inSiDE

inSiDE • Vol. 7 No. 2 • Autumn 2009

Innovatives Supercomputing in Deutschland

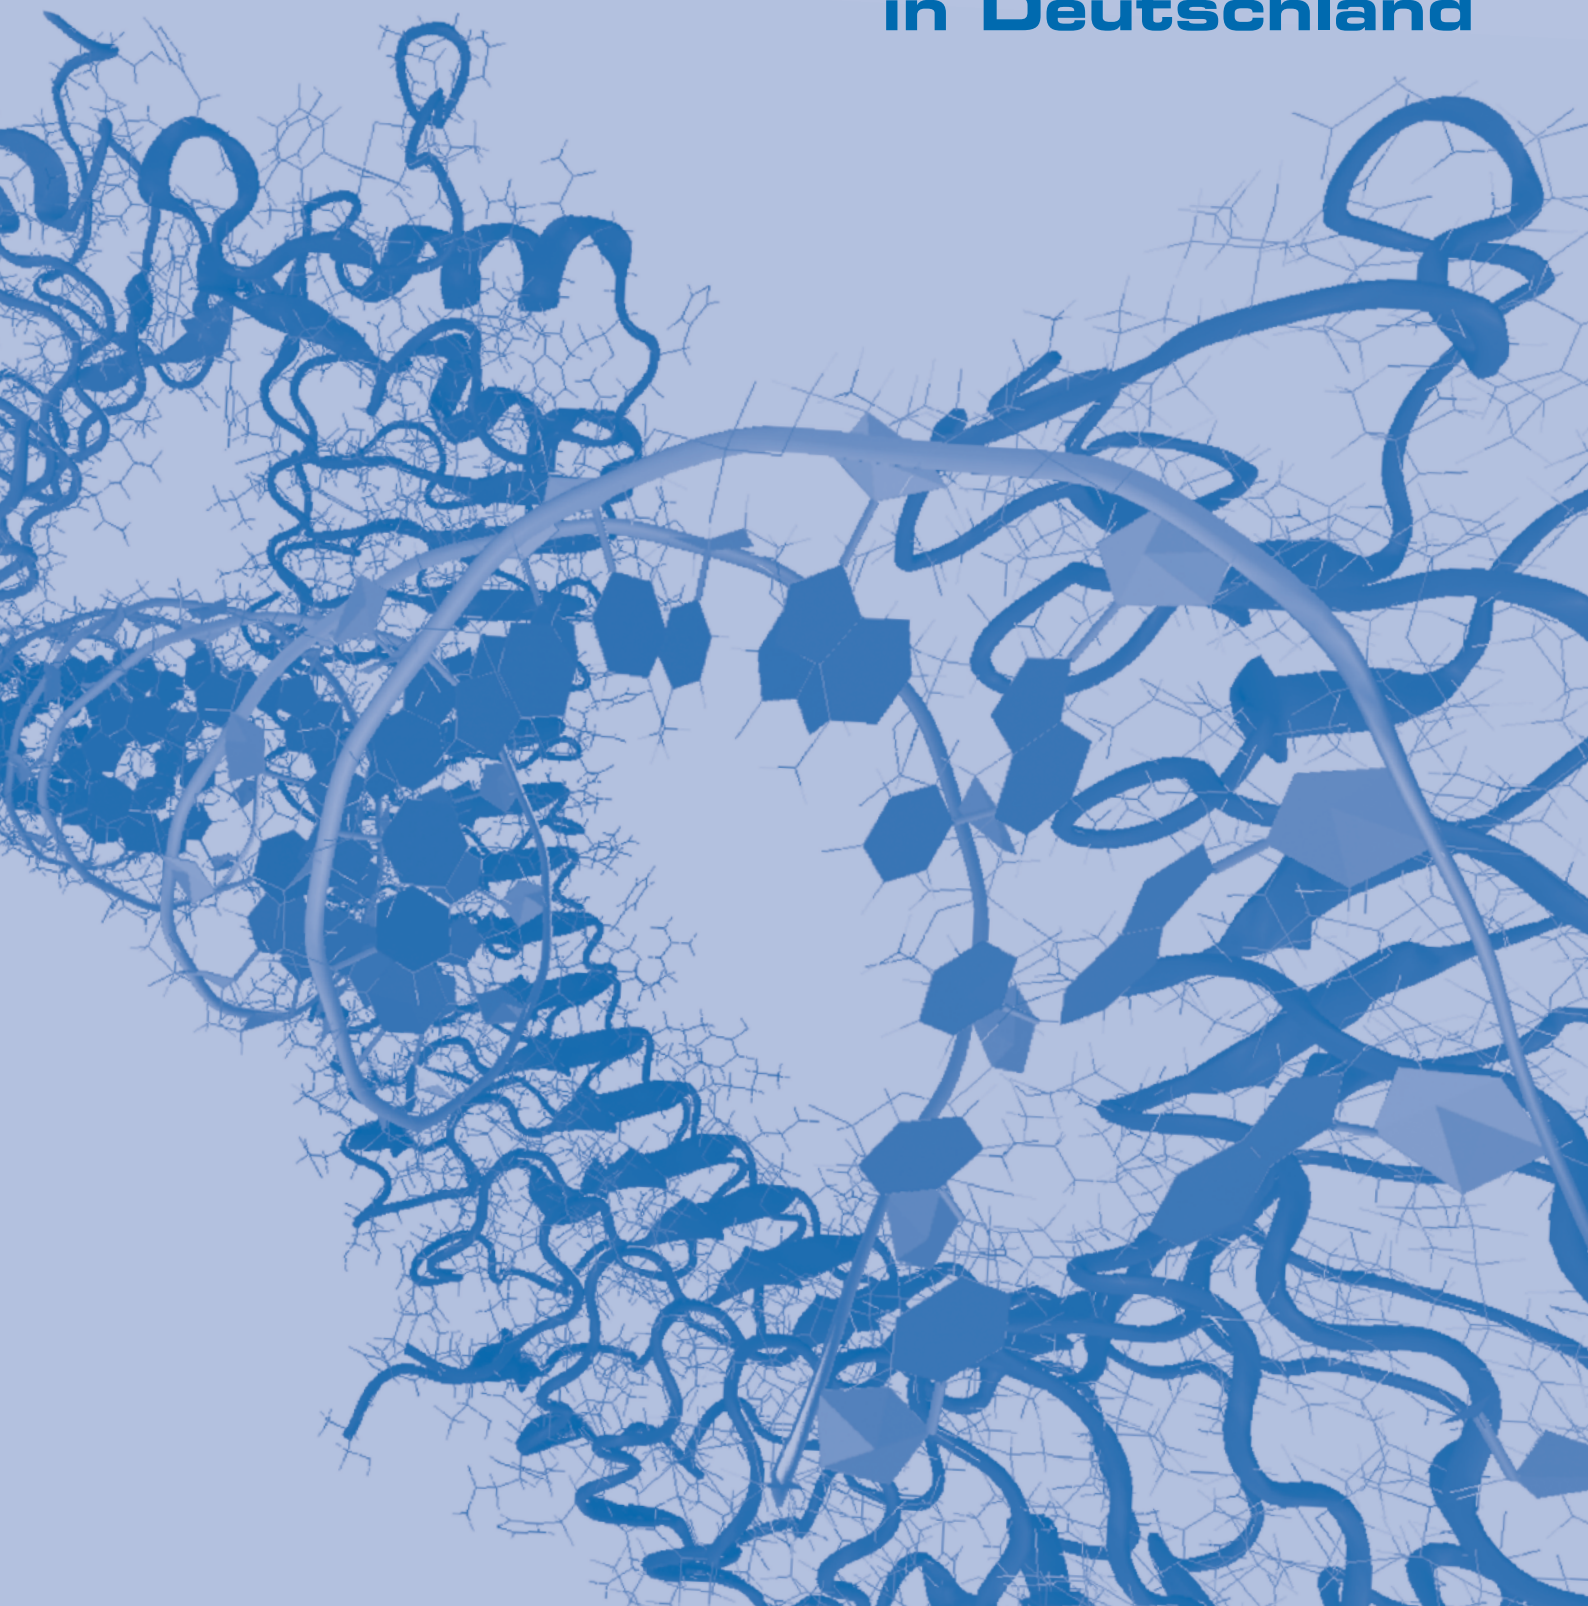

Editorial

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

Welcome to this new issue of InSiDE the journal on innovative Supercomputing in Germany. Although it is only a few months since the last issue German supercomputing has made a great leap forward since.

Three supercomputers for European research were unveiled on May 26, 2009 in Jülich in a ceremony attended by the Federal Minister for Education and Research, Prof. Dr. Annette Schavan, and the Prime Minister of North Rhine-Westphalia, Dr. Jürgen Rüttgers. The three systems include the supercomputer JUGENE, which with a computing power of one petaflop/s is currently the fastest computer in Europe.

This is the first major step to implement the strategy of the Gauss Center for Supercomputing. The High-Performance Computing Center Stuttgart (HLRS) will follow up with a system for which the procurement is currently running. Leibniz Computing Center (LRZ) has launched an infrastructure initiative to prepare for their next step expected to be operational in 2011. Germany will hence be able to launch a new system every year for the next five years to come and will so emphasize its leading role in High-Performance Computing worldwide.

Hardware is only one important building block for High-Performance Computing and simulation. Software and applications are the second important part. Therefore all the three Gauss partners have established strong software programs. In this issue we highlight the Munich Centre of Advanced Computing (MAC). In MAC, seven departments (mathematics, informatics, physics, chemistry, mechanical engineering, civil engineering, and electrical engineering & information technology) of the Technical University of Munich and other Munich research institutions (Ludwig Maximilians Universität, Max-Planck institutes, and the Leibniz Supercomputing Centre (LRZ)) as well as TUM's new international partner KAUST, the King Abdullah University of Science and Technology, join their forces to ensure the sustainable usage of current and future HPC architectures for most relevant and most challenging CSE applications.

In this issue of inSiDE you will find again a comprehensive section about large scale applications. Five contributions from users all over Germany present the variety of simulations that are currently run on German HPC systems. At the same time these papers give a good insight into the future needs for supercomputing in Germany.

The project part in reflects the variety of critical problems that we face in the era of petaflops computing and beyond. S(o)OS (Service oriented Operating Systems) is a European funded project and addresses the problems of executing processes in large scale resources. Another article covers the key aspects of performance analysis in large scale systems. An application supporting framework is presented in an article that relates both to applications and tools. The framework waLBerla is a software package specifically designed for large scale Lattice Boltzmann Method simulations for various CFD applications.

Contents

Inauguration and Start-up: Europe's Fastest Supercomputers

Three supercomputers for European research were unveiled on 26 May 2009 in Jülich in a ceremony attended by the Federal Minister for Education and Research, Prof. Dr. Annette Schavan, and the Prime Minister of North Rhine-Westphalia, Dr. Jürgen Rüttgers. The three systems include the supercomputer JUGENE, which with a computing power of one petaflop/s is currently the fastest computer in Europe.

More than 250 guests attended the inauguration ceremony. At the beginning, it soon became clear that this event is not only of importance for Jülich. All three partners in the Gauss Centre for Supercomputing GCS – formed by the national supercomputing centres in Stuttgart, Garching and Jülich – were proud to present the new petaflop system to the public. JUGENE is the first German supercomputer system to be selected and purchased in the context of the GCS. Therefore, the welcoming address was given jointly by the GCS partners, represented by Prof. Dr. Achim Bachem, Chairman of the Board of Directors of Forschungszentrum Jülich, Prof. Dr. Arndt Bode, Chairman of the Board of Directors of the Leibniz Rechenzentrum in Garching, and Prof. Dr. Michael Resch, Director of the High-Performance Computing Center Stuttgart.

Start-up for the supercomputers. From left to right: Jürgen Rüttgers (Prime Minister of North Rhine-Westphalia), Achim Bachem (Chairman of the Board of Directors, Forschungszentrum Jülich), Annette Schavan (Federal Minister for Education and Research)

"The supercomputer JUGENE will secure Europe independent access to a decisive key technology of the 21st century," said Bachem, who is also Coordinator of PRACE, the European Supercomputing Alliance. PRACE is funded by the EU and will coordinate the creation of a Europe-wide computer infrastructure. Bachem expressed special thanks to the Federal Government and the state of North Rhine-Westphalia for the many years of support and funding. "The supercomputers at Forschungszentrum Jülich show that North Rhine-Westphalia is already one of the leaders in strategically important sectors," said Prime Minister Rüttgers. "We want to make North Rhine-Westphalia the top state for innovation – and Forschungszentrum Jülich with its excellent research will make a major contribution to achieving this goal."

In a round table discussion, representatives of the companies Bull, Sun, Intel and ParTec gave their views on the unique collaboration established in building the supercomputers JUROPA and HPC-FF. JSC director Prof. Dr. Dr. Thomas Lippert expressed the view that in building a computer from best-of-breed parts, one has to combine the heart and soul of the computer: the processors form the heart, the software represents the soul. Yvan Capouet of the European Commission pointed out that great hopes are placed in the new HPC-FF system, which will be exclusively available to European fusion researchers. A short time-lapse video showed the building of the supercomputers in the JSC machine room. In the second discussion, representatives of PRACE, the European Commission and IBM emphasized the importance of JUGENE for European researchers. Finally, Minister Schavan stressed the importance of supercomputing for Germany. "This is a good day for the German Gauss Centre and a good day for Europe as well. Acquiring JUGENE demonstrates Germany's bid for leadership in supercomputing," she said. "This day confirms our strategy that the GCS as a model of partnership in supercomputing is right for Germany." The supercomputers were officially launched by Schavan, Rüttgers and Bachem, who together pushed a lever for the start-up. After that, they visited the machine room of the Jülich Supercomputing Centre.

JUGENE went into production on 1 July 2009, just in time for the new computing-time period in Jülich. In total, there were 75 project proposals for JUGENE requesting more than four times the available computing time. The proposals were reviewed and selected by the NIC Peer Review Board on behalf of the GCS Management Board. The large number of project proposals show that JUGENE has turned out to be a very popular supercomputer. JUGENE was also the first computer available for the first call for large-scale projects by the Gauss Centre for Supercomputing in April, which attracted ten project submissions.

Projects are classified as "large-scale" if they require more than 5% of the available CPU cycles. In the case of JUGENE, this is at least 20 rack months. Two projects were awarded the status of large-scale project: one from the field of fluid dynamics, "Geometrical Properties of Small-Scale Turbulence", by Prof. Dr. Norbert Peters (RWTH Aachen University) with 24 rack months, and one from elementary particle physics, "QCD Simulations with Light, Strange and Charm Dynamical Quarks", by Dr. Karl Jansen (DESY Zeuthen) with 20 rack months.

Among all proposals for the supercomputers in Jülich – JUGENE as well as JUROPA – the NIC Peer Review Board decided to recognize two outstanding projects by designating them NIC Excellence Projects 2009/2010. The projects receiving this distinction are from the field of elementary particle physics, "Lattice QCD with 2 plus 1 flavours at the physical mass point", submitted by Prof. Dr. Zoltan Fodor (University of Wuppertal), and from astrophysics, "The small-scale structure of the universe", submitted by Dr. Stefan Gottlöber (Astrophysical Institute Potsdam).

In front of the first European petaflop system JUGENE in the machine room in Jülich. From left to right: Martin Jetter (Managing Director of IBM Germany), Achim Bachem (Chairman of the Board of Directors, Forschungszentrum Jülich), Annette Schavan (Federal Minister for Education and Research), Jürgen Rüttgers (Prime Minister of North Rhine-Westphalia), Thomas Lippert (Director of Jülich Supercomputing Centre)

• Sabine Höfler-Thierfeldt
• Walter Nadler

Jülich Supercomputing Centre

Munich Centre of Advanced Computing Opened

The Munich Centre of Advanced Computing (MAC) is a new research consortium established at Technische Universität München (TUM) to bundle research activities related to Computational Science and Engineering (CSE) and High-Performance Computing (HPC) – across disciplines, across departments, and across institutions. In MAC, seven TUM departments (mathematics, informatics, physics, chemistry, mechanical engineering, civil engineering, and electrical engineering & information technology) and other Munich research institutions (Ludwig-Maximilians-Universität, Max-Planck institutes, and the Leibniz Supercomputing Centre (LRZ)) as well as TUM's new international partner KAUST, the King Abdullah University of Science and Technology, join their forces to ensure the sustainable usage of current and future HPC architectures for most relevant and most challenging CSE applications.

On July 2, MAC's official opening took place at LRZ, in the presence of Prof. Wolfgang A. Herrmann, president of TUM, Prof. Choon Fong Shih, president of KAUST, various representatives of the participating institutions, and the many scientists and Ph.D. students involved in MAC. The keynote presentation was given by Prof. David Keyes, the new chair of the mathematical and computer sciences and engineering division at KAUST.

Advanced Computing

CSE has become a highly interdisciplinary, challenging, and thriving scientific domain of increasing importance and visibility in research, development, and education. Its essence primarily involves modelling and computations – in particular (numerical) simulations – on computer systems ranging from standard PC to supercomputers. Thus, CSE has always been a driving force of computing technologies, in particular HPC.

However, an increasing gap is observed world-wide between what could be possible due to recent advancements in algorithms, hardware, and networks, and what is really achieved in practice. While computing power and considerations on complexity and accuracy of algorithms were in the focus over the last decades, new limiting factors are now encountered: the need for hardware awareness and the emerging ubiquitous parallelism due to many-core systems; the energy issue; the amount of simulation data, their mere handling, their analysis, and, thus, the extraction and representation of relevant information – often in an interactive way; the increasingly complex process of design, engineering, and maintenance of HPC software, ranging from specialized single-application codes to general problem solving environments; the need of this process for language and tool support and for a substantial professionalization.

Research Concept and Goals

Compared to respective activities existing or emerging world-wide, MAC's main idea, and also its key differentiating aspect, is to address all major aspects of Advanced Computing in a concerted effort: foundations and applications; science and engineering; enabling technologies – some with a longer tradition in CSE (visualization, e.g.), some with rather fresh (data exploration, e.g.) or basically even without relations so far, such as systematic software engineering – and HPC; mathematics and informatics; supercomputers, commodity systems, and special hardware; simulation and software engineering. This provides a fascinating perspective, also at an international scale.

Three overall scientific goals will drive research within MAC: first, advancing the frontier of computing; second, creating synergies by trans-disciplinary research; third, advancing the frontier of each scientific domain in MAC. MAC's main structural goals are to establish and consolidate the Garching campus as the core of a leading centre for Advanced Computing and HPC, and to contribute to improve Germany's (and, thus, the Gauss Centre's) international standing in the field.

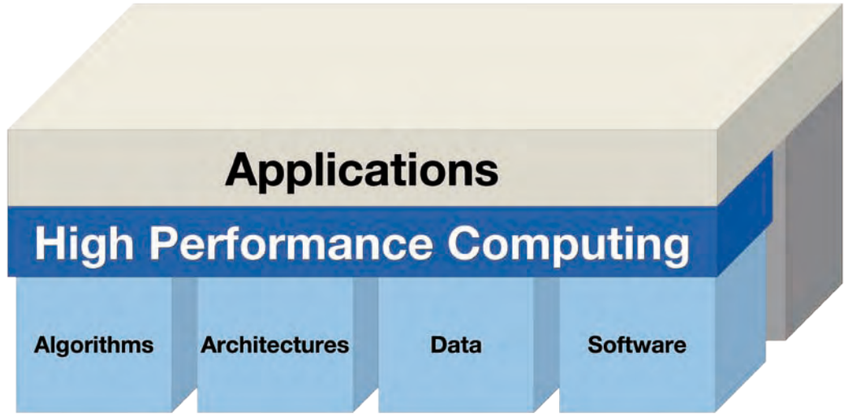

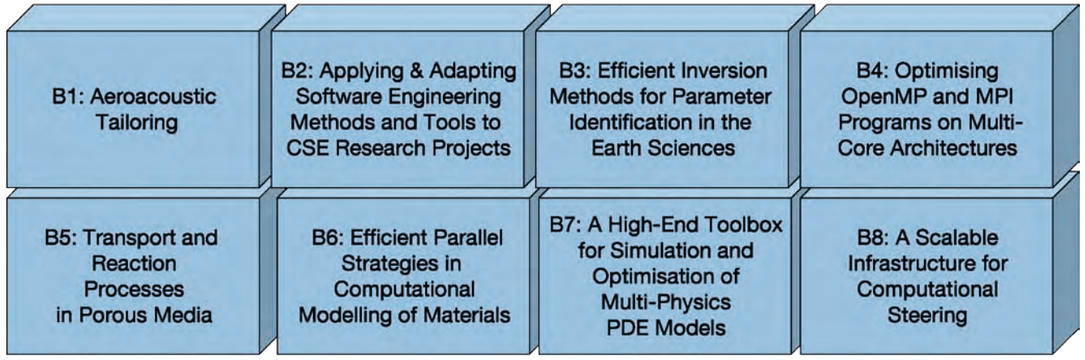

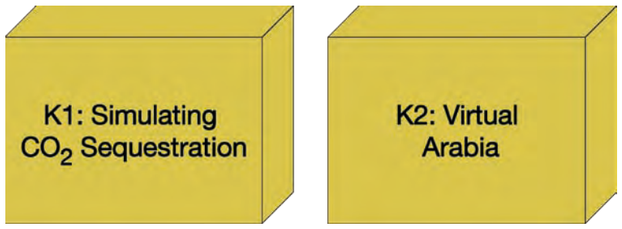

Organization

MAC was designed as an independent research consortium at TUM, kind of a common roof for all computing-related activities of groups from TUM as well as from partner institutions. So far, MAC's funding has been based on two pillars: eight interdisciplinary projects co-financed by the State of Bavaria and by TUM, and two large collaborative projects established in the framework of the strategic partnership of TUM and KAUST:

At present, 22 professors with their groups are participating in MAC. Via the different funding channels, 31 new Ph.D. scholarships as well as 13 new research assistant positions could be created, altogether summing up to a funding volume of about 3 M € per year for the next four years. MAC is led by a board (Profs. Arndt Bode, Hans-Joachim Bungartz (head), and Ernst Rank) and by the coordinator, Prof. Michael Bader.

Research in MAC – Two Examples

In project B2, the partners – Profs. Bernd Brügge, Alfons Kemper, Rüdiger Westermann (all TUM, informatics), Hans-Peter Bunge (LMU, geophysics), and Siegfried Bethke (Max-Planck Institut für Physik) apply software engineering methods and tools to the field of CSE and use the gained experiences to continuously evaluate and improve them. In particular, the applications of a uniform project model will be in the focus. The tool-supported approach of this project model can externalize all artefacts relevant for a project in one central repository. These can be requirements, analysis and design objects, but also project-management related elements such as team members or individual tasks. Also explicit and implicit relations between these artefacts are captured, which enables traceability.

For example, this allows to trace the requirements of an algorithm to the run time complexity of its implementation. As case studies, three CSE research projects will be investigated: MPP's ATLAS participation as an existing large and globally distributed project, Computational Steering as a mid-size, emerging project, and SEISSOL as an existing small project.

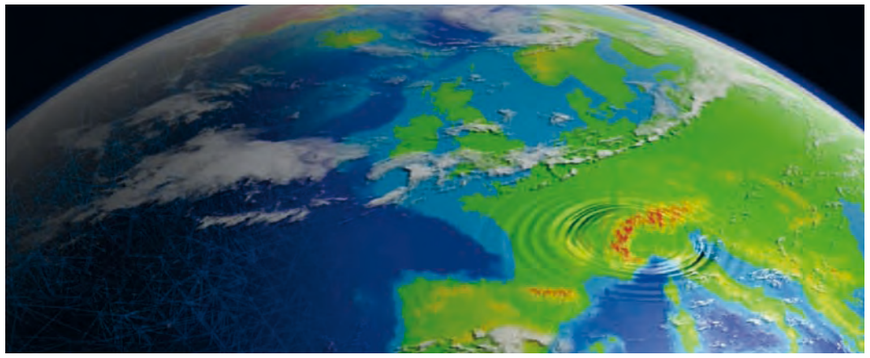

Project B2, result with SEISSOL: Visualization of the propagation of seismic waves after an earthquake event in Central Europe after 100 seconds of propagation time. The faster P-waves have relatively low amplitudes compared to the slower S- and surface wave, which typically have higher damage potential.

Project K1 brings together partners from TUM's math department (Profs. Martin Brokate, Karl-Heinz Hoffmann, and Michael Ulbrich), TUM's informatics department (Prof. Hans-Joachim Bungartz), TUM's civil engineering department (Prof. Michael Manhart), and the LRZ (Prof. Arndt Bode). The goal is to design and investigate novel approaches to modelling and simulation of CO_2 sequestration processes, in particular in the context of enhanced oil recovery.

The project will involve both fine-grain simulations – with all related aspects from multi-phase schemes via numerical algorithmics to high-performance computing issues – and homogenization approaches to efficiently capture the fine-grain effects on the macro-scale. For that, groups with expertise in flow physics, mathematical modelling, numerical analysis, numerical algorithmics, optimization and inverse problems, and High-Performance Computing and HPC systems join their forces. Topics addressed will cover multi-scale modelling and homogenization, fully-resolved pore-scale simulation, constrained optimization of the sequestration process, enhanced numerics and parallelization, and HPC implementation. We expect close relations to various units on KAUST campus in Jeddah (3D Modelling and Visualization Centre, Computational Earth Sciences Centre, HPC facilities with "Shaheen", the current #14 in the Top500 list).

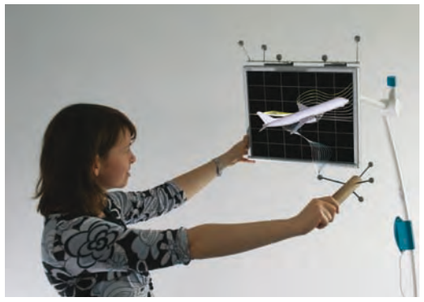

Projects B1, B8, and K2: Interactive Augmented Reality at work

Further Steps

MAC and HPC research in Munich continue to grow. A new cross-institutional research group consisting of one professorship and eight research assistants jointly funded by TUM, LMU, LRZ, and MPG will be established, focussing on software aspects of HPC – from the more established issues of monitoring, tuning, and porting via high-end algorithmics to widely neglected issues of software engineering, software management, as well as validation and verification. This group will complement research in MAC and foster links between LRZ and the various university groups involved in Advanced Computing research in Munich.

For details see www.mac.tum.de

• Hans-Joachim Bungartz
• Michael Bader

Faculty of Informatics, TU München

LRZ: Towards the Next Petaflop System of GCS

Alea iacta est! In July, the federal Bavarian parliament paved the way towards the next Petaflop supercomputer for the Munich branch of the Gauss Centre by committing to finance this system and to extend LRZ's buildings.

At least 135 million Euros will be spent by the Bavarian and the German governments for LRZ's next supercomputer, named "SuperMUC", to double the usable space in the computer building, to extend the power and cooling facilities to several Megawatts, to offer more office space for staff and to install a cave for three-dimensional visualization. Not only will the next supercomputer demand more space than the current system but also the extension of services offered by LRZ to the scientists in Munich and Bavaria.

49.2 million Euros are available for the new buildings. Another 84 million Euros can be spent on SuperMUC in a first installation phase. An additional 48 million Euros will probably follow in a second installation phase of SuperMUC.

In 2011, LRZ intends to replace its present flagship system, its 62-TFlop SGI Altix 4700 by the successor SuperMUC whose computational capacity will be of the order of several Petaflops/s (a Petaflop denoting 10^{15} floating point operations), placing it in the league of the most powerful computers of the world. To enable the scientists using the new system to productively use a significant proportion of this enormous capacity, which will be aggregated by installing on the order of more than hundred thousand processor cores, both LRZ and the prospective vendors of such systems need to make big efforts to achieve good scalability and efficiency of operation with respect to nearly every system component.

LRZ's target for the architecture is a combination of a large number of moderately powerful compute nodes, with a peak performance of several hundred GFlop/s each, and a smaller number of fat compute nodes with a large shared memory. The network interconnect between the nodes must allow for perfectly linear scaling of parallel applications up to the level of 10,000 tasks. By stipulating suitable benchmarks, LRZ intends to assure that it is also possible to go significantly beyond this limit without large cutbacks in performance. Additionally, more than a one-order-of-magnitude increase of permanent storage capacity and bandwidth compared to the present system will be provided.

In order to accomplish the goals set above, LRZ will considerably increase its level of subject-specific project and user support. For very large projects, this will allow systematic improvements of the parallel algorithms deployed in scientific codebases.

"SuperMUC will strengthen the position of the Gauss Centre for Supercomputing in Europe by delivering outstanding compute power and integrating it into the European High-Performance Computing ecosystem," says Prof. Dr. A. Bode, Director of LRZ.

News

Extensions (marked in yellow) of the LRZ facilities: Doubling the Computer Cube

New institute building with Visualization Cave (marked in yellow)

Copyright of Figures: Herzog+Partner, Architekten

Project management and control: Staatl. Bauamt München 2

Autumn 2009 • Vol. 7 No. 2 • inSiDE 11

Identification of Anisotropic Elastic Material Properties from Micro-Fem Simulations for Natural Materials

The biomechanical simulation of bone-implant systems is subject to many studies presented in the open literature. Up to the present day most of these studies consider bone tissue as an isotropic material with inhomogeneous distribution throughout the model even though it is well proven that especially cancellous bone, which fills the joint region of the long bones has to be considered at least as orthotropic [1,2].

The negligence of the correct material properties can be assigned to their minor influence to global solutions in contrast to, for example, the muscle forces. Another problem is the non-existence of proper material models which provide a suitable mapping of orthotropic inhomogeneous material distributions from clinical imaging data to FE-models. The existence of correlations between bone density data which are provided for example by clinical computer tomography and orthotropic elastic constants for bone is shown in [3]. That these experimentally determined correlations can be utilized in an engineering manner to generate orthotropic inhomogeneous material distributions from clinical CT-data suitable for FE-simulations is shown by Schneider et al. [4]. A question which arises from this kind of material mappings and also from the general application of the FE-method with elastic continuum material formulations to micro structured natural materials is the one of resolution. With clinical computer tomography isotropic spatial resolutions up to 0.5 mm are possible. The element length in currently used FE-models for the simulation of bone implant systems, for example of the human femur, varies between 1 mm – 3 mm [2,5,6]. For the experimental determination of elastic constants specimens with an edge length of approximately 10 mm are used [7,8].

The current mapping method proposed by our group calculates the elastic constants by averaging all voxels of a CT-dataset which are enclosed in one element of the corresponding FE-mesh. This means the resulting Representative Volume Element (RVE) is equal to the element resolution of the FE-mesh but not even close to the size of the RVE used to determine the elastic constants from experiments. If one now wants to do non-linear simulations which include plasticity and damage models a refinement of the FE-meshes close to the implant will be necessary what of course will lead to a further decrease of the RVE size.

The aim of this study is to evaluate the influence of the RVE size reduction to the elastic material constants by means of the direct mechanics approach. Furthermore this study is a first step towards the clarification whether it is justified to apply an orthotropic continuum material model to bone tissue if the RVE size decreases to the resolution of clinical imaging data.

The Standard Mechanics Approach

With the standard mechanics approach a procedure is described which is often used as the starting point to the application of homogenization theory [9,10]. Aim of the approach is the determination of averaged elastic properties of micro structured materials on the continuum level.

One perquisite which is essential to this approach is that the material properties of the microstructures in the analyzed RVE are known. In this study the material of the bone microstructures is considered to be homogeneous with isotropic behavior.

To calculate the averaged macroscopic strain in an RVE from the local microscopic strain in the microstructures, the connection

$$\bar{\varepsilon}_{ij} = \frac{1}{|V_{RVE}|} \int_{V_{RVE}} \varepsilon_{ij}\, dV_{RVE} \qquad (1)$$

is given, with:
V_{RVE} the Volume of the RVE;
$\bar{\varepsilon}_{ij}$ the mean strain tensor of the RVE;
ε_{ij} the local strain tensor;

A similar connection is given for the calculation of the averaged macroscopic stress in the RVE form the local microscopic stresses.

$$\bar{\sigma}_{ij} = \frac{1}{|V_{RVE}|} \int_{V_{RVE}} \sigma_{ij}\, dV_{RVE} \qquad (2)$$

With:
$\bar{\sigma}_{ij}$ the mean stress tensor of the RVE;
σ_{ij} the local stress tensor;

To connect the averaged macroscopic strain applied to the RVE with the resulting averaged macroscopic stress one defines

$$\bar{\sigma}_{ij} = \bar{C}_{ijkl}\, \bar{\varepsilon}^{kl} \qquad (3)$$

With $\bar{C}_{ijkl}$ the so called effective stiffness of the RVE which is a tensor of rank four.

Furthermore, to connect the local microscopic strains with the averaged macroscopic strain the local structure tensor M_{ijkl}, also a tensor of rank four, is defined by

$$\varepsilon_{ij} = M_{ijkl}\, \bar{\varepsilon}^{kl} \qquad (4)$$

The development of the correlation between the elastic properties of the microscopic structures to the averaged elastic properties of the RVE is started from the generalized Hooke's law on the microscopic level

$$\sigma_{ij} = C_{ijkl}\, \varepsilon^{kl} \qquad (5)$$

Integrating both sides of equation (5) yields

$$\frac{1}{|V_{RVE}|} \int_{V_{RVE}} \sigma_{ij}\, dV_{RVE} = \frac{1}{|V_{RVE}|} \int_{V_{RVE}} C_{ijkl}\, \varepsilon^{kl}\, dV_{RVE} \qquad (6)$$

Substituting the left hand side with equation (2) and ε^{kl} with equation (4) we get

$$\bar{\sigma}_{ij} = \frac{1}{|V_{RVE}|} \int_{V_{RVE}} C_{ijmn}\, M^{mn}{}_{kl}\, dV_{RVE}\ \ \bar{\varepsilon}^{kl} \qquad (7)$$

Comparing equation (7) to equation (3) one recognizes the relation

$$\bar{C}_{ijkl} = \frac{1}{|V_{RVE}|} \int_{V_{RVE}} C_{ijmn}\, M^{mn}{}_{kl}\, dV_{RVE} \qquad (8)$$

from which the effective stiffness of the RVE can be calculated if the local elastic properties and the function of the local structure tensor for the micro structure of the RVE are known.

Applications

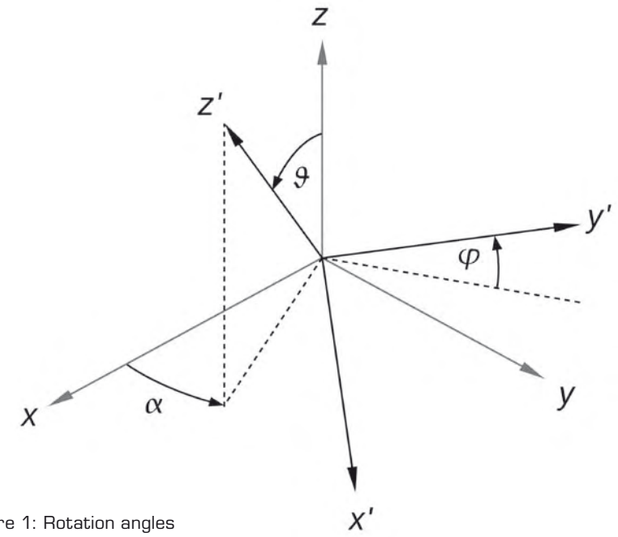

Figure 1: Rotation angles

Determination of Material Symmetries

Since this study deals with natural materials material symmetries in the analyzed specimens cannot be known a priori and so the calculated effective stiffness matrices will be dense in general. To determine whether there are distinguished directions in which material symmetries like orthotropy appear, the calculated effective stiffness tensors were simply transformed step by step with a rotating coordinate system. The coordinate system was rotated with its third direction over a complete hemisphere and according to each rotation step of the third direction the first and second direction was rotated over an angle of π. See Figure 1 for the definition of the rotation angles.

For each transformation of the effective stiffness tensor the sum of squares of the coupling terms being 0 for an ideal symmetric material was evaluated. The minimum of these sums was then assigned the highest achievable symmetry in the material properties of the analyzed RVE.

Bone Specimens

The subject matter to this study is a micro computer tomography data set of a cancellous bone specimen taken from the distal joint region of a cadaverous, primed human femur (see Fig. 2). The size of the micro-CT data set was 817 x 865 x 504 voxels with an isotropic spatial resolution of 0.014 mm. To avoid the disturbance from the cutting blade in the boundary regions of the specimen only 350 voxels in each direction where used from the original data set which results in a physical cube shaped model with an edge length of l = 4.9 mm.

Finite-Element Models

For the extraction of the bone geometry from the CT-data the iso-surface module from the HLRS visualization tool COVISE was used. The module triangulates iso-surfaces by means of the marching-cubes algorithm. The output is then written as an ABAQUS input deck which makes it possible to directly import the bone geometry on triangle basis into a FE pre-processor.

Figure 2: Position and representations of cancellous bone specimen

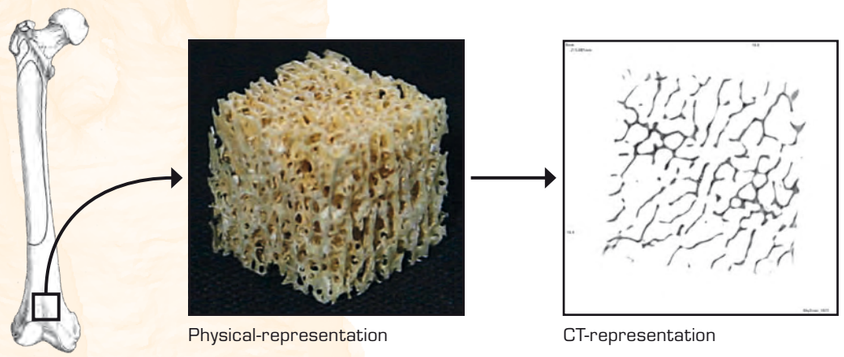

Physical-representation

CT-representation

Since the marching cubes algorithm produces triangles of bad quality not suitable for direct FE-meshing a mesh reconstruction was performed. Furthermore the element edge length was increased during the reconstruction from 0.014 mm to 0.03 mm to reduce the total number of triangles.

After mesh reconstruction and improvement the bone geometry was split in 8 sub-volumes each with an edge length of l =2.45 mm (see Fig. 4). The sub-volumes are specified by indices according to their position on the coordinate axes, e.g. the one next to the origin of the coordinate system in the lower left corner is called $SV_{-x,-y,-z}$.

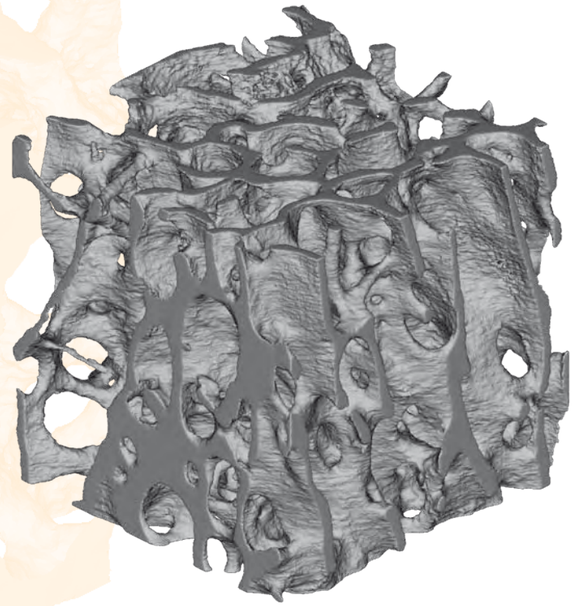

Figure 3: Geometry configuration of complete modelled volume

The boundary conditions were modelled with ABAQUS constraints as follows. One independent master node was placed in the middle of each face of the considered cube. Then the degree of freedom (DOF) of all nodes in one face acting in the normal direction of the face was coupled to the rigid body movements of the corresponding master node. This technique allows a free movement of all nodes in one face in tangential directions while ensuring that the face remains absolutely planar. Prescribed displacements can now simply be placed on the six master nodes, translational as well as rotational ones,

Figure 4: Geometry configuration of sub-volumes

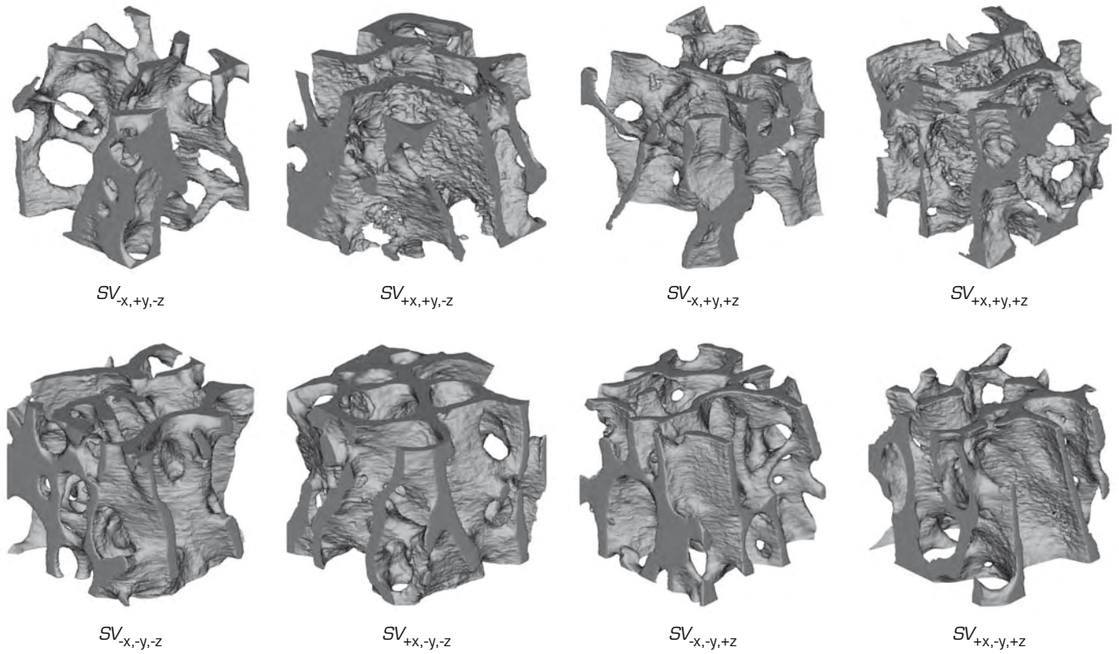

so reproducing the desired average strains needed for the evaluation of equ. (4). As already mentioned the FE-simulations done for this study where carried out with ABAQUS 6.8-4. The homogeneous isotropic material applied to the bone structures was modeled with the elastic constants E_{iso} = 5330 MPa and ν_{iso} = 0.3 according to [11].

The calculations of the effective stiffness matrices and the search for possible symmetries of the material constants where done with a post processing procedure implemented and developed in FORTRAN 2003 and Coarray FORTRAN at the HLRS.

Absolute Value of the Effective Stiffness Matrix Coefficients

As the first result of this study the absolute values of the components E_{ij} of the calculated effective stiffness matrices E are evaluated. Before the evaluation was done, all matrices were transformed so that E_{11} is the lowest and E_{33} is the highest of the three normal stiffness components. In Figure 5 the absolute value of E_{11} and E_{33}, the highest and lowest shear stiffness components, denoted with $E_{S,min}$ and $E_{S,max}$, and the highest coupling term, denoted with $E_{C,max}$, which should be zero in the case of a material with 3 or more major symmetry planes, are presented.

First thing to notice is that the values of the calculated normal and shear stiffness coefficients are comparable to the ones that can be found in the literature for cancellous bone [11]. Regarding the problem of resolution, from Figure 5 it can further be seen that the increase of resolution from 4.9 mm to 2.45 mm reveals a significant variance of the elastic coefficients in V.

Difference between Calculated and Expected Material Behavior

The major characteristic of orthotropic material behavior is the decoupling of normal and shear strains if the material coordinate system is oriented along the three orthogonal, distinguished material directions. This means, that only the non diagonal coefficients of the stiffness matrix, which couple normal to normal strains, are $\neq 0$.

To specify the difference between the calculated effective material behavior and the expected orthotropic material behavior Δo the sum of absolute values of the shear to normal strain coupling coefficients normalized with respect to the highest normal to normal strain coupling coefficient in the considered matrix is evaluated.

As an example for this approach the evaluation for E_V is shown.

$$E_V = \begin{bmatrix} \mathbf{192} & \mathbf{71} & \mathbf{104} & 17 & 4 & 1 \\ \mathbf{71} & \mathbf{245} & \mathbf{95} & 0 & 7 & 19 \\ \mathbf{104} & \mathbf{95} & \mathbf{689} & 2 & 0 & 0 \\ 17 & 0 & 2 & \mathbf{32} & 1 & 4 \\ 4 & 7 & 0 & 1 & \mathbf{56} & 8 \\ 1 & 19 & 0 & 4 & 8 & \mathbf{59} \end{bmatrix} \rightarrow \begin{bmatrix} 0{,}234 & 0{,}060 & 0{,}016 \\ 0{,}001 & 0{,}093 & 0{,}266 \\ 0{,}028 & 0{,}001 & 0{,}006 \\ & 0{,}018 & 0{,}055 \\ & & 0{,}112 \end{bmatrix}$$

$$\rightsquigarrow \Delta o_V = 0.890$$

What can be seen from the comparison of E_V and the sub-volumes in Figure 6 is the fact that not only the absolute values of the stiffness coefficients show variations when increasing the resolution but also the degree of orthotropy in 4 out of the eight sub-volumes is significantly different from the one of *V*. The authors are aware, that from the analysis of one specimen no general conclusions can be drawn. But looking at the total values of Δo_V and Δo_{SV} one can recognize the tendency that a material model made for a lower resolution scale is not necessarily valid for a higher one. Another tendency which can be seen from the analysis of the elastic constants is the increasing difference to the general assumption of bone material being fully orthotropic, that goes along with the level of resolution.

Directions of Major Material Symmetries

The evaluation of the coordinate systems which were found with the search for the minimum of the shear to normal strain coupling coefficients showed that the orientation of the highest normal stiffness coefficient is not significantly different in *V* and six of its sub-volumes. Here the 33-direction is oriented up to the difference of some

Figure 5: Absolute values of selected stiffness components

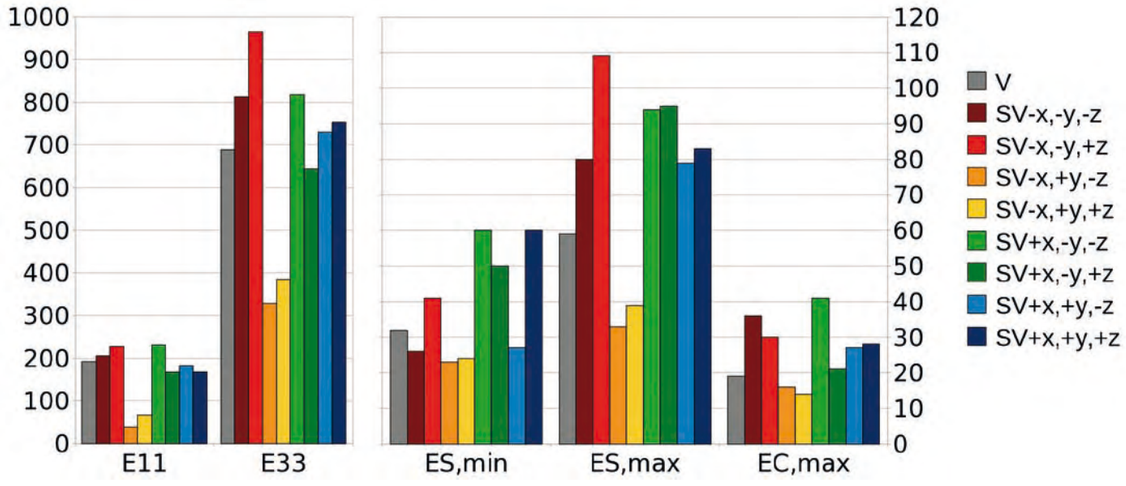

Figure 6: Absolute values of selected stiffness components

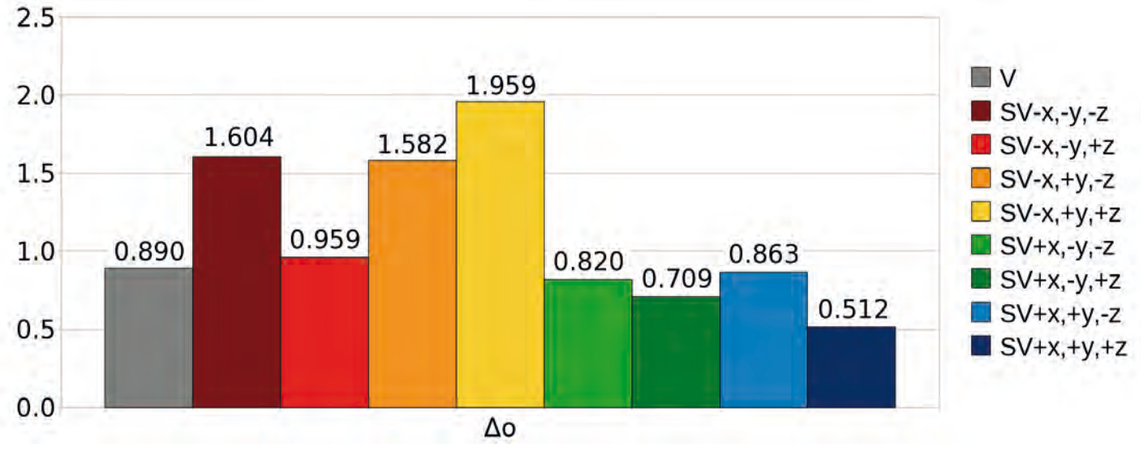

degrees along the global zz-direction. This orientation is the one which was expected at least for *V* since the specimen was indeed cut from an arbitrary position in the knee joint region of the bone but the axes of the cube were oriented along the main load trajectories.

These load trajectories can be found in most joint regions of long bones. They are oriented during the remodeling process of bone tissue towards the directions in which the main loads to the bone are acting. In Figure 7 the orientations of *V* and all sub-volumes are shown with their correct relative position to each other.

What was not expected and kind of surprising in this evaluation is the fact, that in the two sub-volumes $SV_{+x,-y,-z}$ and $SV_{+x,+y,-z}$ the general orientation of the major symmetry directions is still similar to the ones of V and the other sub-volumes but the 11- and 33-direction have switched places. Looking at the orientations of the 11- and 22-direction it can be seen, that only the ones in $SV_{-x,-y,-z}$ and $SV_{-x,-y,+z}$ are oriented similar as the one in V.

This result suggests that not only the variance of the elastic coefficients with the level of resolution but also the variance of the orientation of the major symmetry directions has to be analyzed on the correct resolution level.

Figure 7: Orientations of the major symmetry directions

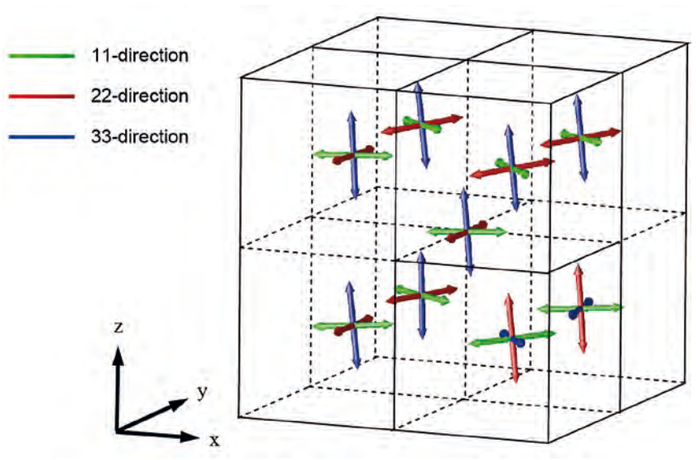

This result is essential to the application of material mapping methods like the one proposed in [4] because there the calculation of the material constants from the scalar field of the CT-density data is independent from the algorithm which determines the orientation of the material symmetry directions. This means only if this algorithm considers the CT-data with the same resolution as the calculation of the material constants does, the mapping method will produce consistent material data.

Conclusions

In this study the standard mechanics approach for the calculation of the effective stiffness matrices of micro structured materials was applied to a cancellous bone specimen at two different RVE resolutions. The resolutions are both significantly smaller than the ones which are used in the literature for the experimental determination of elastic material constants whereas the resolution of the smaller RVE size was comparable to the one which is used as the mesh resolution in current FE-simulations of bone-implant-systems.

It was shown, that on this RVE scale the general assumption of bone material being orthotropic has to be handled with care, since the normal strain to shear strain coupling coefficients are already of significant size compared to the normal to normal strain coupling components.

Furthermore the evaluation of the orientation of the major symmetry planes in the analyzed RVEs showed that there are also variances with increasing resolution. The effect was more observable between the local 11- and 22-directions than between the local 33-directions but nonetheless two out of eight sub-volumes showed a significantly different orientation than the others.

Although it is clear that from the analysis of one bone specimen one cannot draw any general conclusions, it can be said that the results of this study reveal a dependency of the local material constants upon the used RVE size. Furthermore the authors think that the resolution currently used for the FE-simulation of bone implant systems lies in a not very lucky range where the statement can be made:

"An increase of resolution goes along with a decrease of material symmetry."

References

[1] **Trabelsi, N., Yosibash, Z., Milgromb, C.**
"Validation of Subject-Specific Automated p-FE Analysis of the Proximal Femur", Journal of Biomechanics, 42, pp. 234-241, 2009

[2] **Helgason, B., Taddei, F., et al.**
"A Modified Method for Assigning Material Properties to FE-Models of Bones", Medical Engineering & Physics 30, pp. 444-453, 2008

[3] **Rho, J. Y., Hobatho, M. C., Ashman, R. B.**
"Relations of Mechanical Properties to Density and CT-Numbers in Human Bone", Medical Engineering & Physics, 17, pp. 347-355, 1995

[4] **Schneider, R., Faust, G., Hindenlang, U., Helwig, P.**
"Inhomogeneous, Orthotropic Material Model for the Cortical Structure of Long Bones Modelled on the Basis of Clinical CT or Density Data", Computer Methods in Applied Mechanics and Engineering, 198, pp. 2167-2174, 2009

[5] **Helwig, P., Faust, G., Hindenlang, U., Kröplin, B., Eingartner, C.**
"Finite Element Analysis of a Bone-Implant System with the Proximal Femur Nail", Technology and Health Care, 14, pp. 411-419, 2006

[6] **Eberle, S., Augat, P.**
"A Standardized Finite Elemente Model of the Human Femur for Biomechanical Simulations of Proximal Fractures", in: 14th Finite Element Workshop, University of Ulm, ISBN:978-3-9806183-9-7, pp. 98-102, 2007

[7] **Snyder, S., Schneider, E.**
"Estimation of Mechanical Properties of Cortical Bone by Computed Tomography", Journal of Orthopaedic Research, 9, pp. 422-431, 1990

[8] **Ashman, R. B., Rho, J. Y.**
"Elastic Modulus of Trabecular Bone Material", Journal of Biomechanics, 21 (3), pp. 177-181, 1988

[9] **Hill, R.**
"Elastic Properties of Reinforced Solids: Some Theoretical Principles", Journal of the Mechanics and Physics of Solids, 11, pp. 357-372, 1963

[10] **Hollister, S. J., Kikuchi, N.**
"A Comparison of Homogenization and Standard Mechanics Analysis for Periodic Porous Composites", Computational Mechanics, 10, pp. 73-95, 1992

[11] **van Rietbergen, B., et al.**
"Direct Mechanics Assessment of Elastic Symmetries and Properties of Trabecular Bone Architecture", Journal of Biomechanics, 29 (12), pp. 1653-1657, 1996

- Ralf Schneider[1]
- Ulrich Hindenlang[2]
- Michael Resch[1]

[1] University of Stuttgart, HLRS

[2] LASSO Ingenieurs-gesellschaft mbH, Germany

Large-Eddy Simulations for Complex Flow Geometries

The investigation of flows in technical applications (automobile, CPU cooling, ...) becomes increasingly affordable through numerical simulation and larger computers, but the numerical simulation of turbulent flows with methods describing the full features of the flow are even way beyond today's supercomputers as theoretically, statistical disturbances would have to be treated down to a molecular-sized level. As turbulence is critical for technical applications (i.e. mixing or heat transfer), the investigation of these highly statistical flows is very important in today's development of increasingly complex technical applications.

For the numerical calculation of turbulent flows, Large-Eddy Simulations (LES) are entering the stage of industrial usability. The principle of the numerical scheme is the direct simulation of large eddies, which are the energy-containing structures. These resolved eddies are generating smaller and smaller structures as the flow dissipated energy from large to small scales which are below the size of the numerical resolution. The energy content of these small scales has to be modeled in the equations, as they cannot be resolved. This is done with a so-called turbulence model, which takes into account the dissipation rate of isotropic turbulence but neglects the structure of the eddies as they are believed to be generic in nature. In our project at the HLRB at the LRZ München, we try to model the dissipation of the small eddies through tailored dissipation of the numerical model which then behaves as a turbulence model. The technical application of LES usually comes with complex geometries (for example a motor-cycle driver on his/her bike), where the generation of a suitable body-aligned mesh takes up considerable time and the validity of the result is dependent on the quality of the mesh. So we are working on methods where complex obstacles are immersed in very simple meshes, that is, the surfaces defining the geometry cut through the mesh. In following such an approach, meshing is fast, can be fully automated, and is much less a source of uncertainty towards the quality of the result.

Figure 1: Speedup – reference value with 32 cores

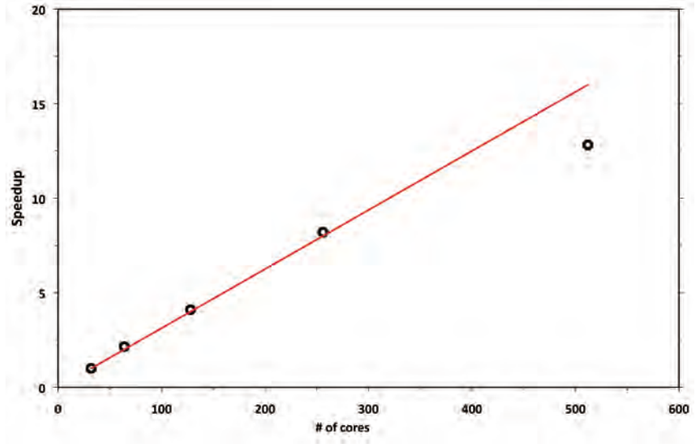

As an example, we have chosen a simplified geometry that exhibit the same complex flow features as the problem cited above: A three-dimensional circular cylinder subjected to a uniform flow at a velocity perpendicular to the cylinder axis. A parameter, which describes the character of the flow regime, is the Reynolds number. The Reynolds number describes the relation of inertial forces to viscous forces. At the considered Reynolds number of Re = 3,900, the flow separates from the cylinder at an angle of about 88° taken from the line of symmetry of the geometry. The large vortices that detach from the cylinder periodically alternating from the top and bottom side of the cylinder very rapidly break down in smaller scales and turbulence develops. The calculations were done on an adaptive, locally refined grid with a total of 7 million cells. Qualitatively good agreement is reached with the results of other calculations: the typical pressure tubes for the p-isosurfaces corresponds with the results reported by Fröhlich [1]. Instabilities in spanwise direction give rise to the elongated structures for the x-vorticity reported by Kravchenko [2] (Fig. 2). The unsteady computation delivers a statistical development of the turbulent wake of the cylinder and was carried out on 512 processors in about 300,000 CPUh with our parallelized code INCA (Solver for the (In)compressible Navier-Stokes equations on Cartesian Adaptive grids [3-6]). INCA was also tested for its parallel performance (Fig. 1). The speedup (reference case: 32 cores) scales almost linearly with the number of processors.

With the implicit LES code INCA, almost arbitrary shapes of boundaries are realizable. In case the time evolution of a complex flow is of interest (e.g. vortex shedding from the back of a vehicle and its noise generation) or the statistical turbulent data are needed for more accurate design purposes, the presented method can deliver dependable results. With the increase in computing power, these type of computations come into reach for the design of high technology products.

Collaborators

Michael Meyer, Stefan Hickel, Antoine Devesa, Christian Stemmer, Nikolaus A. Adams

References

[1] **Fröhlich, J., Rodi, W., Kessler, Ph., Parpais, S., Bertoglio, J. P., Laurence, D.** "Large Eddy Simulation of Flow around Circular Cylinders on Structured and Unstructured Grids", CNRS DFG Collaborative Research Programme, Vieweg, Braunschweig 66, pp. 319-338, 1998

[2] **Kravchenko, A., Moin, P.** "Numerical Studies of Flow over a Circular Cylinder at ReD = 3900", Physics of Fluids 12, p. 403, 2000

[3] **Hickel, S., Adams, N. A.** "Implicit LES Applied to Zero-Pressure-Gradient and Adverse-Pressure-Gradient Boundary-Layer Turbulence", International Journal of Heat and Fluid Flow 29, pp. 626-63, 2008

[4] **Hickel, S., Adams, N. A., Domaradzki, J. A.** "An Adaptive Local Deconvolution Method for Implicit LES", Journal of Computational Physics 213, pp. 413-436, 2006

[5] **Hickel, S., Adams, N. A.** "On Implicit Subgrid-Scale Modeling in Wall-Bounded Flows", Physics of Fluids 19, pp. 105-106, 2007

[6] **Hu, X. Y, Khoom B. C., Adams, N. A. Huang, F. L.** "A Conservative Interface Method for Compressible Flows", Journal of Computational Physics 219, pp. 553-578, 2006

• Christian Stemmer

Department for Aerodynamics, TU München

Figure 2: Instantaneous pressure iso-contours in the turbulent cylinder wake

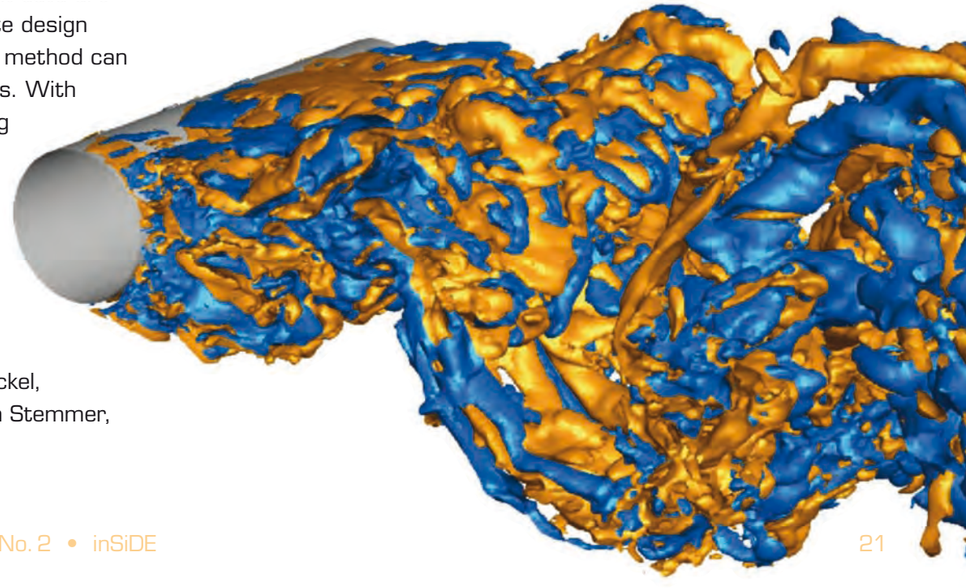

Direct Simulation and Modeling of Micro-Mixing in High Schmidt Number Flows

The prediction of turbulent mixing and reaction processes in fluids is a challenging and important task in engineering and environmental science as it is encountered in many different areas like combustion, chemical engineering or biology. The difficulty lies in the fact that a large number of scales are involved and reactions rates are non linear. Often the reactants are initialized completely segregated, undergo turbulent mixing where the segregation length scale is reduced by a mechanical process and finally molecular diffusion brings the reactants together on the molecular scale at which reaction takes place. As reaction rates are often highly non linear small errors in the prediction of the distribution of the reactants on the molecular scale can lead to the prediction of very inaccurate reaction rates. The smallest length scale at which the reactants are mixed on a molecular level is determined by the molecular diffusion.

When the molecular diffusion is related to the kinematic viscosity of the fluid the dimensionless Schmidt (Sc) number is formed. The relation of the smallest scales of the reactants to the smallest scales in the flow field can be deduced by the theories of Kolmogorov and Batchelor [1] to be $1/\sqrt{Sc}$. In liquids the Schmidt number can easily reach values around Sc = 1,000 which implies that the smallest relevant scales in the scalar field (reactants) is approx. 32 times smaller than the smallest relevant scales in the flow field (Kolmogorov scale). Figure 1 shows a snapshot from the time evolution of the mixing of two species for two different Schmidt numbers, Sc = 1 (left) and Sc = 100 (right), in a 2D temporal mixing layer. In both simulations the flow field is completely identical, only the concentration distribution develops much smaller structures at high Sc.

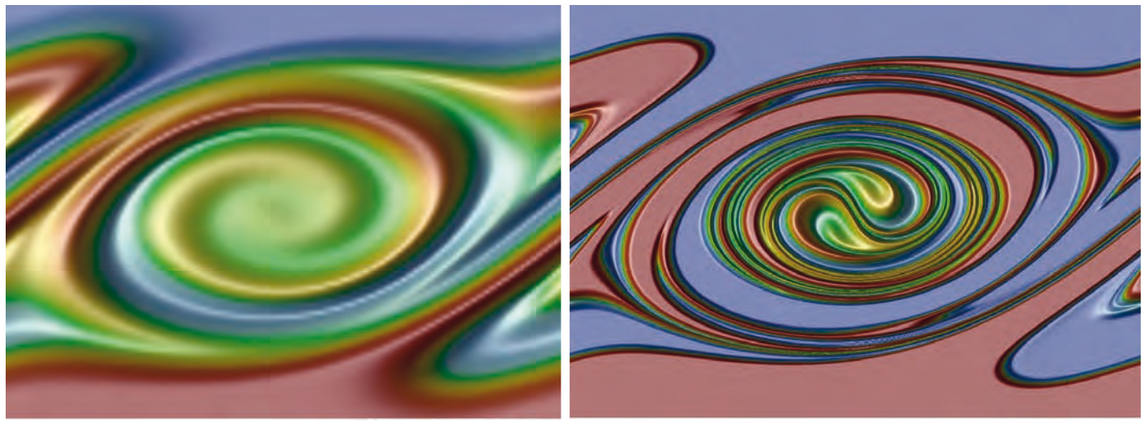

Figure 1: Concentration distribution at different Schmidt numbers, Sc = 1 (left) and Sc = 100 (right) in a 2D temporal mixing layer with same initial conditions

Several strategies are available for the description of mixing in turbulent flows of which Direct Numerical Simulation (DNS) is the most reliable one as all turbulent length and time scales are resolved by the numerical method. However, as the smallest length scales in the scalar concentration fields in liquids are much smaller than the Kolmogorov scale the direct computation of turbulent scalar fields in a liquid is unfeasible. If only scales down to the Kolmogorov scale are resolved by the computational grid, no further turbulence model is required for the flow field, but modeling of micro-mixing, i.e. mixing on scales smaller than the Kolmogorov scale is therefore necessary in such computations. This is illustrated with Figure 1. If one considers the flow field of the 2D mixing layer as a prototypical eddy at the Kolmogorov scale, the reaction rate would be strongly Sc dependent and its prediction would be prone to large errors if the filtered concentration value at the scale of the eddy would be used. Therefor reaction rates can not be described by quantities filtered at the Kolmogorov scale as the sub-Kolmogorov scales dominate the concentration fluctuations of the species involved. Predicting reaction rates in liquid systems requires either modeling of the reaction rate itself or a correct representation of the sub-Kolmogorov fluctuations which can be done in a method based on probability density function in which the reaction term in the convection-diffusion equation can be obtained in closed form. When combining a DNS of the flow field with a stochastic representation of the scalar field, which is termed DNS-FDF method, the problem of modeling the reaction rate reduces to the problem of finding the correct mixing model with its corresponding time scale.

Direct Numerical Simulation of Turbulent Mixing at High Schmidt Number

Analysis of the mixing process in turbulent flows at high Schmidt numbers and the development and validation of a mixing model for the sub-Kolmogorov fluctuations requires detailed data on the dynamics and spectral distribution of production and dissipation of scalar variance for all length scales. Information of this kind can only be supplied in detail by resolving all relevant length and time scales in the flow and scalar field by the numerical grid (full DNS).

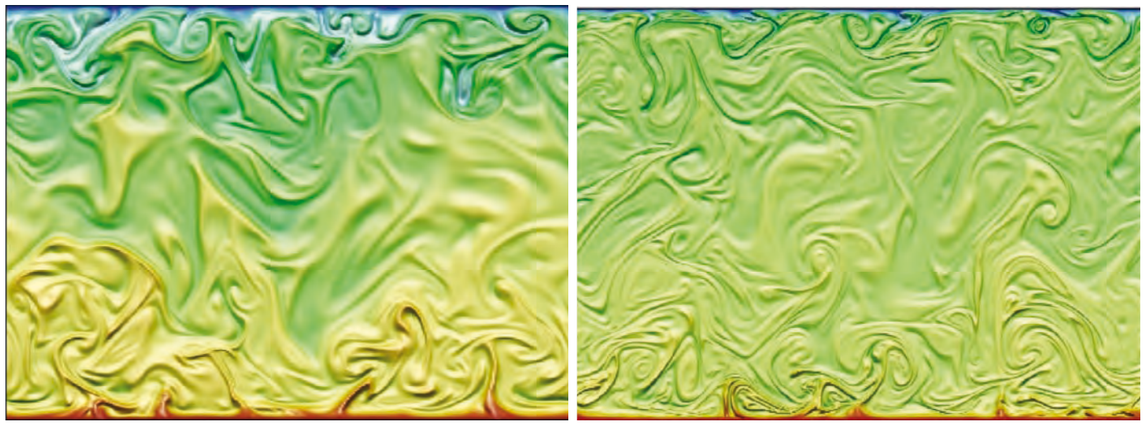

Figure 2: Passive scalar field in turbulent channel flow (Re_tau = 180) in a transverse plane. Sc = 10 (left) and Sc = 49 (right).

As the required resolution increases with Reynolds (Re) and Schmidt number a trade off between available computational resources and simulation parameters was chosen. Simulation of the mixing of a passive scalar up to Sc = 49 in a turbulent channel flow at Re_tau = 180 were performed. They have been obtained by massive computations on the SGI Altix 4700 at the Leibniz Computing Center (LRZ) in Munich. These were run on up to 10^9 grid cells and on up to 210 cores. Figure 2 shows snapshots of the scalar distribution at two different Sc in a plane normal to the main flow direction. It demonstrates the structural change of the scalar field with increasing Sc. A careful grid resolution study renders the results as reliable so that they can be used as reference data. These data gave a number of new insights into and results for passive scalar transport in turbulent channel flow of which a new formulation for the scaling of the heat transfer coefficient with Schmidt number can be considered as the most important [2]. These computations establish results for the highest Schmidt number so far in turbulent channel flow.

Micro-Mixing Model for Filtered Density Function Simulations

In our DNS-FDF method the distribution of the scalar field is described in a stochastic way: what is the probability to encounter a specific concentration in a specific area at a certain time instant? When the specific area is defined by a certain filter width, the resulting probability distribution is termed Filtered Density Function (FDF). The FDF gives information on the distribution of the scalar field within a small volume of the flow field. In the DNS-FDF method this small volume is the grid cell which dimensions are in the order of the Kolmogorov length scale. With this information the non linear reaction rates can be computed in a closed form and no modeling is needed at this point, which is the reason for the attractiveness of the FDF methods for reacting flows. However, the time evolution of the FDF is governed by the turbulent mixing processes. Scalar variance is produced and destroyed, both processes change the shape of the FDF. While the production of scalar variance is also known within the framework of DNS-FDF simulation the dissipation of scalar variance is not and must be modeled. With the detailed DNS data at hand, the processes of dissipation of scalar sub-Kolmogorov variance was studied and a model for the prediction of scalar dissipation rate was developed [3,4]. Figure 3 shows the time evolution of the production (G+V) and destruction (D) of the scalar dissipation rate along a Lagrangian trajectory as predicted by the model compared to the DNS data. The combined effects of these lead to the time evolution of the scalar dissipation rate. This time evolution is shown in Figure 4 as predicted by the model compared to the DNS data.

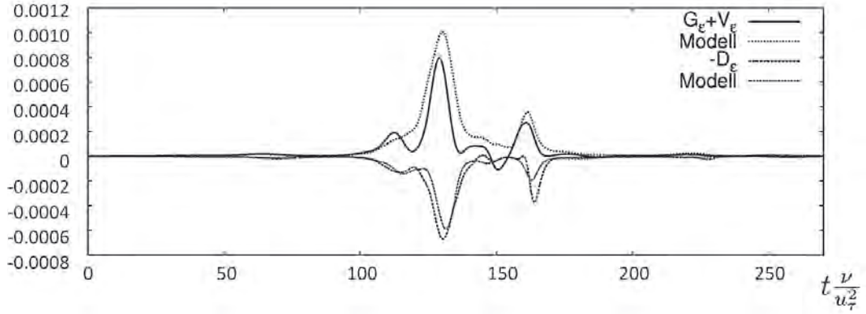

Figure 3: Sub-filter production (G+V) and destruction (D) of scalar dissipation rate along a Lagrangian trajectory in turbulent channel flow. Comparison of DNS data with results from model at Sc = 49

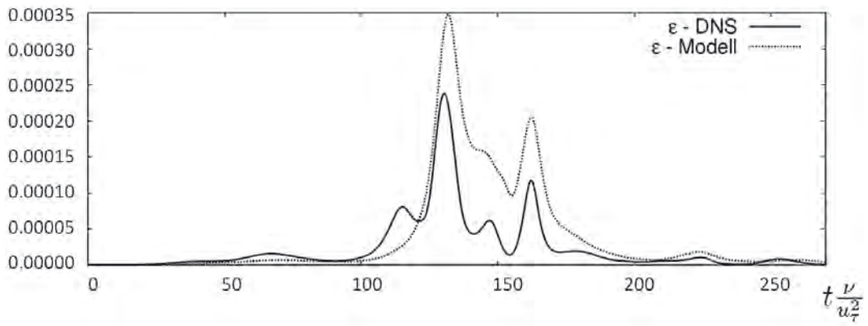

Figure 4: Sub-filter scalar dissipation rate along a Lagrangian trajectory in turbulent channel flow. Comparison of DNS data with model at Sc = 49

Conclusion and Outlook

The detailed comparison of the scalar dissipation rate as predicted by the micro-mixing model with the DNS data shows an excellent model performance [3,4]. Thus the DNS-FDF method provides capability for mixing and reaction processes in turbulent high Schmidt number flows. With the availability of more computational resources additional validation at Sc > 49 would further corroborate the good model performance. As the method is currently limited to turbulent flows at low Reynolds number due to the resolution requirement of the DNS of the flow field, an extension of the micro-mixing model to a LES-FDF approach is the next step.

References

[1] **Batchelor, G. K.**
"Small Scale Variation of Convected Quantities like Temperature in Turbulent Fluid." Part 1: General discussion and the case of small conductivity. Fluid Mechanics 5, pp. 113-133, 1959

[2] **Schwertfirm, F., Manhart, M.**
"DNS of Passive Scalar Transport in Turbulent Channel Flow at High Schmidt Numbers." Journal of Heat and Fluid Flow, 28, pp. 1204-1214, 2007

[3] **Schwertfirm, F.**
"Direkte Simulation und Modellierung des Mikromischens bei hohen Schmidt Zahlen.", München, Dissertation, 2008

[4] **Schwertfirm, F., Manhart, M.**
"Development of a DNS-FDF Approach to Non-homogeneous Non-equilibrium Mixing for High Schmidt Number Flows and Large Eddy Simulation 7", Trieste, Italy, September 2008

• Florian Schwertfirm
• Michael Manhart

Department for Hydromechanics, TU München

Numerical Investigation of a Highly Loaded Axial Compressor Stage with Inlet Distortions

In order to decrease costs during the design and development process of modern aircrafts, high efficient, accurate, and reliable design tools are necessary. Nowadays, powerful tools have got available which are able to predict the main and most relevant flow features around aircraft configurations and inside jet engines at normal flight conditions. However, at flight conditions at high angle of attack and at the border of the flight boundaries the flow is getting very complicated and is prone to massive flow separation. This has a strong impact on the safe operation of the aircraft and might cause critical situations if this is not accounted for during the design process. Additionally, in these situations flow distortions are likely to occur at the inlet of the jet engines which may result in unstable performance behavior of the compressor components inside the jet engine. In the worst case the engines cannot provide enough thrust anymore to ensure a safe operation of the aircraft. A reliable prediction of all relevant flow phenomena gets very challenging for these operating conditions. For this reason it is in the focus of many research activities worldwide.

For simulating the flow physics correctly, the strong and unsteady interaction of the inner jet engine flow and the outer flow around the aircraft becomes very important and needs to be taken into account. Accurate predictions are especially demanded during conditions where high aerodynamic loads are present.

One of these critical phases is the take-off procedure, where highly turbulent air with strongly varying total pressure may enter the jet engine. These inlet distortions increase the risk of compression system instabilities. Inlet distortions may be composed of total pressure, flow angle, and total temperature distortions, depending on the particular configuration. All three disturbances have a reduction of the safe operating range in common. The particular challenge is to predict accurately the creation and the migration of the inlet disturbances and its impact on the compressor and vice versa, since the outer aerodynamics in front of the engine and the flow into the jet engine has to be simulated simultaneously. At present, most of the numerical tools (flow solvers) are specialized either to inner or to outer aerodynamics. Thus one possibility to solve this problem is to couple numerically two flow solvers – one optimized code for the inner, the other one for the outer flow dynamics. This strategy is pursued by the members of a DFG (German Research Foundation) project, where a coupling of two powerful DLR codes is pursued.

The Institute of Jet-Propulsion at the University of the Federal Armed Forces in Munich is one member of the research consortium and concentrates it's research activities on the development of a simulation environment for the prediction of the flow phenomena in the inlet of jet engines with inlet distortions present. One of the first steps is the validation of existing flow solvers for this kind of situations. For this purpose a single axial compressor stage with a steady inlet total pressure distortion has been chosen. It is a transonic compressor stage, which was experimentally investigated for homogenous and disturbed inflow conditions at the Institute of Propulsion Technology of DLR in Cologne.

The CFD (Computational Fluid Dynamics) simulations are performed with the U-RANS (Unsteady Reynolds-Averaged-Navier-Stokes-Equations) code named TRACE developed by the Institute of Propulsion Technology of DLR in Cologne in collaboration with MTU Aero Engines. The flow solver allows the simulation of a multistage three-

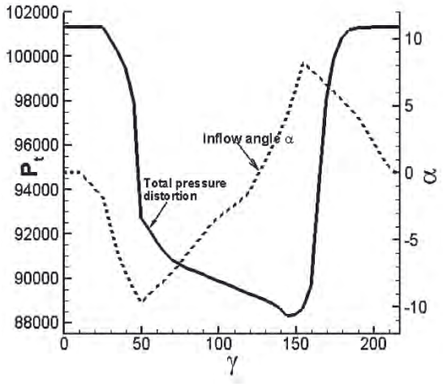

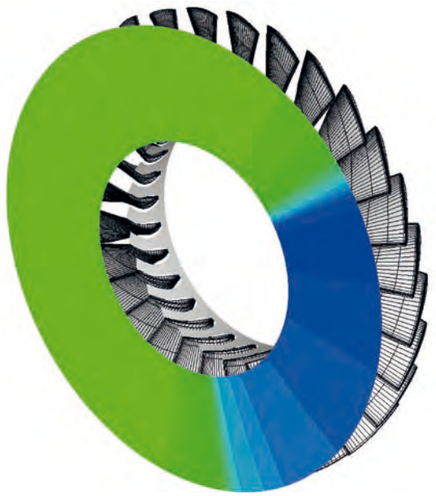

1: Measured and
ed inlet distortion
compressor inlet

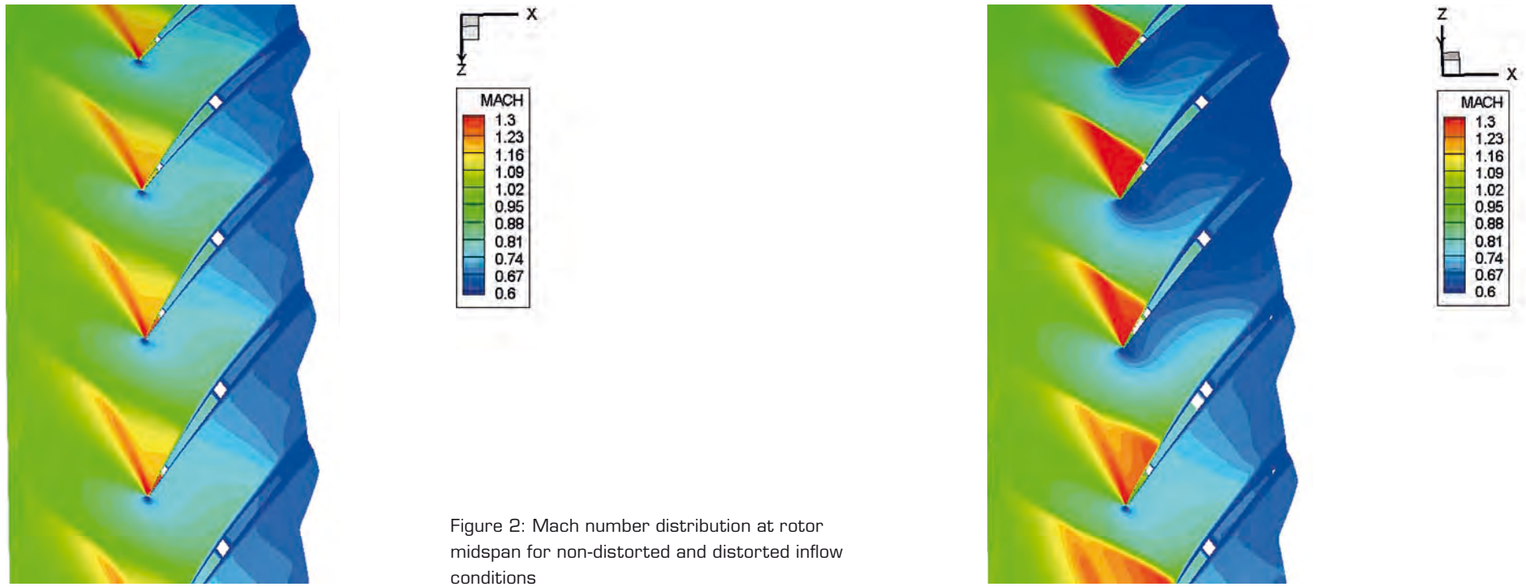

Figure 2: Mach number distribution at rotor midspan for non-distorted and distorted inflow conditions

dimensional, steady, unsteady and transitional turbomachinery flow on unstructured and structured grids.

In Figure 1 one of the inlet pressure distributions under consideration is shown. On the left hand side the experimental measured radial total pressure distribution is shown for a sector from 0° to 220°. The remaining sector from 220° to 360° is not included since the total pressure is constant in this part. Between 50° and 150° a massive drop of pressure level can be observed. On the right hand side the realization of the total pressure distortion in the inlet plane of the compressor rotor is presented as used for the flow simulations.

For flow simulations with distorted inflow conditions, a high grid resolution is requested. For the entire compressor stage a computational grid has been generated consisting of 766 structured blocks and more than 50 million grid points in total. Steady and unsteady simulations are performed with 130 CPUs. The requested memory is 495 GB. With this number of CPUs, a performance of 3,000 iterations per day is achieved. This leads to a runtime of approximately six months for an unsteady run consisting of typically 128 physical time steps on the supercomputer of Leibniz Computer Center in Garching.

An exemplary result of the steady flow simulations is depicted in Figure 2. Here the Mach number distribution at midspan of the blade is shown for non-distorted and distorted inflow conditions. On the left hand side, outside of the distortion sector, the results are very similar to those obtained for homogeneous inflow. In the picture on the right hand side the influence of the distortion becomes apparent: The smaller axial flow velocities due to the smaller total pressure result in higher incidence angles at the rotor leading edge. Hence, an increase of the aerodynamic load combined with stronger normal shock waves occurs. Unsteady simulations for disturbed inflow conditions are currently in progress.

After the simulation of steady total pressure distortions, in a further step unsteady distortions phenomena will be investigated as part of the DFG project indicated above. The Institute of Jet Propulsion is collaborating with Rolls Royce Germany and the University of Darmstadt to establish a new test case with a distortion generator of unsteady behavior. In order resolve the unsteadiness accurately the numerical grid for these kinds of inlet distortions has to be even finer and the time step size has to be properly adjusted. All this means that this test case will require even more computational resources, justified however by the gain of knowledge and better flow prediction capabilities.

After a successful completion of these research activities simulations tools will be available for design purposes and may therefore contribute to higher safety in general aviation.

• Reinhard Niehuis
• Jens Iseler
• Andreas Lesser

Institute of
Jet Propulsion,
University of the
Federal Armed
Forces Munich

Applications

The Molecular Bottle-Brush Controversy

Bottle-brushes are familiar to everybody as kitchen devices for the cleaning of glassware, test tubes in the laboratory, etc. Molecular bottle-brushes are cylindrical macromolecules with a "chemical architecture" that is a nanoscale analogue of macroscopic bottle-brushes: i.e., their diameter is a few nm (nanometer) rather than a few cm, but they also consist of a (more or less stiff) backbone, on which linear macromolecules are densely grafted, so that these side-chain polymers stretch out, perpendicularly oriented away from the backbone.

Such macromolecules can be synthesized, thanks to the "art" of organic chemists, since about 20 years, and find recently a lot of attention due to various potential applications [1]. For instance, if these cylindrical molecules are sufficiently stiff, they can align parallel to each other and behave as liquid crystals (LC), similar as are used in all kind of LC displays. Or one can use them as nanosize "building blocks" in various supramolecular assemblies, functioning as nanotechnological devices. Other applications rest on the property that these molecular bottle-brushes are stimuli-responsive polymers: changing the temperature, or the PH-value of a solution, etc., these macromolecules may exhibit drastic shape changes (e.g., flexible cylinders may collapse to spherical clusters or to pearl-necklace chains), and it has been suggested that such molecular brushes can be useful as sensors at the nanoscale, or as actuators causing a suitable response of a nanodevice [1].

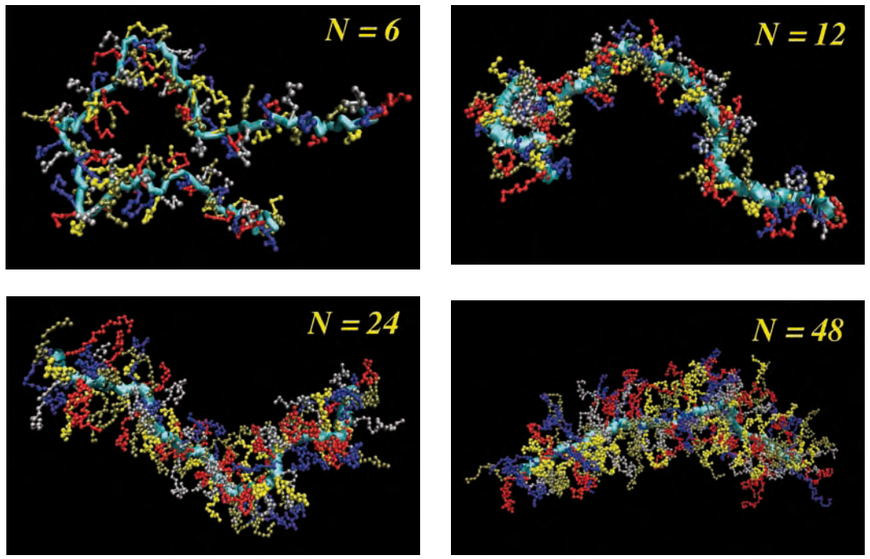

Figure 1: Snapshots of simulated bottle-brush polymers with a number N_b = 131 of effective "monomers" along the backbone. Each backbone monomer has a flexible side chain of N bonds connecting the N monomers grafted to it. Four choices of N are shown: N = 6 (a), N = 12 (b), N = 24 (c) and N = 48 (d). Note that the whole molecule becomes considerably stiffer when the length of the side chains increases, due to the "excluded volume interaction" among the monomers (no two monomers may sit "on top" of each other).

Actually, when such versatile macromolecules are discovered by the scrutinous work of synthetic chemists, one may suspect that nature has made use of such molecular bottle-brushes as well, and indeed useful biomolecular bottle-brush molecules were discovered. In the cartilage joints the aggrecan molecule has bottle brush architecture, and is assumed to be responsible for the excellent lubrication properties of all mammalian joints; friction forces between our bones in these joints are much less than what one observes in artificial devices [2].

Of course, for understanding the function of such macromolecules, good knowledge of their structure is an indispensable prerequisite. Unfortunately, as one can imagine from the snapshot picture of computer simulations of simple models, the structure can be pretty complicated (Fig. 1), and to quantitatively characterize it a multitude of length-scales has to be understood (Fig. 2). Extracting these multiple length-scales from experiment is very difficult, and hence a controversy in the experimental and theoretical literature has arisen (see e.g. [3,4,5,6]).

Controversies in polymer science about the structure of macromolecules are not uncommon. When Staudinger about 100 years ago introduced the very concept of a macromolecule consisting of 1,000 to 100,000 constituents, he was declared almost crazy by most of his contemporaries, who claimed that such large linear molecules can never exist and must break into small pieces. However, now we know that macromolecules exist as stable coils in solution and in dense melts (and Staudinger's insight earned him the 1953 Chemistry Nobel Prize). Another important step is due to Flory, who suggested that these macromolecules do not occur as stretched out linear straight-line like objects, but due to their flexibility rather form random coils, i.e. have a structure of a "fractal".

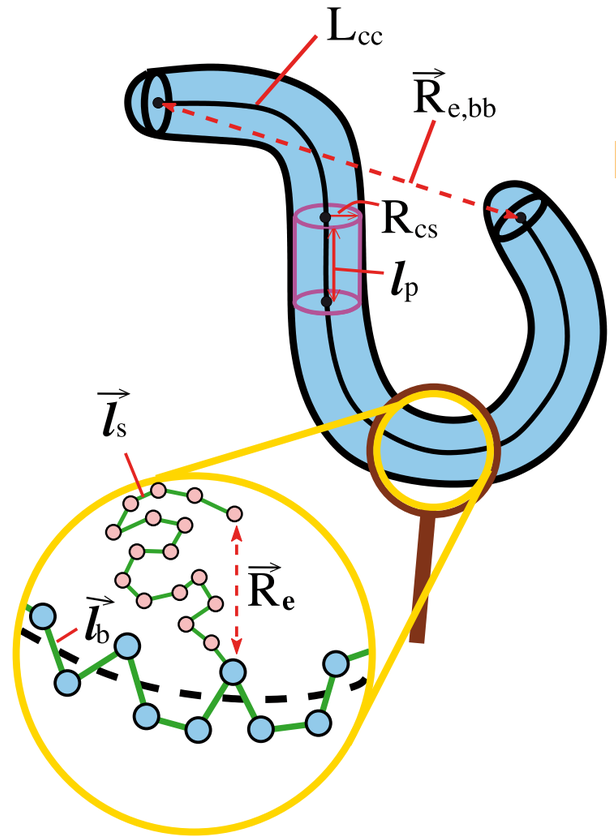

Figure 2: Multiple length scales are needed for the structure characterization of a molecular bottle-brush: a coarse-grained view depicts it as a flexible spherocylinder, upper part, while a "microscope" with better resolution would reveal the atomic description of monomers connected by bond vectors $\vec{l}_b$ along the backbone chain and by bond vectors $\vec{l}_s$ in the side chains (lower part). One needs to know: end-to-end distances of the backbone ($\vec{R}_{e,bb}$) and of the side chains ($\vec{R}_e$), the length L_{cc} of the (coarse-grained) cylinder axis, the persistence length l_p measuring the distance over which this axis basically is a straight line, and the cross sectional radius R_{cs} of the cylinder.

In the melt, the relation between the mass M of a macromolecule and the radius R of the (spherical) volume that it takes is not that of standard "massive" matter in three-dimensional (d = 3) space, $M \propto R^d$. Rather one has $M \propto R^{d_f}$, with a "fractal dimension" $d_f = 2$ (typical for random walk-like structures, RW), while in solution (where excluded volume interactions lead to a "swelling" of the random coil) one has $d_f = 5/3$ (the chain forms then a "self-avoiding random walk", SAW).

Also these ideas were challenged, meander-like configurations of macromolecules etc. were suggested, until finally careful neutron scattering studies of deuterated macromolecules (in melts of protonated chains, or in dilute solutions) settled this controversy. Flory earned a Nobel Prize for Chemistry for his work in 1974, but for a long time the explanation how one could understand the fractal dimensionality d_f for the SAW was controversial. In 1972, this controversy was then settled by deGennes (Physics Nobel Prize 1991), who mapped the problem of SAW exponents to critical exponents in the theory of phase transitions, and exploited mathematical methods such as renormalization group methods and scaling concepts based on the "selfsimilarity" of fractal structures. A small part of a SAW, say a mass $M' = \lambda M$ with $\lambda << 1$, scales also as $M' \propto R'^{d_f}$, so $R' = \lambda^{d_f} R$ – a change of scale for M by λ can be compensated by a change of scale by λ^{d_f} for R, but the character of the structure remains the same.

When one tries to extend such scaling ideas to bottle-brush polymers, then one ends up with predictions such as $R_{cs} \propto N^{3/4}$ (for the cross-sectional radius of the brush) and $l_p \propto N^{15/8}$ (for the persistence length of the brush) [5] but these power laws are neither useful to understand simulations [6] nor experiment: the problem of bottle-brush polymers is a problem of multiple length scales, which all influence each other. Computer simulation is the only promising tool to hope to provide understanding!

Now to tackle such a problem one needs both excellent computational resources and powerful algorithms.

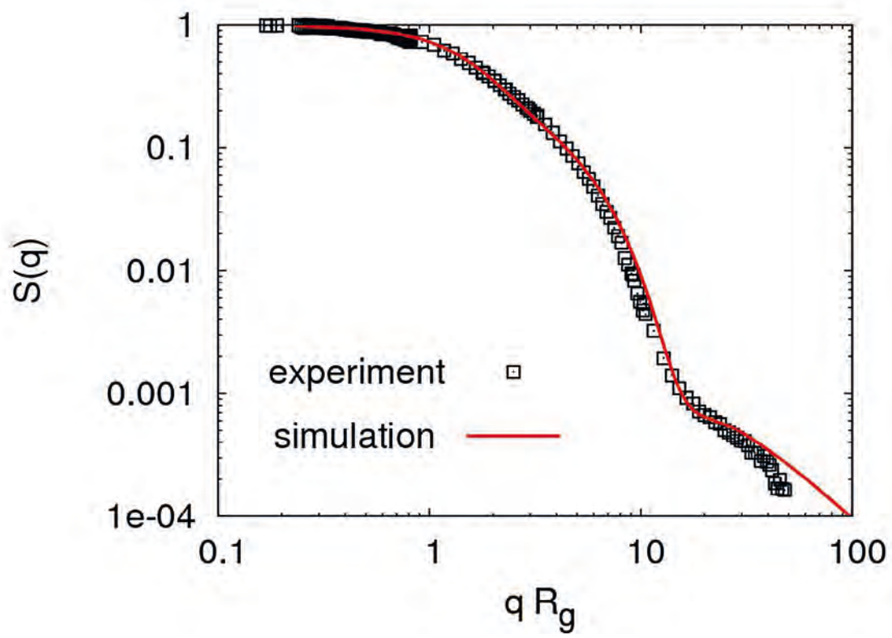

Figure 3: Structure factor S(q) of a bottle brush polymer obtained by scattering experiments [3] with a corresponding simulation [8]

The Mainz group has developed such algorithms, improving further both the so-called "Pruned Enriched Rosenbluth Method" (PERM) [6,7], and the bond fluctuation model where one combines medium-range local moves (to 26 possible sites rather than only the 6 nearest neighbors) with the "pivot" move, where a large part of a side chain (or backbone) is rotated to a new position [8]. As a resource, the cluster of the EU Network of Excellence (NoE) "SOFTCOMP" (SOFT Matter Composites), located at the Jülich Supercomputer Center (JSC), has been used.

This NoE [9] integrates many leading groups with more than 300 scientists working on the science of "soft matter" (polymers, liquid crystals, colloids, microemulsions, etc.) in Europe, and has become a focus point of research on such systems, since due to the complex structure of soft-matter systems, where interactions on many different energy scales matter, and nontrivial structure occurs on many different length scales, one needs to combine the complementary expertise of many different research groups to make progress. Thus, also the theoretical groups contributing to this research effort have combined their efforts, even putting together their computing resource in a joint large Compute Cluster (a PC cluster consisting of 500 CPU cores with a peak performance of 2.5 TFlop), which is operated by the JSC. As a group participating in the NoE and having contributed to create this Computer Cluster, we have performed a large part of our simulations on this facility (and the JUMP computer of the JSC was used as well).

Figure 3 now shows, as a typical example of our simulation results [8], a comparison of a simulation of the structure factor S(q), which describes the intensity of scattering experiments as a function of the scattering wave number q, to a corresponding experiment [3]. The experiments refers to a

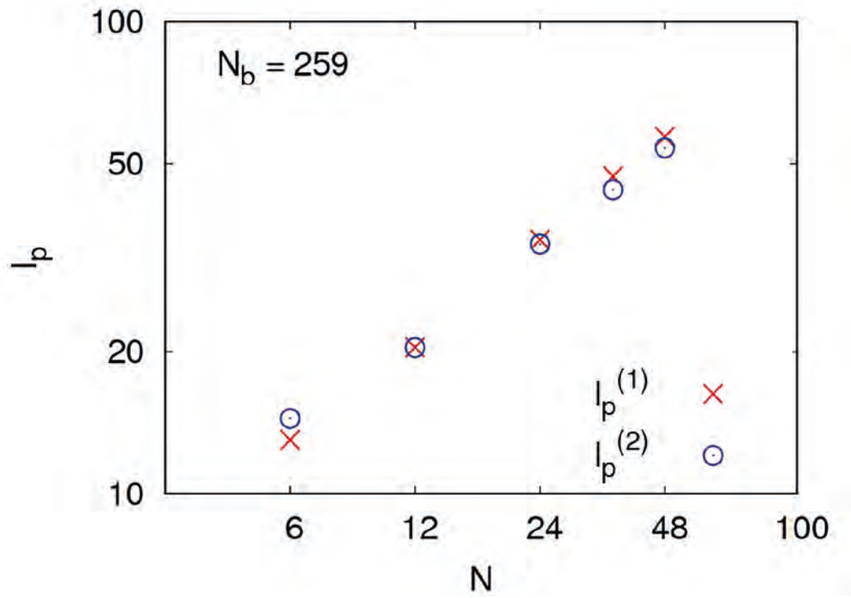

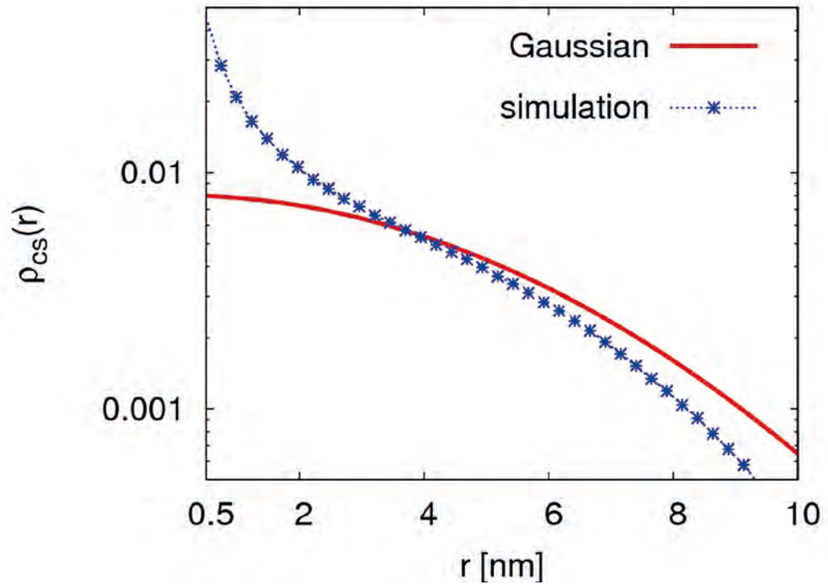

Figure 4: Left: Log-log plot of the persistence length l_p versus side chain length N, for a backbone chain length N_b = 259. The quantities $l_p^{(1)}$, $l_p^{(2)}$ denote two distinct measures of the persistence length (see [8] for technical details). Right: Cross sectional radial density profile p(r) vs. r of the simulated bottle-brush of Figure 3 and the proposed [3] Gaussian profile.

backbone containing 400 monomers, the distance between two successive monomers is about 0.26 nm, and at each monomer a side-chain containing 62 monomers each is grafted. Since each chemical monomer is a chemical group containing several atoms, the macromolecule contains of the order 10^5 atoms, and the total gyration radius of the bottle-brush polymer is about 30 nm – Figure 3 then illustrates that nontrivial structure is observed over more than two decades (in reciprocal space, as well as in real space). The coarse-grained simulation model has been chosen such, that with a suitable adjustment of the number of lattice units per nm, an almost perfect match of the structure factor is achieved. This is achieved by choosing N_b = 259, N = 48 in the simulation model (which hence contains almost 13,000 monomers, which implies a major computational effort) and 1 nm = 3.79 lattice spacings.

Of course, the value of the simulation is not just that it can reproduce experiment, but that one can go beyond experiment. The simulation allows a direct estimation of the persistence length l_p and how it scales with the side chain length (Fig. 4 left), and one can directly "measure" the radial distribution profile from the simulations (Fig. 4 right), rather than extract it indirectly from the scattering experiment [3,4] by controversial and questionable assumptions. Thus, we find that l_p varies somewhat stronger with side chain length in the simulation [8] than in the experiment, but there is also a distinct dependence of l_p on the backbone length, and thus the conflicting numbers found in various experiments can be probably reconciled. The result is intermediate between the strong increase postulated by the power law of the scaling theory [5] and the suggestion that there is no increase of l_p at all [3]. The density profile (Fig. 4 right) is similar to the experimental suggestion [3] for 2 nm < r < 10 nm, but for small distances close to the backbone (r < 2 nm) the simulation shows that the density is much higher, than expected from the postulated [3] Gaussian profile, whereas for long distances it decays somewhat faster. In this way, simulations of soft matter systems can guide the correct interpretation of corresponding experiments, and put pertinent theoretical predictions to a crucial test. Due to the large size and complex structure of these systems a massive computational effort using supercomputers is indispensable for making progress, however.

Acknowledgement

We are grateful to S. Rathgeber for providing us with the original scattering data of Reference [3] that we reproduced in Figure 3.

References

[1] **Li, C., et al.**
Angewandte Chemie, International Edition 43, p. 1101, 2004

[2] **Klein, J.**
Science 323, p. 47, 2009

[3] **Rathgeber, S., et al.**
Journal of Chemical Physics 122, pp. 124-904, 2005

[4] **Zhang, B., et al.**
Macromolecules 39, p. 8440, 2006

[5] **Fredrickson, G. H.**
Macromolecules 26, p. 2825, 1993, 2007

[6] **Hsu, H.-P., et al.**
Macromolecular Theory and Simulations 16, p. 660, 2007

[7] **Grassberger, P.**
Physical Review E56, p. 3682, 1997

[8] **Hsu, H.-P., Paul, W., Binder, K., Rathgeber, S.**
in preparation

[9] **http://www.eu-softcomp.net**

• Hsiao-Ping Hsu
• Wolfgang Paul
• Kurt Binder

Institute for Physics, Johannes Gutenberg University Mainz

S(o)OS – Service-oriented Operating Systems

Processor and network architectures are making rapid progress with more and more cores being integrated into single processors and more and more machines getting connected with increasing bandwidth. Processors become heterogeneous and reconfigurable, thus allowing for dynamic adaptation to specialized needs. In future, thousands of billions of devices may be connected to form a single computing unit.

Current programming models are insufficient to deal with this development, as they are generally too tightly coupled with the underlying device structure and types. Further to this, no current operating system is capable of dealing with the high degree of scalability and potential variability as is requested by future computing environments – be it High-Performance Computing or even local desktop PCs. Thus, in order to make efficient use of the available resources, application developers have to put considerable effort into structuring code and algorithm, as well as its adaptation to the respective resource environment. This expertise is however not available to average developers and such specialization also contradicts the potential flexibility of the resource infrastructure. Instead, an ideal application should make best use of the available infrastructure, exploiting resources specifically dedicated to the code's purpose, but not ignoring resources that nonetheless contribute to the execution performance – ideally in a way that relieves the developer from having to account for the different resource types.

The main obstacle towards heterogeneous scalability and at the same time towards richer programming models consists in an out-dated operating system architecture that takes a monolithic approach towards dealing with applications and resources. Scalable systems, however, need to take the requirements and capabilities of individual segments into consideration and provide dedicated execution support, minimizing communication and management overhead across all layers. The layered, monolithic approach, as well as most of the current typical programming models base on years, if not decades of evolutionary development following the classical hardware "frequency race", i.e. focusing in particular on single-core chipsets in homogeneous environments. With introduction of the "core race" and heterogeneous environments, new requirements have arisen that necessitate a rethinking of these classical approaches. Such new approaches require a more holistic view on all aspects related to distributed code execution in order to ensure efficiency and cater for the various differences between system environments.

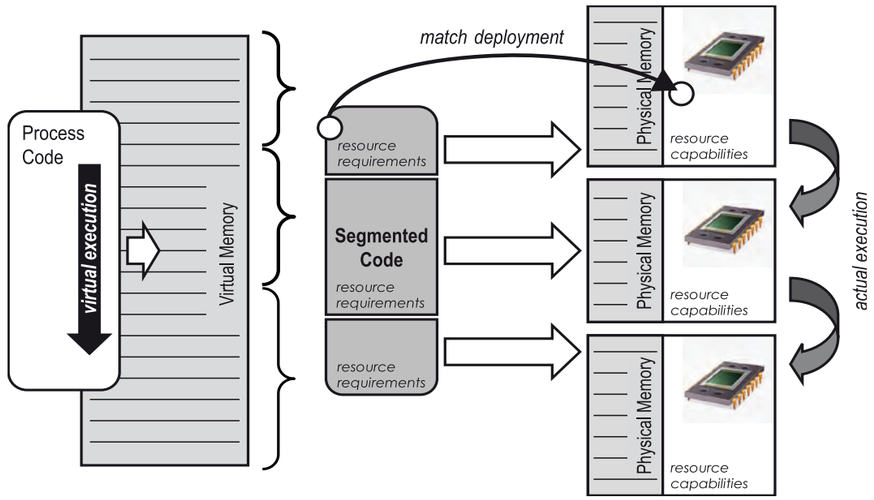

Figure 1: The S(o)OS principle of segmenting and distributing running code

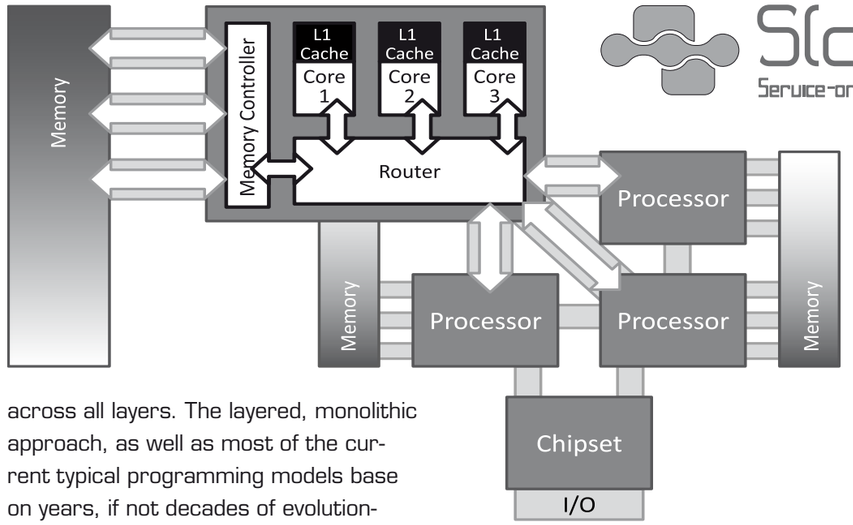

Figure 2: Architecture of a multi-core & -processor system (adapted from Intel QPI)

Since the according implications are far-reaching and require combined expertise from different areas, the EC funded the S(o)OS project which investigates into new operating system models for efficient distributed execution in heterogeneous, large-scale environments that enable new programming models for better control over code- and data distribution and its execution. Coordinated by HLRS, the S(o)OS project consisting of the European Microsoft Innovation Centre, the Institute of Telecommunications Averio, RETIS Lab of the Sant'Anna School of Advanced Studies and CTIT of the University of Twente will develop and analyze foundations for new distributed operating systems drawing from microkernel models, service oriented architectures and the strength of Grids. This will decouple the OS from the underlying resource infrastructure, thus making execution across an almost unlimited number of varying devices possible, independent from the actual hardware.

S(o)OS will allow for automatic distribution of code parts across such a resource fabric by investigating means to execute processes, threads and parallel applications across resources in a way that addresses both code requirements and resource availability, thus improving overall performance. The project aims at enabling even average developers to cope with large, widely distributed infrastructures. We therefore examine means for run-time code analysis, its segmentation and distribution across the infrastructure. This will range from automated assessment to a powerful extension for experienced developers to specify e.g. communication and relationship requirements.

The S(o)OS will run for 3 years under grant agreement number 248465, FP7 ICT-2009.8.1 "Concurrent Tera-Device Computing".

• Lutz Schubert

University of Stuttgart, HLRS

Performance Tools for Petascale Systems

• Felix Wolf

Jülich Supercomputing Centre

The growing number of cores on modern supercomputers imposes scalability demands not only on applications but also on the software tools needed for their development. At the same time, the optimization of parallel codes becomes more difficult due to increasing system complexity, creating a need for advanced and robust performance-analysis technology. The SILC project (Skalierbare Infrastruktur zur automatischen Leistungsanalyse paralleler Codes), funded under the BMBF call "HPC-Software für skalierbare Parallelrechner", therefore further reinforces the development of performance-analysis tools for High-Performance Computing, which has traditionally been a strong discipline in Germany.

Emerged from the Virtual Institute – High Productivity Supercomputing, a collaboration of leading German HPC tool builders, the goal of the SILC project is the design and implementation of a scalable and easy-to-use performance measurement and monitoring infrastructure for supercomputing applications. The infrastructure will serve as a common basis for the performance tools Vampir, Scalasca, and Periscope, which are developed by research groups in Dresden, Jülich, and Munich, respectively. Vampir is an interactive trace browser whose particular strength is the detailed visualization of the interactions between the different processes of a parallel program, offering highly flexible views to the user. Scalasca, which has been specifically designed for large-scale systems, integrates efficient performance summaries with the ability to automatically identify wait states that occur in simulation codes, for example, as a result of unevenly distributed workloads. Whereas the first two tools analyze the performance data postmortem, that is, after the parallel program has been terminated, Periscope characterizes the performance properties of an application and quantifies associated overheads already at runtime. In addition to increasing their scalability, making the three tools rest on this joint basis will allow them to interoperate more easily and simplify their installation. The enhanced tool suite will be used to tune the performance of academic and industrial simulation programs especially from the Gauss-Alliance, preparing them for the new peta-scale era.

The academic project partners in SILC are TU Dresden as the coordinator, the Jülich Supercomputing Centre, RWTH Aachen University, and TU Munich. Aachen, which plays an active role in the OpenMP Architecture Review Board, plans to equip the infrastructure with support for the more advanced features defined in the OpenMP standard. Industrial partners include the GNS mbH, a private company that specializes in services related to metal forming simulations, such as mesh generation for complex structures and finite element analyses, and – as an associate partner – the GWT-TUD GmbH, which is the commercial distributor of Vampir.

For more information see http://www.vi-hps.org/projects/silc/

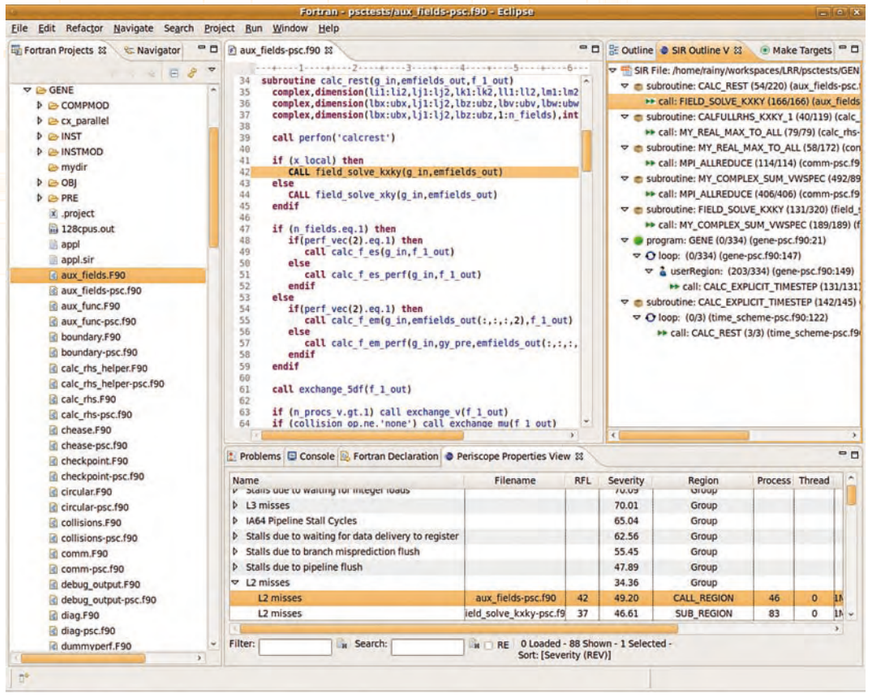

Figure 1: Periscope integrating performance analysis with application development within the Eclipse IDE

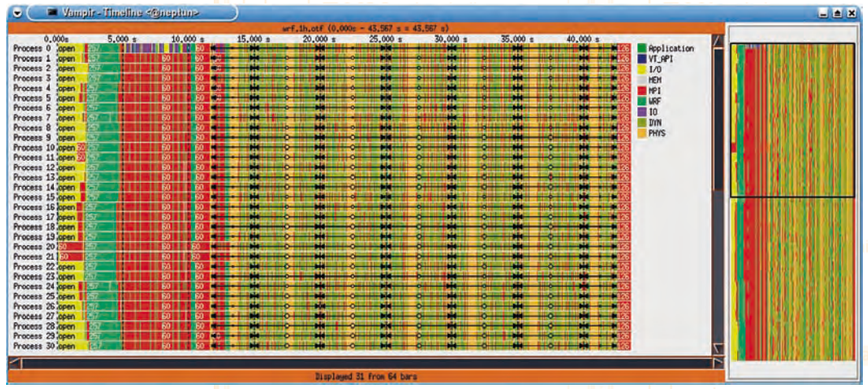

Figure 2: Vampir timeline display zoomed in on a subset of the application processes

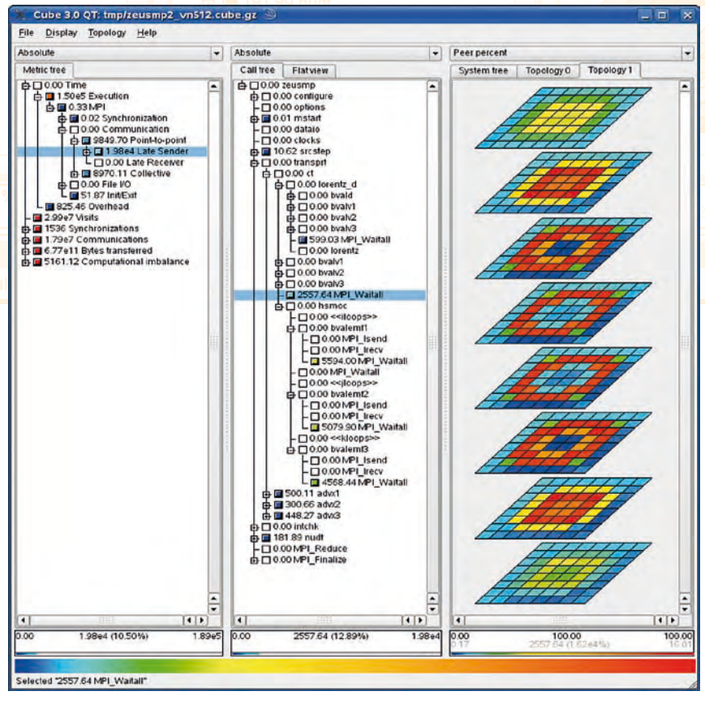

Figure 3: Scalasca showing the distribution of wait states across a three-dimensional process topology

Large Scale Simulations of Realistic Flow Problems

The simulation of many realistic engineering problems requires enormous computing power to cope with the complexity of the simulated scenarios. Especially fluid simulations are amongst the most compute intense engineering applications. In this article we present the state of the art of two currently studied engineering problems in the field of fluid simulation: the simulation of particulate flows and the simulation of free surface flows for the optimization of low temperature fuel cells.

For the simulation of fluids there exist several methods. Our simulations are based on the Lattice Boltzmann Method (LBM) [5], a mesoscopic simulation approach that is very well suited for the simulation of complex geometries (as required for both problems) and that can be efficiently implemented for the purpose of large scale simulations. The basic approach of LBM is the discretization of the computational domain by a finite number of fixed lattice nodes that contain information about the fluid particles at a specific position at a certain time step. Our C++ library for LBM simulations is the waLBerla framework [2], a software package specifically designed for large scale LBM simulations for various CFD applications.

Figure 1: Computational Fluid Dynamics applications of the waLBerla framework

Large Scale Simulation of Particulate Flows

The first example for large scale fluid simulations of engineering problems are particulate flows, such as for instance sedimentation and fluidization processes, where fluidization deals with flow through granular media and their erosion. Simulations of particulate flows are crucial for the modeling of many natural phenomena and for the optimization of related industrial applications. In the chemical industry these flows are essential to the production of key commodity and specialty chemicals, such as polymers and pigments. However, the current understanding of the actual behavior of materials in for instance a fluidized bed is rather limited, but important for the quality of the industrial products. It is very difficult to predict the complex mass and other flow quantities within the bed. For a detailed understanding of the physical phenomena within particulate flows, real world scenarios have to be simulated, which have to deal with millions of interacting particles. To this end, many simulation approaches use particle methods, as for instance molecular dynamics, or particle hydrodynamics, to simulate the immersed particles as point masses.

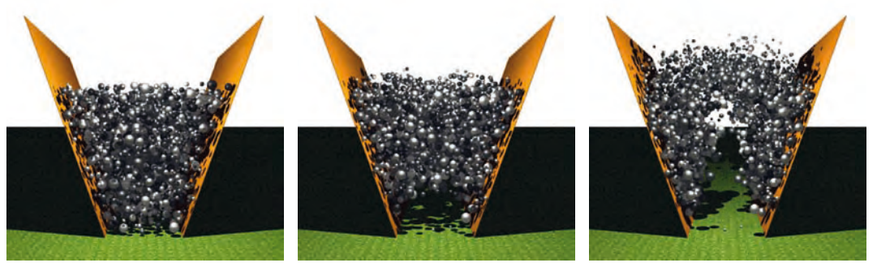

Figure 2: Fluidization simulation of 2,500 individual particles in a hopper

However, these approaches abstract from the real geometry of the particles and for example cannot take all occurring physical effects into account, such as frictional collisions between particles and accurate fluid-particle interactions. A more involved approach taken by us is the full resolution of these objects as three-dimensional, geometrically and individually modeled objects. This approach allows for a more realistic and more accurate two way coupling process between the particles and the fluid: the fluid flow exerts hydrodynamic forces on the particles and the movement of the particles directly affects the fluid. The calculation of hydrodynamic forces acting on the particles can be implemented very efficiently in LBM simulations due to the kinetic origin of the LBM by a simple momentum exchange. However, the accuracy of this approach is directly related to the resolution of the particles. In order to calculate accurate forces, the particles must be resolved with at least six to ten lattice nodes in diameter. This fact makes the requirement for a supercomputer obvious: if we want to simulate several million particles, a realistic simulation requires at least a billion lattice nodes, resulting in a huge computational and memory demand.

Since the immersed particles are considered rigid and undeformable, our approach to simulate particulate flows with accurately resolved particles is to couple the 3D LBM fluid simulation with a multibody dynamics simulation [3]. For this purpose, our physics engine pe [4] is used, which is to our knowledge the only large scale rigid multibody dynamics simulation framework. pe has been used to simulate up to 1.1 billion rigid bodies on up to 9,120 cores.

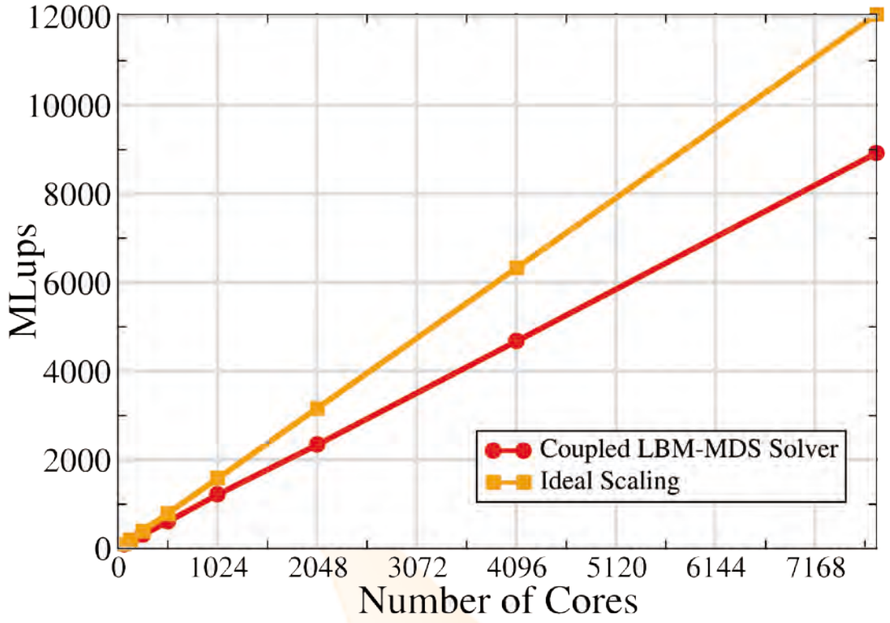

Figure 3: Weak scaling results for the coupled LBM-MDS particulate flow simulation

For our experiments we used the HLRB II supercomputer at the LRZ in Munich with up to 9,728 Intel Itanium II Montecito cores and up to 39 TBytes of main memory. Figure 3 shows some scaling results of a waLBerla simulation of a particulate flow. The performance unit for our measurements are million lattice node updates per second (MLUPS). For this scaling experiment, the number of lattice nodes per core was kept constant at 8 million lattice nodes, where the number of cores was increased up to 7,800 cores. The largest simulation performed in this way involved 35 million immersed and fully resolved particles using a total of 62 billion lattice nodes, a simulation experiment that to our knowledge has never been attempted before. The performance scales nicely with the number of cores and exhibits a parallel efficiency of 74%. This experiment demonstrates that detailed particulate flow scenarios with several million accurately resolved particles are already manageable on an appropriately sized supercomputer.

Large Scale Free Surface Simulations

The second example for a large scale engineering problem is the simulation of free surface and multiphase flows. These kinds of flows are frequently occurring in many important applications. One of the applications of interest are fuel cells, where it is necessary to accurately simulate the mixture of gases and liquids in a porous medium.

There are various types of fuel cells. In high temperature fuel cells, different gases are oxidized in a porous media to create an electric current. Here, we use multi-component extensions to the LBM in order to simulate the gas components, e.g. hydrogen, oxygen and water vapor. In low temperature fuel cells, such as the so-called proton exchange membrane fuel cell (see Figure 4), the operating temperature often is below boiling point, such that not only vaporized, but also liquid water can occur. Since the pores of the gas diffusion layer (GDL), where this reaction takes place, have a size of only several micrometers, liquid water may obstruct the gas flow and decrease the electrical performance of the cell. Moreover, in the long run erosion effects and chemical reactions fatigue the materials and degrade the durability of the fuel cell. In order to overcome such problems, the liquid water has to be evacuated from the GDL as efficiently as possible. Major influence on liquid water transport is imposed by capillary effects due to the pore structure of this layer, such that the water is “pressed out” into a channel. In order to optimize the structure and the material properties of the porous media, analysis and understanding of the behavior of the liquid water inside the GDL is supported by simulations, which is especially complex for low temperature fuel cells, because the capillary effects have to be modeled and the real structure has to be taken into account.

The simulations of the liquid water transport in the GDL of a low temperature fuel cells are carried out by using an extension of the LBM solver that is capable of dealing with liquids that have a free surface to a gas phase. In the waLBerla framework, this type of two-phase simulation is handled by not simulating the gas phase explicitly, but only taking its effect on the liquid phase into account by calculating the gas pressure and the surface tension it exerts on the liquid. By a special treatment at wall boundaries, also the material-specific contact angle is taken into account. The modeling of the capillary forces requires a high resolution in the micrometer scale pores of the porous medium. The simulation of a realistically sized section of the GDL therefore requires several billion lattice nodes in a LBM simulation. Additionally, this resolution influences the time step of the simulation. For realistic simulations, it is necessary to perform hundreds of thousands of time steps on these large lattices.

One major difficulty for free surface simulations of this size is the information exchange about the current gas pressure among the processes. Due to the size of the gas regions, this information exchange may involve an arbitrary number of processes and there might be a considerable decrease of the computation performance. Only recently, an efficient communication algorithm was developed and incorporated in the waLBerla framework to handle this problem as locally as possible and therefore to efficiently solve this kind of problem on a supercomputer [1].

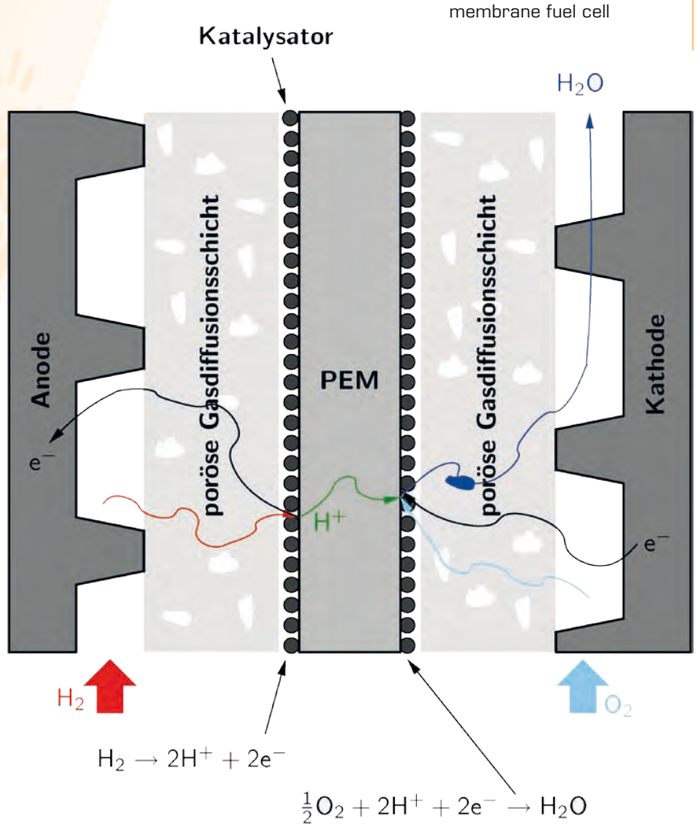

Figure 4: Schematic view of the main components of a proton-exchange membrane fuel cell

Figure 5: Scaling results of several rising bubble scenarios

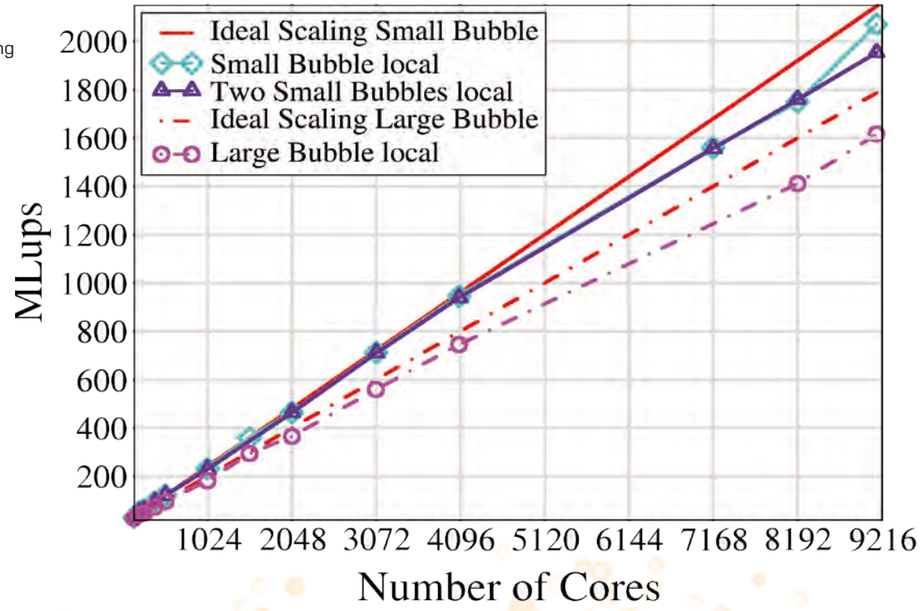

Figure 5 shows scaling results of the free surface method in the waLBerla framework, carried out on the HLRB II. Compared to a former implementation, the optimized, more local communication improves the scaling scenarios tremendously. On HLRB II, we could show that even with up to 9,152 processes the algorithm scales well for different scenarios. The same algorithms can also be used for the simulation of foams. Foams are used in an increasing number of industrial processes since they improve thermal properties, as for instance foams used for isolation purposes, and/or the mechanical stability of materials, as for instance metal foams. Figure 6 gives an impression of such a foam simulation that served as a starting point to understand the nature and behavior of metal foams. Another highly complex application for free surface flows is the simulation of bubbly flows. Using the same methodology as described above it is possible to simulate the dynamic of bubbles under the influence of gravity and other forces. Even the coalescence of multiple bubbles is considered. An impression of a large scale bubbly flow simulation is illustrated in Figure 7 that shows four snap shots of a simulation of 1,000 simultaneously rising bubbles.

Conclusion

Given the supercomputing power available today, it is possible to simulate realistic engineering scenarios such as sedimentation and fluidization processes or a representative sample of a low temperature fuel cell. These simulations will help to gain a deeper understanding of these processes and to gain insight that could hardly be identified by experiments.

Acknowledgments

Special thanks to the LRZ for its support and the computing privileges granted for the performance runs in this paper. This work is partially funded by the European Commission with DECODE, CORDIS project no. 213295, by the Federal Minister for Education and Research under the SKALB project, no. 01IH08003A, as well as by the "Kompetenznetzwerk für Technisch-Wissenschaftliches Hoch- und Höchstleistungsrechnen" in Bayern (KONWIHR) via waLBerlaMC.

Figure 6: Simulation of foam structures

References

[1] **Donath, S., Feichtinger, C., Pohl, T., Götz, J., Rüde, U.**
"Localized Parallel Algorithm for Bubble Coalescence in Free Surface Lattice-Boltzmann Method", Euro-Par 2009, August 25-28, 2009, Delft, Netherlands (Sips, H., Epema, D., Lin, H.-X., eds.), pp. 735-746, 2009

[2] **Feichtinger, C., Götz, J., Donath, S., Iglberger, K., Rüde, U.**
"waLBerla: Exploiting Massively Parallel Systems", Technical Report 08-2, Lehrstuhl für Informatik 10 (Systemsimulation), Friedrich-Alexander-Universität Erlangen-Nürnberg, June 2008

[3] **Götz, J., Feichtinger, C., Iglberger, K., Donath, S., Rüde, U.**
"Large Scale Simulation of Fluid Structure Interaction using Lattice Boltzmann Methods and the 'Physics Engine'", Proceedings of the 14th Biennial Computational Techniques and Applications Conference, CTAC-2008 (Mercer, G. N., Roberts, A. J., eds.), ANZIAM Journal, Vol. 50, October 2008, http://anziamj.austms.org.au/ojs/index.php/ANZIAMJ/article/view/1445 [October 29, 2008], pp. C166-C188

[4] **Iglberger, K., Rüde, U.**
"The pe Rigid Multi-Body Physics Engine", Tech. Report 09-9, Lehrstuhl für Informatik 10 (Systemsimulation), Friedrich-Alexander-Universität Erlangen-Nürnberg, May 2009

[5] **Succi, S.**
"The Lattice Boltzmann Equation for Fluid Dynamics and Beyond", Oxford University Press, 2001

- Klaus Iglberger
- Jan Götz
- Stefan Donath
- Christian Feichtinger
- Ulrich Rüde

Department of Computer Science, System Simulation, University of Erlangen-Nuremberg

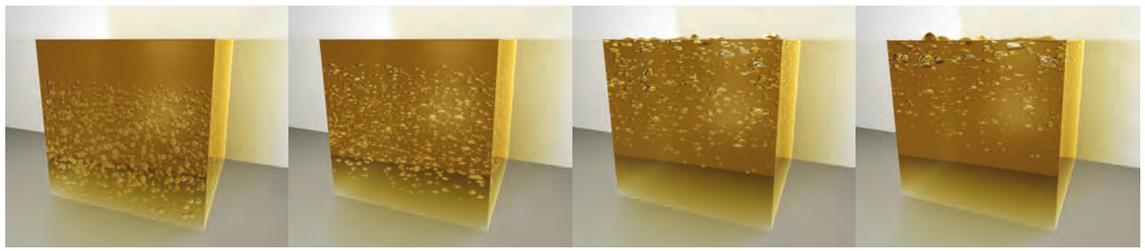

Figure 7: Simulation of 1000 simultaneously rising and coalescing bubbles

First Part of LRZ's PRACE Prototype Starts Operation

On July 30, 2009 the first part of LRZ's PRACE prototype started operation. It is an SGI ICE system based on "Nehalem" Xeon E5540 processors with an Infiniband interconnect. The theoretical peak performance of the 384 cores is 3.9 TFlops; they achieved a Linpack performance of 3.5 TFlops. A number of synthetic benchmarks were executed to evaluate architecture specific aspects of the system. To provide insight in the suitability of the system for efficiently executing scientific codes, the two applications RAxML and GADGET were also ported to the system. RAxML is a program for inference of large phylogenies with maximum likelihood. GADGET is a widely used astrophysics code for cosmological N-body/SPH simulations by Volker Springel, Max-Plank-Institute for Astrophysics, Garching, Germany. Three mathematical kernels from the Euroben benchmark suite have been selected by PRACE to assess the potential of accelerator hardware: a dense matrix-matrix multiply code, a sparse matrix vector multiplication and a 1D FFT code. These will be ported to accelerator devices presently not yet available on the system.

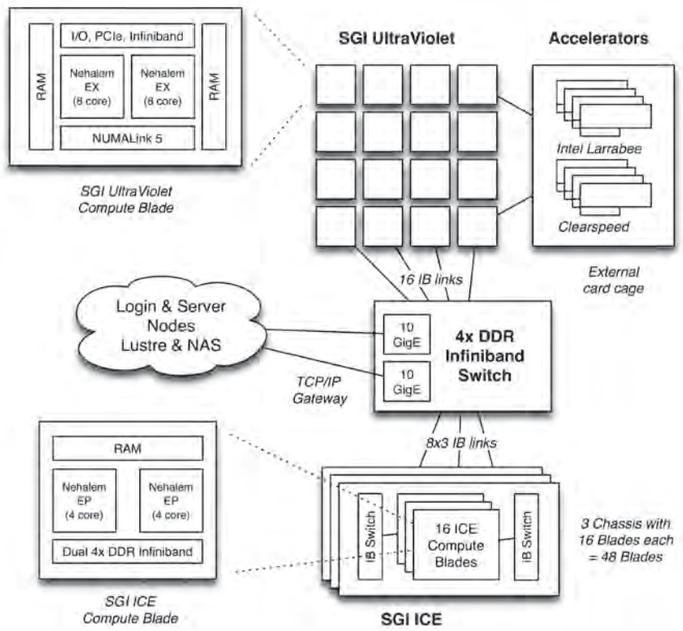

Figure 1: Layout of PRACE Prototype system

A Joint PRACE Prototype

The main objective of the PRACE prototype is the evaluation of a hybrid system architecture containing thin nodes, fat nodes and compute accelerators with a shared file system. For this evaluation, LRZ joined forces with CINES, the "Centre Informatique National de l'Enseignement Supérieur", an important french HPC Centre in Montpellier. The computing centres have agreed to install separate parts of the prototype: CINES added 32 Clearspeed e710 Accelerator Boards to an SGI ICE system comprised of 256 "Nehalem" EP cores while LRZ is in the process of installing a system based on thin nodes (ICE) and fat nodes (UltraViolet). Both segments of the LRZ system will be connected by an Infiniband switch. Additionally, Clearspeed and "Larrabee" cards will be evaluated with respect to their potential as compute accelerators.

A Truly Hybrid System

This joint prototype allows to evaluate a number of innovative components: the new Intel "Nehalem" EP and EX processors, SGI's new NUMAlink 5 interconnect, 4x QDR Infiniband networks,

Figure 2: Proving this is an ICE

the Lustre file system and the modern accelerator devices. It will also allow gaining first experience with hybrid systems that are candidates for future multi-petaflop architectures. LRZ is currently preparing the tender documents for the procurement of its next flagship system, which will be installed in the 2011 timeframe. While hybrid systems have been introduced in HPC already a few years ago to satisfy the various needs of the increasing HPC community, managing these systems is not easy. To use them efficiently requires additional complexity in the scheduler, in the operating and in the programming environment. Much research is presently underway to exploit the full potential of hybrid systems.

An Example for "Green-IT"

The usage of the individual components can be optimized in several ways; one of them is the energy footprint of an application or an application mix. To reduce this footprint, computational accelerators are currently en vogue. Their special design, which avoids all unnecessary "chip intelligence" and big caches, allows suitable applications to run on these devices with low power consumption at tremendous speeds. The Clearspeed accelerator cards have been the most energy efficient accelerator boards. While everyone knows about the potential of using computational accelerators to make HPC "greener", their programming languages, paradigms and tools, which are quite different from the standard C/C++ or Fortran mixed with OpenMP or MPI programming environment, have prevented them from being widely spread at this time. Intel's announcement of the "Larrabee" GPU which consists of standard CPU cores instead of the highly tailored GPU cores of previous devices, received much attention from the HPC community. It is hoped that the regular Intel software tool chain can be used to program these devices for higher programmability, flexibility and performance, which is extremely necessary in environments where many legacy codes, consisting of hundreds of thousands of lines of code, are being used. In HPC software has usually outlived the hardware.

• Iris Christadler

Leibniz Supercomputing Centre

Autumn 2009 • Vol. 7 No. 2 • inSiDE

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

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

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

Contact:
Leibniz Supercomputing Centre

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

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

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

Compute servers currently operated by LRZ are

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
HRLB II: SGI Altix 4700 Intel IA64 19 x 512-way	9,728 Cores 39 TByte	62.3	Capability Computing	German Universities and Research Institutes, DEISA
PRACE Prototype SGI ICE Intel Nehalem 8-way	348 Cores 1.1 TByte	3.8	Capability Computing	German Universities and Research Institutes, PRACE
Linux-Cluster SGI Altix 4700 Intel IA64 256-way	256 Cores 1 TByte	1.6	Capability Computing	Bavarian Universities
Linux-Cluster SGI Altix 3700 BX2 Intel IA64 128-way	128 Cores 0.5 TByte	0.8	Capability Computing	Bavarian Universities
Linux-Cluster Intel IA64 2-, 4- and 8-way	220 Cores 1.1 TByte	1.3	Capacity Computing	Bavarian and Munich Universities
Linux-Cluster Intel Xeon EM64T AMD Opteron 2-, 4-, 8-, 16- and 32-way	3,156 Cores 11.6 TByte	19.1	Capacity Computing	Bavarian and Munich Universities, D-Grid
Linux-Cluster Intel Xeon EM64T AMD Opteron 4- and 8-way	1,188 Cores 2.4 TByte	6.1	Capacity Computing	LHC Tier 2 Computing Grid

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

HLRS

Based on a long tradition in supercomputing at Universität Stuttgart, HLRS was founded in 1995 as a federal Centre for High-Performance Computing. HLRS serves researchers at universities and research laboratories in Germany and their external and industrial partners with high-end computing power for engineering and scientific applications.

Operation of its systems 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 a variety of systems can be provided to its users.

In order to bundle service resources in the state of Baden-Württemberg HLRS has teamed up with the Computing Centre of the University of Karlsruhe and the Centre for Scientific Computing of the University of Heidelberg in the hkz-bw (Höchstleistungsrechner-Kompetenzzentrum Baden-Württemberg).

Together with its partners HLRS provides the right architecture for the right application and can thus serve a wide range of fields and a variety of user groups.

Contact:
Höchstleistungsrechenzentrum
Stuttgart (HLRS)
Universität Stuttgart

Prof. Dr.-Ing. 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 hybrid NEC supercomputer (SX-9 / Intel Nehalem / NVIDIA S1070)
(Photos: HLRS)

Compute servers currently operated by HLRS are

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
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
Intel Cluster	480 Intel cores 256 GByte memory	3	Technical Computing	Research Institutes and Industry
AMD Cluster	288 AMD cores 1.6 TByte memory	3.7	Technical Computing	Research Institutes and Industry

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

JÜLICH
FORSCHUNGSZENTRUM

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

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

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

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

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

Contact:
Jülich Supercomputing Centre (JSC)
Forschungszentrum Jülich

Prof. Dr. Dr. Thomas Lippert
52425 Jülich
Germany

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

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

Compute servers currently operated by JSC are

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
IBM Blue Gene/P "JUGENE"	72 racks 73,728 nodes 294,912 processors PowerPC 450 144 TByte memory	1002.6	Capability computing	German Universities, Research Institutes and Industry
Intel Linux CLuster "JUROPA"	2,208 SMT nodes with 2 Intel Nehalem-EP quad-core 2.93 GHz processors each 17,664 cores 52 TByte memory	207	Capacity and Capability Computing	German Universities, Research Institutes and Industry
Intel Linux CLuster "HPC-FF"	1,080 SMT nodes with 2 Intel Nehalem-EP quad-core 2.93 GHz processors each 8,640 cores 25 TByte memory	101	Capacity and Capability Computing	EU Fusion Community
IBM Cell System "JUICEnext"	35 Blades 70 PowerXCell 8i processors 280 GByte memory	7	Capability computing	German Research School
AMD Linux Cluster "SoftComp"	125 compute nodes 500 AMD Opteron 2.0 GHz cores 708 GByte memory	2.5	Capability computing	EU SoftComp Community
AMD Linux Cluster "JUGGLE"	44 compute nodes 176 AMD Opteron 2.4 GHz cores 352 GByte memory	0.85	Capacity and capability computing	Selected D-Grid Projects

17th EuroMPI Conference at HLRS

EuroMPI 2010 will be held at Stuttgart, Germany on September 12-15, 2010. HLRS has a long tradition in MPI and has contributed to the definition of the standard and to the proliferation of MPI over the years. It was therefore only natural to accept an invitation to host the next EuroMPI meeting at Stuttgart.

The 17th European MPI Users' Group Meeting will be a forum for users and developers of MPI and other message-passing programming environments. Through the presentation of contributed papers, poster presentations and invited talks, attendees will have the opportunity to share ideas and experiences to contribute to the improvement and furthering of message-passing and related parallel programming paradigms.

With HLRS focusing on industrial co-operations and applications EuroMPI will aim at bringing science and industry in the field closer together. The potential of MPI for industrial applications will be one of the issues discussed at the HLRS. Furthermore the key questions of scalability will be addressed. Handling one million cores is a task that is beyond currently available implementations and the workshop will be a good chance to discuss new ideas and strategies.

The link to industry will also be given in the planned social event. The Mercedes Museum at Stuttgart is the perfect place to get an overview of more than 100 years of automobile history embedded in the overall world history. This visit will bring us close to the industrial heart of Stuttgart that beats for cars and everything automotive engineers can imagine.

eeClust Project Started in April

The goal of the eeClust project (Energy-Efficient Cluster Computing), funded by the German Ministry of Education and Research under the call "HPC-Software für skalierbare Parallelrechner", is to determine relationships between the behavior of parallel programs and the energy consumption of their execution on a compute cluster.

Based on this, strategies to reduce the energy consumption without impairing program performance will be developed. Project partners are the University of Hamburg (coordinator), Dresden University of Technology (TUD/ZIH), ParTec Cluster Competence Center GmbH, and the Jülich Supercomputing Centre of Forschungszentrum Jülich GmbH.

In principle, this goal can be achieved when the energy saving mode for as many hardware components as possible can be activated for those periods of time when they are not used. Modern hardware and operating systems already use these mechanisms based on simple heuristics, but without knowledge about the execution behavior of the applications currently executing. This naturally has a high potential for wrong decisions.

The project will develop enhanced parallel programming analysis software based on the successful Vampir (Dresden) and Scalasca (Jülich) software tools which in addition to measuring and analyzing program behavior will be enhanced to also record energy-related metrics.

Based on this new energy efficiency analysis, the users can then insert energy control calls into their applications which will allow the operating system and the cluster job scheduler to control the cluster hardware in an energy-efficient way. The necessary software components will be developed by the University of Hamburg and ParTec.

The effectiveness of the proposed strategy will be evaluated with the help of a small cluster testbed with special energy measurement and control components and synthetic and realistic benchmarks which also need to be developed in the course of the project.

Colloquium in Honour of Horst Rollnik

Prof. Dr. Dr. h.c. mult. Horst Rollnik is one of the founding fathers of the Höchstleistungsrechenzentrum (HLRZ), later renamed John von Neumann Institute for Computing (NIC), that supports scientific research and the peer-reviewed allocation of computing time on the Jülich supercomputers. For many years, he shaped NIC as head of its Scientific Council. Since 1997, he has organized and headed informal meetings of the three national supercomputing centres in Garching (LRZ), Stuttgart (HLRS), and Jülich (JSC). These meetings were the basis for developments that culminated in the creation of the Gauss Centre for Supercomputing.

Prof. Rollnik's contributions to the progress of scientific computing in Germany were honoured in a colloquium on May 14, 2009 at JSC in Jülich. Talks were given by Prof. Bachem, Chairman of the Board of Directors of Forschungszentrum Jülich, by several leaders of earlier NIC research groups (Prof. Gutbrod, Prof. Schilling, Prof. Wolf), and by the former head of the Zentralinstitut für Angewandte Mathematik (ZAM, now Jülich Supercomputing Centre, JSC), Prof. Hoßfeld. They addressed the history of NIC and the research performed and supported through it and its research groups as well as the future of NIC. Prof. Lippert, director of JSC, presented Prof. Rollnik with a board from one of the first supercomputers at Jülich as a souvenir. This June the NIC Scientific Council awarded him the status of an honorary lifetime member.

Some more information and a picture gallery of this event can be found at http://www.fz-juelich.de/nic/events/rollnik_colloquium.html

From left to right: Prof. Häfele, Prof. Rollnik, Prof. Bachem, Prof. Lippert

10th Guest Student Programme at JSC

During summer 2009 JSC hosted its traditional 10-week guest student programme. Within this programme students from the natural sciences, engineering, computer science and mathematics have the opportunity to familiarize themselves with different aspects of scientific computing. Mainly supported by staff members of JSC, the participants worked on various topics in computational science including mathematics, physics, chemistry, software development tools, visualization, grid computing, operating systems and communication. Special emphasis was placed on the use of the newly installed JUGENE and JUROPA supercomputers at JSC, which have raised the attraction of the programme this year. A training course in parallel programming gave an introduction of capabilities and challenges of these HPC systems.

During a concluding colloquium the guest students had the opportunity to present and discuss their work with other students and scientists. More detailed results of each project were collected and published in a JSC report, which is also available online.

It is already the 10th year that young and promising scientists of European universities had the chance to visit JSC and inform themselves about ongoing research in scientific computing. A recent survey highlights the mutual benefits of this programme: professors and participating students alike reports positive impact on their academic research, while JSC scientists also profit from new insights and collaborations.

LRZ Presents Talk and Live Demo at the Ars Electronica Festival

Ars Electronica stands for the world's leading media arts festival, a superlative state-of-the-art museum, and an innovative R&D facility. The Ars Electronica Festival, the Ars Electronica Center – Museum of the Future, and the Ars Electronica Futurelab are big draws that attract visitors, tourists, clients and associates from throughout Upper Austria and around the world.

The festival for digital art and media culture as basis of the Ars Electronica is held annualy in Linz since 1979. This year, the festival took place from 3-8 September with the motto "Human Nature". It is the most important festival for digital art world wide, bringing together about 500 artists and speakers from 25 countries, and more than 100,000 visitors from all over the world.

At this years 30th anniversary of the Ars Electronica Festival, the Leibniz Supercomuting Centre of the Bavarian Academy of Sciences and Humanities presented the demonstration: "In-silico experiments – live from LRZ" as part of the Pixelspaces Symposium of the Ars Electronica Futurelab. At noon, LRZ presented a talk in the panel "Deep Space as platform for projects based on realtime data visualization".

In addition to the talk, the LRZ could present their demo two times as part of the nightline "science cinema". The demonstration consists of a real-time simulation of a protein molecule re-create a real experiment live inside the computer ("in-silico experiment"). There, the molecular dynamic simulation of a protein molecule was started on 500 cores of the supercomputer HLRB II in Garching and simultaneously visualized online in stereoscopic 3D in Linz on the 16 m wide and 9 m high screen of the AEC Deep Space. The simulation is interactive and can be inspected and interacted with in real time. Forces can be applied interactively by the observer and the simulation on the remote supercomputer reacts accordingly. The scientist is no longer a passive observer, but becomes an actor in the molecular simulation. This demonstration allowed the public a glance into the future of life sciences, where supercomputers help scientists to conduct important research and allow us to make the invisible visible.

Workshop on Advanced Debugging with TotalView

On August 12, 2009, Jülich Supercomputing Centre in cooperation with TotalView Technologies organized a full-day workshop on "Advanced Debugging with TotalView". The meeting attracted more than 50 participants from all over Germany from as far as Potsdam or Regensburg.

The morning programme concentrated on basic usage features like startup, data and process control, and data monitoring and visualization. In the afternoon, more advanced features like parallel debugging, remote display, batch debugging, and memory debugging were presented. TotalView provided materials that included a CD that is bootable on a PC or Mac, that provides a Linux-based environment with TotalView (and sample programs) for a hands-on workshop "lab". Participants used this to follow the demonstrations of the presenter on their own laptop and to solve small exercises under guidance of the experts.

The CD contains a temporary license for TotalView valid for a few months, as well as extensive documentation, so that the participants can continue to explore and learn debugging with TotalView features back home.

12th HLRS Results & Review Workshop

The 12th Results and Review Workshop of the HLRS was held October 8-9, 2009 and brought together more than 50 participants from german research institutions, the steering committee and scientific support staff of HLRS. A whole of 41 scientists from all over Germany presented their results which they had achieved over the last year on the supercomputers at HLRS and the Karlsruhe HPC systems as talks and posters at the workshop. Their sophisticated papers were selected in advance from the steering committee out of the yearly supercomputer project reports. State-of-the-art scientific simulations on the supercomputer systems again emphasized the world-class research done at the centers obtaining outstanding results in achieving highest performance for production codes which are of particular interest for both scientists and engineers.

The presentations covered all fields of computational science and engineering ranging from Computational Fluid Dynamics and reacting flows via chemistry, solid state physics, nanotechnology, astrophysics to climate research, structural mechanics and earth sciences with a special emphasis on industrially relevant applications. These outstanding results of research problems in science and engineering are milestones of modern highest level scientific computing also by using exceptional complex models and methods, and therefore provide an excellent overview of recent developments in HPC and simulation techniques.

Every year the steering committee, a panel of twelve top-class scientists and responsible for project proposal reviews, appreciated the high quality of the work carried out as well as the spectacular scientific results and the efficient usage of supercomputer resources by awarding the three most outstanding contributions of the workshop, combining the quality of the paper and the lecture. These three papers were awarded the traditional Golden Spike Award. The laureates of this year are (from left to right):

- Björn Selent from the Institute of Aerodynamics and Gasdynamics of the University of Stuttgart for the work on direct numerical simulation of jet in cross-flow actuators

- Volker Gaibler from the Max-Planck Institute for Extraterrestrial Physics in Garching for the work on numerical models for emission line nebulae in high redshift radio galaxies

- Juliane Schwendike from the Institute for Meteorology and Climate Research of the University of Karlsruhe for the work on modelling convection over West Africa.

34th IDC HPC User Forum at HLRS

For the third time in a row HLRS hosted an IDC User Forum at Stuttgart. The collaboration of HLRS and IDC reflects the common interest in supercomputing applications with a link to the industrial usage of High-Performance Computing systems.

The 34th IDC HPC User Forum was held October 5-6, 2009. About 60 international participants joined and discussed over two days the current problems of High-Performance Computing. The program was focused on the practical aspects of High-Performance Computing.

Tom Sterling from the Lousiana State University – a Faculty Associate at California Institute of Technology, a CSRI Fellow for Sandia National Laboratories, and a Distinguished Visiting Scientist at Oak Ridge National Laboratory – gave a brilliant overview of the programming of supercomputers.

Erich Schelkle – the director of the Automotive Simulation Center Stuttgart (ASCS) – presented an industrial view of the problems in High-Performance Computing. He explained the structure and goals of ASCS and laid out the fundamental technical problems that High-Performance Computing can help to solve in automotive engineering.

A number of talks gave an overview of High-Performance Computing in Europe and US where the presentation about the European flagship project PRACE was of special interest to the audience. The program was enriched by a number of talks about future strategies of international centers and vendors.

For more information see http://www.hpcuserforum.com/

HLRS Scientific Tutorials, Workshop Report and Outlook

With the paradigm-shift to multi-core CPUs, new programming languages evolve. In May 2009, HLRS started with an own course on the partitioned global address space (PGAS) languages **Unified Parallel C (UPC)** and **Co-Array Fortran (CAF)**. The series of PGAS courses will be continued in 2010.

The question of programming on clusters of multi-core nodes was also addressed by HLRS in four tutorials at scientific conferences in 2009, as you can see in the box:

distributed memory parallelization with the Message Passing Interface (MPI) and shared memory multi-threading with OpenMP. This course will be presented twice, in March 2010 at HLRS in Stuttgart and in September 2010 at LRZ in Garching.

Another highlight is the **Introduction to Computational Fluid Dynamics**. The emphasis is placed on explicit finite volume methods for the compressible Euler equations. Moreover outlooks on implicit methods, the extension to the

Tutorials at Scientific Conferences in 2009

Hybrid MPI and OpenMP Parallel Programming. Rolf Rabenseifner, Georg Hager, Gabriele Jost. Half-day, Tutorial No. M-09 at Super Computing 2009, SC09, Portland/Oregon, USA, Nov. 14 - 20, 2009

Application Supercomputing and the Many-Core Paradigm Shift. Alice Koniges, William Gropp, Ewing (Rusty) Lusk, Rolf Rabenseifner, David Eder. Full-day, Tutorial No. S-01 at Super Computing 2009, SC09, Portland/Oregon, USA, Nov. 14 - 20, 2009

Programming Models and Languages for Clusters of Multicore Nodes. Alice Koniges, Rolf Rabenseifner, Gabriele Jost, Georg Hager. Tutorial at SciDAC 2009, tutorial day, San Diego/California, USA, June 19, 2009

Hybrid OpenMP/MPI Programming and other Models for Multicore Architecures. Gabriele Jost, Alice Koniges, Gerhard Wellein, Georg Hager, Rolf Rabenseifner, Ewing (Rusty) Lusk. Tutorial at the 21st International Conference on Parallel Computational Fluid Dynamics (Parallel CFD 2009), Moffett Field/California, USA, May 18 - 22, 2009

Next year, HLRS will have again a number of different courses and workshops. One of the flagships of our courses is the week on **Iterative Solvers and Parallelization**. Prof. A. Meister and Prof. B. Fischer teach basics and details on Krylov Subspace Methods. Lecturers from HLRS give lessons on 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.

The **Parallel Programming Workshop**, Sep. 27 - Oct. 1, 2010 at HLRS, will have three parts: The first two days of this course are dedicated to parallelization with the Message Passing Interface (MPI). Shared memory multi-threading is taught on the third day, and in the last two days, advanced topics are discussed. As in all courses, hands-on sessions (in C and Fortran) will allow users to immediately test and understand the parallelization methods. The course language is English.

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

Fortran for Scientific Computing started the first time at HLRS in October 2006 and had up to now more than 180 participants. Mainly PhD students from Stuttgart and other universities in Germany came to learn not only the basics of programming, but also to get an insight on the principles of developing high-performance applications with Fortran. The Fortran course will be repeated twice in 2010.

Additional courses are planned on **Platforms at HLRS**, and **CUDA Programming**. In the table, you can find the whole HLRS series of training courses in 2010. They are organized at HLRS and also at several other HPC institutions: LRZ Garching, NIC/ZAM (FZ Jülich), and CSCS (Manno, CH), ZIH (TU Dresden), TUHH (Hamburg-Harburg).

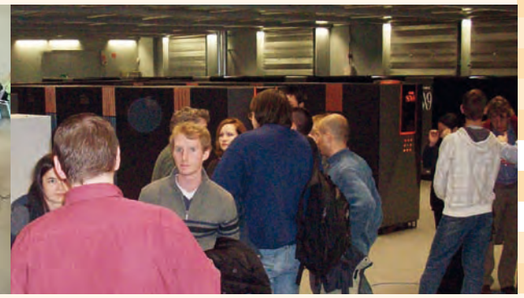

2010 – Workshop Announcements

Scientific Conferences and Workshops at HLRS, 2010
12th Teraflop Workshop (March, not yet fixed)
9th HLRS/hww Workshop on Scalable Global Parallel File Systems (April 26 - 28)
4th HLRS Parallel Tools Workshop (July, not yet fixed)
17th EuroMPI 2010 (Successor of the EuroPVM/MPI Series, September 12 - 15)
High-Performance Computing in Science and Engineering - The 13th Results and Review Workshop of the HPC Center Stuttgart (October 4 - 5)
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 22 - 25)
Introduction to Computational Fluid Dynamics (March 15 - 19, not yet fixed)
Iterative Linear Solvers and Parallelization (HLRS, March 22 - 26)
Platforms at HLRS (HLRS, March 29 - 30)
Unified Parallel C (UPC) and Co-Array Fortran (CAF) (HLRS, May 18 - 19)
4th Parallel Tools Workshop (not yet fixed)
Parallel Programming with MPI & OpenMP (TU Hamburg-Harburg, July 21 - 23)
Parallel Programming with MPI & OpenMP (CSCS Manno, CH, August 10 - 12)
Introduction to Computational Fluid Dynamics (HLRS, September 6 - 10)
Message Passing Interface (MPI) for Beginners (HLRS, September 27 - 28)
Shared Memory Parallelization with OpenMP (HLRS, September 29)
Advanced Topics in Parallel Programming (HLRS, September 30 - October 1)
Iterative Linear Solvers and Parallelization (LRZ, Garching, October 4 - 8)
Parallel Programming with MPI & OpenMP (FZ Jülich, JSC, November 29 - December 1)
Unified Parallel C (UPC) and Co-Array Fortran (CAF) (HLRS, December 14 - 15)
Introduction to CUDA Programming (not yet fixed)
Training in Programming Languages at HLRS
Fortran for Scientific Computing (March 1 - 5 and October 25 - 29)
URLs: http://www.hlrs.de/events/
https://fs.hlrs.de/projects/par/events/2010/parallel_prog_spring2010/
https://fs.hlrs.de/projects/par/events/2010/prog_lang_spring2010/

High-Performance Computing Courses and Tutorials

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

Date & Location
November 26 - 27, 2009
May 17 - 18, 2010
JSC, Research Centre Jülich

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

Webpage
http://www.fz-juelich.de/jsc/neues/termine/supercomputer/

Parallel Programming with MPI, OpenMP and PETSc

Date & Location
November 30 - December 2, 2009
JSC, Research Centre Jülich

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

Webpage
http://www.fz-juelich.de/zam/neues/termine/mpi-openmp/

HLRB and KONWIHR and Linux-Cluster Review, Results and Future Projects

Date & Location
December 8 - 9, 2009
LRZ Building, Garching/Munich

Contents
The workshop covers the scientific projects from the timeframe 2008 - 2009 and serves as a forum for the exchange of experiences between the user communities as well as between users and LRZ. Participants from new projects are also invited for presentations.

Webpage
http://www.lrz-muenchen.de/services/compute/hlrb/reviewworkshop/

Introduction to the PGAS Languages UPC and CAF

Date & Location
January 20, 2010
LRZ Building, Garching/Munich

Contents
Partitioned global address space languages have emerged as an alternative to other parallel programming models, promising a shorter development cycle due to improved programmability while keeping the performance level on par with MPI. This course introduces the parallel facilities integrated into the Fortran language (coarrays) and the C language (unified parallel C), respectively. A hands-on session allows to experiment with the new concepts, using prototype implementations on LRZ's HPC systems.

Prerequisites: Good working knowledge of Fortran and/or C.

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

Programming with Fortran

Date & Location
February 8 - 12, 2010
LRZ Building, Garching/Munich

A video conference to other sites can be set up if there is sufficient interest.

Contents
This course is targeted at scientists with little or no knowledge of the Fortran programming language, but need it for participation in projects using a Fortran code base, for development of their own codes, and for getting acquainted with additional tools like debugger and syntax checker as well as handling of compilers and libraries. The language is for the most part treated at the level of the Fortran 95 standard; features from Fortran 2003 are limited to improvements on the elementary level. Lectures are complemented by 3 hours of hands-on sessions.

Prerequisites: Basic UNIX/Linux knowledge.

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

Scientific 3D Animation with Blender

Date & Location
February 18 - 19, 2010
LRZ Building, Garching/Munich

Contents
This course provides an introduction to visualization of scientific simulation data using the open source Blender 3D animation package (www.blender.org). Participants will be taught techniques to generate impressive and professional quality animations of the data obtained from their own scientific projects. Based on an example project using a protein molecule, all steps from importing the data over geometry correction and material assignment down to illumination, key framing and post-production are discussed.

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

Parallel Programming with MPI, OpenMP and PETSc

Date & Location
February 22 - 25, 2010
ZIH, Dresden

Contents
The focus is on programming models MPI, OpenMP, and PETSc. Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the basic constructs of the Message Passing Interface (MPI) and the shared memory directives of OpenMP. The last day is dedicated to tools for debugging and performance analysis of parallel applications. This course is organized by ZIH in collaboration with HLRS.

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

Fortran for Scientific Computing

Date & Location
March 1 - 5, 2010
HLRS, Stuttgart

Contents
This course is dedicated for scientists and students to learn (sequential) programming scientific applications with Fortran. The course teaches newest Fortran standards. Hands-on sessions will allow users to immediately test and understand the language constructs.

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

High-Performance Computing Courses and Tutorials

Parallel Programming of High-Performance Systems

Date & Location
March 8 - 12, 2010
RRZE Building, Erlangen
LRZ Building, Garching/Munich
(via video conference)

Contents
This course, a collaboration of Erlangen Regional Computing Centre (RRZE) and LRZ, is targeted at students and scientists with interest in programming modern HPC hardware, specifically the large scale parallel computing systems available in Jülich, Stuttgart and Munich. Each day is comprised of approximately 4 hours of lectures and 3 hours of hands-on sessions.

Prerequisites: Good working knowledge of at least one of the standard HPC languages: Fortran 95, C or C++.

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

Introduction to Computational Fluid Dynamics

Date & Location
March 15 - 19, 2010
(Location not yet fixed)

Contents
Numerical methods to solve the equations of Fluid Dynamics are presented. The main focus is on explicit Finite Volume schemes for the compressible Euler equations. Hands-on sessions will manifest the content of the lectures. Participants will learn to implement the algorithms, but also to apply existing software and to interpret the solutions correctly. Methods and problems of parallelization are discussed.

This course is based on a lecture and practical awarded with the "Landeslehrpreis Baden-Württemberg 2003" and organized by HLRS, IAG, and University of Kassel.

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

Advanced Topics in High-Performance Computing

Date & Location
March 22 - 24, 2010
LRZ Building, Garching/Munich
RRZE Building, Erlangen
(via video conference)

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

Prerequisites: Good MPI and OpenMP knowledge as presented in the course "Parallel programming of High-Performance Systems".

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

Iterative Linear Solvers and Parallelization

Date & Location
March 22 - 26, 2010
HLRS, Stuttgart

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

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

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

Platforms at HLRS

Date & Location
March 29 - 30, 2010
HLRS, Stuttgart

Contents
The course is focused on the usage of the HPC platforms at HLRS.

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

Introduction to Molecular Modeling on Supercomputers

Date & Location
April 14 - 15, 2010
LRZ Building, Garching/Munich

Contents
The course gives an introduction into the simulation of molecules based on several software packages (e.g. NAMD, VMD) on the supercomputers at LRZ Garching. This also includes an introduction to the remote visualization services at LRZ as well as hands-on sessions.

Prerequisites: Basic Knowledge of molecular simulations.

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

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

Date & Location
May 18 - 19, 2009
HLRS, Stuttgart

Contents
Partitioned Global Address Space (PGAS) is a new model for parallel programming. Unified Parallel C (UPC) and Co-array Fortran (CAF) are PGAS extensions to C and Fortran. PGAS languages allow any processor to directly address memory/data on any other processors. Parallelism can be expressed more easily compared to library based approches as MPI.

Hands-on sessions (in UPC and/or CAF) will allow users to immediately test and understand the basic constructs of PGAS languages.

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

Usage of High-Performance Systems, Remote Visualization and Grid Facilities at LRZ

Date & Location
June 2, 2010
LRZ Building, Garching/Munich

Contents
Based on hands-on examples, an easy-to-follow introduction to the hard- and software of the cluster systems, the visualization systems and usage of the grid middleware (Globus Toolkit) is given. Grid certificates can be provided if needed.

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

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

Education in Scientific Computing

Date & Location
August 2 - October 8, 2010
JSC, Research Centre Jülich

Contents
Guest Student Programme "Scientific Computing" to support education and training in the fields of supercomputing. Application deadline is April 30, 2010. Students of Computational Sciences, Computer Science and Mathematics can work 10 weeks in close collaboration with a local scientific host on a subject in their field.

Webpage
http://www.fz-juelich.de/jsc/gaststudenten/

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

Publishers

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

Editor & Design

F. Rainer Klank, HLRS klank@hlrs.de
Stefanie Pichlmayer, HLRS pichlmayer@hlrs.de

Authors

Nikolaus A. Adams nikolaus.adams@tum.de
Michael Bader bader@informatik.tu-muenchen.de
Kurt Binder kurt.binder@uni-mainz.de
Hans-Joachim Bungartz bungartz@in.tum.de
Antoine Devesa antoine.devesa@aer.mw.tum.de
Stefan Donath stefan.donath@informatik.uni-erlangen.de
Iris Christadler iris.christadler@lrz.de
Christian Feichtinger christian.feichtinger@informatik.uni-erlangen.de
Jan Götz jan.goetz@informatik.uni-erlangen.de
Stefan Hickel sh@tum.de
Ulrich Hindenlang ulrich.hindenlang@lasso.de
Sabine Höfler-Thierfeldt s.hoefler-thierfeldt@fz-juelich.de
Hsiao-Ping Hsu hsu@uni-mainz.de
Klaus Iglberger klaus.iglberger@informatik.uni-erlangen.de
Jens Iseler jens.iseler@unibw-muenchen.de
Andreas Lesser andreas.lesser@unibw-muenchen.de
Michael Manhart m.manhart@bv.tu-muenchen.de
Michael Meyer michael.meyer@aer.mw.tum.de
Walter Nadler w.nadler@fz-juelich.de
Reinhard Niehuis reinhard.niehuis@unibw.de
Wolfgang Paul wolfgang.paul@uni-mainz.de
Michael Resch resch@hlrs.de
Ulrich Rüde ulrich.ruede@informatik.uni-erlangen.de
Ralf Schneider schneider@hlrs.de
Lutz Schubert schubert@hlrs.de
Florian Schwertfirm f.schwertfirm@muenchen-mail.de
Christian Stemmer christian.stemmer@aer.mw.tum.de
Felix Wolf f.wolf@fz-juelich.de

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