inSiDE

inSiDE • Vol. 6 No.1 • Spring 2008

Innovatives Supercomputing in Deutschland

Editorial

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

Welcome to this first issue of inSiDE in 2008. German supercomputing is keeping the pace of the last year with the Gauss Centre for Supercomputing taking the lead. The collaboration between the three national supercomputing centres (HLRS, LRZ and NIC) is making good progress. At the same time European supercomputing has seen a boost. In this issue you will read about the most recent developments in the field in Germany.

In January the European supercomputing infrastructure project PRACE was started and is making excellent progress led by Prof. Achim Bachem from the Gauss Centre. In February the fastest civil supercomputer JUGENE was officially unveiled at the Jülich Research Centre. The IBM BlueGene/P system will provide German and European researchers with a peak speed of 223 TFlop/s. With this new installation Germany is knocking at the PFlops-door. In March the University of Stuttgart and HLRS together with its partners from industry founded the Automotive Simulation Centre Stuttgart (ASCS). The focus of ASCS is on simulation in automotive engineering that require high performance computing. The ultimate goal is to support automotive industry in meeting the challenging targets in CO_2 reduction as put forward by the European Commission recently.

All these activities are evolutionary steps towards a German and European hardware and software landscape that is able to provide researchers with a competitive cyberinfrastructure which allows them to see eye-to-eye with colleagues in the US and Japan. In addition to these activities over the last months the Federal Ministry of Education and Research (BMBF) has issued a call for proposals for the development of software in the field of high performance computing. The projects within this software framework are expected to start in autumn 2008 and inSiDE will keep reporting on this.

Again we have included a section on applications running on one of the three main German supercomputing systems. In this issue the reader will find six articles about the various fields of applications and how they exploit the various architectures made available. The section not only reflects the wide variety of simulation research in Germany but at the same time nicely shows how various architectures provide users with the best hardware technology for all of their problems.

In the project section we present a look at one of the key projects in Europe for the usage of high performance computers DEISA. This has recently been extended and you will find a report on the new and extended focus here. The close cooperation between the supercomputing community and the Grid community is reflected in the project report on SmartLM which focuses on Grid-friendly software licensing for location independent application execution. The importance of visualization tools to understand the results of simulation is shown in the IRMOS project. It introduces interactive real-time multimedia applications on service oriented infrastructures. A further specifically interesting project focusing on optimal load balance in parallel programs for geophysical simulations is also featured.

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

Contents

Automotive Simulation Centre Stuttgart (ASCS) founded

A step forward in bringing high performance computing to automotive industry was made on March 7th, 2008 at Stuttgart. Private and public partners joined the University of Stuttgart in creating a centre for the collaboration in automotive simulation on high performance computing systems. For the University of Stuttgart this is a further step to strengthen its industrial collaboration and to focus the activities of its High Performance Computing Centre Stuttgart (HLRS) in the field of computational engineering. For automotive industry, software and hardware vendors this is a step closer to public research organizations. The common driving force is the need to substantially reduce CO_2 emissions while sustaining an economically leading position in the field.

Dr. Thomas Weber, member of the board of Daimler AG and of the board of the University of Stuttgart said: "This first transfer centre between the University of Stuttgart and industry will further strengthen our competence in automotive and transportation technology and will further improve the competitiveness of Germany as a location for research and development". Added Wolfgang Dürnberger, member of the board of Porsche AG: "ASCS has a pilot character for the collaboration of research and industry and is setting an example for Europe." Prof. Wolfram Ressel, President of the University of Stuttgart, who had been a driving force behind the ASCS said: "The big interest of industry in such a collaboration again emphasizes the outstanding level of quality and performance of research at the University of Stuttgart."

President Prof. Wolfram Ressel at the foundation of ASCS

ASCS will focus on pre-competitive research in fields that require simulation on high performance computing systems. First projects are currently discussed with the board of ASCS. Among the most important problems are multidisciplinary optimization and combustion. Both require a high level of compute performance and expertise in computational engineering on high performance computers available at the HLRS.

Further topics to be covered by the projects of ASCS will be material science investigations and the usage of simulation and high performance computing in production planning and optimization. By bringing together public and private research ASCS aims at shortening the innovation cycle in automotive simulation technology thus increasing the quality of automotive products. Initial funding of ASCS is provided by the University of Stuttgart and the State of Baden-Württemberg over a period of three years. Industrial partners will bring in additional funding for this first startup period. After that ASCS will work based on public and private funding and will be self-sustaining.

Members of the ASCS contribute to the funding with an annual contribution.

The founding partners of ASCS are: Abaqus Deutschland GmbH, Adam Opel GmbH, Altair Engineering GmbH, Cray Computer Deutschland GmbH, Daimler AG, Dr.-Ing. h.c. F. Porsche AG, Dynamore, Engineous GmbH, Research Institute for Automotive and Engine Design (FKFS), INTES GmbH, the University of Stuttgart, The Virtual Dimension Centre Fellbach, Wilhelm Karman GmbH. Only four weeks after the foundation NEC Europe joined the centre as a full member. Collaboration with IBM and further hardware vendors is currently discussed.

The centre is the first in a series of simulation centres that the State of Baden-Württemberg wants to introduce in order to improve the know-how transfer between research and industry over the next years. Further similar centres are planned for energy simulation, chemical simulation, and biomedical simulation.

Founding members at the foundation ceremony of ASCS on March 7th, 2008

News

JUGENE officially unveiled

On February 22th, 2008, the supercomputer IBM Blue Gene/P JUGENE was officially unveiled at Research Centre Jülich. More than 250 guests attended the inauguration ceremony which found its highlight in the go-ahead given by Prime Minister Dr. Jürgen Rüttgers together with State Secretary Thomas Rachel.

Prof. Bachem, Chairman of the Board of Directors of Research Centre Jülich, gave the welcoming address. “Science and industry need computing power of the highest quality”, said Bachem. “With JUGENE, we have now set another milestone in Jülich for cutting-edge research. Supercomputer simulations have become a key technology. They function as virtual microscopes and telescopes and enhance our understanding in the subatomic region and in the universe, respectively.”

North-Rhine Westphalia’s Prime Minister Rüttgers could not conceal his pride. “Jülich not only operates the fastest civil supercomputer”, he said approvingly, “but it also offers a network of education opportunities and top research in the simulation sciences. There is a vigorous, multi-faceted research and technology landscape in this area that makes North-Rhine Westphalia the number one in innovation.”

Thomas Rachel, State Secretary of the Federal Ministry of Education and Research, explained the significance of the new supercomputer for Germany: “With JUGENE, the German position within Europe has been strengthened. The German Gauss Centre for Supercomputing will take on a leading role in establishing a European supercomputing centre.”

Prof. Thomas Lippert, Director of the Jülich Supercomputing Centre, pointed out how supercomputer users can benefit from the dedicated topic-specific support in Jülich. “JSC‘s staff do not just ensure that JUGENE is ready for use; they also provide expertise on all aspects of the simulation sciences. In simulation laboratories, for example, scientific simulation know-how is provided for the topics of plasma physics, biology, earth system science and the nanosciences.”

In a short discussion, representatives of IBM, Intel and EADS took a look at the importance of supercomputing from industry’s point of view. Video clips on hardware, applications, users, and project partners in science and industry livened up the discussions and interviews.

PRACE Project launched

PRACE, the Partnership for Advanced Computing in Europe, has reached another milestone in its mission to create a European high-performance computing infrastructure. At the end of 2007 the project received a grant of € 10 million from the European Commission towards its total budget of € 20 million for the next two years. On January 29th to 30th, 2008, the kick-off meeting of the project – which is coordinated by Forschungszentrum Jülich – took place at JSC. 74 participants from 14 European countries attended the meeting. Thomas Rachel, Parliamentary State Secretary at the Federal Ministry of Education and Research, opened the event.

PRACE has been established to create a permanent pan-European High Performance Computing service for research. In the preparatory phase, which runs until the end of 2009, the project will establish the basis for a transnational organizational structure for scientific supercomputing in Europe. By bringing together the know-how and resources of the partners, PRACE will provide European researchers with access to supercomputing capacities at a world-class level, transcending those affordable at the national level. This includes a coordinated approach to hardware procurement and potentially a European platform for the development of hardware and software jointly with industry. Close co-operation with national and regional computer centres and scientific organizations will ease access to computing resources at all levels for scientists and engineers from academia and industry.

To achieve these challenging goals, the researchers from the partner organizations firmed up details of the project workplan during the kick-off meeting at Jülich. One task is to define a suitable legal form and organizational structure for the permanent European HPC infrastructure. Key to the success of the project are the technical developments required to enable operation of a distributed supercomputing infrastructure, the scaling and optimisation of application software, and the evaluation of prototypes of future computers. PRACE aims to install a petaflop system as early as 2009.

The following countries are collaborating in the PRACE project: Germany, UK, France, Spain, Finland, Greece, Italy, The Netherlands, Norway, Austria, Poland, Portugal, Sweden, and Switzerland. Germany is represented in PRACE through the Gauss Centre for Supercomputing, which focuses the activities of the three HPC centres in Jülich, Stuttgart, and Garching. Through close co-operation with established European research organizations such as ESF, EMBL, and ESA, PRACE will be embedded into the European Research Area.

The PRACE project receives funding from the European Community’s Seventh Framework Programme (FP7/2007-2013) under grant agreement no. RI-211528. Additional information: www.prace-project.eu

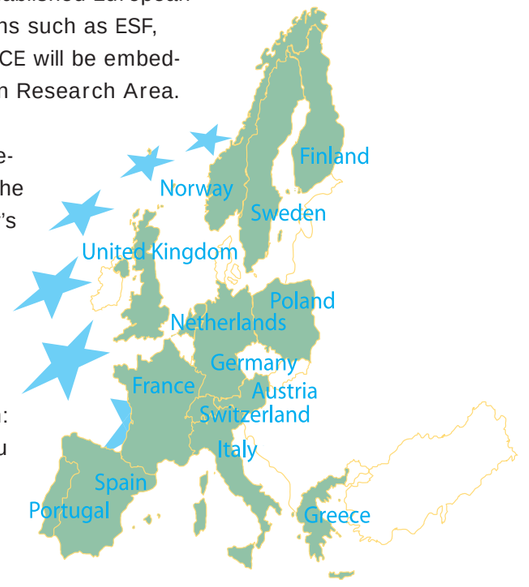

Drag Reduction by Dimples? An Investigation Relying Upon HPC

Introducing a regular arrangement of surface depressions called dimples is a well-known measure to increase heat transfer from a wall. Compared to a smooth wall, the Nusselt number can be significantly enhanced by dimples, whereas the increase of the pressure drop was found to be small. For that purpose, deep dimples with a ratio of depth to print diameter of h / D = 0.2 - 0.5 are typically applied to enhance the convective transport from the wall.

Some years back Russian scientists discovered that apart from heat-transfer enhancement dimples might be useful for drag reduction. Their findings are based on experimental measurements of turbulent flows over surfaces with a regular arrangement of shallow dimples leading to a decrease of the skin-friction drag of a turbulent flow up to 20%. (Note that this is not the “golf ball“ effect!). However, up to now no clear explanation has been provided which can illuminate this effect. Hence some doubts remain regarding the experimental investigations and their outcome. Compared to classical devices for drag reduction, such as riblets, shallow dimples would be advantageous since they are composed of macroscopic structures which are less sensitive to dirt and mechanical degradation. Furthermore, their effectiveness should not depend on the flow direction.

Since the drag reduction promised by dimples would have a tremendous economical influence, it is worth studying it in more detail. For that purpose, experimental investigations (see [1]) as well as direct numerical simulations (DNS) of the turbulent flow inside a channel with dimpled walls were carried out to investigate the physical mechanism.

Direct Numerical Simulations

The direct numerical simulations were carried out with the CFD code LESOCC which was developed for the simulation of complex turbulent flows. LESOCC is based on a 3-D finite-volume method for arbitrary non-orthogonal and non-staggered, block-structured grids [2,3]. The spatial discretization of all fluxes is based on central differences of second-order accuracy. A low-storage multi-stage Runge-Kutta method (second-order accurate) is applied for time-marching. LESOCC is highly vectorized (> 99.8%) and additionally parallelized by domain decomposition using MPI. The present simulations were carried out on the NEC SX-8 machine at HLRS Stuttgart using 8 processors of one node leading to a total performance of about 52.7 GFlop / s. For details about vectorization and parallelization we refer to [4].

In order to investigate the effect of dimples, a classical wall-bounded flow often studied in literature was considered, i.e., a turbulent plane channel flow at Re = 10,935. This test case has several advantages. With smooth non-dimpled walls the flow is homogeneous in streamwise and spanwise directions. This allows to use periodic boundary conditions in both directions and thus avoids the definition of appropriate inflow and outflow boundary conditions. The pressure gradient in the streamwise direction was adjusted such that a fixed mass flow rate was assured. Furthermore, no-slip boundary conditions were used at both walls.

Since the simulation cannot cover the entire experimental set-up, the computational domain consisted of a cutout of the channel maintaining periodic boundary conditions. Three cases were considered:

[L:] Plane channel with multiple dimples at the lower wall

[B:] Plane channel with multiple dimples at both walls

[R:] Plane channel without dimples, i.e., reference case with smooth walls

Two geometries of dimples were tested, both supplied by Inventors Network GmbH, the holder of the patents and representative of the Russian scientists. The results reported in the following refer to the smaller dimple geometry, which was provided as the first choice. The print diameter of these smaller dimples was D = 15 mm and the depth h = 0.75 mm. As shown in Figure 1, multiple shallow dimples (depth to print diameter of h / D = 0.05) were arranged regularly on the surface of the lower and / or upper channel walls.

Using block-structured curvilinear grids, two different grid resolutions were taken into account. The first, denoted coarse grid, consisted of about 5.5 million control volumes (CVs) and thus was not really coarse. For the second grid, denoted fine grid, the number of control volumes was exactly doubled

Figure 1: Channel wall with milled dimples

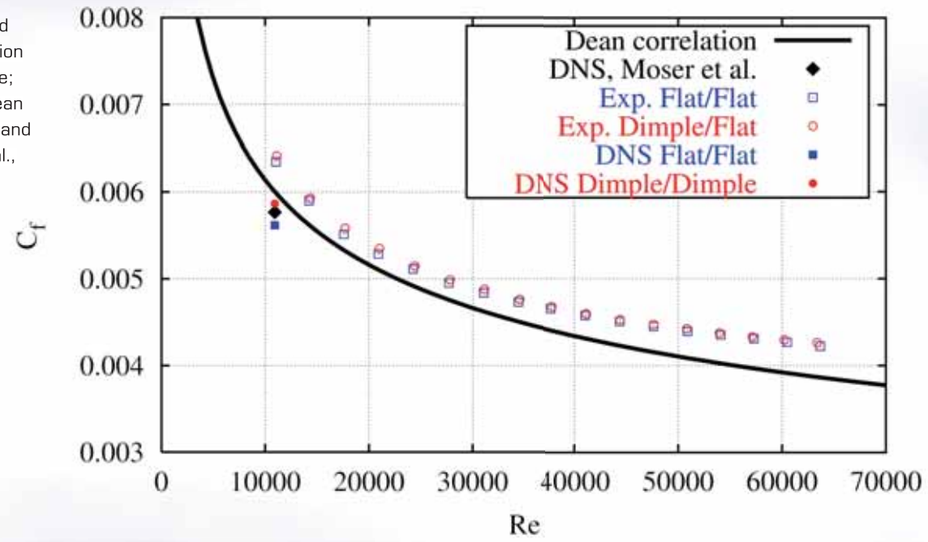

Figure 2: Measured and predicted friction coefficient C_f vs. Re; reference data: Dean correlation (1978) and DNS by Moser et al., 1999

in each direction. In total about 43 million CVs were used and the resolution in each direction was halved yielding an excellent resolution of the viscous sublayer. For the cases with dimples the grid was clinging to the curvilinear geometry of the walls and thus was locally no longer Cartesian. To achieve reliable statistical data, the flow was averaged over long time intervals.

Results

The skin-friction coefficients measured in the channel flow (see [1]) were plotted as a function of the Reynolds number. Figure 2 presents a summary of the experimental results for the smaller dimples together with data points derived from integration of the numerically obtained wall quantities. For comparison, data available in the literature (Dean, 1978; Moser et al., 1999) was added. The experimental skin-friction coefficients for the channel flow with smooth walls and with a dimpled wall on one side almost collapse. They differ by less than 1% through all the Reynolds number range considered. Therefore, it may be concluded that, surprisingly, the dimpled wall did not show any increase in pressure loss, but a significant decrease could also not be detected. The second dimple geometry performed even worse and resulted in a definite increase of the pressure loss in the channel flow. Additionally, an external flow test case was set-up, i.e. a zero-pressure-gradient flat plate boundary layer flow. This again supports the conclusion drawn from the channel flow experiments, namely, that the dimples did not influence drag, neither to the better nor to the worse.

More insight into this not very satisfactory result regarding the integral behavior of the dimpled surface compared to the flat walls was sought using CFD. Since only minor deviations were found between the results obtained on the coarse and the fine grid, in the following only the latter are discussed.

Figure 3 depicts the time-averaged pressure distribution obtained for the case B using a regular arrangement of shallow dimples on both walls. The influence of the dimples on the pressure distribution on the lower wall can clearly be seen. It should be mentioned that the linear pressure gradient in the streamwise direction of the channel is not included in the figure. In Figure 4 the time-averaged wall-shear stress distribution is displayed. Obviously the wall-shear stresses decrease within the dimples, leading to a very small recirculation region at the falling edge (see Figs. 5 and 6). However, at the downstream edge of the dimples, where the fluid flow leaves the surface depression again, large values of the wall-shear stress are found.

In Figure 5 the streamlines at the lower walls defined by the average surface shear stress are displayed in conjunction with the distribution of the time-averaged spanwise velocity close to the wall. Upstream of the dimple and partially also in the dimple the streamlines converge. A tiny recirculation region visible also in Figure 6 is found inside the dimple. The existence of this phenomenon is confirmed by the results on both grids. At the side and downstream of the recirculation region the streamlines diverge again.

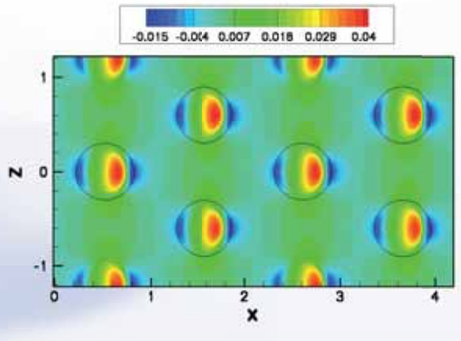

Figure 3: Channel flow with multiple dimples at both walls, time-averaged pressure

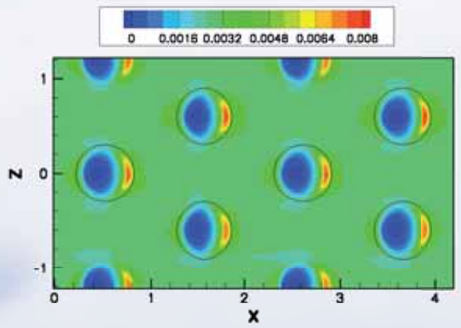

Figure 4: Channel flow with multiple dimples at both walls, time-averaged wall-shear stress

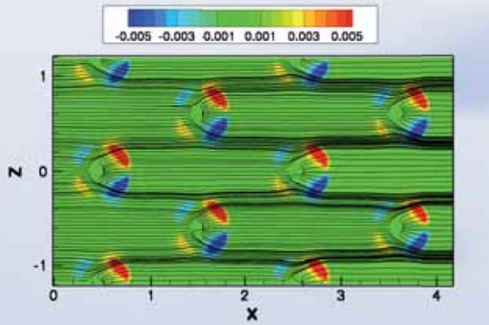

Figure 5: Channel flow with multiple dimples at both walls, wall streamlines

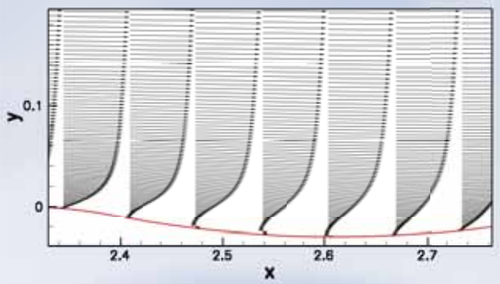

Figure 6: Zoom of the dimple region in a x-y midplane

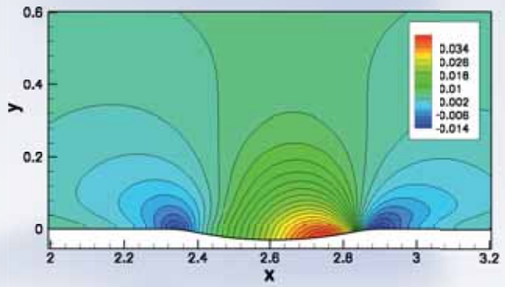

Figure 7: Pressure distribution in the dimple in a x-y midplane

Figures 6 and 7 display the time-averaged flow in a x-y midplane of one dimple. Based on the velocity vectors the reduction of the velocity gradient near the wall and thus of the wall-shear stress is visible. On the other hand, it is obvious that the dimples lead to a modified pressure distribution compared with the case of a smooth wall. At both upstream and downstream regions of the dimple, where the flow either enters or leaves the surface depression, the pressure is slightly decreased. However, more important is the observation that the pressure increases on the rising edge of the dimple yielding a contribution to the overall drag. The strongest flow structures in the time-averaged flow, visualized by iso-surfaces of λ_2 used as a structure-identification method are located at the upstream and downstream edges of the dimple, where the flow enters and leaves the dimple.

In order to evaluate the effect on the drag resistance, histories of the forces acting on the lower dimpled wall of case B and on the non-dimpled smooth wall of case R were investigated. For the non-dimpled wall the situation is simple, since on this surface only the wall-shear stress leads to a skin-friction drag. This force also acts on the lower dimpled wall, where the time-average is marginally smaller on the dimpled wall than on the smooth wall. However, at the dimpled wall the pressure distribution on the wavy surface (Fig. 7) yields a pressure force in the main flow direction (about 5.5% of total force) which contributes to the total drag force. Both effects – the slight decrease of the average shear force and the additional contribution of the pressure force – approximately compensate each other, so that there is no net drag reduction due to the dimples. In fact, the total drag for case B even increases about 4.4% compared to the fine-grid reference case R. For the coarse grid, we found an increase of about 2.0% and 3.8%, respectively for the case L (one dimpled wall) and case B (two dimpled walls). Thus the effect of the dimples is approximately doubled for case B compared with case L and an increase of about 4% is found for case B on both grids. Hence, these simulations confirm the experimental findings [1] that the present arrangement of dimples does not lead to drag reduction but rather to a slight increase of the total drag.

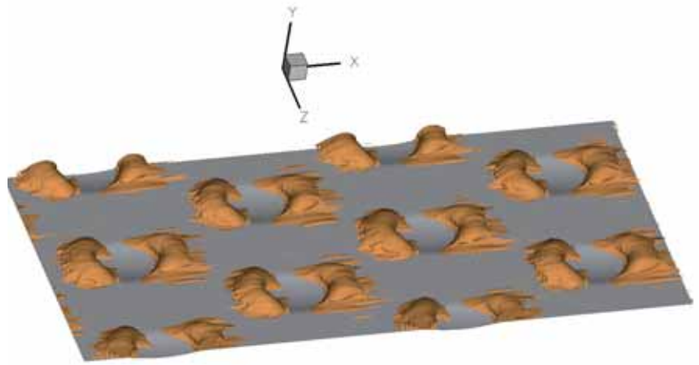

Figure 8: Iso-surfaces of $\lambda_2 = -0.05$

Conclusions

Based on the simulations (and the measurements [1]) described above, the question whether the present arrangement of dimples leads to drag reduction has to be answered in the negative. That is supported by both experimental measurements of an internal and an external flow as well as DNS of channel flows with dimpled and smooth walls. However, a significant increase of pressure loss was also not observed. Integrated over the entire surface the contribution from the shear stresses is marginally decreased by the shallow dimples but the effect on the drag is overcome by the newly appearing pressure force contribution. Concerning prospective applications of dimpled surfaces in engineering, this study implies that it is feasible to achieve an augmentation of heat transfer by shallow dimples without significant additional pressure losses typically encountered with deep dimples.

Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft under contract number BR 1847/9. The computations were carried out on the NEC SX-8 machine at HLRS Stuttgart (grant no.: DIMPLE/12760). All kinds of support are gratefully acknowledged.

References

[1] **Lienhart, H., Breuer, M., Köksoy, C.** Drag Reduction by Dimples? A Complementary Experimental/Numerical Investigation, Int. J. of Heat and Fluid Flow, in press, 2008

[2] **Breuer, M.** Large-Eddy Simulation of the Sub-Critical Flow Past a Circular Cylinder: Numerical and Modeling Aspects, Int. J. for Num. Methods in Fluids, vol. 28, pp. 1281-1302, 1998

[3] **Breuer, M.** Direkte Numerische Simulation und Large-Eddy Simulation turbulenter Strömungen auf Hochleistungsrechnern, Habilitationsschrift, University of Erlangen-Nürnberg, Berichte aus der Strömungstechnik, ISBN: 3-8265-9958-6, 2002

[4] **Breuer, M., Lammers, P., Zeiser, Th., Hager, G., Wellein, G.** Direct Numerical Simulation of Turbulent Flow Over Dimples - Code Optimization for NEC SX-8 plus Flow Results, High Performance Computing in Science and Engineering 2007, Transaction of the HLRS 2007, Oct. 4-5, 2007, HLRS Stuttgart, Germany, pp. 303-318, Springer, ISBN 978-3-540-74738-3, 2008

• Michael Breuer

University of Erlangen-Nürnberg, Institute of Fluid Mechanics

Turbulent Convection in large-aspect-ratio Cells

Convective turbulence is created in a fluid which is heated strongly from below. The resulting complex interplay between buoyancy and gravity is present in many systems, reaching from astro- and geophysical flows to chemical engineering and indoor ventilation (see e.g., [1]). Frequently the lateral extension of convective turbulence exceed the vertical one by orders of magnitude. For example, atmospheric mesoscale layers and the Gulf stream evolve on characteristic horizontal lengths of 500 to 5,000 km in combination with vertical heights of about 1 to 5 km. The resulting aspect ratios are then 100 to 1,000.

The massively parallel supercomputing facilities which are available today allow us to study how the flow structures and the transport properties in convective turbulence change when the convection cell is successively extended in both horizontal directions, while keeping the vertical one fixed. Theoretically, this opens the possibility for the formation of pancake-like large-scale circulation patterns or bigger clusters of hot and cold temperature blobs. But is convective turbulence seizing this opportunity? Indeed, a very recent work by Hardenberg et al. [2] indicates that such structures can be formed. However, the lateral Grid resolution and the degree of convective turbulence were still moderate in this simulation. Systematic numerical studies on the geometry dependence of turbulence are rare, since they become extremely numerically expensive. In particular, the large lateral Grid extensions for the flat cells require computational resources which have become available with the advent of the Blue Gene systems.

Another important open question in this context is the dependence of transport properties on the aspect ratio. We can expect that the turbulent dispersion of tracer particles, such as aerosols in the atmosphere or phytoplankton in the upper ocean, will be different for the lateral and vertical directions. Recall that convective turbulence is inhomogeneous and anisotropic. Planets and stars rotate at different angular frequencies, i.e. in addition to the buoyancy and gravity effects Coriolis forces appear. The effect of rotation on the structure formation and the tracer transport is a further question which will be investigated in our massively parallel high-resolution simulations of turbulent convection.

Parallel Computation

Our direct numerical simulations solve the Boussinesq-Navier-Stokes equations by a pseudospectral method. At the core of this numerical scheme is the Fast Fourier Transform (FFT), which allows for the calculation of different terms of the nonlinear partial differential equations, either in physical or in wavenumber spaces, and also for rapid switching between the spaces. The classical parallel implementation of three-dimensional FFTs uses a slab-wise decomposition of the simulation domain. For a simulation with N^3 Grid points, the method allows a parallelization on up to N processors. The rather small memory size per core on Blue Gene and the large difference in the number of vertical and horizontal Grid points require so-called volumetric FFTs which decompose the three-dimensional volume into iproc × jproc pencils and allow for a parallelization degree of N^2. The prime requirement for a simulation with a large Grid is that the FFT algorithm should also be scalable, i.e. increasing the number of CPUs to solve the problem should also substantially decrease the time-to-answer [3]. Figure 1 shows the scaling properties of our simulation code for an increase in the number of CPUs by more than two orders of magnitude (see inset). The main figure illustrates also that different decompositions iproc × jproc of the total CPU number alter the performance significantly.

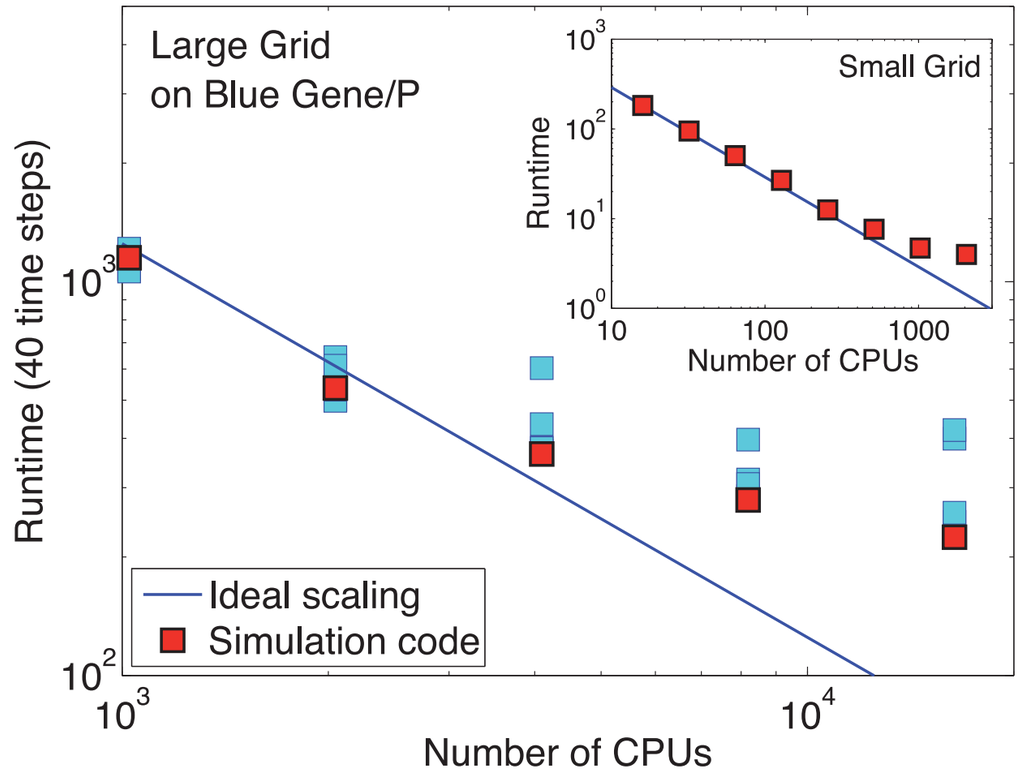

Figure 1: Strong scaling tests of the production code. The cyan boxes mark different decompositions iproc × jproc for a given total number of CPUs. Inset: The number of CPUs has been varied from 16 to 2,048 on a smaller Grid.

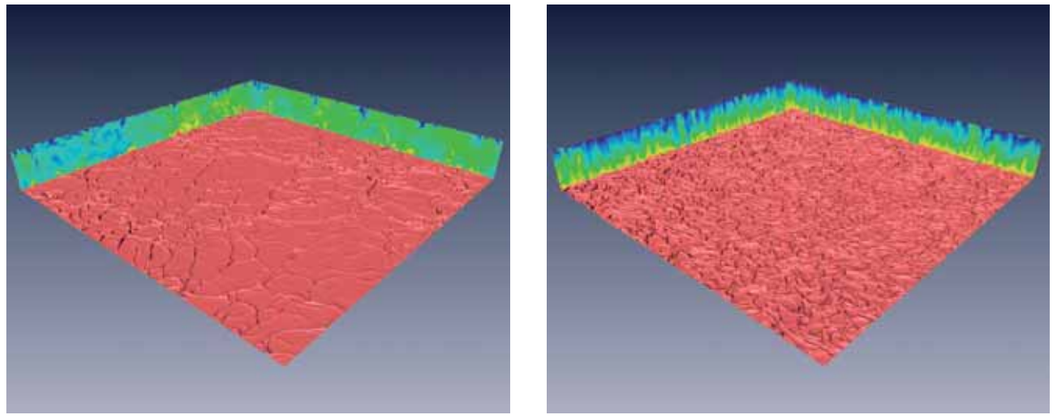

Figure 2: Snapshot of temperature field. An isosurface close to the heating plate and contour plots at two side planes are shown. Left: No rotation. Right: Strong rotation about the vertical axis.

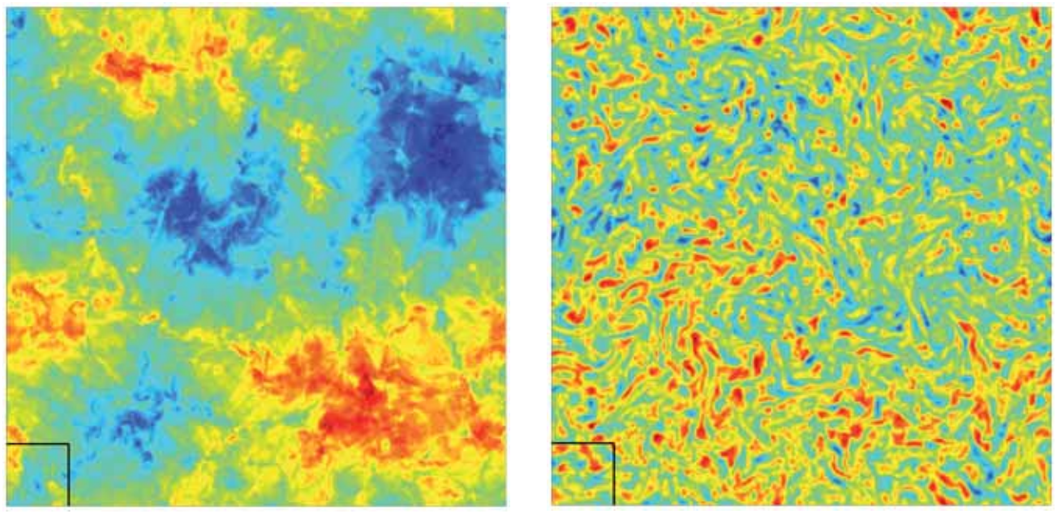

Figure 3: Contours of the vertically averaged temperature field. The black box is the aspect ratio of one cell. The vertical direction points out of the plane. Left: No rotation. Right: Strong rotation.

Results and Outlook

Figure 2 shows two instantaneous snapshots of the full convection cell at an aspect ratio 8. In both cases, the computational Grid contains 2,048×2,048×257 points. The production jobs for these parameter sets were run on 1,024 CPUs. An isosurface of the temperature close to the heating plane is shown. In addition, contour plots of the temperature on two side walls are displayed. The left picture for the non-rotating case indicates the formation of large-scale patches which are visible as ridges of the isosurface. They are known as thermal plumes which detach permanently from the thermal boundary layers at the top and bottom planes. These characteristic elements of convective turbulence can also be seen in the side planes. The situation changes drastically when strong rotation is present. The vertical transport of heat is arranged in columnar structures. Strong rotation prohibits thus the formation of large-scale lateral patterns in our convection cell.

Figure 3 (left) highlights large-scale temperature patterns which have been formed for the non-rotating case. Contours of the vertically averaged temperature at one time instant are shown (red=hot; blue=cold). The figure also underlines clearly that the structures would not be observable in a cell of aspect ratio one which is indicated by the black box. The strong rotation case (Figure 3 right) reflects again the vertical columnar arrangement of the turbulence as shown in Figure 2. An open question, that we are going to study in the future, is whether the large-scale filaments in the non-rotating case prevail when the driving of the turbulent motion by the vertical temperature difference becomes stronger.

As said before, an important aspect for our better understanding of convective turbulence is to unravel the local mechanisms of the transport of heat through the cell. The preferential choice is then to “go with the flow”. This so-called Lagrangian description provides exactly this perspective on the turbulent fluid motion. The velocity and temperature fields are monitored from a local frame of reference which is co-moving with a tracer particle along the streamline. Here, we advect simultaneously up to 1,5 million of such tracers with the flow equations. In convective turbulence, it turns out that this motion becomes qualitatively different in the lateral directions compared to the vertical one. The lateral dispersion is found to be similar to isotropic turbulence and reveals the famous Richardson dispersion law [4].

Acknowledgements

The author thanks Mathias Pütz (IBM Germany) for his help to implement the P3DFFT package by Dmitry Pekurovsky (SDSC) and to optimize the simulation code on the Blue Gene/L. He also thanks Sonja Habbinga and Marc-André Hermanns from the JSC. Financial support by the Deutsche Forschungsgemeinschaft is acknowledged.

References

[1] **Kadanoff, L. P.**
Turbulent heat flow: Structures and scaling. Physics Today 54, 34-39, 2001

[2] **Hardenberg, J. v., Parodi, A., Passoni, G., Provenzale, A., Spiegel, E. A.**
Large-scale patterns in Rayleigh-Bénard convection, Physics Letters A, in press, 2008

[3] **Schumacher, J., Pütz, M.**
Turbulence in laterally extended systems, Proceedings of ParCo 2007, 585-592, 2007

[4] **Schumacher, J.**
Lagrangian dispersion and heat transport in convective turbulence, Phys. Rev. Lett., in press, 2008

• Jörg Schumacher

Technische Universität Ilmenau, Department for Theoretical Fluid Mechanics

Applications

Direct Numerical Simulation of Flame / Acoustic Interactions

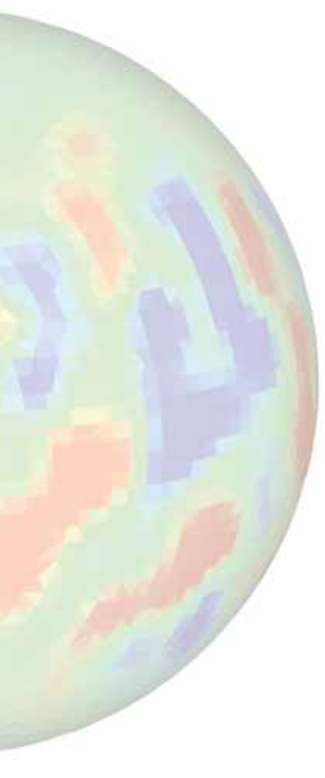

Applications involving turbulent flows with chemical reactions play an essential role in our daily life. All our transportation systems and the overwhelming majority of our energy supply rely directly or indirectly on the combustion of fossil fuels. In order to improve the efficiency of such systems and to develop new approaches and configurations, numerical simulations play an increasing role in research and industry. They provide a cost-efficient complement to experimental testing and prototypes. But, of course, they are only useful if they can accurately reproduce the real physical processes controlling practical configurations.

Depending on the value of a key non-dimensional number, the Reynolds number, every flow can be predicted to be laminar or turbulent. Turbulent flows are found in most practical applications but are much more complex than laminar flows, being chaotic in nature. The most exact numerical description of a turbulent flow field is achieved using the so-called direct numerical simulation (DNS) approach, for which the Navier-Stokes equations are solved as exactly as possible on an extremely fine grid: DNS results are often called “numerical experiments”. Nowadays, DNS is only possible for simple configurations and low-Reynolds number flows even on large computing clusters, because this type of simulation requires enormous computational resources. Furthermore, DNS is associated with complex post-processing and visualization issues, due to the extremely large quantity of raw data delivered by such computations. Nevertheless, DNS has become an essential and well-established research tool to investigate the structure of turbulent flames, since they do not rely on any approximate turbulence models [1]. In the present project our DNS code π^3C is employed to investigate different flames. This three-dimensional DNS code is a Finite-Difference solver written using FORTRAN 95 for a high numerical efficiency. It solves directly the compressible Navier-Stokes equations plus supplementary transport equations for total energy and for the mass fraction of each chemical component. Methane flames are considered in this study with a realistic description of all chemical processes, using 29 chemical species and 141 elementary chemical reactions. The modeling of the chemistry can be simplified when needed using a two-dimensional look-up table, built using the full mechanism [2,3]. In that case, species mass fractions and reaction variables are pre-computed and stored in a large database, used to speed-up the computational procedure.

Using a realistic description of chemistry on a growing number of grid elements rapidly leads to a huge discretized equation system and to an enormous computation time. In this case parallel computations are absolutely necessary. The simulation time can be highly reduced by dividing the numerical domain into smaller sub-domains (a method called Domain Decomposition). Each processor of the parallel supercomputer is then responsible for its own sub-domain and exchanges information with its topological neighbours. The inter-processor communication relies on the Message-Passing Interface (MPI) communication library.

Configuration

The ignition and initial development of a flame inside a turbulent flow is a problem of great interest, both from a fundamental (complex, multi-scale, fully coupled physical process) and from a practical (internal combustion engines, gas turbine re-ignition, security issues...) point of view. The flame is initially perfectly spherical, laminar and centreed in the middle of the three-dimensional numerical domain. Initial conditions correspond to ambient pressure and temperature. As a result of the solution procedure, a turbulent, fully premixed flame expands with time into the numerical domain.

Each side of the computational domain is 1,0 cm long. This small size will be increased in the future by employing even more computer nodes. For the presented results the grid spacing is constant and uniform, equal to 33 µm. This leads to a computational Grid with ca. 27 million Grid points. Such a fine grid-resolution is needed to describe correctly all turbulent structures (vortices) but also to resolve stiff intermediate (short-lived) chemical radicals, which exist only in very narrow reaction fronts.

Results

Various species can be used as an indicator of the flame front in a combustion process. Among them, one specific iso-surface of the main product (carbon dioxide CO_2) is retained here. A typical time-evolution for the expanding turbulent premixed flame is shown in Figure 1. The impact of turbulence on the development of the initially spherical flame is clearly observed in this figure, leading to a considerable wrinkling and deformation of the flame front. As a result, the final structure indeed corresponds to a realistic turbulent flame. The first step of a standard post-processing is usually to extract the instantaneous flame front from the raw data, as illustrated in Figure 1. Many important modeling quantities can then be extracted from the DNS along the flame front. A dedicated Matlab-based library has been developed in our group for this purpose [4]. This post-processing delivers for example the local strain-rate, as shown on top of the flame surface in Figure 1.

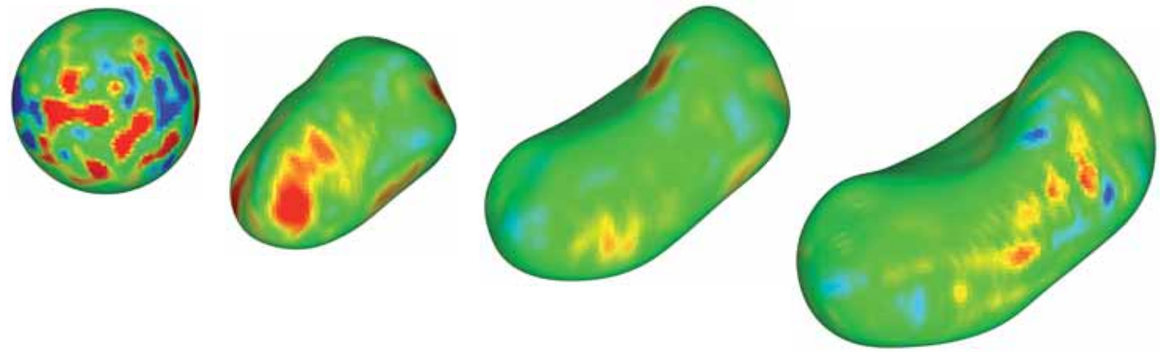

Figure 1: Successive positions of the flame front in time, corresponding respectively to t = 0, 0,3, 0,5, 0,7 ms (from left to right) after flame initialization. The flame surface has been colored by the tangential component of the strain rate.

Then, the unit vector locally perpendicular to the flame surface can be computed. The flame front curvature is given by the divergence of this vector. In Figure 2 (left) the local flame curvature is shown again on top of the instantaneous flame surface. As a complement, the local flame thickness has been also computed at the same time. Several definitions can be used to determine flame thickness. Here, the thermal flame thickness, computed from the maximal temperature gradient has been retained.

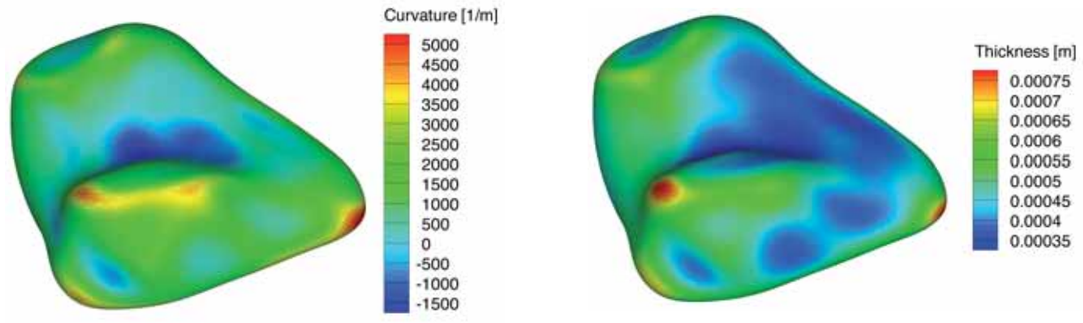

Figure 2: Instantaneous flame front at t=0,7 ms, colored with local mean flame curvature (left) and flame thickness (right). Note that the angle of view is not the same as in Figure 1.

Figure 2 (right) shows this local flame thickness on top of the instantaneous flame front. From Figure 2 a non-negligible positive correlation between curvature and thickness is already visible, which can easily be quantified in a further analysis. Such correlations should clearly be taken into account in simplified turbulent combustion models, such as used for industrial configurations. This exemplifies how DNS results might be employed to investigate, check and eventually improve such models.

Summary

Using powerful parallel supercomputers, accurate physical models and efficient numerical techniques, complex three-dimensional turbulent flames can be computed as “numerical experiments” using Direct Numerical Simulation. Interesting information can be obtained in this way concerning, e.g., the modifications of the local flame structure induced by the turbulence, or concerning flame acoustics (Fig. 3). A further numerical optimization of the DNS code is now needed to employ efficiently an even higher number of nodes. In that manner larger numerical domains and more turbulent conditions (i.e., higher Reynolds numbers) will become accessible to such simulations.

Acknowledgement

The authors would like to thank the Leibniz Supercomputing Centre for providing access to its supercomputers (Project h1121). The financial support of the Deutsche Forschungsgemeinschaft (DFG) for the position of Dr. Hemdan Shalaby (Research Unit #486 “Combustion noise”) has been essential for the success of this project.

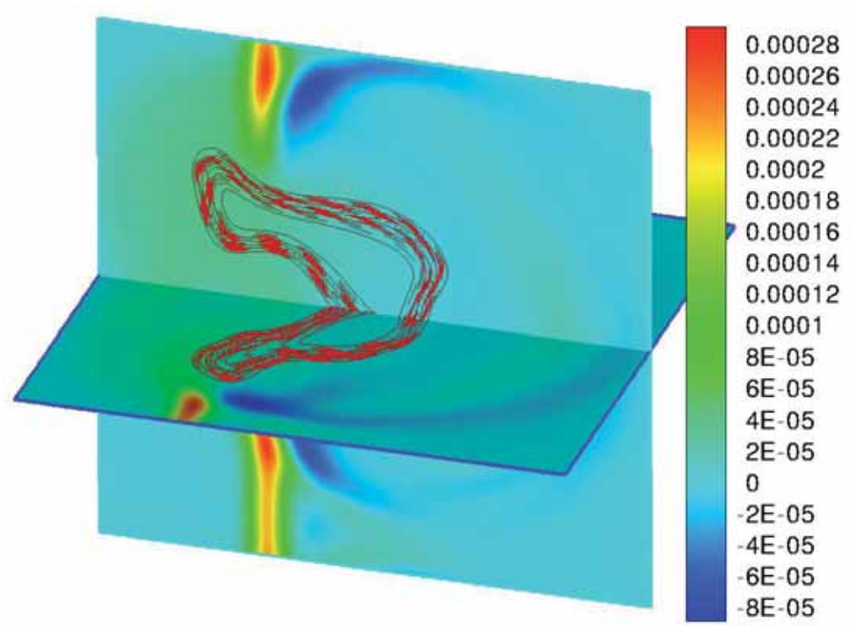

Figure 3: Instantaneous perturbation of the pressure field during the interaction of a turbulent flame with a Gaussian acoustic wave

References

[1] **Hilbert, R., Tap, F., El-Rabii, H., Thévenin, D.** Impact of detailed chemistry and transport models on turbulent combustion simulations. Prog. Energy Combust. Sci. 30, 61–117, 2004

[2] **Fiorina, B., Baron, R., Gicquel, O., Thévenin, D., Carpentier, S., Darabiha, N.** Modeling non-adiabatic partially premixed flames using flame-prolongation of ILDM. Combust. Theory Modeling 7, 449–470, 2003

[3] **Janiga, G., Gicquel, O., Thévenin, D.** High-resolution simulation of three-dimensional laminar burners using tabulated chemistry on parallel computers. In: Proc. 2nd ECCOMAS Conference on Computational Combustion, (Coelho, P., Boersma, B.J., Claramunt, K., Eds.), Delft, The Netherlands, 2007

[4] **Zistl, C., Hilbert, R., Janiga, G., Thévenin, D.** Increasing the efficiency of post-processing for turbulent reacting flows. Comput. Visual. Sci., in press, 2008

• Gábor Janiga
• Hemdan Shalaby
• Dominique Thévenin

University of Magdeburg
Otto von Guericke, Laboratory of Fluid Dynamics and Technical Flows

Massively parallel single and multi-phase Flow in porous Media

The water exchange between the atmosphere and the terrestrial system is strongly influenced by water flow and phase transition processes at the soil surface and within the underlying vadose zone. Although this zone is a relatively small compartment of the biosphere it hosts most terrestrial biological and life supporting processes. Water distribution and dynamics within this zone on the micro-scale (the scale where the porous medium is resolved) determine the mass fluxes and the transport of substances at the macro-scale (which is related to compartment). Thus a proper understanding of the processes on the microscale can help to improve the prediction of the behavior on the macro-scale and is extremely important for application fields like environmental engineering, petroleum industry, irrigation, CO_2 sequestration and many others.

Hysteresis

In multiphase flow systems of two immiscible phases as e.g. air and water in a porous medium like a soil, the flow and transport properties depend on the amount and the spatial distribution of the phases within the pore space. On the pore scale the phase distribution is controlled by the capillary forces depending on pore size, surface tension and wettability. For that reason, the relationship between capillary pressure and liquid saturation (P_c-S_w relationship) is of high importance for the prediction of water flow and solute transport. Unfortunately, this relationship is ambiguous and depends on the preceding wetting and drainage processes. This phenomenon is denoted as hysteresis and was first documented in [7]: It is caused by different pore structures relevant for drainage and wetting processes. While the drainage of a large pore body may be prevented by surrounding small pore throats, the wetting of fine pores above a large pore is hampered by the weak capillary forces in the wide body. Additional hysteresis effects are caused by a difference in advancing and receding contact angles and the inclusion of air in a first wetting process. Due to the inclusion of air the porous medium cannot be completely re-saturated after the first drainage.

Computational Approach

Lattice Boltzmann Methods

In the last two decades the Lattice Boltzmann Method (LBM) has matured as an alternative and efficient numerical scheme for the simulation of fluid flows and transport problems. Unlike conventional numerical schemes based on discretizations of macroscopic continuum equations, the LBM is based on microscopic models and mesoscopic kinetic equations. The fundamental idea of the LBM is to construct simplified kinetic models that incorporate the essential physics of microscopic or mesoscopic processes in a way that the macroscopic averaged properties obey the desired macroscopic equations. The scheme is particularly successful in problems where transport problems are described by an advective and diffusive operator. Especially applications involving interfacial dynamics, complex and/or moving boundaries and complicated constitutive relationships which can be derived from a microscopic picture are suitable for the LBM.

We developed a LBM optimized for incompressible multiphase Stokes flow problems. A detailed description of the model is given in [1,2]. The simulation kernel called “Virtual Fluids” [4] developed at the Institute for Computational Modeling uses block structured grids and the parallelization follows a distributed memory approach using the Message Passing Interface (MPI).

Tomography and Image Processing

To obtain the geometry of a real porous medium we scanned a sand sample of 1,5 cm in diameter and 1 cm in height using X-rays from a synchrotron source. The tomography was carried out at the Hamburger Synchrotron Laboratories (HASYLAB) in Germany. The particle size of the sand material ranged from 0,1 to 0,5 mm. Based on the imaged X-ray attenuation the density distribution of the solid material can be reconstructed. The reconstructed density map was segmented into a black and white image of pore space and solid phase with a voxel resolution of 11 µm. Figure 1 shows the reconstructed porous medium. Details can be found in [6].

Grid Generation

Direct computations of flows through porous media in the literature so far are based on binarized porous media data mapped to uniform Cartesian grids. The tomographic voxel set is directly used as the computational grid and therefore the geometrical representation is usually only first-order accurate due to staircase patterns. We pursue a more elaborate approach, where the geometry is obtained as follows: Starting from a highly resolved tomographic data set we use an isocontour algorithm (marching-cubes) to reconstruct the surface of the porous medium as a set of planar triangles. Then the numerical resolution of the Cartesian grid for the simulation can be chosen independently from the voxel set. Details can be found in [5]. Figure 1 shows a detail of the triangulated surface for the porous medium. The complete geometry consists of 110 million triangles.

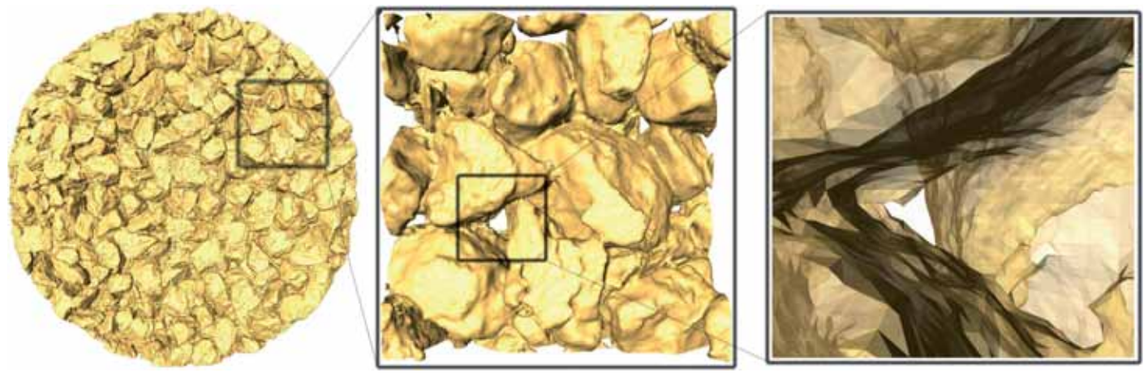

Figure 1: Left: Geometry obtained by an X-ray scan of a porous medium measured at HASYLAB and image processing methods. Right: Detail of the reconstructed surface of the porous medium.

Simulation

Only a few studies have reported simulations of multiphase flow in three-dimensional porous media mainly because of computational limitations. Only with the computational resources delivered by the SGI Altix 4700 at the LRZ in Munich it was possible at all to perform meaningful simulations for such a complex problem. To compute the hysteretic P_c-S_w relationship the numerical simulation is conducted as follows. Initially the entire pore space is filled with the wetting phase (water). A time dependent suction is applied at the bottom, while the top of the sample is connected to a non-wetting phase reservoir (air). The numerical grid resolution for the different simulations was in the range of 200^3 to 800^3. In Figure 2 a snapshot of a hysteresis simulation is given. Air infiltrates at the top of the porous medium which is filled with (non-visible) water. Many other simulations have been performed and the results have been compared to other modeling approaches (e.g. a morphological pore network model) and to experimental data and show a very good agreement [2].

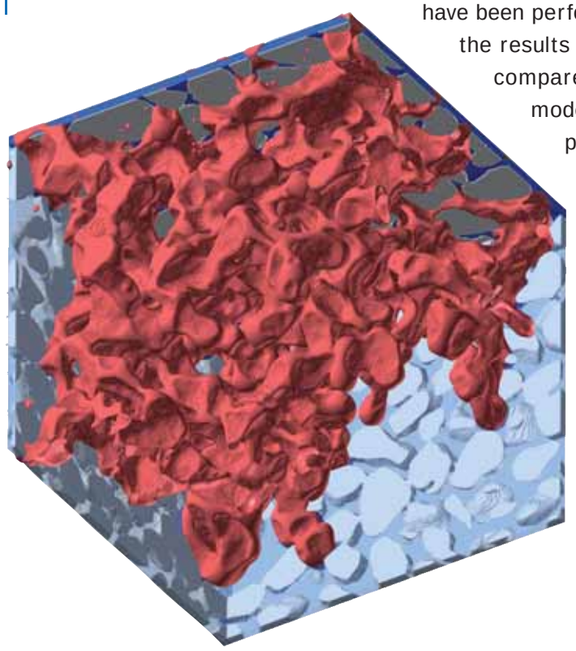

Figure 2: Simulation of air (red) penetration from the top in an initially water (invisible) saturated porous medium

Computational and parallel Performance

The LBM is an explicit time stepping scheme, which operates on regular grids and usually relies only on next-neighbor interaction. Therefore the method is very well suited for massively parallel systems. Also the single core performance is well exploited with LBM. The performance P of an LB implementation is usually measured in Lattice Updates per Seconds (LUPS), which states how many lattice nodes can be updated in one second. For a single core we obtain P = 3,33E6 LUPS (in single precision) which corresponds to 37% of the maximum possible performance.

In Figure 3 the parallel efficiency for a single-phase simulation and different grid sizes in a simple regular geometry with high porosity (0.9) is shown. Even for a relatively small grid size of 240^3 we obtain a very good scaling behavior. Up to a number of 1,000 cores we note a super-linear scaling due to caching effects. In general a very good scaling behavior of the code is observed. Using 4,080 cores for a grid of $4{,}000^3$ we obtain a parallel efficiency of 92%. The performance of a parallel multiphase simulation in a real porous medium depends strongly on the local porosity of the sub-domains. The smaller the variation of the local porosities in the sub-domains the better the load-balancing. For a typical simulation we used 512 cores and a partition of 8×8×8. The minimum and maximum porosity for one partition was 0.25 and 0.53. Nevertheless the parallel efficiency for this problem (70%) was still reasonable.

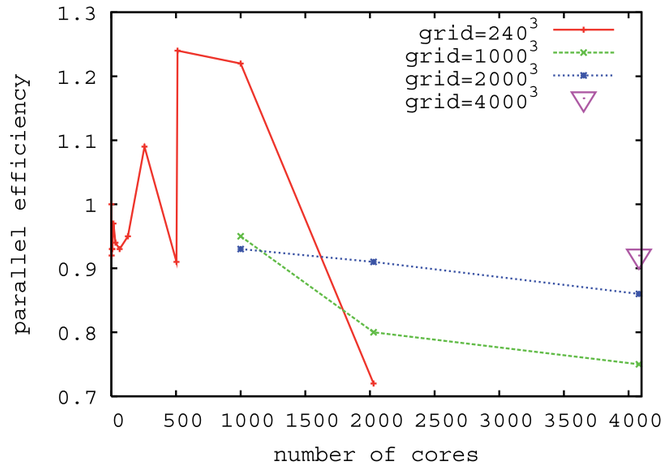

Figure 3: Parallel efficiency of “Virtual Fluids” for a regular single-phase problem and different Grid sizes

Outlook

Presently we are working on algorithms to simulate free surface problems with mesh refinement. Here we follow an adaptive block-wise strategy, where cuboids containing a fixed number of nodes are the leaves of a hierarchical octree data structure and are dynamically (de-)allocated if needed [3].

The LBM method is also very well suited for acoustics problems. The application field we are currently investigating is the determination of sound absorption properties of road surfaces. The geometry is again obtained by tomography methods. Since the number of time steps needed for an acoustics problem is much smaller than for a hydrodynamic simulation we can run simulations on very large grids and thus expect to be able to solve complex real life problems with sufficient accuracy.

References

[1] **Tölke, J., Freudiger, S., Krafczyk, M.** An adaptive scheme for LBE Multiphase Flow Simulations on hierarchical Grids, Computers & Fluids, Volume 35(8-9): 820-830, 2006

[2] **Ahrenholz, B., Tölke, J., Lehmann, P., Peters, A., Kaestner, A., Krafczyk, M., Durner, W.** Prediction of capillary hysteresis in porous material using lattice Boltzmann methods and comparision to experimental data and a morphological pore network model, acc. for publ. in Advances in Water Resources

[3] **Freudiger, S., Hegewald, J., Krafczyk, M.** A parallelization concept for a multi-physics lattice Boltzmann prototype based on hierarchical Grids, Progress in Computational Fluid Dynamics, in press

[4] **Geller, S., Krafczyk, M., Tölke, J., Turek, S., Hron, J.** Benchmark computations based on lattice Boltzmann, Finite Element and Finite Volume Methods for laminar flows, Computers & Fluids, Volume 35(8-9): 888-897, 2006

[5] **Ahrenholz, B., Tölke, J., Krafczyk, M.** Second-order accurate lattice Boltzmann flow simulations in reconstructed porous media, International Journal of Computational Fluid Dynamics, 20 (6): 369-377, 2006

[6] **Kaestner, A., Lehmann, E., Stampanoni, M.** Imaging and image processing in porous media research, acc. for publ. in Advances in Water Resources

[7] **Haines, W.** Studies in the physical properties of soils: V. The hysteresis effect in capillary properties, and the modes of moisture distribution associated therewith, Journal of Agricultural Science, 20: 97-116, 1930

- J. Tölke[1]
- B. Ahrenholz[1]
- M. Krafczyk[1]
- P. Lehmann[2]

[1] Technische Universität Braunschweig, Institute for Computational Modeling in Civil Engineering

[2] Ecole Polytechnique Fédérale de Lausanne, Laboratory of Soil and Environmental Physics

Unravelling the Mechanism How Stars Explode

Applications

Understanding how massive stars end their lives in supernova explosions is one of the most important problems in stellar astrophysics. With the most sophisticated computer simulations carried out to date, researchers at the Max-Planck-Institute for Astrophysics in Garching are making progress in deciphering the complex interplay of hydrodynamics and particle physics that reverses the collapse of the stellar core and initiates the violent disruption of the star.

Astrophysical Context

Roughly five times a second somewhere in the universe the life of a star with more than eight times the mass of our sun is terminated in a gigantic supernova explosion. These cosmic catastrophies are the most violent events after the Big Bang. Within only fractions of a second more energy is released than the sun produces during all its evolution. The hot plasma of the distroyed star, expanding into interstellar space with a tenth of the speed of light, can outshine a whole galaxy for several weeks.

It is these cataclysmic blasts, which we owe our existence to. They enrich the galaxy with carbon, oxygen, silicon, and iron, the building blocks of the earth and of all the creatures on it. Assembled during millions of years of nuclear burning as the dying star has been aging, or forged in the inferno of its final explosion, these chemical elements are disseminated into the galactic medium when the star is disrupted in the supernova event.

Observations of supernova remnants tell us that not all the stellar gas is ejected in the outburst. The core of the star with the size of the moon but more massive than the sun, collapses to an ultracompact object, a neutron star, which has only the diameter of a big city. Such a neutron star is still visible as a point source in many of the gaseous clouds that are left behind as heralds of past explosions.

The Supernova Puzzle

The gravitational binding energy released in the neutron star formation is hundred times more than needed for powering the supernova. But how does this happen? How is the implosion of the stellar core reversed to the explosion of the overlying layers of the star? This is still a puzzle that challenges theorists' intuition and modeling abilities. One of the most popular ideas involves neutrinos as the driving agents. These elementary particles are produced in huge numbers at the extreme conditions in the newly formed neutron star, where the matter is denser than in atomic nuclei and reaches temperatures of several hundred billion degrees. Neutrinos are the leak by which the collapsing stellar core loses its gravitational binding energy. But some fraction of the emitted neutrinos is still able to deposit its energy in the matter surrounding the compact neutron star (see Fig. 1). This energy transfer could be enough powerful to accelerate the supernova shock front and to expel the overlying shells of the star. The question whether this happens or not is a central problem in supernova theory.

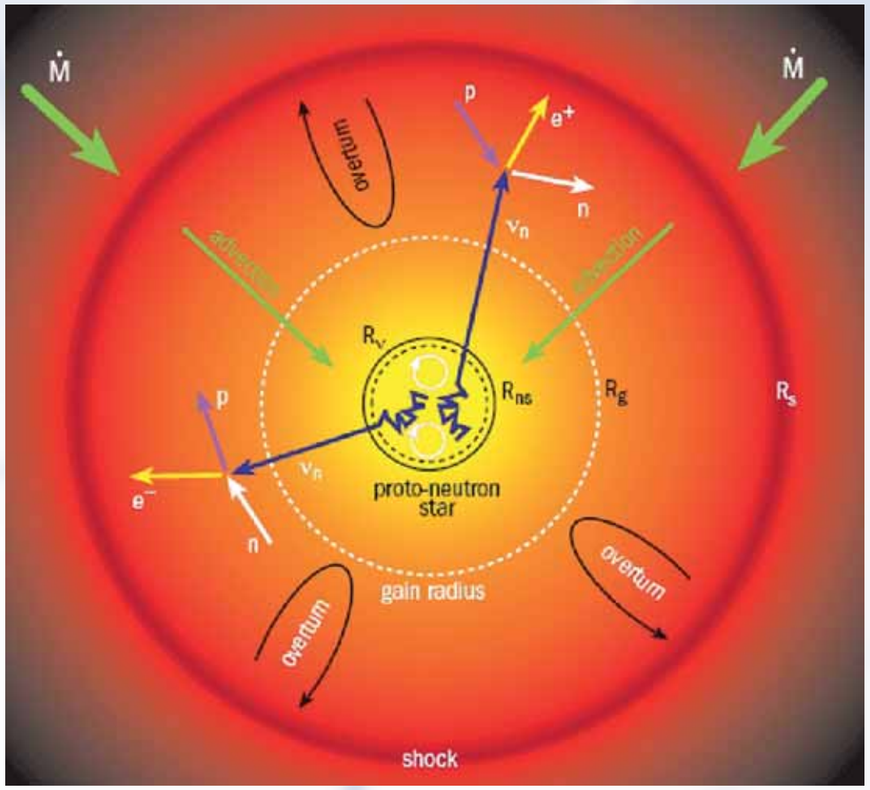

Figure 1: Neutrinos, radiated from the newly formed neutron star at the centre of a collapsing star, deposit the energy to drive the shockwave that causes the supernova explosion of the star.

Goals of this Project

Answering this question would mean a major breakthrough in stellar astrophysics. It would not only allow us to better link the properties of supernova explosions and their remnants to the different types of progenitor stars. It would also bring us closer to an answer of the question whether supernovae are the still mysterious source of rare elements like gold, lead, thorium, and uranium. And it would allow us to more reliably predict the neutrino and gravitational wave signals, which are planned to be measured for future galactic supernovae by a new generation of big experiments, and which are the only ways to observationally probe the processes deep inside the core of a dying star.

Until then, computer models are an indispensable tool for promoting our theoretical insight. At Garching, we have developed numerical codes that allow us to perform, with unprecidented accuracy, simulations of the complex particle physics, nuclear physics, and plasma dynamics that determine the destiny of collapsing and exploding stars.

Computational Challenges

The modeling of supernova explosions is in fact one of the most difficult problems in computational astrophysics. Largely different time scales, varying between microseconds and seconds, and length scales that extend from tens of meters to millions of kilometers, have to be resolved to follow

neutrino interactions, nuclear reactions, turbulent convection, and sound wave propagation in different regions of the collapsing core and of the ejected outer layers of exploding stars. This is computationally extremely demanding: half a second of evolution requires 500,000 time steps and in two spatial dimensions with 500-1,000 radial zones and typically 196 lateral bins needs some 10^{18} floating point operations.

Besides integrating the Euler equations that describe the time-dependent motion of the stellar fluid, one needs to solve the transport of neutrinos in the dense stellar matter. Different from the stellar gas, where very fast electromagnetic and strong interactions ensure that equilibrium is established on dynamical timescales, neutrinos couple with matter only through weak reactions. Thus they require a transport treatment in phase space by the Boltzmann equation and its moment equations. This constitutes a high-dimensional, time-dependent problem – in spherically symmetric models the transport is three-dimensional, for axisymmetric models five-dimensional, and in full generality six-dimensional – and poses the major computational challenge: the neutrino transport consumes the by far dominant amount of CPU time during supernova simulations.

The problem is also hard to be implemented efficiently on massively parallel computers. In particular the neutrino transport module has resisted such efforts so far and currently prevents us from using distributed memory architectures. Because the interaction timescale of neutrinos in neutron star matter is extremely short and the neutrino propagation happens with the speed of light after decoupling, the nonlinear transport equations of neutrinos, which as fermions are subject to phase-space blocking effects, need to be solved with fully implicit time stepping. In our current numerical implementation this leads to big, densly filled matrices that have to be inverted several times on every time level of the calculated evolution. This is computationally very expensive and defies easy parallelization. New algorithms, based on iterative multigrid solvers for hyperbolic systems of equations, are currently under construction but not yet available for full-scale supernova calculations. Their use, however, will be unavoidable in future three-dimensional models of supernovae.

The special needs of our current two-dimensional simulations, i.e., shared-memory nodes (up to 256 CPUs) with powerful processors and the continuous availability of these nodes for many months, are satisfied on different supercomputing platforms in all three national supercomputing centres, each of which has required special code adaptation and optimization: on the NEC SX-8 of the Höchstleistungsrechenzentrum Stuttgart (HLRS), on the sgi Altix 4700 of the Leibniz-Rechenzentrum (LRZ) in Munich, and on the IBM p690 Jump of the John von Neumann Institute for Computing (NIC) in Jülich. Our project also receives support by CPU time within the AstroGrid-D as part of the D-Grid initiative that is funded by the German Federal Ministry of Education and Research (BMBF).

Recent Progress

With the most sophisticated simulations carried out to date, we have learned that the neutrino energy deposition around the newly formed neutron star is supported by different fluid instabilities that take place in the gas flow that continuously adds more matter from the collapsing stellar core to the compact remnant at the centre. The neutrino-heated gas is stirred up by vigorous convective overturn as hot matter becomes buoyant and begins to rise while cooler fluid sinks inward and is partly absorbed into the neutron star (Fig. 2).

In addition to this phenomenon, which has been known already for ten years, we have discovered a global non-radial instability of the gas flow towards the centre. The layer between neutron star and supernova shock front is unstable by a so-called advective-acoustic cycle, which constitutes an amplifying feedback cycle of inward advected vorticity perturbations and outward propagating sound waves. This instability can grow even in conditions where convection remains weak and it can instigate violent secondary convection (Fig. 3). It thus improves the conditions for ongoing strong neutrino energy transfer to the supernova shock. In fact, our simulations show that for stars from about eight solar masses to at least 15 solar masses neutrino energy deposition, supported to different extent by both hydrodynamic instabilities, can initiate and drive the supernova explosion (Figs. 2 and 3).

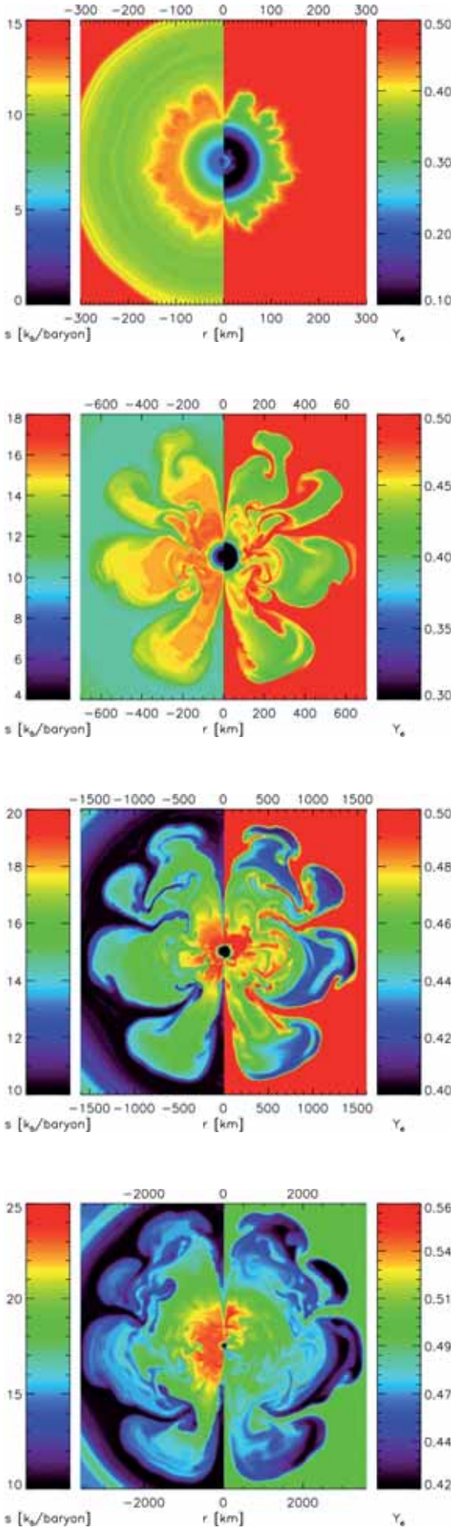

Figure 2:
Gas entropy (left half) and electron-to-nucleon ratio (right half) during the early stages (0,097, 0,144, 0,185, and 0,262 seconds) of the explosion of a star with nine solar masses. Convection causes anisotropies of the ejected gas.

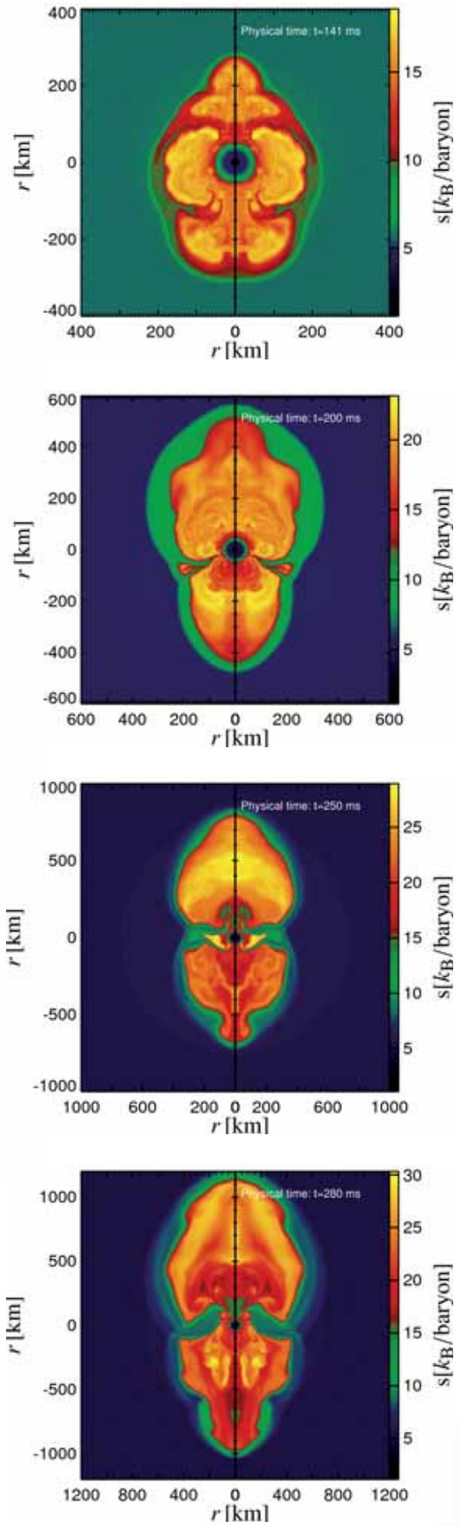

Figure 3:
Gas entropy for the beginning explosion of a star with eleven solar masses. The explosion becomes globally asymmetric because of the generic nonradial instability of the developing supernova shock.

The onset of the explosion thus turns out to be a generically multi-dimensional phenomenon. The highly aspherical initiation of the blast (Fig. 3), even in the absence of rapid rotation, suggests explanations for a variety of observations. The fast motions of many young neutron stars, which are measured to have average velocities of several hundred kilometers per second, some of them even of more than 1,000 km/s, can be explained by the recoil imparted to the compact remnant by the anisotropically ejected supernova gas. The asymmetry of the developing explosion also triggers large-scale mixing instabilities in the outer layers of the disrupted star, thus accounting for the clumpiness seen in many supernovae and the asymmetric appearance of gaseous supernova remnants.

Outlook

Despite the significant progress of our fundamental understanding of the processes that conspire in starting the explosion, many more simulations and in particular long-time simulations are needed to establish the properties of self-consistently calculated explosions. The different structures of stars with different masses require studies of a wider range of supernova progenitors. And the incomplete knowledge of the initial conditions (e.g., of the angular momentum in the stellar core) and of various aspects of the microphysics (e.g., of the equation of state of hot neutron star matter) make it necessary to explore the full range of variability. Ultimately, three-dimensional simulations will have to be performed to confirm our findings of the present two-dimensional models. But for that to be possible, we still need to wait for the next generations of more powerful supercomputers, and we have to advance our modeling tools to massively parallel application.

References

[1] **Buras, R., Janka, H.-Th., Rampp, M., Kifonidis, K.**
Two-dimensional hydrodynamic core-collapse supernova simulations with spectral neutrino transport. Models for different progenitor stars, Astronomy & Astrophysics, 457, 281-308, 2006

[2] **Janka, H.-Th., Marek, A., Müller, B., Scheck, L.**
Supernova explosions and the birth of neutron stars, in Proceedings of International Conference 40 Years of Pulsars: Millisecond Pulsars, Magnetars, and More, McGill University, Montreal, Canada, August 12-17, 2007, Eds. Wang, Z., Bassa, C., Cumming, A., Kaspi, V., AIP Conference Proceedings, Vol. 983 (American Institute of Physics, New York), 369-378, 2008; arXiv:0712.3070

[3] **Kitaura, F.S., Janka, H.-Th., Hillebrandt, W.**
Explosions of O-Ne-Mg cores, the Crab supernova, and subluminous Type II-P supernovae, Astronomy & Astrophysics, 450, 345-350, 2006

[3] **Marek, A., Janka, H.-Th.**
Delayed neutrino-driven supernova explosions aided by the standing accretion-shock instability, Astrophysical Journal, submitted; arXiv:0708.3372, 2007

[3] **Scheck, L., Kifonidis, K., Janka, H.-Th., Müller, E.**
Multidimensional supernova simulations with approximative neutrino transport. Neutron star kicks and the anisotropy of neutrino-driven explosions in two spatial dimensions, Astronomy & Astrophysics, 457, 963-986, 2006

• Hans-Thomas Janka
• Andreas Marek
• Bernhard Müller

Max-Planck-Institut für Astrophysik, Garching

The Aquarius Project: Cold Dark Matter under a Numerical Microscope

A major puzzle in Cosmology is that the main matter component in today's Universe appears to be a yet undiscovered elementary particle whose contribution to the cosmic density is more than 5 times that of ordinary baryonic matter. This Cold Dark Matter (CDM) interacts extremely weakly with regular atoms and photons, so that gravity alone has affected its distribution since very early times, when the Universe was in a nearly uniform state.

When the effects of the baryons can be neglected, the nonlinear growth of structure is a well-posed problem where both the initial conditions and the evolution equations are known. In fact, this is an N-body problem par excellence. The faithfulness of late-time predictions is limited purely by numerical technique and by the available computing resources. Over the past two decades, simulations have already been of tremendous importance for establishing the viability of the CDM paradigm. In particular, simulation predictions for the distribution of matter on large scales have been compared directly with a wide array of observations; so far the paradigm has passed with flying colors.

Given CDM's success in reproducing the main aspects of the large-scale structure of the Universe, it is important to test its predictions on smaller scales, both to test it further and to seek clues to the nature of dark matter. In the Aquarius Project carried out by the international Virgo Consortium on the HLRB II supercomputer at LRZ, we aim to do this by studying the highly nonlinear structure of CDM halos in unprecedented detail. We are especially interested in the innermost regions of these halos and in their substructures, where the density contrast exceeds 10^6 and the astrophysical consequences of the nature of dark matter may be most clearly apparent. Quantifying these consequences reliably through simulations is, however, an acute challenge to numerical technique.

The numerical Challenge

In the Aquarius Project, we have performed the first ever one-billion particle simulation of a Milky Way-sized dark matter halo, improving resolution by a factor of more than 15 relative to previously published simulations of this type. The achieved mass resolution of ~1,700 solar masses is nearly a million times better than that of the largest cosmological simulation of large-scale structure formation carried out to date (the "Millennium Simulation"). Our spatial resolution reaches 20 parsec, which implies a dynamic range of close to 10^7 per dimension within the simulated box-size of more than 400 million lightyears across. This huge dynamic range makes our simulation a microscope for the phase-space structure of dark matter and enables dramatic advances in our understanding of the structure and substructure of dark matter in our galaxy. However, formidable challenges had to be overcome to make this calculation possible. Gravitational timescales are inversely proportional to the square root of the density, so simulating a CDM halo means dealing with a system where different regions evolve on timescales which may differ by factors of thousands. Codes with spatially-dependent, adaptive timestepping are mandatory otherwise the most rapidly evolving regions, which usually include only a tiny fraction of the mass, force timesteps so short that the calculation grinds to a halt.

A second challenge stems from the highly clustered spatial distribution of matter which affects in particular the scalability of parallel algorithms. A CDM halo is a near-monolithic, highly concentrated structure with a well-defined centre and no obvious geometrical decomposition which can separate it into the large number of computationally equivalent domains required for optimal exploitation of the many processors available in high-performance parallel architectures. In addition, gravity couples the dynamics of matter throughout the halo and beyond, requiring efficient communication between all parts of the simulated region. In our calculation, the clustering is so extreme that about one third of all simulation particles collect in a region that encompasses less than a fraction of 10^{-8} of the simulated volume!

Calculation Method and Parallelization Techniques

To make the Aquarius Project possible on the HLRB II, we have developed a major new version of our simulation code, GADGET-3, in order to improve scalability and performance for this extremely tightly coupled problem. GADGET uses a hierarchical multipole expansion (organized in a "tree") to calculate gravitational forces. In this method, particles are hierarchically grouped, multipole moments are calculated for each node, and then the force on each particle is obtained by approximating the exact force with a sum over multipoles. A great strength of the tree algorithm is the near insensitivity of its performance to clustering of matter, its ability to adapt to arbitrary geometries of the particle distribution, and the absence of any intrinsic resolution limitation.

However, there are actually faster methods to obtain the gravitational fields on large scales. In particular, the particle-mesh (PM) approach based on Fourier techniques is probably the fastest approach to calculate the gravitational field on a homogeneous mesh. The obvious limitation of this method is that the force resolution cannot be better than the size of one mesh cell, and the latter cannot be made small enough to resolve all the scales of interest in cosmological simulations. In fact, in our application we would need a mesh with $(10{,}000.000)^3$ cells to deliver the desired resolution with a single PM mesh - storing such a mesh would require several million petabytes!

GADGET therefore uses a compromise between the two methods. The gravitational field on large scales is calculated with a Particle-Mesh (PM) algorithm, while the short-range forces are delivered by the tree, such that a very accurate and fast gravitational solver results. A central role in our parallel code is played by the domain decomposition. It has to split the problem into smaller parts without data duplication, in a way that ensures a good balance of the computational work induced for each processor. GADGET uses a space-filling self-similar fractal, a Peano-Hilbert curve, for this purpose, which is made to become finer in high-density regions.

The domains themselves are then generated by cutting the one-dimensional space-filling curve into N_{cpu} pieces that approximately induce the same computational work-load. The domains are of nearly arbitrary shape but always have a relatively small surface-to-volume ratio.

First Results

As part of our Aquarius Project on the HLRB II, we have carried out extensive numerical resolution tests where we systematically increased the particle number used to simulate the same galaxy. Our primary simulation series culminated in our largest production calculation, which we refer to as C02-2400. This simulation followed about 4,5 billion particles, from a time briefly after the Big Bang to the present epoch, over more than 13 billion years of cosmic evolution. About 1,3 billion particles end up in the virialized region of a single, Milky-Way-sized galaxy, opening up a qualitatively new regime for studying the non-linear phase-space structure of dark matter halos.

An impression of the dynamic range of the simulation is given in Figure 1. What is readily apparent is the fascinating richness of dark matter substructure revealed by the simulation, which resolves several hundred thousand gravitational bound clumps of matter that orbit within the galaxy's potential. However, this extraordinary dynamic range comes at a price. More than 100,000 timesteps on 1,024 cores of the HLRB II and about 3 TB of RAM were required to carry out the simulation. In sum, the total CPU-time needed for completion of the C02-2400 calculation was nearly 4 million hours. The output produced forms a rich dataset of 45 terabytes in size, and will provide an extremely valuable scientific resource for many years to come.

Figure 2 shows spherically averaged density profiles obtained for the different resolutions that we calculated for our "C02" Milky Way halo. The convergence is excellent over the entire range where convergence can be expected based on the numerical parameters of the simulations. For the first time, our simulation series probes directly into a regime where the local logarithmic slope of the density profile of the dark matter cusp becomes shallower than -1. The structure of the cusp is of fundamental importance for our understanding of the CDM model, but has remained a highly contentious issue up to now. Our results demonstrate convincingly that an asymptotic power law of fixed slope apparently does not exist. Instead, the profile continues to become gradually shallower at ever smaller radii.

Our simulations also provide the first accurate and numerically converged measurement of the density profile of dark matter substructures, and they deliver precise predictions for the abundance of dark matter substructures. Using the simulation, we can obtain accurate determinations of the dark matter annihilation signal expected from the Milky Way's halo, which becomes potentially measurable with the launch of the GLAST gamma-ray satellite later this year. The simulation will also help to improve our understanding of galaxy formation and, in particular, the role and physics of satellite galaxies in the Milky Way's halo.

Thanks to the powerful HLRB II computer, the Aquarius Project produced the best resolved calculation of the Milky Way's halo carried out worldwide, and the ongoing analysis delivers many new insights for theoretical astrophysics. However, computational astrophysics in this area has still many exciting challenges in store for the future. Ultimately we would like to repeat our calculation by not only including the dark matter, but also the ordinary baryons that make up the stars, at similar or even better resolution. Carrying out such ultra-highly resolved hydrodynamical cosmological simulations will require further progress in the scalability of our codes, and the use of yet larger numbers of compute cores.

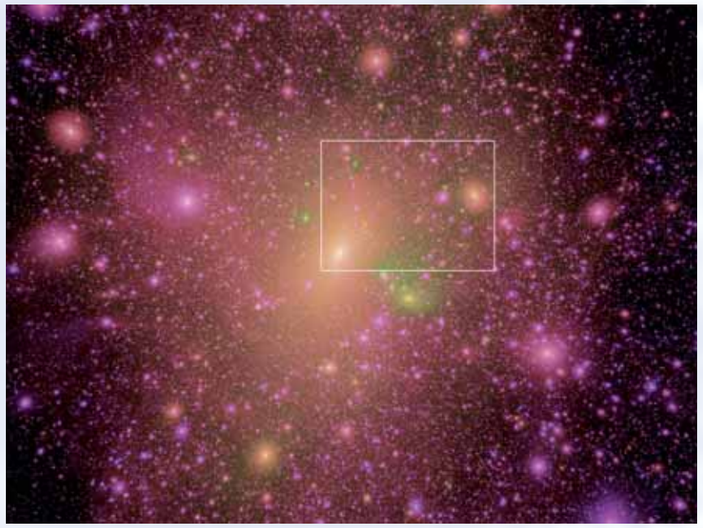

Figure 1: Dark matter distribution in the C02-2400 halo at the present time, showing the virialized region of the main halo and its immediate surroundings. The smaller picture enlarges the marked region by a factor of 4.

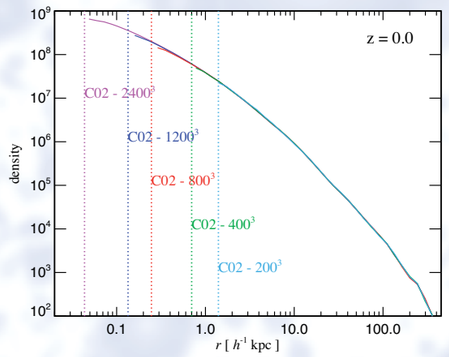

Figure 2: Spherically averaged density profiles of simulations of the same object carried out at different numerical resolution within the Aquarius Project. The dashed vertical lines mark the gravitational resolution limit of the individual calculations.

[1] Barnes, J., Hut, P.
A Hierarchical O(N logN) Force-Calculation Algorithm, 1986, Nature, 324, 446

[2] Davis, M., Efstathiou, G., Frenk, C. S., White, S. D. M.
The evolution of large-scale structure in a universe dominated by cold dark matter, 1985, Astrophysical Journal, 292, 371

[3] Diemand J., Kuhlen M., Madau P.
Dark Matter Substructure and Gamma-Ray Annihilation in the Milky Way Halo, Astrophysical Journal, 657, 262, 2007

[4] Navarro J. F., Frenk C. S., White S. D. M.
The Structure of Cold Dark Matter Halos, Astrophysical Journal, 462, 563, 1996

[5] Springel, V.
The cosmological simulation code GADGET-2, Monthly Notices of the Royal Astronomical Society, 364, 1105, 2005

[6] Springel, V., White S. D. M., Jenkins A., et al.
Simulations of the formation, evolution and clustering of galaxies and quasars, Nature, 435, 629, 2005

- Volker Springel[1]
- Simon D.M. White[1]
- Julio Navarro[2]
- Adrian Jenkins[3]
- Carlos S. Frenk[3]
- Amina Helmi[4]

[1] Max-Planck-Institute for Astrophysics, Garching

[2] University of Victoria

[3] University of Durham, Institute for Computational Cosmology

[4] University of Groningen, Kapteyn Astronomical Institute

Spring 2008 • Vol. 6 No. 1 • inSiDE 35

IRMOS – Interactive Real-time Multimedia Applications on Service-oriented Infrastructures

Traditionally, “real-time” refers to hard real-time systems, where even a single violation of the desired timing behaviour is not acceptable, for example because it leads to total failure, possibly causing loss of human lives. However, there is also a wide range of applications that also have stringent timing and performance needs, but for which some deviations in Quality of Service (QoS) are acceptable, provided these are well understood and carefully managed. These are soft real-time applications and include a broad class of interactive and collaborative tools and environments, including concurrent design and visualization in the engineering sector, media production in the creative industries, and multi-user virtual environments in education and gaming.

Soft real-time applications are traditionally developed without any real-time methodology or run-time support from the infrastructure on which they run. The result is that either expensive and dedicated hardware has to be purchased to ensure good interactivity levels and performance, or that general-purpose resources are used as a compromise (e.g. commodity operating systems and Internet networking) with no way to guarantee or control the behaviour of the application as a result.

IRMOS in General

IRMOS aims to break this mould by enabling “soft real-time” applications to be delivered through value chains that span organizational boundaries by a service oriented infrastructure that enables the real-time interaction of a distributed set of people and resources.

IRMOS is set apart from today’s Service Oriented Infrastructures through the following key features:

1) IRMOS will make it possible to distribute interactive real-time applications across organizational boundaries, instead of having to use dedicated, expensive and collocated hardware at a single site.

2) Businesses will be able to come together quickly and efficiently using IRMOS to identify, agree and deliver real-time applications without the need for protracted manual negotiations or service provisioning.

3) Providers will be able to deliver services that are cost effective and have guaranteed Quality of Service by using IRMOS to give them full control over their resources.

4) IRMOS will give all participants in inter-organization value chains including service providers and consumers the confidence that interactive real-time applications will be delivered in a predictable, reliable and efficient way.

5) IRMOS provides a comprehensive approach that addresses real-time at all levels (network, processing, storage, application, workflow and business) to allow complete end-to-end solutions to be built using a single infrastructure.

Each of these benefits will be demonstrated by using a range of applications including collaborative film post production, virtual and augmented reality application and interactive online learning using shared virtual environments. This broad range of demonstrators will enable the benefits of IRMOS to be promoted to the widest possible audience and achieve significant impact in the industrial, civil and educational sectors.

- Toolbox and best practices for real-time interactive applications
- Semantic descriptions and modelling of application characteristics
- Contribution to standardisation bodies

At the end IRMOS provides a framework that eases usage of services, either pure software-based or hardware-related, across organizational boundaries. For example it provides the possibility to SMEs to gain temporary access to specialized services on a rental basis

Objectives in IRMOS include:

- Facilitate real-time interactivity in SOIs
- Consolidate management and control of the infrastructure and services
- Enable integration between network and application services
- Engineer a platform of services

that are only affordable to maintain with high acquisition costs and specialized staff. In addition the IRMOS provides certain guaranteed environment properties, e.g. in terms of network link quality, which are negotiated on the basis of a service level agreement, which acts like a contract between service-consumer and service-provider.

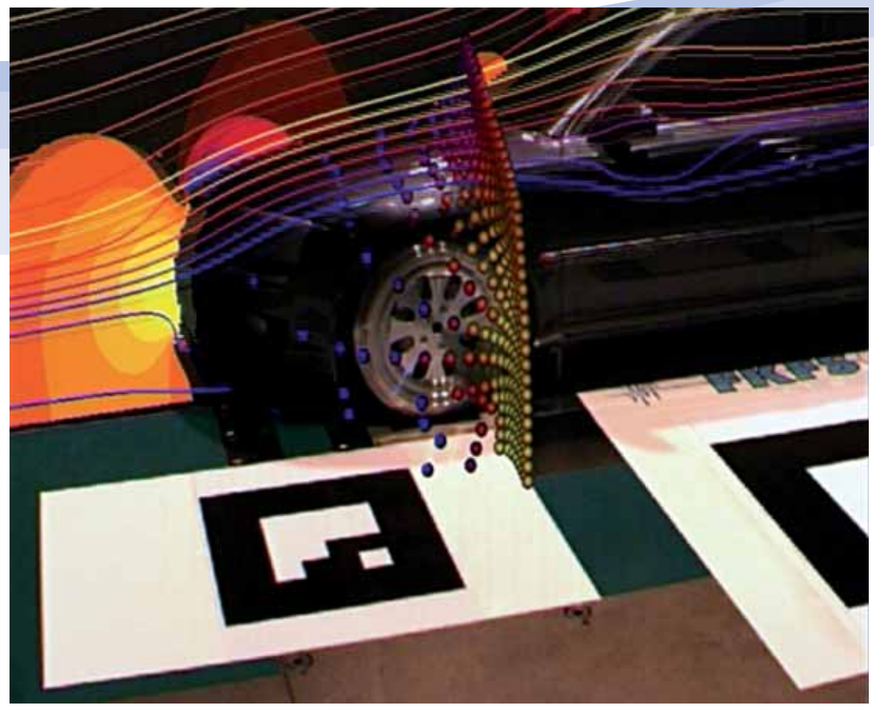

HLRS in IRMOS

The role of HLRS in the IRMOS project is to provide an application scenario for the service-oriented infrastructure, to show its advantages and benefits compared to today's existing solutions, as well as contributing to the development of the framework. The application scenario focuses on Virtual and Augmented Reality in collaborative working sessions. It comprises cross-organizational provision of Augmented Reality as a service to remote locations. Several partners join in a visualization session where one partner site has the capabilities to perform the Augmented Reality. This partner site however does not only provide Augmented Reality locally, but also to all collaboration partners over the internet. The application used in this scenario is called COVISE (which stands for Collaborative Visualization and Simulation environment). In IRMOS a service will be developed that performs video transmission of real-time video data which additionally is synchronized with position and orientation data of the Augmented Reality software. The video stream is then received by the collaboration partners and incorporated into their Virtual Reality environment and overlaid over simulation data based on the position and orientation that was received with the additional synchronized stream. This way partners in the collaboration can share the Augmented Display even though they don't have the hardware equipment at their site, but instead benefit of the so called Remote Augmented Reality where the Augmented Reality service is provided by a specifically equipped remote partner. The COVISE application needs certain adaptations for interfacing with the IRMOS framework to allow the required SLA-negotiation for IRMOS and the service development to be able to transparently utilise the hardware features provided by IRMOS.

Such a transmission of video data as well as metadata is done in real-time in a synchronized way. The constraint of real-time in this context raises requirements considering bandwidth, which has to be large enough for the video data of the appropriate resolution, as well as latency to ensure that all participants of such collaboration more or less see the visualization at the same time. These constraints are the issues to be solved by IRMOS through its framework services that lookup and determine appropriate resources for a gateway service as well as selection of an appropriate network link. These features requested through the application are, once setup, guaranteed for the requested time. With its features for SLA-based QoS, e.g. in terms of latency, bandwidth or even synchronization of multiple streams, as well as monitoring and feedback provision of the acquired link, IRMOS makes the scenario more stable and reliable.

The HLRS application scenario poses high requirements on the IRMOS framework in the area of guaranteed bandwidth, latency and processing power. For the IRMOS framework to be able to provide resources and network links that fulfil these constraints an interface is provided, that allows an application to specify its requirements on a high abstraction layer using a specification language. For that reason framework services are provided by the IRMOS framework, that offer webservices for requirements specification, SLA negotiation, service discovery, resource booking and service monitoring. With its experience in the field of web technologies HLRS provides a good knowledge in that field of technology and also contributes to the IRMOS project to successfully establish such a framework service.

Facts

IRMOS is a project funded by the European Union within the 7th Framework program in the 1st call of ICT. The consortium consists of eleven project partners with a total project budget of 12.9 M. It runs for 36 months and started on 1st of February 2008.

Partners

Xyratex (UK)
University of Stuttgart (DE)
ICCS / NTUA (GR)
Alcatel-Lucent (DE)
SINTEF (NO)
University of Southampton (UK)
Scuola Superiore Sant'Anna (IT)
Telefonica I+D (ES)
Giunti Labs (IT)
Grass Valley (DE)
Deutsche Thomson OHG (DE)

Project Coordinator

Eddie Townsed

Xyratex Technology Limited
Langstone Road
Havant
Hampshire
PO9 1SA
United Kingdom

Tel: +44 (0) 23 9249 6000
Fax: +44 (0) 23 9249 6001

Website: www.irmosproject.eu

• Stefan Wesner

University
of Stuttgart,
HLRS

IRMOS
Interactive Realtime Multimedia Applications on Service Oriented Infrastructures

Projects

DEISA to enhance the European HPC Infrastructure in FP7

In EU FP7, the DEISA Consortium continues to support and further develop the distributed high performance computing infrastructure and its services through the DEISA2 project funded for three years as of May 2008. Activities and services relevant for Applications Enabling, Operation, and Technologies are continued and further enhanced, as these are indispensable for the effective support of computational sciences in the HPC area. The service provisioning model will be extended from one that supports single projects to one supporting Virtual European Communities. Collaborative activities will be carried out with new European and other international initiatives.

Of strategic importance is the co-operation with the PRACE project which is preparing for the installation of a limited number of leadership-class Tier-0 supercomputers in Europe. The key role and aim will be to deliver a turnkey operational solution for a future persistent European HPC ecosystem, as suggested by ESFRI. The ecosystem will integrate national Tier-1 centres and the new Tier-0 centres, as illustrated in the following sketch.

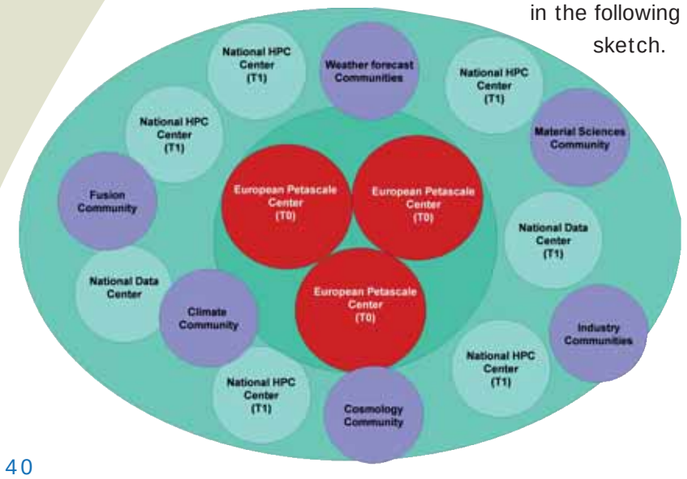

DEISA Background

In spring 2002 the idea emerged to overcome the fragmentation of supercomputing resources in Europe both in terms of system availability and in the necessary skills for efficient supercomputing support. The establishment of a distributed European supercomputing infrastructure was proposed. In May 2004 the DEISA project was started as a EU FP6 Integrated Infrastructure Initiative by eight leading European supercomputing centres. In 2006 DEISA was joined by three additional leading centres. Through the joint efforts, DEISA reached production quality soon after to support leading edge capability computing for the European scientific community. DEISA has also contributed to a raising awareness of the need for a persistent European HPC infrastructure as recommended in the ESFRI report 2006.

DEISA Extreme Computing Projects

The DEISA Extreme Computing Initiative (DECI), launched in 2005, continues to support the most challenging supercomputing projects in Europe which require the special resources and skills of DEISA. A European Call for Extreme Computing Proposals is published annually in spring. By selecting the most appropriate supercomputer architectures for each project, DEISA is opening up the currently most powerful HPC architectures available in Europe for the most challenging projects. This mitigates the rapid performance decay of a single national supercomputer within its short lifetime cycle of typically about 5 years, as implied by Moore's law. So far scientists from 15 different European countries with collaborators from four other continents have benefited.

DEISA2 Essentials

In the follow-up FP7 project DEISA2, the single-project oriented activities (DECI) will be qualitatively extended towards persistent support of Virtual Science Communities. DEISA2 will provide a computational platform for them, offering integration via distributed services and web applications, as well as managing data repositories. Emphasis will be put on collaborations with research infrastructure projects established by the ESFRI, and European HPC and Grid projects. The activity reinforces the relations to other European HPC centres, leading international HPC centres in Australia, China, Japan, Russia and the United States, and leading HPC projects worldwide. For supporting international science communities across existing political boundaries, DEISA2 participates in the evaluation and implementation of standards for interoperation.

Taking care of the operation of the infrastructure and the support of its efficient usage is the task of the service activities Operations, Technologies and Applications. Operations refers to operating the HPC infrastructure and advancing it to a turnkey solution for the future European HPC ecosystem. Technology covers monitoring of existing technologies in use and taking care of new emerging technologies with relevance for the infrastructure. Applications addresses the areas applications enabling for the DECI, Virtual Communities and EU projects, environment and user related application support, and benchmarking for the provision and maintenance of a European Benchmark Suite for supercomputers. In addition, Joint Research Activities aim at an integrated environment for scientific application development, and at the enabling of applications for the efficient exploitation of current and future supercomputers by an aggressive parallelism. Further reading: www.deisa.eu

DEISA Members

- BSC, Barcelona, Spain
- CINECA, Bologna, Italy
- CSC, Espoo, Finland
- ECMWF, Reading, UK
- EPCC, Edinburgh, UK
- FZJ, Jülich, Germany
- HLRS, Stuttgart, Germany
- IDRIS-CNRS, Orsay, France
- LRZ, Garching, Germany
- RZG, Garching, Germany
- SARA, Amsterdam, The Netherlands

Acknowledgment

The DEISA Consortium thanks the European Commission for support through contracts RI-508830, RI-031513 and RI-222919.

• Hermann Lederer
• Stefan Heinzel

Rechenzentrum Garching der Max-Planck-Gesellschaft

Projects

SmartLM

Grid-friendly Software Licensing for Location independent Application Execution

The currently existing licensing models for commercial applications are focussing on software used on compute resources within an administrative domain. Licenses are provided on the basis of named users, hostnames, or sometimes as a site license for the administrative domain of an organization. If we want to use this software in a distributed service oriented infrastructure, using resources that are spread across different administrative domains, that do not host the applications license server, we run into trouble. The licenses usually are bound to hardware within the domain of the user and do not allow access from outside thus enforcing local use of the protected applications only. Grid environments are usually spread across multiple organizations and their administrative domains, and virtualised infrastructures, like utility and cloud computing, hide the underlying hardware and their performance indicators, e.g. CPU type and frequency, that are often used in license agreements. While current mechanisms limit the usage of licensed software in Grid environments and virtualized infrastructures, the increasing usage of these environments and infrastructures make it necessary to overcome these limitations.

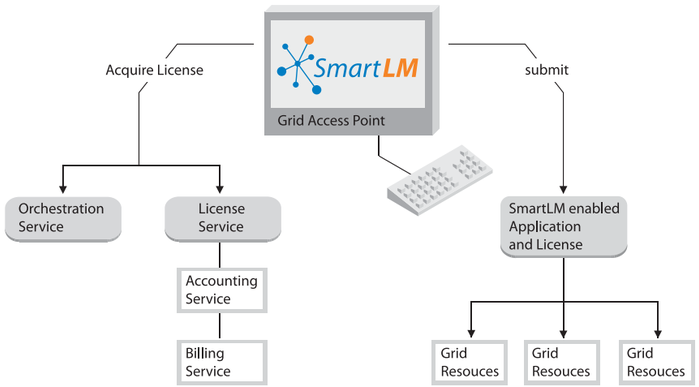

Solution

SmartLM will provide a generic and flexible licensing virtualization technology based on standards for new service-oriented business models. The solution will implement software licenses as Grid services, providing platform-independent access just like other Grid resources and being accessible from resources outside organizational boundaries. Service Level Agreements based on evolving standards will govern licenses. Secure agreements will be used to transport licenses through the Grid and to make them available on the resource to which a user has been granted access for the execution of the application. The license agreement and conditions of use for an application will be reached through negotiation between service providers and service customers. SmartLM will integrate the generic licensing virtualization technology into the major Grid middleware solutions UNICORE and Globus.

New service-oriented business models for this approach will be identified and a number of widely-used license-protected commercial applications will be adapted to be executed under control of the new licensing mechanisms and will become part of a high quality show-case to convince more code-owners to adapt their applications.

Licenses as Grid Services

A promising approach to overcome the limitations of the current monolithic licensing models is to reformulate licenses as manageable Grid resources or even Grid services. This will allow managing and orchestrating licenses in a job or workflow together with other resources like compute nodes, data storage, and network Quality of Service (QoS). SmartLM will provide an orchestration service that is capable to co-allocate different resources and services to be used at the same time, e.g. reservation of compute nodes and the licenses for applications that will be executed in a workflow scheduled to these nodes.

Licenses managed as Agreements

In emerging Grids and virtualized infrastructures the conditions of resource usage reflect and extend the conventional Service Level Agreements (SLAs) which are made today between resource providers and resource consumers. These SLAs define for example the lifetime, resource usage, accounting and billing, but also the guarantee of QoS and penalties or compensations if guarantees are missed. In SmartLM the agreements between service provider and service consumer about using license protected software are expected to transport and propagate the appropriate licenses corresponding to the providers' and consumers' requirements. Encrypting and signing of the agreements will provide the necessary level of security and protection against fraud. All negotiations aiming to lead to a Service Level Agreement, i.e. an electronic contract, on terms of the usage of licensed software under the umbrella of the license contract framework. Agreement based license schema enable brokering of licenses, thus allowing to extend the current business models and to create new ones.

Dynamic Licenses

To support license agreements that change over time and for cases where negotiation between service provider (i.e. license provider) and service consumer is needed to settle an agreement, SmartLM will implement negotiation procedures. Making licenses flexible with respect to lifetime, agreed QoS, pricing, etc. allows for the re-negotiating of service consumption in case of unpredicted extension, reduction, or other changes of the requested and agreed service.

Project Facts

The SmartLM project is partly funded by the European Commission under contract number 216759. The project started in February 2008 and has a duration of 30 months. The project consortium, led by Atos Origin, is formed by Independent Software Vendor (ISV), Business Analyst, Application Service Provider (ASP), academic partners, and public centres: Atos Origin (Spain), Fraunhofer SCAI (Germany), FZ-Jülich (Germany), CINECA (Italy), The 451 Group (UK), INTES (Germany), ANSYS (Germany), LMS International (Belgium), T-Systems (Germany), CESGA (Spain), and Gridcore Aktiebolag (Sweden).

The website of the project offers detailed information: http://www.smartlm.eu/

- Daniel Mallmann[1]
- Wolfgang Ziegler[2]
- Josep Martrat[3]

[1] Forschungszentrum Jülich, Jülich Supercomputing Centre

[2] Fraunhofer Institut SCAI, Sankt Augustin

[3] Atos Origin, Barcelona

Projects

Toward Optimal Load Balance in Parallel Programs for Geophysical Simulations

Understanding and simulating seismic phenomena is of vital importance today. Seissol, the tool for the simulation of the generation and propagation of seismic waves, used for such purpose has been modified to achieve better speed up by introducing the concept of local time stepping. In the original implementation of Seissol, the time step was calculated for all elements being the smallest chosen and communicated to all elements. Thus, regardless of the shape and size of any element, all of them were doing the same calculations in one time iteration. With local time stepping each cell uses its own time measurement. In one time iteration some cells will be updated while some will not. Some subdomains have more cells to update than others, introducing load imbalance.

Standard methods like METIS [1], used to distribute data among processors are no longer capable of decomposing the domains in an efficient manner. We are now looking toward new algorithms that preserve the good qualities of the local time stepping and reduce the load imbalance.

Space Filling Curves (SFC) is, in few words, a line that passes for every point in a discrete domain (Fig. 1). We use the Hilbert space filling curve that has two nice properties. One is its shape that maintains communication at a minimum level. The second one is that we can calculate the order in the curve of any point knowing only its coordinates.

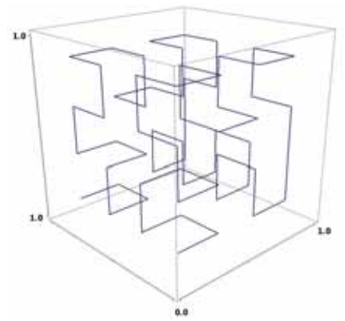

Figure 1: Hilbert SFC of level 2 and 64 subcubes

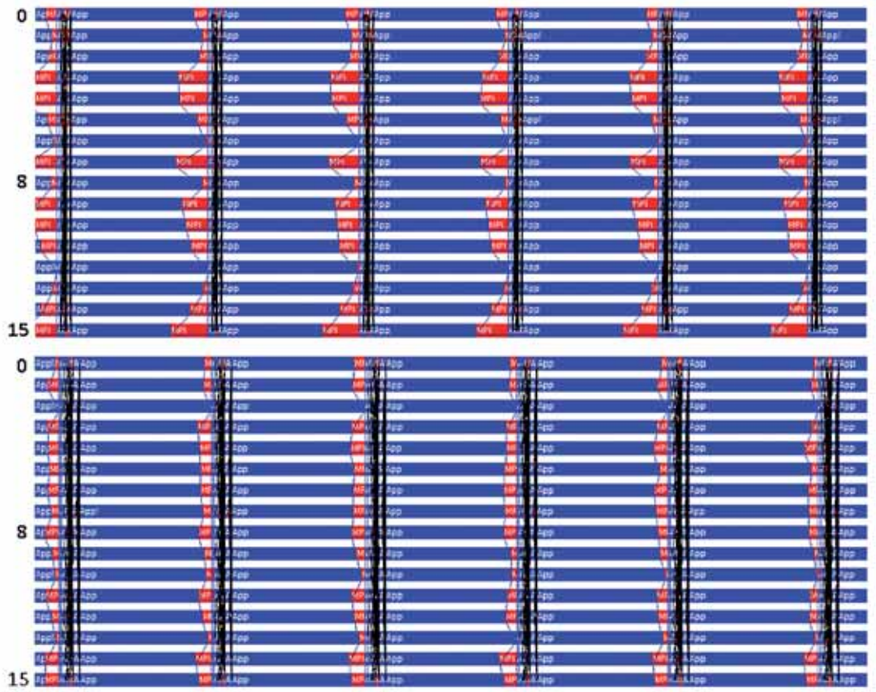

Figure 2: Event time line. Up: Load balance using METIS. Down: load balance using SFC. In blue user code and in red MPI calls, 16 processors.

Generating a space filling curve is a demanding task. Besides we have also to performe a lot of calculations over millions of cells, so brute-force methods are not going to help. A SFC of level 7, normally required, has over 2.0×10^6 points whose information has to be stored somehow (each level increases the number of points by a factor of 8). Instead we use modern programming techniques and efficient algorithms that together with binary representation and Gray code reduce the ratio of calculations from $O(N^2)$ to $O(N)$. In Figure 2 we see the load balance obtained by using SFC compared to METIS.

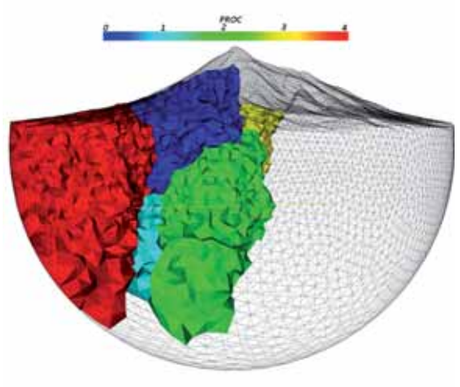

Figure 3: Domain decompoition by METIS, 16 subdomains, visible subdomains: 0 to 4

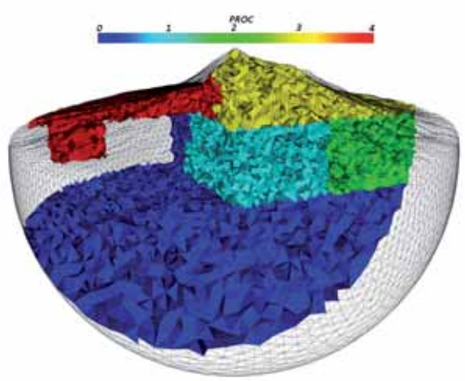

Figure 4: Domain decomposition by SFC, 16 subdomains, visible subdomains: 0 to 4

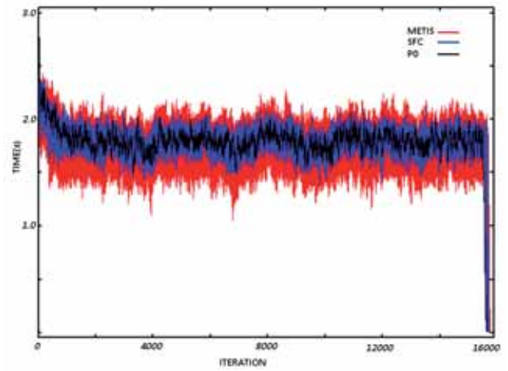

Figure 5: Time per processor vs. Iteration. 16 processors and 15000 iterations

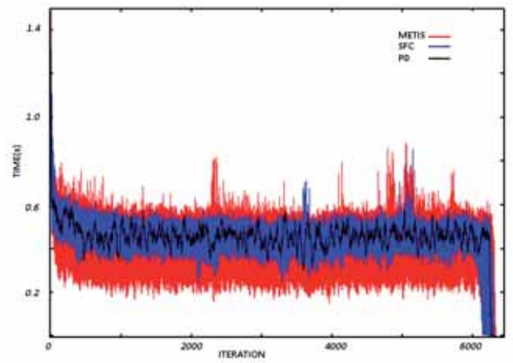

Figure 6: Time per processor vs. Iteration, 64 processors and 6500 iterations

How does it work?

A SFC begins at one corner of our domain and starts collecting elements until they sum up to the load that corresponds to one subdomain.

This process is repeated for the next subdomian. Obviously, subdomains will no longer have the same number of cells. Instead they have a comparable load. In Figure 4 we see a typical domain decomposition produced by SFC, while in Figure 3 the same domain is decomposed by METIS.

Conclusions

In general we have seen an average improvement between 15% and 20%. We are still investigating how the time evolution of the load balance develops.

The time tracing per iteration of all processors during one single simulation generates a spectrum-like graph. The narrower the band, the better the load balance is, as shown in Figure 5 and 6. The more processors are used for the simulation, the more crucial the load balancing problems gets. The domain decomposition made by SFC presents a better load balance for 16 and 64 subdomains. In the later case the load imbalance is dramatically increased by the relatively small number of cells per processor. However, the partition done by SFC presents a better performance. One interesting property of these pictures is, if we trace one single processor during the simulation we will see it moving from top to bottom and vice versa in the band. The same effect was observed either using METIS or SFC decomposition (in Figure 5 and 6 processor 0 is drawn in black as example). That means that a processor might be sometimes the fastest or sometimes might be the slowest. Predicting the load balance statically is a very difficult business. But most importantly, it tells us that if we start considering dynamic load balance we know as a fact that this has to be done much more often during the simulation than we thought. We will see cells jumping from subdomains to subdomains. How to migrate them without creating an excessive overhead is still an open question.

References

[1] **Karypis, G.**
METIS - Family of Multilevel Partitioning Algorithms, http://glaros.dtc.umn.edu/gkhome, 2006-2007

[2] **Dumbser, M., Käser, M., Toro, E.**
An Arbitrary High Order Discontinuous Galerkin Method for Elastic Waves on Unstructured Meshes V: Local Time Stepping and p-Adaptivity, 2000

[3] **Hamilton., C.**
Compact Hilbert Indices, Dalhousie University Technical Report CS-2006-07, July 2006

• Orlando Rivera[1]
• Martin Käser[2]

[1] Leibniz Supercomputing Centre

[2] University of Munich, Department of Earth and Environmental Sciences

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

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

Contact
Leibniz-Rechenzentrum

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

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

Compute servers currently operated by LRZ are

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
HRLB II: SGI Altix 4700 Intel IA64 19 x 512-way	9,728 Cores 39 TByte	62,3	Capability Computing	German Universities and Research Institutes, DEISA
Linux-Cluster SGI Altix 4700 256-way	256 Cores 1 TByte	1,6	Capability Computing	Bavarian Universities
Linux-Cluster SGI Altix 3700 BX2 128-way	128 Cores 0,5 TByte	0,8	Capability Computing	Bavarian Universities
Linux-Cluster Intel Xeon EM64T AMD Opteron 2-, 4-, 8-, and 16-way	2,020 Cores 3,7 TByte	12,3	Capacity Computing	Bavarian and Munich Universities, D-Grid, Bavarian State Library
Linux-Cluster Intel IA64 2-, 4- and 8-way	220 Cores 1,1 TByte	1,3	Capacity Computing	Bavarian and Munich Universities
Linux-Cluster Intel Xeon EM64T 4-way	752 Cores 1,4 TByte	7,2	Capacity Computing	LCG Grid Tier 2
Linux-Cluster Intel IA32	132 Cores 0,3 TByte	0,72	Capacity Computing	LCG Grid Tier 2

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

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

Operation of its systems is done together with T-Systems, T-Systems sfr, and Porsche in the public-private joint venture hww (Höchstleistungsrechner für Wissenschaft und Wirtschaft). Through this co-operation a variety of systems can be provided to its users.

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

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

Contact
Höchstleistungsrechenzentrum
Stuttgart (HLRS)
Universität Stuttgart

Prof. Dr.-Ing. Michael M. Resch
Nobelstraße 19
70569 Stuttgart
Germany
Phone +49-711-685-872 69
resch@hlrs.de
www.hlrs.de

View of the NEC SX-8 at HLRS

Compute servers currently operated by HLRS are

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
NEC SX-8	72 8-way nodes 9,22 TByte memory	12,67	Capability Computing	German Universities, Research Institutes and Industry
TX-7	32-way node 256 GByte memory	0,19	Preprocessing for SX-8	German Universities, Research Institutes and Industry
IBM BW-Grid	498 2-way nodes 8 TByte memory	45,9	Grid Computing	D-Grid Community
Nec SX-9	8 16-way nodes 4 TByte memory	12,8	Grid Computing	D-Grid Community
Intel Nocona	205 2-way nodes 0,62 TByte memory	2,5	Capacity Computing	Research Institutes and Industry
AMD Opteron	194 2-way nodes 1 TByte memory	2,5	Capacity Computing	Research Institutes and Industry
Cray XD1	48 2-way nodes 0,24 TByte memory	0,54	Industrial Development	Industry and Research Institutes

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

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 modeling and computer simulation including their methods. The supercomputers with the required information technology infrastructure (software, data storage, networks) are operated by the Jülich Supercomputing Centre (JSC) and by the Centre 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 Zoltan Fodor and located at the DESY laboratory in Zeuthen and the group Computational Biology and Biophysics, headed by Ulrich Hansmann at the Research Centre Jülich.

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

Contact
John von Neumann Institute
for Computing (NIC)
Jülich Supercomputing Centre (JSC)
Institute for Advanced Simulation (IAS)

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
www.fz-juelich.de/jsc

The IBM supercomputers “JUGENE“ (right) and “JUMP“ (left) in Jülich (Photo: Research Centre Jülich)

Compute servers currently operated by JSC are

System	Size	Peak Performance (TFlop/s)	Purpose	User Community
IBM BlueGene/P “JUGENE”	16 racks 16,384 nodes 65,536 processors PowerPC 450 32 TByte memory	222,8	Capability computing	German Universities, Research Institutes and Industry
IBM pSeries 690 Cluster 1600 “JUMP”	41 SMP nodes 1,312 processors POWER4+ 5,1 TByte memory	9,0	Capability computing	German Universities, Research Institutes and Industry
IBM BladeCentre-H “JULI”	2 racks 56 Blades 224 PowerPC 970 MP cores 224 GByte memory	2,2	Capability computing	Selected NIC Projects
IBM Cell System “JUICE”	12 Blades 24 Cell processors 12 GByte memory	4,8 (single precision)	Capability computing	Selected NIC Projects
AMD Linux Cluster “SoftComp”	125 compute nodes 500 AMD Opteron 2.0 GHz cores 708 GByte memory	2,5	Capability computing	EU SoftComp Community
AMD Linux Cluster “JUGGLE”	44 compute nodes 176 AMD Opteron 2.4 GHz cores 352 GByte memory	0,85	Capacity and capability computing	Selected D-Grid Projects
apeNEXT (special purpose computer)	4 racks 2,048 processors 512 GByte memory	2,5	Capability computing	Lattice gauge theory Groups at Universities and Research Institutes
APEmille (special purpose computer)	4 racks 1,024 processors 32 GByte memory	0,55	Capability computing	Lattice gauge theory Groups at Universities and Research Institutes

The Jülich Supercomputing Centre

Supercomputing is a driving force for the field of simulation sciences, which is the third category of scientific research complementing theory and experiment. The ever increasing complexity of the systems and processes under investigation in science and engineering results at the same time in growing requirements concerning the accessibility of theories to simulation, the accuracy of mathematical modeling, the efficiency of numerical and stochastic methods, and the computational methodology. This includes the performance and scalability of supercomputers, networks, Grid infrastructures, and data centres as well as programming models, software technology and visualization techniques. Furthermore, due to the generation of complex data flows, future large-scale experimental research presents extraordinary challenges for the storage and processing of huge amounts of experimental data. The mission of the Jülich Supercomputing Centre (JSC) is to provide large-scale computational resources and infrastructures for German and European science. Under the name Supercomputing, JSC's work programme forms part of the Helmholtz Association's research field Key Technologies. It is structured into three topics:

- Supercomputer Facility, concerning the provision of super–computer resources of the highest performance and widest scope along with the design and construction of future leadership-class systems in co-operation with European companies.
- Computational Science and Mathematical Methods, providing well-qualified support through Simulation Laboratories, an innovative community-oriented research and support structure, together with cross-disciplinary scientific support groups.
- Grid Technologies and Infrastructures, making Grid technologies available for distributed supercomputing infrastructures, empowering computational scientists and experimentalists to use the European high-performance computing (HPC) infrastructure most effectively.

In the following, we will outline some of the important goals and activities for the near future.

Supercomputer Facility

The prime objective of this topic is the advancement of JSC to a world-class leadership facility. JSC strives to become a European supercomputing centre with

Centres

petaflop/s capability in 2009/2010. As a member of the Gauss Centre for Supercomputing, JSC heads the European project Partnership for Advanced Computing in Europe (PRACE), which is preparing the creation of a pan-European supercomputing service at the highest performance level, its full integration into the European HPC ecosystem and its sustained operation.

For Jülich, the realisation of such ambitious goals implies the cost-efficient and continuous upgrading of JSC's supercomputers, data storage, networking and post-processing capabilities towards a petascale facility, enabling both the highest scalability and broadest general purpose flexibility. Only through a continuous dual modernization process can the sustained provision of supercomputer resources of the highest performance and widest scope be guaranteed.

We can only remain in the forefront of high-performance computing through active research and technological development at the system level. This will be carried out by JSC in close collaboration with leading hardware vendors, software companies and system integrators, embedded in the PRACE initiative. Activities range from the investigation of novel hardware architectures to the design and development of hardware and software for the next generation of leadership-class systems, including cluster computing, networking, security and scientific visualization systems.

Computational Science and Mathematical Methods

An effective utilization of leadership-class supercomputing facilities requires both outstanding technical operation and high-level user support at JSC. To this end, JSC is pursuing the realignment of disciplinary research and cross-sectional activities by its innovative concept of Simulation Laboratories. This measure will ensure the long-term critical mass of the existing research teams and will further intensify cooperation with external academic groups. A Simulation Laboratory is a targeted research and support structure focused on a specific scientific community. Each of these consists of a core group located at a supercomputing centre and a number of associated scientists outside. The first Simulation Laboratories for Jülich and regionally based communities will be created in the fields of computational plasma physics, earth sciences, computational biology and nanosciences. They are complemented by cross-sectional research groups on mathematical modeling and methods and on performance analysis tools that play a key role in the scaling of large application codes to future petascale machines. The topic also comprises the work of the Jülich research group at the John von Neumann Institute for Computing, whose current research theme is Computational Biology and Biophysics.

Grid Technologies and Infrastructures

JSC develops and provides Grid technologies for distributed supercomputing infrastructures based on the UNICORE paradigm. The overall objective is to support national and international e-infrastructures, such as the Gauss Centre for Supercomputing, D-Grid, DEISA, and PRACE, forming the future European high-performance computing infrastructure. JSC collaborates with many partners to achieve standards-based interoperability between the various Grid technologies, thus creating seamless user access to a wide range of e-infrastructures. It designs and implements world-class service-oriented architectures according to emerging standards from the Grid Computing and Web Services domain, accompanied by the identification and implementation of new trends in the distributed systems research area. JSC is a major partner in the operation of the heterogeneous environment of the D-Grid and will continue to sustain the Grid infrastructures for the German e-Science community, including high-level user support and training events in order to strengthen the application sector. The topic also includes JSC's research and development of future communications and networking technologies, both at the wide and local area and at the processor interconnect level.

Partnerships and Alliances

JSC's work has a distinct interdisciplinary and collaborative character. It has an important bridging function for its user communities as well as for the research fields of the Helmholtz Association and partners in national and international science organizations and competence networks.

The simulation activities of JSC are integrated in several formal collaborations, in particular in the newly established Jülich-Aachen Research Alliance (JARA), Section JARA-SIM. The long-term activities and experiences of JSC in the teaching and training of young scientists together with the Aachen University of Applied Sciences are strengthened by its firm commitment to the German Research School for Simulation Sciences at FZJ and RWTH Aachen University, which will be inaugurated in 2008.

JSC was the first German national supercomputer centre founded in 1987 and has established - together with its partners DESY and GSI within the framework of the John von Neumann Institute for Computing (NIC) - a highly regarded peer-review scheme for the provision of supercomputer resources. JSC is a founding member of the European project Distributed European Infrastructure for Supercomputing Applications (DEISA). Together with the national HPC centres Leibniz-Rechenzentrum in Garching (LRZ) and Höchstleistungsrechenzentrum Stuttgart (HLRS), JSC founded the Gauss Centre for Supercomputing (GCS) in 2006. The GCS e.V. represents Germany as a single legal entity in the European supercomputing infrastructure initiative PRACE.

Under its new name, the Jülich Supercomputing Centre will endeavour to continue the internationally recognized work of the Zentralinstitut für Angewandte Mathematik. Prospects of success are very good for the Centre in view of its competence in high-performance computing, its excellent equipment, and its active integration into the European HPC ecosystem.

• Dr. Sabine Höfler-Thierfeldt

Jülich Supercomputing Centre

HPC in Science and Engineering – The 10th Results and Review Workshop of the HLRS

The 10th Results and Review Workshop of the HLRS brought together more than 50 participants from German research institutions, the steering committee and scientific support staff of HLRS. More than 45 sophisticated talk and poster presentations were selected in advance from the steering committee out of the yearly supercomputer project reports and presented at the Results and Review Workshop.

State-of-the-art scientific simulations on the supercomputer systems of HLRS again emphasized the world-class research done at the centre obtaining outstanding results in achieving highest performance for production codes which are of particular interest for both scientists and engineers.

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

Every year the steering committee of the HLRS, a panel of twelve top-class 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 the three most outstanding projects, honoured by the Golden Spike Award. The laureates of the year 2007 and their project titles are:

- Rainer Stauch
  Institute for Technical Thermodynamics, University of Karlsruhe

  ”Ignition of Droplets in a Laminar Convective Environment”

- Janina Zimmermann
  Fraunhofer Institut für Werkstoffmechanik, Freiburg

  ”DFT Modeling of Oxygen Adsorption on CoCr Surfaces”

- Michael Breuer
  Institute of Fluid Mechanics, University of Erlangen-Nürnberg

  ”Direct Numerical Simulation of Turbulent Flow Over Dimples - Code Optimization for NEC SX-8 plus Flow Results”

Transactions of the High Performance Computing Centre Stuttgart

High Performance Computing in Science and Engineering ‘07

The book presents the state-of-the-art in simulation on supercomputers. Leading researchers present results achieved on systems of the Höchstleistungsrechenzentrum Stuttgart (HLRS) for the year 2007.

The reports covers all fields of computational science and engineering ranging from CFD via computational physics and chemistry to computer science with a special emphasis on industrially relevant applications. Presenting results for both vector systems and microprocessor based systems the book makes it possible to compare performance levels and usability of various architectures. The book gives an excellent insight into the potential of vector systems. The book further covers the main methods in high performance computing. Its outstanding results in achieving highest performance for production codes are of particular interest for both scientists and engineers. The book comes with a wealth of coloured illustrations and tables.

Nagel, W. E., Kröner, D., Resch, M. (Eds.)
High Performance Computing in Science and Engineering ‘07, Transactions of the High Performance Computing Centre Stuttgart (HLRS) 2007, Springer, Berlin, Heidelberg, New York 2007

Product Details:
- Hardcover: 676 pages
- 1st edition (January 2008)
- Language: English
- ISBN-10: 3540747389
- ISBN-13: 978-3540747383

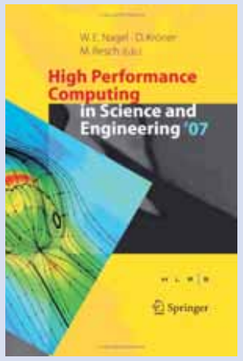

The 8th HLRS-NEC Teraflop Workshop

The eighth Teraflop Workshop is a continuation to the series of workshops held twice a year, alternatively at Tohoku University, Japan and HLRS, Germany as part of the Teraflop-Workbench cooperation between NEC and HLRS. The workshop provides a platform for scientists, application developers, hardware designers and international HPC experts to discuss recent developments and future directions of supercomputing:This year's spring event was held on April 10th to 11th.

The workshop included two hardware sessions with talks on the three major supercomputing projects. The talk on IBM Roadrunner petaflop project which is scheduled to be finished in summer 2008 was given by Cornell Wright from LANL/IBM. He taked about all aspects of the machine including hardware, operation and application software. A general overview of Cray's recent and future HPC products along with an introduction to the Cascade project was given by Wilfried Oed. Hiroaki Kobayashi from Tohoku University, Japan presented their new NEC SX-9 installation along with first experiences on application performance. Tadashi Watanabe from RIKEN, Japan introduced the Japanese "Next Generation Supercomputer" project and underlined the challenges ahead to developing a petaflop system.

The talks on applications reflect the wide range of simulations run on the HLRS hardware installations. Various talks were closely related to industrial applications such as surface technologies (Andreas Scheibe, IFF, University of Stuttgart), the simulation of fluid structure interaction in turbines (Albert Ruprecht, IHS, University of Stuttgart) and the simulation of combustion chambers given by Benedetto Riso (RECOM Services). The last talk was accompanied by a presentation from Ulrich Maas (ITT, University of Karlsruhe) on the details of numerical simulation of combustion.

There was a separate session dedicated to computational fluid dynamics in which Markus Kloker (IAG, University of Stuttgart) reported about DNS of controlled shear flow. Matthias Meinke presented a whole bunch of applications that they are working with at AIA, University Aachen and Dinan Wang (CCRL, St.Augustin) talked about blood flow simulation using Lattice Boltzmann methods. Harald Klimach (HLRS, University of Stuttgart) talked about the issues related with heterogeneous parallelism in order to couple a CFD and a CAA (computational aeroacoustics) application. Nowadays, there is a lot of focus on climate and environmental research. The progress achieved in this field over the past few years was detailed in three talks. Regional climate simulations were presented by Hans-Jürgen Panitz (IMK, FZK). Malte Müller of DKRZ talked about barotropic free oscillations of a global ocean model and Stefan Borowski (NEC, HPCE) discussed performance aspects of the widely used ECHAM5 climate code. Finally, the talk given by Stephan Blankenburg (Theoretical Physics,University Paderborn) focused on molecular recognition and self assembly using the molecular dynamics simulation code VASP.

The success of this workshop will be followed up by the proceedings book "High performance computing on vector systems 2008", to appear in autumn. The next Teraflop Workshop will be held at Tohoku University, Japan, in November 2008.

First Industrial Grids Meeting

The first Industrial Grids Meeting took place on October 25th 2008 in Leinfelden-Echterdngen near Stuttgart. It was organized by the projects AeroGrid, In-Grid, PartnerGrid, and ProGrid. Representatives of all D-Grid projects with a strong stake in industry, especially the 40 industrial project partners, were invited. The main points for discussion were the expectations and requirements for Grid Computing by the industry. Representatives from the organizing projects as well as from the projects BauVO-Grid, BIS-Grid, biz2-Grid, GDI-Grid, and FinGrid met for discussion. The participants were able to discuss a large number of points-of-view covering a broad palette of use cases that ranged over scientific, technical, commerical, financial as well as organizational work areas for the industrial Grid users.

Over the course of the workshop, after setting aside typical specifics, much commonness between the D-Grid projects was demonstrated. Important requirements were to improve usability, integration into the existing IT infrastructures and the possibilities for long-term system usage; improved by the reusability of the services offered. Another focal point of common expectations and requirements is for the organizations to integrate the users into the Grid structure. This comprises the management of roles and rights as well as the service provision of legally binding business processes and the availability of service level agreements. Finally, there are several common requirements regarding the Grid technology and the tools to operate the infrastructure. Examples discussed were applying security to the networks as well as questions in the fields of management and technical monitoring of the systems.

It was agreed to continue the contacts inspired by this exchange of information. There are some intentions to form work groups between the different projects to tackle several themes in the future.

NIC-Workshop "From Computational Biophysics to Systems Biology"

About 120 scientist and students from Germany, Europe and the USA participated in the third annual workshop "From Computational Biophysics to Systems Biology" (CBSB08). The workshop took place from May 19th to May 21st 2007 at the Research Centre Jülich and was organized by the NIC research group "Computational Biology and Biophysics" headed by Prof. Dr. Ulrich H.E. Hansmann. As in past years, experts from physics, chemistry, biology, and computer science discussed how to bridge the different scales in physics-based simulations of cells. Topics included protein folding and misfolding, aggregation, molecular docking, and the modeling of cardiac contractions.

CECAM Tutorial 2008

Twenty scientists from several European and non-European countries attended this year's CECAM tutorial entitled "Programming Parallel Computers" at the Jülich Supercomputing Centre. This was the third tutorial of its kind held in Jülich and is part of the educational programme offered by the Centre Européen de Calcul Atomique et Moléculaire (CECAM). From 11th to 15th February 2008, the participants learned about aspects of parallel programming with MPI and OpenMP. JSC experts were responsible for the mixed programme of lectures and hands-on exercises, which continued the success of the preceding tutorials in 2006 and 2007 and established the tutorial in CECAM's educational programme.

FZJ-CEA Workshop on High-Performance Computing and Simulation

The Jülich Supercomputing Centre hosted a meeting of researchers from Jülich and the French research organization CEA. The subject of the workshop was the future collaboration between the two organizations, particularly in the computational sciences and in supercomputing. Leading scientists from a wide range of scientific and technical disciplines were present and discussed plans for joint research in the fields of materials science and quantum technology, soft matter and biophysics, earth and atmospheric sciences, nuclear safety research, hadron physics and nuclear structure, as well as plasma physics and fusion. It was noted that in these areas of the simulation sciences the two centres are in the forefront of European research, due both to their outstanding science and their excellent computational resources.

CEA and FZJ are members of the national consortia which are preparing for the future European HPC ecosystem. They stated that they are both ready and willing to host future European supercomputers and that they want to start with technical preparations for the evaluation and installation of these petaflop systems. They agreed to coordinate their search for complementary systems, which is a requirement of the international simulation community. They also decided to jointly investigate fundamental technical challenges of extreme parallelism. These challenges include the development of highly scalable communication technology, tools for monitoring and steering energy consumption, and the development of terabyte/s-class I/O systems with the emphasis on open software, interoperability, and hierarchical storage.

Special discussions concerned the fusion experiment ITER, which will be built on the CEA campus at Cadarache. CEA supports Jülich's proposal to host and operate a 100 TFlop/s system required for crucial simulations accompanying the construction phase.

The workshop successfully defined the scientific basis for a bilateral collaboration agreement between FZJ and CEA, which will be presented at the 3rd German-French Research Forum in Paris on 29th February 2008.

HLRS Workshop Report and Outlook

In spring 2008 HLRS had a number of courses an workshops. The total number of participants in these courses from Germany and Europe was about 150.

One of the flagships of our courses is the week on **Iterative Solvers and Parallelization**. Prof. A. Meister and Prof. B. Fischer teach basics and details on Krylov Subspace Methods. Lecturers from HLRS give lessons on distributed memory parallelization with the Message Passing Interface (MPI) and shared memory multi-threading with OpenMP. This course will be repeated in September 2008 at LRZ in Garching.

NEC SX-8 Usage and Programming is a 1,5 day course, dedicated to our NEC SX-8 users to learn about vectorization and optimization of their applications on the Multi-Teraflops-System at HLRS.

The course on **Fortran for Scientific Computing** was given the first time at HLRS in October 2006. The number of participants is still growing from course to course. This year, we had therefore to move into our largest lab.

Mainly PhD students from Stuttgart and other universities in Germany came to learn not only the basics of programming, but also to get insight into the principles of developing high-performance applications with Fortran. The Fortran course will be repeated in October this year.

On the next side you find a summary of HLRS courses for autumn 2008. Detailed information about individual courses providing information about dates, location and contents.

July-December 2008 – Workshop Announcements
Scientific Workshops at HLRS, 2008
2nd HLRS Parallel Tools Workshop (July 7-9)
High Performance Computing in Science and Engineering - The 11th Results and Review Workshop of the HPC Center Stuttgart (September 29-30)
IDC International HPC User Forum (October 13-14)
Parallel Programming Workshops: Training in Parallel Programming and CFD
2nd HLRS Parallel Tools Workshop (HLRS, July 7-9)
Parallel Programming with MPI & OpenMP (CSCS Manno, CH, August 12-14)
Iterative Linear Solvers and Parallelization (LRZ, Garching, September 15-19)
Introduction to Computational Fluid Dynamics (HLRS, September 22-26)
Message Passing Interface (MPI) for Beginners (HLRS, October 6-7)
Shared Memory Parallelization with OpenMP (HLRS, October 8)
Advanced Topics in Parallel Programming (HLRS, October 9-10)
Parallel Programming with MPI & OpenMP (FZ Jülich, ZAM / NIC, November 26-28)
Training in Programming Languages at HLRS
Fortran for Scientific Computing (October 27-31)
URL
http: / / www.hlrs.de / news-events / events /
http: / / www.hlrs.de / news-events / events / 2008 / parallel_prog_fall2008 /
http: / / www.hlrs.de / news-events / events / 2008 / prog_lang_fall2008 /

Workshop Announcement: 2nd HLRS Parallel Tools Workshop 2008

Date & Location
July 7-9, 2008, HLRS, Stuttgart

Developing for current and future processors will more and more require parallel programming techniques for application and library programmers.

This workshop offers to the industrial and scientific user community, as well as the tools developers itself an in-depth workshop on the state-of-the-art of parallel programming tools, ranging from debugging tools, performance analysis and best practices in integrated developing environments for parallel platforms.

Participants and tools developers itself will get the chance to see the strengths of the various tools. Therefore, this workshop is focused on persons who already know about parallel programming, e.g. with MPI or OpenMP. Hands-on sessions will give a first touch and allow to test the features of the tools.

High Performance Computing Courses and Tutorials

Parallel Programming with MPI, OpenMP and PETSc

Date & Location
August 12-14, 2008
CSCS, Manno (CH)

November 26-28, 2008
NIC/JSC, Research Centre Jülich

Contents
The focus is on MPI, OpenMP, and PETSc. Hands-on sessions (C/Fortran) will allow to immediately test and understand the basic constructs of the presented programming models. Course language is English.
Lecturer: R. Rabenseifner (HLRS)

Webpage
http://www.hlrs.de/news-events/external-events
http://www.fz-juelich.de/zam/neues/termine/mpi-openmp

Introduction to Computational Fluids Dynamics

Date & Location
September 22-26, 2008
HLRS, Stuttgart

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

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

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

Message Passing Interface (MPI) for Beginners

Date & Location
October 6-7, 2008, HLRS, Stuttgart

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

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

Shared Memory Parallelization with OpenMP

Date & Location
October 8, 2008, HLRS, Stuttgart

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

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

Advanced Topics in Parallel Programming

Date & Location
October 9-10, 2008, HLRS, Stuttgart

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

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

Fortran for Scientific Computing

Date & Location
October 27-31, 2008, HLRS, Stuttgart

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

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

Iterative Linear Solvers and Parallelization

Date & Location
September 15-19, 2008
LRZ Building, Garching/Munich

Contents
The focus is on iterative and parallel solvers, the parallel programming models MPI and OpenMP, and the parallel middleware PETSc.

Different modern Krylov Subspace Methods (CG, GMRES, BiCGSTAB ...) as well as highly efficient preconditioning and multi Grid techniques are presented in the context of real life applications. Hands-on sessions (C/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
http://www.hlrs.de/news-events/external-events

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

Date & Location
August 5-8, 2008
NIC/JSC, Research Centre Jülich

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

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

Introduction to Programming and Using the IBM Supercomputers

Date & Location
August 11-12, 2008
NIC/JSC, Research Centre Jülich

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

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

CECAM Tutorial: Programming Parallel Computers

Date & Location
February 2009 (tentative)
NIC/JSC, Research Centre 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.org/tutorials.html

inSiDE

inSiDE is published two times a year by
The GAUSS Centre for Supercomputing
(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

Stefanie Pichlmayer hpcpich@hlrs.de

Authors

Benjamin Ahrenholz ahrenholz@cab.bau.tu-bs.de
Michael Breuer mbreuer@lstm.uni-erlangen.de
Carlos S. Frenk c.s.frenk@durham.ac.uk
Stefan Heinzel heinzel@rzg.mpg.de
Amina Helmi ahelmi@astro.rug.nl
Sabine Höfler-Thierfeldt s.hoefler-thierfeldt@fz-juelich.de
Gábor Janiga janiga@ovgu.de
Hans-Thomas Janka thj@mpa-garching.mpg.de
Adrian Jenkins a.r.jenkins@durham.ac.uk
Martin Käser martin.kaeser@geophysik.uni-muenchen.de
Manfred Krafczyk kraft@cab.bau.tu-bs.de
Hermann Lederer lederer@rzg.mpg.de
Peter Lehmann peter.lehmann@epfl.ch
Daniel Mallmann d.mallmann@fz-juelich.de
Andreas Marek amarek@mpa-garching.mpg.de
Josep Martrat josep.martrat@atosresearch.eu
Bernhard Müller bjmuellr@mpa-garching.mpg.de
Julio Navarro jfn@uvic.ca
Orlando Rivera orlando.rivera@lrz.de
Jörg Schumacher joerg.schumacher@tu-ilmenau.de
Hemdan Shalaby shalaby@ovgu.de
Volker Springel volker@mpa-garching.mpg.de
Dominique Thévenin thevenin@ovgu.de
Jonas Tölke toelke@cab.bau.tu-bs.de
Stefan Wesner wesner@hlrs.de
Simon D.M. White swhite@mpa-garching.mpg.de
Wolfgang Ziegler wolfgang.ziegler@scai.fraunhofer.de

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