inSiDE

inSiDE • Vol. 1 No. 2 • Autumn 2003

Innovatives Supercomputing in Deutschland

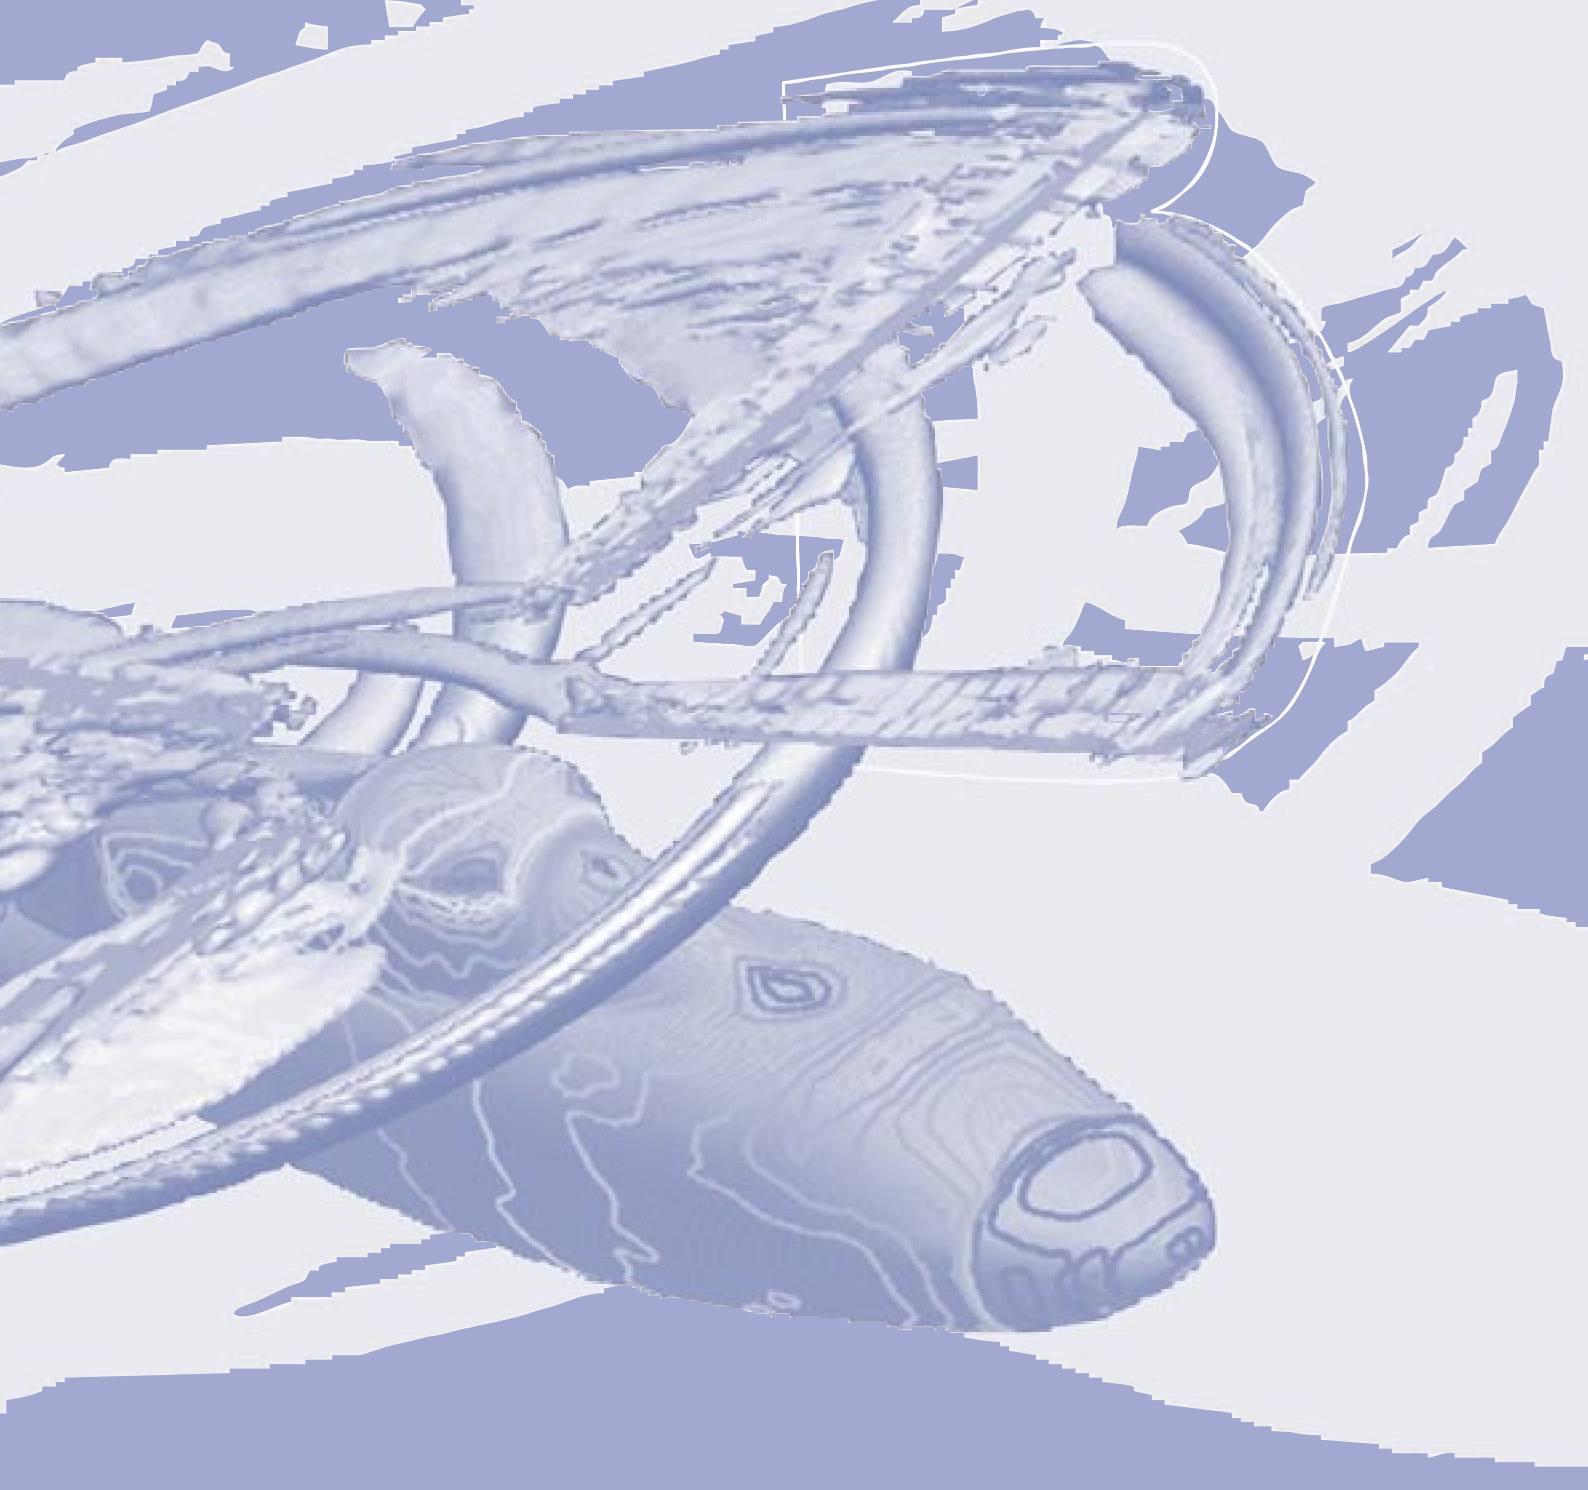

Editorial

With this second issue of inSiDE the German Federal Supercomputing Centers in München (LRZ), Jülich (NIC) and Stuttgart (HLRS) are continuing their initiative to disseminate information about new concepts of supercomputing as are embraced by the German supercomputing research community. Biannually inSiDE presents information on supercomputing in Germany to its national and international readers. Following the enthusiastic and positive response to the first issue, both in Germany and abroad, this second issue is being published at the Supercomputing Conference (SC'03) in Phoenix/AZ, USA.
This second issue starts with application reports which show that high performance computing goes far beyond the activities of the three national centers. Gerrit Schierholz from DESY Zeuthen reports on Nucleon Structure from Lattice QCD. Oliver Wenisch, Joachim Noll and J. Leo van Hemmen from the Department of Physics of the Technische Universität München report on the Numerical Simulation of Neural Network Dynamics in the Primary Visual Cortex. Robert O. Jones and Jaakko Akola from the Forschungszentrum Jülich report on prediction of molecular structures and reaction paths. The HLRS Golden Spike 2003 winners conclude the application section. This award is given each year to the three most outstanding contributions of young researchers using HLRS supercomputers.
The architecture section introduces the new IBM system at Jülich which will be fully installed by the end of 2003 and will then be the fastest available system in Germany. First user experience is given and indicates the potential of the system.
The section on centers opens with an overview of the supercomputing centers in Germany by Friedel Hossfeld. Each center has its own history, its own strengths (and weaknesses) and its own profile. Hossfeld's article is mainly addressed to non-German readers but also the German reader may find new and interesting material. Finally you will again find information about the available supercomputer systems at the three centers and about workshops offered by the centers. All workshops are open to researchers from Germany. Other interested people should contact the responsible center for terms of participation.
One of the key points of discussion in the supercomputing community in Germany during the recent months was a German GRID initiative (called d-grid) that would continue and dramatically extend the successful work of UNICORE. In May a steering committee was established to prepare an initiative that aims at going beyond traditional GRID approaches and embraces the idea of e-science in a distributed cyber-infrastructure. Because of the difficult budgetary situation the project start has been delayed such that it is still too early to report. However, we expect to see some progress in this field in the next year which will then also be reflected in the next issue of inSiDE. This next issue will also cover information about the new computer systems at Jülich and Stuttgart.

Prof. Dr. H.-G. Hegering (LRZ)
Dr. R. Esser (NIC)
Prof. Dr. M. M. Resch (HLRS)

Contents

Numerical Simulation of Neural Network Dynamics in the Primary Visual Cortex

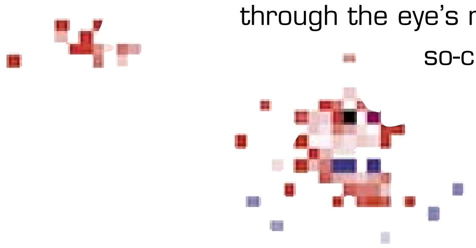

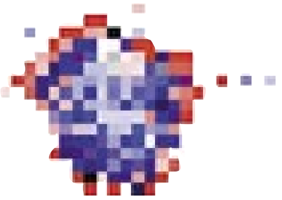

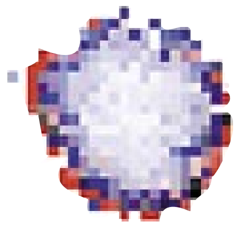

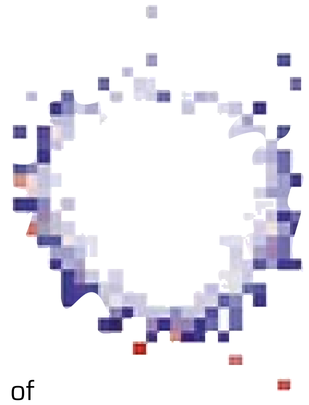

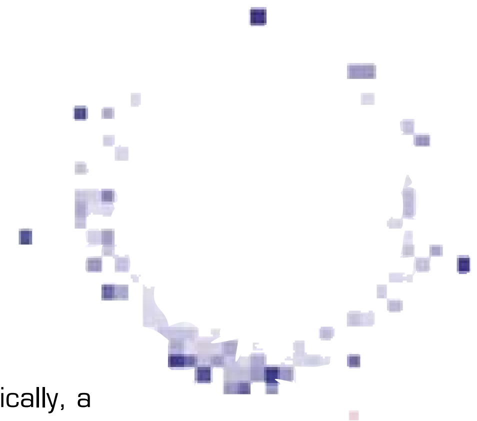

Processing of visual input received through the eye's retina is done in the so-called visual cortex. On the genetic substrate of an individual's neuronal development an activity-driven learning process connects neurons in this part of the brain with each other by means of synapses during early childhood and, in so doing, forms a powerful neuronal network. While the time scale of neuron activation is on the order of a few milliseconds, the synaptic learning process extends over several days. As a consequence, simultaneous numerical simulation of the activity dynamics in a network of 10^5 or more neurons and the much slower learning process is hardly possible even on today's supercomputers.

By using special approximations such as the Spike-response model, we were nevertheless able to efficiently simulate ongoing activity in a large artificial network for primary visual cortex and study various learning paradigms with high temporal resolution. According to the underlying model, the state of a neuron is completely determined by the value of its membrane potential. This quantity varies with the strength of the inputs a neuron receives from other neurons. When the input of a neuron exceeds a certain threshold, it fires a "spike", i.e., a characteristic voltage pulse is induced in its membrane potential. This signal is transmitted to surrounding neurons. At the connecting synapse, the neuron's output either excites or inhibits activity in its target depending on the neuron being excitatory or inhibitory.

For our purpose a model network of the visual cortex has been implemented as a two-layer approximation (cf. zoom inlet of anatomic figure below) in a parallel simulation setup on the Hitachi SR8000 supercomputer of the Leibniz-Rechenzentrum to study neuronal activity dynamics and learning processes with millisecond time resolution.

Typically, a total of about $2 \cdot 10^4$ neurons - connected by approximately $5 \cdot 10^6$ synapses - were simulated using periodic boundary conditions. A short sequence of the activity development is depicted in the text background as a series of snapshots of an inflating activity bubble at 10 millisecond time intervals (excitatory and inhibitory activity are represented in red/blue, respectively). It can be seen that the spreading of activity over the cortex is not an instantaneous process. This is due to the fact that spike signals have to travel along thin cell elongations, cables or axons, over considerable distances to their postsynaptic targets. We have varied the simulated propagation delays to see how they influence the simulation results and found that changing their values has a significant impact on the network dynamics in that it can completely change the corresponding activity scenario.

This is particularly important for the activity-driven learning process in the prenatal phase of cortex development, when structured input from the retina is absent or only very weak. In this case only spontaneous activity due to cortex-intrinsic processes is present, which is nevertheless able to drive the learning process that "primes" the synaptic coupling strengths between the neurons.

Structures in the coupling pattern between neurons represent the physiological basis for the neuronal circuits that allow the brain to filter important visual information such as object movement. Abilities of this kind are taken for granted in every-day life but quite often are not yet established at the time of birth. A popular view is that activity-dependent processes of synaptic self-organization in an early phase of development pave the way for later refinement by learning with structured input after birth.

With the ongoing increase of computing power of high-performance supercomputers it will soon be possible to provide realistic simulations of the neural network contained in 1 cubic millimetre of cortex volume. Though still far from network size and complexity exhibited by, e.g., the human brain, we trust larger system size and greater realism will provide both a true challenge and a most welcome perspective.

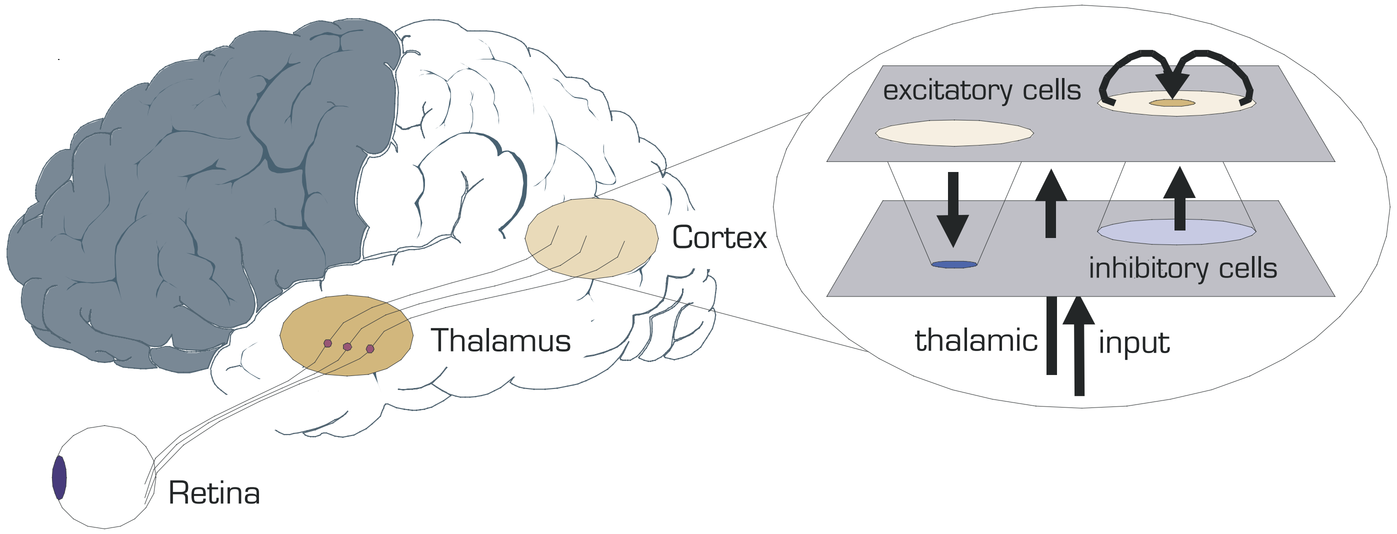

•Oliver Wenisch
•Joachim Noll
•J. Leo van Hemmen

Theoretische Biophysik
Physik Department
Technische Universität München

Nucleon Structure from Lattice QCD

Quantum chromodynamics (QCD) is the fundamental theory of the strong interactions. It binds quarks and gluons, the fundamental building blocks of matter, to nucleons and mesons, and these to nuclei. The forces are so strong that quarks do not appear as free particles in nature. This phenomenon is called quark confinement.

Figure 1 shows a sketch of the proton. The proton can be visualized by three valence quarks (balls) bound together by the exchange of gluons (springs). The gluons strongly interact with each other, which gives rise to the confining force. In addition, gluons may transform into short-lived quark-antiquark pairs, the so-called sea quarks.

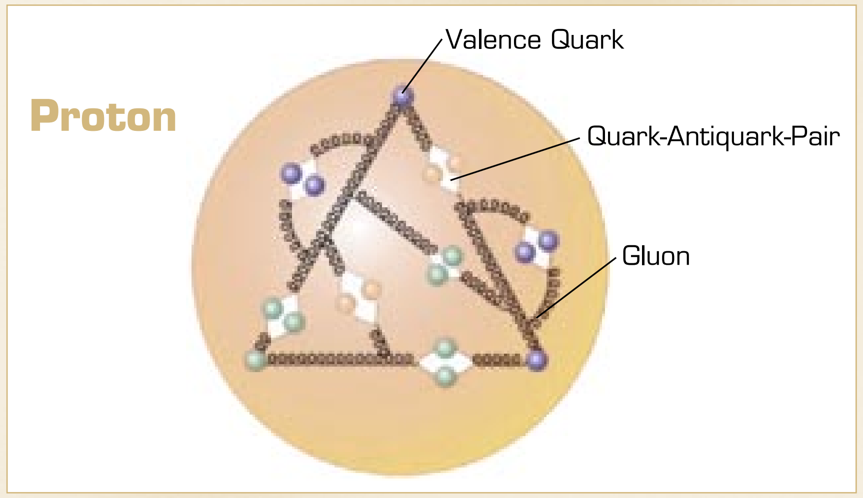

Figure 1: Sketch of a proton

The equations of this theory are so complicated that they cannot be solved by traditional techniques. The modern way to tackle them is through the use of powerful computers. For the computer to solve the theory one approximates space and time by a finite box divided into a discrete set of points, the lattice. This reduces the problem to a finite system of coupled equations which can be solved by standard Monte-Carlo-techniques. Later on one may remove this approximation by letting the lattice spacing go to zero and the box size to infinity.

Probably the biggest challenge in theoretical particle physics is to understand the internal structure of hadrons in terms of quarks and gluons and, in particular, how quarks and gluons provide the mass (binding) and spin of the nucleon. Continuing advances in computing power have now brought us to the point where ab initio calculations of nucleon structure are becoming possible.

A fast moving nucleon of energy E can be viewed as an ensemble of quasi-free quarks (and gluons) contracted to a disc (due to Lorentz contraction), as shown in Figure 2. The quarks are located at a distance r from the center and contribute a fraction xE to the total energy (mass) of the nucleon each.

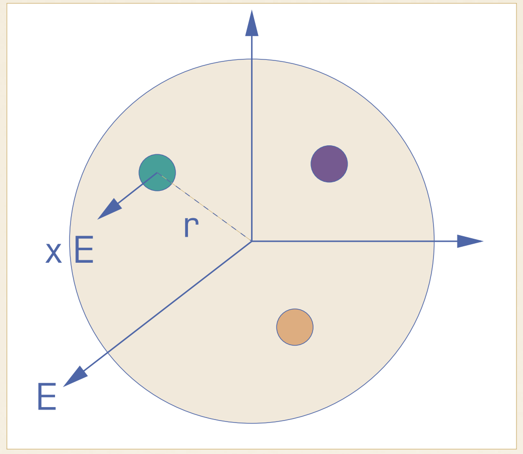

Figure 2: Schematic view of a fast moving nucleon

One of the principal questions is how the energy (mass) of the nucleon is distributed among its constituents, as well as over space. In Figure 3 we show the probability distribution of finding an up quark inside the proton with fractional energy x at a distance r from the center. The plot covers the region $0 \le x \le 1$ and $-1 \le r \le +1$ Fermi (10^{-15}m). We observe that the energetic quarks are limited to a rather narrow region of $r \le 0.3$ Fermi, while the soft quarks spread out to distances of $r = 1$ Fermi This is to say that most of the energy of the nucleon is concentrated in a core of a fraction of a Fermi.

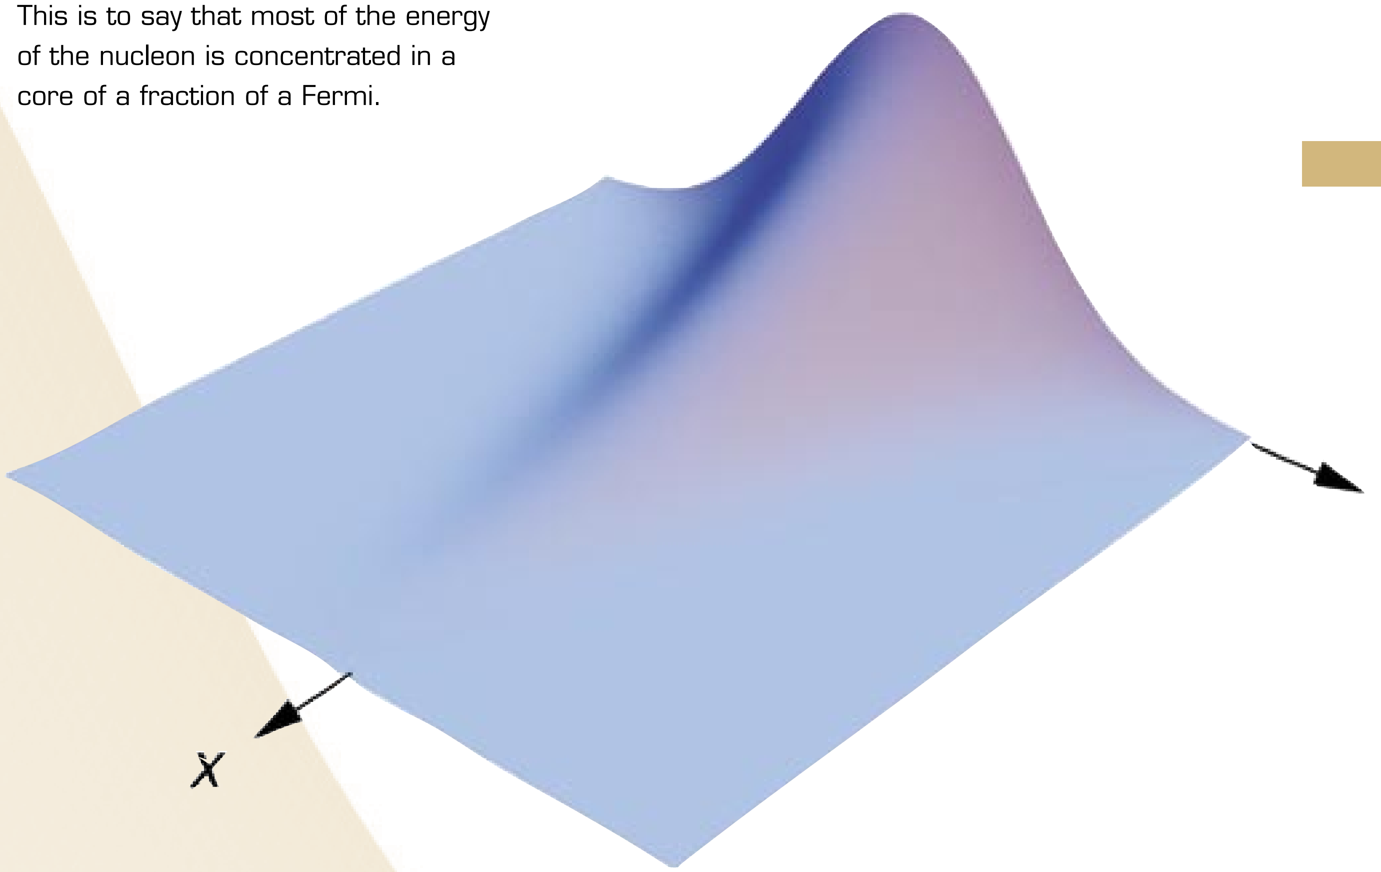

Figure 3: Probability distribution of the up quark inside the proton

Similar pictures can be obtained for the distribution of spin, which will complete our view of the structure of the nucleon.

These calculations, revealing the innermost structure of hadrons, would not have been possible without the use of the world fastest supercomputers, such as the Hitachi SR8000-F1 at LRZ. With the next generation of machines it will, for the first time, be possible to do simulations at realistically small quark masses, which spares us uncontrolled extrapolations to the physical quark mass.

• Gerrit Schierholz

DESY Zeuthen
QCDSF Collaboration

Calculation and Prediction of Molecular Structures and Reaction Paths

ATP molecule

The atomic arrangement of molecules and the changes that take place when molecules react is a central problem in chemistry, molecular and condensed matter physics, and molecular biology. It is not surprising that enormous effort has been expended to calculate these quantities, and supercomputers have played a prominent role.

Geometrical arrangements and reaction paths could be predicted if one could calculate the total energy of a system for all atomic positions. Stable structures correspond to local minima in the energy "surface", and the energy required to perform the transition from one structure to another (the energy "barrier") determines the rate at which the reaction can proceed. Accurate solutions of the Schrödinger equation for the many-electron wave function would lead, in principle, to this information and much more, but this approach is restricted to systems containing few atoms. The description of the forces in the system by adjusting the parameters of a potential function (a "force field") provides an alternative, but the results are sensitive to the way the parameters are fixed.

Experience over the past 20 years or so has shown that the density functional formalism can provide energy differences that have predictive value. The range of applications has been overwhelming: reactions at surfaces, the structure and magnetic properties of metallic multilayers, catalytic reactions in bulk systems, etc.. The formalism states that the energy (and many other properties) can be determined from a knowledge of the electronic density. It results in equations that have a form that is similar to that of the Schrödinger equation, but much simpler to solve numerically. Approximations are still unavoidable for one contribution to the energy (the "exchange-correlation energy"), but there has been significant progress on this front as well.

A recent application of density functional methods is provided by the reaction of water with adenosine 5'-triphosphate (ATP). ATP is the most important energy carrier in cellular metabolism, and each human being produces its own weight in ATP every day. The ATP molecule is shown in Figure 1, where carbon atoms are grey, hydrogen white, oxygen red, nitrogen blue, and phosphorus orange. The situation in biological systems is complicated by the presence of water and many other molecules and ions, and we show in Figure 2 the model system for which we have studied two reactions in an aqueous environment: an associative reaction involving the attack of one water molecule, and a dissociative reaction involving the scission of the terminal bridging P-O bond. Magnesium ions (green) are catalysts in these processes.

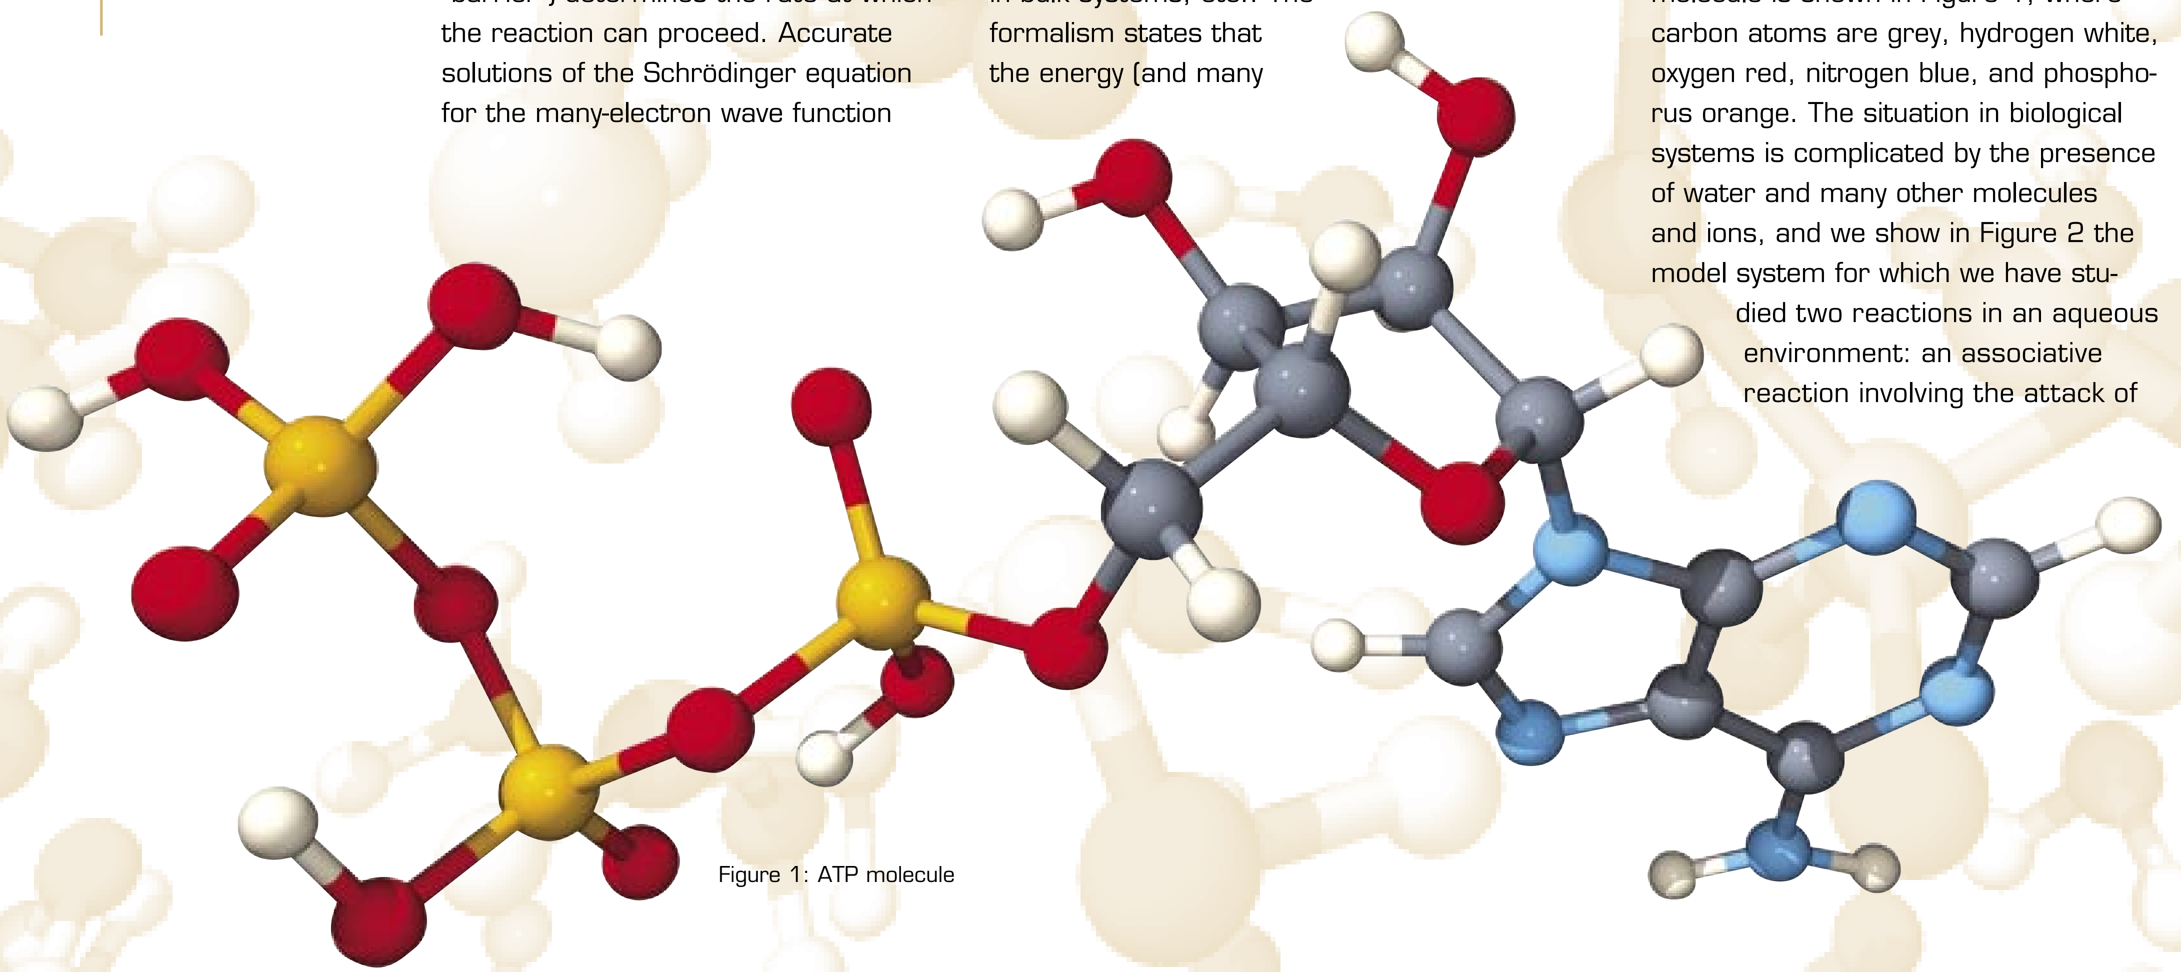

Figure 1: ATP molecule

The calculations were performed with the CPMD program (J. Hutter et al., IBM Zürich and MPI für Festkörperforschung, Stuttgart), which uses a plane wave basis and makes extensive use of the fast Fourier transform (FFT). The program runs on many platforms, and it is used by research groups in many parts of the world. There are, of course, implementations of the density functional approach in many other program packages.

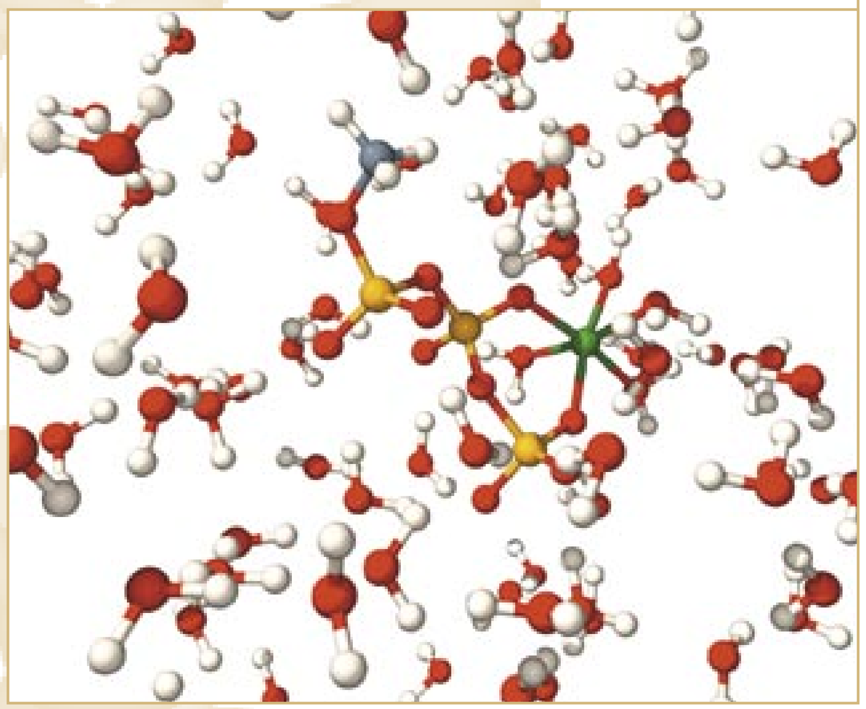

Figure 2:
Chemical reactions of ATP in an aqueous environment

•Robert O. Jones
•Jaakko Akola

Institut für Festkörperforschung (IFF), Forschungszentrum Jülich GmbH

HLRS Golden Spike Award 2003

The sixth Results and Review Workshop of the High Performance Computing Center Stuttgart (HLRS) took place at the HLRS on October 6–7, 2003. Outstanding results were presented by 33 selected projects processed on the supercomputing platforms at HLRS and its partner organisation SSC Karlsruhe. Additional 6 contributions came from the Leibniz Computing Center (LRZ) as part of a cooperation among the German Federal HPC Centers.
Research areas included computational fluid dynamics, reacting flows, physics, astrophysics, solid state physics, geophysics, earth sciences, chemistry, bioinformatics, nanotechnology and computer science. Reflecting the close cooperation of the HLRS with industry, special emphasis was put on the industrial relevance of the presented results and methods.
The workshop was completed by a tutorial on parallel programming models on hybrid systems in which the strengths and weaknesses of several parallel programming models on clusters of shared memory processor nodes were shown.

Golden Spike Award

The steering committee of the HLRS, a panel of twelve leading German scientists and responsible for project proposal reviews, appreciated the high quality of the work carried out in Stuttgart as well as the spectacular scientific results and the efficient usage of supercomputer resources by awarding three selected projects of Golden Spikes of HLRS. The winnig projects are presented on the following pages.

Collisional Dynamics of Black Holes, Star Clusters, and Galactic Nuclei

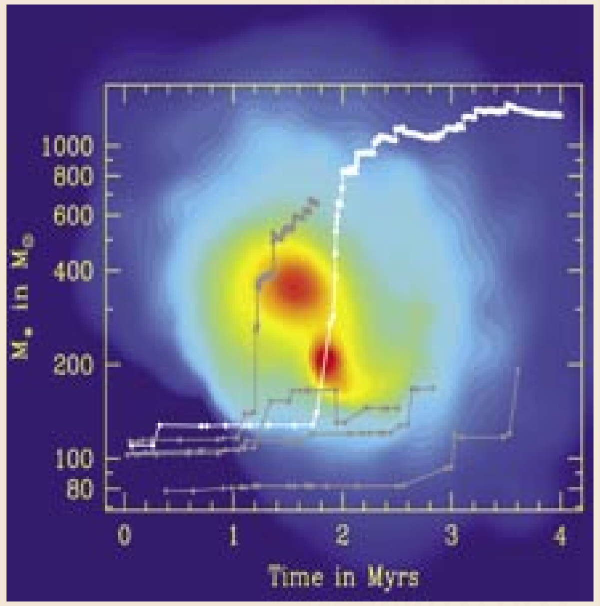

Figure: RunAwayGrowth

Figure 2: CoreCollapse

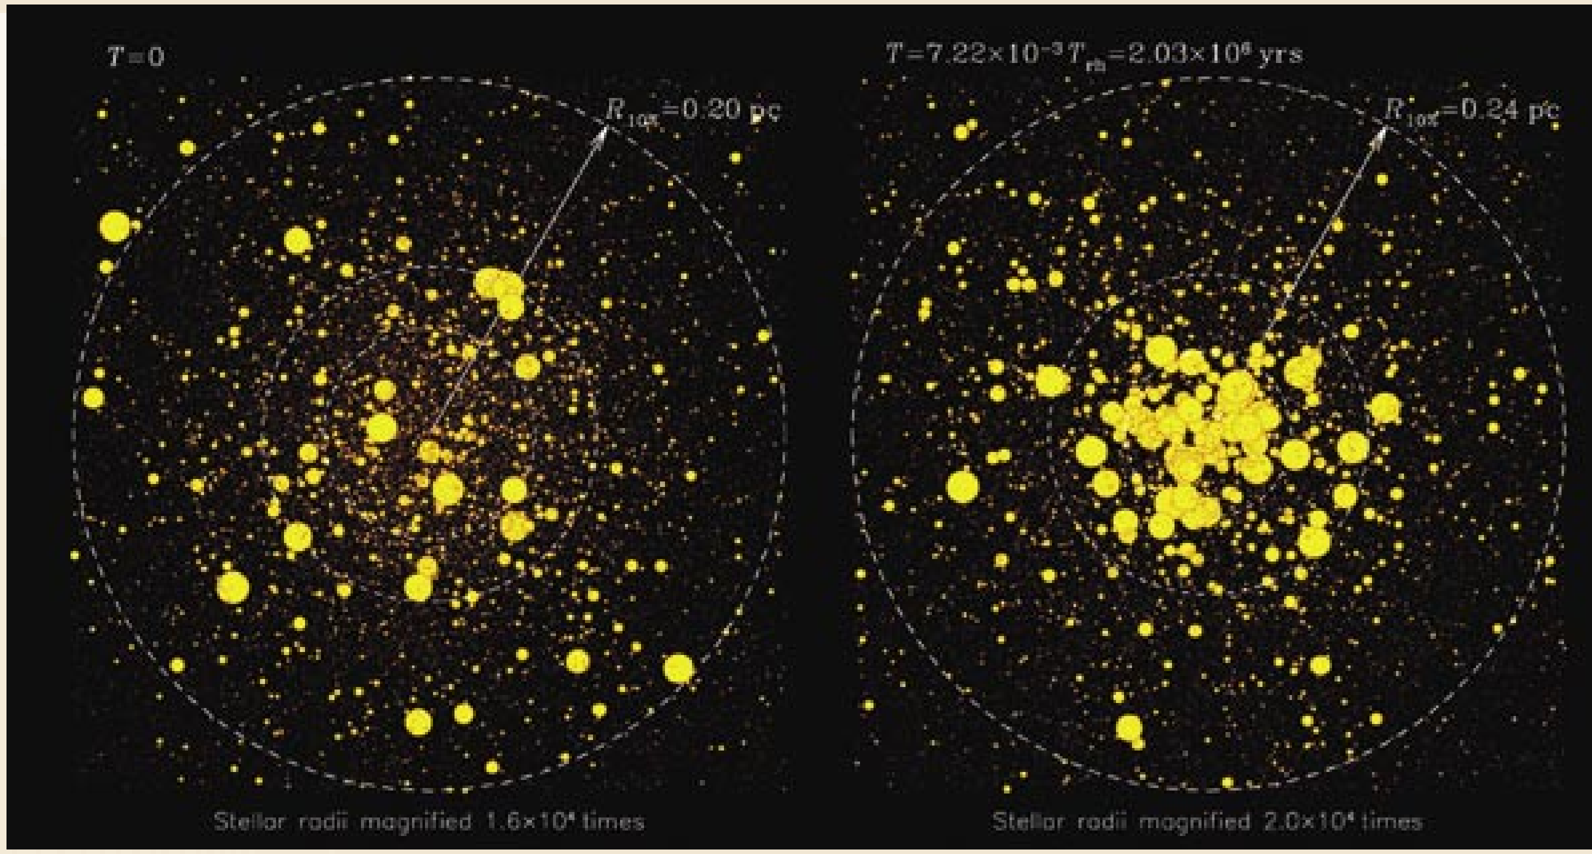

Evolution of the central region of a spherical cluster of stars, as computed with the Monte Carlo algorithm. Left: initial configuration. Right: during core collapse, at the moment massive stars start colliding with each other at the centre. All the stars within a slice containing the centre are depicted. For clarity, their radii are highly magnified. The white circles represent spheres containing 1, 3 and 10% of the total cluster mass (from the centre). One parsec (pc) is 3e16 meters. Note how the massive, large stars concentrate to the centre, a process called “mass segregation”. This is caused by relaxation, i.e. energy exchange between stars due to close gravitational encounters. T_rh is the time after which the energy of an average star orbiting at a distance from the centre enclosing one half of the cluster's mass has changed by of order 100% through relaxation.
In a dense stellar cluster, Massive stars that have sunk to the centre during core collapse collide with each other. Some of the collisions result in mergers which allow for the runaway growth a very massive star. The curves show the mass evolution for a few stars that have undergone a large number of collisions (dots) in a Monte Carlo cluster simulation. The masses are in units of the solar mass (M_sun).

The white line corresponds to a star growing from 120 M_sun to more than 1000 M_sun. Such a star may turn into a massive black hole at the end of its life. The background image is a snapshot from one of the 12000 “Smoothed Particle Hydrodynamics” simulations realised to construct the grid of collision outcomes used in the Monte Carlo algorithm.

•Dr. Marc Freitag

Astronomisches
Rechen-Institut
Universität Heidelberg

Interactive CFD Simulations by Coupling Supercomputers with Virtual Reality

In this joint project of the Chair for Bauinformatik of the Technical University Munich and the Leibniz Computing Center (LRZ) we are implementing an interactively steerable computational fluid dynamics (CFD) simulation. CFD simulations in a Virtual Reality (VR) environment allow a very flexible analysis of complex flow phenomena, supporting the planning process of a building with respect to fluid mechanical aspects. The long-term goal of the project is an online-CFD simulation of CAD-generated virtual rooms. At the present stage we are studying the air flow within a cuboid containing fluid obstacles like cubes, spheres, and cones which the user can freely place within a virtual-reality interface.

Only dynamic online visualization allows the intuitive investigation of the transient nature of a flow field. Hence visualization and steering is performed concurrently on a separate virtual-reality graphics workstation, which itself is connected via PACX-MPI to the simulation kernel running on the supercomputer. With the interactive 3D user interface it is possible to add or remove fluid obstacles into or from the fluid volume at any time during the running simulation. The user can change the size or location of the objects and watch the adapting fluid flow instantaneously. The change in geometry or parameters is sent to the simulation kernel which computes the corresponding changes in fluid data. The underlying CFD simulation is performed by a Lattice-Boltzman kernel.

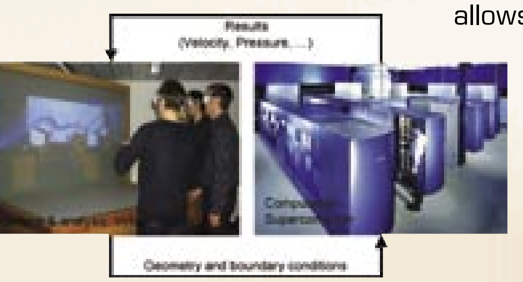

Figure 1: Data exchange between VR interface and CFD simulation on supercomputer

• Petra Hardt[1]
• Siegfried Kühner[1]
• Ernst Rank[1]
• Oliver Wenisch[2]

[1] Lehrstuhl für Bauinformatik TU München
[2] LRZ

This approach to solving the Navier-Stokes equations is particularly well suited for taking advantage of the parallelization capabilities of the LRZ's 2 TFlop high-performance supercomputer Hitachi SR8000. Excellent parallel efficiency is achieved by partitioning the simulation domain along the x-axis into several sub-domains assigned to a corresponding number of nodes of the SR8000 communicating via the MPI message passing library. Within each sub-domain Hitachi's SMP-parallelization COMPAS is applied to split the y-dimension work for 8 CPUs of each computation node. Finally, pseudo-vectorization is used to concurrently process voxels along the z-direction.

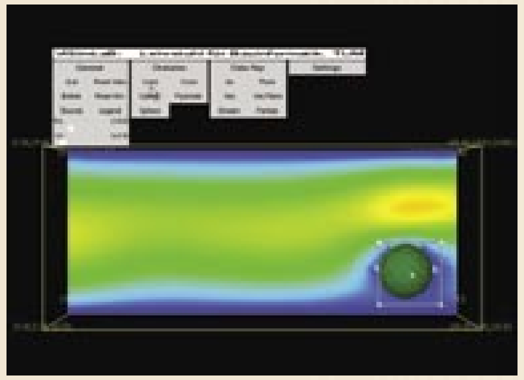

Figure 2: Steady state fluid flow configuration with one obstacle when air is streaming from left to right. Fluid velocity increases from blue (resting air) to red

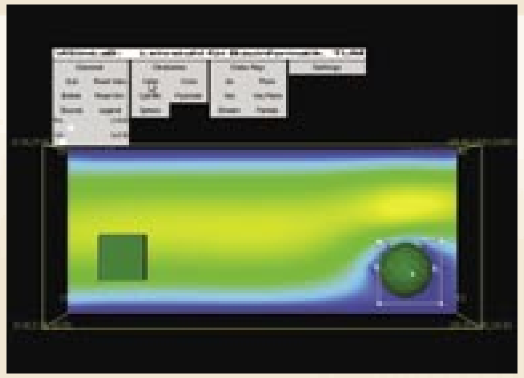

Figure 3: A cube is added to the fluid volume ...

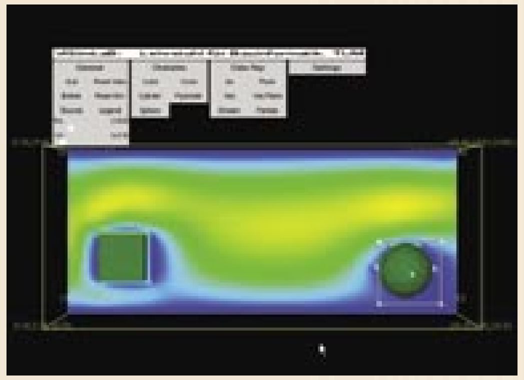

Figure 4: The new steady state is reached within a few seconds

An important feature of our computa tional steering approach is that the simulation kernel need not be interrupted during its ongoing computation for the communication required for user interaction. This offers the possibility of smooth scene updates within the postprocessing on the visualization client and fast computation of the adapting fluid which reaches its new steady state within a few seconds.

Future improvements of our application will cover support for more complex CAD-generated geometries and the incorporation of turbulence effects.

Acknowledgements: We would like to thank Prof. Manfred Krafczyk (TU Braunschweig) and KONWIHR for their support.

Turbulent Convection at Large Aspect Ratios

Convection is a ubiquitous phenomenon. It occurs in many engineering applications, in the atmosphere, the oceans, and on a long time scale even in the mantle of the earth. Most applications deal with turbulent flows. Research on turbulent flows focuses on at least two fairly separated sets of questions, the study of small scales of the flow and its coherent structures.

The interest in the small scales stems from the hope that these are in some sense universal, i.e. identical from one flow to the other. If the behavior at small scales is indeed universal, it should be possible to parameterize its effect on the large scales so that it would not be necessary to spatially resolve the smallest scales of the flow in numerical simulations. So called Large Eddy Simulations (LES) are one class of methods exploiting this idea.

Coherent structures on the other hand are organized features of the flow which are distinguished from the turbulent background. These structures can be of either large or small scales. Prominent examples include stream-wise vortices and thermal plumes.

This project makes a contribution to the study of small scales in turbulent convection by extracting statistical quantities useful for the design and validation of LES methods, but the main focus of this research is on the large scale patterns which persist in the flow in spite of turbulent fluctuations.

We study turbulent convection by direct numerical simulations (DNS) in the simplest possible geometry: A plane layer of fluid heated from below and cooled from above. Such flows contain a wide spectrum of length and time scales. Performing DNS is therefore computationally demanding.

• Thomas Hartlep

Institut für Geophysik
Universität Göttingen

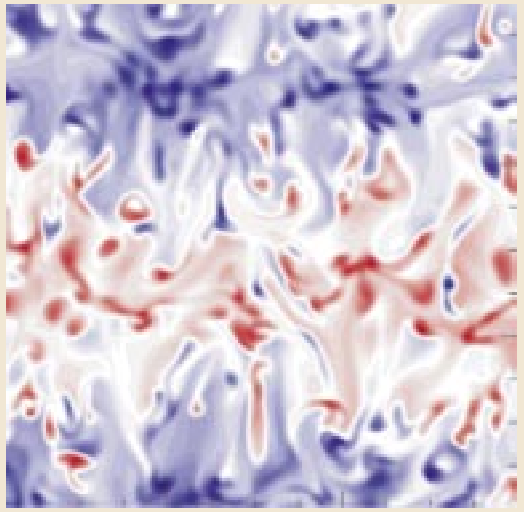

Figure 5: An example of a temperature distribution in the midplane of the layer

The New IBM Supercomputer in Jülich – Configuration and First User Experiences

IBM

Figure 1: IBM p690 node

Introduction

In August 2002 Forschungszentrum Jülich and IBM signed a contract about the delivery, installation, and maintenance of a new high-performance supercomputer. This system which uses IBM POWER-4 microprocessor technology will be installed in a new machine room at Forschungszentrum Jülich's Central Institute for Applied Mathematics (ZAM). It will replace the two "old" 512 processor Cray T3Es installed in 1996 and 1999 respectively, and will reach a peak performance of 8.1 TFLOPS – a factor of 13 more than ZAM's current Cray T3E-1200 capacity.

Configuration

The installation is being performed in three steps: In October 2002 two IBM e-servers p690 with 32 POWER-4 processors (1.3 GHz) each were delivered for tests of the operating system and early application porting. In July 2003 a larger subsystem, consisting of six IBM e-servers p690 with 32 new POWER-4+ processors (1.7 GHz) each was delivered, installed, and offered to scientific computing projects, which have been successfully peer-reviewed by the Resource Allocation Committee of the John von Neumann Institute for Computing (NIC) or the corresponding commission of the Forschungszentrum Jülich. The nodes of the subsystem are connected by Gigabit-Ethernet. As this may produce communication bottlenecks when parallel programs are executed on more than one node, currently only programs requesting 32 or less processors are run. Together with the installation of the final configuration in December 2003 a high performance switch will be available which enables a node-spanning execution of parallel programs. With 37 IBM e-servers p690 (POWER-4+ processors) this system will reach the theoretical peak performance of 8.1 TFLOPS.
Five of the six nodes currently available are configured as compute nodes and are used for the execution of parallel programs. The sixth node (login node) consists of a login partition with several processors and a partition for data management tasks (I/O, backup etc.). Batch jobs and interactive parallel programs are both controlled by the IBM proprietary batch system LoadLeveler.
The architecture and software of the system supports a variety of parallel programming paradigms: message passing with MPI (including the functions "one-sided communication" and "parallel I/O" defined in MPI 2), one-sided communication with IBM LAPI or Cray SHMEM, multi-threading with OpenMP or Posix Threads or any combination of these paradigms (hybrid programming).
Besides the necessary compilers and libraries, which support the different programming models, important software tools are also available, like the parallel TotalView debugger, the performance analysis tools Vampirtrace/Vampir for MPI programs and Guide-View for OpenMP programs or several tools from IBM like the callgraph profiler gprof/Xprofiler, the hardware counter profiling tools hpmcount and hpmlib, the MPI profiling tool MP_profiler and the cache simulator Sigma. The offer is complemented by mathematical software, like the Engineering Scientific Subroutine Library (ESSL) or linear algebra libraries (LAPACK etc.) and application software like CPMD (Car-Parrinello molecular dynamics), LS-DYNA or ANSYS (finite elements; restricted use) or Gaussian03 (ab initio chemistry).

User Experiences

First user experiences show that porting applications from Cray T3E to the new IBM system seems to be no problem at all. Even the data transfer between these two supercomputers including the data conversion was in general straightforward. In most cases the compiler options hot and arch=pwr4 together with the optimization level O3 were sufficient to create well-performing codes. Detailed studies of codes with respect to memory access and the attempt to guide the access to the three cache layers by special environment variables and compiler options often did not improve the performance of the codes much more than the three options given above.
Users report that the performance gain between Cray T3E and IBM p690 is between a factor of 4 and a factor of 10. Measurements of applications from high energy physics (HEP), molecular dynamics (MD), and environmental research (ENV) are compiled in a Kiviat diagram (Figure 3). In this plot we compare runs on a Cray T3E-1200 using 32 processors with corresponding runs on a p690 node (32 processors, 1.7 GHz). On this node large page support was not enabled. Memory affinity, however, was switched on. Of course, up to now we are not able to present any results for codes using more than one node.
It should be mentioned that our experiences very well match with benchmark results of our partners at DESY Zeuthen. They have just recently published a summary of benchmark runs on two different special purpose computers, on a PC cluster, and on a Cray T3E-900, on a Hitachi, and on an IBM p690 system (hep-lat/0309149).

Outlook

We are now looking forward to the upcoming acceptance test of the complete system in November, the installation during the turn of the year in the new machine room and the start of full production at the beginning of February 2004.
For some period, this system will be the largest, made available to the scientific community by the German high-performance computing centres. According to the recommendations of the German Scientific Council, the other two National Supercomputer Centres in Stuttgart and Munich are already preparing the next steps to take the lead in the innovation spiral in high-performance computing in 2005 and 2006.

Figure 2: View of the new machine room in Jülich, being built for the IBM supercomputer

IBM

Figure 3: Kiviat diagram for six applications

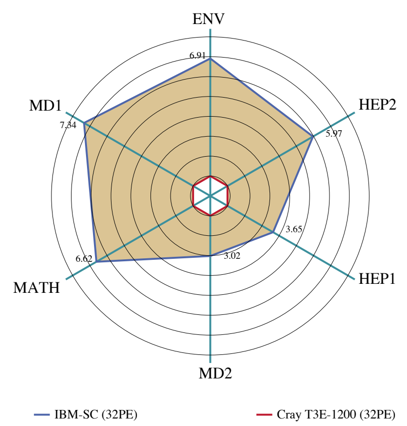

• Norbert Attig

Forschungszentrum Jülich, ZAM

National Supercomputing Centres – Turning the Innovation Spiral

Scientific Supercomputing as a Key Technology

For the first time in computing history, in the 1990s we were able to establish a well balanced pyramid ranging from local-area and wide-area broadband networks, via the diversity of work-station and PC platforms in client-server structures supporting cooperative and even realtime computing, up to the layer of medium-sized innovative com-puter architectures capable to spread computational science and engineering as a key technology over science and industry and also to support methodo-logically and capacity-wise, finally, the apex of the pyramid, the infrastructure of top-level supercomputers with the mission of a nation-wide resource.
Several national initiatives focused much attention and gave terrific technological and scientific impact to a research and development field which developed in parallel with the tremendous increase and ubiquitous distribution of computer capacity over the past five decades: Although born in the 1940s, it has been named Computational Science only in the mid-1980s by the Nobel Price Winner Kenneth Wilson and has been termed in the 1990s Computational Science & Engineering. Computer simulation has grown and established itself as the third category of scientific meth-odology. This ever-innovating discipline fundamentally supplements and com-plements theory and experiment, as the two traditional categories of scien-tific investigation, in a qualitative and quantitative manner while integrating these into the methodological tripod of science and engineering. Being com-parable rather with an experimental discipline, Computational Science and Engineering wastly extends the analytical techniques provided by theory and mathematics; today, in a sense, it is synonymous with inves-tigating complex systems. Its main instrument is the supercomputer; its primary technique is computer simula-tion. Unsolved complex problems in the areas of climate research and weather forecast, chemical reactions and combustion, biochemistry, biology, environment and ecological as well as economic and sociologic systems, or-der-disorder phenomena in condensed-matter physics, astrophysics and cosmology, quantum chromodynamics, and, in particular, aerodynamics have been identified as "Grand Challenges".
The tripod of science and engineering, thus, has proved to provide scientific research and technology with the meth-odological basis and the instrumental laboratory to effectively approach the solutions of complex problems which are critical to the future of science, technology, and society. It will be a crucial factor for the industry in order to meet the requirements of interna-tional economic competition especially in the area of high-tech products. The various strategic position papers and the government technology programs in the USA, in Europe, and in Japan in the 1990s claimed that the timely provision of supercomputers to science and engineering and the ambitious development of innovative supercom-puting hardware and software archi-tectures as well as new algorithms and effective programming tools are an urgent research-strategic response

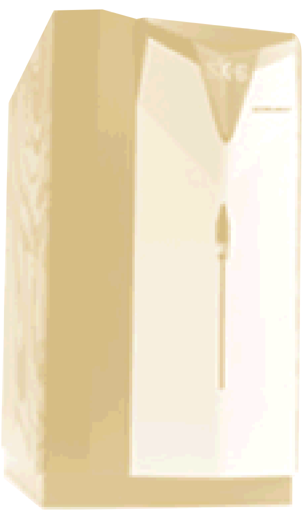

to the grand challenges arising from these huge scientific and technological barriers. Their goals were also to enhance supercomputing by stimulating the technology transfer from univer-sities and research institutions into industry and by increasing the fraction of the technical community which gets the opportunity to develop the skills required to efficiently access the high-performance computing resources.
Thus, scientific supercomputing as a discipline covering the fields of mathe-matical modeling, computer simula-tion, and visualization has become a strategic key technology. The growing complexity of the systems and pro-cesses under investigation in natural sciences and engineering results in simultaneously increasing requirements on the accuracy of mathematical modeling, the efficiency of numerical and non-numerical methods as well as on the performance of computer architectures and networks, and on the software systems, programming models, and programming techniques available for this purpose.
In industry, design optimization for improving products and accelerating development cycles is a major economic factor. Hence, competency in scientific computing has become an indispensable site factor of com-petitiveness. In the USA, strategic funding programmes push innova-tion: Apart from the programs and efforts under the NPACI initiative of the National Science Foundation, the US government has launched a big pro-grammatic offensive for the methodi-cal upgrade of scientific computing, i.e. the DoE's Accelerated Strategic Computing Initiative (ASCI), together with the national research laboratories of Los Alamos, Lawrence Livermore, and Sandia, universities, and indus-tries. Despite or even challenged by the success of the Japanese Earth Simulator Project which conquered the top of the TOP500 list in spring 2002, ASCI research and develop-ment will considerably increase the scientific and technological lead of the USA and enhance the instrumental equipment of their centres with new parallel supercomputers in a nearly unrivaled manner for years. However, these developments will challenge not only system reliability, availability and serviceability to novel levels, but also algorithm design and, in particular, adaptivity, accuracy, and stability of parallel numerical methods. In the future, research and industry will fall behind where such supercomputers of highest class will not be available. Therefore, Germany as a traditionally strong site of science and industry requires long-lasting measures to maintain and enhance scientific super-computing.

Recommendations of the German Science Council

In 1995 and 2000, the Science Council made recommendations with regard to the role of supercomputing for German science and technology. These recommendations clearly underline the significance of supercomputing in order to gain scientific know-how as well as to establish an efficient research infra-structure. At the top of the pyramid, sufficient supercomputing capacity is to be supplied by several national centres.

Access to these centres via efficient data networks must be ensured. Furthermore, the centres have to be embedded in a competence network in order to promote the development of methods, tools, and applications, and the education and training of young scientists and engineers.
In the Year 2000 Recommendation, the German Science Council stated that high performance computing (HPC) is indispensable for top research in the global competition of science, research, and engineering while the demand on HPC capacity tends to become unlimited due to the increasing problem complexity. Therefore, continuous investments are necessary on all levels of the performance pyramid for scientific computing. In addition, networks of competence must support the efficient usage of supercomputers, and HPC education and training must be enhanced, since so far HPC is yet insufficiently integrated into curricula; education must not be limited to PCs and workstations. In particular, efforts in HPC software development need to be enforced. Finally, strategic coordination of investments and procurements is necessary. Therefore, a National Coordination Committee was founded to develop strategies for continuous funding and upgrade scheduling of the HPC centres' resources (and balance potential regional ambitions according to national goals). This Coordination Committee started working in May 2001.
Its tasks and responsibilities are:

- Decisions on investments supported by prospective exploration of HPC demands, strategic advice for federal and "Länder" HPC plans and decisions as well as recommendations for upgrades in infrastructure and staffing.
- Orientational support by position papers and hearings on HPC issues and advice to the Science Council's HBFG Committee.
- Evolution of control and steering models by development and testing of demand-driven self-control mechanisms and investigations of differing accounting models for suitable user profiles.
- Develop a nation-wide concept for HPC provision including all centres independent of funding and institutional type and keeping the list of centres open for change and innovation.

The German Science Council favours in the Recommendation on the establishment and enhancement of national supercomputing centres a significant increase of the innovation frequency of the supercomputer hardware, however, without eliminating the financial and administrative barriers which force the individual centres to stick much too long with the systems. Therefore it is important to guarantee innovation also and primarily via other ways. The Recommendation proposes up to four national supercomputing centres which should be equipped to compete internationally at the apex of the performance pyramid. In the past, already four centres had qualified themselves with internationally outstanding profiles: HLRS in Stuttgart, LRZ in Munich, NIC at the Research Centre Jülich, and ZIB in Berlin. A single centre, which is either carried budget-wise by the Federal Government – like NIC – or according to the so-called HBFG Law by the corresponding federal country and the federal government with equal portions – like HLRS, LRZ, and ZIB – is not in the position today to place itself by its own financial power permanently in the pyramid apex, because in comparison to the long individual installation periods the decay of the once installed supercomputer power relative to the worldwide progress of supercomputer technology is extremely fast.
With respect to the dynamics of the supercomputer technology, effective concepts and procedures had to be developed to increase the flexibility in order to permanently provide science, research and engineering with access to the highest level of computing power. Thus, the intention of the Science Council to place Germany continuously at the apex of the performance pyramid must ground on a spiral of innovation which is steadily turned, with adequate phase shifts, by the three approved national supercomputing centres: HLRS, NIC, and LRZ, by coordinating and timely adapting their supercomputer procurements. In this way, it can be warrented that at any time the utmost novelty of supercomputer architecture and highest performance are available for science and research in Germany. This comprehensive strategy, of course, needs as necessary and sufficient condition and prerequisite a long-term guarantee and stabilization of the required budget frameworks by the federal as well as the local governments – which definitely has not yet been achieved, thus making every new procurement an ever challenging adventure.
In the same sense, as the Science Council outlined in corresponding recommendations, the demand of powerful data networks to support the interconnection of the supercomputing centres requests also the long-term government responsibility to maintain and enhance broadband communications in ever innovative steps in order to promote the key technology of scientific computing in science and industry. A corresponding long-range research program would be beneficial to improve the perspectives and the position of German science and technology as well as industry in the process of establishing enhanced broadband interconnections of supercomputing centres and competence centres to keep scientific computing in this country competitive within the EU and in interaction with the corresponding activity centres in USA and Japan.

Supercomputing Centres as Attractors of Competence

Definitely, at least in Germany, missing (super)computer hardware industry demands supercomputer centres to play the role of crystallization kernels and attractors of competence in scientific computing and in computational science and engineering. Also with respect to the developments in the USA it must be stated, that supercomputer centres are indispensable creators of competence structures as well as attractors of scientific computing. The John von Neumann Institute for Computing (NIC) – whose primary supercomputing resources are provided by the Central Institute for Applied Mathematics (ZAM) at the Research Centre Jülich –, the Supercomputer Centre (HLRS) at the University of Stuttgart, and the Leibniz Computer Centre (LRZ) act as National Supercomputer Centres along the strategic lines of the German Science Council by offering

computing capabilities to scientists and engineers in universities, research institutes, and industry all over Germany, and by promoting competence and skills in scientific applications, mathematical algorithms, and visualization methods through research projects and educational programs. Since the very beginning of their operation, the utilization by science and research of the resources in these supercomputer centres is steered via peer reviewing and controlling procedures certifying the quality and scientific relevance of the projects which get approved access to these supercomputer resources. Steering Committees and Scientific Councils associated with the individual centres play an eminent role. The exchange of computer resources between the centres needs, of course, mutual aggreements on the votes of these committees which must guarantee the necessary exchange of information on the project applications and research proposals. The usage of these resources by the industry is ruled via different special regulations.
It is extremely important that the cooperating centres increase their joint efforts to foster education and train-

ing in the field of scientific supercomputing via lectures, seminars, courses, workshops, colloquia, and symposia as well as by regular information exchange on the technical progress of architectures, software, and algorithms.
They are prominent institutions to provide an efficient national scientific forum for the research groups in the competence centres and the project partners in science and industry in order to promote joint projects, to stimulate innovative research topics and facilitate interdisciplinary cooperations. Therefore, the question is becoming urgent whether and how the national supercomputing centres can be effectively integrated, together with the various competence centres of scientific computing, into a cooperative compound to further push innovative supercomputing methods and applications. The compound will act as a common turntable to strengthen and accelerate the transfer of the capabilities, techniques, and skills in scientific supercomputing into the applications in industry. This will increase mutual benefits and added values, which, finally, protect and perpetuate the high level of supercomputing in Germany in the international competition of science and economy, while the modern concepts and requirements of process modelling and product design, the "virtual laboratory" and "virtual product", are based on the potential of visualization and multimedia technologies, communication techniques, and cooperative work distribution.

Therefore, scientific computing is basically interdisciplinary. This innovative technology requires, in addition to solid knowledge and systematic know-how of the individual scientific and technical disciplines, the complete information-technological instruments of supercomputing and communications. The technology requests large performance steps in the computer architectures and algorithms. There is no doubt that the interaction of modelling, computer simulation and visualization with mathematical methodology and computer science has made the strategic interdiscipline of scientific supercomputing a key technology which definitely contributes to the strength of the international competitiveness of science and economy.
A major step in the support of science and research including industrial research projects by supercomputing capability and capacity was made in 1987 when the so-called HLRZ ("Höchstleistungsrechenzentrum") was founded by the Research Centre Jülich, DESY, and GMD. The Central Institute for Applied Mathematics in Jülich was and is the main carrier of this supercomputing centre which since 1998 is operated as the John von Neumann Institute for Computing (NIC) primarily by the Jülich Research Centre with DESY as partner who provides special purpose computers (of APE type) to the QCD community.
It definitely was a novelty in this country that supercomputer capacity could be provided to scientists all over Germany, based on the approval of peer-reviewed project proposals according to the DFG rules and regulations – a model which then was applied also to the foundation of the other two National Supercomputing Centres in Stuttgart and Munich in the 90-s. The establishment of the Stuttgart supercomputer centre HLRS in tight connection with the local industry Daimler and Porsche and the University of Karlsruhe via the joint venture "hww" with the debis company, now as T-Systems a daughter of the German Telekom, and DLR again represents a novel model of cooperation between science and industry on the strategic field of scientific supercomputing.

Synergy through Grid Technologies

In the past, technical and administrative conditions as well as the federal structure in Germany did not appear to ease the Germany-wide interconnection of computer centres as a natural instrument of scientific cooperation nor of economic exploitation of computer resources. Exceptions have been, on the one hand, the so-called North-German Vectorcomputer Compound (NVV) which, based on a state treatise of three federal countries, managed the joint usage and exchange of high-performance computer resources between ZIB in Berlin, the Regional Computer Centre (RRZN) in Hannover and the University Computer Centre in Kiel in the same manner as nowadays practiced with the two supercomputer subsystems at ZIB and RRZN extending the compound to six federal countries in Northern Germany; on the other hand, NIC in Jülich was able, according to the regulations within the Helmholtz Association of German Research Centres (HGF), to act beyond federal boundaries as a first national supercomputing centre until in the 90-s the two national supercomputing centres in Stuttgart and Munich got approved by the Science Council to offer their resources Germany-wide.
Today, a comprehensive competence structure provides beneficial conditions for the connection of supercomputing centres to exploit new concepts and paradigms of computing like grid computing on broadband communication

Innovation

networks. Thus, these supercomputing centres can

- promote supercomputing and its applications through interaction with other competence centres for scientific computing and their tight integration;
- enter more efficient cooperations through joint projects in scientific supercomputing, mathematical modelling, simulation methodology, and visualization;
- even better support and accelerate, with the compound of the centres as a turntable, the transfer of methods and technologies of this key technology into the application areas and product development in economy and industry;
- increase the efficient usage and, hence, the added value of their diverse and heterogeneous computer architectures through the tight connection via broadband networks and by the development of adequate interfaces and operations concepts;
- create and provide the scientific and technical platform, through the tight cooperation and communication compound of the centres, for future evolutionary steps of broadband communications, as offered by the DFN through G-WiN, and innovations via pilot testbed projects exploring, for instance, optical technology;
- can assure, through adequate financing strategies, mutual agreements, and coordination procedures, the permanent international top position of the supercomputing centres in Germany, thus contributing strongly to the international competetiveness of science and engineering.

The idea of a compound of the supercomputing-centres does not assume that in this way, as has been stated sometimes elsewhere in the past, the large computers in the centres, finally, could be ultimately efficiently utilized by a Germany-wide load-balancing model and, hence, the supposedly misguided local resource allocations get solved. The idea is rather the functional compound which exploits the diversity of supercomputer architectures and the broad and heterogeneous spectrum of large software systems. Technical and architectural reasons as well as licencing and economic criteria and aspects of competence, capability, and technical support make it sensible that the individual centres gradually specialize in (commercial) application software focussing on certain application fields. Such a functional compound of supercomputing centres will overcome the regional barriers which made it difficult in the past to provide resources on a reasonal mutual exchange basis to nonlocal users beyond those boundaries. In the future, one can guarantee such resource exchange via equivalent shares which can be dynamically controlled with respect to their mutual currencies through the already now installed sophisticated accounting systems

One of the major and most decisive factors of success of top research and, as a consequence, of industry and society in the future will be rapid and location-independent access to worldwide distributed data and other information technological (IT) resources and their sensible usage. The basis will be a highly complex, globally interlinked information infrastructure with complex access and administration mechanisms, which is commonly referred to as "grid computing" and, in UK, as "E-science". Grid computing or grid means the distributed IT infrastructure, including its access mechanisms. The structural concept of grid computing results from the necessity of combining supercomputers in a global functional system network of supercomputing centres and the integration of competence centres as well as from the worldwide linking of clusters of thousands of commodity processors.

Grid computing does not only set extreme requirements on broadband communications, but also needs the development of novel software systems for the "seamless access" to the distributed resources. In Germany, interconnection of supercomputers via grid technology was started with the UNICORE system already in 1996. For the global networks also other so-called open systems are being developed worldwide, e. g. Globus Toolkit. The grid or E-science concept additionally comprises application of grid technology to various disciplines, such as high-energy physics and astrophysics, biological sciences and medicine, or geo-research and environmental research, with the integration of large experimental devices (accelerators, satellites, tomographs, telescopes, and others) on the one hand and the setup of regionally distributed mass data archives and their supply for sophisticated information analysis (e.g. of accelerator data – from CERN – or satellite measurements) on the other hand. The initiative for the development of the UNICORE software system was launched in 1996 at ZAM in Jülich and funded by the BMBF through the UNICORE and the UNICORE Plus projects from 1997 to 1999 and from 2000 to 2002, respectively. With this software system, the concept of grid computing is provided with an effective "seamless interface". The software is also serving as a technical basis of several EU-funded grid projects like EUROGRID (European network of supercomputing centres), OpenMolGrid (Molecular Biology), and GRIP, ZAM being an important partner in these projects.

Recently, an initiative has been launched by the Helmholtz Association, in coordination with the other German science and research institutions and universities, in order to create and promote a new infrastructure which, as in other countries, is apt to revolutionize scientific research and development in Germany: the D-Grid. The tight interconnection of the National Supercomputer Centres and the integration of competence centres of scientific computing will definitely be one of the important missions of D-Grid.

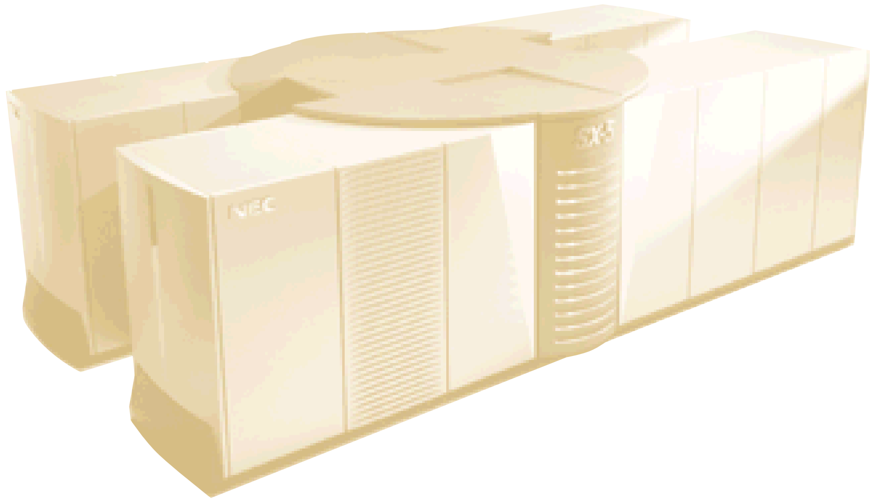

•Prof. Dr. F. Hoßfeld

Zentralinstitut für Angewandte Mathematik (ZAM)
John von Neumann–Institut für Computing (NIC)
Forschungszentrum Jülich

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
High-Performance Systems Department

Dr. Horst-Dieter Steinhöfer
Barer Straße 21
80333 München
Germany
Phone +49 89 28 92 87 79
steinhoefer@lrz.de
www.lrz-muenchen.de

View of the Hitachi SR8000-F1 at LRZ.

Compute servers currently operated by LRZ are

System	Size	Peak Performance Performance from Memory* (GFlop/s)	Purpose	User Community
Hitachi SR8000-F1	168 8-way SMP nodes 1376 GByte memory	2016 247*	Capability computing	German universities and research institutes
Fujitsu/ Siemens VPP700	52 vector processors 144 GByte memory	114 38*	Capability and capacity computing	Bavarian universities
Linux Cluster Intel IA32	48 nodes 88 IA32 processors 78 GByte memory	135 4*	Capacity computing	Munich universities
Linux Cluster Intel IA32 Intel IA64 (2Q2003)	90 IA32 single CPU 16 IA64 4-way CPU 218 GByte memory	693 20*	Capability and capacity computing	Munich universities
IBM pSeries 690 hpc	1 SMP node 8 processors POWER 4 32 GBytes memory	42 2*	Capacity computing	Munich universities

Centers

HLRS

Based on a long tradition in supercomputing at Stuttgart University, 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.

Contact:
HLRS
High-Performance Computing Center
Stuttgart

Prof. Dr. Michael M. Resch
Allmandring 30
70500 Stuttgart
Germany
Phone +49 71 16 85 25 04
resch@hlrs.de
www.hlrs.de

View of the NEC SX-5/32 M2e at HLRS.

Compute servers currently operated by HLRS are

System	Size	Peak Performance (GFlop/s)	Purpose	User Community
NEC SX-5/32 M2e	2 16-way nodes 80 GByte memory	128	Capability computing	German universities, research institutes, and industry
Cray T3E-900/512	512 nodes 64 GByte memory	460,8	Capability computing	German universities, research institutes, and industry
Hitachi SR8000	16 8-way nodes 128 GByte memory	128	Capability and capacity computing	German universities, research institutes, and industry
NEC Azusa	16 IA64 32 GByte memory	51,2	Capacity computing	Universität Stuttgart
IA64 Cluster	8 2-way nodes 32 GBytes memory	57,6	Capacity computing	Universität Stuttgart
IA32 Cluster	24 2-way nodes 48 GBytes memory	230,4	Capacity computing	Universität Stuttgart

NIC

The John von Neumann Institute for Computing (NIC) is a joint foundation of Forschungszentrum Jülich and Deutsches Elektronen-Synchrotron DESY to support supercomputer-aided scientific research and development in Germany. NIC takes over the functions and tasks of the High Performance Computer Centre (HLRZ) established in 1987 and continues this centre's successful work in the field of supercomputing and its applications.

Nationwide 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 Centre for Parallel Computing of DESY at 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, research groups exist for high energy physics and complex systems; another research group in the field of "Bioinformatics" is under consideration.

Education and training in the fields of supercomputing by symposia, workshops, summer schools, seminars and courses.

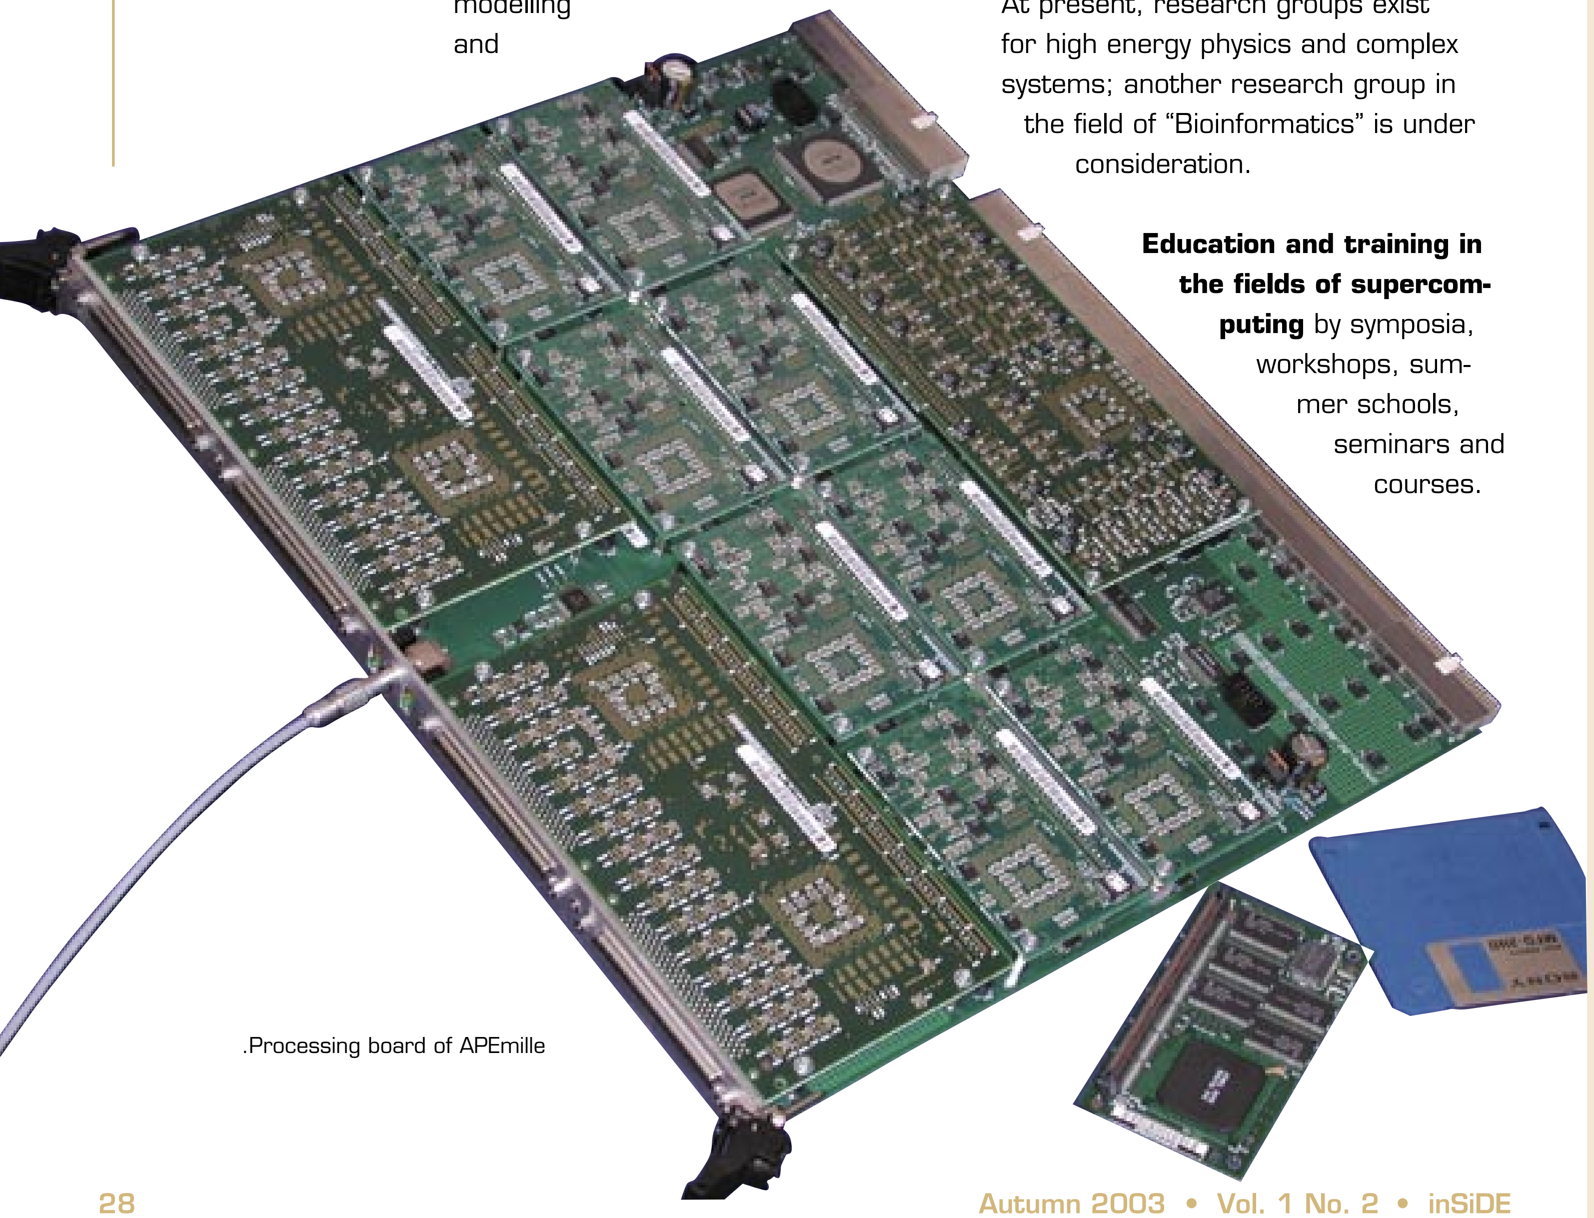

.Processing board of APEmille

The following supercomputers are available for research projects of the communities mentioned below, peer-reviewed by the Resource Allocation Committee of NIC. A more detailed description of the supercomputers can be found on the web servers of the Research Centre Jülich or of the German Electron Synchrotron DESY, respectively:

http://www.fz-juelich.de/zam/CompServ/services/sco.html
http://www-zeuthen.desy.de/ape/html/Installation/

System	Size	Peak Performance (GFlop/s)	Purpose	User Community
IBM pSeries 690 Cluster 1600 (3Q2003 - 1Q2004)	6 SMP nodes 192 processors POWER4+ 384 GBytes memory	1300	Capability computing	German universities, research institutes, and industry
IBM pSeries 690 Cluster 1600 (from 1Q2004)	37 SMP nodes 1120 processors POWER4+ 2240 GBytes memory	8100	Capability computing	German universities, research institutes, and industry
CRAY T3E-1200	512 nodes 262 GBytes memory	614	Capability computing	German universities, research institutes, and industry
CRAY SV1ex	16 CPUs 32 GBytes memory	32	Capability computing	German universities, research institutes, and industry
APEmille (special purpose computers)	4 racks 1124 processors 32 GBytes memory	550	Capability computing	Lattice gauge theory groups at German universities and research institutes

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

Dr. Rüdiger Esser
Wilhelm-Jonen-Straße
52425 Jülich
Germany
Phone +49 24 61 61 65 88
r.esser@fz-juelich.de
www.fz-juelich.de/zam/CompServ/services/sco.html

High-Performance Computing Courses and Tutorials

LRZ
www.lrz.de

Parallel Programming of High-Performance Computers

Date and Time:
October 13–14, 2003,
9:00 am to 5:00 pm (talks) and
October 15, 2003,
9:00 am to 5:00 pm (hands-on sessions)

Location:
Erlangen, RRZE

Contents:
- IA32 and IA64: architecture and systems
- distributed and shared memory programming paradigms
- MPI-1
- MPI-2 survey
- OpenMP

Prerequisites: knowledge of Fortran, C, or C++ as well as general UNIX skills

Special Topics of Fortran90 Programming

Date and Time:
October 16, 2003, 9:00 am to 17:00 pm

Location:
Erlangen, RRZE

Contents:
Focus is on topics of high performance computing.
Prerequisite: a firm grasp of Fortran

Object-oriented Programming with C++

Date and Time:
January 19–22, 2004,
8:30 am to 12:30 pm

Location:
LRZ, lecture room PEP

Contents:
- overview of object-oriented programming
- C versus C++
- I/O in C++
- function and operator overloading
- classes, objects, and inheritance
- templates and the C++ standard library
- introduction to advanced template techniques (template expressions)

This course is primarily aimed at programmers in the field of technical and scientific computing. Prerequisites: average knowledge of C or Java

Basics and Tools of High-Performance Computing

Date and Time:
February 2, 2004, 9:30 am to 18:00 pm

Location:
LRZ, lecture room PEP

Contents:
This course gives an introduction to common tools on high performance computers.
- HPC in Germany
- software packages on high performance computers
- secure connections with SSH und OpenSSH
- introduction to (GNU)make
- source code management with CVS
- basics of performance measurements

Prerequisites: basic knowledge of Unix and at least one programming language

TotalView: A Universal Debugger

Date and Time:
February 4, 2004, 9:30 am to 13:00 pm

Location:
LRZ, lecture room 1st floor

Contents:
Introduction to the usage of the TotalView debugger.
This tool is available on all high performance computers installed at LRZ.
Hands-on sessions will give you additional experience.
Prerequisites: knowledge of Fortran or C as well as Unix

VAMPIR: Diagnosing Parallel Programs

Date and Time:
February 5, 2004, 9:30 am to 12:30 pm

Location:
LRZ, lecture room 1st floor

Contents:
VAMPIR is a tool for diagnosing communication bottlenecks in MPI programs.
Usage of trace libraries las well as evaluation of the generated trace files are explained and trained in a hands-on session.
Prerequisites: MPI programming with Fortran and C

HLRS
www.hlrs.de

Parallel Programming with MPI and OpenMP

Date:
February 17–19, 2004

Location:
Dresden, ZHR

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 ZHR in collaboration with HLRS

Iterative Linear Solvers and Parallelization

Date:
March 8–12, 2004

Location:
Stuttgart, HLRS

Contents:
The focus is on iterative and parallel solvers, the parallel programming models MPI and OpenMP, and the parallel middleware PETSc. Thereby, different modern Krylov Subspace Methods (CG, GMRES, BiCGSTAB ...) as well as highly efficient preconditioning techniques are presented. Hands-on sessions (in C and Fortran) will allow users to immediately test and understand iterative solvers, MPI and OpenMP. This course is organized by HLRS, IAG, WIR, and Uni. Kassel.

Advanced Topics in High-Performance Comuting and Parallel Programming

Date:
March 15–19, 2004

Location:
Stuttgart, HLRS

Contents:
The focus is on vectorization (Mon+Tue), MPI-2 and MPI Performance Optimization (Wed), advanced topics in parallel programming with MPI (Thu), and advanced topics in shared memory parallelization with OpenMP (Fri). Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the constructs.
Seperate registration of individual days is possible

Introduction to Computational Fluids Dynamics

Date:
March 29 – April 2, 2004

Location:
Stuttgart, HLRS

Contents:
Numerical methods in Fluid Dynamics are presented. The focus is on explicit Finite Volume schemes for the compressible Euler equations. Hands-on sessions will allow participants to implement the algorithms, to apply commercial software, and to interpret the solutions correctly. This course is organized by HLRS, IAG, WIR, and Uni. Kassel, and is based on a computational practical awarded with the "Landeslehrpreis Baden-Württemberg 2003" (held at Uni. Stuttgart)

Iterative Linear Solvers and Parallelization

Date:
September 13–17, 2004

Location:
University of Kassel

Contents:
The focus is on iterative and parallel solvers, the parallel programming models MPI and OpenMP, and the parallel middleware PETSc. Thereby, different modern Krylov Subspace Methods (CG, GMRES, BiCGSTAB ...) as well as highly efficient preconditioning techniques are presented in the context of real life applications. Hands-on sessions (in C and Fortran) will allow users to immediately test and understand the basic constructs of iterative solvers, the Message Passing Interface (MPI) and the shared memory directives of OpenMP. This course is organized by Uni Kassel, HLRS, IAG, and WIR

HLRS Parallel Programming Workshop

Date:
October 11–15, 2004

Location:
Stuttgart, HLRS

Contents:
The focus is on programming models MPI and OpenMP, domain decomposition, parallelization with PETSc, and optimization. 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)

NIC
www.fz-juelich.de/nic

Grid Computing Workshop Access Grid Presentation in the frame of SC Global 2003

Date:
November 20, 2003

Location:
ZAM (Video Conference Room), Research Centre Jülich

User Course "Programming and Usage of the System IBM pSeries 690 Cluster 1600" Common Presentation of ZAM and IBM

Date:
January 14–16, 2004

Location:
Research Centre Jülich

Official Opening of the New IBM Supercomputer Facility

Date:
February 16, 2004

Location:
Research Centre Jülich

NIC Symposium 2004 Result Presentation of Invited NIC Users

Date:
February 17–18, 2004

Location:
Research Centre Jülich

Winter School "Computational Soft Matter: From Synthetic Polymers to Proteins"

Date:
February 29 – March 6, 2004

Location:
Gustav-Stresemann-Institut, Bonn, organized by NIC, Research Centre Jülich

inSiDE

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

Publishers:
Prof. Dr. Heinz-Gerd Hegering, LRZ
Dr. Rüdiger Esser, NIC
Prof. Dr. Michael M. Resch, HLRS

Editor:
F. Rainer Klank, RUS/HLRS
klank@rus.uni-stuttgart.de

Design:
Katharina Schlatterer
kschlatterer@rus.uni-stuttgart.de

Authors:
Jaakko Akola
j.akola@fz-juelich.de
Norbert Attig
n.attig@fz-juelich.de
Marc Freitag
freitag@ari.uni-heidelberg.de
Petra Hardt
hardt@bv.tum.de
Thomas Hartlep
thomas.hartlep@
physik.uni-goettingen.de
J. Leo van Hemmen
lvh@ph.tum.de
Friedel Hoßfeld
f.hossfeld@fz-juelich.de
Robert O. Jones
r.jones@fz-juelich.de
Siegfried Kühner
kuehner@bv.tum.de
Joachim Noll
jnoll@physik.tu-muenchen.de
Ernst Rank
rank@bv.tum.de
Gerrit Schierholz
Gerrit.Schierholz@desy.de
Oliver Wenisch
owenisch@physik.tu-muenchen.de