inSiDE

inSiDE • Vol. 4 No. 1 • Spring 2006

Innovatives Supercomputing in Deutschland

Editorial

Welcome to this new issue of inSiDE the bi-annual German Supercomputing information journal. The first half of 2006 is full of good news for our community. The LRZ at Munich has opened its new facilities at Garching. With its 36x36x36 m^3 cube for computing facilities it is now ready to welcome its next generation HPC system to be installed in mid 2006. NIC at Jülich has installed a new large Blue Gene/L system and is hosting Europe's largest of these IBM systems. Other centers are lining up for new installations and we will see some interesting decisions this year.

The installation of new systems has given rise to new and improved methods in various application fields. This issue is, therefore, focussing on new results and applications. Peter Lammers et al report on their usage of the Lattice Boltzmann Method on the HLRS NEC SX-8. Up to nearly 6 TFLOP/s can be achieved with the new method on the new system. Inga Mahle et al from the TU Munich are reporting on direct numerical simulations. The code is still using the Hitachi SR8000 system but first results already indicate how such an application will benefit in the future from massively parallel systems introduced at LRZ. The same can be said about the contribution by Ari Seitsonen (CNRS) and Herbert Over (University of Giessen) who have brought catalysis to the supercomputer and are looking forward to the usage of even larger systems. The prediction of properties of mineral surfaces is covered in a paper by Rossitza Pentcheva from the University of Munich. First results of the Blue/Gene system at Jülich can be found in an article about realistic descriptions of materials with strong correlations. The method shows very good speedup up to 2048 processors.

As announced in our last issue inSiDE starts to present information about the German e-science initiative D-Grid. D-Grid was started in September 2005 and is now in a phase to create very interesting and positive results which we want to present in inSiDE. We kick off a series of articles with a description of the engineering Grid project InGrid that is part of D-Grid. Further projects and an overall description of D-Grid will follow in later issues.

Continuing our effort to report on European projects this issue describes a project started only recently. As a result of the last call on Grid activities in the Framework Program 6 of the European Commission BEinGrid wants to exploit European Grid middleware by creating a toolset repository of Grid services from across the Grid research domain and to use these services to deliver a set of successful business experiments that stimulate the early adoption of Grid technologies across the European Union.

As always this issue includes information about 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.

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

Contents

Recent Performance Results of the Lattice Boltzmann Method

Abstract

We show performance results for the lattice Boltzmann method on three different architctures. The benchmarked architectures include an IBM Power5, the SGI Altix 3700 (LRZ Munich) and the NEC SX-8 (HLRS Stuttgart) respectively. The application used is the well known lattice Boltzmann solver BEST. Furthermore we use a modified lattice Boltzmann solver using a boundary-fitted cartesian mesh mesh to compare the performance of indirect addressing on the NEC SX-8 with the predecessor model SX-6.

Lattice Boltzmann Method

Nowadays the lattice Boltzmann method is applied to a wide range of scientific and and engineering disciplines including chemical and process engineering, bio and environmental processes and medical applications. Preferably it is used when highly complex geometries are involved. The first author of this artikle is intensively using the lattice Boltzmann method for simulations of flow control in wall bounded turbulence, turbulent drag reduction and control of transition [1]. Among chemical engineering the flow solver BEST was developed at the Lehrstuhl für Strömungsmechanik (LSTM), Universität Erlangen-Nürnberg for this purpose. Here, the BEST solver is used for most of the benchmarks their results are shown in this text. The benchmark itself (a plane channel) is motivated by wall bounded turbulence.

The lattice Boltzmann method consists of a discretized kinetic equation for the one particle distribution function f,

$$f_\alpha(x+c_\alpha,t+1)=f_\alpha(x,t)-\frac{1}{\tau}\{f_\alpha(x,t)-f_\alpha^{eq}(x,t)\}, \quad (1)$$

for which an appropriate phase velocities lattice needs to be specified. In BEST, the 3D spatial cartesian grid is coupled to the D3Q19 streaming lattice [2] with 19 mesoscopic variables c_α. τ is a relaxation parameter that determines the rate at which the particle interaction drives f to the equilibrium state f^{eq}. f^{eq} is appropriately chosen to satisfy the Navier-Stokes equation and depends on the hydrodynamic quantities like density, pressure and the velocity fields which are calculated from the moments of f.

In the benchmark codes, the algorithm is implemented just as given by eq. (1). That means a collision step (r.h.s) in which basically the equilibrium state function f^{eq} is calculated is followed by a propagation step (l.h.s) in which the mesoscopic variables f_α are streamed along the lattice directions c_α. In order to reduce the memory traffic both steps are combined. As the streaming step implicates a data dependency, it cannot be vectorized in a straightforward way. In the implementation this is by-passed by using two arrays of f for time step t and $t+1$ which are used in an alternating way. For cache based machines the inner loop of the main kernel is broken up in pieces. Especially the propagation step is done for pairs of the densities f_α.

For parallelization the overall spatial grid is block-wise distributed to the processors in BEST. At the boundaries, the propagating densities are first copied to a ghost layer and then sent to the adjacent processor block. The data exchange is realized by MPISendrecv which is at least the optimal choice on the SX network for this communication pattern. For more details of the implementation see [1] and [3].

Performance

In Figure 1, 2 and 3 the single and parallel performance of the algorithm implemented in BEST are shown for an IBM p575, an SGI Altix3700Bx2 and the NEC SX-8. The measurements on the Power5 are done by IBM [4]. In each case the CPU efficiency depending on the domain size per process is plotted. The efficiency is related to the peak performance (vector peak performance for the SX).

First we will look for the single processor performance given by the dotted black graphs first. The achievable performance per process of the Power5 CPU (Figure 1) is in the range between 16,5 and a maximum of 25 % efficiency. This corresponds to 1,25 and 1,9 GFlop/s. Significant outliers due to cache effects can be observed for inner loop lengths of 16, 24 and 32.

The Itanium2 (Figure 2) achieves a maximum of 36 % efficiency corresponding to 2,3 GFlop/s. Performance drops significantly when the problem size exceeds the cache size. Further increasing the problem size, compiler-generated prefetching takes over and leads to gradual improvement up to a final level of 2,2 GFlop/s. In contrast, the performance of the vector system in Figure 3 increases with increasing vector length and saturates at an efficiency of close to 75 %, i.e. at a single processor application performance of 11,9 GFlop/s.

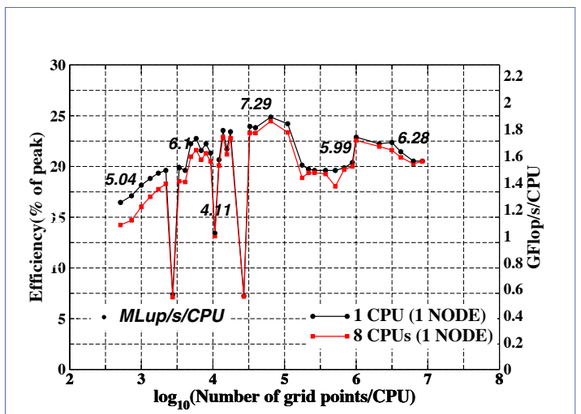

Figure 1: Left: Efficiency, GFlop/s and MLup/s of the lattice Boltzmann solver BEST depending on the domain size and the number of processors for up to 8 CPUs of IBM p575

For the parallel scalability analysis we focus on weak-scaling scenarios. Here, ideal scalability means that all curves would collapse into one.

In Figure 1 the results for a node IBM p575 with 8 CPUs are available. In contrast to the Power4 (results not shown here) the p575 provides sufficient memory bandwidth inside one node now.

In the SGI Altix two processors have to share one memory connection. Consequently when using two processors single CPU performance drops to 1,29 GFlop/s. Prefetching is still important but not as efficient as for one

Applications

processor. Further, in Figure 2 weak scaling results for up to 120 Itanium2 CPUs on an SGI Altix3700 are given and in Figure 3 the NEC SX-8 results with up to 576 CPUs. For the lattice Boltzmann application we see on the 576 processor NEC SX-8 system a maximum sustained performance of 5,7 TFlop/s. The same performance level would require at least 6400 Itanium2 CPUs on an SGI Altix3700. As seen in 3 the scalability of the NEC SX-8 system is not perfect. Parasitic non synchronous operating system processes show influences on aggregated waiting times. This effect becomes of course smaller by increasing per process system size. The load balancing itself is ideal for this benchmark case and can not cause any problems.

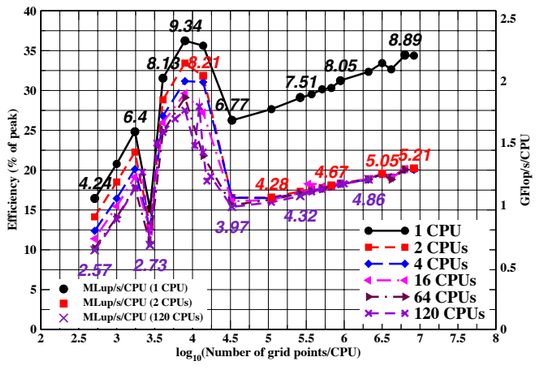

Figue 2: Efficiency, GFlop/s and MLup/s of the lattice Boltzmann solver BEST depending on the domain size and the number of processors for up to 120 CPUs of a SGI Altix3700Bx2

Finally we would like to focus in this context on the performance of indirect array addressing on the SX-8. In BEST the calculation is done on a block structured regular grid. But for highly complex geometries it is of course more appropriate to store only the cells inside the flow domain. For this, an index is needed to select the fluid cells and the information about adjacent nodes has to be stored in index lists.

In Figure 4 the performance on the SX-6 and 8 of such an implementation of the lattice Boltzmann method is plotted. The code used for the measurement is a collaborative development of the RRZE, the Institut für Computeranwendungen im Bauingenieurwesen (CAB), Universität Braunschweig, Section Computational Science, University of Amsterdam, Department of Medical Physics and Clinical Engineering, University of Sheffield, NEC CCRL and HLRS. The test cases are identical with the cases used for the previous plots.

Till now only a serial code version is available. This explains the slightly better performance for small grids. In the results shown so far, the communication dominates the calculation for small grids. The single CPU efficiency of the SX-6 is at a level of 60 %. On the other hand the efficiency of the SX-8 is still 50 % but 10 % lower than on the SX-6. The reason is the higher memory latency of the recent system. Therefore it is more important on the SX-8 to prefetch the vector gather/scatter instructions for the indirect load/stores and hide the latency by other instructions. This can typically be done by unrolling simple loops. The present loop is complex enough for hiding these latencies partially.

Outlock

We too look forward to the performance of the new HLRB II system, an SGI Altix 4700. On this system the performance drop between one and two CPUs should not occur because of the better memory connection. It would also be interesting to have results for the IBM Blue Gene in Jülich although this architecture is not aiming at flow simulations.

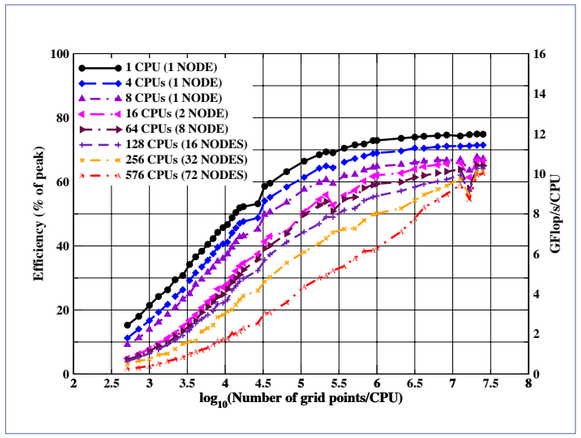

Figure 3: Efficiency, and GFlop/s of the lattice Boltzmann solver BEST depending on the domain size and the number of processors for up to 72 nodes 576 CPUs of a NEC SX-8

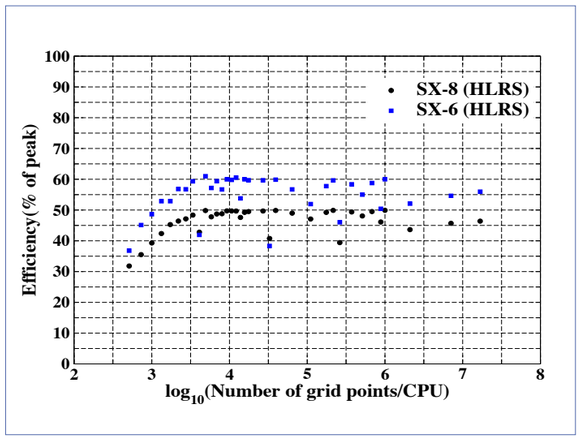

Figure 4: Performance of an unstructured implementation of the lattice Boltzmann kernel on the SX-8 in comparison to the SX-6

References

1. Lammers, P.
Direkte numerische Simulationen wandgebundener Strömungen kleiner Reynoldszahlen mit dem lattice Boltzmann Verfahren. Dissertation, Universität Erlangen-Nürnberg (2005)

2. Qian, Y.H., d'Humières, D., Lallemand, P.
Lattice BGK models for Navier-Stokes equation Europhysical Letter 17 (1992), 479-484

3. Wellein, G., Zeiser, T., Hager, G., Donath, S.
On the single processor performance of simple lattice Boltzmann kernels. Computers & Fluids (2005) in Press, corrected Proof, available online December 20, 2005

4. Pichlmeier, J.
IBM, private communication (2005)

• Peter Lammers
• Uwe Küster

Höchstleistungsrechenzentrum Stuttgart, Germany

Direct Numerical Simulation (DNS) and Large Eddy Simulation (LES) of Turbulent Reacting Shear Layers

Hydrogen combustion presently plays an important role in transportation and power applications like rocket engines. This will remain so in the near future. In order to gain full understanding in all processes involved, it is important to investigate simple configurations, like turbulent temporally developing and reacting shear layers, that serve as generic test cases for any non-premixed combustion device working with hydrogen. In a first step, we have chosen a simple combustion model with just one global, infinitely fast chemical reaction, which enabled us to perform a high-resolution DNS.

We simulate two gas streams with opposite flow directions, between which a turbulent shear layer develops and grows over time. Figure 1 shows the density distribution in a vertical plane of the three-dimensional domain: The upper stream, the velocity of which is directed to the right, is air, while the lower stream is a mixture of hydrogen and nitrogen. The Mach number in this test case is low to reflect the situation in combustion chambers. The use of an infinitely fast chemical reaction permits relating the chemical species to just one scalar, the so-called mixture fraction. Therefore, only the field of this non-reactive quantity has to be computed instead of the distribution of all chemical species, and a DNS with a high number of grid points, e.g. 64 million for the computation in Figure 1 can be performed. As a DNS resolves all spatial and temporal scales down to the smallest ones, the results of such a computation are of great value for many detailed statistical investigations and the development of improved subgrid-scale models for LES.

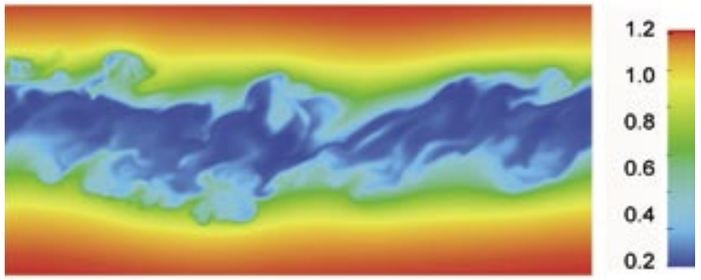

Figure 1: Instantaneous density field in [kg/m^3] of the DNS with infinitely fast chemistry, vertical cut through the domain

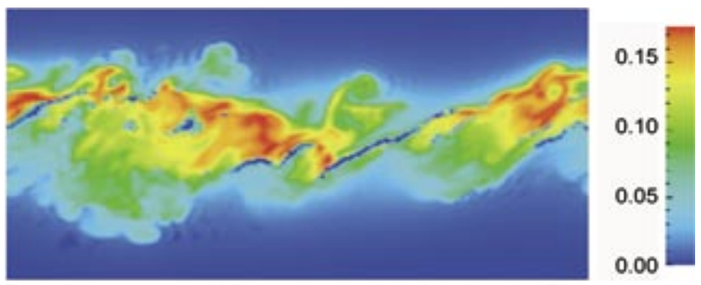

Figure 2: Instantaneous H^2O mass fraction field of the flamelet LES with infinitely fast chemistry, vertical cut through the domain

Since the direct numerical simulation already took as much as 82,000 CPU-hours on 192 processors of the Hitachi SR8000 at LRZ, it is too expensive to refine the combustion model as this would further increase the computational cost. When, on the other hand, the smallest scales are not resolved but modeled, as it is done in an LES, coarser grids and larger time steps can be used reducing the computational requirements drastically. Instead, more complicated combustion models can be applied, like finite rate chemistry, which has been implemented in our LES code. We considered 9 species and 19 chemical reactions. Yet even with an LES grid that has about two orders of magnitude fewer grid points than the DNS, it is not feasible with the present computational resources to integrate the transport equations of all these species in 3D. Therefore we used a flamelet approach, which is based on the assumption that the thin turbulent flame consists of many laminar flames, the so-called flamelets. These flamelets are computed in one-dimensional space which makes it possible to apply the detailed reaction scheme. The computation of the flamelets is done before the actual LES and the results are stored in a database. During the LES, the flamelet data, like the species mass fractions, are recovered from this database.

One of our main objectives concerning the evaluation of the flamelet LES is to investigate the influence of detailed diffusion, namely the Soret and Dufour effects, which represent the molecular transport of species due to temperature gradients and the transport of heat due to species gradients, respectively. Due to computational restrictions they have mostly been neglected in combustion simulations, however they are of particular importance for hydrogen chemistry. The one-dimensional nature of the flamelets permitted us to compute the flamelet database twice, one time with these thermodiffusive effects and one time without them. Then, we used these databases in two different LES simulations. Figure 2 shows an instantaneous field of the H_2O mass fraction for the computation with detailed diffusion. It is clearly visible that this simulation has a lower resolution than the DNS in Figure 1. The blue spots within the turbulent shear layer are places where the flame is extinguished due to heat release, respectively its modeled equivalent, namely, high scalar dissipation rate (Figure 3) and no H_2O is produced. One of the most striking differences that we found between detailed and simplified diffusion was the value of this extinction limit. For simplified diffusion the scalar dissipation rate at which extinction occurs is more than double the value obtained with detailed diffusion. This not only has local influence, but affects mean profiles as well.

Because the interpolation of the flamelet quantities has to be done at each time step and at all grid points, we have, up to now, only been able to perform flamelet LES. Flamelet DNS is one type of computation that we are planning to perform on the next generation system at LRZ (HLRB-II). It will help us to further validate our LES results and to gain more insight into detailed diffusion effects.

• Inga Mahle
• Jörn Sesterhenn
• Rainer Friedrich

Fachgebiet
Strömungsmechanik,
TU München,
Garching, Germany

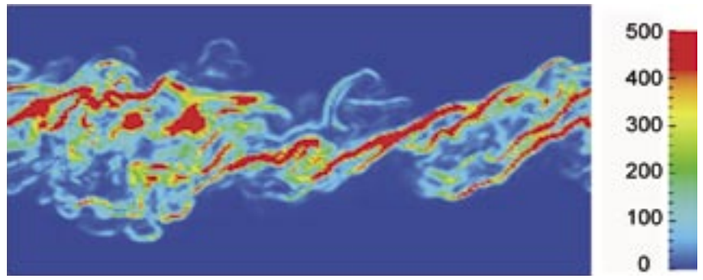

Figure 3: Instantaneous scalar gradient field in [s^{-1}] of the flamelet LES with infinitely fast chemistry, vertical cut through the domain, values above the extinction limit (400 s^{-1}-1000 s^{-1}) are shown in red and are partially beyond the scale

Chemistry by Numbers: Catalysis goes Supercomputing

In our every-day life chemical processing by catalysis – increasing the reaction rate by imposing the reactants to a catalyst, which is not consumed in the process – forms an important field of investigation. In order to find the optimal design it is of utmost importance to understand the reaction process on the atomic scale. Hence, not only sophisticated experiments are essential, but computational quantum chemistry provides additional assistance by confirming or correcting the empirical data and giving insight of the atomic details – a request for supercomputing is born!

The best-suited calculational method for surface studies is the density functional theory (DFT) in the Kohn-Sham (KS) formulation. It combines sufficient accuracy with a manageable computing effort, treating the electronic structure explicitly with the electron density as the basic variable. The Kohn-Sham equations not only provide the total energy of the system, but also the forces acting on the ions, allowing for structural optimization or even more realistic molecular dynamics simulations. The discretization of the Kohn-Sham equations by expanding the electronic orbitals into a basis set leads to a class of numerical problem which can be solved using modern HPC equipment. Most applications in surface science are done using plane waves as the basis set. Their advantage is that the matrix elements in the equations can be cast into diagonal form employing fast Fourier transforms (FFT) to switch between real and reciprocal space. However, the need for supercomputing becomes evident from the vast number of basis functions needed: A case of "only" about 100-200 atoms can require a search for the optimal value of 106-109 unknown coefficients – which even have to be defined self-consistently! Luckily the problem consists of relatively sparse matrices. In parallel computers the two major operations needed are global sums and all-to-all communication in linear algebra (BLAS libraries) and FFT routines. Scaling up to thousands of processors has been demonstrated using the CPMD code [1].

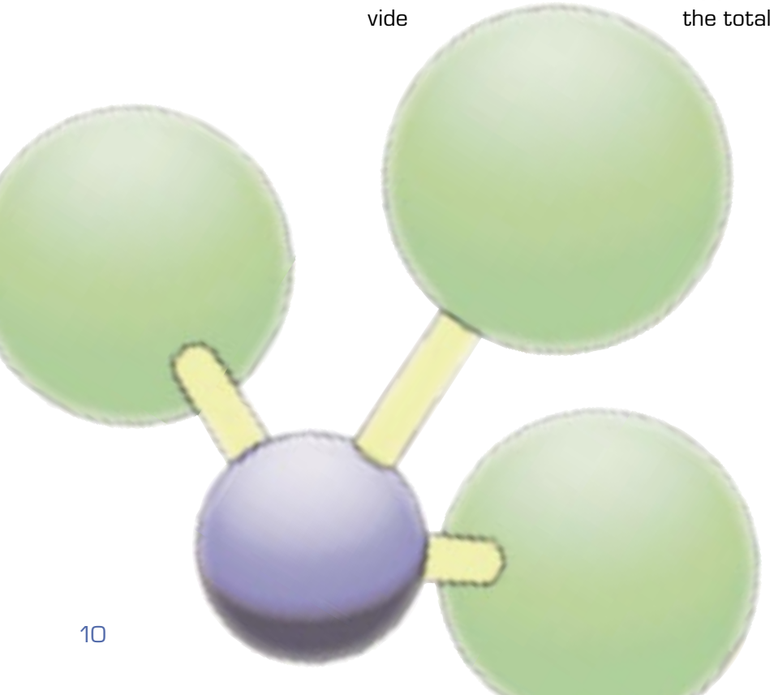

A classic example of catalysis is the three-way Rh-Pd-Pt catalyst in the car exhaust. It converts the exhaust gases coming from the engine into environmentally less hazardous end products, for example by oxidising carbon monoxide (CO) into carbon dioxide (CO_2). This process also serves as a model for catalytic reaction due to its simplicity.

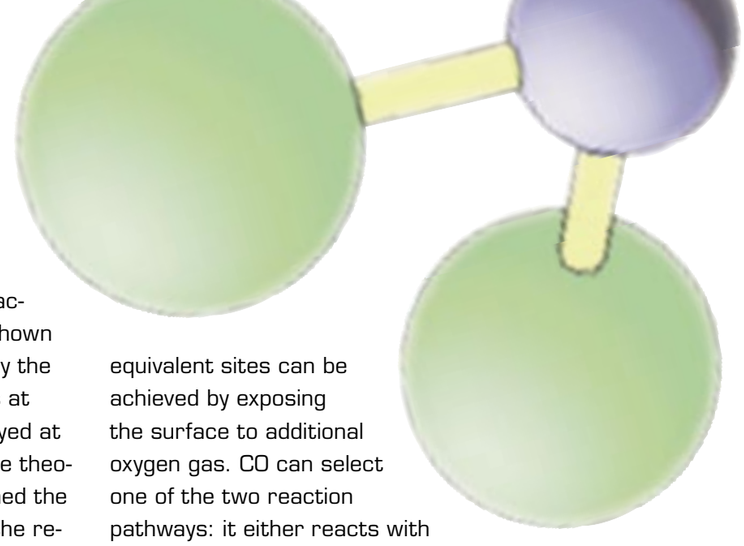

Ruthenium (Ru) is another transition metal. Its surfaces are initially a bad catalyst for this reaction, but its efficiency can be enhanced tremendously by a massive oxygen dose. The active phase on the surface was shown to be ruthenium dioxide (RuO_2) by the group of Herbert Over (who was at the time of this discovery employed at Fritz-Haber-Institut in Berlin). The theoretical calculations have confirmed the surface structure and clarified the reaction mechanisms [2 and references therein].

The most common surface orientation of ruthenium dioxide, RuO_2(110), exposes two coordinatively unsaturated sites: The 5-fold (6-fold in bulk) 1f-cus-Ru atoms and the 2-fold (3-fold in bulk) Obr oxygen atoms are shown in Figure 1. The reactant CO binds to the surface at the 1f-cus-Ru site. A surface termination with an over-stoichiometric concentration of oxygen O_{cus} bound to the equivalent sites can be achieved by exposing the surface to additional oxygen gas. CO can select one of the two reaction pathways: it either reacts with the added, weakly bound (1-fold, binding energy 83 kJ/mol) O_{cus} or the more strongly (2-fold, binding energy 245 kJ/mol) bound O_{br}.

Whereas intuition would lead us to think that the former reaction channel dominates, it was shown experimentally that both reactions occur at the surface. Thus we again decided to perform computations to study the reaction mechanism using DFT.

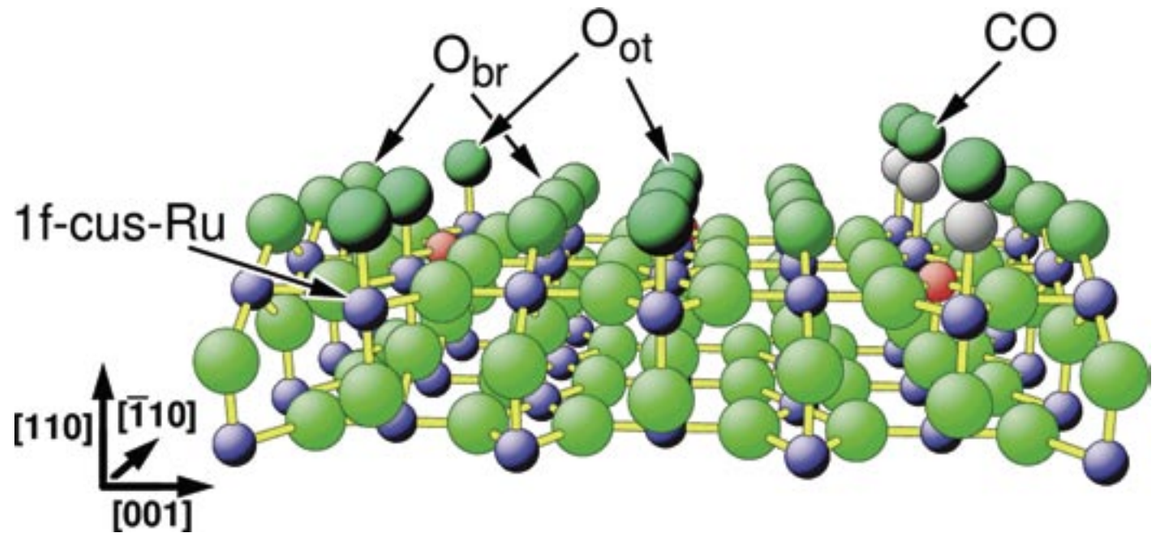

Figure 1: Ball and stick model of the RuO_2(110) surface with adsorbed on-top O (O_{ot}) atoms and on-top CO molecules. Large and small balls represent oxygen and ruthenium atoms, respectively

In our calculations for the reactions we used the distance between the carbon of CO and the reacting oxygen as the reaction coordinate. The energy along the reaction path is shown in Figure 2. The transition state is at the maximum of the curve, and the energy difference between the initial and transition state is the activation barrier E_{act} for the reaction.

Since the reaction rate depends exponentially on the activation barrier, from the two similar reaction barriers obtained from Figure 2 we can conclude that the reaction probabilities would be similar along the two reaction pathways as opposed to our intuitive guess. What is the reason for this, even if the end state is energetically very different? We note that the local geometry at the transition states of the CO-O react complex are similar along both paths. Further it is CO which has to be activated for the reaction to take place: Since the CO can be displaced towards either of the reacting oxygens with a similar energy cost both pathways are active at reaction conditions on the catalyst.

As a conclusion we want to express the fruitfulness and enhanced productivity in scientific output by the continuous inter-play between the experiments and theory. The computing time at LRZ Munich has been essential to achieve these results on the theoretical side.

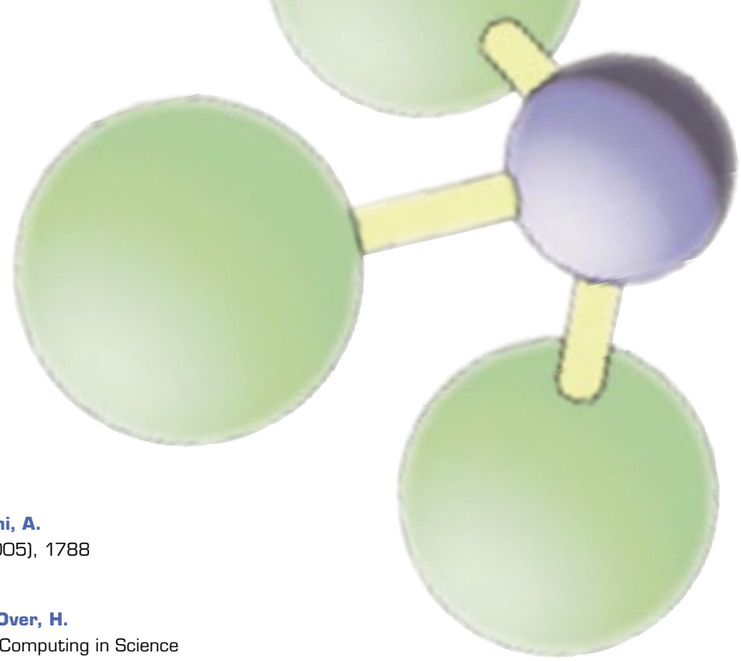

References

[1] **Hutter, J. and Curioni, A.**
ChemPhysChem 6 (2005), 1788
http://www.cpmd.org

[2] **Seitsonen, A.P. and Over, H.**
in "High Performance Computing in Science and Engineering in Munich 2002", Springer-Verlag (2003), p. 171 ff, and references therein

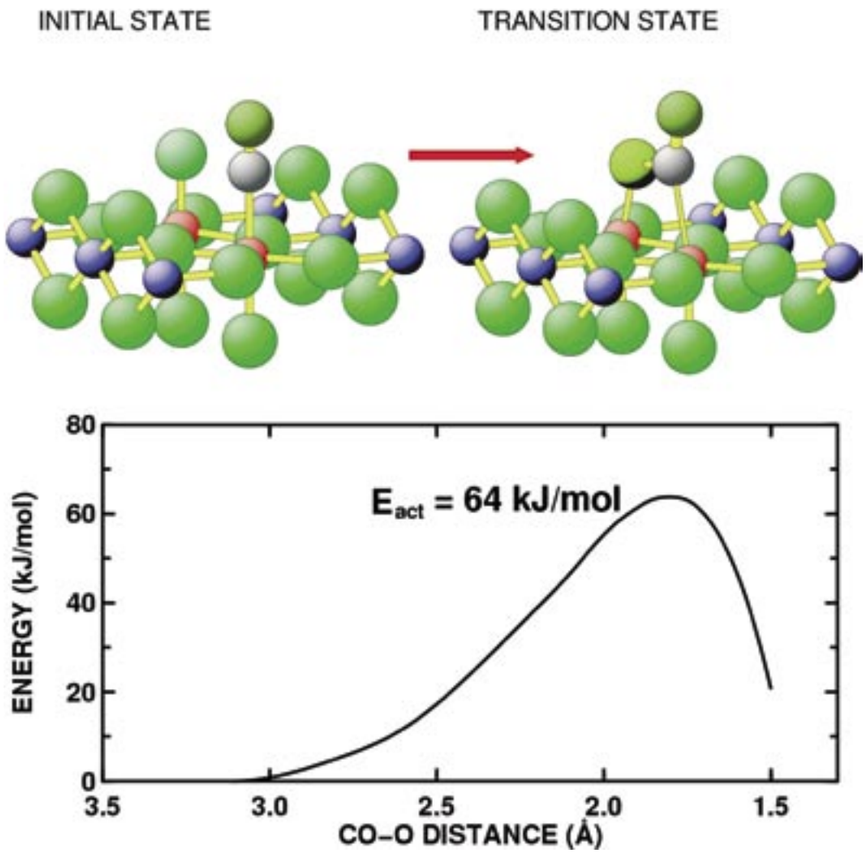

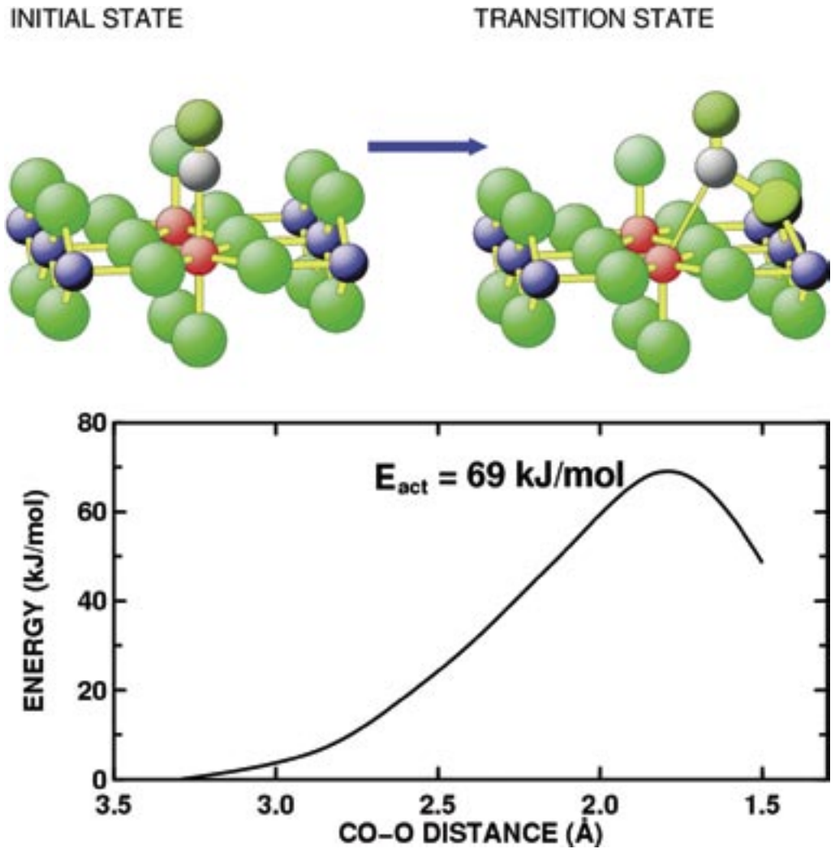

Figure 2:
(Above) The initial and transition states for the two reaction paths.
(Below) The energy along the reaction path for the reaction CO-O_{ot}–left and CO-O_{br}–right

• Ari P. Seitsonen

IMPMC, CNRS & Université Pierre et Marie Curie

• Herbert Over

Institut für Physikalische Chemie, Justus-Liebig-Universität Giessen, Germany

Predicting the Stability and Properties of Mineral Surfaces at Ambient Pressure

A multitude of important reactions both in nature and industry (e.g., catalysis) take place on mineral surfaces, some examples being the adsorption and dissociation of water, the adsorption and reduction of heavy metals, or ammonia synthesis, to mention only a few. In order to understand and control the reactivity of these surfaces and/or tailor their electronic and magnetic properties, it is indispensable to know the surface structure and morphology at ambient conditions. Hence, a major problem in modern surface science is that the vast majority of experimental techniques are restricted to UHV (ultra high vacuum). Furthermore, the insulating nature of most of the oxides limits the use of imaging techniques like scanning tunneling microscopy. Quantitative diffraction analyses for oxide surfaces are scarce and by far less satisfactory than for metal or semiconductor surfaces.

On the side of theoretical modeling, density-functional theory (DFT) in the past decades has established itself as a powerful tool for prediction of the properties and stability of technologically relevant materials. Last but not least this was recognized by the Nobel Prize for Chemistry given to Walter Kohn in 1998 for the development of DFT. However, this approach is restricted to T=0 and p=0. Only recently it has become possible to extend the predictive power of DFT to finite pressures and temperatures by combining DFT with concepts from thermodynamics in the framework of ab initio atomistic thermodynamics.

A further major challenge for the quantum mechanical modeling of mineral surfaces consists in the large system sizes due to their complex structure, which makes such calculations computationally extremely demanding.

Magnetite is one of the most important minerals in geophysics and mineralogy since it records information on the earth's magnetic field, and plays a role in the orientational capablilites of microorganisms and birds. Moreover, it has recently attracted a lot of attention as a potential material for use in spintronic devices. Still, the surface termination of magnetite (001) has been subject of a controversial debate in the literature. We performed systematic DFT-calculations employing the full-potential linearized augmented-plane wave method as implemented in the WIEN2k code for a variety of stochiometric and non-stoichiometric terminations, including a full structural optimization of each system. The surface is modeled by supercells containing on the average 100 atoms and 1000 electrons resulting in matrix dimensions for the generalized eigenvalue problem of 18000x18000. Due to the high computational cost such calculations have only recently become feasible through the implementation of fine-grain parallelization schemes. After porting and optimizing the fine-grain parallel version of the WIEN2k code in-collaboration with the Leibniz Computing Center Munich (LRZ) on the Hitachi SR8000, we were able to achieve a performance of 5 GFlops per 8-processor-node (640 MFlops per CPU). This amounts to 42% of the theoretical peak performance of the machine and is a very good result for this type of code, making it one of the best performing codes on the SR8000.

It is typically assumed that the excess charges on a polar surface can be compensated by surface reconstructions. The latter are commonly understood as the ordering of surface vacancies. Contrary to this, the surface phase diagram of Fe_3O_4(001) computed within the framework of ab initio atomistic thermodynamics [cf. Figure 1] reveals that a modified bulk termination, which was hither to ignored on the basis of simple electrostatic arguments, is the most stable termination over a broad range of oxygen pressures. The results indicate a novel stabilization mechanism on polar oxide surfaces, where the surface periodicity is achieved by a wave-like distortion of the surface layer. This unusual stabilization mechanism is accompanied by dramatic changes in the electronic properties e.g., a half-metal to metal transition from bulk to the surface which may be relevant for future applications. Experimental results obtained by scientists in Munich, Aachen and Konstanz support the theoretically predicted model. The close collaboration between theory and experiment plays an important role in this project: For example, the structural information from DFT is used as a starting point for quantitative diffraction analyses and vice versa. Current and future investigations include the influence of correlation effects beyond the generalized gradient approximation (GGA) as well as the adsorption of water and the impact of a humid atmosphere on the surface stability and electronic properties of oxide surfaces. Another aspect of relevance for technological applications are novel charge and magnetic order states realized in oxide interfaces. The recent resource extension at LRZ (e.g., the newly installed SGI Altix 4700) will be invaluable to tackle these numerically intensive and challenging problems.

• Rossitza Pentcheva

Department für Geo- und Umweltwissenschaften, Universität München, Germany

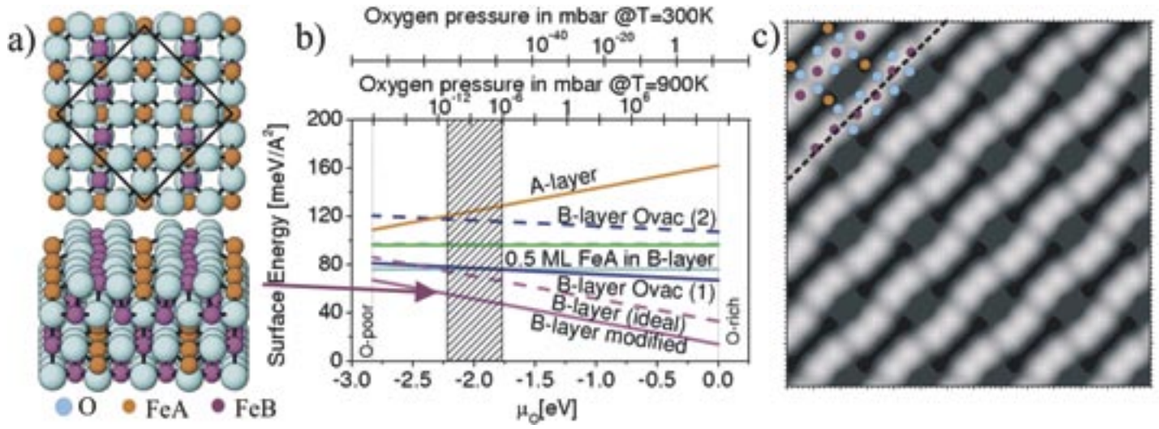

Figure 1: (a) Jahn-Teller distorted modified bulk termination of Fe_3O_4(001) predicted as lowest energy configuration over a broad range of oxygen pressures from the calculated surface phase diagram (b); (c) STM-simulation of the surface layer

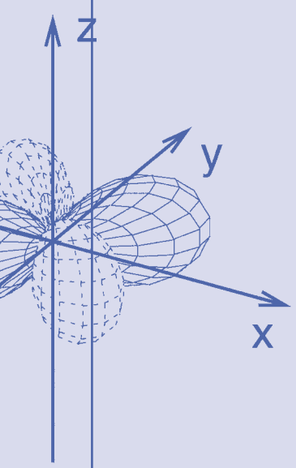

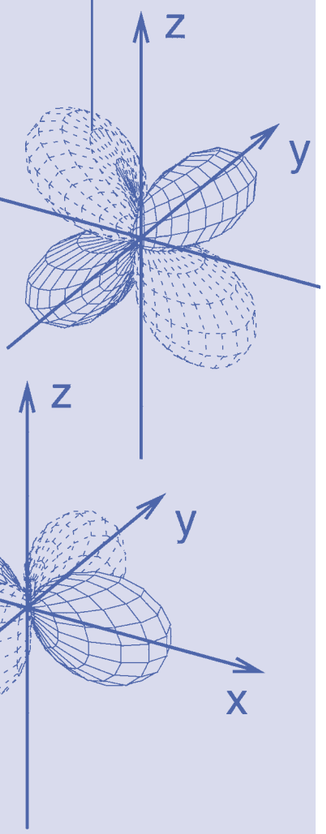

New Horizons for the Realistic Description of Materials with Strong Correlations

Electronic structure theory is the basis of modern technologies such as electronics and computing. Electronic properties of materials are determined by quantum mechanics. Thus, by solving the Schrödinger equation, we should be able to predict the properties of real materials, or even design new ones with superior qualities. Unfortunately, solving this equation is not easy at all. The essential complication comes from the inherent quantum many-body nature of the problem. As a result, a brute-force solution is impossible, except in the simplest cases. As an illustration let us consider a single atom of iron. Having 26 electrons, its wave function is a function of 26 times 3 coordinates. Neglecting spin, already an extremely crude representation of this function at merely 10 values of each variable would thus require storage of 10^{78} numbers. Even after reducing this number by exploiting symmetries, there is simply not enough matter available in our galaxy for building the required memory.

Given this example, electronic structure theory seems a hopeless enterprise. Nevertheless, it is a thriving discipline. This is largely due to density functional theory. In practice, this approach drastically simplifies the many-body problem by assuming that the electrons retain their individuality and experience the other electrons via a static mean field. In this picture electrons occupy states that extend over the whole crystal, forming the band structure of the material.

For many important classes of materials a density functional description fails, however, even qualitatively. Striking effects like the breakdown of the Fermi-liquid picture at the Mott metal-insulator transition, heavy Fermion behavior, exotic one-dimensional Luttinger phases, or high-temperature superconductivity cannot be addressed by such a simple approach. All these materials are strongly correlated. This means that the repulsion between the electrons is so strong that the electrons lose their individuality, and the single-particle picture breaks down. Because of the strength of the interaction non-perturbative many-body techniques have to be used, so that powerful computers are essential for reliable calculations. And still, calculations are restricted to quite small model systems. This means that the full Hamiltonian of a crystal has to be approximated by a small lattice Hamiltonian, which describes only (few of) the strongly correlated electrons. All other electrons have to be included in the calculation in an average way.

The modern approach to solving the many-body problem is dynamical mean-field theory (DMFT). It reduces the lattice Hamiltonian to a correlated impurity embedded in a self-consistent dynamical medium, which mimics the other lattice sites. This approximation simplifies the problem significantly. Still, a complicated quantum impurity problem remains, which has to be solved, e.g., with quantum Monte Carlo (QMC) or the Lanczos method. With DMFT, it was possible, for the first time, to understand the physics of the Mott transition. In a Mott insulator the electronic band structure loses its meaning. Instead, physics becomes more local and it is more appropriate to think about the electrons as occupying atomic-like orbitals.

In this strongly correlated regime we find a number of fascinating ordering phenomena. Most well known is antiferromagnetism, where spins on neighboring lattice sites point in opposite directions. When there are many correlated orbitals, a similar ordered phase can exist: occupied orbitals on neighboring sites point in different directions, as illustrated in Figure 1. This directionality can give rise to highly anisotropic transport properties. Coupling of spin and orbital degrees of freedom can make transport properties strongly dependent on magnetic fields. Such a mechanism is believed to be the basis of the colossal magneto-resistance effect (CMR). Like the giant magneto-resistance (GMR), this effect, once understood, holds the promise of, e.g., another vast increase in hard-disk capacity.

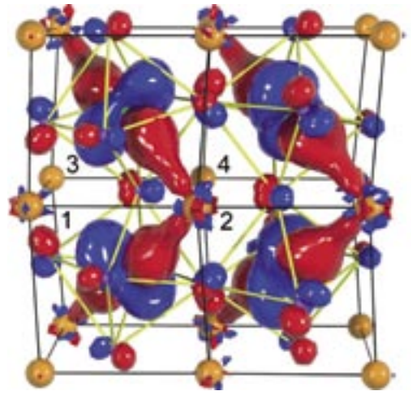

Figure 1: Orbital ordering in the Mott insulator $LaTiO_3$. The displayed Wannier orbitals are the occupied states in this system, as obtained from NMTO+DMFT calculations using quantum Monte Carlo

Complicated spatial patterns like orbital-ordering, however, can not be described by a single-site approach such as DMFT, which assumes that all lattice sites are equivalent. In order to add the required spatial degrees of freedom the single impurity of DMFT has to be replaced by a cluster of sites. This approach is accordingly called cluster DMFT (CDMFT). Unfortunately, treating a cluster instead of a single site increases the already high computational cost of a calculation even further: the required CPU time rises (at least) as the third power of number of sites in the cluster.

With an efficient parallelization of the QMC solver and of the DMFT self-consistency loop, we can however exploit the spectacular increase in performance offered by massively parallel machines like the new Blue Gene/L system in Jülich called JUBL. Finally it is possible to reach reasonably low temperatures, where so far calculations of low temperature physics often had to be done at about 1000 K. For lack of computer time, uncontrolled approximations had to be introduced in the model Hamiltonians. Now it is possible to check these approximations by explicit calculations. In short, calculations are becoming significantly more reliable and thus gain predictive power. Alternatively, it is now also possible to go beyond the single-site approximation

of DMFT and study, e.g., the physics of orbital ordering. In the foreseeable future it should even be possible to combine these two advances and simulate realistic Hamiltonians using reasonably large clusters at the experimental temperatures.

It might not be too surprising that Blue Gene/L is very suitable for large Monte Carlo simulations, as the communication when taking statistics is somewhat limited. But also the second main DMFT-solver, the Lanczos method, can benefit from the new architecture. This might, at first, be unexpected: In the Lanczos method, the full ground state vector of a many-body system is handled. Thus the method is limited by the available main memory. The principal problem for a distributed memory implementation is that the central routine of a Lanczos code, the application

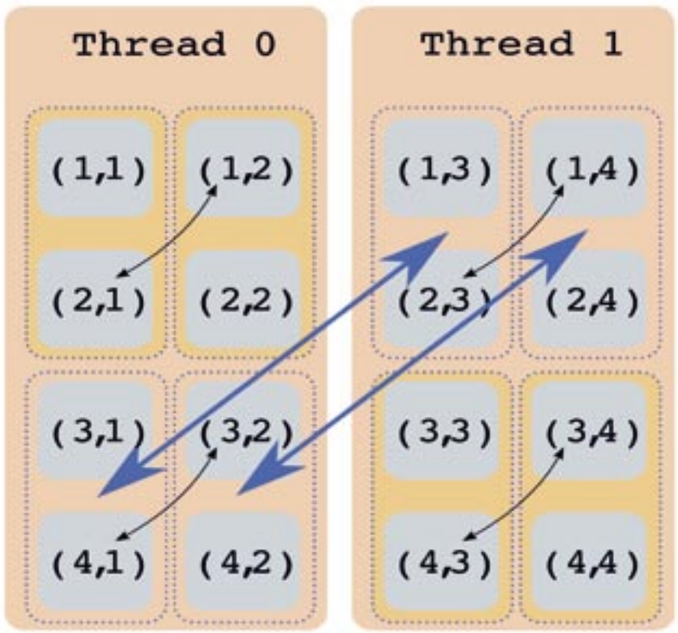

Figure 2: Scheme of the transpose operation that makes memory access thread-local when calculating the operation of the Hamiltonian on a state-vector. The communication (blue arrows) is realized by a call to MPI_Alltoall. The small black arrows indicate the local operations needed to complete the matrix transpose

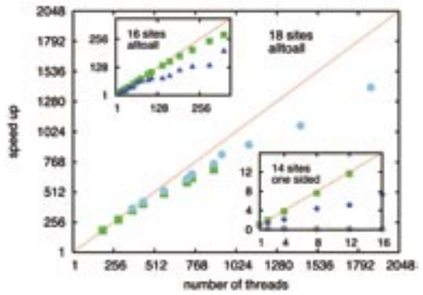

Figure 3: Speedup of our Lanczos code on IBM Blue Gene/L JUBL (green (CO mode) and turquoise (VN mode) symbols) and IBM Regatta JUMP (blue and grey symbols) for different problem sizes

of the Hamiltonian to the many-body state leads, due to the kinetic energy term, to very non-local memory access patterns. Thus, a naive implementation, using one-sided communication to access the required vector elements, gives extremely poor performance, even a speed-down (see lower right panel of Figure 3). We can, however, create an efficient MPI implementation by using a simple but important observation: in the kinetic term of the Hamiltonian the electron-spin is not changed. Thus, writing the many-body vector as a matrix $v(i_\uparrow,i_\downarrow)$, where the indices label spin-configurations, we find that the hopping term only connects vector elements that differ in one index. Hence, storing entire slices $v(i_\uparrow:)$ on one processor, the kinetic term for the spin-down electrons is local to that thread. After transposing v, the same is true for the hopping of the spin-up electrons. Therefore, the efficient implementation of the sparse-matrix-vector product, which is central to a Lanczos code, depends on the performance of the matrix transpose, which can be implemented by MPI_Alltoall. And, as can be seen from Figure 3, global communication is indeed very efficient on Blue Gene/L: the main plot shows the speedup for a calculation, where, in each iteration, a state vector of about 18 GB has to be moved across the machine. Comparing with the IBM Regatta system JUMP in Jülich, it is interesting to note that per CPU the code runs only about twice as fast on JUMP than on JUBL, despite the difference in clock speeds. In addition, JUBL shows a better speedup when going to larger numbers of processors (see upper left panel of Figure 3). Finally, speedup is with the number of CPUs, such that the code can fully exploit the second CPU on each node (virtual node mode). It thus turns out that Blue Gene/L is not only the ideal machine for running large Monte Carlo simulations but is also extremely well suited for the Lanczos method, which can efficiently take advantage of the very large, however distributed, memory.

With these methods, JUBL opens the path to many-body calculations of unprecedented complexity. This will lift most of the limitations that usually force us to make uncontrolled approximations when constructing model Hamiltonians, enabling many-body physics to leap into the real word.

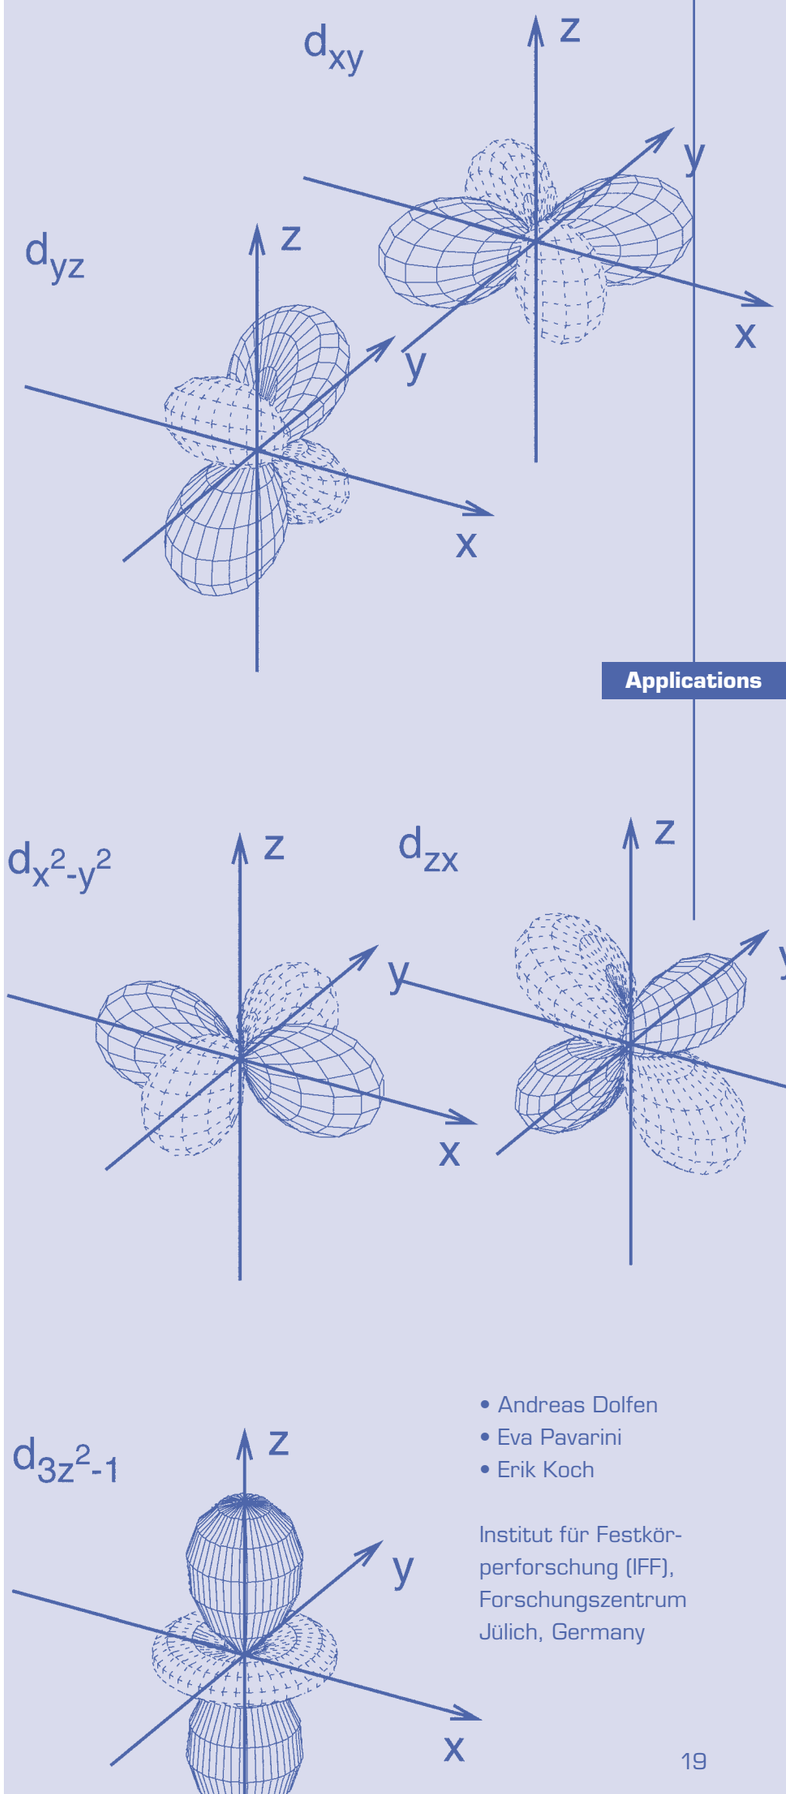

- Andreas Dolfen
- Eva Pavarini
- Erik Koch

Institut für Festkörperforschung (IFF), Forschungszentrum Jülich, Germany

InGrid – Innovative Grid Developme nts for Engineering Applications

On September 1, 2005 the InGrid project started under the leadership of HLRS. Its aim is to connect engineers in a German wide Grid. The InGrid project is part of the German D-Grid initiative connecting computers, software, data and scientists all over Germany.

Modern information society is faced with a multitude of possibilities for efficient communication and easy access to large sets of data and information as well as high performance computing systems. The opportunities in science and industrial development have increased dramatically. At the same time, complexity and thus difficulties in handling and use of distributed and dynamic systems have increased dramatically. Therefore, the German Ministry for Research and Education (BMBF) is funding in the D-Grid initiative the development and provision of an integrated working environment, which permits the transparent use of a broad spectrum of IT achievements for research and development in science and industry.

InGrid and D-Grid

D-Grid consists of a Grid infrastructure project and currently 6 community projects in the areas of engineering sciences (InGrid), medicine (MediGrid), climate research (C3 Grid), high-energy physics (HEP Grid), astrophysics (AstroGrid-D), and text-based disciplines (TextGrid). In the engineering community project InGrid, a Grid environment will be developed to enable modeling, simulation, and optimization of engineering applications from areas such as foundry technology, mental forming, groundwater flows, turbine simulation, and fluid-structure interaction.

InGrid is focusing on the question, how engineering work can be supported by Grid technology. InGrid builds on the idea that the abilities of man and machine have to be engaged in a most synergetic way. In particular, adequately interlaced Grid-based engineering has a substantial significance for future technical products and product development processes. The development of a Grid based computational engineering community demonstrates technologically advanced applications of engineering work in research as well as in operational innovation management.

The community project InGrid will enable engineering projects for Grid-based application and efficient use of common compute and software resources. Grid-Computing (the word is taken from the electrical power grid) means distributed computing, not only in a local network, but across system- and software borders. Researchers and developers should be able to work on complex scientific questions, independent of their local environment of computers, programs, data and information. InGrid provides a Grid environment for scientific as well as industrial engineering applications. The flexible use of Grid technologies will combine the competences in modeling, simulation and optimization and allow for the common, efficient use of distributed resources.

Work Packages

Five typical applications – foundry technologies, metal forming technologies, groundwater flow and transport, turbine simulation and fluid-structure interaction -- are considered as showcases in order to cover the three central areas of computationally intensive engineering applications, that are coupled multi-scale problems, coupled multi-discipline problems, and distributed simulation-based optimization. The project is concerned with basic research as well as application oriented research to establish a Grid-based, networked software infrastructure for the support of scientific engineering work.

The problem categories in engineering today are optimal design, optimal system control, design of hybrid systems, and optimization of multi-level production systems. These categories are subtasks of virtual prototyping and collaborative engineering. By use of Grid technologies, necessary resources such as processors, memory systems, software tools, information systems and quick and safe network connections are combined to fit the demand for co-operation between researchers, developers, manufacturer, providers, and users. The virtual prototyping is supported by internet-based collaborative methods, mapping of engineering workflows on a Grid environment, integration of specific technical data sources, and building of portals for selected engineering scenarios. Today's usual development and research processes overcome the classical borders of engineering and economical-organizational ranges in the sense of the product life cycle management. Thus, knowledge-based systems for the optimization of engineering workflows and support of decision-making processes are to be developed as well as co-operation, business, safety and confidence models for scientific engineering applications.

Organizational Aspects

InGrid is carried out by 8 partners from academies and industry. The partners are Institute of Fluid Mechanics and Hydraulic Machinery (IHS) – HLRS, University of Stuttgart, Department of Mathematics and Computer Science – Philipps-University Marburg, Institute of Business Information Technology – University of Siegen, Institute of Scientific Computing and Algorithms (SCAI) – Fraunhofer-Gesellschaft, Access e.V. Materials + Processes – RWTH Aachen, WASY Gesellschaft für wasserwirtschaftliche Planung und Systemforschung mbh – Berlin, and T-Systems – Solutions for Research GmbH (SfR).

More information about the InGrid project can be found on the web-site http://www.ingrid-info.de. Most of the information provided here was taken from this webpage. Information on the D-Grid initiative in general are available at http://www.d-grid.de

• Sabine Roller

Höchstleistungs-rechenzentrum Stuttgart, Germany

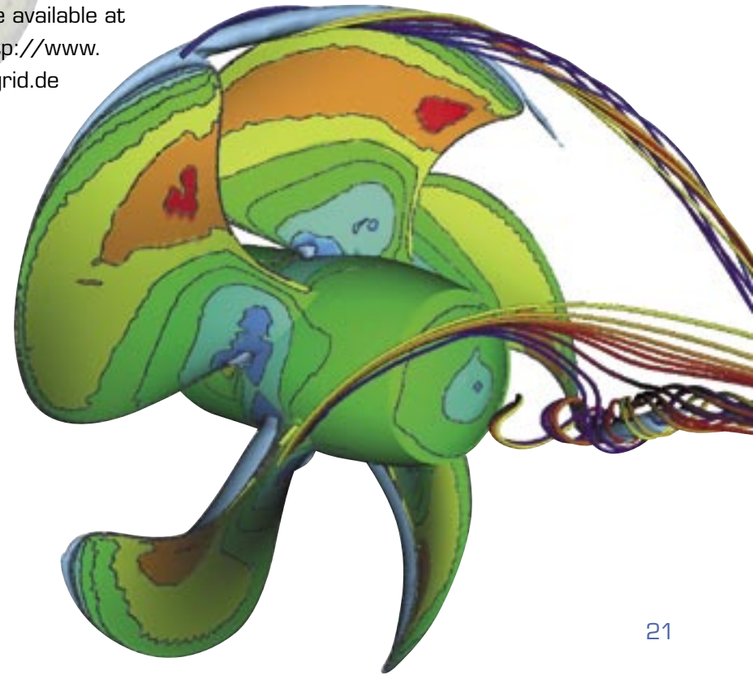

Projects

BEinGRID – Business Experiments in Grid

A large number of research projects exploring the field of Grids have produced a wide variety of different solutions. Very often the solutions targeting a specific application or problem domain provided a highly specialized middleware. This has lead to a diversity of approaches for "Grids". The research community has realized with projects like EGEE or DEISA an infrastructure across Europe targeting mainly scientific users. In parallel to this development projects like Grid for Industrial Applications (GRIA) or Grid based Application Service Provision (GRASP) have delivered a more business oriented flovour of Grids. Key additions are the support of Service Level Arguments (SLA) and Cross Organisational Accounting.

The major paradigm shift in the Grid community moving from proprietary protocols and solutions to Service Oriented Architectures (SOA) and Web Services as basis for the realization of the Open Grid Service Architecture (OGSA) concept has been taken up by a range of European Research Projects aiming for the so called Next Generation Grid (NGG). These projects such as SIMDAT, Akogrimo and NextGrid while sharing common concepts and underlying basis middleware are addressing again only a small variety of application domains and provide solutions or platforms that are not prepared to be taken up by a broad application community in an easy way.

Grid technology has reached a critical phase in its adoption. So far most Grid toolkits are only usable for Grid experts or application developers or users with decent knowledge in distributed computing. In order to make Grid the tool for eScience and industry that it is promised to be, the necessary basis for an easy take up of Grid solutions by service providers e.g. for computational or data resources but also for commercial service providers integrating their solutions and applications into an infrastructure a further consolidation of Grid middleware is needed.

BEinGRID – Business Experiments in Grid

The mission of BEinGRID initiative is to Exploit European Grid middleware by creating a toolset repository of Grid services from across the Grid research domain and to use these services to deliver a set of successful business experiments that stimulate the early adoption of Grid technologies across the European Union. The main objective of the project consists in fostering the adoption of the called "Next Generation Grids" by means of the realization of 19 business experiments in key sectors such as multimedia, finance, engineering, chemistry, gaming, environmental science, logistics, etc on different Grid middleware solutions (GTv4, UNICORE/GS, gLite, GRIA, WS-* on Axis Apache). In short, a business experiment is a real Grid pilot application that addresses a concrete business case and in which the main actors of the economic sector are represented: end-user, service provider, Grid service integrator. The involvement of the full value chain is considered a catalysing aspect to demonstrate Grid potential and to capitalize on the derived benefits and to convince the potential early adopters about the benefits of using this leading edge technology.

BEinGRID will design and set up a toolset repository of Grid middleware, using existing research results of state of the art technology and toolkits as initial input, developing necessary adaptations of missing pieces (glue software), for empowering Grid take-up.

BEinGRID will catalyse European business adoption of Grid services and lead to widespread take up of these important technologies for the benefit of all.

Selected Business Experiment – Environmental Science

The BEinGRID project covers 19 business experiments that are realized on the basis of different middleware and that are steered by a group of core project partners that are aiming to provide consultancy and guidance and

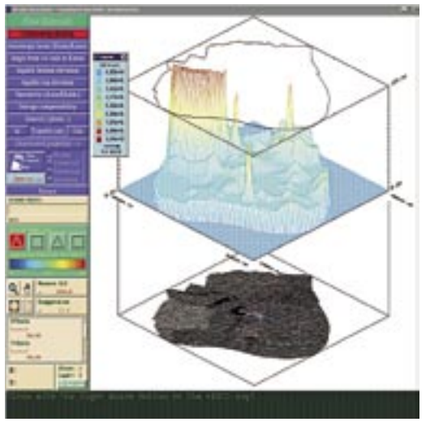

Figure 1a

Figure 1b

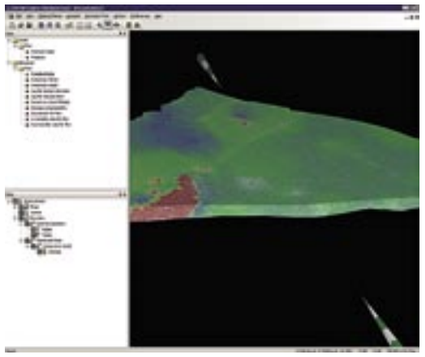

Figure 1c

Conductivity parameter distribution in the FEFLOW GUI for 2D (Figure 1a) and for 3D (Figure 1b) and using FEFLOW Explorer, the free 3D-FEFLOW-file-viewer, in a cut view (Figure 1c)

develop the necessary glue to build from the solutions from the experiments valuable repository artefacts.

One of the experiments is coordinated by the High Performance Computing Center (HLRS) in collaboration with other partners from Germany namely University Siegen, WASY GmbH, and the Linksniederrheinische Entwässerungsgenossenschaft and one end-user from China the Beijing Water Authority (BWA). The goal of this experiment is the realization of a large scale multidisciplinary compute grid, enabling the end user co-operation partner to generate cost-effective and optimized solutions even for complex problems of water management, addressing both business and society aspects in the sector of environmental engineering. From the business point of view it will be possible to perform more complex computing tasks, directly influencing the decision-making process by providing better forecast possibilities for the environmental state, which on the other hand leads to a decision directly affecting the environment. The grid based infrastructure to be developed will provide a cutting-edge technology with its benefits in business and society and will have sustainable great value. Not only will the basic services for simulation and optimization be provided by the grid infrastructure. Grid based backup-mechanisms and data stores which have to be accessed are mandatory for a reasonable utilization and therefore will have to be developed and integrated as well. The research activities will have to face the subjects of persistence and reliability, as well as global access to the grid infrastructure and the development of new business models for collaboration in a multidisciplinary international grid-economy. By reducing total cost of ownership, administrative overhead and computation time for the end user co-operation partner, the experiment and its results will have an impact on both the economic and ecologic sector of water management. The final objective in providing this infrastructure is to establish a sustainable economic relationship between end users and service providers, thus leading to a better position in protecting the environment and using resources more efficiently. The project will be based on UNICORE or its successor UNICORE/GS.

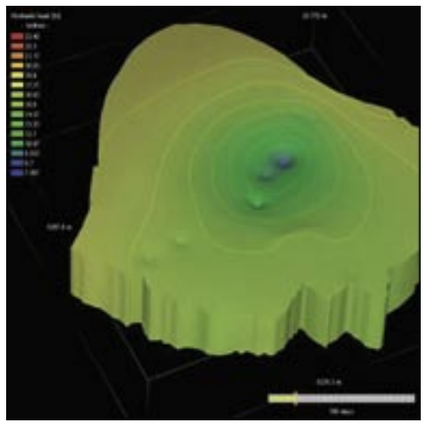

Figure 2a

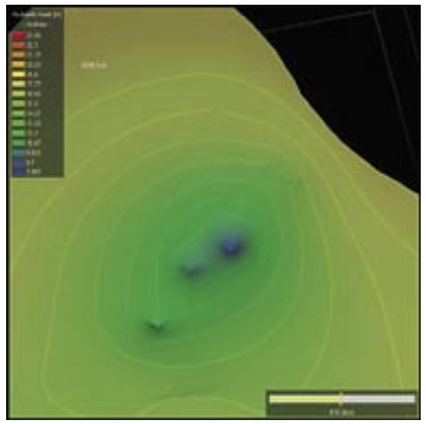

Figure 2b

Model results (water table) in a far (Figure 2a) and in a zoom view (Figure 2b). Both screenshots were taken from a avi-file showing transient water table development.

Approach

The key question is how the core project of BEinGRID is establishing co-operation with the large number of independent business experiments. A problem that is somehow similar to the organization of the D-Grid initiative where a supporting project is collaborating with a range of community projects. In BEinGRID this collaboration is organized that the core project is setting up a development infrastructure and provide ground work on best practices and toolkits for the different experiments. During the execution of the project the support is organized in two different so called “cross-activities” in the technical and business area. The technical cross activities are structured in 7 main S&T areas namely trust & security, architecture and interoperability, service management, data management, VO management and portals. These activities will analyze and steer all experiments and drive the process for taking up the results from the business experiments and consolidate them by developing glue components as part of the components development activity.

Project Facts

The BEinGRID project has started in June 2006 and is realized by a consortium of 75 partners across Europe with an overall budget of 24,7 M€. The project includes 13 partners from Germany covering universities, research centers and industrial partners. Three partners from Germany, Fraunhofer SCAI, T-Systems and HLRS are member of the Core Team of the project. Further information can be found at http://www.beingrid.com

• Stefan Wesner

Höchstleistungs-
rechenzentrum
Stuttgart, Germany

Projects

Leibniz Computing Center of the Bavarian Academy of Sciences (Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften, LRZ) in 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

View of the Hitachi SR8000-F1 at LRZ

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

System	Size	Peak Performance (GFlop/s)	Purpose	User Community
SGI Altix 4700 **(installation: June 2006)**	4096 processors 17.5 TByte	26,200	Capability computing	German universities and research institutes
Hitachi SR8000-F1 8-way **(until June 2006)**	168 nodes 1344 processors (+168 service procs) 1.3 TByte memory	2,016	Capability computing	German universities and research institutes
SGI Altix 64-way	64 processors 256 GByte memory	410	Tests and porting	German universities and research institutes
SGI Altix 128-way	128 processors 512 GByte memory	820	Capability computing	Bavarian universities
Linux Cluster Intel IA64 2-way	68 nodes 136 processors 816 GByte memory	870	Capability and capacity computing	Bavarian universities
Linux Cluster Intel IA64 4- and 8-way	19 nodes 84 cores 250 GByte memory	443	Capacity computing	Munich universities
Linux cluster Intel IA32 Intel&AMD EM64T	154 nodes 192 processors 320 GByte memory	850	Capacity computing	Munich universities

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 center 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 Center of the University of Karlsruhe and the Center 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
70500 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 (GFlop/s)	Purpose	User Community
NEC SX-8	72 8-way nodes 9,22 TB memory	12,670	Capability computing	German universities, research institutes, and industry
TX-7	32 way node 256 GByte memory	192	Preprocessing	German universities, research institutes, and industry
Intel Nocona Cluster	205 2-way nodes 240 GB memory	2,624	Capability and capacity computing	Research institutes, and industry
Cray Opteron	129 2-way nodes 512 GByte memory	1,024	Capability and capacity computing	Research institutes, and industry
Cray XD1	8 12-way nodes 96 GByte	500	Industrial development	Research institutes, and industry

The John von Neumann Institute for Computing (NIC) is a joint foundation of Forschungszentrum Jülich, Deutsches Elektronen-Synchrotron DESY, and Gesellschaft für Schwerionenforschung GSI to support supercomputer-aided scientific research and development. Its tasks are:

Provision of supercomputer capacity for projects in science, research and industry in the fields of modelling and computer simulation including their methods. The supercomputers with the required information technology infrastructure (software, data storage, networks) are operated by the Central Institute for Applied Mathematics (ZAM) in Jülich and by the Center for Parallel Computing at DESY in Zeuthen.

Supercomputer-oriented research and development in selected fields of physics and other natural sciences, especially in elementary-particle physics, by research groups of competence in supercomputing applications. At present, two research groups exist: the group Elementary Particle Physics, headed by Karl Jansen and located at the DESY laboratory in Zeuthen and the group Computational Biology and Biophysics, headed by Ulrich Hansmann at the Research Center Jülich.

Education and training in the fields of supercomputing by symposia, workshops, school, seminars, courses, and guest programmes.

The IBM supercomputers "JUBL" (top) and "JUMP" (bottom) in Jülich (Photo: Research Center Jülich)

The following supercomputers are available for research projects of the communities mentioned below, evaluated by the Peer Review Board of NIC. A more detailed description of the supercomputers can be found on the web servers of the Research Center Jülich and of the German Electron Synchrotron DESY, respectively:
www.fz-juelich.de/zam/nic/en
www-zeuthen.desy.de/main/html/home

System	Size	Peak Performance (GFlop/s)	Purpose	User Community
IBM Blue Gene/L "JUBL"	8 racks 8192 nodes 16384 processors PowerPC 440 4 TByte memory	45,875	Capability computing	German universities, research institutes, and industry
IBM pSeries 690 Cluster 1600 "JUMP"	41 SMP nodes 1312 processors POWER4+ 5,1 TByte memory	9,000	Capability computing	German universities, research institutes, and industry
apeNEXT (special purpose computer)	4 racks 2048 processors 512 GByte memory	2,500	Capability computing	Lattice gauge theory groups at universities and research institutes
APEmille (special purpose computer)	4 racks 1024 processors 32 GByte memory	550	Capability computing	Lattice gauge theory groups at universities and research institutes

Contact
John von Neumann Institute for Computing (NIC)
Central Institute for Applied Mathematics (ZAM)

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

High Performance Computing Courses and Tutorials

LRZ
www.lrz.de

Iterative Linear Solvers and Parallelization

Date

October 9-13, 2006

Location

New LRZ building, Munich/Garching

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, BiCG-STAB ...) as well as highly efficient preconditioning techniques are presented in the context of real life applications. Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the basic constructs of iterative solvers, the Message Passing Interface (MPI) and the shared memory directives of OpenMP. This course is organized by University of Kassel, HLRS, IAG, and LRZ.

Webpage

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

NIC
www.fz-juelich.de/nic

User Course The IBM Supercomputers JUMP and JUBL in Jülich: Programming and Usage

Date

August 10-11, 2006

Location

NIC/ZAM, Research Center Jülich

Contents

This course gives an overview on the IBM supercomputers JUMP and JUBL in Jülich. Especially new users will learn how to program and use these systems efficiently. The following topics are discussed in detail: system architecture, usage model, compiler, tools, monitoring, MPI, OpenMP, performance optimization, mathematical software, and application software.

Webpage

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

Education in Scientific Computing

Date

August 7-October 13, 2006

Location

NIC/ZAM, Research Center Jülich

Contents

Guest Students' Programme "Scientific Computing" to support education and training in the fields of supercomputing. Application deadline was April 30, 2006; the event is already fully booked.

Webpage

http://www.fz-juelich.de/zam/gaststudenten

Parallel Programming with MPI, OpenMP, and PETSc

Date

November 27-29, 2006

Location

NIC/ZAM, Research Center Jülich

Contents

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

Presented by Dr. Rolf Rabenseifner, HLRS

Webpage

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

CECAM Tutorial Programming Parallel Computers

Date

January 22-26, 2007

Location

NIC/ZAM, Research Center Jülich

Contents

This tutorial provides a thorough introduction to scientific parallel programming. It covers parallel programming with MPI and OpenMP. Lectures will alternate with hands-on exercises.

Webpage

http://www.cecam.fr/indexphp?content=activities/tutorial

HLRS
www.hlrs.de

Parallel Programming with MPI, OpenMP and PETSc

Date

August 23-25, 2006

Location

Manno (CH), CSCS

Contents

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

Webpage

www.hlrs.de/news-events/external-events

Introduction to Computational Fluids Dynamics

Date

September 18-22, 2006

Location

Stuttgart, HLRS

Contents

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

Webpage

www.hlrs.de/news-events/events

Parallel Programming with MPI, OpenMP and PETSc

Date

September 25-27, 2006

Location

Stuttgart, HLRS

Contents

The focus is on programming models MPI-1, OpenMP, and PETSc. It includes also an overview on MPI-2. 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 if required.

Webpage

www.hlrs.de/news-events/events

Advanced Topics in Parallel Programming

Date

September 28-29, 2006

Location

Stuttgart, HLRS

Contents

Topics are MPI-2 parallel file I/O, OpenMP tools and tuning, hybrid mixed model MPI+OpenMP parallelization, domain decomposition of structured and unstructured grids and with particle based applications, parallel numerics, object oriented parallel programming with C++, and Grid computing. Course language is English if required.

Webpage

www.hlrs.de/news-events/events

Fortran for Scientific Computing

Date

October 3-13, 2006

Location

Stuttgart, HLRS

Contents

This introduction to Fortran is taught with lectures and hands-on sessions. This course is organized by HLRS and Department for Computational Physics/ University of Stuttgart.

Webpage

www.hlrs.de/news-events/events

New IBM Blue Gene/L system JUBL inaugurated at NIC in Jülich

On March 6, 2006 a new era in supercomputing has been started in Jülich: The German Federal Ministry for Education and Research (BMBF) Undersecretary of State Thomas Rachel, together with North Rhine-Westphalia Innovation Minister, Professor Andreas Pinkwart inaugurated the fastest supercomputer in Europe with an output of 46 TeraFlop/s (trillions of floating point operations per second). Among the purely scientifically oriented research establishments the Research Center Jülich then stood at first place worldwide.

The new supercomputer JUBL (Jülich Blue Gene/L) was built by IBM. Its new technology is considered as key to fundamental insights in biology, chemistry, physics and climatic research. For materials scientists, nano-technologists and energy researchers Blue Gene opens windows to completely new applications.

Simulations with supercomputers have become indispensable for research and development, and in particular interdisciplinary questions. JUBL strengthens this key authority of the Research Center Jülich and opens a new dimension in computing for physics, material research and nano-technology. In addition, life sciences, particularly biology and medicine as well as environmental research, for instance the analysis of pollutant propagation in soil and the atmosphere, are main application areas of the Jülich computers. Through the John von Neumann Institute for Computing (NIC), more than 200 European groups of researchers use the computer capacity in the Research Center Jülich. Whether climatic and earth system research or biophysics and bio computer science: Many computations for science and applications are already today "Made in Jülich".

The NIC is pursuing a dual future computing concept. The new Blue Gene towers in the modern Jülich machine room are supplemented by JUMP, the already existing supercomputer, a constellation of 41 IBM p690 shared memory systems. Since applications in the future will rely on JUBL for the highest computing requirements, users will now have more opportunity to utilize JUMP for their data intensive problems.

NIC Winter School and NIC Symposium

Continuing the tradition of the bi-annual Winter Schools on topics in the computational sciences, preferably in Theoretical Chemistry, NIC hosted the School "Computational Nanoscience: Do It Yourself!" from February 14 to 22, 2006. The scientific programme of this School was directed to PhD students and PostDocs from Computational Materials Science and Chemistry. Although the number of participants was restricted to 50 due to limited computer lab resources, more than 90 applications from 22 mainly European countries were submitted, which required a strong selection process. The lectures and hands-on sessions were concentrated mainly on the understanding and handling of the four application codes

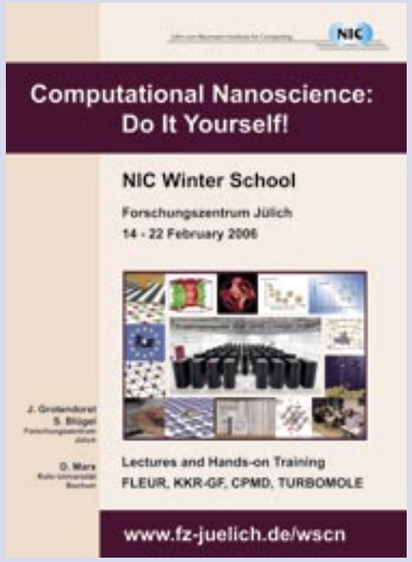

FLEUR, KKRGF (Solid State and Surface Physics), CPMD (chemical reactions and dynamics), and TURBOMOLE (structure and properties of large molecules and clusters). Lecture Notes of the School have been published as Volume 31 of the NIC Series, ISBN 3-00-017350-1.

Activities

The 3rd NIC Symposium was held at the Research Center Jülich from March 1 to March 3, 2006. It was attended by more than 100 scientists from the different research disciplines, which use the supercomputers in Jülich for their computational studies. Together with the meeting an extended proceedings volume (NIC Series Volume 32, ISBN 3-00-017351-X) was published, which provides an overview on the broad spectrum of computational science projects from Astrophysics, Chemistry, Elementary Particle Physics, Materials Science, Condensed Matter Physics, Polymers, Earth and Environment, Computer Science and further topical research areas.

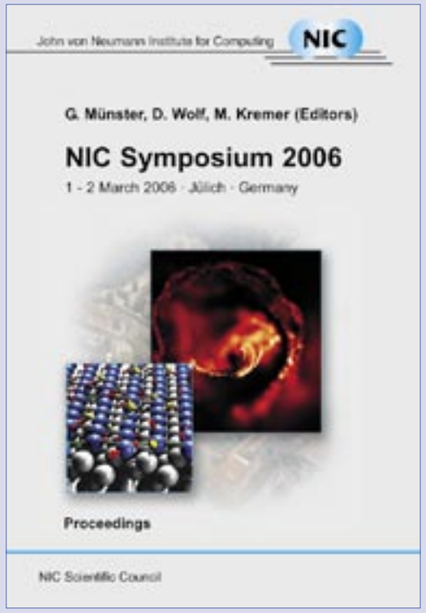

Both volumes are available either as hard copies (please send corresponding order requests to nic@fz-juelich.de) or as PDF or HMTL files on the web (http://www.fz-juelich.de/nic-series/volume31 or volume32).

Jülich Welcomes Computational Scientists from Central Europe

Scientists and science managers from Central Europe and Germany met for a two-day workshop in Jülich in January 2006, in order to discuss opportunities of collaboration with German scientists and to assess the requirements of the computational sciences in a larger Europe. Strengthening this important research area is seen as a key component of the ongoing international scientific competition and has a strong influence on the economic development of a country. Whereas in the new EU countries much less high-end computing facilities are available than in Western Europe, many renowned research groups there need such systems for their simulations in physics, chemistry, biology, or engineering.

"We want the brightest brains in Europe to have access to the best resources for their research", Thomas Lippert, Director of the John von Neumann Institute for Computing (NIC), told the participants from Poland, the Czech Republic, the Slovak Republic, Hungary, Estonia, Cyprus, and Germany in his welcome address. "A big common effort in the area of supercomputing is needed to maintain Europe's position in the growing competition with the USA and Japan."

The Research Center Jülich, as part of its promotion of the computational sciences, wants to support these groups and – in the framework of the NIC – offers them 600.000 processor hours per year on its JUMP supercomputer. The allocation of computing time follows the same procedure as for German users and requires a successful peer review of a scientific proposal. The European science network meanwhile is powerful enough to provide sufficiently fast access also from the new EU member states.

During the workshop, interested groups had the opportunity to get first-hands information about the technical prerequisites and organizational access conditions. The workshop will contribute to the continually increasing number of NIC users from outside Germany, in particular from the new EU member states, and underlined the need for a powerful European supercomputing infrastructure.

LRZ has moved to its New Location

Leibniz Computing Center has moved to its new location on the Garching Campus, 10 km north of Munich. The new LRZ building consists of three sections. The compute cube is a 36x36x36 meter building which contains the new high end system, the Linux cluster, servers, networking devices, backup and archiving facilities, and the air conditioning and power supplies. The institute building provides office space for more than 150 staff. An auditorium and lecture rooms can be found in the third section.

While the relocation of all other systems has been completed, the installation of the new high-end computing system at LRZ will start in June. The new SGI Altix 4700 will have a peak performance of 26.2 trillion floating point operations per second (26.2 TFlop/s) and is comprized of 4,096 Intel Itanium2 processors. A substantial upgrade of the system is planned for 2007.

Figure: New LRZ building with compute cube (1), institute building (2), lecture rooms (3)

Close to the LRZ building the Faculties of Mathematics and Informatics of the Technical University Munich are located, and within a distance of a few hundred metres is the Computing Center of the Max Planck Society. This opens up numerous new opportunities for collaborations.

The new address of LRZ is

Leibniz Computing Center
Boltzmannstraße 1
85748 Garching/München
Germany
Phone: +49-89-35831-8784

The 4th HLRS-NEC Teraflop Workshop

The fourth Teraflop Workshop continues the series of workshops which are held alternating in Japan and at HLRS twice a year within the Teraflop Workbench co-operation between NEC and HLRS. The workshop provides a meeting platform for scientists, application developers, hardware designers and international HPC experts to discuss recent developments and future directions of supercomputing. This years spring event held on March 30 and 31 focussed on the applications perspective of high-end computer systems. Speakers like Hiroaki Kobayashi (Tohoku University) and Leonid Oliker (LBNL) discussed leading computational methods and performance issues on leading vector and scalar platforms. At the second day, experts like Thomas Schulthess (ORNL) and Alice Koniges (LLNL) talked about terascale and and the direction to petascale simulations in material science and other fields.

In additional sessions, computational scientists from different fields presented their current work. Many of them were already able to provide results of their applications running with teraflop performance on the NEC SX-8 system. Some of the scientists have already shown new solutions which led to new insights and developments which have not been possible before.

The next workshop will be held at Tohoku University, Japan,
November 20-22, 2006

The 5th HLRS/hww Workshop on Scalable Global Parallel File Systems, HNF Europe Spring Meeting 2006 and OpenIB/RDMA Track

April 3-5, 2006,
HLRS Stuttgart, Germany

Representatives from science and industry interested in high performance storage solutions met at HLRS during April 3-5, 2006, for the fifth annual Workshop on Scalable Global Parallel File Systems. Under the motto of "Getting on Top of the Storage CUBE", this year's three-day event tried to leverage contributions coming from three major research and development areas: Object storage devices and associated file systems, unified 40-Gbit/s communication systems as well as standardization efforts performed by the OpenFabrics Alliance on the InfiniBand architecture. More than 150 participants followed a total of 50 presentations that have been on the workshop agenda.

The opening address on Monday morning was given by Prof. Dr. Michael Resch, the HLRS director. Kent Winchell, deep computing architect at the IBM Systems & Technology Group, delivered the keynote speech on the challenge of exponentially increasing storage in a "Data Intense World". Most of the pioneers of the OSD standard movement were on Monday afternoon's speakers' list, as Harriet Coverston, distinguished engineer responsible for storage systems at Sun Microsystems remarked correctly. This was true for Dr. Peter Braam (CFS Lustre) represented by Jeffrey Denworth, Dr. Garth Gibson (Panasas), Gautham Sastri, Terrascale, and Julian Satran and Dr. Dalit Naor, IBM Haifa.

Tuesday morning saw a series of simultaneous presentations demonstrating the maturity, performance and user experience of the major parallel file systems in the market, IBM GPFS, Panasas Active Scale at RRZK Cologne, HP Scalable File Share at SSCK Karlsruhe, CFS Lustre at CEA Paris, as well as ADIC's StorNext. Engenio and FalconStor talked about their storage solutions incorporating InfiniBand host interfaces which may form the beginning of an all open fabrics cluster design space.

By honoring HNF Europe's 10th year anniversary, Tuesday afternoon provided the venue for unified 40 Gbit/s communication architectures, as presented by Thomas Brenner, Alcatel SEL. His presentation was based on research jointly carried out by Alcatel, SSCK and HLRS with financial support from the state of Baden-Württemberg. The aim of this project has been to better understand the compensation methods for various limiting effects in very high speed optical paths. Since 12 months, SSCK and HLRS are connected by a 40 Gbit/s production link, and part of this work may be applied to the DEISA project in order to enable distributed supercomputing applications. Cisco Systems leveraged the HLRS workshop as a forum to announce their Isola family of Fibre Channel switches, featuring very high port capacities for 4 Gbit/s FC ports as well as 10 Gbit/s Inter Switch Links. In fact, the Isola MDS 9513 beta test was carried out at HLRS end of last year, and on February 23rd the first 10 Gbit/s long-distance FC Link was demonstrated between HLRS and the University of Hohenheim.

On Tuesday night, the Stuttgart Museum of Fine Arts hosted the workshop attendees for a special tour of the contemporary Fritz Winter exhibition and a very memorable dinner at the exclusive window front of its Cube Restaurant on fourth floor: The top of the Storage CUBE!

Wednesday 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. There are plans to form a European Chapter of the OpenFabrics Alliance which is underscored by a substantial presence of OpenFabrics at various conferences in Europe including ISC 2006, Dresden. Finally, HLRS provided insight into its center on many levels: parallel file systems (NEC GFS, IBM GPFS, Terrascale TerraGrid), parallel applications within the Teraflop Workbench, simulation steering and virtual reality over 40 Gbit/s networks, as well as personal tours of the center.

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

The 3rd Russian-German School on Parallel Programming

The 3rd Russian-German School on Parallel Programming using High Performance Computation Systems will be held on August 28-September 8, 2006, in the Institute of Computational Technologies of the Siberian Branch of the Russian Academy of Sciences (ICT SB RAS) in Novosibirsk, Russia.

The aim of the School is to invite young researchers to discuss modern parallel computing techniques, gain experience in using up-to-date hardware and exchange experience in solving practical problems on high performance computation systems of different architectures.

The program of the school consists of two courses. The first course is an introductory course into the technologies of parallel computing and their using in numerical modelling. The advanced course is designed for a detailed study of the methods and techniques of high performance computing covering advanced topics in parallel programming.

The courses will be run by the leading scientists from the High Performance Computing Center Stuttgart. The leading scientists from the institutes of the Siberian Branch of the Russian Academy of Sciences will give the invited lectures.

The Student Session is planned, where participants will have a possibility to give presentations on their works in the field of parallel programming. Abstracts of the accepted presentations will be published in the Student Session Proceedings.

New Books

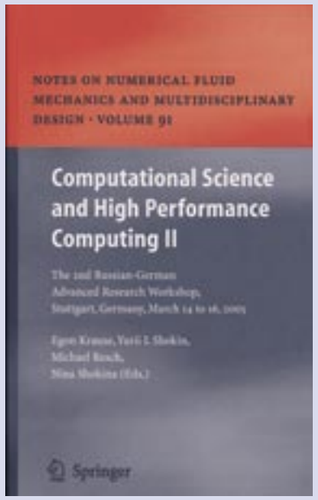

Computational Science and High Performance Computing II

Krause, Shokin, Resch, Shokina (Eds.), Springer, 2006, ISBN 3-540-31767-8

This volume contains 27 contributions to the Second Russian-German Advanced Research Workshop on Computational Science and High Performance Computing presented in March 2005 at Stuttgart, Germany. The workshop was organized by the High Performance Computing Center Stuttgart (HLRS) and the Russian Institute for Computational Technologies (ICT SB RAS).

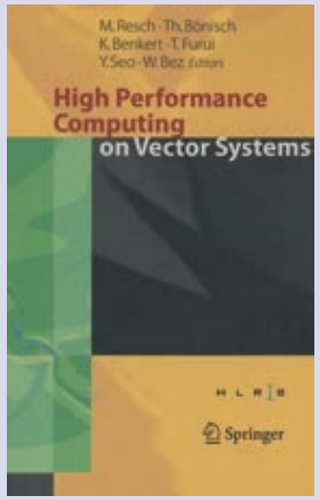

High Performance Computing on Vector Systems

Resch, Bönisch, Benkert, Furui, Seo, Bez (Eds.), Springer, 2006, ISBN 3-540-29124-5

The book presents the state of the art in HPC and simulation on modern supercomputers covering architectural trends in general and specifically the future of vector-based systems and heterogeneous architectures. The contributions cover among others computational fluid dynamics, physics, chemistry, astrophysics, and biology. Innovative fields like multiphysics simulations and material science are presented. Furthermore, the use of supercomputers in the growing field of medical simulations is shown. All papers were presented at the second Teraflop Workbench Project Workshop held at the Höchstleistungsrechenzentrum Stuttgart (HLRS) in March 2005.

inSiDE

inSiDE is published two times a year by
The German National Supercomputing
Centers HLRS, LRZ, and NIC

Publishers

Prof. Dr. Heinz-Gerd Hegering, LRZ
Prof. Dr. Dr. Thomas Lippert, NIC
Prof. Dr. Michael M. Resch, HLRS

Editor

F. Rainer Klank, HLRS
klank@hlrs.de

Design

Julia Schlatterer
schlatterer@hlrs.de

Authors

Andreas Dolfen
a.dolfen@fz-juelich.de
Prof. Dr. Rainer Friedrich
r.friedrich@lrz.tu-muenchen.de
Dr. Erik Koch
e.koch@fz-juelich.de
Uwe Küster
kuester@hlrs.de
Dr. Peter Lammers
p.lammers@hlrs.de
Dr. Inga Mahle
inga.mahle@gmx.de
Prof. Herbert Over
herbert.over@phys.chemie.uni-giessen.de
Dr. Eva Pavarini
e.pavarini@fz-juelich.de
Dr. Rossitza Pentcheva
pentcheva@lrz.uni-muenchen.de
Dr. Sabine Roller
roller@hlrs.de
Dr. Ari P. Seitsonen
ari.p.seitsonen@iki.fi
Dr. Jörn Sesterhenn
joern.sesterhenn@alumni.ethz.ch
Stefan Wesner
wesner@hlrs.de

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