inSiDE

inSiDE • Vol. 7 No. 1 • Spring 2009

Innovatives Supercomputing in Deutschland

Editorial

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

Welcome to the German Petaflops era! While this issue of inSiDE goes to print, the Jülich Supercomputing Centre (JSC) will inaugurate Europe's first Petaflops system and make it available to its German and European users. The IBM BlueGene at Jülich has about 295.000 processors and a main memory of about 144 TB. Its highly power efficient concept is a tribute to the current GreenIT discussion and a substantial contribution to solve environmental problems in supercomputing. This is a true milestone for German supercomputing.

But supercomputing in Germany has more to offer than hardware. The German Federal Ministry of Education and Research (BMBF) will launch a second call for proposals for the development of HPC software soon. A description of several projects initiated in the first phase of this funding initiative can be found in the project section of this issue of inSiDE.

At the same time the Gauss Alliance has established a strong organizational structure. A short report describes the purpose, partners and strategy of the alliance in more details. Gauss Centre and Gauss Alliance will work closely together to strengthen HPC in Germany.

The installation of a Petaflops system at the Jülich Supercomputing Centre is also an important and positive signal for the European process of consolidation in supercomputing. Together with its European partners, the German Gauss Centre for Supercomputing is pushing towards a truly European and internationally competitive solution. The new system at Jülich takes the lead in this effort and will be followed by further European systems.

A first step towards an increase of the number of European systems has recently been made in the European project PRACE (Partnership for Advanced Computing in Europe). The prototype installations of the PRACE are now in place and will show first results soon. Details on the progress of PRACE in its first project year are reported in this issue of inSiDE.

But the goal of supercomputing is in new and excellent solutions for users in research and industry. Hence in this issue you will again find a comprehensive section about large scale applications. Eight contributions from users from all over Germany present the variety of simulations that are currently run on German HPC systems. At the same time these papers provide a good insight into the future needs for supercomputing in Germany.

The project part in this issue is a showcase of federal funding for HPC software development. Four projects (IMEMO, VisPME, STEDG, ScaFaCoS) are presented that were initiated through funding by the "HPC-Software Initiative" of the German Federal Ministry of Education and Research (BMBF). An additional project (HERMES) is funded by the BMBF through the scheme "Research for Civil Security". HERMES very well reflects the impact of HPC beyond the traditional scientific communities.

As usual, this issue includes information about new systems and events in supercomputing in Germany over the last months and gives an outlook of workshops in the field. Readers are invited to participate in these workshops.

JUGENE Becomes First Petaflop System in Europe

Forschungszentrum Jülich proudly announces the procurement of the first supercomputer in Europe capable of one petaflop or one thousand trillion operations per second. The installation of the IBM Blue Gene/P system will be completed in June 2009. "Supercomputers of this performance level are universal key technology instruments for solving the most complex and urgent scientific problems in many areas," commented Professor Achim Bachem, Chairman of the Board of Directors of Forschungszentrum Jülich. "Scientists of all disciplines use supercomputers to identify climate change, conduct research about protein folding in cells, and discover how semiconductors work or how fuel cells can be improved." "With speeds of over a petaflop/s, this new supercomputer provides the processing ability of more than 200,000 standard desktop computers," explained Professor Thomas Lippert, director of the Jülich Supercomputing Centre. "In addition to raw power, this new system will be among the most energy-efficient in the world."

This new system will help to ensure that Forschungszentrum Jülich continues to play a leading role in the global high-performance computing research community. It is the first German supercomputer system that was selected and purchased in the context of the Gauss Centre for Supercomputing (GCS) and is funded by the German Federal Ministry of Education and Research (BMBF) and the Ministry for Innovation of North Rhine-Westphalia.

The system is an extension of the existing JUGENE machine and will have the characteristics:

JUGENE system characteristics	
Type:	IBM Blue Gene/P 72 racks
Peak performance:	1 petaflop/s
Processors:	294,912 (PowerPC 450 at 850 MHz)
Compute nodes:	73,728 (4-way SMP, 2 GB memory)
Total memory:	144 TB
IO nodes:	600 (connecting to external 10 GigE)
Communication networks:	3D Torus (fully DMA capable)
Network latency:	160 ns
Bandwidth:	188 TB/s

The history of leadership-class systems for capability computing applications at FZJ started in 2005 with a one-rack Blue Gene/L, which had been extended to eight racks in 2006 and replaced by a 16-rack Blue Gene/P system – called JUGENE – with 65,536 processors and a peak performance of 222.8 TFlop/s in 2007 (inSiDE Vol. 5 no. 2, p. 46). With this system it became possible amongst other things to calculate the masses of the hadron spectrum in an ab-initio simulation with unprecedented precision – one of the ten scientific breakthroughs 2008, selected by the editorial boards of Science and Nature; further details can be found in a separate article in this volume and in the original paper [1].

Now, with JUGENE's increased capacity, Jülich is excited to see further ground-breaking scientific highlights. These might be achieved by the newly established large-scale projects within the GCS. A special call for large-scale projects has been opened up to promote highly scalable parallel applications, which require more than 5% of the potentially available CPU cycles of JUGENE. The GCS centres guarantee that large-scale projects can fully utilize the assigned resources, if necessary at the expense of normal or smaller projects. Furthermore, the advanced training of the users will be done by high-level Blue Gene scaling workshops, as successfully practiced in the past; the next workshop is planned for autumn 2009.

JUGENE is fully integrated in the Jülich supercomputer infrastructure. Its I/O nodes are connected to the fileserver JUST by four FORCE 10 E1200i switches with a maximum of 896 ports, based on 10-Gigabit Ethernet. The fileserver itself has also been upgraded. It now consists of 28 IBM Power 6 servers and 18 storage controllers with a maximum bandwidth of 20 GB/s. An increase of the disk capacity to 6 PB and a bandwidth of 67 GB/s will follow in autumn 2009. To allow the build up of such a system, major infrastructure changes had to be made to the Jülich machine room, originally opened in 2004. Although the Blue Gene/P has a small footprint of 130 m^2 and is one of the most power-efficient supercomputers in the world, the energy provisioning was enlarged to a maximum of 4.5 MW, of which 2.5 MW are necessary for JUGENE including cooling. In addition to that, water cooling had to be brought into the machine room for the newly developed hydro-cooling elements that replaced the air plenums between the Blue Gene/P racks thus reducing the air conditioning requirements by 91%. It should also be noted that the installation includes about 23 km of copper cables (Blue Gene/P data cables) and 21 km fibre cables (10-Gigabit Ethernet), plus 7.5 km Cat5 cables (1 Gigabit Ethernet) and is a masterpiece in logistics and coordination.

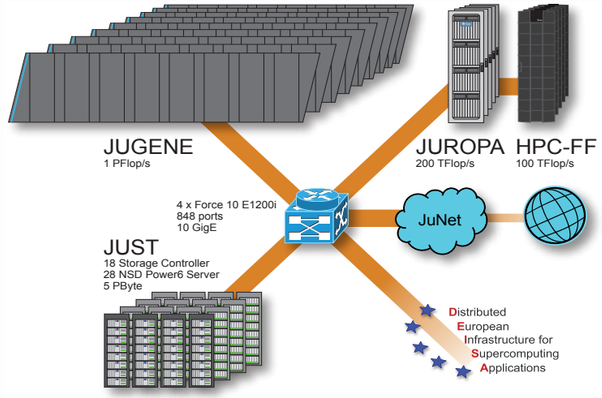

Supercomputer infrastructure at Forschungszentrum Jülich

Currently, it is planned to enter production mode with the full system at the end of June 2009. With this supercomputer Forschungszentrum Jülich will keep its European pole position in the Top 500 list of supercomputer sites, a position which it has held since November 2007.

• Jutta Docter
• Michael Stephan
• Norbert Attig

Forschungszentrum Jülich, Jülich Supercomputing Centre

References

[1] Dürr, S., et al.
Ab Initio Determination of Light Hadron Masses, Science Vol. 322, no. 5905, pp. 1224–1227, 2008

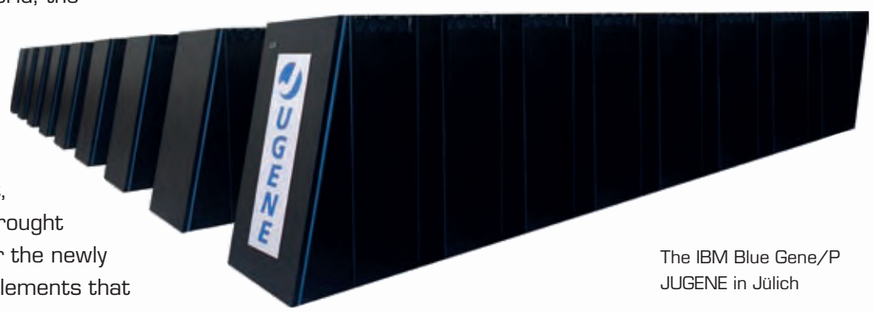

The IBM Blue Gene/P JUGENE in Jülich

PRACE – A Very Successful First Year

PRACE, the Partnership for Advanced Computing in Europe [1], is preparing the creation of a persistent, pan-European Research Infrastructure that will provide a world-class HPC service to enable world-class science. The objective of the two year project is the completion of the necessary legal, administrative, and technical work that will allow the permanent Research Infrastructure to commence operation in 2010. PRACE reached its first major milestone by successfully completing the mid-term review in early March. Highlights of the of the first year include: evaluation of applications of the current users of European HPC systems; selection of promising architectures for potential petaflop systems and deployment of prototypes; establishing PRACE as an internationally renowned actor in the global HPC scene; identifying training requirements of users wishing to exploit novel systems. To serve the users and to support their applications is the driving force for PRACE. The set of systems that PRACE intends to deploy must provide for current applications while breaking ground for new applications that target to exploit unprecedented performance levels of novel architectures.

Applications

PRACE analysed close to 70 applications from 24 systems – a snapshot of current European HPC usage – with the following objectives:

- to map them to architecture specifications that were used to select the prototype systems
- select a subset what will serve as the nucleus of a benchmark suite to evaluate the prototypes and the future production systems
- port the applications to most prototypes
- optimize and scale suitable applications

Once the Research Infrastructure is in place, users will be in a position to start exploiting the new systems in a very effective way and benefit from the experience gathered in the project.

Prototypes

Based on the analysis of applications PRACE identified a comprehensive set of architectures and selected six related prototypes [2] – the IBM BlueGene/P at FZJ, the CRAY XT5 at CSC, the IBM Power 6 at Sara, and the IBM Cell at BSC – have been installed and have been evaluated using synthetic benchmarks to measure system performance, I/O bandwidth, and communication characteristics. The NEC SX-9 at HLRS is installed. The Intel components for HLRS's hybrid system and the Intel systems at CEA and FZJ will be installed in the very near future. The execution of the application benchmarks is in progress on the first systems.

Dissemination

The full impact of PRACE for European Science and industry can only be achieved if it serves informed users. Already in its first year PRACE became a well established actor internationally through numerous presentations, press releases, and participation in key events, like ISC08 in Dresden, ICT2008 in Lyon, and SC08 in Austin, TX. PRACE created an award for the best paper on petascaling by a student or young European scientist at the annual ISC conference. The first winner was Dominik Göddeke, TU Dortmund [3]. This will be continued to encourage research in HPC related topics.

Training

An equally important aspect is the knowledge how to use the systems, harness the 100thousand-fold parallelism, and exploit the opportunities of novel or heterogeneous architectures in systems containing Cell processors or GPUs. PRACE conducted a survey among Europe's top HPC users about their training needs. Over 90% stated that they need advanced training. In addition, this survey raised great interest in the US and Asia where organizations requested to use it in to identify their training needs. To teach programming for the prototype architectures, PRACE already organized a Summer School in Stockholm and a Winter School in Athens including hands-on usage of the prototypes. These two events attracted already far more applicants than could be accommodated.

Work in Progress

During the remainder of the project duration, application related work PRACE will continue to sale and optimize important codes, evaluate and enhance libraries and tools, and finalize the benchmark suite. The prototype evaluation will be completed, resulting in a compendium describing the systems and theirs characteristics. This will be used to guide the procurement of the production systems. Access to the prototypes to port and evaluate applications can be requested following a lightweight process described at www.prace-project.eu/prototype-access. Software to manage the distributed PRACE systems and to integrate them into the European HPC ecosystem is largely based on the developments done in the DEISA projects [4].




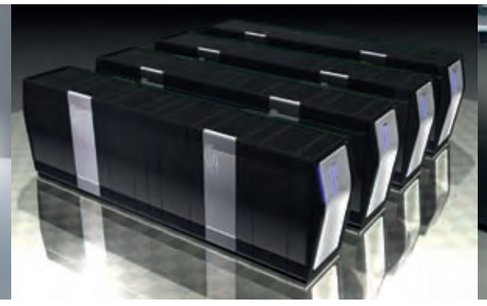


PRACE Prototypes

It will be enhanced and deployed at the partner sites and remain a lasting result of EC funded research and development.

PRACE will continue active dissemination and presence at conferences and exhibitions. Education and training will become a permanent task of the Research Infrastructure. In 2009 five additional workshops will be held at partner sites. In addition, the evaluation of technologies was started and STRATOS, the PRACE advisory group on Strategic Technologies, was created as the initial step in a continuous process to watch promising hardware and software components for future multi-petaflop/s systems and to foster collaborations with industry to develop HPC competence in Europe.

Equally important is the completion of the legal and administrative work in preparation for a pan-European legal entity. This includes: deciding the legal form; finalizing the governance structure; confirming the funding and usage models; strengthening the links within the HPC ecosystem; establishing a peer-review process; creating a procurement strategy and process.

The Future

PRACE made significant progress in its first year. Over 250 participants from all 16 partners contributed to the achievements. Nearly 100 participated in the all-hands meeting in February in Jülich to finalize and fine-tune the activities for 2009. The work will not stop when the PRACE project ends. It will be carried forward into the legal entity, enhanced where needed and complemented by further R&D projects. PRACE is meeting the numerous challenges to create a truly European HPC service that is designed to stay and to produce world-class science.

References

[1] The PRACE Project Receives Funding from the European Community's Seventh Framework Programme under Grant Agreement no RI-211528. Additional information can be found under www.prace-project.eu

[2] See inSiDE Vol. 6 No. 2, p. 4 ff.

[3] See www.mathematik.tu-dortmund.de/lsiii/static/showpdffile_TurekGoeddekeBeckerBuijssenWobker2008.pdf

[4] DEISA, Distributed European Infrastructure for Supercomputer Applications, was funded in part und grant numbers 508803, 031513, RI-222919

PRACE all-hands meeting

Gauß-Allianz Officially in Place

Two years after the Gauss Centre has been created to coordinate the European supercomputing activities, Germany took the next step to shape the future of high performance computing (HPC) on a national level. On December 3rd, 2008 representatives of the nations leading supercomputing centers met in Bonn to create the Gauß-Allianz. While the Gauss Centre was successfully addressing the needs on the top end, it was clear from the beginning that an additional layer of support was required to maintain the longevity with a network of competence centers across whole Germany. This gap is now addressed by the Gauß-Allianz, where regional and topical centers team up to create the necessary infrastructure.

The founding members are Gauss Centre for Supercomputing (GCS), Center for Computing and Communication of RWTH Aachen University, Norddeutscher Verbund für Hoch- und Höchstleistungsrechnen (HLRN) consisting of Konrad-Zuse-Centre for Information Technology Berlin (ZIB) and Regional Computer Centre for Lower Saxony (RRZN), Center for Information Services and High Performance Computing (ZIH) at TU Dresden, Rechenzentrum Garching (RZG) of the Max-Planck Society, Deutscher Wetterdienst (DWD), German Climate Computing Centre (DKRZ), and Steinbuch Centre for Computing (SCC) as full members and Goethe Centre for Scientific Computing (G-CSC) at Frankfurt University, Forschungszentrum Computational Engineering, Kompetenzgruppe Wissenschaftliches Hochleistungsrechnen der Technischen Universität Darmstadt, Paderborn Center for Parallel Computing (PC2), Regionales Rechenzentrum Erlangen (RRZE), and the German Research and Education Network (DFN) as associated members.

The mission is to coordinate the HPC related activities of the members. With the provision of versatile computing architectures and by combining the expertise of the participating centres, this creates the ecosystem necessary for computational science. Strengthening the research and increasing the visibility to compete on an international level are further goals of the Gauß-Allianz. It is especially an important milestone to maintain the leading role of Germany in the European supercomputing activities. With the foundation of the Gauß-Allianz, the HPC community creates the organizational framework to continue the successful collaboration of the past and implements the recommendations given in the BMBF study "HPC in Germany – Reasons for a Strategic Alliance". On the other hand BMBF recognizes the important role of HPC in science and industry, where sufficient computing power and advanced computational methods are inevitable to drive research and development. This demand also motivates the research in HPC that is necessary to continue the exponential growth in compute power while maintaining and improving usability, productivity and efficiency.

GA Gauß Allianz

Gauss Centre Looks Back on a Successful Year 2008

In its annual report the Gauss Centre for Supercomputing (GCS) can look back on a very successful year 2008. This holds for many aspects, especially for

- a successful research project application guaranteeing the funding of Petascale Computer Systems for the GCS member institutions
- a successful series of scientific workshops and HPC related training courses
- an active and leadership involvement in several PRACE work packages
- the intensive participation in preparing the foundation of the Gauß-Allianz as a German association (Gemeinnütziger e.V.).

In order to be successful in the competition for becoming an European TierO-site within the planned European HPC Ecosystem GCS has to qualify both with respect to research in the area of Petascale Computing as well as to guaranteeing the development of an adequate Petascale Computing infrastructure. GCS succeeded in progressing a lot in 2008 concerning both aspects. Based on several scientific use cases and good scientific results we were successful in convincing politicians and ministries to support GCS's position as a Principal Partner on behalf of Germany within the PRACE project. This means that the German National Ministry for Research and Education (BMBF) promized to provide up to 200 Million € for GCS for the purpose of Petascale Computing under the additional conditions that

- GCS will develop a national Petacomputing concept that is jointly agreed upon by the three national supercomputer centres Garching (LRZ), Jülich (JSC) and Stuttgart (HLRS) and meets the plans and expectations of PRACE
- the German states (Bundesländer) Baden-Württemberg, Bayern and Nordrhein-Westfalen are willing to supply the same amount of money as the corresponding national supercomputing center within its domain gets proportionately from GCS as part of the BMBF money

In October 2008 a corresponding contract (Verwaltungsabkommen) between the four responsible ministries was signed and at the same time GCS applied for the project PetaGCS. The long title of the application is "Beschaffung und Betrieb von Supercomputern für das Gauss Centre for Supercomputing (GCS) als Beitrag zum nationalen Versorgungskonzept für TierO/1 im Rahmen eines europäischen HPC-Ökosystems" (Procuring and Operating Petacomputing Systems for GCS as a German national contribution to a European HPC ecosystem). The application contains GCS's plans concerning mutually agreed system architectures, research focus, support structures, peer review procedures, access policies, and governance structures. Cooperation contracts between GCS as a formal association and its member centers will ensure that the BMBF funding will be deployed according to GCS's plans, and a steering committee jointly nominated by the four responsible ministries will monitor whether GCS is meeting the expectations of the research project and the funding ministries.

In December 2008 GCS obtained the approval for a first project phase: the sum 128 Million € was granted by the German Federal Ministry of Education and Research. Additional funding of the same amount will be contributed by the local Federal States (Bundesländer). According to the GCS plans, JSC (Jülich) was appointed to prepare the procurement of the first GCS Petacomputing systems to be installed still in 2009, Stuttgart and Garching will follow in later years. Details of the first system will be published in separate articles in GCS's publication "InSiDE". By founding GCS, by providing an appropriate funding by politics and ministries, and by unbureaucratically searching for adequate administrative and organizational solutions Germany succeeded in making a big step forward with respect to its competitiveness within the European HPC landscape and achieving the top position within PRACE principal partners. I thank all persons, ministries and institutions that have been involved in the GCS success story so far.

Also in 2008, with strong cooperation of GCS, Gauß-Allianz was founded, a German association comprising especially HPC centers of Tier 2 in the well-known HPC supply pyramid. This step complements the endeavors of GCS to take care for a broader awareness of the scientific and economic chances of HPC and for a lasting support of HPC systems and infrastructures including the development of HPC relevant software systems.

So indeed, GCS can look back on a successful year 2008. Of course, a lot of still open activities remain for 2009 and the following years. The preparatory phase of the PRACE project gradually comes to an end setting the course for important decisions concerning the European HPC ecosystem such as form and location of legal entity, governance structure, access and review policy, financing structure, etc.. Also, the project PetaGCS is still in its infancy, but we are confident that we can act in accordance with our plans and that politics will provide lasting support because of the great importance of simulation and of scientific as well as high performance computing for our national economy.

- Heinz-Gerd Hegering

Chairman of GCS

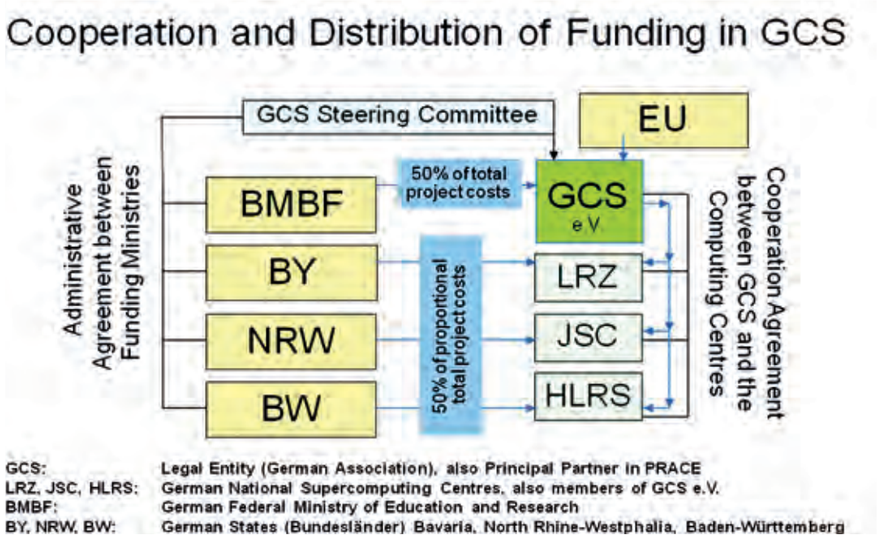

Selecting an Appropriate Computational Platform for Supporting the Development of New Catalyst Carriers

In the chemical process industries the majority of heterogeneously catalyzed gas-phase reactions are carried out in tubular fixed-bed reactors that contain a so-called "random packing". This term refers to an irregular arrangement of individual (i.e. not interconnected) catalyst particles, which in turn can be assigned to the class of unconsolidated structures. However, owing to this inhomogeneous and incoherent geometrical structure, such conventional fixed-bed reactors exhibit several drawbacks [1] that result from the non-uniform distribution of the velocity field and the concentration field, respectively. The choice of optimal values for the tube diameter and the particle size is a complex trade-off where a variety of – partly incompatible – requirements such as a sufficient heat transfer rate, an acceptable pressure drop over the packing, adequate mechanical stability of the particles and negligible (internal) mass transport limitations due to pore diffusion have to be considered. A promising alternative is the use of consolidated structures (e.g. ceramic foams as shown in Figure 2 or monoliths) as catalyst support since these structures may eliminate the aforementioned drawbacks, featuring potentially enhanced mass and heat transfer characteristics and a lower pressure drop compared to conventional fixed-bed configurations [2]. These features are particularly favorable in situations which involve high flow rates and/or strongly exothermic or endothermic reactions. Despite these clear advantages, the application of consolidated catalytic supports in the chemical process industries is still limited to very few examples. The main reasons for this are the higher manufacturing costs and the necessity for – and sometimes lack of experience with – different handling methods for the loading, unloading and sealing of the reactor units compared to the conventional unconsolidated counterparts [3].

Figure 1: Sustained aggregated performance on different systems as function of the number of MPI processes for a fixed domain size of only 100^3 (or 259^3 in case of NEC SX-8/9 as noted in the legend). Thin dash-dotted lines indicate ideal scaling.

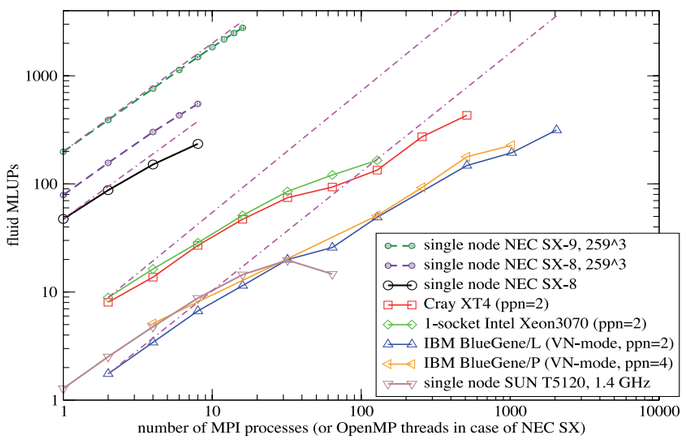

In order to profit from the superior mass and heat transfer characteristics, it is first of all essential in the manufacturing of the consolidated structures to be able to design a desired geometry in a reproducible manner. In addition, it is desirable to have the possibility to perform a "geometry tuning" in order to control the textural properties of the support and thus adjust specific transport characteristics. However, as of today there is no quantitative understanding of the structural influence of e.g. local foam structures (consolidated support) on the fluid dynamics and the heat and mass transport (and thus on the reactor performance), hence further research in this direction is needed [4] which is the motivation for our work.
For randomly packed fixed-bed reactors, we have already shown that the integral reactor performance can be influenced significantly by local transport phenomena [5,6,8,9]. These local processes can hardly be measured and cannot be described by conventional reactor models that rely on homogenization approaches. The same is expected to hold true for foams, because although they belong to the class of consolidated structures, foams and random packings have a complex geometry in common. Therefore, it is not sufficient to calculate (or correlate) the transport characteristics of a foam structure using integral values such as the global porosity. It is in fact necessary to perform spatially resolving numerical simulations that use the "real" 3-D geometrical structure as a basis.

Methods

In our work, we combine experimental and numerical methods in order to obtain a better insight into the local transport phenomena, allowing for a more fundamental understanding of the relationship between support structure and reactor performance.

As an example of a consolidated structure, we used a cellular SiC foam with a global porosity of 65% and compared the foam with a randomly packed bed of 2 mm spherical quartz beads. The measurements were carried out for different flow rates in a flanged stainless-steel tube at room temperature. The samples were placed in a cylindrical case (3 cm diameter, 8 cm length) with a steel grid at the bottom, which was fixed between in the middle of the tube by a stainless-steel packing. A carbon mat was used to seal the samples to the walls.

Unlike integral quantities such as the global porosity and the integral pressure drop, the details of the geometry and of the flow field cannot easily be investigated using conventional experimental methods. However, as pointed out before, the local transport phenomena are essential for the integral reactor performance. Therefore we revert to numerical simulations of the flow field in the 3-D geometry. The details of the geometrical structure are either generated numerically or supplied by an X-ray scan of the foam structure. For the latter case this allows us to perform numerical flow simulations and experimental investigations in the identical geometry. Numerically, the geometrical structure of random packings of spheres can be generated efficiently by Monte Carlo methods [5,6,8,9]. As a first approximation, foams can be seen as "inverse" packings and thus be generated by similar means. For the 3-D flow simulations on the pore-scale level we apply an advanced lattice Boltzmann flow solver [9]. As resolving the 3-D geometry results in huge computational domains, high performance computing and the selection of appropriate compute resources is essential.

Figure 2 left: SiC foam samples with different pore density. Manufacturer: Fraunhofer IKTS, Dresden/Germany

Flow Results and Discussion

The measured pressure drop for the spherical bead packing was in agreement with literature data [7]. Interest-

Figure 2 right: Synthetically generated sphere packing with a tube-to-particle diameter ratio of 5

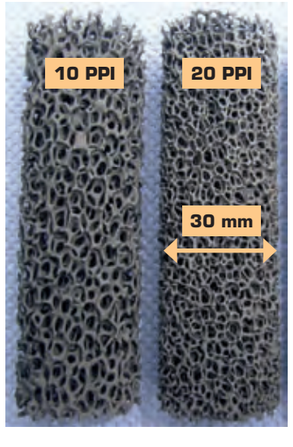

ingly, the pressure drop of the foam is – despite its higher global porosity – up to about four times higher than that of the packing [12]. This clearly illustrates the need for a local investigation of the flow field. The results of the numerical flow simulation show that the high pressure drop can be attributed to the local cellular foam structure (Fig. 3 and 4). Because the open windows between the cells are just a few µm in size, a very inhomogeneous flow distribution with more than 20 times the superficial velocity can be observed at these positions. These first very promising results will be complemented in the future by a more detailed comparative study of different geometrical structures. A comparison of experiment and simulation for identical geometries will also be performed. A long term effort is the simulation of the local mass transport inside the complex geometries, which will allow for the investigation of e.g. the residence time distribution and mixing properties. As final step on a long term perspective, the simulation of coupled chemical reactions inside the 3-D geometrical structure is desirable.

Figure 3: Comparison of 2-D slices of the simulated 3-D flow field in the foam and a periodic sphere packing at different Reynolds numbers

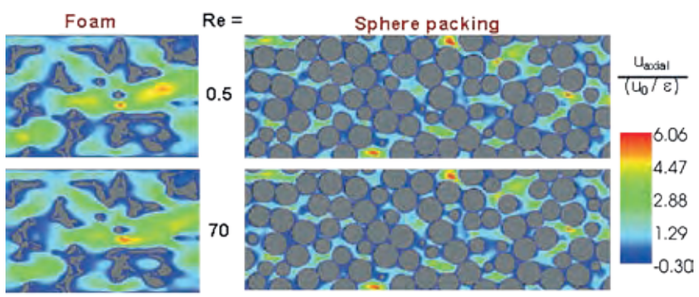

Computational Aspects

Recent Top500 lists are dominated by commodity clusters and massively parallel IBM BlueGene systems. But vector systems are still alive, in particular the NEC SX series [10,11]. Using GPUs for general purpose computing is another emerging trend, however, not considered here. The lattice Boltzmann method is a recent method from computational fluid dynamics which has its roots in a highly simplified gas-kinetic description, i.e. a velocity-discrete Boltzmann equation with appropriate collision term. For the present investigations, we use the ILBDC code which is based on a sparse list-based data structure [10] holding only the data of fluid cells and required adjacency information of the underlying 3-D Cartesian lattice with 19 discrete velocities per cell (D3Q19 model). This results in a data transfer of 528 bytes per cell update if cache line read-for-ownership (RFO) is required, or 376 bytes otherwise. The applied two-relaxation-time (TRT) collision operator requires about 200 Flops per cell update. Performance numbers are generally given in million lattice site updates per second (MLUPs).

Figure 1 shows the scalability of the code on different parallel computers for strong-scaling experiment with a challenging fixed total domain size of only 100^3 (or 259^3) as function of the number of MPI processes (or OpenMP threads in case of the NEC SX systems). On all clustered systems, the reported numbers are for the case where all cores of the nodes have been utilized – although using only some cores per node usually gives better performance – as users usually have to pay on a per-node basis. Generally, three different groups of systems can be distinguished: "green systems" with low power CPUs, clusters with commodity CPUs and "high end systems" using vector CPUs. Up to 32 MPI processes, the single one-socket Sun T5120 system with Sun's new multi-core multi-threaded Niagara2 processor [13] can keep up with 16 BlueGene/L or 8 BlueGene/P nodes – which probably are much more expensive and consume in total much more power than this single high-end workstation server. The performance of the single-socket Intel Xeon system with Infiniband interconnect is quite comparable with the Cray XT4. As the dashed lines show, there is a significant performance loss with increasing number of MPI processes although the lines look quite straight in the double-logarithmic plot. The NEC SX-8 is a class of its own and the NEC SX-9 even quadruples the performance. Scalability and sustained performance are not directly connected to each other. It is well known that slow code scales much better than a code with highly optimized single processor performance owing to the ratio of computation to communication time. Similar rules apply when comparing parallel computers with slow and powerful processors/single nodes. The results presented above clearly demonstrate for the ILBDC solver that the NEC SX systems are a class of their own. However, clustered systems which state-of-the-art CPUs scale quite well if a good interconnect (DDR Infiniband or better) is used. For a fixed total amount of work, the per-core efficiency gradually drops as the number of cores is increased, although the sustained aggregated performance still increases. Other applications, e.g. the molecular dynamics code Amber, will not scale at all beyond just a few dozens of MPI processes. The single processor or single node performance of the IBM BlueGene systems is (as expected) very low. Using huge numbers of BlueGene nodes, the performance level of faster systems can be reached in the case of ILBDC, however, domain decomposition and load balancing becomes more and more complicated as the number of partitions increases. For other applications (e.g. Amber) a BlueGene sys-

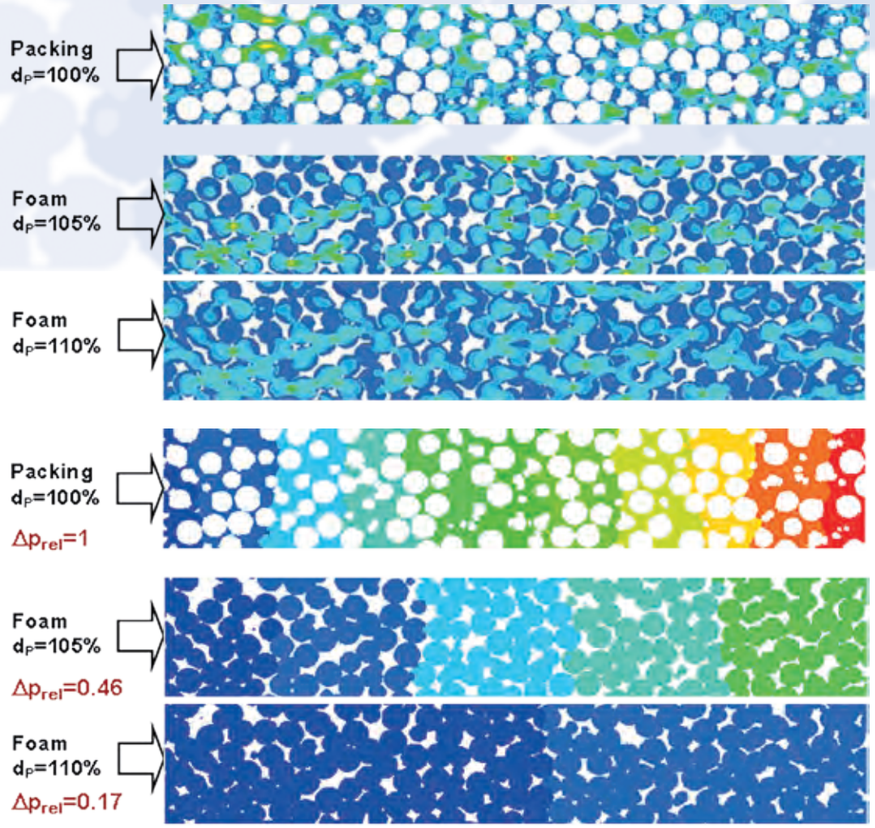

Figure 4a: Comparison of the velocity field of the sphere packing and foams of slightly different window openings (and porosity). The same colour scale is used for all three plots.

Figure 4b: Comparison of the pressure drop of the sphere packing and foams of slightly different window openings (and porosity). The same colour scale is used for all three plots.

tem will never reach the performance level of even simple commodity clusters owing to scalability constraints. Except on the NEC SX systems, the experience with a hybrid approach (i.e. OpenMP within a node and MPI between nodes) was rather disappointing [10]. Pure MPI with the explicit separation of memory accesses not only gave advantages in the case of ccNUMA nodes. A lot of research will be required to see whether pure MPI, hybrid approaches or totally new concepts (e.g. PGAS languages) are the way to go for massively parallel computers consisting of SMP nodes with multi-socket many-core CPUs.

- Thomas Zeiser[1]
- Georg Hager[1]
- Gerhard Wellein[1]
- Amer Inayat[2]
- Wilhelm Schwieger[2]
- Tobias Heidig[3]
- Hannsjörg Freund[3]

[1] Regionales Rechenzentrum Erlangen, University of Erlangen-Nuremberg

[2] Lehrstuhl für Chemische Reaktionstechnik, University of Erlangen-Nuremberg

[3] Max-Planck-Institut für Dynamik komplexer technischer Systeme, Magdeburg

Acknowledgments

The scaling benchmarks presented in this report were measured on the NEC SX systems at the High Performance Computing Centre Stuttgart (HLRS) within the project lba-diff, on the IBM BlueGene L/P systems at the Jülich Supercomputing Centre (JSC), the IBM BlueGene P system of the Rechenzentrum Garching (RZG) of the Max Planck Society, the Cray XT4 system of the National Energy Research Scientific Computing Center (NERSC) at the Lawrence Berkeley National Laboratory, the Sun Niagara2 cluster at RWTH Aachen and RRZE's own resources. This work was carried out within the Bavarian framework of KONWIHR-II with additional financial support from BMBF through the project SKALB (grant 01IH08003A). The 1-D list based code is the result of joined work of several partners within the International Lattice Boltzmann Development Consortium (ILBDC).

References

[1] Cybulski, A., Moulijn, J. A.
Catal. Rev., 36(2), 179–270, 1994

[2] Kreutzer, M. T., Kapteijn, F., Moulijn, J. A.
Catal. Today, 111(1-2), 111-118, 2006

[3] Boger, T., Heibel, A. K., Sorensen, C. M.
Ind. Eng. Chem. Res., 43(16), 4602-4611, 2004

[4] Reitzmann, A., Patcas, F. C. Kraushaar-Czarnetzki, B.
Chem.-Ing.-Tech., 78(7), 885-898, 2006

[5] Freund, H., Zeiser, T., Huber, F., Klemm, E., Brenner, G., Durst, F., Emig, G.
Chem. Eng. Sci., 58(3-6), 903-910, 2003

[6] Freund, H., Bauer, J., Zeiser, T., Emig, G.
Ind. Eng. Chem. Res., 44(16), 6423-6434, 2005

[7] Zampieri, A.
Development of MFI-type Zeolite Coatings on SiSiC Ceramic Monoliths for Catalytic Applications, PhD Thesis, University of Erlangen-Nuremberg, 2007

[8] Freund, H.
Ortsaufgelöste Simulation von Transportprozessen in durchströmten Festbetten, Shaker Verlag, Aachen, 2008

[9] Zeiser, T.
Simulation und Analyse von durchströmten Kugelschüttungen in engen Rohren unter Verwendung von Hochleistungsrechnern, Verlag Dr. Hut, München, 2008

[10] Zeiser, T., Hager, G., Wellein, G.
Vector Computers in a World of Commodity Clusters, Massively Parallel Systems and Many-core Many-threaded CPUs: Recent experience based on an Advanced Lattice Boltzmann Flow Solver, in: Nagel, W. E., Kröner, D.B., Resch, M. M. (Editors), High Performance Computing in Science and Engineering '08, Springer, Berlin, 2009

[11] Zeiser, T., Hager, G., Wellein, G.
The World's Fastest Single-core Processor: Some Performance Results from the NEC SX-9, LSPP Workshop at IPDPS Conference, Rome, May 2009

[12] Freund, H., Inayat, A., Bauer, J., Zeiser, T., Schwieger, W.
3D Simulation of the Local Flow Field in Ceramic Foam Structures, AIChE 2008 Annual Meeting, Philadelphia, PA, USA, November 2008

[13] Hager, G., Zeiser, T., Wellein, G.
Data Access Characteristics and Optimizations for Sun UltraSPARC T2 and T2+ Systems, Parallel Processing Letters, 18(4), 471-490, 2008

Hybrid Techniques to Meet the Challenge of Simulating Complex Turbulent Flows

Flows in engineering applications are predominantly turbulent and complex. The complexity is generated by complex geometries and complex physical mechanisms at work manifesting themselves in mean flow, spectra, and their energy transfer. Such mechanisms include but are not limited to flow separation and the formation of large coherent structures as occur in the examples of the asymmetric air intake in Figure 1 and the hill flow in Figure 4. For "simple" turbulent flows, such as fully attached flows, jets and wakes, the "complicated physics" are excluded or can be modeled adequately for many engineering purposes. As a consequence statistical turbulence models employed to solve the Reynolds-Averaged Navier-Stokes (RANS) equations are fairly successful in predicting the mean flow behavior of these kinds of flows at reasonable computational costs. However, for complex flows, the semi-deterministic Large Eddy Simulation (LES), which does not rely on modeling the entirety of the turbulent fluctuations, has clear superiority over RANS methods. Indeed, with the availability of ever increasing computing power, researchers successfully applied LES in recent years to tackle a range of non-trivial problems [1]. Unfortunately, high resolution requirements, particularly in wall-bounded flows, and the need to determine time averages explicitly by executing a large number of time steps are making LES still too expensive for routine application in engineering. The dilemma faced today by many practitioners of Computational Fluid Dynamics (CFD) hence is that, for a given complex turbulent flow, the economical RANS approach may be unsatisfactory whereas LES would simply be unaffordable. This

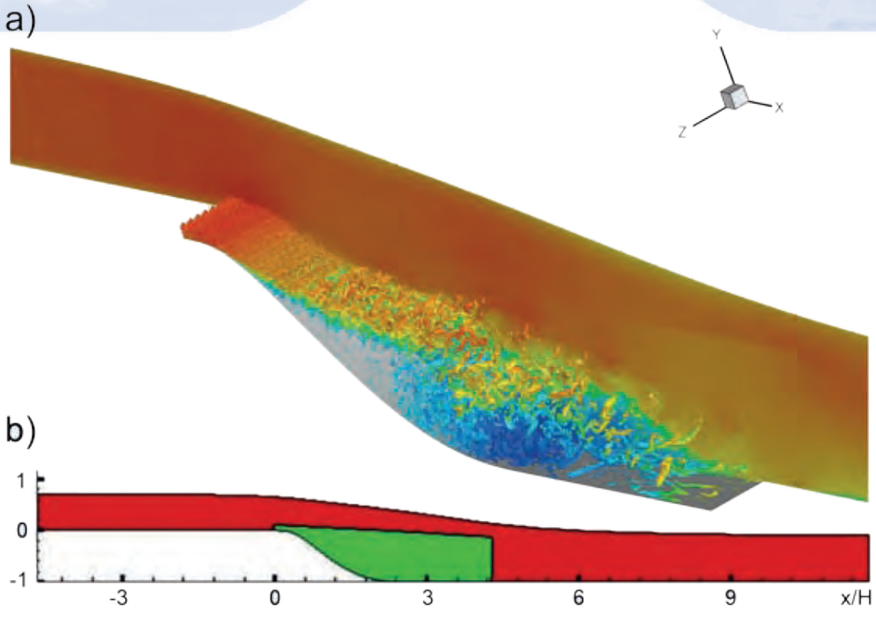

Figure 1: Application of segregated LES/RANS coupling for the simulation of an air intake: (a) Instantaneous contours of vortex identification criterion Q; (b) setup of embedded LES: 2D RANS zone (red) and 3D LES zone (green)

is the motivation to construct so-called hybrid LES/RANS methods that combine both methodologies in a single simulation. The underlying philosophy is "to perform LES only where it is needed while using RANS in regions where it is reliable and efficient" [2].

Segregated LES/RANS Modeling

Several distinct approaches to devise LES/RANS hybrids exist and a host of individual methods has been proposed having met with varying degrees of success and popularity. In [2], the fundamental principles of various hybrid methods are formulated and categorized, supplemented with characteristic examples and with assessments. Here, we will illustrate the promising class of "segregated LES-RANS modeling" pursued in our own work. More details can be found in [3]. Segregated modeling is based on decomposing the entire domain into clearly identifiable regions for RANS and LES before starting the simulation. The connection between the distinct zones during the simulation is then established via explicit coupling of the solution at the interfaces. The aim of the approach is to compute all models in their regime of validity: steady RANS for flows with stationary statistics and unsteady LES with high resolution where it is needed. Therefore one can choose the best-suited turbulence modeling for each sub-domain without considering compatibility issues and without fear of inconsistencies in their use. In particular, the occurrence of a so-called gray zone where the resolved motion has to accomplish some quasi-physical transition between LES and RANS character is an often encountered disadvantage with non-segregated approaches. The price to pay is the need for comparatively complex coupling conditions.

For block-structured solvers employing the concept of ghost (or halo) cells, routines for data exchange at block boundaries are required anyway. The related routines lend themselves to a straightforward implementation of the LES/RANS coupling when the interface is positioned at block boundaries. This is illustrated in Figure 2 for a Finite Volume method as was used in the present project and the streamwise velocity as the flow quantity to be coupled. Note that across the interface only statistics can be continuous. Consequently, the mean velocity is exchanged at the LES-RANS interface. In the RANS domain the mean is obtained as result of solving the governing equation, whereas an explicit averaging operation has to be carried out in the LES domain. In addition, the LES domain requires realistic unsteady data in its ghost cells as boundary condition. This data has to be consistent with the physical processes at the interface and must not violate the (sometimes intricate) turbulence modeling assumptions on both sides of the boundary. Generating such data at the interface, i.e. the fluctuations to be added to the RANS prediction of the mean, is a key issue in segregated LES/RANS modeling.

Inappropriate coupling leads to wanting results in the LES and/or RANS subdomains. This is demonstrated in Figure 3 for turbulent channel flow. Straight-forward direct exchange of the computed variables at the interface yields an unsteady, in the vicinity of the interface degraded solution in the RANS domain. On the other hand, matching only mean values and adding unphysical or even no fluctuations to the LES ghost cells causes artificial reflections contaminating the LES results. Finally, with a physically consistent coupling a flow field undistinguishable from a full LES can be attained – even directly at the interface. How to construct viable fluctuations depends strongly on the type of interface but also on the sophistication level of the segregated model. Suitable techniques for the different cases are presented in [3] and the references therein.

Types of LES/RANS Interfaces

In general, one can distinguish three types of interfaces by the direction of the mean flow with respect to the LES/RANS boundary. In Figure 1 all three types are present: The (green) LES zone has an upstream and downstream boundary, i.e. an "inflow-type" and "outflow-type" interface, respectively, and an interface with more or less tangential alignment of the LES and RANS boundary with the direction of the mean flow. Of course, in other applications not all of these situations need to appear. The type of interface has strong repercussions on what information can be used to construct the fluctuations. The arguably easiest case is the outflow-type, since time-resolved information from the LES side of this boundary can be used to "transport" fluctuations into the LES ghost cell. Using a one-dimensional convection equation yielded the very

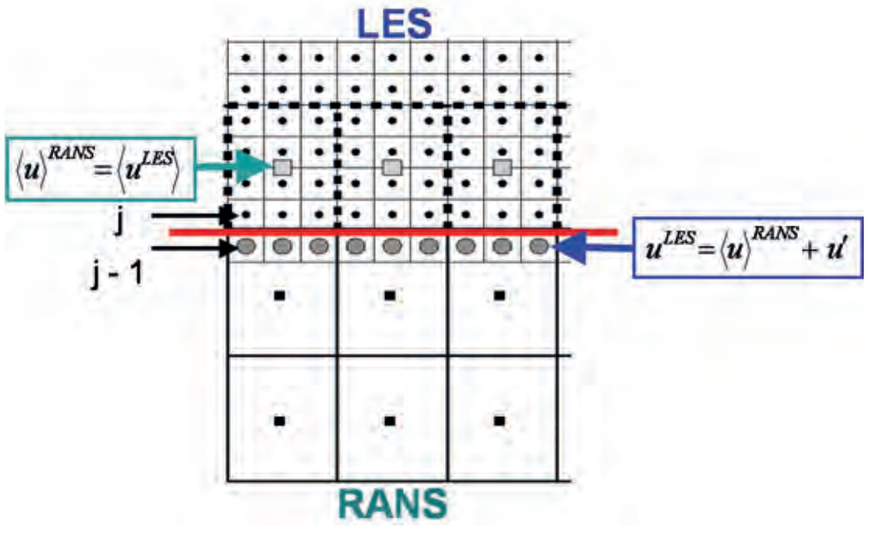

Figure 2: Implementation of coupling conditions for the streamwise velocity in a Finite Volume solver using ghost cells at the block boundary: RANS on coarser grid provides mean flow data for LES ghost cells on a finer grid, fluctuations if needed are generated according to the type of interface; explicitly averaged LES data are assigned to the RANS ghost cells; the red line indicates the block boundary.

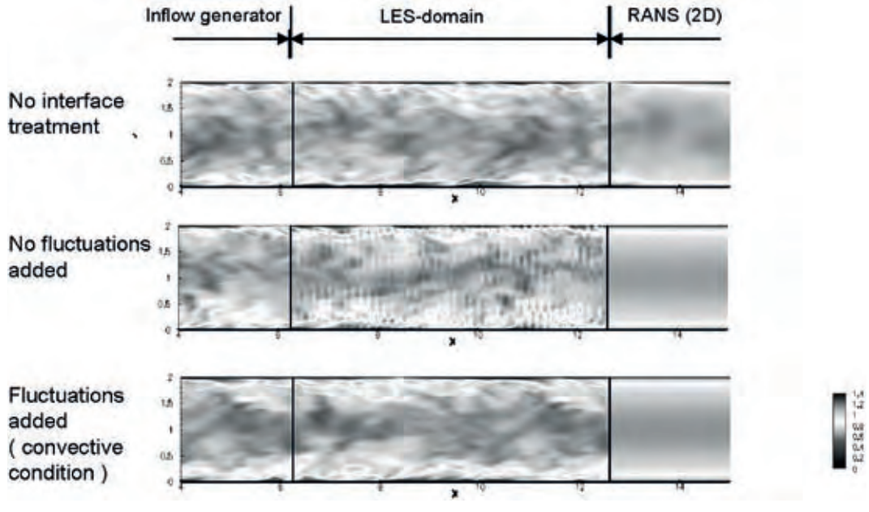

Figure 3: Illustration of the effect of the LES/RANS interface for turbulent channel flow using streamwise velocity contours; top: no interface treatment results in a degraded unsteady RANS zone, center: interface without (or wrong) fluctuations yields contamination of LES zone with artificial reflections, bottom: physically consistent coupling

good results shown in Figure 3. However, for the incompressible hill flow with a challenging positioning of the interface depicted in Figure 4, it was also necessary to decouple the pressure fields at the boundary and then explicitly ensure mass conservation at the LES/RANS interface (see [4] and references therein for details).

At the inflow-type interface, the same situation prevails as for traditional LES inflow boundaries. Therefore, adjusted techniques from the literature can be employed [5]. However, only the part of such methods dealing with the fluctuations can be taken into account since the mean flow profiles are provided by the RANS calculation. For the air intake in Figure 1, a database of fluctuations from canonical turbulent boundary layers in conjunction with a rescaling technique to match Reynolds and Mach numbers at the interface location as predicted by the RANS calculation worked well (see [3] and references therein). For the tangential interface, if the boundary is located in uncritical regions, as with the air intake in Figure 1, a simple copying of fluctuations from the LES domain into the corresponding ghost cells often suffices. However, for more challenging locations, e.g. if the RANS zone is placed between the LES and a wall in close proximity, the tangential boundary can be difficult to handle and is still subject to ongoing research. The approach followed in the current project is the use of statistical information provided by the RANS part of the simulation to rescale fluctuations copied from the LES side of the interface [3,4]. Preliminary results are promising but need further scrutiny.

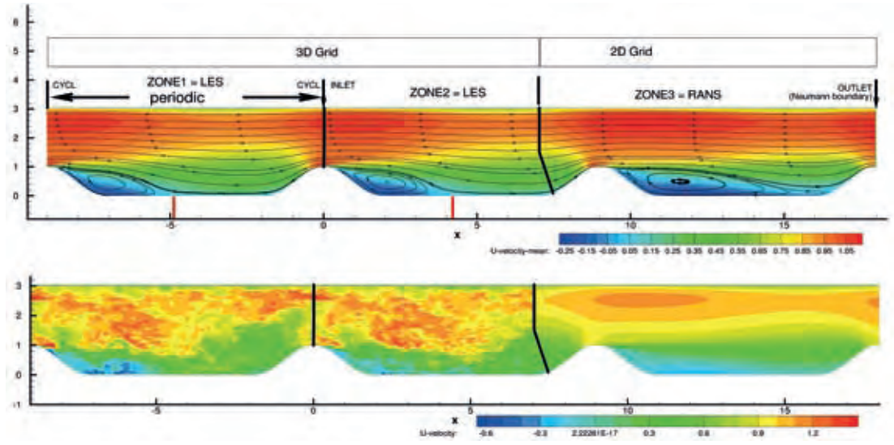

Figure 4: Segregated LES/RANS coupling for the flow over periodic hills with the LES-to-RANS interface at x ≈ 7; mean (top) and instantaneous (bottom) streamwise velocity

Flow Solver and High-performance Computing

The simulations presented here, except for the air intake in Figure 1, were performed with the Finite Volume code LESOCC2 developed at the University of Karlsruhe and described in [6]. This FORTRAN 95 program solves the incompressible, three-dimensional, time-dependent, filtered and/or Reynolds-Averaged Navier-Stokes equations on body-fitted, collocated, curvilinear, block-structured grids. The code has been used in numerous studies and substantial experience has been gained with respect to its numerical properties and its performance on various hardware platforms including the NEC-SX8, SGI-Altix, and several Linux clusters. Parallelization is accomplished via domain decomposition and MPI for the data transfer. To achieve better performance on a given hardware, different implementations of the computationally most intensive algorithms exist. The user can decide between highly vectorized versions and unvectorized versions of subroutines or can choose to trade off between memory requirements and number of operations in order to fit a run on a specific machine or possibly take advantage of cache effects. Scaling of LESOCC2 on the HP-XC4000 (the Linux cluster used for the present investigation) was tested for up to 512 processor cores and satisfactory results were obtained [4]. To develop the hybrid techniques presented here, several parameter studies were necessary and, in addition, a few larger simulations have been computed to obtain reference data. For the parameter studies, typically 16 to 32 processors were used and the efficiency of the parallelization was over 90%. A single run required between 2,000 and 3,000 CPU hours and about 100 cases had to be computed. Reference simulations of additional test cases for the ongoing project (not shown here) are employing between 112 and 256 processor cores.

Conclusions

The coupling of Large Eddy Simulation (LES) with statistical turbulence models (RANS models) has arguably become the main strategy to drastically reduce computational cost for making LES affordable in a wide range of complex industrial applications. One such method is segregated LES/RANS modeling as illustrated above. With a segregated approach, turbulence models are operated under conditions they were intended for; on the other hand, the user is required to define a priori where LES and where RANS are to be performed. The development of smart interfaces is challenging and not completed yet. Even though the LES/RANS coupling is intended to save tremendous amounts of computational costs, the LES part of the hybrid simulation still requires high LES-type resolution. As a consequence, even with efficient hybrid techniques, high-performance computing will still be necessary but more and more complex turbulent flows should become accessible to simulation.

Acknowledgements

The present project is part of a French-German cooperation funded by Deutsche Forschungsgemeinschaft (DFG) and ONERA, Châtillon, in the framework of the DFG/CNRS research group FOR 507 “LES of Complex Flows”. The provision of computing time by SCC Karlsruhe is gratefully acknowledged.

References

[1] **Fröhlich, J.**
Large Eddy Simulation turbulenter Strömungen, Teubner, 2006

[2] **Fröhlich, J., von Terzi, D. A.**
Hybrid LES/RANS Methods for the Simulation of Turbulent Flows. Progress in Aerospace Sciences, volume 44, pp. 349–377, 2008

[3] **von Terzi, D. A., Mary, I., Fröhlich, J.**
Segregated LES/RANS Coupling Conditions for the Simulation of Complex Turbulent Flows in Brun, C., Juvé, D., Manhart, M., Munz, C.-D. (editors), Numerical Simulation of Turbulent Flows and Noise Generation, volume 104 of Notes on Numerical Fluid Mechanics and Multidisciplinary Design, pp. 231-252, Springer, 2009

[4] **von Terzi, D. A., Fröhlich, J., Rodi, W.**
Hybrid Techniques for Large-Eddy Simulations of Complex Turbulent Flows in Nagel, W. E., Kröner, D. B., Resch, M. M. (editors), High Performance Computing in Science and Engineering `08, pp. 317–332, Springer, 2009

[5] **Nolin, G., Mary, I., Ta, P.-L.**
LES Inflow Conditions for Turbulent Boundary Layers of Complex Flows Using Database, AIAA paper 2006-2888, 2006

[6] **Hinterberger, C.**
Dreidimensionale und tiefengemittelte Large-Eddy-Simulation von Flachwasserströmungen. PhD Thesis, University of Karlsruhe (TH), 2004

• Dominic von Terzi[1]
• Ivan Mary[2]
• Jochen Fröhlich[3]

[1] Institut für Thermische Strömungsmaschinen, University of Karlsruhe (TH)

[2] ONERA, Châtillon, France

[3] Institut für Strömungsmechanik, TU Dresden

Simulation of a Complete Viral Shell

Applications

Viral shells are protein assemblies that protect the genetic material inside. The icosahedral structure of many animal and plant viruses allows internal pressures of more than 60 atm. These viral capsids behave extremely stable and highly elastic upon external forces and rupture of the shells occurs during indentations with an atomic force microscopy (AFM) tip of more than 30% in capsid height. Only a fundamental understanding of the capsid's mechanical properties and their distribution on the viral shells can finally answer the questions how viral assembly and infection proceeds, and what are the driving forces of the observed structural changes during maturation and capsid breakage.

We studied the mechanical properties of the fully solvated shells of Southern Bean Mosaic Virus (SBMV) and Human Rhinovirus (HRV) 16 on atomic length scale by extended force-probe (FP) molecular dynamics simulations. Both shells consist of 60 identical triangular subunits comprising the complete capsid. Structural units are built up from 5 subunits surrounding the 5-fold symmetry axes (pentamers) and 6 subunits around the 3-fold symmetry

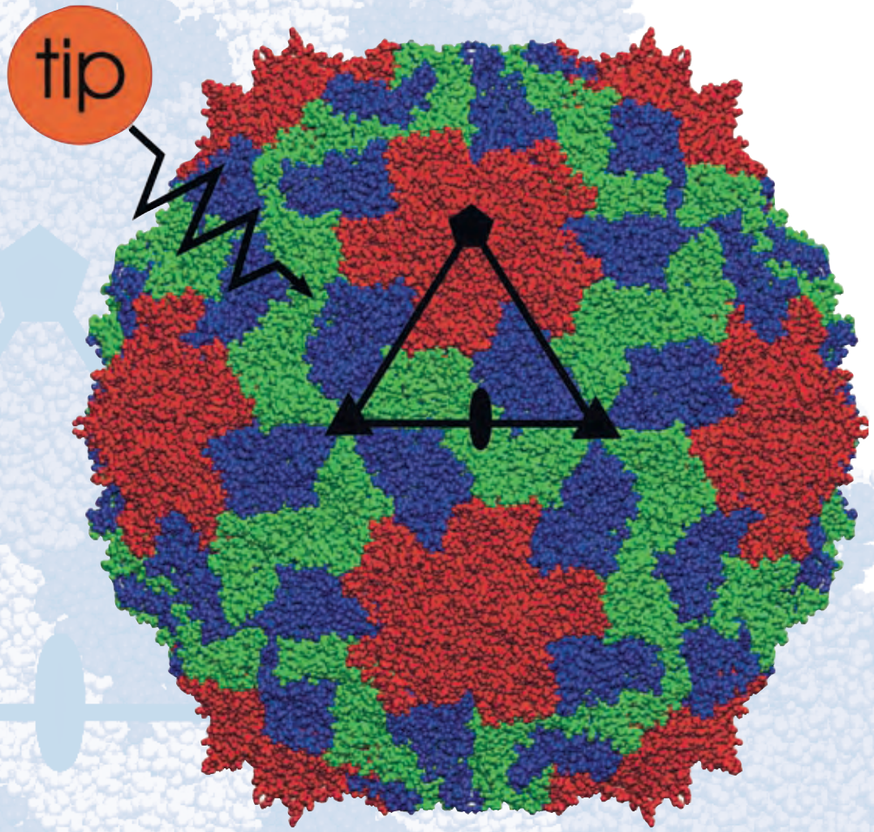

Figure 1: Shell of SBMV consisting of protein A (red), B (blue), C (green). The tip-sphere is shown as orange sphere. One subunit is marked (triangle), the black symbols denote the 5-, 3- and 2-fold symmetry axes.

axes (hexamers). All simulations were carried out on 64 processors of the Altix-SGI machines, using the software package GROMACS-4.0. To facilitate direct comparison with atomic force microscopy measurements, a Lennard-Jones sphere, which served as a model of the AFM tip, was attached to a "virtual" spring and pushed with different velocities towards and through the capsid protein at 19 different positions on the triangular subunit of the viral surface (Fig. 1). This simulation technique offers the unique opportunity to probe the mechanical properties of the capsids' internal surfaces. Thus, the tip-sphere was additionally placed inside the shells and pushed towards the inner surfaces of SBMV and HRV 16.

The simulation system of HRV 16 comprised over 4,200,000 atoms, the system of SBMV more than 4,500,000 atoms and are therefore two of the largest biomolecular systems simulated so far. In total, our simulations sum up to 0.5 µs in length which can only be handled with up-to-date supercomputer power.

Our simulations showed that the capsids of SBMV and HRV 16 behaved highly elastic upon indentation with the tip-sphere. Only the amino acids close to the indentation position deformed plastically, whereas the complete shell did not deform at all (Fig. 2).

A detailed picture of the spatial distribution of elastic constants that describe the stiffness of the material, and yielding forces needed to cause rupture and determine the shell's stability, was obtained for the surfaces of SBMV and HRV 16 (Fig. 3 and 4). An inhomogeneous distribution of both, elastic constants and yielding forces was found.

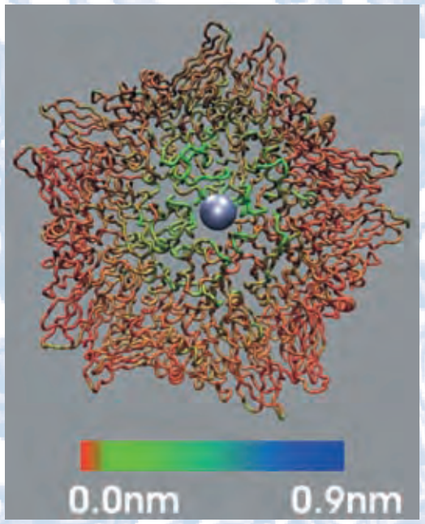

Figure 2: Indentation of the tip-sphere (grey) along the 5-fold symmetry axis of SBMV through the pentamer (shown as tube representation). The deformation of the amino acids is color-coded from red (no deformation) to blue (high deformation). The snapshot was taken at the yielding point at which the maximum force acting on the tip-sphere was obtained.

Applications

Our simulations showed a weak stiffness and stability, obtained as small elastic constants and yielding forces, respectively, at the subunit center of the inner and outer capsid surfaces of both shells, SBMV and HRV 16. The least stiff position of the HRV 16 capsid was found at the inner and outer pentamer center, as well as the inner pentamer center of SBMV. In contrast, the pentamer center on the outer shell of SBMV exhibited the largest elastic constant. The most stiff and stable position on the capsid of HRV 16 was found along the 2-fold symmetry axis. Large elastic constants and yielding forces along the 2-fold symmetry axis of SBMV were only found when the capsid was indented from the inside. On the outer capsid surface, the 2-fold symmetry axis exhibited a small stiffness and stability.

Although the assembly of subunits for SBMV and HRV 16 subunits is very similar, both capsids exhibit very different mechanical properties. The stiffness of hexamers and pentamers varies and depends on details on the atomic structure. Therefore no

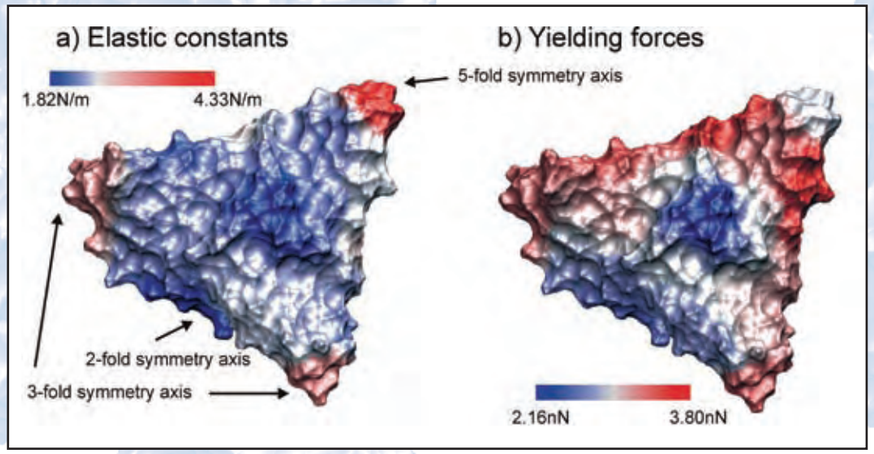

Figure 3: Subunit of SBMV. Color-coded distribution of elastic constants and yielding forces obtained from indentations of 19 positions on the outer subunit surface.

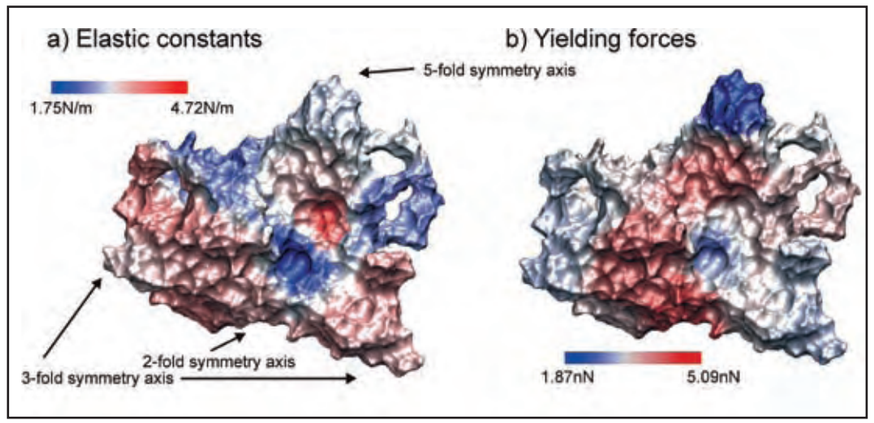

Figure 4: Subunit of HRV 16. Color-coded distribution of elastic constants and yielding forces shown as in Figure 2.

general picture for the distribution of mechanical properties on icosahedral capsids can be drawn. Our atomistic simulations clearly show that atomic detail cannot be neglected, because the atomic structure is a major determinant of the viral shell's mechanical properties, together with the geometrical arrangement of the subunits to an icosahedron.

The first step in viral infection of bean leaves with SBMV particles is the removal of the 180 calcium ions from the capsid structure. To study possible changes of the mechanical properties of SBMV without calcium compared to the coat protein with ions, the calcium ions were removed from the viral shell with subsequent equilibration of 32 ns with MD simulations. Subsequently, force-probe simulations were performed as described before for the capsid with ions. No significant change in mechanical properties was observed compared to the complete capsid with ions, with the exception of a marked softening along the 5-fold symmetry axis. Due to the fact that the pentamer center exhibited the highest stability when ions were present in the structure but became the least stable capsid position of the surface of SBMV without calcium, we suggest that dematuration and, thus, capsid rupture might occur at the pentamer centers. Therefore, the pentamer center might act as a possible port for RNA release after calcium removal.

The next systems studied with force-probe simulations were three different structural capsid mutants of HRV 16. The first mutation was the removal of the pocket factor from the hydrophobic pocket in protein VP1. Docking of the ICAM-1 receptor to protein VP1 and emptying of the pocket was proposed to cause a destabilization of the capsid. The HRV 16 structure includes 12 zinc ions located at the pentamer center and block a possible channel for RNA release along the 5-fold symmetry axis. If the removal of the zinc ions could cause a structural change of the capsid that becomes able to release the RNA, remains unclear. Additionally, removal of protein VP4 was proposed as the first step in viral infection. Therefore, the mechanical properties of the HRV 16 capsid without pocket factor, VP4 and zinc ions were studied in order to investigate changes in elasticity and fracture behavior from which a possible port for RNA release could be proposed. All three capsid mutants were indented with the tip-sphere at the same 19 grip point positions.

We obtained modifications in the mechanical properties on the inner and outer surfaces of the three HRV 16 capsid mutants compared to the initial structure, whereas variations in elastic constants and yielding forces were small compare to the complete HRV 16 shell. The largest changes were obtained along the symmetry axes, e.g. the pentamer center became stiffer after removal of the pocket factor, the 2-fold symmetry axis of the capsid without zinc was strengthened. These variations in mechanical properties compared to the complete HRV 16 capsid were not strong enough to corroborate a possible pathway for RNA release. Therefore, the proposed mutations do not result in dematuration of HRV 16. We suggest that additional mutations or combinations of the studied mutations are necessary to destabilize the capsid structure and result in a release of the RNA. In contrast to SBMV which becomes infections after the removal of calcium in an alkaline environment, docking of external receptors to the capsid surface might also be necessary to cause RNA release of HRV 16.

To summarize, mechanical properties of viral shells were investigated on atomic length scales for the first time. The simulations show that small changes of the capsid structure can result in significant changes of mechanical properties from which possible RNA ports could be proposed.

References

[1] **Roos, W. H., Ivanovska, I. L., Evilevitch, A., Wuite, G. J. L.**
Cell. Mol. Life Sci. 64, 1484-1497, 2007

[2] **Hess, B., Kutzner, C., van der Spoel, D., Lindahl, E.**
J. Chem. Theory Comput. 4, 435-447, 2008

[4] **Zink, M., Grubmüller, H.**
Biophys. J., in press, 2009

• Mareike Zink
• Helmut Grubmüller

Max Planck Institute for Biophysical Chemistry, Department for Theoretical and Computational Biophysics

Applications

The Origin of the Mass of the Visible Universe

The topic of this article is the origin of the mass of the visible universe. We attempt to understand the mass of ordinary matter without mass at all. Actually we want to explain all mass without mass, and the mechanism described below will explain about 95% of it. It is important to emphasize that we are speaking here about ordinary, visible matter. Dark matter or dark energy are not the topic of the present discussion. It is a popular opinion that the basic goal of LHC at Cern, Geneva (Switzerland) is to clarify the origin of mass. This is believed to be achieved by the so-called Higgs mechanism or some other ultimate mass generating mechanism. Indeed, LHC is the largest scientific equipment of mankind. It is in a 27 km long tunnel and the order of magnitude of the total costs is close to 10 billion dollars. Even Google changed its logo when the LHC was turned on (see figure 1).

It is important to note, however, that this ultimate mechanism is responsible only for the mass of the leptons and for the mass of the quarks. Interestingly enough, this is just a tiny fraction of the visible mass (such as stars, the earth, the reader and anything else, which is built up from atoms). Why is that? The electrons of ordinary matter are almost massless. The mass of an electron is just about 1/2000 of the mass of a nucleon (proton or neutron). The nucleus of an atom contains protons and neutrons, which are built up from gluons and quarks. Gluons are exactly massless and the quarks of ordinary matter are also almost massless. Actually, the vast majority of the visible mass (more than about 95%) comes through another mechanism. This mechanism and this 95% will be the main topic of the present work.

Figure 1: The Google logo at the time of the turn-on of the LHC. Given this coincidence, the depicted event in the center of the picture is probably a collision and not an explosion.

The surprising fact is that usually the mass of "some ordinary thing" is just the sum of the mass of its constituents (up to tiny corrections). The origin of the mass of the visible universe is dramatically different. As we have discussed, the protons and neutrons are made up of massless gluons and almost massless quarks. Their masses are much larger than the masses of their constituents. One can imagine an example as depicted in figure 2.

Let us discuss a hypothetical example which illustrates the order of magnitude of the mass generation associated with the nucleons of ordinary matter. The three cherries illustrate the quarks. They have a mass of 5 grams each. In real life the quarks of ordinary matter have a typical mass of 5 MeV (MeV is million electron Volt is the usual unit used in particle physics). The stem is supposed to be massless. In this example the hypothetical bound state of three cherries would have a mass of 1 kilogram, i.e. 1,000 grams. In real life the mass of the proton is indeed about 1,000 MeV. This analogy illustrates the dramatic effects which are at the core of the mass generation of the visible universe.

The paper is structured as follows. After this introductory section the basic idea of quantum chromo dynamics (QCD) is introduced in the first section. In the next section we show how the most important tool for particle physics, namely quantum field theory is constructed. The basic interactions of such a theory are discussed. Since the theory can not be solved analytically we solve it numerically. Then we present the technique, namely lattice gauge theory, for that. Finally we show the results and summarize the article. The original work has been presented in Ref. [1]. Discussions on this work can be found in [2,3].

Quantum Chromo-dynamics

Quantum chromo-dynamics is the theory of strong interactions. These interactions bind quarks and gluons into nucleons. Each quantum theory has its classical origin. For instance quantum mechanics is just a quantum formulation of the everyday Newtonian mechanics. Before we can discuss the quantum features of chromo-dynamics it is useful to study the classical version of it. The classical chromo-dynamics is a field theory. It means that the physics of such a theory is described by fields, thus by variables, which depend on space and time. A typical field theory is meteorology. The temperature, the pressure, the humidity and the speed of wind depend on space and time. The classical equations of meteorology describe the time evolution of those variables (they are given by the Navier-Stokes equations). Another typical field theory is classical electromagnetism. The electric and magnetic fields depend on space and time. The theory describes their time evolution (Maxwell equations).

Classical chromo-dynamics is just a generalization of electromagnetism (electro-dynamics). In electro-dynamics we have one electric charge. It can be negative or positive. In chromo-dynamics we have three different charges, we call them "colors" (hence the wording "chromo"). Each of them can be negative or positive. Clearly, they are not real colors, we just call them so, because by mixing all three of them we obtain a color neutral combination (in analogy to the real world in which a mixture of the basic colors green, red, and blue yields a neutral white color).

In the quantized version of electrodynamics (quantum electro-dynamics or QED) photons transmit the interactions. Gluons are similar to photons, they transmit the strong interactions between color-charged quarks. The wording "strong" reflects a relative ordering, namely the fact that the coupling of gluons to color-charged particles are much stronger that that

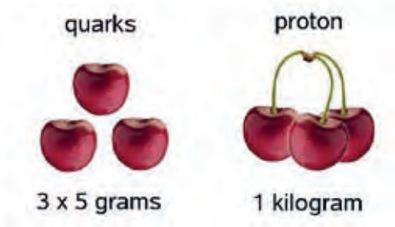

Figure 2: The three cherries – acting as quarks of a nucleon – have masses of 5 grams each. In this hypothetical example they obtain a mass of 1 kilogram when a massless stem binds them together.

of the photons to electric charges. It is important to note that chromo-dynamics has a richer structure than electro-dynamics. Gluons carry a color charge, while photons do not carry an electric charge themselves. This has very deep theoretical consequences.

Despite the appealing features of this similarity there is one important difference. Though one can nicely formulate the theory of the strong interactions with the help of quark and gluonic degrees of freedoms, these particles were never observed experimentally as freely moving particles. They appear always in bound states. The fundamental degrees of freedoms have never been seen.

The basic equations of classical electro-dynamics can be obtained – essentially unambiguously – by using one single symmetry, called gauge symmetry. In order to fulfill this symmetry we assume that the basic equations remain the same if the quark fields are multiplied by a complex phase factor $\exp(i\theta)$. Here θ can depend on space and time. This symmetry condition fixes the form of the interaction. The equations of classical chromo-dynamics can be obtained analogously. Since we have three different color charges the symmetry condition implies that the equations remain unchanged if one multiplies the quark fields by a three-by-three matrix-like phase factor (these matrices are called SU(3) matrices).

Particle Physics and Quantum Field Theory

The basic tool to understand particle physics is quantum field theory. Quantum field theory is a quantum theory and a field theory. The most important difference of a quantum theory compared to a classical theory is that some observables are not commuting, they are treated as operators. Examples for commuting quantities are the numbers in multiplications (e.g. 5*6 = 6*5, the order of the numbers under multiplication is not important). The classical example for a non-commuting operation is rotation. It does make a difference if we turn our face from south to west and then lay down with face down, or we look to south, lay down with face down and then turn our face to west. In quantum mechanics multiplying the position and momentum operators or multiplying them in the opposite order yields different results. This difference leads to important consequences. In classical physics the energy of a positively charged and negatively charged binary system (such as a hydrogen atom with a positively charged proton and a negatively charged electron) can be "anything", it is a continuous variable. In quantum mechanics the energy spectrum is quantized, the energy can take only some definite values.

In field theories it is not the position and the momentum which are treated as operators, but the field variables. As a consequence we have well defined energy levels, which vary in space and time. This variation can be seen as energy packages, which move in space as time goes on. This is exactly our intuitive picture of a particle. An energy package (which eventually has other properties such as charge or spin) moves as time goes on. The energy package, which is related to the electromagnetic field is the photon.

The symmetry properties (gauge invariance) are already defined the underlying equations. Quantizing the theory (treating the variables as operators) does not change the equations. It merely leads to the picture of the moving energy package, thus adding the particle interpretation. The classical equations describe the interactions between the fields, or in the quantized form they describe the interactions between the particles.

In quantum electro-dynamics we have two types of particles. Photons and electrons (or its positively charged partner the positron). There is only one type of interaction, which is shown in figure 3. On the left hand side we see an event at LEP, CERN (Switzerland). It was a huge accerelator (actually the same tunnel was used for it as it will be used after 2009 for LHC). One of the four experiments was Delphi. A typical event was an electron-positron pair production. (Or equally often muon-antimuon, or tau-antitau production. They are just particles with exactly the same properties like electrons and positrons, just somewhat heavier. We call them together "leptons".) Occasionally, an outgoing lepton emits a photon. This process is the basic interaction in quantum electrodynamics. And indeed, we see three tracks in the detector. The particle coming out to the right emitted a photon upwards. Since an outgoing positron is equivalent to an incoming electron the observed process can be graphically represented as shown on the right hand side of this figure. An electron comes in, emits a photon and goes out. This way of graphically representing various processes turned out to be very useful. We call them Feynman graphs. The strength of the interaction can be characterized by the fine structure constant. Its value is $\alpha = 1/137$, and indeed the three particle LEP processes are observed about in 1% of all QED processes. It is important to note that this is the only elementary QED process. All other interactions are built up from this single process. The predictive power of the theory can be already seen by the fact that the symmetry condition unambiguously defines the interaction, which describes for instance the angle distribution of the three particle QED process shown in the figure.

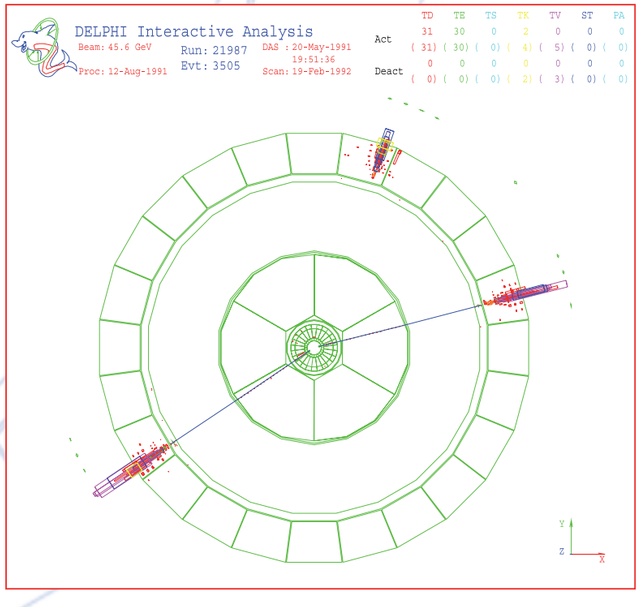

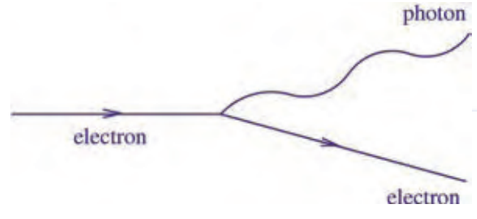

Figure 3: Three particle QED event at LEP (left panel). The Feynman graph of the single elementary process in QED (right panel).

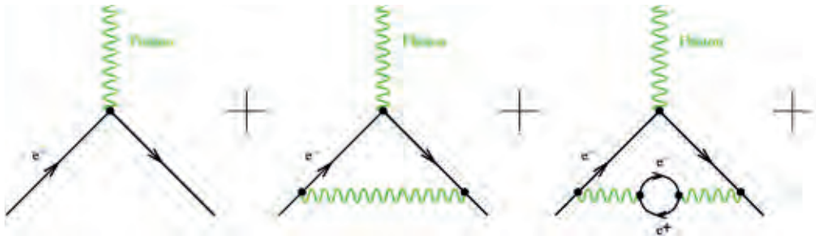

Figure 4: Feynman graphs contributing to the magnetic moment of the electron.

The smallness of the coupling provides us with a theoretical tool, which can be seen as a Taylor expansion in α. One example for such an expansion is the magnetic moment of the electron. Figure 4 shows the various Feynman diagrams that contribute to this process. On the first graph an incoming electron interacts with the external magnetic field (we are in a quantum theory, thus all fields are treated as particles, in this case as a photon). After the interaction it goes out. On the second graph the incoming electron emits a photon, then interacts with the external magnetic field and finally re-absorbs the emitted photon. On the third graph the emitted photon creates an electron-positron pair, which re-annihilates into a photon, before the latter is re-absorbed by the electron.

Since the fine structure constant is a small number (α = 1/137) each term is about 1% of the previous one. Adding up all possible graphs one obtain a very impressive 13 digit agreement with experiment. This is the most precise prediction in science. The experimental value is μ_e=2.00231930443622(14), whereas the theoretical calculation based on the above graphical expansion (obviously with many more terms) gives μ_e=2.0023193044352(16). The agreement is more than attractive.

The same LEP accelerator produces also QCD events. The left panel of figure 5 shows such an event of the Opal experiment. This elementary process is more complicated than the QED process. Instead of outgoing single particles jets are observed. Jets are groups of hadronic particles with similar momentum. Hadrons are composite particles built up from quarks and gluons. In the detector we see three jets. These jets correspond to the elementary process shown on the right panel. An incoming quark emits a gluon and goes out. This sort of three jet event appears in about 10% of all QCD like events. Since this is a much larger ratio than observed in QED processes we conclude that the interaction is much stronger (hence the name strong interaction). A very deep feature of the theory – the duality of asymptotic freedom and infrared slavery – is already present in this three jet event. At large energies the coupling gets smaller and smaller, this phenomenon is called asymptotic freedom, the particles behave as almost free particles. For small energies – we call it infrared regime – the coupling gets so large that the quarks are confined, a phenomenon called infrared slavery.

Lattice QCD

It is certainly a surprising fact that the degrees of freedom of the theory do not appear experimentally. There is no sign of the same physical content. The theory contains massless gluons and almost massless quarks. We detect none of them, they are confined. We detect instead composite particles: protons, pions. It is important to mention that the nucleon (the most important hadron) is several hundred times heavier than the quarks. This underlines the question we asked in the introduction. How and when was the mass generated. Before we study the question in more detail it is illustrative to summarize the most important qualitative features of the mass generation. Though the illustration is clearly incomplete, it contains most of the important ingredients of the more complete calculation. In the early universe the temperature was very high. The high temperature means motion. The motion was diluted by the expansion of the early universe. Nevertheless a small fraction of this motion remained with us confined in protons. As a consequence the kinetic energy inside the proton is observed as the mass of the particle (energy and mass is equivalent as Einstein's famous equation predicts: $E=mc^2$).

As we have seen the coupling in the strong interaction is quite large and it is particularly large for energies equivalent to typical hadron masses. No perturbative method (Taylor expansion in the coupling) works in this regime. For this sort of questions another systematic approach can be applied, namely lattice QCD. Instead of using a continuous space-time variable we put our equations on a space-time grid, a so-called lattice. This procedure is quite often used in various field theories. For instance weather forecast is done similarly. Aircrafts measure the pressure, temperature and wind velocity and scientists use the data, restricted to a grid, as input field variables to solve the Navier-Stokes equations. In lattice QCD

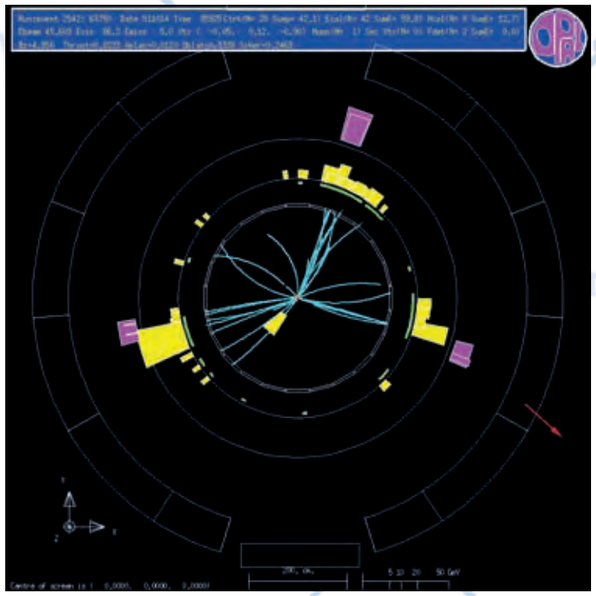

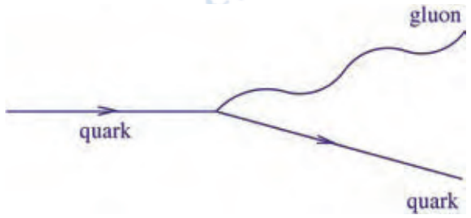

Figure 5: Event of the Opal experiment (left panel). Three jets correspond to the elementary process (right panel).

we put the equations on the lattice and solve them numerically. Such a lattice is shown in figure 6.

In order to be able to resolve the structure of the proton the lattice has to be fine enough. Typically 0.1 fm (thus 0.00000000000001 mm) is used for that purpose. The lattices have about 50 points in each direction. Since each lattice point carries dozens of variables, the numerical treatment of such a system is quite difficult. Mathematically it corresponds to a one billion dimensional integral. Clearly advanced techniques and supercomputers are needed. Such a high performance supercomputer is the Blue-Gene system in Jülich (see figure 7).

With the help of lattice techniques and supercomputers typical vacuum field configurations are generated. For the classical theory the vacuum is just a trivial configuration with vanishing field strengths. In the quantum theory, however, the vacuum is fluctuating around this classical configuration. Typically a few hundred of those configurations are enough to calculate various observables with a few percent accuracy. In a full numerical analysis special attention should be payed to control systematics. The two most important ingredients of such an analysis are the quark masses and the lattice spacings. The quark masses should be close enough to their physical values. The lattice spacing should be small enough to carry out a controlled continuum extrapolation (extrapolating the numerical results to zero lattice spacing).

In our approach the discretized form of the theory was chosen in such a way that the continuum result could be reached very fast. Effects due to non-vanishing lattice spacings are quite small. Since these effects strongly influence how much computer time is needed for a full calculation, our choice of action (discretized form of the theory) is one of the most important ingredients to reach our goals.

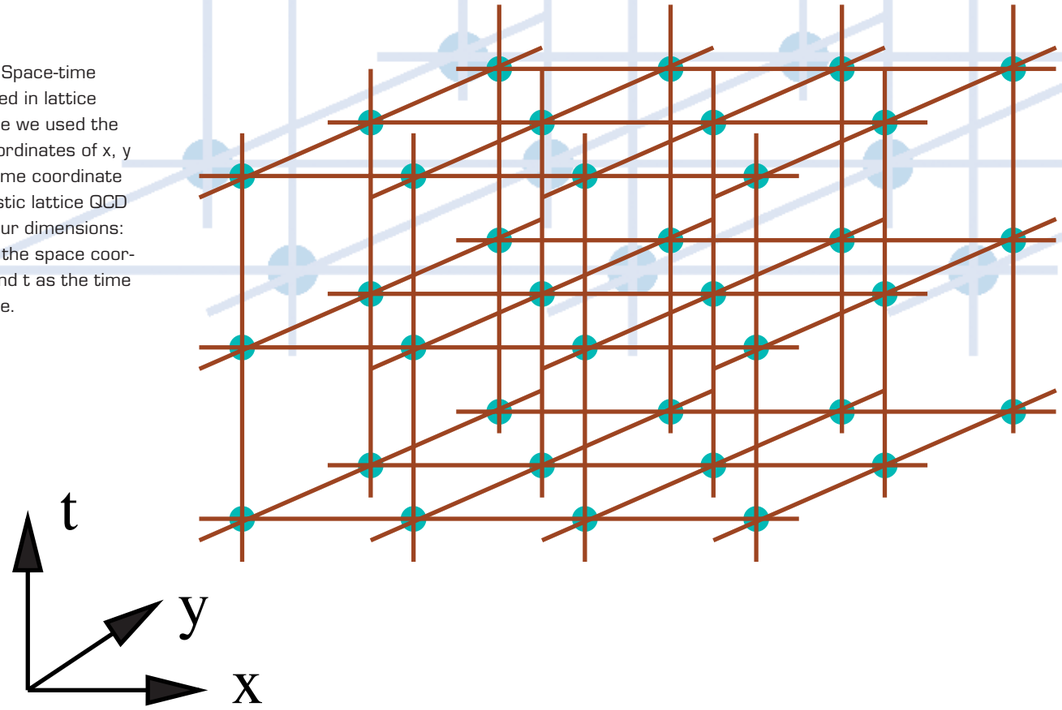

Figure 6: Space-time lattice used in lattice QCD. Here we used the space coordinates of x, y and the time coordinate t. In realistic lattice QCD we use four dimensions: x, y, z for the space coordinates and t as the time coordinate.

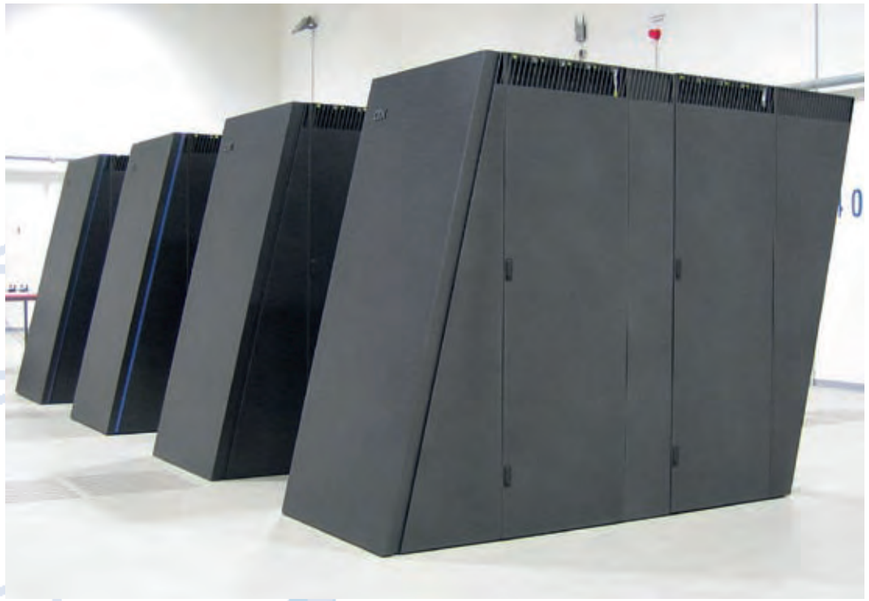

Figure 7: The Blue-Gene supercomputer in Jülich

Another important ingredient of such a large scale computation is the implementation of the algorithm most suited for the hardware enviroment. As it turned out our code is one of the codes with high scalability. It scales from one midplane of the Blue-Gene system to 16 Racks without any loss of performance. This scaling study has been done in the strong scaling scheme. A strong scaling analysis means that the problem size is fixed, however it is distributed on more and more processors. This scheme is much more sensitive to the optimization of the code than the so-called weak scaling analysis, for which the size of the system is fixed for one processor and more and more processor is taken. Thus the system under study increases, too. In the strong scaling scheme the communication overhead increases with the number of processors involved in the analysis. For the weak scaling the communication overhead stays constant (in leading order). The particularly attractive feature of our implementation is that strong scaling is achieved with a very high sustained performance. Figure 8 shows our strong scaling curve. As it can be seen about 37% performance is reached and the scaling is essentially linear.

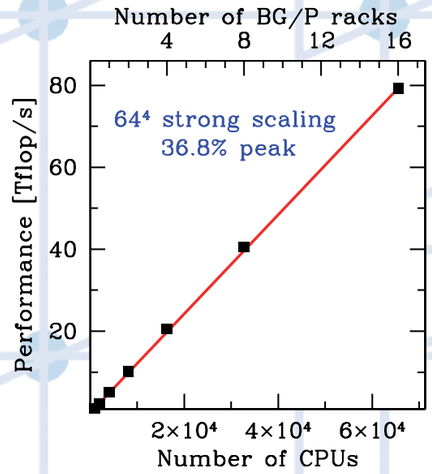

Figure 8: Scaling curve of our lattice QCD code. A linear strong scaling is obtained with a very high sustained performance of about 37%.

The Mass of the Nucleon and Other Hadrons

Once the typical vacuum configurations are generated they can be used to calculate the masses of hadrons. One puts a composite hadron state into the vacuum and looks how it propagates. The details of its propagation provides us with the masses.

The Budapest-Marseille-Wuppertal Collaboration determined the masses of the low lying hadrons. These hadrons are composed of light and strange quarks. Light quarks (up and down) have masses of around 5 MeV, whereas the strange mass is much heavier, it is about 100 MeV. Figure 8 shows the masses of a dozen hadrons which are composed of up, down and strange quarks. In the first column the hadrons (pion and kaon) are composed of a quark and antiquark and have zero angular momentum. The hadrons in the second column (rho and K-star) are composed of a quark and antiquark and have one unit of angular momentum. The third column shows the baryon octet. They are composed of three quarks and have angular momentum 1/2. The most prominent member of this group is the nucleon (proton or neutron), which gives the vast majority of the mass of the visible universe. The fourth column shows the baryon decuplet, thus hadrons built up from three quarks with an angular momentum of 3/2.

The horizontal lines show the experimental values of the hadron masses. Since about half of these hadrons are unstable, the grey bands show the width of these resonances.

We have three input parameters. The up and down quark masses (we treat them degenerate) are given by the pion mass. The kaon mass is particularly sensitive to the strange mass, therefore we use it to set the strange mass. Since QCD predicts only dimensionless combinations, an overall scale is needed. We use the mass of the Xi baryon. These input parameters are represented by open blue symbols in figure 9.

The predictions of our calculations are shown by the red symbols with error-bars. As it can be seen there is a perfect agreement between the QCD results and the experimental values.

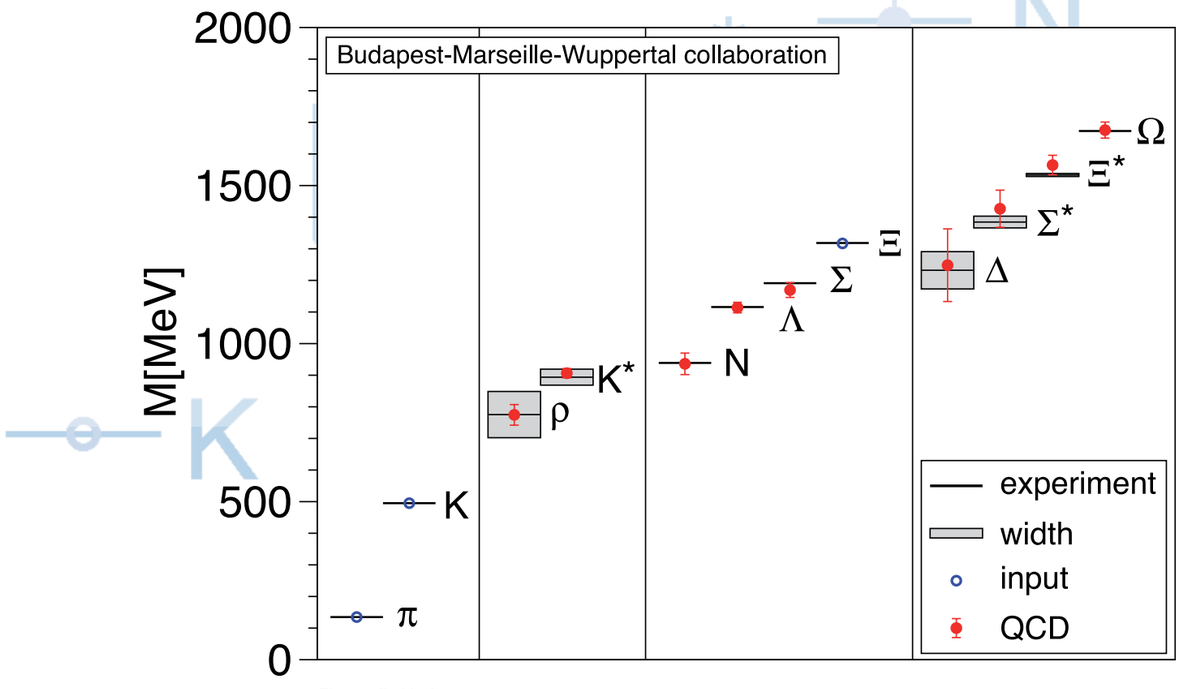

Figure 9: Hadrons spectrum

Conclusions

In this article we discussed a recent result of lattice QCD. A detailed introduction was presented and the most important concepts of quantum field theory and particle physics were summarized. The similarities between quantum electrodynamics and quantum chromodynamics were discussed. Since the theory can not be treated analytically in the low energy region, we introduced the concept of lattice QCD. Using this technique the hadron spectrum has been determined with a few percent accuracy.

References

[1] **Durr, S., et al.**
Ab Initio Determination of Light Hadron Masses, Science 322: 1224-1227, 2008.

[2] **Wilczek, F.**
Particle Physics: Mass by Numbers, Nature 456: 449-450, 2008

[3] **Kronfeld, A. S.**
The Weight of the World is Quantum Chromodynamics, Science 322: 1198-1199, 2008

[4] The author thanks S. Durr for careful reading of the manuscript

• Zoltan Fodor

University of Wuppertal

Heavy Quark Physics on Large Lattices

It is widely expected that the modern theory of elementary particle interactions – the so-called standard model – is only a low-energy limit of a more general theory that unifies all known interactions including gravity. Deviations from predictions of the standard model are expected to occur at small distances, which can be probed in collisions at very high energy, or, alternatively, manifest themselves in quantum fluctuations that accompany the decays of hadrons containing heavy quarks. Discrepancies between theory and experiment can be indications for new physics and are attracting a lot of attention.

The task of our project is to calculate two types of observables with high precision: the so-called decay constants and form factors. The decay constants are essentially probability amplitudes for finding a heavy meson in a state where the heavy quark and the light antiquark, which make up the meson, are at the same space-time point. The form factors describe the probability that the light quark produced in the decay of the heavy quark recombines with an antiquark producing a hadron of a given type in the final state. They depend on the energy release during the decay, usually referred to as the momentum transfer. The theoretical basis for the calculation is QCD (Quantum Chromodynamics), a quantum field theory which describes the interaction of the quarks through the exchange of gluons, the quanta of the so-called color gauge field. For the calculation of the above observables an approach beyond perturbation theory is required, and we have adopted lattice QCD for this purpose. In lattice QCD, space and time are substituted by a four-dimensional Euclidean lattice. The quantum fluctuations of the color gauge fields are simulated using a Monte Carlo method, and a number of "pictures" of gauge field configurations generated during the simulations is then stored. In a second step, for each of these configurations quark-antiquark pairs are created at a given time, propagate through the color field, and the amplitude for annihilating the pair at a later time is measured. From this amplitude (a so-called two-point function) one can extract meson masses and decay constants. For the evaluation of form factors an additional "current" has to be inserted at some time between creation and annihilation resulting in a three-point function.

Simulations are performed for several values of the quark masses, which do not necessarily agree with the physical values. Final results are then found by extra- and interpolation in the quark masses. It is necessary to compare results from different simulations, in particular with different lattice spacings, to ascertain that the systematic errors are under control. The goal of our project is therefore to contribute to this effort by lattice simulations of systems containing heavy quarks using a very large and fine lattice of size $40^3 \times 80$, with a lattice spacing of about 0.04 fm. Very fine lattices such as this are at present only treatable in the so-called quenched approximation, where quantum fluctuations due to virtual quark-antiquark pairs are neglected. We work with non-perturbatively improved Wilson quarks and cover a large range of masses. For the light quarks we take masses around the strange quark while the heavy quark mass ranges from values around the charm quark mass up to 3/4 of the bottom quark mass. The calculations require programs to 1) generate the gauge field configurations, to 2) calculate the quark propagators (i.e. columns of the inverse of the lattice Dirac operator) and to 3) construct the meson propagators. We have applied the publicly available MILC code [1] with some modifications to generate the gauge field configurations. Our Fortran based code for steps 2 and 3 uses a fast MPI and OMP parallel code for inverting the lattice Dirac operator, which in our case is a large sparse complex matrix of rank 61,440,000. We have implemented the so-called bicgstab algorithm, incorporating a conjugate gradient routine written by H. Stüben [2].

Our result for the decay constant of a meson containing a charm and an anti-strange quark [3] is shown in figure 1 together with results from other recent quenched lattice calculations. The figure suggests that lattice spacing effects may indeed be large on coarser lattices. Our result is lower but within errors in acceptable agreement with the unquenched result of Ref. [7], supporting the view that the lattice results are systematically lower than the experimental value. Of particular interest in our calculation are form factors of semi-leptonic decays of D mesons (consisting of a charm quark and a light antiquark) and B mesons (consisting of a bottom quark and a light antiquark) since they were previously not available from fine lattices. Results for the form factors f_0 and f_+ are shown for a strange light quark and three values of the mass of the heavy quark in figure 2 [8]. The extrapolation or interpolation to the physical masses still needs to be finished, but the results seem to be in agreement with recent unquenched calculations.

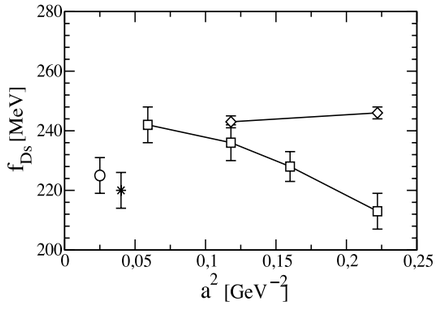

Figure 1: Decay constant of the D_s meson from quenched simulations with improved Wilson quarks at various squared lattice spacings a^2 given in GeV^{-2}. The star denotes our result. Results from other investigations are shown as diamonds [4], squares [5] and a circle [6].

References

[1] http://www.physics.utah.edu/~detar/milc/

[2] **Stüben, H.**
Hadron Masses from Dynamical, Non-perturbatively O(a) Improved Wilson Fermions, Nuclear Physics (Proceedings Supplement) 94, 273, 2001

[3] **Ali , A., et al.**
Decay Constants of Charm and Beauty Pseudoscalar Heavy-light Mesons on Fine Lattices, Physics Letters B652, 150, 2007

[4] **Bowler, K. C., et al.**
Decay Constants of B and D Mesons from Non-perturbatively Improved Lattice QCD, Nuclear Physics B619, 507, 2001

[5] **Jüttner, A., Rolf, J.**
A Precise Determination of the Decay Constant of the D_s Meson in Quenched QCD, Physics Letters B560, 59, 2003

[6] **Jüttner, A.**
Ph.D. thesis, arXiv:hep-lat/0503040

[7] **Follana, E., Davies, C. T. H., Lepage, G. P., Shigemitsu, J.**
High Precision Determination of the π, K, D and D_s Decay Constants from Lattice QCD, Physical Review Letters 100, 062002, 2008

[8] **Ali Khan, A., et al.**
Decays of Mesons with Charm Quarks on the Lattice, Proceedings of Science, LATTICE2007, 343, 2007

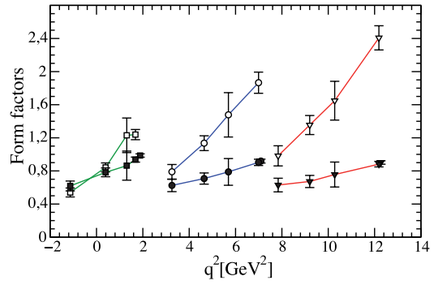

Figure 2: Form factors f_0 (filled symbols) and f_+ (open symbols) for semi-leptonic decays at three values of the heavy quark mass shown as functions of the momentum transfer q^2. The corresponding meson masses are 2.1 GeV (squares), 3.4 GeV (circles) and 4.2 GeV (triangles).

• Arifa Ali Khan

Taiz University, Yemen, Faculty of Science, Department of Physics

Powers of R: Supercomputing for Everybody

Efficient and fast solving of large scale numerical problems is only possible through massive parallelism. However, it is hard to write code that parallelizes well, especially when very large data-sets have to be analyzed and no predefined model is available. For this, dynamic languages are well suited since they are easy to learn and use and, furthermore, allow complex data-manipulation with only a few lines of code. The price for this flexibility is typically bad performance. This article demonstrates how to use the dynamic language R for straight forward data analysis on a supercomputer and especially how to accelerate R code and how to achieve strong scaling up to hundreds of cores.

What is R?

R is a dynamic language for numerical computing and graphics with a strong affinity to statistics. R is available as Free Software under the terms of the GNU General Public License (GPL). It compiles and runs on a wide variety of UNIX and Linux platforms, Microsoft Windows and MacOS. R is a fully featured programming language and much of the system itself is written in R. Advanced users can link and call C, C++ and Fortran code at run time. R has its roots in statistics, but its extensibility, ease-of-use and powerful graphics makes it ideal for users looking for a fast, easy and robust environment for data analysis and numerics. R can easily be extended with more than 1,700 additional packages available through the Internet that can be installed with the command `install.package` ("name") and then loaded with the function `library (name)`.

An important remark about the equal sign in R: `x<-1` is equivalent to `x=1` but the latter should be used wherever optional parameters are expected like in function calls. Use the arrow form for explicit calculations and the equal sign form for definitions.

R has a very simple syntax. Almost everything is a function. Functions are defined by the function keyword and stored in a named variable. Functions can be used as arguments to other functions and even return functions.

```
function(x,y,z,...) {R commands}      define an unnamed function object
myfunc = function() {...}             define a function and store it as
                                      myfunc which can then be used as
                                      myfunc()
```

Here are some important builtin functions of R:

```
c(x,y,z,...)                          concatenate items to a list
runif (npoints)                       uniform random numbers between 0 and 1
ifelse (condition, yes, no)           if-then-else as a function
system.time (command)                 calculate elapsed system time of command
```


Here are some important matrix and array functions:

```
array (X, dims=c(nx, ny))             make matrix with dimension nx times ny
t(x)                                  transpose x
sum(x)                                calculate total sum of all elements of x
colSums(x)                            calculate sum for all columns of x
unlist(x)                             rearrange multidimensional array x to
                                      one-dimensional array
```

The following functions take a function and apply it to each element of an array:

```
lapply(X, func)                       apply function func on linear array X
mclapply(X, func, mc.cores=n)         apply function func for package multicore
clusterApplyLB(cluster, X, func)      apply function func for package snow
```

Example: Fractal Dimension

To demonstrate the merits of using a dynamic language like R, we chose a simple yet non-trivial example: Computation of the fractal dimension of a set of points. Figure 1 shows a few typical examples with fractal dimension 0, 1, 2 and 3. To calculate the fractal dimension of a given set of points, one counts the number of points inside a sphere of given radius `r` (so-called pointwise correlation dimension). The number of points inside the sphere is proportional to r^n where the exponent n is a measure of the fractal dimension of the set. To obtain an unambiguous number, one calculates the derivative with respect to r of the logarithm of n.

The derivative is approximated by the difference quotient. For ideal situations, the fractal dimension is independent of the chosen radius r. However, in real world situations with finite set size, one has to choose r and can only calculate an approximation to the fractal dimension which depends on r.

The fractal dimension has to be calculated for each point in the set (Figure 1). To determine whether a neighboring point falls inside the sphere, distances to every other point in the set have to be calculated. This makes the algorithm computationally expensive, yet easily parallelizable since the calculations for each point are independent from each other.

Using R

We start by generating a random set of a million points (`np=10**6`) in three dimensional space (`ndim=3`). The built-in function `runif` produces a matrix x of dimension 10^6 times 3 (the matrix has 10^6 columns and 3 rows) and consists of uniformly distributed random floating point numbers between 0 and 1. Next, two functions `rho` and `fracdim` are

```
## data definitions
np = 10**6
ndim = 3
x = array(runif(np*ndim),c(ndim,np))
```

defined. The function `rho` calculates the number of points within the sphere of radius `r` and `fracdim` returns the fractal dimension as described above with radius `r1` and `r2` chosen by the user to calculate the difference quotient. Each column of x represents the coordinates of one point of the set. In the function `rho` the index i chooses one point, subtracts the coordinates of this point from all the other points, squares the result and adds up all three elements of the difference vector (`colSums`). Thus for each point the relative distance to the chosen point with index i is computed. Now it is decided whether the distance is smaller than a threshold r which gives a logical true/false value. These values are then added up for all points of the set, where true is represented as 1 and false as 0. Finally, one obtains the number of points in a sphere around the chosen point. This function rho is further used to calculate the fractal dimension for two chosen radii `r1` and `r2`. Ideally the fractal dimension should be independent of `r1` and `r2`, but in real cases this is seldom realized. One has therefore to select sensible values for `r1` and `r2` which we choose here as 0.1 and 0.2, respectively.

```
## function definitions
rho = function(i,x,r) sum(colSums((x-x[,i])**2) < r**2)
fracdim = function(i,x,r1,r2) (log(rho(i,x,r1))-log(rho(i,x,r2)))/(log(r1)-log(r2))
```

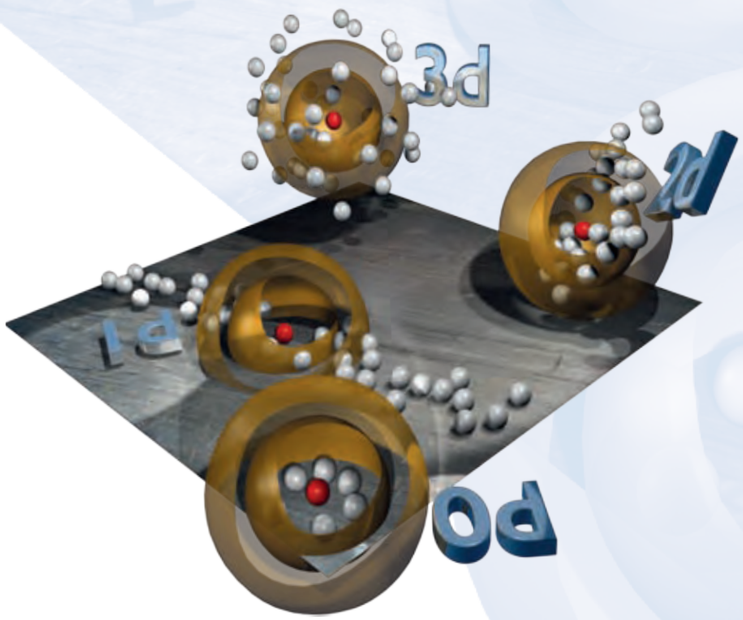

Figure 1: Calculating the pointwise fractal dimension of a set of points: The figure shows four sets of points with dimension zero, one, two and three. To determine the fractal dimension of a selected point (red), one counts the number of surrounding points (white) in different sized spheres and calculates the scaling behavior. The exponent of the scaling of the number of points for different radii is the local fractal dimension. The ordinary (topological) dimension of a point set is then the average for all points of the set.

E.g. for a plane-like structure the number of points scales with the square of the radius of the enclosing sphere, which means a scaling exponent of 2 which is also the topological dimension of a plane.

Easy Parallelization

Now, as a first test, we calculate the fractal dimension for 1,000 points only. To do this, instead of a for-loop, we apply the function fracdim to the index list (`1:1,000`), which can be done by using the list apply function lapply. Since we also want to benchmark the different approaches, we measure the elapsed system time with the function system.time.

```
#  compute fractal dimension for first 1000 points only
      fraclist <- lapply(1:1000, function(i) fracdim(i,x,0.1,0.2))

#  measure elapsed time for first 1000 points
      system.time(fraclist <- lapply(1:1000, function(i) fracdim(i,x,0.1,0.2)))

#  print mean fractal dimension for first 1000 points
      print(mean(unlist(fraclist)))
```

On a 2.6 GHz Opteron, this will take about 17 minutes on one core. How can we speed this up? The simplest trick is to use the package multicore with which the application of the function to the indexed list is automatically distributed to the number of cores chosen by the user (keyword `mc.cores`). For example, we choose 16 cores. The call to `lapply` must be replaced by `mclapply`. The same computation of the fractal dimension is now 16 times faster.

```
## multicore
library(multicore)

fraclist <- mclapply(1:1000, function(i) fracdim(i,x,0.1,0.2), mc.cores=16)
print(mean(unlist(fraclist)))
```

Not fast enough for you? If all the cores of one machine are maxed out, you can distribute the work to a compute cluster using the package snow (simple network of workers). The function `makeSOCKcluster` allocates the cluster and returns a list of worker nodes and `clusterExport` distributes the data. Lapply is now replaced by `clusterApplyLB`. The suffix LB stands for load balancing which means that the work load is balanced to all available cluster nodes. In our example we deploy a cluster of 32 Altix nodes. Of course we need to reserve the cores first, executing in a batch job we would have to specify the node names of the individual machines in order to use cores on different machines. The cluster can be heterogeneous and different architectures can easily be combined, given that R is available on the target platforms.

```
##snow
library(snow)
cl=makeSOCKcluster(32)
clusterExport(cl,list("x","np","ndim","rho","fracdim"))

fraclist <- clusterApplyLB(cl, 1:1000, function(i) fracdim(i,x,0.1,0.2))
print(mean(unlist(fraclist)))
```

Benchmarks

Now we focus on the scaling behavior of the presented R code. Do we get a speedup of the calculation when we use more cores or is the computing time bounded by the network communication? Figure 2 shows the elapsed time as a function of the number of processor cores in a log-log plot. Both the multicore package and the snow package scale very linearly up to about 100 cores. That means we observe an acceleration of the program by a factor of 100 with respect to the single core variant. For larger numbers of cores a saturation of the computing time can be observed which is due to Amdahl's Law, which states that for many cores the startup and communication overhead becomes dominant over the computation time per core. This could be overcome by enlarging the problem size and thus the computing time per core. In the present example the computing time per core is of the order of 2 seconds for 500 cores which is a lower limit for the efficiency of parallel computation. The computations were performed on one partition of the SGI Altix 4700 of the LRZ which consists of 510 cores. For the multicore package an SMP machine (like one partition of the Altix) is needed, while the snow package can access as many cores as are available, even on multiple machines, i.e. in the case of the Altix a total of 9,728 cores.

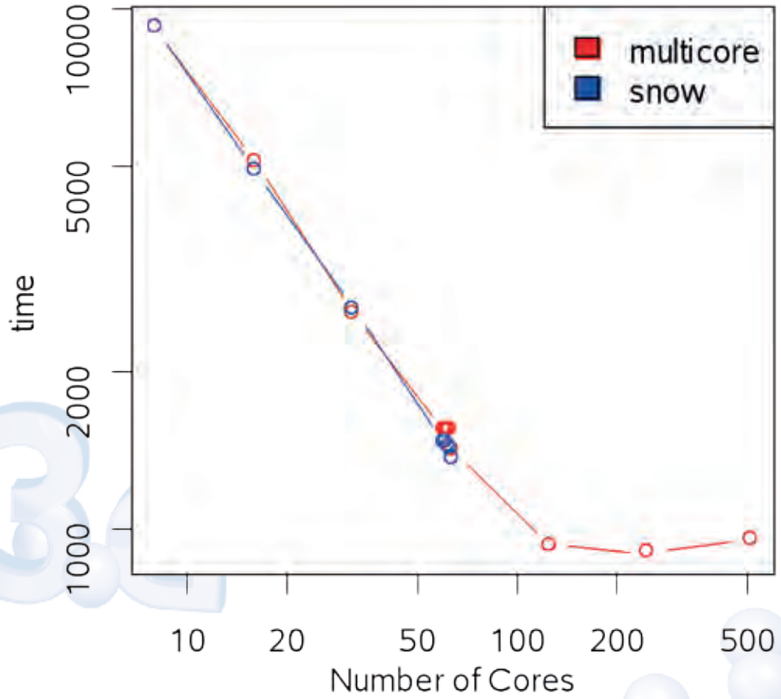

Figure 2: Scaling of the presented R code to determine the fractal dimension of each point in a point set with a million points with the multicore (blue) and snow (red) packages for parallel computing. The computing time scales with one over the number of cores until the communication overhead becomes too large (Amdahl's Law). Due to the decrease of the per-core problem size we observe a plateau of the computing time at around 200 cores.

Conclusion

R is a very powerful dynamic language and its high abstraction makes it easy to quickly test new models against very large data sets. Incorporating the multicore and snow packages enable strong scaling of R code without the need for time consuming low-level programming. This demonstrates how massively parallel supercomputing can be combined with high abstraction, making large-scale data analysis both feasible and easy.

Links

http://www.r-project.org/

http://www.lrz-muenchen.de/services/compute/visualization

References

[1] **Grassberger, P., Procaccia, I.** Measuring the Strangeness of Strange Attractors, Physica D: Nonlinear Phenomena 9 (1-2): 189-208, 1983

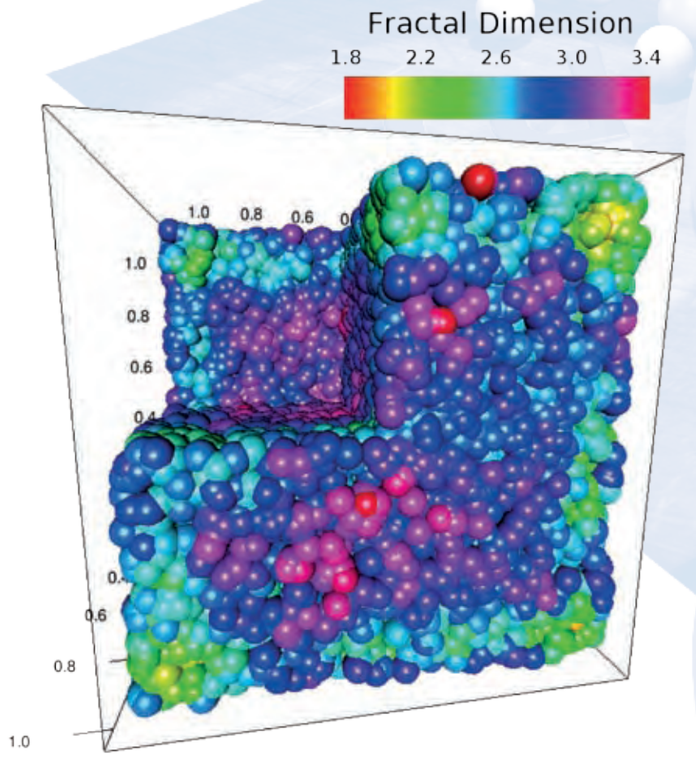

Figure 3: Visualization of the fractal dimension of a point set with 50,000 elements. The data set consists of uniformly distributed data points in 3D (unit cube). The color coding represents the computed pointwise fractal dimension of the data set. The color plot has been generated using the R package "rgl" with the following command:

```
plot3d(x[1,],x[2,], x[3,], type="s", size=1, col=hsv(fraclist))
```

• Ferdinand Jamitzky
• Helmut Satzger

Leibniz Supercomputing Centre

Simulating Molecular Switches at Work

A wide range of cellular functions depends on the propagation of chemical signals from the cell surface into the interior. Through these signals, the cell becomes aware of its ever-changing environment, and is able to respond in appropriate ways. For example, these signals may trigger the transcription of DNA, or activate an immune response. They may also induce the cell to generate another signal for a neighboring cell, as in the neuronal synapses. It is therefore not surprising that many human diseases originate from a deficiency in the cell's ability to transduce these external stimuli correctly. For example, cancer often develops in tissues where cells are no longer able to propagate signals for self-suicide (or apoptosis).

By and large, signal transduction is mediated by proteins, organized in complex interaction networks. Some of these proteins reside in the cell membrane and therefore serve as the cell's primary sensors, or "receptors". Many others are within the cell interior, where they diffuse in a crowded, jelly-like environment, constantly interacting with each other and with those in the membrane. An essential characteristic of most protein interaction networks is that they are dynamic – or in other words, regulated. This allows the cell to amplify, inhibit or even shut down the propagation of chemical signals whenever necessary. At the molecular level, the dynamical nature of signaling networks is possible because many of these proteins act as molecular switches. That is to say, their chemical (or enzymatic) activity is turned on or off depending on their location at a given time, and on which other proteins they can interact with in that context. Another important class of signaling proteins is often referred to as "scaffolds"; these cannot carry out any chemical work, but they help to bring together others that do, and therefore contribute to the spatial and temporal regulation of signal transduction networks. Finally, selective evolution is thought to have driven combinations of these signaling proteins to fuse into larger macromolecular complexes, where, for example, enzyme domains are combined with scaffolding domains, and so on.

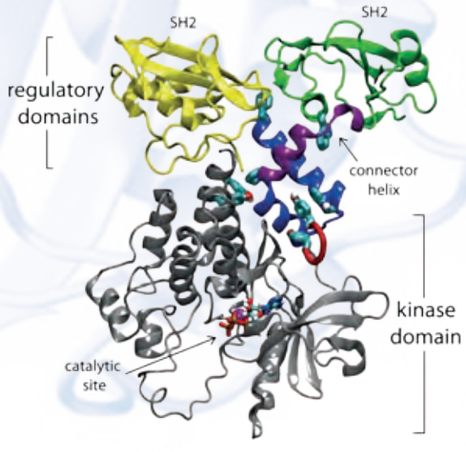

Figure 1: Cartoon representation of the non-receptor tyrosine kinase ZAP-70 in the auto-inhibited state (Deindl, et al., Cell 2007, 129: pp. 735-46)

A particularly important class of signaling proteins are the non-receptor tyrosine kinases (NRTK). These proteins propagate signals by virtue of their ability to phosphorylate other proteins (i.e. attach a phosphate group to the side chain of a tyrosine amino acid), and thus switch them from the "off" to the "on" state or vice versa. Among the NRTKs, those in the so-called Src and Syk families are especially important, for example, for our adaptive immune system. From a molecular perspective, they are also interesting, first, in that they belong to the class of signaling proteins where enzymatic domains are fused with scaffolding domains, which serve to direct these proteins to specific locations in the cell. In addition, these scaffolding modules also regulate the enzymatic activity of the protein by influencing, in direct but subtle ways, the atomic structure and dynamics of the catalytic domains.

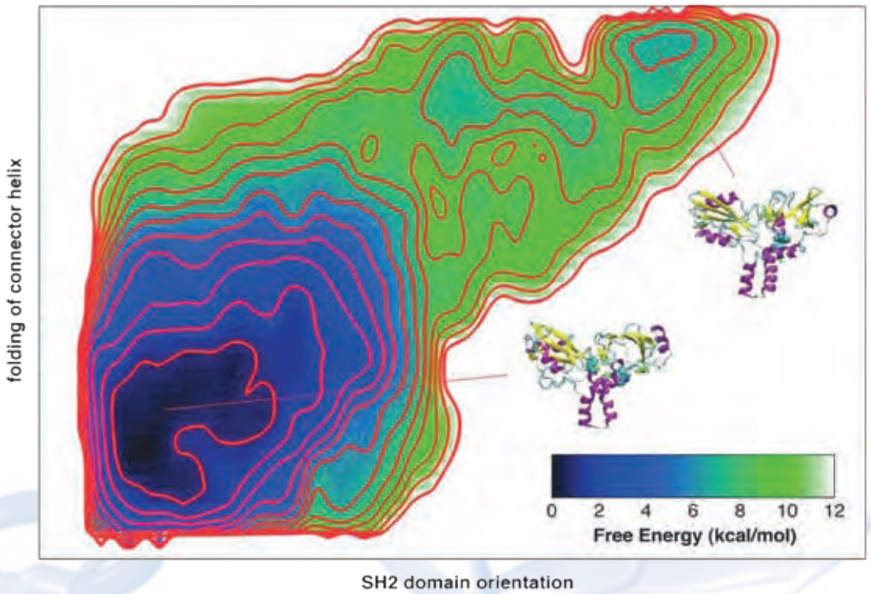

Figure 2. Free-energy surface of the SH2 tandem from ZAP-70, connecting the kinase- and ITAM-bound conformations. The surface was computed via umbrella-sampling molecular dynamics simulations of the tandem in explicit solvent.

Our understanding of these self-regulatory mechanisms is, however, largely sketchy – and such lack of understanding will hinder our ability to interfere or modulate these signaling pathways, e.g. in biomedical applications. In my group at the Max Planck Institute of Biophysics in Frankfurt am Main, we are employing large-scale molecular dynamics simulations to gain insights into these complex mechanisms of self-regulation. The figure below (Figure 1) depicts the structure of one of these enzymes, namely ZAP-70 – which Peter J. Bond, an EMBO postdoctoral fellow in my group, is currently studying. This NRTK from the Syk family contains two regulatory / scaffolding modules, referred to as SH2 domains. At present time, our work is focused on the analysis of the structural and dynamical properties of this tandem of SH2 domains. In particular, we try to elucidate the ways in which the tandem is influenced by its association with kinase domain, or with its scaffolding target (known as ITAM motifs) in the intracellular side of the T-cell membrane receptors. In practice, our approach is centered on the computation of the conformational free-energy landscape of the tandem (Figure 2). These computations entail hundreds of independent simulations, each of which is used to assess the response of the structure to a given conformational bias. This is akin to analyzing the stability of, say, a bridge, by modeling its response to the application of pressure in specific points. To optimize the realism of these molecular models, our simulations include tens of thousands of protein and solvent atoms. These free-energy calculations are therefore extremely demanding from a computational standpoint, but fortunately they are also highly parallelizable. Therefore, our research would not be possible without computational resources such as the HLRB II at LRZ.

• José Faraldo-Gómez

Max Planck Institute of Biophysics, Frankfurt am Main

Structural Relaxation at Interfaces of Transition Metal Oxides

Interfaces between different transition metal oxides attract considerable attention due to their technological and fundamental importance: electronic states at such interfaces can display physical properties which differ dramatically from those in the respective bulk materials. This is especially relevant for ultrathin perovskite oxide films grown on dielectric substrates. In the nm-thick film, the interface discontinuities due to the lattice mismatch with the substrate and the polarity of the termination planes induce a new type of a mixed electronic-lattice reconstruction near the interface which include: (1) the modification of lattice constants, (2) the local atomic displacements with consequent static electric polarization in the interfacial region, and (3) a "reconstruction" of the state of the electron subsystem in this region.

In several technologically important heterostructures, such a mixed reconstruction leads to novel phenomena which are not found in their bulk constituents. The promiment example is the metal-insulator transition in the heterostructure of $LaAlO_3$ (LAO) film and $SrTiO_3$ (STO) observed upon the increase of the film thickness [1]. The goal of our current work is the theoretical engineering of model superlattices which comprise key features of technologically important classes of heterostructures like $LaAlO_3$-$SrTiO_3$.

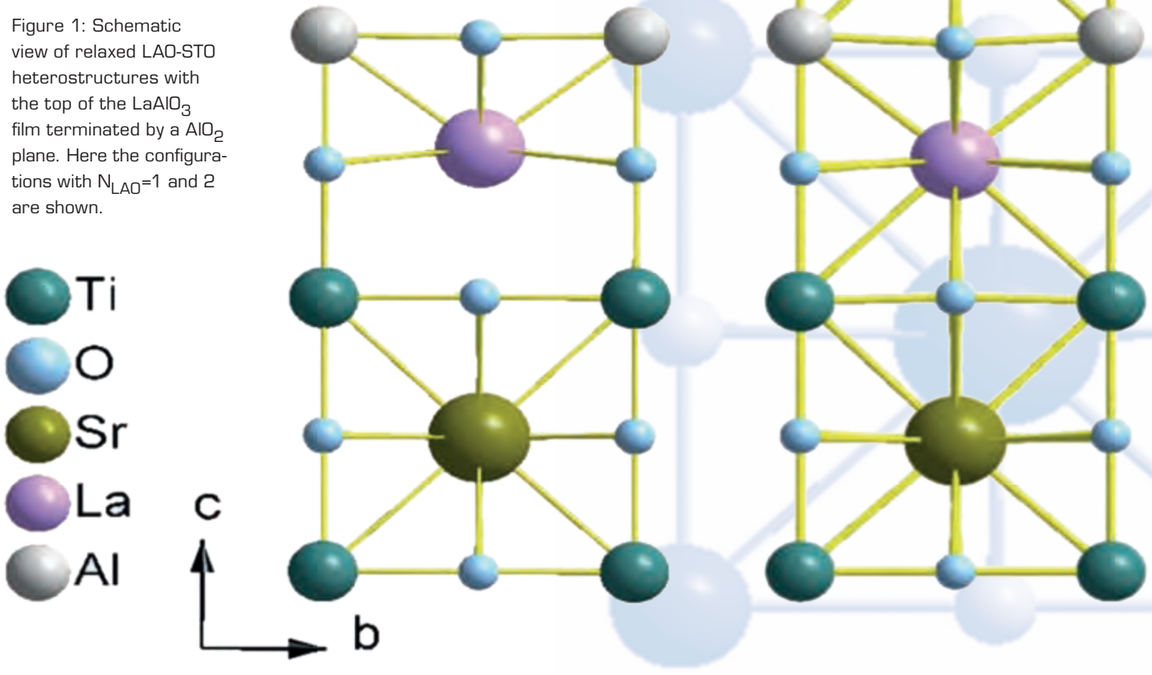

Figure 1: Schematic view of relaxed LAO-STO heterostructures with the top of the $LaAlO_3$ film terminated by a AlO_2 plane. Here the configurations with N_{LAO}=1 and 2 are shown.

For the studies of a $LaAlO_3$-$SrTiO_3$ bilayer, we generated a supercell containing a vacuum slab on the top of a bilayer containing several unit-cell-thick $SrTiO_3$ and $LaAlO_3$. The correct choice of model systems allows us to shed light on the nature of the mechanisms responsible for the appearance of new electronic properties in these compounds.

Description of the Computational Methods

Our analyses of novel electronic states are based on the density functional theory within LSDA and LSDA+*U* implementations [2]. As appears in our studies, the physical and chemical properties of unrelaxed or even partially relaxed systems strongly differ from the behavior observed in the experimentally prepared samples. As a consequence, the full procedure of structural and atomic relaxation is required for each theoretically engineered supercell which even for rather small supercells of 10-20 atoms is an extremely demanding computational task. In our project, the main amount of the computer time is used for high-performance theoretical engineering of multilayer oxide supercells and for structural optimization.

The main part of the resources was absorbed by computationally demanding structural optimization and LDA+*U* calculations of the optimized superlattices. Specifically, for superlattices of CuO_2-$SrTiO_3$ and $LaAlO_3$-$SrTiO_3$ containing from 18 to 33 atoms in a supercell, a single run of LDA+*U* included from 50 to 80 iteration steps until the final convergence with respect to the electronic charge and total energy could be reached. With 10-20 processors used for the parallel k-points calculations, each converged LDA+*U* run usually required from 10,000 to 35,000 CPU hours. In addition, in each superlattice, the relaxation of the interface distances could be achieved by performing up to 15 single LSDA runs which required about 10,000 CPU hours. Due to a wide range of superlattices analyzed in our project ($LaAlO_3$-$SrTiO_3$ sandwiches and bilayers, CuO_2-$SrTiO_3$ superlattices), such computationally demanding calculations required about 700,000 CPU hours in the first two years after beginning the project.

Structural Relaxation and Metal-insulator Transition at the Interface between $SrTiO_3$ and $LaAlO_3$

To study the bilayers STO-LAO with an AlO_2 plane at the top surface (Figure 1), we have generated supercells with the number of LAO unit cells (N_{LAO}) varying from 1 to 3. In these supercells, the structural optimization results in large opposite distortions of cations and anions in the LaO planes by values up to 0.3 Å and only small atomic displacements by about 0.05 Å in the STO substrate (see Figure 1). In experiments, STO-LAO bilayers with N_{LAO}=3 are still insulating whereas LAO films with $N_{LAO}\geq4$ have been found to be metallic. In the present LDA-calculations, due to the underestimation of the energy gap by LDA, the "critical thickness" of the LAO layer is decreased by 2 and the metallic state occurs already at N_{LAO}=2 [3]. We can conclude that the polarization in the LAO layer due to structural relaxation competes with the electric field caused by the polarity catastrophe and in this way establishes the insulating state in STO-LAO bilayers below the LAO critical thickness. In contrast, above the critical thickness the polarity compensation through the surface and interface charging leads to the metallic properties [3,4,5].

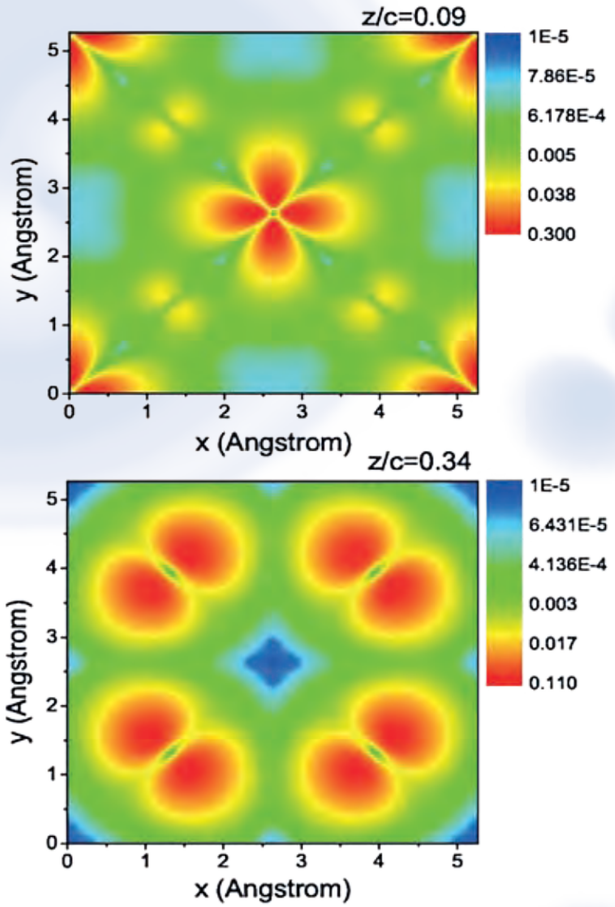

Figure 2: Electron density contours in the three unit-cell thick LAO film deposited on $SrTiO_3$. Here the system is in the metallic state. The maps show the electron density in the AlO_2 surface layer (z/c3=0.34) and in the interface TiO_2 layer (z/c3=0.09), c3=15.27 Å.

In the considered bilayers, two different charging processes can be distinguished. Near the positively charged LaO layer at the TiO_2-LaO interface, the compensation produces the additional negative charge accumulated at Ti $3d_{xy}$ 'orbitals. The high Ti $3d_{xy}$ electron density can be observed in the charge density contour of the interface TiO_2 plane in the metallic state (Figure 2, top panel with N_{LAO}=3). In contrast to this, in an insulating system with one unit-cell thick LAO film, the electrons in TiO_2 are located on the 2p valence orbitals of oxygens (Figure 2, bottom panel with N_{LAO}=1).

In distinction to the interface, the negative polarity near the surface AlO2 plane leads to additional holes distributed over the 2p orbitals of the surface oxygen. The charge density is calculated from the generation of charge in the energy range between -0.5 eV and 0.8 eV with respect to the Fermi level. In the metallic state, the charge distribution in Figure 2 (bottom panel) is spatially extended with substantial charge density in the regions between O 2p orbitals, whereas the Al ions in the center of the contour remain unoccupied by the conduction holes. In distinction to this, the charge distribution in the insulating state (Figure 3, bottom panel) has more localized character, typical for O 2p valence electrons.

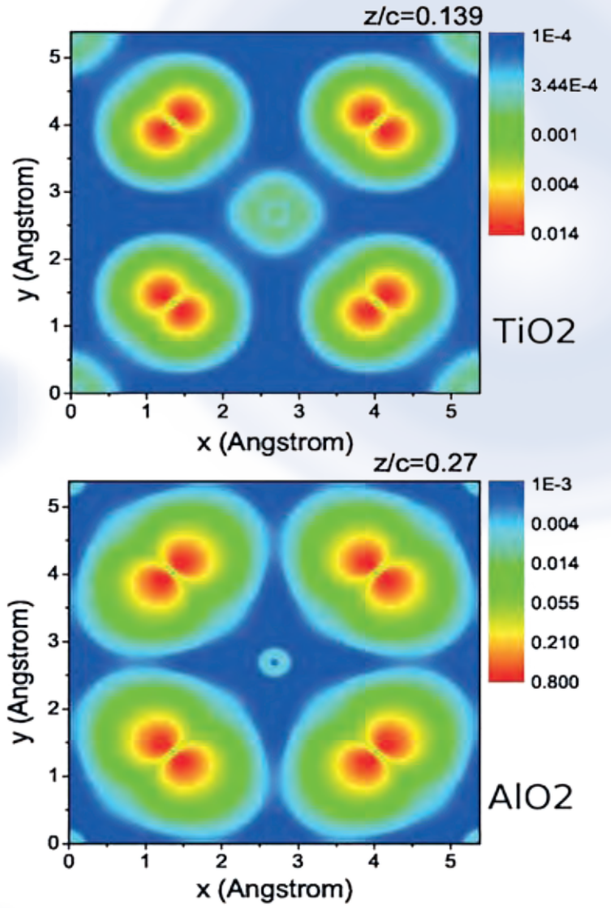

Figure 3: Electron density contours in a one unit-cell thick LAO film deposited on $SrTiO_3$. Here the system is in the insulating state. The maps show the electron density in the AlO_2 surface layer (z/c1=0.27) and in the interface TiO_2 layer (z/c1=0.139), c1=7.69 Å.

Recently, indirect indications for positive surface charge in STO-LAO have been discussed in [6] where such a positive surface charge is compensated by oxygen vacancies. In other experimental conditions, the surface hole charge can be also neutralized by surface oxidation or by near-surface bonding with negatively charged radicals which can be a possible reason for the insulating state usually found on the AlO_2 surfaces of LAO-STO bilayers.

References

[1] **Thiel, S., Hammerl, G., Schmehl, A., Schneider, C. W., Mannhart, J.**
Science 313, 1942, 2006

[2] **Blaha, P., Schwarz, K., Madsen, G. K. H., Kvasnicka, D., Luitz, J.**
WIEN2K, An Augmented Plane Wave and Local Orbitals Program for Calculating Crystal Properties, ISBN 3-9501031-1-2, TU Wien, Austria, 2001

[3] **Pavlenko, N., Kopp, T.**
arXiv: cond-mat/0901.4610, 2009

[4] **Pentcheva, R., Pickett, W. E.**
Phys. Rev. B 78, 205106, 2008

[5] **Pentcheva, R., Pickett, W. E.**
arXiv: cond-mat/0812.0550, 2008

[6] **Cen, C., Thiel, S., Hammerl, G., Schneider, C. W., Andersen, K. E., Hellberg, C. S., Mannhart, J., Levy J.**
Nature Materials 7, 298, 2008

• Natalia Pavlenko
• Thilo Kopp

Center for Electronic Correlations and Magnetism, Universität Augsburg

Innovative HPC Methods and Application to Highly Scalable Molecular Simulation

Molecular simulation allows entering regions beyond the domain of classical continuum approaches. For various application fields, like process engineering, material sciences, bio- and nanotechnology, atomistic scale simulations based on Molecular Dynamics (MD) and Monte-Carlo (MC) are regarded as a key to substantial further progress. Despite this well-known potential of molecular methods for industrial research and development, there is still limited usage due to the extreme computing time needed. On the other hand, computer hardware is continuously more powerful, still following Moore's law. In contrast to the frequency increase battle of the past, the actual processor development focuses on parallelism and manycore architectures are supposed to deliver a power-efficient performance gain.

Running a simulation on a large number of homogeneous (e.g. x86) or even heterogeneous (e.g. CellBE) cores, which are connected in a hierarchical structure over a CPU, a node or the whole system, requires an adaptation of the code to get a maximized performance. Parallelism also comes to fore on the core itself, supporting SIMD instructions or vectorization. This also has to be taken into account during code development. Massively parallel systems with a large number of processing elements also boost the importance of load balancing issues. Furthermore, software engineering aspects are important: Efficient code still needs to be readable and expandable to be open for future needs.

Project Goals

The Lehrstuhl für Thermodynamik (LTD) of Technische Universität Kaiserslautern, the Lehrstuhl für Thermodynamik und Energietechnik (ThEt) of Universität Paderborn, the Lehrstuhl für Informatik mit Schwerpunkt Wissenschaftliches Rechnen (SCCS) of Technische Universität München, the High Performance Computing Centre Stuttgart (HLRS) and the Fraunhofer Institut für Techno- und Wirtschaftsmathematik (ITWM) joined to meet the challenge of creating and improving highly scalable numerical molecular simulation software. The software packages will be optimized for massively parallel manycore systems featuring a hybrid OpenMP+MPI parallelization, but also addressing new architectures, like the Cell Broadband Engine. Dynamic load balancing algorithms for parallelization based on spatial decomposition are developed to deal with the very heterogeneous particle distributions which are typical for the application domain. These algorithms use e.g. kD-trees, space filling curves as well as graph based approaches to optimize the domain decomposition for equally distributed load, but also try to minimize communication. Concurrently software development based on Partitioned Global Address Space (PGAS) languages will re-implement already existing functionality and provide novel tools, which are a foundation for future work.

Simulation Software

In the center of the project are the two simulation packages called ms2 and ls1, which are the result of preliminary work of a subset of the IMEMO project partners, being enhanced during the project.

ms2 is a software package to get new insights into a large range of thermophysical properties of arbitrary rigid molecules mixtures. The program, initiated in 2001 at the Universität Stuttgart, was already successfully used for three Industrial Fluid Property Simulation Challenges (www.ifpsc.org), allowing for a title of champions in 2007 [1]. ms2 implements MD as well as MC simulation with several classical ensembles to determine static, entropic and transport properties. The reasonably object oriented Fortran90 code is parallelized for distributed memory architectures using MPI and with all relevant loops being vectorized, it is also suitable for vector machines.

ls1 is a MD simulation tool for time dependent nano-scale processes for arbitrary rigid molecule mixtures with a focus on large systems which aims to close the gap between the nanoscopic and the microscopic scale. The C++ code was initiated in 2005 at the Universität Stuttgart to study nucleation and flow phenomena and targets massively-parallel distributed memory architectures using MPI.

Applications

Some exemplary applications of these software packages for an atomistic numerical simulation will give a short overview about the direction of the ongoing work. Taking a set of molecular models for 78 real fluids and their binary mixtures as a basis [2,3], the fluid phase coexistence was studied for ternary mixtures using the ms2 program. The phase separation properties for the ternary system consisting of methane, ethane, and carbon dioxide are shown in Figure 1 for a temperature of 230 K and a pressure of 4.65 MPa, which is near the triple point temperature of carbon dioxide (217 K) and above the critical temperature of methane (191 K). Ethane and carbon dioxide were modeled using the quadrupolar two-center Lennard-Jones potential, which predicts the aceotropy for this subsystem [2], and methane was represented by the Lennard-Jones (LJ) potential. The agreement with experimental data [5] is excellent. A systematic comparison with respect to all 33 ternary mixtures for which experimental data are available confirms that this is generally the case [4].

Vapor to liquid nucleation for systems containing up to ten million particles was studied using the ls1 program. In a closed system, the nucleation process is instationary, since the vapor is continuously depleted by the aggregation of monomers to small droplets. This leads to a high uncertainty for the nucleation rate obtained from canonical ensemble MD simulation [6,7]. The precision can be significantly increased by a new molecular simulation method for

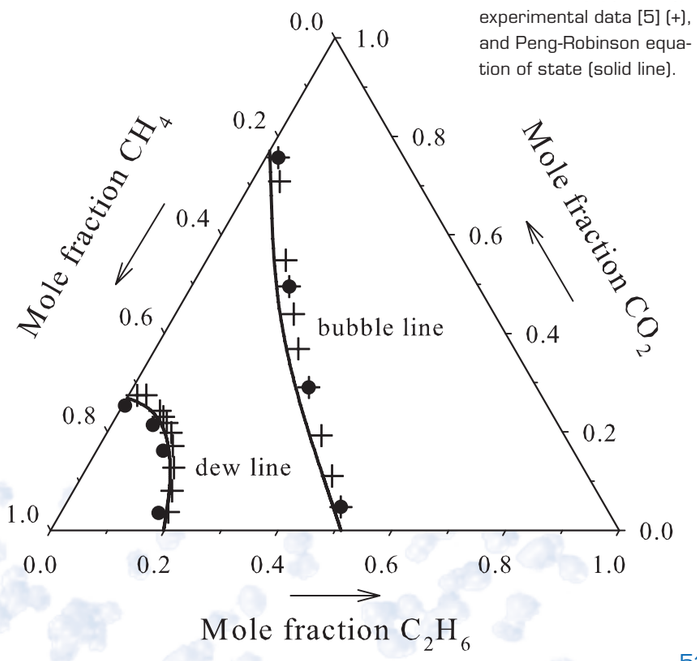

Figure 1: Ternary vapor-liquid equilibrium phase diagram for the mixture CH_4 + CO_2 + C_2H_6 at 230 K and 4.65 MPa: simulation data [4] (bullets), experimental data [5] (+), and Peng-Robinson equation of state (solid line).

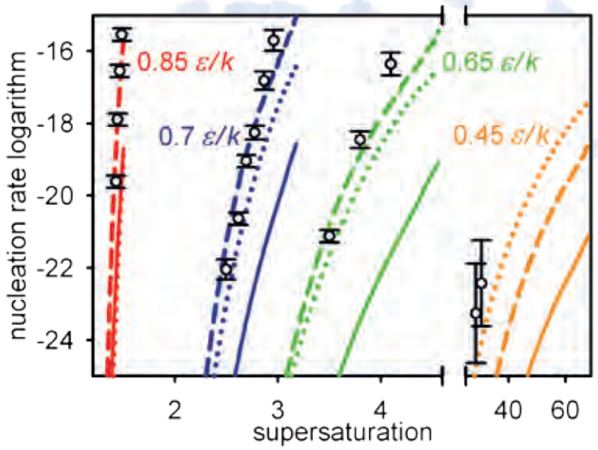

Figure 2: Nucleation rate over supersaturation for the t. s. LJ fluid at T = 0.45, 0.65, 0.7, and 0.85 ε/k_B, wherein ε is the energy parameter of the t. s. LJ model, according to CNT (solid line), the SPC modification [7] (dashed line), as well as the HCSF (dotted line) in comparison with present simulation results (circles)

Figure 3: Pressure drop δp in terms of the flow velocity v_z and the channel length δz (top) as well as slip velocity in terms of v_z (bottom), for Poiseuille flow of liquid methane at T =166.3 K as well as v_z = 10 (circles) and 30 m/s (bullets), as a function of the channel diameter, compared with Darcy's law (solid line)

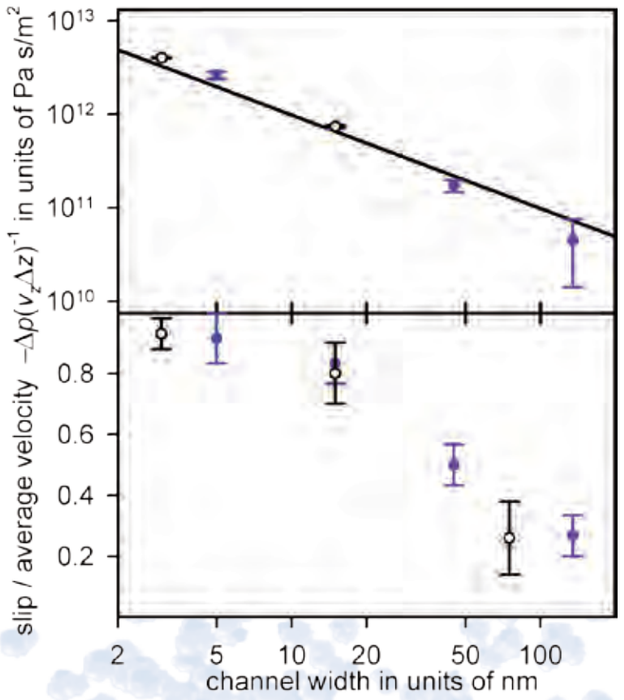

sampling both nucleation kinetics and steady-state properties of the nucleating supersaturated vapor. The idea behind this approach is to simulate the production of droplets up to a given size for a specified metastable state. This is achieved by extending MD simulation in the grand canonical ensemble with an "intelligent being" that continuously removes all large droplets: McDonald's demon [8]. A series of simulations for the truncated and shifted LJ (t. s. LJ) fluid, an accurate model for noble gases and methane, shows that the classical nucleation theory (CNT) underpredicts the nucleation rate by two orders of magnitude [Fig. 2]. The Hale critical scaling formalism (HCSF) [9], which was introduced to explain experimental data on low-temperature nucleation, leads to better agreement for the t. s. LJ fluid near and below the triple point temperature (T_3 =0.65 ε / k_B). This proves for low temperatures that the validity of HCSF extends up to extremely high nucleation rates which are experimentally inaccessible but can be studied by molecular simulation.

Due to their anisotropy, graphite nanochannels and carbon nanotubes are of particular interest. On the nanometer length scale, continuum approaches like the Navier-Stokes equation break down [10]. It is known for Poiseuille flow with a channel width below 2 nm that the slip velocity, which is often neglected, reaches almost the same magnitude as the flow velocity [11].

The ls1 program was applied to the Poiseuille flow of methane, represented by the t. s. LJ fluid, through a graphite channel, where a rescaled version of the Tersoff [12] potential was used for the carbon atoms. The interaction between fluid molecules and wall atoms was also modeled by the LJ potential [13].

As figure 3 shows, a transition from high to low slip velocities occurs at a channel diameter of about 40 nm. Darcy's law, stating that the pressure drop is inversely proportional to the cross-sectional area of the channel, is valid down to the molecular length scale.

Acknowledgement

IMEMO is a project supported by the German Federal Ministry of Education and Research (BMBF) within the program "IKT 2020 - Forschung für Innovationen" in the call "HPC-Software für skalierbare Parallelrechner". The project runs for 36 months and started on 1st of October 2009. The consortium consists of five project partners.

Partners

High Performance Computing Centre Stuttgart (HLRS) | Lehrstuhl für Thermodynamik und Energietechnik (ThEt), Universität Paderborn | Lehrstuhl für Thermodynamik (LTD), Technische Universität Kaiserslautern | Lehrstuhl für Informatik mit Schwerpunkt Wissenschaftliches Rechnen (SCCS), Technische Universität München | Institut für Techno- und Wirtschaftsmathematik (ITWM), Fraunhofer Gesellschaft, Kaiserslautern

in collaboration with

BASF SE | Bayer Technology Services GmbH (BTS) | Evonik Industries AG | IBM Deutschland Entwicklung GmbH

References

[1] Eckl, B., Vrabec, J., Hasse, H.
On the Application of Force Fields for Predicting a Wide Variety of Properties: Ethylene Oxide as an Example, Fluid Phase Equilibria 274: 16-26, 2008

[2] Stoll, J., Vrabec, J., Hasse, H.
Vapor-liquid Equilibria of Mixtures Containing Nitrogen, Oxygen, Carbon Dioxide, and Ethane, AIChE Journal 43: 212, 2003

[3] Vrabec, J., Huang, Y.-L., Hasse, H.
Molecular Models for 267 Binary Mixtures Validated by Vapor-liquid Equilibria: a Systematic Approach, Fluid Phase Equilibria 279: 120, 2009

[4] Huang, Y.-L., Vrabec, J., Hasse, H.
Prediction of Ternary Vapor-liquid Equilibria for 33 Systems by Molecular Simulation, Journal of Physical Chemistry B, submitted 2009

[5] Knapp, H., Yang, X., Zhang, Z.
Vapor-liquid Equilibria in Ternary Mixtures Containing Nitrogen, Methane, Ethane and Carbondioxide at Low Temperatures and High Pressures, Fluid Phase Equilibria 54: 1, 1990

[6] Yasuoka, K., Matsumoto, M.
Molecular Dynamics of Homogeneous Nucleation in the Vapor Phase, The Journal of Chemical Physics 109: 8463, 1998

[7] Horsch, M., Vrabec, J., Hasse, H.
Modification of the Classical Nucleation Theory Based on Molecular Simulation Data for Surface Tension, Critical Nucleus Size and Nucleation Rate, Physical Review E 78: 011603, 2008

[8] McDonald, J. E.
Homogeneous Nucleation of Vapor Condensation. II. Kinetic aspects, American Journal of Physics 31: 31, 1962

[9] Hale, B. N.
Application of a Scaled Homogeneous Nucleation-rate Formalism to Experimental Data at T << Tc, Physical Review A 33: 4156, 1986

[10] Karniadakis, G., Beskok, A., Aluru, N.,
Microflows and Nanoflows: Fundamentals and Simulation, Springer, New York, 2005

[11] Sokhan, V. P., Nicholson, D., Quirke, N.
Fluid Flow in Nanopores: Accurate Boundary Conditions for Carbon Nanotubes, The Journal of Chemical Physics 117: 8531, 2002

[12] Tersoff, J.
Modeling Solid-state Chemistry: Interatomic Potentials for Multicomponent Systems, Physical Review B 39: 5566, 1989

[13] Battezzati, L., Pisani, C., Ricca, F.
Equilibrium Conformation and Surface Motion of Hydrocarbon Molecules Physisorbed on Graphit, Journal of the Chemical Society: Faraday Transactions 2, 71: 1629, 1975

- Martin Bernreuther[1]
- Christoph Niethammer[1]
- Martin Horsch[2]
- Jadran Vrabec[2]
- Stephan Deublein[3]
- Hans Hasse[3]
- Martin Buchholz[4]

[1] University of Stuttgart, HLRS

[2] Universität Paderborn, ThEt

[3] Technische Universität Kaiserslautern, LTD

[4] Technische Universität München, SCCS

Projects

Visualization in Parallel Manycore Environments

Foreseeable changes in the development of hardware architectures in the domain of high performance computing allow for considerable increase of the complexity of scientific simulations. Single workstations are replaced by visualization clusters to provide adequate performance for post processing and visualization of huge data sets. These clusters comprise multiple nodes that are connected by a high bandwidth, low latency network infrastructure such as Infiniband. The nodes themselves are parallel systems, composed of multiple processors containing multiple cores. These cores have access to shared memory, often with non-uniform latency and bandwidth.

Each of these nodes contains one or more graphics cards that can be used for rendering. Modern graphics cards can also be used for highly complex parallel computations that are not directly related to rendering. Utilizing these graphics cards for postprocessing introduces additional latency due to the need for copying data between the host main memory and the comparably small amount of memory available on the graphics board. To make this infrastructure available for interactive visualization to the users, they have to be able to access the cluster remotely from their workstation which is typically connected by a low bandwidth, high latency network. For maximum utilization of the visualization infrastructure by a visualization environment based on the data-flow paradigm, several requirements have to be met. The data management has to be adapted to the distributed memory model with hierarchical memory access. The algorithms have to be optimized for parallel hybrid distributed/shared memory architectures and have to be integrated with the data management. The algorithms have to be adapted to make use of GPUs and other accelerators such as FPGAs. The rendering has to be able to utilize several GPUs to drive immersive environments such as CAVEs and provide remote rendering services with low latency and high detail. Due to the complexity of the post processing, an automatic scheduling of the required processes has to be provided. To make optimal use of the available resources and to assure tolerable latency for user interaction, the scheduling algorithms have to take into account load balancing, a given decomposition of the data and the amount of communication between the compute nodes.

The High Performance Computing Centre Stuttgart (HLRS) is the leader of a consortium of research institutes and industry collaborating in this research project that is partially funded by the German Federal Ministry of Education and Research (BMBF)

- to create a flexible framework for parallel and distributed visualization that is applicable throughout the different fields

- to develop scheduling strategies that are able to re-order the processing of data and dynamically adapt the accuracy of algorithms and resolution of the data

- to provide algorithms for visualization of data sets from different application areas in distributed manycore environments

- to assure the applicability of the developed concepts and the usability of the visualization environments with the help of users from different domains

Parallelization of Algorithms for Clusters of Manycore Systems

For the first time, the availability of massively parallel manycore architectures – both within a visualization system and as remote computing resources – allows for the interactive utilization of post-processing algorithms such as particle tracing even for extremely complex phenomena. One goal of this project is the development of methods for highly-interactive particle tracing in complex flow fields and their implementation in a software framework in order to offer engineers a tool for the intuitive exploration of flow phenomena. Besides offering support for structured cartesian grids, this requires all developed approaches to work with unstructured, potentially time-varying grids as well.

Realtime interactive postprocessing requires extremely low latency which can be achieved by employing massively parallel stream processing units available on the latest generation graphics boards. To be able to use GPUs efficently, partitioning and resampling of the data is necessary due to the limited amount of memory available. To be able to utilize remote visualization clusters for rendering of polygonal and volume data, lossy as well as lossless low-latency compression schemes running directly on the GPU have to be developed, possibly exploiting intra-eye coherence in stereo image streams.

Outlook

The visualization systems COVISE and ViSTA FlowLib that are enhanced and extended as a result of this project are available to the scientific community. The partners involved are collaborating with researchers and industry of different application domains to interpret large, transient data obtained from simulation and measurements. As a result of this project, tools will be made available to foster the understanding of processes in the domain of engineering (e.g. reduction of CO_2 emmisions in coal combustion) and medicine (e.g. formation of edema due to apoplexia) which, due to their complexity, cannot be intuitively analyzed by the means of interactive visualization so far.

Facts

VisPME is a project partially funded by the German Federal Ministry of Education and Research (BMBF) within the program “IKT 2020 – Forschung für Innovationen” in the call “HPC-Software für skalierbare Parallelrechner”. The consortium consists of seven project partners. The project runs for 36 months and started on 1st of January 2009.

Partners

High Performance Computing Centre Stuttgart (HLRS) | Regional Computing Centre of the University of Cologne | Center for Computing and Communication of RWTH Aachen University | Max Planck Institute for Neurological Research | Aachen Institute for Computational Engineering Science | RECOM Services GmbH | NVIDIA GmbH

• Florian Niebling

University of Stuttgart, HLRS

STEDG – Highly Efficient, Scalable Software for the Simulation of Turbulent Fluid Flows in Complex Geometries

Numerical simulation of fluid flows is nowadays an indispensable method for research and development in all areas of engineering. Over the years, it became a key technology to improve economic, environmental and security behaviour of products, thereby vitaly strengthening the competitiveness of German industry. Nevertheless, it often uses simplified methods to obtain results within the given turn-around times. The simulation of turbulent flows is often based on Reynolds-averaged Navier-Stokes equations (RANS). Here, the algorithms give stationary solutions, averaged in time. On the other hand, time-dependent phenomena often play an important role in industrial applications. Research and development departments carry out first calculations, based e.g. on Large-Eddy Simulations (LES). But the commercial software packages used work with numerical methods suited for traditional stationary simulations with turbulence model. They have a broad application spectrum, but are much to costly for instationary processes.

The efficient simulation of instationary fluid flows with higher order turbulence modelling requires an overall renewing of all components: numerical methods, algorithms and their efficient implementation on modern hardware architectures. While several groups work intensively on new numerical methods, the changes in hardware of future computer generations remain mostly unconsidered. Current developments in hardware architecture of high performance computers require dramatic changes in the programming model of simulation codes to be able to make use of these new architectures.

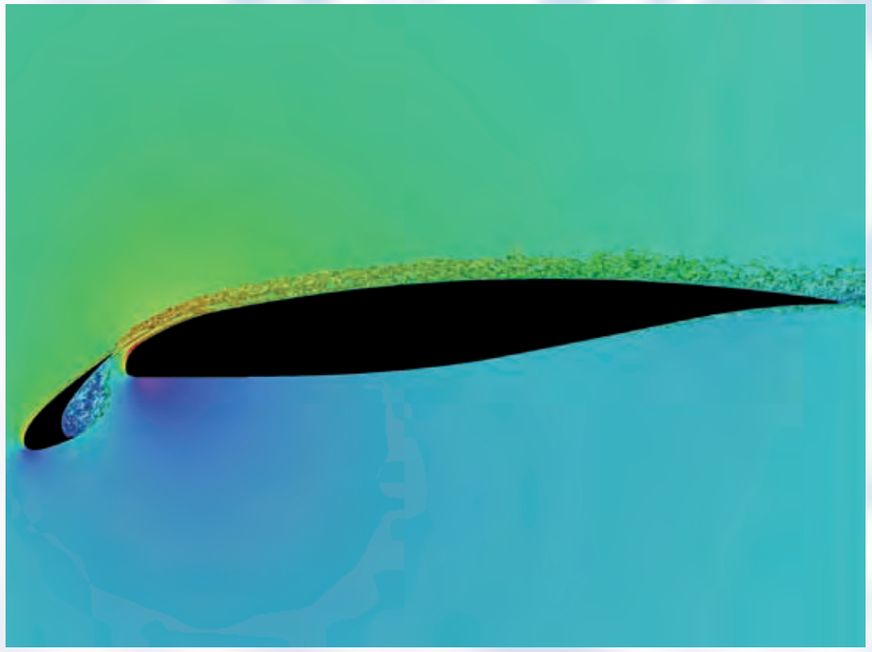

Figure 1: Turbulent flow field around a high lift configuration. Plot shows Mach number and vortex structures in shear layers.

Necessary is a paradigm change in the comprehension of the users: in the past, users could expect an increase of performance of their simulations directly with the performance of the hardware. In the future, this is not directly the case. The power and performance of computer systems is no longer increasing by higher clock frequencies of single processors, but through an more and more increasing number of processors – with respect to increasing number of nodes as well as increasing number of cores within one node. Future architectures will be systems of shared memory nodes, what is currently hardly used by application codes. Also the increasing gap between processor speed to memory bandwidth is hardly ever a topic discussed at user conferences.

The STEDG project adresses the changes, software development has to catch up to keep track with hardware development. Current new numerical methods – by structure well suited for these computer architectures – are further developed to efficient flow simulation tools. On the one hand, future numerical methods and algorithms, on the other hand their mapping to current (multi-core) and future (many-core) hardware architectures are considered. They are applied to real-life applications from industrial research and development (R&D) and their usability for industrial development cycles correlated with current and future developments.

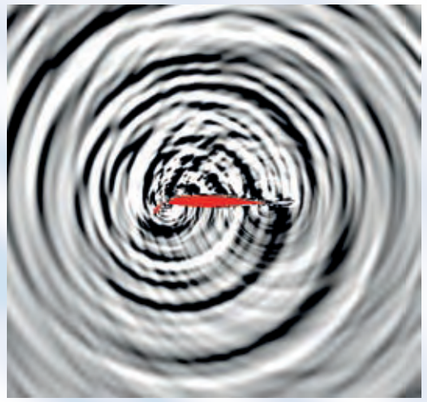

Figure 2: Sound emittance of a wing profile in high lift configuration. Gray scales show the amplitude of acoustic waves.

Facts

The STEDG project funded and supported by the German Federal Ministry for Education and Research (Bundesministerium für Bildung und Forschung, BMBF) in the HPC Software Initiative call “HPC-Software für skalierbare Parallelrechner”, develops highly efficient, scalable simulation tools for industry-relevant turbulent flows in complex geometries.

Partners

The project is lead by HLRS with partners from academy, Institute for Aerodynamics and Gasdynamics (IAG), University of Stuttgart, and Aerodynamics Institute (AIA), RWTH Aachen, as well as industry, Robert Bosch GmbH, Stuttgart, and Trumpf Werkzeugmaschinen GmbH + Co. KG, Ditzingen. The associated partner EADS Deutschland GmbH, EADS Innovation Works, München, supports the project.

• Sabine Roller

University of Stuttgart, HLRS

Projects

ScaFaCoS – when Long Range Goes Parallel

• Godehard Sutmann

Forschungszentrum Jülich, Jülich Supercomputing Centre

What makes difference between an ideal gas and real matter is that there are interactions between particles. Due to a variety of types of interactions, substances might have completely different structural or dynamical properties. It is the interplay between attraction and repulsion between particles which makes the world so diverse in nature. Taking the viewpoint of classical mechanics, matter is composed of particles, described as force-centers, which form a potential energy landscape. Considering the work, which is needed to bring a particle to a given location on this map, characterizes its energy. The force onto this particle is then simply given by the gradient at the particles' position, which is responsible that it moves away from a given position, thus changing the energy landscape and inducing a collective response of the particle system. This simplified description is taken as a basis for various types of classical simulation techniques, which are based on so called force-fields, which are nothing else than analytical approximations to various types of interactions between particles. These interactions are usually classified on the one hand into bonded and non-bonded and on the other hand into short- and long-range interactions. Although bonded interactions, like stretching, bending or torsional deformations of bonds between sites, define what is classically called a molecule, the non-bonded interactions are responsible for the dynamical interplay between atoms and molecules, inducing collective behavior. In the extreme case of astrophysical objects, like planets, stars or galaxies it is the long-range nature of gravitation, which is responsible for the dynamics and large scale structure formation.

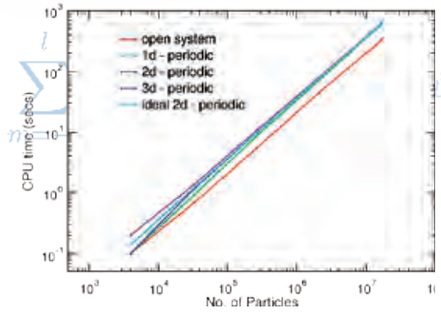

Figure 1: Scaling behavior of the FMM method with the number of particles for different geometries. As can be seen, the method has complexity O(N) for 1D-, 2D- and 3D-periodic systems as well as for systems with open boundaries.

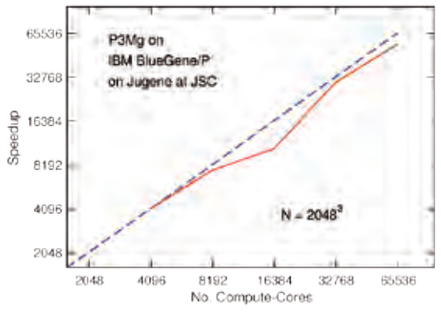

Figure 2: Parallel scaling on the Jülich IBM BlueGene/P for the long range part of the P3Mg method (Particle-Particle Particle-Multigrid), which was developed at JSC.

Relying on the framework of classical mechanics and knowing the interactions between objects, it is straight forward to set up Molecular Dynamics (MD) or Monte Carlo (MC) simulations to study these systems. With the advent of large scale parallel computers, like the IBM BlueGene/P in Jülich, it is nowadays in principle possible to simulate systems of billions of particles. Although this is still far beyond Avogadros number, it brings simulations closer to length scales of laboratory experiments. However, if long range interactions play an important role for systems of interest, like gravitational systems or non-equilibrium Coulomb interaction driven plasmas, it would be necessary to take them properly into account. Both, gravitational and electrostatic potentials are the solution of the linear Poisson equation, which means that the total solution of a many-body system is built up of a linear superposition of individual contributions. This means, however, that for particle systems all pair-contributions between particles have to be taken into account. Since in a system, composed of N objects, there are N(N-1)/2 particle pairs, this implies a computational complexity of $O(N^2)$.

Various methods were developed to overcome this computational bottleneck, which all have in common to split the computing domain into short-range and long-range contributions. The short-range part, which has on average a constant number of neighbors (and therefore is O(N)) is calculated in a pair-wise fashion, whereas the long-range part is subject to different kinds of techniques, depending on the algorithm. In periodic systems, where so called Ewald-summation techniques, like the Particle-Particle Particle-Mesh method (P3M), are applied, fast Fourier transform techniques (FFT) may be used to reduce the computational work to O(Nlog(N)). Other techniques, like the Fast Multipole Method (FMM, O(N)) or Barnes-Hut tree methods (O(Nlog(N))) group far away particles together to pseudo-particles or multipole expansions, in order to reduce the complexity, induced by the number of pair contributions. Finally, multigrid techniques, which have complexity O(N), use different grid hierarchies, on which the Poisson equation is solved subject to charges, which are transferred from the particle positions to equidistant grid nodes. In solving the Poisson equation this method has the advantage that it can also take into account a prescribed dielectric background.

All these methods have in common that they solve the Coulomb problem for a prescribed error tolerance very efficiently, but that the implementation into a computer program may be very demanding. This is even more a fact, if the target platform is a massively parallel computer and that the program should run and scale on up to thousands of processors. Within the Call "HPC Software für skalierbare Parallelrechner" (HPC software for scalable parallel computers) of the German Ministery for Education and Science (BMBF), the grant proposal ScaFaCoS (Scalable Fast Coulomb Solvers) from Jülich Supercomputing Centre (JSC) was succesful, which aims to concentrate the most promising methods for solving the Coulomb problem in classical atomistic systems, in a parallel library which may be linked to existing programs. This will on the one hand strongly facilitate code development, since programs can easily be extended to efficient, but sophisticated algorithms. On the other hand it allows to couple these methods to programs, which have already well developed parallel methods, but lack some functionalities, e.g. Coulomb solvers for 1D- or 2D-periodic systems, high accurate solvers or solvers for systems embedded in a dielectric continuum.

Within the German network project ScaFaCoS, formed by University groups from Bonn, Chemnitz, Stuttgart and Wuppertal, by Research Centre groups from JSC at Research Centre Jülich, Fraunhofer Institute SCAI St. Augustin and Max Planck Institute for Polymer Research Mainz, as by industrial partners from BASF, Cognis and IBM, physicists, chemists, mathematicians and computer scientists will work in common to bring together their expertise to realize a parallel Open-Source library, which will run on target platforms like the Jülich IBM BlueGene/P to help the scientific research community to satisfy their needs.

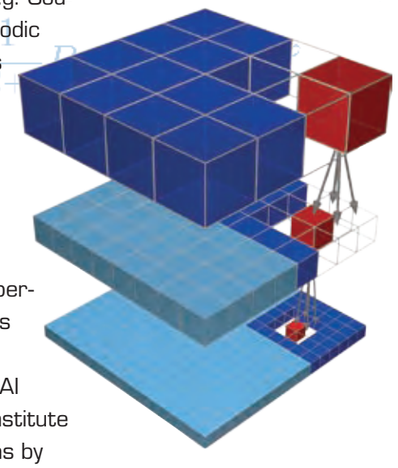

Figure 3: Schematic of the hierarchical calculation of the long range interactions in the fast multipole method (FMM), which works on different levels of coarse scales to sum up energy contributions in O(N) complexity (cmp. Fig. 1).

Hermes – an Evacuation Assistant for Mass Events

BMBF Research Programme

As part of the German Government's high-tech strategy, the Federal Ministry of Education and Research (BMBF) has launched a programme on "Research for Civil Security". The field of "Protecting and Saving Human Life" in this programme includes funding for the Hermes Project. The aim of the collaborative project is to improve safety for people in large multifunctional buildings and also at big events by exploring the effectiveness of an evacuation assistant. In this project, the Institute for Advanced Simulation (IAS) of the Forschungszentrum Jülich is producing a real-time simulation for the reliable prediction of the course of an evacuation of a crowd of more than 60,000 people, combining mathematical models of pedestrian dynamics with a parallel computer.

Motivation and Concept

The trend towards large multifunctional buildings and also the dimensions of public events today make new demands on the quality of security concepts. In case of emergency, all attendees present must be able to evacuate the danger area rapidly. Although this is generally ensured by the application of building regulations, in case of overcrowding or if some of the emergency exits are blocked this can result in a dangerous crush and long queues.

Figure 1: Concert in a multifunctional area (Source: LTU arena)

In order to ensure the safety of the visitors it is therefore necessary to develop new strategies and tools capable of determining the number of persons present and also providing support for optimum evacuation of a danger area. The evacuation assistant couples the simulation core to the fire alarm system installed in the building and is also able to count the number of persons by means of automated image processing thus making it possible for the first time to obtain simulation results based on the actual risk situation.

The LTU arena in Düsseldorf will provide a venue for testing the evacuation assistant. The example of this multifunctional area with a capacity of 66,000 spectators will show how crowds of people at big events can be guided – also considering the actual risk situation – so that optimum use can be made of the emergency exits. A test system of the evacuation assistant will be installed in the LTU arena in 2011.

Modelling and Parallelization

The aim of the Hermes Project is to produce a real-time simulation for a reliable forecast of the evacuation behaviour. This forecast is based on "microscopic" models of pedestrian dynamics which calculate the individual movement of the people through the building by means of a system of coupled differential equations. These models resemble the approach used in molecular dynamics. However, pedestrians are represented as self-driven particles in a continuous space, in which they interact with each other and with the spatial geometry (walls, doors, obstacles). Previous research in this field has concentrated on the qualitative reproduction of self-organization phenomena such as queuing at bottlenecks or formation of lanes in bidirectional flows of people and have only been validated in a rudimentary manner. Apart from the above-mentioned modelling approaches, the parallelization concepts normally applied in molecular dynamics must be adapted.

The task of IAS is, firstly, to enhance previous approaches for modelling pedestrian dynamics in such a way that flows of people can be described in both a qualitatively and also quantitatively correct manner in any configuration (corridors, stairs, bottlenecks). In cooperation with the Universities of Wuppertal and Cologne, validation experiments will be carried out under laboratory conditions with up to 300 people. Secondly, IAS will undertake the efficient implementation of the models thus obtained. The runtime demands made on forecasts for evacuating more than 60,000 people in the complex spatial geometry of a building mean that the application of state-of-the-art processor technology such as system-on-chip (SoC) and floating point accelerator is absolutely indispensable. The use of Cell processors promises to speed up the simulations. However, the load balancing concepts from molecular dynamics need to be expanded and adapted to the cell architecture. The parallel application of several such systems will ultimately make it possible to automatically estimate the reliability of the forecasts. Furthermore, the simulation will be fine-tuned on the basis of the real person count.

Figure 2: Schematic diagram of the evacuation assistant

Partners Cooperating in the Hermes Project

- Forschungszentrum Jülich GmbH, Institute for Advanced Simulation (IAS)
- Bergische Universität Wuppertal, Abteilung Bauingenieurwesen, Lehrstuhl für Baustofftechnologie und Brandschutz
- Imtech Deutschland GmbH & Co. KG
- Multifunktionsarena Immobiliengesellschaft mbH & Co. KG (LTU arena)
- PTV Planung Transport Verkehr AG
- TraffGO HT GmbH
- Rheinische Friedrich-Wilhelms-Universität Bonn, Geographisches Institut
- Universität zu Köln, Institut für Theoretische Physik
- Vitracom AG

Associated Partners

- Feuerwehr Düsseldorf
- Landesamt für Zentrale Polizeiliche Dienste NRW und Polizeipräsidium Düsseldorf
- Special Security Service SSSD GmbH

• Stefan Holl
• Armin Seyfried

Institute for Advanced Simulation, Forschungszentrum Jülich

Spring 2009 • Vol. 7 No. 1 • inSiDE 61

New HPC Systems at HLRS

HLRS has decided to upgrade its existing hardware infrastructure in 2008. The systems became operational in April 2008. HLRS both increased its vector performance and its scalar performance. The total performance of systems available now is about 200 TFLOP/s. These two upgrades were decided late in 2008 with the support of the federal government and the ministry of science of the state of Baden-Württemberg.

HLRS first decided to enrich its hardware environment by bringing in a new type of MPP systems. The newly introduced Cray XT5m system is the first installation of Cray (see also page 76 of this inSiDE) and is part of a new collaboration between HLRS and Cray.

The XT5m comes with 896 AMD Shanghai cores. Its peak performance is about 9 TFLOP/s and the main memory is 1.8 TB. The Cray XT5m interconnect is the proprietary SeaStar network. It allows for a closer coupling of the processors. Thus a sustained performance in the range of 1-2 TFLOP/s can be expected.

This system is especially interesting for engineering applications and is currently also tested for industrial applications. The prize-performance ratio and the low total cost of ownership will make the Cray XT5m an attractive system for industrial customers of HLRS. At the same time HLRS and its users have a chance to explore in detail the potential of this new architectural concept.

The second change in architecture is the update of the NEC SX-8 system. HLRS already installed 8 nodes of a new SX-9 system as part of the German D-Grid project.

To this 4 additional nodes of NEC SX-9 were added. In total the vector performance at HLRS has increased to 19.2 TFLOP/s. The estimated sustained performance is about 6-9 TFLOP/s. Considering the growing importance of clusters HLRS has added an NEC Intel cluster. The system has 700 nodes with a total of 1400 Intel Nehalem processor. To this HLRS has added 64 NVIDIA graphics cards. Overall main memory size is 9 TB. Overall peak performance is 133 TFLOP/s. The cluster interconnect network is Infiniband.

New Systems at HLRS			
	NEC SX-9	NEC Intel Cluster	Cray XT5m
Proc. Type:	Vector	Intel Nehalem	AMD Shanghai
Interconnect:	IXS	Infiniband	SeaStar
# cores:	192	5600	896
Total Performance [TFLOP/s]:	19.2	133	9
Total Memory [TB]:	8	9	1.8

The expected sustained performance is in the range of 6-10 TFLOP/s.

The vector system and the cluster system will be closely coupled. This will allow running coupled simulations on the overall system. Aeroacoustics and fluid-structure interaction simulations will substantially benefit from this.

With these two new systems HLRS can provide its users with competitive performance for all fields of applications. At the same time users have a chance to evaluate various architectures. This will help HLRS to make a better hardware decision for its next big investment step planned for 2011 and 2013.

JUROPA – Jülich Research on Petaflop Architectures

Following a dual track systems strategy, Jülich has ever since aimed at complementing their highly specialized leadership-class systems with general-purpose supercomputers, enabling to tackle scientific applications from many different areas. Lately, this strategy lead to the development of an innovative 200 Teraflops general-purpose supercomputer system named JUROPA. To develop the next generation of general purpose cluster computing as a research activity it is necessary to bring together the most experienced companies and research site to combine and integrate the best processor, server, interconnect, cluster software and I/O and storage technology.

Partners in this European co-development project are Forschungszentrum Jülich/JSC, Bull, Intel, Mellanox, Novell, ParTec and Sun: Intel, Mellanox and ParTec will cooperate together with Forschungszentrum Jülich to establish a cluster-architecture based on Intel technology, Mellanox interconnects and Parastation Cluster middleware. As integrator Bull was chosen to implement the new system. Sun contributes with the further development and the optimization of the SUN Lustre parallel file system with end-to-end data integrity and supports the optimization of the ParaStation communication and management software stack on the SUN-IB-communication network.

Applications on JUROPA originate from various disciplines like chemistry, many particle physics, elementary particle physics, biology/biophysics, material science, soft matter and others. One important new source of applications is the European fusion community (EFDA – European Fusion Development Agreement) who became interested in the JUROPA concept and signed a contract on using a 100 Teraflops JUROPA cluster of their own named HPC-FF. HPC-FF follows the same design principles as JUROPA and will be operated by Jülich Supercomputing Centre. HPC-FF and JUROPA are closely coupled to each other through a common high-bandwidth/low-latency network. Both systems can be operated as one big 300 Teraflops system.

JUROPA and HPC-FF system characteristics

	JUROPA	HPC-FF
Processor:	Xeon X5570 (Nehalem-EP)	Xeon X5570 (Nehalem-EP)
Processor clock speed:	2.93 GHz	2.93 GHz
Processors per node:	2	2
Cores per node:	8	8
Logical CPUs (SMT on):	16	16
Memory per node:	24 GB	24 GB
Node technology:	Sun Blade SB6048	Bull NovaScale R422-E2
Total number of nodes:	2,208	1,080
Total peak performance:	207 Teraflop/s	101 Teraflop/s

JUROPA is built up from innovative components combined into a “best-of-breed” HPC cluster. Intel Xeon X5570 processors (code name: Nehalem-EP) form the computational basis of the cluster nodes (Table 1). Each node comprises two of them, providing 8 CPU cores per node attached to 24 GB of main memory. Running at 2.93 GHz clock speed, each node reaches a peak performance of approximately 94 Gigaflops. 2,208 compute nodes yield a total peak performance of 207 Teraflops.

The nodes are connected via Infiniband/QDR with fat-tree topology. First performance measurements show network latencies around 1.5 µs on the MPI level and a sustained bandwidth of over 2,800 MB/s per link and direction. It is expected that future optimizations will further improve these results.

JUROPA, consisting of SUN racks

JUROPA and HPC-FF share a common Lustre storage pool providing a total of 850 TB usable disk space. The Lustre storage pool will be used for a fast scratch file system striped across all available disks as well as for moderately striped home file systems. Access to data residing on Jülich’s storage cluster JUST running GPFS will be provided through GPFS gateway nodes. They act as a bridge between the Lustre file server running over Infiniband and the GPFS file server running over 10-Gigabit-Ethernet. A total of 12 Login nodes form the interface to the JUROPA users. They will allow for interactive program development and testing as well as post processing.

References

http://www.fz-juelich.de/jsc/juropa/

HPC-FF, consisting of Bull racks

• Ulrich Detert
• Klaus Wolkersdorfer

Forschungszentrum Jülich, Jülich Supercomputing Centre

Leibniz Supercomputing Centre of the Bavarian Academy of Sciences (Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften, LRZ) in Garching near Munich provides national, regional and local HPC services.

Each platform described below is documented on the LRZ WWW server; please choose the appropriate link from www.lrz.de/services/compute

Contact
Leibniz-Rechenzentrum

Dr. Horst-Dieter Steinhöfer
Boltzmannstraße 1
85748 Garching/München
Germany
Phone +49-89-3 58 31-87 79
www.lrz.de

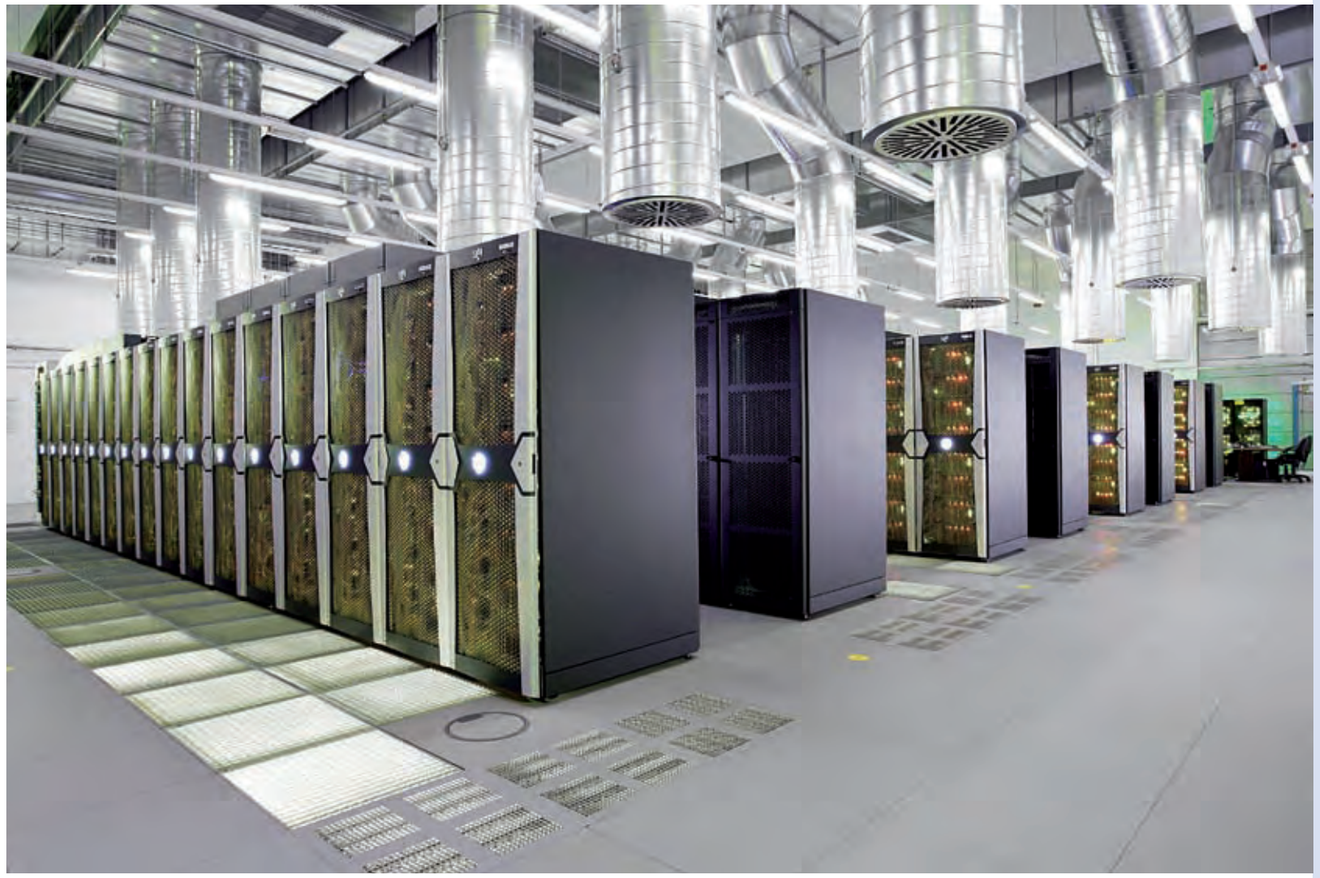

View of „Höchstleistungsrechner in Bayern HLRB II", an SGI Altix 4700
Foto: Kai Hamann, produced by gsiCom

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
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

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 +49-711-685-872 69
resch@hlrs.de
www.hlrs.de

View of the NEC SX-8 at 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	152	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/hw-access

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

The IBM supercomputer "JUGENE" 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

The 8th HLRS/hww Workshop on Scalable Global Parallel File-Systems, HNF Europe Spring Meeting 2009 and OpenFabrics

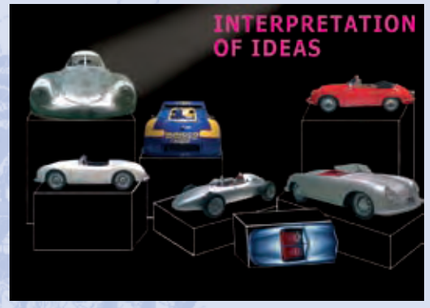

Representatives from science and industry interested in high performance storage solutions did meet at HLRS during April 27-29, 2009, for the eighth annual Workshop on Scalable Global Parallel File Systems. Under the motto of "Interpretation of Ideas", this year's three-day event tried to illuminate the idea of a reunification of the main disciplines of calculating machines: namely computing, storage, communication and visualization. More than 130 participants did follow a total of 35 presentations that have been on the workshop agenda.

The opening address on Monday morning was given by Prof. Dr. Michael Resch, the HLRS director. James Reaney, Research Director at BlueArc, delivered the keynote speech on Unified Scalable File Systems and pNFS. His conclusion was that research on OSD-based file systems during the recent years has led to the first open standard on parallel file systems, the IETF parallel NFS, or pNFS. Dr. Peter Braam, Braam Research, opened a discussion on Modularity for clustered File Data Management, which will lead to new systematic interfaces that will enable applications to directly control underlying file systems via specific APIs. This was in line with the talk of Dave Fellinger, CTO of DataDirect Networks, on Petaflop/s Storage Systems. Here the storage control subsystems are based on parallel processors by themselves enabling the direct execution of storage protocols as well as of full-fledged applications with high storage demand. Harry Hulen from IBM Houston gave an outlook on HPSS version 8 which will feature redundant DB2 databases

for HPSS metadata arbitration and future content management. There were more talks on Monday dealing with the Convergence at the Data Center as seen by Yaron Haviv, CTO of Voltaire, speaking for the OpenFabrics Alliance. And Herbert Grau, owner of Grau Data, promoted the idea of an Open Source Archive that will be of great help in developing the foundations for long-term national archives.

Tuesday morning saw a series of simultaneous presentations on Unified Computing and Data Management, starting out with Cisco Systems' message on the Unified Computing System, aka California Server, given by Gilles Chekroun. Dr. Reza Rooholamini did illuminate Dell's High Performance Computing Strategies followed by a series of talks on large scale parallel file systems contributed by Grau Data, Panasas and the Fraunhofer Gesellschaft. Also OpenFabrics started with a series of interesting talks on InfiniBand converged fabrics for HPC storage, virtualization of storage hardware and software, the Windows HPC Server 2008, as well as design constraints for low-latency communication platforms.

On Tuesday afternoon, the High Performance Networking Forum (HNF) Europe provided the venue for the forthcoming 100-Gbit/s all-optical communication architectures in next generation research networks. This was a follow-up on talks in the previous years on a wavelength specific switching (WSS) infrastructure for the German HPC centers that could enable and/or accelerate a new set of distributed scientific applications which revolve around

a network-centric data management scheme. There were more talks in parallel sessions on Unified Wire protocols and applications that are emerging from the ongoing standardization work at various IEEE 802.3 committees towards unified data center fabrics. Tuesday night's agenda started with an invitation to closely study Ferdinand Porsche's most famous "Four Interpretations of one Idea" at the Porsche Museum - who has contributed the "Interpretation of Ideas" motto. From there the bus moved on to the Stuttgart Museum of Fine Arts for a reception at o.T. Lounge and a very memorable dinner at the Cube Restaurant on fourth floor.

Wednesday again saw a parallel track organized by the OpenFabrics Alliance. These sessions covered various InfiniBand issues from physical layer and kernel problems to storage and message passing. Bob Curran, IBM GPFS lead, gave an update on new developments with GPFS and the Scale Out File Services solution. Finally, HLRS provided insight into its center on many levels: a first multi-GB/s demo system featuring the HPSS version 7.1 GPFS/HPSS Interface (GHI), the Xtreem File System for potential use within the Teraflop Workbench, augmented reality and simulation steering to create hybrid prototypes, as well as private tours of the center.

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

LRZ-Workshop "Molecular Modelling on Supercomputers"

More than 50 scientists and students from Munich and Erlangen (via videolink) participated in a half day workshop on molecular modelling on supercomputers, held on March 9th, 2009 at the LRZ in Garching. The presentations covered many interesting topics, such as complex dynamics of reactive systems, reactions of biomolecules with surfaces, data mining on conformations of helical proteins and the registration of biomolecules to electron tomography imaging in structural systems biology.

Key speaker was Prof. Klaus Schulten from the University of Illinois at Urbana-Chamapaign. He is director of the Theoretical and Computational Biophysics Group at UIUC and he is famous for his molecular dynamics simulation software NAMD (NAnoscale Molecular Dynamics) and VMD (Visual Molecular Dynamics). NAMD is a molecular dynamics code that scales to tens of thousands of processor cores and was recently ported to GPGPUs (CUDA). Additional speakers were Prof. Irmgard Frank from the University of Hannover, Prof. Dieter Langosch (TU Munich), Dr. Christina Scharnagl (TU Munich), Dr. Kay Gottschalk (LMU Munich), Dr. Elisabeth Villa (MPI for Biochemistry) and Dr. Jen Hsin (UIUC).

There were interesting discussions how to bridge the different important scales from single molecule first principle quantum mechanics over force-field calculations up to coarse-grained molecular dynamics. Special focus of the talks of the workshop was the utilization of the new peta-flops-machines in light of millisecond simulation timescales.

10th Teraflop Workshop

On March 16-17, 2009, HLRS and NEC held the 10th Teraflop Workshop on High Performance Computing. Topics discussed ranged from application simulation via operating and management software to hardware developments. Also European strategies concerning HPC have been covered and three major supercomputing centres, running different large scale installations by different vendors, presented their view on supercomputing.

The workshop opened with experiences from the latest installations of a NEC SX-9 systems at Tohoku University in Sendai, Japan, and at Deutscher Wetterdienst (DWD) in Offenbach, Germany. The second session concentrated on applications. The first talk addressed necessities and requirements of applications running at German national supercomputing centres. The talk shared HLRS' view on supercomputing as hybrid supercomputing, a view expressed by the present installation of a hybrid vector/scalar/GPGPU system by HLRS and NEC. The second talk was given by a former HPC Europa guest, Mrs. Ferrari of Trento University, discussing a 3D SPH scheme for free surface flow. The session was completed by Prof. Kästner, a Junior Professor in the SimTech Cluster of Excellence at Stuttgart University. His talk discussed biochemical simulations using QM/MM techniques. During the next session, Barcelona Supercomputing Center (BSC) shared its view on supercomputing with the audience. BSC has large installations fom IBM, namely Mare Nostrum and the new cell-based prototype in PRACE.

The two talks addressed BSCs hardware operational view and HPC applications on BSC systems. The afternoon was completed by a view on going forward with GPU computing, given by CEA as running site and CAPS enterprise.

The second day started with the NEC-HLRS SX-Linux project. The aim of the project is on porting a common Linux operating system to the vector system SX-9, to ease porting and usage of applications. Two further application-related talks on "Graph domain decomposition for heterogeneous cluster programming" and "Direct Numerical Simulations of Turbulent Shear Flows" completed this session. The third talk of the European supercomputing centers was given by the Edinburgh Parallel Computing Centre (EPCC), concentrating on the one hand on "EPCC Experiences with the Cray XT4" as well as on "FPGAs and HPC".

Inauguration of Cray XT5m

At a technical workshop co-hosted with the High Performance Computing Centre Stuttgart (HLRS) at the University of Stuttgart, Germany, Cray Inc. announced the launch of the new Cray XT5m line of midrange supercomputers, making the company's world-leading High Performance Computing (HPC) technology available to a broader set of users. HLRS has decided to complement its set of supercomputers with a system based on standard processor technology but using a special network interconnect. Due to the long collaboration of HLRS and Cray HLRS is now the first Cray XT5m customer.

The Cray XT5m system builds on the success of the Cray XT5 product, whose installations include the world's most powerful supercomputer for open science at Oak Ridge National Laboratory. The Cray XT5m supercomputer is a massively parallel processing (MPP) system that delivers performance, efficiency and manageability unrivaled in its price range with capabilities that previously were only available to the world's largest research facilities. "The Cray XT5m will broaden our customer base while extending the Cray experience to new users and new market segments," said Cray CEO and President Peter Ungaro. "Cray XT5m users can now benefit from the proven technology of our larger Cray XT5 systems, offering unrivaled performance and usability, but at lower price points. We are very excited to be launching this system in Europe with HLRS as their partnership with the automotive industry is an excellent example of being able to leverage Cray supercomputers to a new set of customers and applications." The CrayXT5m system incorporates a version of the Cray SeaStar(TM) network specially designed and optimized for systems with peak performance of less than 100 teraflops. This provides Cray XT5m systems with superior

bandwidth, upgradeability and manageability at price points similar to those of commodity clusters. Offered with up to six cabinets, the Cray XT5m series features Quad-Core AMD Opteron(TM) processors and a Cray SeaStar-based 2D torus interconnect. The Cray Linux Environment(TM) enables the use of a wide range of open source tools as well as streamlined porting of a broad set of applications from independent software vendors (ISVs). "The Cray XT5m system supports our mission of providing researchers and scientists from both the automotive industry and the public sector with leading edge supercomputing technology," said Michael Resch, director of HLRS. "Cray has an excellent track record of developing innovative supercomputing systems that meet a wide range of research needs. We are excited to be the first Cray XT5m customer and partner to open a new era in MPP computing together with Cray. With the Cray XT5m, we get a leading edge MPP system and the power of a Cray supercomputer that supports our scientific endeavors as well as our financial objectives."

Successful UNICORE and Supercomputing Workshop 2009

On March 18th, 2009 the UNICORE and Supercomputing Workshop took place at the Deutscher Wetterdienst (DWD) in Offenbach near Frankfurt, Germany. The objective of the workshop was to present and demonstrate interesting results of UNICORE usage in Supercomputing to major stakeholders in Supercomputing from Germany and Europe as well as to members of the UNICORE Forum (www.unicore.eu/forum).

About 40 people participated in the workshop and listened to the seven speakers. The workshop opened with an overview on UNICORE 6 followed by presentations from the life-science domain, the exploitation of DEISA – the European Supercomputing Grid e-infrastructure – and presentations from the French Atomic Energy Commission CEA and T-Systems SfR about the usage and adaptation of UNICORE in industrial environments. The detailed program and all presentations as PDF can be found at www.unicore.eu/events/supercomputing-workshop-2009/

3rd VI-HPS Tuning Workshop 2009

From February 16th to 20th, 2009, JSC organized the 3rd Tuning Workshop of the Virtual Institute – High Productivity Supercomputing (VI-HPS), following the model of two previous workshops hosted by RWTH Aachen and TU Dresden during 2008.
28 participants from Germany, Bulgaria, Denmark and the UK took advantage of the extended five-day format to learn how to make effective use of tools for verifying correctness and tuning performance of parallel programs. Those who had prepared, and tested their own MPI and OpenMP application codes on the workshop systems in advance, particularly appreciated the opportunity to analyse them with coaching from the VI-HPS tools developers.

The local JUGENE BlueGene/P and JUMP p5-575 cluster systems were particularly popular, where applications could be tuned with tips provided by experts from IBM, along with the Intel Xeon cluster and other systems made available for the workshop by partner sites RWTH Aachen and TU Dresden. With the latest Scalasca and Vampir toolsets installed for use on all of the systems, participants could not only apply them to analyse and visualize their parallel application communication and synchronization but also directly compare the analyses from different systems. Where supported, measurements and analyses also included hardware counter metrics obtained with the PAPI library, and the MARMOT tool was used to verify message-passing correctness of the initial version and at each stage of the tuning process.

A few participants were successful in getting their codes scaling to efficiently use one quarter or more of the entire BlueGene/P system, while others achieved more modest performance and scalability improvements during the course of the workshop, but left with ideas for future implementation and investigation. An interesting development from previous hands-on workshops, several codes combining OpenMP with MPI were among the predominantly MPI-based applications analyzed, and using effective hybrid parallelization they were able to achieve notably better performance on identical numbers of cores.

Further VI-HPS Tuning Workshops are being planned for the future in Germany, along with a range of hands-on and regular tutorials featuring VI-HPS tools at various international conferences: for details see www.vi-hps.org/training/

Winter School on Multiscale Simulation Methods

From March 2nd to 6th, 2009, the Winter School "Multiscale Simulation Methods in Molecular Sciences" was held at Jülich Supercomputing Centre. It continued the successful series of biennial Winter Schools organized by JSC. The focus of this year's school was on simulation methods for materials, soft matter and biomatter covering various length and time scales.
The programme was complemented by lectures on methodological foundations (Molecular Dynamics, Monte Carlo techniques, DFT and electronic structure theory, embedding methods), mathematical aspects of multiscale approaches as well as practical sessions on parallel computing.

The scientific programme was drawn up by Johannes Grotendorst, Norbert Attig, Stefan Blügel (Forschungszentrum Jülich), and Dominik Marx (Ruhr-Universität Bochum). Leading scientists in computational physics, chemistry, biology and mathematics presented stimulating lectures for 50 PhD students and young postdocs from 20 different countries. In two poster sessions the young scientists presented topics from their own research.

Highlights of the meeting were the keynote lectures of Aiichiro Nakano, University of Southern California, on "Large spatiotemporal-scale materials simulations on Petaflops computers" and Klaus Schulten, University of Illinois at Urbana-Champaign, on "Application of residue-based and shape-based coarse graining to biomolecular simulations".

Lecture notes of the School have been published as volume 42 of the NIC series. The volume is available either as hard copy or as PDF files on the web (http://www.fz-juelich.de/nic-series/volume42/).

International Exascale Software Project Launched

The mission of the "International Exascale Software Project" (IESP) is to lay the foundation for supercomputing with Exaflop performance. The IESP aims at mobilizing the global HPC community in order to combine and coordinate the collective efforts towards innovative software for Peta- and Exascale computers on an international level. For this purpose, the IESP will hold a series of three workshops in the US, Europe, and Japan during 2009. The first, invitation-only workshop took place on April 7th and 8th in Santa Fé, New Mexico. The 69 attendees, 22 from Europe, included participants from industry and academia, as well as government agencies, with expertise in a range of critical areas. The next IESP workshop will be organized on June 28th and 29th, 2009 in Paris, shortly after ISC 2009, followed by a workshop in Japan in fall 2009. Among the goals for the first meeting was the assessment of short, medium and long-term needs of applications for Peta/Exascale systems as well as a draft for a roadmap for software on extreme-scale systems. Breakout sessions explored ways how laboratories, universities, and vendors can work together in order to create Exascale software, including a discussion of existing R&D plans with focus on new programming models and tools, extreme scalability, many-core processors, heterogeneity, and performance. A central topic was the international coordination of appropriate funding strategies. Presentations from Satoshi Matsuoka (Tokyo Institute of Technology), Horst Simon (Lawrence Berkeley National Laboratory), and Thomas Lippert (Gauss Centre for Supercomputing) summarized the situation with regard to HPC soft- and hardware developments in Japan, US, and Europe, respectively. The discussion emphasized the need to include the requirements of science and engineering right from the beginning. In the US, the initiative is supported by all major funding agencies (NSF, DOE ASCR, DOE NNSA). In Europe, a consortium with the members EDF, GENCI, INRIA (France), EPSRC (UK), JSC (DE), BSC (ES), and NCF (NL) is starting the "European Exascale Software Initiative" (EESI) with the aim to build and promote the European position within the IESP. Further goals of the EESI are to identify grand challenge applications from academia and industry that will benefit from Petaflop capabilities in 2010 and Exaflop systems around 2018, with a strong economical, societal and/or environmental impact, to identify critical software issues for Peta/Exascale systems, to propose coordinated research actions for Peta/Exascale software at the European level, and to build an coordinated international program in education and training for the next generation of computational scientists.

More information about the IESP including whitepapers and slides of the presentations of the first workshop can be found at www.exascale.org.

PROSPECT: Towards a European Technology Platform

European supercomputing centres should gain more influence on the development and use of future petaflop architectures. This has become common understanding and motivates the efforts of leading supercomputing centres in Europe to set up joint ventures together with HPC vendors which may succeed in accomplishing European requirements in the international HPC competition. The joint venture PROSPECT, the acronym stands for Promotion of Supercomputing and Petacomputing Technologies, is a first consortium to realize these goals. It is an open European interest group with partners from research institutes, universities and industry in the field of high-performance computing (HPC), supercomputing and numerical simulations in a diversity of scientific and engineering fields. On December 15th, 2008, PROSPECT became a legal entity: the association PROSPECT e.V. was founded in Munich. The Jülich Supercomputing Centre, the Barcelona Supercomputing Centre and the Leibniz-Rechenzentrum Garching will chair the association. The aim of PROSPECT is to strengthen the European position in the worldwide HPC competition with regard to the development and usage of future petaflop supercomputers by science and industry. PROSPECT is already a member of the Advisory Group for Strategic Technologies (STRATOS) of the Partnership for Advanced Computing in Europe (PRACE). Within STRATOS, partners of the PRACE project will set up joint projects in cooperation with industrial partners involving consortia like PROSPECT. The aim is to gain access to and influence HPC technology roadmaps worldwide and facilitate contributions from European industry to these technologies. The preparation of a European Technology Platform for HPC is an important long-term goal of PROSPECT.

HLRS Workshop Report and Outlook

With the paradigm-shift to multi-core CPUs, new programming languages evolve. In October 2008, HLRS invited Dr. Andrew A. Johnson to start with a 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 this year.

In spring 2009, HLRS had a number of different courses and workshops. The total number of participants in these courses from Germany and Europe was about 175. One of the flagships of our courses is the week on **Iterative Solvers and Parallelization**. Prof. A. Meister and Prof. B. Fischer teach basics and details on Krylov Subspace Methods. Lecturers from HLRS give lessons on distributed memory parallelization with the Message Passing Interface (MPI) and shared memory multi-threading with OpenMP. This course will be repeated in September 2009 at LRZ in Garching.

NEC SX-9 Usage and Programming is a two day course, dedicated to our NEC SX-9 users to learn about vectorization and optimization of their applications on the newly upgraded system at HLRS (see page 62).

Fortran for Scientific Computing started the first time at HLRS in October 2006 and had up to now more than 150 participants. Mainly PhD students from Stuttgart and other universities in Germany came to learn not only the basics of programming, but also to get in insight on the principles of developing high-performance applications with Fortran. The Fortran course will be repeated in October this year.

In the second half of 2009, we also repeat the **Introduction to Computational Fluid Dynamics** (in September at HLRS). The emphasis is placed on explicit finite volume methods for the compressible Euler equations. Moreover outlooks on implicit methods, the extension to the Navier-Stokes equations and turbulence modeling are given. Additional topics are classical numerical methods for the solution of the incompressible Navier-Stokes equations, aero-acoustics and high order numerical methods for the solution of systems of partial differential equations.

The **Parallel Programming Workshop**, October 12th - 16th at HLRS, has 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.

In the table, you can find the whole HLRS series of training courses in the second half of 2009. They are organized at HLRS and also at several other HPC institutions:
LRZ Garching, JSC (FZ Jülich), and CSCS (Manno, CH).

July-December 2009 – Workshop Announcements
Scientific Workshops at HLRS, 2009
High Performance Computing in Science and Engineering – The 12th Results and Review Workshop of the HPC Center Stuttgart (October 8-9)
IDC International HPC User Forum (October 5-6)
Parallel Programming Workshops: Training in Parallel Programming and CFD
Parallel Programming with MPI & OpenMP (CSCS Manno, CH, August 11-13)
Iterative Linear Solvers and Parallelization (LRZ, Garching, September 14-18)
Introduction to Computational Fluid Dynamics (HLRS, September 21-25)
Message Passing Interface (MPI) for Beginners (HLRS, October 12-13)
Shared Memory Parallelization with OpenMP (HLRS, October 14)
Advanced Topics in Parallel Programming (HLRS, October 15-16)
Unified Parallel C (UPC) and Co-Array Fortran (CAF) (HLRS, October 19)
Parallel Programming with MPI & OpenMP (FZ Jülich, JSC, November 30 – December 2)
Training in Programming Languages at HLRS
Fortran for Scientific Computing (October 26-30)
URL
http://www.hlrs.de/events/
https://fs.hlrs.de/projects/par/events/2009/parallel_prog_fall2009/
https://fs.hlrs.de/projects/par/events/2009/prog_lang_fall2009/






High Performance Computing Courses and Tutorials

Introduction to Molecular Modelling

Date & Location
Juli 21-22, 2009
LRZ Building, Garching/Munich

Contents
The course gives an introduction into the simulation of molecules based on several software packages (both classical and quantum-mechanical: e.g. NAMD, VMD and CPMD) on the supercomputer at LRZ Garching.

The course includes an introduction to the remote visualization services at LRZ as well as hands-on sessions.

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

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

Date & Location
August 4-7, 2009
JSC, Research Centre Jülich

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

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

Introduction to Programming and Using the IBM Supercomputers

Date & Location
August 10-11, 2009
JSC, Research Centre Jülich

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

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

Parallel Programming with MPI, OpenMP and PETSc

Date & Location
August 11-13, 2009
CSCS, Manno (CH)

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 CSCS in collaboration with HLRS.

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

Iterative Linear Solvers and Parallelization

Date & Location
September 14-18, 2009
LRZ Building, Garching/Munich

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

Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the basic constructs of iterative solvers, the Message Passing Interface (MPI) and the shared memory directives of OpenMP.

This course is organized by University of Kassel, HLRS, IAG, and LRZ.

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

Introduction to Computational Fluid Dynamics

Date & Location
September 21-25, 2009
HLRS, Stuttgart

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".

This course is organized by HLRS, IAG, and University of Kassel.

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

Parallel I/O and Portable Data Formats

Date & Location
October 5-7, 2009
JSC, Research Centre Jülich

Contents
This course will introduce the use of MPI parallel I/O and portable, self-describing data formats, such as HDF5 and netCDF. Participants should have experience in parallel programming in general, and either C/C++ or Fortran in particular.

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

Message Passing Interface (MPI) for Beginners

Date & Location
October 12-13, 2009
HLRS, Stuttgart

Contents
The course gives an full introduction into MPI-1. Further aspects are domain decomposition, load balancing, and debugging. An MPI-2 overview and the MPI-2 one-sided communication is also taught.

Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the basic constructs of the Message Passing Interface (MPI).

Course language is English (if required).

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

High Performance Computing Courses and Tutorials

Shared Memory Parallelization with OpenMP

Date & Location
October 14, 2009
HLRS, Stuttgart

Contents
This course teaches shared memory OpenMP parallelization, the key concept on hyper-threading, dual-core, multi-core, shared memory, and ccNUMA platforms.

Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the directives and other interfaces of OpenMP. Tools for performance tuning and debugging are presented.

Course language is English (if required).

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

High-Performance Computing with GPGPUs

Date & Location
October 14-15, 2009
LRZ Building, Garching/Munich

Contents
Heterogeneous GPGPU computing offers great potential to accelerate applications substatially. This programming workshop includes hands-on sessions and gives examples of application and an introduction to the use of CUDA and the Portland Group Fortran Compiler, amongst others.

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

Advanced Topics in Parallel Programming

Date & Location
October 15-16, 2009
HLRS, Stuttgart

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

Course language is English (if required).

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

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

Date & Location
October 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/

Parallel Programming with Dynamic Languages

Date & Location
October 20-21, 2009
LRZ Building, Garching/Munich

Contents
The course teaches the use of dynamic languages like python and R for parallel programming of supercomputers, especially on the large machines in Jülich, Stuttgart and Garching. Hands-on sessions with example applications are given.

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

Fortran for Scientific Computing

Date & Location
October 26-30, 2009
HLRS, Stuttgart

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

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

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/

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

Arifa Ali Khan arifa.ali-khan@physik.uni-regensburg.de
Norbert Attig n.attig@fz-juelich.de
Martin Bernreuther bernreuther@hlrs.de
Martin Buchholz buchholm@in.tum.de
Ulrich Detert u.detert@fz-juelich.de
Stephan Deublein stephan.deublein@mv.uni-kl.de
Jutta Docter j.docter@fz-juelich.de
José Faraldo-Gómez jose.faraldo@mpibp-frankfurt.mpg.de
Zoltan Fodor fodor@bodri.elte.hu
Hannsjörg Freund freund@mpi-magdeburg.mpg.de
Jochen Fröhlich froehlich@ict.uni-karlsruhe.de
Helmut Grubmüller hgrubmu@gwdg.de
Georg Hager georg.hager@rrze.uni-erlangen.de
Hans Hasse hans.hasse@mv.uni-kl.de
Tobias Heidig info@mpi-magdeburg.mpg.de
Stefan Holl st.holl@fz-juelich.de
Martin Horsch martin.horsch@uni-paderborn.de
Amer Inayat amer.inayat@crt.cbi.uni-erlangen.de
Ferdinand Jamitzky ferdinand.jamitzky@lrz.de
Thilo Kopp thilo.kopp@physik.uni-augsburg.de
Ivan Mary imary@onera.fr
Florian Niebling niebling@hlrs.de
Christoph Niethammer niethammer@hlrs.de
Natalia Pavlenko natalia.pavlenko@physik.uni-augsburg.de
Sabine Roller roller@hlrs.de
Helmut Satzger helmut.satzger@lrz-muenchen.de
Wilhelm Schwieger schwieger@rzmail.uni-erlangen.de
Armin Seyfried a.seyfried@fz-juelich.de
Michael Stephan m.stephan@fz-juelich.de
Godehard Sutmann g.sutmann@fz-juelich.de
Dominic von Terzi terzi@ifh.uni-karlsruhe.de
Gerhard Wellein gerhard.wellein@rrze.uni-erlangen.de
Klaus Wolkersdorfer k.wolkersdorfer@fz-juelich.de
Jadran Vrabec jadran.vrabec@upb.de
Thomas Zeiser thomas.zeiser@rrze.uni-erlangen.de
Mareike Zink mzink@gwdg.de

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