L., Kuczera, K., Lau, F.T.K., Mattos, C., Michnick, S., Ngo, T., Nguyen, D.T., Prodhom, B., Reiher, III, W.E., Roux, B., Schlenkrich, M., Smith, J.C., Stote, R., Straub, J., Watanabe, M., Wiorkiewicz-Kuczera, J., Yin, D., and Karplus, M. (1998) All-atom

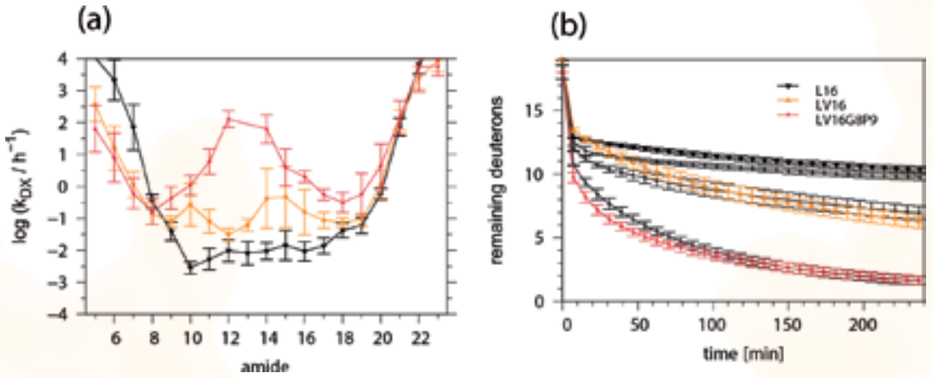

Figure 5: Site-specific helix dynamics for TMD model peptides with increasing content of helix-destabilizing residues (for sequences see Fig. 4) (a) Local deuterium-hydrogen exchange rates calculated from MD. Exchange competence is determined by the fraction of open intrahelical hydrogen-bonds and the presence of water molecules in close contact to the backbone amide groups. Both properties can be reliably retrieved from MD, as is demonstrated by the close correspondence between experimental and calculated exchange kinetics in (b). Experimental values are shown in dark grey, MD-derived values are shown in colors.

- Dieter Langosch[1]
- Matthias Mörch[1]
- Christina Scharnagl[2]

[1] Technische Universität München, Lehrstuhl für Chemie der Biopolymere

[2] Technische Universität München, Fakultät für Physik

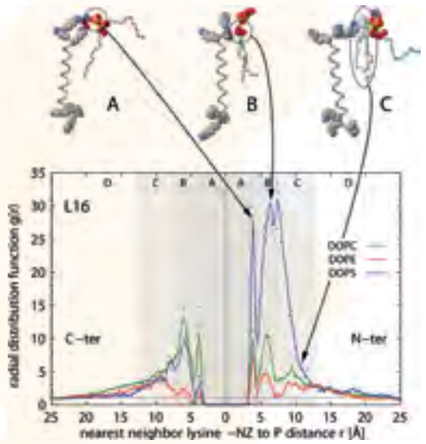

Figure 6: The local lipid composition around a transmembrane helix differs drastically from the bulk values and depends on the orientation of the TMD (N-terminal, C-terminal) as well as on the lipid type (DOPC, DOPS, DOPE). Interaction modes are exemplified by lipids complexed with the charged residues (lysine, grey and blue spheres) flanking the hydrophobic core (grey trace) of oligo-leucine L16: (A) hydrogen-bonding to the lipid phosphate group (orange), (B) hydrogen-bonding to lipid carbonyl groups (red), (C) indirect interaction of lipids in the second shell.

Empirical Potential for Molecular Modeling and Dynamics Studies of Proteins. J. Phys. Chem. B 102, 3586-3616.

[3] Feller, S., and MacKerell, Jr., A.D. (2000)
An Improved Empirical Potential Energy Function for Molecular Simulations of Phospholipids, J. Phys. Chem. B 104, 7510-7515.

[4] Phillips, J.C., Braun, R., Wang, W., Gumbart, J., Tajkhorshid, E., Villa, E., Chipot, C., Skeel, R.D., Kalé, L., and Schulten, K. (2005).
Scalable Molecular Dynamics with NAMD. J. Comp. Chem., 1781-1802.

[5] Quint, S., Widmaier, S., Minde, D., Langosch, D., and Scharnagl, C. (2010).
Residue-Specific Side-Chain Packing Determines Backbone Dynamics of Transmembrane Model Helices. Biophys. J. 99, 2541-2549.

[6] Pester, O., Barret, P., Hornburg, D., Hornburg, P., Pröbstle, R., Widmaier, S., Kutzner, C., Dürrbaum, M., Kapurniotu, A., Sanders, C.R., Scharnagl, C., and Langosch, D. (2013).
The Backbone Dynamics of the Amyloid Precursor Protein Transmembrane Helix Provides a Rationale for the Sequential Cleavage Mechanism of γ-Secretase. J. Am. Chem. Soc. 135, 1317-1329.

Simulation of 3-D seismic wave propagation in a synthetic Earth

Long-standing questions in the study of Earth's deep interior are about the origin of seismic heterogeneity and the nature of flow in the mantle. Understanding the dynamic behaviour is important as the flow drives plate tectonics and controls the way the Earth looses its heat. Thus, it is a crucial factor in tectonic modelling or in simulations of the geodynamo and the thermal evolution of the Earth. A better understanding of these aspects is of great societal importance. For example, the continuous drift of tectonic plates relative to each other results in a build up of stress at the plate boundaries. This stress can eventually exceed the yield stress of rock thus leading to (often disastrous) earthquakes.

Most of our knowledge on deep Earth structure and flow comes from the analysis of recordings of seismic waves that travel through the Earth after large earthquakes. Similar to medical tomography, seismic tomography allows us to "see" the present day elastic structure of Earth's mantle in 3-D. During the last two decades, a variety of such 3-D models were built from different datasets. However, seismic tomography can only provide a limited resolution; that is, only a blurred and distorted image of Earth's structure can be obtained. This is a consequence of errors in the data and the non-uniform distribution of earthquakes as well as seismic receiver locations, which mainly cover continental regions. A second problem inherent to any tomographic inversion is that an infinite number of models will fit the data equally well, a problem known as non-uniqueness of the solution.

Geochemical observations also offer important constraints on Earth's structure and evolution, providing compelling evidence for compositional heterogeneity within the mantle. Geophysical techniques, including seismology, can provide detailed information on such heterogeneity. However, despite substantial improvements in seismological methods and a dramatic increase in the amount and quality of available data, interpretations of images from seismic studies remain hampered by trade-offs between thermal and chemical effects. As a consequence, the distribution of chemical heterogeneity within Earth's mantle and, accordingly, its role in governing mantle dynamics, remains poorly understood.

Imaging Earth's interior structure at higher resolution and the convergence towards consistent models serving as a 3-D seismic reference is currently one of the most pressing challenges of modern global seismology today. So far, tomographic studies only used the arrival times of a very limited set of seismic phases from all of the information that is contained in a seismogram. In addition, most studies were based on ray theory, thus completely neglecting the wave character of seismic disturbances.

In order to improve conceptual models of mantle flow, the major challenges in seismology today are to efficiently mine the wealth of information contained in seismic waveforms and to constrain the relative contributions of thermal anomalies and compositional variations to the observed seismic heterogeneity.

High expectations to gain new insight currently lie within numerical simulations of wave propagation through complex three-dimensional structures.

Modern computational tools for seismic wave propagation incorporate a large range of physical phenomena and are able to produce synthetic datasets that show a complexity comparable to real observations. Also, computing whole waveform synthetic seismograms at relevant frequencies became feasible on a routine basis in recent years thanks to rapidly growing computational resources. However, it has long been not clear how to introduce geodynamic considerations into seismological forward simulations in an efficient and consistent manner, and how to benefit from expensive large-scale simulations for our understanding of deep Earth structure and dynamics. This was the motivation to develop a novel method, in which we generate synthetic 3-D mantle structures based on dynamic flow calculations that serve as input models in the simulation of seismic wave propagation.

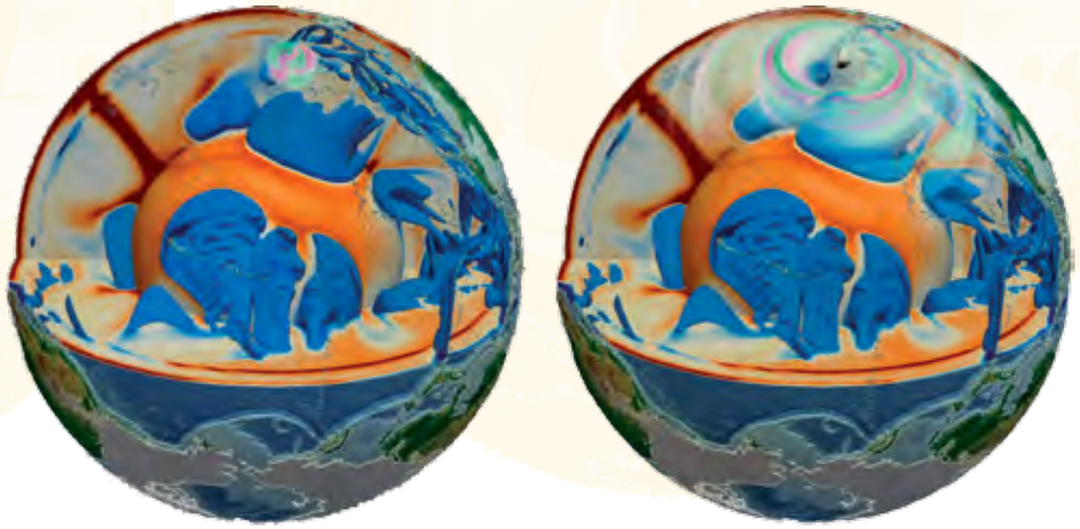

Figure 1: Snapshots of the three-dimensional wavefield in one of our geodynamic models. 3-D global wave propagation was simulated for an earthquake in the Fiji Islands region using a spectral element technique. The wavefield is depicted by green and magenta colours together with the shear wave velocity variations in the model, for which vertical cross-sections and iso-surfaces are shown on a blue to brownish colour scale ranging from -2 to 2 per cent. Surface topography is also shown for parts of the globe for geographic reference [1]

Here, we present the results of our new multi-disciplinary approach that combines forward modelling techniques from geodynamics, mineral physics and seismology. The thermal state of Earth's mantle at present-day geologic time is predicted by 3-D high-resolution mantle circulation models using a finite-element method. The temperature field is then mapped to seismic velocities. For this task, we take advantage of recent progress in describing the state of dynamic Earth models in terms of elastic properties through thermodynamically self-consistent models of mantle mineralogy. The predicted models of seismic heterogeneity are then implemented in a spectral element code for the simulation of 3-D global wave propagation. Using state-of-the-art techniques to solve the wave equation in 3-D heterogeneous media, this approach allows us to capture the correct physics of wave propagation.

Both, the geodynamic as well as the seismic simulations require large-scale high-performance calculations. The computational resources provided through the supercomputing platform HLRB II at the Leibniz Supercomputing Centre (LRZ) allowed for the first time to simulate seismic wave propagation in synthetic Earths; that is, we are now able to compute synthetic seismograms completely independent of seismic observations. This means that we can test geodynamic hypotheses directly against seismic observations, which may serve as a complementary tool to tomographic inversions. More specifically, it is for the first time possible to study frequency-dependent waveform effects, such as wavefront healing and focusing/defocusing in mantle structures with realistic length-scales; that is, in a physically consistent manner.

Results

One specific question that we addressed with our joint forward modelling approach is the origin of two large regions of strongly reduced seismic velocities in the lowermost mantle (the so-called African and Pacific superplumes). Several seismological observations are typically taken as an indication that the superplumes are being caused by compositional variations and that they are “piles” of material with higher density than normal mantle rock. However, a large number of recent geodynamic, mineralogical and also seismological studies argue for a strong thermal gradient across the core-mantle boundary (CMB) that might provide an alternative explanation through the resulting large lateral temperature variations. We tested the hypothesis whether the presence of such a strong thermal gradient in isochemical whole mantle flow is compatible with geophysical observations.

We have computed synthetic seismograms and a corresponding dataset of traveltime variations for one of our synthetic mantle models assuming a pyrolite composition for the mantle mineralogy. Synthetic seismograms for periods down to 10 seconds were computed for the predicted structure using 486 cores on HLRB II. Altogether, we simulated the seismic wavefield for 17 earthquakes distributed evenly over the globe. The wavefield of each earthquake was "recorded" by a very large number of virtual seismic stations in order to achieve a relatively homogeneous illumination of our model even with a low number of seismic sources. In total, we obtained each ~350,000 cross-correlation traveltime measurements at a dominant period of 15 s for compressional (P) and shear (S) waves.

The traveltimes of observed seismic waves, to which we compared our synthetic dataset, show a peculiar behaviour of their statistics as a function of the turning depth of the waves in the mantle: The standard deviation of P-wave traveltime variations stays almost constant with depth, while that of the S-wave traveltimes increases strongly towards the CMB. This increase in case of S-waves is particularly strong below a depth of around 2000 km (cf. Fig. 2). Such a difference between P- and S-waves is not expected in a chemically homogeneous mantle in the framework of ray-theory, which forms the backbone of seismology.

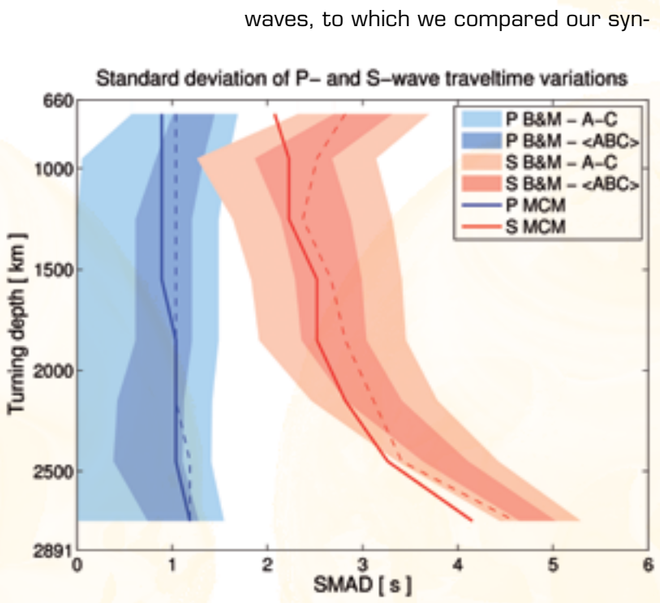

Figure 2: Comparison of the standard deviation (SMAD = scaled median average deviation) of traveltime variations in our geodynamic model to that of the observations. Intermediate and light shaded areas show the range of values inferred from the data [2]. Blue lines: simulated P-wave traveltime variations. Red lines: same for S-waves. Solid and dashed lines show SMAD curves for two different measurement techniques [1].

Using full wavefield simulations, we find, however, that the standard deviations of P- and S-wave traveltime variations in our geodynamic model show the same peculiar behaviour as the observations. This is a remarkable result in light of the isochemical nature of our mantle circulation model and highlights the importance of taking the correct physics of wave propagation into account in the interpretation of long-period seismic data. Most important, our comparison shows that isochemical whole mantle flow with strong core heating and a pyrolite composition is capable of explaining the statistics of seismic observations. The standard deviations of our synthetic P- and S-wave traveltimes do not only show different trends with depth, but are also matching those of the observations well in terms of their magnitude. While this finding does not necessarily mean that there is no chemical heterogeneity present in the lower mantle, it shows that complex large-scale variations in chemical composition are not required by the dataset studied here.

On-going Research / Outlook

One particular aspect that we currently strive to clarify in the elastic representation of geodynamic heterogeneity is the suspected existence of a mineral phase named post-perovskite (the high pressure polymorph of magnesium silicate MgSiO3). The elastic properties of the mantle will be affected by the existence of this phase, which would cause strong shear wave velocity variations across the phase transition while the velocity of compressional waves would change only slightly. To better constrain the relative importance of thermal anomalies and variations in chemical composition for generating seismic heterogeneity, the comparison presented here needs to be repeated with post-perovskite included in the mineralogical conversion.

Seismic datasets are rapidly growing not only due to an increasing number of seismic stations, but also due to the fact that traveltime measurements are now starting to be done at multiple frequencies. In this study, we have focused on a single frequency band. In future, it will become important to understand also the multi-frequency datasets from a forward modelling perspective complementary to using them for tomographic inversions. A limitation that we still face is that memory requirements are huge in order to simulate seismic wave propagation up to the highest frequencies used in seismological studies on a global scale (around 1 Hz; that is, 1 s shortest period). SuperMUC, the successor system to HLRB II will certainly help to bring us closer to the goal of covering the whole range of frequencies relevant for deep Earth studies.

Acknowledgements

Bernhard Schuberth was supported by a Marie Curie Intra European Fellowship within the 7th European Community Framework Programme [FP7/2007-2013] under grant agreement no. 235861. Guust Nolet and Christophe Zarolie received support from the European Research Council (ERC Advanced grant 226837). The authors thank the Leibniz Supercomputing Centre for access to computing resources on HLRBII and the DEISA Consortium (www.deisa.eu), co-funded through the EU FP6 project RI-031513 and the FP7 project RI-222919, for support within the DEISA Extreme Computing Initiative.

References

[1] **Schuberth, B.S.A., Zaroli, C. & Nolet, G., Geophys.** J. Int., 188(3), 1393-1412, [2012]

[2] **Bolton, H. & Masters, G., J. Geophys.** Res., 106(B7), 13,527–13,540, [2001]

[3] http://www.bernhard-schuberth.de

- Bernhard Schuberth
- Christophe Zaroli
- Guust Nolet

Observatoire de la Côte d'Azur, Géoazur, Université de Nice Sophia-Antipolis, CNRS

Applications

HD(CP)²: High Definition Clouds and Precipitation for Climate Prediction

The lack of understanding of cloud and precipitation processes is arguably the foremost problem of climate simulations. The BMBF initiative "High Definition Clouds and Precipitation for Climate Prediction" or "HD(CP)²" for short will focus German community efforts on the understanding of clouds and precipitation processes. By addressing this topic, a significant reduction in the uncertainty of climate change projections, both on global and regional scales, is expected. The project consist of three elements: (i) model development and simulation, (ii) observation, and (iii) synthesis to advance cloud-precipitation prediction. The first phase will be three years to create a high resolution prototype model for a series of summer hindcasts over Germany, to collect data to allow for model evaluation, and to start the synthesis of the produced data. In a potential second phase, the model domain could be extended and complemented by other regions in the world that will allow for a deeper analysis of cloud and precipitation related projects.

Background of the project

Clouds and precipitation are an essential element of the climate system, but they are difficult to account for in theories and models [1]. Clouds significantly influence the atmospheric radiative forcing and modulate many feedback processes of the Earth system. Interactions with the terrestrial biosphere, sea-ice interactions, and biogeochemical interactions are critically dependent on the cloud-precipitation representation [2], [3]. Atmospheric clouds and precipitation are small- and meso-scale phenomena, which cannot be modeled explicitly on the global scale with current computer architectures and hence are parametrized in global climate models. Global climate simulations today show significant biases in precipitation patterns which are related to insufficient parametrization of these processes.

Aim of the project

HD(CP)² is a coordinated initiative to improve our understanding of cloud and precipitation events, to increase the quality and number of available observations, and to evaluate and improve our modeling abilities. One group of project partners will develop a highly scalable regional weather model which will be used to provide a series of ultra-high resolution summer-season hindcasts over Germany. This model will start on top of the ongoing efforts of MPI-M and DWD to develop the next generation climate model ICON. The HD(CP)² target region for this regional model is illustrated in Figure 1. This figure also shows the density and diversity of the observational network. A second major aim of the project is to improve, homogenise, and understand the multitude of available observations of cloud and precipitation processes in the target region. The network provides ground based measurements that can be combined with satellite remote sensing measurements to evaluate the simulations and to obtain statistics of cloud-precipitation processes. Combining simulation results and observational data from networks and supersites all over Germany will help to improve cloud and precipitation parametrization, allowing for improved climate prediction. The HD(CP)² simulations will also help to understand current limits on cloud-precipitation modeling and parametrization skills. The scale of the simulations allows the explicit representation and analysis of entirely new phenomena. For example, in summer months the shallow cumulus cloud regime, which has been shown to be particularly relevant for climate sensitivity, can be investigated. Surface-forced precipitation regimes, whose poor representation underlies much of the uncertainty in current representations of the hydrological cycle, can also be studied.

Challenges

The scales of clouds range from several micrometers, which is the size of a single cloud particle, to hundreds of kilometers, which is the dimension of a frontal system. It is impossible to numerically represent cloud-precipitation processes on scales that span so many orders of magnitude. Some authors have come to speak of the grid spacing range between 1 and 5 km as the grey zone, or Terra Incognita (Figure 2). In this zone many of the key interactions are only partially resolved. Aiming on a horizontal grid spacing of approximately 100 meters, the new HD(CP)² model will help to explore the grey scales of cloud and precipitation modeling. By developing the capacity to evaluate simulations of clouds and precipitation on such scales, new frontiers will be opened. Better understanding of clouds and precipitation will provide a basis for improved climate predictions. The HD(CP)² project will bring theoretical and observational work together, to a new level of understanding of the role of precipitation and clouds in the climate process. HD(CP)² will bring German research efforts to the forefront of high resolution climate modeling and climate data observation.

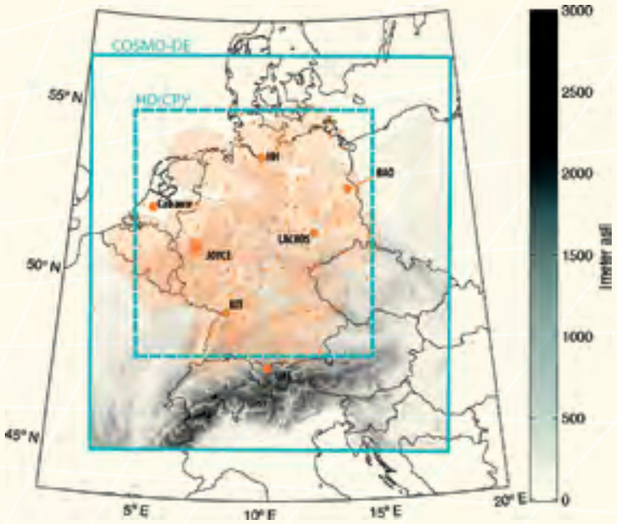

Figure 1: Map showing a possible HD(CP)² simulation domain, and the density of observational networks. The model domains are outlined in blue, the orange elements show different measurement instruments of the observational network.

Project Partners

- Computer Science, University of Cologne (CS)
- Deutscher Wetterdienst (DWD)
- Engineering Mathematics and Computing Lab, Karlsruhe Institute of Technology (KIT)
- German Climate Computing Centre (DKRZ)
- Institute of Atmospheric Physics, DLR (DLR-IPA)
- Institute for Energy and Climate, Research Centre Jülich (FZJ-IEK)
- Institute for Geophysics and Meteorology, University of Cologne (IGM)

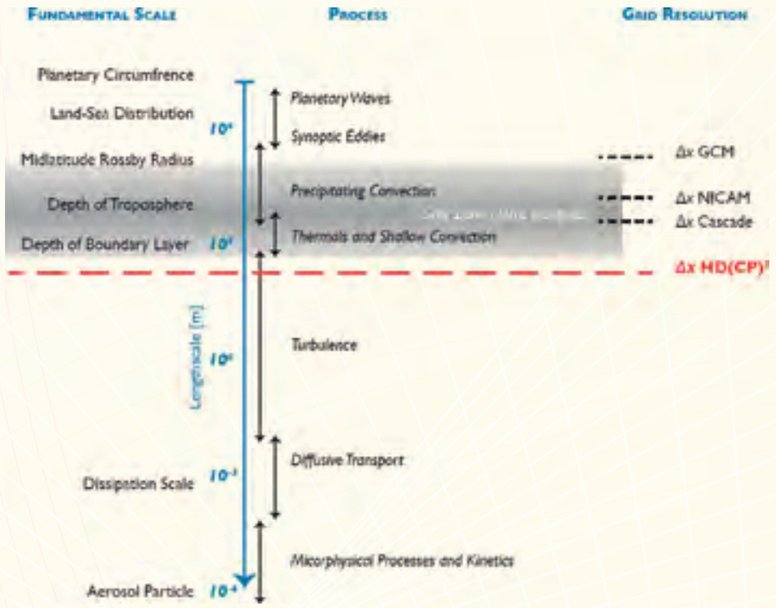

Figure 2: Diagram showing fundamental scales and processes in the atmosphere and the grid spacing employment in state of the art modeling versus grid spacing proposed for this project.

- Catrin Lehmann[1]
- Joachim Biercamp[2]
- Lars Hoffmann[1]
- Florian Rauser[3]
- Björn Stevens[3]

[1] Jülich Supercomputing Centre

[2] German Climate Computing Centre

[3] MPI for Meteorology, Hamburg

- Institute for Meteorology, University of Leipzig (LIM)
- Institute for Meteorology and Climate Research, Karlsruhe Institute of Technology (KIT)
- Institute for Meteorology and Climatology, University of Hannover (IMuK)
- Institute for Mathematics, University of Mainz (JGU Mainz)
- Institute of Physics and Meteorology, University of Hohenheim (IPM)
- Institute for Space Sciences, Free University of Berlin (FUB)
- Jülich Supercomputing Centre, Research Centre Jülich (FZJ-JSC)
- Leibniz Institute for Tropospheric Research, Leipzig (IfT)
- Ludwig-Maximilians University Munich (LMU)
- Max Planck Institute for Meteorology (MPI-Met)
- Meteorological Institute, University of Bonn (MIUB)
- Meteorological Institute, University of Hamburg (UHH)
- Richard Aßmann Observatory (DWD)
- Zuse Institute Berlin (ZIB)

References

[1] Jacob, C.
Accelerating progress in global atmospheric model development through improved parametrization Bull. Amer. Meteor. Soc., 91, pp. 869-875, 2010

[2] Hohenegger, C., Brockhaus, P., Bretherton, C. S., Schär, C.
The soil moisture-precipitation feedback in simulations with explicit and parametrized convection J. Climate, 22(19), pp. 5003-5020, 2009

[3] Kay, J. E., L'Ecuyer, T., Gettelman, A., Stephens, G., O'Dell, C.
The contribution of cloud and radiation anomalies to the 2007 arctic sea ice extent minimum Geophys. Res. Let., 35, L08 503, 2008

How the ScalaLife project integrates software developers, computing centres and life science researchers

Scalable Software Services for Life Science

The ScalaLife project (Scalable Software Services for Life Science) started in September 2010 and develops new hierarchical parallelization approaches explicitly based on ensemble and high-throughput computing for new multi-core and streaming/GPU architectures, establishes open software standards for data storage and exchange. The project implements, documents, and maintains such techniques in pilot European open-source codes such as the widely used GROMACS & DALTON, as well as a new application for ensemble simulation (DISCRETE).

ScalaLife created a Competence Centre for scalable life science software to strengthen Europe as a major software provider and to enable the community to exploit e-Infrastructures to their full extent. This Competence Network provides training and support infrastructure, and establishes a long-term framework for maintenance and optimization of life science codes.

ScalaLife for software developers

The project offers software developers efficient and scalable methodologies for compute intensive codes. Experts within the ScalaLife project can perform code analysis and profiling for codes that have been approved by the Competence Centre. Furthermore, the utilization of accelerators and alternative platforms (GPGPUs, FPGAs) is being investigated and ScalaLife can help with porting.

Among the experts in the centre are

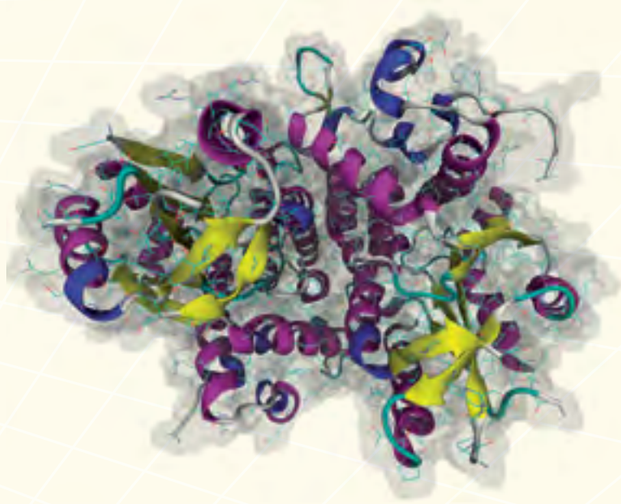

Figure 1: DISCRETE test case: Study of allosteric effects in the dynamics of proteins forming complexes.
The aim is to study the dynamics of several proteins of biological interest that form complexes with small proteins, specifically how the dynamics of a certain protein are changed when it forms a complex with another protein. This change in the dynamics can result in an allosteric effect, i.e. it can lead to a change in the activity of the protein. In some cases this mechanism can work like a switch between the inactive/active state of a protein. Using DISCRETE allows to run long molecular dynamics trajectories, taking advantage of its speed compared with standard MD codes. Contact: Ramon Goni (BSC).

the main developers and maintainers of several important software packages (so called pilot applications):

- GROMACS (widely used, possibly fastest code for MD simulations)
- DALTON (widely used package for QM and hybrid-QM/MM simulations)
- DISCRETE (new package for Discrete MD simulations)
- ERGO (memory efficient QM package)
- MUSIC (library for managing neuro simulator setups)
- XMIPP (advanced image processing software for electron microscopy research)

The goals of the project are being achieved by

- Analyzing the properties of the pilot applications and developing scalable techniques
- Optimizing the interfaces and data formats used for life science applications
- Integrating new concepts into the pilot applications, maintaining, and releasing them
- Validating results by analyzing their impact on applications running on European e-Infrastructures
- Widely disseminating results and train members of the European life science software communities, and
- Collecting and making widely available the knowledge and best practices developed by the project to the community and providing expert services

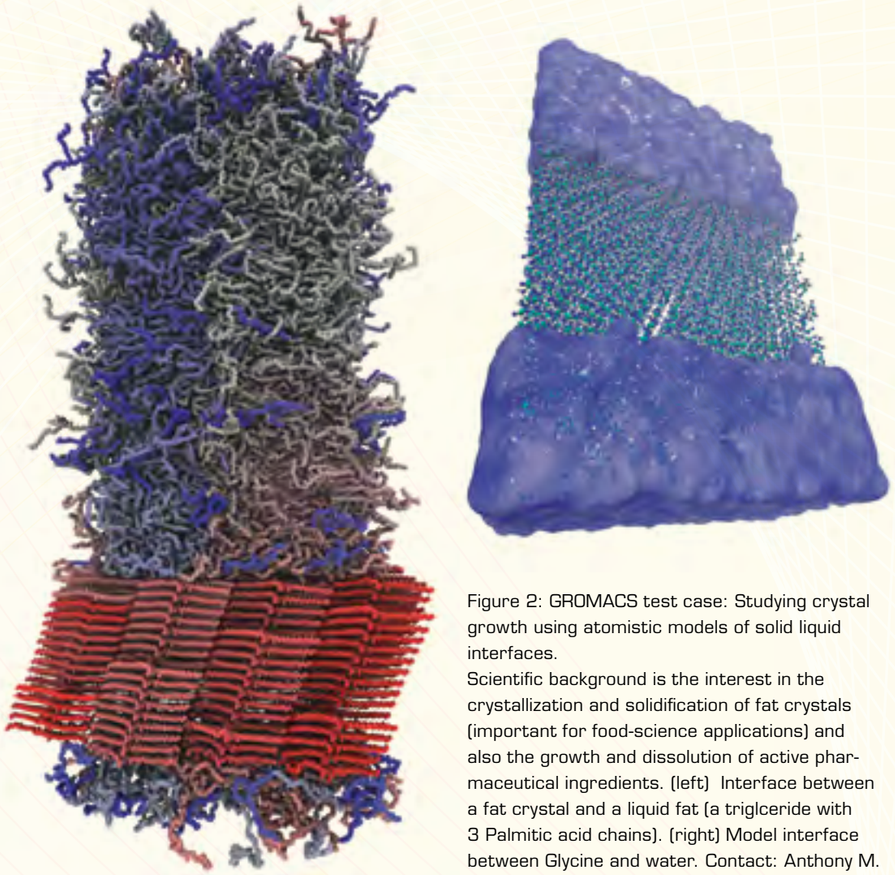

Figure 2: GROMACS test case: Studying crystal growth using atomistic models of solid liquid interfaces.
Scientific background is the interest in the crystallization and solidification of fat crystals (important for food-science applications) and also the growth and dissolution of active pharmaceutical ingredients. (left) Interface between a fat crystal and a liquid fat (a triglceride with 3 Palmitic acid chains). (right) Model interface between Glycine and water. Contact: Anthony M. Reilly (TUM).

ScalaLife Partner Centres

ScalaLife project partners include major players in the European high performance community: PDC, Center for High Performance Computing at KTH, Sweden, LRZ, Leibniz Supercomputing Centre, Germany, BSC, Barcelona Supercomputing Center, Spain, and OeRC, Oxford e-Science Research Center, UK.

The project partners offer an automated validation suite, performance profiling, application agnostic standards for data storage and exchange and an extensive knowledge base with best practices guides.

ScalaLife for researchers

Right from the start, ScalaLife approached alpha users from the European life science communities. The alpha users provided real world test cases for the pilot applications (see Figure 1-5). These test cases are at the heart of the automated validation suite. Through the validation suite, the alpha users received immediate feedback both from the computing centres and the software developers on how to optimize the configuration files and input files for faster and more robust simulations.

The Competence Centre as one-stop shop for the Life Sciences

The ScalaLife Competence Centre (www.scalalife.eu) provides a knowledgebase, tutorials, best practices guides, and forums.

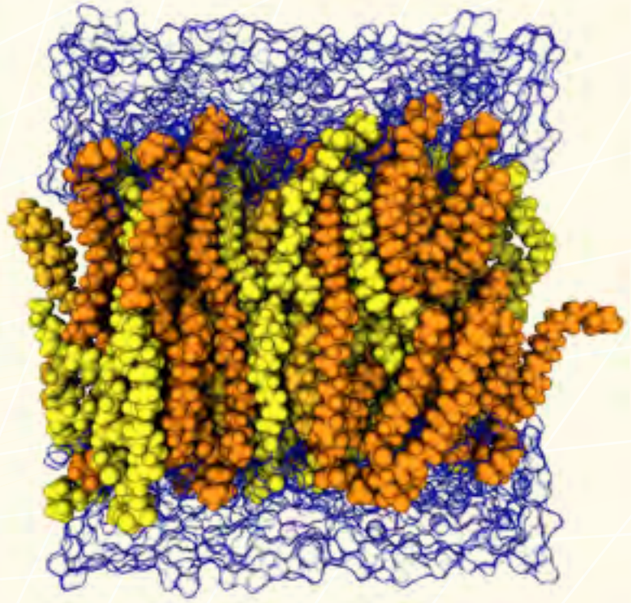

Figure 3: GROMACS test case: Modeling of mixed lipid bilayers.
The interest to model bilayers consisting of mixtures unsaturated, single-saturated and polyunsaturated lipids and cholesterol comes from the fact that such mixtures often show rich phase behavior with the formation of domains which differ in structure or composition. This domain formation may be related to rafts in biomembranes which is hypothesized as a possible pathway of cell signaling processes. It is known that affinity between polyunsaturated lipids (e.g., 18:8/22:6 PC lipid) and cholesterol is pure which can drive the phase separation process. The study of such domain formations requires molecular dynamics simulations of large model systems and very long simulation times since the molecules need to have time to find their places in the formed domains. The system studied in this example consists of a mixture of single-unsaturated (18:0/18:1w9 PC) lipids, polyunsaturated (18:0/22:6w3 PC) lipids, and cholesterol, fully hydrated by water. Contact: Alexander Lyubartsev.

Research groups can apply for exclusive support by application experts at the Competence Centre. After that, they will receive help with issues concerning the deployment of software codes and optimal utilization of available hardware resources. Additional support with analysis of the computational output is also available.

The Competence Centre provides the following services:

- Code analysis/profiling of new software packages

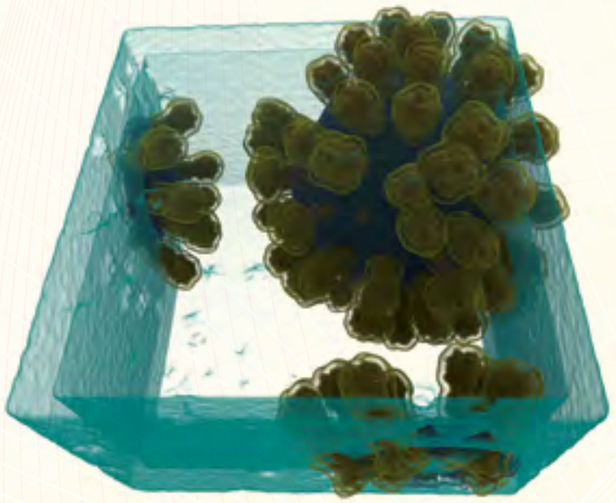

Figure 4: GROMACS test case. Large scale coarse-grained simulations of an influenza virion (5.1 million coarse-grained particles). The aim of this simulation is to study the dynamics of the envelope of a realistically sized influenza virion (84 nm diameter). Specifically to investigate the diffusion and possible aggregation of the main membrane proteins present in the envelope (hemagglutinin, neuraminidase and the influenza M2 channel). Contact: Philip Fowler.

- Benchmarking of user tests cases, e.g. for the use in proposals
- User clinics for improving the setups of simulated models
- Guidance on efficient software usage
- Asynchronous support by experts via discussion forums/mailing lists
- Support to centres and resource providers on efficient configuration of the applications
- Second level support to users via the resource providers

An automated validation suite to test and benchmark the ScalaLife software packages was developed, based on the DEISA validation suite.

In Figures 1-5 are shown some of the test cases from alpha users that were included in the validation suite. Figure 6 shows the scaling behavior of GROMACS with the influenza virion test case (Fig. 4).

Conclusion

ScalaLife successfully launched the Competence Centre pilot at the end of project year 1 and a main objective of the ScalaLife project is to create a Competence Centre, which will harness the accumulated knowledge in order to support the computational Life Science communities in Europe. Through training and provisioning of support infrastructures, the centre is establishing a long-term framework for the maintenance and optimization of life science codes.

This objective is very important for the Life Science software ecosystem in Europe. The community will highly benefit from efficient exchange of the knowledge acquired from the work in different projects (EU-funded ones or not), and especially from connecting software developers with experts in HPC software engineering, code optimization and method/algorithms development. It is also important to point out that improving the software efficiency directly translates to monetary savings for resource providers such as HPC centres.

Aknowledgement

The ScalaLife project receives funding from the EC's Seventh Framework Programme (FP7/2007-2013) under grant agreement No. RI-261523.

We thank the scalaLife alpha user community for their continuing support of the project: P. Fowler, G. Portella, A. Laaksonen, A. Lyubartsev, C. Scheurer, J. Weber, L. Frediani, A. Reilly, X. Li

Reference

http://www.scalalife.eu

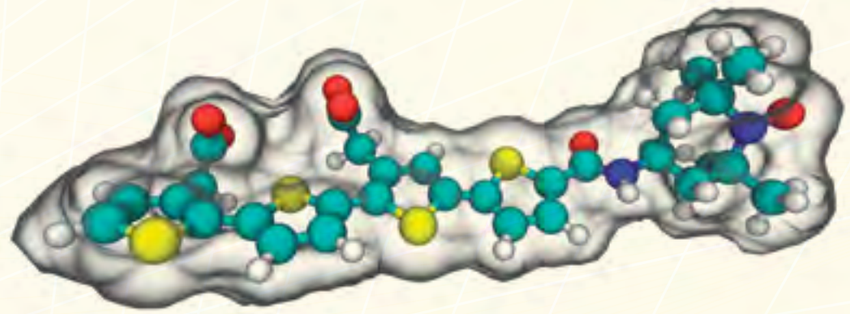

Figure 5: DALTON test case. Fluorescent spin label bound to insulin beta sheet model. Development of novel fluorescent spin labels suitable for EPR imaging of Alzheimer's disease require knowledge of they EPR spin Hamiltonian parameters in free and bound states. QM/MM calculations can accurately determine electronic g-tensors and hyperfine coupling constants of spin labels and in this way guide rational design of spin labels for EPR imaging. Contact: Xin Li and Zilvinas Rinkevicius.

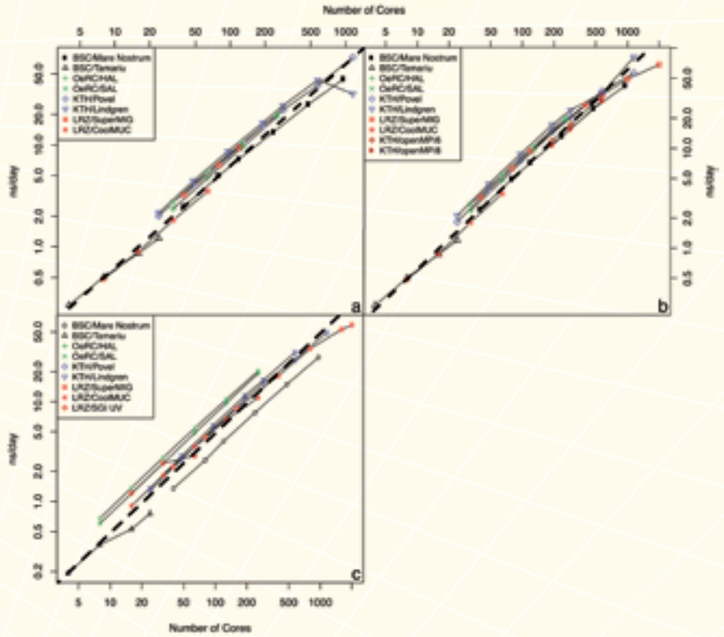

Figure 6: Scaling behavior of the molecular dynamics code GROMACS when running the validation suite with the input files from the influenza virion (see Figure 4). The Figure shows the accumulated results when running the validation suite on different architectures at the participating computing centres. a) the first version of the validation suite was used in 2010 to get a baseline. b) and c) show the results from running the validation suite again in 2011 and 2012, each time with updated versions of the software packages used.

- Helmut Satzger
- Rossen Apostolov
- Lilit Axner
- Mihai Duta
- Judit Gimenez
- Ramon Goni
- Berk Hess
- Ferdinand Jamitzky
- Erwin Laure
- Albert Solernou
- Momme Allalen
- Christoph Bernau
- Zilvinas Rinkevicius
- Modesto Orozco
- Josep Lluis Gelpí
- Agusti Emperador

JUNIPER takes aim at Big Data

Motivation

The trend towards Cloud technology, the advances of Internet technology, and the proliferation of mobile computing (in the context of "Internet of Things") are just a few factors contributing to the massive amounts of data collected by means of the modern IT solutions, which is frequently referred in the Science and Technology as a "Big Data" problem. The age of Big Data presents a variety of challenges in how to store, access, handle, stream, process, and analyze massive amounts of structured and unstructured data effectively. In the near term, the major issue with the Big Data oriented applications will be to find new ways to leverage the technology available today to improve the performance and scalability, reduce response time, figure out more efficient ways to deliver services to end users, and meet the growing demand for virtualized (web and cloud) Big Data services. Using High Performance Computing infrastructures is essential in dealing with these issues.

Project Description

JUNIPER (Java platform for hIgh PErformance and Real-time large scale data management) is an EU FP7-ICT project, started in December 2012, aiming to establish a development platform for new-generation data-demanding applications. JUNIPER will exploit the power of heterogeneous high performance (supercomputer) and on-demand (cloud) computing and storage infrastructures in order to help Big Data applications provide real-time performance guarantees. Creation of more effective data processing environments, also in terms of the power consumption, and offering real-time guarantees of the data services are tremendous challenges that JUNIPER addresses. Guided by requirements of enterprise Big Data applications, the project will exploit synergies between the major parallelization technologies (such as MPI, MapReduce, COMPSs) and elaborate new paradigms in data-centric parallel processing that will balance flexibility and performance of data processing applications in heterogeneous computing architectures. JUNIPER will deliver a promising platform for tackling today's and also tomorrow's Big Data challenges and embrace new opportunities for the Java programming language to reinforce its place as the major technology for development of the effective and scalable solutions for data-intensive problems.

The JUNIPER platform (see Figure 1) will empower the developers to transcend the limitations of today's data processing environments by leveraging the full-backed high performance computing technology. The advantages of using JUNIPER will result in spawning much more users and increasing data services' transaction rate.

A related issue for the JUNIPER platform is how to simplify the development process and minimize the time-to-market costs of data-intensive applications. The key of the JUNIPER technology is that it will allow the application developers predict the characteristics of the data-centric application workflows, including performance, scalability, throughput, etc., on the available hardware resources, and vice versa, will make suggestions for the computing platform required to gain the specified characteristics. Guided by these requirements, JUNIPER will elaborate methodologies for modelling (e.g. based on UML) application work- and data-flow and will develop necessary tools for this.

Project Partners

Apart from the High Performance Computing Center Stuttgart (HLRS, Germany), the following organizations participate at the project:

- The Open Group (United Kingdom) - coordinator
- Aicas GmbH (Germany)
- University of Your (United Kingdom)
- Scuola Superiore Sant'Anna (Italy)
- Softteam (France)
- petaFuel (Germany)
- Luminis (The Nertherlands)

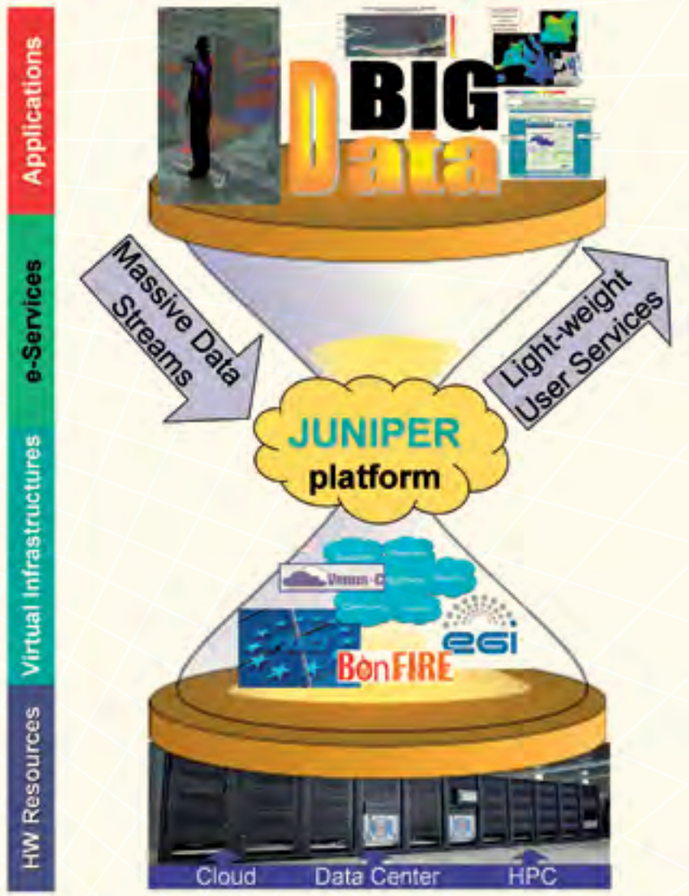

Figure 1: JUNIPER – a development platform for high performance and real time Big Data applications.

- Alexey Cheptsov
- Bastian Koller

University of Stuttgart, HLRS

Projects

ECO_2Clouds – Experimental Awareness of CO_2 in Federated Cloud Sourcing

Project Outline

Ecological implications of cloud-based IT infrastructures are creating a critical gap in the current state of the art in research and business. The ECO_2Clouds project investigates strategies that can ensure effective application deployment on the cloud infrastructure, reduce energy consumption and by association CO_2 emissions and furthermore the costs of the whole execution. The need for novel deployment strategies becomes more evident when an application spans multiple clouds. Cloud providers operate under different regulatory frameworks and cost structures in relation to environmental policies and energy value-chains. In addition, the optimization of key assets like virtual machines, applications and databases is constrained by a set of requirements such as quality, privacy and cross-platform service level agreements.

ECO_2Clouds will provide a challenging and innovative approach to cloud computing service delivery by:

- Developing extensions and mechanisms for cloud application programming interfaces to collect eco-efficient data at the physical cloud infrastructure and virtual machine levels
- Investigating the key environment and the needed quality and cost parameters in order to deploy virtual machines in multi-cloud environments
- Developing evaluation mechanisms and optimization algorithms to assess different parameter configurations and their influence in energy-efficient cloud sourcing and application deployment strategies.

The carbon-aware mechanisms will be integrated into the Future Internet Research and Experimentation (FIRE) facility BonFIRE [1] to test, validate and optimize the ecometrics, models and algorithms developed. Furthermore, the availability of eco-efficient deployment strategies improves the FIRE offering.

ECO_2Clouds currently comprises 3 geographically distributed testbeds across Europe which offer heterogeneous cloud resources including the possibility to measure precisely the energy consumption of each physical host. As the entire infrastructure providers are also part of the BonFIRE project, the resources can be accessed seamlessly by describing an application energy profile which is parsed by the ECO_2Clouds scheduler in order to submit the extended experiment descriptor to the BonFIRE control mechanisms which are mainly based on the Open Cloud Computing Interface (OCCI). So ECO_2Clouds will reuse the components of the BonFIRE project: the implemented scheduling algorithm will extend the experiment descriptors by application profiles which represent the application itself and the desired power consumption of the whole execution. Figure 1 gives details about the participating infrastructure providers of ECO_2Clouds with the possiblity of energy aware deployment and execution of applications.

ECO_2Clouds will develop and test an application deployment strategy which will primarily consider the environmental implications of deploying applications on multi or federated clouds. Among the multiple factors that will be considered in a successful application deployment strategy, ECO_2Clouds will conduct experimentally-driven research to investigate innovative practices for sustainable federated cloud sourcing. This will be done by examining key model parameters (ecological, quality and cost related dimensions) and proposing an optimized deployment model for cloud infrastructures.

ECO_2Clouds will address the following challenges for federated clouds:

- To collect and expose carbon footprint concerns at the infrastructure and virtual machine level
- To incorporate carbon footprint concerns into a federated cloud deployment strategy
- To develop optimization mechanisms that can enable optimal utilization of federated cloud resources
- To develop intelligent adaptation mechanisms that can perform changes during a running application based on its energy consumption

Experimental case studies will assess and compare the degree of energy efficiency and CO_2 footprint reduction achieved by ECO_2Clouds. As several parts of the BonFIRE stack are used, all case studies can be executed with and without the support of optimization and adaptation strategies to analyse the results in performance and energy consumption in detail.

The whole software and component stack is detailed in Figure 2 which describes the interaction to the BonFIRE project as well.

Work of HLRS

Within the two year duration of ECO_2Clouds, HLRS is responsible for various tasks, starting with providing an efficient and scalable monitoring infrastructure for all the required data, offering two case studies for the improvement of the ECO_2Clouds scheduler as well as providing a dedicated cloud infrastructure to the project.

The main part of the HLRS work is implementing the monitoring infrastructure including a data mining approach to capture long term deployment information. In the past, various monitoring solutions were developed and implemented at HLRS, for example in the GAMES [2] and OPTIMIS [3] projects which will be reused for the monitoring framework in ECO_2Clouds. These solutions mainly rely on the monitoring tool Nagios [4] whereas BonFIRE already provides a complete Zabbix [5] monitoring framework. Hence, the monitoring solutions need to be adapted in order to overcome different monitoring approaches as well as potential monitoring overhead.

For the two case studies, the ones used in the GAMES project are introduced again. The first case study deals with Finite Element (FEM) simulations

Figure 1: Schematic Infrastructure of ECO_2Clouds

for malicious bones which represent a high performance computing (HPC) application. It needs to be ported into the cloud environment for the execution and will produce high loads for the scheduler calibration, configuration and improvement. The second case study is dealing with modern e-Business applications: various benchmarks like Linpack as well as I/O intensive operations were executed to simulate a common e-Business application behaviour. For that purpose, load balancing methods using intelligent deployment approaches will also be examined to distribute the workload and interact with the ECO_2Clouds scheduler.

The third part of the work is providing the mandatory cloud infrastructure to execute deployments and obtain the precise energy consumption data. For that purpose, the already offered infrastructure of BonFIRE will be extended with power distribution units (PDU). In addition, HLRS will host the central access points of ECO_2Clouds, an application portal and the scheduler, as well.

- Michael Gienger
- Axel Tenschert

University of Stuttgart, HLRS

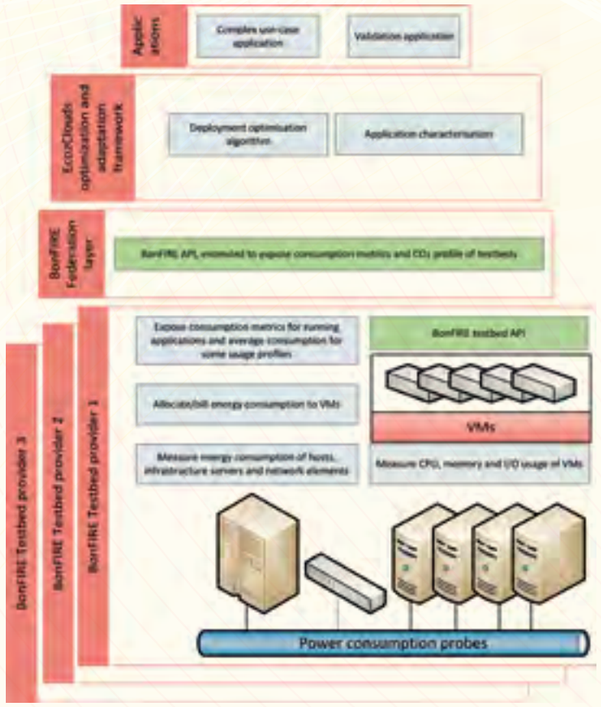

Figure 2: The ECO_2Clouds Stack

References

[1] BonFIRE Project Website: http://www.bonfire-project.eu/, Last visited: Feb 28, 2013.

[2] GAMES Project Website: http://green-datacenters.eu/, Last visited: Feb 28, 2013.

[3] OPTIMIS Project Website: http://www.optimis-project.eu/, Last visited: Feb 28, 2013.

[4] Nagios Website: http://www.nagios.org/, Last visited Feb 28, 2013.

[5] Zabbix Website: http://www.zabbix.com/, Last visited Feb 28, 2013.

Key Facts

ECO_2Clouds is a small or medium-scale project funded by the European Commission within the 7th Framework programme. The consortium consists of 6 project partners and was started at the beginning of September 2012 and will run until the end of August 2014.

ECO_2Clouds Members

- ATOS Spain SA (ATOS)
- University of Manchester (UNIMAN)
- EPCC, The University of Edinburgh (EPCC)
- Politecnico di Milano (POLIMI)
- Institut National de Recherche en Informatique et en Automatique (INRIA)
- Universität Stuttgart (USTUTT-HLRS)

ECO_2Clouds Project Website

http://www.eco2clouds.eu/

Model Based Cloud Application Development using PaaSage

Cloud computing is a popular concept in ICT. The concept of infinitely scalable elastic resource changing without complex systems administration and paying only for resources used is attractive for both high- and low-end users and providers. These benefits are however not immediately realisable and many providers are currently disappointed by the capabilities of cloud system, especially if they simply try to move their application to the cloud and expect these benefits to arise “by magic”.

The primary obstacle already arises from the diverging understanding of “clouds” in the first instance. As the EC expert group has noted in the 2012 report [1], the key characteristics behind clouds are

(1) Resource provisioning on demand

(2) Optimisation of resource utilisation

(3) Serving a dynamic number of stakeholders

Effectively, this means that the application / service is instantiated and hosted as many times as needed to satisfy the quality (availability) demands. The aim is always to reduce the cost for offering this service. For simple services, such as search engines or book sales, this may be considerably straight-forward, but the problems begin when the users share content and potentially even face time-constraints, such as in auction houses.

Modern applications expose way more complicated requirements towards the infrastructure than these simple cases, including e.g. partially shared data, various computation and visualisation tasks, and mobile end-users. Porting such applications into cloud environments whilst still maintaining all non-functional properties, such as response time, minimal resource utilisation etc. is a highly complex task in itself. Software engineering approaches still build up on classical principles (single user, single instance etc.) and therefore cannot help the developer in building up his cloud-dedicated application. At the same time, experience and expertise in this context is still very low and will only grow slowly over time.

The PaaSage project [2] addresses this concern by providing the necessary tools and means to support the full application lifecycle in clouds (cf. Figure 1).

Application Development

The primary tools enable developers to generate applications that exploit cloud specific characteristics. To this end, PaaSage regards applications as module-based workflows that need to satisfy varying non-functional properties throughout – these include aspects such as:

- Response time
- Data size
- Availability
- Sharing state
- Security
- etc.

Depending on circumstances, these factors apply to multiple modules, respectively to different sections of the application and thus have to be aligned with each other and broken down ("decomposed") according to the scope of the section covered.

Obviously, this decomposition is impacted not only by the non-functional requirements, but more importantly by the infrastructure characteristics and the usage context, such as number of users, range of dynamicity etc. To this end, PaaSage builds up an expert system that allows storing typical decomposition patterns as best use recommendations.

Generation and Deployment

Once the concrete properties per module have been identified, the application can principally be broken down into a set of services according to these aspects / features. PaaSage will support this process by injecting modes of communication and control features into the code, thus enabling execution according to these properties on the destination cloud platform / infrastructure.

This step must respect the specific conditions of destination platform and requirements, whilst ideally remaining platform independent. In other words, the modules generated should principally be executable in any cloud environments. To this end, PaaSage, together with MODAClouds (3) and Artist (4) will generate a common modelling and development language for clouds that enables the developer, respectively the developing toolkit, to generate platform-independent code that can be converted to different destination environments.

Execution an Adaption

Conversion alone is however insufficient, as one implication from decomposing the application and its properties consists in the direct and indirect relationships between the individual modules. This includes on one hand, the direct communication links between modules as foreseen by the developer: as the number and location of instances may vary dynamically at runtime, additional mechanisms need to be foreseen that cater for this dynamicity. As opposed to classical distributed systems, this also includes a dynamic n-to-m relationship in the communications.

On the other hand, the behaviour itself, i.e. aspects such as when to replicate, when to relocate and where etc. all depend on the conditions arising in various other modules and environments. For example, the response delay in one module may lead to the scale out in various other sections, or adding a further user may generate multiple replications etc. To this end, control communication needs to be added to the application, that is not directly foreseen by the application itself.

PaaSage will provide an execution wrapper that controls the module behaviour under various conditions. This wrapper will thereby link to other instances and monitoring sources according to the specifications given by the developer.

The consortium brings together research experts from different aspects of cloud computing, including software and services, High Performance Computing, programming models and systems development environments, with industrial representatives from cloud technology provisioning and usage.

PaaSage has started in October 2012 with a duration of 4 years and an overall budget of 8.4 Million Euros. PaaSage consists of 18 partners from 9 European countrys. It is partially funded by the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 317715.

Contact Persons

- Bastian Koller, HLRS
- Anthony Sulistio, HLRS

References

[1] L. Schubert and K. Jeffery (2012). Advances in Clouds - Research in Future Cloud Computing. European Commission, Brussels. Available online: http://cordis.europa.eu/fp7/ict/ssai/docs/future-cc-2may-finalreport-experts.pdf

[2] PaaSage project: http://www.paasage.eu/

[3] MODAClouds project: http://www.modclouds.eu

[4] Artist project: http://www.artist-project.eu

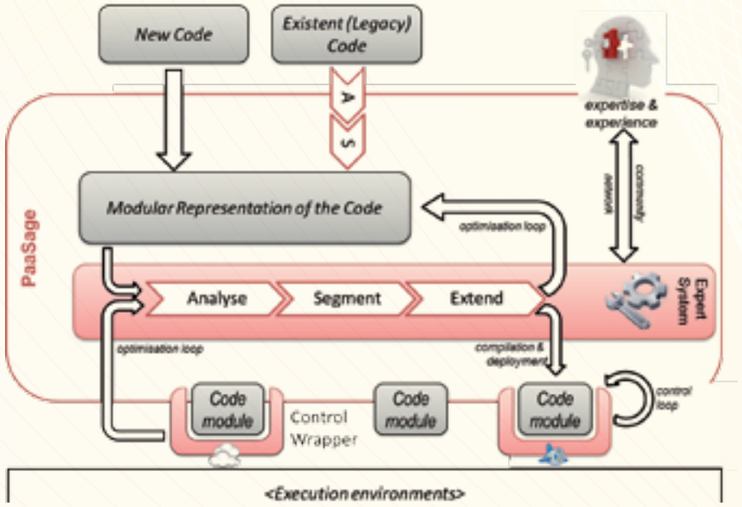

Figure 1: schematic view of the PaaSage system

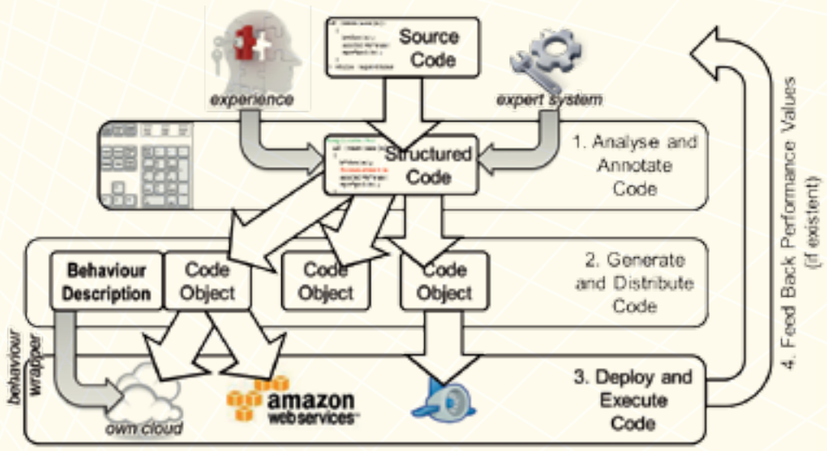

Figure 2: application lifecycle in PaaSage

• Bastian Koller

University of Stuttgart, HLRS

Inauguration of HLRS CAVE

As supercomputing technology becomes more complex, simulation and visualization have become increasingly important tools for researchers to gain scientific insight. As a result, the High-Performance Computing Center Stuttgart (HLRS) has built a user-friendly, adaptable visualization infrastructure for its users. The new HLRS building, inaugurated in October of 2012, offers world-class visualization to accompany world-class computation.

Through the state-of-the-art Cave Automatic Virtual Environment (CAVE), the high resolution visualization wall and other visualization tools, HLRS offers HPC users additional science and research tools to gain real insight from the ever-growing volumes of increasingly complex computing results. The CAVE is a fully immersive 3D environment and uses five rear projection surfaces for visualizations, allowing teams of researchers to step inside their simulations. Using a special tracking system, the visualizations are oriented dynamically and accurately in real-time in relation to a viewer's position and direction of view. With a special mouse, he or she can navigate through their results.

Virtual reality allows researchers to work with large data sets of simulation data in real time. Through the dynamic adaptation of perspective in the CAVE, a user working inside gets the feeling of actually moving inside of a "virtual world". Simulation models and data sets can be enriched with completely simulated nature environment consisting of forests, mountains, valleys, river beds, and probable interventions in nature and landscape – such as the placing of a power plant into said environment - can be illustrated in real time and its effects on it can be observed instantaneously.

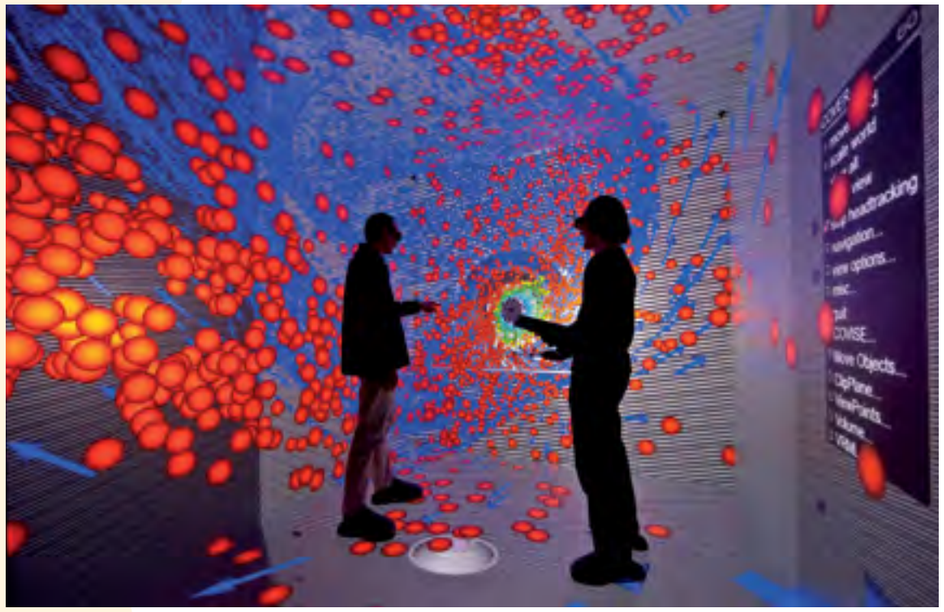

Figure 1: ReCom - Visualization of particle and air flow inside a furnace

In addition, the HLRS visualization team developed the Collaborative Visualization and Simulation Environment (COVISE), an adaptable software package for researchers trying visualize anything from supernovae to the building blocks of matter. COVISE is widely used in the visualization community, as it allows researchers to easily modify simulation parameters and plug in different tools. In addition, COVISE can be used to connect visualization environments remotely, giving a research in one part of Europe the opportunity to present findings to a group of researchers in another place.

Figure 2: Ephesos - Visualization of the Paulus cave on the HLRS high-resolution tiled display

System Specifications

The new HLRS CAVE has 5 projection surfaces–3 walls, ceiling and floor–forming a cube with an edge length of 2.7 m. The projection technology extends over 3 stories. Single chip DLP projectors with two separate input signals for left and right images render a resolution of 1200 x 1200 pixels.

The visualization wall at HLRS is comprised of 9 professional HD 3D-Displays (slim bezel) in a 3x3 setup. Each display is driven by one Render-Node.

The CAVE and the visualization wall is operated by a cluster of 22 Intel Sandybridge nodes which come with 11 K5000 and 11 K6000 NVIDIA Quadro GPUs and QDR Infiniband interconnect.

Figure 3: Forbach - Forbach and surroundings rendered in the CAVE

Projects

RECOM Services, in cooperation with energy providers and power plant manufacturers uses HLRS supercomputers for designing and modernizing coal power plants worldwide. Their simulation of the Altbach power station allows research teams to plug in different variables and decide how modernization can be most effective in preventing corrosion to the plant, keeping emissions at the lowest possible level, and producing electricity efficiently. In their visualization, engineers can visualize temperature distributions and gas concentrations anywhere in the furnace. In this visualization, they can study complex flow fields through massive particle animations of coal powder and fresh air as well as the chemical reactions during combustion and exhaust gas cleaning.

Researchers at the HLRS have created an entire replica of Forbach, a town in the Black Forest, its Rudolf Fettweis power station, and surroundings. In this simulation, researchers can not only go above and below ground to observe how the dam, underground turbines, and reservoir work together, but also go to any house in the area and observe how construction and modernization projects may affect citizens, wild life and the environment in general.

One of the biggest challenges for high-performance computing is keeping machines at optimal operating temperature. HLRS computer scientists are working as part of a consortium called the CoolEmAll project, which works toward simulating computational resources at work, calculating where the majority of heat builds up, and remediating energy waste through more efficient cooling. By using visualization, staff at various HPC centers can make informed decisions about how the layout of a computer room will affect the ability to cool it efficiently as well as deciding on the best cooling methods.

Archeologists from the Technical University of Vienna have been exploring the Paulus cave in Ephesos for more than 13 years now. The analysis of the wall paintings, which have up to five layers, each depict different religious subjects. These layers presumably resulted from religious and cultural changes in society over time, and are very difficult to analyze without destroying historical testimony. The most famous painting in the cave shows one of the oldest well-preserved illustrations of the apostle Paulus. A high-detailed visualization of the cave with a depth of 15 m located at the northern slope of Bülbülda helps researchers to get better insight into background story of contemporary church history.

By integrating existing facilities across Europe, the Visionair project (VISION Advanced Infrastructure for Research) aims to create a world-class research infrastructure for conducting state-of-the-art research in visualization, thus significantly enhancing the attractiveness and visibility of the European Research Area (ERA). With over 20 members participating, VISIONAIR offers facilities for virtual reality, scientific visualization, ultra high definition, augmented reality and virtual services. The group's current scientific challenges are focused on issues such as climate evolution, environmental risks, molecular biology, health, and energy. These issues require the management of increasingly complex and voluminous information, thus calling for the development of ever more powerful visualization methods and tools.

• Wolfgang Schotte

University of Stuttgart, HLRS

Figure 4: CoolEmAll - Interactive simulation of airflow and temperature distribution using a tangible interface

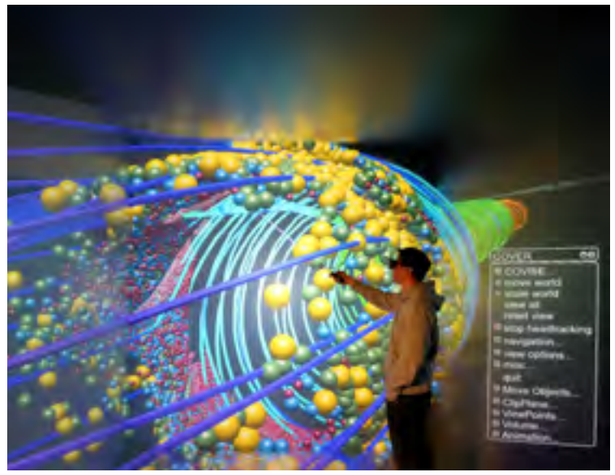

Figure 5: Visionair - Large CFD data rendered in real-time

Systems

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, national and international 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
85748 Garching near Munich
Germany

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

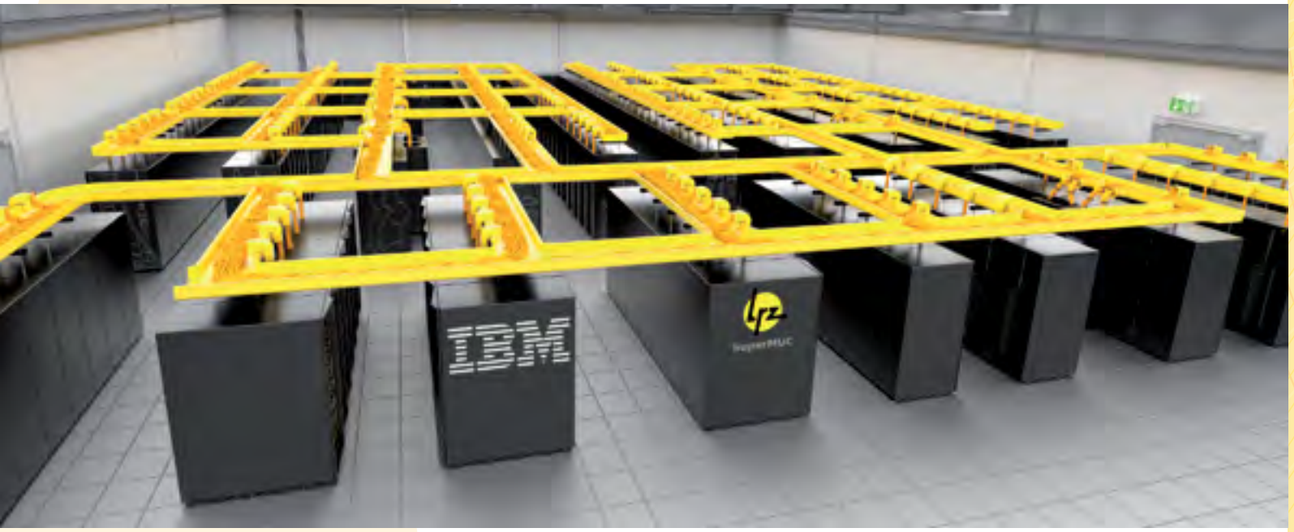

Picture of the Petascale system SuperMUC at the Leibniz Supercomputing Centre.

Compute servers currently operated by LRZ

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
IBM System x iDataPlex "SuperMUC"	18 islands 512 nodes with 2 Intel Sandy Bridge EP processors each 147,456 cores 288 TByte main memory	3.185	Capability Computing	German universities and research institutes, PRACE (Tier-0 System)
IBM BladeCenter HX5 "SuperMIG"	205 nodes with 4 Intel Westmere EX each 8,200 cores 52 TByte main memory	78	Capability Computing	German universities and research institutes, PRACE
Linux-Cluster	510 nodes with Intel Xeon EM64T/ AMD Opteron 2-, 4-, 8-, 16-, 32-way 2,030 Cores 4.7 TByte	13.2	Capacity Computing	Bavarian and Munich Universities, LCG Grid
SGI Altix ICE	64 nodes with Intel Nehalem EP 512 Cores 1.5 TByte memory	5.2	Capacity Computing	Bavarian Universities
SGI Altix Ultraviolet	2 nodes with Intel Westmere EX 2,080 Cores 6.0 TByte memory	20.0	Capability Computing	Bavarian Universities
Megware IB-Cluster "CoolMUC"	178 nodes with AMD Magny Cours 2,848 Cores 2.8 TByte memory	22.7	Capability Computing, PRACE prototype	Bavarian Universities
MAC research cluster	64 Intel Westmere Cores, 528 Intel Sandy Bridge Cores, 1,248 AMD Bulldozer Cores, 8 NVIDIA GPGPU cards, 8 ATI/AMD GPGPU cards	40.5	Testing accelerated architectures and cooling technologies	Munich Centre of Advanced Computing (MAC), Computer Science TUM

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

HLRS

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 SICOS BW 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. hc. Dr. h.c.
Michael M. Resch
Nobelstraße 19
70569 Stuttgart
Germany

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

View of the HLRS Cray XE6 "HERMIT"

Compute servers currently operated by 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 Hybrid Architecture	12 16-way nodes SX-9 with 8 TByte main memory + 5,600 Intel Nehalem cores 9 TB memory and 64 NVIDIA Tesla S1070	146	Capability Computing	German Universities, Research Institutes and Industry, D-Grid
IBM BW-Grid	3,984 Intel Harpertown cores 8 TByte memory	45.9	Grid Computing	D-Grid Community
Cray XT5m	896 AMD Shanghai cores 1.8 TByte memory	9	Technical Computing	BW Users and Industry

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

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

JÜLICH FORSCHUNGSZENTRUM

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

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

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

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

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

Contact:
Jülich Supercomputing Centre (JSC)
Forschungszentrum Jülich

Prof. Dr. Dr. Thomas Lippert
52425 Jülich
Germany
Phone +49-24 61-61-64 02
th.lippert@fz-juelich.de
www.fz-juelich.de/jsc

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"	3,288 SMT nodes with 2 Intel Nehalem-EP quad-core 2.93 GHz processors each 26,304 cores 77 TByte memory	308	Capacity and Capability Computing	European Universities, Research Institutes and Industry, 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
IBM Cell System "QPACE"	1,024 PowerXCell 8i processors 4 TByte memory	100	Capability Computing	QCD Applications SFB TR55, PRACE

Activities

First JUQUEEN Porting and Tuning Workshop

The Jülich Supercomputing Centre (JSC) recently finished installing JUQUEEN to a final size of 28 racks of the latest Blue Gene architecture. As with any new system, new challenges arise for its users to unlock the full potential of the architecture. Designed to improve the users' expertise in using JUQUEEN, this first workshop was held from 4th to 6th of February 2013 at JSC. As a PRACE Advanced Training Centre (PATC) course it attracted 46 participants from eight European countries, all with active scientific projects on JUQUEEN. The goals of the workshop were to make the participants more familiar with the system installed at JSC and to provide tools and ideas to help with porting their codes, analysing the performance, and in improving the efficiency.

The program of the workshop included in-depth talks on very specific features of the Blue Gene/Q architecture alongside introductory talks to get the participants started. Topics covered were the Blue Gene/Q hardware, best practices for programmers, performance tools and debuggers as well as guidelines for OpenMP usage and parallel I/O. The specialists' talks explored the memory system including methods for prefetching and atomic operations for the level 2 cache. They also elaborated on transactional memory, speculative execution, intrinsics for vectorisation, and low-level networking. The main part of the workshop was then spent in hands-on sessions with the users' codes. These hands-on sessions were supervised by almost 20 members of staff from JSC's Simulation Laboratories and cross-sectional teams (Application Optimisation, Performance Analysis, Mathematical Methods and Algorithms) as well as from IBM.

Among the many promising improvements that had been identified were those of parallel I/O and a hybrid programming model. Some participants gained an immediate speed-up of several tens of percent with OpenMP adjustments or by better understanding the vendor's compilers. The workshop also helped JSC and IBM staff to appreciate the needs and requirements of users on JUQUEEN much better. At the same time, a closer collaboration between the participants and the Simulation Laboratories at JSC could be initiated.

The participants' general perception of the workshop seemed very positive and inspiring. Especially the hands-on sessions with direct support by the hard- and software experts were received well and helped solving specific problems. To quote two of the participants on what they enjoyed on the workshop: "I really liked, the hands on sessions. I would need to spend so much more time if I were to do all these things on my own - so much perhaps, that I just would not do it at all" and "The school was with very useful content and well organized from a team with excellent professional expertise". So due to the success of this workshop and expected future demand it is planned to stage similar workshops as a regular event.

The slides to the talks can be found on the web at http://www.fz-juelich.de/ias/jsc/jqws13.

• Dirk Brömmel

Jülich Supercomputing Centre

Figure 1: Participants of the First JUQUEEN Porting and Tuning Workshop (Source: Forschungszentrum Jülich GmbH)

HPC Changing Europe`s Industrial Landscape - 5th PRACE Industrial Seminar, Stuttgart

Figure 1: The attendees of the seminar

Already being a tradition in the European HPC landscape the PRACE Industrial Seminars serve as a forum to develop PRACE's working relations with industry.

This year the GCS partner HLRS organized the seminar which took place in the small but lovely city of Bad Boll in the Stuttgart region during two half-days on April 15th and 16th 2013. In Bad Boll, the 78 seminar attendees from 12 countries listened to presentations about Europe's most exiting HPC applications used in industry and were informed about PRACE's future plans for its industrial relations. Almost 40 speakers contributed to a variety of plenary talks and special topics sessions.

Figure 2: Sergi Girona, Chair of the PRACE Board of Directors

In the last four years PRACE had held Executive Industrial Seminars in Amsterdam, Toulouse, Stockholm and Bologna. In total, the five Industrial Seminars till now have attracted as many as 142 different companies of various business and industrial profiles.

The organizers were glad to call up leading decision makers to embrace the strategic importance of HPC and the issue of accessing EU and national supercomputing infrastructure.

The first keynote speaker, Christoph Gümbel of Porsche, gave a presentation about the "Strategic importance of High Performing Computing for vehicle development using Porsche as an example".

The second keynote speaker, Suzy Tichenor from the US Oak Ridge National Laboratory, spoke about "Accelerating Competitiveness with Leadership Computing". Ms. Tichenor pointed out the interest of the US HPC organization in a deeper cooperation with its European partners.

Figure 3: Suzy Tichenor, Oak Ridge National Laboratory, who gave a presentation about industry support in the U.S.

Leonardo Flores Anover from the European Commission explained the actual EU HPC strategy [1] which combines three elements: (a) developing the next generations of HPC towards exascale; (b) providing access to the best facilities and services for both industry and academia, and (c) achieving excellence in HPC applications. In his view these three elements are not independent and should work in synergy. Furthermore Mr. Flores Anover pointed out that the European Commission has proposed an ambitious programme reflecting the European Union`s support to research and innovation for the coming years. This programme covers research, technological development, demonstration and innovation for the seven year period 2014-2020. The European Union expects that the support to the HPC strategy will be properly reflected in the Horizon 2020 programme.

In Bad Boll the seminar attendees from 56 companies were informed about PRACE's Open R&D Industry Access. Twice a year, European companies or companies having significant research activity located in Europe, can apply to a PRACE Call for Proposals with scientifically innovative projects of which they are able to publish the results at the end of the award period. More information can be found at [2].

In order to meet the different prerequisites and requirements of large companies and SMEs (small and medium sized companies), the details of PRACE's R&D Industry Access were presented and discussed in two separate parallel sessions. In his presentation on "HPC Provisioning for Industrial Corporates" Alfred Geiger of T-Systems pointed out that for industry Tier-0 or Tier-1 systems are generally uneconomic and that they should be provided as public infrastructure, with analogy to roads or airports.

PRACE's main effort in the next year will be to strengthen the support for SMEs. Therefore, a special HPC Adoption Programme for European SMEs, SHAPE, will be launched with a pilot project in June 2013. The programme aims to equip the SMEs with the expertise necessary to take advantage of the pos-

Figure 4: Leonardo Flores Anover, European Commission, explains the EU HPC strategy

sibilities opened by HPC and increasing their competitiveness. More than twenty companies from the SME industry sector paid attention to the idea of a low-threshold HPC support which aims to overcome the barriers to HPC adoption as lack of knowledge and resources and high entry costs of implementing new technologies.

A seminar highlight was awarding the winner of the second round of PRACE's Competition for the Most Innovative HPC Industrial Application. This year the prize was won by the French research institute CERFACS for developing a massively parallel CFD tool for Large Eddy Simulations of gas turbine engine combustors. The Second Prize was awarded to Termo Fluids of Spain and the Third Prize went to DigiCortex of Germany.

The parallel session "Technology Watch" provided an overview of the PRACE efforts in the technical evaluation of prototype computing platforms together with the port of necessary system software and programming environments

Figure 5: Catherine Riviere, Chair of the PRACE Council, congratulates the winner, Bijan Mohammadi of CERFACS

to those platforms. Another topic were Novel Programming Techniques including the development of auto-tuning runtime environments, the exploration of new scalable numerical algorithms and the development of technologies to help programmers to accelerate the process of application porting.

In order to demonstrate the concrete achievements of PRACE's industry cooperation four Success Stories were presented: Marco Stupazzini of Munich Re delivered insight into a cooperative research programme of the Politecnico di Milano/CINECA, Italy, and Munich Re in which the high performance code SPEED for simulating earthquakes is developed.

The Spanish SME Termo Fluids S.L. benefited from PRACE Tier-0 computational resources within the project "Broadening the scalability of Termo Fluids code" and was presented by Ricard Borrell.

Alessandro Chiarini, Metronic, gave a talk about the RT3S project involving Medtronic Italy, ANSYS, POLIMI and SCS-CINECA. The aim of this project is to develop safer stents by simulating the fatigue fracture of these small medical devices. The fourth Success Story was delivered by Pierre Spatz of Murex.

Furthermore, there were four special topic sessions with many stimulating short presentations. Marco Stenta, Syngenta, emphasized the advantages of using open-source software in industry. Yves Tourbier of Renault, Jan Blickwede of Volkswagen, Michael Hoffmann presenting Dallara and Albrecht Gehring, Lauer & Weiss, delivered insight into the state-of-the-art of using HPC in the automotive industry.

Norbert Kroll of the DLR, Yoon Ho of Rolls-Royce and Gregor Veble of Pipistrel represented the aerospace industry. And in "Emerging Applications" Claudio Arlandini and Eva Casoni presented applications, not yet ready for industrial production, which performance analysis, code development and optimization is supported by PRACE experts.

Next year's PRACE Industrial Seminar will take place in Barcelona on the 22nd and 23rd of May 2014, just a few weeks before the end of PRACE's 3rd Implementation Phase. At this 6th Industrial Seminar the results of the pilot project and the details of SHAPE will be presented.

Last but not least the seminar offered the attendees excellent opportunities to get informally acquainted. For the participants who arrived at the seminar on Sunday a guided tour to the Mercedes Benz engine plant in Stuttgart-Untertürkheim was scheduled on Monday morning.

After a busy first seminar day the attendees had the choice to visit the Museum Hauff of the Prehistoric World, the biggest private museum of its kind in Germany managed in the third generation as an example of private entrepreneurial spirit.

The social dinner in a well-known typical swabian and tweedy restaurant gave the attendees the chance to talk about the perceptions they had won during the seminar work.

About PRACE

The Partnership for Advanced Computing in Europe (PRACE) is an international non-profit association with its seat in Brussels. The PRACE Research Infrastructure provides a persistent world-class high performance computing service for scientists and researchers from academia and industry in Europe. The computer systems and their operations accessible through PRACE are provided by 4 PRACE members (BSC representing Spain, CINECA representing Italy, GCS representing Germany and GENCI representing France). For more information, see www.prace-ri.eu

References

[1] High-Performance Computing: Europe's place in a Global Race, Brussels, February 15, 2012

[2] http://www.prace-ri.eu/Industry-Access

[3] http://prace-ri.eu/PRACE-Industrial-Seminar-2013

• Bärbel Große-Wöhrmann
• Thomas Weitzel
• Thomas Bönisch

University of Stuttgart, HLRS

Workshop on Blue Gene Active Storage

The term active storage refers to computer architectures which comprise a storage sub-system that integrates significant processing power. The Jülich Supercomputing Centre plans in collaboration with IBM to realize such a concept for its Blue Gene/Q system JUQUEEN. As part of this project a workshop [1] was organized in January to bring computer architects and application developers together.

The architecture of I/O sub-systems is a particular challenge when progressing towards exascale machines. Already today it is difficult to maintain a reasonable balance between compute performance and performance of the I/O sub-system. In practice, this gap is widening and systems are moving away from Amdahl's rule of thumb for a balanced performance ratio, namely a bit of I/O per second for each instruction per second. Additionally, traditional disk based storage systems do not perform particularly well in case a large number of random I/O requests are performed.

Technology

Active storage is not a new concept (see, e.g., [2]). It has the potential of mitigating the mentioned problems as data processing is moved closer to the data. This approach therefore is additionally promising as it helps to reduce energy consumption due to data transport. Performance in terms of bandwidth and in particular I/O access rates as well as energy efficiency can additionally be improved by using non-volatile memory technologies like flash memory.

Fitch et al. [3] analysed already in 2010 the vision of an active storage concept based on the (at that time) emerging Blue Gene/Q architecture and solid state storage devices. The latter are integrated inside the compute racks and form one part of what becomes a tiered storage system. While the internal storage due to the high costs of suitable flash memory, i.e. SLC NAND flash, is limited in terms of capacity it provides high bandwidth and – compared to disk technology – very high I/O access rates. The second tier, an external storage system based on traditional technologies, i.e. disk, continues to be available to provide large storage capacity.

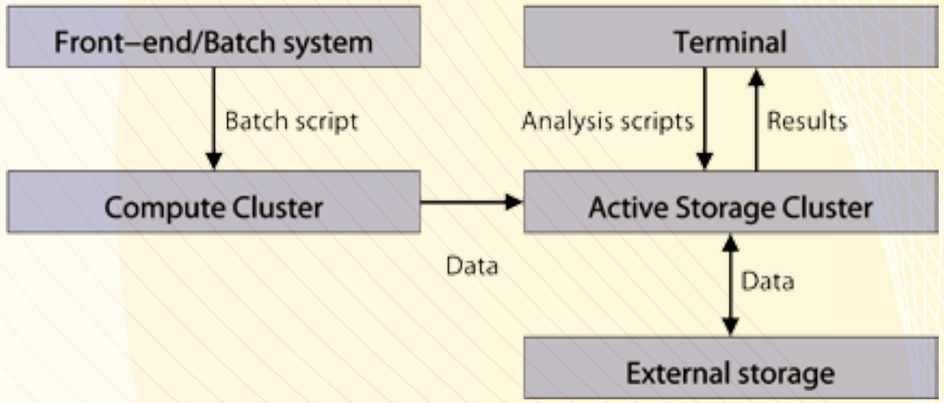

Figure 1. Work-flow of the neuronal network simulator NEST using active storage.

Such a tiered storage architecture comprising active storage can be used in different ways. For instance, in case of applications where the amount of data generated is too large to be written to external storage systems the processing capabilities of the active storage enables data post-processing such that the remaining amount, which needs to be written to disk, is significantly reduced. Other use cases are out-of-core computations, where main memory capacity limitations are mitigated by temporary swapping data to storage, or multi-pass analysis, where multi-terabyte data-sets are randomly accessed many times.

Applications

Neuronal network simulators are an example for applications which plan to exploit BGAS for data post-processing. Markus Diesmann (Forschungszentrum Jülich) presented the vision of the developers of the simulator NEST (see Fig. 1). Fast storage being available will allow to not only write (and later analyse) information about spike events, but also about membrane potentials and synaptic weights. This information would otherwise be much too large to be written. The size of the simulated networks is mainly limited by the available memory. Since non-volatile memory could provide additional memory space, future node architectures comprising large amounts of such memory would enable simulation of realistic neural tissue models as shown by James Kozloski (IBM). Such an out-of-core computing ansatz could also be used by other applications. Stefan Blügel and Paul Baumeister (Forschungszentrum Jülich) considered this approach for calculations based on density functional theory.

Multi-pass analysis is a use case which occurs in genetic epidemiology or radio astronomy. Paolo Bientinesi (RWTH Aachen) analysed the computational requirements of genome association studies, which are used to examine common genetic variants of different individuals to identify variants associated with a trait. For different regions of the genome measured data have to be processed many times. In a related field, genomics, active storage allows to deal with the exploding amount of data generated by next-generation sequencing methods. David Carrera (Barcelona

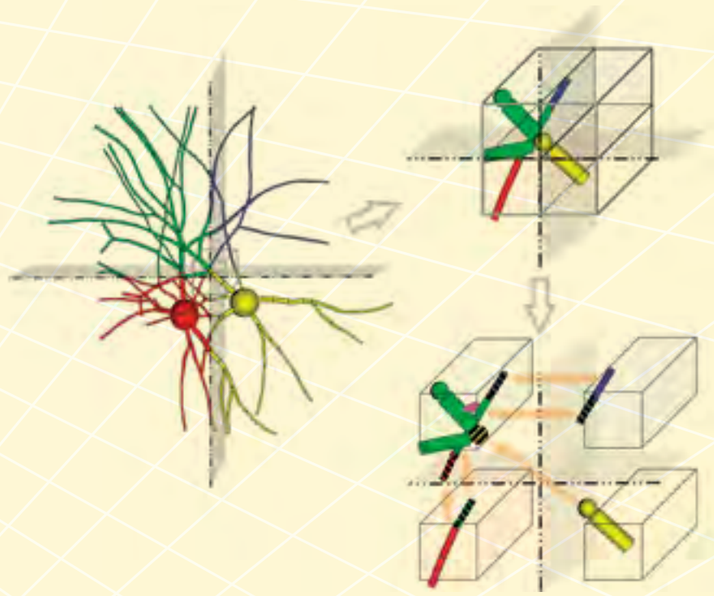

Figure 2. Brain tissue containing two neurons needs to be volume decomposed for parallel simulations. Non-volatile memory can help to keep these volumes large. (© 2011 Kozloski and Wagner, Frontiers Media)

• Dirk Pleiter
• Marcus Richter

Jülich Supercomputing Centre

Supercomputing Center) presented his Parallel In-Memory Database, where the active storage is used to implement a key-value store.

Active storage concepts could also be utilized by climate science applications, as pointed out by Nathanael Hübbe (University of Hamburg), to implement lossless compression and thus reduce the growing amount of data written to and read from large capacity, external storage systems.

Managing increasing data volumes is also a challenge for research in astronomy and radio astronomy. David Champion (MPI for Radio-Astronomy) explained how to search for pulsars in petabytes of data generated by planned surveys of the universe with high time resolution. Astronomy is traditionally data driven, as pointed out by Alex Szalay (Johns Hopkins University), who had been one of the architects of the archive of the Sloan Digital Sky Survey project. By making data accessible and by enabling any scientist to process this data, such archives turn into unique research instruments. Active storage concepts can help to bridge the gap between large capacity data services and HPC.

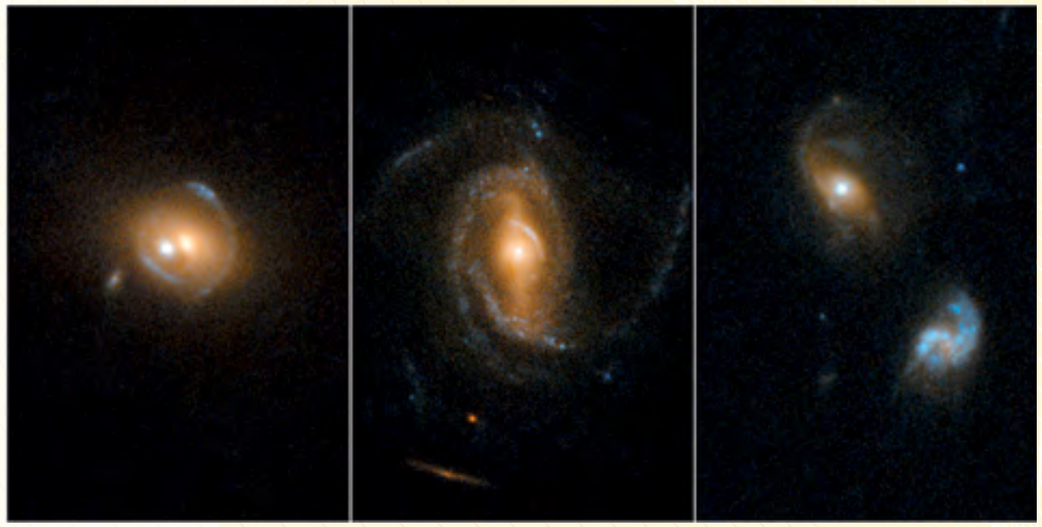

Figure 3. A galaxy-quasar combination found using the Sloan Digital Sky Survey (Credit: NASA, ESA/Hubble and F. Courbin (Ecole Polytechnique Federale de Lausanne, Switzerland)).

References

[1] http://www.fz-juelich.de/ias/jsc/EN/Expertise/Services/Documentation/presentations/presentation-bgas_table.html

[2] A. Acharya et al., "Active Disks: Programming Model, Algorithms and Evaluation," 5th conference on Computing frontiers, 2008.

[3] B. Fitch et al., "Blue Gene Active Storage," HEC FSIO Workshop, 2010.

International Workshop on Cooperative Quantum Dynamics and Its Control (CQDC2012)

• Kristel Michielsen

Jülich Supercomputing Centre

The "International Workshop on Cooperative Quantum Dynamics and Its Control (CQDC2012)" took place from 29 to 31 October 2012 at the rotunda of the Jülich Supercomputing Centre. The goal of the workshop was to discuss the possibility of essentially new quantum phenomena, and to discuss the observations and/or operations that would be required for their realization, detection, and understanding.

About 40 researchers from Germany, France, Switzerland, the Netherlands, Canada, the United States and Japan participated in the workshop. Recent developments in quantum annealing, open quantum systems, quantum computer hardware, quantum dynamics of spin systems, equilibration and thermalization of quantum systems, decoherence, entanglement and related topics were highlighted in talks and posters.

Quantum dynamics and novel quantum states in systems of many particles or spins, many of which have been beyond the reach of experimental realization, are now getting more and more realistic due to recent advancements in experimental techniques, such as the synthesis of molecular magnets, nano-engineering of quantum dots, time-resolved measurements with ultra-short pulses, and optical lattices of cold atoms.

In particular, many attempts have been made for characterizing the wave function from the viewpoint of entanglement, which is often accomplished by the aid of supercomputers through the method of direct numerical solution of the evolution equation and/or through quantum Monte Carlo methods. In addition, new numerical methods, such as the tensor-network variational approximation, are making many previously hard problems now tractable.

The exploitation of quantum effects, requiring coherent control of dissipative dynamics and entanglement control, is expected to have profound implications for future emerging information technologies. Managing and designing complex quantum systems with specified behaviour for quantum information processing requires a deep understanding of the cooperative behaviour of their components. Unraveling this behaviour necessitates an intensive collaboration between theoreticians and experimenters. The workshop successfully presented an overview of the current research on various topics very closely and less closely related to quantum information processing.

NIC Workshop: Hybrid Particle-Continuum Methods in Computational Materials Physics

From March 4 to 7, 2013, the NIC Workshop "Hybrid Particle-Continuum Methods in Computational Materials Physics" was held at Jülich Supercomputing Centre (JSC). It was organized jointly by the NIC research group "Computational Materials Physics" and the Institute of Advanced Simulation. The goal of the workshop was to foster the exchange of ideas between the communities working on complex fluids and on complex solids. 65 participants could learn about the new developments in hybrid particle-continuum methods from 15 invited and 10 contributed talks as well as from 30 posters.

During the workshop, particular emphasis was placed on continuum-mediated interactions between particles as well as on the adaptive and non-adaptive coupling between particle-based and continuum-based descriptions of materials. The subjects covered included the modeling of hydrodynamic interactions between particles in complex fluids or environments, through coarse-grained descriptions of biological systems, to the coupling of atomically represented regions with various continuum-based theories for fluids and solids. Special aspects were long time scale properties of systems with slow collective dynamics, the development of efficient adaptive resolution algorithms, and the coupling of quantum-mechanically treated regions with continuum descriptions.

The proceedings of the workshop were edited by Martin Müser, Godehard Sutmann and Roland Winkler (Forschungszentrum Jülich). The volume is available either as hard copy or as PDF file on the web http://juser.fz-juelich.de/record/132949/files/

• Martin Müser

Jülich Supercomputing Centre

NVIDIA Application Lab at Jülich

Accelerating scientific HPC applications using GPUs has become popular. For many applications it has indeed proved to be a very successful approach. Nevertheless, there is still a lot to learn about algorithms and methodologies for porting applications. Enabling more scientific applications for GPU-based architectures is a core goal of the NVIDIA Application Lab at Jülich. The lab is jointly operated by Jülich Supercomputing Centre (JSC) and NVIDIA since July 2012. For JSC it is yet another step to further strengthen links to vendors that are actively pursuing an exascale strategy.

During its first half year of operation the main focus of the lab was to establish a broad application portfolio encompassing computational neuroscience, high-energy physics, radio astronomy, data analytics and others. In common with each other, the applications have a high level of parallelism, ideally with few dependencies between tasks or task groups. One example is an application from the JuBrain project developed at the Jülich Institute for Neuroscience and Medicine INM-1 [1]. The project will result in an accurate, highly detailed computer model of the human brain. This atlas is created by reconstructing fibre tracks from pictures of thousands of slices (see Fig. 1). The process of mapping the pictures, called registration, requires repeated computation of a metric that measures how pixels of two pictures map to each other, a computationally expensive process that maps well to the GPU.

In applications from experimental physics, a natural data decomposition may be according to how the data leaves the detector. For instance, in high-energy physics experiments data is generated at such extremely high data rates that data for different time slices has to be distributed to different processing devices. Another example is search for pulsars in radio astronomy, where data sets from different beams and different measurements have to be repeatedly processed resulting in thousands or even millions of Fast Fourier Transforms.

Large quantities of data which need further processing are created not only by experiments and observatories, but increasingly from "computational experiments" such as Monte Carlo simulations of protein folding. Here clustering, a standard method of data analytics, is applied to identify regions of similar objects in multi-dimensional data sets. At the lab we have shown that sub-space clustering algorithms can be very efficiently implemented on GPUs opening the path to analysis of high-dimensional data sets in other areas as well [2].

For many applications, using a single GPU is not sufficient, either because more computing power is required, or because the problem size is too large to fit into the memory of a single device. This forces application developers to not only consider parallelization at device level, but also to manage an additional level of parallelism. Depending on the application this may be challeng-

ing since on most architectures device-to-device communication through the network is implemented as split transactions via the host processor. Software solutions like CUDA-aware MPI implementations can help mitigate this problem, but ultimately better hardware support is needed to interconnect GPU and network devices.

The goal of the lab is not solely to provide service to application developers and achieve performance improvements for their applications. To improve future architectures and their usability for scientific applications it is necessary to better understand how well these applications map onto such architectures. The introduction of a new architecture, Kepler, or more specifically, the GK110 GPU, is a good opportunity to learn about the effects of architectural changes by means of comparison with the previous architecture. How can application kernels cope with a significant increase in compute performance when bandwidth to device memory becomes only moderately larger? Do other changes in the memory hierarchy allow compensating for the increased flops-per-byte ratio? First experience shows that a careful analysis of the utilization of all levels of the memory hierarchy helps to make the application use the new architecture as efficiently as its predecessor. A better understanding at this level helps not only application developers but also processor and systems architects to improve GPU-based architectures for scientific computations.

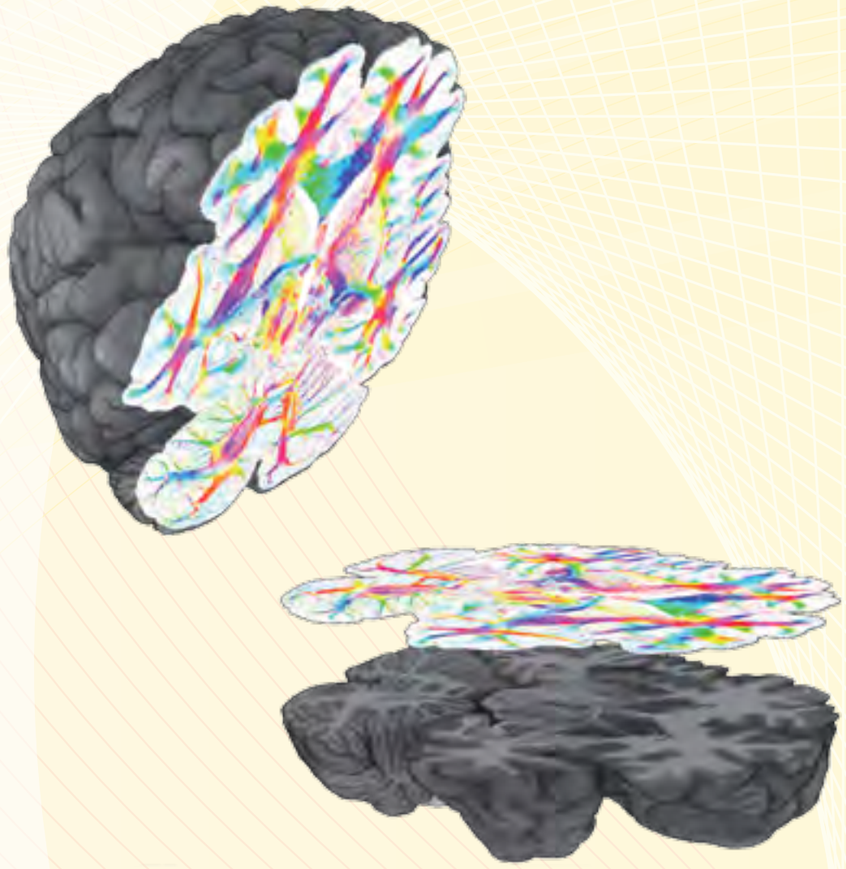

Figure 1: Many cuts of a human brain are required to build an accurate, highly detailed computer model of the human brain [1].

Acknowledgements

We would like to thank all application developers that have been involved in the lab so far, in particular: Markus Axer, Andreas Herten, Marcel Huysegoms, Jan Meinke (all Forschungszentrum Jülich) and David Champion (MPI for Radioastronomy).

References

[1] http://www.fz-juelich.de/JuBrain/EN

[2] **A. Adinetz, J. Kraus, J. Meinke, D. Pleiter,** „GPUMAFIA: Efficient Subspace Clustering with MAFIA on GPUs," submitted to Euro-Par 2013.

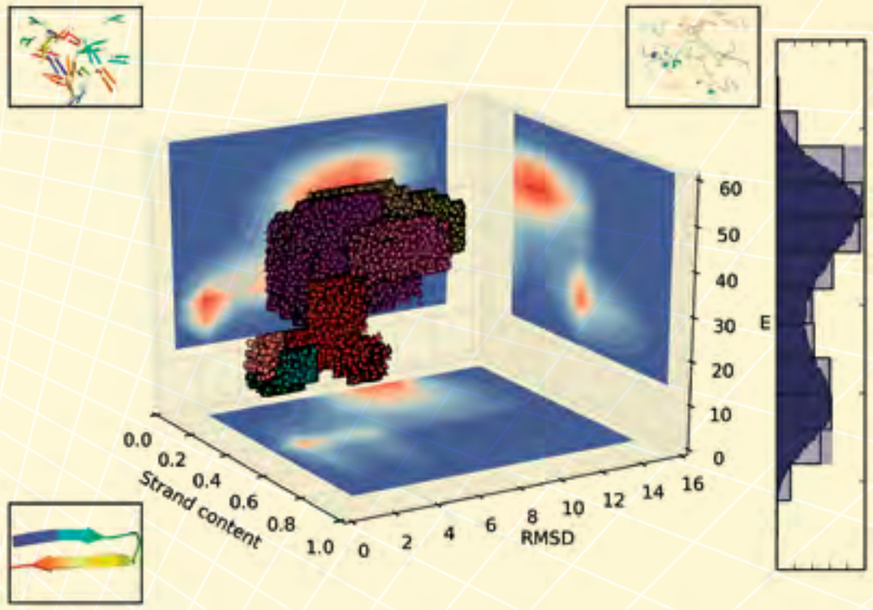

Figure 2: Clusters identified in data obtained by Monte Carlo simulations of protein folding using MAFIA [2].

• Andrew Adinetz
• Jiri Kraus
• Dirk Pleiter

Jülich Supercomputing Centre

Activities

Simulation Lab Neuroscience launched

How the human brain works is one of the biggest scientific questions of our time. A better understanding of the brain's structure and function, in health and disease, will lead to better diagnoses and treatments of brain diseases. It may also lead to revolutionary new brain-inspired computing technologies, given that the brain is able to compute, think, and learn in a fault-tolerant way using only the power of a conventional light bulb. To investigate the vast complexity of the 100 billion-neuron network in our brain, neuroscientists increasingly rely on supercomputer simulations. As the brain models developed by the scientists are becoming more and more realistic and complex in scale, more and more computing power and memory are needed to simulate them. Ultimately, the simulation of brain models at the scale of the full human brain amounts to an exascale challenge.

The Simulation Laboratory Neuroscience [1] at the Jülich Supercomputing Centre, which was officially launched in an opening event on 14-15 January 2013, will support the computational neuroscience community in leveraging high-end supercomputing resources such as JUQUEEN for their research. Like all other Simulation Labs in Jülich, the support mission of the Simulation Lab Neuroscience is grounded on its own research work, with a special focus on the methodological aspects. Under the supervision of Prof. Abigail Morrison about ten neuroscientists, computer scientists, mathematicians and physicists work together and with external groups on highly interdisciplinary tasks such as building large-scale models and databases of brain structure, function and connectivity, developing simulation, database and virtual reality technology, developing algorithms and workflows for data acquisition and analysis, and porting and optimizing scientific codes for supercomputers.

Figure 1: From left to right: Prof. Th. Lippert (Director of the Jülich Supercomputing Centre), Prof. A. Morrison (Head of the Simulation Lab Neuroscience, INM-6), Prof. K. Amunts (Director of the Institute of Neuroscience and Medicine, INM-1), Prof. S. Schmidt (Member of the Board of Directors of Forschungszentrum Jülich)

The Simulation Lab Neuroscience is a central element of the Helmholtz Portfolio Theme "Supercomputing and Modeling for the Human Brain" [2], through which it receives the majority of its funding. As part of the Portfolio Theme it will also contribute to the European FET Flagship "Human Brain Project" [3]. Additional funding for the Simulation Lab comes from the Jülich Aachen Research Alliance (JARA) [4]. A special relationship has been established with the German National Bernstein Network Computational Neuroscience [5]. The Bernstein Network is funded by the Federal Ministry of Education and Research (BMBF) and connects more than 200 research groups in Germany and beyond. The Simulation Lab Neuroscience contributes its expertise in simulation and database technology to the network as the "Bernstein Facility for Simulation and Database Technology".

In the opening event the Simulation Lab Neuroscience was officially introduced to the scientific community. The event was attended by representatives of the Bernstein Network, the Helmholtz Association, and JARA-HPC. It gave an overview of the existing and planned activities combining neuroscience and HPC in Jülich and beyond, and featured a series of neuroscientific talks by renowned international speakers from the field.

Figure 2: Official opening of the Simulation Lab Neuroscience at the Jülich Supercomputing Centre

References

[1] http://www.fz-juelich.de/ias/jsc/EN/Expertise/SimLab/slns/_node.html

[2] http://www.fz-juelich.de/JuBrain/EN/Helmholtz%20Portfolio.html

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

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

[5] http://www.nncn.de

• Boris Orth

Jülich Supercomputing Centre

12th HLRS/hww Workshop on Scalable Global Parallel File Systems

Representatives from science and industry working in the field of global parallel file systems and high performance storage solutions did meet at HLRS from March 18th to March 20th, 2013, for the twelfth annual HLRS/hww Workshop on Scalable Global Parallel File Systems. More than 90 participants did follow a total of 24 presentations that have been on the workshop agenda.

Thomas Bönisch, HLRS, opened the workshop with an opening address on Monday morning.

In the keynote talk, Brent Welch, CTO Panasas, discussed technology trends in the storage area. He explained current and future developments in the hardware area and their impact on file system and storage software developments. Brent Gorda, head of Intel High Performance Data Division and former CEO of Whamcloud explained the current status and future directions of the Lustre development. The following talk by Josh Judd, CTO Warp Mechanics, introduced a Lustre Solution using ZFS based OSS systems.

In the first presentation of the Monday afternoon session, Peter Braam, Xyratex, illuminated EIOW, a framework for exascale IO. This framework will tackle the storage issues of the exascale era by providing additional methods to pass more information about the IO from the application to the system level. Puneet Chaudhary, senior storage expert IBM, explained the GPFS Storage Server and current developments in GPFS. To complete the file system sessions, the current status and the roadmap of FhGFS was presented by Sven Breuner, FhGFS lead developer at Fraunhofer ITWM.

The last session on Monday was focusing on Storage Architectures. Didier Gava, Netapp, presented technologies to efficiently handle huge amounts of data on large scale storage systems. The talks on Cray storage solutions by Cray and Intelligence Driven Security by RSA/EMC concluded the first workshop day.

The first Tuesday morning session covered storage solutions using tape technology. Isabel Schwerdtfeger, IBM, introduced the High Performance Storage System HPSS. Ulrich Lechner, Grau Data provided information about the newly developed XtremeStore Storage Management Solution and Peter Schenke, IBM, covered the Linear Tape File System (LTFS) including use cases.

In the following session, Thomas Leibovici, CEA, provided a nice overview about open source tools for the operation of file systems at large scale. This covered several CEA developments in that field which eases the administration and usage of large file systems in many cases. Stefan Radtke, EMC, provided insight into the Isilon solution and its new support for HDFS.

The first session on Tuesday afternoon was dedicated to networking technologies. Rene Raeber, Cisco, covered topics around the datacenter network evolution. In the second presentation the Brocade view to the datacenter network convergence was provided by Christopher Feussner.

In the second afternoon session, big data topics and the architecture of Lustre storage solutions have been covered by DDN and Xyratex.

The Wednesday sessions have covered cloud topics and the Scalable IO Extensions project SIOX. Thomas Uhl, Grau Data, presented the newly developed cloud storage Data Space and Axel Tenschert, HLRS, showed the current developments at HLRS in the cloud field in his presentation named: Data Storage as a Service Using CDMI. In addition, Uwe Sauter, HLRS, provided information about the newly developed tool for flexible block IO over Infiniband.

Holger Mickler from the SIOX project gave an overview about the whole development and the current SIOX roadmap. Details of the SIOX architecture have been provided by Michaela Zimmer in the following talk. In the last presentation of the workshop, Xuan Wang, HLRS, gave insight into the IO optimization using the SIOX framework.

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

• Thomas Bönisch

University of Stuttgart, HLRS

HLRS Scientific Tutorials and Workshop Report and Outlook

• Rolf Rabenseifner
University of Stuttgart, HLRS

HLRS has installed **HERMIT,** a Cray XE6 system with AMD Interlagos processors and 1 PFlop/s peak performance and extended with an XC30 system. We strongly encourage you to port your applications to these architectures as early as possible. To support such effort we invite current and future users to participate in the special **Cray XE6/XC30 Optimization Workshops**. With these courses, we will give all necessary information to move applications to this Petaflop system. The Cray XE6 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. These 4-day courses in cooperation with Cray and multi-core optimization specialists are Oct. 28-31, 2013 and also in spring 2014.

ISC and SC Tutorials

Georg Hager, Gabriele Jost, Rolf Rabenseifner: **Hybrid Parallel Programming with MPI & OpenMP**. Tutorial 9 at the International Supercomputing Conference, ISC'13, Leipzig, June 16-20. 203.

Georg Hager, Jan Treibig, Gerhard Wellein: **Node-Level Performance Engineering.** Tutorial 2 at the International Supercomputing Conference, ISC'13, Leipzig, June 16-20. 203.

Rolf Rabenseifner, Georg Hager, Gabriele Jost: **Hybrid MPI and OpenMP Parallel Programming.** Half-day Tutorial at Super Computing 2012, SC12, Salt Lake City, USA, November 10-16, 2012.

Alice E. Koniges, Katherine Yelick, Rolf Rabenseifner, Reinhold Bader, David Eder: **Introduction to PGAS (UPC and CAF) and Hybrid for Multicore Programming**. Full-day Tutorial at Super Computing 2012, SC12, Salt Lake City, USA, November 10-16, 2012.

Programming of Cray XK7 clusters with GPUs is taught in **OpenACC Programming for Parallel Accelerated Supercomputers – an alternative to CUDA from Cray perspective** on Apr. 29-30, 2013.

These Cray XE6/XC30 and XK7 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 teaches 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 University of Siegen. It is again scheduled in September/October at the university of Siegen and in Spring 2014 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**, Sep. 23-27, 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. This includes MPI-2 functionality, e.g., parallel file I/O and hybrid MPI+OpenMP, as well as the upcoming MPI-3.0. 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 December 2013, mainly visited by PhD students from Stuttgart and other universities 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 and November 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 Supersmith, the HLRS also continues with contributions on hybrid MPI & OpenMP programming with **tutorials** at conferences; see the box

on the left page, which includes also a second tutorial with Georg Hager from RRZE.
In the table below, 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 ZIMT (Siegen).

2013 – Workshop Announcements
Scientific Conferences and Workshops at/by HLRS
11th HLRS/hww Workshop on Scalable Global Parallel File Systems (March 18-20)
7th ZIH+HLRS Parallel Tools Workshop (September 3-4 in Dresden)
High Performance Computing in Science and Engineering - The 16th Results and Review Workshop of the HPC Center Stuttgart (Sep. 30 – Oct. 1)
IDC International HPC User Forum (October 7-8)
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 18-22)
Iterative Linear Solvers and Parallelization (HLRS, March 11-15)
Cray XE6/XC30 Optimization Workshops (HLRS, April 16-19, and October 28-31)
GPU Programming using CUDA (HLRS, April 24-26, and October 21-23)
OpenACC Programming for Parallel Accelerated Supercomputers – an alternative to CUDA from Cray perspective (HLRS, April 29-30)
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 2-6)
Message Passing Interface (MPI) for Beginners (HLRS, September 23-24)
Shared Memory Parallelization with OpenMP (HLRS, September 25)
Advanced Topics in Parallel Programming (including MPI-3.0) (HLRS, September 26-27)
Introduction to Computational Fluid Dynamics (ZIMT Siegen, September/October)
Scientific Visualisation (HLRS, October 14-15)
Industrial Services of the National HPC Centre Stuttgart (HLRS, October 16 and 24)
Parallel Programming with MPI & OpenMP (FZ Jülich, JSC, November 25-27)
Training in Programming Languages at HLRS
Fortran for Scientific Computing (March 4-8 and December 2-6)
URLs: www.hlrs.de/events/
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 SuperMUC - the new Petaflop Supercomputer at LRZ

Date & Location:
July 08 - 11, 2013
LRZ, Garching near Munich

Contents:
This four-day workshop gives an introduction to the usage of SuperMUC, the new Petaflop class Supercomputer at LRZ. The first three days are dedicated to presentations by Intel on their software development stack (compilers, tools and libraries); the remaining day will be comprised of talks and exercises delivered by IBM and LRZ on usage of the IBM-specific aspects of the new system (IBM MPI, LoadLeveler, HPC Toolkit) and recommendations on tuning and optimizing for the new system.

Prerequisites:
Participants should have good knowledge of HPC-related programming, in particular MPI, OpenMP and at least one of the languages C, C++ or Fortran.

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

Parallel Programming with MPI, OpenMP and PETSc

Date & Location:
July 29 - 31, 2013
TU Hamburg-Harburg

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

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

Introduction to Parallel Programming with MPI and OpenMP

Date & Location:
August 06 - 09, 2013
JSC, Forschungszentrum Jülich

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

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

Iterative Linear Solvers and Parallelization

Date & Location:
September 02 - 06, 2013
LRZ Building, University campus Garching, near Munich, Boltzmannstr. 1

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

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

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

Advanced Fortran Topics

Date & Location:
September 16 - 20, 2013
LRZ, Garching near Munich

Contents:
This course is targeted at scientists who wish to extend their knowledge of Fortran beyond what is provided in the Fortran 95 standard. Some other tools relevant for software engineering are also discussed. Topics covered include

- object oriented features
- design patterns
- generation and handling of shared libraries
- mixed language programming
- standardized IEEE arithmetic and exceptions
- I/O extensions from Fortran 2003
- parallel programming with coarrays
- source code versioning system (subversion)

To consolidate the lecture material, each day's approximately 4 hours of lectures are complemented by 3 hours of hands-on sessions.

Prerequisites:
Course participants should have basic UNIX/Linux knowledge (login with secure shell, shell commands, simple scripts, editor vi or emacs). Good knowledge of the Fortran 95 standard is also necessary, such as covered in the February course at LRZ.

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

Message Passing Interface (MPI) for Beginners

Date & Location:
September 23 - 24, 2013
Stuttgart, HLRS

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

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

Shared Memory Parallelization with OpenMP

Date & Location:
September 25, 2013
Stuttgart, HLRS

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

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

Advanced Topics in Parallel Programming

Date & Location:
September 26 - 27, 2013
Stuttgart, HLRS

Contents:
Topics are MPI-2 parallel file I/O, hybrid mixed model MPI+OpenMP parallelization, MPI -3.0 parallelization of explicit and implicit solvers and of particle based applications, parallel numerics and libraries, and parallelization with PETSc. Hands-on sessions are included. Course language is English (if required).

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

GCS – High Performance Computing

Introduction to Computational Fluids Dynamics

Date & Location:
September/October, 2013
University of Siegen

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 ZIMT University of Siegen.

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

Scientific Visualisation

Date & Location:
October 14 - 15, 2013
Stuttgart, HLRS

Contents:
This two day course is targeted at researchers with basic knowledge in numerical simulation, who would like to learn how to visualize their simulation results on the desktop but also in Augmented Reality and Virtual Environments. It will start with a short overview of scientific visualization, following a hands-on introduction to 3D desktop visualization with COVISE. On the second day, we will discuss how to build interactive 3D Models for Virtual Environments and how to set up an Augmented Reality visualization.

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

Industrial Services of the National HPC Centre Stuttgart (HLRS)

Dates & Location:
October 16, 2013, and
October 24, 2013
Stuttgart, HLRS

Contents:
In order to permanently assure their competitiveness, enterprises and institutions are increasingly forced to deliver highest performance. Powerful computers, among the best in the world, can reliably support them in doing so.
This course is targeted towards decision makers in companies that would like to learn more about the advantages of using high performance computers in their field of business. They will be given extensive information about the properties and the capabilities of the computers as well as access methods and security aspects. In addition we present our comprehensive service offering - ranging from individual consulting via training courses to visualization. Real world examples will finally allow an interesting insight into our current activities..

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

Cray XE6/XC30 Optimization Workshop

Date & Location:
October 28 - 31, 2013
Stuttgart, HLRS

Contents:
HLRS installed HERMIT, a Cray XE6 system with AMD Interlagos processors and a performance of 1 PFlop/s. Currently, the system is extended by a Cray XC30 system. We invite current and future users to participate in this special course on porting applications to our Cray architectures. The Cray XE6 and Cray XC30 will provide our users with a new level of performance. To harvest this potential will require all our efforts. We are looking forward to working with you on these opportunities.

Authors

Andrew Adinetz
a.adinets@fz-juelich.de
Momme Allalen
allalen@lrz.de
Rossen Apostolov
rossen@kth.se
Lilit Axner
lilit@kth.se
Christoph Bernau
bernau@lrz.de
Joachim Biercamp
biercamp@dkrz.de
Kurt Binder
kbinder@uni-mainz.de
Thomas Bönisch
boenisch@hlrs.de
Doris Breuer
doris.breuer@dlr.de
Dirk Brömmel
d.broemmel@fz-juelich.de
Alexey Cheptsov
cheptsov@hlrs.de
Mihai Duta
mihai.duta@oerc.ox.ac.uk
Wolfgang Eckhardt
eckhardw@in.tum.de
Agusti Emperador
agusti.emperador@
irbbarcelona.org
Josep Lluis Gelpi
jlgelpi@gmail.com
Michael Gienger
gienger@hlrs.de
Judit Gimenez
judit@bsc.es
Gerhard Gompper
g.gompper@fz-juelich.de
Ramon Goni
ramon.goni@bsc.es
Gordon Groskopf
groskopf@iag.uni-stuttgart.de
Alexander Heinecke
heinecke@in.tum.de
Herlina Herlina
Herlina.Herlina@kit.edu
Berk Hess
hess@kth.se
Lars Hoffmann
l.hoffmann@fz-juelich.de
Sabine Höfler-Thierfeldt
s.hoefler-thierfeldt@fz-juelich.de
Ferdinand Jamitzky
jamitzky@lrz.de
Markus J. Kloker
kloker@iag.uni-stuttgart.de
Bastian Koller
koller@hlrs.de
Christina Korntreff
c.korntreff@fz-juelich.de
Jiri Kraus
jkraus@nvidia.com
Boris Kubrak
Boris.Kubrak@brunel.ac.uk
Dieter Langosch
biopolymere@lrz.tum.de
Erwin Laure
erwinl@pdc.kth.se
Catrin Lehmann
c.lehmann@fz-juelich.de
Kristel Michielsen
k. michielsen@fz-juelich.de
Matthias Mörch
matthias.moerch@mytum.de
Martin Müser
m.mueser@fz-juelich.de
Walter Nadler
w.nadler@fz-juelich.de
Guust Nolet
nolet@geoazur.unice.fr
Modesto Orozco
modesto.orozco@irbbarcelona.org
Boris Orth
b.orth@fz-juelich.de
Wolfgang Paul
Wolfgang.Paul@physik.uni-halle.de
Susanne Pfalzner
spfalzner@mpifr-bonn.mpg.de
Dirk Pleiter
d.pleiter@fz-juelich.de
Ana-Catalina Plesa
ana.plesa@dlr.de
Rolf Rabenseifner
rabenseifner@hlrs.de
Florian Rauser
florian.rauser@zmaw.de
Michael M. Resch
resch@hlrs.de
Marcus Richter
m.richter@fz-juelich.de
Zilvinas Rinkevicius
rinkevic@TheoChem.kth.se
Helmut Satzger
satzger@lrz.de
Christina Scharnagl
christina.scharnagl@tum.de
Wolfgang Schotte
schotte@hlrs.de
Bernhard Schuberth
mail@bernhard-schuberth.de
Mathieu Solar
Mathieu.Solar@physik.uni-halle.de
Albert Solernou
albert.solernou@oerc.ox.ac.uk
Tilman Spohn
tilman.spohn@dlr.de
Björn Stevens
bjorn.stevens@mpimet.mpg.de
Axel Tenschert
tenschert@hlrs.de
Peter Virnau
Virnau@uni-mainz.de
Regina Weigand
weigand@hlrs.de
Alexander Winkler
winkalex@uni-mainz.de
Roland G. Winkler
r.winkler@fz-juelich.de
Jan Wissink
Jan.Wissink@brunel.ac.uk
Leonid Yelash
Yelash@uni-mainz.de
Christophe Zaroli
c.zaroli@unistra.fr

If you would like to receive inside regularly, send an email with your postal address to klank@hlrs.de
Web page: http://inside.hlrs.de/